GA32-0030-2
File No. S370-07

Systems

IBM 3850 Mass Storage System (MSS) Installation Guide

IBM

In this publication, any reference to an IBM program product or Installed User Program (IUP) is not intended to state or imply that only these programs may be used; any functionally equivalent program may be used instead. This publication has references to the following IBM programs:

JDCA - Job Data Compression Aid, IUP 5796-PHN

RACF - Resource Access Control Facility, Program Product 5740-XXH

SDMA - Storage Device Migration Aid, IUP 5796-PHP

OS/VS JCL Editor Program, IUP 5796-PDC

Third Edition (May 1977)

This edition is a revision of the previous edition numbered GA32-0030-1 and includes Technical Newsletter GN32-0013.

Each technical change is marked by a vertical bar to the left of the change.

The information contained in this publication is subject to significant change. Any such changes will be published in new editions or technical newsletters. Before using this publication in connection with the operation of IBM systems, consult the *System/360 and System/370 Bibliography*, GA22-6822, and the latest *System/370 Advanced Function Bibliography*, GC20-1763, for the editions that are applicable and current.

Requests for copies of IBM publications should be made to your local IBM representative or to the IBM branch office that serves you.

A form for reader's comments is at the back of this publication. Address any additional comments to Information Planning and Editing, Dept. 21E, P. O. Box 1900, Boulder, Colorado 80302.

About this manual...

This publication is a guide for installing the IBM 3850 Mass Storage System (MSS). It documents the composite knowledge of the Mass Storage System Support Group and the experience gained with installed systems. It is to be used by the installation team (as defined in the first chapter) and IBM people responsible for installing the Mass Storage System. Every attempt has been made to clarify complex areas, and to provide recommendations or alternate solutions to many of the data migration and conversion problems. We will update this publication as we accumulate experience. Your input is invited. Please use the last page of this publication, entitled 'Your Comments, Please...'. It is there for your convenience.

Because the chapters vary in complexity, you may find the following summary a useful guide for locating the information desired.

- **"Overview"** describes planning and installation in general, beginning with the prerequisites, and includes project management considerations. The subsequent chapters describe each topic in more detail.
- **"Planning a New Data Environment"** introduces a new data environment. This chapter explores the selection of data, its migration to mass storage volumes, and the space/volume standards for the new Mass Storage System data environment. Examples from internal installations are included as an overview of the data planning and migration process.
- **"Logical Concepts, Addresses, and Notations"** describes the basic logical concepts and the physical units that are used to "create" virtual drives. It introduces the Mass Storage Control Table Create program notation that the succeeding chapters use to identify physical and logical addresses, connections, and configuration.
- **"Logical Connections"** describes in detail the five basic connections and the three key data sets used in a Mass Storage System configuration. The chapter then extends the basic connections to include multi-Mass Storage Facilities, attached processors, and multiprocessor configurations.
- **"Related Logical and Physical Planning"** describes unit address considerations for the Mass Storage Control Table Create program and SYSGEN, including unit control blocks and unit control words. It also describes shared DASD and the string-switching of real drives and summarizes physical component layouts, cable lengths, and I/O device priorities.
- **"Staging Drive Group Space"** describes how the Least Recently Used (LRU) service function and Allocation routines control page activity, and how to specify LRU parameters to the Mass Storage Control Table Create program.
- **"MSC Table Create Program Worksheets Description"** summarizes the table create process and describes how to use the worksheets as input to the Mass Storage Control Table Create program. It uses as an example a minimum Mass Storage System configuration for the XYZ Company.
- **"Mass Storage Control Table Create Examples"** uses the worksheets to describe planning and execution of two Mass Storage Control Table Create programs for the fictitious XYZ Company. An attached processor with asymmetric connections to the Mass Storage Control and a

multiprocessor with symmetric and asymmetric connections to the Mass Storage Control are used as examples.

- **"Data Security and Recovery Planning"** describes these new considerations for the 3850, relates them to current procedures, and recommends backup methods for customer Mass Storage Control, MSVC, and application data sets.
- **"Installation"** summarizes the installation in a general sequence.
- **"Appendix A: Mass Storage Control Table Create Program Worksheets"** contains copyrighted worksheets that describe and develop input specifications for the Mass Storage Control Table Create program. For your convenience, the worksheets can be reproduced for the sole purpose of creating program inputs in connection with configuring apparatus manufactured by IBM.
- **"Appendix B: Worksheets for the XYZ Company, Configuration B"** contains completed worksheets and Mass Storage Control Table Create program output for a multihost (attached processor and multiprocessor systems) Mass Storage System configuration. The multiprocessor has symmetric connections to the Mass Storage Control and is described in the chapter "Mass Storage Control Table Create Examples".
- **"Appendix C: Worksheets for the XYZ Company, Configuration C"** contains the worksheets (and the Mass Storage Control Table Create program output) required to change the symmetric connection between the multiprocessor and the Mass Storage Control to an asymmetric connection. The changes are described in the chapter "Mass Storage Control Table Create Examples".

Summary of Amendments

Changes have been made throughout this publication to clarify the text and figures, add new material, and integrate the Mass Storage System enhancements. The major changes are summarized as follows:

- Attachment of System/370 Model 148 is documented in the "Overview" and "Installation" chapters.
- Allocation to mounted volume support is described in the "Planning a New Data Environment" chapter.
- The attachment of the IBM 3350 disk drives is described in the "Logical Concepts, Addresses, and Notations" and "Logical Connections" chapters.
- New subjects, "Shared DASD", "String-Switching Real Drives", and "Mass Storage Facility Model Change" have been added to the "Related Logical and Physical Planning" chapter.
- Preferred inactive space has been included in the "Staging Drive Group Space" chapter.

A minimum Mass Storage System configuration is described in Mass Storage Control Table Create Worksheet Description. This chapter has been revised and should be reviewed in its entirety.

- An AP with asymmetric connections to the Mass Storage Control, and an MP with symmetric and asymmetric connections to the Mass Storage Control are described in the "Mass Storage Control Examples" chapter. This chapter has been revised and should be reviewed in its entirety.

- The Mass Storage Control Table Create Program worksheets are revised and are contained in Appendix A.

Required Publications

- *Introduction to the IBM 3850 Mass Storage System (MSS),* GA32-0028, which introduces the product and basic vocabulary.
- *OS/VS Mass Storage System (MSS) Planning Guide,* GC35-0011, which expands the basic concepts and introduces the new Access Method Services commands.
- *IBM 3850 Mass Storage System (MSS) Principles of Operation,* GA32-0029, which describes the operation of the basic components and the subsystem.
- *OS/VS Mass Storage System (MSS) Services: General Information,* GC35-0016, which describes the volume control functions, the Inventory data set, the Journal data set, and the use of the new Access Method Services commands.
- *OS/VS Mass Storage System MSS) Services: Reference Information,* GC35-0017, which describes the syntax of the Mass Storage System Access Method Services commands.
- *OS/VS Mass Storage Control Table Create,* GC35-0013, which describes the Mass Storage Control Table Create program and its commands.

Note: The IBM 3850 Mass Storage System MSS is supported by OS/VS1, OS/VS2 SVS, OS/VS2 MVS, VM/370, and JES3 as described in the "Overview" chapter.

The System Control Program information and references in this publication are OS/VS1 and OS/VS2 MVS as noted. Refer to the OS/VS2 SVS, VM/370, and JES3 publications for mass storage related information.

Contents

Figures

Overview

How to plan for and install the IBM 3850 Mass Storage System (MSS) may be a familiar project because of your experience gained from installing new tape or disk devices, or from converting tape applications to disk. Data migration to this new storage system hierarchy, which is capable of storing and managing all or most of your data, is similar. However, the planning and implementing aspects of a system-managed environment are new.

Prerequisites

The 3850 Mass Storage System contains two major subsystems: the 3851 Mass Storage Facility (MSF), and the 3333/3330 Disk Storage and Controls. Up to four System/370 host processors can attach to the Mass Storage System. The four host processors can be any combination of uniprocessors, attached processors (AP, and multiprocessors (MP). The product is supported beginning with:

- OS/VS1 Release 4
- OS/VS2 SVS Release 1.7
- OS/VS2 MVS Release 3
- VM/370 Release 4
- JES3 Data and Device Management Support of Mass Storage System Selectable Unit (SU). OS/VS2 MVS Release 3.7 and JES3 Release 2 Selectable Unit are prerequisites.

The 3330 Model 1, 2, or 11 disk drives can be used as staging drives with the Mass Storage System. The number of drives varies, depending on your needs. One or more 3830 Model 3 Storage Controls or System/370 Models 158 and 168 Integrated Storage Controls (ISC) that have the Staging Adapter feature are required to control the staging drives and, optionally, real drives. The 3830 Model 2 Storage Controls, the Models 158 ISC and the 168 ISC can be field-changed by adding a Staging Adapter.

System/370 Models 158, 168, 158AP, 168AP, 158MP, and 168MP can attach to a Mass Storage System by using the ISC with the Staging Adapter and/or the 3830 Model 3 Storage Controls. System/370 Models 145, 148, 155II, 158 (and 158 submodel 2 in Japan only) 165II, 168, 158MP, 158AP, 168AP, 168MP, and the 3033 Processor can attach to a Mass Storage System by using a 3830 Model 3 Storage Control. The Staging Adapter function cannot be installed on a System/370 Model 145 or 148 ISC.

Planning and installing these prerequisites should be part of the Mass Storage System installation plan. Physical layout, System Control Program (SCP) version, catalog usage, data set naming conventions, data set placement, program conversion, OS/VS standards, and System/370 virtual storage requirements can be planned to ease the transition to the Mass Storage System.

Related Conversions

Related conversions include those required by a system-managed data environment, by a System Control Program change, and by other major equipment changes scheduled for approximately the same time. Major changes should be made before the Mass Storage System installation; minor changes can be made concurrently with the Mass Storage System installation, if they are carefully planned.

Data migration to a system-managed environment requires broadening the base of your data discipline. With manual controls, an installation can have duplicate volume or data set names, uncataloged data sets, operator intervention, programmer decisions, and so on. With data under control of the Mass Storage System, the objective is to mount more volumes and execute more jobs with minimum operator intervention. To do so, data must be uniquely identified, preferably cataloged, and data set job decisions preplanned so that the system can maintain a continuous flow.

To protect your programming investment, try to limit the required data migration changes to the job control language (OS/VS JCL). If your current programs are device-independent, they probably can be used without being redesigned, and production code probably can be used without recompilation or relink-editing. Time-dependent and space-optimized programs are excluded.

Tape-oriented applications are likely to have the oldest programs that contain device-dependent code. Such programs can be identified and converted before the Mass Storage System installation. Because the Mass Storage System mass storage volumes are compatible with 3330-1 volumes, testing tape application programs and device-dependent code can be done by using a System/370 processor that has 3330-1 disk drives. When the Mass Storage System is installed, the Access Method Services CONVERTV command can be used to transfer the data from the 3330-1 volumes to the mass storage volumes.

Program Products and changes coupled to either a new System Control Program or a new version, such as IMS-VS, CICS-VS, TSO, TCAM, VTAM, and NCP versions must also be considered. Although program conversion and data migration for related changes can be concurrent with Mass Storage System planning, they should be tested and be stable before the Mass Storage System installation begins.

Planning

After the Mass Storage System is installed, you should review Program Products application in this new environment. If you are considering data-base applications, their design can include historical data as well as current transactions.

Because the Mass Storage System introduces a new concept in data processing and management, planning is especially important. For the first time, it is feasible to place most of your data into a system-managed environment, making it readily available for processing.

Selecting the Applications

The normal starting point is to consider the applications that are already operational as the first candidates for transfer to the Mass Storage System. If your application systems (and associated data) are isolated, data set migration planning is simplified. However, it is not unusual for data sets to be used by several application systems, one or more of which may contain device-dependent programs. Consider these factors when you assign the priorities to the application (and data set) migration sequence.

Application systems are comprised of jobs that are identified by unique job control cards. Each application contains one or more jobs and programs that access several data sets. Identify the data sets to be transferred by answering these questions:

- How many are there?
- What is their frequency of use?
- How are they used?
- Which applications use them?
- How large are they?
- Which ones are critical?
- What is their retention period?
- What are their characteristics?
- What OS/VS JCL must be changed, added or deleted?
- What standards should be added or changed?

Analyze your current production over an extended period of time to determine the cyclic use of data sets. The System Management Facilities (SMF) in the MFT, MVT, OS/VS1 and OS/VS2 System Control Programs has options that collect data set information, which are your best source.

SMF data provides information to determine data set characteristics. Application programs are needed to reduce the SMF data to produce "where used" and cross-reference listings by job to provide input for planning your data migration. Two IBM Installed User Programs are available to assist you in these data reductions. They are the SMF Job Data Compression Aid (JDCA), Program Number 5796-PHN, and the Storage Device Migration Aid (SDMA), Program Number 5796-PHP. The interdependencies of data sets, programs, and applications may require several schedule iterations before the final migration sequence is determined.

System Support

Mass Storage System programming support is provided in OS/VS1, OS/VS2 SVS, OS/VS2 MVS, and JES3 System Control Programs. Because the mass storage volume is supported as a virtual 3330 volume (UNIT=3330V), data migration from tape or disk to the Mass Storage System can be controlled primarily by OS/VS JCL changes. This compatibility is made possible by the OS/VS Scheduler which recognizes 3330V as a mass storage volume. It passes the Mount command to the Mass Storage System Communicator (MSSC) for execution by the Mass Storage Control (MSC) rather than by the operator. The MSSC includes the Mass Storage Volume Control (MSVC) functions to assist physical control and space management of mass storage volumes.

These new programming support functions, or extensions to existing functions, recognize or include the Mass Storage System:

- The Mass Storage Control Table Create program, which precedes SYSGEN, for the creation of the logical system and the Mass Storage Control tables that control the Mass Storage System.
- SYSGEN of the real and virtual drives
- New OS/VS JCL, to identify the new device (3330V) and the mass storage volume groups (for example MSVGP=PAYROLL).

- Access Method Services, to create and modify mass storage volumes and volume groups. These include the Trace and Tune commands to analyze the Mass Storage System activity and adjust selected parameters.
- VSAM support
- The Mass Storage System Communicator, which translates SCP requests into Mass Storage Control orders. It includes the Mass Storage Volume Control functions.
- System Data Analyzer (SDA) reports, to analyze the Mass Storage System data cartridge activity in SYS1.LOGREC.

Current data access methods are supported for direct access data sets, and include BSAM, QSAM, BPAM, BDAM, VSAM, EXCP, and XDAP. ISAM can be used with the Mass Storage System but with performance degradation. ISAM data sets should be considered for conversion to VSAM and use of the VSAM Compatibility Interface.

Other inclusions or extensions are console commands and messages, and service functions. The new console commands messages provide operator communication with the Mass Storage System and space and volume information.

Service functions for installation, operator control, recovery, and data migration include new Access Methods Services functions for the Mass Storage System, and extensions to existing OS/VS utilities.

Data sets on mass storage volumes can be created and accessed from a TSO terminal (OS/VS2 only). You must evaluate the response time of the Mass Storage System to determine whether to use TSO for current, batch, or archival data sets, and how to manage them.

JES3, (OS/VS2 MVS prerequisite) manages virtual units and data on mass storage volumes for multiple CPUs connected to one or more Mass Storage Systems.

Because the Mass Storage System operates as an entity and manages its own I/O system, a System Data Analyzer (SDA) program is provided to analyze the usage and error statistics of the subsystem components, including the recording media.

The Mass Storage System is supported by Program Products and programs that support the 3330 Models 1, 2, or 11 with OS/VS1, OS/VS2 SVS and OS/VS2 MVS standard access methods for direct-access data sets.

Configuration

The Mass Storage System is the interconnection of the Mass Storage Facility and the 3333/3330 Disk Storage and Controls. Each subsystem must be physically and logically configured as a total system. The physical configuration is developed by analyzing the current installation workload over a predetermined time in terms of the number of bytes transferred, the number of data sets opened, and the probability of data reuse. If you anticipate equipment changes before the Mass Storage System is installed, the Mass Storage System physical configuration must be adjusted for any new, extended, or additional applications and/or anticipated changes in scheduling.

After the physical configuration is complete, the various components must be logically connected. These connections are the input to the Mass Storage

Control Table Create program when generating the Mass Storage Control tables. The Mass Storage Control tables determine:

- Which CPU requested the data
- Where the data should be staged
- The primary and alternate paths between system and subsystem components
- Which staging drive group to use
- The pages within the staging drive group
- The aging rate of data contained in the staging drive group
- The eligibility of the data for reuse or destaging by the least-recently used (LRU) algorithm

Reliability, performance, backup, capacity, and recovery considerations are other factors that can influence the physical, as well as the logical, configuration.

SYSGEN planning must be included in the logical configuration process to make sure that sufficient unit addresses and nonshared subchannels are available to address the real and virtual drives that you define to the Mass Storage Control Table Create program. Although the storage capability of a virtual system is large, practical limits must be determined and established that correlate with your installation needs.

Physical planning and layout are part of the configuration process. Initial and final layouts must be prepared to accurately position the Mass Storage Facility, especially if multiple CPUs and associated components need to be repositioned to attach to the Mass Storage Facility. Cable lengths should be carefully reviewed.

Remote maintenance is a recommended customer option and requires an acoustically-coupled telephone connection to the Mass Storage Control, ISC with Staging Adapter, and the 3830 Model 3 Storage Control through a portable video terminal. Plan for multiple telephone jacks or one telephone with a line cord that will reach each unit. The portable video terminal extends the CE panel function to a remote location for diagnostic purposes. Reserve a telephone for this purpose.

Security

Data privacy and data security, especially when related to computer systems, are currently receiving increasing attention. As you place more of your data under direct computer control, it becomes more important that the relationship of data privacy and data security to computer systems be better understood.

Data privacy is a business problem to be resolved by public concern or management policies and cannot be resolved by a computer system.

Data security is a technical problem that can be resolved by:

- Security procedures
- Loyal personnel
- Physical protection
- Audited practices

- A secure computer system
- A concerned attitude toward security

Comprehensive data security cannot be achieved. The degree of data security wanted varies with each customer, depending on experience, and exposure to problems experienced by others.

The greatest problems affecting customer data security are unintentional data modification and accidental physical data destruction, which can happen when prescribed procedures are bypassed for the sake of expediency or convenience. These problems are easier to solve than malicious and intentional efforts to copy, steal, or destroy data.

Data security requires:

- Identifying the person, device, or component that has access to a system resource.
- Authorizing access to a system resource only as previously defined.
- Auditing every access to a system resource for reinforcement and evaluation of security measures.
- System integrity, its ability to function according to specifications and protect itself from unauthorized accesses that compromise or bypass security controls.

Because all of these measures are not available today, plan for a level of data security that is consistent with your overall security management plan. Also consider reviewing your installation security management and incorporating the changes that the Mass Storage System requires.

Data security concerns have been recognized and addressed in the Mass Storage System. Internally, checks are made to verify that a component has reacted as directed. For example, whenever a data cartridge is accessed for processing, the Mass Storage Control reads the cartridge label to verify that data cartridge's location.

For internal security, the additional Access Method Services commands can:

- 'Move' a mass storage volume from active to inactive inventory (STOREV) within the Mass Storage Facility so that it is not available for program access, and 'return' it (ADDV) when needed.
- Re-create a mass storage volume from an inactive copy (RECOVERV).

For external security, the additional Access Methods commands can remove a mass storage volume from the 3851 for storage in an off-site location using the EJECTV command or the EJECT option of several other commands.

Mass storage volumes can be protected by these volume attributes:

- DASDERASE, which erases the data in the staging drives after it has been destaged to the mass storage volume.
- READONLY, which protects common data that is accessed by many users from extraneous Write commands.
- EXCLUSIVE, which mounts a mass storage volume for only one CPU to protect data from simultaneous modification by a program in another CPU.

OS/VS2 MVS Release 2 removed all known integrity exposures. We will accept APARs if a program successfully uses a system interface to bypass store or fetch protection, or to obtain control in an unauthorized state.

OS/VS1, OS/VS SVS and OS/VS2 MVS offer the Authorized Program Facility (APF) to restrict the use of the Access Method Services to authorized individuals.

Data sets can be protected by passwords. VSAM contains extended password protection for cataloging and indexing data sets using master, control, update, and read passwords. Passwords can be entered to retrieve, update, insert, or delete data sets, depending on the password level.

The Mass Storage Control passes error recovery activity information to SYS1.LOGREC. The Trace command can be used to collect activity information in the Mass Storage Control tables. Its Volume Activity Reports can provide audit information on a mass storage volume basis. A user-written report based on System Management Facilities (SMF) data is another potential audit source.

You can submit physical protection RPQs for Mass Storage Facility modifications with specifications for a customer-supplied fire detection and suppression system, and for Mass Storage Facility cover locks. Because the Mass Storage Facility does not need operator intervention, other than cartridge entry and exit, consider a separate badge or key-locked enclosure for its physical protection.

The external elements of a data security system, both tangible and intangible, must be part of your security management plan. A controlled physical environment, trained and loyal personnel, and a concerned attitude toward security are necessary to make security procedures function in normal day-to-day operations. Other aspects of internal data security, described under "Recovery Planning", must be considered when data is unavailable because of error conditions.

Recovery

Recovery action is required to restore or re-create lost or damaged data using predefined and tested procedures. It ranges from recovery of an unreadable record to the reconstruction of the total system. Data recovery and data backup are the basis for recovery management. Data recovery is provided by the hardware and/or software algorithms to *restore* the original data, with or without manual intervention. Data backup is provided by safely-guarded copies of critical data needed to *re-create* the original data, using recovery programs and manual intervention.

The Mass Storage System has been designed to handle data recovery with and without operator intervention. It can be configured with alternate paths:

- **At the system level,** with one or more CPUs and alternate paths to the Staging Adapters.
- **At the DASD subsystem level,** with one or more Staging Adapter pairs and alternate paths between the data in staging drive groups and the Mass Storage Facility.
- **At the Mass Storage Facility level,** with models B2 or larger (or two model A2s or larger) and alternate Mass Storage Control paths to the CPUs, and between the data recording controls (DRC) and data recording devices (DRD).

The Mass Storage Facility uses an error-correction code that permits correcting up to 32 bytes of data in a 208-byte block without retrying the read operation. Following a read error, the Staging Adapter can successively adjust the DRD read amplifier gain, realign the tape for potential skew or tracking errors, and the Mass Storage Control can select a different DRD, DRC, and alternate Staging Adapter path to retry a read operation. Each unsuccessful attempt is logged for the Customer Engineer. The Mass Storage Control marks a subsystem component unusable and notifies the operator accordingly when statistical data pinpoints the failing component.

Depending on the model of Mass Storage Facility, multiple Mass Storage Facility components such as accessors and controls, DC power supplies and AC sequencers, pneumatic supplies, and functional components provide backup to single-component failures.

Preventive maintenance and unscheduled repair actions can be performed on subsystem components while the Mass Storage Facility is in operation, which degrades subsystem performance only to the extent necessary to bypass the components being serviced.

The capability of the Mass Storage System to recover data and/or provide alternate paths to data does not remove the need for data backup because the possibility of errors, malicious intent, and disaster remain. Critical data, including data needed to reconstruct your entire data base, must have safely-guarded data backups.

The Mass Storage System offers new approaches to data backup. You can continue to use tape or disk backup with OS/VS utilities, use CONVERTV command (3330-1 only) for backup on disk, or eject and safeguard the mass storage volumes. A secondary data backup can be maintained within the Mass Storage System by using the Access Method Services to deactivate or copy mass storage volumes within the Mass Storage Facility. Evaluate your backup methods and develop a recovery plan for the Mass Storage System that meets your needs.

Installation

You need checkpoints to make certain that all tasks being performed are on schedule. Include these checkpoints in the planning and the implementation phases. Provide for contingencies, alternative plans, and a fall-back position in case problems occur.

Preinstallation Testing

Preinstallation test allowance hours are available for Mass Storage System-related conversions and data migration. These include Mass Storage Control Table Create program for configuration planning, program and data discipline changes, converting device-dependent programs, creating generation data groups, analyzing data sets, changing access methods (ISAM to VSAM), and testing OS/VS JCL changes for tape application by reading and writing data to 3330-1 disks.

The number of available preinstallation test allowance hours depends on the model of Mass Storage Facility. Contact your IBM representative for details.

Physical Installation

The Mass Storage System installation can be completed in two phases. The first is the Staging Adapter additions to the installed 3830 Storage Controls, Models 158 ISC, and 168 ISC; the second is the Mass Storage Facility installation.

Installed 3830 Storage Controls can be field-changed from a Model 2 (or Model 1 by RPQ) to a Model 3, and the Staging Adapter feature can be added to the Models 158 ISC and 168 ISC. Depending on the number of Storage Controls to be modified, potential scheduling conflicts, and the availability of IBM resources, the first phase can begin up to 60 days before the Mass Storage Facility installation.

Installing Staging Adapters is a major task, ranging from 35 to 45 CE hours (20 to 25 elapsed hours) per change. After modification, the 3830 Storage Control or ISC can continue to function in a nonstaging mode, but the number of interfaces available for connection to a CPU is reduced by one (3830 only), the 3340 Direct Access Storage Facility can no longer be attached, and the address format is changed.

The Mass Storage Facility installation time varies depending on the model and work schedules but usually requires 7 to 12 work days. Much of this activity can be completed without interrupting your production. Checking out the system requires approximately 15 to 25 system hours in the final installation stages.

At the completion of the second phase, all of the modified ISC and/or 3830 Storage Controls are permanently changed to function in a staging mode. The 3333/3330 drives attached to the modified ISC and/or 3830 can function in a real mode, but only if they have been defined as real or convertible-to-real to the Mass Storage Control Table Create program.

Other Mass Storage System-related changes can be part of the total physical installation task. Temporary 3830 Storage Controls may be needed to minimize disruption during the first phase. Additional processor storage, channels, and disk drives may be installed during either phase. It might be even necessary to allow time for major relocations in the machine room, especially when more than one CPU is to attach to the Mass Storage System. Physical changes, including interim and final layouts, must be part of the physical planning activity.

On-Site Testing

The initial installation of a Mass Storage System qualifies for an on-site installation and test allowance beginning on the Date of Installation. The maximum allowance varies from three to six weeks, depending on the model of Mass Storage Facility. The Date of Installation is the day following completion of the Mass Storage System physical installation.

For an initial Mass Storage System configuration with two Mass Storage Facilities, the on-site installation and test allowance for each Mass Storage Facility runs concurrently. If the two Mass Storage Facilities are of unequal size, the on-site installation and testing allowance for the smaller Mass Storage Facility ends first, and its rental begins before the larger Mass Storage Facility. No additional on-site testing is provided for a Mass Storage Facility model change.

There are several tasks to be performed before you begin data migration. The first is to initialize the Mass Storage System and execute a regression test on current non-Mass Storage System programs using the OS/VS System Control Program generated for the Mass Storage System. This ensures that you can continue normal production during the on-site testing. Whether or not problems occur, your fall-back plan should be tested at this time.

Computer operations people need to become familiar with the Mass Storage System. They must be able to identify and address the various Mass Storage System components, understand how to vary a 3330 from staging use to nonstaging use, and learn the correct responses to Mass Storage System messages. They must also be able to load data cartridges, and initialize staging disks. They should be given time to learn how to control a system-managed environment on an exception basis. Access Method Services commands should be tested before production. Mass storage volumes must be named and mass storage volume groups created. Data should be copied from tapes and disks to and from mass storage volumes for education, testing, and backup purposes. Data cartridges and volumes should be entered and removed from the Mass Storage Facility as part of the familiarization process.

Sufficient tests should be planned to ensure that the logical configuration described by the Mass Storage Control Table Create program matches the physical configuration. The number of tests depends on the configuration complexity.

Initial testing can begin with a single CPU, a single Staging Adapter, and a minimum number of 3330 drives. As each logical test is successfully completed, additional system components can be varied online and tested until all paths and addresses have been verified.

In the planning phase, applications or sets of jobs will be selected for the pilot testing, parallel operations, and transfer to the Mass Storage System. As these tasks are executed, the data security and recovery management plans associated with these applications should also be tested. Inactive mass storage volume copies should be used to successfully re-create 'destroyed' active mass storage volumes. Data stored in a safely-guarded location on tape, disk, or mass storage volumes should be entered into the system to test the programs and procedures.

When the intial Mass Storage System testing is complete, pilot applications can be tested and parallel operations started. The on-site installation and test allowance ends when the Mass Storage System becomes operational or when the maximum time allowance for the Mass Storage Facility model is reached, whichever occurs first.

Data Migration

Data migration is the process of transferring data sets from disk and tape to Mass Storage System data cartridges for use in the Mass Storage Facility. Data sets to be transferred to the Mass Storage System are selected during the planning cycle. Data set and program interdependencies determine the migration schedule and affect the migration technique selected for each data set.

The techniques to transfer data from tape or disk volumes to mass storage volumes include data set copy, volume copy (3330-1 CONVERTV command only), and data set transfer during normal processing (old master in from tape or disk, new master out to the Mass Storage Facility). Data migration during normal processing is preferred because the workload can be gradually transferred to the Mass Storage System as knowledge of, experience, and confidence in the installation increase. This approach also reduces the CPU and personnel resources required because no additional CPU time is needed, and the effort is spread over a longer period of time.

Part of the data must be transferred to the Mass Storage System using copy functions. This data transfer includes data sets used for testing, fall-back plans, and contingencies. Data sets created by special or periodically-run programs, or updated in place, and data that is reorganized because of device-dependent program changes should also be copied.

The CONVERTV command can be used to copy 3330-1 volumes to mass storage volumes and simultaneously update the catalog to reflect the unit change. To get data set copies, you can use the OS/VS utilities (IEBGENER, IEHMOVE, IEBCOPY) and the Access Method Services.

The migration objective should be to transfer the active data first. This permits you to remove the displaced equipment and gain operational benefits sooner. You may decide to copy tape generations, data backup, and archived data sets on an exception basis, as needed, or as resources permit.

The quick transfer of a large amount of data to the Mass Storage System requires good controls. The transfer process, whether automated or manual, should keep track of how much data has been transferred, and how much remains.

Tuning

The Mass Storage System is an I/O system that manages its own resources to mount and stage data as it is requested, by up to four hosts. The Mass Storage Control uses the Mass Storage Control Table Create LRU parameters that you supply, and communicates with the OS/VS Mass Storage System Communicator modules to evenly distribute the workload across all of the staging drive groups. As your workload, experience, and knowledge grow, you might want to review the Mass Storage System component activities and redistribute the workload, optimize component performance, modify the configuration, or change the Mass Storage Control Table Create parameters.

Existing monitors and tracing functions such as SMF and GTF (generalized trace facility), which are valuable tools in analyzing SCP activity, are limited with respect to the Mass Storage System because its activity is known only to the Mass Storage Control. Although the Mass Storage Control passes error recovery information to SYS1.LOGREC, it does not usually record Mass Storage System activity.

A Trace command in the Access Method Services program allows you to start and stop recording 3850 activity in the Trace X and Trace Y areas of the Mass Storage Control tables. This recording information can be copied to a mass storage volume and the Mass Storage System Trace Report programs can interpret and reformat the data into six different reports. The programs provide distribution reports and activity information on mass storage volumes, by Staging Adapter and data recording device.

Depending on system activity, consider Trace On as a normal mode of operation, at least initially. Provision should be made to continually copy Trace data to mass storage volumes. The volume activity and distribution information is required for LRU parameters, and the volume activity is needed for volume and space management.

Management

Management must assemble a team that will be responsible for planning and installing the Mass Storage System. Management must also assess the installation mission, policies, workload, and procedures to determine which elements are affected and which changes should be made in anticipation of the new subsystem installation.

Functional Changes

Installing the Mass Storage System causes functional changes in your operation. Data in the subsystem is available to the CPU from the Mass Storage System on demand, thereby reducing the physical handling and volume management workload. The subsystem task execution and volume and space management must be planned, then managed on an exception basis to reduce the volume and space concerns for the application programmer, just as virtual storage reduced processor storage management concerns.

Decide whether space management is to be an operational or system programming task. Volume and space management for mass storage volumes is similar to that of the current 3330, except there are more mass storage volumes.

Tools for planning and controlling volume and space management are included in the system support. They are the Mass Storage Volume Control functions, Access Method Services commands, and the OS/VS utilities. Within the guidelines set by your systems programming staff, the Mass Storage System and the MSVC functions automatically perform these tasks.

Space management must intervene when exceptions occur, or when volume space is exceeded. Routine volume and space maintenance should be planned to minimize manual intervention by anticipating and prescribing procedures for exception-handling.

The degree of functional changes varies with each installation. Plans must be made to train people for new tasks as these changes occur.

Education

Education is the most important part of the installation plan. Without sufficient knowledge about the Mass Storage System, you cannot plan and execute an installation plan smoothly. The Mass Storage System installation is equivalent to installing a major new system, not simply adding or replacing a system component.

Three IBM courses are available: a one-day Introduction to the Mass Storage System, a four-day Mass Storage System Installation Planning Workshop, and a four-day Mass Storage System Internals and Recovery class. The latter two are fee education courses.

The objective of the Introduction to the Mass Storage System course is to provide the student with a general understanding of the Mass Storage System as an I/O subsystem. This includes:

- Mass Storage System principles of operation
- New recording technology and media
- Functional component descriptions
- The concepts of virtual volumes and virtual drives
- How data is staged and destaged
- The OS/VS interfaces to the Mass Storage System
- The product relationship to tape and disk storage devices

With this background, the student has an appreciation of potential Mass Storage System applications, can recognize the required operational changes, and knows about the available system support.

The objective of the Mass Storage System Installation Planning Workshop is to provide the installation team with detailed technical information about installing and using the Mass Storage System. They learn to:

- Logically design and describe the Mass Storage System using the Mass Storage Control Table Create program.
- Select the appropriate system support functions, generate the SCP, and use OS/VS utilities.
- Choose and implement migration and conversion techniques.
- Develop a base installation plan.

Other subjects include data security options, Trace reports and their interpretation, use of the Tune command for optimizing subsystem activity, messages and responses, and the operational and functional changes that occur with the installation of the 3850.

The objective of the Mass Storage System Internals and Recovery class is to provide the installation team with detailed technical information about the internal operation of the Mass Storage System and its effect on possible recovery actions. This includes:

- Mass Storage System Communicator (MSSC) general flow diagram
- Mass Storage System Inventory and Journal data sets
- Mass Storage Control operation and tables
- Recovery analysis and procedures

Task Distribution

The installation team must plan and implement several tasks that require skills in systems programming, operations, space management, and application programming.

Figure 1 shows a suggested workload distribution of planning and installation tasks, which constitute the first level of planning. Assignment of these tasks varies with each installation, depending on the organization of the installation team and skills of its members. Task measurement also varies depending on how much change is required in each functional area.

A more detailed checklist of migration and conversion plans might include:

- Applications
- Availability and recovery
- Documentation
- Education
- Equipment modification and installation
- Operations analysis
- Order Data Cartridges
- Performance—measurement and tuning
- Physical facilities
- Physical and logical configuration
- SCP migration
- SCP-related migrations
- Security
- Volume and space management

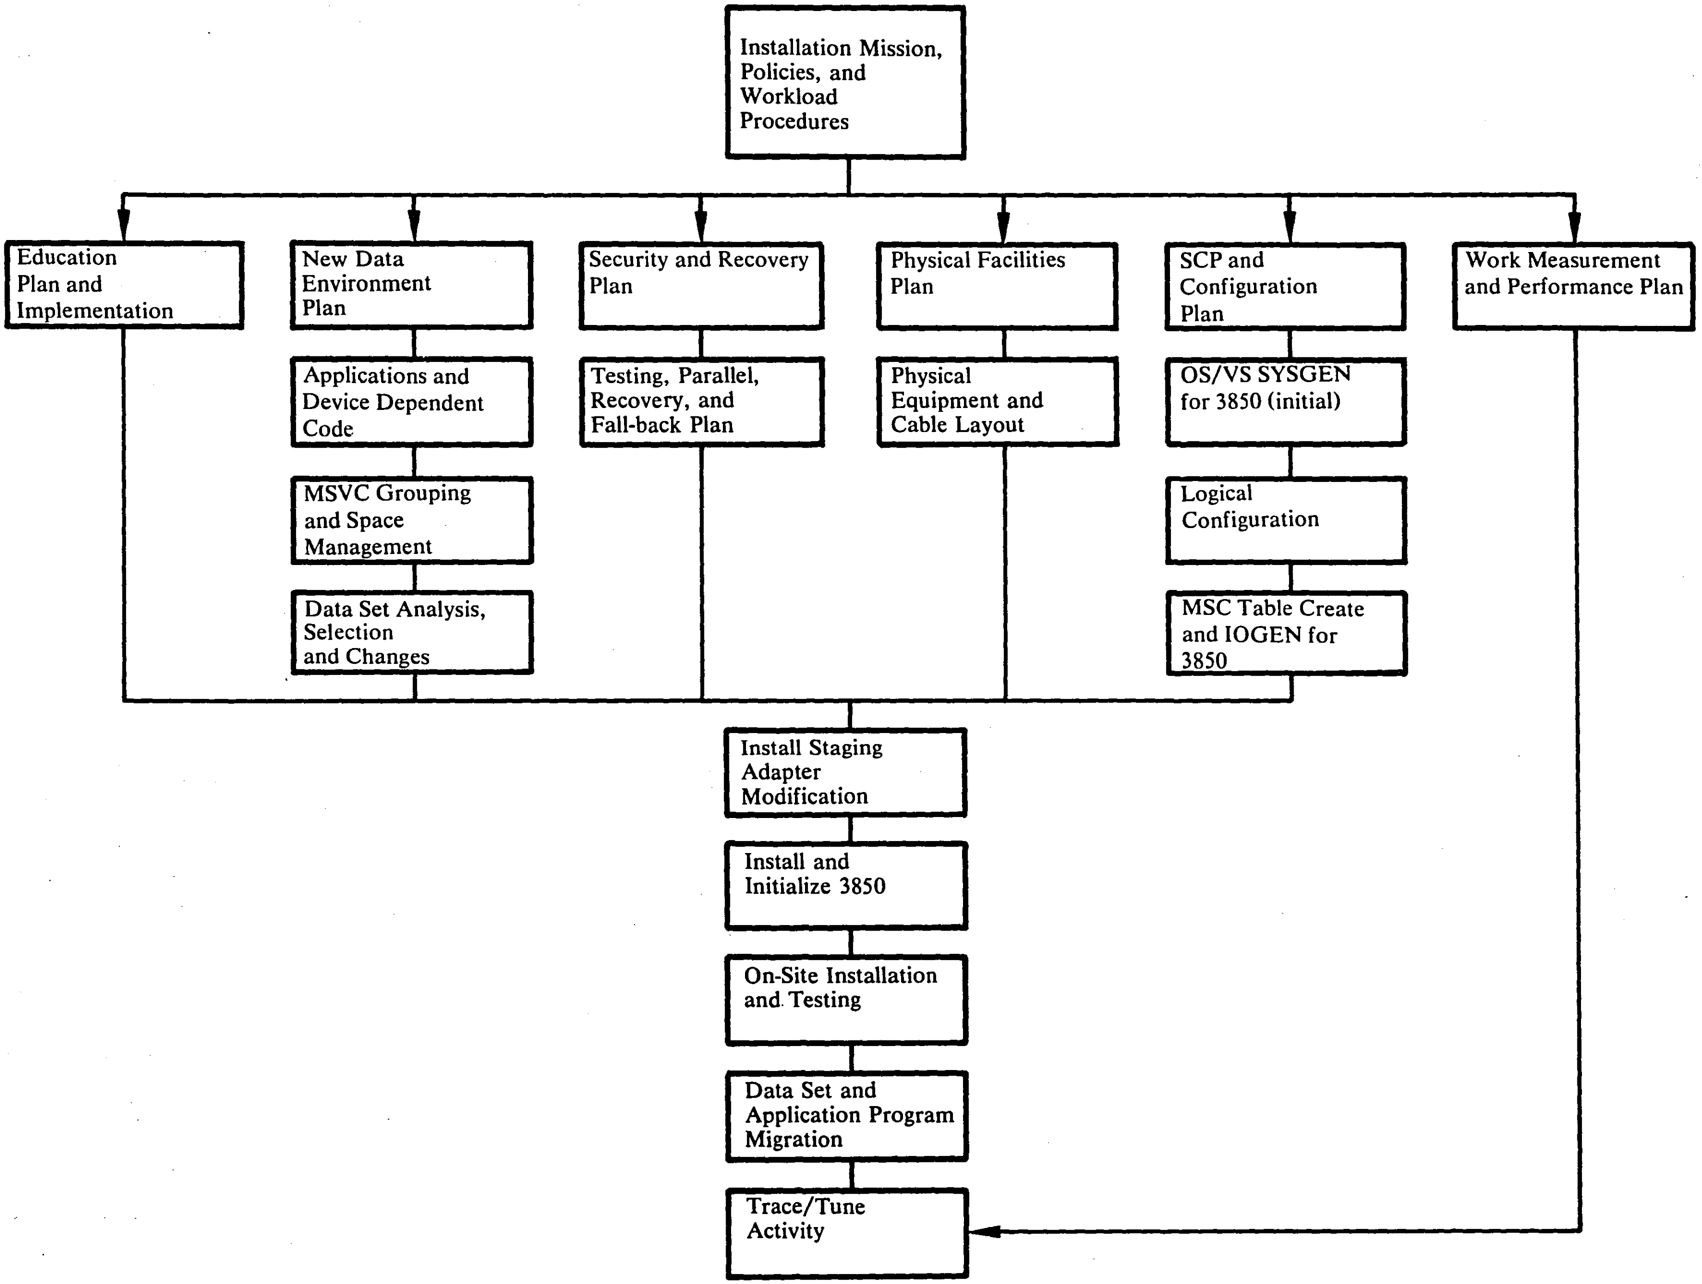

| **Figure 1. Planning and Installation Tasks**

Major Checkpoints

The following checkpoints summarize Mass Storage System installation planning. Adapting these steps into a documented installation plan is the key to a successful Mass Storage System installation.

- Project Organization

  After the Mass Storage System is ordered, management establishes project objectives, assigns a project team, and develops major checkpoints. Education starts now.

- Related Activities

  You should review the System/370 and OS/VS prerequisites for the Mass Storage System. Changes related to an SCP change are part of the integrated migration and conversion plan.

- General Planning

  Review the present data discipline and begin making the required changes. Application programs and related data sets should be identified. Conversion of device-dependent programs begins.

- Detailed Planning

  Review the physical and logical Mass Storage System configurations before confirming the order for a Mass Storage System. At this time, you should make a detailed analysis of data sets, review of mass storage volume and VSAM attributes, review of MSVC functions and grouping strategy, determining migration methods and schedules, developing system data management procedures, and new security and recovery management methods.

- Preinstallation Testing

  This includes testing such things as data discipline changes, converted device-dependent programs, and ISAM to VSAM changes.

- Installation - Phase 1

  Review the physical facilities. Finish the major equipment relocations, cabling, addresses, and SYSGENs. Approximately 60 days before the Mass Storage System installation, install the Staging Adapters.

- Installation - Phase 2

  Assemble the Mass Storage System, cable and test the units. Run a final Mass Storage Control Table Create program and SYSGEN before initializing the Mass Storage System. Activate the Staging Adapters.

- Customer On-Site Testing

  Begin the on-site installation and testing. Complete the operator training, regression tests, system management of data, pilot and parallel application testing, and data copying.

- Data Migration

  Transfer data to the Mass Storage System during normal production. Copy data into the Mass Storage System during low production hours. Remove displaced equipment and document all changes.

- Normal Operations

  Use trace analysis to adjust the Tune LRU parameters. Make volume and space management changes as determined by your experience. Develop new and extended applications in DB/DC and text retrieval.

Planning a New Data Environment

The IBM 3850 Mass Storage System introduces a new data environment. The Mass Storage System is a hierarchical extension of DASD; its maximum configuration offers the equivalent of over 4700 DASD volumes. To control this amount of online storage requires you to examine the way you manage DASD space.

Some key DASD data management concerns, such as data set selection and placement, are not critical in a tape environment where one tape volume is usually dedicated to one data set. Tape data sets are sequential, making their placement in a tape volume irrelevant. Tape space management is not critical because space is usually managed by removing the label and "scratching" the volume.

Therefore, the Mass Storage System might bring a significant change to your data processing operation. System control of data requires you to reinforce existing procedures, and implement some new conventions and standards. As you plan your new data environment, closely examine the selection, conversion, and management of the data that is to reside in the Mass Storage System.

Data Set Analysis

To best select data sets for the Mass Storage System, study the data set characteristics for your current data environment. Evaluate such items as data interchange requirements, job processing 'windows,' and application data set interdependencies.

The following paragraphs describe parameters for you to consider when selecting data sets, sources of information about your installation's data, and considerations for the initial Mass Storage System data migration.

Characteristics

You should know these things about your current data sets:

- Data set name
- Access method(s) used
- Data set size
- Block size
- Frequency of use (both input and output)
- Number of unique uses
- Device type

Identifying data set characteristics is not easy; the names of uncataloged data sets might not be unique. Data set size is important for space allocation in the Mass Storage System, and a given data set's size can vary greatly between uses, because the same data set might be referenced by a different access method each time. The large number of data sets in most installations can make identifying their characteristics laborious.

Much of this data set information can be made available by using the System Management Facilities (SMF). SMF recording is available in OS/MFT, OS/MVT, OS/VS1, OS/VS2 SVS and OS/VS2 MVS. The following example assumes that you are using SMF. To capture SMF Types 14 (input data set) and 15 (output data set) records:

- SMF must be generated in the operating system.
- MAN=ALL, OPT=2, DSV=2, or DEV=3 must be specified at IPL time or be specified in the SYS1.PARMLIB as a default member.

Include the "typical" processing cycle in the SMF sampling period; SMF reports information on data sets that are processed during this period. The information collected during the sampling period can be used to process the SMF data and prepare a data set characteristics report. See *OS/VS System Management Facilities* for SMF Types 14 and 15 record descriptions and preparing report programs.

"Where Used"

Because transferring data sets to the Mass Storage System is easier to plan and control by their application, a 'where used' list is a valuable tool for planning the Mass Storage System data migration sequence. This list should identify the jobs that use specific data sets, and help you determine application interdependencies.

A 'where used' list can be prepared from SMF Types 14 and 15 records sorted by job name and job end time, within the data set name. If job names are not unique, job end time can help you identify specific data set uses.

Another way to identify unique jobs is by using the job accounting information if it is available in the OS/VS JCL. SMF Types 4 (step termination) and 5 (job termination) records provide this information. Be sure you identify input and output uses in each case.

You can use a cross-reference list of data sets used by job as a checklist during the data migration process. To make a cross-reference list, sort the 'where used' listing into a different sequence; that is, data set name within job name and job end time.

SMF supplies information only for the data sets that are closed during the sampling period. Other data set information can be found in PROCLIB listings, librarian control logs, and system catalogs. Be sure to include any new applications that you are developing.

Instead of writing your own SMF data reduction programs, you may find it convenient to obtain licenses for two IBM Installed User Programs, the SMF Job Data Compression Aid (JDCA), Program Number 5796-PHN, and the Storage Device Migration Aid (SDMA), Program Number 5796-PHP. These programs provide the materials outlined above. In addition, they contain routines for data set filtering and for the production of summary reports that will be helpful in tracking the progress of your conversion effort after the planning phase is complete.

Data Set Selection

Data sets must be selected for both initial testing and production. A pilot data migration can test procedures, changes, and the Mass Storage System.

For Initial Testing

To test initial data migration to the Mass Storage System, select data sets that are not critical. Be aware of job processing time requirements that must be met if the test fails. Make the initial data migration workload large enough to test the Mass Storage System thoroughly. Use a representative sample of your data environment that includes your access methods, data set sizes, and data set usage. If possible, exclude data sets that have application interdependencies. To limit the number of new procedures initially, do not choose data sets that are used for data interchange or those that require removal from the Mass Storage System. Instead, try to find data sets that are in a stable application.

For Production

Use what you know about your data sets and processing requirements to select data sets for the Mass Storage System. All standard OS/VS access methods can be used with the Mass Storage System. The staging and destaging attributes for each access method are described in the *Introduction to the IBM 3850 Mass Storage System* publication.

Basic Sequential Access Method (BSAM)

If you plan to transfer tape data sets to the Mass Storage System, BSAM is the access method that you need. If installation jobs specify DISP=MOD in the OS/VS JCL for extending a BSAM data set, only the mass storage volume cylinders previously written for that data set are staged when the data set is opened.

For sequential data sets that reside on multiple mass storage volumes, only the data set extents on the initial mass storage volume are staged. Data set extents on succeeding mass storage volumes are staged at end-of-volume time. Remember, mass storage volumes are in a 3330-1 format, and tape data sets that previously required multiple volumes can now reside on a smaller number of mass storage volumes.

Indexed Sequential Access Method (ISAM)

ISAM data sets can be processed by the Mass Storage System but only in the cylinder fault mode. The ISAM data sets you plan to transfer to the Mass Storage System should be converted to VSAM and processed using the ISAM/VSAM interface. Conversion of ISAM data sets to VSAM is relatively simple, and is described in the *OS/VS Virtual Storage Access Method (*VSAM*) Programmer's Guide.* Be sure to identify and convert all uses of each data set.

Note: Using the BIND attribute for a mass storage volume that contains an ISAM data set does not bypass staging in a cylinder fault mode. The data set is staged cylinder by cylinder and bound as each cylinder fault is processed. The staged cylinder data remains bound until the volume is demounted.

Basic Partitioned Access Method (BPAM)

When a BPAM data set is opened, the entire data set is scheduled for staging, even though access to it is usually for only one of its members. Investigate its usage and potential staging activity if you select BPAM data sets for the Mass Storage System.

Basic Direct Access Method (BDAM)

BDAM data sets are staged when opened. Multivolume BDAM data sets are completely staged before processing continues. If the selected staging drive group does not have enough space available at staging time, the Mass Storage Control validates the unstaged pages and stages them when they are accessed.

Execute Channel Program (EXCP)

Data sets accessed using EXCP are scheduled for staging when opened. If you specify DSORG=PS, you must maintain the device-dependent portion of the data control block so that the System Control Program can write a file mark for output data sets. This portion includes the track balance (DCBTRBAL) field that contains a two-byte binary number indicating the number of bytes remaining on the current track. It also includes the full disk address (DCBFDAD) that indicates the location of the last cylinder specified in the current record. The address takes the form MBBCCHHR. All of these fields are written in the Format 1 DSCB, and are used for staging by the OPEN routines. Staging is done up to and including the last cylinder specified by these fields if the data set is reopened for OUTPUT, INOUT, or OUTIN.

Virtual Storage Access Method (VSAM)

VSAM provides the greatest number of staging and destaging options for the 3850. VSAM data set attributes can specify staging the entire data set when opened, staging in cylinder fault mode, or staging with BIND. By using the VSAM DESTAGEWAIT option, destaging can be synchronized with program execution. Different data set attributes can be individually assigned to both the index and data elements of a VSAM key-sequential data set. VSAM attributes are in addition to those for mass storage volumes. The Access Method Services CREATEV and MODIFYV commands record these attributes in a VSAM user catalog.

All IBM SCP access methods other than ISAM, issue an Acquire order when a data set is opened. Any other data organization, such as OS CVOLS, operates in cylinder fault mode.

Size and Usage

Size alone is not reason enough to select a data set for residence in the Mass Storage System; you must also consider its usage. Online data sets used in DB/DC applications should probably remain on real disks to satisfy response time requirements. However, data sets to be used for slow-response online applications may be appropriate choices for Mass Storage System residence.

Data sets that occupy less than one 3330 cylinder and are heavily used can be more effectively stored on real disk because:

1. The cost to store them on real disks is reasonable,
2. They contribute to higher cartridge delivery rates in the Mass Storage Facility, and
3. They are often uniquely used for tables or TSO data sets.

If your objective is to eliminate mounting as many tapes and disks as possible, the ratio of data set size to usage can help identify prime data sets for transfer to the Mass Storage System. The higher the ratio, the better candidate the data set is for transfer to permanently mounted real disk or to the Mass Storage System. Transferring these prime data sets reduces operational costs and the amount of tape and DASD equipment required to process this data. This analysis may convince you that it is best to use real DASD for a lot of small data sets.

Unless automated access to them is necessary, large multivolume sequential files belong on tape. Large sequential data sets that have been converted to VSAM KSDS (key sequential data set) offer two operation options: skip sequential and current sequential.

Intersystem Usage

You probably use several of your data sets for more than one job and application. Therefore, when you select data sets for the Mass Storage System, examine all of their uses because the required OS/VS JCL and program changes must be made for each use. Data set 'where used' information helps you identify interrelated data usage. The transfer of a specific data set to the Mass Storage System might not be possible until its related applications are scheduled for production.

Interchange and Off-Site Storage

Installing the Mass Storage System does not change your data interchange and off-site storage requirements, which might dictate that specific data sets remain on tape for compatibility purposes. Tape is still the standard medium for data interchange. Off-site data storage usually provides data backup. If you plan to store data cartridges off-site, consider arranging access to a backup system on which to process them.

Data Set Placement

The tasks involved in transferring tape-resident data to mass storage volumes are similar to those for moving it to real DASD volumes. However, there is a difference when considering data set placement within a volume. Placement in a real volume achieves I/O overlap and improves DASD performance. Placement in a virtual volume has no effect on I/O overlap because the data set on a mass storage volume (not a bound volume) can be staged to one or more staging drives within a staging drive group. I/O overlap is determined by a staging algorithm.

Data sets staged at OPEN time and those staged by cylinder fault can be placed on the same mass storage volume. However, you should understand the difference in operation of and the amount of time that the unit control block (UCB) is busy between a Start I/O command to a real drive and a Start I/O that cylinder faults to a virtual drive. A Start I/O command issued for data on a real drive is queued on the UCB until all previous Start I/O commands have been executed. The UCB is busy for each Start I/O command until the requested data is transferred to the CPU.

A Start I/O command issued for data on a virtual drive functions in the same manner as on real drives if the data has been staged. If a cylinder fault occurs during the Start I/O, the operation and the time the UCB is busy depends on the cylinder fault option specified in the CREATE command of the Mass Storage Control Table Create program. If the option is:

- CYLFAULT=YES, the cylinder fault request is queued but the UCB is available to other requests in the queue while the cylinder fault is in process. When at least one cylinder fault for that volume has been satisfied, the Staging Adapter notifies the CPU by an attention interrupt.

All cylinder fault requests are then placed on the associated logical channel queue to be retried.

- CYLFAULT=NO (default), the cylinder fault request is placed in channel command retry and the UCB remains busy until the cylinder fault is satisfied and the data transferred to the CPU. While the UCB is busy, the CPU cannot access any other data on the staging drive until the cylinder fault is satisfied and the UCB becomes available.

In a multihost environment, the CYLFAULT=NO option must be used until all attached hosts have installed the Selectable Unit (SU) that supports cylinder fault processing with attention interrupts. The SU can be installed one host at a time and the option changed to CYLFAULT=YES when all hosts have had the SU installed.

Volume and Space Management

Volume and space management requires a close study of several considerations before you plan and implement a grouping strategy for your installation. Although the major considerations are summarized here, you should read the chapter "Considerations for the Use of Groups" in the *OS/VS Mass Storage System (MSS) Services: General Information* publication. It describes in more detail the allocation to mounted volume support at Early Allocation Exit (before the unit is selected) and Late Allocation Exit (after the unit is selected) times.

Volume Attributes

Most real volumes remain mounted during production hours, and their volume attributes can be examined, assigned, and deleted. The Mass Storage System, however, can issue many mounts and demounts of mass storage volumes during normal operation as data sets are accessed. Therefore, you must have a good working knowledge of volume attributes because they change as new volumes are mounted. For details concerning volume attributes, see *OS/VS2 Job Management* or *OS/VS1 JCL Services*.

Do not confuse volume mount and use attributes with mass storage volume attributes. Volume attributes include the mount attributes (permanently resident, reserved, and removable) and the use attributes (private, public, and storage) and are assigned when volumes are mounted. Mass storage volume attributes (BIND, DASDERASE, EXCLUSIVE, and READONLY) are assigned to individual mass storage volumes when they are created (CREATEV), and are permanent unless changed (MODIFYV) by the Access Method Services commands. Remember, the volume mount and use attributes that apply to real volumes also apply to mass storage volumes.

There are three methods for assigning the volume mount and use attributes:

- VATLST or PRESRES
- Operator mount
- Job allocation

VATLST or PRESRES defines the mount and use attributes for selected volumes to be mounted at IPL time. When VATLST or PRESRES has defined the mount and use attributes for all real DASD volumes to be mounted, the SCP issues a command to mount all of the mass storage volumes in the list, and to assign the mount and use attributes.

It is important that you understand how volume mount and use attributes affect the demounting of mass storage volumes. If data on a mass storage volume is not relinquished with the DESTAGEWAIT option, it is not scheduled for destaging until the mass storage volume is demounted. OS/VS2 MVS demounts a private and removable volume when its use count is 0 and either the job has terminated or another volume needs the disk drive. OS/VS1 demounts a private and removable volume when its use count is 0 and the job has terminated. OS/VS2 MVS and OV/VS1 demount a public and removable volume when its use count is 0. Tape volumes default to private and removable. However, if a tape volume is converted to a mass storage volume, the default is to public and removable if new specific requests are used.

With or without the allocation to mounted volume support, mass storage volume group (MSVGP) requests are assigned a volume use attribute of private. At Early Allocation Exit time, MSVGP requests are allocated to volumes mounted as private and reserved, private and removable, or with no mount or use attributes. Allocation to mounted volume affects the use attribute of a removable volume as follows:

- An MSVGP request is treated as a private request.
- A nonspecific volume request without MSVGP or VOL=PRIVATE specified is selected from SYSGROUP and is treated as a public request.
- A new specific request (other than VOL=REF to a MSVGP request) without the VOL=PRIVATE specification is treated as a public request.
- An old request or a new VOL=REF (referring to a MSVGP request) with or without VOL=PRIVATE specification is treated as a private request. It is treated as a public request if VOL=PRIVATE is not specified in a request with a mixed esoteric name specified, or if one or more of the requested volumes are mounted as public.

With OS/VS1, a reserved volume can be changed from public to private by the Mount command without unloading the volume. When this occurs, the volume is eligible for allocation at Early Allocation Exit time without any information in MSVC about the space that had been used when the volume was public. The space information is updated if the reserved volume is unloaded when its use attribute is changed from public to private.

Two examples of volume state assignments to mass storage volumes follow:

- Storage volumes can be mounted as storage-reserved or storage-permanent by using VATLST or PRESRES. The operator-issued Mount command can also be used to assign the storage-reserved attributes. By having storage volumes mounted, temporary data sets allocated in cylinder increments do not require mounting, staging, or destaging. Only the cylinders allocated on the storage volumes occupy staging drive space, which is eventually deallocated by the LRU alogrithm or by demounting the storage volume.
- Storage volumes can be mounted as private and removable so that MSVGP requests can be allocated to them. These attributes are assigned by the operating system for all virtual requests except mixed esoteric names and new specific requests.

Volume Selection By User

You may select a 3330V volume by making a specific request, or defer the selection to the Mass Storage Volume Control (MSVC) by making a nonspecific request with or without the MSVGP specification. The space on volumes allocated by specific request is usually managed very closely. It is impractical and costly to manage virtual volumes the same way you would real volumes; there are too many. The MSVC functions contained in the Mass Storage System Communicator MSSC modules enable you to make nonspecific requests and to manage the volumes by groups.

Volume Selection By MSVC

When a nonspecific request is made, MSVC selects a mounted volume (if a candidate is available) at Early Allocation Exit time, or an unmounted volume at Late Allocation Exit time.

To select a mounted volume at Early Allocation Exit time:

- MSVGP must be other than SYSGROUP
- The MSVGP parameter CONCURRENTUSERS must be nonzero
- UNIT must be 3330V (or an esoteric name with only 3330V or 3330V unit address).

The objective of the MSVC volume selection algorithm at Early Allocation Exit time is to apply the following criteria against all general use volumes in the group in order to find the most suitable volume. These criteria include:

- High risk of job failure
- Potential job wait
- CONCURRENTUSERS value
- Retention period
- Sharability
- Mounted or scheduled for mounting on the requesting host
- Mounted to another host
- Unmounted

If UNIT=3330V is used, the probability of selecting a suitable volume is improved. If UNIT= esoteric name or 3330V unit address is used, then volumes mounted to another host must be considered as high risk volumes because they may not be available to the requesting host. The volumes mounted to the requesting host must be considered as wait volumes because the requested job may be delayed.

If a suitable mounted volume is not available, an unmounted volume is selected. The nonspecific request then becomes a specific request to Allocation. If the volume selected is not suitable to Allocation, the job fails.

The Late Allocation Exit selects 3330V volumes for nonspecific requests with:

- The MSVGP parameter CONCURRENTUSERS set to zero
- MSVGP=SYSGROUP
- UNIT= mixed esoteric name
- No MSVGP group name (selected from SYSGROUP)
- MSVGP and VOL=PRIVATE (OS/VS1 only)
- Secondary volumes of a multivolume MSVGP data set at EOV

The objective of the MSVC volume selection algorithm at Late Allocation Exit is to serially search the unmounted general use volumes in the group for one with sufficient space. The volume selected for the nonspecific request is then mounted:

- Private and removable
  1. From the MSVGP if specified
  2. From SYSGROUP if MSVGP is not specified and VOL=PRIVATE
  3. From SYSGROUP, if neither MSVGP nor VOL=PRIVATE is specified, and a mounted public or storage volume is not available for a permanent data set
- Public and removable from SYSGROUP if MSVGP is not specified and a mounted public or storage volume is not available for a temporary data set.

If the selected volume is not suitable for allocation, the volume request is repeated until an acceptable volume is selected.

In the JES3 environment, a final check is made of the MSVC volume selection. If the selection is not acceptable to JES3, the MSVC selection is rejected. This occurs if the volume is:

- Set up in a staging drive group exclusive to another host
- Set up for exclusive use by another job
- Required by a "barrier" priority job waiting to complete the set up
- Made unavailable by a JES3 operator command
- Not recognized because the unit on which the volume is mounted or the host to which it is mounted, is not managed by JES3
- Mounted to another host and JES3 does not have a virtual unit address or a path available to the staging drive group.

If JES3 rejects the volume, MSVC presents its next suitable volume. If the final selection is not suitable for allocation, the job fails.

Volume Grouping

Mass storage volume grouping helps you manage mass storage volumes. Its major advantages permit:

- Grouping a number of mass storage volumes that have the same characteristics under one name; that is, MSVGP=groupname. A large number of mass storage volumes can then be divided into manageable units.
- Collecting data concerning a group of volumes for control and reporting purposes.
- Separating sensitive data or special uses of data.

Real volumes can be allocated by specific requests because they are usually few enough to be individually managed. As you obtain more volumes, it becomes increasingly difficult to maintain space and volume serial numbers for specific allocations. Volume grouping makes many of the new data set requests nonspecific. The allocation routine then searches for a group of volumes instead of only one.

There are two types of volumes within groups: restricted-use and general-use. If a volume is assigned to a group without being specified as general or restricted-use, the volume defaults to general-use.

The MSVC can select a general-use volume from a group when a nonspecific volume request is made to the group. The initial nonspecific request is usually made by specifying DISP=(,CATLG) and MSVGP=groupname. Subsequent allocations to this data set are specific requests that use the catalog.

Specifying restricted-use indicates that the mass storage volume cannot be selected by the MSVC function for nonspecific requests. New data sets must be specifically allocated to restricted-use volumes with the VOL=SER parameter or the catalog. Restricted-use volumes can be used to separate sensitive, archival, and special use data. This separation permits a volume to be part of a group for control and reporting purposes, but reserves its use for other purposes.

Grouping Strategies

Before group assignments can be made, you must decide whether 3850 space management is to be centralized or decentralized. Centralized space management means that people in a central control area manage all groups of mass storage volume space. These people interface with the various user departments to establish and assign groups, run reports, and manage the space on their respective mass storage volume groups. Centralized control offers better control by managing space across all volumes using the same rules. However, its practice may be contrary to your organizational structure and user departments might want to manage their own data and volumes.

Decentralized space management means that user departments or functions manage their own groups of mass storage volume space. In this case, the central control establishes and assigns the groups to the users. The users then assume the responsibility of managing the data and the volumes in the group. This responsibility can extend from indicating changes in volume status reports to complete space management using the available system support. Problems that the user department can not resolve are taken to the central control. A disadvantage of decentralized space management is that space can be wasted.

Some common methods of grouping are mentioned here; yours may be different. Although each method is described separately, this does not preclude combining methods for a better grouping strategy.

By Application

This method places all major applications in unique groups. Several small applications can be placed into a common group. All data sets are divided into specific application groups such as payroll, billing, and inventory. This strategy lends itself to decentralized control because groups are separated by application and, consequently, by user department. It would then be natural to delegate data or space management responsibility to the application areas for their respective groups.

Advantages:

- Definition is usually easy.
- Reporting group space usage and volumes status for control reasons is simpler because each group represents a unique application.

- Group standards and specifications might be easier to establish within an application area.
- All data sets related to a given application can be easily transported.

Disadvantages:

- Data set size can vary considerably, which wastes space and causes fragmentation.
- Each application may require subgroups with different group parameters. This increases the number of groups and the number of total volumes required.

By Data Set Size

This method assigns each data set to a group, based on its size. This strategy lends itself to centralized control because each group contains data sets from various user departments.

Advantages:

- Space fragmentation is minimized because data sets that have the same primary and secondary extents are allocated to the same group.
- The total number of groups is probably reduced because the subgroups by data set size that might be required with other grouping strategies are consolidated.

Disadvantages:

- Reporting is difficult because multiple functions and applications share groups. Data set naming conventions are required.
- Managing and/or transferring applications is difficult because the applications are contained in many groups.
- You must categorize all data sets by the space they require. Executing each job will probably require several different groups.
- Data set size can vary within processing cycles.

By Retention Period

This method categorizes all data sets by their retention periods, and assigns groups to each category. For example, you can separate all data sets into retention periods of 5, 30, and 60 days, and assign GROUP05, GROUP30, and GROUP60 respectively. As described in *OS/VS Mass Storage System (MSS) Services: General Information*, CREATEG and MODIFYG commands are used to define and modify group retention periods.

New volumes that have been assigned to a group, have not yet been assigned retention periods. As they are selected to satisfy allocation requirements, MSVC assigns the respective group retention period to the volume. MSVC searches for a volume that has a retention period equal to or greater than that of the data set that is being allocated. A new volume can be selected if the group retention period is equal to or greater than the data set retention period.

A volume expires when its retention period has elapsed, and all data sets contained on an expired volume are considered to be expired. The volume can be specified as scratch, and reassigned a new retention period. However, if specific allocations are made to mass storage volumes in a retention period group, unexpired data sets might be present on an expired volume because retention period checking is bypassed. The MSVC does not check retention periods for specific allocations. However, for nonspecific allocations, retention period is one of the conditions considered by MSVC in selecting volumes to satisfy the request.

At Early Allocation Exit time, a mounted volume can be selected as a primary volume if it has enough space for the primary extent (rounded to the nearest cylinder) after the reserved space has been subtracted from the running count. If the volume is selected as a secondary volume (multivolume request), then it must have enough space for the secondary extent (rounded to the nearest cylinder) without considering the reserved space. If the volume meets all selection requirements, including retention period, allocation is complete.

At Late Allocation Exit, MSVC sequentially searches a logical list of volumes in GROUP05 for a retention period equal to or greater than that of the data set. GROUP05 may have volumes with retention periods of 5, 4, 3, 2, 1, and 0 days. When it finds a volume, MSVC checks for free space to satisfy one primary extent and 15 secondary extents. If the volume meets all selection requirements including retention period, allocation is complete. If not, the next volume in the group is selected. If a volume with enough space for the primary and secondary extents can not be found, MSVC searches for a volume that can hold only the primary extent.

It is possible for MSVC to select all new volumes in a group on the same day. If this happens, they all have the same retention period, and all expire at the same time, making it difficult for MSVC to satisfy allocations to this group. Assigning groups by retention periods lends itself to centralized control because user department data sets are allocated across many groups.

Be aware that the retention period parameter (LABEL=RETPD=nn) can cause a data set to be kept on an abnormal step termination if DISP=(NEW,CATLG,DELETE) is coded and the retention period has not expired. This action is the same for 3330 or 3330V volumes. See *OS/VS1 JCL Services* or *OS/VS2 JCL* for information about how to use these parameters.

Advantage:

- Space management is easier because as the volume expires, all data sets on the volume also expire and can be deleted, assuming that specific allocations are not allowed.

Disadvantages:

- Reporting is difficult because multiple functions and applications share groups. Data set naming conventions are required.
- Managing and/or transferring applications is difficult because the applications are contained in many groups.
- You must categorize all data sets by the retention period. Executing each job will probably require several different groups.
- Deleting an entire volume when it expires requires strict standards to which everyone must adhere.

By Function or User Department

This method categorizes all data sets by function or user department. Groups are then assigned to these categories. Several groups may be assigned to one department, each having subcategories determined by the department. This strategy lends itself to decentralized control because the responsibility for maintaining the groups can be easily defined to the user department.

Advantages:

- Definition is usually easy.
- Reporting group space usage and volumes status for control reasons is simpler because each group represents a unique application.

- Group standards and specifications might be easier to establish within an application area.
- All data sets related to a given application can be easily transported.
- More responsibility for maintaining the group can be delegated to the department or functional level than to the application level.

Disadvantages:

- Data set size can vary considerably, which wastes space and causes fragmentation.
- The number of subgroups required for each function or user department may cause the total number of groups required to be unnecessarily high.

By Special Areas

This method defines groups for special data sets, such as those for security, off-site storage, shared, or exclusive use. Each customer has his own specific area for specialized group control.

Security groups might consist of a set of volumes separated from the others because they contain confidential data. Perhaps some volumes are taken off-site frequently; these can be grouped to make moving them in and out of the Mass Storage System easier. Certain other volumes, especially those containing tape data, should be accessed by only one CPU at a time; these can be placed in a group that has the exclusive mass storage volume attribute specified.

Usability Guidelines

OS/VS JCL users should be aware of the following guidelines when using the MSVGP parameter for groups other than SYSGROUP. These are not requirements but supplements to the MSVC algorithm, which increase the probability of successfully selecting a suitable volume.

Space Fragmentation

If space fragmentation is not controlled, MSVC might select a volume with sufficient but fragmented space. DADSM must satisfy the primary extent with no more than five noncontiguous extents. To reduce fragmentation consider using:

- Group retention dates
- Group space defaults
- Common or multiple primary and secondary cylinder space allocation. For example (4,4) or (8,4).

MSVC Space Accounting

MSVC maintains a running count of available space for each general-use volume in a mass storage volume group other than SYSGROUP. MSVC does not account for space used by:

- Specific or VOL=REF to a specific request that allocates new data sets on a general-use volume
- VSAM space allocated by the Access Method Services DEFINE command and extended by VSAM EOV.

The RESERVEDSPACE parameter, used for secondary allocation and VOL=REF requests for a new mass storage volume, should be planned accordingly. The default is 40%.

Using Access Method Services

Do not execute the Access Method Services commands to create, copy, or rename general-use volumes in a mass storage volume group while jobs being executed are selecting volumes from the same group. An incompletely created, copied, or renamed volume may be selected for a mass storage volume request and fail in Allocation.

Multihost Environment

Because the activity of one host is unknown to another host (except with JES3), care should be taken in a multihost environment when two or more hosts access the same mass storage volume group concurrently.

To allow for space accounting discrepancies that may occur, reserved space should be planned in order to provide a space reservoir. Space discrepancies may occur because of the elapsed time between changes in the available space count and the actual allocation and deallocation of space in two attached hosts.

For example, the use of Access Method Services commands in host A to a mass storage volume group while host B is making nonspecific requests to that group may cause problems. A SCRDSET GROUPA, UNCATALOGED command in host A can delete the new data set being created in host B if the Access Method Services command is executed before the data is cataloged.

The allocation to mounted volume support assumes that the host requesting space in a mass storage volume group has a path, both logical and electrical, to the staging drive group where the volume is mounted. However, a mount conflict may occur. Assume that volume 10, selected at Early Allocation Exit time for host A, had been previously selected for host B but the Mount command has not been executed. If the mount to host B occurs first on a nonsharable unit or the volume is mounted with the EXCLUSIVE mount attribute, the job in host A fails.

OS/VS1 Users

To achieve the maximum benefit from the allocation to mounted volume support, OS/VS1 users should avoid the use of esoteric names with the MSVGP parameter. OS/VS1 does not utilize data reuse for esoteric name requests or for multivolume data sets.

Information Sources

Regardless of the space management strategy you use, you might have to list VTOC information for mass storage volumes. VTOC information can be obtained using IEHLIST or the Access Method Services LISTDSET command to determine which data sets should be "scratched". IEHPROGM or the Access Method Services SCRDSET command can be used to delete data sets on mass storage volumes using UNIT=3330V.

Although IEHLIST and the LISTDSET command provide vital data set information, most of the space management information comes from data that is unique to the Mass Storage System. *OS/VS Mass Storage System (MSS) Services: General Information* describes the programs that report information contained in the Inventory and Journal data sets. Also, use catalogs with mass storage volume groups because the number of virtual mass storage volumes within the Mass Storage System is large. Well-designed catalogs contribute to system performance, and are another information source. If you have organized your mass storage volume groups by application or function, a catalog for each mass storage volume group, or series of groups, is a reasonable structure.

Standards and Procedures

When you establish a set of standards and procedures, do it early and include methods for changing them with minimum impact on your regular production.

Catalog Structure and Usage

After transferring tape data sets to the Mass Storage System, there is no longer a one-to-one relationship of data sets to volumes. Catalog structure and usage applies to the Mass Storage System the same way it does to real DASD, except that more data is under system control. OS CVOLs and VSAM user catalogs help manage the data capacity provided by the Mass Storage System.

Consider data set naming conventions for improving control of application data sets, and future guidance to application programmers. Using qualified data set names (such as user department or application) and cataloging facilities provide the most flexible means of uniquely defining and tracking data.

A standard for assigning volume serial numbers is valuable because the Mass Storage System enforces the use of unique volume serials for all mass storage volumes by allowing only one mass storage volume with a given serial number to be active at a time. Any other mass storage volume with the same serial number is automatically placed on a transient volume list, which means that it cannot be mounted. This restriction applies whether the configuration has one or two Mass Storage Facilities, and whether the mass storage volumes are accessible by only one, two, three, or four hosts.

The *OS/VS2 Conversion Notebook*, and *OS/VS1 Planning and Use Guide*, provide additional considerations for data set naming, catalog structure, and catalog usage.

Job Control Language (OS/VS JCL) Parameters

Analyzing OS/VS JCL usage for data sets and applications helps you define OS/VS JCL standards. Consider the use of catalogs, VTOC, and mass storage volume grouping to minimize the external OS/VS JCL parameters required for job execution. OS/VS JCL standards make the parameter's interpretation and conversion easier, by either manual or automated methods.

UNIT

The UNIT parameter provides device information for OS/VS allocation routines. In a Mass Storage System environment the options are:

- UNIT=3330V, or a virtual unit address
- Esoteric names that have same-as-current usage

The UNIT parameter is a prime candidate for using symbolic subparameters, which simplify future data transitions. For more information about symbolic subparameters, see *OS/VS1 JCL Services* or *OS/VS2 JCL Reference Manual.*

DISP

In a DASD environment, the DISP parameter controls data set disposition in normal and abnormal job step execution. Consider using DELETE as a disposition subparameter for the Mass Storage System.

Using DELETE as the second disposition subparameter in a tape environment has no effect on the data but does inform the operator to delete the tape. However, such tapes may be retained in your tape library because manual procedures override the specified disposition in special circumstances. When tape data sets are moved to the Mass Storage System, they are deleted if the DELETE subparameter is used. Using DELETE as the third disposition subparameter when a data set is created controls data set disposition for an abnormal job step execution.

SPACE

If the SPACE parameter is specified for tape data sets, the system ignores it. Specifying this parameter in current OS/VS JCL during preinstallation activity might help simplify the conversion to the Mass Storage System.

To minimize unnecessary staging and destaging operations, specify data set space allocation in cylinders. Space allocation by number of blocks, using the ROUND option, maintains device independence. Use of the RLSE option is not recommended because it tends to increase space fragmentation and does not relinquish staging drive space.

Space may be specified for data sets on mass storage volumes even if groups are defined. If space is not specified by the OS/VS JCL, a new data set added to a volume in a group uses the group default space allocation.

MSVGP

This parameter is used to request space for new data sets on a general-use mass storage volume in a specified group, that is MSVGP=groupname You can request a specific volume either by VOL=SER or have MSVC select one by using MSVGP, but not both. The MSVGP SPACE parameter is intended for use with UNIT=3330V, a specific 3330V unit address or an esoteric name containing 3330V.

If the MSVGP parameter is added to the OS/VS JCL for a 3330 unit prior to data migration, a real volume mounted on that unit is marked as private. It will not be used for selecting general-use mass storage volumes until the unit is changed to 3330V.

MSVGP = (groupname,ddname)

A new data separation parameter facilitates the separation of Master In and Master Out data sets. The MSVGP positional parameter is coded in the Master Out DD statement where the ddname identifies the Master In data set. The first OS/VS JCL example in the section "Generation Data Groups" illustrates the use of this new parameter. The two DD statements must be in the same job step. The data separation parameter is not supported by OS/VS2 dynamic allocation.

To force data separation from a Master In data set contained within a previous step, specify a VOL=REF request to the DD statement in the previous step and then specify data separation from the VOL=REF request.

Generation Data Groups

Generation data groups provide a method for managing cataloged data sets that are related by a common name, and are maintained in the sequence of their creation. Use generation data groups whenever possible because they can provide automatic space reclamation. If a data set within this group is defined with the NOEMPTY default and SCRATCH options, the oldest generation is uncataloged and deleted (if the volume is mounted) when the new generation is created. A DD statement referencing the oldest generation is required in the creation step to ensure that the volume is mounted for deletion. OS/VS2 MVS dynamic allocation does not support generation data groups. The following example shows the OS/VS JCL required:

```
//JOBONE  JOB
//DDREF DD DSN=xxx.yyy(-4),DISP=OLD  (oldest)
//DDIN  DD DSN=xxx.yyy(0),DISP=(OLD,KEEP)  (most  recent)
//DDOUT DD DSN=xxx.yyy(+1),UNIT=3330V,
           MSVGP=(AA,DDIN),DISP=(,CATLG)  (newest)
```

Data set xxx.yyy was created with five generations using the NOEMPTY and SCRATCH options. Because they are specific requests, DDREF and DDIN are allocated first. With allocation to mounted volume support, a volume in group AA, other than the volume for DDIN, is allocated. Data separation is achieved by the DDIN parameter in MSVGP=(AA,DDIN). The oldest generation referenced by DDREF will be uncataloged and deleted because of the NOEMPTY and SCRATCH options. The newest generation may be created on a volume referenced by DDREF if it is the most suitable volume in group AA.

Another technique is to rotate generation data groups among the same set of volumes. Assume that the space manager is requested to create a new structure for data to be placed in the Mass Storage System containing four generations. The space manager builds the appropriate catalog structure for the generation data group, selects two volumes to contain the four generations, and makes catalog entries as if the four generations existed. The intent is to assign the data set generations to the two volumes in a "flip-flop" manner as shown in Figure 2.

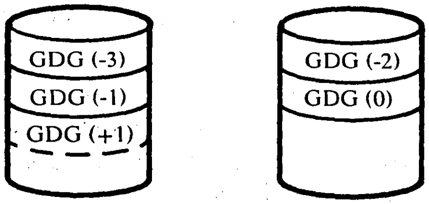

Figure 2. Rotation of Generation Data Groups

The OS/VS JCL required is:

```
//JOBTWO  JOB
//DDREF DD  DSN=xxx.yyy(-3),DISP=SHR,
            VOL=PRIVATE,UNIT=(,,DEFER) (oldest)
//DDIN  DD  DSN=xxx.yyy(0),DISP=SHR,
            VOL=PRIVATE  (most  recent)
//DDOUT DD  DSN=xxx.yyy(+1),
            DISP=(NEW,CATLG,DELETE),UNIT=3330V,
            SPACE=(CYL,(8,8)),
            VOL=(PRIVATE,REF=*.DDREF),DCB=(xxxx)  (newest)
```

The DEFER option in DDREF is required for the first three generation cycles to prevent the System Control Program from mounting the volume and searching for the oldest generation (-3) which has been cataloged but not created. The REF parameter in DDOUT ensures that the newest generation (+1) is created on the same volume that contained the oldest generation.

This rotation procedure ensures that successive generations are not on the same volume. It requires additional work to set up and sufficient space to contain an extra data set generation. The extra space is required because the newest generation is created on the referenced volume before the oldest generation is removed by the SCRATCH option. This technique has the advantage of better control because the volume location is known for all generations.

When a DD statement references all generations of a generation data group data set on multiple real volumes, the first volume is mounted and succeeding volumes are enqueued to the requesting host with UNITAFFINITY (OS/VS2 MVS only). Because mount conflicts can occur in a virtual volume environment when more than one program or host references the same mass storage volume group, the generation data group ALL support has been changed for 3330V units. When a DD statement references all generations of a data set on multiple volumes, all volumes are mounted and enqueued to the requesting host. Be aware that a waiting for volumes condition can occur in a multihost environment until the generation data group volumes are dequeued.

Procedure Libraries (PROCLIB)

When the same set of OS/VS JCL statements is used repeatedly with little or no change, consider using a procedure library. Procedure libraries save programming time, reduce the possibility of processing errors, and provide installation control of OS/VS JCL through the operating system facilities. Steps can be established in a procedure library to specifically test applications that are migrating to the Mass Storage System, without changing production OS/VS JCL.

Run Books and Procedures

Existing setup procedures must be revised for the Mass Storage System environment. For example, the tape librarian may not have to select tapes for job execution. Run books should include operator guidance for normal job execution, as well as restart and recovery action for abnormal job execution. Job processing procedures must be current and accurate. The Mass Storage System changes a people-managed data environment into a system-managed data environment, and requires that all actions, either normal or abnormal, be defined.

Program Conversion Considerations

Review application programs that are migrating to the Mass Storage System to determine if program changes are required to ensure proper execution. The factors to consider during the review are described below.

Data Control Block (DCB)

If DCB parameters have been specified in the program code, certain subparameters must be removed. For example, DEVD=TA, TRTCH, DEN, and certain parameter options in OPTCD must be removed because they are not valid for DASD. Most DCB parameters should not be coded in a program, but should be specified in external OS/VS JCL on the DD statement for the creation of the data set.

Device Dependencies

Programs that use disk-addressing algorithms might require changes. For example, if a BDAM randomizing algorithm depends on the 2314 architecture in terms of records per track, tracks per cylinder, or cylinders per module, the program must be changed.

In general, programs that run on 3330 DASD do not require modification for the Mass Storage System. If programs are written at the execute channel program (EXCP) level, make sure that the channel command word (CCW) strings do not contain device-dependent subparameters, such as write tape mark to tape device.

Any user-written EXCP for DASD must conform to IBM standards for Record Zero (R0). Any attempt to write the home address or count-key-data for R0 on virtual mass storage volumes will cause a command reject. If modifications have been made to IOS in the operating system, you must ensure that the first command of a command chain is a Seek command.

Other factors that require examination are sort exits, SYNAD routines, and label-processing routines. Tape-dependent logic in programs (for example, read backwards) must be modified. A program in which DEVTYPE macros are specified must be able to accept the 3330V device type. Emulator programs that work with real DASD will work for 3330V.

Consider program accounting or measurement routines. SMF records have been changed to reflect the 3330V device type. Programs that analyze SMF data might have to be changed. For more information regarding SMF records, see *OS/VS System Management Facilities SMF.*

Planning Examples

The two internal production installations summarized in this section were planned and installed several months before the first customer shipment of the Mass Storage System. Therefore, references indicating that MSVC does not select a mounted volume are valid. The examples have been retained to give you different views of the data planning process.

Example 1

Example 1 is for an istallation that uses OS/VS2 MVS. The installation took place before the final system and program testing for customer shipments were completed. During the first phase, 15-20% of the production workload was identified and isolated for extensive testing and debugging. The information presented here is based on the experience gained during this first phase, during which approximately 20% of the workload was scheduled for migration to the Mass Storage System.

Initial Conditions

The data processing department in this example is an administrative service center for a manufacturing complex. The center's program applications range from finance to manufacturing, and they process approximately 1000 jobs per day. Their tape library contains about 12,000 volumes and 3500 active data sets.

When the installation planning began, the system was running under OS/MVT with this equipment:

- Two System/370 Model 155s (two megabytes each)
- Twenty-eight 2420-7 tape drives
- Thirty-two 3330-1 disk drives
- Miscellaneous unit record devices (2540, 3211, 1403)

In addition, a third Model 155II processor (two megabytes) had been installed to support miscellaneous testing requirements. This CPU was selected for the installation of OS/VS2 and the Mass Storage System. Production work was transferred from the other two systems for initial testing and eventual production. This approach was selected because the Mass Storage System and OS/VS2 MVS Release 3 were installed several months before the first Mass Storage System customer shipment. This equipment was attached to the Model 155II:

- One 3851-B2 Mass Storage Facility
- Two 3830-3 Storage Controls
- Two 3333-11 Disk Storage and Controls (four drives)
- Two 3330-11 Disk Storage units (four drives)
- Four 2420-7 Magnetic Tape Units

During the planning phase, these areas were analyzed:

- Data sets, for placement in the Mass Storage System
- Volume and space management techniques
- Standards and procedures
- Conversion techniques

Data Set Analysis

A detailed analysis of data set usage was done using data set characteristic and "where used" reports. The information sources used to prepare these reports included an automated scheduling system, OS/VS JCL (PROCLIB), and a two-week sample of SMF data. The results for tape data sets showed:

- 1600 permanent tape data sets in use
- 96% of all tape data sets contained fewer than 30 megabytes of data
- 95% of data set usage (number of times used) came from data sets that constituted only 3% of the total number of tape bytes transferred
- 70% of the data sets were in generation data groups
- 86% of the data sets had standard labels
- 25% of data set usage came from data sets having less than one megabyte of data
- 27% of the data sets contained less than one megabyte of data, but constituted only 1% of the total number of tape bytes transferred

The results for DASD data sets showed:

- 1762 permanent DASD data sets in use
- 99% of all DASD data sets contained fewer than 30 megabytes of data

- 89% of data set usage came from data sets having fewer than 30 megabytes of data. These data sets constituted 90% of the total number of DASD bytes transferred. track, tracks per cylinder, or cylinders per module, the program must be changed.
- 10% of the data sets were ISAM
- 45% of the sequential data sets were in generation data groups
- 23% of the data set usage came from data sets having less than one megabyte of data
- 48% of the data sets contained less than one megabyte of data, but constituted only 2% of the total number of DASD bytes transferred

This data set information helped to decide where to place the data sets in the Mass Storage System, and was useful for data management and control in the new environment.

Data Set Selection

To be a candidate for placement in the Mass Storage System, a data set had to meet these guidelines:

- All Mass Storage System data sets must be cataloged so that data set information can be obtained without manual procedures or VTOC listings.
- Use standard OS/VS labels. This is an installation standard
- Mass storage volumes are not to be exchanged between data centers. Analyze data sets that must be exchanged, and then either leave them on tape or copy them from a mass storage volume to tape before exchanging them.
- ISAM data sets must be converted to VSAM before placing them in the Mass Storage System. They will remain on 3330s until conversion.
- No data sets larger than 100 cylinders (approximately 30 megabytes). A 30-megabyte cutoff reduces manual handling of data sets (95% for tape and 89% for DASD), and produces staging times that are acceptable. Analyze data sets larger than 30 megabytes for placement in the Mass Storage System based on their frequency of use, system resource requirements, and possible system redesign. Consider re-creating some data sets in a VSAM KSDS organization to allow "update-in-place". Data sharing reuse can make a large data set a valid candidate for the Mass Storage System.
- To avoid unnecessary staging or destaging operations, all Mass Storage System data sets must use space allocation in cylinder increments only.
- No data sets smaller than one cylinder. Smaller data sets should be transferred to 3330s. If all data sets smaller than one cylinder are placed on 3330-11s, approximately 27% of the tape data sets, and 48% of the DASD data sets, can be on 3330-11s. This requires approximately one and one-half 3330-11 disk packs to contain the 434 tape data sets, and the 845 DASD data sets.

