M. Haupt-81

IBM® Technical Newsletter

This Newsletter No.	GN32-0011
Date	September 27, 1976
Base Publication No.	GA32-0028-2
File No.	S370-07
Previous Newsletters	None

Introduction to the Mass Storage System IBM 3850 Mass Storage System (MSS)

This Technical Newsletter provides replacement pages for the subject publication. These replacement pages remain in effect for subsequent releases unless specifically altered. Pages to be inserted and/or removed are:

47–54

A change to the text or to an illustration is indicated by a vertical line to the left of the change.

Summary of Amendments

Miscellaneous editorial and technical changes have been made on the pages to be inserted or replaced. 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, MSS Support Center, Department 25E, Boulder, Colorado 80302

Printed in U. S. A.

System Availability and Data Integrity

Computer-based systems have become increasingly important to businesses, governments and other organizations. For this reason, increasing attention has been given to systems availability and data integrity.

System availability is having the system when you want it. Elements of system availability are:

- *Reliability* of the components.
- *Continuity* of system functions
- *Serviceability* of the components

Reliability means that the system stays up. When components become unreliable, they may disrupt continuity of system functions. When that happens, the serviceability of the components becomes an important consideration.

Reliability is the measure of the probability—at best an estimate—that the system will do what it is designed to do for a given period of time.

Continuity involves answers to:

- How often will a malfunction occur?
- What effect will malfunctions have on system capability?
- How long will it take to fix?
- How often will preventive maintenance be performed?
- When will engineering improvements be installed?

Serviceability questions are phrased as how easy and fast is a system to fix and maintain and what effect does restoring one element have on another.

Data Integrity means the ability of the system to store, maintain, update, and move data without alteration due to a malfunction. Important data integrity considerations are:

- *Recovery* of data after system failures.
- *Data Security* from unauthorized access, change, or exposure.

Recovery is an automated function in the sense of returning to a normal state after a temporary error in the system. Recovery is also the process of manually restoring a system when it is unable to recover automatically.

Data security refers to protecting data from unauthorized disclosure, modification or destruction by accidental or intentional means.

System Perspective

To look at these aspects of a specific computer-based system, both IBM and the user must have the same understanding and perception of the system. A *system* is not merely an interconnection of computer hardware units, software, and firmware. A system should be perceived—by everyone concerned with it—as a dynamic combination of resources:

- Hardware
- Software
- Firmware
- Data
- Operators
- Application programs
- Normal operating procedures
- Contingency procedures
- Computer room environment
- Management

Note that a key concept is that all systems continually undergo change. A system running a payroll at ten o'clock, for example, is different from the "same" system running linear programs at eleven o'clock. Even the hardware changes, because what was a critical unit for the payroll (for example, a printer) is possibly not even used in the linear-programming application.

When a question is asked regarding the systems availability or integrity characteristics of a system, the questioner is really asking "management" questions about the control of a full set of resources to do his job. The key to meeting security, integrity, and system availability requirements is not only good hardware and systems design, but also includes good availability management, contingency procedures, well trained operators, and intelligent planning.

IBM's 3850 Mass Storage System, properly employed with good systems design, provides an opportunity for improved resource management over conventional tape/DASD in these areas:

1. Tape handling is completely automated. This reduces problems caused by manual handling as dirty leaders, damaged tape, bent reel flanges, etc.
2. Improvements in the recording quality of the tape, reduction in the recording head-to-tape separation, and advances in error correction code provide not only improved data density on the cartridge but more reliable reproduction of the data compared to tape/DASD.
3. Because much more data can be under system(s) control when compared to tape/DASD, normal operating occurrences such as misfiled, mislabeled or lost reels of tape and mismounts are nearly eliminated.
4. Contention for a unique tape data set is also reduced because the data, once staged, can be shared among two, three or four operating systems from DASD, assuming that the data set has been qualified as sharable.
5. Extensive hardware replication and sophisticated, automatic error recovery procedures are integral to the design of the Mass Storage Facility.
6. The Mass Storage System has been designed so that most maintenance can be accomplished concurrently with productive use of the system.

