SBAdmin

System Backup Administrator

USER GUIDE

Version 5.3

STORIX™

SOFTWARE

Trademarks and Copyrights

Publicly Available Software

This product either includes or is developed using source code that is publicly available:

AESCrypt*	Rijndael and Cipher Block Feedback mode (CFB-128) encryption/decryption algorithms	Copyright 1999,2000 Enhanced Software Technologies Inc. http://aescrypt.sourceforge.net/
BusyBox	Single executable containing tiny versions of many common UNIX utilities.	Copyright 1989, 1991 Free Software Foundation, Inc. http://busybox.net/cgi-bin/cvsweb/busybox/
LILO	LInux boot LOader	Copyright 1999-2003 John Coffman. Copyright 1992-1998 Werner Almesberger. http://freshmeat.net/projects/lilo/
parted	Linux partition editor for EFI GPT partition tables	Copyright (C) 1999, 2000, 2001, 2002, 2003, 2005 Free Software Foundation, Inc.
sfdisk	Linux partition editor for MSDOS partition tables	Copyright (C) 1995 Andries E. Brouwer (aeb@cwi.nl)
Tcl	Open source scripting language	Copyright Regents of the University of California, Sun Microsystems, Inc. http://tcl.sourceforge.net
Tk	Tk graphics toolkit	Copyright Regents of the University of California, Sun Microsystems, Inc. http://tcl.sourceforge.net
Xpdf	PDF Document viewer (for AIX)	Copyright 1996-2003 Glyph & Cog, LLC. http://www.foolabs.com/xpdf
Yaboot	IBM CHRP Bootloader for Linux on pSeries	Copyright 2001-2003 Ethan Benson Copyright 1999-2001 Benjamin Herrenschmidt Copyright 2001-2003 Peter Bergner http://penguinppc.org/projects/yaboot/

*Encryption Software

System Backup Administrator Backup Data Encryption Feature has a cryptographic component, using ***Advanced Encryption Standard (AES)*** "Rijndael" encryption algorithm in Cipher Block Feedback (stream) mode (CFB-128), supporting 128, 192 and 256-bit keys.

It is not for export or redistribution to any of what are called the "T-10 Terrorist States" as determined by the U.S. Department of State. System Backup Administrator Backup Data Encryption Feature has been registered with U.S. Bureau of Information and Security and is distributed under Export Control Classification Number (ECCN) 5D992. This encryption item is authorized for export and re-export under section 742.15 (B)(2) of the Export Administration Regulations (EAR).

Table of Contents

1. Getting Started

Supported Operating Systems & Hardware

As of the time of this publication, the software is supported on the following systems:

AIX: All IBM *RS/6000*, *pSeries*, *OpenPower* and *JS/20* systems running AIX Version 4.3 and later (currently 5.3).

Linux: ***X86***: All distributions which run on *Intel 32-bit* based processors and 64-bit processors capable of running 32-bit software (includes *AMD*, *Opteron* and *Athlon*-based systems). Linux kernel levels 2.4 and glibc 2.2.2 and higher are required. Support is provided for Linux LVM Library version 1.0 and higher, and Software Raid Devices (meta-disks) when installed.

PPC: All distributions supported on 32-bit or 64-bit systems with *PowerPC CHRP* hardware. Linux kernel levels 2.4 and higher, and glibc 2.3.3 and higher are required. Support is provided for Linux LVM Library version 1.0 and higher, and Software Raid Devices (meta-disks) when installed.

IA64: All distributions supported on *Intel* 64-bit IA64 (*Itanium*) processor-based systems, including *HP Integrity* systems. Linux kernel levels 2.4 and higher, and glibc 2.3.3 and higher are required. Support is provided for Linux LVM Library version 1.0 and higher, and Software Raid Devices (meta-disks) when installed.

Software and License Requirements

Installation of the software provides the graphical user interface and application programs for administering the backups of the administrator system itself. If the *Network Administrator* license is installed, administration of client system backups and the backup media servers may also be performed from the administrator system. It is also necessary to install a subset of the software onto each system that will act as either a backup media server or client.

The following table describes each license type:

Desktop Edition	This option provides backups to both disk directories and tape drives. This option is only for personal (non-commercial) use. Options provided are those typically required for a home system. This option includes **Full System**, **Filesystem** and **Directory Backups**, with some additional features, such as exclude lists and auto-verify. ***Desktop Edition is not available for AIX or Linux/PPC systems***
Workstation Edition	This more advanced option provides all available backup and recovery features for local system backups to tape or disk directory, and is available for commercial use. This includes all available features needed for standalone system backups. Backup types include Full System, **Filesystem**, **Directory**, **Logical Volume**, **Meta-disk** (Linux) and **Partition** (Linux). Many additional features more commonly used in a commercial environment, such as incremental backups and tape libraries, are also provided.

Network Administrator	This license is installed only onto the system from which network backups will be centrally administered. This system may also be a client or backup server, but this is not a requirement. This option includes all features of the *Workstation Edition*, but allows backups of other systems configured with a *Client/Server* license to be managed by the local *Network Administrator*. A license key is required for the *Network Administrator*, which also defines the number of clients and/or backup servers which may be managed by the administrator. Although no license key is required for each of the clients or backup servers, the client/server software must be installed and configured on each system before they may be managed by the Network Administrator.
Client/Server	Must be installed on each system which will be a client or backup media server. A client/server license for the local system is included with the *Network Administrator* license. If installed separately, this client must be controlled by a *Network Administrator*. No license key is required on the client or backup server since the number of supported clients and backup servers are defined by the Network Administrator license. Backup management features, such as scheduling and history reporting are provided only by the Network Administrator.
Backup Data Encryption Feature	This optional license may be added to a *Network Administrator* or *Workstation Edition* to enable AES data encryption support for all backups. If used with a Network Administrator, a license is purchased for the number of *Clients* for which backup data should be encrypted.
Windows (SMB) Data Backup Feature	This optional license may be added to a *Network Administrator* to allow SMB (Win) Share data to be backed up. This feature requires at least one ***Linux*** *Client/Server* system acting as the “SMB host”. SMB File Sharing Protocol is used on a variety of systems, most commonly Microsoft Windows and Mac OS X. This license is purchased for the number of SMB systems for which data contained in “shares” (or “shared folders”) are to be backed up.

Evaluation License Key

All license options and features above, except the **Client/Server** require a license key. This key is unique to each system that the software is installed onto, and must be obtained from Storix. Wherever a license key is required, the user may type the word “**trial**” for a free 30-day evaluation of all features of the software.

Software Installation and Configuration

The following instructions may be used to install the software from either installation images downloaded from the Storix Software web site (http://www.storix.com) or from a System Backup Administrator installation CDROM:

Downloading and Installing from the Web Site

1. Select the software package you wish to download from the web site based on your ***operating system type, machine type*** and desired ***software configuration***.

Be sure to download the file in BINARY. Some browsers will recognize the ".tar" extension of the file and ask you if it should open the file or expand it. You should NOT do so, but select to save it to disk

2. Change to the /tmp directory:

```
cd /tmp
```

3. Extract the contents of the file. Note that this does not extract the software, but only the installation program files and install image:

```
tar -xvf IMAGEFILE.tar
```

(where *IMAGEFILE*.tar is the name of the downloaded file)

4. Run the installation program by typing:

```
./stinstall
```

Installing from CDROM

1. Mount the cdrom by typing:

   a. On ***AIX*** systems: `mount -v cdrfs -r /dev/cd0 /mnt`

   b. On ***Linux*** systems: `mount –t iso9660 –r /dev/cdrom /mnt`

2. Run the installation program by typing the following, then follow the instructions provided:

```
/mnt/stinstall
```

3. When complete, unmount the CDROM by typing:

```
umount /mnt
```

Updating the Software

To update the software connected to the internet, you can automatically check, download and apply updates directly from the Storix Web Server by selecting Help→Download Software Updates from the user interface. A screen similar to the following will appear (note the option for client updates only appears on *Network Administrators*):

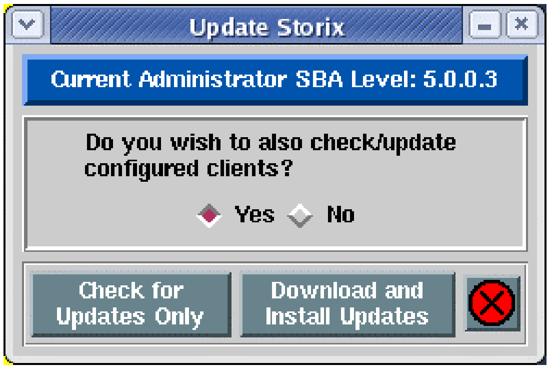

Update Storix

Current Administrator SBA Level: 5.0.0.3

Do you wish to also check/update configured clients?

● Yes ○ No

You will have an option of checking for updates only and/or downloading and installing updates. If the system is a Network Administrator, you will have an additional option of automatically applying updates to configured clients.

If the system cannot contact the Storix Web Server directly, you may apply updates by re-installing the software using the same instructions used to initially install the software (shown above). When you re-install the software onto the network admin system using the "**stinstall**" command described above, you will be asked if you wish to install the new software level onto configured clients and servers.

NOTE **Re-installing the software will replace existing program files, but WILL NOT OVERWRITE current configuration or history files.**

Starting the Software

To access the graphical user interface, also referred as the "***backup administrator***", type:

```
sbadmin &
```

from within an *xterm* window. If you wish to run the application on a display attached to a different host (perhaps even a PC running an Xwindows emulator), type:

```
sbadmin –display hostname:0 &
```

(where *hostname* is the host name of the remote system). It may also be necessary to provide access to the application to write to the display by first typing "`xhost +`" within an *xterm* window on the remote system.

When starting the administrator software, the Main Screen will appear.

For *AIX* Systems:

Many options performed by the administrator may also be run on each client or server using SMIT. Instructions for using the SMIT options are provided from the SMIT menus by selecting the **Help Menu** option or selecting the **Help** button (graphical SMIT) or the F4 key (ASCII SMIT) at each menu or entry field.

To access the SMIT menus, simply type:

```
smit storix
```

at the command line. As this document is intended to provide instructions on the Backup Administrator graphical user interface used on the administrator system, no further instructions are provided in this document for the SMIT options. Refer to the SMIT help panels for additional information on the menus and options provided.

For systems configured as a ***Client/Server***, ***Network Administrator*** or ***Workstation Administrator***, options are provided for performing backups, verifies and restores from the AIX SMIT menus. However, all configuration and maintenance options must be performed using the graphical user interface on the administrator system. Configuration options are not provided in the SMIT menus on each client or backup server because they would be overridden by any changes made using the graphical application on the administrator system.

Enabling Optional Features

Optional features, such as ***Backup Data Encryption*** and ***Windows (SMB) Data Backups*** may be enabled after the *Network Administrator* or *Workstation Edition* has been installed. To enable these features, select File→Preferences→Software License from the menu bar on the Main Screen. Refer to Software License in the Preferences section for details on viewing and changing the license options.

2. Introduction

System Backup Administrator (SBAdmin) is designed to simplify the administration of backups on the local system as well as client backups in a networked environment (when the *Network Administrator* license is installed). It does so by combining powerful backup tools with an easy-to-use graphical interface for administering backups of an unlimited number of client systems from a central system. Backups created by the *Backup Administrator* application may include single directories or entire systems that may be used to reinstall the source system or another system with an entirely different disk configuration. Backups may be automated through the use of a backup scheduler and queuing system, and client systems may be installed from backups on a network server.

This document will provide a description of all of the functions of the Backup Administrator, and will include instructions for performing common tasks. For additional detailed information on each option within the application, you may get on-screen help by simply clicking the right mouse button over the object in question.

This document is intended only to provide instructions on the use of the *Backup Administrator* (graphical) interface used with the *Network Administrator* , *Workstation Edition*, and *Desktop Edition* licenses. Some tasks may be performed directly on a client or server using **SMIT** menus (on ***AIX*** systems), or running commands from the command line.

> **The remainder of this document provides instructions on the use of the *Backup Administrator* graphical user interface. The *Network Administrator* is used in the examples throughout this guide. Options which are not applicable to the Workstation or Desktop Editions are noted.**
>
> **The Commands chapter is provided for information on running commands at the command line, some of which may be used to perform backups, verifies and restores from clients without using the network administrator. It also describes a number of commands which may be used to perform backup administrator-related functions.**

Terminology

It is important to understand the relationship between the different systems that will interact with the Backup Administrator software:

- **Admin System** - This is the system running the Backup Administrator software. When using a *Network Administrator*, all backup servers, clients, and backup options are configured and maintained from the admin system, and the admin system will centrally perform all tasks for the servers and clients, including scheduling and run the backup jobs, monitoring backups, performing verifies and restores, and even recreating volume groups and filesystems. For a ***Standalone System*** running *the Workstation* or *Desktop Edition*, the admin system, single client and server are assumed to be the local system.

- **Backup Server** - This is the server on which the backup media is attached, sometimes referred to as a backup media server. Backup media may be a tape drive, set of tape drives, tape autoloaders, or disks on which backup files will be saved. Any system on the network may act as a backup server, and multiple backup servers may be used. Select this link for information on configuring a backup server.

  Note: When using *Workstation* or *Desktop Edition*, the admin system will always act as the backup server. Therefore, references to the backup server in this manual refer also to the admin system.

- **Client** - This is the system from which backups will be made. The admin system or any backup server may also be configured as a client, since they also need to be backed up. Any client may also be configured as a server. A client will be defined as either an ***AIX*** or ***Linux*** (UNIX) client, or as an ***SMB (Windows) client***. Select this link for detailed information on adding or removing a backup client.

Note: When using a *Workstation* or *Desktop Edition*, the admin system itself is assumed the only client. Therefore, references to the client in this manual refer also to the admin system

With the *Network Administrator*, the backup clients and servers, as well as the tape device on the backup servers may be displayed on the main screen of the application. The application will constantly monitor the status of the clients, servers and devices, and the icons on the screen will represent whether or not the system or device is available.

Additional terms are commonly used in this document and in the application:

- **Backup Profiles** - Any number of backup profiles may be created, which will contain the backup defaults to be used when performing a backup job. This prevents the need to answer the same questions repeatedly when configuring backup jobs. At least one backup profile must be created for each type of backup to be performed. Select this link for detailed information on adding or removing a backup profile.

- **Backup Jobs** - A backup job will contain all the information needed to perform a backup, including the client(s) to backup, the server to backup to, and the specific device or directory on the backup server to use. A backup profile will be assigned to the job, which will provide most of the common backup defaults. The information in the profile, however, may be customized for each job. A backup job is identified by a Job ID and may be scheduled to run only upon demand, only once at a specific date and time, or scheduled to run on a regular basis. A backup job may contain one or more clients. If multiple clients are included in a single job, the data for all clients is appended to the same tape (or set of tapes), or stored in the same backup files (if written to disk). When writing backup to tape, multiple backup jobs may also be appended to the same tape or set of tapes. Select this link for additional information on creating, scheduling and running backup jobs.

- **Job Queues** - The Backup Administrator provides a queuing system that prevents multiple backup jobs from attempting to write to the same devices at the same time. A queue is defined for each device (and one for disk backups) on each backup server for which a backup job is scheduled. Backup jobs are added to the queues when they are run. The queues may be displayed in the main screen of the application, providing an easy glance at the queue contents and the status of queued jobs, and action buttons for manipulating the queued jobs. The jobs may be started, stopped, removed from the queue or placed on hold. Running jobs may be monitored, displaying the backup progress and/or the backup output messages. Select this link for more detailed information on backup queues and how to manipulate backup jobs in the queue.

Understanding Backup Media

Tape Backups

A tape backup may consist of a single tape in a single tape drive, multiple tapes from a single tape drive, or multiple tape from multiple tape drives. For simplicity, the term "*tape*" or "*tape backup*" may refer to any of these.

When writing backups to tape, each filesystem and logical volume is stored in a separate backup file, allowing the tapes to be quickly forwarded to the desired data for faster restores.

A tape might contain a single backup job, and the job might contain only a single client. The tape may also contain multiple backup jobs, each containing one or more clients. A single client backup on the tape is identified by its backup sequence number. The backup sequence number begins with 1 (the first client backup on the tape) and is incremented for each additional client backup performed to the same tape.

The *Backup Administrator* keeps track of the contents of a tape. At any time, the administrator may display or print the backup label, which contains a list of the client backups and corresponding sequence numbers.

It is usually a good idea to print the backup label and store it with the backup tape. If the printed label is lost, the Backup Label ID may be read from the tape and the label information may again be displayed or printed.

A backup tape may also be identified by a Tape Label ID. If desired, the user must write a *unique* tape label id to each tape that will be used with the Backup Administrator. Often tapes come with physical tape labels with a unique tape id printed on it. This label may be physically applied to the tape and the tape label id may be written to the tape media using the Backup Administrator. After doing so, that tape label ID will be associated with any backups written to that tape. The backup label may be displayed given the tape label ID and the tape label IDs used with a backup will be displayed within the backup label.

The *Backup Administrator* backup retention policy ensures that you do not accidentally write over a prior backup by reading the label from the tape before each backup is performed to the beginning of the tape. If the backup label is current, the backup will fail with an error message before the tape is overwritten. Tapes may be overwritten only after the tape is expired. By manually expiring a tape, the label information is removed from the database and the tape may be reused. The administrator may also set the overwrite policy to allow current backup tapes to be overwritten. If so, the tape label will be automatically expired when a new backup is written at the start of the tape. The global overwrite policy may be explicitly overridden for each backup job.

Multiple tape drives may be combined into a single "virtual device", providing increased performance and capacity over a single tape drive. There are three types of virtual devices for performing sequential, parallel or multi-copy backups. Virtual devices may also be configured as a Sequential Autoloader or a Random Tape Library. Refer to Types of Virtual Devices as described in detail in the Virtual Devices section for a complete description.

Disk File Backups

Any backup may be written to a disk file on the backup server or local system. This includes portable devices such as USB disks and RAID arrays such as SAN-attached disks. With disk backups, each filesystem or logical volume within the backup is stored in a different file, so access to the data is much faster than from tape, where it is usually necessary to rewind and forward a tape to a particular backup and filesystem to restore select data.

When a backup server is configured, one or more directories may be designated for storing backups. A different directory may be specified for storing System Backups than other backup types. System Backups are usually used for performing complete system installations of either the local system, or of other systems on the network (*Network Administrator* only).

System Backups may also be written to a directory on disk in a way that allows the disk to serve as system installation media. In other words, you can boot from, then reinstall the system from, locally-attached (or even SAN-attached) disks, allowing for a complete system backup and recovery using only a local disk. More information can be found in Configuring Spare Disks for System Backup/Recovery.

Each backup job will be assigned a unique Backup ID, and each client backup within the job will have a unique backup sequence number. Unlike tape backups, each disk backup will have a unique name, so there is no danger of overwriting a prior backup. Instead, the user must expire disk backups manually to prevent excessive use of disk space. When doing so, not only the backup labels, but also the actual backup files are removed from the disk. The administrator may also set the overwrite policy so that a disk backup over a certain number of days old is automatically expired and removed when the same backup job is re-run. This will prevent the filling up filesystems containing disk backups while still ensuring that the latest backups are kept on file.

Backup status, output, and label information may be displayed or disk backups just as with tape backups.

Understanding Backup Types

There are many types of backups that may be performed using SBAdmin. The backup type is configured into the backup profile, which is why you must have at least one profile setup for each type of backup you want to perform. The backup types are as follows:

1. System Backup - This backup contains the operating system and optionally all user data. User data may be only files in mounted filesystems, or may also contain raw data found in logical volumes, or partitions or meta-disks (***Linux***). It is possible to reinstall the entire system from a System Backup, or even use the backup of one client to install another. Select files, directories, logical volumes and volumes groups, and even raw data may be restored from a System Backup. For information the system installation process, refer to the ***SBAdmin System Recovery Guide***.

   AIX: The system backup contains the *rootvg* volume group, and may optionally contain some or all of the other volume groups on the system. If the backup is performed to tape, then this tape is also configured to boot to the **System Installation process**. The System backup may also be configured in the backup profile to be a "**Power System Backup**", meaning that all data is backed up and restored as raw logical volume data. Power System backups are much faster than when backing up individual files, but it is only possible to restore an entire filesystem or logical volume, not individual files, from a Power System Backup.

2. Volume Group Backup - This backup is typically used to separately backup the LVM volume groups that are not part of the operating system. Files or logical volumes within the volume group backup may also be backed up *incrementally*, including only files or logical volumes that have changed from a prior backup. The backup may contain one or more volume groups, and an entire volume group may be recreated and/or restored from the backup. Individual files, directories, filesystems or raw logical volumes may be recreated and/or restored.

   Volume groups may also be backed up *incrementally*, including only files or raw logical volumes that have changed from a prior backup.

   Volume Group backups are only available for ***AIX*** and for ***Linux*** systems with LVM installed.

3. Filesystem Backup - This backup will contain one or more filesystems on the system. The filesystems may be built on any logical volume, partition or meta-disk. Files within the filesystems may also be backed up *incrementally*, including only files that have changed from a prior backup. From this backup, select files, directories or filesystems may be recreated and/or restored.

4. Logical Volume Backup - This backup may include one or more "raw" logical volumes. From this backup, only an entire logical volume may be recreated and/or restored. Logical volume backups are only available for ***AIX*** and for ***Linux*** systems with LVM installed.

5. Directory Backup - This is the only backup type common to most other backup applications. It includes any number of directories and files, and select files and directories may be restored.

6. Partition Backup (***Linux*** systems only) - This backup may include one or more "raw" partitions typically containing non-filesystem data. From this backup, only an entire partition may be restored.

7. Meta-disk Backup (***Linux*** systems only) – Meta-disks are often referred to as ***Software RAID*** devices and either ***MD*** or ***multi-disk*** devices. This option is only available of Software RAID support is installed. This option will provide the ability to backup specific meta-disks, regardless of the type of device the meta-disk is built on. Meta-disks may be created on disks, partitions, logical volumes, and even other meta-disks.

8. SMB (Windows) Shares – The ***SMB File Sharing Protocol*** is commonly used for remote sharing of data found on Windows or Mac OS X systems with other operating systems. SBAdmin supports backing up of these "shared folders" using a ***Linux*** "*SMB host*" running the **SAMBA** *client* software. This also requires the "**smbfs**" filesystem support is available on the SMB host (typically available by

default). From an SMB data backup, specific files or directories may be restored, either to the same or different share, the same or different ***SMB client***, or even to a directory on a ***Linux*** or ***AIX*** client. Note that SMB Share Backups require the optional *SMB (Windows) Data Backup Feature* is installed.

It is possible to later restore specific data contained within the backup. It is not necessary to restore the entire backup. A System Backup, for instance, may contain multiple *volume groups*, each of which may contain *raw logical volumes* and *filesystems*, each of which may contain various *directories*, which each contain multiple *files*. It is therefore possible to restore one or more files, directories, logical volumes, filesystems, volume groups, meta-disks (***Linux***), partitions (***Linux***), or the entire system from a System Backup!

3. The Backup Administrator User Interface

The Backup Administrator User Interface is used for all configuration options, including servers, clients, virtual devices, jobs, profiles, etc. It is also used for the monitoring of job queues, displaying job status, backup output messages, and backup history.

After all backup jobs are configured and scheduled, they will continue to run even if the Backup Administrator interface is *not* running. Backup jobs may also be manually started, monitored or controlled from the command line when the interface is not running, and can be monitored or controlled after the administrator is restarted.

Ordinarily, messages regarding the status of the backup jobs are reported on the screen. If, however, the Backup Administrator interface is not running when a job is run, the status messages will be reported using an alternate notification method, which may be defined by the user.

The Main Screen

The following is a sample of the **Main Screen**, which appears when the application is first started. The options at the top of the screen (File, Configure, etc) are called contained in the menu bar. Click on any of the menu bar options to display a pull-down menu of options in each category. When selecting an option from the menu bar, a new screen, or window, will appear with additional optional options that apply to the menu selection.

The remainder of the screen will vary depending on the "Display" option chosen from the menu bar:

- **The Clients, Servers & Devices** display (shown below) is available only if the *Network Administrator* license is installed. In this example, several clients, servers and devices have already been configured. The application continually checks the availability of the systems, and displays an icon that represents both the client system type (**A**=***AIX***, **L**=***Linux, S=SMB [Windows]***) and whether or not the system is available (**Green**=available, **Red**=not available). Tape devices, virtual devices, and backup directories are shown. Devices will appear red if the device is un-configured or unavailable.

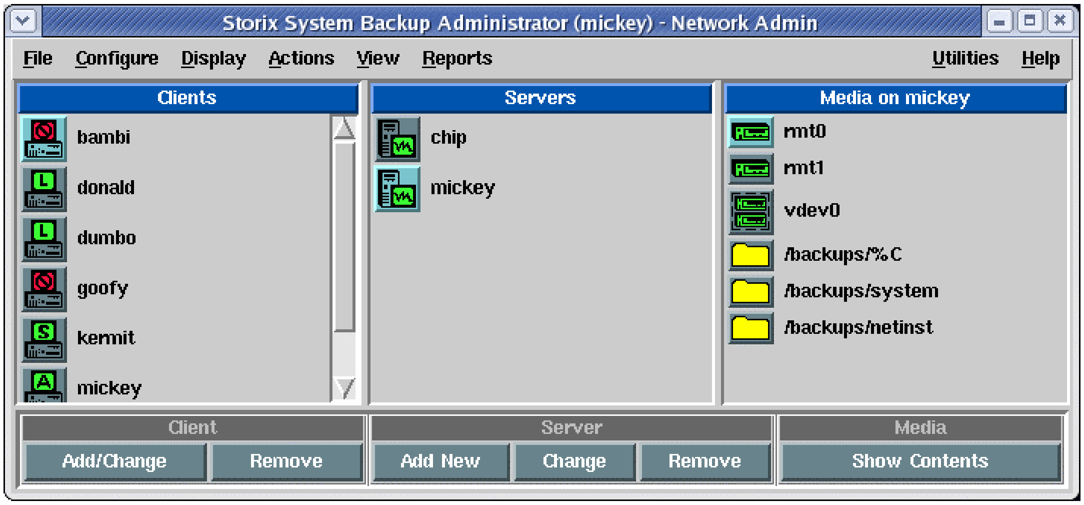

A client may be selected by clicking the left mouse button on the icon next to the client hostname. Likewise, a server may be selected by clicking the mouse button on the server icon. When you click on a server, a list of backup devices and directories configured for that server will appear. The *selected* client, server, and device or directory will appear with a blue background.

The action buttons at the bottom of the screen apply to the selected client, server, or device or directory. They provide a shortcut to performing the same tasks that can be performed from various options within the menu bar.

- **The Job Information Display** provides a quick glance at the jobs that are configured. The left-most display area shows the job icons. The green calendar represents a job is that is scheduled. The red calendar represents jobs that are not scheduled to run. By clicking on a job icon, the job information and schedule information, if any, is displayed in the right two display areas, and the icon background is changed to blue.

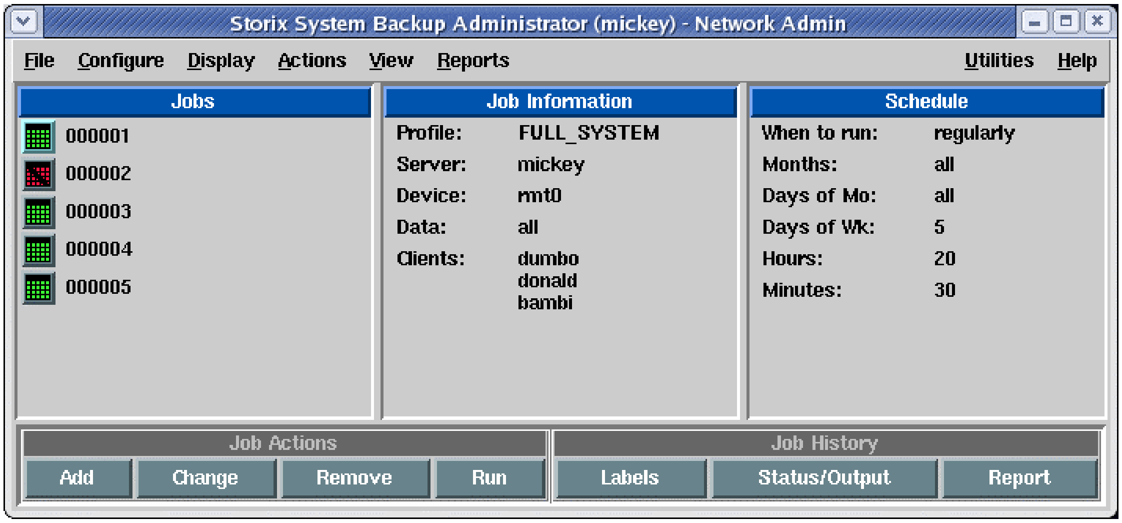

Jobs

🟩 000001
🟥 000002
🟩 000003
🟩 000004
🟩 000005

Job Information

Profile:	FULL_SYSTEM
Server:	mickey
Device:	rmt0
Data:	all
Clients:	dumbo donald bambi

Schedule

When to run:	regularly
Months:	all
Days of Mo:	all
Days of Wk:	5
Hours:	20
Minutes:	30

The action buttons at the bottom apply to the selected job. They are shortcuts for various job-related functions. The ***Job Actions*** buttons perform the same job operations available from the Action menu on the menu bar. The Run Now button will place the selected job in the queue (even if it is scheduled to run at another time), and it will be run as soon as the server and device assigned to the job are available. Each of these functions is described in detail in the section Schedule or Run Backup Jobs. The ***Job History*** buttons may be used to view backup labels, status/output messages or a history report for previously run jobs.

- **The Job Queue Display** provides a look at the jobs that are currently in the queues. A queue is shown in the left-most display area, which consists of the backup server and the device name. When you click on a queue, the selected queue is highlighted in blue, and the jobs in the selected queue are displayed in the middle display area.

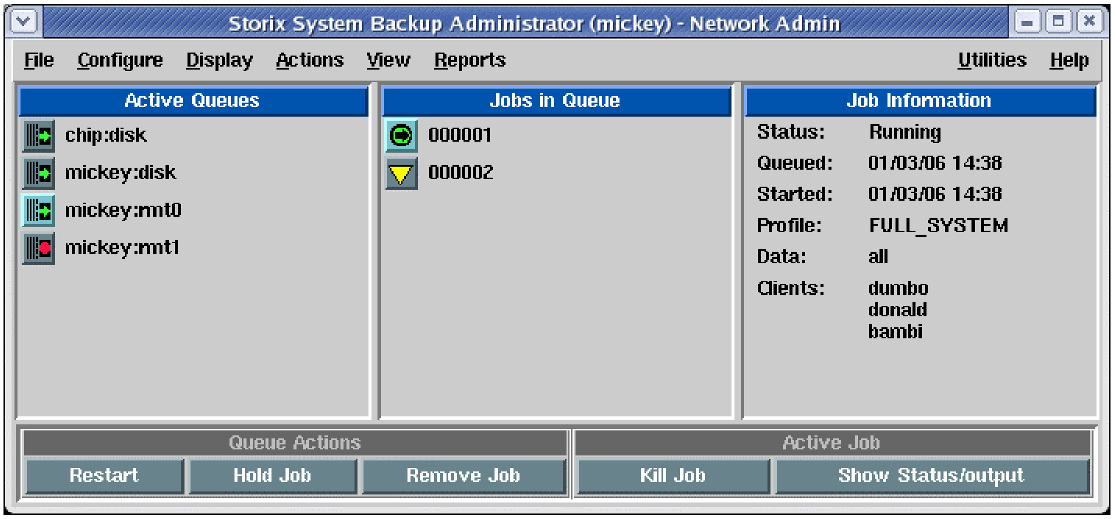

Active Queues

chip:disk
mickey:disk
mickey:rmt0
mickey:rmt1

Jobs in Queue

● 000001
▼ 000002

Job Information

Status:	Running
Queued:	01/03/06 14:38
Started:	01/03/06 14:38
Profile:	FULL_SYSTEM
Data:	all
Clients:	dumbo donald bambi

You may then click on a particular job to display the job information, including the status of the job. Both the queue and job icons represent the status of the job. The ***Queue Actions*** buttons at the bottom of the screen may be used to manipulate the selected job. The ***Active Job*** buttons include the ability to *kill* a running job or display the status or output messages of a running or failed job. All of these functions and a list of any possible icons or status messages are described in detail in the Job Queues section.

Closing and Iconifying Windows

Two common icons appear on most screens, or windows. They are:

cancel button

iconify button

After making changes to information on any screen, use the cancel button to cancel the changes and close the window. If you would like to temporarily remove the window from the screen without saving or canceling the changes, use the iconify button. The window will then appear as an icon on the display with the SBA logo. To restore the iconified window to its normal size, double click on it.

These icons do not appear on the Main Screen. From the Main Screen, you should always use the File→Exit option on the menu bar to exit the application, and you may use the icons in the title bar for other window manager functions, such as iconifying the window.

Buttons, Messages, Lists and Entry Fields

A button is a rectangular object that usually has a raised surface. A button is pressed by moving the pointer over the button and pressing the left mouse button. If a button is active, you will see the background color of the button change when the mouse is moved over it. If it is inactive, the background will not change and the image or text on the button will usually appear grayed out. If a response is required by the user, a messagebox will appear with a button available for each valid response. If a default response is available (may be selected by pressing **Enter** on the keyboard rather than clicking the button), this option's button will appear with a groove around it. The following is an example of a popup message with option buttons:

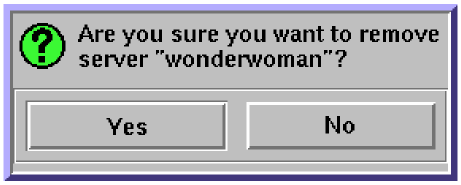

Are you sure you want to remove server "wonderwoman"?

Radiobuttons are special buttons that are grouped so that only one button may be selected at a time for a group of options. By clicking on a radio button, any previous button in the group will be deselected and the new button will be selected (indicated in red). The following is an example of a group of radio buttons:

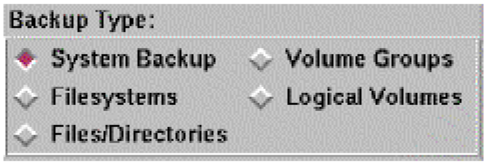

Backup Type:

◆ System Backup ◇ Volume Groups
◇ Filesystems ◇ Logical Volumes
◇ Files/Directories

Listboxes display a list of options, and the user may select either one or more options, depending on the type of selection. If a single option is requested, clicking on an option will highlight that option, and de-select any prior option made. If more than one option is allowed, clicking on an option will highlight (select) it, and clicking on an already selected item will de-select it. The following is an example of a listbox with a list of clients. In this case, more than one option is selected:

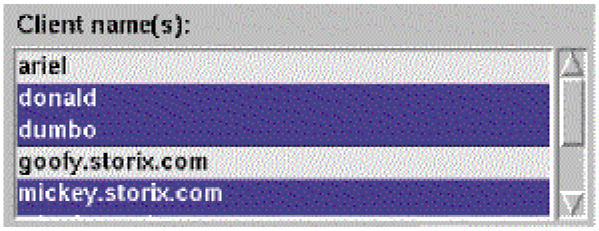

Client name(s):

ariel
donald
dumbo
goofy.storix.com
mickey.storix.com

Note that the above list contains a vertical scrollbar to the right. This scrollbar will appear only if the list exceeds the height of the box. To display options beyond the bottom of the box, click on and hold the scrollbar as you move the cursor down, *dragging* the scrollbar downward.

An entry field allows the user to type data, rather than select an option. Entry fields may be either enabled (user may enter data), or disabled (user may not type in the field). If an entry field is disabled, it will appear with a darker background. An entry field is either disabled because the option does not apply to the current operation, or because limited selections are available. If limited selections are available, a pop-up list of options will be provided. Pop-up lists are displayed by selecting the arrow button to the right of the entry field. If this arrow is available, it may be used to display a list of available options and select from the list (like a listbox). After selecting an option from a popup list, the selection will appear in the entry field. The following are examples of entry fields. The first is disabled and provides a popup list, the second allows the user to enter data:

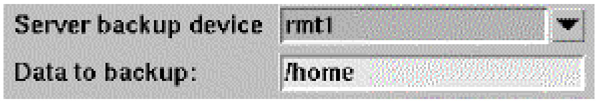

Server backup device rmt1 ▼

Data to backup: /home

4. Configuring Users

To install and initially configure your software license for SBAdmin, you must be logged onto the system as the "***root***" user. You may continue to perform all other backup configuration and administration tasks using SBAdmin while logged in as root.

However, it is often desirable to allow other users on the system without root access to perform the backup configuration and administration tasks, without giving them root access to other applications and system functions.

NOTE **It is not necessary to configure other users if the person performing the backup tasks is logged on as *root*.**

By using this option, you may indicate which users, while logged into the system, may perform SBAdmin functions. These functions will include:

- Starting the ***administrator*** (**sbadmin** program)
- Performing all tasks within the administrator (configuration, backup, restore, etc)
- Execution of SBAdmin commands at the command line (as described in the Commands section)

Note that the user may perform only the above operations, and there is no way for the user to execute commands outside of SBAdmin, either local or remote, except for those commands indicated in the Commands section.

Adding a User

To add a user, select Configure→Users from the menu bar. A screen such as the following example will appear:

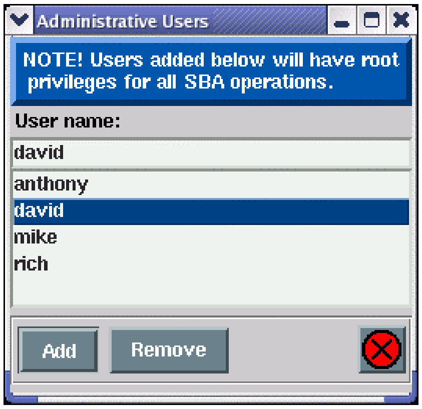
Administrative Users

NOTE! Users added below will have root privileges for all SBA operations.

User name:

david

anthony
david
mike
rich

From this screen, simply type the name of the user to add in the entry field and press the Add button. When finished, press the **Cancel** button at the bottom.

Removing a User

Note that removing a user from the list will not prevent the user from logging into the system. Click Configure→Users from the menu bar. Select the name of the user to remove from the list and press the Remove button.

When finished, press the **Cancel** button at the bottom.

5. Configuring Clients

This section only applies to the *Network Administrator* license option.

A client is defined as any system that will be backed up using the Backup Administrator. If backups are to be performed of the backup administrator itself, or any backup servers, then they should also be configured as clients. There are two client types, **AIX/Linux clients** and **SMB (Windows) clients**. Only if the ***Windows (SMB) Data Backup Feature*** is installed will the SMB client options be presented.

Any number of clients may be added to the administrator as long as the total number of *unique* clients and servers does not exceed the number of clients licensed to the network administrator. Note that the administrator itself also includes a client license, so it may be configured as a client or server without using one of your additional client/server licenses.

Adding a Client

Any client may be added to the administrator by simply adding its hostname. However, the number of clients which may be added is dependent on the number of clients the administrator is licensed for. Also, any client hostname may be added, but the client is only accessible to the administrator after the software has been installed and configured onto the client system as well.

If the ***Windows (SMB) Data Backup Feature*** is not installed on the admin system, then only one option is available for configuring clients. If the optional features is installed, then you will have a separate option for configuring ***Linux/AIX*** or ***Windows (SMB) clients***.

To add a client, select one of the following form the menu bar::

- Configure→Clients (if no Windows client option presented) to configure Linux or AIX clients
- Configure→Linux/AIX Clients,
- Configure→Windows (SMB) Clients.
- Click the Add Client button at the bottom of the Main Screen when the Clients, Servers & Devices are displayed to add or change a Linux or AIX client.

Configuring a Linux or AIX Client

After selecting the appropriate option above, the following window will be displayed:

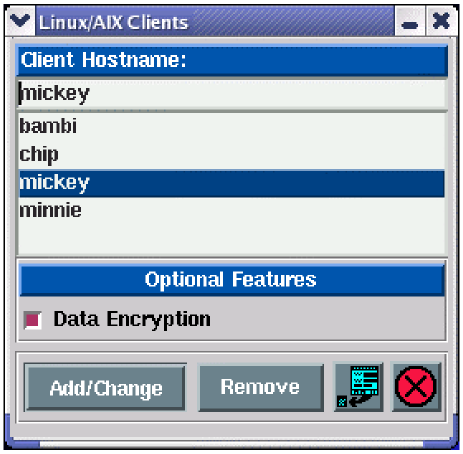

Linux/AIX Clients

Client Hostname:

mickey

bambi
chip
mickey
minnie

Optional Features

■ Data Encryption

This example shows several clients already configured. To add a new client, enter the hostname of the client in the entry field at the top. Note that the hostname you enter may be a simple hostname (i.e. *ariel*) or a full domain name (*goofy.storix.com*) and must be known to the admin system.

Configuring a Windows (SMB) Client

Upon selecting the corresponding option from the menu bar, a screen similar to the following will appear:

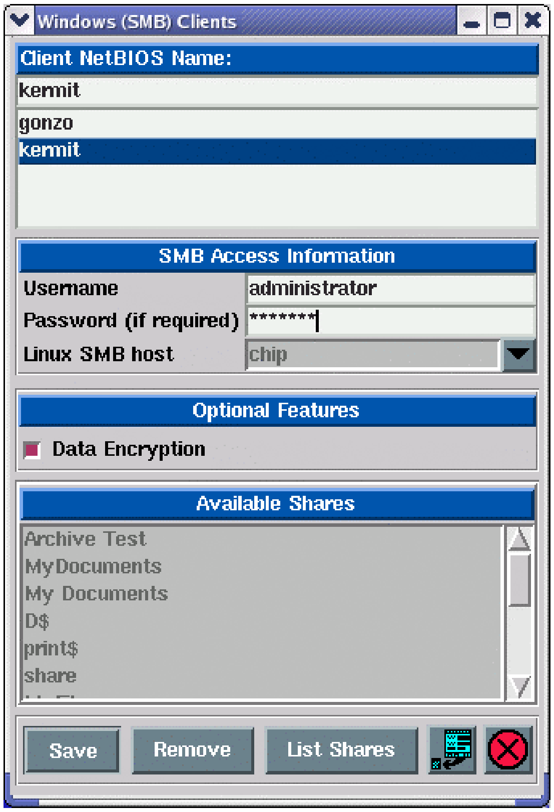

Windows (SMB) Clients

Client NetBIOS Name:

kermit

gonzo
kermit

SMB Access Information

Username: administrator
Password (if required): *******
Linux SMB host: chip

Optional Features

■ Data Encryption

Available Shares

Archive Test
MyDocuments
My Documents
D$
print$
share

Configuring an ***SMB client*** is similar to Linux or AIX clients, except that no software is actually installed on the client itself. Instead, SBAdmin will backup the data from the SMB client by accessing a shared resource, typically called a “shared folder” on ***Windows*** or ***Mac OS X*** systems. Note that this feature may be used on any system that uses the *SMB File Sharing Protocol*.

Access to the SMB client's shared data must be obtained from an "*SMB host*". The SMB host is a Linux system that has **SAMBA** (SMB client software) installed. The SMB host must also have "smbfs" filesystem support available, which is usually available by default on most Linux distributions.

To configure the SMB client within SBAdmin, you must assign it an SMB host, which may be any other configured Linux client. The SMB client's data will be backed up through the SMB host.

Access to an SMB share, depending on the configuration of the SMB system itself, usually requires both a username and a password. Therefore, this information must be entered as well, and will be used to gain access to the share from the SMB host.

The SMB usernames and passwords stored are only on the administrator in encrypted form, and are not viewable by any other user, nor passed in un-encrypted form over the network to the SMB host.

It is important to note, however, that the username and password are sent from the SMB host to the SMB client by the "smbclient" (SAMBA) command. Depending on your version of smbclient, this transfer of the username and password may be unencrypted.

To select the **SMB host**, use the drop-down button to the right of the entry field. Select any ***Linux*** system with **SAMBA** client support and **smbfs** filesystem support installed. You may alternatively select "**Use backup server**". This will indicate that the backup server configured for the backup job will also act as the SMB host. This is useful if you use different backup servers, even for the same client, and want to reduce network traffic by having the SMB data accessed directly by the system where the backup data is stored.

To check that the SMB host can access the shared folders of the SMB client, after entering the correct **username** and **password**, press the List Shares button at the bottom of the screen. If successful, a list of available shared folders that can be included in backup jobs will be displayed in the Available Shares box.

Enabling Backup Data Encryption for a Client

The **Data Encryption field** will be enabled only if the ***Backup Data Encryption Feature*** is installed. If so, you may select this button to indicate that data may be encrypted when backing up this client. Any type of data, for any client type, may be encrypted using 128, 192, or 256-bit AES encryption. Encryption is configured for specific clients according to the number of clients your encryption license, if any, supports. You may only select this button for the number of clients your encryption license supports.

Enabling data encryption for a client does not cause all backups to be encrypted automatically. It only designates which clients will support encryption. For clients that support encryption, the encryption option becomes available when configuring backup jobs.

To encrypt data for a client, each client must have at least one configured Encryption Key. The encryption key must be a 32, 48 or 64-byte hexadecimal number, depending on the number of bits of encryption used. An encryption key will be given a user-defined Encryption Key ID, and you may have as many Key IDs as you like. You will later select which Key ID to use when performing a particular backup.

To prevent encryption keys from ever being transmitted across the network, the encryption keys may not be configured from within the GUI interface, and client keys may not be configured from the network admin system. Instead, you must run the **stkeys** command on each client for which encryption is to be used. Refer to stkeys in the ***Commands*** section, and the Encrypt data field in the backup job configuration for additional information.

Press the Save button to add or change the client settings. After adding a client, its icon will immediately appear on the Main Screen when Clients, Servers & Devices are displayed. If the software has not been configured on the client, or if the client was not configured using the correct hostname of the admin system, the client icon will appear in red. If the software is installed and configured properly on the client, the icon will appear green to indicate that the client is accessible to the admin system.

Removing a Client

A client may be removed from the system only if it is not assigned to any backup jobs. If it is, you will be informed so, and you must remove or change the job to remove the client from the list of clients to backup.

To remove a client, either:

- Select a client on the Main Screen when Clients, Servers & Devices are displayed, then click the Remove Client button at the bottom of the screen, or

- Click Configure→Clients from the menu bar. Select the name of the client to remove from the list and press the Remove button.

The client icon will be removed from the Main Screen when Clients, Servers & Devices are displayed.

6. Configuring Servers, Backup Devices & Directories

A backup media server, also referred to simply as the ***backup server*** or just ***server***, is defined any system to which backups will be sent. The backups may be stored onto tape drives attached to the server or saved in directories on the disks of the backup server. Any system may be a backup server, including any client or the admin system. A backup server is usually also defined as a client since it too must be backed up periodically.

> NOTE **The following section is used only when the *Network Administrator* license is installed. Refer to the *Configuring Backup Devices and Directories on a Standalone System* section below for *Workstation, Desktop* and *Personal Edition* licenses.**

Adding a Server

A new server may be added to the system by either:

1. Selecting Configure→Servers from the menu bar, or

2. Pressing the Add Server button at the bottom of the Main Screen when Clients, Servers & Devices are displayed.

After doing so, the **server selection screen** such as the following will be displayed:

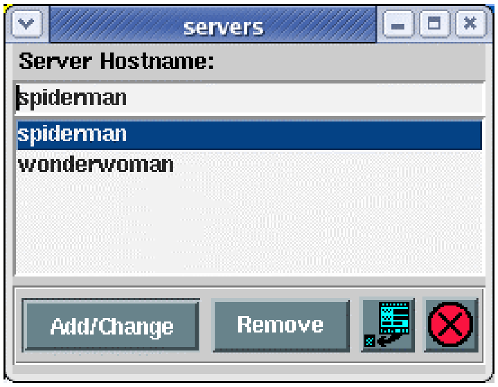

To add a new server, enter a new server name in the entry field at the top, then click the Add/Change button. Additional options will be provided on the **server options screen** as shown in the following example:

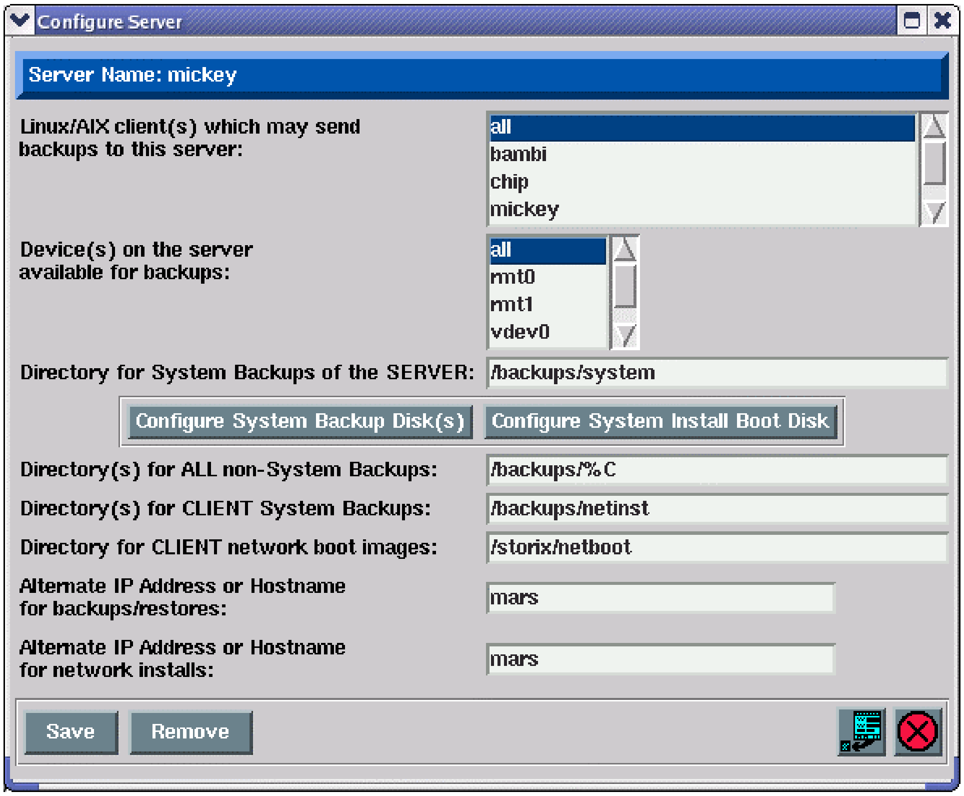
Configure Server
Configure System Backup Disk(s)
Configure System Install Boot Disk
Save
Remove

Server Name: mickey

Linux/AIX client(s) which may send backups to this server: all (selected), bambi, chip, mickey

Device(s) on the server available for backups: all (selected), rmt0, rmt1, vdev0

Directory for System Backups of the SERVER: /backups/system

Directory(s) for ALL non-System Backups: /backups/%C

Directory(s) for CLIENT System Backups: /backups/netinst

Directory for CLIENT network boot images: /storix/netboot

Alternate IP Address or Hostname for backups/restores: mars

Alternate IP Address or Hostname for network installs: mars

Assigning Clients to a Backup Server

In the first listbox, you must select one or more client that will be assigned to this server. Only clients selected in this box will have permission to backup to the server, and only those clients listed will appear as client options when configuring backup jobs.

You may alternatively select "***all***". If you do so, all specific client selections will be de-selected. If all clients are to be permitted to backup to the server, all configured clients will be displayed when a client options list is provided.

Making Backup Devices Available

The second listbox shows all of the tape devices and virtual devices available on the server. You need to select at least one option. The selected devices will be made available as backup device options in other functions.

You may alternatively select "***all***". If you do so, all specific device selections will be de-selected. If all devices on the server are to be made available as backup devices, the server will be queried each time a backup device list is needed in other functions and all devices on the server will be shown. Note that if you select "***all***", any devices added to the server later will automatically appear as a backup option without having to change the server information.

Specifying Backup Directories

You may specify a number of different directories for storing backups for both the local system and other clients. You may enter one or more directories in each field except for local System Backups, which will accept only one entry. You may also enter the same directory in each field if you do not want the backups to be placed in different directories:

- **Directory for System Backups of the SERVER** – Enter here a directory where you would like full system backups of the *local server* system to be stored. Since these backups usually take a lot of space (up to the size of all data on the system), it is a good idea that you create a large separate filesystem for storing these backups.

  If you want to be able to re-install the entire system from a system backup stored on a local disk, then you must configure a spare hard disk (or portable or SAN-attached disk) for system backups. To do so, select the button **Configure System Backup Disk(s)**. This option is described in detail in the section Configuring Spare Disks for System Backup/Recovery below.

- **Directory(s) for all non-System Backups** – Here you can specify where you would like all non-System Backups to be written, either for the local system or other clients. As with the directory for client System Backups, you can use the %C notation to cause the backups to be written to a separate sub-directory for each client. You may enter any number of directories, separated by spaces.

- **Directory(s) for CLIENT System Backups** – A separate directory is usually designated for client system backups since a client may be installed from this server over the network. When doing so, a list of backups available in the directory list provided here will be presented to the client from the System Installation Menus as available installation media.

  If you include the characters "**%C**" in a directory name, those characters will be replaced with the ***hostname*** of the client when backups are performed. For instance, if you accept the default directory of **/backups/%C**, a backup job containing clients **ariel** and **dumbo** will be placed in directories **/backups/ariel** and **/backups/dumbo** respectively. This will make it easier to differentiate one client system backup from another when reinstalling clients over the network. The same applies to *SMB client* backups, except that the directory name will also include the SMB host. In this case, an SMB client **winclient1** using SMB host "**mickey**" will be stored in the **directory /backups/mickey/winclient1**.

- **Directory for CLIENT network boot images** – This is the directory where network boot images will be stored for booting AIX or Linux client systems over the network for system recovery. Common network boot images may be created for similar systems, or a separate boot image may be created for each client, all of which will be stored in this directory.

NoTE **The network boot image directory must be accessible by the *tftp* server, required for serving network boot images. You must make note of this directory also when enabling the clients for network boot.**

Setting Up Alternate Network Adapters

Two entry fields are used for entry of optional *IP Addresses* or *Hostnames* pointing to alternate network adapters on the server. The first entry field is used to set the network adapter to use for backups and restores, and the second is used to set the network adapter for network boots and network installs (see ***Network Boot/Install Configuration*** in the ***SBAmin AIX System Recovery Guide***).

If no entry is made in these fields, the same network adapter used to communicate with the admin system will be used to communicate with the clients. By selecting alternate network adapter, that network adapter may be used for communication between this server and any clients it is backing up or restoring data to, or when re-installing the client over the network. This is particularly useful if a different network is available for communication between the clients and the backup server that is not available to the Admin System.