Volume and Space Management

The data processing department chose to retain its central control over the mass storage volumes as it had previously done with the 2314 and 3330 volumes. This centralized space management department is responsible for:

- Grouping

- Volume creation
- Volume space fragmentation control
- Space reclamation

Volume and space management encompasses many areas, including organizational structure, use of generation data groups, and programming tools to control mass storage volumes.

Grouping Volumes by Space: The following groups of mass storage volumes were defined for production work on the basis of space requirements:

- Group I contains data sets that require from one to seven cylinders. Primary and secondary extents of one cylinder are specified.

  SPACE = (CYL,(1,1))

- Group II contains data sets that require from 8 to 39 cylinders. Primary and secondary extents of eight cylinders are specified.

  SPACE = (CYL,(8,8))

- Group III contains data sets that require from 40 to 100 cylinders. A primary extent of 40 cylinders, and a secondary of eight cylinders are specified.

  SPACE = (CYL,(40,8))

All volumes within each group were general-use volumes created with the default options of NOBIND, NODASDERASE, READWRITE, and SHARED. A restricted-use volume can be created in a given application area to ensure that no other application allocates data on that particular volume. The restricted-use option is a control tool used by space management to remove volumes from the MSVC's 'availability list'. This prevents selection of a given volume from a group in response to a nonspecific request. Restricted-use volumes are only available for specific requests (old or shared). Additional groups were created, which contained volumes for testing and system programming.

The initial approach to grouping was chosen for these reasons:

- MSVC does not select a currently mounted volume to satisfy an MSVGP request. It is possible to have a job in allocation with the MSVGP parameter requesting a volume from a group that currently has all its volumes mounted. This causes abnormal job termination.

  Note: Installation Example 1 was planned and implemented before the allocation to mounted volume support was available.

By creating larger groups, the number of volumes in each group always exceeds the number of online virtual UCBs. The above situation cannot cause abnormal job termination. Instead, the allocation recovery routine is called when no online UCBs are available.

For example, if groups I, II, and III above each were to contain 100 mass storage volumes, no more than 99 virtual UCBs would be online at any one time. This prevents possible job termination but might call the allocation recovery routine.

- The data processing department chose to align data sets on page boundaries. Because the mass storage volume VTOC is always on cylinder 0 (using CREATEV and assuming a VTOC extent of one cylinder), the next seven cylinders must have a "dummy" data set created to initially force staging page alignment. The above grouping strategy (except group I, for which 'page splitting' should be minimal) maintains this page alignment.

- The specific space default allocations for groups I, II, and III were made because data set sizes tend to cluster at these particular points, as observed during data set analysis.
- Space requests of similar size within a group also tend to minimize fragmentation on individual mass storage volumes within the group.

Generation Data Groups: Generation data groups are useful tools for space management, to help meet grouping and space request requirements. When a generation data group is to be converted to the 3850, the space manager builds a catalog structure for the GDG. The DELETE option is specified so that as a new generation is created, the oldest generation is uncataloged and deleted from the VTOC.

Three important points are:

- Generation data groups are normally used as Master In/Master Out. Two generations are utilized in a given job step: the input (generation 0) and the new generation (+1) being created. The allocation routine resolves requests for old data sets first. Use the data separation parameter MSVGP to achieve volume separation.
- After the new generation is created, the oldest generation is uncataloged and the data set is deleted. This feature of processing allows the system to reclaim space on mass storage volumes without using VTOC listings. A DD statement is required to mount the volume that contains the oldest generation data set to be deleted.
- Over 70% of the tape-resident data sets, and 45% of the sequential DASD data sets, are generation data groups. This provides system-managed space reclamation.

  Note: Installation Example 1 was planned and implemented before the allocation to mounted volume support was available.

Non-VSAM Data Sets: This category includes non-generation data groups sequential, partitioned, and direct data sets. The OS/VS JCL DD statements specify the appropriate MSVGP for creating these data sets the same way they did for the generation data groups. Reclaiming space for these data sets is done by normal OS/VS JCL disposition processing (OLD, DELETE, KEEP), where possible. If this is not possible, list and analyze the VTOCs, then delete data sets as required.

VSAM Data Sets: VSAM requires preallocation of data sets in a predefined VSAM data space; these data sets can be on any number of mass storage volumes. Because all VSAM data set information is contained in the VSAM catalogs, you can analyze VSAM data sets without using VTOC listings.

MSV Inventory Data Sets: The MSV Inventory data set helps control mass storage volumes. A space threshold of 70% generates a message when that amount of the total space on general-use volumes within the group is depleted. At this point, a LISTMSVI Space Usage Report is run for this group to determine the largest extent, the total number of extents, and the amount of fragmentation associated with each volume. The VTOC is listed to display the data set information, and the appropriate actions are taken to free space. These actions might be a simple audit followed by deletion of unneeded data sets, or a compression of the volume. If the volume has no space, other than fragmented space, audit the VTOCs of the volumes within the group. Space can be made available by deleting data sets, or by adding volumes to the group.

Standards and Procedures

A review showed that most existing standards were valid but that they required better enforcement. Some new standards were introduced. The standards related to the Mass Storage System are: cataloging and OS/VS JCL parameters. Catalogs reflect the VTOCs of all mass storage volumes. VTOCs and catalogs are scheduled for periodic comparisons. Any data sets not cataloged are deleted; any data sets cataloged but nonexistent are uncataloged.

The job control language parameter standards are:

- All unnecessary parameters must be removed from all OS/VS JCL statements. DCB information, for example, is allowed only on those DD statements that create a data set.
- To take advantage of 3330 capabilities, data sets must be reblocked to approximately 4096 bytes, and region sizes must be increased where necessary.
- Unit parameters are required for:
  1. UNIT=2400-3 for all tape data sets
  2. UNIT=3330 for all real DASD permanent data sets
  3. UNIT=SYSDA for all real DASD temporary data sets
  4. UNIT=3330V for all 3850 data sets
- Disposition parameters are required for:
  1. DISP=(NEW,CATLG,DELETE) for new data sets
  2. DISP=(OLD,KEEP) for retained old data sets
  3. DISP=(OLD,DELETE,KEEP) for old data sets not to be retained

Because 3330V data sets have the same characteristics as real DASD, it is now possible to share previous tape data sets with several jobs. Consequently after specific analysis is done of data set usage in relation to job scheduling, it might be advantageous to specify DISP=SHR for some data sets.

- Group defaults must be used. SPACE is coded in the OS/VS JCL in only exception cases.
- The three groups I, II, III of mass storage volumes must be coded MSVGP=groupname.
- Retention periods must be assigned and maintained.
- The conversion to the Mass Storage System is an opportunity to redefine data set naming conventions. Previously, the high-level DSN qualifier was the same for all data sets. This left little flexibility for user catalogs. The new data set naming convention is xxx.yyyyyyy, where xxx is the application code uniquely identifying an application area and yyyyyyy is the unique identification within an application. This method allows flexible catalog structures, and the high-level qualifier readily associates a data set with an application area for cross-reference and space auditing procedures.

- All OS/VS JCL must reside in a single procedure library.
- The scheduling system run books must be updated to reflect run procedures. In particular, the rerun and restart procedures for all jobs must be clearly defined. Procedures no longer can rely on operator visual checks to determine whether a certain tape is mounted. Differences between tape and DASD restart procedures, such as using the third parameter on DISP, must be clearly defined.

Conversion Techniques

The conversion areas addressed were:

- Sharing data and catalogs between the OS/MVT and OS/VS2 MVS systems
- OS/VS JCL modifications
- Data movement to the Mass Storage System

Catalog and Data Sharing: The data set analyses showed where the data sets were used, which is the basis for managing catalog usage and data migration. Data sets were divided into two categories: one for migration to the Mass Storage System, and one for those that could not then migrate because of interdependencies with other applications running under OS/MVT. As data sets were transferred to the Mass Storage System, they were assigned new names and cataloged in a VSAM user catalog under OS/VS2 MVS. Those not transferred, retained their old names and positions in the OS CVOL catalog. Figure 3 shows this catalog structure.

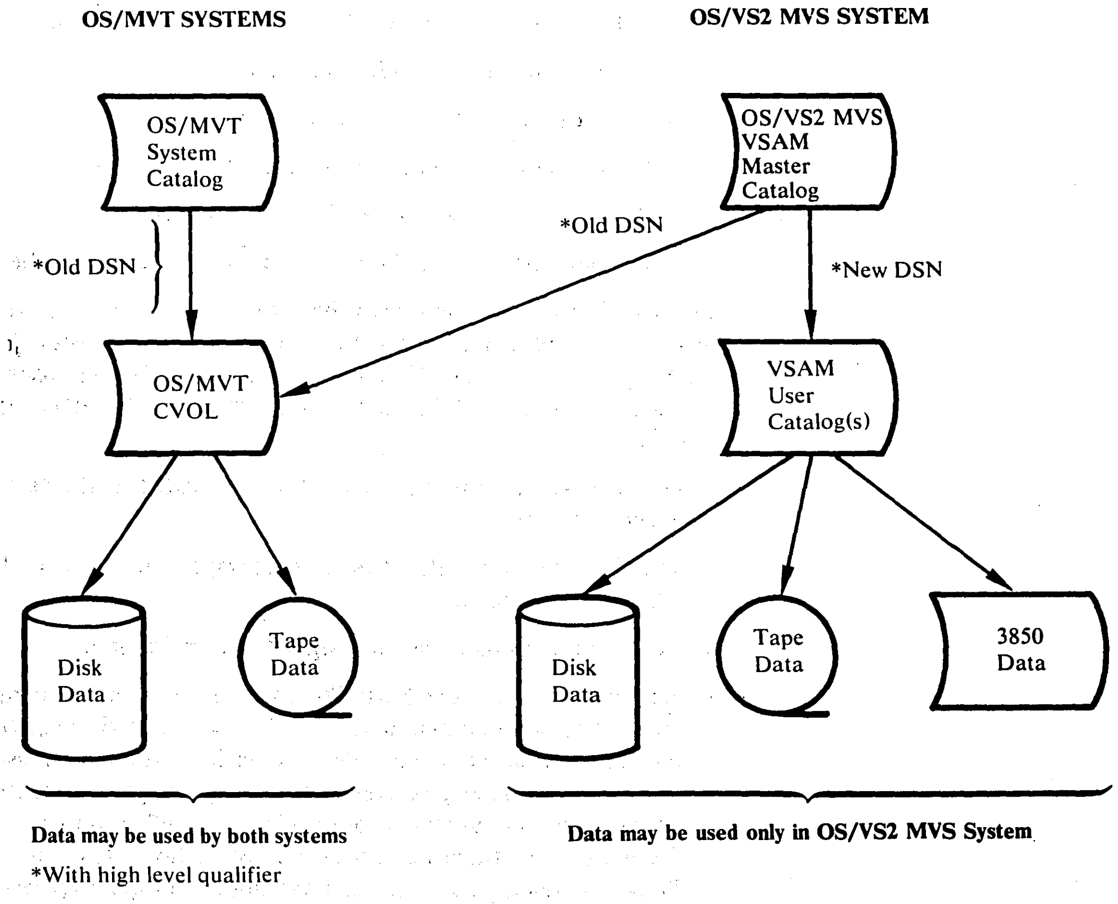

Figure 3. Catalog Structures

OS/VS JCL Modifications The procedures used for modifying OS/VS JCL were:

- Copy the ones affected in the OS/MVT PROCLIB to a new PROCLIB under the OS/VS2 system. Modifications were made by using TSO and a batch text-editing Installed User Program (IUP) (OS/VS JCL Editor Program 5796-PDC). All changes met the established standards.
- Make the changes in two separate passes. Pass 1 changed the DD statements to reflect the new data set names and the new unit (3330V); pass 2 changed the INPUT DD statements to reflect the new input data set names.

Data Migration: There are three ways to move data to the Mass Storage System:

- Master In/Master Out
- Copy (IEBGENER)
- CONVERTV

Because all staging drives were 3330-11s, the CONVERTV command could not be used. Therefore, data was transferred to the Mass Storage System by using Master In/Master Out.

During the first pass with OS/VS JCL changes, the jobs were run using the current data set locations from the OS CVOL catalog; this is Master In. The output, or new version of the data set, is transferred to the Mass Storage System and recorded in the OS/VS2 MVS VSAM user catalog under the new data set name; this is Master Out. Pass 2's subsequent OS/VS JCL change to the INPUT DD statement sent the allocation routine to the OS/VS2 MVS VSAM user catalog to get the data set location. At this point, the VSAM catalog directs both input and output to the Mass Storage System. Some data sets were copied from the OS/MVT system to the Mass Storage System using IEBGENER; these were few compared to the number of data sets converted using Master In/Master Out.

Example 2

Example 2 is for an installation that uses OS/VS1. The installation also took place before the final system and program testing for customer shipments were completed. As in the first example, a portion of the production workload was identified for testing during the first phase of a two-phase plan. The following pages describe the experience gained.

Initial Conditions

The data processing department's mission in this example is to collect and process special hand-written paper documents. An average daily input of 65,000 documents is scanned by an IBM 1287 Optical Character Reader (OCR). The applications that use this input are predominantly batch. The queued sequential access method (QSAM) is used extensively to process master files, transaction files, and exception files, which resided primarily on tape reels before the Mass Storage System was installed. This is the reason for the Master In/Master Out system design of the primary batch applications. Although the primary system workload is batch, there is some document correction being done online using IBM 2260 display terminals and DASD-resident data sets.

This medium-sized OS/VS1 operation runs three shifts per day, five days per week. The existing tape library contained approximately 4,800 reels and

4,000 active data sets. Between 200 and 250 jobs are run each day. When the Mass Storage System installation planning began, the production system was using OS/MVT on a System/370 Model 155 (one megabyte). A Model 158 replaced the Model 155 for two reasons: to obtain the Dynamic Address Translation Facility, and to support a planned application extension to one of the primary applications. The Model 158 was used to convert from OS/MVT to OS/VS1 Release 3 before installing the Mass Storage System.

The production system, running under OS/VS1 Release 3, had this equipment:

- One System/370 Model 158 (two megabytes)
- Fourteen 3420-5,8 magnetic tape drives
- Sixteen 2314 disk drives
- Four 3330-11 disk drives
- Miscellaneous unit record devices (1287, 1419, 2260, 3211, 2540)

A System/370 Model 145 (512KB) was installed adjacent to the Model 158 production system to test the prerelease of OS/VS1 Release 4 SCP and the 3850. The following equipment was attached to the Model 145:

- One 3851-B1 Mass Storage Facility
- Two 3830-3 Storage Controls
- One 3830-2 Storage Control
- Fourteen 3330-1 disk drives
- Two 3420-5 magnetic tape drives
- One 1403 printer
- One 2540 card read/punch

The installation activities are described relative to the following requirements:

- Data set analysis to determine which data sets should migrate to the Mass Storage System
- Establishment of new standards and procedures to permit an orderly migration to a disciplined data environment managed by the Mass Storage System
- Application conversion for a smooth transfer of work to the Mass Storage System

Data Set Analysis

A detailed analysis of tape and DASD data sets was completed early; it required examining the block size, data organization, data set size, and frequency of use for each data set. It also examined the relationship of each data set to each job and job step in the application by using data set "where used" information. Sufficient data set information was obtained from two sources: the user-developed Job Information System, and SMF records that had been collected over a four-week period. The tape data set analysis showed:

- 2100 permanent data sets were active during the four-week period.
- 98% of all permanent tape data sets were smaller than 30 megabytes.

- 97.5% of the permanent data set usage (number of times used) came from data sets constituting only 82% of the total number of tape bytes transferred.
- 3.5% of the data sets were in generation data groups.
- 2% of the data sets were multiple tape volumes (reels).
- 24% of the data set usage came from data sets smaller than one megabyte.
- 35% of the data sets were smaller than one megabyte in size but constituted only 1% of the total number of tapes bytes transferred.
- 39% of all tape data transferred was to, or from, tape SORTWORK files.

The DASD data set analysis showed:

- 1550 permanent DASD data sets were active during the four-week period.
- 100% of the DASD data sets were smaller than 30 megabytes.
- 98% of the data set usage came from data sets smaller than 20 megabytes. These data sets constituted 81% of the total number of DASD bytes transferred.
- 2.3% of the data sets were ISAM.
- 0.5% of the sequential data sets were in generation data groups.
- 80% of the data set usage came from data sets smaller than one megabyte.
- 85% of the data sets were smaller than one megabyte but constituted only 6% of the total number of DASD bytes transferred.
- 32% of all DASD data transferred was to, or from, DASD SORTWORK files.

Combining the tape and DASD data set analyses showed:

- 1400 permanent data sets (1.7 billion bytes) were created, which had a retention period of zero (0) days.
- 130 permanent data sets (240 megabytes) were created, which had a retention period between one and two days.
- 250 permanent data sets (736 megabytes) were created, which had a retention period between three and ten days.
- 330 permanent data sets (2.3 billion bytes) were created, which had a retention period between 15 and 30 days.
- 960 permanent data sets (9.9 billion bytes) were created, which had a retention period between 45 and 92 days.
- 23 data sets (1.3 billion bytes) were created, which had a retention period between 180 and 999 days.

Data Set Selection

The information collected from the data set analysis was used to establish the following criteria for placing application data sets in the Mass Storage System:

- Tape data sets smaller than 100 megabytes were planned for transfer to the Mass Storage System.
- DASD data sets larger than what one 3330-1 cylinder can hold were planned for transfer to the Mass Storage System.
- SORTWORK data sets (although they were temporary data sets) were allocated to 3330V devices.
- ISAM data sets were to remain on nonstaging DASD devices until they were converted to VSAM.
- Special data sets called "control files" (tables, certain master and transaction files, and check point areas), which are small and frequently accessed, were to remain on tape or DASD.
- Data sets that had to be exchanged were to remain on tape or DASD.

Standards and Procedures

This data processing operation is well-organized and managed by a centralized group of scheduling and control people. A user-developed Job Information System documents all job data set information. Some of the standards include using:

- PROCLIB for all production OS/VS JCL
- Symbolic OS/VS JCL parameters in most procedures
- Stringent naming conventions for all procedure names, job names, step names, dd names, data set names, and volume serial numbers
- Rigidly observed expiration dates on all data sets (which permits the timely and systematic reclamation of space)
- Password protection for most master files
- Protection of certain Access Method Services commands in an APF-protected library
- The OS/VS1 system catalog for access to all production data sets that are not temporary
- Device-independent programming techniques
- Centralized creation and maintenance of all OS/VS JCL in the PROCLIB (includes enforcing OS/VS JCL coding standards and conventions)

The OS/VS JCL Editor Program IUP (5796-PDC) was used to help maintain consistent coding of the OS/VS JCL in the PROCLIB, and to help modify the DD statement parameters that define data sets for the 3850. The keyword parameters that specify UNIT=value and MSVGP=volume are the most widely modified DD statement parameters.

A set of DD statements such as:

```
//NEWMAST   DD   DSN=MASTER.&D1,DISP=(,CATLG,CATLG),
//               UNIT=TAPE9,...,
//               LABEL=RETPD=10
//OLDMAST   DD   DSN=MASTER.&D2,DISP=OLD
```

might be changed to:

```
//NEWMAST   DD   DSN=MASTER.&D1,DISP=(,CATLG),
//               UNIT=&DEVTYP,...,
//               LABEL=RETPD=10,MSVGP=&APPLGP10
//OLDMAST   DD   DSN=MASTER.&D2,DISP=OLD
```

After changing TAPE9 to &DEVTYP, another symbolic parameter in the OS/VS JCL was used for the job. By providing the symbolic parameter card:

```
//          DD    DEVTYP=TAPE9,
```

the job ran as before. In this case, the MSVGP= parameter is *not* used during allocation. The job allocates the new master data set to mass storage volumes selected by Mass Storage Control by coding:

```
//          DD    DEVTYP=3330V
```

The MSVGP parameter is inserted in the DD statement on the basis of the RETPD parameter value. Mass storage volume grouping is described later in this example. The DD statement for the old Master data set remains unmodified. The system catalog is used to determine volume serial and unit information for the old master.

The previous DD statements illustrate the data set naming convention used. The &D1 and &D2 notations are symbolic Julian calendar dates. They represent the creation dates of the new and old master data sets, respectively. The centralized control people prepare the symbolic OS/VS JCL parameter cards and schedule all production jobs in the system. This data set naming convention allows you to establish a GDG-like environment for all production data sets. That is, all versions of a data set are differentiated by the date qualifier in the data set name. The procedure described in *OS/VS Mass Storage System (MSS) Services: General Information* was used to protect the critical Access Method Services commands from unauthorized users.

Mass Storage Volume Management

New procedures were developed for mass storage volumes, which apply to volume numbering, mass storage volume groups, nongrouped mass storage volumes, and space reclamation.

All mass storage volumes are assigned serial numbers by the Access Method Services CREATEV command. The serial number format is MSnnnn, where nnnn is a decimal number from 0001 to 9999. Mass storage volume groups are defined primarily on the basis of their application, and secondarily on their data set retention periods. Hence, symbolic mass storage volume (MSV) group names are defined as follows:

- The group name SYSGROUP contains 10 mass storage volumes. Data sets with retention periods from zero to two days are allocated to this group, and belong to all applications. Mass storage volumes in this group are assigned a volume retention period of three days. Space occupied by expired data sets is reclaimed daily by deleting the entire volume on which the data set resides.
- The group name &APPL.GP10 contains approximately 15 mass storage volumes. Data sets with retention periods from 3 to 10 days are allocated to this group. The application qualifier &APPL (its use is optional) permits separating data sets into groups based on application. The volume retention period is 15 days. Space is reclaimed daily by deleting expired volumes.
- The group name &APPL.GP30 contains 38 mass storage volumes. Data sets with retention periods from 11 to 30 days are allocated to this group. The volume retention period is 35 days. Space is reclaimed twice a week by deleting expired volumes.

- The group name &APPL.GP90 contains 140 mass storage volumes. Data sets with retention periods from 31 to 90 days are allocated to this group. The volume retention period is 95 days. Space is reclaimed twice a week by deleting expired volumes.
- The group name &APPL.GP99 contains 24 mass storage volumes. Data sets with retention periods from 91 to 999 days are allocated to this group. The volume retention period is 999 days. Space is reclaimed weekly by deleting expired data sets.

By supplying a symbolic application parameter (for example, &APPL=PROL, in the payroll application) in the EXEC JCL statement, the output data sets created by the job can be allocated to the groups that are created especially for that application; that is, PROLGP10, PROLGP30, PROLGP90, and PROLGP99. The grouping technique emphasizes two major space management objectives:

- To eliminate space fragmentation using the release parameter
- To reclaim space by volume instead of by data set

The Access Method Services LISTMSVI command is used to display expired mass storage volumes. A utility that lists the VTOC of each expired volume is executed to determine whether free space still exists on the volume. If free space exists, the MODIFYV command is used to change the volume expiration date; if not, the entire volume is deleted and MODIFYV is used to nullify the volume retention period. The mass storage volume is then "empty", and is a candidate for volume selection in its assigned group.

For each of the above group names, these attributes and characteristics are defined:

- NOBIND
- NODASDERASE
- SHARED
- READWRITE
- Release unused space
- Space threshold of 25%
- Default space parameters (CYL,(8,8))

Sometimes, for data sets containing 30 to 100 megabytes, the group's default space parameters are overridden by the DD statement SPACE parameter. In this case, all space requests are coded by using the CYL subparameter. Track or block allocations are not allowed.

These *nongrouped* mass storage volumes are used, thusly:

- Ten mass storage volumes are defined in PRESRES as permanently mounted storage volumes, and are used for temporary data sets (for example, SORTWK).
- Twelve mass storage volumes (one for each application programmer) are defined as programmer test volumes. The owner of the mass storage volume is responsible for volume management.
- Fifteen mass storage volumes are defined as system support volumes. They are used by the staff system programmers.

Application Conversion

The primary conversion activities included:

- Analyzing, changing, and testing OS/VS JCL procedures
- Analyzing, changing, and testing application programs
- Transferring data sets from tape and DASD to mass storage volumes
- Changing and documenting operational procedures

The plan was to do the testing on the Model 145 and the production on the Model 158. The plan had five phases: preunit test (Model 145), unit test (Model 145), parallel test (Model 145), environmental test (Model 158), and production (Model 158).

Preunit Test: The preunit test was further divided into four areas: training the installation team, system support preparation, equipment preparation, and documentation and test procedures.

The system programmers, I/O control clerk, and other members of the installation team attended the Mass Storage System Installation Planning Workshop. Operators are given on-site training and use the *Operators Library: IBM 3850 Mass Storage System Under OS/VS* as the basic instruction material.

System support preparation is the job of the systems programmer and the I/O control people. Their primary tasks are:

- Creating the Mass Storage Control tables using the Mass Storage Control Table Create program
- Completing the OS/VS1 system generation
- Coding and testing the OS/VS JCL, the control statements for the Access Method Services, the OS/VS utilities, the MSVC Journal and MSV Inventory data sets, and the Mass Storage System Trace function
- Modifying and testing user-written OS/VS1 exit routines and utilities for the Mass Storage System
- Modifying the OS/VS JCL for the jobs to be tested on the unit

* Setting up the APF-protected libraries

Equipment preparation requires:

- Loading the data cartridges needed for unit testing
- Creating mass storage volumes and groups using the CREATEV and CREATEG commands
- Initializing the staging disk packs using IEHDASDR

Documentation and test procedures include obtaining the required OS/VS1 Release 4 and Mass Storage System publications for defining, using, and maintaining the system.

Unit Test: The unit test ensures that all batch jobs run in a single-initiator mode, and tests the OS/VS JCL, programs, and data results, using the Mass Storage System. A unit test log, maintained by the system operator, is used to show the results of each job. This log helps track OS/VS JCL errors and program problems. The tape and DASD data sets used in the unit test were transferred from the Model 158 to mass storage volumes, and cataloged in the OS/VS1 system catalog. Existing documentation, including scheduling, operation procedures, and run books, was updated as each application was tested.

After it is tested, an application is ready for the parallel test environment.

Parallel Test: The parallel test ensures that:

- Applications achieve the same results with the Mass Storage System as they did with tape and DASD.
- All system operators receive the maximum amount of hands-on training with each application.
- Application documentation correctly reflects the Mass Storage System environment.
- The mass storage grouping strategy and space reclamation techniques are usable.

An application begins its parallel test at the beginning of the monthly application cycle. The current OS/VS JCL, programs, and "control data sets" were copied to the Model 145, and installed for use during the parallel test. The input data for the application was used first in the production system (Model 158), and then in the test system (Model 145).

Environmental Test: The environmental test ensures that the applications tested during the parallel test on the Model 145 can run on the Model 158 concurrently with the jobs that use the 1419 and 1287 magnetic and optical scanners. These tests were run during weekends so that they would not interfere with the regular production.

Preparing for the environmental test included:

- Doing a Mass Storage Control Table Create program update to substitute the Model 158 identification
- Transferring the Mass Storage System from the Model 145 to the Model 158
- Doing an IPL of the Model 158 from the OS/VS1 system-residence volume used during the parallel test
- Initiating a job mixture similar to the Model 158 production

Production: When all environmental testing was completed, the Mass Storage System remained connected to the Model 158 for normal use. Testing continues for optimizing the Mass Storage System operation. This requires the system programmer to collect the Trace information, interpret the results, and change the Mass Storage System Table Create program parameters or modify the system configuration.

Access Method Services commands are used to:

- Analyze the Mass Storage System staging and destaging operations
- Analyze the Mass Storage System component utilization, including the DRDs, DRCs, Staging Adapters, and staging drives
- Analyze mass storage volume use
- Modify the Mass Storage System Table Create program parameters if necessary
- Modify the mass storage volume groups if necessary

Trace and Tune are normal operation activities that are performed on a regular schedule or whenever there are significant system or workload changes.

Logical Concepts, Addresses, and Notations

Configuring a Mass Storage System means specifying a physical system, a logical system, and the interconnections between the Mass Storage Facility, 3330, and CPUs. This guide assumes that the physical configuration was done using the available programming tools or by a manual process.

Logical configuration means planning 3330 staging drive usage as a collection of physical and logical entities, called staging drive groups, and assigning a virtual unit address range to each group. Also, the Mass Storage System physical component interconnections must be planned so that the physical and logical systems function as specified.

Staging Drive Groups

A CPU can access data on a mass storage volume only after the data has been staged to a virtual drive and becomes a virtual volume in a staging drive group. Figure 4 shows the staging drives represented as a pair of staging drive groups. The data path from the Mass Storage Facility to the staging drive group pair is through the Staging Adapter buffer (path B-D), and the data path from the staging drive group pair to the CPU is a direct path through the Staging Adapter (path C-D).

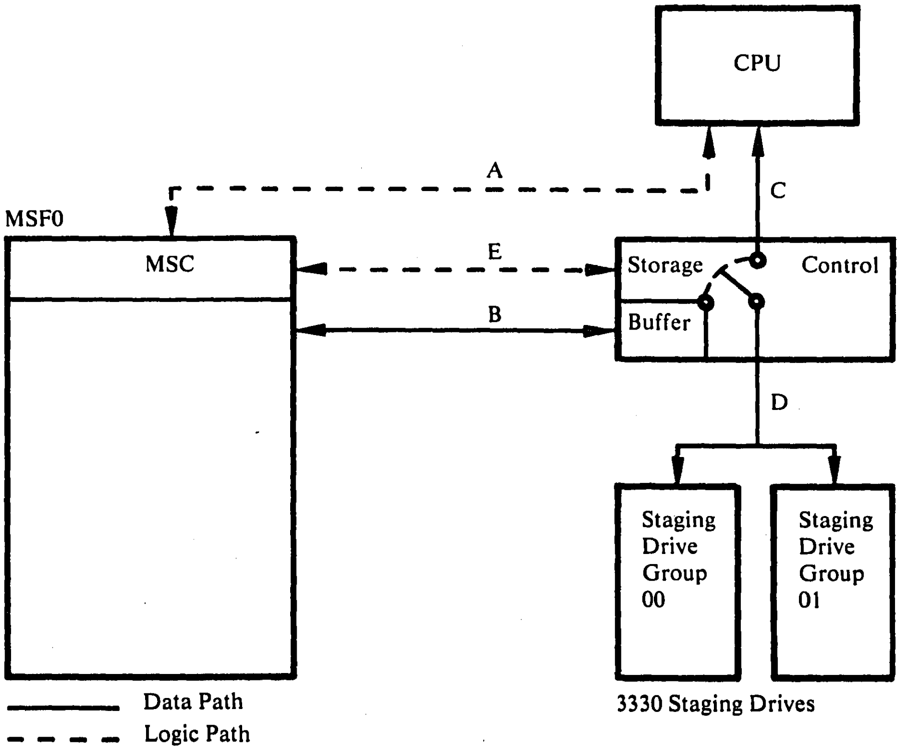

Figure 4. Staging Drive Groups

Mass Storage Control management of the 3330 staging drives permits the "creation" of a virtual disk storage subsystem that has a larger capacity than the physical disk storage subsystem. The internal and external management of the 3330 staging drives is based on staging drive groups.

A staging drive group consists of up to 800 megabytes of staging drive space that is divided into eight-cylinder space increments called "pages". The size of a staging drive group can range from a minimum of 51 pages (one 3330 Model 1 drive), or 101 pages (one 3330 Model 11 drive), to a maximum of 408 pages (eight 3330 Model 1 drives) or 404 pages (four 3330 Model 11 drives).

The Mass Storage System uses:

- Real, convertible-to-real (CONVREAL), and staging drives
- Real, virtual, and virtual shared unit addresses
- Virtual devices
- Real, virtual, and mass storage volumes
- Staging and nonstaging disk packs

IBM 3330 Models 1 or 11 drives can be used as real, convertible-to-real, and staging drives; IBM 3350 Model 2 disk drives can be used only as real (nonstaging) devices. Real unit addresses always map to the same real drives, or to the same convertible-to-real drives in a real mode (varied online as real). In real mode, the unit control block and unit address format are identical to a real drive. The function of the drive (not the address) is convertible-to-real.

Virtual unit addresses map to a staging drive group and create the virtual devices used by the CPU. Pages that are Mass Storage Control-allocated for the mass storage volume mounted on the virtual drive, might be on one or more staging drives within the staging drive group. This includes pages on convertible-to-real drives when they are in staging mode (varied online as staging).

Real volumes are volumes that are either operator-mounted on 3330-1 or 3330-11 real drives or convertible-to-real drives that are in real mode, or they are volumes that are permanently mounted on 3350 Model 2 disk drives.

A virtual volume is a logical 3330-1 volume "created" when the data (or a portion thereof) from a mass storage volume is staged to pages on one or more staging drives in a staging drive group.

For staging use, the operator must mount disk packs, which are formatted for staging, on staging drives or convertible-to-real drives that are varied online as staging. For nonstaging use, the operator must mount disk packs, which are formatted for real use, on real drives or convertible-to-real drives that are varied online as real.

The illusion of a virtual drive is created when the System Control Program issues an order to mount a mass storage volume on a virtual unit address. In reality, the virtual drive exists logically on one or more staging drives that contain the virtual volume. The System Control Program sends a Mount order to the Mass Storage Control specifying a mass storage volume serial number and a virtual unit address. The Mass Storage Control then stages cylinder 0 (volume label) of the mass storage volume to a page in the staging drive group associated with the selected virtual unit address. At

OPEN time, the Mass Storage Control stages the requested data set extents to pages in the same staging drive group as it did the volume label.

The CPU cannot distinguish a virtual from a real volume. Each virtual drive has a unit control block (UCB) that has one bit to indicate that the UCB is virtual. The System Control Program "sees" the Mass Storage System as a large number of virtual 3330-1 volumes mounted on virtual drives.

The number of virtual drives in a Mass Storage System is determined by the number of virtual unit addresses (or virtual UCBs) associated with all staging drive groups. The number of staging drive groups is determined by the number of Staging Adapters and your input to the Mass Storage Control Table Create program. Up to two staging drive groups can be specified for each Staging Adapter or Staging Adapter pair.

The total number of available OS/VS addresses or UCBs, and the corresponding number of nonshared subchannels or unit control words (UCW), determine the total number of virtual drives that can be "created" for each CPU. Your input to the Mass Storage Control Table Create program specifies the distribution of the virtual unit addresses to each CPU for access to each staging drive group.

As Figure 4 shows, the CPU issues a mass storage volume Mount order to the Mass Storage Control (in MSF0) by using path A. The requested data or VTOC is staged to one or more pages in Mass Storage Control-selected staging drive group using path B-D. The CPU accesses the data using path C-D. The Mass Storage Control communicates with the Staging Adapters and the Mass Storage Control tables using path E and E-D respectively.

Be sure you understand how staging drive groups, pages, and virtual drives interact, because they are the elements used internally and externally to manage, utilize, and adjust the staging drive space accordingly. Internally, the Mass Storage Control uses path A to pass "activity" and "reuse" information to the System Control Program so that it can distribute the workload among all staging drive groups. The System Control Program also uses the "activity" information to "sort" the staging drive group names (SDG00, SDG01, SDG02, and so forth) into a "least mounts/most available space" sequence list based on least recently used (LRU) activity. When the System Control Program needs a mass storage volume mounted, it selects a virtual unit address from those available in the first staging drive group in the list. If all virtual drives in the first staging group are busy, the System Control Program selects a virtual unit address from the second staging drive group in the list, and so forth. A mass storage volume is considered to be mounted when the volume label (and usually the VTOC) in cylinder 0 has been staged to a page in its staging drive group.

If the CPU-requested data was previously mounted and demounted but is still available in a staging drive group that is known to the Mass Storage Control, the System Control Program moves the staging drive group name that contains the data to the top of the staging drive group name list. A virtual unit address is selected from that staging drive group and the requested data (called data reuse) is immediately available to the CPU.

The number of active pages within a staging drive group is controlled by defining lower and upper page deltas, which the Mass Storage Control uses to calculate upper and lower space thresholds. When the upper space threshold is exceeded, the amount of free and inactive pages remaining in a staging drive group is minimal. If the upper space threshold is exceeded during allocation, the LRU algorithm is called; that is, the virtual volume with the largest number of oldest pages within each successive LRUGROUP (pages grouped by age) is marked for deallocation (changed cylinders are destaged before becoming inactive), under control of the LRU algorithm, until the new space request is satisfied and/or the lower threshold is reached. The lower threshold sets a minimum reserve of active pages.

The addition or deletion of pages in a staging drive group can be controlled by varying staging drives online or offline. The amount of staging and destaging operations and the optimizing of the Mass Storage System are directly affected by the rate at which pages are aged (using LRUCLOCK), the page grouping for LRU (LRUGROUP), and the upper and lower space thresholds.

Staging Adapters

The 3830-2 Storage Control or the Integrated Storage Control, in the Model 158 or 168 is modified for use with the Mass Storage System by installing the Staging Adapter feature. Figure 5 shows the prerequisite and optional special features for an ISC and 3830 before and after the Staging Adapter installation.

Hereafter, reference to a Staging Adapter indicates one 3830 Model 3 Storage Control or a logical half of one ISC with Staging Adapter.

Regardless of whether a Staging Adapter is operated in a staging or nonstaging mode after the Staging Adapter installation, the system considerations are:

- Fewer control units can be connected to a channel because each Staging Adapter interface responds to 64 of the 256 unit addresses available on a block multiplexer channel.
- A Staging Adapter requires a 3830-3 Type 23 Flexible Disk, and uses the 3830-3 unit address format for real and virtual devices. An IOGEN might be required because a Staging Adapter unit address must be plugged as hex 0, 4, 8, or C, as described in the following section.
- Staging Adapter interface A is reserved for the connection to a Mass Storage Control port. This reduces the maximum number of interfaces available on a Staging Adapter (3830 Model 3 only) for connection to a block multiplexer channel (See Figure 5).
- The 3340 Direct Access Storage Facility cannot be connected to a Staging Adapter.

Each Staging Adapter manages the data paths for up to two staging drive groups, as illustrated in Figure 4. This provides a staging drive group pair with one stage and/or destage path and up to three interfaces to a CPU channel. Two Staging Adapters can be paired to manage the data paths for a staging drive group pair (3333s string-switched to two Staging Adapters). This connection provides a staging drive group pair with two stage and/or destage paths and up to six interfaces to CPU channels. Each channel interface can respond to a maximum of 64 device addresses.

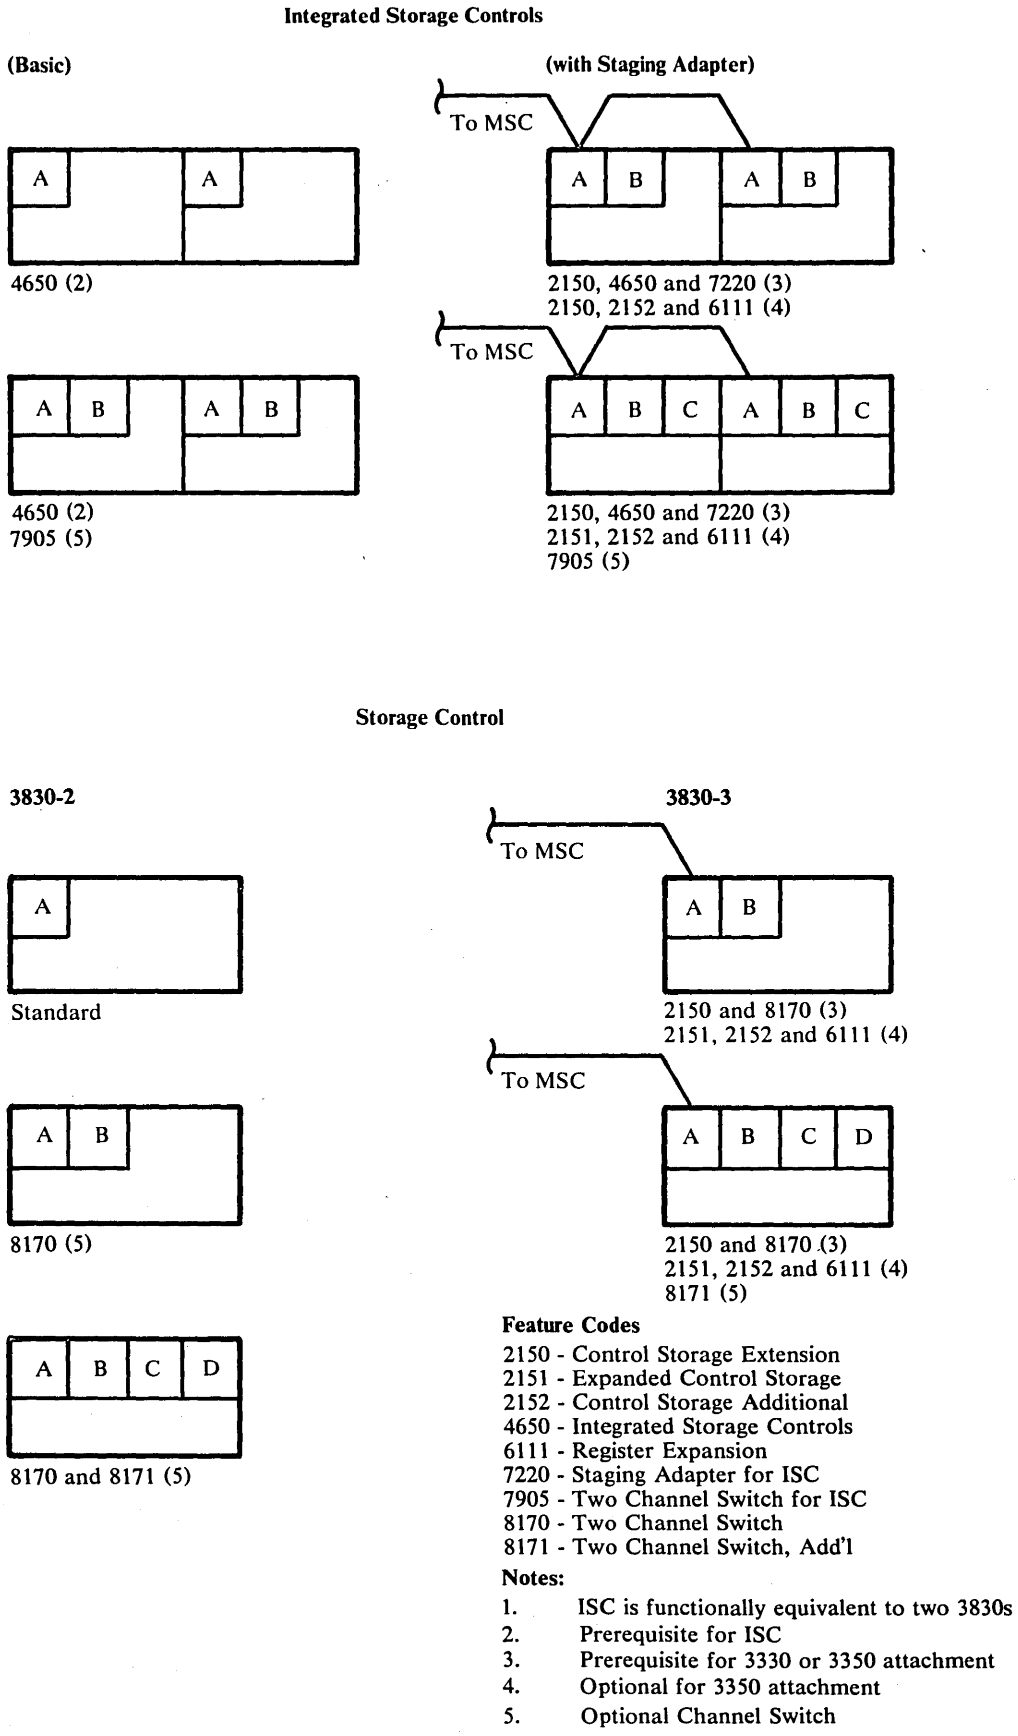

Feature Codes

2150 - Control Storage Extension
2151 - Expanded Control Storage
2152 - Control Storage Additional
4650 - Integrated Storage Controls
6111 - Register Expansion
7220 - Staging Adapter for ISC
7905 - Two Channel Switch for ISC
8170 - Two Channel Switch
8171 - Two Channel Switch, Add'l

Notes:

1. ISC is functionally equivalent to two 3830s
2. Prerequisite for ISC
3. Prerequisite for 3330 or 3350 attachment
4. Optional for 3350 attachment
5. Optional Channel Switch

| **Figure 5. ISC and 3830 Interfaces**

Each Staging Adapter or Staging Adapter pair can connect up to four strings of eight drives for a maximum of 32 drives. Each string must be controlled by a 3333 Model 1 or 11 or by a 3350 Direct Access Storage A2 or A2F and Control and their corresponding disk storage drives, 3330 Models 1, 2, or 11 or 3350 Models B2 or B2F and C2 or C2F. The 3333/3330 Models 1, 2, or 11 can be used as real, convertible-to-real, or staging drives. The 3350 Model 2 drives can only be used as real drives in native (nonstaging) mode. In addition, 3330 and 3350 series drives cannot be intermixed on the same string.

Unit Address Format

Each host connected to the 3850 uses real and virtual unit addresses to access real (3330) or virtual (3330V) drives using primary or alternate channel paths. The Mass Storage Control Table Create program specifies the unit addresses known to the Mass Storage Control, and creates the IODEVICE cards for input to Stage 1 of the SYSGEN.

As shown in Figure 6, bits 16-31 of an I/O instruction designate the channel, control unit, and device to be used in an input/output operation. In the 3830-3 unit address format, bits 16-23 designate the channel; bits 24 and 25 designate the Staging Adapter.

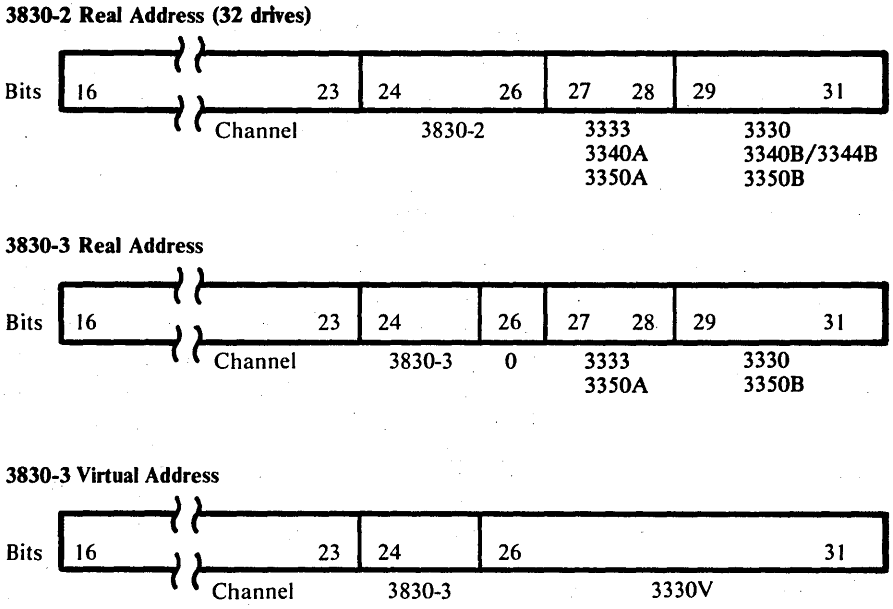

Figure 6. Unit Address Formats

If a device is specified as virtual (or virtual shared), bits 26-31 designate the virtual device. This provides a maximum of 64 virtual devices (00-3F), on each Staging Adapter's B, C, and D interfaces.

If a device is specified as real (or convertible-to-real), bit 26 must be zero, bits 27-28 designate the 3333, and bits 29-31 designate the real device. A Staging Adapter can address a maximum of 32 real devices.

If a Staging Adapter is operated in 3830-2 mode (nonstaging), the control unit address (bits 24-25) must be binary 00, 01, 10, or 11, to prevent an Interface Control Check. The Customer Engineer refers to these four address bit combinations as hex 0, 4, 8, and C, respectively, as if bits 26 and 27 were zero.

Because channel bits 16-19 are zero, a 12-bit unit address format is commonly used for virtual and real devices, that is:

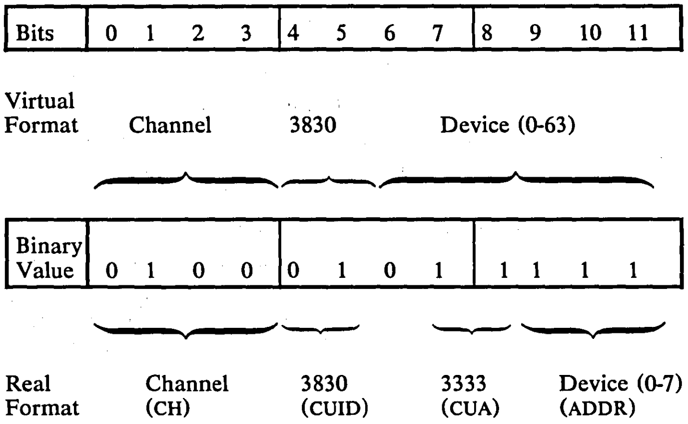

Unit address 45F is shown above in binary notation. The input to the Mass Storage Control Table Create program is device address 1F on channel 4. If the device address is specified as virtual or virtual shared, the virtual bit is set in the unit control block. If the device address is specified as real or convertible-to-real (CONVREAL), the virtual bit is not set and bit 6 above (or bit 26 of the I/O instruction) must be zero.

The Mass Storage Control Table Create program notation for the unit address is CH (for channel), CUID (control unit identification for the Staging Adapter channel interface), CUA (control unit address for the 3333), and ADDR (address for the real or virtual device).

In this publication, reference to a unit address includes bits 20-31 (12-bit format); reference to a device address (input to the Mass Storage Control Table Create program) includes bits 26-31 (6-bit format).

A block multiplexer channel can connect to a maximum of four Staging Adapter interfaces. The 256 unit addresses available on a channel are distributed to the four Staging Adapter interfaces as follows:

Staging Adapter Interface Address	00	01	10	11
Virtual or Real Unit Address	00-1F	40-5F	80-9F	C-DF
Virtual Unit Address	20-3F	60-7F	A-BF	E-FF

Unit Address Examples

Figure 7 is used in the following sections to describe the mapping of unit addresses for real, convertible-to-real, virtual devices to real, and staging drives on primary or alternate channel paths. Two CPUs connected to three channel interfaces on one Staging Adapter are used to map 30 of the available 192 (3X64) unit addresses.

Real unit addresses consistently map to the same real or convertible-to-real drives as a function of electrical paths; virtual unit addresses map to pages available in a particular staging drive group as a function of table translations. Each unit address has a single use that determines one channel path, either primary or alternate, to a real, convertible-to-real, virtual, or shared drive.

Real Unit Addresses

Real drive A and convertible-to-real drive C are addressed only by CPU1 using real unit addresses 100 and 101 on primary channel 1, or 200 and 201 on alternate channel 2.

Real drive D and convertible-to-real drive F are addressed only by CPU2 using real unit addresses 244 and 245 on primary channel 2. No alternate channel path is available.

Real drive B and convertible-to-real drive E are shared (FEATURE=SHARED on IODEVICE card) by both CPUs. The two drives are addressed by CPU1 using real unit addresses 102, and 103 on primary channel 1, or 202 and 203 on alternate channel 2 and by CPU2 using real unit addresses 242 and 243 on primary channel 2. Because FEATURE=(SHARED) is specified, Reserve and Release commands are issued by each System Control Program when volumes mounted on these drives are accessed for exclusive use.

Virtual Unit Addresses

Staging drives G and H show 12 virtual volumes mounted on 12 virtual drives. Because staging drives G and H are both specified in the staging drive group 07, any of the virtual or virtual shared addresses shown can map to either staging drive. The Mass Storage Control dynamically allocates pages on any staging drive in staging drive group 07 when it mounts a mass storage volume.

Virtual drives *g, h, i,* and *j* on staging drive H are addressed only by CPU1 using virtual unit addresses 118 and 119 on primary channel 1, and virtual unit addresses 218 and 219 on primary channel 2. Note that channel 2 is not an alternate channel. Virtual unit addresses 218 and 219 are unique addresses for two different virtual drives.

Virtual drives *k* and *l* on staging drive H and virtual drives *e* and *f* on staging drive G are addressed only by CPU2 using virtual unit addresses 27E, 27F, 277, and 278 on primary channel 2.

Virtual drives (*c* and *d*) on staging drive G are addressed only by CPU1 using virtual unit addresses 114 and 115 on primary channel 1, or 214 and 215 on alternate channel 2.

Virtual drives *a* and *b* on staging drive G are shared (FEATURE=SHARED on IODEVICE card) by both CPUs. CPU1 addresses these two drives using virtual unit addresses 110 and 111 on primary channel 1, or 210 and 211 on alternate channel 2. CPU2 addresses these two drives using virtual unit addresses 273 and 274 on primary channel 2. Reserve and Release commands are issued by each CPU when mass storage volumes mounted on these drives are accessed for exclusive use.

Shared Unit Addresses

There are logical differences between real shared and virtual shared drives. For example, a real volume mounted on real shared drive B in Figure 7 is addressed as 102 by CPU1 and 242 by CPU2. Although each CPU uses a different channel address and Staging Adapter interface address, the real device address must be the same, that is 02. An electrical path must be available.

However, a virtual volume can be shared as long as each CPU has a virtual shared unit address available in the same staging drive group on which to mount the mass storage volume. A virtual volume can appear to be mounted on two different virtual drives at the same time, one for each CPU. For example, CPU1 addresses a virtual shared volume on virtual drive *a* as either 110 or 111; CPU2 addresses the same virtual shared volume on virtual drive *a* as either 273 or 274. Note that each CPU uses a unique address; that is, the channel, Staging Adapter interface, and virtual device addresses are all different. The path to the single virtual shared volume is determined by Staging Adapter table translations that map the two virtual shared addresses to the common data in the staging drive group.

As with real volumes, only one CPU can physically access the data at one time. When either CPU accesses a shared volume mounted on either a real shared drive or virtual shared drive for exclusive use, the Reserve and Release commands are used. Note that only the virtual volume is reserved rather than the entire staging pack.

Unit Addresses for Convertible-to-Real Drives

As the name implies, the function of a convertible-to-real drive can be converted from staging to real, and back again. Because its basic function is staging, a convertible-to-real drive (CONVREAL) must be specified as a member of a staging drive group.

Convertible-to-real drives E and F, and staging drives G and H in Figure 7 are specified as members of staging drive group 07. If drive E is varied online as real, it can function as a second real shared drive or as backup to real shared drive B. If drive E is varied online as staging, then any of the virtual and virtual shared unit addresses shown on staging drives G and H can map to staging drive E,

When convertible-to-real drives C, E and F are varied online as staging, their real unit addresses (101, 103, 243, and 245) must be varied offline in both CPUs. The operator should be cautioned not to select these offline real unit addresses if they are presented by dynamic allocation because the System Control Program will issue a request to mount a disk pack on the convertible-to-real drive that is being used for staging.

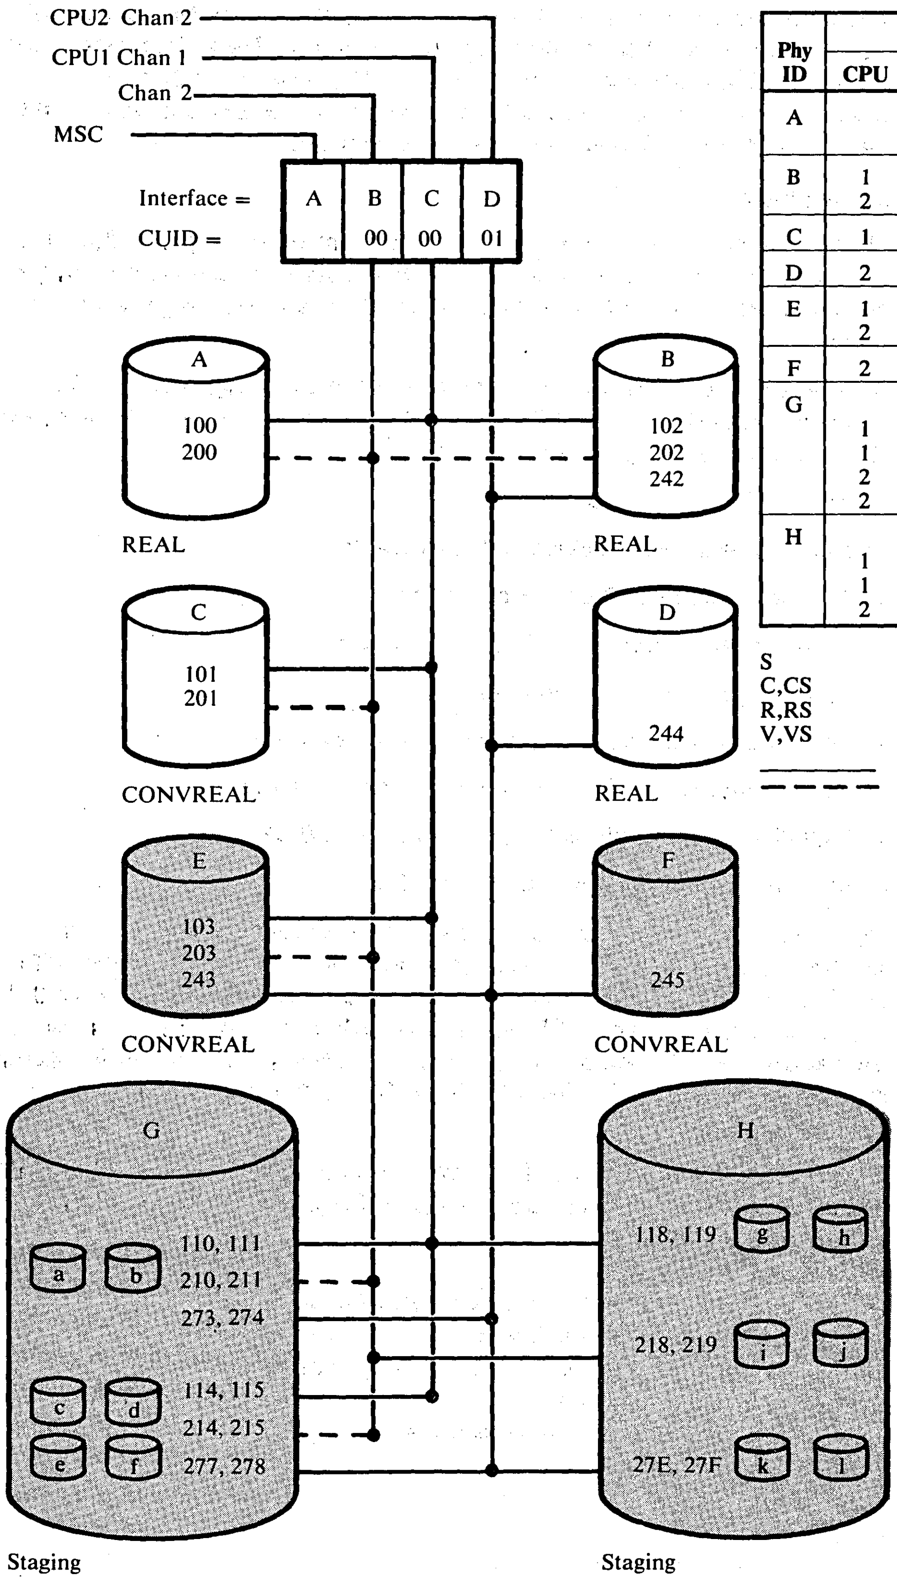

Phy ID	Unit Address			Drive	SDG
	CPU	PRI	ALT		
A		1	200	R	
		100			
B	1	102	202	RS	
	2	242		RS	
C	1	101	201	C	06
D	2	244		R	
E	1	103	203	CS	07
	2	243		CS	
F	2	245		C	07
G				S	07
	1	110-111	210-211	VS	
	1	114-115	214-215	V	
	2	273-274		VS	
	2	277-278		V	
H				S	07
	1	118-119		V	
	1	218-219		V	
	2	27E-27F		V	

S	Staging
C,CS	CONVREAL, CONVREAL Shared
R,RS	REAL, REAL Shared
V,VS	Virtual, VIrtual Shared
———	Primary path
– – –	Alternate path

▨ = staging drive group 07

| **Figure 7. Real and Virtual Unit Addresses**

The Mass Storage Control's primary and secondary tables (and their two alternate areas) must be mounted on staging drives. These four drives can not be specified CONVREAL. If possible, at least one nonshared convertible-to-real drive should be specified in a configuration for use with the Access Method Services CONVERTV command.

Unit Addresses for Other Control Units

Plan unit addresses carefully when different control unit types and models are connected to the same channel. Control units that respond to a range of unit addresses should be so plugged that their address ranges do not overlap.

Although we do not recommend connecting the following control units to the same channel, this example illustrates unit address ranges. Assume that the units are addressed in the order listed.

- 3830-2 (basic 16 device addresses)
- 3830-3 (64 device addresses on interface B)
- 3830-2 (expanded 32 device addresses)
- Mass Storage Control (single device address)

A valid assignment of unit addresses is:

- 3830-2 interface address 000, device addresses 00-0F
- 3830-3 interface address 01, device addresses 40-7F
- 3830-2 interface address 10, device address 80-9F
- Mass Storage Control device address C0

Note that both 3830-2s can be upgraded to a 3830-3 without creating address conflicts caused by overlapping addresses ranges. In addition, the 3830-2 real device addresses are within the real address range of a 3830-3; that is, either 00-1F or 80-9F, so that the real devices can be ad assigning new addresses. The Mass Storage Control address is beyond the 64 address range of the potential 3830-3, that is 80-BF, so that it is not affected by the upgrading.

The channel unit control words (UCW) should also be considered. They are discussed in a later chapter called "Related Logical and Physical Planning".

Alternate Channel Unit Addresses

For the alternate channel (OPCHAN) unit address to function, the Staging Adapter interface address on the primary and alternate channels must be the same. For CPU1 in Figure 7, the primary path is to interface C, and the alternate path is to interface B. Both interfaces have the same address, that is, CUID=00. Also, the alternate channel number must be higher than the primary channel number.

Summary

In summary, data paths are determined by electrical connections for real and convertible-to-real drives (varied online as real) and by table translations for virtual and virtual shared drives. When a CPU addresses a virtual or virtual shared drive, the actual path is determined by the Staging Adapter tables. The virtual unit address and virtual cylinders requested by the CPU are mapped to a real unit address and real cylinders of the staging drive containing the requested data. This Staging Adapter function is controlled by the Mass Storage Control and is unknown to the System Control Program.

Logical Annotation

This section describes the notation and identification of subsystem components. The logical configuration connections used in Figure 8 are similar to the Mass Storage Control Table Create Program Worksheet 1 in Appendix A, and defines five basic Mass Storage System interconnections as Connections I, II, III, IV, and V used in *OS/VS Mass Storage Control Table Create* program notation.

Notation and Identification

The CPUID is a 10-digit manufacturer's identification number in the form *abbbbbcccc* that indicates the plant of manufacture (*a*), the CPU serial number (*bbbbb*), and the CPU model (*cccc*). For example, a System/370 Model 158 can have a CPUID of 0200250158 if manufactured in Poughkeepsie, New York, or a CPUID of 510025158 if manufactured in Montpellier, France. The CPUID is wired on a circuit card and can be obtained by a read CPUID instruction. The customer engineer can obtain the CPUID for you from the Initial Microprogram Load record in the SYS1.LOGREC data set (OS/VS1 or OS/VS2 MVS) or you can display the CPUID at the console (OS/VS2 MVS only) with a DISPLAY CONFIGURATION command (DM=CPU). The CPUID specified as input to the Mass Storage Control Table Create program must be identical to the CPUID wired internally.

MSCINT specifies the primary, and ALTINT specifies the alternate, channel path to the Mass Storage Control shown as Connection I in Figure 8. A CPU can have one primary and one alternate channel path to the Mass Storage Control using a block multiplexer channel, or only one primary channel path using the byte multiplexer channel.

MPCPU is a manufacturer's identification for the second CPU in a Model 158, or Model 168 multiprocessor (MP). It is specified when (CPU0 and CPU1) are attached to the Mass Storage Control with two symmetric primary connections. With the multiprocessor connected symmetrically, CPU0 and CPU1 can respond to the Mass Storage Control either as a single host in a multiprocessing mode, or as two hosts in a uniprocessing mode.

A Mass Storage System configuration can have a maximum of two Mass Storage Facilities, identified as MSF0 and MSF1. The size of each Mass Storage Facility is defined by the decimal digit in the model number, such as the 1 in models A1 and B1, the 2 in models A2 and B2, and so on.

A Mass Storage Facility can have up to four data recording controls logically identified as DRC0 for Mass Storage Facility Models A1 and B1, DRC0 and DRC1 for Mass Storage Facility Models A2 and B2, and so on. Each DRC has two external connectors that permit connection to two different Staging Adapters, or any four DRCs (in either MSF0 or MSF1) can share a common connection to one Staging Adapter. CUxx (xx is a binary value) identifies each of the four DRC interfaces connected to a single Staging Adapter (see Figure 14, Example B). These addresses are specified to the Mass Storage Control Table Create program.

Connection II from the Mass Storage Control is the data path to the Mass Storage Control tables and the communication path to the accessor control pair (MSFID) in MSF0 (and MSF1) and to the channel A interface of each attached Staging Adapter. Although not shown in Figure 8, with the Mass Storage Control twin port feature installed, there is a Connection II communication-only path between the Mass Storage Control port 1 in MSF0 and MSF1 and channel A interface of each attached Staging Adapter. The standard Mass Storage Control port can address eight 3850 components. The twin port feature adds a second port that can address eight more. Each of the 16 logical components is defined by one hex digit (0-F). Hex 0 and 1 are reserved for the Staging Adapter SA=0 (and optionally SA=1) that controls the Mass Storage Control tables in staging drive group 00 and can only be used for this purpose. If Staging Adapter SA=1 is logically defined in a configuration, it must be *paired* with Staging Adapter SA=0, and the 3333 Disk Storage and Controls must be string-switched (a 3333 feature) to the Staging Adapter *pair* SA=0 and SA=1.

MSFID defines the accessor control pairs connected to Mass Storage Control port 0 of MSF0 and MSF1. It must be a value in the range of hex 2-7. This guide uses MSFID=7 for the accessor control pair in MSF0, and MSFID=6 for the accessor control pair in MSF1. The digit 6 can also be used to define a Staging Adapter (SA=6) in a configuration that has one Mass Storage Facility and seven Staging Adapters on the standard Mass Storage Control port 0.

Each Staging Adapter can have up to three interfaces connected to CPU channels. They are defined as CI=B, CI=C, and CI=D. The Model 158 and Model 168 ISC with Staging Adapter can have up to two interfaces defined as CI=B and CI=C.

Staging Adapter interface A is reserved for connection to the Mass Storage Control port. Each Staging Adapter is defined by hex digits 0-F (SA=0 through SA=F). The Staging Adapter block in Figure 8 represents either a 3830 Model 3 or a logical half of an ISC with Staging Adapter.

Up to four Staging Adapter interfaces (CUID) from two or more Staging Adapters can connect to a single CPU block multiplexer channel. The CUID (control unit identification) is a binary value (00-11) used to logically define the four Staging Adapters connected to a single CPU channel. For example, CUID=00 defines interface 00. *Only two of the three interfaces on one Staging Adapter can be connected to channels on the same* CPU *because of a Staging Adapter reset timing limitation.*

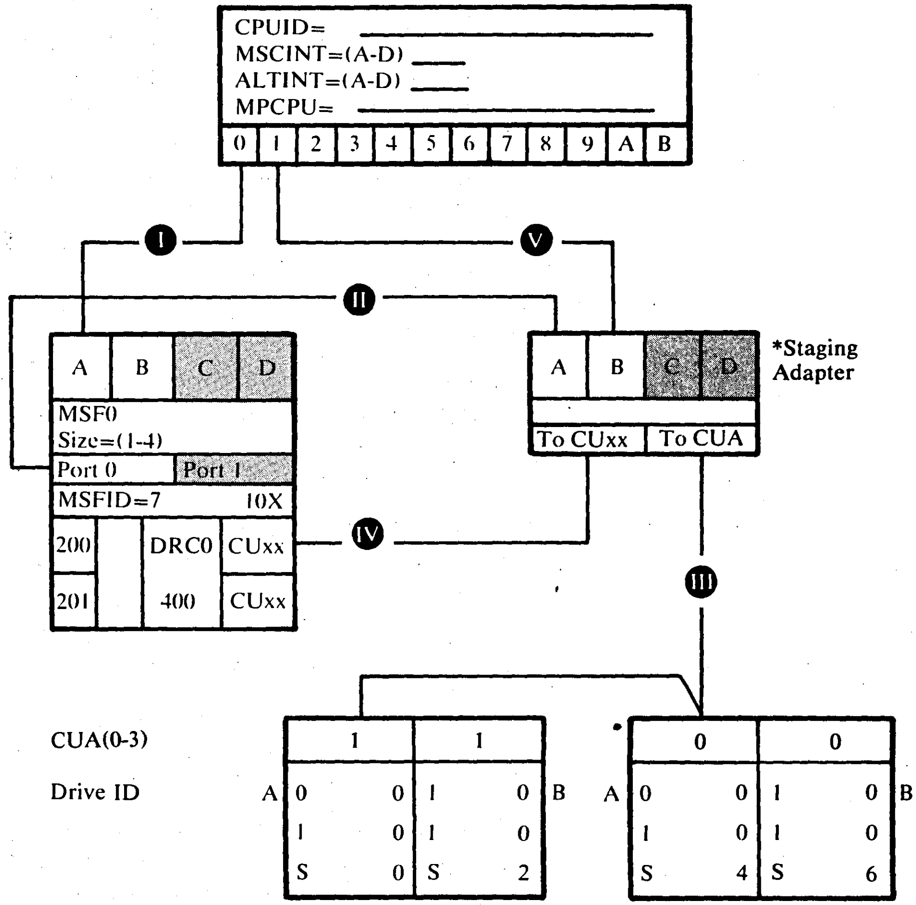

DASD Parameters

Physical			Logical
DEVn	(0-7)	(00-27)	SDGxx
Model	(1,11)	(0,1)	LCUz
Drive	*	(0-7)	LUAy

*S - Staging
R,RS - REAL, REAL Shared
C,CS - CONVREAL, CONVREAL Shared

▨ = Optional Feature

*Model 3 or logical half of ISC with Staging Adapter.

Note: The four drives in the minimum configuration do not include space for system-residence, paging, MSVI, or other system functions.

Figure 8. Logical Connections I-V with SSIDs

CUA (control unit address) is the single-digit (0-3) logical parameter used to logically define each of the four possible 3333 Disk Storage and Controls connected to a Staging Adapter. For example, CUA=0 indicates that the channel interface is to be plugged as address 00. Up to 32 physical drives are controlled by four 3333s. Sixteen logical 3330-1s can be specified in two staging drive groups. This requires 8 to 16 physical drives, depending on the 3330-1 and 3330-11 model mixture. The 8 to 16 physical drives used for staging in the two staging drive groups can all be specified CONVREAL, except for the four drives in staging drive group 00 containing the Mass Storage Control tables. The remaining physical drives can function only as real drives.

Mass Storage Control Table Create Program Worksheet 1 uses six parameters to physically and logically define the 3330 and 3350 drives. The physical parameters are DEVn, Model, and Function; the logical parameters are SDGxx, LCUz, and LUAy.

- DEVn defines the logical address plug inserted in the socket associated with each physical drive. The decimal value of n is between 0 and 7, that is, DEV0-DEV7.
- Model defines the drive as a 3333/3330 Model 1 or 2 (3330-1), a 3333/3330 Model 11 (3330-11) or a 3350 (3350).
- Drive defines each physical drive as REAL (R), CONVREAL (C), or staging (S). Shared drives (FEATURE=SHARED) are defined as RS or CS, that is REAL shared or CONVREAL shared.
- SDGxx defines up to 28 staging drive groups. The maximum number of staging drive groups that can be configured is 16 (see Figure 26). Each Staging Adapter or Staging Adapter *pair* manages one or two staging drive groups (see Figure 26). Decimal values of xx are 00-27.
- LCUz (logical control unit z) defines the two staging drive groups connected to each Staging Adapter or Staging Adapter *pair*. The decimal value of z is 0 or 1.
- LUAy (logical unit address y) defines up to eight logical 3330 Model 1 disk drives within each staging drive group (LCU0 and LCU1). The decimal value of y is 0-7.

Note: REAL, CONVREAL, DEVn, SDGxx, LCUz, and LUAy are Mass Storage Control Table Create notations.

Subsystem Identification SSID

Real drives and convertible-to-real drives varied online as real are addressed by real unit addresses; virtual drives are addressed by virtual unit addresses. Staging drives and other Mass Storage System components are addressed by subsystem identifiers or SSIDs. The operating system uses the SSID in S-type commands to address a Mass Storage System component. The 12-bit SSID format is:

Bits	0 3	4 8	9 11
	Component	SDGxx	LUAy

If bits 0-3 are zero, the Mass Storage System component is a staging drive. Bits 4-8 define up to 28 staging drive groups (0-27). Bits 9-11 define the logical unit address (LUAy) of the eight logical 3330-1 drives within the staging drive group.

For example, the SSID for staging drive LUA3, which is the fourth (binary 011) logical staging drive in staging drive group 07 (binary 00111), is hexadecimal 03B as shown below:

Bits	0 1 2 3	4 5 6 7	8 9 10 11
	Components	SDG07	LUA3
Binary Value	0 0 0 0	0 0 1 1	1 0 1 1
Hex	0	3	B

Drive H in Figure 9 illustrates this example. Figure 9 shows one staging drive (D) in staging drive group 06 (SDG06/LCU0), and four staging drives (E, F, G, and H) in staging drive group 07 (SDG07/LCU1) as it appears on Mass Storage Control Table Create Worksheet 1. The Address Table shows the drive and the physical, logical, and unit addresses for each CPU.

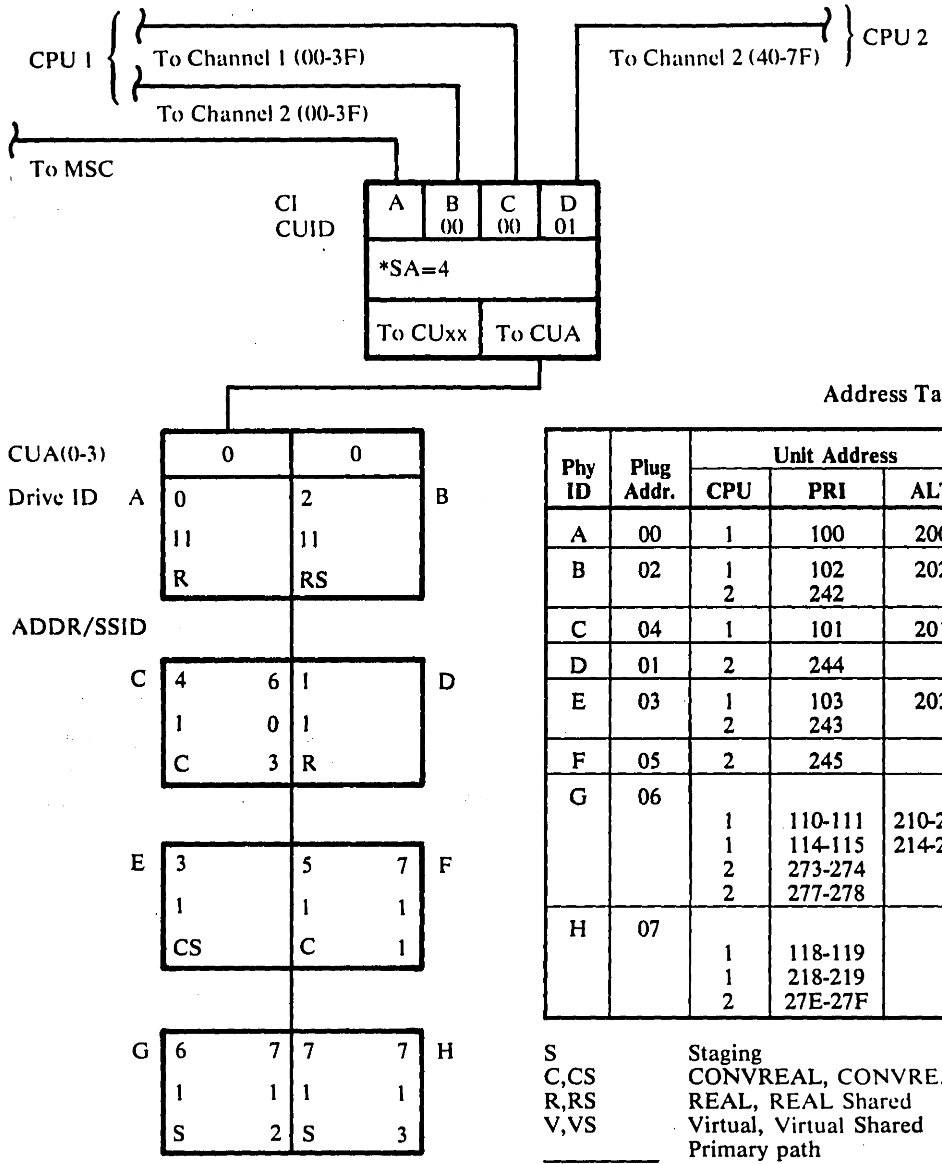

Address Table

Phy ID	Plug Addr.	Unit Address			Drive	SSID	SDG
		CPU	PRI	ALT			
A	00	1	100	200	R		
B	02	1	102	202	RS		
		2	242		RS		
C	04	1	101	201	C	033	06
D	01	2	244		R		
E	03	1	103	203	CS	038	07
		2	243		CS		
F	05	2	245		C	039	07
G	06				S	03A	07
		1	110-111	210-211	VS		
		1	114-115	214-215	V		
		2	273-274		VS		
		2	277-278		V		
H	07				S	03B	07
		1	118-119		V		
		1	218-219		V		
		2	27E-27F		V		

S	Staging
C,CS	CONVREAL, CONVREAL Shared
R,RS	REAL, REAL Shared
V,VS	Virtual, Virtual Shared
————	Primary path
— — — —	Alternate path

DASD Parameters

Physical			Logical
DEVn	(0-7)	(00-27)	SDGxx
Model	(1,11)	(0,1)	LCUz
Drive	*	(0-7)	LUAy

* S - Staging
R,RS - REAL, REAL Shared
C,CS - CONVREAL, CONVREAL Shared

Figure 9. Staging Drive Groups with SSIDs

Figure 10 shows the SSIDs for all Mass Storage Facility components, Staging Adapters, and the staging drives in the first 16 staging drive groups. The SSID for a staging drive varies depending on the 3330 model and the Mass Storage Control Table Create program specification.

The Mass Storage Control Table Create worksheet notations already described are used in the following chapters. Other notations are defined as they are introduced.

MSF Component	MSF0 SSID	MSF1 SSID
Accessor Controls 1 and 2	100	110
Accessor Control 1	101	111
Accessor Control 2	102	112
DRD 0	200	210
DRD 1	201	211
DRD 2	202	212
DRD 3	203	213
DRD 4	204	214
DRD 5	205	215
DRD 6	206	216
DRD 7	207	217
DRC 0	400	410
DRC 1	401	411
DRC 2	402	412
DRC 3	403	413

Staging Adapter	
Notation	SSID
SA=0	800
SA=1	810
SA=2	820
SA=3	830
SA=4	840
SA=5	850
SA=6	860
SA=7	870
SA=8	880
SA=9	890
SA=A	8A0
SA=B	8B0
SA=C	8C0
SA=D	8D0
SA=E	8E0
SA=F	8F0

Staging Drive	
Group	SSID
SDG00	000-007
SDG01	008-00F
SDG02	010-017
SDG03	018-01F
SDG04	020-027
SDG05	028-02F
SDG06	030-037
SDG07	038-03F
SDG08	040-047
SDG09	048-04F
SDG10	050-057
SDG11	058-05F
SDG12	060-067
SDG13	068-06F
SDG14	070-077
SDG15	078-07F

Notes:

1. Maximum of 14 Staging Adapters. Logical 6 and 7 optionally reserved for MSFID.
2. Mass Storage Control Table Create program accepts SDG00 to SDG27.

Figure 10. MSF, Staging Adapter, and Staging Drive SSIDs

Logical Connections

This chapter describes in detail first the five basic logical connections of a Mass Storage System configuration with uniprocessors and a single Mass Storage Facility, then miscellaneous connections, and the physical and logical location of key Mass Storage System data sets. Second, it extends some of the basic connections to the more complex environment of multi-Mass Storage Facilities, attached processors, and multiprocessors. Several figures reinforce the text.

At the end of the chapter is a summary of the logical connections and configuration requirements.

Basic Connections

Figure 8 in the previous chapter shows a minimum logical Mass Storage System configuration (a Mass Storage Facility Model A1, one Staging Adapter, and two 3333s) and labels the five basic connections as:

- Connection I, CPU to Mass Storage Facility
- Connection II, Mass Storage Facility to Staging Adapter
- Connection III, Staging Adapter to Disk Storage and Controls
- Connection IV, Staging Adapter to Mass Storage Facility
- Connection V, CPU to Staging Adapter

Note: The four drives in the minimum configuration do not include space for system-residence, paging, MSVI, or other system functions.

Connection I, CPU to Mass Storage Facility

Connection I is the communication path between the CPU and the active Mass Storage Control in MSF0 or MSF1. Figure 11 illustrates the minimum and maximum CPU-to-Mass Storage Facility connections for uniprocessor systems. Example A is a typical connection with one primary channel path on a byte multiplexer channel (CPU channel 0). Example B shows two CPUs, one with a primary and an alternate channel path on a block multiplexer channel (CPU channels 3 and 4). A maximum of four CPUs can be connected to the Mass Storage Control as shown in Example C, each with one primary channel path. Alternate channel paths on byte or mixed byte and block multiplexer channels are not supported by OS/VS1 or OS/VS2 MVS.

Because the Mass Storage Control can function only in a multiplexer mode, its interfaces cannot connect to a selector channel or to the 2870 Byte Multiplexer Selector Subchannel (Feature Code 6990-6993).

A Mass Storage Facility Model B has two Mass Storage Controls and requires two positions on the byte or block multiplexer channel to which it is connected. Therefore, a unique address must be defined for each Mass Storage Control in the SYSGEN.

Connection I is established by the Mass Storage Control Table Create CPUCONFG command, as specified on the Mass Storage Control Table Create Program Worksheet 7. An IODEVICE card is required to specify the unit address for each Mass Storage Control and include Mass Storage System support in the SYSGEN.

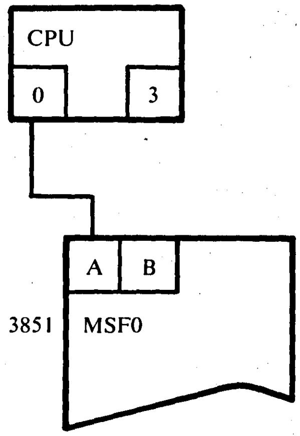

Example A. Byte Multiplexer Channel (primary connection only).

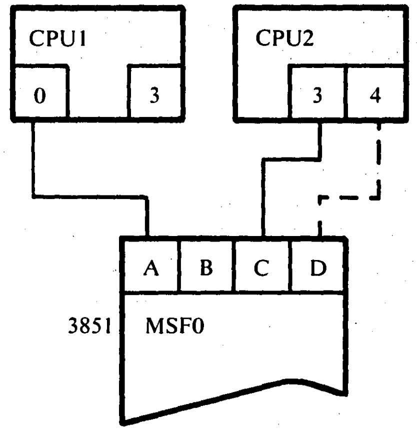

Example B. Block Multiplexer Channels (primary and alternate connection).

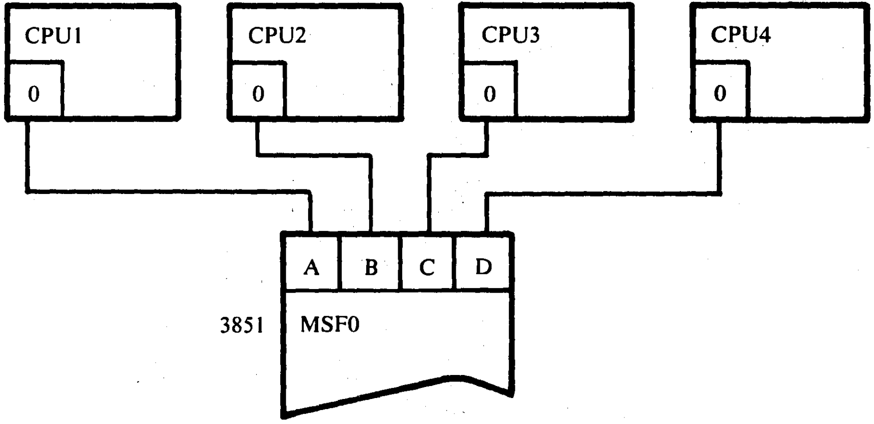

Example C. Maximum of four CPUs may attach.

— — — — Alternate Path

Figure 11. Connection I

Connection II, Mass Storage Facility to Staging Adapter

The Mass Storage Control communicates with the Staging Adapters and Mass Storage Facility accessor control pairs using the paths from port 0 (standard) and port 1 (optional), which are identified as Connection II. (See "Miscellaneous Connections" in this chapter for more details about the accessor control connections.) The Mass Storage Control uses the port 0 path to read and write its primary and secondary tables. Each Mass Storage Control port can address up to eight Mass Storage System components (such as Staging Adapters and accessor controls), serially connected on a standard interface cable.

The Mass Storage Control addresses up to 16 logical Mass Storage System components as hex 0-F, with hex 0-7 on the standard port 0, and hex 8-F on the optional port 1. Logical address 0 is reserved for the first logical Staging Adapter (SA=0) on port 0, and logical address 1 is reserved for the second logical Staging Adapter (SA=1) on port 0. These two logical Staging Adapters connect to the staging drives that have the Mass Storage Control primary and secondary tables in staging drive group 00. The rest of the logical addresses (2-F) are specified for the remaining Staging Adapters. Of these, only addresses 2-7 can be specified for the Mass Storage Facility accessor controls (MSFID).

Figure 12 illustrates one Mass Storage Facility (MSF0, MSFID=7) connected to six Staging Adapters (SA=0 through SA=5), for a total of seven logical components. Logical address 6 is used in this publication for the second Mass Storage Facility (MSF1, MSFID=6). The twin port feature is required whenever the total number of Mass Storage Facility components exceeds eight. The two logical halves of the ISC with Staging Adapter count as two components.

Connection II is established by the logical addresses specified for the accessor controls (on Worksheet 7) and Staging Adapters (on Worksheet 2).

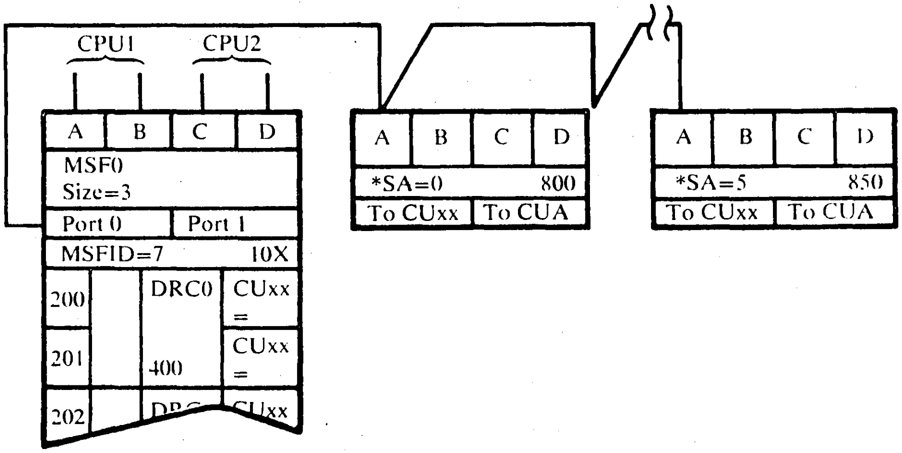

*Staging Adapter in 3830-3 or ISCSA.

Figure 12. Connection II

Connection III, Staging Adapter to Disk Storage and Controls

Connection III is the portion of the stage and/or destage data path that exists between the Staging Adapter buffer and the 3330 staging drives. The CPU uses this same portion (as illustrated in Figure 4) to access data that is on either real or staging drives.

A Staging Adapter has two ports, one for connection to the 3333 or 3350 (Connection III to CUA), and one for connection to the DRCs in the Mass Storage Facility (Connection IV to CUxx). Each port serially connects to its four devices on a standard interface cable. Figure 13 shows Staging Adapter Connections III without (Example A) and with (Example B) string-switched Disk Storage and Controls (the string-switch feature is required on all 3333s and is optional on the 3350-A2).

