International Technical Support Organization

SG24-4563-00

IBM RAMAC Array Family Additions (RAMAC 2)

October 1995

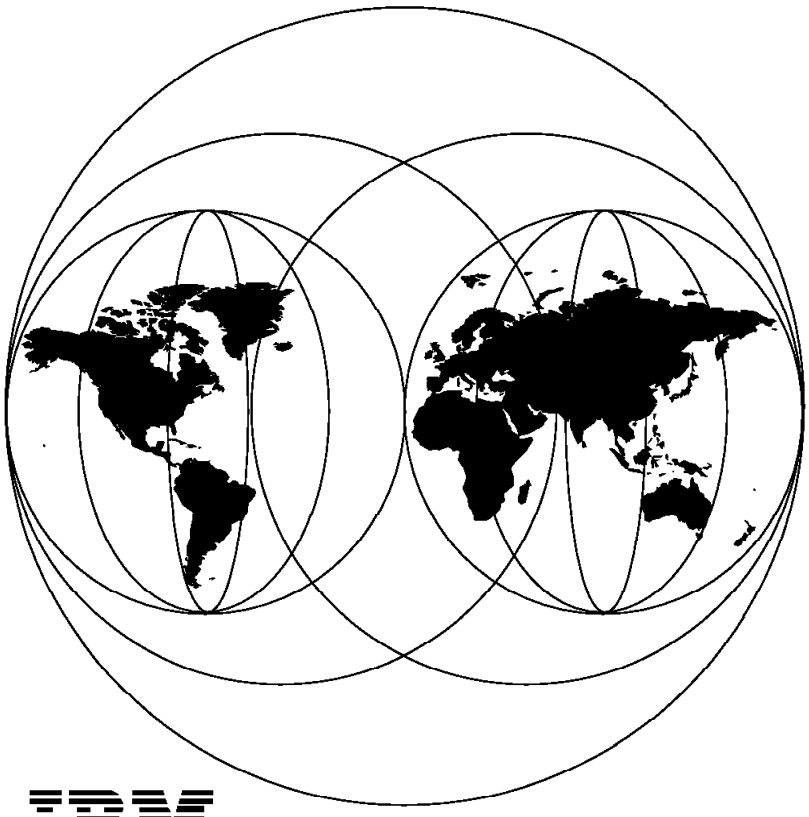

IBM

International Technical Support Organization

San Jose Center

International Technical Support Organization

SG24-4563-00

IBM RAMAC Array Family Additions (RAMAC 2)

October 1995

Take Note!

Before using this information and the product it supports, be sure to read the general information under “Special Notices” on page vii.

First Edition (October 1995)

This edition applies to the IBM RAMAC Array Family of products that use RAMAC 2 drawers.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address given below.

An ITSO Technical Bulletin Evaluation Form for reader's feedback appears facing Chapter 1. If the form has been removed, comments may be addressed to:

IBM Corporation, International Technical Support Organization
Dept. 471 Building 070B
5600 Cottle Road
San Jose, California 95193-0001

Abstract

This document describes the features and functions of the IBM RAMAC Array Family enhancements. It provides information about the IBM RAMAC 2 Array Subsystem and the IBM RAMAC 2 Array DASD that can be used to help customers plan and prepare their implementation. The technology and design changes made to the RAMAC 2 drawer are described, as well as data migration, installation, upgrade options, and DASD configuration considerations.

This document is written for storage administrators, systems programmers, and other technical support staff involved with and interested in storage subsystems.

(100 pages)

Contents

Special Notices

This publication is intended to help storage administrators, systems programmers, and other technical support staff evaluate and plan for the implementation of the IBM RAMAC Array family of products. The information in this publication is not intended as the specification of any programming interfaces that are provided by the IBM RAMAC Array family. See the PUBLICATIONS section of the IBM Programming Announcement for the IBM RAMAC Array family for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM′s product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM′s intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer′s ability to evaluate and integrate them into the customer′s operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

AIX	AIX/ESA
DB2	DFSMS
DFSMS/MVS	DFSMSdfp
DFSORT	ECKD
ESA/370	ESA/390
ESCON	IBM
MVS/DFP	MVS/ESA
MVS/SP	MVS/XA
NetView	Service Director
VM/ESA	VM/XA

VSE/ESA

The following terms are trademarks of other companies:

Windows is a trademark of Microsoft Corporation.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

C-bus is a trademark of Corollary, Inc.

Other trademarks are trademarks of their respective companies.

Preface

This document is intended to help customers plan for and evaluate the IBM RAMAC Array Family of products. It describes the features and functions of the IBM RAMAC 2 Array Subsystem and the IBM RAMAC 2 Array DASD.

This document is intended for storage administrators, systems programmers, and other technical support staff involved in the planning, evaluating, and implementing of storage subsystems.

Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this document.

- *Using the IBM RAMAC Array Subsystem in an MVS, VM, or VSE Environment*, GC26-7005-00
- *IBM RAMAC Array Subsystem Introduction*, GC26-7004
- *IBM RAMAC Array Subsystem Reference*, GC26-7006
- *ICKDSF R16 User's Guide and Reference*, GC35-0033
- *IBM 3990 Storage Control Introduction*, GA32-0098
- *Using the IBM RAMAC Array DASD in an MVS, VM, or VSE Environment*, GC26-7013
- *IBM 3990 Operations and Recovery Guide*, GA32-0253
- *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide*, GA32-0100

International Technical Support Organization Publications

- *IBM RAMAC Array Family*, GG24-2509
- *Extended Count Key Data and Nonsynchronous I/O*, GG24-3571
- *IBM 3990 Storage Control ESCON Features Presentation Guide*, GG24-3803

A complete list of International Technical Support Organization publications, with a brief description of each, may be found in:

International Technical Support Organization Bibliography of Redbooks, GG24-3070.

To get a catalog of ITSO technical publications (known as “redbooks”), VNET users may type:

```
TOOLS SENDTO WTSCPOK TOOLS REDBOOKS GET REDBOOKS CATALOG
```

How to Order ITSO Technical Publications

IBM employees in the USA may order ITSO books and CD-ROMs using PUBORDER. Customers in the USA may order by calling 1-800-879-2755 or by faxing 1-800-284-4721. Visa and Master Cards are accepted. Outside the USA, customers should contact their local IBM office.

Customers may order hardcopy ITSO books individually or in customized sets, called GBOFs, which relate to specific functions of interest. IBM employees and customers may also order ITSO books in online format on CD-ROM collections, which contain books on a variety of products.

How to Obtain a Copy of the Full-Size Foils

The full-size foils of this presentation are available from your IBM representative. IBM employees may obtain a package of materials for this presentation by typing the following command from their local VM user IDs:

```
TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET SG244563 PACKAGE
```

This package may vary and may contain LIST3820 foils in various formats and DRW files, as the author determines.

Acknowledgments

This project was designed and managed by:

Dave McAuley
International Technical Support Organization, San Jose Center

The authors of this document are:

Dave McAuley, ITSO, San Jose
Roland Wolf, IBM Germany

This publication is the result of a residency conducted at the International Technical Support Organization, San Jose Center.

Thanks to Dave Foote, Steve Gerdt, Barrie Harding, Scott Karstetter, Mike Koester, Vern Legvold, Gary Morain, and Jeff Steffan of the IBM Storage Systems Division, and Ed Grochowski of the IBM Almaden Research Center for the invaluable advice and guidance provided in the production of this document.

Thanks are also due to the many other collaborators and reviewers who contributed to the production of this document. In particular, we would like to thank Jack Flynn, Mike Harrison, Bruce McNutt, Chris Saul, and Bob Shomler of the IBM Storage Systems Division, and Maggie Cutler, ITSO Technical Editor.

Introduction

Acknowledgments

This presentation material is the result of a residency project run at the International Technical Support Organization, San Jose Center, from March 13 through April 7, 1995.

Project leader

Dave McAuley, ITSO - San Jose

Resident

Roland Wolf, IBM Germany

Copyright IBM Corp. 1994

ITSC - San Jose

Audience

This presentation is written for:

- Customer personnel
- IBM technical specialists
- IBM client representatives

 ITSC - San Jose

Agenda

- Overview of the RAMAC 2 Array Family enhancements
- Enhanced high-capacity disk drive
- Enhanced RAMAC 2 Array drawer
- RAMAC 2 Array DASD specifics
- RAMAC 2 Array Subsystem specifics
- Configuration options
- Performance
- Operational considerations
- Installation and migration
- Software
- Ordering considerations
- Bibliography
- Summary

 ITSC - San Jose

.

Agenda

This presentation describes the features and functions of the IBM RAMAC 2 Array Family enhancements.

- **Overview of the RAMAC 2 Array Family enhancements**

  This section of the book provides an overview of the enhancements for the RAMAC 2 Array Family.

- **Enhanced high-capacity disk drive**

  This section describes the most important enhancement—the B23 drawer model with twice the capacity of a B13 drawer.

- **Enhanced RAMAC 2 Array drawer**

  In this section all of the enhancements applied to the B23 drawers are discussed.

- **RAMAC 2 Array DASD specifics**

  LIC changes and Dynamic Sparing specifics for the RAMAC Array DASD are discussed in this section.

- **RAMAC 2 Array Subsystem specifics**

  This section reviews several enhancements to the RAMAC Array Subsystem.

- **Configuration options**

  The flexible addressing schemes supported by the 3990-6 for RAMAC 2 are described in this section.

- **Performance**

  This section reviews the various changes that have been applied to the RAMAC 2 drawers and the storage controls to preserve RAMAC performance for the high-capacity B23 drawers.

- **Operational considerations**

  DEVSERV and service information messages (SIMs) have been changed to report details of the real device type in addition to the device type being emulated. This section discusses those changes.

- **Installation and migration**

  This section discusses the installation steps and the role of ICKDSF for RAMAC devices. Environmental data for RAMAC subsystems is presented.

- **Software**

  The installation of RAMAC devices is transparent to software. Some PTFs are recommended for performance reasons, however. This section reviews those PTFs and other recommendations for RAMAC devices.

- **Ordering considerations**

  This section provides a list of feature codes for the RAMAC Array Family.

- **Bibliography**

  This section lists manuals to consult for more information about the RAMAC Array Family and related products.

- **Summary**

  This section summarizes the value of the enhancements to the RAMAC Array Family.

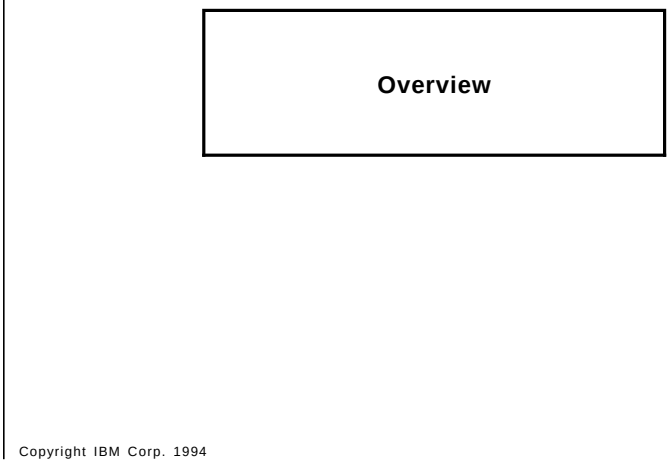

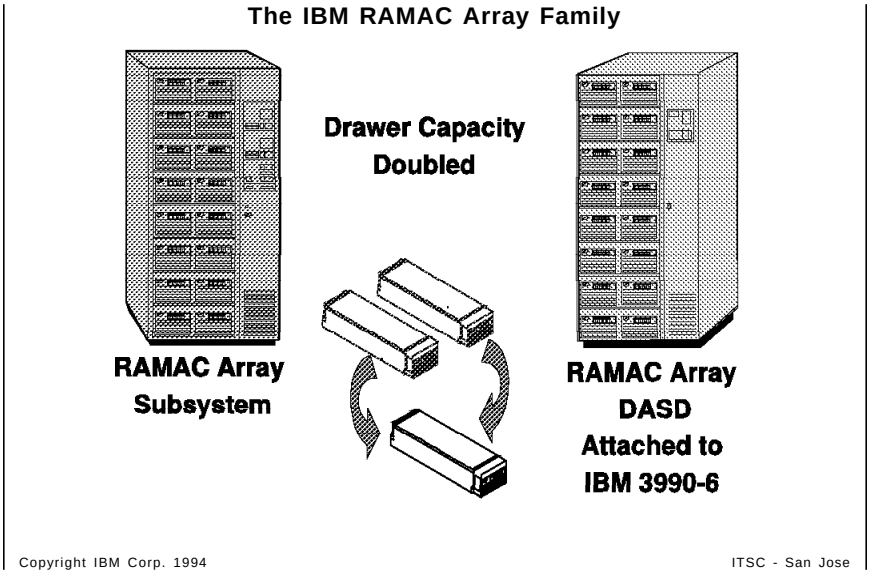

OVER-6 The IBM RAMAC Array Family

The RAMAC Array Subsystem and the RAMAC Array DASD, both members of the RAMAC Array Family, can be equipped with new RAMAC 2 drawers that have significantly improved data capacity over the previous RAMAC drawers. The RAMAC 2 drawers have double the capacity of the RAMAC drawers. This increase in capacity allows for up to 180 GB in a single rack.

- **IBM RAMAC Array Subsystem**

  The RAMAC Array Subsystem provides up to 180 GB of fault-tolerant storage and all storage control functions in a single physical unit.

  The RAMAC Array Subsystem offers fault tolerance through redundant hardware and RAID-5 data storage techniques. The combination of the integrated storage control and 3.5-inch high capacity disks results in

significant environmental (floor space, power, cooling) savings compared with existing 3380, 3390, or RAMAC strings.

- **IBM RAMAC Array DASD**

  The RAMAC Array DASD offers up to 180 GB of fault-tolerant storage in 3390 model 3 format in a single rack. If the RAMAC Array DASD rack is equipped with only RAMAC drawers, it attaches to IBM′s 3990 model 3 or model 6 storage control. If the rack has some RAMAC 2 drawers, it only attaches to IBM′s 3990 model 6 and takes advantage of all of the functions offered by the IBM 3990 model 6 storage control. The RAMAC Array DASD uses RAID-5 data storage techniques and has redundant hardware to greatly improve data availability.

- **IBM RAMAC Array Drawer**

  The RAMAC Array Drawer is used in both the RAMAC Array Subsystem and the RAMAC Array DASD. The drawer contains four disk drives, a nonvolatile drawer cache, and multiple microprocessors to direct the independent functions of the drawer. The drawer can contain four 2 GB disk drives. Drawers with four 2 GB disk drives are called RAMAC drawers in this publication. The drawers can also contain four 4 GB disk drives. Drawers with four 4 GB disk drives are called RAMAC 2 drawers in this publication. The drawer performs volume emulation, RAID-5 function, and drawer level caching. A mix of RAMAC and RAMAC 2 drawers can be used in a RAMAC Array DASD as well as in a RAMAC Array Subsystem.

  Each drawer in a RAMAC Array Subsystem or a RAMAC Array DASD configuration can be considered an independent parallel processing array capable of delivering significant function but at the same time working closely with its subsystem attachment.

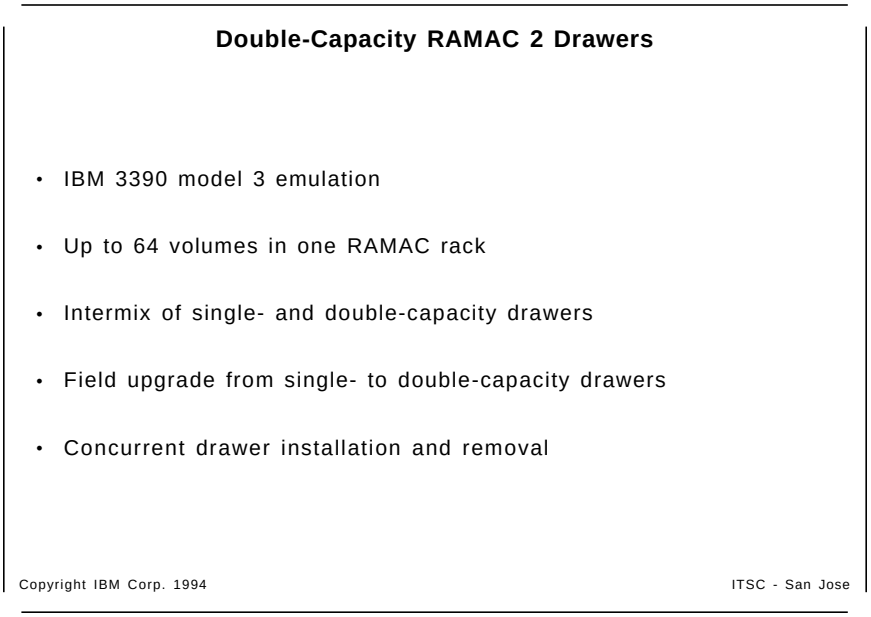

OVER-7 Double-Capacity RAMAC 2 Drawers

The use of double-capacity disk drives in RAMAC 2 drawers results in doubling the number of volumes emulated. Increasing the number of logical volumes is a simple way of providing additional capacity.

- **IBM 3390 model 3 emulation**

  If a RAMAC Array Subsystem is equipped with RAMAC drawers, it offers storage in two different device geometries: 3390 model 3 or 3380 model K attached to a 3990-2 control unit. If the RAMAC Array Subsystem is equipped with one or more RAMAC 2 drawers, only the 3390 model 3 emulation is supported.

  For the RAMAC Array DASD, only 3390 model 3 emulation is available for RAMAC and RAMAC 2 drawers.

- **Up to 64 volumes in one RAMAC rack**

  The RAMAC 2 drawer emulates four logical 3390 model 3 volumes. The address mapping of an IBM 3990 model 6 and a RAMAC Array Subsystem is enhanced to support up to 64 volume addresses in a single rack.

- **Intermix of single- and double-capacity drawers**

  RAMAC and RAMAC 2 drawers can be mixed in the same rack for both the RAMAC Array DASD and RAMAC Array Subsystem.

- **Field upgrade from single- to double-capacity drawers**

  Installed RAMAC drawers can be upgraded in the field to RAMAC 2 drawers. For the RAMAC Array DASD, dual copy can be used to nondisruptively move the data before the upgrade is applied.

- **Concurrent drawer installation and removal**

  Additional drawers can be added and removed concurrently with I/O operations.

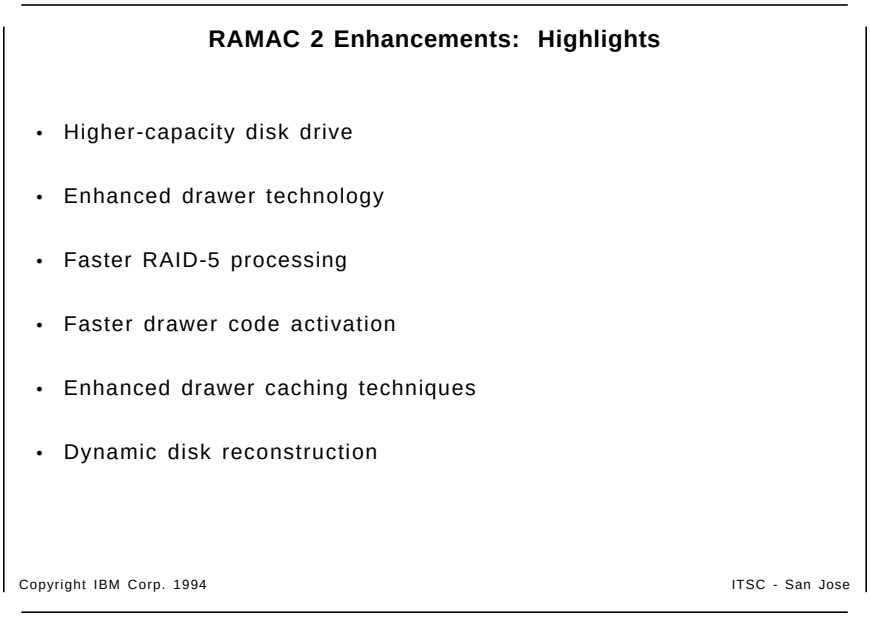

OVER-8 RAMAC 2 Enhancements: Highlights

This section discusses the highlights of the new RAMAC 2 drawers.

- **Higher-capacity disk drive**

  The capacity of the disk drive is doubled to 4 GB through the use of zoned sectors and improvements in areal density. Zoned sectors enable more efficient use of the disk area for data storage.

- **Enhanced drawer technology**

  Some components of the drawer are enhanced for better performance. A faster SCSI path manager provides faster data transfer between the disk drives and the drawer cache. The faster SCSI manager is required to support the faster disk in the RAMAC 2 drawer.

- **Faster RAID-5 processing**

  The XOR process to calculate the RAID-5 parity is improved by a new chip, which operates in a fast burst mode when calculating parity.

- **Faster drawer code activation**

  Drawer microcode activation has been improved. If the drawer is already active, it is no longer necessary to run all of the diagnostics normally required at initial power on. In addition, the microprocessors in the drawer can be initialized with their new code directly from drawer cache instead of waiting for the code to be written to the disk drives first. Both of these enhancements reduce the time taken to download and activate new drawer microcode.

- **Enhanced drawer caching techniques**

  RAMAC 2 drawers implement enhanced caching algorithms that result in more effective use of drawer resources. Disk drive usage is reduced during destage operations by using techniques that avoid multiple arm movements during random data transfer from drawer cache to disk.

- **Dynamic disk reconstruction**

  With dynamic disk reconstruction, you can take out a defective disk drive and replace it with a new disk drive as application I/O continues to the data in the drawer. The RAMAC RAID-5 architecture ensures that data is protected during the removal and repair of the disk drive. Data is reconstructed automatically on the newly replaced disk.

Enhanced High-Capacity Disk Drive

Copyright IBM Corp. 1994 ITSC - San Jose

Areal Density

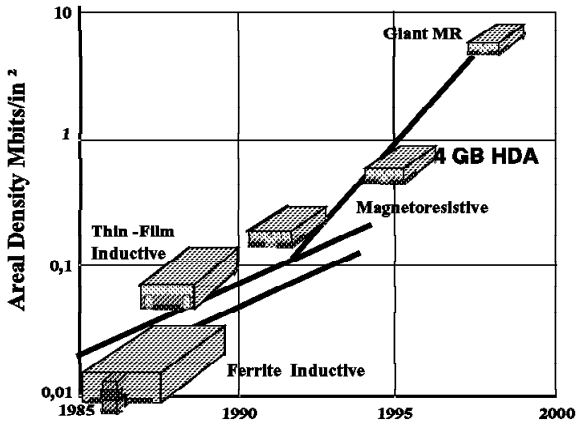

Copyright IBM Corp. 1994 ITSC - San Jose

HDA-10 Areal Density

Foil HDA-10 shows the evolution of disk recording technology. IBM is the leader in disk areal density and has developed several disk products with areal densities of up to 865 Mbit/sq in. Areal density research and development is vital to the disk manufacturing industry because it results in smaller form-factor disk drives, which are cheaper to manufacture, package, and operate. Table 1 on page 10 lists some of the disk drives available to OEMs.

Table 1. IBM Disk Drives and Areal Densities			
Product	**Capacity (MB)**	**Areal Density (Mbit/sq in.)**	**Disk Platters**
DFHS-31080	1126	544	2
DFHS-34320	4512	544	8
DFMS-32600	2657	578	4
DFMS-34320	4320	578	7
DFMS-35250	5318	578	8
DCHS-38700	8700	790	9
DCMS-310800	10800	865	10

The disk drives in the RAMAC 2 drawer have an areal density of 544 Mbit/sq in. and a capacity of 4512 MB. The disk drives have 8 platters and use all of the 16 heads and disk surfaces. This usage is possible because of the embedded servo track technique, which frees up one head and one disk surface for user capacity.

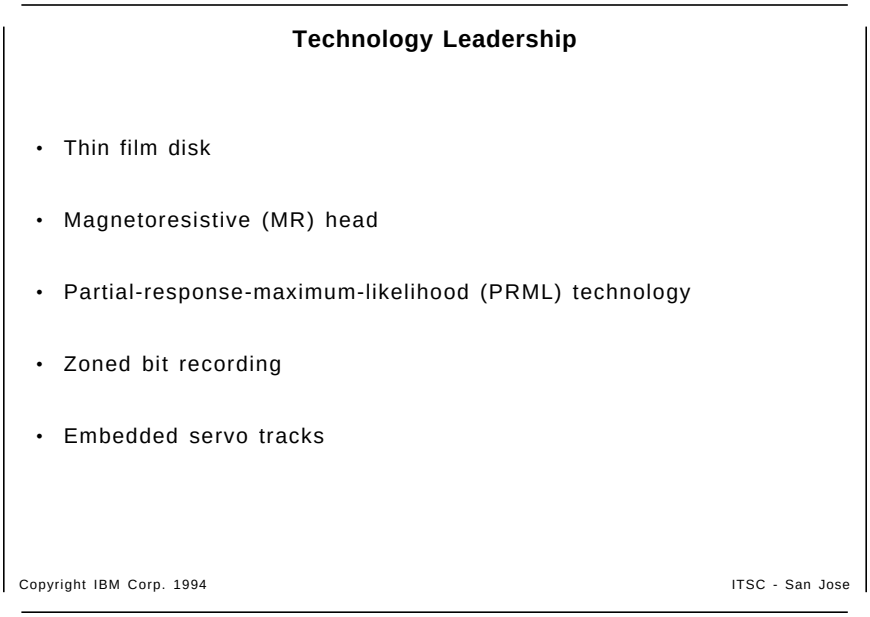

HDA-11 Technology Leadership

High capacity 3.5 in. form-factor disks must have a stable and proven technology base to be used in the RAMAC 2 Array family of products. Foil HDA-11shows the major technologies that are used.

- **Thin-film disk**

  When the bits are packed closer together, the signals from the disk get weaker. Therefore the head must be as close to the surface as possible, and the surface must be very smooth. Thin-film technology provides the necessary smooth disk surface.

- **Magnetoresistive (MR) head**

  The MR head technology is the basis for high-capacity disks. A small sensor that utilizes the MR effect detects very weak magnetic fields. The high areal densities attainable with MR heads have enabled IBM to produce disk drives that permit storage capacity with a minimum number of heads and disk components.

  IBM produced the first commercial MR-head-based disk drives in 1991 and now uses MR head technology across all of its disk drive product lines.

- **Partial-response-maximum-likelihood (PRML) technology**

  As the bit density increases, the signals received by the disk drive channel become more difficult to interpret. The determination of whether a signal reflects a binary one or zero is more complex because the bits are packed closer together on the disk surface. The conventional peak detection channel to read the data is not reliable enough in such a situation. IBM uses the PRML channel, which extracts data from the rotating disk and supplies the data to the interface. The PRML channel samples the waveform generated by the signal from the disk and calculates the most likely bit pattern. The PRML channel is a high-speed digital channel whose accuracy avoids a lot of ECC processing that would otherwise be necessary because of errors in translating waveforms.

