AS/400e

Hierarchical Storage Management

SC41-5351-01

AS/400e

Hierarchical Storage Management

SC41-5351-01

Note

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

Second Edition (May 1999)

This edition applies to version 4, release 4, modification 0 of IBM Operating System/400 (product number 5769-SS1) and to all subsequent releases and modifications until otherwise indicated in new editions. This edition applies only to reduced instruction set computer (RISC) systems.

© Copyright International Business Machines Corporation 1998, 1999. All rights reserved.
Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

© Copyright IBM Corp. 1998, 1999

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described 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:

IBM Director of Licensing
IBM Corporation
500 Columbus Avenue
Thornwood, NY 10594
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Software Interoperability Coordinator
3605 Highway 52 N
Rochester, MN 55901-7829
U.S.A.

© Copyright IBM Corp. 1998, 1999

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement or any equivalent agreement between us.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

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

Application System/400
AS/400
AS/400e
IBM
Magstar
Operating System/400

OS/400
SQL
400

Microsoft, Windows, Windows NT, and the Windows 95 logo are registered trademarks of Microsoft Corporation.

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

Other company, product, and service names may be trademarks or service marks of others.

About Hierarchical Storage Management (SC41-5351-01)

The Backup, Recovery, and Media Services (BRMS) product has been restructured to provide three separately installed, keyed compliant components consisting of a standard product, a Network feature, and an Advanced Functions feature. The Advanced Functions feature contains all of the heirarchical storage management (HSM) components, including the capability to perform archive, migration, and dynamic retrieval operations. This book provides instruction on how to set up, implement, and use the hierarchical storage management components (migration, archive, and dynamic retrieval) in your business environment. This book assumes the use of the standard Backup Recovery and Media Services for AS/400 (BRMS) product with HSM, though you need not have BRMS installed on your AS/400 to implement HSM in your business environment.

Who should read this book

This book helps the system administrators and application programmers to design and implement a hierarchical storage management solution to meet their business needs. This can include the introduction of new applications or customization of existing programs.

The book is also intended to help IBM technical specialists in the AS/400 systems management area to understand and assist with the implementation of hierarchical storage management with BRMS.

An understanding of the overview of BRMS functions and capabilities and the AS/400 architecture and file construction are assumed. Some basic knowledge of application design would be helpful.

AS/400 Operations Navigator

AS/400 Operations Navigator is a powerful graphical interface for Windows 95/NT clients. With AS/400 Operations Navigator, you can use your Windows 95/NT skills to manage and administer your AS/400 systems.

- You can work with basic operations (messages, printer output, and printers), job management, system configuration, network administration, security, users and groups, database administration, file systems, and multimedia.
- You can schedule regular system backups, work with Interprocess Communication through application development, and manage multiple AS/400 systems through a central system by using Management Central. You can also customize the amount of Operations Navigator function that a user or user group can use through application administration.
- You can create a shortcut to any item in the explorer view of Operations Navigator. For example, you can create a shortcut either to **Basic Operations** or to the items that are listed under **Basic Operations** (**Messages**, **Printer Output**, and **Printers**). You can even create a shortcut to an individual printer or use a shortcut as a fast way to open the item.

Figure 1 on page xii shows an example of the Operations Navigator display:

© Copyright IBM Corp. 1998, 1999

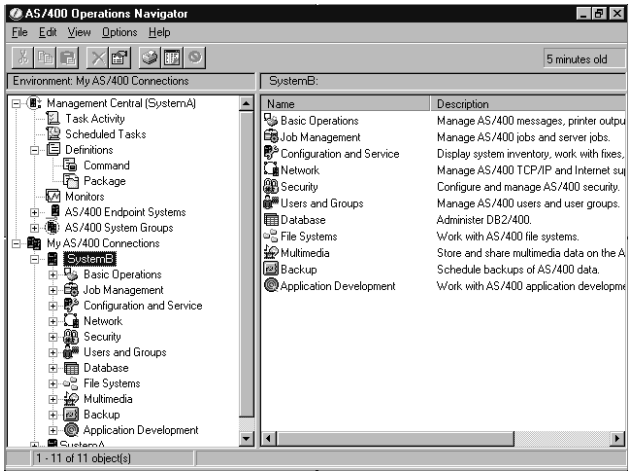

Figure 1. AS/400 Operations Navigator Display

IBM recommends that you use this new interface. It has online help to guide you. While we develop this interface, you will still need to use either of the following to do some of your tasks:

- Graphical Access (which provides a graphical interface to AS/400 screens). Graphical Access is part of the base Client Access.
- A traditional emulator such as PC5250.

Installing Operations Navigator subcomponents

AS/400 Operations Navigator is packaged as separately installable subcomponents. If you are upgrading from a previous release of AS/400 Operations Navigator, only those subcomponents that correspond to the function that is contained in the previous release will be installed. If you are installing for the first time and you use the **Typical** or **Minimum** installation options, the following options are installed by default:

- Operations Navigator base support
- Basic operations (messages, printer output, and printers)

To install additional AS/400 Operations Navigator subcomponents, either use the **Custom** installation option or use **selective setup** to add subcomponents after Operations Navigator has been installed:

1. Display the list of currently installed subcomponents in the **Component Selection** window of **Custom** installation or selective setup.
2. Select AS/400 Operations Navigator and click **Details**.
3. Select any additional subcomponents that you want to install and continue with **Custom** installation or selective setup.

Note: To use AS/400 Operations Navigator, you must have Client Access installed on your Windows 95/NT PC and have an AS/400 connection from that PC. For help in connecting your Windows 95/NT PC to your AS/400 system, consult *Client Access for Windows 95/NT - Setup*, SC41-3512-05 .

Accessing AS/400 Operations Navigator

To access Operations Navigator after you install Client Access and create an AS/400 connection, do the following:

1. Double-click the **Client Access** folder on your desktop.
2. Double-click the **Operations Navigator** icon to open Operations Navigator. You can also drag the icon to your desktop for even quicker access.

Prerequisite and related information

Use the AS/400 Information Center as a starting point for your AS/400 information needs. It is available in either of the following ways:

- The Internet at this uniform resource locator (URL) address:
  `http://publib.boulder.ibm.com/html/as400/infocenter.html`
- On CD-ROM: *AS/400e Information Center*, SK3T-2027-01 .

The AS/400 Information Center contains browsable information on important topics such as Java, program temporary fixes (PTFs), and Internet security. It also contains hypertext links to related topics, including Internet links to Web sites such as the AS/400 Technical Studio, the AS/400 Softcopy Library, and the AS/400 home page.

How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information. If you have any comments about this book or any other AS/400 documentation, fill out the readers' comment form at the back of this book.

- If you prefer to send comments by mail, use the readers' comment form with the address that is printed on the back. If you are mailing a readers' comment form from a country other than the United States, you can give the form to the local IBM branch office or IBM representative for postage-paid mailing.
- If you prefer to send comments by FAX, use either of the following numbers:
  - United States and Canada: 1-800-937-3430
  - Other countries: 1-507-253-5192
- If you prefer to send comments electronically, use this network ID:
  - IBMMAIL, to IBMMAIL(USIB56RZ)
  - RCHCLERK@us.ibm.com

Be sure to include the following:

- The name of the book.
- The publication number of the book.
- The page number or topic to which your comment applies.

Summary of Changes

This section provides a summary of the enhancements made to the hierachical storage management (HSM) components, and to the HSM manual, for the V4R4 release. These changes are indicated within the manual by the | symbol, which is called a change bar. Where specified, you must have BRMS installed to use the enhanced function or option.

Software Enhancements

Following are brief descriptions of the software enhancements recently made to HSM components.

Frequency of Use Count

The *Days used (per month) less than* parameter allows you to indicate whether the system should include an object in an archive operation based on the object's frequency of use. You set the frequency of use count based on the number of times an object is used in a month. You can find this parameter on the BRMS archive policy and Change Archive Control Group Attributes displays.

Support for the Dynamic Retrieval of *QHST Files

You can now use the *QHST special value for archive operations. In addition, you can dynamically retrieve *QHST files.

The Addition of Parallel Device Support for Archive and Retrieve Operations

You can use parallel device support if you have BRMS installed on your AS/400. You can use parallel device support to process a single large library or object in parallel by using multiple devices. In a parallel operation, BRMS splits the library or object across the devices that you specify. Parallel support can reduce your archive and retrieve window. You can find this parameter on the archive and retrieve policy displays, on the Change Archive Control Group Attributes display, and on the Set Retrieve using BRM (SETRTVBRM) command.

Book Enhancements

Following are brief descriptions of the new additions to the HSM manual.

Chapter 5: Establishing Security for Hierarchical Storage Management (HSM) Components

BRMS now provides the following security options:

- The functional usage model allows you to customize security access by user for selected functions and functional components. You can give some users access to use and view selected functions, and others to change specific items such as control groups and archive lists.

© Copyright IBM Corp. 1998, 1999

- The Set User Usage for BRM (SETUSRBRM) command allows you to specify usage by job category. The categories are *ADMIN and *OPERATOR. Each category offers access at different levels, based on the job responsibilities usually assigned to people in these roles.
- The BRMS log now contains a new area that is called *SEC. *SEC contains security-related messages such as those sent when a user does not have access to a function or a component.

Chapter 8: Balancing High and Low-Used Data within an Archive ASP

This chapter provides information on the various methods you can use to balance data within an ASP. These methods include HSM balancing, capacity balancing, and usage balancing in an archive ASP. This chapter also provides information on how to implement each method.

Appendix A: Summary of Reports

This appendix provides brief descriptions of the HSM-related reports you can generate if you have BRMS installed on your AS/400.

Appendix B: Customizing the User Interface for Dynamic Retrieval

This appendix provides information on how to use the BRMS Object Retrieval Exit Program to customize the messages that are sent to users when dynamic retrieval occurs. It also provides information on how to design the program, and how to register it with BRMS.

Chapter 1. Introduction to Hierarchical Storage Management

Hierarchical storage management (HSM) provides an automatic way of managing and distributing data between the different storage layers to meet user needs for accessing data while minimizing the overall cost. The concept of hierarchical storage management involves the placement of data items in such a way as to minimize the cost of storing your data, while maximizing the accessibility.

This book describes the high-level principles of hierarchical storage management and relates the system architecture model of the AS/400 to the hierarchical storage management model. It describes migration of data to and from the different disk technologies using user ASPs. The book also describes the changes to existing archival and dynamic retrieval functions (disk to tape, disk to disk, and vice versa), and the benefits you can gain from these changes. In addition, this book provides information on the type of objects that you can consider for migration, archival and dynamic retrieval. You can also find information about the kinds of issues you should consider before the design and implementation of your hierarchical storage management solution.

This book emphasizes the use of HSM with Backup Recovery and Media Services (BRMS). However, you need not have BRMS installed on your system to implement HSM in your business environment. If you are a non-BRMS customer, you can either develop your own HSM application, or use the CL commands that are available to balance data within an archive ASP. See *A Practical Approach to Hierarchical Storage Management on the AS/400: SG24-4450-01* for information on how to develop your own HSM application.

Hierarchical Storage Management Benefits

The principal features that hierarchical storage management provides include the following:

- **Reduced total cost of storage**

  Data that is accessed less frequently resides on lower cost storage.
- **Performance benefits**

  Data residing on lower cost storage, frees up high performance storage, possibly increasing overall system performance.
- **Tape automation reduces operating costs**

  If you use tape automation and BRMS, you can design a solution to reduce operational costs due to less manual intervention as most tape operations are unattended.
- **Improved system availability**

  You do not need to save historical data repeatedly.
- **Transparent to applications**

  No application changes required when data has been moved to a different storage technology.
- **New applications that are enabled**

  Printed information or information in file drawers can now be stored inexpensively on the system.
- **Customized hierarchical storage management applications**

© Copyright IBM Corp. 1998, 1999

You can write customized applications such as BRMS, with the hierarchical storage management APIs that are provided by OS/400. See *A Practical Approach to Hierarchical Storage Management on the AS/400: SG24-4450-01* for more information on this topic.

These features, in turn, give rise to significant cost savings and revenue opportunities to your business including:

- **Reduction in total disk cost and tape hardware cost** derived from the more effective use of expensive disk as opposed to less expensive tape storage
- **Increased revenue** to service delivery, derived from increased system availability as save times and restore times will be reduced
- **Reduced cost of labor** derived from reduction in both skills that are needed and time that is allocated to tasks that are involved in the management and cleanup of system disk space
- **Increased customer satisfaction** derived from the enabling of new applications that traditionally could only run with a very large amount of online storage space
- **Increased customer service** derived from faster and more error-free response to requests for archived data

General Principles of Hierarchical Storage Management

The following section provides a detailed overview of how HSM works and the benefits you can gain from its use.

The Cost of Data Storage Components

We will take each component of a typical computer system that is used to store data and rank them in the order of cost of storage. In descending order one might draw up the following list:

1. CPU registers **MOST EXPENSIVE**
2. Cache memory
3. Main memory
4. Extended memory (typically IOP controllers)
5. Primary storage (typically high performance disks)
6. Secondary storage (typically compressed disks)
7. Optical library
8. Tape library
9. Vaulted tape (offline storage) **LEAST EXPENSIVE**

The expense of each data storage medium is normally expressed as a cost per megabyte (MB) of data stored.

The Speed of Access of Data Storage Components

We can rank these same components by their speed of access. Speed of access is typically the number of seconds that it might take to transfer a word (that is, two bytes) of data from this component to the CPU ready for processing. The list may well look like this:

1. CPU registers **FASTEST**
2. Cache memory

3. Main memory
4. Extended memory (typically IOP controllers)
5. Primary storage (typically high performance disks)
6. Secondary storage (typically compressed disks)
7. Optical library
8. Tape library
9. Vaulted tape (offline storage) **SLOWEST**

Note: HSM with BRMS does not currently support optical storage devices.

It may come as no surprise, then, that the data storage components with the fastest access are the most expensive. If we assume that to maximize the return on our computing investment we wish to deliver maximum performance (or transaction throughput) for minimum investment, we need to balance the amounts of differing levels of resource with overall cost in our system.

Assume that we have a poor performing interactive AS/400 application. User response times are too long. A possible solution may be to allocate more system memory to be available to that application. To do this, you may need to purchase extra memory for the AS/400 system. This allows you to access more of the application's active data directly from fast memory rather than slower disk.

The result is that to increase the performance of that application you must spend more money on its resources. Given that the application has a fixed total active data set size, we are effectively placing a larger proportion of that data set in a more expensive data storage component.

The application may work with a 200MB data set (stored on disk), of which 2MB may have been "active" at any one time. If we only have 1MB of main memory allocated to this application, a significant amount of data access work will come directly from the disk. This method slows the speed at which the application may work.

The effect is that we pay more money to attach more of the faster data storage components and transfer as much of the data as possible to the fast access areas. Making the best use of the more expensive components also involves the intelligent transfer of data across the interfaces between the storage levels.

The Hierarchical Storage Management Diagram

Figure 2 on page 4 shows the way in which we might distribute the data across a typical system. The general theory is to place the most frequently accessed data in the higher-speed access storage components.

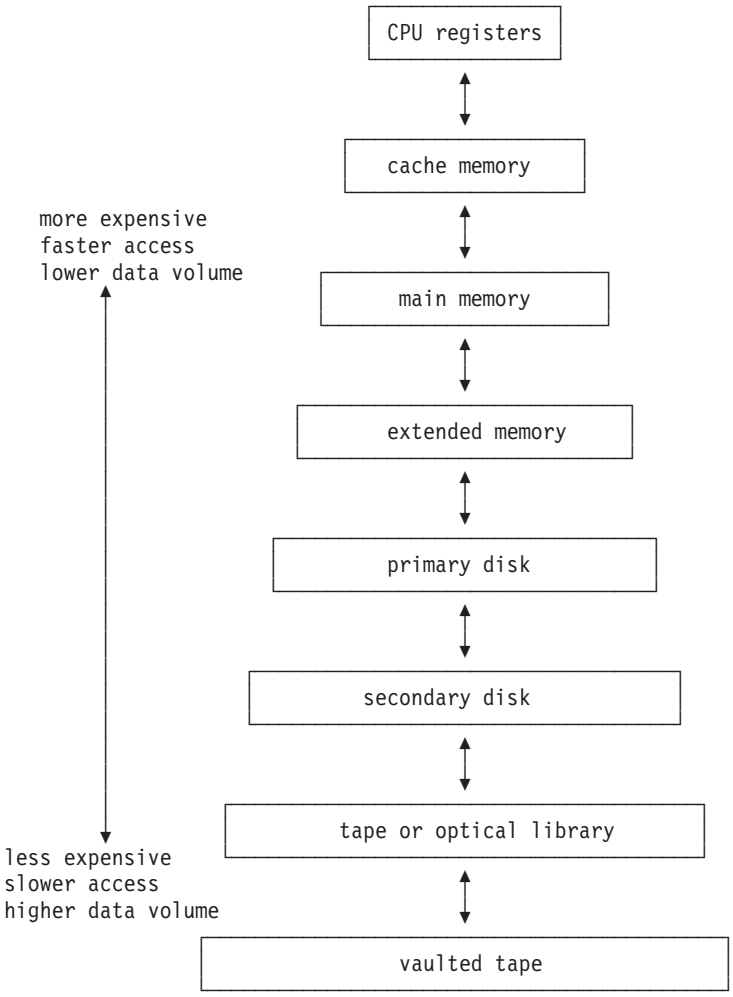

Figure 2. Hierarchical Storage Management Diagram. Shows the hierarchy of storage components in terms of expense, speed, and typical volume of data stored.

Expense constraints, however, may limit the amount of high performance storage that we attach to a system. Most of the system data is stored in the lower performance data storage components. The increasing width of the hierarchical layers, as we descend from the high cost components to the low cost ones, represents the increasing size of the data space that is attached to the system.

It is not always the same data that resides in the high performance storage components. Data is moved among the different layers according to the current system needs. The key to successful and seamless hierarchical storage management lies in the management and distribution of data across the different layers. Specific system functions will handle each boundary between layers (or “interface”) and the transfer of data from one layer to another (across the interface).

The AS/400 single level storage implementation manages the migration of data across the memory-to-disk interface. This implementation uses caching algorithms to ensure maximum performance of the disk storage layer by minimizing the number of data requests that must come directly from disk. This is often referred to as **disk paging**. Thus, we see that OS/400 provides hierarchical storage management on the AS/400 on all layers down to the disks.

Hierarchical Storage Management Overview

Hierarchical storage management can be especially useful at sites that maintain large volumes of historical information but do not require rapid access to that information. The following figure illustrates how hierarchical storage management provides an automated, transparent management system that distributes data across a storage hierarchy.

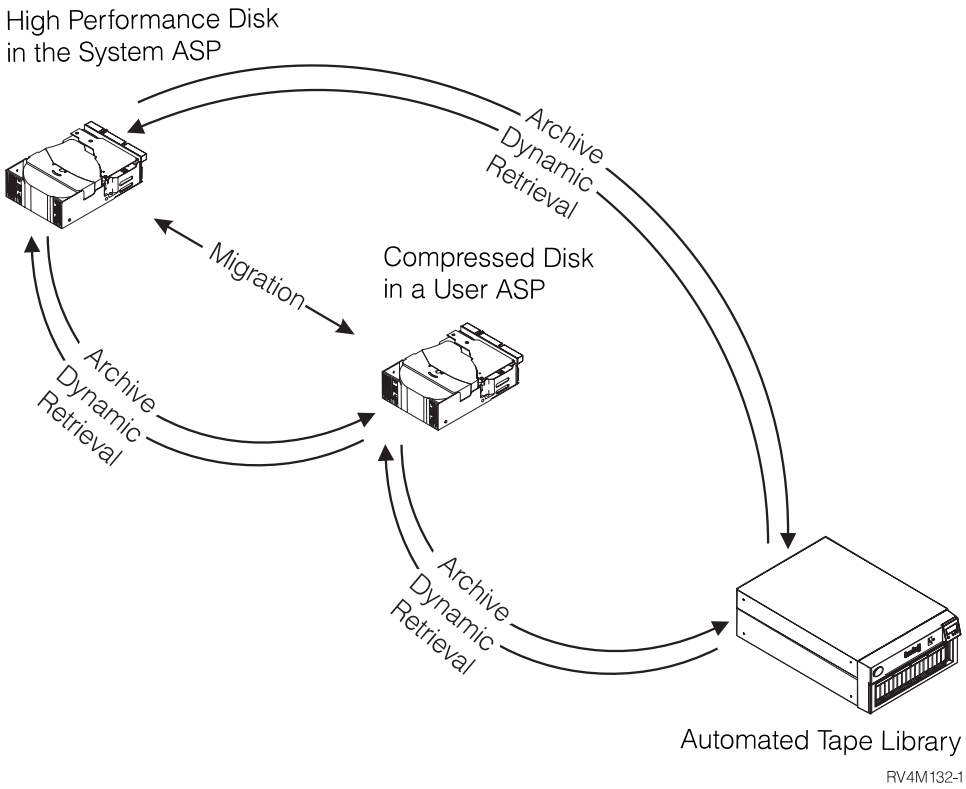

Figure 3. Hierarchical Storage Management Overview. Shows the different storage components of migration, archive, and dynamic retrieval.

Migration is an operation in which selected data is physically moved, or migrated, to different auxiliary storage pools (ASPs). The use of ASPs differentiates disk units of differing physical characteristics. It is an effective technique for moving data that is infrequently used. Once you migrate the data, applications can access the data directly from the lower cost storage. Currently, BRMS migration supports the automatic movement of data between ASPs.

Archiving sets up an interface from disk to tape or from disk to disk (save files) allowing us to move inactive data to a less expensive form of storage. This action frees up storage space. BRMS will archive data and track information about objects that you have archived.

Once you archive the data, applications that access the data will automatically move it back to disk when it is accessed. This is known as **dynamic retrieval**.

The difference between migration and archiving is the bidirectional interface that dynamic retrieval provides. The dynamic retrieval interface is where restoring the data back to disk automatically when it is accessed can make data available for

processing again. A feature that is called the Media and Storage Extensions or MSE accomplishes this within BRMS. See “Software Requirements” on page 10.

The ability to transfer data across the disk and tape interface in both directions has always existed but never before has it been system controlled, that is, automatic. With the introduction of **automated tape library devices,** the transition to full automation is complete because you do not have to wait for operators to mount and demount tapes.

Hierarchical Storage Management Planning

Overall system considerations may occur when planning an HSM solution. These include:

- User ASPs are required to migrate data or move spooled files.
- Disk compression is limited to user ASPs. For a list of disk compression limitations, see “Disk Compression” on page 8.
- Migration of data is done on a user library level (not file level) and root folder level (not document level) to maintain application transparency. Any library that starts with ’Q’ will not be migrated. These libraries are assumed to be system libraries and are not required for recovery. For a list of other migration limitations, see “Migration Limitations” on page 13.
- You archive data at a database file level and member level, or at a document or a stream file level. You can retrieve these object types automatically on demand (dynamically) when applications need them. With the archiving and dynamic retrieval functions, your tape storage method becomes an extension of disk storage.
- When moving spooled files from one ASP to another, you must specify an output queue and library. Application changes may be required to locate the new library and new output queue for the spooled files.
- You can archive spooled files to tape, but currently you cannot retrieve them dynamically.
- You may want to consider scheduling hierarchical storage management functions to occur on non-peak times. This can be achieved by specifying start and end times on the STRMGRBRM and STRARCBRM command.
- An OS/400 Info APAR is available, II11238, and contains information regarding hierarchical storage management recovery and other considerations.

Chapter 2. Hardware Considerations and Software Considerations

Components of a hierarchical storage management solution include:

- Hardware features
  - User ASPs, typically disk compression
  - Automated tape library
- Software requirements
  - Backup Recovery and Media Services for AS/400 and Media and Storage Extensions (MSE) for custom applications
  - OS/400 APIs and custom applications

You can use these components in any combination to customize your business needs.

Hardware Features

The following hardware features may be considered when planning for your HSM solution.

User Auxiliary Storage Pools (ASPs)

An **auxiliary storage pool (ASP)** is a group of units that is defined from all the disk units that make up auxiliary storage. ASPs provide the means of isolating objects on a specific disk unit, or disk units. Isolating the objects prevents the loss of data due to a disk failure on other disk units not included in the ASP.

The system automatically creates **system ASP (ASP 1)** and includes unit 1 as well as all other configured disks that are not assigned to a user ASP. The system ASP contains all system objects for the OS/400 licensed program and all user objects that are not assigned to a user ASP.

Now with HSM, you can use user ASPs to store historical or infrequently used data. You can assign a user ASP to any disk unit, which includes compression-capable disk units, to increase capacity, and yet be accessible from your applications.

Grouping together a physical set of disk units and assigning them to an ASP creates a **user auxiliary storage pool**. You can configure user ASPs 2 - 16. They can contain libraries, documents, and isolated objects of certain types. The system allows two types of user ASPs:

- Non-library user ASP

  Contains journals, journal receivers, and save files whose libraries are in the system ASP. These libraries are not eligible for HSM migration.
- Library user ASP

  Contains libraries and document library objects (DLOs). Library user ASPs are recommended because the recovery steps are easier than with non-library user ASPs.

To configure a user ASP, you can use System Service Tools (SST). For more information about user ASPs, see the *Backup and Recovery* book.

© Copyright IBM Corp. 1998, 1999

Disk Compression

Disk compression is a technology that increases the apparent capacity of disk storage devices by encoding the data to take up less physical storage space on disk. Disk compression is performed in the disk subsystem controller and does not affect the system processor. The compression and decompression of data is performed automatically on each read command and each write command and is transparent to applications.

You can use **disk compression** with an automated tape library to provide a total solution, or any combination of the components. Disk compression allows you to trade performance for storage at a lower cost per MB. It is a method of increasing the physical capacity of individual disk storage devices.

Disk compression characteristics:

- An average 2x compression ratio, with up to 4x achievable depending on type of data
- Compressed disk performance may be slower
- Disk compression is limited to user ASPs
- Hardware data compression is on the disk controller (compression-capable IOP)
- Customer on/off compression option per disk arm
- Compressed disk can be Redundant Array of Independent Disks (RAID) protected or mirrored

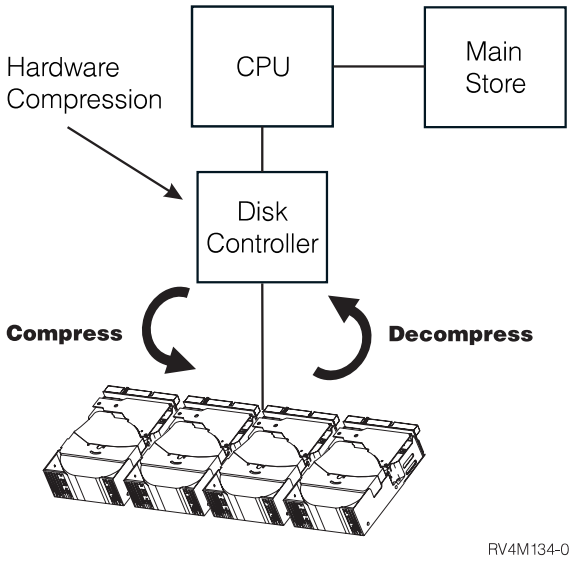

Figure 4. Example of Disk Compression

Restrictions and limitations

- Disk compression is allowed only on non-configured units. If you need to perform this operation on a configured unit, perform a remove operation of the configured unit and then start compression on the non-configured unit.
- You cannot add disk units that are using disk compression to the *system* ASP. You can only add compressed units to a *user* ASP.
- You cannot use disk compression on the disk unit for the load source.

- You must attach internal disk units to a compression-capable IOP. Certain types and models of disk units, 1GB, 2GB, 4GB, 8GB, and 17GB drives, are candidates for disk compression.
- If you are using mirroring, disk compression status must be the same for each disk unit in the mirrored pair. Both disk units in the mirrored pair must either be compressed or not compressed.

Note: Be aware that using disk compression could cause some performance degradation on your system

You can use System Service Tools (SST) to start disk compression on any non-configured compression capable disk unit. You can use Dedicated Service Tools (DST) to start compression, stop compression, remove a disk unit (non-configure), or add a disk unit to an ASP. For further information, see the disk compression section of the *Backup and Recovery* book.

Automated Tape Library

The combination of BRMS with an automated tape library system provides a total automated solution for all your tape requirements. When you use tape automation and BRMS, you can design a solution to reduce operational costs due to less manual intervention because most tape operations are unattended. A reduction in capital cost due to archiving and dynamic retrieving functions enables you to increase the quantity of online (on-disk) data that can be moved to less expensive tape media.

Archival and dynamic retrieval function is supported on all types of automated tape libraries that are supported on the AS/400, due to the high speed or fast access capabilities. We recommend the IBM Magstar MP tape library technology, such as type 3570 and 3575.

The IBM Magstar MP 3570 tape library models are designed to provide a mid-range tape solution with a lower price than the 3590. The 3570 Tape Subsystem uses a new and unique data cartridge. The capacity is 5 GB per cartridge and up to 15 GB per cartridge with LZ1 compaction. The 3570 Tape Library is designed to operate with two 10-cartridge magazines that provide access to 100-300 GB of data.

IBM developed the **Magstar MP 3575 Tape Library Dataserver** specifically for the mid-range environment. This technology combines automated tape handling with high reliability and fast data access. The supported models provide for a range of 300 GB to 1.6 TB of data that is not compressed.

You can implement an automated tape library with a media management application such as BRMS. If used with BRMS, it provides a set of user-level commands for using tape libraries. If you do not use BRMS and are planning on developing your own media management application, OS/400 provides a tape library command interface with APIs in which the OS/400 Media and Storage Extensions (MSE) provides additional media management functions. For more information on APIs, see the *System API Reference* . For more information on implementing an automated tape library, see the *Automated Tape Library Planning and Management* book.

Software Requirements

Hierarchical storage management is part of the AS/400 operating system, OS/400. OS/400 provides APIs that allow you to implement HSM. However, to implement a full HSM solution, you should first install Media and Storage Extensions (MSE), BRMS, and the Advanced Functions feature (HSM).

MSE allows an application, such as BRMS, to dynamically retrieve objects that you saved with storage freed. Each time the AS/400 accesses an archived object, it calls a user exit program that is provided by an application, such as BRMS, through the MSE exit point. BRMS then determines that an archived version of the object should be restored to the system to satisfy the OS/400 request. With the addition of HSM, which includes migration and archive with dynamic retrieval function, BRMS provides a complete, integrated solution for AS/400 storage management. After you install BRMS, you can install the Advanced Functions (HSM) feature.

Installing MSE

You must install the Media and Storage Extensions (MSE) feature on your AS/400 *before* you install BRMS. If you do not, BRMS cannot register the exit points and your saves will not complete. Take the following steps to install MSE on your system:

1. Enter GO LICPGM from a command line.
2. At the Work with Licensed Programs display, select option 11 (Install Licensed Programs).
3. Select Product 5769SS1, Media and Storage Extensions (MSE), and press Enter.
4. Press Enter again to confirm your selection.
5. At the Install Options display, type in the name of your installation device as requested. Press Enter to process the installation.
6. Review the job log to ensure that the installation completed successfully. Use the Display Job Log (DSPJOBLOG) command to review the job log.

You also need to ensure that the QSYS2 library is in your system library list. Use the Display System Values command (DSPSYSVAL QSYSLIBL) to check. If required, be sure to add the library to your list.

When you install MSE (OS/400 feature), you create OS/400 exit points. When you install BRMS, BRMS registers itself with MSE. These exit points track save and restore activities by using OS/400 commands. BRMS commands also use these exit points.

Note: If you load BRMS before MSE, the exit points are not created. In this case, you need to run INZBRM(*DATA) after loading MSE to register BRMS programs with the exit programs.

Installing BRMS

To install BRMS on your AS/400, perform the following steps:

1. Enter GO LICPGM from a command line.
2. At the Work with Licensed Programs display, select option 11 (Install Licensed Programs).

3. Select Product 5769BR1, option *BASE, to install the standard Backup, Recovery, and Media Services for the AS/400 product. Then press Enter.
4. Press Enter again to confirm your selection.
5. At the Install Options display, type in the name of your installation device as requested. Then press Enter to process the installation.
6. Review the job log to ensure that the installation completed successfully. Use the Display Job Log (DSPJOBLOG) command to review the job log.

During installation, BRMS performs the following tasks:

- Creates tape drive and media class entries that are based on the installed tape drives.
- Creates default BRMS objects.
- Creates the QBRM and QUSRBRM libraries, which contain all BRMS-related objects and information.
- Registers BRMS with the MSE feature of OS/400.

For more information about BRMS, see the *Backup Recovery and Media Services for AS/400* book.

Installing the Advanced Functions feature (HSM)

To install HSM on your AS/400 system, take the following steps:

1. Enter GO LICPGM from a command line.
2. At the Work with Licensed Programs display, select option 11 (Install Licensed Programs).
3. Select Product 5769BR1, option 2, to install the BRMS Advanced Functions feature. Then press Enter.
4. Press Enter again to confirm your selection.
5. At the Install Options display, type in the name of your installation device as requested. Then press Enter to process the installation.
6. Review the job log to ensure that the installation completed successfully. Use the Display Job Log (DSPJOBLOG) command to review the job log.

You are now ready to implement your HSM solution in a BRMS environment.

Chapter 3. Hierarchical Storage Management Migration

This chapter deals with what migration is, the implementation, the types of data that you may consider for migration. This chapter also provides details on how to set up and use the BRMS functions, groups, and policies that are provided to enable migration.

What is Migration?

Migration is an operation where you physically move selected data to different auxiliary storage pools (ASPs). A typical use is to move data from fast, high performance disk (system ASP) to slower or compressed disk (user ASP). This results in saving space on fast disk. Except for possible changes in access times, data that has been migrated is still fully available to any application that was able to access it before the data was migrated. However, now you will access it from the migrated area. If required, data will be moved from slow disk to fast disk. You can perform migration operations with or without the control of migration policies by using BRMS.

Migration functions include:

- Migrating **user libraries** from one ASP to another ASP
- Migrating **root folders** and its contents from one ASP to another ASP
- Moving **spooled files** from output queues in one or more ASPs to an output queue in another ASP.

  Note: A command to move spooled files is available in BRMS (Move Spooled Files using BRM (MOVSPLFBRM)). This command moves selected spooled files to a specific library/qualified output queue. You can move spooled files from one ASP to another (for example, from system ASP to user ASP) to free up disk space in the system ASP). See “Using the MOVSPLFBRM Command” on page 40 for more information about this command.

Migration Limitations

When you plan the migration of data, remember:

1. System libraries and system folders must remain in the system ASP.
2. User libraries to be migrated must not be in use and cannot be in the library list of any active job on the system, such as CHGLIB or CLRLIB.
3. Libraries containing journal objects and journal receivers will not be migrated.
4. Libraries containing files with database dependencies outside the library will not be migrated.
5. Libraries that begin with ’Q’ are considered to be system libraries, and are not migrated.
6. A library that is the current library (*CURLIB) is not allowed to be migrated.
7. When migrating libraries, the QSYSWRK subsystem must be active.
8. Program (*PGM) objects in a library will be placed in library QRPLOBJ. Another copy of each *PGM object will be moved with the library to the target ASP.

© Copyright IBM Corp. 1998, 1999

9. You can migrate first level root folders and their underlying documents. The API creates a temporary save file during the migration.
10. Though you can migrate library objects that you archived with storage free, you cannot migrate first level root folders that you archived with storage free.
11. Root folders that were restored with the RSTLICPGM (Restore License Program) command should not be migrated.
12. Root folders that contain more than 99 subfolders cannot be moved.
13. When migrating folders, the user must be enrolled in the system distribution directory.
14. After the migration of user libraries and root folders, the following attributes are changed:
    - The date that is last used will be set to blank.
    - The change date and time will be set to the date and time of the migration.
    - The days used count will be set to zero.
    - The date use count reset will be set to blank.
    - The restore date and time will be set to the date and time of the migration.
15. The user profile must have *OBJAUT authority to the user libraries and root folders.
16. When moving spooled file entries to a user ASP, you should specify an output queue and library. You might have to make application changes if you move the spooled file to a different output queue and library.
17. Only one migration operation (MGRBRM, STRMGRBRM, or MOVSPLFBRM command) is allowed at a time.
18. Job queue and output queue entries are migrated if the queue is not allocated.
19. Data queue entries are the responsibility of the user as migrating the libraries may cause the data queue entries to be lost.

