Windows NT Backup and Recovery with ADSM

Bruno Friess, Marie-Christine Gamet, Barry Kadleck
Patrick Randall

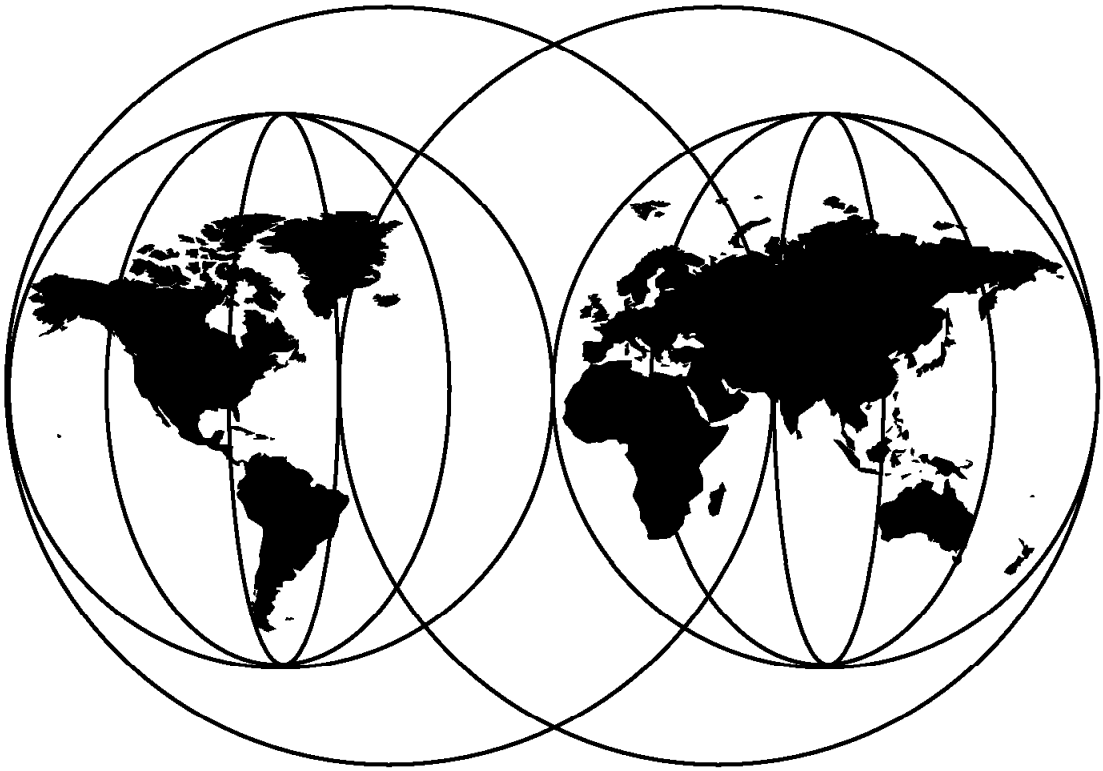

International Technical Support Organization

http://www.redbooks.ibm.com

SG24-2231-00

International Technical Support Organization

SG24-2231-00

Windows NT Backup and Recovery with ADSM

May 1998

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix B, “Special Notices” on page 165.

First Edition (May 1998)

This edition applies to

- Version 3, Release 1 of ADSTAR Distributed Storage Manager for MVS, Program Number 5655-A30, for use with the MVS and OS/390 operating system
- Version 3, Release 1 of ADSTAR Distributed Storage Manager for AIX, Program Number 5765-C43, for use with the AIX operating system, Version 4
- Version 3, Release 1 of ADSTAR Distributed Storage Manager for HP-UX, Program Number 5639-D92
- Version 3, Release 1 of ADSTAR Distributed Storage Manager for Sun Solaris, Program Number 5639-D91
- Version 3, Release 1 of ADSTAR Distributed Storage Manager for Windows/NT, Program Number 5639-C59, for use with the Windows/NT operating system

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Department QXXE Building 80-E2
650 Harry Road
San Jose, California 95120-6099

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

Contents

Preface

This book explores the requirement for backup solutions to protect the Microsoft NT Server and some of its common applications against failures from minor to disaster level in severity. Common methods of backup are outlined, and the advantages and disadvantages of the use of IBM′s ADSTAR Distributed Storage Manager (ADSM) to protect against the failure types are covered in some detail.

Where the use of ADSM is appropriate, examples are given of how to set up the solution, and those solutions are tested and the results shown. The book is intended for anyone who has a need to back up important data on NT server systems where the backup data must be held safely and separate from the working copies and with absolute integrity.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization San Jose Center.

Patrick Randall is a Distributed Storage Software Specialist at the International Technical Support Organization, San Jose Center. He has written seven redbooks on ADSM, teaches IBM classes worldwide on all areas of distributed storage, and is a consultant in disaster and business recovery. Before joining the ITSO in July 1996, Pat worked in IBM UK′s Business Recovery Services as a Solutions Architect.

Bruno Friess is a member of the ADSM Level 2 support in EMEA. He has six years of experience in system administration and programming in Windows NT and related Microsoft software. Bruno holds a degree in computer science with a focal point on human-computer interaction and natural language research. His areas of expertise include operating systems and databases.

Marie-Christine Gamet is an independent Security and Backup/Recovery Consultant based in Paris, France. She has seven years of experience in the information security field. Marie-Christine has worked with IBM for three years as a Business Partner. She is president of the IPSO (VM/ESA and ADSM User Group) in France. She is a qualified CERSSI (European Certificate of Security and Quality for Enterprise Information Systems).

Barry Kadleck is a member of the ADSM Technical Support team, based in Newcastle upon Tyne, United Kingdom. He has worked at IBM for 12 years, initially in Hursley Laboratories, moving to a field support role, working with many customers in Scotland and northern England. Barry spent the last three years working with ADSM and other client/server products. He has written two other Red Books on LANRES. He holds a degree in electronic engineering.

Thanks to the following people for their invaluable contributions to this project:

Tim Mortimer
Alan Tippett
David Wray
International Technical Support Organization, San Jose Center

Cindy Jiang
IBM ADSM Development, San Jose

Mike Collins
Don Moxley
IBM ADSM Development, Tucson

Dale McInnis
IBM Toronto

Karl Grose
University of California, Berkeley

and to:

Maggie Cutler
ITSO-San Jose Center Editor

Comments Welcome

We want our redbooks to be as helpful as possible. Should you have any comments about this or other redbooks, please send us a note at the following address:

redbook@vnet.ibm.com

Your comments are important to us!

Part 1. Introduction

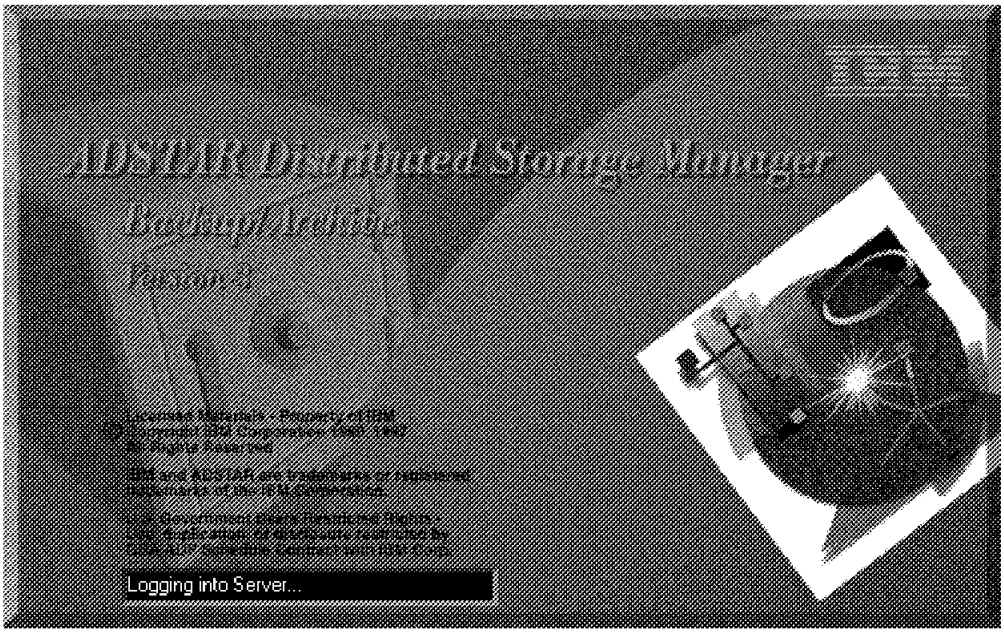

Chapter 1. ADSM Overview

In this chapter we provide a brief overview of ADSTAR Distributed Storage Manager (ADSM). We look at its components and describe key functions such as scheduling, policy management, and the Disaster Recovery Manager (DRM) feature.

1.1 What Is ADSM?

ADSM, IBM′s solution to enterprisewide distributed storage management, is a client/server program product. It provides highly automated, centrally scheduled, network-based backup and archive functions for workstations and local area network (LAN) file servers. ADSM supports a wide variety of IBM and non-IBM clients and servers (see Figure 1) and addresses the need for customer asset protection and data availability for distributed environments.[1]

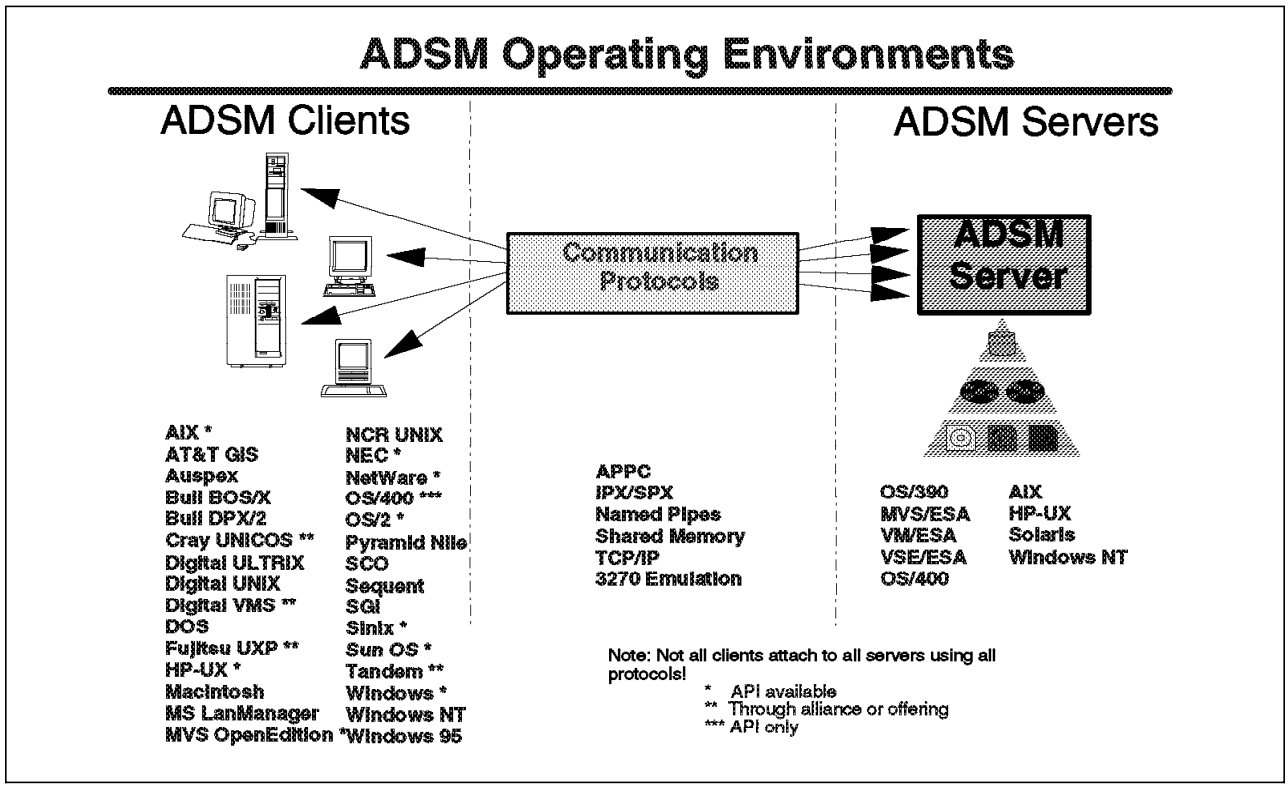

Figure 1. ADSM Platforms

[1] The material in this chapter is summarized from the following sources:

- Chapter 1 of *Using ADSM to Back Up Databases*, SG24-4335
- Chapters 1 through 5 of *ADSM for AIX: Advanced Topics*, SG24-4601
- Chapter 1 of *Disaster Recovery Manager Administrator′s Guide and Reference*, GC35-0238

1.2 Main Components

In this section we look at the main components of ADSM—the backup/archive client, administrative client, and server and briefly review the application client.

1.2.1 Backup/Archive Client

The backup/archive client runs on the workstation and, depending on the platform, provides both a graphical user interface (GUI) and a command line interface (CLI). Although all clients are similar, each has the look and feel of the platform on which it runs. Thus users can back up or restore files, using an interface with which they are familiar.

The main functions of ADSM are backup and restore. You can back up all of your files (full), specific files (selective), or only those files that have changed since your last backup (incremental). You can specifically include or exclude certain files from being backed up.

The file compression provided on the client platforms reduces network traffic and the amount of storage required on the server to store the files.

ADSM′s cross-user restore and cross-platform restore provide you with significant flexibility. Cross-user restore enables you to authorize someone else to restore your files. Cross-platform restore enables you to restore your file on a platform different from the platform on which it was backed up. For example, you could back up your file from a DOS workstation but then restore it to an OS/2 workstation. Cross-platform restore can be extremely useful when you migrate to new workstation platforms, or even if you happen to work at a different office one day that has different workstations. You will still have access to the data you backed up!

A separate archive/retrieve function is also part of ADSM. This function provides a way for you to store files that you may not use but have to retain for long-term storage. Archive is also useful as a way of reducing the disk space on your workstation. You can archive files for long-term storage and erase the original files from your workstation to create room for more active files and applications.

1.2.2 Administrative Client

As shown in Figure 2 on page 5, the ADSM administrative client has functions that allow an administrator to control and monitor server activity, define storage management policies for workstation files, and set up schedules to provide backup and archive services.

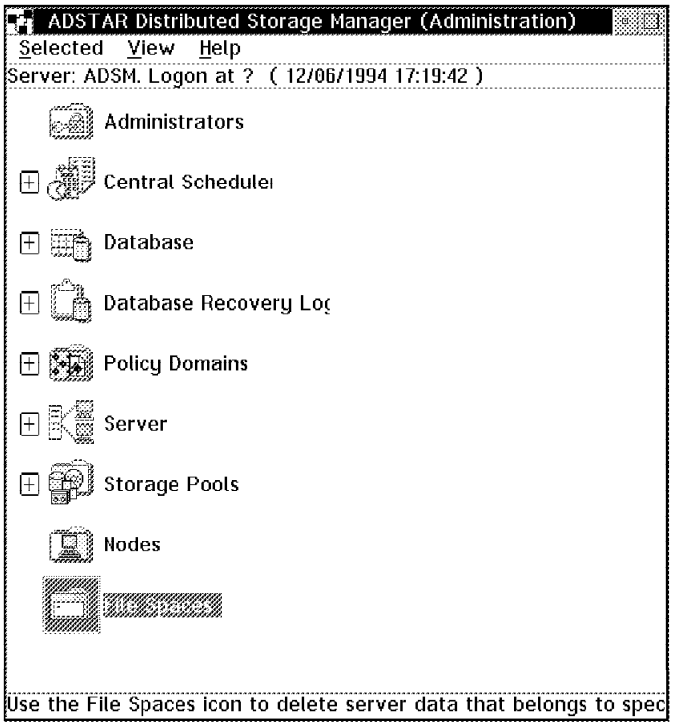

Figure 2. ADSM Administrative Client GUI

An administrative client is a program that enables administrators to control and monitor the server through administrative commands. The administrative program can be installed on a programmable workstation (PWS), personal computer, or mainframe. An administrative client passes commands through an administrative command line. In some cases, a GUI has been added to the administrative client code.

ADSM provides a hierarchical structure to the authority you can grant an administrator. Thus you can establish as flexible an administration scheme as you would like while still providing control over your system. The ADSM administrator with overall authority is called the *system administrator*. Other administrators are called *policy, storage, operator*, or *analyst administrators*, depending on which part of the system they control. Their administrative tasks are separated into logical categories, such as controlling management policies, storage pools and databases, operation of the server, and analysis of certain server events.

Dividing the administrative authority according to logical categories of tasks is not the only way of granting authority. You can also divide the administrative authority by organization. You can give the logical categories of authority to a department, but only for the data that belongs to that department. For example, you can give a department policy and storage authority for the policy domain and storage pools that it owns.

1.2.3 Server

The server component provides storage resources and services for the backup/archive clients. Users can back up or archive their files onto server storage resources such as disk, tape, or optical devices that the ADSM server policy manages and monitors. .

Figure 3 shows the two key components of the ADSM server: the storage pools where the client files are actually stored, and the database that serves as an inventory or index to the client files within the storage pools. The database consists of the database space and the recovery log. The recovery log keeps track of all changes made to the database.

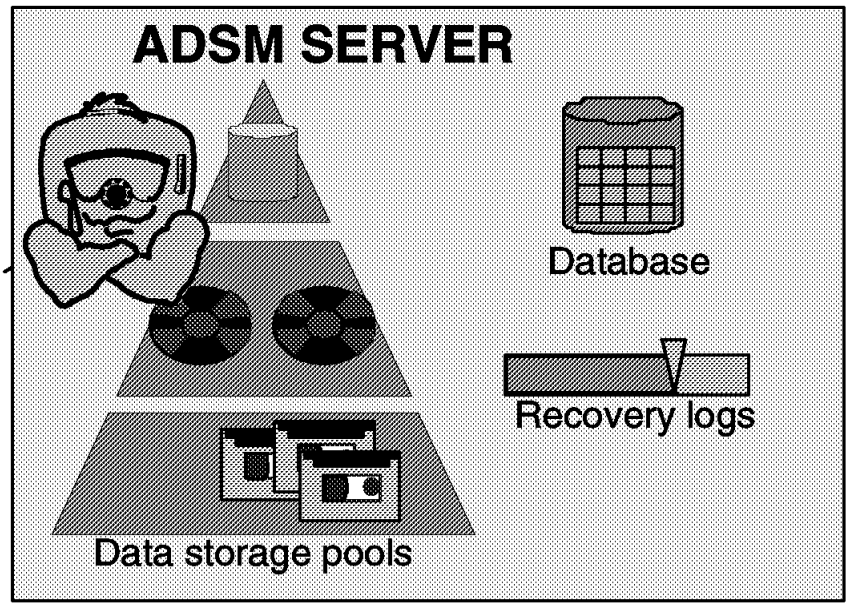

Figure 3. ADSM Server Components

The storage pools contain the client files that have been backed up, archived, or migrated.

You can use a hierarchy of storage media to define the storage pools. The pools can contain disk storage, optical devices, and tape devices. Each ADSM server platform supports a different set of storage media, so please verify the devices that are supported in your environment.

You can move data automatically through the storage hierarchy onto less expensive media with ADSM′s migration function. Additional management functions are provided, such as reclamation and collocation for tape management.

The ADSM server is multitasking, so multiple clients can back up data concurrently.

The ADSM database is the heart of the server. It is critical to the operation of ADSM because it contains file location information as well as policy and scheduling information. The following information is stored in the database:

- Information about registered client nodes
- Policies assigned to those client nodes

- Schedules and their association with client nodes
- Event records, such as whether a schedule successfully completed
- The activity log that contains the messages generated by the server
- Information about ADSM volumes
- The data storage inventory, that is, the information used to locate files that reside in storage pools
- Disaster recovery information (if the DRM feature is installed on the ADSM server. DRM is currently supported on the AIX and MVS ADSM server platforms.)

The database has all of the features associated with a database management system. Because the database is critical, many features are built into ADSM to help maintain the availability, integrity, and performance of the database. Two of these features are the recovery log and mirroring.

A recovery log is used to help maintain the integrity of the database. It keeps track of all changes made to the database, so that if a system outage were to occur, a record of the changes would be available in the log. When a change to the database occurs, the recovery log is updated with some transaction information before the database is updated. Thus uncommitted transactions can be rolled back during recovery so that the database remains consistent.

Mirroring is the process of writing the same data to multiple storage devices at the same time. The administrator can configure the server so that up to three copies of the database and recovery logs are maintained at all times. This mirroring capability provides nondisruptive and immediate recovery from physical failures on database and recovery log volumes.

If a mirrored volume encounters a media failure, the server automatically places the failing volume offline and continues database operations, using the other mirrored copies. Once the failed disk is replaced and made available to the server, it is automatically synchronized with the intact copies.

The mirroring facility improves database performance. The mirrored copies are treated equally; there is no concept of primary copy and alternate copies. Therefore, the server reads from the database copy that is on the device with the best response time.

Another server function, export/import, creates a self-describing copy of specified server information. Information that can be exported includes:

- Administrator information
- Client node definitions
- Policy information
- Backup and archive data

Export/import is useful for migration and conversion, workload balancing, and cloning of information.

ADSM provides extensive ADSM server database and storage pool backup facilities. Incremental backups are provided as well as a mechanism for offsite backups to aid in disaster recovery.

1.2.4 Application Client

The application client is a software application that runs on a workstation and uses the ADSM application programming interface (API) to back up, archive, restore, or retrieve objects from an ADSM server.

The application client program enables other IBM and non-IBM products to use the storage management services of ADSM. The application client allows applications to back up or archive valuable data in any format that an application programmer specifies.

The number of ways of using the API is unlimited. You can use it to improve the handling of nonfile data in the enterprise, such as databases or image volumes. You could, for example, provide extensions to the existing ADSM backup and restore functions to meet your user′s needs or write a virtual tape device driver so that other applications can use ADSM transparently.

The API is available for the C programming language.

1.3 Functions

In this section we look at the ADSM backup and restore, archive and retrieve, central scheduling, and policy management functions. Then we review the disaster recovery features and functions of ADSM.

1.3.1 Backup and Restore

ADSM backup and restore are shown in overview in Figure 4 on page 9. The backup process creates a copy of a client file on the ADSM server. The backup process also backs up the directory in which the file resides.

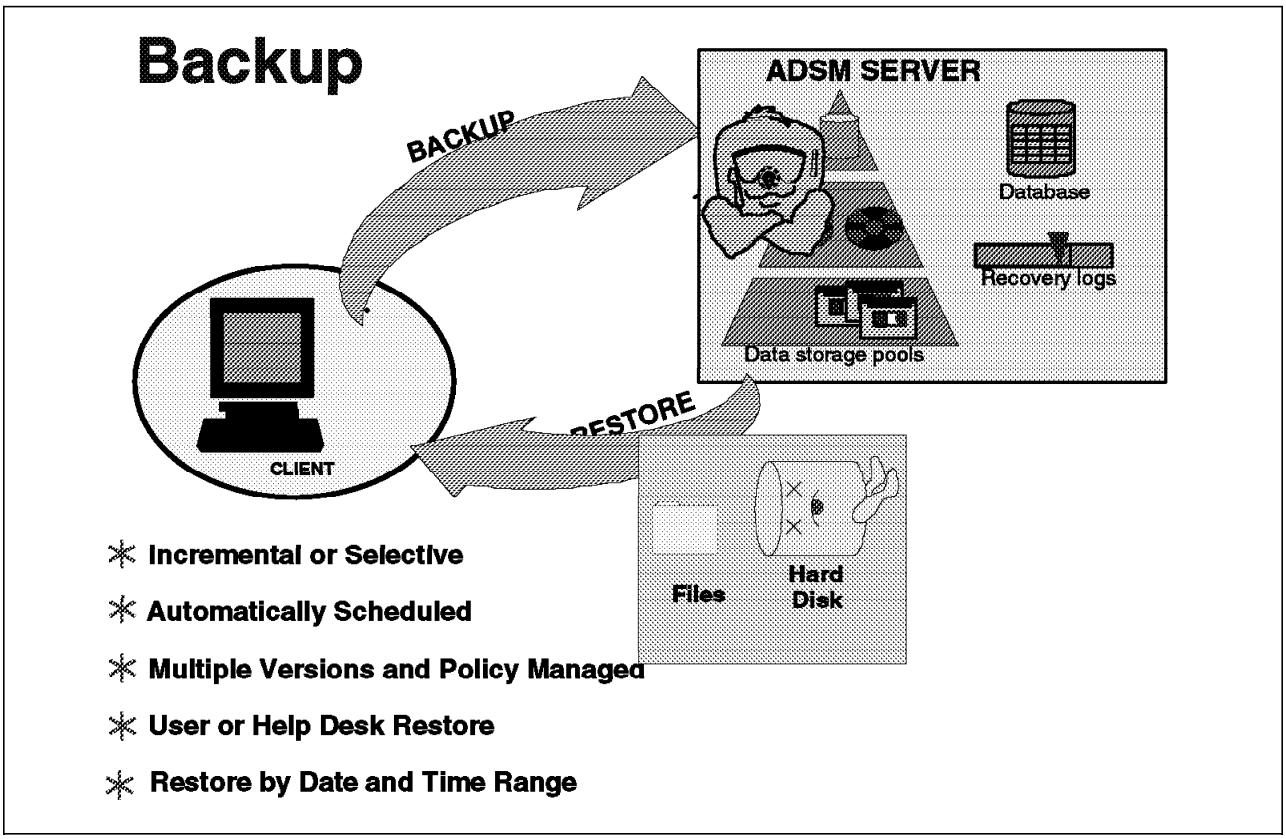

Figure 4. ADSM Backup and Restore

Incremental backup sends to the server the files that have changed since the last backup. The first time an incremental is done, all files are sent to the ADSM server. This is a full backup. ADSM determines that a file has changed if any of the following has changed: file size, date and/or time stamp, file owner, file group, file permission, or attribute change time.

Selective backup specifies which files a user wants to back up. A selective backup can consist of a single file, or a user can select a directory or subdirectory tree to back up. Because wildcards are allowed in the specification, there is great flexibility in file selection.

The files are backed up according to policies that the administrator has predefined. The policies define, for example, how many backup versions should be retained in the ADSM storage pools, how long to retain those versions, and whether to back up files that are in use.

Restore is the process of copying a backup version from the server to the client. This process is system assisted; that is, the system performs the restore for the user. The user does not have to call the ADSM administrator to request restoration of the file.

1.3.2 Archive and Retrieve

The ADSM archive and retrieve process is shown in overview in Figure 5. The archive process creates a copy of a client file on the ADSM server. As with backup, archived files are managed on the basis of policies; however, the archive function does not have a concept of versioning. You can archive multiple versions of a file by invoking the archive function multiple times. In other words, each archived copy is treated as a separate file, not as multiple versions of a single file.

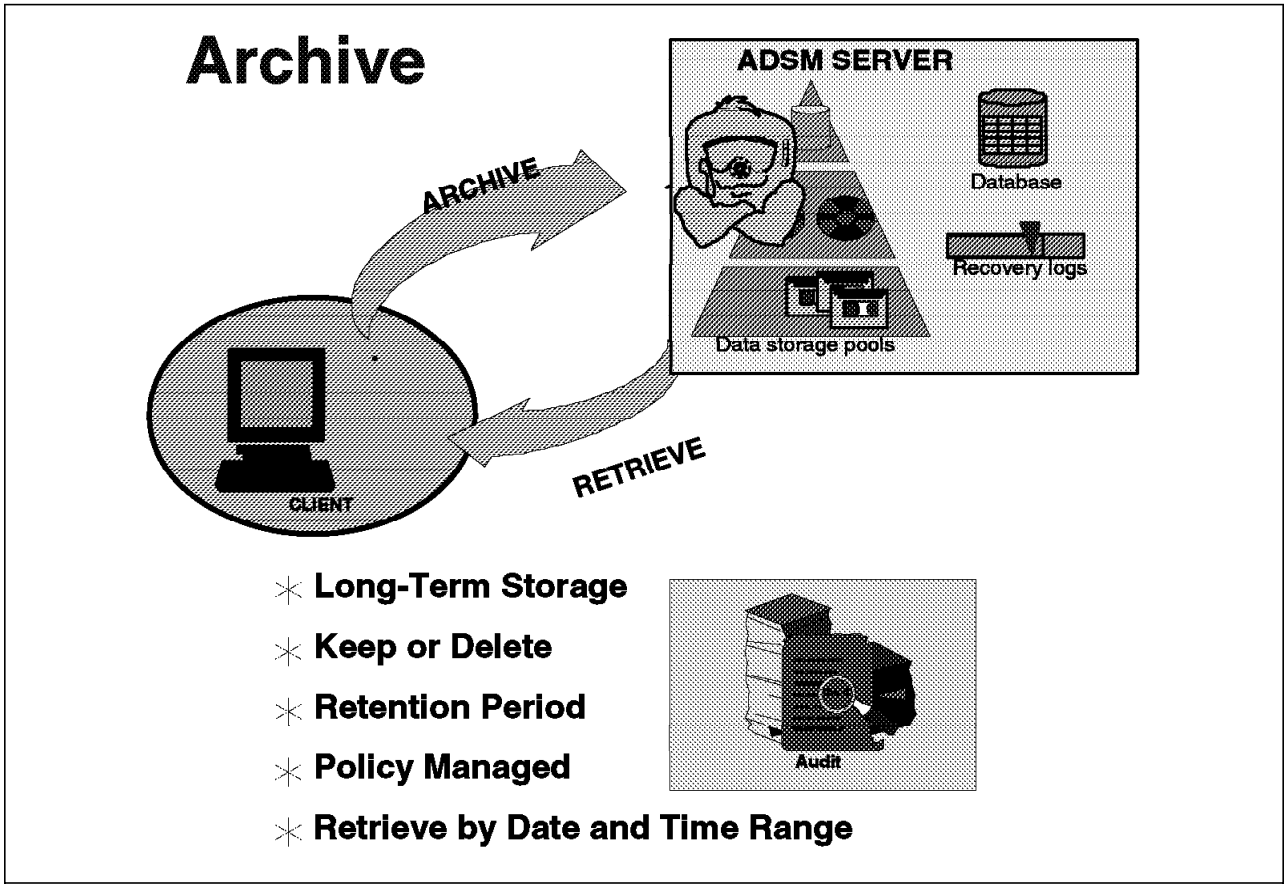

Figure 5. ADSM Archive and Retrieve

A user can save a description of an archived file so that it will be easy to retrieve the file if multiple files are archived with the same file name.

The key difference between backing up a file and archiving a file is that the user can erase the original file after archiving it. The archived version is expected to be retained for a long time. Erasing the original file does not affect the retention period for the archived file.

1.3.3 Central Scheduling

ADSM central scheduling is shown in overview in Figure 6 on page 11. This facility automates the initiation of client backup, archive, restore, and retrieve, as well as ADSM server administrative operations. It also can schedule any client operating system command and ADSM client macros. New clients can be easily associated with schedules in a nondisruptive manner. The central scheduler consists of client and server processes that cooperate to execute the scheduled functions.

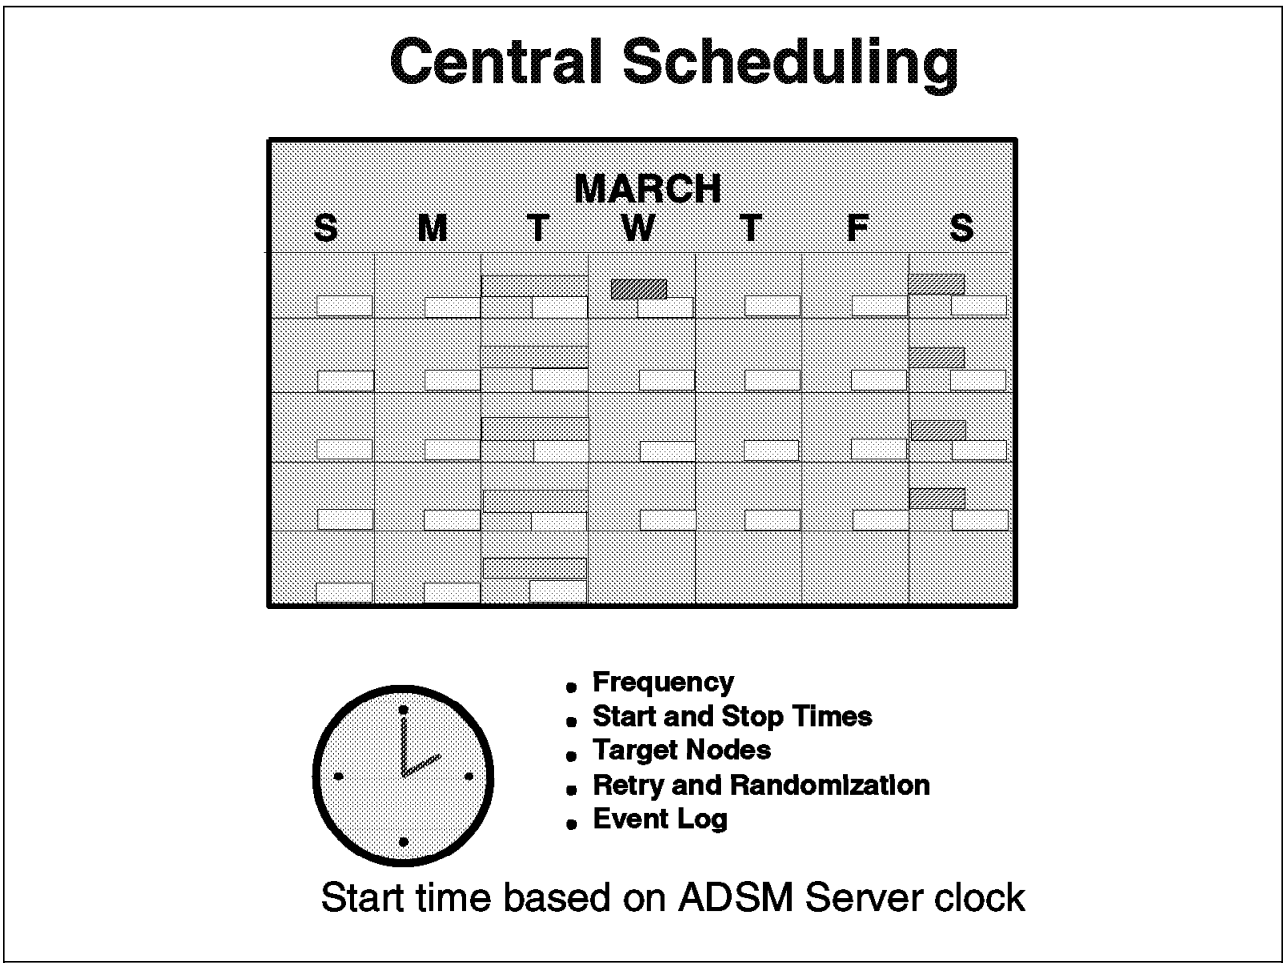

Figure 6. ADSM Central Scheduling

The administrator is responsible for defining and maintaining the schedules and has the authority to prioritize clients so that clients that contain more important data are given preferential treatment.

A schedule event log is maintained in the server database. Whenever a schedule process starts or fails, an event record is written to the log. An administrator can query the log to determine whether scheduled events completed successfully or not.

1.3.4 Policy Management

ADSM enables you to manage the backup and archive process according to policies you establish for your enterprise. The granularity of control that you have is down to the file level. You can decide on how granular you want your policies to be. You can establish an overall system policy, policies by department or organizational structure, or policies by user or file name. Policy management makes ADSM a true system-managed storage implementation. The elements of policy management are discussed below.

1.3.4.1 Policy Domain

A policy domain is a group of clients that are working according to the same set of policy needs. A policy domain provides a logical way of managing backup and archive policies for a group of client nodes. There is no limit to the number of policy domains that can be defined on an ADSM server. Policy domains can be used to provide standard storage management policies to most users, group together clients that have similar storage management requirements, limit the number of clients to be managed by a single policy administrator, and restrict the number of management classes to which users have access.

1.3.4.2 Policy Set

Each policy domain can contain one or more policy sets. A policy set contains one or more management classes. A policy domain can have more than one policy set, but only one policy set can be activated at any one time. Each policy set contains a default management class and can contain any number of additional management classes. Policy domain and policy set information is stored in the server database.

1.3.4.3 Management Class

Policy sets contain one or more management classes. Management classes contain a backup copy group and/or an archive copy group or no copy group.

You can think of management classes as a Service Level Agreement you have with your clients on how their backup and archive data will be handled. There is a concept of binding the management class to the file when it is backed up or archived. Thus the management class is associated with that file. You can rebind a file with a new management class. Users can use the default management class or explicitly select a management class that is within the active policy set to which they have access.

1.3.4.4 Copy Group

Copy groups are where you specify the parameters that will control the generation and expiration of backup and archive data. There is a copy group for backup and one for archive. In the current version of ADSM, all copy groups are named STANDARD. Copy group information is stored in the ADSM server database. Figure 7 on page 13 illustrates some of the options for controlling your ADSM file copies. (Figure 7 also shows the HSM Space Management options. These options are not really relevant to restores but are listed for completeness.)

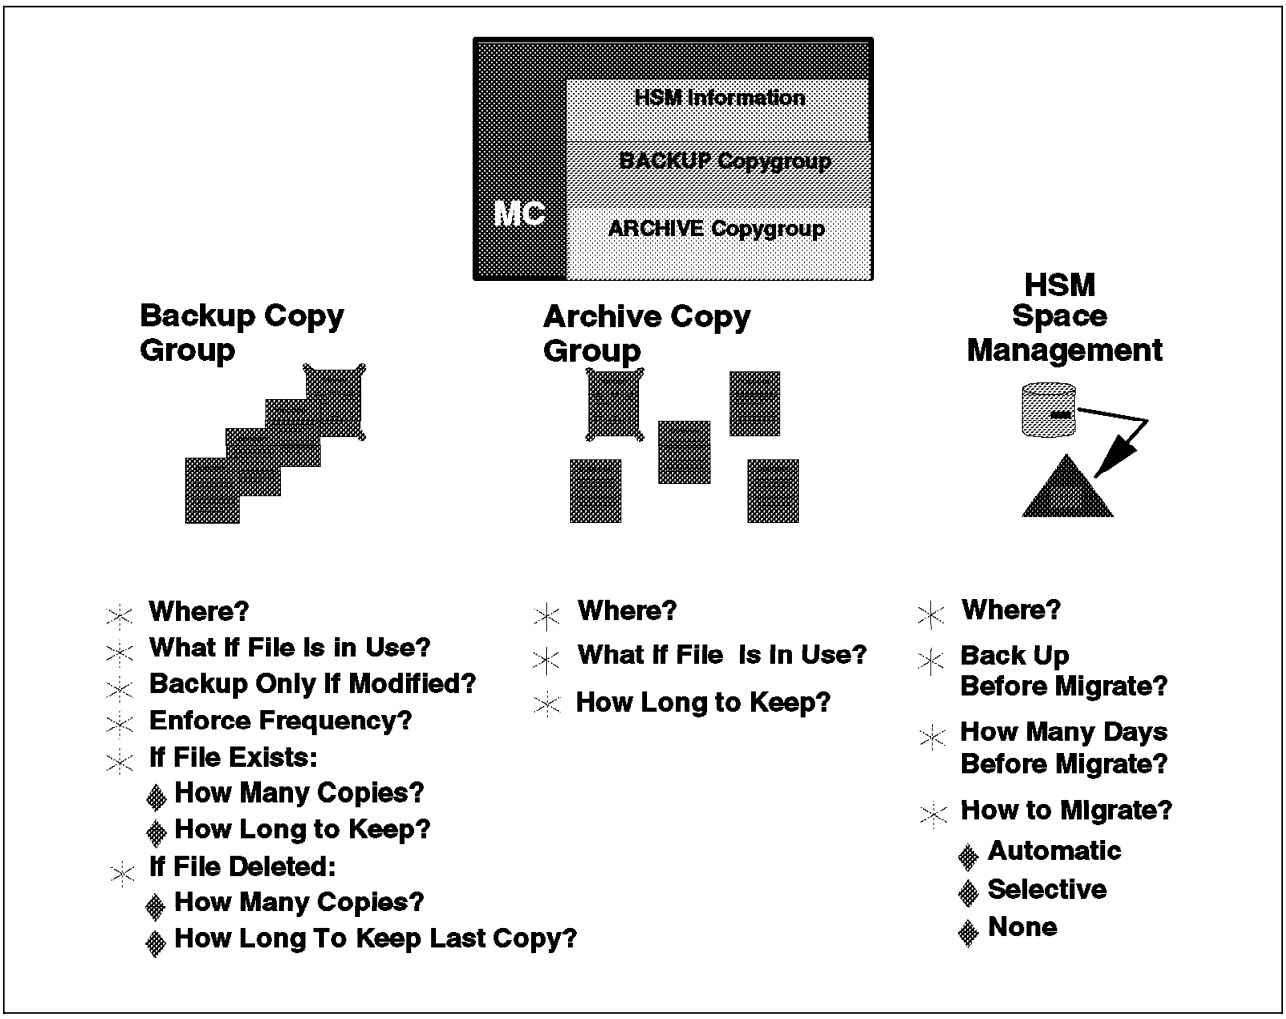

Figure 7. ADSM Copy Group Parameters

Backup, archive and space management operations can be controlled at all levels from drive through directory, down to individual files. Once bound to a management class, the file copies in ADSM are managed individually according to their management class copy group attributes. By changing the attributes of a management class, all files bound to that class are managed as well, right down to a file granularity. This enables tremendous flexibility of control. Some examples of copy group parameters are described below.

Destination specifies the name of the storage pool where the server stores the backed up or archived files.

Frequency for a backup file specifies the minimum number of days that must elapse between incremental backups. This parameter is not used for selective backups. Frequency for an archive file is always command (CMD) as archives are only on command. A file is archived only when a client issues an archive command or chooses archive from the GUI.

Versioning applies only to backup files. You can specify two different parameters to tell ADSM how many versions of a backup file you want it to maintain. The *Version data exists* parameter specifies the maximum number of different backup versions the server retains for files and directories that exist on the client

workstation. The most current backup version is called the *active version.* All other versions are called the *inactive versions.*

When the maximum number of versions is exceeded, the server rolls off the oldest version. The *Version data deleted* parameter specifies the maximum number of different backup versions the server retains for files and directories that have been erased from the client workstation.

The *retention period* parameter specifies how long to retain the backed up and archived files. There are two retention parameters for backed up files that correspond to the two types of versioning, and there is one retention parameter for the archived files. *Retain extra versions* specifies how many days the server retains the inactive backup versions when the original file no longer exists on the client's workstation. *Retain only version* specifies how many days the server retains the backup versions it has of a file when the original file has been deleted from the workstation. *Retain version* specifies the number of days an archived copy remains in data storage.

With the *mode* parameter you can specify file backup depending on whether the file has changed since the last backup. This parameter applies to incremental backups, not selective backups. The options for mode are *modified* and *absolute*. *Modified* indicates that you want to back up the file only if it has changed. *Absolute* indicates that you want to back up the file regardless of whether it has changed. For archive files, the mode is always absolute.

Serialization specifies how files or directories are handled if they are modified during the backup or archive process. The serialization parameter has four options: static, shared static, shared dynamic, and dynamic:

- **Static**

  Static specifies that if a file or directory is modified during the backup or archive process, ADSM will not back up or archive the file. The static mode is not supported on the DOS platform.

- **Shared static**

  Shared static specifies that ADSM will retry the backup operation as many times as specified in the client options file. The default is four retries. If the file or directory is modified during each backup or archive attempt, ADSM will not back up or archive the file.

- **Shared dynamic**

  Shared dynamic specifies that if a file is modified during a backup or archive attempt, ADSM will back up or archive the file only on its last retry.

- **Dynamic**

  Dynamic specifies that even if the file is modified during the backup or archive attempt, ADSM will back up or archive the file anyway. No retries are required.

1.3.5 Disaster Recovery and ADSM

Critical ADSM elements that must be considered when planning for disaster recovery of an ADSM server are listed below. We do not cover these elements in any detail in this book as they are covered very well in the redbook entitled *ADSM Server for Windows NT Configuration and Recovery*, SG24-4878.

- ADSM database

The database manages information about the location of client backup, archive, and migrated files residing in storage pools, and it records information for ongoing server operations. The database also records storage volume locations for backed up, archived, and migrated files.

- Recovery log

  The recovery log is used to maintain a consistent database image by recording changes made to the database as a transaction proceeds. A transaction is any exchange between a client and the server. If a transaction completes successfully, database changes are committed, and permanent changes are made to the database. If a transaction fails, database changes are undone by removing them from the database and recovery log.

- Storage pools

  The storage pools contain the actual backup, archive, or migrated copies of client files.

- Copy pools

  Copy pools are used to provide backup copies of the objects held in the ADSM storage pools. Although copy pools are, in effect, ″backups of backups,″ you must plan for their use for data that will be required after a site disaster. Copy pool volumes can be taken offsite, either manually or by the Version 3 server-to-server communications function.

1.3.6 Disaster Recovery Manager

DRM is available for ADSM Version 2 servers on AIX, MVS, NT, HP, and Sun servers. DRM automatically generates a server disaster recovery plan. It also manages offsite disaster recovery media and stores client machine recovery information. It is the latter that we focus on in this section as it allows the centralized organization of information that will be needed in the event that client machines are lost in a disaster or stolen.

You can use DRM to store the following ADSM client information:

- Identity and priority of ADSM clients, according to application or business needs
- ADSM client machine information
  - Business priority and machine location
  - Association of one or more ADSM node definitions with a machine
  - Machine characteristics (such as machine type, RAM, hard drives, and network adapters)
  - Recovery instructions
- Boot media requirements
- Associations with recovery media

Once this information about the ADSM client is defined and stored at the ADSM server, it becomes part of the disaster recovery plan and hence is available to assist in the bare metal restore of ADSM clients. For example, in a disaster recovery situation, once the ADSM server has been restored, the queries listed below could be made to obtain information about how to recover the ADSM client hardware platform. These queries must be issued before the ADSM client code is used for data restoration.

Note: The query examples given below are taken from Chapter 1 of the *Disaster Recovery Manager Administrator's Guide and Reference,* GC35-0238, for ADSM on MVS.

When the ADSM administrator issues this command:

```
QUERY MACHINE BUILDING=20.21 FORMAT=DETAILED
```

ADSM displays a list of client machines in building 20.21 and indicates their restore priority:

```
            Machine Name: DILPER.RZ.UNI-KARLSRUHE.DE
        Machine Priority: 1
                Building: 20.21
                   Floor: 2
                    Room: 206
            ADSM Server?: No
             Description: Rocky's Server
               Node Name: DILPER
     Recovery Media Name: DASBOOT
        Characteristics?: Yes
  Recovery Instructions?: Yes
```

To determine the location of the boot media for a particular machine, the ADSM administrator would issue this command:

```
QUERY RECOVERYMEDIA DASBOOT
```

ADSM displays the following information in response:

```
Recovery Media Name  Volume Names  Location  Machine Name
-------------------  ------------  --------  ------------
DASBOOT              VOL1 VOL2     KERKER    DILPER.RZ.UNI-KARLSRUHE.DE
```

To determine the machine-specific recovery instructions for the machine, the ADSM administrator issues this command:

```
QUERY MACHINE DILPER.RZ.UNI-KARLSRUHE.DE FORMAT=RECOVERYINSTRUCTIONS
```

ADSM displays the following information in response:

```
  Recovery Instructions for DILPER.RZ.UNI-KARLSRUHE.DE.
  Primary Contact: Alpir S. Bacher (home: ++49-721-608-4040)
  etc...
```

To determine the hardware requirements for machine DILPER.RZ.UNI-KARLSRUHE.DE, the ADSM administrator issues this command:

```
QUERY MACHINE DILPER.RZ.UNI-KARLSRUHE.DE FORMAT=CHARACTERISTICS
```

ADSM displays the following information in response:

```
Intel Pentium 160MHz,64MBytes RAM,SCSI adapter 1GByte hard drive
SCSI CDROM, 3COM EtherlinkIII network adapter
Partitioning of Hard Drives:
BootManager
C: DRIVE_C logical partition 200 MBytes
D: DRIVE_D logical partition 450 MBytes
E: DRIVE_E logical partition 200 MBytes
F: DRIVE_F logical partition 150 MBytes
etc...
```

The recovery of ADSM clients presupposes the existence of a recovery plan to re-create the hardware, operating system, and communications environment necessary to run the ADSM client code before that client's file systems can be restored from the ADSM server storage pool.

Although the DRM feature does not contain automation to aid in the recovery of the underlying ADSM client hardware, operating system, and communications software, it does permit at the ADSM server the storage of such instructions as client recovery information. Hence, in a disaster recovery situation, after you have used DRM to restore the ADSM server, you can also use it to retrieve information about how to undertake the bare metal restore of the ADSM clients.

Of course, the DRM feature is not a prerequisite to using the bare metal restore techniques we describe in this book. Information about the steps necessary to restore the ADSM client platform in the stages before using the ADSM client code (namely, restoring the hardware, partitioning hard drives, and restoring the operating system software, communications, and ADSM client code), although conveniently stored at the ADSM server by using DRM, can also be stored elsewhere (for example, on hardcopy stored offsite, or in a file on another machine).

Chapter 2. Typical Customer Scenarios

In this chapter we examine a number of typical customer scenarios, looking at ways in which Windows NT is installed and how ADSM may best be installed to provide maximum backup resilience.

For each scenario we highlight the exposures to the various system and application components and provide pointers on which procedures should be followed, and in which order, to recover from errors as minor as a user deleting a file by mistake and as major as the complete loss of the site.

We recommend that you take some time to decide which example is closest to your environment before reading the rest of the book.

2.1 Using Simple, Native NT Backup

In this scenario, Figure 8 on page 20, each computer′s backup must be set up and controlled separately. The NT server should be backed up each day with the NT backup program. Different versions or retention periods are foreign words for this ″Backup Solution″, and the manual maintenance of the tape management grows with the daily growth of data. Tape management is particularly difficult if versioning is required, as versions of a file must be a manually recovered.

If applications are installed on this system, additional recovery is required.

Each attached client (Windows 95, OS/2, or others) must be backed up separately. There must be a backup operator, who checks the current status of each backup and is responsible for the tape management.

