SBAdmin User Guide

Version 8.2

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 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/
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
DropBear	A Smallish SSH 2 Server and Client	Copyright 2002, 2003 Matt Johnston http://www.matt.ucc.asn.au/dropbear/dropbear.html
GRUB	Grand Unified Bootloader (GNU GRUB)	Copyright 1989, 1991 Free Software Foundation, Inc. http://www.gnu.org/software/grub/grub.html
Lighttpd	Secure, fast, compliant and flexible web-server	Copyright 2004 Jan Kneschkle, incremental http://www.lighttpd.net
OpenSSL	Toolkit implementing Secure Socket Layer	Copyright 1998-2008 The OpenSSL Project Copyright 1995-1998 Eric A. Young, Tim J. Hudson http://www.openssl.org
Xpdf	PDF Document viewer (for AIX)	Copyright 1996-2003 Glyph & Cog, LLC. http://www.foolabs.com/xpdf
bpgetfile	RPC Bootparams client (for Solaris)	Copyright 2000 Rensselaer Polytechnic Institute, Department of Computer Science
parted	GNU parted	Copyright 2007 Free Software Foundation, Inc. http://www.gnu.org/software/parted
ELILO	Linux boot loader for EFI/x86_64 based systems	Copyright 2000-2003 Hewlett Packard Co. Copyright 2006-2010 Intel Co. ftp://ftp.hpl.hp.com/pub/linux-ia64
btrfs-progs	Btrfs utilities programs	Copyright 2007 Oracle Copyright 2012 STRATO AG http://www.btrfs.wiki.kernel.org

*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

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

AIX: All IBM *RS/6000*, *System p*, System i, *OpenPower* and *JS/20* systems running AIX Version 5.1 and later (currently 7.1).

Solaris: **x86 and x86_64**: All Solaris 9 versions 9/05 and later (32 and 64-bit platforms), and Solaris 10 versions 1/06 and later (32 and 64-bit platforms), Solaris 11 Express (64-bit platforms), and Solaris 11 version 11/11 and later (64-bit platforms).

SPARC: All Solaris 9 versions 9/05 and later, Solaris 10 versions 11/06 and later, Solaris 11 Express, and Solaris 11 versions 11/11 and later. Includes *sun4u* and *sun4v* platforms.

Linux: ***x86 and x86_64***: All distributions which run on *Intel 32-bit* based processors and 64-bit processors software (includes *AMD*, *Opteron* and *Athlon*-based systems). Linux kernel levels 2.4 and glibc 2.2.5 and higher are required. Support is provided for Linux LVM Library version 1.0 and higher, and Software Raid Devices (meta-disks) when installed. UEFI is supported on x86_64 systems running 2.6.21 or later kernel levels, CONFIG_EFI enabled in the kernel, and support for creating VFAT filesystems.

PPC (IBM PowerLinux): All distributions supported 64-bit systems with *PowerPC CHRP* hardware. Linux kernel levels 2.6.16 and higher, and glibc 2.4.2 and higher are required. Support is provided for Linux LVM Library version 1.0 and higher, and Software Raid Devices (meta-disks) when installed.

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 Edition* 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:

Workstation Edition	The *Workstation Edition* license provides all available backup and recovery features for backups using local tape, local disk, or a local NFS mount. This includes all available features needed for standalone system backups. Backup types include **Full System**, **Filesystem**, **Directory**, **Logical Volume** (AIX/Linux), **Meta-disk** (Solaris/Linux), **Partition** (Linux), **Slice** (Solaris) and **ZFS Volume** (Solaris). Many additional features more commonly used in a commercial environment, such as incremental backups and tape libraries, are also provided.

Network Edition	The *Network Edition* license is only installed onto the system from which network backups will be centrally managed. This system is known as the *Network Administrator*. Typically, there is a single *Network Administrator* in a backup environment. 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 remote systems configured with a *Client/Server* license to be managed by the local *Network Administrator* system. A license key is required on the *Network Administrator* system, 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 the Network Administrator may manage them.
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 Edition* license. If installed separately, this client must be controlled by a *Network Administrator* system. 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 Edition* license. Backup management features, such as scheduling and history reporting are provided only on the *Network Administrator* system.
TSM Edition	This license is installed onto a system from which only ***TSM (Tivoli Storage Manager)*** backups will be managed. This edition differs significantly from the types of backups and backup devices that are supported by the *Network* and*, Workstation Editions*, and are therefore documented in a separate user guide. If you want to extend the *Network Edition* to support backups to TSM servers, then refer to the **TSM Client Feature** below. For *TSM Edition*, please refer to the ***System Backup Administrator TSM Edition User Guide*** for details on this license and compatible features.
Backup Data Encryption Feature	This optional license may be added to a *Network Edition, TSM Edition* or *Workstation Edition* to enable AES data encryption support for all backups. If used with *Network* or *TSM Edition*, a license is purchased for the number of *Clients* for which backup data should be encrypted.
TSM Client Backup Feature	This optional license may be added to a *Network Edition* license to add the ability to perform ***TSM client (node)*** backups to a ***TSM server***. While the *TSM Edition* license allows only **Full System Backups**, adding this option to a *Network Edition* will allow you to perform any type of SBAdmin backup to a TSM server, while continuing to utilize tape and disk devices as backup media on other SBAdmin backup servers.

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 software packages downloaded from the Storix Software web site (http://www.storix.com) or from a System Backup Administrator installation CDROM:

Downloading and Installing from the Website

1. Select the software package you wish to download from the website 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`

   c. On ***Solaris*** systems: Normally, a CDROM will automatically be mounted to the ***/cdrom/cdrom*** directory when inserted. If this is the case, replace ***/mnt*** with /***cdrom/cdrom*** in the following commands. If the cdrom is not auto-mounted, type:

      ```
      mount -F hsfs –o ro /dev/dsk/c1t0d0s0 /mnt
      ```
      (where c1t1d0s0 is an example of your cdrom drive name).

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

Updates to new versions (i.e. 7.2 to 8.1) will require that you follow the instructions above as though you installed the software for the first time. For most other updates, the administrator system and all configured clients and servers may be updated from the user interface.

To update the software with access to the internet, you may automatically check, download and apply updates directly from the Storix website 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 and TSM Editions*):

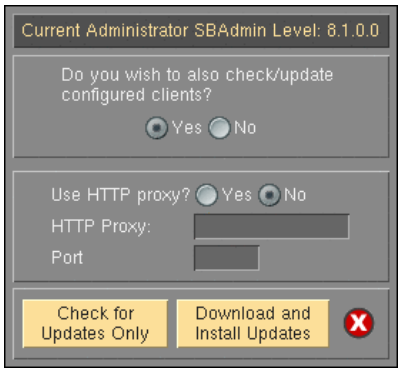

You will have an option of checking for updates only and/or downloading and installing updates. If using *Network Edition*, you will have an additional option of automatically applying updates to configured clients.

If the system cannot contact the Storix website 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 administrator system using the "**stinstall**" command described above, you will have the option to install the new software level onto configured clients and servers.

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

Starting the Software

There are three interfaces available for performing SBAdmin operations:

1. **Graphical-User (Xwindows) Interface (GUI)** - This is the interface described in this user guide.

   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.

   If you are not logged into the system as the "root" user, then you are prompted for a SBAdmin username and password to start the application. When SBAdmin was first installed, an "**admin**" username was configured, and you were prompted for the password for this user. If, however, you are logged onto the system as "root", you will only be prompted for the username if more than one is configured.

   If multiple *groups* are configured, the application will automatically start with the group the current user was last logged into.

When multiple groups or users configured, you can specify the group or user/password at the command-line with the following flags:

```
# sbadmin –G group –U user/passwd
```

For more information see the SBAdmin Commands Reference.

Refer to Configuring Users and Configuring Groups chapters for more information.

When starting the GUI interface, the Main Screen will appear.

Web-Based Interface – This interface is accessed through a web browser and is designed to be similar in use and function to the GUI interface. By default, the web interface is not configured when installing the software. Therefore, you must have enabled the web interface when installing the software on the *Administrator System* in order to access it.

To start the web interface, you must point your browser to the hostname or IP address of the administrator system and the port number (default 8080) configured when you installed the software (note the use of *https* instead of *http*):

```
https://adminhost:8080
```

You will be taken to a page where you must enter your SBAdmin username and password to access the application. When SBAdmin was first installed, an "**admin**" username was configured, and you were prompted for the password for this user.

If you forgot the username or password, you can log into the admin system as the "root" user and use the Xwindows (GUI) or command-line interface ('stuser' command) to change the password of the SBAdmin user.

Refer to the ***SBAdmin Web Interface Install Guide*** for additional details on installing, configuring and starting this interface.

2. **Command-Line Interface (CLI)** – This refers to running commands provided with this software at the shell prompt. Commands may only be run on a *client* or *backup server* system when logged on as the ***root*** user. On the administrator system, commands can be run when logged onto the system as a non-root user, but you will be prompted for the username and password of a configured SBAdmin user with the proper authority to run the command.