- **Zoned bit recording**

  The disks in the RAMAC drawer have the same amount of data stored on the inner and outer areas of the disk surface. Because an outer rim track is longer than an inner rim track, the areal density in an outer rim is lower than it is in an inner rim. The areal density is lower in an inner track because, during rotation, the outer rim travels faster than the inner rim. If the data were packed as close on the outer regions of the disk surface as they are on the inner regions of the disk surface, the read/write channel would not be able to support the data rate from the device. This method of storing data on the disk wastes space because the outer tracks contain “gaps” designed to maintain uniform capacity across the disk and thus a uniform media transfer rate. The disks in a RAMAC 2 drawer are divided into zones with about the same areal density. More data fits in an outer zone than in a zone near the center. This method of storage affects the transfer rate, which is higher in an outer zone for a disk in a RAMAC 2 drawer.

- **Embedded servo track**

  Conventional disk drive designs use a single, dedicated servo surface for maintaining read/write head track positioning. This design requires performance-robbing thermal calibration routines to position the heads and a dedicated surface and read/write head. Therefore in traditional disk drive designs, although 16 heads and surfaces are available, only 15 are used for data.

  The disk drives in RAMAC 2 drawers eliminate this restriction by distributing the servo function into each sector to ensure linear, high-speed data transfers. Thus a disk surface and head are saved.

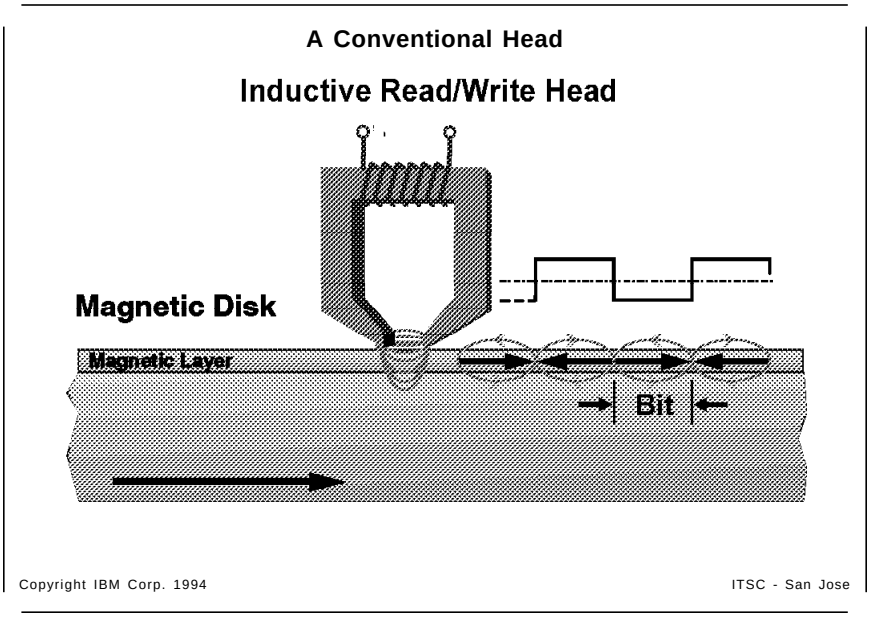

HDA-12 A Conventional Head

Most vendors still use the old inductive head technology where the same head is used for reading and writing data. However, reading and writing have contrary physical requirements. Optimal writing requires a head with only a few thick wire coils, a high current, and slow rotation of the disk. Optimal reading, however, requires a head with many thin wire coils and a high rotation speed to produce a high induction signal. The conventional head is a compromise between these requirements.

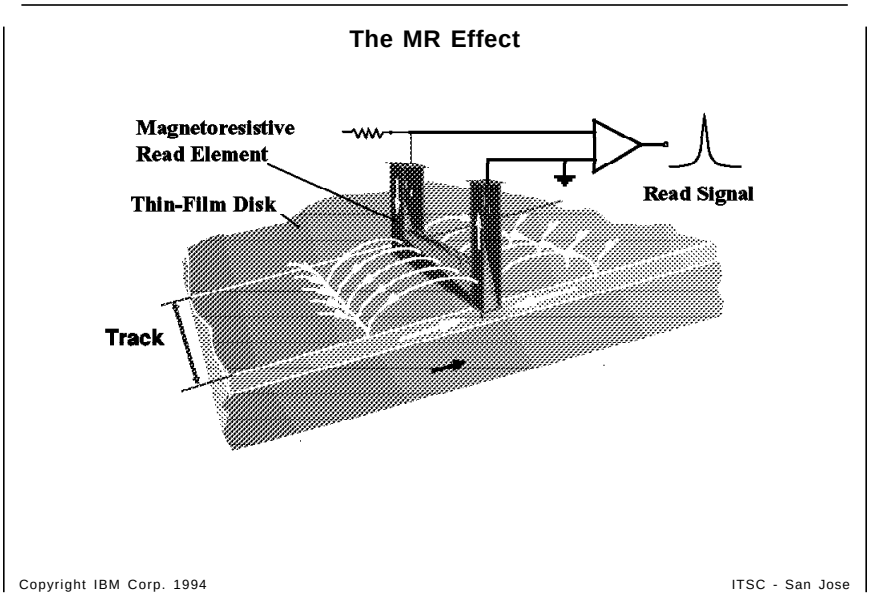

HDA-13 The MR Effect

IBM was the first manufacturer to use read heads that exploit the MR effect in its products. The read head uses an NiFe alloy film that changes resistance in the presence of a magnetic transition, or flux reversals that represent data—the magnetoresistive effect. A current sent through the read element changes as the magnetic field changes. This change of the current represents the data signal. With the MR effect, very small changes of a magnetic field can be detected, thus allowing dense packaging of the data.

IBM scientists have developed films that exhibit the *giant MR* (GMR) effect. GMR structures are even more sensitive than MR structures, and one can expect areal densities greater than 1 Gbit/sq in. in the future.

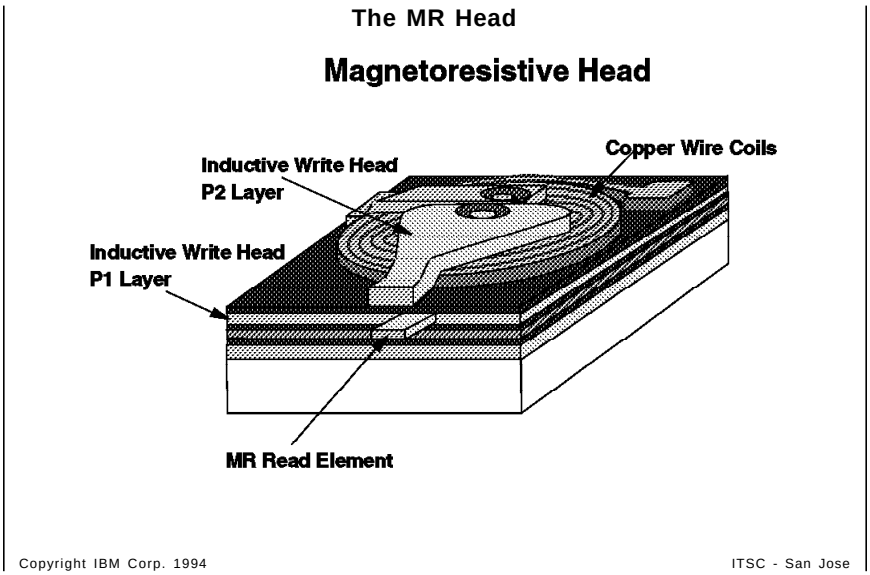

HDA-14 The MR Head

Foil HDA-14 presents a view of the MR head.

The foil shows how the read and the write elements are combined to form the MR head. The write element is a thin-film inductive head but now optimized for writing. The MR read element is integrated in the write head, thus giving rise to the term *merged MR head*.

One can see the wire coils and the two poles of the write magnet. Note how small the read element is compared to the write element. Shielding layers protect the MR element from other magnetic fields. The second shield (P2) also functions as one pole of the inductive head. The optimized design of the MR head results in a component that is easier to build than an inductive head that has to perform both read and write operations. The MR head requires fewer copper coils, material layers, and head-tolerance controls, resulting in higher process yields and reduced manufacturing and end-user costs.

The higher areal densities are achieved through increased linear density (bits per inch along a track) and increased track density (number of tracks per inch).

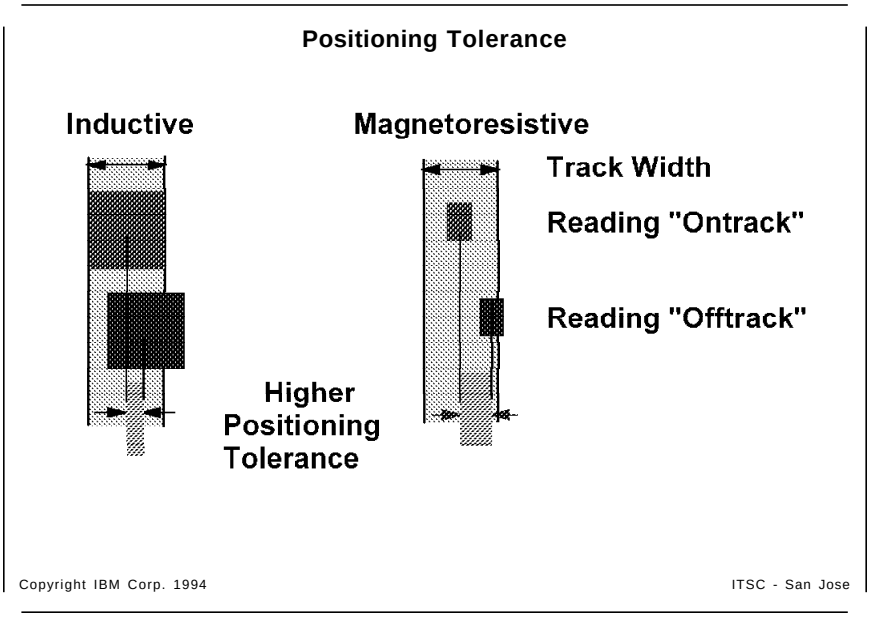

HDA-15 Positioning Tolerance

The write element writes the data in a relatively wide track, whereas the read element is very small. This arrangement allows for some flexibility to position the arm. The read element is narrower than the write element to improve off-track performance. As a rotary actuator moves across the disk surface, the heads are skewed with respect to the tracks, causing an offset between the centerline of the write head and the read head. Optimum performance is achieved when both heads are aligned with the center of the track, but this is not possible because of the head skew. The write head is wider than the read head, so it has less positioning tolerance, and center-of-track alignment “priority” is given to the write head. Thus the narrower read head is partially offtrack during write operations. The dimensions of the MR read head are such that the offtrack operation does not affect the reliability of the drive because repositioning of the head, if required, can execute faster than the inductive head.

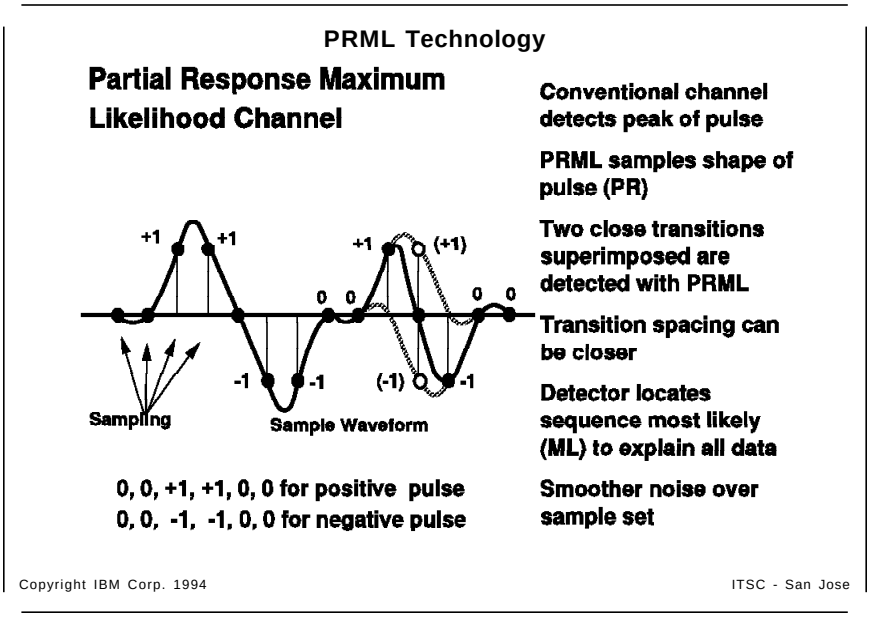

HDA-16 PRML Technology

The conventional read head uses a peak detection technology to analyze the waveforms that represent data coming off the disk. However, as the bit density increases, the waveforms interfere because of their close proximity to each other and can sometimes appear as a single waveform, thus resulting in data loss. Sometimes the waveforms cancel each other out because the data is so densely packed, and thus it is not possible to distinguish the peak of the waveform. Waveform peak detection technology has therefore reached its practical limits during this decade as the need for higher areal densities pushes for major advances in disk drive channel technology. The PRML channel used by the disk drives in RAMAC and RAMAC 2 drawers supports the highest areal densities available today.

The PRML channel samples the waveforms that represent the data as they come off the disk. They use the sampled data points of the waveforms rather than the waveforms themselves to construct the data stream. The result is a dramatic increase in the ability of a drive to correctly resolve the read pulses representing data stored at very high linear densities. IBM enhances the PRML with a unique digital filter, which produces lower soft error rates and enables higher data throughput.

Improving the areal density of the disk drive will be to no avail if the data cannot be read and recorded accurately. The PRML channel makes the high areal densities available today both practical and cost effective while sustaining the highest levels of reliability and performance.

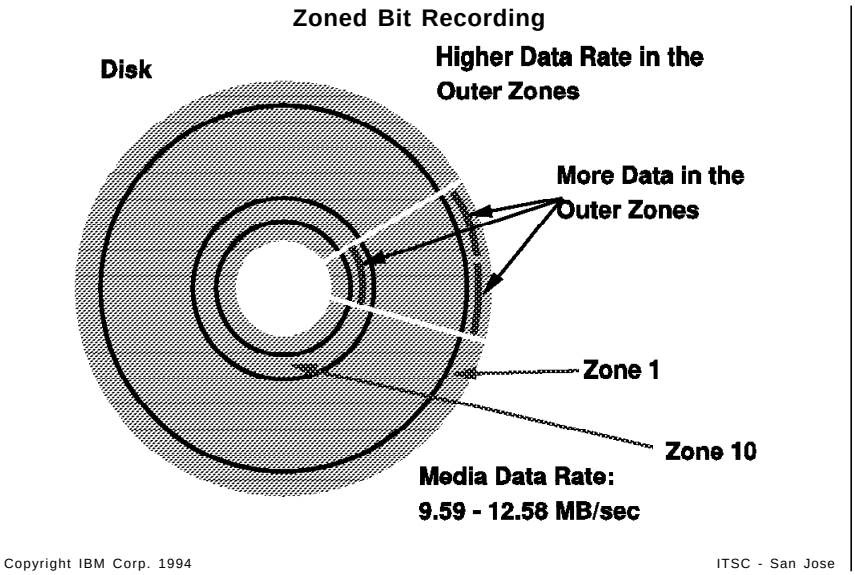

HDA-17 Zoned Bit Recording

The disks in the RAMAC 2 drawer are partitioned into recording bands or zones. There are 10 zones on each disk. Zone 1 is located nearest the outer diameter of the disk, and zone 10 is located nearest the inner diameter of the disk.

Previously the areal density was not the same in the inner and the outer diameter regions of the disk, but in the RAMAC disk the zones have approximately the same areal density. Thus more data can be stored in an outer zone than an inner zone, and the disk surface is better utilized for storing data.

Because the radial speed of the outer zones is about twice as high as the radial speed of the inner zones, the transfer rate is higher in the outer zones. The media transfer rate varies between 9.59 MB/sec (zone 10) and 12.58 MB/sec (zone 1), with an average of 12.07 MB/sec.

Approximately 43% of all data stored on the disk is stored in zone 1.

Device Characteristics

Configuration

Capacity	4512 MB
Form factor	3.5 in. (95 mm)
Number of disks	8
Number of heads	16
Areal density (max.)	544 Mbits/sq in.
Interface	Fast/wide SCSI-2
Bytes per sector	256 - 744
Disk type	Thin film
Channel	PRML
Actuator type	Rotary VCM
Head type	Magnetoresistive
Size	Height 41.3 mm, width 101.6 mm, depth 146.0 mm
Weight	0.82 kg

 ITSC - San Jose

HDA-18 Device Characteristics

Foil HDA-18 shows the characteristics of the disk drives used in RAMAC 2 drawers. Compared with the disk drives used in RAMAC drawers, the capacity is doubled.

The disk drives used in RAMAC 2 drawers are 3.5 in. form factor with a capacity of 4.51 GB. Because RAMAC is a RAID-5 device, 25% of the disk capacity is used for parity information. When the space required for storing microcode and subsystem status is subtracted, the usable capacity of the four drives in the RAMAC 2 drawer is 11.35 GB. The disk drives transfer data over the SCSI interface in 16-bit mode. (This is referred to as wide mode; 8 bit transfers are referred to as narrow mode.) The sector size used for RAMAC 2 disks is 688 bytes—the same as for RAMAC disks.

The size and weight of the disk drives in RAMAC 2 drawers are the same as the size and weight of the disk drives in RAMAC drawers.

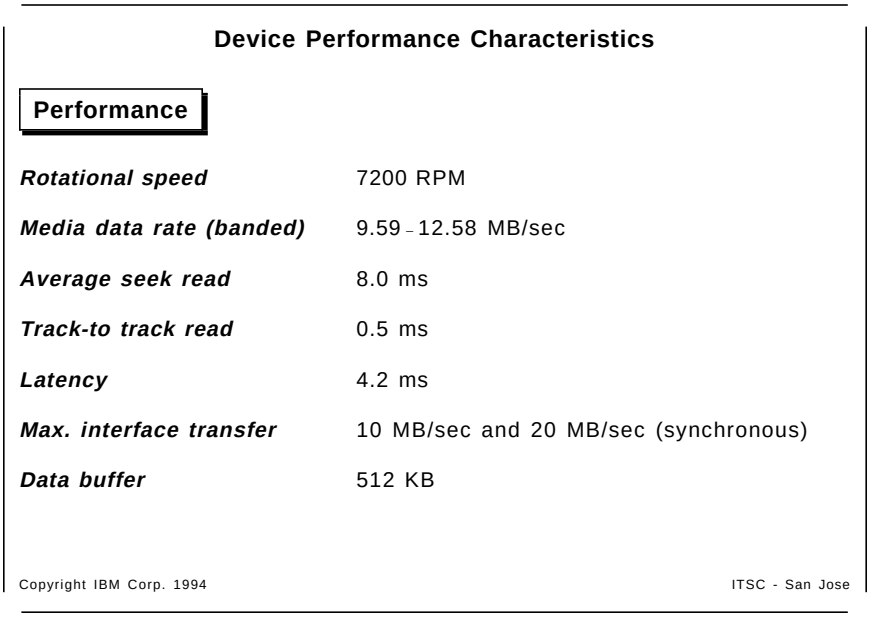

HDA-19 Device Performance Characteristics

Foil HDA-19 shows the performance characteristics of the disk drives used in the RAMAC 2 drawers. The rotational speed is increased by 33% from 5400 RPM to 7200 RPM. Thus the latency time is reduced from 5.6 ms to 4.2 ms. The media data transfer rate depends on the location of the data. It ranges from 9.59 MB/sec (zone 10) to 12.58 MB/sec (zone 1) This is about twice the 5.2 MB/sec transfer rate of a disk drive in a RAMAC drawer.

The disk drives in the RAMAC 2 drawer are configured for a 16 bit wide synchronous data transfer rate of 20 MB/sec. This is the time required to transfer data between the SCSI bus and the disk's data buffer. The disk drive buffer is 512 KB and consists of eight 64K segments.

These improvements together with other drawer enhancements ensure that the higher capacity drives provide a level of performance that supports the four-volume RAMAC 2 drawer.

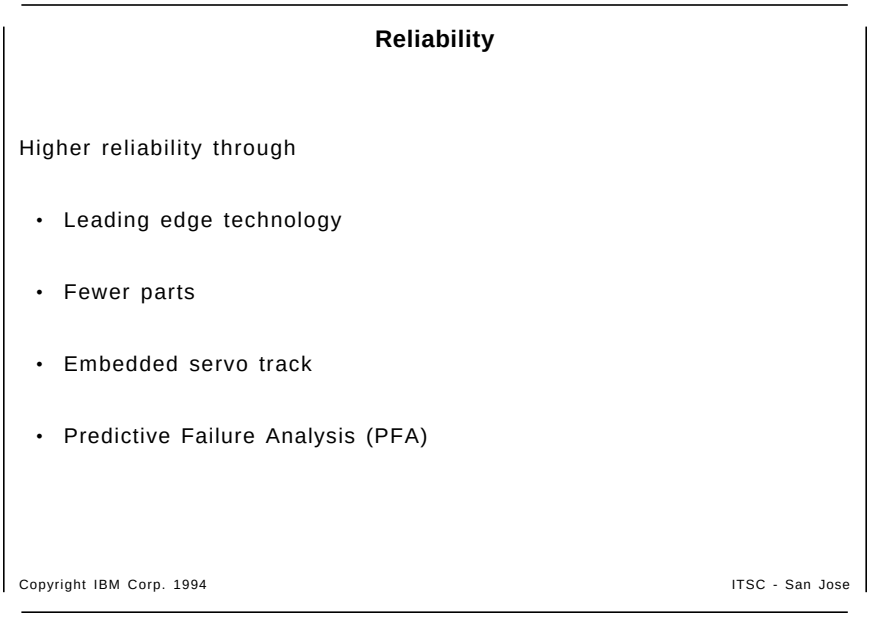

HDA-20 Reliability

Foil HDA-20 shows the building blocks for the high reliability of the disk drives in the RAMAC 2 drawers.

- **Leading edge technology**

  The use of such state of the art technology as MR heads and PRML channels enables steady reading and writing of densely packed data with very low error rates.

- **Fewer parts**

  IBM's technology leadership results in more efficient component packaging which in turn results in fewer electronic circuits and parts. With fewer parts you have fewer failures.

- **Embedded servo track**

  The use of embedded servo tracks keeps the MR head on track and renders complicated thermal calibration routines obsolete.

- **Predictive Failure Analysis (PFA)**

  IBM applies Predictive Failure Analysis (PFA) features in the disk drives in the RAMAC 2 drawers. PFA routines monitor the drives by reviewing error logs and performing a suite of automatic self-diagnostic tests that measure changes in heads and disks and other critical components, such as the read channel. When specific thresholds are exceeded, PFA can alert the user to the need for drive replacement, before a drive fails.

Enhanced RAMAC 2 Array Drawer

 ITSC - San Jose

Device Emulation

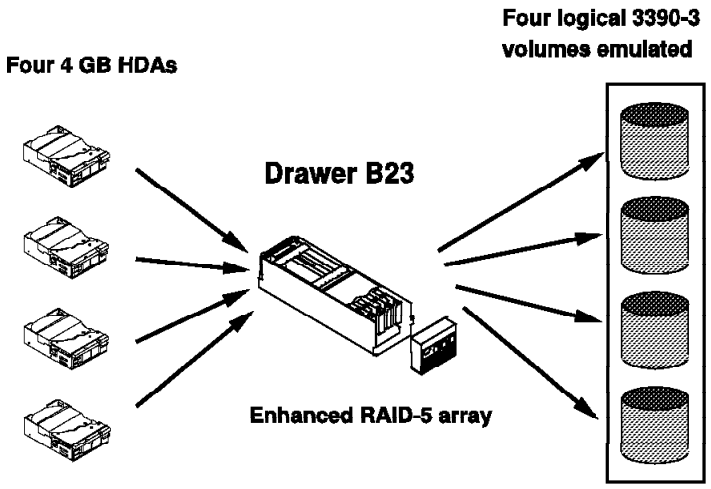

 ITSC - San Jose

Drawer-22 Device Emulation

A RAMAC B13 drawer emulates two logical 3390 model 3 volumes or three 3380-K volumes with a drawer capacity of 5.675 GB. A RAMAC 2 B23 drawer emulates four logical 3390 model 3 volumes with a drawer capacity of 11.35 GB.

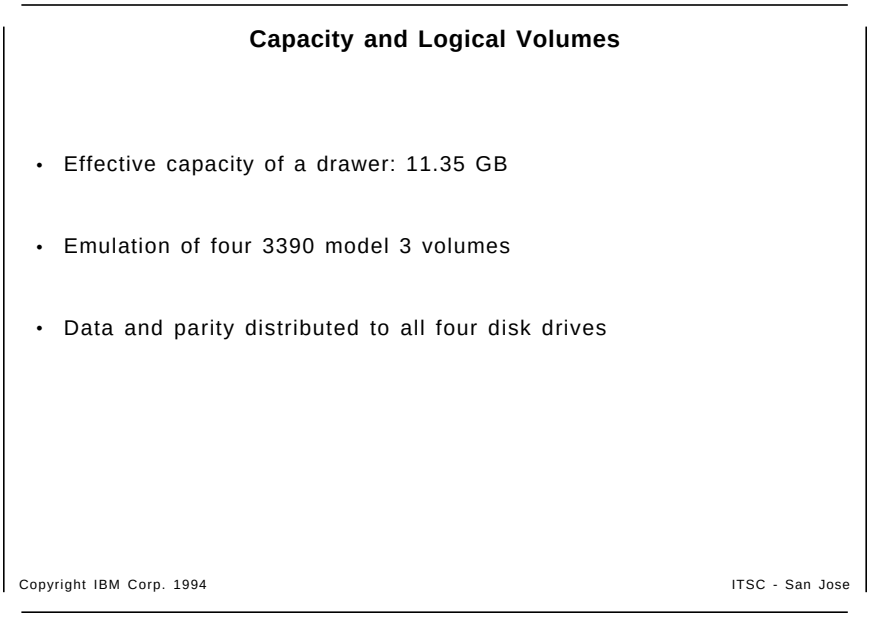

Drawer-23 Capacity and Logical Volumes

The use of 4 GB disk drives in the RAMAC 2 drawers doubles the capacity from 5.675 GB to 11.35 GB. Therefore instead of emulating two logical volumes, the RAMAC 2 drawer emulates four logical volumes. CKD emulation and the physical-logical mapping relationship implemented in RAMAC are the same for RAMAC 2.

- **Effective capacity of a drawer: 11.35 GB**

  Each disk drive has a capacity of 4 GB. With four disk drives in a drawer the total capacity of a RAMAC 2 drawer is 16 GB. About one-quarter of the drawer capacity is needed to store the RAID-5 parity data. The formatted capacity available for storing data is 11.35 GB after subtracting space reserved on the drives for microcode storage and subsystem status.