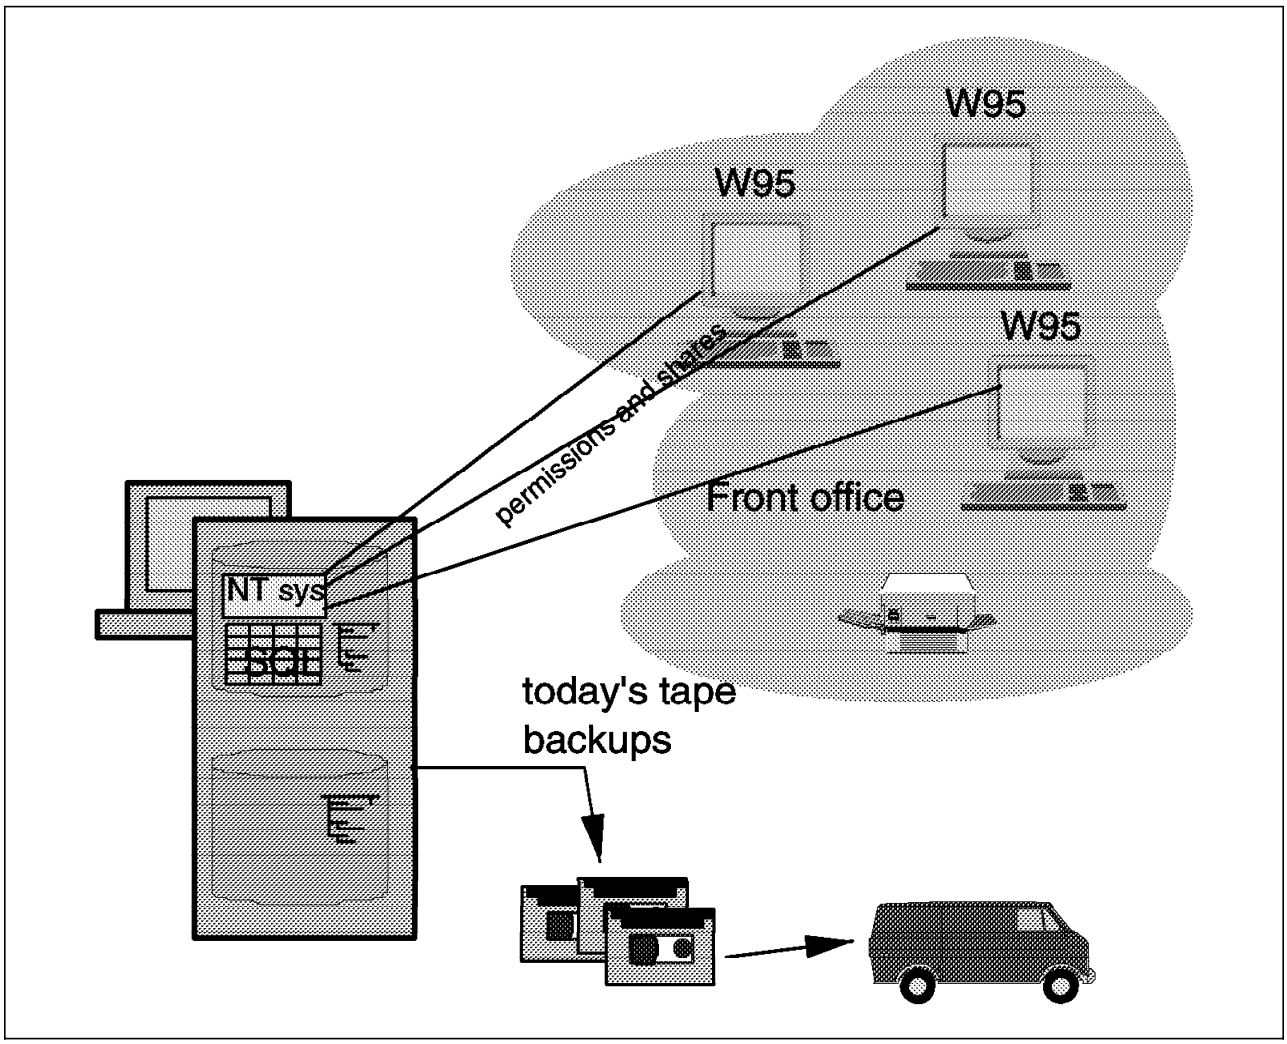

Figure 8. Simple, Native NT Backup

For availability of your system refer to:

- 3.1, “NT Boot Process” on page 37
- 3.2, “Useful Availability Tools” on page 42 and the Microsoft Windows NT Resource Kit
- 3.3, “The NT File System” on page 50
- 3.4, “Configuring NT for Availability” on page 55 and the Microsoft Windows NT Resource Kit

For recovery of your system refer to:

- 4.3, “System Recovery (Bare Metal Restore)” on page 76
- 4.4.2, “Synchronization of the PDC and BDCs” on page 85
- Using the NT Backup program for system recovery (see the Microsoft NT system documentation)

For the application recovery of your system, refer to:

- Application literature from Microsoft, Lotus, IBM, and other vendors

2.2 ADSM Single Server Edition

ADSM is an intelligent backup solution combined with integrated tape management. In the next two examples, Figure 9 on page 22 and Figure 10 on page 23, the ADSM Server is located on the machine that works as the application and network server as well. The network clients typically work with files served from the NT server but (although not shown in the examples) could, by running the ADSM client, make their backups to this ADSM server as well. For clarity, both examples show an ADSM storage pool hierarchy as disk devices only. The storage pools could also include offline media such as tape. The important point is that such media do not leave the site and are part of the primary storage pool — available for onsite recovery.

2.2.1 Backup/Archive Client

In Figure 9 on page 22, we show backups being made to another physical hard drive. By following the procedures listed below, it would be possible to recover from the loss of the primary disk. However, in this example, there is no provision for recovery from a site disaster. (See 2.3, “Single Server with Site Disaster Recovery” on page 24).

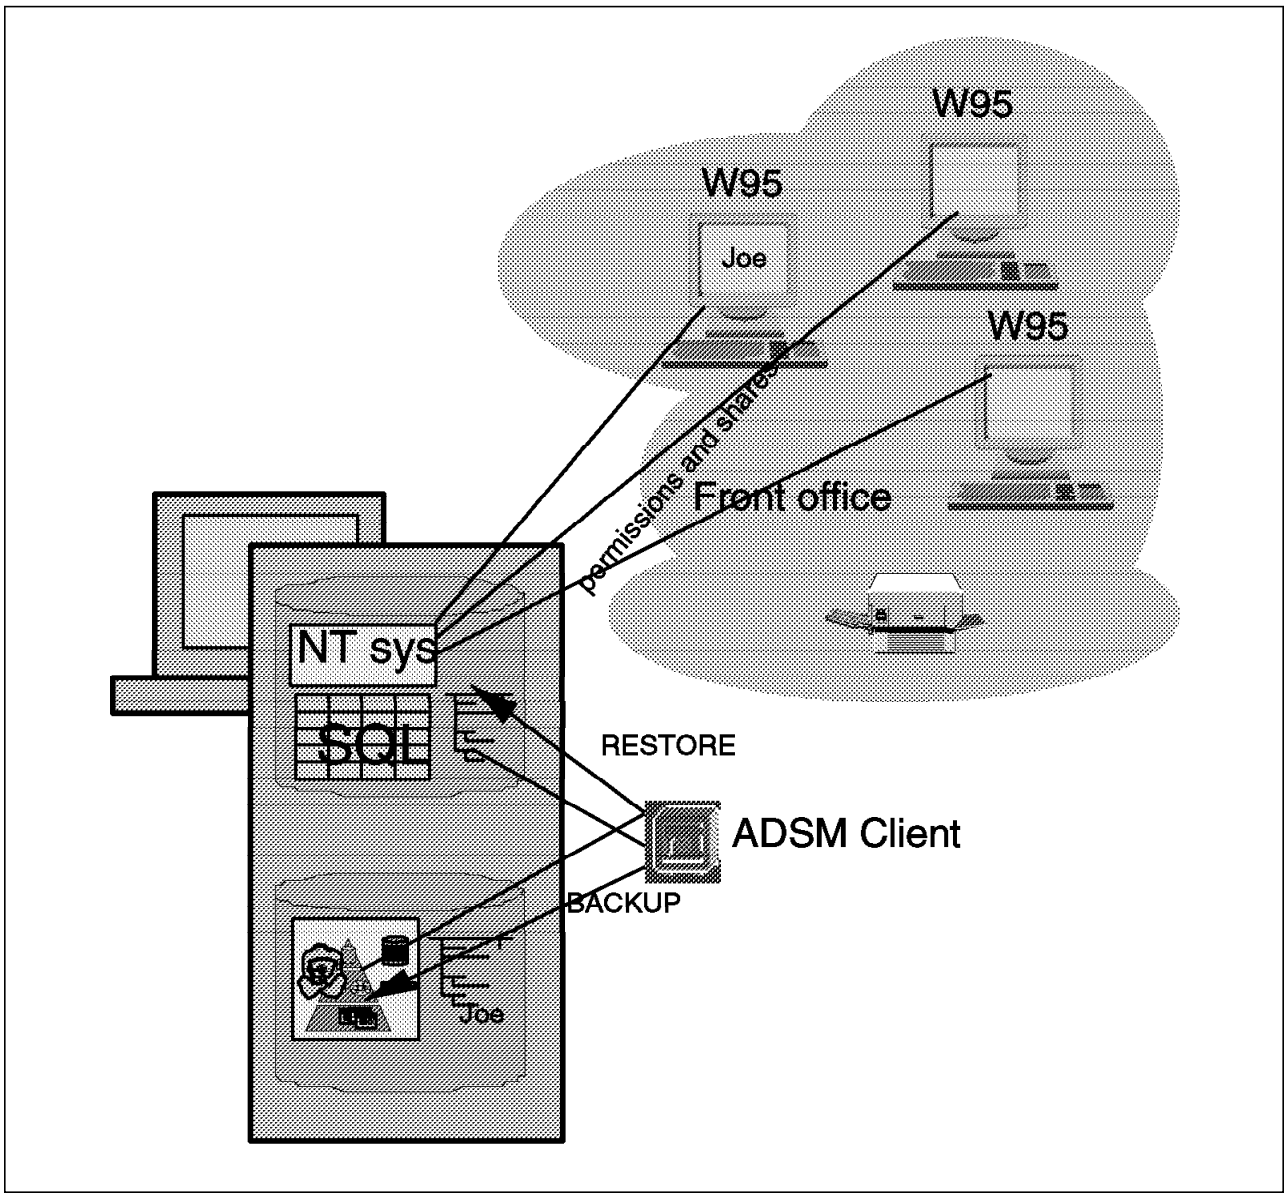

Figure 9. ADSM Single Server Edition. Backup and Restore of the File Systems and Registry

2.2.2 Online Database Backup with ADSMConnect Agents

In addition to the file system backup shown in 2.2, “ADSM Single Server Edition” on page 21, you can make online backups of a database, such as Microsoft SQL or Lotus Notes, using an ADSMConnect Agent. This is illustrated in Figure 10 on page 23. Again, there is no provision for recovery from a site disaster. (See 2.3, “Single Server with Site Disaster Recovery” on page 24).

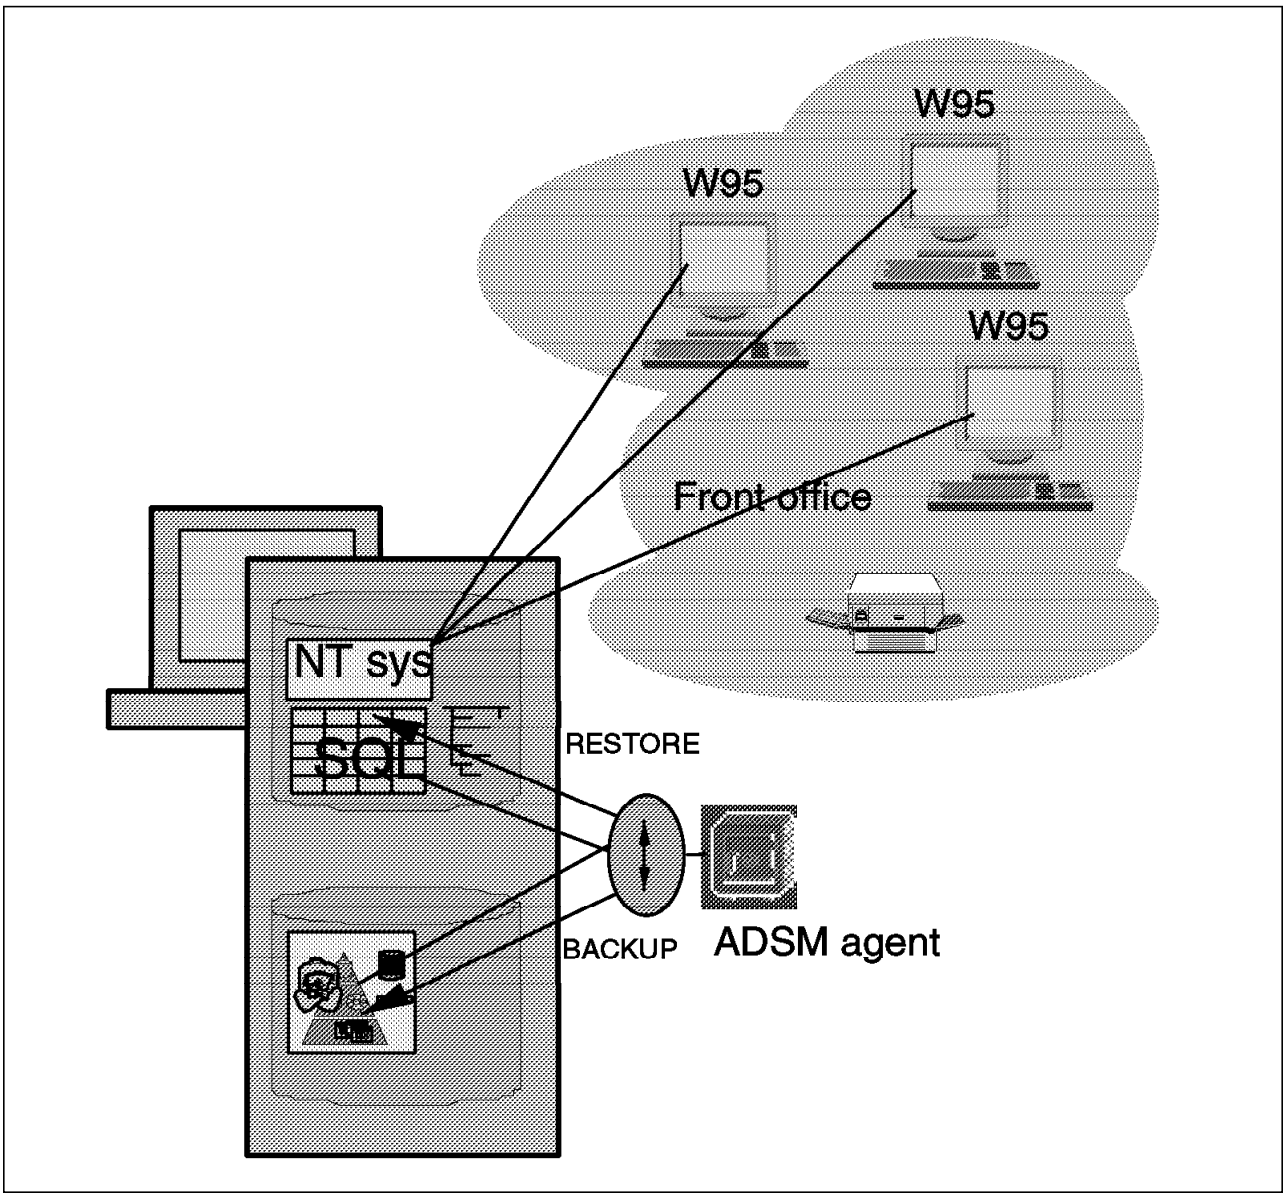

Figure 10. ADSM Single Server Edition - Online Database Backup

For availability of the above systems, refer to:

- 3.1, “NT Boot Process” on page 37
- 3.2, “Useful Availability Tools” on page 42 and the Microsoft Windows NT Resource Kit
- 3.3, “The NT File System” on page 50
- 3.4, “Configuring NT for Availability” on page 55 and the Microsoft Windows NT Resource Kit

For recovery of your system refer to:

- Chapter 4, “System Recovery with ADSM” on page 69
- 4.4.2, “Synchronization of the PDC and BDCs” on page 85
- Using the NT backup program for system recovery (see Microsoft NT system documentation)

- *ADSM Server for Windows NT Configuration and Recovery*, SG24-4878
- *ADSM Client Disaster Recovery*, SG24-4880

For the application recovery of your system, refer to:

- Chapter 5, “Recovering Lotus Notes” on page 95
- Chapter 6, “Microsoft SQL Server Backup” on page 103
- Chapter 7, “Recovering Microsoft Exchange Server” on page 125
- Chapter 8, “Recovering Microsoft Access” on page 145
- Chapter 9, “Recovering DB2 for NT” on page 151
- Application literature from Microsoft, Lotus, IBM, and other vendors

2.3 Single Server with Site Disaster Recovery

This environment, shown in Figure 11 on page 25, is the same as described in 2.2.2, “Online Database Backup with ADSMConnect Agents” on page 22, but you want to protect your location against a loss of the whole site. For this purpose you can use ADSM on its own or ADSM DRM for offline tape management. Additional ADSM functions protect against single failures and a site disaster. The most important of these are the copy pool and database backups taken to media that is stored offsite, away from the primary storage pools and ADSM database.

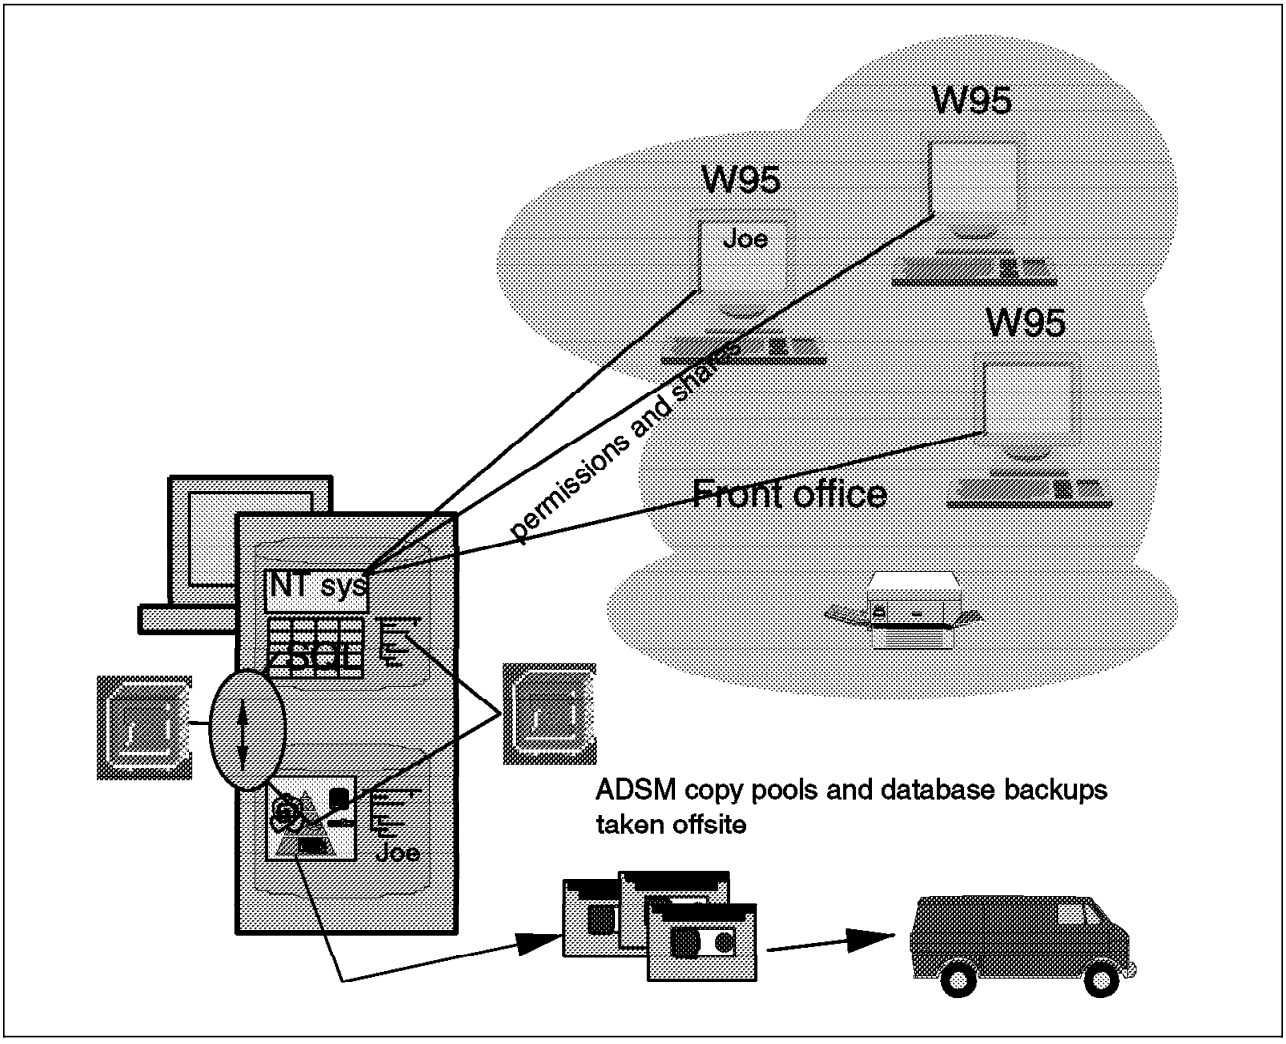

Figure 11. ADSM Server Backed Up to Offsite Media

For availability of your system, refer to:

- 3.1, “NT Boot Process” on page 37
- 3.2, “Useful Availability Tools” on page 42 and the Microsoft Windows NT Resource Kit
- 3.3, “The NT File System” on page 50
- 3.4, “Configuring NT for Availability” on page 55 and the Microsoft Windows NT Resource Kit

For recovery of your system refer to:

- Chapter 4, “System Recovery with ADSM” on page 69
- 4.4.2, “Synchronization of the PDC and BDCs” on page 85
- Using the NT backup program for system recovery (see the Microsoft NT system documentation)
- *ADSM Server for Windows NT Configuration and Recovery*, SG24-4878
- *ADSM Client Disaster Recovery*, SG24-4880

For the application recovery of your system refer to:

- Chapter 5, “Recovering Lotus Notes” on page 95
- Chapter 6, “Microsoft SQL Server Backup” on page 103
- Chapter 7, “Recovering Microsoft Exchange Server” on page 125
- Chapter 8, “Recovering Microsoft Access” on page 145
- Chapter 9, “Recovering DB2 for NT” on page 151
- Application literature from Microsoft, Lotus, IBM, and other vendors.

2.4 Additional NT Servers for Application Servers

This example, shown in Figure 12 on page 27, introduces additional NT server machines that handle applications to lessen the load on the PDC.

2.4.1 Nonsite Disaster Recovery

As the ADSM server is now separate from the application servers, their individual recovery is much easier in the case of a disaster that does not affect the entire site.

2.4.2 Site Disasters

In the case of a site disaster, the NT PDC machine must be recovered first, as otherwise no other machine can use LAN resources. When the NT system on the PDC is recovered, the ADSM server is recovered from offsite copy pool and database backups. You can then start additional restore processes on the PDC machine to restore other partitions or applications. Then you can recover any NT systems on attached NT servers that are also damaged. The sequence for the recovering each system is first the NT system including the ADSM client and then the other disk partitions and the important application. The time to restore applications is the most important consideration. We recommend that you test the NT system recovery regularly to understand the steps required and the time needed.

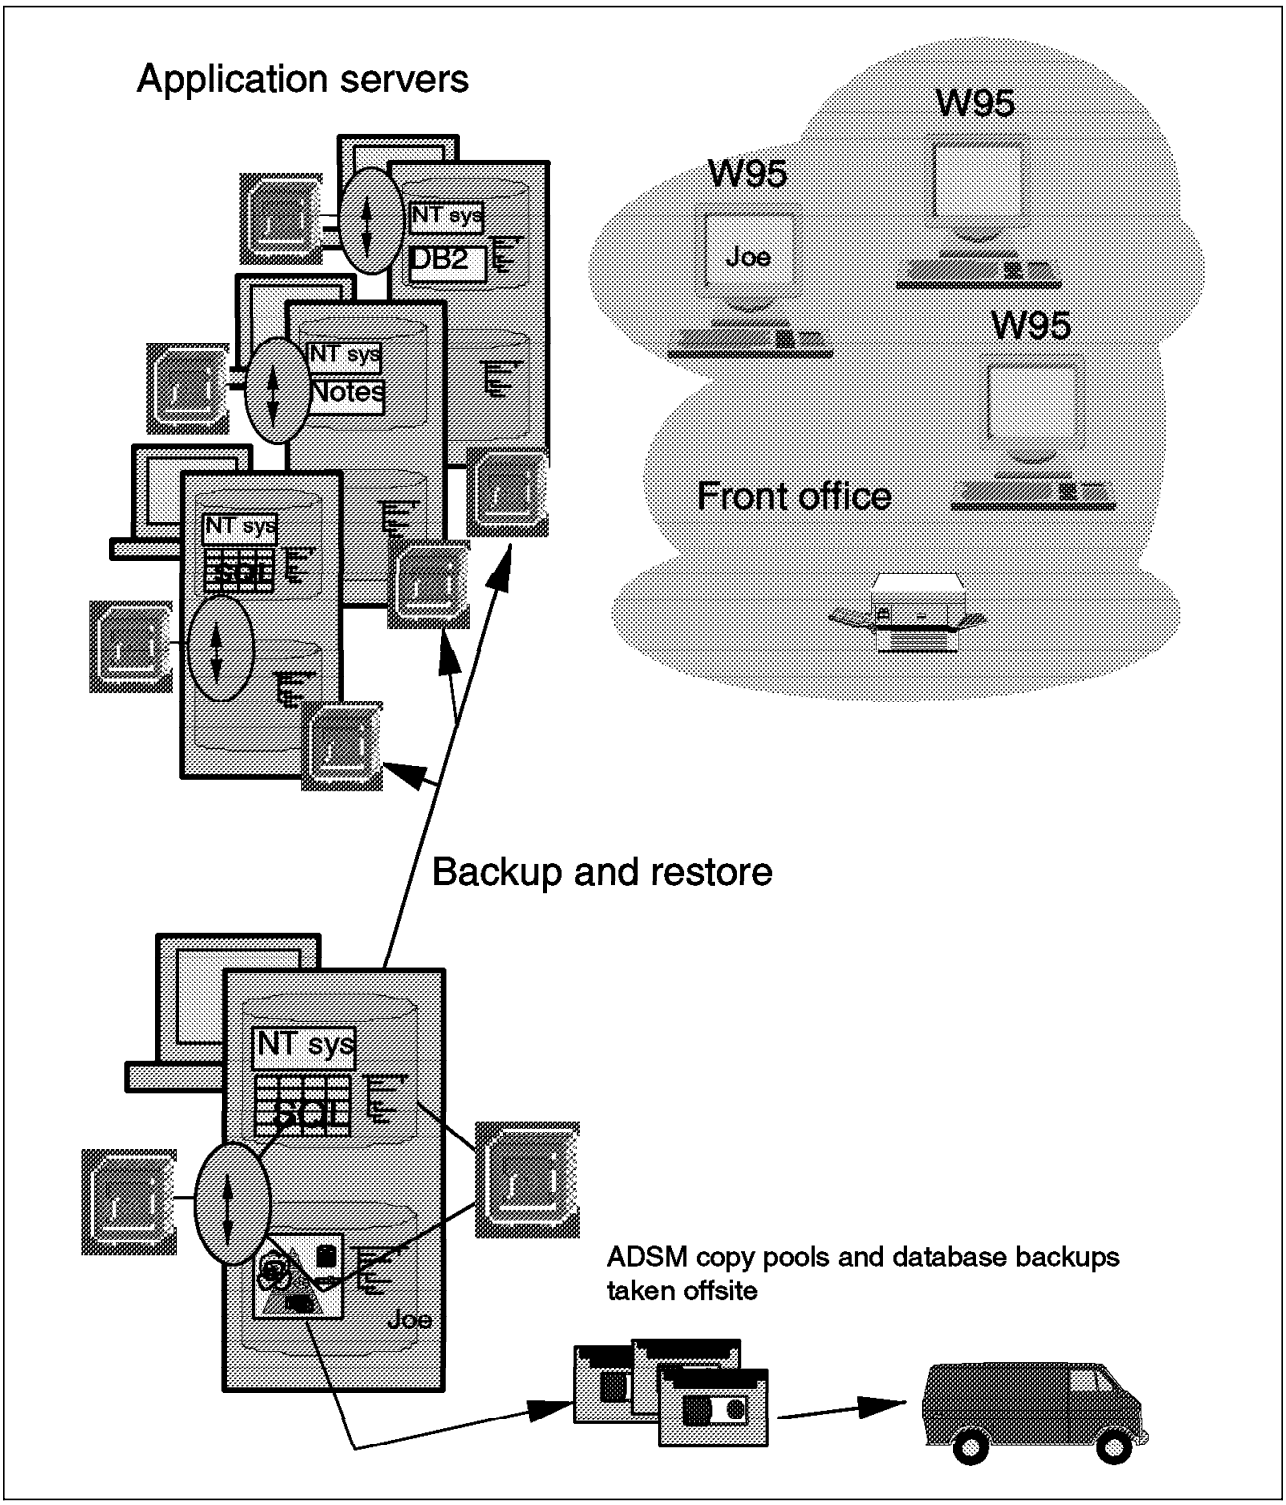

Figure 12. Additional NT Servers As Application Servers

For availability of your systems, refer to:

- 3.1, “NT Boot Process” on page 37
- 3.2, “Useful Availability Tools” on page 42 and the Microsoft Windows NT Resource Kit
- 3.3, “The NT File System” on page 50

- 3.4, “Configuring NT for Availability” on page 55 and the Microsoft Windows NT Resource Kit

  You will need to set up a repair partition, on either a separate hard disk or removable drive, for each NT system in the environment.

For recovery of your system, refer to:

- Chapter 4, “System Recovery with ADSM” on page 69
- 4.4.2, “Synchronization of the PDC and BDCs” on page 85
- Using the NT backup program for system recovery (see the Microsoft NT system documentation)
- *ADSM Server for Windows NT Configuration and Recovery*, SG24-4878
- *ADSM Client Disaster Recovery*, SG24-4880

For the application recovery of your system refer to:

- Chapter 5, “Recovering Lotus Notes” on page 95
- Chapter 6, “Microsoft SQL Server Backup” on page 103
- Chapter 7, “Recovering Microsoft Exchange Server” on page 125
- Chapter 8, “Recovering Microsoft Access” on page 145
- Chapter 9, “Recovering DB2 for NT” on page 151
- Application literature from Microsoft, Lotus, IBM, and other vendors

2.5 Separate, Onsite ADSM Server Machine

If the NT environment grows considerably or is mixed with additional and different platforms, a separate NT machine as a dedicated backup server might be necessary (see Figure 13 on page 29). The performance considerations for a stand-alone backup server and high availability are now important factors in the backup scenario. Site recovery depends on your ability to protect the ADSM server with suitable offsite copies of ADSM storage pools and the ADSM database.

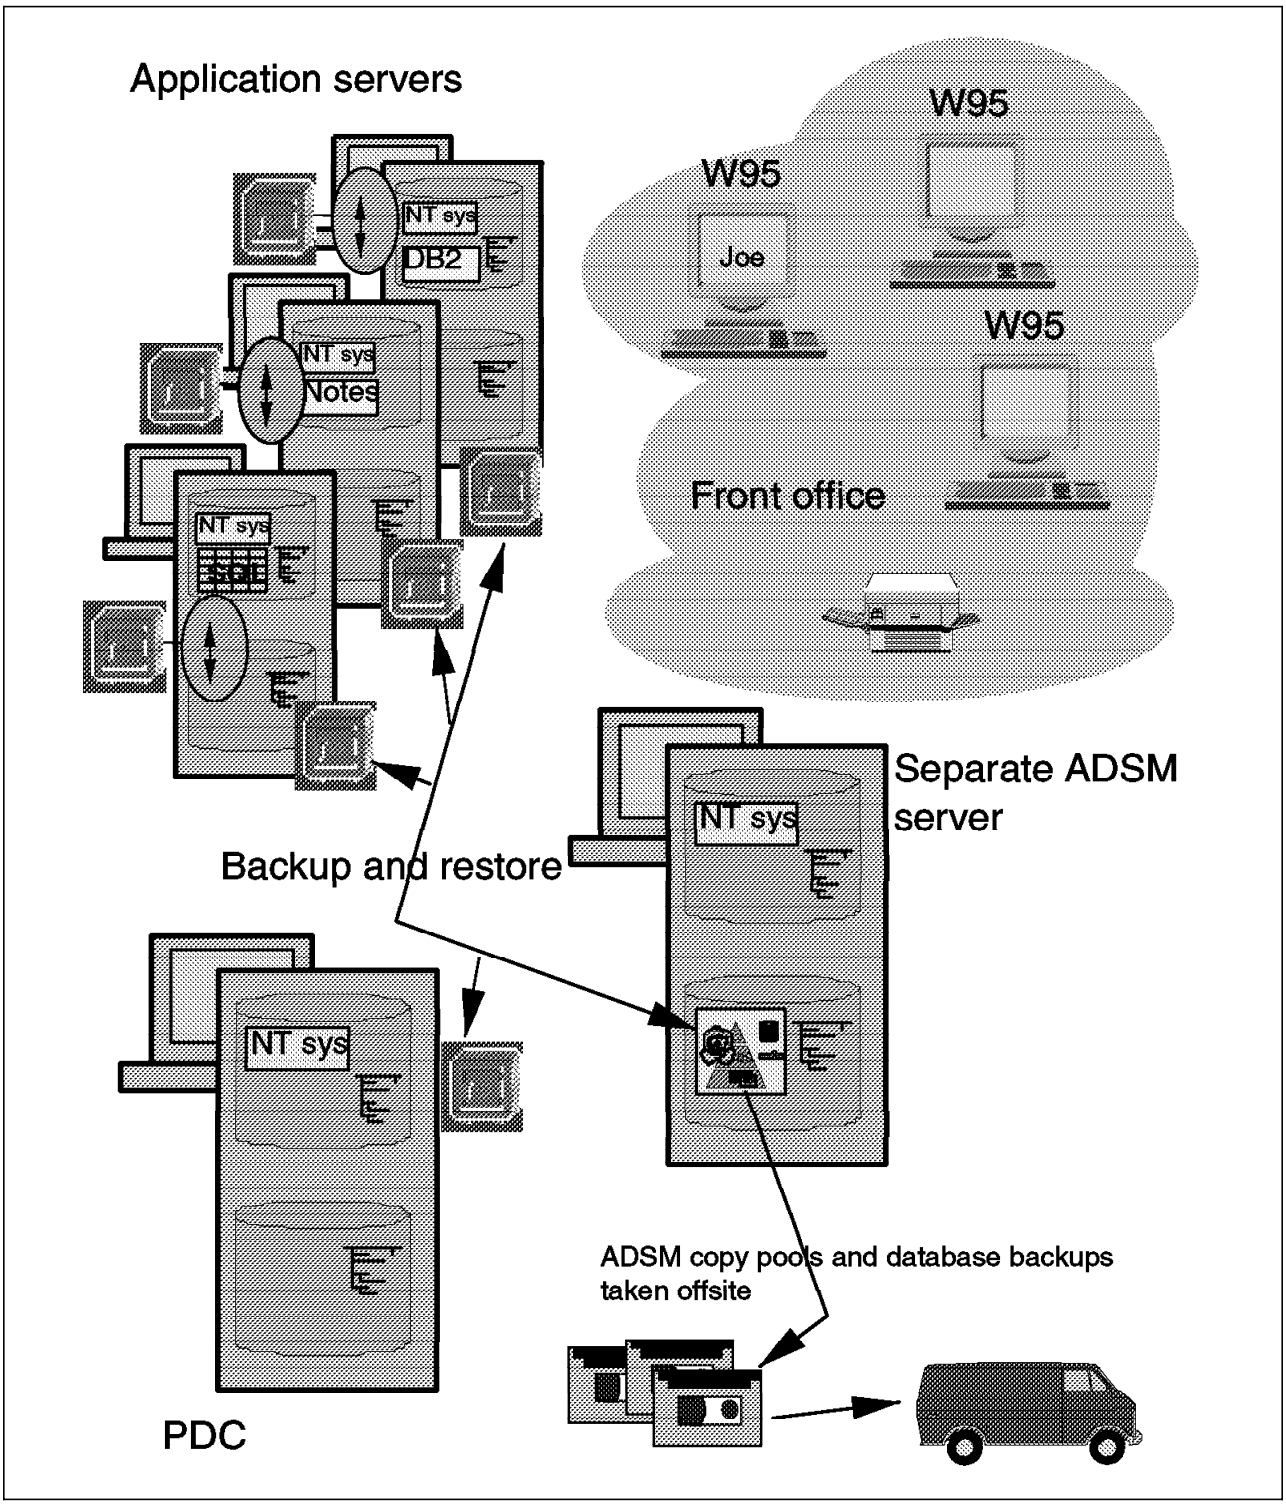

Figure 13. Separate, Onsite ADSM Server

For availability of your systems, refer to:

- 3.1, “NT Boot Process” on page 37
- 3.2, “Useful Availability Tools” on page 42 and the Microsoft Windows NT Resource Kit
- 3.3, “The NT File System” on page 50

- 3.4, “Configuring NT for Availability” on page 55 and the Microsoft Windows NT Resource Kit

  You will need to set up a repair partition, on either a separate hard disk or removable drive, for each NT system in the environment.

For recovery of your system, refer to:

- Chapter 4, “System Recovery with ADSM” on page 69
- 4.4.2, “Synchronization of the PDC and BDCs” on page 85
- Using the NT backup program for system recovery (see Microsoft NT system documentation)
- *ADSM Server for Windows NT Configuration and Recovery*, SG24-4878
- *ADSM Client Disaster Recovery*, SG24-4880

For the application recovery of your system, refer to:

- Chapter 5, “Recovering Lotus Notes” on page 95
- Chapter 6, “Microsoft SQL Server Backup” on page 103
- Chapter 7, “Recovering Microsoft Exchange Server” on page 125
- Chapter 8, “Recovering Microsoft Access” on page 145
- Chapter 9, “Recovering DB2 for NT” on page 151
- Application literature from Microsoft, Lotus, IBM, and other vendors.

2.6 Dedicated Remote ADSM Server Machine

If you have a fast wide area network or a bandwidth-on-demand link to a remote site, you can site the separate ADSM server remote from the data that is being backed up (see Figure 14 on page 31). The remote server must still be backed up in case of it′s failure but its backups can be local to it. The offsite disaster recovery capability is provided by the backup copies themselves being moved offsite.

Availability and recovery considerations are the same as those in 2.5, “Separate, Onsite ADSM Server Machine” on page 28, and are not repeated here. The primary consideration is the performance of the offsite link. It must be capable of containing the daily backup loads and fast enough to allow the recovery from a complete site loss, maybe to an alternative location. As part of your disaster recovery plan for your site, you should consider providing this capacity either in standby or as part of a contract with a business recovery service.

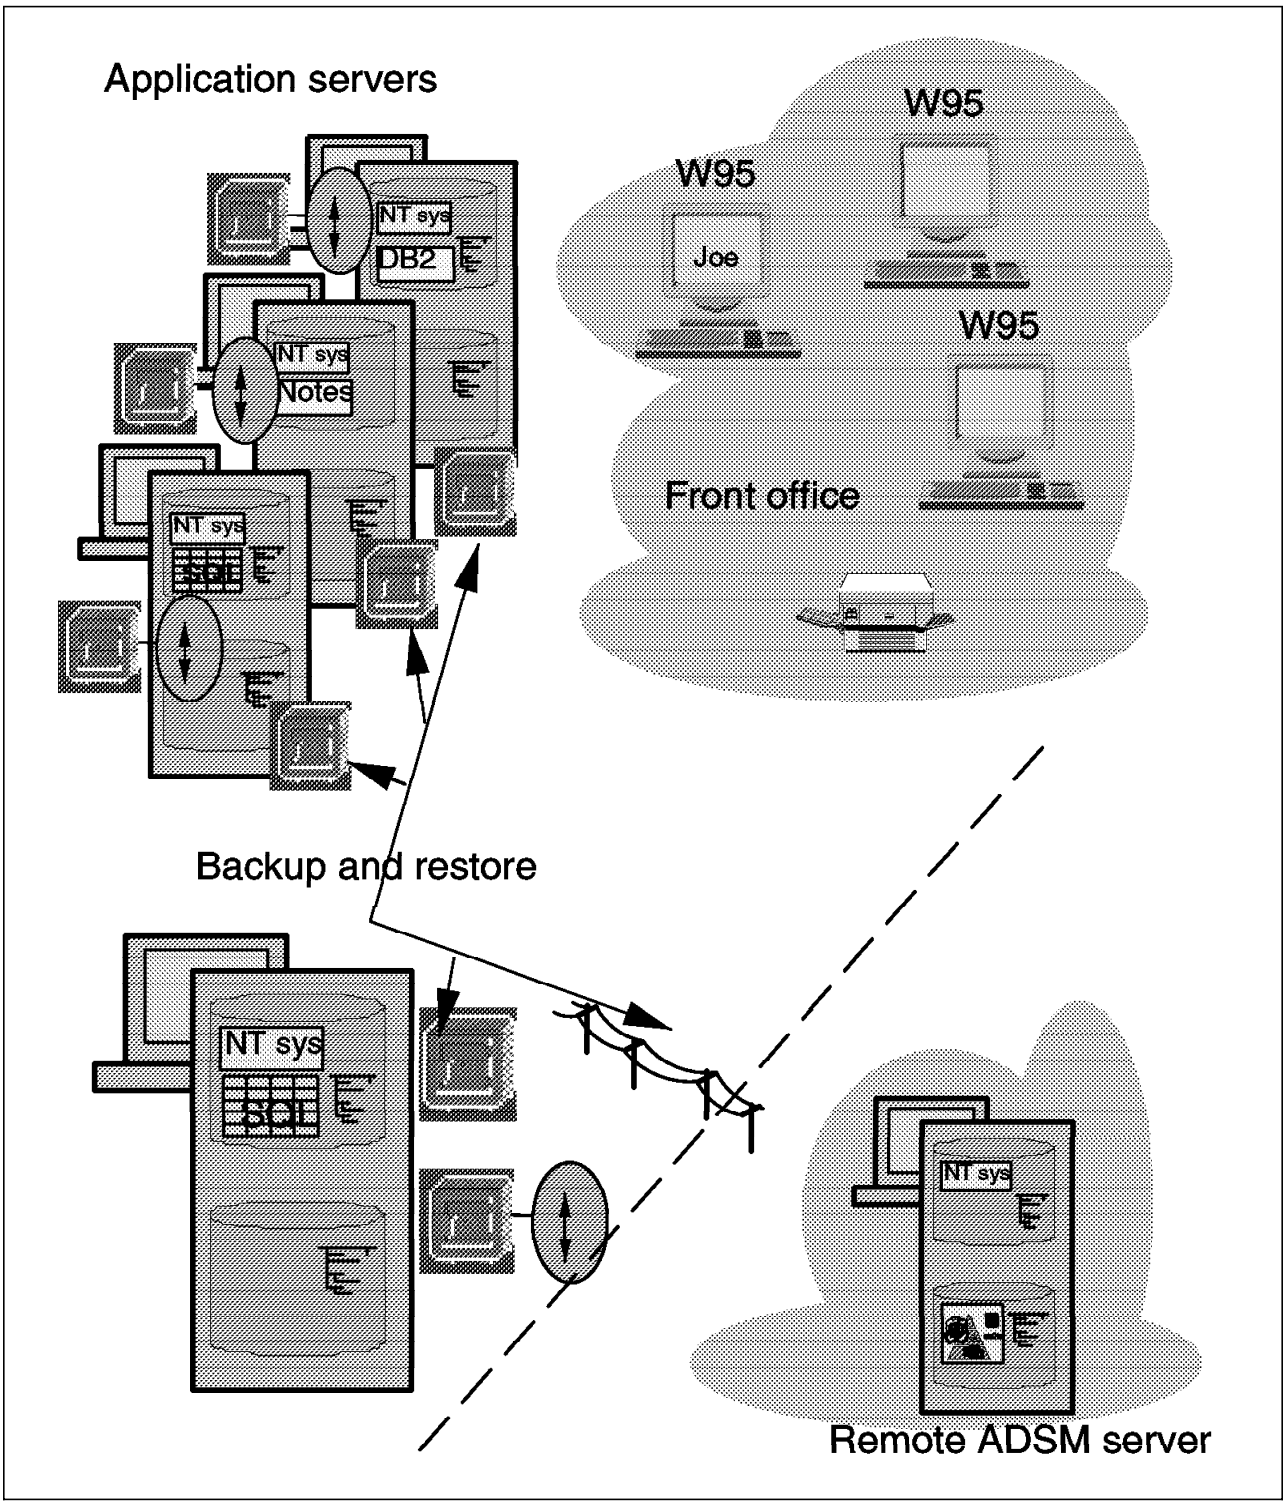

Figure 14. Separate, Offsite ADSM Server

2.7 Local ADSM Servers and Server-to-Server Communications

For high availability purposes you can use a second ADSM server as part of a highly available clustered multiprocessing (HACMP) solution. The second server sleeps and waits until a failure occurs.

If tight backup windows require a fast backup and offsite copies to be maintained automatically without tape handling, you may want to use the new feature of ADSM in Version 3, server-to-server Communications. This feature is described in detail in a new IBM Redbook to be available in mid 1998 - *ADSM Server-to-Server Implementation and Operation*, SG-24-5244, - shown in Figure 15 on page 33.

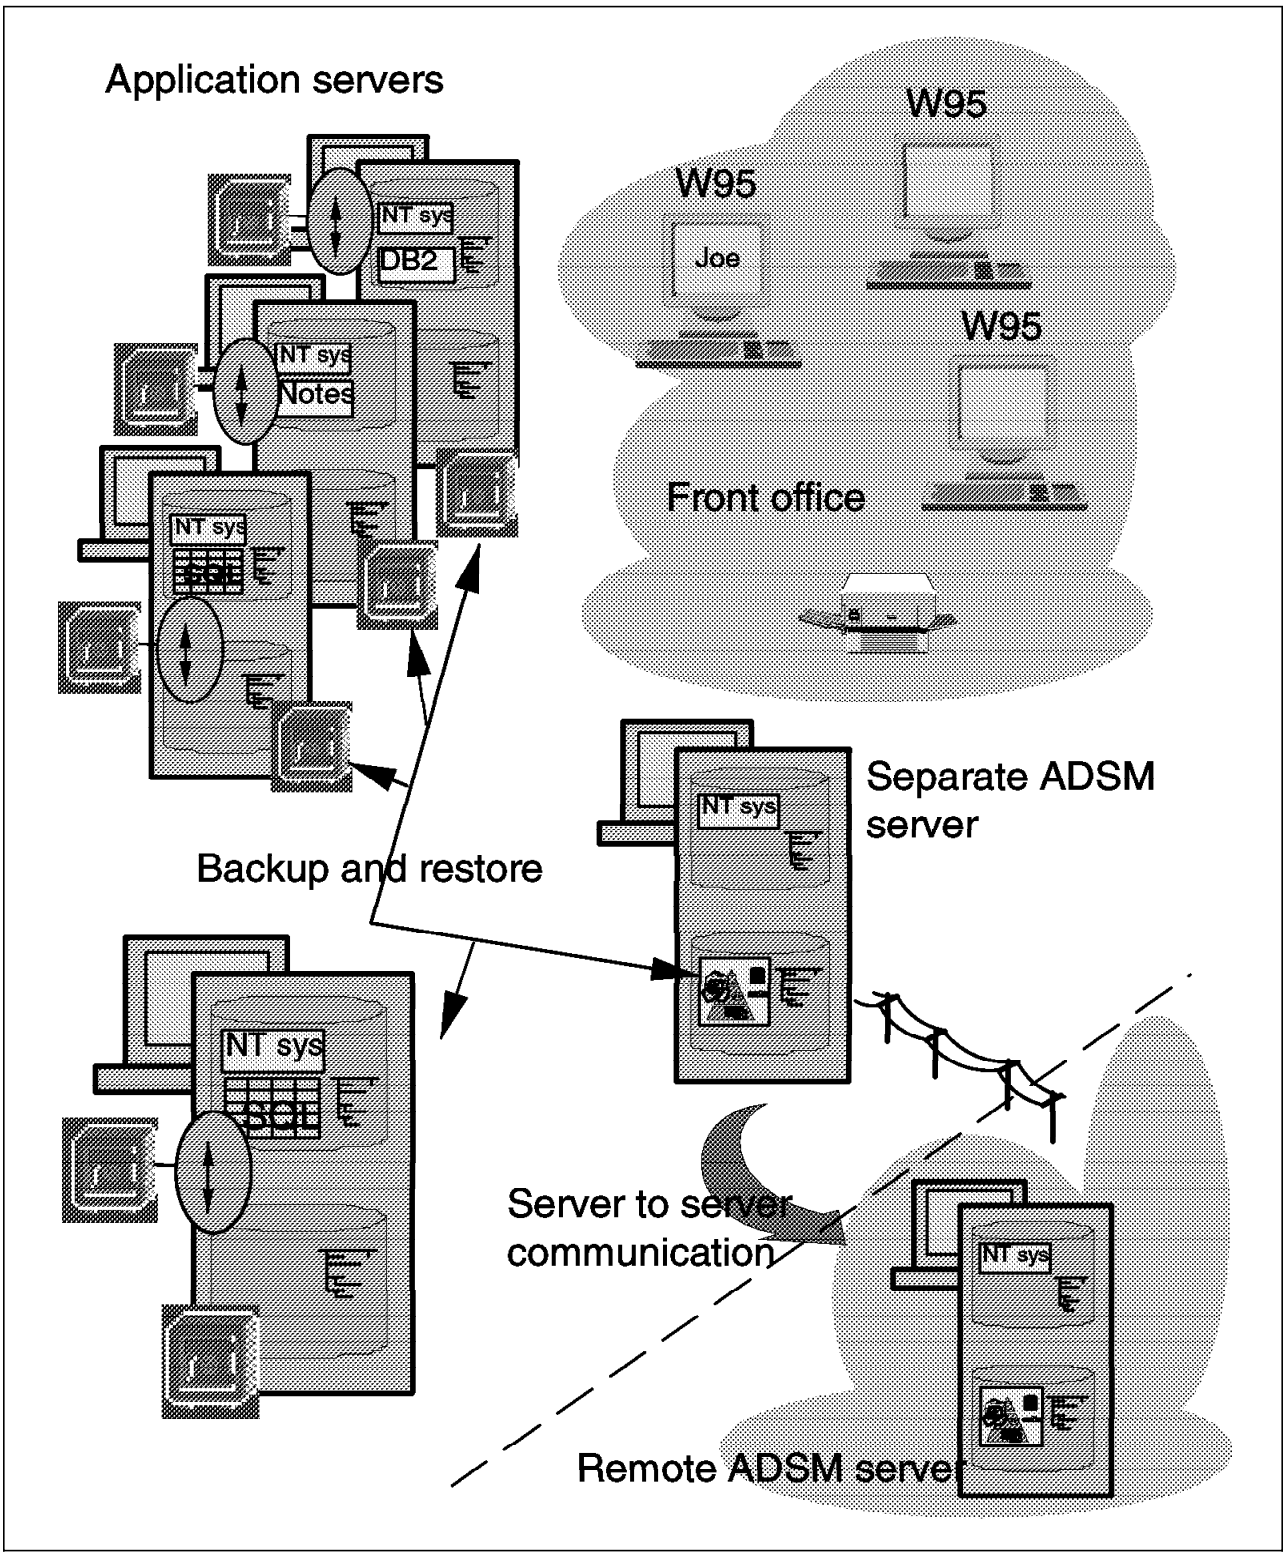

Figure 15. Server-to-Server Communications Used for Offsite Backup

For availability of your systems, refer to:

- 3.1, “NT Boot Process” on page 37
- 3.2, “Useful Availability Tools” on page 42 and the Microsoft Windows NT Resource Kit

- 3.3, “The NT File System” on page 50
- 3.4, “Configuring NT for Availability” on page 55 and the Microsoft Windows NT Resource Kit

  You will need to set up a repair partition, on either a separate hard disk or removable drive, for each NT system in the environment.