   From the admin system, you can run any command without being prompted for the username and password by providing the authentication credentials at the command line. For example, to start a job you would use the "`strunjob JobID`" command. To provide the username at the command line, supply the `–U username/password` flag as the first argument to this or any other command:

   ```
   strunjob –U userid/password JobID
   ```

   When SBAdmin was first installed, an "**admin**" username was configured, and you were prompted for the password for this user.

   Refer to the **SBAdmin Commands Reference** for details on each available command.

Enabling Optional Features

Optional features, such as ***Backup Data Encryption***, and ***TSM Integration***, may be enabled after the admin 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.

Initial TSM Setup

This document describes TSM client backups only as used with the optional *TSM Client Backup Feature* for *SBAdmin Network Edition*. If using *SBAdmin TSM Edition*, refer to the SBAdmin TSM Edition User Guide instead.

TSM Server

The TSM Server software must be at level 5.2 or later. Refer to the TSM documentation for instructions on checking and updating the TSM server.

SBADMIN Management Class

Before any backup may be performed to TSM, you must define a new management class on the TSM server called SBADMIN. All SBAdmin backups will be stored under this management class. The management class must be defined to disallow versioning of backup objects. The management class must be created using the TSM Integrated Solutions Console or using the following command within ***dsmadmc***:

```
DEFINE MGMTCLASS domain_name policy_set SBADMIN
```

Next, you must turn off file versioning in the ***BACKUP*** **copygroup** of the management class. This is done by setting the number of versions of each file to be kept to “1” and days to retain inactive versions to “0”:

```
DEFINE COPYGROUP domain_name policy_set SBADMIN TYPE=BACKUP \
  DESTINATION=backup_pool VEREXISTS=1 RETonly=0
```

Once the **copygroup** and management class are defined we need to activate the **policyset** with the command:

```
ACTIVATE POLICYSET domain_name policy_set
```

TSM API Client

The *TSM API Client* software is normally installed on each node when the *TSM Backup/Archive Client* software is installed. This is required by SBAdmin. To check if the API client is installed:

On AIX: Ensure the *tivoli.tsm.client.api.32bit* fileset is installed at level 5.2 or later:

```
lslpp –l tivoli.tsm.client.api.32bit
```

On Linux: Ensure the TIVsm-API package is installed at level 5.2 or later:

```
rpm –qa | grep TIV
```

On **Solaris**: Ensure the *TIVsmCapi* package is installed at level 5.2 or later:

```
pkginfo –l TIVsmCapi
```

2. Introduction

System Backup Administrator (SBAdmin) is designed to simplify the administration of backups on the local system or client backups in a networked environment. It does so by combining powerful backup tools with an easy-to-use *graphical* or *web-based* interface for administering backups, boot media, clients, servers and backup devices. Backups created by the *Backup Administrator* application may include single directories or entire systems (operating system and data). Backups of the entire system (referred to as a ***Full-System backup***) can be used to reinstall the source system or another system with an entirely different storage 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 (GUI) or rolling over (Web-interface) the object in question.

Most examples and screen shots used in this document were created using the *Backup Administrator* (graphical) interface used with the *Network Edition* or *Workstation Edition* licenses. Most tasks may also be performed using the Web-based Interface which will look very similar to the GUI. Major differences between the two interfaces will be noted in this document and are also documented in the ***SBAdmin Web Interface Installation Guide***. Most tasks may also be performed directly on a client or server using commands from the command line.

The remainder of this document provides instructions on the use of the *System Backup Administrator* (SBAdmin) graphical user interface. The *Network Edition* is used in the examples throughout this guide. Options which are not applicable to the Workstation Edition are noted.

Most instructions shown here may also be performed using the web-based interface and any compatible web browser. The concepts are the same, but exact instructions differ from when using the GUI interface. When using the web-based interface, refer to the *SBAdmin Web Interface Installation Guide.*

The *SBAdmin Commands Reference Guide* is also available for information on running commands at the command line. Most configuration and maintenance options, as well as performing backups from both the network administrator or from the client systems, may be performed using commands from the command-line.

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 *Network Edition*, 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 running the backup jobs, monitoring backups, performing verifies and restores, and even recreating volume groups and filesystems. For a ***Standalone System*** running *the Workstation 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 or directories, tape autoloaders or libraries. 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 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 an ***AIX, Solaris*** or ***Linux*** (UNIX) client. Select this link for detailed information on adding or removing a backup client.

  Note: When using a *Workstation 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 Edition*, the backup clients and servers, as well as the configured devices 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 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 upon demand, 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 set of backup files (if written to disk). When writing backups 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 each directory 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

Before you can write any SBAdmin backups, you must first configure one or more backup “devices”. A backup device will consist or one or more tape drives or one or more directories on the server. You may name a device anything you wish, but for simplicity, it’s recommended you use the tape drive name (for a single-tape device) or directory path (for single directory devices).

Tape Devices

A tape backup device may consist of a single tape in a single tape drive, multiple tapes from a single tape drive, or multiple tapes 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 or raw storage device (partition, logical volume, metadisk, ZFS device, etc) is stored in a separate backup image. This allows 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 appended to the same tape.

The *Backup Administrator* keeps track of the contents of a tape. At any time, the admin 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 can 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 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 admin may also set the backup retention policy (also known as 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 device, providing increased performance and capacity. There are three types of tape devices you can configure for performing sequential, parallel or multi-copy backups. Tape devices may also be configured as a Sequential Autoloader or a Random Tape Library. Refer to Types of Devices as described in detail in the Devices section for a complete description.

Also, when configuring a client (*Network Edition*), you may specify a tape drive name to be used for local System Backups of that client (often referred to as SBTAPE). This backup device option will then appear when you choose to configure a backup job to run locally on the client, requiring no server or network traffic. See Clients for more details.

Directory Devices

Any backup may be written to a directory on the backup server. In addition, full system backups may be written to the local client's configured System Backup Disk (SBDIR). 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 you configure a directory device for a server (or local system for *Workstation Edition*), you can name the device anything you wish. Usually, you can simply use the directory itself for the name. You may also

choose to use a name such as "*ClientSystemBackups*" if the directory will be used solely for System Backups of clients, which may be used as network install images.

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 admin may also set the backup retention policy (also known as the *overwrite policy*) so that a disk backup that has aged over a certain number of days, or exceeded a certain number to retain, is automatically expired and removed when the same backup job is re-run. These options will prevent filesystems containing disk backups from filling up while ensuring that the latest backups are kept on file.

When you configure a directory device, you can specify multiple directories on the server as the destination. When doing so, a backup that fills the first filesystem will be automatically continued onto the next directory. This assumes, of course, the directories are in separate filesystems, providing the same type of functionality as a *sequential* tape device with multiple drives.

When configuring a client (*Network Edition*), you can also configure a dedicated disk or disks for System Backups of that client (often referred to as a local system backup disk, or SBDIR). This will allow a System Backup of the client without the use of a server or network traffic, and the client can restore data or be completely reinstalled from this local disk backup. In other words, you can boot and reinstall the system from locally-attached (or even SAN-attached) disks, providing complete system backup and recovery using only a local disk. You can also use this disk as portable media, which can be directly attached and used to install other systems. More information can be found in Configuring Local System Backup Disks.

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

NFS Backups

Backups may be written to locally mounted NFS shares. Writing to a locally mounted NFS share minimizes network in that only the backup client and the NFS server are creating network traffic. This feature requires that each backup client, and the admin system, be able to mount the NFS share read-write. It is also not expected that any SBAdmin software is installed on the physical NFS server. More importantly, the configuration of the NFS server that is exporting the share is the sole responsibility of the user.

When using *Network Edition*, you will configure the NFS share to be used by the backup clients using a Shared NFS Server. When using *Workstation Edition*, you will configure a NFS Local System Backup Device.

Backup status, output, and label information may be displayed and managed for backups to NFS just as with directory devices.

TSM Backups

Any backup may be written to a *TSM server*. With TSM backups, each filesystem or logical volume within the backup is stored in a different TSM *object*, so access to the data is much faster than reading the entire backup.

Refer to the *SBAdmin TSM Edition User Guide* to view only the options available when using *TSM Edition*.

The ***TSM Client Backup Feature (for Network Edition)*** provides all backup types and options of a TSM client (node) to a TSM server. In this case, TSM is simply an added backup media option to the existing administrator license. When the ***TSM Edition*** is installed, however, the TSM server will be the only backup media option, and only System Backups are supported. This provides a more cost-effective means for users of TSM to simply add system backup and adaptable system recovery (ASR) to their daily backup routine.

Each backup job will be assigned a unique Backup ID, and each client backup within the job will have a unique backup sequence number. Each TSM backup will have a unique name on the server, so there is no danger of overwriting a prior backup. Instead, the user must expire TSM backups manually to prevent excessive use of disk space. When doing so, not only the backup labels, but also the actual backup objects are removed from the TSM server. The administrator may also set the overwrite policy so that a TSM backup that has aged over a certain number of days, or exceeded a certain number to retain, is automatically expired and removed when the same backup job is re-run. These options will help limit the amount of space on the TSM server required to store backups while ensuring that the latest backups are kept.

Backup status, output, and label information may be displayed for TSM 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 (**AIX/Linux**), partitions (**Linux**), meta-disks (**Linux/Solaris**), disk slices (**Solaris**) or ZFS volumes (**Solaris**). 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 the beginning of a tape, then the tape is also configured to boot to the **System Installation process**.