For a more detailed list of the Move Library to ASP (QHSMMOVL) and Move Folder to ASP (QHSMMOVF) API limitations, see the *System API Reference* book.

Migration Considerations

You should develop standards or policies that control the libraries, root folders, or spooled files that are being migrated or moved. These policies can include:

1. **Library structure - Is it suitable for migration?**

   Your ability to migrate libraries depends on what object types are in the library. See the list of migration limitations to determine which libraries do not qualify for migration. For example, if your library consists of source files, programs, or physical files, you may consider migrating this library. If it contains journal and journal receivers, network database dependencies, then the library is not suitable for migration.

2. **Length of time a library, root folder, or spooled files have been inactive**

   The dormancy criteria for migration are date last used and change date. Updates to a file or member may not cause the date last used of a library to be updated. Specifying the parameters *ANY, *ALL, or percent value in the migration policy, ASP class attributes, or migration control group attributes causes BRMS to do file level searching for a candidate to migrate.

3. **Age of a library, root folder, or spooled files**

   You may consider migration for items that have aged for a certain period of time. The age criterion for migration of moved spooled files is the creation date.

4. **Size of a library, root folder, or spooled files**

   You may consider migration for items that have reached a certain size to free up some storage space from the source ASP. You can specify size in the migration policy, ASP class attributes, or migration control group attributes when migrating libraries or folders or when using the MOVSPLFBRM command.

5. **Retention period**

   You may consider migrating libraries or root folders or moving spooled files to a slower ASP (demote) temporarily for a certain period of time.

   Based on access requirements to the data, you may consider migrating the data back up to faster ASP (promote) if operations to the data require faster access. After a certain inactivity limit, consider moving them back (demote) to slower ASPs.

   To permanently store the data for legal or taxation purposes, you may consider archiving the data to tape. See “Chapter 4. Archival and Dynamic Retrieval” on page 43 for more information.

6. **Ensure enough storage space in target ASP**

   BRMS will ensure that there is enough storage space prior to migrating libraries or root folders. When you move root folders, you create a temporary save file on the target ASP, which is then deleted after a successful migration of the root folder.

7. **When will migration or moving take place?**

   Migrating user libraries and root folders, or moving spooled files does not require restricted state. However, you should run these operations in batch during non-peak hours.

How Does Migration Work?

The migration function moves data from its existing auxiliary storage pool (ASP) through a save process and a restore process while preserving private authorities to the objects. The migration functions of BRMS provide policies, control groups, commands, and reports that enable the user to define what, when, and how items are to be moved across ASP classes in the ASP storage hierarchy. The basic functions for planning a migration strategy using BRMS is as follows:

1. **Identify ″sets″ of data suitable for migration**

   Establish suitable migration criteria for each ″set″ of data, grouping by common characteristics. Identify any job that is associated with the various data sets.

2. **Establish a migration policy**

   The migration policy allows you to specify migration characteristics that you can use as values in the migration control group or ASP class attribute definitions. However, you can override the values that are specified with a control group or an ASP class.

3. **Establish an ASP class based on ASP attributes**

   When you use auxiliary storage pool (ASP) classes, you can identify ASPs that are used in the migration operation. You can add, remove, change, or display ASPs in an ASP class, work with ASPs that belong to a specified class, or change ASP class attributes.

4. **Build a migration control group**

   A migration control group is a list of libraries or first level root folders that share common migration characteristics. You can create multiple control groups and adjust each control group’s attributes where they vary from the migration policy.

5. **Add the migration jobs to the scheduler** (if required). You can perform automatic migration operations through the use of BRMS commands in conjunction with OS/400 job scheduling.

BRMS supports a hierarchical policy structure, which allows you to set global control group attributes by using the system policy and the migration policy. You can establish a set of the most commonly used attributes. You can set the policies with these attributes and even override these values (where necessary) at the control group level.

Migration commands are:

- Start Migration using BRM (STRMGRBRM)
- Migrate using BRM (MGRBRM)
- Move Spooled Files using BRM (MOVSPLFBRM)

The STRMGRBRM and MGRBRM commands are defined later in sections “Using the STRMGRBRM Command” on page 34 and “Using the MGRBRM Command” on page 35.

The Move Spooled Files using BRM (MOVSPLFBRM) command moves selected spooled files to a specified library/qualified output queue. See “Using the MOVSPLFBRM Command” on page 40 for more information about this command.

Setting Up Migration Using BRMS

The first step in setting up migration is to configure auxiliary storage pools (ASPs) to the system. For a brief description of an ASP, see “User Auxiliary Storage Pools (ASPs)” on page 7. See the *Backup and Recovery* book for instruction on how to set up an ASP.

After configuring ASPs, do the following:

1. Setting up policies to control the overall HSM strategy.
2. Setting up ASP classes.
3. Setting up ASP descriptions and assign them to corresponding ASP classes.
4. Setting up control groups and scheduling of migration jobs.
5. Starting migration operations.
6. Checking logs.

Setting Up the Migration Policy

Like the other BRMS services, you use a single policy to store a set of values that can serve as default values for an ASP class or a migration control group. The Change Migration Policy display allows you to specify migration characteristics that you can use as values in the migration control group attributes or ASP class attribute definitions. You can override values that are specified in the migration policy in a migration control group or use the values specified in the migration policy. Likewise, you can override values that are specified in the migration policy in the ASP class attributes or use the values specified in the migration policy. Press F9 to review the system policy before specifying migration policy values. You can also reach the Change Migration Policy display by using the Work with Policies using BRM (WRKPCYBRM TYPE(*MGR)) command.

Work with Migration Policy

The migration policy has no default values, so you are responsible for specifying all values. To set up the migration policy use the WRKPCYBRM *MGR command or specify these values:

1. Type `GO BRMS`
2. Select option 5 (Migration)
3. Select option 1 (Migration Planning)
4. Select option 1 (Work with migration policy)
5. Select option 1 (Change migration policy) to set the default values. You can use the values in the migration policy in control groups or an ASP class description by specifying the special value of *MGRPCY or values unique to them.

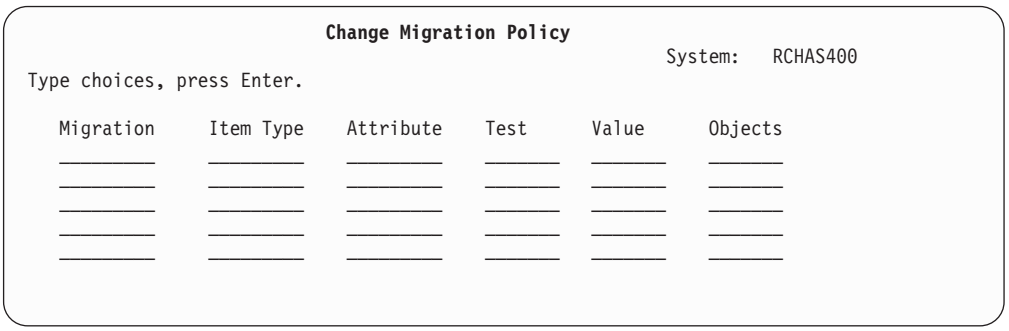

Figure 5. Change Migration Policy Display

Following is an explanation of each field:

- **Migration**

  Specifies whether items are to be migrated to a specific ASP class, or to an ASP class that is higher or lower in the ASP class hierarchy. ASPs that are members of an ASP class with a high hierarchy number reside on less expensive or slower storage devices than ASPs that reside in an ASP class with a lower hierarchy number. Two special values are available to indicate whether data is to be moved higher or lower in the ASP class hierarchy. *DEMOTE indicates that the data is to move to a higher ASP class (lower value, so therefore demoted to less expensive or slower storage devices). *PROMOTE indicates that the data is to move to a lower ASP class (higher value, so therefore promoted to faster disk). You can review the System Policy information by pressing F9 (System Policy).

- **Item Type**

  Specifies the type of object to migrate. Some examples are libraries, folders, or all libraries and folders that meet migration criteria. It can also be used to specify ″and″ logic or ″or″ logic by using the special values *AND or *OR. Special values *ALL, *LIB, and *FLR are left justified, and logic operators *AND and *OR are right justified.

- **Attribute**

  The attributes column allows you to specify the characteristics that you want to use as a migration criteria for selected libraries, first level folders, and objects that they contain. These characteristics that are each given a special value include size, age, creation date, date last used, and change date. Depending on the attribute that you select, a date, a number of days, or a number of

megabytes is specified in the Value column. The Test column indicates the comparison criteria between the attribute and the value, such as equal to, less than and so on.

If you specified the size special value, *SIZE, in the Attribute column, the migration criteria is the size that is calculated in megabytes. The way this size is calculated depends on whether there is a value that is specified on the corresponding line in the Objects column. If the Objects column is blank, the size of the entire library or first level folder is calculated. If you specify a percent in the Objects column, the size of each object in the library or first level folder is calculated.

For example, if you had specified *SIZE in the Attribute column, greater than (*GT) in the Test column, 300 in the Value column, and 50% in the Objects column, BRMS would calculate the size of all the objects in the library or folder. If 50% of the objects were greater than 300 megabytes in size, migration operations would proceed. If not, migration would not proceed.

- **Test**

  In the Test column you can specify the type of comparison that is to be performed between the value specified in the Attribute field and the Value field. Based on the results of the comparison and other migration criteria, libraries or first level folders may or may not be migrated.

  There are two types of test criteria; those that deal with dates, and those that deal with size or age. The following is a list of each test and valid uses of the test:

  1. Test types *AFTER, *BEFORE, *IS, and *ISNOT are date tests. You can use them when you specify *CRTDATE, *LASTUSE, *LASTCHG, or *LASTMGR in the Attribute column. You must specify *CURRENT or a date in the Value column.
  2. Test types *EQ, *NE, *LE, *GE, *LT, or *GT are size or number of days tests. You can use them when you specify *SIZE, *AGE, *LASTUSE, *LASTCHG, or *LASTMGR in the Attribute column. You must specify a number of megabytes for *SIZE or number of days for *AGE, *LASTUSE, *LASTCHG, or *LASTMGR in the Value column.

- **Value**

  The *Value* field contains a date, a number of days, or number of megabytes (if *SIZE is used as an attribute) that is compared to a value specified in the *Attribute* field. The special value *CURRENT represents the current date. Date values are also valid for *CRTDATE, *LASTUSE, *LASTCHG, and *LASTMGR special values when you use them in the attribute column.

- **Objects**

  The Objects column allows you to specify whether the attribute, test, and value specifications on a line apply to the library or first level folder or to one or more objects in the library or first level folder. If the Objects field is blank, the test criteria for the line apply to the entire library or first level folder. If you specify a value, the test criteria apply to one or more objects in the library or folder. You can specify a value as a number, a percent, or several special values.

Usage Considerations for Migration and Item Type Entries

There are rules that govern the types of entries that you make in the *Migration* and *Item Type* fields. Most of these are enforced through error messages that you receive when you make an improper entry to one of these displays. The entries that will be discussed in the *Migration* column are *PROMOTE and *DEMOTE. The

entries that will be discussed in the *Item type* column are *ALL, *LIB, and *FLR which are special values and *OR and *AND, which are logical operators.

Following are a list of entry rules. Because these entry rules only concern the *Migration* and *Item Type* fields, these examples show data only in those fields. When you want to process a migration control group, however, you need to fill out all of the required fields.

1. You can have one occurrence of the migration special values *PROMOTE or *DEMOTE with any one of the item type special values in the Change Migration Class Attributes display. For example, you could have the following and not receive an error message:

```
                  Change Migration Control Group Attributes       RCHAS400

Group  . . . . . . . . . . . . .  *MGRGRP

Type choices, press Enter.

    Migration   Item Type  Attribute  Test      Value      Objects
    *PROMOTE     *ALL
    *PROMOTE     *LIB
    *PROMOTE     *FLR

                                                                      Bottom
F3=Exit    F6=Insert    F9=Migration policy
F10=Work with ASP classes    F12=Cancel
```

*Figure 6. Using the *PROMOTE special value*

You could have any one of these entries, any two of these entries, or all three, and they would be correct.

The following example shows an incorrect entry with two uses of the item type special value *ALL with the *DEMOTE special value. To correct this entry, you would have to remove or change one of the *DEMOTE *ALL entries.

```
                  Change Migration Control Group Attributes       RCHAS400

Group  . . . . . . . . . . . . .  *MGRGRP

Type choices, press Enter.

    Migration   Item Type  Attribute  Test      Value      Objects
    *DEMOTE      *ALL
    *DEMOTE      *LIB
    *DEMOTE      *ALL




                                                                    Bottom

F3=Exit    F6=Insert    F9=Migration policy
F10=Work with ASP classes    F12=Cancel
```

*Figure 7. Incorrect use of the *DEMOTE special value*

2. You can have multiple occurrences of item type logical operators with each migration and item type combination. When you use a logical operator, you do not need an entry in the *Migration* column. In fact, when you press Enter, you blank out the entry. For example, if you enter the following:

```
                  Change Migration Control Group Attributes       RCHAS400

Group  . . . . . . . . . . . . .  *MGRGRP

Type choices, press Enter.

    Migration   Item Type  Attribute  Test      Value      Objects
    *DEMOTE      *ALL
    *DEMOTE      *OR




                                                                    Bottom

F3=Exit    F6=Insert    F9=Migration policy
F10=Work with ASP classes    F12=Cancel
```

Figure 8. Multiple occurances of item type logical operators

The resulting entry would appear as:

```
                  Change Migration Control Group Attributes       RCHAS400

Group  . . . . . . . . . . . . .  *MGRGRP

Type choices, press Enter.

    Migration   Item Type  Attribute  Test      Value      Objects
    *DEMOTE     *ALL_____
    _______     ______*OR

                                                                    Bottom

F3=Exit    F6=Insert    F9=Migration policy
F10=Work with ASP classes    F12=Cancel
```

Figure 9. Multiple occurances after pressing Enter

Following is an example that uses multiple logical operators:

```
                  Change Migration Control Group Attributes       RCHAS400

Group  . . . . . . . . . . . . .  *MGRGRP

Type choices, press Enter.

    Migration   Item Type  Attribute  Test      Value      Objects
    *PROMOTE     *LIB__
    ________     ___*OR
                 ___*AND
    *PROMOTE     *FLR__
    ________     ___*AND
    *PROMOTE     *FLR__
    ________     ___*AND
    ________     ___*OR

                                                                    Bottom

F3=Exit    F6=Insert    F9=Migration policy
F10=Work with ASP classes    F12=Cancel
```

Figure 10. Using multiple logical operators

Migration operations that would normally promote or demote an item from one ASP to another can do so only if the item resides in an ASP that has an associated ASP class. The item would be placed in an ASP that has an associated ASP class. These types of migration operations require BRMS to understand the ASP storage hierarchy (which is established using ASP classification).

Aside from promotion or demotion, migration operations that would normally migrate an item from one specific ASP to another can do so only if the item resides in an ASP with a low storage threshold that allows migration. The item would be placed in an ASP with a high storage threshold that allows migration. This is known to BRMS

only when you view the ASPs by using the WRKASPBRM display and have values for low and high storage threshold values other than *NONE.

Source ASP considerations:

- If the source ASP has a low storage threshold of *NONE, you cannot migrate from the ASP.
- If the source ASP has a low storage threshold of a specified percentage, migration from the source ASP occurs only if the current utilization exceeds this value.
- If the source ASP has a low storage threshold of *ASPCLS, migration from the source ASP occurs only if the current utilization exceeds the low storage threshold of the ASP class to which this ASP belongs. See “Changing ASP Class Attributes” on page 24.

Target ASP considerations:

- If the target ASP has a high storage threshold of *NONE, you cannot migrate into the ASP.
- If the target ASP has a high storage threshold of *SYS, you can migrate into the target ASP as long as you do not exceed the ASP storage threshold for the OS/400 operating system.
- If the target ASP has a high storage threshold of a specified percentage, you can migrate into the target ASP as long as you do not exceed the utilization limit.
- If the target ASP has a high storage threshold of *ASPCLS, you can migrate into the target ASP. The utilization (percentage) of the ASP must not exceed the high storage threshold of the ASP class to which this ASP belongs. See “Changing ASP Class Attributes” on page 24.

Omitting Items from Migration

The second selection on the Migration Policy menu allows you to specify items that you do not want to include in migration. You can add or remove folders or libraries from the list of items that you want to exclude from migration. The list specified here only applies when you are processing the Start Migration using BRM (STRMGRBRM) command, not the Migrate using BRM (MGRBRM) command. Two function keys, F19 and F20, assist you in selecting libraries or folders respectively to omit from migration.

```
                  Work with Items to Omit from Migration
                                                              System:   RCHAS400
Type options, press Enter.
  1=Add   4=Remove

Opt     Type     Migration Item
 _      _____    ___________
 _      *LIB     B*
 _      *FLR     PAYROLL



                                                                          Bottom
Type command, press Enter.
===>
F3=Exit   F4=Prompt    F5=Refresh   F9=Retrieve   F12=Cancel
F19=Select libraries   F20=Select folders
```

Figure 11. Work with Items to Omit from Migration Display

In the preceding example, all libraries that begin with the letter ″B″ are omitted from migration. The folder named PAYROLL is omitted also.

Setting Up ASP Classes

If you have already configured ASPs to the system, you can assign a classification to the ASPs that you want to manage through HSM. ASP classifications have management characteristics that you define. For instance, you could set up an ASP classification for high performance type data, high capacity data, dormant objects, and so on.

Work with ASP Classes

Select option 3 (Work with ASP classes) from the Migration Planning menu or issue WRKCLSBRM *ASP command.

```
                          Work with ASP Classes           RCHAS400

 Type options, press Enter
  1=Add       2=Change        4=Remove       5=Display   6=Work with ASP
  8=Change attributes
    Class                   -Threshold-
Opt Name        Hierarchy    Low  High  Text
1   _________
_   SYSTEM            1        20    60  High Performance disks
_   MEDIUM            5        30    70  High Capacity disks
_   COMPRESS         10        50    80  Dormant Objects disk
```

Figure 12. Work with ASP Classes Display

In the Work with ASP Classes display, three classes have been set up as an example. You can set up any number of classes. The ASP class is user defined. You can define the class name, hierarchy, thresholds, and accompanying text. To

add an ASP class, enter a 1 (Add) in the Option column and press Enter. You see the Add ASP Class display.

```
                              Add ASP Class                    RCHAS400

ASP class  . . . . . . . . . . :   _________     Name
Storage hierarchy priority . . :   ___           1 - 99
Low storage threshold  . . . . :   *NONE        0 - 99, *NONE
High storage threshold . . . . :   *NONE        1 - 95, *NONE, *SYS
Text . . . . . . . . . . . . . :   *NONE_____________________________
```

Figure 13. Add ASP Class Display

An ASP class defines attributes that are based on the following:

- **ASP class**

  Specify the ASP class name that you want to add. ASP classes are user defined.
- **Storage hierarchy priority**

  The priority values are inversely related to their relative position in the hierarchy. Larger priority values indicate a lower position in the hierarchy. Low priority values indicate a higher position in the hierarchy. You cannot have multiple classes with the same priority in the storage hierarchy.
- **Low storage threshold**

  When the storage utilization of an ASP exceeds the low storage threshold, you **can** migrate objects from the ASP. A special low storage threshold value of *NONE indicates that the ASP class does not allow migration of items out of ASPs of this class.
- **High storage threshold**

  When the storage utilization of an ASP is below the high threshold, you can migrate objects into the ASP. A high storage threshold value of *NONE indicates that ASP class does not allow migration of objects into ASPs of this class. A high storage threshold value of *SYS indicates that you should use the value of the ASP's operating system storage threshold for the threshold value. Use the System Service Tools to establish the storage threshold value for the operating system.

Changing ASP Class Attributes

Each ASP class that you set up has attributes associated with it. These attributes are the rules that determine when and where items selected for migration, that reside in ASPs of a particular ASP class, are migrated. Whenever you add an ASP class, you add a default set of attributes for the ASP class. You can use the default attributes, or change them to reflect the rules you want to assign to the ASP class.

When you select option 8 (Change attributes), you display the default entries for the ASP class. You can use the values from the migration policy by specifying *MGRPCY. However, if you use *MGRPCY, the special values in the Item type, column, the attributes, tests, values, and objects for this line must be blank.

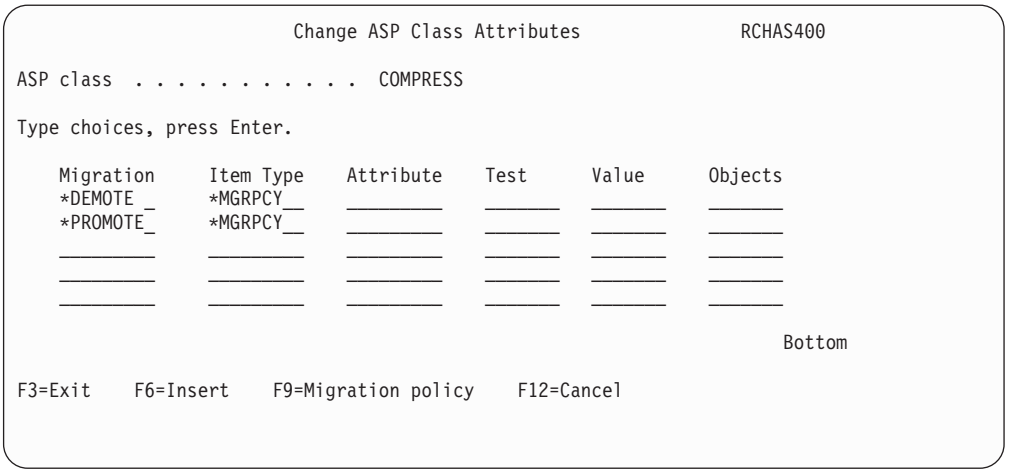

Figure 14. Change ASP Class Attributes Display with Defaults

The fields and associated definitions in the Change ASP Class Attributes are identical to those found in the Change Migration Policy display and in the Change Migration Control Group Attributes. Use F9 (Migration policy) to review the current values in the migration policy.

Setting Up ASP Descriptions

Once you have set up your ASP classes, classified them in terms of their position in the migration hierarchy, and entered the migration rules that you want to use, you must describe each ASP that you want to manage and assign it to an ASP class. As stated earlier, you are not setting up an ASP, but you are describing the characteristics of an ASP that you have already set up to use in HSM. Information includes ASP name, ASP class assignment, high and low storage thresholds, if different from the ASP class to which the ASP is assigned, and text for the ASP. BRMS also determines the eligibility of an ASP to participate in migration that is based on the content of the ASP.

Work with ASPs

To work with ASP descriptions, select option 4 (Work with ASPs) from the Migration Planning menu. You can also reach this display by processing the Work with ASP Descriptions (WRKASPBRM) command.

The Work with ASP Description display has three views. Press F11 to see each view. The first view shows the class. The second view shows the utilization of each ASP and their threshold level. The third view shows the text description.

The first view shows the ASP class name, the ASP number, its class, the hierarchy, and whether or not the ASP is eligible to support migration operations.

```
                    Work with ASP Descriptions              RCHAS400

 Type options, press Enter
  1=Add       2=Change        4=Remove

Opt  Name       Number  Class      Hierarchy   Eligible
_   _________
_   *SYSTEM          1  SYSTEM             1       *YES
_   ASP02            2  MEDIUM             5       *YES
_   ASP03            3  COMPRESS          10       *YES
2   *NONE            4  *NONE                      *YES

F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve   F11=Show utilization
F12=Cancel
```

Figure 15. Work with ASP Description - By Class

The following statements are true for this example:

- ASP 1 is always the system ASP, and the special value *SYSTEM indicates ASP 1. The *SYSTEM entry was created during BRMS installation after detecting the system ASP.
- ASPs 2 and 3 have been named and are assigned to an ASP class.
- ASP 4 is configured on the system, but has not been named or assigned to an ASP class as indicated by the special value *NONE. You cannot identify an ASP with *NONE for a name except through use of an ASP number. When BRMS automatically configures ASP descriptions, these values are assigned:
  - A default ASP class value of *NONE
  - A low storage threshold value of *NONE
  - A high threshold value of *SYS
  - Text of ″Entry created by BRMS configuration″

To assign ASP 4 a class name, for example, take option 2 (Change).

```
                          Change ASP Description
  ASP Number . . . . . . . . . .  4

Type choices, press Enter.

  ASP name . . . . . . . . . . .   *NONE_____       Character
  ASP class  . . . . . . . . . .   *NONE           Name, *NONE, F4 for list
  Low storage threshold  . . . .   *NONE           0-99, *NONE, *ASPCLS
  High storage threshold . . . .   *SYS            1-95, *NONE, *ASPCLS, *SYS
  Eligible for migration . . . .   *YES            *YES, *NO
  Text . . . . . . . . . . . . .   Entry created by BRMS configuration
```

Figure 16. Changing an ASP Description

Many of the fields that are displayed above have already been previously described. The following fields require further description.

ASP name

Enter the name of the ASP that you want to change. ASP names are user-defined. You cannot add an ASP with a special value such as *NONE or *SYSTEM. However, you can change the descriptions that are associated with these special values.

ASP class

Specify the ASP class that is associated with the ASP name that you are adding or changing. If you want to select a class from an ASP class list, place the cursor within the ASP class field and press F4. The special value *NONE indicates that the ASP is not assigned to an ASP class.

Low storage threshold

Specifying a low storage threshold on the ASP description will override the value that is specified in the ASP Class. The special value *ASPCLS indicates that the system will use the low storage threshold value from the ASP class to which this ASP has been assigned.

High storage threshold

Specifying a high storage threshold on the ASP description will override the value that is specified in the ASP Class. The special value *ASPCLS indicates that the system will use the high storage threshold value from the ASP class to which this ASP has been assigned.

Eligible for migration

You can specify whether the ASP that you are adding or changing is eligible for migration. An indication is provided for each ASP which is used to show whether the ASP supports migration operations. Non-library user ASPs are not supported for migration purposes and are not eligible for hierarchical storage management. BRMS determines ASP eligibility that is based on ASP content. If the ASP contains a library or is capable of containing a library, it will be considered eligible for inclusion in the ASP hierarchy and for migration operations. An ASP can change its eligibility since an ASP that is empty can be made into a non-eligible, object only ASP by adding a journal receiver or save file to it. During migration operations, BRMS resets ASP eligibility if the ASP has been made ineligible due to a change in content. You see these changes when BRMS presents the Work with ASP Descriptions display. In these cases the ASP eligibility attribute is set to indicate that it is no longer eligible, and a message is placed in the BRMS log to indicate that the ASP is ineligible.

You can name and classify an ineligible ASP, and generally change any other attribute without restriction except for the eligibility attribute. These other attributes are ignored until the ASP becomes eligible. You can change the eligibility attribute under the following conditions:

- ASP eligibility is changed from *YES to *NO to make it ineligible; this includes *SYSTEM.
- ASP eligibility can be changed to make it eligible if the ASP is not a non-library user ASP.

Text You can specify the text that describes the ASP that you are adding. You can specify the special value *NONE which indicates that there is no text that is associated with the ASP or the text that you want to use.

Note: If you remove an ASP from BRMS ASP descriptions with option 4 (Remove), it will no longer show up on the Work with ASP Descriptions display. You will need to use option 1 (Add) to add it back into the BRMS ASP descriptions.

You see the second view of the Work with ASP Description display when you press F11 (Show utilization). Figure 17 on page 28 shows you the second view. In addition to the name and number of the ASP, you see all threshold information and current utilization information.

```
                        Work with ASP Descriptions              RCHAS400

 Type options, press Enter
  1=Add       2=Change        4=Remove
                              ----Threshold----  ----Current Utilization-----
 Opt  Name         Number  Low    High   Sys  %Used    Size   Avail  Ovrflw
 _   _________
 _   *SYSTEM          1      30     70   90   72.6   12583    1735  NO
 _   ASP02            2      20     60   90   36.6    2200    2068  NO
 _   ASP03            3      50*    70   90    3.6    1967    1895  NO
 _   *NONE            4      *NONE  *SYS 90   22.2   17119    17074 NO

F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve  F11=Show text   F12=Cancel
```

Figure 17. Work with ASP Description Display - By Utilization

Many of the fields that are displayed above have already been described previously. The following fields require further description.

Threshold values

Notice that the system places an asterisk ″*″ after the low or high threshold values. This reminds you that this value is not the same as that specified in the ASP class. You see the system threshold value. You can change this value by using the disk configuration options of the System Service Tool (SST). In the previous example the system threshold is set at 90%.

Percent used

The ASP media that is used as a percentage of the capacity size of the ASP is displayed.

ASP size

The total space in megabytes of auxiliary storage on the storage media that is allocated to the ASP is displayed. This is the sum of the unit media capacity fields for (1) the units allocated to the ASP or (2) the mirrored pairs in the ASP.

ASP availability

The number of bytes of secondary storage space that is not currently assigned to objects or internal machine functions, and therefore, is available for allocation in the ASP. Note that a mirrored pair counts for only one unit.

Space overflow

Whether allocations for unprotected data in the ASP have exceeded the unprotected space capacity and overflowed into the area normally used for allocation of protected data. A value of Yes indicates that such overflow has occurred; a value of No indicates that it has not. This status is set when the ASP unprotected space overflow event is signaled. It is reset automatically on the subsequent IPL of the machine. Because unprotected storage is used primarily for allocation of temporary objects which are automatically deallocated as part of the IPL process, the overflowed allocations are freed up at IPL. This action provides for the automatic reset of the overflow condition.

You see the final view of the Work with ASP Descriptions display when you press F11 (Show text). You see the text that describes each ASP. If you make no changes to the text description, you see the default text ″Entry created by BRM configuration.″

```
                          Work with ASP Descriptions
                                                              System:   RCHAS400
Type options, press Enter.
  1=Add   2=Change   4=Remove


 Opt  Name        Number  Text

  _   __________
  _   *SYSTEM         1   High Performance ASP
  _   ASP02           2   High Capacity ASP
  _   ASP03           3   Dormant ASP
  _   *NONE           4   Entry created by BRM configuration

                                                                        Bottom
Type command, press Enter.
===>
F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve   F11=Show class   F12=Cancel
```

Figure 18. Work with ASP Descriptions - By Text Description

Setting Up Migration Control Groups

Use **migration control groups** to group items for migration. Migration control groups are lists of libraries or first level folders that you group together because they share common migration characteristics. You can set up any number of control groups to accomplish your migration plan. The migration control group is like other BRMS control groups in that you can perform various tasks. You can create control groups, edit entries, copy, delete, display, add to schedule, change attributes, and select subsystems or job queues to process for one or more control groups. You perform these functions by using the Work with Migration Control Groups display.

Work with Migration Control Groups

To work with migration control groups select option 2 (Work with migration control groups) from the Migration Planning menu. You also can use the Work with Control Groups (WRKCTLGBRM *MGR) command.

```
                    Work with Migration Control Groups              RCHAS400

Position to . . . . . _________    Starting characters

Type options, press Enter
  1=Create    2=Edit entries    3=Copy       4=Delete    5=Display
  6=Add to schedule    8=Change attributes   9=Subsystems to process
    Control       Last
Opt Group        Migrated     Text

_   _________
2   *MGRGRP        *NONE    Entry created by BRM configuration
__  MGRLIB       04/12/98   Control group for libraries
__  MGRFLR       03/10/98   Control group for folders
__  MGRMEDIUM    04/01/98   Medium capacity control group
__  MGRCOMPRESS  03/01/98   High capacity control group
```

Figure 19. Migration Control Groups Display

The default migration control group name is *MGRGRP. The default migration control group is created during BRMS installation, and contains no entries. The last migrated field displays the date that the control group was last processed for migration. The special date *NONE indicates that the migration control group has never been migrated.

Use option 2 (Edit entries) to set up migration characteristics for the default control group. To change attributes of a control group, specify option 8. You can choose to use migration policy values through the use of the special value *MGRPCY or ASP class attributes by specifying the ASP class name. You also can use values that are unique to this control group. Press F9 to view migration policy values.

Editing Control Group Entries

In the following example, select the default migration control group *MGRGRP by selecting Option 2 (Edit entries) from the Work with Migration Control Groups display. On this display you can add, change, or remove migration items.

```
                    Edit Migration Control Group Entries            RCHAS400

Group . . . . . : *MGRGRP

Text  . . . . . . Entry created by BRM configuration _____________

Type information, press Enter.

        Migration     Item
Seq     Items         Type
60__    _________
  10    MYLIB*_____   *LIB
  20    YOURLIB____   *LIB
  30    MYFLR*_____   *FLR
  40    YOURFLR____   *FLR
  50    *EXIT______
                                                                  Bottom
F3=Exit                  F5=Refresh             F10=Change item
F11=Display exits        F12=Cancel             F24=More keys
```

Figure 20. Edit Migration Control Group Entries Display

You see the following display when you press F11 (Display exits). Only part of the exit command is displayed.

```
                  Edit Migration Control Group Entries

Group . . . . . . . . . . . : *MGRGRP

Text  . . . . . . . . . . . . Entry created by BRM configuration

Type information, press Enter.

        Migration
Seq     Items         Exit command

  10    MYLIB*
  20    YOURLIB
  30    MYFLR*
  40    YOURFLR
  50    *EXIT         SNDMSG MSG('Migration processing for the default mig
                                                                      Bottom
F3=Exit                   F5=Refresh             F10=Change item
F11=Display exits         F12=Cancel             F24=More keys
```

Figure 21. Edit Migration Control Groups display

You see the following display when you position the cursor on the sequence line 50 and press F10 (Change Item).

```
                      User Exit Maintenance                       RCHAS400

Type command, press Enter.

  Sequence number . . . . . . . :    50
  Where used  . . . . . . . . . :  *EXIT

  Command . . . . . . . . . . . .  SNDMSG MSG('Migration processing for the
  default migration control group is complete') TOUSR(*SYSOPR)_____________
  ___________________________________________________________________________
  ___________________________________________________________________________
```

Figure 22. User Exit Maintenance displays

Following are descriptions of key fields in the above displays.

Sequence

Sequence numbers indicate the order in which you process the entries. The lowest number indicates that the entry is to be processed first, and the highest number indicates that the entry is to be processed last. The numbers do not have to be consecutive (1, 2, 3, and so on).

Migration items

Migration items are libraries or first level folders that you want to include when processing this migration control group. You can specify a specific folder, a generic folder, a specific library, or a generic library. In addition to libraries and first level folders, you can specify the special value *EXIT.

The special value *EXIT, allows you to perform additional processing at selected sequences in the control group. Several function keys are available to assist you in entering information for this control group. If you

press F11 (Display exits), you see the second view of the Edit Migration Control Group Entries display. This view shows all or a portion of exit processing that you may have entered. If you want further information about the exit processing, or you want to enter exit processing, place the cursor on the sequence number of either view of the display that contains the *EXIT special value and press F10 (Change item). You see the User Exit Maintenance display.

Note: You cannot include any items in the migration items list that begin with the letter Q. You should never include QBRM or QUSRBRM in the list of items to migrate. You cannot specify the libraries QSYS, QSRV, QTEMP, QSPL, QRPLOBJ, QSPLxxxx, and QRECOVER.

Item type

The type of item can be one of two special values, *LIB (library) or *FLR (folder).

Exit command

You can specify most commands in the user exit command. Where you position the *EXIT keyword and the associated exit command in the list of migration items is important. See *Backup Recovery and Media Services for AS/400* book for special considerations for using *EXIT as a migration item.