For recovery of your system refer to:

- Chapter 4, “System Recovery with ADSM” on page 69
- 4.4.2, “Synchronization of the PDC and BDCs” on page 85
- Using the NT backup program for system recovery (see the Microsoft NT system documentation)
- *ADSM Server for Windows NT Configuration and Recovery*, SG24-4878
- *ADSM Client Disaster Recovery*, SG24-4880
- Wolf Pack server clustering information from Microsoft
- *ADSM Server-to-Server Implementation and Operation*, SG-24-5244

For the application recovery of your system, refer to:

- Chapter 5, “Recovering Lotus Notes” on page 95
- Chapter 6, “Microsoft SQL Server Backup” on page 103
- Chapter 7, “Recovering Microsoft Exchange Server” on page 125
- Chapter 8, “Recovering Microsoft Access” on page 145
- Chapter 9, “Recovering DB2 for NT” on page 151
- Application literature from Microsoft, Lotus, IBM, and other vendors

Part 2. System Recovery

Chapter 3. NT System Availability

In this chapter we look at configuring Windows NT for availability and preparing for recovery with ADSM. We describe the Windows NT boot process, the NT file system (NTFS), and the configuration of an alternate NT system image. We provide examples of configuring a disk repair partition and configuring repair partitions on removable media.

3.1 NT Boot Process

The Windows NT boot process is a relatively straightforward two-stage process. The NT loader sets up the environment for loading an NT boot image. The NT loader is located in the *boot* partition. The boot partition is typically on a primary disk partition. However, it also can be on removable media such as floppy disks.

The NT system is loaded from the *system* partition. The system partition is typically the same partition as the boot partition; however, it can also be located on any fixed disk partition. Multiple system partitions can be created for a single PC. The BOOT.INI file definitions provide the link between the boot and system partitions. Thus multiple NT systems can be installed on a single PC.

When considering how to recover a Windows NT system, we first need to understand the NT boot process and its restrictions. When an NT system is started, it is loaded in two stages. First the NT loader is booted, and second, an NT system image is loaded (Figure 16).

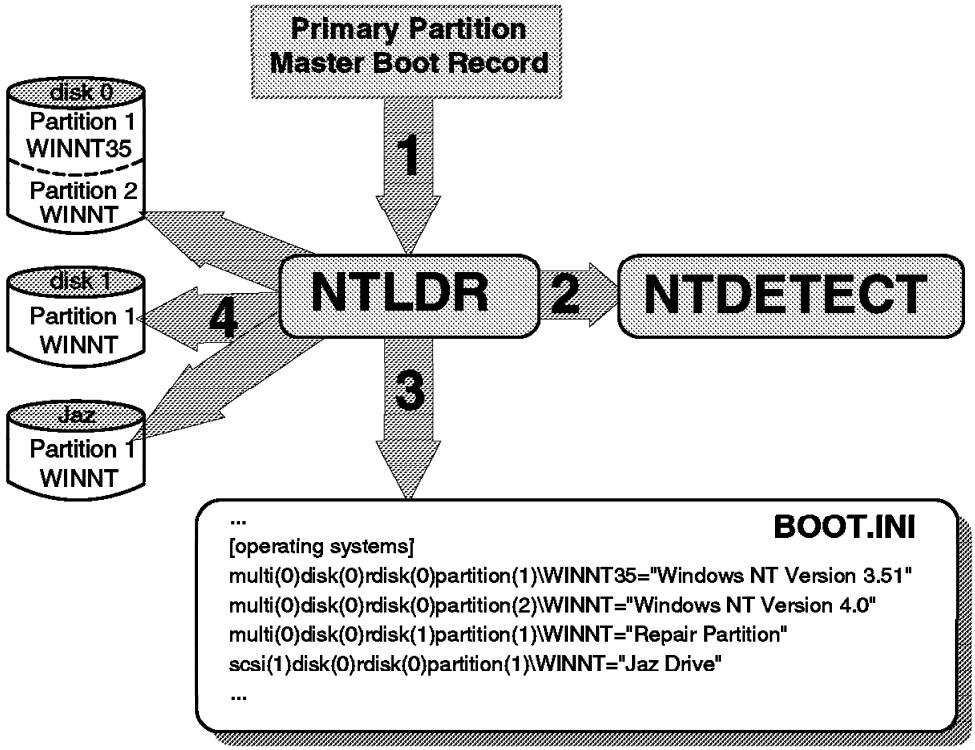

Figure 16. Windows NT Boot and Load Sequence

The Windows NT boot process consists of four steps:

1. The hardware basic input/output system (BIOS) reads the NT master boot record (MBR) located in the first 512-byte sector of the primary disk partition. The MBR contains the bootstrap code that starts the NT loader (NTLDR).
2. NTLDR invokes NTDETECT, which compiles a list of hardware configurations on the PC. The list is used during the NT load process and is stored in the registry when NT has been loaded.
3. NTLDR reads the BOOT.INI file (Figure 18 on page 40), which contains details of the loadable NT images and their physical locations. NTLDR displays a list of the NT systems, enabling you to choose which system to load. This step completes the boot stage of the process.
4. Having determined the hardware configuration, NTLDR attempts to load the system image selected in step 3. If the hardware configuration detected by NTDETECT matches the configuration being booted, Windows NT continues to load.

This Windows NT boot process uses two NT partitions: system and boot.

3.1.1 NT System Partition

The NT system partition is typically the active primary partition on the first physical disk in a system. It is the partition from which the hardware BIOS attempts to load. You must format the system partition by using NT with any of the supported Windows NT file systems: file allocation table (FAT), high performance file system (HPFS) (only Windows NT 3.x; NT 4.0 does not support HPFS), or NTFS. This formatting by NT creates the master boot record (MBR), which, when activated, initializes the NT loader.

The NT system partition must contain the following files in its root directory:

- NTLDR
- NTDETECT.COM
- NTBOOTDD.SYS
- BOOT.INI

These files are hidden system files and therefore are not usually visible with either File Manager or a dir command. The attrib command can be used to reset these attributes so that the files can be viewed. The functions of NTLDR and NTDETECT are explained in 3.1, “NT Boot Process” on page 37. BOOT.INI and NTBOOTDD.SYS are used to configure and access an NT boot partition. NTBOOTDD.SYS is an optional SCSI device driver used during the NT loading process. Its purpose is described in 3.1.3, “BOOT.INI” on page 40. The file is not necessary if a SCSI disk is not in use.

Useful tip

When NT is installed, it uses the first primary partition on the first physical disk as its system partition (typically the C: drive). Although this is a sensible configuration for day-to-day use, you can create an alternative system partition. NT can use any bootable media for a system partition, including floppy disks.

3.1.2 NT Boot Partition

An NT system image is located and loaded from an NT boot partition, typically the same partition used for the system partition. When NT is installed, the operating system code is installed in a predefined directory structure known to NT as the SystemRoot directory.

Figure 17, shows the directory structure of the C: drive. The SystemRoot directory in this example is Winnt. Having located this directory during the boot and load processes, the system continues to load. The operating system files are primarily located in the system32 directory and its subdirectories. User information, including the settings of the user environment, is in the Profiles directory. For each user who is locally logged on, there is a subdirectory with the same name as the user. The most important information, the registry, is located in the system32config subdirectory.

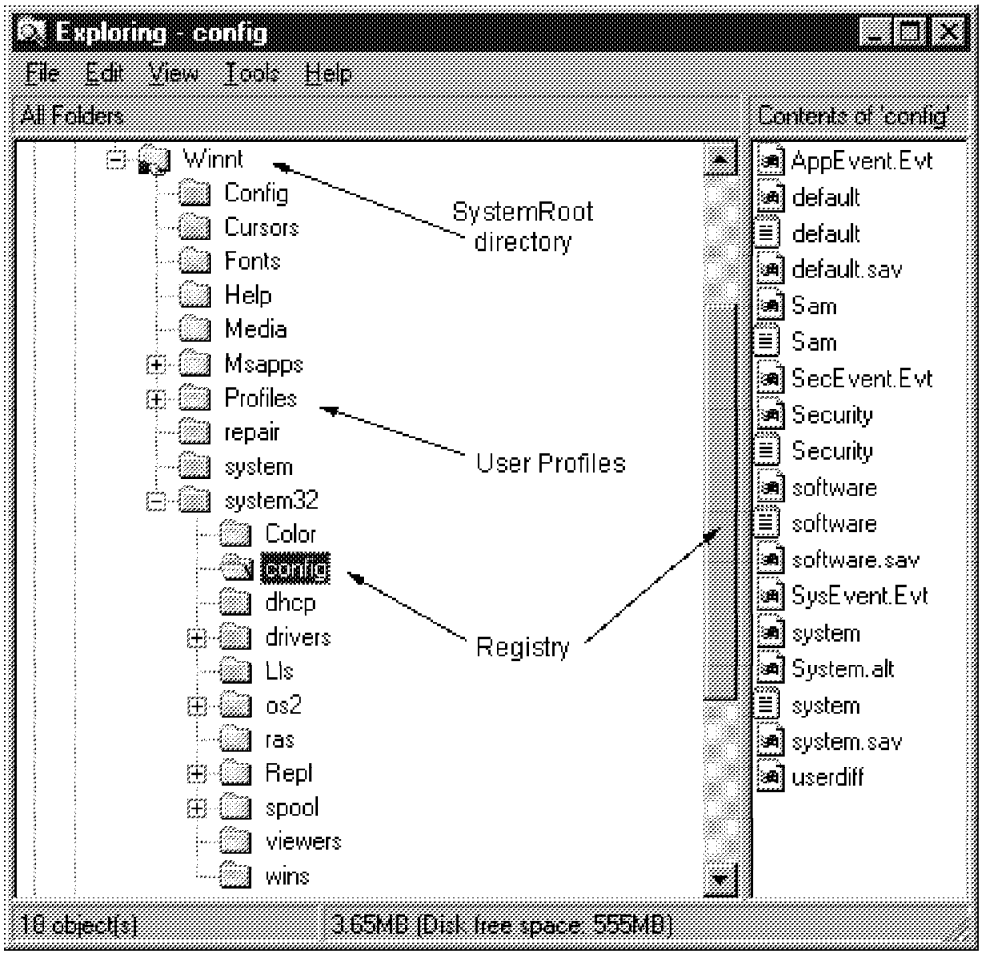

Figure 17. Windows NT System Partition

The directory structure must be on a direct-access fixed disk. NT cannot load from a floppy disk or sequential media such as tape. In a typical NT installation, all of the NT boot and system files are in the first primary partition on the first physical disk (usually the C: drive). In this case the boot and system partitions are the same.

The link between the NT boot and system partitions is the BOOT.INI file. This file provides the information required by NTLDR to locate a boot partition and the

SystemRoot directory structure. In 3.1.3, “BOOT.INI” on page 40 we look at the contents of this file.

Useful tip

The boot and system partitions are usually the same; however, they also can be located on separate disks and partitions. The difference in location enables great flexibility. A system partition can load NT from any number of boot partitions.

3.1.3 BOOT.INI

BOOT.INI is a text file containing details of NT boot partitions and the disk partitions where they are located. During the boot process, NTLDR displays some of the information in the BOOT.INI file. Thus the user can choose which NT system to load (assuming that more than one is installed).

The file consists of two parts. The [boot loader] part defines the default system image and a timeout value in seconds after which the default system is loaded. If the user moves the cursor to another image, the timer stops and an alternative system can be selected. The [operating systems] part defines the installed system images and their boot partition location (see Figure 18).

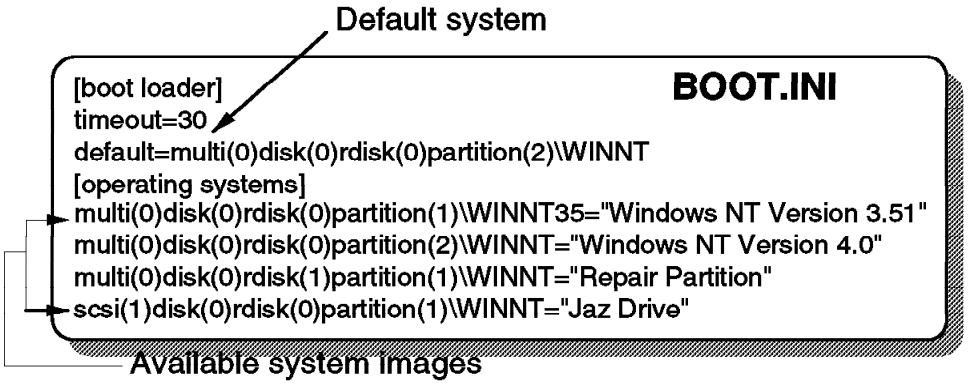

Figure 18. Sample BOOT.INI Contents

The structure of the BOOT.INI file is based on the Advanced RISC Computing (ARC) specifications. The specifications define adapters, disks, partitions, and paths to locate an operating system image. Each line in [operating systems] defines an NT boot partition and its physical location. Each system image is defined with one of the following definitions:

multi()disk()rdisk()partition()path=″Description″

or

scsi()disk()rdisk()partition()path=″Description″

The definition NT uses depends on the controller and device driver used to access the physical disks. The parameters in the two definitions have different meanings.

3.1.3.1 multi() Definitions

The NT installation process by default uses multi() definitions to define the location and path of the NT system image. The definition specifies that NTLDR should use the hardware BIOS to access the disk and load the NT system. This support is limited to the first disk controller in the system (either IDE or SCSI) and to the first two physical disks installed on that adapter (four with a dual channel IDE controller).

multi() This parameter tells NTLDR to use the hardware BIOS. The value is always 0, to represent the first disk controller.

disk() This parameter is not directly used with multi() definitions but must be included. It is always set to 0.

rdisk() This parameter is used to define the physical disk number attached to the controller. It is always set between 0 and 3.

partition() This parameter is used to define the partition number on the physical disk defined with the rdisk() parameter. It is the only parameter with a start number of 1. All other parameters start with 0.

path This parameter defines the SystemRoot directory from which the NT system image is loaded. The directory is typically WINNT35 (for an NT 3.51 system) or WINNT (for an NT 4.0 system). All of the system files required to load an NT system are in this directory and its subdirectories. This is the NT boot partition.

"Description" This parameter indicates the system to be loaded. It is displayed by NTLDR when prompting the user to choose a system to load.

The following two examples of multi() definitions are taken from Figure 16 on page 37:

```
multi(0)disk(0)rdisk(0)partition(1)\WINNT35="Windows NT Server 3.51"
multi(0)disk(0)rdisk(0)partition(2)\WINNT="Windows NT Server 4.0"
```

The first example is for an NT Server 3.51 system. It is located in the Winnt35 directory in the first partition—a primary partition in this case. This partition is on the first disk, rdisk(0), attached to the first disk controller, multi(0). The second example is for an NT Server 4.0 system located in the Winnt directory but this time in the second partition, a logical volume on the same physical disk as the previous definitions.

Note: These examples are actually for an IDE adapter. Because the IDE adapter is the first adapter in the system supported by the hardware BIOS, a multi() definition has been used. The NT installation process uses multi() definitions for the first adapter, whether it is IDE or SCSI.

3.1.3.2 scsi() Definitions

A scsi() definition differs from multi() in that, rather than relying on the hardware BIOS, it loads and uses a SCSI device driver to access the disks. This driver, NTBOOTDD.SYS, is created in the system partition when NT is installed on a SCSI-attached disk. This support is not limited to the first two disks on the first controller as with the multi() definitions. A scsi() definition can be used with any SCSI adapter and disk supported by Windows NT. The parameters for scsi() definitions, while having the same names, have different meanings.

scsi() This parameter tells NTLDR to use the NTBOOTDD.SYS device driver to access the disk. The value of scsi() is the SCSI adapter number: 0 represents the first SCSI controller, 1 represents the second, and so on. The actual numbering and ordering of the adapters depend on the individual adapter and NTBOOTDD.SYS.

disk() This parameter defines the SCSI ID of the physical disk.

rdisk() This parameter defines the SCSI logical unit number (LUN) of the disk. In most cases the LUN is 0.

partition() This parameter defines the partition number on the physical disk defined with the disk() parameter.

path This parameter is the same as the \path parameter in a multi() definition. It defines the directory containing the NT system image.

Here is an example (taken from Figure 16 on page 37) of a scsi() definition:

```
scsi(1)disk(0)rdisk(0)partition(1)\WINNT="Jaz Drive"
```

The example is for an NT Server 4.0 system, described as a *repair partition*. It is accessed using the second adapter, a SCSI adapter, in the system (scsi(1)). The disk uses a SCSI ID of 0 and an LUN of 0 (disk(0) and rdisk(0)). The path to the NT system image again is the Winnt directory located in the first partition (partition(1)) on the disk.

In this scenario the NTBOOTDD.SYS device driver is used to access the adapter and disk. The NT install process creates this driver in the system partition during installation. The driver is not a generic SCSI device driver; it is a copy of the device driver for the adapter being used. In the example above, NTBOOTDD.SYS is a copy of the Microchannel SCSI adapter driver, SPOCK.SYS. If an Adaptec 1540 SCSI controller were used, NTBOOTDD.SYS would be a copy of the AHA154X.SYS device driver. This is only visible through the file size of the NTBOOTDD.SYS.

Note: The SCSI ID in this example is actually a logical ID used by the NT device driver for PS/2 Microchannel SCSI adapters. The real SCSI ID of the disk is 6, but NT interprets it as a logical ID of 0. This behavior is common to all SCSI adapters.

3.2 Useful Availability Tools

In this section we describe how you can increase the availability of your NT server by using some helpful disk tools and practices:

For an NT server installed on a FAT file system, we show how you can prepare a DOS boot diskette.

The emergency recovery diskette is a helpful utility to repair the NT system files or the registry.

With the DISKSAVE.EXE program and the NT Disk Administrator you can save the MBR and the boot sector of the physical disks. Both tools can save and restore the sectors needed to boot the machine, and both can initialize the first step of a full restore of the machine. You do not need to use these partition-saving tools if the machine uses only one partition per physical disk, because the complimentary metal oxide semiconductor (CMOS) part of the hardware BIOS includes all information about the drives.

3.2.1 Bootable DOS Diskette

In some cases it is right to install an NT system on a FAT file system. The availability is higher than that of an NTFS system because you can boot from a normal DOS diskette and edit, delete, and copy files on the production server. The FAT file system is also a little bit faster than NTFS because the security overhead of the latter does not exist. However, even though the NT system runs on FAT, you can protect it with a security mechanism. The main difference is that you cannot set permissions for locally logged on users. If you have locally logged on users and you need permissions for some files or directories, FAT cannot hold your needs. If you install the NT server as file server and put the machine in a secure room, there is no limitation. You can set permissions on the shared resources on FAT because the permission set for shared resources is part of the operating system and not part of the file system.

A bootable DOS diskette could help in case of any problems with the NT server. Format a diskette as a boot diskette with any DOS operating system. We recommend that you add the following programs as a minimum on that diskette:

- FORMAT.EXE
- FDISK.EXE
- DISKSAVE.EXE (see 3.2.3, “DISKSAVE and the NT Disk Administrator” on page 44)
- DOS editor
- SCANDISK.EXE or CHKDSK.EXE
- Virus scanner and cleaner

Additional programs from other vendors could also be helpful.

3.2.2 Emergency Recovery Diskette

The emergency recovery diskette (ERD) is first mentioned during the NT setup procedure. The NT system saves all system files, such as the registry, on one diskette.

With the RDISK.EXE program you can create new ERD diskettes or update an older version each time you install or change important parts of the environment.

The diskette contains the following files:

```
 Volume in drive A is ERD
 Volume Serial Number is E0B6-AFEB
 Directory of A:\

SETUP    LOG        49,616  04-10-97 10:53a SETUP.LOG
SYSTEM   _          80,835  04-10-97 10:54a SYSTEM._
SOFTWARE _         124,700  04-10-97 10:54a SOFTWARE._
SECURITY _           4,725  04-10-97 10:55a SECURITY._
SAM      _           3,827  04-10-97 10:55a SAM._
DEFAULT  _          17,662  04-10-97 10:55a DEFAULT._
NTUSER   DA_        14,676  04-10-97 10:55a NTUSER.DA_
AUTOEXEC NT            438  11-17-96  6:38p AUTOEXEC.NT
CONFIG   NT          2,510  04-10-97  3:30a CONFIG.NT
         9 file(s)        298,989 bytes
         0 dir(s)       1,157,120 bytes free
```

Those files that end with an underline sign (_) are compressed during the save process. Use the DECOMP.EXE program to decompress them.

The SETUP.LOG file includes all file names of the NT operating system with a checksum (CRC). Changes in the sizes of these files will be detected.

Important

The RDISK.EXE program copies only the uncompressed files to the `Winnt\repair` directory and then to the diskette. It does not copy the current registry. You must start RDISK.EXE with the /S parameter to copy the registry to the `Winnt\repair` directory and the diskette. Or you can click on **Update** in the ERD program.

Use the ERD diskette for any problems with the NT boot process. Insert the normal install boot diskette to start the system from diskette. The setup program asks for the second diskette, and the menu for the normal installation appears. If you select **R** for repair, NT asks for the ERD diskette and tries to restore the files to the production partition. The entries in the SETUP.LOG file with the CRC must be comparable to the files on the disk. Any differences result in an error message. NT then installs a fresh copy of the files from the install media, the CD-ROM. In most cases the problems disappear after a reboot of the system.

3.2.3 DISKSAVE and the NT Disk Administrator

For the restore of an NT machine, information about the primary and logical partitions—including the size and the drive letter assigned to them—is known as a partition table. The table is stored on the physical disk in the first sector (sector 0) with a length of 512 bytes. The MBR contains code that the hardware BIOS of an x86 machine uses to read the partition table and locate the boot sector of any operating system. If the MBR is damaged, the system is no longer startable and displays a blank screen after the POST routine or messages such as “invalid partition table.” The MBR and the boot sector can also reside on a diskette, which opens up the possibility of booting from a rescue partition.

The boot sector contains the code for loading the operating system or a multiboot program such as OS/2 Boot Manager. A damaged boot sector also prevents the system from booting. Error messages are issued to provide further information about the problem.

If you are using more than one partition per physical disk, there will be one MBR per logical partition. The boot sector is now divided into one partition boot sector per partition. It is possible to save and restore the partition table for each partition. The most important partition is the startup partition, because if that information is lost, a boot is not possible. All other partitions can be rebuilt with the information stored in the operating system.

Back up the MBR to a disk every time you change partition information for primary partitions or an extended partition. You should also back up a partition boot sector when you format a volume, install Windows NT on the volume, or convert a volume from FAT to NTFS.

Useful tip

If the MBR or the boot sector seems damaged, it is possible that a virus destroyed the information in them. With the restore of a clean copy of the MBR and the boot sector, all virus code is deleted, and the original sectors are rebuilt. Be careful to use only the correct copy of these critical sectors. Restoring the wrong sectors will damage the system.

Both Microsoft DISKSAVE and the NT Disk Administrator are suitable for saving and restoring the partition table. In case of a primary partition loss without any additional partitions, DISKSAVE will be the better program to use, because it can be started from a bootable DOS diskette. The NT Disk Administrator program can start only from a running NT system.

The DISKSAVE.EXE program is part of the Microsoft Windows NT Resource Kit. The tool reads the important sector information and saves it as binary image files. In case of any trouble with the MBR or the boot sector, the information can be restored with the tool.

The program is located on the Resource Kit CD in the i386\FAULTTOL directory, or you can download it from Microsoft's ftp server. Additional information is available in the readme file or in the Microsoft Developer Network (MSDN).

Note: You can subscribe to the MSDN free of charge, at http://www.microsoft.com/msdn/subscribe/online.htm

This program will not run on a Windows NT system. Boot from a DOS system diskette and then start the program.

If you do not have a repair partition or any other partition on your NT machine, do not use the NT Disk Administrator because it can be started only from a running NT system. If there is trouble with the partition table or a boot sector, there is no way to start this program from the normal production partition. In this case the DISKSAVE program is the better choice.

The NT Disk Administrator enables you to create, change, or delete all partition and disk information. For this work you must be the local administrator of the NT machine.

In the NT Disk Administration **Partition/Configuration** menu there is a sub-menu for saving or restoring the partition information. The principle is the same as that of the DISKSAVE program. The NT Disk Administrator saves all partition information as a binary file on a diskette.

In the next section we show how you use the NT Disk Administrator to create additional availability functions for the disk environment.

Not for repairing disk information, but as additional information source you can use the SHOWDISK.EXE program included in the Windows NT Resource Kit. The output shows the mapping of the drive letters to the physical partition and the physical sectors that hold the data. You can print out this important information, or you can redirect the output to a file, which enables you to compare this file with a file created after the system restore. Incompatible drive letter assignments and sector sizes become apparent when you compare the two files.

Here is the sample output of a machine with two physical disks and some partitions:

```
Opening \SYSTEM\DISK successful

Disk Registry Information Size........ 272
Operating System Version.............. 3
Checksum.............................. 0x0
Dirty Shutdown?....................... 0
<reserved 1>.......................... 0x0
<reserved 2>.......................... 0x0
<reserved 3>.......................... 0x0
Disk Info Offset...................... 0x2c
Disk Info Size........................ 228
FT Info Offset........................ 0x110
FT Info Size.......................... 0
FT Stripe Width....................... 0
FT Pool Size.......................... 0
Name Offset <not implemented>......... 0x0
Name Size <not implemented>........... 0

General Disk Information:
Number of Disks....................... 2
<reserved>............................ 0

 Disk #0:
      Number Of Partitions.......... 1
      <reserved>.................... 0x0
      Signature..................... 0x395fe6c4

      Partition #1:
           FT Type............... Not a Fault Tolerance Partition
           FT State.............. Healthy
           Starting Offset....... 0x7e00
           Length................ 1707213312
           FtLength.............. 0
           <reserved 1>.......... 0x0
           <reserved 2>.......... 0x0
           Drive Letter.......... C
           Assign Drive Letter?.. Yes
           Logical Number........ 1
           Ft Group.............. Not an FT Partition
           Modified?............. Yes
           <reserved char 0>..... 0x0
           <reserved char 1>..... 0x0
           <reserved char 2>..... 0x0
```

```
Disk #1:
     Number Of Partitions.......... 3
     <reserved>.................... 0x0
     Signature..................... 0x395fe6c5

     Partition #1:
          FT Type............... Not a Fault Tolerance Partition
          FT State.............. Healthy
          Starting Offset....... 0x1ffe00
          Length................ 850493952
          FtLength.............. 0
          <reserved 1>.......... 0x0
          <reserved 2>.......... 0x0
          Drive Letter.......... D
          Assign Drive Letter?.. Yes
          Logical Number........ 1
          Ft Group.............. Not an FT Partition
          Modified?............. Yes
          <reserved char 0>..... 0x0
          <reserved char 1>..... 0x0
          <reserved char 2>..... 0x0
     Partition #2:
          FT Type............... Not a Fault Tolerance Partition
          FT State.............. Healthy
          Starting Offset....... 0x32d1fe00
          Length................ 524321280
          FtLength.............. 0
          <reserved 1>.......... 0x0
          <reserved 2>.......... 0x0
          Drive Letter.......... E
          Assign Drive Letter?.. Yes
          Logical Number........ 2
          Ft Group.............. Not an FT Partition
          Modified?............. Yes
          <reserved char 0>..... 0x0
          <reserved char 1>..... 0x0
          <reserved char 2>..... 0x0
     Partition #3:
          FT Type............... Not a Fault Tolerance Partition
          FT State.............. Healthy
          Starting Offset....... 0x5212fe00
          Length................ 315818496
          FtLength.............. 0
          <reserved 1>.......... 0x0
          <reserved 2>.......... 0x0
          Drive Letter.......... F
          Assign Drive Letter?.. Yes
          Logical Number........ 3
          Ft Group.............. Not an FT Partition
          Modified?............. Yes
          <reserved char 0>..... 0x0
          <reserved char 1>..... 0x0
          <reserved char 2>..... 0x0
```

The last useful disk utility we want to mention is the DISKMAP.EXE program for displaying the physical information about a specified disk and the location of the MBR and the extended boot records (EBRs):

```
Cylinders  HeadsPerCylinder SectorsPerHead BytesPerSector MediaType
      820               64             63            512      12
TrackSize = 32256, CylinderSize = 2064384, DiskSize = 1692794880 (1614MB)

Signature = 0x395fe6c5
     StartingOffset    PartitionLength StartingSector PartitionNumber
            2096640          850493952            63               1
          852622848          524321280            63               2
         1376976384          315818496            63               3

MBR:
         Starting              Ending       System    Relative    Total
  Cylinder Head Sector  Cylinder Head Sector   ID      Sector      Sectors
         1    0    1        1023   15   63     0x05        1008    3305232
         0    0    0           0    0    0     0x00           0          0
         0    0    0           0    0    0     0x00           0          0
         0    0    0           0    0    0     0x00           0          0

EBR: (sector 1008)
         Starting              Ending       System    Relative    Total
  Cylinder Head Sector  Cylinder Head Sector   ID      Sector      Sectors
         1    0    1         412   63   63     0x05        3024    1661184
         0    0    0           0    0    0     0x00           0          0
         0    0    0           0    0    0     0x00           0          0
         0    0    0           0    0    0     0x00           0          0

EBR: (sector 4032)
         Starting              Ending       System    Relative    Total
  Cylinder Head Sector  Cylinder Head Sector   ID      Sector      Sectors
         1    1    1         412   63   63     0x07          63    1661121
       413    0    1         666   63   63     0x05     1664208    1024128
         0    0    0           0    0    0     0x00           0          0
         0    0    0           0    0    0     0x00           0          0

EBR: (sector 1665216)
         Starting              Ending       System    Relative    Total
  Cylinder Head Sector  Cylinder Head Sector   ID      Sector      Sectors
       413    1    1         666   63   63     0x07          63    1024065
       667    0    1         819   63   63     0x05     2688336     616896
         0    0    0           0    0    0     0x00           0          0
         0    0    0           0    0    0     0x00           0          0

EBR: (sector 2689344)
         Starting              Ending       System    Relative    Total
  Cylinder Head Sector  Cylinder Head Sector   ID      Sector      Sectors
       667    1    1         819   63   63     0x07          63     616833
         0    0    0           0    0    0     0x00           0          0
         0    0    0           0    0    0     0x00           0          0
         0    0    0           0    0    0     0x00           0          0
```

3.2.4 Additional Disk Protection

Although our intention is not to discuss programs that protect against physical disk failures, we do discuss NT availability and therefore want to mention some hardware and software enhancements for NT disk usage. Many books and the Windows NT Resource Kit contain detailed information about this topic.

Using the NT Disk Administrator, you can install some native NT features to protect against disk failures.

3.2.4.1 Disk Mirroring

Disk mirroring is used most frequently because it is powerful and easy to use. Every operation to disk is directed to two different disks in the system. If you write a file to disk, the file is written twice, once on each disk. Such mirroring causes a performance degradation for all disk operations. Mirroring also requires additional space on the physical disk, which is another disadvantage, because you need a second disk space with the same size. In Figure 19 you can see how the mirrored disk is embedded in the disk environment. The NT system uses the same disk size and the same drive letter for the production partition and the mirror partition. If the production partition is damaged, you can rebuild it with the mirror. This is covered well in standard Microsoft NT documentation.

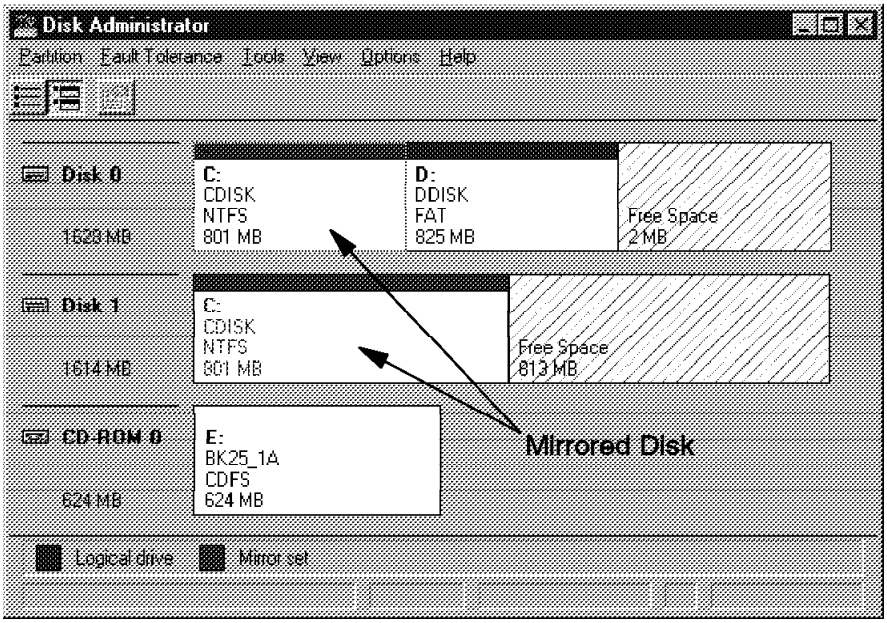

Figure 19. Disk Mirroring in Windows NT

You can compensate for the reduced performance if your machine has a second disk adapter and a second disk. In this case the information is sent to two adapters, and they are responsible for the file operations.

3.2.4.2 Stripe Sets

A stripe set is based on nearly the same principle as RAID, but it is implemented as a software solution. The amount of disk space is the same as that required for a hardware RAID system.

Note: For all software enhancements, refer to the Windows NT or Windows NT Resource Kit documentation. For hardware RAID information, refer to the manufacturer's documentation.

The enhancements discussed above protect against physical disk failures only. They cannot be used to restore errors or deletions of files. Defects in the registry or in the partition table are also not protected.

3.3 The NT File System

As a secure network server, NT includes many mechanisms for user permissions, auditing events, and object protection. These mechanisms are available only for the NTFS; a FAT file system cannot store the additional information.

Any backup program, such as ADSM, must have as a minimum read permission for every file. Knowledge of the permissions for files and directories is useful for understanding why some files cannot be backed up.

In this section we look at the NT security mechanisms, special and standard permissions for files, and the permissions for directories.

3.3.1 Security Mechanisms

A Windows NT machine with a FAT file system has four DOS attributes: S (system); H (hidden); R (read only); and A (archive). The NTFS has the same DOS attributes as well as attributes for file system security provided through the discretionary access control list (DACL). By default NT creates a file or a directory with Full Control (All). This is an important feature for ease of operation in a peer-to-peer network environment. If files or directories include information that all users should not be able to view or read, the NT mechanism of setting permissions can be used. Setting permissions on files is a way of providing security for files other than the normal DOS attribute of read-only. A permission for a file or a directory is set by the owner of the file and can be changed only by the owner. You can set permissions on one file, several files, or a whole directory. The NT security information indicates which user is allowed to access or read a file and which user is the owner of a file.

For NT all users are numbers. If the NT administrator creates a new user, the NT system creates a unique number through a hashing algorithm—the security identifier (SID; see Figure 20). Some SIDs are marked as “well-known” and therefore cannot be deleted. The well-known SIDs can be the administrators of the NT machines or other users.

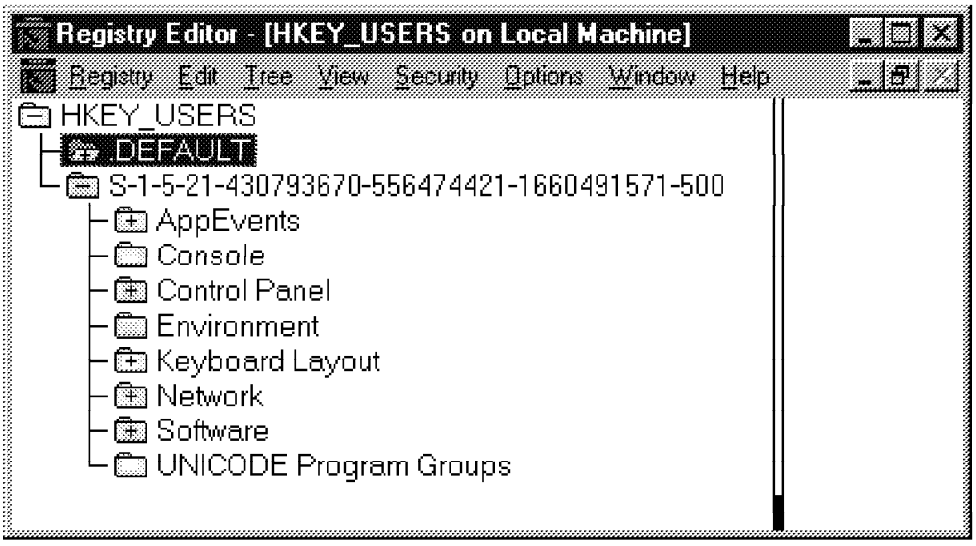

Figure 20. Registry Entry for Security Identifier

Because of the hashing algorithm, there will never be two SIDs that are the same. If you create a new user, grant the user some resources, and delete the user, it is not possible to create the same user with the same SID. When you delete a user, NT security issues a warning:

Each user account is represented by a unique identifier which is independent of the username. Once this user account is deleted, even creating an identically named user account in the future will not restore access to resources which currently name this user account in the access control list.

Every object in the NT environment with a name (such as files, printers, users, or devices) has its own access control list (ACL). Each object's ACL includes access control elements (ACEs). An ACL describes the operations that are allowed on objects and therefore differs from a DACL.

Note: ACLs are often called system ACLs.

One user or one group of users represents one ACE for one object. If the owner of a printer grants printing rights for a new user, the ACL of the printer adds a new ACE entry with the SID of the new user. The new user can then print the object, in this case, the printer.

There are three ACEs in NT:

- Access allowed
- Access denied
- System audit

Access allowed represents the permissions granted to one user or a group of users for the object within the ACL.

Access denied represents the permissions denied to one user or a group of users.

System audit is a logging mechanism that registers events for the object.

Note: The user or group of users is not mentioned by name, because in the ACE only SIDs are allowed.

3.3.2 Special and Standard Permissions for Files

You must distinguish between special and standard permissions for files. Special permissions describe in detail which rights the user has for one or more files. Standard permissions are combinations of special permissions, and they facilitate the administrative work.

The special permissions are:

R - read

W - write

X - execute

D - delete

P - change permission

O - take-ownership permission

To set or delete permissions for files in the Windows Explorer, select the file, using the right mouse button. Then select **Properties** and click on **permissions** on the security folder.

You can set the attributes if you are the owner of the file or if you have administrative rights. In all other cases, the following message appears on the screen:

```
you only have permission to view the
current security information on ...
```

Note: Special file permissions are comparable to UNIX file permissions with different names and separation for both user or group permissions.

Figure 21 shows the NT file permissions and the OEMNXPDL.INF file. In the Type of Access field you can select **Read**, **Change**, **Full Control**, **No Access**, or **Special Access** for special permissions.

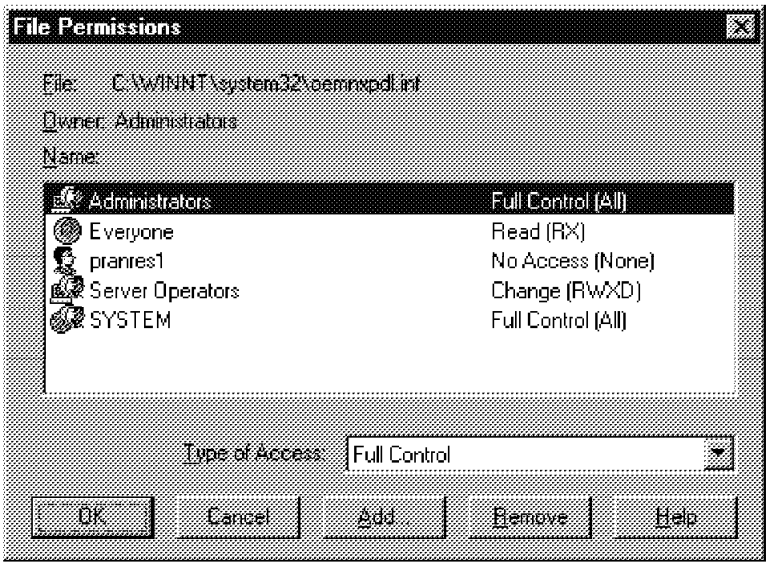

Figure 21. NT File Permissions

If you select **Special Access**, the Special Access window appears (Figure 22 on page 53), showing the special permissions set for the user. Click the check boxes to add or delete rights for the user.

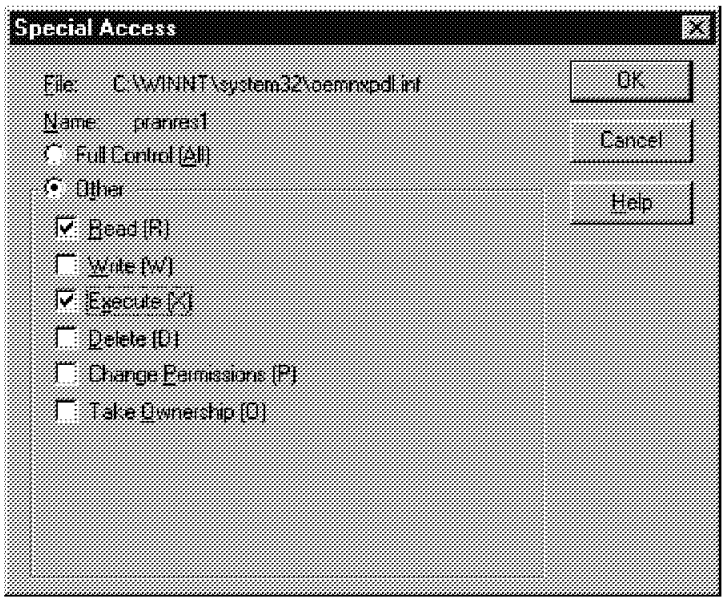

Figure 22. NT Special Permissions

3.3.3 Permissions for Directories

There are only standard permissions for directories. It is difficult to separate the permissions for directories from the permissions for files in the directory. If you grant a user read access for a directory, you must also grant read access for the files and subdirectories in that directory.

The directory permissions are:

(RX)() List

(RX)(RX) Read

(WX)() Add and no access

(RWX)(RX) Add and read

(RWXD)(RWXD) Change

(RWXDPO)(RWXDPO) Full control

()() No access

The process of setting or deleting permissions in directories is similar to the process of setting or deleting permissions in files. In the Windows Explorer, click with the right mouse button on a directory name, select **Properties**, and click on **permissions** on the security folder. The Directory Permissions window appears (Figure 23 on page 54).

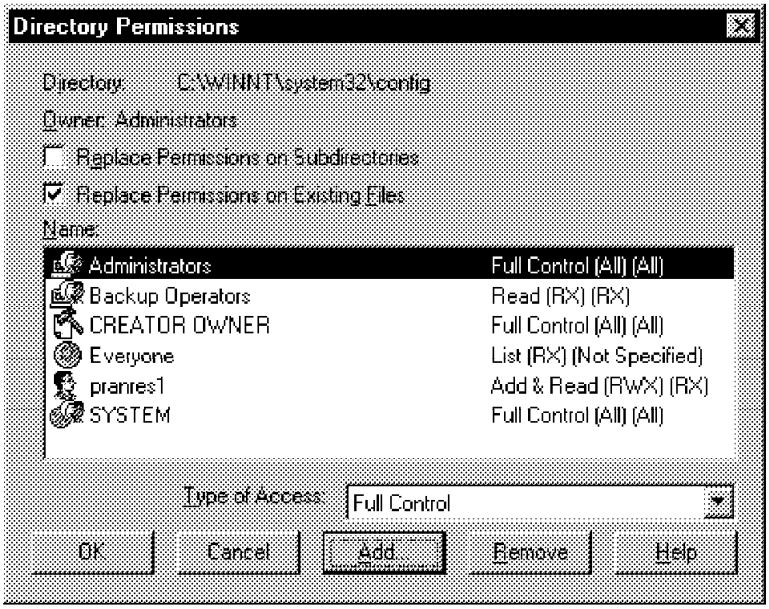

Figure 23. NT Directory Permissions

The dialog windows are a little different for directories than for files. One more entry appears for the file permission setting in the selected directory.

When you click on the **Replace Permissions on Subdirectories** checkbox in the Directory Permissions window (Figure 23), all files in all subdirectories will have the same permissions. The permissions change only for those files that you own.

Figure 24 shows the special directory access permissions for the Winnt directory.

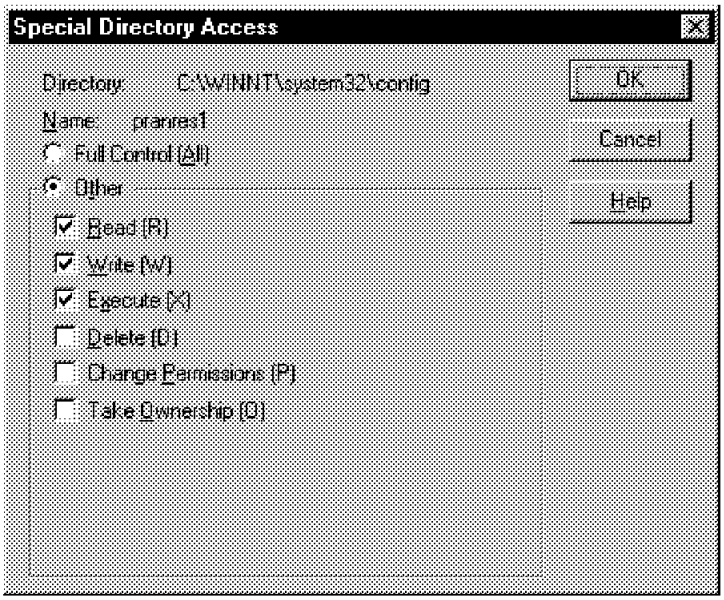

Figure 24. NT Special Directory Access Permissions

Note: Use caution when setting permissions in the Winnt directory or any of its subdirectories. If you remove the entry that gives ″Everyone - Full Control″ and

you have no other users with administrative rights for the directory, you cannot start the NT system.

3.4 Configuring NT for Availability

Now that we have established the “rules” governing the NT boot and load processes and their restrictions, we can look at configuration examples for NT systems. For systems that justify the effort, particularly NT servers, we recommend that you prepare for the time when the system will not load. An inability to load could result from a hardware failure (processor, adapter, or disk drive) or a system failure such as registry corruption or loss of a critical system file.

For these configuration examples we recommend that you configure an alternative, or “repair” NT system partition. This repair partition can be loaded and used to repair whatever has caused the problem. Having a repair partition available is particularly important when using the NTFS. With a FAT file system, you can boot a DOS system from diskette; thus you can access these FAT drives on the system. With NTFS, you do not have this option. The only system that can access an NTFS drive is another NT system.

In this section we review the system information you should save and two options for creating a repair partition:

- Installing a disk repair partition
- Installing a repair partition on removable media

These examples are based on using the NT install process to create and configure the repair partitions.

3.4.1 Saving the System Information

To reinstall an NT system on the same machine, we recommend that you save the system information for each computer in your environment. You will need the information for the NT installation. In case of a disaster, all of the offline data may be secure, but you will lose all of the system information.

You can save the system information in an ASCII text file stored on a diskette or in a printout stored at a remote site.

For our computers we saved the following information in a printout:

```
     System Definition Worksheet

LAN Details for TCP/IP Host:    NOGALES
====================================

Token Ring 16/4 ISA card

Token Ring Address : 400052002096
TCP/IP Hostname    : NOGALES
TCP/IP Address     : 9.1.150.142
Subnet Mask        : 255.255.240.0
SOCKS Server       : socks.almaden.ibm.com
                     9.1.40.40
Gateway Router     : 9.1.150.254
Names Server       : 9.1.72.196
Hostname           : nogales.almaden.ibm.com

Communication Mgr.
SNA 3270 Def.      : SCO2096

LAN Requester
Domain             : SIGD0471
Machine Name       : NOGALES

Hosts
Mozart      9.39.128.19
Haydn       1.90.156.40
Bach        7.34.22.128
```

For the NT installation the information about the kind of the server is important. If you install NT as:

1. Primary domain controller (PDC), you need the following information:
   - Domain name
   - Administrator name and password
2. Backup domain controller (BDC), you need the following information:
   - Name of the PDC
   - BDC machine name
3. Stand-alone server, no additional information is needed.

If you have installed other features such as the Windows Internet Naming Service (WINS), add all of the appropriate information to the above list.

Useful Tip

When you install a production server with NT, we recommend that you install it right after you install a repair partition for disaster preparation. The production server installation will be faster, because you have all of the required information fresh in your mind. In 3.5, “Using ADSM to Create Repair Partitions” on page 64, we describe how you can use ADSM to do a clone installation, which will bypass the long installation procedure. If you do not install the repair partition directly or have NT already installed, it is more important to have the above system information saved in a secure place.

3.4.2 Installing a Disk Repair Partition

The NT boot and system partitions can be located in separate disk partitions. You can take advantage of this flexibility by creating an alternate boot partition in another disk partition. Install another copy of NT in a different disk partition and update the system partition to enable loading of the repair system.

Note: It is also possible to install a repair partition after a system crash to enable a full restore process, but in this case you must have all of the required system information (see 3.4.1, “Saving the System Information” on page 55) and of course you have to go through a complete install with the accompanying delay to your recovery.

Once you have installed the boot partition and updated the system partition, you can copy the updated system partition files to an NT formatted diskette to create an alternate bootable NT loader (see Figure 25).