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 (**AIX**) and/or restored from the backup. Individual files, directories, filesystems or raw logical volumes may be recreated (**AIX**) and/or restored.

   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 (**AIX/Linux**), partition (**Linux**), slice (**Solaris**), or meta-disk (**Solaris/Linux**). ***ZFS filesystems*** are also supported on **Solaris** systems. Files within the filesystems may also be backed up *incrementally*, including only files that have changed from a prior backup. Specific files directories or filesystems may be recreated (**AIX**) and/or restored from this backup.

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 (**AIX**) 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. Slice Backup (***Solaris*** systems only) – This backup may include one or more “raw” disk slices typically containing non-filesystem data. From this backup, only an entire disk slice may be restored.

8. Meta-disk Backup (***Linux*** and **Solaris** systems only) – Meta-disks are often referred to as ***Software RAID*** devices and ***MDs***, ***multi-disks*** and ***meta-devices***. This option is only available for Linux if Software RAID support is installed on the system, and only available on Solaris if Solaris Volume Manager is installed on the system. 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, slices, and even other meta-disks.

9. ZFS Pool – Also called "***zpools***", this backup is typically used to separately backup the ZFS pools that are not part of the operating system. Files or *ZFS volumes* within the pool may also be backed up *incrementally*, including only files within filesystems or volumes that have changed from a prior backup. The backup may contain one or more zpools, and an entire zpool may be restored from the backup. Individual files, directories, filesystems or raw volume data may be restored.

   Zpool backups are only available for ***Solaris*** systems with ZFS installed.

10. ZFS Volume – This backup may include one or more "raw" ZFS volumes. From this backup, only an entire volume may be restored. ZFS volume backups are only available for ***Solaris*** systems with ZFS installed.

It is possible to later restore select 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, partitions, ZFS volumes 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, devices, jobs, profiles, etc. It is also used for the monitoring of job queues, displaying job status, backup output messages, generating reports and backup history.

Scheduled backup jobs will continue to run even if the SBAdmin 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 SBAdmin 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.

This document provides examples using the graphical user interface (GUI). Also, the *Network Edition* is illustrated since many features do not appear when using the Workstation Edition.

Also, the options which appear on the screen will differ depending on the user level of the user running the application. To illustrate all features and options, the examples here assume *System Admin* authority. Refer to Configuring Users for more information.

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-left of the screen (File, Configure, etc) are 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.

At the top-right of the screen is the status bar. This contains indicators that will show green, yellow or red, indicating if there is backup job in the queue that is either running, pending (waiting) or failed, respectively. It also contains a button to view the log of backup status messages not already displayed.

The remainder of the screen will vary depending on the "Display" tab chosen:

- **The Clients, Servers & Devices** display tab (shown below) is available only if the *Network Edition* or *TSM Edition* 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*=*Solaris***) and whether or not the system is available (**Green**=available, **Red**=not available). Devices (tape and directory) are also shown. Tape devices will appear red if the device is unavailable on the server.

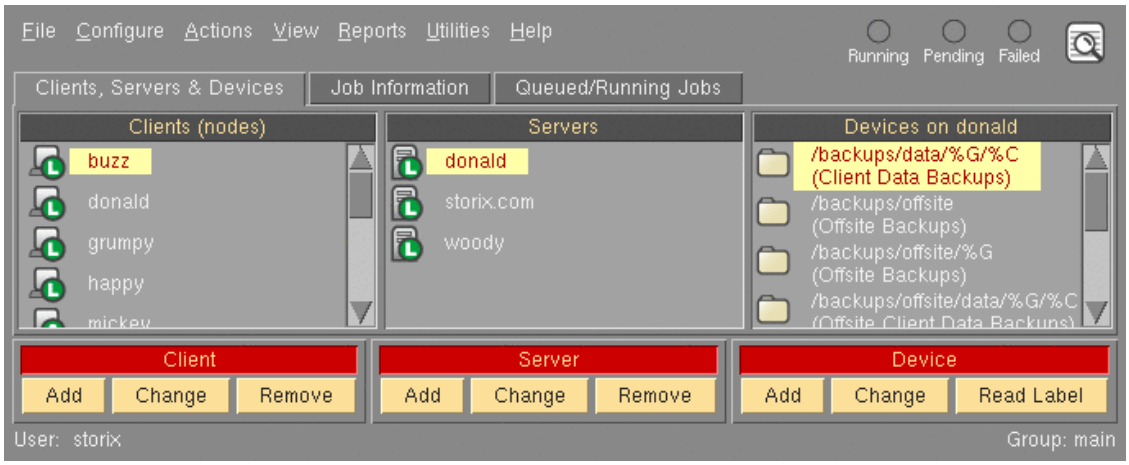

A client may be selected by clicking the left mouse button on the icon next to the client hostname. Likewise, a server may be selected by clicking the mouse button on the server icon. When you click on a server, a list of backup devices configured for that server will appear. The *selected* client, server,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

The action buttons at the bottom of the screen apply to the selected client, server or device. They provide a shortcut to performing the same tasks that can be performed from various options within the menu bar. The **Add** or **Change** buttons refer to the item selected above (client, server or device). If you want to display the backup labels for all backups stored on the server, you must select a server and a device, then press the "**Read Label**" button.

- **The Job Information** display tab provides a quick glance at the jobs that are configured. The left-most display area shows the job icons. The green clock over the calendar icon represents a job is that is scheduled. The blue pause over the calendar icon 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 highlighted.

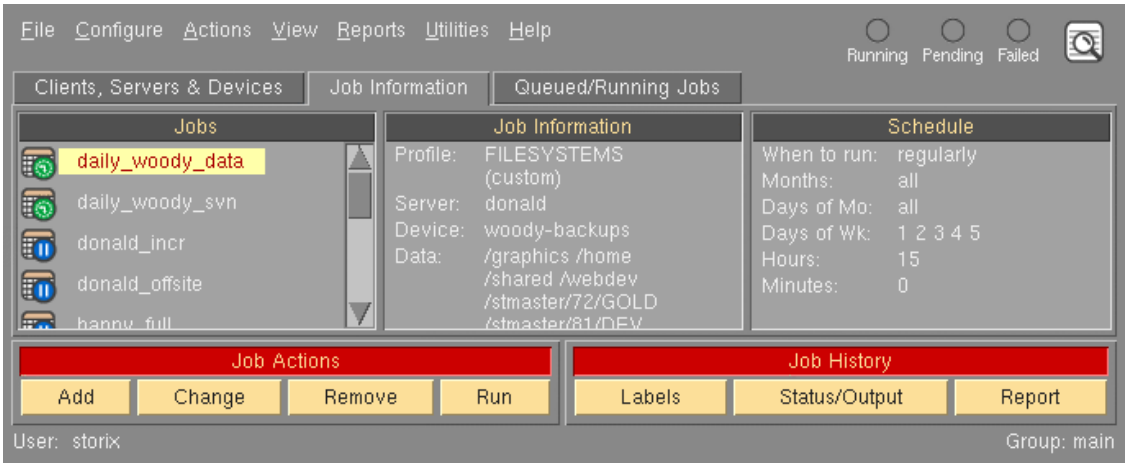

The action buttons at the bottom apply to the selected job. They are shortcuts for various job-related functions. The Run button will place the selected job in the queue (even if it is scheduled to run at another time), and it will be run as soon as the server and device assigned to the job are available. 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 Queued/Running Jobs** display tab 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, and the jobs in the selected queue are displayed in the middle display area.

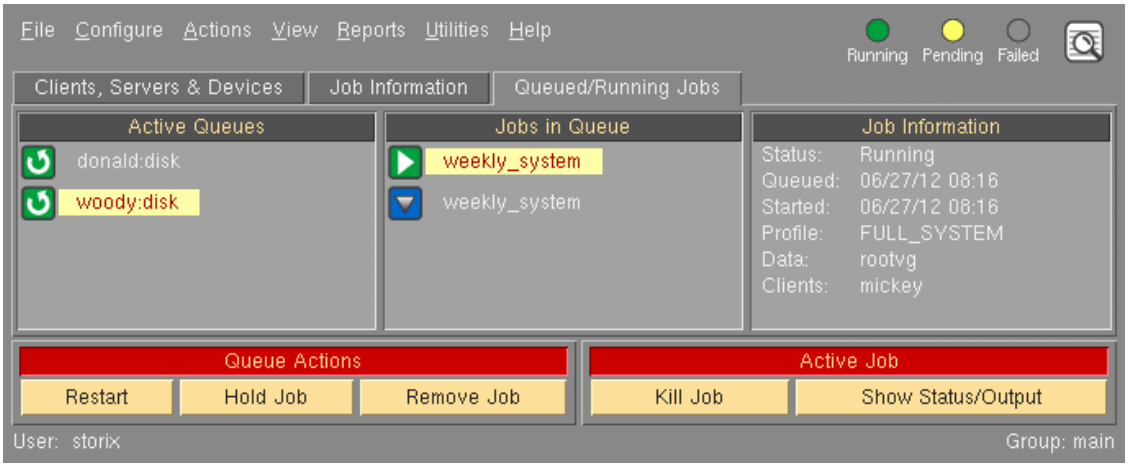

You may then click on a particular job to display the job information, including the status of the job. Both the queue and job icons represent the status of the job. The ***Queue Actions*** buttons at the bottom of the screen may be used to manipulate the selected job. The ***Active Job*** buttons include 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 Windows

A common icon which appears at the bottom of each window is:

Cancel button

After making changes to information on any screen, use the cancel button to cancel the changes and close the window. Avoid using the window-manager button (usually at the top-right of the window) to close windows as this does not always perform the entire cleanup needed. The cancel button does 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 minimizing the window.