Two other function keys that are useful but which are not shown here are the F19 (Select libraries) and F20 (Select folders) keys. You can use these function keys to display libraries or folders respectively and select those that you want to include in the migration control group as migration items.

Changing Control Group Attributes

For each migration control group you set up, you can specify the migration attributes (rules) that you want to use for this control group. Do this by selecting Option 8 (Change attributes) from the Work with Migration Control Groups display. The migration rules can be any of the following:

- Unique for this control group
- Based on rules that are specified in the migration policy through the use of the special value *MGRPCY
- Based on rules that are specified in the ASP class attributes through the use of the special value *ASPCLS. The *ASPCLS special value uses the ASP attributes from the ASP in which the object (specified in the control group) resides.

Press F9 to review the migration policy, and press F10 to review the ASP class attributes.

```
                  Change Migration Control Group Attributes        RCHAS400

Group  . . . . . . . . . . . . .  *MGRGRP

Type choices, press Enter.

    Migration     Item Type    Attribute    Test      Value      Objects
    *DEMOTE _     *LIB_____    *CRTDATE_    *BEFORE   05/01/99   *ANY___
    *PROMOTE_     *LIB_____    *LASTUSE_    *LT____   10_____    80%____
    _________     _____*AND    *SIZE____    *LE____   300____    _______
    COMPRESS_     *ASPCLS__    _________    _______   _______    _______
    _________     _________    _________    _______   _______    _______
    _________     _________    _________    _______   _______    _______
    _________     _________    _________    _______   _______    _______

                                                                      Bottom

F3=Exit    F6=Insert    F9=Migration policy
F10=Work with ASP class    F12=Cancel
```

Figure 23. Change Migration Control Group Attributes Display

The fields for the Change Migration Control Group Attributes and the associated definition are identical to those found in the Change Migration Policy display and in the Change ASP class attributes display. Review Figure 5 on page 17 for the definitions.

An explanation of this example is as follows:

1. DEMOTE criteria
   - Demote libraries (moved lower in the ASP hierarchy) that contain any objects that you created prior to 05/01/99.
2. PROMOTE criteria
   - Libraries are to be promoted (moved higher in the ASP hierarchy) which have been used in less than ten days and 80% of the objects in the library meet the criteria.
   - An ″AND″ criteria also specifies that to promote libraries selected in the first criteria, the library size must be less than or equal to 300 megabytes.
3. ASP COMPRESS criteria
   - Libraries or folders are to be migrated based on the ASP class from the ASP in which the library or folder resides. It will migrate the library or folder to ASP class COMPRESS if the ASP which currently contains the library or folder has an entry, COMPRESS.

For more information of the options available, see the *Backup Recovery and Media Services for AS/400* book for more details.

Starting Migration Operations

You can use the following commands in the migration process:

- Start Migration using BRM (STRMGRBRM)
- Migrate using BRM (MGRBRM)

Using the STRMGRBRM Command

The Start Migration using BRM (STRMGRBRM) command allows you to specify a migration control group to process. You can start the migration immediately, at a scheduled time, or you can schedule it using the system job scheduler. Processing can be batch or interactive. You can type the command on any command line or go through the Migration menu.

The default value for the *Option* parameter on the STRMGRBRM command is *REPORT, which produces the Migration Candidate report. The *REPORT value allows you to process a report for each item in the control group and review what will be migrated prior to performing the migration operation. To perform a migration for the items in the Migration Candidate report, or to restart a migration that was interrupted, you can specify the special value *RESUME.

When you process the STRMGRBRM command with the *MIGRATE value in the *Option* field, BRMS creates two reports, the Migration Candidate report and the Migration report. The Migration report creates a migration candidate file for each item in the control group. This file contains information about the items in the control group at the time that you processed the STRMGRBRM command. You also see information about ASPs, and about the size of the items. The Migration Candidate report that you created produces candidate reports and performs migration operations for the control group. If you specify *RESUME for a control group, the migration candidate file that was created for the control group when you processed the STRMGRBRM command is used. A new migration candidate file is not created when you specify *RESUME, even though items in the control group or ASP characteristics may have changed.

Notes:

1. If you use the special value *EXIT in a migration control group and the resulting processing changes the characteristics of items in a control group or an ASP, BRMS does not record the changes in the migration candidate file. Consider not using the *EXIT special value in migration control groups if the resulting processing alters control group or ASP characteristics.
2. The control group *EXIT should not use this command because item processing results are unpredictable.
3. The Migration Candidate report is not associated with a job or user ID. This report is only associated with the name of a control group. If job A processes a control group, and job B subsequently processes the same control group, then the candidate report that is created for job B overwrites the report created for job A. If you want to resume a control group, the migration candidate report created for job B will be used for the resume operation, even if it was job A that was interrupted.

```
                    Start Migration using BRM (STRMGRBRM)

Type choices, press Enter.

Control group  . . . . . .                *MGRGRP
Run option . . . . . .  .   *REPORT       *REPORT, *MIGRATE, *RESUME
Schedule time:
  Beginning time . . . . . .  *IMMED       hhmm, *IMMED
  Ending time  . . . . . .    *NONE        hhmm, *NONE
Submit to batch  . . . . . .  *YES         *CONSOLE, *YES, *NO
Job description  . . . . . .  *USRPRF      Name, *USRPRF
  Library  . . . . . .  .  .              Name, *LIBL, *CURLIB
Job queue  . . . . . .  .  .  *JOBD        Name, *JOBD
  Library  . . . . . .  .  .              Name, *LIBL, *CURLIB
End interactive jobs . . . . . .  *SYSPCY  *YES, *NO, *SYSPCY
Delay  . . . . . .  .  .  .  .  .  *SYSPCY  0-999, *SYSPCY




                                                                      Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 24. The Start Migration using BRM (STRMGRBRM) command

This report uses printer file QP1AMGR.

Using the MGRBRM Command

The Migrate using BRM (MGRBRM) command allows you to migrate a library or first level folder to a specified ASP. BRMS migration processing uses this command to request the migration of a specific library or folder as needed. You can specify the ASP to which you want to migrate the library or folder.

When you use the MGRBRM command, BRMS ignores low storage threshold constraints for the ASP from which the item is being moved. However, BRMS does honor the target ASP's high storage threshold.

```
                         Migrate using BRM (MGRBRM)

Type choices, press Enter.

To asp . . . . . .  .  .              Name, *SYSTEM, 1, 2, 3, 4...
Type . . . . . .  .  .  .  .  .  *LIB  *LIB, *FLR
Library  . . . . . .  .  .  .  .      Name
Folder . . . . . .  .  .  .  .  .     Name









                                                                      Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 25. Migrate using BRM (MGRBRM) command

This report uses printer file QP1AMGR.

Adding Migration Control Groups to the Job Scheduler

To add a migration control group to the job scheduler, take the following steps:

1. At the Work with Migration Control Groups display, type a 6 (Add to schedule) in front of the control group you want to work with and press Enter.

```
                      Add Job Schedule Entry (ADDJOBSCDE)

 Type choices, press Enter.

 Job name . . . . . . . . . . . . > QBRMMGR       Name, *JOBD
 Command to run . . . . . . . . . > STRMGRBRM CTLGRP(MGRLIB) OPTION(*REPORT) SBM
JOB(*NO)




 Frequency  . . . . . . . . . . . > *WEEKLY       *ONCE, *WEEKLY, *MONTHLY
 Schedule date, or  . . . . . . . > *NONE         Date, *CURRENT, *MONTHSTR...
 Schedule day . . . . . . . . . . > *ALL          *NONE, *ALL, *MON, *TUE...
                + for more values
 Schedule time  . . . . . . . . . > '00:01'       Time, *CURRENT



                                                                        More...
 F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
 F13=How to use this display        F24=More keys
```

Figure 26. The Add Job Schedule Entry (ADDJOBSCDE) display

2. Confirm that the command in the *Command to run* field indicates the correct control group job that you want the scheduler to process. Make corrections if needed.
3. Review and change the frequency, day, date, and time parameters as needed.
4. Press Enter to add the job to the scheduler and return to the Work with Migration Control Groups display.

To review the migration jobs currently in the job scheduler, press F7 (Work with BRM Job Schedule Entries) from the Work with Migration Control Groups display.

Working with Migration Activity

BRMS maintains a historical record of migration activity. This includes what, when, and where items were migrated as well as the success or failure of migration operations.

You can display migration activity by selecting Option 5 (Display migration activity) from the Migration menu. You see the Display Migration Activity display where you can review the migration log, work with migration exceptions, or work with migration information. Options 1 and 2 allow you to work with the BRMS log, and Option 3 allows you to work with migration information.

Working with Migration Information

To work with migration information, select Option 3 (Work with migration information) from the Display Migration Activity menu. You see the Work with Migration Information (WRKMGRIBRM) display where you can select whether you want to work with library or first level folder information. Once you select one of the special values, either *LIB (libraries) or *FLR (folders), you see additional parameters (see Figure 27). After tailoring the command, press Enter to go to the Work with Migration Information display. If you issue WRKMGRIBRM command, it defaults to *LIB unless you prompt and specify *FLR.

```
                Work with Migration Information (WRKMGRIBRM)

Type choices, press Enter.

Item type  . . . . . .   *LIB          *FLR, *LIB
```

```
                Work with Migration Information (WRKMGRIBRM)

Type choices, press Enter.

Item type  . . . . . .   *LIB          *FLR, *LIB
Library  . . . . . .   *ALL          Name, generic*, *ALL
From ASP . . . . . .   *ALL          Name, *ALL, *SYSTEM, 1, 2...
To ASP . . . . . .   *ALL          Name, *ALL, *SYSTEM, 1, 2...
From ASP class . . . . . .   *ALL          Name, *ALL
To ASP class . . . . . .   *ALL          Name, *ALL
Select dates:
  From date  . . . . . .   *BEGIN        Date, *CURRENT, *BEGIN, nnnnn
  To date  . . . . . .   *END          Date, *CURRENT, *END, nnnnn
Migration status . . . . . .   *ALL          *ALL, *NOERROR, *ERROR
Sequence . . . . . .   *DATE         *DATE, *NAME, *FROMASP...
Entries to be displayed first  .   *LAST         *LAST, *FIRST
Output . . . . . .   *            *, *PRINT
```

```
                     Work with Migration Information
                                                        System:   RCHAS400
Position to date . . .   ________       Starting date

Type options, press Enter.
  4=Remove   5=Display

     Migrated      Item                          -----Start-------
Opt  Item          Type  From ASP    To ASP      Date     Time      Duration
 _   MYLIB         *LIB  *SYSTEM     ASP02       04/12/98 14:02:15  00:03:03
 _   YOURLIB       *LIB  ASP02       ASP03       04/01/98 14:12:15  00:23:03
```

Figure 27. Work with Migration Information (WRKMGRIBRM) Command

The following is an explanation of the previous display.

Migrated items

The name of the library or folder that has been migrated from one ASP to another is displayed.

Item type
The type of item that has been migrated is displayed. The possible special values are *LIB, which indicates that the item is a library and *FLR, which indicates that the item is a folder.

From ASP
The ASP number, name, or special value for the ASP from which the library or folder was migrated is displayed. The special value *SYSTEM indicates that the library or folder was migrated from the system ASP. A number (2-16) is the number of the ASP from which the information was migrated. An ASP name indicates the name of the ASP from which the library or folder was migrated.

To ASP
The ASP number, name, or special value for the ASP to which the library or folder was migrated is displayed. The special value *SYSTEM indicates that the library or folder was migrated to the system ASP. A number (2-16) is the number of the ASP to which the information was migrated. An ASP name indicates the name of the ASP to which the library or folder was migrated.

Start date and time
The starting date and time for each library or folder that has been migrated from one ASP to another is displayed. The Date field and the Time field are expressed in job date and time format.

Duration
The duration is the difference in time between when the migration process for a library or folder began and when it completed. The duration is expressed in job time format.

Displaying Migration Information

Once you have displayed migration information, you have two options. You can display detailed information about the migrated folder or library (Option 5). You can remove information about the migrated folder or library (Option 4). Following is an example of the Display Migration Information display for a folder. You should review the online help information for each field in the display.

```
                    Display Migration Information

  Migrated item  . . . . . :   MYLIB
  Item type  . . . . . . . :   *LIB
  Size . . . . . . . . . . :   1830

  From ASP number/name . . . :   01/*SYSTEM
    ASP class  . . . . . . . :   SYSTEM
  To ASP number/name . . . . :   02/ASP02
    ASP class  . . . . . . . :   MEDIUM

  Start date . . . . . . . . :   04/12/98
    Time . . . . . . . . . . :   14:02:15
  End date . . . . . . . . . :   04/12/98
    Time . . . . . . . . . . :   14:05:18

  Control group  . . . . . . :   *NONE
  Job  . . . . . . . . . . . :   QPADEV000B/STEVEN/167768
  Error  . . . . . . . . . . :
```

Figure 28. Displaying Migration Information for a Folder

Displaying the Migration Log

If you select Option 1 (Display log), from the Display Migration Activity display, you see the Display Log for BRM (DSPLOGBRM) display. The default for the TYPE parameter is *MGR (migration) which displays only migration entries. The BRMS log is a detailed description of how items are processed within BRMS and is your first source of information when researching problems.

Option 2 (Display migration exceptions) also takes you to the DSPLOGBRM command where the TYPE parameter is specified as *MGR and the SEV parameter is specified as 40. The value 40 indicates that an operation has failed.

The following display illustrates how migration log items are displayed with all severity codes.

```
                      Display Log for BRM (DSPLOGBRM)

Type choices, press Enter.

Type . . . . . . . . . . . . . . > *MGR          *ALL, *ARC, *BKU, *MED...
Time period for log output:
  Start time and date:
  Beginning time . . . . . . . .   *AVAIL        Time, *AVAIL
  Beginning date . . . . . . . .   *CURRENT      Date, *CURRENT, *BEGIN
  End time and date:
  Ending time  . . . . . . . . .   *AVAIL        Time, *AVAIL
  Ending date  . . . . . . . . .   *CURRENT      Date, *CURRENT, *END
Severity . . . . . . . . . . . .   00            00-99
Output . . . . . . . . . . . . .   *             *, *PRINT
```

Figure 29. Displaying Migration Information

Using the MOVSPLFBRM Command

The Move Spooled Files using BRM (MOVSPLFBRM) command allows you to move selected spooled files to a specified library-qualified output queue. A run option of *REPORT is provided to allow the user to review the Move Spooled Files using BRM report prior to moving the the selected spooled files.

Note: All spooled files reside on a system ASP even if the output queue resides in a different ASP. To move spooled files to the ASP where the output queue resides, you need to specify *OUTQASP on the SPLFASP parameter on the CRTOUTQ command.

BRMS uses the ASP attribute of the spooled file for the specified output queue to determine whether or not spooled files will actually move from one ASP to another when moved from one output queue to another.

- If a spooled file that moves from one output queue to another does not actutally move from one ASP to another, BRMS does not check the ASP high storage threshold.
- If a spooled file that moves from one output queue to another would move from one ASP to another, then BRMS determines if the target ASP has sufficient space to accomodate the spooled file. BRMS does this without exceeding the high storage threshold as determined by the ASP high storage threshold attribute before moving the spooled file.
- If the spooled file cannot be moved without exceeding the threshold, BRMS shows that the file was not moved. BRMS includes it in summary section detail that indicates the number of files and the amount of data that could not be moved.

```
                  Move Spooled Files using BRM (MOVSPLFBRM)

Type choices, press Enter.

Option . . . . . .   *REPORT       *REPORT, *MOVE
To output queue  . . . . . .       Name
  Library  . . . . . .             Name
From output queue  . . . . . .     *ALL          Name, generic*, *ALL
  Library  . . . . . .               *ALL        Name, *ALL
File . . . . . .   *ALL             Name, *ALL
Job name . . . . . .   *ALL         Name, *ALL, *
User . . . . . .   *ALL             Name, *ALL
User data  . . . . . .   *ALL       Name, *ALL
From ASP . . . . . .   *ALL         Name, *ALL, *SYSTEM, 1, 2...
Select create date:
  From date  . . . . . .   *BEGIN   Date, *CURRENT, *BEGIN, nnnnn
  To date  . . . . . .   *END       Date, *CURRENT, *END, nnnnn
Select last used date:
  From date  . . . . . .   *BEGIN   Date, *CURRENT, *BEGIN, nnnnn
  To date  . . . . . .   *END       Date, *CURRENT, *END, nnnnn
                                                            More...
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

```
                  Move Spooled Files using BRM (MOVSPLFBRM)

Type choices, press Enter.

Select size:
  Size type  . . . . . .   *KB       *KB, *MB, *GB, *PAGES
  Minimum size . . . . . .   0       0-999999
  Maximum size . . . . . .   *NOMAX  0-999999, *NOMAX









                                                              Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 30. The Move Spooled Files using BRM (MOVSPLFBRM) command

This report uses printer file QP1AMSF.

Note: See “Chapter 7. Hierarchical Storage Management Backup Strategy with BRMS” on page 123 for considerations about performing backups after the movement of spooled files.

Chapter 4. Archival and Dynamic Retrieval

This chapter provides a description of how you implement hierarchical storage management with BRMS archiving (using save with storage freed) and dynamic retrieval (automatic recall of data). It then discusses various application design considerations that you should be aware of to aid in the design and implementation of your hierarchical storage management solution.

What Is Archiving?

Archiving in its original form, is a technique for saving valuable disk space by selecting infrequently used objects, saving them to tape, and then deleting them from disk. In previous releases, the concept of archiving using save with storage freed was introduced. The storage for the object is freed as part of the save operation. Archiving of data to tape or to a save file on secondary disk saves disk space on primary (fast) disk because it moves the object to a less expensive form of storage. BRMS tracks the information about the objects that you archive. When you recall the object by using dynamic retrieval, you restore the save file to the place from which it was saved. Full tape automation allows you to perform archiving quickly and effortlessly while appearing seamless in its operation.

The archival function allows you to archive objects to tape. However, only the following object types can be archived and dynamically retrieved:

- Database files, including QHST files (member level)
- Source files (member level)
- Document library objects
- Stream files

Archiving Limitations

Limitations for archiving with dynamic retrieval are:

- You cannot dynamically retrieve archived documents by using integrated file server commands.
- You cannot dynamically retrieve files in libraries by using integrated file server commands.

This list is not exhaustive. You should only use this list as a reference.

Archiving Considerations

You should develop standards or policies that control the objects being archived. These may include:

- **Application Structure - Is it suitable for archiving?**

  Your ability to archive can depend on your application design. Usually, archiving uses standard system commands such as SAVOBJ or SAVLIB. This requires that the object is accessible to such commands so individual records, for example, will not be able to be archived from a file until they are isolated in a separate object.
- **Length of time object has been inactive before archiving**

© Copyright IBM Corp. 1998, 1999

Dormancy criteria for archiving are the number of days of inactivity. If you set the value for the number of days too small, you will spend a large amount of time and resource archiving objects to tape. If the number of days is too large, there may be only very small space savings by implementing the archival process. For some very large objects, you may need to set the number of days to be zero. In the situation where you need to retrieve them, you can run a job against them to have the objects archived again as soon as possible.

- **Retention period**

  Some objects may be archived permanently for legal or taxation purposes. There is the expectation that you will not need these objects, but you need to track them in case they are required. Other objects may be staged on tape for deletion. As such, no object is ever deleted; rather it is archived to tape for a nominated period. If you make no restore request, the tape expires, and the object is 'lost' when you use the tape again. Save files also have an expiration date, and if you make no restore request, the save file expires, and the object is 'lost'.

- **Number of copies**

  If you are archiving data offline, you should make a copy of the tape. This will ensure that you have two copies to be **stored in different locations**. This is particularly important in case of media errors on one copy as there will not be a copy of the data on your AS/400.

- **Which objects will be "storage freed"**

  You need to decide whether to keep or delete the object description when you save the object with storage freed. The advantage of keeping just the object description for those objects is the ability for dynamic retrieval. Dynamic retrieval will be available only for objects that are archived in this way.

- **Maximum time to recover**

  As with any other aspects of providing a service to users of the system, you should develop service level agreements for your archive strategy. If data is not immediately available, how long will it take to retrieve? Users will be more willing to archive their data if they can be assured of a prompt retrieval service. They are more likely to cooperate if you can demonstrate that storing large amounts of data is less expensive on tape rather than disk.

- **Storage media (including ASPs for save files)**

  You can archive objects to an ASP as a save file, and then delete the save file after a period of time by setting the save file to *Expired. Also you can archive the objects directly to tape by using the SAVSAVFBRM command, or you can use a combination of both.

- **Length of time object will remain in a save file before copying to tape**

  If full staging is required, you might choose to archive to a save file on secondary (compressed) disk which could then be archived to tape. You can keep the save file after you copy the data for a small amount of time, a few days for instance, to more rapidly fulfill any restore request soon after the archive has been done. This would improve user satisfaction with having the data removed from the system.

- **Who will control the archiving?**

  If during the process of removing the data from the system, insufficient details are kept about the objects removed, it becomes a difficult task to manage and restore these objects. In this case, archival may require an operator dedicated to these tasks.

  If archival and restore have an easy-to-use interface, which tracks details of the object and the tape volumes to which they have been archived, this allows for greater flexibility as it can be used by most system users. Should the operations

staff remain the sole managers of the archive and restore process? Should knowledgeable end users be able to specify which of their objects are candidates for archival, and process the restore when the objects are required? Alternatively, should archival be run unattended? If you are using an automated tape library, you can run the archive and the restore without operator intervention.

- **When will archiving take place?**

  When is the best time of the day, or night, to archive data? If the objects are not being used, it does not really matter if the objects are archived during the day, except for the possible performance overhead of the save job. When you archive the data, you should specify a suitable priority.

How Does Archiving Work?

You can do archiving with OS/400 commands or with BRMS. BRMS uses standard OS/400 save and restore commands for archive and retrieval activity. The actual archiving of an object is achieved using a standard OS/400 save of the object with the “Storage” parameter set to *FREE. Within this publication this will be known as save with storage freed. BRMS tracks the information about the objects that you archive.

Use of Save with Storage Freed

Archiving works by saving objects, then deleting the entire object or just the object content from the system. The important characteristic of save with storage freed is that the object description is left on the system. This object description uses very little storage space. The description acts as a place-holder for the object in the system, while indicating that the data portion is on tape or in a save file.

Figure 31 on page 46 shows the makeup of an AS/400 object and how that object may look after being saved with storage freed. The object description contains only a small amount of data, which describes the object, including object name, object type, library name, security information, and so on. You find this information when processing commands such as Display Object Description (DSPOBJD), Display File Description (DSPFD), and Display File Field Description (DSPFFD). Even after save with storage freed, the object description remains in the original library for reference.

It is the data portion that contains all the real processable data (for example all the records in a physical file), and which therefore constitutes the majority of the object size. When objects are saved with storage freed, the data portion is deleted from the system after successful completion of the save.

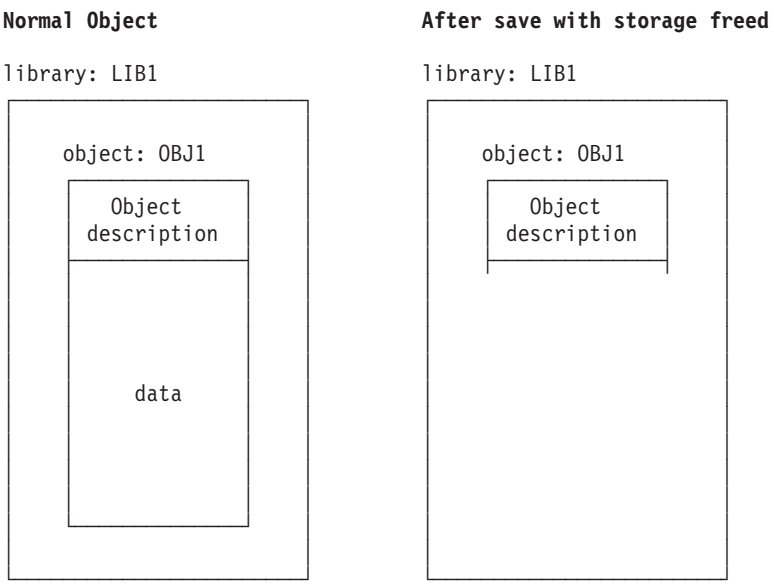

Figure 31. AS/400 Objects Before and After Save with Storage Freed

The *Storage* parameter on OS/400 save operations specifies whether to free system storage that is occupied by data. The data portion of the storage can include files, programs, Structured Query Language (SQL) packages, journal receivers, service programs, modules, and documents. The system frees only the data portion of the objects, not the descriptions of the objects. BRMS allows you to specify whether to include these types of objects as candidates for the archive process. If you elect to include them, you can specify that they should have freed storage, or you can delete the entire object.

BRMS processes the objects in two steps. First, it saves the object to the volume by using the STG(*KEEP) value to ensure that the media content information is successfully updated. Then it updates the media content information. If this completes successfully, and if you opted not to retain the object description, BRMS deletes the objects from the disk. If you chose to retain the object description, BRMS can perform another save of the same objects with the STG(*FREE) value specified. When this operation completes, BRMS deletes the save file.

BRMS allows you to specify whether to include other object types as candidates for the archive process. BRMS deletes the object once it successfully archives the object and updates the media content information.

When an object is accessed, the system starts a search for the object. See “Operations that Call or Do Not Call Dynamic Retrieval” on page 52 for more information on the types of operations that may or may not start a retrieve operation. Either the current job’s library list will be searched, or the object will be referred to with its library name, and the object description will be found. If the data portion is missing (saved with storage freed), BRMS checks its inventory of archived objects to see whether BRMS has archived the object . If the object is found in the BRMS archive inventory, then BRMS may then call the restore of that object.

Note:

If an object has been saved with storage freed by any means other than BRMS, an inventory of the archived objects is not performed. In this case BRMS cannot locate and restore the object without manual intervention. You receive an OS/400 message that indicates that the object was not found.

This situation can arise when you issue an OS/400 SAVOBJ command by using the STG(*FREE) parameter. Dynamic retrieval will not be possible for this object.

If the object that has been saved with storage freed is not one of the supported object types for the dynamic retrieval function, then BRMS will be aware that the object has been archived. BRMS can assist an operator to locate the correct tape. Because of the unsupported object type, automatic initiation of a restore operation is not possible. The user or job that was attempting to access the unsupported object will receive an error condition indicating that the object has been saved with storage freed — and nothing else. The user or operator must then consult the BRMS archive inventory to locate the object ready for a manually initiated restore. You can do this from the Work with Saved Objects (WRKOBJBRM) display or by using the Restore Object using BRM (RSTOBJBRM) command.

You may decide to write extra code into an application to support certain objects that are not supported for dynamic retrieval by BRMS. The additional application code would need to manage the types of OS/400 error messages returned for the objects required, and then be able to interrogate the BRMS inventory and initiate a BRMS retrieve operation.

The save with storage freed solution is very simple in its design and operation. The only alternative available to save with storage freed is to delete the object entirely. It may be considered impractical to use a solution that deletes the object entirely. The save with storage freed solution appears far more integral, secure, and simple.

Setting Up Archiving with BRMS

The objects that you select to archive are identified by entire auxiliary storage pool (libraries only), library, folder, stream files, or as lists of objects that are known as archive lists. A BRMS archive control group includes the ASPs, libraries, folders, stream files, or lists. Each control group has parameters that control such things as the amount of time the object must have been inactive to be selected for archive and so on. These parameters are known as the control group attributes. These details may also be set in the archive policy which sets the defaults to use, unless specifically overridden at control group level.

The important point to remember when setting up archive policies for dynamic retrieval in BRMS is to set the control group entries to allow the objects to be saved with storage freed. You can set this particular parameter as a default value in the archive policy and also in each individual control group's attributes. This parameter is explained in more detail later.

The BRMS archive control groups used for save with storage freed will also default to saving the access paths of the file members. This is a performance

consideration, and you may change it. This means that the object takes longer to save (archive) but eliminates the need for a potentially lengthy access path rebuild on restore (retrieve).

The process we will follow is:

1. **Identify ″sets″ of data suitable for archiving**

   Establish suitable archive criteria for each ″set″ of data, grouping by common characteristics. You can identify this set of data by entire ASP, library, folder, stream files or as lists of objects that are known as archive lists. Lists are groups of objects, folders, spooled files, or directories that are grouped for processing in an archive control group.
2. **Establish an archive control group**

   Control groups are lists of objects and libraries that share common archive characteristics. You can create multiple control groups and adjust the attributes for each control group where they vary from the archive policy.
3. **Establish an archive policy**

   The set of values in the archive policy serves as default values in the archive control group. You can override these values in the individual archive control groups. These values override the values that are set in the system policy.
4. **Create a media policy and incorporate special media classes**

   Operator-assisted archive provides information about tape media that are candidates for the operation and prevents the erroneous use of active media prior to expiration.
5. **Check the logs for results**

   Displaying archive activity provides access to information about archived data, estimated and actual direct access storage savings, activity logs, and archive media content information.

For more information, refer to “Setting Up Archive with Dynamic Retrieval Using BRMS” on page 56. For a full description of how to implement archiving with BRMS, refer to the *Backup Recovery and Media Services for AS/400* book .

Duplicating Your Archive Tapes

The implication of archiving an object is that it is saved and then deleted or storage freed in one operation. The ability to check whether the save to tape is successful before you delete it is reduced. Despite the extensive error checking and correction routines of modern tape device technology, the only true test of a successful save is to check whether the object can be read successfully in its entirety.

Also, it is possible that eventually all other copies of an object that have been saved in the normal backup procedure will expire, leaving only one copy (on tape) of the archived object. This will be the most up-to-date copy.

So the data loss exposure created by data archiving is two-fold:

1. There is limited verification of a successful save before deletion of the object.
2. There will eventually be only one copy of the object in existence.

For these reasons we recommend that you **make duplicate tape copies of archived data immediately** and then move them to an offsite storage location. To make duplicates of your tapes:

1. Type `WRKMEDIBRM SAVTYPE(*ARC)` from the command line and press Enter.
2. Select option 6 (Work with media) to work with the media volumes that are used in the archive operation.
3. Select option 14 (Duplicate data) to duplicate a tape or set of tapes. If a duplicate volume exists, you see an asterisk (*) character in the Dup column. The duplicate option takes you to the DUPMEDBRM command (depending on whether you selected an individual volume or a set). Use the tape automation to select scratch volumes that are automatically mounted for this process.

Here is another way to duplicate your tapes.

1. Specify a media policy that says Duplicate Media *YES.
2. Run the `DUPMEDBRM *SEARCH` command, and BRMS will find the tapes that need to be duplicated. This automatically marks the tape to be duplicated when it is created.
3. Assign a move policy to the output volumes so that they are moved to an offsite location. Do this after you complete the duplication process.

You can find more information on duplicating your archive tapes in the *Backup Recovery and Media Services for AS/400* book.

Dynamic Retrieval

This section details how retrieve operations are performed. The discussion includes batch and interactive controls as well as the device to be used and criteria for determining objects to be retrieved. Some objects that were archived through BRMS by using save with storage freed can now be retrieved automatically when the objects are needed (opened).

Note: To retrieve archived objects, you must have specified *YES in the *Retain object description* field in either the archive policy or the archive control group. For a description of the Retain object description field, see “Setting Up Your Archive Policy” on page 60.

How Does Dynamic Retrieval Work?

Figure 32 on page 50 shows the dynamic retrieval process.

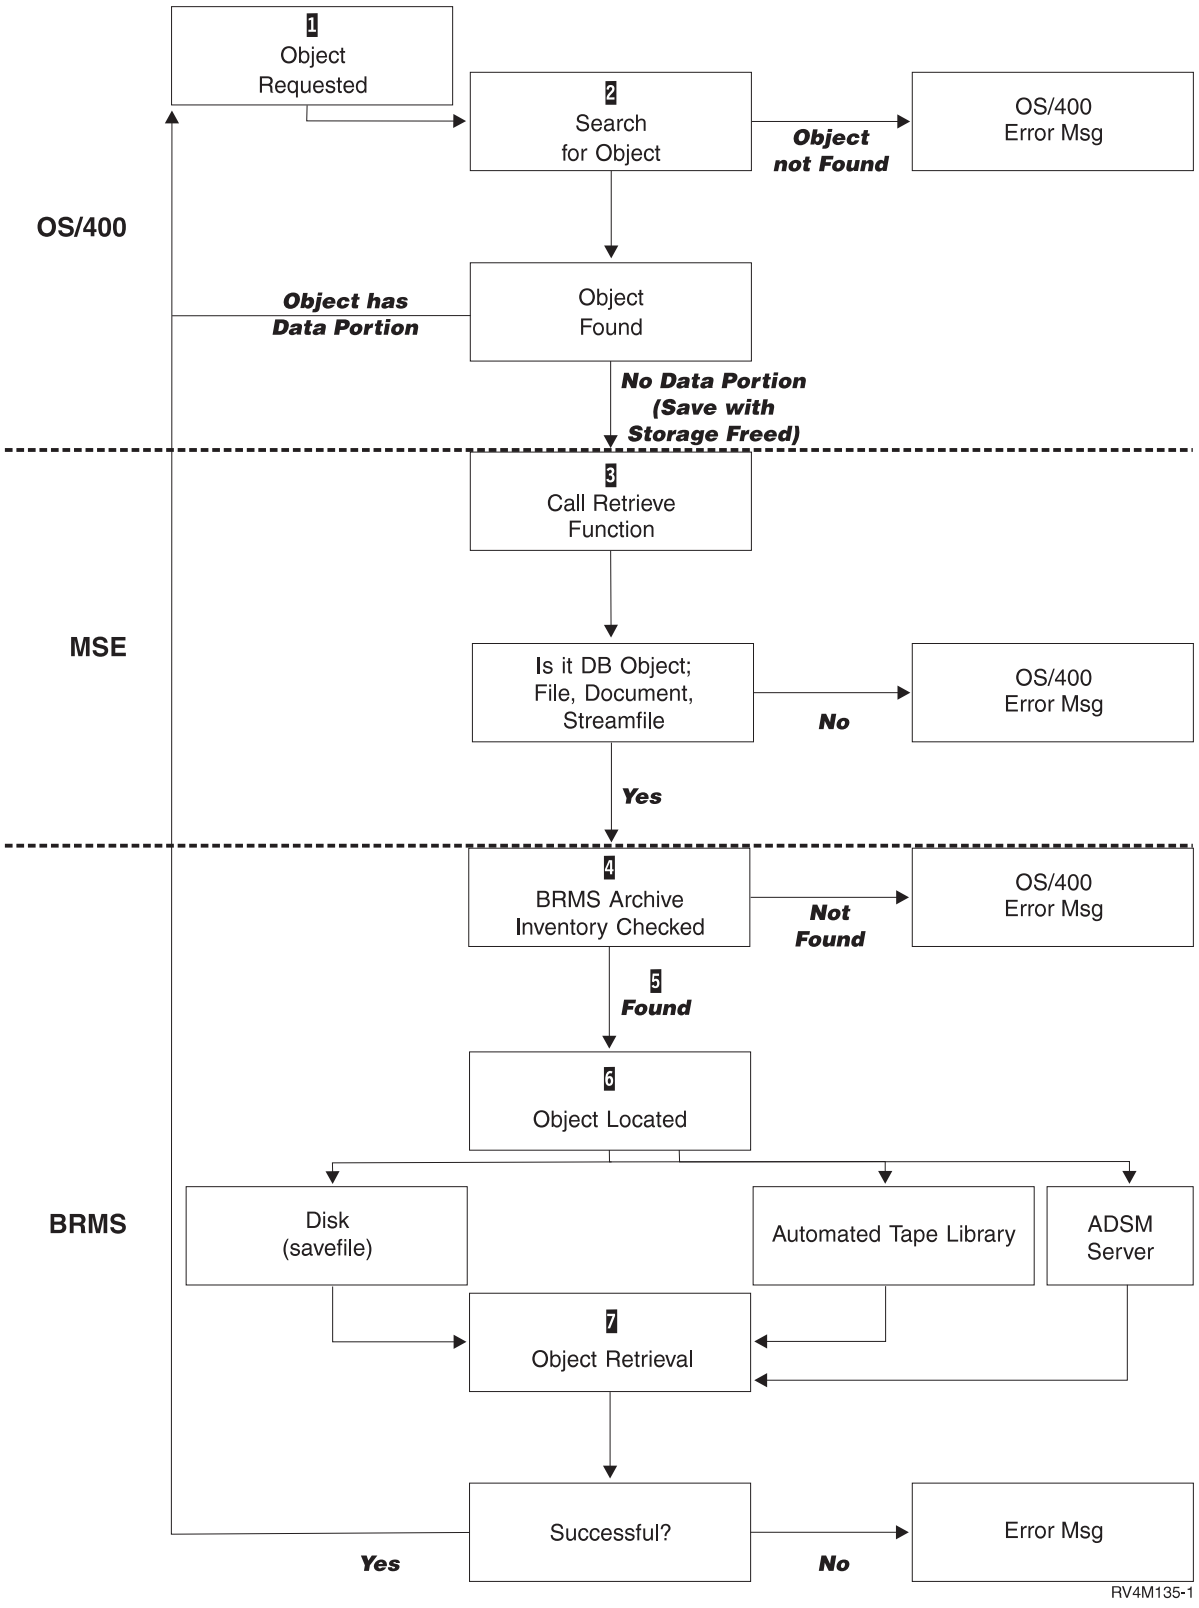

Figure 32. Dynamic Retrieval Process

These are the steps for the dynamic retrieval process:

1. Request an operation on an object.

Any operation that requires the polling of an object (see “Operations that Call or Do Not Call Dynamic Retrieval” on page 52).