Mass Storage System Availability

Description of MSS availability must start with a review of the general architecture of the system. The following schematic (Figure 14) shows a Mass Storage System consisting of one 3851 Mass Storage Facility. The MSS is composed of

eight major groups of components indicated by the columns of boxes. The numbers in the schematic indicate the maximum number of each type of component that can be included in an MSS containing one Mass Storage Facility.

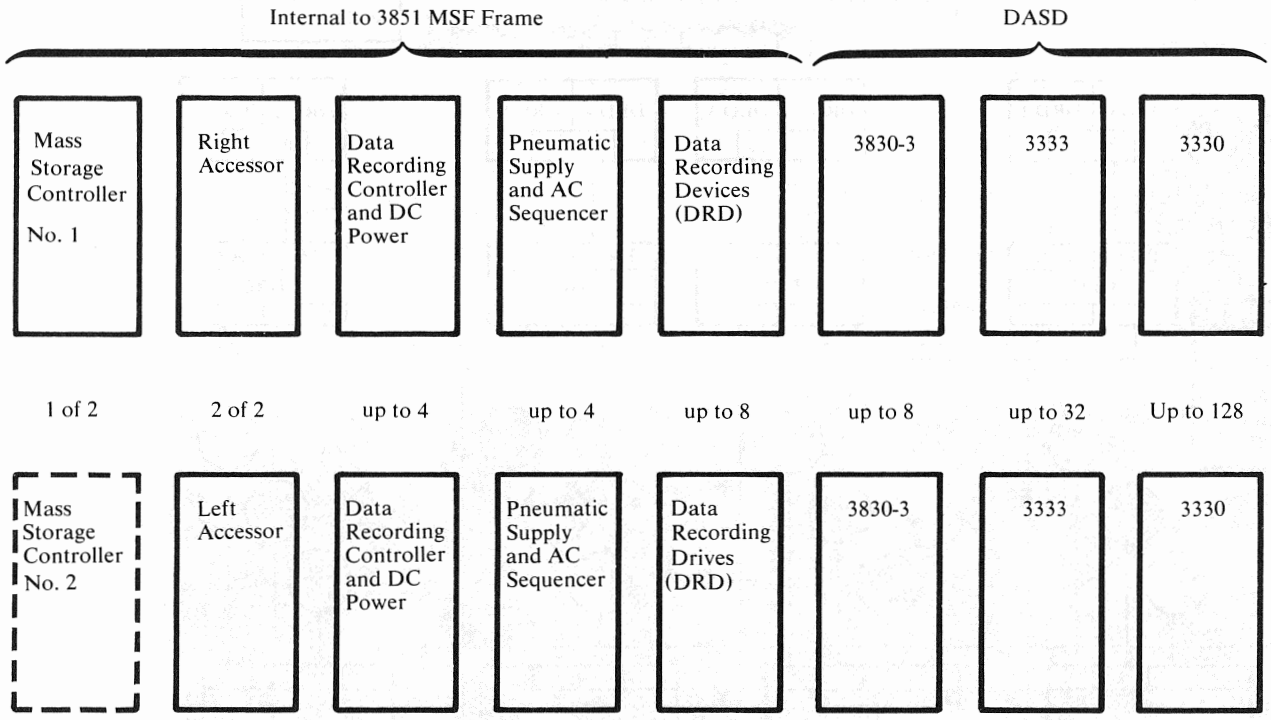

Figure 14. IBM 3850 Mass Storage System

The MSS is extensively replicated. Any cartridge can be mounted on any data recording device which is accessible to the 3830-3 selected by the MSC. Therefore, with alternate paths in a system configured for availability, the impact of any single unit failure and many multiple unit failures usually will not cause the system to stop. Rather, it will operate at a degraded data rate until the failure is corrected.

Mass Storage Facility—Redundancy in Data Paths

Two accessors and their associated controls and power supplies are included in each MSF. If one fails, an error recovery program in microcode in the MSF causes the functioning accessor to push the failing accessor into a "garage" to await service. The remaining accessor continues to operate, and the effect in Models 2, 3, and 4 is to slightly reduce cartridge access rate. The failing accessor can normally be serviced in its "garage".