- **Emulation of four 3390 model 3 volumes**

  The data storage is used to emulate four logical 3390 model 3 volumes with a capacity of 2.83 GB each. The RAMAC Array DASD emulates the 3390 model 3. The RAMAC Array Subsystem with RAMAC drawers allows emulation of 3380-K and 3390-3 DASD. For RAMAC 2 drawers in a RAMAC Array Subsystem only 3390-3 emulation is available.

- **Data and parity distributed to all four disk drives**

  Data and parity are distributed across the four disk drives for all of the logical volumes in the drawer. The logical track mapping is the same as that used for RAMAC drawers.

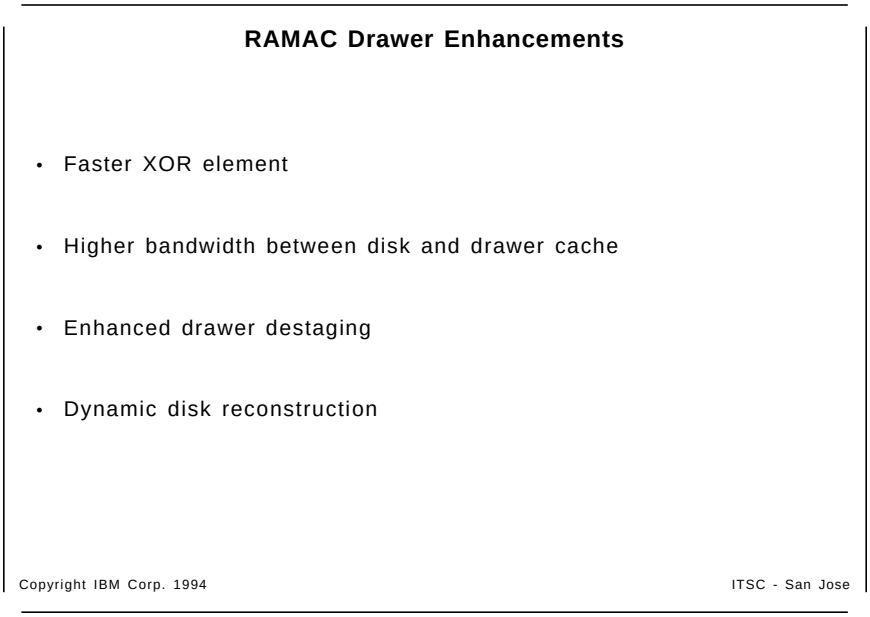

Drawer-24 RAMAC Drawer Enhancements

Foil Drawer-24 lists the RAMAC 2 drawer enhancements.

- **Faster XOR element**

  In a RAID-5 implementation there is a lot of parity processing. For the high data traffic expected for the RAMAC 2 high-capacity drives, the cycle time of the XOR hardware and hence the performance of the parity processing element are increased.

- **Higher bandwidth between disk and drawer cache**

  The SCSI rate has been increased to support the faster media transfer rate of RAMAC 2 drives. Because of the faster transfer capability of the disk drives used in the RAMAC 2 drawer, the drawer′s port-SCSI manager (PSM) has been enhanced to handle faster disk-to-drawer cache and drawer cache-to-disk data transfers.

- **Enhanced drawer destaging**

  When writes are performed to more than one logical volume in a drawer, data ultimately has to be destaged to the disk drives. Because the four logical volumes are spread over all of the disk drives, there will be multiple destages to the same physical disks, which results in the disk head moving rapidly back and forth to satisfy the I/Os.

  The RAMAC 2 drawer microcode groups destages by disk drive location. This grouping of destages by disk device reduces the ping-pong effect of the head moving back and forth and increases overall performance by optimizing disk usage.

- **Dynamic disk reconstruction**

  With dynamic disk reconstruction you can take out a defective disk drive and replace it with a new disk drive, and the data is reconstructed on the new disk drive automatically. During the repair RAMAC′s RAID-5 architecture regenerates the data to satisfy application requests.

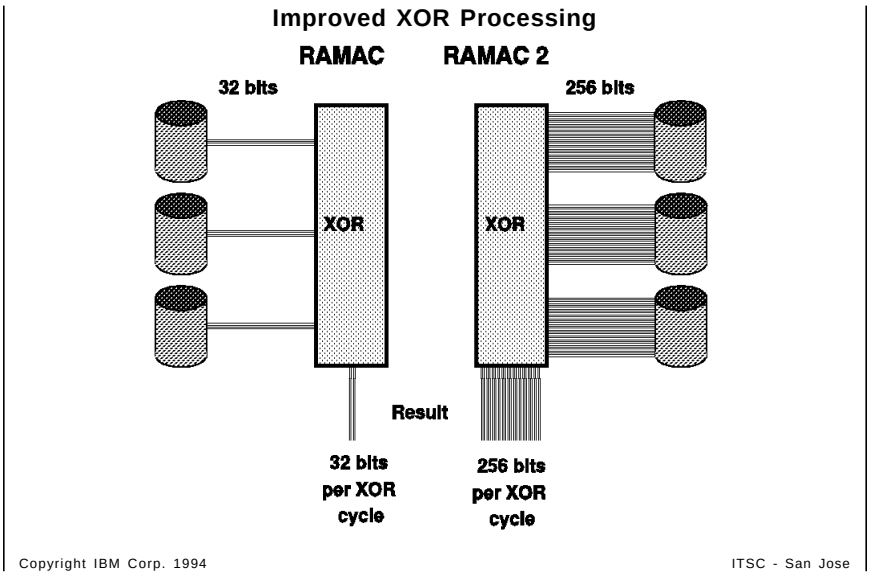

Drawer-25 Improved XOR Processing

With RAID-5 implementations, parity information must be stored on disk. The parity calculation is performed frequently, and hence the time taken to produce a parity result is crucial for a RAID-5 implementation. To cope with the high data transfer rate expected for a high-capacity disk, the parity execution unit in a RAMAC 2 drawer is enhanced.

Parity is calculated by an XOR operation of corresponding data blocks (688 bytes) from three disk drives. Previously 32 bits were processed in each parity processing cycle in the RAMAC drawer. Now 256 bits are processed in each cycle by the XOR element in the RAMAC 2 drawer. This processing results in fewer cycles to deliver a parity calculation for the array and improves the performance of the RAMAC 2 drawer. The RAMAC 2 drawer has to cope with higher I/O rates than the RAMAC drawer because its double-capacity disk drives emulate four logical volumes compared with two logical volumes of the RAMAC drawer.

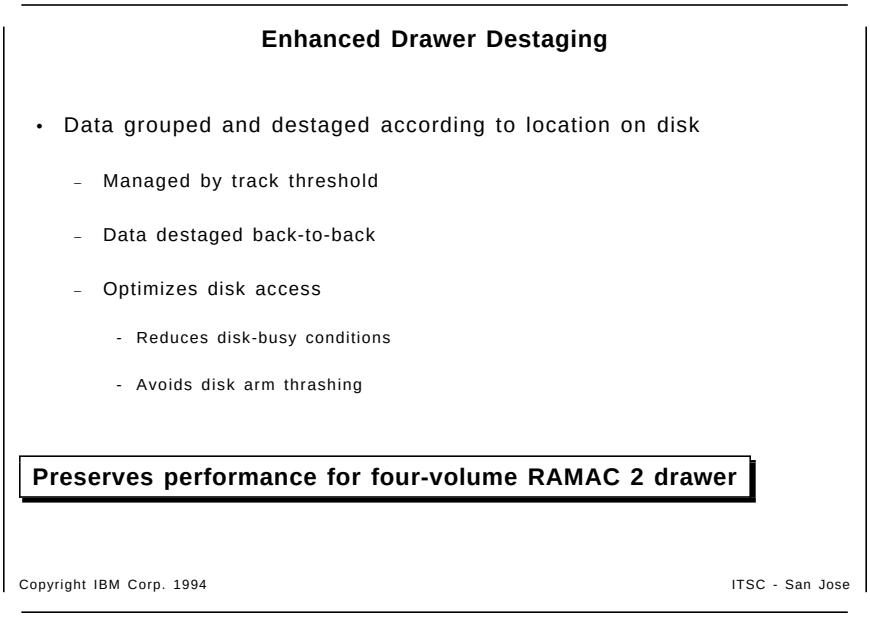

Drawer-26 Enhanced Drawer Destaging

The second generation of IBM′s RAMAC Array Family is here. The RAMAC 2 Array DASD and the RAMAC 2 Array Subsystem bring high-capacity, 4 GB disk drives to an enhanced RAID-5 storage device. Increasing the capacity of the disk drives to reduce the cost of storage is not enough – conventional storage subsystem techniques must be revisited if high-capacity RAID-5 devices are to be feasible. RAMAC 2 implements a change to traditional cache management methods, as we discuss below.

- **Data grouped and destaged according to location on disk**

  The enhanced RAMAC 2 drawer destaging is designed to optimize drawer resources and provide improved performance for the high-capacity RAMAC 2 array. The enhanced destaging uses an awareness of the “home” location on disk of data eligible for destage. The RAMAC 2 drawer arranges the data such that multiple blocks are written out to the same disk when a certain threshold is reached. Before initiating the destage, the RAMAC 2 drawer ensures that the data is organized such that the data is written back-to-back in a single “sweep” of the actuator arm in one direction. This method avoids having to reposition the disk arm for subsequent destaging.

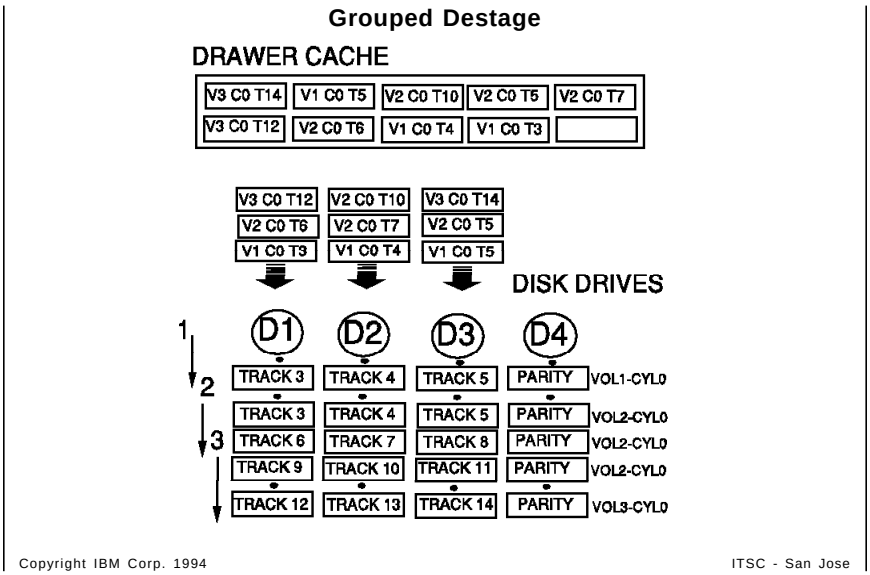

Drawer-27 Grouped Destage

The RAMAC 2 drawer destaging enhancements overcome the problem of inefficient disk use caused by random destaging. Rather than initiate a destage from drawer cache to the disk drives on the basis of frequency of use of the data, the RAMAC 2 drawer sets a track threshold for each disk. When the threshold is reached it groups the data by disk drive location and writes the data out in such a way as to avoid repeated arm repositioning.

The foil depicts the improved technique. Notice that the three tracks being destaged to each of the disk drives have been arranged to optimize disk arm movement. We see the data being sorted by logical volume to eliminate arm movement or optimize it when it is unavoidable. Tracks 7 and 10 are destaged back-to-back.

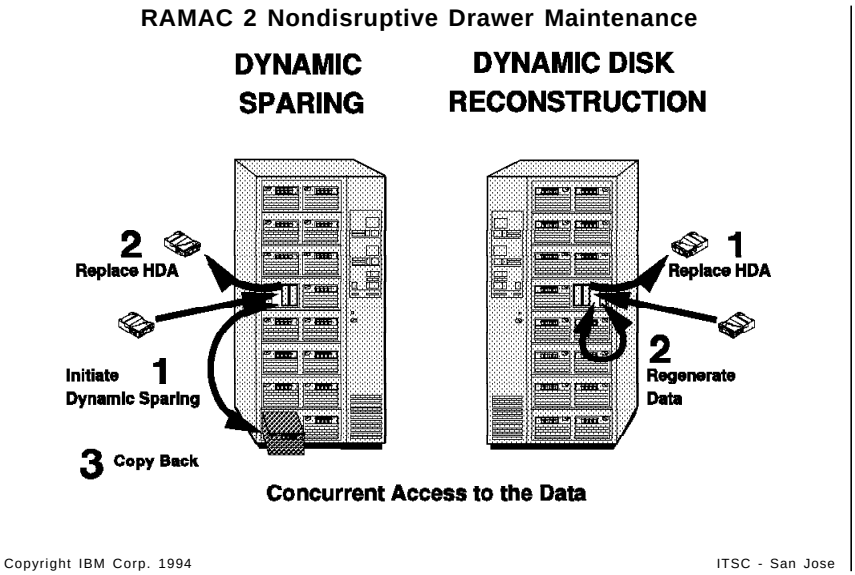

Drawer-28 RAMAC 2 Nondisruptive Drawer Maintenance

The RAMAC Array Family is designed to provide nondisruptive drawer maintenance. Thus customers′ applications can continue to access data when drawer repairs are taking place. This function is delivered by Dynamic Sparing copying the data from the failing drawer to a spare drawer. After the copy process is complete, the CE can repair the failed drawer component and optionally copy the data back from the spare drawer. Dynamic Sparing is the preferred nondisruptive drawer maintenance support function for all drawer component repairs because it produces a second copy of the data on a fully functional RAID-5 device.

Dynamic disk reconstruction offers an alternative nondisruptive drawer maintenance procedure for RAMAC 2 drawers for disk drive repair only. The data of the replaced disk can be reconstructed from the data and the parity of the other disks within the drawer itself. This function uses the hot-pluggable feature of the RAMAC 2 disk drives that allows a disk drive to be removed from a drawer independent of the other three disk drives. The RAID-5 capabilities of the RAMAC 2 drawer allow continuous access to the three-drive drawer by regenerating the data for the disk drive that has been removed. When a new drive is hot-plugged back into the drawer, the reconstruction process begins. The RAMAC 2 drawer reconstructs the data on the new disk and thus restores the full RAID-5 capability. Note that data regeneration is different from reconstruction. Both RAMAC and RAMAC 2 drawers perform dynamic data regeneration, but only the RAMAC 2 drawer performs dynamic reconstruction.

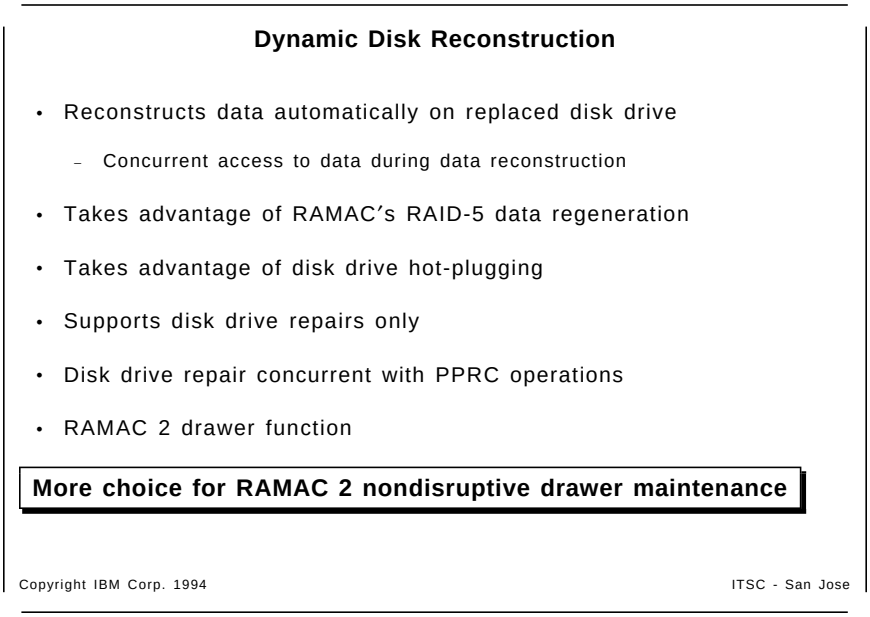

Drawer-29 Dynamic Disk Reconstruction

- **Reconstructs data automatically on replaced disk drive**

  When the CE selects the dynamic disk reconstruction option, he or she removes, repairs or replaces, and reinstalls the disk drive. The contents of the replaced disk drive are automatically reconstructed concurrently with normal drawer operations. During data reconstruction, host I/Os can still access data in the drawer. Starting from cylinder 0, data reconstruction progresses toward the end of the disk. Depending on to which track or cylinder the I/O is directed, normal RAID-5 processing or RAID-5 data regeneration takes place.

- **Takes advantage of RAMAC′s RAID-5 data regeneration**

  If a disk drive malfunctions, the RAID-5 function of the RAMAC 2 drawers regenerates the data for the failed drive such that application access to the data continues. If Dynamic Sparing is configured for the subsystem, the drawer contents are automatically copied to the spare drawer when the error is first detected. This provides the best protection for the data because a second copy is rapidly secured on another RAID-5 device. If Dynamic Sparing is not configured for the rack, the CE can use dynamic disk reconstruction for nondisruptive drive repair.

- **Takes advantage of disk drive hot-plugging**

  Dynamic disk reconstruction takes advantage of RAMAC 2′s hot-pluggable disk drives. Hot pluggable disk drives that are part of a disk array can be removed independently of the other disk drives in the array and reinstalled at a later time. Hot-plugging is a physical characteristic of the drive itself and works with the dynamic disk reconstruction function of the RAMAC 2 drawer to offer customers another nondisruptive drawer maintenance option.

- **Supports disk drive repairs only**

  If there is a problem with one of the logic or memory cards in the drawer, it is not possible to replace the card without taking the drawer out of service. Therefore the data must be copied to a spare drawer before any repair

action. Dynamic disk reconstruction provides nondisruptive drawer maintenance for disk drive repairs only.

- **Disk drive repair concurrent with PPRC operations**

  Dynamic Sparing, dual copy, and peer-to-peer remote copy (PPRC) are mutually exclusive for a volume or drawer. The operation that comes first gets the required resources. A PPRC pair or a dual copy pair must be terminated before Dynamic Sparing can initiate. With dynamic disk reconstruction, termination of the pairs is no longer necessary for disk drive failures.

- **RAMAC 2 drawer function**

  Dynamic Sparing is a cooperative function that involves the RAMAC or RAMAC 2 drawer and the control unit function because data is copied across drawer boundaries. Dynamic disk reconstruction is a RAMAC 2 drawer function only—the control unit has no involvement in the reconstruction process, which is confined to a RAMAC 2 drawer.

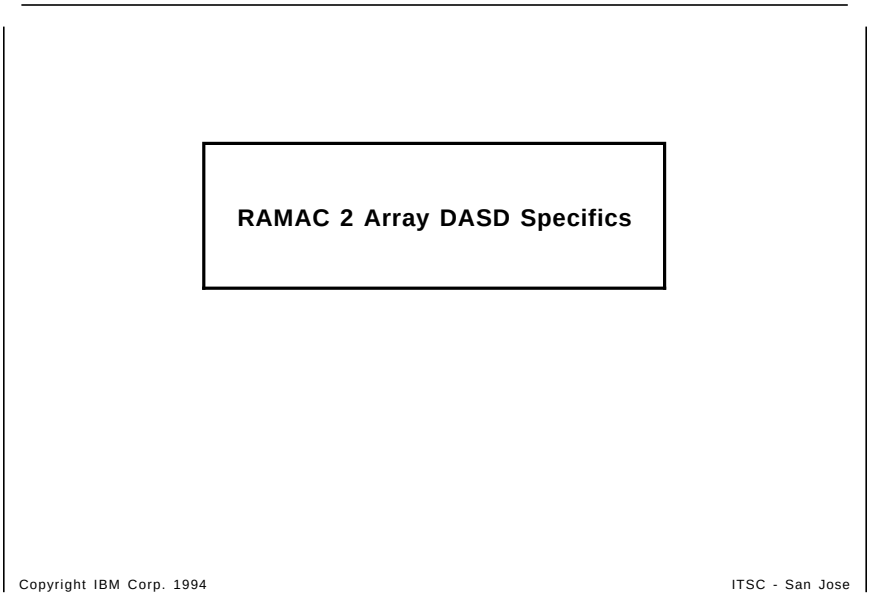

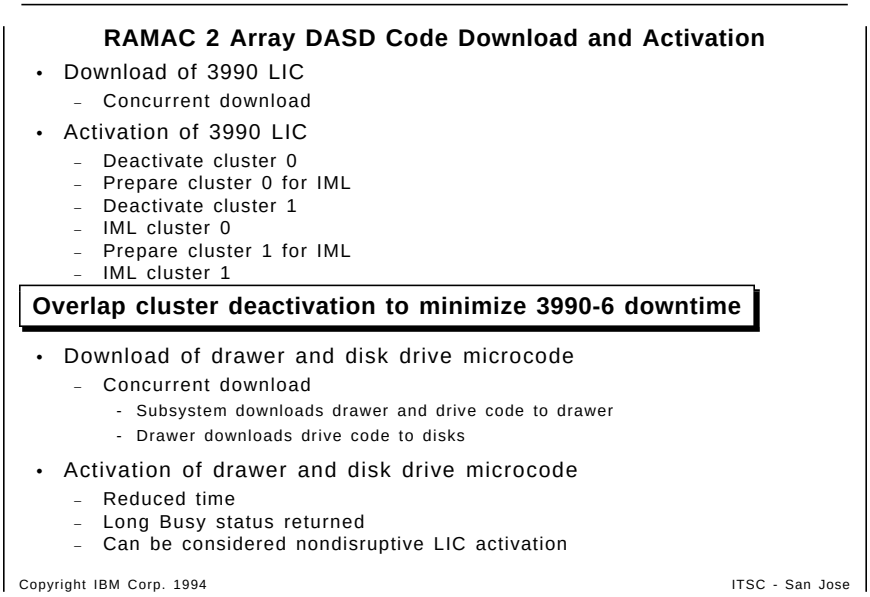

ArrayDAS-31 RAMAC 2 Array DASD Code Download and Activation

You must distinguish between subsystem, drawer, and drive microcode. You must further distinguish between code download and code activation.

- **Download of 3990 LIC**

  The download of LIC to the 3990-6 storage control is nondisruptive because it takes place as the 3990-6 processes application I/Os. The 3990-6 has two hard disks that contain different levels of LIC.

- **Activation of 3990 LIC**

  To activate the 3990 LIC you deactivate one cluster, after taking the paths to that cluster offline.

  While preparing the first cluster for LIC activation, the second cluster is servicing I/Os. Before IMLing the first cluster, the second cluster must be

deactivated. Thus neither cluster is available for a period of time. The first cluster is IMLed with the new RAMAC 2 supporting LIC, followed by the second cluster. This staggering of cluster deactivation and activation minimizes 3990-6 downtime.

The deactivation of both clusters occurs only when the RAMAC 2 supporting LIC is installed for the first time. Subsequent RAMAC 2 drawer or rack upgrades are nondisruptive.

- **Download of drawer and drive microcode**

  The download of code to the RAMAC drawers is nondisruptive because the 3990 and RAMAC drawers can interleave I/O processing with the transfer and storage of the code in drawer cache and on the four disk drives.

- **Activation of drawer and disk drive microcode**

  Drawer code activation used to take about 1 minute. However, if the drawer is already active, it is not necessary to run all of the diagnostics required for an initial power on. Previously, the code was first written from the cache to the disk drives and then loaded into the drawer processors. In RAMAC 2 drawers the code is first activated for the drawer processors from cache and then written to disk. This change reduces the drawer and drive microcode activation time. During this time a Long Busy status is returned to the host. If the host (operating system and application) supports the Long Busy status, the drawer and drive LIC activation can be considered to be concurrent.

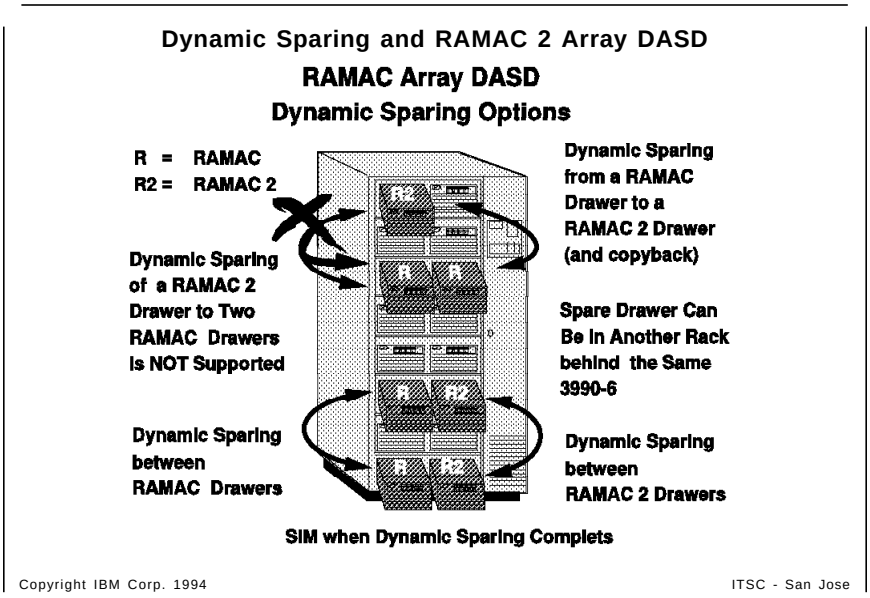

ArrayDAS-32 Dynamic Sparing and RAMAC 2 Array DASD

RAMAC 2 Array DASD configurations support the coexistence of RAMAC and RAMAC 2 drawers in the same rack. RAMAC 2 Array DASD configurations also support the attachment of two strings, one of which may contain RAMAC drawers, and the other, RAMAC 2 drawers. The ability to intermix RAMAC and RAMAC 2 drawers behind the same 3990-6 has implications for Dynamic Sparing.

Dynamic Sparing can take place between two RAMAC drawers and between two RAMAC 2 drawers. Dynamic Sparing can also occur from a RAMAC drawer to a RAMAC 2 drawer. However, Dynamic Sparing is not supported between a RAMAC 2 drawer and two RAMAC drawers. Dynamic Sparing operates at the drawer level, and the capacity of the target drawer must be the same or greater than that of the source drawer. The spare drawer can be in another string behind the same subsystem. The storage control tries to find a drawer with the same capacity or higher as the failing drawer. If it finds one, Dynamic Sparing starts. Otherwise a SIM is issued to indicate that Dynamic Sparing is not possible. In this case the CE must identify another RAMAC 2 drawer as a spare drawer or set up dual copy pairs to other volumes in the same subsystem.