2. Perform a search for an object.

   The library list for the current job is searched to locate the object description, or the object is located because the request was qualified with a library name.

   If the object description cannot be found, the process is ended here with an escape message. If the object description is found, but the requester requires the data portion to be present and it is not, the retrieve function is called.

3. Call the retrieve function, and check the validity of the object type.

   See “Operations that Call or Do Not Call Dynamic Retrieval” on page 52 for the types of operations that would call the retrieve function. Several object types that have the ability to be archived cannot be dynamically retrieved. Currently, database members, files, documents, and stream files are supported for dynamic retrieval. This is performed using MSE (Media and Storage Extensions), an optional feature of OS/400.

4. Check the BRMS archive inventory.

   MSE passes control to BRMS, and a search of the list of archived objects is performed. If the object is not found, the original OS/400 message is sent to the requester.

5. Notify the requester.

   Depending on the retrieve method and the type of job running, the user or the system operator message queue is notified of the intention to restore the object. See “Retrieval Methods” on page 81 for more details about the retrieve method.

6. Locate and restore the archived copy of the object.

   If the ’referenced’ object is found on tape and an automated tape library is present and operational, the tape may be loaded automatically. The object is restored in the normal way under BRMS control. Otherwise, an operator must respond to a mount message from BRMS.

   If the ’referenced’ object is found on disk as a save file, the object is restored to the referred to area. If the ’referenced’ object is found on the Adstar Distributed Storage Management (ADSM) server, the object is restored under BRMS control.

7. Requester operation continues as though the object had always been on the system.

   If the retrieve operation is completed immediately, the requester may continue business as usual (with only a slight delay). The reference to the object is automatically retried by OS/400 with the aid of MSE. This time it will work as normal.

   If the retrieve operation is not completed immediately, the requester is notified (effectively failing the operation), and must retry the operation at a later time. A media error or operator who takes the cancel reply to the tape mount message can cause a restore failure. Your application will need error handling so that it can recover from these conditions to allow a retry of the operation later when the completion message is sent.

   Note: If the operation fails because the object type is not supported by BRMS, you receive an OS/400 message. One cause might be that the object type is a *PGM type object that has been archived with storage freed. Another cause might be that the object is not in the BRMS archive inventory. The save may have been performed with storage freed outside

of BRMS. It will not be apparent that BRMS has been consulted at all. This situation will require standard application “object has had its storage freed” error processing.

Operations that Call or Do Not Call Dynamic Retrieval

The basic rule of thumb for understanding that what type of operation will initiate the retrieve function revolves around what is known as an “open”. This function may be an attempt to access the data portion of an object. This table contains a list of operations that call the dynamic retrieval function. This table is **not** exhaustive. Use this table only as a starting point.

Table 1. Dynamic Retrieval Operations that Call or Do Not Call Dynamic Retrieval

Dynamic Retrieval type	Operation or Command	Call Dynamic Retrieval	Does not Call Dynamic Retrieval
Database	Any type of database open whether explicit (running a CL command) or implicit (startup of a program). Any operation qualifies that sets up a database file or member for read or update by setting an open data path in the job’s process access group.	X	
	DSPOBJD - Displaying the object description only touches the description part of the object and, therefore, no attempt is made to access the data part that is freed.		X
	OPNDBF (Open Database File)	X	
	OPNQRYF (Open Query File)	X	
	DSPPFM (Display Physical File Member)	X	
	RTVMBRD (Retrieve Member Description)	X	
	DSPFD and DSPFFD - Display of the file description or file field description fetches only the file description data.		X
	CPYF (Copy File)	X	
	SNDNETF (Send Network File)	X	
	CHGOBJD, CHGOBJOWN, or CHGOBJAUD - Changing of the object description, owner, or audit level affects only the object description.		X
	CHGPFM - Changing the member information affects only the file description.		X
	RNMOBJ and RNMM - Renaming objects (including libraries) or physical file members has more considerations than not just accessing the actual data. Rename does not call a retrieval because it will not touch the object data portion. However, renaming may prevent a retrieve from ever happening again as BRMS refers to only the object and member names for retrieve operations. Renaming breaks the link between the object description (on the system) and the object data (on tape).		X
	MOVOBJ - Object movement commands do not cause a database retrieval, although similar restrictions apply when using MOVOBJ with archived objects as with RNMOBJ and RNMM.		X

Table 1. Dynamic Retrieval Operations that Call or Do Not Call Dynamic Retrieval (continued)

Dynamic Retrieval type	Operation or Command	Call Dynamic Retrieval	Does not Call Dynamic Retrieval
	CHKOBJ - Check object verifies only the existence of the object and verifies the authority of the user to the object before trying to access it. This does not involve reading any data records.		X
	ADDPFM and RMVM - Adding or removing members affects only the member attributes of a file that are stored in the file description. When adding a member, the new member description joins the others within the file as they have not been deleted by the archive with save with storage freed. A retrieve of other members does not affect any new ones. When removing a member, the member description is deleted and BRMS will no longer be able to retrieve it because the reference (by name) in BRMS to the archived tape copy no longer exists. A file member that is removed no longer initiates a dynamic retrieval operation if an open is attempted for it.		X
	CHGPF - Changing physical file attributes affects the object description which may or may not cause a dynamic retrieval to occur.	X	X
	DLTF - When deleting a file, all members are removed and no retrieve is necessary. Of course, this means that if access to this file is attempted later, it fails.		X
	RCLSTG - Reclaim storage operation requires access to the file member but bypasses the retrieve operation.		X
	DSPLOG - Even though it seems logical that the system history log files are needed when using the Display Log (DSPLOG) command, any of these files that have been archived with storage freed are simply bypassed, as if they did not even exist.		X
	CRTxxxPGM - Any program compile that refers to an archived database file does not actually access the data which may or may not cause a dynamic retrieval.	X	X
	CRTxxxPGM - Program source code in a source file member requires access to the data (source statements) and will be retrieved if a compile is performed on it.	X	
	Query/400 - Either the “Specify File Selections” part of defining a query or the actual running of that query over a file initiates a retrieval.	X	
	SQL - File selection during an interactive SQL query setup or running SQL statements on a file initiates a retrieval.	X	
	DFU - The Data File Utility (DFU) initiates a retrieval when performing a file selection during a temporary or permanent DFU program build, or while starting a DFU program.	X	
	Options from PDM - When using Programming Development Manager (PDM), certain actions on file members may or may not cause a dynamic retrieval to occur.	X	X

Table 1. Dynamic Retrieval Operations that Call or Do Not Call Dynamic Retrieval (continued)

Dynamic Retrieval type	Operation or Command	Call Dynamic Retrieval	Does not Call Dynamic Retrieval
	Client Access file transfer - Client Access file transfer calls a normal database open prior to the transfer of records down to the PC. However, it initiates a dynamic retrieval operation based on the retrieval method previously defined.	X	
	Start or end journaling - Starting or ending journaling of a file does not touch the data at all. Only the receiver, journal, and object description are updated.		X
	Journal changes - When applying and removing journal changes to or from an archived file, the file is automatically retrieved. It is not necessary to open the file in preparation. A consideration when applying or removing journal changes is when the journal entry for the storage free operation is included in the block of sequence numbers to be processed in the apply or remove. Typically this occurs during an RMVJRNCHG operation where the “Starting sequence number” parameter contains *LAST and the “Ending sequence number” parameter contains a number that is related to the point in time to which the file changes must be rolled back. The journal entries in this range include the BRMS archive (and storage free) operation. An RMVJRNCHG command cannot roll back this type of operation and will fail. You need to use the Display Journal (DSPJRN) command to select a range of journal entries that do not include the storage free operation. When applying journal changes to multiple files, the same problems that are listed in “Multiple Physical Files behind a Logical File” on page 90 occur. Even though you know exactly which physical file members you will be using, you still do not know which ones will need retrieving. If prediction of performance is really critical, you may attempt a strategy that opens all the required files first, to check that they are online.	X	
Document Library Objects (DLO)			
	Viewing a DLO	X	
	Printing a DLO	X	
	Copying a DLO	X	
	Moving a DLO	X	
	Changing details of a DLO	X	
	Changing contents of a DLO	X	
	Integrated file system or hierarchical file system (HFS) open APIs		X
	Security changes of a DLO (ADDDLOAUT, CHGDLOAUT, EDTDLOAUT, DSPDLOAUT, RTVDLOAUT)	X	
	Renaming a DLO	X	
	Direct existence of a DLO (CHKDLO, DSPDLONAM, RTVDLONAM)		X

Table 1. Dynamic Retrieval Operations that Call or Do Not Call Dynamic Retrieval (continued)

Dynamic Retrieval type	Operation or Command	Call Dynamic Retrieval	Does not Call Dynamic Retrieval
	Indirect existence of a DLO (DIR, WRKDOC, DSPFLR)		X
	Search of a DLO object		X
	Save of a DLO object		X
	Restore of a DLO object		X
	Reclaim of a DLO object		X
	Reorganize a DLO object		X
	Deletion of a DLO object		X
	Dump of a DLO object		X
Stream Files			
	Checking out the object (CHKOBJ)	X	
	Accessing the data of the object (open(), creat(), MOV, CPY, or CPYFRMSTMF, CPYTOSTMF)	X	
	Adding a new name to the object (RNM, ADDLNK, link(), rename(), Qp01RenameKeep(), or Qp01RenameUnlink())	X	
	Indirect existence of a stream file - Stream File Objects		X

Additional dynamic retrieval considerations are:

- Many CL commands can be set up in a simple CL program with corresponding Close File (CLOF) commands if you wish to retrieve a number of files together before an application starts.
- A database open that uses a DDM file will qualify for dynamic retrieval on the remote (target) system. However, it initiates a dynamic retrieval operation that is based on the Retrieve Policy. See “Setting Up the Retrieve Policy” on page 76 for more information.
- Certain PC applications sometimes use an open of a file for simply an existence check. There are no distinguishing features for this opening of a file so disabling dynamic retrieval is not possible.

Re-Archiving Retrieved Objects

In some cases a collection of data that has been archived and then retrieved may subsequently need to be re-archived differently to data that has not been archived at all.

Example:

You have a parts inventory system. You have archived the parts list file from last year because your catalog of parts has been refreshed for this year. The old catalog will rarely be used again, but while you are running down your stocks from the old catalog you may want to keep it online. After a continuous 90 days of inactivity, we can be sure that it is not being used, and it can be automatically archived. Then a customer calls and asks for a discontinued part. Your system tells you that it is discontinued. The customer is desperate for this part. You decide to check last year's catalog to see if you have any of these really old parts still lying around.

At this point you retrieve the old catalog file. Having performed the search using the old catalog you are confident that you will probably not need this catalog again unless some really exceptional circumstances occur. Why wait for another 90 days for it to archive? Why not re-archive this *retrieved file* after 5 days of inactivity?

BRMS supports archiving at multiple levels that are based on the date last used, change date, or both (whichever is the later). There is no differentiation between objects that have never been archived and retrieved objects.

If you wish to enable such function, create your own program that works with the BRMS retrieve exit point to generate a list of retrieved objects. You may have your own special archive control group that has a different inactivity level specified from the regular control group. The list of retrieved objects could be used against a list of candidate objects for this special treatment to generate the list of objects to be included in your special control group.

You may also attempt to create a method of differentiating between objects that have been retrieved for read only purposes and those that are retrieved for update. BRMS does not differentiate between these conditions. If you are sure that the file member has not been updated, you may consider just freeing the storage of the file member after the necessary dormancy period has elapsed. This requires that you save to a temporary save file with STG(*FREE) and then delete the save file. To do this, establish an instant lock on the file member to allow only read transactions. You can also set authority such that no one could update it (and be sure that no one signed on with *ALLOBJ authority). This is clearly not a simple task. You cannot use the last updated date because the restore operation changes this when you retrieved the file member. Because it is only a date (and not a time), there is no way of judging whether an update took place on the same day as the restore.

Setting Up Archive with Dynamic Retrieval Using BRMS

The first part of archiving for dynamic retrieval is the identification of candidate objects that you wish to archive. This can be done at differing levels, depending on how granular you want to be. The levels are:

- ASP level (Only libraries will be archived)
- Library, folder, or directory level
- Object level (Files, members, DLOs, and stream files)

When you archive by ASP or library, folder level and directory level give you the opportunity to specify a large number of candidates without much typing. You can

archive by ASP, specifying an ASP number as a special value in the archive control group, for example, *ASP03. When you specify *ASP03, only libraries in the specified ASP are included. No other object types are included. If you archive by library, all members in objects in the library will be archived. If you archive by folder, all documents in subfolders in the folder will be archived. Likewise, if you archive by directory, all stream files in subdirectories in the directory will be archived. These levels are also useful in producing Archive Candidate reports, which assist in estimating potential space savings.

Typically object and member-level archive candidate selection is most often used. With these, you will have to group these candidate file members into lists of objects with something in common. You may make good use of the generic selection functions within BRMS if (for example) the file members are all in the same file.

Archive lists consist of grouping objects that have similar archival and retrieval characteristics into individual lists. These lists will be placed into a BRMS archive control group. It is the control group that specifies the archiving parameters (such as dormancy levels and media policies). Thus it is feasible to split groups of objects that currently have similar archive and retrieve characteristics into several individual archive lists that allow additional flexibility for future changes.

This topic discusses setting up the necessary archiving details in BRMS to allow the implementation of dynamic retrieval by:

1. Setting up archive lists
2. Setting up an archive policy
3. Setting up archive control groups

Setting Up Archive Lists

Archive lists contain groups of items, such as objects or folders, that share similar characteristics. You locate the lists inside an archive control group, where BRMS processes them during the archive operation.

Archive lists can consist of the following:

- *OBJ for files, members
- *FLR for folders, documents
- *LNK for stream files
- *SPL for spooled files

Notes:

1. Though you can create a spooled file list to archive, you should be aware that HSM does not currently support the retrieval of spooled files.
2. You must archive *LNK lists in directories other than QSYS.LIB and QDLS.

The lists you create can contain one type of item (for example, just object (*OBJ) lists. Or it can contain a variety of types (*OBJ, *FLR, *LNK). You can use the Work with Lists (WRKLBRM) command to add, change, view, copy, or print a list.

Adding an Object List

To create an archive list, take the following steps:

1. Type WRKLBRM at a command line and press Enter. This takes you to the Work with Archive Lists display.

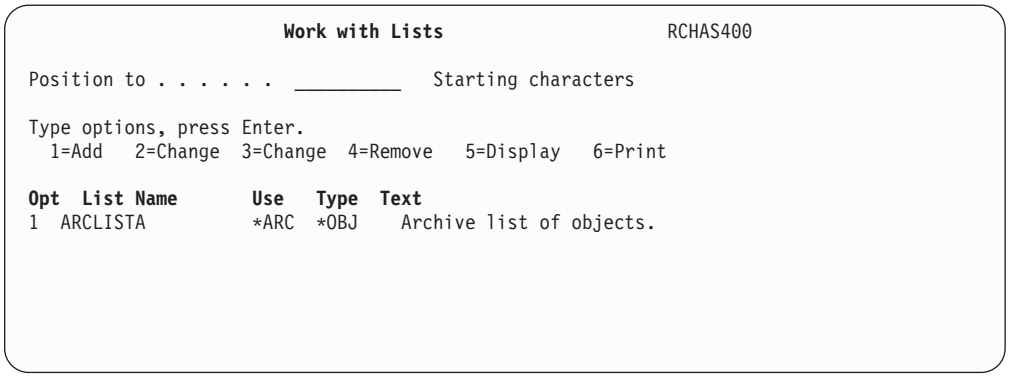

Figure 33. Work with Lists

2. Type a 1 (Add) in the *Opt* field.
3. Type the name of the list you want to create in the *List name* field.
4. Type *ARC in the *Use* field.
5. Indicate the type of list you want to create in the *Type* field.
6. Add descriptive text about the list you want to create in the *Text* field.
7. Press Enter to get to the next display. On the previous display, you indicated that you want to create an object (*OBJ) list. Thus, in this case, pressing Enter takes you to the Add Object List display.

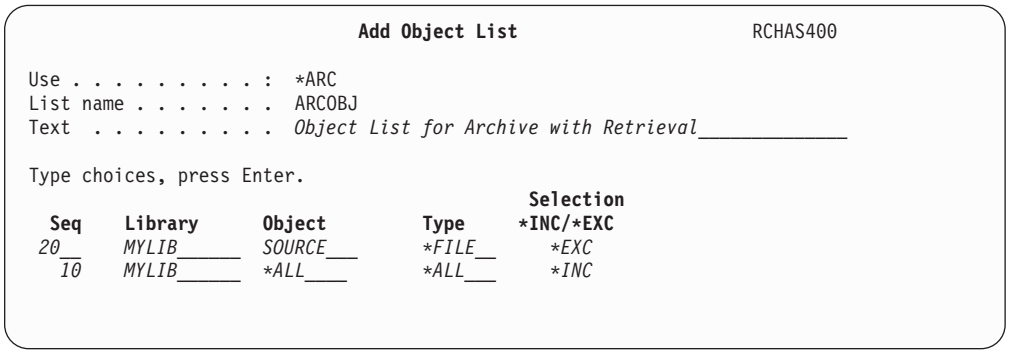

Figure 34. Archive Lists

8. At this display, type the names of the libraries and objects you want to archive in the *Library* and *Object* fields.
9. Type the order in which you want them to be archived in the list in the *Seq* field.
10. In the *Selection *INC/*EXC* field, indicate whether you want to include or exclude the library or object from the archive.
11. Press Enter twice to save the list and return to the Work with Archive Lists display.

    Note: If you press F3 or F12, you lose all the changes you just made.

Note: To include a complete library of objects as candidates, type the library name as an entry in an archive control group rather than including it in an archive list. If you include the library in an archive list, you then need to add an entry in an archive control group, which is a two-step process. Remember that if

you include whole libraries or ASPs, you could easily archive object types, such as programs, logical files, journal receivers, and so on that are not supported for dynamic retrieval. These objects can archive without problems, but cannot be dynamically retrieved. When you try to retrieve an unsupported object, the system simply reports that it saved the object with storage freed, and the user program ends in error. For a list of objects that can be archived for dynamic retrieval, see “What Is Archiving?” on page 43.

Adding a Folder List

To add a folder list, take the following steps:

1. Follow steps 1-4 in the “Adding an Object List” on page 57 section.
2. In the *Type* field, specify *FLR to add a folder list.
3. Press Enter to get to the Add Folder List display.

```
                              Add Folder List                  RCHAS400

Use . . . . . :  *ARC
List name . . . . . .  ARCLISTB
Text  . . . . . .

Type choices, press Enter.                                      Sel Subflr
                                                               *INC/ *YES/
Seq  Folder                                                     *EXC  *NO

10   GNTEST                                                     *INC  *YES
20   GNTEST/GNSUB1/GNSUB2/GNSUB3                                *EXC  *NO









F3=Exit    F5=Refresh    F12=Cancel
```

Figure 35. Add Folder List display

4. On this display, indicate the name of the folder you want to archive in the *Folder* field.
5. Specify the order in which you want BRMS to archive the folder in the *Seq* field.
6. Indicate whether you want BRMS to include or exclude the folder in the *Sel *INC/*EXC* field.
7. Indicate whether you want BRMS to archive any subfolders within the folder in the *Subflr *YES/*NO*.
8. Press Enter twice to add the list and return to the Work with Archive List display.

Adding a Link List

Use a link list to save Integrated File System objects. Creating a link list takes two steps. First you add the link list, and then you define the path names.

1. At the Work with Lists display, type in the name of the link list you want to create, and specify *LNK in the type field. Then press Enter. This takes you to the Add Link List display.

```
                          Add Link List                    RCHAS400


List name . . . . . .  ARCLNK
Text  . . . . . .  Link List for Archive with Retrieval_____________________
```

Figure 36. Add Link List display

2. Type in a name and a description of the link list you want to create and press Enter. This returns you to the Work with Lists display.
3. To define the path names, type a 2 (Change) in front of the link list you just created and press Enter. This takes you to the Change Link List display.

```
                          Change Link List                 RCHAS400

List name . . . . . .  ARCLNK
Usage type . . . . . . *ARC   *BKU, *ARC
Objects:
  Name  . . . . . :  MYDIR/SUBDIR
  Include or Exclude. . . . :  *INCLUDE

Directory Subtree . . . . . :  *ALL

Text  . . . . . :  Link List for Archive with Retrieval__________
```

Figure 37. Change Link List (CHGLNKLST)

4. In the *Usage type* field, specify *ARC for archive lists, and *BKU for non-archive list. The job will fail if the value in this field is not correct.
5. Define the path name in the *Name* field, and indicate whether you want to include or omit the objects from the archive list.
6. Review and change the other parameters as needed.
7. Press Enter twice to save your changes and return to the Work with Lists display.

Setting Up Your Archive Policy

The archive policy is the system-wide default set of controls that governs the behavior of all archive control groups when you run them. You can override any of the parameters in the archive policy within each individual control group (using the control group's attributes), but the archive policy serves as a system standard.

You can access the archive policy directly by using the `WRKPCYBRM *ARC` command. Figure 38 on page 61 shows the parameters that you can set.

```
                    Change Archive Policy                      RCHAS400


Type choices, press Enter.

Media policy . . . . . . . . . . . . . . *SYSPCY___  Name, F4 for list
Archive devices  . . . . . . . . . . . . *SYSPCY___  Name, F4 for list
                                         __________
                                         __________
                                         __________
Include:
  ASP storage threshold. . . . . . . . . 0__          *ASP, 0-99 %
  Date type for archival . . . . . . . . *BOTH__     *BOTH, *CHANGE, *USE
    Inactivity limit . . . . . . . . . .  365    0-9999 days
  Object size greater than . . . . . . . 0__          0-999999 MB
  Objects able to be freed . . . . . . . *YES        *YES, *NO
    Retain object description  . . . . . *YES        *YES, *NO
  Objects not able to be freed . . . . . *YES        *YES, *NO
  Days used (per month) less than. . . . 0           0-31
```

```
                    Change Archive Policy                  RCHAS400


Type choices, press Enter.

Default weekly activity  . . . . . . . . *******     SMTWTFS(*)
Sign off interactive users . . . . . . . *SYSPCY     *YES, *NO, *SYSPCY
Sign off limit . . . . . . . . . . . . . *SYSPCY     0-999 minutes, *SYSPCY
Save access paths  . . . . . . . . . . . *YES        *YES, *NO
Save contents of save files  . . . . . . *YES        *YES, *NO
Data compression . . . . . . . . . . . . *DEV        *DEV, *YES, *NO
Data compaction  . . . . . . . . . . . . *DEV        *DEV, *NO
Target release . . . . . . . . . . . . . *CURRENT    *CURRENT, *PRV
Clear  . . . . . . . . . . . . . . . . . *NONE__     *NONE, *ALL, *AFTER
Append to media  . . . . . . . . . . . . *NO_        *YES, *NO
End of tape option . . . . . . . . . . . *UNLOAD     *UNLOAD, *REWIND, *LEAVE
Use optimum block size . . . . . . . . . *DEV        *DEV, *YES, *NO
                                                                More....
```

```
                    Change Archive Policy                      RCHAS400


Type choices, press Enter.

IPL after archive  . . . . . . . . . . . *SYSPCY     *YES, *NO, *SYSPCY
 How to end  . . . . . . . . . . . . . . *SYSPCY     *CNTRLD, *IMMED, *SYSPCY
 Delay time, if *CNTRLD  . . . . . . . . *SYSPCY     Seconds, *NOLIMIT
 Restart after power down  . . . . . . . *SYSPCY     *YES, *NO, *SYSPCY
 IPL source  . . . . . . . . . . . . . . *SYSPCY     *PANEL, A, B, *SYSPCY
Retrieved object retention . . . . . . . *NOMAX      0-9999, *NOMAX
                                                                  Bottom
```

Figure 38. Change Archive Policy Display

This example shows the default parameters that are based on the system policy. Use the following recommendations for all archive control group attributes that will archive objects for use with dynamic retrieval. If you do not perform any other kind of archiving, follow these recommendations for the archive policy, and allow all your archive control groups to default to the archive policy. If non-retrieve archive control groups also exist, you must take care with this approach. It may be best to specifically set all these parameters within each archive control group and leave the archive policy as general as possible.

Refer to the online help for a description of all the parameters that are shown. However, for the purpose of dynamic retrieval, take particular note of the **Include, Default weekly activity,** and **Save access path** parameters.

Include parameter

The *Include* parameter contains the following elements: ASP storage threshold, Data type for archival, Inactivity limit, Object size greater than, Objects able to be freed, Retain object description, and Objects not able to be freed. Following are brief summaries of each parameter.

ASP storage threshold
: Specifies the percent storage utilization that if exceeded, will allow objects that meet all other archive criteria to be archived. Setting this percentage value high suppresses archive operations from an ASP. When the storage utilization of an ASP is less than this percent, objects in this ASP are not archived. The special value, *ASP, is provided to refer to BRMS's ASP descriptions. A value of *NONE means that no objects from the ASP will be archived.

Date type for archival
: For objects you can be more specific about the date that you wish to use for dormancy qualification. This parameter allows you to specifically use the *last used date*, the *last changed date*, or both dates to determine the number of days of inactivity.

Inactivity Limit
: Specifies the number of days of inactivity before an object is archived. The inactivity limit uses the most recent date, whether change or use date, in determining candidates for archive. For example, you may have changed an object a year ago, but displayed it for review 30 days ago. In this case, the last change date is 4/30/98 and the last use date is 3/30/99. If you run the archive system on 4/30/99 with an inactivity limit of 60 days, then the object would not archive because its last use date (3/30/99) was less than 60 days ago.

Object size greater than
: Specifies the size in megabytes that an object must exceed before it can be a candidate for archival. Only objects that meet this size specification and all other archive criteria will be archived.

Objects able to be freed
: Specifies whether to include objects whose storage can be freed during an archive operation. Object type *FILE, *PGM, *JRNRCV, *SQLPKG, *DLO, and stream files have descriptions that you can keep after an archive operation has freed the storage taken by an object. You cannot keep descriptions for other types of objects. To use the BRMS dynamic retrieval function, you must specify ***YES** to this parameter to enable the save of objects with storage freed.

Retain object description
: You must specify ***YES** to this parameter if you want to use the dynamic retrieval function. This parameter initiates the save with storage freed operation. If you specify ***NO**, you delete the object description and the object when the object is archived.

Objects not able to be freed
: It is advisable to specify ***NO** for this parameter. The *NO value ensures that it is less likely that you will archive an object that cannot be retrieved dynamically. Remember that there still may be a significant number of

objects that can be saved with storage freed, but the BRMS dynamic retrieval function does not support them.

Days used (per month) less than
Specifies whether to include objects in the archive operation based on the frequency of object use. If you specify zero in this field, then BRMS ignores this attribute in the selection criteria for archive.

Default Weekly Activity

The seven separate fields in the *Default weekly activity* field indicate the days of the week on which you want BRMS to perform archive activities. The value in this field on this display is the default weekly activity value that is used for all archive activities unless changed by you. The default value instructs BRMS to process an archive each of the seven days of the week (SMTWTFS). You can change this value on the Archive Policy display, or on one of the archive control group displays.

You can run the archives less frequently than your backups. When you do this, you may increase the availability of the system, or you may increase the amount of data that you archive in one operation. You do not necessarily need the system in a quiesced state to perform archiving. By definition, the objects to be archived will not be in use. Special conditions may apply (for example immediate archive, or an exclusive lock on the entire library prevents objects from archive). The more data that is archived in one run (and therefore on a single set of tapes), the more likely that fragmentation of tapes will occur. If you keep archiving frequently, you will have only a few objects on each tape. This, in turn, could lead to a great deal of tape space that is wasted.

Save Access Paths

Specify ***YES** to this parameter for all objects that may be retrieved. This is so that a lengthy access path rebuild phase does not inhibit the performance of the retrieve operation (especially if in *NOTIFY or *VERIFY mode).

The main recommendation for your archive activity plan is that you **commence archive immediately after your backups have completed**. This will minimize the impact of archive media errors because a backup copy of the data exists.

Setting Up Archive Control Groups

After you create your archive lists, you are then in a position to build some archive control groups. You will have set your archive policy to reflect the most desirable run time options for all of your archive control groups. If the only archiving that you are performing is archiving for dynamic retrieval, it is possible that you have set the archive policy to be most suitable for your archive needs. However, you can override these values at a control group level.

To create an archive control group:

1. Type `WRKCTLGBRM *ARC` from the command line.

```
                    Work with Archive Control Groups              RCHAS400

Position to . . . . . _________    Starting characters

Type options, press Enter
  1=Create    2=Edit entries    3=Copy        4=Delete   5=Display
  6=Add to schedule   8=Change attributes   9=Subsystems to process ...
                                        Weekly
    Control     Inactivity  Media       Activity
Opt Group         Limit     Policy      SMTWTFS  Text
1   MAINARC
_   *ARCGRP      *ARCPCY    *ARCPCY     *ARCPCY  Entry created by BRM configura
_   TELESALES    *ARCPCY    *ARCPCY     *ARCPCY  Telesales order files
_   TESTFILES         30    TESTARC           *  Test files in work libs
_   HISTLOGS           7    *ARCPCY           *  History Logs
```

Figure 39. Work with Archive Control Groups Display

2. The default control group (*ARCGRP) contains values that are established in the archive policy (*ARCPCY). Press F9 from the Work with Archive Control Group display to see these values. Select option 1 to create a new control group.

```
                  Create Archive Control Group Entries            RCHAS400

Group . . . . . : MAINARC
Default activity  . . . . . *ARCPCY
Text  . . . . . . . . . . . Main Archive Control Group_____________________

Type information, press Enter.

                                Weekly
        Archive       List      Activity
Seq     Items         Type      SMTWTFS
40__    ARCLNK_____   *LNK
  10    ARCOBJ_____   *OBJ      *DFTACT
  20    ARCFLR_____   *FLR      *DFTACT
  30    STOCKLIST__   *OBJ      * * * *
```

Figure 40. Create Archive Control Group Entries Display

3. Type the names of the archive lists as previously created, or type names of libraries, folders or directories that you regard as suitable candidates for archiving. When you specify *ASPnn, leave the List Type blank. Press F19 to display a list of libraries. In each case you are simply listing available *candidates* for archive. The selection of objects that actually does get archived is performed at run-time, using the list of archive candidates (entered here) and the Include Criteria specified in the control group attributes (Option 8) or the archive policy.
4. Specify the days in which the archive is to be performed for the Weekly Activity. If you do not specify a value, this field uses the default value of *DFTACT, which is taken from the default activity specified on the control group attributes display. Each of the control group's attributes may also default to the archive policy, and some of these may also default to the system policy.
5. Press Enter after you have completed entering all of the items in your control group. Then press F3 to save the control group entry and to return to the Work with Archive Control Groups display.

6. We suggest that you print a report of the available candidates for archiving within the control group by using the STRARCBRM command.
   a. Type `STRARCBRM` from the command line and press F4.
   b. Type in the name of your archive control group. The default value for the run option is *REPORT.
   c. Press Enter to generate the archive candidate report. The archive candidate report is a list of objects that *would have been* archived if you had taken the run option of *ARCHIVE. You cannot print the report if none of the files in your lists are dormant yet. In this case you receive a message that says that the report did not contain any data. Check the entries in this report against your expectations.

Option 6 (Add to schedule)

Select option 6 from the Work with Archive Control Groups display allows you to send an archive control group to the job scheduler for processing. You can select the values for the frequency, date and time parameters. The archive control group has now been placed into the OS/400 job scheduler for automatic processing. To work with scheduled jobs, press F7 (Work with BRM scheduled jobs) to see the BRM jobs currently in the job scheduler.

Option 8 (Change attributes)

Select option 8 from the Work with Archive Control Groups display to change any of the attributes of this control group.

```
                  Change Archive Control Group Attributes

Group  . . . . . . . . . . . . . . . . : SAMPLE

Type information, press Enter.

Media policy . . . . . . . . . . . . . . SAMPLE      Name, F4 for list
Archive devices  . . . . . . . . . . . . TAPMLBBRM   Name, F4 for list



Parallel device resources:
  Minimum resources  . . . . . . . . . . *NONE       1-32, *NONE, *AVAIL
  Maximum resources  . . . . . . . . . .             1-32, *AVAIL, *MIN

                                                                    More...
F3=Exit   F4=Prompt   F12=Cancel
```

Figure 41. Change Archive Control Group Attributes Display

The parameters on the Change Archive Control Group Attributes display are identical to the parameters on the Change Archive Policy display, with one exception (the parallel device parameters). You can change any of the attributes on this display to override the values set in the archive policy. The changes you make on the Attributes display apply only to the control group you are working with at that time.

You can find more information on how to use the new parallel device support parameters in the Additional Options for Processing Your Archive Control Groups section of this chapter. Refer to the online help for descriptions of the other parameters or see “Setting Up Your Archive Policy” on page 60 for more information.

To archive to a save file, specify a media policy that defines the save file information and specify *NONE in the *Archive device* parameter.

Option 9 (Subsystems to process)

You can specify subsystem information by control group from the Work with Archive Control Groups display. Use the Subsystems to Process display to add or remove subsystems that you want to start or end automatically before and after control group processing. In this example, BRMS ends the subsystems at the start of one control group (SAMPLE) and restarts them after saving another control group (SAMPLE2).

To get to the Subsystems to process display, take the following steps:

1. On the Work with Archive Control Groups display, place a 9 (Subsystems to process) next to the control group you want to work with.
2. Press Enter.

```
                        Subsystems to Process                  RCHAS400

Use . . . . . :  *ARC
Control group . . . . :  SAMPLE2


Type choices, press Enter.
                              End
 Seq    Subsystem     Library   Option     Delay       Restart
 10     QBATCH        *LIBL     *NONE                  *YES
 20     QCNTL         *LIB3     *CNTRLD    300         *YES




                                                        Bottom
```

Figure 42. The Subsystems to Process display

3. Type a number that reflects the order in which you want the subsystem job processed in the *Seq* field.
4. Type the name of the subsystem you want to process in the *Subsystem* field.
5. Type the name of the library that contains the subsystem in the *Library* field.
6. Specify whether you want BRMS to automatically restart the subsystem at the end of control group processing in the *Restart* field.