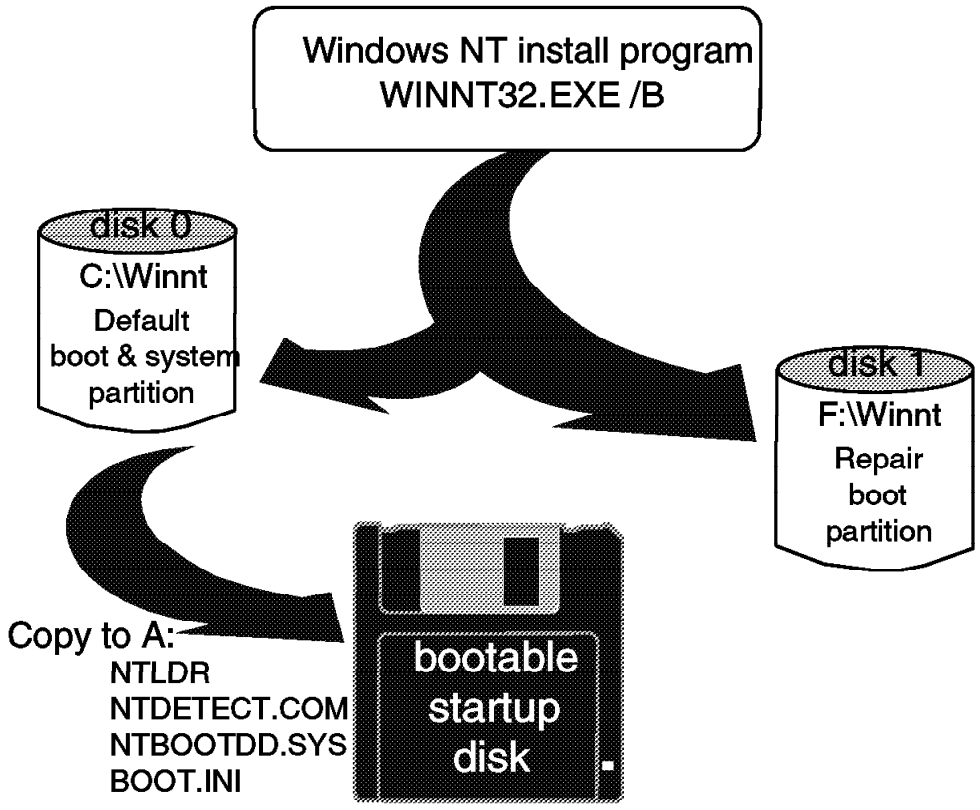

Figure 25. Creating a Disk Repair Partition

To create a disk repair partition:

1. Define the disk partition to be used
2. Install an NT system image and update the system partition
3. Create a bootable startup diskette

3.4.2.1 Defining the Disk Repair Partition

Before you install the NT disk repair partition, you must decide where to put it. For maximum availability, the NT disk repair partition should be installed on another physical disk, and, if possible, connected to a different disk controller. It may be a dedicated partition created with the NT Disk Administrator, or an existing partition with sufficient free space. With a dedicated partition the system can be isolated from other applications and users. However, a dedicated partition uses another drive letter on the system. Using an existing partition saves the trouble of having to define another drive, but it also exposes the repair system to potential damage from other applications or users.

Approximately 100 MB of disk space are required for a complete NT 3.51 system (the server or workstation version). For NT 4.0, a minimum disk space of 128 MB is required for the installation program. We recommend that you have a minimum of 150 MB available, because NT requires a permanent swap space on disk for the memory pages, and disk space is required for a restore program such as ADSM.

3.4.2.2 Installing an NT System

The actual installation of an NT system is straightforward. Two installation programs are available on the NT installation media. The first program, WINNT.EXE, is used for installing NT on an existing DOS system. The second program, WINNT32.EXE, is used for installing NT on an existing NT system. For our installation we used WINNT32.EXE. The install program has a number of options. For the purpose of installing our repair system, the /B option was very useful. It enabled us to install NT without creating installation recovery diskettes. These diskettes are created by default during the NT installation process and contain a subset of the NT system code. They can be used to assist in recovering a corrupted NT installation. However, for our purpose of installing an additional NT system, the diskettes were optional, so we used the /B option to bypass their creation.

Useful Tip

The /ox parameter creates three NT installation disks, which are required for a new installation with a CD. For the other parameters, please refer to the WINNT32.EXE′s online help.

When WINNT32.EXE is executed, it detects whether an NT system is installed. If an NT system is installed, the program asks you whether you want to upgrade it or install a new copy of NT. If you choose to install a new system, you are prompted for a drive location. If you specify the drive letter of your selected target disk partition and a directory name, the installation process continues. If you reinstall NT to create the repair partition, you have to do a complete install, which includes answering all configuration questions about machine name, domain name, network protocols, and other parameters. It is unnecessary to configure the repair partition as an exact clone of the normal NT system if you want to use the repair partition for the restore process only. Although the new computer name appears in the domain list, it can be deleted after the restore process is finished.

When the installation is complete, two tasks will have been completed:

- An NT boot partition will have been created in the specified disk partition. This will include a new SystemRoot directory structure (3.1.2, “NT Boot Partition” on page 39) containing all of the NT system files.
- The BOOT.INI file in the system partition will have been updated with the location and a basic description of the new boot partition.

At this stage, look at the BOOT.INI file (3.1.3, “BOOT.INI” on page 40). If an NT 4.0 workstation is installed, the installation will add two additional boot partition entries to the BOOT.INI file; for example:

```
scsi(1)disk(0)rdisk(0)partition(2)\WINNT="Windows NT Workstation Version 4.0"
scsi(1)disk(0)rdisk(0)partition(2)\WINNT=...
       ..."Windows NT Workstation Version 4.0 [VGA mode]" /basevideo /sos
```

These entries are the pointers to the newly created boot partition. The NT installation program creates two entries for each system. The first entry is the default system entry that will load the system on the basis of the configuration defined during its installation. The second entry is created for situations where the system will not load, because of either a bad device driver or missing system file. The /basevideo option loads the system, using a simple VGA device driver. The /sos option displays the individual NT system files on the screen as they are loaded and runs a checkdisk (CHKDSK.EXE) on the disk partitions. This option can be useful if a missing system file, file system problem, or bad device driver is preventing the system from loading.

The multi() and scsi() parameters for these two entries are created by the installation program, and it is not necessary to cover those again. It is sensible, however, to update the boot partition description. The installation program describes each boot partition by its level and type only: workstation or server. With multiple systems installed, it is advisable to update this description so that it is clear to users which boot partition they are loading. We changed the description of our repair partition to something meaningful, such as “Repair Partition,” by editing BOOT.INI, which is located in the root directory of the boot partition, typically the C: drive. As previously stated, this is a system and read-only file. Use the attrib command to reset these attributes before you edit the file and to set them back after you edit the file. Finally, while editing BOOT.INI, ensure that the default system, loaded after the boot timer expires, is the correct production system. The default system is defined in the [boot loader] section of BOOT.INI.

The repair partition is now ready for use. When the NT system is shut down and restarted, the new repair partition will be displayed on the NT loader selection screen:

```
OS Loader V4.00

Please select the operating system to start:


Windows NT Workstation 3.51
Windows NT Workstation 4.00
Windows NT Repair Partition

Use ↑ and ↓ to move the highlight of your choice
Press Enter to choose.
```

The newly created repair partition can be loaded by selecting it and pressing Enter.

3.4.2.3 Creating a Diskette Boot Partition

The repair partition can be used if the normal partition will not load. However, the NT loader still must started in the boot partition. Thus, the hardware BIOS must be able to access the NT master boot record on the first physical disk to invoke the NT loader. If the disk has failed, the MBR will be inaccessible and even though you have created a repair partition on another disk, it cannot be loaded because the system startup partition has been lost.

Therefore you have to create an emergency bootable system on a diskette. As mentioned in 3.1, “NT Boot Process” on page 37, the boot and system partitions can be located on different disks. The boot partition must be on a fixed disk, but the system startup partition can be on other media such as a diskette. Creating a bootable startup diskette is a simple process:

1. Create an NT bootable diskette by formatting it with NT. You can do this from File Manager or from a command prompt with the format command. Both of these will create the NT MBR on the diskette.
2. Using File Manager, copy the following files from the root of the system partition (C: drive) to the diskette:
   - NTLDR
   - NTDETECT.COM
   - NTBOOTDD.SYS
   - BOOT.INI

   Of these files, NTBOOTDD.SYS is optional and is required only if a SCSI controller is being used. Once you have copied these files, you can boot the diskette at system startup. A boot partition can then be loaded from any fixed disk on the system.

3.4.3 Installing a Disk Repair Partition on Removable Media

The disk repair partition installation process described in 3.4.2, “Installing a Disk Repair Partition” on page 57 works if the repair partition disk is accessible. However, to prepare for situations where the disk is inaccessible, or a permanent disk repair partition is not feasible, a repair partition can be created on removable media.

NT cannot be loaded from diskette. This is an NT restriction, discussed in 3.1, “NT Boot Process” on page 37. The system partition files can be copied to diskette to create a diskette boot loader. However, the boot partition must be on a fixed disk, that is, a device that NT understands as being “fixed” and “direct access.” This excludes sequential media such as tape drives. NT does, however, view removable cartridge devices as disks. Popular examples of these devices are Iomega Zip and Jaz drives, and SyQuest ezflyer and SyJet drives. These devices provide direct access and can be formatted as a FAT file system or an NTFS. The NT installation program allows you to install an NT system on these devices. The combination of a diskette-based NT loader and a removable boot partition allows the installation and configuration of a truly removable NT system (Figure 26 on page 61). With an ADSM system ready to run from this removable system, you can recover to another hard disk.

To use a ZIP drive with 100 MB of free space for an NT 4.0 installation, you must decrease the WINDOWS NT setup parameter in the first diskette of the NT Setup

diskette set. The file name is TXTSETUP.SIF. In this file you will find a **SetupData** section. In this section change the **valueFreeDiskSpace** to about 96 MB. This prevents the system from complaining about lack of installation space on the partition when installing NT4.0. With a minimum installation, you have about 4 MB of space left on the ZIP drive partition (including the TCP/IP network connection). There is still some space to reduce the number of Megabytes the NT system needs. You could reduce the swap file and delete some other files that are not needed for the recovery partition. This additional space is needed only if you plan to put the ADSM client locally onto the recovery partition. We recommend using a network-installed version if you are using an external ZIP drive.

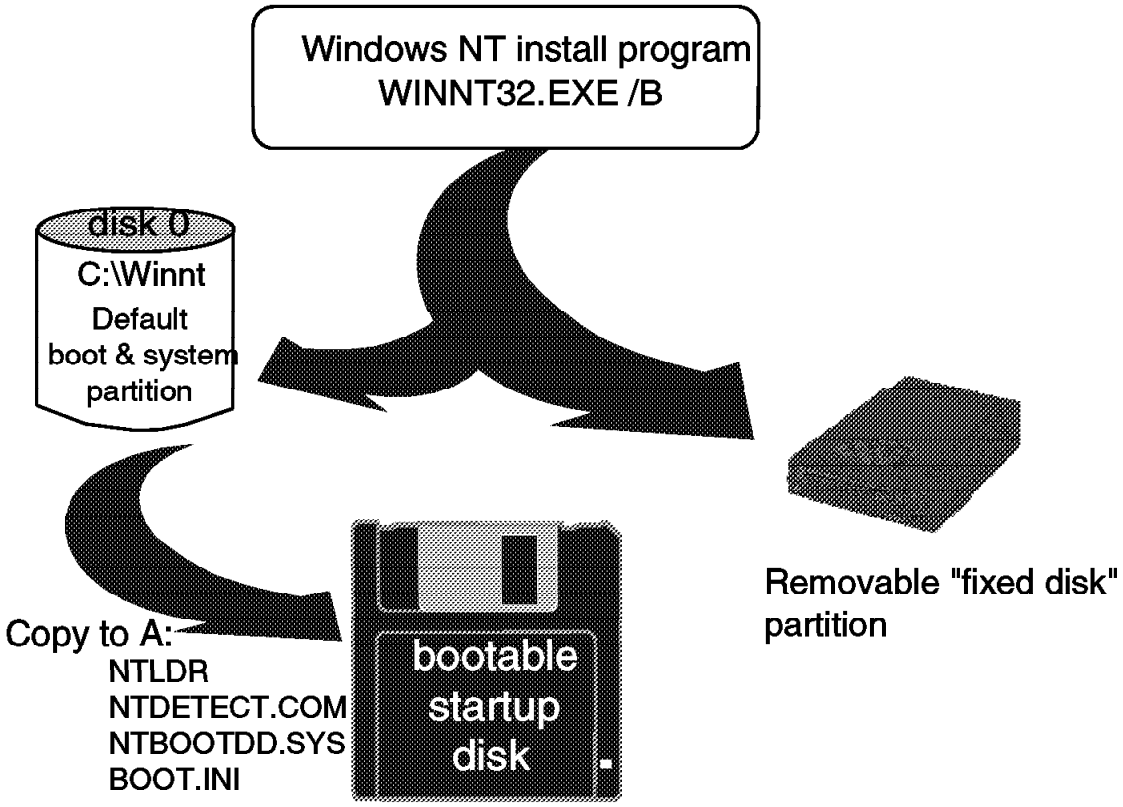

Figure 26. Creating a Repair Partition on Removable Media

For our example, shown in Figure 26, we use an Iomega Jaz drive. This device has cartridges of 1 GB capacity, which is sufficient for installation of an NT server or workstation system. First we discuss the hardware configuration considerations for an Iomega Jaz drive and then explain how to install an NT system.

Note: The Iomega Zip drive may be sufficient for installation of an NT 3.51 server or workstation, but for an NT 4.0 server or workstation you must have at least 128 MB of free disk space.

3.4.3.1 Removable Drive Configuration

Windows NT supports a variety of removable drive types. Details of all currently supported hardware can be found in the Windows NT Hardware Compatibility List (HCL), which you can find on the following Microsoft WWW and ftp servers:

- http://www.microsoft.com
- ftp.microsoft.com

Basically, NT supports a variety of SCSI-attached devices including an Iomega Jaz drive. In addition to a supported device, a SCSI adapter that is supported by NT is also required. The SCSI adapter must have on-board hardware BIOS, be able to automatically identify and configure SCSI devices when connected, and have an NT device driver.

The SCSI adapter could be an existing SCSI adapter used for other attached devices or a separate adapter specifically used for this task. If an existing adapter is used, care must be taken in defining the SCSI IDs. The Zip drive, for example, can only be configured with a SCSI ID of either 5 or 6, through a switch on the rear of the device. Such a restriction may conflict with other devices and dictate that a separate SCSI adapter be used.

After the Jaz drive is connected and the machine is powered on, the hardware automatically detects and configures the device. When NT is started, it detects that a new drive is available and runs a checkdisk operation on it during the load process. NT then assigns the next available drive letter to the new drive.

New SCSI adapter?

If a new SCSI adapter is also installed for the removable drive, the required NT device driver may not be installed. You have to install the device driver before the drive can be used. Use the Windows NT Setup program, which has an option to “Add/Remove SCSI Adapters.” You must install this option with NT before connecting the new removable device and installing the new NT system partition.

3.4.3.2 Repair System Configuration

The NT Disk Administrator views a Jaz drive as a physical drive with a single formatted 1021 MB primary partition (Figure 27 on page 63). Jaz cartridges are preformatted as a FAT file system. They can also be formatted as NTFS. However, for the purpose of creating an emergency repair partition, FAT is the best choice.

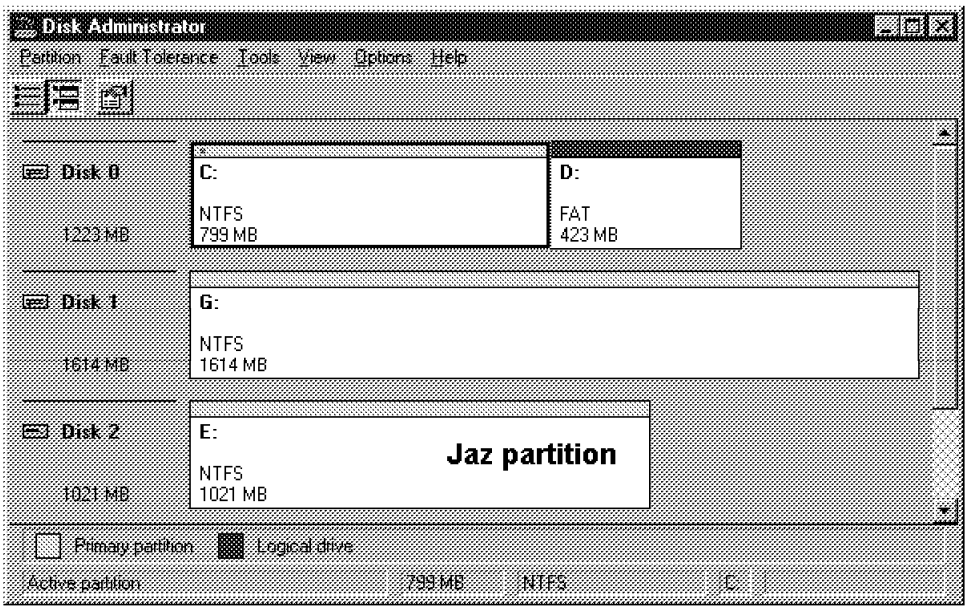

Figure 27. NT Disk Administrator Showing Jaz Partition

Once the drive is configured to NT, the installation process is the same as for a fixed disk. NT updates the BOOT.INI with the location of the system partition on the Jaz drive. When the NT installation is finished, you can load the newly created NT system in the same way as with any other installed system. It is important to update the system description in the BOOT.INI file because the installed system is removable. As with the fixed disk example, you can create a diskette boot partition by copying the NT loader files and BOOT.INI to an NT formatted diskette. The combination of this diskette and a removable cartridge system partition creates a truly portable NT system that can be removed and kept in a secure location for future use. The BOOT.INI file on the primary disk partition should be updated to remove the pointers to the partition that has been removed.

SCSI configurations

Exercise caution when using SCSI adapters for booting NT images, particularly if you are using more than one type of adapter. The installation process copies the appropriate SCSI device driver for the adapter card to the system partition and renames it NTBOOTDD.SYS. Only one NTBOOTDD.SYS can be in the system partition, and it will be a copy of the driver for the adapter used during the last installation. If a previously installed NT partition is used with a different driver, it will no longer load, as its “personalized” NTBOOTDD.SYS will have been replaced by the subsequent installation. A sensible precaution is to copy and rename NTBOOTDD.SYS before installing it with a different adapter. In this way it is always possible to correct any problems after the installation.

3.4.4 Using Repair Partitions

A repair partition is used primarily in system recovery when the normal NT system will not load. Having an alternate loadable partition enables you to access the corrupted or damaged primary partition and, it is hoped, resolve the problem.

Note: You can change or repair files on an NTFS partition only from a running NT system.

In most cases, system load problems will be caused by a corrupted registry or lost system file. Loading a repair partition enables you to inspect such problems and potentially resolve them. One of the problems for administrators and users of NT is that a number of files, including the registry files, are opened and locked during the NT load process (see 3.3.1, “Security Mechanisms” on page 50). It is nearly impossible to perform any remedial action on a running NT system; however, running from a repair partition changes this situation. If the repair partition is installed on the F: drive, NT will run from that drive, and all locked system files will be on that drive. If the primary partition is the C: drive, the repair partition treats all system files, usually locked, as ordinary files. If the problem is a corrupted registry or a missing device driver, these files can be copied, renamed, and generally treated like any other files. Lack of protection exposes the system to potential damage. Use care when working with system files. However, the ability to use a repair partition prevents the need to reinstall NT and, in the case of a disk failure, a repair partition enables you to quickly get an NT system running again. Once you have replaced the failed disk, you can use the repair partition to configure the disk partitions and restore the data.

3.5 Using ADSM to Create Repair Partitions

To install a repair partition, you must create a new system image on a fixed disk or removable media. Therefore, you must answer all questions about the machine, such as disk configuration, video card, and install directory, and all questions about the network, such as IP address, domain, and network card. This takes time. See 3.4.1, “Saving the System Information” on page 55 for more details.

A faster and easier way to create an image of an NT system is to copy the whole operating system from the production partition. NT has no information about drive letters in any file of the system partition—only partition descriptions such as `harddisk0/partition1` are used. Other operating systems, such as Windows 95, store information about drive letters in several settings and therefore cannot copy the information to another location.

Note: A DOS command such as copy or xcopy cannot copy an NT system to another partition because several files are not readable from a running NT system.

In conjunction with ADSM you can install a repair partition quickly. You must complete the following steps:

1. Install the NT server
2. Format the repair partition with NTFS
3. Install ADSM
4. Back up the NT server incrementally (including the registry)

5. Restore by subdirectory branch to another partition

The installation of the NT server is a relatively straightforward procedure, described in *NT - Getting Started* from Microsoft.

With the NT Disk Administrator you must prepare another partition for the repair partition. We recommend a disk space of 128 MB or more. Format the disk space with the Explorer or with this command:

```
format x: /FS:NTFS /V:label
```

The drive must be labeled for use with ADSM, and the file system should be the same as the production partition (typically NTFS, as in our example). The drive letter x is only a placeholder and can be reset to the assigned drive letter from the NT Disk Administrator.

For ADSM, ensure that you have the recommended resources on disk and in memory. Install the ADSM software by answering the questions about the path and the environment settings. After the installation you must not reboot the computer to start ADSM. Change the DSM.OPT file for your environment. The information about the ADSM server, client name, communication protocol, and include/exclude list must be correct.

Note: See 4.1.2.2, “Exclude/Include Lists” on page 72 for the options file content for backup of the operating system.

Check the connection to the server and start the incremental backup through the GUI or with this command:

```
dsmc incr ... -backreg=yes
```

You must have administrative rights to back up the registry. When the backup is finished, you can start the restore process without making any changes. From the GUI select **Restore by subdirectory branch** or start the restore with this command:

```
dsmc res \ x:\*.* -regres=yes -subdir=yes
```

For the destination of the restore you must select the prepared partition. For details on the correct restore of the registry, refer to Chapter 4, “System Recovery with ADSM” on page 69.

We did not start an additional installation process of the original NT media, so we had to insert the new NT system manually into the BOOT.INI file of the startup partition.

When you reboot the system, you can start the copied partition from the start menu.

Useful Tip

In Chapter 4.6, “Migration of a ñ System” on page 86, we use this method to migrate a Windows NT system from a small computer to a larger computer, thus avoiding a total reinstallation.

3.6 Multiple Versions of the NT Registry

In this section we explain how to handle multiple versions of the NT registry and how you can restore an older version of the files.

In the ADSM client you can choose whether ADSM should show only the active files or both active and inactive files. The files with a dot in the icon represent the inactive, the older versions of the files, see Figure 28.

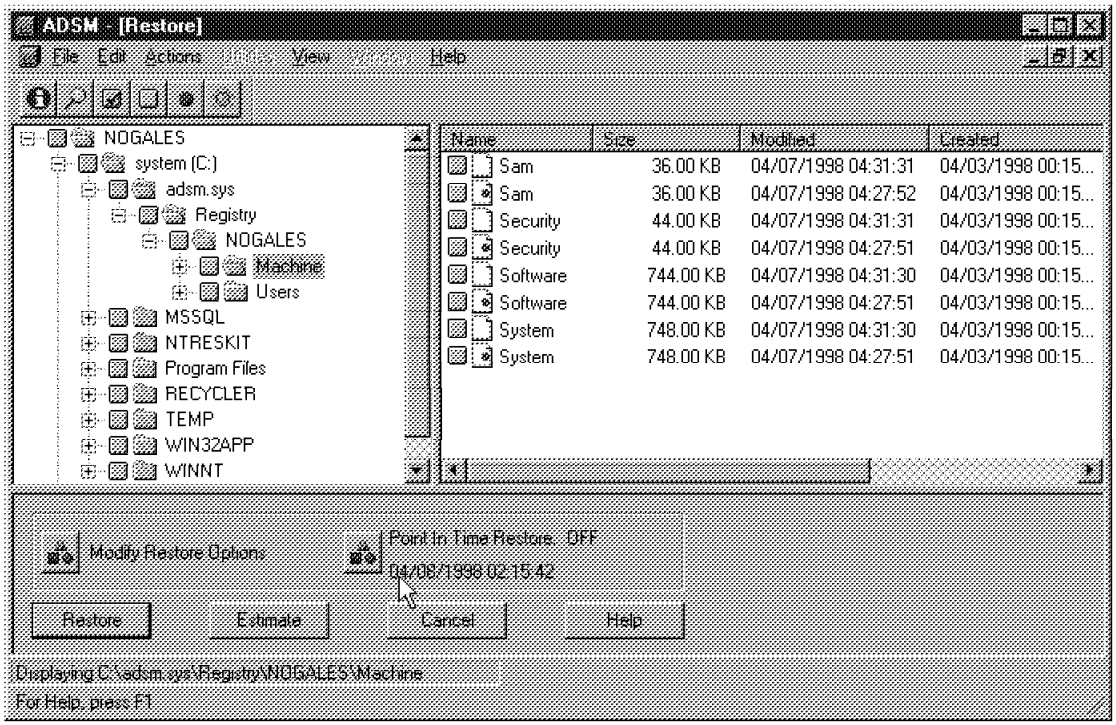

Figure 28. ADSM Client GUI Showing Inactive Files

In some circumstances it might be helpful to restore a special version of the registry.

Note: It could be dangerous to restore an older version because each NT program writes values into the registry. If you restore a registry version and a program does not find the appropriate registry entries, the program might not function correctly.

With the ADSM REGREST command, you can restore only the most recent version of the registry. It is not possible to restore an inactive version of the registry. The manual process of restoring older versions is straightforward:

1. Select the files in the ADSM GUI. Be careful to select files from one backup only. Do not mix the files by selecting files from different backups.
2. Restore the files to their original location, which in this case is the adsm.sys directory.
3. You will get a warning that the files already exist, but replace them all.
4. Start the registry restore. From within the ADSM GUI use the **Restore Registry** command in the **Utilities** menu or from the ADSM CLI use REGBACK ENTIRE.

5. Be sure that you check the **Activate Key after Restore** box in the dialog window.
6. The ADSM client tries to restore the latest version of the files into the `adsm.sys` directory, but this time, do not allow replacement of the files on your disk. This will guarantee that the older files will remain on the disk.
7. The last dialog window that appears confirms, that the registry restore has completed and the restored version has been activated as the current registry.
8. Reboot the machine for the changes to take effect.

Chapter 4. System Recovery with ADSM

This chapter looks at how ADSM can be used in conjunction with repair partitions to provide full system recovery capability. We cover the following topics:

- ADSM setup considerations
- Preparing for Windows NT recovery with ADSM
- System recovery (bare metal restore)
- Additional system recovery (network drives, PDC, and BDC)
- Considerations for an ADSM server on the same machine as data that is being backed up
- Migration of a Windows NT system

4.1 ADSM Setup Considerations

A current backup of all system files required to boot and load the system is the basic requirement for a full system recovery. An ADSM client running on an NT system cannot restore the system back to itself, because the user who made the backup and the user who wants to make the restore have different SIDs. As described in 3.1.1, “NT System Partition” on page 38, it is not possible to re-create the same user, although the new user has the same name, password, and rights. Because of incorrect user rights, if you try to restore the system files over the running NT system, you will receive error messages such as ″Access denied.″ The NT system protects the system files, and even the same backup user with the same SID could not restore the system files.

However, the ADSM client can back up the system files from an active and running NT system and can perform a full recovery using an alternative repair partition. To ensure that an up-to-date back up is always available, the ADSM client should be used to perform regular incremental backups of the primary partition and should be configured to back up the NT registry.

In this section we look at:

- Where to install the ADSM client
- Client options file
- ADSM backup archive client execution mode

4.1.1 Where to Install the ADSM Client

Configuring the ADSM client in the described manner ensures that the backups required for a full system recovery are performed on a regular basis. If recovery using a repair partition is required, the up-to-date backup inventory can be restored by an ADSM client running from a repair partition and thus recovering the primary partition (see Figure 29 on page 70).

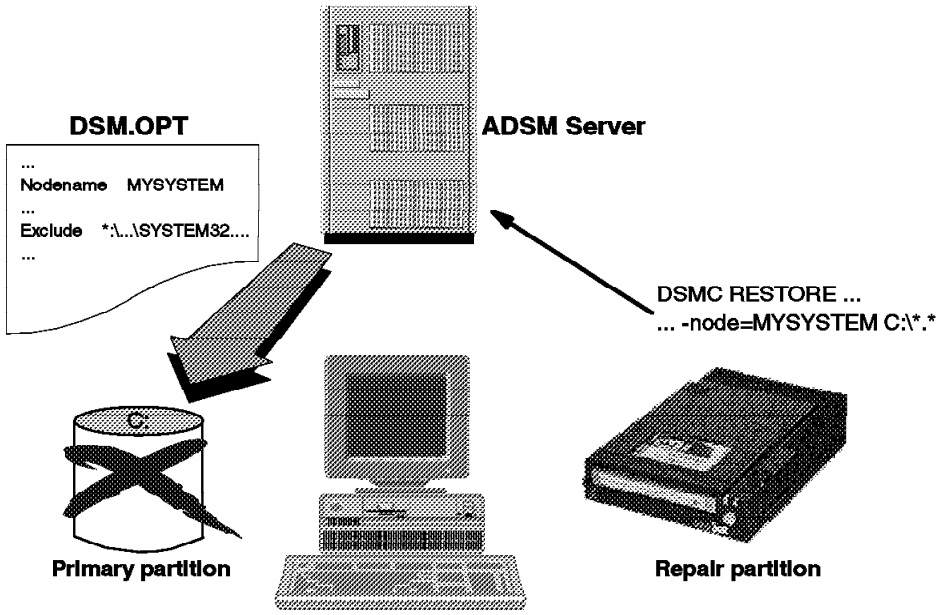

Figure 29. Primary Partition Recovery with ADSM

Assuming that the repair partition has been created on a removable device and a diskette boot loader was created, it is possible to boot and load an NT system image, even if the primary partition is completely lost. This repair system will appear just like any other NT system to the remainder of the network. This is very important as it gives you the ability to access other systems and resources on the network. With NT running on a repair partition, the original primary partition and the disk on which it is to be reinstalled are just another disk partition to the NT system. All the system files that were locked by NT while it was running on the original primary partition can now be restored by using an ADSM client running from the repair partition.

Creating a working repair partition on removable media enables you to access, copy, or restore an NT system image that would otherwise be protected by the system. For restoring the system with ADSM, access to an ADSM client is also required. Two ways to gain access to an ADSM client are:

1. Install the ADSM client on the repair partition
2. Use a network shared ADSM directory

The first option is the simplest. However, it requires that the ADSM client be installed on each repair partition if it is to be used. It also takes up valuable space on the repair partition, which may be at a premium if a removable device such as a Zip drive is being used. An alternative in an NT network environment is to set up an ADSM client directory on another machine as a shared resource for the network (see Figure 30 on page 71).

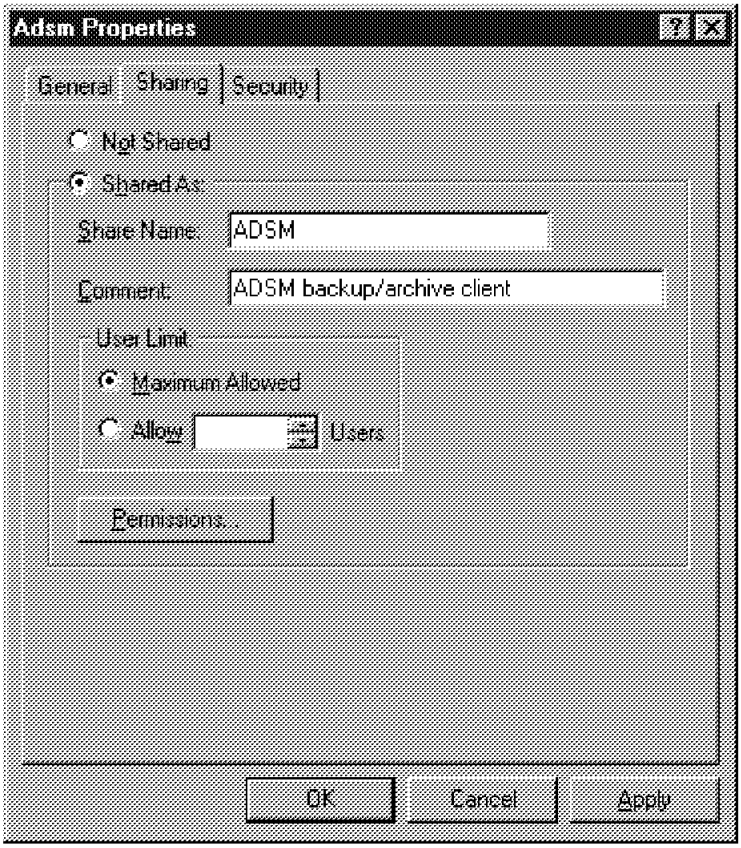

Figure 30. Setting Up Network Share for an ADSM Client

Figure 30 shows the ADSM backup/archive client directory on an NT system with a share name of ADSM. In addition to this network name, permissions can also be defined to control access to specific users, or groups of users such as ADSM administrators. With this technique, the ADSM client can be run as a networked application from another NT system elsewhere on the network, thereby removing the requirement to install the ADSM client code on individual repair partitions.

ADSM as a network application Be careful when running ADSM as a

networked application. When invoked from a network share it will use the client options file, DSM.OPT, located in that shared directory. These options, for that other system, may not be appropriate for the system you are running. They can be overridden by specifying the options as parameters when starting the ADSM client. An example of this is using the -NODENAME parameter to set the ADSM nodename to that of the system you are running.

4.1.2 Client Options File

The client options file, DSM.OPT, controls how the client works and which files are backed up. The basic configuration options such as nodename and connectivity options are not covered here. We review instead the NT registry backup option and the Exclude/Include file filter lists.

4.1.2.1 Registry Backup

The most important NT system resource that must be backed up is probably the registry. The registry consists of a number of files in the system32\config directory that make up data structures that are often called *hives*.

The hives are locked by NT when it starts and cannot be accessed by other applications such as the ADSM client. To back up these files, NT provides an API that applications can use. This API is used by the native NT registry backup; the restore commands, REGBACK.EXE and REGREST.EXE; and the ADSM client. Use of this API enables the ADSM client to perform online backup and registry recovery. This is controlled in the ADSM client by the BACKUPRegistry option in the client options file. The option is set to Yes by default, so the registry is backed up every time an incremental backup is performed. Backup can be turned off by setting the BACKUPREG option to No. Do not set BACKUPREG to No if you want to fully recover your system.

Important

The user account running ADSM must have administrative rights to back up the registry. As described in the ADSM readme file, it is untrue that the Backup Files, Restore Files and Manage Auditing and Security Log rights are sufficient. A “normal user” in NT, that is, one who has no administrative rights, can do incremental backups if the BACKUPREG NO option is set in DSM.OPT.

Registry backup is a two-stage process. First, the ADSM client uses the NT API to create a backup copy of the registry files. It stores this backup copy in the adsm.sys directory. This directory and its subdirectories mirror the structure of the actual registry directory. Second, the ADSM client backs up the backup copies in adsm.sys as if they were ordinary files. This two-stage process provides for the file system a local backup copy of the registry as a copy stored on the ADSM server.

4.1.2.2 Exclude/Include Lists

The Exclude/Include lists exclude or include files from being backed up depending on how they are defined. When installation is complete, the ADSM NT client has a default set of Exclude/Include statements. These defaults ensure that all of the necessary files for a system recovery are backed up. The following are sample default statements:

```
..
Exclude *:\...\pagefile.sys
Exclude *:\...\system32\config\*.*
Exclude *:\...\system32\config\...\*
..
```

The NT virtual memory paging file, pagefile.sys should always be excluded from backup because it is usually very large and backup is unnecessary. When an NT system starts and cannot find its paging file, it creates a new one based on the system settings defined in the Control Panel application.

The system32\config directory must always be excluded because it contains active registry files that cannot be backed up. Attempting to back up the directory will result in “access denied” errors. The adsm.sys directory structure, however, must not be excluded because it contains the registry backup files created by ADSM and required for a registry recovery.

Because of the huge filespace required for backup, we recommend excluding the Windows recycle bin. In a typical NT installation, the reserved space for the recycle bin is about 10% of the disk space. With a 1 GB disk, the recycle bin can hold up to 100 MB of disk space. As an additional exclude line, type `Exclude *:\...\recycler\*.*` in `DSM.OPT`.

On a Windows 95 machine, the recycle bin is in the `Recycled`

directory. The directory has the `SH` attributes and is therefore not visible with the `DIR` command on NT or Windows 95.

4.1.3 ADSM Backup/Archive Client Scheduling

The best way to ensure that regular backups take place is to configure the ADSM scheduler to run as an NT service (Figure 31).

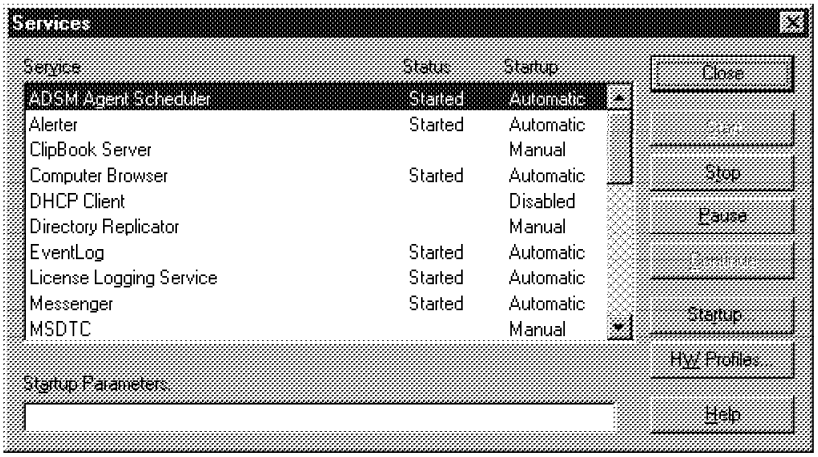

Figure 31. ADSM Scheduler Service

This ensures that the client scheduler cannot be canceled accidentally by the user; that it runs without a locally logged on user; and that it starts automatically when the machine is booted. In addition, a schedule should be set up on the ADSM server to perform regular incremental backups, ideally, on a daily basis. The actual operation of the scheduled backups is controlled by client configuration options defined in the `DSM.OPT File`.

The ADSM scheduler can be started as a program or as a service on an NT server. Started as a service, the advantages are:

- Logged on as a system account
- Protected from other applications
- Protected from the locally logged on user
- Hidden for the locally logged on user (runs in the background)
- Faster, because the service will not be paged to disk (NT right “Lock pages in memory”)
- Located in the secure part of the working memory, because the operating system is responsible for the memory allocation and no user can move the program to another place in memory

It is better to start the ADSM scheduler as service than a program. As a program, the ADSM schedule window remains on the working desktop at all times.

If you start the ADSM scheduler as a service, set the startup type to **Automatic** in the Control Panel.

4.2 Preparing for Windows NT Recovery with ADSM

In this section, we describe how to prepare for recovering a Windows NT system with ADSM.

Note: ADSM DRM provides these steps automatically.

Imagine that an accident, a fire, for example, has destroyed all your machines and it is impossible for you to return to your office because the site has been severely damaged.

Your company needs to continue its business activity, but all you have prepared is the offsite ADSM backup, which is not enough.

Therefore you also must:

1. Prepare the Jaz, WORM, or CD-ROM with a repair partition and the ADSM directory (see Chapter 4, “System Recovery with ADSM” on page 69).

   For a more comfortable recovery, you have to put all SCSI drivers on the repair disk, so that the NT boot process can complete without errors. This is needed only if you have changed the hardware.

2. Perform an ADSM database backup

   Some ADSM files contain ADSM server configuration information that changes all time. You must put these files on magnetic support (a floppy disk) to complete the ADSM directory that is on the bootable repair unit.

3. Create a directory in your repair disk, containing all the *.INI files made by the setup program using the -r parameter. With ADSM, you will be able to restore the complete configuration of your NT server. But, in some special cases, you may need to reinstall some software.

4. At the backup site keep a copy of your backup and recovery plans. Keep a copy of this redbook and *ADSM Server for Windows NT Configuration and Recovery Examples* (SG24-4878) with your own documentation.

4.2.1 Perform ADSM Database Backup

If your ADSM database is not too large, you can restore it with a full database backup. For large databases, use incremental backups.

You can back up the database every day, after all of the other backup processes. You can use the ADSM Administrative GUI by selecting **Database Backup** or you can plan an administrative command schedule:

```
DSMC BACKUP DB devclass=8MM type=full
```

If your backup strategy plan is to use a specific volume, you can specify the volume name with the BACKUP DB command.

4.2.2 ADSM Configuration File Backup

You will need these files to perform the recovery of the ADSM server:

1. **DSMSERV.OPT**

   The DSMSERV.OPT file contains the server options. It is a text file that can be edited if you need some information about your ADSM server. Some ADSM parameters such as Uselargebuffers, must be the same for the backup and for recovery processes, so you must back up the DSMSERV.OPT file.

   To back up this file, select the Server Properties window from the Administrative GUI, select **Backup** and click on the **Save** button.

2. **Device configuration file**

   The ADSM device configuration file is named with the DEVCONFIG option in DSMSERV.OPT, so check the option setting. We used the default name, DEVCNFG.OUT. This file contains the commands necessary to re-create the server device configuration.

   To back up this file, use the administrative command schedule:

   ```
   DSMC backup devconfig file=DEVCNFG.OUT
   ```

3. **Volume history file**

   The volume history file contains details of all sequential volumes used within storage pools, for database backups and for server exports. So this file is very important, because all you have at the recovery site are sequential volumes. The file is vital for recovering the server database because you can check the latest database backup volume.

   The name of the volume history file is specified in the DSMSERV.OPT, the default is VOLHIST.OUT.

   To back up this file, use the administrative command schedule:

   ```
   DSMC backup volhistory file=VOLHIST.OUT
   ```

4.2.3 ADSM Configuration Macro

You can write an ADSM macro and a command file and run them with the NT **AT** scheduler or the ADSM scheduler:

1. **ADSM macro**

   ```
   query db f=d
   query dbf f=d
   query logv f=d
   query volhist type=dbb
   query option
   query sta
   query dbb f=d
   ```

   Name the macro SOS_MACR and schedule the macro with an out file name such as SOS_MACR.OUT:

   ```
   dsmadmc -id=admin -pass=pw macro SOS_MACRO > SOS_MACR.OUT
   ```

2. **Command file**

```
       * * *    ADSMTOOL.CMD   * * *
C:
CD\win32app\ibm\bacserv
COPY DSMSERV.OPT A:\DSMSERV.OPT
COPY SOS_MACR.OUT A:\SOS_MACR.OUT
DSMC backup devconfig file=A:\DEVCNFG.OUT
DSMC backup volhistory file=A:\VOLHIST.OUT
```

This command file writes four files to a floppy disk. Keep the floppy disk with the sequential volumes at the recovery site.

4.3 System Recovery (Bare Metal Restore)

Having discussed how to configure and use a recovery partition with ADSM, in this section we look at the steps involved in recovering a hard drive containing an NT system partition and another data partition. Depending on the reason for performing the recovery, a number of steps may be involved. Our example assumes that the physical disk that held the primary partition has been replaced because of a defect. The steps involved in recovering the system are:

1. Boot from the repair partition
2. Create partition(s) on disk and assign drive letter(s)
3. Format the primary partition
4. Restore the system partition from the last incremental backup
5. Restore the registry
6. Shut down and reboot the system
7. Recover the remaining partitions

4.3.1 Boot from Repair Partition

You boot from the previously configured repair partition, using the prepared boot diskette. Use either an alternative disk partition or a partition on removable media. This repair partition provides you with an NT system from which you perform the remaining steps.

Note: Remember that when the recovery system is running, the active registry is that of the recovery system. Commands that work with the active registry should take this into account.

4.3.2 Create Partitions

After a disk replacement, you can redefine the disk's partition information, using the NT Disk Administrator tool. You define the replacement primary partition and any logical volumes defined within the extended partition by using the **Partition** pull-down menu of the Disk Administrator (Figure 32 on page 77).

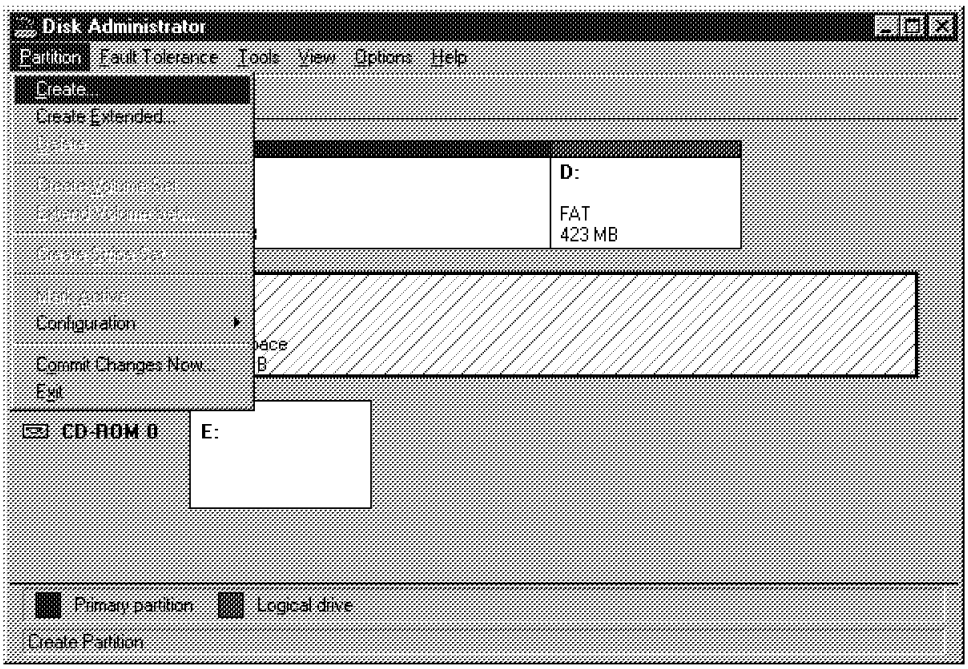

Figure 32. Defining Partitions with the Disk Administrator

Alternatively, you can use the DISKSAVE.EXE program as described in 3.2, “Useful Availability Tools” on page 42 to restore the disk’s MBR complete with the partition table.

When the partitions have been defined, they must be assigned with drive letters. Drive letters are assigned either automatically by NT or from the Disk Administrator **Tools** pull-down menu. Take particular care with drive letters when adding new disks. NT assigns drive letters by default, starting with C, on the following basis: primary partitions first, followed by logical volumes in extended partitions. This dynamic allocation can cause problems for applications that depend on a specific path to resources. If you use Disk Administrator to assign drive letters, static drive letters can be used. These letters can be assigned to partitions and logical volumes in whatever sequence is required. They are retained when new disks are added.

Saving partition information

Disk Administrator enables you to save the partition definitions to diskette through the Configuration option on the **Partition** pull-down menu. You can also restore previously saved definitions; however, be careful because a restore of partition definitions overwrites all partition information for the entire system.

4.3.3 Format the Primary Partition

Once you have defined new partitions on the disk and assigned drive letters, you must format them. Use the **Tools** pull-down menu from Disk Administrator or the FORMAT command from the NT command prompt. For example, the following command formats the C: drive with NTFS, labels it C_DRIVE, and creates the MBR on the drive, making it bootable as an NT boot partition:

```
FORMAT C: /FS:NTFS /V:C_DRIVE
```

4.3.4 Restore the System Partition

Having re-created a bootable primary partition, you can restore the most recent ADSM backup for the system. For our example we use the command line ADSM client, DSMC.EXE. Our example assumes that the ADSM client is being run as a networked application from a shared drive elsewhere on the network. If you have an ADSM client on the repair partition itself, you can of course omit the nodename.

The following ADSM command restores the entire C: drive:

```
DSMC> RESTORE  -NODENAME=MYSYSTEM  -SUBDIR=YES  C:\*
```

The -NODENAME parameter is used here to identify this client session to the ADSM server with the correct NODENAME for the system being restored. This is important if you are running ADSM as a networked application. If -NODENAME is not specified, the client would use the NODENAME coded in DSM.OPT, which is located in the ADSM directory being accessed on the network. Specifying the -NODENAME parameter ensures that the ADSM client running on the repair partition correctly identifies itself to the ADSM server. The second parameter, -SUBDIR, tells ADSM to restore all subdirectories of the specified restore object, which is in this case the root directory of the C: drive. Note: If there are some files on the disk, you must use the replace=all option as well.

The above example assumes that the new primary partition being restored has the same volume label as before. ADSM actually substitutes the drive letter specified (C:) with a “filespace” name. The filespace name is the volume label for the drive and is used by the ADSM server to uniquely identify file systems on client workstations. If you have labeled the new partition with a different label, a restore with a drive letter specification will not work. ADSM will not be able to resolve the identity of the backup filespace on the server. An alternative method of restoring in this situation is to use the ADSM server filespace name rather than a drive letter:

```
DSMC> RESTORE  -NODENAME=MYSYSTEM  -SUBDIR=YES  {"filespace"}\*
```

The filespace name can be determined by first using a QUERY FILESPACE command. The output lists the client filespaces held on the ADSM server.

Once you have performed this restore, all files required for the NT boot and system partitions should have been restored. This includes the boot partition files in the root directory, NTLDR, NTDETECT.COM, BOOT.INI, NTBOOTDD.SYS, and the contents of the SystemRoot directory and subdirectories that make up the system partition. At this stage the restored drive should be a fully functional boot partition. When booted, it would invoke the NT loader and display the list of system partitions on the basis of the contents of BOOT.INI.

Note: Do not start the restored system until the registry has been fully recovered.

4.3.5 Restore the Registry

The previous step restored the last registry backup to the adsm.sys directory and subdirectories. ADSM uses the adsm.sys directory to hold its latest backup version of the various registry files. When you run the ADSM client from the repair partition, this registry cannot be automatically restored. To restore the registry the command line or GUI options work with the active registry. In this case, the active registry is that of the repair partition. At this state of the recovery, still running on a repair partition, the backup registry files must be manually copied to their correct directory locations. If locally logged on users are in session when the ADSM backup is run, their profiles have to be restored too.

4.3.5.1 Restoring the Registry Hives

The registry is a hierarchical database containing all system, hardware, software, user, and desktop settings. It is logically organized into four separate data structures, or hives. These hives are physically stored in a number of files on the NT system partition. When an NT system partition is installed, a directory named system32\config is created. The registry files are placed in that directory. The important files in System 32\config are:

- Sam
- Security
- Software
- System

These files contain the information stored in the registry hives. Other files, located in %SystemRoot%\Profiles, define the desktop settings for all the users that have logged on to the NT system locally. The system32\config\ directory also contains a file named Default, which is the default desktop.

Without these files in the system32\config directory, an NT system will not load. It will stop with an error, indicating that the registry is corrupted. When ADSM backs up the registry, it creates copies of these files in the adsm.sys directory and its subdirectories. You can use these copies from the recovery partition to recover the registry of he failed the failed system. Figure 33 on page 80 shows the contents of the adsm.sys and the system32\config directory structure. The files in the adsm.sys directory are exact copies of the real registry.