For RAMAC and RAMAC 2 mixed configurations, we recommend that each subsystem have a RAMAC 2 drawer defined as a spare drawer.

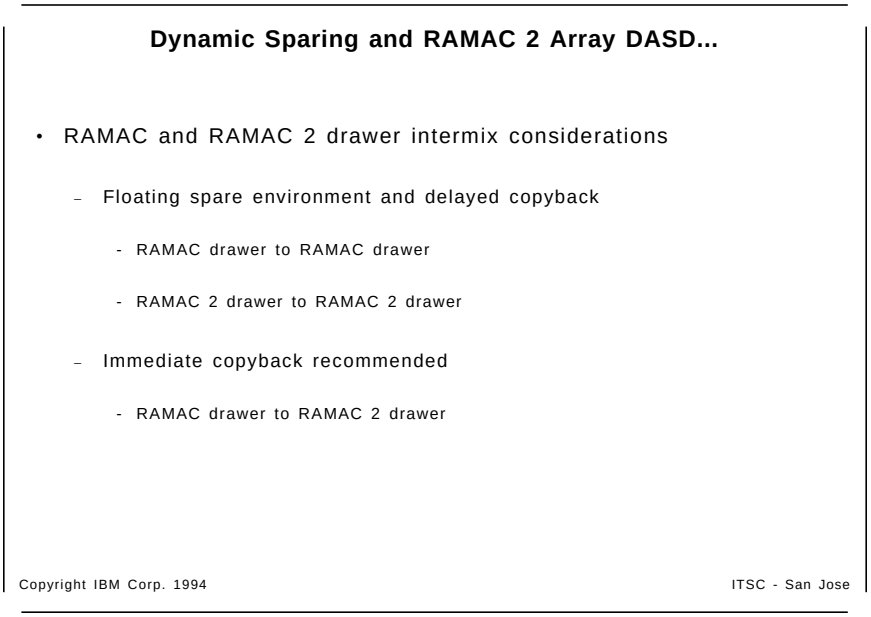

ArrayDAS-33 Dynamic Sparing and RAMAC 2 Array DASD...

The foil lists several important considerations when using Dynamic Sparing in a mixed RAMAC and RAMAC 2 drawer configuration.

- **Floating spare environment and delayed copyback**

  The 3990-6 monitors and tracks the changing DDC and CCA relationship when Dynamic Sparing is used in a floating spare configuration. Floating spare configurations are viable when RAMAC drawers are being spared to RAMAC drawers, and RAMAC 2 drawers are being spared to RAMAC 2 drawers. The delayed copyback function is also viable when “like” drawers are used in Dynamic Sparing configurations.

- **Immediate copyback recommended**

  When RAMAC drawers are spared to RAMAC 2 drawers, we strongly recommend that, after repair, the data be copied back immediately. The new spare is a two-volume drawer, and if a four-volume drawer subsequently malfunctions, Dynamic Sparing will not start. Performing an immediate copyback results in the subsystem reverting back to having a four-volume spare drawer that provides maximum protection for all data.

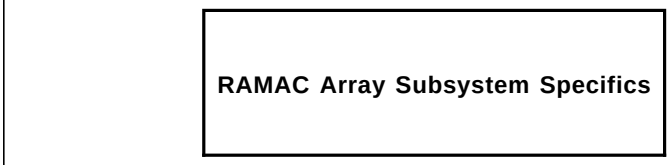

RAMAC Array Subsystem Code Download and Activation

- Download of Subsystem microcode
  - Concurrent download
  - Code stored on two IML log volumes
- Activation of Subsystem microcode
  - Deactivate one cluster controller
    - Other cluster sustains workload
  - Activate code for this cluster
  - IML overlaps for both array controllers
  - Subsystem busy for about 55 seconds
- Download of drawer and drive microcode
  - Concurrent download
- Activation of drawer and drive microcode
  - Reduced time for microcode activation
  - Long Busy status returned
  - Can be considered nondisruptive LIC activation

ArraySub-35 RAMAC Array Subsystem Code Download and Activation

- **Download of Subsystem microcode**

  The RAMAC Array Subsystem microcode is stored on reserved tracks of the disk drives in two drawers. The drives are addressed as logical volumes, and the two logical volumes containing the RAMAC Array Subsystem microcode are called the IML log volumes. For availability, the IML log volumes must be in separate drawers. The download of new code to the IML log volumes is nondisruptive.

- **Activation of Subsystem microcode**

  To activate the new code, the CE pushes the install button. This action results in the microcode automatically deactivating one cluster controller. Access to the data is still possible through the other cluster controller. The new code is then activated by an IML of one cluster controller, which

overlaps the IML of the second cluster controller. The activation process is phased, and the RAMAC Array Subsystem keeps the second cluster controller active for as long as possible during the installation and initialization phases of the first controller. There is a time period, approximately 55 seconds, when neither controller is available to service I/Os. During this time, the RAMAC Array Subsystem presents a control unit (CU) busy signal to the host.

- **Download of drawer and drive microcode**

  The download of microcode to the drawers is nondisruptive.

- **Activation of drawer and drive microcode**

  Drawer code activation used to take about 1 minute. However, if the drawer is already active, it is not necessary to run all of the diagnostics required for an initial power on. Previously, the code was first written from the cache to the disk drives and then loaded into the drawer processors. In RAMAC 2 drawers the code is first activated for the drawer processors from cache and then written to disk. This change reduces the drawer and drive code activation time. During this time a Long Busy status is returned to the host.

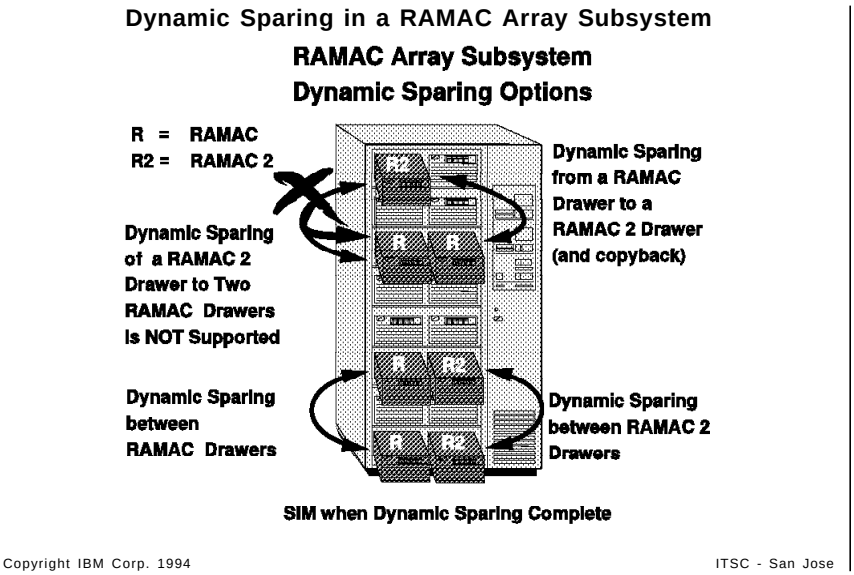

ArraySub-36 Dynamic Sparing in a RAMAC Array Subsystem

RAMAC 2 Array Subsystem configurations support the coexistence of RAMAC and RAMAC 2 drawers in the same rack. The ability to intermix RAMAC and RAMAC 2 drawers in the same RAMAC Array Subsystem has implications for Dynamic Sparing.

Dynamic Sparing can take place between two RAMAC drawers and between two RAMAC 2 drawers. Dynamic Sparing can also occur from a RAMAC drawer to a RAMAC 2 drawer.

For RAMAC and RAMAC 2 mixed configurations, we recommend that each subsystem have a RAMAC 2 drawer defined as a spare drawer.

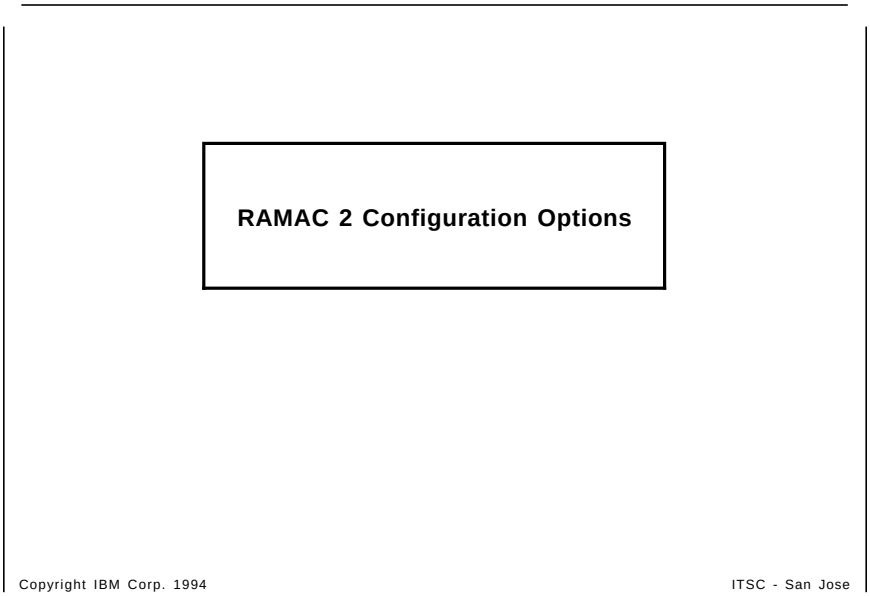

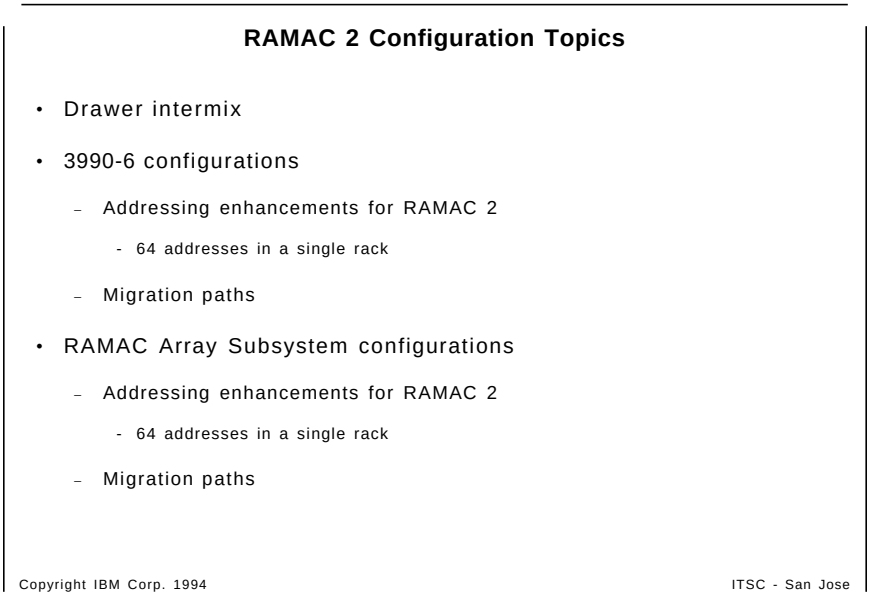

Config-38 RAMAC 2 Configuration Topics

This foil reviews the configuration and migration topics of interest for the RAMAC 2 Array DASD and RAMAC 2 Array Subsystem.

- **Drawer intermix**

  RAMAC drawers and RAMAC 2 drawers can be intermixed in a RAMAC 2 rack. This intermixing allows an easy migration to the high-capacity drawers and avoids the need for total drawer upgrade and replacement. The intermixing of RAMAC and RAMAC 2 drawers in the same rack is supported by both the RAMAC 2 Array DASD and the RAMAC 2 Array Subsystem.

- **3990-6 configurations**

  The 3990-6 storage control supports 64 addresses in a single rack containing RAMAC 2 drawers. The 3990-6 uses different address mapping schemes to

support various configurations and nondisruptive migration to a RAMAC 2 environment after the initial installation of the RAMAC 2 supporting LIC. The intermix of RAMAC and RAMAC 2 drawers is also supported through the addressing enhancements.

- **RAMAC Array Subsystem configurations**

  The RAMAC Array Subsystem supports the RAMAC 2 drawers by implementing a single address mapping scheme, which can recognize up to 64 devices in a single rack. The RAMAC Array Subsystem model 3 allows intermixing of RAMAC and RAMAC 2 drawers in either of its logical subsystems, but only allows one device emulation mode per logical subsystem. RAMAC 2 drawers only emulate 3390-3 device geometry.

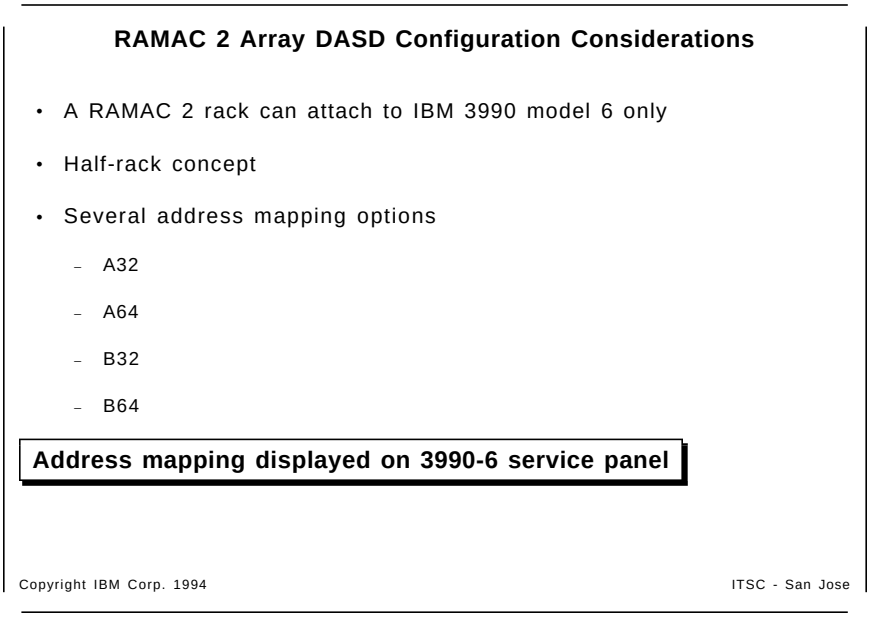

Config-39 RAMAC 2 Array DASD Configuration Considerations

The foil lists several topics for consideration when implementing RAMAC 2 configurations.

- **A RAMAC 2 rack can attach to IBM 3990 model 6 only**

  A rack with RAMAC drawers can attach to the 3990-6 or 3990-3 storage controls. A rack that contains a RAMAC 2 drawer or a mix of RAMAC and RAMAC 2 drawers can attach to an IBM 3990 model 6 only.

- **Half-rack concept**

  If there are RAMAC 2 drawers in a rack, and the number of drawers exceeds eight, the rack must be the only string attached to the 3990-6. If an existing string of 3390 or RAMAC devices is attached to the 3990-6, any second string that contains at least one RAMAC 2 drawer must not contain more than eight drawers.

- **Several address mapping options**

  The configurations discussed above arise from the address mapping schemes that the IBM 3990-6 uses. Four address mapping schemes are available (see foils Config-42 to Config-45), each of which allows for a different migration path.

  The address mapping used by the 3990-6 is displayed on the service panel as either A32, A64, B32, or B64.

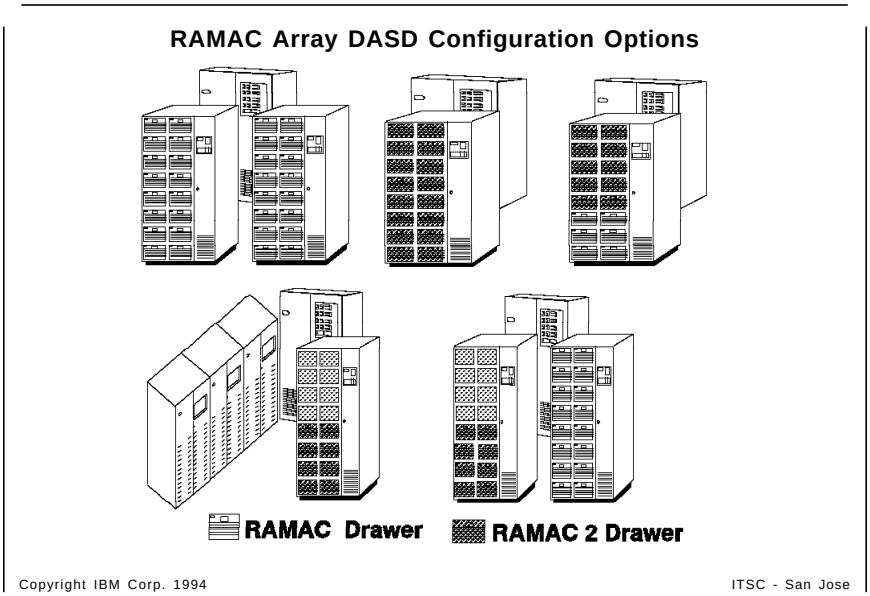

Config-40 RAMAC Array DASD Configuration Options

Foil Config-40 shows the various configuration options for the RAMAC Array DASD.

The attachment of two racks containing only RAMAC drawers is, of course, still a valid configuration option as depicted by the top-left configuration on the foil. This configuration option offers 180 GB of storage capacity in two racks attached to a 3990-6.

Moving clockwise, the second configuration depicted on the foil is a 3990-6 attached to a single rack containing 16 RAMAC 2 drawers. This configuration option offers the same 180 GB of storage capacity as the previous configuration, with half the number of racks.

The third configuration option depicts a single rack containing a mix of RAMAC and RAMAC 2 drawers and attached to a 3990-6. This configuration may arise when customers have a single rack partially populated with RAMAC drawers and want to increase capacity by adding RAMAC 2 drawers to the existing rack. This option is valid as long as the rack is the only string attached to the 3990-6. Customers can upgrade or replace their RAMAC drawers with RAMAC 2 drawers as they progress to the second option—the ideal configuration of a single, fully populated RAMAC 2 rack.

Continuing clockwise on the foil, the fourth configuration option has two strings attached to a 3990-6. One string is a fully populated RAMAC rack with 32 devices giving 90 GB of storage capacity. The second string is a half-rack of RAMAC 2 drawers. In this configuration, the 3990-6 protects against any violation of the 64 device address threshold by recognizing the addressability potential of the second string. The second string contains RAMAC 2 drawers and, if fully populated, would result in more than 64 device addresses. The 3990-6 configuration data therefore limits the number of drawers allowed in the second string to a maximum of eight. This configuration could arise when customers have an existing, fully populated rack of RAMAC drawers that they want to keep rather than replace. The second string offers customers an additional 90 GB of

storage capacity in a RAMAC 2 half-rack configuration. Thus, 3990-6 facilitates RAMAC to RAMAC 2 drawer replacement, and RAMAC and RAMAC 2 coexistence.

The fifth configuration option shows a 3990-6 with a full string of 3390 DASD attached. This string could be any model of 3390 DASD providing 32 devices addresses. This configuration is similar to the fourth configuration in terms of the number of RAMAC 2 drawers that can be installed in the second string. Because the 3390 string already occupies 32 device addresses, the rack can have up to eight drawers of RAMAC 2 installed.

In the fourth and fifth configurations, it is also possible to have RAMAC 2 drawers intermixed with RAMAC drawers in the half-rack. The maximum number of drawers in each case is still eight because the 3990-6 will protect against any possible breach of the device address thresholds.

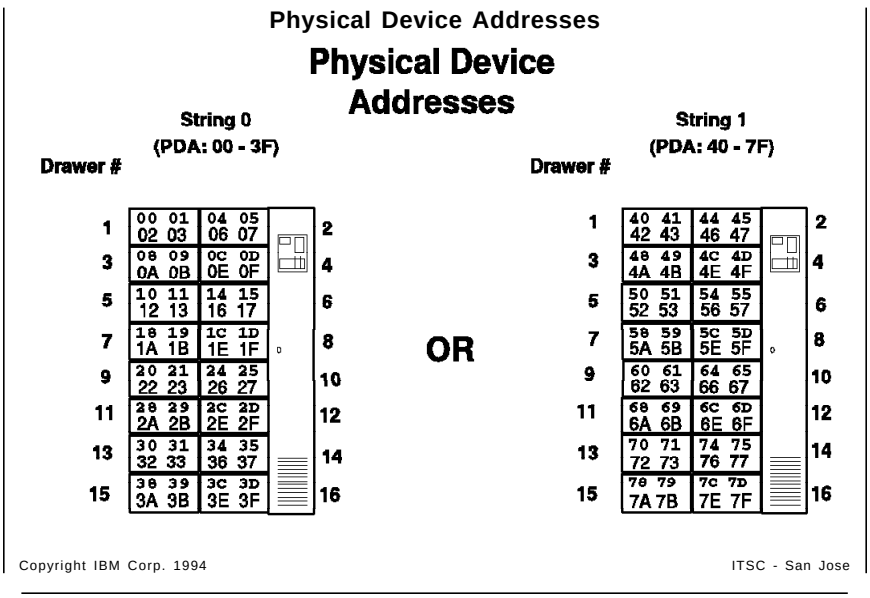

Config-41 Physical Device Addresses

This foil shows the physical device addresses (PDAs) or director to device connection (DDC) addresses used by the storage control to address the logical device on the controller interface. DDC addresses differ from the channel connection addresses (CCAs) used by the host to address the logical devices and are related to the position of the drawer in the rack. The DDC addresses do not change, although in dual copy or Dynamic Sparing situations, the relationship between the CCA and the DDC addresses will change. When that occurs, the 3990 and the RAMAC Array Subsystem manage and track the changes. The four DDC addresses in each drawer are contiguous. Even if RAMAC drawers are emulating two logical volumes, there are still four DDCs per drawer. For the RAMAC drawer, the two lowest DDC addresses are used. For example, if RAMAC drawers were installed in drawer positions 9 and 10 in string 0, the DDC addresses used will be 20, 21, 24, and 25.

RAMAC Array DASD Address Mapping A32

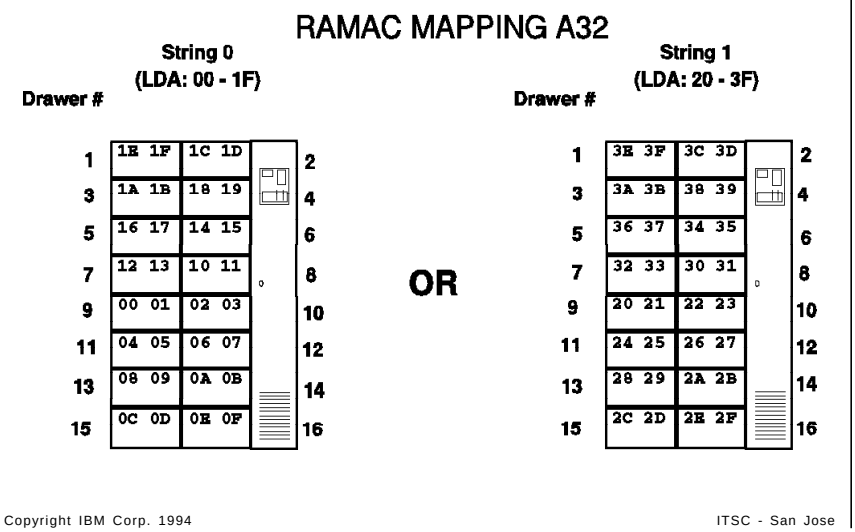

 ITSC - San Jose

Config-42 RAMAC Array DASD Address Mapping A32

Foil Config-42 shows the logical device addresses (LDAs) or CCAs for RAMAC drawers as used today. The IBM 3990 model 6 uses address mapping A32 when 3990 DASD strings are attached or there are racks with only RAMAC drawers. This address mapping scheme does not support RAMAC 2 drawers because there are no addresses for the four logical volumes in the drawer. The CCAs are used by the host to address the devices and are predefined according to drawer location. The CCAs increment in line with the drawer installation sequence and are contiguous within each drawer. For example, drawer number 8 is installed after drawer number 16. The CCAs of the logical volumes in these drawers are 0E, 0F, 10, and 11 (hexadecimal). It can be seen from the foil that the logical device addressing in string 1 continues from string 0, that is, the first CCA in string 1 is contained in drawer number 9, so 32 addresses are already reserved for string 0. If string 1 was the only string installed, the addressing would start at 20 (hexadecimal) and increment through 3F. Thus, a gap of 32 device addresses has been made to support the installation of string 0 at a later time.

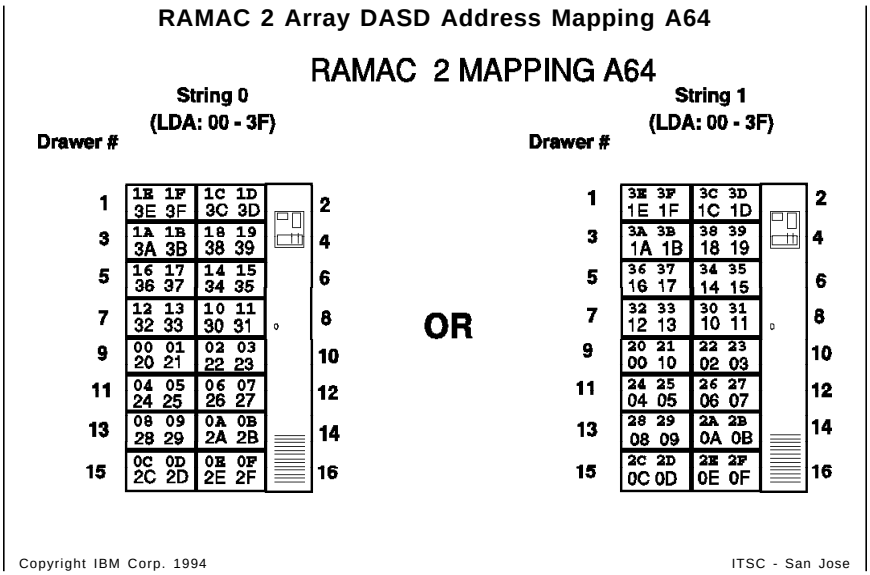

Config-43 RAMAC 2 Array DASD Address Mapping A64

For simplicity we call this addressing scheme “mapping A64.” Mapping A64 is used when RAMAC 2 drawers are in the rack and no other string is attached to the 3990-9. Mapping A64 is designed for customers who want to upgrade from RAMAC drawers to RAMAC 2 drawers in a single rack configuration and thus take advantage of cost-effective, single-rack, 180-GB configurations. Customers retain their existing 3990-6 and rack and upgrade their RAMAC drawers with RAMAC 2 drawers nondisruptively, using dual copy for migration assistance.