Except in Model 1, each data recording device is cabled to its primary DRC and to another DRC (Figure 15). Each pair of data recording devices is cabled to its own power supply and pneumatic system. If power or pneumatics fail, only that pair of DRDs is disabled. In the models 2, 3 and 4, the system continues to operate with a potentially degraded data rate.

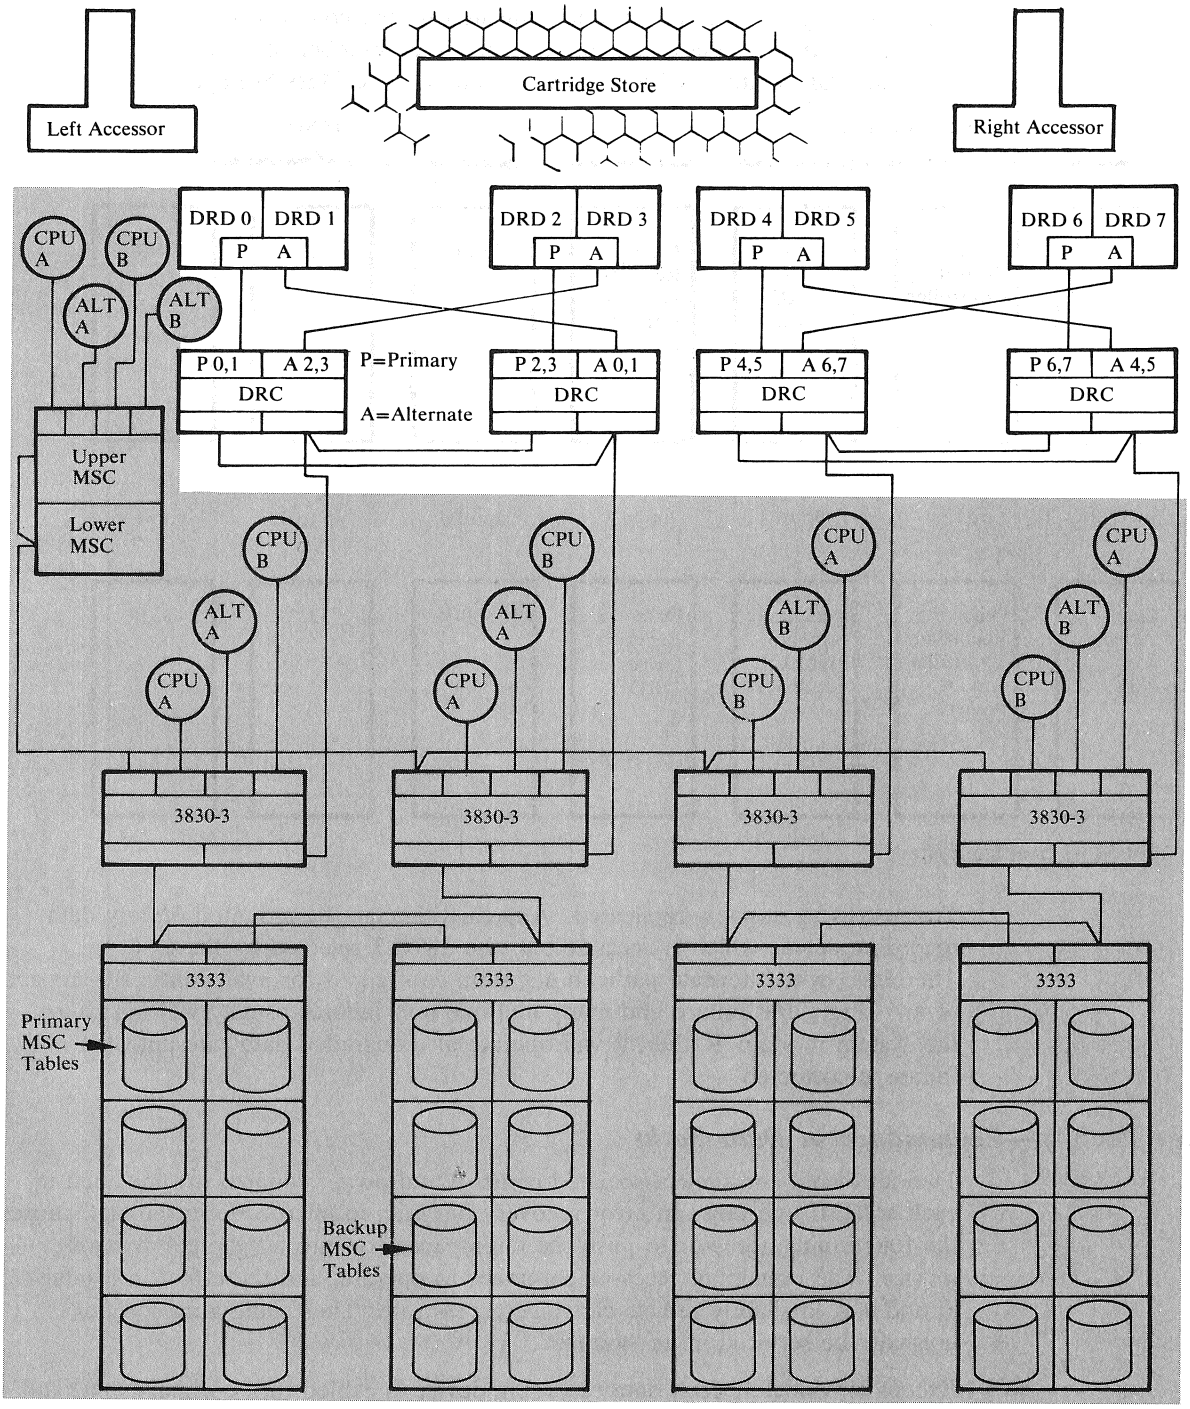

Notes:

1. Each 3830-3 is accessible to two Data Recording Controls.
2. One host system can be down and the other is operational.
3. The configuration can lose any one 3333 and still be operational.

Figure 15. Mass Storage Facility Data Path Redundancy

If the primary DRC fails, the switch to the backup DRC is automatic.

In turn, each data recording control can have access to two 3830 Model 3 Storage Controls and each Storage Control can attach to up to four DRCs. DRCs also have independent power supplies, so that if one power supply fails it does not disable the other DRCs in an MSF.

For most situations, reconfiguration is automatic through error recovery procedures in the MSC microcode without dependency on the host system(s). The failing unit is marked unusable and the operator is notified. He may vary it offline. The Mass Storage System is designed to permit maintenance and check-out on the unit varied off line without interrupting the rest of the storage system.

Mass Storage Facility—Redundancy in Control Path

Each Mass Storage Facility contains one Mass Storage Control (MSC) and can optionally contain a second one (Figure 16). (In the event of an MSC failure, switchover to the backup MSC is automatic.) Only one is active at a time. Both MSCs are powered and storage is loaded at power-on time of the MSF. The MSC which has the lower port becomes the primary and completes an initial microprogram load (IML). The backup MSC loops in a master dispatcher loop until it receives one of a select group of commands.

When an MSC failure is detected by a host, operating system error recovery procedures issue commands to switch from the on-line MSC to the backup MSC. The backup MSC then completes the same IML as the first MSC. Upon completion of the IML, a Device End is issued. The operating system will then issue the INITIALIZE and HOST READY FOR MESSAGES orders. Upon completion of these orders, the CCW that failed is retried, now using the backup MSC.

This entire switchover procedure occurs in seconds and will not interrupt jobs already in process on the host CPUs. Control messages that were in process at the time of failure of the first MSC will be reestablished by the backup.

This switchover procedure operates identically between the two MSCs in a "B" series Mass Storage Facility or between each of the MSCs in two "A" series Mass Storage Facilities.

Customer Engineering action is necessary to check out the failing MSC, take appropriate repair action, and again load its storage. This is done without interrupting the use of the Mass Storage System.

DASD

When configured for availability, the DASD allows for alternate data paths between each Data Recording Control and 3830-3 and from there to each 3333 Disk Storage and Control. Failure to properly configure may cause an unnecessary reduction in available data paths, thereby reducing the storage system's capacity to deliver data when operating in a degraded mode or even disabling the MSS if the Mass Storage Control tables are not accessible.