Without the 3333 String Switch Feature, each Staging Adapter manages up to 1600 megabytes of staging space (two staging drive groups, each with a maximum of eight logical 3330-1 drives) in addition to 16-24 real drives. This provides one stage and/or destage path (Connections III and IV) and up to three CPU access paths (Connection V) to each staging drive group pair through a single Staging Adapter.

With the 3333 string-switch feature, each Staging Adapter pair manages up to 1600 megabytes of staging space and 16-24 real drives. This provides two stage and/or destage paths (Connections III and IV) and up to six CPU access paths (Connection V) to each staging drive group pair through two Staging Adapters. String-switching reduces the potential number of staging drive groups but provides multiple stage, destage, and CPU access paths to each staging drive group pair.

With the appropriate configuration and connections, an inoperable DRC or a Staging Adapter that has been varied offline can degrade system performance but not prevent CPU or Mass Storage Facility access to any shared staging drive group. Depending on system activity, data can be staged and/or destaged to string-switched staging drive groups while the CPUs are simultaneously accessing other data in the same staging drive group, provided the two sets of data do not reside on the same virtual volume.

As shown in Figure 13, Example B, if one 3333 controlling staging drives is string-switched to two Staging Adapters (SA=0 and SA=1), then all 3333s controlling staging drives connected to either Staging Adapter (SA=0 or SA=1) must be string-switched. The 3350 can be accessed through Staging Adapter (SA=1) and optionally through SA=0.

A Staging Adapter pair can consist of two 3830-3s, one 3830-3 and a logical half of one ISC with Staging Adapter, or two logical halves of one ISC with Staging Adapter. IBM does not support the use of a 3830-2 as a member of a pair that controls staging drives. See the chapter "Related Logical and Physical Planning" for connecting real drive strings.

Connection III is established by the Mass Storage Control Table Create DASD command, as specified on Mass Storage Control Table Create Program Worksheet 2.

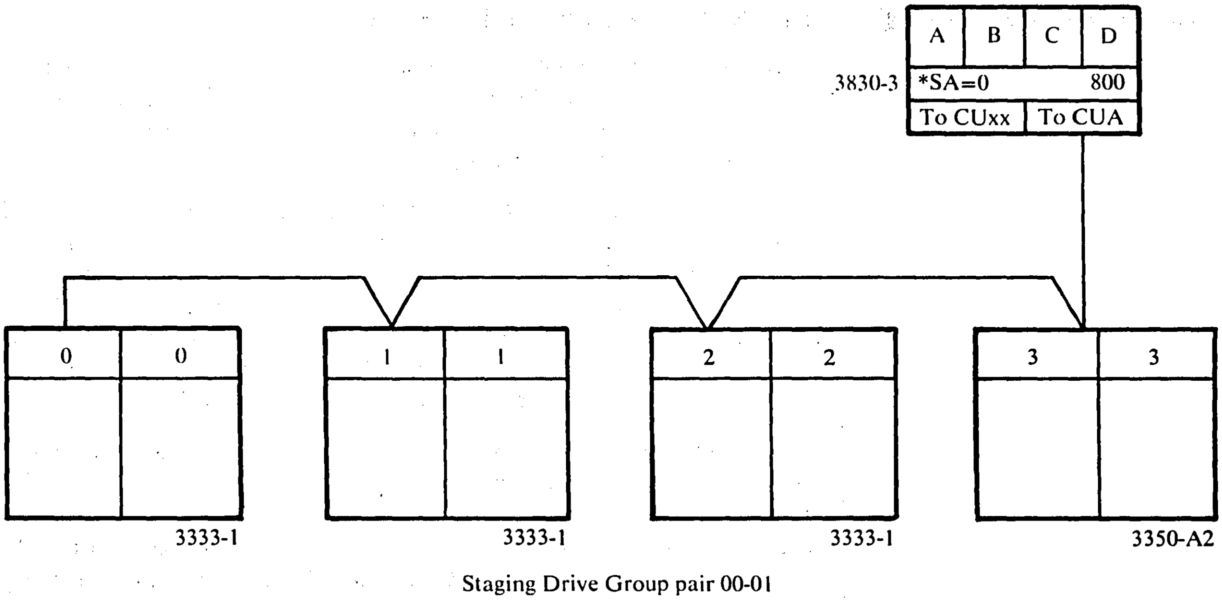

Example A. Staging Adapter with Unshared 3333s and 3350 Disk Storage and Controls

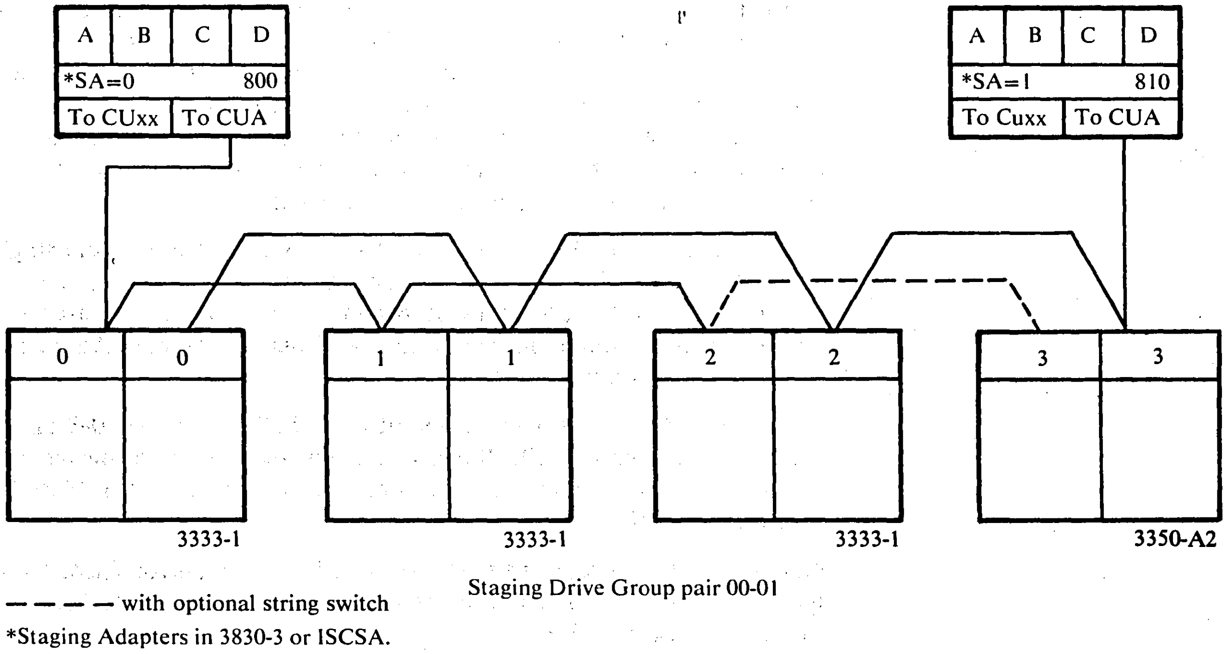

Example B. Staging Adapter Pair Sharing 3333s and 3350 Disk Storage and Controls

Figure 13. Connection III

Connection IV, Staging Adapter to Mass Storage Facility

Connection IV is the stage and/or destage data path between the DRC and the Staging Adapter buffer. The lower port of a Staging Adapter (Connection IV to CUxx) can serially connect to a maximum of four DRCs in one or two Mass Storage Facilities.

Figure 14 shows some of the possible DRC connections. Each DRC has two connectors and can connect to a single Staging Adapter (Example A), or to two Staging Adapters (Example C). DRC0 in Example C can operate with only one Staging Adapter at a time. Up to four DRCs can share one Staging Adapter (Example B). If more than one DRC is connected to a Staging Adapter, only one DRC-to-Staging Adapter connection can be active (reading or writing) at a time. The DRC connector addresses are plugged by the Customer Engineer as they were specified to the Mass Storage Control Table Create program.

If two DRCs in one or two Mass Storage Facilities are connected to a Staging Adapter pair that have string-switched staging drive groups, both DRC-to-Staging Adapter paths can be active at the same time.

The Staging Adapter connection to DRC0 in Example A of Figure 14 can be extended to the second connector using a unique address, that is, 01, 10, or 11. This provides an alternative connection in the event of a cable connector or address decoding problem. However, only one DRC0 connector path is active at a time.

Connection IV is established by the Mass Storage Control Table Create LOWERCON command, specified on Mass Storage Control Table Create Program Worksheet 3.

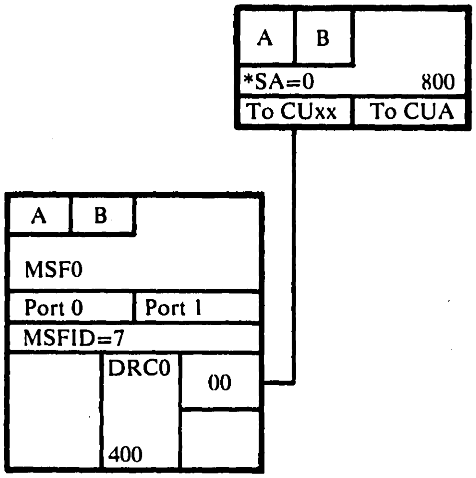

Example A. DRC connected to one Staging Adapter

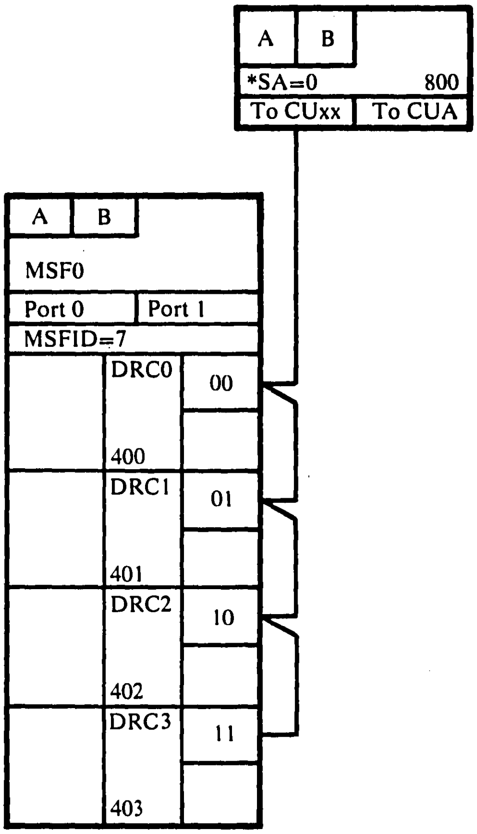

Example B. Four DRCs on a Shared DRC path to one Staging Adapter

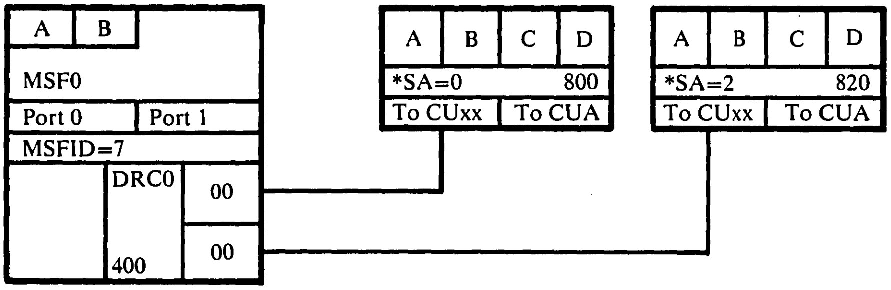

*Staging Adapters in 3830-3 or ISCSA.

Example C. DRC Connected to two Staging Adapters

Figure 14. Connection IV

Connections III and IV, String-Switched Staging Drive Group Pairs

Up to this point, the connections have been examined separately. Consider Connections III and IV together when you plan stage and destage paths to string-switched staging drive group pairs.

Figure 15 shows the recommended connections from the DRCs to Staging Adapters that have string-switched staging drive group pairs. It shows two stage and destage paths from MSF0 through Staging Adapter pair SA=0 and SA=1 to staging drive group pair 00-01 with DRC0 connecting to SA=0 and DRC1 connecting to SA=1. The other two stage and destage paths from MSF0 are through Staging Adapter pair SA=2 and SA=3 to staging drive group pair 02-03 with DRC0 connecting to SA=2 and DRC1 connecting to SA=3.

If you specify primary and alternate CPU channel-to-Staging Adapter paths and an alternate control unit, or multiple primary paths, staging drive group pair 00-01 can normally be accessed through Staging Adapter SA=0 or SA=1, and staging drive group pair 02-03 can normally be accessed through Staging Adapter SA=2 or SA=3.

With the connections as shown in Figure 15, stage or destage activity to staging drive group pairs 00-01 and 02-03 can be shared through DRC1 if DRC0 is not operating. If both connections of DRC0 had been attached to Staging Adapter SA=0 and SA=1, then the loss of DRC0 would eliminate all stage or destage activity to staging drive group pair 00-01.

The Mass Storage Facilities must have a stage and destage path to every staging drive group, either through one Staging Adapter or through a Staging Adapter pair.

Connection III is established by using the DASD command, as specified on Mass Storage Control Table Create Program Worksheet 2. Connection IV is established by using the LOWERCON command, as specified on the Mass Storage Control Table Create Program Worksheet 3.

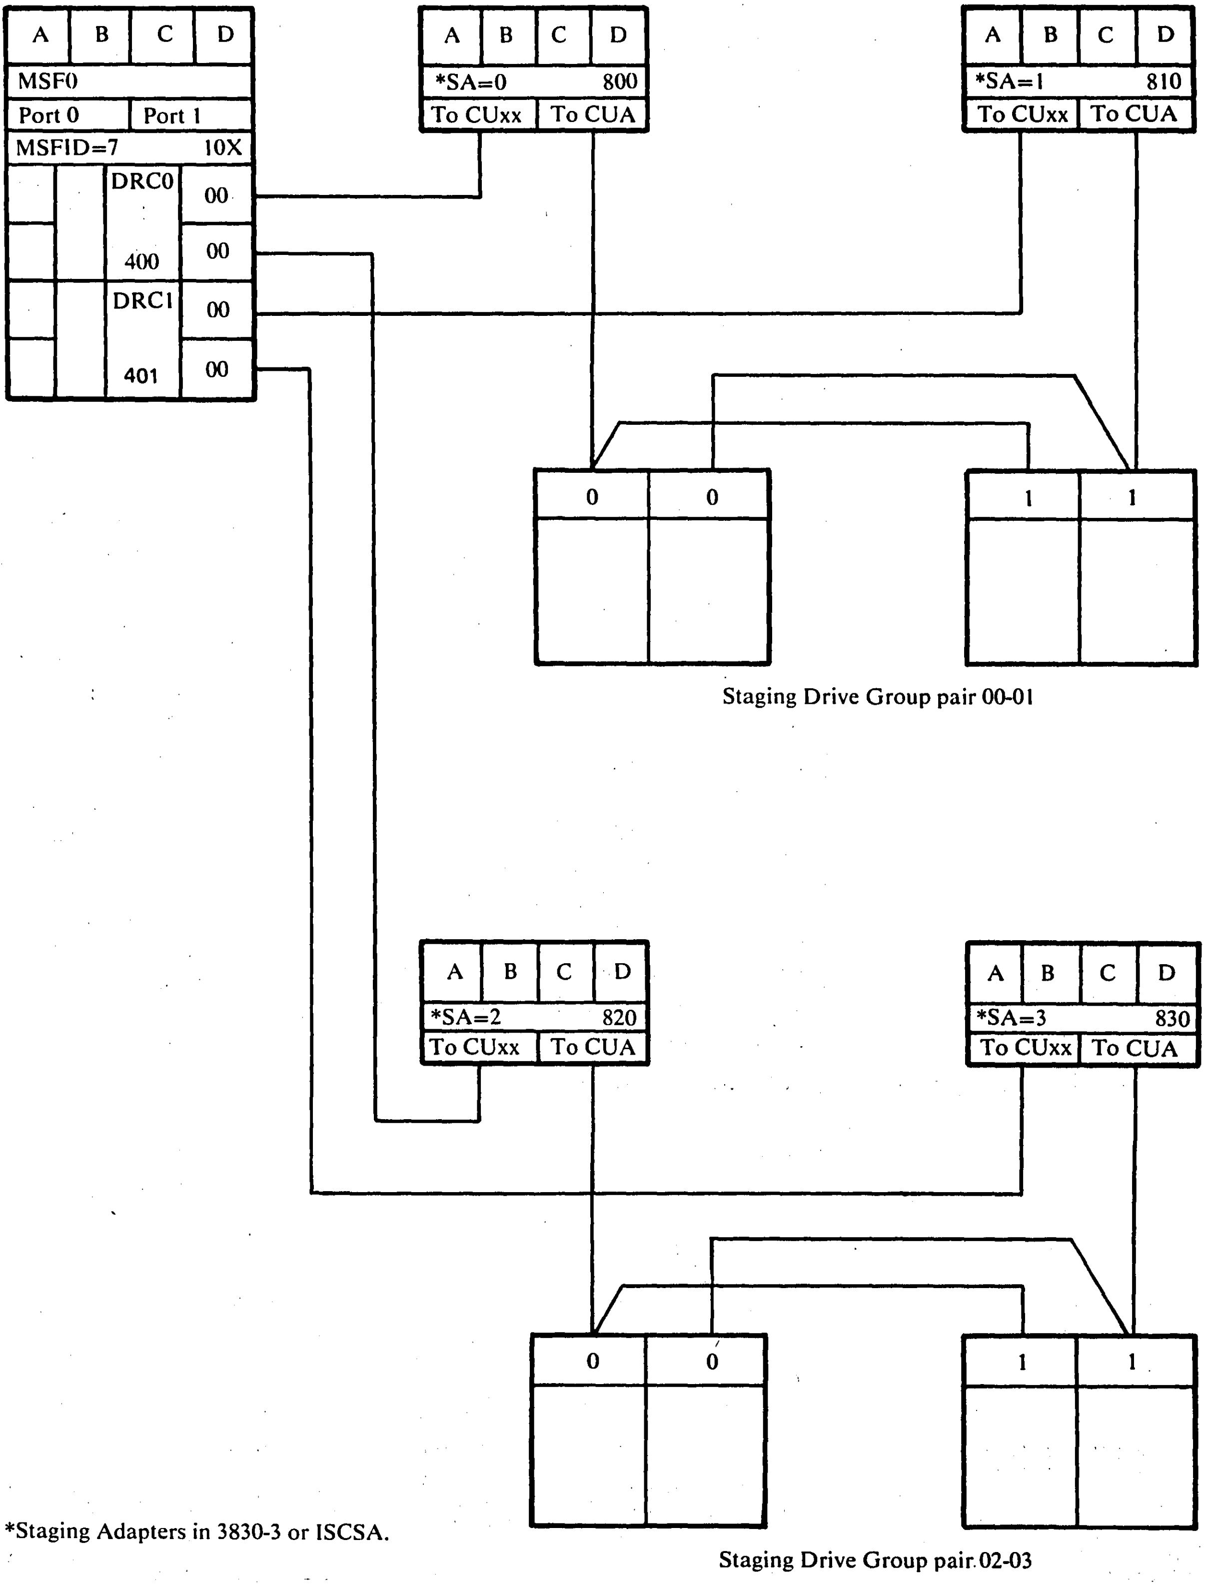

*Staging Adapters in 3830-3 or ISCSA.

Figure 15. Connections III and IV

Connection V, CPU to Staging Adapter

Connection V is the path used by the CPU channel to access virtual volumes created when a mass storage volume is mounted on staging drives, or to access real volumes that are mounted by the operator on real drives.

Every CPU that is connected to a Staging Adapter must also be connected to the Mass Storage Control. Except in JES3 environments, every CPU must have access to every staging drive group. A CPU can address up to 64 virtual or real drives on each Staging Adapter interface (primary or alternate) that is connected to one of its channels.

A single Staging Adapter provides up to three CPU paths (primary or alternate) to the data in a staging drive group pair, or up to six paths using Staging Adapter pairs. The maximum number of paths that you plan may be affected by:

- 3830 Model 3 limitations that permit only two interfaces to be connected to channels in a single CPU
- OS/VS1 and OS/VS2 MVS limitations on the maximum number of alternate paths

If more than two interfaces of a Staging Adapter are connected to a single CPU, pressing the System Reset and/or Load keys more than once may be required, because a timeout can occur before the Staging Adapter interfaces are reset.

The alternate path selection logic and the number of alternate paths are different between OS/VS1 (three alternate paths) and OS/VS2 MVS (one alternate path). For example, in Figure 16, Example F, OS/VS2 MVS in CPU1 uses the alternate channel 2 path through the alternate control unit when the channel is busy but does not use it when the control unit is busy unless FEATURE=(ALTCTRL) has been specified on the IODEVICE card. OS/VS1 uses the alternate path on either a channel busy or control unit busy condition without any specification. With either System Control Program, the alternate channel path must have a higher address than the primary channel path, and the Staging Adapter interface address (CUID) must be the same for both paths.

Figure 16 illustrates the variation in primary and alternate paths. Example A shows a primary path (100-13F) with an alternate path (200-23F) to address up to 64 drives. Example C appears identical but is actually two primary paths (100-13F and 240-24F) to address up to 128 drives. Example B has two primary paths (100-13F and 140-17F), each with an alternate (200-23F and 340-37F) to address up to 128 drives.

Multiple primary paths to a staging drive group pair increase the maximum number of unit addresses available, as shown in Examples B, C (128 addresses), and D (256 addresses). All connections shown in Figure 16, except Example E, can be used in OS/VS1 or OS/VS2 MVS. Examples F and A show typical connections with and without string-switched staging drives; Example F illustrates a special case, because staging drive group 00 contains the Mass Storage Control tables. Note that both CPUs are connected so that each has direct access to the Mass Storage Control tables through Staging Adapter interface B. The interface is reserved for this purpose. Refer to the chapter "Data Security and Recovery Planning" in this manual for a description of the direct access to the Mass Storage Control tables through interface B of Staging Adapters 0 and 1 and optionally through interfaces C and D.

Connection V is established by the Mass Storage Control Table Create UPPERCON command, as specified on Mass Storage Control Table Program Worksheet 5.

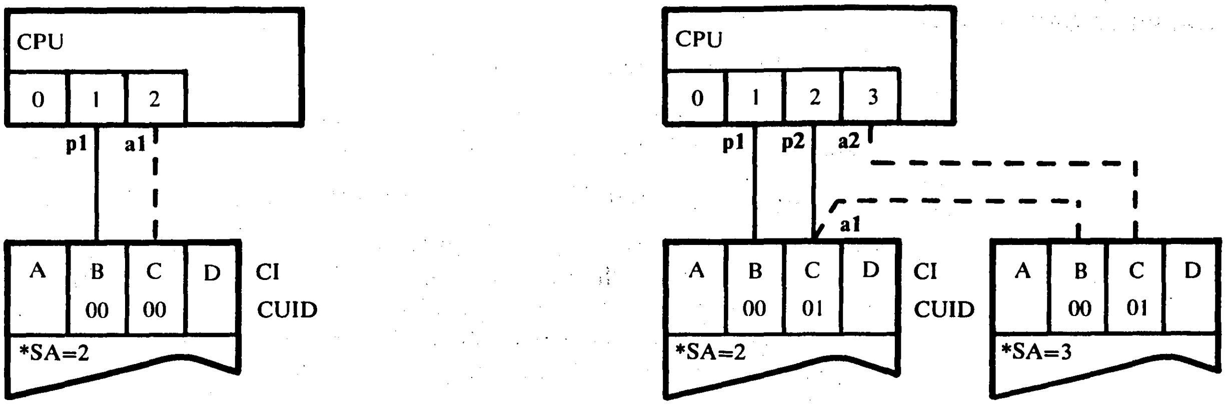

Staging Drive Group pair 02-03

Example A. Primary and Alternate Path

Staging Drive Group pair 02-03

Example B. Two Primary and Two Alternate Paths

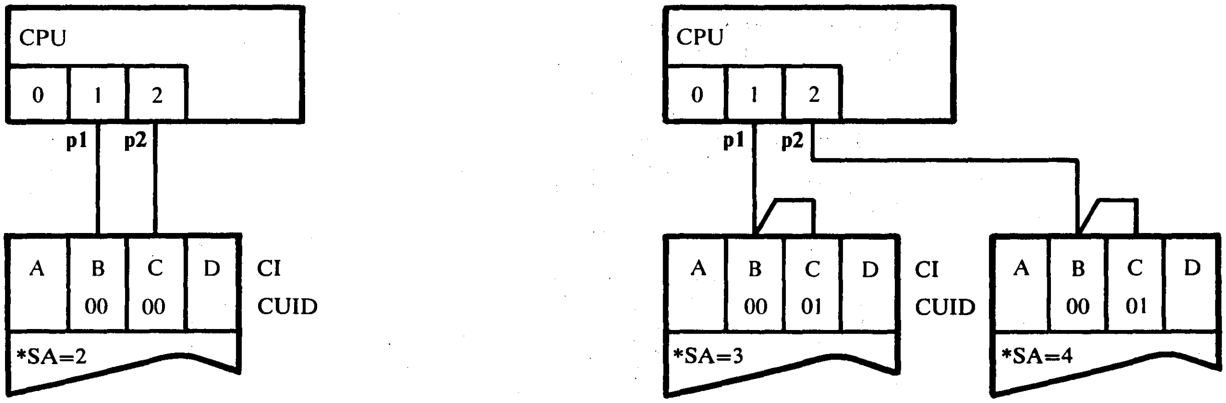

Staging Drive Group Pair 03-04

Example C. Two Primary Paths

Staging Drive Group pair 03-04

Example D. Four Primary Paths

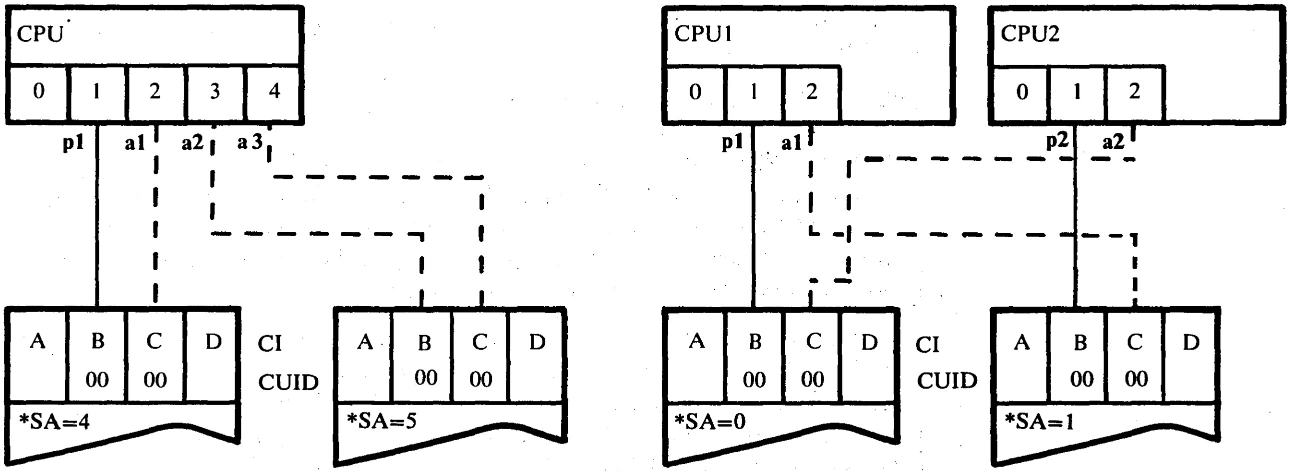

Staging Drive Group pair 05-06

*Staging Adapter in 3830-3 or ISC

Example E. Primary and Three Alternate Paths (OS/VS1 only)

Staging Drive Group pair 00-01

— — — Alternate Path

Example F. Shared Staging Drive Groups

Figure 16. Connection V

Miscellaneous Connections

The miscellaneous connections are:

- Mass Storage Control to the accessor control pair on port 0 of a Mass Storage Facility Model A or B
- Data Recording Control (DRC) to the Data Recording Devices (DRD)

Figure 17 illustrates ports 0 and 1 connections. Port 0 is cabled to the accessor control within the Mass Storage Facility at the manufacturing plant. The Staging Adapters are cabled by the customer engineer when the Mass Storage Facility is installed. The logical address (0-7 or 8-F) determines which port the Mass Storage Control uses to address a particular component.

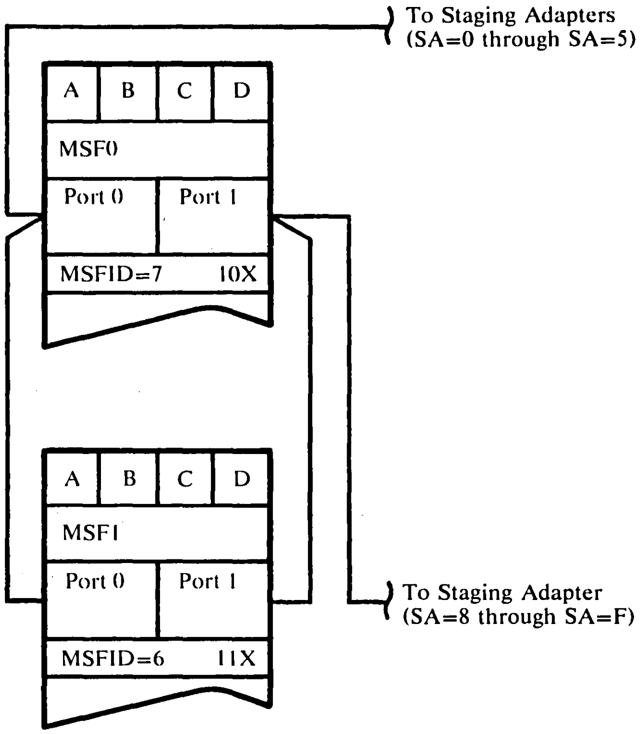

Note: As shown, MSF0 is specified as secondary (9121); MSF1 is specified as primary (9120).

Figure 17. Accessor Control Connections

With two Mass Storage Facilities, the port 0 cable from MSF0 must also connect to port 0 of MSF1. Similarly, the port 1 (optional feature) cable from MSF0 must also connect to port 1 of MSF1. MSF0 has cable-in and cable-out connections on both ports because it is specified as secondary (9121). MSF1 has only cable-out connections.

Figure 18 shows the DRC to DRD connections for all Mass Storage Facility models. The DRC and DRD addresses shown are the SSIDs for MSF0. These are standard connections that the Customer Engineer completes during installation, or modifies during a model change. They are indicated on the Mass Storage Control Table Create Connections Report by SSID.

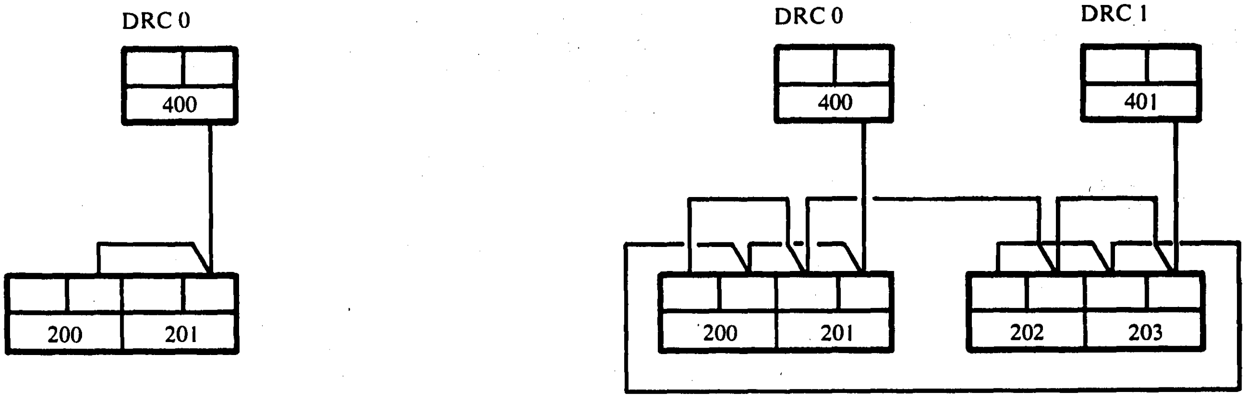

Example A. Model A1 or B1

Example B. Model A2 or B2

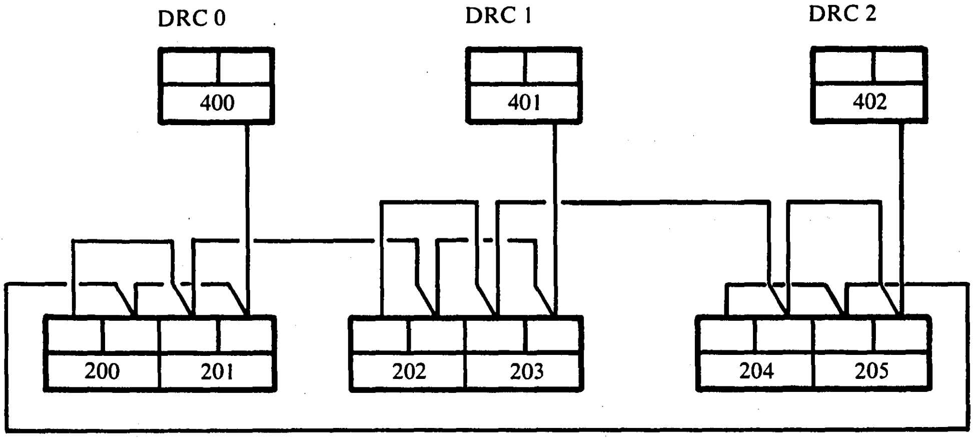

Example C. Model A3 or B3

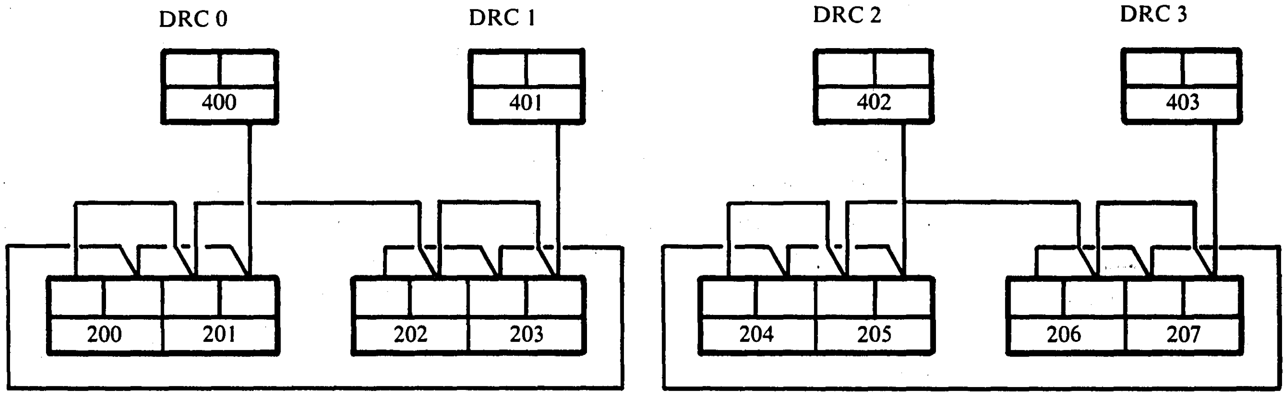

Example D. Model A4 or B4

Figure 18. DRC to DRD Connections

Except in Mass Storage Facility Models A1 and B1, every DRC is connected to four DRDs. The method of connecting a DRC to four DRDs is different for Models A3 and B3 from that for Models A2, B2, A4, and B4. As shown in Example D, if the Mass Storage Control selects DRC0 (SSID 400) in a Mass Storage Facility Model A3 or B3 as a stage or destage path, it has a choice of four DRDs (SSIDs 200, 201, 204, and 205) in which to load the data cartridge. If a DRC failure occurs, these connections provide an alternate DRC path to the same data without having to move the data cartridge to another DRD.

Data Sets For the Mass Storage System

Three key data sets are introduced with the Mass Storage System:

- Mass Storage Control tables
- Mass Storage Volume Inventory (MSVI)
- Mass Storage Volume Control Journal (MSVCJRNL)

The location of the Mass Storage Control tables is specified, but the Inventory and Journal data set locations are installation options.

Figure 19 shows the logical and physical locations of the primary, secondary, and alternate Mass Storage Control tables. The names of physical drive 3333-0, physical drive 0, physical drive 3333-1, and physical drive 1 in *OS/VS Mass Storage Control Table Create* refer to the logical address plugs in the drives that contain the Mass Storage Control tables. The Mass Storage Control Table Create program uses SDGxx, LCUz, and LUAy as logical references. The table below shows these references and the suggested labels for the logical address plugs.

MSC Tables	Logical Location			Physical Location			Logical Address Plug ADDR/SSID
	SDGxx	LCUz	LUAy	*SA=	3333	Drive	
Primary	0	0	0	0	0	0	00/000
Secondary	0	0	4	0	1	0	08/004
Alternate 1	0	0	2	0	0	1	01/002
Alternate 2	0	0	6	0	1	1	09/006

*Staging Adapter in a 3830-3 or an ISCSA.

SA=1 is optional.

Figure 19. Mass Storage Control Tables Locations and Addresses

A logical address plug must be inserted in each drive. Labels should be made to display the real drive address or the ADDR and SSID of the convertible-to-real or staging drive. The SSID depends on the staging drive group and logical unit address assigned to the drive.

The chart below shows the unit address for the real or convertible-to-real drives connected to each 3333 Disk Storage and Control.

3333 (CUA)	Unit Address
0	00-07
1	08-0F
2	10-17
3	18-1F

The Mass Storage Control tables are shown as they might appear on Mass Storage Control Table Create Program Worksheet 1. Figure 20 shows table locations without string-switched staging drive groups; Figure 21 shows them with. These two figures also show the Mass Storage Control tables on either a 3330-1 or a 3330-11 drive. Note that the SSIDs for the Mass Storage Control tables must be the same whether you use 3330-1 or 3330-11 drives.

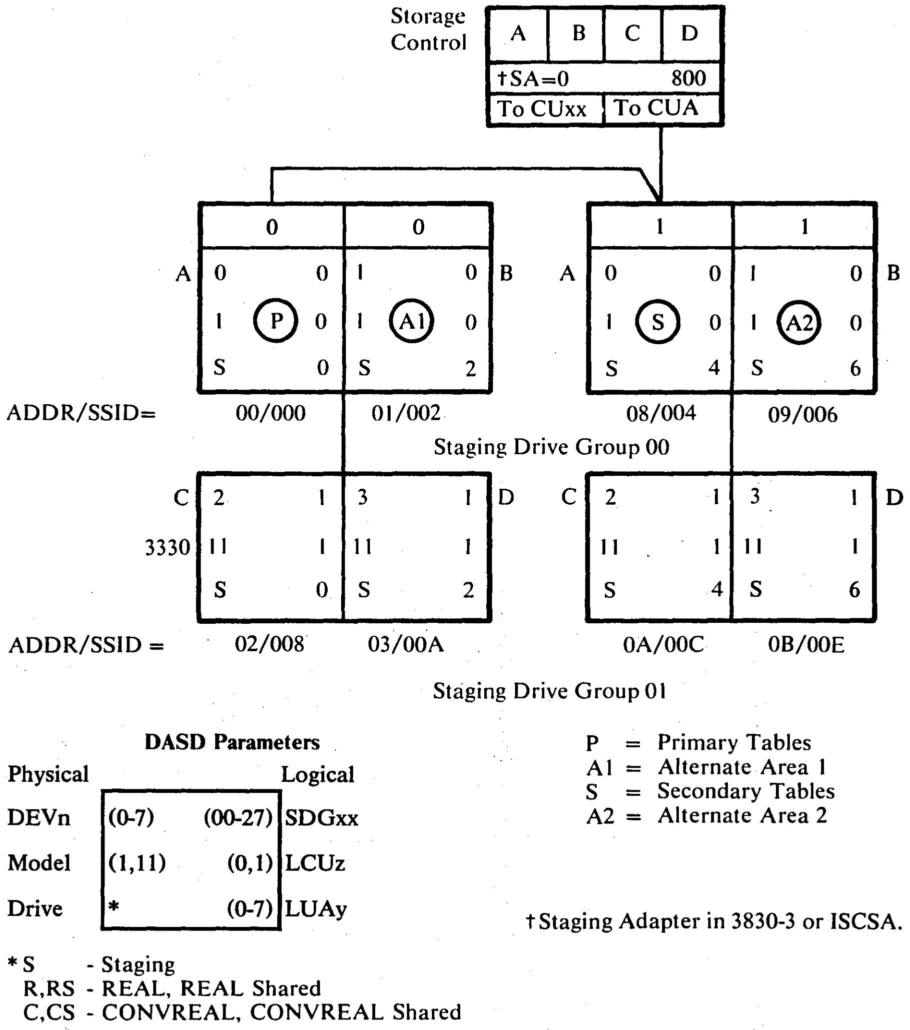

Figure 20. Physical and Logical Mass Storage Control Table Locations in a Staging Drive Group *without* String-Switching

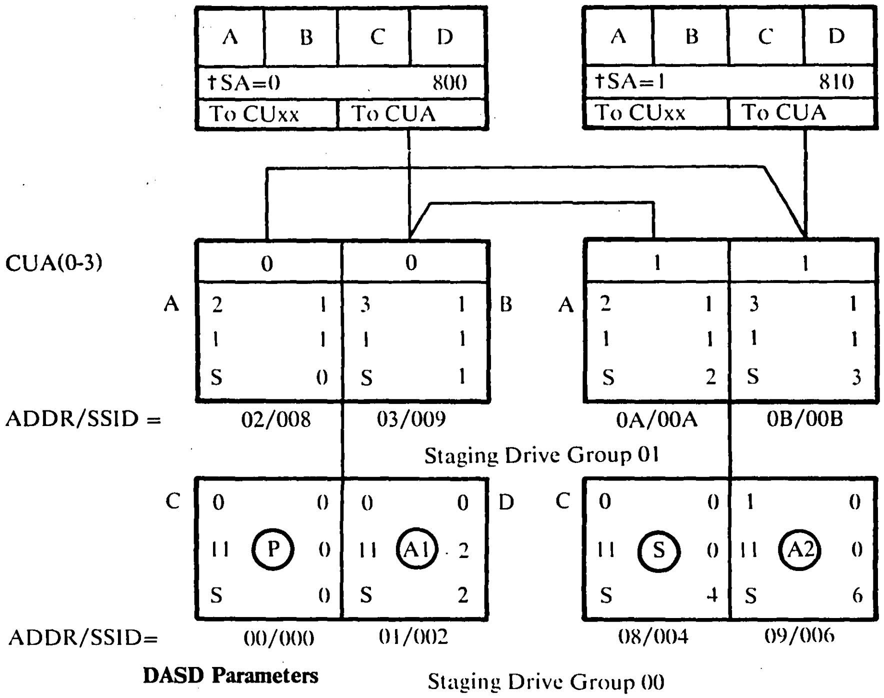

Physical			Logical
DEVn	(0-7)	(00-27)	SDGxx
Model	(1,11)	(0,1)	LCUz
Drive	*	(0-7)	LUAy

*S - Staging
R,RS - REAL, REAL Shared
C,CS - CONVREAL, CONVREAL Shared

P = Primary Tables
A1 = Alternate Area 1
S = Secondary Tables
A2 = Alternate Area 2

† Staging Adapter in 3830-3 or ISCSA.

Figure 21. Physical and Logical Mass Storage Control Table Locations in a Staging Drive Group *with* String-Switching

The Mass Storage Control Table Create program identifies the logical address plugs as DEV0 through DEV7. These are converted into the ADDR and SSID shown in the two figures and as they should appear on the logical address plug labels.

For a minimum Mass Storage System configuration, the Mass Storage Control tables are placed under a single Staging Adapter (Figure 20). We recommend using a Staging Adapter *pair* (Figure 21) to increase accessibility to the Mass Storage Control tables by providing up to three alternate CPU access paths and an alternate stage and/or destage path for Mass Storage Facility Models A2, B2 and above.

Thirty-two cylinders of space (virtual pages 25, 26, 27, and 28) are required on each of the four staging drives for the Mass Storage Control tables and alternate areas. The remaining pages on the four drives are available for use as staging drive space. These four staging drives must be defined as staging. For additional information on direct access to the Mass Storage Control tables, see the "Data Security and Recovery" section of this manual.

In large CPU configurations (such as three Model 158s or two 168s), consider dedicating the primary Mass Storage Control table staging drive to Mass Storage Control use by varying offline the staging drive that contains the primary Mass Storage Control tables. This is usually SSID 000 but can be 002, 004, or 006 depending on the current location of the primary Mass Storage Control tables. In this way, the Mass Storage Control can continue to access the primary Mass Storage Control tables and VTOC but does not use the remaining pages for staging drive space. The remaining pages on the secondary Mass Storage Control table, alternate 1, and alternate 2 staging drives can be used for staging on all configurations with minimal interference. The secondary Mass Storage Control tables are accessed only to write information.

In a multi-CPU configuration, the Inventory and Journal data sets, and the VSAM user catalog in which they are referenced, must permanently reside on real or virtual shared drives. To prevent losing MSVC information, place the MSV Inventory and MSVC Journal data sets on different volumes. For best performance, do not place them on volumes that also contain the operating system, page data set, or user catalogs.

It is better to use a real drive under control of a 3830 Model 2 or 3 for shared data sets because OS/VS dynamic device reconfiguration (DDR) can select another 3830, 3333 or 3330, if a failure occurs. In a multi-CPU environment, the data sets should be on real drives, especially if the data sets are to be shared through a Storage Control or a Staging Adapter pair, because the reserve and release logic interaction is less for real volumes than for virtual volumes.

System and paging volumes must be on real drives. They can be under the control of a Staging Adapter, depending on the system and staging activity. Customers that do a lot of compiling, whose systems do a lot of paging, and who anticipate a lot of staging should consider using a 3830 Model 2 Storage Control for controlling the system and paging volumes. High Storage Control utilization indicates that you need separate Storage Controls. If the above activities are moderate, or their peaks do not coincide, consider placing these volumes under a Staging Adapter.

If installation analysis indicates that a single 3830 Model 2 Storage Control must be retained, consider using a real or convertible-to-real drive on a Staging Adapter for backup purposes.

Multi-Mass Storage Facility Connections

This section examines Connections II, III, and IV for Mass Storage System configurations that have multiple Mass Storage Facilities. Figures 25-28, which show Connections III and IV have a table to relate the Staging Adapters or Staging Adapter pairs to the staging drive groups. The 3830 and 3330 representations in these four figures are different from those used previously to simplify the illustration of staging drive groups with and without string-switching. The connections in these multi-Mass Storage Facility configurations are more complex because every staging drive group pair must have a stage and/or destage path from both Mass Storage Facilities (MSF0 and MSF1).

Connection II, Mass Storage Facility to Staging Adapter

There are two Connection II paths for Mass Storage System configurations that have more than eight Mass Storage System components: one from port 0 and one from port 1. Note also that both Mass Storage Facilities are connected to all 16 components so that the Mass Storage Control in either Mass Storage Facility is able to manage the Mass Storage System. Therefore, both Mass Storage Controls must have identical features.

Connection II (Figure 22) from port 0 of MSF0 and MSF1 is the data and communication path to Staging Adapters SA=0 and SA=1. This portion of the path controls the Mass Storage Control tables in staging drive group 00 and its length cannot exceed 150 feet (46 meters). However, the remainder of the path can extend another 150 feet (46 meters) to Staging Adapters SA=2 through SA=5. Connection II from port 1 of MSF0 and MSF1 is a communication path to Staging Adapters SA=8 through SA=F. Figure 22 shows the maximum Mass Storage System configuration: two Mass Storage Facilities (MSF0 and MSF1) and 14 Staging Adapters (SA=0–SA=5 and SA=8–SA=F), a total of 16 components.

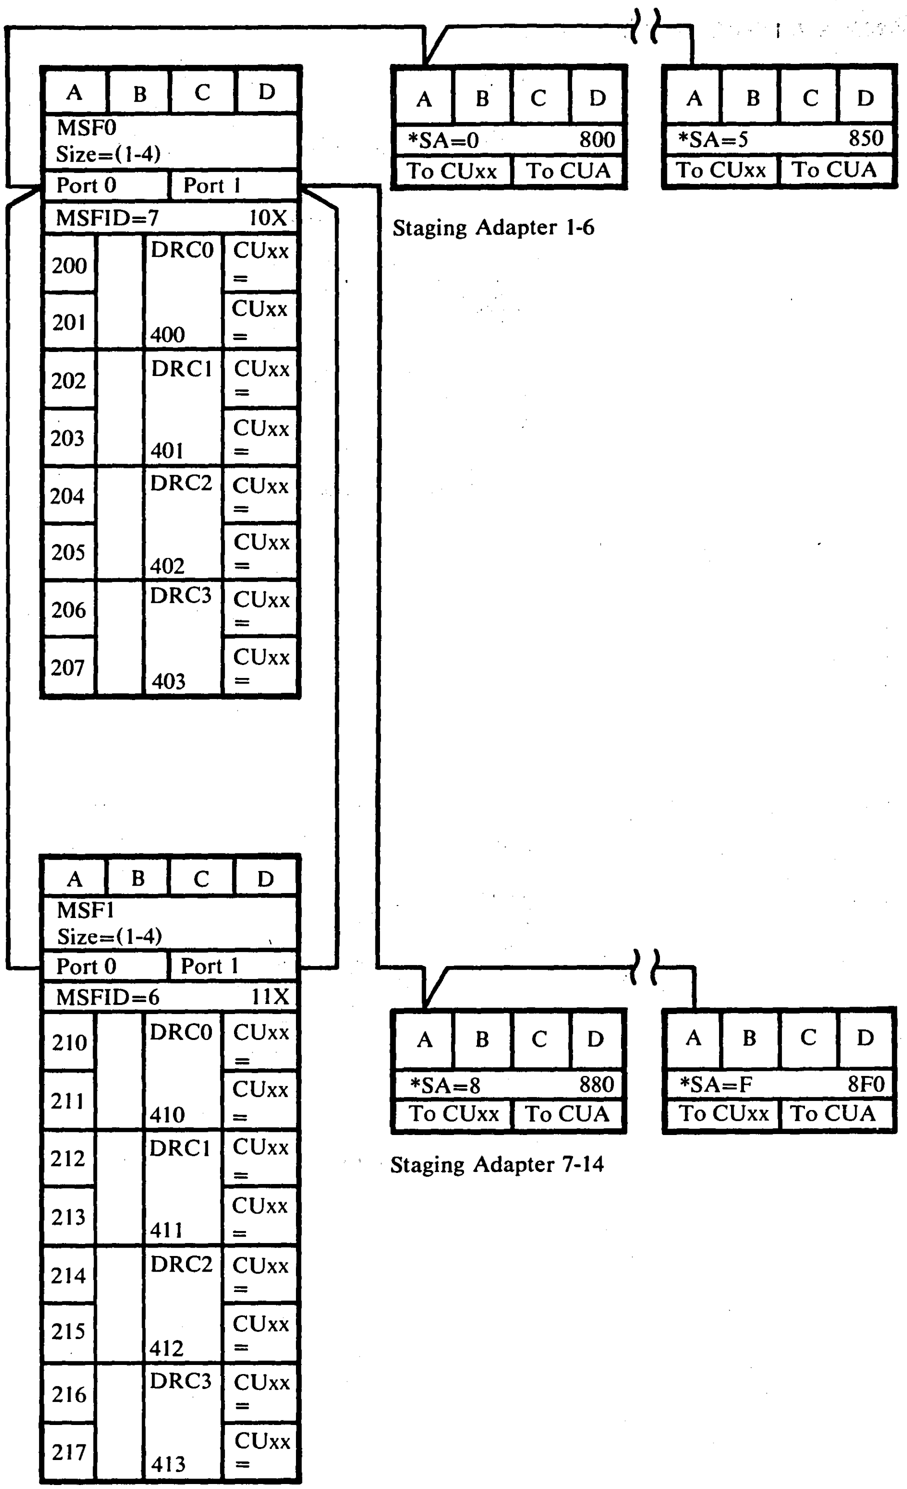

*Staging Adapter in 3830-3 or ISCSA.

Note: As shown, MSF0 is specified secondary (9121); MSF1 is specified primary (9120).

Figure 22. Connecting Two Mass Storage Facilities to Staging Adapters

Basic Connections III and IV

Figure 23 shows the basic connections from MSF0 and MSF1 to staging drive group pair 00-01. Without string-switching, the staging drive group has only one stage and/or destage path to data in MSF0 and MSF1 through the Staging Adapter.

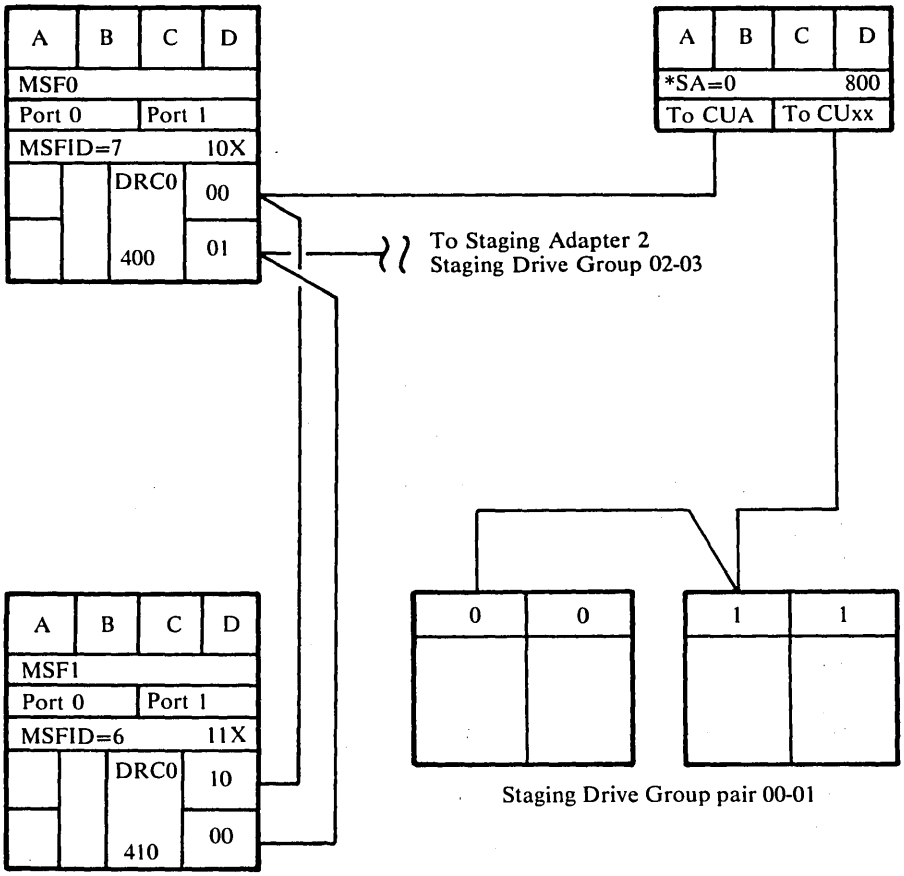

*Staging Adapter 3830-3 or ISCSA.

Figure 23. Connecting Multiple Mass Storage Facilities to Staging Drive Group Pairs

Figure 24 shows the basic connections to MSF0 and MSF1 from one string-switched staging drive group. With string-switching, a staging drive group pair has two stage and destage paths to data in MSF0 and MSF1 through two Staging Adapters.

When the System Control Program selects a virtual unit address (UNIT=3330V) on which to mount a mass storage volume in a staging drive group, it does so without knowing which Mass Storage Facility contains the mass storage volume. If both Mass Storage Facilities are not connected to all staging drive groups, the System Control Program can select a virtual unit address in a staging drive group that does not have a stage or destage path to the Mass Storage Facility containing the mass storage volume.

To illustrate, refer to Figure 24 and assume that the DRC0 upper connector (address 01) in MSF1. is disabled, thus removing the stage and/or destage path from MSF1 to the Staging Adapter SA=1 and staging drive group pair 00-01. Mass storage volume MSV001 resides in MSF1 and the System Control Program selects a virtual unit address on which to mount the mount the volume in either staging drive group 00 or 01. In this case, the Mass Storage Control cannot mount the mass storage volume because the stage or destage path from MSF1 to staging drive group pair 00-01 is disabled. Because the System Control Program cannot recognize this condition and cannot select a virtual unit address in another staging drive group, the job is aborted. This problem can be circumvented for jobs that have not been initiated by varying offline all virtual unit addresses in staging drive groups 00 and 01.

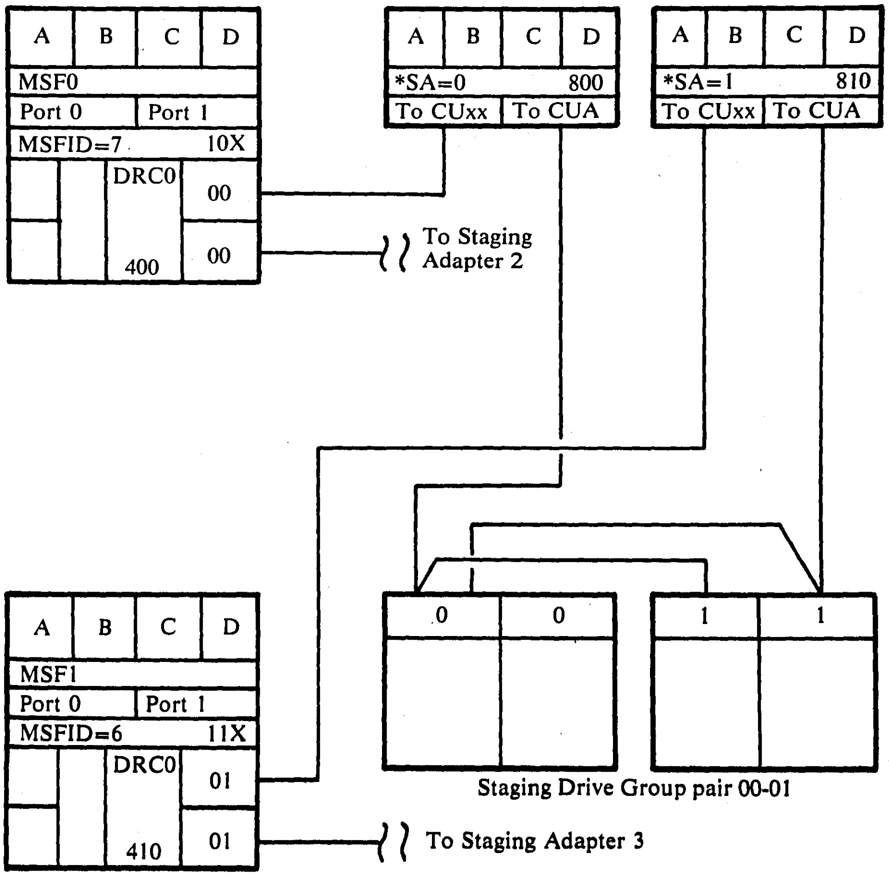

*Staging Adapter in 3830-3 or ISCSA.

Figure 24. Connecting Multiple Mass Storage Facilities to a Staging Drive Group with String-Switching

Connecting Equal Size Mass Storage Facilities

The following description extends the basic multi-Mass Storage Facility configurations to illustrate Mass Storage System configurations with Mass Storage Facilities of equal size.

The Mass Storage Control Table Create program permits identifying up to 28 staging drive groups (SDG00-SDG27). Sixteen staging drive groups with a maximum of 128 logical 3330-1s can be connected and meet all configuration requirements (see Figure 26). Staging drive group 00 must be used and must contain the Mass Storage Control tables, but the other 15 possible staging drive groups do not need to be consecutively numbered.

Figure 25 illustrates connections for two Mass Storage Facility Model A4s with eight Staging Adapters. Because staging drive group pair 00-01 is string-switched, it has two stage and destage paths, and the Mass Storage Control has two paths to its tables, through Staging Adapters SA=0 and SA=1. Staging Adapter SA=0 and SA=1 are each connected to a data recording control in MSF0 and MSF1. A data recording control or a Staging Adapter that is varied offline does not isolate staging drive group pair 00-01 from either Mass Storage Facility.

To obtain the maximum number of staging drive groups, the other six pairs (SDG02-03 through SDG12-13) each have a single stage or destage path and one data recording control connection to each Mass Storage Facility.

In Figure 25, up to eight DRC-to-Staging Adapter paths can be active concurrently.

Figure 26 illustrates connections for two Mass Storage Facilities Model A4 with the maximum 14 Staging Adapters. Twelve Staging Adapters (SA=0-SA=5 and SA=8-SA=D) are connected to support six string-switched staging drive group pairs (SDG00–SDG05 and SDG06–SDG11). Two Staging Adapters (SA=8 and SA=F) are connected to support two staging drive group pairs without string-switching (SDG12-13 and SDG14-15).

In Figure 26, up to eight DRC-to-Staging Adapter paths can be concurrently active.

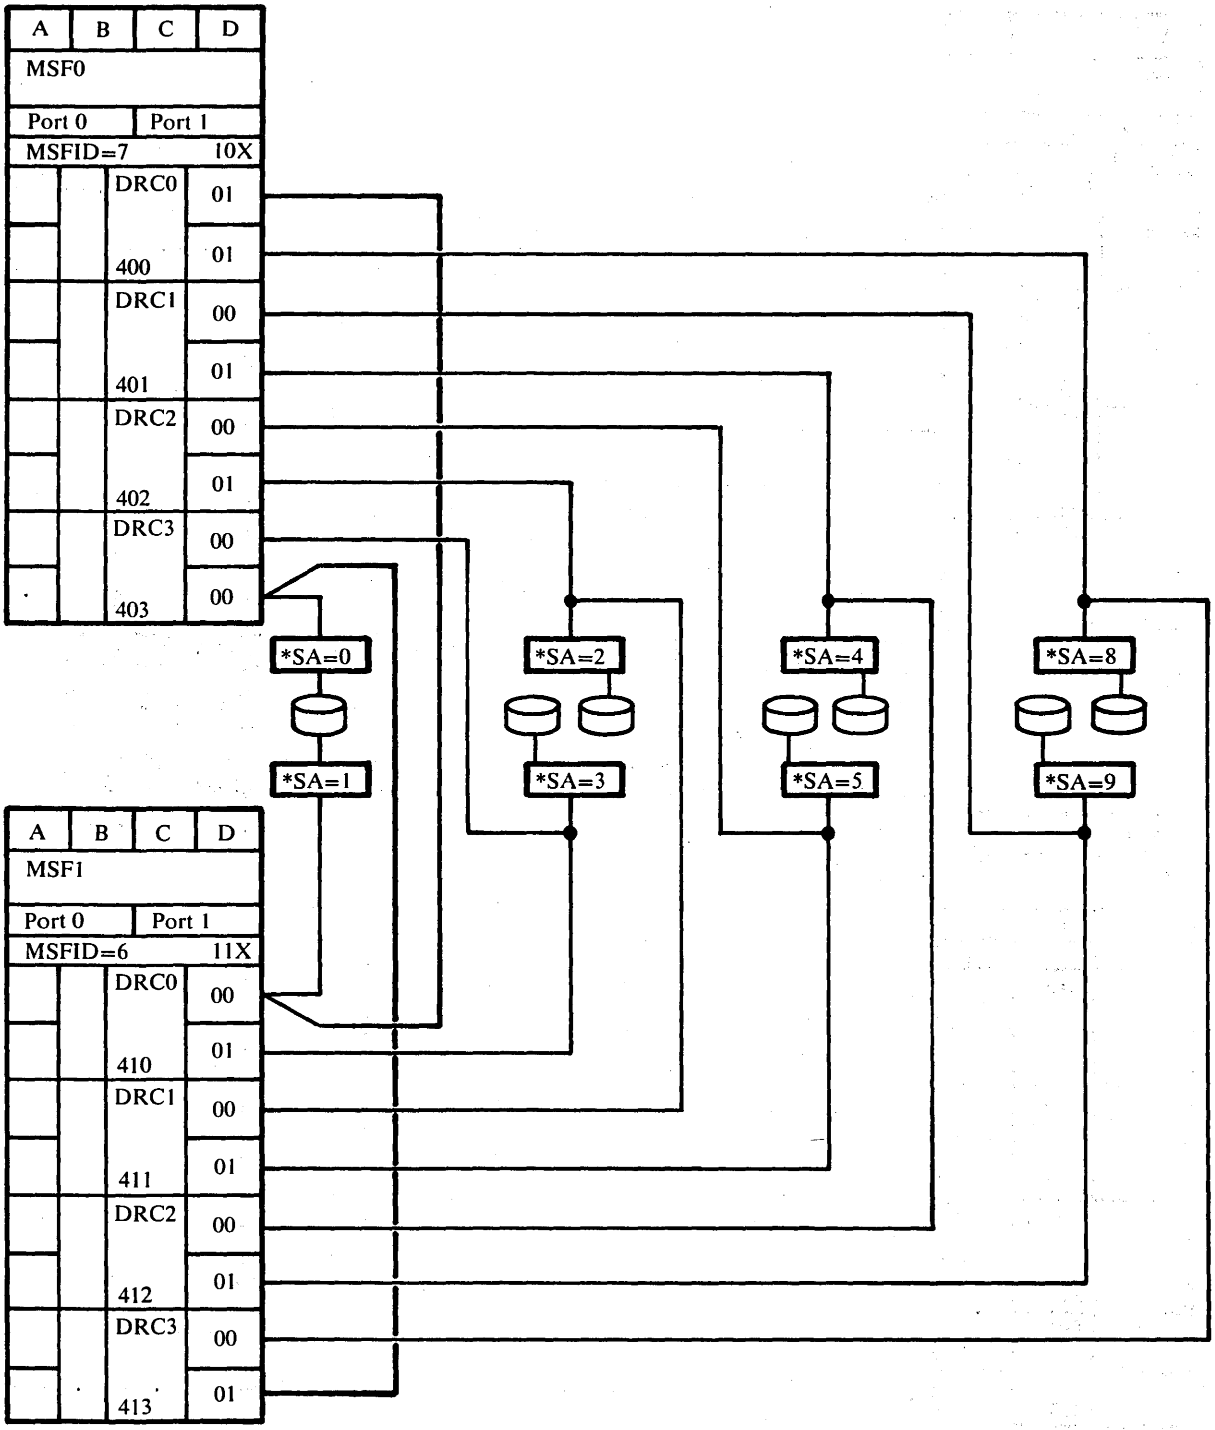

*Staging Adapter in 3830-3 or ISCSA.

Staging Drive Group Table

Staging Adapter	0-1	2	3	4	5	8	9
Staging Drive Group	00-01	02-03	04-05	06-07	08-09	10-11	12-13

Figure 25. Connecting Two 3851-A4s to Eight Staging Adapters

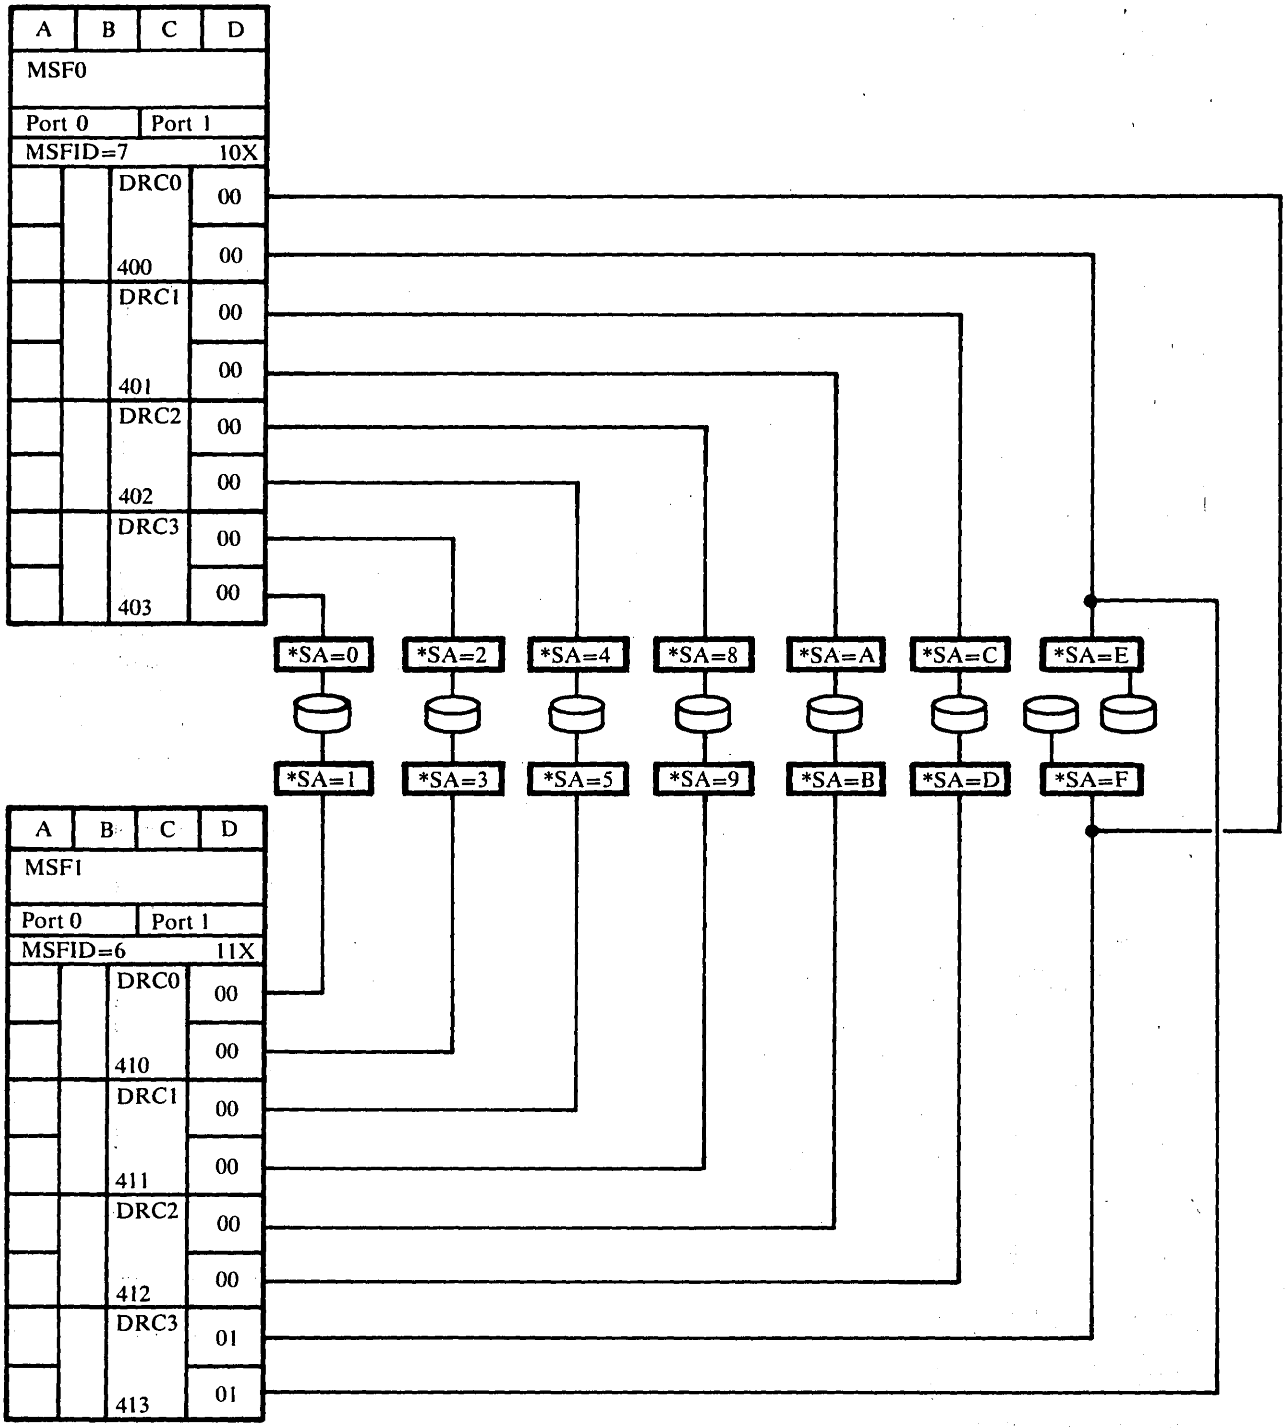

*Staging Adapter in 3830-3 or ISCSA.

Staging Drive Group Table

Staging Adapter	0-1	2-3	4-5	8-9	A-B	C-D	E	F
Staging Drive Groups	00-01	02-03	04-05	06-07	08-09	10-11	12-13	14-15

Figure 26. Connecting Two 3851-A4s to Fourteen Staging Adapters

Connecting Unequal Size Mass Storage Facilities

Figure 27 illustrates two Mass Storage Facilities of different models connected to four Staging Adapters that support three staging drive group pairs. Staging drive group pairs 02-03 and 04-05 are not string-switched. They each have one stage and/or destage path from MSF0 and MSF1.

Because staging drive group pair 00-01 is string-switched, there are two stage and/or destage paths from MSF0 and MSF1, and two paths from the Mass Storage Control to its tables through Staging Adapters 0 and 1. Note also that both stage and destage paths are connected to two DRCs one in MSF0 and one in MSF1, for maximum availability. The loss of a single DRC or Staging Adapter that has been varied offline does not isolate staging drive group pair 00-01 from either Mass Storage Facility.

A total of four DRC-to-Staging Adapter paths can be concurrently active in the Figure 27 configuration.

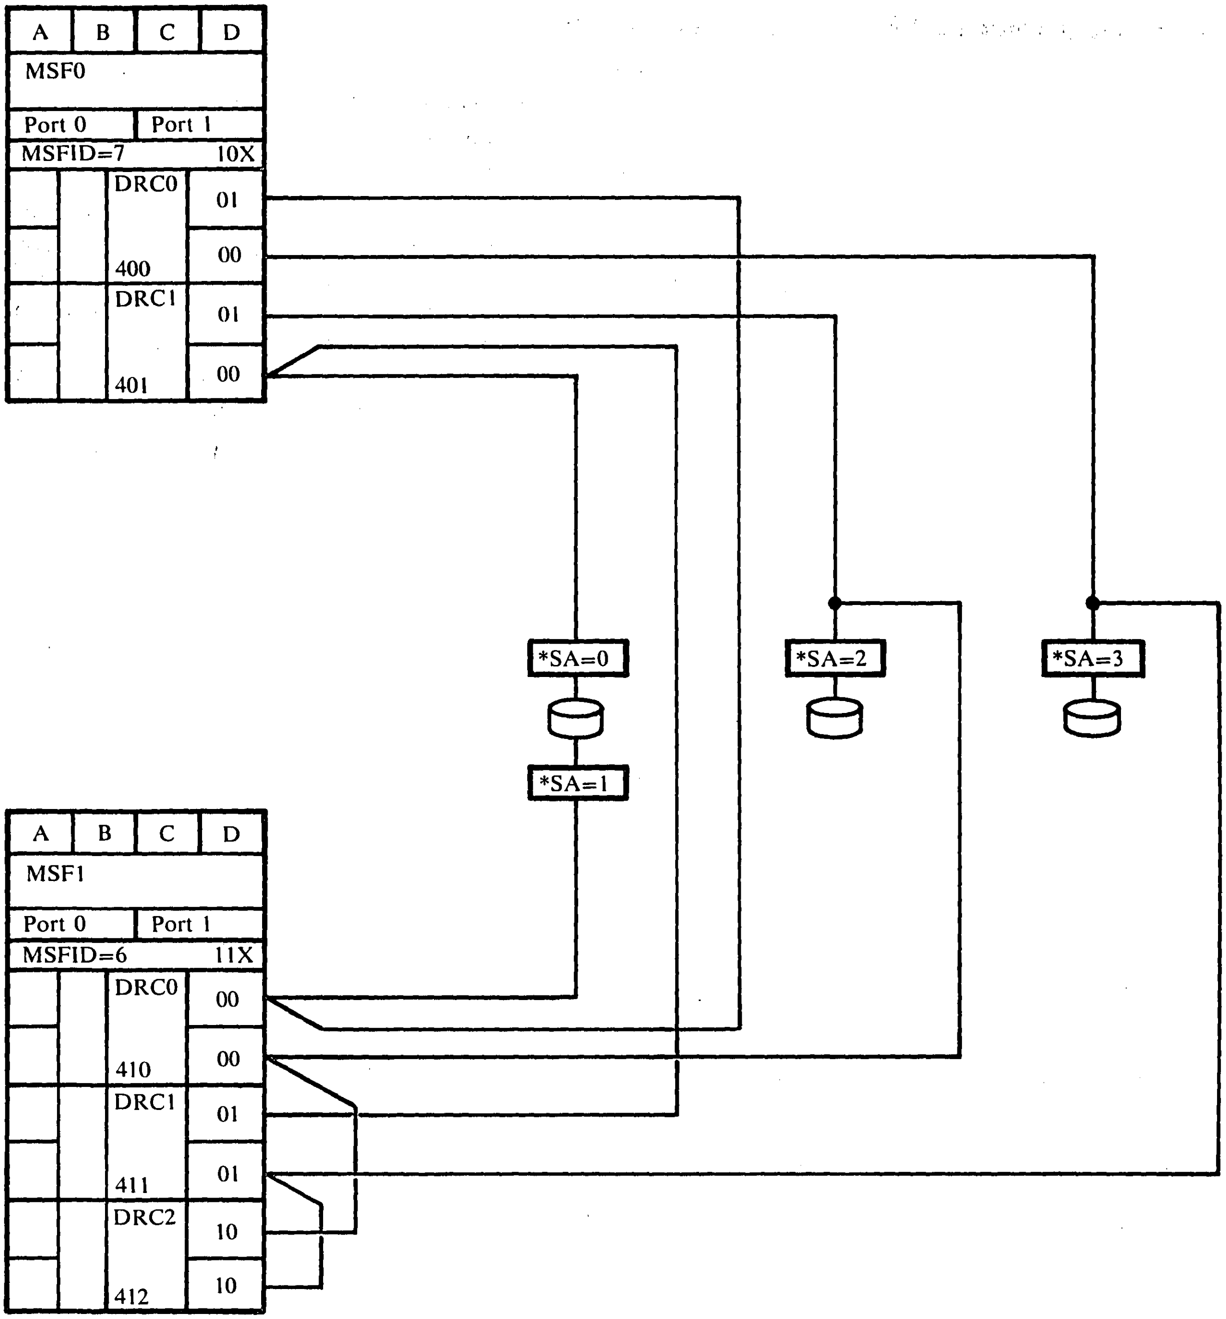

*Staging Adapter in 3830-3 or ISCSA.

Staging Drive Group Table

Staging Adapter	0-1	2	3
Staging Drive Group	00-01	02-03	04-05

Figure 27. Connecting 3851-A2 and 3851-A3 to Four Staging Adapters

Figure 28 illustrates the same two Mass Storage Facility models shown in Figure 27 but this time connected to eight Staging Adapters that support four shared staging drive group pairs with string-switching. Every staging drive group pair has two stage and/or destage paths from MSF0 and MSF1. Note that there are two DRC connections from MSF1 to staging drive group pair 00-01 (through DRC0 or DRC2) and staging drive group pair 06-07 (through DRC1 or DRC2).

A total of five DRC-to-Staging Adapter paths can be concurrently active in the Figure 28 configuration.

The illustrations and descriptions give you basic reference material for connecting, and perhaps adjusting, your configuration for maximum data capacity, performance, and sharing among several CPUs.

Attached Processor Connections

An attached processor (AP) system consists of two tightly coupled processors, a CPU and an APU (attached processor unit) with an asymmetrical connection to the Mass Storage Control. The APU uses the storage and channels of the CPU but does not attach directly to the Mass Storage System. The CPU connects to the Mass Storage Control through the use of a byte or block multiplexer channel and is specified to the Mass Storage Control Table Create program as a uniprocessor. This support is provided only by OS/VS2 MVS.

The initial release of attached processor support under OS/VS2 MVS required that it be specified as a multiprocessor to the Mass Storage Control Table Create program. This method requires two channel interfaces on the Mass Storage Control; one with a physical connection to the CPU and one with a a logical specification to the APU. These two interfaces are described to the Mass Storage Control Table Create program with two CPUCONFG commands that specify the CPU (CPUID), APU (MPCPU), and the respective Mass Storage Control channel interfaces (MSCINT). Alternate channel paths (ALTINT) are not permitted.

An AP system can be connected to the Mass Storage System asymmetrically as a uniprocessor or symmetrically as a multiprocessor. If your AP system is installed as a multiprocessor system, there is no immediate need to change unless the logical Mass Storage Control interface is needed to attach other processors. If a change is made, new Mass Storage Control tables must be created.

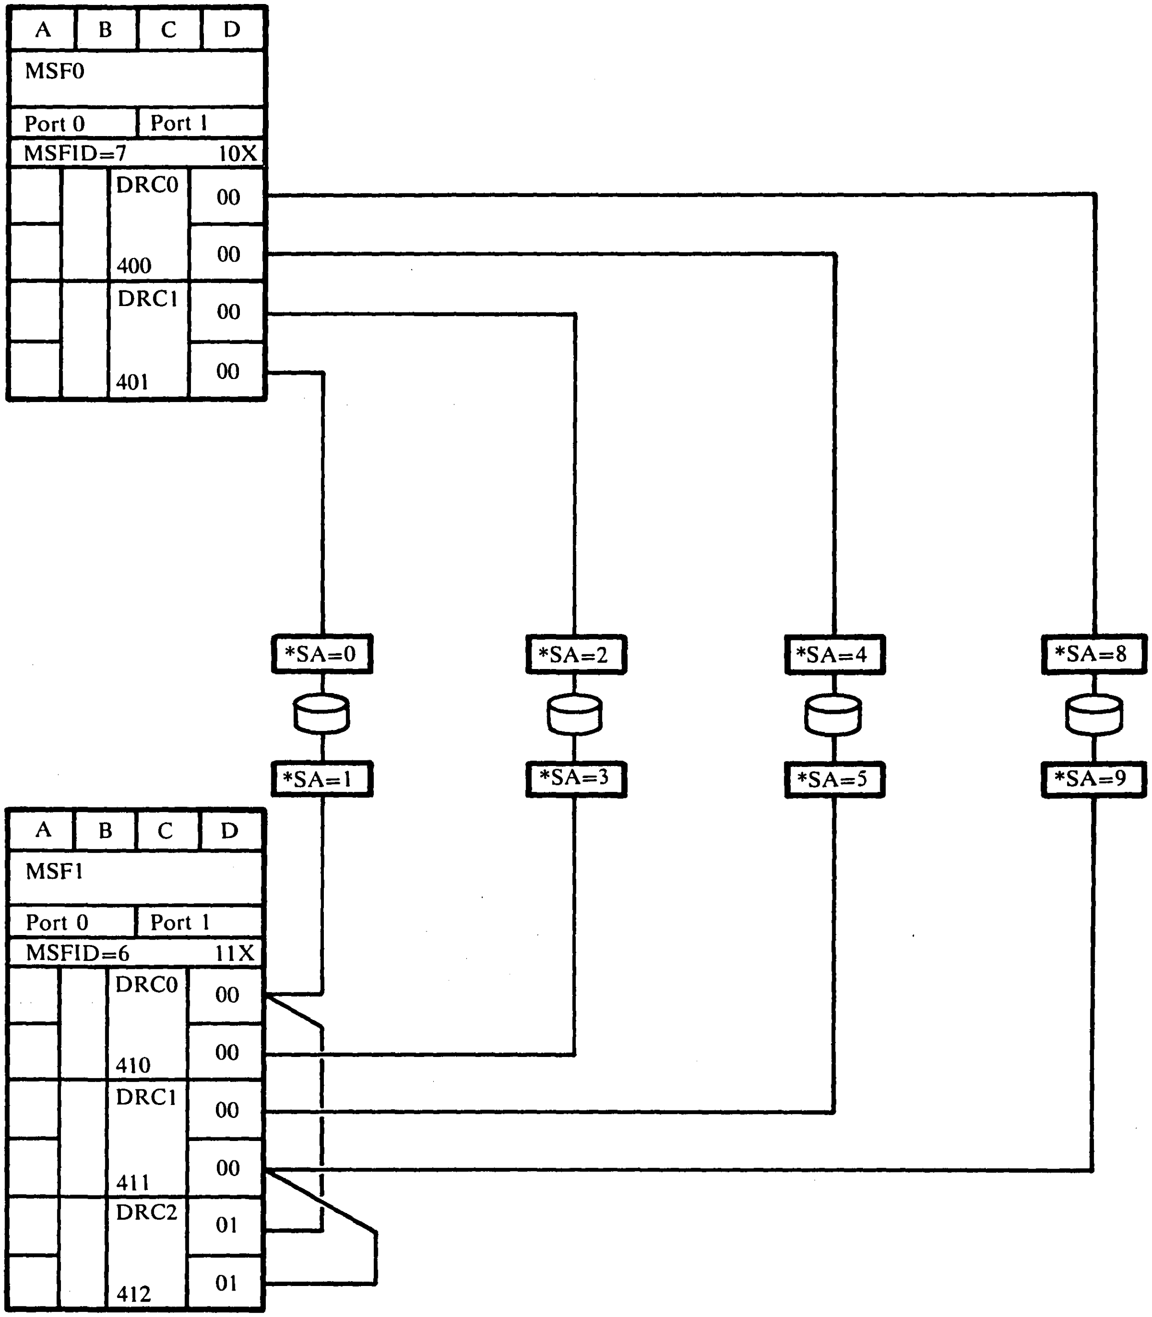

*Staging Adapter in 3830-3 or ISCSA

Staging Drive Group Table

Staging Adapter	0-1	2-3	4-5	8-9
Staging Drive Group	00-01	02-03	04-05	06-07

Figure 28. Connecting 3851-A2 and 3851-A3 to Eight Staging Adapters

Multiprocessor Connections

There are additional unit address, path, and IODEVICE parameter considerations to review when you plan multiprocessor (MP) system connections to the Mass Storage System. A multiprocessor system consists of two tightly coupled processors, CPU0 and CPU1, that can be connected to the Mass Storage Control symmetrically or asymmetrically.

OS/VS2 MVS supports the multiprocessor pair connected to the Mass Storage Control symmetrically using two Mass Storage Control interfaces, or only one of the multiprocessor pair connected to the Mass Storage Control asymmetrically using one Mass Storage Control interface. Either method of connecting may require the IODEVICE card parameter FEATURE=(SHARED) or FEATURE=(SHAREDUP). These IODEVICE card parameters are described in the following sections and are used in the chapter "Mass Storage Control Table Create Examples." We recommend you plan multiprocessor configurations with symmetrical connections so as to provide the maximum number of paths for each processor.

Connection I, CPU to Mass Storage Facility

Figure 29 shows Connection I for CPU0 and CPU1 with the symmetric (A) and asymmetric (B) paths to the Mass Storage Control. With symmetric paths, CPU0 and CPU1 are specified to the Mass Storage Control Table Create program as a multiprocessor pair (CPUID and MPCPU). Each CPU has a primary path to a Mass Storage Control interface (MSCINT) but not an alternate path.

With asymmetric paths, CPU0 or CPU1 is specified to the Mass Storage Control Table Create program as a uniprocessor (CPUID) with one primary path to the Mass Storage Control interface. Although an alternate path to a Mass Storage Control interface (ALTINT) can be specified, the asymmetric method of connection frees a Mass Storage Control interface for connection to another CPU.

With two Mass Storage Controls, every attached CPU must connect to the same primary and alternate Mass Storage Control on the same channel as shown in Examples C and D of Figure 29. In addition, both Mass Storage Controls (one Mass Storage Facility Model B or two Model A's) must have identical features and connections so that either Mass Storage Control can communicate with all the Mass Storage System components.

Connection I is established by the Mass Storage Control Table Create CPUCONFG command, as specified on Mass Storage Control Table Create Program Worksheet 7. IODEVICE cards must be provided for the unit addresses of the Mass Storage Controls.