Notice how the additional 32 logical device addresses have been overlaid on the original 32 device addresses defined for RAMAC drawers. For example, RAMAC 2 drawer number 5 in string 0 contains logical device addresses 16, 17, 36, and 37. Recall from the previous foil that RAMAC drawer number 5 contained logical device addresses 16 and 17. Device addresses 16 and 17 are located in the same drawer for both RAMAC and RAMAC 2. This address mapping scheme enables an application using RAMAC addresses 16 and 17 to continue using those addresses when the drawer is upgraded to a RAMAC 2 drawer. If you compare this foil with the previous foil, you can see that every RAMAC 2 drawer has preserved the address location of the original RAMAC devices.

Mapping A64 is a powerful and flexible address mapping scheme that supports nondisruptive migration to the new RAMAC 2 environment if only one RAMAC rack is attached to the 3990-6.

Depending on which string the RAMAC rack was installed before the installation of RAMAC 2 drawers, two addressing schemes are available. If the rack was string 0, drawer 9 would have addresses 00, 01, 20, and 21. If the rack was string 1, drawer 9 would have addresses 20, 21, 00, 01.

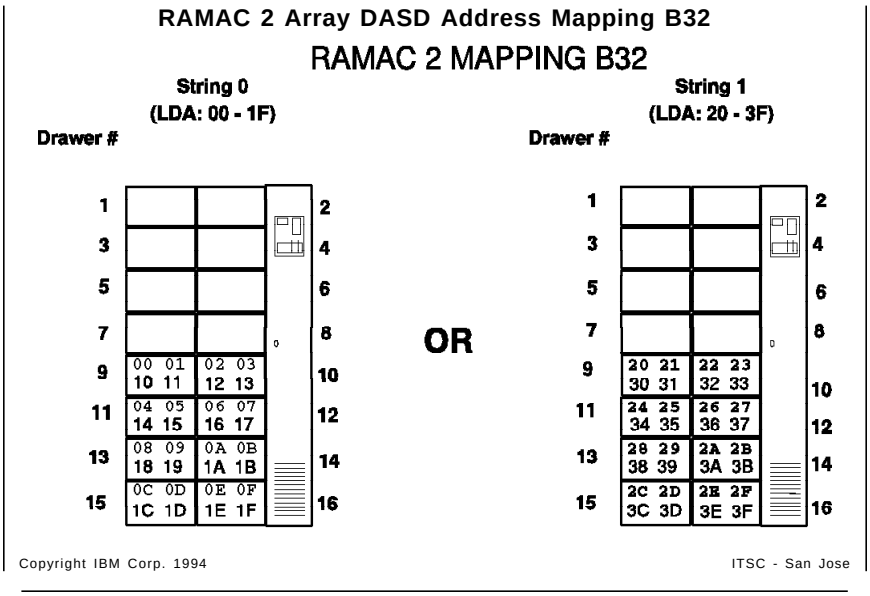

Config-44 RAMAC 2 Array DASD Address Mapping B32

RAMAC 2 address mapping B32 is used when customers want to add RAMAC 2 capacity to an existing RAMAC Array DASD subsystem. Thus an existing string of RAMAC will coexist with a new string of RAMAC 2. This mapping is an alternative for customers who do not want to replace RAMAC drawers with RAMAC 2 drawers in the same rack. Because the 3990-6 protects the 64 device address threshold, the second string of RAMAC 2 must have a maximum of eight drawers. This half-rack concept allows coexistence of RAMAC and RAMAC 2.

The address range is contiguous within the half-rack, although not contiguous within each drawer. Thus, 32 contiguous device addresses are supported in this half-rack configuration. The contiguous addressing, half-rack configuration is important because this mapping scheme assumes that 32 other RAMAC device addresses already exist—either as string 0 or string 1. If the existing RAMAC rack is string 0, the half-rack of RAMAC 2 will be string 1 and will use the 32 devices addresses starting at 20 (hexadecimal) and incrementing through 3F. Notice how the RAMAC 2 string 1 addresses “leave a gap” of 32 devices. However, if the existing RAMAC string is installed as string 1, its device range will be 20 to 3F, leaving the first 32 addresses to be used for the RAMAC 2 half-rack.

RAMAC 2 Array DASD Address Mapping B64

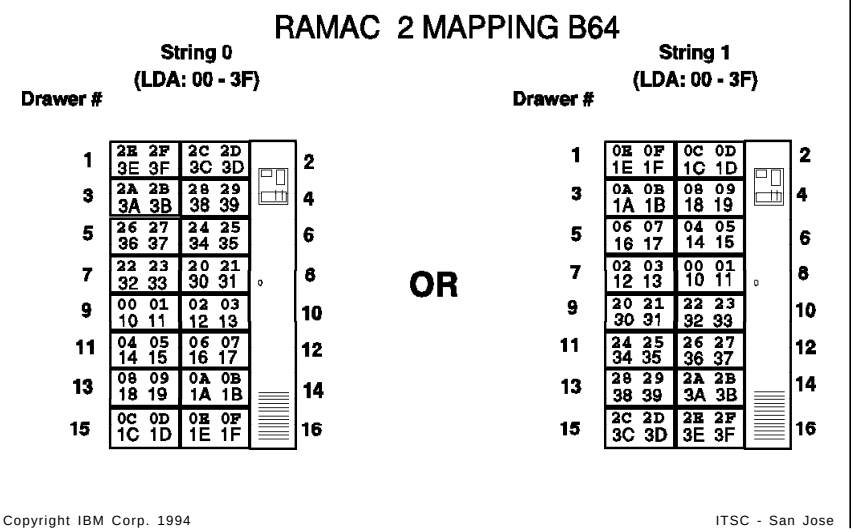

 ITSC - San Jose

Config-45 RAMAC 2 Array DASD Address Mapping B64

Address mapping B64 enables customers to move to an all-RAMAC 2 configuration from the previous, mixed RAMAC and RAMAC 2 configuration. Recall that address mapping B32 allowed a string of RAMAC to coexist with a half-string of RAMAC 2. Customers wanting to move to a single-rack, 180-GB configuration of all-RAMAC 2 from the address mapping B32 environment can do so without impacting the addressability of their existing RAMAC 2 drawers.

Mapping B64 defines the additional 32 addresses “on top of” the existing RAMAC 2 device addresses whether the RAMAC 2 string is string 0 or string 1. The device addresses are now contiguous on an upper-rack and lower-rack basis. Any existing RAMAC string would have to be deinstalled before migrating to a full rack of RAMAC 2 in this way. Address mapping B64 allows for up to 64 addresses in a single rack. If the original half-rack was installed as string 0, addresses in the newly populated full rack will range from 00 to 1F in the lower eight drawers and 20 to 3F in the upper eight drawers. If the original half-rack was installed as string 1, addresses in the newly populated full rack will range from 20 to 3F in the lower eight drawers and 00 to 1F in the upper eight drawers.

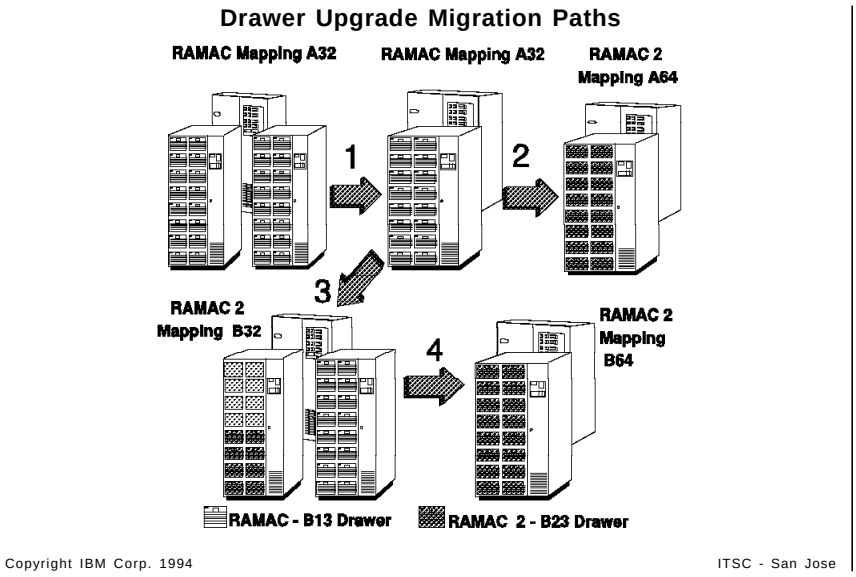

Config-46 Drawer Upgrade Migration Paths

Foil Config-46 summarizes the upgrade migration paths from a RAMAC Array DASD configuration to a RAMAC 2 Array DASD configuration. Starting from a RAMAC Array DASD configuration with two strings, you choose to first remove the second string (step 1) and then upgrade RAMAC drawers to RAMAC 2 drawers in the existing rack (step 2). In this case the 3990-6 switches nondisruptively from RAMAC mapping A32 to RAMAC 2 mapping A64.

Alternatively, you may want to keep the RAMAC string for a while and add RAMAC 2 drawers to the configuration instead of upgrading existing drawers (step 3). This involves a migration from RAMAC mapping A32 to RAMAC 2 mapping B32.

Before more than eight drawers can be installed in the RAMAC 2 string, the RAMAC string must be removed to pave the way for a full string of RAMAC 2 giving 180 GB in a single rack. Migration to this configuration results in moving from RAMAC 2 mapping B32 to RAMAC 2 mapping B64 (step 4).

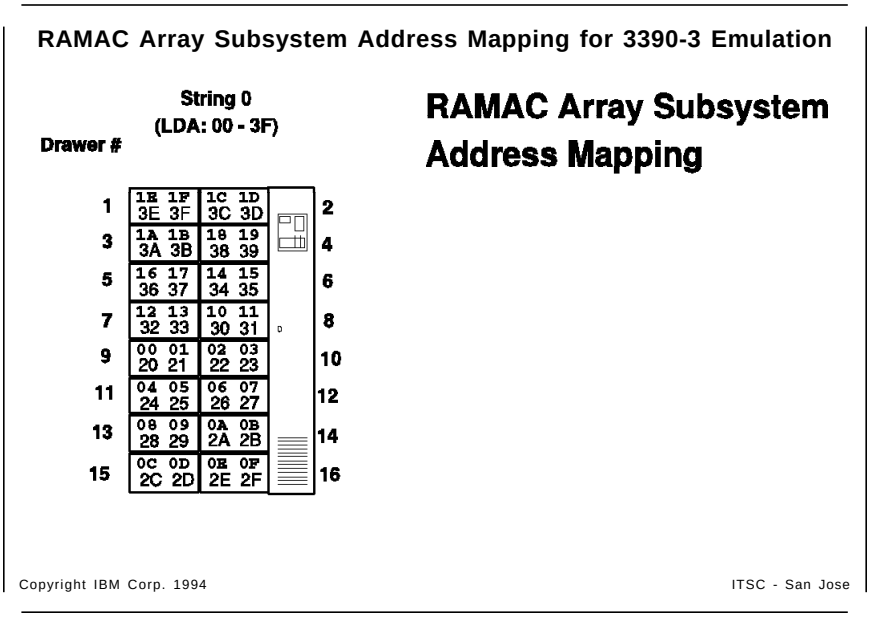

Config-47 RAMAC Array Subsystem Address Mapping for 3390-3 Emulation

There is only one address mapping for 3390-3 emulation in a RAMAC Array Subsystem configuration. You can always add new RAMAC 2 drawers in the normal sequence or replace existing RAMAC drawers in any sequence. Note, however, that you get consecutive addresses if you replace RAMAC drawers starting with drawer number 9, followed by 10, and so on.

Notice how the RAMAC Array Subsystem “overlays” the additional two addresses in the four-volume drawer on the existing two addresses in the two-volume drawer. Thus data can be moved off the two-volume drawer, the drawer can be upgraded to a four-volume drawer, and the data moved back again. Because the addresses in the two-volume drawer have remained the same, the application can continue as normal after the data has been restored to the new four-volume drawer.

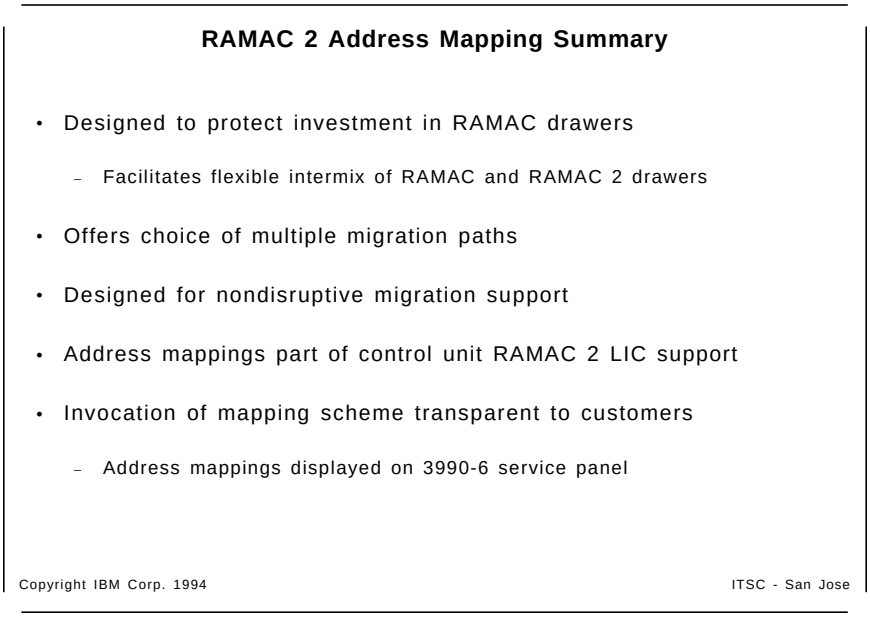

Config-48 RAMAC 2 Address Mapping Summary

IBM addresses the challenge of delivering flexible, high-capacity, RAID-5 storage solutions within the framework of the RAMAC Array Family design by enabling RAMAC and RAMAC 2 drawer intermix in the same rack. Drawer intermix is vital if you want to grow your existing RAMAC configuration to include RAMAC 2 technology.

- **Designed to protect investment in RAMAC drawers**

  Why does the 3990-6 implement multiple addressing schemes? To enable customers to take advantage of the cost-effective improvements of RAMAC 2 while retaining existing RAMAC DASD and provide low-risk migration capability to the new RAMAC 2 environment. Customers who are using RAMAC drawers and require additional RAMAC capacity can install both types of drawers in the same rack or in the same subsystem. Because RAMAC and RAMAC 2 can be flexibly intermixed, customers can exploit the latest RAMAC family technology and design while using existing technology and design.

- **Offers choice of multiple migration paths**

  The address mapping schemes provide customers with a choice in how to migrate to the RAMAC 2 environment. There are address mapping schemes to support drawer upgrades, RAMAC and RAMAC 2 coexistence within the same rack, RAMAC and RAMAC 2 coexistence within the same subsystem, and different all-RAMAC 2 full-rack device addressing. The customer's original configuration before migration determines the address mapping.

- **Designed for nondisruptive migration support**

  The address mapping options carefully take into account the existing device address definitions by overlaying the additional two addresses in a four-volume drawer with those of the two-volume drawer. The original addresses of a two-volume drawer occupy the same location when the drawer is upgraded to a four-volume drawer. Thus customer applications can access data continuously while the drawers are upgraded. Dual copy provides a nondisruptive drawer upgrade assist by copying the data to

another drawer before the upgrade kit is applied. The data can then be copied back to the newly upgraded drawer. Because dual copy is a logical volume copy function, it is independent of drawer type.

- **Address mappings part of control unit RAMAC 2 LIC support**

  The address mapping schemes are included as part of the 3990-6 LIC enhancement that supports RAMAC 2. To take advantage of the flexible addressing schemes for RAMAC 2 all that installations need do is download and activate the RAMAC 2 supporting LIC. Installation personnel are aware of the address mapping in operation at all times because it is displayed on the 3990-6 service panel.

- **Invocation of mapping scheme transparent to customers**

  When the service representative installs RAMAC 2 drawers in a rack attached to the 3990-6, he or she configures the 3990-6 in accordance with the installation requirements of the customer. If the customer is upgrading drawers, or installing a second string of RAMAC 2 with an existing string of RAMAC, the service representative uses the CE panel to validate and configure the subsystem. The invocatio of the addressing scheme required is transparent to the customer.

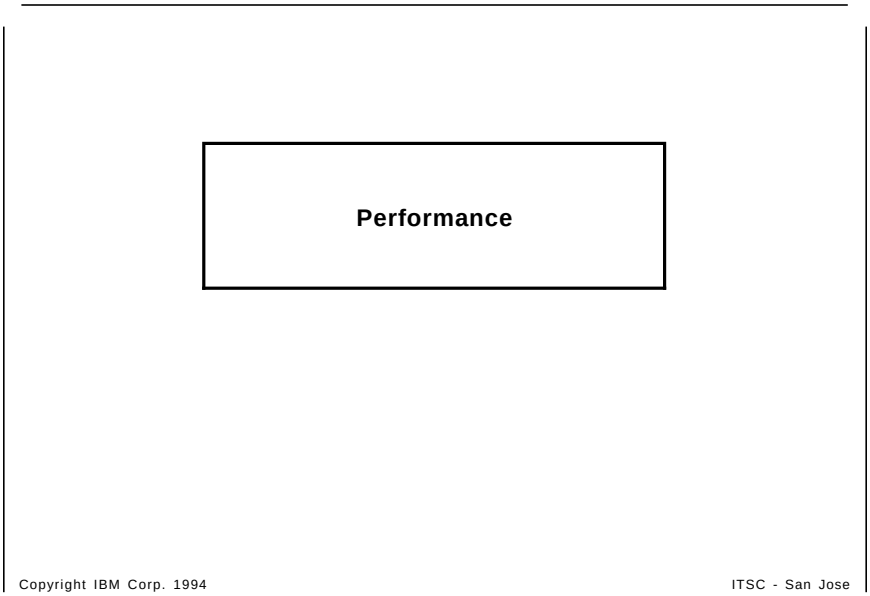

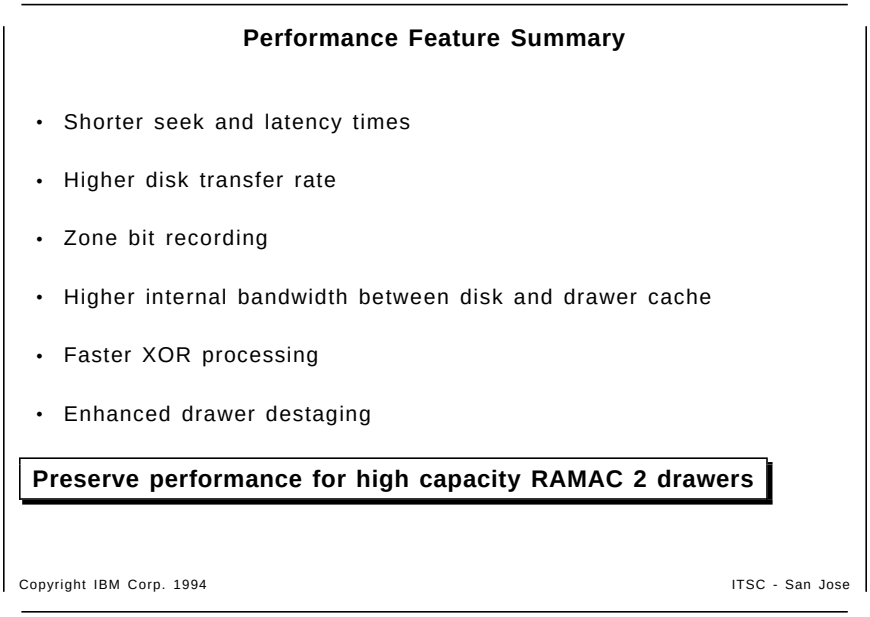

Perform-50 Performance Feature Summary

Foil Perform-50 summarizes the design changes that have been made to preserve the performance of the RAMAC 2 high-capacity, enhanced RAID-5 drawer.

- **Shorter seek and latency times**

  The disk rotates at a higher speed of 7200 RPM while the data is packed more closely together. Faster disk rotation leads to reduced latency time.

- **Higher disk transfer rate**

  The disk media transfer rate is about twice as high for a RAMAC 2 disk compared to a RAMAC disk.

- **Zone bit recording**

The transfer rate is even higher in the outer zones as more data is stored in an outer track when the relative speed between the MR head and the disk surface is higher.

- **Higher internal bandwidth between disk and drawer cache**

  To cope with the higher disk media transfer rate the SCSI path manager has a higher bandwidth for a faster stage and destage of data between the disk and the drawer cache.

- **Faster XOR processing**

  The heart of a RAID-5 operation, the calculation of parity by an XOR operation, is enhanced by faster hardware.

- **Enhanced drawer destaging**

  The efficiency of the RAMAC 2 drawer destaging process has been enhanced by new techniques that group destage data together and perform “back-to-back” writes to the disks to avoid excessive disk usage. The data is organized such that the destage can take place in a single sweep of the disk in one direction.

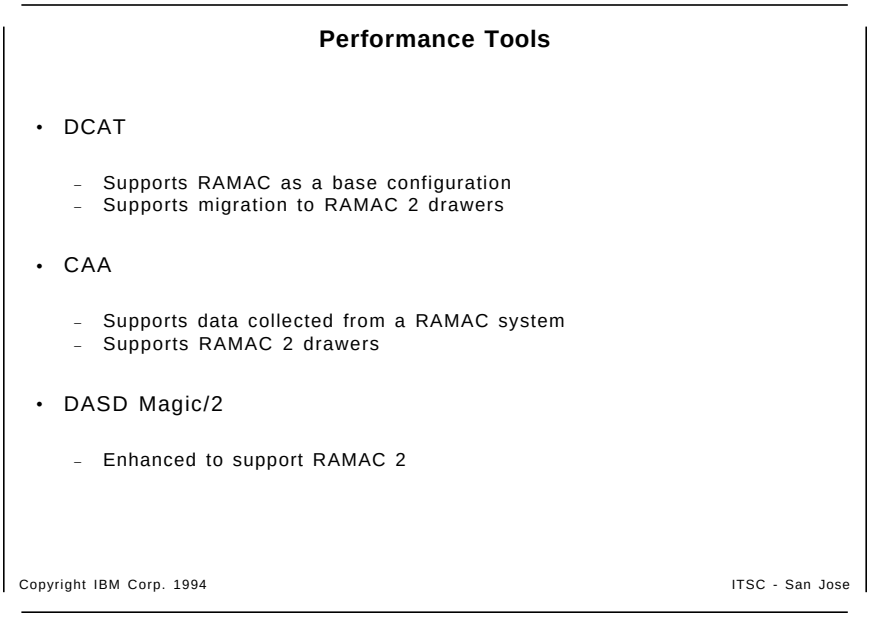

Perform-51 Performance Tools

Some tools are available to model DASD and control unit performance.

- **DCAT**

  DCAT supports RAMAC as a base configuration. Thus you can collect data from a subsystem with RAMAC devices. DCAT supports migrations from older devices to RAMAC 2 drawers and from RAMAC drawers to RAMAC 2 drawers. DCAT is available on HONE.

- **CAA**

  The Cache Analysis Aid (CAA) program simulates caching behavior. CAA supports data collected from a RAMAC system. It supports drawer cache simulation for the RAMAC 2 drawers. CAA is available from MKTTOOLS.

- **DASD Magic/2**

  DASD Magic/2 is an OS/2 tool for high-level DASD and control unit analysis. DASD Magic/2 is enhanced to support RAMAC 2 drawers. The tool is available from MKTTOOLS.

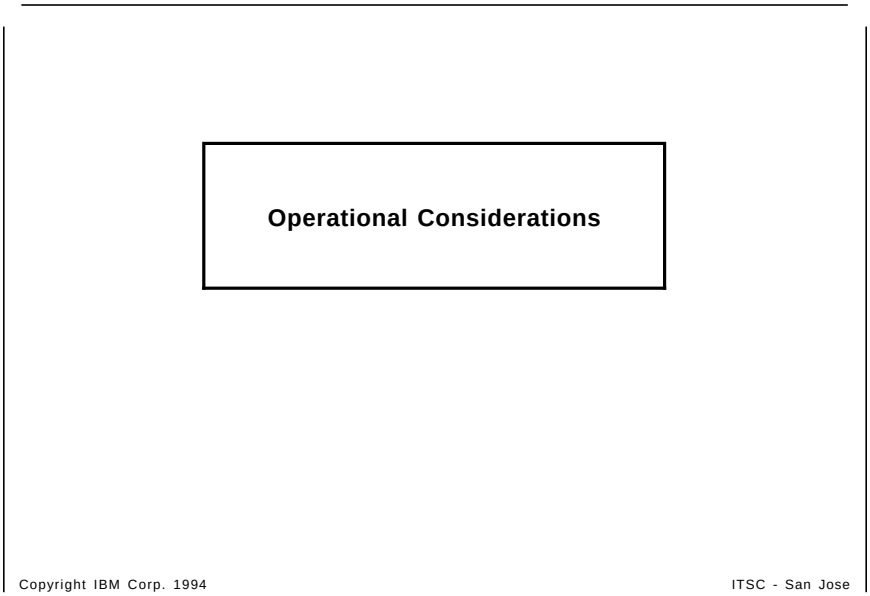

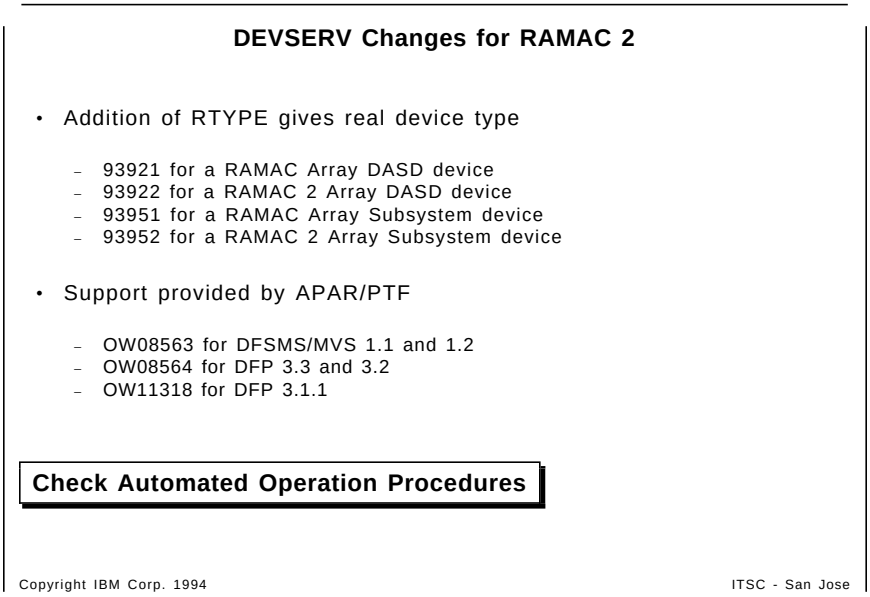

Oper-53 DEVSERV Changes for RAMAC 2

The availability of RAMAC 2 devices is accompanied by enhanced monitoring capabilities in the DEVSERV operator command. The enhancements enable customers to identify which devices are RAMAC and which are RAMAC 2.