4. Users

When you first installed SBAdmin, an "**admin**" user was created and you were prompted to provide a password for this user. The admin user is given authority to all (*System Admin*) functions within the SBAdmin application, including configuring other users.

If there is only one user configured, and you are logged onto the system as "root", you are logged into the application under this user by default. You will not need to provide a username or password at the command-line.

However, for the web-based application, you must always provide a username and password.

You may configure one or more users, each with permission to perform specific tasks. Each user will be assigned to a default group, but a user may be allowed access to multiple *groups* (see Configuring Groups).

User Levels (Roles)

By configuring users you will be able to limit permissions and roles within the software. Configuring users at different levels is useful if multiple people are accessing the administrator, and security policy dictates what access each person should have. The following are the four types of SBAdmin users that may be configured:

- Backup User **-** limited to monitoring backups and running backup jobs that are already configured by a privileged user.
- Backup Admin **-** allowed to configure backup settings and backup jobs. They are also able to monitor backups and run backup jobs.
- Group Admin - allowed all *Backup Admin* functions as well as configure application settings, clients and servers within their *group*.
- System Admin **-** allowed all access and may configure all backup functions as well as application settings, clients, servers and groups.

Adding a User

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

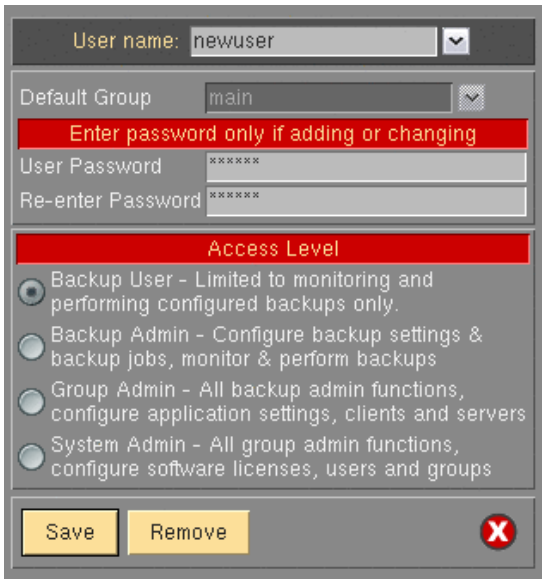

From this screen, type the user name in the **User name:** entry field. When adding a new user you must also specify the **Default Group, User Password**, and **Select the access level for this user** section. Press the **Save** button to add the newly configured user.

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

- **User name** – This field indicates the username within the SBAdmin program. This does not need to be a user defined on the Unix/Linux system.
- **Default Group** – This is the group that the user will be logged into when launching the SBAdmin interface. To allow this user to access other groups, refer to Groups section.
- **User Password** – These fields are for specifying or changing the user password. User passwords are encrypted and do not need to correspond with system passwords.
- **Select the access level for this user** – This selection will determine what functions the user may perform. For more information, see the User Levels (Roles) section.

Removing a User

Click Configure→Users from the menu bar. Select the name of the user to remove from the list and press the **Remove** button. Any configurations this user may have made will remain intact; however the user will no longer be able to log into the Administrator.

> **You must leave at least one user configured with System Admin access.**

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

Changing a User

Click Configure→Users from the menu bar. Select the name of the user to change from the list. Here you may enter a new password into the **User Password:** and **Re-enter Password** fields and or **Select the access level for this user**. To change the user press the **Save** button.

Users with access lower than Group Admin will only have the ability to change their password. All other fields will be disabled.

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

Changing your User Information

Only a user with System Admin access can add, change or remove other users. Other users can change their own information by selecting Configure→User Information or File→User Information (for users with only Backup User access).

You will see the same screen shown above, but will only be able to change your Default Group and your Password. Simply select a new default group (if more than one available), or enter a new password in the fields provided and press the **Save** button.

5. Groups

To configure Groups you must be logged into SBAdmin as a User with “System Admin” privilege.

When SBAdmin was first installed the group “main” was created. The “main” group will function as the default group and no further group configuration is required. You may choose to configure groups to assist with the organization or security of your backup environment.

Groups are used to allow a single Administrator to organize and manage Clients and Servers; and may be configured to allow or restrict certain User access. Groups are also necessary when configuring servers that will share or limit client access based on group ID. For examples, see the Using Groups section of this guide.

Adding a Group

To add a group, select Configure->Groups from the menu bar. A similar looking screen will appear:

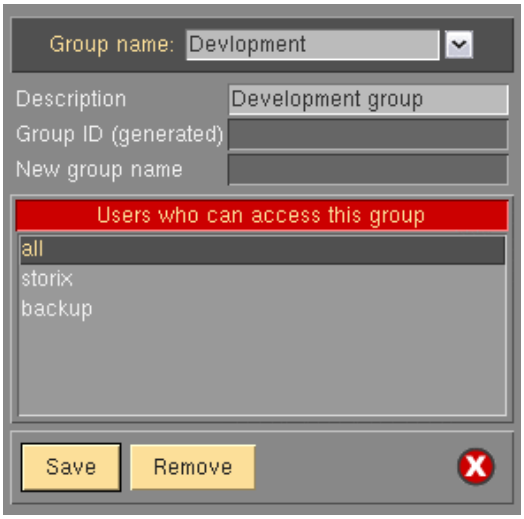

From this screen, simply type the name of the group to add in the **Group name:** entry field and optionally a **Description** of the group. The **Group ID** field will be automatically populated with a unique value to be associated with the group. The **New group name** field is only used when changing a group. Select any users who should have access to this group from the **Users who can access this group** box. Then press the **Save** button.

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

- **Group name: -** This field defines the name of the group to be added.
- **Description –** This is an optional field and is used to describe the group.
- **Group ID –** This field will be automatically populated when adding a group. It is a unique identifier and will not change if the group name is later changed.
- **New group name –** This field is only used when changing a group.
- **Users who can access this group –** This field allows you to specify one or more users that have permission to launch the Administrator under this group. For more information on users and user roles please see Configuring Users.

Changing a Group

Click Configure→Groups. Select the name of the group you wish to change from the **Group name:** drop-down arrow to the right. You may now edit the **Description** and **New group name** fields or select/deselect users from the **Users who can access this group** box. Once you have made the appropriate changes press the **Save** button to update the group.

> **The "Group ID" field will never change. Once a group has been configured this value will always be used to identify the group.**

Removing a Group

Click Configure→Groups from the menu bar. Select the name of the group to remove from the list and press the **Remove** button. When finished, press the **Cancel** button at the bottom.

> **A group may only be removed after all clients and servers have been removed from the group.**

Switching Groups

Select File→Switch Group from the menu bar and select the name of the group you would like to switch to. Then press the Switch button. This will cause the Administrator to close and re-open under the new group. Clients, servers, media, jobs and queues will all update to reflect the settings in the new group.

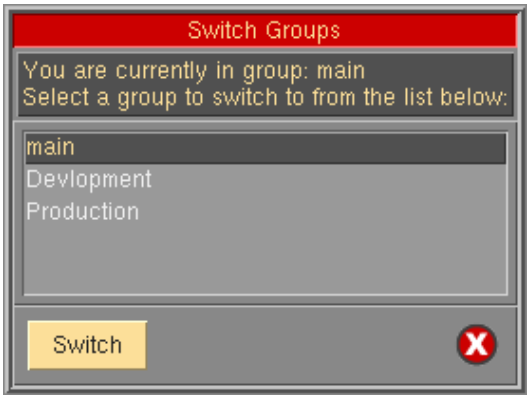

> **Only one instance of the SBAdmin graphical interface may run per group. You may launch multiple instances of the interface with different groups using the –G command. The following command illustrates how to launch the interface for the group "main".**
>
> **# sbadmin -G main**

Using Groups

The following diagram illustrates how two groups may be configured on a single Administrator to isolate clients and servers from two separate environments. For security reasons clients from the *Development* group will not have access to backups performed by clients of the *Production* group and vice-versa.