Each 3333 should be cabled to a pair of 3830 model 3 Storage Controls and each 3830-3 should be cabled to a pair of data recording controls.

Figure 17 shows an example of a complex configuration involving two hosts, four 3830-3's, four 3333's and 32 disk drives.

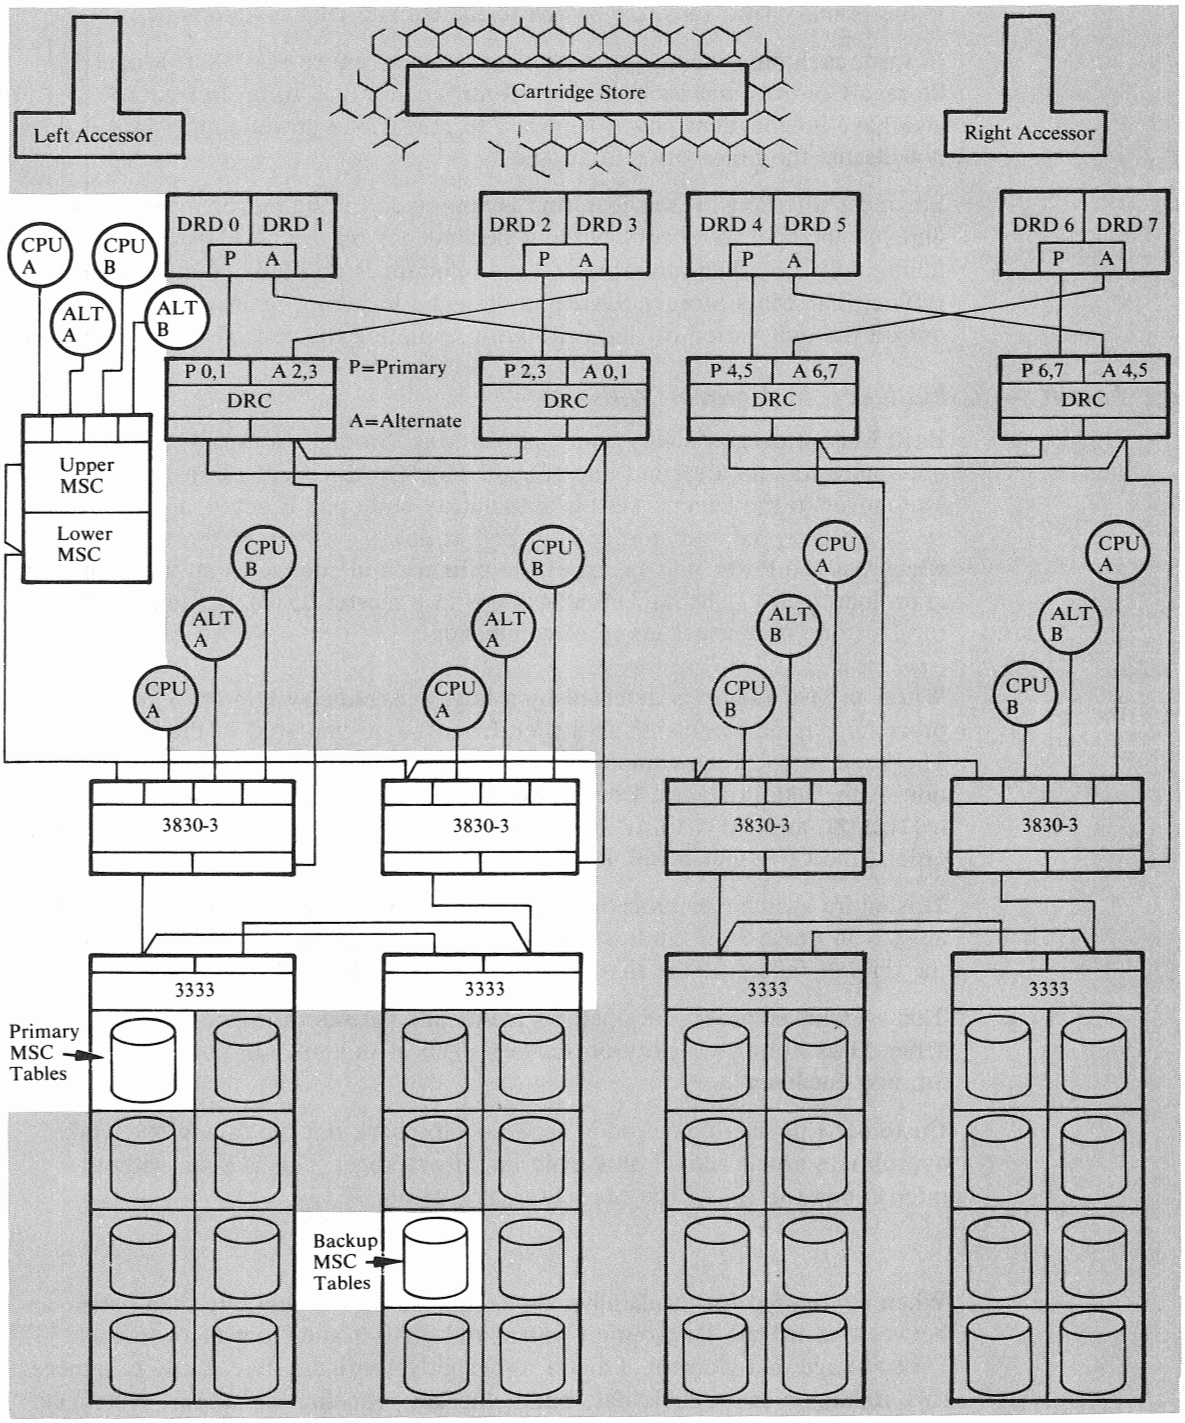

Notes:

1. Each 3830-3 is accessible to two Data Recording Controls.
2. One host system can be down and the other is operational.
3. The configuration can lose any one 3830-3 and with manual intervention still be operational.
4. The configuration can lose any one 3333 and still be operational.