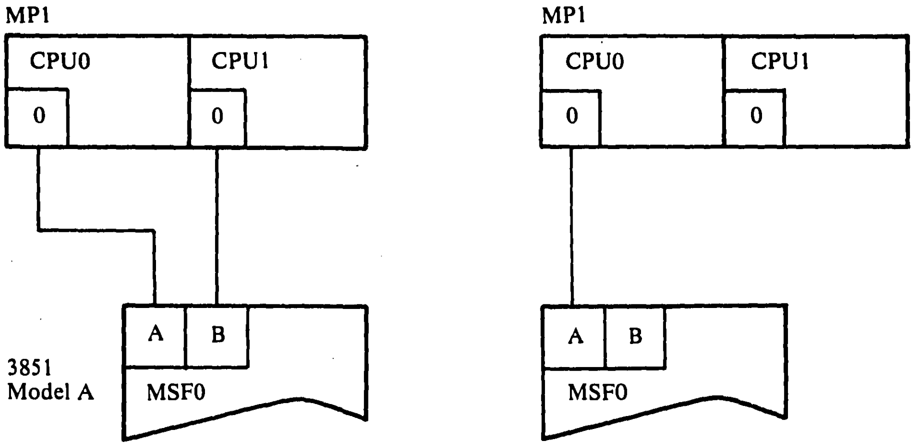

Example A. Symmetrical MP Connection I

Example B. Asymetrical MP Connection I

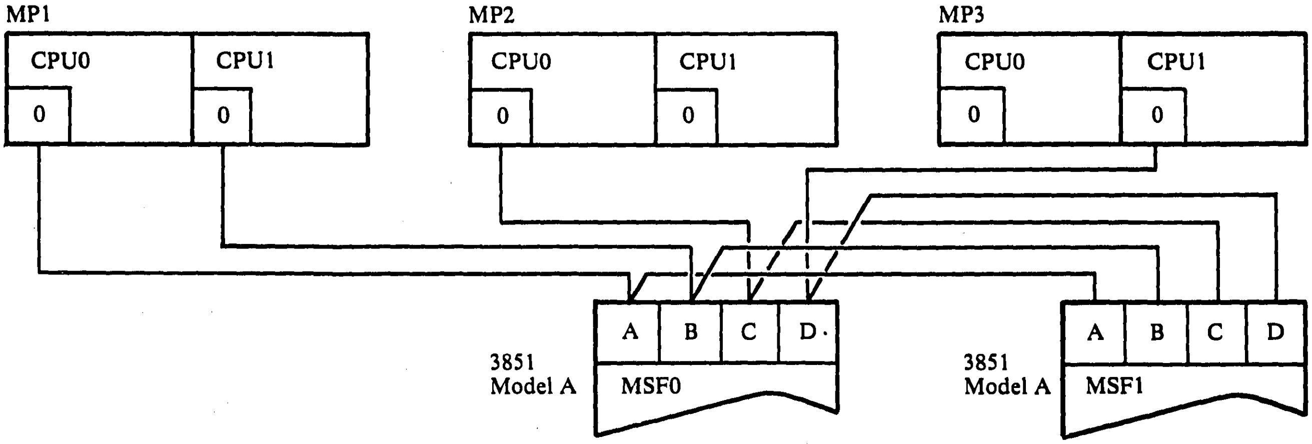

Example C. Every CPU must connect to the same MSC interfaces on the same channel

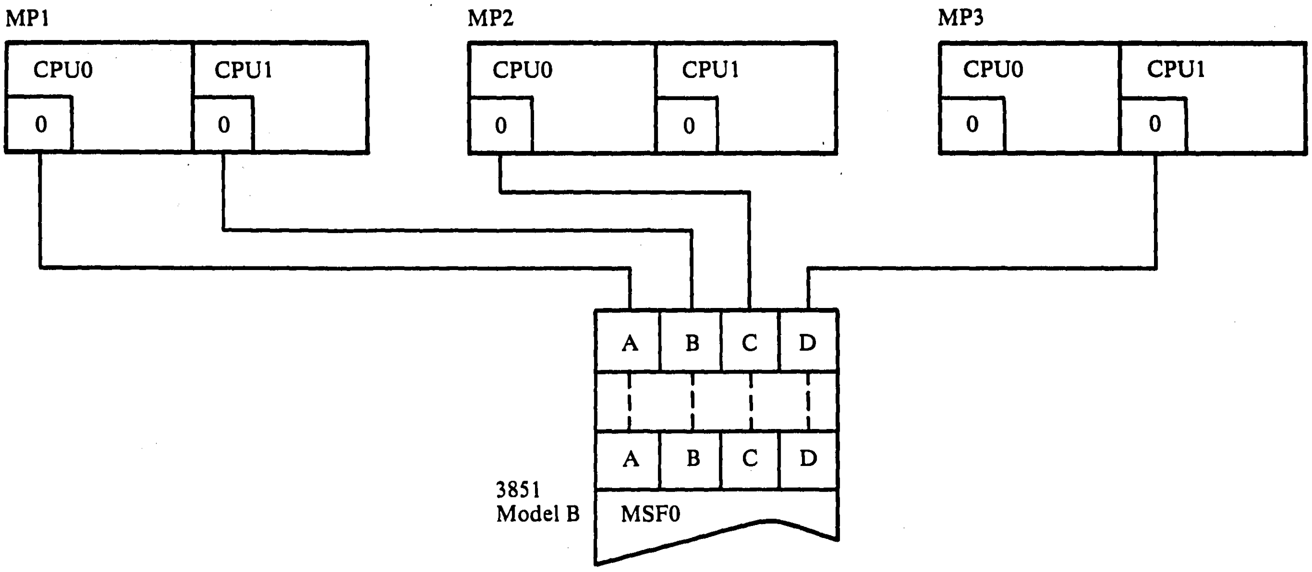

Example D. At least one CPU must connect to the MSC.

Figure 29. Connection I for Multiprocessors

Connection V, CPU to Staging Adapter

Example C in Figure 30 shows the recommended method of connecting the multiprocessor pair CPU0 and CPU1 to the Staging Adapters. With this method, each CPU has an alternate channel path through an alternate Staging Adapter for either multiprocessor or uniprocessor operation.

The OS/VS2 MVS logic for the alternate path and control unit busy condition is essentially the same as described in the section "Connection V, CPU to Staging Adapters (for Uniprocessors)." In Example C of Figure 30, if a channel busy condition is found on primary channel 1 of CPU0, OS/VS2 MVS attempts to use alternate channel 2 of CPU0. If both paths of CPU0 are busy, OS/VS2 MVS switches to CPU1 and tries primary channel 1 and alternate channel 2 in sequence. If a control unit busy condition is encountered, the path through an alternate control unit is tried only if FEATURE=(ALTCTRL) has been specified for the unit address in use.

In normal multiprocessor mode, all system resources are controlled by a single OS/VS2 MVS System Control Program. Virtual and real units accessed by both CPU0 and CPU1 must have identical unit addresses. In uniprocessor mode, both CPU0 and CPU1 must have a System Control Program to control their own resources. If CPU0 and CPU1 share virtual units that are not shared with other CPUs connected to the Mass Storage System, the IODEVICE cards created by the Mass Storage Control Table Create program should be changed to FEATURE=(SHAREDUP). This parameter permits the device reserve and release logic to be bypassed in multiprocessor mode and used in uniprocessor mode.

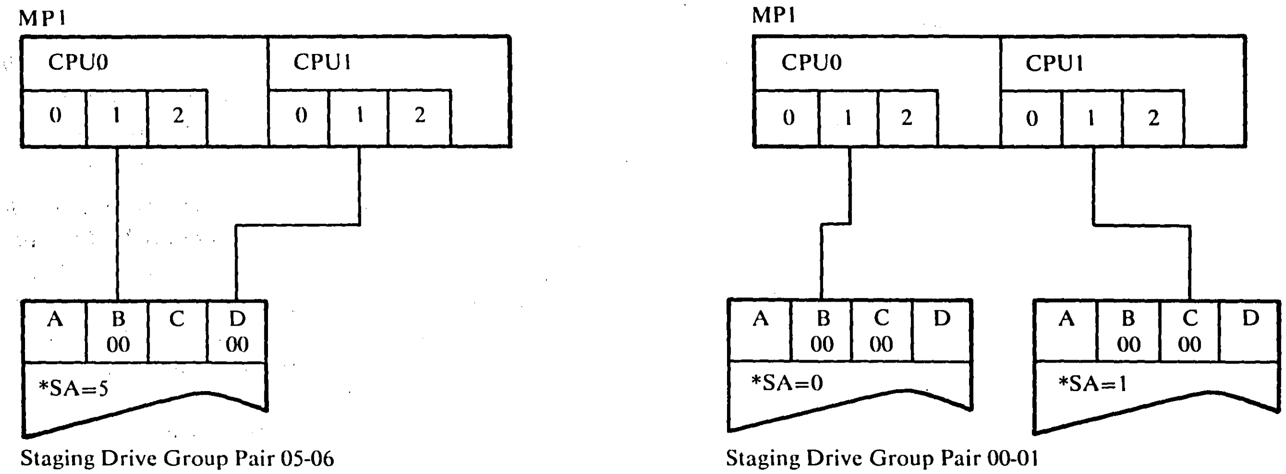

Example A. MP without string-switching

Example B. MP with string-switching

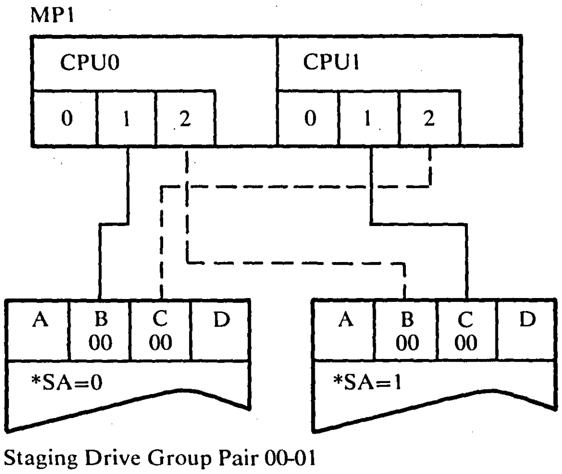

Example C. MP with string-switching and Alternate Paths

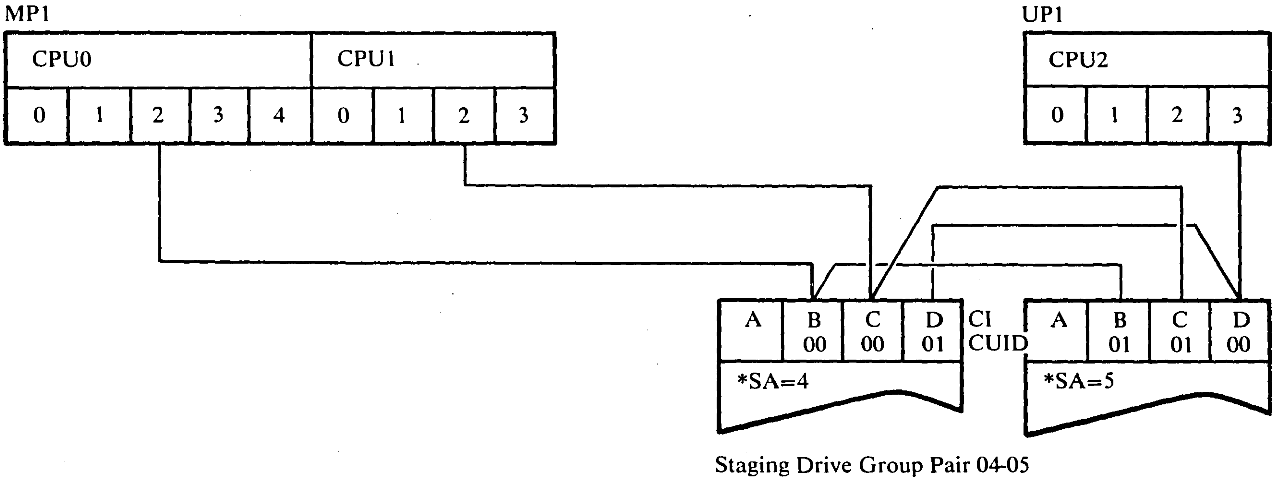

Example D. MP and UP Multihost

*Staging Adapter in 3830 or ISC.

Figure 30. Connection V for Multiprocessors

If you are planning to run a multiprocessor system in uniprocessor mode, specify virtual shared unit addresses. Change FEATURE=(SHARED) to FEATURE=(SHAREDUP) only if there are no other CPUs connected to the Mass Storage System using the same virtual shared unit addresses.

If unit addresses 120-124 in staging drive group 00 are specified as virtual, rather than virtual shared, both CPUs can mount and access virtual volume A on address 120. If the CPUs are disassociated (uniprocessor mode), only one CPU can mount and access virtual volume A on address 120 because the virtual unit address is not shared.

If unit addresses 120-124 are specified as virtual shared, either CPU0 or CPU1 can mount and access virtual volume A on address 120 in multiprocessor mode. If the CPUs are disassociated, both CPUs can mount volume A because they both have virtual shared addresses. However, CPU0 can mount volume A on its virtual shared unit address 120 and CPU1 can mount volume A on its virtual shared unit address 124 (or 120, 121, 122, and 123). The only requirement for sharing volume A is that both CPUs have a virtual shared unit address available in staging drive group 00. (This description assumes that volume A is not defined with the EXCLUSIVE volume attribute.)

Figure 30 shows several Multiprocessor-to-Staging Adapter connections, both symmetrical and asymmetrical. All connections shown can be used but be careful when you define addresses for CPU0 and CPU1.

Example D of Figure 30 shows connections that provide the CPUs with up to 128 unit addresses. The first 64 addresses (00-3F) are selected through Staging Adapter channel interface CUID 00, and the second 64 addresses (40-7F) through Staging Adapter channel interface CUID 01. In this example, all CPU paths to staging drive group pairs 04 and 05 are primary paths.

Connection V is established by the Mass Storage Control Table Create UPPERCON command, as specified on the Mass Storage Control Table Create Program Worksheet 5.

Summary

The logical connections described apply to a single Mass Storage System having one or two Mass Storage Facilities, and connected to as many as four CPUs. Unless otherwise specified, Staging Adapter refers to a 3830 Model 3 or a logical half of an ISC with Staging Adapter.

Minimum Configuration

The minimum 3850 configuration is:

- One 3851 Mass Storage Facility Model A1
- One 3830 Model 3 Storage Control or Model 158 or Model 168 ISC with Staging Adapter
- Two 3333 Disk Storage and Controls, Model 1 or 11 (four drives)

  Note: The four drives in the minimum configuration do not include space for system-residence, paging, Mass Storage Volume Inventory, or other system functions.

Connection I

- Connection I is from the host to the Mass Storage Facility. A maximum of four hosts can be connected to the primary Mass Storage Control (Mass Storage Facility Model A) and to the alternate Mass Storage Control (Mass Storage Facility Model B or two Mass Storage Facility Model A's) from either the byte or block multiplexer channel. The four hosts can be uniprocessors (CPU), attached processors (CPU and APU), or multiprocessor (CPU0 and CPU1) systems.

  An attached processor can be logically specified to a Mass Storage Control as a multiprocessor with symmetric connections (physical from the CPU and logical from the APU) or as a uniprocessor with asymmetric connections.

  A multiprocessor can be logically specified as a multiprocessor with symmetric connections (from CPU0 and CPU1) to the Mass Storage Control, or as a uniprocessor with asymmetric connections (from CPU0 or CPU1) to the Mass Storage Control.

- In a configuration with two Mass Storage Facilities, the same primary and alternate Mass Storage Control channel interfaces must be connected to a CPU on the same primary channel (MSCINT) or on the same alternate channel (ALTINT).
- Two Mass Storage Controls (3851 Model B or two 3851 Model A's) require two control positions on the byte or block multiplexer channel to which they are connected.
- A CPU can have one primary and one alternate channel path to the Mass Storage Control on a block multiplexer channel, or only one primary channel path on a byte multiplexer channel. A multiprocessor (CPU0 and CPU1) with symmetric paths to the Mass Storage Control cannot have an alternate path to the Mass Storage Control.
- Every host connected to a Staging Adapter must connect to the Mass Storage Control.

Connection II

- Connection II is from the Mass Storage Facility to the Staging Adapter. The Mass Storage Control can address Mass Storage System components logically as 0-7 (seven Staging Adapters and one accessor control pair, or six Staging Adapters and two accessor control pairs) on port 0, and 8-F (Staging Adapters only) on port 1 (optional feature).
- The Mass Storage Control tables are accessed through Staging Adapters 0 and 1 (optional path with 3333 string-switching). These two Staging Adapter logical addresses are reserved for this use.
- The Mass Storage Control table drives must be defined as staging drives in staging drive group 00 on SSIDs 000, 002, 004, and 006 (logical address plugs 00, 01, 08, and 09).
- If Staging Adapter 1 is used in a Mass Storage System configuration, it must have a path to the Mass Storage Control tables and it must be paired with Staging Adapter 0.

- In a configuration with two Mass Storage Facilities, port 0 on MSF0 must be serially connected to port 0 on MSF1, and port 1 on MSF0 must be serially connected to port 1 on MSF1. This permits either Mass Storage Control to address all 3850 components on either port.

  Feature code 9120 determines the primary Mass Storage Facility with the cable-in connections and feature code 9121 determines the secondary Mass Storage Facility with the cable-in and cable-out connections.

- The primary Mass Storage Control and the alternate Mass Storage Control must have identical features and connections (two-channel switch additional or twin port) so that all 3850 components can continue to operate if the SCP switches Mass Storage Controls.

Connection III

- Connection III is from a Staging Adapter to a 3333 or 3350. A Staging Adapter can connect up to a maximum of four 3333s or 3350s.
- A Staging Adapter pair can connect to the same 3333 if the 3333 has the string-switch feature. If the 3333 that controls staging or convertible-to-real drives is string-switched to a Staging Adapter pair, then all 3333s that control staging or convertible-to-real drives connected to either Staging Adapter must be string-switched to the other Staging Adapter in the pair.
- A Staging Adapter pair can consist of two 3830-3s, one 3830-3 and a logical half of an ISC with Staging Adapter, or two ISCs with Staging Adapter. IBM does not support a 3830-2 Storage Control as half of a Staging Adapter pair that controls staging drives.

Connection IV

- Connection IV is from the Staging Adapter to the Mass Storage Facility. A Staging Adapter can connect to a maximum of four DRCs in one or two Mass Storage Facilities.
- If string-switching is used, each Staging Adapter in a pair can connect to a different DRC for maximum availability.
- With two Mass Storage Facilities, the Staging Adapter connections can be planned so that both Mass Storage Facilities have a stage and/or destage path to every staging drive group through one Staging Adapter or a Staging Adapter pair (using string-switched 3333s).

Connection V

- Connection V is from the CPU to the Staging Adapter. Every host connected to a Staging Adapter must also be connected to the Mass Storage Control.
- An attached host must have a logical path to every staging drive group if shared access to all volumes is required.
- OS/VS1 supports one primary and three alternate channel paths to real or virtual devices.

- OS/VS2 MVS supports one primary and one alternate channel path to real or virtual devices.
- The alternate channel path address must be a higher address than that of the primary channel path.
- All Staging Adapter interfaces (B, C, and D) respond to a maximum of 64 unique addresses (00-3F) which can all be specified as virtual. Only the first 32 unique addresses (00-1F) can be specified as REAL or CONVREAL.
- A maximum of four Staging Adapter interfaces on a channel respond to the maximum 256 addresses; that is, Staging Adapter 00 (CUID) responds to 00-3F, 01 responds to 40-7F, 10 responds to 80-BF, and 11 responds to C0-FF.
- Other tape and disk control units on the same channel with a Staging Adapter must have addresses beyond the range of the Staging Adapter and preferably within a REAL address range; that is, 00-1F, 40-5F, 80-9F, or C0-DF.
- Every unit address, or UCB, specified in the OS/VS SYSGEN for a virtual device should have a corresponding virtual or virtual shared address specified to the Mass Storage Control Table Create program.
- A nonshared subchannel, or UCW, must be available for every unit address, or UCB, specified to the OS/VS SYSGEN and Mass Storage Control Table Create program.
- One CPU (with one or more block multiplexer channels) cannot connect to more than two interfaces on one Staging Adapter. If more than two interfaces of a Staging Adapter are connected, then more than one system reset may be required to reset them.
- The interface addresses (CUID) on a Staging Adapter are specified as 00, 01, 10, or 11, depending on their use. The Customer Engineer might refer to these addresses as 0, 4, 8, or C.
- If two Staging Adapter interfaces are connected to primary and alternate channel paths from one CPU, their addresses (CUID) must be the same.
- If two Staging Adapter interfaces are connected to only primary channel paths from two CPUs, their addresses (CUID) can be either the same or different.
- If four Staging Adapter interfaces (two or more Staging Adapters) are connected to one CPU channel, their addresses (CUID) must be different.
- A multiprocessor can have an alternate channel path to a Staging Adapter but it must originate from the same multiprocessor CPU as the primary channel path.
- If an alternate channel path is specified from both multiprocessor CPUs, the alternate channel path address must be the same for both multiprocessor CPUs.
- If both multiprocessor CPUs access the same real or virtual unit, the unit addresses (bits 16-31) must be identical.

- If you plan uniprocessor operation for a multiprocessor, virtual addresses accessible to both CPU0 and CPU1 must be specified as virtual shared. If a virtual unit address is not shared with another CPU in a multihost environment, its IODEVICE card can be changed from FEATURE=(SHARED) to FEATURE=(SHAREDUP).
- The Staging Adapter configuration can be asymmetrical for a multiprocessor.

Related Logical and Physical Planning

System Control Program Generation

An operating system that meets immediate and short-term needs is the goal of most installations. Because the number of OS/VS unit addresses needed to identify all of the installed, transitional, and final real and virtual devices is large, addresses must be specified with care. Also, addresses for 3850 real and virtual devices must be clearly understood before you plan the 3850 SYSGEN and resolve any address conflicts between installed real 3330 or 3350 drives and planned virtual 3330 drives on the same channel.

Until the 3850 is installed, the modified Storage Controls function in a nonstaging mode that addresses only real 3330 drives, but uses the 3830-3 unit address format. Installed real 3330 drives, planned for future staging use, can be defined as convertible-to-real (CONVREAL) drives if sufficient unit addresses are available. After the 3850 is installed, these convertible-to-real 3330 drives can be varied online as real 3330 drives. Gradually, as testing and data migration progress, the convertible-to-real 3330 drives can be varied online as staging drives that are known only to the Mass Storage Control.

Support for the 3850 can be added only by a new SYSGEN with an IODEVICE card specifying UNIT=3851. A control program must be generated (3850 SYSGEN) to include the 3851 device and the associated 3850 program modules. The Mass Storage Control Table Create program can be run prior to the SYSGEN or at a later time with an IOGEN to include the 3850 real and virtual unit addresses. Consider using one of these options when you plan your 3850 configuration:

- An inital Mass Storage Control Table Create program run and 3850 SYSGEN containing all known installed, transitional, and final device addresses.
- An initial 3850 SYSGEN containing the known address changes and a Mass Storage Control Table Create program run and IOGEN immediately before the 3850 installation.
- An inital SYSGEN without the 3850, containing the known address changes, and a Mass Storage Control Table Create program run and 3850 SYSGEN immediately before the 3850 installation.

Review the IODEVICE cards that the Mass Storage Control Table Create program created as input to the Stage 1 of SYSGEN. Also, prepare an IODEVICE card for the Mass Storage Control for each CPU.

The Mass Storage Control Table Create program does not generate information on IODEVICE cards for ALTCTRL, OFFLINE, and SHAREDUP. See the *OS/VS1* or *OS/VS2 System Generation Reference* for additional information.

Unit Control Blocks

You can SYSGEN more unit control blocks (UCB) than are immediately required by a CPU to avoid an IOGEN as new and additional devices are installed. In multi-CPU installations, SCP maintenance can be reduced by generating one SCP that has a set of I/O addresses for all CPUs. Those I/O addresses not required by a specific CPU should be:

- Generated offline in the SYSGEN
- Or varied offline with the HMASPZAP program
- Or varied offline by the operator

You may decide to SYSGEN more virtual UCBs for virtual devices than are specified to the Mass Storage Control Table Create program. In doing so, problems might occur when virtual UCBs are generated within or beyond the unit address range of an installed Storage Control, and which are not specified to the Mass Storage Control Table Create program.

If you are using OS/VS2 MVS, the additional unit addresses can be generated offline by specifying OFFLINE=YES in the IODEVICE cards. If you are using VS1, there is no offline option but the HMASPZAP program can be used to vary the unit addresses offline in the disk-resident program code.

The following chart shows the status of real and virtual unit addresses after a TEST I/O command has been issued during IPL.

				Unit Status		
SCP	**Generated Unit Addresses**		**NIP***	**Volume Mounted**	**Intervention Required**	**Unit Not Available**
VS2	Real	Offline	Std.	Rdy - Offline	Rdy - Offline	Rdy - Offline
		Online	Std.	Read VOLSER (2)	Nrdy - Online	Rdy - Offline
		Online	Opt. (1)	Read VOLSER (2)	Nrdy - Offline	Rdy - Offline
	Virtual	Offline	Std.	Rdy - Offline (2)	Rdy - Offline	Rdy - Offline
		Online	Std.	Read VOLSER (2)	Nrdy - Online	Nrdy - Offline
		Online	Opt. (1)	Read VOLSER (2)	Nrdy - Offline	Rdy - Offline
VS1	Real		Std.	Read VOLSER (2)	Nrdy - Online	Nrdy - Online
			Opt. (1)	Read VOLSER (2)	Nrdy - Offline	Nrdy - Offline
	Virtual		Std.	Read VOLSER (2, 3)	Nrdy - Online	Nrdy - Online
			Opt. (1)	Read VOLSER (2, 3)	Nrdy - Online	Nrdy - Offline

* NIP=Nucleus Initialization Program

Notes:

1. DEVSTAT option can be specified in OS/VS2 MVS SYSGEN.

   Equivalent of DEVSTAT (SMARTNIP) can be specified in OS/VS1 SYSGEN or the operator can reply DEVSTAT=n to message IEA101A.

2. If unit is shared and busy and the operator reply to message IEA20I is continue, the unit address is set Nrdy - offline.
3. VOLSER is not read by NIP but is read after MSS initialization.

When NIP is used with the DEVSTAT option, virtual unit addresses generated online are varied offline if intervention required status is returned. Your initialization procedure should contain in-stream Vary On commands for the virtual unit addresses that you want online.

A potential problem exists with additional unit addresses when the SCP has allocated all of the online devices and calls for allocation recovery. To satisfy allocation, the operator receives a message containing a list of offline units to be varied online. If the operator mistakenly varies online a UCB for a virtual device that is not contained in the Mass Storage Control tables, the job aborts because the device is unknown to the Mass Storage Control.

This problem can be avoided by specifying virtual devices to the Mass Storage Control Table Create program for all virtual UCBs generated. If too many virtual UCBs are generated, specify them as offline and document a procedure to ensure that these difficulties do not recur.

UCBs for convertible-to-real drives should also be specified as offline to avoid possible conflicts between real and staging drives. To illustrate this problem, assume that physical drive C (Staging Adapter 10 on channel 2) is specified to the Mass Storage Control Table Create program as CONVREAL (SSID=012) in staging drive group 02. The System Control Program is running with real unit address 242 varied offline and SSID 012 varied online as a staging drive.

A System Control Program failure occurs and NIP (not SMARTNIP) is used to IPL the system because the UCB for the convertible-to-real drive is specified as online. NIP does not change that status. When operation is resumed, the UCB for the convertible-to-real drive is allocated to mount a real volume.

The System Control Program issues an operator message that requests mounting a disk pack on physical drive C, which is still being used as a staging drive. The operator cannot complete the mount operation and must terminate the job or IPL the system again.

On OS/VS2 MVS, the conflict between real and staging drives can be avoided by adding OFFLINE=YES to the IODEVICE cards created by the Mass Storage Control Table Create program for convertible-to-real drives. On OS/VS1, the conflict can be avoided by using the HMASPZAP program to vary the UCBs of the convertible-to-real drives offline in the disk-resident program code.

Unit Control Words

CPU or channel features might be required to increase the number of nonshared subchannels and channel unit control words. One OS/VS unit control block (UCB) is required for each real, convertible-to-real, and virtual drive. One channel unit control word (UCW) is required for each unit address, primary and alternate, sent to that channel.

The range of nonshared subchannels is:

- A minimum of 16, (feature code 9491) to a maximum of 512 (feature code 9522) in multiples of 16 on System/370 Model 145. Each UCW requires eight bytes of control storage (or processor storage if requirements exceed 32,768 bytes).
- A minimum of 64 with options for 128 (feature code 9581), 256 (feature code 9582), or 512 (feature code 9583) on System/370 Models 145-3 or 148.
- Variable with the size of processor storage on System/370 Model 155II; that is, 96 subchannels for 262K; 160 for 393K; 224 for 524K; 252 for 786K; and 480 for 1 megabyte and above.
- 480 on a System/370 Model 158 or 158-3. These are available on up to five block multiplexer channels (two standard and up to three optional). If the second byte multiplexer channel is not installed on a Model 158-3, a maximum of 736 nonshared subchannels is available.
- 56 on the 2880 Block Multiplexer Channel Models 1 and 2 System/370 Models 165II, 168, 168AP, and 168MP, with up to 56 nonshared subchannels on each block multiplexer channel. See the following section for details.
- 256 per block multiplexer channel on the 3033 Processor.

The block multiplexer channels on System/370 Models 145, 155II, and 158 share a pool of nonshared subchannels, assigned in groups of eight unit control words. The channel assigns a subchannel group of eight unit control words to a block of eight unit control blocks that contain addresses with a low-order bit value of 000 to 111.

Because of the group and block relationship of unit addresses to UCWs, use consecutive blocks of unit addresses so that all nonshared subchannels can be used, but do not use more unit address blocks than there are nonshared subchannel groups.

For example, assume that the 32 consecutive unit addresses 100-11F are specified on channel 1, and that four groups (32 nonshared subchannels) are assigned to the four blocks (32 unit addresses). All 32 subchannels are used because the unit addresses and unit control words are consecutive.

Now assume that the four nonconsecutive unit addresses 100, 108, 110, and 118 are specified on channel 1. Four groups (32 nonshared subchannels) are assigned to the four blocks (four devices). Only four of the 32 subchannels are used but eight consecutive unit control words are assigned beginning with each unit address.

The total number of real, virtual, and nonexistent unit control blocks in the SYSGEN must not drain the pool of nonshared subchannel groups for the System/370 Models 145, 155II, and 158.

On a Model 145 or 155II block multiplexer channel, the channel reverts to selector mode and "hangs" the 3830 Model 3 Storage Control or ISC with Staging Adapter if a nonshared subchannel is not available. In that case, you must IPL to recover.

On a Model 158, the block multiplexer channel serially shares the last available nonshared subchannel and queues the unit addresses. If a nonshared subchannel is not available, a busy condition occurs. The 158 block multiplexer channel operates in a degraded mode, but it does not revert to selector mode and does not "hang" the 3830 Model 3 Storage Control or ISC with Staging Adapter.

IBM 2880 Block Multiplexer Channel

On a Model 165II, 168, 168AP, or 168MP, the standard 2880 Block Multiplexer Channel provides 256 unit addresses, 00-FF, available on one shared subchannel and up to 56 nonshared subchannels. During installation, the Customer Engineer plugs a circuit card for each group of eight nonshared subchannels. One nonshared unit control word (UCW) is required for each unit addrress recognized by a control unit. If seven groups are plugged for nonshared operation, 57 channel programs can operate concurrently (56 nonshared devices on the nonshared subchannels and 1 shared device on the shared subchannel). If a nonshared UCB is not plugged for each nonshared unit address, undefined errors can occur when the address is used.

All eight groups can be plugged for nonshared operation. If the conditions described in the *IBM System/370 Model 168 Functional Characteristics* manual are met, up to 64 nonshared devices can operate concurrently.

One Staging Adapter can be connected to a 2880 Block Multiplexer Channel with 64 nonshared subchannels plugged (eight groups) and meet the conditions described in the *IBM System/370 Model 168 Funtional Characteristics* manual by:

- Specifying the channel interface address as 00
- Assigning the 64 UCWs as 00-3F.
- Specifying the 64 device addresses as 00-3F. If not all 64 addresses are required, addresses 00-07 must be specified.

With the extended unit control words (feature codes 3851 and 3852) installed, the capability of the 2880 Block Multiplexer Channel to operate nonshared subchannels can be extended, in groups of eight, to 256.

Mass Storage Control Addresses

A unit address and a nonshared unit control word must be specified for each Mass Storage Control, either primary or alternate, whether or not it is connected to a byte or block multiplexer channel. If the Mass Storage Control is connected to a block multiplexer channel with other devices, the unit address assigned to the Mass Storage Control must be outside the address ranges of the other devices.

Unit Control Word Example

The following example illustrates the unit address ranges, unit control word assignments, and potential conflicts you should consider when different control unit types are connected to the same block multiplexer channel. Assume block multiplexer channel 5 has:

- One 3830-3 with 46 virtual and real devices
- One 3830-2 with 6 real devices
- A 3851 Model B2 Mass Storage Facility (primary and alternate Mass Storage Control)
- Eight tape units

The 54 nonshared devices require 54 unit addresses and 88 nonshared UCWs. The latter are plugged in groups of eight and assigned as follows:

- 3830-3 Storage Control, 46 unit addresses, 500-52D, and 64 UCWs, 00-3F
- 3830-2 Storage Control, 6 unit addresses, 540-545 and 16 UCWs, 40-4F
- Mass Storage Controls, 2 unit addresses 560-561, and 8 UCWs, 60-67.

The eight tape units are assigned unit addresses 580-587 but require only one shared UCW on the shared subchannel. The extended unit control word feature is required on the 2880 Block Multiplexer Channel.

The Mass Storage Control unit addresses, 560-561, are beyond the address range of the 3830-2 Storage Control with the expanded 32 device addresses, that is 540-54F or 540-55F. However, if the 3830-2 is upgraded to a 3830-3, an address conflict occurs because Mass Storage Control unit addresses 560-561 are within the 3830-3 Storage Control range, which is 540-57F. A better address choice for the Mass Storage Controls is unit addresses 5A0-5A1 with nonshared UCWs A0-A7. Note that nonshared UCWs A2-A7 can be used by assigning unit addresses 5A2-5A7 to other single address devices, such as reader and printer control units, which are connected on the block multiplexer channel 5.

Shared DASD

A direct access device is sharable between two or more CPUs provided that each CPU has a unit address for, and an electrical path to, the device. Each CPUs access to the device is synchronized so that if one CPU is accessing the device, a busy condition is presented when another CPU attempts to access the same device. However, without additional controls, there are data integrity exposures if one CPU changes data or a VTOC while another CPU is updating data.

To provide additional controls, specify the FEATURE=(SHARED) command for the unit address of the shared devices on the IODEVICE cards. With the FEATURE=(SHARED) command, the access of each CPU to a shared device is synchronized as before but now the using program can reserve the device for exclusive use; for example, a VTOC update operation. Reserving the shared device ensures that the device is available only to the CPU channel path (primary or alternate) that issued the Reserve command. If any other sharing CPU attempts to reserve the shared device, a busy condition is presented and the reserve request is queued. When the issuing program completes its update, it issues a Release command. However, the device remains reserved to the CPU if there are other outstanding Reserve commands for that device. When all outstanding Reserve commands are completed, the device is released to the other sharing CPU.

Because sharing CPUs are unaware of each other's activity, the using program is responsible for issuing Reserve and Release commands as required for data integrity. Do not assume that Reserve and Release commands are automatically issued because the FEATURE=(SHARED) command is specified. However, there are certain System Control Program modules such as Linkage Editor, DADSM, and CATLG, which if shared, automatically use Reserve and Release command logic for their own updating operations.

Each sharing CPU must have an electrical path to the shared device that can be switched either by:

- Channel switching with two or four channel switches on a single control unit to the device, or
- Channel switching with two or four channel switches on two control units and string-switching on the 3333 or 3350.

There is a logical difference between sharing a real volume and sharing a virtual volume. To share a real volume, the volume must be mounted on a real device that is shared by one or more CPUs. To share a virtual volume, the virtual volume must be mounted on a virtual shared device in a staging drive group that is available to each sharing CPU. Therefore, to access a real or virtual shared volume, each sharing CPU must have:

- The real unit address of the electrical path that always maps to the same real device.
- The virtual shared unit address in the same staging drive group. Because the same virtual volume is logically mounted on each CPU's unique virtual shared unit address, the Staging Adapter must translate each CPU's virtual shared unit address to the staging drive real address and the virtual cylinders to the real cylinders where the shared data exists.

Because a virtual volume can be mounted on more than one virtual shared drive, the Mass Storage Control manages the Reserve and Release commands for Staging Adapter pairs. Before one of the Staging Adapters can reserve a virtual shared drive to one of its interfaces, the Mass Storage Control checks the tables of the other Staging Adapter for a Reserve command issued for any virtual shared device that has the same virtual volume mounted. This is in contrast to a 3830 Model 2 Storage Control pair, where either Storage Control can determine whether a device is reserved by testing a hardware latch in the 3333 or 3350.

If the FEATURE=(SHARED) command is desired in a multihost environment:

- It must be added to the IODEVICE card that is generated by the Mass Storage Control Table Create program for the real and convertible-to-real drives.
- It is punched in the IODEVICE cards by the Mass Storage Control Table Create program for the addresses that are specified as virtual shared (VS) in the UPPERCON command.

A virtual volume cannot be shared unless it is mounted on virtual shared addresses. Therefore, the mount operation must be externally directed to virtual shared devices or all virtual unit addresses must be specified as virtual shared. The mount operation can be controlled by the specific 3330V unit addresses or use of esoteric names for all virtual and virtual shared addresses or by permanently mounting the volumes in PRESRES or VATLST. If virtual and virtual shared unit addresses are mixed within a staging drive group, there is no assurance that a shared virtual volume will be mounted on a virtual shared unit address. The Mass Storage System Communicator (MSSC) module in OS/VS selects the online unit addresses for mounting as they become available. Therefore, a shared volume can be mounted on either a virtual or virtual shared address. However, if a virtual volume is mounted to a CPU, the Mass Storage Control checks that the virtual unit address is shared and within the same staging drive group before mounting the virtual volume to another CPU.

The need for shared DASD should not be confused with the sharing of data sets in a multiprogramming environment. Data sets can be shared:

- Between tasks of one job that is executing in one CPU
- Between two or more jobs executing in one CPU with DISP=SHR specified in the OS/VS JCL

  Note: DISP=OLD does not cause the System Control Program to issue a Reserve or Release command to a shared device.
- Between two or more jobs that are executing in two or more CPUs. This environment requires Reserve and Release commands when mutually exclusive operations use shared devices.

  Note: The OS/VS JCL parameters DISP=OLD or DISP=SHR in one System Control Program have no effect in the other sharing System Control Program.

Unless plans are made to control the mounting of shared virtual volumes, virtual shared addresses should be specified if volume sharing in a multihost environment is desired.

Connecting Real Drive Strings

Real drives strings, either 3333/3330 or 3350A/3350B, connected to a Staging Adapter pair as shown in Figure 13 of the "Logical Connections" chapter are defined as being within an MSS environment. Real drives connected to a Staging Adapter and and a 3830 Model 2 (or ISC pair) as shown in Figures 31 and 31.1 are defined as being outside an MSS environment. Because of potential data integrity exposures and unsolicited interrupt problems, IBM does not recommend the string-switching of drives between an MSS and non-MSS environment. This means that 3333s controlling staging drives and convertible-to-real drives cannot be string-switched to a Staging Adapter and 3830 Model 2 (or ISC) pair; they must be string-switched to a Staging Adapter pair. In addition, any host that is connected to a Staging Adapter must also be connected to the Mass Storage Control.

Real drives can be string-switched to Staging Adapter and 3830 Model 2 (or ISC) pairs with a Staging Adapter controlling the staging drives as shown in Figure 31 or, with a Staging Adapter pair controlling the staging drives as shown in Figure 31.1. The Real Unit Address Table in Figures 31 and 31.1 shows the primary and alternate logical paths for the two CPUs; one in an MSS environment and one outside an MSS environment.

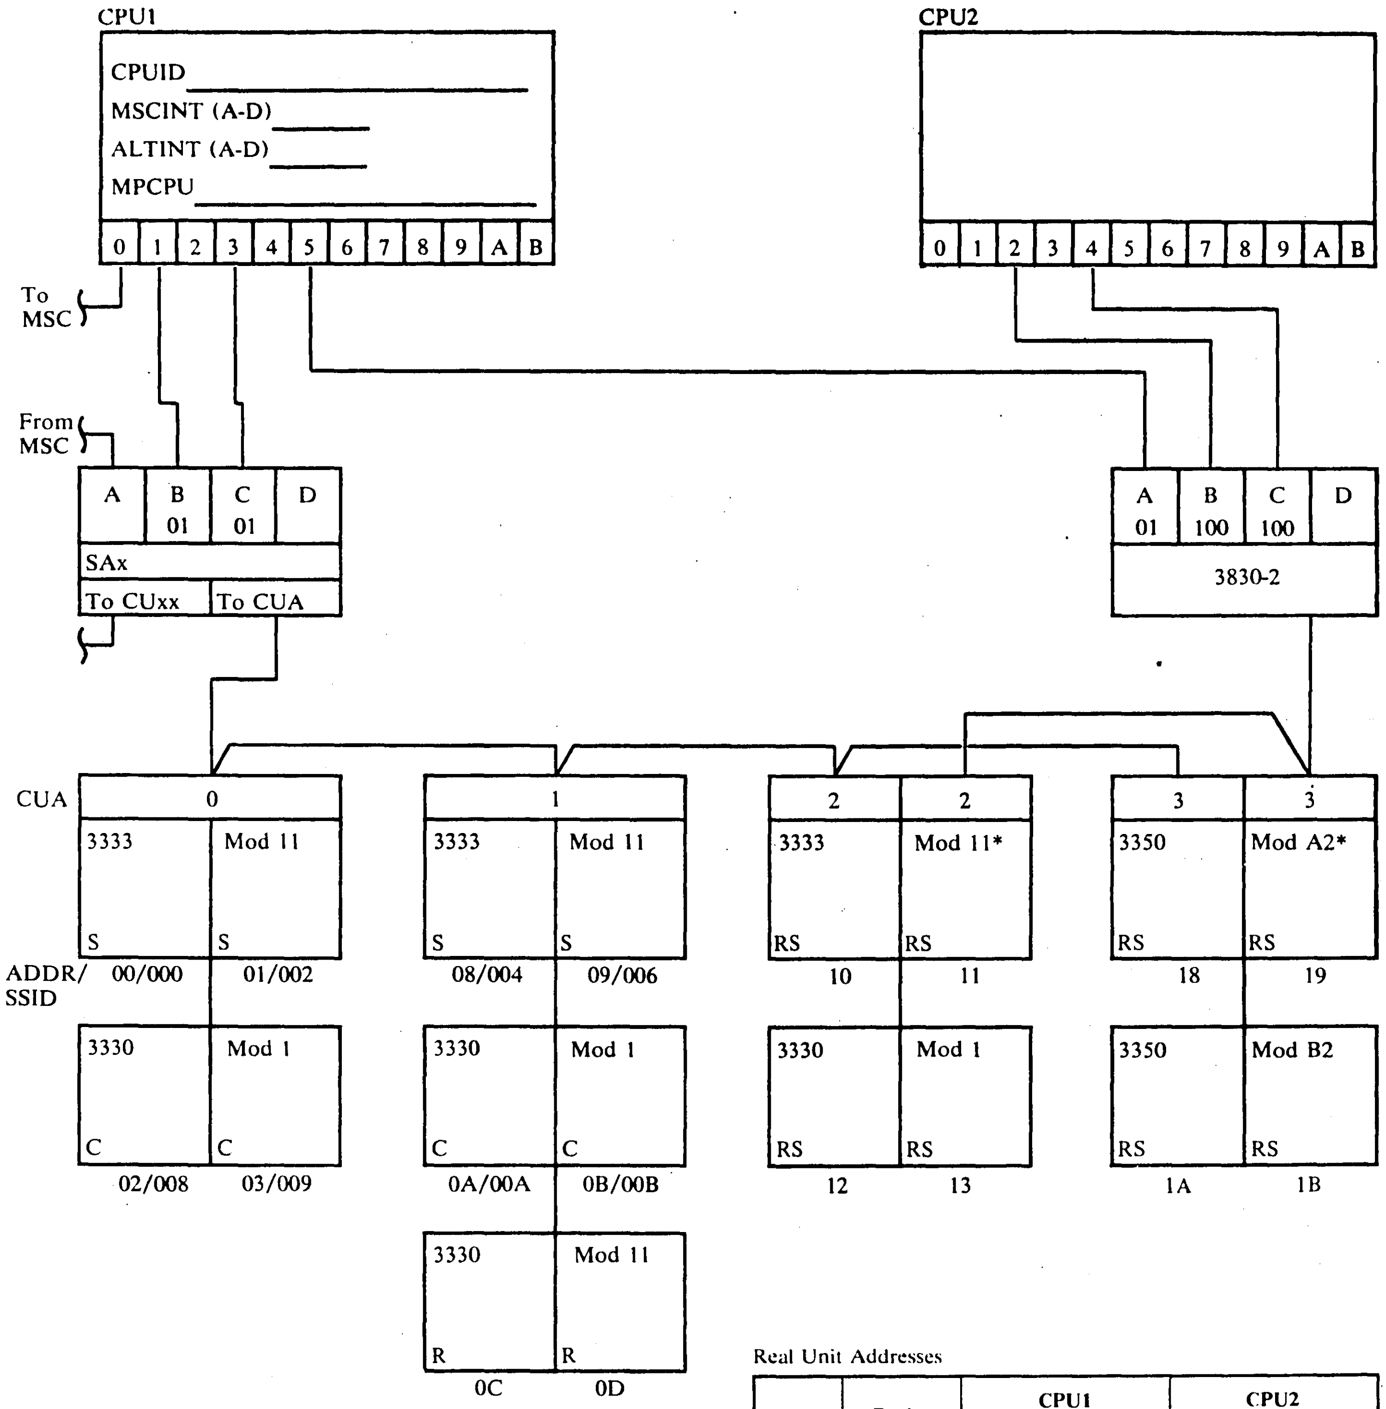

Real Unit Addresses

String	Device Address	CPU1		CPU2	
		Primary	Alternate	Primary	Alternate
1	00-03	Staging	Staging	**	**
2	08-0B	Staging	Staging	**	**
2	0C-0D	14C-14D	34C-34D	**	**
3	10-13	150-153	550-553	290-293	490-493
4	18-1B	158-15B	558-55B	298-29B	498-49B

Legend:

* - String Switch Feature
** - No Path Permitted
C - CONVREAL
R - Real
RS - Real Shared

Figure 31. Connecting Real Drive Strings–Example 1

As shown in Figure 31, CPU2 can have only electrical and logical paths to real drive strings 3 and 4 through the 3830 Model 2. CPU1 has electrical and logical paths to all of the drives (real, convertible-to-real, and staging) through the Staging Adapter. However, the alternate channel path for CPU1 to real drive strings 3 and 4 is through the 3830 Model 2. Real devices 0C-0D, 10-13, and 18-1B are specified to the Mass Storage Control Table Create program for CPU1 as if the 3830 Model 2 did not exist. FEATURE=(SHARED) must be added to the IODEVICE cards for real devices 10-13 and 18-1B.

As shown in Figure 31.1, CPU2 can have only electrical and logical paths to real drive strings 3 and 4 through the 3830 Model 2 and the ISC. CPU1 has electrical and logical paths to all drives (real, convertible-to-real, and staging) through the Staging Adapter pair. The primary channel path to real drive string 3 for CPU1 is through Staging Adapter *x* and to real drive string 4 it is through Staging Adapter *y*. However, CPU1's alternate channel path to real drive string 3 is through the 3830 Model 2 and to real drive string 4 through the ISC. Real devices 0C-0D, 10-13, and 18-1B are specified to the Mass Storage Control Table Create as if the 3830 Model 2 and ISC did not exist. Although the CUA parameter is the same for real drive strings 3 and 4 (CUA=2), unique values can be specified with a corresponding change in the real unit address. FEATURE=(SHARED) must be added to the IODEVICE cards for real devices 10-13 and 18-1B.

The disk and control unit configurations shown in Figures 31 and 31.1 were developed only to illustrate several methods of connecting and addressing real drive strings. If access to real drives or real drive strings through a Staging Adapter is required, the overall effect of the additional workload on the Mass Storage System should be considered.

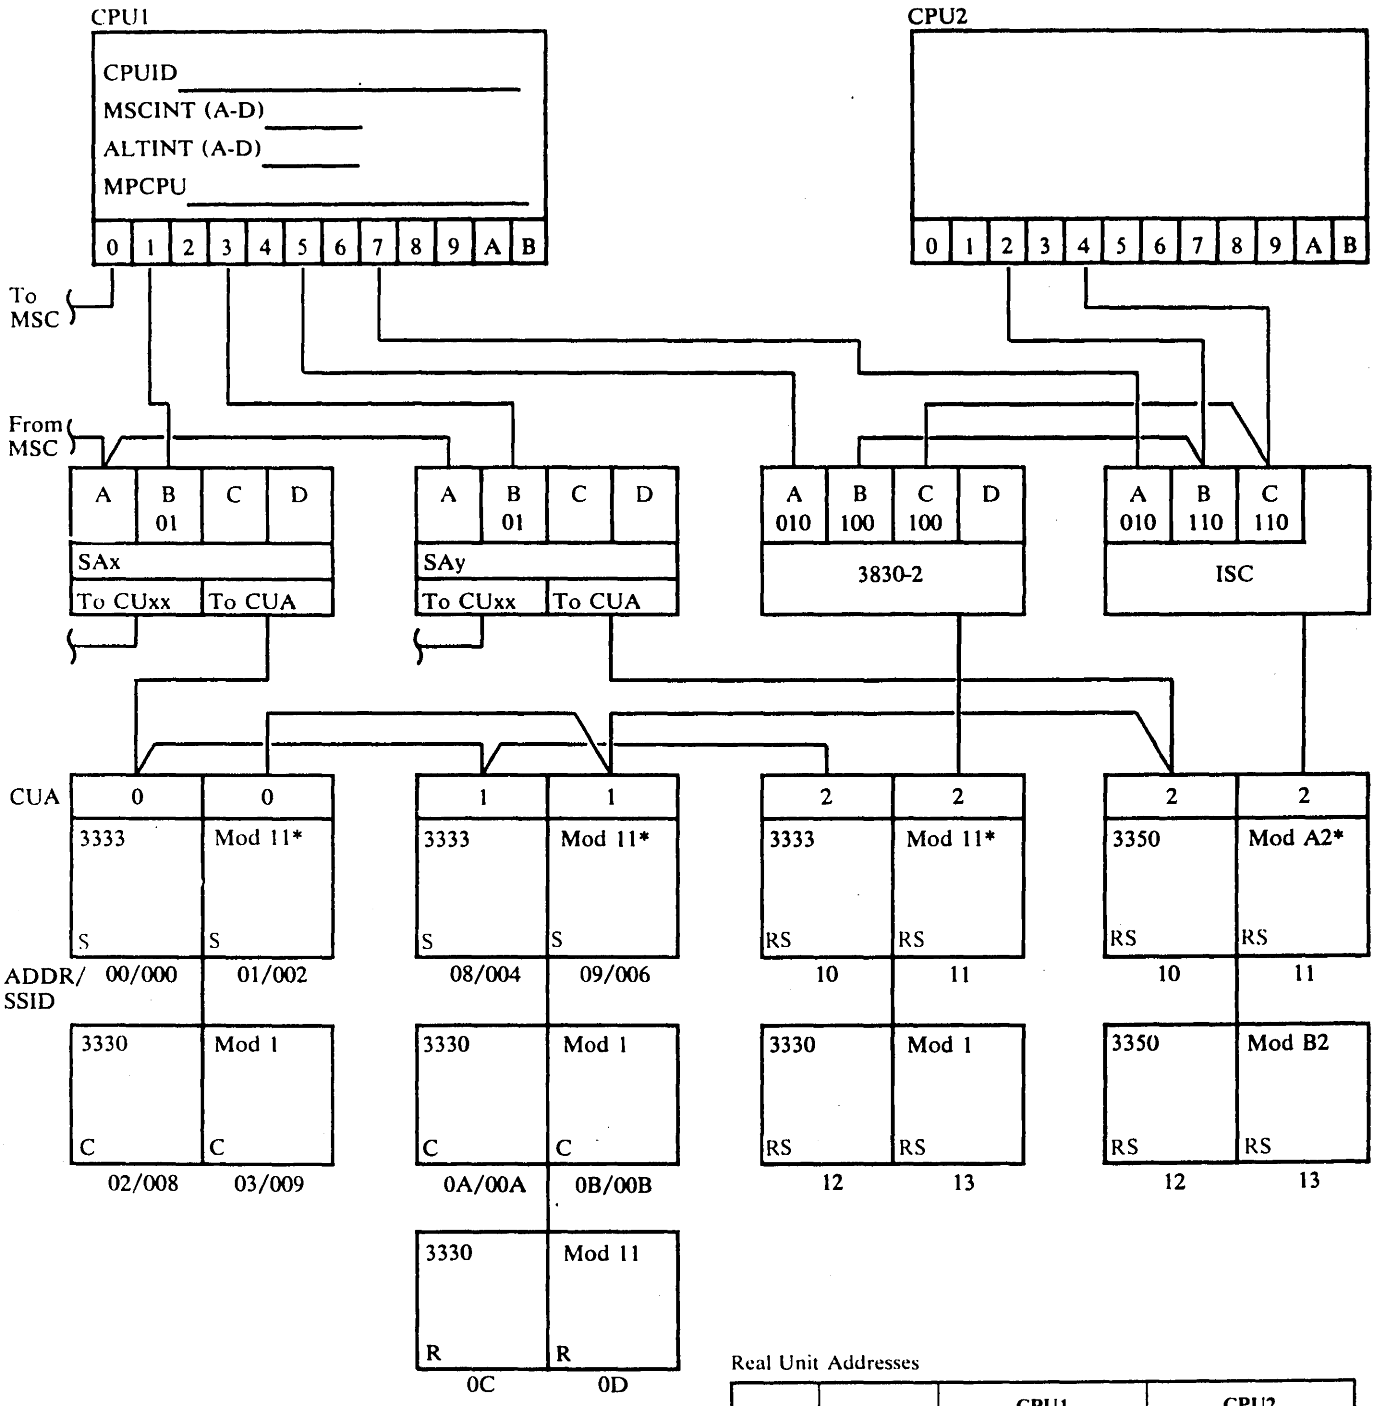

Real Unit Addresses

String	Device Address	CPU1		CPU2	
		Primary	Alternate	Primary	Alternate
1	00-03	Staging	Staging	**	**
2	08-0B	Staging	Staging	**	**
2	0C-0D	14C-14D	34C-34D	**	**
3	10-13	150-153	550-553	290-293	490-493
4	10-13	350-353	750-753	2D0-2D3	4D0-4D3

Legend:
* - String Switch Feature
** - No Path Permitted
C - CONVREAL
R - Real
RS - Real Shared

Figure 31.1 Connnecting Real Drive Strings—Example 2

Physical Planning Considerations

See *IBM System/370 Installation Manual - Physical Planning* for detailed physical planning and cable length information. The following description highlights physical layouts that might be restricted by critical cable lengths, and shows them as they appear in the Mass Storage Control Table Create Program Worksheet 1. These simplified diagrams are only an overview, but they demonstrate the need for careful and detailed physical planning.

Cabling Summary

Figure 32 shows Connections I-V and a table of maximum cumulative cable lengths for each connection in a Mass Storage System configuration with either single or multiple Mass Storage Facilities. These lengths are not necessarily the maximum distances between Mass Storage System components but rather a cumulative total of the "X" lengths of all devices on a particular path. "X" lengths are point-to-point distances between the cable entries to each component.

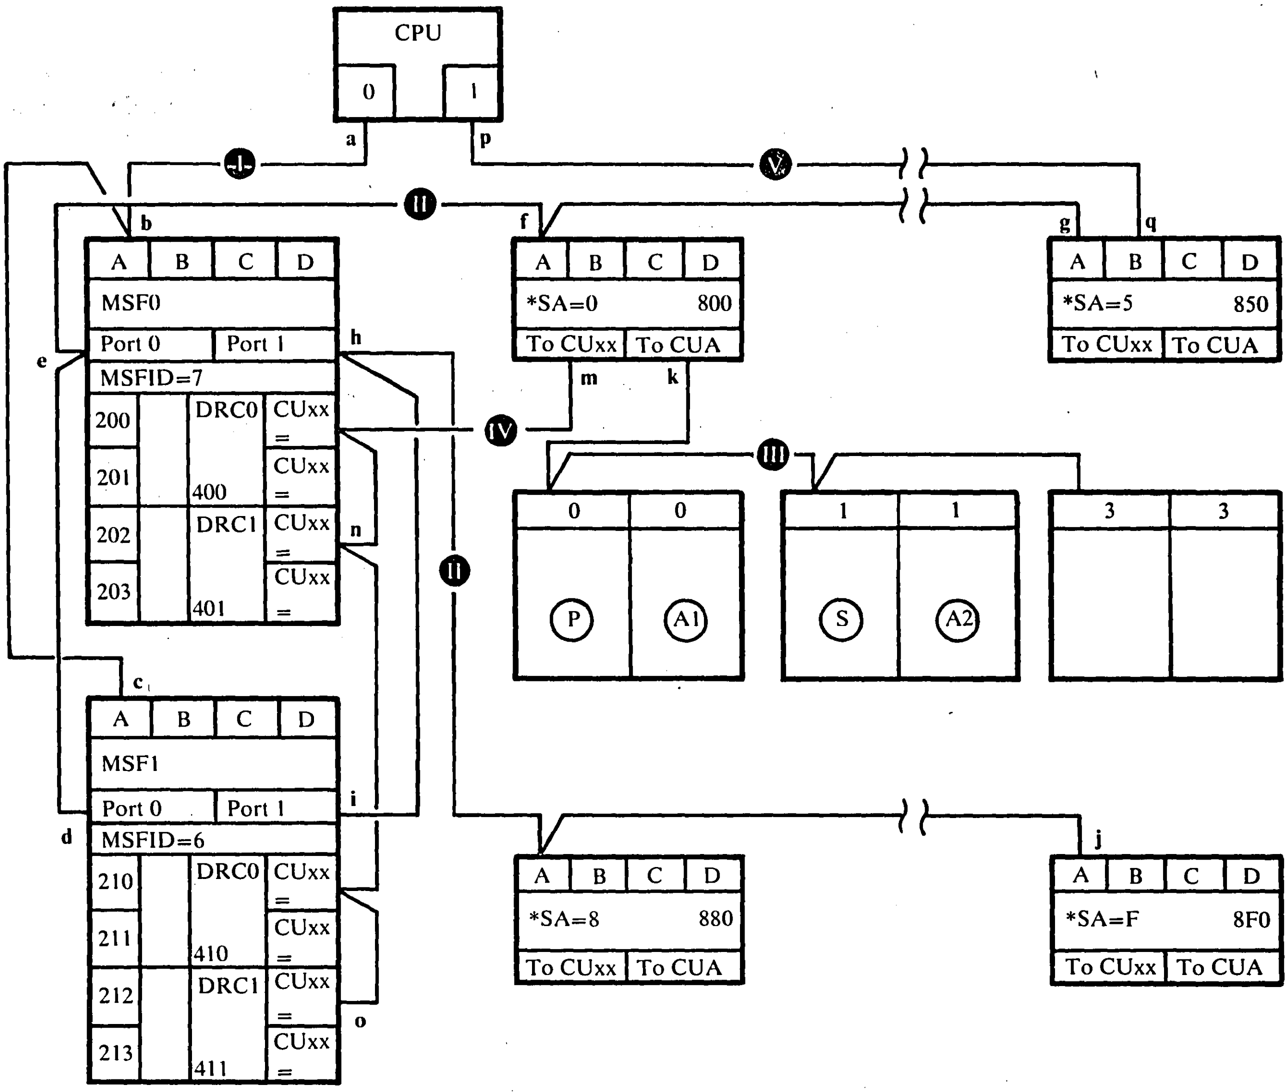

*Staging Adapter in 3830-3 or ISCSA.

P = Primary Tables
A1 = Alternate Area 1
S = Secondary Tables
A2 = Alternate Area 2

Maximum Cumulative Cable Lengths					
Single 3851 Connection		**Multi-3851 Connection**		**Feet**	**Meters**
I	a-b	I	a-c	200	61
II	e-f	II	d-f	150	46
II	e-g	II	d-g	300	91
II	h-j	II	i-j	300	91
III	k-l	III	k-l	200	61
IV	m-n	IV	m-o	200	61
V	p-q	V	p-q	150	46

Note: As shown, MSF0 is specified as secondary (9121) and MSF1 is specified as primary (9120).

Figure 32. Connections I-V, Maximum Cumulative Cable Lengths

Connection I from the CPU to the 3851 has a maximum cumulative cable length of 200 feet (61 meters). This includes the point-to-point "X" lengths from the CPU to all byte multiplexer channel devices plus MSF0 and MSF1. The maximum cable length possible between MSF0 and MSF1 is affected by this connection.

Connection II from the Mass Storage Control ports to the last Staging Adapter is a maximum cumulative length of 300 feet (91 meters), but the distance from Mass Storage Control port 0 to the Staging Adapter (SA=0 or SA=1) that controls the Mass Storage Control tables can be only 150 feet (46 meters). Connection II from Mass Storage Control port 1 does not have this restriction because the Staging Adapter that controls the Mass Storage Control tables cannot be connected to port 1.

Connection III from the Staging Adapter to the 3333 Disk Storage and Control has a maximum cumulative length of 200 feet (61 meters).

Normally, the maximum cumulative cable length available to attach control units to a channel is 200 feet (61 meters), but Connection V is limited to 150 feet (46 meters) because the 3830 Models 2 and 3 Storage Controls have high data transfer rates.

Single Mass Storage Facility Layout

Figure 33 shows a simplified physical layout for a single Mass Storage Facility with restrictive cable lengths for either a 3830 Model 3 Storage Control (Example A), or an ISC with Staging Adapter (Example B). Although the distance from the CPU to the Mass Storage Facility is affected by the cumulative "X" lengths for all devices on the Connection I channel path, it can also be affected by Connection II if the Mass Storage Control tables are under the control of an ISC with Staging Adapter.

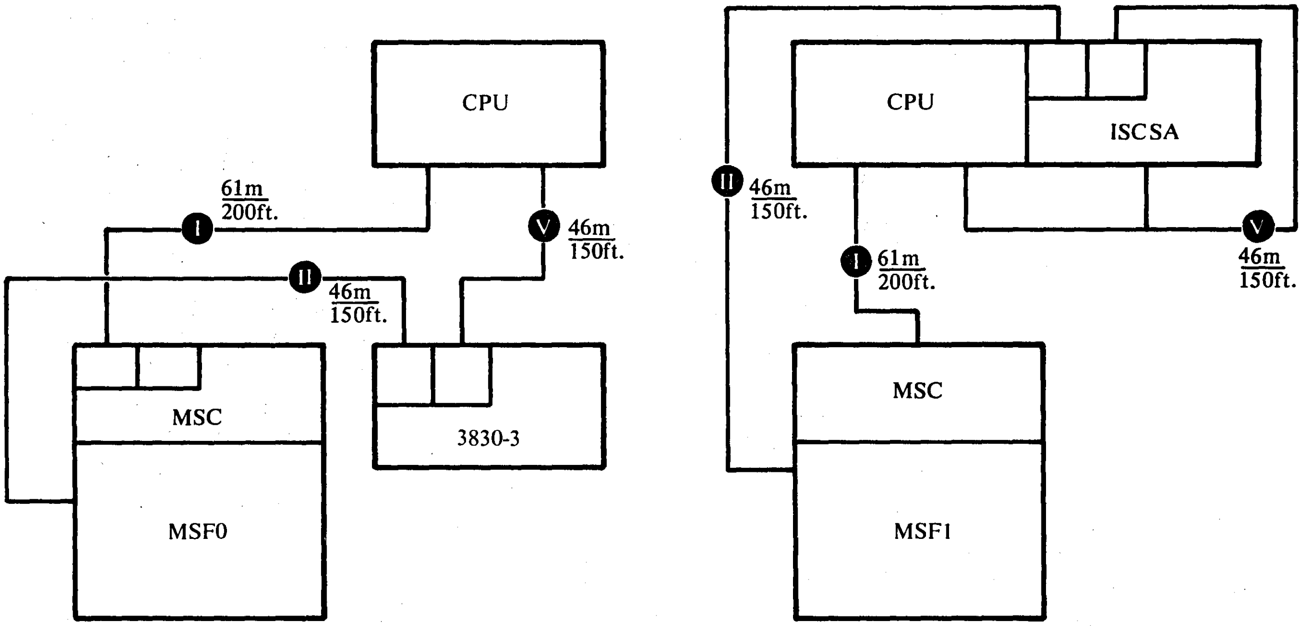

Example A. With 3830-3

Example B. ISC with Staging Adapter.

Figure 33. Single Mass Storage Facility Layout

Multi-Mass Storage Facility Layout

Figure 34 shows three simplified layouts of multi-Mass Storage Facility configurations. The layouts illustrate the factors that affect the distances between the CPU, Mass Storage Facilities, and the Staging Adapters (3830 Model 3 Storage Controls) that control the Mass Storage Control tables. These factors are:

- Primary or secondary Mass Storage Facility specifications (9120 and 9121)
- Maximum length of Connection I
- Maximum length of Connection II

The primary and secondary specificationss determine which Mass Storage Control is initialized first at Initial Microprogram Load time (with both Mass Storage Controls enabled) and the Mass Storage Control port changes that must be made for multi-Mass Storage Facility operation. The primary Mass Storage Facility has only cable-out connections in the two Mass Storage Control ports. The secondary Mass Storage Facility has cable-in and cable-out connections in the two Mass Storage Control ports. As shown in the three examples, the primary or secondary specifications affect the point-to-point cabling of Connection II.

Connection I affects the distance between the CPUs and the Mass Storage Facilities, and the distance between the Mass Storage Facilities.

Connection II, from the Mass Storage Control to Staging Adapters 0 and 1, is limited to 150 feet (46 meters) because of the data flow between the Mass Storage Control and its tables. Connection II, from the Mass Storage Control to Staging Adapters 2 through 13, can be 300 feet (92 meters).

Example A in Figure 34 illustrates the minimum separation between Mass Storage Facilities. Connection II must be cabled from MSF0 (primary) to MSF1 (secondary) and then doubled back to the Staging Adapter.

Example B increases the separation between the Mass Storage Facilities by changing the primary and secondary specification. Connection II is then cabled from MSF1 (primary) to MSF0 (secondary) and to the Staging Adapter without doubling back. MSF0 can be located at any point between MSF1 and the Staging Adapter.

The maximum separation between Mass Storage Facilities is shown in Example C. Here, the primary and secondary specification is not important because the Mass Storage Control cable-in connection is not used. Each Mass Storage Facility has in effect its own Connection II (150 feet or 46 meters) to the Staging Adapter but the distance between the Mass Storage Facilities is controlled by connection I. In this example, the select-in and select-out path (but not the data path) from the Mass Storage Control in MSF1 use Connection II from MSF1 to MSF0 to the Staging Adapter, a maximum distance of 450 feet (137.3 meters).

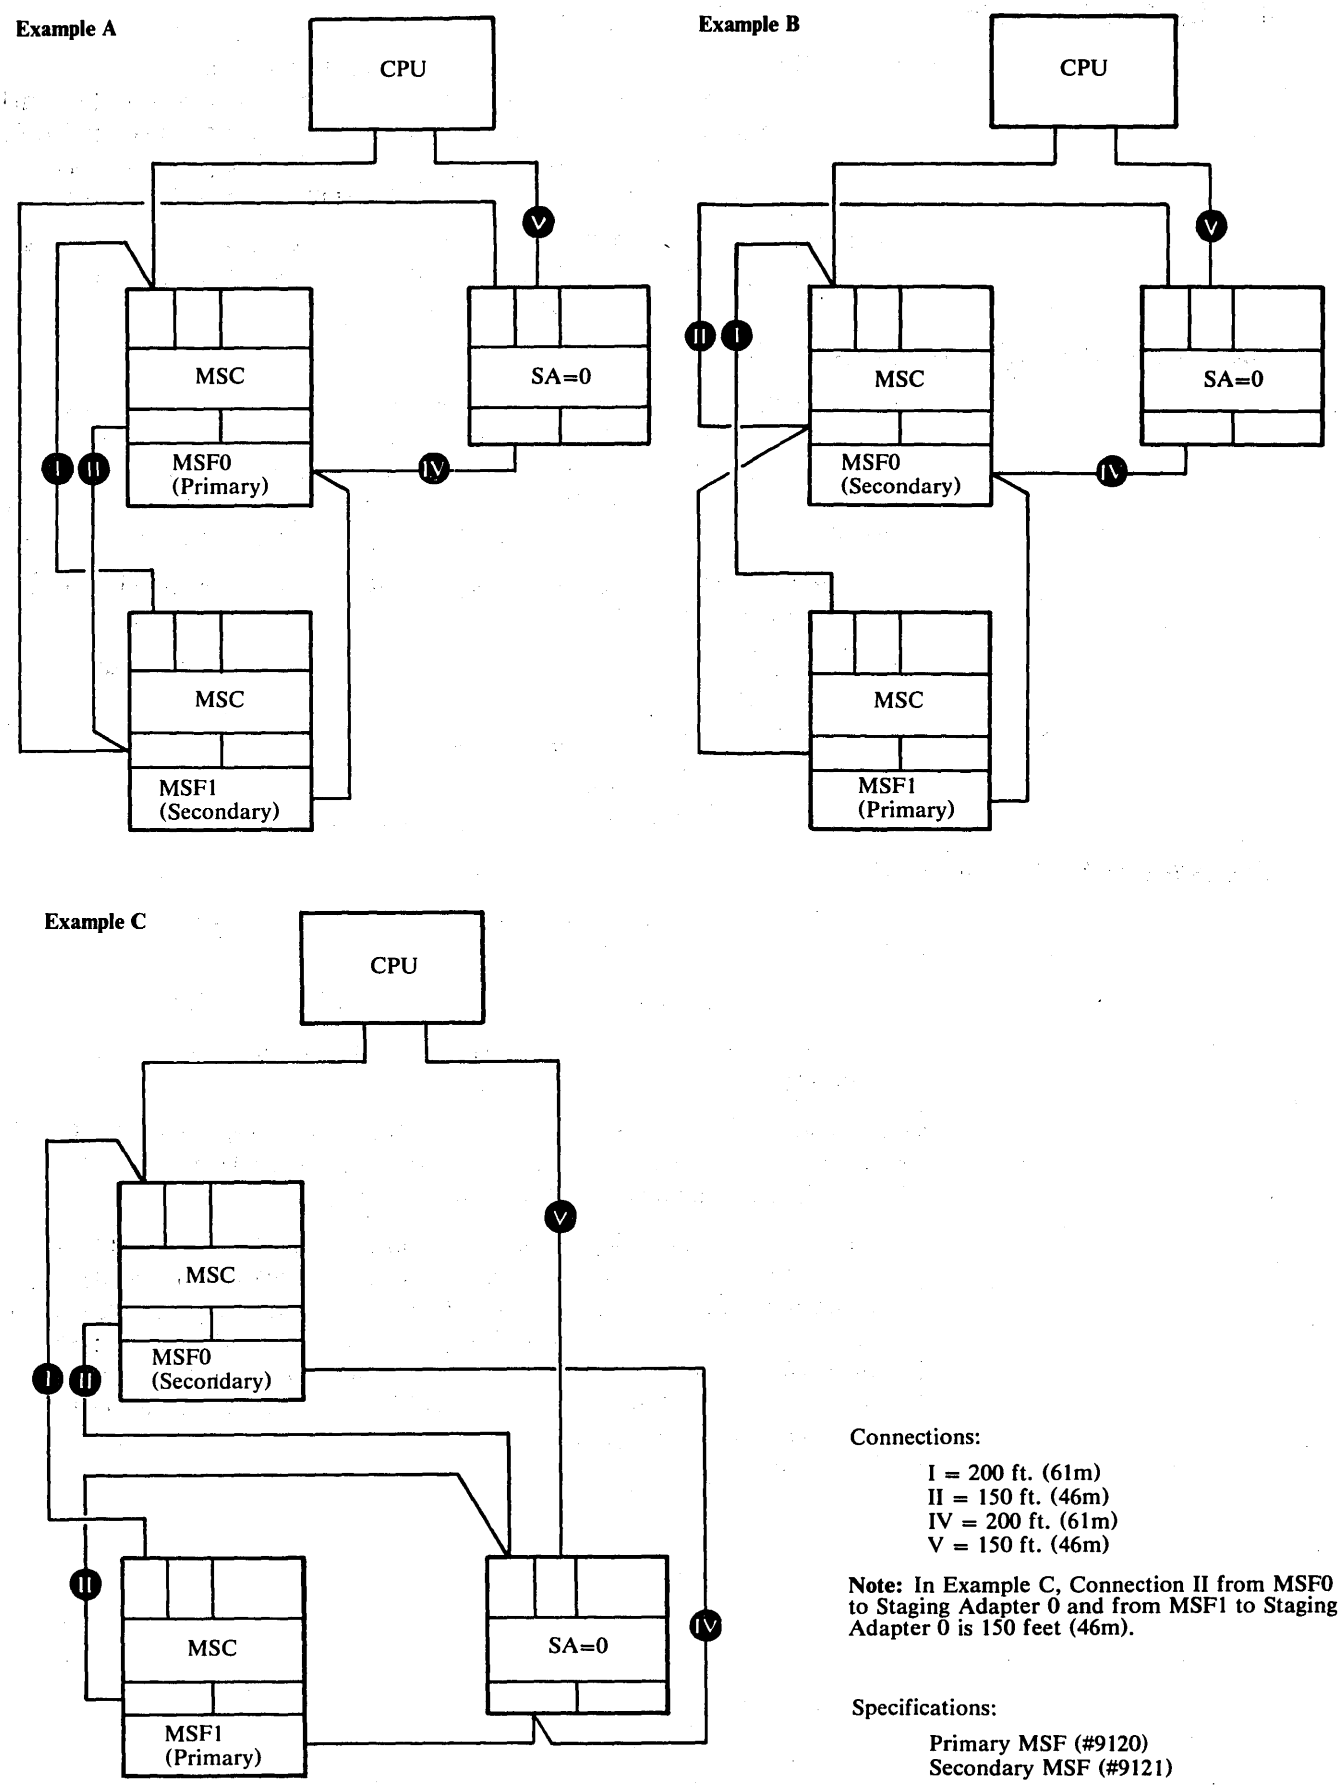

Connections:

I = 200 ft. (61m)
II = 150 ft. (46m)
IV = 200 ft. (61m)
V = 150 ft. (46m)

Note: In Example C, Connection II from MSF0 to Staging Adapter 0 and from MSF1 to Staging Adapter 0 is 150 feet (46m).

Specifications:

Primary MSF (#9120)
Secondary MSF (#9121)

| **Figure 34. Multi-Mass Storage Facility Layout**

Portable Terminal 2 (PT-2)

A portable video terminal is used for local and remote (optional) maintenance. It can be used for diagnostic purposes to display the CE panel and microprogram at a remote location. The terminal connects to the 3851, ISC with Staging Adapter, and to the 3830 Model 3. Remote maintenance is a customer option and requires a telephone reserved for this purpose.

Either a telephone with a long cord, or multiple telephone jacks is required so that the telephone is within 25 feet (7.6 meters) of all terminal cable connectors in the 3851 and Staging Adapters.

Mass Storage Facility Without Raised Floor

The Mass Storage Facility has been designed for installation on a raised floor for safety, cleanliness, appearance, and convenience. For these reasons, we recommend a raised floor. However, there are installations where a raised floor is not practical or possible as can occur when the Mass Storage Facility is installed in an existing vault for security reasons.

Any Mass Storage Facility model can be installed without a raised floor if these additional factors are considered:

- Cable paths are planned to avoid normal operation and service areas
- Cable ramps are planned for the safety of personnel and protection of cables when layouts must pass through traffic areas
- The 3851 may have to be lifted slightly during cable installation to permit passage of large cable connectors beneath it

Mass Storage Facility Model Changes

When planning a Mass Storage Facility model change, two items must be considered:

- Removal of the data cartridges from the Mass Storage Facility
- Execution of Mass Storage Control Table Create program with the UPDATE parameter

The removal of data cartridges is not required for a Mass Storage Facility model change but you may want to consider it for security reasons. The two methods of removal and reentry are:

- Remove all data cartridges and mass storage volumes using the Access Method Services EJECTC command with the ALL option and the STOREV command with the EJECT option. When the data cartridges are reentered at the cartridge access station of the new model Mass Storage Facility, they will be stored in new locations according to the LOADMAP option you specify in the Mass Storage Control Table Create program MSFn command, or according to the default LOADMAP for the new model Mass Storage Facility.

- Remove all data cartridges using the Access Methods Services MODIFYC command with the DIRECTEJECT option. However, the Mass Storage Control tables are not changed and when the data cartridges are reentered through the cartridge access station of the new model Mass Storage Facility, they will be stored in their previous locations. Therefore, the Mass Storage Control tables must be current and must contain the correct location and status of all data cartridges.

If the Mass Storage Control Table Create program with UPDATE option is run prior to the model change, there are two items to consider:

- The updated Mass Storage Control tables cannot be used with the old model Mass Storage Facility because the accessor must reach the reserve cells during the Initial Microprogram Load of the Mass Storage Facility. The location of the reserve cells varies depending on the model of Mass Storage Facility.
- No changes can be made to the data cartridge inventory or volume status after the table update because the physical inventory would not be synchronous with the updated Mass Storage Control tables.

When changing Mass Storage Facility models, execute the Mass Storage Control Table Create program with the NEW and PRTONLY options. Decide whether to remove the cartridges and the method of removal. If you decide to remove the cartridges, follow these steps:

1. Hold all jobs that require the Mass Storage System and allow the work that is in process in the Mass Storage System to complete.
2. Vary all staging drives offline.
3. When all staging drives are offline, eject all data cartridges.
4. Execute the Mass Storage Control Table Create program with the UPDATE parameter.
5. Update all Mass Storage System checkpoints.
6. Complete the model change.
7. Initialize the new Mass Storage Facility using the updated Mass Storage Control tables.
8. If data cartridges were removed in step 3, enter them and resume operation.

Input/Output Priority Sequence

Channel capabilities are affected by the priority or sequence in which input/output control devices are attached to the channels. A summary is presented here for the 3830 Model 3 Storage Control and the 3851 Mass Storage Facility. See the *IBM System/370 Installation Manual - Physical Planning* or consult with your IBM Installation Planning Representative for details.

The 3830 Model 3 Storage Control is a Class 1 priority device (one that is subject to data overrun), and must be attached to a block multiplexer channel. Class 1 devices with the highest data rates are normally attached to the lowest numbered channel, that is channel 1. The attachment sequence within a class is in order of decreasing device data rates.

The 3851 Mass Storage Facility is a Class 3 priority device (one that does not require channel service to be in synchronization with device operation), and can be attached to either a byte or block multiplexer channel. Input/output control devices are attached to a byte multiplexer channel in class sequence, that is Class 1, Class 2, and Class 3. The attachment sequence within a class is in order of increasing critical wait time intervals.

Control devices are addressed by the channel through an attachment cable that contains select-out and select-in lines. The priority sequence of a control device is determined by its connection to either the select-out or select-in line and its physical position along the attachment cable.

Within the Mass Storage System, there is a similar input/output priority from the Mass Stroage Control ports to the accessor controls and the Staging Adapters. The accessor controls are on port 0 and are first in priority on the select-out lines. The priority of the other Mass Storage System components can be established at installation time. If all Mass Storage System components are sequentially connected to the select-out lines with MSF0 specified as secondary (9121), and MSF1 specified as primary (9120), the priority sequence is:

- MSF1 accessor control (MSFID=6), MSF0 accessor control (MSFID=7), SA=0, SA=1,...... SA=6 on Mass Storage Control port 0.
- SA=8, SA=9, ...SA=F on Mass Storage Control port 1.

Staging Drive Group Space

The status of staging drive group pages changes as mass storage volumes are mounted and demounted, as staging drives are varied online and offline, and as pages are acquired and relinquished. When the number of active pages exceeds the upper threshold in a staging drive group, the LRU algorithm is called to identify the mass storage volume that has the largest number of oldest pages for temporary deallocation.

This chapter reviews both normal mount and demount operations and LRU page activities. We have not covered all examples, but have provided sufficient information to understand how the LRU parameters specified to the Mass Storage Control Table Create program relate to staging drive group space management.

To understand space allocation, LRU, and deallocation in a staging drive group, picture a staging drive group with two 3330-1 staging drives or 102 free pages, as shown in Figure 35. Assume that you do not want the LRU algorithm to lower the number of active (bound and unbound) pages below 60 (lower threshold), and you want the LRU algorithm to be called when the current active page count exceeds 80 (upper threshold). These thresholds are determined by specifying page deltas (ACTPAGE=30,40) to the Mass Storage Control Table Create program. Note that the lower page delta and upper page delta is specified relative to the total number of pages on one logical 3330-1 volume. From these parameters, the Mass Storage Control calculates the lower and upper thresholds for each staging drive group, and recalculates them as staging drives are varied online or offline or as mass storage volumes are mounted or demounted with the volume bind attribute.

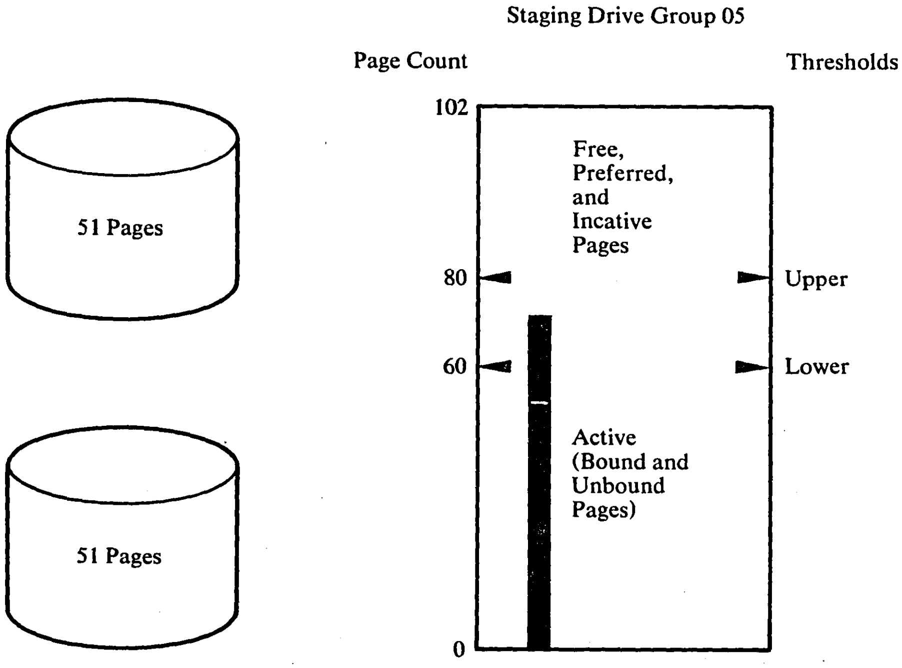

Figure 35. Staging Drive Group Page Thresholds

For staging drive group 05 in our example, the lower and upper thresholds are twice as large as the lower and upper deltas, that is, 60 and 80 pages. When normal mount and demount operations; in staging drive group 05 are balanced, volumes are demounted and pages become inactive at approximately the same rate as they are acquired.

The active page count fluctuates but does not exceed 80. Sufficient free and inactive pages are always available for new space allocations (free, preferred, or inactive) or data reuse (inactive).

When staging drive group 05 is selected for a mass storage volume mount operation, the Space Allocation routine searches for enough pages to contain the data set extents that are to be staged (request). Free (initial page condition), preferred (deleted and unused), and inactive read-only (demounted or LRU), pages are immediately available for allocation and are searched first, as indicated in Figure 36. Inactive written pages might not be immediately available because the destaging of cylinders that contain altered data is not complete and these cylinders are searched last.

Page Flow

As shown in Figure 36, space in a staging drive group is made up of six types of pages:

- Free—pages that become available for space allocation when a staging drive is varied online or when a demounted and destaged volume is remounted on another staging drive group.
- Preferred Inactive—pages that become available for space allocation because the data is no longer needed or the pages were not used to satisfy the original request. They include:
    1. Demounted bound volume pages that are unused or that are secure (volume attribute).
    2. Inactive staged—pages from a demounted but not yet destaged volume that has been remounted on another staging drive group.
    3. Secure pages (volume attribute) that were made inactive by the LRU algorithm or were demounted at end-of-job time.
    4. Active unbound pages that have been relinquished with the inhibit destage attribute or contain deleted data sets.
- Inactive destaged—pages that become available for allocation or data reuse when a mass storage volume is demounted.
- Inactive staged—pages that become available for new allocations or data reuse when the upper threshold for the staging drive group is exceeded and the LRU algorithm temporarily deallocates the oldest pages.
- Active unbound—pages that are allocated and acquired when a mass storage volume is mounted.
- Active bound—pages that are allocated, acquired, and bound (VSAM data set BIND attribute) when a mass storage volume is mounted.

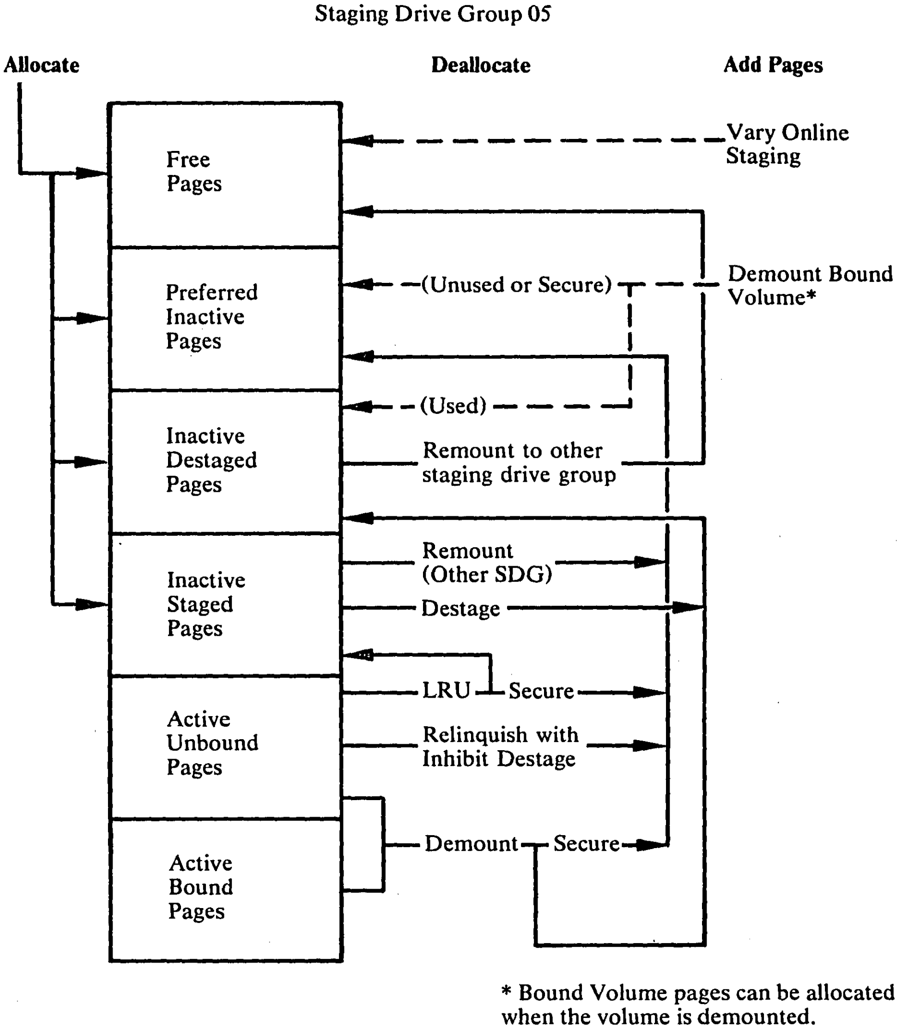

Figure 36. Page Flow

Pages on bound mass storage volumes are not available for space allocation and are not included when the Mass Storage Control calculates lower and upper thresholds.

Figure 36 shows allocation and the "flow" of pages from an active to a free, preferred, or inactive state. This normal flow occurs as mass storage volumes are demounted.

If the new active page count (request plus the active page count) exceeds the upper threshold, the LRU algorithm is called to reduce the new active page count by making the oldest active pages inactive. A time stamp on each page in the staging drive group determines the oldest pages. The LRU algorithm selects only the oldest unbound active pages, as shown in Figure 36.

Each Staging Adapter maintains a 16-position clock that runs consecutively from time-stamp position '0' to 'F' under control of the Mass Storage Control master clock. All Staging Adapter clocks advance at the same time, at intervals of 128 seconds, to 16,354 seconds as specified to the Mass Storage Control Table Create program by LRUCLOCK or as modified by the Access Methods Services Tune command.