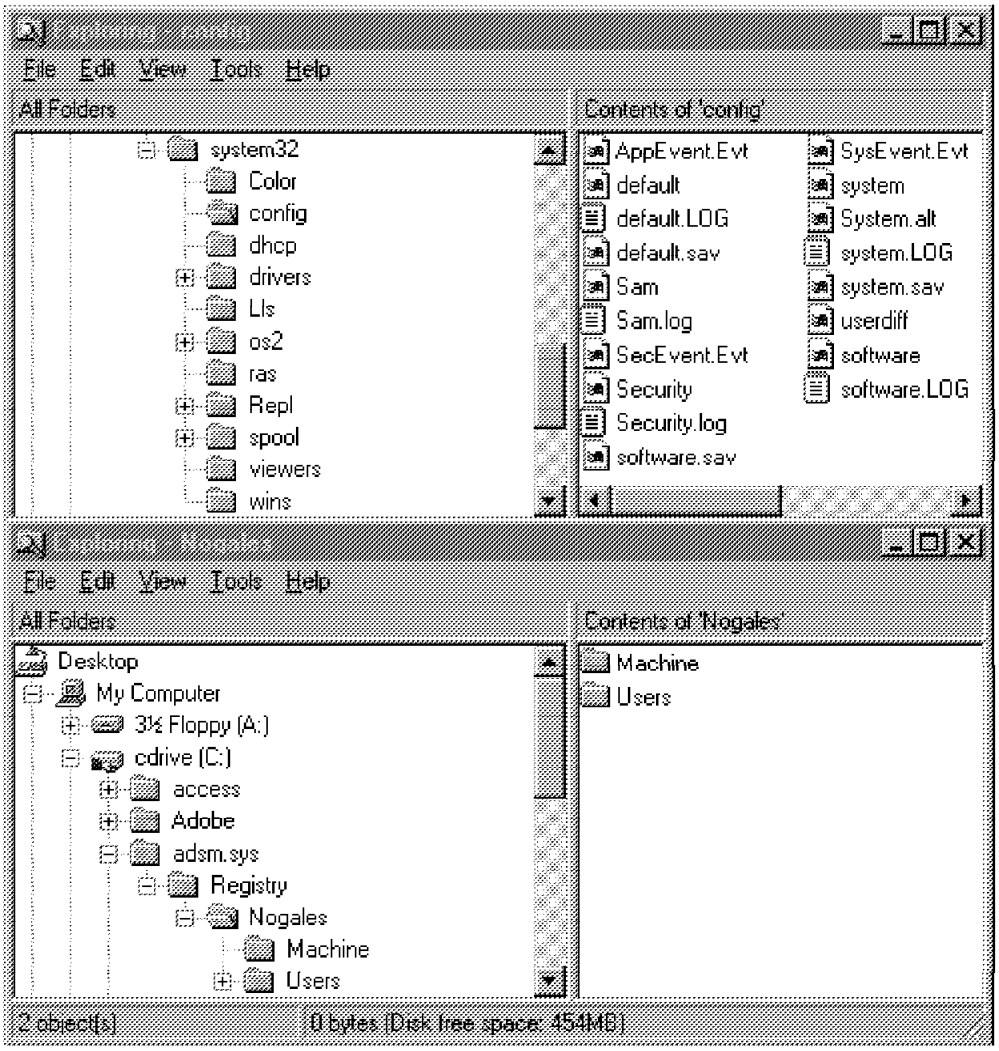

Figure 33. Registry and ADSM.SYS Directory Structures

The registry\"machine name" directory identifies the name of the NT system, in this case "Nogales". The last copy of the registry is included in the registry\nogales\machine subdirectory. The four files have the SHR file attributes and contain all of the production system's hives. The registry\nogales\users subdirectory contains the files for the profiles, but only for the default settings and the person who performs the backup. This person is the only one locally logged on at backup time. All other users are remote users logged in through the network. There can be no more locally logged on users because NT is not yet a multiuser operating system. Before the system can successfully start, these profile files must be copied to their original position. Using the NT Explorer or the NT command line, you must perform the following steps:

1. COPY C:\ADSM.SYS\REGISTRY\NOGALES\MACHINE*.* C:\WINNT\SYSTEM32\CONFIG
2. COPY C:\ADSM.SYS\REGISTRY\NOGALES\USERS\DEFAULT C:\WINNT\SYSTEM32\CONFIG

These two copy commands restore the original production system's entire registry. If the ADSM backup was run as an NT service by the ADSM scheduler, and a user was not logged on locally, you have no additional steps or copy commands to perform. You can reboot from the restored system disk and then proceed to 4.3.7, "Recover the Remaining Partitions" on page 82 to recover any additional partitions.

4.3.5.2 Restoring the Locally Logged On User (If Required)

If a user was locally logged on at the time of the last backup, there are some additional steps to be performed. The ntuser.dat and ntuser.dat.log files have the special NT attribute “deny-read” or “no access” from the operating system. These files are part of the registry and they contain the settings of the user profile, such as background color/picture, desktop/environment, and mouse pointers. During the backup process, ADSM extracts copies of these files from the registry and stores them in the adsm.sys directory. You can recover the user profile settings by copying the ntuser.dat and ntuser.log files as described below.

ADSM Client Runs As a Service, but a User Is Logged On: Assume that the locally logged on ID for backup was JOE and the registry directory after the ADSM restore is c:\winnt\profiles\JOE.

The c:\adsm.sys\registry\NOGALES\users\JOE directory has two files, JOE and JOE000.KEY, which represent the ntuser.dat and ntuser.dat.log files. There are several key points here. The directory represents the current user logged on. In this situation the local logged on userid is JOE. When ADSM runs as a service and there are no locally logged on users, the directory exists but there are no locked files.

This scenario is confusing because, to restore the proper profiles, you must copy the files to the proper place. Unless you follow the steps below, when you sign on as JOE, NT will assign a default profile and you will get a different desktop. A new directory would be created in the Profiles as JOE.000.

Note: Even if you try to use the DSMC REGBACK USER CURSER command, you are not going to be able to restore the previous desktop settings. The moment you log on to the system, NT has already created the JOE userid and a restore of the profile with the above command will restore the original profile to this new user.

To restore the proper profiles, issue these commands:

```
copy c:\adsm.sys\registry\NOGALES\user\JOE\*.*
 c:\winnt\profiles\JOE
rename c:\winnt\profiles\JOE\JOE000
 c:\winnt\profiles\JOE\ntuser.dat
rename c:\winnt\profiles\JOE\JOE000.key
 c:\winnt\profiles\JOE\ntuser.dat.log
```

This ensures that the proper profile is set for JOE before logging on to the system.

Now let us discuss how to copy the registry files directly with the ADSM client. The registry files that need to be copied were previously restored from the ADSM server to the adsm.sys directory. As you can see from the above, with an understanding of the registry structure, it is quite simple to restore files from the ADSM server directly to their correct locations. You can restore manually or within an ADSM macro by using restore commands as in these examples:

```
1. RESTORE C:\ADSM.SYS\REGISTRY\NOGALES\MACHINE\*.* C:\WINNT\SYSTEM32\
CONFIG
2. RESTORE C:\ADSM.SYS\REGISTRY\NOGALES\USERS\DEFAULT C:\WINNT\SYSTEM32\
CONFIG
3. RESTORE C:\ADSM.SYS\REGISTRY\NOGALES\USER\JOE\JOE000C:\WINNT\
PROFILES\NTUSER.DAT
```

```
4. RESTORE C:\ADSM.SYS\REGISTRY\NOGALES\USERS\JOE\JOE000.KEY\WINNT\
PROFILES\NTUSER.DAT.LOG
```

4.3.6 Shut Down and Reboot the System

Once the primary partition is restored and the registry files are copied the system can be shut down and rebooted. The system should first boot the NT loader and then start the system. If everything has been done correctly, the system will start correctly and be in the same state it was in after the last ADSM backup. In our case, as the second partition had yet to be restored, any references to the D: drive caused ″not found″ messages. If the system is an NT domain controller, it should be resynchronized with the backup domain controller.

4.3.7 Recover the Remaining Partitions

The previous sequence of actions recovers an NT system primary partition, typically the C: drive. If other logical volumes existed on the recovered disk, they now have to be recovered as well. This is a straightforward process once the NT production system is running and synchronized with the rest of the network.

To restore the second partition on our failed hdisk 0, we used the ADSM command line backup/archive client:

```
REStore -SUBDir=YES D:\*
```

After this restore, another reboot should find the system in the same state it was in at the last incremental backup.

4.4 Additional System Recovery

In this section we cover backup and recovery of network mounted drives, synchronization of the PDC and BDCs, and Windows NT workstation considerations.

4.4.1 Network Mounted Drives

The Microsoft NT Server provides a simple network file system (NFS). You can connect to the server; create a folder, which will be the file system; and assign a drive letter for the file system.

Note: We use the acronym NFS even though it is not a real network file system (as, for example, on UNIX machines); however, the idea of sharing file systems over the network is the same. The creation of an NFS and the mount process are different on the two machine architectures.

On the NFS you can grant and deny user rights as on a local disk. There are some restrictions in a multidomain environment and some new details for the backup and restore of these mounted drives.

4.4.1.1 Creation of an NFS

With the Windows Explorer from NT or Windows 95, you can search in the Network Neighborhood for the machine on which you want to create an NFS. It is not enough to select a machine and create a folder for sharing some files because when you start the machine the next time there will be no reconnection. The way to create an NFS is to:

1. Connect to the machine with the Explorer
2. Create a new folder
3. Set the permissions for this folder
4. Assign a drive letter to this folder
5. Check the field **Reconnect at logon** in the dialog window (see Figure 34)

The connection can be done through the Windows Explorer or directly on the command line. The advantage of using the command line is that you directly assign a drive letter to the connection. The disadvantage is that you cannot tell the operating system to reconnect this sharing every time you start the machine (NT can do this with the parameter /PERSISTENT:YES). The command can be, for example:

```
NET USE H: \\NOGALES\CDRIVE
```

The correct syntax of all NET commands can be read in the online help. In this example you connect a drive, CDRIVE, on machine NOGALES to drive letter H:. This information about the connection will be lost on the next start. You can include all NET commands in the local AUTOEXEC.BAT file for the reconnection.

Instead of creating a new folder you can use any existing folder on the server, but you must have control of the user rights for its directory to change any permissions. It is also possible that you would only want to use an existing folder that has another owner as an NFS drive (for example, a folder for a group that only has read access to this folder).

If you use Explorer to build the connection, the dialog window in Figure 34 appears.

Figure 34. Assigning the Drive Letter

From a Windows 95 machine you cannot see or set any permissions of an NFS drive. Only the general information about the drive is visible.

4.4.1.2 Permission Settings for NFS

After creating the NFS, you can grant or deny user rights for it from any NT machine. The principle is the same as described in 3.3, “The NT File System” on page 50, although the file system is on a remote machine. The default permissions for a new folder on the remote machine are:

Creator Owner Full Control (All)(All)

Everyone Change (RWXD)(RWXD)

PRANRES\Administrators Full Control (All)(All)

PRANRES\JOE Special Access (All)(Not Specified)

PRANRES\Server Operators Change (RWXD)(RWXD)

System Full Control (All)(All)

In this example we use PRANRES as the name of the domain and JOE as the name of the user who has created the new folder. User JOE must be a known user of the domain server and must have the special write permission for the shared directory on the server. Because the permissions are cumulative, user JOE has full control over the folder because he is also the Creator Owner of the folder.

In this context there can be a problem with setting the permissions. If user JOE is a member of domain PRANRES and a member of domain PRANDOM, for example, he can create the folder and assign the drive letter as described in 4.4.1.1, “Creation of an NFS” on page 83. However, when user JOE changes some of the permissions, it is possible that he will not be able to access any NFS files because of the different SIDs on the different domain controllers. User JOE on domain PRANRES is not the same user JOE on PRANDOM, even though he has the same name and password and is the same person.

The problem arises because the user creates the folder on domain server PRANDOM, but the domain controller of PRANDOM sets the default permission to PRANDOM\JOE, and not to PRANRES\JOE. Windows NT treats user JOE as a different user.

4.4.1.3 Backup and Restore Recommendations

After assigning the drive letter to the NFS, there is no problem with backing up the NFS with ADSM. The NFS drive must have a label like a normal physical disk.

In the ADSM backup client, the NFS appears as an ordinary disk, and you can select the drive for an incremental backup or a backup by tree. The only difference is the added underscore (_) and the letter of the assigned drive. For additional information, ADSM shows as the drive type the acronyms RMT for Remote and NTFS or HPFS386 for the file system type; for example:

```
H: CDRIVE_H:  RMT-NTFS
```

All files for which you have at least read permission are backed up. All other files with a No Access permission are not backed up and appear in the report as a failure with the text “Access denied.”

When the backup is completed, you can see the filespace of the ADSM server on the bottom of the ADSM backup client GUI window (see Figure 35 on page 85).

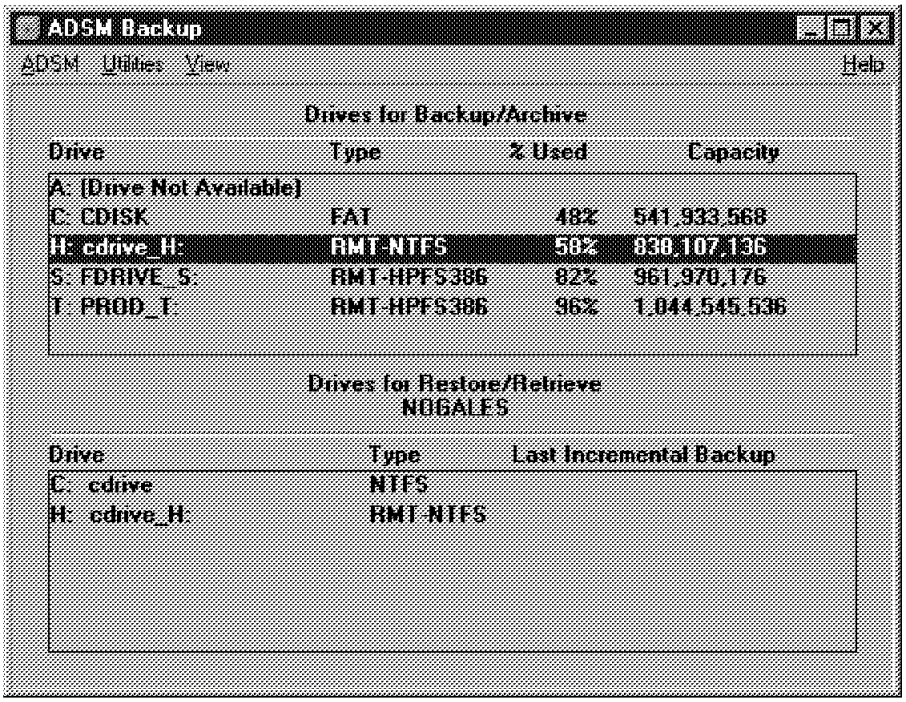

Figure 35. ADSM Backup of an NFS Drive

If you want to restore the file to the same location, the NFS must be mounted and have the same drive letter. Otherwise ADSM returns the error Access denied and aborts the current restore process.

4.4.2 Synchronization of the PDC and BDCs

In 4.3, “System Recovery (Bare Metal Restore)” on page 76 we describe the way to recover a Windows NT system after a disaster. In this section we explain the additional steps to synchronize the domain database, if you are using the principle of the BDC.

The PDC and one or more BDCs share and replicate the domain user database. If the PDC is damaged and loses some information from this database, it is easy to replicate the database with the BDCs. Because the work for a BDC is minimal, we recommend that you install as a minimum one BDC in your NT environment if you have more than one NT server. In this way the administrator can work with the domain database, even if the PDC is not available.

If the PDC is defective, you can start the Server Manager on the BDC and promote this machine to be the PDC of the domain. The BDC searches for the PDC, and, if it cannot find the machine, it automatically starts all necessary services. There is no need to stop any machine in the network, because the BDC starts the work as PDC immediately. To inform all other BDCs, you must invoke **Synchronize Entire Domain** in the **Computer** menu.

Note: There is no data replication between the PDC and the BDC. All shared drives or directories of the lost PDC are inaccessible. The principle of the PDC and BDC is not a mirror concept based on the file system.

When the PDC is recovered and you start the machine, you will have two PDCs for the same domain. During the boot process, the recovered PDC remarks on this and returns an error message. However, the machine starts successfully,

and you must promote it as the only PDC in the environment. The current PDC, which was the BDC before, accepts this request if you are the administrator and you degrade this machine again to the BDC service. Synchronize the entire domain again, and the environment will work as before.

4.4.3 Windows NT Workstation

To recover an NT workstation, use the procedure described in 4.3, “System Recovery (Bare Metal Restore)” on page 76.

The NT workstation can be treated like the NT server for disaster recovery backup and recovery, but many problems experienced with NT server recovery are not encountered with the NT workstation. The main difference between the two is that the NT workstation does not support the domain concept. All created users or accounts are local to the machine, so it is not possible to log on at an NT workstation over the network from another machine.

Both the NT workstation and server have a problem with the permission set and the registry. The security mechanism is the same for both the workstation and server, and both exhibit the same stable behavior.

Because there is no domain database, you cannot install programs with client support, such as Microsoft′s SQL Server, Microsoft′s Exchange, or Lotus Domino, on an NT workstation.

4.5 ADSM Server on Same System As Data

The examples in this book typically have the backup data on a server that is remote to the system to be recovered. However, some simple or entry-level installations install their ADSM server on the same machine as the data to be backed up. Although this arrangement provides a degree of protection, mostly in versioning, there are clearly some issues to consider to obtain a reasonable level of availability. Carefully read the redbook *ADSM Server for Windows NT Configuration and Recovery Examples* (SG24-4878) for details on how to back up an ADSM server to removable or remote media. With ADSM Version 3 server-to-server communication, a remote server can be used for backups of storage pools, databases, and other files.

Critical Windows NT machines should be protected with additional software or hardware extensions, as described in 3.2.4, “Additional Disk Protection” on page 48. A RAID, stripe set, or stripe set with parity protects against any hardware failure of the disks, but the software availability implementations can also degrade performance.

4.6 Migration of a Windows NT System

In this section we explain how to migrate an existing NT server to another machine or to a bigger physical disk and discuss some areas to watch out for when using this approach.

4.6.1 Reasons for Migration

The time needed for the NT system install procedure is the shortest of all installations on an NT machine. However, if the machine has been running for some time, there are many applications and user settings on the machine as well as a lot of distributed data. Reinstalling this information is time consuming, so we describe how you can use ADSM to migrate an NT system with all applications and data.

With NT you can copy or restore a whole system to another partition because all applications are installed using the ARC name convention. Changing the system from the C: drive to the D: drive will not cause a problem. The only thing to change is the new-location entry in the BOOT.INI file. It is not possible to migrate the NT system into another directory. The %SystemRoot% variable of the NT system always points to a set partition on a set disk. Although you can change the variable in the BOOT.INI file, some applications store the variable in the registry, using a method that conflicts with the change in the BOOT.INI file.

You can migrate your NT system if:

- The partition or disk has no free space.
- You install an additional operating system.
- You replace the old disk with a newer and bigger one.
- You replace the whole computer.

Make sure you have the newest backup of the whole system before you install additional disks or replace hardware. The migration depends on where you want to restore the image of the system, so we present the procedure for two cases. You want to migrate the current NT system to another place, and:

- The current NT system is still available.
- The current NT system is no longer available.

For both cases we recommend that you have all installation media ready for use, because in some circumstances NT asks for the original installation CD.

4.6.2 Backup and Migration within the Same Machine

The prerequisite for this case is a (full) backup of the NT system, including the registry and all profiles. This backup must be done by the system administrator, who is the only person allowed to back up the registry. The connection to the ADSM server must be available at all times. The location of the ADSM client does not matter for this case; it could be locally started through a network.

Steps for backup and migration:

1. Ensure that the environment contains one NT server with an internal physical disk and has one existing NT system on the startable partition.
2. Migrate the NT system to another partition or a new installed disk (see Figure 36 on page 88).
   a. Be aware that the new partition or disk must be as a minimum the same size as the current NT system.
   b. Restore the NT system with the DSMC RES command including the restore and copy of the registry hives and the users profile.
   c. Change or add the entries in the BOOT.INI file.

If you have migrated the NT system to a new disk, the entry in the registry for the disk parameter is incorrect. However, this is not a handicap to successful migration.

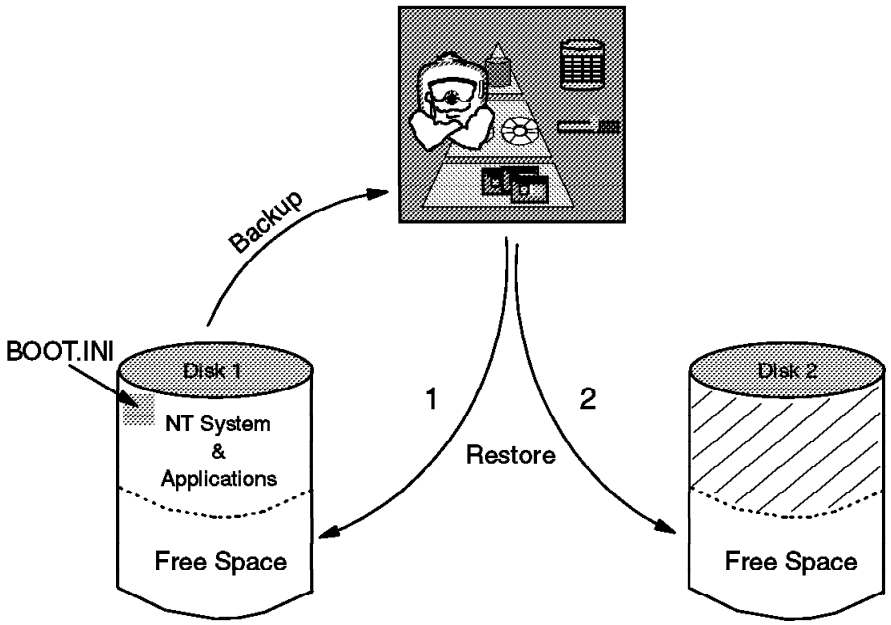

Figure 36. Migration of the NT System within the Same Machine

Note: This migration is based on the principle described in 3.5, “Using ADSM to Create Repair Partitions” on page 64.

After rebooting the machine you can select the new partition. If you include all applications in this backup and restore process, you can delete the old partition without any problems. Leave the old partition on the machine until you have finished the test of the new partition.

4.6.3 Migration to Another Machine

The prerequisite for this case is a (full) backup of the NT system, including the registry and all profiles. The location of the ADSM server does not matter.

Changing the machine during migration can cause many problems. We recommend that you do not change the main architecture of the hard disk controller. If the current NT system is on a machine with an IDE controller and you want to migrate to a machine with a SCSI controller, you must change all entries in the BOOT.INI file, add the SCSI device driver (see 3.4, “Configuring NT for Availability” on page 55), and possibly solve all resulting problems from the registry entries.

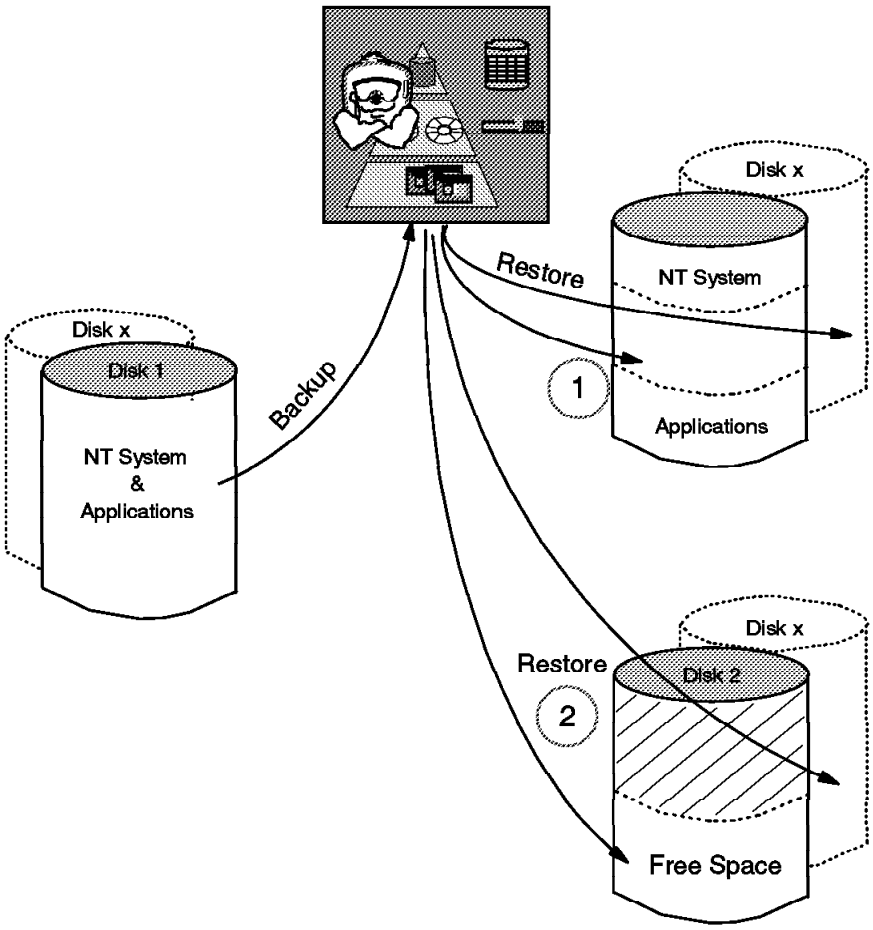

Figure 37. Migration of the NT System to Another Machine

In Figure 37 you can see the possible migration directions. The target machine can either have an NT system or be empty. In the examples that follow we assume that you do not change the hard disk architecture.

4.6.3.1 Migration As Additional NT System

Steps for the migration:

1. Make sure the new partition or disk is at least the same size as the backed up NT system.
2. Boot from the target machine's NT system.
3. Install and connect the ADSM client to the ADSM server.
4. Restore the NT system with the DSMC RES command
5. Add an entry in the BOOT.INI file for the newly installed NT system and add one entry for the same partition with the /sos /basevideo parameter.
6. Restart the machine.
7. Boot the new partition with the [VGA mode].
8. Some error messages regarding network connection problems appear; disregard them.

9. When the NT system has finished the startup, log in as the administrator of the backed up NT system.
10. You must change all settings about the network (for example, IP address and gateway), the VGA card, and additional cards (such as sound, modem, or I/O cards). Ideally you have all this information on a printout from the already installed NT system of the target machine.
11. Log off and reboot the machine.
12. Test the system and the applications.

In most cases the migration completes without any problems, but there is no guarantee. Test the environment carefully before you start the new NT system as the production system.

Another way to handle the hardware changes in the registry is to export the hkey_local_machine hive in the registry of the target machine.

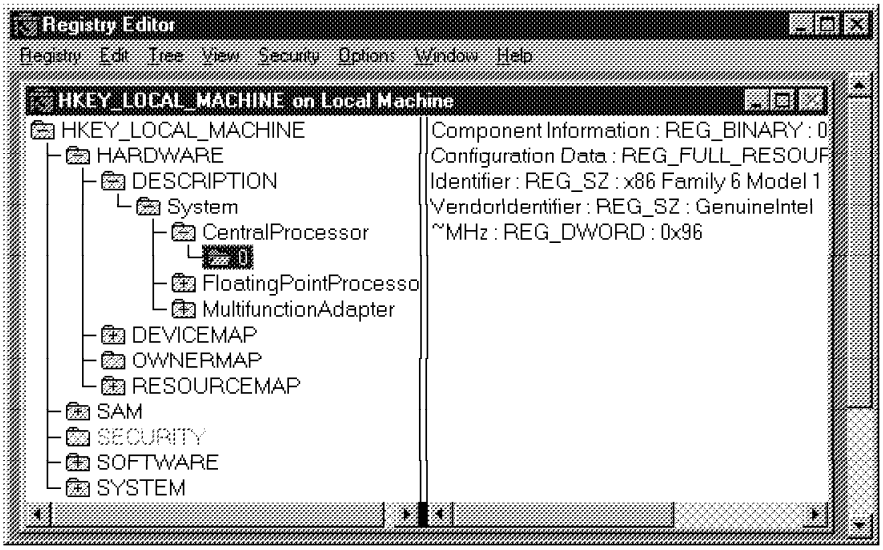

Figure 38. Registry Hive for the Local Machine

After the restore with ADSM, you start the new partition as described and import the registry hive into the registry of the new NT system. For this approach you must have a thorough knowledge of the registry components. The hive of the local machine contains all security and software information, which is different from that of the target machine.

4.6.3.2 Migration to an Empty Machine

Steps for the migration:

1. Make sure that the new partition or disk is at least the same size as the backed up NT system.
2. You must install a minimum NT system and an ADSM client to do the restore. You must have the ADSM information about the node name, the server name, and the connection used. When you install the NT system, be sure to read all machine information.
3. Boot from the NT system of the target machine.
4. Restore the backed up NT system with the DSMC RES command

5. Add an entry in the BOOT.INI for the newly installed NT system and add one entry for the same partition with the /sos /basevideo parameter.
6. Restart the machine.
7. Boot the new partition with the [VGA mode].
8. Some error messages regarding network connection problems appear; disregard them.
9. When the NT system has finished the startup, log in as the administrator of the backed up NT system.
10. You must change all settings of the network (for example, IP address and Gateway), the VGA card, and additional cards (such as sound, modem, or I/O cards). Ideally you have all this information on a printout from the already installed NT system of the target machine.
11. Log off and reboot the machine.
12. Test the system and the applications

Part 3. Application Recovery

Chapter 5. Recovering Lotus Notes

In this chapter we describe the structure of the Lotus Notes database. We briefly discuss how you can use ADSM to back up and recover the basic structure (.nsf files) and individual notes through the ADSMConnect Agent for Lotus Notes.

Note: This topic is covered in more detail in the redbook *Using ADSM to back up Lotus Notes* - SG-24-4534-01, available in HTML on the ITSO web pages, www.redbooks.ibm.com .

5.1 Data Characteristics

Lotus Notes files are stored by default in the directory structure shown in Figure 39.

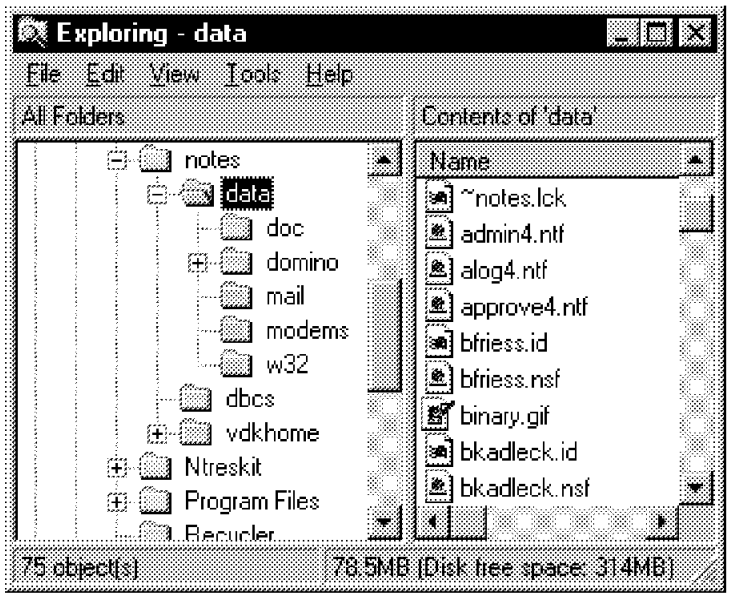

Figure 39. Lotus Notes Directory Tree

It is this basic structure that you must protect from data loss. The .nsf files shown in Figure 39 are the databases of Lotus Notes documents. The .id files are the userid files for each authorized Lotus Notes user. They contain an encrypted form of the logon password, along with other important user data. Users may choose to keep a copy of the .id files on a diskette, so that they can gain access to their data from a PC other than their base machine. They also may carry a copy of the .dsk files, enabling them to re-create their desktop at the remote machine.

The fact that users may have copies of these files in no way lessens the need for you to provide adequate backups of the data.

5.2 Using ADSM to Back Up Lotus Notes Data

You use the basic ADSM backup/archive client to back up, archive, restore, and retrieve an entire Lotus Notes database (see the .nsf file structure described in 5.1, “Data Characteristics” on page 95), and in most cases this will be the fastest option. There is, however, a good case for using both the basic ADSM backup/archive client and the ADSM Lotus Notes agent to perform backup, because they provide different, overlapping functions:

- The Lotus Notes agent can restore individual notes within a database (.nsf), but it is slower when used to restore the entire database
- The basic ADSM client cannot back up an incremental copy of a database (.nsf), but it is faster to use for backup and restore of an entire database file.
- The Lotus Notes agent has no archive function.
- The Lotus Notes agent has its own ADSM options file and should have its own node name. Thus you can provide different management classes for the incremental backups produced by the agent.
- Both the Windows NT server and the Lotus Notes agent have different ADSM file spaces, so there is no correlation within ADSM of backups produced by the Lotus Notes agent and the basic ADSM client. Therefore you cannot restore a database (.nsf) file with the agent if it was backed up by using the ADSM client.

5.2.1 Installing the ADSMConnect Lotus Notes Agent

Follow the installation instructions provided with your ADSMConnect Lotus Notes agent. Some points worth noting:

- You must include the path for the Lotus Notes dynamic link library (DLL), nnotes.dll, and the path for the ADSMConnect agent DLLs, dsmnoted.dll, adsm32.dll, and notesdlg.dll, in your PATH environment variable.
- You must include user variables to set DSMI_CONFIG, DSMI_DIR and DSMI_LOG. See the example in Figure 40 on page 97.

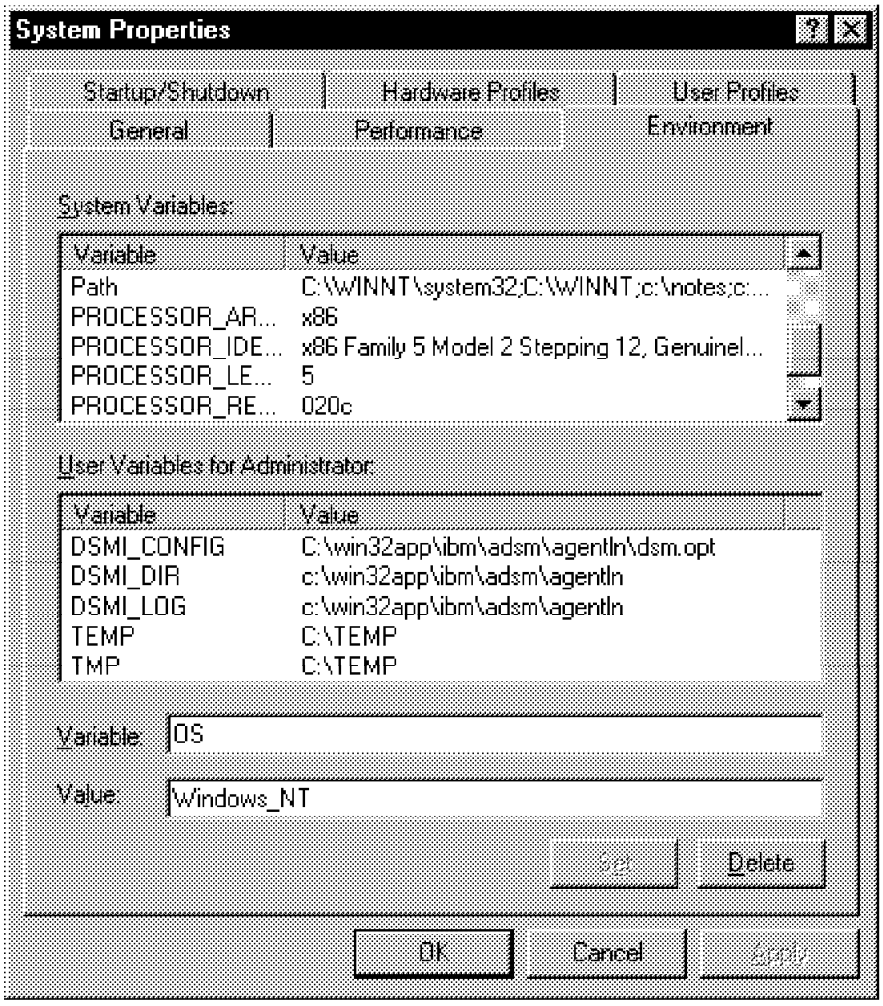

Figure 40. Environment Variables for ADSMConnect Lotus Notes Agent

When you install the ADSMConnect Lotus Notes agent, three new functions appear on the Actions pull-down in the Lotus Notes client window. Figure 41 shows a portion of a Lotus Notes client, with the functions added by the agent.

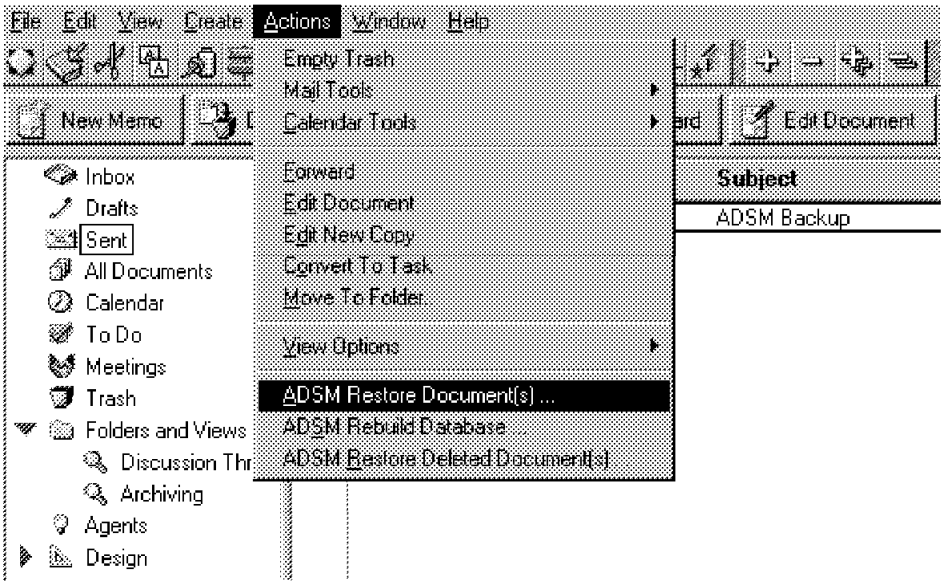

Figure 41. Lotus Notes Client: Actions Pull-Down Menu

Note: In earlier versions of Lotus Notes, the added functions appear on the Tools pull-down menu.

5.2.2 Using the ADSMConnect Lotus Notes Agent

To back up all Lotus Notes databases, issue the following command:

```
DSMNOTES INCR x:\notes\data\* -SUBDIR=Y
```

The above example assumes that you have your .nsf files in subdirectories of notes\data. If all .nsf files are in the notes\data, you can omit the optional -SUBDIR=Y stanza. You can opt to back up a single .nsf file by substituting its name in place of the wildcard (*) character in the command. There are other options available which are outside the scope of this book. Please consult the *ADSMConnect for Lotus Notes User Guide* for further information and full syntax.

Note: See A.2.1.2, “Lotus Notes Include/Exclude Sample File” on page 161 for an example of the include/exclude list syntax for the ADSMConnect Lotus Notes agent dsm.opt file.

During dsmnotes execution, progress is shown in the command window as follows:

```
 .
 .
Backing up database c:\notes\data\bfriess.nsf
Documents sent: 16
Documents sent: 32
Documents sent: 48
Documents sent: 64
Documents sent: 80
Documents sent: 96
Documents sent: 112
Documents sent: 122
Backing up database c:\notes\data\bkadleck.nsf
Documents sent: 16
Documents sent: 32
Documents sent: 48
```

```
Documents sent: 64
Documents sent: 80
Documents sent: 96
Documents sent: 112
 .
 .
```

Statistics from the DSMNOTES command are sent to a file named log.dsm. The following is an example of the log.dsm file's contents:

```
08/13/97 04:03:32 PM  dsmnotes incr start
 .
 .
08/13/97 04:03:33 PM  Database: c:\notes\data\bfriess.nsf
08/13/97 04:03:39 PM  Database: c:\notes\data\bkadleck.nsf
08/13/97 04:03:46 PM  Database: c:\notes\data\busytime.nsf
 .
 .
08/13/97 04:04:05 PM  Database: C:\NOTES\DATA\BFRIESS.NSF
08/13/97 04:04:05 PM  Documents backed up: 122, documents deleted: 0
08/13/97 04:04:05 PM  Database: C:\NOTES\DATA\BKADLECK.NSF
08/13/97 04:04:05 PM  Documents backed up: 121, documents deleted: 0
08/13/97 04:04:05 PM  Database: C:\NOTES\DATA\BUSYTIME.NSF
08/13/97 04:04:05 PM  Documents backed up: 6, documents deleted: 0
 .
 .
08/13/97 04:04:05 PM  Total number of documents backed up: 602
08/13/97 04:04:05 PM  Total number of documents deleted: 0
08/13/97 04:04:05 PM  Total number of bytes sent: 6255.6 KB
08/13/97 04:04:05 PM  ANS0900I dsmnotes completed
```

You can manage data from the Lotus Notes agent with the basic ADSM administration client, by manipulating the management class for the client named in the &connect. Lotus Notes agent dsm.opt file.

5.3 Restoring Lotus Notes Documents

Figure 41 on page 98 shows the new functions added to the Actions pull down menu by the ADSMConnect Lotus Notes agent installation:

1. ADSM Restore Document(s) ...
2. ADSM Rebuild Database ...
3. ADSM Restore Deleted Document(s) ...

We discuss each of these in turn.

5.3.1 Restoring Documents from the ADSM Server

Selecting **ADSM Restore Document(s) ...** starts the dialog shown in Figure 42 on page 100.

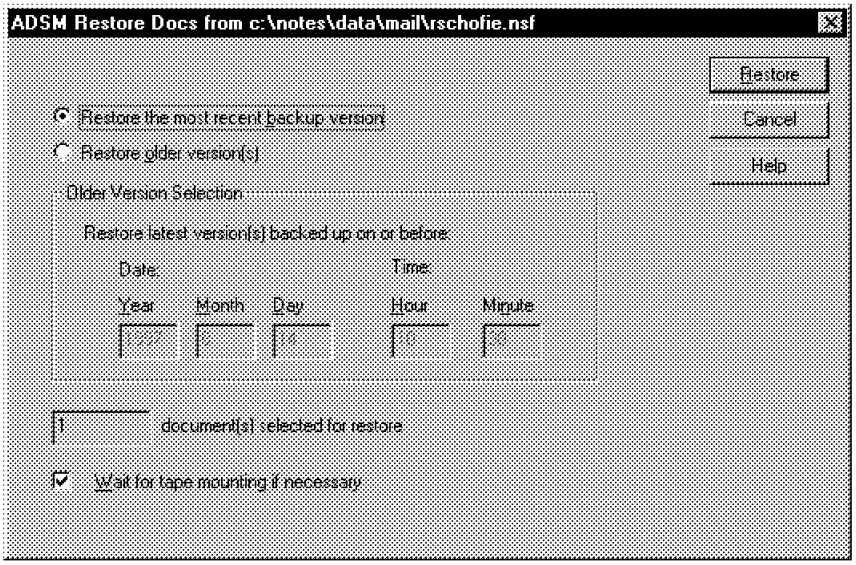

Figure 42. Dialog for Restoring Documents

From this dialog box, you can choose to restore the most recent copy of the ADSM backup (the default) or an older version. If you select **Restore older version(s)**, ADSM restores those documents that were backed up on or before the date and time you supply in the **Older Version Selection** boxes.

5.3.2 Rebuilding the Database

In the Rebuild Database dialog (Figure 43 on page 101) you can restore an entire database, either to a new file database or merged into an existing database. It is a simple matter of supplying the required information in the dialog box.

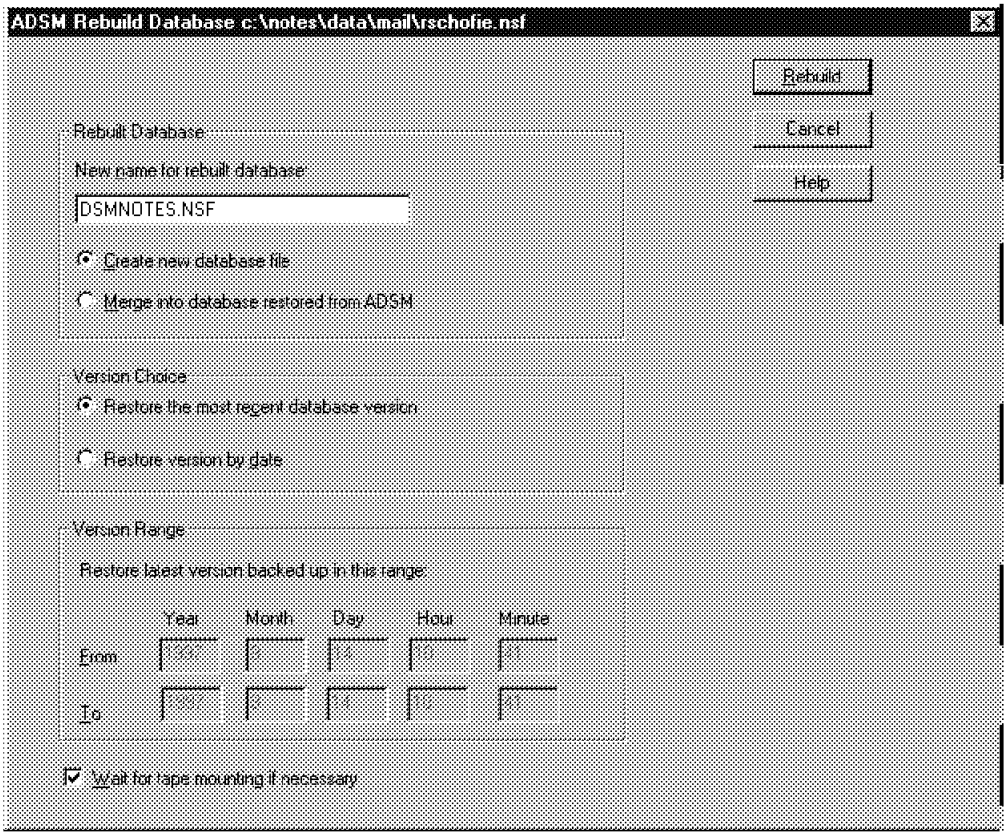

Figure 43. Dialog for Rebuilding Database

5.3.3 Restoring Deleted Documents

Restoration of deleted documents relies on another function of Lotus Notes. When you delete a document from the database, it is not immediately deleted but is turned into a stub file. This stub usually is not viewable. At a time specified by Lotus Notes administration options, the stub files are purged from the database. You can use the dialog shown in Figure 44 if the stubs still exist (that is, they have not yet been purged by Lotus Notes). If a purge has taken place, you must first restore a copy of the database taken before the purge (5.3.2, “Rebuilding the Database” on page 100).

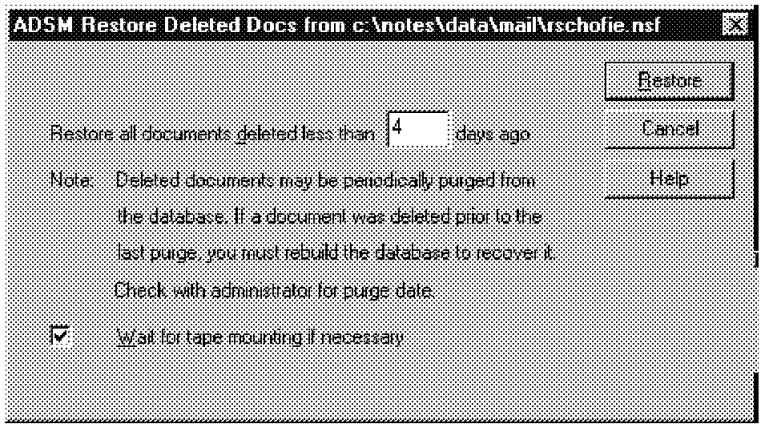

Figure 44. Dialog for Restoring Deleted Documents

Here, the only setting required is the number of days for deleted documents to be restored. If you enter a zero in this field, all notes deleted but not purged will be restored.

Chapter 6. Microsoft SQL Server Backup

The relational database management system (RDBMS) SQL Server is totally embedded in the NT system architecture. The server software runs as an NT service, communicates directly with the NT kernel, and uses the registry extensively for storing system information.

In this chapter we briefly describe the structure of the SQL Server from the perspective of the backup operator. We assume that you are familiar with the SQL Server work approach.

The backup of the SQL Server can be done online or offline. For the offline backup you must shut down the SQL Server. Offline backup is the securest way of performing the database backup, because there are no active transactions to get lost. In most companies, however, shutting down the SQL Server is not feasible.

We show the following offline and online backup methods:

1. Offline
   - Using the ADSM backup client
2. Online
   - Using an SQL dump to disk
   - Bulk copy of the database devices
   - Using the ADSM/SQL Agent

For additional information about the SQL Server, refer to the SQL Server documentation or any of the commercial publications.

6.1 SQL Server DBMS Structure

Before you can back up and restore databases in the SQL Server, you must know where the information is stored. Figure 45 on page 104 shows the key components of the SQL Server system. The components that you must consider for backup and recovery are:

- The system database devices (Master, MSDBData, MSDBLog)
- The user-defined database devices
- Backup dump devices

The SQL Server uses a system variable for the home directory of the installation location. The %SqlRoot% variable has the same function as the %SystemRoot% variable of the NT system. During the setup the system databases are created in the `%SqlRoot%\DATA` directory. The initial backup device, diskdump, is located in the `%SqlRoot%\BACKUP` directory.

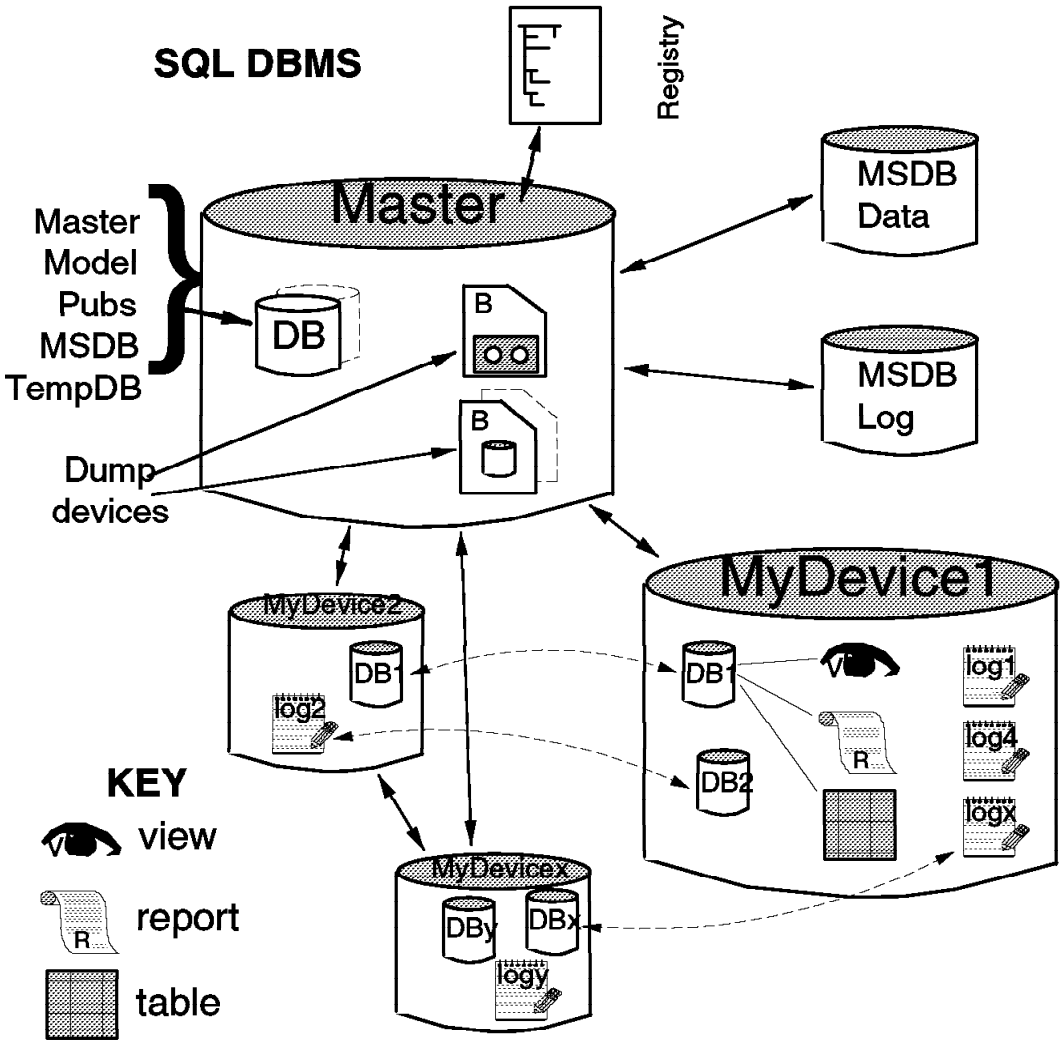

Figure 45. SQL Server DBMS Structure

A database device is a logical unit with a predefined size on a physical disk. When you define a new database device, the SQL system asks for the size of the device and the location. We recommend that you locate all database devices in the same folder on the disk to minimize the administrative and backup work. All information about the name and the size of the device must be recorded on paper, because you need this information in case of a disaster. It is a good idea to record all structure information and all changes in the design.