> **Clients may ONLY belong to one group but a server can be shared among multiple groups. In this example *SharedServer* is a server configured under group "Production" but allows access to "Development" group clients.**

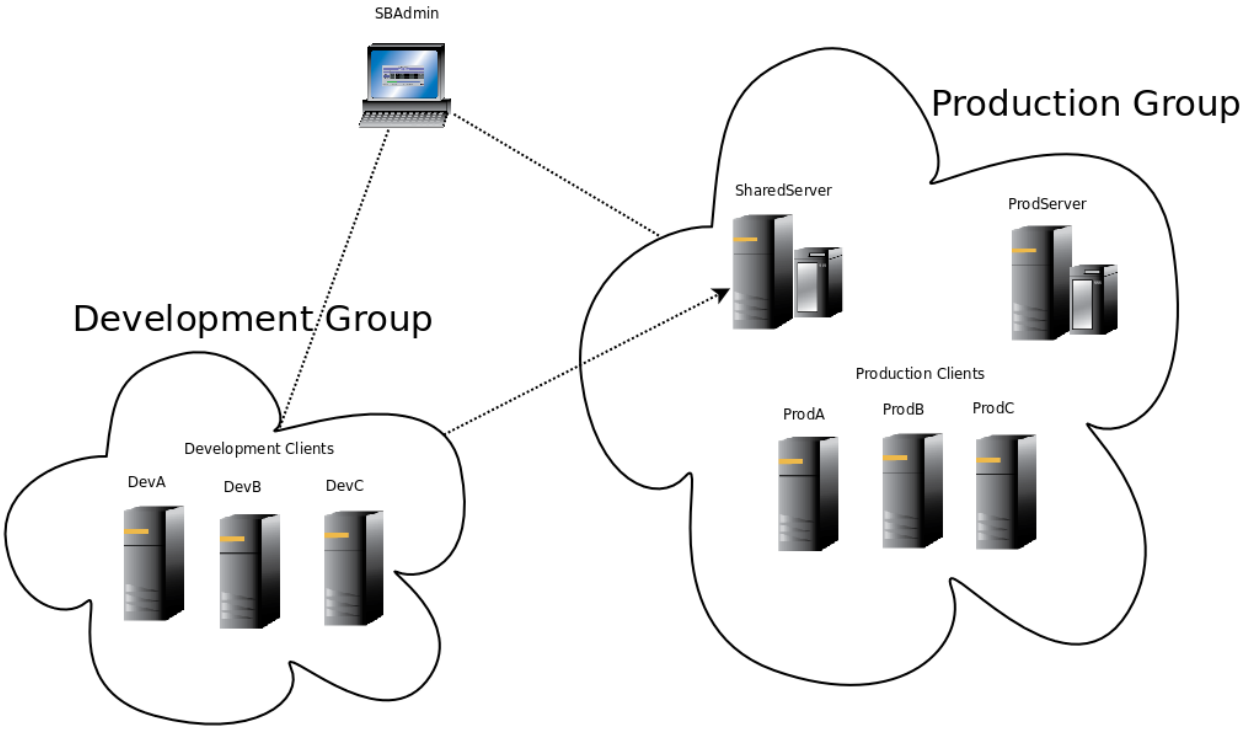

- **SharedServer** is a server configured by the Production group to allow access from the Development group. This means clients from either group may backup to it. However, only the Production group may make changes to it.
- **ProdServer** is a server configured to only allow clients from the Production Group access.

With the above configuration, the SBAdmin will configure two groups: **Development** and **Production**. After switching to group **Production,** the clients and servers may be added and configured. See Configure Clients and Configure Servers for more details. The following screen shows the Administrator interface configured under group **Production**:

In this example, the *SharedServer* has a directory device configured containing both %G and %C notation. This means only the original client and group will have access to the backup. See Configuring Servers for further detail.

After the first group is configured, the SBAdmin interface can either be switched to group **Development** or a second instance of the interface can be started (by typing "*`sbadmin -G Development`*"), allowing simultaneous access to both groups. You can also type "*`sbadmin`*" and enter a username and password when prompted to start the application using the default group for that user.

The following is an example of the interface running under group **Development**. Notice the color scheme has been changed for this group, to help differentiate between the groups.

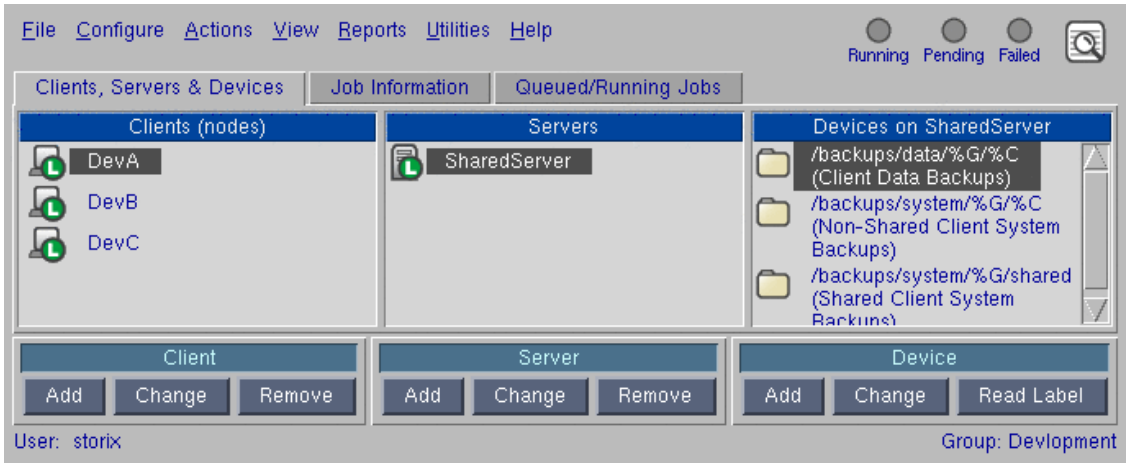

Also note that the ***SharedServer*** has been added to this group even though it was initially configured and owned by the Production group. This is only possible if the ***SharedServer***, when configured in the **Production** group, specifically allowed access to the **Development** group. See Configuring Servers for further detail.

6. Clients

This section only applies to the *Network Edition* 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, they should also be configured as clients.

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 and/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.

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

- Configure→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, Solaris or AIX client.

Configuring a Linux, Solaris or AIX Client

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

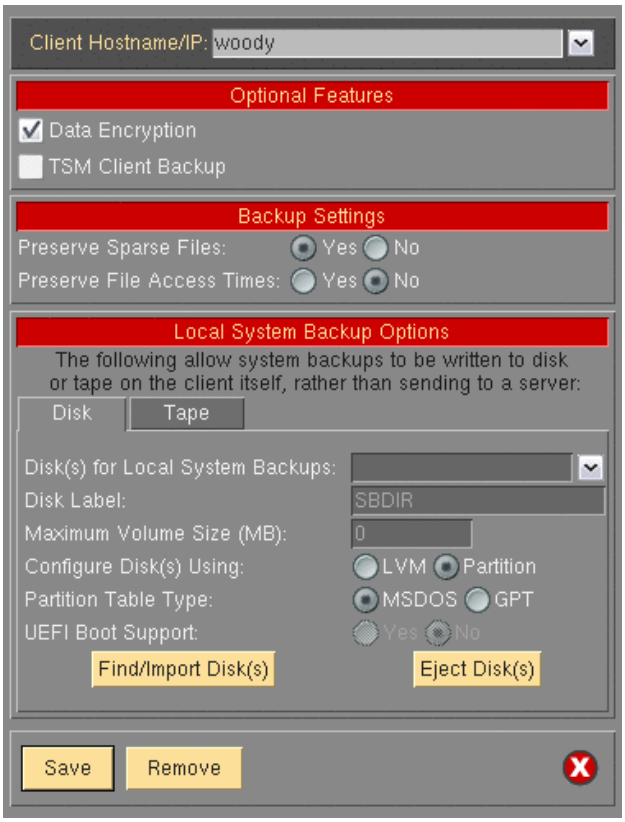

To add a new client, enter the hostname of the client in the entry field at the top. Note that the hostname you enter may be a simple hostname (i.e. *LinuxServer*), a full domain name (*LinuxServer.storix.com*) or IP Address as known to the admin system. Then select any optional features that apply to this client. The optional features will only be selectable if you have installed a software license for that feature. If so, you may select the option for only the number of clients the feature is licensed for. When you've made your selections, select the **Save** button.

When adding the client you may also want to configure a disk, NFS mount, or tape **Local System Backup Device**. See Configure a Local System backup to create a backup job directly to this client's locally attached disk, NFS mount, or tape drive.

To change an existing client, you may type the client name in the **Client Hostname/IP** field, or use the arrow button to the right of the field to select from a list of configured clients. You may then select or deselect optional features, or remove the client.