When pages are first acquired and subsequently each time they are accessed, the Staging Adapter time-stamps the pages with its clock value and records the event in its tables. Not all Staging Adapter clocks necessarily have the same clock value, because some are advanced by the LRU algorithm to artificially age the pages. A page in constant use does not age because its time-stamp is continually updated with the current clock value. Staging Adapters pairs that are connected to string-switched staging drive groups always have the same clock value because the LRU algorithm advances the clocks in both Staging Adapters.

Before the LRU algorithm examines pages in a staging drive group, the Mass Storage Control normalizes the time-stamp values so that the oldest pages are stamped '0', and the newest pages 'F'. The normalized time-stamped pages are grouped into four (or five implied) time groups called slice ranges, as determined by the LRUGROUP parameter. For example, LRUGROUP=5,4,3,2 creates five slice ranges, that is I (five time-stamps), II (four time-stamps), III (three time stamps), and IV (two time stamps), and an implied V (two time stamps). LRUGROUP=1 creates one slice range, that is I (one time stamp).

Figure 37 shows the number of time-stamped pages (active unbound) on four mass storage volumes (A, B, C, and D) grouped in each of the five slice ranges. Slice range I contains the pages with the five oldest time-stamps; slice range II contains the pages with the next four oldest time-stamps; and so on. The LRUGROUP parameter defines 14 of the 16 possible time-stamps. The two newest time-stamps (E and F) are in the implied slice rang The LRU algorithm uses the slice ranges to select the mounted mass storage volumes that have the largest number of oldest pages for deallocation, beginning with slice range I.

Slice Ranges	I					II				III			IV		V*	
LRUGROUP=5,4,3,2	**5**					**4**				**3**			**2**			
Time Stamps (Normalized)	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
Mass Storage Volume A (17 active unbound pages)	4	0	0	2	3	4	1	2	1	0	0	0	0	0	0	0
Mass Storage Volume B (25 active unbound pages)	0	1	3	3	4	0	2	0	0	0	0	0	3	4	2	3
Mass Storage Volume C (14 active unbound pages)	0	0	0	0	0	1	0	0	0	0	0	0	3	5	2	3
Mass Storage Volume D (24 active unbound pages)	0	4	1	3	2	2	1	0	0	1	1	1	0	1	4	3
Time Stamps (unnormalized)	7	8	9	A	B	C	D	E	F	0	1	2	3	4	5	6

*Implied time slice V. (Use not recommended.)

Figure 37. LRU Volume Selection

Assume that the LRU algorithm is called at time-stamp time 6. The unnormalized time stamps are shown in the bottom row of Figure 37. At time-stamp time 6, the pages time-stamped 6 are current, and the pages time-stamped 7 are the oldest. The time-stamps are normalized, as shown at the top of Figure 37, with the current pages time-stamped F, and the oldest pages 0.

Figures 37 and 38 describe the selection of a mass storage volume for LRU activity. Assume that staging drive group 05 (Figure 38) has four mass storage volumes (A, B, C, and D) mounted on 70 active pages when the Mass Storage Control is ordered to acquire space for a 115-cylinder data set. There are only 32 free and inactive pages remaining in staging drive group 05. The Space Allocation routine determines that 15 pages (120 cylinders) added to the active page count of 70 equals a new active page count of 85, which exceeds the upper threshold.

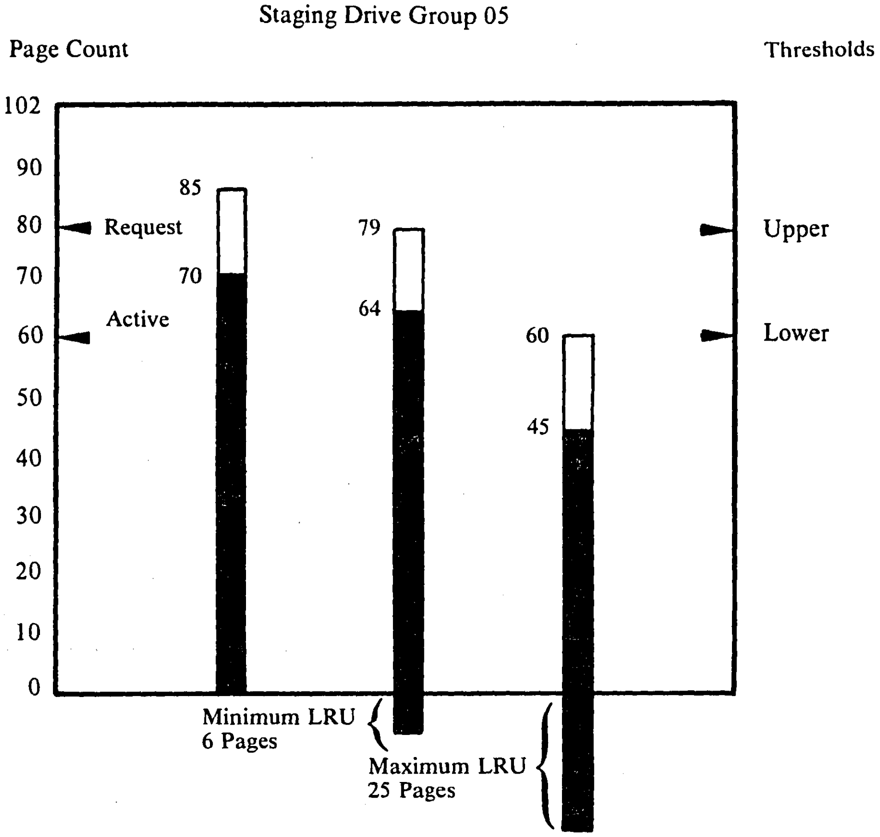

Figure 38. Minimum and Maximum LRUs—Example 1

The Space Allocation routine acquires the 15 pages and calls the LRU algorithm. The LRU algorithm normalizes the time-stamp values in staging drive group 05 and groups them by slice range, as shown in Figure 37.

The LRU algorithm does two things: lowers the new active page count (85) to a count below (79) the upper threshold (80), as shown in Figure 38, and searches slice ranges I-IV to find mass storage volumes with at least six active unbound pages (minimum LRU in this example). Lowers the new active page count (79) to the lower threshold (60), as shown in Figure 38, and searches slice ranges I-III to find mass storage volumes with at least 19 active unbound pages.

Slice range I in Figure 37 has 30 unbound active pages, ranging in age from normalized time stamp 0 to 4. For volume selection, all pages in a slice range are considered as having the same age. The mass storage volumes with the the largest number of oldest slice range I pages, in descending order are B (11 pages), D (10 pages), A (9 pages), C (0 pages). They are identified for deallocation in that order. Although only four of the nine pages on mass storage volume A are required for the maximum LRU of 25 pages (11 + 10 + 4), all nine are deallocated. Pages on cylinders that contain changed data must be destaged before they are available for new allocations.

To illustrate how slice ranges affect the selection of mass storage volumes, assume LRUGROUP=6,2,2,6, and review the volume selection in Figure 37. Slice range I now has 37 unbound pages ranging in age from normalized time stamp 0 to 5, and the mass storage volumes with the largest number of oldest slice range I pages, in descending order are A (13 pages), D (12 pages), B (11 pages), and C (1 page). In this variation, there are enough slice range I pages on volumes A and D to satisfy the maximum LRU of 25 pages (13 + 12).

In Figure 39 for example, the active page count (70) includes the active unbound (45) and active bound (25) pages. A space request is issued for a 320-cylinder data set (40 pages) that causes the new active page count (110) to exceed the remaining free, preferred, and inactive pages (32) available in staging drive group 05. The LRU algorithm is called to deallocate eight pages in slice ranges I to IV to satisfy the space request. These pages are queued for destaging (if required) and made available for allocation and subsequent staging.

The LRU algorithm does two other things: First, it deallocates sufficient active unbound pages in slice ranges I to IV to lower the active page count below the upper threshold. As shown in Figure 39, an additional 24 pages must be deallocated for a total of 31 pages.

Second, it lowers the new active page count to the lower threshold by deallocating active unbound pages in slice ranges I to III. But, because staging drive group 05 contains 25 active bound pages, this cannot be attained. However, an additional four pages can be deallocated for a total of 45 pages, as shown in Figure 39.

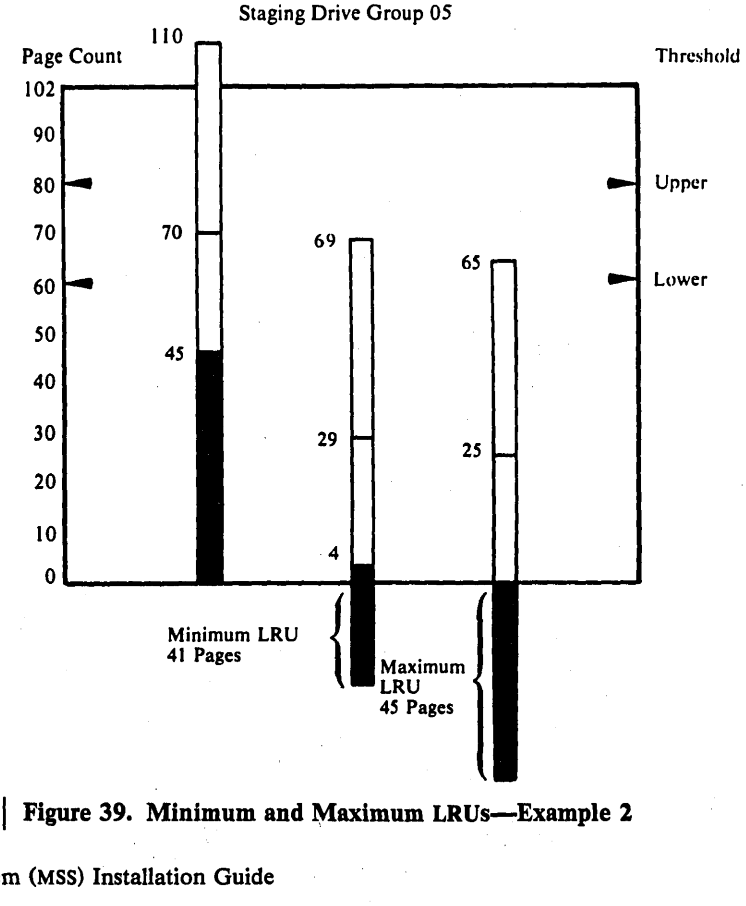

Figure 39. Minimum and Maximum LRUs—Example 2

If the eight pages are not available in staging drive group 05 for the new data set, the 32 available pages are allocated and acquired with the Inhibit Stage attribute (new data set). The 32 pages are marked as valid in the Staging Adapter tables and are available for use. When a new data set is created, a cylinder fault occurs if the CPU attempts to write a page that has not been acquired. The Mass Storage Control recognizes the cylinder fault and attempts a resolution. By this time, the pages in the staging drive group have been aged a minimum of one time stamp by the LRU algorithm or by the normal advancement of the Staging Adapter clock.

If the cylinder fault cannot be resolved, the job aborts; if space is available, a page is allocated. The Mass Storage Control treats this situation as a page fault but does not recognize that the original acquire order was for a new data set. When the page fault is resolved, the old data is staged even though it was not used.

The total staging drive capacity of a Mass Storage System configuration should be designed to 95% of the installation's peak requirement so that sufficient free and inactive pages can be normally available for allocation. If the LRU algorithm must search slice range IV to meet its initial objective, no attempt is made to meet subsequent objectives because all pages in slice ranges I-III have been deallocated. If no space can be found, the job aborts. Note that the pages in implied slice range V are not available to the LRU algorithm.

LRU Parameters

Three LRU parameters (ACTPAGE, LRUGROUP and LRUCLOCK) control the aging rate and number of active pages in each staging drive group. These parameters are initially specified to the Mass Storage Control Table Create program, but they can be subsequently adjusted using the Tune order if the analysis of Trace activity indicates adjustment is necessary.

ACTPAGE

ACTPAGE specifies the lower and upper deltas for each staging drive group. The lower delta can range from 0 to 50, and the upper delta from 1 to 51, but the upper delta must always be a higher number than the lower delta. We recommend a lower page delta of 30 and an upper page delta of 40 per logical 3330-1 as reasonable initial values.

The Mass Storage Control uses the page deltas to calculate the actual lower and upper thresholds of each staging drive group. These actual thresholds change as staging drives are varied online or offline, or as mass storage volumes are bound or unbound. The Mass Storage Control considers the 3330-11 staging drive equal to two logical 3330-1s.

Unless UNIT=SDGxx is used to direct the mounting of volumes, the same initial ACTPAGE values can be used for all staging drive groups. The Mass Storage System Communicator attempts to balance the workload across all staging drive groups by selecting the one that has the most space and least number of volume mounts.

LRUGROUP

LRUGROUP specifies the grouping of the normalized time-stamped pages in slice ranges I-IV and implied slice range V, as described earlier. We do not recommend the use of the implied slice range V, as shown in Figure 37, because it keeps highly active pages and newly staged pages from being selected for deallocation (and destaging). As a result, jobs can abend because the LRU algorithm is unable to find sufficient pages for deallocation.

A large number of time stamps in slice range I increase the number of least-active pages available for LRU deallocation on each selected mass storage volume. We recommend LRUGROUP=6,2,2,6 as an initial value because it provides the LRU algorithm with:

1. Six time-stamps in slice range I, which potentially reduces the number of mass storage volumes that must be initially selected, as described previously.
2. Ten time-stamps in slice ranges I-III, which lowers the new active page count to the lower threshold. Note that this excludes the six newest and most-active pages in slice range IV. For example, if LRUCLOCK=256, these pages are excluded for approximately 25 minutes.
3. All 16 time-stamps in slice ranges I-IV, which lowers the new active page count below the upper threshold.

LRUCLOCK

The LRUCLOCK parameter values range from 128 seconds to 16,384 seconds. The chosen value determines the rate at which the Mass Storage Control clock advances all of the Staging Adapter clocks. Figure 40 shows the time in hours, minutes, and seconds for one time stamp and for 16 time-stamps (one complete cycle). A single LRUCLOCK value is selected for a Mass Storage System with the CREATE command. A reasonable initial value can be estimated by adding the average time (seconds) that data sets are in use to the desired number of seconds of retention time (1800), and dividing this sum by 16 (the number of LRU clock segments). The SMF summary data used to develop the Mass Storage System configuration can also be used to calculate the average times based on the average job step time, the average number of initiators, and the total SMF analysis time. You may decide to choose LRU parameters based on a subset of this SMF data, which represent a particular shift or time period.

The LRU algorithm's job is to select the oldest pages (those not recently used) for deallocation (and destaging). Continuously used pages do not age because they are time-stamped with the current Staging Adapter clock value each time they are accessed. If pages are aged too fast, those at the end of a sequential file can become eligible for destaging before they are used; however, pages that age too slowly may not get destaged when they should. When the LRU clock advances from 'F' to '0', an old page cannot be distinguished from a new one.

Pages become inactive and available for allocation when mass storage volumes are demounted. The LRU algorithm is intended as a safety valve to be called only when unusual peak conditions occur. We recommend an initial parameter LRUCLOCK value of 256, 512, or 1024.

	One Time-Stamp			Sixteen Time-Stamps		
LRUCLOCK=	**Hours**	**Minutes**	**Seconds**	**Hours**	**Minutes**	**Seconds**
128	0	2	8	0	34	8
256	0	4	16	1	8	16
512	0	8	32	2	16	32
1024	0	17	4	4	33	4
2048	0	34	8	9	6	8
4096	1	8	16	18	12	16
8192	2	16	32	36	24	32
16384	4	33	4	72	49	4

Figure 40. LRUCLOCK Values and LRU Time-Stamp Intervals

Mass Storage Control Table Create Program Worksheets Description

To execute the Mass Storage Control Table Create program involves:

- Planning the logical configuration
- Coding the input parameters that describe the logical configuration unit addresses and the staging drive groups
- Comparing the logical connections to the physical configuration
- Reviewing the punched and printed Mass Storage Control Table Create program output and separating the punched output by CPUs
- Creating an IODEVICE card for each Mass Storage Control and modifying IODEVICE cards for parameters ALTCTRL, SHARED, SHAREDUP, and OFFLINE=YES
- Removing, modifying, and adding the IODEVICE cards created for direct access to the Mass Storage Control tables
- Integrating all IODEVICE cards as input to stage 1 of a SYSGEN for each System Control Program

OS/VS Mass Storage Control Table Create must be used in conjunction with this chapter. It describes the Mass Storage Control Table Create program and its commands, which are in alphabetic order. However, the worksheets in this publication contain the commands and are numbered in the order of use.

Mass Storage Control Table Create program tasks are done in this general sequence:

1. Lay out of the Mass Storage System components to physically and logically "create" the staging drive groups (Worksheet 1)
2. Planning the real unit addresses for the real and convertible-to-real drives (Worksheet 4)
3. Planning the virtual unit addresses for the virtual drives associated with the staging drive groups (Worksheet 4)
4. Connecting the Mass Storage System components (Worksheet 1)
5. Writing the input specifications for the Mass Storage Control Table Create program (Worksheets 2, 3, 5, 6, and 7)

Appendix A contains seven copyrighted Mass Storage Control Table Create Program Worksheets, which you can use as a prototype for your input or for planning and documenting your Mass Storage System configuration. We suggest that you keep the worksheets as masters and make copies as required for input and/or documentation purposes.

Planning a Minimum System Configuration

We use completed Mass Storage Control Table Create program worksheets and Mass Storage Control Table Create program reports to describe how to plan a logical configuration. In this example, the minimum system possible, which we shall call Configuration A, is installed at a branch location of the fictitious XYZ company and consists of:

- One System/370 Model 158

- One 3851 Mass Storage Facility Model A1
- One 3830 Model 3 Storage Control (Staging Adapter)
- Two 3333-11 Disk Storage and Controls (4 drives)
- One 3330-11 Disk Storage Unit (2 drives)

Worksheet 1, Logical Configuration

Worksheet 1, in Appendix A, contains representations of four CPUs, six Staging Adapters, six 3333 Disk Storage and Controls (12 drives), eighteen 3330 Disk Storage units (36 drives), and one Mass Storage Facility Model 4. The components are identified using Mass Storage Control Table Create program notation except for Model, Drive, and miscellaneous labels. The range of values, either alphabetic, binary, decimal, or hexadecimal, is included (in parentheses) for most parameters and subparameters.

Planning your logical configuration begins by identifying the Mass Storage System physical components. For configurations with two Mass Storage Facilities, or those with a large DASD configuration, two copies of Worksheet 1 may be needed. No connections are planned at this time.

The 3830/3333/3330 components provide the physical base on which the logical components are "created". To establish each staging drive group, the physical and logical parameters are planned concurrently. Several factors influence the size, number, and connections of each staging drive group. A major factor is the total number and proportion of 3830s, 3333s, and 3330s.

If serviceability is the primary configuration requirement, the 3333s are string-switched to Staging Adapter pairs, which provide two stage and destage paths but reduces the maximum number (28) of staging drive groups. Consider this method of connection for the Mass Storage Control tables in staging drive group 00.

If staging drive capacity is the primary configuration requirement, the 3333s are connected to each Staging Adapter without string-switching, which provides only one stage and destage path but allows the maximum number of staging drive groups.

Form trial staging drive groups and then adjust them to meet all your configuration requirements. Each Staging Adapter or Staging Adapter pair manages one or two logical groups of staging drives, that is, LCU0 and LCU1. Each logical group contains as many as eight logical 3330-1 drives, each with a logical unit address, that is, LUA0 through LUA7.

Each staging drive group pair that you form resides on 8 to 16 physical 3330 drives, depending on the mix of 3330-1s and 3330-11s. If the physical drive is a 3330-11, it contains two logical 3330-1s and is identified by even logical unit addresses, that is, LUA0, LUA2, LUA4, and so on.

Staging drive group 00 has special requirements because it contains the Mass Storage Control tables and their alternate areas. It must have a minimum of two 3333 Disk Storage and Controls (four drives) connected to one Staging Adapter (SA=0), or be string-switched to a Staging Adapter pair (SA=0 and SA=1). The four drives must be defined as staging. The Mass Storage Control Table Create program does not accept a CONVREAL specification for these four drives.

REAL (R) 3330-1, 3330-11, or 3350 drives can also be connected to the Staging Adapters, but are not part of a staging drive group. Shared drives on the logical configuration are identified as REAL shared (RS) or CONVREAL shared (CS) to flag them for SYSGEN address planning, that is, FEATURE=(SHARED) on IODEVICE card.

Figure 41 shows Configuration A on Worksheet 1. After your staging drive groups have been identified, the physical (DEVn, Model, and Drive) and logical (SDGxx, LCUz, and LUAy) parameters are entered vertically for each physical drive used, as shown in the lower left-hand corner example on Worksheet 1.

The special requirements of staging drive group 00 have determined the logical parameters for the four drives on which the Mass Storage Control tables reside, which must be specified as staging. The two real drives are planned for system use, that is, system residence, paging, MSV Inventory and MSVC Journal data sets.

Connections III and IV are simple connections. Note that both connectors on DRC0 have been used and are uniquely addressed.

MSC Table Create Program Worksheet 1

Worksheet 1, Part A

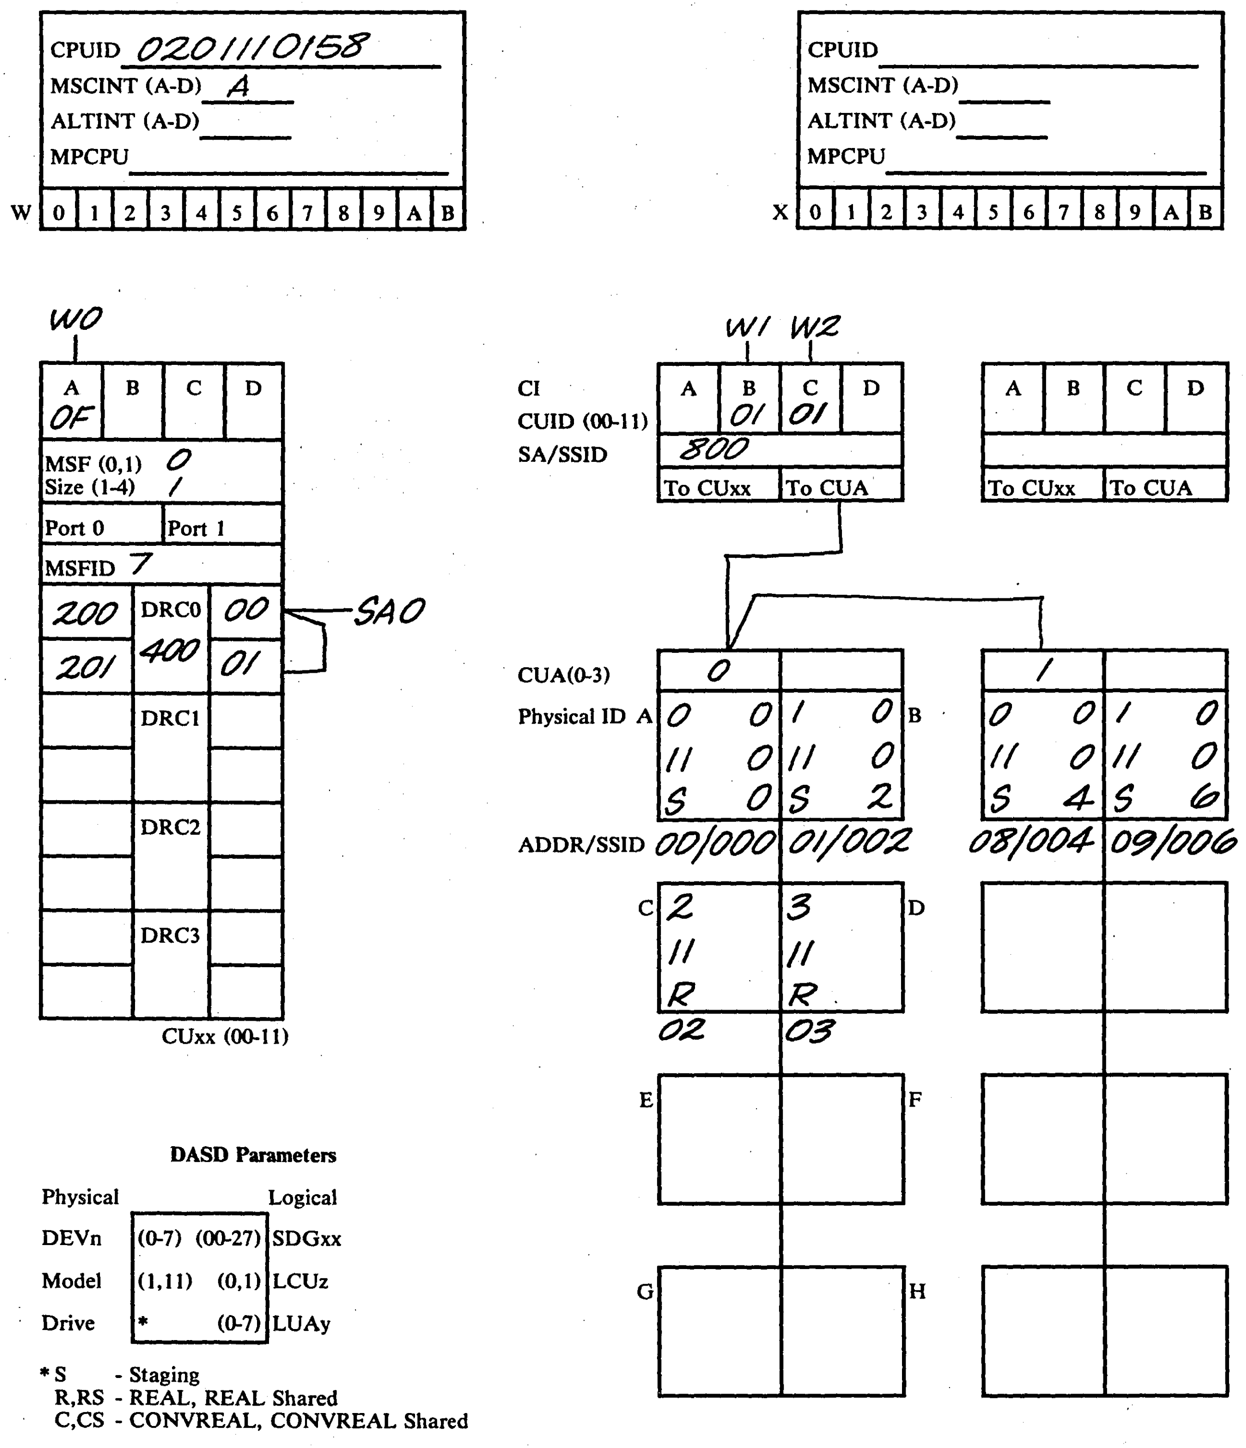

CPUID 0201110158
MSCINT (A-D) 4
ALTINT (A-D)
MPCPU

W	0	1	2	3	4	5	6	7	8	9	A	B

CPUID
MSCINT (A-D)
ALTINT (A-D)
MPCPU

X	0	1	2	3	4	5	6	7	8	9	A	B

W0

A	B	C	D
0F			

MSF (0,1) 0
Size (1-4) 1
Port 0 Port 1
MSFID 7

200	DRC0	00
201	400	01
	DRC1	
	DRC2	
	DRC3	

SA0

CUxx (00-11)

W1 W2

	A	B	C	D		A	B	C	D
CI									
CUID (00-11)		01	01						
SA/SSID	800								
	To CUxx		To CUA			To CUxx		To CUA	

CUA(0-3)	0			1	
Physical ID A	0 0 11 0 S 0	1 0 11 0 S 2	B	0 0 11 0 S 4	1 0 11 0 S 6
ADDR/SSID	00/000	01/002		08/004	09/006
C	2 11 R	3 11 R	D		
	02	03			
E			F		
G			H		

DASD Parameters

Physical			Logical
DEVn	(0-7)	(00-27)	SDGxx
Model	(1,11)	(0,1)	LCUz
Drive	*	(0-7)	LUAy

*S - Staging
R,RS - REAL, REAL Shared
C,CS - CONVREAL, CONVREAL Shared

Figure 41. Sample of Worksheet 1

Worksheet 2, DASD Command

The physical and logical parameters indicated on Worksheet 1 for each physical drive are the direct input to the DASD command. Make sure you delete unused parameters to conform to the syntax described in *OS/VS Mass Storage Control Table Create*.

Worksheet 2 in Figure 42 has been completed for Configuration A. Note that real drives DEV2 and DEV3 are not specified as staging drive group members. If the drives had been specified as CONVREAL, they could be members of staging drive group 01 (LCU1, LUA0, and LUA1). Note that the 3333 is specified as a 3330 device.

The DASD command establishes Connection III and the Staging Adapter portion of Connection II.

Worksheet 3, LOWERCON Command

The LOWERCON command establishes the Staging Adapter lower port connections (to CUxx) to the Data Recording Controls (DRC or CUxx) in the Mass Storage Facility. This is the Connection IV that was planned on Worksheet 1.

Worksheet 3 in Figure 42 shows the single LOWERCON command required for Configuration A.

MSC Table Create Program Worksheet 2

DASD: Specifying the Drives Attached to a DASD Controller (3333 or 3350).

Required: One DASD command per 3333 or 3350 (CUA). Max = 32 commands.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create* for SA and SSID.
DEVn = 3330-1, 3330-11, or 3350.
Delete: [SIDD=______,LCU,______,LUA,______,] if function is real REAL, or DEVn = 3350. Delete [REAL,] if function is staging (S) or CONVREAL (C).

DASD SA= 0 ┌ ,CUA= 0 ,
DEV0= 3330-11 =SSID=SDG 0 ,LCU 0 ,LUA 0 ,REAL,
DEV1= 3330-11 =SSID=SDG 0 ,LCU 0 ,LUA 2 ,REAL,
DEV2= 3330-11 ~~=SSID=SDG______,LCU______,LUA______~~,REAL,
DEV3= 3330-11 ~~=SSID=SDG______,LCU______,LUA______~~,REAL┌
DEV4=______=SSID=SDG______,LCU______,LUA______,REAL,
DEV5=______=SSID=SDG______,LCU______,LUA______,REAL,
DEV6=______=SSID=SDG______,LCU______,LUA______,REAL,
DEV7=______=SSID=SDG______,LCU______,LUA______,REAL

DASD SA= 0 ┌ ,CUA= 1 ,
DEV0= 3330-11 =SSID=SDG 0 ,LCU 0 ,LUA 4 ~~,REAL~~
DEV1= 3330-11 =SSID=SDG 0 ,LCU 0 ,LUA 6 ~~,REAL~~
DEV2=______=SSID=SDG______,LCU______,LUA______,REAL,

MSC Table Create Program Worksheet 3

LOWERCON: Describing the Connection Between a Staging Adapter and a Data Recording Control.

Required: One LOWERCON command per Staging Adapter (3830-3 or ISCSA). Max = 14 commands.

Restrictions: DRC value must be in range 0-3 depending on 3851 size.

LOWERCON SA= 0 ,
CU00=MSF 0 ,DRC 0 ,
CU01=MSF 0 ,DRC 0 ┌
CU10=MSF______,DRC______,
CU11=MSF______,DRC______

LOWERCON SA=______,
CU00=MSF______,DRC______,
CU01=MSF______,DRC______,
CU10=MSF______,DRC______,
CU11=MSF______,DRC______

Figure 42. Sample of Worksheets 2 and 3

Worksheet 4, Planning Device Addresses

A device address must be specified to the Mass Storage Control Table Create program for each REAL, CONVREAL, and staging drive that is available to a CPU. Worksheet 4, as shown in Appendix A, has been designed so that you can plan unit addresses (12-bit format) and translate them directly to device addresses (6-bit format). The number of virtual and virtual shared unit addresses specified for a CPU determines how many virtual drives it has available on which to mount virtual volumes. The distribution of these unit addresses to staging drive groups can affect the planning and use of your staging drive groups, especially if you plan to use esoteric names so that units can direct staging of virtual volumes to specific staging drive groups.

Each Staging Adapter channel interface (B, C, and D) can recognize up to 64 unit addresses. For example, the 64 unit addresses recognized by Staging Adapter 2, channel interface (CI) C with control unit address (CUID) 10 or channel (CH) 4 are 480-4BF. These are specified to the Mass Storage Control Table Create program as: SA=2, CI=C, CUID=10, CH=4,,, ADDR=00-3F. Specifying all of the available 64 unit addresses is not practical because 3333s string-switched to Staging Adapter pairs provide up to 384 virtual units (64 addresses on 6 unique channel interfaces) in one staging drive group. You can analyze your requirements to determine the maximum number of virtual volumes that need to mounted using this formula:

1. The number of active initiators multiplied by the estimated average number of data sets per job step equals the total number of open data sets.
2. The total number of open data sets divided by the estimated number of open data sets per mass storage volume equals the number of virtual units (or unit control blocks) required for a CPU.

For efficient use, the two address tables in Worksheet 4 should be made to logically represent two Staging Adapter channel interfaces logically connected to a staging drive group pair from primary and alternate channel paths from a CPU. Although these two channel interfaces may not be physically adjacent in the same Staging Adapter, placing them logically adjacent makes it easier to plan alternate unit addresses. Space is provided on the worksheet to indicate the use of all addresses and the staging drive group of the virtual and virtual shared unit addresses. Note the planned use of both virtual and real unit addresses, because you must add certain IODEVICE card parameters for real unit addresses and possibly modify those cards for virtual unit addresses.

As noted on Worksheet 4, device addresses 00, 01, 08, and 09 on channel interface B of Staging Adapters 0 and 1 are reserved for direct access to the Mass Storage Control tables. The Mass Storage Control Table Create program automatically punches out IODEVICE cards for these unit addresses. Channel interfaces C and D of Staging Adapters 0 and 1 can also be used as direct access paths to the Mass Storage Control tables. However, you must advise the Customer Engineer so that he can plug the device addresses on Staging Adapter 0 and 1 for direct access. You must reserve device addresses 00, 01, 08, and 09 as real units on channel interfaces C and D if you choose this option, because the Mass Storage Control Table Create program does not punch the IODEVICE cards. See

"Direct Access to Mass Storage Control Tables" in the chapter "Data Security and Recovery Planning".

Unit addresses for real and convertible-to-real drives should be specified first because they represent the electrical paths to the physical devices. Note that the first 32 unit addresses on each Staging Adapter channel interface (CUID) are real. Every drive, whether specified as REAL, CONVREAL, or staging should have its electrical path represented by a unit address, either real or virtual. This ensures that the Online Test Executive Program (OLTEP) can address every drive using the real unit address, or resetting the virtual bit in the virtual UCB and using that unit address as real. The Mass Storage Control Table Create Host Connections Report prints ***** for each drive that is not represented by a real or virtual unit address. At least one electrical path must be represented in each DASD string. This makes it necessary to move the logical address plug (and the Customer Engineer's disk pack) within the DASD string to run diagnostic routines on each drive. Moving pack and plug might not be practical in a busy environment.

The remaining unit addresses are available for specification as virtual or virtual shared addresses. Depending on the number of unit control words available on a 2880 Block Multiplexer Channel, you may need to use consecutive unit addresses beginning with 00. See 2880 Block Multiplexer Channel and Unit Control Word Examples in the chapter "Related Logical and Physical Planning" for details.

At least one worksheet is required for every two Staging Adapter interfaces. Connection V, between the CPU block multiplexer channels and the Staging Adapter channel interfaces, is determined as you plan the unit addresses. These connections are drawn on Worksheet 1.

Sixty-four unit addresses are available on Staging Adapter interface B. Refer to Figure 43 and plan the unit addresses needed for Configuration A. With channel interface address 01, the unit address range is 140-17F. Select the real unit addresses first. Real unit addresses 140, 141, 148, and 149 (device addresses 00, 01, 08, and 09) are reserved for direct access to the Mass Storage Control tables. The two real drives in the first DASD string require real unit addresses 142 and 143 (device addresses 02 and 03).

MSC Table Create Program Worksheet 4

Addresses: Planning Device Addresses (00-3F) for UPPERCON and IODEVICE input.

Required: One worksheet per two Staging Adapter channel interfaces.

*** Distribution of Device Addresses 00–FF**

CUID	00	01	10	11
Real or Virtual	00-0F	40-4F	80-8F	C0-CF
	10-1F	50-5F	90-9F	D0-DF
Virtual	20-2F	60-6F	A0-AF	E0-EF
	30-3F	70-7F	B0-BF	F0-FF

Drive:
V, VS - Virtual, Virtual Shared
R, RS - REAL, REAL Shared
C, CS - CONVREAL, CONVREAL Shared
A - Alternate Path

Unit Address Formats

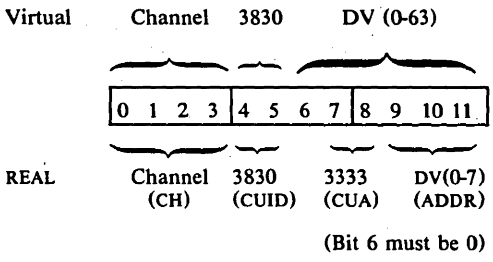

Note: SA-0, SA-1 ADDR 00, 01, 08, and 09 for direct access to MSC tables:

- Reserved on CI-B
- User option on CI-C and D

CPUID 020 111 0158

SA 0 CI B
CUID 01 CH 1

Address		Drive	SDGxx
*Unit	Dv		
140	00	MSC	
	01		
	02	R	
	03		
	04	V	00
	05		
	06		
	07		
	08	MSC	
	09		
	0A	V	00
	0B		
	0C		
	0D		
	0E		
	0F		
150	10		
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
160	20		
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

CPUID

SA 0 CI C
CUID 01 CH 2

Address		Drive	SDGxx
*Unit	Dv		
240	00	MSC	
	01		
	02	A	
	03		
	04	A	00
	05		
	06		
	07		
	08	MSC	
	09		
	0A	A	00
	0B		
	0C		
	0D		
	0E		
	0F		
250	10		
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
260	20		
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

Figure 43. Sample of Worksheet 4

The maximum number of mounted virtual volumes required for Configuration A is anticipated to be 36. Unit address ranges 144-147 and 14A-169 are selected as the 36 virtual unit addresses.

Channel interface C on channel 2 is planned as the alternate path to the real and virtual drives. The corresponding unit addresses on channel 2 (242-243, 244-247, and 24A-269) are to be selected as alternates.

Unit addresses 142-143, 144-147, and 14A-169 translate to device addresses 02-03, 04-07, and 0A-29 on Staging Adapter channel interface address (CUID=01) and channel 2. The virtual device addresses are to be specified in staging drive group 00.

The remaining 22 unit addresses 16A-17F are unused. Because they are within the unit address range recognized by CUID=01, they should not be used for any other device on channel 1.

Worksheet 5, UPPERCON Command

Worksheet 5 contains two UPPERCON commands. One specifies the primary channel path, and the other (optional) specifies the alternate channel path. If the ALTHPATH command is specified for the virtual, virtual shared, CONVREAL, or REAL devices, the second UPPERCON command must be used. The Mass Storage Control Table Create program expects an UPPERCON command for each Staging Adapter channel interface used in the configuration. Although CONVREAL and REAL drives can be specified to the program, they cannot be specified as shared. The FEATURE=(SHARED) parameter must be added to the IODEVICE cards and punched by the program.

Figure 44 shows Worksheet 5 for Configuration A. Because ALTPATH is a subparameter, it must be specified for each device type. The parameters and subparamters required for Worksheet 5 can be entered directly from or readily determined from Worksheet 4. Worksheet 5 establishes Connection V for up to two Staging Adapter channel interfaces.

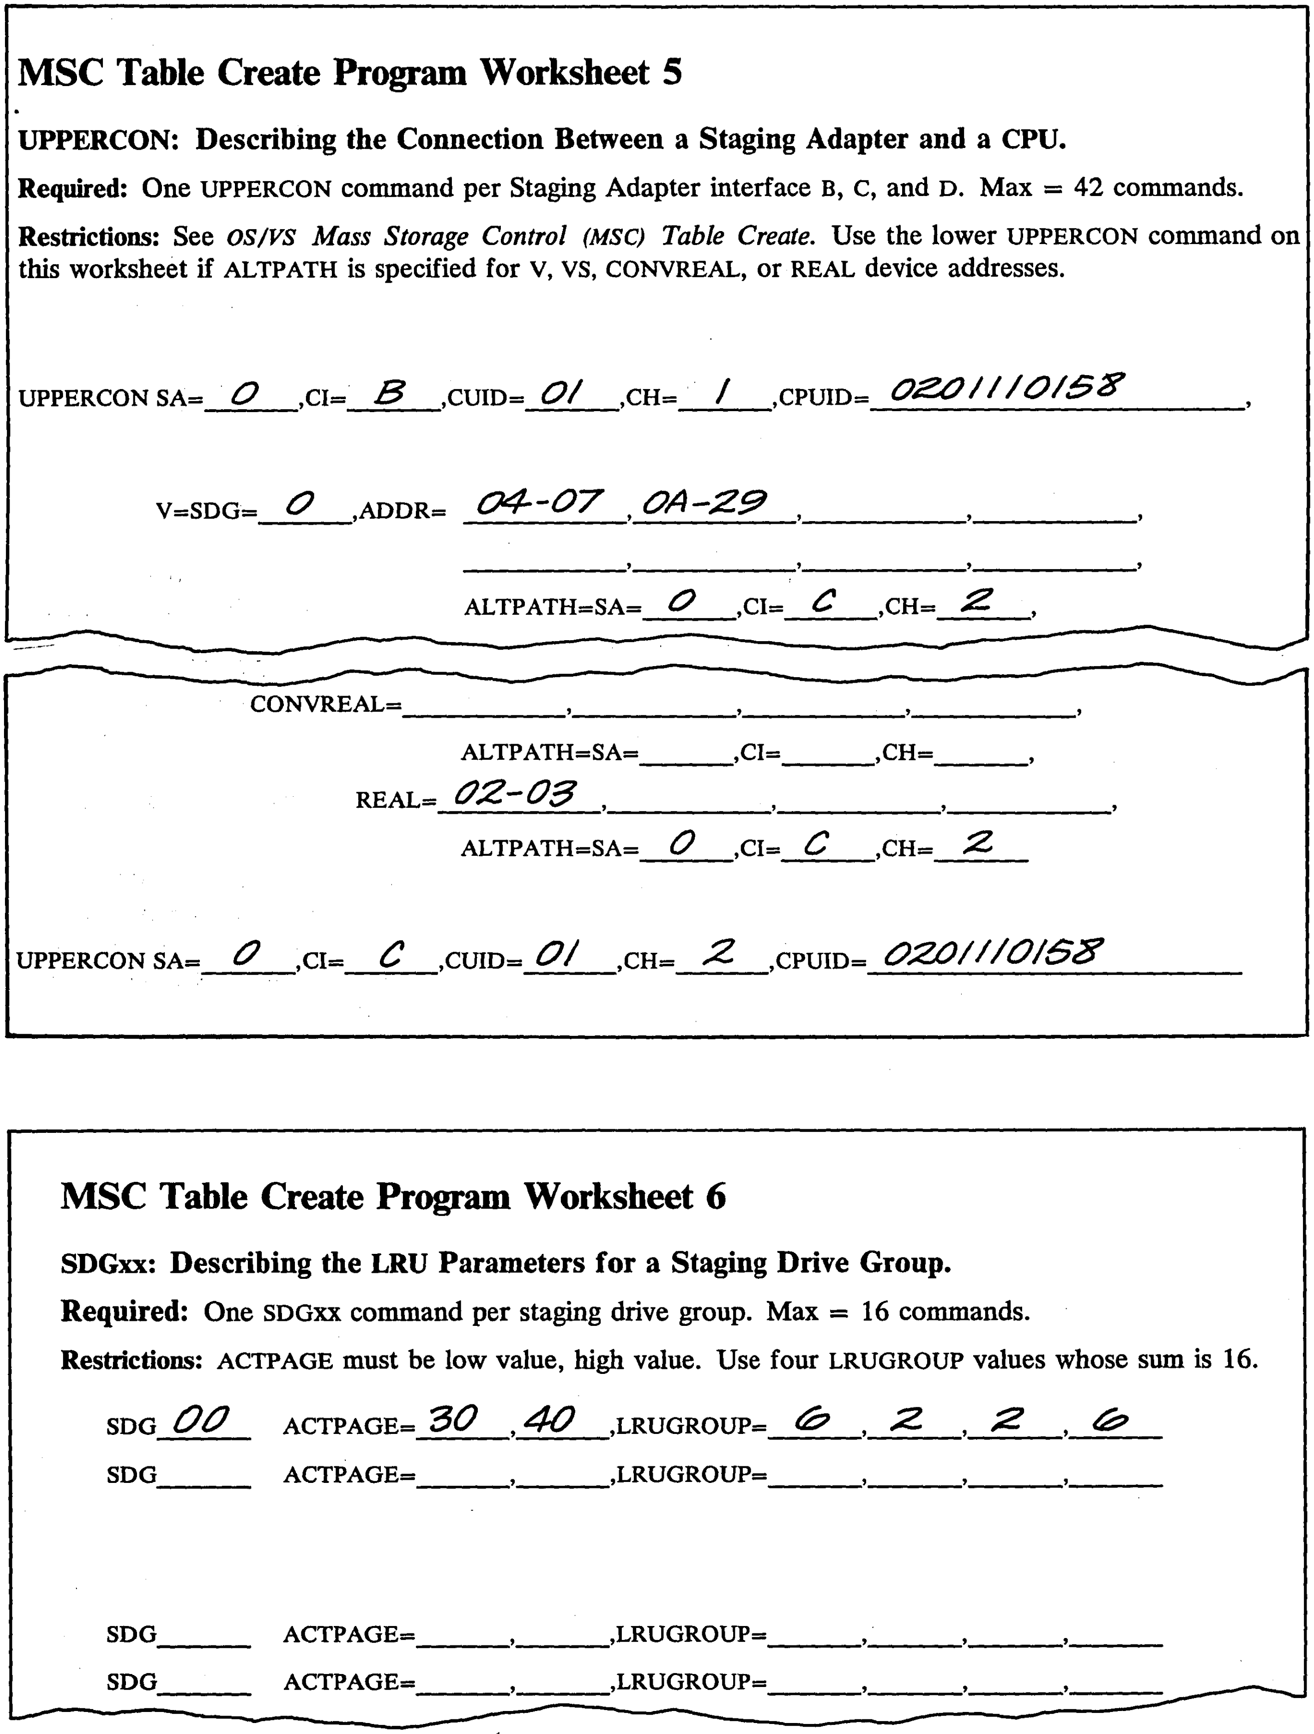

MSC Table Create Program Worksheet 5

UPPERCON: Describing the Connection Between a Staging Adapter and a CPU.

Required: One UPPERCON command per Staging Adapter interface B, C, and D. Max = 42 commands.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create.* Use the lower UPPERCON command on this worksheet if ALTPATH is specified for V, VS, CONVREAL, or REAL device addresses.

UPPERCON SA= 0 ,CI= B ,CUID= 01 ,CH= 1 ,CPUID= 0201110158 ,

V=SDG= 0 ,ADDR= 04-07 , 0A-29 ,______,______,

______,______,______,______,

ALTPATH=SA= 0 ,CI= C ,CH= 2 ,

CONVREAL=______,______,______,______,

ALTPATH=SA=______,CI=______,CH=______,

REAL= 02-03 ,______,______,______,

ALTPATH=SA= 0 ,CI= C ,CH= 2

UPPERCON SA= 0 ,CI= C ,CUID= 01 ,CH= 2 ,CPUID= 0201110158

MSC Table Create Program Worksheet 6

SDGxx: Describing the LRU Parameters for a Staging Drive Group.

Required: One SDGxx command per staging drive group. Max = 16 commands.

Restrictions: ACTPAGE must be low value, high value. Use four LRUGROUP values whose sum is 16.

SDG 00 ACTPAGE= 30 , 40 ,LRUGROUP= 6 , 2 , 2 , 6

SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

Figure 44. Sample of Worksheets 5 and 6

Worksheet 6, SDGxx Command

Three LRU parameters are specified to the Mass Storage Control Table Create program: ACTPAGE, LRUGROUP, and LRUCLOCK. The upper and lower page deltas (ACTPAGE) and page grouping by age (LRUGROUP) can be individually specified for each staging drive group. The three LRU parameters and the internal management of staging drive group space are described in the chapter "Staging Drive Group Space". LRUCLOCK is specified in the CREATE command. Worksheet 6 for Configuration A is in Figure 44. All four LRUGROUP values (6, 2, 2, 6) are specified and their sum is equal to 16. This ensures that all pages (slice ranges I through IV) in the staging drive group are available to the LRU algorithm for deallocation. If a sum less than 16 is used, pages in the implied slice range V are not available. As a result, jobs may abend because space cannot be found in the selected staging drive group.

Worksheet 6 contains 14 SDGxx commands. ACTPAGE and LRUGROUP are specified in one SDGxx command.

Worksheet 7, CPUCONFG, MSFn, and CREATE Commands

CPUCONFG Command

Connection I between each CPU and the Mass Storage Controls is established by the CPUCONFG command. The input comes from Worksheet 1.

The CPUID and MPCPU parameters determine whether the connecting CPU is a logical uniprocessor (UP) or multiprocessor (MP). An alternate channel path can be established for a uniprocessor using the ALTINT parameter. Note that the MPCPU and ALTINT parameters are mutually exclusive.

The OPSYS parameter (OS/VS1, OS/VS2, JES3 or SVS) specifies the System Control Program for each CPU or MPCPU and is used by the Mass Storage Control Table Create program to check alternate channel and multiprocessor connections. If the parameter is omitted, the default is OS/VS2 with MPCPU specified, or the System Control Program on which the Mass Storage Control Table Create program is run. Worksheet 7 contains the four possible CPUCONFG commands. Figure 45 shows the CPUCONFG command for Configuration A.

MSC Table Create Program Worksheet 7

CPUCONFG: Describing the CPU Configuration.

Required: One CPUCONFG command per attached CPU. Max = 4 commands.

Restrictions: MPCPU and ALTINT are mutually exclusive. OPSYS=VS1, VS2, SVS, or JES3.
Delete: [,MPCPU], [,ALTINT], or [OPSYS] if not required.

CPUCONFG CPUID=0201110158 ,MSCINT= A ~~,ALTINT=~~ ,OPSYS= VS1
~~MPCPU=~~

CPUCONFG CPUID=________,MSCINT=______,ALTINT=______,OPSYS=______,
MPCPU=________

CPUCONFG CPUID=________,MSCINT=______,ALTINT=______,OPSYS=______,
MPCPU=________

CPUCONFG CPUID=________,MSCINT=______,ALTINT=______,OPSYS=______,
MPCPU=________

MSFn: Describing the Mass Storage Facility

Required: One MSFn command per 3851 MSF. Max = 2.

Restrictions: See OS/VS *Mass Storage Table Create* for MSFn and MSFID.
Delete: [,LOADMAP] if option not required.

MSF0 SIZE= 1 ,MSFID=7~~,LOADMAP=~~
MSF1 SIZE=______,MSFID=6,LOADMAP=________

CREATE: Describing the Type of Table Create Execution

Required: One CREATE command per execution.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create* for VOLID.
Delete: [NEW, or UPDATE,], [PRNTONLY], and [CYLFAULT=] if not required.

CREATE NEW,~~UPDATE,~~PRNTONLY,
LRUCLOCK= 256 ,VOLID= XYZA 00,
CYLFAULT= YES

Figure 45. Sample of Worksheet 7

MSFn Command

The MSFn command specifies the size of each Mass Storage Facility and the logical address (MSFID) of the accessor control pair in each Mass Storage Facility. The accessor control pair is specified on Worksheet 7 as MSFID=7 in MSF0, and as MSFID=6 in MSF1. If you prefer, MSFID can be any single digit from 2-7. Make sure that a Staging Adapter and accessor control pair are not identified with the same digit. MSFID determines a portion of Connection II from the Mass Storage Control port.

If you prefer to write your own cell-loading program, you must provide the MSFn command with the program name (exit-name) as an optional parameter. See *OS/VS Mass Storage Control Table Create* for further information.

Worksheet 7 contains the two MSFn commands. Figure 45 shows the MSFn command for Configuration A.

CREATE Command

The CREATE command establishes this Mass Storage Control Table Create program generation as NEW or UPDATE. If UPDATE is specified, all mass storage volumes must be demounted prior to the program execution and the existing primary Mass Storage Control tables pack must be mounted on a real drive.

LRUCLOCK is specified in the CREATE command. The chapter, "Staging Drive Group Space", describes its use and recommends initial values.

The VOLID of the primary Mass Storage Control table pack is specified in the CREATE command.

The PRNTONLY parameter is useful during your first execution of the program, because it allows you to verify the printed output before the Mass Storage Control tables are written. If a SYSPUNCH DD card is included in the OS/VS JCL stream, the program punches IODEVICE and UNITNAME cards.

The CYLFAULT parameter specifies whether channel command retry (CYLFAULT=NO, the default parameter) or cylinder fault attention interrupt (CYLFAULT=YES) should be used.

Figure 45 shows the CREATE command for Configuration A.

Program Printouts

Worksheets 2, 3, 5, 6, and 7 are the source of input to the Mass Storage Control Table Create commands. The output is easier to use if you list the input commands in alphabetic order. The printed output consists of:

- Input/output listings
- Host Connections Report
- Connections Report

Input/Output Listings

Figure 46 shows the Mass Storage Control Table Create input/output listing for Configuration A. Scan the input statements to be sure they are correct. Read the warning messages and make sure you understand them.

```
                    MSC TABLE CREATE  (VS1)
ICG103I  INITIALIZE PHASE COMPLETED
ICG109I  0000 WARNING ERRORS, 0000 TERMINATING ERRORS
ICG110I  HIGHEST RETURN CODE WAS 00
                            CREATE NEW,PRNTONLY,
                                 LRUCLOCK=256,VOLID=XYZA00,CYLFAULT=YES
                            CPUCONFG CPUID=0201110158,MSCINT=A,OPSYS=VS1
                            DASD SA=0,CUA=0,
                                 DEV0=3330-11=SSID=SDG0,LCU0,LUA0,
                                 DEV1=3330-11=SSID=SDG0,LCU0,LUA2,
                                 DEV2=3330-11=REAL,
                                 DEV3=3330-11=REAL
                            DASD SA=0,CUA=1,
                                 DEV0=3330-11=SSID=SDG0,LCU0,LUA4,
                                 DEV1=3330-11=SSID=SDG0,LCU0,LUA6
                            LOWERCON SA=0,
                                 CU00=MSF0,DRC0,
                                 CU01=MSF0,DRC0
                            MSF0 SIZE=1,MSFID=7
                            SDG00 ACTPAGE=30,40,LRUGROUP=6,2,2,6
                            UPPERCON SA=0,CI=B,CUID=01,CH=1,CPUID=0201110158,
                                 V=SDG=00,ADDR=04-07,0A-29,
                                 ALTPATH=SA=0,CI=C,CH=2,
                                 REAL=02,03,
                                 ALTPATH=SA0=CI=C,CH=2
                            UPPERCON SA=0,CI=C,CUID=01,CH=2,CPUID=0201110158
ICG104I  SCAN PHASE COMPLETED
ICG109I  0000 WARNING ERRORS, 0000 TERMINATING ERRORS
ICG110I  HIGHEST RETURN CODE WAS 00
ICG334I  NO ALTERNATE SA PATH TO MSC TABLES
ICG322I  WARNING - SA 0 INTERFACE D NOT USED
ICG105I  ANALYZE PHASE COMPLETED
ICG109I  0002 WARNING ERRORS, 0000 TERMINATING ERRORS
ICG110I  HIGHEST RETURN CODE WAS 04
ICG401I  WARNING - NOT USING CURRENT LEVEL OF MSC TABLES AS INPUT
ICG402I  UNABLE TO OPEN SYSPUNCH. OUTPUT LISTED IN MESSAGE DATA SET
                   I2011101 IODEVICE ADDRESS=(140,2),UNIT=3330,MODEL=11
                   I2011102 IODEVICE ADDRESS=(142,2),UNIT=3330,OPTCHAN=2,MODEL=11
                   I2011103 IODEVICE ADDRESS=(144,4),UNIT=3330V,OPTCHAN=2
                   I2011104 IODEVICE ADDRESS=(148,2),UNIT=3330,MODEL=11
                   I2011105 IODEVICE ADDRESS=(14A,32),UNIT=3330V,OPTCHAN=2
                   U2011100 UNITNAME NAME=SDG00,UNIT=((144,4),(14A,32))
```

Figure 46. Input/Output Listing

For Configuration A, there is only one Staging Adapter and channel interface D is not used. Therefore, messages ICG334I and ICG322I are valid but of no concern. Message ICG401I occurs whenever the Mass Storage Control Table Create program is run with the NEW,PRNTONLY parameters.

The IODEVICE cards for Configuration A were printed but not punched. They are the input to your system generation (SYSGEN) and must be carefully checked. Note that the program automatically punched the IODEVICE cards (I2011101 and I201104) for direct access to the Mass Storage Control tables. If an alternate path to the Mass Storage Control tables is desired, you must inform the Customer Engineer so that the necessary changes can be made to the circuit card in Staging Adapter 0 for direct access to the Mass Storage Control tables. In addition, you must add the OPTCHAN=2 option to the IODEVICE cards. An IODEVICE card must be punched for the Mass Storage Control in the Mass Storage Facility Model A1. If device address 20 is selected for the Mass Storage Control, the format is:

```
IODEVICE ADDRESS=(020),UNIT=3851
```

Host Connections Report

Figure 47 shows the Host Connections Report for Configuration A. The report verifies the virtual and real device addresses that map to a specific physical drive from each CPU connected to a Staging Adapter channel interface.

For Configuration A, the addresses of the two 3333 Disk Storage and Controls (DASD Controller on the report) are 00 and 01 (CUA input to the

DASD command). No addresses are ever shown on the Host Connections Report for the Mass Storage Control tables packs (SSIDs 000, 002, 004, and 006). Device plugs 2 and 3 under DASD controller 00 are the two real devices, but they are not members of a staging drive group and do not have an SSID. Device plugs 0 and 1 under DASD controller 01 are logical address plugs 08 and 09, as shown on Worksheet 1.

```
                    MSC TABLE CREATE (VS1)
HOST CONNECTIONS REPORT
DASD CONTROLLER 00
                            ******* STAGING ADAPTER 0  *******

    DEVICE   DEVICE  SSID     CPUID       CPUID       CPUID
     PLUG     TYPE          0201110158  0201110158    *****
                              OS/VS1      OS/VS1

      0      3330-11  000     *****       *****       *****
      1      3330-11  002     *****       *****       *****
      2      3330-11  ***     142-R       242-A       *****
      3      3330-11  ***     143-R       243-A       *****
DASD CONTROLLER 01
                            ******* STAGING ADAPTER 0  *******

    DEVICE   DEVICE  SSID     CPUID       CPUID       CPUID
     PLUG     TYPE          0201110158  0201110158    *****
                              OS/VS1      OS/VS1

      0      3330-11  004     *****       *****       *****
      1      3330-11  006     *****       *****       *****
```

Figure 47. Host Connections Report

Connections Report

Figure 48 shows the Connections Report for Configuration A. This report is used by the Customer Engineer as an aid in physically cabling and setting addresses in the hardware to match the specified configuration expected by the Mass Storage Control. It identifies the upper and lower connections to the Staging Adapters and the Mass Storage Facility.

```
                     MSC TABLE CREATE (VS1)
CONNECTIONS REPORT
CONNECTION FOR CARTRIDGE STORE  MSFO, SSID 100, MSFID = 7
CONNECTIONS FOR SA 0 SSID 800

   INTERFACE A TO MSC, PORT 0, CONTROL UNIT ADDRESS 000
   INTERFACE B TO CPU 0201110158, CHANNEL 1, CONTROL UNIT ADDRESS 01
   INTERFACE C TO CPU 0201110158, CHANNEL 2, CONTROL UNIT ADDRESS 01
   INTERFACE D IS NOT USED
     DASD  (CTL-I)
           CONTROLLER 00           CONTROLLER 01           CONTROLLER 10            CONTROLLER 11

           PLUG   TYPE    SSID     PLUG   TYPE    SSID     PLUG   TYPE    SSID      PLUG   TYPE    SSID
             0   3330-11   000       0   3330-11   004       0                        0
             1   3330-11   002       1   3330-11   006       1                        1
             2   3330-11   ***       2   *****     ***       2                        2
             3   3330-11   ***       3   *****     ***       3                        3
             4   *****     ***       4   *****     ***       4                        4
             5   *****     ***       5   *****     ***       5                        5
             6   *****     ***       6   *****     ***       6                        6
             7   *****     ***       7   *****     ***       7                        7
     DRC/DRD (CTL-I/ADT)
           DRD CONTROLLER ADDRESS 00 TO MSF 0, DRC 0,A
                 DRC 0    SSID 400
                 DRD 00   SSID 200
                 DRD 01   SSID 201
           DRD CONTROLLER ADDRESS 01 TO MSF 0, DRC 0,B
                 DRC 0    SSID 400
                 DRD 00   SSID 200
                 DRD 01   SSID 201
           DRD CONTROLLER ADDRESS 10 IS UNUSED
           DRD CONTROLLER ADDRESS 11 IS UNUSED
```

Figure 48. Connections Report

For Configuration A, Staging Adapter 0 channel interfaces A, B, C, and D are planned or not used as indicated in the Connections Report. The six DASD units are the four drives containing the Mass Storage Control tables, SSIDs 000, 002, 004, and 006, in staging drive group 00 and the two real drives. Plugs 0 and 1 under controller 01 correspond to logical address plugs 08 and 09 on Worksheet 1.

Connections from the Staging Adapter lower port to the data recording control and from the data recording control to the data recording devices are the last set of connections shown. The DRC0, A and DRC0, B connections refer to the two connectors on data recording control 0 (SSID 400). The connections from the data recording control to the data recording device. are shown as standard connections in Figure 18, and are not specified to the Mass Storage Control Table Create program.

Mass Storage Control Table Create Examples

This chapter describes the logical configuration input to the Mass Storage Control Table Create program for the XYZ Company. They have a System/370 Model 158 Multiprocessor (MP) and a System/370 Model 168 Attached Processor (AP), both running under OS/VS2 MVS. We describe two methods of connecting MP and AP systems to the Mass Storage Control, which are supported only by OS/VS2 MVS.

The XYZ Company selected the following system components after analyzing their current and projected data use and application requirements.

- Model 158MP system requirements (these components not shown on Worksheet 1 in Appendix B):

  Integrated Storage Control for controlling real drives

  One 3333-11 for system residence and paging

  One 3330-11 for the Inventory and Journal data sets to be shared with Model 168AP.

- Model 168AP system requirements (these components not shown on Worksheet 1 in Appendix B):

  Integrated Storage Control for controlling real drives

  One 3333-11 for system residence and paging

  One 3330-11 for application data sets

- Mass Storage System components (these components are shown on Worksheet 1 in Appendix B):

  Four 3830-3s for controlling staging and real drives

  One 3333-1 for the Mass Storage Control tables drives

  One 3333-11 for the Mass Storage Control alternate tables drives

  One 3333-11 for use as staging drives

  Four 3330-11s for use as staging and real drives

  Three 3330-1s for use as staging drives

  Two 3350s for use as real drives

Magnetic tape and card I/O devices are not considered in this example or shown on the worksheets.

Convertible-to-real drives are included to provide backup for critical system, Mass Storage System, and application data sets. Although the 3350 drives are string-switched with staging drives, they can be used only as real (nonstaging) drives.

Configuration B

In this example, the connections to the Mass Storage Control are asymmetrical from Model 168AP, and symmetrical from Model 158MP. The 168AP consists of two tightly coupled processors, CPU and APU; the 158MP consists of two tightly coupled processors, CPU0 and CPU1. If both processors in the pair are physically connected to the Mass Storage Control, the connections are symmetrical; if only one is physically connected to the Mass Storage Control, the connections are asymmetrical.

Although only the CPU is physically connected and logically specified to the Mass Storage Control, the APU can process data in the Mass Storage System when the maintenance switch is in the APU position.

Using Worksheet 1

The Mass Storage System components identified on Worksheet 1, as shown in Appendix B.

- Model 158MP, represented by the first two CPU blocks (W and X), CPUIDs 0202220158 and 0203330158
- Model 168AP, represented by the last two CPU blocks (Y and Z), CPUIDs 0204440168 and 0205550168. Because the AP connections to the Mass Storage Control are asymmetrical, no reference to the APU (CPUID 0200555168) is required.
- Four 3830-3s, represented by the first four Staging Adapter blocks. All have the two-channel switch additional feature.
- Twenty-four 3330 and 3350 drives. Sixteen for staging and eight for nonstaging. Eight of the staging drives are convertible-to-real.

The staging drives are distributed to staging drive goups as follows:

- Six drives to staging drive group 00. Two are 3330-11s and four are 3330-1s, for a total of 404 pages.

  The XYZ Company plans to use 341 pages for staging and reserve 63 pages for the Mass Storage Control. The reserved pages include 51 pages for the Mass Storage Control primary tables on SSID 000, and four pages each on SSID 002, 004, and 006 for the Mass Storage Control secondary tables and their two alternate areas. The drive containing the Mass Storage Control primary tables is reserved by varying its SSID (000) offline as staging.

The Mass Storage Control Table Create program does not permit the four Mass Storage Control tables (device addresses 00, 01, 08, and 09 on channel interface B of Staging Adapters 0 and 1 to be specified as REAL or CONVREAL. Their function must be specified as staging.

- Six drives to staging drive group 01. Two are 3330-11s and four are 3330-1s for a total of 404 pages. Two of the 3330-1s and the two 3330-11 drives are specified as CONVREAL. These convertible-to-real drives, varied online as staging, can be used as backup for the same model drives that contain the Mass Storage Control tables with only a pack and plug change within their respective strings.
- All four drives in staging drive group 02 are 3330-11s for a total of 400 pages. These are specified as CONVREAL and can be used as backup to the 3330-11 real drives when they are varied online as REAL.

The physical DASD parameters, DEVn, Model and Drive, should be completed for each drive, as shown on Worksheet 1. (Note that the 3333 is specified to the program as a 3330.) Be careful in designating DEVn (the logical address plug) because of location requirements of the primary and secondary Mass Storage Control tables and their alternate areas. Refer to Figures 19, 20, and 21 and their corresponding descriptions for additional information. Although drives A and B in the first string under Staging

Adapters 0 and 1 are designated DEV0 and DEV1, you can name any two drives in the string as DEV0 and DEV1 as long as the physical and logical parameters correspond to the address plug.

The logical DASD parameters can now be assigned. Each Staging Adapter or Staging Adapter pair controls up to two staging drive groups, LCU0 and LCU1. To identify LCU0 and LCU1 pairs under different Staging Adapters, a staging drive group number between 00 and 27 is assigned. The Mass Storage Control tables must be on staging drive group 00, but all other staging drive group designations are not restricted and can be used in any sequence. Within each staging drive group there are up to eight logical 3330-1 staging drives, LUA0-LUA7. The value of y in LUAy must be even for 3330-11 drives because each 3330-11 is the equivalent of two 3330-1 drives. Although the Mass Storage Control tables are fixed as LUA0, LUA2, LUA4, and LUA6 whether they are on 3330-1 or 3330-11 drives, this does not affect the total number of logical 3330-1 drives. All three staging drive groups for Configuration B have eight logical 3330-1 drives.

Using Worksheet 2

The physical and logical DASD parameters on Worksheet 1 are the input to the DASD commands on Worksheet 2. Four DASD commands are required to describe Configuration B, one for each 3330 or 3350 disk drive and associated control unit.

The physical DASD Drive parameter on Worksheet 1 identifies each drive as REAL (R), CONVREAL (C), staging (S), REAL shared (RS or CONVREAL shared CS. Use this parameter as a reminder that some drives may require additional IODEVICE cards or IODEVICE card changes. Only the REAL parameter is used in the DASD command.

Each DEVn is identified as a member of a staging drive group (SSID=SDGxx, LCUz, LUAy) or as a real drive. A CONVREAL drive is a staging drive group member; a REAL drive is not. Two Staging Adapter pairs are used in Configuration B. Under Staging Adapter pair SA=2 and SA=3, one string of 3330-11 staging and nonstaging drives are string-switched with 3350 nonstaging drives. Although 3330-1 and 3330-11 drives can be mixed within a string, they cannot be mixed with 3350 drives, because the 3350 can function only in native (nonstaging) mode.

Delete the unused portions of each Mass Storage Control Table Create command, including the comma that ends each command.

Using Worksheet 3

The 3851 Model B2 has two data recording controls, (DRC0 and DRC1); each with two connectors. A Staging Adapter can address a maximum of four data recording controls in one or two Mass Storage Facilities using up to four different logical connection addresses, CU00, CU01, CU10 and CU11. The four connectors on the Model B2 for Configuration B each connect to a different Staging Adapter using logical address CU00. This provides every data recording device (SSIDs 200-203) with two stage and destage paths to each staging drive group. The Staging Adapter to DRC connections shown on Worksheet 1 are specified to the program in four LOWERCON commands, one for each Staging Adapter.

Using Worksheet 4

The CPU to Staging Adapter Connection V and the unit address for virtual shared VS, REAL shared RS, and CONVREAL shared CS drives for each host is planned concurrently. The XYZ Company plans to specify 124 virtual UCBs (or virtual drives) as follows:

- 64 UCBs for Model 158MP, 22 each in staging drive groups 00 and 01, and 20 in staging drive group 02. Four of the virtual drives, two each in staging drive groups 00 and 01, are to be permanently mounted in VATLST for the exclusive use of Model 158MP. All others are to be shared with Model 168AP.
- 60 UCBs for Model 168AP. Assign 20 each to staging drive groups 00, 01, and 02.

Access to the three staging drives groups are planned for each CPU. Worksheet 1 shows that each CPU has a primary channel path to each staging drive group through one Staging Adapter, and an alternate channel path through the other Staging Adapter of the pair. In Configuration B, CPU0 and CPU1 of Model 158MP use the opposite Staging Adapters for their primary and alternate paths to each staging drive group.

Now that the DASD Drive parameters, CPU to channel interfaces, and the number of virtual UCBs are known, the real and virtual unit addresses can be planned on Worksheet 4. Each pair of 64 device addresses is coded as the primary and alternate channel interfaces from one CPU to a staging drive group *pair*. Worksheet 4A in Appendix B shows the primary (SA=0, CI=B, CUID=00, CH=4) and alternate (SA=1, CI=D, CUID=00, CH=6 channel interfaces from Model 158MP CPU0 (CPUID=0202220158) to staging drive groups 00 and 01.

The real unit address for REAL and CONVREAL drives are assigned first because they are determined by the electrical paths. The Mass Storage Control Table Create program uses DEV0 through DEV7 for logical address plug values under each 3333 or 3350 (CUA). Labels should be used to identify the maximum 32 logical address plugs under a Staging Adapter (or pair) as:

DEVn	CUA=0	CUA=1	CUA=2	CUA=3
0	00	08	10	18
1	01	09	11	19
2	02	0A	12	1A
3	03	0B	13	1B
4	04	0C	14	1C
5	05	0D	15	1D
6	06	0E	16	1E
7	07	0F	17	1F

In Configuration B, all of the real addresses for REAL and CONVREAL drives have primary and alternate addresses and are shared by Models 158MP and 168AP, except device 05. Device 05 under Staging Adapters 0 and 1 is a nonshared CONVREAL drive that will be used to transfer data from 3330-1 disk packs to mass storage volumes using the Access Method Services CONVERTV command. Refer to the *OS/VS Mass Storage System (MSS) Services: Reference Information* publication for CONVERTV requirements.

Device addresses 00, 01, 08, and 09 on interface B of Staging Adapters 0 and 1 are assigned by the Mass Storage Control Table Create program and are not specified. Use the physical DASD Drive parameter noted on Worksheet 4 to identify the REAL shared (RS) and CONVREAL shared (CS) drives that require real unit addresses. Device addresses 02-04 and 0A-0B are CONVREAL shared drives; device addresses 0E-0F are REAL shared drives. The remaining 52 addresses are available for assignment to the 44 virtual drives in staging drive groups 00 and 01. Contiguous unit addresses are assigned to use all of the unit control words on the channel. Virtual unit addresses 40C and 40D (device addresses 0C and 0D) can be used by OLTEP as real unit addresses (if drives 40C and 40D are varied offline) for diagnostic routines on staging drives E and F in the second string.

Device addresses 00, 01, 08, and 09 on channel interfaces B, C, and D are planned for direct access to the Mass Storage Control tables. The Customer Engineer has been advised so that he can plug the Staging Adapters as required. The IODEVICE card changes are shown in Figure 49

Worksheets 4A through 4F show the unit addresses assigned for each of the 12 channel interfaces on the four Staging Adapters.

Using Worksheet 5

The unit addresses planned on Worksheet 4 must be specified as device addresses to the UPPERCON command on Worksheet 5. Refer to Worksheets 4A and 5A in Appendix B to see how the unit addresses are specified as address ranges to the program. If an alternate path is desired, specify ALTPATH with each optional V, VS, CONVREAL, and REAL parameter. If ALTPATH is used with any optional parameter, the second UPPERCON command at the bottom of the worksheet must be completed, as shown on Worksheet 5A.

Worksheets 5A through 5F correspond to Worksheets 4A through 4F. For consistency with System Control Program usage, SDGxx notation 00, 01, 02 are used. However, SDGxx notations 0, 1, and 2 can be used with the Mass Storage Control Table Create program.

Using Worksheet 6

There are three staging drive groups in Configuration B, each requiring one SDGxx command. The ACTPAGE values for the lower and upper deltas are 30, 40. The LRUGROUP parameter value is 6,2,2,6 for all staging drive groups. Consider using LRUGROUP=6,2,2,6 in your configuration, as described in the chapter "Staging Drive Group Space".

The XYZ Company plans to run with the Trace function on and periodically analyze Mass Storage System activity. This will help them establish an activity base and give them an understanding of how the activity is distributed among the staging drive groups as their workload increases.

Using Worksheet 7

CPUCONFG Command

Connection I between each CPU and the Mass Storage Control channel interface is specified by the CPUCONFG command. Although there are four CPUs, only three of them are physically connected and specified: Model 158MP CPU0 and CPU1 (CPUID 0202220158 and 0203330158) and Model 168AP CPU (CPUID 0204440168). Because Model 168AP is asymmetrically connected to the Mass Storage Control, its APU is not connected or specified, physically or logically, to the program.

Worksheet 7 in Appendix B shows the three CPUCONFG commands for the XYZ Company; the OPSYS for each CPU must be OS/VS2.

MSFn Command

One MSFn command is required for the 3851 Model B2, and the SIZE parameter is 2. Because the XYZ Company plans to use the standard data cartridge loading pattern of the program, the LOADMAP parameter is deleted.

CREATE Command

During system checkout, the program for the XYZ Company is run with the NEW and PRNTONLY parameters. No IODEVICE cards are punched and no Mass Storage Control tables are created, although the VOLID is entered as XYZA00.

Because both hosts are at the required System Control Program level, the CYLFAULT parameter is "yes". This allows cylinder faults to be queued and processed when attention interrupts occur rather than keeping the UCB busy until the cylinder is staged.

Program Input

Use Worksheets 2, 3, 5, 6, and 7 to punch the input cards for the program. To make the output listing easier to use, consider putting the input cards in alphabetic order of the commands. Although the OPSYS parameter for CPUs is OS/VS2, the XYZ Company ran the Mass Storage Control Table Create program under OS/VS1.

Program Output

After an error-free run of the Mass Storage Control Table Create program, carefully check the listing of commands again to be certain they correctly describe Configuration B. Check the Host Connections Report to ensure the unit addresses shown are correct. Note that the only virtual unit addresses shown on this report are those that map to specific physical devices and provide an electrical path for OLTEP. It is important that each physical device have a virtual or real unit address available for OLTEP use.

The Connections Report lists most of the physical connections by logical name, and includes:

- Connection III, Staging Adapter to Disk Storage and Control.
- Connection IV, Staging Adapter to Mass Storage Facility
- Connection V, CPU to Staging Adapter

The data recording control-to-data recording device connections are identified by an SSID list. The actual connections are shown in Figure 18.

Changes to the IODEVICE cards punched by that the Mass Storage Control Table Create program are shown in Figure 49.

Host	Sequence	Unit Addresses	Drive	Notes
158MP	01	400-401	RS	1, 2, 13
	02	408-409	RS	1, 2, 13
	12	400-401		12
	13	402-404	CS	3, 5
	New	405	C	7
	14	406-407	VS	2
	16	408-409		12
	17	40A-40B	CS	2, 3
	18	40C-40D	VS	2
	20	40E-40F	RS	2
	21	410-433	VS	2, 4
	New	434-435	VS	8, 14
	29	540-543	CS	2, 3
	29	544-545	RS	2
	30	546-547	VS	2
	32	548-54B	RS	2
	33	54C-55D	VS	2
168AP	New	100-101	RS	10, 13
	03	102-104	CS	2, 6
	New	105	C	9
	04	106-107	VS	2
	New	108-109	RS	11, 13
	06	10A-10B	CS	2, 3
	07	10C-10D	VS	2
	09	10E-10F	RS	2
	10	110-133	VS	2
	23	240-243	CS	2, 3
	New	244-245	RS	2
	24	246-247	VS	2
	26	248-24B	RS	2
	27	24C-25D	VS	2

Notes

1. Add: ,OPTCHAN=6
2. Add, or change ,FEATURE=(SHARED,ALTCTRL)
3. Add: ,OFFLINE=YES
4. Change: IODEVICE ADDRESS=(410,36)
5. Change: IODEVICE ADDRESS=(402,3)
6. Change: IODEVICE ADDRESS=(102,3)
7. New: IODEVICE ADDRESS=(405,1),UNIT=3330,OPTCHAN=6, FEATURE=(ALTCTRL),OFFLINE=YES
8. New: IODEVICE ADDRESS=(434,2),UNIT=3330V,OPTCHAN=6, FEATURE=(SHAREDUP,ALTCTRL),OFFLINE=YES
9. New: IODEVICE ADDRESS=(105,1),UNIT=3330,OPTCHAN=6, FEATURE=(ALTCTRL),OFFLINE=YES
10. New: IODEVICE ADDRESS=(100,2),UNIT=3330,OPTCHAN=3, FEATURE=(SHARED,ALTCTRL)
11. New: IODEVICE ADDRESS=(108,2),UNIT=3330,MODEL=11,OPTCHAN=3, FEATURE=(SHARED,ALTCTRL)
12. Delete IODEVICE card. The Mass Storage Control Table Create program punches two sets of IODEVICE cards for Mass Storage Control tables on channel interface B if Staging Adapters 0 and 1 are specified.
13. New VATLST entry as reserved and private for direct access to the Mass Storage Control tables by Models 158MP and 168AP.
14. New VATLST entry as reserved and private for exclusive use by Model 158MP.

Figure 49. IODEVICE Card Changes, Configuration B

The SSIDs for all the staging drives are assigned by the Mass Storage Control Table Create program, as shown on the Host Connections Report. Place the ADDR or ADDR/SSID on the logical address plugs and add them to Worksheet 1. Make copies of Worksheets 1 and 4 available to the computer operations people and to the Customer Engineer, so that they can readily determine the configuration, SSIDs, addresses, and staging drive groups.

Because Staging Adapters 0 and 1 are specified, the program punches two sets of IODEVICE cards. Figure 49 shows the IODEVICE card changes by sequence number. This number is in columns 7 and 8 of the IODEVICE card, as shown on the program listing. Interface B of Staging Adapters 0 and 1 is connected to the 158MP, therefore, one set of IODEVICE cards is deleted as shown in Figure 49, Sequences 12 and 16. The OPTCHAN and FEATURE=(SHARED,ALTCTRL) must be added to the other set of IODEVICE cards for direct access to the Mass Storage Control tables, as shown by Sequences 01 and 02. A set of IODEVICE cards must be punched for Model 168AP to directly access the Mass Storage Control tables on interface C of Staging Adapters 0 and 1 as shown by the new IODEVICE cards for unit addresses 100-101 and 108-109.

Two IODEVICE cards must be added with FEATURE=(SHAREUP,ALTCTRL) for virtual unit addresses 434-435 and 244-245. The SHAREDUP parameter allows reserve and release logic for shared devices to be bypassed when CPU0 and CPU1 of Model 158MP are associated as a tightly coupled pair. All other virtual unit addresses are shared with Model 168AP and reserve and release logic may be required and must not be bypassed. The parameter OFFLINE=YES has been added to the IODEVICE addresses of all CONVREAL drives.

Also, two IODEVICE cards must be added with FEATURE=(ALTCTRL) for real unit address 405 and 105. These are the addresses of device 05 under Staging Adapters 0 and 1, which is the device to be used for the CONVERTV operation. An IODEVICE card must be punched for the two Mass Storage Controls in the 3851 Model B2. The Mass Storage Control IODEVICE card for Model 158MP is:

```
IODEVICE ADDRESS=(040,2),UNIT=3851
```

The card for Model 168AP is:

```
IODEVICE ADDRESS=(0E0,2),UNIT=3851
```

Most of the IODEVICE card changes consists of adding the ALTCTRL parameter. OS/VS2 MVS does not attempt to use an alternate path on a control unit busy condition unless ALTCTRL has been specified. Except for unit addresses 400-401 and 108-109, the IODEVICE cards for Model 158MP are identified by CPUID 0203330158.

After the IODEVICE cards changes shown in Figure 49 have been made, they should be separated by CPU, and merged with the UNITNAME cards and other system IODEVICE cards that may be required. These cards are the input to the OS/VS2 MVS SYSGEN or IOGEN for each CPU.

Configuration C

In this example, the connections are asymmetrical from both Models 158MP and 168AP. The 158MP is connected to the Mass Storage Control as a uniprocessor with a physical connection and logical reference only to CPU0 (CPUID 0202220158); the 168AP is connected to the Mass Storage Control (as it was in Configuration B) with a physical connection and reference only to CPU (CPUID 0204440168).

The components and basic addresses are identical to Configuration B. Only the changes required to the connections from the Model 158MP to the Mass Storage Control are described. Appendix C contains only the worksheets that required changes.

Using Worksheet 1

Because the physical and logical identifications of components are the same, the procedure described for Configuration B is applicable here. Connection X1 to the Mass Storage Control from CPU1 (CPUID 0203330158) channel 0 has been eliminated, as well as the alternate channel connections W6 and X6 to Staging Adapters 0. This is described in "Using Worksheet 4."

Using Worksheet 2

The distribution and function of the 3330 and 3350 drives remain the same. The DASD commands on Worksheet 2 are shown in Appendix B only.

Using Worksheet 3

There are no changes in the LOWERCON commands. Worksheet 3 is shown in Appendix B only.

Using Worksheet 4

The virtual, REAL and CONVREAL shared address requirements remain the same. However, the number of 158MP connections to the Staging Adapters is reduced. In Configuration B, CPU0 and CPU1 had a primary and alternate path to each staging drive group, and a primary channel path to the Mass Storage Control. In Configuration C, CPU0 and CPU1 have only a primary channel path to each staging drive group and only CPU0 has a primary channel path to the Mass Storage Control. The 158MP CPU0 and CPU1 alternate channel connections to the Staging Adapters are eliminated.

The program requires that all MP connections to the Staging Adapters be specified *as if they were connected to the same MPCPU that is connected to the Mass Storage Control.*

The use of alternate interface D on Staging Adapters 0, 1, 2, and 3 is also eliminated. Primary interface B on Staging Adapters 0, 1, 2, and 3 remains the same. Thus, Appendix C shows only Worksheets 4A through 4D. Worksheets 4E and 4F remain the same and are shown only in Appendix B.

The eight examples in Figure 50 show the different connections to CPU0 and CPU1, and the logical specification of each that is made to the Mass Storage Control Table Create program for each CPU0.

- Examples 1, 2, and 3 are those described in the *OS/VS Mass Storage Control Table Create* publication.
- Example 4 adds an alternate channel to the CPU0 of Example 1 and shows its logical specification to CPU0.

- Example 5 adds an alternate channel to the CPU1 of Example 2 and shows its logical specification to CPU0.
- Examples 6 and 7 show other variations in primary channel connections.
- Example 8 shows the physical connections of Configuration B (symmetrical connection) and the logical connections of Configuration C (asymmetrical) to the Mass Storage Control. Note that CPU0 and CPU1 alternate channel paths are eliminated, and the CPU1 primary paths are logically specified as CPU0 alternate channel paths on nonexistent channel F.

	Physical Connections to									Logical Specification to				
	CPU0					CPU1				CPU0				
Ex.	**Path**	**SA**	**CI**	**CUID**	**CH**	**SA**	**CI**	**CUID**	**CH**	**Path**	**SA**	**CI**	**CUID**	**CH**
1	PRI	6	B	01	2					PRI	6	B	01	2
2	PRI					6	C	01	3	PRI	6	C	01	3
3	PRI	7	B	10	4					PRI	7	B	10	4
	PRI					7	D	10	4	ALT	7	D	10	F
4	PRI	6	B	01	2					PRI	6	B	01	2
	ALT	7	B	01	3					ALT	7	B	01	3
5	PRI					6	C	01	3	PRI	6	C	01	3
	ALT					6	D	01	4	ALT	6	D	01	4
6	PRI	2	B	00	4					PRI	2	B	00	4
	PRI					2	C	01	4	PRI	2	C	01	4
7	PRI	2	B	00	4					PRI	2	B	00	4
	PRI					3	B	01	4	PRI	3	B	01	4
8	PRI	0	B	00	4					PRI	0	B	00	4
	ALT	1	D	00	6					Connection Removed				
	PRI					1	B	00	4	ALT	1	B	00	F
	ALT					0	D	00	6	Connection Removed				
	PRI	2	B	01	5					PRI	2	B	01	5
	ALT	3	D	01	6					Connection Removed				
	PRI					3	B	01	5	ALT	3	B	01	F
	ALT					2	D	01	6					

Notes:

1. CPU0 is asymmetrically connected to the Mass Storage Control.
2. CH F is a nonexistent logical channel on CPU0.
3. ▨ alternate channel paths that are eliminated in Configuration C with the asymmetric connection to the Mass Storage Control.
4. Both CPU0 and CPU1 in Example 6 have primary channel connections to a unique set of virtual unit addresses.

Figure 50. Physical and Logical Channel Paths for CPU0 and CPU1

When the channel (CH) and control unit identification (CUID) are unique on CPU0 and CPU1, the connection is logically specified to CPU0 without change. When the CH and CUID are the same on CPU0 and CPU1, the connection is logically specified to CPU0 as an unused or nonexistent channel.

For the XYZ Company, the MPCPU connected to the Mass Storage Control is CPU0 (CPUID 0202220158). This means that the Staging Adapter primary channel paths from CPU1 that are identical to CPU0 must be specified as CPU0 alternate channel paths to the Mass Storage Control. Because OS/VS MVS allows only one alternate channel path, no other channel path can be specified, thus eliminating the CPU0 and CPU1 alternate channel paths.

Using Worksheet 5

Although the unit addresses to Models 158MP and 168AP have not changed, the number of Staging Adapter interfaces to the three staging drive groups has been reduced. Configuration B had four primary and four alternate channel interfaces specified; Configuration C has only the four primary Staging Adapter channel interfaces to the three staging drive groups; two from CPU0 and two from CPU1. These are shown on Worksheets 4A to 4D in Appendix C.

The two primary channel paths from CPU0 are specified to the Mass Storage Control Table Create program as primary; the two paths from CPU1 are specified as alternate channel paths *from CPU0 on nonexistent channel F*. This specification is required so as to establish interface B on Staging Adapters 1 and 3 as a valid path in the Mass Storage Control's Configuration Table. Channel F appears in the Configuration Table but is never used because the IODEVICE cards containing the OPTCHAN=F parameter must be removed. OS/VS2 MVS uses the CPU0 and CPU1 paths to the three staging drive groups as primary channel paths.

The four UPPERCON commands, two primary and two alternate, are shown on Worksheets 5A and 5C in Appendix C. Although Worksheets 5B and 5D are blank, they are included to illustrate the UPPERCON ccommand changes that are necessary. when the MP symmetric connections are changed to asymmetric. Worksheets 4A to 4D correlate with the UPPERCON commands on Worksheets 5A to 5D.

Using Worksheet 6

No changes are required in the SDGxx commands. Worksheet 6 is shown in Appendix B only.

Using Worksheet 7

CPUCONFG Command

Two CPUCONFG commands are needed, one for each host CPU, but the one for Model 158MP specifies only CPUID 0202220158. The MPCPU parameter is not used. Only the CPUCONFG command portion of Worksheet 7 is in Appendix C.

MSFn Command

No changes are required in the MSFn command, which is shown in Appendix B only.