NoTE **Although alternate network adapters may be set in the server configuration, they will NOT be used by default. For the alternate adapter to be used, you must select the *Use Alternate IP/Hostname* option when configuring backup jobs or *Network Boot/Install Configuration* in the *SBADmin AIX System Recovery Guide*).**

In the ***Alternate IP or Hostname for Backups/Restores*** field, you may enter any hostname or IP address for the server that is known to the client. This may be any network type that can communicate using **TCP/IP** and includes the SP High Speed Switch networks.

In the ***Alternate IP or Hostname for Network Boot/Installs*** field, you may enter any hostname or IP address for the server with which it may boot the client. This includes only adapter types supported by the ***AIX*** network boot process (ethernet, token-ring and FDDI networks only).

> NOTE **The High Performance Switch (HPS) networks on IBM SP systems are not supported by *AIX* for network boots and installs.**

When all entries are complete, press the Save button to save the server information and close the window.

Changing a Server

The information for an existing server may be changed by either:

1. Selecting Configure→Servers from the menu bar, or
2. Selecting a server icon from the Main Screen when the Clients, Servers and Devices are displayed, then pressing the Change Server button at the bottom of the screen.

If selected from the menu bar, the server selection screen will appear. If a server was selected from the Main Screen, the server options screen will appear with the prior settings for the server. Simply add or change any of the information on the screen, then press the Save button at the bottom to save the changes.

Removing a Server

A server may be removed from the system only if there are no jobs currently assigned to it. If there are jobs assigned, you will be informed so, and you must remove or change the job to use a different server before the server may be removed.

To remove a server, either:

Select Configure→Servers from the menu bar. Select the server to remove from the listbox (see the server selection screen), then select the Remove button, or

Select a server icon from the Main Screen when the Clients, Servers and Devices are displayed, then press the Remove Server button at the bottom of the screen.

Configuring Backup Devices and Directories on a Standalone System

As shown above, you specify the possible backup devices and directories for each server when using the *Network Administrator*. For a standalone system using the *Workstation* or *Desktop* license, you may use the option **Configure→Backup Devices/Directories** to define these options.

The backups may be stored onto tape drives attached to the system or saved in one or more directories on disk. Default options are automatically setup when you install the software. These allow for the configuration of all tape drives for use as backup devices and the “/backups” directory for storing disk backup images.

To change the backup devices or directories that will appear as backup device options in other part of this application, select Configure→Backup Devices/Directories from the menu bar

After doing so, a screen such as the following will be displayed:

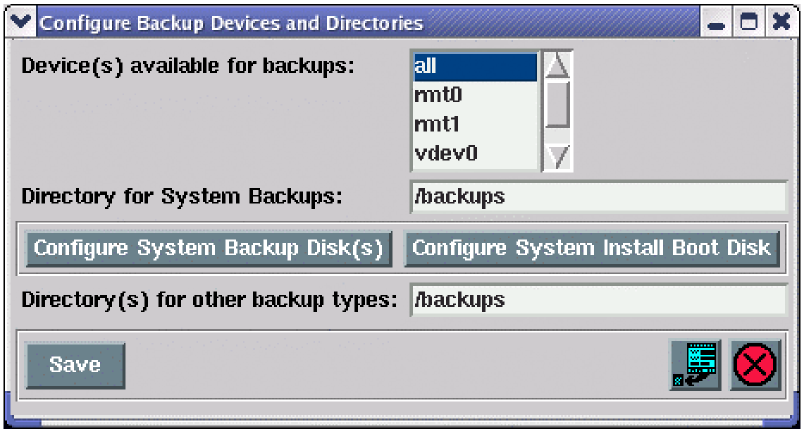

Configure Backup Devices and Directories

Device(s) available for backups: all, rmt0, rmt1, vdev0 (all selected)

Directory for System Backups: /backups

Configure System Backup Disk(s)
Configure System Install Boot Disk

Directory(s) for other backup types: /backups

Making Backup Devices Available

The listbox shows all of the tape devices and virtual devices (*Workstation Edition* only) available on the system. You need to select at least one option. The selected devices will be made available as backup device options in other functions.

You may alternatively select "***all***". If you do so, all specific device selections will be de-selected. If all devices on the system are to be made available as backup devices, the system will be queried each time a backup device list is needed in other functions and all devices on the system will be shown. Note that if you select "***all***", any devices added to the system later will automatically appear as a backup option without having to change this option.

Specifying Backup Directories

You may specify a number of different directories for storing backups for both the local system and other clients. You may enter one or more directories in each field except for local System Backups, which will accept only one entry. You may also enter the same directory in each field if you do not want the backups to be placed in different directories:

- **Directory for System Backups** – Enter here a directory where you would like full system backups of the system to be stored. Since these backups usually take a lot of space (up to the size of all data on the system), it is a good idea that you create a large separate filesystem for storing these backups.

  If you want to be able to re-install the entire system from a system backup stored on a local disk, then you must configure a spare hard disk (or portable or SAN-attached disk) for system backups. To do so, select the button **Configure System Backup Disk(s)**. This option is described in detail in the section Configuring Spare Disks for System Backup/Recovery below.

- **Directory(s) for other backup types** – Here you can specify where you would like all non-System Backups to be written.

  When your entries are completed, press the Save button to save the entries and close this window.

Configuring Spare Disks for System Backup/Recovery

This option is available for all license types and allows you to use a spare hard disk, portable disk, or even SAN-attached disks for full system recovery media. When using a *Network Administrator* license, this option is configurable for each server under the **Configure→Servers** option. For other standalone system licenses, this is found under **Configure→Backup Devices/Directories**.

If using a *Network Administrator*, you must configure the backup directory for each server since the directory will serve as backup media. However, only clients are backed up. Therefore, if you want to backup a client to its own local directory, you must configure the system as both a client and server.

When the screen is displayed where you would specify the devices and directories for backups, two additional buttons appear, which are used for configuring spare disks for system backup and recovery:

Configure System Backup Disk(s)

This option will configure one or more disks to contain system backups which will be recognized as system installation devices by the ***System Installation*** process. All contents of the disks, if any, will be overwritten by this process, and the disks will be configured with partitions (***Linux***), and LVM logical volumes. Using logical volumes, a single filesystem will be created, which will span all of the disks selected.

When selecting this button, a screen similar to the following is shown:

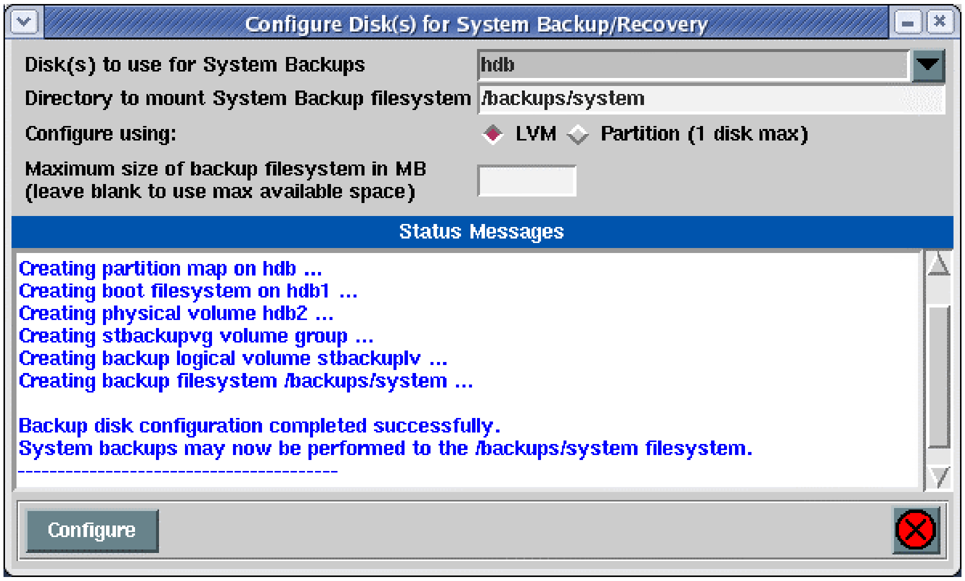

Configure Disk(s) for System Backup/Recovery

Disk(s) to use for System Backups: hdb

Directory to mount System Backup filesystem: /backups/system

Configure using: ● LVM ○ Partition (1 disk max)

Maximum size of backup filesystem in MB (leave blank to use max available space):

Status Messages

```
Creating partition map on hdb ...
Creating boot filesystem on hdb1 ...
Creating physical volume hdb2 ...
Creating stbackupvg volume group ...
Creating backup logical volume stbackuplv ...
Creating backup filesystem /backups/system ...

Backup disk configuration completed successfully.
System backups may now be performed to the /backups/system filesystem.
--------------------------------------------
```

In the Disk(s) to use for System Backups field, select the arrow to the right to display and select one or more disks to use. Note that it may be necessary to select more than one disk if the System Backup data is too large to fit on a single disk.

If no disks appear in the list, then there are no spare disks available on the system. Those disks which appear in the list are those believed to be unused by other data. Select one or more disks from the list.

Be absolutely sure the disk you select does not contain any needed data! Using this option will overwrite the entire contents of the disk!

Disks containing partitions (*Linux*) are considered available for System Backups if the partitions are not mounted filesystems. Disks configured as LVM Physical Volumes are considered available if they are either not assigned to a Volume Group or the Volume Group was exported.

Enter a directory where the backup filesystem should be mounted. If you entered a Directory for (local)

System Backups in the previous screen, that entry will be displayed here. If you change the directory in this field, it will be changed in the previous screen as well. After entering a backup directory, that directory will appear in the list of devices/directories to backup to (see **Configure Jobs**) when performing a **System Backup**.

On ***Linux*** systems, there are two methods in which the disk may be configured for system backup/recovery (***AIX*** systems are entirely *LVM*-based and therefore have only one option). In either case, a small amount of disk space at the beginning of each disk is reserved for making the disk bootable. The contents of the remainder of the disk depend on the option selected:

- **Configure using *LVM*** – If Logical Volume management is installed on Linux, then it may be used to configure one or more disks into a *Volume Group*. Within this volume group, a *Logical Volume* is created, containing the backup filesystem. This filesystem may be limited in size using the **Maximum Size of Backup Filesystem** field, or it can span the entire volume group, and therefore more than one physical disk.

- **Configure using *Partition*** – For systems that do not have LVM, a single disk may be configured for system backup/recovery by creating the backup filesystem in a disk partition. A single partition will be created for the entire disk. Since this partition may not span multiple disks, the disk used in this case must be large enough to write an entire system backup.

When you have finished your selections, press the Configure button at the bottom of the screen. The messages indicating the progress of this configuration process will be displayed in the Status Messages portion of the screen.

Configure System Install Boot Disk

Once one or more disks have been configured for system backups, you can also make one of the disks boot to the ***System Installation*** process. This will allow you to perform a full system recovery from a system backup written to the local disk without the need of other boot media. By selecting to boot from this disk in your system firmware, the System Installation process will appear, from which you can select to restore from a local System Backup on a locally-attached disk. You will also be able to select a local tape device or remote tape or disk backup (if using ***Network Administrator***) to restore from if you do not want to use a backup on this disk.

> NOTE
>
> **Although this process will allow you to configure virtually any disk to boot to the System Installation process, not all *system firmware* (built into your hardware) will recognize the disks as bootable. After successfully performing this option, you should always test the boot disk by selecting to boot from it within the firmware. Note that the System Installation menus will appear on the screen, but no information on the system will be changed without selecting the backup media and selecting to continue the system installation.**

Making a disk bootable is relatively the same as creating any other boot media within SBAdmin. Therefore, refer to the ***SBAdmin (AIX or Linux) System Recovery Guide*** for details on **Creating System Installation Boot Media for further details**.

7. Backup Profile

A backup profile is used to set default backup selections commonly used when performing different types of backups. Assigning a backup profile to a backup job alleviates the need to repeatedly answer the same questions every time a new job is added.

At least one profile must be created for each backup type to be performed. When the software is first installed, a set of pre-defined backup profiles, one for each backup type, is automatically installed. These profiles are not required and may be removed if desired.

After creating a single backup profile, any options selected for that profile, except the backup type, may be customized for each backup job it is assigned to. It is therefore only necessary to create a single backup profile for each backup type, but you may want to create different profiles for a single backup type to prevent having to change the options for different jobs.

Adding a Backup Profile

A new profile may be added by selecting Configure→Backup Profiles from the menu bar. When you do so, a **profile selection screen** similar to the following example will appear:

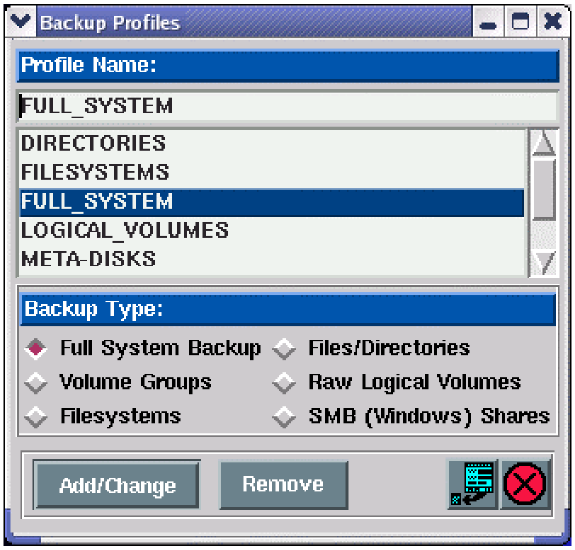

Backup Profiles

Profile Name:

FULL_SYSTEM

DIRECTORIES
FILESYSTEMS
FULL_SYSTEM
LOGICAL_VOLUMES
META-DISKS

Backup Type:

● Full System Backup ○ Files/Directories
○ Volume Groups ○ Raw Logical Volumes
○ Filesystems ○ SMB (Windows) Shares

When the software is initially installed, a set of pre-defined profiles are automatically configured (as shown above). There is one pre-defined profile for each type of backup. You may choose to edit any of these profiles, delete them, or add new ones of your own. To edit an existing profile, select the profile in the **Profile Name** listbox, then press the Add/Change button.

The Backup Type options may vary depending on the operating system support and additional features enabled. Refer to Operating Systems Support and Enabling Optional Features for more information.

To add a new profile, enter a new profile name in the entry field at the top of the screen, then select the type of backup for this profile by pressing one of the buttons in the section below the profile names. A profile name may consist of any characters except a colon (:) or space (spaces will be changed to underscores).

You may select only one backup type for each profile. If, for instance, you want to create a profile for performing volume group backups, check the "Volume Group" button. The options which follow will then only be applicable to that backup type.

To continue, click the Add/Change button. Additional options will be provided on the **profile options** screen.

The following is an example of a **System Backup** profile when both ***Linux*** and ***AIX*** client support is enabled:

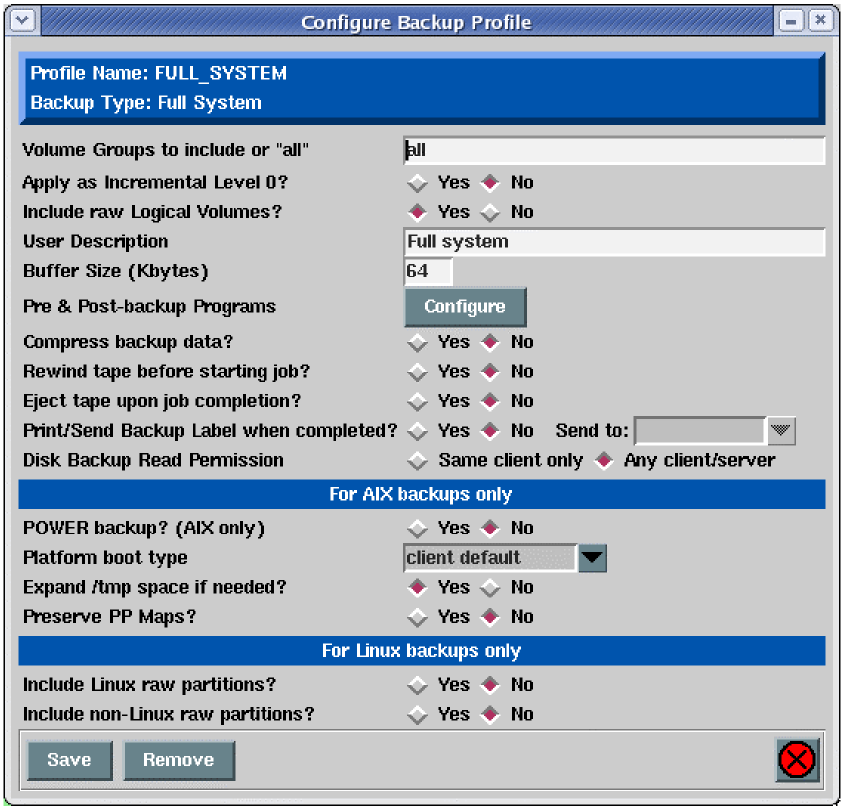

Configure Backup Profile

Profile Name: FULL_SYSTEM
Backup Type: Full System

Volume Groups to include or "all"	all
Apply as Incremental Level 0?	◇ Yes ◆ No
Include raw Logical Volumes?	◆ Yes ◇ No
User Description	Full system
Buffer Size (Kbytes)	64
Pre & Post-backup Programs	Configure
Compress backup data?	◇ Yes ◆ No
Rewind tape before starting job?	◇ Yes ◆ No
Eject tape upon job completion?	◇ Yes ◆ No
Print/Send Backup Label when completed?	◇ Yes ◆ No Send to:
Disk Backup Read Permission	◇ Same client only ◆ Any client/server
For AIX backups only	
POWER backup? (AIX only)	◇ Yes ◆ No
Platform boot type	client default
Expand /tmp space if needed?	◆ Yes ◇ No
Preserve PP Maps?	◇ Yes ◆ No
For Linux backups only	
Include Linux raw partitions?	◇ Yes ◆ No
Include non-Linux raw partitions?	◇ Yes ◆ No

The options that appear may differ depending on the type of backup selected, since not all options apply to all backup types. Also, the options which appear will differ depending on the operating system support that is enabled. Not all options on the screen may apply to all client system types.

The following is an example of an **SMB (Windows) Share** profile when ***Windows (SMB) Data Backup Feature*** is enabled:

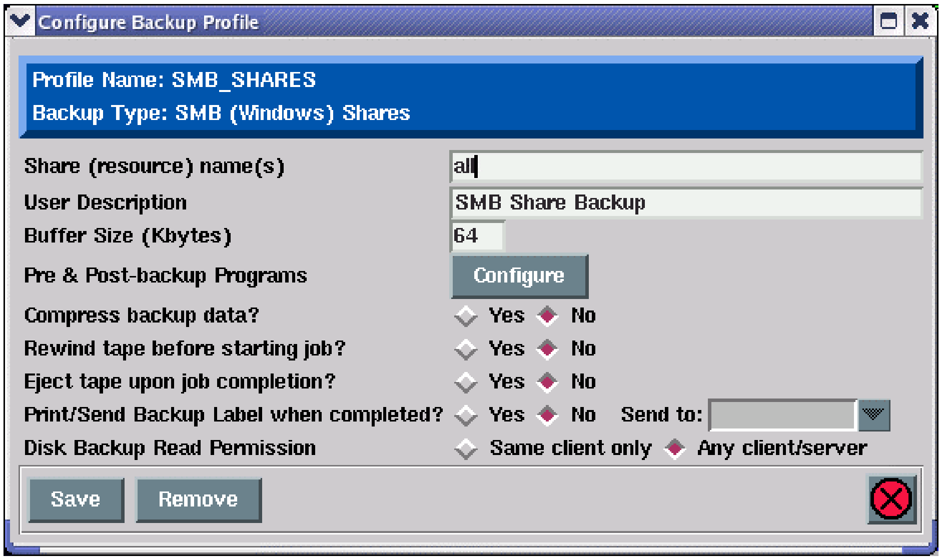

Configure Backup Profile

Profile Name: SMB_SHARES
Backup Type: SMB (Windows) Shares

Share (resource) name(s)	all
User Description	SMB Share Backup
Buffer Size (Kbytes)	64
Pre & Post-backup Programs	Configure
Compress backup data?	◇ Yes ◆ No
Rewind tape before starting job?	◇ Yes ◆ No
Eject tape upon job completion?	◇ Yes ◆ No
Print/Send Backup Label when completed?	◇ Yes ◆ No Send to:
Disk Backup Read Permission	◇ Same client only ◆ Any client/server

Use QuickHelp at any time to display a description or instructions for a particular option. Also note that a profile will be assigned to each backup job. Since all settings shown do not always apply to all backup jobs, any of the options you see here may also be customized for each backup job. Refer to Configure a Backup Job for more information.

Pay attention to the options for rewinding or ejecting the tape. If you want the backup jobs using this profile to always start at the beginning of a tape, select the option "***Rewind tape before starting job***". However, if you want the backup to always be appended to the end of the last backup performed to the tape, deselect this option. If you want to protect this or any other application from overwriting a backup once is complete, you can check the "***Eject tape***" option to automatically eject the tape from the drive at the end of a backup. This option is also handy if you are using a sequential autoloader and want each new job to start at the beginning of the next tape rather than being appended to the current tape.

Specifying the Data to Backup

The description of the first field on the screen will differ based on the backup type you selected for this profile. In this field, you may enter the ***data to backup***. This information is not required at this time and may be filled in when configuring the backup job later. The type of data to enter in this field will differ depending on the backup type. For instance, if this is a **Volume Group** or **System Backup** profile, you may enter a list of volume group names or type "**all**" to include all volume groups on the system. Likewise, if this is a **Filesystem** profile, you may enter a list of filesystems, etc. In addition to the "all" option, you may also enter a list of options to *exclude*. For example, to include all volume groups in a volume group backup EXCEPT the "rootvg" and "tempvg" volume groups, type:

```
all –rootvg –tempvg
```

If you want to exclude all volume groups on a **System Backup** you may leave this option blank. Leaving this option blank does have different effect depending on the type of client the backup is performed on. On an ***AIX*** system, leaving this option blank will still include the **rootvg** volume group (required on a base system). It will also include all volume group definitions of currently defined volume groups but will not backup the data within the excluded volume groups. On a ***Linux*** system, leaving this option blank will exclude all LVM data including their definitions and data.

If any items within the data list do not apply to a client, the item will simply be ignored. For example, using a filesystem profile containg "/var /tmp /home", a client without a /tmp filesystem will only backup "/var" and "/home".

After making all selections, save the profile by pressing the Save button at the bottom. The information will be saved and the window will be closed.

Pre-backup and Post-backup Programs

Within the backup profile, you may configure a program to run on either the client or server, before and/or after the backup command or backup job runs. You can also select to have programs execute before and after the creation of snapshots used for backups.. This program, either a *pre-backup* program or *post-backup* program, is a custom program which exists on one or more clients or servers, and may perform any operation, such as starting and stopping database programs, forcing users to log off the system, resetting tape library devices, etc. To configure pre- or post-backup programs, press the Configure button next to the ***Pre & Post Backup Programs*** field. When doing so, the following screen will appear:

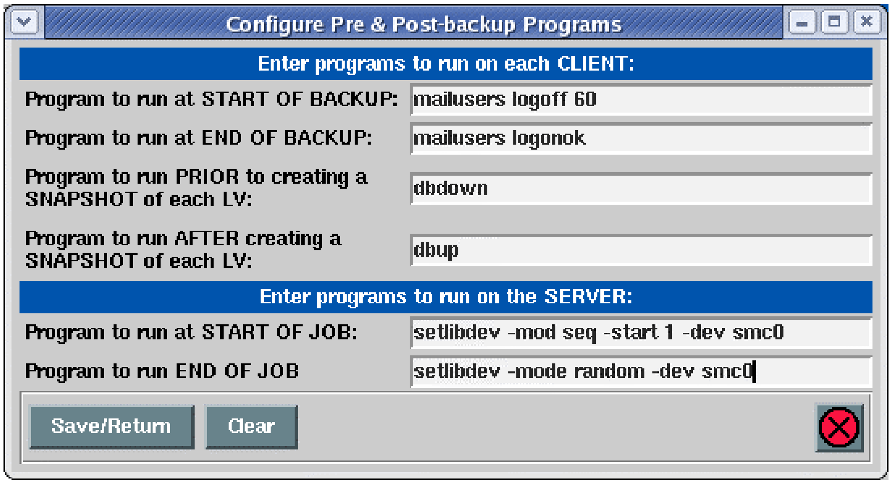

Configure Pre & Post-backup Programs

Enter programs to run on each CLIENT:

Program to run at START OF BACKUP: mailusers logoff 60

Program to run at END OF BACKUP: mailusers logonok

Program to run PRIOR to creating a SNAPSHOT of each LV: dbdown

Program to run AFTER creating a SNAPSHOT of each LV: dbup

Enter programs to run on the SERVER:

Program to run at START OF JOB: setlibdev -mod seq -start 1 -dev smc0

Program to run END OF JOB setlibdev -mode random -dev smc0

Note that the options for running programs prior to or after creation of snapshots is only available with backups of data contained in LVM logical volumes.

The pre-backup and post-backup programs will be executed with **ROOT USER** authority. Therefore, they must be placed in the ***DATADIR*/custom** directory by the root user on the client (where *DATADIR* is the directory you selected on each system when Backup Administrator was configured - i.e. */storix*). The *custom* directory is owned by the root user and only the root user on each system has the ability to add files to this directory. The commands placed in the custom directory may be shell scripts or binary programs and must have execute permission.

To configure a pre-backup or post-backup program, simply add the name of the program to the profile in either of the Pre-backup Program or Post-backup Program fields. Do not enter the full path name of the program, only the file name. The program is assumed to be in the ***DATADIR*/custom** directory. You may also add optional arguments to the command, separated by spaces.

Client Pre & Post Backup Programs

When a backup job using a profile containing client pre-backup or post-backup programs is run, the system will attempt to execute the specified program on each client before or after that client backup is performed. If the program does not exist on any client or is not executable, it will be ignored. Otherwise, it will be executed and one of the following actions will be taken depending on the exit code of the program:

	Pre-backup Program	**Post-backup Program**
Exit code 0	Client will be backed up and the job will continue normally.	Job will continue normally.
Exit code 1	Client will not be backed up and the backup job will be terminated with an error message	Job will terminate with an error.
Exit code 2	Client will not be backed up. If there are other clients to backup, the job will continue normally. However the job will complete with warning messages.	Job will continue normally. However, the job will complete with warning messages.
Exit code 3 or higher	Client will be backed up and the job will continue normally. However, the job will complete with warning messages.	Job will continue normally. However, the job will complete with warning messages.

NOTE **A post-backup program will be executed even if the backup command that precedes it fails. This is necessary in case the post-backup program must record information about the backup or restart processes that were stopped by the pre-backup program, etc.**

Pre & Post Snapshot Programs

When a backup job using a profile containing pre-snapshot and post-snapshot programs is run, the system will attempt to execute the specified program on each client before and/or after each logical volume snapshot is created.

> NOTE **For AIX systems, snapshots may only be created for mirrored logical volumes. If a logical volume is not mirrored, no snapshot will be created and the pre and post-snapshot programs will not apply to that logical volume.**

The program will only be executed for *LVM logical volumes* to be included in the backup, if Snapshot Backups have been configured (for the client if using Network Administrator), and the **Backup Job** is configured to perform snapshot backups.

> NOTE **The program names provided will be executed before a snapshot is created for *each logical volume*. Therefore, the program must be intelligent enough to recognize the name of the logical volume or filesystem the snapshot is being created for at that time and act accordingly (or do nothing). Refer to Creating Pre and Post Backup Programs below for more information.**

If the specified program does not exist on any client or is not executable, it will be ignored. Otherwise, it will be executed and one of the following actions will be taken depending on the exit code of the program:

	Pre-snapshot Program	**Post-snapshot Program**
Exit code 0	The snapshot of the logical volume will be created and the backup will continue.	Backup will continue normally.
Exit code 1	The snapshot will not be created for the logical volume and the backup will terminate.	The backup will terminate with an error.
Exit code 2	The snapshot will not be created for the logical volume, and the backup will continue using the active (online) data.	The backup will terminate with an error.
Exit code 3 or higher	A warning message will appear, but the snapshot will be created and the backup will continue normally.	The backup will terminate with an error.

Backup Server Pre & Post Backup Job programs

> NOTE **Server pre and post backup programs are only available with Network Administrator.**

When a backup job using a profile containing a server pre-backup or post-backup job program is run, the system will attempt to execute the specified program on the server before the first client backup (pre) or after the last client backup (post). This allows you to perform operations such as initializing tape libraries before backups are performed to the backup server. If the program does not exist on any client or is not executable, it will be ignored. Otherwise, it will be executed and one of the following actions will be taken depending on the exit code of the program:

	Pre-backup Program	Post-backup Program
Exit code 0	Job will continue normally.	Job will complete successfully.
Exit code 1	No clients will be backed up and job will terminated with an error	Job will terminate with an error.
Exit code 2 or higher	Client backups will continue. When backup are complete, job will terminate with a warning message.	Job will complete with warning messages only.

NoTE **A post backup job program will be executed <u>even a client backup fails or another error occurs</u>. This is necessary in case the post-backup program must record information about the backup or restart processes that were stopped by the pre-backup program, etc.**

Creating Pre & Post Backup Programs

A customized program may perform any function on the system since it is run under *root user* authority. Any arguments or flags may be provided to the command. The same script may be called with arguments that tell the script how to proceed. For example:

`mypreprogram –kill`	may be used to log off users and
`mypreprogram –warn`	may warn users of the backup only, or
`mypreprogram –kill 60`	may warn users, then log them off after 20 seconds, etc.

In many cases, it is desirable for the program to have certain information about the backup job. The program may want to display or save information about the backup job in another application or file, or a post-backup program may need to respond differently depending on whether the backup was successful or not. Every program will have access to the following environment variables:

STX_SERVER	The name of the backup server
STX_DEVICE	The name of the device on the server (or directory)
STX_JOBID	The Job ID
STX_BACKUPID	The Backup ID
STX_EXITCODE	The exit code of the backup command or job
STX_SNAPLVNAME	The logical volume for which a snapshot is created.
STX_SNAPFSNAME	The filesystem name (mount point) of the snapshot LV.. This will show a dash "-" if the logical volume is not a filesystem.

The **STX_EXITCODE** variable is only used in client or server post-backup/job programs. For client programs, this indicates the success or failure of the backup. On servers, indicates the success or failure of the overall backup job.

The software is installed with sample script programs that may be used for any client or server pre-backup, post-backup or pre/post snapshot program. The programs are called “**prepost.sample**” and "**prepostsnap.sample**" and will simply display the values of all of the above variables when the backup job is run. You may edit or view the contents of this script file (contained in the ***DATADIR*/custom** directory), which contains additional details on the use of this option.

Incremental/Differential Backups

An *incremental backup* is on in which the only data to be included in the backup is that which has changed since the prior incremental backup level. An incremental backup level can be from 0 to 9, where 0 is a “full incremental” backup from which all other levels are based. Levels 1 through 9 indicate that only data that has changed since the last ***prior-level*** backup should be included.

Differential backups are also incremental backups, except that backups include a cumulative list of files that have changed since a certain time. This is achieved by running the same incremental level backup repeatedly, this backing up the same files that changed since the last prior-level (or level 0) backup, along with any additional files that have changed since the last time the same incremental level backup was run. The result is that the backup gets continually larger each time it is run, until a prior level (or level 0) backup is run again.

Incrementals may be performed for both **Volume Group** and **Filesystem** backups. Backing up a volume group incrementally is the same as performing the same level of incremental backup for every filesystem and logical volume within. If a volume group contains raw logical volumes (those with no filesystems built on them) AND you are including raw logical volumes in the backup, then the raw logical volumes will be included in their entirety if they have been written to since the last prior-level backup.

Incremental Backup Examples

1. Consider the following backup schedule:

Monday	**Level 0**
Tuesday	**Level 1**
Wednesday	**Level 2**
Thursday	**Level 3**
Friday	**Level 4**

   a. On Monday, all of the data in the specified filesystem or volume group will be backed up, and the volume group or the next level of backups will be based.

   b. On Tuesday, only the files or logical volumes that have changed since Monday’s backup will be included in the incremental level 1 backup.

   c. On Wednesday, only files backed up since the last ***prior-level backup*** (level 1) will be included in this backup.

   d. Likewise on Thursday and Friday.

   e. On the following Monday, a new incremental level 0 is performed, backing up all data regardless of prior incrementals, and thus making the filesystem(s) will be flagged as having had a full incremental backup performed. This is the backup from which other incremental levels obsolete.

2. In a second example, consider the following backup schedule, which is often referred to as ***differential*** backups since we’re effectively backing up the differences between a filesystem now versus a specific day in the past :

First day of the Month	**Level 0**
Each Friday night	**Level 4**
Each other weekday	**Level 7**

   a. On the first day of every month, regardless of the day of the week, a full incremental backup is performed.

   b. The next day, an incremental level 4 will be performed (if Friday) or an incremental level 7 will be performed (if Monday through Thursday)

   In this example, keep in mind that it is not necessary to perform a level 1 backup after a level 0, since each level (1-9) will backup the data from the last ***prior-level*** backup performed, even if it was several levels prior. Therefore, if your last level was 0 (full), then either a level 4 or a level 7 will backup the same data. However, if your last level was 4, a level 7 will always backup files changed only since the last level 4.

   In addition, the same backup level will be performed several times in a row. Since all data will be backed up since the last ***prior*** level, your last backup of the same level will become obsolete.

3. This example is a ***differential*** backup, where all backups are based on the most recent level 0 (full)_ backup that was performed:

Every Friday night	**Level 0**
Monday through Thursday night	**Level 1**

   a. Every Friday night, a full backup level 0) is performed

   b. On every other night, a level 1 backup is performed. The result is that, each day, all files that have been created or changed since the Friday night backup will be backed up. The size of the backup will grow each day until after the next Friday night backup is performed.

Restoring from Incremental Backups

There are a few things to remember when restoring from incremental backups in order to get your data back to the most recent state:

a. Always start by restoring from your most recent incremental level 0

b. Always restore full Volume Groups or Filesystems from incremental backups. If you choose to restore a directory from a Filesystem backup, even if that directory is the filesystem mount point, you will restore all files from the backup, but will not re-apply all changes such as re-removing files which no longer existed when that incremental backup level was performed.

c. Restore incremental levels in the order they were performed ONLY if the next incremental level to restore is more recent than the last. For instance, if you performed a level 1 backup most recently, do not restore a level 2 backup which is older.

d. When you perform the same incremental backup level multiple times without performing a lower-level, restore only the most recent backup of that level.

In the first backup example above, you must restore each backup, starting with level 0 in the order of each backup level, stopping when you encounter a backup level that is older than this predecessor. If your level 1 backup was most recent, then you will need to restore only level 0 and 1. If your level 4 was most recent, you will need to restore all levels 0 through 4.

In the second example, you are ensured never to have to restore more than three backups to get your data up-to-date. This convenience comes with some complication when restoring. First, you must or course always start by restoring your last level 0. Then, if there was a higher level backup performed after your level 0, restore it next (it could be a 4 or 7 depending on what day is the first day of the month). Lastly, if you restored a level 4 and there was a level 7 backup performed after your level 4, restore it next.

Changing a Backup Profile

The information for an existing profile may be changed by selecting Configure→Backup Profiles from the menu bar. The profile selection screen will appear. Select a profile from the list, then press the Add/Change button at the bottom. The profile options screen will then appear with the prior settings for the profile. Simply add or change any of the information on the screen, then press the Save button at the bottom to save the changes.

Removing a Profile

A profile may be removed from the system only if it is not assigned to any backup jobs. If it is assigned to a job, you will be informed so, and you must remove or change the job to use a different backup profile before the current profile may be removed.

To remove a profile, select Configure→Backup Profiles from the menu bar, then select the profile to remove from the listbox (see profile selection screen) and select the Remove button.

8. Tape Libraries and Autoloaders

> **This features described in this chapter are not available with Desktop Edition.**

A tape libraries and sequential autoloaders (also known as “stackers”) are devices that contain one or more tape drives and are capable of moving tapes between the tape drives and various tape storage slots. Although a sequential autoloader can only operate in sequential mode, most libraries can be configured either as a sequential autoloader or a random library.

To use a tape library, you must first create a virtual device and specify whether it is an autoloader or random tape library. By configuring the library or autoloader as a virtual device, you will also be able to take advantage of other virtual device features. You may not backup directly to a library (since technically a library is a changer, not a backup media), but must backup to a virtual device, which in turn uses the library definition to change volumes as needed.

Sequential Autoloaders

Most tape libraries have what is called a "*sequential mode*". When set to this mode, the loader will automatically detect when a tape has been ejected from the drive and will automatically replace that tape with the next sequential tape in the storage slot. In this case, the software does not communicate with the stacker device, only to the tape drive.

System Backup Administrator takes advantage of the sequential mode of the autoloaders by ejecting a tape at the end of a volume and simply waiting for the loader to do its job. As soon as a new tape is inserted, the backups will continue with no need to bother the user.

> **Autoloaders will recognize when a tape has been ejected and will load the next tape in sequence, however it is required to have a tape loaded in the drive before starting a backup.**

Any virtual device may be setup to use an autoloader. If so, the tape changes are expected to take place automatically in each drive, regardless of the number of drives assigned to the virtual device. To separate autoloaders may be used, each with a single tape drive, or a single autoloader with two tape drives can be used. In either case, the backups would perform basically the same.

If you have an autoloader with a single tape drive, you must create a Sequential Virtual Device for that tape drive alone if you want it to be identified as an autoloader. Note, however, that the only difference between backing up to the sequential virtual device and backing up directly to the tape drive is that volume prompts will not appear on the screen when a sequential virtual device identified as an autoloader is used.

Random Libraries

A random library is a tape library that does not automatically load and stack tapes. Tape movement is performed manually by using sets of commands that the library driver can understand. These commands vary depending on the operating system and the library, but generally the functions are to move tapes from their current location in the library to the drive and back again.

SBAdmin can be configured to utilize a random library so that backups and restores can be performed as if the library was a sequential autoloader. When a tape is ejected, SBAdmin will execute the necessary commands to remove the tape from the drive and insert the next tape in the library as specified by the library configuration. One advantage of using a random library instead of an autoloader is the ability to start a backup without having a tape in the library.

Random tape libraries are not supported by the system installation process. If you have a multi-volume backup to install from, you will need to insert volume one of the backup, then set your library to work in sequential mode. Otherwise, you will need to change tapes manually when prompted by the installation process.

Tape libraries are available in many different hardware configurations. The number of tape slots and the number of tape drives are not always the same even for a particular brand and model of library. The configuration of all tape libraries used by SBAdmin, must be defined prior to their use. Before you configure a library, it is best to get an understanding of the two main classifications of libraries used by SBAdmin – Single Drive Libraries and Multiple Drive Libraries.

Single Drive Libraries

A single library is defined as a library with either only one tape drive or it is configured to use only one of the multiple tape drives available. The following picture illustrates a single-drive library configuration or a two-drive library that has been split into two separate single-drive library configurations.

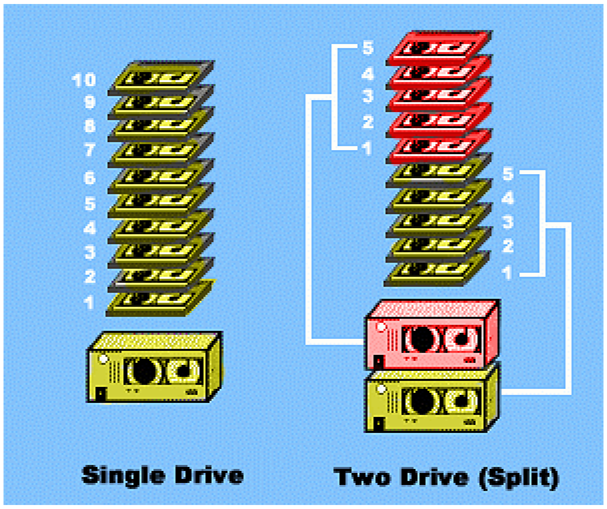

Multiple Drive Libraries

A multiple drive library configuration can use two or more drives which are used concurrently for a single backup job. The following picture illustrates the three different virtual device configurations, each using a two-drive library. In each case, the same tapes are assigned to the same drives, but the numbers indicate the volume numbers of the backup if all tapes were used. See the section on virtual devices for more information on multi-drive configurations.

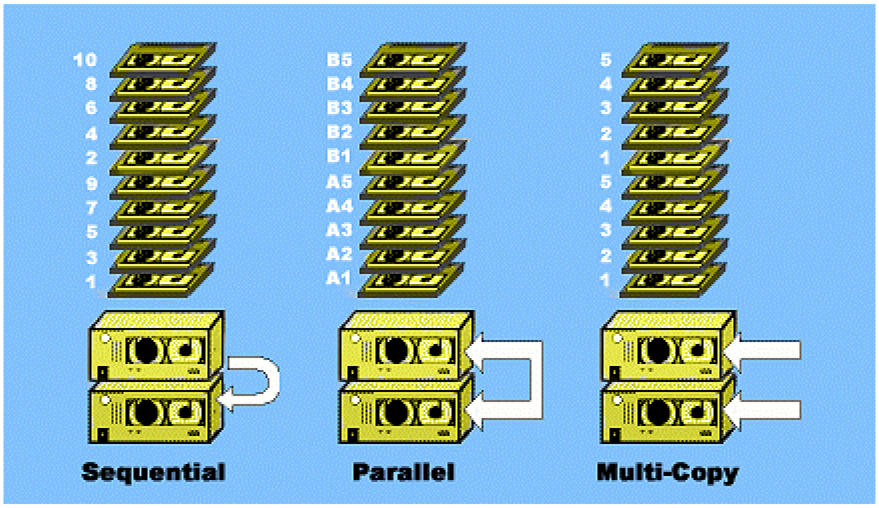

Using Multiple Drives in a Library Independently

If you configure your library to use 2 drives, then both drives will be used during a backup. Depending on your virtual device configuration, the backup data may be written to the drives sequentially, in parallel, or by writing a separate copy to each drive.

You may also want each drive in the library to act independently, allowing separate backups to be performed to each. If using a random library, you will need to create a separate library name for each drive, and assign a set of tapes to each of the drives. If using a sequential library, the library must support changing the tapes in the drives independently (usually referred to as “split-sequential mode”).

In either case, this will be the same as if you had two single-drive libraries sitting side-by-side.

When configuring separate random libraries, be sure you always assign different tape slot positions to different drives. Refer to the predefined library definitions with the suffix of “drive1” and “drive2” for examples.

Configuring Random Tape Libraries

A tape library must be configured within SBAdmin, so that the software knows the number of drives, the number of tapes to assign to each drive, and the commands used to move the tapes. In an effort to make configuring a random tape library an easy process for users, SBAdmin includes a number of predefined library definitions that may fit your environment. Additional Library Profiles can be created or existing profiles can be adjusted to fit a particular need.

To configure a Random Tape Library, select Configure→Random Tape Libraries from the menu bar. When you do so, a Configure Tape Library screen similar to the following example will appear.

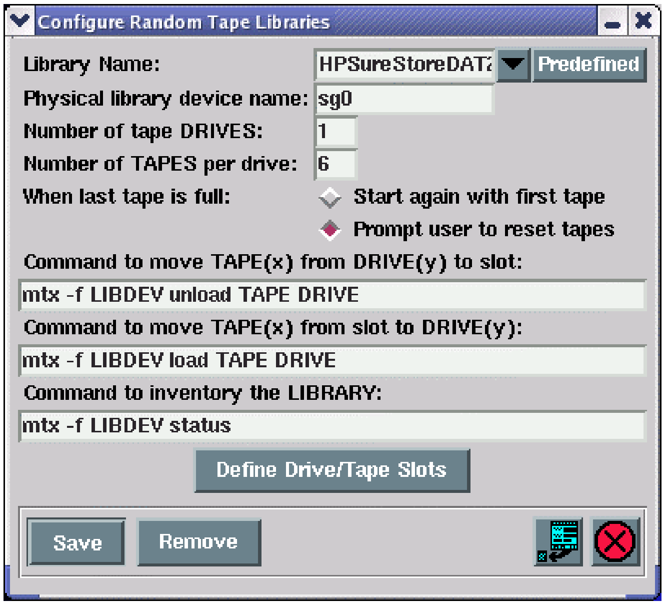

Library Name: HPSureStoreDAT

Physical library device name: sg0

Number of tape DRIVES: 1

Number of TAPES per drive: 6

When last tape is full: ○ Start again with first tape ● Prompt user to reset tapes

Command to move TAPE(x) from DRIVE(y) to slot:

```
mtx -f LIBDEV unload TAPE DRIVE
```

Command to move TAPE(x) from slot to DRIVE(y):

```
mtx -f LIBDEV load TAPE DRIVE
```

Command to inventory the LIBRARY:

```
mtx -f LIBDEV status
```

Select the **Predefined** button and search for the library that best fits your environment. After making a selection, the configuration of that predefined library is displayed. You may choose to keep the same **Library Name** or change it to any name you choose.

You may then change the additional fields as needed.

When using a predefined library, it is often necessary to change the name of the physical device because SCSI devices are named by the system in detection order.

Note the following keywords that are used in the **Command** fields:

Standard Library Commands

The command string used to move tapes in the library contain specific keywords that are replaced when the command is called. SBAdmin predefines "**tapeutil**" (***AIX***) and "**mtx**" (***Linux***) as typical tape library utilities. Refer to the section below if you plan to use different commands.

Although SBAdmin recognizes "tapeutil" and "mtx" as standard library commands, these commands are not supplied by SBAdmin. The "tapeutil" command is installed with the IBM Atape driver (on AIX) and "mtx" is a free Linux command which you may need to download from a free software site if it is not already installed.

The following is a list of **Keywords** that SBAdmin will replace with values when the command is executed:

LIBDEV: This variable will be replaced with the physical library device name.

DRIVE: This variable will be replaced with the library element or slot position of the drive (See Define Tape/Drive Slots).

TAPE: This variable will be replaced with the library element or slot position of the tape (See Define Tape/Drive Slots).

Custom Library Commands

You may choose to use or create different tape utility commands than the standard "**mtx**" and "**tapeutil**" commands that SBAdmin recognizes. However, you must add the names of the commands to execute to a configuration file on the server (if remote) to which the library is attached. To add a new tape library command, edit the ***/storix*/config/library_cmds** file (where ***/storix*** is replaced with your data directory if configured differently), and add the name of the library command. You may not insert the full pathname of the command, so you should copy or link your command to the /usr/bin directory to be sure it is found in the standard command search path.

The variables listed below are *optional* and can be used to create custom scripts to run in the place of your standard library utilities:

> **VDEV**: This keyword will be replaced with the name of the virtual device (the device your library is assigned to).
>
> **TAPEDEV**: This keyword will be replaced with the physical tape device name known by the system. (i.e. st0, rmt0). For libraries with more than one drive, the TAPEDEV will reference the specific drive a tape is being moved to or from.
>
> **SERVER**: This keyword will be replaced with a server name, if your virtual device is remote.
>
> **CLIENT**: This keyword will be replaced with a client name when performing a backup of a client (only if Network Administrator license is used).
>
> **BACKUPID**: This keyword will be replaced with the Backup ID number when a backup is being performed.
>
> **JOBID**: This keyword will be replaced with the current job id number when a backup is being performed.

For example, if you create a script called “**mytapeutility**”, place it in /usr/bin and add it to the **/storix/config/library_cmds** file), you may specify this command in the Library Configuration screen as:

```
mytapeutil get LIBDEV DRIVE TAPE VDEV TAPEDEV SERVER /tmp/liblog
```

And “**mytapeutility**” could be a script such as the following:

```
#!/bin/sh
action=$1
libdev=$2
drivenum=$3
tapenum=$4
vdevname=$5
tapedevname=$6
server=$7
log=$8

if [ $action = get ]
then cmd=”mtx –f $libdev load $tapenum $drivenum”
     echo “Moving tape #$tapenum to drive #$drivenum” >>$log
else cmd=”mtx –f $libdev unload $tapenum $drivenum”
     echo “Returning tape #$tapenum from drive #$drivenum” >>$log
fi
echo “Server is $server, device is $tapedevname ($vdevname)” >>$log
echo “Executing: $cmd” >> $log
$cmd
exit $?
```

Define Drive/Tape Slots

On the Library Configuration screen, select the **Define Drive/Tape Slots** button. When you do so, a Define Library Drive and Tape Slots screen similar to the following example will appear.

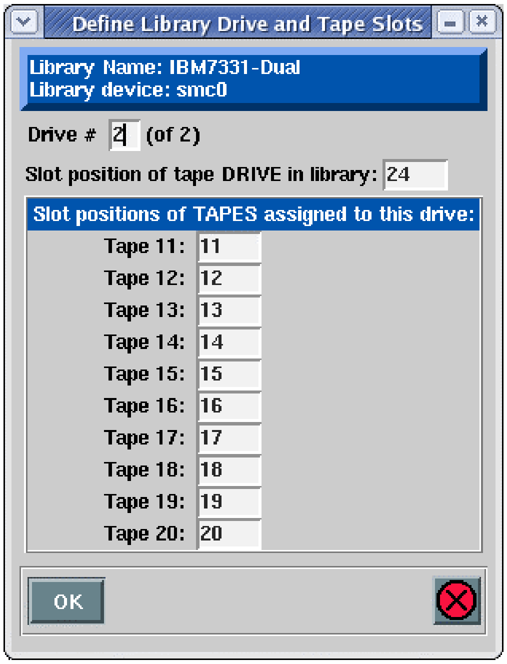

Library Name: IBM7331-Dual
Library device: smc0

Drive # 2 (of 2)

Slot position of tape DRIVE in library: 24

Slot positions of TAPES assigned to this drive:

Tape 11:	11
Tape 12:	12
Tape 13:	13
Tape 14:	14
Tape 15:	15
Tape 16:	16
Tape 17:	17
Tape 18:	18
Tape 19:	19
Tape 20:	20

The slot positions referred to are the ***physical slot*** or ***element location*** that the library uses to reference the positions of tapes and drives. The location you provide will determine what tape position is assigned for each tape used with SBAdmin. In the above example, SBAdmin's *tape number* "1" for *drive number* "1" is referenced by the library as slot position (or element address) "32".

> NoTE
>
> **If the library contains more than one drive, you may create a separate library name for each drive (allowing different backups to be performed simultaneously) or multiple drives may be configured with a single library name (allowing the drives to be used concurrently by the same backup process).**
>
> **In either case, you must be sure that you do not define the same tape slot positions for both drives! When doing so, SBAdmin will attempt to use the same tapes in both drives and will fail.**

Be sure to define the tape slots for both drives if using a 2-drive library. Never enter the same slot position in more than one field, else SBAdmin will try to use the same tape for different volume numbers of the same backup.

9. Virtual Devices

> NoTE **This features described in this chapter are not available with Desktop Edition.**

A virtual device is used to group one or more physical ***tape drives*** into a single "*virtual*" device. The virtual device may then be used just like any single tape drive. Using a virtual device provides added functionality, capacity and performance depending on the type of virtual device configured.

Types of Virtual Devices

There are three types of virtual devices. They are:

- Sequential Virtual Device - This device may contain one or more tape drives. If multiple drives are used, the backup will start on the first drive and automatically continue on the next drive when the first tape is full. The user is prompted to change tapes only when the tape in the last drive is full. Backups created to a sequential virtual device may be restored using any single drive, provided all drives used to create the backup were of the same type. A single drive may also be configured as a virtual device, however, there would only be an advantage if you are using a sequential autoloader or random library. If so, the virtual device will eject the tape when it becomes full, and the autoloader or random library will automatically change the tape, allowing the backup to continue unattended.

- Parallel Virtual Device - This device must consist of two or more tape drives. The data in the backup will be evenly spread, or *striped*, across all of the tape drive, allowing the backup to complete in a fraction of the time it would take to write to a single drive. The same virtual device number (or one containing the same number and types of drives) of drives must be used to restore data from the backup.

- Multi-copy Virtual Device - This device must consist of two or more tape drives. When backup are sent to this device, the same data is written to all drives, providing multiple copies of the same backup in about the same time it would normally take to make a single copy to a single drive. Any copy of the backup may be read from any single tape drive.

How Data is Stored on Virtual Devices

If a virtual device is used to read a backup made by a virtual device, it is necessary that the tapes be placed in the corresponding tape drives in the same fashion as they were when they were backed up. This does not apply to a multi-copy virtual device backup since each set of tapes from a single drive are independent copies of the same backup, and may only be read by a single device or a sequential virtual device.

The following illustrations show how the data is saved on each of the different virtual devices. Note how the tape volumes are numbered for each virtual device type when data spans a different physical tape.

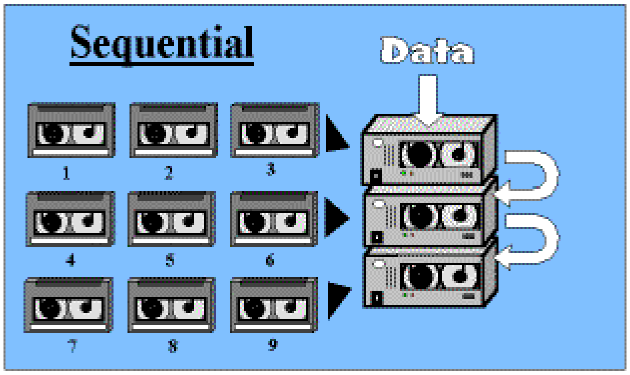

The sequential data is written to the first tape in the first drive until it fills up. Then, the backup continues onto the next device in the list, etc. Only when all tapes in all devices are filled will the user be prompted to change volumes. Then, the backup continues on the first drive and so forth. The volumes are numbered in sequential order. Assuming all drives are of the same type, the backup will be identical to a backup written to a single tape drive, so restores may be done either with the same virtual device or from any one of the tape drives.

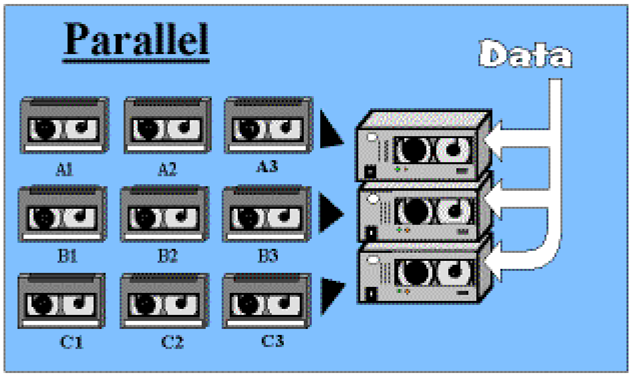

The data is split into multiple buffers (one for each device) of the same size, and the data is sent to all three devices at the same time. The user will be prompted to change the volumes in each device as it fills up, which may not necessarily be at the same time, particularly when using different types of tape drives in the same parallel device. The volumes are numbered with a letter representing the tape drive (A, B, C..) and a number representing the volume in each drive (1, 2, 3...). Data from a parallel virtual device backup may only be verified or restored using a virtual device with the same number and types of drives.