Configure a TSM Client (node)

If the client to be added is a TSM node and you want to back this client up to a TSM server, select the **TSM Client Backup** button. Keep in mind, however, that the number of TSM clients which may be added is dependent on the number of TSM clients the administrator is licensed for. When selecting this option, additional options will appear on the screen as shown below:

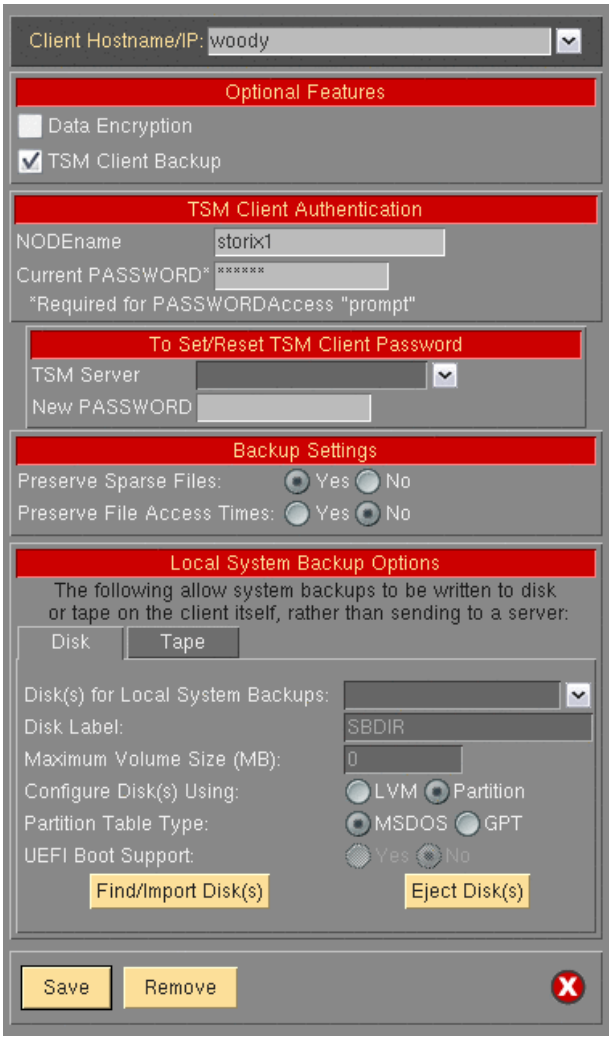

For most options, you can simply use the defaults shown. Other options are TSM-specific and described in detail in the TSM documentation. However, you can also use the right mouse button over any option to display the **QuickHelp** containing more information

- **TSM Node Name and Password**

  You must enter the nodename of the client in the **NODEname** field. This is the name of the client system as registered with the TSM server. The password of the client will be require if the *PasswordAccess* option of the TSM server is set to "*prompt*", since the password must be provided with each command executed between the client and server. This password will be stored on the client, in a protected file and in non-textual form, for use by SBAdmin commands.

- **Set or Reset the Node's Password**

  This screen can also be used to set or reset the password of the node on the TSM server by selecting the *TSM server* in the drop-down list, and entering a new password in the **New PASSWORD** field. In this case, you must also enter the **Current PASSWORD**, regardless of the *PasswordAccess* option of the server.

Enabling Backup Data Encryption for a Client

The **Data Encryption** option 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. 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 Reference Guide***, and the Encrypt data field in the backup job configuration for additional information.

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 and the default. 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 resolve this issue, you may select the **Preserve Sparse Files** "*No*" option so that all backups of this client will be created 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.

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.

Preserve File Access Times

When a file is backed up it is opened for reading and the access time in the inode table is updated. While this does not affect the modification time, it does change the access time. To prevent this behavior, you may select the **Preserve File Access Time** "*Yes*" option to retain the original access time during a backup.

This option is only valid for Linux and Solaris systems and is not available for AIX.

Local System Backup Options

The options in this section may be used to configure devices on the client that may be used as backup media for system backups of the client. These devices would be used in lieu of writing to a device on a backup server. Only *system backups* may be written to these devices. Refer to Creating a Local System Backup in the job configuration for more details. Supported local system backup device types are disk and tape, and the configuration of each type may be viewed by selecting the appropriate tab.

You may configure a system backup device of each type. However, each type may only have one configuration.

Disk(s) for Local System Backups

The options in this tab may be used to configure one or more disks directly attached to the client for full System Backups. You can then perform a system backup of the client to its own disk(s), boot and re-install from those disks with no need of a network server. The disk(s) may also be moved to any other client and used to install, or clone, other systems.

Complete details on these options are shown in the Local System Backup Disks section.

Note that this option will create a special directory-based device with the name "**SBDIR**". This device will then be available when you configure a backup job and indicate the backup is to a local disk device. Refer to Creating a Local System Backup in the job configuration for more details.

Tape for Local System Backups

A client may perform a System Backup to its own direct-attached tape drive. This allows any client with a tape drive to back up to itself without the use of a server or any network traffic. This tape may then be used to reinstall this client. You can move the tape or tape drive to a server to make the backup available to any client, or you can move the tape or tape drive to any other client to allow then to perform a local system recovery.

Since the tape drive names may differ for each client, you will select in this field the name of the tape drive to configure. Only tape drives that are configured and available on the client will be shown. Although multiple tape drives may exist, you may only select one drive per client for system backups.

Note that this option will create a special tape device with the name "**SBTAPE**". This device will then be available when you configure a backup job and indicate the backup is to a local tape device. Refer to Creating a Local System Backup in the job configuration for more details.

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 prompted to remove the client from all configured backup jobs as part of removing the client. If you choose not to remove the client from all jobs then 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 and press the **Remove** button.

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

7. Servers

A backup media server, also referred to simply as the ***backup server, SBAdmin server***, or just ***server***, is defined as any system to which backups will be sent to and stored. This type of server will have SBAdmin software installed, and all network data transfer between a client and the server will be performed using SBAdmin network communication programs.

If you wish to use a ***locally mounted NFS share*** on a backup client to store backups, then you may do so by configuring a Shared NFS Server, as described in the Adding a Shared NFS Server section below.

If the ***TSM Client Backup Feature*** is installed, then a *TSM server* may also be used for all backups. To configure a TSM server, skip to Adding a TSM Server below for details.

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.

> **The following section is used only when the *Network Edition* license is installed. Refer to the *Backup Devices* chapter below for *Workstation Edition* licenses.**

Adding a Backup Media Server

A new server may be added to the system by:

1. Selecting Configure→Servers→Backup Media Servers from the menu bar.
2. When *Network Edition* is installed, press the **Add Server** button at the bottom of the Main Screen when Clients, Servers & Devices are displayed.

A screen similar to the following will be displayed:

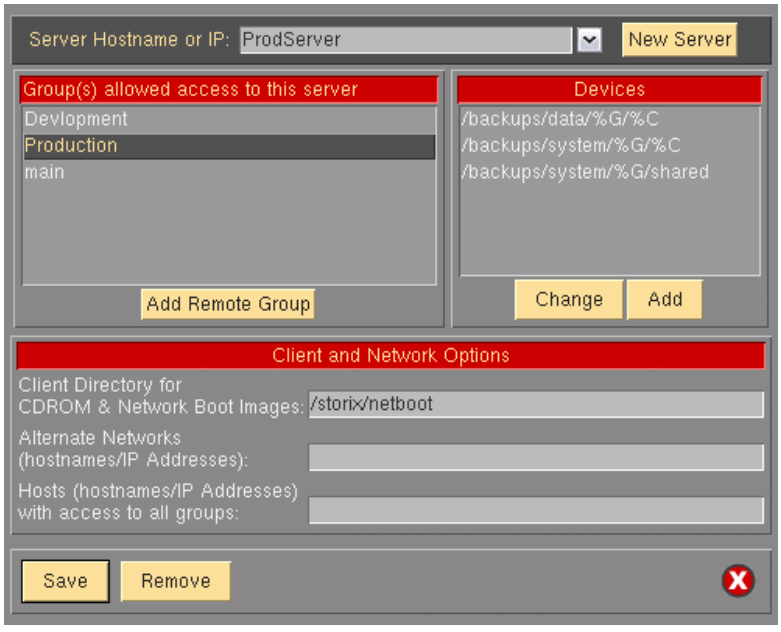

In the **Server Hostname or IP** field you will need to enter the hostname or IP address of the server you would like to configure, and then press the **New Server** button. If this server is not already defined by

another group you may continue to Group(s) allowed to access this server. If this server is already defined, continue with the next paragraph.

Adding a Server Configured in Another Group

When adding a server that is defined in another group, you may be presented with the following message.

> *"This server was previously configured from another group. Although you have been granted access to define or un-define this server to your group, you will not be able to change the settings."*

You will be able to define this server and run backup jobs to it. If your Admin System has access to the server you will need to switch groups in order to make changes to the server settings. If your Admin System does not have access to this server you will need to contact the Administrator of the server to make any changes to the server.

When adding a server that is defined in another group you may get the following message, where ***GroupID*** is replaced with your actual group id.

> *"This system is currently defined as a server, but your group id **GroupID** is not allowed access. Permission for your group must be granted from the controlling Admin System."*

In this case you will not be able to add this server. You must run the *Admin System* under the group which owns (originally configured) the server, and then grant access to this new group. See Switch groups for changing groups using the Administrator and Sharing a Server Between Administrator Environments.

Group(s) allowed access to this server

In the first listbox, you must select one or more groups that will be allowed access to this server. Only clients in the selected groups will have permission to write backups to the server, and only clients within these groups may be selected when configuring backup jobs for this server.

Note that the groups listed here by name are groups defined on this Administrator system. If Group IDs appear in the list, they are currently defined Remote Groups allowed access, but defined to other Administrator systems.

Adding Access to a Remote Group

Any group defined by name to this Administrator interface is considered a local group. Since a server defined to this admin system (and possibly one or more local groups) can also be made available to other SBAdmin systems, you must define which groups on that admin system will also have access. Groups defined on a SBAdmin system other than this one are considered *remote groups*.