CREATE Command

No changes are required for the CREATE command, which is shown in Appendix B only.

Program Input

Although the AP and MP asymmetric connections to the Mass Storage Control are supported by OS/VS2 MVS, the Mass Storage Control Table Create program can be executed on an OS/VS1 system. However, the OPSYS parameter must specify VS2 or JES3.

Program Output

The Mass Storage Control Table Create program output for Configuration C is in Appendix C. This section describes the differences between the outputs for Configurations B and C.

The Host Connections Report for Configuration C does not show CPUID 0203330158. The unit addresses that have prefix F correspond to the CPU1 device plugs, and they are listed in the correct Staging Adapter channel interface column, but as alternates.

The Connections Report for Configuration C is almost identical to Configuration Bs'; the difference is in the number of Staging Adapter interfaces connected to Model 158MP.

Figure 51 shows the IODEVICE card changes. The IODEVICE card for the nonexistent alternate channel F (OPTCHAN=F) must be removed. Model 158MP requires only FEATURE=(SHARED) instead of FEATURE=(SHARED,ALTCTRL), because there are no alternate channel paths. The remaining changes are as described for Configuration B.

As before, one set of IODEVICE cards for direct access must be deleted for Model 158MP, and one set must be added for 168AP. The changed IODEVICE and UNITNAME cards should be separated by CPU along with other IODEVICE cards that may be required. They are the input to the OS/VS2 SYSGEN or IOGEN for each CPU.

Technical Considerations

Configuration C for the XYZ Company permits both the 158MP and 168AP to access data in the three staging drive groups when they are associated as tightly coupled multiprocessors. If 158MP CPU0 and CPU1 are disassociated, the Staging Adapter channel interfaces connected to CPU1 should be disabled to avoid any potential problems resulting from unsolicited interrupts or loss of data.

Input specified to the Mass Storage Control Table Create program should adhere to the procedures described in this chapter and in the *OS/VS Mass Storage Control Table Create* publication. Although the program can be executed on an OS/VS1 or OS/VS2 MVS system, the OPSYS parameter in the CPUCONFG command must be specified as VS2 or JES3; any other designation allows alternate channel paths to the Mass Storage Control or to the Staging Adapters, which are not valid under OS/VS2 MVS.

Host	Sequence	Unit Addresses	Drive	Notes
158MP	01	400-401	RS	2,12
	02	402-404	CS	1, 2, 3
	New	405	C	6
	03	406-407	VS	1
	05	408-409	RS	2, 12
	06	40A-40B	CS	1, 2, 3
	07	40C-40D	VS	1
	09	40E-40F	RS	1,2
	10	410-433	VS	1, 5
	New	434-435	VS	7, 13
	21	F00-F01		11
	22	F08-F09		11
	23	540-543	CS	1, 2, 3
	23	544-545	RS	1,2
	24	546-547	VS	1
	26	548-54B	RS	1,2
	27	54C-55D	VS	1
168AP	New	100-101	RS	9, 12
	12	102-104	CS	3, 4
	New	105	C	8
	13	106-107	VS	
	New	108-109	RS	10, 12
	15	10A-10B	CS	4
	16	10C-10D	VS	
	18	10E-10F	RS	4
	19	110-133	VS	
	29	240-243	CS	3, 4
	29	244-245	RS	4
	30	246-247	VS	
	32	248-24B	RS	4
	33	24C-25D	VS	

Notes

1. Remove: ,OPTCHAN=F
2. Add: ,FEATURE=(SHARED)
3. Add: ,OFFLINE=YES
4. Add or change: ,FEATURE=(SHARED,ALTCTRL)
5. Change IODEVICE=(410,36)
6. New: IODEVICE ADDRESS=(405,1),UNIT=3330,OPTCHAN=6,OFFLINE=YES
7. New: IODEVICE ADDRESS=(434,2),UNIT=3330V,FEATURE=(SHAREDUP)
8. New: IODEVICE ADDRESS=(105,1),UNIT=3330,OPTCHAN=6, FEATURE=(ALTCTRL),OFFLINE=YES
9. New: IODEVICE ADDRESS=(100,2),UNIT=3330,OPTCHAN=3, FEATURE=(SHARED,ALTCTRL)
10. New: IODEVICE ADDRESS=(108,2),UNIT=3330,MODEL=11,OPTCHAN=3, FEATURE=(SHARED,ALTCTRL)
11. Delete IODEVICE card. The Mass Storage Control Table Create program punches two sets of IODEVICE cards on channel interface B if Staging Adapters 0 and 1 are specified.
12. VATLST entry is reserved and private for direct access to Mass Storage Control tables by Models 158MP and 168AP.
13. Add to VATLST entry as reserved and private for exclusive use by Model 158MP.

Figure 51. IODEVICE Card Changes, Configuration C

Data Security and Recovery Planning

Planning for data security and for data recovery is related. Security planning involves the physical and logical protection of data; recovery planning ensures that there are recovery and backup procedures for the data protected.

Data Security

Depending on how secure a computer installation is to be, the measures you take to protect data might include site, building, computer room, and people, and you are already familiar with such measures. What we present here describes data security considerations for the Mass Storage System, as illustrated in Figure 52.

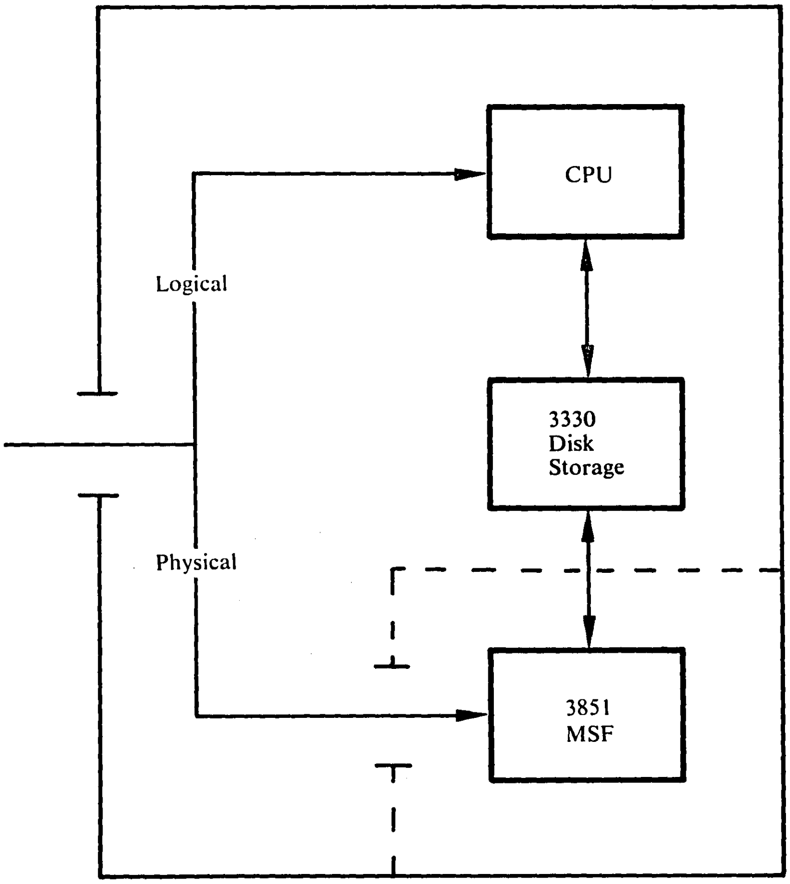

Figure 52. Logical and Physical Data Protection

The physical and logical data security considerations for a Mass Storage System configuration cover:

- The active volume inventory within the Mass Storage System
- The inactive volume inventory within the Mass Storage Facility
- The active or inactive volume inventory outside the Mass Storage Facility, but within the computer room

If the Mass Storage Facility contains data that needs to be moved off site as mass storage volumes, you should review the labeling, shipping, and handling precautions listed in the *IBM 3850 Data Cartridge Care and Handling Instructions* publication.

As shown in Figure 52, data protection measures for the Mass Storage System must address both the logical and physical access to the data. You security plan should cover the possibility of destruction, disclosure, and modification of data by either accidental and/or intentional means, whichever presents the greatest hazard to your installation.

Logical Data Protection

The active volume inventory is the data that is available for staging, or is staged and in use. Its protection, after it is staged, is dependent on the System Control Program level of control, and you can use several types of programs to access it, such as authorized or unauthorized System Control Program services and user programs, the Access Method Services, or the OS/VS utilities.

The Mass Storage System Communicator (MSSC) requires that the Access Method Services modules be kept in Authorized Program Facility (APF) protected libraries. These libraries can be either SYS1.LINKLIB or other password-protected APF libraries. Complete information about the logical data protection items summarized here can be found in *OS/VS Mass Storage System (MSS) Services: Reference Information.*

Hardware, software, and microcode for the Mass Storage System are designed with safeguards so that requested data can be retrieved. The Mass Storage Control verifies data cartridge retrievals by reading the XYZ address on the cartridge label. The System Control Program reads the VTOC to verify the volume serial number of the mass storage volume.

The mass storage volume READONLY, DASDERASE, and EXCLUSIVE attributes provide logical protection for the active volume inventory on staging drives or in the Mass Storage Facility. Use the Access Method Services MODIFYV command in conjunction with READONLY to change the mass storage volume attribute to READWRITE when required.

The Access Method Services STOREV, ADDV, COPYV, and RECOVERV commands "move" volumes between the active and inactive volume inventories (both within the Mass Storage Facility), without requiring actual data cartridge moves. The Mass Storage Control updates its tables to reflect the status of each volume. If a program attempts to access a volume listed in the inactive volume inventory, a message tells the operator the volume location, and the job aborts.

The Access Method Services ADDV, COPYV, EJECTC, EJECTV, and STOREV commands move data cartridges in and out of the Mass Storage Facility. Regardless of whether the data cartridges are stored in the computer room adjacent to the Mass Storage Facility, in an on-site vault, or in an off-site vault, the correct handling precautions must be observed.

The Access Method Services command MODIFYC is used for data cartridge recovery. It can:

- Rewrite a data cartridge label or fields within the label
- Move a data cartridge from one location to another, including the cartridge access station, with or without the use of a data recording device

The data cartridge location changes are not maintained in the Mass Storage Control tables, because one of the uses of the command is to make the data cartridge label agree with the Mass Storage Control tables. Careless use of the MODIFYC command can cause data integrity problems. We recommend that this command, and other critical Access Method Services commands, be considered for placement in password-protected or APF-protected libraries to restrict its use.

The VSAM data set ERASE attribute ensures that deleted data is erased from the data cartridges that make up a mass storage volume. SCRATCHV disassociates the two cartridge serial numbers from the volume serial number, but does not erase the data.

Data protection at the System Control Program level for data sets on mass storage volumes is the same as for data sets on 3330 volumes. If data security has high priority in your installation, the System Control Program used with the Mass Storage System should be OS/VS2 MVS. It is the most secure operating system and supports the Resource Access Control Facility (RACF) for the protection of data sets on mass storage volumes. IBM accepts APARs if unauthorized programs can successfully use a system interface to bypass storage protection, or obtain system control in an unauthorized program state. Password protection is available for data sets on mass storage volumes and we recommend it for the Mass Storage Control tables.

VSAM catalog integrity must be recognized by your computer operations people. You cannot change the serial number of a volume containing VSAM data sets in the VSAM catalog, and a valid copy cannot be substituted for an original using the Access Method Services commands with the NEWSERIAL parameter. Therefore, your backup procedures should not require moving VSAM mass storage volumes that have identical volume serial numbers from one Mass Storage System to another, or require substituting a valid copy if the original VSAM mass storage volume is lost *without* ejecting the original copy or maintaining a record for the ejected volume. See *OS/VS Mass Storage System (MSS) Services: General Information* for further information. Use this integrity factor of the VSAM catalog to reinforce your data protection.

In a Mass Storage System the operator can no longer control the manipulation of mass storage volumes; this is now handled by Access Method Services commands. For VSAM volumes, this approach requires you to be able to verify authorization to access volumes before allowing its manipulation. Today, security checking requires either an individual VSAM data set master password or a VSAM catalog master password for dump and restore operations. This security check is extended to all VSAM volume manipulations through additional Access Method Services commands. Because a catalog master password allows full VSAM volume access, a single password supplied by the command language or by the operator can give full access to volume manipulation in the case of protected data sets.

Physical Data Protection

Physical data protection in your installation might begin at the site entrance, the building entrance, or the computer room entrance. If you protect your data using a locked or controlled-access tape library and computer room, you may want to extend this protection to the Mass Storage System, as controlled tape data is transferred to the Mass Storage Facility. We have found that controlling access to a computer room eliminates unnecessary interruptions and improves data security. Sometimes, a chance observation of a volume serial number, a program name, or a password by an unauthorized person can affect data security, therefore, access to the computer room should be limited to those people who have a need to enter it.

The Mass Storage Facility in Figure 52 is shown in a separate room indicated by the broken line. Depending on the degree of data security you want, your computer room might have walls extending above the false ceiling and below the raised floor; it might be locked and have electronic motion detectors above the ceiling and below the floor; or it might be just locked. Because the Mass Storage Facility requires only minimal operator attention for entering or removing data cartridges, a separate enclosure offers good physical protection. If you decide on a separate enclosure, you should plan for adequate physical facilities, service requirements, and future system expansion.

Two RPQs are available to physically protect data in the Mass Storage Facility. One provides internal Mass Storage Facility modifications and specifications for a user-supplied fire suppression system, and the other provides Mass Storage Facility external cover locks. Your IBM representative can supply you with detailed information on both. The RPQ numbers are:

- RPQMF6167 Fire Suppress–3851 Base
- RPQMF6191 Cover Locks–3851 Base

Fire Suppress or Cover Locks can be installed individually, but there may be co-requisite RPQs depending on the model of Mass Storage Facility.

The *IBM System/370 Installation Manual—Physical Planning* contains information on building, environment, media, and machine requirements, which you should review and consider when developing a data security plan.

Data Recovery

Knowledge and experience gained by your computer installation people regarding data recovery can be applied to the Mass Storage System. This section helps you apply and extend that knowledge and experience to the Mass Storage System by describing operations pertinent to potential problem areas, and includes a number of 'what to do' and 'how to do' considerations and guidelines.

Recovery planning includes both data backup and data recovery. Data backup is the process of creating or retaining information as the basis of a fall-back plan. Data recovery is the execution of that plan when data is destroyed or unavailable. This planning should determine what, and when, action must be taken, and who should take the action.

The Mass Storage System provides features and innovations for recovering data lost through system problems and errors. You should understand what these features are in order to plan adequately for the required level of data backup, and also consult:

- *Operator's Library: IBM 3850 Mass Storage System (MSS) Under OS/VS*
- *OS/VS Mass Storage System (MSS) Services: General Information*
- *OS/VS Mass Storage System (MSS) Services: Reference Information*

Assigning the responsibility of data backup and recovery planning to a member of your installation team who has a good understanding of the Mass Storage System and of the system support, its interaction with OS/VS, and your requirements, is a good beginning.

Job and System Restart

Moving a large portion of your computer installation's workload to a system-managed environment causes changes in daily operating procedures, because the Mass Storage System brings a new environment, new potential failure points, and new restart actions.

Review the restart procedures for each application system that accesses data in the Mass Storage System, to better understand the status and availability of a data set when an error occurs. Data in the Mass Storage System environment is controlled primarily by the Mount, Demount, Acquire, and Relinquish orders, which the CPU sends to the Mass Storage Control.

To put the planning task into the context of a general system flow, the System Control Program and Mass Storage System functions shown in Figure 53 are summarized below. These functions are needed to access an existing data set, or to create a new data set on mass storage volumes:

- *Allocation* determines the quantity and type of volumes required by the user program utilizing the OS/VS JCL parameters, the specified MSVGP, or the catalog. New nonspecific requests can be allocated to mounted or unmounted volumes if the allocation to mounted volume support is in your System Control Program and CONCURRENTUSERS is not zero. The problems that might be encountered here are common ones, such as OS/VS JCL and catalog errors. The Mass Storage System introduces a new unit (3330V) and a new OS/VS JCL parameter (MSVGP) for nonspecific volume requests.
- *Mount Equalization* provides the Space Allocation routine with a list of staging drive groups. The one listed at the top is the best candidate to satisfy mount equalization, based on input from the Mass Storage Control to the Mass Storage Communicator module in the System Control Program. The sampling frequency, in the SYS1.PARMLIB member of MVIKEY00, can be adjusted. If the mass storage volume is already mounted in a multiple CPU environment, or if data from it is already staged but known only to the Mass Storage Control, the staging drive group containing the mounted volume or staged data is moved to the top of the list. The Space Allocation routine then selects a virtual unit associated with that staging drive group.

  The mount equalization record is part of the Verification Table maintained by the Mass Storage Control. If the sampling frequency is incorrect, data reuse effectiveness is reduced but recovery is not a problem. The Mass Storage Control's Trace data can be used to detect this problem and determine the effectiveness of data reuse.

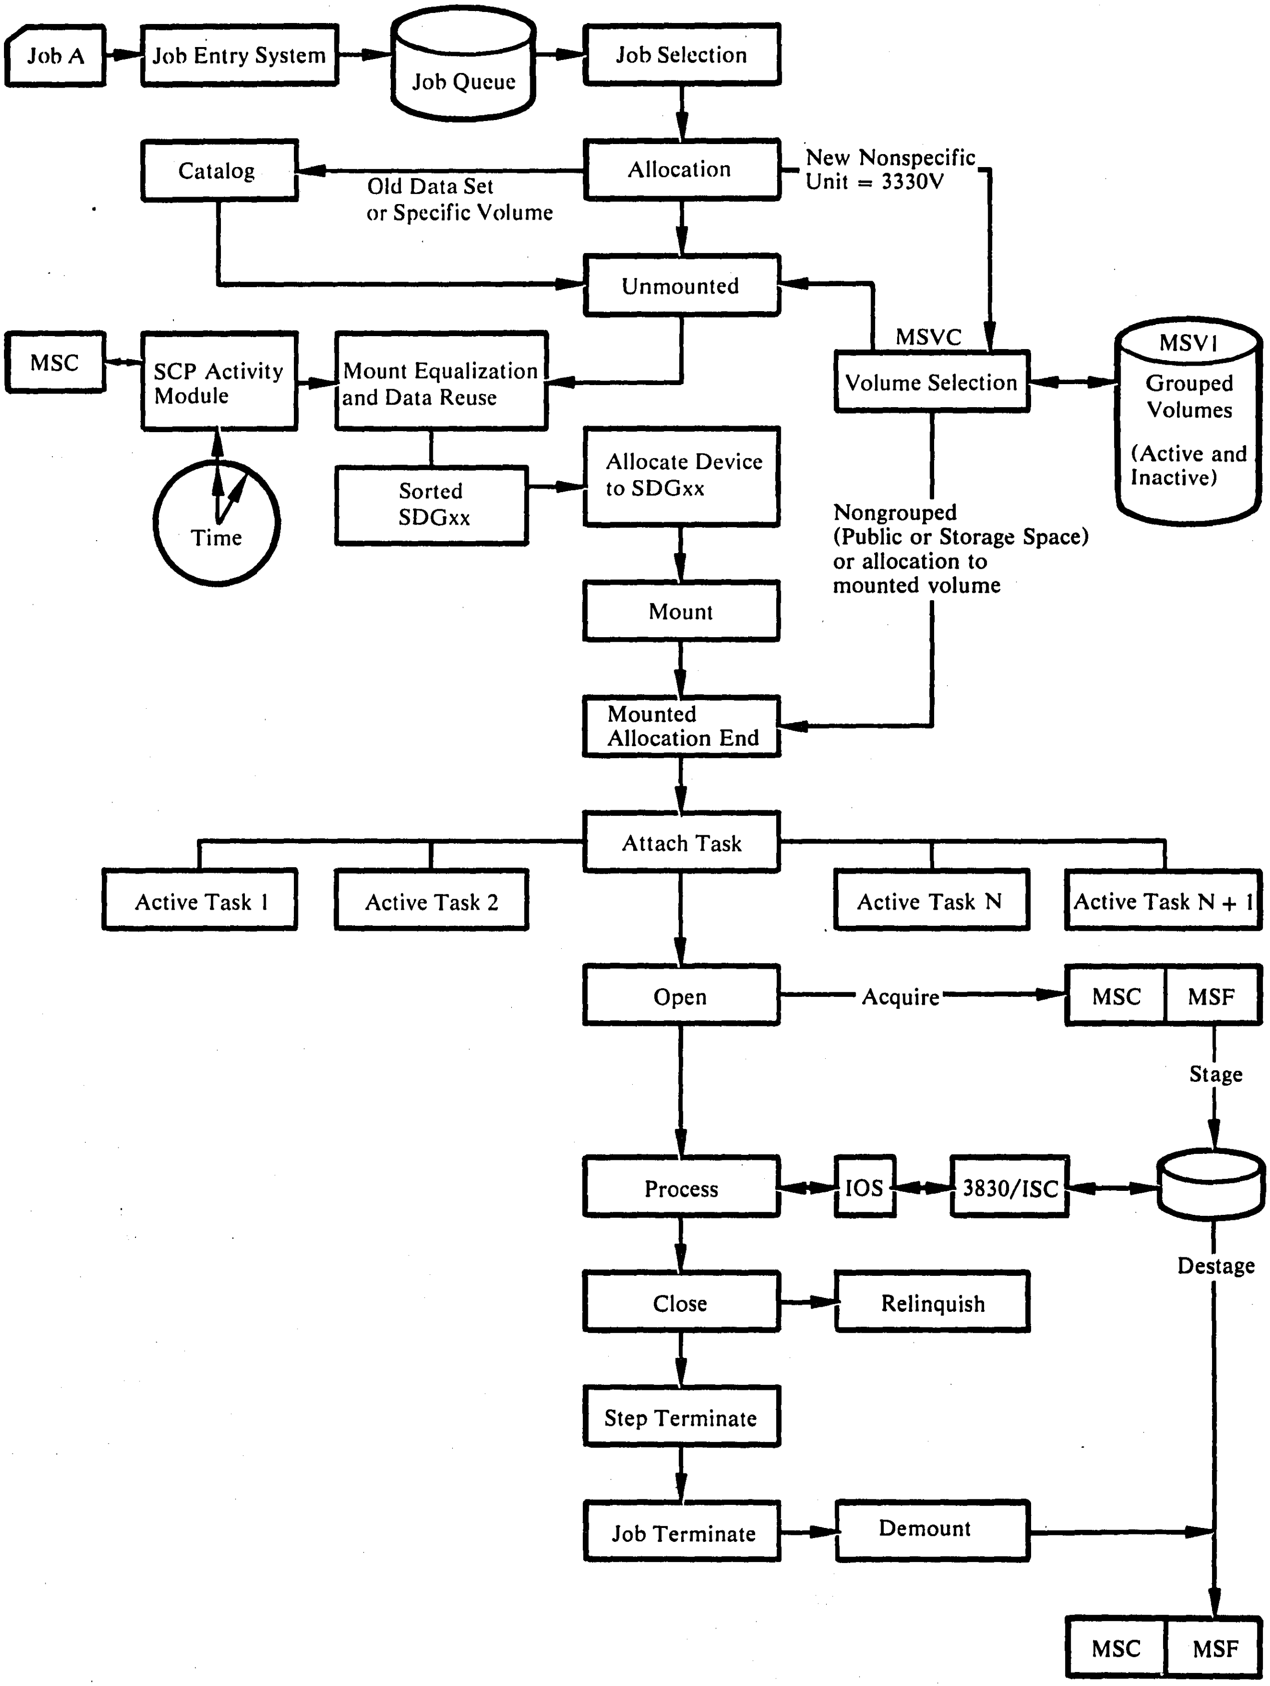

Figure 53. System Flow

- The *Mount* order is issued to the Mass Storage Control by the Space Allocation routine and specifies the volume serial number and virtual unit address, if the mass storage volume is not mounted. It is recorded in the MSVI data set.

  The Mass Storage Control validates (or stages and validates) cylinder 0 in the Staging Adapter tables. It then simulates a pack-change interrupt by causing the Staging Adapters to send an unsolicited Device End to the primary and alternate channel paths attached to the CPU that requested the mount operation.

  The System Control Program verifies that the volume is mounted by reading its label, and posts its serial number to the virtual unit control block. The Mount order causes the Mass Storage Control's dynamic tables to be referenced and updated, and the controlled tables to be referenced. Controlled, dynamic, and static Mass Storage Control tables are described under "Mass Storage Control Tables" later in this chapter. If staging space is not available, the Mass Storage Control honors the Mount order and processes the cylinder fault when the volume label is read.

  Mount operation problems can be caused by data cartridge load, unload, read, and write failures, by disk write errors, and by table synchronization errors between MSVI data sets and the Mass Storage Control tables. The recovery action for volume verification errors in the Mass Storage System environment is similar to your current one, because the data cartridge label area contains both the data cartridge and the mass storage volume serial numbers.

- *Open* routines access the VTOC to locate extents for BDAM, BPAM, BSAM, ISAM, or QSAM data sets. They process the mass storage volume VTOC the same way as the 3330 VTOC. The VSAM catalog contains the extents for VSAM data sets. An Acquire order is issued to the Mass Storage Control, which specifies the data set cylinder ranges to be staged from the specified mass storage volume. For BDAM and BPAM, the entire data set is staged. Because staging occurs in cylinder increments, at least one cylinder is always staged for old data sets. For ISAM, the data set is not staged except as the result of cylinder faults. For BSAM or QSAM, only the data set extents on the first mass storage volume of a multivolume data set are staged. For VSAM, the staging depends on the specified option.

  The Mass Storage Control reads its tables to determine if any of the cylinders requested are already staged. When the staging from any one page is completed, the Staging Adapter tables are validated for the acquired cylinders and the data is available for processing.

  For a new data set, the Acquire order is issued with the Inhibit Stage attribute and no staging occurs unless the allocation is not on cylinder boundaries. If space is specified in number of tracks, entire cylinders are staged to ensure that another data set occupying the the same cylinder is not lost if those cylinders of data are destaged. The Mass Storage Control allocates space in the staging drive group containing the volume label and marks each cylinder valid. This action causes the Mass Storage Control tables to be referenced and updated in the same way as for the Mount order.

  During open, potential problems are data cartridge loading, staging, and unloading errors.

- *Close* routines process the Relinquish order for VSAM data sets with the DESTAGEWAIT option. Relinquish causes the Mass Storage Control's dynamic tables to be accessed and modified. Potential problems might include incorrect data cartridge access, disk read errors, and cartridge write errors.
- *Job Termination* causes mass storage volumes associated with the job to be demounted for the issuing CPU, if there are no other users of the volume and if it is marked private. The MSVI data set is accessed to reflect the demount operation, and to update the space information.

  The Mass Storage Control schedules for destaging all cylinders containing changed data as a result of a Relinquish or a Demount order. When DELETE is the data set disposition parameter, the Relinquish order is issued with the Inhibit Destage attribute and no destaging occurs.

Mass Storage Facility Manual Intervention

If a load or unload failure occurs causing a data cartridge to become unavailable, all unstaged data on that mass storage volume is unavailable, and a console message informs the operator that a failure has occurred. Subsequent job steps that try to access the mass storage volume fail when attempting to mount the volume, and a Customer Engineer is needed to retrieve the data cartridge. The situation is analogous to manual intervention on a 3330 drive. Operation procedures that call for manual intervention might have to be replaced by problem identification procedures to determine the appropriate action for the Mass Storage System. In some cases, it might be advantageous to initiate data recovery procedures, not because data has been lost, but because you want to reduce production delays.

Staging Drive Manual Intervention

When a staging drive becomes unavailable, the virtual drives mapped to it also become unavailable, as well as the data on the virtual volumes (staged portions of mass storage volumes). The Mass Storage Control issues an operator message that indicates whether the problem is associated with the 3330 drive or with the 3333 Disk Storage and Control. If the problem is in the drive, the operator can move the disk pack and logical address plug to another drive under the same 3333; if the problem is in the 3333, the operator can move the pack to another drive under a different 3333 connected to the same Staging Adapter. In either case, the operator should issue the VARY SSID, OFFLINE,S command, which causes Mass Storage Control to destage written data and clear the pack so that it can be removed, making that drive available.

Asynchronous I/O Errors

I/O errors in the Mass Storage System can occur during or after the execution of a program. Disk read errors that occur during a destage operation are usually detected after the execution of a program (VSAM DESTAGEWAIT option excepted). If these errors are identified before the next scheduled run, corrective action can be started before they affect computer operations.

Current Data Backup

You may want to make changes to your current methods to take advantage of certain Mass Storage System functions or make adjustments to them to compensate for differences that have been introduced, as described below.

Generation Data Sets

Generation data sets are commonly used in cyclic tape applications to provide data backup. The process consists of reading an old master file, updating the records with new transactions, and writing a new master file. The new transactions and the old master file are retained so that the process can be repeated if necessary. For critical data sets copies can also be made, especially if reruns are time-consuming or expensive.

When you use this method, the cost of a failure covers only the lost time it causes plus the time required to rerun the application to the point of failure. The cost of the backup data covers only the storage cost of retaining down-level generations of data sets.

Data Set Copy

This method consists of making a direct copy of a data set after it has been created or updated. It is used for data sets that are primarily read-only.

When you use this method, the cost of a failure covers the copy, re-creating the active data set from the copy, lost time due to the failure, and rerun and copy time.

Checkpoint and Journal

This method is used to back up "update-in-place" data. A checkpoint or copy of the data set is recorded in the user's Journal data set. When the data set is updated or modified, the changes are recorded in the Journal data set and, in case of failure, the current level of the data set can be re-created by rerunning the Journal data against the copy.

When you use this method, the recovery cost includes the processing of the Journal data set to find the changes that must be applied to the copy, and the rerun time to make the change. The data backup cost includes the copy and the additional resources required for the Journal data set, lost time, and maintaining the copies and Journal data set.

Data Backup in a Storage Hierarchy

The methods for data backup we have just described are also applicable to the Mass Storage System, except that the Mass Storage System dynamically manages the staging space and records its activity in the Mass Storage Control tables. In case of failure, the status and contents of the staging space are known only to the Mass Storage Control. The Mass Storage System environment is similar to a disk environment, where multiple data sets exist per volume, and contrary to a tape environment, where generally one data set exists per volume.

In the Mass Storage System, data set recovery is keyed to the fact that methods and procedures exist, which can establish the status and location of each data set after a failure occurs.

Failure Modes

Failure modes include media, equipment, and loss of synchronization between the MSVI data set and the Mass Storage Control tables. These failures can affect data sets, tables, catalogs, VTOCs, or MSVI, and they can cause data to be lost or unavailable.

A Data Check due to a permanent read error can result in lost data, depending on whether the media failure occurred on a data cartridge or on a staging disk. A data cartridge read failure results in lost data and the need for recovery. A staging disk read failure can result in lost data, depending on the situation, and usually involves some recovery action.

A permanent write error can occur on a data cartridge or on a staging disk. Write errors on a data cartridge should never cause data to be lost, because the data remains on the staging disk until recovery action corrects the condition. Write errors on a staging disk, on the other hand, can result in lost data if the error is not detected during the write operation. Some recovery action is always associated with a permanent write error, even if no data is lost.

Equipment failures affect the availability of data but do not necessarily cause the data to be lost unless the tape or disk is destroyed. Depending on how critical the application is, immediate data recovery action might be more advisable than waiting for data to become available.

Data Cartridge Permanent Read Error: This error is detected by the Staging Adapter when it is attempting to stage data, and occurs when the Mass Storage Control and Staging Adapter exhaust all of the programmed recovery actions and must stage the data with error. Because the Staging Adapter maintains only a beginning-of-track index, it can no longer recognize the previous record boundaries. The error track is flagged and contains a single concatenated record. If the CPU attempts to access any of the data based on the previous record boundaries, an error return code is sent to the user program.

The worst-case error condition is when multiple tracks are flagged in error, because portions of each track might be on the same 4096-byte data stripe, although errors can involve less than 4096 bytes.

A data stripe contains up to 4096 bytes of user data composed of records (count, key, and data) transmitted to and from a staging drive, plus an identification byte for each record. Each data stripe is divided into 20 segments, each containing up to 208 bytes of user data plus error correction and synchronization bits. The Mass Storage Control can resynchronize on a segment error so that only a maximum of 208 bytes of user data are lost.

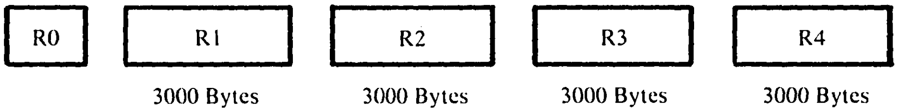

Example A. Track (t) cylinder (c) on staging drive (12,000 bytes).

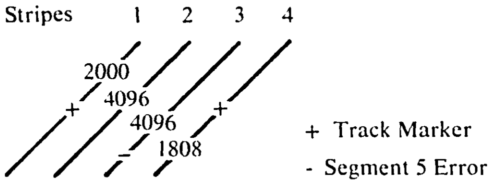

Example B. Track (t) cylinder (c) data stripes 1-4 on data cartridge (12,000 bytes).

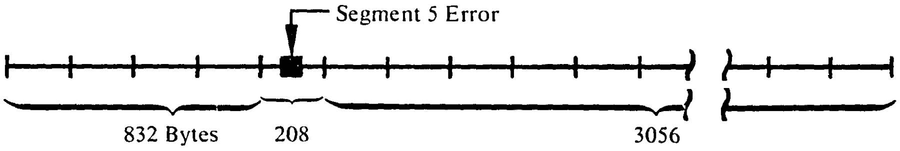

Example C. Data stripe 3 (4096 bytes) showing 208 byte segments and error.

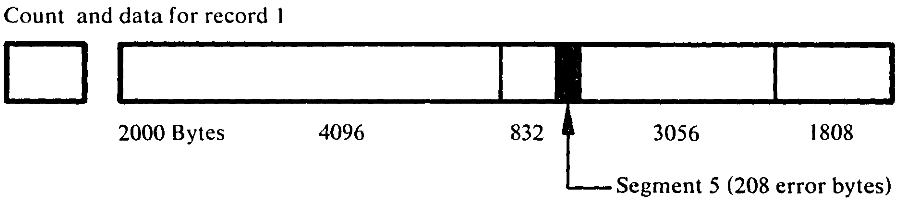

Example D. Track (t) cylinder (c) on staging drive after stage with error.

Figure 54. Data Cartridge Permanent Read Error

In Figure 54, Parts A and B illustrate a case where four 3000-byte records are recorded on a staging drive track and destaged without error to four data stripes. During subsequent stage operations, a permanent read error (media failure) occurs on stripe 3. Part C shows the 20 segments of stripe 3 data and the segment in which the error occurred.

In this case, the track is "staged with error" and the Mass Storage Control:

- Builds a new count field with the identification byte inverted
- Concatenates the data and stages a record indicating an error to a staging drive

Part D shows the 208 bytes of data in error, and the remaining 11,792 bytes of error-free data in the concatenated record. The record also has data inserted by the Staging Adapter, such as track separators and the ID bytes, which identify the beginning of count and data fields. Normally, permanent read errors are recovered using the data set backup copy and any of the methods previously described in this chapter. But, if the error track must be recovered without journals, copies, or generation data sets, it must be reconstructed from the concatenated record, which requires documentation or knowledge of the record and field layouts.

The reconstruction process involves:

- Display of the error track
- Input to rewrite the error track
- Salvage of any recognizable data
- Insertion of the correct or substitute data for the portion of data lost or in error

  Note: Assignment of an alternate track is not required.

An environmental data record of the error is written in SYS1.LOGREC. If the user program does not access the error track after the failure occurs, no record of the error exists except in SYS1.LOGREC. If the program does not modify any data in the cylinder, no destaging occurs and the data on the cartridge remains exactly as it was originally written. If the program writes data on other than the error track within a cylinder, the concatenated record is destaged from DASD without a disk read error. As the cylinder data is destaged, a data cartridge permanent write error occurs on the damaged data stripe 3. It is demarked and the data is written on the next data stripe (stripe 4), which prevents compounding the error.

Each subsequent usage of the cylinder data results in the error track being staged (and destaged if the cylinder data is written) until it is corrected or accessed by the user program. We recommend the use of SDA to display the environmental data provides detection of the error in case the user program does not access the error track.

Because the Access Method Services COPYV command can be used to copy this error record to back up volumes, use the REPAIRV command to initiate recovery action on all stage and destage errors when they occur.

Review user program exits that analyze error records to ensure that they can accept the Mass Storage System concatenated error record and to determine if the program reconstruction logic needs to be changed.

Data Cartridge Permanent Write Error: The Staging Adapter detects data cartridge permanent write errors during a destage operation. The normal action is to demark the error stripe and write the next data stripe. If demarking stripes reduces the number of stripes available to zero, the data cartridge becomes unusable for this destage operation, and error messages are sent to the operator and to SYS1.LOGREC. If the Relinquish order is issued with the Delayed Response attribute, a message is also sent to the user program. The REPLACEC command is used to complete the destaging on a new data cartridge and to copy all other cylinders from the old cartridge to the new cartridge. The Mass Storage Control tables are updated to reflect the new cartridge serial number, and the old cartridge is ejected from the Mass Storage Facility. If secure data is involved, follow the procedures for controlling such data on the old cartridge as described earlier. No data is lost because of this error situation.

Damaged Data Cartridge: A hardware malfunction that causes damage to a data cartridge can also cause data loss. Depending on the amount and type of damage, a tape leader or plastic shell can be repaired or replaced, and the data recovered. Your planning should cover the possibility of a data cartridge and all of its data can being destroyed, a situation we assume in the following description.

The effect of equipment failures on a user program depends on when the failure occurs. When a data cartridge is damaged as a result of a Mount or Acquire order, the program either aborts before execution, as it begins, or at Open time.

When the failure occurs as a result of a Demount order, the program has been executed. In either event, the recovery action depends on the extent of the damage.

A different condition exists when the data cartridge containing the mass storage volume VTOC cannot be loaded to identify the lost data. VSAM catalogs identify data for VSAM data sets, but other catalogs can identify lost data only at the volume level for non-VSAM data sets. Therefore, your plan of action is different for the two instances. Use of backup volumes and full volume recovery, similar to the techniques in use today, is your simplest approach, thus ensuring that all of the data sets associated with a mass storage volume can be identified. First, provide a replacement data cartridge so that the mass storage volume can be reestablished; activate the mass storage volume copy; and remove the damaged cartridge. Second, you must bring the data sets up to the current level, making sure that backup data (down-level generations or journals) is not written on the same mass storage volume as the original data.

Use the COPYV command to make mass storage volume copies. The copy volume is always created as an inactive volume and recorded in the MSVI data set. The two cartridges making up the mass storage volume are inventoried in the Transient Volume List maintained by the Mass Storage Control. Neither MSVI nor the Mass Storage Control allow the copy to be mounted or modified without deactivating the original volume and activating the copy. Use Access Method Services commands to move the volumes between the active and inactive volume inventories. You can retain the copy volume in the Mass Storage Facility itself or remove it and store it somewhere else.

If your computer installation is large and operates in a dynamic environment, you may want to consider data set recovery for individual application systems. The common requirements in this case are that you identify the lost data and notify the responsible user. You should also incorporate recovery techniques into the application system. To provide positive identification of lost data, we recommend the collection of SMF data set records and Mass Storage Control Trace data. This kind of information can help you create a data base that can be used for data recovery purposes.

Staging Drive Read Errors: Read errors that occur on staging drives and are detected by the System Control Program are treated as normal disk errors, except that additional sense bytes are collected for the staging drives. Here are several disk read error examples:

1. A program is writing and reading data on a mass storage volume with no intervening destage operations (LRU activity) when the System Control Program detects an uncorrectable read error. It can either accept the data containing the error, or abend. The Staging Adapter reports the error to the CPU (Unit Check), and records it in its Page Status Table. A second access to the real page causes the data written on the cylinders to be destaged. The real page is saved by staging the data to another page in the staging drive group. The action is the same as if a cylinder fault had occurred; the difference is that the Mass Storage Control marks the real page in error as unusable. To the CPU, the defective DASD track

appears as being repaired. When it is convenient for you, vary the affected staging drive SSID offline, reinitialize the staging disk pack, and reclaim the error page. If the final disposition is to delete the data, the Mass Storage Control will never encounter any disk read errors.

2. A program writes data on, but does not read data from, a mounted mass storage volume. The program terminates normally unless the Relinquish order with the Delayed Response attribute (VSAM DESTAGEWAIT option) is issued. With or without these options, the Mass Storage Control detects the disk read error when the data is read for destaging.
3. Existing data is staged to the staging drive without a disk write error being detected (the Mass Storage Control does not use Write Checks). The program accesses the data and gets a read disk error, which it can accept, bypass, abort, or take any current disk read option. A second access to any record on that real page causes the same action as described in the example 1, except that the data in the cylinders has not been modified. Destaging is not required and the real page is saved by staging the data from the mass storage volume to another real page in the same staging drive group.

A cylinder of data is the smallest increment handled by the Mass Storage Control for stage and destage operations. When a staging drive read error occurs during a destage operation, the error is logged in SYS1.LOGREC, the operator is notified, and the data is destaged with error. If a program has issued the Relinquish order with Delayed Response attribute, it is also notified so that it can take action. On a subsequent stage of the data that contains the error, the program encounters the error condition as it originally existed, but now it is on a different staging drive page. If the error is not in a DASD count field, you can apply any data recovery methods except alternate track assignment; if the error is in a DASD count field, the destaged data's record length is incorrect. In that case, recovery methods include formatting of the track.

Staging Drive Write Errors: Because the Mass Storage Control does not use Write Checks, any staging drive errors encountered during writing are caused by the inability to read the home address or count fields. If such errors are encountered, operations such as mount, cylinder fault, or acquire, which initiate the staging, will fail.

For example, the Mount order is completed (Device End presented) before cylinder 0 is completely staged or the read count field error is detected. Therefore, a normal pack-change interrupt is presented without validating cylinder 0 or presenting Unit Check. When the System Control Program attempts to read the volume label, the Staging Adapter encounters a cylinder fault condition during the seek operation because cylinder 0 was not validated. The Staging Adapter then presents Unit Check. If staging is initiated by a cylinder fault, the read count field error is detected during staging and Unit Check is presented. If staging is initiated by an Acquire order, the read count field error is detected during staging and a message is sent to the operator.

In all cases, the Mass Storage Control records the volume error in the queue control block (QCB) of the Schedule Queue Table. To recover, use IEHDASDR to assign an alternate track. The failing staging disk pack must be disabled and moved to a real drive. When the track assignment is complete, the staging disk pack can then be returned and made ready. The Mass Storage Control then retries the QCB that is in error.

Staging drive errors encountered by a program can be treated in the same way as you treat disk errors today.

Staging Drive Equipment Failures: In the current disk environment, equipment failures which do not damage the disk, can be bypassed by moving a disk pack from drive to drive within the same 3333 string, or across 3333 strings. In the Mass Storage System, moving staging disk packs within the same 3333 string requires moving the disk pack plus the logical address plug. Packs can also be moved across 3333 strings, if the 3333s are attached to the same Staging Adapters, and require moving only the disk pack and using the reconfiguration capability discussed next.

These conditions are required for reconfiguration:

1. The Mass Storage Control must receive a pack-change interrupt.
2. The staging disk pack SSID in the volume label must not match the drive SSID.
3. The SSID of the "moved from" and "moved to" drives must be offline or in neutral status.
4. The staging disk pack must not contain either the primary or the secondary Mass Storage Control tables.

The first host to access the failed staging drive receives an Equipment Check that causes the Staging Adapter to mark the drive as unusable. In turn, this causes the Mass Storage Control to place the staging drive in neutral status. Subsequent host accesses to the unusable drive receive Intervention Required.

After the staging disk pack has been moved within (or across) a 3333 string, it is varied online. All outstanding requests for intervention are satisfied and the Mass Storage Control retries any stage or destage operations that were interrupted.

You can always reconfigure the staging drive group, but but a staging drive must be made available before the staging disk pack can be moved. A real drive cannot be used as the "moved to" drive. The net result of the error is that only one program is affected, and the number of stage and destage operations temporarily increases while moving the data from the varied-offline staging drive to the spare or substitute drive. If more than one pack has to be moved within a staging drive group, the stage and destage operations could become more complex.

If head-disk interference damages the disk surface, data is lost, thereby making data recovery more difficult. You can still recover from this situation, as described under "Mass Storage Control Tables," later in this chapter. The effect of losing VTOCs and other data within a staging drive page in this type of failure is the same as with real disk, except that a staging drive can contain VTOCs for several mass storage volumes.

Complete protection from this type of failure can justify a data recovery base built from SMF and Trace data, and we suggest you take these actions to protect your data against compound failures:

- Continuously collect all Trace and SMF data.
- Checkpoint the Mass Storage System and update it on a periodic schedule, as described later in this chapter. Analysis of Trace or SMF data can identify low-activity periods when these checkpoints can be done with minimum interference with your operation. When the Mass Storage Control is reinitialized without disturbing its tables, the status of all staged data is known without affecting data reuse.
- As part of the checkpoint procedure, copy the primary and secondary Mass Storage Control tables to their alternate areas using the Mass Storage Control Table Create program, or IEHDASDR without copying the dynamic tables. Either creates the alternate tables so that they can be used as fall-back in extreme error situations.
- Make offline copies of the Mass Storage Control tables from their alternate areas and store them in a secure area.
- Keep an audit trail of all data cartridge entries and exits and of all mass storage volume definitions, deletions, renamings, and related activities. The fall-back tables previously saved do not contain any of these changes.
- Retain all Mass Storage System messages; they provide a data recovery base for handling any error conditions that might occur, with a minimum loss of time and data.

Information Sources

Because the Mass Storage Control handles staging space and mass storage volumes, and controls hardware status, it depends on the its tables and extensive microcode control that is external to the System Control Program. SDA, SMF, and the Trace function are information sources for collecting information about these activities. The messages at the operator console constitute another source of information. With the 3850 Usability Enhancements (OS/VS1 Release 5) or (OS/VS2 ICR), selected sense data is converted to prose before it is routed to the operator. See the following publications for related information on the SCHEDULR and IODEVICE macros, and Operator Control functions.

- *OS/VS1* or *OS/VS2 System Programming Library: System Generation Reference*
- *Operator's Library*: *OS/VS1* or *OS/VS2 Display Console*

3850 Console Exit Routine

We recommend a hard copy console, hereafter referred to as a 3850 Console. to list WTO/WTOR operator and space management messages to be used for recovery and audit trails. A user-written exit routine is required to support it. The following publications contain the documentation needed for this task:

- *OS/VS1 Planning and User's Guide*, or *OS/VS2 Planning Guide for Release 2* contains information about the WTO/WTOR macros, the user exit structure, and the system generation options in the section on Message Routing Exit Routines.
- *OS/VS1* or *OS/VS2 Supervisor Services and Macro Instructions* contains definitions of the routing and descriptor codes in the section on Message Routing for Multiple Operators' Consoles.

The user-written exit routine must analyze the WTO/WTOR messages and add a unique installation-assigned code (13, 14, and 15 are available for customer use) in the routing code field of the message. Adding a unique code is preferred to replacing the released routing code. In addition, the MSFMSG parameter in the SYS1.PARMLIB member, MVIKEY00, which controls the routing of the MSVC space management messages should be initialized with the same routing code specified in the WTO/WTOR exit routine. A decision should be made as to whether the MSVC messages should also be recorded in the MSVCJRNL data set.

We recommend scanning for and routing the following message IDs and information:

- "ICBddd", and "ICB ddd". All messages with an ID of ICB, issued by the Mass Storage System Communicator and MSVC components, should be routed to the Mass Storage System Console. Note that some messages IDs may be offset with two heading blanks. The "ddd" suffix is a three-digit decimal number, 000-399 for Mass Storage System Communicator tasks and 400-499 for MSVC tasks. ICB5xxx represents Formatted Sense messages, where xxx is two digits followed by an alphabetic character. ICB9xxx represents Mass Storage System Communicator initialization messages, which include data cartridge entry messages.
- "IDC083E", "IDC783E", and "IDC791E". These three messages are issued by the Access Method Services that support the Mass Storage System. Because they are currently the only "IDC" operator messages issued, a scan for "IDC" is sufficient.
- "IEA000Iddd", "IEA000I SWITCH", and "IEA000I SYS SWITCH. If the "IEA000I" message is not for a "SWITCH" or "SYS SWITCH", then the "ddd" field must be checked for a Mass Storage Control or 3330V address. This check is done by scanning the UCB lookup table or comparing stored constants in an appropriate exit routine table.

The following publications contain detailed information on messages, which you may wish to consult:

- *OS/VS Message Library: Routing and Descriptor Codes* describes the codes for all message IDs
- *OS/VS Message Library: Mass Storage System (MSS) Messages* describes the message IDs of ICB, and the reason codes
- *OS/VS Message Library: VS1 System Messages* describes the message IDs of IDC and IEA for OS/VS1
- *OS/VS Message Library: OS/VS2 System Messages* describes the message IDs of IDC and IEA for OS/VS2

A sample of a user-written exit routine, IEECVXIT, can be found in SYS1.SAMPLIB.

Using the 3850 Console

In installations where the Mass Storage System is attached to a single CPU, APU, or MPCPU, all operator and space management messages can be routed to a single 3850 Console using a WTO/WTOR exit routine. In installations where the Mass Storage System is attached to multiple CPUs, APUs or MPCPUs, JES3 offers the only released support for routing messages from multiple hosts to a single 3850 Console.

If provisions are made to connect the 3850 Console to the primary host, the following procedures can be used to maintain a hard copy history:

- A Mass Storage System WTO/WTOR exit routine must be included in the host control program
- The MSFMSG parameter in the SYS1.PARMLIB member, MVIKEY00, which controls the routing of space management messages, must be coded in all hosts to route these messages to MSVCJRNL. For the primary host, which has the 3850 Console, these space management messages must be routed to the 3850. Console as well as the MSVCJRNL.
- All Access Method Services commands that affect the data cartridge inventory must be executed on the primary host.
- All Mass Storage System operator commands that can optionally be entered on any host console must be entered only on the primary host console.

The 3850 Console listing together with the MSVCJRNL, provides a chronological history of all routable messages, with the exception of job-related operator messages, which must be entered on the secondary hosts. The more important job-related messages are recorded in the SYS1.LOGREC for each host. System Data Analyzer (SDA) reports can also be used to list them from either SYS1.LOGREC or from IFCEREP history, or both.

Usability Tools

The skills, knowledge, and planning required to manage and operate the Mass Storage System are comparable to those required for the System Control Program, except that more effort is spent on the planning phase.

We have provided several new tools to assist problem determination and correction. We have added a trace table to the MSSC resident control block to store Mass Storage Control orders for analysis. A function called SVC

126 Request Block Trace stores all bytes of the last order plus the first 16 bytes of the previous 16 orders. In addition, programs such as IEHDASDR, IEBGENER, HMASPZAP or AMASPZAP, IEBCOMPR, and IEHPROGM can access the Mass Storage Control tables.

The Access Method Services commands used by the Mass Storage System are:

- AUDITMSS audits the presence or absence of data cartridges in the Mass Storage Facility.
- CHECKMSS checks portions of the Mass Storage Control tables, Staging Adapter tables, and MSVI data set for consistency.
- COMPARET compares the Mass Storage Control primary and secondary tables.
- COPYT copies the primary Mass Storage Control controlled tables.
- DUMPMSS dumps the Mass Storage Control and Staging Adapter tables.
- LISTDSET lists information for non-VSAM data sets.
- LISTMSF lists the data cartridge and mass storage volume contents of the Mass Storage Facility.
- LISTMSVI lists information about mass storage volumes and mass storage volume groups.
- MODIFYC provides the capability to change data cartridge labels and move data cartridges from one location to another with or without the use of a data recording device.
- NULLIFYC removes all records pertaining to a damaged data cartridge.
- RECOVERV restores data from a copy to an active mass storage volume.
- REMOVER deletes one or more mass storage volume records from the MSVI data set.
- REPAIRV provides the capability to locate, list, and correct stage and/or destage errors.
- REPLACEC replaces damaged or defective data cartridges assigned to a mass storage volume
- SWAPT exchanges the primary and secondary designations of the Mass Storage Control tables.
- TRACE starts and stops the recording of Mass Storage System activity and prints formatted reports.
- TUNE changes the LRU parameters.

These commands and their use are described in the *OS/VS Mass Storage System (MSS) Services: General Information* and in the *OS/VS Mass Storage System (MSS) Services: Reference Information.* Functions described in these publications can be used to back up and recover data, for example, COPYV and RECOVERV.

Mass Storage System Data Sets

The Mass Storage System's microcode logic is distributed in the Mass Storage Control and the Staging Adapters. The status of all operations is maintained in sets of tables similarly distributed. The Mass Storage Control's primary and secondary tables and their alternate areas reside on staging drives in staging drive group 00. They are similar in importance and function to the OS/VS system-residence packs.

The Mass Storage Control is responsible for writing and maintaining the contents of the Staging Adapter tables with one exception: each Staging Adapter keeps track of all cylinders under its control, which have had data written on them. This information is then transmitted to the Mass Storage Control on request.

Mass Storage Control Tables

Figure 55 shows the primary and secondary Mass Storage Control table names, and their alternate areas. They are on two 3333 Disk Storage and Controls, that is, four drives, which is the minimum number of staging drives in a Mass Storage System configuration.

The Mass Storage Control manages the Mass Storage System resources and dynamically responds to requests for mass storage volumes and data set extents issued by each CPU. All of the information needed by the Mass Storage Control is contained in the primary and secondary Mass Storage Control tables. All table information is written on, and read from, the primary tables; the secondary tables are only read at initialization time, but they are written whenever the primary tables are written. If both Staging Adapters 0 and 1 are available, all write operations to the primary Mass Storage Control tables are through Staging Adapter 0, and all write operations to the secondary Mass Storage Control tables are through Staging Adapter 1.

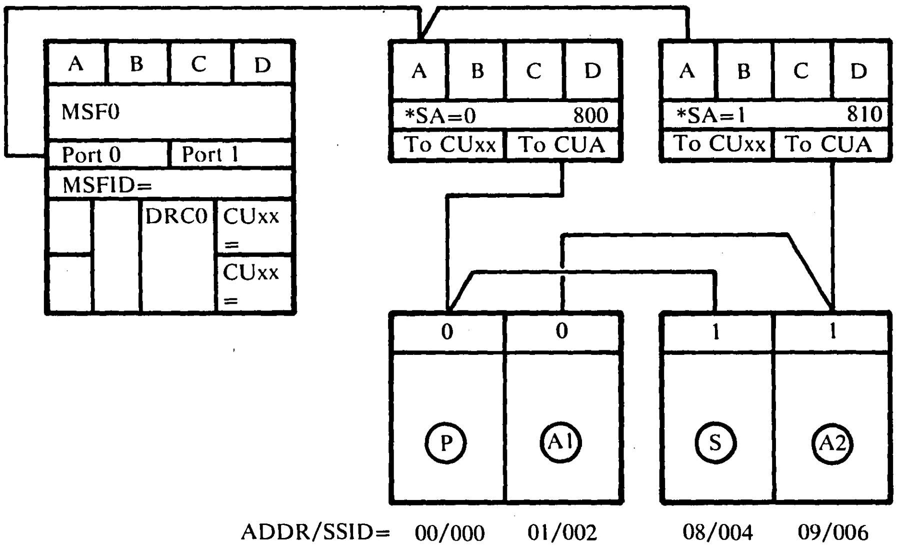

Configuration Table
Diagnostic/EC/Overlay Table
Cell Map
Message Buffer
Mounted Volume Table
Recovery Journal
Page Status Table Image
Schedule Queue Table
Scratch Cartridge List
Staging Drive Group
Trace Areas
Transient Volume List
Virtual Volume Address
VVA/VOLID Cross Reference
Verification Table
Volume Inventory Table

P= Primary MSC Table read and written by the MSC.

S= Secondary MSC Table written by the MSC.

A1 and A2= Allocated alternate areas for primary or secondary tables depending on recovery options.

*Staging Adapter in 3830-3 or ISCSA.

Figure 55. Mass Storage Control Table Names

Figures 56 and 57 list the Mass Storage Control and Staging Adapter tables and their status relative to data recovery, that is:

- *Controlled* tables change as a result of an action affecting the data cartridge or Mass Storage Volume Inventory data set. This action can include manual entries and exits, and the Access Method Services commands. These tables can be restored to current status if you maintain an audit trail. You create the audit trail from checkpoints of the Mass Storage Control tables, the MSVI, and by collecting the messages issued when scratch cartridges are entered into the system. The Diagnostic/EC/Overlay Table is changed only when engineering changes (ECs) are applied to the microcode. These ECs should also be retained for audit trail purposes.
- *Dynamic* tables change with system activity. Most Mass Storage System activity, such as the Mount, Acquire, Relinquish, and Demount orders, LRU, and space allocation, are included in this category.
- *Static* tables that contain relatively stable information and they are not critical to data recovery. This category includes configuration and mount equalization information that can be set to the initial Mass Storage Control Table Create program values without losing data integrity. Operations such as vary on, or values of the Tune command might need to be repeated or reentered.

For further information on table content, see *IBM 3850 Mass Storage System (MSS) Principles of Operation.*

Mass Storage Control Table Names	Information		MSC Table Create Update Action		
	Type	*Status	Reconstructed	Copied	Reinitialized
Configuration Table	Configuration	Static	X		
Diagnostic/EC/Overlay Table	Microcode	Controlled		X	
Cell Map	Configuration	Controlled	X	X	
Message Buffers	Activity	Dynamic			X
Mounted Volume Table	Activity	Dynamic			X
MSC Recovery Journal	Activity	Dynamic			X
Page Status Table Image	Activity	Dynamic			X
Schedule Queue Table	Activity	Dynamic			X
Scratch Cartridge List	Media	Controlled		X	
Staging Drive Group Table	Configuration	Dynamic	X		
Trace Areas	Activity	Dynamic			X
Transient Volume List	Media	Controlled		X	
Virtual Volume Address Table	Activity	Dynamic			X
VVA/VOLID Cross-Reference Table	Activity	Dynamic			X
Verification Table	Control	Static	X		
Volume Inventory Table	Media	Controlled		X	

*See definition of terms in text

Figure 56. Mass Storage Control Table—Type, Status, and Update

Staging Adapter Table Names	Information	
	Type	*Status
Virtual Address Directory	Activity	Dynamic
Virtual Volume Information	Activity	Dynamic
Page Status Table	Activity	Dynamic
Least Recently Used	Activity	Dynamic
Address Using	Activity	Dynamic
Logical to Real	Configuration	Static

*See definition of terms in text

Figure 57. Staging Adapter Tables—Type and Status

Direct Access to Mass Storage Control Tables

The Mass Storage Control tables can be accessed directly by the CPU through interface B of Staging Adapters 0 and 1, and optionally through interfaces C and D. The Mass Storage Control Table Create program automatically provides the IODEVICE cards (device addresses 00, 01, 08, and 09) for direct access through interface B. If direct access is desired through interfaces C and D, IODEVICE cards must be punched for device addresses 00, 01, 08, and 09 on their respective channels.

The Customer Engineer must plug the circuit card for Staging Adapter 0 and 1 for one of the following direct access options:

- Interface B only
- Interfaces B and C
- Interfaces B, C, and D

The Mass Storage Control Table Create program creates a two extent data set on each tables pack, called SYS1.MSCTABLE.volser, with the CC HH RR in hexadecimal notation as follows:

- Extent 1 is 07 01 01 - 07 01 10 and contains the Verification Table.
- Extent 2 is 08 00 01 - 27 12 20 and contains the remaining Mass Storage Control tables.

The Format 5 DSCB (Data Set Control Block) shows that the rest of the space on each volume is unavailable.

The Mass Storage Control Table Create program writes the SYS1.MSCTABLE.volser data set on the physical cylinders shown in Example A of Figure 58, but the VTOC points to the virtual cylinders shown in Example B. Direct access to the Mass Storage Control tables can only be done through the virtual cylinders that are mapped to the real cylinders by the Staging Adapter. Access requests to any cylinders outside the range of the SYS1.MSCTABLE.volser data set receive Command Reject. Movement of a tables pack to a 3330 real drive to access the tables results in a count field error.

To prevent destruction of the Mass Storage Control tables, the Mass Storage Control Table Create program turns on the read/write protection indicator in Format 1 DSCB. Use IEHPROGM to update the PASSWORD data set with the appropriate key before to the first data set is opened. An abnormal termination occurs when an attempt is made to open a data set without assigning a password.

PRESRES (OS/VS1) and VATLST (OS/VS2 MVS) must be updated to include the four Mass Storage Control table volumes. The entries should be specified as reserved and private. Only VOL=SER can be used to access non-3330V volumes in PRESRES and/or VATLST. The OS/VS JCL to access the SYS1.MSCTABLE.volser data set must include the following DD statement:

```
//MSC DD  DSN=SYS1.MSCTABLE.volser,
//        DISP=OLD,
//        UNIT=3330,VOL=SER=volser,
//        DCB=(RECFM=F),BLKSIZE=264,
//        DSORG=DA,PS,[OPTCD=[A][W]]
```

In addition to the usability tools, we have tested the following utilities for use with the Mass Storage Control tables.

- IEBCOMPR to compare the primary and secondary tables using data set names
- IEHDASDR to dump, restore, and print selected table areas using CHRs
- IEBGENER to dump and restore tables using data set names
- IEHPROGM to update or verify the password protection
- HMASPZAP or AMASPZAP to dump or alter tables using CHRs

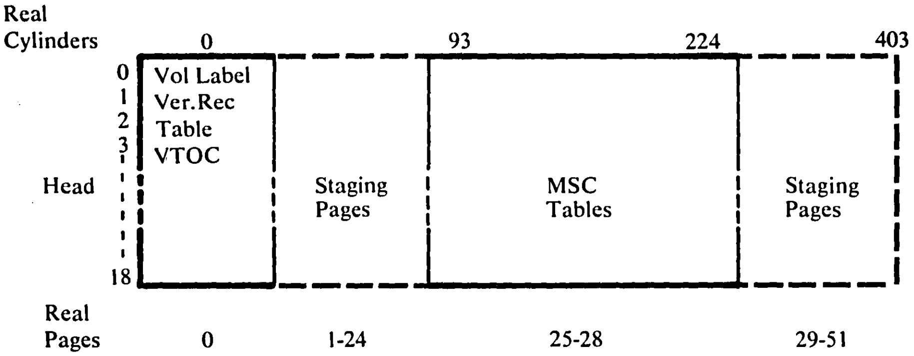

A. Physical View

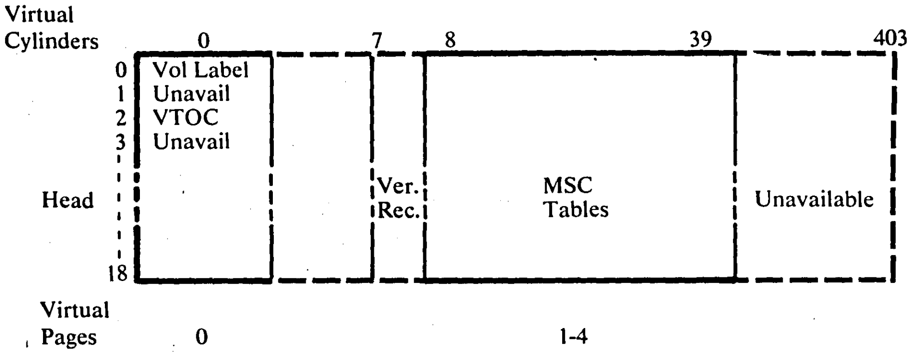

B. Logical View

Figure 58. SYS1.MSCTABLE.volser Data Set

Initial Mass Storage Control Tables Backup

The preparation for backing up the Mass Storage Control tables begins at installation time, and includes:

- Executing IEHDASDR to format the staging space on each staging disk pack. The first four characters of the volume serial number on all staging volumes must be the same, and are chosen by you. The last two characters of the volume serial number must be the same as the last two digits of the staging drive SSID, that is, 00 for the primary (P), 02 for the alternate primary (A1), 04 for the secondary (S), and 06 for the alternate secondary (A2).
- Executing the Mass Storage Control Table Create program twice to format a table area on each of the four tables staging volumes with the Verification Table on track 1 of real cylinder 0. The 32 cylinders used for the tables are on virtual pages 1, 2, 3, and 4 (real pages 25, 26, 27, and 28).
- Retaining the Mass Storage Control Table Create program output, worksheets, and Host Connections Report as a base record to be used by you, the Customer Engineer and the Program Support Representative. All microcode overlays and ECs are written on the disk packs by the Customer Engineer using an Online Test (OLT) program. Retain the input cards or flexible disks containing the changes for data backup and for updating the backup copies of the Mass Storage Control tables. To save time initially, the Customer Engineer can install the microcode on the primary and secondary tables packs before the second Mass Storage Control Table Create program run.
- Mounting staging volumes xxxx00, xxxx02, xxxx04, xxxx06, on staging drive SSIDs 000, 002, 004, and 006 respectively.
- Initializing the MSVI and MSVCJRNL data sets. Details about this task is given in *OS/VS Mass Storage System (MSS) Services: General Information.*
- Initializing the Mass Storage Facility, the Staging Adapters IML, and the system IPL.
- Loading data cartridges and defining mass storage volumes (and mass storage volume groups). All manual data cartridge entries or Access Method Services ejection, creations or deletions, and additions or storing of mass storage volumes cause changes in the primary and secondary Mass Storage Control tables. These changes appear as entries in the MSVI and MSVCJRNL data sets, except for the scratch cartridge entries or exits. The MSVC must be activated when the above activities are being performed.

At this point, as shown in Figure 59, the six data sets or backup areas have been created and operations can proceed.

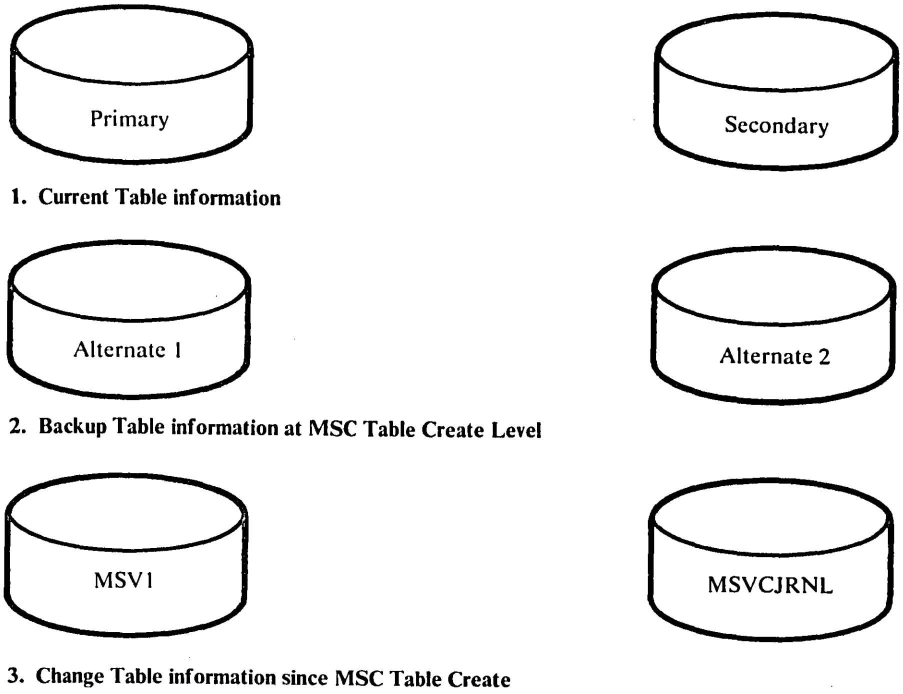

Figure 59. Mass Storage Control Tables Backup

Mass Storage System Checkpoint

This section describes a Mass Storage System checkpoint procedure, that is, all mass storage volumes demounted, all changed data destaged, and all data cartridges in their cells. This can be accomplished by a HALT S,LONG command from the primary host and a HALT S,SNAP command from all other hosts. Demounting of all mass storage volumes can be accomplished by holding back all job classes (or 3850 job classes) and allowing all other jobs to complete. Destaging of data can be forced by varying all staging drive SSIDs offline. You may elect to do a checkpoint once and keep it current by updating only those portions that are affected by changes.

A Mass Storage System checkpoint includes:

- Executing the Access Method Services CHECKMSS, LISTMSF, and LISTMSVI commands
- Deactivating the MSVCJRNL data set to ensure no update activity is in progress, and copying it on magnetic tape. We recommend that you use IEHDASDR, because absolute track pointers are used in the MSVCJRNL data set.
- Copying the MSVI data set to magnetic tape using the EXPORT command. You now have two levels of data backup. By using the IMPORT command, the latest MSVI backup can be used for data recovery, or else the previous MSVI backup, if the latest MSVJRNL backup is used to update the MSVI data set.

- Printing the space management messages and nullifying the MSVCJRNL data set.
- Copying the primary and secondary Mass Storage Control tables to their respective alternate areas using the Mass Storage Control Table Create program with the UPDATE option. Copying the primary and secondary tables is a precaution to ensure that a valid copy exists. The UPDATE option initializes the dynamic tables, as shown in Figure 56, so that they cannot be restored to some previous activity level and cause data integrity problems. The tables packs must be mounted on real drives (or convertible-to-real drives varied online as real).
- Copying the alternate Mass Storage Control tables to magnetic tape or disk using IEHDASDR. This additional precaution creates an offline copy for storage in a safe location.

You now have a complete Mass Storage System checkpoint and have verified that all four tables packs can be read. The following conditions exist:

- All data cartridges are in their cell locations.
- All changed data has been destaged.
- The primary and secondary tables have been verified and can be read without error.
- The alternate areas contain copies of the primary and secondary tables, (controlled and static). If either alternate area is used as primary, all inactive data on the staging drives is unknown to the Mass Storage Control and cannot be used by the data reuse function, because the dynamic tables have been reinitialized.
- An offline copy of the Mass Storage Control tables has been saved.
- The MSVI data set has been checkpointed.
- The space management messages printed from the MSVCJRNL data set are an audit trail of all defined mass storage volumes.
- A record has been made of all data cartridges entered in the Mass Storage Facility and their locations.
- All 3330V volumes are demounted to OS/VS.
- All inactive data on staging drives is known to the Mass Storage Control and can be used by the data reuse function.

If all staging drive SSIDs are varied offline, use the opportunity to assign alternate tracks to staging packs that have real pages marked unusable. The VARY SSID, OFFLINE, S command clears the staging pack and the Mass Storage Control tables entries for that pack. Therefore, if data is destroyed during the alternate track assignment or the staging drive must be replaced, the Mass Storage Control does not attempt to reuse the data on the staging pack (or its replacement) when the drive is varied back online.

Checkpoint Update

Once a complete Mass Storage System checkpoint is established, it can be kept current by updating only the portions that change. This requires control of activities that affect the checkpoint so that an update can take place after each change.

The Mass Storage Control controlled tables can be updated using IEHDASDR to copy the appropriate primary and secondary tables to their alternate areas. The controlled table names and their beginning and ending cylinder and head addresses (CC HH) in hexadecimal are:

• Scratch Cartridge List Volume Inventory Table Transient Volume List	13 02 → 1A 07
• Cell Map	1B 01 → 1B 01
• Diagnostic/EC/Overlay	1B 06 → 1B 0A
	1C 08 → 1F 12

As an additional precaution, the alternate areas can be copied to magnetic tape using IEHDASDR:

• Volume Label	00 00 → 00 00
• Volume Table of Contents	00 02 → 00 02
• MSC Tables	07 01 → 07 01
	08 00 → 27 12

Execute the Access Method Services LISTMSF and LISTMSVI commands. Optionally, checkpoint the MSVI and MSVCJRNL data sets.

The Mass Storage Control is not affected by the MSVC mass storage volume groups and the associated Access Method Services commands, such as CREATEG, MODIFYG, and SCRATCHG. They only affect the Inventory data set and do not cause changes in the Mass Storage Control tables. Figure 60 shows the Access Method Services commands and the data sets affected by each command. See *IBM 3850 Mass Storage System (MSS) Principles of Operation* for a more detailed description.

Function		MSVI	MSVCJNRL	SCL	VIT	TVL	CMT	Note
ADDV		X	X		X	X		
COPYV		X	X	X		X	X	1
CREATEV		X	X	X	X			
EJECTV		X	X			X	X	
MODIFYV	(NEWSERIAL)	X	X		X			
	(OWNER)	X	X					
	(GROUP)	X	X					
	(GEN/RSTR)	X	X					
	(BIND/NOBIND)	X	X		X	X		2
	(EXCL/SHARED)	X	X		X	X		2
	(RONLY/RW)	X	X		X	X		2
	(DERAS/NDERAS)		X	X		X	X	2
	(DESCRIPTION)		X	X				
	(BKUPNO)		X	X				
	(NULLIFY)		X	X				
RECOVERV		X	X					
REPLACEC		X	X	X	X		X	
SCRATCHV		X	X	X	X	X		3
STOREV		X	X		X	X	X	1
CONVERTV								
REMOVEVR		X	X					
CREATEG		X	X					
MODIFYG		X	X		X			4
SCRATCHG		X	X					
EJECTC				X			X	
LISTMSVI								
LISTMSF								
TRACE								5
TUNE								6

MSVI =Mass Storage Volume Inventory
MSVCJRNL =Mass Storage Volume Control Journal
SCL =Scratch Cartridge List
VIT =Volume Inventory Table
TVL =Transient Volume List
CMT =Cell Map Table

Notes:

1. Cell Map Table is affected only if the volume is ejected.
2. Transient Volume List is affected if the volume has copies.
3. Either table may be affected depending on whether it is an active or inactive volume.
4. The Volume Inventory Table is modified only if the volume attributes are changed.
5. Affects Trace area usage by Mass Storage Control.
6. Affects Verification and Staging Drive Group Tables.

Figure 60. 3850 Data Affected by Access Method Services Commands

As you develop partial checkpoint procedures, you might use these guidelines:

- Checkpoint a specific table after an operation affecting that table is completed. The operation cannot take place at the same time as the checkpoint.
- Control the entry, exit, and definition of mass storage volumes at a central location, and checkpoint the tables after the operations are completed.
- Execute a complete Mass Storage System checkpoint before reconfiguring the Mass Storage System, or after ECs are applied.
- Retain all message information as second-level data backup in case the primary and secondary tables, and their checkpointed alternate areas are all somehow destroyed.

These checkpoint procedures give you many levels of protection. In addition, the Mass Storage Control tables can be completely or partially recovered as follows:

- If the primary or the secondary table is lost, complete recovery is possible using the COPYT command to create a new primary or secordary table, as described in the *OS/VS Mass Storage System (MSS) Services: General Information*.
- If the primary and secondary tables are lost, partial recovery (minus the dynamic tables) is possible using the Mass Storage Control Table Create program with the UPDATE option to create new primary and secondary tables from the alternate areas.
- If the primary and secondary tables, and their alternate areas are lost, partial recovery (minus the dynamic tables) is possible using IEHDASDR to create new primary and secondary tables, and alternate areas, from the offline tape copies.

It may be necessary to change the volume label with the desired VOLID, and rename the table data set in the Volume Table of Contents to match the VOLID.

Loss of the dynamic tables requires other recovery actions, which are described under "Compound Data Checks".

If data recovery is required, do the following to restore the MSVI data set:

- Inactivate the MSVC functions.
- Use the IMPORT command to restore the MSVI data set from the the backup copy.
- Restore the MSVI data set to current status using the latest MSVJRNL data set and the MSVI function called RESTORE.
- Activate the MSVI functions.

Resuming Normal Operation

When Mass Storage System work is stopped for an update, or a HALT S,LONG or HALT S,SNAP command is issued, subsystem operation can be resumed without a system IPL, and Trace data should be collected immediately after operation is resumed.

Mass Storage Control Tables Verification at IML Time

The first 64 words of the microprogram are loaded from the Type 23 Flexible Disk (23FD) under control of the Mass Storage Control each time there is a power-on sequence and the Mass Storage Control's CE panel Operation Mode switch is in the Normal, Forced Logging, or CE Normal position. The microprogram load can also be initiated manually by operating the Operator's panel IML switch if the Mass Storage Control has been reset by the MSC Clear switch. Only one Mass Storage Control can be in CE mode at a time.

If both Mass Storage Controls are enabled, Mass Storage Control 1 begins the IML first. If the Mass Storage System configuration has two Mass Storage Facility Model A's or one Mass Storage Facility Model B, a switch is made to the alternate Mass Storage Control if an error occurs.

After the IML has been completed, the primary Mass Storage Control establishes a communication link with its table volumes through port 0. The Mass Storage Control must have a primary table (normally SSID 000) and a secondary table (normally SSID 004) in staging drive group 00.

The Mass Storage Control reads the Verification Table on SSIDs 000, 004, and possibly 002 and 006, to positively identify the current primary and secondary tables. Each Verification Table contains a primary and a secondary address of the Mass Storage Control tables. Positive identification of the current primary and secondary Mass Storage Control tables takes place when the Mass Storage Control finds two Verification Tables with the same primary and secondary addresses. The search begins on SSID 000. Assume that the Verification Table on 000 contains the primary address of SSID 000 and a secondary address of SSID 004. The Mass Storage Control then proceeds to read the Verification Table on SSID 004; if it contains the primary address SSID 000 and the secondary address SSID 004, a positive identification has been made. If the Verification Tables do not agree, the Mass Storage Control searches the four SSIDs until it finds two Verification Tables that are identical.

During IML, this process ensures that the correct primary and secondary tables have been located. The Mass Storage Control then verifies the status of all components in the Mass Storage System configuration, which remains the same as it was when the system was powered off: a unit that had been varied offline continues in that state until it is varied online, a system power off or an IML does not change the unit status.

Staging Adapter IML

Each Staging Adapter initiates an Initial Microprogram Load (IML) when it is powered on. Device addresses 00, 01, 08, and 09 on interface B of Staging Adapters 0 and 1 are always staging drives. All other physical drives are initialized as real drives until the Mass Storage Control sets bits in the Staging Adapter tables to indicate they are staging drives. Each staging drive VOLID is read to verify that it is a correct staging volume. if it is not, the Mass Storage Control places the drive's SSID in neutral status. The level of the 23FD, tables, and overlays on the primary tables pack is also verified.

System Control Program Initialization

At the system's Initial Program Load (IPL) time, the System Control Program locates the UCBs for the Mass Storage Control and allocates them to the Mass Storage System Communicator. The Mass Storage System

Communicator then determines the primary Mass Storage Control (and optionally, the alternate Mass Storage Control) and issues the Initialize order. Both OS/VS1 and OS/VS2 MVS require that the primary and secondary Mass Storage Control tables packs be available. If the System Control Program is OS/VS2 MVS, an Associate or Disassociate command is issued depending on whether the MPCPUs are in a multiprocessor or uniprocessor mode. If the System Control Program so indicates, the PURGE command is issued to the Mass Storage Control and the VOLID in all the UCBs is set to zero.

If the Mass Storage Control initialization fails, the IPL can continue without the Mass Storage System. If initialization does not fail, the System Control Program establishes the address of the Attention routine, attaches the MSSC Message Task, and issues the Host Ready for Messages order.

Mass Storage Control Failures

The Mass Storage Control table information is protected from read or write errors, Equipment Checks, or partial updates due to power or equipment failures by writing the information in the secondary tables and alternate areas. A recovery journal is contained in both tables, and before any table entry is changed, the old record is written in this journal. At the same time that a Mass Storage Control order is being executed, the Schedule Queue Table contains a schedule queue block for each item of work to be accomplished. After successful execution, the primary and secondary Mass Storage Control tables are updated and the Recovery Journal entry is nulled.

If a Mass Storage Control failure occurs, the alternate Mass Storage Control (or the primary Mass Storage Control after repair) determines the work status and ensures the integrity of the tables using the Recovery Journal and the Schedule Queue Table.

Permanent Data Check

Any permanent read error on the primary or secondary tables causes an interrupt and requires corrective action. It can occur during IML or normal operation.

Permanent Read Error During IML

If a permanent read error occurs during IML, it must be corrected before the operation can be completed. The system IPL can continue but must be completed so that the tables' corrective procedures can be used. The four Mass Storage Control table addresses allow access to the drives containing the tables through interface B (and optionally interfaces C and D) of Staging Adapters 0 and 1 without requiring the Mass Storage Control to be operational.

Device addresses 00, 01, 08, and 09 on interface B of Staging Adapters 0 and 1 are always staging drives; all other drives connected to a Staging Adapter accessible on interfaces B, C, and D are initialized as real drives. When the Mass Storage Control finishes verifying its tables, it sets the virtual bit in the Staging Adapter table if the drive address is specified as virtual. When the Mass Storage Control cannot complete the IML operation, all drives initialized as real can be used as real until the Mass Storage Control becomes operational. If a failure occurs immediately following a Mass Storage System checkpoint, the alternate tables are duplicate images. You can use IEHDASDR to correct the failing disk track, or an alternate track assignment can be made.

You have the option to use a copy from the Mass Storage Control secondary tables. This is possible only if the failing track is not track 1 of virtual cylinder 7 (real cylinder 0) that contains the Verification Tables. Another option you have is to make the tables' alternate areas the new Mass Storage Control primary and secondary tables, but only if you do it immediately following the Mass Storage System checkpoint.

You can use IEHDASDR to correct failures in the Verification Tables. Because the information therein is static, a backup data set can be prepared using the Mass Storage Control Table Create program output for each tables pack.

Permanent Data Error During Operation

When a permanent read or write failure occurs in the Mass Storage Control tables, the order being executed at the time fails, and any new orders sent to the Mass Storage Control also fail, except for COPYT.

Mass Storage Control primary tables failure. Copy the current secondary tables to alternate area 1 or to alternate area 2. These are then designated as the new primary tables.

Mass Storage Control secondary tables failure. Copy the current primary tables to alternate area 1 or to alternate area 2. These are then designated as the new secondary tables.

As soon as the Mass Storage System is back in operation and production workload permits, the checkpoint's update procedures should be executed. But, should another failure occur before a checkpoint is made, it could complicate data recovery.

Compound Data Errors

If multiple data errors on the tables areas occur simultaneously, such as a permanent read error while the tables are being copied, the Mass Storage System is disabled. IEHDASDR is then required to reconstruct a copy of the tables from the four tables and the backup copies.

If it becomes necessary to use the alternate tables (and they are current and valid except for the dynamic tables), all currently mounted and active cylinders (staged and valid) must be identified. You can use several approaches to do this, depending on the extent of the damage to the primary and secondary tables. Assuming both sets of dynamic tables are destroyed, you must:

- Identify all data that has been created or modified since the last Mass Storage System checkpoint.
- Verify the validity of the data in the cartridge. If you cannot positively determine that the data and its associated VTOC have been destaged, the verification is not positive.

You can determine the status of the data by answering the following questions:

1. Which volumes were mounted?
2. Which cylinders were staged?
3. Which cylinders were modified?
4. Which cylinders were destroyed?
5. Has the VTOC been destaged?

6. Which data sets occupied the staged, modified, or destroyed cylinders?

7. Who owns and uses the data sets?

8. Were the data sets created or updated?

If Trace data is collected on a periodic basis, the first five questions can be answered for the events that preceded the last data collection. If SMF data is available, questions 1, 2, 6, 7, and 8 can be answered, and question 3 can be deduced.

There are two potential levels of data recovery. If the dynamic tables are completely recovered, the status of the data can be determined: the only loss is data physically destroyed. If the dynamic tables are only partially recovered, the status of only some of the data can be determined: the loss is data physically destroyed plus data whose status is unknown. If Trace data or the Mounted Volume Table information is available, the dynamic tables can be recovered.

Once the tables have been restored, the MSVI data set and system UCBs must be synchronized with the Mass Storage Control tables. If the fall-back is to a point that does not restore the dynamic table information, then further recovery action is required. An IPL-with-purge operation can assist the recovery action by forcing the System Control Program to demount all mass storage volumes. The MSVC dynamically recovers space information as the mass storage volumes are mounted and demounted. If the mass storage volumes created and stored agree, these data sets are in synchronization.

An alternate 3333 path to the Mass Storage Control tables is provided in the minimum Mass Storage System configuration, which requires at least two 3333s. If spare staging drives are available, a disk pack and plug can be moved within the same 3333 string. If the 3333 is the problem, the COPYT command must be executed.

Power Failure

Because the Mass Storage Control maintains its tables on disk and is able to recover from partial updates, the only status information lost as a result of a power failure is within the Staging Adapter tables. The Mass Storage Control has the capability of restoring these tables as the individual Staging Adapters are varied back online, and the lost status information is limited to the cylinder written indicators in the Page Status Table. The Mass Storage Control assumes that every cylinder containing data that could have been written, has been written. The net result is that read-only, as well as written, data cylinders are destaged. No data loss occurs here, but the destaging workload is higher for some period of time following the power failure as the volumes are demounted.

Configuration Control

Introducing the Mass Storage System into your operating environment affects regression testing because microcode and configuration changes need to be tested. Deciding what the amount of regression testing should be relative to the degree of change is a matter that can only be learned with experience.

If job stream is modified to test the Mass Storage System functions, operations, and performance, then new releases of the Mass Storage Control or Staging Adapter microcode can be tested for correct operation. We require that all microcode releases and ECs be applied to your installation. We also require that you install new 23FDs within 90 days after their availability. The new microcode should be regression-tested against a job stream before you release it for production.

Emergency repairs are handled on an individual basis and coordinated by the Mass Storage System Support Center. Such repairs are tested and distributed to all users as required.

The Mass Storage Control's Trace reports can be used to compare the regression test performance before and after changes to the configuration or to the microcode are made. Although a change may not be made for performance reasons, verification of Mass Storage System component activity before and after the change can be of assistance in identifying problems.

Data Cartridge Management

This topic describes the physical management of the Mass Storage System data cartridge. The *IBM 3850 Data Cartridge Care and Handling Instructions* should be used in developing your plan, which should include the periodic execution of the LISTMSF command, especially when there is an unusual amount of data cartridge entry and exit activity.

Note that nine cell locations in a Mass Storage Facility are reserved for special uses, whose number, use, and addresses are:

- Six cells are reserved for storing CE scratch cartridges at cell addresses 01F7, 02F7, 03F7, 04F7, 05F7, and 06F7.
- Two cells are reserved for storing temporarily data cartridges if their assigned cells are full. For the right accessor, the cell address is 00F7.

  For the left accessor, the cell address depends on the Mass Storage Facility model: that is, address 00EA for models A1 and B1, 00CE for models A2 and B2, 00B2 for models A3 and B3, or 0096 for models A4 and B4.
- One cell, at cell address 07F7, is reserved for storing temporarily a data cartridge found in a data recording device during a vary online or initialize operation.

The total number of cell locations that you can use for your data cartridges is:

697 for Mass Storage Facility Models A1 and B1

2035 for Mass Storage Facility Models A2 and B2

3373 for Mass Storage Facility Models A3 and B3

4711 for Mass Storage Facility Models A4 and B4

Based on our installation experience, we recommend that some cell locations be reserved for your anticipated ejection and entry requirements.

Data Cartridge Care

Once they enter the Mass Storage Facility, data cartridges are automatically handled by the accessors and the data recording devices.

Whenever data cartridges or mass storage volumes are removed from the Mass Storage Facility, the labeling, handling, inspection, damage repair, and environmental requirements described in *IBM 3850 Data Cartridge Care and Handling Instructions* must be followed.

Do not enter damaged data cartridges into the Mass Storage Facility. If a damaged data cartridge contains important data, consult the Customer Engineer to determine which method you should use to copy the data cartridge.

Data Cartridge Performance

You should plan to monitor cartridge performance using the System Data Analyzer (SDA) problem program ISDASDA0 that runs under OS/VS1 or OS/VS2 MVS. It provides you with information about Mass Storage System data cartridge usage and tape errors. It also formats, and in some cases analyzes, the environmental error sense data logged during the Mass Storage System operation. It can be used as an aid to anticipate, diagnose, or isolate Mass Storage System failures. SYS1.LOGREC is the input to SDA, as described in *OS/VS1 and OS/VS2 System Programming Library: SYS.LOGREC Error Recording.*

The errors logged in SYS1.LOGREC provide a record of all equipment failures, selected software failures, and selected system conditions. The information is written by the system recording routines and can be retrieved by the IFCEREP0 service aid, and they are included as standard programs in OS/VS SYSGEN. Each record on SYS1.LOGREC contains complete and specific information for the device and type of failure, or the statistical condition that caused the record to be written.

Exception conditions in the Mass Storage System are also logged in SYS1.LOGREC. For example, for some types of Mass Storage System failures, an error threshold must be exceeded before the errors are logged. When a data cartridge error threshold is exceeded during a load operation, the usage statistics for the cartridge are transferred to SYS1.LOGREC.