A database in the SQL Server consists of tables, views, rules, and procedures. Every user database contains all of the user data as well as a subset of the system tables with information specific to that database, such as allowed users and/or groups.

In the internal storage of one database, the SQL Server uses *segments*, another key component of the DBMS, which are subsets of the database. Each database can have up to 32 segments. In Figure 45, one segment of database DB1 is located in database device MyDevice1, and another segment of database DB1 is located in MyDevice2. Segments provide a flexible way of assigning objects to particular database devices.

When you create a database device and have to predefine the size, the size for the included databases is limited. The size could be changed later, but for various reasons, such as performance or security, it is better to locate parts of a database on another drive.

Each database contains transaction logs, which are really system tables called *syslogs*. Syslogs automatically record every transaction issued by each user of the database. By default, the transaction logs for a database are stored on the same device as the database. You can store the log file of a database in the same database device or in another device. In Figure 45 on page 104 the transaction log of database DB1 is local on the same database device, MyDevice1. But the transaction log of database DB2 is located on MyDevice2. Run the sp_helpdb script to find out whether your logs are on the same device as your database. If they are, you do not have to back them up separately.

The distributed location of the database segments and the transaction logs could cause some problems with the backup.

All system information is stored in internal databases, which are located in the master database device. In Figure 45 the databases in the master database device are:

- Master
- Model
- Pubs
- MSDB
- TempDB

The master database controls the user databases and the operation of the server as a whole. It contains system tables that keep track of server information, such as users, databases, storage space, and locks. System tables can also be thought of as the data dictionary or system catalog. The model database is used as a template for creating user databases. The Pubs database includes a simple demo database with tables and views.

Internal objects and procedures are stored in the MSDB. The TempDB is comparable to a temporary folder, in which everyone can store temporary data.

The SQL system makes heavy use of the NT resources. Use of the NT registry is not limited to the install process of the program. Important system information is stored and changed during the run time of the SQL Server. When you back up the SQL Server, you must always have a backup of the registry; otherwise the SQL Server becomes inoperable.

After the installation of the SQL Server there is only one backup dump device, diskdump. You can create more backup devices in the SQL Server for the online backup of the databases. The storage medium could be a disk or tape. If you select a disk, you must install and name a file. We recommend that you locate the file in the `%SqlRoot%\BACKUP` directory because all backup devices are in this directory.

6.2 Offline Backup with the Backup/Archive Client

Our example shows how to back up the SQL Server databases and logs directly to ADSM. The SQL Server locks its database devices while it is running, so ADSM cannot access the associated database files while the SQL Server is up and running. To use this type of backup, the SQL Server must be stopped first (see Figure 46).

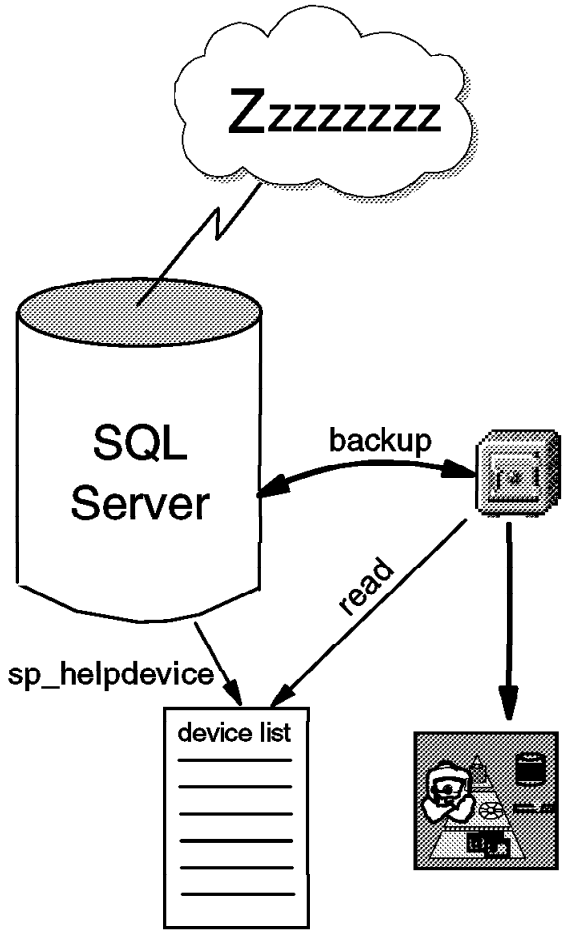

Figure 46. Offline Backup Using ADSM

The advantage of this approach is that no additional filespace is required to store intermediate dump files. One disadvantage of this approach is that the SQL Server must be shut down to perform either a backup or restore operation. Another disadvantage is that a particular database device may contain more than one logical database or only a portion of a database. Care must be taken with this approach to back up the correct combination of files in order to properly restore their associated database.

6.2.1 Preparatory Steps

First you must determine which files the SQL Server is using as database devices. Use the "sp_helpdevice" stored procedure:

```
isql -U SA -Q "exec sp_helpdevice" -P
```

This stored procedure lists all devices defined to the SQL Server:

```
device_name    physical_name                    description
-------------- ------------------------------   -----------
diskdump        nul                              disk, dump device
diskettedumpa   a:sql\table.dat              diskette, 1.2 MB,
                                                 dump device
diskettedumpb   b:sql\table.dat              diskette, 1.2 MB,
dump device
dumpdev1        C:\MSSQL\BACKUP\dumpdev1.DAT disk, dump device
dumpdev2        C:\MSSQL\BACKUP\dumpdev2.DAT disk, dump device
dumpdev3        C:\MSSQL\BACKUP\dumpdev3.DAT disk, dump device
master          C:\MSSQL\DATA\MASTER.DAT     special, default
                                                 disk, physical disk, 25MB
MSDBData        C:\MSSQL\DATA\MSDB.DAT       special, physical
                                                 disk, 6 MB
MSDBLog         C:\MSSQL\DATA\MSDBLOG.DAT    special, physical
                                                 disk, 2 MB
Sample          C:\MSSQL\DATA\Company.DAT    special, physical
                                                 disk, 7 MB
```

You are concerned only with the devices with a description of ″special.″ These are the devices that the SQL Server uses to hold the databases. The ″physical_name″ column contains the actual NTFS or FAT file that represents the database device. In this example, the following four files contain SQL Server databases:

- C:\MSSQL\DATA\MASTER.DAT
- C:\MSSQL\DATA\MSDB.DAT
- C:\MSSQL\DATA\MSDBLOG.DAT
- C:\MSSQL\DATA\Company.DAT

You also need to know which databases are in each database device. Use the ″sp_helpdb″ stored procedure for each database. An example of this command and its output are shown later in 6.3, “Using the SQL Dump” on page 110. Because we are backing up files and not databases with this approach, it is important that a database be restored in its entirety when the database devices are restored. For instance, in this example, the C:\MSSQL\DATA\MSDB.DAT file cannot be restored without also restoring the C:\MSSQL\DATA\MSDBLOG.DAT file at the same time. You should also be concerned with database devices that contain more than one database. The process of restoring a single database will not be possible in that case.

6.2.2 Backing Up the SQL Server Database Devices to ADSM

Follow these steps:

1. Make sure that no one is using the SQL Server. Use the ″sp_who″ stored procedure:

```
isql -U SA -Q ″exec sp_who″ -P
```

The SQL Server returns the following table:

```
spid   status     loginame hostname    blk   dbname     cmd
------ ---------- -------- ----------- ----- ---------- ----------------
1      sleeping   sa                     0   master     MIRROR HANDLER
2      sleeping   sa                     0   master     LAZY WRITER
3      sleeping   sa                     0   master     CHECKPOINT SLEEP
4      sleeping   sa                     0   master     RA MANAGER
13     runnable   sa       LARSON_SERVER 0   master     SELECT
```

2. Stop the SQL Server. Use the net stop command for both the SQLExecutive and MSSQLServer services:

```
C:\> net stop SQLExecutive
The SQLExecutive service is stopping.
The SQLExecutive service was stopped successfully.

C:\> net stop MSSQLServer
The MSSQLServer service is stopping.
The MSSQLServer service was stopped successfully.
```

3. Now the SQL Server is completely shut down and the database devices can be backed up with the ADSM backup/archive client. Selectively back up each individual file. Use the following ADSM commands, which can be shortened by using a macro or by effectively masking the file names:

```
dsmc sel C:\MSSQL\DATA\MASTER.DAT
dsmc sel C:\MSSQL\DATA\MSDB.DAT
dsmc sel C:\MSSQL\DATA\MSDBLOG.DAT
dsmc sel C:\MSSQL\DATA\Company.DAT
```

4. Start the SQL Server, using the net start command.

```
C:\> net start MSSQLServer
The MSSQLServer service is starting..
The MSSQLServer service was started successfully.

C:\> net start SQLExecutive
The SQLExecutive service is starting.
The SQLExecutive service was started successfully.
```

6.2.3 Restoring SQL Server Database Devices from ADSM

An individual database can be restored by restoring all database devices that contain the database components. Ensure that these database devices are restored together. In the example above, the MASTER database is on a single database device, and the MSDB database has its data on one device and its log on another. To restore the MASTER database, you simply need to restore the device on which it resides. The MSDB database must have both devices restored together.

In the event of a media failure or total system failure, you must take into account other considerations beyond simply restoring database devices. Review the Microsoft documentation to determine the additional actions that may be required.

Follow these steps to restore the SQL Server database device from ADSM:

1. Be certain that the SQL Server is stopped. Use the net stop commands as indicated in step 2 of 6.2.2, “Backing Up the SQL Server Database Devices to ADSM” on page 107.
2. Restore the required database device or devices to their original location.
3. Restart the SQL Server, using the net start commands, as indicated in step 4 of 6.2.2, “Backing Up the SQL Server Database Devices to ADSM” on page 107.
4. The following batch file can be used to completely restore the SQL Server in our sample environment:

```
net stop SQLExecutive
net stop MSSQLServer
dsmc res -rep=yes C:\MSSQL\DATA\MASTER.DAT
dsmc res -rep=yes C:\MSSQL\DATA\MSDB.DAT
dsmc res -rep=yes C:\MSSQL\DATA\MSDBLOG.DAT
dsmc res -rep=yes C:\MSSQL\DATA\Company.DAT
net start MSSQLServer
net start SQLExecutive
```

Restoring the database devices can cause a problem. If you need to restore two database devices for the restore of one database, and the second database device includes a part of another database, this additional database could be damaged. In Figure 47 database DB2 is located in database devices MyDevice1 and MyDevice2. The transaction log of database DB2 is stored in MyDevice2. If you want to restore database DB2, the above commands show that a restore of database devices MyDevice2 and MyDevice2 is enough. But when you restore these two database devices, database DB3 will be damaged because only parts of it are restored. The transaction log of DB3, located in MyDevice3, is not usable for the rebuild of the transactions for DB3.

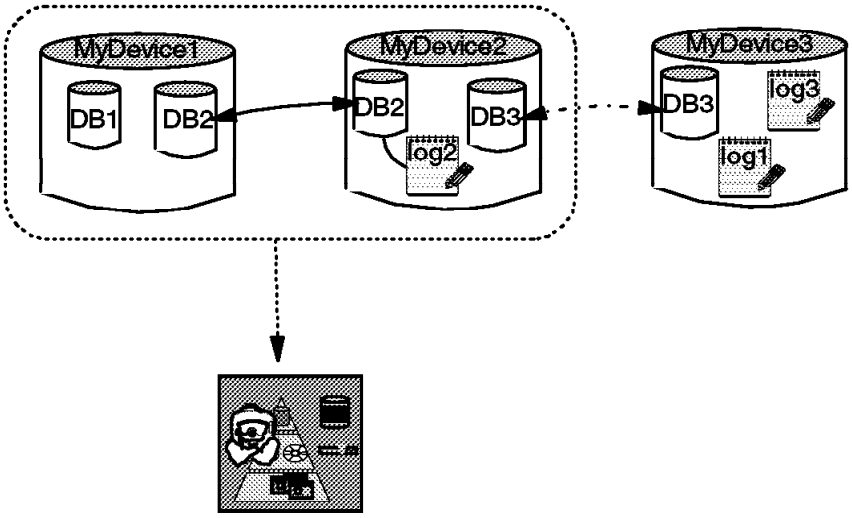

Figure 47. Restore of Database Devices

6.3 Using the SQL Dump

Our example shows how to back up and restore the SQL Server databases with ADSM online. The SQL Server provides both a dump utility and a load utility, which can work with external files it calls ″dump devices.″ These utilities must be run while the SQL Server is running. Once the dump utility is run, you can use ADSM to back up the output of the operation without stopping the SQL Server. Restore operations can also be run while the SQL Server is running. The SQL Server does not keep its dump devices locked once the dump utility is complete, so ADSM can either back up or restore these files while the SQL Server is up and running.

The major advantage of this approach is that no knowledge of the underlying file structure of the SQL Server is necessary, and individual databases can be dumped and loaded independently. Another advantage of this approach is that both the backup and restore operations can take place without shutting down the SQL Server. One further advantage is that the SQL Server itself keeps track of the dump activity in the master database, and a GUI is provided for assistance in the recovery operation. The disadvantage of this approach is that additional file space is required to store intermediate dump files.

6.3.1 Preparatory Steps

First you must determine which files the MS SQL Server is using as dump devices. Use the sp_helpdevice stored procedure:

```
isql -U SA -Q ″exec sp_helpdevice″ -P
```

The stored procedure lists all devices defined to the SQL server:

```
device_name    physical_name                 description
-------------- ----------------------------- -----------
diskdump        nul                          disk, dump device
diskettedumpa   a:sql\table.dat              diskette, 1.2 MB, dump
                                                 device
diskettedumpb   b:sql\table.dat              diskette, 1.2 MB, dump
                                                 device
dumpdev1        C:\MSSQL\BACKUP\dumpdev1.DAT disk, dump device
dumpdev2        C:\MSSQL\BACKUP\dumpdev2.DAT disk, dump device
dumpdev3        C:\MSSQL\BACKUP\dumpdev3.DAT disk, dump device
master          C:\MSSQL\DATA\MASTER.DAT     special, default disk,
                                                 physical disk, 25 MB
MSDBData        C:\MSSQL\DATA\MSDB.DAT       special, physical disk,
                                                 6 MB
MSDBLog         C:\MSSQL\DATA\MSDBLOG.DAT    special, physical disk,
                                                 2 MB
Sample          C:\MSSQL\DATA\Sample.DAT     special, physical disk,
                                                 7 MB
```

You are concerned only with the devices with a description of ″disk, dump device.″ These are the devices that the SQL Server uses to contain the database backups. In this example, the following three files can be used to contain SQL Server databases and log backups:

- C:\MSSQL\BACKUP\dumpdev1.DAT
- C:\MSSQL\BACKUP\dumpdev2.DAT

- C:\MSSQL\BACKUP\dumpdev3.DAT

The next step in this approach is to determine the type of dumps that can be made with the database dump. Only a database that has its data and log stored separately can use the ″dump transaction″ format of the dump command. To determine if a database can use the dump transaction command, use the sp_helpdb stored procedure:

```
isql -U SA -Q ″exec sp_helpdb master″ -P
```

The resulting output is:

```
name              db_size       owner
----------------- ------------- ------
master                 17.00 MB sa

device_fragments  size          usage
----------------- ------------- -------------
master                  3.00 MB data and log
master                 14.00 MB data and log
```

Note the comments in the usage column. It states that the data and log are stored together. Therefore the master database, or any other database organized in this way, must be completely dumped and cannot use the dump transaction command. Here is the result of issuing the sp_helpdb command against a database that stores its log separately from its data:

```
name              db_size       owner
----------------- ------------- ------
msdb                    8.00 MB sa

device_fragments  size          usage
----------------- ------------- -------------
MSDBData                6.00 MB data only
MSDBLog                 2.00 MB log only
```

6.3.2 Full Backup of SQL Server Databases to ADSM

All databases can be backed up in their entirety through the dump command, although there are various forms of the dump command. Before a complete backup of a database, it is useful to use the ″dump transaction″ command with the ″truncate_only″ option to clear the log of inactive entries. To use the dump feature to back up a database, issue the dump database command:

```
C:\mssql\backup>isql -U SA -Q ″DUMP DATABASE SampleDB
                              TO DUMPDEV1 WITH INIT″ -P
Database ′SampleDB′ (76 pages) dumped to file <1> on device ′DUMPDEV1′.

C:\mssql\backup>
```

The WITH INIT option on the command creates the dump as the only logical dump within this dump device. Without this option, or with the WITH NOINIT

option, new dumps are appended to the file. If this is the first dump to this dump device, you must use the WITH INIT option; otherwise the SQL Server returns an error message indicating that the device is not ready. Once the command completes, ADSM can either selectively or incrementally back up the file housing DUMPDEV1.

6.3.3 Incremental Backup of SQL Server Databases to ADSM

The dump transaction command dumps only the log portion of the database. The output of the dump transaction command is typically much smaller than a full dump database command:

```
C:\mssql\backup>isql -U SA -Q "DUMP TRANSACTION SampleDB TO DUMPDEV1" -P
Database 'SampleDB' log(4 pages) dumped to file <2> on device
'DUMPDEV1'.
```

In our example, the WITH NOINIT option is used. Thus the output of the dump utility is appended as file 2 in the dump device. Dump device DUMPDEV1 now contains both a full dump and a transaction dump. Additional transaction or database dumps can be appended to the dump device in this way. For small or medium databases, keeping all of the dumps together may be a good option. However, for large databases, placing a database dump on one dump device and all transaction dumps on another minimizes the amount of data that ADSM has to back up.

Transaction dumps can be restored only after the restore of the previous database dump and all other intervening transaction dumps. Therefore, all transaction dumps since the last database dump must be made available to fully restore up through the last transaction dump.

6.3.4 Restoring the SQL Server Database Devices with ADSM

Typically, most database restores will be made from the dump devices themselves. If one database is lost, you cannot restore the database by using the information about the dump that is stored in the master database. You must restore the database directly from the dump device. In the case of a disaster affecting the server hard disk or other dump device media, ADSM is simply used to restore the dump devices so that a standard recovery can be made. If a point-in-time restore is required, the entire file set can be recovered from the appropriate-level ADSM backups and then a restore made from the recovered file sets through the standard SQL restore process.

If it is important to be able to go back several levels, consider the number of versions maintained by ADSM.

6.3.5 Dump Strategies

The internal dump feature of the SQL Server provides an online backup of all databases. But there are some disadvantages to this approach. If one database is created for use in single-user mode, and the system wants to perform the backup with the dump, it must be warranted that no user makes a transaction. It is easy to change the database usage mode to multiuser mode, but the SQL Server must be stopped and restarted.

The installation program creates some diskette dump devices, but, considering the huge amount of disk space of a normal SQL database, these are not useful devices.

6.3.5.1 Bulk Copy

To save disk space or tape space for the dump of the databases, the SQL Server provides another dump process: bulk copy. The bulk copy program (BCP) is a command line utility that copies SQL Server data to or from a file in a user-defined format.

Bulk copy uses the same API functions of the server as the basic dump, but it copies only the real data stored in the database device. In Figure 48 database device MyDevice1 includes data such as tables, views, and reports with a real disk space of 18MB. The size of the database device is predefined as 25MB, so the device includes free space of 7MB.

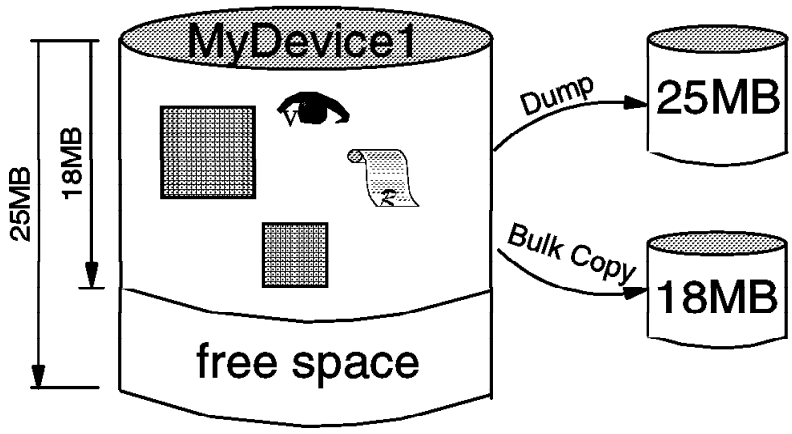

Figure 48. Bulk Copy vs. Dump

With a database dump, the space in the dump device will be 25MB size. With the bulk copy, the space in the dump device will be 18MB.

You can create a batch job for a bulk copy of all databases in the SQL Server as a backup. When you use the character mode format, you can import parts or all of the data into another SQL Server and/or onto another platform, such as OS/2.

Note: You can increase the size of a database, using the Enterprise Manager, but you cannot decrease it. Using bulk copy is the only way to decrease the size of a database.

6.4 Using the ADSMConnect Agent for SQL Server

The new ADSMConnect Agent for SQL Server is an optional product for online backups of the SQL Server. We use the short name, ′SQL Agent′, for this product.

The SQL Agent communicates with the SQL Server through the SQL Server DB-Library API (dblib) and communicates with the ADSM Server through the ADSM API, as shown in Figure 49 on page 114.

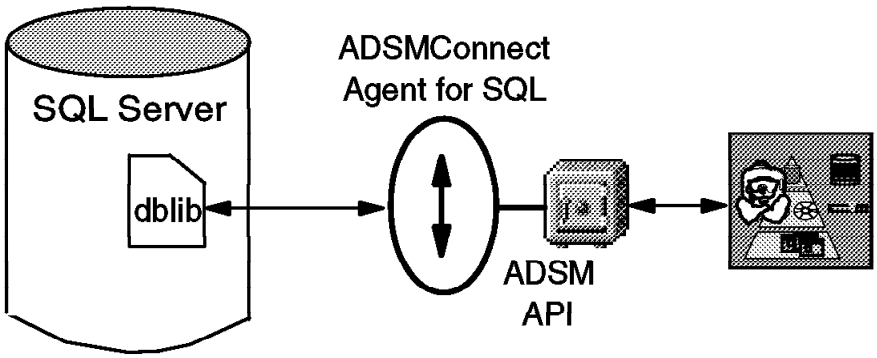

Figure 49. SQL Agent Communication

The SQL Agent (as a command line or GUI version) backs up and restores SQL databases running on the local machine. Multiple instances of the SQL Agent can be run at the same time. The command line program can be scheduled through the ADSM central scheduler. The GUI displays all of the SQL servers running on that machine as well as any number of user-defined ADSM servers it knows. The user must choose a database or transaction log and an ADSM Server to perform a backup or restore operation.

The GUI facilities direct manipulation of Agents, SQL Servers, databases, and ADSM Servers by using context-sensitive pop-up menus where the user selects an object and then clicks on the right mouse button to get a menu of operations that can be used with that object. Drag and drop is enabled between appropriate objects as well. For example, any database can be dropped on any ADSM Server to initiate a backup. Drag and drop will also allow copying of the options file representing an ADSM Server from one SQL Agent to another. The GUI is implemented as an Explorer-like single document interface with a tree view on the left side, a splitter bar, and a tabbed dialog with two tabs (backup and restore) on the right, (see Figure 50 on page 115.)

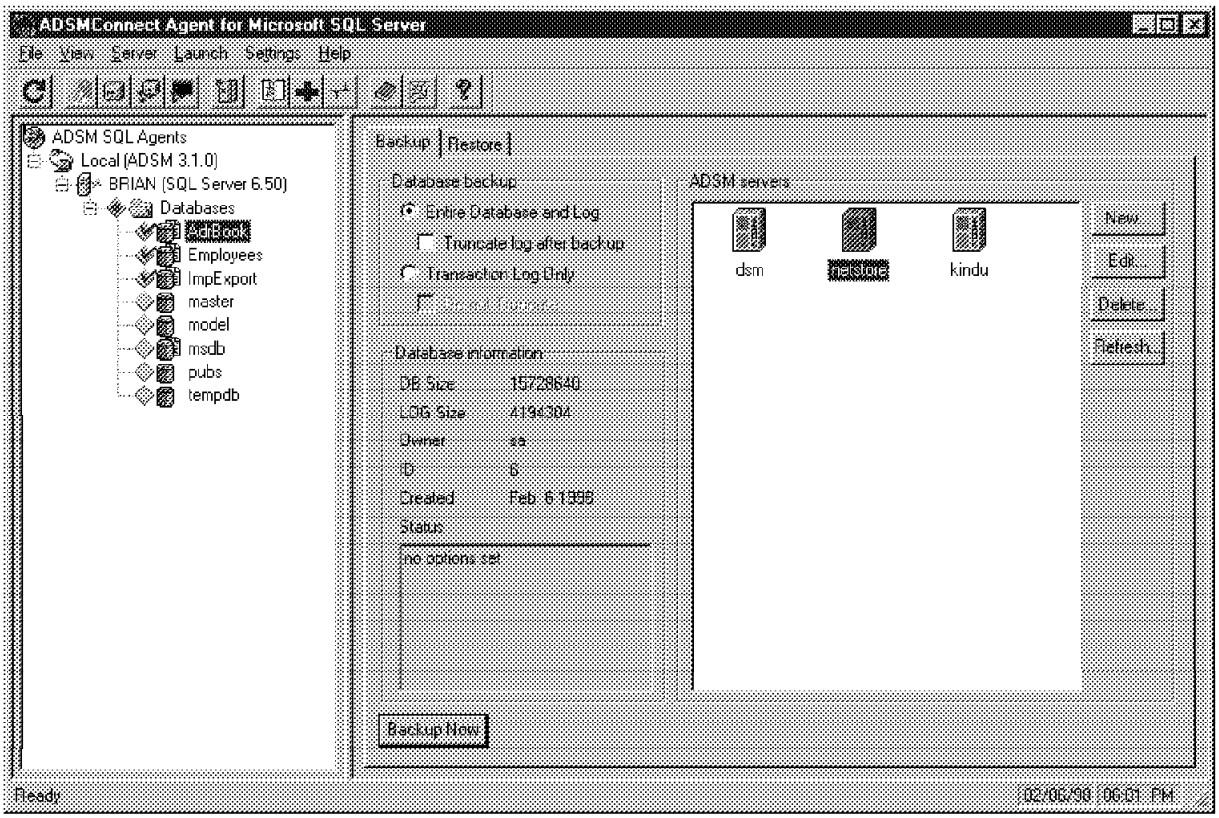

Figure 50. The SQL Agent GUI

All controls and fields resize with the parent window, and the splitter can be used to provide optimal viewing of the tree. Tree traversals drive the connection state of the SQL Agent. When the user clicks on the plus sign for the first time, either a connection with an SQL Agent or SQL Server is made or the objects in a folder are enumerated and displayed. After the initial expansion, the user can choose **refresh** to update the display.

6.4.1 Connection to the ADSM Server

In the *ADSMConnect Agent for Microsoft SQL Server Installation and User's Guide*, you can read everything about the installation and the command syntax.

After a successful installation, you must define an ADSM Server to which you want to connect. The Agent has one extra options file for every defined ADSM Server. The name of the options file is the same as the name of the ADSM Server. In our example the ADSM Server has the name Kindu.

Figure 51 on page 116 shows the GUI version of the options file. You can change all values in the GUI or directly in the options file.

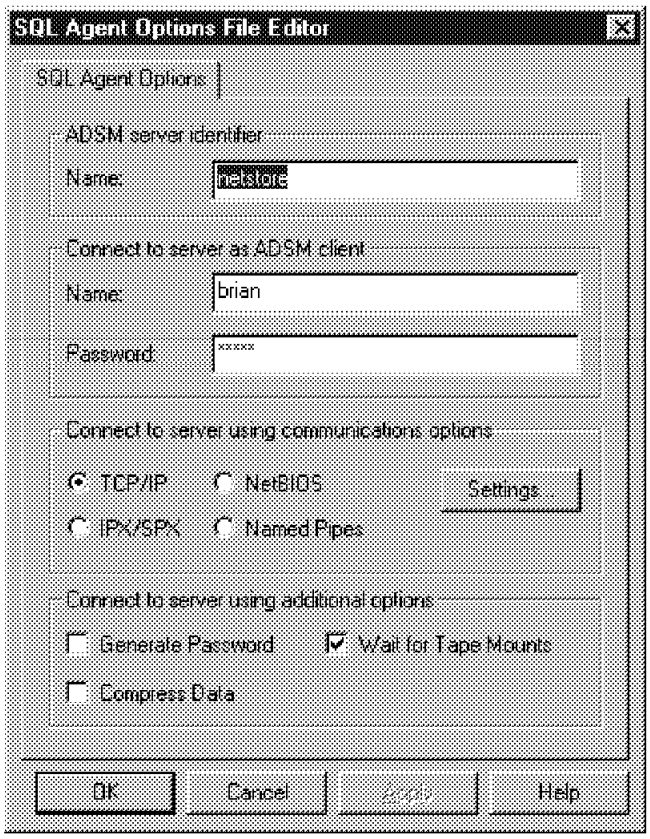

Figure 51. Settings for an ADSM Server

In the TCP/IP connection dialog you must enter the IP address of the ADSM server. (see Figure 52.)

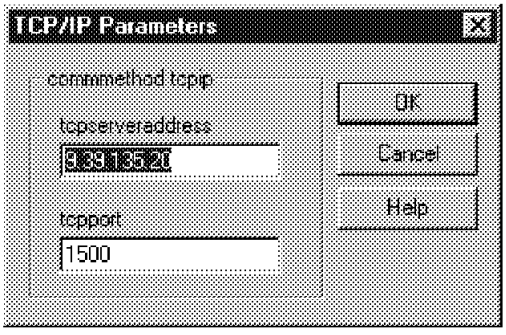

Figure 52. TCP/IP Configuration

6.4.2 Online Backup of the SQL Server

When you open the GUI of the Agent, you see all current databases of all SQL Servers. The agent is designed for the backup of a single database, multiple databases, or transaction logs. You can make full backups or incremental backups. Select the databases in the tree on the left side, the ADSM Server on the right side, and click on **Backup Now** to start the backup. The progress and the result of the operation appears in the Backup Progress window (Figure 53 on page 117).

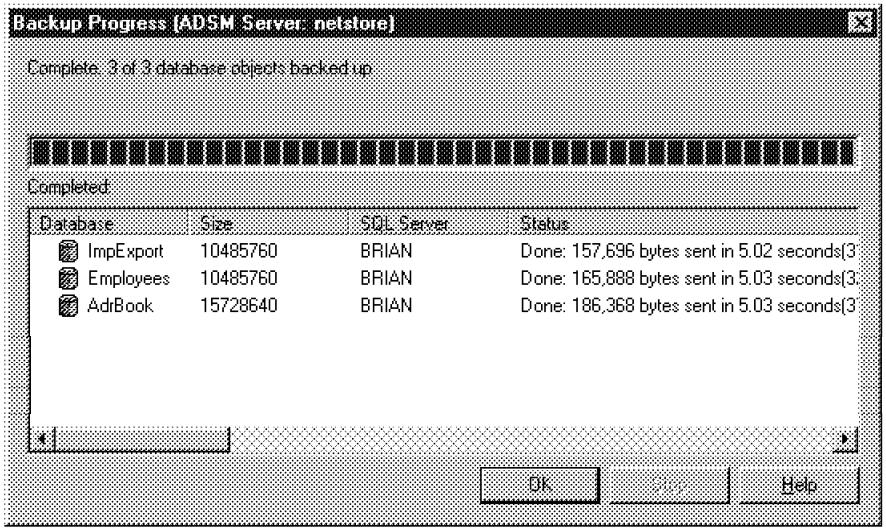

Figure 53. Backup of Databases with the SQL Agent

You can use the command line interface of the Agent as well to perform the backup. The command to back up database 'company' of SQL Server BRIAN in our example is:

```
sqldsmc /backupincr:company /adsmnode:sqlclient
/adsmpwd:mypasswd
/adsmoptfile:kindu.opt /mountwait
```

Additional information about the node name and the password can be supplied as additional parameters. Or, all needed information can be stored in the options file. In our example, this file is KINDU.OPT and you can start the SQL Agent with the `/adsmoptfile` parameter.

For a detailed description of the command line interface and all possible options, refer to the online help or the *ADSMConnect Agent for Microsoft SQL Server Installation and User's Guide*.

The result of the backup appears in the file space overview in the ADSM Server with the platform WinNTDB and the file space type of API:NTSQL.

6.4.3 Restore Process

Each time you start the SQL Agent, the program looks for the SQL Server and checks whether it is up and running. Otherwise the Agent asks whether it should start the SQL Server service. This behavior is important because the SQL Agent needs the current database information to show the current SQL Server contents in the database tree. Figure 54 on page 118 shows the contents of the SQL Server BRIAN.

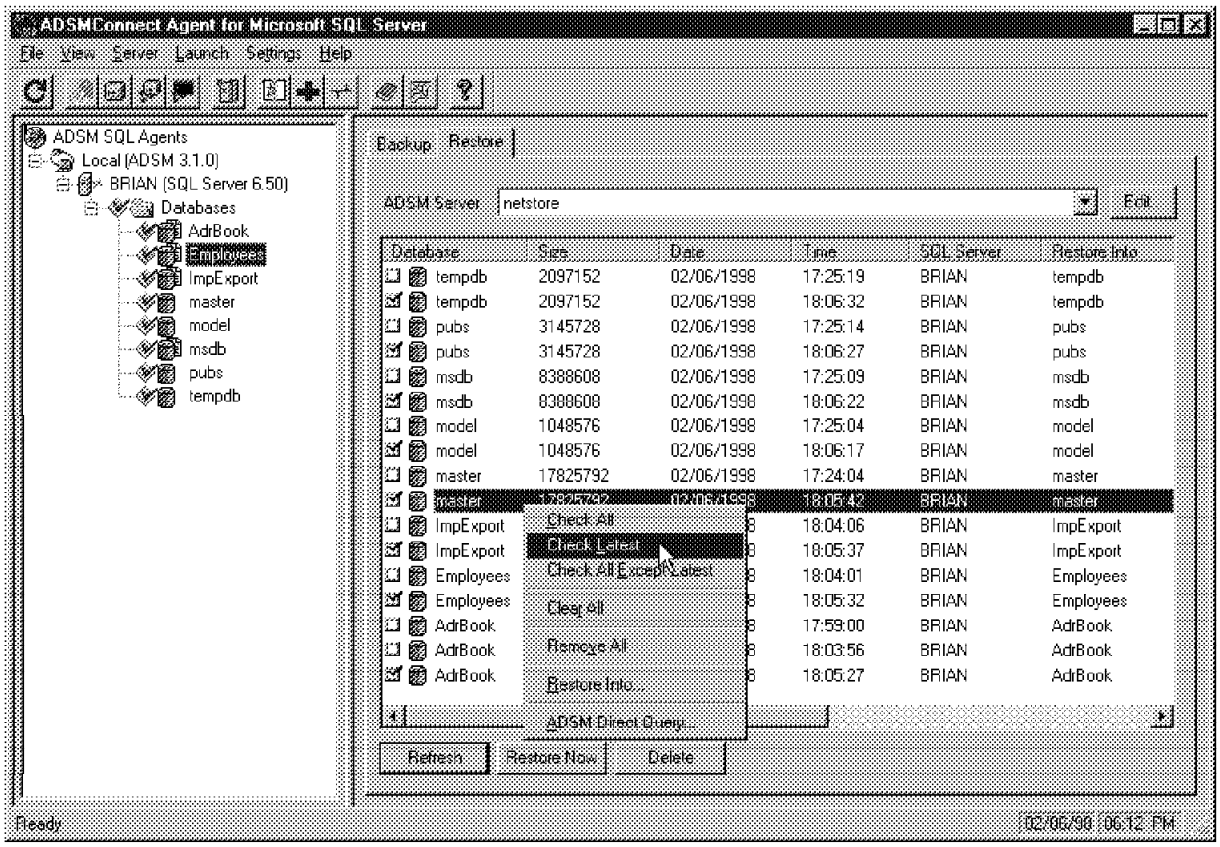

Figure 54. The SQL Agent Restore Window

When you select one or more databases on the left and the correct ADSM server and click on **Refresh**, the Agent shows all existing backups of the selected databases on the ADSM Server. The **Refresh** button is available only when you select at least one database.

The right side of the SQL Agent screen (Figure 54) shows the backups of database company. Different icons indicate different types of backup. A database has a square checkbox and a symbol for a database. A transaction log has a diamond checkbox and a symbol for a page.

To start a restore of one or more database objects, which could be full or incremental backups, click the appropriate checkbox and click on **Restore Now** or use the helpful right mouse button menu by choosing for example, **Check Latest**.

If you want to restore a database into another database or into another database device, click on the database with the right mouse button (see Figure 54). If you click on the right mouse button, the Restore Database Options dialog window appears (Figure 55 on page 119) showing the current databases of the SQL Server.

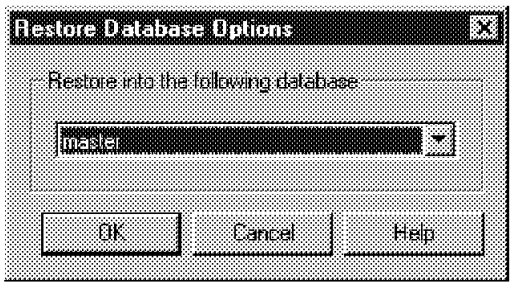

Figure 55. Restore into Another Database

Figure 54 on page 118 shows the progress and the success of the restore. If any problems occur during the restore process, the Agent displays an error number and provides a short description of the problem in the Restore Progress window.

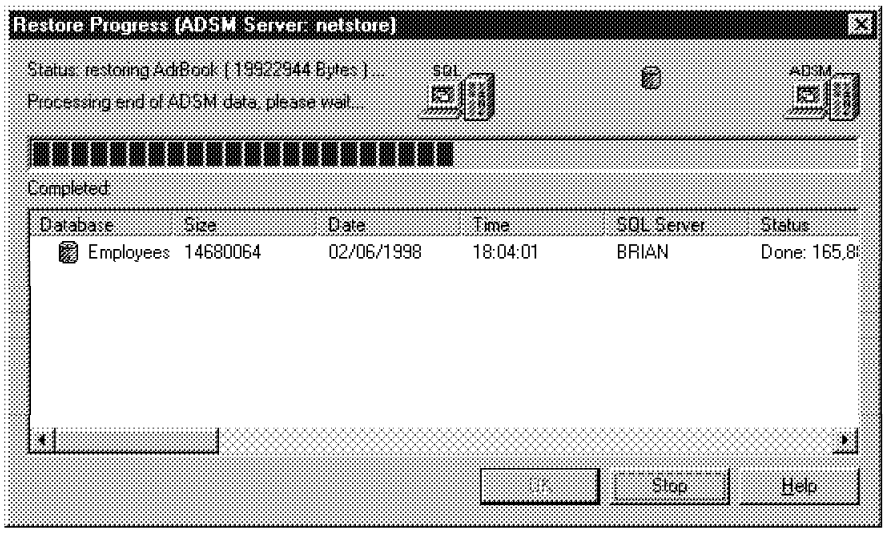

Figure 56. Restore Successfully Finished

6.4.4 Accidental Data Loss

In this section we show how the Agent restores damaged database objects, an entire database, and the entire SQL Server.

Unfortunately there is no way of backing up only parts of an SQL Server database, such as a table, view, or report. You must always back up at least one database. If an error occurs in one table, for example, a user deletes an important entry, it is not possible to select only that table for a restore. Such a corrupted database is called a suspect database in the SQL community. However, the SQL Server offers help for such a problem through the transaction log, because you can start a rollback for the transaction.

As we discussed at the beginning of this chapter, the deletion of an entire database causes more work for the restore. The deleted database does not appear in the current database information of the Enterprise Manager. As result, the deleted database does not appear in the dump device restore or in the SQL Agent.

For the restore of a deleted database you must perform one more step. First you must create a new database with the same name and the same size. If you do not know the name or the size of the database, you must create a new database that is bigger than the old one and rename the new one after the successful restore.

After the creation of a new database, you can use the dump device restore to restore directly from the dump device.

When you use the SQL Agent you must create a new database manually in the Enterprise Manager as well, before you can start the restore of the database. With the Agent you have more information about the deleted database. When you are in the restore window of the SQL Agent (see Figure 54 on page 118), you can click on the right part with the right mouse button. In the context menu that pops up, there is an entry to start an ADSM query on the ADSM Server. This query provides details of the existing NTSQL filespaces.

In the ADSM Direct Query window you can enter a pattern including * to view the databases stored on the ADSM Server (see Figure 57). The Agent lets you view or delete any database or transaction log objects that are stored on the selected ADSM Server, but it will not let you restore it until you create it on the SQL Server side. In the Agent, you can see the size and the name of the deleted database. Thus it is easier to create a new database in SQL Server with the same information.

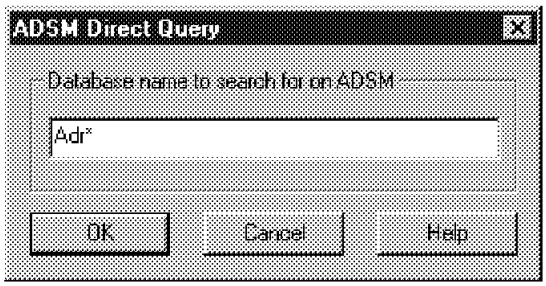

Figure 57. ADSM Direct Query

The last possible accidental data loss is the loss of the whole SQL Server system. In preparation for such a disaster, we recommend that you make a full backup of the %SqlRoot% directory and then run regular incremental backups with the ADSM backup/archive client. If you want or must do online backups, we recommend that you make a full backup of all databases with the SQL Agent. A daily incremental backup, perhaps started with the ADSM scheduler, will finish off your good preparation for disaster recovery.

6.5 Recovering the SQL Server

In this chapter so far we describe how you can backup and restore databases and database devices. Because all settings and system values are stored in the Master database, a loss of this important database will render the SQL Server inoperable. It is fundamental to protect this database with data redundancy techniques and consecutive backups. The Master database does not change every day, because only events such as Create Database, Alter Database, and Disk Mirror will change it. We recommend prohibiting user-defined objects in this database and being aware of the system procedures that modify it.

In this section we explain how you can recover the SQL Server system if the Master database is damaged or the whole system is lost. The procedure consists of:

1. Recovering the Master database from the SQL Server setup
2. Starting the SQL Server in single-user mode
3. Restoring the original Master database from an existing backup

First you must create a new Master database by using the SQL Server setup program to make the Server restartable. You can restore the Master database only if you start the SQL Server in single-user mode. For the restore process you can use three different approaches: SQL dump (see 6.3, “Using the SQL Dump” on page 110), bulk copy (see 6.3.5.1, “Bulk Copy” on page 113), or the ADSMConnect Agent for SQL Server. (see 6.4, “Using the ADSMConnect Agent for SQL Server” on page 113).

6.5.1 Recovering the Master Database

Start the SQL Server setup program from the install media and click on **Continue**, to get to the Options dialog box (Figure 58). In this box, select **Rebuild Master Database**.

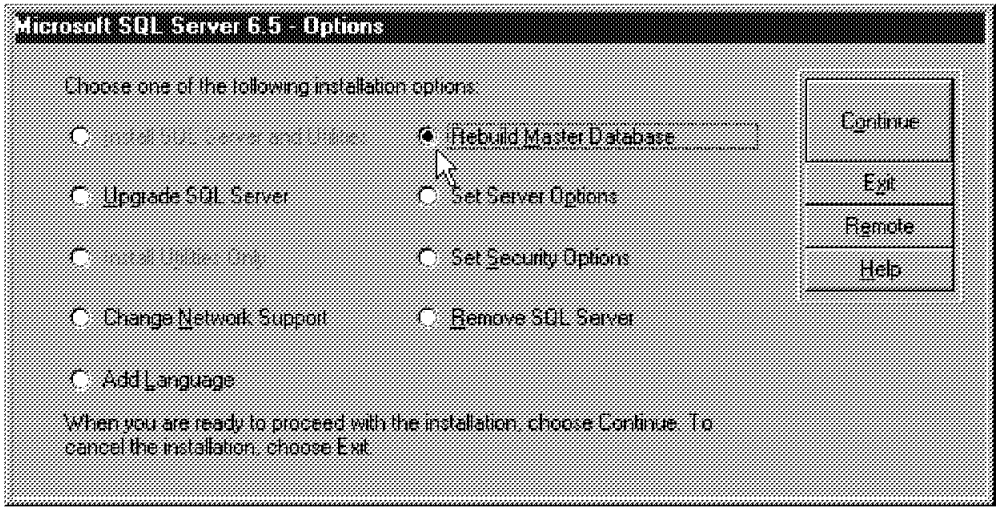

Figure 58. Options Dialog Box

The next window of the setup program warns you that a rebuild will replace the previous version of the Master database but that is exactly what you must do to recover the system. The Rebuild MASTER Database window (Figure 59 on page 122) will appear, regardless of whether there is a damaged Master database file or no file at all.

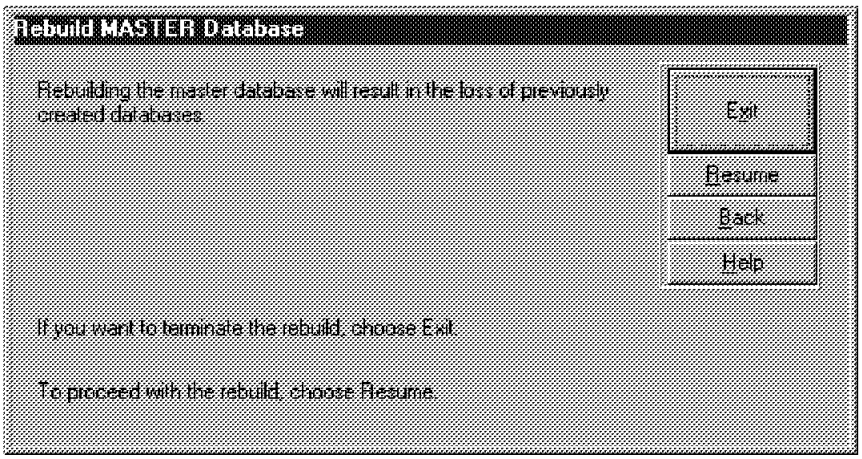

Figure 59. Rebuild MASTER Database Window

Enter the size of the Master database for the rebuild and choose a value that is equal to or greater than the previous value. After some minutes you will have a new Master database file on the disk, and you can leave the setup program.

Now you must start the SQL Server in single-user mode for the restore.

6.5.2 Starting the SQL Server in Single-User Mode

Single-user mode is a startup mode of the SQL Server that restricts connections. Only one user can connect to the SQL Server. The Checkpoint mechanism, which guarantees that completed transactions are regularly written from the cache to the database device, will not start.

The MSSQLServer service in NT is preconfigured to start in multi-user mode, and you must change this in the startup parameters. Open the Control Panel window, and click on the services icon. Insert the ″-m″ parameter in the startup parameters box to specify single-user mode.

Start the service by clicking on the **Start** button or through the SQL Service Manager. Alternatively, you can start the command line program:

```
C:\> sqlservr /dc:\sql65\data\master.dat /m
```

The /d parameter specifies the path to the Master database file and the /m selects the single-user mode startup.

6.5.3 Restoring the Original Master Database

If you are using SQL dump or bulk copy, you must reconfigure the dump device, because this information was included in the Master database. Once more you see how important it is to keep track of all system values and settings regarding the size and names of all devices. As soon as the dump device is available, you can start the restore of the last dump you have made.

After replacing the Master database you must stop the SQL Server, change the startup mode to multiuser, and restart the original system.

If you are using the SQL Agent, it is not necessary to reconfigure a dump device. Just start the SQL Agent as usual and restore the recent Master database file.

After the load and restore of the Master database file, you get the following error message in the SQL Agent:

```
Error: SQL Server Errors encountered DB-LIB message: Possible
network error: Write to SQL Server Failed. Connection broken.
```

You can ignore this message, because the SQL Server tries to prevent an inconsistent state and therefore disconnects all current users, including the API. Do not forget to stop the SQL Server after the successful restore process and restart the SQL Server in multiuser mode.

Note: If you are starting the SQL Agent and try to restore the Master database while the SQL Server is running in multiuser mode, you will get the following error message:

```
Error: SQL Server Errors encountered DB-LIB message: LOAD DATABASE must
be used in single user mode if trying to restore the Master database.
```

During the rebuild of the Master database, the setup program drops and rebuilds the MSDB database without any notification as well. This database includes all the scheduling information, which will be lost. You must restore this database directly after the restore of the Master database. This MSDB database restore can be done in single-user or multiuser mode.

For the MSDB database there could be also backups of transaction logs. Restore and apply every transaction log to the original MSDB database in the SQL Enterprise Manager.

Chapter 7. Recovering Microsoft Exchange Server

The Microsoft Exchange Server is software for handling mail, organizing meetings, and managing the enterprise directory. In this chapter we explain how to configure ADSM to back up the Microsoft Exchange Server and prepare for recovery.

Note:

At the time of writing, the ADSMConnect Agent for Microsoft Exchange Server is about to be released. You should refer to the documentation that accompanies this agent for additional methods of backing up Microsoft Exchange Server. The ITSO will shortly be running a residency to explore and document its usage. This will appear first on the ITSO website as HTML. It is anticipated to be released in hardcopy in September 1998.

7.1 Data Characteristics

In this section we show you how to determine the data that must be backed up with ADSM in order to recover it.

7.1.1 General View

The Microsoft Exchange Server manages private and public folders. Users can copy the private and public files and store the copies on their personal computers or in the Microsoft Exchange Server.

The Microsoft Exchange Server manages these files:

Private file The PRIV.EDB file, on the Exchange Server, contains the mail of each user of the Microsoft Exchange Server application.

Public file The PUB.EDB file, on the Exchange server, contains all of the public files. The public files enable you to manage information and update it quickly and easily.

Personal files The *.PST files are copies of mail or data in the Microsoft Exchange Server database. So, personal files are replications of information that is also in the Exchange server. The *.PST files can be on the personal computer disk or the server disk.