Since this SBAdmin system does not know the names of the remote groups (and the names may conflict with their own), the remote groups must be defined using their *group IDs*. A group ID is a 16-digit Hex number and can be located using the Configure→Groups option from the remote SBAdmin interface.

On the Administrator that currently has the server defined, select the **<Add Remote Group ID>** option from the **Groups allowed access to this server** field. The following window will be displayed.

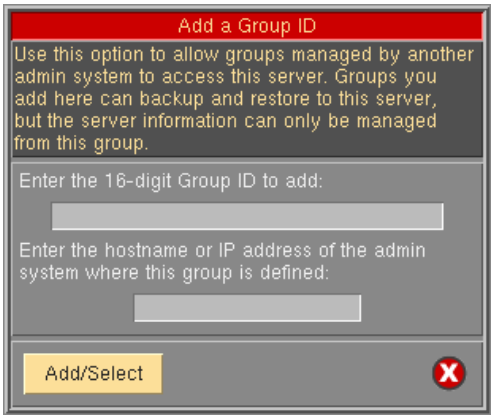

Here you will need to enter the **Group ID** as defined on the remote Administrator that will be accessing the server (see Configuring Group to determine Group ID), and the **hostname or IP** of the remote Administrator.

Now that the server has been shared on the first Administrator, the remote Administrator will have to add the shared server. To do so, refer to Adding a Server. Once the shared server is configured on the remote Administrator, jobs may be configured to write to the shared server (see Creating a Backup Job for more information).

Configuring Server Devices

The second listbox shows all of the *devices* (tape or directory) configured on the server. You may add, change or remove devices by selecting a device and the appropriate button below. This is simply a shortcut to the same configuration screen which appears when selecting Configure→Devices from the menu bar

Client Directory for CDROM and Network Boot Images

When a bootable CDROM image or a network boot image is created for a client, allowing it to be booted from CDROM or over the network, the images must be created on the client, but may stored on the boot server. You can later select images from this directory to copy to a CDROM drive or configure for network boot of a client.

Alternate Networks

This field allows you to enter the optional *IP Addresses* or *Hostnames* pointing to alternate network adapters on the server. These "*alternate networks*" may be used by clients to backup, restore, or boot from the server when the primary network (that used between the admin system and server) is not available to the client, or to simply move the network backup data traffic off of the primary network.

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. To add alternate networks for the server, type the IP addresses or hostnames, separated by spaces.

After adding entries in this field, you will be able to select that network for the server whenever creating a backup job, restoring data, booting or re-installing a client from this server.

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* (see the *SBAdmin System Recovery Guide*).

Hosts with Access to All Groups

When using a Shared Server, it is often useful to define a client with access to all groups. If a remote Administrator copies a client's backup to a shared server, the original client will not have access to the backup on the Shared server (since he is not part of that group). A client with access to all groups will be the only way to restore from these backups the Shared Server.

> **It is not necessary to have this hostname or IP defined as a client on any Administrator system, and entering a client here will not configure the client on the Administrator.**

By entering one or more hostnames or IP addresses in this field, that system will be granted access to the backups of all groups on the server. *This should be used cautiously*! The host specified here does not need to be a client or server configured on this Admin System, but will be able to access backups originally performed by any group on the server.

Note that, if you use a hostname, the hostname will be looked up (using TCP/IP name resolution) and replaced with an IP address when the server information is saved.

Adding a Shared NFS Server

If you would like to store backup data on a NFS share mounted locally on the client, then you will configure a **Shared NFS Server**. Any or all of the clients may then use this server configuration for backup storage. You may also configure a physical NFS server multiple times, which may be necessary when clients require different mount options to mount the share.

The shared NFS server provides the ability to perform a backup of a client, and store the backup on a NFS share that is mounted on that client. Thereby minimizing network traffic to be between the client and the nfs server. It will be necessary for the client to have r*ead-write mount capabilities of the share* prior to running a backup. If you wish to centrally mount the NFS share on a single system and have clients write to share in that manner, then you should configure a backup media server.

> **It is not expected that SBAdmin software is installed on the physical NFS server, and the configuration of the physical NFS server (i.e. security, exports) is the sole responsibility of the user.**

A new server may be added by selecting Configure→Servers→NFS Servers from the menu bar. After doing so, a screen similar to the following will be displayed:

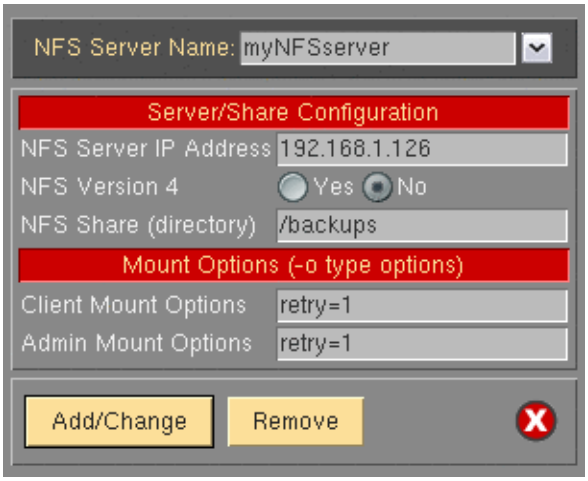

NFS Server Name

The NFS Server Name field is a unique name that you will use to identify this server's configuration. You may configure any number of servers, and may specify the same physical NFS server multiple times by using a different name.

NFS Server IP Address

In the NFS Server IP Address field, enter the IP address of the NFS server that is exporting the share.

NFS Version 4

The NFS Version 4 option allows you to indicate whether the NFS share that is to be mounted is a NFS v4 share. This is necessary because depending on the operating system type, special handling may be required for NFS v4.

NFS Share (directory)

In the NFS Share (directory) field, enter the directory path of the share on the NFS server to be mounted on the client. You may use *%C notation* (where %C will be replaced with the name of the backup client) but the %C must be the last directory in the NFS share path (i.e. /backups/nfs/%C). Also, when using %C notation, the admin must be able to mount read-write the parent directory of the %C directory (i.e. /backups/nfs), while the backup client itself only needs to be able to mount read-write the full path.

NFS Client Mount Options

In the Client Mount Options field, enter any mount options *(-o type)* to the *mount* command required to mount the NFS share on a client. The options should be comma separated.

NFS Admin Mount Options

In the Admin Mount Options field, enter any mount options *(-o type)* to the *mount* command required to mount the NFS share on the admin. The options should be comma separated. The admin needs to be able to mount the share to perform administrative operations like applying retention policies and %C notation.

When done entering required information into the fields, click the **Add/Change** button to save the configuration.

Adding a TSM Server

Any TSM server, and any number of TSM servers may be configured when the ***TSM Client Backup Feature*** is installed. This is not limited by the number of client licenses installed. In addition, a single physical TSM server may be accessed using different TSM server configurations. For example, one server (*tsmserver-comp*) may be configured to backup data using TSM compression, while the same server (*tsmserver-nocomp*) may be configured to backup without using data compression.

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

1. Selecting Configure→Servers→TSM Servers from the menu bar, or
2. When *TSM Edition* is installed, press the **Add Server** button at the bottom of the Main Screen when Clients & Servers are displayed.

After doing so, a screen similar to the following will be displayed:

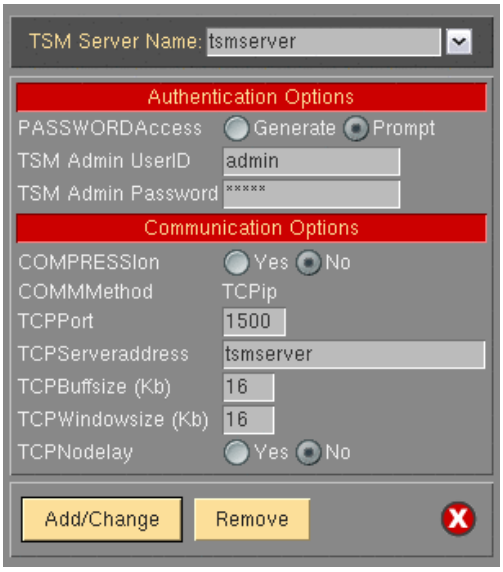

You should refer to your TSM documentation for information on the additional fields which are normally specified in your TSM system user options (***dsm.sys***) file. You may use the right-click button over any field to show the *QuickHelp* information on each field, so information is not detailed here. However, there are some special considerations:

TSM Server Name

You can call the server anything you like, since it is the entry in the *TCPServeraddress* field, not the server name, which determines how the server will be contacted. You can also configure multiple servers, each with a different name, that use the same *TCPServeraddress* entry.

TSM Admin User ID/Password

You must enter the name and password of a TSM administrative user already configured on the *TSM server*. This administrative user must have been configured with either *System or Policy authority*. The administrative user ID entered will be used only by this Administrator application to perform backup management tasks on behalf of the clients. This information is never sent to or saved on the TSM client systems.

PASSWORDAccess

If this field is set to "**generate**", it is assumed that you have already set the password on the client using another TSM application (or have run another type of TSM backup from this client), and the server used had this option set to "generate". In doing so, an encrypted password file was created on the *TSM client*, and will be used by this application also. In this case, you do not need to enter a **Current PASSWORD** when configuring the TSM clients.