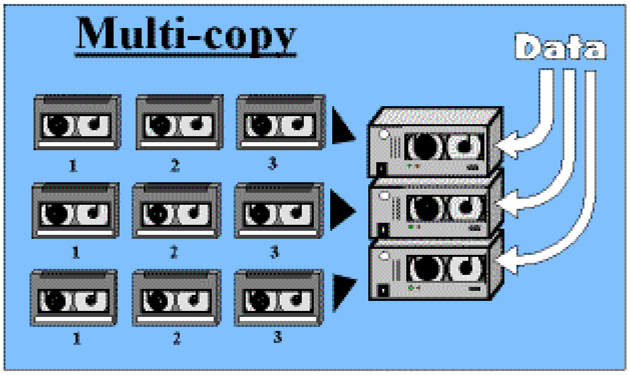

An identical copy of the same data is written to each of the drives at the same time, normally at about the same time it would take to write to a single device. Since the same data is written to each device, the volumes are numbered as though each is a single-device backup. Since a multi-copy virtual device backup looks identical to multiple sequential backups, each of the backups can only be read using one drive at a time.

Add a Virtual Device

A new virtual device may be added by selecting Configure→Virtual Devices from the menu bar. When you do so, a **virtual device selection screen** similar to the following example will appear:

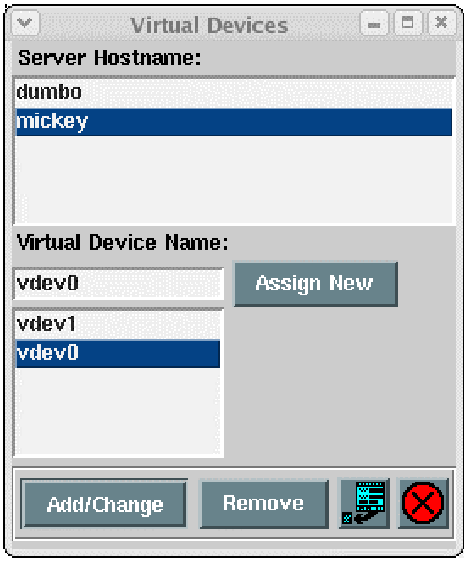

Server Hostname:
dumbo
mickey

Virtual Device Name: vdev0
vdev1
vdev0

To add a new virtual device, you must first select the backup server the tape drives are attached to. Then, you may either type a new virtual device name in the entry field, or press the Assign New button to generate the next sequential virtual device name for the server. A virtual device name must begin with the "*vdev*" prefix. Once you have made your selections, press the Add/Change button to continue. Additional options will be provided on the **virtual device options screen** as shown in the following example:

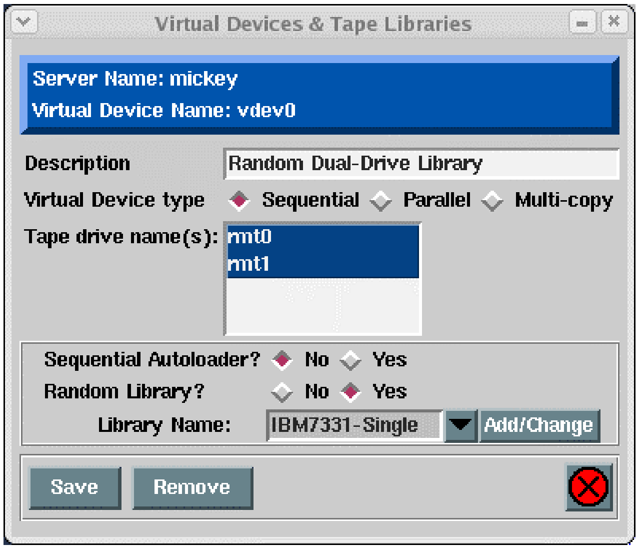

Server Name: mickey
Virtual Device Name: vdev0

Description: Random Dual-Drive Library
Virtual Device type: ● Sequential ○ Parallel ○ Multi-copy
Tape drive name(s): rmt0, rmt1

Sequential Autoloader? ● No ○ Yes
Random Library? ○ No ● Yes
Library Name: IBM7331-Single

From this screen, you may enter a description or simply use the default description displayed. Select the ***virtual device type*** and whether or not an ***autoloader*** or ***random library*** is used. The listbox will contain a list of the physical devices currently available on the server. For a sequential virtual device, you may select one or more options. For other virtual device types, you must select at least two physical devices.

Sequential Autoloader

Select the “Yes” button if the selected tape drives are contained in a sequential autoloader or “tape changer”. This may be also be a tape library set to sequential mode. When you backup or restore using a virtual device configured as an autoloader, no volume prompts will appear on the screen at end of volume, but the tape will be ejected and the process will wait for a new tape to be inserted. No commands are issued to the autoloader device, but the process expects that the autoloader will automatically insert the next tape as needed. The backup or restore will continue automatically when the new tape is inserted. Note that, when beginning a backup or restore, the first tape containing the backup to read or write must already be inserted before the process starts.

Random Library

Select the “Yes” to indicate the tape drives are contained in a random tape library. When backing up or restoring from a random tape library, the tape will be ejected at the end of volume, and the administrator will issue the commands needed to return the tape to its original slot in the library, then grab and insert the next tape number. The backup or restore will continue automatically when the new tape is inserted. When starting a backup or restore, if the tape is not already inserted in the drive, the administrator process will grab and insert the tape automatically. The tape to grab must be set in the Set/Reset Next Tape for Backup/Restore option.

After selecting “Yes”, the drop-down list will be enabled. Here you must select the name of the tape library configuration to use. If you have not already configured a tape library, you may do so now by pressing the Add/Change button to the right. The Configuring a Random Tape Library screen will appear, from which you can view, add or remove tape library definitions.

After making all selections, save the virtual device by pressing the Save button at the bottom. The information will be saved and the window will be closed.

If the Clients, Servers & Devices is displayed on the Main Screen and the server for which the virtual device was configured is selected, the new virtual device name and image will automatically appear on the screen.

Change a Virtual Device

The information for an existing virtual device may be changed by selecting Configure→Virtual Devices from the menu bar. On the virtual device selection screen, you can simply select the server and the virtual device to change from each of the listboxes, then press the Add/Change button at the bottom. The virtual device options screen will then appear with the prior settings for the virtual device. Simply add or change any of the information on the screen, then press the Save button at the bottom to save the changes.

Remove a Virtual Device

A virtual device may be removed only if it is not currently assigned to any backup jobs. If it is, you will be informed so, and you will need to either remove or change the job to backup to a different device.

To remove a virtual device, select Configure→Virtual Devices from the menu bar. the virtual device selection screen will then appear. Select the server on which the virtual device is configured. The virtual devices configured on the server are then displayed. Select the virtual device to remove from the listbox, then press the Remove button.

10. Exclude Lists

Exclude lists are used to exclude certain files, directories, or devices (such as partitions or logical volumes) from backup jobs. You may create any number of different exclude lists, and assign *one or more* exclude lists to a particular backup job. You may also select which clients the exclude list will apply to. This allows you to use an exclude list for a job, but still have it only apply to certain clients if multiple clients are backed up by the same job.

Note that you may also select certain data to include or exclude on each backup when configuring a backup job (depending on the backup type). You can specify, for instance, the filesystems to include on a filesystem backup (or all filesystems except certain ones). Using an exclude list as described in this section, however, will provide the ability to exclude specific files or directories within the filesystems.

Exclude lists may be used to exclude files, directories, entire filesystems or *device data* (such as partitions or logical volumes) from various backups. *Wildcard* characters (*) in exclude list entries may also be used to exclude may files or directories matching a certain pattern.

Device names may also be added to the exclude list. A device name may be an LVM *logical volume*, *meta-disk* (software RAID) device name, or disk *partition*. The data within the device will only be excluded if it is ***not*** used for a filesystem. To exclude a filesystem, you must exclude the filesystem mount point (directory).

Using Wildcards

If you wish to exclude a directory, all files within the directory as well as any sub directories will also be excluded. A ***wildcard*** (*) may be used in an exclude list entry for files and directories. For instance, having ***/usr/local/*.old*** in the exclude list will exclude all files in the /usr/local directory with a ".old" extension. Wildcards in the exclude list work the same as at the command line. For example, typing "`ls /usr/local/*.old`" will yield the same list of files that will be excluded if */usr/local/*.old* is in the exclude list. You may specify multiple wildcards in the same string. For example, "*/*/local/x*.old*" will exclude files starting with an "x" and ending with ".old" in the /usr/(*anydir*)/local directory.

NOTE **You may not use other special characters in exclude list entries, even if they exist in the names of the files to exclude. Those characters are $, +, ? and ^, which have special meaning to the system.**

Adding an Entry to the Exclude List

Select Configure→Exclude Lists from the menu bar to display the following **exclude list screen**:

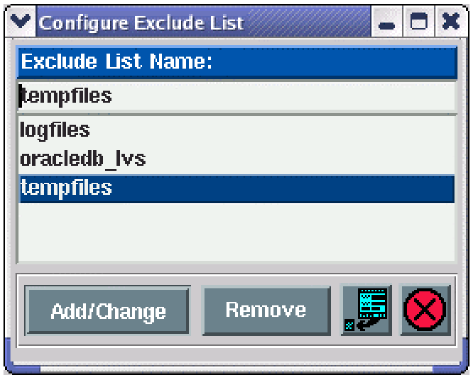

You may enter a new exclude list name in the entry field at the top of the screen, or select an existing exclude list name from the listbox below to change the exclude list. When you have made your selection, press the

Add/Change button to continue. A new window such as the following example will appear:

The *Clients* listbox will only appear when using Network Administrator.

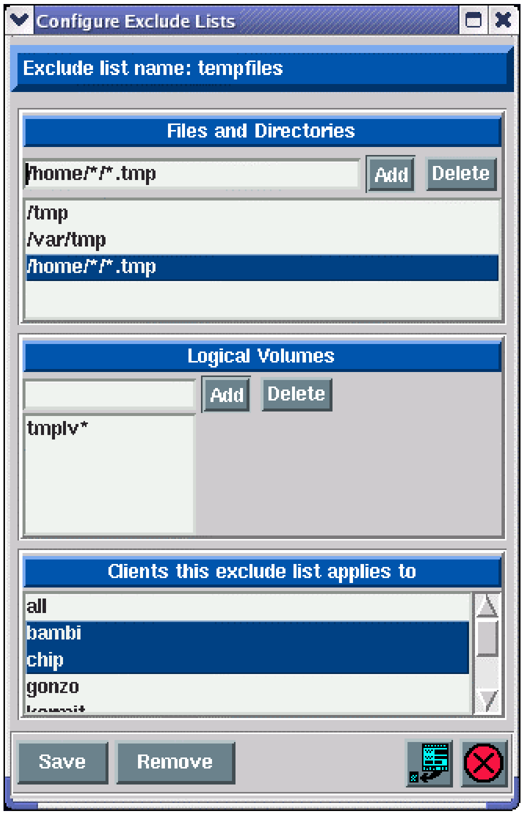
Configure Exclude Lists

Exclude list name: tempfiles

Files and Directories

/home/*/*.tmp

/tmp
/var/tmp
/home/*/*.tmp

Logical Volumes

tmplv*

Clients this exclude list applies to

all
bambi
chip
gonzo
kermit

To exclude files or directories, type the file or directory name (or wildcard string) in the entry field under the ***Files and Directories*** heading. To add a *logical volume, partition* (***Linux***) or *meta-disk* (***Linux***) to the exclude list, enter the device name (do not prefix with /dev) in the entry field under the appropriate heading. Note that the heading will only show **Logical Volumes** if only ***AIX*** support is enabled, and only **Partitions/Meta-disks** if no LVM support is enabled for ***Linux*** systems.

In the last listbox, you may select "**all**" to apply this exclude list to all *clients* (when assigned to a backup job), or select individual clients the exclude list should apply to.

Press Enter or select the Add button next to the corresponding entry field to add the item to the list.

When all selections have been made, press the Save Changes button at the bottom of the screen to save the entries and exit. To undo all changes made, press the cancel button at the bottom.

Removing Entries from the Exclude List

To remove an entry from the exclude list, display the exclude list screen by selecting Configure→Exclude Lists from the menu bar and selecting the exclude list to change. Then, to remove a file or directory entry, select on the item in the ***Files and Directories*** listbox and press the Remove button next to the file or directory entry field. Likewise, to remove a logical volume from the list, select the item in the ***Logical Volumes*** listbox and press the Remove button next to the logical volume entry field. When you have removed all desired selections, press the Save/Done button at the bottom of the screen to save the remaining entries and exit. To undo all changes made, press the cancel button at the bottom.

To remove an entire exclude list, select Configure→Exclude Lists from the menu bar, select the exclude list from the listbox, and press the Remove button at the bottom of the screen. Note that, when removing an exclude list that is assigned to current backup jobs, the exclude list will be removed from the job configuration. Yu will be informed if the exclude list is assigned to any jobs before proceeding.

11. Backup Jobs

A backup job must be created before any backup may be performed by the admin system. A job is not required when running backup directory from the client using the stbackup command. The job information will identify the backup server and one or more clients to backup (when using the *Network Administrator*), the backup profile, and the device or directory on the backup server to send the backup to. If the backup is to be scheduled to run either at a later time or on a regular basis, the dates and times are also added to the backup job information. Temporary backup jobs which are run only once may also be set to be automatically deleted once the job has completed.

Before configuring a backup job, when using a *Network Administrator*, you must first have configured at least one client to backup and a backup server to backup to (even if the client and server are the same). There must also be at least one backup profile for the type of backup to be performed (several sample profiles come installed with the software). On the job configuration screen, you can customize the selected backup profile to apply changes which apply only to the job, if desired.

NOTE **Clients may only be assigned to a backup job which uses a backup profile compatible with the operating system type of the client. For instance, an *AIX* client cannot be added to a backup job using a Raw Partition backup profile since *AIX* systems do not support partitions. Windows (SMB) Clients may only perform SMB (Windows) Share backups.**

Creating a Backup Job

To create a backup job, either

1. Select Configure→Backup Jobs from the menu bar, or

2. Select the Add Job button when the Job Information is displayed on the Main Screen.

The following backup **job entry screen** will be displayed:

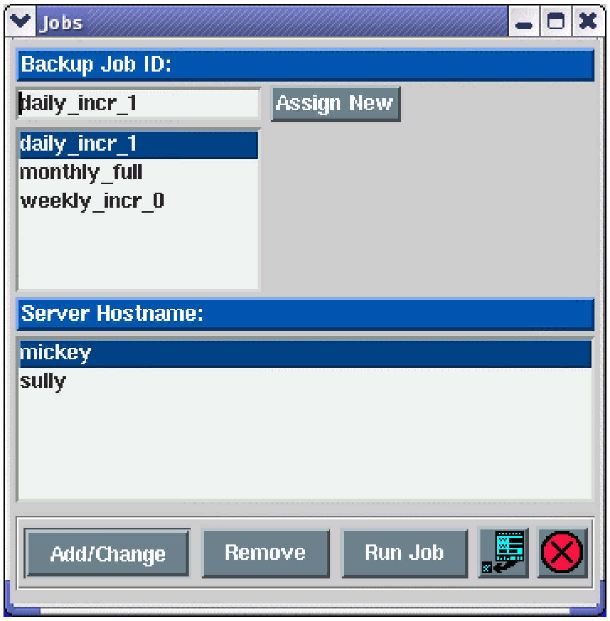

daily_incr_1

daily_incr_1
monthly_full
weekly_incr_0

mickey
sully

First, enter the Backup Job ID in the entry field at the top. The Job ID is used as unique identifier for this job, and may consist of any letters or numbers except for a colon (:) or space (spaces will be automatically replaced with an underscore). You may also press the Assign New button to generate a new sequential 6-digit number. This number begins with *000001* when the software is installed and will be increased sequentially each time you press this button, even if a prior number was not used.

Next, if you are using a *Network Administrator*, you must select a server from the ***Server Hostname*** listbox. The backup will be sent to a backup device on the selected server.

When you have made your selections, press the Add/Change button to proceed. The following is an example of the **job options screen** which appears next:

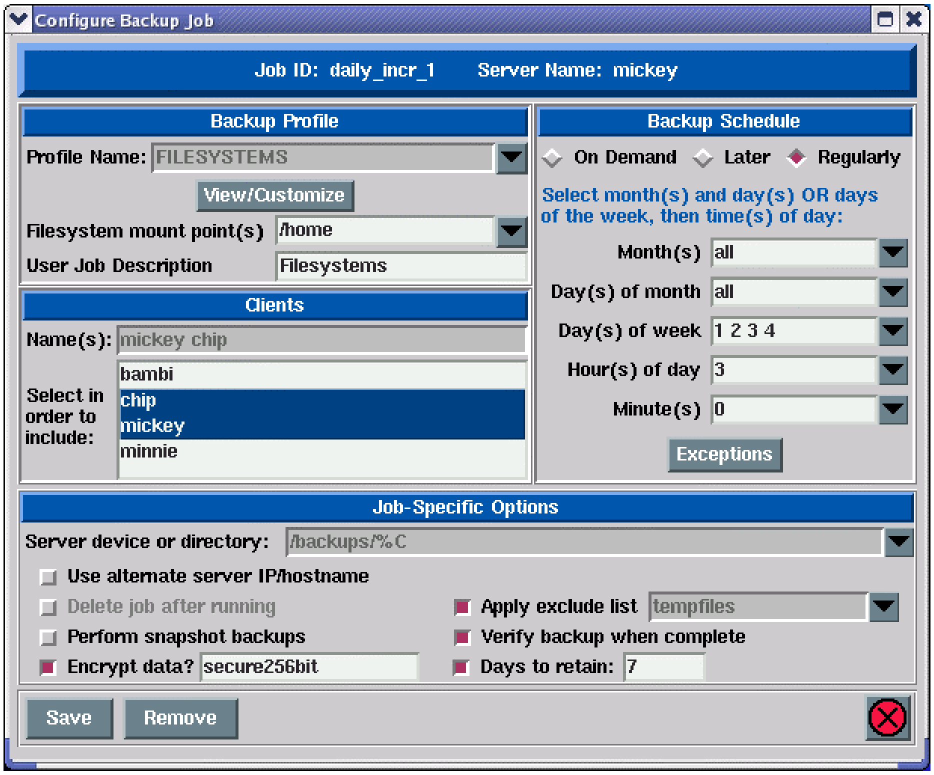

Job ID: daily_incr_1 Server Name: mickey

Backup Profile

Profile Name: FILESYSTEMS

Filesystem mount point(s) /home

User Job Description Filesystems

Clients

Name(s): mickey chip

Select in order to include: bambi, chip, mickey, minnie

Backup Schedule

○ On Demand ○ Later ● Regularly

Select month(s) and day(s) OR days of the week, then time(s) of day:

Month(s) all

Day(s) of month all

Day(s) of week 1 2 3 4

Hour(s) of day 3

Minute(s) 0

Job-Specific Options

Server device or directory: /backups/%C

□ Use alternate server IP/hostname

□ Delete job after running

□ Perform snapshot backups

☒ Encrypt data? secure256bit

☒ Apply exclude list tempfiles

☒ Verify backup when complete

☒ Days to retain: 7

The selected Job ID, server and option of when the job will run appears at the top. Note that you may use QuickHelp anywhere on this screen for specific instructions or information on a specific option.

When using the *Desktop* or *Workstation Edition*, no client or server options will appear. Other fields on the screen will be enabled or disabled (grayed-out) depending on whether the option is applicable given the other selections.

If running a *Network Administrator*, you must make one or more selections from the ***Clients*** listbox. The selections will be displayed in the **Name(s)** entry field to indicate the order in which the client backups will be performed. If you want to change the order of the backups, just de-select and re-select the clients in the listbox.

Selecting the Data to Backup

The ***[Data to Backup]*** field description will be one of the following, based on the backup type defined by the selected backup profile:

Volume Group name(s)
Filesystem mount point(s)
Logical Volume(s)
Partition (PP) name(s)
Meta-disk name(s)
Share (resource) name(s)

Files, Directories or @Flist

The data in this field will be filled in automatically from the selected backup profile if provided there. You may change the data to backup by entering one or more options, separated by spaces, in this field. Note that this will not change the original data in the original backup profile. Refer to the Data to Backup in the **Backup Profiles** section for additional information on the contents of this field.

You may press the arrow button to the right of this field to list the options available for the selected backup type. However, since it is undesirable to query every client (if multiple are selected), only options for the first client in the list will be displayed. To list all current filesystems on the client when performing a ***Filesystem*** backup, press the arrow button to display and select from the list. This button does not apply to ***File/Directory*** backups, as the time and resources it takes to display a complete file or directory list would be considerable.

Note that, since the backup job may contain multiple clients, not all of the items in the data list need to apply to all clients. If an item in the list does not exist on any of the clients, it will simply be ignored when the backup is run.

Selecting the Backup Media

Press the arrow button next to the ***[Server/backup] device or directory*** field to select from a list of devices and directories configured for this server. If the backup type from the profile is a *System Backup*, directories configured for both local and client ***System Backups*** will be shown. For all other backup types, the directories configured for other (non-System) backups will be shown. Refer to Configuring Backup Directories for details.

Additional Options

Answers to the following question buttons may be used to override the default actions taken during a backup:

- ***Use alternate server IP/hostname***: This option only appears for *Network Administrators*, and is only enabled if an ***alternate IP/hostname*** was configured for the backup server. To set the alternate IP address or hostname for a server, refer to the server configuration.

  By default, the client will use its default network to reach the server based on the server's hostname and routing information configured on the client. It may at times be desirable for the client to send backup data to the server using a different network than the default. For instance, if there are multiple networks available for reaching the server from the client, or if you wish to offload the heavy backup data traffic onto a different network than other applications are using, you can choose to backup using an alternate network. The alternate network may use a different network adapter on the client, or may route through a different gateway to reach the server. For SP Systems with High-Speed Switch networks, this is particular useful in allowing nodes to backup across the switch network to other nodes. Refer to the SP System Information for additional information.

- ***Delete job after running***: This option is only available when a backup job has been configured to run "*Later*", or once-only. If so, you may also select, using this checkbutton, to have the job configuration removed from the system upon completion of the backup job. This is useful if you are creating temporary backup jobs that are never to be used again.

- ***Perform snapshot backups***: This option is not available for *Desktop Edition* licenses, and only available if snapshot backups have been configured for one or more of the selected clients. By default, all backups are performed using the active (online) copy of a filesystem or logical volume (even when snapshot backups have been configured). To create snapshots of each logical volume before backing it up, check this button.

Refer to Snapshot Backups for details on the configuring filesystems and logical volumes to be backed up using offline mirror copies.

- ***Encrypt data:*** This option is only available if a ***Backup Data Encryption Feature*** license is installed and encryption support is enabled for at least one of the clients selected above. Refer to Enabling Encryption Support in the client configuration to add encryption support for a client. After selecting this button, the entry field to the right will become available. In this field, you must enter the encryption key ID which has been configured on the client. You may not save the job information with this option selected until you have entered the valid name of an encryption key for each selected client.

  For information on configuring encryption keys on the client, refer to Enabling Backup Data Encryption for a Client and the stkeys command.

- ***Apply exclude list***: This option is only available if there is at least one exclude list configured, which applies to at least one of the selected clients. If you select this button, indicating that you wish to use an exclude list, the arrow button to the right will be enabled. You may press the arrow button to select one or more *exclude list name(s)* to use, which will be shown in the box. Click outside the list to complete the selections. To perform the backup without excluding any data, simply un-check this button.

  Note that exclude lists are cumulative, meaning that you can select multiple lists, and the entries in all lists will be combined into a single list when the backup job is performed. Any entries (files, directories, or devices) that do not exist on one or more of the selected clients. If this is the case, that exclude list item will simply be ignored.

- ***Verify backup when complete***: If you want to automatically verify a backup by re-reading the data on the backup media once the backup completes, check this button. Note, however, that an automatic verify will not be performed if you are using a single tape drive or Sequential Autoloader and the backup has spanned more than one tape volume. This is because user-intervention would be required to begin the verify at the first volume. However, if you are using a Random Tape Library, the first tape will be automatically re-inserted into the drive before the verify begins. When a verify process ends (unless you specified to rewind at end of backup in the profile), the tape will be set to the end of the backup data for this job to allow for additional jobs to be appended, if desired.

- ***Days to retain:*** This field will show the word “default”, indicating that the default retention period will be used for this backup. The default retention period may be configured using the Backup Retention Policy option. If you enter a number in this field, it will represent the minimum number of days this backup must be retained before it may be overwritten by another backup. For instance, if you have jobs scheduled to run only once a month, you might want to force a 27-day retention period to protect them from accidental overwriting, while another backup, performed weekly, needs only a 6-day retention period.

  Note that this will NOT allow automatic expiration and overwriting of the backup job if the global policy is set to never allow automatic expiration when overwriting. Note also that an entry in this fields does NOT mean that the backup will be automatically expired after the specified number of days, only that it may be automatically expired if you attempt to write another backup to the beginning of this media.

  For backup jobs written to disk, this will prevent the backup disk files from being removed should the same backup job be run at an earlier time.

Selecting/Customizing the Backup Profile

You must assign a backup profile to the job. The profile will determine the type of backup to be performed as well as the specific backup options which apply to the backup type. Refer to the Data to Backup in the **Backup Profiles** section for additional information. After selecting a profile, the ***Data to Backup*** and ***User Backup Description*** fields will be filled in automatically from the profile information. You may override the profile data by simply changing the information in those fields. This will not change the information in the original profile.

If you want to change any of the default backup settings from the profile, you may select the View/Customize button. This will display the profile options screen and allow you to make any changes that will apply only to this job. You may use this option, for instance, to set the tape to be rewound and ejected at the end of this job even though other jobs that use this profile will not rewind or eject the tape. You can also use this option, for example, to change only the incremental backup level, so that all incremental backups, even those at different levels, can use a single backup profile.

> NOTE **When selecting an SMB (Windows) Shares, the client listbox will be populated only with SMB (Windows) Clients. Likewise, if any other backup profile type is selected, the client listbox will be populated only with Linux or AIX clients.**

Scheduling the Backup

The Backup Schedule box to the right of the screen contains entry fields for backups that are to be scheduled. You need to indicate in the section when the backup should be performed:

1. **Upon Demand –** Selecting this option will save the job information but only run when you choose to do so manually. When selecting this options, all other options in this box will be disabled.

2. **Later -** The job will be run only once at a specified date and time. You will need to enter in the remaining fields a single date and time the backup should run.

3. **Regularly -** The job will be scheduled to run on a regular basis on specific days and times. You may enter multiple options in each of the date and time fields to have the backup run multiple days per week, only on certain days of the week, or even multiple times in a single day. When this option is selected, you may also press the Exceptions button to specify certain days, contrary to your backup schedule, on which the backup should NOT run. Refer to Configuring Backup Exceptions section for more details.

If you set the backup to run only "**Upon Demand**", all other fields in this section will be grayed out and no entries will be accepted. Otherwise, you must enter information into these fields indicating when the backup is to be run. The easiest way to enter the data into these fields is by pressing the arrow to the right of each field and selecting from the popup list.

If the backup is to run "**Later**", only one option may be selected from each list.

If the backup is to run "**Regularly**", more than one option may be selected in each field, and there will be an "*all*" option at the top of the ***Month*** and ***Day of Month*** fields, and an "*any*" option will appear for ***Days of Week*** field. Selecting "*all*" in both the month and day of month fields indicates the job should run on all days of all months. Select "*any*" for the day of week field to indicate that the job should run on any day of the week. Otherwise, the job will run only on the days of week indicated. Note that, if you make an entry in the ***Days of Week*** field and the ***Days of Month*** field is not set to "all", then the job will be run on the specified days of the month **only** if they occur on the specified days of the week.

Changing a Backup Job

To change information for an existing backup job, either:

1. Select Configure→Backup Jobs from the menu bar, then select the Job ID from the listbox and press the Add/Update button, or

2. If the Job Information is displayed on the Main Screen, select the icon for the job to change and press the Change Job button at the bottom of the screen.

The job options screen will appear. Make all desired changes to the information on the screen, then press the Save button to save the changes and close the window.

Removing a Backup Job

To remove a backup job, the job may not currently be in a job queue. A job will only be in a job queue if it is currently running, waiting to be run, has been placed on hold, or had previously failed.

To remove a backup job, either:

1. Select Configure→Backup Jobs from the menu bar, then select the Job ID from the listbox and press the Remove button, or

2. If the Job Information is displayed on the Main Screen, select the icon for the job to remove and press the Remove Job button at the bottom of the screen.

Running a Backup Job on Demand

Any backup job, whether it is currently scheduled or not, may be run at any time. There are several ways to start a job running:

1. Select Configure→Backup Jobs or Actions→Run a Backup Job from the menu bar, then select the Job ID from the listbox and press the Run Now button, or

2. If the Job Information is displayed on the Main Screen, select the icon for the job to run and press the Run Job button at the bottom of the screen, or

3. If the job is currently at the top of a job queue but is not running because it had previously failed or was placed on hold, display the job queues on the Main Screen, select the queue in which the job is placed, and then press the Start Job button.

For the first two options, "running" the job actually just places the job in the job queue. If there are no other jobs in the same queue, the job will start running immediately. When a job is added to the queue, it will be run immediately if there are no other jobs queued to the same device on the same server (except that disk file backups on a server may run simultaneously). If another job is running to the same device, this job will be placed in a “Pending” state until the prior job finishes. If a prior job had failed, it will remain in the queue and block other jobs from starting. The failed job must therefore be either restarted or removed from the queue to allow jobs behind it to start.

Adding a Job to the Queue from the Command Line

Even if the Backup Administrator user interface is not running, scheduled jobs will automatically be placed in the queue at their scheduled times, and the queues will be processed and jobs in each queue will be run on a first-come first-serve basis. It is also possible to manually add jobs to the queue without using the Backup Administrator interface. To add a job to the queue, refer to the stqueue command.

Running a Backup Job from the Command Line

It is possible to run a backup job from the command line, bypassing the job queues, by using the strunjob command (refer to the strunjob command syntax). The Backup Administrator user interface need not be running. Note that the job will start immediately and may interfere with other jobs writing to the same devices since the queues are not used. If you wish to add the job to the queue from the command line, so that it will run only when the backup server and devices are available, refer to the section Adding a Job to the Queue from the Command Line.

12. Holidays

The features described in this chapter are not available when using Desktop Edition.

There may be days of the year, even days of the month, that you don't want any of your backup jobs to run. On holidays, for instance, there may have been no activity on the system, and there may not be anyone available to insert new backup cartridges in the tape drive. You may schedule ***Backup Exceptions*** or "Holidays" from performing backups.

There are actually two ways to do this, on a job-by-job basis, or for all backup jobs. To set exceptions for all backup jobs, select Configure→Holidays from the menu bar. To set exceptions for a specific job, press the Exceptions button in the **Backup Schedule** section of the Job Configuration Screen. The respective screens will be displayed as follows:

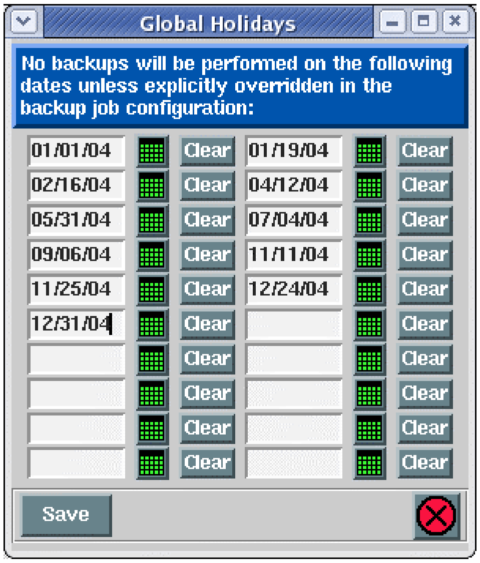
Global Holidays

No backups will be performed on the following dates unless explicitly overridden in the backup job configuration:

01/01/04	01/19/04
02/16/04	04/12/04
05/31/04	07/04/04
09/06/04	11/11/04
11/25/04	12/24/04
12/31/04	

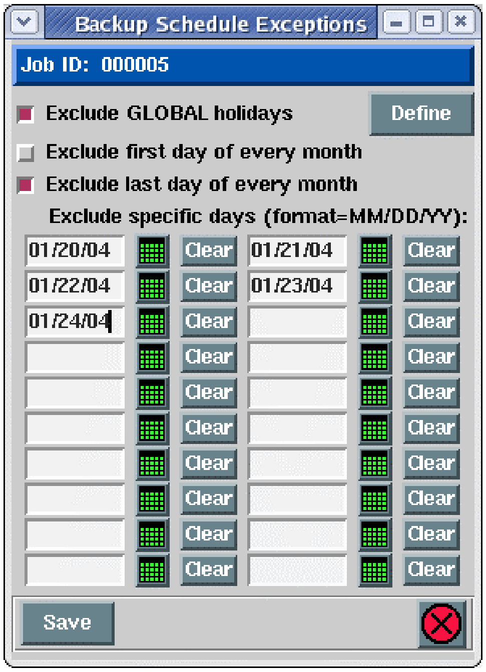
Backup Schedule Exceptions

Job ID: 000005

■ Exclude GLOBAL holidays

□ Exclude first day of every month

■ Exclude last day of every month

Exclude specific days (format=MM/DD/YY):

01/20/04	01/21/04
01/22/04	01/23/04
01/24/04	

In the first screen, you may enter one or more dates on which ALL backup jobs will be excluded from running. In the second screen, you may enter additional dates in the date fields that will be excluded for this job in addition to those excluded on global holidays. Since it may be cumbersome to enter the dates by hand, and since the dates may be dependent on the day of the week, it is useful to have a calendar handy. By pressing the calendar icon next to each date field, a calendar will appear such as the following:

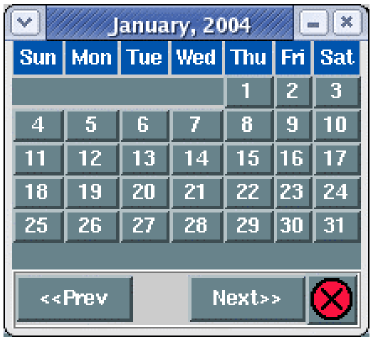
January, 2004

Sun	Mon	Tue	Wed	Thu	Fri	Sat
				1	2	3
4	5	6	7	8	9	10
11	12	13	14	15	16	17
18	19	20	21	22	23	24
25	26	27	28	29	30	31

If you press a specific date on the calendar, that date will be automatically inserted into the date field and this window will close. You may press either the "<<PREV" and "NEXT>>" buttons to change the calendar to the previous or next month and select a date from that calendar.

In the ***Backup Schedule Exceptions*** (by job) window, the following options may also be selected:

1. **Exclude GLOBAL holidays**: This box is always checked by default, meaning that global holidays apply to this job as well. If you un-check this box, then the global holidays which apply to other jobs will not apply to this job, and the job will therefore run on those holidays if the job schedule permits. You may press the Define button next to this field to bring up the ***Global Holidays*** window and make changes to the global holidays if desired.

2. **Exclude the fist day of every month**: Since there may be a monthly backup set to run on the first day of the month, you may not need this job to run on the same day. If not, select this box, which is the same as adding the first day of every month in the date fields.

3. **Exclude the last day of every month**: Since there may be a monthly backup set to run on the last day of the month, you may not need this job to run on the same day. If not, select this box, which is the same as adding the last day of every month in the date fields.

When all selections and entries have been made, press the Save button to save the dates and options. The backup job (or jobs) will no longer run on the specified dates.

13. Configuring Snapshot Backups

The options described in this chapter are not supported when using Desktop Edition. Snapshot feature is available only for *mirrored* logical volumes on *AIX*, and only for data contained in LVM logical volumes on *Linux* systems.

SBAdmin provides an option of creating a "*point-in-time*" backup of data contained in logical volumes. This is typically referred to as a **snapshot backup**. Although the feature is available for both AIX and Linux systems, the process differs to some extent:

- For ***Linux***, an LVM *snapshot logical volume* is created for each logical volume to be backed up. This snapshot LV is generally smaller than the original LV, but large enough to contain any changes which occur to the original logical volume for the duration of the backup. As the backup is performed, original data to be changed by another process is first copied to the snapshot LV, and the data from the snapshot LV is backed up in place of the changed data. When the backup is complete, the snapshot LV is simply removed. Any process which reads or writes data to the logical volume (or filesystem within) during the backup will use the most up-to-date data, while the backup contains only the original data as it was when the backup began.

- For ***AIX***, snapshots are accomplished by splitting off a copy of a *mirrored* logical volume so that this *offline* copy may be used to perform a backup independent of the primary copy. This allows a backup of the data at a point-in-time while other applications continue to update the primary copy. When the backup is complete, the offline copy is put back online and only the *physical partitions* that became *stale* (because primary copy was changed) are updated, or *synchronized*, the primary copy.

Any logical volume on the system (containing any type of data including filesystem data) may be backed up as a snapshot. However, for ***AIX*** systems, the logical volume must first have been *mirrored*. Logical volume mirroring is explained in detail in the AIX system management documentation.

IMPORTANT NOTES FOR AIX SYSTEMS:

1. **If logical volumes have been mirrored for increased availability, then availability is no longer ensured when a single mirrored copy is taken offline for backups (since only one active copy remains). It is therefore recommended, but not required, that logical volumes use two mirrored copies (3 copies total) so that the online data retains mirrored availability when one copy is taken offline for backups.**
2. **Splitting off a mirrored copy of a logical volume can take several seconds, depending on the size of the logical volume. During that time users may create, delete and extend files within a filesystem. Those types of transactions require changes to the filesystem *metadata*, some of which may be only partially completed after the mirrored copy is taken offline. As the filesystem integrity of the offline copy cannot be guaranteed, this may result in filesystem errors during the backup, but will have no affect on the remaining online copy(s). Although unlikely, there is a chance of unexpected backup or system errors caused by reading data from an inconsistent filesystem. It is therefore recommended, when possible, that any applications which modify filesystem data be temporarily suspended while the mirrored copies are taken offline.**
3. **Should the system fail for any reason (such as a power loss) while a split-mirror backup is in process, the mirrored copies will remain offline after the system is rebooted. SBAdmin records information about mirrors that have been taken offline during a backup in order to rebuild and resync the mirrors in case of a system failure. Should the system fail during a backup, check to be sure that all logical volumes are in sync (the "`lsvg –l VGname`" command will show "stale" if they are not). If mirrors are stale, use the utility Resync Split-Mirrors after a System Failure to bring the mirrors back online.**

Enabling Snapshot Backups

Snapshot backups are configured on each client on which a backup will be performed. On each client, you may specify each logical volume for which a snapshot is created, or you may indicate that all logical volumes will use snapshot backups, when possible.

Although logical volumes and filesystems may be configured to allow snapshot backups, a snapshot will not be created by default when a backup job is run. You must also select to *Perform Snapshot Backup?* from the job configuration screen before a snapshot backup will be performed.

Options for configuring snapshot backups differ slightly between ***AIX*** and ***Linux*** systems. Therefore, select one of the following from the menu bar, depending on the client system type (or local system type if not a *Network Administrator*):

Configure→Snapshot Backups→Split-Mirror Backups (AIX)→Configure Split-Mirror Backups

Configure→Snapshot Backups→Snapshot LV Backups (Linux)

The configuration screen will appear as in the following samples:

For Linux

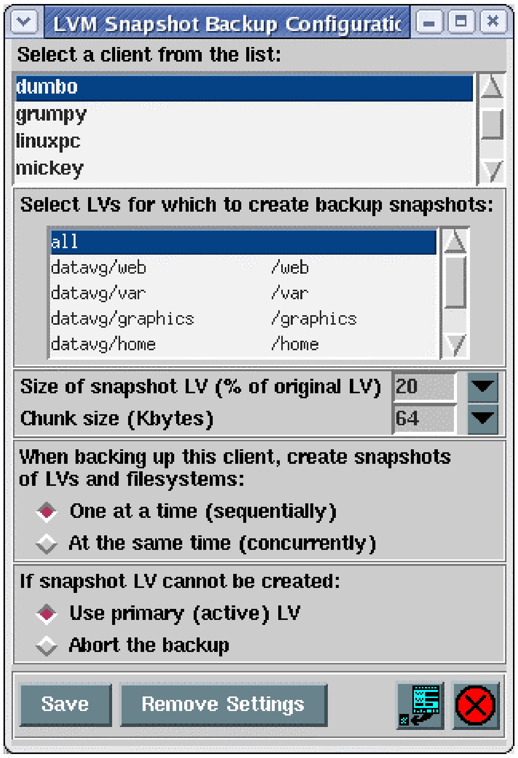

LVM Snapshot Backup Configuratic

Select a client from the list:

- dumbo
- grumpy
- linuxpc
- mickey

Select LVs for which to create backup snapshots:

```
all
datavg/web          /web
datavg/var          /var
datavg/graphics     /graphics
datavg/home         /home
```

Size of snapshot LV (% of original LV) 20

Chunk size (Kbytes) 64

When backing up this client, create snapshots of LVs and filesystems:

● One at a time (sequentially)
○ At the same time (concurrently)

If snapshot LV cannot be created:

● Use primary (active) LV
○ Abort the backup

For AIX

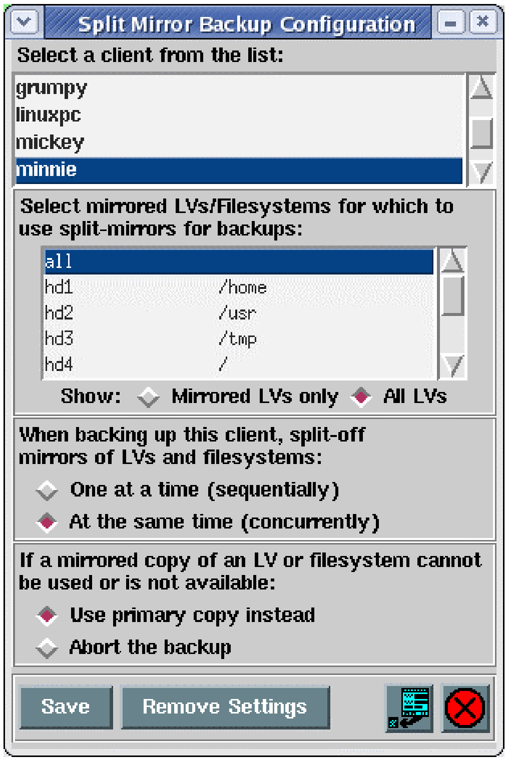

Split Mirror Backup Configuration

Select a client from the list:

- grumpy
- linuxpc
- mickey
- minnie

Select mirrored LVs/Filesystems for which to use split-mirrors for backups:

```
all
hd1          /home
hd2          /usr
hd3          /tmp
hd4          /
```

Show: ○ Mirrored LVs only ● All LVs

When backing up this client, split-off mirrors of LVs and filesystems:

○ One at a time (sequentially)
● At the same time (concurrently)

If a mirrored copy of an LV or filesystem cannot be used or is not available:

● Use primary copy instead
○ Abort the backup

If using *Workstation or Desktop Edition*, the **Client** listbox does not appear. If using a *Network Administrator*, you must select the client for which to configure snapshots in the first listbox. When doing so, a list of logical volumes on that system will be displayed in the second listbox.

For AIX systems, the logical volume list will contain only mirrored logical volumes. If you wish to expand this list to include both mirrored and non-mirrored logical volumes on the system, press the button labeled Show: All LVs. Logical volumes that are mirrored will be shown with an asterisk (*) to the right of the logical volume name.

You must either select the individual logical volumes for which snapshots may be created, or select “***all***”. If “all” is selected, snapshots will be permitted for all logical volumes (or all mirrored logical volumes on ***AIX***).

> NOTE **Again, you must also select to *Perform Snapshot Backup?* from the job configuration screen before a snapshot of a logical volume or filesystem will be created.**

The options which follow indicate the action the backup process should take when performing snapshot backups:

1. **Size of Snapshot LV (% of original LV)** :

   This option is applicable to ***Linux*** systems only. Indicate the size of the snapshot as a percentage of the original LV by using the arrow button to list and select a percentage. The minimum size will depend on the amount of data that is changed within the original LV while the backup is in progress. It is very important to create the snapshot large enough that it does not run out of space, as this will result in a failure of the entire backup.

2. **Chunk size (Kbytes)** :

   This option is applicable to ***Linux*** systems only. Use the arrow to the right of the entry field to list and select from a valid chunk size. Valid sizes are from 4 Kbytes to 1024 Kbytes (1 megabyte).

   A “chunk” is the unit in which the original logical volume will be divided when tracking changes to the LV when a snapshot is used. Each time a chunk is changed for the first time, the original chunk is copied to the snapshot LV in its entirety, then referred to in place of the original by the backup process.

   When determining the best chunk size to use, there is a trade-off: The larger the chunk, the fewer writes to the original LV it will take to fill up the snapshot LV (since larger chunks of data must be copied, even when only a small piece of data is changed). The smaller the chunk, the more individual copies must occur as the original data is changed, which may have a greater impact on system performance during the backup. The default of 64 Kbytes is sufficient for most purposes.

3. **When backing up this client, create snapshots (*split-off mirrors*) of LVs and filesystems**:

   a. ***One at a time (sequentially)***. Select this option if a snapshot should be created individually when the data in that logical volume (or filesystem) is to be backed up. When the backup of this LV completes, the snapshot is removed (*resyncd*). This option is recommended if there is no relational data between different logical volumes and filesystems that must be backed up at the same *point-in-time*.

      - For ***AIX***, splitting off one logical volume copy at a time increases the availability of the data in case of a hardware failure since all other mirrors remain intact.

      - For ***Linux***, less disk space is required since only one snapshot LV is created at a time. Also, when creating and removing snapshots one at a time, the snapshot exists for a lesser time, reducing the amount of data written to it, thereby decreasing the possibility of running out of space in the snapshot.

   b. ***At the same time (concurrently)***. Select this option if a snapshot of all logical volumes to be included in a backup should be created at the same time. This is important if there is relational data between different logical volumes and filesystems that require that the data from all logical volumes be backed up from the same *point-in-time*. For the reasons described above, this option is not recommended if there is no relational data between different logical volumes.

4. **If a snapshot cannot be created (*mirrored copy cannot be used*):**

   Indicate the action which should be taken if the snapshot backup cannot be performed for a specific logical volume:

a. ***Use primary (active/copy) instead.*** If this option is selected, then the failure to create a snapshot of the logical volume or filesystem will result in the backup using the original (online) copy without a snapshot. The result would be the same as if snapshot backups were not configured for this logical volume or filesystem.

b. ***Abort the backup.*** Select this option if the client backup should abort when a snapshot cannot be created.

   If **concurrent** snapshot backup is performed, all snapshots will be removed (***Linux***) or mirrors will be resynchronized (***AIX***), and the backup of the client will terminate, but the job will continue processing other client backups, if any.

   If **sequential** snapshot backups are performed, no snapshots will exist at this point, but there may have already been some data written to the backup media. Therefore, both the backup and job will terminated, preventing other backups from continuing to write to the backup media.

Possible issues preventing a logical volume or filesystem snapshot from being created include:

For AIX:

1) Another snapshot backup already has a logical volume copy offline.

2) A logical volume copy was previously split off during a prior snapshot backup and a system failure had occurred. If this is the case, you should use the option Resync Split-Mirrors after a System Failure to bring the copies back online.

3) You specified specific logical volumes when configuring split-mirror backups, but a selected logical volume has only one copy (not mirrored).

4) The logical volume mirrors are not in sync. This is an LVM state that indicates a mirror was likely not updated due to a disk failure and was later recovered. If this is the case, you will need to use the AIX "***syncvg***" command to resync the mirrored copies before a split-mirror backup may be performed.

For Linux:

1) A snapshot LV already exists for the logical volume. Another snapshot backup may not have removed the snapshot due to a program failure, or another (non-SBAdmin) process may have created a snapshot LV.

2) There may not be enough space in the volume group to create the snapshot logical volume. If this is the case, you need to either expand the volume group, remove other unused logical volumes, or select to create smaller snapshots using the Size of Snapshot LV option.

When all selections have been made, press the Save button. The settings for the selected client will be saved and you may then select a different client for which to configure snapshot backups.

If you wish to remove prior settings for a client, select the Remove Settings button. If selected, the prior configuration will be removed and no backups performed on that client will use snapshot backups, even if the backup job configuration indicates that snapshots should be used.

Resync Split-Mirrors After a System Failure

Note **This option is only applicable to *AIX* systems.**

On ***AIX*** systems, SBAdmin tracks at all times the logical volumes and filesystems that have been taken offline for *snapshot* backups. Should the system fail for any reason, such as a power outage, the mirrored logical volumes that had been split-off at the time of the failure will remain offline after the system is rebooted. Also, the LVM information in the device configuration database will usually be inconsistent with the LVM information on the disks. This process will make the LVM data consistent, then re-join and resynchronize all mirrored copies.

To re-join and resync offline mirrored copies, select the option Configure→Snapshot Backups→Split-Mirror Backups (AIX)→Resync Split-Mirrors After a System Failure. When doing so, the following message will appear:

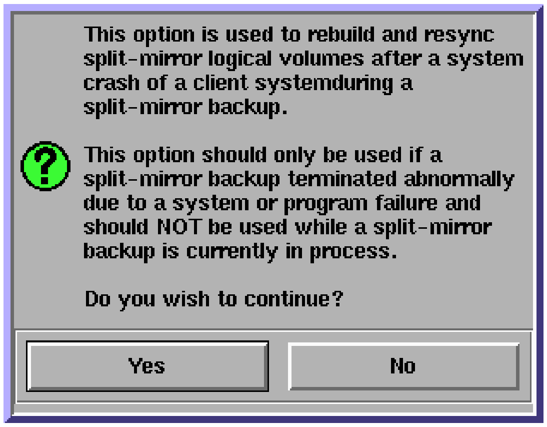

This option is used to rebuild and resync split-mirror logical volumes after a system crash of a client systemduring a split-mirror backup.

This option should only be used if a split-mirror backup terminated abnormally due to a system or program failure and should NOT be used while a split-mirror backup is currently in process.

Do you wish to continue?

As stated, this option should not be used if a snapshot backup is currently in process. If you wish to continue, press the Yes button.

You must then select client for which to resynchronize split-mirrors from the list provided. After doing so, another window similar to the following will appear which shows the progress of the resync procedure:

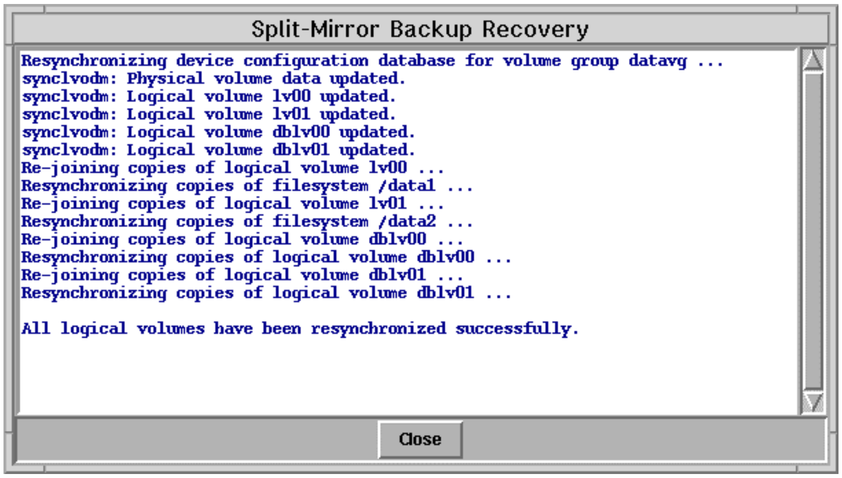

Split-Mirror Backup Recovery

```
Resynchronizing device configuration database for volume group datavg ...
synclvodm: Physical volume data updated.
synclvodm: Logical volume lv00 updated.
synclvodm: Logical volume lv01 updated.
synclvodm: Logical volume dblv00 updated.
synclvodm: Logical volume dblv01 updated.
Re-joining copies of logical volume lv00 ...
Resynchronizing copies of filesystem /data1 ...
Re-joining copies of logical volume lv01 ...
Resynchronizing copies of filesystem /data2 ...
Re-joining copies of logical volume dblv00 ...
Resynchronizing copies of logical volume dblv00 ...
Re-joining copies of logical volume dblv01 ...
Resynchronizing copies of logical volume dblv01 ...

All logical volumes have been resynchronized successfully.
```

Some errors from the ***synclvodm*** command may appear in the window. These can usually be ignored as they indicate the LVM data in the device configuration database is inconsistent with the volume group information on the disks. They may be ignored since these inconsistencies are repaired by this process.

When this process is complete, you may press the Close button to close this window.

14. Job Queues

When jobs are run, the are actually placed in a ***job queue***. A job queue will exist for each device on each server, and a "disk" queue will exist for each server with backups directories defined. The job queues are used to prevent multiple jobs from attempting to write to the same device at the same time. The jobs in a particular queue will be run in the order in which they were placed in the queue.