Offline files The *.OST files are personal copies of public files from the Exchange server. They are used to copy mailbox folders or ″favorite″ public folders to a laptop computer. The user is connected to the server, and Microsoft Exchange synchronizes the data between the server and the *.OST files on the laptop.

The following four components are installed during installation of the Microsoft Exchange Server and must run all the time:

- **Directory**

  The directory contains information about the organization of the Microsoft Exchange Server. The information in the directory includes addresses,

mailboxes, distribution lists, public folders, and other servers. Microsoft Exchange Server administrators manage the directory.

- **Information store**

  The information store consists of *private information* for users' mailboxes and *public information* for public folders. The private and public information stores are two databases. When a message or a folder is created, the Microsoft Exchange Server writes information in log files and works in the server memory to increase performance. This explains why you have to stop the Microsoft Exchange Server before you back it up with ADSM.

- **Message transfer agent (MTA)**

  The MTA is used for routing messages between Exchange servers or between an Exchange server and the Internet.

- **System attendant**

  The system attendant performs all administrative functions, such as creating users and setting up the distribution lists.

The Internet Mail Connector, Microsoft Mail Connector, and X400 Connector are optional components of the Microsoft Exchange Server.

The Microsoft Exchange Server provides mail distribution.

Microsoft Exchange Server data can be organized in several ways. Figure 60 on page 127 shows the Microsoft Exchange Server organization with a single server.

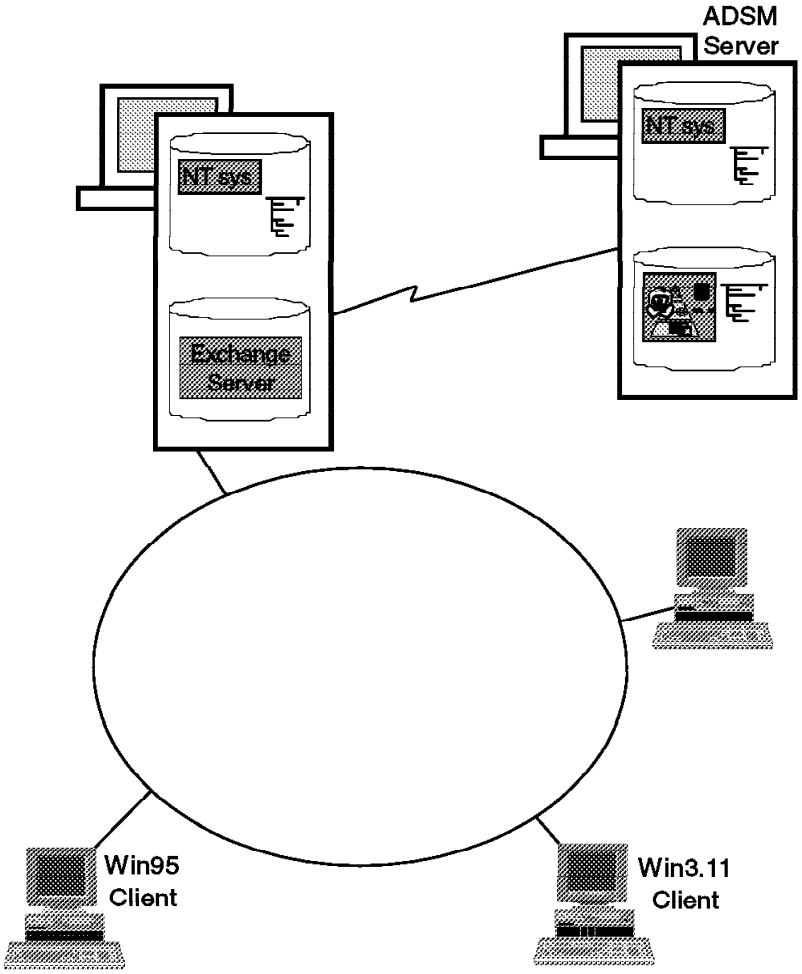

Figure 60. Microsoft Exchange Server Organization: One Server

In this organization, there is only one copy of the PUB.EDB and PRIV.EDB files. In case of damage, you must use the backup copy to restore the Microsoft Exchange Server data of all users in the enterprise.

Figure 61 on page 128 shows an organization with two Microsoft Exchange Servers. You can use more than two servers, and they can be located in the same or in different places.

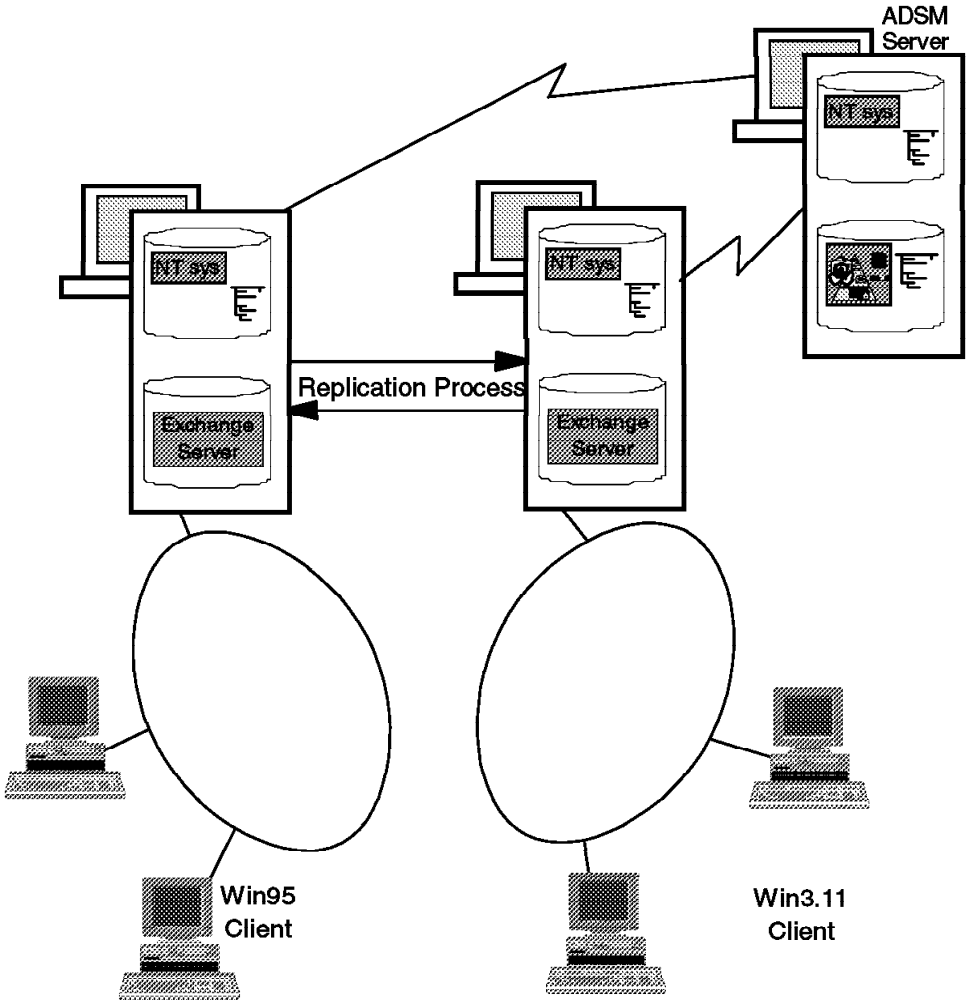

Figure 61. Microsoft Exchange Server Organization: Two Servers

In this case, the Microsoft Exchange Server replicates the data between the directories in the two servers and sends messages to the Exchange server to which the mail is addressed.

When a company uses many servers, the administrator can configure the Microsoft Exchange Server differently: All servers can be either public or private servers with replication, or the private information store can be separated from the public information store. The organization you choose depends on the amount of the traffic.

The Microsoft Exchange Server documentation recommends using NT Backup to back up the Microsoft Exchange Server. NT Backup runs well, but it does not allow you to develop a backup strategy as ADSM does. NT Backup performs the backup without stopping Microsoft Exchange Server services.

With ADSM, you need an API to back up the Microsoft Exchange Server without stopping its services. Until the API becomes available, you have to stop the Microsoft Exchange Server services before you back up the data.

The paths that you have to back up are listed under the Database Paths on the Properties window of the Microsoft Exchange Server administration application (see Figure 62 on page 129).

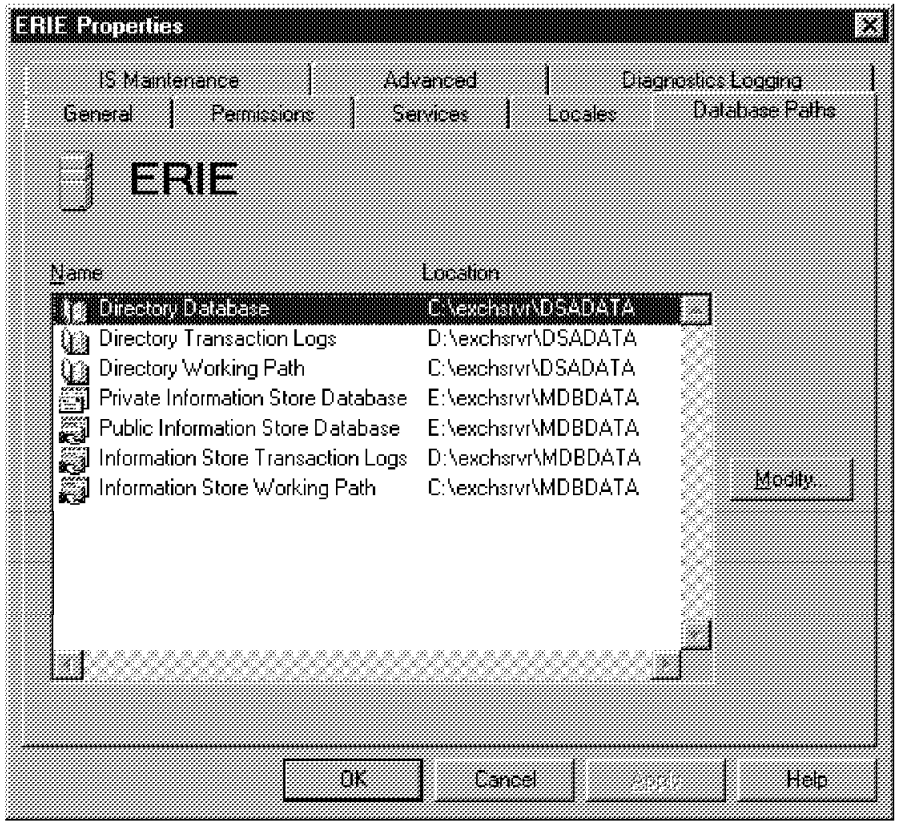

Figure 62. Database Paths for Microsoft Exchange Server Data

Backing up Microsoft Exchange Server data is critical, because the Microsoft Exchange Server manages all of the information about important events in your enterprise such as appointments, meetings, and strategic documents.

7.1.2 Backup and Restore Strategies

Your ADSM backup strategy must include these files:

- **Directory**

  It is essential to back up the directory of each Exchange server you use. The directory is changed each time you change your Exchange organization (for example, when you add a new server, create a new mailbox, or set up a distribution list). So, you have to back up the Exchange directory whenever you change it.

  The directory contains some server-specific information, so you have to maintain a backup of each server directory.

- **Information store**

  Each information store is backed up as one object.

You only have to back up the directory and the information store. The other components contain temporary data. If you back up and restore the MTA, you will create duplicate messages!

7.1.3 Backing Up an .INI File

After loss of a server disk, you must reinstall the Microsoft Exchange Server in the same configuration, so you should create an EXCHANGE.INI file, which contains the installation parameters. You can back up and restore this file with ADSM.

You can create the .INI file during the Microsoft Exchange Server installation. Use the setup program with the -r parameter.

7.1.4 Attention Please...

If you are testing the Microsoft Exchange Server application and the ADSM backup software, be careful with the server Options Permissions. As shown in Figure 63, do not check the Delete primary Windows NT account when deleting mailbox box.

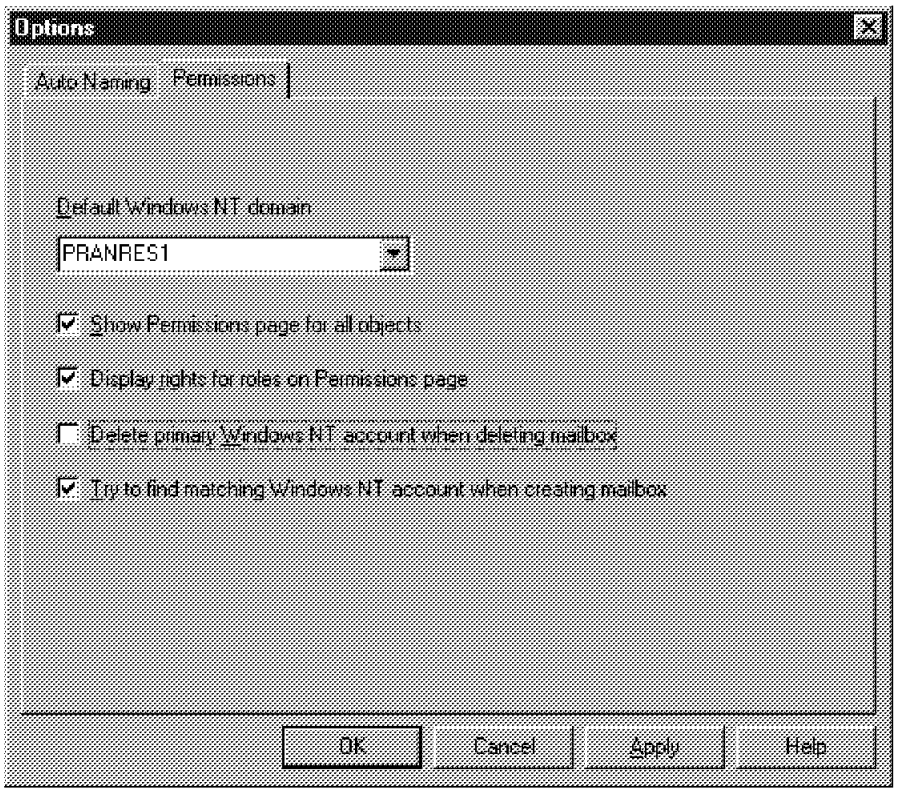

Figure 63. Microsoft Exchange Server: Server Options Permissions

7.2 ADSM Considerations

If you try to back up the Microsoft Exchange Server with the services open, you get these messages in the DSMSCHED.LOG:

```
09:48:38 Normal File--> 1,581,056 E:\exchsrvr\MDBDATA\PRIV.EDB  ** Unsuccessful **
09:48:38 ANS4228E Sending of object 'E:\exchsrvr\MDBDATA\PRIV.EDB' failed
09:48:38 ANS4090E Access to the specified file or directory is denied
09:48:38 Normal File--> 1,581,056 E:\exchsrvr\MDBDATA\PUB.EDB  ** Unsuccessful **
09:48:38 ANS4228E Sending of object 'E:\exchsrvr\MDBDATA\PUB.EDB' failed
09:48:38 ANS4090E Access to the specified file or directory is denied
```

Below we discuss two methods of backing up the Microsoft Exchange Server:

- Using **ADSM** functions only
- Using both ADSM and the NT at command.

Table 1 lists the disavantages and avantages of each method. Your backup plan will determine which method you choose.

Table 1. Using ADSM or ADSM and the NT at Command

	ADSM Functions	**ADSM and NT at Command**
	Use preschedulecmd and postschedulecmd in the DSM.OPT file.	Use at command to plan the backup.
Disadvantages	All ADSM processes for this NT Client stop the Exchange services. The Exchange services are stopped while ADSM takes the backup. You have to write an NT script to stop all Microsoft Exchange Server and other application processes that should be stopped.	Uses a process other than ADSM. An NT administrator could stop the process without telling the ADSM administrator. If you change the Microsoft Exchange Server database paths, you have to change the NT commands file used for backing up the Microsoft Exchange Server data.
Advantages	Uses only one process. All information about the ADSM backup is in the DSMSCHED.LOG, so it is very simple for the ADSM administrator to control the backup of the Exchange server.	Stops the Exchange services for the duration of the Exchange data backup. Can have different strategies for backing up Microsoft Exchange Server data and the data of other applications.

Before we discuss the backup methods, let us consider when to back up Microsoft Exchange Server data. The Microsoft Exchange Server administrator determines the number of days that the Microsoft Exchange Server data should be kept and the IS maintenance process properties (see Figure 64 on page 132).

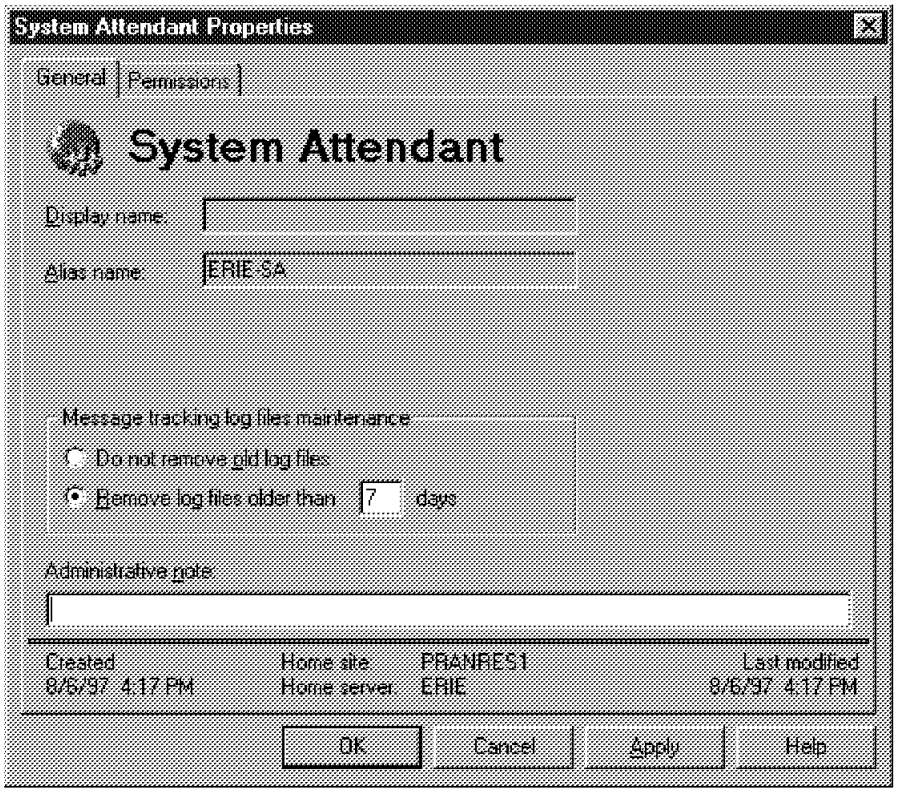

Figure 64. Microsoft Exchange Server Remove Log Files Parameter

You should backup the Microsoft Exchange Server when the IS maintenance process is not running and with a frequency lower than indicated on the System Attendant Properties window in the Remove log files older than ... days radio button.

7.2.1 Using ADSM Functions Only

In the ADSM DSM.OPT file for the NT Client, you can use two parameters:

- PREschedulecmd

  Specifies a command to process before running a schedule.

- POSTschedulecmd

  Specifies a command to process after running a schedule.

For more explanation, see *ADSTAR Distributed Storage Manager: Using the Microsoft Windows Backup-Archive Clients Version 2* (SH26-4056).

Put these two parameters in the DSM.OPT file:

```
************************************************************************

PRE  C:\STOPEXCH.CMD
POST C:\STAREXCH.CMD

************************************************************************
```

With the PREschedulecmd and POSTschedulecmd parameters, you can only use one command, so you have to write a preprocessing program that shuts down all of your NT applications and a postprocessing program that opens them.

ADSM runs the preprocessing job and waits for it to end before it makes the backup, so you have to test your preprocessing job because a failure in this job makes the backup process fail.

This process closes all Microsoft Exchange Server services:

```
   ***   STOPEXCH.CMD   ***

NET STOP MSEXCHANGEMTMI /y
NET STOP MSEXCHANGEMTA /y
NET STOP MSEXCHANGEIS /y
NET STOP MSEXCHANGESA /y
NET STOP MSEXCHANGEDS /y
```

Next, ADSM backs up the files (as indicated in DSMSCHED.LOG):

```
09:48:38 Normal File--> 1,581,056 E:\ exchsrvr\MDBDATA \PRIV.EDB  Sent
09:48:41 Normal File--> 1,581,056 E:\ exchsrvr\MDBDATA \PUB.EDB  Sent
```

On completion, the next process starts the Microsoft Exchange Server services:

```
***   STAREXCH.CMD   ***

NET START MSEXCHANGEIS /y
NET START MSEXCHANGEMTA /y
NET START MSEXCHANGEDS /y
```

7.2.2 Using ADSM and the NT ″at″ Command

If you use the ADSM scheduler to plan your backup but do not use the PREschedulecmd and POSTschedulecmd parameters, which stop the NT Exchange services, you will find the following errors in the DSMSCHED.LOG:

```
09:48:38 Normal File--> 1,581,056 E:\exchsrvr\MDBDATA\PRIV.EDB  ** Unsuccessful **
09:48:38 ANS4228E Sending of object ′E:\exchsrvr\MDBDATA\PRIV.EDB′ failed
09:48:38 ANS4090E Access to the specified file or directory is denied
09:48:38 Normal File--> 1,581,056 E:\exchsrvr\MDBDATA\PUB.EDB  ** Unsuccessful **
09:48:38 ANS4228E Sending of object ′E:\exchsrvr\MDBDATA\PUB.EDB′ failed
09:48:38 ANS4090E Access to the specified file or directory is denied
```

The process does not back up the Microsoft Exchange Server. So you have to make a CMD file, as shown below. The dsmc command is used to back up all database paths. Check under Database Paths in the Microsoft Exchange Server Properties window (see Figure 62 on page 129).

```
   *** ADSMEXCH.CMD ***

NET STOP MSEXCHANGEMTMI /y
NET STOP MSEXCHANGEMTA /y
NET STOP MSEXCHANGEIS /y
NET STOP MSEXCHANGESA /y
NET STOP MSEXCHANGEDS /y
c:
cd\win32app\ibm\adsm\baclient
dsmc s -subdir=yes c:\exchsrvr\dsadata\*.*   | more >  adsmexch.log
dsmc s -subdir=yes d:\exchsrvr\dsadata\*.*   | more >> adsmexch.log
dsmc s -subdir=yes e:\exchsrvr\mdbdata\*.*   | more >> adsmexch.log
```

```
dsmc s -subdir=yes d:\exchsrvr\mdbdata\*.*   | more >> adsmexch.log
dsmc s -subdir=yes c:\exchsrvr\mdbdata\*.*   | more >> adsmexch.log
cd\
NET START MSEXCHANGEIS /y
NET START MSEXCHANGEMTA /y
NET START MSEXCHANGEDS /y
```

Schedule this job with the at command or the WINAT tool.

The at command schedules commands and programs to run on a computer at a specified time and date. The Schedule service must be running to use the **at** command:

```
at :computername time /interactive
    /every:date,... ] /next:date,... "command"
```

For **date** specify one or more days of the week or one or more days of the month (using numbers 1 through 31). Separate multiple date entries with commas. If the date is omitted, the current day of the month is assumed.

The **command** is the Windows NT command, program (.EXE or .COM file), or batch program (.BAT or .CMD file) to be run. When the command requires a path as an argument, use the absolute path, that is, the entire pathname beginning with the drive letter. If the command is on a remote computer, specify the server and sharename, rather than a remote drive letter. You can use quotation marks around the command, whether you are using the at command at the command line or in a batch file.

If the at command does not work, either the Schedule service is not running or it does not have all of the required rights. In this case, go to Configuration/Service window, select **Schedule** and the **Settings Control panel**. For Startup Type select **Automatic** and for LogOnAs select **System Account** (see Figure 65).

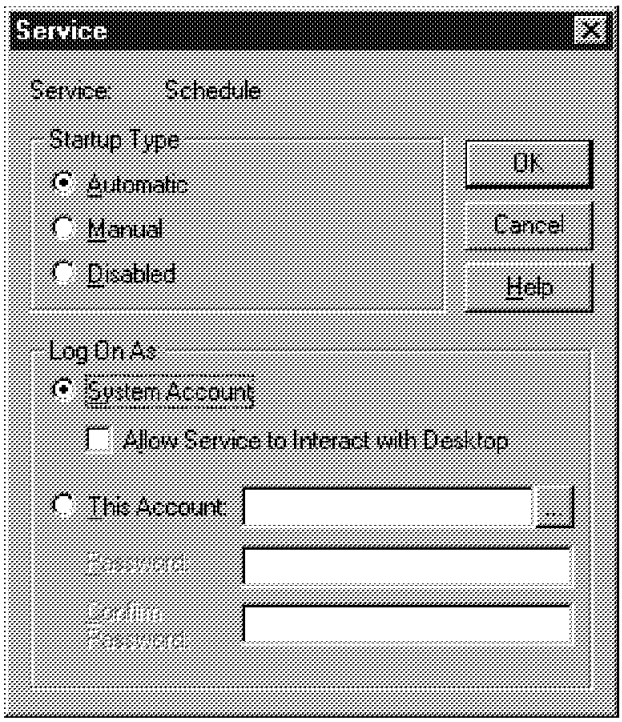

Figure 65. Windows NT Service Schedule Startup Option

To schedule the ADSMEXCH.CMD job, issue this command:

```
AT 10:00PM /interactive /every :M,T,W,Th,F,S,Su C :\ADSMEXCH.CMD
```

The ADSMEXCH.CMD command creates the ADSMEXCH.LOG in the DSMC.EXE directory:

```
ADSTAR Distributed Storage Manager
Command Line Backup Client Interface - Version 2, Release 1, Level 0.6
(C) Copyright IBM Corporation, 1990, 1996, All Rights Reserved.

Selective Backup function invoked.

Session established with server KINDU: AIX-RS/6000
  Server Version 2, Release 1, Level 0.12
  Data compression forced on by the server
  Server date/time: 08/01/1997 10:17:00   Last access: 08/01/1997 10:13:42

Normal File-->       1,581,056 E:\exchsrvr\MDBDATA\PRIV.EDB ..... Sent
Normal File-->       1,843,200 E:\exchsrvr\MDBDATA\PUB.EDB ..... Sent
Selective Backup processing of 'E:\EXCHSRVR\MDBDATA\*.EDB' finished with 0 failures.


Total number of objects inspected:        2
Total number of objects backed up:        2
Total number of objects updated:          0
Total number of objects rebound:          0
Total number of objects deleted:          0
Total number of objects failed:           0
Total number of bytes transferred:    751.4 KB
Data transfer time:                    0.07 sec
Data transfer rate:                10,734.44 KB/sec
Average file size:                  1,672.1 KB
Compression percent reduction:           78%
Elapsed processing time:            0:00:04
```

7.3 Accidental Loss of Data Integrity

In this section, we show you how to restore the Microsoft Exchange Server data after a loss of integrity. Table 2 on page 136 explains how to recover according to your Microsoft Exchange Server organization.

As mentioned previously, a loss of Microsoft Exchange Server data is a loss of the dir.edb for the directory or the PUB.EDB and the PRIV.EDB files for the information stores.

Before you can recover the data, you must select the best approach that will restore the latest version of these files.

Table 2. Restoring the Data after Loss of Integrity			
	One Server	**Two or More Servers with Replication of All Public Files**	
Directory restoration	Use ADSM to restore the dir.edb file.	Use ADSM to restore the dir.edb file and NT for the duplication process.	
Data restoration	Use ADSM to restore directory and *.EDB files.	**If all folders are duplicate**: Use ADSM to restore the Microsoft Exchange Server information store. Use the Microsoft Exchange Server for the duplication process.	**If only certain information is duplicate**: Use ADSM to restore programs and *.EDB files. Let Exchange do the recovery... and evaluate the data lost.

Before restoring The Microsoft Exchange Server, take a few minutes to write these two small programs, which will help you:

- **STOPEXCH.BAT** to stop all Microsoft Exchange Server services:

```
     * * *  STOPEXCH   * * *
NET STOP MSEXCHANGEMTMI /y
NET STOP MSEXCHANGEMTA  /y
NET STOP MSEXCHANGEIS   /y
NET STOP MSEXCHANGESA   /y
NET STOP MSEXCHANGEDS   /y
```

- **STAREXCH.BAT** to start some Microsoft Exchange Server services:

```
     * * *  STOPEXCH   * * *
NET STOP MSEXCHANGEMTA  /y
NET STOP MSEXCHANGEDS   /y
```

7.3.1 Restoring the Exchange Server Directory

The Microsoft Exchange Server directory has three components (see Figure 66 on page 137): the dir.edb, the directory database, and the TEMP.EDB and EDB.CHK, the database change logs for changes made to the directory for the last seven days (if you specified 7 days in the System Attendant Properties window).

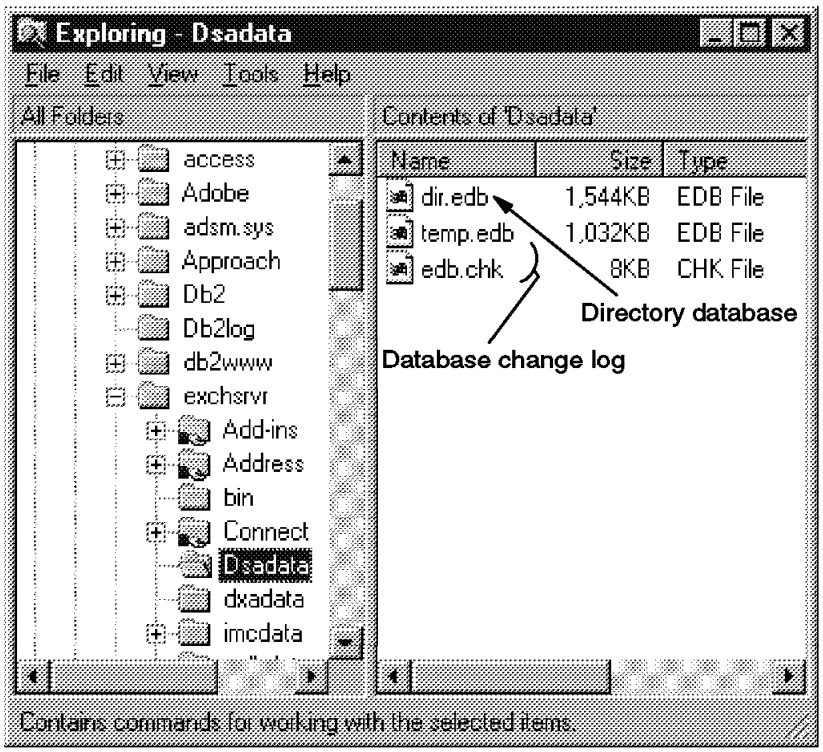

Figure 66. Exchange Server Directory

7.3.1.1 Loss of Microsoft Exchange Server Directory Integrity

When you start the Microsoft Exchange Server services, the following message is displayed on your screen:

```
System error 1068 occurred
```

If Microsoft Exchange Server directory integrity has been lost, you can restore a dir.edb file backed up with ADSM and let the Microsoft Exchange Server recover the changes made to the directory since the last backup. Therefore the backup must be less than seven days.

Here are the steps to recover the Microsoft Exchange Server directory:

1. Use the NT Control/Services window or the STOPEXCH.BAT program to stop all Microsoft Exchange Server services.
2. Go to the ADSM command line application.

   Use the ADSM restore function:

   ```
   DSMC Restore -subdir=yes -todate=080397
       C:\exchsrvr\dsadata\*.edb
   ```

   The -todate parameter enables you to select the date of the last correct database.

   You also can use the ADSM GUI.

   The -latest parameter enables you to restore the last backup of the Microsoft Exchange Services server directory.

   See A.1, “ADSM Restore Commands” on page 157 for more information about the ADSM DSMC restore function.

3. Use the STAREXCH.BAT command or the NT Control/Services window to restart Microsoft Exchange Server services. It takes a few minutes to rebuild the directory.

7.3.1.2 Loss of Exchange Directory Data

The Exchange application is installed on the Microsoft Exchange Server. Use the NT Control/Services window or the STOPEXCH.BAT program to stop all Exchange services.

To recover the directory follow these steps:

1. Go to the ADSM command line application or use the ADSM GUI.

   Use the ADSM restore function:

   ```
   DSMC Restore -subdir=yes -latest
       C:\exchsrvr\dsadata\*.edb
   DSMC Restore -subdir=yes -latest
       D:\exchsrvr\dsadata\*.edb
   ```

   The -latest parameter enables you to restore the last backup of the Microsoft Exchange Server directory.

   See A.1, “ADSM Restore Commands” on page 157 for more information about the ADSM DSMC Restore function.

2. Use the STAREXCH.CMD or the Control\Services window to restart the Microsoft Exchange Server services.

7.3.2 Restoring the Exchange Information Store

Because the Microsoft Exchange Server directory is now restored, the Microsoft Exchange Server clients can make connections, but the public folders (see Figure 67) contain only the Favorites and All Public Folders default folders.

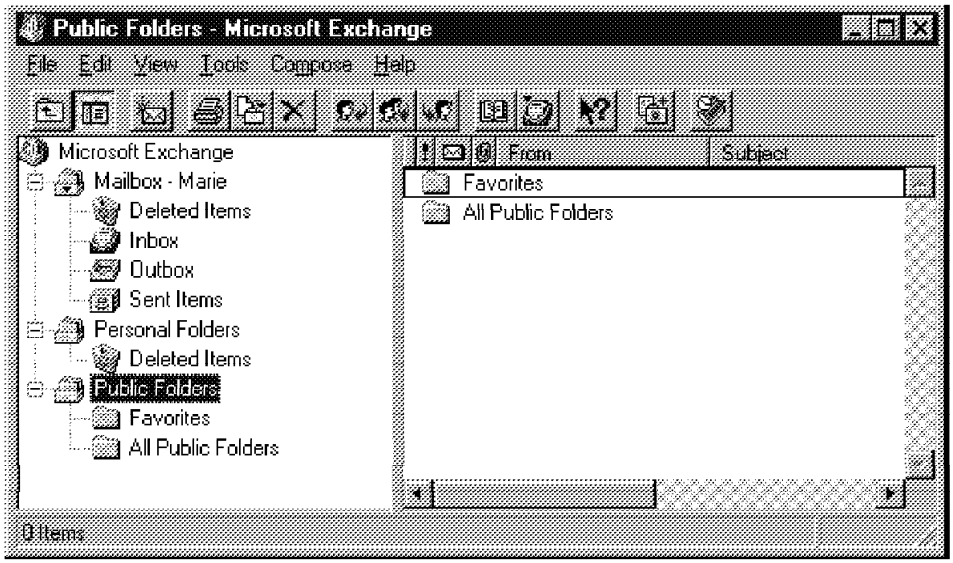

Figure 67. Microsoft Exchange Server Client Public Folders

To restore the Exchange information store, follow these steps:

1. Stop the Microsoft Exchange Server services, using the STOPEXCH.BAT command or the NT Control/Services window.

2. Use ADSM to restore the data:

```
DSMC Restore -subdir=yes -latest
    C:\exchsrvr\mdbdata\*.*
DSMC Restore -subdir=yes -latest
    E:\exchsrvr\mdbdata\*.*
```

Because the Microsoft Exchange Server had made a connection, ADSM found the default PRIV.EDB and PUB.EDB files. You have to replace these files:

```
ANS4098I Specified directory branch structure has been restored
file E:\EXCHSRVR\MDBDATA\PRIV.EDB exists
Do you want to replace it (Yes/No) ?
```

Reply **Y**. It is also possible to use the replace=yes parameter in the DSMC ADSM command (see A.1, “ADSM Restore Commands” on page 157 for more information about the ADSM DSMC restore function).

3. Restart the Microsoft Exchange Server services, using the STAREXCH.BAT command or the NT Control/Services window.

When you restart the Microsoft Exchange Server services, if Microsoft Exchange Server has some more recent information in its logs, it will take a few minutes to re-create the information store.

7.3.3 Exchange Duplication Process

If you delete a directory replication between two sites because one Microsoft Exchange Server is down, do not run the DS/IS consistency adjuster before the end of the ADSM restoration. This will result in loss of data.

7.4 Recovering a Lost Data Disk

If you lose a data disk, follow these steps to recover it:

1. Estimate the damage.
2. Install the NT system and ADSM client.
3. Install the Microsoft Exchange Server software:
   a. Use ADSM to restore the EXCHANGE.INI file and use the Setup /q C: exchange.ini command to install the Microsoft Exchange Server

      or

   b. Use the setup command and, during the installation, choose the same name for the organization and site.
4. Use ADSM to restore all of the Microsoft Exchange Server data. You can use either the command line or the GUI. Restore all database paths that contain Microsoft Exchange Server data.

The program restores the Microsoft Exchange Server data:

```
    * * *   RESTEXCH.CMD    * * *
NET STOP MSEXCHANGEMTMI /Y
NET STOP MSEXCHANGEMTA  /Y
NET STOP MSEXCHANGEIS   /Y
NET STOP MSEXCHANGESA   /Y
NET STOP MSEXCHANGEDS   /Y
C:
```

```
cd\win32app\ibm\adsm\baclient
dsmc res -subdir=y -latest -rep=all C:\exchsrvr\dsadata\*.* ] more>rest.log
dsmc res -subdir=y -latest -rep=all C:\exchsrvr\mdbdata\*.* ] more>>rest.log
dsmc res -subdir=y -latest -rep=all D:\exchsrvr\dsadata\*.* ] more>>rest.log
dsmc res -subdir=y -latest -rep=all D:\exchsrvr\mdbdata\*.* ] more>>rest.log
cd\
NET START MSEXCHANGEIS /Y
NET START MSEXCHANGEMTA /Y
NET START MSEXCHANGEDS /Y
```

5. Use the NT Control/Services windows to verify all Microsoft Exchange Server services.

7.5 Disaster Recovery

In this section, we explain how to restore the Microsoft Exchange Server after a disaster.

Before making the restore, notice that:

- You can only restore the directory to a computer with the same Windows NT domain. If you do not, you cannot access the information store (see 7.5.1, “Restoring the Application in the Same Server”).
- An information store can be restored to a different Microsoft Exchange Server (see 7.5.2, “Restoring the Data in Another Server”).

7.5.1 Restoring the Application in the Same Server

Follow these steps after severe damage such as a lost disk or machine:

1. Install NT on the computer that has the same name as the old server.
2. Install the NT ADSM client on this computer and make the connection with the ADSM server by setting the appropriate parameters in the DSM.OPT file.
3. Install the Microsoft Exchange Server programs on the new computer, as shown in 7.4, “Recovering a Lost Data Disk” on page 139.
4. Restore the private and public information stores from the last ADSM backup:

```
DSMC Restore -subdir=Y -lastest C:\exchsrvr\dsadata\*
DSMC Restore -subdir=Y -lastest C:\exchsrvr\ddbdata\*
DSMC Restore -subdir=Y -lastest D:\exchsrvr\dsadata\*
DSMC Restore -subdir=Y -lastest E:\exchsrvr\mdbdata\*
```

5. Run the DS/IS consistency adjustment.

7.5.2 Restoring the Data in Another Server

In an emergency, you might have to restore the information store backed up by ADSM on an alternative server. The process shown here creates a *.PST file that a user can use with the Microsoft Exchange Server client.

Attention: To restore files with ADSM from another node, you must be using the same file system (NTFS to NTFS, FAT to FAT).

1. You have to find another server that is not replicating the existing organization. The server must use the ADSM NT client. Install the Microsoft Exchange Server on this new server, and enter an Organization Name and a

Site Name (see Figure 68 on page 141). Use the same names as the failed server.

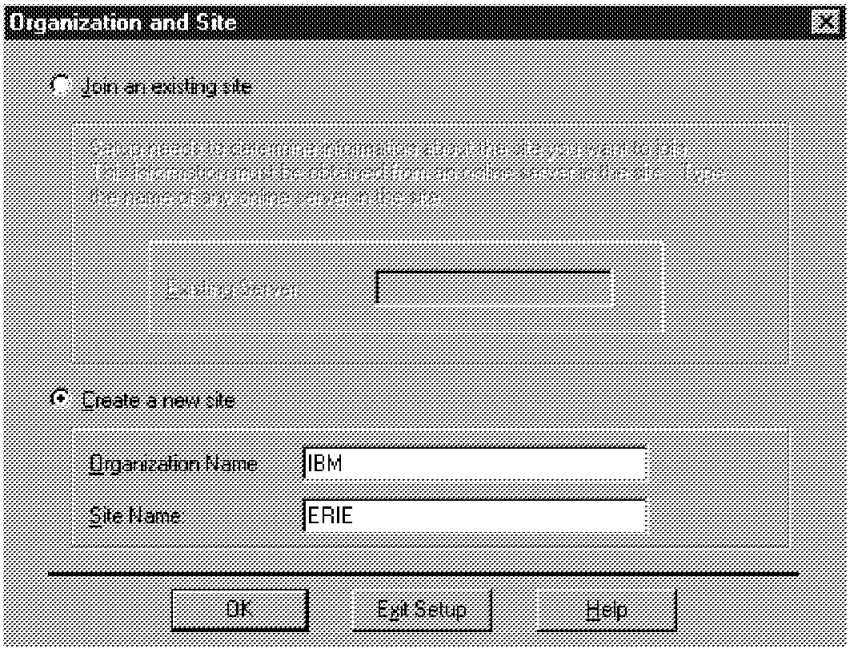

Figure 68. Microsoft Exchange Server setup option

2. Use ADSM to restore the Microsoft Exchange Server data from the old computer.

   Before the restore, you have to authorize the new server to restore the data from the old server:

   ```
   DSMC set access backup * ERIE
   ```

   To restore the data, use the ADSM GUI or command line:

   ```
   DSMC Restore -latest -fromnode=NODE -subdir=yes
        sourcefilespec destinationfilespec
   ```

   Node is the name of the ADSM client of the old server. You must know its password.

   sourcefilespec is the path and the name of files that you need to restore.

   destinationfilespec is the path and file name where you want to place the restored files:

   ```
   DSMC Restore -latest -fromnode=EXCHSERV1 -subdir=yes
        C:\exchsrvr\*
   DSMC Restore -latest -fromnode=EXCHSERV1 -subdir=yes
        E:\exchsrvr\mdbdata\*   D:\exchsrvr\mdbdata\*
   ```

3. In the Microsoft Exchange Server Administrator, select the server and open its property pages. Select **All Inconsistencies** and run the DS/IS adjustment program (see Figure 65).

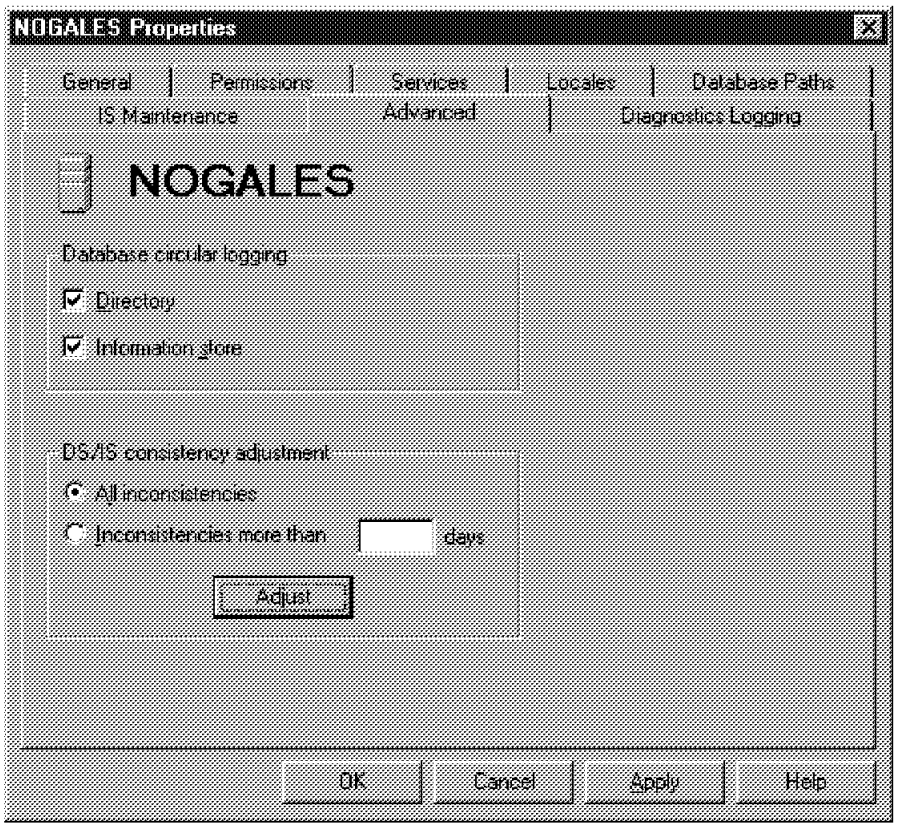

Figure 69. Microsoft Exchange Server DS/IS Adjustment Process

4. If you want to restore data from a private mailbox, you must grant yourself permission. If you want to restore a public folder, permission is not necessary. Move the data to a *.PST file and give it to the user.

7.6 Validating the Restore

You must validate the restore before users can use the Exchange application. Try to connect your own personal folders and the public information.

Some Microsoft Exchange Server data could be lost, for example, all data created after the last ADSM backup.

Evaluate the loss of data (see Figure 70 on page 143).

Send a Microsoft Exchange Server message to all users to explain which data has been lost.

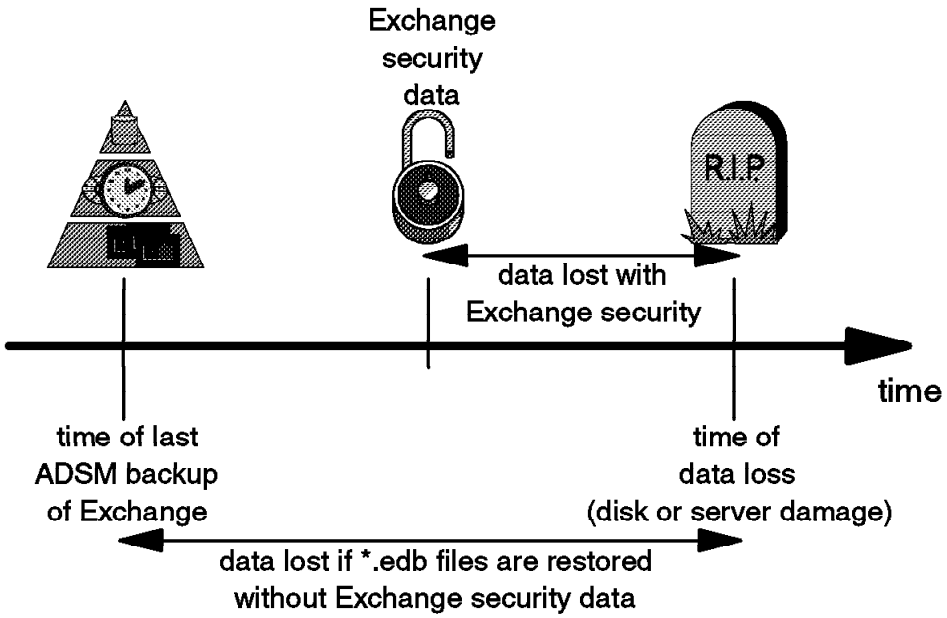

Figure 70. Evaluate Loss of Microsoft Exchange Server Data

If you only did an ADSM restore of all of the data, all mail and meeting data since the last ADSM backup will be lost.

If you have to perform an ADSM restore and a DS/IS adjustment, the data lost will depend on the IS Maintenance parameters (see Figure 71 on page 144).

Recovering the lost data is now a manual process!

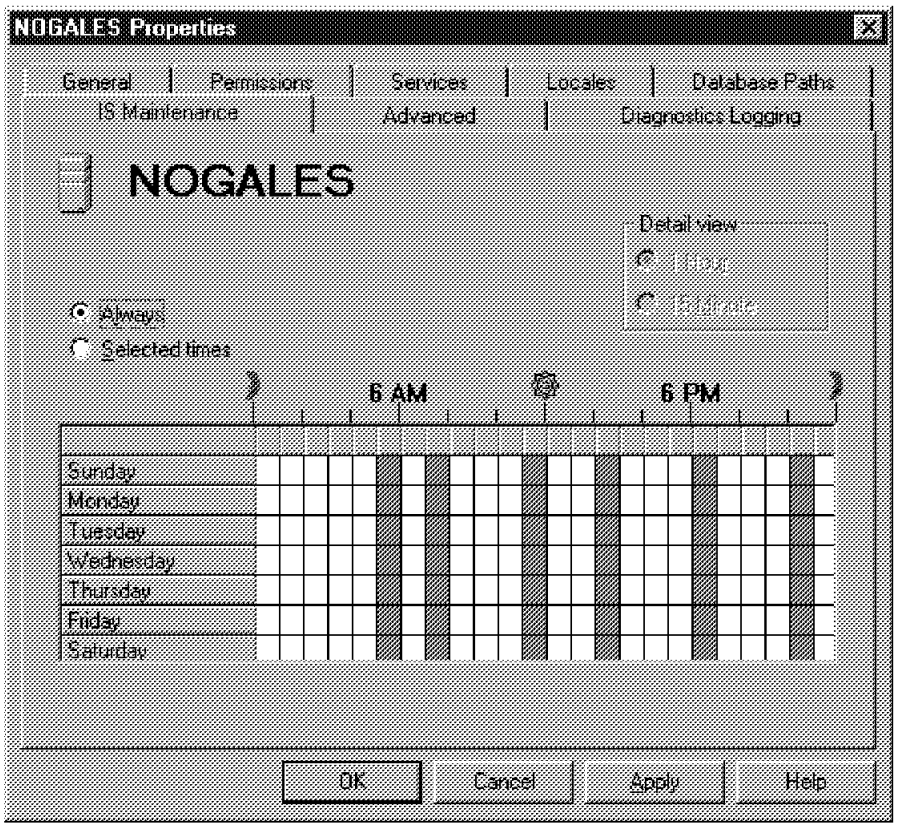

Figure 71. Windows NT Service Schedule Startup Option

Chapter 8. Recovering Microsoft Access

Microsoft Access, a relational database system, is very popular in small offices because of the fast and easy implementation of applications. Microsoft Access works in a logical divided environment, supports SQL, and uses a very fast algorithm for the queries. The Open Database Connectivity (ODBC) feature is also a part of Microsoft Access and could be used to connect to database servers, such as the MS SQL Server and IBM DB/2.

Microsoft Access is easy to implement in a multiuser environment, but there are some dangers. The need for backup in a network version of Microsoft Access is much important than in a single-user environment.

In this section we show some working aspects of Microsoft Access and explain how you can use ADSM to make incremental backups of an online database.

8.1 Data Characteristics

Microsoft Access can run in a multiuser environment on all Windows-supported networks, such as MS Lanmanager, Windows for Workgroups, NT, or Novell. The connection to another database must be implemented on the application level, because connections on the transport level (for example, TCP/IP) are not supported. In this section we use Version 2.0 of the program because it is more popular and the difference between it and version 7.0 is mostly in appearance to users. All examples are the same on both versions, but the screenshots may be different.