   Note: If you do not specify an IPL, BRMS automatically restarts the subsystems after save processing completes. If you do specify an IPL, BRMS restarts subsystems after an IPL.
7. Review and change the other parameters as needed.
8. Press Enter to save and to return to the Work with Archive Control Groups display (if you press F3, you exit without saving).

To delete an entry from the Subsystems to Process display, position the cursor on the item you want to delete. Then space through the sequence number, and press Enter.

Before processing your subsystem requests, you need to ensure that the media policy for the control groups you want to process is appropriate for this request.

Note: You can run multiple control groups serially (one after another) when using the subsystems to process feature. The first control group ends the subsystems, and the last one restarts them.

Option 10 (Job Queues to process)

Use the Job Queues to Process display to add or remove job queue items from a list of job queues. These are job queues you want to hold or release before and after control group processing. Take the following steps to get to the Job Queues to Process display:

1. Select F23 (More options) from the Work with Archive Control Groups display.
2. Place an option 10 (Job queues to process) next to the control group you want to work with and press Enter.

   In this example, the *ARCGRP control group specifies one job queue to process.

```
                          Job Queues to Process                  RCHAS400

Use . . . . . :  *ARC
Control group . . . . :  SAMPLE


Type choices, press Enter.

  Seq    Job queue     Library        Hold       Release

   10    QBATCH        *LIBL          *YES       *YES




                                                                    Bottom

F3=Exit    F5=Refresh    F12=Cancel
```

Figure 43. The Job Queues to Process Display

3. In the *Seq* field list the order in which you want BRMS to hold or release the job queues.
4. Type the job queue name in the *Job queue* field.
5. At the *Library* parameter, you can specify the library in which the job queue resides or use the default library, *LIBL.
6. Specify whether you want BRMS to hold or release the job. In this case, BRMS holds the QBATCH job queue during the archive, which means that no new batch jobs can start. BRMS automatically releases the job queue after the archive or IPL (if you specify IPL in the control group) completes.
7. Press Enter to save and return to the Work with Archive Control Groups display.

To delete a job queue, position the cursor on that entry, space through the sequence number, and press Enter. Press Enter again to return to the Work with Archive Control Groups display.

Using the *LOAD and *EXIT Special Values in Archive Control Groups

Use the *LOAD and *EXIT special values in archive control groups to instruct the system to perform special operations. You can use the *LOAD and *EXIT special values as items in an archive list.

Using the *LOAD Special Value: Use the *LOAD special value instructs BRMS to send a message to the operator to load a new tape before processing the next archive item. The *LOAD special value allows you to segregate selected archive items from regular control group processing.
In this situation, BRMS instructs the operator to load a new volume before

```
                  Create Archive Control Group Entries   RCHAS400

Group . . . . . . . . . . . : SAMPLE
Default activity  . . . . . *ARCPCY
Text  . . . . . . . . . . . *NONE

Type information, press Enter.

                              Weekly
        Archive       List    Activity
Seq     Items         Type    SMTWTFS

  10    LIB01                 *DFTACT
  20    *LOAD                 *DFTACT
  30    LIB02                 *DFTACT









                                                                  Bottom
F3=Exit                  F5=Refresh             F10=Change item
F11=Display exits        F12=Cancel             F24=More keys
```

*Figure 44. The *LOAD special operation as an entry within a control group*

processing LIB02.

Note: If you use *LOAD for MLB media, the tape loads automatically so the operator does not receive a message to load a new tape.

Using the *EXIT Special Value: The *EXIT special operation allows you to perform user command processing. Consider using *EXIT at the following times:

- Prior to performing pre-processing activities for the control group (such as processing subsystems, and job queues).
- During sequential control group processing (which is based on the numbers you assign).
- After performing post-processing activities for the control group (such as processing subsystems and job queues).

```
                    Edit Archive Control Group Entries       RCHAS400

Group . . . . . : SAMPLE
Default activity  . . . . . *ARCPCY
Text  . . . . . *NONE

Type information, press Enter.

                              Weekly
        Archive       List    Activity
Seq     Items         Type    SMTWTFS

  10    *EXIT                 *
  20    LIB01                 * *
  30    *EXIT                   *  *
  40    LIB02                 *    *
  50    *EXIT                   *  *
  60    LIB03                   *
  70    *EXIT                 *

                                                              Bottom
F3=Exit                   F5=Refresh              F10=Change item
F11=Display exits         F12=Cancel              F24=More keys
```

*Figure 45. The *EXIT special operation as an entry in a control group*

In this example, BRMS archives libraries on Sunday, Tuesday, and Friday. On Sunday, BRMS processes two *EXIT operations, items number 10 and number 70. BRMS processes the first *EXIT operation before control group pre-processing activities because it is the first item (sequence number 10) in the list of items to process that day. The second Sunday *EXIT operation (sequence number 70) is the last item in the list of items to process that day. Thus, BRMS processes this *EXIT operation after completing the post-processing activities for that control group.

BRMS also processes two *EXIT operations on Tuesday, sequence numbers 30 and 50. Because neither of the exits is first or last in the list of items to process that day, BRMS processes them in the order in which you sequenced them. So, on Tuesday, BRMS processes LIB01 (sequence number 20) first, followed by *EXITs 30 and 50. BRMS processes LIB03 (sequence number 60) last.

On Friday, BRMS processes *EXIT operation 30 before performing control group pre-processing activities, followed by LIB02 (sequence number 40). BRMS processes *EXIT operation 50 after completing the post-processing activities for that control group.

Setting Up Media Classes for Archive

You can create an entirely separate media class for use only with archiving. When you create a separate media class, you can use different tape drives for archiving or use a different tape format, or you can avoid having to share your archive tapes with other systems or backup jobs. To create a media class:

1. Type `WRKCLSBRM *MED` from the command line and press Enter.

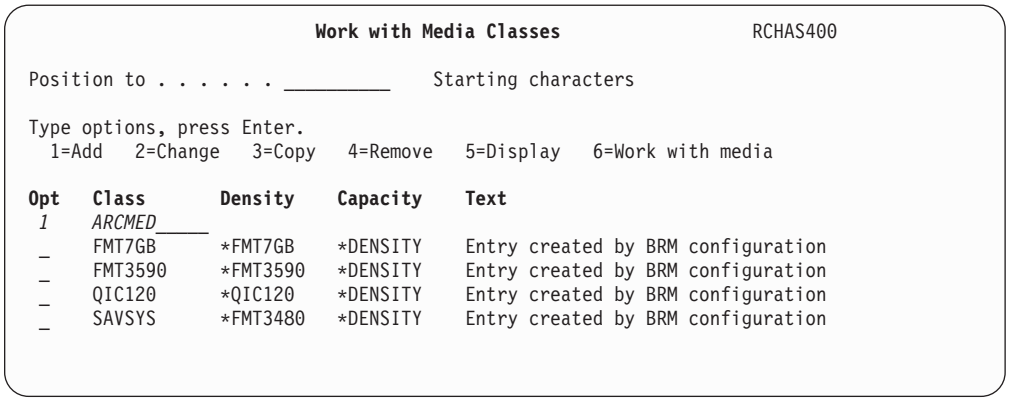

Figure 46. Work with Media Classes Display

2. Select option 1 to create a new media class.

   Media classes define the types of physical media that are used for archive. Typical physical media are any removable storage media available on the system.

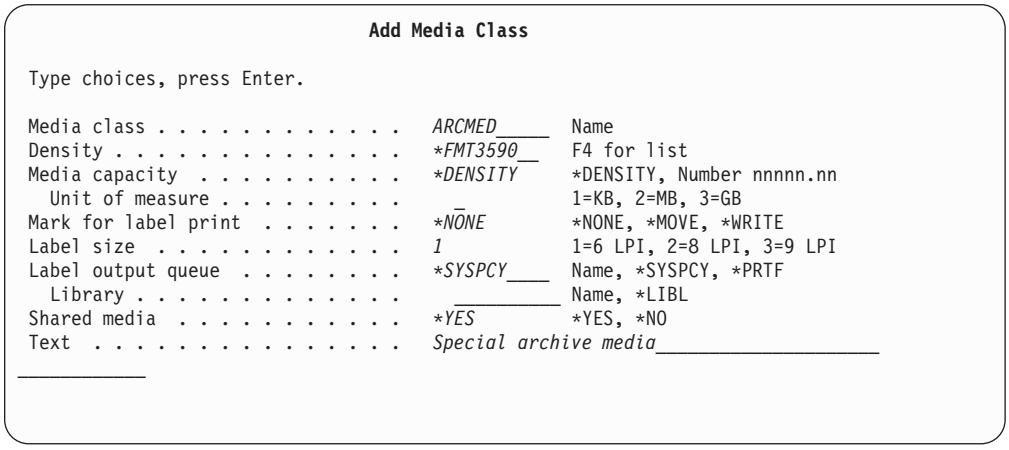

Figure 47. Add Media Class Display

3. Press Enter to create the media class and return to the Work with Media Classes display.
4. Press Enter once more to leave this display.

Setting Up Move Policies for Archive

You may want to create special move policies for your archive-with-retrieve tapes. It is unlikely that these tapes will cycle through the locations in the same way as regular save tapes, or even non-retrieve type archive tapes.

You might need to create at least two move policies.

The first should be used for the “active” tapes that contain the “active” data that has been archived, and could be required for retrieve at any moment. **This set will be the copy of the original archive set.** Because you created the duplicate set after the original set of tapes, BRMS will regard these tapes as the latest versions. You

must use the duplicate set as the “active” set. Keep them close to the drives (for example within your tape library device or a tape rack). This move policy should keep the tapes close at hand, that is, near to the tape drives that will be used for retrieve operations. If you have an automated tape library device, you may set the move policy such that the tapes remain in the library at all times.

The second move policy should be for the original tapes of which we recommend that you make a duplicate after every archive operation. You should send these tapes immediately to a different site from that in which the “active” archive copies are stored.

You may repeat this pairing of move policies for each different physical location in which you are likely to store archive tapes.

To create a move policy:

1. Type `WRKPCYBRM *MOV` command and press Enter.

```
                          Work with Move Policies                RCHAS400

Position to . . . . . _____________    Starting characters

Type options, press Enter.
  1=Create   2=Change   3=Copy   4=Delete   5=Display

Opt    Policy        Text
 1     ARCMOV_______
 _     OFFSITE       Entry created by BRM configuration
```

Figure 48. Work with Move Policy Display

2. Select option 1 to create a new move policy and type in the name of that policy.

```
                             Create Move Policy                  RCHAS400

Move policy  . . . . . .  ARCMOV
Home location  . . . . .  *SYSPCY____  Name, *SYSPCY, F4 for list
Use container  . . . . .  *NO_        *YES, *NO
Verify moves . . . . . .  *YES        *YES, *NO
Calendar for working days  . . .  *ALLDAYS__  Name, *ALLDAYS, F4 for list
Calendar for move days . . . . .  *ALLDAYS__  Name, *ALLDAYS, F4 for list
Text . . . . . .  Archive tapes only_____________________________

Type choices, press Enter.

  Seq    Location      Duration
   10    ATL01_______  180____
   20    OFFSITE_____  180_____
   30    TAPMLB01____  90_____
   40    VAULT_______  30_____
                                                                  Bottom

F3=Exit    F4=Prompt    F5=Refresh    F12=Cancel
```

Figure 49. Create Move Policy Display

3. Define the sequence and duration of all the locations. Press Enter to create the new move policy and to return to the Work with Move Policies display.

Setting Up Media Policies for Archive

You need to create a separate archive media policy for every different retention period for your archived data. The retention period indicates how long a piece of data on a tape or in a save file should stay active from the day that the data was written to that tape or to the save file. When the retention period is exceeded, the data on the tape or save file is expired. The storage media then becomes available for any new backup or archive operation.

In the media policy you also establish the link to the media class that you want to use (possibly created in “Setting Up Media Classes for Archive” on page 69), and the move policy that is appropriate (created in “Setting Up Move Policies for Archive” on page 70). If you are using an automated tape library, you may also specify the name of the required library location in the Storage location parameter. This will help BRMS select the correct tape drives to use (that is the ones in the tape library) when performing archive operations or retrieve operations.

To create a media policy:

1. Type `WRKPCYBRM *MED` command and press Enter.

```
                    Work with Media Policies                    RCHAS400


Position to . . . . . _________    Starting characters

Type options, press Enter
 1=Create   2=Change   3=Copy   4=Delete   5=Display

Opt Policy       Text

 _  _________
 2  ARCHIVAL     Entry created by BRM configuration
 _  FULL         Entry created by BRM configuration
 _  INCR         Entry created by BRM configuration
 _  SAVSYS       Entry created by BRM configuration
```

Figure 50. Work with Media Policies Display

2. Select option 1 to create a new media policy, or select option 2 to change the media policy.

```
                        Change Media Policy

  Media policy . . . . . :   ARCHIVAL

Type choices, press Enter.

  Retention type . . . . . .  2           1=Date, 2=Days,
                                          3=Versions, 4=Permanent
    Retain media . . . . . .  1725_____   Date, Number
  Move policy  . . . . . .  *NONE____     Name, *NONE, *ADSM, F4
  Media class  . . . . . .  *SYSPCY__     Name, *SYSPCY, *ADSM, F4
  Storage location . . . . .  *ANY_____   Name, *ANY, F4 for list
  Save to save file  . . . .  *NO_        *YES, *NO
    ASP for save files . . .  01          Name, *SYSTEM, 1-16
    Save file retention type  4           1=Date, 2=Days,
                                          3=Permanent, 4=None
      Retain save files  . .  *NONE____   Date, Number, *NONE
    ASP storage limit  . . .  90          *SYS, 1-99
  Secure media . . . . . .  *NO           *YES, *NO, *ADSM
  Text . . . . . .  Entry created by BRM configuration________
  _________

 F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel
```

Figure 51. Add or Change Media Policy Display

Refer to the online help information for an explanation of the parameters. To archive to a save file, specify *YES in the Save to save file parameter and fill in the other pertinent ASP parameters.

Expiration of a tape or save file with archived data is effectively deleting that data. Every time an object is retrieved, it is restored from tape or save file and used on the system. At a later date it will be archived again, and the original copy will become redundant, or “expired”. For a tape that contains several archived objects, as time passes, more and more of the archived objects will “expire”. This is all wasted space on the tape, and can only be reclaimed when BRMS expires the entire tape.

Starting Archive Operations

The Start Archive using BRM (STRARCBRM) command allows you to select a control group and begin the archive process or produce a report of archive candidates. You can start the archive immediately, at a scheduled time, or you can schedule it using the system job scheduler. Processing can be batch or interactive.

To start an archive operation, issue the STRARCBRM command, and press Enter. The default value for the OPTION parameter is *REPORT, which produces an Archive Candidates Report. The special values *LOAD and *EXIT control group entries are ignored when you choose the *REPORT option. There are four Archive Candidate Reports:

- The report for the Archive DLO Candidate report is written to printer file QP1A8ARF
- The report for the Archive Object Candidate report is written to printer file QP1AARC
- The report for the Archive Stream File Candidate report is written to printer file QP1A1ARC

- The report for the Archive Spooled File Candidate report is written to printer file QP1AOQ.

To perform an archive, issue the command STRARCBRM and change the OPTION parameter to *ARCHIVE. You will receive an Archive report that is based on the Candidate report, previously created, for each archive operation that you run.

Use the Work with Media Information (WRKMEDIBRM) command to validate the archive occurred specifying the SAVTYPE *ARC parameter. This command displays archived information based on specified libraries, date ranges and sequences. The Work with Media Information command shows the date and time each library was saved, the type of save, the volume, expiration date and the number of objects that were saved and not saved. The command gives you the option to display or print the report. The report that is produced is the Media Information report and is written to printer file QP1AHS.

Additional Options for Processing Your Archive Control Groups

If you use a single tape device, all of your archive operations are **serial** operations. In a serial operation, BRMS processes your archive jobs sequentially, one at a time, one after another. If, however, you have multiple tape devices, BRMS can perform concurrent and parallel archive operations. Concurrent and parallel archive operations can considerably reduce your backup window and streamline your backup operations. Following are brief descriptions of each.

Concurrent Archiving Operations: In a concurrent archive operation, you send multiple archive jobs to multiple tapes devices to process at the same time (concurrently). For example, in a concurrent archive strategy, you might send one library or group of libraries to one tape device, and another set of libraries to different tape device. Concurrent archive support also allows multiple archive operations from the same auxiliary storage pool (ASP). Anything that you can archive concurrently, you can retrieve concurrently, thereby reducing the amount of time it takes to retrieve multiple libraries or objects.

Parallel Archiving Operations: A parallel archive operation processes one archive item (a control group, library, or other object) across multiple devices at the same time. Due to the processing overhead that is required for each item archived, you should use parallel support only for large objects or large libraries.

For a concurrent archive, you must manually determine the size of the data groups you want BRMS to process. When you specify parallel support, the system balances the data as evenly as possible so that the devices complete processing at the same time.

You should, whenever possible, use the same **number** of devices to perform the parallel retrieve that you used for the parallel archive. However, you can retrieve an item that you archived with parallel support by using only one device. To perform a large number of parallel retrieves from fewer devices, you should use a tape library because extra media mounts may be required.

Note: You can find more information on how to perform concurrent and parallel retrieve operations in the “Performing Parallel and Concurrent Retrieve Operations” on page 80 section of this chapter.

Parallel Support Restrictions: Before you set up a archive strategy using parallel support, review the following restrictions:

- Due to OS/400 restrictions, you can use parallel support for only one library, or the objects from one library, at a time.
- Limit your parallel operations to not only just one library or group of objects, but to large libraries or objects. This means that BRMS cannot process *ALLUSR or other special values with parallel resources. Parallel processing begins when the data streams into the tape devices. If you process small objects or libraries in a parallel archive, your archive window could increase due to the serial nature of the archive setup. That is because the set up of each tape, including volume, header, and trailer labels, is not done during a parallel operation.
- BRMS cannot support parallel archives of DLO or Integrated File System objects at this time.
- While parallel operations can work with multiple stand-alone tape drives, the use of tape libraries makes retrieval simpler.
- BRMS does not archive media information for libraries and obejcts archived in parallel. Therefore, you must run the Save Media Information using BRM (SAVMEDIBRM) command after each group of items that you archive in parallel.

When to Use Parallel and Concurrent Support: Due to OS/400 restrictions, you can gain the greatest benefit by designing a strategy that combines concurrent and parallel archive support. You should, if possible, use concurrent archives to reduce archive windows when saving folders, documents, spooled files, and Integrated File System data. Use parallel support when you want to reduce your archive window for large libraries and large objects.

How to Set Up a Parallel Archive Operation: To use parallel support to process an existing control group, take the following steps:

1. At the Work with Archive Control Groups display, type option 8 (Change attributes) in front of the control group you want to work with, and press Enter. This takes you to the Change Archive Control Group Attributes display.

```
                  Change Archive Control Group Attributes

Group  . . . . . . . . . . . . . . . . : SAMPLE

Type information, press Enter.

Media policy . . . . . . . . . . . . . . *ARCPCY     Name, F4 for list
Archive devices  . . . . . . . . . . . . *ARCPCY     Name, F4 for list



Parallel device resources:
  Minimum resources  . . . . . . . . . . 2            1-32, *NONE, *AVAIL
  Maximum resources  . . . . . . . . . . 4            1-32, *AVAIL, *MIN

                                                                      More...
F3=Exit   F4=Prompt   F12=Cancel
```

Figure 52. Processing a control group with parallel device resources

2. At the *Parallel device resources* parameter, indicate the minimum and maximum number of device resources you want to allocate to the processing of this

control group. In this example, we instructed BRMS to archive the contents of the SAMPLE control group to a minimum of 2 devices and a maximum of 4 devices.

3. Press Enter to save and exit.

Though this example processes a single control group, you can also use parallel support to archive multiple control groups.

Use a job scheduler to assist you in scheduling parallel and concurrent archives.

Setting Up Dynamic Retrieval with BRMS

This topic describes the day-to-day activities that will help you to control the BRMS dynamic retrieval function to suit your needs. You can find more information in "Dynamic Retrieval" on page 49. This topic discusses the following items:

1. Setting up Retrieve Policy
2. Setting up Retrieve Controls for a particular job
3. Retrieval Methods

Setting Up the Retrieve Policy

A retrieve operation is initiated for each object that has been opened and found to be saved with storage freed. The retrieve policy within the BRMS policy structure guides the retrieve operation. This retrieve policy governs many attributes that control the way in which a retrieve is performed, and *it takes effect across the entire system*. That is, the retrieve policy is system wide.

You may not want to retrieve the data for all your applications in the same manner. For additional flexibility, you can use the `SETRTVBRM` command to override the system-wide retrieve policy settings, at the job level. When you issue `SETRTVBRM`, the new settings apply for all retrieve operations initiated after that point, and until the job ends, or until you issue another `SETRTVBRM`. "Setting the Retrieve Controls for a Particular Job" on page 79 explains the SETRTVBRM command in detail.

To set the Retrieve Policy, issue WRKPCYBRM *RTV.

```
                      Change Retrieve Policy                    RCHAS400

Type choices, press Enter.

Retrieve device  . . . . . .  *MEDCLS     Name, F4 for list



Retrieve confirmation:
  Interactive operation  . . . . . .  *VERIFY     *VERIFY, *NOTIFY, *DELAY..
  Batch operation  . . . . . .  *NOTIFY     *NOTIFY, *DELAY, *VERIFY..
Retrieve authorization . . . . . .  *OBJEXIST   *OBJEXIST, *READ, *UPD...
End of tape option . . . . . .  *REWIND     *REWIND, *LEAVE, *UNLOAD
Option . . . . . .  *ALL        *ALL, *NEW, *OLD, *FREE
Allow object differences . . . . . .  *NONE       *NONE, *ALL
ASP high storage threshold . . . . . .  *SYS        *ASP, *SYS, 1-95
Retrieved object retention . . . . . .  0           0-9999, *ARC, *NOMAX
Extend retention on usage  . . . . . .  0           0-9999
Reset days used count. . . . . .  *NO         *NO, *YES

                                                                   More...
F3=Exit    F4=Prompt    F5=Refresh    F9=System policy
F12=Cancel
```

```
                      Change Retrieve Policy                    RCHAS400

Type choices, press Enter.

Parallel device resources:
    Minimum resources. . . . . .  *SAV      1-32, *SAV, *NONE, *AVAIL
    Maximum resources. . . . . .            1-32, *AVAIL, *MIN









F3=Exit    F4=Prompt    F5=Refresh    F9=System policy
F12=Cancel
```

Figure 53. Change Retrieve Policy Display

The following options are supported:

Retrieve Device

The device that is used for the retrieve operation. *MEDCLS means that BRMS determines the media class on which the requested item is saved. You can specify up to four devices at this parameter.

When media containing the file is located in a Media Library device (MLB), BRMS limits its choice of *MEDCLS devices to those that are at the MLB location.

Retrieve Confirmation

Specify how objects requested for retrieval are processed, either batch or interactive.

You can find a full description of these parameters in “Retrieval Methods” on page 81.

Retrieve Authorization

The authority option tells BRMS what level of authorization to an object is necessary before the accessing user will be able to retrieve the object. This authorization level is checked, and if met, the retrieve operation is performed. If the authorization level is not met, a message is sent indicating that the object was not restored and that it cannot be used until restored. BRMS tracks the unretrieved file and indicates that an *AUTHORITY failure occurred. Use the RSMRTVBRM command and the Resume Retrieve display to identify files that were unable to be retrieved due to authority failures. You can request that the retrieve operation for one or more of them be performed or canceled.

For many enterprises, users only have use or update authority to objects. If the authority level is *UPD, at open time BRMS automatically retrieves the archived objects for those users that have at least update authority. In doing so, BRMS allows the object to be retrieved, without having to grant users who would access the file with *OBJEXIST authority to enable dynamic retrieval.

You may decide to downgrade the user’s required authority level for a restore (for example to *OBJOPR). What this does is grant a user limited existence (or creation) rights to certain objects, under certain conditions, when they would previously only have been able to use the object. There is an inverse relationship here. The lower you specify the required authority, the more authority you are granting.

End of tape option

This is identical to the OS/400 standard save/restore end of tape option parameters. The default is *REWIND.

If you are using an automated tape library or even a drive with an automatic cartridge loader, you increase the level of automation if you specify *UNLOAD. This removes the current cartridge from the drive, which leaves it available for the next operation.

Option Specifies which objects are restored, depending on whether the object exists on the system.

Allow object differences

This parameter indicates if a restore operation can tolerate object differences. The values of *NONE and *ALL are supported and have exactly the same meaning as they have for the ALWOBJDIF parameter on the RSTOBJ command. The default value is *NONE.

ASP high storage threshold

Specifies the ASP high storage threshold value that cannot be exceeded when retrieving an object.

Retrieved object retention

Specifies how long objects that have been retrieved, are kept on the system. After the object retention period has passed, the storage associated with the object is freed.

Extend retention on usage

Specifies the number of days to extend the retention of a retrieved object if it is used again after the retrieval operation.

Reset days used count
Specifies whether the object days used count should be reset on a successful dynamic retrieval.

Parallel Device Resources
Specifies the minimum and maximum number of devices you want to allocate to a parallel retrieve operation, if any. You can find more information on performing parallel retrieves in the “Performing Parallel and Concurrent Retrieve Operations” on page 80 section of this chapter.

It should be noted that *all* retrieve operations are constrained by the Storage Threshold, which is a high water mark for auxiliary storage pool (ASP) utilization. You express this value through the System Service Tools (STRSST) ASP threshold. (See *Backup and Recovery* for more details.) BRMS does not restore a file if doing so would cause the ASP’s storage threshold to be exceeded. If the storage threshold were to be exceeded, messages are sent indicating that the file was not restored and that it cannot be used until restored. BRMS tracks the unretrieved file to indicate that a *STORAGE failure occurred. Use the RSMRTVBRM (Resume Retrieve Using BRM) command and the Resume Retrieve display to identify files that were unable to be retrieved due to DASD space constraints. You can request that the retrieve operation for one or more of them be performed or canceled.

Setting the Retrieve Controls for a Particular Job

You may want to override the values that are set by the Retrieve Policy (for the whole system) for a particular job. To do this you must issue the `SETRTVBRM` (Set Retrieve Controls for BRM) command within the job that requires the override.

The controls you specify with SETRTVBRM remain in effect for your job until they are reset (for example the job ends) or otherwise changed with another SETRTVBRM command. To see control values that are currently in effect, type `SETRTVBRM` on a command line and press F4.

```
                  Set Retrieve Controls for BRM (SETRTVBRM)

Type choices, press Enter.

Retrieve device  . . . . . .   *SAME         *SAME, *MEDCLS, TAPMLBBRM...
               + for more values
Parallel device resources:
  Minimum resources  . . . . .   *SAV          1-32, *SAV, *NONE, *AVAIL
  Maximum resources  . . . . .   *MIN          1-32, *MIN, *AVAIL
Retrieve confirmation:
  Interactive operation  . . . .   *SAME       *SAME, *RTVPCY, *VERIFY...
  Batch operation  . . . . . .   *SAME         *SAME, *RTVPCY, *VERIFY...
Retrieve authorization . . . . .   *SAME       *SAME, *RTVPCY, *OBJEXIST...
End of tape option . . . . . .   *SAME         *SAME, *RTVPCY, *REWIND...
Option . . . . . .   *SAME         *SAME, *RTVPCY, *ALL, *NEW...
Allow object differences . . . .   *SAME       *SAME, *RTVPCY, *NONE, *ALL
ASP high storage threshold . . .   *SAME       *SAME, *RTVPCY, 1-95, *ASP...
Retrieved object retention . . .   *SAME       *SAME, *RTVPCY, 0-9999...
Extend retention on usage  . . .   *SAME       *SAME, *RTVPCY, 0-9999
Reset days used counter  . . . .   *SAME       *SAME, *RTVPCY, *YES, *NO
                                                                    Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 54. Set Retrieve Controls Display

The parameters shown on this display are identical to those on the retrieve policy. You can override any of the retrieve policy parameters on this display. If you want to perform a parallel retrieve operation, you should specify that option on this display, rather than on the retrieve policy display. You are likely to change the parameters on this display to suit individual jobs, whereas you are not likely to change the retrieve policy often. You can find more information on parallel retrieve in the following section. You can review the“Setting Up the Retrieve Policy” on page 76 for descriptions of the parameters on this and the retrieve policy display.

Performing Parallel and Concurrent Retrieve Operations

“Performing Parallel and Concurrent Retrieve Operations” You can retrieve data by using serial, concurrent, or parallel retrieval operations. If BRMS performed your archives as serial or concurrent backups, BRMS can retrieve them either serially or concurrently. On the other hand, however, you can perform parallel retrieves only if BRMS used a parallel operation to perform the archive. Both concurrent and parallel restore operations require multiple tape drives.

Following are brief summaries of the characteristics of, and the differences between, serial, concurrent and parallel retrieves.

Serial Retrieval

This straightforward retrieve method is the simplest way to restore data that was saved by multiple control groups. In a serial retrieve operation, BRMS issues one retrieve job at a time to one tape drive. BRMS performs serial retrieves in a sequential manner, by retrieving one tape and one control group after another. By default, BRMS considers every archive and retrieve a serial operation unless otherwise specified by you.

Concurrent Retrieval

In a concurrent retrieve operation, you send multiple retrieve jobs to multiple tapes drives to process at the same time (concurrently). You, not BRMS, set up the concurrent retrieve. To do this, try to equally balance the size of the jobs you send to process concurrently so that they end at the

same time. Concurrent retrieve operations can reduce your down time after a system failure by allowing you to retrieve multiple libraries or objects at the same time.

Parallel Retrieval
A parallel retrieve operation processes one retrieve item (a control group, library, or other object) across multiple devices at the same time. You specify the item you want to retrieve, and BRMS splits it into equal parts to process concurrently on multiple devices. By using parallel resource support, you can reduce down time during a retrieval operation.

To perform a parallel retrieve operation, you first need to meet the following criteria:

- You can use parallel retrieve support only on items that BRMS archived with parallel support.
- You must have multiple tape drives available to run at the same time.

BRMS should perform parallel retrieves on the same devices that were used to process the archives. However, you can issue parallel retrieves to a smaller number of resources. If you plan to perform a large number of parallel retrieves, you should use a tape library if you need to retrieve to fewer devices. In that case, BRMS may require additional media mounts.

Note: For performance reasons, you should use parallel archive and retrieve options only for large control groups, objects, or large libraries.

You can specify the use of parallel resources during retrieve processing from either the retrieve policy or the SETRTVBRM command.

Retrieval Methods

There are several different modes of operation for performing a retrieve with BRMS. This section describes these modes and their application. To set up the retrieval method, you must either change the BRMS retrieve policy, or issue a SETRTVBRM command within the job that you wish to change. The retrieve policy and SETRTVBRM command support options for separate batch and interactive controls.

Retrieve operations involve retrieving archived objects from tape or disk when requested by a program or other operation. The value that you set up in the retrieve policy determines how archived objects requested in interactive or batch processing are handled.

For interactive operations you can set the following values:

***VERIFY**
Send a message for each object that is retrieved before continuing.

***NOTIFY**
Retrieve the object immediately and notify with a message.

***DELAY**
Postpone the retrieve until a later time.

***SBMJOB**
Submit the retrieve job as a batch job.

***NONE**
Do not have BRMS perform retrieve operations.

For batch operations you can set the following values:

***VERIFY**
Send a message for each object that is retrieved before continuing.

***NOTIFY**
Retrieve the object immediately and notify with a message.

***DELAY**
Postpone the retrieve until a later time.

***NONE**
Do not have BRMS perform retrieve operations.

When you specify *VERIFY in **batch mode**, you must monitor the system operator message queues frequently. A batch job waiting for an operator message reply is a frequent cause of batch throughput problems.

Responding to a Retrieve Operation

This topic details the messages that you may see while witnessing a retrieve operation and any responses that you may need to give to inquiry messages received as part of that retrieve operation.

****VERIFY:*** By default, a program message is displayed to the user as shown in Figure 55 on page 83. You see the first display, and you see the “Additional Message Information” display only if you press the Help key. You see important additional information in the second display that includes object size and ASP utilization. This information may influence your decision to initiate the retrieve immediately.

```
                          Display Program Messages

Job 016022/CRAIG/TAMLINLUJ started on 04/30/98 at 11:33:06 in subsystem QINT
Retrieving PAYMASTFIL in library PAYROLL. (C G I S)









Type reply, press Enter.
  Reply . . .   _____________________________________________________________
_____________________________________________________________________________

F3=Exit   F12=Cancel
```

```
                       Additional Message Information

Message ID . . . . . :   BRM1822       Severity . . . . . :   99
Message type . . . . . :   Inquiry
Date sent  . . . . . :   04/30/98      Time sent  . . . . . :   18:14:51

Message . . . . :   Retrieving PAYMASTFIL in library PAYROLL. (C G I S)
Cause . . . . . :   Access to suspended object PAYMASTFIL member PAYDEC94 in
  library PAYROLL type *FILE is requesting that the object be restored to the
  system. The size of the object is 51.798 megabytes. The object will be
  restored to ASP 1 which is currently 84.53 percent utilized. When complete
  the approximate ASP utilization will be 89.14 percent.
Recovery  . . . :   Type a valid reply for the restore of the object.
Possible choices for replying to message . . . . . . . . . . . . . . . :
  G -- Continue the operation.
  C -- Cancel the operation.
  I -- Ignore the request and delay the retrieve operation. To resume a
    delayed retrieve operation at some later time use the RSMRTVBRM command.

                                                                More....
```

```
                       Additional Message Information

Message ID . . . . . :   BRM1822       Severity . . . . . :   99
Message type . . . . . :   Inquiry

  S -- Submit the retrieve operation for batch processing. The current job
    will receive an indication that the object's data was not found.
Technical description . . . . . :   Access to a suspended object has
  caused BRMS to attempt to retrieve the object from archives. If the object
  is a physical file then only the requested member for that file will be
  restored.
```

*Figure 55. Retrieve *VERIFY Messages*

The valid responses to the program message screen are:

G Go - The retrieve begins immediately and the application is suspended waiting for it to complete. You should not use the End request (System Request, option 2) function during this time.

C Cancel - This option will return two main messages. BRM1823 will be added to the job log. This message indicates that the object was archived

and that the BRMS retrieve request was canceled. Also the standard OS/400 CPF4102 message is sent to enable the application to respond.

S Submit Job - The retrieve request is submitted to the job queue that is specified in the user's job description. Again, the same two messages are sent for the application to handle. A message will be later sent to inform the user that the retrieve operation is complete.

I Ignore and Delay - The retrieve request is added to the list of file members to be retrieved at a later time. Again, the same two messages are sent for the application to handle. A message is sent later to inform the user that the retrieve operation is complete.

NOTIFY: You see the status message that is shown in Figure 56. The job waits until the restore is complete. If you specify the immediate (GO) option with *VERIFY mode, do not use the End job (System Request, option 2) function during this time.

For *NOTIFY and the immediate (GO) option of *VERIFY, errors (such as media errors) and tape mount messages are reported to the system operator, or the BRMS notification message queue. If an error is severe and the operation is canceled, all messages (including those responded to by the system operator) are added to the user's job log. The application is sent the original OS/400 CPF4102 message.