The Job Queue Display

Job queues may only be displayed or manipulated from the Main Screen. The following is an example of the **Job Queue Display**, which may be shown by selecting Display→Job Queues from the menu bar:

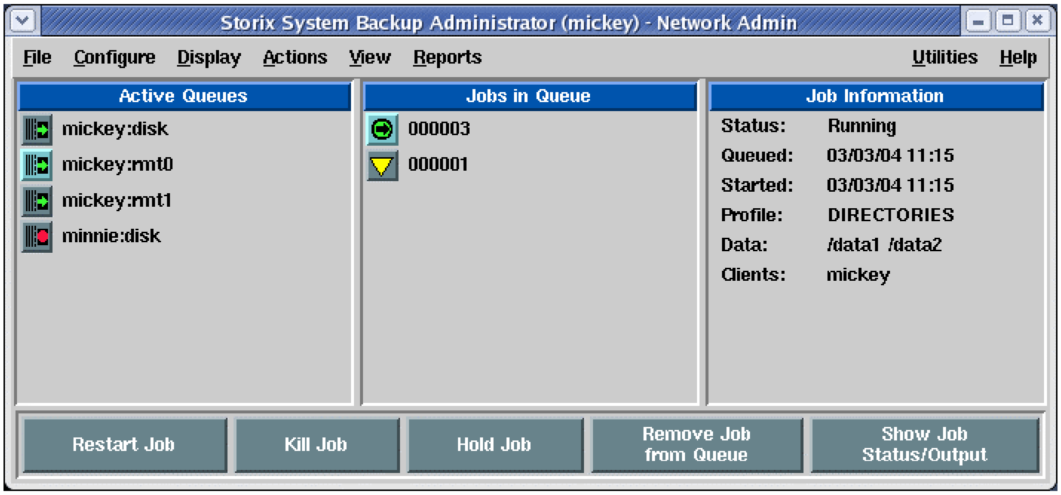

mickey:disk
mickey:rmt0
mickey:rmt1
minnie:disk

000003
000001

Status: Running
Queued: 03/03/04 11:15
Started: 03/03/04 11:15
Profile: DIRECTORIES
Data: /data1 /data2
Clients: mickey

The left-most display area contains the names of the job queues *for which at least one job exists*. If there are no jobs either running or stopped in a queue, the queue will not be displayed. The name of the queue contains the server and device (or "disk") if the *Network Administrator* is used, or just the device name if using *Workstation* or *Desktop Edition*. To display the jobs within a queue, select the icon corresponding to the desired queue. When doing so, the selected queue will be highlighted in blue and the jobs in that queue will be listed in the center display area.

The center area contains the jobs currently in the queue. The jobs are place, and will be run, in the order they were added to the queue. To show a summary of the job information for a job, click on the icon corresponding to the Job ID. The job information will appear in the display area to the right and the selected job icon will be highlighted in blue.

The action buttons at the bottom of the screen will apply to the selected queue or selected job.

Icons on the Job Queue Display

The icons for the queues and jobs display a symbol representing the status of the queue or job. The following is a list of possible status icons that may appear:

🟩 A queue in which a job is currently running

🟥 A queue in which a job has failed (click on queue and job icons to see why)

- ➡ A job that is currently running
- ▼ A pending job (waiting for a prior job to complete)
- 🚫 A job that has failed (click on icon to display job information)
- ✋ A job placed on hold by the user

The status of a queue or queued job is checked every few seconds and the icons are automatically updated with the new status, if changed. When a job has completed successfully, the icon for the job is removed from the screen. Once the last job in a queue has completed successfully, the queue icon is also removed from the screen.

Monitoring Backups

The Job Message Screen

Since many jobs run automatically after being scheduled to run at a certain date and time, there may not always be a person watching the screen when a job is started. Therefore, as jobs are run, the queuing system keeps an updated list of messages on the screen, showing which jobs have run, which have completed, and which jobs have failed (and why). The following is an example of this **job status message screen**:

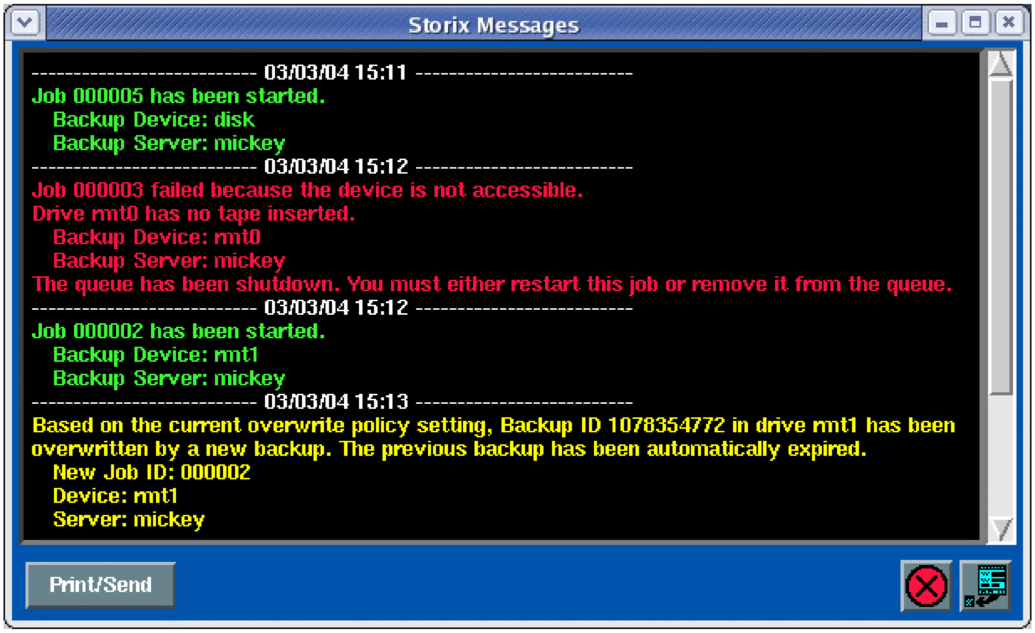

```
--------------------------- 03/03/04 15:11 -------------------------
Job 000005 has been started.
   Backup Device: disk
   Backup Server: mickey
--------------------------- 03/03/04 15:12 -------------------------
Job 000003 failed because the device is not accessible.
Drive rmt0 has no tape inserted.
   Backup Device: rmt0
   Backup Server: mickey
The queue has been shutdown. You must either restart this job or remove it from the queue.
--------------------------- 03/03/04 15:12 -------------------------
Job 000002 has been started.
   Backup Device: rmt1
   Backup Server: mickey
--------------------------- 03/03/04 15:13 -------------------------
Based on the current overwrite policy setting, Backup ID 1078354772 in drive rmt1 has been
overwritten by a new backup. The previous backup has been automatically expired.
   New Job ID: 000002
   Device: rmt1
   Server: mickey
```

A scrollbar is provided to the right to scroll up and down the messages. This screen may only be displayed if the Backup Administrator application is running. If not, mail messages are sent to the ***root user*** when a job has failed, so the system administrator will know to check the job queues to fix the problem and re-run the job.

If a job completes successfully, or if it fails after having begun the backup, a Label button will appear within the text of the message. By pressing this button, the media label will be displayed, which provides a summary of the contents of the media, both for the current job and any prior jobs, if any. Refer to the Backup Labels section for a sample and information on the label contents.

This window appears each time a new message is posted. Once the window is closed (using the cancel button), it will not be possible to view the previous message contents. However, the message window will reappear with any new messages that are posted.

The Backup Status Screen

A detailed status report of a job that is currently running, or one which has failed may be displayed at any time by pressing the Status Report button at the bottom of the Job Queue Display. The status screen for the currently selected job will be displayed such as the following example:

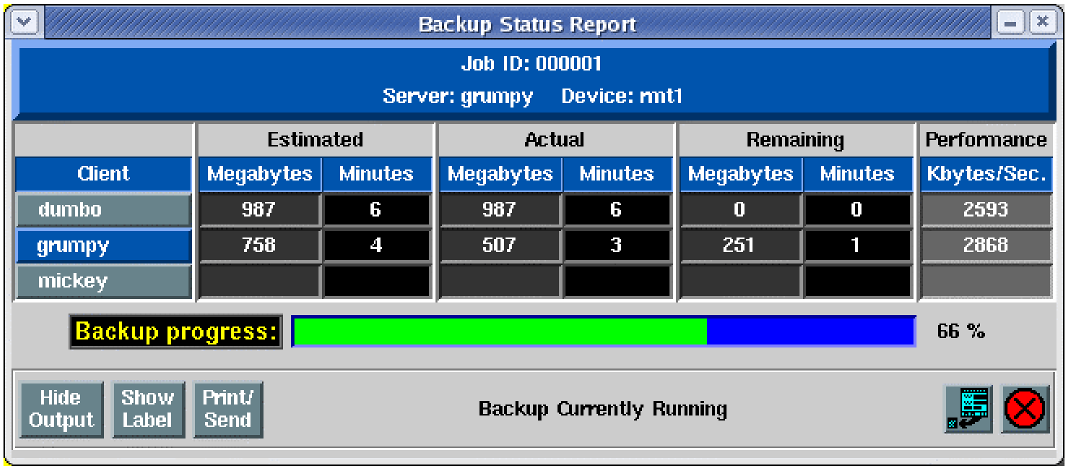

Backup Status Report

Job ID: 000001

Server: grumpy   Device: rmt1

	Estimated		Actual		Remaining		Performance
Client	Megabytes	Minutes	Megabytes	Minutes	Megabytes	Minutes	Kbytes/Sec.
dumbo	987	6	987	6	0	0	2593
grumpy	758	4	507	3	251	1	2868
mickey							

Backup progress: 66 %

Backup Currently Running

The Job ID, server (if *Network Administrator*) and device are shown at the top of the screen. The middle section will contain a set of boxes for **each client** in the job. If using *Desktop or Workstation Edition*, the **Client** column will not appear, and only one progress line will be shown. The corresponding client is indicated in the button at the far left. These client buttons may be used to display the progress bar or backup output for a particular client.

Next to each client button is a list of values, indicating the approximate progress of the backup. This shows the **estimated** time and size of the backup, the **actual** time elapsed and amount of data written so far, and the **remaining** time and data to be written. Note that these values apply to each corresponding client. If a client backup has not yet started, its progress values will not be shown.

The progress bar is seen below the client information and shows a graphical representation of the percent of the backup that has completed. Again, this applies only to the selected client backup. To view the status bar for a different client backup, press the desired client button.

At the bottom of the status screen are more buttons for displaying additional information:

- The Show Output button is described in the Backup Output Display section below.

- The Show Label button will allow you to view the current contents of the media label, which will include only those client backups that have completed, as well as any prior jobs written to the same media, if any. Refer to the Backup Labels section for a sample and information on the label contents.

- The Print/Send button will allow you to send a report to the printer with the contents of this window as well as the Backup Output Display for all clients assigned to this job.

- The Show Verify Status button will only appear if you selected to automatically verify the backup data when the backup completed in the job settings. If the verify was performed, or is in progress, the verify progress is already shown, and the button will appear as Show Backup Status instead. When selected, the progress bar will change from ***Backup progress*** to ***Verify status***, and vice-versa, and the corresponding progress values will be displayed in the section above. If the option to automatically verify the backup was not selected with configuring the backup job, this button will not be shown.

Use the cancel button on the lower right corner to close this window. The information will continue to be updated and may be redisplayed at any time, **even after the backup job has completed**.

The Backup Output Display

The Show Output button on the bottom left corner of the status report screen will display the backup messages for the *selected client*. These might include status messages, warnings or error messages. Any time a backup job fails after the backup begins, this is the button you would use to find out why. The following is a sample output screen:

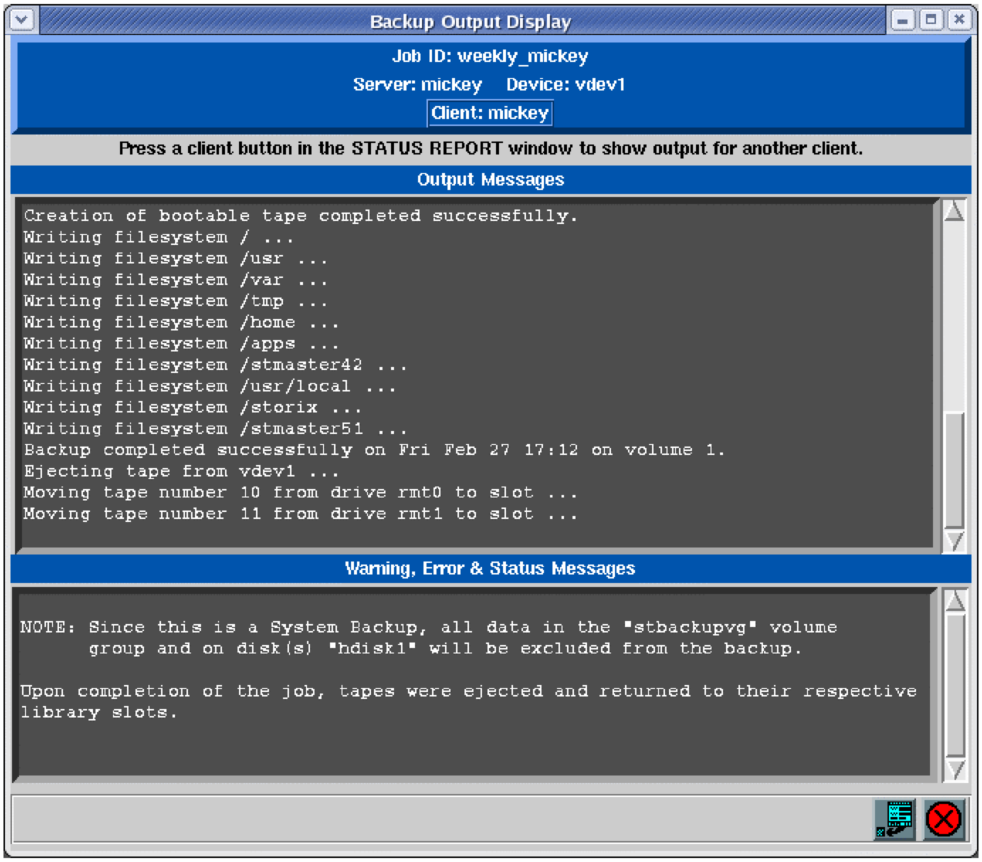

Backup Output Display

Job ID: weekly_mickey

Server: mickey Device: vdev1

Client: mickey

Press a client button in the STATUS REPORT window to show output for another client.

Output Messages

```
Creation of bootable tape completed successfully.
Writing filesystem / ...
Writing filesystem /usr ...
Writing filesystem /var ...
Writing filesystem /tmp ...
Writing filesystem /home ...
Writing filesystem /apps ...
Writing filesystem /stmaster42 ...
Writing filesystem /usr/local ...
Writing filesystem /storix ...
Writing filesystem /stmaster51 ...
Backup completed successfully on Fri Feb 27 17:12 on volume 1.
Ejecting tape from vdev1 ...
Moving tape number 10 from drive rmt0 to slot ...
Moving tape number 11 from drive rmt1 to slot ...
```

Warning, Error & Status Messages

```
NOTE: Since this is a System Backup, all data in the "stbackupvg" volume
      group and on disk(s) "hdisk1" will be excluded from the backup.

Upon completion of the job, tapes were ejected and returned to their respective
library slots.
```

Scrollbars are provided to the right of each display panel in case the output exceeds the size of the panel. The Job ID, server, device, and client are shown at the top of the screen. To view the backup output for a different client, select the desired client button on the status report screen. You may press the cancel button at the bottom to close this window. It may be redisplayed at any time, **even after the backup has completed**.

Manipulating Backup Jobs

To change the status of a job in the queue, you must select the queue and job on the Job Queue Display screen. The action buttons at the bottom of the screen then apply to the selected jobs. The following functions may be performed:

Kill a Running Job

Jobs that are currently ***running*** may be killed, or *canceled*, by selecting the Kill Job button from the Main Screen when the Job Queues are displayed. A signal is sent to the job telling it to terminate. Depending on the current backup operation being performed, this may take a little time. Once the job has been killed, a message will appear on the Job Message Screen indicating that the backup has been terminated.

> NOTE **If the backup was being performed to tape, the tape will be rewound after a job is killed. This is necessary to prevent any future jobs from being appended to the same tape. The tape should be removed from the drive immediately if there are prior successful backup jobs on the same tape that need to be preserved.**

Place a Job on Hold

A job which is currently in the ***pending*** state may be placed on hold by pressing the Hold Job button from the Main Screen when the Job Queues are displayed. When a job is placed on hold, it will not run when any prior backups complete, but will remain in the queue waiting to be manually started.

Restart a Job

A job which is either on ***hold***, had previously ***failed***, or had been ***killed***, may be started, or restarted, by pressing the Start Job button from the Main Screen when the Job Queues are displayed. Jobs that are restarted after they have failed or had been killed will restart from the beginning of the job, even if one or more of the client backups had completed.

Remove a Job from the Queue

Any job, except a running job, may be removed from the queue by pressing the Remove Job from Queue button from the Main Screen when the Job Queues are displayed. After doing so, the selected job is removed from the queue and its icon will disappear. If this was the last remaining job in the queue, the queue icon will disappear as well.

> NOTE **Removing a job from the queue does not delete the job itself. The job will remain on file and can be scheduled or run manually at another time.**

15. Backup Labels

A backup label is generated for each backup that is started at the beginning of a tape as well as for any backups stored to disk files. These labels are used to keep track of the contents of the backup for use when verifying or restoring data at a later time. The backup label contains a summary of the contents of the backup media, which may include multiple backup jobs and multiple client backups (if *Network Administrator*) within each job. Also, for each backup, status information is recorded, including the backup time, size of the backup and the output of the backup commands. This backup information is kept on file for as long as the backup label is also available.

Note that the backup media may contain multiple tape volumes. If a new backup job or multiple client backups within a job are appended to an existing backup tape, that backup information is appended to the same backup label.

Backup labels are not the same as Tape Labels. A tape label is a unique identifier assigned to each individual tape, allowing the backup label information to be obtained given a tape label id. The tape label IDs for tapes used within a backup are also shown in the backup label. Note, however, that tape labels must be placed on the tape before they are used in a backup. Refer to the option Write a Tape Label ID to a Tape in the Utilities section for details on tape labels.

The following is an example of a backup label for a tape containing multiple backup jobs, each containing multiple client backups:

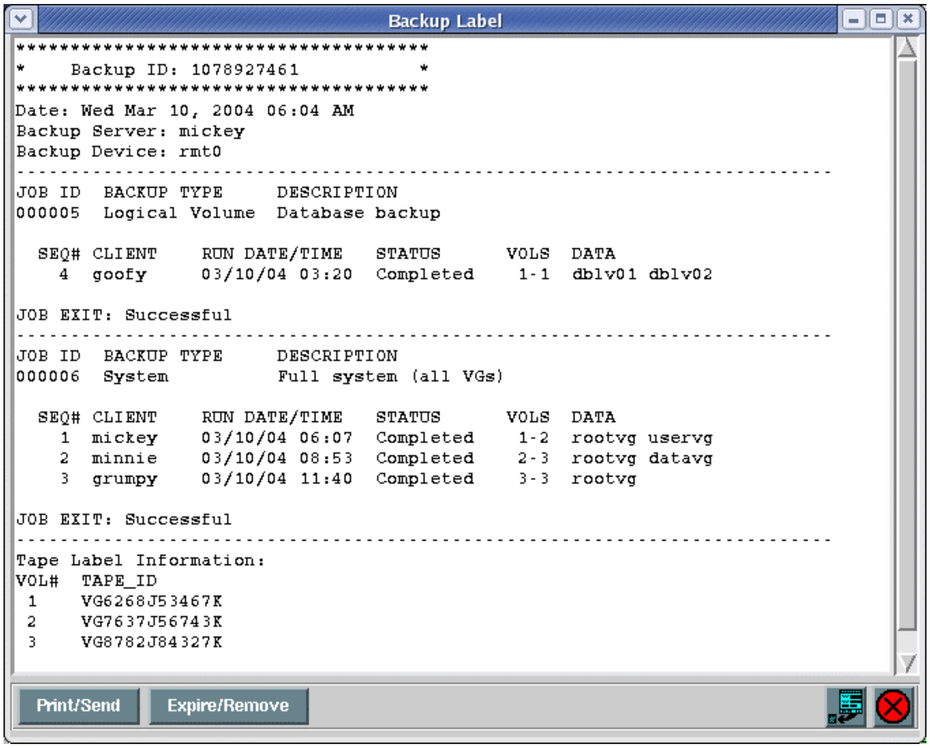

```
**************************************
*    Backup ID: 1078927461           *
**************************************
Date: Wed Mar 10, 2004 06:04 AM
Backup Server: mickey
Backup Device: rmt0
---------------------------------------------------------------------------
JOB ID  BACKUP TYPE     DESCRIPTION
000005  Logical Volume  Database backup

  SEQ# CLIENT    RUN DATE/TIME   STATUS      VOLS  DATA
    4  goofy     03/10/04 03:20  Completed    1-1  dblv01 dblv02

JOB EXIT: Successful
---------------------------------------------------------------------------
JOB ID  BACKUP TYPE     DESCRIPTION
000006  System          Full system (all VGs)

  SEQ# CLIENT    RUN DATE/TIME   STATUS      VOLS  DATA
    1  mickey    03/10/04 06:07  Completed    1-2  rootvg uservg
    2  minnie    03/10/04 08:53  Completed    2-3  rootvg datavg
    3  grumpy    03/10/04 11:40  Completed    3-3  rootvg

JOB EXIT: Successful
---------------------------------------------------------------------------
Tape Label Information:
VOL#  TAPE_ID
 1    VG6268J53467K
 2    VG7637J56743K
 3    VG8782J84327K
```

The Backup ID appears at the top. This ID is a unique identifier generated automatically for each label and is also stored on the backup media itself. This way, it is possible to read the Backup ID from the backup media and reference its contents in the label information. Also at the top of the label is the date the label was first created, and the server and device the backups were written to.

The above backup label contains two backup jobs. Job *000005* contains a backup of database logical volumes *dblv01* and *dblv02* on client *dumbo*. Job *000006* contains a full system backup of clients *dumbo*, *mickey* and *minnie*. This backup label, therefore, contains a total of four backups, stacked together onto three tape volumes.

The Tape Label ID for each volume is shown at the bottom. The tape label ids will be shown in the backup label if a previous backup containing tape labels overwritten by this backup, or if the option Write a Tape Label ID to Tape was used prior to writing this set of backups.

Use the Print button to send a copy of the backup label to the printer. You will always know the contents of the tape without reading it if you have a copy of the label with each backup tape.

The Expire/Remove button is used to expire, or remove, the backup label from the system. This should be done only when the tape will be discarded or reused. Refer to Expiring a Backup Label below for details.

Automatically Printing Backup Labels

After a backup job completes, the backup label created or associated with that job may be automatically sent to any printer queue configured on the admin system. This may be accomplished by setting an option in the Backup Profile configuration for the profile assigned to the job. Note that you must have configured the printer queue in AIX before using this option.

To print backup labels upon completion of a backup job, follow these steps:

1. Select Configure→Backup Profiles from the menu bar.
2. Select the profile name to change, then press the Add/Change button.
3. For ***Print Backup Label upon completion***, press the button to indicate "**Yes**".
4. Next to the ***Print queue*** field, press the down-arrow button to list and select a printer queue.

> NoTE **If you want to print only the backup labels for certain backup jobs, you may also customize the backup profile for a job instead of setting a printer queue for all jobs using the profile. Refer to Selecting/Customizing a Backup Profile in the Job Configuration section for details.**

View Backup Labels

Because it is often desirable to view a backup label, there are many places within the application where the backup label may be displayed:

1. A label for any completed backup may be displayed at any time by selecting View→Backup Labels from the menu bar on the Main Screen. Since there are many ways to search for the desired label, this option is explained in detail below.
2. When a backup completes or fails, a message is displayed in the job message screen. If this screen is not already visible, it will be displayed automatically any time a job message is posted. If the backup job completed successfully or failed after the backup had started, a Label button will appear on the message screen. When pressed, the label for the tape containing the backup is displayed on the screen.
3. When displaying status of a backup that has completed or is still in progress, a Show Label button is provided at the bottom of the status report screen. By pressing this button, the label for the media on which the backup is being placed is displayed. In this case, the label will not contain information for backups that are still running.
4. When displaying the status of a job that is being verified or a backup that is being restored, a Show Label button is provided at the bottom of the status report screen. By pressing this button, the label for the media being read is displayed.

A history of backup labels is stored on the admin system, and may be displayed by selecting View→Backup Labels from the main menu bar. Several options are available for finding the backup label you want to display:

View by Backup ID

Select View→Backup Labels→By Backup ID from the menu bar. A list of all labels will be displayed as shown below:

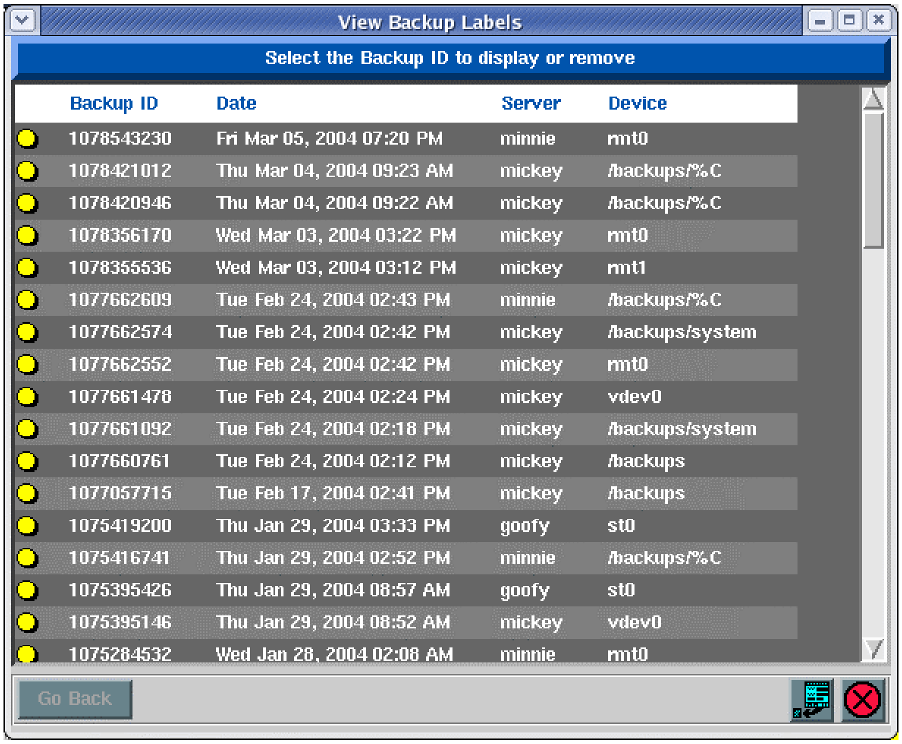

	Backup ID	Date	Server	Device
●	1078543230	Fri Mar 05, 2004 07:20 PM	minnie	rmt0
●	1078421012	Thu Mar 04, 2004 09:23 AM	mickey	/backups/%C
●	1078420946	Thu Mar 04, 2004 09:22 AM	mickey	/backups/%C
●	1078356170	Wed Mar 03, 2004 03:22 PM	mickey	rmt0
●	1078355536	Wed Mar 03, 2004 03:12 PM	mickey	rmt1
●	1077662609	Tue Feb 24, 2004 02:43 PM	minnie	/backups/%C
●	1077662574	Tue Feb 24, 2004 02:42 PM	mickey	/backups/system
●	1077662552	Tue Feb 24, 2004 02:42 PM	mickey	rmt0
●	1077661478	Tue Feb 24, 2004 02:24 PM	mickey	vdev0
●	1077661092	Tue Feb 24, 2004 02:18 PM	mickey	/backups/system
●	1077660761	Tue Feb 24, 2004 02:12 PM	mickey	/backups
●	1077057715	Tue Feb 17, 2004 02:41 PM	mickey	/backups
●	1075419200	Thu Jan 29, 2004 03:33 PM	goofy	st0
●	1075416741	Thu Jan 29, 2004 02:52 PM	minnie	/backups/%C
●	1075395426	Thu Jan 29, 2004 08:57 AM	goofy	st0
●	1075395146	Thu Jan 29, 2004 08:52 AM	mickey	vdev0
●	1075284532	Wed Jan 28, 2004 02:08 AM	minnie	rmt0

This list could become very lengthy if there are a lot of labels on file. To display the detailed label information, click on the button to the left of the desired Backup ID.

View by Tape Label ID

This option is not available when using Desktop Edition.

To display the backup label in which a physical tape was used, select View→Backup Labels→By Tape Label ID from the menu bar. A list of tape labels currently associated with backup labels is displayed. Only tape labels for which the tape ID was written to the tape prior to its use within a backup will be shown as in the following example:

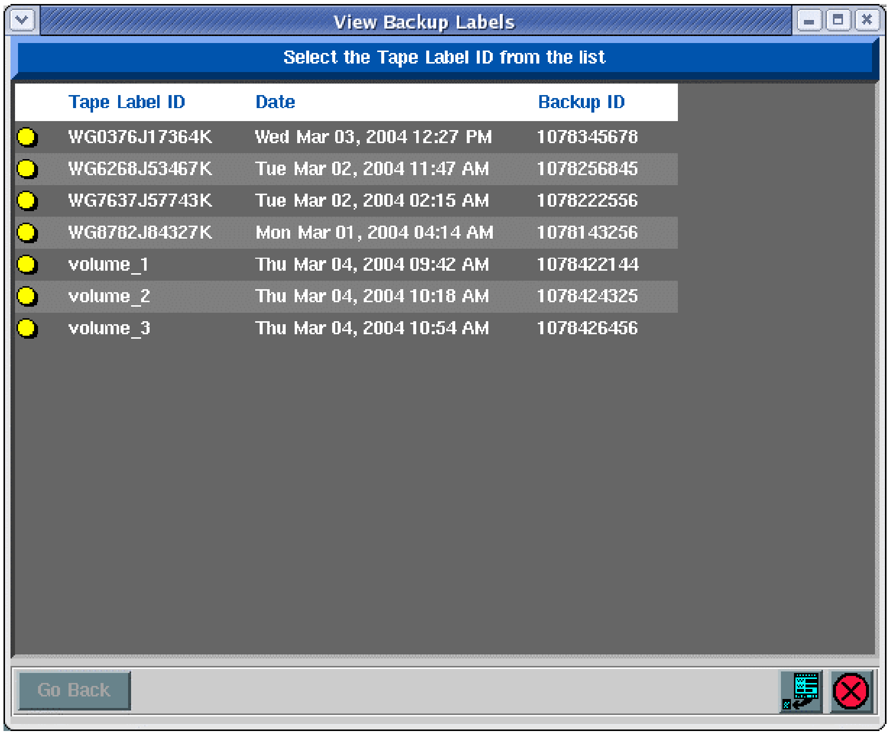

Select the Tape Label ID from the list

	Tape Label ID	Date	Backup ID
●	WG0376J17364K	Wed Mar 03, 2004 12:27 PM	1078345678
●	WG6268J53467K	Tue Mar 02, 2004 11:47 AM	1078256845
●	WG7637J57743K	Tue Mar 02, 2004 02:15 AM	1078222556
●	WG8782J84327K	Mon Mar 01, 2004 04:14 AM	1078143256
●	volume_1	Thu Mar 04, 2004 09:42 AM	1078422144
●	volume_2	Thu Mar 04, 2004 10:18 AM	1078424325
●	volume_3	Thu Mar 04, 2004 10:54 AM	1078426456

To display the backup label, click on the button to the left of the desired Tape Label ID.

View by Server

NOTE **This option is only available when using Network Administrator.**

Select View→Backup Labels→By Server from the menu bar. A list of servers and backup devices or directories is displayed as in the following example:

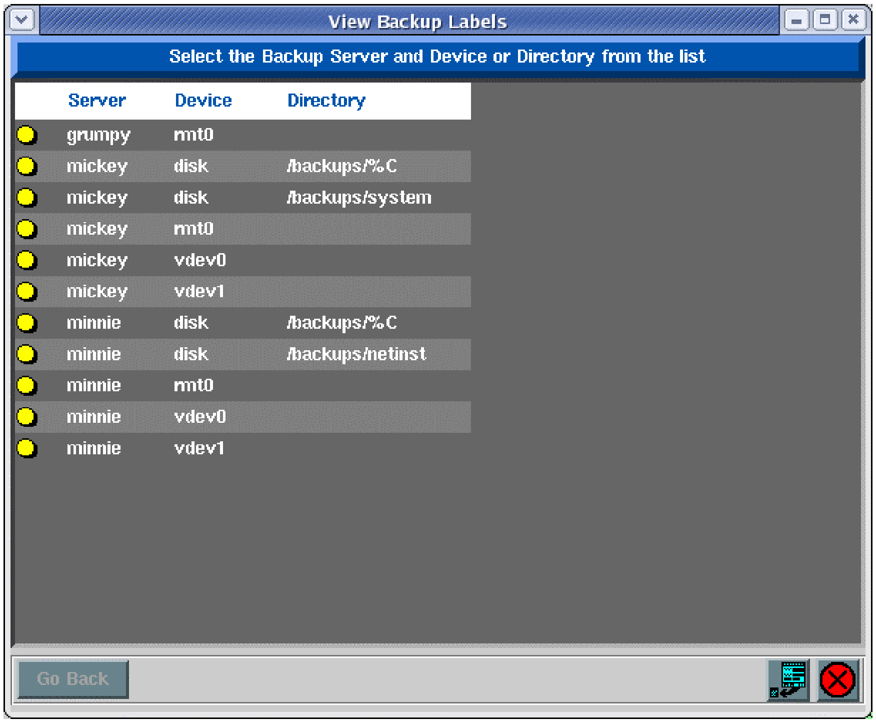

View Backup Labels

Select the Backup Server and Device or Directory from the list

	Server	Device	Directory
🟡	grumpy	rmt0	
🟡	mickey	disk	/backups/%C
🟡	mickey	disk	/backups/system
🟡	mickey	rmt0	
🟡	mickey	vdev0	
🟡	mickey	vdev1	
🟡	minnie	disk	/backups/%C
🟡	minnie	disk	/backups/netinst
🟡	minnie	rmt0	
🟡	minnie	vdev0	
🟡	minnie	vdev1	

Select the server and device (or directory) option from the list by clicking on the button to the left. Once you do so, a list of backups for the selected server and device is displayed, as shown below:

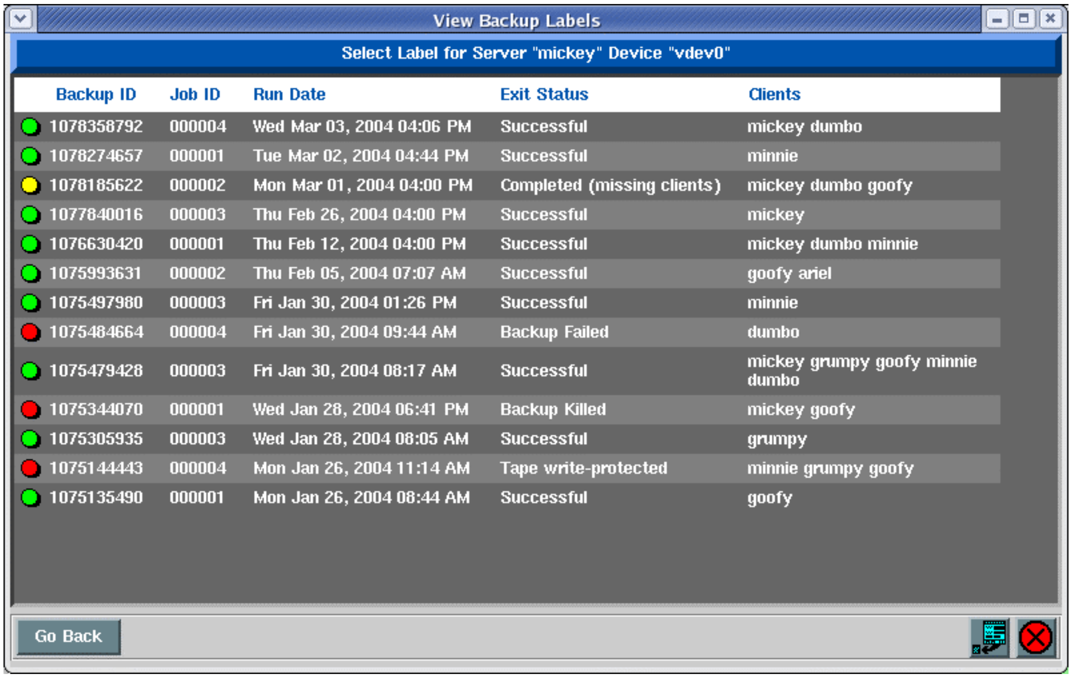

View Backup Labels

Select Label for Server "mickey" Device "vdev0"

	Backup ID	Job ID	Run Date	Exit Status	Clients
🟢	1078358792	000004	Wed Mar 03, 2004 04:06 PM	Successful	mickey dumbo
🟢	1078274657	000001	Tue Mar 02, 2004 04:44 PM	Successful	minnie
🟡	1078185622	000002	Mon Mar 01, 2004 04:00 PM	Completed (missing clients)	mickey dumbo goofy
🟢	1077840016	000003	Thu Feb 26, 2004 04:00 PM	Successful	mickey
🟢	1076630420	000001	Thu Feb 12, 2004 04:00 PM	Successful	mickey dumbo minnie
🟢	1075993631	000002	Thu Feb 05, 2004 07:07 AM	Successful	goofy ariel
🟢	1075497980	000003	Fri Jan 30, 2004 01:26 PM	Successful	minnie
🔴	1075484664	000004	Fri Jan 30, 2004 09:44 AM	Backup Failed	dumbo
🟢	1075479428	000003	Fri Jan 30, 2004 08:17 AM	Successful	mickey grumpy goofy minnie dumbo
🔴	1075344070	000001	Wed Jan 28, 2004 06:41 PM	Backup Killed	mickey goofy
🟢	1075305935	000003	Wed Jan 28, 2004 08:05 AM	Successful	grumpy
🔴	1075144443	000004	Mon Jan 26, 2004 11:14 AM	Tape write-protected	minnie grumpy goofy
🟢	1075135490	000001	Mon Jan 26, 2004 08:44 AM	Successful	goofy

Note that the color of the button represents whether the backup was successful (green) or failed (red). The Job ID, date and time, and a list of clients on the media is displayed for each label in the list. To display a label, click the button next to the Backup ID. The label will be displayed (see above). If you want to return to the server and device display screen, press the Go Back button. Otherwise press the cancel button to close this window.

View by Job ID

Select View→Backup Labels→By Job ID from the main menu bar if you want to select the label to display from a list of Job IDs.. The following screen will be displayed:

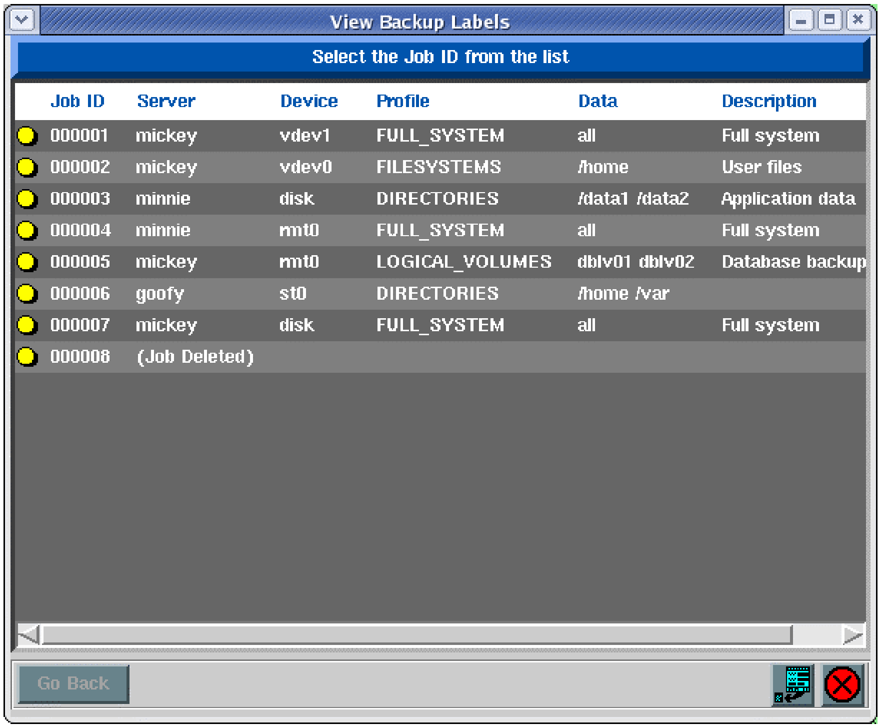

Select the Job ID from the list

	Job ID	Server	Device	Profile	Data	Description
🟡	000001	mickey	vdev1	FULL_SYSTEM	all	Full system
🟡	000002	mickey	vdev0	FILESYSTEMS	/home	User files
🟡	000003	minnie	disk	DIRECTORIES	/data1 /data2	Application data
🟡	000004	minnie	rmt0	FULL_SYSTEM	all	Full system
🟡	000005	mickey	rmt0	LOGICAL_VOLUMES	dblv01 dblv02	Database backup
🟡	000006	goofy	st0	DIRECTORIES	/home /var	
🟡	000007	mickey	disk	FULL_SYSTEM	all	Full system
🟡	000008	(Job Deleted)				

Select the desired job. An additional list will display, showing the dates the job has been run:

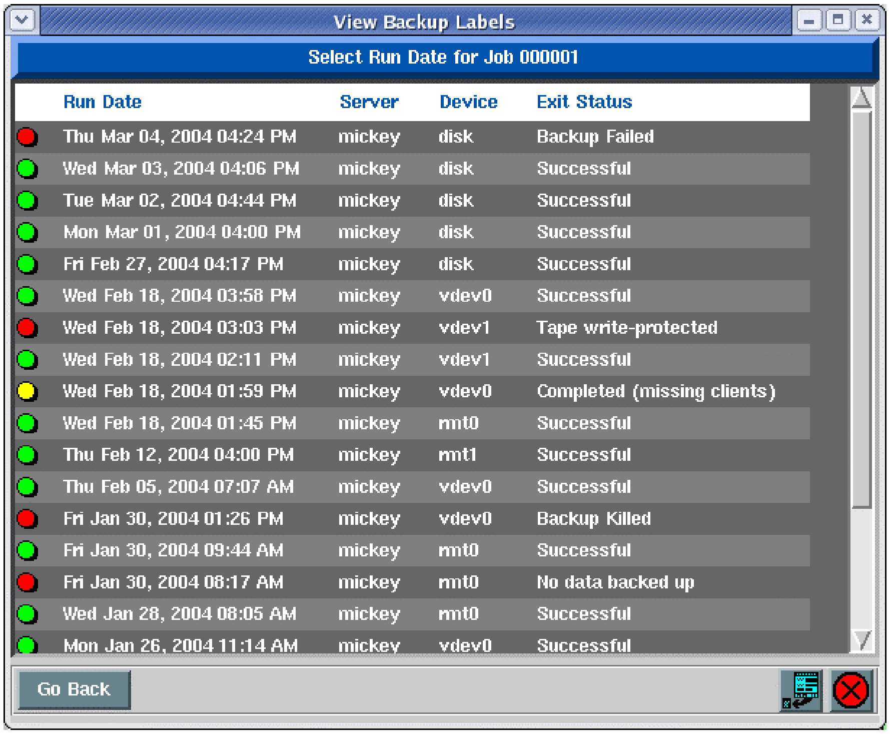

Select Run Date for Job 000001

	Run Date	Server	Device	Exit Status
🔴	Thu Mar 04, 2004 04:24 PM	mickey	disk	Backup Failed
🟢	Wed Mar 03, 2004 04:06 PM	mickey	disk	Successful
🟢	Tue Mar 02, 2004 04:44 PM	mickey	disk	Successful
🟢	Mon Mar 01, 2004 04:00 PM	mickey	disk	Successful
🟢	Fri Feb 27, 2004 04:17 PM	mickey	disk	Successful
🟢	Wed Feb 18, 2004 03:58 PM	mickey	vdev0	Successful
🔴	Wed Feb 18, 2004 03:03 PM	mickey	vdev1	Tape write-protected
🟢	Wed Feb 18, 2004 02:11 PM	mickey	vdev1	Successful
🟡	Wed Feb 18, 2004 01:59 PM	mickey	vdev0	Completed (missing clients)
🟢	Wed Feb 18, 2004 01:45 PM	mickey	rmt0	Successful
🟢	Thu Feb 12, 2004 04:00 PM	mickey	rmt1	Successful
🟢	Thu Feb 05, 2004 07:07 AM	mickey	vdev0	Successful
🔴	Fri Jan 30, 2004 01:26 PM	mickey	vdev0	Backup Killed
🟢	Fri Jan 30, 2004 09:44 AM	mickey	rmt0	Successful
🔴	Fri Jan 30, 2004 08:17 AM	mickey	rmt0	No data backed up
🟢	Wed Jan 28, 2004 08:05 AM	mickey	rmt0	Successful
🟢	Mon Jan 26, 2004 11:14 AM	mickey	vdev0	Successful

Note that the color of the button represents whether the backup was successful (green) or failed (red). The Job ID, date and time, and a list of clients on the media is displayed for each label in the list. To display a label, select a specific run date from the list. The label will be displayed (see above). If you want to return to the job display screen to select a different job, press the Go Back button. Otherwise press the cancel button to close this window.

View by Client

This option is only available when using Network Administrator.

Select View→Backup Labels→By Client from the main menu bar if you want to select the label to display from a list of backups performed by client. This option is particularly useful is you want to know the last time certain data was backed up from a client. After selecting this option, a list of clients and each backup type that the client has performed is displayed similar to the following example:

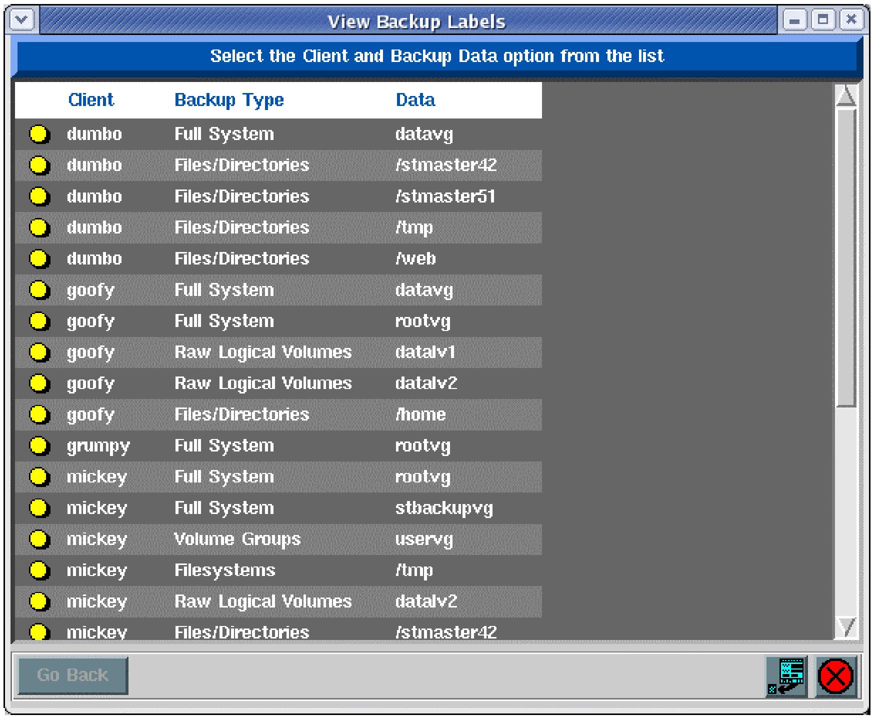

Select the Client and Backup Data option from the list

	Client	Backup Type	Data
●	dumbo	Full System	datavg
●	dumbo	Files/Directories	/stmaster42
●	dumbo	Files/Directories	/stmaster51
●	dumbo	Files/Directories	/tmp
●	dumbo	Files/Directories	/web
●	goofy	Full System	datavg
●	goofy	Full System	rootvg
●	goofy	Raw Logical Volumes	datalv1
●	goofy	Raw Logical Volumes	datalv2
●	goofy	Files/Directories	/home
●	grumpy	Full System	rootvg
●	mickey	Full System	rootvg
●	mickey	Full System	stbackupvg
●	mickey	Volume Groups	uservg
●	mickey	Filesystems	/tmp
●	mickey	Raw Logical Volumes	datalv2
●	mickey	Files/Directories	/stmaster42

Select the button next to the client and backup type you wish to display. An additional list of specific backup dates for the selected client and backup data will be shown:

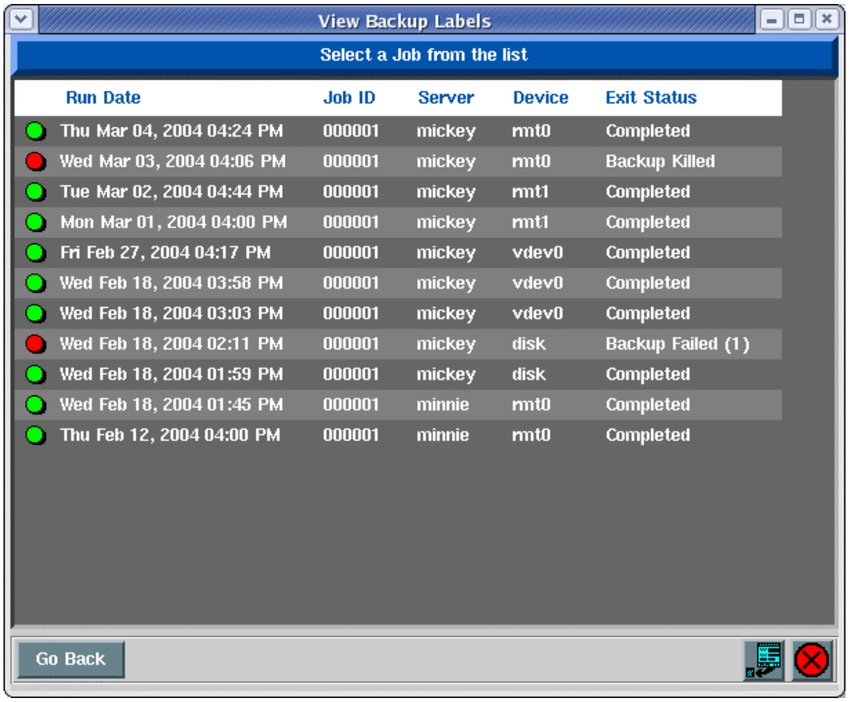

	Run Date	Job ID	Server	Device	Exit Status
🟢	Thu Mar 04, 2004 04:24 PM	000001	mickey	rmt0	Completed
🔴	Wed Mar 03, 2004 04:06 PM	000001	mickey	rmt0	Backup Killed
🟢	Tue Mar 02, 2004 04:44 PM	000001	mickey	rmt1	Completed
🟢	Mon Mar 01, 2004 04:00 PM	000001	mickey	rmt1	Completed
🟢	Fri Feb 27, 2004 04:17 PM	000001	mickey	vdev0	Completed
🟢	Wed Feb 18, 2004 03:58 PM	000001	mickey	vdev0	Completed
🟢	Wed Feb 18, 2004 03:03 PM	000001	mickey	vdev0	Completed
🔴	Wed Feb 18, 2004 02:11 PM	000001	mickey	disk	Backup Failed (1)
🟢	Wed Feb 18, 2004 01:59 PM	000001	mickey	disk	Completed
🟢	Wed Feb 18, 2004 01:45 PM	000001	minnie	rmt0	Completed
🟢	Thu Feb 12, 2004 04:00 PM	000001	minnie	rmt0	Completed

Note that the color of the button represents whether the backup was successful (green) or failed (red). The Job ID, date and time, and a list of clients on the media is displayed for each label in the list. To display a label, select a specific run date from the list. The label will be displayed (see above). If you want to return to the client list to select a different client, press the Go Back. Otherwise press the cancel button to close this window.

Read from Media

If you have a backup tape with no physical label and are unsure of its contents, the backup label may be read from the media and displayed on the screen. This option may also be used to view the backup label associated with a disk backup image file. To do so, follow these steps:

1. Insert the tape in the drive, then select View→Backup Labels→Read from Media from the menu bar. A list of servers is displayed in a listbox.

2. Select a server from the list. The tape drives, virtual devices and backup directories for the server are listed.

3. Select the device or directory from the list.

4. If you selected a directory from the list, a list of backup images in the selected directory are displayed. Select a backup image to read.

5. Select the Continue button.

The tape is read and the label information will be displayed if it exists. If a Backup ID exists on the tape, but the label information for the label is not found, an error message will appear. This indicates that the label of the tape was expired, so no detailed information on the tape contents is available.

The Backup Sequence Number

A backup sequence number, often referred to as simply the “*backup number*”, is associated with each client backup on a backup label. This number is incremented for each client backup on the media, regardless of the type of backup or the backup contents. If multiple backups, from the same or different clients, are appended to a backup tape, the backup sequence number is incremented for each new backup.

The backup sequence number is only incremented when a new backup job is appended to an existing backup, and is incremented by one for each client backed up within the job. When backing up to disk image files, each backup job always begins a new label, and therefore starts with backup sequence number 1. If a new tape backup starts at the beginning of a tape volume, a new backup label is started at backup sequence number 1.

Normally the user does not need to know the backup sequence number as this is used internally for quickly forwarding to data on the tape when performing restores. However, when a system is to be reinstalled from a **System Backup** after booting from a local tape, the user must know the backup sequence number of the backup to restore from. If there is only one backup on the tape, or if the System Backup to be restored from is the first backup on the tape, the user need not know the backup sequence number as the default value is always 1.

Expiring a Backup

Since backup tapes are usually reused after a certain amount of time, or are discarded after they have aged, it is necessary to get rid of the backup label and backup status information when the backup is no longer valid. Disk image backups may also become obsolete and need to be occasionally removed from the disk to free space on the server. This is referred to as "expiring" a backup.

By default, the Backup Retention Policy prevents tapes associated with a current backup label from being overwritten by new backup jobs. When a backup is expired, the label information is destroyed and the tape may be overwritten. The overwrite policy also determines if new disk backup jobs should overwrite an existing backup or create an additional backup image.

> **Once a backup label has been expired, it will not be possible to verify or restore data from this backup using the Backup Administrator application. However, you will still be able to reinstall a system from a System Backup even if it has been expired. If a backup has been expired or the label history has been inadvertently removed from the system, it is still possible to rebuild this information. Refer to Rebuild (unexpire) a Backup Label for details.**
>
> **Very important note: If you expire a backup that was written to disk, rather than tape, the actual disk backup will be removed from the backup server. You are given ample warning before the backup is removed, and once it has been remove it will no longer be possible to access that data.**

Manually Expiring a Backup

To manually expire a backup, first perform any one of the various methods to view the backup label . Then select the Expire/Remove button at the bottom of the screen.

Automatic Expiration of Backups

The Backup Retention Policies determines if and when an old backup may be overwritten by a new backup. Any time an old backup is overwritten by a new one, the previous backup label must be expired as the data the label points to will no longer exist.

For tape backups, if the ***Tape Overwrite/Retention Policy*** has been set to allow current labels to be overwritten by new backup jobs, the backup being overwritten will be **automatically expired**, allowing the tape to be overwritten by a new backup. This policy may allow any backup to be automatically expired and overwritten, or only backups that are older than a certain number of days.