Figure 16. Mass Storage Facility Control Path Redundancy

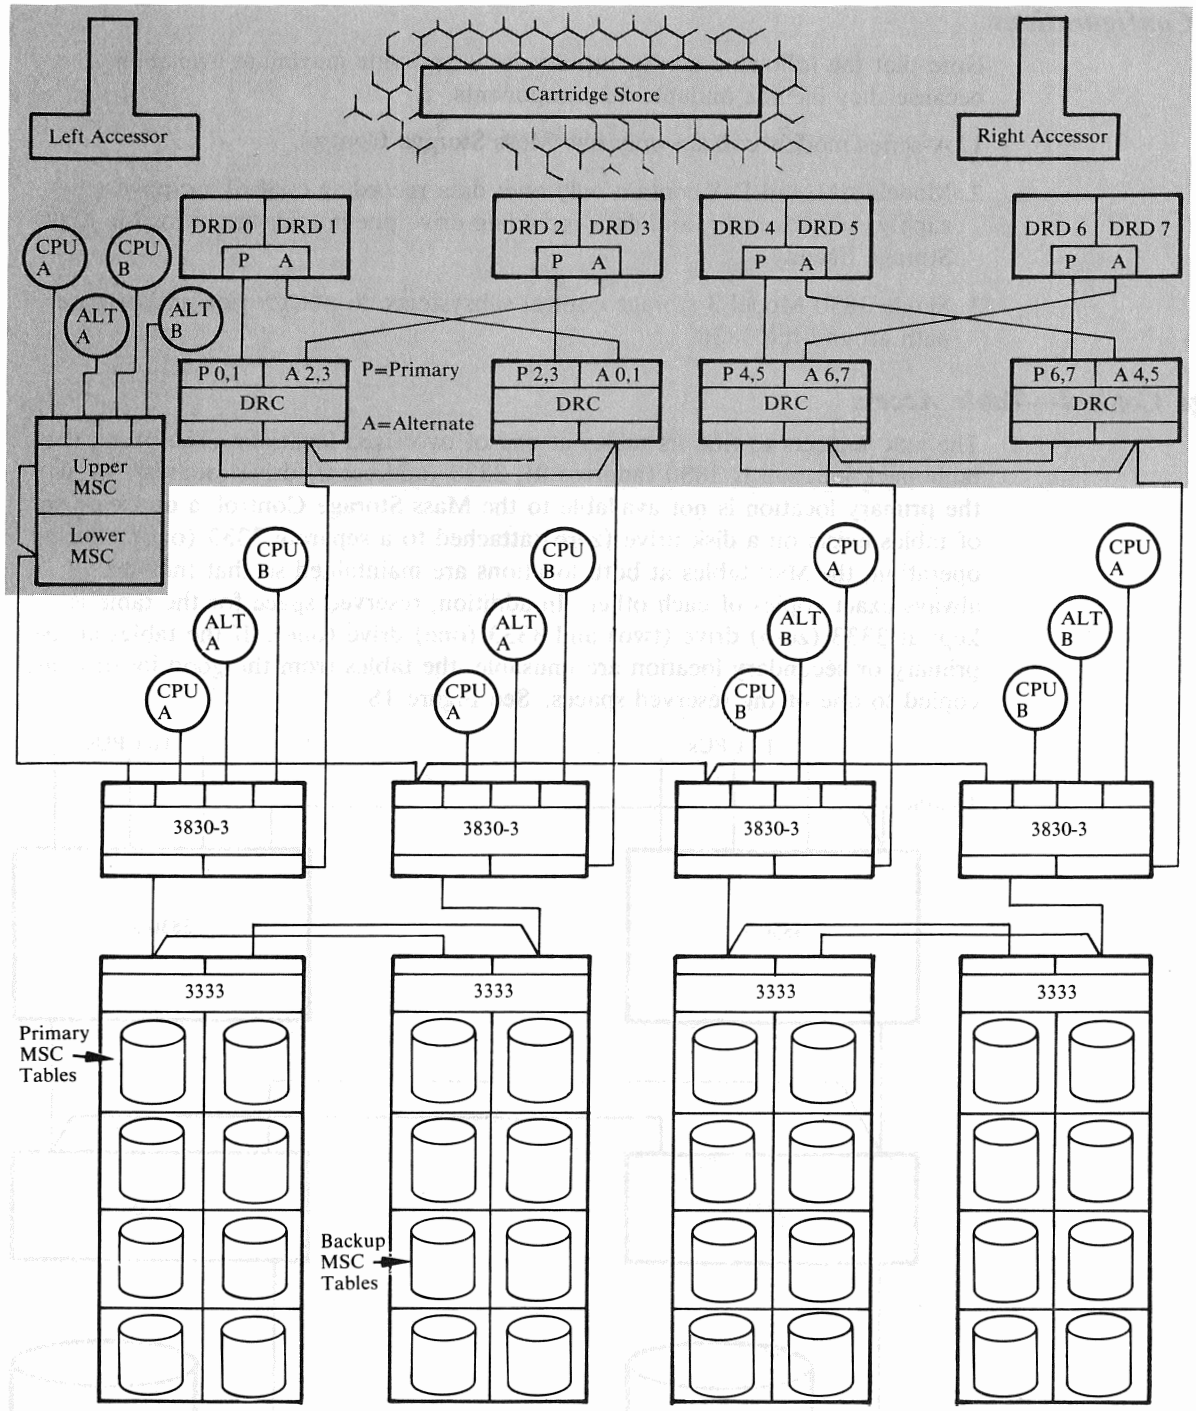

Notes:

1. Each 3830-3 is accessible to two Data Recording Controls.
2. One host system can be down and the other is operational.
3. The configuration can lose any one 3830-3 and with manual intervention still be operational.
4. The configuration can lose any one 3333 and still be operational.

Figure 17. Complex Mass Storage System DASD Configuration

Unduplexed Configurations

Note that the following configurations do not provide maximum availability because they include unduplexed components.