```
                                                                    Bottom
Parameters or command
===> dsppfm PAYROLL/PAYMASTFIL_____________________________________________
F3=Exit          F4=Prompt             F5=Refresh            F6=Create
F9=Retrieve      F10=Command entry     F23=More options      F24=More keys
Retrieve object PAYMASTFIL for library PAYROLL in progress.
```

*Figure 56. Retrieve *NOTIFY Message*

SBMJOB: Message BRM1824 is added to the user job log to inform the user that the retrieve job is submitted. The open request (the application) is sent message CPF4102 to handle as an error, and enables the user to retry the function at a later time.

When the batch job is complete, the user is informed with the status of the restore. If the restore were to fail, the application has already reacted. Do not try that function again until the cause of the error is fixed.

DELAY: Message BRM1823 is added to the user job log to inform the user that the retrieve job is submitted. The open request (the application) is sent message CPF4102 to handle as an error, and enables the user to retry the function at a later time.

When the retrieval is later resumed, the user is informed with the status of the restore. As for the *SBMJOB option, if the restore were to fail, the application has already reacted. Do not try that function again until the cause of the error is fixed.

Failed Retrieve Operations: In general if a retrieve operation fails due to exceptions other than *STORAGE or *SECURITY, then it is the submitter's responsibility to retry (or re-submit) the retrieve. There is no implication that the failed retrieve will then be converted to a *DELAY type retrieve.

Check the BRMS log frequently. Use option *RTV to check on retrieve operations.

Using the BRMS Log

The primary method of auditing all BRMS activity is through the BRMS log. You access this log through the DSPLOGBRM command.

The DSPLOGBRM (Display Log using BRM) command supports the display of log entries that record the occurrence, success, and failure of BRMS operations. The same log concept that is used throughout BRMS can also be used to track retrieve operations. Log entries are categorized by type to show which operation caused the log entry. The entry type “*RTV” is supported for retrieve type operations, and is used to record whether these are successful or unsuccessful. A user can search all BRMS log entries by using type of *RTV to audit retrieve operations.

Simply issue the `DSPLOGBRM *RTV` command to list all the available log entries for retrieve operations.

```
04/30/98                Display BRM Log Information                RCHAS400
14:27:42                                        Position to . . . . 04/30/98

Object retrieve for PAYROLL1 not successful.
------------------------------ 04/26/98 ----------------------------------
Retrieve object QRPGSRC for library CRAIG in progress.
Object retrieve for QRPGSRC not successful.
Retrieve object 93TRANFILE for library ACCTSPAY in progress.
Retrieve of 93TRANFILE in library ACCTSPAY was successful.
------------------------------ 04/27/98 ----------------------------------
Retrieve object PAYROLL1 for library CRAIG in progress.
Object retrieve for PAYROLL1 not successful.
------------------------------ 04/28/98 ----------------------------------
Retrieve object QHST94296C for library QSYS in progress.
Retrieve of QHST94296C in library QSYS was successful.
Retrieve object QHST94294A for library QSYS in progress.
Retrieve of QHST94294A in library QSYS was successful.
                                                                    More...
Press Enter to continue.
```

Figure 57. BRMS Log Showing Retrieve Entries

Controlling Retrieve Operations Using the RSMRTVBRM Command

A retrieve operation may fail (or not even be started) because of authority problems, or the ASP high threshold. Other retrieve operations may be initiated later in the *DELAY retrieve mode. In any of these cases the retrieve operation will enter into a deferred state. Further control of these deferred retrieves is needed.

The RSMRTVBRM (Resume Retrieve using BRM) command facilitates the recovery of delayed or otherwise unsuccessful retrieve operations. This command allows the user to work with or print a list of files for which retrieve operations are pending. This may be due to reasons such as the following:

- Retrieve policy, user, or system operator specified that the retrieve operation be delayed
- ASP to contain the retrieved file would have exceeded its storage utilization limit
- User accessing the archived file did not have appropriate authority to perform a retrieve operation

Enter the RSMRTVBRM command and press F4, followed by F9 for all parameters. You see a display similar to that shown in Figure 58.

```
                    Resume Retrieve using BRM (RSMRTVBRM)

Type choices, press Enter.

Retrieve select  . . . . . .   *ALL          *ALL, *DELAY, *SECURITY...
Action . . . . . . . . . . .   *RETRIEVE     *RETRIEVE, *REPORT
Confirm retrieval  . . . . .   *YES          *YES, *NO
Auxiliary storage pool . . . .   *ALL         Name, *ALL, *SYSTEM, 1, 2...
Item type  . . . . . . . . .   *LIB          *LIB, *FLR, *LNK
Library  . . . . . . . . . .   *ALL          Name, *ALL
Retrieve device  . . . . . .   *RTVPCY       Name, *RTVPCY, *MEDCLS
              + for more values
Parallel device resources:
  Minimum resources  . . . . .   *SAV          1-32, *SAV, *NONE, *AVAIL
  Maximum resources  . . . . .   *MIN          1-32, *MIN, *AVAIL
End of tape option . . . . . .   *RTVPCY       *RTVPCY, *REWIND, *LEAVE...
Option . . . . . . . . . . .   *RTVPCY       *RTVPCY, *ALL, *NEW, *OLD...
Allow object differences . . . .   *RTVPCY       *RTVPCY, *NONE, *ALL


                                                                      Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 58. The Resume Retrieve using BRM (RSMRTVBRM) Command

The following parameters are supported:

Retrieve Select

The type of condition that caused the retrieve to be delayed. During the retrieve operation several things could have occurred to cause an object to be marked for restore at a later time. Values supported for this option are:

- ***ALL**

  Use of this value selects objects whose retrieve operation was postponed due to any of the reasons that are described below.
- ***DELAY**

  This value indicates that the retrieve operation was delayed through specification on the retrieve policy, the SETRTVBRM command, or a delay reply to the *VERIFY option. All will cause the object to be marked as delayed. Use of this value selects these delayed objects to be retrieved.
- ***SECURITY**

  The restore of the object for a user failed because of insufficient authority to the object. An authorized user can use this value to cause the restore of objects which were previously unretrievable due to insufficient authorization.
- ***STORAGE**

  When attempting to retrieve an object, there was not enough space on DASD to contain it without exceeding the ASP threshold. Later when space on DASD is available to contain the object, use this value to cause the restore of objects which were previously unretrievable due to space constraints.

Confirm Retrieval
When running this command, determine if the user sees a confirmation display. This parameter is ignored in batch.

Allowed values are *YES and *NO. The default is *YES, and this will present a confirmation screen. Note that semantically speaking this parameter is actually “Present confirm screen?” rather than the “Confirm Retrieval” description.

Auxiliary Storage Pool ID
Specifies the auxiliary storage pool into which you want to retrieve items. Allowed values are *ALL, *SYSTEM, ASP-number and ASP-name.

Item type Specifies which type of item is displayed on the Confirm Retrieve display prior to retrieve operations.

Library Library name can qualify selection of objects. Allowed values are *ALL or a specific library name.

Parallel Device Resources
If you specified parallel support on the original retrieve job (that you now want to resume), you need to specify parallel support on this display as well.

Use the remainder of the parameters to specify how the retrieve is performed, once you have selected the required objects. These parameters are exactly the same as the parameters in the retrieve policy.

Note that there is no additional retrieve confirmation parameter for this function. It is assumed that when you select the required objects to be retrieved from this display, that you are confirming the retrieve operations, and therefore the retrieves will continue as if in *NOTIFY mode.

If you specify *YES to the Confirm Retrieval parameter (and assuming that this is not a batch operation), you see a confirm display as shown in Figure 59.

```
                          Confirm Retrieve                        RCHAS400

Retrieve select . . . :  *ALL

Type options, press Enter.  Press F16 to confirm all.
  1=Confirm    4=Remove    5=Display

Opt   Library    Object      Member      Volume   Asp  Size (M)  User

__    QUSRSYS    MYFILE      MYFILE      V00001    01     1.05   BILL
__    GLLIB      LEDGER      LEDGER      V00888    02   225.55   TONY
__    PAYLIB     PAYROLL     PAYROLL01   V00999    01    15.05   JIM
```

Figure 59. Confirm Retrieve Display

The Confirm Retrieve display presents a list of files for which retrieve operations were delayed or unsuccessful. You can select and retry, ignore or cancel the retrieve operation for one or more of the files listed. Use option 1 (Confirm) to select and retry the retrieve operation. Use option 4 (Remove) to cancel the retrieve operation. Just leave the option column blank to ignore the retrieve operation. You can run the operation at a later time.

Messages Sent After Retrieves

The RSMRTVBRM function will send a completion message to the initiator of the delayed retrieve. Messages are also sent to the requester when the retrieves are submitted to batch (*SBMJOB). In this topic we are referring to retrieves in *DELAY mode, or retrieves that were suspended due to potential storage threshold overflows or security violations.

Multiple Retrieve Requests for the Same File Member

When a retrieve operation is in delayed mode, a flag associated with the file member in question is set. The RSMRTVBRM command scans the BRMS records to find all files that are “marked” for delayed retrieve. If a file member is set for delayed retrieve and a second request to retrieve it (in delayed mode) is sent, then BRMS checks that file member. BRMS establishes that it is already set for retrieve, and then adds the user profile name of the second requester to its “users to notify” list. A second request for delayed retrieve (even if generated by an authority or storage exception) will not cause a second entry in the RSMRTVBRM display. When the file member is eventually retrieved, all requesters will be notified.

If the second retrieve request is actually a successful *NOTIFY, *VERIFY, or *SBMJOB operation, then BRMS marks that file member as having been retrieved. BRMS removes the file member from the list of file members to be resumed at a later time. The names of the users on delay list will be notified that the successful retrieve operation has occurred.

RSMRTVBRM Submitted to Batch

To increase automation, schedule the submission of the RSMRTVBRM command to occur regularly. Choose a time when there will be little use of the tape drives and other system resources (such as processor and memory).

The automatic submission of this command to batch implies that you would not use the confirm display. You will have to decide which retrieve operations to select for initiation at the time that you submit the command to the scheduler. Consider the following suggestions for inclusion criteria:

- Library *ALL: specify *ALL or library name if you know exactly which libraries contain archived file members.
- Auxiliary Storage Pool ID: You may have an application that you know resides in a specific ASP. When you specify this ASP number, you are ensuring that you do not automatically start any unusual or unexpected retrieve operations. In general, however, you should specify *ALL.
- Retrieve Select *DELAY: in most cases you should automatically initiate only retrieve operations that were purposefully delayed. Using the *ALL parameter would include retrieves that have been stopped because of storage or security considerations. In both of these cases, further operator or administrator action may be required before these retrieves can take place.

Even when submitting the resume retrieve request to batch, have a person check for other file members that may be waiting to be retrieved. These would include file members that you did not include in your selection criteria in the batch submission of the command. It also includes any “failed” retrieves due to authority or storage exceptions. It is not recommended that you include these in the batch submitted command. For example, if you regularly run `RSMRTVBRM *DELAY *RETRIEVE *NO`, then there may be other file members waiting to be retrieved (such as those that have *STORAGE or *SECURITY conditions) which will not be retrieved. You do not need

to perform an interactive “double check” as regularly as the automatic scheduling of the command in a batch environment. There may be further action needed before retrying the failed retrieves (for example: ASP clear up, or security adjustments). You could use the *REPORT option of the RSMRTVBRM command to print any pending retrieves *after* you have run the *DELAY retrieve operation in batch.

This scheduling of the batch submission may be performed for your entire system. If you choose to set the system-wide Retrieve Policy to use a retrieve mode of *DELAY, you are effectively queueing up most of the system retrieve requests until a suitable time for significant tape activity. This approach is useful for good balancing of system resources (in particular your tape devices), but reduces the responsiveness of the retrieve system. It could also lead to a stocking up of a large number of retrieve requests, and the time window for completing all of them may affect other necessary tape activity.

Using RSMRTVBRM Interactively

If you plan to use RSMRTVBRM interactively, then you may or may not choose to use the confirm display.

Without the confirm display, you remove the necessity to make decisions for each file member. However, you will still need to specify your include parameters. Use the same ones as listed for the batch submission above. Note, however, that calling this command interactively without the confirm display will only buy you a possible improved performance and simply tie up an interactive session for an indeterminate amount of time. You can accomplish the same goal by creating a special high priority batch job.

With the confirm display, you gain much more flexibility. You may choose to do any of the following:

1. Use the display to understand which retrieve operations have been delayed due to storage or security exceptions. Take necessary action to resolve these issues, and retry the retrieves.
2. Use the information on the confirm display to help you roughly estimate the restore times and storage level effects of retrieving certain file members. On this basis choose which ones to initiate.
3. Identify which file members should not be retrieved and take action to prevent this from occurring. Your actions can include canceling the retrieve or removing or excluding the file member from your BRMS archive lists.

In each case you must have an administrator who is qualified to make such decisions and take appropriate action to run the RSMRTVBRM command with the confirm display.

Retrieval Performance

The performance of the retrieve function varies according to what you retrieve, whether access paths were saved, the file size, if using multiple physical files with a logical file, and which retrieve mode to use.

Saving Access Paths when Archiving

When you use BRMS to archive data that uses archive with save with storage freed, the option to save the access paths of the file defaults to ***YES**. This is to

avoid lengthy access path rebuilds at the time of retrieve. Having this parameter set to *YES is particularly important when using retrieve modes *NOTIFY and *VERIFY. If you set this parameter to *NO, the access path rebuild time could add to the user wait time. While you can still override the save access path parameter, leave it at *YES to save the access paths and improve the retrieve performance.

File Size

The retrieve of large files takes a longer time than smaller ones. It may be appropriate to break your large files into several smaller ones. Choosing how to divide the file may not be easy. You will have to address the following points:

1. Where are the logical boundaries?

   You may be able to break down the file into groups of records with a common theme. But how would you introduce a group of records to the main file again?
2. How transaction based is the application?

   If the records tend to expire in groups, you may have an opportunity to establish break points.
3. Can the application stand a change of file name?

   Should we group common records in different files?
4. Can the application stand a change of member name?

   Should we group common records in different members within the same file?
5. How normalized are the data entities?

   Can we reduce the size of the file by further normalizing its structure, that is, by splitting the record fields into different files?

Be sure you break files appropriately. If, having split a large file into smaller components, you still require the retrieval of all of those components to satisfy a single data request, then you have made things worse. You will have to perform multiple restores, possibly from multiple volumes (maybe even stored in different locations), adding all the pre- and post-processing and tape mounting overheads for each restore operation.

Multiple Physical Files behind a Logical File

Users should note the following points where a join logical file causes the retrieve of multiple physical files for a single database request:

- Fragmentation

  If the various file members under the logical file have been archived at different times, the archived tape copies will be spread across many different volumes. It is possible that other logical files or processes use these file members, which leads to different dormancy levels. The requested operation will need many separate tape mounts plus additional time for processing the restore for each file member. This may affect performance.
- Predicting retrieval size

  The nature of the retrieve function is to handle one file member restore operation at a time. When one retrieve is being handled, BRMS cannot look forward to predict the next retrieves that will be processed; not even if they are obvious. BRMS will not be able to display *any* application knowledge to predict the incoming retrieve operations. The result is that for a given complex (multiple file) operation, you cannot predict either the total size of all the members to be restored or the time it will take to complete. Therefore, the performance is not predictable.

When you respond to a *VERIFY mode message, you will not know whether you can wait for all the retrieve operations for this request. You will not know whether the request will be forced to end because of imminent ASP overflow. This is despite the fact that for each individual member retrieve you will receive a message that identifies the total restore size. It is the number of other members that are needed for this request that are unknown.

- Access path rebuild times

  Access through a multiple format or join logical file may cause the retrieve of several physical file members. Because the retrieve operation is performed separately and independently of any other operations for each physical file member, an access path rebuild for each physical file member retrieved will be experienced. Thus, a situation could occur where a string of access path rebuilds is performed when one final rebuild would suffice. You cannot override the access path rebuild to *DELAY because you cannot predict which physical file members (if any) will be retrieved. This situation does not apply if you are retrieving in *DELAY mode (see “Retrieval Methods” on page 81 for retrieve mode details), and use the RSMRTVBRM confirm display.

  You should be aware of the potential performance implications of the multiple access path rebuilds that will be caused by archiving physical files under a multiple format logical file.

Which Retrieve Mode to Use for Interactive Applications

Interactive applications that are considered critical to your business may not be good archive candidates. However, if there are *performance critical* applications to your business, it is logical to assume that the best performance for your interactive application can be achieved by using the *NOTIFY mode for retrieving objects. In this way a retrieve is performed immediately at interactive priorities, without waiting for a reply to a message.

This would certainly be true for an application with a fixed logical flow of activities that must be performed to complete a unit of work. If any of the sub-units of this piece of work are temporarily stalled, you have no option but to wait for completion of that sub-unit.

For example: the unit of work may be the processing for a customer that places an order. The first sub-unit may be to retrieve the customer’s details. The second may be to retrieve the stock details of the item required, and the third to create an order.

Customer places an order:

1. Retrieve customer details
2. Check stock levels
3. Create order

The following assumptions may apply to this simplified scenario:

1. The order cannot be created (third sub-unit) until stock data can be retrieved for the part (second sub-unit) required. There must be some data in stock to be able to honor that order.
2. The order cannot be created until the customer data can be retrieved for the customer (first sub-unit). The customer details are needed to fill in certain parts of the order.
3. Sub-unit one and sub-unit two are totally independent of one another.
4. You cannot actually place an order until the order file member is online.

5. You do not know which order file member to open until you have the customer details. You do not know whether to open an order file until you have the part details.

The logical flow of the unit of work implies that you cannot start sub-unit three until both one and two have completed. You can do nothing else while waiting for the order file to be retrieved. Thus, it seems sensible to use *NOTIFY for creating the order.

However, if sub-unit one causes a retrieve operation, then it is sensible to be getting on with something else while the retrieve is being performed. Use the *SBMJOB mode to submit the retrieve to batch and then attempt to run sub-unit two. Similarly, you could submit a retrieve from sub-unit two in batch in the *SBMJOB mode, while you run sub-unit one.

Of course, life is not that simple.

- You may have to change the logic of the application to allow backing out of sub-units to return to them later.
- The performance (productivity) penalties for *SBMJOB include:
  1. The user may not return to a sub-unit as soon as the retrieve completion message is sent.
  2. Retrieve jobs may run in lower priority.

  It is, of course, possible to create a special batch environment for retrieve jobs to speed their performance. BRMS uses the job queue that is named in your job description. You can create a special job queue and refer to it in your job description to change it from the default that is associated with the user profile.
- The retrieve mode used is typically set at the job level. It is possible to alter the retrieve mode for the entire job by issuing the Set Retrieve for BRM (SETRTVBRM) command before each file open operation. This may have unexpected results on other activities within your job (for example with group jobs). It may also require changes to your application.

  You may, however, decide that you are not affecting performance at all by submitting sub-unit three to batch. In this case you may set the mode for this entire job to *SBMJOB.

In summary, it is not always best for productivity (or application performance—depending on how you look at it) to use *NOTIFY. In some cases *SBMJOB may be more appropriate.

Which Retrieve Mode to Use for Batch Jobs

While the *VERIFY mode offers best control of your system resources by forcing a decision to be made for every possible retrieve operation, it does have some considerations:

1. In general, you must be sure that the people who have to respond to the *VERIFY messages are informed enough to make the correct decisions. What you need is a clear overview of all the system components and their relation to each other. For example, one might need to have a good idea of system size, ASP maps, how large an object is, what sort of files are important, what applications are doing, and which applications are important.
2. When you specify *VERIFY in **batch mode**, you must monitor the system operator message queues frequently. A batch job waiting for an operator message reply is a frequent cause of batch throughput problems.

Chapter 5. Using BRMS to Secure Your Hierarchical Storage Management (HSM) Operations

As an HSM administrator, you understand the importance of protecting user and system data from deletion, distortion, and theft. The BRMS Functional Usage Model allows you to customize user access to the HSM functions (archive, retrieve, and migration) and functional components. For example, you can give one user authority to change a specific archive control group, and another the authority only to use and view it. You can also use the functional usage model to grant all users equal access to a particular function or functional component. You must access the BRMS functional usage model through the Operations Navigator interface. You can find instructions on how to use the Operations Navigator with the functional usage model later in this chapter. If Operations Navigator is not installed on your AS/400 and you want to secure your HSM operation, see the ″Other Security Options″ section of this chapter for instruction on how to use other BRMS and OS/400 security options.

BRMS provides effective control over user access to BRMS functionality. Use OS/400 security options to prevent users from causing accidental or intentional damage to your files or system. However, because BRMS ships its databases with *PUBLIC *USE authority (which allows use but not change), irreparable damage is unlikely to occur.

How the Functional Usage Model Works

Use the Functional Usage Model when you want to customize user access to certain functions or when you want to grant all users the same access to a certain function. The functional usage model provides two levels of security for each HSM function and functional component, and specific item (such as control group or backup list):

ACCESS OR NO ACCESS
: In the functional usage model, a user either has access to a BRMS function or component, or does not have access to it. If a user has access to a function or component, that user can use and view it. If a user does not have access to that function or component, then that user cannot use or view it. At this basic level of access, a user can process a specific control group in a backup operation, but does not have the ability to change it.

SPECIFIC CHANGE OR NO CHANGE
: Another level of access allows a user to change a **specific** function, component, or item. For example, to change a backup list, a user must have access to a *specific* backup list. Similarly, to change a media policy, a user must have access to a *specific* media policy. The functional usage model provides lists of existing items (control groups, backup lists, media and move policies, and so on) from which you can grant access to specific items.

With the functional usage model, you can give a user both types of access (so that the user can both use and change a particular function, component, or item). Or you can give a user just one type of access, for example, access to use, but not change, a particular function, component, or item.

Note: The functional usage model works in conjunction with Operations Navigator. If you want to secure your HSM operation, and do not have Operations

© Copyright IBM Corp. 1998, 1999

Navigator installed on your system, see the ″Other Security Options″ section of this chapter for more information. Also note that though you need Operations Navigator to use the functional usage model, you do not need it to use the SETUSRBRM command.

Archive Function

In the archive area, the following usage levels are available:

Basic Archive Activities
This is the basic level of archive usage. Those with Basic Backup Activities access can use and view the archive policy, control groups, and lists. Use access also allows these users to process archive activities by using control groups (through the STRARCBRM command). A user without access to basic backup activities cannot see archive menu options or command parameter options.

Archive Policy
Users with access to the Archive Policy can change the archive policy, in addition to using and viewing it. Users without access to the archive policy cannot change it.

Archive Control Groups
Users with access to the Archive Control Groups can change specific archive control groups, in addition to using and viewing them. You can find a list of all of your existing control groups under the archive control groups heading in Operations Navigator. You can grant a user to any number of specific control groups. Users without access to the archive control groups cannot change them.

Archive Lists
Users with access to the Archive Lists can change specific archive lists, in addition to using and viewing them. You can find a list of all of your existing backup lists under the backup lists heading in Operations Navigator. You can grant a user access to any number of specific backup lists. Users without access to a backup list cannot change it..

Migration Function

In the migration area, the following usage levels are available:

Basic Migration Activities
This is the basic level of migration usage. Those with Basic Migration Activities access can use and view the migration policy, control groups, and lists. Use access also allows these users to process archive activities by using control groups (through the STRMGRBRM command), or by using commands (MGRBRM, MOVSPLFBRM). A user without access to basic backup activities cannot see migration menu options or command parameter options.

Migration Policy
Users with access to the Migration Policy can change the migration policy, in addition to using and viewing it. Users without access to the migration policy cannot change it.

Migration Control Groups
Users with access to the Migration Control Groups can change migration control groups, in addition to using and viewing them. You can find a list of all of your existing control groups under the migration control groups

heading in Operations Navigator. You can grant a user to any number of specific control groups. Users without access to the migration control groups cannot change them.

Migration Information
Users with access to Basic Migration Activities can display migration history information (by using the Work with Migration Info (WRKMGRIBRM) display. Users with Migration Information access can change the history (for example, by deleting an entry).

Retrieve Function

In the retrieval area, the following functional usage levels are available:

Basic Retrieval Activities
This is the basic level of retrieval usage. Those with Basic Retrieval Activities access can use and view the retrieval policy. Users with basic access can also process a retrieval through the Set Retrieve Controls for BRM (SETRTVBRM) command, though they cannot change any of its parameters. A user without access to basic retrieval activities cannot see or use retrieval-related menu options or command parameter options.

Retrieve Policy
Users with access to the Retrieval Policy can change the retrieve policy, in addition to using and viewing it. Users with access to the retrieve policy can also change the parameters on the SETRTVBRM and Resume Retrieve using BRM (RSMRTVBRM) commands. Users without access to the retrieve policy cannot change it.

Note: See the *Backup Recovery and Media Services for AS/400* book for information on how to apply the functional usage model to the standard BRMS functions and functional components.

Implementing the Functional Usage Model

Use the AS/400 Operations Navigator interface to access the functional usage model feature. To get to a list of HSM functions and components, perform the following steps from your Operations Navigator window:

1. Highlight your system name and, from the File menu, select Application Administration. If you have Operations Navigator installed, but do not see the Application Administration feature, then you need to reinstall it.
2. When the Application Administration dialog appears, click the Host Applications tab.
3. From the Host Applications window select the Backup, Recovery, and Media Services for AS/400 product.
4. Click the BRMS icon to expand the tree structure to see the different functions and components available. You can see the HSM functions and components listed along with the standard BRMS functions and components.

Following is a summary of what you can see or do on the Host Applications display:

- In the **Function** column, you can see the names of the HSM functions and functional components.
- If you highlight a box in the **Default Usage** column, you are granting all users access to that function or functional component. An X in a Default Usage box grants all users access to the function or component to the left of the box.

- If you type a checkmark in the **All Object Usage** column, you are indicating that a user or number of users have *ALLOBJ authority at the user level. The OS/400 operating system provides a special *ALLOBJ security value, which allows user access to all objects on the AS/400. A user with (*ALLOBJ) authority automatically has complete access to all of the HSM functions and components that you mark. If you do not type a checkmark in this box, then BRMS ignores the users *ALLOBJ authority and requires customized access.
- To customize user access, click on the **Customize** button on the bottom right of the display. From the Customize display, you can select individual users by name and tailor their access as needed. When you customize user access to any of the functions or functional components, an X appears in the **Customize Access** column on the Host Applications display.

Registering New HSM Activities with the Functional Usage Model

Once you establish the functional usage model for HSM users, BRMS begins registering new entries with the OS/400 security system. Each time you create a new HSM control group, list, or related entry, BRMS adds the item to the usage model. When this happens, BRMS records the name of the new item in the log. You can use the DSPLOGBRM *SEC command to review the log message. BRMS registers each of the new items with the default usage level that you specified in the system policy.

To get to the Change System Policy display, take the following steps:

1. Type GO BRMSYSPCY at a command line and press Enter.
2. Select option 1 (Display or change system policy).

```
                          Change System Policy


Type choices, press Enter.
Media policy . . . . . . . . . . . . . . FULL          Name, F4 for list
Devices  . . . . . . . . . . . . . . . . TAPMLB04      Name, F4 for list
Home location for media  . . . . . . . . *HOME         Name, F4 for list
Media class  . . . . . . . . . . . . . . CART3490E     Name, F4 for list
Sign off interactive users . . . . . . . *NO           *YES, *NO
Sign off limit . . . . . . . . . . . . .  30          0-999 minutes
Output queue . . . . . . . . . . . . . . *PRTF         Name, *PRTF
Library  . . . . . . . . . . . . . .                   Name, *LIBL
Day start time . . . . . . . . . . . .  0:00:00      Time
Media monitor  . . . . . . . . . . . . . *YES          *YES, *NO
Shared inventory delay . . . . . . . . .  60          30-9999 seconds
Auto enroll media  . . . . . . . . . . . *NO           *NO, *YES
Trace  . . . . . . . . . . . . . . . . . *NO           *NO, *YES
Default usage. . . . . . . . . . . . . . *NO           *NO, *YES
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel    OS/400
```

Figure 60. The Change System Policy display

3. The value in the *Default usage* field specifies whether to give default access to the specific BRMS object that you are creating, for example, a control group, archive list, or policy. functional authority to each new activity that is registered through BRMS. The values for this prompt are *YES and *NO. If you set the *Default usage* value to *NO, only you can access that control group. No other user can use it. If you specify *YES, users can access the control group at the usage level they were given when the administrator set up the Functional Usage Model.

4. Press Enter to save your changes and exits.

Working with the SETUSRBRM Command

As indicated above, use the functional usage model to ***customize*** access by user or when you want to grant all users access to a certain function. Use the SETUSRBRM command to grant system operators and administrators access to BRMS-defined activities. BRMS bases these default operator and administrator categories on the kinds of activities that are usually performed by users in these roles.

To get to the Set User Usage for BRM (SETUSRBRM) display, take the following steps:

1. Type SETUSRBRM at a command line and press F4 to prompt the Set User Usage for BRM display.

```
                     Set User Usage for BRM (SETUSRBRM)

Type choices, press Enter.

User . . . . . .                    Character value, *DEFAULT
Usage  . . . . . .   *OPERATOR      *OPERATOR, *ADMIN









                                                                    Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 61. The SETUSRBRM command

2. In the *Usage* field, select either the *OPERATOR or the *ADMIN value. An *OPERATOR can access the following activities:
   - Basic backup activities
   - Basic archive activities
   - Basic migration activities
   - Basic media activities
   - Basic movement activities

   If you select the *ADMIN value, the user named in the *User* field has use and change access to all of the BRMS and HSM functions and functional components.
3. Press Enter to save your changes.

Other Security Options

If you have BRMS installed on your system, but do not have Operations Navigator, you can use the following options to secure selected aspects of your HSM operation. You can also use these security options in combination with the SETUSRBRM command to set user access.

Retrieve and Migration Authority

The retrieve policy allows you to set the authority level requirements for users to initiate a retrieve. You can find this parameter (Retrieve authorization) on the Set Retrieve Controls for BRM (SETRTVBRM) command. Any user with authority to access the SETRTVBRM command can alter the value of the *Retrieve authority* parameter for their job.

For example, you might choose to alter the parameter to allow users with *USE authority to retrieve that object. You can do that by setting the parameter to *USE, which allows the user to create objects (using retrieve) that they ordinarily could only read.

You might think that this is not a correct use of the security parameter. You can assume that an archived object already exists on the system if you include the archive tapes as part of system storage. Thus, a retrieve operation simply *moves* an object from one part of single level storage (tape) to another (disk). In this case, we are conceptually talking about a storage management *move* task rather than an OS/400 *create* task. In this case, because it is necessary to move an object from tape to disk to enable read access, you can justify the apparent create operation even for a user who only has read access to an object (that is, *USE authority).

Either way, the fact remains that users may change this parameter as they want. The administrator *should not rely on this parameter to restrict the availability of the retrieve function*.

If you want to restrict the use of the retrieve function, then you need to restrict access to the SETRTVBRM command. You can do this by using standard OS/400 security access controls for the command object (type *CMD). This also restricts user ability to change other retrieve controls such as the retrieve mode to use. When you restrict authority to the SETRTVBRM command, you also need to restrict the WRKPCYBRM command to inhibit users from changing the retrieve policy. You can find information on how to secure the WRKPCYBRM command later in this section.

Restore Options

When restoring a file member as part of a retrieve operation, you want to protect your system from overwriting a recreated file member, or from importing incorrect versions of it. Remember that the BRMS retrieve function works with a name oriented inventory. When you rename, delete, or recreate file members, you can cause unpredictable results.

You can use the *Options* parameter in the retrieve policy and the SETRTVBRM command to reduce the chances of retrieving inappropriate data. The *Option* parameter allows you to specify which objects to retrieve, depending on whether the objects exist on the system.

The *Allow Object Differences* parameter can assist you in preventing the restore of deleted and subsequently recreated file members. If you set the parameter to *NONE, then BRMS cross-checks the create timestamp and owner information before allowing the retrieve. Thus, if you submit a delayed retrieve after recreating a member, then the retrieve operation fails.

Securing the Retrieve and Migration Functions

Effective change management is an important part of your quality process. Securing the retrieve function can help you manage and control your disk capacity and your tape activity. You can use the Grant Object Authority (GRTOBJSUT) command to grant users access to these BRMS, and OS/400 commands. Use the Remove Object Authority (RMVOBJAUT) command to remove access to these commands.

To secure the retrieve function, take the following steps:

1. Secure the WRKPCYBRM *RTV and *MGR commands

   To prevent users from adjusting the systemwide retrieve policy, use the standard OS/400 security functions to restrict authority to the WRKPCYBRM command (object type *CMD). This prevents unauthorized users from using any part of the command.

   You can consider implementing this as part of a global restriction to all BRMS commands. Remember to identify the key personnel that require access to these commands before you revoke the authority.
2. Secure the SETRTVBRM, STRMGRBRM, and MGRBRM commands

   Revoke authority of all non-BRMS administrative personnel from the SETRTVBRM, STRMGRBRM, and MGRBRM commands. Remember that this also removes their ability to alter other parameters such as the retrieve mode.
3. Set up users with an initial program

   If certain users need to occasionally alter retrieve or migration policy parameters, consider the following provisions:

   a. Include SETRTVBRM or STRMGRBRM and MGRBRM commands in the initial program for the required users.
   b. When compiling this initial program, set the run authority of the program to *OWNER. This adopts the authority of the program object owner. You can change this parameter after compiling it by using the CHGPGM command.
   c. Compile the initial program under a user profile that has authority to the SETRTVBRM, STRMGRBRM, or MGRBRM commands. (You can change this parameter after compiling it by using the CHGOBJOWN command.)
   d. You can then restrict the authority to the program object itself.

Chapter 6. Suggested Hierarchical Storage Management Strategies

This chapter provides examples of how to implement a hierarchical storage management strategy with BRMS. Use these examples as a starting point to tailor your HSM operations to best suit your specific business needs.

The examples provide guidance in developing an HSM strategy, and then provide assistance on how to implement the strategies through the use of BRMS functions, control groups, and policies. The setup required for the following strategies may not be repeated in this chapter. See the appropriate references for more information.

The following list may require additional setup:

1. Disk compression - See the disk compression section of the *Backup and Recovery* book.
2. User ASPs - See the user ASP section of the *Backup and Recovery* book.
3. Automated Tape Library - See the *Automated Tape Library Planning and Management* book.
4. If you are using BRMS:
   - Set up ASPs managed by BRMS, see “Setting Up ASP Classes” on page 23.
   - Set up a migration policy and control groups, see “Setting Up Migration Using BRMS” on page 16.
   - Set up an archive policy and control groups, see “Setting Up Archive with Dynamic Retrieval Using BRMS” on page 56.
   - Set up a retrieve policy and control groups, see “Setting Up Dynamic Retrieval with BRMS” on page 76.

Suggestions:

1. You may want to run a report first to see a list of candidates for migration and archive operations prior to actually performing the migration or archive operation. Use WRKSPLF to view the report.
2. To automate migration or archive operations, you may consider adding them to the job scheduler you have on your system.
3. When listing the set of criteria to use in the migration policy, ASP class attributes, or the migration control group attributes, the order of the criteria listed is important. For each item in the migration control group, BRMS will evaluate its eligibility for migration using the criteria starting from the top of the list to the bottom of the list. The migration will be based on the first criteria it can satisfy. When planning, consider the following:
   - Put the higher priority criteria on the top of the list.
   - When you specify a greater than value criteria, list the order in the migration criteria from the largest number to the smallest number.
   - When specifying a less than value criteria, the order in the migration criteria should be listed from the smallest number to the largest number.
4. When defining the migration criteria, specifying a criterion on the Object column, such as *ANY, *ALL, % value, or number of objects will cause BRMS to satisfy the condition by looking at the attributes from the object or document level. Specify a criterion for this field for more accurate processing. If you require that all the processing is done at an object or document level, then you should

© Copyright IBM Corp. 1998, 1999

specify a criterion in the Object field on every migration criteria. If you leave the Object field blank, BRMS will look at the attributes in the library or root folder level only.

5. When you specify the most recent use of an object as a criterion for the migration rules, specify both the **date last used** and the **change date**. The date last used is set to blanks on a scratch install and after the migration of a library or a folder. Not all commands that cause an update changes the date last used. This tip applies to archival function as well. When specifying the inactivity criteria for archive, indicate *BOTH for the Date type parameter.
6. You can specify the migration criteria in any of the following ways:
   - When the criteria applies to the items in a control group, specify this in the control group attributes. For example, demote library MYLIB to a user ASP when the library is older than 30 days old.
   - When the criteria applies to the ASP class, specify this in the ASP class attributes. For example, demote all the libraries to a user ASP that meets the criteria in the migration control group that are older than 30 days.
   - When the criteria apply to the overall migration operation, then change the migration policy. The migration control group overrides the ASP class attributes, and the ASP class attributes override the Migration Policy.
7. Only one migration operation (MGRBRM, STRMGRBRM, and MOVSPLFBRM commands) is allowed at one time.
8. Run your migration operation during non-peak hours.

Examples of Strategies

Strategies include:

1. Example 1 - Archiving members of a file to tape
2. Example 2 - Archiving documents and stream files to tape
3. Example 3 - Moving spooled files to a user ASP
4. Example 4 - Migrating a user library to user ASP
5. Example 5 - Migrating root folders to a user ASP
6. Example 6 - Staging migration of user libraries to user ASPs that is based on migration control group attributes
7. Example 7 - Staging migration of user libraries to user ASPs that is based on ASP class attributes
8. Example 8 - Migrating root folders to a user ASP, then archiving documents to tape
9. Example 9 - Migrating user libraries to a user ASP, then archiving files and members to tape.
10. Example 10 - Migrating user libraries back to a high performance ASP.

We recommend that you run a report prior to running your HSM examples. To do this use STRMGRBRM *REPORT for migration, MOVSPLFBRM *REPORT for spooled file, and the STRARCBRM *REPORT for archive by specifying the control groups you set up.

Example 1

Archiving members of a file to tape based on frequency of usage.

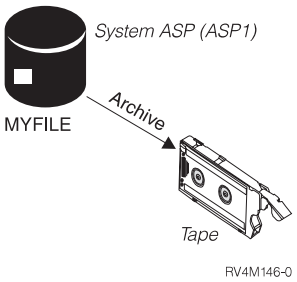

Figure 62. Example 1

This example demonstrates how to archive members of a file to tape. The members of this file have been used less than twice per month.

1. Define an archive object list that is called OBJLST.