For disk image backups, expiring the backup label also means removing the actual backup image files from the disk on the server. The ***Disk Backup Retention Policy*** may be set to automatically expire and remove disk backups over a certain number of days old when the same backup job is run again. When set, the prior backup label will be expired and the prior backup image files are automatically removed from the server. If this policy is not set, then a new backup is created in addition to the prior backup.

16. Backup Job Status & Output History

The job status and backup output, which may be displayed while a backup is running, is kept on file as long as the backup label for the job exists. It is therefore possible to view this information long after the backup has completed. The screens which appear are identical to those that may be displayed while the backup job is running, as shown in the following sample screens:

The following is the Backup Status Report Screen which appears when selecting the Status Report button on the Job Queues Display or when viewing ***Job Status/Output*** for a completed job (as described throughout this section):

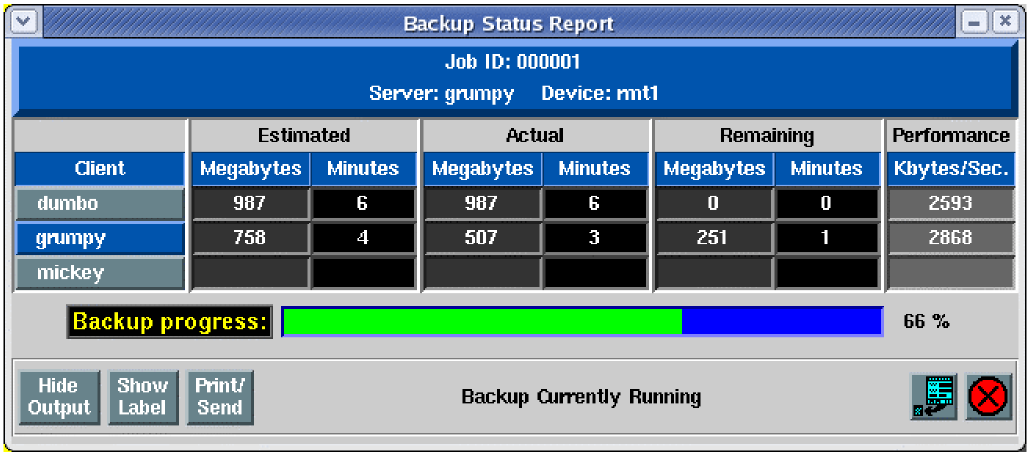

Job ID: 000001

Server: grumpy   Device: rmt1

	Estimated		Actual		Remaining		Performance
Client	Megabytes	Minutes	Megabytes	Minutes	Megabytes	Minutes	Kbytes/Sec.
dumbo	987	6	987	6	0	0	2593
grumpy	758	4	507	3	251	1	2868
mickey							

Backup progress: 66 %

Backup Currently Running

The following is the Backup Output Display which is displayed when the Show Output button is pressed on the Backup Status Report screen above:

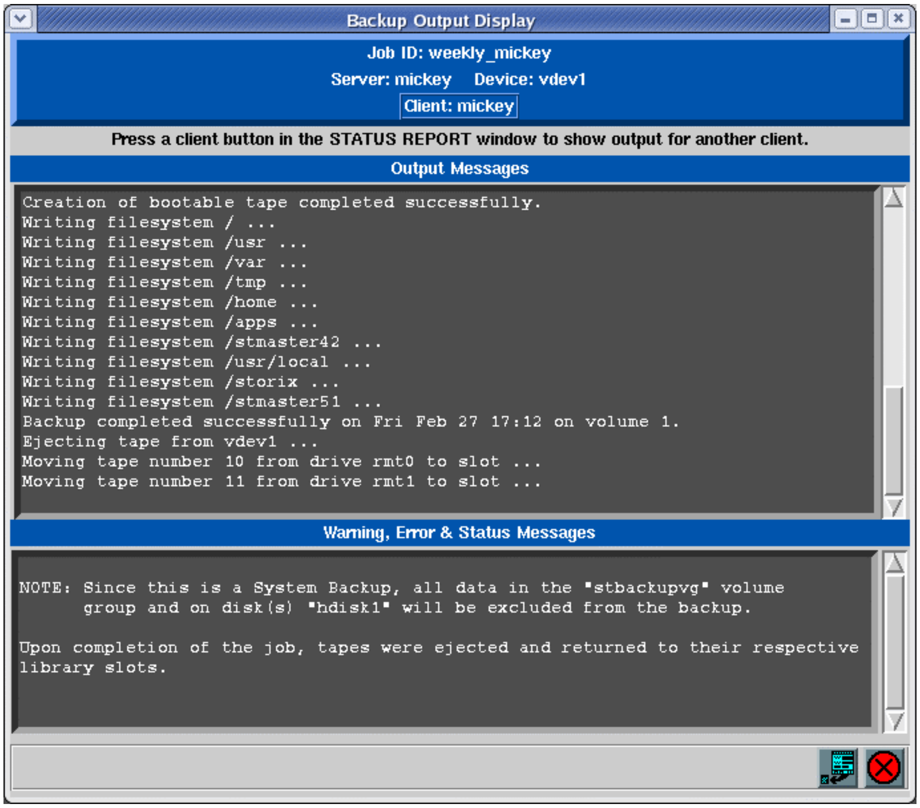

Job ID: weekly_mickey

Server: mickey   Device: vdev1

Client: mickey

Press a client button in the STATUS REPORT window to show output for another client.

Output Messages

```
Creation of bootable tape completed successfully.
Writing filesystem / ...
Writing filesystem /usr ...
Writing filesystem /var ...
Writing filesystem /tmp ...
Writing filesystem /home ...
Writing filesystem /apps ...
Writing filesystem /stmaster42 ...
Writing filesystem /usr/local ...
Writing filesystem /storix ...
Writing filesystem /stmaster51 ...
Backup completed successfully on Fri Feb 27 17:12 on volume 1.
Ejecting tape from vdev1 ...
Moving tape number 10 from drive rmt0 to slot ...
Moving tape number 11 from drive rmt1 to slot ...
```

Warning, Error & Status Messages

```
NOTE: Since this is a System Backup, all data in the "stbackupvg" volume
      group and on disk(s) "hdisk1" will be excluded from the backup.

Upon completion of the job, tapes were ejected and returned to their respective
library slots.
```

The desired job status information may be obtained in each of the following ways:

View by Server

> NoTE **This option is only available when using Network Administrator.**

Select View→Backup Status/Output→By Server from the menu bar. A list of servers and backup devices or directories is displayed as in the same screen available when selecting to View Backup Labels by Server.

Select the server and device (or directory) option from the list by clicking on the button to the left. Once you do so, a list of backups performed to the selected server and device is displayed, as in the example View Backup labels by Server.

Note that the color of the button represents whether the backup was successful (green) or failed (red). The Job ID, date and time, and a list of clients on the media is displayed for each label in the list. To display the Backup Status Report, select a specific run date from the list. The Status Report Screen will be displayed (see above). If you want to return to the server list to select a different server and device, press the Go Back. Otherwise press the cancel button to close this window.

To show the backup output display, select the Show Output button on the status report screen.

View by Job ID

Select View→Job Status/Output→By Job ID from the main menu bar if you want to select the backup status to display from a list of Job Ids. After selecting this option, a list of configured Jobs and corresponding job information is displayed similar to the screen shown when you select to View Backup Labels by Job ID.

Select the desired job. An additional list will display, showing the dates the job has been run, as seen in the display View Backup Labels by Job ID.

Note that the color of the button represents whether the backup was successful (green) or failed (red). The Job ID, date and time, and a list of clients on the media is displayed for each label in the list. To display the Backup Status Report, select a specific run date from the list. The Status Report Screen will be displayed (see above). If you want to return to the job display screen to select a different job, press the Go Back button. Otherwise press the cancel button to close this window.

To show the backup output display, select the Show Output button on the status report screen.

View by Client

> NoTE **This option is only available when using Network Administrator.**

Select View→Job Stauts/Output→By Client from the main menu bar if you want to select the job to display from a list of backups performed by client. After selecting this option, a list of clients and each backup type that the client has performed is displayed similar to the example when you select to View Backup labels by Client.

Select the button next to the client and backup type you wish to display. An additional list of specific backup dates for the selected client and backup data will be shown, as seen in View Backup Labels by Client.

Note that the color of the button represents whether the backup was successful (green) or failed (red). The Job ID, date and time, and a list of clients on the media is displayed for each label in the list. To display the Backup Status Report, select a specific run date from the list. The Status Report Screen will be displayed (see above).

Note that the display will include all client backups in the job, not just the selected client. The selected client button on the status screen will be automatically selected, however, so you can show the backup command output for the client by pressing the Show Output button.

If you want to return to the client list to select a different client, press the Go Back. Otherwise press the cancel button to close this window.

17. Verify a Backup

After a backup job has complete, it is often a good precaution to verify the backup to ensure the data on the backup media is complete and readable. The verify process reads all of the data on the backups and verifies it is in the correct format. The backup job may have included multiple clients. For tape backups, there may also be multiple jobs stacked on the same tape or set of tapes.

If you selected to automatically Verify Backups When Complete within the Backup Job configuration, then the backup data was verified at that time, and it is generally not necessary to verify again. However, if you are using the *Personal Edition*, where this option is not available, or do not auto-verify as part of the backup process, you may do so at a later time by following the steps in this section.

It will be possible to select each client backup on the media that you want to verify, even those from different jobs.

Selecting What to verify

To verify a backup job, perform the following steps:

1. Select Actions→Verify Backup Jobs from the menu bar.

2. If using a *Network Administrator* license, a list will pop up showing the configured backup servers. Select the backup server on which the backup was written.

3. A list of devices and/or directories for the selected backup server will appear. Click on the device or directory onto which the backup was written, then press the Continue button. The information about the backup will then be read from the media.

4. If you selected to verify a backup written to a disk directory, you will be provided a list of backup jobs in the selected directory, similar to the following example:

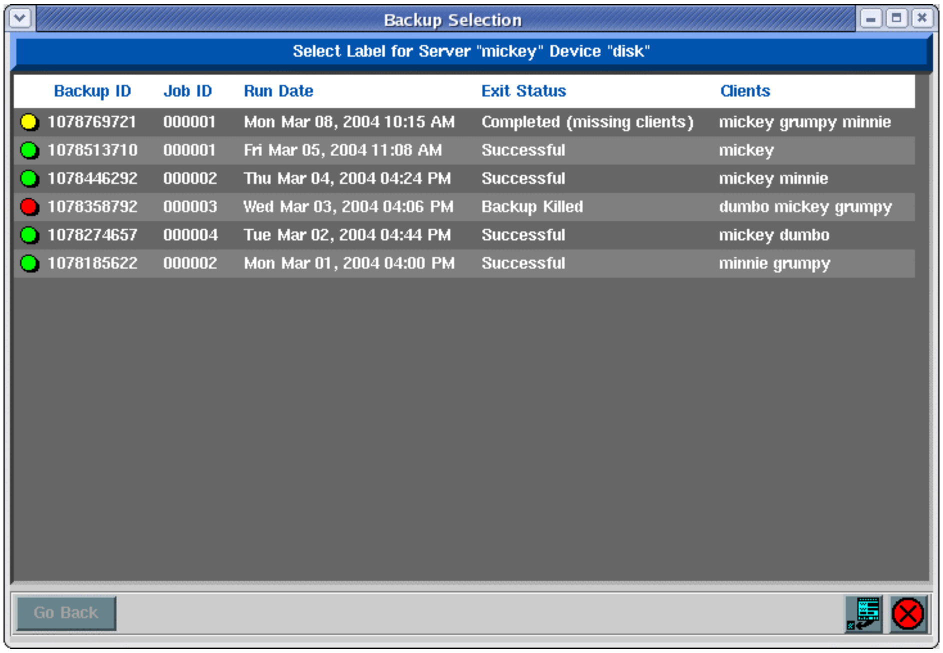

Select Label for Server "mickey" Device "disk"

	Backup ID	Job ID	Run Date	Exit Status	Clients
🟡	1078769721	000001	Mon Mar 08, 2004 10:15 AM	Completed (missing clients)	mickey grumpy minnie
🟢	1078513710	000001	Fri Mar 05, 2004 11:08 AM	Successful	mickey
🟢	1078446292	000002	Thu Mar 04, 2004 04:24 PM	Successful	mickey minnie
🔴	1078358792	000003	Wed Mar 03, 2004 04:06 PM	Backup Killed	dumbo mickey grumpy
🟢	1078274657	000004	Tue Mar 02, 2004 04:44 PM	Successful	mickey dumbo
🟢	1078185622	000002	Mon Mar 01, 2004 04:00 PM	Successful	minnie grumpy

Select the specific backup job to verify by clicking on the button to the left of the desired job.

5. Next, a screen will appear with a list of backups on the media. For disk backups, this list will contain all of the backups within the selected job. For tape backups, there may be multiple jobs on the media. In this case, the list will contain all of the backups, even those from different jobs. The information about the backup will be preceded by the ***backup sequence number***, starting with 1 and ending with the last backup on the media.

   The following is a sample of this screen:

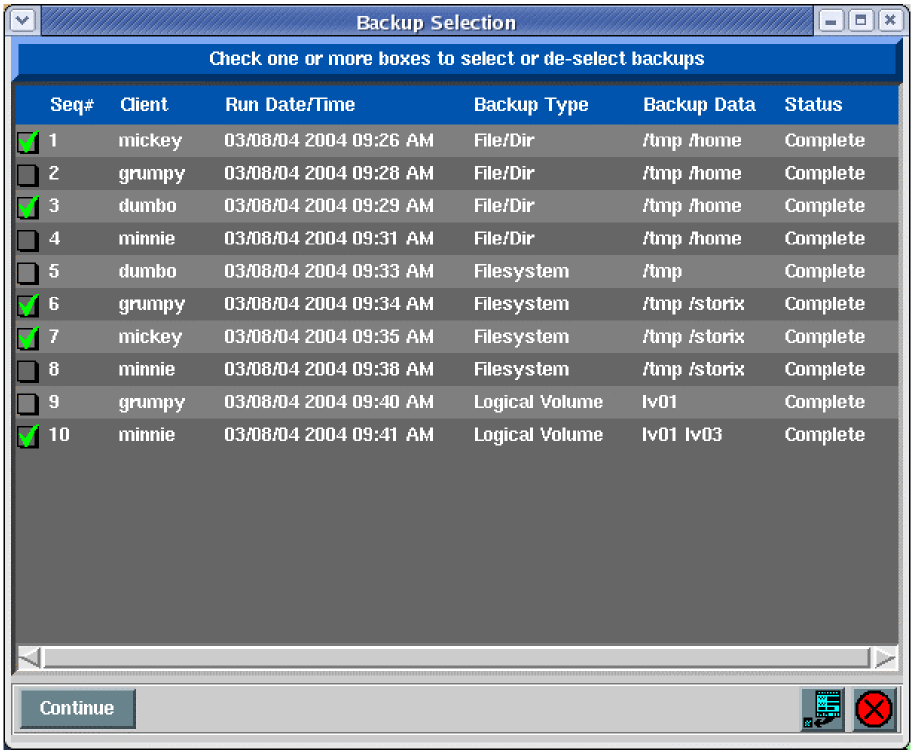

Check one or more boxes to select or de-select backups

	Seq#	Client	Run Date/Time	Backup Type	Backup Data	Status
☑	1	mickey	03/08/04 2004 09:26 AM	File/Dir	/tmp /home	Complete
□	2	grumpy	03/08/04 2004 09:28 AM	File/Dir	/tmp /home	Complete
☑	3	dumbo	03/08/04 2004 09:29 AM	File/Dir	/tmp /home	Complete
□	4	minnie	03/08/04 2004 09:31 AM	File/Dir	/tmp /home	Complete
□	5	dumbo	03/08/04 2004 09:33 AM	Filesystem	/tmp	Complete
☑	6	grumpy	03/08/04 2004 09:34 AM	Filesystem	/tmp /storix	Complete
☑	7	mickey	03/08/04 2004 09:35 AM	Filesystem	/tmp /storix	Complete
□	8	minnie	03/08/04 2004 09:38 AM	Filesystem	/tmp /storix	Complete
□	9	grumpy	03/08/04 2004 09:40 AM	Logical Volume	lv01	Complete
☑	10	minnie	03/08/04 2004 09:41 AM	Logical Volume	lv01 lv03	Complete

   You may select any one or more backups to verify by clicking on the box to the left of the desired selection and a checkmark will appear in the box. If you wish to de-select an option, simply click the box again and the checkmark will disappear. When all selections have been made, click the Continue button at the bottom of the screen to begin the verify.

Displaying the Status and Output of the Verify

The verify will begin, and the status report screen, as shown below, will appear automatically. Listed on the screen will be a status line for backup previously selected. Information pertaining to the progress and performance of the verify will be updated for each line as the corresponding backup is being read. If not all of the backups on the media were selected, the process may ***fast-forward*** over certain backups before reading the next. Fast-forwarding a tape backup is much faster than reading through all the data.

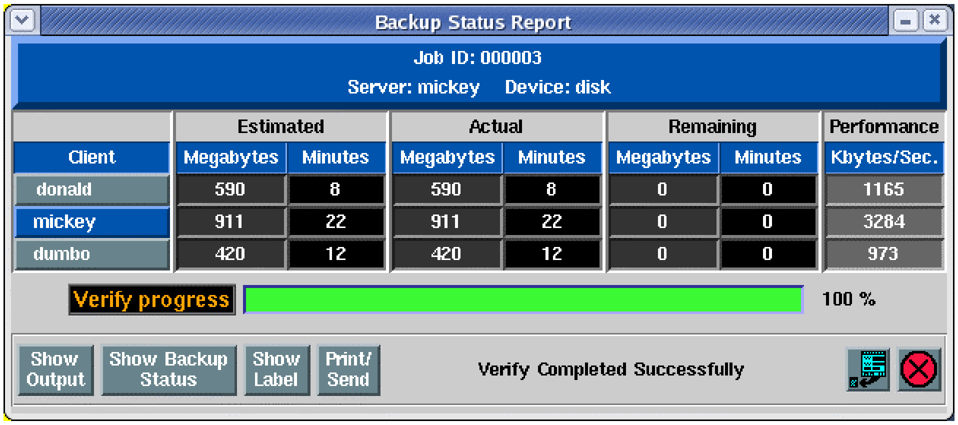
Backup Status Report

Job ID: 000003

Server: mickey   Device: disk

	Estimated		Actual		Remaining		Performance
Client	Megabytes	Minutes	Megabytes	Minutes	Megabytes	Minutes	Kbytes/Sec.
donald	590	8	590	8	0	0	1165
mickey	911	22	911	22	0	0	3284
dumbo	420	12	420	12	0	0	973

Verify progress 100 %

Verify Completed Successfully

Note that this screen may not be closed as long as the verify is running. It will remain on the screen after the verify completes until it is closed by the user. Once the screen is closed, the verify status and output messages may not be redisplayed.

To view the output of the verification of a particular backup, first select the *Backup Sequence Number:Client* button (only the *Backup Sequence Number* appears if not using a *Network Administrator*), then press the Show Output button at the bottom of the screen. An output screen similar to the following will then appear, showing the status of the verify:

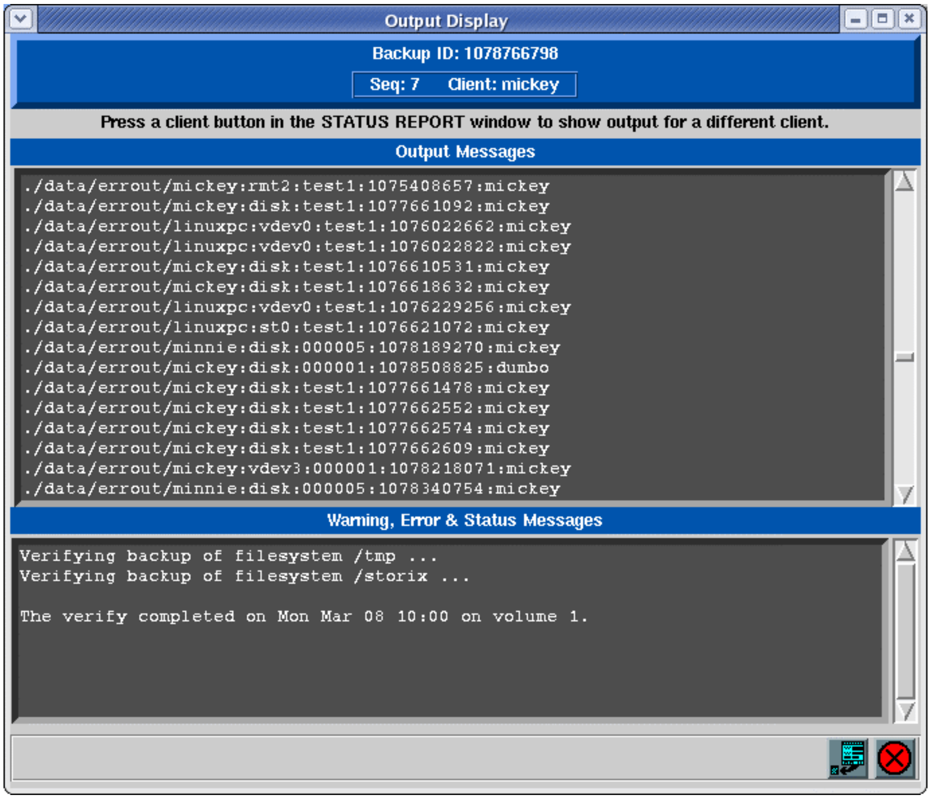
Output Display

Backup ID: 1078766798

Seq: 7   Client: mickey

Press a client button in the STATUS REPORT window to show output for a different client.

Output Messages

```
./data/errout/mickey:rmt2:test1:1075408657:mickey
./data/errout/mickey:disk:test1:1077661092:mickey
./data/errout/linuxpc:vdev0:test1:1076022662:mickey
./data/errout/linuxpc:vdev0:test1:1076022822:mickey
./data/errout/mickey:disk:test1:1076610531:mickey
./data/errout/mickey:disk:test1:1076618632:mickey
./data/errout/linuxpc:vdev0:test1:1076229256:mickey
./data/errout/linuxpc:st0:test1:1076621072:mickey
./data/errout/minnie:disk:000005:1078189270:mickey
./data/errout/mickey:disk:000001:1078508825:dumbo
./data/errout/mickey:disk:test1:1077661478:mickey
./data/errout/mickey:disk:test1:1077662552:mickey
./data/errout/mickey:disk:test1:1077662574:mickey
./data/errout/mickey:disk:test1:1077662609:mickey
./data/errout/mickey:vdev3:000001:1078218071:mickey
./data/errout/minnie:disk:000005:1078340754:mickey
```

Warning, Error & Status Messages

```
Verifying backup of filesystem /tmp ...
Verifying backup of filesystem /storix ...

The verify completed on Mon Mar 08 10:00 on volume 1.
```

If the verify is of any backup type other than a *Logical Volume Backup* or a *Power System Backup*, a complete list of files on the backup will be displayed as each file is read. This screen may be closed and redisplayed at any time, even after the verify completes, as long as the Verify Status Report screen has not been closed.

In addition to the job output, summary information for the selected backup may be displayed by selecting the *Backup Sequence Number:Client* button (or just *Backup Sequence* Number button if not using *Network*

Administrator), then pressing the Backup Info button. A screen similar to the following example will appear. Simply press the Dismiss button to close this window.

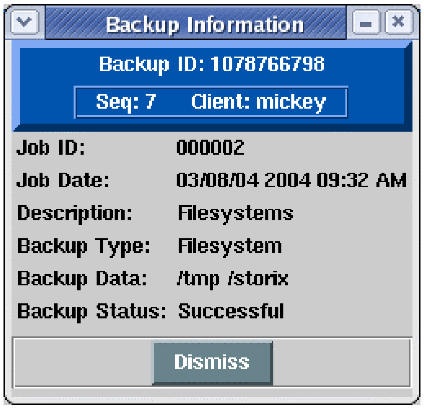

Backup ID: 1078766798

Seq: 7 Client: mickey

Job ID: 000002
Job Date: 03/08/04 2004 09:32 AM
Description: Filesystems
Backup Type: Filesystem
Backup Data: /tmp /storix
Backup Status: Successful

18. Recreate Volume Groups, Logical Volumes or Filesystems

The options described in this chapter are supported for *AIX* systems only. Due to the complexity of *Linux* configurations in building logical volumes onto meta-disks onto partitions, etc, this features is not available for Linux at this time.

When to Use These Options

Due to various system problems, it may be necessary to recreate a filesystem or even an entire volume group that had to be removed from the system due to a failed disk drive or other problem. Since changes frequently occur to the system configuration, such as the expansion of filesystems, and moving or striping of logical volumes across disks, it is often not known the proper sizes and locations of the logical volumes and filesystems needed to restore the data properly. This information is stored on the backup media, however, and these options provide an automated way of recreating the volume groups, logical volumes and filesystems exactly as they were previously without prior knowledge.

Use one of these options to recreate the volume groups, logical volumes and/or filesystems into which you will later restore the data using the option Restore Data from a Backup.

It is sometimes also desirable to replicate a volume group configuration from one system onto another. This option will allow you to use the information stored on a backup to create or recreate volume groups, filesystems or logical volumes on another system, while changing the locations and sizes of the filesystems and logical volumes to accommodate the new system.

In addition, a volume group or logical volume may be recreated on the same system from which it originated, even if the original volume group or logical volume still exist. This is handy for being able to restore prior data to the system and still keep the current copy available. This is accomplished by assigning a different volume group or logical volume name(s) to the new volume group or logical volumes created.

Important: This is the only option in the Backup Administrator that must run a user interface on the client (although the client system need not have a graphical display). In order to have the user interface (which is running on the client) display on the admin system, the client must have AIXwindows installed. If AIXwindows cannot be found on the client, an appropriate message will be displayed and you may not continue. You must either install AIXwindows on the client or rebuild the volume groups, logical volumes or filesystems manually on the client.

Recreate Volume Groups

To recreate volume groups, you must have accessible a System or Volume Group Backup containing the desired volume groups you with to create.

To recreate a volume group, perform the following steps:

1. Select Actions➔Recreate Volume Groups from the menu bar.

2. If using a *Network Administrator*, a list will pop up showing the configured backup servers. Select the backup server on which the backup was written.

3. A list of devices and/or directories for the selected backup server will appear. Click on the device or directory onto which the backup was written, then press the Continue button. The information about the backup will then be read from the media.

4. If you selected to restore from backup written to a disk directory, you will be provided a list of backup jobs in the selected directory, similar to the following example:

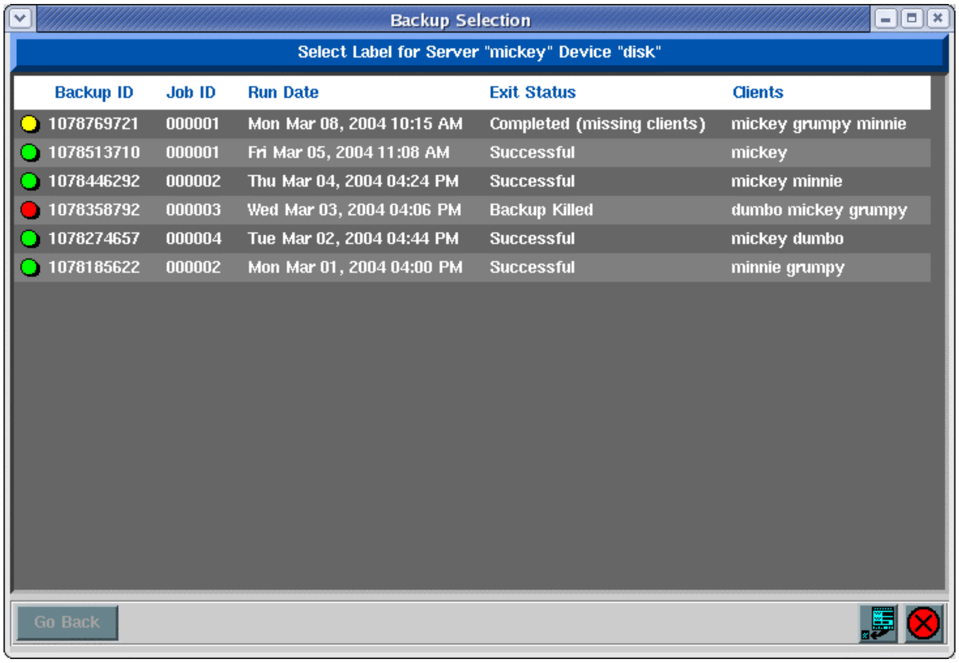

Select Label for Server "mickey" Device "disk"

	Backup ID	Job ID	Run Date	Exit Status	Clients
🟡	1078769721	000001	Mon Mar 08, 2004 10:15 AM	Completed (missing clients)	mickey grumpy minnie
🟢	1078513710	000001	Fri Mar 05, 2004 11:08 AM	Successful	mickey
🟢	1078446292	000002	Thu Mar 04, 2004 04:24 PM	Successful	mickey minnie
🔴	1078358792	000003	Wed Mar 03, 2004 04:06 PM	Backup Killed	dumbo mickey grumpy
🟢	1078274657	000004	Tue Mar 02, 2004 04:44 PM	Successful	mickey dumbo
🟢	1078185622	000002	Mon Mar 01, 2004 04:00 PM	Successful	minnie grumpy

Select the specific backup job from which to restore by clicking on the button to the left of the desired job.

5. Next, a screen will appear with a list of backups on the media. For disk backups, this list will contain all of the backups in the selected job. For tape backups, there may be multiple jobs on the media. In this case, the list will contain all of the backups, even those from different jobs. The information about the backup will be preceded by the ***backup sequence number***, starting with 1 and ending with the last backup on the media.

   The following is a sample of this screen:

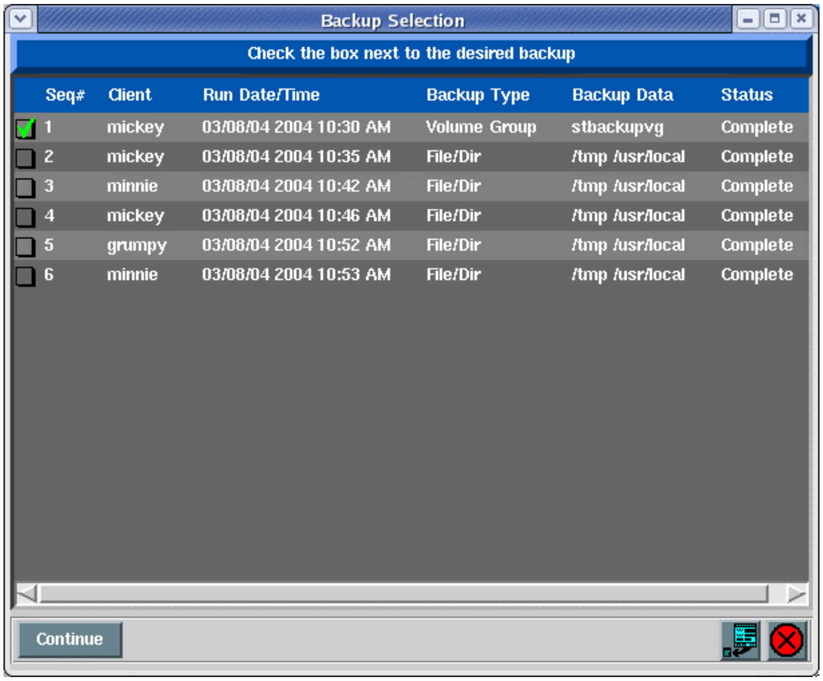

Backup Selection

Check the box next to the desired backup

	Seq#	Client	Run Date/Time	Backup Type	Backup Data	Status
☑	1	mickey	03/08/04 2004 10:30 AM	Volume Group	stbackupvg	Complete
□	2	mickey	03/08/04 2004 10:35 AM	File/Dir	/tmp /usr/local	Complete
□	3	minnie	03/08/04 2004 10:42 AM	File/Dir	/tmp /usr/local	Complete
□	4	mickey	03/08/04 2004 10:46 AM	File/Dir	/tmp /usr/local	Complete
□	5	grumpy	03/08/04 2004 10:52 AM	File/Dir	/tmp /usr/local	Complete
□	6	minnie	03/08/04 2004 10:53 AM	File/Dir	/tmp /usr/local	Complete

You may select the backup from which to recreate the volume groups by clicking on the box to the left of the desired selection and a checkmark will appear in the box. Only one selection may be made. If you select a different option, the checkmark will be removed from the previous selection. After making your selection, click the Continue button at the bottom of the screen to continue.

6. Next, a screen similar to the following will appear:

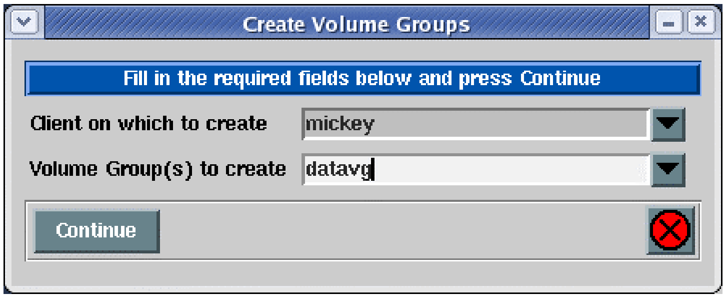

Create Volume Groups

Fill in the required fields below and press Continue

Client on which to create: mickey

Volume Group(s) to create: datavg

If using a *Network Administrator*, the ***Client on which to create*** field will show the original client from which the backup was made. The backup information may be used to create the volume group(s) on a different client by selecting the arrow button to the right of this field and selecting a different client from the list.

To list and select the volume group(s) that are defined on the backup and select one or more to create from the list, select the arrow next to the ***Volume Group(s) to create*** field.

7. When all selections have been made, press the Continue button at the bottom of the screen. A new screen similar to the following will appear and the LVM data on the media will be retrieved and checked for consistency with the current system configuration:

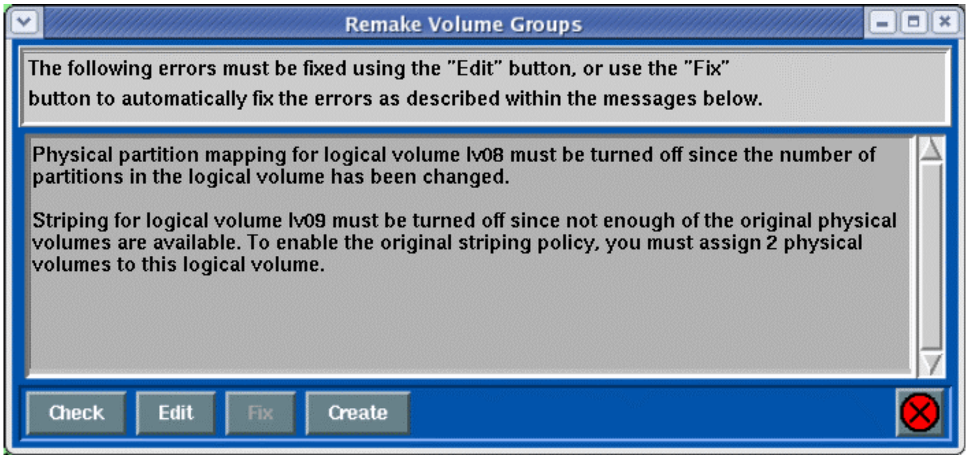

The following errors must be fixed using the "Edit" button, or use the "Fix" button to automatically fix the errors as described within the messages below.

Physical partition mapping for logical volume lv08 must be turned off since the number of partitions in the logical volume has been changed.

Striping for logical volume lv09 must be turned off since not enough of the original physical volumes are available. To enable the original striping policy, you must assign 2 physical volumes to this logical volume.

If there are changes required to make the selected volume group fit onto the current system, the Edit and Fix buttons will become available. If there are no problem found, the Create button will be available.

a. The Check button may be used to check the LVM information again. This is automatically performed when you initially display this screen and any time you change the volume group, logical volume or filesystem information.

b. The Edit button may be used to change any of the volume group, logical volume or filesystem information defined on the backup in order to make the volume group conform to the current system configuration. This may include changing the volume group or logical volume names, selecting different disks on which to build the volume group, etc. This editing process is identical to that which is available during a system installation, and is described in detail in the section ***Change the Volume Group, Logical Volume and Filesystem Information*** in the ***SBAdmin AIX System Recovery Guide***. After following the instructions in that section, press the **ESC (escape) key** on that screen to exit and save changes.

c. The Fix button may be used if there were non-fatal errors that can be automatically repaired. For instance, if there is only one physical volume available, and a logical volume is striped, the striping would need to be turned off to create the logical volume as this required at least two physical volumes. The errors described in the messages section of the window indicate if and what changes would automatically be made if the Fix button is selected.

d. The Create button will become available only after all errors, both fatal and non-fatal, have been fixed (either using the Fix button or by editing the volume group, logical volume or filesystem information using the Edit button). When you select this button, the volume group and all of its logical volumes and filesystems will be created as defined and the messages will be updated to reflect the progress and completion of the process as follows:

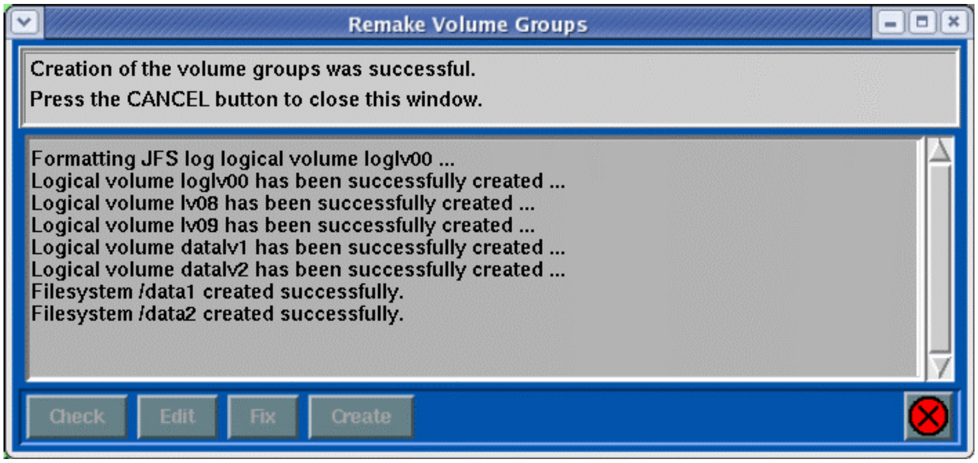

Creation of the volume groups was successful.
Press the CANCEL button to close this window.

```
Formatting JFS log logical volume loglv00 ...
Logical volume loglv00 has been successfully created ...
Logical volume lv08 has been successfully created ...
Logical volume lv09 has been successfully created ...
Logical volume datalv1 has been successfully created ...
Logical volume datalv2 has been successfully created ...
Filesystem /data1 created successfully.
Filesystem /data2 created successfully.
```

Recreate Logical Volumes or Filesystems

To recreate logical volumes or filesystems, you must have accessible a System, Volume Group, Logical Volume or Filesystem Backup containing the desired logical volumes or filesystems you with to create.

To recreate a logical volume or filesystem, perform the following steps:

1. Select Actions→Recreate Logical Volumes or Filesystems from the menu bar.

2. If using a *Network Administrator*, a list will pop up showing the configured backup servers. Select the backup server on which the backup was written.

3. A list of devices and/or directories for the selected backup server will appear. Click on the device or directory onto which the backup was written, then press the Continue button. The information about the backup will then be read from the media.

4. If you selected to restore from backup written to a disk directory, you will be provided a list of backup jobs in the selected directory, similar to the example used to Select Jobs when using the option to Recreate Volume Groups.

   Select the specific backup job from which to restore by clicking on the button to the left of the desired job.

5. Next, a screen will appear with a list of backups on the. For disk backups, this list will contain all of the backups in the selected job. For tape backups, there may be multiple jobs on the media. In this case, the list will contain all of the backups, even those from different jobs. The information about the backup will be preceded by the ***backup sequence number***, starting with 1 and ending with the last backup on the media.

5. A screen will be displayed, similar to the example used to Select Backup when using the option to Recreate Volume Groups.

   You may select the backup from which to recreate the logical volume or filesystem by clicking on the box to the left of the desired selection and a checkmark will appear in the box. Only one selection may be made. If you select a different option, the checkmark will be removed from the previous selection. After making your selection, click the Continue button at the bottom of the screen to continue.

6. Next, a screen similar to the following will appear:

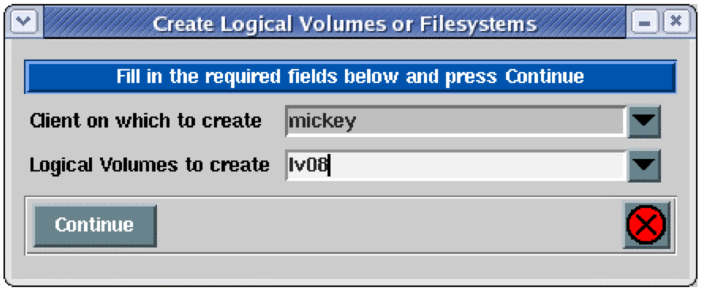
Create Logical Volumes or Filesystems
Fill in the required fields below and press Continue
Client on which to create: mickey

Logical Volumes to create: lv08

If using a *Network Administrator*, the ***Client on which to create*** field will show the original client from which the backup was made. The backup information may be used to create the logical volumes or filesystems on a different client by selecting the arrow button to the right of this field and selecting a different client from the list.

To list and select the logical volumes (filesystems) that are defined on the backup and select one or more to create from the list, select the arrow next to the ***Logical Volumes to create*** field. If recreating filesystems, this will show the ***Filesystems to create***. However, when selecting the filesystem from the list, the corresponding logical volume will be placed in the entry field.

7. When all selections have been made, press the Continue button at the bottom of the screen. A new screen similar to the following will appear and the LVM data on the media will be retrieved and checked for consistency with the current system configuration:

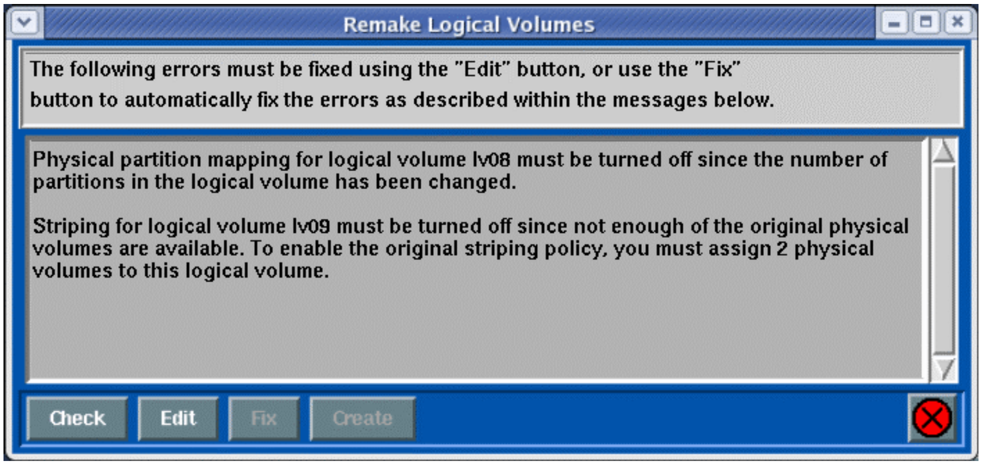

The following errors must be fixed using the "Edit" button, or use the "Fix" button to automatically fix the errors as described within the messages below.

Physical partition mapping for logical volume lv08 must be turned off since the number of partitions in the logical volume has been changed.

Striping for logical volume lv09 must be turned off since not enough of the original physical volumes are available. To enable the original striping policy, you must assign 2 physical volumes to this logical volume.

If there are changes required to make the selected volume group fit onto the current system, the Edit and Fix buttons will become available. If there are no problem found, the Create button will be available.

a. The Check button may be used to check the LVM information again. This is automatically performed when you initially display this screen and any time you change the volume group, logical volume or filesystem information.

b. The Edit button may be used to change any of the volume group, logical volume or filesystem information defined on the backup in order to make the volume group conform to the current system configuration. This may include changing the volume group or logical volume names, selecting different disks on which to build the volume group, etc. This editing process is identical to that which is available during a system installation, and is described in detail in the section ***Change the Volume Group, Logical Volume and Filesystem Information*** in the ***SBAdmin AIX System Recovery Guide***. After following the instructions in that section, press the **ESC (escape) key** on that screen to exit and save changes.

c. The Fix button may be used if there were non-fatal errors that can be automatically repaired. For instance, if there is only one physical volume available, and a logical volume is striped, the striping would need to be turned off to create the logical volume as this required at least two physical volumes. The errors described in the messages section of the window indicate if and what changes would automatically be made if the Fix button is selected.

d. The Create button will become available only after all errors, both fatal and non-fatal, have been fixed (either using the Fix button or by editing the volume group, logical volume or filesystem information using the Edit button). When you select this button, the volume group and all of its logical volumes and filesystems will be created as defined and the messages will be updated to reflect the progress and completion of the process as follows:

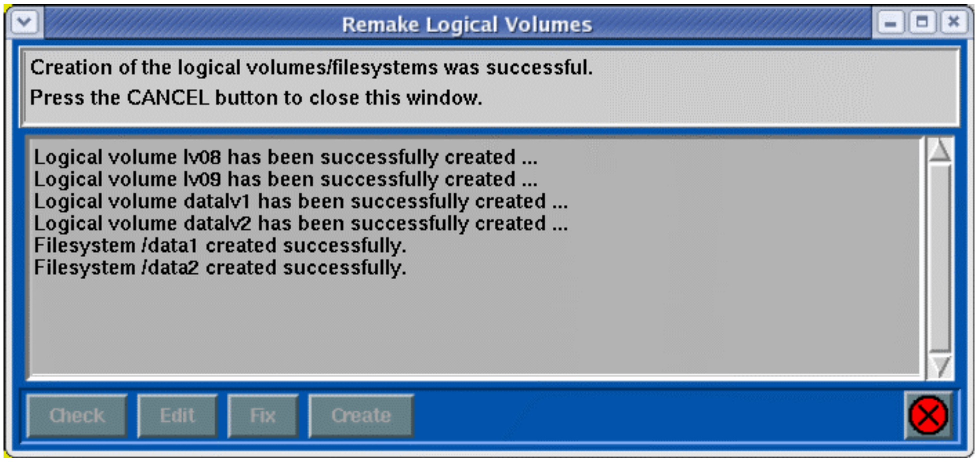

Creation of the logical volumes/filesystems was successful.
Press the CANCEL button to close this window.

```
Logical volume lv08 has been successfully created ...
Logical volume lv09 has been successfully created ...
Logical volume datalv1 has been successfully created ...
Logical volume datalv2 has been successfully created ...
Filesystem /data1 created successfully.
Filesystem /data2 created successfully.
```

19. Restore Data from a Backup

Data may be restored from a backup device or directory on any backup server to any client using the Backup Administrator. Using a *Network Administrator* license, a backup taken from one client may also be restored to another client, unless it is a disk backup and, for security reasons, you chose (in the backup profile) not to allow a client to read a backup on the backup server's disk that belonged to a different client.

Any type of data contained on a backup may be restored. A System Backup, for instance, may contain multiple *volume groups*, each of which may contain *raw logical volumes* and *filesystems*, each of which may contain various *directories*, which each contain multiple *files*. It is therefore possible to restore one or more files, directories, logical volumes, filesystems, volume groups, or the entire system from a System Backup!

> **Restoring data from a backup is not the same as *reinstalling* a client from a System Backup. This is a different process which is described in detail in the section *Installing from a System Backup* in the *SBAdmin System Recovery Guide*.**

Selecting the Backup to Restore From

To restore data from a backup, perform the following steps:

1. Select Actions➔Restore Data from a Backup from the menu bar.

2. If using a *Network Administrator*, a list will pop up showing the configured backup servers. Select the backup server on which the backup was written.

3. A list of devices and/or directories for the selected backup server will appear. Click on the device or directory onto which the backup was written, then press the Continue button. The information about the backup will then be read from the media.

4. If you selected to restore from backup written to a disk directory, you will be provided a list of backup jobs in the selected directory, similar to the following example:

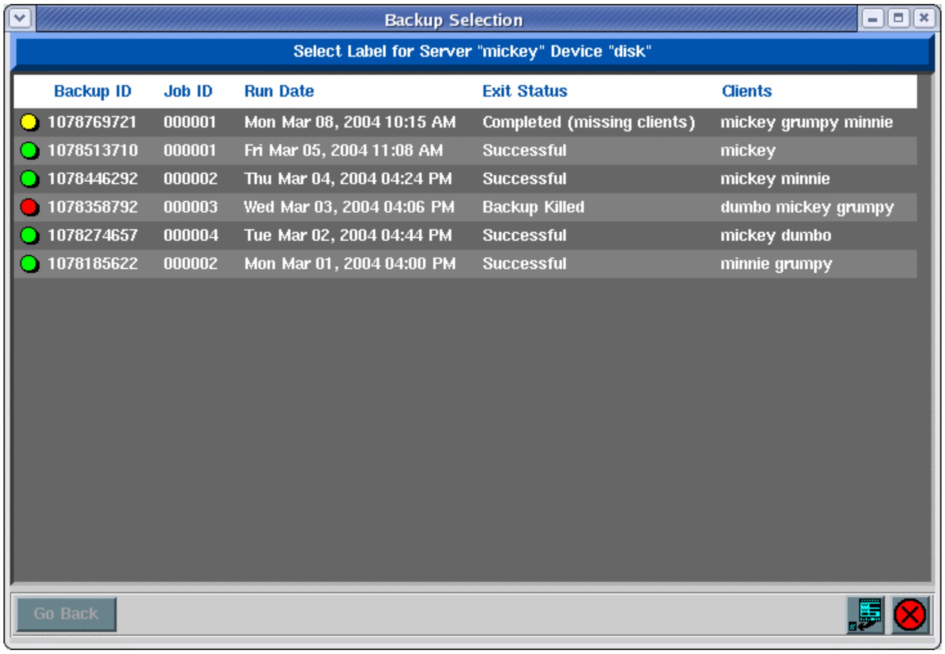

Backup Selection

Select Label for Server "mickey" Device "disk"

	Backup ID	Job ID	Run Date	Exit Status	Clients
🟡	1078769721	000001	Mon Mar 08, 2004 10:15 AM	Completed (missing clients)	mickey grumpy minnie
🟢	1078513710	000001	Fri Mar 05, 2004 11:08 AM	Successful	mickey
🟢	1078446292	000002	Thu Mar 04, 2004 04:24 PM	Successful	mickey minnie
🔴	1078358792	000003	Wed Mar 03, 2004 04:06 PM	Backup Killed	dumbo mickey grumpy
🟢	1078274657	000004	Tue Mar 02, 2004 04:44 PM	Successful	mickey dumbo
🟢	1078185622	000002	Mon Mar 01, 2004 04:00 PM	Successful	minnie grumpy

Select the specific backup job from which to restore by clicking on the button to the left of the desired job.

5. Next, if there are multiple backups on the selected media, a screen will appear with a list of backups. For disk backups, this list will contain all of the client backups in the selected job. For tape backups, the list will contain all of the backups, even those from different jobs. The information about the backup will be preceded by the ***backup sequence number***, starting with 1 and ending with the last backup on the media. The following is a sample of this screen:

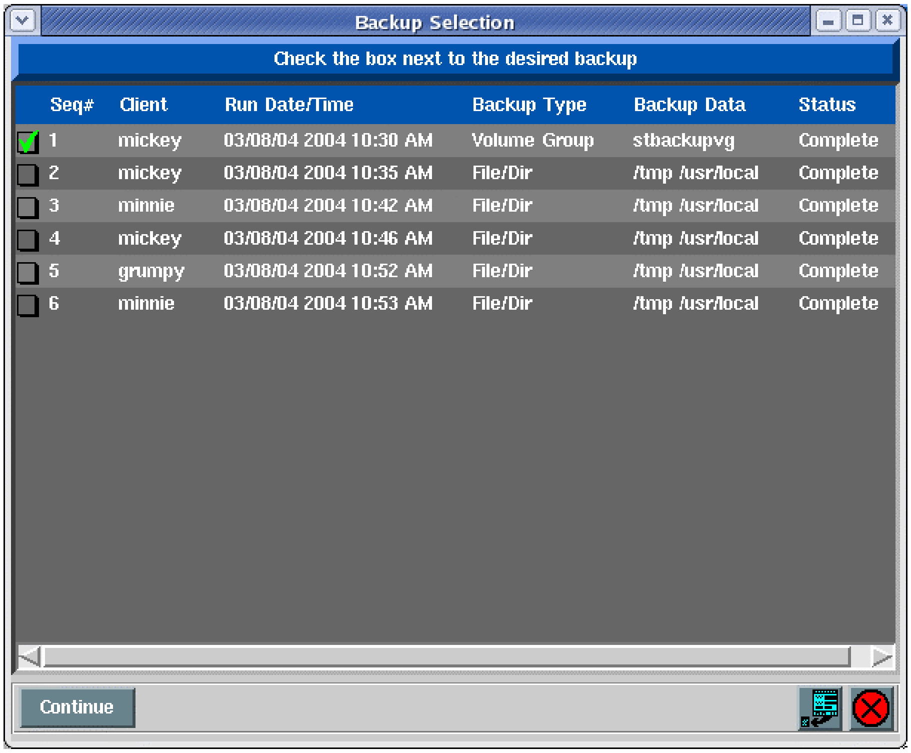

Check the box next to the desired backup

	Seq#	Client	Run Date/Time	Backup Type	Backup Data	Status
☑	1	mickey	03/08/04 2004 10:30 AM	Volume Group	stbackupvg	Complete
□	2	mickey	03/08/04 2004 10:35 AM	File/Dir	/tmp /usr/local	Complete
□	3	minnie	03/08/04 2004 10:42 AM	File/Dir	/tmp /usr/local	Complete
□	4	mickey	03/08/04 2004 10:46 AM	File/Dir	/tmp /usr/local	Complete
□	5	grumpy	03/08/04 2004 10:52 AM	File/Dir	/tmp /usr/local	Complete
□	6	minnie	03/08/04 2004 10:53 AM	File/Dir	/tmp /usr/local	Complete

You may select the backup from which to restore by clicking on the box to the left of the desired selection and a checkmark will appear in the box. Only one selection may be made. If you select a different option, the checkmark will be removed from the previous selection. After making your selection, click the Continue button at the bottom of the screen to continue.