Microsoft Access is not a real database management system (DBMS), but the database engine and the database can be in separate locations. You can install Microsoft Access as a local program with the advantage of better performance and faster access to the network database, or you can install it as a network program with a small local part of the program. In the latter case, the program and data must be loaded over the network, which increases network traffic but facilitates administrative work and database security.

Because the path of the program is stored in the initialization files, Microsoft Access cannot be copied to another location. Therefore, a restore onto another drive is not possible. The ARC name convention is not included in Microsoft Access, Version 2.0 or 7.0 as it is in Windows NT. Although Version 7.0 is a real 32-bit application, it is not a real NT application, it is designed for Windows 95.

In the Workgroup Administrator program, you must first create a workgroup definition file, which will include the information about the database environment and the security of the database. This file is the SYSTEM.MDA in a defined directory.

After a workgroup is created, users can be allowed to join it. Through the security settings in Microsoft Access, one administrator is responsible for all tables, forms, and reports and can grant or deny access to any object in the database.

In the Microsoft Access options you must change some values for the multiuser environment (see Figure 72 on page 146). The **Default Record Locking** item tells Microsoft Access how to treat the edit requests from other users. You can set

the value to No locks, All Records or Edited Records. The locking information is stored in a file with the same name of the database file (MDB), but with the suffix LDB. The LDB file is important only during work with the open database and can be re-created if it becomes damaged. The settings for the refresh interval are important for the update of all open views. To use a database in a multiuser environment, the value for the **Default Open Mode** item must be Shared and not Exclusive.

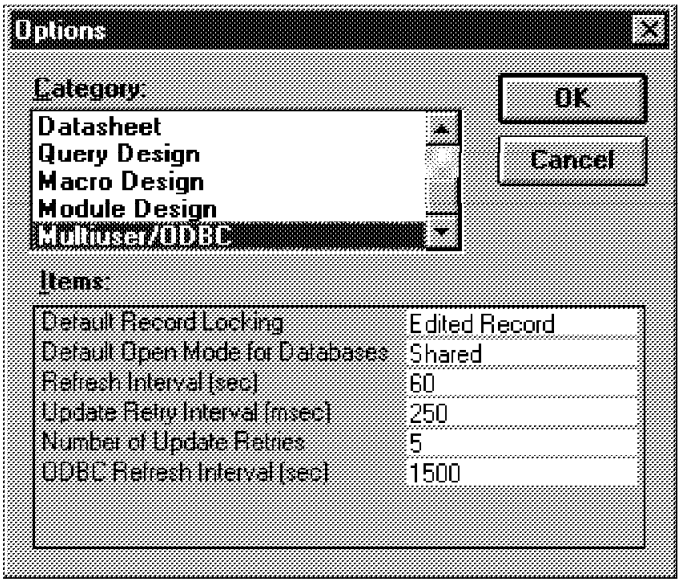

Figure 72. Multiuser Settings

Unfortunately, there is no setting for how long a user can open a record for editing. If a user opens a record, it is locked for all other users for additional editing. Other users have only a refresh interval which determines how often the information in the record will be updated in their view. They cannot start their own editing until the first user closes the record manually.

Microsoft Access has no log file or redo component. All user changes to the database appear directly in the database, when the user finishes editing. Because of the missing log file, it is very difficult to make a backup. The backup includes only a ″copy″ of the database at the moment of the backup. All changes made during the backup are not in the database and are not logged in a file.

Note: In other database systems you can restore the last backup and the log file. After that you must start a rollback of the log file to make all changes.

Now we must distinguish two types of access to a shared database object:

- A group can share and work in one database
- Every user has his or her own database and attached tables

If a group shares one database, there is only one file, the MDB file, which includes all tables, forms, and reports. This file is unique in the network. As a result, there is only one SYSTEM.MDA and one MDB file in the network, including all information.

If every user has his or her own database and uses the attach command to embed one or more tables from a database server or another Microsoft Access

client, he or she will have one SYSTEM.MDA file and one MDB file local on his or her machine. The locking mechanism depends on the setting of the server.

You must prepare different backup plans for these two database sharing approaches and include all needed files. In section 8.2, “ADSM Considerations” we show how ADSM can help you.

8.2 ADSM Considerations

An advantage of Microsoft Access is the one-file-use for all information. The MDB file contains all tables, forms, reports, and tablespaces. The number of files that must be backed up depend on how you share the database.

For the functionality of the locally installed Microsoft Access program, the MSACC20.INI file for Version 2 and the MSACC70.INI file for Version 7.0 in the `Winnt\` directory hold all environment information. For a reinstall of the program, this file is obligatory.

Note: The file exists in a locally installed Microsoft Access environment ***and*** in a network-based environment.

In the multiuser options setting you can only set a refresh value for the view updates. All changes are added when the user save the changes manually. Microsoft Access does not use a log file, so the only file for the database backup is the MDB file. The changes for all open records during the backup are not included in the backup. This missing information is usually stored in a log file, but unfortunately Microsoft Access has no log mechanism.

The locking information stored in the LDB file is not important and must not be backed up.

Again we must distinguish two types of sharing. If you are a member of a workgroup and use one MDB file for all users, that file exists only once in the network (Figure 73 on page 148). Thus you must back up the file with an ADSM client on this machine. The settings of the programs of all Microsoft Access client machines must be backed up with a locally installed ADSM client.

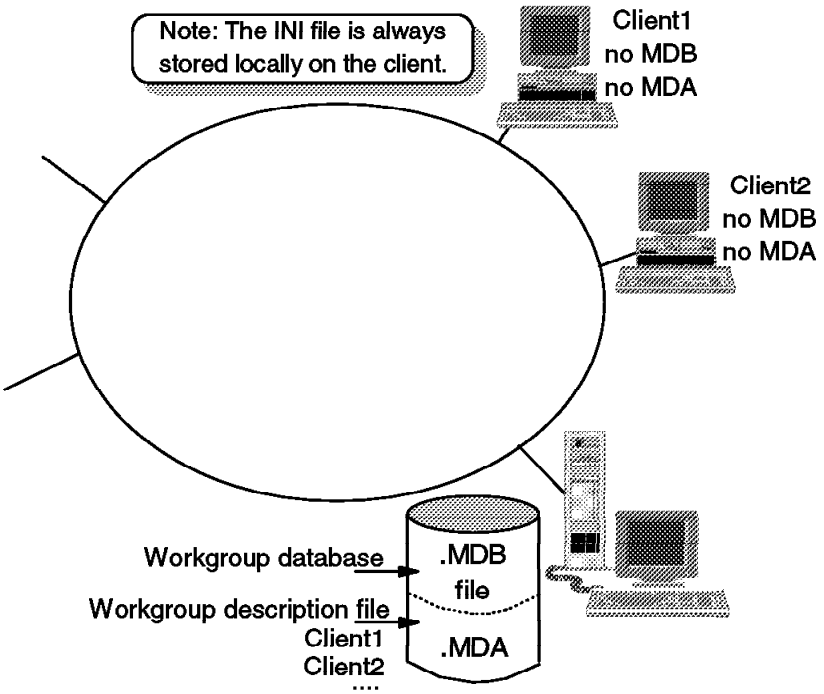

Figure 73. Workgroup Database

If you have your own database located on your own machine, you can share an attached table from a defined ODBC source, such as an MS SQL or a DB2 table. In this case you must back up your local setting file in the `Winnt\` directory and the local database file with the extension MDB (Figure 74 on page 149). The backup of the attached table must be done on the machine where the table is located.

The backup of a Microsoft Access database contains an image of the database at the moment of the backup. This behavior facilitates a point-in-time restore. We recommend that you make daily backups, always during the same timeframe, and store multiple copies in the backup pool of ADSM. If an error occurs in the database, you can restore the whole database from one specific backup.

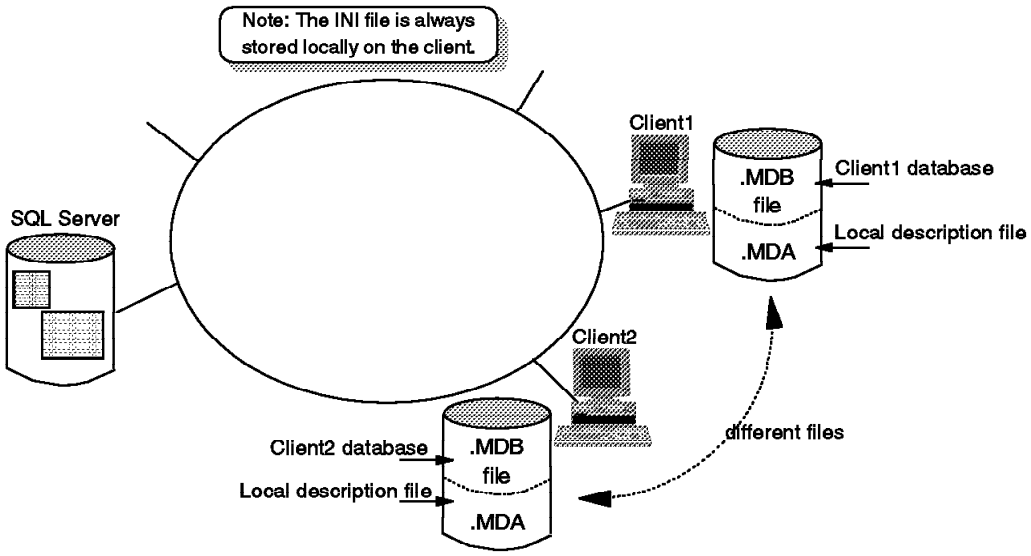

Figure 74. Working with Attached Tables

For smaller errors, such as deletion of one record by a user, the process of restore is difficult. It is not possible to restore any detail information of a Microsoft Access database with ADSM. You can only restore one image of the whole database.

If you have a lot of changes since the last backup and a user deletes one important record, you can restore that record by using Microsoft Access:

1. Restore the whole image of the database to another filename
2. Open the restored database with Microsoft Access
3. Export the record required
4. Close the restored database
5. Open the production database
6. Import the record required

Note: If you restore the database to the original place on disk with the same name, all changes since the last backup are lost; only the one deleted record is restored.

In case of a disaster, you have only the last backup of the database. All changes since the last backup are lost, including the changes of all open records during the backup process.

Chapter 9. Recovering DB2 for NT

DB2 is a relational database management system available on many different platforms. DB2 is well known as the relational database of choice for MVS mainframe systems and is now available on all major IBM and many non-IBM platforms.

Like the other products in the DB2 family, DB2 for Windows 95 and Windows NT supports SQL, which has become the industry standard way to define, update, and control data in a relational database. DB2 for Windows 95 and Windows NT also provides utilities that help maintain the contents of the database, including a built-in API for database backup/restore through ADSM. In this chapter, “DB2” refers solely to DB2 for Windows 95 and Windows NT.

9.1 DB2 File Structure

Under Windows NT, DB2 table spaces and recovery logs are stored in directories as shown in Figure 75.

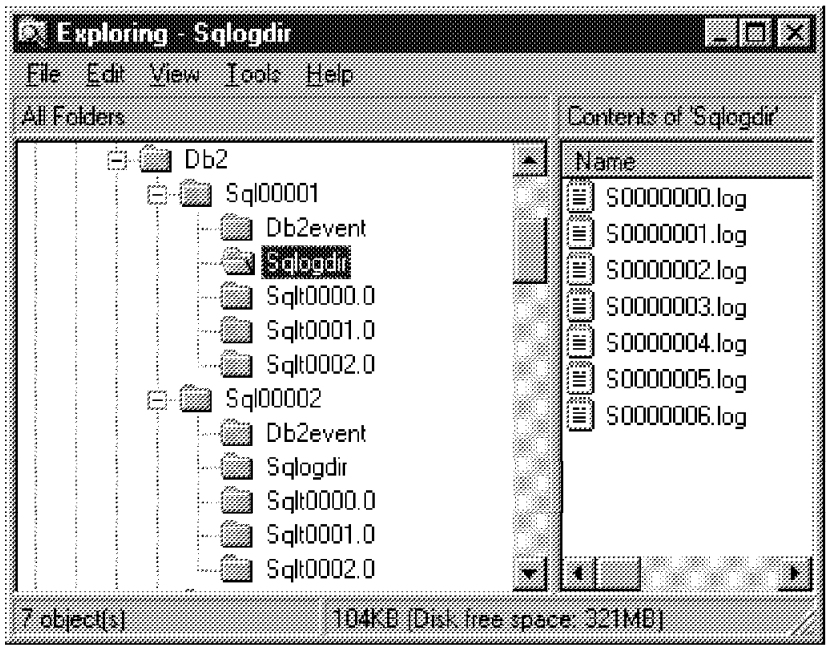

Figure 75. DB2 Filespaces in NT

For offline backup, it is this tree structure that you recover with ADSM. DB2 provides facilities for both offline and online backup. You may also perform offline backups during normal NT server backup, by including this tree structure in the include/exclude list in the DSM.OPT file.

9.2 Online Backup Using ADSM

It is easy to use ADSM as the backup medium for DB2 because the required APIs are built into the DB2 product. Figure 76 shows the DB2 Database Director Tree View with its SAMPLE database and associated table spaces and tables.

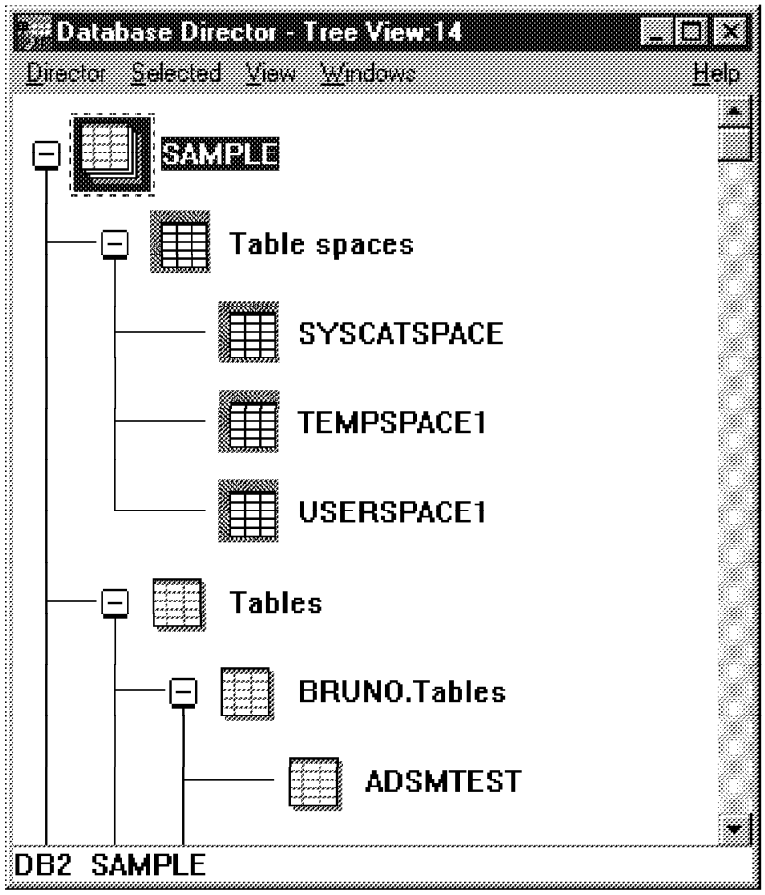

Figure 76. DB2 Database Director-Tree View

From the tree view in Figure 76, a right mouse click on an object brings up a window to access the backup command (Figure 77 on page 153).

Note

You must have DB2 SYSADM, SYSCTRL, or SYSMAINT authority to use the backup command.

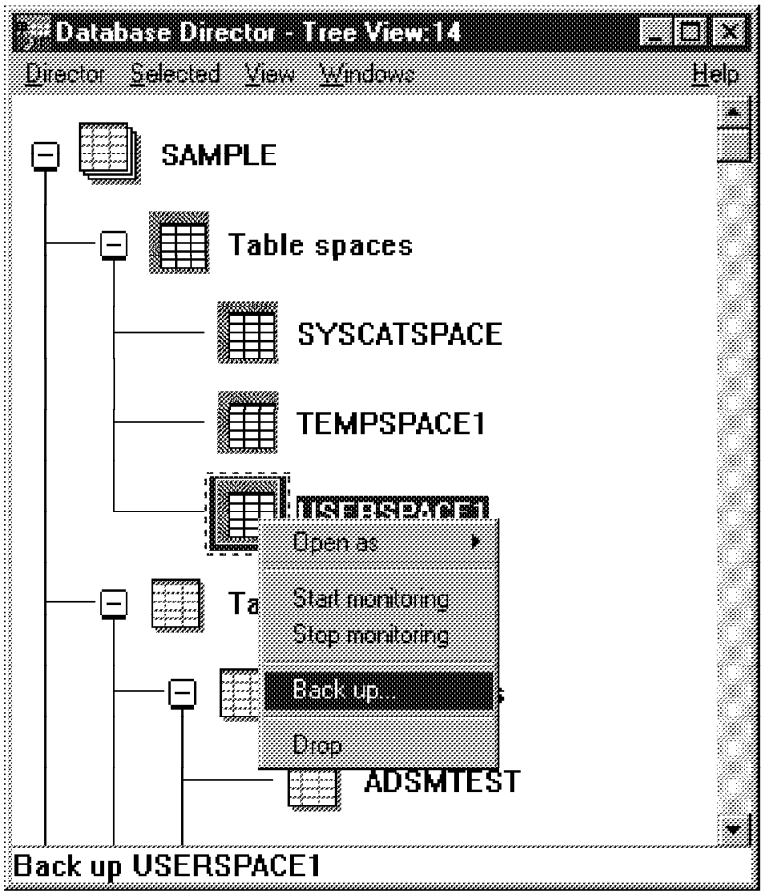

Figure 77. DB2 Database Director: Backup Table Space

Select Back up**...** to bring up the window shown in Figure 78 on page 154. In this window select **ADSM** as the backup medium and the **Online** backup.

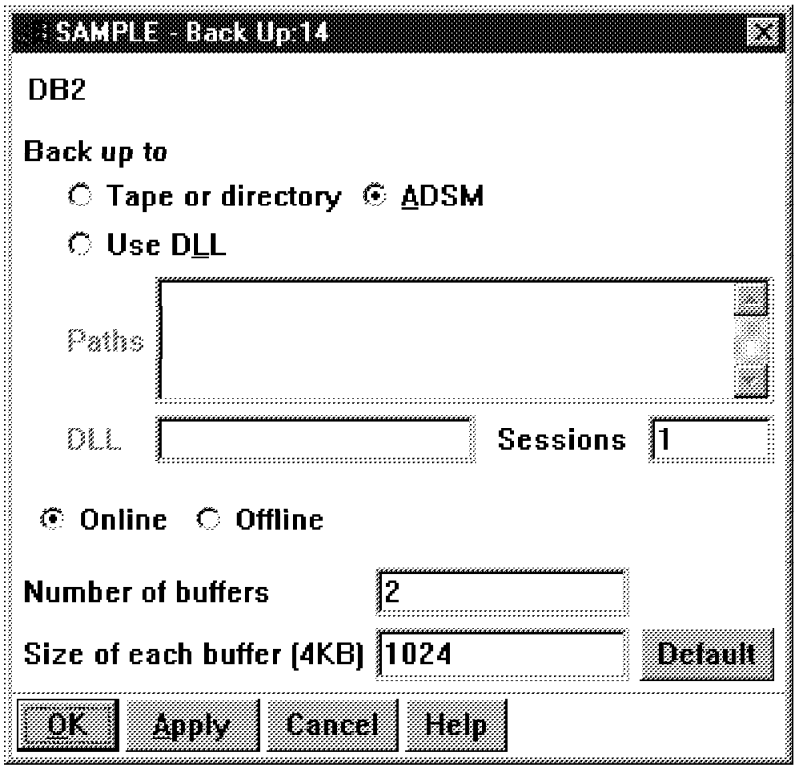

Figure 78. DB2 Database Director: Backup Command Window

Once you have submitted the backup job, you can monitor its progress through the DB2 Jobs Details View window. Figure 79 shows that the table space backup for SAMPLE has completed successfully.

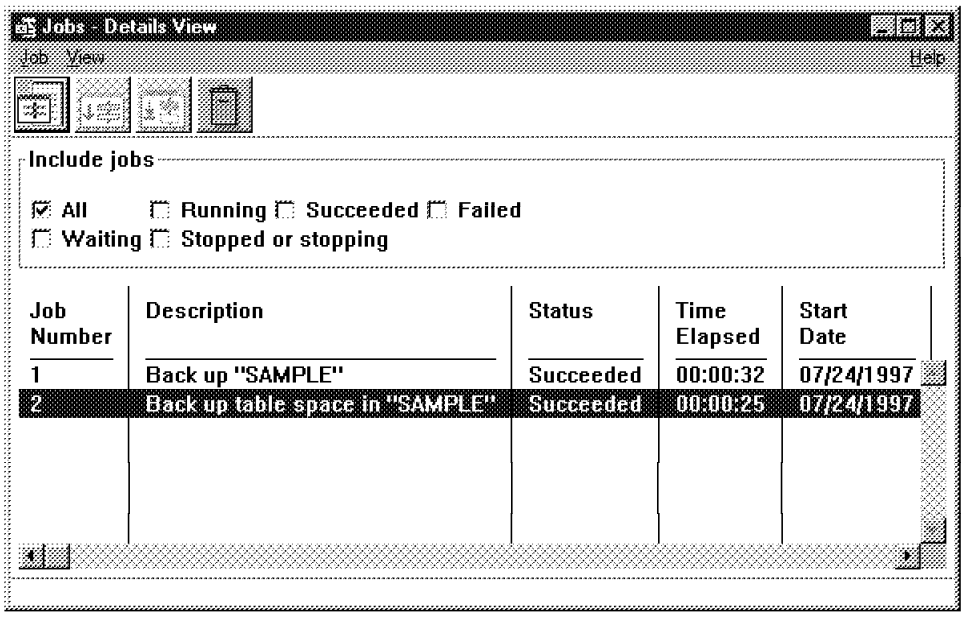

Figure 79. DB2 Database Director: Jobs - Details View List

To view the status of the log, double-click on the job shown in Figure 75 or select the job and press enter. The Job Output window shows the completion message and the timestamp for the backup (Figure 80 on page 155).

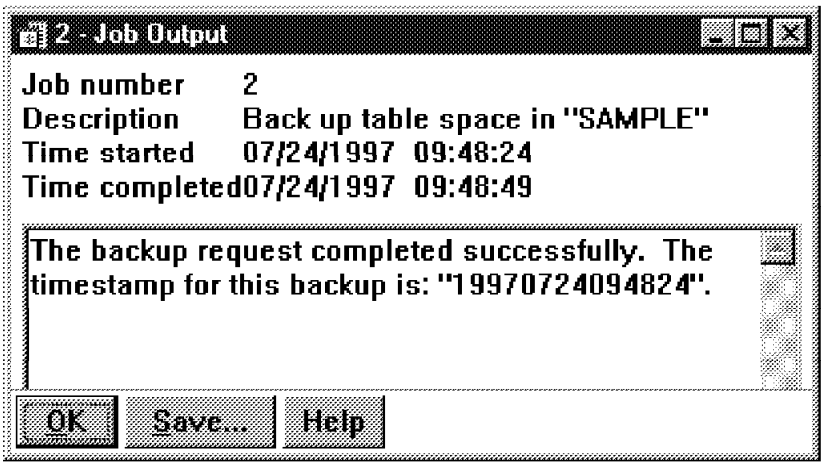

Figure 80. DB2 Database Director: Job Output

Similarly, you can restore the data to a point in time by using the GUI shown in Figure 81. Apply roll forward to the logs to complete the process.

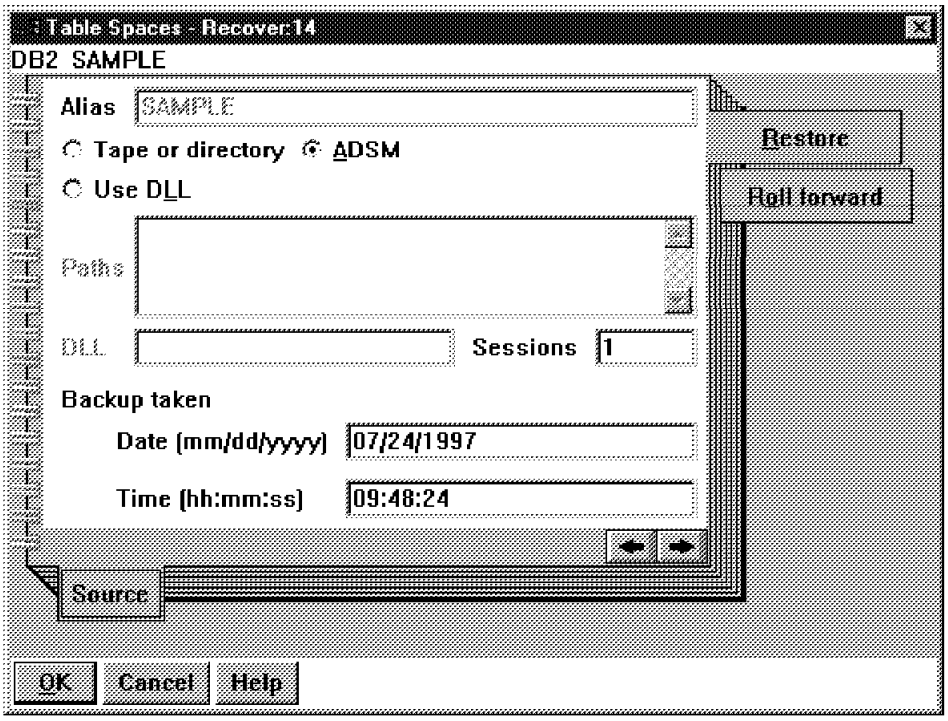

Figure 81. The DB2 Restore GUI

Appendix A. Additional ADSM Usage Information

A.1 ADSM Restore Commands

In this section we describe the necessary ADSM commands for the restore process and give some additional information. In this book we use only some of the options, because the assumption is the loss of the NT system.

Note: A complete list of command line parameters is available in the ADSM documentation.

In some locations in this book we show how to restore the whole NT system with ADSM. The complying command in the ADSM client is:

```
DSMC> RESTORE -NODENAME=nogales -SUBDIR=YES C:\*
```

The default for the SUBDIR parameter is NO. Therefore this parameter must always be set to YES. As a result the command restores all subdirectories including the empty directory for the registry and other important system directories.

The destination in this example is C:, which is the startable partition of the machine. Because of the ARC name convention, used by the NT system, it is easy to restore the NT system to any drive letter.

The value for the NODENAME option must be the name of the node that has performed the backup. If you want to restore the NT system from another ADSM node, you must use the FOMNODE parameter. In this case however, you must have the granted right of the previous node. For example, the backup is from the node ERIE and you want to restore the files to the node NOGALES. You must use the FROMNODE=ERIE parameter and you must know the password. The user of node ERIE or the ADSM administrator must allow you to restore the files. The complying ADSM command for this operation is SET ACCESS.

The above example assumes that the drive letter is the same as that of the backed up drive. ADSM actually substitutes the drive letter with a filespace name. The filespace name is the label of the partition and is used by the ADSM server to uniquely identify file systems on client workstations. If you have labeled the target partition for the restore with a different label, you must use the name of the filespace instead of the drive letter.

With this ADSM command

```
DSMC> QUERY FILESPACE -FROMNODE=erie
```

you can see the names of all filespaces for client node ERIE. You must have the the granted rights for this node as well.

The correct ADSM command for the restore of a named filespace is:

```
DSMC> RESTORE -NODENAME=nogales -SUBDIR=YES {c_drive}\*
```

In this example we use c_drive as the name of the filespace. If you want to define the target of the restore, you can add the location in the restore command. For example, restore to the F: drive with:

```
DSMC> RESTORE -NODENAME=nogales -SUBDIR=YES {c_drive}\* F:\
```

The options for the restore command could also be located in the DSM.OPT file of the ADSM client. But if you use ADSM as a network program, be sure you are using the correct options file.

We also use the **-REPlace** option, which determines what happens when files of the same name as the restore files already exist in the destination directory:

- -replace=prompt

  (Default). ADSM prompts you for your choice for overwriting existing files.

- -replace=no

  ADSM does not overwrite files.

- -replace=yes

  ADSM overwrites existing files, except for read-only files.

- -replace=all

  ADSM overwrites existing files, even if they are read-only.

A.2 ADSM Point-in-Time Restore

In this section we briefly examine the use of ADSM to restore all files to a point in time as opposed to trying for the latest copies.

Figure 82 shows what happens to files backed up by ADSM.

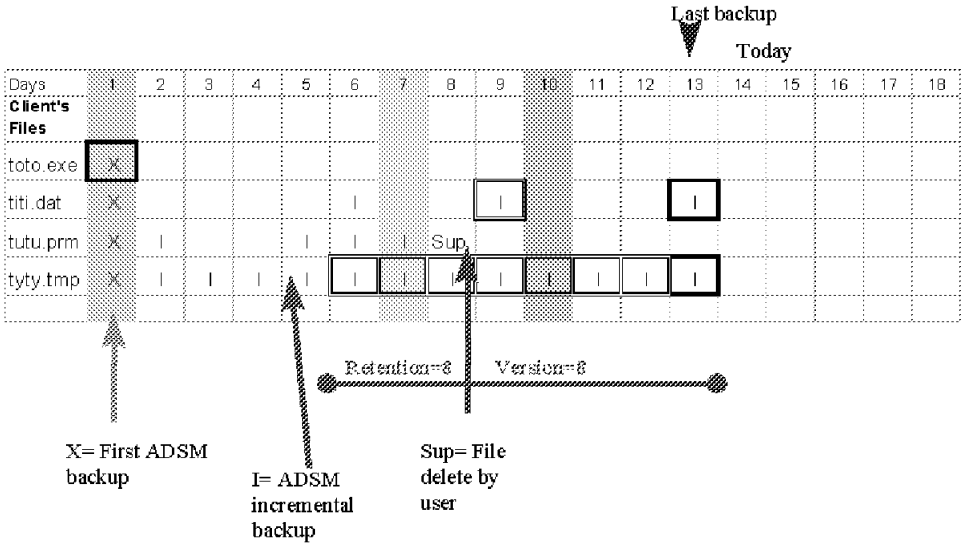

Figure 82. Different Versions of Backed Up Files

The different cases are:

- File toto.exe never changes. ADSM keeps the first backup as the active version.
- Files titi.dat and tyty.tmp sometimes change. ADSM keeps their last backup as the active version and the other backups as the inactive versions. The number of versions for these files depends on the Backup retention days specified in the Policy Domain Properties window of the ADSM Administrator

(see Figure 84 on page 160) and the Number of backup versions to keep (see Figure 83 on page 159).

- File tutu.prm is deleted on the eighth day. The number of backup versions depends of the Number of backup versions to keep if client data is deleted (see Figure 83).

A point-in-time restore, as shown in Figure 82 on page 158, must recover the files like this:

- On day 7 the toto.exe file backed up on day 1, the titi.dat file backed up on day 6, the tutu.prm file backed up on day 7, and the tyty.tmp file backed up on day 7 must be restored.
- On day 10 the toto.exe file backed up on day 1, the titi.dat file backed up on day 9, and the tyty.tmp file backed up on day 9 must be restored. The tutu.prm file must not be restored.

  Most backup software will not restore the tutu.prm correctly.

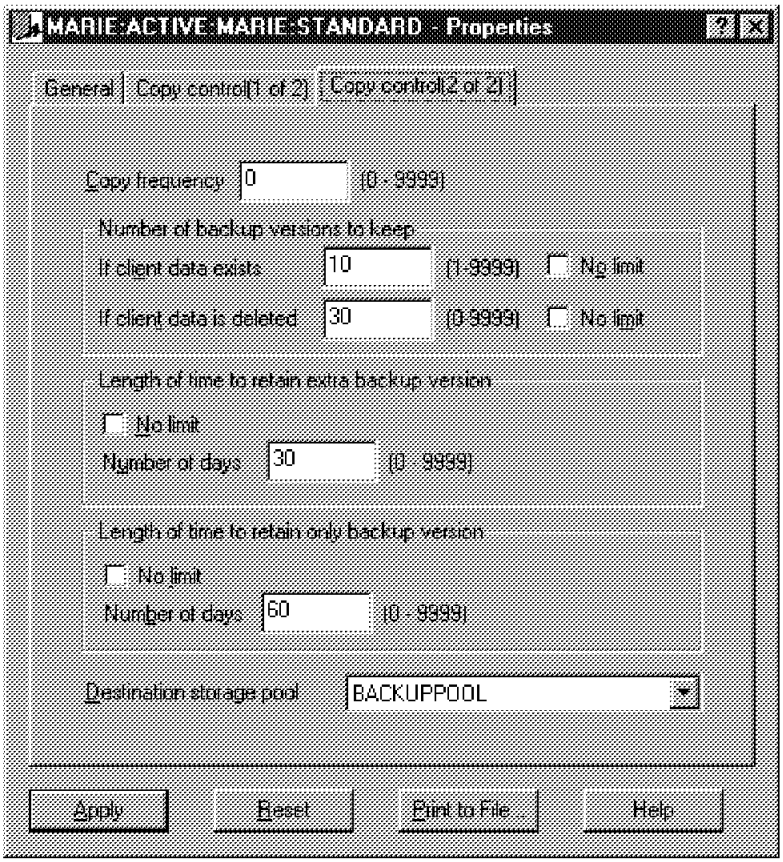

Figure 83. ADSM Number of Versions Parameters

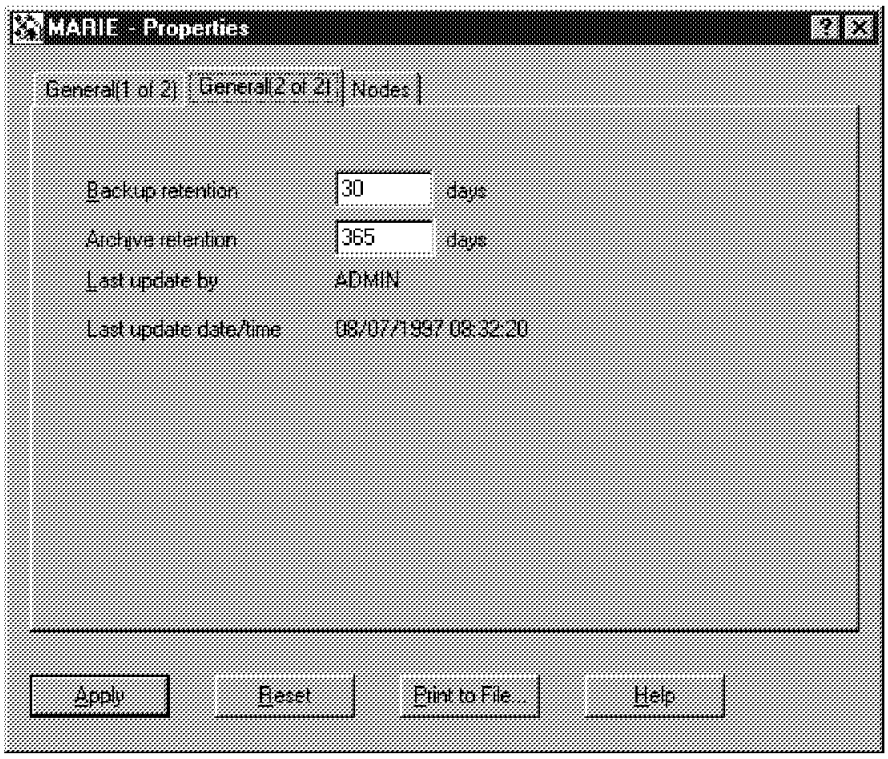

Figure 84. ADSM Retention Time Parameter

A point-in-time-recovery (or true image recovery) with ADSM depends on two parameters:

1. Number of versions to retain (see Figure 83 on page 159)
2. Retention time (see Figure 84)

In Figure 82 on page 158, the current date is day 14. The point-in-time restore is between day 6 and day 13.

If you want to do a point-in-time restore for a path, you have to use an alternative path for the recovery or delete the existing files in the path before the restore. Use the ADSM -todate and -totime parameters to select the appropriate version of the files.

A.2.1 Include/Exclude Sample Files

In this section we explain the include/exclude options to use in the DSM.OPT file of each ADSM NT client.

A.2.1.1 NT System Include/Exclude Sample File

The following are the default statements is the DSM.OPT file for the NT client:

```
Include *:\*.*
Exclude *:\...\pagefile.sys
Exclude *:\...\system32\config\*.*
Exclude *:\...\system32\config\...\*
```

We recommend excluding the Windows recycle bin, because of the huge file space to backup. In a typical NT installation, the reserved space for the recycle bin is about 10% of the disk space. On a 1 GB disk, the recycle bin can take up

to 100 MB of disk space. You can write an additional exclude line in the DSM.OPT file:

```
Exclude *:\...\recycler\*.*
```

For a Windows 95 machine the recycle bin is in the Recycled directory. The directory has the attributes **SH** and therefore is not viewable with the DIR command on NT or Windows 95.

A.2.1.2 Lotus Notes Include/Exclude Sample File

We recommend put these exclude options in DSM.OPT when you use the ADSMConnect Agent for Lotus Notes:

```
exclude *:\notes\data\readme.nsf\*
exclude *:\notes\data\help4.nsf\*
exclude *:\notes\data\helplt4.nsf\*
```

A.2.1.3 SQL Include/Exclude Sample File

The SQL Server is located in the %SqlRoot% directory, typically it is X:\MSSQL. A full backup of this directory is a minimum step for disaster recovery. As described in Chapter 6, “Microsoft SQL Server Backup” on page 103, the databases of the SQL server are locked. Therefore you must exclude these files from the ADSM backup/archive client:

```
exclude *:\mssql\data\*
```

If you are using the dump facility of the SQL server, you must include the backup directory, as shown below:

```
include *:\mssql\backup\*
```

A.2.1.4 Microsoft Exchange Server Include/Exclude Sample File

For the backup of Microsoft Exchange Server data, you have to put these include/exclude lines in the DSM.OPT files of each Exchange server:

```
exclude "*:\exchsrvr\mtadata\*.*"
include "*:\exchsrvr\dsadata\*.*"
include "*:\exchsrvr\mdbdata\*.*"
```

A.3 Using the ADSM Central Scheduler

This section shows an example on how to use the ADSM Central Scheduler Service to automate online backups with the ADSMConnect Agent.

Once the ADSMConnect Agent has been registered to an ADSM server, the procedure involves the following steps:

1. On the ADSM Server
   - Define a schedule to execute a Windows NT command file in the policy domain to which the ADSMConnect Agent is registered.
   - Associate the Agent node to the defined schedule.
2. On the Application Server where the ADSMConnect Agent is installed
   - Install the ADSM scheduler client as a Windows NT service for the Agent.
   - Define a command file that contains the Agent commands to do the desired backup.
   - Start the scheduler service

A.3.1 Example of Central Scheduler Service to Automate Backups

The example below uses the following assumptions:

- The Agent is registered to an ADSM server with a node name of **mars** and a password of **marspswd** ..
- The event to be scheduled is a daily incremental backup. The backups are to begin between 9:00 and 9:15 pm.

This method is flexible because you can define a command file with any set of commands you choose. This allows you to use the same method to schedule any ADSMConnect Agents on Windows NT.

On the ADSM server:

Enter the following command to define the schedule. You can enter this command on the server console or from an administrative client.

```
def sched agents daily_incr desc="Agent Daily Incremental
 Backup" action=command
objects="c:\incr.cmd" priority=2 starttime=21:00
```

ADSM displays this message:

```
ANR2500I Schedule DAILY_INCR defined in policy domain AGENTS.
```

To associate the Agent node to this schedule, issue the following command:

```
define association agents daily_incr mars
```

ADSM displays this message:

```
ANR2510I Node MARS associated with schedule DAILY_INCR in policy
domain AGENTS.
```

At this point, a schedule has been defined on the ADSM server that runs a command file that is called c:\incr.cmd. The schedule starts around 9:00 pm. The schedule is re-executed once a day and can start on any day of the week.

On the Application Server:

This example assumes that you have installed the ADSM client in the directory: c:\win32app\ibm\adsm\baclient and the ADSMConnect Agent in the directory: c:\win32app\ibm\adsm\agent.

Note: For the ADSMConnect Agent for MS SQL the directory will be c:\win32app\ibm\adsm\agentsql, for the ADSMConnect Agent for MS

It is also assumed that the options files in each of these directories has been updated so that the communication parameters point to the ADSM server.

1. Login using a Windows NT account that has administrative privileges.
2. Open a Windows NT command prompt window.
3. In the window, issue the following command:

   ```
   cd /d c:\win32app\ibm\adsm\baclient
   ```

   Note: If an ADSM scheduler service is already installed on your machine (for the regular backups of the Windows NT system), you should use a different node name from the regular ADSM backup client.
4. In the window, issue the following command:

```
dsmcutil inst /name:"ADSM Agent Scheduler" /node:mars
 /password:marspswd /autostart:yes
  /clientdir:c:\win32a/optfile:c:\win32app\ibm\adsm\agent
  \dsm.opt
```

An example of the output is shown below:

```
ADSM Windows NT Scheduler Service Configuration Utility
     Command Line Interface Version 1.00.d
     Last Updated 9-11-1997
     Command: Install and Configure ADSM Scheduler Service
     Machine: MARS (Local Machine)
     Installing ADSM Client Service:
     Machine : MARS
     Service Name : ADSM Agent Scheduler
     Client Directory : c:\win32app\ibm\adsm\baclient
     Automatic Start : Yes
     The service was successfully installed.
     Creating Registry Keys ...
     Updated registry value 'ImagePath' .
     Updated registry value 'EventMessageFile' .
     Updated registry value 'TypesSupported' .
     Updated registry value 'OptionsFile' .
     Updated registry value 'EventLogging' .
     Updated registry value 'ClientNodeName' .
     Updated registry value 'ADSMClientKey' .
     Generating registry password ...
     Authenticating password with ADSM for node MARS ....
     Password authentication successful.
     The Registry password for node MARS has been updated.
```

Note that the options file that is defined for the Agent is used by the scheduler when validating the node and password.

If you see the following message:

```
A communications error occurred connecting to the ADSM server
```

You should ensure that the dsm.opt file contains entries that point to the correct ADSM server.

5. The Exchange Agent must be running under the Exchange Site Services account in order to be able to access the Exchange backup API. Account information can be specified using the services applet in the control panel.
6. Now you must create the incr.cmd file. In the next section we show the incr.cmd file for each ADSMConnect Agent.
7. At this point the Central Scheduler Service is installed, but has not been started. To start the service, issue the following command :

   ```
   net start "ADSM Agent Scheduler"
   ```

   The following output is displayed:

   ```
   The ADSM Agent Scheduler service is starting.
   The ADSM Agent Scheduler service was started successfully.
   ```

   Note that because the /autostart:yes option is used, the ADSM scheduling service is automatically started each time. Your system is now ready to run automatic daily incremental backups of the application databases.

Appendix B. Special Notices

This publication is intended to help NT administrators prepare for the recovery of their systems using ADSM. The information in this publication is not intended as the specification of any programming interfaces that are provided by Adstar Distributed Storage Manager (ADSM). See the PUBLICATIONS section of the IBM Programming Announcement for Adstar Distributed Storage Manager (ADSM) for more information about what publications are considered to be product documentation.

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

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

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

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, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

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

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

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

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

IBM® S/390®

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

Java and HotJava are trademarks of Sun Microsystems, Incorporated.

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

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

Pentium, MMX, ProShare, LANDesk, and ActionMedia are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries.

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.

Appendix C. Related Publications

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

C.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see “How to Get ITSO Redbooks” on page 171.

C.1.1 ADSM Redbooks

Book Title	Publication Number
General Topics	
ADSM Concepts	SG24-4877
ADSM Version 2 Presentation Guide	SG24-4532
ADSM Version 3 Technical Guide	SG24-2236
A Practical Guide to Network Storage Manager	SG24-2242
ADSM Advanced Implementation Experiences	GG24-4221
Using ADSM Hierarchical Storage Management	SG24-4631
Client Disaster Recovery: Bare Metal Restore	SG24-4880
Server Books	
ADSM Server for Windows NT Configuration and Recovery Examples	SG24-4878
Getting Started with ADSM/6000	GG24-4421
ADSM for AIX: Advanced Topics	SG24-4601
AIX Tape Management	SG24-4705
ADSTAR Distributed Storage Manager/6000 on 9076 SP2	GG24-4499
ADSM for MVS: Recovery and Disaster Recovery	SG24-4537
ADSM for MVS: Using Tapes and Tape Libraries	SG24-4538
Getting Started with ADSM/2	GG24-4321
ADSM for OS/2: Advanced Topics	SG24-4740
Setting Up and Implementing ADSTAR Distributed Storage Manager/400	GG24-4460
ADSM/VSE Implementation Guide	SG24-4266
Specific Client Books	
Getting Started with ADSM NetWare Clients	GG24-4242
Getting Started with ADSM AIX Clients	GG24-4243
ADSM API Examples for OS/2 and Windows	SG24-2588
Windows NT Backup and Recovery with ADSM	SG24-2231
ADSM with Other Products	
Using ADSM to Back Up Databases	SG24-4335
Using ADSM to Back Up Lotus Notes	SG24-4534
Hierarchical Storage Management for NetWare: ADSM and AvailHSM Implementation	SG24-4713
Using ADSM to Back Up OS/2 LAN Server and Warp Server	SG24-4682
Backup, Recovery, and Availability with DB2 Parallel Edition on RISC/6000	SG24-4695

C.1.2 Tivoli Redbooks

Book Title	Publication Number
TME 10 Cookbook for AIX Systems Management and Networking Applications	SG24-4867
Understanding Tivoli’s TME 3.0 and TME 10	SG24-4948

C.1.3 General Interest Redbooks

Book Title	Publication Number
General Topics	
Guide to Sharing and Partitioning IBM Tape Library Data Servers	SG24-4409

C.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
Application Development Redbooks Collection	SBOF-7290	SK2T-8037
Personal Systems Redbooks Collection	SBOF-7250	SK2T-8042

C.3 ADSM Product Publications

Book Title	Publication Number
ADSM General Information	GH35-0131
ADSM V2 Installing the AIX Server and Administrative Client	SH35-0136
ADSM V2 for HP-UX Quick Start	GC35-0256
ADSM V2 Installing the MVS Server and Administrative Client	SH26-4043
ADSM V2 for Sun Solaris Quick Start	GC35-0262
ADSM V2 for AIX Administrator's Guide	SH35-0134
ADSM V2 for HP-UX Administrator's Guide	GC35-0257
ADSM V2 for Sun Solaris Administrator's Guide	GC35-0263
ADSM V2 for MVS Administrator's Guide	SH26-4039
ADSM V2 for AIX Administrator's Reference	SH35-0135
ADSM V2 for HP-UX Administrator's Reference	GC35-0258
ADSM V2 for MVS Administrator's Reference	SH26-4040
ADSM V2 for Sun Solaris Administrator's Reference	GC35-0264
ADSM V2 Messages	SH35-0133
ADSM V2 Device Configuration	SH35-0137
ADSM V2 Installing the Clients	SH26-4049
ADSM V2 Using the UNIX HSM Clients	SH26-4030
ADSM V2 Using the UNIX Backup-Archive Client	SH26-4052
ADSM V2 Using the OS/2 Backup-Archive Client	SH26-4053
ADSM V2 Using the Microsoft Windows Backup-Archive Clients	SH26-4056
ADSM V2 Using the Novell NetWare Backup-Archive Client	SH26-4055
ADSM V2 Using the Apple Macintosh Backup-Archive Client	SH26-4051
ADSM V2 Reference Cards for the Backup-Archive Clients	SX26-6013

C.4 ADSM Online Product Library

All of the ADSM publications are available in online readable format on the CD-ROM listed below. The ADSM library is also available on the following CD-ROMs:

CD-ROM Title	Publication Number
ADSM Online Product Library	SK2T-1893
MVS Base Collection Kit	SK2T-0710
VM Base Collection Kit	SK2T-2067
AS/400 Base Collection Kit	SK2T-2171

C.5 Tivoli Publications

Book Title	Publication Number
TME 10 Tivoli/Plus for ADSM User's Guide	GC31-8405
Tivoli/Courier Documentation Kit	SK2T-6046
Tivoli Enterprise Console Documentation Kit	SK2T-6050
Tivoli/Sentry Documentation Kit	SK2T-6052
Tivoli/Management Platform Documentation Kit	SK2T-6058
Tivoli/ADE Documentation Kit Volume 1	SK2T-6062
Tivoli/ADE Documentation Kit Volume 2	SK2T-6063
Tivoli/ADE Documentation Kit Volume 3	SK2T-6064
Tivoli/ADE Documentation Kit Volume 4	SK2T-6065

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at URL http://www.redbooks.ibm.com.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **PUBORDER** — to order hardcopies in United States
- **GOPHER link to the Internet** - type GOPHER.WTSCPOK.ITSO.IBM.COM
- **Tools disks**

  To get LIST3820s of redbooks, type one of the following commands:

  ```
  TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
  TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
  ```

  To get lists of redbooks:

  ```
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LISTSERV PACKAGE
  ```

  To register for information on workshops, residencies, and redbooks:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1996
  ```

  For a list of product area specialists in the ITSO:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
  ```

- **Redbooks Home Page on the World Wide Web**

  http://w3.itso.ibm.com/redbooks

- **IBM Direct Publications Catalog on the World Wide Web**

  http://www.elink.ibmlink.ibm.com/pbl/pbl

  IBM employees may obtain LIST3820s of redbooks from this page.

- **REDBOOKS category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL
- **Internet Listserver**

  With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibmlink.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** (Do not send credit card information over the Internet) — send orders to:

	IBMMAIL	**Internet**
In United States:	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada:	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America:	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+45) 4810-1170 - Swedish

- **Mail Orders** — send orders to:

IBM Publications	IBM Publications	IBM Direct Services
Publications Customer Support	144-4th Avenue, S.W.	Sortemosevej 21
P.O. Box 29570	Calgary, Alberta T2P 3N5	DK-3450 Allerød
Raleigh, NC 27626-0570	Canada	Denmark
USA		

- **Fax** — send orders to:

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
(+45) 48 14 2207 (long distance charge)	Outside North America

- **1-800-IBM-4FAX (United States)** or **(+1)001-408-256-5422 (Outside USA)** — ask for:

  Index # 4421 Abstracts of new redbooks
  Index # 4422 IBM redbooks
  Index # 4420 Redbooks for last six months

- **Direct Services** - send note to softwareshop@vnet.ibm.com
- **On the World Wide Web**

Redbooks Home Page	http://www.redbooks.ibm.com
IBM Direct Publications Catalog	http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Internet Listserver**

  With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibmlink.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank).

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

- Invoice to customer number
- Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

DO NOT SEND CREDIT CARD INFORMATION OVER THE INTERNET.

Index

E

F

H

I

L

M

N

O

P

R

S

T

U

V

W

Z

ITSO Redbook Evaluation

Windows NT Backup and Recovery with ADSM
SG24-2231-00

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

- Use the online evaluation form found at http://www.redbooks.com
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redeval@vnet.ibm.com

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes____ No____

If no, please explain:

_____

_____

_____

_____

What other redbooks would you like to see published?

_____

_____

_____

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

_____

_____

_____

_____

_____

SG24-2231-00
Printed in the U.S.A.