- **Addition of RTYPE gives real device type**

  The DEVSERV command is enhanced by the addition of RTYPE in the same line of the command output as the CACHE status. This field gives the real device type when the DTYPE represents the emulated type. Possible values are 93921 and 93922 for RAMAC Array DASD, and 93951 and 93952 for RAMAC Array Subsystem devices. For RAMAC Array Subsystem devices, the CFW, TC, DFW, PIN, and DC-STATE fields are not shown.

- **Support provided by APAR/PTF**

  The foil shows the APAR numbers whose PTFs provide the enhanced DEVSERV command support for the real device type.

The change of the DEVSERV command may influence automated operations procedures. Review the procedures for any changes required.

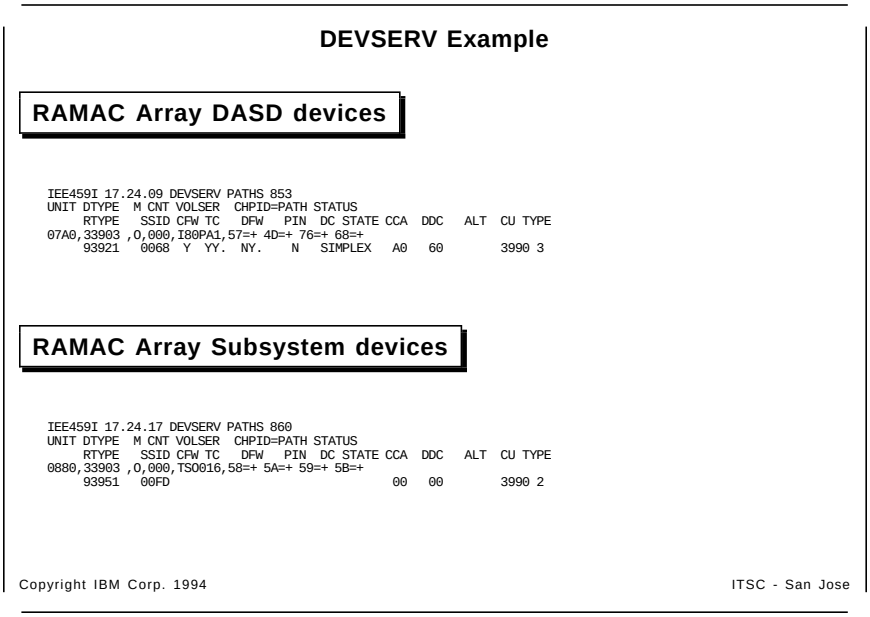

Oper-54 DEVSERV Example

Foil Oper-54 shows the output of a DEVSERV command issued to RAMAC devices behind an IBM 3990 storage control and a RAMAC Array Subsystem, respectively.

In the first example the emulated device (DTYPE) is a 3390 model 3, and the real device type (RTYPE) is a RAMAC Array DASD RAMAC drawer (93921). In the second example the real device type is a RAMAC Array Subsystem RAMAC drawer (93951).

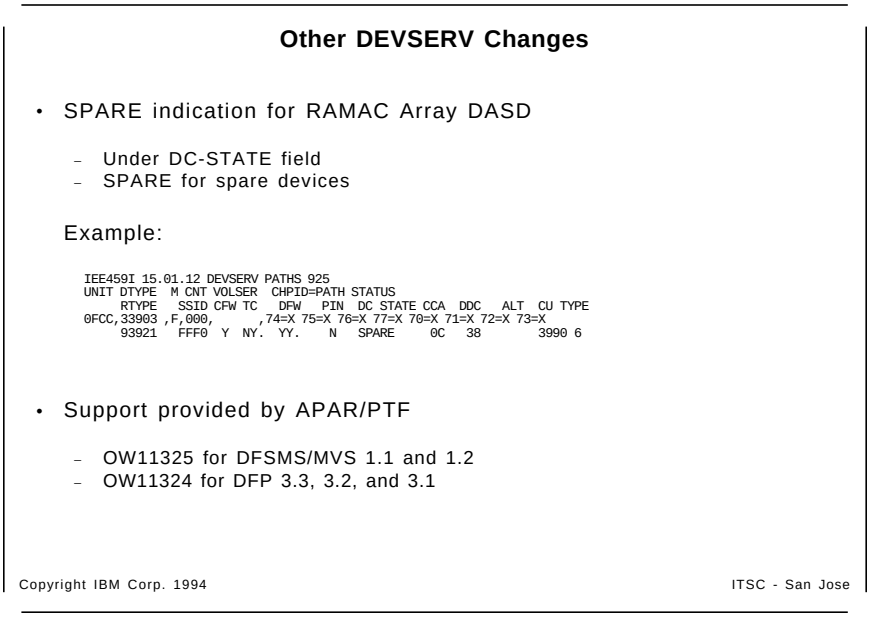

Oper-55 Other DEVSERV Changes

The DEVSERV report has been enhanced to provide additional status information that is useful in monitoring storage subsystems configured for Dynamic Sparing.

- **SPARE indication for RAMAC Array DASD**

  Under the DC-STATE field, DEVSERV shows SPARE for spare devices for RAMAC Array DASD only. The DEVSERV command will not, however, track the DDC-CCA address relationship after Dynamic Sparing starts.

- **Support provided by APAR/PTF**

  The foil shows the APAR numbers whose PTFs provide the enhanced DEVSERV command support for spare volumes.

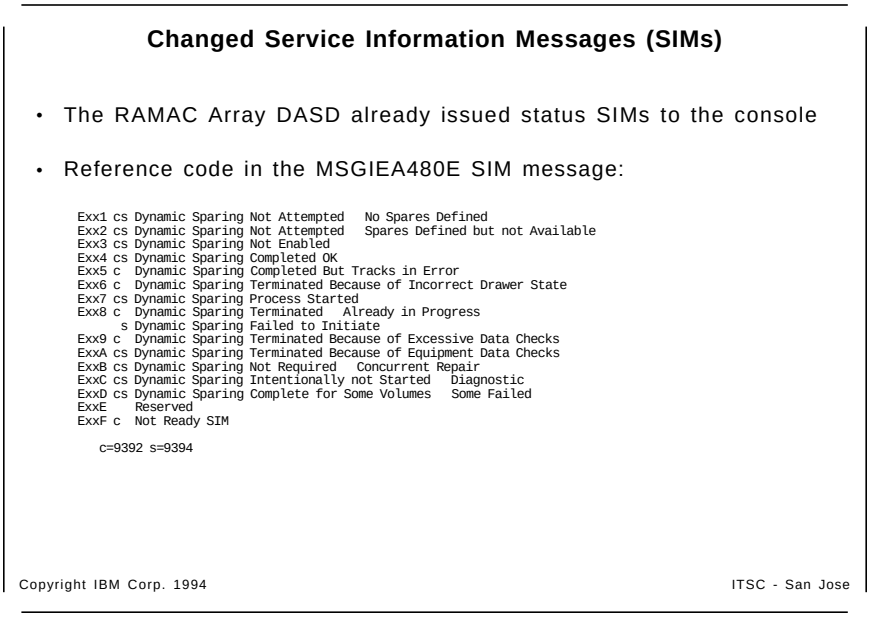

Oper-56 Changed Service Information Messages (SIMs)

A SIM flows to the service director and to EREP (recorder file) with bits modified to indicate the initiation and completion of sparing. These are level 0 SIMs and, unless the VPD is changed, they do not flow to the console. The default VPD setting for SIM reporting is level two and above. If you make this change you will see MSGIEA480E, which is the standard SIM for MVS. The VPD default for the RAMAC Array Subsystem is to report all SIMs to the console. For the RAMAC Array DASD, service SIMs must be enabled in the 3990 storage control.

- **The RAMAC Array DASD already issued status SIMs to the console**

  The RAMAC Array DASD already has code that supports SIMs to the console about the status of Dynamic Sparing.

- **Reference code in the MSGIEA480E SIM message:**

  A digit in the reference code of the SIM represents the Dynamic Sparing status. The foil shows the reference codes and their meaning. The status codes differ in some cases for a RAMAC Array DASD and a RAMAC Array Subsystem.

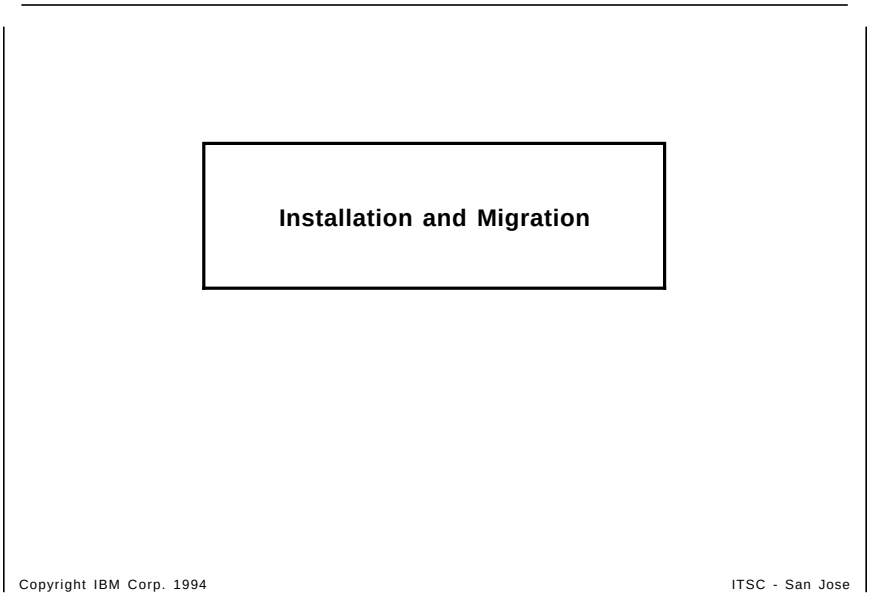

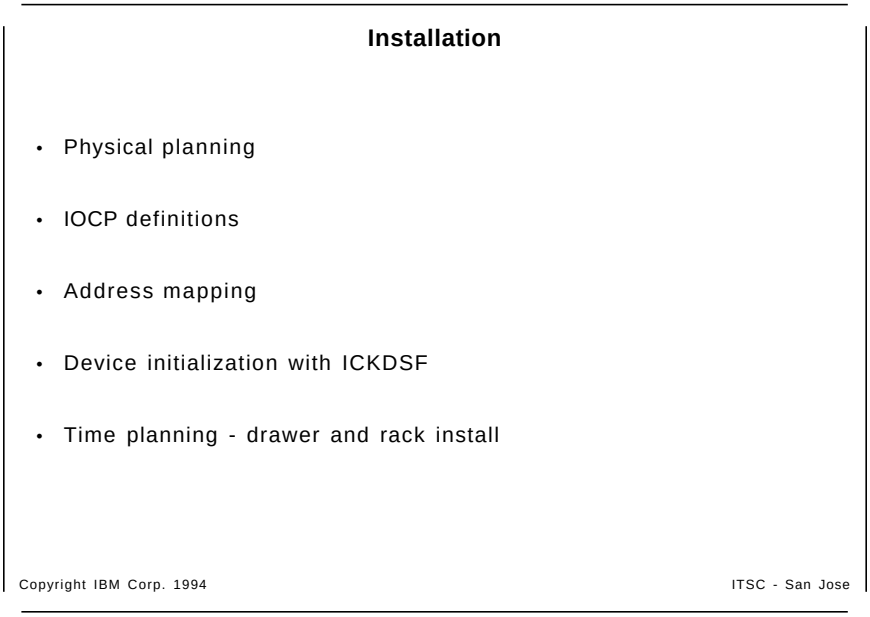

Install-58 Installation

The installation process of a RAMAC Array DASD or a RAMAC Array Subsystem includes the tasks listed on the foil.

- **Physical planning**

  Although a RAMAC 2 drawer has twice the capacity of a RAMAC drawer the size and weight are the same for both. In addition the size and environmental characteristics of a RAMAC rack have not changed. Therefore physical planning has not changed for RAMAC 2 devices, except that now fewer racks, fewer drawers, and less power are required for the same capacity.

- **IOCP definitions**

  Not much has changed with regard to the IOCP definitions for RAMAC 2 drawers. Where you once specified 32 addresses, for example, for a RAMAC Array Subsystem IOCP definition with RAMAC drawers, you should specify 64 addresses for a RAMAC Array Subsystem that has at least one RAMAC 2 drawer. Four addresses are reserved for each drawer position. If there is only a RAMAC drawer, two addresses are not used. If the first RAMAC 2 drawers are installed starting from position 9, the address range of available devices is consecutive.

- **Address mapping**

  With RAMAC 2 drawers in a RAMAC Array DASD you should be familiar with the address mapping used by the storage control (see foils Config-41 through Config-45). It is important to be able to identify the real drawer if there is a problem on a logical device.

- **Device initialization with ICKDSF**

  Before a device can be used it must be initialized by ICKDSF. For a RAMAC device do not run an ICKDSF job with the INSTALL command. RAMAC devices come preinstalled from the factory. All you have to do is run an ICKDSF INIT job that labels the devices and creates a VTOC (see foil Install-60). If the volume is being initialized for control program (CP) use in a VM environment the CPVOLUME command should be used. If the volume is being initialized for AIX use, the AIXVOL command should be used.

- **Time planning - drawer and rack install**

  The time needed to install a RAMAC system with RAMAC 2 drawers is about the same as for a system with RAMAC drawers. Estimated physical installation time is 6.2 hours for each rack shipped with 16 drawers. For field-installed drawers, estimated install time is 0.5 hours per single drawer. These time estimates assume that the storage control support for the RAMAC 2 drawers has been installed.

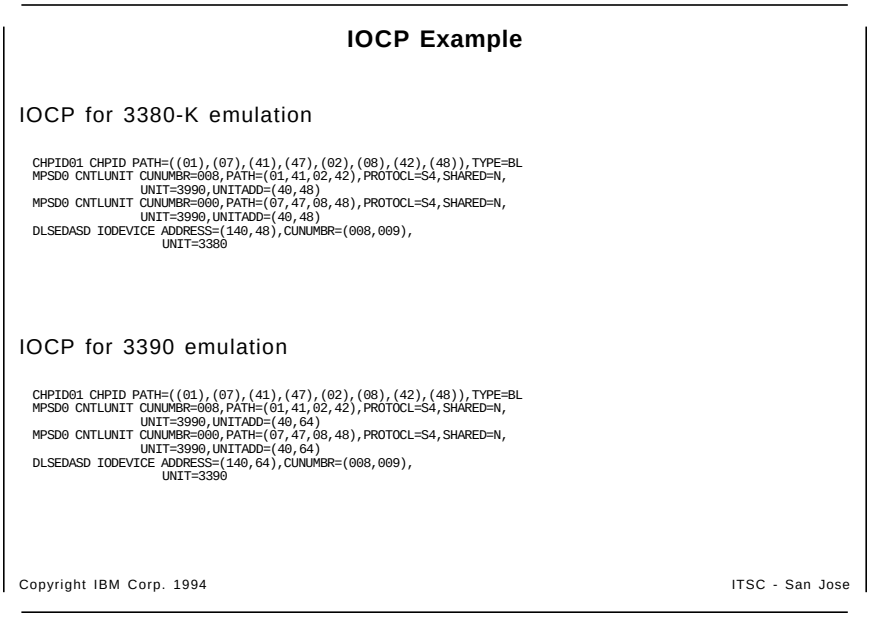

Install-59 IOCP Example

Changes are not required in the IOCP definitions for RAMAC 2 drawers if 64 addresses have been defined. If only 32 addresses are defined for a RAMAC Array Subsystem, change the number of addresses to 64 if RAMAC 2 drawers are used. The 3380-K IOCP example on the foil is valid only for a RAMAC Array Subsystem model 1 or model 2 with RAMAC drawers. In a RAMAC Array Subsystem model 3 with two controllers, one-half of the rack can be equipped with RAMAC 2 drawers and one-half with RAMAC drawers. In this latter case, however, you should specify only 24 addresses for 3380-K.

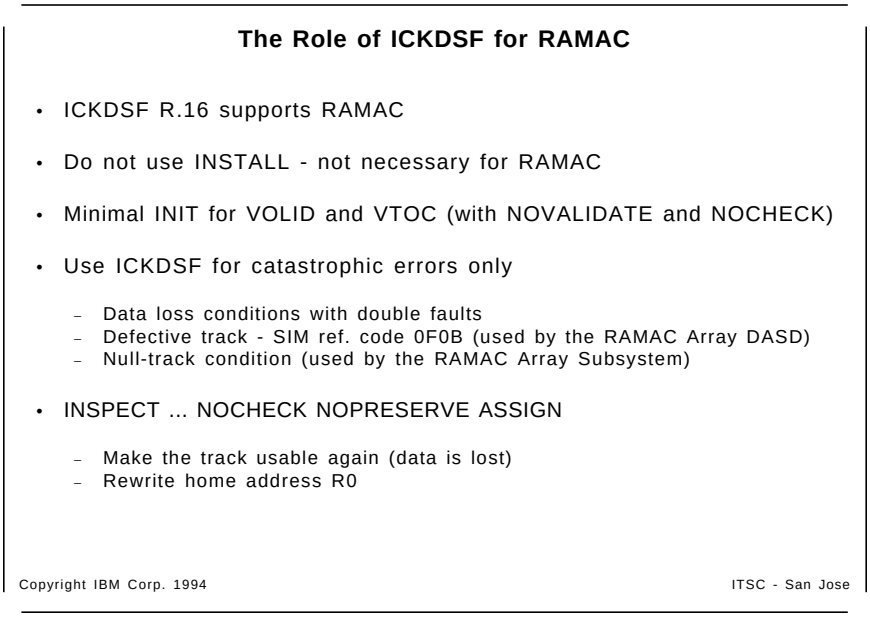

Install-60 The Role of ICKDSF for RAMAC

- **ICKDSF R.16 supports RAMAC**

  ICKDSF Release 16 should be used for any RAMAC device. This release issues a read device characteristics command, recognizes the real RAMAC device type, and acts accordingly.

- **Do not use INSTALL - not necessary for RAMAC**

  A RAMAC device need not be installed with the INSTALL command. A RAMAC drawer is delivered preformatted from the factory.

- **Minimal INIT for VOLID and VTOC**

  Only a minimal INIT step is required to label the device and define the VTOC and VTOC index. The NOVALIDATE and NOCHECK options should be specified. NOVALIDATE and NOCHECK are the default options. For control program (CP) use in a VM system the RAMAC devices must be initialized with the ICKDSF CPVOLUME command.

- **Use ICKDSF for catastrophic errors only**

  The drawer microcode performs most of the media maintenance. A correctable data check SIM should never be reported for a RAMAC device. The automatic media maintenance capability of the drawer leaves ICKDSF with the task of dealing with severe data loss conditions only. These data loss conditions can occur with double faults when the drawer LIC cannot correct the problem. A RAMAC Array DASD can indicate such a defective track condition with a SIM reference code of 0F0B. A RAMAC Array Subsystem signals a null-track condition.

- **INSPECT ... NOCHECK NOPRESERVE ASSIGN**

  In case of a severe problem you must run ICKDSF with the INSPECT command and the NOCHECK (do not try to read the data), NOPRESERVE (do not try to preserve the data), and ASSIGN options. If you request PRESERVE on a RAMAC Array Subsystem, ICKDSF tries to preserve the data if possible; on a RAMAC Array DASD, ICKDSF rejects the command.

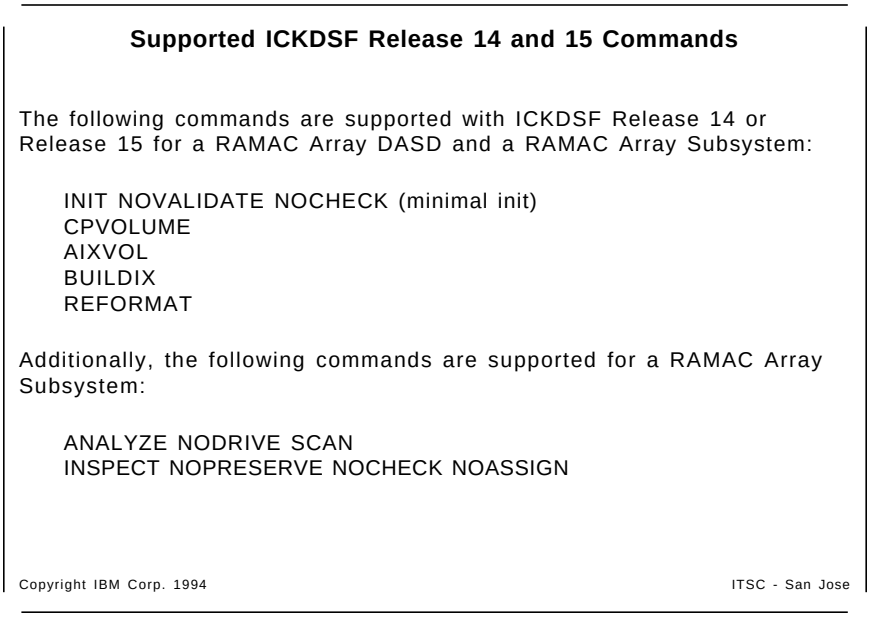

Install-61 Supported ICKDSF Release 14 and 15 Commands

Foil Install-61 lists those ICKDSF commands that are allowed with ICKDSF Release 14 or Release 15 for RAMAC devices. The validity of these commands for RAMAC devices allows you to use previous versions of ICKDSF for operating system functions. Use of any other command or command and parameter combination, particularly for media maintenance on a RAMAC Array DASD, requires ICKDSF Release 16.

ICKDSF Release 15 went out of service on 9/30/95.

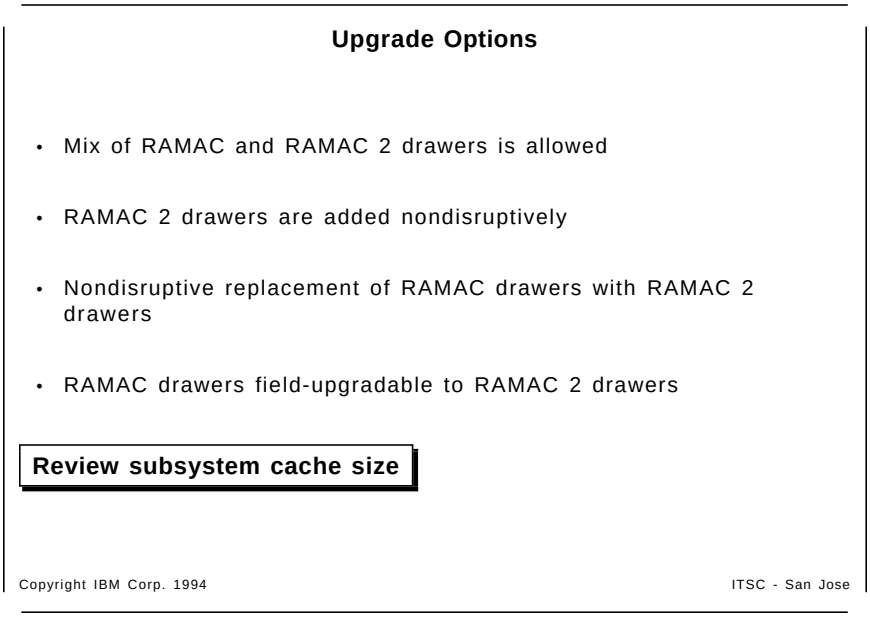

Install-62 Upgrade Options

The migration from a RAMAC to a RAMAC 2 configuration involves control unit microcode upgrade, drawer upgrades, or drawer replacement. The foil lists some upgrade considerations.

- **Mix of RAMAC and RAMAC 2 drawers is allowed**

  A migration from RAMAC drawers to the high-capacity RAMAC 2 drawers is an easy task. Changes to the software are not required. RAMAC and RAMAC 2 drawers can be mixed in a rack as long as you observe the addressing requirements of the IBM 3990 model 6.

- **RAMAC 2 drawers are added nondisruptively**

  After activation of the RAMAC 2 LIC, the installation of new RAMAC 2 drawers is nondisruptive.

- **Nondisruptive replacement of RAMAC drawers with RAMAC 2 drawers**

  Existing RAMAC drawers can be replaced with RAMAC 2 drawers. With the RAMAC Array DASD, you can use dual copy to save and copy back the data. Dual copy enables concurrent drawer replacement.

- **RAMAC drawers field-upgradable to RAMAC 2 drawers**

  It is possible to field-upgrade RAMAC drawers to RAMAC 2 drawers. The four disk drives and three cards of the drawer must be changed.

With an upgrade to RAMAC 2 or the installation of additional RAMAC 2 drawers you add storage capacity to the subsystem. You should review the cache size of the subsystem. Is the cache size large enough to sustain the higher load that is expected with the increase in capacity?

Field Reformat Option

A RAMAC Array Subsystem equipped with RAMAC drawers can be reformatted in the field.

- Reformat from 3380-K track format to 3390 model 3 format
- Reformat from 3390 model 3 track format to 3380-K format
- Copy data to tape or other DASDs on another storage control
- Not an option for the RAMAC 2 drawers
- Two different track formats in a 9394 model 3
- CE performs reformat

 ITSC - San Jose

Install-62 Upgrade Options

Previously the track format of a RAMAC Array Subsystem was fixed by the orders feature code. Now microcode is available that allows a CE to reformat the drawers in a RAMAC Array Subsystem to another track format. Two reformat options are available.

- **Reformat from 3380-K track format to 3390 model 3 format**

  You can reformat the drawer from a 3380-K track format to a 3390 model 3 track format. This option allows a migration to the track format of the future.

- **Reformat from 3390 model 3 track format to 3380-K format**

  There is also the reverse option where you reformat the drawers from a 3390 model 3 track format to a 3380-K format. This option will be of interest for those customers who provide, for example, backup services for others.

- **Copy data to tape or other DASDs on another storage control**

  All data on the logical volumes of a RAMAC rack that is to be reformatted must be saved to tape or other DASDs on other subsystems. A reformat erases all data in the rack. For more information on how to migrate to 3390 track format refer to the *3390 Direct Access Storage Migration Guide*, GG24-3373.

- **Not an option for the RAMAC 2 drawers**

  The 3380-K track format is not supported for a RAMAC 2 drawer. Therefore reformat is not on option.

- **Two different track formats in a 9394 model 3**

  A RAMAC Array Subsystem 9394 model 3 has two storage controls in a rack, one in the upper half and one in the lower half of the rack. The drawers in the upper and lower half can emulate different track formats. One-half could be equipped with RAMAC 2 drawers emulating 3390 devices, and the other half could be all RAMAC drawers in 3380-K emulation mode.