Selecting Restore Options

The following screen will appear, which provides additional options for restoring data:

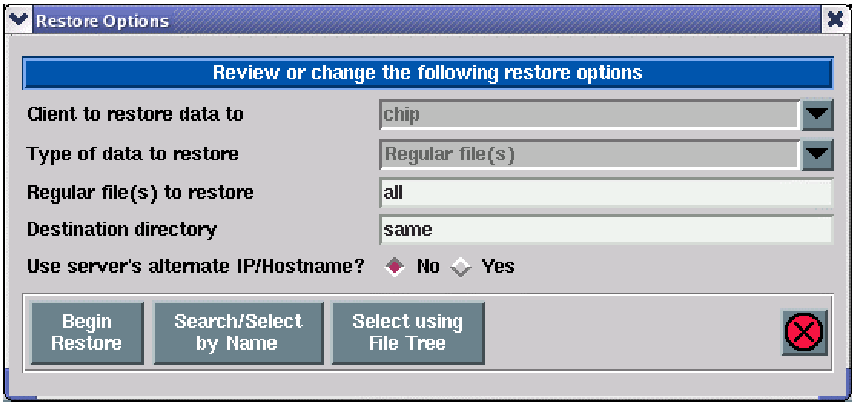

Review or change the following restore options

Client to restore data to	chip
Type of data to restore	Regular file(s)
Regular file(s) to restore	all
Destination directory	same
Use server's alternate IP/Hostname?	● No ○ Yes

Each of the fields is described in detail:

1. ***Client to restore data to***: This option appears only on *Network Administrator* systems. The client from which the backup originated will be displayed. If you wish to restore the data to a different client, press the arrow button to the right of the client name to display a list of clients and select from the list. If this is a disk backup (storied in a directory on the server) and the backup profile did not allow a different client to read the data, the client may not be changed.

2. ***Type of data to restore***: By default, the type of data to restore will equal the type of backup. However, it is possible to restore different types of data, including volume groups, logical volumes, filesystems, directories or individual files. To restore a different type of data than that shown, select the arrow button to display a list of restore data types allowed for this type of backup and select from the list.

3. ***[Data] to restore***: This label will indicate the ***restore data type*** selected in the previous field. You may type one or more options to restore (i.e. a list of volume groups if restoring volume groups), each option separated by spaces. You may also click on the arrow button to display a list of options to restore. If restoring files or directories, this list could be quite long. When selecting this button, a new window will appear from which you may select specific items to restore or search the list for specific patterns. Refer to Selecting Data to Restore below for details. You may also use wildcards in the names of files and directories to restore. Refer to Restoring Files or Directories Using Wildcards for details.

4. ***Destination [option]***: This label will show either directory or logical volume, depending on the restore data type. If restoring a logical volume (from a Logical Volume or Power System Backup), you may type the name of a different logical volume (which must already exist) to restore the data to. If restoring from any other backup type, you may select the directory into which the data will be restored. For more details on how the files will be restored to the new destination, refer to Restoring Data to a New Destination below.

5. ***Use alternate server IP/hostname***: This option appears only on *Network Administrators* and only if the selected backup server has an alternate IP/hostname configured. Select this button if you want to use an alternate network adapter to restore the data from the server. Refer to Using an Alternate Network below for more details.

When all desired selections have been made, press the Begin Restore button at the bottom of the screen to begin the restore.

Backup Types and Restore Data Types

As mentioned earlier, it is possible to restore various types of data, depending on the backup type. The table below indicates what type of data may be restored from each backup. Note that any type of data may be restored to a different compatible destination, even on a different client.

Backup Type	Restore Data Type(s)	Destination Type(s)
System Backup	Volume Groups Filesystems Directories Regular Files Logical Volumes Meta-disks (Linux) Partitions (Linux)	Volume Group Filesystem, Directory or SMB share Directory or SMB share Directory or SMB share Logical Volume Meta-disk Partition
Volume Group	Volume Groups Filesystems Directories Regular Files Logical Volumes	Volume Group Filesystem, Directory or SMB share Directory or SMB share Directory or SMB share Logical Volume
Filesystem	Filesystems Directories Regular Files	Filesystem, Directory or SMB share Directory or SMB share Directory or SMB share
Directory	Directories Regular Files	Directory or SMB share Directory or SMB share
Logical Volume	Logical Volume	Logical Volume
Meta-disk (Linux)	Meta-disk (Linux)	Meta-disk (Linux)
Partition (Linux)	Partition (Linux)	Partition (Linux)
SMB (Windows) Share	SMB Windows Share Directories Regular Files	Directory or SMB share Directory or SMB share Directory or SMB share

Restoring (Copying) Data Between AIX, Linux and Windows

As you can see from the above, even a backup of an ***SMB (Windows) Share*** may be restored to the same or different share, on the same or different client, or it may be restored to a directory on an ***AIX*** or ***Linux*** client! Likewise, Filesystem and Directory backups of ***AIX*** or ***Linux*** clients may also be restored to an ***SMB share*** on an *SMB (Windows) client*!

Selecting Data to Restore

There are different ways of selecting the data to restore from the Restore Options Screen, depending on the type of data being restored:

1. If you have selected to restore restoring Volume Groups, Filesystems, Logical Volumes, Meta-disks (Linux), Partitions (Linux), or SMB Shares (Windows), then an arrow button will appear to the right of the ***[Data] to Restore*** field. By pressing this button, the list of data items of the selected type will be read from the backup, and you can select one or more items from the list.

2. If you have selected to restore either Directories or Regular Files, the arrow button next to the [Data] to Restore field will disappear, and new buttons will appear at the bottom of the screen instead, labeled Search/Select by Name and Select using File Tree. Those options are explained in the next sections below.

3. Lastly, you may simply enter the data to restore in the field. You can enter one or more items, separated by spaces. If am item, such as a filename, contains spaces, you must enter that filename surrounded by quotes to preserve the space in the filename. Note that you can use wildcards (*) to restore multiple files with similar names or locations. Refer to Restoring Files Using Wildcards below.

Search/Select by Name

When restoring directories and regular files, you may press this button to view a complete list of files, select one or more files or directories from the list, select a group of files or directories, or search the list using a string or characters or wildcards (*). When pressing this button, a screen similar to the following will appear:

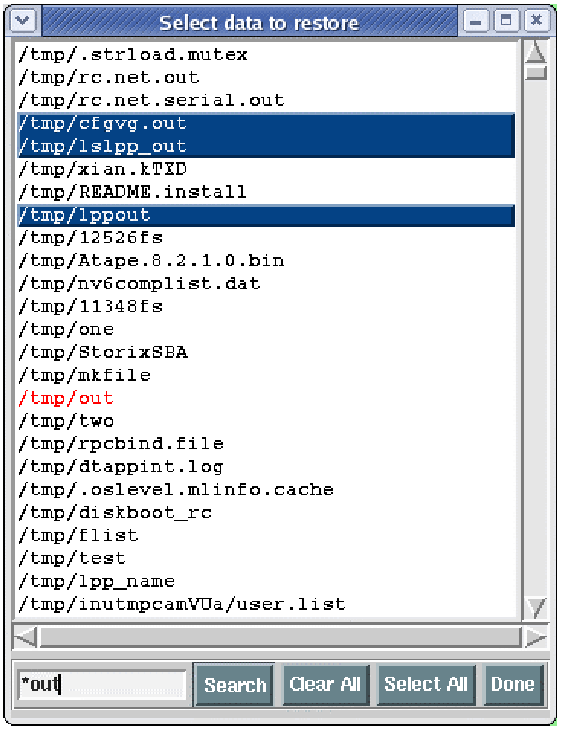

```
/tmp/.strload.mutex
/tmp/rc.net.out
/tmp/rc.net.serial.out
/tmp/cfgvg.out
/tmp/lslpp_out
/tmp/xian.kTID
/tmp/README.install
/tmp/lppout
/tmp/12526fs
/tmp/Atape.8.2.1.0.bin
/tmp/nv6complist.dat
/tmp/11348fs
/tmp/one
/tmp/StorixSBA
/tmp/mkfile
/tmp/out
/tmp/two
/tmp/rpcbind.file
/tmp/dtappint.log
/tmp/.oslevel.mlinfo.cache
/tmp/diskboot_rc
/tmp/flist
/tmp/test
/tmp/lpp_name
/tmp/inutmpcamVUa/user.list
```

*out

From this window, you may:

1. Click a specific entry to highlight and select that entry to restore.

2. Click and drag the mouse over a number of entries to highlight and select all those entries.

3. Click any highlighted entry to de-select that item to restore.

4. Enter a search pattern in the box at the lower-left corner of the window and press the Search button to find the next occurrence of that pattern. The next entry found that matches the search pattern will be highlighted in red. S search pattern can be any character string which may also include wildcard characters, or ***asterisks* (*)**. An asterisk in a search pattern may match any number of other characters in the list item.

5. Press the Clear All button to de-select any highlighted entries.

6. Press the Select All button to select “all” entries and return to the previous screen.

7. When all specific entries have been selected, press the Done button. You will be returned to the Restore Options Screen, and the selected list of files will appear in the ***[Data] to Restore*** field.

Select Using File Tree

Also, when selecting to restore regular files or directories, you can press this button to view a drop-down file-tree list of files or directories, and select from the list. When pressing this button, the backup media will be read, and a list of directories will appear, which may or may not contain the regular files, depending on which you selected to restore (viewing only directories will save much time and memory).

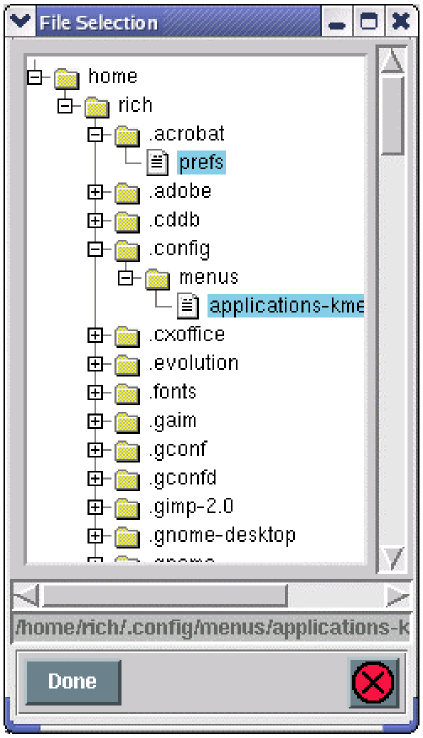

/home/rich/.config/menus/applications-k

From this window, you may:

1. Click on any folder or file icon to select that directory or file. Note that when selecting a directory, all files and directories beneath become un-selected as they will be restored automatically as part of their parent directory. Click a selected folder or file to de-select. Note that the full path of selected files or directories will appear in the box below the file tree.

2. Click the plus-sign (+) next to a directory to open the directory and view and select from the files or directories beneath. The plus (+) sign will turn to a minus (-) sign. Clicking the minus sign will close the directory, but any files or directories selected within will remain selected.

3. Double-click on a folder icon will open the folder just as pressing the plus (+) sign.

4. When you have selected all desired files and directories, press the Done button at the bottom of the screen. You will be returned to the Restore Options Screen, and the selected list of files will appear in the ***[Data] to Restore*** field.

Restoring Files or Directories Using Wildcards

There may be many files on a backup containing similar names that you want to restore without having to select each and every file, which may exist in different directories. To do so, you may use wildcards in the filenames. A wildcard is denoted by an ***asterisk* (*)** in one or more parts of the name. For example:

```
/home/*/*.gif
```

will result in all files in a sib-directory of the **/home** filesystem containing "**.gif**" at the end of the name. This would result in files such as:

```
/home/anthony/mom.gif
/home/michelle/candy.gif
```

but will NOT result in files such as:

```
/home/picture.gif
/home/anthony/myfiles/picture.gif
```

because these files are not in a single sub-directory of **/home** as indicated by the wildcard filename (**/home/*/*.gif**). To restore these files you would need to also include "**/home/*.gif**" and "**/home/*/*/*.gif**" in the list of files to restore.

To understand the use of wildcards in the restore, you need only understand how to list files on the system. Any files that are listed on the system when you type:

```
ls /home/*/*.gif
```

would be restored when using this same notation in the list of files to restore.

Restoring Data to a New Destination

When restoring files or directories from a System, Volume Group, Filesystem or File/Directory backup, you may enter a new destination directory in the ***Destination [option]*** field. When restoring a single logical volume, partition or meta-disk to restore, you may enter a new device name into which to restore the data.

If restoring a single filesystem, or specific files or directories from a System, Volume Group or Filesystem backup and you want to restore to a different directory, the files will be restored relative to the original filesystem mount point. For example, if you are restoring data from the **/data1** filesystem into the **/data2** directory, the **/data1/info/stuff** file will be restored to **/data2/info/stuff**.

If restoring multiple filesystems from a System, Volume Group or Filesystem backup, the files from each filesystem will be restored to different directories under the new destination directory. This is to protect against the same filename from different filesystems being restored to the same location. For example, when restoring the **/data1** and **/data2** filesystems to the **/datanew** directory, the files will be restored to **/datanew/data1** and **/datanew/data2** respectively.

If restoring from a File/Directory backup, the data will be restored relative to the file's full path name. For example, if restoring the **/data1/info/stuff** file to the **/data2** directory, the resulting file will be **/data2/data1/info/stuff**.

When restoring a single logical volume, the new logical volume name must already exist, may not currently be in use by any process, and must have been created at least as large as the original logical volume.

Using an Alternate Network to Restore from the Server

When using a *Network Administrator*, it may at time be desirable to have the client restore the data using a different network to communicate with the server than is used by default. For instance, if there are multiple networks available for reaching the server from the client, or if the client cannot communicate with the server using the default network (defined by the server's hostname and network routing configuration of the client), you can choose to restore using the alternate network.

For *AIX* SP Systems with High-Speed Switch networks, this is particular useful in allowing node data to be restored across the switch network from another node acting as the backup server. Refer to the SP System Information for additional information.

If an alternate network IP Address or hostname was defined for the server you are restoring from, an addition option will appear on the restore options screen above, "Use Alternate Server IP/Hostname?". If you want to use the alternate network to perform the restore, simply select "Yes" for this option. Note that this option will not appear if there was no alternate IP address or hostname setup for the server. To set the alternate IP address or hostname for a server, refer to the server configuration.

Displaying the Status and Output of the Restore

The restore will begin, and the status report screen, as shown below, will appear automatically. Listed on the screen will be a status line for the backup previously selected. Information pertaining to the progress and performance of the restore will be updated as the data from the backup is read. If the backup selected was not the first backup on the media, the process will need to ***fast-forward*** over the prior backups before reading the data. Fast-forwarding a tape backup is much faster than reading through all the data.

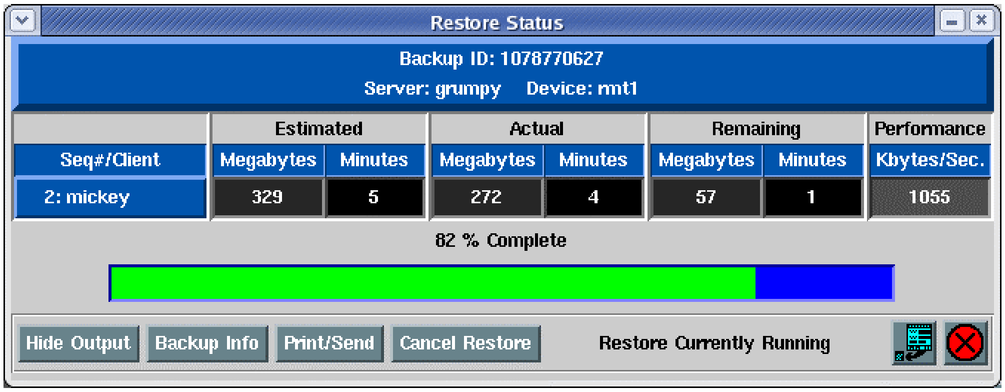

Restore Status

Backup ID: 1078770627

Server: grumpy Device: rmt1

	Estimated		Actual		Remaining		Performance
Seq#/Client	Megabytes	Minutes	Megabytes	Minutes	Megabytes	Minutes	Kbytes/Sec.
2: mickey	329	5	272	4	57	1	1055

82 % Complete

Restore Currently Running

Note that this screen may not be closed as long as the restore is running. It must remain on the screen after the restore completes, after which time it may be closed by pressing the cancel button. Once the screen is closed, the restore status and output messages may not be redisplayed.

To view the output of the restore process, press the Show Output button at the bottom of the screen. An output screen similar to the following will then appear, showing the output and status messages of the restore:

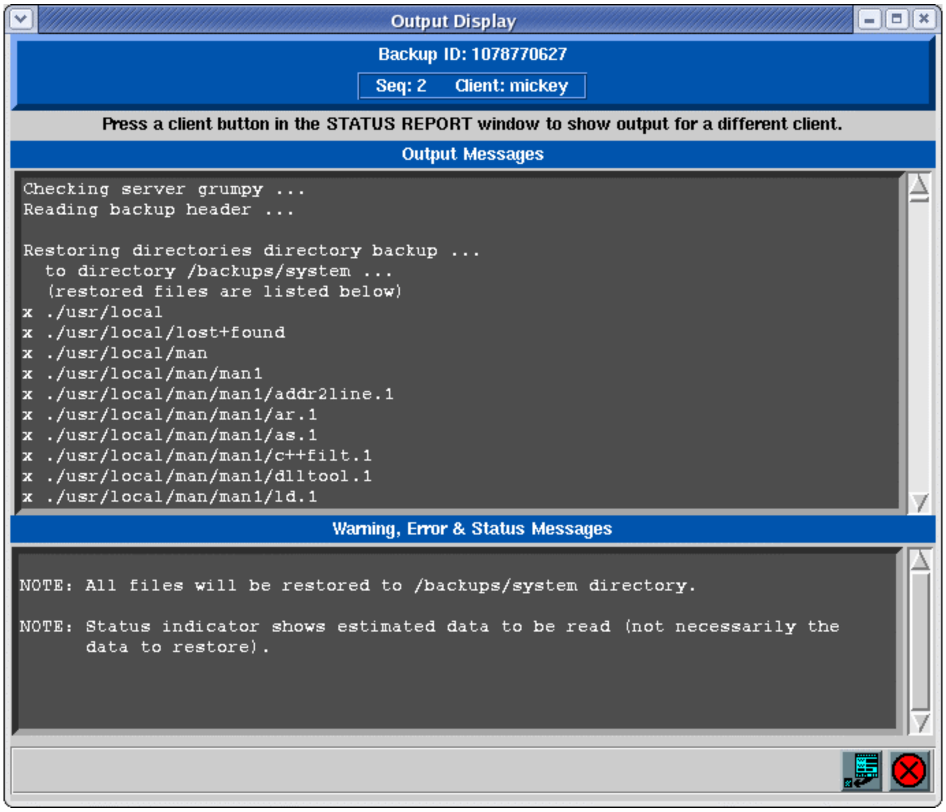

Backup ID: 1078770627

Seq: 2 Client: mickey

Press a client button in the STATUS REPORT window to show output for a different client.

```
Checking server grumpy ...
Reading backup header ...

Restoring directories directory backup ...
  to directory /backups/system ...
  (restored files are listed below)
x ./usr/local
x ./usr/local/lost+found
x ./usr/local/man
x ./usr/local/man/man1
x ./usr/local/man/man1/addr2line.1
x ./usr/local/man/man1/ar.1
x ./usr/local/man/man1/as.1
x ./usr/local/man/man1/c++filt.1
x ./usr/local/man/man1/dlltool.1
x ./usr/local/man/man1/ld.1
```

```
NOTE: All files will be restored to /backups/system directory.

NOTE: Status indicator shows estimated data to be read (not necessarily the
      data to restore).
```

If the restore is of any backup type containing filesystem data, the files will be listed on the screen as they are restored. For **Logical Volume**, **Partition** (***Linux***), **Meta-disk** (***Linux***) or **Power System Backup** (***AIX***), only one message is displayed as each raw device data is restored. This screen may be closed and redisplayed at any time, even after the restore completes, as long as the Restore Status Report screen has not been closed.

In addition to the restore output, summary information for the selected backup may be displayed by selecting the Backup Info button. A screen similar to the following example will appear. Simply press the Dismiss button to close this window.

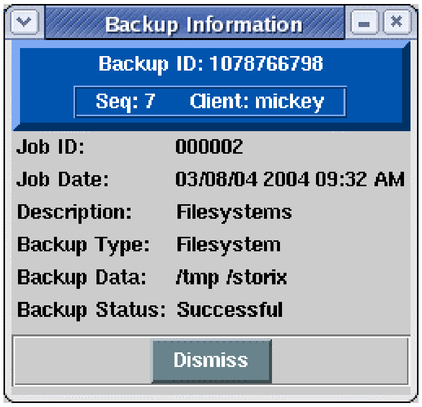

Backup ID: 1078766798

Seq: 7 Client: mickey

Job ID: 000002

Job Date: 03/08/04 2004 09:32 AM

Description: Filesystems

Backup Type: Filesystem

Backup Data: /tmp /storix

Backup Status: Successful

20. Copying Backups to Different Media

> NOTE **The features described in this chapter are not available when using Desktop Edition.**

This feature may be used to copy any backup from a local or remote (when using Network Administrator) system to any backup media (disk or tape) on the local or another remote system. When copying a backup, the data within the backup is unchanged, this providing you with two working copies from which to restore from.

Common uses

This option may be used to serve many purposes, for example:

1. **Backup *staging*** - Perform backups to a local disk, then offload the backup to tape. Backups to local disk often takes less time. If your data is unavailable to users during the backup “window”, this may reduce the downtime. The backup may later be copied to tape while users are back online since the backup data does not change when copied to new media.

2. **Copy backups to *offsite* server** – When complete, a local and remote copy of the backup will exist, increasing the availability of the system by keeping off-site backups. Backups over the network may also take longer, thus increasing the downtime of the local system if users cannot work during the backup process. When copying backups, much less system resource is used, and users may work without affecting the backup data.

3. ***Stacking* multiple backups onto tapes** - Multiple backups of the local system or different clients can be copied to the same tape device, thereby consolidating them all on the same backup “label’ (refer to the User Guide for an explanation of backup labels). The tape device may be local or remote. The destination backup may use multiple volumes when writing to tape, and multi-volume backups can be automated by using sequential autoloaders or random tape libraries.

To use this option, select Actions→Copy Backups to Different Media. When doing so, a screen similar to the following will appear:

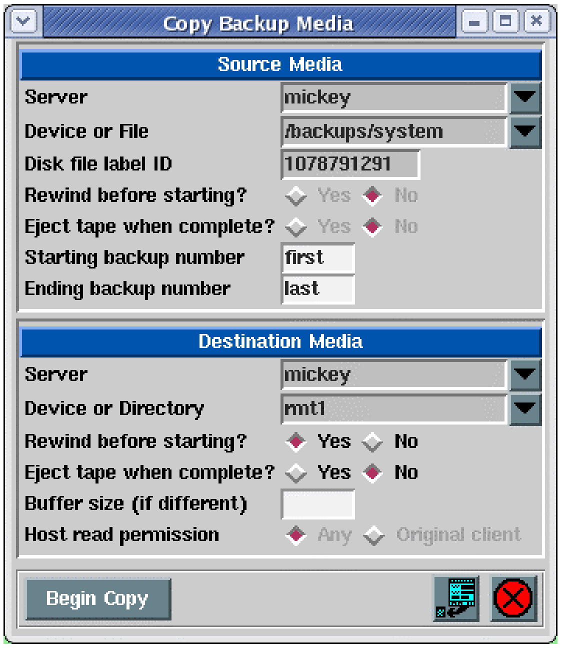

Source Media

Server	mickey
Device or File	/backups/system
Disk file label ID	1078791291
Rewind before starting?	○ Yes ● No
Eject tape when complete?	○ Yes ● No
Starting backup number	first
Ending backup number	last

Destination Media

Server	mickey
Device or Directory	rmt1
Rewind before starting?	● Yes ○ No
Eject tape when complete?	○ Yes ● No
Buffer size (if different)	
Host read permission	● Any ○ Original client

This screen is broken into two sections, one for the source backup and one for the destination media. If using *Network Administrator*, you must specify a source and destination server, otherwise, these fields do not appear.

Source Media

The source media may be any disk backup or tape backup. Use the arrow to the right of the Device or File entry field to select a device or directory to copy from. Only directories containing current backups will be shown. If you select a directory from the list, the backup labels which exist in that directory will be displayed such as in the following example:

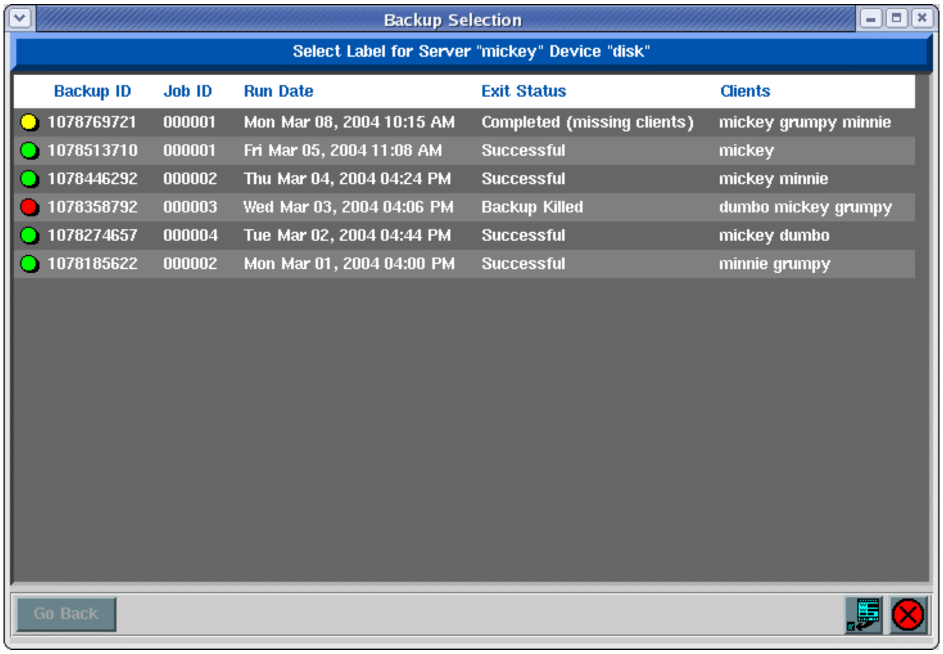

Select Label for Server "mickey" Device "disk"

	Backup ID	Job ID	Run Date	Exit Status	Clients
🟡	1078769721	000001	Mon Mar 08, 2004 10:15 AM	Completed (missing clients)	mickey grumpy minnie
🟢	1078513710	000001	Fri Mar 05, 2004 11:08 AM	Successful	mickey
🟢	1078446292	000002	Thu Mar 04, 2004 04:24 PM	Successful	mickey minnie
🔴	1078358792	000003	Wed Mar 03, 2004 04:06 PM	Backup Killed	dumbo mickey grumpy
🟢	1078274657	000004	Tue Mar 02, 2004 04:44 PM	Successful	mickey dumbo
🟢	1078185622	000002	Mon Mar 01, 2004 04:00 PM	Successful	minnie grumpy

If the backup contained multiple backup sequence numbers, you may select the starting and ending backup numbers to copy in the Starting backup number and Ending backup number fields. This is useful if, for instance, you created a backup of multiple clients, but want to copy only one client backup in the list to tape. Another example for *non-network administrator* systems, would be if you appended a daily backup to the same tape each day, but want to create new backup media which only contains one or more days from the tape.

> NOTE
>
> **You will not be able to use a virtual device configured as a random tape library as a source device. This is because the administrator is only capable of tracking volume changes to one random tape library at a time, and tape libraries are more likely to be used for destination devices. If you want to use a tape drive in a random tape library as a source device, use the tape device name instead of the virtual device name, and you will be prompted to change tapes, if required.**

If copying from tape, you may also indicate whether the source tape should be rewound before starting the backup and/or ***rewound*** and ***ejected*** at the end of the backup. If you select to copy a backup number which is prior to the current position of the tape, the tape will be automatically rewound and forwarded, if necessary, to the start of the backup number to copy.

Destination Media

Any backup may be copied from tape to disk directory, from one tape drive (or virtual device) to another, or from one directory to another. If using *Network Administrator*, the selected backup may be copied from any server to any other server (including the local system).

Stacking backups to tape

If copying to a tape device, you may indicate if you want to ***rewind*** before starting the backup and if the tape should be rewound and ***ejected*** at the end of the backup. If you do not rewind at the start of the backup, you may append the source backup to the end of the destination media (if the destination media is currently at the of volume. The destination backup label will be appended with the selected source backup(s).

You may also alter the buffer size of the backup by entering a buffer size (in Kbytes) in the Buffer size field. This is quite useful in increasing the performance of backups when writing to different media. For example, the default 64K buffer size may be adequate when you wrote your original disk backup file, but when copying to a high-speed tape drive, a higher buffer size (i.e. 256K) may provide much greater backup performance. To use the same buffer size for the destination as was used for the source, leave this field blank.

If using *Network Administrator*, and the destination backup is written to a disk directory, you may also change whether only the original client host or any host may read the backup data by making the appropriate selection in the Host read permission field. If not using *Network Administrator*, this field will not appear.

When your selections are complete, press the **Begin Copy** button. A dialog will appear asking you to confirm, and after doing so a **Copy Backup Status** screen similar to the following will be shown:

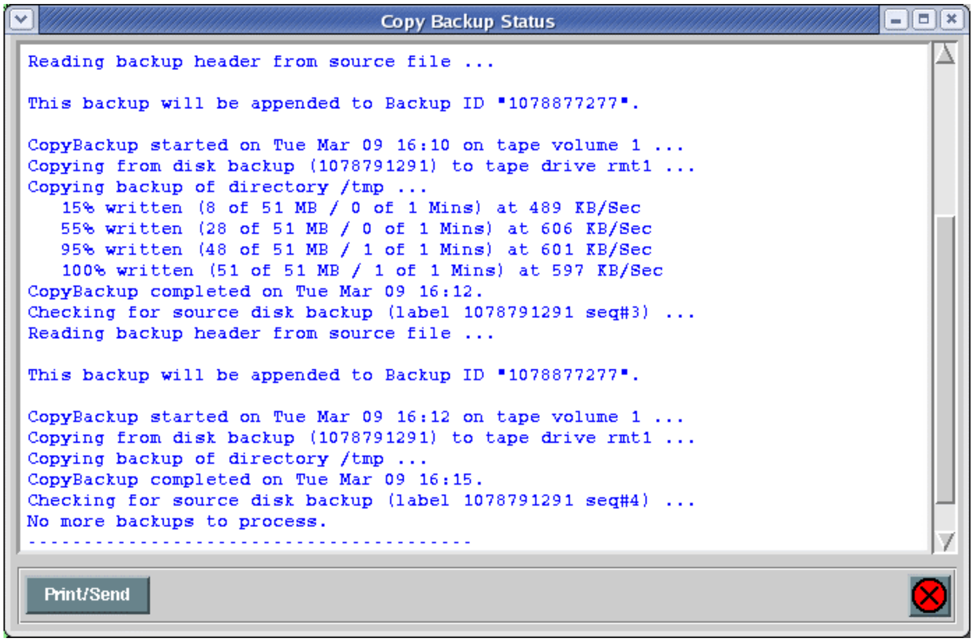

```
Reading backup header from source file ...

This backup will be appended to Backup ID "1078877277".

CopyBackup started on Tue Mar 09 16:10 on tape volume 1 ...
Copying from disk backup (1078791291) to tape drive rmt1 ...
Copying backup of directory /tmp ...
   15% written (8 of 51 MB / 0 of 1 Mins) at 489 KB/Sec
   55% written (28 of 51 MB / 0 of 1 Mins) at 606 KB/Sec
   95% written (48 of 51 MB / 1 of 1 Mins) at 601 KB/Sec
   100% written (51 of 51 MB / 1 of 1 Mins) at 597 KB/Sec
CopyBackup completed on Tue Mar 09 16:12.
Checking for source disk backup (label 1078791291 seq#3) ...
Reading backup header from source file ...

This backup will be appended to Backup ID "1078877277".

CopyBackup started on Tue Mar 09 16:12 on tape volume 1 ...
Copying from disk backup (1078791291) to tape drive rmt1 ...
Copying backup of directory /tmp ...
CopyBackup completed on Tue Mar 09 16:15.
Checking for source disk backup (label 1078791291 seq#4) ...
No more backups to process.
----------------------------------------
```

Canceling the Operation

You may not close this window. However, if you return the prior Copy Backup Media screen, you may press the Cancel button in that window to terminate the operation.

21. User Preferences

In this section, options that affect the overall operation or appearance of the application are discussed. To change the user preferences, select File→Preferences on the menu bar.

Software License

This option may be used to display a screen used to reconfigure your Administrator license or add or change Optional Features. To view or change the license information, select File→Preferences→Software License from the menu bar. A screen similar to the following will appear:

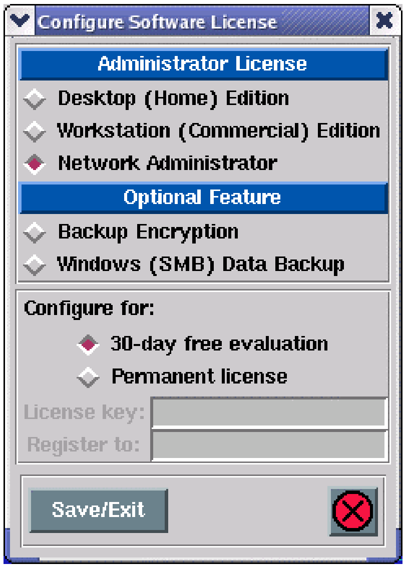

Configure Software License

Administrator License

○ Desktop (Home) Edition
○ Workstation (Commercial) Edition
● Network Administrator

Optional Feature

○ Backup Encryption
○ Windows (SMB) Data Backup

Configure for:

● 30-day free evaluation
○ Permanent license

License key:
Register to:

It may be desirable to change the license for a number of reasons:

1. You installed one license type and now want to change to another, such as to take advantage of additional features of another administrator license.

2. You installed an evaluation (trial) license, and upon expiration, want to now installed a permanent (purchased) license key.

3. You purchased a Network Administrator for 10 clients, and now want to add support for another 6 clients.

Administrator License

When the software is initially installed, you indicated the type of administrator you would be installing. This was either ***Desktop***, ***Workstation*** or ***Network Administrator***. You also entered a license key, which matched the administrator license type, and also indicated the number of clients supported by a Network Administrator and the expiration time (if any) of the license.

Optional Features

Also, there are additional features available that require their own software license:

1. **Backup Encryption**: The license key will indicate the number of clients that are to support backup data encryption. When installing this license, you will be able to apply data encryption to the

number of clients the license supports. If you have a Workstation Edition license installed, a Backup Encryption license is a single license for the local system.

2. **Windows (SMB) Data Backup**: The license key will indicate the number of ***SMB clients*** (i.e. ***Windows*** or ***Mac OS X***) that the application will support. This option is only available when you have installed a *Network Administrator* license.

When entering this function, your current **Administrator License** information is displayed. If you select a different **Optional Feature** from the list, the license information for that optional feature is displayed.

Any Administrator license or optional feature may be installed for a trial period, but may only be configured for a trial period one time. After that, a permanent license key is required for the administrator and each optional feature, which may be obtained from ***Storix***.

After selecting either the **Administrator License** or the **Optional Feature**, you may add or change the license information. This includes changing from a trial to a permanent license (or vice-versa), entering a permanent license key, and the name of the person or company the product is registered to (also provided by ***Storix*** and must match the license key).

when you have completed all entries, press the Save/Exit button. The software will be reconfigured. In some cases, the administrator (***sbadmin***) must be terminated and restarted for the changes to appear.

Operating Systems Support

> NOTE **This option will not appear if using Workstation or Desktop Edition, and the only operating system-specific options which will appear in the application will be those applicable to the local operating system.**

By default, when the Backup Administrator software is first installed, only support for clients running the same operating system as the *Network Administrator* system is enabled. If, for instance, the *Network Administrator* is running ***AIX***, the only options that appear in the application will be applicable to ***AIX*** systems. If, however, this ***AIX*** *Network Administrator* will be supporting ***Linux*** clients, then you will want to add into the application those options that are applicable to ***Linux*** systems also. You may later turn off support for client operating systems that will not be managed by the *Network Administrator*.

For example, only ***AIX*** systems currently support Split-Mirror Backups, so by turning off AIX support, this option will no longer appear on the menu bar. In another example, only ***Linux*** systems support Partition and meta-disk backups. By turning off Linux support, the Partition and Meta-disk backup types will no longer appear when configuring backup profiles.

To change the default settings as described above, select File→ Preferences → Operating Systems Support on the menu bar. The following screen will appear:

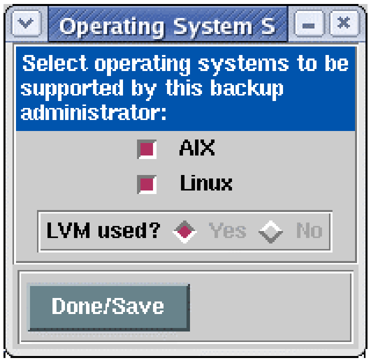

Operating System S

Select operating systems to be supported by this backup administrator:

☑ AIX

☑ Linux

LVM used? ◆ Yes ◇ No

To enable or disable support for a particular operating system, simply select or de-select the button next to the corresponding operating system type.

The **LVM used?** option is only available if you do not have ***AIX*** support enabled (since LVM is always assumed when AIX is used). If you are using only ***Linux***, you may select whether or not LVM options (such as *volume group* and *logical volume* backups, *snapshot LV backups*, etc) should be made available.

> **If you are not using LVM (Logical Volume Manager) on Linux, you should consider doing so since it provides much more flexibility in disk management than partition-based filesystems.**

When finished, press the **Done/Save** button.

Backup Retention Policy

The backup retention policy (also referred to as the ***overwrite policy***) determines whether or not a new backup should be allowed to write over (thereby destroying) a current backup. A current backup is defined as one with a label currently on record. The default policy for tape backups is to prevent accidental overwriting by requiring the user to manually ***expire*** a current backup before the same tape may be reused. The default policy regarding disk backups is to keep all disk backups on file unless explicitly expired (and removed) by the user.

This option allows you to define the ***global*** backup retention policy. This will apply to all backup jobs unless you explicitly change the backup retention policy for a particular backup job. Refer to Changing the Backup Retention Period in the Configuring a Backup Job section for more details on overriding the global backup retention policy.

To change the default settings as described above, select File→ Preferences → Backup Retention Policy on the menu bar. The following screen will appear:

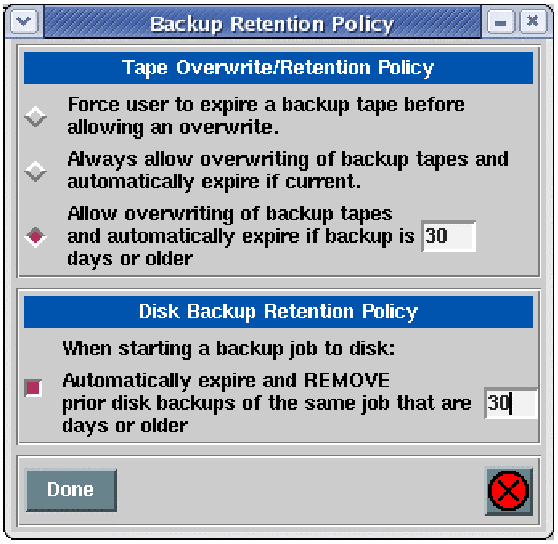

Backup Retention Policy

Tape Overwrite/Retention Policy

◇ Force user to expire a backup tape before allowing an overwrite.

◇ Always allow overwriting of backup tapes and automatically expire if current.

◆ Allow overwriting of backup tapes and automatically expire if backup is 30 days or older

Disk Backup Retention Policy

When starting a backup job to disk:

☑ Automatically expire and REMOVE prior disk backups of the same job that are 30 days or older

Tape Backups

Before any backup is performed to tape, the backup label is read. If a prior backup exists, a check is made to see if the backup label is still on file. If so, the setting applied here will determine the action that will be taken:

1. If the option "***Force user to expire a backup tape before allowing an overwrite***" is selected, no data will be written to the tape and the backup will fail with an error message. In order to overwrite the backup, the user must expire the backup manually. Refer to Expiring a Backup for the steps to expire a backup.

2. If the option "***Allow overwriting of backup tapes and automatically expire if current***" is selected, the backup will proceed to the tape after all record of the prior backup is removed from the system (backup is expired). A message is sent to the *root user*'s mail indicating that the prior backup has been expired.

3. If the option "***Allow overwriting of backup tapes and automatically expire if backup is __ days or older***" is selected, the backup will proceed ONLY if the backup to be overwritten is at least the specified number of days old. If the current backup is less than the specified number of days, the backup will not be overwritten and the user will be required to expire the backup manually before proceeding. Refer to Expiring a Backup for the steps to manually expire a backup.

Disk Backups

Note that backups written to disk are never overwritten by a new backup since each backup has a different filename. A new copy of a backup will be written to the server's disk each time a job is re-run. This option is used to free disk space on the server by automatically removing old backups when new backups are started.

A single checkbox is provided which, when checked, will cause the automatic removal of disk backups when the same job is run again. You must specify the number of days the disk backups should be retained. If, for example, you enter a value of "30" days, then when backup job "000001" is run, any prior disk backups created by job 000001 that are over 30 days old will be expired and automatically removed from the disk. If you select a value of "0", then the previous disk backup will be remove every time the same job is run again.

If this checkbox is not selected, then no disk backups will be removed when a new backup is started. In this case, you will need to manually expire all disk backups you do not wish to retain in order to free up disk space on the backup server. Refer to Expiring a Backup for the steps to manually expire a backup.

Press the Done button to save your selection and exit this function.

Backup Status Reporting

This feature is not available for Desktop Edition. The defaults shown in the sample screen below will be used with Desktop Edition.

Because scheduled backups may be running even when the Backup Administrator application (or even Xwindows) is not running, it is necessary to provide a method by which the system administrator is informed of the status of backups. These status messages include indications of when backup jobs are started and completed, as well as any errors or warning messages that occur prior to, during, or after the completion of a backup. The messages will also include notifications of the automatic expiration and overwriting of prior backups (as determined by the Backup Overwrite/Retention Policy).

By default, the messages will appear only on the screen, if available, and if the screen is not available (Backup Administrator is not running), the messages will be sent to the *root user*'s mail. This option will allow you to change the default method of notification.

To change the settings, select File→Preferences→Backup Status Reporting on the menu bar. The following screen will appear:

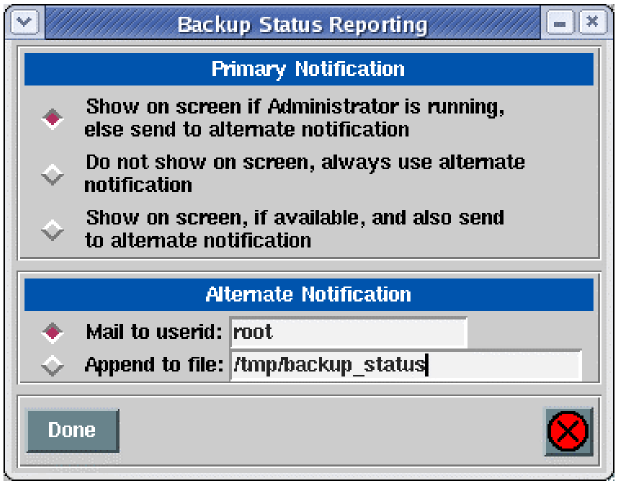

Backup Status Reporting

Primary Notification

◆ Show on screen if Administrator is running, else send to alternate notification

◇ Do not show on screen, always use alternate notification

◇ Show on screen, if available, and also send to alternate notification

Alternate Notification

◆ Mail to userid: root

◇ Append to file: /tmp/backup_status

Primary Notification

One of the options in this section must be selected to determine where backup status messages should be reported:

1. By default, the option "***Show on Screen if Administrator is running, else send to alternate notification***" is selected. If this option is used, messages will be reported on the screen if the "*sbadmin*" program is running. If not, the messages will be reported using the alternate notification method indicated in the next section.

2. The option "***Do not show on screen, always use alternate notification***" may be selected if you do not want messages reported on the screen. In this case they will always be reported using the alternate notification method. Selecting this option is equivalent to using the first option when the Administrator is not running.

3. If you want messages always reported on the screen (when Administrator is running) and also sent using the alternate notification method, select the option "***Send to screen, if available, and also use alternate notification***".

Alternate Notification

In this section, you will select how backup status messages should be handled when the alternate notification method is used. The alternate notification method will be used any time the Administrator is not running (cannot be displayed on the screen) and/or when the second or third options of the Primary Notification are selected.

By default, messages will be sent to the *root* user's mail when the alternate notification is used. If you want a different user to receive the mail messages, select "***Mail to userid***", then enter the user id in the corresponding entry box. The user ID entered may be a local user (i.e. "mary") or a user on another host (i.e. "scooter@adminsys").

If, rather than sending mail, you want messages to be appended to a text file on the admin system, select the option "***Append to file***" and enter the name of the file in the corresponding entry box. If the file does not already exist, it will be created when the first message is written. If it already exists, messages will be appended to the bottom of the file. Messages in this file will look similar to the following:

```
---------------------------------------------------------------
SBA JOBSTART: 000003
June 17 16:50:04 PDT 1999
Job 000003 has been started.
   Backup Device: rmt1
   Backup Server: spiderman
---------------------------------------------------------------
SBA JOBERR: 000003
June 17 16:53:44 PDT 1999
Job 000003 cannot be written to the tape.
   Backup Device: rmt1
   Backup Server: spiderman
   Error Message: The tape currently in the drive contains a current backup label
(929523610). The overwrite policy does not allow overwriting of this backup. Please
either expire this backup or change the overwrite policy to allow overwriting of
current backups.

The queue has been shut down. You must either restart the queue to re-run the entire
job, or delete the job from the queue.
---------------------------------------------------------------
SBA JOBSTART: 000003
June 17 16:58:21 PDT 1999
Job 000003 has been started.
   Backup Device: rmt1
   Backup Server: spiderman
---------------------------------------------------------------
SBA JOBOK: 000003
June 17 17:45:12 PDT 1999
Job 000003 completed successfully.
   Backup Device: rmt1
   Backup Server: spiderman
   Backup ID: 929665124
```

For easier identification of important messages, all messages in this file contain a header indicating the message type. These include:

ERROR	general error
INFO	general info
VOLCHG	tape volume change requested
JOBSTART	a job has started
JOBOK	a job completed successfully
JOBWARN	a job completed successfully with a warning message
JOBERR	a job terminated with an error

Additional message types may also be used. As a dashed line separates all messages in the file, the following command may be used to extract all job error messages:

```
grep -p^- JOBERR /tmp/backups
```

The message types indicated above (i.e. "JOBERR") will also appear in the subject line of mail messages if mail is used as the alternate notification method.

Server/Device Error Handling

This feature is not available for Desktop Edition. The defaults shown in the sample screen below will be used with Desktop Edition.

When a job starts, either through the scheduler or manually, the software will verify that data can be written to the specified server or device. If unavailable or ***write-protected***, by default, a notification message will be sent using the Backup Status Reporting Policy , the job will fail and the queue will be shutdown. This prevents subsequent jobs using that device from failing and essential places those jobs in a waiting status until the failed job is removed from the queue. This option allows you to change this default behavior.

To view or change the error handling setting, select File→Preferences→Device Error handling from the menu bar. When you do so, a **Device Error Handling screen** similar to the following example will appear:

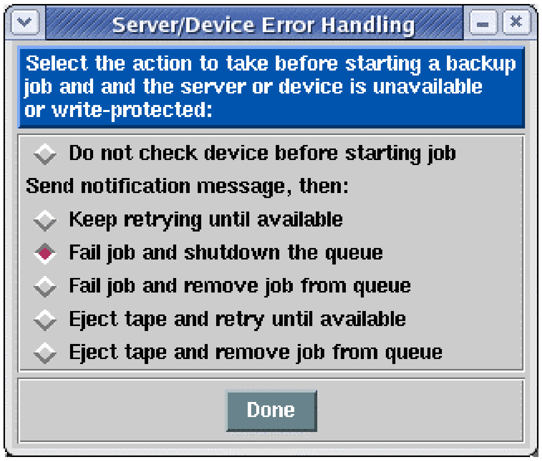

Server/Device Error Handling

Select the action to take before starting a backup job and and the server or device is unavailable or write-protected:

○ Do not check device before starting job

Send notification message, then:

○ Keep retrying until available
● Fail job and shutdown the queue
○ Fail job and remove job from queue
○ Eject tape and retry until available
○ Eject tape and remove job from queue

By default, the backup server and device availability is checked before a backup job is started. The first option, **Do not check device before starting job**, allows a backup to begin without first checking the availability. This may be preferable, for instance, when you have a pre-backup program which initializes the backup device or inserts a tape. This way, the device will be made available by the backup process and not checked for availability before starting it.

The additional options indicate what should happen if the server or device is unavailable when the job is pre-processed. Use the QuickHelp feature to obtain a detailed description of each option.

Select the radio button that best fits your needs or environment. Once you have made your selection press Done and this new error handling behavior will be applied for all devices and servers.

Report Preferences

Reports and notifications can be printed, sent to an email address, or appended to a text file. This option is used to set up preferences for each.

NOTE **The *Email* and *File* options described in this section are not available for Desktop Edition.**

To edit configure or change these options, select Report Preferences from the File→Preferences menu. A screen will appear similar to the following example:

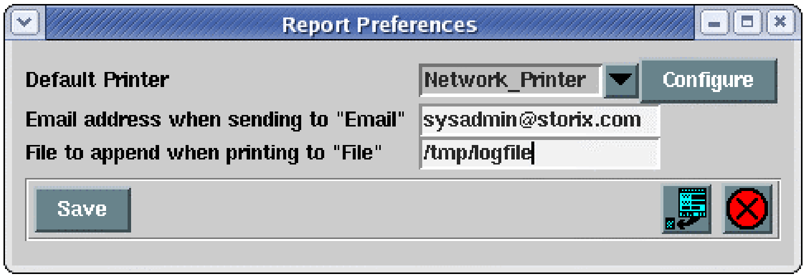

Report Preferences

Default Printer: Network_Printer

Email address when sending to "Email": sysadmin@storix.com

File to append when printing to "File": /tmp/logfile

The option for setting up a printer to use with SBAdmin varies between ***AIX*** and ***Linux*** systems. The above example is from a Linux system. An AIX system will not have a Configure button. Instructions for each are described separately:

Default Printer (*AIX*)

It is assumed that AIX systems will always have AIX printer *queues* already setup. This option is used to select the default printer queue which will be automatically selected when using any of the "Print/Send" options within the application. The printer queue must already have been set up on the system. To select the default printer to use, select the printer queue from the pull down.

Default Printer (*Linux*)

Linux systems provide a variety of ways to configure printers and supply numerous commands that may be used to submit files or jobs to the printers or queues. Therefore, this option will allow you to select a printer definition along with the command used to send data to the printer.

You should first configure your printer or printer queue using your Linux system administration utilities. Be sure to test the command by typing it at the command line to send something to the printer before adding the command to the SBAdmin Preferences..

To make a print option available in the SBAdmin report options, select from the pull down menu and select a printer from the list, or press Configure to setup a printer that has not already been configured. When you press the Configure button, a screen will appear like the example below:

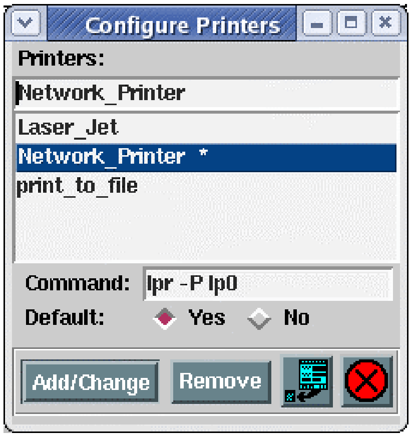

Printers: Network_Printer

Laser_Jet
Network_Printer *
print_to_file

Command: lpr -P lp0
Default: ● Yes ○ No

Enter the name of your printer in the box at the top of the screen. Note that this may be any name you choose, not necessarily the name of the printer queue as defined to Linux. The name you choose will be presented when you select a printer from any of the other SBAdmin printer list options.

In the **Command** box, enter the command used to submit a job to this printer. The name of the file (which is temporarily generated by SBAdmin) will be added to the END of this command.

Press Add/Change to add this printer to the list.

To remove a currently defined printer, select it in the listbox, then press the Remove button.

Sending Reports to an Email Address

This option is used to designate the email address used when **Email** is chosen for the "Print/Send" options within the application. To set or change the email address, enter the email address into the text field labeled: Email address when sending to “Email”. After doing so, an “**Email**” option will appear when selecting to Print/Send any SBAdmin report.

Appending Reports to a File

This option is used to designate the path to the file that is appended to when **File** is chosen for the “Print/Send” options within the application. To change the file path, enter the fill pathname of the file into the text field labeled: File to append when printing to “File”. After doing so, a “**File**” option will appear when selecting to Print/Send any SBAdmin report. If the specified file does not exist when printing to the **File** option, the file will be automatically created. Any parent directories of the file must already exist. Press the Save button when finished.