```
                      Add Object List
                                                          System:   RCHAS400
List name . . . . . OBJLST__
                                        Selection
 Seq   Library      Object   Type       *INC/*EXC
 10    MYLIB        MYFILE   *ALL          *INC
 __    _______      ______   ____          ____
```

Figure 63. Add Object List Display

   This is basically a list of objects you want to archive.

2. Create an archive control group that is called ARCOBJ.

```
                    Create Archive Control Group Entry
                                                          System:   RCHAS400
Group . . . . . ARCOBJ__
                                 Weekly
        Archive        List      Activity
 Seq    Items          Type      SMTWTFS
 10     OBJLST         *OBJ         * *
 __     _______      ______      _______
```

Figure 64. Create Archive Control Group Entry Display

   The control group points to the object list, OBJLST, and specifies how often to run the archive control group.

3. Change the archive control group attributes for ARCOBJ.

```
                  Change Archive Control Group Attributes

Group  . . . . . . . . . . . . . . . . : ARCOBJ

Type information, press Enter.

Include:
  ASP low storage threshold  . . . . . . *ARCPCY     *ASP, 0-99%, *ARCPCY
  Date type for archival . . . . . . . . *BOTH       *BOTH, *CHANGE, *USE...
    Inactivity limit . . . . . . . . . . *ARCPCY     0-9999 days, *ARCPCY
  Object size greater than . . . . . . . *ARCPCY     0-999999 MB, *ARCPCY
  Objects able to be freed . . . . . . . *ARCPCY     *YES, *NO, *ARCPCY
    Retain object description  . . . . . *ARCPCY     *YES, *NO, *ARCPCY
  Objects not able to be freed . . . . . *ARCPCY     *YES, *NO, *ARCPCY
  Days used (per month) less than. . . . 2          0-31, *ARCPCY









                                                                        More...
F3=Exit   F12=Cancel
```

Figure 65. Change Archive Control Group Attributes Display

4. Issue STRARCBRM CTLGRP(ARCOBJ) OPTION(*ARCHIVE) command.
5. Issue DSPJOBLOG command to verify that the archive completed successfully.

Members in MYFILE that have been used less than twice per month will now be archived.

Example 2

Archiving documents and stream files to tape based on inactivity limit.

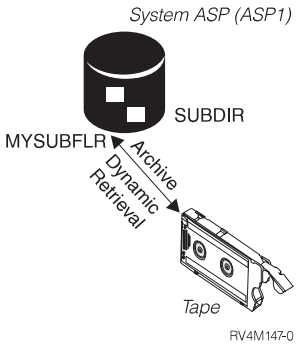

Figure 66. Example 2

This example demonstrates archiving office documents in MYSUBFLR and stream files in directory SUBDIR to tape that have not been used and changed in the last 3 months.

1. Define an archive folder list that is called FLRLST.

```
                      Add Folder List
                                                          System:   RCHAS400
List name    FLRLST__
                                        Selection   Subfolder
 Seq    Folder                          *INC/*EXC   *YES/*NO
 10     MYFOLDER/MYSUBFLR                 *INC        *YES
 __     ____________                      ____        ____
```

Figure 67. Add Folder List Display

This list will archive all documents in subfolder, MYSUBFLR, and in folder MYFOLDER.

2. Define an archive link list that is called LNKLST.

```
                      Add Link List
                                                          System:   RCHAS400
List name    LNKLST__

Text         Archive Link List____________________
```

Figure 68. Add Link List Display

This is a two step process, the first step you define the name and text. The second step you will need to change the entry.

3. To change the entry, use option 2 and include the directories you want to archive.

```
                      Change Link List
                                                          System:   RCHAS400
List name    LNKLST__

Objects
   Name . . . . . :  MYDIR/SUBDIR
   Include. . . . . :  *INCLUDE
Directory Subtree . . :  *ALL
```

Figure 69. Change Link List Display

This list will archive all stream files in subdirectory, SUBDIR, and in directory MYDIR.

4. Create an archive control group that is called DOCLNK.

```
                      Create Archive Control Group Entry
                                                          System:   RCHAS400
Group   DOCLNK__
                                  Weekly
        Archive       List        Activity
 Seq     Items        Type        SMTWTFS

 10     FLRLST        *FLR        *DFTACT
 20     LNKLST        *LNK        *DFTACT
 __     _______      ______       _______
```

Figure 70. Create Archive Control Group Entry Display

The control group points to both the folder list, FLRLST, and the directory list, LNKLST.

5. Change the archive control group attributes for DOCLNK.

Figure 71. Change Archive Control Group Attributes Display

```
                  Change Archive Control Group Attributes

Group  . . . . . . . . . . . . . . . . : DOCLINK

Type information, press Enter.

Include:
  ASP low storage threshold  . . . . . . *ARCPCY     *ASP, 0-99%, *ARCPCY
  Date type for archival . . . . . . . . *BOTH       *BOTH, *CHANGE, *USE...
    Inactivity limit . . . . . . . . . . 90          0-9999 days, *ARCPCY
  Object size greater than . . . . . . . *ARCHPCY    0-999999 MB, *ARCPCY
  Objects able to be freed . . . . . . . *ARCPCY     *YES, *NO, *ARCPCY
    Retain object description  . . . . . *ARCPCY     *YES, *NO, *ARCPCY
  Objects not able to be freed . . . . . *ARCPCY     *YES, *NO, *ARCPCY
  Days used (per month) less than. . . . 31          0-31, *ARCPCY









                                                                      More...
F3=Exit   F12=Cancel
```

When you specify *BOTH for the Date type, the date last used and the change date determine the days of inactivity. In this case, any document in MYFOLDER/MYSUBFLR and any stream files in MYDIR/SUBDIR will be archived if inactive for 90 days.

6. Issue STRARCBRM CTLGRP(DOCLNK) OPTION(*ARCHIVE) command.
7. Issue DSPJOBLOG command to verify that archive completed successfully.

When doing operations such as viewing documents in MYFOLDER/MYSUBFLR or when checking out (CHKOUT) stream files in MYDIR/SUBDIR, BRMS will dynamically retrieve from tape to disk. For a list of commands that may or may not cause a dynamic retrieval to occur, see “Operations that Call or Do Not Call Dynamic Retrieval” on page 52.

Example 3

Moving spooled files to a user ASP based on user and creation date.

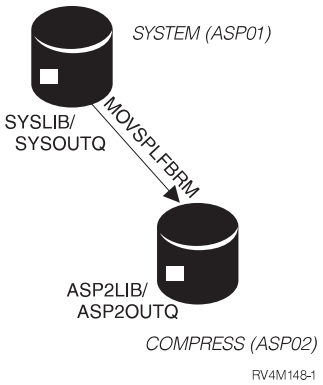

Figure 72. Example 3

This example moves spooled files from the system ASP to ASP02 based on user and creation date. Run the MOVSPLFBRM *REPORT command for a list of candidates.

1. Create an output queue ASP2OUTQ into library ASP2LIB.

   This is assuming that library ASP2LIB is in user ASP, ASP02. You must create the output queue, ASP2OUTQ, with *OUTASP in the Spooled File ASP parameter.
2. Define a user ASP that is called COMPRESS. To set up ASP class and description, see “Setting Up ASP Classes” on page 23.

   This will define the attributes such as the priority, low and high thresholds.
3. Issue the MOVSPLFBRM command.

```
                  Move Spooled Files using BRM (MOVSPLFBRM)

Type choices, press Enter.

Option . . . . . .   *MOVE          *REPORT, *MOVE
To output queue  . . . . . .   ASP2OUTQ       Name
  Library  . . . . . .   ASP2LIB_     Name
From output queue  . . . . . .   SYSOUTQ       Name, generic*, *ALL
  Library  . . . . . .   SYSLIB       Name, *ALL
File . . . . . .   *ALL           Name, *ALL
Job name . . . . . .   *ALL           Name, *ALL, *
User . . . . . .   FINANCE        Name, *ALL
User data  . . . . . .   *ALL           Name, *ALL
Auxiliary storage pool . . . . .   *SYSTEM        Name, *ALL, *SYSTEM, 1, 2...
Select create date:
  From date  . . . . . .   04/01/98       Date, *CURRENT, *BEGIN, nnnnn
  To date  . . . . . .   04/08/98       Date, *CURRENT, *END, nnnnn
Select last used date:
  From date  . . . . . .   *BEGIN         Date, *CURRENT, *BEGIN, nnnnn
  To date  . . . . . .   *END____       Date, *CURRENT, *END, nnnnn
```

```
                  Move Spooled Files using BRM (MOVSPLFBRM)

Type choices, press Enter.

Select size:
  Size type  . . . . . .   *KB           *KB, *MG, *GB, *PAGES
  Minimum size . . . . . .   0            0-999999
  Maximum size . . . . . .   *NOMAX       0-999999, *NOMAX
```

Figure 73. Move Spooled Files using BRM (MOVSPLFBRM) Command

This will move spooled files from the system ASP (ASP01) to user ASP (ASP02) if the spooled file was created by user FINANCE on or between the dates of 04/01/98 and 04/08/98. Depending on your application, you may need to specify the new library and new output queue in a user ASP.

4. This will create a report that will show every spooled file that was moved. Issue the command WRKSPLF. The printer file is QP1AMSF.

Example 4

Migrating a user library to a user ASP (demoted) based on size, date last used and change date.

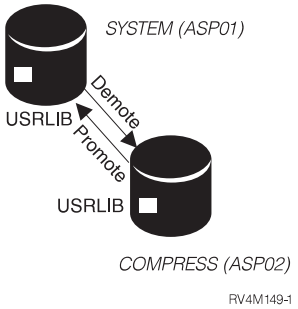

Figure 74. Example 4

Libraries generated by an application can be migrated to a user ASP that is called COMPRESS based on size and if date last used or change date is greater than 90 days.

1. Define a user ASP that is called COMPRESS. To set up ASP class and description, see “Setting Up ASP Classes” on page 23.
2. Create a migration control group that is called MGRLIB.

```
                      Create Migration Control Group Entry
                                                          System:   RCHAS400
Group . . . . . . . . . . :   MGRLIB

          Migration         Item
 Seq      Items             Type
 10       USRLIB            *LIB
___       _______           ______
```

Figure 75. Create Migration Control Group Entry Display

This migration control group will migrate library USRLIB.

3. Change the attributes for the migration control group, MGRLIB.

```
                        Change Migration Control Group Attributes
                                                          System:   RCHAS400

 Group. . . . . . . . . . . . .   MGRLIB___

 Type choices, press Enter.

   Migration     Item Type    Attribute    Test      Value      Objects
   *DEMOTE _     *LIB_____    *SIZE____    *GT____   300____    80%____
   _________     _____*AND    *LASTUSE_    *GT____   90_____    80%____
   _________     ______*OR    *LASTCHG_    *GT____   90_____    80%____
   _________     _________    _________    _______   _______    _______
   _________     _________    _________    _______   _______    _______
```

Figure 76. Change Migration Control Group Attributes Display

This will migrate library, USRLIB, if 80% of all the objects in library USRLIB are greater than 300 MB and where 80% of all the objects in USRLIB have not been used or changed in the last 90 days.

4. Issue STRMGRBRM CTLGRP(MGRLIB) OPTION(*MIGRATE) command.
5. Issue DSPJOBLOG command to verify that migration completed successfully.
6. To promote the library back to ASP01, SYSTEM, change the migration control group attributes for MGRLIB.

```
                        Change Migration Control Group Attributes
                                                          System:   RCHAS400

Group. . . . . . . . . . . . .   MGRLIB___

Type choices, press Enter.

  Migration     Item Type    Attribute    Test      Value      Objects
  *PROMOTE_     *LIB_____    *LASTCHG     *LT____   7_______   *ANY___
  *DEMOTE _     *LIB_____    *SIZE____    *GT____   300____    80%____
  _________     _____*AND    *LASTUSE_    *GT____   90_____    80%____
  _________     ______*OR    *LASTCHG_    *GT____   90_____    80%____
  _________     _________    _________    _______   _______    _______
  _________     _________    _________    _______   _______    _______
```

Figure 77. Change Migration Control Group Attributes Display

7. Run the STRMGRBRM *MIGRATE command again. This will promote library USRLIB to ASP01, SYSTEM, if any objects in the library were changed in the last 7 days.

Example 5

Migrating root folders to a user ASP based on age.

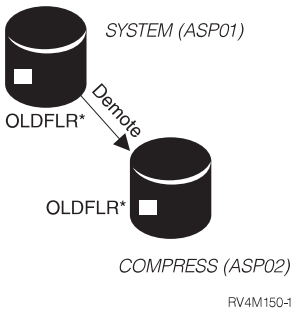

Figure 78. Example 5

The system ASP is exceeding the system high threshold of 85% and migrating folders that start with OLDFLR* that are older than 60 days old will help decrease storage on system ASP.

1. Define a user ASP that is called COMPRESS. To set up ASP class and description, see “Setting Up ASP Classes” on page 23.
2. Create a migration control group that is called MGRFLR.

```
                      Create Migration Control Group Entry
                                                          System:   RCHAS400
Group . . . . . :   MGRFLR

          Migration         Item
 Seq       Items            Type
 10       OLDFLR*           *FLR
___       _______           _____
```

Figure 79. Create Migration Control Group Entry Display

This migration control group will migrate all root folders that start with OLDFLR*.

3. Change the attributes for the migration control group, MGRFLR.

```
                        Change Migration Control Group Attributes
                                                          System:   RCHAS400

 Group. . . . . .   MGRFLR___

 Type choices, press Enter.

   Migration     Item Type    Attribute    Test      Value      Objects
   *DEMOTE _     *FLR_____    *AGE____    *GT____   60_____    *ALL___
   _________     _________    _________    _______   _______    _______
   _________     _________    _________    _______   _______    _______
   _________     _________    _________    _______   _______    _______
```

Figure 80. Change Migration Control Group Attributes Display

This will migrate any root folder that starts with OLDFLR* only if all objects within the folder have an age greater than 60 days.

4. Issue STRMGRBRM CTLGRP(MGRFLR) OPTION(*MIGRATE) command.
5. Issue DSPJOBLOG command to verify that migration completed successfully.

Example 6

Staging migration of user libraries to ASPs based on migration control group attributes.

Note: For the remainder of the examples, we have introduced another user ASP, ASP02. ASP02 will now have ASP Class name MEDIUM, and ASP03 will have ASP Class name COMPRESS to show the staging of objects between ASP's.

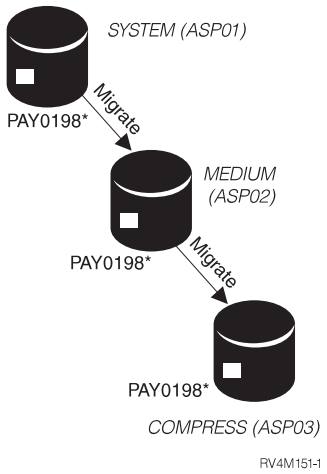

Figure 81. Example 6

This example will migrate PAY0198 libraries from SYSTEM ASP to user ASP, MEDIUM, when any objects in the libraries are greater than or equal to 7 days old yet less than 14 days old. Then when the objects are between 14 and 21 days old, they will be migrated from user ASP MEDIUM to user ASP COMPRESS.

1. Define user ASPs
   - Define ASP class and attributes for user ASPs, MEDIUM and COMPRESS. See “Setting Up ASP Classes” on page 23.
2. Create a migration control group, MGRPAYRLL.

```
                      Create Migration Control Group Entry
                                                          System:   RCHAS400
Group . . . . . . . . . . :   MGRPAYRLL

          Migration         Item
 Seq       Items            Type
 10       PAY0198*          *LIB
___       _______           ______
```

Figure 82. Create Migration Control Group Entry Display

This will create a migration control group for all PAY0198 libraries.

3. Change the migration control group attributes for MGRPAYRLL.

```
                        Change Migration Control Group Attributes
                                                          System:   RCHAS400

Group. . . . . . . . . . . . .   MGRPAYRLL___

Type choices, press Enter.

  Migration     Item Type    Attribute    Test      Value      Objects
  MEDIUM___     *LIB_____    *AGE_____    *GE____   7______    *ANY___
  _________     _____*AND    *AGE_____    *LT____   14_____    *ANY___
  COMPRESS_     *LIB_____    *AGE_____    *GE____   14_____    *ANY___
  _________     _____*AND    *AGE_____    *LT____   21_____    *ANY___
  _________     _________    _________    _______   _______    _______
```

Figure 83. Change Migration Control Group Attributes Display

This migration control group will migrate (demote) objects to ASP class MEDIUM if any objects in libraries PAY0198 have a creation date that is greater than or equal to 7 days old and less than 14 days old. It will then migrate PAY0198 libraries to ASP class COMPRESS if any objects in the library have a creation date that is between 14 and 20 days old.

4. Issue STRMGRBRM CTLGRP(MGRPAYRLL) OPTION(*MIGRATE) command.
5. Issue DSPJOBLOG command to verify that migration completed successfully.

Example 7

Staging migration of user libraries to user ASPs based on ASP class attributes.

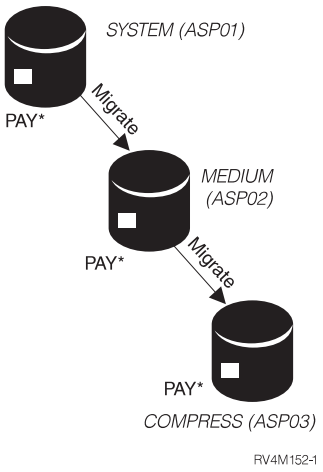

Figure 84. Example 7

This example is similar to Example 6. However, the migration is based on ASP attributes versus control group attributes. This example will migrate libraries starting with PAY* to user ASP, MEDIUM, if any of the objects within PAY* libraries are greater than or equal to 7 days old yet less than 14 days old. Then when the objects are between 14 and 21 days old, they will be migrated from user ASP MEDIUM to user ASP COMPRESS.

1. Define user ASPs

- Define ASP class and attributes for user ASPs, MEDIUM and COMPRESS. See “Setting Up ASP Classes” on page 23.

2. Change the ASP Class attributes for ASPs SYSTEM and MEDIUM.

```
                        Change ASP Class Attributes              RCHAS400

ASP class  . . . . . .  SYSTEM

Type choices, press Enter.

    Migration     Item Type    Attribute    Test      Value      Objects
    *DEMOTE__     *LIB_____    *AGE_____    *GE____   7_____     *ANY___
    _________     _____*AND    *AGE_____    *LT____   14____     *ANY___
    _________     _________    _________    _______   _______    _______
```

Figure 85. Change ASP Class Attributes Display

```
                        Change ASP Class Attributes              RCHAS400

ASP class  . . . . . .  MEDIUM

Type choices, press Enter.

    Migration     Item Type    Attribute    Test      Value      Objects
    DEMOTE___     *LIB_____    *AGE_____    *GE____   14____     *ANY___
    _________     _____*AND    *AGE_____    *LT____   21____     *ANY___
    _________     _________    _________    _______   _______    _______
```

Figure 86. Change ASP Class Attributes Display

Define the attributes for ASP classes SYSTEM and MEDIUM. Any libraries that reside in the SYSTEM ASP will migrate from the SYSTEM ASP to ASP MEDIUM that meet the *AGE criteria. Any libraries that reside on ASP MEDIUM will migrate from ASP MEDIUM to ASP COMPRESS that meet the *AGE criteria.

3. Create a migration control group, MGRPAYR1.

```
                      Create Migration Control Group Entry
                                                          System:   RCHAS400
Group . . . . . :   MGRPAYR1

          Migration         Item
 Seq      Items             Type
 10       PAY*              *LIB
___       _______           ______
```

Figure 87. Create Migration Control Group Entry Display

As previously defined in example 6, this migration control group MGRPAYR1 will migrate libraries that start with PAY*.

4. Change the attributes for the migration control group, MGRPAYR1.

```
                        Change Migration Control Group Attributes
                                                          System:   RCHAS400

Group. . . . . . . . . . . .   MGRPAYR1___

Type choices, press Enter.

  Migration     Item Type    Attribute    Test      Value      Objects
  *DEMOTE__     *ASPCLS__    _________    _______   _______    _______
  _________     _________    _________    _______   _______    _______
```

Figure 88. Change Migration Control Group Attributes Display

The control group attributes will migrate (demote) the libraries PAY* based on the ASP class atributes.

5. Issue STRMGRBRM CTLGRP(MGRPAYR1) OPTION(*MIGRATE) command.
6. Issue DSPJOBLOG command to verify that migration completed successfully.

Libraries PAY* will migrate from SYSTEM ASP to MEDIUM ASP if any objects in the libraries are between 7 and 13 days old and from SYSTEM ASP to COMPRESS ASP if any objects in the libraries are between 14 and 21 days old.

Example 8

Migrating root folders to a user ASP based on creation date and then archiving the documents to tape based on inactivity.

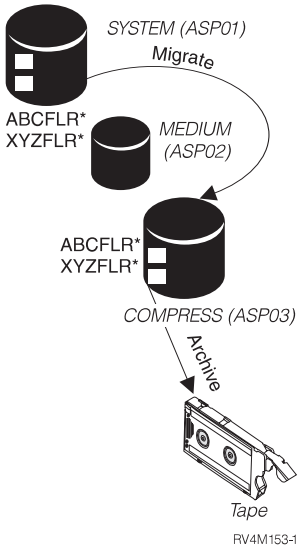

Figure 89. Example 8

This example will demonstrate migrating user root folders directly to ASP COMPRESS that were created prior to 05/31/99 and then archiving the documents to tape that have not been changed in the last 9 months.

1. Create a migration control group that is called MGRFNANCE.

```
                    Create Migration Control Group Entry
                                                          System:   RCHAS400
Group . . . . . . . . . . :   MGRFNANCE

         Migration         Item
 Seq      Items            Type
 10      ABCFLR*           *FLR
 20      XYZFLR*           *FLR
```

Figure 90. Create Migration Control Group Entry Display

This will generate a list of migration candidates of root folders that start with ABCFLR* and XYZFLR*.

2. Define the attributes for control group, MGRFNANCE.

```
                        Change Migration Control Group Attributes
                                                          System:   RCHAS400

 Group. . . . . . . . . . . . .   MGRFNANCE__

 Type choices, press Enter.

   Migration     Item Type    Attribute    Test      Value      Objects
   COMPRESS_     *FLR_____    *CRTDATE_    *BEFORE   05/31/99   *ANY___
   _________     _________    _________    _______   _______    _______
   _________     _________    _________    _______   _______    _______
```

Figure 91. Change Migration Control Group Attributes Display

ANY facilitates the migration of root folders that start with ABCFLR and XYZFLR* for any documents in the folders that you created before 05/31/99. To be eligible for movement, the folders must contain documents that are not storage freed.

3. Issue STRMGRBRM CTLGRP(MGRFNANCE) OPTION(*MIGRATE) command.
4. Issue DSPJOBLOG command to verify that migration completed successfully.
5. Define an archive folder list that is called ARCFLR.

```
                      Add Folder List
                                                          System:   RCHAS400
List name . . . . . ARCFLR__
                                      Sel Subflr
 Seq   Folder                        *INC/ *YES/
                                     *EXC  *NO

 10    ABCFLR/SUBFLR1                *INC  *YES
 20    XYZFLR                        *INC  *YES
 30    XYZFLR/SUBFLR1                *EXC  *YES
```

Figure 92. Add Folder List Display

This is basically a list of documents that you want to archive to tape. In this example, three entries make up the archive list. The first entry (sequence 10)

instructs BRMS to include documents in ABCFLR/SUBFLR1 when processing the list. The second and third entries (sequence 20 and 30) instruct BRMS to include all the documents in XYZFLR but exclude the documents in XYZFLR/SUBFLR1 when processing the archive list.

6. Create an archive control group that is called FLRGRP.

```
                      Create Archive Control Group Entry
                                                            System:   RCHAS400
Group . . . . . . . FLRGRP__
                                  Weekly
        Archive        List       Activity
 Seq    Items          Type       SMTWTFS
 10     ARCFLR         *FLR       * * * *
 __     _______       ______      _______
```

Figure 93. Create Archive Control Group Entry Display

The control group points to the folder list, ARCFLR, and specifies how often the archive control group is run.

7. Change the archive control group attributes for FLRGRP.

```
                           Change Archive Control Group Attributes

  Group. . . . . . . . . . . . . . . : FLRGRP___

  Date type for archival . . . . . :   *CHANGE     *BOTH, *CHANGE, *USE
    Inactivity limit . . . . . . . :   270         0 - 999999 MB, *ARCPCY
```

Figure 94. Change Archive Control Group Attributes Display

This will archive all documents in folders ABCFLR/SUBFLR1 and XYZFLR* excluding XYZFLR/SUBFLR1 that have not been changed in the last 9 months, to tape.

8. Issue STRARCBRM CTLGRP(FLRGRP) OPTION(*ARCHIVE) command.
9. Issue DSPJOBLOG command to verify that the archive completed successfully.
10. Issuing a command to update a document (see “Operations that Call or Do Not Call Dynamic Retrieval” on page 52) causes a dynamic retrieval to occur.

Example 9

Migrating user libraries to a user ASP and then archiving files and members to tape based on inactivity and age.

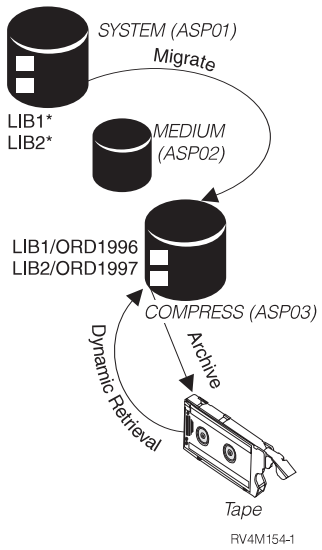

Figure 95. Example 9

This example will demonstrate:

- Migrating user libraries directly to user ASP COMPRESS, when the libraries have not been used or changed for the last 30 days.
- Archiving to tape the files and members that have not been used or changed in 90 days.

1. Define a migration control group that is called MGRORDER.

```
                      Create Migration Control Group Entry
                                                          System:   RCHAS400
Group . . . . . . . . . . . :   MGRORDER

          Migration         Item
 Seq      Items             Type
 10       LIB1*             *LIB
 20       LIB2*             *LIB
```

Figure 96. Create Migration Control Group Entry Display

   This migration control group will migrate libraries that match the names that are specified in the control group.

2. Define the attributes for migration control group, MGRORDER.

```
                    Change Migration Control Group Attributes
                                                          System:   RCHAS400

 Group. . . . . . . . . . . .   MGRORDER___

 Type choices, press Enter.

   Migration     Item Type    Attribute    Test      Value      Objects
   COMPRESS_     *LIB_____    *LASTUSE_    *GE____   30_____    *ANY___
   _________     ______*OR    *LASTCHG_    *GE____   30_____    *ANY___
   _________     _________    _________    _______   _______    _______
```

Figure 97. Change Migration Control Group Attributes Display

This will migrate all libraries that start with LIB1* and LIB2* from the SYSTEM ASP to ASP COMPRESS if any of the objects in the libraries have not been used or changed in the last 30 days.

3. Issue STRMGRBRM CTLGRP(MGRORDER) OPTION(*MIGRATE) command.
4. Issue DSPJOBLOG command to verify that migration completed successfully.
5. Now, define an archive object list that is called ORDER.

```
                      Add Object List
                                                          System:   RCHAS400
List name . . . . . ORDER__
                                     Selection
 Seq   Library         Object   Type   *INC/*EXC

 10    LIB1            ORD1996  *FILE  *INC
 20    LIB2            ORD1997  *FILE  *INC
```

Figure 98. Add Object List Display

This is a list of objects that you want to archive to tape. In this example, the archive list contains two entries. The first line (sequence 10) instructs BRMS to include all the members in file ORD1996 in library LIB1. The second entry (sequence 20) instructs BRMS to include all the members in file ORD1997 in library LIB2.

6. Create an archive control group that is called ARCORDER.

```
                      Create Archive Control Group Entries
                                                          System:   RCHAS400
Group . . . . . . . ARCORDER__
                                    Weekly
        Archive       List          Activity
 Seq     Items        Type          SMTWTFS
 10      ORDER        *OBJ               *
 __     _______      ______      _______
```

Figure 99. Create Archive Control Group Entry Display

The control group points to the object list, ORDER, and has a weekly activity to specify the archive operation to run on Saturdays only.

7. Define the attributes for archive control group, ARCORDER.

```
                          Change Archive Control Group Attributes

Group. . . . . . . . . . . . . . : ARCORDER___

Date type for archival . . . . . :   *BOTH      *BOTH, *CHANGE, *USE
  Inactivity limit . . . . . . . :    90         0 - 999999 MB, *ARCPCY
```

Figure 100. Change Archive Control Group Attributes Display

   This will archive to tape all the members in files ORD1996 and ORD1997 that have not been used or changed in the last 90 days.
8. Issue STRARCBRM CTLGRP(ARCORDER) OPTION(*ARCHIVE) command.
9. Issue DSPJOBLOG command to verify that the archive completed successfully.
10. To dynamically retrieve, see “Operations that Call or Do Not Call Dynamic Retrieval” on page 52.

Example 10

Migrating user libraries back to high performance ASP based on usage.

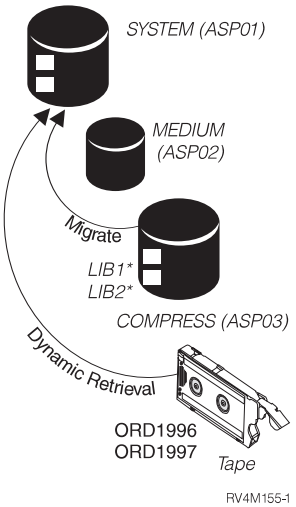

Figure 101. Example 10

In Example 9, we demonstrated migrating libraries that start with LIB1 and LIB2 directly to user ASP COMPRESS. After the libraries were migrated, we archived certain objects such as ORD1996 and ORD1997 to tape. In this example, we are assuming the objects are still archived to tape, and will show how to migrate these libraries that start with LIB1 and LIB2 back to SYSTEM ASP if the libraries have been used and changed in the last 7 days.

1. Change the attributes for the control group, MGRORDER.

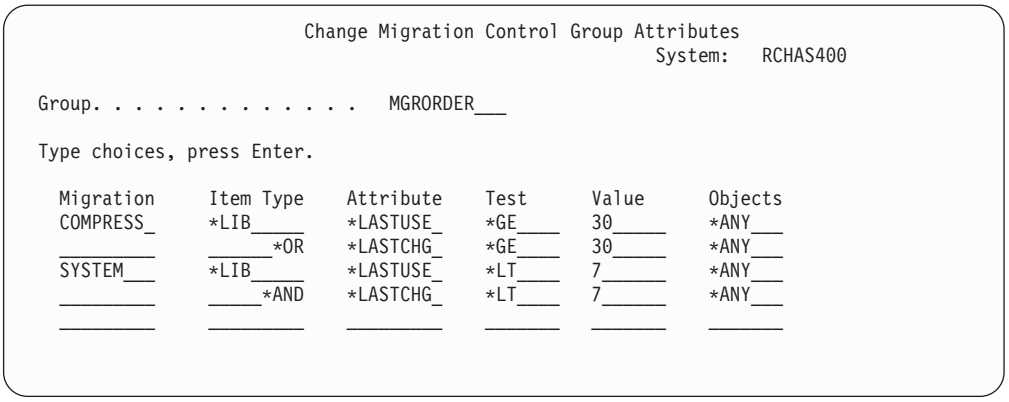

Figure 102. Change Migration Control Group Attributes Display

This will migrate all the libraries from the migration control group, MGRORDER, back to SYSTEM ASP if the libraries have been used and changed in the last week.

2. Issue STRMGRBRM CTLGRP(MGRORDER) OPTION(*MIGRATE) command.
3. Issue DSPJOBLOG command to verify that migration completed successfully.
4. After the library has been migrated back to SYSTEM ASP, making an update to files or members that are archived will cause a dynamic retrieval to occur and automatically or dynamically be restored from tape to the SYSTEM ASP.

Chapter 7. Hierarchical Storage Management Backup Strategy with BRMS

Your ability to ultimately recover any part of your system depends on a solid backup strategy. After you have implemented the HSM strategy that is suitable for your business, it is important to perform a backup.

Below is a list of considerations for your backup strategy:

- When migrating data from one ASP to another, it is important that you save the ASPs that are used during the migration. This applies to both the source ASP and the target ASP.
  - When migrating only libraries, you can use the *ASP special value in your backup control group. The *ASP special value will save only user libraries.
  - When migrating only root folders, you can use the special value *DLO*nn* (where *nn* is the ASP number) in your backup control group. The *DLO*nn* special value will save all folders, documents, and mail in a specified ASP.
- When data has been archived to tape, it may not be necessary to do another backup of this data because the data is already written to tape. Saving only the stubs to tape will reduce your backup time.
- When moving spooled files from one ASP to another (from one output queue to another), it is important that the source output queue and the target output queue are saved. To do this, include the type *SPL in your backup list. Specify the source library and output queue as one entry and the target library and output queue as another entry in the type *SPL backup list.

For more detailed information on how to plan your backup strategy and recovery strategy, see the *Backup Recovery and Media Services for AS/400* book.

© Copyright IBM Corp. 1998, 1999

Chapter 8. Balancing an Auxiliary Storage Pool

The balancing function improves system performance by balancing disk utilization across all of the disk arms in an ASP. You can use the auxiliary storage pool (ASP) balancing function with one or more ASPs. This chapter provides information on the different methods you can use to balance your ASPs for better performance.

The following sections provide more information on how each method works.

Types of Balancing

You can use one of the following methods to balance your ASPs:

- Hierarchical storage management balancing
- Balancing by capacity
- Balancing by usage

Hierarchical Storage Management Balancing

HSM balancing redistributes the high and low use data on each unit in an archive ASP. The system moves the high use data to high performance units and the low use data to low performance units. The ASP must contain the correct combination of high and low performance units to perform this type of balancing. After the balancing process completes, the system clears the trace information. Use hierarchical storage management (HSM) balancing for archive ASPs that contain compressed disk units.

Balancing by Capacity

Balancing by capacity evenly distributes the data on the disk units within an ASP across all of the units. Instead of some units containing the majority of the data, each unit has an equal percentage of used and unused space. This type of balancing is particularly useful when you add disk units to the ASP.

Balancing by Usage

Balancing by usage redistributes use data on each unit within the ASP to balance the arm utilization of each unit within the specified ASP. Balance your ASPs by the usage method when an ASP contains large capacity disk units.

Implementing the Method You Choose

To balance an ASP, take the following steps:

1. If you want to use the usage or HSM balancing method, you must first run the Trace ASP Balance (TRCASPBAL) command. This command starts a trace function that collects statistics on data in the ASP you want to balance. When the trace command completes processing, you can move to the next step. If you are using the capacity method, you do not have to run the trace command and can go directly to the next step.
2. Type STRASPBAL at a command line and press Enter. This takes you to the Start ASP Balance (STRASPBAL) display.

© Copyright IBM Corp. 1998, 1999

Note: You must have *ALLOBJ authority to use this command.