- **CE performs reformat**

  The reformat of the drawers is a task for the customer engineer. Reformatting is performed from the RAMAC service panel.

Power and Cooling Requirements for 64 Volumes

Model	Total GB	kVA	GB/ kVA	kW	GB/kW	kBTU /Hr	GB/ kBTU /Hr
3990-3 3380-K	121	26.2	4.6	27.4	4.4	87.3	1.4
3990-3 3390-2	121	20.7	5.8	18.8	6.4	63.7	1.9
3990-6 3390-3	182	18.7	9.7	17.3	10.5	58.5	3.1
3990-6 3390-9	545	12.5	43.8	11.1	49.1	37.9	14.4
9343 9345	96	5.4	17.8	5.0	19.2	17.1	5.6
3990-6 9392-1	182	8.9	20.5	8.1	22.4	27.7	6.6
3990-6 9392-2	182	5.1	35.7	4.6	39.5	15.9	11.4
9394 9395-1	182	7.6	24.0	7.0	26.2	23.7	7.7
9394 9395-2	182	3.8	48.0	3.5	52.0	11.9	15.3

 ITSC - San Jose

Install-64 Power and Cooling Requirements for 64 Volumes

Foil Install-64 compares entire subsystems with the maximum number of storage devices (64) attached to the storage controller. kVA is the product of the effective values of voltage and current (kilovolt x ampere), sometimes called “apparent power.”

The RAMAC Array DASD and the RAMAC Array Subsystem equipped with RAMAC 2 drawers have the lowest power and cooling requirements.

Power and Cooling Requirements per Gigabyte

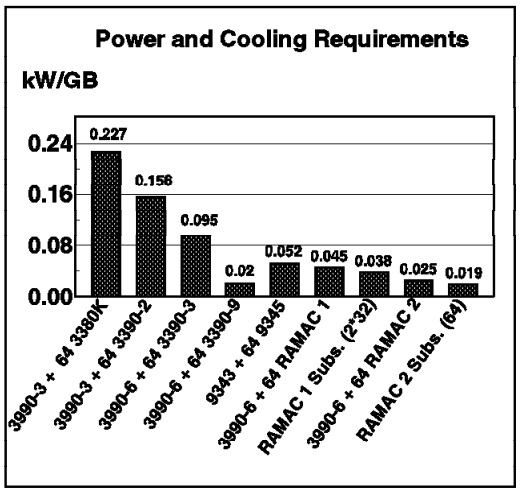

 ITSC - San Jose

Install-65 Power and Cooling Requirements per Gigabyte

Foil Install-65 compares the power and cooling requirements on a kW/GB basis. The RAMAC Array DASD and the RAMAC Array Subsystem equipped with RAMAC 2 drawers have the lowest power and cooling requirements per gigabyte. The total power required to operate a RAMAC Array DASD with RAMAC 2 drawers, for example, is less than 30% of the 3390-3 configuration.

Floor Space Requirements

Comparison for 182 GB

Model	Number of Vols	Number of Units	Required Square Meters	Required Square Feet
3990-3 3380-K	96	26	43.7	471.4
3990-3 3390-2	96	11	17.2	185.5
3990-6 3390-3	64	7	10.9	117.1
3990-6 3390-9	20	3	4.8	52.2
9343 9345	128	4	5.9	63.1
3990-6 9392-1	64	3	4.8	51.3
3990-6 9392-2	64	2	3.3	35.4
9394 9395-1	64	2	2.9	31.6
9394 9395-2	64	1	1.5	15.8

ITSC - San Jose

Install-66 Floor Space Requirements

Foil Install-66 compares entire subsystems with a full complement of storage devices attached to the storage controller (3990 and 3390-9 includes only two 3390-9 units in the subsystem). RAMAC Arrays have the full complement of B23 RAMAC 2 drawers. One-half service clearance is included in the calculations.

The RAMAC Array DASD and the RAMAC Array Subsystem equipped with RAMAC 2 drawers have the smallest floor space requirements.

Floor Space Requirements for 182 GB

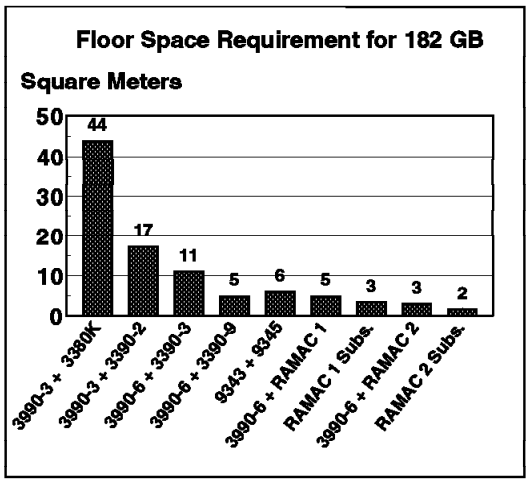

ITSC - San Jose

Install-67 Floor Space Requirements for 182 GB

Foil Install-67 compares the floor space requirements for various 182 GB configurations. The small footprint of RAMAC Arrays with RAMAC 2 drawers is striking. The required floor space for a RAMAC Array DASD with RAMAC 2 drawers, for example, is less than 30% of the 3390-3 configuration.

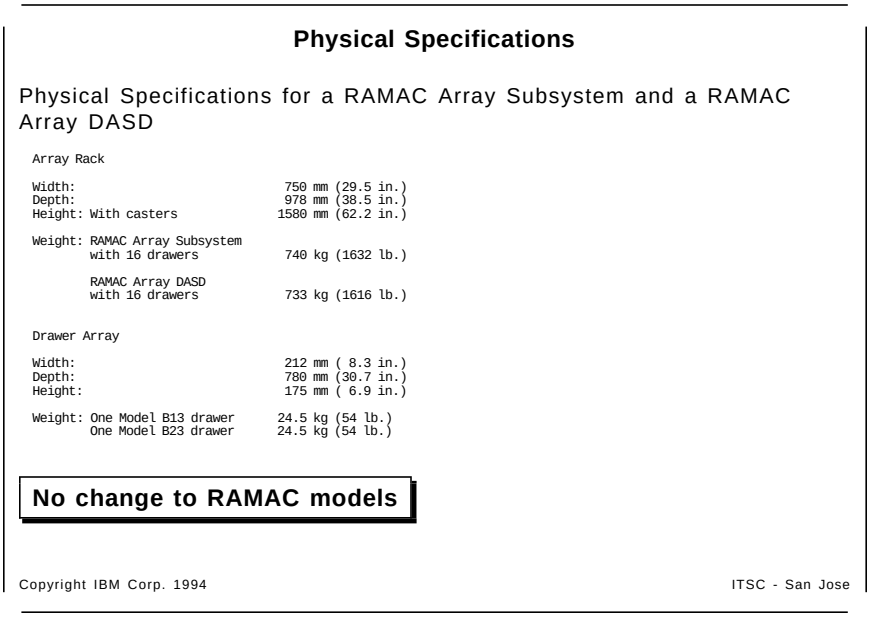

Install-68 Physical Specifications

Foil Install-68 shows the physical specifications for a RAMAC Array Subsystem model 2 or a RAMAC Array DASD rack. Each configuration is equipped with 16 B23 RAMAC 2 drawers.

The physical specifications are the same if the rack is equipped with B13 RAMAC drawers.

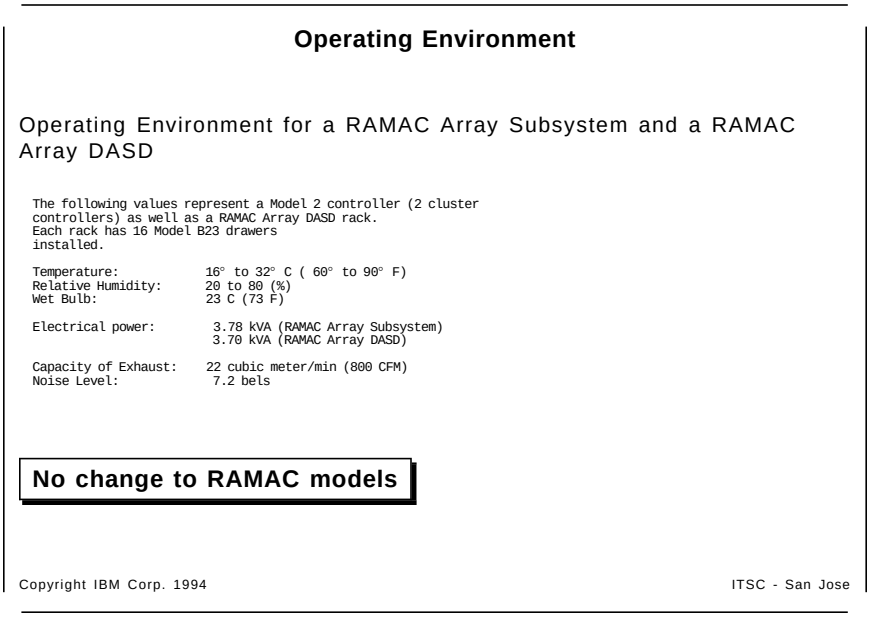

Install-69 Operating Environment

The foil shows the operating environment for a RAMAC Array Subsystem model 2 or a RAMAC Array DASD rack. Each configuration is equipped with 16 B23 RAMAC 2 drawers.

The environmental data are the same if the rack is equipped with B13 RAMAC drawers.

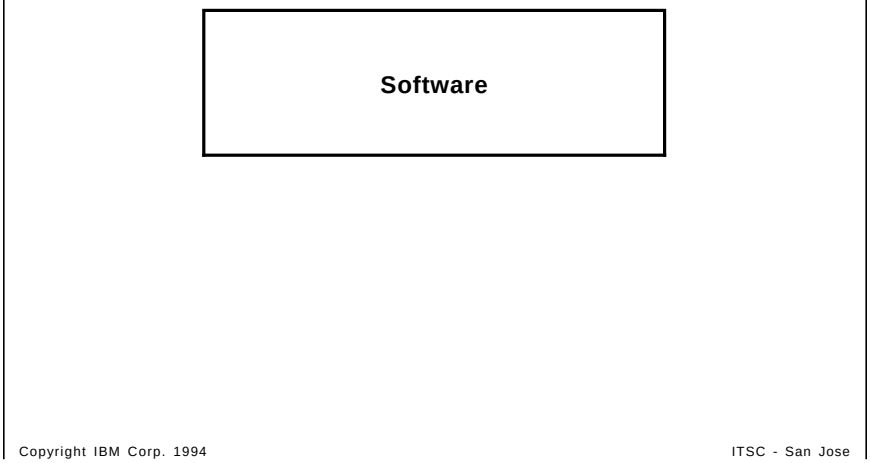

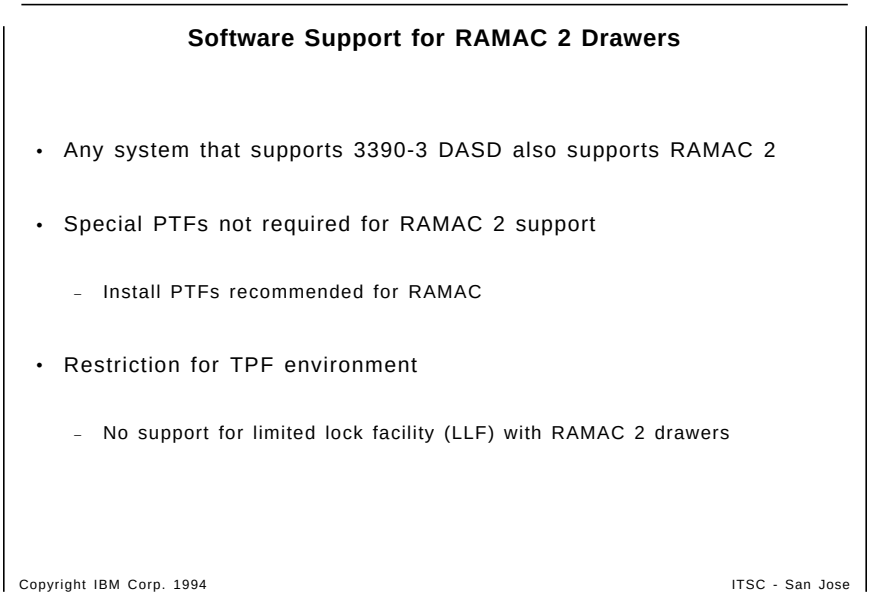

Softw-71 Software Support for RAMAC 2 Drawers

The installation of RAMAC 2 drawers is totally transparent to the software.

- **Any system that supports 3390-3 DASD also supports RAMAC 2**

  Any operating system that supports 3390 model 3 DASD attached to an IBM 3990 model 6 or an IBM 3990 model 2 also supports RAMAC 2 drawers in a RAMAC Array DASD or a RAMAC Array Subsystem.

  The following environments support RAMAC 2 drawers emulating 3390-3 DASD:

  – MVS operating environments

    - MVS/ESA SP Version 5 Release 1.0
    - MVS/ESA SP Version 4 Release 3.0
    - MVS/ESA SP Version 4 Release 2.2

- MVS/ESA SP Version 4 Release 2.0
- MVS/SP Version 2 Release 2.3 (discontinued service)
- MVS/SP Version 2 Release 2.0 (discontinued service)

- DFSMS/MVS Version 1 Release 2.0
- DFSMS/MVS Version 1 Release 1.0
- MVS/DFP Version 3 Release 3.2
- MVS/DFP Version 3 Release 3.1
- MVS/DFP Version 3 Release 3.0
- MVS/DFP Version 3 Release 2.1
- MVS/DFP Version 3 Release 2.0
- MVS/DFP Version 3 Release 1.1
- MVS/DFP Version 3 Release 1.0
- MVS/DFP Version 2 Release 4.0

– VM operating environments

- VM/ESA Release 2.2
- VM/ESA Release 2.1
- VM/ESA Release 2.0 (discontinued service)
- VM/ESA Release 1.1 (discontinued service)
- VM/XA SP Release 2.1 (discontinued service)

– VSE operating environments

- VSE/ESA Version 1 Release 3
- VSE/ESA Version 1 Release 2

– TPF operating environments

- TPF Version 4 Release 1
- TPF Version 3 Release 1

– AIX operating environments

- AIX/ESA Version 2 Release 2
- AIX/ESA Version 2 Release 1

• **Special PTFs not required for RAMAC 2 support**

PTFs are not required for RAMAC 2 drawer use. You might, however, want to install the PTFs that provide, for example, the real device type recognition that enables VM or the DEVSERV command in MVS to show not only the emulated but also the real device type. Installation of the recommended performance PTFs for RAMAC will result in more effective RAMAC and RAMAC 2 configurations.

• **Restriction for TPF environments**

Limited lock facility (LLF) is provided for use in transaction processing facility (TPF) environments. LLF supplies a symbolic locking mechanism for TPF such that a record is “locked” to one logical channel path for update integrity purposes. A logical static switch allows only one path to each DASD volume, such that data integrity is maintained when more than one host attempts to simultaneously update the same record.

RAMAC Array Subsystems or RAMAC Array DASDs operating in a TPF environment that uses LLF must be configured with parallel channels and have B13 drawers only. Support for LLF is not provided in a RAMAC Array Subsystem model 1 or 2 that has B23 drawers or in a model 3 rack partition that has B23 drawers. ESCON channels are not allowed on the controller in this environment. LLF is not supported in a RAMAC Array DASD with RAMAC 2 drawers.

Recommended PTFs for MVS and Related Products

Product	APAR	PTF	Description
MVS/ESA 4.2 MVS/ESA 4.3 MVS/ESA 5.1	OW09268	UW14956 UW14957 UW14958	Performance fix for formatting a JES2 SPOOL volume
EREP 3.5	IR27996	UR43179	Support for 9394 and 9395
ICKDSF Rel. 16	PN62330	UN67696	Fixes various problems
DFSORT	PN58749	UN68125	Format write performance
IMS/ESA 5.1 IMS/ESA 4.1 IMS/ESA 3.1	PN65578 PN65577 PN65575	UN71690 UN71372 UN71371	Performance fix for OSAM

 ITSC - San Jose

Softw-72 Recommended PTFs for MVS and Related Products

Although RAMAC devices can be used without any PTFs, some PTFs are recommended for better performance or to allow software to recognize the real device type.

MVS APARs for DFP and DFSMS/MVS

APAR for DFP	APAR for DFSMS/MVS	Description
OW09237	OW09476	SAM and EOV support for 9391 DASD
OW08564	OW08563	RAMAC real device type
OW10855	OW09702	Wrong message may be issued with dual copy
OW11403	OW11403	C.U.I.R VARY PATH rejected when spare defined
OW03645	OW03644	Invalid RMF shared disconnect time for RAMAC Array Subsystem
	OY63044	Command reject message during logical restore from 9394
	OW08402	Looping between AOM and Dev. Service Exit

 ITSC - San Jose

Softw-73 MVS APARs for DFP and DFSMS/MVS

Foil Softw-73 lists some recommended APARs for DFP and DFSMS/MVS. Check with your CE for the latest PTFs for these APARs.

Format Write PTFs for DFP and DFSMS/MVS

FMID	PTFs
HDP3320	UW09929, UW10233, UW11679, UW11281, UW11830
HDP3321	UW10234
HDP3330	UW09930, UW10235, UW11282, UW11680, UW11831
HDP3332	UW10236
JDZ1110	UW09928, UW10232, UW11284, UW11770, UW11833
JDZ11B0	UW09927, UW10231, UW11283, UW11769, UW11832

 ITSC - San Jose

Softw-74 Format Write PTFs for DFP and DFSMS/MVS

Write performance is crucial for a RAID-5 implementation. RAMAC drawers can reduce the write penalty if the drawer gets some information about the I/O, for example, if it is a sequential write. The foil lists the recommended PTFs that address format write performance for DFP and DFSMS/MVS.

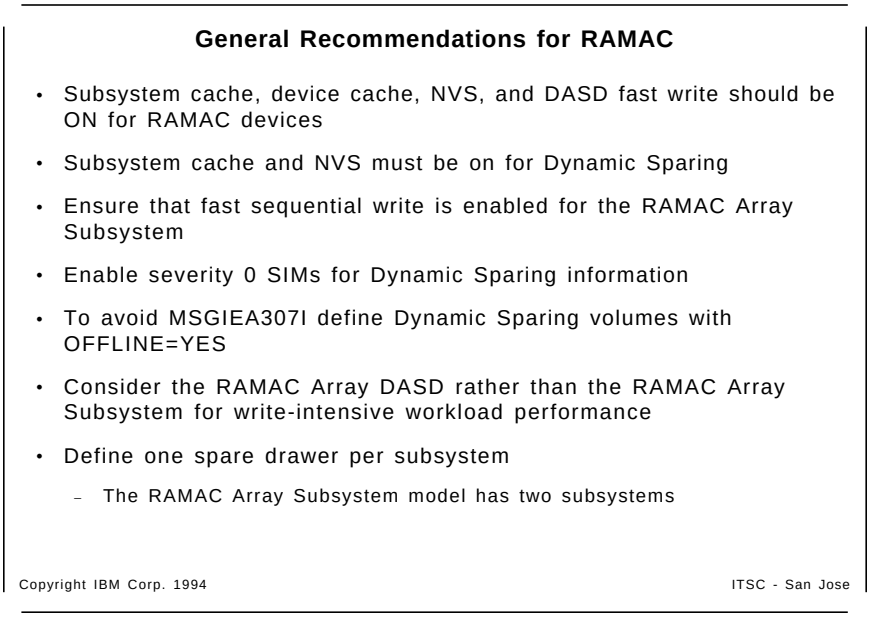

Softw-75 General Recommendations for RAMAC

The foil lists several general recommendations for implementing RAMAC Array DASD and RAMAC Array Subsystem configurations.

- **Subsystem cache, device cache, NVS, and DFW should be on for RAMAC devices**

  For RAMAC Array DASD attached to a 3990 storage control, subsystem cache, volume caching, nonvolatile storage (NVS), and DFW should be on to get the best performance. The RAMAC drawers are designed to work with subsystem cache.

- **Subsystem cache and NVS must be on for Dynamic Sparing**

  Subsystem cache and NVS are required during a Dynamic Sparing operation.

- **Ensure that fast sequential write is enabled for the RAMAC Arra Subsystem**

  The RAMAC Array Subsystem fast sequential write function provides special support to handle sequential writes at high speed. This optional function uses some cache (8 bytes per track) and is enabled by default. Disabling this function causes applications that require sequential writes to experience substantially lower throughput and increased response time.

  Fast sequential write should be enabled for RAMAC 2 configurations.

- **Enable severity 0 SIMs for Dynamic Sparing information**

  The SIMs for Dynamic Sparing are considered service SIMs; therefore, the VPD setting at the 3990 storage control and RAMAC Array DASD or at the RAMAC Array Subsystem must be “00” in order for the messages to be displayed at the console. This setting presents all service SIMs for all attached devices to the console.

- **To avoid MSGIEA307I define Dynamic Sparing volumes with OFFLINE=YES**

  Message IEA307I is issued for the addresses associated with the drawers defined for Dynamic Sparing. This message is also received for any attempt to vary a Dynamic Sparing volume online. To bypass this message at IPL,

you can specify that the volumes come up offline with the OFFLINE=YES option in the MVSCP. Alternatively, apply the optional PTF associated with APAR OW09398 or OW09400 for DFP or DFSMS/MVS, respectively.

- **Consider the RAMAC Array DASD rather than the RAMAC Array Subsystem for write-intensive workload performance**

  In the RAMAC Array Subsystem the drawer cache acts as NVS. Channel end and device end are returned when the data is in two places—the storage control cache and the drawer cache. In an IBM 3990 storage control the two copies of the data are placed in the cache and the NVS of the storage control. Storing data in the NVS of an IBM 3990 model 6 is faster than propagating the data down to the drawer cache first and then returning device end. Therefore write-intensive workloads will have better performance on a RAMAC Array DASD attached to an IBM 3990 model 6.

- **Define one spare drawer per subsystem**

  To enable nondisruptive drawer maintenance for all repair situations, define one spare drawer per subsystem. This arrangement supports Dynamic Sparing, which is used to provide continuous access to data during all drawer repairs. Remember that the RAMAC Array Subsystem model 3 supports two logical subsystems in a single rack. Therefore one spare drawer should be defined for each subsystem.

VM PTF Recommendations

Product	APAR	PTF	Description
ICKDSF Release 16	PN62330	UN67694 old inst. UN67695 VMSES/E	CPVOLUME format performance improvement
EREP 3.5	VM58564	UV90731	RAMAC Array Family support
EREP 3.5	VM59009	UV58809	9394 and 9395 support
VM/ESA 1.2.2 VM/ESA 1.2.1 VM/ESA 1.2.0	VM58671 VM58942 VM58942	UM26755 UM26891 UM26890	Enhancement for RAMAC - actual device type in SIMs

Softw-76 VM PTF Recommendations

Foil Softw-76 shows some recommended APAR and PTF numbers for VM.

The PTF for APAR VM58671 provides the recognition of the real (not emulated) device type for SIMs. This PTF might already be installed in your system to recognize a RAMAC device type. The PTF includes code that recognizes a RAMAC 2 drawer.

The other PTFs provide RAMAC device support for EREP and other improvements for ICKDSF.

VM PTF Recommendations for RAMAC Array DASD

Product	APAR	PTF	Description
VSE/VSAM for VM 2.2.0 VSE/VSAM for VM 2.1.0	VM58884 VM58884	UV90734 UV90733	R0 performance fix for 9391 DASD
VM/ESA 1.2.1 VM/ESA 1.2.0 VM/ESA 1.1.1	VM57925 VM57925 VM57925	UM25750 UM25649 UM25748	Recognition of 3990-6 at IPL time
VM/ESA 1.2.2 VM/ESA 1.2.1	VM59119 VM59119	UM27058 UM27057	R0 performance fix for CMS FORMAT of 9391 DASD
VM/ESA 1.2.2 VM/ESA 1.2.1	VM59200 VM59200	OPEN OPEN	R0 performance fix for TDSK clearing on 9391 DASD (includes RAMAC Array Subsystem MONITOR data).

 ITSC - San Jose

Softw-77 VM PTF Recommendations for RAMAC Array DASD

Foil Softw-77 shows some recommended APAR and PTF numbers for a RAMAC Array DASD in a VM environment. PTFs for the RAMAC Array DASD include performance fixes for VSE/VSAM for VM and VM/ESA and recognition of the 3990-6.

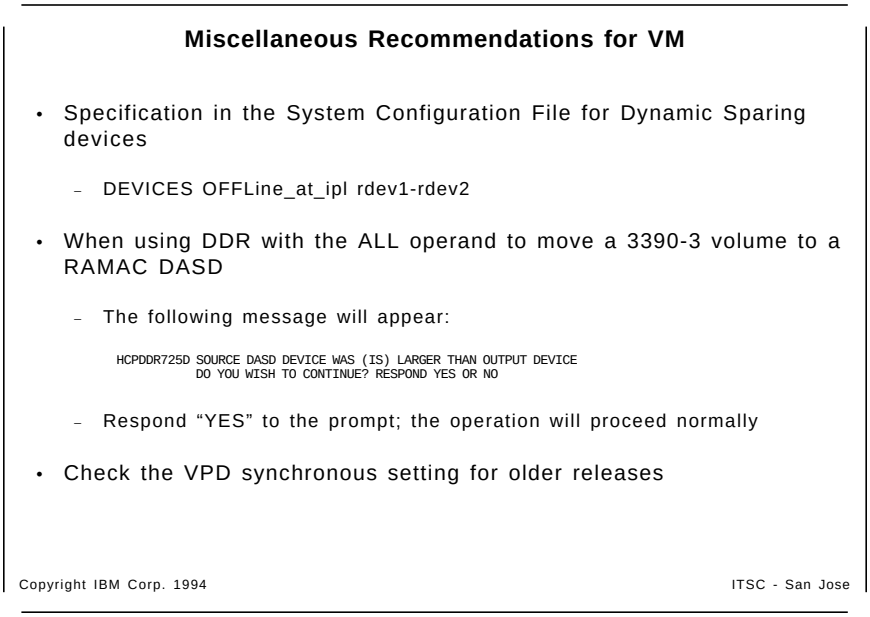

Softw-78 Miscellaneous Recommendations for VM

The foil lists several recommendations when using RAMAC in a VM environment.

- **Specification in the System Configuration File for Dynamic Sparing devices**

  Spare drawers in a RAMAC Array DASD can produce HCP... messages at IPL time, or when a spare volume is varied online during normal processing, or when there is I/O to a spare volume. To avoid this problem at IPL, specify in the System Configuration File that these volumes are offline.

- **When using DDR with the ALL operand to move a 3390-3 volume to a RAMAC DASD**

  When using DDR with the ALL operand to move a 3390-3 volume to a RAMAC Array DASD, a message will appear indicating that the source device is larger than the output device. The DDR utility is reacting to the fact that the alternate tracks on the 3390-3 are not present on the RAMAC Array DASD output device. Ignore this message and respond with “YES” to the prompt to continue. To avoid this message, specify the actual cylinder extents.