Fonts & Colors

The font, font size, and colors used by the graphical user interface may be changed to suit your preferences. The selections made will apply to all screens within this application. Several color “themes” are available. To change these preferences, select Fonts & Colors from the File➔Preferences menu. The following screen will appear:

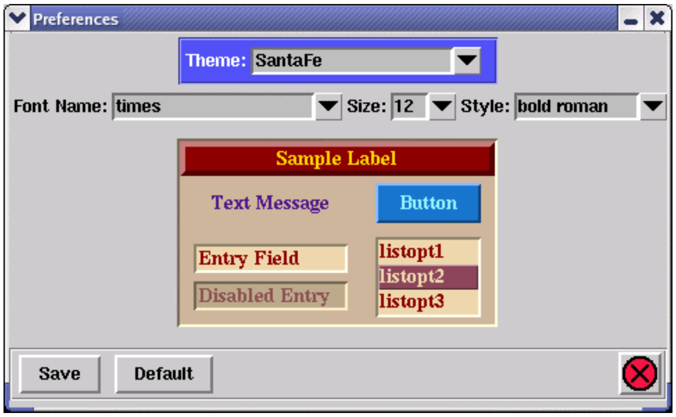
Preferences

Theme: SantaFe

Font Name: times Size: 12 Style: bold roman

This screen will always be displayed using the ***Classic*** color theme, even when another theme has been applied to other screens. To display a different color theme in the **Sample** section of the display, select one of the following from the Theme drop-down list:

Aquamarine

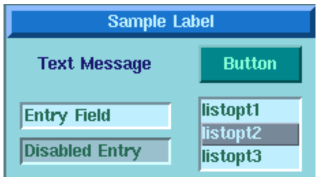

Carnival

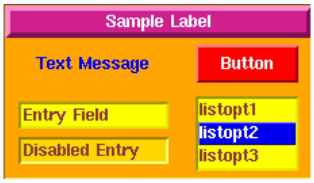

GolfCourse

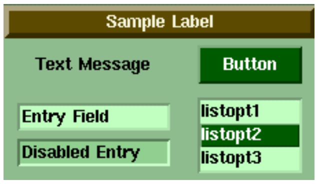

MardiGras

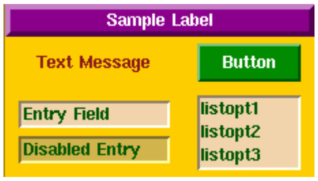

PowderPuff

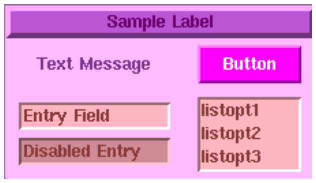

SantaFe

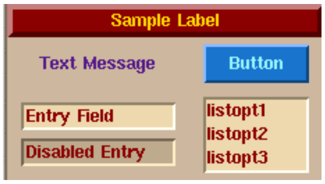

Winter

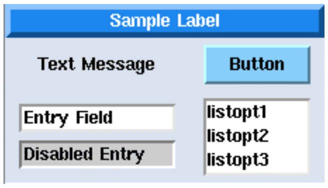

To change the Font Name, Size or Style, click the arrow button next to the desired selection, then select an option from the list. To preview your selections before saving them, press the Preview button. The sample box will be changed to show your selections.

Once you're satisfied with your selections, press the **Save** button to save the changes. If, after having previously saved difference sections, you want to return the screen to the default font and colors used when the application is first installed, select "***Classic***" from the Theme drop-down, then **Save** the settings again.

As soon as you save your settings, a confirmation dialog box will appear, and when selecting to continue, all windows except the Main Screen will be closed. The Main Screen will then be updated to reflect your selection. All windows opened from this point will display the selected settings.

Sound On/Off

This simple option will allow you to select whether or not you wish to bear a "beep" whenever the Backup Administrator reports a message on the screen that requires attention. To do so, select File→Preferences→Sound On/Off from the menu bar. You may then select "Yes" or "No" indicating whether or not the bell should ring.

Network Interface

This option is only available when using a Network Administrator.

By default, the admin system will use the network adapter associated with the default *hostname* of the system. If the system has multiple network adapters, you may choose a different adapter to use when sending data to and from clients or servers by selecting a different *hostname* or *IP address* associated with the desired network adapter. To do so, select File→Preferences→Network→Network Interface from the menu bar. After selecting this option, a small screen will appear where you may enter the hostname or IP address. If you want the system to go back to using the default adapter, simply remove any prior entry.

The network adapter selected will be used to pass information between the admin system and clients or backup servers, such as backup status messages, command output, and for polling the system availability. It is NOT used to pass the actual backup data, which is sent directly from the clients to the backup servers <u>even if the admin system is the backup server</u>. To configure the network interface used for sending backup data, refer to the <u>alternate hostname</u> options in the server configuration.

Important! After setting this value, some or all of the clients or backup servers may show as unavailable when the <u>Clients, Servers and Devices</u> are displayed on the Main Screen. If this should occur, it means that the client or server does not have the admin system defined using the alternate network adapter. To resolve this problem, do one of the following:

On *AIX* systems: Type "`smit storix`" on each client or backup server and select Set or Change Network Administrator. Set the Network Administrator Hostname or IP Address field to the hostname or IP address of the admin system as known by the client using the selected network interface.

On *Linux* systems: Edit the ***/storix*/config/admin_servers** file (where ***/storix*** is the data directory you chose when you installed the software), and either change the existing admin system hostname, or add the new admin system hostname on a line by itself.

Network Timeout

This option is only available when using a Network Administrator.

For a Network Administrator to perform any operations on a client or server, from querying it's availability to starting a backup job, it must execute a remote command. By default, if the admin system cannot contact the client or server within 10 seconds, it is assumed that the client or server is unavailable. This is adequate in most cases. However, if your network is slow to respond, perhaps due to slow hostname resolution, you may need to increase this value. To do so, select File→Preferences→Network→Network Timeout from the menu bar. A screen similar to the following will appear:

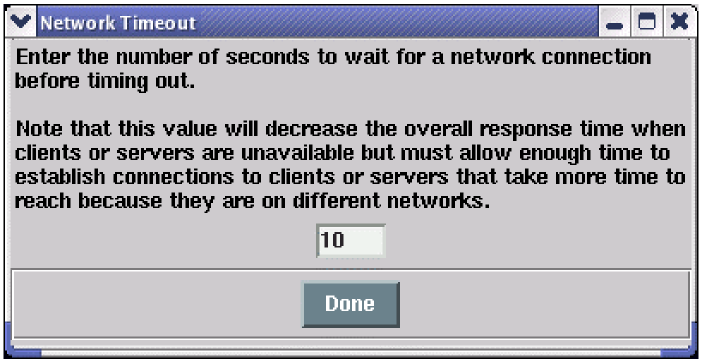

Enter the number of seconds to wait for a network connection before timing out.

Note that this value will decrease the overall response time when clients or servers are unavailable but must allow enough time to establish connections to clients or servers that take more time to reach because they are on different networks.

10

To change the value, simply enter the new number of seconds in the box provided. As indicated on the screen, when increasing the value, it will take longer to determine that a system is unavailable, so many user interface updates may take longer. It is not advisable to increase this value to more than 30 seconds, depending on the total number of clients and servers configured.

Press Done when you have made your selection.

Changing this option will affect only the timeout value when the Network Administrator contacts the clients or servers. It has no affect on the default (10-second) timeout the clients use to contact the server when running backups. To change the default timeout on the client, edit the */.stdefaults* file on the client and change the *SOCK_TIMEOUT* value to the desired number of seconds.

22. Reports

When you select Reports from the menu bar, you may further select from a list of reports that may be viewed and/or printed. Refer to Report Preferences for information on configuring printers, files and email addresses for reporting. Each time a report option is selected, a screen similar to the following will be presented:

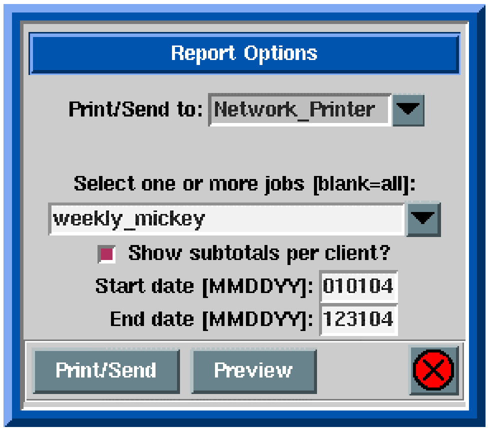

Print/Send to: Network_Printer

Select one or more jobs [blank=all]: weekly_mickey

■ Show subtotals per client?

Start date [MMDDYY]: 010104

End date [MMDDYY]: 123104

The above example is used when print a "**Backup Job History**" report. The option at the top, Printer queue, as well as the Print/Send and Preview buttons at the bottom of the screen are provided for all report options. The other options will vary for each report option selected.

You may select the Print/Send button to generate the report and send it directly to the specified printer, file or email address, or you may use the Preview button to generate the report in a window such as the following example (**Backup Job History** report):

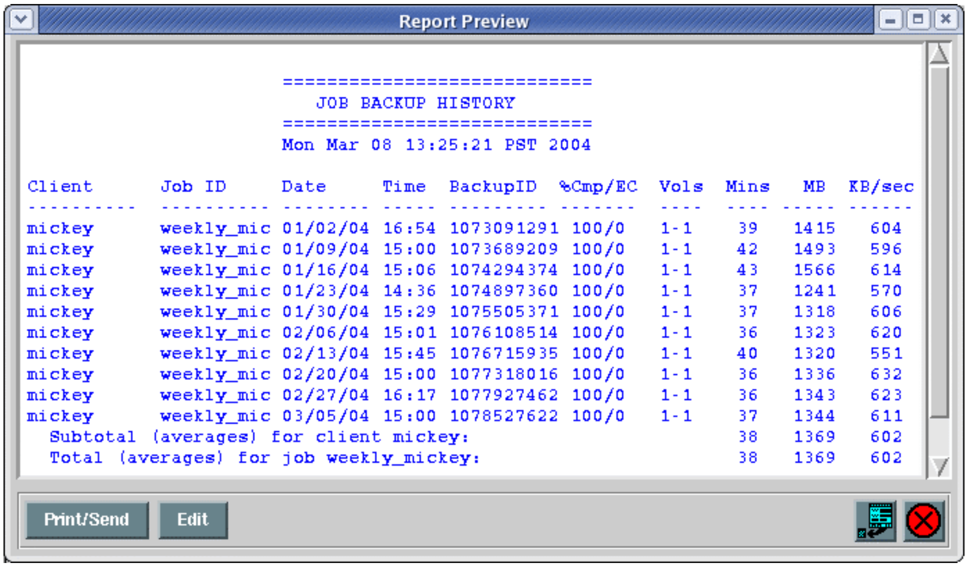

```
                         ============================
                            JOB BACKUP HISTORY
                         ============================
                         Mon Mar 08 13:25:21 PST 2004

Client      Job ID     Date     Time  BackupID   %Cmp/EC Vols Mins  MB   KB/sec
----------  ---------- -------- ----- ---------- ------- ---- ---- ----- ------
mickey      weekly_mic 01/02/04 16:54 1073091291 100/0   1-1   39  1415   604
mickey      weekly_mic 01/09/04 15:00 1073689209 100/0   1-1   42  1493   596
mickey      weekly_mic 01/16/04 15:06 1074294374 100/0   1-1   43  1566   614
mickey      weekly_mic 01/23/04 14:36 1074897360 100/0   1-1   37  1241   570
mickey      weekly_mic 01/30/04 15:29 1075505371 100/0   1-1   37  1318   606
mickey      weekly_mic 02/06/04 15:01 1076108514 100/0   1-1   36  1323   620
mickey      weekly_mic 02/13/04 15:45 1076715935 100/0   1-1   40  1320   551
mickey      weekly_mic 02/20/04 15:00 1077318016 100/0   1-1   36  1336   632
mickey      weekly_mic 02/27/04 16:17 1077927462 100/0   1-1   36  1343   623
mickey      weekly_mic 03/05/04 15:00 1078527622 100/0   1-1   37  1344   611
  Subtotal (averages) for client mickey:                        38  1369   602
  Total (averages) for job weekly_mickey:                       38  1369   602
```

From the **Preview** window you may scroll up and down the report, then print or send the report by selecting the Print/Send button, or you may also edit the contents of the report in order to add your own comments. To edit the report, select the Edit button. The color of the text will change and you will be allowed to click-on and make changes to the text. The Edit button will change to "Save" which, when pressed, will save the changes and disable the editor.

In the remainder of this section, a brief description of each report option is provided.

Clients & Servers

This option is only available when using a Network Administrator.

Select Reports→Clients & Servers to print a list of the clients and servers configured on the system. Refer to the main Reports section above for details on the Print and Preview options. When selecting this option, an additional option is provided:

- **Show Server's virtual device details**: Check this box if you want to also show a list of the virtual device configuration, if any, for each server listed.

Backup Profiles

Select Reports→Backup Profiles to print a list of the profiles configured on the system. Refer to the main Reports section above for details on the Print and Preview options. When selecting this option, an additional option is provided:

- **Include customized job profiles**: This option is not available if using *Desktop Edition*. Check this box if you want to print a list of the profiles that have been customized for particular jobs. If not checked, only the original job profiles will be included.

Exclude Lists

Select Reports→Exclude Lists to display or print a list of configured exclude lists. Refer to the main Reports section above for details on the Print and Preview options. When selecting this option, no additional options are provided. The list will contain each exclude list name, along with a list of files, directories and devices which are excluded, and the list of clients (or "all) that the exclude list applies to.

Backup Jobs

Select Reports→Backup Jobs to print a list of the backup jobs configured on the system. Refer to the main Reports section above for details on the Print and Preview options. When selecting this option, no additional options are provided. The list will contain all jobs in the system, whether set to run once, regularly or on-demand. If set to run at a certain time or times, the schedule will be included in the report.

Network Install Clients

This option is only available when using Network Administrator

Select Reports→Network Install Clients to print a list of the clients that have been configured for network installation. Refer to the main Reports section above for details on the Print and Preview options. The report will contain all information pertaining to the *network boot* and *network installation* for each client. The process will also attempt to determine if the client is currently configured for network boot on the server system, and will include an appropriate message (client is ready to boot, *boot server* is unavailable, or client is not configured on the boot serer). When selecting this option, an additional option is provided:

Include only clients currently ready for network boot: Check this box if you want the list to include only those clients that are currently ready to be network booted from a boot server. If the boot server cannot be contacted to determine the status of the network boot configuration, the client information will be listed regardless.

After a client is configured for network boot (see ***Network Boot/Install Configuration*** in the ***SBAdmin System Recovery Guide***), the boot configuration is updated on the boot server. If the network boot is disabled, the client network boot and install configuration is removed from the boot server but retained on the admin system for future use. If not checked, the list will include all network install client configurations, whether the client is currently ready for booting or not. If checked, the boot server assigned to the client will be checked to see if the client is currently configured for booting, and the client configuration will not be listed only if configured on the boot server.

Backup History

To print a Backup History Report showing the dates, times and backup statistics for each client backup, select Reports→Backup History from the main menu bar. A further option is provided for either running the report in the order of client or job ID.

> NOTE **When using Workstation or Desktop Edition, this report is always run in order of backup job. The option of running by client is not available.**

Running the report by client: Select Reports→Backup History→By Client. You may select one or more clients for which to print the report or, by not specifying any clients, the report will be reported for all clients. The report may be printed even for clients that are no longer configured by manually typing the client name in the entry box. You may optionally select to print subtotals by Job ID, detailing the average megabytes, number of minutes, and Kbytes per second for each job under which the client bas been backed up. These averages will also be shown for each client in the list.

Running the report by Job ID: Select Reports→Backup History→By Job ID. You may select one or more Job IDs for which to print the report or, by not specifying any job IDs, the report will be reported for all Job IDs. The report may be printed even for jobs no longer configured by manually typing the job ID in the entry box. You may optionally select to print subtotals by client, detailing the average megabytes, number of minutes, and Kbytes per second for each client within the job. These averages will also be shown for each job in the list.

You may also select starting and ending dates for this report. If provided, the report will only include backups that occurred within that date range.

Backup Expiration Report

> NOTE **This option is not available when using Desktop Edition.**

To print a report showing the backup labels past their expiration dates, select Reports→Backup Expiration Report from the main menu bar. You will have the option of showing all backups, even if they are not past their expiration date, backups past their expiration dates as of today's date, or backups that will be past their expiration date as of a specified date.

This resulting report will tell you what backups are past their expiration date and may be expired. Of course, any backup may be expired manually (See expiring a backup), and if your overwrite policy is set to allow any overwriting of backups and you did not specify a , then the backups will always be expired when they are overwritten.

A backup will be shown on this report if any of the following are true:

1. You are listing all backups, regardless of their expiration date.

2. There is no backup retention period specified in the job settings.

3. The retention period in the job settings (number of days) has passed since the backup was performed.

4. You specify a reporting date in the future at which time the backup will have expired.

Note that the changing the backup retention period for a job will not change the retention period of backups that have already been performed.

23. Utilities

This section provides instruction on the use of the utilities that are not typically used on a day-to-day basis but provide useful features or the ability to tailor the behavior of the application.

Create System Installation Media

System Installation media is bootable media which may be used to boot the system to the **SBAdmin System Installation** process. Numerous options are available for creating system boot media, depending on the operating system type:

- For ***AIX***, boot media types are **CDROM images**, **tapes**, **hard disks**, and **network boot images**. When a **System Backup** of an AIX system is written to the beginning of a tape, the tape is automatically made bootable for the client's system type. You may also specify in the backup profile the type of system for which to create the boot tape. Refer to the Backup Profiles for additional information on the ***platform type*** for bootable tapes.

- For ***Linux***, boot media types are **CDROM images**, **diskettes**, **hard disks** and **network boot images**. Although the network boot images are created and copied to the boot server, some *bootloader* configuration must be manually performed by the user. This is automated on AIX systems, but is more difficult to do on Linux systems due to the number of different boot loaders and configuration file formats that are available with different Linux distributions.

CDROM images are *ISO9660 format filesystems*, which may be burned to a CD writer using any number of third party applications. For most Linux systems, you can use the "***cdrecord***" command, and on AIX systems you can use the "***cdwrite***" command. These software applications must be installed separately (not provided by SBAdmin), and you must refer to the instructions with the individual application for detailed instructions.

Hard Disks are made bootable only after configuring the hard disk as a ***System Backup Disk***. This option is available when configuring **Servers** (Network Administrator) or **Backup Devices and Directories** (Desktop Workstation Editions) in the Configure System Backup Disk section. Using a hard disk as a boot/recovery device is very handy, especially when storing the ***System Backup*** on the disk, because a system can be booted and reinstalled from a spare disk (such as portable USB or SAN-attached disk) without needing any network or other external (i.e. cdrom) boot media.

Network Boot Images allow a client system to be booted over the network from a ***Boot Server***. This option will create the images and copy them to the boot server. You can create a single boot image for compatible sstems (i.e. same OS release and hardware type), or a separate boot image for each client. A separate option is used to ***Enable a Client for Network Boot***, which is described in the ***SBAdmin System Recovery Guide***.

Creation of each of these media types is described in more detail in the respective sections of the Recovery Guide for each operating system. Refer to the ***SBAdmin (AIX or Linux) System Recovery Guide*** for details.

Node Front Panel

> NOTE **This option will only appear on *AIX* SP Systems, and if the Network Administrator is also the SP Control Workstation.**

This option provides a graphical representation of an **RS/6000** front panel. IBM SP system nodes do not have font panels, keyboards, displays, etc, themselves. Using this graphical window, hardware controls may be executed on the nodes such as powering on and off the system, displaying a console, and even a 1-button *network boot*. Refer to SP System Information for detailed information on this feature.

Write a Tape Label ID to a Tape

> **This option is not available when using the Desktop Edition.**

A Tape Label ID is a unique identifier for each tape that is used with SBAdmin. Tape labels are not required in order to use a tape for a backup, but having a tape label will make it easier to determine the contents of a tape and track which tapes belong together in a set.

For SBAdmin to track the contents by tape labels, the tape label id must be physically written to the tape before it is used for any backups. A physical adhesive tape label often comes with tapes that contains a unique tape identifier. You may use this tape id, if any, or you may create your own id. Tape IDs may contain up to 16 characters, but may not include colons (:) or spaces.

Two write a tape label id to a tape, select Utilities→Write a Tape Label ID to a Tape. A screen similar to the following will appear:

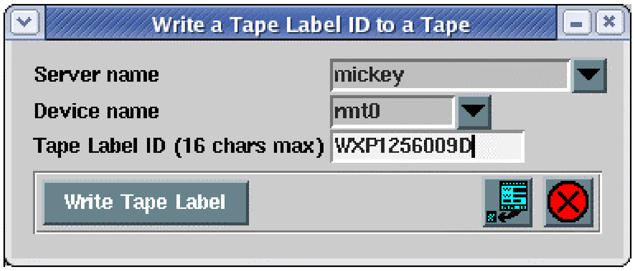

Server name: mickey
Device name: rmt0
Tape Label ID (16 chars max): WXP1256009D

You must press the arrow keys next to each entry field to list and select the backup ***server name*** (if *Network Administrator* used) and the ***device*** in which the tape is inserted. Next, type the tape label id in the ***Tape Label ID*** field. When all entries have been made, press the Write Tape Label button.

First, the tape will be read to ensure that there is not already a current backup on the media. Because this process will write a new label, overwriting any previous backup contents, you may not overwrite a current backup. If a current backup is found on the tape, you will be given the option of automatically expiring this backup and overwriting the tape contents with the new tape label.

The process will then write the tape label to the tape, which usually takes only a few seconds. A message will appear when the process is complete.

> **Once a tape label ID has been written to a tape, it should never again be necessary to use this option again for the same tape, as the tape label id is always reused, even when overwriting a previous backup with a new backup.**

If you ever need to read the tape label ID from the tape, you can use the option Perform a Tape Operation and select the ***Read Backup Header*** option. The tape label ID, as well as other backup information, if any, will be displayed on the screen.

Once a backup has been written onto a tape that has a tape label, the Backup Label will show the tape label IDs of each tape volume that makes up the entire backup.

Perform a Tape Operation

This option provides a number of useful features for performing tape operations, such as rewinding, ejecting, checking and reading information from a tape. To use these options, select Utilities→Perform a Tape Operation from the main menu bar. When doing so, a screen similar to the following will display:

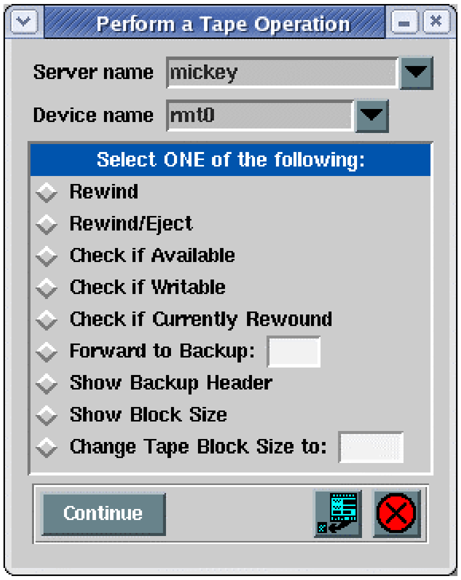

Perform a Tape Operation

Server name: mickey

Device name: rmt0

Select ONE of the following:

- ◇ Rewind
- ◇ Rewind/Eject
- ◇ Check if Available
- ◇ Check if Writable
- ◇ Check if Currently Rewound
- ◇ Forward to Backup:
- ◇ Show Backup Header
- ◇ Show Block Size
- ◇ Change Tape Block Size to:

You must press the arrow keys next to each entry field to list and select the backup ***server name*** (if *Network Administrator* used) and the ***device*** in which the tape is inserted. To perform an operation, press the radio button next to the desired option and press the Continue button at the bottom of the screen. Each option is described below:

1. **Rewind**: Rewinds the tape in the device

2. **Rewind/Eject**: Rewinds, then ejects the tape from the device.

3. **Check if Available**: Displays a message indicating whether or not the device is available and a tape is inserted.

4. **Check if Writable**: Displays a message indicating whether or not the device is available, a tape is inserted and whether or not the write-protect tab on the tape has been set.

5. **Check if Currently Rewound**: Displays a message indicating whether or not the tape is currently rewound, or at beginning of media.

6. **Forward to Backup**: To use this option, you must also enter a backup sequence number in the field to the right, or you may enter the word “end” to forward to the end of the backup. You may insert any volume of the backup prior to or including the start of the backup you are forwarding to (or the last tape volume if forwarding to the end). After forwarding to the end of the last backup on the media, may append additional backups to the same tape and backup label.

7. **Show Backup Header**: Reads the backup header on the tape and displays the header information, which includes the backup id, tape label (if any), backup date, volume number, client, job id, backup type, etc. Note that this differs from showing the backup label since the output of this option pertains only to this tape. Included in the display will be information showing the current position of the tape within the backup.

8. **Show Block Size**: Displays the current physical block size setting for the tape drive. For SBAdmin backups, the tape block size will always be changed to 0 (variable) before a backup is performed, and it will remain set to 0 after the backup completes, since the block size must be set the same during a restore as it was during the backup.

9. **Change Tape Block Size to**: You must enter a block size in the field to the right when using this option. SBAdmin backups are always performed using a variable (0) physical block size setting. If the tape drive block size was set to any other value by another application or when the drive was reconfigured, you will need to set the block size to 0 again before you can read an SBadmin backup.

Perform Tape Library Operations

The options described in this section are not available when using Desktop Edition.

These options are used to perform a move operation, display an inventory of the media within a library, and to display or set the tapes within the library to use in the next backup or restore process.

The Tape Position Number used by these options indicate the tape used by the particular drive number, as you configured in the Random Library Configuration screen, under Define Drive/Tape Slots. This is NOT the library element address, but a tape position number, starting at 1 (for drive 1) and ending with the total number of tape slots configured.

For example, if you have a dual-drive library with 10 tapes assigned to each, the tape slot position numbers for drive 1 would be 1-10 and the tape slot position numbers for drive 2 would be 11-20.

Set/Reset Next Tape for Backup/Restore

SBAdmin always keeps track of the last tape that was used for a backup or restore operation. After you physically replace the tapes in the library, it will be necessary to inform SBAdmin that it should start again with the first tape in the stack. Also, after a backup is performed, you will need to reset the library back to the first backup tape (if the volume was changed) before a verify or restore operation can be performed.

This option is used to set the next tape number in the library that will be used for the next backup or restore process. To use the option, select Utilities→Perform Tape Library Operations→Set/Reset Next Tape for Backup/Restore from the menubar. After doing so, the following will be displayed:

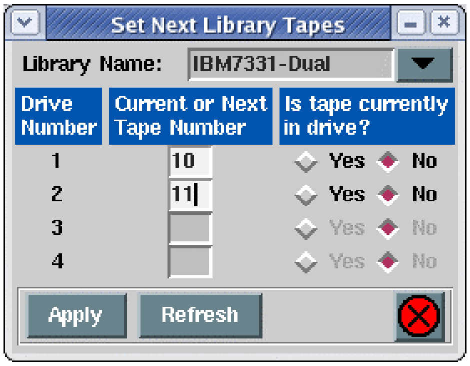

Set Next Library Tapes

Library Name: IBM7331-Dual

Drive Number	Current or Next Tape Number	Is tape currently in drive?
1	10	○ Yes ● No
2	11	○ Yes ● No
3		○ Yes ● No
4		○ Yes ● No

Select the **Library Name** by pressing the arrow button to the right of the entry field. Entry boxes will become available for the number of drives that are configured for this library. Next to the **Drive Number**, enter the **Next Tape to Use**, which must correspond to a tape position *for that drive*.

For example, if you have a dual-drive library configured for 10 tapes per drive, you would enter a *Tape Position Number* from 1 to 10 for Drive 1 and from 11 to 20 for Drive 2.

Move Tapes in Library

Use this option to move tapes from a library tape slot to a tape drive or vice-versa. To begin, select Utilities→Perform Tape Library Operations→Move Tapes in Library from the menu bar. The following screen will be displayed:

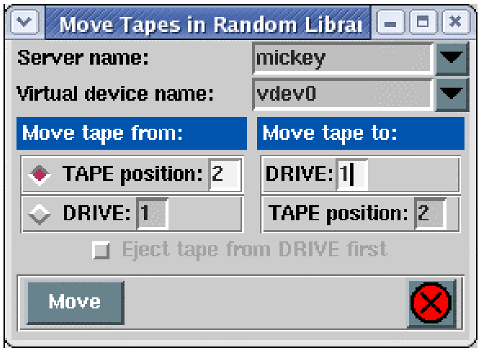

Server name: mickey

Virtual device name: vdev0

Move tape from:

● TAPE position: 2

○ DRIVE: 1

Move tape to:

DRIVE: 1

TAPE position: 2

□ Eject tape from DRIVE first

Select the server name and the virtual device name by pressing the button to the right of each entry field. Only virtual devices you previously configured as a random tape library will be displayed.

In the "**Move tape from**" column, select a radio button indicating whether you want to move from a **Tape Position** or a **Drive**. Enter the *tape position number* or drive number that you want to move from in the entry box to the left, then the tape position number or drive number to "**Move tape to**" in the entry field to the right.

If you are moving from a tape drive to a tape slot position, you should also indicate whether the tape is currently inserted in the drive or not. If the tape has already been ejected, but is still sitting in the drive door, select "No" to this option. If the tape is currently inserted in the drive, then the library cannot move the tape until it is ejected. In this case, select "Yes" to this option.

After making your selections, press the **Move** button to begin the operation. If any of the entries are not valid according to the library configuration, an error message will be displayed. Also, if a move operation error occurs, the message will be displayed.

Display Library Media Inventory

This option may be used to display the media inventory of the library using the **Command to Inventory** configured in the Configure Random Library configuration screen. If no command to inventory the library was defined in the library configuration, you will not be able to perform this operation.

To perform this operation, select Utilities→Perform Tape Library Operations→Display Library Media Inventory from the menu bar. When doing so, the following screen will display:

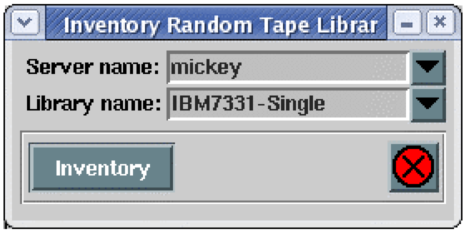

Server name: mickey

Library name: IBM7331-Single

Select the server name (if *Network Administrator* is used) and the library name by pressing the arrow buttons to the right of each field, then press the **Inventory** button at the bottom of the screen. After doing so, a new window will appear showing the command to execute and the output of that command, such as:

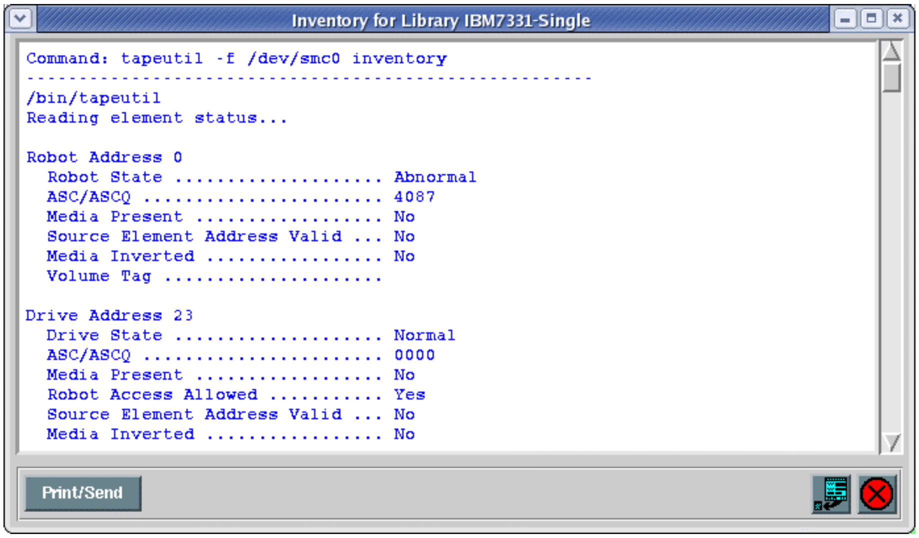

```
Command: tapeutil -f /dev/smc0 inventory
------------------------------------------------------
/bin/tapeutil
Reading element status...

Robot Address 0
  Robot State .................... Abnormal
  ASC/ASCQ ....................... 4087
  Media Present .................. No
  Source Element Address Valid ... No
  Media Inverted ................. No
  Volume Tag .....................

Drive Address 23
  Drive State .................... Normal
  ASC/ASCQ ....................... 0000
  Media Present .................. No
  Robot Access Allowed ........... Yes
  Source Element Address Valid ... No
  Media Inverted ................. No
```

Any error messages that occur will also be displayed in that window. Simply press the ⊗ (close) button when done.

Change Backup/Restore Settings

NoTE **The options described in this section are not available when using Desktop Edition.**

Some settings may be altered which affect the say backups are created or restored. These include the following:

Sparse File Handling

A sparse file is a file in which blocks of data have been written non-sequentially, leaving unallocated blocks in the middle of a file. If the sparseness of a file is not preserved when restoring, the file will be expanded to include all blocks in the middle of the file, often causing a filesystem to inadvertently run out of space.

Preserving sparseness in files is usually desirable. This is sometimes a problem, however, if your files were pre-allocated using NULL characters. If a file is created and all blocks are allocated by writing nulls, or "0"s, throughout the file, the file appears identical to a sparse file on the backup. Since files containing null blocks are indistinguishable from sparse files, the blocks are not retained upon restore. The affect is that a file created at a large size could be restored to a very small size.

To resolves this issue, you may select to create the backup without preserving the sparseness of files. Therefore, if a file was pre-allocated using NULL blocks, the null blocks will also be restored. Note that, when using this option, a truly sparse file (created without pre-allocating blocks by writing nulls) will be interpreted a large file of null blocks, and will be expanded upon restore in order to retain the null blocks. This will often cause the filesystem to run out of space since a file that was once very small is restored quite large.

NoTE **If a backup is created by preserving sparseness, which is the default, then the backup files may not be restored to another system of a different operating system type. If you want to restore a backup to a different operating system type, then you should turn OFF sparse file handling BEFORE creating the backup.**

To change whether or not sparseness should be preserved, select Utilities→Change Backup/Restore Settings→Sparse File Handling from the menu bar. A screen similar to the following will be displayed:

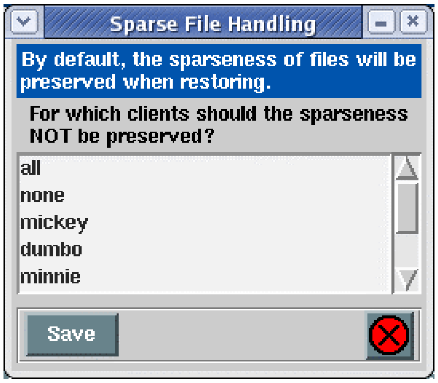

By default, the sparseness of files will be preserved when restoring.

For which clients should the sparseness NOT be preserved?

all
none
mickey
dumbo
minnie

Since sparseness of files is preserved by default (null blocks are discarded), you must select the clients for which the sparseness should not be preserved (null blocks restored). You may select "***All***" if sparse files should not be preserved for any client, or "***None***", to restore the sparse file handing by default. Otherwise, select one or more clients from the list and press the Save button to save the changes.

Read Error Handling

This section provides information on the options for controlling the way read errors are handled when reading from the backup media. Although the Backup Administrator itself provides a reliable backup, the media on which the backup is placed can sometimes become corrupt. These options will allow you to select how much work the application should try to recover from errors reading from a corrupt tape. To change these settings, select Utilities→Backup/Restore Settings→Read Error Handling from the menu bar. The following screen will appear:

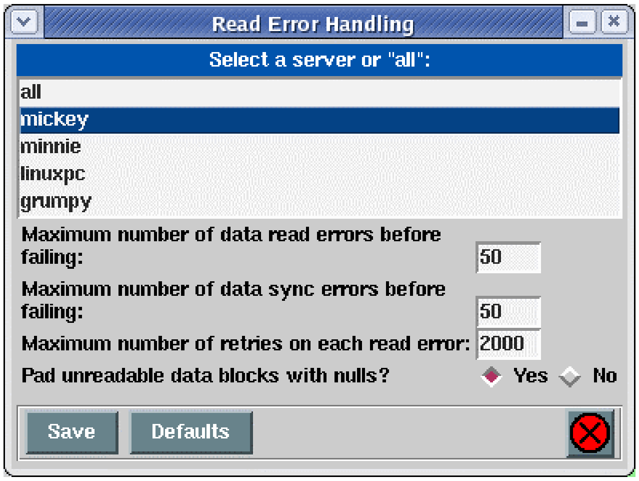

Select a server or "all":

all
mickey
minnie
linuxpc
grumpy

Maximum number of data read errors before failing: 50

Maximum number of data sync errors before failing: 50

Maximum number of retries on each read error: 2000

Pad unreadable data blocks with nulls? ● Yes ○ No

The following is an explanation of each field:

1. **Maximum of data read errors before read failing**: When a read error is encountered, the media device driver, it will, by default, attempt to retry the read up to the number of times specified in the field "***Maximum number of retries on each read error***" below. If the application is unable to read the data, a read error is produced, and the process will either skip the missing data entirely or pad the missing data with NULL bytes, as defined by the field "***Pad missing data blocks with nulls?***".

This option allows you to specify the maximum number of read errors that are produced before the backup aborts. You may specify any number up to 32768 in this field, or you may use a zero (0) to indicate that the reading should abort after the first read error.

2. **Maximum number of data sync errors before read failing**: An individual read of the backup is performed for each buffer, defined by the buffer size of the backup. At the beginning of each buffer is a special key that is used to ensure that the data is being read at the correct point. A "data sync" error occurs when the key is not encountered when reading the data, or the key has an incorrect sequence number.

   When a sync error occurs, the process will either skip the missing data altogether or pad the missing data with NULL bytes, as defined by the field "***Pad missing data blocks with nulls?***".

   This field determines the maximum number of sync errors that may occur before the reading aborts. The value of this field may be any number up to 32768. Using a value of zero (0) indicates that the reading should abort after the first sync error.

3. **Maximum number of retries on each read error:** When a read error occurs, the process will, by default, attempt to re-read the same buffer of data up to the number of times specified by this field. The reading will abort when a read error occurs and has been retried the number of times indicated. You may enter a number up to 32768. An entry of zero (0) indicates that no retries should be attempted.

   > Note
   >
   > **Most tape devices, including 8MM tape drives, will return an error very quickly when a read error occurs, and will allow retries to be attempted from the same data location. Others, such as DDS 4MM tape drives, take up to 2 minutes to return from a read error. These tape devices also do not allow read retries, but will still take 2 minutes to return from an attempt. Therefore, for these, and similar devices, you will want to set this value to zero (0) since retries are not supported, and any attempts will appear to pause the reading indefinitely.**

4. **Pad missing data blocks with nulls?**: When a data sync error occurs, assuming the reading is setup to continue, the missing data will be padded with NULL bytes by default (the field is set to **yes**). This is to ensure that, although the data has been altered, it remains in the correct alignment.

   > Note
   >
   > **It is very important for the data to remain in the correct alignment when restoring raw logical volume backups. If you do not pad sync errors with NULL bytes, all of the data following the error would be restored to a different location than expected. Volume Group, Filesystem and File/Directory backups use an underlying restore command that is capable of resynchronizing when there is missing data in the data stream. Therefore, the value of this field is less relevant when restoring from these backup types. However, the restore command may sometimes fail when it encounters too large a stream of NULL bytes. In this case, it may be advisable to change this value to "no".**

When all changes have been made, press the Save button to save changes and exit this function. If you are unsure of your changes and want to return to the system defaults, press the Defaults button. After doing so, all data in the fields will be replaced with the defaults and you must then press Save to save them.

Network Settings

> Note
>
> **This option is only available when using a Network Administrator.**

By tuning certain network parameters, it may be possible to increase the performance of backups and restores performed when using this application. This option makes it possible to set certain values that affect network performance during backups and restores without affecting network performance of other processes using the same network.

> NOTE
>
> **The changes will only affect the backup and restore processes within this and will have no affect on other system network operations. These settings will override the system default settings or any prior settings changed with the "no" command. The changes applied here will not apply when installing a system from a System Backup because it is not possible for the client to query the values set during the installation process.**
>
> **Important note: You should not change the network settings using this option unless you are familiar with network tuning. Also, it is imperative that the same settings be applied to both the backup server and clients or network backups may lock up. For IBM SP nodes used as backup servers, refer to the SP Section for details on recommended network settings to enhance network performance when using the SP High Performance Switch network.**

The system defaults should be acceptable in most situations. The settings appropriate for achieving the greatest performance on different systems and networks vary widely, so no specific guidance can be given here.

To change the network settings, select Utilities→Backup/Restore Settings→Network Settings from the menu bar. The following screen will appear:

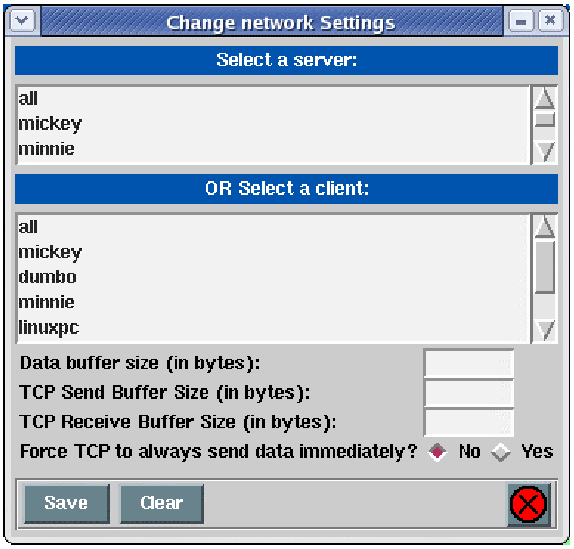

Select a server:
all
mickey
minnie

OR Select a client:
all
mickey
dumbo
minnie
linuxpc

Data buffer size (in bytes):
TCP Send Buffer Size (in bytes):
TCP Receive Buffer Size (in bytes):
Force TCP to always send data immediately? ● No ○ Yes

Select the backup server **or** the client to apply the settings to from one of the corresponding list boxes. You may select "**all**" in either box to apply the settings to all backup servers or all clients respectively. Always remember to apply the same changes to the server as you do the clients!

The following is an explanation of each of the settings fields:

1. **Data buffer size**: This value represents the size of the buffer of data that is written to the network socket in a single write operation.

2. **TCP Send Buffer size (in bytes)**: This value indicates the TCP "***send buffer***" size and is equivalent to the tcp_sendspace parameter of the AIX **no** command. If the send buffer size is greater than 64 Kbytes, the RFC1323 TCP parameter will automatically be enabled, which is equivalent to the rfc1323 parameter of the AIX **no** command.

3. **TCP Receive Buffer size (in bytes)**: This value indicates the TCP "***receive buffer***" size and is equivalent to the tcp_recvspace parameter of the AIX **no** command. If the receive buffer size is greater than 64 Kbytes, the RFC1323 TCP parameter will automatically be enabled, which is equivalent to the rfc1323 parameter of the AIX **no** command.

4. **Force TCP to always send data immediately?**: Select "**yes**" if TCP packets should send immediately. Otherwise, a value of "**no**" indicates that small amount of data should be collected into single packets before being sent.

To remove all prior settings for the select server or client (or “all” if selected), press the Clear button. When doing so, any customized settings will be removed and the default system network settings will be used by default.

When all entries are completed, press Save to apply the changes.

Rebuild (unexpire) a Backup Label

You were sure you'd never need that backup again and didn't want that old backup label cluttering up the system. After taking all the warnings into account you removed the backup label only to find that the backup last night didn't run because you mistakenly scheduled it for noon instead of midnight. Now, the janitor spilled cleaning fluid all over the disk drive and you have to restore your data to the spare disk you cleverly kept in the file cabinet. Your only backup tape is the one you expired yesterday.

As stated in the many warnings you received when you pushed that "expire" button, it is not possible to restore from a tape once the backup label has been removed. But there is hope. Actually, it’s really no problem at all. This option will read through the contents of the backup and rebuild the label, one backup at a time. Once that is accomplished, you may restore from the tape just as you could before making this terrible blunder.

To rebuild the backup label, select Utilities→Rebuild (unexpire) a Backup Label from Tape from the menu bar. You will be prompted for the backup server (if *Network Administrator* used) and the tape device to read. After doing so, a screen will appear showing the status of the process such as the following:

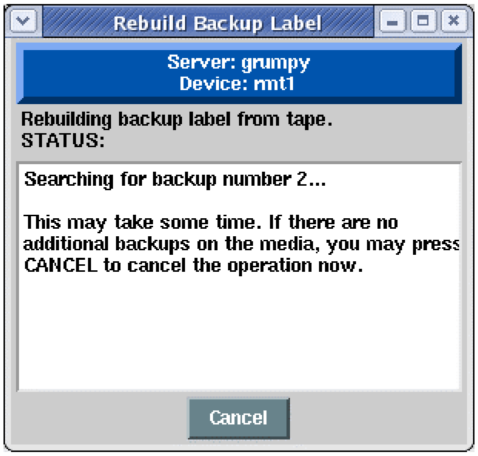

Server: grumpy
Device: rmt1

Rebuilding backup label from tape.
STATUS:

Searching for backup number 2...

This may take some time. If there are no additional backups on the media, you may press CANCEL to cancel the operation now.

When the process is complete, the completion status will be displayed and a Show Label button will appear, with which you can display the Backup Label or press the Cancel button to return.

Change Access Permission of a Disk Backup

This option is only available when using a Network Administrator.

When a backup is performed, one of the options of the backup profile allows you to specify the read permission of backups written to disk. This is because a disk file is, by default, readable by any user, local or remote. Since the disk backup files may contain confidential information, this is often not advisable.

The Backup Administrator automatically makes the contents of the file readable only by the ***root user*** of the local or remote system, whether the file itself is readable or not. In the backup profile, you may also specify whether only the client from which the backup originated may read the contents, or any client with access to the backup server. If you plan to allow the backup of one client to be restored or installed onto another client system, you must permit other clients access to the backup. If you did not do so when the backup was created, this option will allow you to change your mind.

To change the client access to a backup, select Utilities→Change Access Permission of a Disk Backup. You will be asked to select the backup server and the directory to which the desired backup was written. You may also be asked to specify the specific client from which the backup originated. Then, you are asked to select the specific backup from a list. After selecting the desired backup, a screen similar to the following will appear:

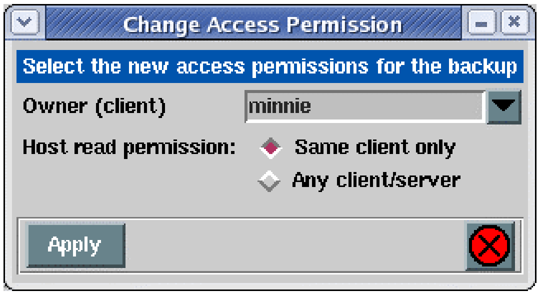

Select the new access permissions for the backup

Owner (client): minnie

Host read permission: ● Same client only ○ Any client/server

You may change the **Owner (client)** of the backup. This will have no relevance if the access permission of the backup allows "**Any client/server**" to read it. However, if the access permission allows "**Same client only**", then only the original client, or the ***Owner***, may read it.

After you have made your selections, press the Apply button to save your changes and exit this function.

24. Using Backup Administrator on an IBM SP System

> **This chapter is only applicable to AIX systems. This section provides instruction on options and messages, as well as any differences in operation of the application on SP systems running AIX. These differences will only occur when the SP Control Workstation (CWS) is configured as the Network Administrator.**

Node Front Panel

The option **Utilities→Node Front Panel** may be selected to display a graphical interface for manipulating the hardware controls of an SP node. When selected, a graphical window such as the following will be displayed:

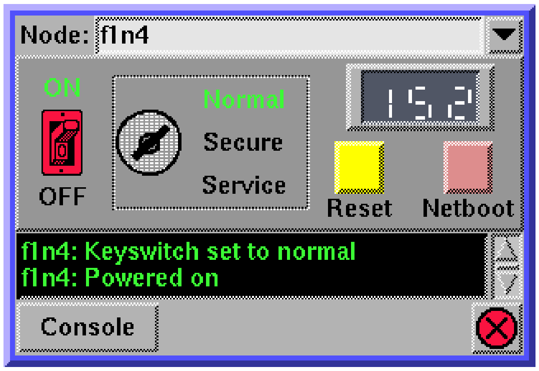

Node: f1n4

f1n4: Keyswitch set to normal
f1n4: Powered on

This interface may be used to perform the following node hardware functions:

- **Power ON or OFF the system**
- **Change the KEY SWITCH**
- **RESET (press the reset button)**
- **Perform a NETWORK BOOT (see below)**
- **Open a CONSOLE TERMINAL (TTY)**

In addition, the value in the **System LED** is displayed if not blank. Also, a scrollable message box at the bottom of the window will be updated to reflect all hardware controls selected from this screen. The size of this message box may be increased by vertically resizing the window.

To select the node to display or change, either type the node name in the **Node** field at the top of the screen, or press the arrow button to the right of this field and select the node from the list. The display will be updated to reflect the current settings of the node's power, key switch, and LED value. The display will be updated every few seconds in case any of these values are changed by an outside source.

> **IMPORTANT: As if you are physically performing the operation directly to the hardware, there are no visible prompts provided after you select a label or icon on the screen. If, for instance, you press the power OFF label, the node will be IMMEDIATELY powered off.**

Select any of the following labels or buttons to perform the hardware operations:

Power On: Press the "**ON**" label above the power switch to turn on the power to the node. When pressed, the power switch will change to the on (up) position, and the ON label will be displayed in green.

Power Off: Press the "**OFF**" label beneath the power switch to turn off the power to the node. When pressed, the power switch will change to the off (down) position, and the OFF label will be displayed in red.

Change Key Switch: The current key switch position of the node will be displayed in green. Press either the **Normal**, **Secure**, or **Service** labels to change the key switch to the new position. When selected, the key switch will point to the new position and the position text will be highlighted in green.

Reset: Press the **Reset** button to reset the system. This is equivalent to pressing the reset button on the front of a standalone system, which will normally begin rebooting the system in the mode specified by the key switch position.

Open a Console (tty): Press the **Console** button at the bottom of the screen to open a console to the node. A window will be displayed and, if the node is operational, a login prompt will appear.

Network Boot: Press the **Network Boot** button to perform a network boot of the node. A network boot must be performed to boot and install a node over the network. You must first have configured the node for network boot/install using the option ***Enable/Disable Network Installation of a Client*** in the ***SBAdmin System Recovery Guide***. There is no physical "Network Boot" button on the front panel of a standalone system, but a button is provided here as a shortcut for performing **manual node conditioning**. Manual node conditioning is the manual process of booting a node using the supplied PSSP software. This process will perform all of the same steps with a single button press by doing the following:

1. A read-only console window will be opened for the node so that you may monitor the process of the BOOTP process
2. The key switch is changed to the SECURE position
3. The node is powered off, if necessary, and powered on again
4. When the LED stops at "200", the key switch is changed to the SERVICE position
5. The reset button is pressed
6. When the BOOTP menus appear, the menu options are selected and prompts are responded to automatically for performing a broadcast boot from the primary ethernet adapter.
7. The key switch is changed to the NORMAL position
8. The read-only console window is closed and a new console window is displayed to which the user may respond to prompts as presented.

After performing the network boot, the installation process should begin or the installation and maintenance menu should appear, depending on how the network boot/install for the node has been configured.

Supported PSSP Code Levels

At the time of this publication, SBAdmin for AIX supports **PSSP levels 2.3, 2.4 and 3.1** and has been fully tested using the following combination of PSSP and AIX code levels:

Node PSSP level	Node AIX Level	CWS PSSP Level
2.3	4.2.1	3.1
2.4	4.3.1	2.4
2.4	4.3.2	3.1
3.1	4.3.2(+)	3.1

These are the latest AIX and PSSP levels tested by Storix at the time this section of the user guide was created. Although later PSSP and AIX versions are available at this time, there have been no report of problems with later levels. Although untested at PSSP levels earlier than 2.3, all features should work the same for 2.1 and 2.2 as for 2.3. Although not guaranteed, Storix will make a best effort to support code levels 2.1 and 2.2.

Testing included full system backup of a node and installation from the backup onto a different node, and was considered successful when the node being installed was automatically customized with the new node information configured in the SDR on the control workstation.

Configuring a Node for Network Installation

The process for configuring a node for network installation is identical to that of any other standalone AIX system as outlined in the section option ***Enable/Disable Network Installation of a Client*** in the ***SBAdmin System Recovery Guide***. However, when running Backup Administrator on the Control Workstation, you will be see the following additional prompts after configuring a node for network install:

```
The client to be installed is an SP node.

To retain the node information, the node must be customized after
the installation completes. This requires setting the boot
response to "customize" for this node.

Do you wish to setup the node to customize now?
```

Unless you plan to reboot this node without reinstalling it, you should always answer "**yes**" to this question. The node's ***boot response*** will be set to customize, which tells the control workstation to automatically re-customize the node the next time it is booted. If you answer "**no**" to this question, the node will NOT be customized with the node information after it is reinstalled, and will need to be done manually by running the command "**spbootins -r customize -s yes *framenum nodenum* 1**" on the control workstation.

```
The "setup_server" process must be run before the installation of
the node is started. This needs to be run only once after all
nodes to be installed are setup to customize. Run setup_server
now?
```

As indicated, the "**setup_server**" command must be run (on the control workstation) before initiating a network boot/install of the node. However, if there are several nodes to be installed, this command need only be run once after all nodes are configured. The command normally takes a few minutes. However, if a network installation of a node has never been performed using PSSP, running setup_server may cause PSSP to create network boot images and SPOTs, which could take an hour or more. If you had selected not to run setup_server but wish to do so, you can either reconfigure the client for network install (takes only a few seconds) or simply run "**setup_server**" at the command line (on the control workstation).

Backing up Nodes Across the SP High Performance Switch

The SP High Performance Switch (HPS) network may be used to perform backups and restores, either to disk files on a node, or to tape drives attached to nodes. When a high-speed switch network is available, backups and restores, which use considerable bandwidth, will likely run faster and cause less network interference then when using slower networks.

To use the HPS network for backups, you must first setup the Alternate IP Address or Hostname for Backups/Restores in the server configuration. There, you would enter the hostname or IP address (as known by the client) of the server's high-speed switch network adapter. After doing so, backups and restores performed to this backup server (node) will take place across the switch whenever you select to use the alternate network connection.

To backup using the alternate network selected for the server, you must select the button labeled "Use Alternate Server IP/Hostname" when configuring the backup job. If the button is not selected, the primary network connection used by the Admin System will be used in performing the backup. Note that this button will be disabled if there was no alternate IP address or hostname setup for the server. To set the alternate IP address or hostname for a server, refer to the server configuration.

You must not select the high-speed switch adapter for network boots/installs, as AIX currently does not support network booting or installing across this network type.

Tuning Performance of the High Performance Switch

If you are using the SP High Performance Switch (HPS) for backups and restores, you may use the AIX “no” command or you can change the Network Settings in the Backup Administrator in order to tune network attributes for better performance. Changing the “no” options will apply to all network communications while the Backup/Restore settings will only apply to backup or restore data transfers between the client and backup server when using this application.

Certain “no” settings are recommended for nodes with an HPS regardless of the application. The Network Settings should only be changed when using the HPS as your backup network (backing up from node-to-node) and when using a high-speed backup device (such as a high-speed tape or disk drive in excess of 2 MB/sec).