The SDA program accumulates data cartridge error statistics on units and cartridges and lists them in the Data Cartridge Statistics Report. Whenever a data cartridge serial number appears in the SDA report, it must be analyzed to find out why it is there. A failing data recording device can cause many data cartridges to be included in the Data Cartridge Statistics Report. The value of the report lies in the ability to identify failure trends of cartridges or components. Your objective should be to identify data cartridges approaching a ‘zero stripes available’ condition before a permanent failure occurs.

Each data cartridge accumulates usage and error statistics in its cartridge label area, by data recording device, which includes:

- Read usage and soft read checks
- Write usage and soft write checks
- Data cartridge loads and load errors

The information in the 256-byte cartridge label area can be obtained using the Access Method Services AUDITMSS command with the READLABEL option. The data cartridge label format is described in the *IBM 3850 Mass Storage System (MSS) Principles of Operation.*

We recommend replacing data cartridges before a ‘stripes available zero’ condition occurs. The cost of a data cartridge is small when compared to the value of the data that could be lost and to the cost and inconvenience involved in recovery time. SDA should be run regularly as part of your cartridge management. The level of cartridge management depends on whether the data is critical or not, and on the cost of data recovery.

Compared to the time and cost of reruns, cartridge replacement is a bargain. To preserve critical data sets and jobs that are long-running, the time and effort expended in cartridge management are justified. The measurement of cartridge management effectiveness is the time spent in reruns and recoveries due to cartridge problems.

Installation

This chapter steps you through a Mass Storage System installation, discusses specific paths, considers possible alternative paths, and gives a planning and the installation summary.

Preinstallation Testing

Most of the preinstallation testing concerns Mass Storage System-related installation tasks. The 3330 Model 1 Disk Storage can be used to test the logic of new programs, changes to device-dependent programs, or existing tape programs. When these programs are tested on the Mass Storage System, an additional OS/VS JCL change is necessary to reflect the unit change from 3330 to 3330V, and to assign the mass storage volumes to mass storage volume groups.

Preinstallation test time can be used to analyze SMF data for configuration modification and data set planning purposes. In this case, unique volume identification is required, and data set cataloging is preferred to take full advantage of the Mass Storage System managed environment. Preinstallation test time can also be used to plan and test these changes.

Generation data groups offer a method of volume space allocation and reclamation that is primarily system-controlled. Changes in generation data group support in OS/VS might require testing for restart purposes. If generation data groups are not in use, preinstallation test time can be used to make and test such changes.

ISAM data sets transferred to the Mass Storage System can be accessed only in a cylinder fault mode. Another use of the preinstallation test time might be to convert ISAM data sets to VSAM. The application programs can access the converted VSAM data sets without a program change using the VSAM program interface.

Consider a preliminary execution of the Mass Storage Control Table Create program at the IBM Data Center. It can help you verify your Mass Storage System configuration. The program can be executed on an OS/VS1 or OS/VS2 system whether the OPSYS parameter is VS1, VS2, SVS, or JES3.

Preinstallation test allowance hours cannot be converted to post-installation test allowance days, nor can they be used to increase the on-site installation and test allowance. Subsequent model upgrades of an installed Mass Storage Facility qualify for a preinstallation test time allowance based on the difference between models.

The physical installation of the Mass Storage System and related changes can be completed in two phases. Phase 1 includes relocating existing systems, installing related units, and modifiying the 3830 Model 2 Storage Controls and/or Models 158 and 168 Integrated Storage Controls. Phase 2 covers the Mass Storage System installation.

Phase 1

Interim and final physical planning layouts should identify any prerequisite equipment relocations. In a multi-CPU environment, relocations might involve moving all CPUs within the allowable cable length distances of the Mass Storage Facility and the Staging Adapters.

To provide space for the Mass Storage Facility, installed system components might have to be relocated. If interim relocations are required, give careful thought to existing and future cable lengths and unit address changes to minimize the need for additional changes later. In some cases, this may mean ordering additional cable sets for temporary use.

Provisions can be made to install the Staging Adapters on the 3830 Model 2 Storage Controls and/or Models 158 and 168 Integrated Storage Controls up to 60 days before installation of the Mass Storage System. The field change requires an estimated 25 elapsed hours on a 3830 Model 2 Storage Control, and an estimated 20 elapsed hours on a Model 158 or Model 168 ISC. These times vary depending on how the changes are scheduled. If several 3830 Model 2 Storage Controls are to be modified, consider a peak load 3830 Model 2 Storage Control for a minimum 30-day period as an alternate for the Storage Control being modified. Depending on your configuration and addresses, it may also be possible to temporarily recable a small group of 3330 drives to another 3830 Model 2 Storage Control and free a 3830 Model 2 Storage Control as an alternate during the change.

When the Staging Adapters are installed:

- The 3340 Direct Access Storage Facility cannot be connected to the 3830 Model 3 Storage Controls
- Staging Adapter interface A is reserved for connection to the Mass Storage Control
- The number of 3830 Model 3 Storage Control interfaces available for connection to a CPU is reduced by one
- The 3830 Model 3 interface addresses must be 00, 01, 10, or 11
- An IOGEN may be required to reflect the address changes

The number of ISC interfaces available for connection to a CPU is not affected, because an additional interface is included with the Staging Adapter for the ISC.

At first, the modified Storage Controls operate in nonstaging mode and address the 3330 drives only as real (nonstaging) drives, but when the Mass Storage System is installed, the modified Storage Controls are activated to function in staging mode. The same drives can be addressed as real if they are so specified to the Mass Storage Control Table Create program. If you plan to add processor storage, channels, or 3330 Disk Storage as part of the Mass Storage System installation, schedule these changes to be completed during low installation activity.

If you plan to use installed 3333 Disk Storage and Controls, the system attachment feature codes for the 3333 must be reviewed and replaced with the appropriate codes as shown below:

	Without Staging Adapter		**With Staging Adapter**	
3333 Attached to:	**Base**	**With 8150**	**Base**	**With 8150**
158 ISC	9584	9594	9587	9597
168 ISC	9585	9595	9588	9598
3830	9581	9591	9589	9599

Note: Two system attachment features are required when the string-switch feature (8150) is specified.

No physical changes are made to the 3333. These 95xx codes identify the Customer Engineer's diagnostics programs and microfiche documentation required for maintenance of the 3333.

Of these changes, additional 3330 Disk Storage is most likely to be installed concurrently with the Mass Storage System, but equipment ordered to minimize interruptions during the Storage Control modifications must be installed earlier. Address changes or additions caused by these installations must be included in the SYSGEN planning.

Phase 2

The initial assembly and unit testing of the Mass Storage Facility can be done with little or no disruption to your operation. The Mass Storage Control Table Create program and Mass Storage System SYSGEN operations must be completed before the Mass Storage System is initialized, because they are used to complete system testing. Seven to twelve work days are required to install and test the Mass Storage Facility, depending on its model and other installation factors. Schedule phase 2 with the help of your Customer Engineer.

In the latter stages of installation, 12 to 20 hours of system time are needed for final checkout of a Mass Storage System configuration that has one Mass Storage Facility. This includes time for the Online Standalone Executive Program (OLSEP) and Online Test Executive Program (OLTEP) diagnostic tests. The time necessary to identify and fix problems discovered during final checkout is not included in this estimate.

Installing the Mass Storage System includes cabling all components and activating the staging function in all Staging Adapters. Your operation is interrupted to make the final cabling changes and execute system tests. At least one 3830 Model 3 or ISC with Staging Adapter must be installed before the Mass Storage System can function.

The Customer Engineer will load the CE scratch cartridges but is not permitted to load customer data cartridges.

An Installation Verification Program (IVP) is available and should be used to verify that the newly generated System Control Program operates with the basic Mass Storage System functions. Its member name in SYS1.SAMPLIB is MSSIVP. The IVP uses ISAM. If you do not have ISAM in your system, the IVP must be modified. An IBM Program Support Representative can assist you with the IVP. However, your test plan should extend beyond the IVP to ensure that all of the logical and physical paths created by your Mass Storage Control Table Create program are verified.

On-Site Testing

On-site installation and test time allowance begins on the Date of Installation. The Date of Installation is the day following the completion of the initial physical installation by IBM Field Engineering. A provision is made for on-site installation and test time allowance, because testing depends on the Mass Storage System configuration, its interaction with the attached system, and the unique characteristics of each customer's data base. Because it is not practical to duplicate the operational environment of each customer at an IBM Data Center, IBM provides an on-site test period of three to six weeks.

After initialization, you begin the on-site testing by executing the SAMPLIB test cases. Follow this by operating all or most of the 3330 drives in a real mode for your normal regression testing. The regression tests should check out the newly created SYSGEN or IOGEN for the Mass Storage System using the Storage Controls that are to function in a staging mode (Staging Adapters). Operations people must learn about staging and nonstaging modes of operation in a mixed environment, so that with the correct procedures, testing of the Mass Storage System can be accomplished concurrently with normal production operation.

The system's management of data must be tested to help prepare people to manage data by using programs, and an adjustment period is needed for the transition from a people-managed to a system-managed data environment.

Thoroughly test all physical and logical components and paths in the system that you specified to the Mass Storage Control Table Create program to ensure they are correct. The application testing and parallel runs can be used to make these checks.

Operator Training

When the Mass Storage System is first initialized, the operator, space manager, and the system programmer need to become familiar with the Mass Storage System operator actions and responses before regression testing begins. All operator actions should be performed by the operator under supervision of the systems programmer. The operator's basic knowledge of staging and real drives; convertible-to-real drives; logical address plugs having both unit addresses and SSIDs; messages; responses, and so on, needs reinforcing in a new data environment. The systems programmer should check out the Mass Storage System unit addresses and make sure that the generated System Control Program, and logical and physical system respond as planned.

It is important that the operator perceive the Mass Storage System as a system and understand how changes to the physical system affect the logical system. For instance, the operator must understand:

- That two cartridges make up a mass storage volume, logically equivalent to a 3330-1 volume, and that they must be handled accordingly outside the Mass Storage System
- The significance of varying virtual unit addresses (virtual drives) and real unit addresses (real and convertible-to-real drives) online and offline
- When and where to mount or demount a staging pack or a real pack
- The physical system and its real unit addresses for real and convertible-to-real drives
- The logical system and its virtual unit addresses and SSIDs
- The procedures to vary convertible-to-real drives between staging and nonstaging use

During this learning period, data cartridges can be loaded into the Mass Storage Facility in preparation for the data handling and application testing. The *IBM 3850 Data Cartridge Care and Handling Instructions* should be made available to the people who will handle the data cartridges.

The operator must also be made familiar with publications that document operation, messages, and recovery actions associated with the Mass Storage System, and understand which actions

- Are his direct responsibility
- Require the services of an installation systems programmer and/or space manager
- Require the services of an IBM Customer Engineer and/or Program Support Representative PSR

The basic publications required near the operator's console include:

- *Operator's Library: IBM 3850 Mass Storage System (MSS) under OS/VS*
- *OS/VS Mass Storage System (MSS) Services: General Information*
- *OS/VS Mass Storage System (MSS) Services: Reference Information*
- *OS/VS Message Library: Mass Storage System (MSS) Messages*
- *OS/VS Message Library: VS1 System Messages* or *OS/VS Message Library: OS/VS2 System Messages*

These are in addition to installation procedures.

Regression Testing

Although the System Control Program release should be stable and in production, an IOGEN or new SYSGEN that includes the Mass Storage System components, unit addresses, and key data sets, is required immediately before installation.

The object of the regression test is to make sure that the newly generated system functions correctly. The test procedure varies with each installation, but most installations have some type of job stream for testing purposes, usually followed by the initial use of the new System Control Program during low activity or noncritical work periods. You can begin using the Mass Storage Facility with normal production or selected applications, but make sure the following units are online:

- The Mass Storage Facility
- At least one data recording control and data recording device in in each Mass Storage Facility
- At least one accessor control (SSID 101 or SSID 102 in MSF0, SSID 111 or SSID 112 in MSF1)
- At least one Staging Adapter with access to the Mass Storage Control tables in staging drive group 00 (SSID 800 or SSID 810)
- All four staging drives that contain the Mass Storage Control tables and their alternate areas (SSIDs 000, 002, 004, and 006)

When problems occur, the normal fall-back procedure is to document the problem and IPL the previous System Control Program. Follow the same procedure with the Mass Storage System. If the new System Control Program appears to be relatively error-free, normal production can be resumed.

System Management of Data

The management of mass storage volumes and space is controlled by predefined MSVC parameters and procedures. External control occurs on an exception basis or during periodic volume and space maintenance. All changes must be made using the OS/VS utilities and Access Method Services.

Begin with the initialization of staging packs and mass storage volumes. Data cartridges must be loaded and input prepared to create the mass storage volumes, their attributes, and the groups to which they are assigned. When the basic preparation is complete, a minimum number of Mass Storage System components, unit addresses, and paths are varied online.

Use the Access Method Services to create the mass storage volumes and mass storage volume groups and place them into the active inventory. The OS/VS utilities and the Access Method Services CONVERTV command are used to copy data into the Mass Storage System. Initial tests should also copy data to and from the Mass Storage Facility and compare the results, as well as data in and out of the Mass Storage Facility to test backup data and safeguard procedures. Mass storage volumes in active inventory should be restored from backup copies, either from the inactive inventory within the Mass Storage Facility or from copies stored outside and reentered for restoration.

Data Cartridge Entry

For better space management, the majority of data cartridges should be entered and volumes created at a time when the Mass Storage System is not busy otherwise the operation might have a noticeable effect on normal production. Also, you should schedule any large number of changes to data cartridges during low-activity periods.

To ensure that the data cartridges are correctly entered and controlled, the data cartridge serial numbers stamped on the tape should be compared with the console printout. Control books and Access Method Services reports should be used for their documentation and control.

The initial loading of cartridges should be monitored, because some temporary loading failures may occur. If a data cartridge fails to load, the Mass Storage Control varies the data recording device offline and uses the next one in sequence. Before the operator varies the data recording device back online, have a Customer Engineer manually remove the data cartridge from it. If the data cartridge has been not damaged, the cartridge will probably load successfully on a second attempt.

Testing the Mass Storage System

By this time, you have tested some of the Mass Storage System functions. The complexity of testing should be gradually increased until all unit addresses, paths, and functions have been thoroughly tested.

Application testing begins with the initial applications that were selected for parallel testing and data migration. Preferably, these should be individual applications that can be easily restarted or checkpointed. Consider also that the initial testing should be done with a single CPU connected to a minimum Mass Storage System configuration and without involving current operations. If this is possible, the testing is simplified and problems can be resolved quicker.

When this testing is completed, it should be repeated with concurrent operation and the results checked again. In addition, simulated data recovery tests should be planned and executed.

In a multi-CPU environment, the next step is to repeat the initial test series with more than one CPU accessing the minimum Mass Storage System configuration. If data sharing is planned for 3330 and 3330V volumes, the shared environments should also be tested.

The Mass Storage System configuration can be gradually enlarged by varying additional components online until all unit addresses and paths are tested. The objective is to determine if any discrepancies exist between physical and logical addresses so that they can be corrected before production begins.

During this period, selected data can be copied concurrently with production and testing. Transfer of data from the existing tape and/or disk to the Mass Storage System is permitted during the on-site installation and test period. Data archiving and contingency are fall-back plans that you should consider. At the end of the on-site installation and testing period, the Mass Storage System should be completely tested, and some data migration completed in preparation for productive use.

Recognize that the Mass Storage System introduces a new system-managed data enviroment. Time is required to adapt current procedures and new procedures that may be required in your installation. The testing of the Mass Storage System should not be limited to the subsystem but should be extended to include all related areas.

Data Migration

Data can be transferred from existing tape and/or disk to the Mass Storage System data cartridges during production or copied during low-activity periods. In the initial migration cycle, data can be selectively transferred by reading the present tape or disk during normal production and writing the new data set on the Mass Storage System data cartridges. This activity should include daily jobs so that experience can be accumulated regarding data generations as soon as possible.

The data set migration must be managed on a daily basis to make certain that OS/VS JCL changes are applied as required. During the initial migration, additional controls may be necessary on a temporary basis to control parallel data sets on tape and/or disk and Mass Storage System data cartridges.

The migration rate depends on the number of problems encountered. The primary objective is to move the active data planned for migration as soon as possible, consistent with safe and recoverable procedures.

Because all data cannot be transferred during production on a Master In/Master Out basis, the various data set and volume copy programs can be used to transfer data during low-activity periods. Infrequently used data that could be transferred on a Master In/Master Out basis may be copied because the normal scheduled execution occurs after the planned migration period ends.

If the CONVERTV command is used to migrate 3330 volumes to the Mass Storage System, make certain that the 3330 volumes are written with a standard record zero (R0). Record zero is normally written on the disk pack at the plant of manufacture. The Write R0 command should be limited to assigning alternate tracks using standard OS/VS utility programs. If the CONVERTV command detects a R0 with a nonzero key field, or a data field that is not equal to eight bytes, the destage operation is terminated. Action must then be taken to delete the target mass storage volume from the MSVI.

If a large amount of data is being transferred, plan Mass Storage System checkpoints, SDA and Trace analysis, and collect SMF data for future data recovery and performance optimization.

Documentation

In addition to the normal schedule and program and data set documentation, on-going procedures, such as volume and space management, Mass Storage System Console, SDA, and Trace listings should be regularly maintained to achieve an optimum level of Mass Storage System operation.

Documentation at the console should include:

- A Mass Storage System physical and logical configuration with unit addresses and paths by CPU, staging drive group, and component. Copies of the Mass Storage Control Table Create program output and Worksheets 1 and 4 provide this information.
- Commonly used system messages, procedures, and actions
- Publications that document the operation, messages, and recovery actions
- Special instructions, such as who to contact in the event an EC is to be installed or the Mass Storage Control tables must be restored, where the type 23FDs are stored, availability of keys to locked cabinets, and so on
- Console logs that can be used as audit trails as well as overviews of system operation and problems

Planning and Installation Summary

The various planning and installation tasks described in this guide are summarized here, listed in the approximate order of execution. For management, this summary can be used to develop checkpoints for monitoring the Mass Storage System installation. For the installation team, it is a detailed list of new installation considerations and tasks.

Please remember that all tasks are not applicable to every installation. Our intent is to help you develop a program that fits to your needs.

General Planning

Review of Tasks:

- Review mission, policies, workload, and procedures.
- Develop a general installation plan.
- Identify the Mass Storage System installation team.
- Assign team responsibilities.

Training (IBM and Installation Offerings):

- Introduction to Mass Storage System (IBM)
- Installation Planning Workshop (IBM)
- Mass Storage System Internal Operations and Recovery (IBM)
- Applications programmers (installation)
- Operations personnel (installation)
- Users (installation)

Prerequisites:

- System/370 Model 145, 148, 155II, 158, 165II, 168, 158AP, 168AP, 158MP, 168MP, or 3033 Processor
- Disk Storage (3330-1, 2, or 11)
- OS/VS2 SVS, OS/VS2 MVS, OS/VS1 or JES3, plus VSAM catalog

Conversions Related to the New System Control Program:

- System Control Program functions such as VSAM and TSO
- TP functions such as TCAM, NCP, and VTAM
- DB/DC programs such as IMS-VS, CICS-VS, and GIS-VS
- TEXT and/or RETRIEVAL programs such as ATMS and STAIRS

Review of Existing Methods:

- Labeling
- Cataloging
- Generation Data Groups
- Procedure Library
- OS/VS JCL
- Device-Dependent Code
- Installation Standards

Select Application Programs For Data Migration:

- Analyze SMF or manually-collected data set information.
- Select data sets for migration.
- Project new or extended applications.
- Estimate projected Mass Storage System workload.

Start Application Programs Analysis:

- Identify unique jobs (OS/VS JCL sets).
- Identify associated data sets.
- Identify device-dependent programs:
  1. DCB parameters, such as DEVD=TA, TRTCH=, DEN, OPTCD, RECFM=D, and DEVICE
  2. BDAM randomizing algorithms
  3. CCWs with WTM
  4. Tape logic, such as READBACKWARDS
  5. SYNAD routines
  6. Label-processing routines
  7. Accounting routines

Detailed Planning

Mass Storage System Configuration Verification:

- Review the physical configuration.
- Develop the logical configuration.
- Run a preliminary Mass Storage Control Table Create program.

Data Set Classification and Analysis:

- By use, type, and attributes
  1. Tape or disk
  2. Temporary or permanent
  3. Exclusive or shared
  4. Backup, archive, or safeguard
  5. Unmounted or mounted
  6. Size and activity
  7. Interchange
  8. Critical or noncritical
  9. Scheduled use
  10. Generation data groups
  11. Update-in-place
- Mass storage volume attributes and use
  1. BIND or NOBIND
  2. EXCLUSIVE or SHARED
  3. DASDERASE or NODASDERASE
  4. READONLY or READWRITE

- VSAM data set attributes and use
  1. Stage or no stage
  2. Stage and bind
  3. Asynchronous or synchronous destage
- Data set considerations for the Mass Storage System
  1. Cataloging, labeling, and naming conventions
  2. Blocking
  3. Password protection
  4. VSAM candidates
  5. OS/VS JCL parameters
  6. Generation data groups
  7. Creation and expiration dates
  8. Backup and archive procedures
- Migration techniques
  1. Cataloged procedures (symbolic parameters)
  2. Cataloged data sets
  3. Generation data groups
  4. Text editing
- Migration methods
  1. Master In/Master Out (by data set during normal processing)
  2. OS/VS utility (IEBGENER, IEBCOPY, IEHMOVE)
  3. CONVERTV (by 3330-1 volume only)

Mass Storage Volume Control:

- Volume grouping strategy
  1. By functional or department users
  2. By data set size and attributes
  3. By volume retention date
  4. By combinations or other special areas
  5. Temporary grouping for migration
- Define mass storage groups and volumes
  1. Group and default parameters
  2. Volume names, attributes, group copies, and so on
- MSV Inventory and MSVC Journal data sets
  1. Define requirements
  2. Plan of data backup and recovery procedures
- Define data cartridge entry, exit, handling, and control procedures.

Develop Job and Data Set Migration Schedules

- Test pilot applications (jobs)
- Job and data set interdependencies
- Migration schedule iterations
- OS/VS JCL/data set procedures and management

System-Managed Environment Considerations:

- Volume and space management
  1. Allocation and reclamation of space and volumes
  2. Maintenance of space and volumes
  3. Mass storage volume group space
- Functional changes
  1. Operations
  2. Librarians
  3. Users
- Operations requirements
  1. Managing real and virtual drives
  2. Access Method Services
  3. OS/VS utilities
  4. Mass Storage System messages

Security and Recovery:

- Physical site, building, and room security
- Physical and logical Mass Storage System security
- Data Recovery procedures
- Data backup procedures
- New Mass Storage System data sets checkpoint and update procedures
  1. Mass Storage Control tables
  2. MSVI
  3. MSVCJRNL
- Media management
- Password protection for Mass Storage Control tables and Access Method Services

Physical Facilities Plan:

- Environment and space considerations
- Interim and final layouts
- Cable lengths and unit address changes

Preinstallation Testing

SMF Data Reduction for Current Data Set Activity:

- Final review and adjustment
- Projection of future changes

System Data Methods:

- OS/VS JCL, labeling, and cataloging
- Generation data groups
- Device-dependent programs

Related Changes:

- ISAM to VSAM
- Program Products

Mass Storage System Applications:

- Mass Storage Control Table Create program and Mass Storage System SYSGEN
- OS/VS JCL for pilot test (using 3330 in place of 3330V)
- New and extended application testing

Preinstallation Allowance:

- By Mass Storage Facility type and serial

  1. Allowance difference between model upgrades
  2. IBM Data Center or customer site

Installation Phase 1

Major Equipment Relocations:

- CPUs for Mass Storage Facility attachment
- Tape and/or disk to provide space for the Mass Storage System

3850 System Control Program Generation:

- Plan real and virtual addresses
- Check for 3830-3 and 3830-2 address conflicts
- Options

  1. Mass Storage Control Table Create program and Mass Storage System SYSGEN
  2. Mass Storage System SYSGEN, followed by Mass Storage Control Table Create program and/or IOGEN
  3. SYSGEN, followed by Mass Storage Control Table Create program and Mass Storage System SYSGEN

Install Storage Control Modifications:

- ISC and/or 3830 (1-14 units)
- Before Mass Storage Facility Installation (0-60 days)
  1. Function in nonstaging (real) mode
  2. Support changes (3340 removed and number interfaces reduced)
  3. Staging Adapter interface addresses restricted to 00, 01, 10, or 11
- Review installed 3333 feature codes (95xx)

Related Installations:

- Additional channels
  1. For capacity or addresses
  2. Unit control word features on 2880 or CPU block multiplexer channel
- Additional 3830/3333/3330/3350s
  1. Peak load (during Staging Adapter modifications)
  2. Staging drive capacity
- Additional processor storage

Installation Phase 2

Physical Mass Storage System Installation:

- Complete Mass Storage Control Table Create program and IOGEN (by customer)
- Assemble, cable, and test units in the Mass Storage Facility
- Test the Mass Storage System (using customer System Control Program)
  1. Vary staging drives online
  2. Load CE scratch cartridges
  3. Complete diagnostic tests
  4. Activate staging function on all modified Storage Controls
  5. Run Installation Verification Procedure (MSSIVP in SAMPLIB)
- Estimated installation and/or test time (working days)
  1. A1 or B1 (7-8)
  2. A2 or B2 (8-9)
  3. A3 or B3 (10-11)
  4. A4 or B4 (11-12)

Customer On-Site Testing

On-Site Installation And Test Allowance:

- Applies to initial Mass Storage Facility type and serial (3-6 weeks maximum)
- No additional on-site installation and test allowance for Mass Storage Facility upgrades

Operator Training

- Staging and nonstaging and SSID
- Basic operation and messages
- Problem determination
- Recovery procedures
- SAMPLIB tests

Regression Test:

- System Control Program and applications

Continue Production with Tested System Control Program

- Test Mass Storage System concurrently with production
- Vary from nonstaging to staging environment

Test System Management of Data:

- Initialize staging packs
- Load data cartridges for testing
- Create mass storage volumes and groups for testing
- Copy and/or eject cartridges in/out/within Mass Storage System
- Test data archive/safeguard/backup

Data Cartridge Entry

- Test data cartridge entry, exit, handling, and control procedures
- Load and monitor data cartridge entry

Test Mass Storage System Logical System:

- Initial pilot test
  1. Vary minimum number of CPUs and Mass Storage System components online.
  2. With or without concurrent production
- Progressive testing
  1. Vary additional CPUs and Mass Storage System components online.
  2. With concurrent production
  3. Test data sharing and multi-CPU environment.

Copy Data:

- Backup and initial migration
- OS/VS utilities and CONVERTV command

On-Site Installation and Test Allowance Ends:

- When Mass Storage System production begins, or
- When maximum allowance is reached

Data Migration

Production Cycles:

- Master In/Master Out
- Concurrent with productive use
- Moves active data
- Daily, weekly, and monthly jobs

Data Copy Cycles:

- Copy data with OS/VS utility or CONVERTV command
- Plan additional CPU time
- Move less active and inactive data
- More generation and backup data
- Quarterly, annual, and special jobs

Documentation:

- Schedules
- Run books
- Programs
- Mass Storage System configuration and unit addresses
- Operation, message, and recovery manuals

Normal Operations

Trace and Tune Activities:

- Configuration changes
- Identify Problems
- Performance and optimization
- LRU parameters

Space and Volume Maintenance:

- Modify groups
- Add or delete
- Exceptions
- Data cartridge requirements

New and Extended Application Development:

- Online interactive batch
- Large data bases with low response requirements

Appendix A: MSC Table Create Program Worksheets

This section contains the Mass Storage Control Table Create Program Worksheets 1 through 7. These copyrighted worksheets are used to describe, develop, and document your Mass Storage Control Table Create program. They may be reproduced for the sole purpose of creating program inputs in connection with configuring apparatus manufactured by IBM.

Note: Copies of Worksheets 1 and 4 that document the configuration, unit address, and paths should be readily available to installation and support personnel.

MSC Table Create Program Worksheet 1

Worksheet 1, Part A

CPUID ____________________
MSCINT (A-D) ________
ALTINT (A-D) ________
MPCPU ____________________

W	0	1	2	3	4	5	6	7	8	9	A	B

CPUID ____________________
MSCINT (A-D) ________
ALTINT (A-D) ________
MPCPU ____________________

X	0	1	2	3	4	5	6	7	8	9	A	B

A	B	C	D
MSF (0,1) Size (1-4)			
Port 0		Port 1	
MSFID			
	DRC0		
	DRC1		
	DRC2		
	DRC3		

CUxx (00-11)

CI	A	B	C	D	A	B	C	D
CUID (00-11)								
SA/SSID	To CUxx		To CUA		To CUxx		To CUA	

CUA(0-3)

Physical ID A B

ADDR/SSID

C D

E F

G H

DASD Parameters

Physical			Logical
DEVn	(0-7)	(00-27)	SDGxx
Model	(1,11)	(0,1)	LCUz
Drive	*	(0-7)	LUAy

*S - Staging
R,RS - REAL, REAL Shared
C,CS - CONVREAL, CONVREAL Shared

Worksheet 1, Part B

CPUID ____________________
MSCINT (A-D) ______
ALTINT (A-D) ______
MPCPU ____________________

Y	0	1	2	3	4	5	6	7	8	9	A	B

CPUID ____________________
MSCINT (A-D) ______
ALTINT (A-D) ______
MPCPU ____________________

Z	0	1	2	3	4	5	6	7	8	9	A	B

A	B	C	D
To CUxx		To CUA	

A	B	C	D
To CUxx		To CUA	

A	B	C	D
To CUxx		To CUA	

A	B	C	D
To CUxx		To CUA	

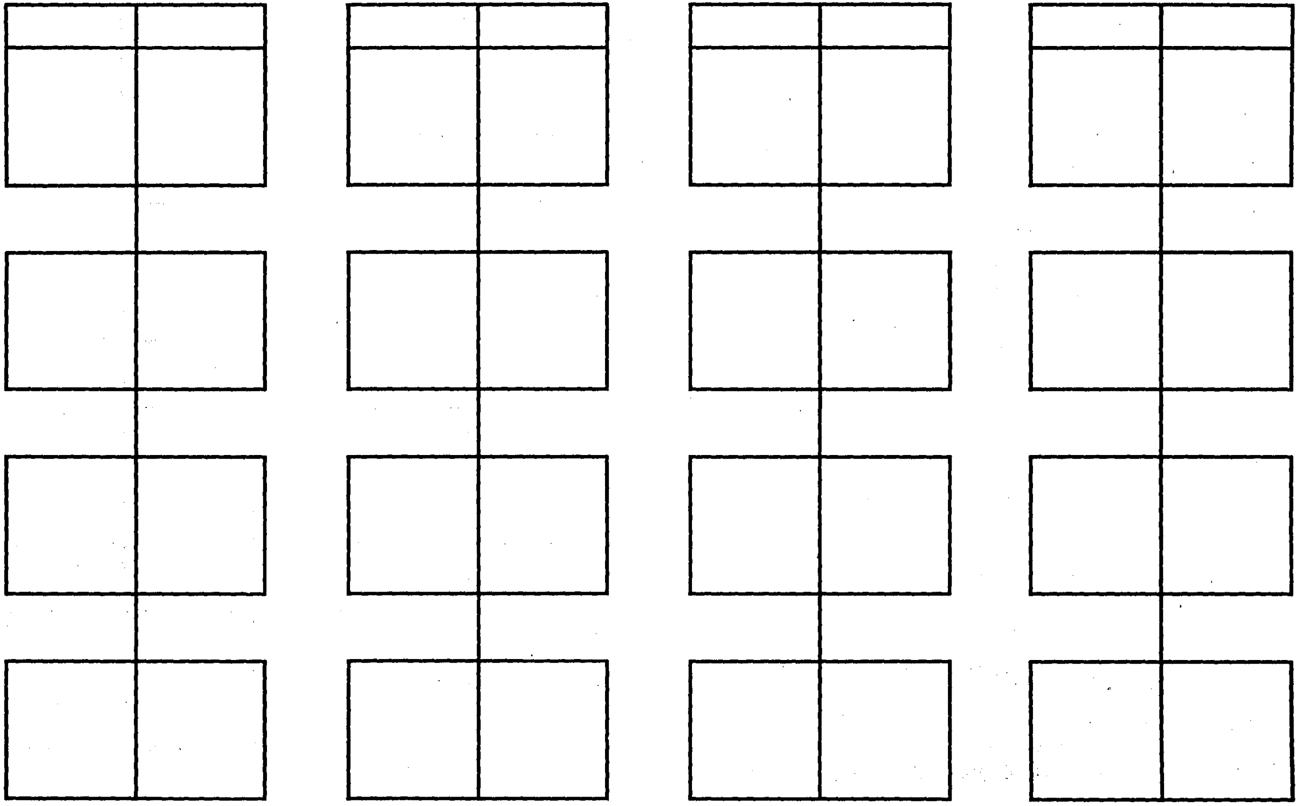

MSC Table Create Program Worksheet 2

DASD: Specifying the Drives Attached to a DASD Controller (3333 or 3350).

Required: One DASD command per 3333 or 3350 (CUA). Max = 32 commands.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create* for SA and SSID.
DEVn = 3330-1, 3330-11, or 3350.
Delete: [SSID=______,LCU,______,LUA,______,] if function is real REAL or DEVn = 3350. Delete [REAL,] if function is staging (S) or CONVREAL (C).

DASD SA=________,________,CUA=________,

DEV0=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV1=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV2=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV3=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV4=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV5=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV6=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV7=________________=SSID=SDG________,LCU________,LUA________,REAL

DASD SA=________,________,CUA=________,

DEV0=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV1=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV2=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV3=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV4=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV5=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV6=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV7=________________=SSID=SDG________,LCU________,LUA________,REAL

DASD SA=________,________,CUA=________,

DEV0=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV1=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV2=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV3=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV4=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV5=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV6=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV7=________________=SSID=SDG________,LCU________,LUA________,REAL

DASD SA=________,________,CUA=________,

DEV0=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV1=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV2=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV3=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV4=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV5=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV6=________________=SSID=SDG________,LCU________,LUA________,REAL,
DEV7=________________=SSID=SDG________,LCU________,LUA________,REAL

MSC Table Create Program Worksheet 3

LOWERCON: Describing the Connection Between a Staging Adapter and a Data Recording Control.

Required: One LOWERCON command per Staging Adapter (3830-3 or ISCSA). Max = 14 commands.

Restrictions: DRC value must be in range 0-3 depending on 3851 size.

LOWERCON SA=_______,
CU00=MSF_______,DRC_______,
CU01=MSF_______,DRC_______,
CU10=MSF_______,DRC_______,
CU11=MSF_______,DRC_______

LOWERCON SA=_______,
CU00=MSF_______,DRC_______,
CU01=MSF_______,DRC_______,
CU10=MSF_______,DRC_______,
CU11=MSF_______,DRC_______

LOWERCON SA=_______,
CU00=MSF_______,DRC_______,
CU01=MSF_______,DRC_______,
CU10=MSF_______,DRC_______,
CU11=MSF_______,DRC_______

LOWERCON SA=_______,
CU00=MSF_______,DRC_______,
CU01=MSF_______,DRC_______,
CU10=MSF_______,DRC_______,
CU11=MSF_______,DRC_______

LOWERCON SA=_______,
CU00=MSF_______,DRC_______,
CU01=MSF_______,DRC_______,
CU10=MSF_______,DRC_______,
CU11=MSF_______,DRC_______

LOWERCON SA=_______,
CU00=MSF_______,DRC_______,
CU01=MSF_______,DRC_______,
CU10=MSF_______,DRC_______,
CU11=MSF_______,DRC_______

LOWERCON SA=_______,
CU00=MSF_______,DRC_______,
CU01=MSF_______,DRC_______,
CU10=MSF_______,DRC_______,
CU11=MSF_______,DRC_______

LOWERCON SA=_______,
CU00=MSF_______,DRC_______,
CU01=MSF_______,DRC_______,
CU10=MSF_______,DRC_______,
CU11=MSF_______,DRC_______

MSC Table Create Program Worksheet 4

Addresses: Planning Device Addresses (00-3F) for UPPERCON and IODEVICE input.

Required: One worksheet per two Staging Adapter channel interfaces.

*** Distribution of Device Addresses 00–FF**

CUID	00	01	10	11
Real or Virtual	00-0F	40-4F	80-8F	C0-CF
	10-1F	50-5F	90-9F	D0-DF
Virtual	20-2F	60-6F	A0-AF	E0-EF
	30-3F	70-7F	B0-BF	F0-FF

Drive:
V, VS - Virtual, Virtual Shared
R, RS - REAL, REAL Shared
C, CS - CONVREAL, CONVREAL Shared
A - Alternate Path

Unit Address Formats

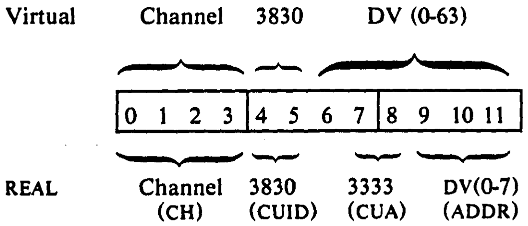

(Bit 6 must be 0)

Note: SA = 0, SA = 1 ADDR 00, 01, 08, and 09 for direct access to MSC tables:

- Reserved on CI = B
- User option on CI = C and D

CPUID ____________

SA _____ **CI** _____
CUID _____ **CH** _____

Address			
***Unit**	**Dv**	**Drive**	**SDGxx**
	00		
	01		
	02		
	03		
	04		
	05		
	06		
	07		
	08		
	09		
	0A		
	0B		
	0C		
	0D		
	0E		
	0F		
	10		
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
	20		
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

CPUID ____________

SA _____ **CI** _____
CUID _____ **CH** _____

Address			
*** Unit**	**Dv**	**Drive**	**SDGxx**
	00		
	01		
	02		
	03		
	04		
	05		
	06		
	07		
	08		
	09		
	0A		
	0B		
	0C		
	0D		
	0E		
	0F		
	10		
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
	20		
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

MSC Table Create Program Worksheet 5

UPPERCON: Describing the Connection Between a Staging Adapter and a CPU.

Required: One UPPERCON command per Staging Adapter interface B, C, and D. Max = 42 commands.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create.* Use the lower UPPERCON command on this worksheet if ALTPATH is specified for V, VS, CONVREAL, or REAL device addresses.

UPPERCON SA=______,CI=______,CUID=______,CH=______,CPUID=________________________,

V=SDG=______,ADDR= __________,__________,__________,__________,
__________,__________,__________,__________,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG=______,ADDR=__________,__________,__________,__________,
__________,__________,__________,__________,
ALTPATH=SA=______,CI=______,CH=______,

V=SDG=______,ADDR= __________,__________,__________,__________,
__________,__________,__________,__________,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG=______,ADDR=__________,__________,__________,__________,
__________,__________,__________,__________,
ALTPATH=SA=______,CI=______,CH=______,

CONVREAL=__________,__________,__________,__________,
ALTPATH=SA=______,CI=______,CH=______,
REAL=__________,__________,__________,__________,
ALTPATH=SA=______,CI=______,CH=______

UPPERCON SA=______,CI=______,CUID=______,CH=______,CPUID=________________________

MSC Table Create Program Worksheet 6

SDGxx: Describing the LRU Parameters for a Staging Drive Group.

Required: One SDGxx command per staging drive group. Max = 16 commands.

Restrictions: ACTPAGE must be low value, high value. Use four LRUGROUP values whose sum is 16.

SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______
SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______
SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______
SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______
SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______
SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______
SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______
SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______
SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

MSC Table Create Program Worksheet 7

CPUCONFG: Describing the CPU Configuration.

Required: One CPUCONFG command per attached CPU. Max = 4 commands.

Restrictions: MPCPU and ALTINT are mutually exclusive. OPSYS=VS1, VS2, SVS, or JES3.
Delete: [,MPCPU], [,ALTINT], or [OPSYS] if not required.

CPUCONFG CPUID=______________________,MSCINT=________,ALTINT=________,OPSYS=________,
MPCPU=______________________

CPUCONFG CPUID=______________________,MSCINT=________,ALTINT=________,OPSYS=________,
MPCPU=______________________

CPUCONFG CPUID=______________________,MSCINT=________,ALTINT=________,OPSYS=________,
MPCPU=______________________

CPUCONFG CPUID=______________________,MSCINT=________,ALTINT=________,OPSYS=________,
MPCPU=______________________

MSFn: Describing the Mass Storage Facility

Required: One MSFn command per 3851 MSF. Max = 2.

Restrictions: See OS/VS *Mass Storage Table Create* for MSFn and MSFID.
Delete: [,LOADMAP] if option not required.

MSF0 SIZE=______,MSFID=7,LOADMAP=________________

MSF1 SIZE=______,MSFID=6,LOADMAP=________________

CREATE: Describing the Type of Table Create Execution

Required: One CREATE command per execution.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create* for VOLID.
Delete: [NEW, or UPDATE,], [PRNTONLY], and [CYLFAULT=] if not required.

CREATE NEW,UPDATE,PRNTONLY,
LRUCLOCK=_____________,VOLID=_____________00,
CYLFAULT=_____________

Appendix B: Worksheets for the XYZ Company, Configuration B

This Appendix contains completed Worksheets 1, 2, 3, 4A-4F, 5A-5F, 6, and 7 for the fictitous XYZ Company, Configuration B. This example, described in the chapter "Mass Storage Table Create Examples," uses a Model 158MP with symmetric connections, and a Model 168MP with asymmetric connections to the Mass Storage Control. The Mass Storage Table Create program output for Configuration B is also included in this appendix.

MSC Table Create Program Worksheet 1

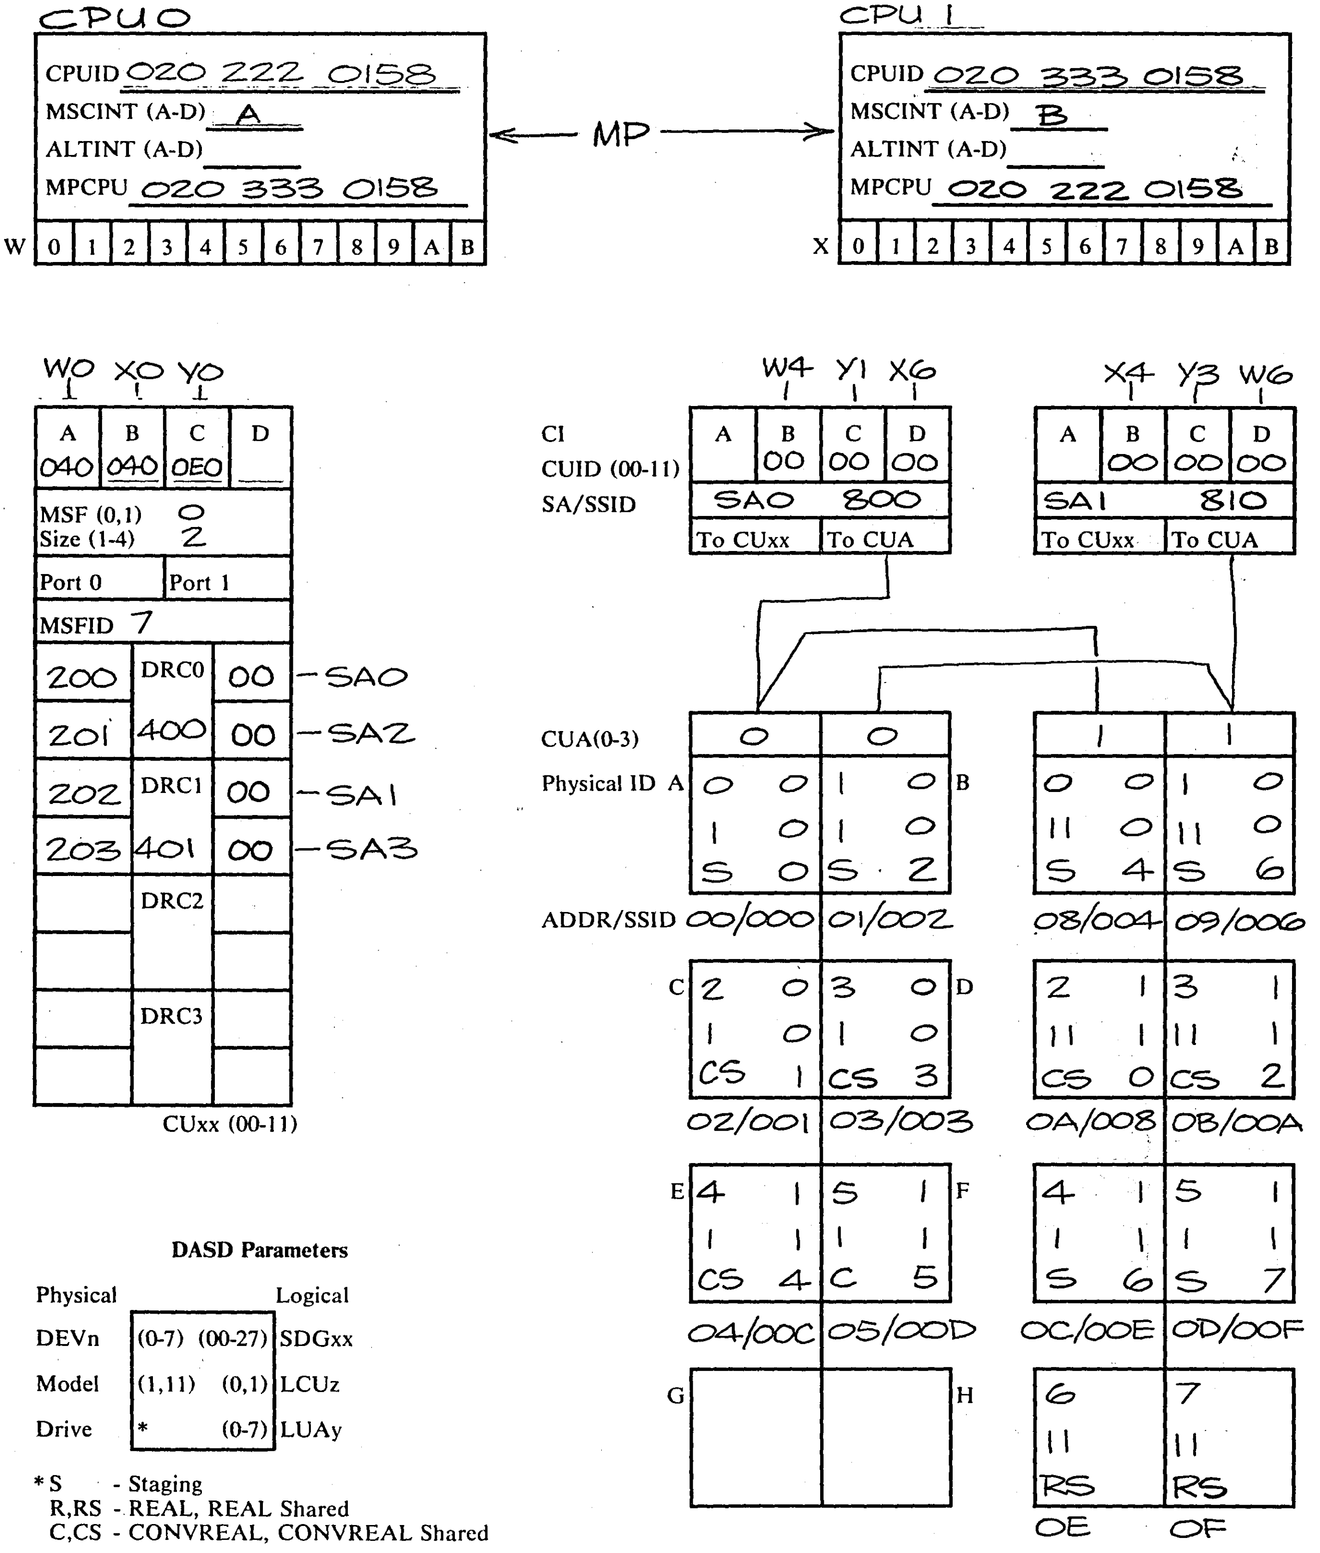

CPU 0

CPUID 020 222 0158
MSCINT (A-D) A
ALTINT (A-D) ____
MPCPU 020 333 0158

W	0	1	2	3	4	5	6	7	8	9	A	B

CPU 1

CPUID 020 333 0158
MSCINT (A-D) B
ALTINT (A-D) ____
MPCPU 020 222 0158

X	0	1	2	3	4	5	6	7	8	9	A	B

A	B	C	D
040	040	0E0	

MSF (0,1) 0
Size (1-4) 2

Port 0	Port 1

MSFID 7

200	DRC0	00	—SA0
201	400	00	—SA2
202	DRC1	00	—SA1
203	401	00	—SA3
	DRC2		
	DRC3		

CUxx (00-11)

	A	B	C	D	A	B	C	D
CI								
CUID (00-11)		00	00	00		00	00	00
SA/SSID	SA0 800				SA1 810			
	To CUxx		To CUA		To CUxx		To CUA	

CUA(0-3)	0	0	1	1
Physical ID	A 0 0 1 0 S 0	B 1 0 1 0 S 2	0 0 11 0 S 4	1 0 11 0 S 6
ADDR/SSID	00/000	01/002	08/004	09/006
	C 2 0 1 0 CS 1	D 3 0 1 0 CS 3	2 1 11 1 CS 0	3 1 11 1 CS 2
	02/001	03/003	0A/008	0B/00A
	E 4 1 1 1 CS 4	F 5 1 1 1 C 5	4 1 1 1 S 6	5 1 1 1 S 7
	04/00C	05/00D	0C/00E	0D/00F
	G	H	6 11 RS	7 11 RS
			0E	0F

DASD Parameters

Physical			Logical
DEVn	(0-7)	(00-27)	SDGxx
Model	(1,11)	(0,1)	LCUz
Drive	*	(0-7)	LUAy

*S - Staging
R,RS - REAL, REAL Shared
C,CS - CONVREAL, CONVREAL Shared

Worksheet 1, Part B

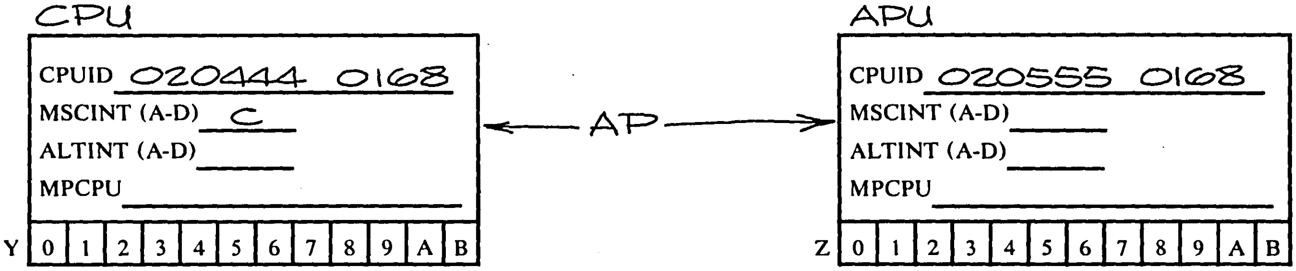

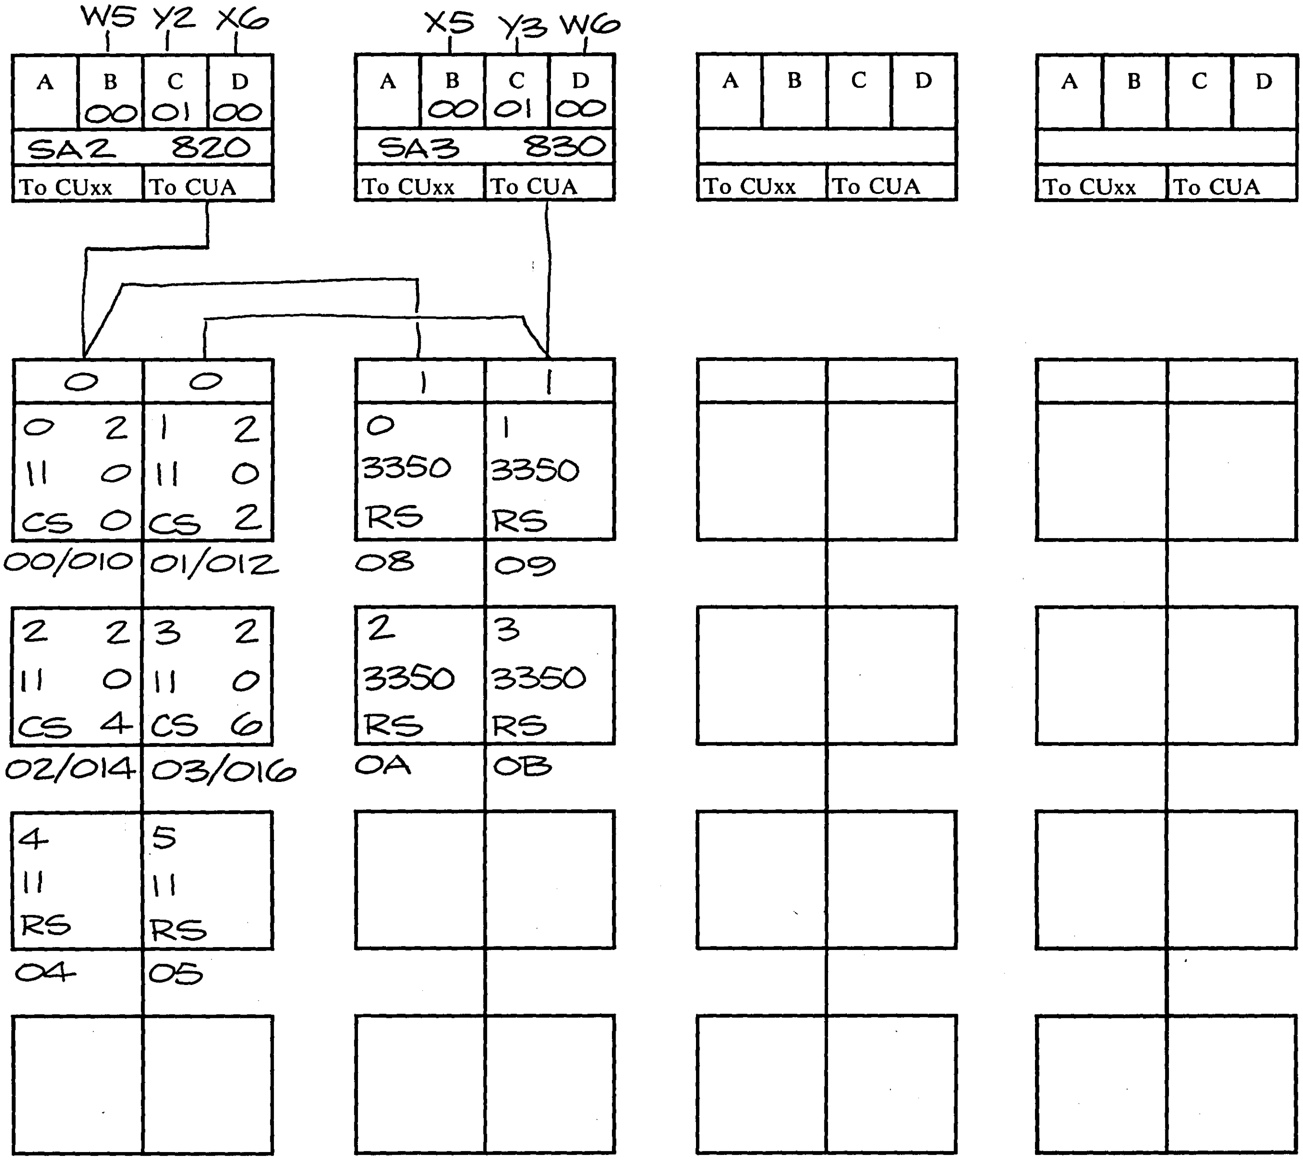

MSC Table Create Program Worksheet 2

DASD: Specifying the Drives Attached to a DASD Controller (3333 or 3350).

Required: One DASD command per 3333 or 3350 (CUA). Max = 32 commands.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create* for SA and SSID.
DEVn = 3330-1, 3330-11, or 3350.
Delete: [SSID=______,LCU,______,LUA,______,] if function is real REAL, or DEVn = 3350. Delete [REAL,] if function is staging (S) or CONVREAL (C).

DASD SA= 0 , 1 ,CUA= 0 ,

DEV0= 3330-1 =SSID=SDG 0 ,LCU 0 ,LUA 0 ~~,REAL~~

DEV1= 3330-1 =SSID=SDG 0 ,LCU 0 ,LUA 2 ~~,REAL~~

DEV2= 3330-1 =SSID=SDG 0 ,LCU 0 ,LUA 1 ~~,REAL~~

DEV3= 3330-1 =SSID=SDG 0 ,LCU 0 ,LUA 3 ~~,REAL~~

DEV4= 3330-1 =SSID=SDG 1 ,LCU 1 ,LUA 4 ~~,REAL~~

DEV5= 3330-1 =SSID=SDG 1 ,LCU 1 ,LUA 5 ~~,REAL,~~

DEV6=______=SSID=SDG______,LCU______,LUA______,REAL,

DEV7=______=SSID=SDG______,LCU______,LUA______,REAL

DASD SA= 0 , 1 ,CUA= 1 ,

DEV0= 3330-11 =SSID=SDG 0 ,LCU 0 ,LUA 4 ~~,REAL~~

DEV1= 3330-11 =SSID=SDG 0 ,LCU 0 ,LUA 6 ~~,REAL~~

DEV2= 3330-11 =SSID=SDG 1 ,LCU 1 ,LUA 0 ~~,REAL~~

DEV3= 3330-11 =SSID=SDG 1 ,LCU 1 ,LUA 2 ~~,REAL~~

DEV4= 3330-1 =SSID=SDG 1 ,LCU 1 ,LUA 6 ~~,REAL~~

DEV5= 3330-1 =SSID=SDG 1 ,LCU 1 ,LUA 7 ~~,REAL~~

DEV6= 3330-11 =~~SSID=SDG______,LCU______,LUA______~~,REAL,

DEV7= 3330-11 =~~SSID=SDG______,LCU______,LUA______~~,REAL

DASD SA= 2 , 3 ,CUA= 0 ,

DEV0= 3330-11 =SSID=SDG 2 ,LCU 0 ,LUA 0 ,REAL,

DEV1= 3330-11 =SSID=SDG 2 ,LCU 0 ,LUA 2 ,REAL,

DEV2= 3330-11 =SSID=SDG 2 ,LCU 0 ,LUA 4 ,REAL,

DEV3= 3330-11 =SSID=SDG 2 ,LCU 0 ,LUA 6 ,REAL,

DEV4= 3330-11 =~~SSID=SDG______,LCU______,LUA______~~,REAL,

DEV5= 3330-11 =~~SSID=SDG______,LCU______,LUA______~~,REAL,

DEV6=______=SSID=SDG______,LCU______,LUA______,REAL,

DEV7=______=SSID=SDG______,LCU______,LUA______,REAL

DASD SA= 2 , 3 ,CUA= 1 ,

DEV0= 3350 =~~SSID=SDG______,LCU______,LUA______~~,REAL,

DEV1= 3350 =~~SSID=SDG______,LCU______,LUA______~~,REAL,

DEV2= 3350 =~~SSID=SDG______,LCU______,LUA______~~,REAL,

DEV3= 3350 =~~SSID=SDG______,LCU______,LUA______~~,REAL,

DEV4=______=SSID=SDG______,LCU______,LUA______,REAL,

DEV5=______=SSID=SDG______,LCU______,LUA______,REAL,

DEV6=______=SSID=SDG______,LCU______,LUA______,REAL,

DEV7=______=SSID=SDG______,LCU______,LUA______,REAL

MSC Table Create Program Worksheet 3

LOWERCON: Describing the Connection Between a Staging Adapter and a Data Recording Control.

Required: One LOWERCON command per Staging Adapter (3830-3 or ISCSA). Max = 14 commands.

Restrictions: DRC value must be in range 0-3 depending on 3851 size.

LOWERCON SA= 0 ,
CU00=MSF 0 ,DRC 0 ┌
CU01=MSF______,DRC______,
CU10=MSF______,DRC______,
CU11=MSF______,DRC______

LOWERCON SA=______,
CU00=MSF______,DRC______,
CU01=MSF______,DRC______,
CU10=MSF______,DRC______,
CU11=MSF______,DRC______

LOWERCON SA= 2 ,
CU00=MSF 0 ,DRC 0 ┌
CU01=MSF______,DRC______,
CU10=MSF______,DRC______,
CU11=MSF______,DRC______

LOWERCON SA=______,
CU00=MSF______,DRC______,
CU01=MSF______,DRC______,
CU10=MSF______,DRC______,
CU11=MSF______,DRC______

LOWERCON SA= 1 ,
CU00=MSF 0 ,DRC 1 ┌
CU01=MSF______,DRC______,
CU10=MSF______,DRC______,
CU11=MSF______,DRC______

LOWERCON SA=______,
CU00=MSF______,DRC______,
CU01=MSF______,DRC______,
CU10=MSF______,DRC______,
CU11=MSF______,DRC______

LOWERCON SA= 3 ,
CU00=MSF 0 ,DRC 1 ┌
CU01=MSF______,DRC______,
CU10=MSF______,DRC______,
CU11=MSF______,DRC______

LOWERCON SA=______,
CU00=MSF______,DRC______,
CU01=MSF______,DRC______,
CU10=MSF______,DRC______,
CU11=MSF______,DRC______

MSC Table Create Program Worksheet 4 A

Addresses: Planning Device Addresses (00-3F) for UPPERCON and IODEVICE input.

Required: One worksheet per two Staging Adapter channel interfaces.

*** Distribution of Device Addresses 00–FF**

CUID	00	01	10	11
Real or Virtual	00-0F	40-4F	80-8F	C0-CF
	10-1F	50-5F	90-9F	D0-DF
Virtual	20-2F	60-6F	A0-AF	E0-EF
	30-3F	70-7F	B0-BF	F0-FF

Drive:
- V, VS - Virtual, Virtual Shared
- R, RS - REAL, REAL Shared
- C, CS - CONVREAL, CONVREAL Shared
- A - Alternate Path

Unit Address Formats

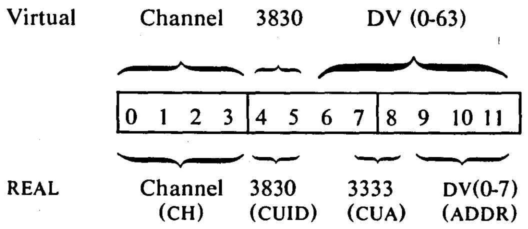

(Bit 6 must be 0)

Note: SA = 0, SA = 1 ADDR 00, 01, 08, and 09 for direct access to MSC tables:

- Reserved on CI = B
- User option on CI = C and D

CPUID 0202220158

SA 0	CI B
CUID 00	CH 4

Address		Drive	SDGxx
*Unit	Dv		
400	00	MSC	
	01		
	02	CS	
	03		
	04		
	05	C	
	06	VS	0
	07		
	08	MSC	
	09		
	0A	CS	
	0B		
	0C	VS	0
	0D		
	0E	RS	
	0F		
410	10	VS	0
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
420	20	VS	1
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
430	30		
	31		
	32		
	33		
	34}	SHARE	DUP
	35}		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

CPUID 0202220158

SA 1	CI D
CUID 00	CH 6

Address		Drive	SDGxx
*Unit	Dv		
600	00	MSC	
	01		
	02	A	
	03		
	04		
	05		
	06	A	0
	07		
	08	MSC	
	09		
	0A	A	
	0B		
	0C	A	0
	0D		
	0E	A	
	0F		
610	10	A	0
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
620	20	A	1
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
630	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

MSC Table Create Program Worksheet 4 B

Addresses: Planning Device Addresses (00-3F) for UPPERCON and IODEVICE input.

Required: One worksheet per two Staging Adapter channel interfaces.

*** Distribution of Device Addresses 00–FF**

CUID	00	01	10	11
Real or Virtual	00-0F	40-4F	80-8F	C0-CF
	10-1F	50-5F	90-9F	D0-DF
Virtual	20-2F	60-6F	A0-AF	E0-EF
	30-3F	70-7F	B0-BF	F0-FF

Drive:
V, VS - Virtual, Virtual Shared
R, RS - REAL, REAL Shared
C, CS - CONVREAL, CONVREAL Shared
A - Alternate Path

Unit Address Formats

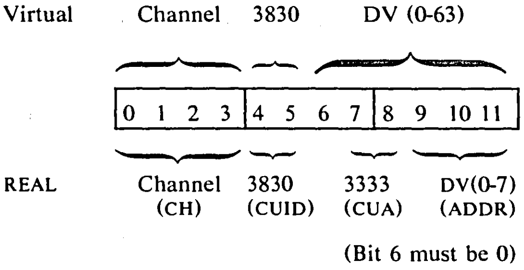

Note: SA=0, SA=1 ADDR 00, 01, 08, and 09 for direct access to MSC tables:

- Reserved on CI=B
- User option on CI=C and D

CPUID 0203330158

SA 1	CI B		
CUID 00	CH 4		

Address		Drive	SDGxx
*Unit	Dv		
400	00	MSC	
	01		
	02	CS	
	03		
	04		
	05	C	
	06	VS	0
	07		
	08	MSC	
	09		
	0A	CS	
	0B		
	0C	VS	0
	0D		
	0E	RS	
	0F		
410	10	VS	0
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
420	20	VS	1
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
430	30		
	31		
	32		
	33		
	34	SHAREDUP	
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

CPUID 0203330158

SA 0	CI D		
CUID 00	CH 6		

Address		Drive	SDGxx
* Unit	Dv		
600	00	MSC	
	01		
	02	A	
	03		
	04		
	05		
	06	A	0
	07		
	08	MSC	
	09		
	0A	A	
	0B		
	0C	A	0
	0D		
	0E	A	
	0F		
610	10	A	0
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
620	20	A	1
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
630	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

MSC Table Create Program Worksheet 4 C

Addresses: Planning Device Addresses (00-3F) for UPPERCON and IODEVICE input.

Required: One worksheet per two Staging Adapter channel interfaces.

*** Distribution of Device Addresses 00–FF**

CUID	00	01	10	11
Real or Virtual	00-0F	40-4F	80-8F	C0-CF
	10-1F	50-5F	90-9F	D0-DF
Virtual	20-2F	60-6F	A0-AF	E0-EF
	30-3F	70-7F	B0-BF	F0-FF

Drive:
V, VS - Virtual, Virtual Shared
R, RS - REAL, REAL Shared
C, CS - CONVREAL, CONVREAL Shared
A - Alternate Path

Unit Address Formats

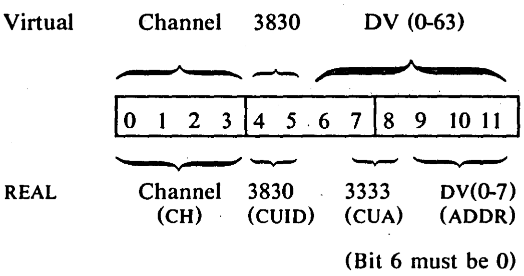

Note: SA=0, SA=1 ADDR 00, 01, 08, and 09 for direct access to MSC tables:

- Reserved on CI=B
- User option on CI=C and D

CPUID 0202220158

SA 2 CI B
CUID 01 CH 5

Address		Drive	SDGxx
*Unit	Dv		
540	00	CS	
	01		
	02		
	03		
	04	RS	
	05		
	06	VS	2
	07		
	08	RS	
	09		
	0A		
	0B		
	0C	VS	2
	0D		
	0E		
	0F		
550	10		
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
	20		
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

CPUID 0202220158

SA 03 CI D
CUID 01 CH 6

Address		Drive	SDGxx
* Unit	Dv		
640	00	A	
	01		
	02		
	03		
	04	A	
	05		
	06	A	2
	07		
	08	A	
	09		
	0A		
	0B		
	0C	A	2
	0D		
	0E		
	0F		
650	10		
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
	20		
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

MSC Table Create Program Worksheet 4 D

Addresses: Planning Device Addresses (00-3F) for UPPERCON and IODEVICE input.

Required: One worksheet per two Staging Adapter channel interfaces.

*** Distribution of Device Addresses 00–FF**

CUID	00	01	10	11
Real or Virtual	00-0F	40-4F	80-8F	C0-CF
	10-1F	50-5F	90-9F	D0-DF
Virtual	20-2F	60-6F	A0-AF	E0-EF
	30-3F	70-7F	B0-BF	F0-FF

Drive: V, VS - Virtual, Virtual Shared
R, RS - REAL, REAL Shared
C, CS - CONVREAL, CONVREAL Shared
A - Alternate Path

Unit Address Formats

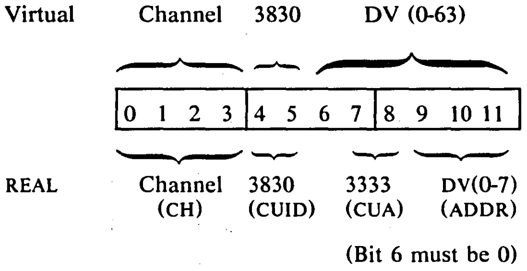

Note: SA = 0, SA = 1 ADDR 00, 01, 08, and 09 for direct access to MSC tables:

- Reserved on CI = B
- User option on CI = C and D

CPUID 0203330158

SA 3 CI B
CUID 01 CH 5

Address		Drive	SDGxx
*Unit	Dv		
540	00	CS	
	01		
	02		
	03		
	04	R	
	05		
	06	VS	2
	07		
	08	RS	
	09		
	0A		
	0B		
	0C	VS	2
	0D		
	0E		
	0F		
550	10		
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
	20		
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

CPUID 0203330158

SA 2 CI D
CUID 01 CH 6

Address		Drive	SDGxx
*Unit	Dv		
640	00	A	
	01		
	02		
	03		
	04	A	
	05		
	06	A	2
	07		
	08	A	
	09		
	0A		
	0B		
	0C	A	2
	0D		
	0E		
	0F		
650	10		
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
	20		
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

MSC Table Create Program Worksheet 4 E

Addresses: Planning Device Addresses (00-3F) for UPPERCON and IODEVICE input.

Required: One worksheet per two Staging Adapter channel interfaces.

*** Distribution of Device Addresses 00–FF**

CUID	00	01	10	11
Real or Virtual	00-0F	40-4F	80-8F	C0-CF
	10-1F	50-5F	90-9F	D0-DF
Virtual	20-2F	60-6F	A0-AF	E0-EF
	30-3F	70-7F	B0-BF	F0-FF

Drive:
V, VS - Virtual, Virtual Shared
R, RS - REAL, REAL Shared
C, CS - CONVREAL, CONVREAL Shared
A - Alternate Path

Unit Address Formats

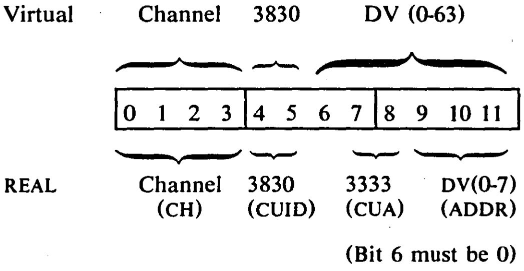

Note: SA = 0, SA = 1 ADDR 00, 01, 08, and 09 for direct access to MSC tables:

- Reserved on CI = B
- User option on CI = C and D

CPUID 020444-0168

SA 0 **CI** C
CUID 00 **CH** 1

Address			
***Unit**	**Dv**	**Drive**	**SDGxx**
100	00	MSC	
	01		
	02	CS	
	03		
	04		
	05	C	
	06	VS	0
	07		
	08	MSC	
	09		
	0A	CS	
	0B		
	0C	VS	0
	0D		
	0E	RS	
	0F		
110	10	VS	0
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
120	20	VS	1
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
130	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

CPUID 0204440168

SA 1 **CI** C
CUID 00 **CH** 3

Address			
***Unit**	**Dv**	**Drive**	**SDGxx**
300	00	MSC	
	01		
	02	A	
	03		
	04		
	05		
	06	A	0
	07		
	08	MSC	
	09		
	0A	A	
	0B		
	0C	A	0
	0D		
	0E	A	
	0F		
310	10	A	0
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
320	20	A	1
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
330	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

MSC Table Create Program Worksheet 4 F

Addresses: Planning Device Addresses (00-3F) for UPPERCON and IODEVICE input.

Required: One worksheet per two Staging Adapter channel interfaces.

*** Distribution of Device Addresses 00–FF**

CUID	00	01	10	11
Real or Virtual	00-0F	40-4F	80-8F	C0-CF
	10-1F	50-5F	90-9F	D0-DF
Virtual	20-2F	60-6F	A0-AF	E0-EF
	30-3F	70-7F	B0-BF	F0-FF

Drive:
V, VS - Virtual, Virtual Shared
R, RS - REAL, REAL Shared
C, CS - CONVREAL, CONVREAL Shared
A - Alternate Path

Unit Address Formats

Virtual: Channel (0 1 2 3), 3830 (4 5), DV (0-63) (6 7 8 9 10 11)

0	1	2	3	4	5	6	7	8	9	10	11

REAL: Channel (CH) (0 1 2 3), 3830 (CUID) (4 5), 3333 (CUA) (6 7), DV(0-7) (ADDR) (9 10 11)

(Bit 6 must be 0)

Note: SA = 0, SA = 1 ADDR 00, 01, 08, and 09 for direct access to MSC tables:

- Reserved on CI = B
- User option on CI = C and D

CPUID 0204440168

SA 2 CI C
CUID 01 CH 2

Address			
*Unit	Dv	Drive	SDGxx
240	00	CS	
	01		
	02		
	03		
	04	RS	
	05		
	06	VS	2
	07		
	08	RS	
	09		
	0A		
	0B		
	0C	VS	2
	0D		
	0E		
	0F		
250	10		
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
	20		
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

CPUID 0204440168

SA 3 CI C
CUID 01 CH 3

Address			
*Unit	Dv	Drive	SDGxx
340	00	A	
	01		
	02		
	03		
	04	A	
	05		
	06	A	2
	07		
	08	A	
	09		
	0A		
	0B		
	0C	A	2
	0D		
	0E		
	0F		
350	10		
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
	20		
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

MSC Table Create Program Worksheet 5 A

UPPERCON: Describing the Connection Between a Staging Adapter and a CPU.

Required: One UPPERCON command per Staging Adapter interface B, C, and D. Max = 42 commands.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create.* Use the lower UPPERCON command on this worksheet if ALTPATH is specified for V, VS, CONVREAL, or REAL device addresses.

UPPERCON SA= 0 ,CI= B ,CUID= 00 ,CH= 4 ,CPUID= 020 222 0158 ,

V=SDG=______,ADDR= __________,__________,__________,__________,
__________,__________,__________,__________,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG= 0 ,ADDR= 06-07 , 0C-0D , 10-1F ,__________,
__________,__________,__________,__________,
ALTPATH=SA= 1 ,CI= D ,CH= 6 ,

V=SDG=______,ADDR= __________,__________,__________,__________,
__________,__________,__________,__________,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG= 1 ,ADDR= 20-35 ,__________,__________,__________,
__________,__________,__________,__________,
ALTPATH=SA= 1 ,CI= D ,CH= 6 ,

CONVREAL= 02-05 , 0A-0B ,__________,__________,
ALTPATH=SA= 1 ,CI= D ,CH= 6 ,
REAL= 0E-0F ,__________,__________,__________,
ALTPATH=SA= 1 ,CI= D ,CH= 6

UPPERCON SA= 1 ,CI= D ,CUID= 00 ,CH= 6 ,CPUID= 020 222 0158

MSC Table Create Program Worksheet 5B

UPPERCON: Describing the Connection Between a Staging Adapter and a CPU.

Required: One UPPERCON command per Staging Adapter interface B, C, and D. Max = 42 commands.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create.* Use the lower UPPERCON command on this worksheet if ALTPATH is specified for V, VS, CONVREAL, or REAL device addresses.

UPPERCON SA= 1 ,CI= B ,CUID= 00 ,CH= 4 ,CPUID= 020 333 0158 ,

V=SDG=______,ADDR= __________,__________,__________,__________,
__________,__________,__________,__________,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG= 0 ,ADDR= 06-07 , 0C-0D , 10-1F ,__________,
__________,__________,__________,__________,
ALTPATH=SA= 0 ,CI= D ,CH= 6 ,

V=SDG=______,ADDR= __________,__________,__________,__________,
__________,__________,__________,__________,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG= 1 ,ADDR= 20-35 ,__________,__________,__________,
__________,__________,__________,__________,
ALTPATH=SA= 0 ,CI= D ,CH= 6 ,

CONVREAL= 02-05 , 0A-0B ,__________,__________,
ALTPATH=SA= 0 ,CI= D ,CH= 6 ,
REAL= 0E-0F ,__________,__________,__________,
ALTPATH=SA= 0 ,CI= D ,CH= 6

UPPERCON SA= 0 ,CI= D ,CUID= 00 ,CH= 6 ,CPUID= 020 333 0158

MSC Table Create Program Worksheet 5C

UPPERCON: Describing the Connection Between a Staging Adapter and a CPU.

Required: One UPPERCON command per Staging Adapter interface B, C, and D. Max = 42 commands.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create.* Use the lower UPPERCON command on this worksheet if ALTPATH is specified for V, VS, CONVREAL, or REAL device addresses.

UPPERCON SA= 2 ,CI= B ,CUID= 01 ,CH= 5 ,CPUID= 020 222 0158 ,

V=SDG=______,ADDR= ______,______,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG= 2 ,ADDR= 06-07 , 0C-1D ,______,______,
______,______,______,______,
ALTPATH=SA= 3 ,CI= D ,CH= 6 ,

V=SDG=______,ADDR= ______,______,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG=______,ADDR=______,______,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

CONVREAL= 00-03 ,______,______,______,
ALTPATH=SA= 3 ,CI= D ,CH= 6 ,
REAL= 04-05 , 08-0B ,______,______,
ALTPATH=SA= 3 ,CI= D ,CH= 6

UPPERCON SA= 3 ,CI= D ,CUID= 01 ,CH= 6 ,CPUID= 020 222 0158

MSC Table Create Program Worksheet 5D

UPPERCON: Describing the Connection Between a Staging Adapter and a CPU.

Required: One UPPERCON command per Staging Adapter interface B, C, and D. Max = 42 commands.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create.* Use the lower UPPERCON command on this worksheet if ALTPATH is specified for V, VS, CONVREAL, or REAL device addresses.

UPPERCON SA=3,CI=B,CUID=01,CH=5,CPUID=020 333 0168,

V=SDG=______,ADDR=______,______,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG=2,ADDR=06-07,0C-1D,______,______,
______,______,______,______,
ALTPATH=SA=2,CI=D,CH=6,

V=SDG=______,ADDR=______,______,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG=______,ADDR=______,______,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

CONVREAL=00-03,______,______,______,
ALTPATH=SA=2,CI=D,CH=6,
REAL=04-05,08-0B,______,______,
ALTPATH=SA=2,CI=D,CH=6

UPPERCON SA=2,CI=D,CUID=01,CH=6,CPUID=022 333 0158

MSC Table Create Program Worksheet 5E

UPPERCON: Describing the Connection Between a Staging Adapter and a CPU.

Required: One UPPERCON command per Staging Adapter interface B, C, and D. Max = 42 commands.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create.* Use the lower UPPERCON command on this worksheet if ALTPATH is specified for V, VS, CONVREAL, or REAL device addresses.

UPPERCON SA= 0 ,CI= C ,CUID= 00 ,CH= 1 ,CPUID= 020 444 0168 ,

V=SDG=______,ADDR= ______,______,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG= 0 ,ADDR= 06-07 , 0C-0D , 10-1F ,______,
______,______,______,______,
ALTPATH=SA= 1 ,CI= C ,CH= 3 ,

V=SDG=______,ADDR= ______,______,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG= 1 ,ADDR= 20-33 ,______,______,______,
______,______,______,______,
ALTPATH=SA= 1 ,CI= C ,CH= 3 ,

CONVREAL= 02-05 , 0A-0B ,______,______,
ALTPATH=SA= 1 ,CI= C ,CH= 3 ,
REAL= 0E-0F ,______,______,______,
ALTPATH=SA= 1 ,CI= C ,CH= 3

UPPERCON SA= 1 ,CI= C ,CUID= 00 ,CH= 3 ,CPUID= 020444 0168

MSC Table Create Program Worksheet 5 F

UPPERCON: Describing the Connection Between a Staging Adapter and a CPU.

Required: One UPPERCON command per Staging Adapter interface B, C, and D. Max = 42 commands.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create.* Use the lower UPPERCON command on this worksheet if ALTPATH is specified for V, VS, CONVREAL, or REAL device addresses.

UPPERCON SA=_2_,CI=_C_,CUID=_01_,CH=_2_,CPUID=_020 444 0168_,

V=SDG=______,ADDR= ______,______,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG=_2_,ADDR=_06-07_,_0C-1D_,______,______,
______,______,______,______,
ALTPATH=SA=_3_,CI=_C_,CH=_3_,

V=SDG=______,ADDR= ______,______,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG=______,ADDR=______,______,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

CONVREAL=_00-03_,______,______,______,
ALTPATH=SA=_3_,CI=_C_,CH=_3_,
REAL=_04-05_,_08-0B_,______,______,
ALTPATH=SA=_3_,CI=_C_,CH=_3_

UPPERCON SA=_3_,CI=_C_,CUID=_01_,CH=_3_,CPUID=_020 444 0168_

MSC Table Create Program Worksheet 6

SDGxx: Describing the LRU Parameters for a Staging Drive Group.

Required: One SDGxx command per staging drive group. Max = 16 commands.

Restrictions: ACTPAGE must be low value, high value. Use four LRUGROUP values whose sum is 16.

SDG_0_ ACTPAGE=_30_,_40_,LRUGROUP=_6_,_2_,_2_,_6_
SDG_1_ ACTPAGE=_30_,_40_,LRUGROUP=_6_,_2_,_2_,_6_

SDG_2_ ACTPAGE=_30_,_40_,LRUGROUP=_6_,_2_,_2_,_6_
SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______
SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______
SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______
SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______
SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______
SDG______ ACTPAGE=______,______,LRUGROUP=______,______,______,______

MSC Table Create Program Worksheet 7

CPUCONFG: Describing the CPU Configuration.

Required: One CPUCONFG command per attached CPU. Max = 4 commands.

Restrictions: MPCPU and ALTINT are mutually exclusive. OPSYS=VS1, VS2, SVS, or JES3.
Delete: [,MPCPU], [,ALTINT], or [OPSYS] if not required.

CPUCONFG CPUID=0202220158,MSCINT=A~~,ALTINT=~~________,OPSYS=VS2,
MPCPU=0203330158

CPUCONFG CPUID=0203330158,MSCINT=B,ALTINT=________,OPSYS=VS2,
MPCPU=0202220158

CPUCONFG CPUID=0204440168,MSCINT=C~~,ALTINT=~~________,OPSYS=VS2
MPCPU=____________________

CPUCONFG CPUID=____________________,MSCINT=________,ALTINT=________,OPSYS=________,
MPCPU=____________________

MSFn: Describing the Mass Storage Facility

Required: One MSFn command per 3851 MSF. Max = 2.

Restrictions: See OS/VS *Mass Storage Table Create* for MSFn and MSFID.
Delete: [,LOADMAP] if option not required.

MSF0 SIZE=2,MSFID=7~~,LOADMAP=~~________________

MSF1 SIZE=________,MSFID=6,LOADMAP=________________

CREATE: Describing the Type of Table Create Execution

Required: One CREATE command per execution.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create* for VOLID.
Delete: [NEW, or UPDATE,], [PRNTONLY], and [CYLFAULT=] if not required.