```
                        Start ASP Balance (STRASPBAL)

Type choices, press Enter.

Auxiliary storage pool ID  . . .   *ASP          1-16, *ALL
               + for more values
Balance type . . . . . . . . . .   *HSM          *CAPACITY, *USAGE, *HSM
Time limit . . . . . . . . . . .   *NOMAX        1-9999 minutes, *NOMAX









                                                                      Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

3. Fill in the following fields:
   - In the *Auxiliary storage pool* field, specify the name of the ASP you want to balance. You can type in the number of the ASP you want to balance, or you can specify the *ASP name, as in this example. *ASP instructs the system to balance all of the ASPs known to your system.
   - In the *Balance type* field, specify the balancing method you want to use. This example uses the *HSM method.
   - In the *Time limit* field, indicate the amount of time each balancing job can take. This example uses the *NOMAX default value, which instructs BRMS to take as much time as it needs to balance the ASPs. You can also specify time in minutes in this field.
4. Press Enter to balance your ASPs.

Notes:

1. If using the HSM method, use the STRASPBAL command for archived ASPs, rather than highly utilized ASPs.
2. If needed, you can use the End ASP Balance (ENDASPBAL) command to end your balancing operation before the system reaches the time limit you specified.

Chapter 9. Problem Analysis

This chapter describes what to look for if problems occur and how to recover from the problems.

There are different ways to find out if a problem occurs when running an operation. If a command fails, do one of the following:

1. Display the job log information by using the DSPJOBLOG command on the system. Use F10 to display the detailed message. The messages include secondary text that describes the recommended course of action.
2. If you are using BRMS, use the DSPLOGBRM command. The display log for BRMS displays the log that BRMS has created as a result of processing a BRMS operation.

Recovery from an Abnormal End During Migration

When an abnormal end occurs during the migration of a library, do the following:

During the migration of libraries, the library that is being moved is renamed to a library that is called QHSMLIB*nnn*. When you use this naming convention, *nnn* is an incremented number. A *USRSPC object is also created, which is named QHSMLIB*nnn*. This object is put into library QUSRSYS. It contains the name of the library that you are moving in the object description text.

1. To determine if the renamed library (QHSMLIB*nnn*) exists, look for the *USRSPC in QUSRSYS:
   - WRKOBJ OBJ(QUSRSYS/QHSMLIB*) OBJTYPE(*USRSPC)

     There will be a *USRSPC object, QHSMLIB*nnn* in QUSRSYS with the same name as the library being moved in the object description text.
2. If the renamed library QHSMLIB*nnn* does not exist, determine if the library was migrated to the target ASP.
   - DSPOBJD OBJ(MYLIB) OBJTYPE(*LIB) and specify option 5 (Display full attributes)
     - If MYLIB library is not in the target ASP, then try the migration again
     - If MYLIB library is in the target ASP, the migration was successful. No further action is necessary.
3. If the renamed library (QHSMLIB*nnn*) does exist, use the DSPLIB command to compare the number of objects in the QHSMLIB*nnn* to the library that you are moving. The library now resides in the target ASP.
   - If the number of objects in QHSMLIB*nnn* is less than the number of objects in the library that you are moving, then do the following:
     - DLTUSRSPC USRSPC(QUSRSYS/QHSMLIB*nnn*)
     - DLTLIB LIB(QHSMLIB*nnn*). When you complete the migration, no further action is necessary.
   - If the number of objects in QHSMLIB*nnn* is greater than or equal to the number of objects in the library that you are moving, then do the following:
     - DLTLIB LIB(MYLIB)
     - RNMOBJ OBJ(QHSMLIB*nnn*) OBJTYPE(*LIB) NEWOBJ(MYLIB)
     - DLTUSRSPC USRSPC(QUSRSYS/QHSMLIB*nnn*)
     - Retry the migration
4. If you encounter problems, contact your support representative.

© Copyright IBM Corp. 1998, 1999

When an abnormal end occurs during the migration of a folder, you may need to do one of the following:

- Recover from a previous backup of the root folder.
- Recover the root folder from a save file that remains on the system.

See Info APAR II11238 for additional details.

Appendix A. Summary of Hierarchical Storage Management (HSM) Reports

Following are brief summaries of the HSM-related reports that you can generate through BRMS. The summaries also include printer source file information for each report. You must have BRMS installed to access these reports.

Archive DLO Candidate

The Archive DLO Candidate report lists all documents that meet the selection criteria as a result of processing an archive control group. You can generate this report by using the Start Archive using BRM (STRARCBRM) command. You should run the STRARCBRM with the *Run Option* value at *REPORT **prior** to running an actual archive job. This allows you to verify that the system is archiving the correct documents.

The following should be noted about this report:

- It is important that you specify a value in either the *Inactivity date* or the *Number of days used* field when you select which documents you want to archive. BRMS bases its selection of items for this report on the date or number of days that you specify in one of these fields, and on other archive specifications (unless you specifically excluded them from the archive).
- You receive an Archive Folder Summary report along with the Archive DLO Candidate report. This report shows all of the lists, folders, and subfolders that the archive control group contains. The Archive Folder Summary report lists the total number of documents as well as the amount of disk unit space that these documents occupy. The report expresses the disk unit space in bytes.

The printer source file for this report is QP1A8ARF.

Note: The DLO Archive report is another name for this report.

Archive Folder List

The Archive Folder List report lists all the folders in an archive list. You can generate this report with the Work with List using BRM command. Be sure to set the value in the *Type of usage* field to *ARC (WRKLBRM TYPE(*ARC)).

The following should be noted about this report:

- The report indicates whether to include or exclude the folders in the list in the archive job.
- The report indicates whether to consider all of the subfolders within the folder in the archive job.

The printer source file for this report is QP1AAF.

Archive Integrated File System Object Candidate

The Archive Integrated File System Object Candidate Report lists the integrated file system objects that you processed inside an archive control group. You generate this report by using the STRARCBRM command with a *Run option* of *REPORT. Always run this report **prior** to performing an actual archive. This allows you to verify that the system archives the correct objects.

© Copyright IBM Corp. 1998, 1999

The printer source file for this report is QP1A1ARC.

Archive Object Candidate

The Archive Object Candidate Report lists those objects that are selected for archive as a result of processing an archive control group. Run the Archive Object Candidate Report by using the STRARCBRM command with a *Run Option* of *REPORT. Always run the Archive Object Candidate Report prior to performing an actual archive to verify that the system is archiving the correct objects.

The following should be noted about this report:

- The *Inactivity date* field is a critical field in selecting which objects to archive. BRMS bases its selection of items for this report on this date and other archive specification (unless you specifically excluded them from the archive criteria).
- The Archive Objects Candidate Report lists the total number of objects as well as the amount of disk unit space that these objects occupy. The report expresses the disk unit space in bytes.

The printer source file for this report is Q1AARC.

Archive Object List

The Archive Object List report lists all the objects in an archive list. You can generate this report by using the Work with Lists using BRM command with the *Type of usage* field set to *ARC (WRKLBRM TYPE(*ARC)). The report indicates whether to include or exclude the objects in this list from the archive activity.

The printer source file for this report is QP1AAO.

Archive Spooled File Candidate

You can generate the Archive Spooled File Candidate report by setting the *Run option* value on the STRARCBRM command to*REPORT. This report lists all of the spooled files that BRMS archived as part of the control group that you specified in the command. This report lists the inactivity date for the archive, the size of each spooled file, and the total size of all the spooled files.

The printer source file for this report QP1AOQ.

Archive Spooled File List

Use the Work with Lists using BRM (WRKLBRM) command to generate the Archive Spooled File List report. Set the value in the *Type of usage* field to *ARC. This report lists all of the spooled files that exist in a list. The report identifies each list item, its sequence number, and the selection criteria.

The printer source file for this report is QP1AAQ.

Display ASP Information

The Display ASP Information report provides statistical information on all of the auxiliary storage pools (ASPs) that exist on your system.

The printer source file for this report is QP1ADV.

Migration Item Candidate

You can automatically generate this report by using the Start Migration using BRM (STRMGRBRM) command with the *Run option* value set to *REPORT (the default). This report lists the migration candidates.

The printer source for this report is QP1ACH.

Move Spooled Files using BRM

You can generate this report by using the Move Spooled Files using BRM (MOVSPLFBRM) command. This report lists the spooled file candidates for movement.

The printer source for this report is QP1AHSF.

Object Archive

The Object Archive report lists the objects that BRMS archives as part of an archive control group. Use the STRARCBRM command with a *Run Option* of *REPORT to produce this report.

The following should be noted about this report:

- The *Inactivity date* field is a critical field in selecting objects to archive. In this example, the run date is 5/19/95, and the inactivity date is 4/19/95. The value for the *Inactivity limit* field for this control group was 30 days, which is the difference between 5/19/95 and 4/19/95. BRMS includes any object in the PROGLIB and GLLIB libraries that have a *Change date* or *Last used date* prior to 4/19/95 in this archive job.
- The Object Archive report lists the total number of objects as well as the amount of disk unit space that these objects occupy. The report expresses the disk unit space in bytes.

The printer source file for this report is QP1AARC.

Retrieve Request

You can generate the Retrieve Request report by using the RSMRTVBRM command. To do so, set the value in the *Action* field to *REPORT. This report provides a list of all the libraries and special values that BRMS selected during retrieve processing. BRMS sequences this report by ASP number.

The printer source file for this report is QP1A1SX.

Spooled File Archive

You can generate the Spooled File Archive report by running the STRARCBRM command with the *Run option* value set to *REPORT. This report lists all spooled files that BRMS archived as part of the control group that you specified in the command. The report lists the inactivity date, the size of each spooled file, and the total size of the combined spooled files.

The printer source file for this report is QP1AOQ.

Appendix B. Customizing the User Interface for Dynamic Retrieval

A user exit program, if registered to an BRMS exit point, allows the system administrator to customize the messages that are sent to users when dynamic retrieval occurs. The current default message indicates only that the job is in progress. With the BRMS Object Retrieval Exit program, you can customize your message displays to give users more specific and useful information about the retrieval job they want to perform. For example, you might want to create a message display that provides users with location and retrieval time information about the object that they want to retrieve. Or you might create messages for problem situations. For example, if BRMS cannot access the file the user wants to retrieve, you can create an informational display that suggests possible reasons for the failure. In this appendix, you can find information on how to set up and register your new user exit program.

About the BRMS Object Retrieval Exit Program

The BRMS Object Exit Retrieval Program provides the capability to retrieve objects that were saved in BRMS with storage freed. BRMS calls the user exit program each time the IBM Operating System/400 Version 4 (OS/400) attempts to access an archived object. At that time, BRMS determines if it should restore the archived object to satisfy the OS/400 request. This exit gives the user exit program a chance to influence the decision to restore the object. It also provides a custom interface that you can tailor to meet specific needs.

When you put the user exit program in charge of the confirmation message, it decides the most appropriate action to take for the archived object at that time. The exit program then returns the information to BRMS, indicating that BRMS should take one of the following options:

- Proceed as normal.
- Retrieve immediately.
- Submit the job to batch.
- Delay the restore.
- Cancel the job.

To use this exit program, you must confirm that the default value in the *Retrieve confirmation* field on the Retrieval Policy is *VERIFY. This instructs BRMS to send a message for each object that you want BRMS to retrieve.

Setting Up the User Exit Program

You can create a user exit program to be as simple or as complex as required for your operation. In a simple program, you can display the fields you create on a screen that matches the interface of OS/400 or any common application. This means that the end-user sees an interface that does not look like an error condition, which is how the default appears.

A more complex program can display different screens to different users. Some complex programs can even make decisions for users under program control, which may or may not include the use of displays. For example, suppose that you want to retrieve an object on a tape that is stored in a 3494 tape library, and that the file

© Copyright IBM Corp. 1998, 1999

member size is below 10MB. In this case, you could return the G (Continue the operation) to BRMS, which simulates user activity for this task. Instead, however, the program is actually making an informed decision for the user.

Another example may be to present a display only for more knowledgeable users. You could define this in a file that is maintained and reviewed by an administrator.

A Sample User Exit Program

The following CL program provides an example of how you can code an exit program. In this example, the program issues messages regarding object retrieval. You can write similar programs to issue messages regarding streamfiles and documents.

The CL program created in the following example is a simple program that contains descriptive statements. We explain the actual processing in just a few lines of CL. You can make the program more complex by adding logic appropriate to your requirements.

Note: This example does not include any error coding. In a production environment, you need to add error logic.

If you are familiar with CL, most of this program will be self-explanatory. The only part that might be a little difficult to understand is the file member size. The program stores this in two fields, the size, and a multiplier. If the object size is greater than 1GB, you must multiply the size value by 1024 to get the actual size (although this is rounded).

```
PGM            PARM(&OBJDINF &MEDINF  &RETCD)
/*****************************************************************************/
/*                                                                           */
/*     Dynamic Retrieval Processing Program for confirmation messages        */
/*     Program Name: DRTVEXIT                                                */
/*                                                                           */
/*****************************************************************************/
/*  Parameters                                                               */
/*****************************************************************************/
DCL VAR(&OBJDINF) TYPE(*CHAR) LEN(145) /* Object description info            */
DCL VAR(&MEDINF) TYPE(CHAR) LEN(9999)  /* Media information                  */
DCL VAR(&RETCD) TYPE(*CHAR) LEN(1)     /* Return code                        */
                                                                             */
*****************************************************************************/
/*  Breakdown of Object Description Information                              */
/*****************************************************************************/
DCL &OBJDBIN     *CHAR   4          /* Length (bin) of object description     */
DCL &MEDBIN      *CHAR   4          /* Length (bin) of media information      */
DCL &CTRLBIN     *CHAR   4          /* Length (bin) of return code            */
DCL &OBJDLEN     *DEC  (15 0)       /* Length (bin) of object description     */
DCL &MEDLEN      *DEC  (15 0)       /* Length (bin) of media information      */
DCL &CTRLLEN     *DEC  (15 0)       /* Length (bin) of return code            */
DCL &OBJNAME     *CHAR  (10)        /* Object name                            */
DCL &OBJLIB      *CHAR  (10)        /* Object library                         */
DCL &MBRNAME     *CHAR  (10)        /* Member name                            */
DCL &OBJOWNER    *CHAR  (10)        /* Object owner                           */
DCL &SAVDATE     *CHAR   7          /* Save date                              */
DCL &SAVTIME     *CHAR   6          /* Save time                              */
DCL &SAVRLS      *CHAR   6          /* Save release                           */
DCL &MEDCLS      *CHAR  10          /* BRMS media class                       */
DCL &FILESEQ     *CHAR   4          /* Tape file sequence number - decimal    */
DCL &MBRSIZE     *CHAR   4          /* Member size, Bytes - binary            */
DCL &MBRMULT     *CHAR   4          /* Member size multiplier - binary        */
DCL &MBRMULTDEC  *DEC (15 0)        /* Member size multiplier - decimal       */
DCL &ASPN        *CHAR   2          /* ASP number                             */
DCL MBRTEXT      *CHAR  50          /* Member text                            */
*****************************************************************************
/*  Breakdown of Volume Information                                          */
*****************************************************************************
DCL &BINVOLS     *CHAR   4          /* Number (bin) of volumes in list        */
DCL &NUMVOLS     *DEC  (6 0)        /* Number (dec) of volumes in list        */
DCL &VOLID       *CHAR   6          /* Volume ID in list                      */
DCL &VOLLOC      *CHAR   10         /* Volume location in list                */
DCL &VOLRET      *DEC  (6 0)        /* Volume retrieval time (HHHHMM) in list */
DCL &VOLRETC     *CHAR   6          /* Volume retrieval time - character      */
*****************************************************************************
```

Figure 103. (Part 1 of 4). Simple exit program for confirmation processing

```
*****************************************************************************/
/* Details of Return Code Information                                        */
*****************************************************************************/
                                          /*  Return code 0-4                 */
                                                        /* 0 *VERIFY         */
                                                        /* 1 *NOTIFY         */
                                                        /* 2 *DELAY          */
                                                        /* 3 *SBMJOB         */
                                                        /* 4 *CANCEL         */
                                                                             */
*****************************************************************************
 /* User Prompt Screen Definition                                            */
*****************************************************************************
   DCLF           FILE(SCREEN)               /*  User prompt screen          */
*****************************************************************************
/*  Other Working Variables                                                  */
*****************************************************************************
 DCL &MEDOFFSET  *DEC  (10 0)   /* Space offset counter in volume list       */
 DCL &COUNTER    *DEC  ( 6 0)   /* Volume counter in volume list             */
 DCL &MBRSIZEN   *DEC  (15 0)   /* Member size as a number                   */
                                                                             */
*****************************************************************************
/*  Process Object Description Parameter                                     */
*****************************************************************************
CHGVAR  &OBJDBIN  %SST(&OBJDINF 1 4)        /* Change length of object       */
                                            /* Description from a char to    */
                                            /* A numeric (bin) field         */
CHGVAR  &MEDBIN   %SST(&OBJDINF 5 4)        /* Change length of media        */
                                            /* Information from a char to    */
                                            /* A numeric (bin) field         */
CHGVAR  &CTRLBIN  %SST(&OBJDINF 9 4)        /* Change length of return code  */
                                            /* From a char to numeric field  */
CHGVAR  &OBJDLEN  %BIN(&OBJDBIN)            /* Put into a decimal field      */
                                            /* Object description is 145     */
CHGVAR  &MEDLEN   %BIN(&MEDBIN)             /* Put into a decimal field      */
                                            /* Media info is a variable length */
CHGVAR  &CTRLLEN  %BIN(&CTRLBIN)            /* Put into a decimal field      */
                                            /* Return code is always 1       */
*****************************************************************************
```

Figure 104. (Part 2 of 4) Simple exit program for confirmation processing

```
*******************************************************************************
CHGVAR  &OBJNAME  %SST(&OBJDINF 13 10)      /* Object name                     */
CHGVAR  &OBJLIB   %SST(&OBJDINF 23 10)      /* Object library                  */
CHGVAR  &MBRNAME  %SST(&OBJDINF 33 10)      /* Member name                     */
CHGVAR  &OBJOWNER %SST(&OBJDINF 43 10)      /* Object owner                    */
CHGVAR  &SAVDATE  %SST(&OBJDINF 53  7)      /* Save date CYYMMDD               */
CHGVAR  &SAVTIME  %SST(&OBJDINF 60  6)      /* Save time HHMMSS                */
CHGVAR  &SAVRLS   %SST(&OBJDINF 66  6)      /* Saved release                   */
CHGVAR  &MEDCLS   %SST(&OBJDINF 72 10)      /* BRMS media class                */
CHGVAR  &FILESEQ  %SST(&OBJDINF 82  4)      /* Tape file sequence number       */
CHGVAR  &MBRSIZE  %SST(&OBJDINF 86  4)      /* Member size                     */
CHGVAR  &MBRMULT  %SST(&OBJDINF 90  4)      /* Member size multiplier          */
CHGVAR  &ASPN     %SST(&OBJDINF 94  2)      /* ASP number                      */
CHGVAR  &MBRTEXT  %SST(&OBJDINF 96 50)      /* Member text                     */
                                                                               */
CHGVAR  &MBRSIZEN %BIN(&MBRSIZE)            /* Convert bin to number           */
CHGVAR  &MBRMULTDEC %BIN(&MBRMULT)          /* Convert bin to number           */
IF  (&MBRMULTDEC *EQ 1024) CHGVAR &MBRSIZEN (&MBRSIZEN * &MBRMULTDEC)          */
                                            /* Multiply by 1024 if large       */
*******************************************************************************
CHGVAR &BINVOLS %SST(&MEDDINF 1 4)          /* Change number of vols to bin    */
CHGVAR &NUMVOLS %BIN(&BINVOLS)              /* Change number of vols to dec    */
CHGVAR &MEDOFFSET 5                         /* Position at first volume        */
CHGVAR &COUNTER &NUMVOLS                    /* Position at first volume        */


                  /*************************************/
                  /* This example is concerned only with  */
                  /* the first volume, and so passes      */
                  /* only once through this code. If you  */
                  /* want data for multiple volumes, you  */
                  /* can use the &COUNTER variable */
                  /* to process the data in a loop.       */
                  /*************************************/

CHGVAR &VOLID   %SST(&MEDINF &MEDOFFSET 6)                                     */
CHGVAR &MEDOFFSET (&MEDOFFSET + 6)          /* Increment ptr beyond Vol ID     */
CHGVAR &VOLLOC  %SST(&MEDINF &MEDOFFSET 10)                                    /*
CHGVAR &MEDOFFSET (&MEDOFFSET + 10)         /* Increment ptr beyond location   /*
CHGVAR &VOLRET  %SST(&MEDINF &MEDOFFSET 6)                                     /*
IF     COND(&VOLRET *EQ '    ') +                                              /*
       THEN(CHGVAR &VOLRET 000000)                                             /*
       ELSE(CHGVAR &VOLRET &VOLRETC)                                           /*
CHGVAR &MEDOFFSET (&MEDOFFSET + 6)          /* Increment ptr beyond rtv tm     /*
*******************************************************************************
```

Figure 105. (Part 3 of 4) Simple exit program for confirmation processing

```
*******************************************************************************
/*               Main Application Logic - Simple Example                       /*
*******************************************************************************
                                                                              /*
CHGVAR  VAR(RETCD) VALUE('1')                /* Set drf to retrieve            /*
IF      COND(&VOLID *EQ '*SAVF ') THEN(DO)   /* If SAVF change
        CHGVAR   VAR(*VOLLOC) VALUE('*SAVF') /* Variables for                  /*
        CHGVAR   VAR(*VOLRET) VALUE('000000') /* the user                      /*
ENDDO                                                                         /*
                                                                              /*
SNDRCVF                                       /* Display the screen           /*
                                                                              /*
IF     COND(*IN03 *EQ '1') THEN(CHGVAR +                                      /*
       VAR(*RETCD) VALUE('4'))                /* If F3, *CANCEL               /*
                                                                              /*
ENDPGM                                                                        /*
*******************************************************************************/
```

Figure 106. (part 4 of 4) Simple exit program for confirmation processing.

Display File Image

In the sample program, you display a screen to the user that explains their choices. Following is an example of the kind of display you could create for your program:

```
                              Accessed File Not Available

             The file you accessed is not currently stored on disk.

             You can initiate a retrieval of the file by entering a
             response below. The following information can assist you
             in deciding a course of action:

File:. . . . . .PAYMAST
Library:. . . . . PAYLIB1
Member:. . . . . .DEC99HIST
Text:. . . . . .Payroll December 1999 Master File Updates
Member size: . . . . . .153,694
ASP for restore: . . . . . .01
Tape Volume ID: . . . . . T02190
Volume Location: . . . . . .OFFSITEVLT
Location Retrieval Time: . .02:30

Options:   1  Retrieve now and wait
           2  Administrator to retrieve later - retry application later
           3  Retrieve in batch - retry application later

Selection: 1
F3=Cancel retrieval
```

Figure 107. Sample File Image

The fields used in this example include:

- OBJNAME
- OBJLIB
- MBRNAME
- MBRTEXT
- MBRSIZEN
- ASPN
- VOLID
- VOLLOC
- VOLRET

Fields not shown, but also available for use include:

- OBJOWNER
- SAVDATE
- SAVTIME
- SAVRLS
- MEDCLS
- FILESEQ

Note: In the sample program, we did not qualify the display file and did not give a library name. To ensure the successful operation of your exit program, you should put the display file in the common library in the library list (such as

QGPL). Or, alternatively, you could change the program to specify the name of the library in which to find the file. You can do this on the DCLF statement.

Display File Source

Following is an example of a display file source (DDS) you can use for this screen format:

```
*********************************Beginning of Data*********************************
...+... 1...+... 2...+... 3...+... 4...+... 5...+... 6...+... 7                  /*
   A*%%TS  SD  19991118  183421  SAMPLE    REL-V4R4M0   5763-PW1                 /*
   A*%%EC                                                                       /*
   A                                      DSPSIZ(24 80 *DS3)                    /*
   A         R BRMSEXIT                                                         /*
   A*%%TS  SD  19991118  183421  SAMPLE    REL-V4R4M0   5763-PW1                 /*
   A                                     CF03(03 'Function Key 3')              /*
   A                                   1 26'Accessed File Not Available'        /*
   A                                      DSPATR(HI)                            /*
   A                                   3 12'The file you have accessed is not -  /*
   A                                       currently stored on disk.'           /*
   A                                   5 12'You can initiate a retrieval of the  /*
   A                                       file by entering a'                  /*
   A                                   6 12'response below. The following inform- /*
   A                                       ation may help you to'               /*
   A                                   7 12'decide your course of action:'      /*
   A                                   9  2'File: . . . . . .'                  /*
   A                                  10  2'Library: . . . . . '                /*
   A                                  11  2'Member: . . . . . .'                /*
   A                                  12  2'Text: . . . . . .'                  /*
   A                                  13  2'Member size (bytes): . . . . '       /*
   A                                  14  2'ASP for restore: . . . . . '        /*
   A                                  15  2'Tape volume ID: . . . . . .'        /*
   A                                  16  2'Volume location: . . . . . '        /*
   A                                  17  2'Location retrieval time: . . '      /*
```

Figure 108. DDS source for customized confirmation panel (Part 1 of 2)

```
   A          OBJNAME         10A  0   9  30DSPATR(HI)                          /*
   A          OBJLIB          10A  0  10  30DSPATR(HI)                          /*
   A          MBRNAME         10A  0  11  30DSPATR(HI)                          /*
   A          MBRTEXT         50A  0  12  30DSPATR(HI)                          /*
   A          MBRSIZEN        15Y 00  13  30DSPATR(HI)                          /*
   A                                          EDTCDE(1)        /*
   A          ASPN             2A  0  14  30DSPATR(HI)         /*
   A          VOLID            6A  0  15  30DSPATR(HI)                          /*
   A          VOLLOC          10A  0  16  30DSPATR(HI)                          /*
   A          VOLRET           6Y 00  17  30DSPATR(HI)                          /*
   A                                          EDTWRD('  0:  ') /*
   A                                  19  3'Options:'                          /*
   A                                  19 12'1'                                  /*
   A                                       DSPATR(HI)                           /*
   A                                  19 14'retrieve now and wait'              /*
   A                                  20 12'2'                                  /*
   A                                  20  DSPATR(HI)                           /*
   A                                  20 14'Administrator to retrieve later --  /*
   A                                     retry application later'               /*
   A                                  21 12'3'                                  /*
   A                                       DSPATR(HI)                           /*
   A                                  21 14'Retrieve in batch - retry           /*
   A                                     application later'    /*
   A                                  23  6'Selection:'                         /*
   A          RETCD             1A B  23 17VALUES('1' '2' '3')                  /*
   A                                       DSPATR(HI)                          /*
   A                                  24 10'F3=Cancel Retrieval'               /*
   A                                       COLOR(BLU)          /*
*******************************************************************************
```

Figure 108. DDS source for customized confirmation panel (Part 2 of 2)

Setting the Retrieve Policy

After you create the program exit and displays you want BRMS to use, confirm that the *Retrieve confirmation* field on the Retrieval Policy is *VERIFY. This instructs BRMS to send the message you created for each object that you want BRMS to retrieve. To get to the BRMS retrieve policy, take the following steps:

1. Type GO BRMPCY at a command line and press Enter. This takes you to the Policy Administration menu.
2. Select option 5 (Retrieve policy) from the Policy Administration menu.

```
                        Change Retrieve Policy                    RCHAS400

Type choices, press Enter.

Retrieve device  . . . . . .  *MEDCLS     Name, F4 for list



Retrieve confirmation:
  Interactive operation  . . . . . .  *VERIFY     *VERIFY, *NOTIFY, *DELAY..
  Batch operation  . . . . . .  *NOTIFY     *NOTIFY, *DELAY, *VERIFY..
Retrieve authorization . . . . . .  *OBJEXIST   *OBJEXIST, *READ, *UPD...
End of tape option . . . . . .  *REWIND     *REWIND, *LEAVE, *UNLOAD
Option . . . . . .  *ALL        *ALL, *NEW, *OLD, *FREE
Allow object differences . . . . . .  *NONE       *NONE, *ALL
ASP high storage threshold . . . . . .  *SYS        *ASP, *SYS, 1-95
Retrieved object retention . . . . . .  0           0-9999, *ARC, *NOMAX
Extend retention on usage  . . . . . .  0           0-9999
Reset days used count. . . . . .  *NO         *NO, *YES

                                                                    More...
F3=Exit    F4=Prompt    F5=Refresh    F9=System policy
F12=Cancel
```

Figure 109. Setting the retrieve policy

3. There are two elements that are associated with the *Retrieve confirmation* parameter.
   - The default value for the *Interactive operation* field is *VERIFY. When you want BRMS to send the confirmation message you created before processing the retrieve, you need to ensure that this default value is in place.
   - The default value for the *Batch operation* field is *NOTIFY. When you want BRMS to send a confirmation message before processing a batch retrieval operation, you need to change this value to *VERIFY.
4. Review and change the other parameters as needed.
5. Press F3 to exit and save your changes.

How to Register Your New Program

Before you can use your new exit program with BRMS, you first need to register it with the OS/400. To do so, take the following steps:

1. Type WRKREGINF at a command line and press Enter. This takes you to the Work with Registration Information display.

```
                    Work with Registration Information

Type options, press Enter.
  5=Display exit point   8=Work with exit programs

                          Exit
     Exit                 Point
Opt  Point                Format    Registered  Text
 8   QIBM_A1A_RETR_INF    RTVI0100     *YES     BRM Services/400 object retri
     QIBM_A1A_TAPE_INF    MEDI0100     *YES     BRM Services/400 media inform
     QIBM_A1A_TAPE_MOVE   MEDM0100     *YES     BRM Services/400 media moveme
     QIBM_QGW_NJEOUTBOUND NJEO0100     *YES     Network Job Entry outbound ex
     QIBM_QHQ_DTAQ        DTAQ0100     *YES     Original Data Queue Server
     QIBM_QJO_DLT_JRNRCV  DRCV0100     *YES     Delete Journal Receiver
     QIBM_QLZP_LICENSE    LICM0100     *YES     Original License Mgmt Server
     QIBM_QMF_MESSAGE     MESS0100     *YES     Original Message Server
     QIBM_QNPS_ENTRY      ENTR0100     *YES     Network Print Server - entry
     QIBM_QNS_CRADDACT    ADDA0100     *YES     Add CRQ description activity
     QIBM_QNS_CRCHGACT    CHGA0100     *YES     Change CRQ description activity
                                                                    More...
Command ===>
F3=Exit   F4=Prompt   F9=Retrieve   F12=Cancel
```

Figure 110. The Work with Registration Information display

2. Type an 8 (Work with exit programs) in front of the exit point name you want to work with (in this case, QIBM_A1A_RETR_INF), and press Enter.

```
                            Work with Exit Programs

Exit point:   QIBM_A1A_RETR_INF        Format:   RTVI0100

Type options, press Enter.
  1=Add   4=Remove   5=Display   10=Replace

              Exit
            Program     Exit
Opt          Number     Program        Library
 1                       DRTVEXIT       QGPL

                                                                       Bottom
Command
===>
F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve   F12=Cancel
```

Figure 111. The Work with Exit Programs display

3. At the Work with Exit Programs display, type a 1 (Add) in the *Opt* field, and the name of your new exit program in the *Exit program* field. Then type the name of the library in which you want to house the program in the *Library* field.
4. Press Enter. This takes you to the Add Exit Program (ADDEXITPGM) display.

```
                        Add Exit Program (ADDEXITPGM)

Type choices, press Enter.

Exit point . . . . . > QIBM_A1A_RETR_INF
Exit point format  . . . . . > RTVI0100      Name
Program number . . . . . > 1             1-2147483647, *LOW, *HIGH
Program  . . . . . > DRTVEXIT      Name
  Library  . . . . . >  QGPL         Name, *CURLIB
Threadsafe . . . . . .   *UNKNOWN      *UNKNOWN, *NO, *YES
Multithreaded job action . . . .   *SYSVAL       *SYSVAL, *RUN, *MSG, *NORUN
Text 'description' . . . . . .   Our confirmation exit program for BRMS.









                                                                      Bottom
F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display        F24=More keys
```

Figure 112. The Add Exit Program display

5. At this display, confirm that the *Exit point* and *Program* name fields are correct. You should also add a brief description of the program in the *Text description* field.
6. Review and change the remaining parameters as needed.
7. Press Enter to process the registration.

Bibliography

The following is a list of related information that may help you as you use this book.

The books below are listed with their full title and base order number.

- *Automated Tape Library Planning and Management*, SC41-5309-02 , provides information about tasks that can be performed with an automated tape library (ATL). It describes recommended methods for designing and using automated tape libraries. It compares automated tape library devices that are currently available.
- *Backup and Recovery*, SC41-5304-03 , contains information about planning a backup and recovery strategy. Other topics include different types of media available to save and restore system data, save and restore procedures, and disk recovery procedures. It also describes how to install the system again from backup, how to plan for and set up user auxiliary storage pools (ASPs), mirrored protection, and checksums, along with other availability recovery topics. It also provides information about journaling, save-while-active, and DASD compression.
- *Backup Recovery and Media Services for AS/400*, SC41-5345-01 , provides information on how to develop and implement a backup and recovery strategy using the Backup, Recovery, and Media Services for AS/400 product.
- *CL Reference (Abridged)*, SC41-5722-03 , provides the application programmer or programmer with a description of the AS/400 control language (CL) and its commands.
- *System API Reference*, SC41-5801-03 , provides information on how to create, use, and delete objects that help manage system performance, use spooling efficiently, and maintain database files efficiently. This manual also includes information on creating and maintaining the programs for system objects and retrieving OS/400 information by working with objects, database files, jobs, and spooling.

© Copyright IBM Corp. 1998, 1999

Index

Special Characters

A

B

© Copyright IBM Corp. 1998, 1999

C

D

E

F

H

I

J

L

M

N

O

P

R

S

T

U

W

Readers' Comments — We'd Like to Hear from You

AS/400e
Hierarchical Storage Management

Publication No. SC41-5351-01

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? ☐ Yes ☐ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name _____ Address _____

Company or Organization _____

Phone No. _____

Readers' Comments — We'd Like to Hear from You
SC41-5351-01

Cut or Fold
Along Line

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

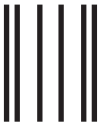

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

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

POSTAGE WILL BE PAID BY ADDRESSEE

IBM CORPORATION
ATTN DEPT 542 IDCLERK
3605 HWY 52 N
ROCHESTER MN 55901-7829

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

Cut or Fold
Along Line

Part Number: 43L1567

Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

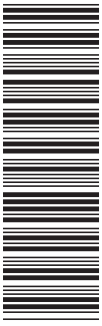

43L1567

SC41-5351-01

Spine information:

AS/400e

OS/400 Hierarchical Storage Management V4R4