In addition, you will need to change the *spoolsize* and *rpoolsize* attributes of the switch adapter (*css0*). The recommended settings are based on the other network settings as well as the amount of network traffic that is expected to and from the clients or servers.

The following table shows the recommended settings based on customer experiences and on IBM recommendations as described in the *SP Administration Guide* (subject to change):

Recommended AIX “**no**” settings (should be applied to all nodes)	
sb_max	1310720
thewall	16384
udp_sendspace	65536
upd_recvspace	655360
Recommended AIX “**no**” settings ONLY if you want settings to apply to all network traffic (for backup/restore traffic only, do not set these values – refer to the next section)	
tcp_sendspace	655360
tcp_recvspace	655360
rfc1323	1
Recommended Network Settings if you want settings to apply to only backup and restore processes (some values will override values set in the prior section)	
Data Buffer Size	524288
TCP Send Buffer Size	655360
TCP Receive Buffer Size	655360
Force TCP to send data immediately?	Yes
Recommended switch adapter settings (set with “**chgcss**” command). Set these to either the recommended value shown or to any multiple of 655360.	
rpoolsize (clients)	4194304
spoolsize (clients)	4194304
rpoolsize (backup servers)	12582912
spoolsize (backup servers)	4194304

Additional Considerations

The following are additional things to consider when using Backup Administrator for AIX in an SP environment:

1. **Disable network boot after installation of a node**: Some PSSP levels automatically set the default boot sequence, even in "*normal*" mode, to boot first from the ethernet network adapter. If a boot server, other than the control workstation, was used to perform a network install of the client, the client will be again booted from the network, possibly re-initiating the network installation of the node. If the control workstation was used as the boot server for the node, the network boot of the client is normally disabled automatically after the network install completes. However, you should do so again manually to ensure the network boot is disabled. Refer to the option ***Enable/Disable Network Installation of a Client*** in the ***SBAdmin System Recovery Guide*** for details.

2. **PSSP 2.3 Network Install:** At PSSP 2.3, node customization occurs at the end of the installation process prior to the system reboot. At that time, network communications are setup only to communicate with the boot and install servers. If the control workstation requires a different route than that of the boot and install servers, the customization will not occur correctly. However, the initial setup of the customization will occur, and the customization may be completed after the installation and reboot is complete by running the **/usr/lpp/ssp/install/bin/pssp_script** command from the command line.

25. Commands

Although all operations may be performed from either the *Backup Administrator*, and most from the *SMIT* user interface, it is sometimes desirable to run commands at the command-line to perform certain tasks. This section lists the commands that may be run from the command line. Some commands are used on the *Admin System* and others on the *Client* or *Server* system, as indicated. All commands listed here may be run *without* specifying the full pathname.

In this section, the following command syntax conventions are used:

bold	text in **bold** font is the command name (type as shown)
normal	text in a normal font is typed as shown (but may be optional)
italics	replace the *italicized* word with a value (which may be optional)
[*argument*]	*argument* is optional
{*arg1* | *arg2*}	either *arg1* or *arg2* must be supplied

Unless otherwise noted, these commands may only be executed by the ***root user***, or a user defined (by the root user) using the administrator interface. Refer to Configuring Users for more information.

sbadmin

Purpose:

Start the System Backup Administrator user interface

Use on:

Administrator

Syntax:

sbadmin [**-d** *display*]

Options:

display	Specifies the display on which to place the application (i.e. *hostname:0*)

Details:

This is the command used to start the System Backup Administrator (SBAdmin) user interface. The interface will be shown on the display indicated by the DISPLAY environment variable, if set, or on the default display (*hostname*:0 or unix:0). If the **–d** option is provided, the interface will be shown on the specified host and display and the DISPLAY variable will be ignored.

stbackup

Purpose:

Run a backup outside of a job

Use on:

Administrator or Client

Syntax:

stbackup [-**cehmnOPrx**][-**b** *bufsize*][-**l** *level*][-**N** *networkType*][-**p** *platformType*][-**s** *server*]
{-**d** *device*|*directory*}{-**t** *backupType*} **dataopt** …

Options:

-b *bufsize* Buffer size of backup data. Default is 64 Kbytes.

-c Compress backup data before sending to server or device

-d *device/dir* Device name (i.e. “rmt0” or "st0") or directory name if creating a disk image backup.

-D *description* Backup description. Any text up to 60 characters, excluding colons (:), may be used. If the text contains spaces, surround the entire text string by double quotes.

-e Eject tape at end of backup

-h Set disk image file permissions to as to allow only host from which backup was made to read the backup data. This has no affect on tape backups.

-l *level* Incremental backup level (0-9). If not specified, no incremental level will apply. Level 0 applies to System, Volume Group and Filesystem backups, and levels 1-9 apply only to Volume Group and Filesystem backups.

-m When recreating logical volumes from the backup, create LVs using the original physical partition placement, when possible. (***AIX*** only)

-n Do not include raw logical volumes on a System or Volume Group backup. This has no affect on File/Directory, Filesystem or Logical Volume backups. (***AIX*** only)

-N *nettype* Network adapter type of network device support to be included on the backup tape. Only one network type may be selected, allowing the user to boot from this tape and install from a network server. This option applies only to ***System Backups*** of clients in a *Network Administrator* environment. (***AIX*** only)

-O Allow overwriting of another SBadmin backup. If this flag is not specified and another backup exists on the tape, the command will fail and no data will be written.

-p *platform* Platform type of boot tape. This option applies only to System Backups and determines the type of system the tape will boot on. (***AIX*** only)

-P Create a ***Power System Backup***. This causes all logical volumes, even filesystems, to be backed up as raw logical volume data. (***AIX*** only)

-r Rewind tape before starting backup.

-s *server* If backup is performed to a device on a remote server, this option indicates the backup server name. This option only applies to client systems in a *Network Administrator* environment.

-t *type* Specifies the type of backup to be performed. This option is required and must be one of the following:

D	File/Directory backup
F	Filesystem backup
L	Logical Volume backup
M	Meta-disk (Software RAID) backup (***Linux*** only)
P	Raw Partition backup (***Linux*** only)
S	System Backup
V	Volume Group backup

-T *tapelabel* Specifies the tape label ID tapelabel should be applied to the first tape when writing to tape. Only has affect when the backup starts at the beginning of the tape.

-x Suppress progress indicator output.

-X ***AIX*** only: Indicates that the /tmp filesystem should be automatically increased in size as needed to create a tape boot image. Only applicable when writing backup to tape.

dataopt … Specify one or more elements to backup. The type of element specified must match the type of backup as specified by the **–t** *type* option.

Details:

The **stbackup** command is used to perform backups from the command line. These backups run independently of the Backup Administrator, therefore the administrator has no knowledge of the backup. Although the backup administrator is not controlling nor recording information about the backup, the backup may later be imported into the Backup Administrator using the option Rebuild (unexpire) a Backup Label. This will be necessary if you plan to use the administrator to restore from these backups in the future.

Note: Running **stbackup** from the command line does not communicate with the Backup Administrator. If you are using the administrator, you may want to create a backup job, then use either the stqueuejob or strunjob command instead, which will make a record of this backup, backup label, history information, etc.

The last argument to the command, "**dataopt**", indicates the data to include on the backup. This is dependent on the type of backup to perform, as specified with the **–t** option. The following shows the backup types and the type of data to specify as the **dataopt** argument:

Type	Description	Systems	Data to Specify
S	System Backup	any	Volume Group name(s) or "all"
V	Volume Group	any	Volume Group name(s)
F	Filesystem	any	Filesystem mount point(s) (i.e. "/home /data")
L	Logical volume	any	Logical volume name(s) (i.e. "lv00 lv01")
M	Meta-disk	Linux	Meta-disk name(s) (i.e. "md0 md1")
D	File/Directory	any	Directory names (i.e. "/home/sam /usr/local/bin")

P	Partition	Linux	Raw partition names (i.e. "sda3 sdb5")

Volume Group and Logical volume backups are only available on Linux systems if LVM (Logical Volume Manager) is installed. Meta-disk backups are only available on Linux if meta-disk (Software RAID) support is installed.

System Backups: When performing a ***System Backup*** (**-S**), the **dataopt** is a list of volume groups to include on the backup. For ***AIX*** systems, the **rootvg** volume group is always included, so it need not be specified. In that case, if no volume groups are specified, only the rootvg volume group is included. On ***Linux*** systems, this option only applies if LVM is implemented on the system and there are volume groups configured. If no volume groups are specified, then no volume group data (including filesystems and raw logical volumes) will be included on the backup. Whether AIX or Linux, you may specify "**all**" to include all available volume group data. Note, however, that raw logical volume data (logical volumes that do not include filesystems) will also be included unless you specify the **–n** option.

The **–P** (Power) option may be specified with the ***System Backup*** type (***AIX*** systems only), and indicates that all filesystems and logical volumes in the specified volume groups should be backed up as raw logical volume data. This will make both backups and restores run faster, but you will only be able to restore entire filesystems and logical volumes, not individual files or directories. Also, if the filesystem is only partially full, then this option could cause both backup and restore to run longer since it backs up the entire logical volume even when only partially used.

The **–l** (incremental level) may be specified to indicate that only files which have changed since the last *prior level* backup should be included. Level 0 indicates that all files should be backed up and is used as a base for all subsequent incremental backup levels. Level 0 may be specified with a System Backup to indicate that this backup should act as a full incremental from which latter incremental levels may be applied.

Backups may run faster when using a larger buffer size (**-b** option) than the default of 64 Kbytes. If the tape device supports a larger than 64 Kbyte buffer, increasing this value will often cause backups to run faster. Try buffer sizes that are a power of 2, such as 128K, 256K and 512K. This value will only support up to 1024 Kbytes (1 Megabyte).

If a server is specified with the **–s** option, you must have first enabled access to the specific *device* (or "all" devices) or *directory* to this client (or "all" clients) on the server. Refer to Configuring Servers to configure access using the *Network Administrator* or use the **Configure Backup Servers** options from the **SMIT** menus (***AIX***).

The **stbackup** command will exit with a completion message and one of the following exit codes:

0	Success
1	Backup failed
2	Pre-backup failure (no data written)
6	Non-bootable System backup created (warning - backup completed)
7	Cannot overwrite tape
9	Internal program error
99	Killed by user (Ctrl-C)

stcheck

Purpose:

Check the network communication between any systems running System Backup Administrator

Use on:

Network Administrator, Client or Server

Syntax:

stcheck [**-**myaddr {*MyIPaddress* | *MyHostname*}] {*Hostname* | *IPAddress*}

Options:

MyIPddress Specifies the IP address of the local network adapter to use to contact the client or server
MyHostname Specifies the hostname of the local network adapter to use to contact the client or server
Hostname Hostname of the client or server to contact
IPAddress IP Address of the client or server to contact

Details:

This command may be used to determine communication errors. If, for instance, a client or server shows as unavailable from the Main Screen of the Backup Administrator user interface, or if an error such as “Cannot contact server” occurs during normal operations, this command can help determine the cause.

In the following example, the command was used to check the communications problems with a client (mickey) that did not have the calling host (spiderman) defined as an admin system:

```
Communications check for mickey
  Actual name is mickey.storix.com
  Socket port is 4156
  Network interface is 192.1.1.101
  Timeout is 5 seconds
Contacting remote host...
Remote host responded:
  My (actual) name is mickeymouse
  My command path is /storix/bin
  My port number is 4156
  Your address is 192.1.1.1
  Your hostname (as I know it) is spiderman
  Your port number is 3181

Your host (spiderman) is not a valid Admin System for mickey.
If it is the admin system, you must add your hostname
to the /storix/config/admin_servers file on mickey.

Your host (spiderman) is a valid client for server mickey.
   Valid devices are: all
   Backup directories: /backups /backups/spiderman
   Network Install directories: /backups /backups/spiderman
```

Should a communication error occur, one or more possible reasons will be listed along with the steps to check or repair the problem.

stclient

Purpose:

Add, delete or list configured clients

Use on:

Network Administrator

Syntax:

stclient [**-aAlr**] [**client**]

Options:

-a	Add or change the specified client
-A	List configured clients while checking and displaying system availability
-E opt	Enable (opt=1) or disable (opt=0) encryption support for this client
-r	Remove the specified client
-l	List configured clients
-v	Show client features enabled (used with –l)

Details:

The **stclient** command is used to add, remove and list clients configured on the Network Administrator.

To list clients, use the **–l** option. By default, only client names are shown. If the **–v** option is also provided, the client features that are enabled are also included. Currently this includes only whether or not encryption support is enabled for the client. If so, “ENC” will appear after the client name. Also, ***SMB clients*** will be shown in the report (SMB in feature list), but may not be added or changed with this command.

If you want to also show the client availability as well as the system type, use **–A** with **-l**. This will indicate the client availability as either **available** or **down**, and will display the system (*AIX* or *Linux*) and machine type (*i386*, *ppc*) or “**unknown**” if the client cannot be contacted. If no client name is provided, all clients are listed, indicating whether or not the client is currently available. Availability is determined by whether or not the network admin can communicate with the client. If not, you can use the **stcheck** command to help isolate the problem.

stcopybackup

Purpose:

Copy an existing backup to different server, directory or device.

Use on:

Backup Server

Syntax:

```
stcopybackup [-er][-b bufsize] [-h hostperm][-l inputBackupID][-L outputBackupID]
             [-n backupNumber][-s server ]{-d InputDeviceOrDirectory}
             {-D outputDeviceOrDirectory}
```

Options:

-b *bufsize* Buffer size to use for new backup. Default is to keep same buffer size of original backup.

-d *device/directory* Source device or directory containing backup to copy

-D *device/directory* Destination device or directory to write new backup to

-e Eject the tape from the output device when copy is complete

-h *hostperm* If writing to a disk file, indicates host read permission (a=any host, h=original host only). Default is “a” if copying from tape, and defaults to the original host permission setting if copying from disk to disk.

-l *backup ID* Backup ID to read if copying from a disk backup

-L *backup ID* If writing to a disk backup and appending to an existing backup label, specify the backup id of the backup to append. Not used when appending tape backups.

-n *seqnum* Backup sequence number to copy from source if there are multiple backups on the media.

-r Rewind the output tape device before starting the copy

-s *server* Server name if output device or directory is on a remote server

Details:

This command is used on the backup media server containing the original backup to copy. Any backup may be copied from tape or disk directory to any other tape or directory, either on the local system or to another server.

Important Note: Since this command is run on a backup server, not the administrator, no backup label or history information is created for the output backup. If this information is required, you must use the ***Utility*** option to ***Rebuild a Backup Label*** from the output media, or use the **struncopy** command on the *administrator* instead.

If writing the backup to a different server, the local system must have been configured as a valid client for the destination server. Use the **stserver** command on the *network administrator* system to list or configure clients for a server.

This command will only copy one backup at a time to the output device or directory. Even if the source device or *backup label* contained multiple backups, only one may be copied at one time. You may, however, append multiple backups to the output device (or backup label if writing to disk). To indicate which source backup to copy (from a *backup label* containing multiple clients or backup jobs), use the **–n** *seqnum* option. To copy all backups from source to destination, run this command multiple times, in creating the *seqnum* option by one each time.

Stacking backups to tape

If copying to a tape device, you may indicate if you want to ***rewind*** before starting the backup and if the tape should be rewound and ***ejected*** at the end of the backup. If you do not rewind at the start of the backup, you may append the source backup to the end of the destination media (if the destination media is currently at the of volume. The destination backup label will be appended with the selected source backup.

Buffer Size

You may alter the buffer size of the backup by entering a buffer size (in Kbytes) using the **–b** *bufsize* option. This is quite useful in increasing the performance of backups when writing to different media. For example, the default 64K buffer size may be adequate when you wrote your original disk backup file, but when copying to a high-speed tape drive, a higher buffer size (i.e. 256K) may provide much greater backup performance. To use the same buffer size for the destination as was used for the source, do not specify a new buffer size.

Host Read Permission

If using *Network Administrator*, and the destination backup is written to a disk directory, you may also change whether only the original client host or any host may read the backup data by using the –h hostperm option. If copying from tape to disk, the default is to allow all hosts (**-h a**) to read the backup. If copying from disk to disk, the default is to leave the original setting unchanged. To allow only the original host which wrote the backup to read it, use **–h h**.

stjob

Purpose:

Add, delete or list configured backup jobs

Use on:

Network Administrator

Syntax:

stjob [**-alrv**][**-A** *opt*][**-E** *enckeyid*][**-R** *opt*][**-S** *opt*][**-V** *opt*][**-b** "*datalist*"][**-c** "*clients*"]
[**-d** "*device/dir*"][**-D** "*descriptio*n"][**-e** "*schedule*"][**-n** *days*][**-s** *server*][**-X** *exlist*]
{**-w** *when*} {**-p** *profile*} jobid

Options:

-a	Add or change the specified job
-A *opt*	Use alternate server hostname/IP address (1=yes, 0=no)
-b *data*	Data list to backup
-c *clients*	Client list to backup (if Network Admin)
-d *device/dir*	Device or directory to write backup to
-D *desc*	Backup job description (overrides description in backup profile)
-e *schedule*	Backup schedule (if **–w** is *later* or *regularly*) in ***cron*** format
-E *keyID*	Enable encryption with supplied encryption key id
-l	List configured jobs
-n *days*	Days to retain (overrides global overwrite policy
-r	Remove the specified job
-p *profile*	Backup profile name (required if adding a new job)
-R *opt*	Remove backup job after run successfully (1=yes, 0=no)
-s *server*	Server to write backup to (if Network Admin)
-S *opt*	Use snapshot backups with this job? (1=yes, 0=no)
-v	Show detailed information on jobs or named job. use with **–l** option.
-V *opt*	Perform auto-verify after successful backup (1=yes, 0=no)
-X exlist	Enable and use specified exclude list name, if any, when job is run.
-w *when*	When to run this backup (o=onDemand, l=later (once), r=regularly (**-e** *schedule* required)

Details:

The **stjob** command is used to add, remove, change and list backup jobs configured on the Network Administrator.

To list currently configured jobs, use the **–l** option. By default, only Job Ids are shown. To show all of the details of a job, use the **–v** flag also. This will display a paragraph of information for each configured job showing all options configured.

If you want to configure a new job, use the **–a** option and specify the *Job ID* at the end of the command. You will be required to specify a *backup profile* using the **–p** flag. If no list of data to backup is provided in the profile, then you must also specify the **–b** *datalist* option. All other options will be written with a default value. After adding a job, be sure to list it with the **–vl** options to display the full list of settings.

NOTE: Be careful to use quotes around any argument to flags which contain multiple words.

Scheduling a job

The default is to configure a job to run *on demand* (**-wo**). If a job is to be run only once, but at a specific time (*later*), or on a regular basis, specify the **-wl** or **–wr** options respectively. You will then be required to add the **–e** *schedule* option, where *schedule* is a cron-style entry in the following format:

Minute Hour DayOfMonth Month DayOfWeek

Each of the above fields are required and must be separated by a single space. You may specify any number of options for each field, each separated by commas. For example:

```
0  17  *  *  1,2,3,4,5
```

indicates that at 5:00 PM (minute 0 of hour 17) on every day of the month (*) of every month of the year (*), the backup should run, but only on Monday through Friday (1,2,3,4,5).

Changing a job

If a job currently exists, you can change specific options by specifying only that option and its new argument. For example, to turn on auto-verify for job “myjob1”, use the following command:

```
stjob –aV1 myjob1
```

stkeys

Purpose:

Create encryption keys on a client for use in encrypting and decrypting backup data.

Use on:

Client

Syntax:

stkeys [-**flr**] **-k** *KeyID*

stkeys [-**a**] -**b** *numbits* {-**h** *HexKey* | -**t** *textKey*} -**k** *KeyID*

Options:

-a	Add a new encryption key
-b *numbits*	Number of encryption bits (128, 192, 256)
-f	Force keys to be removed without prompting (use with -r)
-h *HexKey*	Encryption key in Hex (length must be at least bits/4)
-k *KeyID*	Encryption Key ID (up to 20 chars)
-l	List configured keys
-r	Remove specified keyid
-t *textKey*	Encryption key in ASCII text (length must be at least bits/8)

Details:

This command is available on every client, and is used to configure the encryption keys used in encrypting and decrypting backup data. The encryption keys are either 128, 192 or 256 bits,.

NOTE: This command is always used on the client to configure the keys in place of doing so from the administrator application. For security reasons, this prevents any network traffic that might contain the encryption key itself.

Once an encryption key is created and used for a backup, the same key MUST be used to decrypt the data! If they key is changed or lost, there is no way to restore the backup data again (and please don't call Storix, because even they won't be able to help you)!

The configured keys are stored in a local file on the client, and are referenced using only the ***Encryption Key ID*** you assign to it. This prevents the need to remember the entire encryption key, and also prevent the need to send the encryption key from one system to another when using a Network Administrator.

Encryption keys may be specified either as a hexidecimal number or as a text string that the command will convert to a hexadecimal number for you. A text strinhg may contain any alpha or numberic characters, as well as any punctuation characters except for quite ("), colon (:) or back-tick (`). Using a text string, such as "My-Encryption-Key!" will help you to remember the key in case it is lost or changed, but will limit the range of hexidecimal characters that are generated.

Note that the encryption key (or text string) that is entered must be the followintg number of characters, based on the encryption bits setting:

Bits	Hex String (length)	ASCII text string (length)
128	32	16
192	48	24
256	64	32

stprintlabel

Purpose:

Send a backup label to the printer, email address, or append to file.

Use on:

Administrator

Syntax:

stprintlabel [**-q** *printer* | "**Email**" | "**File**"] { **-b** *backupid* | **-t** *tapelabelid* }

Options:

-q *printer*	Name of print queue. If not specified, label will be sent to standard output.
Email	The report will be sent to the email address specified in Report Preferences.
File	The report will be appended to the file specified in Report Preferences.
-b *backupid*	Specifies a Backup ID if printing backup label given a backup ID.
-t *tapelabelid*	Specifies a Tape Label ID when printing backup label given a tape label ID.

Details:

This option will format and output a backup label to the specified printer queue, email address, append to a file. If no queue is specified, the backup label will be sent to standard output.

The name of the print queue, email address, or file name to append must be defined in Report Preferences prior to choosing the option to output the backup label.

Specify either a backup id (-**b**) or tape label id (-**l**) for which to format and print the backup label. Although specifying a tape label id, the entire backup label will be shown, including a list of other tape label IDs used in the backup.

stqueue

Purpose:

Performs various operations on a job queue, such as listing, adding and deleting jobs, etc.

Use on:

Administrator

Syntax:

stqueue –L
stqueue –A *JobID*
stqueue {**-D**|**-H**|**-K**|**-R**|**-S**} {**-i** *queueid*} *JobID*

Options:

-A	Indicates the specified JobID should be added to the queue.
-L	List all jobs currently in queue. Each job will be listed on one line.
-D	Delete the specified JobID from the queue.
-H	Place the specified JobID on hold, preventing this and other jobs submitted after it to be run until the queue is restarted.
-K	Kill the specified JobID if currently running.
-R	Restart the specified JobID. Use this if a previous job failed or was placed on hold.
-S	Displays the progress indicator, command output, status and error messages for a job that is currently running or has failed.
-i queueid	Queue ID (as displayed with **–L** flag) of the specific job to add, delete, change or show. Use this option if JobID appears multiple times in the queue to indicate the specific occurrence of the job.
JobID	Job ID of the job to add, delete, change, or for which to display job command output.

Details:

This command provides all of the functions of adding jobs to the job queues, or manipulating or displaying jobs currently in the queue. This command is particularly useful to administrators that must access the backup admin system remotely in order to handle a backup job that failed and cannot use the graphical administrator application because no Xwindows server is available.

Use the **–L** option to list all queues. This will show the current status of each job, whether currently running, pending, on hold, or failed.

All options except **–L** require a *JobID*. When a JobID is specified, the queue name (server:device) is retrieved from the job information configured using the **sbadmin** application.

Adding a Job to the Queue

Use the **–A** option to add a pre-defined backup job to the queue. If there are other jobs in the same, the added job will be run after all prior jobs in the same queue complete. Note that, if a prior job had failed, the queue will not process any new jobs until the failed job is either restarted or removed. Refer to the Job Queues section for information on monitoring and changing the status of job queues using the Backup Administrator user interface.

When jobs are run from the queue and the Backup Administrator interface is not running, no job status messages will be reported on the screen. Status messages, in this case, will always be reported using the Alternate Notification method (mailed to a user or appended to a text file). The backup output and progress information may later be displayed after running the Backup Administrator and selecting the View Backup Status/Output options.

Note that the **stqueue** command will return as soon as the job is added to the queue. The exit code of this command will be 0 if the job was queued successfully, or 1 if the job could not be queued (usually because the JobID supplied is invalid). If, for instance, the server for the job is not available, the **stqueue** command will succeed, but the job will fail after it is executed by the queuing system, and will remain in the queue until it is restarted or removed.

Killing a Running Job

To kill a job that is currently running, use the **–K** flag. This will sent a kill signal to the backup process. It may take some time for the job to stop since it may be necessary to complete the current operation before the process will die.

When a job configured to write to tape is killed, it will remain in the queue in a “failed” state, which will prevent other jobs in the same queue from starting. Jobs that write to disk image files will be automatically removed from the queue, since “disk” queues will allow jobs to run simultaneously anyway.

When a job is killed, the tape drive, if used, will be rewound to prevent any other backup jobs from being appended to an incomplete backup. If the Tape Overwrite/Retention Policy does not allow overwriting of current (unexpired) backups, any subsequent jobs started after removing the failed job from the queue will fail with an overwrite protection error. If this is the case, use the stremovelabel program to expire the failed backup label, then restart the queue to allow other jobs to continue.

Removing a Job from the Queue

Use the **–D** flag to delete a job from the queue. Removing a job from the queue will not delete the original job information, but only removes it form the queue. It can be resubmitted using the **–A** option if desired.

You cannot delete a job from the queue if it is currently running. To delete a job already running, first kill it with the **–K** option.

Displaying Job Output

If a job has failed, it may be necessary to display the output messages of the backup command to determine the cause. To do so, specify the **–S** option. The progress indicator (indicating the percentage of backup data written to the media), output and error messages of the backup command will be displayed to standard output.

If this option produces no output, then the backup command had not yet started. If the job had failed, it is due to a pre-backup error, such as a tape overwrite protection error, or because a device or server was unavailable. In these cases, the message indicating the problem was sent using the notification process defined in the preferences options.

stremovelabel

Purpose:

Remove a backup label. Also removes the associated backup if written to disk.

Use On:

Administrator

Syntax:

stremovelabel [-f] *backupid*

Options:

-f	If the backup label to remove is a disk backup, this flag is required to force the removal of the backup from disk. If not supplied, disk backup labels cannot be removed. If the backup was to tape or virtual device, this flag is ignored.
backupid	The *Backup Label ID* of the label to be removed.

Details:

The **stremovelabel** command provides a way to remove a backup label, and associated disk backup images from the command line or from a script. If, for instance, you have a routine which watches for errors in the backup status reporting, that routine can automatically expire the backup label associated with the backup to allow other backups to be written over the same tape, or to free up the disk space used by a disk backup.

Important note: Once a backup label has been expired, it will not be possible to verify, restore, or install a client from this backup. If a backup has been expired or the label history has been inadvertently removed from the system, it is still possible to rebuild this information. Refer to Rebuild (unexpire) a Backup Label from Tape for details.

strestore

Purpose:

Verify or restore data from a local or remote device to the local system

Use on:

Administrator or Client System

Syntax:

strestore [–**fpv]** -**s** *server* [-**D** *destination*][-**S** "*seqlist*"]
-**d** {*device* | *directory* -**L** *backupid*} {-**t** *datatype*} *datalist* …

Options:

-**c** smbclient — SMB client to restore data to (if this is an SMB host)

-**d** *device* — *device* on the *server* (i.e. "rmt0") to restore from if tape or virtual device

-**d** *directory* — *Directory* on the *server* if restoring from a disk backup. If supplied, you must also supply the **–l *backupid*** option.

-D *destination* — Logical volume name or directory into which data should be restored.

-f — Indicates files should be listed while verifying or restoring.

-**L** *backupid* — The *Backup ID* of the backup to restore if restoring from a disk backup. Not used if restoring from tape or virtual device.

-**p** — Indicates the tape is already positioned at the beginning of the desired backup. Prevents rewinding the tape before attempting the restore.

-**s** *server* — *Server* on which the backup media is attached. If not used, a local device or directory is assumed. This option is used only on clients in a *Network Administrator* environment.

-**S** "*seqlist*" — If verifying from tape or virtual device, you may specify a list of backup *sequence numbers* to verify. If restoring, you may specify only one. The default is sequence number 1 if this flag is omitted. Display the backup label for a list of backup sequence numbers on the tape.

-**t** *datatype* — Specify the *type* of data to restore from the backup. If omitted, the entire contents of the backup will be restored. *Datatype* must be one of:
-**V** (volume group)
-**F** (filesystem)
-**L** (logical volume)
-**D** (directory)
-M (meta-disk) (***Linux*** only)
-**R** (regular file)
-P (partition) (***Linux*** only)
-S (SMB [Windows] share) - use with **–c**

-**v** — Indicates a verify is performed. If omitted, a restore is performed.

datalist … — Indicates one or more data elements of the specified *datatype* to be verified or restored. You may specify "all" to indicate all data of the corresponding *datatype* is to be restored.

Details:

The **strestore** command is used to restore data from any SBAdmin Backup. When restoring data, the data will be restored to the system from which the command is executed.

Specify the **–v** flag to perform a verify of the backup by reading through the contents. If the **-v** flag is not specified, a restore of the data is assumed.

The *datatype* must be specified using the **–t** flag. The type of data to restore must be one of the following:

Type	Description	Systems	Data to Specify
V	Volume Group	any	Volume Group name(s)
F	Filesystem	any	Filesystem mount point(s) (i.e. "/home /data")
L	Logical volume	any	Logical volume name(s) (i.e. "lv00 lv01")
M	Meta-disk	Linux	Meta-disk names (i.e. "md0 md1")
D	File/Directory	any	Directory names (i.e. "/home/sam ..")
P	Partition	Linux	Raw partition names (i.e. "sda3 sdb5")
S	SMB Share	Win/Mac	Share (shared folder) name

You must specify the *datalist* as shown above which corresponds to the *datatype* you are restoring. Any elements of the datalist that are multiple words must be surrounded by double quotes (i.e. `strestore –d st0 –tD /shared "/root/User Docs" /mail`).

Volume Group and Logical volume restores are only available on Linux systems if LVM (Logical Volume Manager) is installed. Meta-disk restores are only available on Linux if meta-disk (Software RAID) support is installed.

You may also specify a backup sequence number with the **–S** flag if there is more than one backup stacked on the media. If the backup was created using a backup job which contained multiple clients, each client backup will be a separate backup sequence number. If you performed multiple backups to the same media without rewinding, then each new backup will comprise a new backup sequence number. The data will be read from the backup corresponding to the specified backup sequence number.

If the tape is already positioned to the start of the backup sequence number you wish to read, you may also specify the **–p** flag in place of the **–S** flag to indicate that you wish to read from the backup at the current tape position.

You may specify either a tape or virtual device using the **–d** flag (i.e. "rmt0"). If you are restoring from a disk image backup, you must specify the directory in which the backup was written as well as the **–L** flag followed by the ***Backup ID***. Refer to the Backup Labels section for details on the backup IDs.

If you want to restore the data to a different *destination* than it was originally read from, specify the **–D** flag followed by the destination. The destination must be a logical volume name (i.e. "lv00") for logical volume restores, or a directory name for all other restores. For all restores, except when restoring single filesystems, the data will be restored to the new directory relative to the original full pathname of the files. For instance, the */home/roger* directory, when restored to the */tmp* directory, will be restored as */tmp/home/roger*. This prevents files by the same name, but in different directories, from being restored over one another.

When restoring only a single filesystem to a new *destination*, the files will be restored relative to the original filesystem mount point. In this case, a file "*/home/file1*" in the /home filesystem, when restored to the */home1* directory (or filesystem), will be restored as */home1/file1*. This allows filesystem data to be moved from one filesystem to another.

The **strestore** command will exit with a completion message and one of the following exit codes:

0	Successful
1	Error reading or writing backup data
2	Error occurred prior to reading or writing data
3	Completed with warnings – one or more files may not have been restored
9	Internal program error
99	Killed by user (Ctrl-C)

struncopy

Purpose:

Copy backups between servers, directories or tape devices.

Use on:

Workstation or Network Administrator

Syntax:

struncopy [**-eErR**][**-b** *bufsize*] [**-h** *hostperm*][**-l** *inputBackupID*][**-n** *starting*SeqNum]
[**-N** *endingSeqNum*][**-s** *inputServer*][**-S** *outputServer*]
{**-d** *InputDeviceOrDirectory*} {**-D** *outputDeviceOrDirectory*}

Options:

-b *bufsize* — Buffer size to use for new backup. Default is to keep same buffer size of original backup.

-d *device/directory* — *Source* device or directory containing backup to copy

-D *device/directory* — *Destination* device or directory to write new backup to

-e — Eject tape from the *input* device when copy is complete

-E — Eject tape from the *output* device when copy is complete

-h *hostperm* — If writing to a disk file, indicates host read permission (**a**=any host, **h**=original host only). Default is "**a**" if copying from tape, and defaults to the original host permission setting if copying from disk to disk.

-l *backup ID* — Backup ID to read if copying from a disk backup

-L *backup ID* — If writing to a disk backup and appending to an existing backup label, specify the backup id of the backup to append. Not used when appending tape backups.

-n *startingSeqnum* — Beginning backup sequence number to copy from source if there are multiple backups on the media.

-N *endingSeqnum* — Ending backup sequence number to copy from source if there are multiple backups on the media.

-r — Rewind the input tape device before starting the copy

-R — Rewind the output tape device before starting the copy

-s *inputServer* — Server name if input device or directory is on a remote server

-S *outputServer* — Server name if output device or directory is on a remote server

Details:

The **struncopy** command is used on an administrator system to copy backups from one server to another, from any backup media type to another. Any backup may be copied from tape or disk directory to any other tape or directory. Backups originating from different backup media may even be appended onto the same output media.

When copying a disk backup, you must specify the directory where the backup is stored with the **–d** *directory* option, and the backup id with the **–I** *backupid* option. When a new backup is created (either by writing to the beginning of a tape or writing a backup to disk), a new backup ID is generated automatically. If a backup is appended to an existing tape backup, this backup will be appended to the same label, and therefore will use the same backup ID as the previous backups on the tape.

By default, all backups will be copied from the source media to the destination media. This includes all clients (if backup was made from a *Network Administrator* and included multiple clients), or multiple jobs (if multiple jobs were appended to the source tape media). Every ***Backup Label*** contains at least one ***backup sequence number***, starting with 1 and ending with the last backup written (one for each client backup appended to the same label/media). To determine the backup sequence numbers within a backup label, use the command:

stprintlabel –b *backupID*

If copying to a tape device, you may indicate if you want to ***rewind*** before starting the backup and if the tape should be rewound and ***ejected*** at the end of the backup. If you do not rewind at the start of the backup, you may append the source backup to the end of the destination media (if the destination media is currently at the of volume). The destination ***backup label*** will be appended with the selected source backup. This is commonly referred to as ***stacking backups*** to tape.

You may alter the **buffer size** of the backup by entering a buffer size (in Kbytes) using the **–b** *bufsize* option. This is quite useful in increasing the performance of backups when writing to different media. For example, the default 64K buffer size may be adequate when you wrote your original disk backup file, but when copying to a high-speed tape drive, a higher buffer size (i.e. 256K) may provide much greater backup performance. To use the same buffer size for the destination as was used for the source, do not specify a new buffer size.

If using *Network Administrator*, and the destination backup is written to a disk directory, you may also change whether only the original client host or any host may read the backup data by using the **–h** *hostperm* option. If copying from tape to disk, the default is to allow all hosts (**-h a**) to read the backup. If copying from disk to disk, the default is to leave the original setting unchanged. To allow only the original host which wrote the backup to read it, use **–h h**.

You may use a ***virtual device*** configured as a ***random tape library*** for the output device, but not the input device. This is because the command is only able to track the tape positions of one library at a time. Therefore, if you need to copy from a random library, you will need to specify only the tape device name as the input device. You will be prompted to change the volumes manually on the source device, but the destination device, if a random library, will change tapes automatically. Note that virtual devices configured as ***sequential autoloaders*** may be used for either source or destination devices.

strunjob

Purpose:

Run a backup job in the foreground

Use on:

Administrator

Syntax:

strunjob *JobID*

Options:

JobID Indicates the Job ID previously configured using the Backup Administrator user interface that you wish to run

Details:

The **strunjob** command will perform the backups for all clients configured for the job just as it would when executed from the job queues. The job status information, command output and error messages, and backup label information will be saved for future viewing using the View Backup Labels from the Main Screen. The command output and error messages will also be sent to standard output and standard error, shown on the screen by default. If you wish to save this output to a file, you may type the command as follows:

```
strunjob JobID > filename 2>&1
```

The **strunjob** command will continue running until all clients in the job have been backed up or until an error causes the command to terminate. The command will exit with one of the following return codes, indicating the status of the job:

0	Job completed successfully
1	Job terminated because a client backup failed. Since partial data has been written to the media, no additional client backups have been started. Failure of a client backup may be due to a media write error or a network failure.
2	Job terminated due to an error occurred while saving the backup command output, status information, or label information.
3	Job terminated due to an error in pre-processing tasks. This might include such things as a backup device already in use or the backup server or network unavailable.
4	The client backups completed, but a post-processing error occurred. This might include the inability to update the backup label or history information after the backup completes, or an error ejecting the tape from the drive when the job requires the tape to be ejected when complete.
5	The client backups completed, but a minor error occurred in post-processing that should not affect the backup history or label, nor should it affect other jobs appended to the same media.
6	The job completed successfully, but one or more clients in the job were not included in

	the backup. The reason may be because the client or network was unavailable or because the client could not reach the server over the network. This error will also occur if a pre- backup program was set to run on the client and failed.
7	Job terminated without writing any data to the server because the backup media was either write-protected, or contained a prior backup label that could not be overwritten due to the tape overwrite/retention policy setting.
99	The command terminated because the process was killed (SIGINT or Control-C)

When running the **strunjob** command at the command line, the status message of the jobs will not be reported using the backup status reporting method defined in the Backup Administrator. The only exception is when previous backup labels are expired based on the backup overwrite/retention policy defined.

strunrest

Purpose:

Verify or restore data from any server to any client

Use on:

Network Administrator

Syntax:

strunrest [**–vP**] –**s** *server* –**f** {*device* | *directory* –**l** *backupid*} {-**t** *datatype*}[-**d** "*datalist*"]
[-**D** *destination*][-**h** althost][-**S** "*seqlist*"]

Options:

-c *client*	*Client* to restore the data to. Only used if restoring data.
-**d** "*datalist*"	If *datatype* is specified, you should also specify the *list of data* to verify or restore of the specified *datatype*. If this argument is omitted, then all data of the specified *datatype* will be verified or restored.
-D *destination*	Destination logical volume name or directory to restore data.
-**f** *device/dir*	Tape or virtual device on the *server* (i.e. "rmt0") if restoring from tape. Specify a directory on the server if restoring from a disk backup. If restoring from a disk image backup, you must also supply the **–l** *backupid* option.
-**l** *backupid*	The *Backup ID* of the backup to restore if restoring from a disk backup. Not used if restoring from tape or virtual device.
-**p** *curSeqnum*	If the tape is already positioned at the beginning of a specific backup, you may specify the *current sequence number* of that backup to prevent rewinding and re-forwarding of the tape.
-**P**	Indicates the *progress indicator* should be shown, displaying the progress of the verify or restore. The progress indicator is send to standard error.
-**s** *server*	*Server* on which the backup media is attached
-**S** "*seqlist*"	If verifying from tape or virtual device, you may specify a list of backup *sequence numbers* to verify. If restoring, you may specify only one. The default is sequence number 1 if this flag is omitted. Display the backup label for a list of backup sequence numbers on the tape.
-**t** *datatype*	Specify the *type* of data to restore from the backup. If omitted, the entire contents of the backup will be restored. *Datatype* must be one of: -**V** (volume group), -**F** (filesystem), -**L** (logical volume), -**D** (directory) or –**R** (regular file).
-**v**	Indicates a verify is performed. If omitted, a restore is performed

Details:

The **strunrest** command may be used to either verify or restore the contents of a backup. This command is only run on the admin system, although the backup data may exist on any server and may be restored to any client.

You must specify both the *server* and *device* to read, and the *client* to restore to (if restoring). If you are restoring from a disk backup, and the backup was created with read permission only by the original client from which the backup was made, only the original client may be restored to. This is to prevent the backup file from being read by other hosts. However, if you wish to change the permission of the backup file to allow it to be restored to a different host, use the option Change Read Permission of a Disk Backup.

Note that the *datatype* and *datalist* arguments must correspond. For instance, if you select to restore a filesystem (**-t F**), then you must supply a list of filesystems to restore ("*/home /tmp*") as they are defined on the backup. If any of the filesystems supplied do not exist on the backup, no restore will take place. Note also that you must surround multiple restore options with double quotes.

You may also specify a sequence list (*seqlist*) using the **–S** flag if there is more than one backup stacked on the media. If the backup was created using a backup job which contained multiple clients, each client backup will be a separate backup sequence number. If you performed multiple backups to the same media without rewinding, then each new backup will comprise a new backup sequence number. When verifying backups, you may specify one or more sequence numbers, surrounded by double-quotes. When restoring data, you may specify only one backup sequence number. The data will be read from the backup, or backups, corresponding to the specified backup sequence number list.

Important note: If you are restoring a volume group or filesystem from an incremental backup level 0 and you do not specify a *datalist* to restore, then all files will be restored. If the incremental level is 0, all files currently in the corresponding filesystems will be removed before the restore takes place. This is to ensure that the filesystem, when completed, will contain ONLY the files that existed when the backup was created. If you do not want to remove existing files before restoring, then you should not restore a filesystem or volume group. Instead, select "directory" as the data type (**-t D**), then specify the list of directories to restore.

If you supply both the **–L** and **–P** flags, you should redirect either standard output or standard error to a file. Otherwise, both will be shown on the screen and the progress indicator data will be intermixed with the file list.

If you want to restore the data to a different *destination* than it was originally read from, specify the **–D** flag followed by the destination. The destination must be a logical volume name (i.e. "lv00") for logical volume restores, or a directory name for all other restores. For all restores, except when restoring single filesystems, the data will be restored to the new directory relative to the original full pathname of the files. For instance, the */home/roger* directory, when restored to the */tmp* directory, will be restored as */tmp/home/roger*. This prevents files by the same name, but in different directories, from being restored over one another.

When restoring only a single filesystem to a new *destination*, the files will be restored relative to the original filesystem mount point. In this case, a file "*/home/file1*" in the /home filesystem, when restored to the */home1* directory (or filesystem), will be restored as */home1/file1*. This allows filesystem data to be moved from one filesystem to another.

The strunrest command will exit with a completion message and one of the following exit codes:

0	Success
1	Failed verifying or restoring data

9	Internal program error
99	Killed by user (Ctrl-C)
101	The **strunrest** command failed after verify/restore started
102	Syntax error calling the **strunrest** command
103	The **strunrest** command failed with an error prior to verify/restore starting

stserver

Purpose:

Add, change, list or remove server information

Use on:

Network Administrator

Syntax:

stserver [**-aAlrv**][**-c** “*clients*”][**-B** *directory*][**-C** “*dirs*”][**-D** “*dirs*”]
[**-L** *dir*][**-d** “*devices*”][**-i** *ipaddr*][**-I** *ipaddr*] [*server* …]

Options:

-a Add or change the selected *server*

-**A** Show server availability (when **–l** is used)

-B *directory* Directory for network boot images

-c “*clients*” *Clients* to allow access to write to this server (default is “all”)

-C *“directories”* Directories for CLIENT System Backups

-D “*directories*” *Directory(s)* for regular (non-system) backups

-d “*devices*” Tape drives or virtual devices clients may write to (default is “all”)

-i *ipaddr* Alternate network IP address for clients to use when sending backup data to this server

-I *ipaddr* Alternate network IP address for clients to use when client performs a network install from this server

-l List servers. If *server* provided, list only specified server.

-L *directory* *Directory* for LOCAL system backups

-r Remove specified *server*

-v Show verbose listing of servers, or specified *server*

Details:

The **stserver** command is used to add, remove, change and list servers configured on the Network Administrator.

To list currently configured servers, use the **–l** option. By default, only server names are shown. To show all of the attributes of a server, use the **–v** flag also.

If you want to also show the server availability as well as the system type, use **–A** with **-l**. This will indicate the server availability as either available or down, and will display the system (*AIX* or *Linux*) and machine type (*i386*, *ppc*) or “unknown” if the server cannot be contacted.

Availability is determined by whether or not the network admin can communicate with the server. If not, you can use the **stcheck** command to help isolate the problem.

If you want to configure a new server, use the **–a** option and specify the *server* name at the end of the command. All attributes will be set to their default values (or none) unless explicitly set. After adding a server, be sure to list it with the **–vl** options to display the full list of settings.

NOTE: Be careful to use quotes around any argument to flags which contain multiple words.

Changing a server

If a server currently exists, you can change specific options by specifying only that option and its new argument. For example, to set the alternate network adapter (hostname) to use for backups, use the following command:

```
stserver –ai 192.168.1.1 buserver
```

Removing a server

To remove a currently configured, server, specify **–r** and the *server* name, such as:

```
stserver –r buserver
```

26. SMIT Options

The information described here is supported only on *AIX* systems.

Although all backups, verifies and restores may be performed from the Admin System, it may sometimes be desirable to perform an operation directly from the client itself. This is necessary, for instance, if the administrator application or system is unavailable and it is necessary to restore data to a client.

Many options are available using the **AIX System Management Interface Tool (SMIT)**. These options are used only performing operations on the client itself. There are no configuration options available when the client is configured for use under a network administrator as the configuration options must be performed from the administrator system.

To access the SMIT menus, type "`smit`" and select ***Storix System Backup Administrator***, or just type "`smit storix`" at the command line. If *AIXwindows* is running, the graphical version of SMIT will appear by default. If AIXwindows is not running, the text (ASCII) version will appear instead.

Since detailed instructions are provided in the SMIT menus and options themselves, they are not provided here. Press the ***F1 key*** (ASCII SMIT) or ***Help button*** (graphical SMIT) at any SMIT menu option or entry field for detailed information or instructions on that option.

Hint: If you prefer using the text (ASCII) version of SMIT over the graphical version (as we do), type "`smitty`" instead of "`smit`".

The SMIT options provided on the main menu are as follows:

Software License Maintenance: These options are for configuring the software on the client. After initially configuring the software on the client, these options need not be used again unless you wish to change the license type installed on the system or other software configuration set during the software installation.

Set or Change Network Administrator: Use this option if the hostname or IP address of the network administrator has changed. This option defines the network administrator that is given permission to run backups on the local client or server system.

Client Configuration Options: Use this option to define the backup servers and local directories that may be used to backup this client from the SMIT menus. Even if servers are already defined for the client from the administrator system, it is necessary to define them using this option for use under the SMIT menus. An option is also provided here for defining the local directories that may be used for backups using SMIT.

Perform a Backup of this System: Use this option to perform a backup of the local system to a local tape or to a directory or server defined in the **Client Configuration Options** above.

Verify Backup Data: Use this option to verify a backup previously created. The backup may be a local disk image or tape backup or a backup on the disk or tape of a defined backup server.

Restore Data from a Backup: Use this option to restore data from a backup. The backup may be a local disk image or tape backup or a backup on the disk or tape of a defined backup server.

Help: Use this option for general help information on all of the options provided in SMIT. You may also press the ***F1 (Help)*** key at any SMIT menu option or entry field for detailed information or instructions on that option.

Note that most commands performed from SMIT may also be typed at the command line on the client. To obtain the command and options, fill out all of the entry fields required, then press the ***F6*** or “***Command***” button to show the actual command being performed.

27. Network Security

SBAdmin was created with safeguards in place to prevent breaches in security without disrupting the security and integrity of the remaining network. This section outlines the flow of network traffic, the security measures that have been implemented, and what steps need to be taken by security personnel to insure that your software will function properly between network firewalls.

TCP/IP Ports

SBAdmin configured with a *Network Administrator* license communicates via the **Transmission Control Protocol/Internet Protocol (TCP/IP).** This communication is handled through two different ports, the **Dataport** and the **Statusport**. By default, the SBAdmin uses port numbers 5026 and 5027 which are registered with the ***Internet Assigned Numbers Authority*** (previously used 8191 and 8192). These ports numbers are determined during the installation of the software and can be changed by the user at that time. If you need to change the port numbers used, simply reinstall the software and update the port numbers at that time. If you change your port numbers, previously made boot images on tapes, CDs, or floppies will attempt to communicate through the old port numbers. It is advised to create your boot media/images after changing your port numbers.

NoTE **It is very important that all Administrators, Servers and Clients using System Backup Administrator are configured to use the same port numbers. You can verify this by checking in the** /.stdefaults **file for the following entries:**

```
DATAPORT=5026
STATPORT=5027
```

These two ports are *listening ports* and must be open to incoming TCP/IP traffic from other systems within your SBAdmin network. SBAdmin uses the ports specified above to transfer backup data, status messages, and to run remote commands. Only the SBAdmin network daemon process "**strexecd**" can properly answer requests on these ports. Any other process attempting to open these ports will receive a connection error.

Network Firewalls

When a backup or restore is performed remotely, commands are initiated between the *Admin* and *Client* as well as the *Client* and *Server*. The network communications on these ports are setup automatically when SBAdmin is installed on any system. If you have a network firewall between any of your systems utilizing SBAdmin, you will need to open the communication on these ports, or select other port numbers to use that are allowed by the firewall.

Some firewalls will close inactive ports after a certain period of time. It is advisable to turn off this timeout, if possible. When performing a remote backup, volume prompt messages are sent over the network, and no other communication takes place until a new tape volume is inserted. If the next tape is not inserted before the firewall timeout, the firewall may close the ports. SBAdmin will continue the backup, but no further messages will appear and SBAdmin will not receive the exit status of the command. Although the backup usually completes successfully, SBAdmin will appear to have hung.

Remote Command Execution

SBAdmin is the only application that can communicate over the SBAdmin ports. In addition, only specific commands can be run remotely. All attempts to run remote commands are checked for authenticity as follows:

NoTE **In the following, $STXPATH designates your SBAdmin data directory chosen when installing the software (default is /storix), and $STXINSTPATH is the SBAdmin application directory (/opt/storix for Linux or /usr/lpp/storix for AIX).**

1. The IP address of the sender is checked to see if it is a valid admin system. Valid admin systems are specified in the **$STXPATH/config/admin_servers** file when SBAdmin is installed onto a client or server system.

2. The IP address of the sender is checked to see if it is a valid client (if communicating to a server). The **$STXPATH/config/remote_access** file determines the permitted hosts, and this file is created by the network administrator system and copied to each server when changes are made to the server information or new clients are added to the configuration.

3. The command to execute is checked to ensure it is not a wrapper. For instance, no commands containing sub-commands such as "command1; command2" or "command1 $(command2)" may be executed.

4. The command to execute is checked to ensure it does not contain an absolute pathname. Only the command name to execute must exist on the system in the **$STXINSTPATH/bin** directory.

5. The command to execute is checked to see if it a permitted remote command. Permitted commands are listed in the **$STXINSTPATH/config/remote_cmds** file. Programs listed here may not have a leading PATH, but the commands themselves must exist in the **$STXINSTPATH/bi**n directory.

6. For user-customized pre and post-backup commands, the commands must exist in the **$STXPATH/custom** directory, must be writeable only by root and must be executable.

7. When executing command to read or write tape drive (**stio** or **sttape** command) or disk image backup files, the **$STXPATH/config/remote_access** file is checked to ensure the specified device or directory is accessible by the calling host.

Note that all of the above configuration files and directories may only be written by the root user on the system.

Encryption Keys

Encryption keys are entered on the client system using the stkeys command. This prevents the encryption keys from being passed across the network in any form. The encryption keys are stored in a file on the client system, unreadable by any user other than "root", and neither the file, nor the information therein is ever sent over the network.

SMB (Windows) Usernames and Passwords

Because an SMB Host (Linux system) is required to remotely access the data in a shared folder on an SMB (Windows) client, SBAdmin uses the Linux "***smbclient***" command installed with the ***SAMBA*** client software. The *smbclient* command requires the username and password of the SMB client to list the SMB shares available. The "mount" command is used to mount a remote filesystem using the "***smbfs***" filesystem type. This mount command also requires the username and password be used to remotely access the SMB shared data. Therefore, the username and password of the SMB share may be passed across in the network in un-encrypted form between the SMB Host and the SMB Client. This process is beyond the control of SBAdmin.

When configuring an SMB Client, you enter the username and password, which are saved in a protected file on the Network Administrator system. Each time a command must be executed to list or copy files to or from the SMB Client, the username and password are passed from the Network Administrator to the SMB Host, but only in encrypted format. The SMB client's usernames and password are never saved on the SMB host.

28. Getting Help

QuickHelp

If you are uncertain of the use of a particular button, listbox or entry field, you may at any time move the cursor over the object in question and press the right mouse button. A popup message will appear on top of the object with information on its use and any options, warnings or special instructions that might apply. Information is provided for every selectable object in the application. After reading the message provided, you may click any mouse button anywhere on the screen to remove the QuickHelp message and continue as usual.

Always use the QuickHelp as your first step in understanding or resolving a problem!

User Guide

This user guide may be displayed at any time from the Backup Administrator user interface by selecting Help→User Guide from the menu bar. This user guide contains links so that clicking on any underlined text will move you immediately to the referenced section of the text.

When selecting this option, a PDF viewer application will be started in order to open the user guide file, which is in *Portable Document Format* (PDF format). The viewer provided with this application is a simple viewer to save space, and does not provide all the functions of some larger PDF viewers. If you prefer to use another PDF viewer, simply set the VIEWER environment variable to the name of the viewer application file before starting the Backup Administrator. If the specified program exists, the user guide will be opened using this program.

Communications Errors

A tool is available to help diagnose problems in communicating between the *Network Administrator* and clients or backup servers. If the client or server icon on the Main Screen shows a red symbol (indicating the client or server is unavailable), or if an error occurs such as “Client host may not be contacted”, run the **stcheck** command to help determine the cause. Refer to the stcheck command for details.

Storix Support

Should you encounter a problem using **System Backup Administrator** or have any questions, numerous support options are available. Select **Help→Storix Support** from the menubar to display current information on support options available to you. This will provide you with links to obtaining online and telephone support, hints and tips, etc.

Index