CREATE NEW,~~UPDATE~~PRNTONLY,
LRUCLOCK=256,VOLID=XYZB 00,
CYLFAULT=YES

MSC Table Create Example (Part 1 of 6)

```
                   MSC TABLE CREATE (VS1)
ICG103I INITIALIZE PHASE COMPLETED
ICG109I 0000 WARNING ERRORS, 0000 TERMINATING ERRORS
ICG110I HIGHEST RETURN CODE WAS 00
                            CREATE NEW,PRNTONLY,
                                 LRUCLOCK=256,VOLID=XYZB00,CYLFAULT=YES
                            CPUCONFG CPUID=0202220158,MSCINT=A,OPSYS=VS2,
                                 MPCPU=0203330158
                            CPUCONFG CPUID=0203330158,MSCINT=B,OPSYS=VS2,
                                 MPCPU=0202220158
                            CPUCONFG CPUID=0204440168,MSCINT=C,OPSYS=VS2
                            DASD SA=0,1,CUA=0,
                                 DEV0=3330-1=SSID=SDG0,LCU0,LUA0,
                                 DEV1=3330-1=SSID=SDG0,LCU0,LUA2,
                                 DEV2=3330-1=SSID=SDG0,LCU0,LUA1,
                                 DEV3=3330-1=SSID=SDG0,LCU0,LUA3,
                                 DEV4=3330-1=SSID=SDG1,LCU1,LUA4,
                                 DEV5=3330-1=SSID=SDG1,LCU1,LUA5
                            DASD SA=0,1,CUA=1,
                                 DEV0=3330-11=SSID=SDG0,LCU0,LUA4,
                                 DEV1=3330-11=SSID=SDG0,LCU0,LUA6,
                                 DEV2=3330-11=SSID=SDG1,LCU1,LUA0,
                                 DEV3=3330-11=SSID=SDG1,LCU1,LUA2,
                                 DEV4=3330-1=SSID=SDG1,LCU1,LUA6,
                                 DEV5=3330-1=SSID=SDG1,LCU1,LUA7,
                                 DEV6=3330-11=REAL,
                                 DEV7=3330-11=REAL
                            DASD SA=2,3,CUA=0,
                                 DEV0=3330-11=SSID=SDG2,LCU0,LUA0,
                                 DEV1=3330-11=SSID=SDG2,LCU0,LUA2,
                                 DEV2=3330-11=SSID=SDG2,LCU0,LUA4,
                                 DEV3=3330-11=SSID=SDG2,LCU0,LUA6,
                                 DEV4=3330-11=REAL,
                                 DEV5=3330-11=REAL
                            DASD SA=2,3,CUA=1,
                                 DEV0=3350=REAL,
                                 DEV1=3350=REAL,
                                 DEV2=3350=REAL,
                                 DEV3=3350=REAL
                            LOWERCON SA=0,
                                 CU00=MSF0,DRC0
                            LOWERCON SA=2,
                                 CU00=MSF0,DRC0
                            LOWERCON SA=1,
                                 CU00=MSF0,DRC1
                            LOWERCON SA=3,
                                 CU00=MSF0,DRC1
                            MSF0 SIZE=2,MSFID=7
                            SDG0 ACTPAGE=30,40,LRUGROUP=6,2,2,6
                            SDG1 ACTPAGE=30,40,LRUGROUP=6,2,2,6
                            SDG2 ACTPAGE=30,40,LRUGROUP=6,2,2,6
                            UPPERCON SA=0,CI=B,CUID=00,CH=4,CPUID=0202220158,
                                 VS=SDG=0,ADDR=06-07,0C-0D,10-1F,
                                 ALTPATH=SA=1,CI=D,CH=6,
                                 VS=SDG=1,ADDR=20-35,
                                 ALTPATH=SA=1,CI=D,CH=6,
                                 CONVREAL=02-05,0A-0B,
                                 ALTPATH=SA=1,CI=D,CH=6,
                                 REAL=0E-0F,
                                 ALTPATH=SA=1,CI=D,CH=6
                            UPPERCON SA=1,CI=D,CUID=00,CH=6,CPUID=0202220158
                            UPPERCON SA=1,CI=B,CUID=00,CH=4,CPUID=0203330158,
                                 VS=SDG=0,ADDR=06-07,0C-0D,10-1F,
                                 ALTPATH=SA=0,CI=D,CH=6,
                                 VS=SDG=1,ADDR=20-35,
                                 ALTPATH=SA=0,CI=D,CH=6,
                                 CONVREAL=02-05,0A-0B,
                                 ALTPATH=SA=0,CI=D,CH=6,
                                 REAL=0E-0F,
                                 ALTPATH=SA=0,CI=D,CH=6
                            UPPERCON SA=0,CI=D,CUID=00,CH=6,CPUID=0203330158
```

```
                            UPPERCON SA=2,CI=B,CUID=01,CH=5,CPUID=0202220158,
                                 VS=SDG=2,ADDR=06-07,0C-1D,
                                 ALTPATH=SA=3,CI=D,CH=6,
                                 CONVREAL=00-03,
                                 ALTPATH=SA=3,CI=D,CH=6,
                                 REAL=04-05,08-0B,
                                 ALTPATH=SA=3,CI=D,CH=6
                            UPPERCON SA=3,CI=D,CUID=01,CH=6,CPUID=0202220158
                            UPPERCON SA=3,CI=B,CUID=01,CH=5,CPUID=0203330158,
                                 VS=SDG=2,ADDR=06-07,0C-1D,
                                 ALTPATH=SA=2,CI=D,CH=6,
                                 CONVREAL=00-03,
                                 ALTPATH=SA=2,CI=D,CH=6,
                                 REAL=04-05,08-0B,
                                 ALTPATH=SA=2,CI=D,CH=6
                            UPPERCON SA=2,CI=D,CUID=01,CH=6,CPUID=0203330158
                            UPPERCON SA=0,CI=C,CUID=00,CH=1,CPUID=0204440168,
                                 VS=SDG=0,ADDR=06-07,0C-0D,10-1F,
                                 ALTPATH=SA=1,CI=C,CH=3,
                                 VS=SDG=1,ADDR=20-33,
                                 ALTPATH=SA=1,CI=C,CH=3,
                                 CONVREAL=02-05,0A-0B,
                                 ALTPATH=SA=1,CI=C,CH=3,
                                 REAL=0E-0F,
                                 ALTPATH=SA=1,CI=C,CH=3
                            UPPERCON SA=1,CI=C,CUID=00,CH=3,CPUID=0204440168
                            UPPERCON SA=2,CI=C,CUID=01,CH=2,CPUID=0204440168,
                                 VS=SDG=2,ADDR=06-07,0C-1D,
                                 ALTPATH=SA=3,CI=C,CH=3,
                                 CONVREAL=00-03,
                                 ALTPATH=SA=3,CI=C,CH=3,
                                 REAL=04-05,08-0B,
                                 ALTPATH=SA=3,CI=C,CH=3
                            UPPERCON SA=3,CI=C,CUID=01,CH=3,CPUID=0204440168
ICG104I SCAN PHASE COMPLETED
ICG109I 0000 WARNING ERRORS, 0000 TERMINATING ERRORS
ICG110I HIGHEST RETURN CODE WAS 00
ICG105I ANALYZE PHASE COMPLETED
ICG109I 0000 WARNING ERRORS, 0000 TERMINATING ERRORS
ICG110I HIGHEST RETURN CODE WAS 00
ICG401I WARNING - NOT USING CURRENT LEVEL OF MSC TABLES AS INPUT
ICG402I UNABLE TO OPEN SYSPUNCH. OUTPUT LISTED IN MESSAGE DATA SET
                   I2022201 IODEVICE ADDRESS=(400,2),UNIT=3330                             20222001
                   I2022202 IODEVICE ADDRESS=(408,2),UNIT=3330,MODEL=11                    20222002
                   I2044403 IODEVICE ADDRESS=(102,4),UNIT=3330,OPTCHAN=3                   20444003
                   I2044404 IODEVICE ADDRESS=(106,2),UNIT=3330V,OPTCHAN=3,                *20444004
                                  FEATURE=(SHARED)                                         20444005
                   I2044406 IODEVICE ADDRESS=(10A,2),UNIT=3330,OPTCHAN=3,MODEL=11          20444006
                   I2044407 IODEVICE ADDRESS=(10C,2),UNIT=3330V,OPTCHAN=3,                *20444007
                                  FEATURE=(SHARED)                                         20444008
                   I2044409 IODEVICE ADDRESS=(10E,2),UNIT=3330,OPTCHAN=3,MODEL=11          20444009
                   I2044410 IODEVICE ADDRESS=(110,36),UNIT=3330V,OPTCHAN=3,               *20444010
                                  FEATURE=(SHARED)                                         20444011
                   I2033312 IODEVICE ADDRESS=(400,2),UNIT=3330                             20333012
                   I2033313 IODEVICE ADDRESS=(402,4),UNIT=3330,OPTCHAN=6                   20333013
                   I2033314 IODEVICE ADDRESS=(406,2),UNIT=3330V,OPTCHAN=6,                *20333014
                                  FEATURE=(SHARED)                                         20333015
                   I2033316 IODEVICE ADDRESS=(408,2),UNIT=3330,MODEL=11                    20333016
                   I2033317 IODEVICE ADDRESS=(40A,2),UNIT=3330,OPTCHAN=6,MODEL=11          20333017
                   I2033318 IODEVICE ADDRESS=(40C,2),UNIT=3330V,OPTCHAN=6,                *20333018
                                  FEATURE=(SHARED)                                         20333019
                   I2033320 IODEVICE ADDRESS=(40E,2),UNIT=3330,OPTCHAN=6,MODEL=11          20333020
                   I2033321 IODEVICE ADDRESS=(410,38),UNIT=3330V,OPTCHAN=6,               *20333021
                                  FEATURE=(SHARED)                                         20333022
                   I2044423 IODEVICE ADDRESS=(240,6),UNIT=3330,OPTCHAN=3,MODEL=11          20444023
                   I2044424 IODEVICE ADDRESS=(246,2),UNIT=3330V,OPTCHAN=3,                *20444024
                                  FEATURE=(SHARED)                                         20444025
                   I2044426 IODEVICE ADDRESS=(248,4),UNIT=3350,OPTCHAN=3                   20444026
                   I2044427 IODEVICE ADDRESS=(24C,18),UNIT=3330V,OPTCHAN=3,               *20444027
                                  FEATURE=(SHARED)                                         20444028
                   I2033329 IODEVICE ADDRESS=(540,6),UNIT=3330,OPTCHAN=6,MODEL=11          20333029
                   I2033330 IODEVICE ADDRESS=(546,2),UNIT=3330V,OPTCHAN=6,                *20333030
                                  FEATURE=(SHARED)                                         20333031
                   I2033332 IODEVICE ADDRESS=(548,4),UNIT=3350,OPTCHAN=6                   20333032
                   I2033333 IODEVICE ADDRESS=(54C,18),UNIT=3330V,OPTCHAN=6,               *20333033
                                  FEATURE=(SHARED)                                         20333034
                   U2044400 UNITNAME NAME=SDG00,UNIT=((106,2),(10C,2),(110,16))            20444001
                   U2044401 UNITNAME NAME=SDG01,UNIT=((120,20))                            20444002
                   U2033300 UNITNAME NAME=SDG00,UNIT=((406,2),(40C,2),(410,16))            20333003
                   U2033301 UNITNAME NAME=SDG01,UNIT=((420,22))                            20333004
                   U2044402 UNITNAME NAME=SDG02,UNIT=((246,2),(24C,18))                    20444005
                   U2033302 UNITNAME NAME=SDG02,UNIT=((546,2),(54C,18))                    20333006
```

MSC Table Create Example (Part 3 of 6)

```
                     MSC TABLE CREATE (VS1)
HOST CONNECTIONS REPORT
DASD CONTROLLER 00
                          ******* STAGING ADAPTER 0  *******     ******* STAGING ADAPTER 1  *******
    DEVICE  DEVICE    SSID   CPUID       CPUID       CPUID          CPUID       CPUID       CPUID
     PLUG    TYPE          0202220158  0204440168  0203330158     0203330158  0204440168  0202220158
                             OS/VS2      OS/VS2      OS/VS2         OS/VS2      OS/VS2      OS/VS2

      0     3330-1    000    *****       *****       *****          *****       *****       *****
      1     3330-1    002    *****       *****       *****          *****       *****       *****
      2     3330-1    001    402-C       102-C       602-A          402-C       302-A       602-A
      3     3330-1    003    403-C       103-C       603-A          403-C       303-A       603-A
      4     3330-1    00C    404-C       104-C       604-A          404-C       304-A       604-A
      5     3330-1    00D    405-C       105-C       605-A          405-C       305-A       605-A
DASD CONTROLLER 01
                          ******* STAGING ADAPTER 0  *******     ******* STAGING ADAPTER 1  *******
    DEVICE  DEVICE    SSID   CPUID       CPUID       CPUID          CPUID       CPUID       CPUID
     PLUG    TYPE          0202220158  0204440168  0203330158     0203330158  0204440168  0202220158
                             OS/VS2      OS/VS2      OS/VS2         OS/VS2      OS/VS2      OS/VS2

      0     3330-11   004    *****       *****       *****          *****       *****       *****
      1     3330-11   006    *****       *****       *****          *****       *****       *****
      2     3330-11   008    40A-C       10A-C       60A-A          40A-C       30A-A       60A-A
      3     3330-11   00A    40B-C       10B-C       60B-A          40B-C       30B-A       60B-A
      4     3330-1    00E    40C-V       10C-V       60C-A          40C-V       30C-A       60C-A
      5     3330-1    00F    40D-V       10D-V       60D-A          40D-V       30D-A       60D-A
      6     3330-11   ***    40E-R       10E-R       60E-A          40E-R       30E-A       60E-A
      7     3330-11   ***    40F-R       10F-R       60F-A          40F-R       30F-A       60F-A
DASD CONTROLLER 00
                          ******* STAGING ADAPTER 1  *******     ******* STAGING ADAPTER 0  *******
    DEVICE  DEVICE    SSID   CPUID       CPUID       CPUID          CPUID       CPUID       CPUID
     PLUG    TYPE          0203330158  0204440168  0202220158     0202220158  0204440168  0203330158
                             OS/VS2      OS/VS2      OS/VS2         OS/VS2      OS/VS2      OS/VS2

      0     3330-1    000    *****       *****       *****          *****       *****       *****
      1     3330-1    002    *****       *****       *****          *****       *****       *****
      2     3330-1    001    402-C       302-A       602-A          402-C       102-C       602-A
      3     3330-1    003    403-C       303-A       603-A          403-C       103-C       603-A
      4     3330-1    00C    404-C       304-A       604-A          404-C       104-C       604-A
      5     3330-1    00D    405-C       305-A       605-A          405-C       105-C       605-A
DASD CONTROLLER 01
                          ******* STAGING ADAPTER 1  *******     ******* STAGING ADAPTER 0  *******
    DEVICE  DEVICE    SSID   CPUID       CPUID       CPUID          CPUID       CPUID       CPUID
     PLUG    TYPE          0203330158  0204440168  0202220158     0202220158  0204440168  0203330158
                             OS/VS2      OS/VS2      OS/VS2         OS/VS2      OS/VS2      OS/VS2

      0     3330-11   004    *****       *****       *****          *****       *****       *****
      1     3330-11   006    *****       *****       *****          *****       *****       *****
      2     3330-11   008    40A-C       30A-A       60A-A          40A-C       10A-C       60A-A
      3     3330-11   00A    40B-C       30B-A       60B-A          40B-C       10B-C       60B-A
      4     3330-1    00E    40C-V       30C-A       60C-A          40C-V       10C-V       60C-A
      5     3330-1    00F    40D-V       30D-A       60D-A          40D-V       10D-V       60D-A
      6     3330-11   ***    40E-R       30E-A       60E-A          40E-R       10E-R       60E-A
      7     3330-11   ***    40F-R       30F-A       60F-A          40F-R       10F-R       60F-A
DASD CONTROLLER 00
                          ******* STAGING ADAPTER 2  *******     ******* STAGING ADAPTER 3  *******
    DEVICE  DEVICE    SSID   CPUID       CPUID       CPUID          CPUID       CPUID       CPUID
     PLUG    TYPE          0202220158  0204440168  0203330158     0203330158  0204440168  0202220158
                             OS/VS2      OS/VS2      OS/VS2         OS/VS2      OS/VS2      OS/VS2

      0     3330-11   010    540-C       240-C       640-A          540-C       340-A       640-A
      1     3330-11   012    541-C       241-C       641-A          541-C       341-A       641-A
      2     3330-11   014    542-C       242-C       642-A          542-C       342-A       642-A
      3     3330-11   016    543-C       243-C       643-A          543-C       343-A       643-A
      4     3330-11   ***    544-R       244-R       644-A          544-R       344-A       644-A
      5     3330-11   ***    545-R       245-R       645-A          545-R       345-A       645-A
DASD CONTROLLER 01
                          ******* STAGING ADAPTER 2  *******     ******* STAGING ADAPTER 3  *******
    DEVICE  DEVICE    SSID   CPUID       CPUID       CPUID          CPUID       CPUID       CPUID
     PLUG    TYPE          0202220158  0204440168  0203330158     0203330158  0204440168  0202220158
                             OS/VS2      OS/VS2      OS/VS2         OS/VS2      OS/VS2      OS/VS2

      0     3350      ***    548-R       248-R       648-A          548-R       348-A       648-A
      1     3350      ***    549-R       249-R       649-A          549-R       349-A       649-A
      2     3350      ***    54A-R       24A-R       64A-A          54A-R       34A-A       64A-A
      3     3350      ***    54B-R       24B-R       64B-A          54B-R       34B-A       64B-A
```

MSC Table Create Example (Part 4 of 6)

```
                    MSC TABLE CREATE (VS1)
HOST CONNECTIONS REPORT
DASD CONTROLLER 00
                                    ******* STAGING ADAPTER 3  *******      ******* STAGING ADAPTER 2  *******

     DEVICE  DEVICE   SSID     CPUID       CPUID       CPUID          CPUID       CPUID       CPUID
      PLUG    TYPE          0203330158  0204440168  0202220158     0202220158  0204440168  0203330158
                              OS/VS2      OS/VS2      OS/VS2         OS/VS2      OS/VS2      OS/VS2

       0     3330-11   010     540-C       340-A       640-A          540-C       240-C       640-A
       1     3330-11   012     541-C       341-A       641-A          541-C       241-C       641-A
       2     3330-11   014     542-C       342-A       642-A          542-C       242-C       642-A
       3     3330-11   016     543-C       343-A       643-A          543-C       243-C       643-A
       4     3330-11   ***     544-R       344-A       644-A          544-R       244-R       644-A
       5     3330-11   ***     545-R       345-A       645-A          545-R       245-R       645-A
DASD CONTROLLER 01
                                    ******* STAGING ADAPTER 3  *******      ******* STAGING ADAPTER 2  *******
     DEVICE  DEVICE   SSID     CPUID       CPUID       CPUID          CPUID       CPUID       CPUID
      PLUG    TYPE          0203330158  0204440168  0202220158     0202220158  0204440168  0203330158
                              OS/VS2      OS/VS2      OS/VS2         OS/VS2      OS/VS2      OS/VS2

       0     3350      ***     548-R       348-A       648-A          548-R       248-R       648-A
       1     3350      ***     549-R       349-A       649-A          549-R       249-R       649-A
       2     3350      ***     54A-R       34A-A       64A-A          54A-R       24A-R       64A-A
       3     3350      ***     54B-R       34B-A       64B-A          54B-R       24B-R       64B-A
```

MSC Table Create Example (Part 5 of 6)

```
                   MSC TABLE CREATE (VS1)
CONNECTIONS REPORT
CONNECTION FOR CARTRIDGE STORE  MSF0, SSID 100, MSFID = 7
CONNECTIONS FOR SA 0 SSID 800
   INTERFACE A TO MSC, PORT 0, CONTROL UNIT ADDRESS 000
   INTERFACE B TO CPU 0202220158, CHANNEL 4, CONTROL UNIT ADDRESS 00
   INTERFACE C TO CPU 0204440168, CHANNEL 1, CONTROL UNIT ADDRESS 00
   INTERFACE D TO CPU 0203330158, CHANNEL 6, CONTROL UNIT ADDRESS 00
     DASD  (CTL-I)
           (SHARED SA 1)           (SHARED SA 1)
           CONTROLLER 00           CONTROLLER 01           CONTROLLER 10           CONTROLLER 11

           PLUG   TYPE    SSID     PLUG   TYPE    SSID     PLUG   TYPE    SSID     PLUG   TYPE    SSID
             0   3330-1   000        0   3330-11   004         0                       0
             1   3330-1   002        1   3330-11   006         1                       1
             2   3330-1   001        2   3330-11   008         2                       2
             3   3330-1   003        3   3330-11   00A         3                       3
             4   3330-1   00C        4   3330-1    00E         4                       4
             5   3330-1   00D        5   3330-1    00F         5                       5
             6   *****    ***        6   3330-11   ***         6                       6
             7   *****    ***        7   3330-11   ***         7                       7
     DRC/DRD (CTL-I/ADT)
          DRD CONTROLLER ADDRESS 00 TO MSF 0, DRC 0,A
                DRC 0    SSID 400
                DRD 00   SSID 200
                DRD 01   SSID 201
                DRD 10   SSID 202
                DRD 11   SSID 203
          DRD CONTROLLER ADDRESS 01 IS UNUSED
          DRD CONTROLLER ADDRESS 10 IS UNUSED
          DRD CONTROLLER ADDRESS 11 IS UNUSED
CONNECTION FOR CARTRIDGE STORE  MSF0, SSID 100, MSFID = 7
CONNECTIONS FOR SA 1 SSID 810
   INTERFACE A TO MSC, PORT 0, CONTROL UNIT ADDRESS 001
   INTERFACE B TO CPU 0203330158, CHANNEL 4, CONTROL UNIT ADDRESS 00
   INTERFACE C TO CPU 0204440168, CHANNEL 3, CONTROL UNIT ADDRESS 00
   INTERFACE D TO CPU 0202220158, CHANNEL 6, CONTROL UNIT ADDRESS 00
     DASD  (CTL-I)
           (SHARED SA 0)           (SHARED SA 0)
           CONTROLLER 00           CONTROLLER 01           CONTROLLER 10           CONTROLLER 11

           PLUG   TYPE    SSID     PLUG   TYPE    SSID     PLUG   TYPE    SSID     PLUG   TYPE    SSID
             0   3330-1   000        0   3330-11   004         0                       0
             1   3330-1   002        1   3330-11   006         1                       1
             2   3330-1   001        2   3330-11   008         2                       2
             3   3330-1   003        3   3330-11   00A         3                       3
             4   3330-1   00C        4   3330-1    00E         4                       4
             5   3330-1   00D        5   3330-1    00F         5                       5
             6   *****    ***        6   3330-11   ***         6                       6
             7   *****    ***        7   3330-11   ***         7                       7
     DRC/DRD (CTL-I/ADT)
          DRD CONTROLLER ADDRESS 00 TO MSF 0, DRC 1,A
                DRC 1    SSID 401
                DRD 00   SSID 202
                DRD 01   SSID 203
                DRD 10   SSID 200
                DRD 11   SSID 201
          DRD CONTROLLER ADDRESS 01 IS UNUSED
          DRD CONTROLLER ADDRESS 10 IS UNUSED
          DRD CONTROLLER ADDRESS 11 IS UNUSED
```

MSC Table Create Example (Part 6 of 6)

```
CONNECTION FOR CARTRIDGE STORE  MSF0, SSID 100, MSFID = 7
CONNECTIONS FOR SA 2 SSID 820
   INTERFACE A TO MSC, PORT 0, CONTROL UNIT ADDRESS 010
   INTERFACE B TO CPU 0202220158, CHANNEL 5, CONTROL UNIT ADDRESS 01
   INTERFACE C TO CPU 0204440168, CHANNEL 2, CONTROL UNIT ADDRESS 01
   INTERFACE D TO CPU 0203330158, CHANNEL 6, CONTROL UNIT ADDRESS 01
     DASD  (CTL-I)
           (SHARED SA 3)           (SHARED SA 3)
           CONTROLLER 00           CONTROLLER 01           CONTROLLER 10           CONTROLLER 11

           PLUG   TYPE    SSID     PLUG   TYPE    SSID     PLUG   TYPE    SSID     PLUG   TYPE    SSID
             0   3330-11  010        0   3350                0                       0
             1   3330-11  012        1   3350                1                       1
             2   3330-11  014        2   3350                2                       2
             3   3330-11  016        3   3350                3                       3
             4   3330-11  ***        4   *****               4                       4
             5   3330-11  ***        5   *****               5                       5
             6   *****    ***        6   *****               6                       6
             7   *****    ***        7   *****               7                       7
     DRC/DRD (CTL-I/ADT)
          DRD CONTROLLER ADDRESS 00 TO MSF 0, DRC 0,B
                DRC 0    SSID 400
                DRD 00   SSID 200
                DRD 01   SSID 201
                DRD 10   SSID 202
                DRD 11   SSID 203
          DRD CONTROLLER ADDRESS 01 IS UNUSED
          DRD CONTROLLER ADDRESS 10 IS UNUSED
          DRD CONTROLLER ADDRESS 11 IS UNUSED

CONNECTION FOR CARTRIDGE STORE  MSF0, SSID 100, MSFID = 7
CONNECTIONS FOR SA 3 SSID 830
   INTERFACE A TO MSC, PORT 0, CONTROL UNIT ADDRESS 011
   INTERFACE B TO CPU 0203330158, CHANNEL 5, CONTROL UNIT ADDRESS 01
   INTERFACE C TO CPU 0204440168, CHANNEL 3, CONTROL UNIT ADDRESS 01
   INTERFACE D TO CPU 0202220158, CHANNEL 6, CONTROL UNIT ADDRESS 01
     DASD  (CTL-I)
           (SHARED SA 2)           (SHARED SA 2)
           CONTROLLER 00           CONTROLLER 01           CONTROLLER 10           CONTROLLER 11

           PLUG   TYPE    SSID     PLUG   TYPE    SSID     PLUG   TYPE    SSID     PLUG   TYPE    SSID
             0   3330-11  010        0   3350                0                       0
             1   3330-11  012        1   3350                1                       1
             2   3330-11  014        2   3350                2                       2
             3   3330-11  016        3   3350                3                       3
             4   3330-11  ***        4   *****               4                       4
             5   3330-11  ***        5   *****               5                       5
             6   *****    ***        6   *****               6                       6
             7   *****    ***        7   *****               7                       7
     DRC/DRD (CTL-I/ADT)
          DRD CONTROLLER ADDRESS 00 TO MSF 0, DRC 1,B
                DRC 1    SSID 401
                DRD 00   SSID 202
                DRD 01   SSID 203
                DRD 10   SSID 200
                DRD 11   SSID 201
          DRD CONTROLLER ADDRESS 01 IS UNUSED
          DRD CONTROLLER ADDRESS 10 IS UNUSED
          DRD CONTROLLER ADDRESS 11 IS UNUSED
                   MSC TABLE CREATE (VS1)
ICG106I BUILD PHASE COMPLETED
ICG109I 0002 WARNING ERRORS, 0000 TERMINATING ERRORS
ICG110I HIGHEST RETURN CODE WAS 04
ICG107I END OF JOB - MSC TABLE CREATE
ICG110I HIGHEST RETURN CODE WAS 04
```

Appendix C: Worksheets for the XYZ Company, Configuration C

This appendix contains changes to Worksheets 1, 4A-4D, 5A-5D, and the CPUCONFG command portion of Worksheet 7 for the XYZ Company. These are the worksheets that required changes when the Model 158MP symmetric connections to the Mass Storage Control were changed to asymmetric. The Mass Storage Table Create program output for Configuration C is also included in this appendix.

MSC Table Create Program Worksheet 1

Worksheet 1, Part A

CPU 0

CPUID 020 222 0158

MSCINT (A-D) A

ALTINT (A-D)

MPCPU

W	0	1	2	3	4	5	6	7	8	9	A	B

← MP →

CPU 1

CPUID 020 333 0158

MSCINT (A-D)

ALTINT (A-D)

MPCPU

X	0	1	2	3	4	5	6	7	8	9	A	B

W0 Y0

A	B	C	D
040		0E0	

MSF (0,1) 0

Size (1-4) 2

Port 0	Port 1

MSFID 7

200	DRC0	00	-SA0
201	400	00	-SA2
202	DRC1	00	-SA1
203	401	00	-SA3
	DRC2		
	DRC3		

CUxx (00-11)

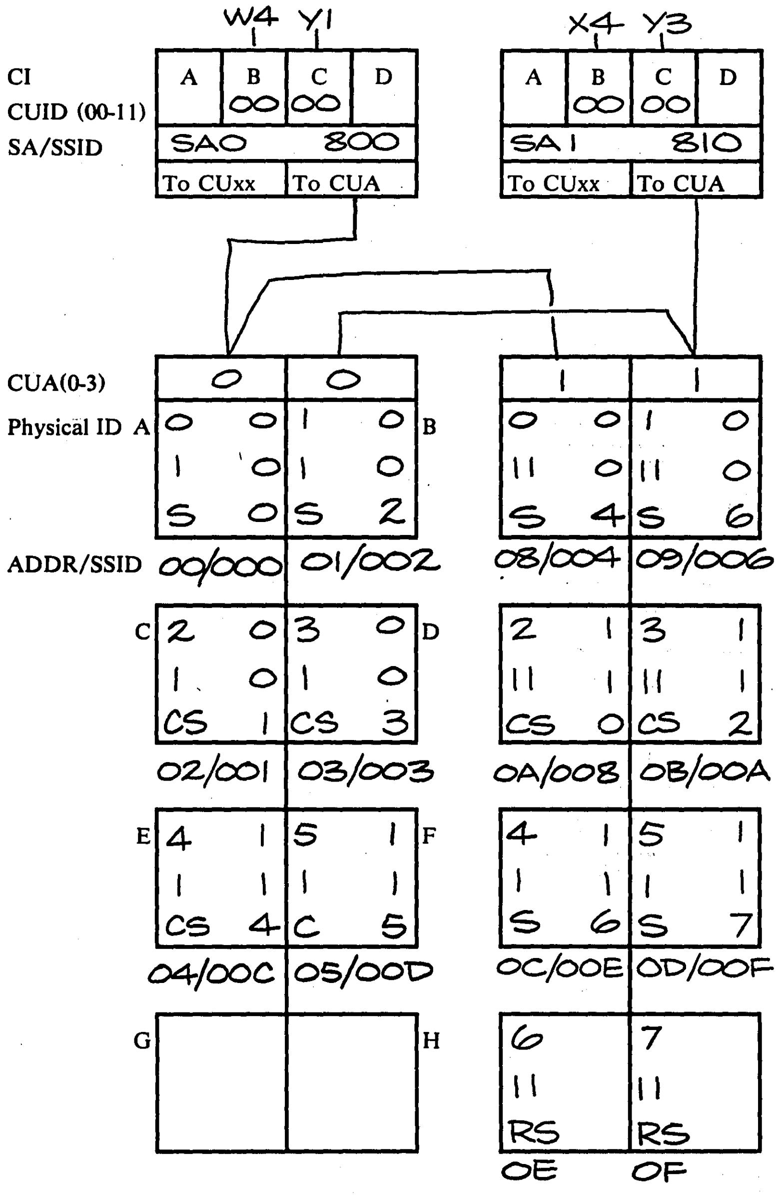

DASD Parameters

Physical			Logical
DEVn	(0-7)	(00-27)	SDGxx
Model	(1,11)	(0,1)	LCUz
Drive	*	(0-7)	LUAy

*S - Staging

R,RS - REAL, REAL Shared

C,CS - CONVREAL, CONVREAL Shared

Worksheet 1, Part B

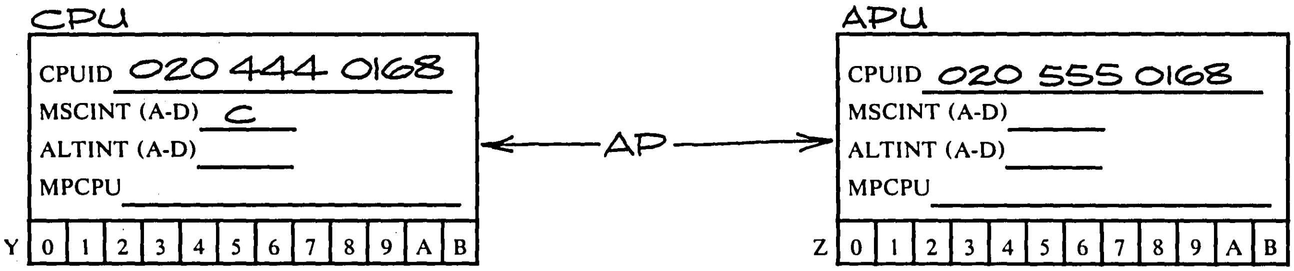

CPU

CPUID 020 444 0168

MSCINT (A-D) C

ALTINT (A-D)

MPCPU

APU

CPUID 020 555 0168

MSCINT (A-D)

ALTINT (A-D)

MPCPU

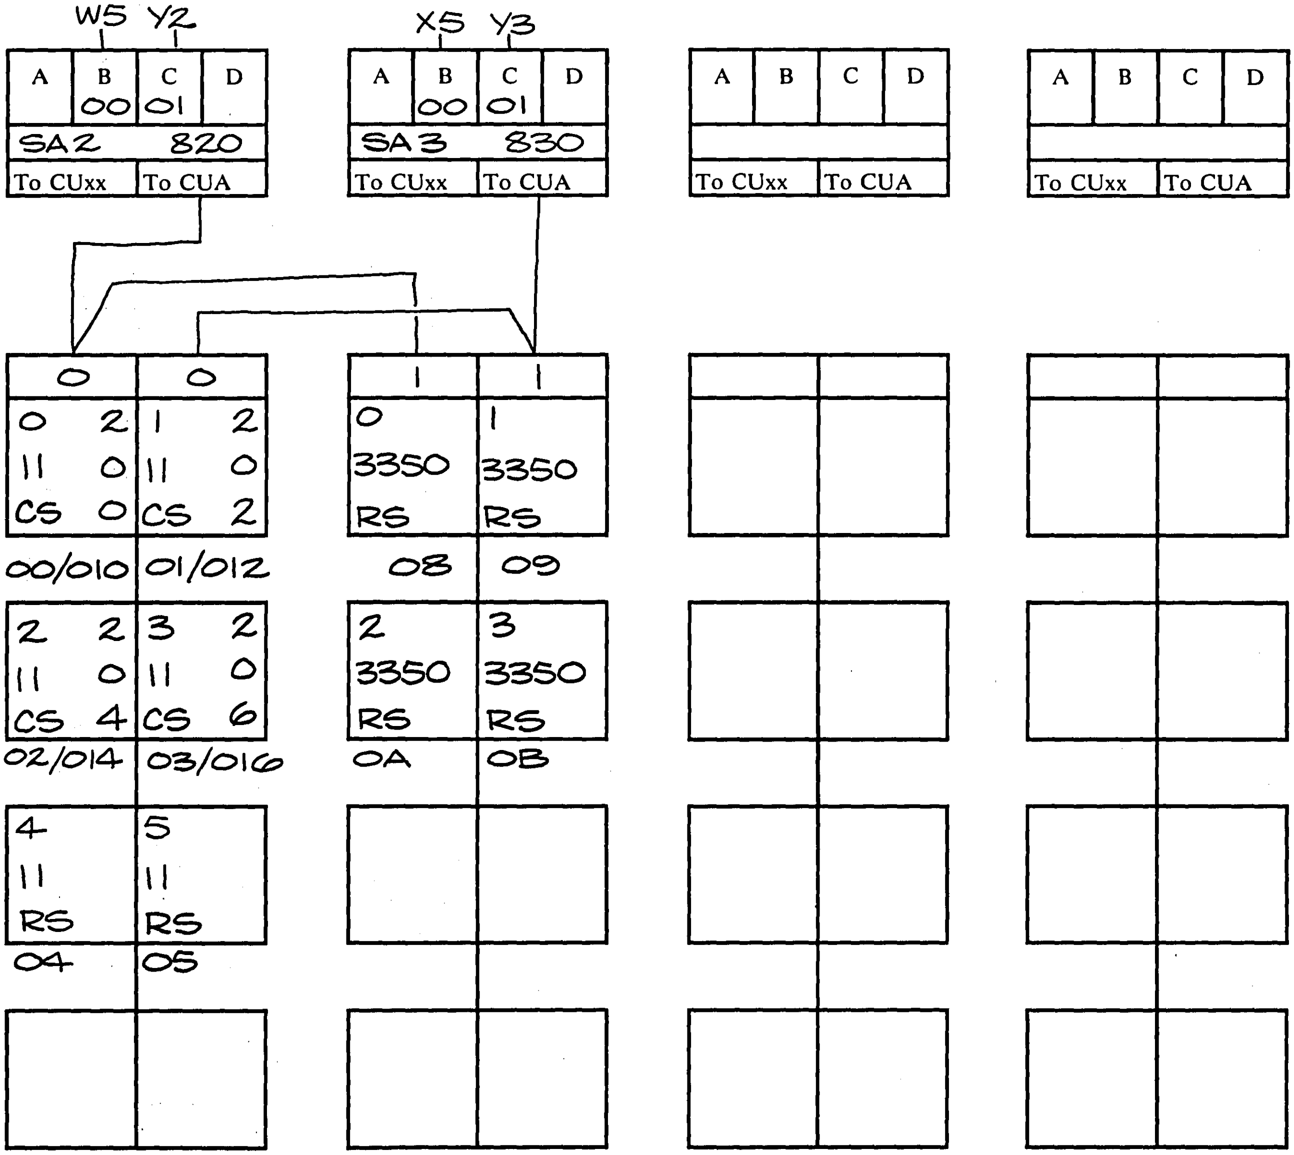

© Copyright International Business Machines Corporation 1975, 1976, 1977
A License is Granted to Users of IBM Equipment to Reproduce this Worksheet
for Configuring Apparatus Manufactured by IBM.

MSC Table Create Program Worksheet 4 A

Addresses: Planning Device Addresses (00-3F) for UPPERCON and IODEVICE input.

Required: One worksheet per two Staging Adapter channel interfaces.

*** Distribution of Device Addresses 00–FF**

CUID	00	01	10	11
Real or Virtual	00-0F	40-4F	80-8F	C0-CF
	10-1F	50-5F	90-9F	D0-DF
Virtual	20-2F	60-6F	A0-AF	E0-EF
	30-3F	70-7F	B0-BF	F0-FF

Drive:
V, VS - Virtual, Virtual Shared
R, RS - REAL, REAL Shared
C, CS - CONVREAL, CONVREAL Shared
A - Alternate Path

Unit Address Formats

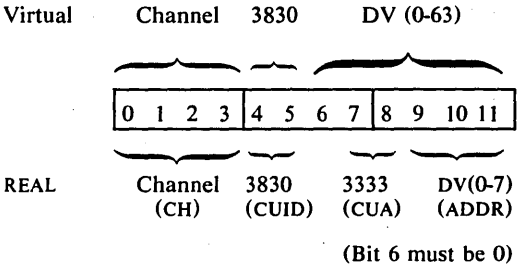

Note: SA = 0, SA = 1 ADDR 00, 01, 08, and 09 for direct access to MSC tables:

- Reserved on CI = B
- User option on CI = C and D

CPUID 0202220158

SA 0 **CI** B
CUID 00 **CH** 4

Address			
***Unit**	**Dv**	**Drive**	**SDGxx**
400	00	MSC	
	01		
	02	CS	
	03		
	04		
	05	C	
	06	VS	0
	07		
	08	MSC	
	09		
	0A	CS	
	0B		
	0C	VS	0
	0D		
	0E	RS	
	0F		
410	10	VS	0
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
420	20	VS	1
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
430	30		
	31		
	32		
	33		
	34 }	SHAREDUP	
	35 }		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

CPUID ____________

SA ______ **CI** ______
CUID ______ **CH** ______

Address			
*** Unit**	**Dv**	**Drive**	**SDGxx**
	00		
	01		
	02		
	03		
	04		
	05		
	06		
	07		
	08		
	09		
	0A		
	0B		
	0C		
	0D		
	0E		
	0F		
	10		
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
	20		
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

MSC Table Create Program Worksheet 4 B

Addresses: Planning Device Addresses (00-3F) for UPPERCON and IODEVICE input.

Required: One worksheet per two Staging Adapter channel interfaces.

*** Distribution of Device Addresses 00–FF**

CUID	00	01	10	11
Real or Virtual	00-0F	40-4F	80-8F	C0-CF
	10-1F	50-5F	90-9F	D0-DF
Virtual	20-2F	60-6F	A0-AF	E0-EF
	30-3F	70-7F	B0-BF	F0-FF

Drive:
V, VS - Virtual, Virtual Shared
R, RS - REAL, REAL Shared
C, CS - CONVREAL, CONVREAL Shared
A - Alternate Path

Unit Address Formats

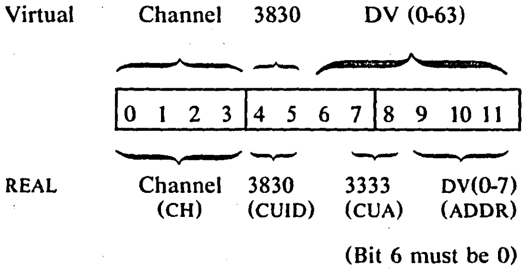

Note: SA = 0, SA = 1 ADDR 00, 01, 08, and 09 for direct access to MSC tables:

- Reserved on CI = B
- User option on CI = C and D

CPUID 020 333 0158

SA 1 CI B
CUID 00 CH 4

Address		Drive	SDGxx
*Unit	Dv		
400	00	MSC	
	01		
	02	CS	
	03		
	04		
	05	C	
	06	VS	0
	07		
	08	MSC	
	09		
	0A	CS	
	0B		
	0C	VS	0
	0D		
	0E	RS	
	0F		
410	10	VS	0
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
420	20	VS	1
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
430	30		
	31		
	32		
	33		
	34	SHARE	DUP
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

CPUID ____________

SA _____ CI _____
CUID _____ CH _____

Address		Drive	SDGxx
*Unit	Dv		
	00		
	01		
	02		
	03		
	04		
	05		
	06		
	07		
	08		
	09		
	0A		
	0B		
	0C		
	0D		
	0E		
	0F		
	10		
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
	20		
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

MSC Table Create Program Worksheet 4C

Addresses: Planning Device Addresses (00-3F) for UPPERCON and IODEVICE input.

Required: One worksheet per two Staging Adapter channel interfaces.

*** Distribution of Device Addresses 00–FF**

CUID	00	01	10	11
Real or Virtual	00-0F	40-4F	80-8F	C0-CF
	10-1F	50-5F	90-9F	D0-DF
Virtual	20-2F	60-6F	A0-AF	E0-EF
	30-3F	70-7F	B0-BF	F0-FF

Drive: V, VS - Virtual, Virtual Shared
R, RS - REAL, REAL Shared
C, CS - CONVREAL, CONVREAL Shared
A - Alternate Path

Unit Address Formats

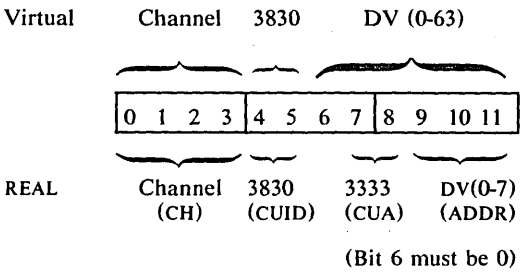

Note: SA=0, SA=1 ADDR 00, 01, 08, and 09 for direct access to MSC tables:

- Reserved on CI=B
- User option on CI=C and D

CPUID 020 222 0158

SA 2 CI B
CUID 01 CH 5

Address		Drive	SDGxx
*Unit	Dv		
540	00	CS	
	01		
	02		
	03		
	04	RS	
	05		
	06	VS	2
	07		
	08	RS	
	09		
	0A		
	0B		
	0C	VS	2
	0D		
	0E		
	0F		
550	10		
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
	20		
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

CPUID

SA CI
CUID CH

Address		Drive	SDGxx
*Unit	Dv		
	00		
	01		
	02		
	03		
	04		
	05		
	06		
	07		
	08		
	09		
	0A		
	0B		
	0C		
	0D		
	0E		
	0F		
	10		
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
	20		
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

MSC Table Create Program Worksheet 4 D

Addresses: Planning Device Addresses (00-3F) for UPPERCON and IODEVICE input.

Required: One worksheet per two Staging Adapter channel interfaces.

*** Distribution of Device Addresses 00–FF**

CUID	00	01	10	11
Real or Virtual	00-0F	40-4F	80-8F	C0-CF
	10-1F	50-5F	90-9F	D0-DF
Virtual	20-2F	60-6F	A0-AF	E0-EF
	30-3F	70-7F	B0-BF	F0-FF

Drive:
V, VS - Virtual, Virtual Shared
R, RS - REAL, REAL Shared
C, CS - CONVREAL, CONVREAL Shared
A - Alternate Path

Unit Address Formats

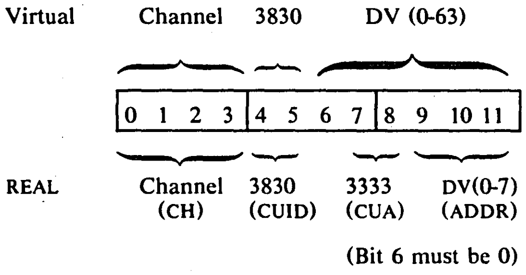

Note: SA = 0, SA = 1 ADDR 00, 01, 08, and 09 for direct access to MSC tables:

- Reserved on CI = B
- User option on CI = C and D

CPUID 020 333 0158

SA 3 CI B
CUID 01 CH 5

Address			
*Unit	Dv	Drive	SDGxx
540	00	CS	
	01		
	02		
	03		
	04	R	
	05		
	06	VS	2
	07		
	08	RS	
	09		
	0A		
	0B		
	0C	VS	2
	0D		
	0E		
	0F		
550	10		
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
	20		
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

CPUID ____________

SA ____ CI ____
CUID ____ CH ____

Address			
* Unit	Dv	Drive	SDGxx
	00		
	01		
	02		
	03		
	04		
	05		
	06		
	07		
	08		
	09		
	0A		
	0B		
	0C		
	0D		
	0E		
	0F		
	10		
	11		
	12		
	13		
	14		
	15		
	16		
	17		
	18		
	19		
	1A		
	1B		
	1C		
	1D		
	1E		
	1F		
	20		
	21		
	22		
	23		
	24		
	25		
	26		
	27		
	28		
	29		
	2A		
	2B		
	2C		
	2D		
	2E		
	2F		
	30		
	31		
	32		
	33		
	34		
	35		
	36		
	37		
	38		
	39		
	3A		
	3B		
	3C		
	3D		
	3E		
	3F		

MSC Table Create Program Worksheet 5 A

UPPERCON: Describing the Connection Between a Staging Adapter and a CPU.

Required: One UPPERCON command per Staging Adapter interface B, C, and D. Max = 42 commands.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create.* Use the lower UPPERCON command on this worksheet if ALTPATH is specified for V, VS, CONVREAL, or REAL device addresses.

UPPERCON SA= 0 ,CI= B ,CUID= 00 ,CH= 4 ,CPUID= 020 222 0158 ,

V=SDG=______,ADDR= ______,______,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG= 0 ,ADDR= 06-07 , 0C-0D , 10-1F ,______,
______,______,______,______,
ALTPATH=SA= 1 ,CI= B ,CH= F ,

V=SDG=______,ADDR= ______,______,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG= 1 ,ADDR= 20-35 ,______,______,______,
______,______,______,______,
ALTPATH=SA= 1 ,CI= B ,CH= F ,

CONVREAL= 02-05 , 0A-0B ,______,______,
ALTPATH=SA=______,CI=______,CH=______,
REAL= 0E-0F ,______,______,______,
ALTPATH=SA=______,CI= B ,CH= F

UPPERCON SA= 1 ,CI= B ,CUID= 00 ,CH= F ,CPUID= 020 222 0158

MSC Table Create Program Worksheet 5B

UPPERCON: Describing the Connection Between a Staging Adapter and a CPU.

Required: One UPPERCON command per Staging Adapter interface B, C, and D. Max = 42 commands.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create.* Use the lower UPPERCON command on this worksheet if ALTPATH is specified for V, VS, CONVREAL, or REAL device addresses.

UPPERCON SA=______,CI=______,CUID=______,CH=______,CPUID=________________________,

V=SDG=______,ADDR= __________,__________,__________,__________,

__________,__________,__________,__________,

ALTPATH=SA=______,CI=______,CH=______,

VS=SDG=______,ADDR=__________,__________,__________ ______,

__________,__________,______ ____,

ALTPATH=SA=______,CI=

V=SDG=______,ADDR= ______ ______,__________,

— ___,__________,__________,

______,CI=______,CH=______,

VS=SDG= ______,__________,__________,__________,

__________,__________,__________,__________,

ALTPATH=SA=______,CI=______,CH=______,

Note: This worksheet (correlates to 4B) is blank because the connection from channel 4 on CPUID 0202220158 to channel interface D on Staging Adapter 1 used in Configuration B has been removed.

CONVREAL=__________,__________,__________,__________,

ALTPATH=SA=______,CI=______,CH=______,

REAL=__________,__________,__________,__________,

ALTPATH=SA=______,CI=______,CH=______

UPPERCON SA=______,CI=______,CUID=______,CH=______,CPUID=________________________

MSC Table Create Program Worksheet 5 C

UPPERCON: Describing the Connection Between a Staging Adapter and a CPU.

Required: One UPPERCON command per Staging Adapter interface B, C, and D. Max = 42 commands.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create.* Use the lower UPPERCON command on this worksheet if ALTPATH is specified for V, VS, CONVREAL, or REAL device addresses.

UPPERCON SA= 2 ,CI= B ,CUID= 01 ,CH= 5 ,CPUID= 020 222 0158 ,

V=SDG=______,ADDR= ______,______,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG= 2 ,ADDR= 06-07 , 0C-1D ,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

V=SDG=______,ADDR= ______,______,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG=______,ADDR=______,______,______,______,
______,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,

CONVREAL= 00-03 ,______,______,______,
ALTPATH=SA=______,CI=______,CH=______,
REAL= 04-05 , 08-0B ,______,______,
ALTPATH=SA=______,CI=______,CH=______

UPPERCON SA= 3 ,CI= B ,CUID= 01 ,CH= F ,CPUID= 020 222 0158

MSC Table Create Program Worksheet 5 D

UPPERCON: Describing the Connection Between a Staging Adapter and a CPU.

Required: One UPPERCON command per Staging Adapter interface B, C, and D. Max = 42 commands.

Restrictions: See *OS/VS Mass Storage Control (MSC) Table Create.* Use the lower UPPERCON command on this worksheet if ALTPATH is specified for V, VS, CONVREAL, or REAL device addresses.

Note: This worksheet (correlates to 4D) is blank because the connection from channel 6 on CPUID 0203330158 to channel interface D on Staging Adapter 0 (used in Configuration B) has been removed.

UPPERCON SA=______,CI=______,CUID=______,CH=______,CPUID=____________________,

V=SDG=______,ADDR=__________,__________,__________,__________,
__________,__________,__________,__________,
ALTPATH=SA=______,CI=______,CH=______,

VS=SDG=______,ADDR=__________,__________,__________ __________,
__________,__________,__________,
ALTPATH=SA=______,CI=

V=SDG=______,ADDR=__________ __________,__________,
__ __________,__________,__________,
______,CI=______,CH=______,

VS=SDG=__ __________,__________,__________,__________,
__________,__________,__________,__________,
ALTPATH=SA=______,CI=______,CH=______,

CONVREAL=__________,__________,__________,__________,
ALTPATH=SA=______,CI=______,CH=______,
REAL=__________,__________,__________,__________,
ALTPATH=SA=______,CI=______,CH=______

UPPERCON SA=______,CI=______,CUID=______,CH=______,CPUID=____________________

MSC Table Create Program Worksheet 7

CPUCONFG: Describing the CPU Configuration.

Required: One CPUCONFG command per attached CPU. Max = 4 commands.

Restrictions: MPCPU and ALTINT are mutually exclusive. OPSYS=VS1, VS2, SVS, or JES3.
Delete: [,MPCPU], [,ALTINT], or [OPSYS] if not required.

CPUCONFG CPUID=020 222 0158,MSCINT=A,ALTINT=________,OPSYS=VS2
MPCPU=________________

CPUCONFG CPUID=0204440168,MSCINT=C~~,ALTINT=~~________OPSYS=VS2
MPCPU=________________

CPUCONFG CPUID=________________,MSCINT=________,ALTINT=________,OPSYS=________,
MPCPU=________________

CPUCONFG CPUID=________________,MSCINT=________,ALTINT=________,OPSYS=________,
MPCPU=________________

MSC Table Create Example (Part 1 of 6)

```
                  MSC TABLE CREATE (VS1)
ICG103I INITIALIZE PHASE COMPLETED
ICG109I 0000 WARNING ERRORS, 0000 TERMINATING ERRORS
ICG110I HIGHEST RETURN CODE WAS 00
                            CREATE NEW,PRNTONLY,
                                 LRUCLOCK=256,VOLID=XYZB00,CYLFAULT=YES
                            CPUCONFG CPUID=0202220158,MSCINT=A,OPSYS=VS2
                            CPUCONFG CPUID=0204440168,MSCINT=C,OPSYS=VS2
                            DASD SA=0,1,CUA=0,
                                 DEV0=3330-1=SSID=SDG0,LCU0,LUA0,
                                 DEV1=3330-1=SSID=SDG0,LCU0,LUA2,
                                 DEV2=3330-1=SSID=SDG0,LCU0,LUA1,
                                 DEV3=3330-1=SSID=SDG0,LCU0,LUA3,
                                 DEV4=3330-1=SSID=SDG1,LCU1,LUA4,
                                 DEV5=3330-1=SSID=SDG1,LCU1,LUA5
                            DASD SA=0,1,CUA=1,
                                 DEV0=3330-11=SSID=SDG0,LCU0,LUA4,
                                 DEV1=3330-11=SSID=SDG0,LCU0,LUA6,
                                 DEV2=3330-11=SSID=SDG1,LCU1,LUA0,
                                 DEV3=3330-11=SSID=SDG1,LCU1,LUA2,
                                 DEV4=3330-1=SSID=SDG1,LCU1,LUA6,
                                 DEV5=3330-1=SSID=SDG1,LCU1,LUA7,
                                 DEV6=3330-11=REAL,
                                 DEV7=3330-11=REAL
                            DASD SA=2,3,CUA=0,
                                 DEV0=3330-11=SSID=SDG2,LCU0,LUA0,
                                 DEV1=3330-11=SSID=SDG2,LCU0,LUA2,
                                 DEV2=3330-11=SSID=SDG2,LCU0,LUA4,
                                 DEV3=3330-11=SSID=SDG2,LCU0,LUA6,
                                 DEV4=3330-11=REAL,
                                 DEV5=3330-11=REAL
                            DASD SA=2,3,CUA=1,
                                 DEV0=3350=REAL,
                                 DEV1=3350=REAL,
                                 DEV2=3350=REAL,
                                 DEV3=3350=REAL
                            LOWERCON SA=0,
                                 CU00=MSF0,DRC0
                            LOWERCON SA=2,
                                 CU00=MSF0,DRC0
                            LOWERCON SA=1,
                                 CU00=MSF0,DRC1
                            LOWERCON SA=3,
                                 CU00=MSF0,DRC1
                            MSF0 SIZE=2,MSFID=7
                            SDG0 ACTPAGE=30,40,LRUGROUP=6,2,2,6
                            SDG1 ACTPAGE=30,40,LRUGROUP=6,2,2,6
                            SDG2 ACTPAGE=30,40,LRUGROUP=6,2,2,6
                            UPPERCON SA=0,CI=B,CUID=00,CH=4,CPUID=0202220158,
                                 VS=SDG=0,ADDR=06-07,0C-0D,10-1F,
                                 ALTPATH=SA=1,CI=B,CH=F,
                                 VS=SDG=1,ADDR=20-35,
                                 ALTPATH=SA=1,CI=B,CH=F,
                                 CONVREAL=02-05,0A-0B,
                                 ALTPATH=SA=1,CI=B,CH=F,
                                 REAL=0E-0F,
                                 ALTPATH=SA=1,CI=B,CH=F
                            UPPERCON SA=1,CI=B,CUID=00,CH=F,CPUID=0202220158
                            UPPERCON SA=2,CI=B,CUID=01,CH=5,CPUID=0202220158,
                                 VS=SDG=2,ADDR=06-07,0C-1D,
                                 ALTPATH=SA=3,CI=B,CH=F,
                                 CONVREAL=00-03,
                                 ALTPATH=SA=3,CI=B,CH=F,
                                 REAL=04-05,08-0B,
                                 ALTPATH=SA=3,CI=B,CH=F
                            UPPERCON SA=3,CI=B,CUID=01,CH=F,CPUID=0202220158
                            UPPERCON SA=0,CI=C,CUID=00,CH=1,CPUID=0204440168,
                                 VS=SDG=0,ADDR=06-07,0C-0D,10-1F,
                                 ALTPATH=SA=1,CI=C,CH=3,
                                 VS=SDG=1,ADDR=20-33,
                                 ALTPATH=SA=1,CI=C,CH=3,
                                 CONVREAL=02-05,0A-0B,
                                  ALTPATH=SA=1,CI=C,CH=3,
                                 REAL=0E-0F,
                                 ALTPATH=SA=1,CI=C,CH=3
                            UPPERCON SA=1,CI=C,CUID=00,CH=3,CPUID=0204440168
                            UPPERCON SA=2,CI=C,CUID=01,CH=2,CPUID=0204440168,
                                 VS=SDG=2,ADDR=06-07,0C-1D,
                                 ALTPATH=SA=3,CI=C,CH=3,
                                 CONVREAL=00-03,
                                 ALTPATH=SA=3,CI=C,CH=3,
                                 REAL=04-05,08-0B,
                                 ALTPATH=SA=3,CI=C,CH=3
                            UPPERCON SA=3,CI=C,CUID=01,CH=3,CPUID=0204440168
```

MSC Table Create Example (Part 2 of 6)

```
ICG104I SCAN PHASE COMPLETED
ICG109I 0000 WARNING ERRORS, 0000 TERMINATING ERRORS
ICG110I HIGHEST RETURN CODE WAS 00
ICG322I WARNING - SA 0 INTERFACE D NOT USED
ICG322I WARNING - SA 1 INTERFACE D NOT USED
ICG322I WARNING - SA 2 INTERFACE D NOT USED
ICG322I WARNING - SA 3 INTERFACE D NOT USED
ICG105I ANALYZE PHASE COMPLETED
ICG109I 0004 WARNING ERRORS, 0000 TERMINATING ERRORS
ICG110I HIGHEST RETURN CODE WAS 04
ICG401I WARNING - NOT USING CURRENT LEVEL OF MSC TABLES AS INPUT
ICG402I UNABLE TO OPEN SYSPUNCH. OUTPUT LISTED IN MESSAGE DATA SET
                   I2022201 IODEVICE ADDRESS=(400,2),UNIT=3330                             20222001
                   I2022202 IODEVICE ADDRESS=(402,4),UNIT=3330,OPTCHAN=F                   20222002
                   I2022203 IODEVICE ADDRESS=(406,2),UNIT=3330V,OPTCHAN=F,                *20222003
                                  FEATURE=(SHARED)                                         20222004
                   I2022205 IODEVICE ADDRESS=(408,2),UNIT=3330,MODEL=11                    20222005
                   I2022206 IODEVICE ADDRESS=(40A,2),UNIT=3330,OPTCHAN=F,MODEL=11          20222006
                   I2022207 IODEVICE ADDRESS=(40C,2),UNIT=3330V,OPTCHAN=F,                *20222007
                                  FEATURE=(SHARED)                                         20222008
                   I2022209 IODEVICE ADDRESS=(40E,2),UNIT=3330,OPTCHAN=F,MODEL=11          20222009
                   I2022210 IODEVICE ADDRESS=(410,38),UNIT=3330V,OPTCHAN=F,               *20222010
                                  FEATURE=(SHARED)                                         20222011
                   I2044412 IODEVICE ADDRESS=(102,4),UNIT=3330,OPTCHAN=3                   20444012
                   I2044413 IODEVICE ADDRESS=(106,2),UNIT=3330V,OPTCHAN=3,                *20444013
                                  FEATURE=(SHARED)                                         20444014
                   I2044415 IODEVICE ADDRESS=(10A,2),UNIT=3330,OPTCHAN=3,MODEL=11          20444015
                   I2044416 IODEVICE ADDRESS=(10C,2),UNIT=3330V,OPTCHAN=3,                *20444016
                                  FEATURE=(SHARED)                                         20444017
                   I2044418 IODEVICE ADDRESS=(10E,2),UNIT=3330,OPTCHAN=3,MODEL=11          20444018
                   I2044419 IODEVICE ADDRESS=(110,36),UNIT=3330V,OPTCHAN=3,               *20444019
                                  FEATURE=(SHARED)                                         20444020
                   I2022221 IODEVICE ADDRESS=(F00,2),UNIT=3330                             20222021
                   I2022222 IODEVICE ADDRESS=(F08,2),UNIT=3330,MODEL=11                    20222022
                   I2022223 IODEVICE ADDRESS=(540,6),UNIT=3330,OPTCHAN=F,MODEL=11          20222023
                   I2022224 IODEVICE ADDRESS=(546,2),UNIT=3330V,OPTCHAN=F,                *20222024
                                  FEATURE=(SHARED)                                         20222025
                   I2022226 IODEVICE ADDRESS=(548,4),UNIT=3350,OPTCHAN=F                   20222026
                   I2022227 IODEVICE ADDRESS=(54C,18),UNIT=3330V,OPTCHAN=F,               *20222027
                                  FEATURE=(SHARED)                                         20222028
                   I2044429 IODEVICE ADDRESS=(240,6),UNIT=3330,OPTCHAN=3,MODEL=11          20444029
                   I2044430 IODEVICE ADDRESS=(246,2),UNIT=3330V,OPTCHAN=3,                *20444030
                                  FEATURE=(SHARED)                                         20444031
                   I2044432 IODEVICE ADDRESS=(248,4),UNIT=3350,OPTCHAN=3                   20444032
                   I2044433 IODEVICE ADDRESS=(24C,18),UNIT=3330V,OPTCHAN=3,               *20444033
                                  FEATURE=(SHARED)                                         20444034
                   U2022200 UNITNAME NAME=SDG00,UNIT=((406,2),(40C,2),(410,16))            20222001
                   U2022201 UNITNAME NAME=SDG01,UNIT=((420,22))                            20222002
                   U2044400 UNITNAME NAME=SDG00,UNIT=((106,2),(10C,2),(110,16))            20444003
                   U2044401 UNITNAME NAME=SDG01,UNIT=((120,20))                            20444004
                   U2022202 UNITNAME NAME=SDG02,UNIT=((546,2),(54C,18))                    20222005
                   U2044402 UNITNAME NAME=SDG02,UNIT=((246,2),(24C,18))                    20444006
```

MSC Table Create Example (Part 3 of 6)

```
                   MSC TABLE CREATE (VS1)
HOST CONNECTIONS REPORT
DASD CONTROLLER 00
                           ******* STAGING ADAPTER 0  *******     ******* STAGING ADAPTER 1  *******
    DEVICE  DEVICE   SSID     CPUID       CPUID       CPUID          CPUID       CPUID       CPUID
     PLUG    TYPE          0202220158  0204440168     *****       0202220158  0204440168     *****
                             OS/VS2      OS/VS2                     OS/VS2      OS/VS2

      0     3330-1    000     *****       *****       *****          *****       *****       *****
      1     3330-1    002     *****       *****       *****          *****       *****       *****
      2     3330-1    001     402-C       102-C       *****          F02-A       302-A       *****
      3     3330-1    003     403-C       103-C       *****          F03-A       303-A       *****
      4     3330-1    00C     404-C       104-C       *****          F04-A       304-A       *****
      5     3330-1    00D     405-C       105-C       *****          F05-A       305-A       *****
DASD CONTROLLER 01
                           ******* STAGING ADAPTER 0  *******     ******* STAGING ADAPTER 1  *******
    DEVICE  DEVICE   SSID     CPUID       CPUID       CPUID          CPUID       CPUID       CPUID
     PLUG    TYPE          0202220158  0204440168     *****       0202220158  0204440168     *****
                             OS/VS2      OS/VS2                     OS/VS2      OS/VS2

      0     3330-11   004     *****       *****       *****          *****       *****       *****
      1     3330-11   006     *****       *****       *****          *****       *****       *****
      2     3330-11   008     40A-C       10A-C       *****          F0A-A       30A-A       *****
      3     3330-11   00A     40B-C       10B-C       *****          F0B-A       30B-A       *****
      4     3330-1    00E     40C-V       10C-V       *****          F0C-A       30C-A       *****
      5     3330-1    00F     40D-V       10D-V       *****          F0D-A       30D-A       *****
      6     3330-11   ***     40E-R       10E-R       *****          F0E-A       30E-A       *****
      7     3330-11   ***     40F-R       10F-R       *****          F0F-A       30F-A       *****
DASD CONTROLLER 00
                           ******* STAGING ADAPTER 1  *******     ******* STAGING ADAPTER 0  *******
    DEVICE  DEVICE   SSID     CPUID       CPUID       CPUID          CPUID       CPUID       CPUID
     PLUG    TYPE          0202220158  0204440168     *****       0202220158  0204440168     *****
                             OS/VS2      OS/VS2                     OS/VS2      OS/VS2

      0     3330-1    000     *****       *****       *****          *****       *****       *****
      1     3330-1    002     *****       *****       *****          *****       *****       *****
      2     3330-1    001     F02-A       302-A       *****          402-C       102-C       *****
      3     3330-1    003     F03-A       303-A       *****          403-C       103-C       *****
      4     3330-1    00C     F04-A       304-A       *****          404-C       104-C       *****
      5     3330-1    00D     F05-A       305-A       *****          405-C       105-C       *****
DASD CONTROLLER 01
                           ******* STAGING ADAPTER 1  *******     ******* STAGING ADAPTER 0  *******
    DEVICE  DEVICE   SSID     CPUID       CPUID       CPUID          CPUID       CPUID       CPUID
     PLUG    TYPE          0202220158  0204440168     *****       0202220158  0204440168     *****
                             OS/VS2      OS/VS2                     OS/VS2      OS/VS2

      0     3330-11   004     *****       *****       *****          *****       *****       *****
      1     3330-11   006     *****       *****       *****          *****       *****       *****
      2     3330-11   008     F0A-A       30A-A       *****          40A-C       10A-C       *****
      3     3330-11   00A     F0B-A       30B-A       *****          40B-C       10B-C       *****
      4     3330-1    00E     F0C-A       30C-A       *****          40C-V       10C-V       *****
      5     3330-1    00F     F0D-A       30D-A       *****          40D-V       10D-V       *****
      6     3330-11   ***     F0E-A       30E-A       *****          40E-R       10E-R       *****
      7     3330-11   ***     F0F-A       30F-A       *****          40F-R       10F-R       *****
```

MSC Table Create Example (Part 4 of 6)

```
DASD CONTROLLER 00
                           ******* STAGING ADAPTER 2  *******     ******* STAGING ADAPTER 3  *******
    DEVICE  DEVICE    SSID     CPUID       CPUID       CPUID          CPUID       CPUID       CPUID
     PLUG    TYPE           0202220158  0204440168     *****       0202220158  0204440168     *****
                              OS/VS2      OS/VS2                     OS/VS2      OS/VS2

      0     3330-11   010     540-C       240-C        *****          F40-A       340-A       *****
      1     3330-11   012     541-C       241-C        *****          F41-A       341-A       *****
      2     3330-11   014     542-C       242-C        *****          F42-A       342-A       *****
      3     3330-11   016     543-C       243-C        *****          F43-A       343-A       *****
      4     3330-11   ***     544-R       244-R        *****          F44-A       344-A       *****
      5     3330-11   ***     545-R       245-R        *****          F45-A       345-A       *****
DASD CONTROLLER 01
                           ******* STAGING ADAPTER 2  *******     ******* STAGING ADAPTER 3  *******
    DEVICE  DEVICE    SSID     CPUID       CPUID       CPUID          CPUID       CPUID       CPUID
     PLUG    TYPE           0202220158  0204440168     *****       0202220158  0204440168     *****
                              OS/VS2      OS/VS2                     OS/VS2      OS/VS2

      0     3350      ***     548-R       248-R        *****          F48-A       348-A       *****
      1     3350      ***     549-R       249-R        *****          F49-A       349-A       *****
      2     3350      ***     54A-R       24A-R        *****          F4A-A       34A-A       *****
      3     3350      ***     54B-R       24B-R        *****          F4B-A       34B-A       *****

DASD CONTROLLER 00
                           ******* STAGING ADAPTER 3  *******     ******* STAGING ADAPTER 2  *******
    DEVICE  DEVICE    SSID     CPUID       CPUID       CPUID          CPUID       CPUID       CPUID
     PLUG    TYPE           0202220158  0204440168     *****       0202220158  0204440168     *****
                              OS/VS2      OS/VS2                     OS/VS2      OS/VS2

      0     3330-11   010     F40-A       340-A        *****          540-C       240-C       *****
      1     3330-11   012     F41-A       341-A        *****          541-C       241-C       *****
      2     3330-11   014     F42-A       342-A        *****          542-C       242-C       *****
      3     3330-11   016     F43-A       343-A        *****          543-C       243-C       *****
      4     3330-11   ***     F44-A       344-A        *****          544-R       244-R       *****
      5     3330-11   ***     F45-A       345-A        *****          545-R       245-R       *****
DASD CONTROLLER 01
                           ******* STAGING ADAPTER 3  *******     ******* STAGING ADAPTER 2  *******
    DEVICE  DEVICE    SSID     CPUID       CPUID       CPUID          CPUID       CPUID       CPUID
     PLUG    TYPE           0202220158  0204440168     *****       0202220158  0204440168     *****
                              OS/VS2      OS/VS2                     OS/VS2      OS/VS2

      0     3350      ***     F48-A       348-A        *****          548-R       248-R       *****
      1     3350      ***     F49-A       349-A        *****          549-R       249-R       *****
      2     3350      ***     F4A-A       34A-A        *****          54A-R       24A-R       *****
      3     3350      ***     F4B-A       34B-A        *****          54B-R       24B-R       *****
```

MSC Table Create Example (Part 5 of 6)

```
                   MSC TABLE CREATE (VS1)
CONNECTIONS REPORT
CONNECTION FOR CARTRIDGE STORE  MSF0, SSID 100, MSFID = 7
CONNECTIONS FOR SA 0 SSID 800
   INTERFACE A TO MSC, PORT 0, CONTROL UNIT ADDRESS 000
   INTERFACE B TO CPU 0202220158, CHANNEL 4, CONTROL UNIT ADDRESS 00
   INTERFACE C TO CPU 0204440168, CHANNEL 1, CONTROL UNIT ADDRESS 00
   INTERFACE D IS NOT USED
     DASD  (CTL-I)
           (SHARED SA 1)           (SHARED SA 1)
           CONTROLLER 00           CONTROLLER 01           CONTROLLER 10          CONTROLLER 11
           PLUG   TYPE    SSID     PLUG   TYPE    SSID     PLUG   TYPE    SSID    PLUG   TYPE    SSID
             0   3330-1    000       0   3330-11   004       0                      0
             1   3330-1    002       1   3330-11   006       1                      1
             2   3330-1    001       2   3330-11   008       2                      2
             3   3330-1    003       3   3330-11   00A       3                      3
             4   3330-1    00C       4   3330-1    00E       4                      4
             5   3330-1    00D       5   3330-1    00F       5                      5
             6   *****     ***       6   3330-11   ***       6                      6
             7   *****     ***       7   3330-11   ***       7                      7
     DRC/DRD (CTL-I/ADT)
          DRD CONTROLLER ADDRESS 00 TO MSF 0, DRC 0,A
                DRC 0    SSID 400
                DRD 00   SSID 200
                DRD 01   SSID 201
                DRD 10   SSID 202
                DRD 11   SSID 203
          DRD CONTROLLER ADDRESS 01 IS UNUSED
          DRD CONTROLLER ADDRESS 10 IS UNUSED
          DRD CONTROLLER ADDRESS 11 IS UNUSED
                   MSC TABLE CREATE (VS1)
CONNECTIONS REPORT
CONNECTION FOR CARTRIDGE STORE  MSF0, SSID 100, MSFID = 7
CONNECTIONS FOR SA 1 SSID 810
   INTERFACE A TO MSC, PORT 0, CONTROL UNIT ADDRESS 001
   INTERFACE B TO CPU 0202220158, CHANNEL F, CONTROL UNIT ADDRESS 00
   INTERFACE C TO CPU 0204440168, CHANNEL 3, CONTROL UNIT ADDRESS 00
   INTERFACE D IS NOT USED
     DASD  (CTL-I)
           (SHARED SA 0)           (SHARED SA 0)
           CONTROLLER 00           CONTROLLER 01           CONTROLLER 10          CONTROLLER 11
           PLUG   TYPE    SSID     PLUG   TYPE    SSID     PLUG   TYPE    SSID    PLUG   TYPE    SSID
             0   3330-1    000       0   3330-11   004       0                      0
             1   3330-1    002       1   3330-11   006       1                      1
             2   3330-1    001       2   3330-11   008       2                      2
             3   3330-1    003       3   3330-11   00A       3                      3
             4   3330-1    00C       4   3330-1    00E       4                      4
             5   3330-1    00D       5   3330-1    00F       5                      5
             6   *****     ***       6   3330-11   ***       6                      6
             7   *****     ***       7   3330-11   ***       7                      7
     DRC/DRD (CTL-I/ADT)
          DRD CONTROLLER ADDRESS 00 TO MSF 0, DRC 1,A
                DRC 1    SSID 401
                DRD 00   SSID 202
                DRD 01   SSID 203
                DRD 10   SSID 200
                DRD 11   SSID 201
          DRD CONTROLLER ADDRESS 01 IS UNUSED
          DRD CONTROLLER ADDRESS 10 IS UNUSED
          DRD CONTROLLER ADDRESS 11 IS UNUSED
```

MSC Table Create Example (Part 6 of 6)

```
                   MSC TABLE CREATE (VS1)
CONNECTIONS REPORT
CONNECTION FOR CARTRIDGE STORE  MSFO, SSID 100, MSFID = 7
CONNECTIONS FOR SA 2 SSID 820
   INTERFACE A TO MSC, PORT 0, CONTROL UNIT ADDRESS 010
   INTERFACE B TO CPU 0202220158, CHANNEL 5, CONTROL UNIT ADDRESS 01
   INTERFACE C TO CPU 0204440168, CHANNEL 2, CONTROL UNIT ADDRESS 01
   INTERFACE D IS NOT USED
     DASD  (CTL-I)
           (SHARED SA 3)           (SHARED SA 3)
           CONTROLLER 00           CONTROLLER 01           CONTROLLER 10           CONTROLLER 11

           PLUG   TYPE    SSID     PLUG   TYPE     SSID    PLUG   TYPE     SSID    PLUG   TYPE     SSID
             0   3330-11  010        0   3350                0                       0
             1   3330-11  012        1   3350                1                       1
             2   3330-11  014        2   3350                2                       2
             3   3330-11  016        3   3350                3                       3
             4   3330-11  ***        4   *****               4                       4
             5   3330-11  ***        5   *****               5                       5
             6   *****    ***        6   *****               6                       6
             7   *****    ***        7   *****               7                       7
     DRC/DRD (CTL-I/ADT)
          DRD CONTROLLER ADDRESS 00 TO MSF 0, DRC 0,B
                DRC 0    SSID 400
                DRD 00   SSID 200
                DRD 01   SSID 201
                DRD 10   SSID 202
                DRD 11   SSID 203
          DRD CONTROLLER ADDRESS 01 IS UNUSED
          DRD CONTROLLER ADDRESS 10 IS UNUSED
          DRD CONTROLLER ADDRESS 11 IS UNUSED
                   MSC TABLE CREATE (VS1)
CONNECTIONS REPORT
CONNECTION FOR CARTRIDGE STORE  MSFO, SSID 100, MSFID = 7
CONNECTIONS FOR SA 3 SSID 830
   INTERFACE A TO MSC, PORT 0, CONTROL UNIT ADDRESS 011
   INTERFACE B TO CPU 0202220158, CHANNEL F, CONTROL UNIT ADDRESS 01
   INTERFACE C TO CPU 0204440168, CHANNEL 3, CONTROL UNIT ADDRESS 01
   INTERFACE D IS NOT USED
     DASD  (CTL-I)
           (SHARED SA 2)           (SHARED SA 2)
           CONTROLLER 00           CONTROLLER 01           CONTROLLER 10           CONTROLLER 11

           PLUG   TYPE    SSID     PLUG   TYPE     SSID    PLUG   TYPE     SSID    PLUG   TYPE     SSID
             0   3330-11  010        0   3350                0                       0
             1   3330-11  012        1   3350                1                       1
             2   3330-11  014        2   3350                2                       2
             3   3330-11  016        3   3350                3                       3
             4   3330-11  ***        4   *****               4                       4
             5   3330-11  ***        5   *****               5                       5
             6   *****    ***        6   *****               6                       6
             7   *****    ***        7   *****               7                       7
     DRC/DRD (CTL-I/ADT)
          DRD CONTROLLER ADDRESS 00 TO MSF 0, DRC 1,B
                DRC 1    SSID 401
                DRD 00   SSID 202
                DRD 01   SSID 203
                DRD 10   SSID 200
                DRD 11   SSID 201
          DRD CONTROLLER ADDRESS 01 IS UNUSED
          DRD CONTROLLER ADDRESS 10 IS UNUSED
          DRD CONTROLLER ADDRESS 11 IS UNUSED
                   MSC TABLE CREATE (VS1)
ICG106I BUILD PHASE COMPLETED
ICG109I 0002 WARNING ERRORS, 0000 TERMINATING ERRORS
ICG110I HIGHEST RETURN CODE WAS 04
ICG107I END OF JOB - MSC TABLE CREATE
ICG110I HIGHEST RETURN CODE WAS 04
```

Glossary

The following terms are used throughout this book.

accessor: The component in the Mass Storage Facility that transports cartridges between the cells, the data recording devices, and the cartridge access station.

accessor control: The component in the Mass Storage Facility that decodes and sequences messages from the Mass Storage Control and directs the motion of the accessor.

acquire: To allocate space on a staging drive and to stage data from a cartridge to the staging drive.

active mass storage volume: See active volume.

active volume: A mass storage volume residing within the Mass Storage Facility and available for mounting by the operating system.

attention interrupt: A signal from the Mass Storage Control to the CPU that a message is waiting for the CPU.

base mass storage volume: See base volume.

base volume: A mass storage volume that can have copies or duplicates.

BIND: (1) An attribute of a data set that keeps the data set on one or more staging drives until the data set is released by the user regardless of the length of time or the demands for space. (2) An attribute of a mass storage volume that reserves an entire staging pack for the mass storage volume whenever the volume is mounted and keeps the volume mounted until an explicit DEMOUNT is issued.

cartridge: The storage medium of the Mass Storage System, consisting of a container with magnetic media wound around a spool inside it.

Cartridge Access Station: An opening on the Mass Storage Facility where cartridges are manually loaded or ejected.

cartridge label: An area on the magnetic storage media that contains the cartridge identification and other information about the cartridge.

cartridge serial number: A unique number that identifies a cartridge; recorded magnetically and visibly on the media.

Cartridge Store: The part of the Mass Storage Facility that consists of the cells, the accessors, and the accessor controls.

cell: A hexagonal compartment within the Mass Storage Facility where a cartridge is stored.

cell cube: A block of 32 cells, four X addresses by four Y addresses, by two Z addresses.

convertible drive: A drive that can be designated to be either a staging drive or a real drive.

copy mass storage volume: See copy volume.

copy volume: An inactive mass storage volume that is an exact reproduction of another mass storage volume. Both volumes have the same volume serial number.

cylinder fault: A condition that occurs when the operating system requires data that has not been staged. The cylinder fault causes a cylinder of data to be staged.

DASDERASE: An attribute of a mass storage volume that causes binary zeros to be written on the staging drive after data from the mass storage volume has been destaged.

data cartridge: See cartridge.

data recording control: The component of the Mass Storage Facility that starts and stops data recording devices, encodes and decodes, and assists with error recovery. The abbreviation is DRC.

data recording device: The unit in the Mass Storage Facility that reads and writes data on the cartridge media. The abbreviation is DRD.

default mass storage volume group: The collection of mass storage volumes that belong to a mass storage volume group defined by the Mass Storage System Communicator. The name of the group is always SYSGROUP.

delayed response: An indication from the Mass Storage Control that a Mass Storage Control I/O is finished.

destage: To move data from a staging drive to a mass storage volume.

destage error: The result of a permanent read error when transferring data from the staging pack to the data cartridge.

DRC: See data recording control.

DRD: See data recording device.

duplicate mass storage volume: See duplicate volume.

duplicate volume: An inactive mass storage volume that has the same volume serial number as another mass storage volume and is not a copy.

eject: To move a cartridge from the Mass Storage Facility to the cartridge access station.

EXCLUSIVE: An attribute of a mass storage volume that allows only one CPU at a time to access the mass storage volume.

Extended Group Coded Recording: The technique is used to encode the data on the media in a data cartridge. This technique includes error correction code. The abbreviation is E/GCR.

E/GCR: See Extended Group Coded Recording.

general-use mass storage volume: See general-use volume.

general-use volume: A mass storage volume that is assigned to a mass storage volume group and can be used to satisfy nonspecific requests for a mass storage volume.

group: See mass storage volume group or staging drive group.

IML: See Initial Microprogram Load.

inactive mass storage volume: See inactive volume.

inactive volume: A mass storage volume that is not available for mounting by the operating system.

Initial Microprogram Load: The action of loading a microprogram. The abbreviation is IML.

Integrated Storage Control: A feature on the Model 158 or 168 processors that control the 3330 Disk Storage. With the addition of the Staging Adapter, the Integrated Storage Control can control staging drives. See also Staging Adapter.

Inventory data set: See Mass Storage Volume Inventory data set.

Journal data set: See Mass Storage Volume Control Journal data set.

journalling: Recording transactions against a data set so that the the data set can be reconstructed by applying the transactions in the journal against a previous version of the data set.

Least Recently Used: An algorithm that determines the order in which active staged pages must be destaged. The algorithm ensures the staging drive group will always have the amount of allocatable space defined by the space manager.

loading pattern: The order in which cells are filled with cartridges that are entered in the Mass Storage Facility through the cartridge access station.

loosely-coupled: A connection of more than one CPU such that the CPUs share only channels.

Mass Storage Control: A microprogrammed portion of the Mass Storage Facility that passes information to the accessor control, and controls data and space on staging drives. The Mass Storage Control is abbreviated MSC.

Mass Storage Control Table Create: A program that builds the Mass Storage Control tables. The abbreviation is MSCTC.

Mass Storage Control tables packs: A direct access storage pack that contains Mass Storage Control tables.

Mass Storage Control twin port: A feature of a Mass Storage Control that allows the Mass Storage Control to address a total of 16 of the following: Mass Storage Facilities, 3830 Model 3 Storage Controls, or Integrated Storage Controls that have the addition of the Staging Adapter. Each Integrated Storage Control counts as two.

Mass Storage Facility: The component of a Mass Storage System that contains the storage media and the facilities for accessing the media. The abbreviation is MSF.

Mass Storage System: The name for the entire storage system, consisting of the Mass Storage Facility and all devices that are defined to the Mass Storage Control. The abbreviation is MSS.

Mass Storage System Communicator: A program that handles communication between system control programs and the Mass Storage Control. The Mass Storage Volume Control functions are an integral part of the Mass Storage System Communicator. The abbreviation is MSSC.

mass storage volume: A direct access storage volume residing on two associated cartridges.

Mass Storage Volume Control functions: A collection of functions that reside in the Mass Storage System Communicator and are designed to assist the space manager in managing mass storage volumes and mass storage volume groups. The abbreviation is MSVC.

Mass Storage Volume Control Inventory data set: Same as Mass Storage Volume Inventory data set.

Mass Storage Volume Control Journal data set: A data set that contains messages to the space manager and information used to rebuild the Mass Storage Volume Inventory data set.

mass storage volume group: A collection of mass storage volumes. The space manager can define a group, and the Mass Storage System Communicator defines a default mass storage volume group. The name of the parameter used in job control language for a group is MSVGP.

Mass Storage Volume Inventory data set: A data set that describes mass storage volumes and mass storage volume groups. The abbreviation is MSVI.

MSC: See Mass Storage Control.

MSCTC: See Mass Storage Control Table Create.

MSF: See Mass Storage Facility.

MSS: See Mass Storage System.

MSSC: See Mass Storage System Communicator.

MSVC: See Mass Storage Volume Control.

MSVI: See Mass Storage Volume Inventory data set.

nonstaging drive: Same as real drive.

page: The unit of space that is allocated on a staging drive. The page consists of 8 cylinders.

path: A hardware connection known to the operating system that permits the movement of data within the hardware.

placeholder record: A temporary base or copy volume record that the Mass Storage Volume Control functions create and add to the Inventory data set during the operation of some of the Mass Storage System Access Method Services commands.

primary CPU: The CPU in a multi-CPU system configuration that has the responsibility of processing unsolicited messages from the Mass Storage Control.

real drive: A drive that is attached to a storage control (3830 Model 3) or an Integrated Storage Control with a Staging Adapter but is defined to the operating system and to MSS as a DASD drive, not to be used for staging.

relinquish: To free space on a staging drive. It can cause data to be destaged.

restricted-use mass storage volume: See restricted-use volume.

restricted-use volume: A mass storage volume assigned to a mass storage volume group and used only by requests specifying the mass storage volume identification.

scratch: To remove the information about a mass storage volume from the Mass Storage Volume Inventory data set and put the identification of both cartridges on a list of scratch cartridges.

scratch cartridge: A cartridge that is not part of a mass storage volume.

scratch data cartridge: See scratch cartridge.

SHARED: An attribute of a mass storage volume that allows more than one CPU at a time to access the mass storage volume.

solicited message: A message from the Mass Storage Control to the CPU that is expected by the CPU.

space manager: The person who is responsible for managing space on mass storage volumes.

stage: To move data from a cartridge to a staging drive.

stage error: The result of a permanent read error while trying to read a particular stripe of a data cartridge.

Staging Adapter: (1) An addition to a System/370 Model 158 or 168 Intergrated Storage Control (ISC) feature that enables the ISC to operate in a Mass Storage System. (2) A 3830 Model 2 Storage Control that has been modified to operate in a Mass Storage System. The modification changes the designation of the storage control to a 3830 Model 3 Storage Control.

staging drive: A 3330 Model 1, 2, or 11 that is designated by the Mass Storage Control Table Create program to receive data from a Mass Storage Facility.

staging drive group: A collection of staging drives for space management and recovery. It is created by the user with the Mass Storage Control Table Create program.

staging effective data rate: An amount of data transferred between the data recording devices and the staging drives in one second. The amount of data is normally averaged over an hour.

staging pack: A 3336 Disk Pack that has been initialized to receive data from a Mass Storage Facility.

Storage Control: The 3830 Model 3, the direct access storage device control unit in the Mass Storage System that controls the transfer of data during staging and destaging operations. Also see Staging Adapter.

stripe: The portion of the cartridge media that is accessible to a given head position.

subsystem identification: An identification on each device in the Mass Storage System. The devices include staging adapters, staging drives, Mass Storage Facility, data recording devices, and data recording controls. The abbreviation is SSID.

System Data Analyzer: A program that analyzes collected data about hardware errors in the Mass Storage System.

system effective data rate: An amount of data transferred between the staging drives and the CPU in one second. The amount of data is normally averaged over an hour.

tables pack: See Mass Storage Control tables pack.

tightly-coupled: A connection of more than one CPU such that the CPUs share main storage and communicate directly with one another.

trace: A monitor in the Mass Storage Control that records data about Mass Storage System activity and staging and destaging. The data describes completed Mass Storage System functions from the activity schedule queues plus time stamps.

twin port: See Mass Storage Control twin port.

unsolicited message: A message from the Mass Storage Control to the primary CPU that is not requested or expected by the primary CPU.

virtual drive: A direct access storage device that does not physically exist. It exists logically on one or more staging drives.

virtual unit address: An address for a virtual drive that consists of the channel address, the Staging Adapter address and the device address. The virtual unit address can be assigned to any staging drive group. Each staging drive can have more than one virtual unit address, but only one real unit address.

virtual unit control block: A unit control block that contains a virtual unit address.

virtual volume: The data from a mass storage volume while it is located on a staging drive.

volume: See mass storage volume or virtual volume.

volume group: See mass storage volume group or default mass storage volume group.

WRITEINHIBIT: An attribute of a mass storage volume that prevents writing on the mass storage volume. It means the same as read-only.

Bibliography

Required Publications

IBM 3850 Mass Storage System (MSS), Principles of Operation GA32-0029

Introduction to the IBM 3850 Mass Storage System (MSS), GA32-0028

OS/VS Mass Storage Control Table Create, GC35-0013

OS/VS Mass Storage System (MSS) Planning Guide, GC35-0011

OS/VS Mass Storage System (MSS) Services: General Information, GC35-0016*

OS/VS Publications

Operator's Library: IBM 3850 Mass Storage System (MSS) under OS/VS, GC35-0014

OS/VS Mass Storage System (MSS) Services: Reference Information, GC35-0017*

OS/VS Mass Storage System (MSS) Services Logic, SY35-0015

OS/VS Message Library: Mass Storage System (MSS) Messages, GC38-1000

OS/VS Message Library: Routing and Descriptor Codes, GC38-1004

OS/VS System Management Facilities (SMF), GC35-0004

OS/VS Utilities, GC35-0005

OS/VS VSAM Programmer's Guide, GC26-3838

OS/VS1 Publications,

Operator's Library: OS/VS1 Display Console, GC38-0255

OS/VS1 Access Method Services, GC26-3840

OS/VS1 JCL Reference, GC24-5099

OS/VS1 JCL Services, GC24-5100

OS/VS1 System Programming Library: SYS1.LOGREC Error Recording, GC28-0668

OS/VS1 Mass Storage System Communicator (MSSC) Logic, SY35-0012

OS/VS Message Library: VS1 System Messages, GC38-1001

OS/VS1 Planning and User's Guide, GC24-5090

---

*Can be ordered in a set as GB0F-3579.

OS/VS1 Supervisor Services and Macro Instructions, GC24-5103

OS/VS1 System Programming Library: System Generation Reference, GC26-3791

OS/VS2 Publications

Operator's Library: OS/VS2 Display Console, GC38-0260

OS/VS2 Access Method Services, GC26-3841

OS/VS2 Conversion Notebook, GC28-0689

OS/VS2 JCL Reference, GC28-0692

OS/VS2 System Programming Library: Job Management VS2 Release 3, GC28-0627

OS/VS2 Mass Storage System Communicator (MSSC) Logic, SY35-0013

OS/VS Message Library: OS/VS2 System Messages, GC38-1002

OS/VS2 Planning Guide for Release 2, GC28-0667

OS/VS2 Supervisor Services and Macro Instructions, GC28-0683

OS/VS2 System Programming Library: SYS1.LOGREC Error Recording, GC28-0677

OS/VS2 System Programming Library: System Generation Reference, GC26-3792

Related Publications

DB/DC MSS Planning Information: IMS/VS, CICS/VS, and GIS/VS, GH20-9048

IBM 3850 Data Cartridge Care and Handling Instructions, GA32-0031

IBM System/370 Installation Manual—Physical Planning, GC22-7004

IBM System/370 Model 168 Functional Characteristics, GA22-7010

Index

The entries in this index are shown as they appear in the text of the book, that is, capitalized words are capitalized in text, lowercase words are lowercase in text, etc. When more than one reference is given for an entry, the primary reference is listed first.

GA32-0030-2

IBM 3850 Mass Storage System (MSS) Installation Guide Printed in U. S. A. GA32-0030-2

IBM®
International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, New York 10604
(U. S. A. only)

IBM World Trade Corporation
821 United Nations Plaza, New York, New York 10017
(International)

IBM Technical Newsletter

This Newsletter No.	GN32-0017
Date	August 15, 1977
Base Publication No.	GA32-0030-2
File No.	S370-30
Previous Newsletters	None

IBM 3850 Mass Storage
System (MSS)
Installation Guide

This Technical Newsletter provides replacement pages for the subject publication. These pages remain in effect for subsequent releases unless specifically altered. Pages to be inserted and/or removed are:

Pages to be Removed	Attached Pages to be Inserted*
cover, edition notice	cover, edition notice
v, vi	v, vi
ix–xii	ix–xii
1, 2	1, 2
27, 28	27–28.1
73, 74	73, 74
105, 106	105, 106
111, 112	111, 112
115–118	115–118
161, 162	161, 162
213, 214	213, 214
283–back cover	283–back cover

*If you are inserting pages from different newsletters or supplements and *identical* page numbers are involved, always use the page with the latest date. The page with the latest date contains the highest level of information.

A change to the text is indicated by a vertical line to the left of the change.

Summary of Amendments
This Technical Newsletter incorporates miscellaneous editorial and technical changes. Each technical change is marked by a vertical line to the left of the change.

Note: *Please file this cover letter at the back of the manual to provide a record of changes.*

IBM Corporation, Information Planning and Editing, Dept. 21E, Boulder, Colorado 80302

Printed in U. S. A.