- **Check the VPD synchronous setting for older releases**

  For a RAMAC Array Subsystem the synchronous setting must be changed for some older VM releases from the default setting (which is nonsynchronous) to synchronous-1. For details see foil Common-89.

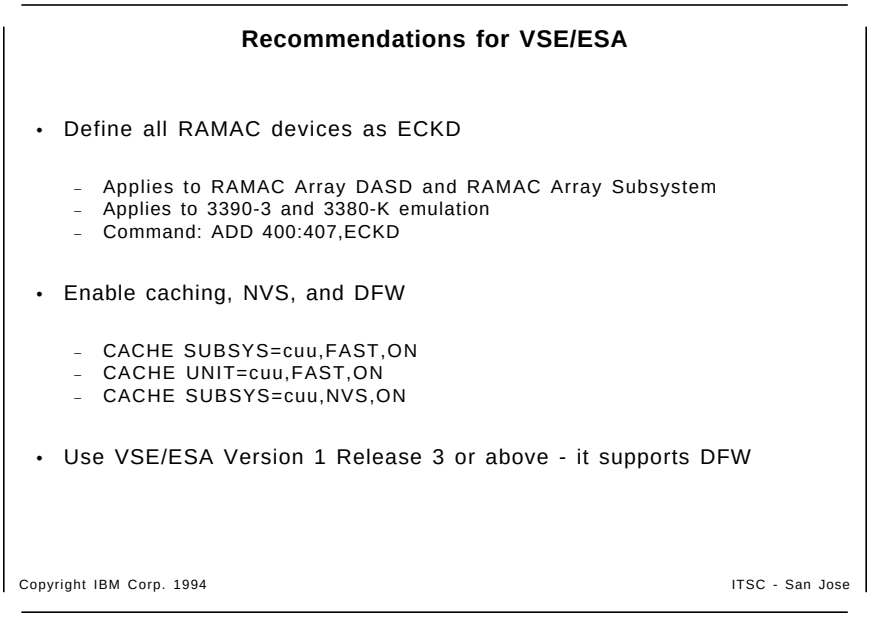

Softw-79 Recommendations for VSE/ESA

The foil lists several recommendations when using RAMAC in a VSE/ESA environment.

- **Define all RAMAC devices as ECKD**

  The ADD command to define devices to VSE has an ECKD parameter for devices that support ECKD channel command words (CCWs). For RAMAC Array DASD or a RAMAC Array Subsystem specify ECKD for 3390-3 and 3380-K emulated devices for VSE releases that support ECKD I/O (VSE/ESA 1.2 and above). ECKD channel programs provide the best performance with RAMAC devices.

- **Enable caching, NVS, and DFW**

  RAMAC is designed to work in collaboration with subsystem cache. Therefore you will get the best performance with the cache functions, particularly DASD fast write (DFW), enabled.

- **Use VSE/ESA Version 1 Release 3 or above - it supports DFW**

  VSE/ESA Version 1 Release 3 provides support for DFW. You should be at this VSE level or above to get the real benefit of RAMAC.

VSE I/O Hints

Product	Sequential	RegDataForm	RecordCache
VSE/SP 4.1	-	-	-
VSE/ESA 1.1	x	-	-
VSE/ESA 1.2	x	-	-
VSE/ESA 1.3	x	VSAM only	VSAM
VSE/ESA 2.1	x	x	VSAM
Beneficial			
3990-3/RAMAC	yes	no	no
3990-6/RAMAC	yes	yes	yes
RAMAC Subs.	yes	yes	no

- VSE CKD/ECKD conversion routines cannot set individual caching bits
- VSE/ESA 2.1 has broadest scope of RAMAC exploitation

 ITSC - San Jose

Softw-80 VSE I/O Hints

The RAMAC drawer can reduce the write penalty if the software provides some hints about the nature of the I/O. The *sequential* flag tells the drawer that sequential I/O occurs. In this case the drawer does not have to read the old data and the old parity to do the write. Three tracks and the parity can be destaged together.

VSE/SP does not provide any software hints for the drawer. Starting with VSE/ESA 1.1 the sequential bit flag is set.

The regular data format flag is required for the subsystem to use record caching. VSE/ESA 1.3 is the first release that supports this regular data format flag for VSAM I/O.

- **VSE CKD/ECKD conversion routines cannot set individual caching bits**

  VSE can convert CKD channel programs to ECKD channel programs. The conversion routines, however, do not set the flags that indicate sequential I/O or a regular record format.

- **VSE/ESA 2.1 has broadest scope of RAMAC exploitation**

  The higher the VSE release the better the cache exploitation and thus the better the RAMAC performance.

RAMAC Array Subsystem ADDs in VSE

	3990-2/3380-K Emul.		3990-2/3390-3 Emul.
VSE Release	**Parallel**	**ESCON**	**Parall.+ESCON**
VSE/SP 4.1	3380 sync2	not supp.	not supp.
VSE/ESA 1.1	3380 sync2	not supp.	not supp.
VSE/ESA 1.2.0	3380 sync2	not supp.	not supp.
VSE/ESA 1.2.0 + PTF for APAR DY41099	ECKD	ECKD	ECKD
VSE/ESA 1.2.1-3	ECKD	ECKD	ECKD
VSE/ESA 1.3	ECKD	ECKD	ECKD
VSE/ESA 2.1	ECKD	ECKD	ECKD

- Always use *ADD cuu,ECKD* in VSE/ESA.
- Do not use *ADD cuu,3380,EML* in VSE/ESA. It forces CKD channel programs.

 ITSC - San Jose

Softw-81 RAMAC Array Subsystem ADDs in VSE

Foil Softw-81 lists the supported device types for RAMAC devices in a VSE environment.

- **Always use** *ADD cuu,ECKD* **in VSE/ESA**

  The recommended device type for RAMAC devices in VSE/ESA is always ECKD. This holds for the 3390-3 emulation as well as the 3380-K emulation.

- **Do not use** *ADD cuu,3380,EML* **in VSE/ESA. It forces CKD channel programs**

  Do not use the *3380,EML* parameter except for a temporary compatibility test for a vendor program. The *3380,EML* parameter combination forces CKD channel programs, which provide lower performance than ECKD channel programs. The RAMAC 2 LIC for the RAMAC Array Subsystem provides improved CKD performance.

VSE PTF Recommendations

Product	APAR	PTF	Description
VSE/ESA 1.3	DY43335	UD49324 UD49317 UD49325 UD49332	Performance improvement for 9391 DASD
VSE/ESA 1.3	DY43414	UD49333	Sequential destaging improvement
VSE/ESA	DY43312		Fix for various problems; enhance ECKD performance
ICKDSF Rel. 16	PN62330	UN67699	Fix for various problems
EREP 3.5	DY43462	UD49343 UD90366	Support for 9394 and 9395

ITSC - San Jose

Softw-82 VSE PTF Recommendations

Foil Softw-82 shows some recommended APAR and PTF numbers for VSE. It is not necessary to install any PTFs for RAMAC, but some PTFs are recommended for best RAMAC performance.

Ordering Considerations

 ITSC - San Jose

Ordering Considerations for RAMAC Array Subsystem

Description	Machine	Model	Feature	
Array Controller	9394			
Dual Cluster Controller		001		
Single Line Cord, Single Phase Power				
64 MB cache			0601	
256 MB cache			0604	
4 ESCON channels			3001	
4 parallel channels			3007	
4 parallel to 4 ESCON field upgrade			3401	
Field reformat to 3990 2/3380 K emulation			7280	(B13 drawers only)
Field reformat to 3990 2/3390 3 emulation			7293	
Power and cable completion group (>8 drawers)			4111	
Dual Cluster Controller		002		
Dual line cord, 3 phase power				
64 MB cache			0601	
256 MB cache			0602	
1024 MB cache			0644	
4 ESCON channels			3001	
4 ESCON + 4 parallel channels			3004	
8 parallel channels			3008	
8 parallel to 4 ESCON field upgrade			3801	
8 parallel to 4 ESCON + 4 parallel field upgrade			3804	
Field reformat to 3990 2/3380 K emulation			7280	(B13 drawers only)
Field reformat to 3990 2/3390 3 emulation			7293	
Power and cable completion group (>8 drawers)			4121	(1)
Power and cable completion group (>8 drawers)			4142	(2)

 ITSC - San Jose

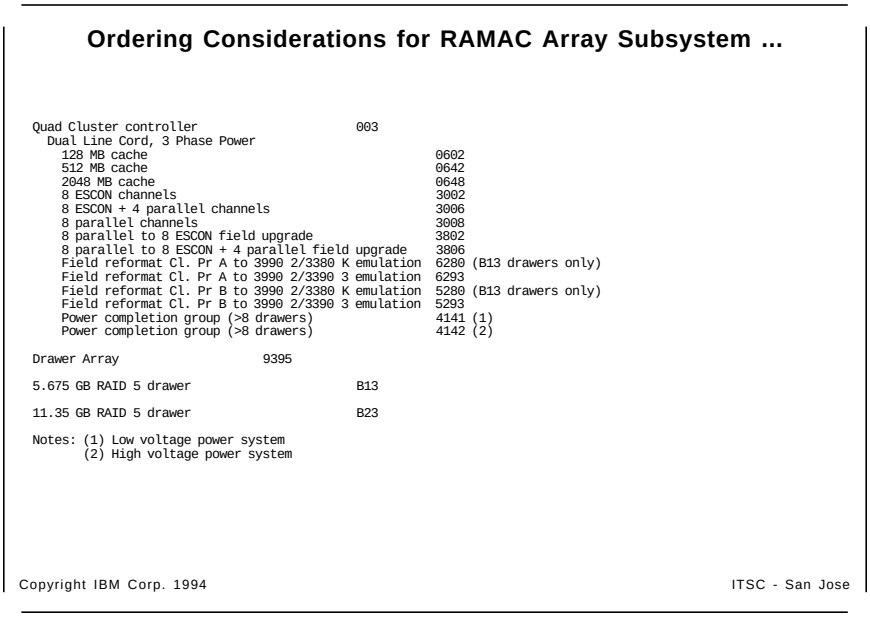

Order-84 Ordering Considerations for RAMAC Array Subsystem

Foil Order-84 lists the feature codes available for a RAMAC Array Subsystem. Note the new field reformat feature codes. A RAMAC 2 drawer has the model code B23.

RAMAC and RAMAC 2 Array Subsystem Considerations

 ITSC - San Jose

RAMAC and RAMAC 2 Array Subsystem Considerations

- VPD options
  - Synchronous settings
  - Fast sequential write
  - Long Busy
- Enhanced format write

 ITSC - San Jose

Common-87 RAMAC and RAMAC 2 Array Subsystem Considerations

The next four foils discuss topics that are common to both RAMAC and RAMAC 2 drawers attached to the RAMAC Array Subsystem. These topics relate to the control unit function of the RAMAC Array Subsystem. Although they are independent of drawer type, they are included here as a refresher.

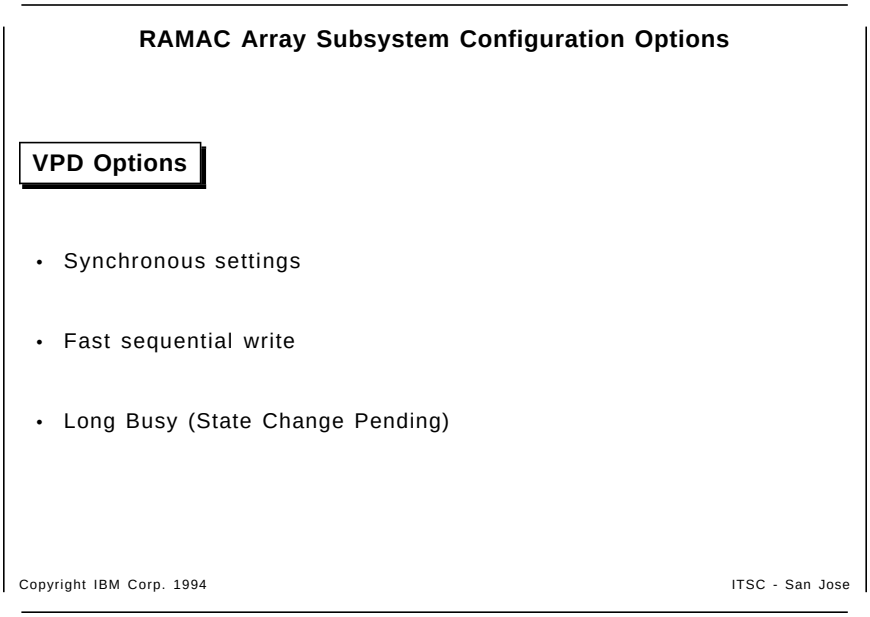

Common-88 RAMAC Array Subsystem Configuration Options

The configuration options listed on the foil are normally set by the service representative.

- **Synchronous settings**

  Although the RAMAC Array Subsystem normally operates in nonsynchronous ECKD mode, there might be compatibility problems with older software. Therefore, for compatibility and performance reasons, there are two other synchronous settings.

- **Fast sequential write**

  The performance of sequential writes can be improved by keeping all R0 records in cache. R0 records occupy some cache storage, however, so the benefit of fast sequential write must be weighed against the effect of a smaller usable cache.

- **Long Busy (State Change Pending)**

  For subsystem events that require more than 15 seconds, the subsystems use Long Busy (State Change Pending). Subsystem actions that require longer than 15 seconds include some internal recovery actions and nondisruptive engineering change (EC) installation.

  One exception to DASD′s “use Long Busy if greater than 15 seconds” rule is concurrent installation of an EC to the RAMAC Array Subsystem itself (not to the drawers). In this case the subsystem does not have logic to set Long Busy, so it returns Control Unit Busy to I/Os received during the time required for the EC activation process. RAMAC 2 Array Subsystem EC activation can take up to 55 seconds.

  Another reason some installations may want a longer MIH time for RAMAC Array Subsystems is that the subsystems can be used with some older operating system releases that do not support State Change Pending (Long Busy). The RAMAC Array Subsystem provides a VPD bit that can be set to prevent the subsystem from presenting Long Busy. Systems that run with VPD set to not present Long Busy will receive a device busy signal during

subsystem events that would have returned Long Busy, and this device busy could cause an MIH action. The RAMAC Array Subsystem events that generate Long Busy are activation of RAMAC drawer microcode and RAMAC drawer cache manager recovery.

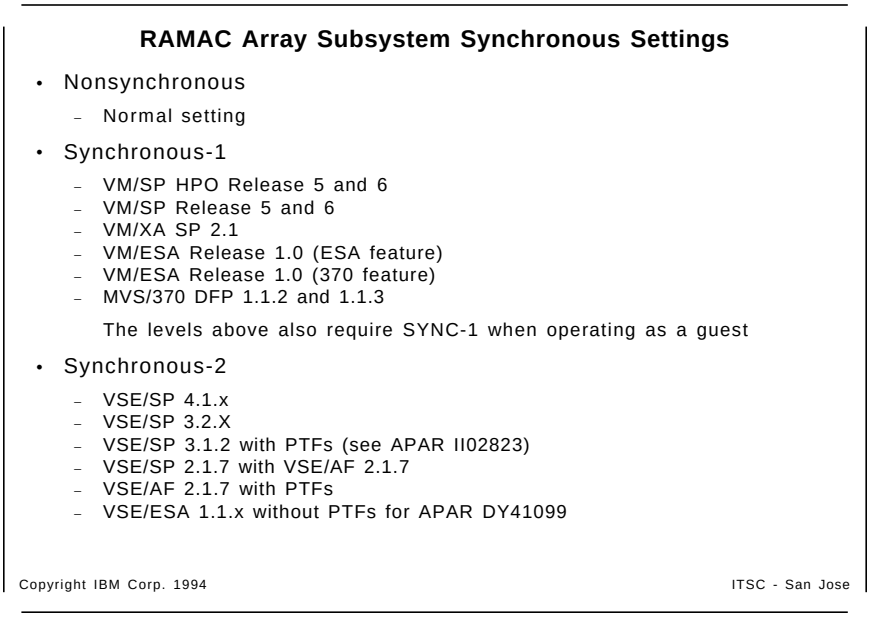

Common-89 RAMAC Array Subsystem Synchronous Settings

The synchronous settings available in the RAMAC Array Subsystem are normally set at installation time by the customer engineer and can be changed at a later time. A change in the synchronous VPD setting requires an IML of the RAMAC Array Subsystem.

- **Nonsynchronous**

  Nonsynchronous is the normal and default setting unless changed by the service representative. Nonsynchronous provides the best performance.

- **Synchronous-1**

  Synchronous-1 is required for older VM releases. The foil lists the operating systems that require a synchronous-1 setting.

- **Synchronous-2**

  Synchronous-2 affects operations on parallel channels only. For a RAMAC Array Subsystem with both ESCON and parallel channels with synchronous-2 set in the VPD, the ESCON channels will still operate in synchronous-1 mode. The synchronous-2 setting applies to older VSE releases as listed on the foil.

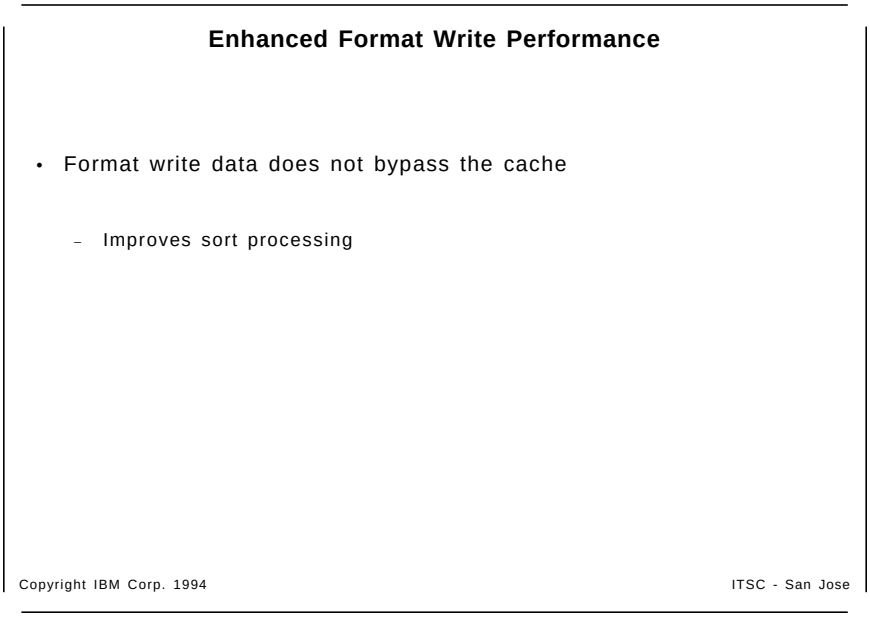

Common-90 Enhanced Format Write Performance

Format write performance has been enhanced for the RAMAC Array Subsystem.

- **Format write data does not bypass the cache**

  In the previous implementation of the RAMAC Array Subsystem, format write data bypassed the cache. Some workloads, such as sort, do format writes and read the data soon afterward. To improve the performance of such workloads, format write data no longer bypasses the cache. A copy of the data is in the cache and available for a subsequent read operation.

This improvement in format write performance comes with the LIC that supports RAMAC 2 drawers. It also applies to RAMAC drawers.

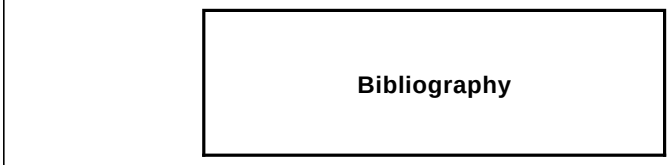

RAMAC Array DASD

- *IBM RAMAC Array DASD Introduction*, ***GC26-7012***
- *Using the IBM RAMAC Array DASD in an MVS, VM or VSE Environment*, ***GC26-7013***
- *IBM RAMAC Array Family*, ***GG24-2509***
- *IBM 3990 Storage Control Introduction*, ***GA32-0098***
- *IBM 3990 Operations and Recovery Guide*, ***GA32-0253***
- *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide*, ***GA32-0100***
- *Extended Count Key Data and Nonsynchronous I/O*, ***GG24-3751***
- *IBM 3990 Storage Control ESCON Features Presentation Guide*, ***GG24-3803***
- *ICKDSF R16 User's Guide and Reference*, ***GC35-0033***

Biblio-92 RAMAC Array DASD

Foil Biblio-92 lists those manuals where more information about the RAMAC Array DASD, the IBM 3990 storage control, and related products can be found.

RAMAC Array Subsystem

- *IBM RAMAC Array Subsystem Introduction*, ***GC26-7004***
- *Using the IBM RAMAC Array Subsystem in an MVS, VM or VSE Environment*, ***GC26-7005***
- *IBM RAMAC Array Family*, ***GG24-2509***
- *IBM RAMAC Array Subsystem Reference*, ***GC26-7006***
- *RAMAC Array Subsystem Service Guide*, ***SY27-7507***
- *Extended Count Key Data and Nonsynchronous I/O*, ***GG24-3751***
- *3390 Direct Access Storage Migration Guide*, ***GG24-3373***
- *ICKDSF R16 User's Guide and Reference*, ***GC35-0033***

 ITSC - San Jose

Biblio-93 RAMAC Array Subsystem

Foil Biblio-92 lists those manuals where more information about the RAMAC Array Subsystem and related topics can be found.

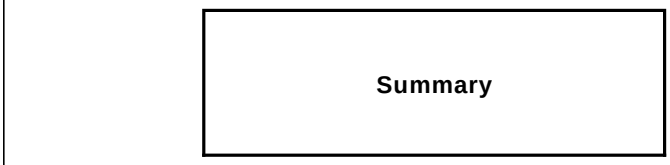

 ITSC - San Jose

The IBM RAMAC Array Family - Summary

- High availability
  - High reliability - disk and control unit technology
  - Fault-tolerant design
  - Enhanced RAID-5
- Cost effectiveness
  - High-capacity drawers
  - Investment protection
  - Improved environmentals
- High performance
  - Multilevel cache
  - Distributed drawer intelligence
  - Improved algorithms
  - Faster hardware
- Continuous operations
  - Nondisruptive upgrades, installation, and maintenance

 ITSC - San Jose

Summary-95 The IBM RAMAC Array Family - Summary

- **High availability**

  The RAMAC Array DASD and the RAMAC Array Subsystem are designed for 100% availability. This high availability is enabled through the use of redundant systems, RAID-5 design, and state of the art technologies. Redundancy is implemented in both the drawers and the storage control units. Thin-film disk, MR heads, and PMRL channels are the building blocks for disk drives that have the highest MTBF rating on the market for 3.5-in. disk drives. In case of a disk drive failure, Dynamic Sparing saves data to a spare drawer while maintaining concurrent access to the data, thus making a RAID-5 implementation even more reliable. The drawer LIC is enhanced to deal with difficult situations such as when more than one error occurs in the drawer.

- **Cost effectiveness**

  RAMAC configurations with RAMAC 2 drawers provide twice the storage capacity of RAMAC configurations with RAMAC drawers. This increase in storage capacity reduces costs and offers an excellent price-to-performance ratio.

  At the same time RAMAC 2 preserves the investment in current equipment. RAMAC 2 drawers can be mixed with RAMAC drawers in the same rack. The RAMAC Array DASD rack and the RAMAC Array Subsystem rack and control unit functions have not changed. A field upgrade option is available to replace RAMAC drawers with RAMAC 2 drawers.

  With the capacity doubled on the same footprint while the power and cooling requirements do not change, you get an improved kW per GB ratio. This improvement lowers operating costs and saves money.

- **High performance**

  RAMAC delivers high performance and high availability. High performance results from such factors as multilevel cache implementation, DFW, and the distribution of RAID-5 tasks to many independently operating RAID-5 drawers.

  To deliver equal or even better performance for the high-capacity RAMAC 2 drawers some enhancements have been applied to the LIC and hardware. The drawer LIC can now even better recognize certain patterns in the data and considers these patterns in its destaging policy to avoid the RAID-5 write penalty.

  The drawer hardware has faster chips for a higher internal bandwidth to cope with a higher data rate of a double-capacity drawer. The parity calculation, a crucial element of RAID-5 processing, is considerably faster because of the improved cycle time of the enhanced XOR chip.

- **Continuous operations**

  Migration to the new RAMAC 2 drawers with their attractive price to performance ratio is very easy. The new drawers are totally transparent to the software. The installation of new RAMAC 2 drawers is nondisruptive for the RAMAC Array Subsystem and is nondisruptive for the RAMAC Array DASD after the initial installation of the RAMAC 2 supporting LIC. RAMAC and RAMAC 2 drawers can be mixed in the same rack. A field upgrade option is available to replace nondisruptively (with Dynamic Sparing) RAMAC drawers with RAMAC 2 drawers.

Index

P

V

X

ITSO Technical Bulletin Evaluation RED000

International Technical Support Organization
IBM RAMAC Array Family Additions
(RAMAC 2)
October 1995

Publication No. SG24-4563-00

Your feedback is very important to help us maintain the quality of ITSO Bulletins. **Please fill out this questionnaire and return it using one of the following methods:**

- Mail it to the address on the back (postage paid in U.S. only)
- Give it to an IBM marketing representative for mailing
- Fax it to: Your International Access Code + 1 914 432 8246
- Send a note to REDBOOK@VNET.IBM.COM

Please rate on a scale of 1 to 5 the subjects below.
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction	____		
Organization of the book	____	Grammar/punctuation/spelling	____
Accuracy of the information	____	Ease of reading and understanding	____
Relevance of the information	____	Ease of finding information	____
Completeness of the information	____	Level of technical detail	____
Value of illustrations	____	Print quality	____

Please answer the following questions:

a) If you are an employee of IBM or its subsidiaries:

Do you provide billable services for 20% or more of your time? Yes____ No____

Are you in a Services Organization? Yes____ No____

b) Are you working in the USA? Yes____ No____

c) Was the Bulletin published in time for your needs? Yes____ No____

d) Did this Bulletin meet your needs? Yes____ No____

If no, please explain:

_____

_____

What other topics would you like to see in this Bulletin?

_____

_____

What other Technical Bulletins would you like to see published?

_____

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

Name	Address
Company or Organization	
Phone No.	

Fold and Tape **Please do not staple** Fold and Tape

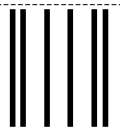

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

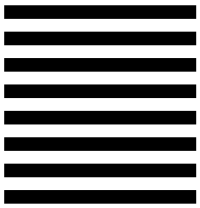

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM International Technical Support Organization
Department 471, Building 070B
5600 COTTLE ROAD
SAN JOSE CA
USA 95193-0001

Fold and Tape **Please do not staple** Fold and Tape

Printed in U.S.A.