1. A-series models contain only one Mass Storage Control.
2. Models A-1 and B-1 contain only one: data recording control, DC power supply, AC sequencer, and data recording drive pneumatic supply for the Mass Storage Facility.
3. Single 3830 Model 3 storage control subsystems do not provide an alternate path around the 3830.

Mass Storage Control—Table Access

The MSC expects to find its tables at one of two fixed locations. The primary table pack location is 3830 (address 0), 3333 (address 0) drive (address 0). If the primary location is not available to the Mass Storage Control, a duplicate set of tables exists on a disk drive (zero) attached to a separate 3333 (one). During operation, the MSC tables at both locations are maintained so that they are always exact copies of each other. In addition, reserved space for the table is kept at 3333 (zero) drive (two) and 3333 (one) drive (one). If the tables at the primary or secondary location are unusable, the tables from the good location are copied to one of the reserved spaces. See Figure 18.

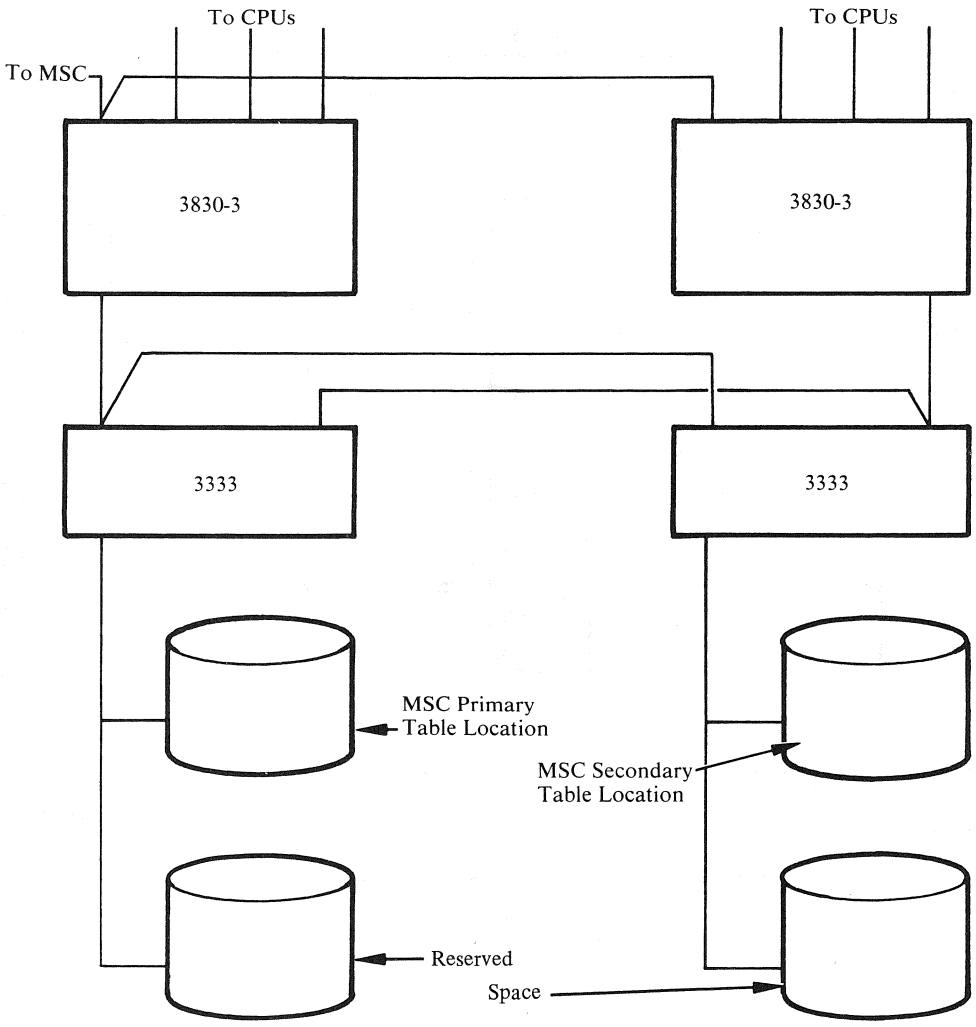

Figure 18. MSC Duplicate Table Pack Locations