If this field is set to "**prompt**", the client password must be provided each time contact is made with the TSM server. The password is stored on this TSM *Admin System* for future use, and is also sent to the client and stored in an encrypted and protected file for use by SBAdmin commands. Normally, the password is contained in the client user options file (***dsm.opt***). However, SBAdmin does not use this file but supplies the password to the TSM server with each command. When using this option, you will need to enter a **Current PASSWORD** for each TSM client you configure.

COMMMethod

This will be displayed as “TCPIP”, which is the only option supported. The only other TSM option is “shared” (shared memory), which can only be used when backing up the server itself. The shared memory option provides no significant performance increase and is not supported by SBAdmin.

COMPRESSion

A selection of “yes” simply indicates that TSM will compress backup data on the client before sending to the server. It is important to note, however, that SBAdmin provides its own compression options (see Configuring a Backup Profile). If you choose to use SBAdmin compression, then no TSM compression will be used *regardless of your selection here*. If you want to use TSM compression, you should select “yes” in this field, and do *not* indicate to use compression within your backup profile (the default). Different compression schemes work best on different types of data and sizes of files. You should therefore experiment with using both TSM and SBAdmin compression options to determine which provides the best compression with the least impact on the client system performance.

To add a new server, enter a name in the entry field at the top, then add or change any of the fields on the screen, then click the **Save** button.

Changing a Server

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

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

If selected from the main screen, the server options screen will appear with the prior settings for the server. If not, select the server by typing the name (or IP address) at the top, or use the arrow button to the left of the entry field to select from a list of configured servers. 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.

Only the admin system and group that originally *created* a server may remove it. Doing so makes it inaccessible to any client or group with prior access.

If you *defined* the server (previously created by another group) to your group, then removing it will only remove access to your group, and will not affect access by other groups.

To remove a server, either:

Select Configure→Servers→Backup Media Servers from the menu bar, then select the server by typing the name (or IP address) at the top, or use the arrow button to the left of the entry field to select from a list of configured servers. To remove the server, then press the **Remove** button at the bottom of the screen.

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 main screen.

8. Backup Devices

Before any backups may be performed, you must define the backup devices to be used on the local system (*Workstation Edition*) or on each server (*Network Edition*). This includes both tape devices (standalone drives, libraries or autoloaders) as well as backup directories on disk.

A backup "device" can be named anything you wish, but typically a single tape drive would be simply called by its physical device name (i.e. "st0" or "rmt0") and a single directory device would be called the directory itself (i.e. /backups). You may also wish to name a device "Tape1" or "SharedDirectory", etc. The name that you choose will be how the device will appear in all of the selections, such as when selecting the device to use when configuring a backup job.

Add a Backup Device

To add a new device, select one of the following:

1. Configure→Backup Devices from the menu bar.
2. When **Clients, Servers and Devices** are shown in the Main Screen (Network Edition), select a server, then press the **Add** button directly below the **Devices** pane.
3. Configure→Servers (Network Edition), select a device from the **Device** listbox, then press the **Add** button below.

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

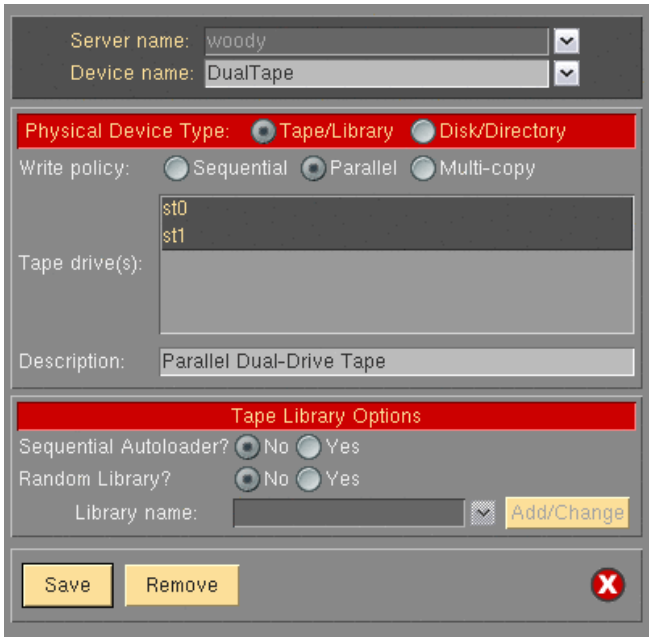

From this window, you must first select the server (Network Edition). Then, to add a new device, type a device name in the **Device Name** field. To change or remove an existing device, select it using the arrow button to the right.

> **Be sure to type a name that will easily identify the device. When using a single tape drive or directory, the device name (i.e. "st0" or "rmt0") or directory name itself (i.e. "/backups/shared") will usually suffice.**

By default, the **Physical Device Type** field will be set to "*Tape/Library*". If you want to create a directory-based device for writing backups to disk on the server, select "*Disk/Directory"* instead. The remaining options will differ depending on the device type you are configuring:

Tape Devices

The physical tape drives available on the server will be shown in the listbox. Most often a single tape drive is used independently. In this case, just select a single tape drive from the list.

Tape Write Policy

When using a single tape drive that is not a part of an autoloader or tape library, only the *Sequential* write policy is used. The write policies are defined as follows:

- Sequential - 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 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 sequential device. However, there would only be an advantage if you are using a sequential autoloader or random library. If so, the 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 - 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 drives, allowing the backup to complete in a fraction of the time it would take to write to a single drive. The same parallel device (or one containing the same number and types of drives) must be used to restore data from the backup.

- Multi-copy - This device must consist of two or more tape drives. When backups 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.

When reading data from a backup made with a multi-drive device, it is necessary that the tapes be placed in the corresponding tape drives in the same fashion as they were during the backup. This does not apply to a *multi-copy* 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* device (you can still use a multi-copy device to perform the restore, but only the first drive is read).

The following illustrations show how the data is saved on each of the different multi-drive devices. Note how the tape volumes are numbered for each write policy when data spans a different physical tape.

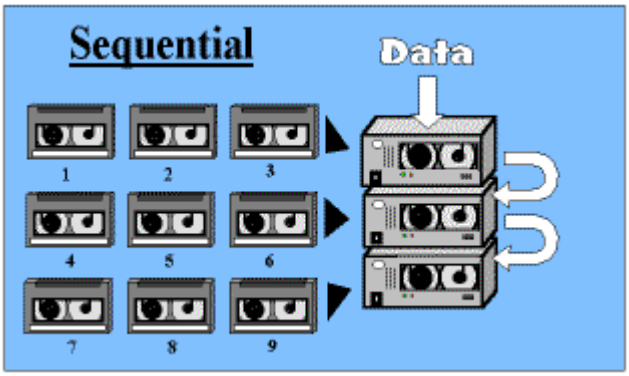

The sequential data is written to the first tape in the first drive until it fills up. Then, the backup continues onto the next drive in the list, etc. Only when all tapes in all drives are filled will the user be prompted to change volumes. Then, the backup continues again with the first drive and so 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 sequential device or from a single-drive device.

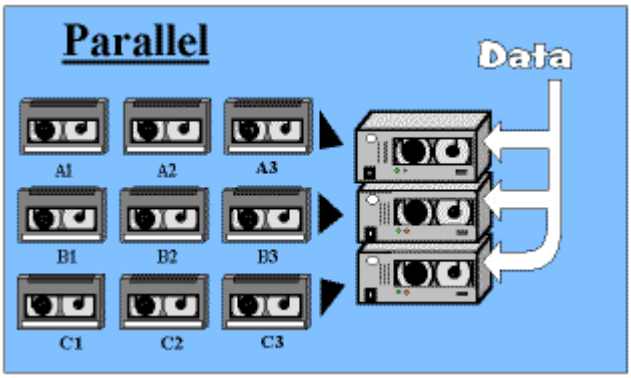

The data is split into multiple buffers (one for each drive) of the same size, and the data is sent to all three drives at the same time. The user will be prompted to change the volumes in each drive as it fills up, which may not necessarily be at the same time as other drives, particularly when using 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 device backup may only be verified or restored using a parallel device with the same number and types of drives.

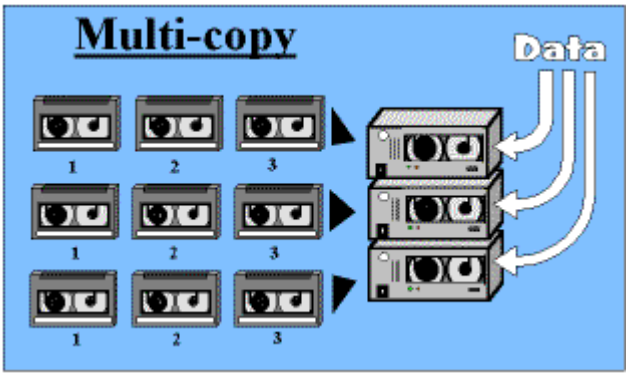

An identical copy of the same data is written to each of the drives at the same time, normally at about the same time it would take to write to a single drive. Since the same data is written to each drive, the volumes are numbered as though each is a single-drive backup. Since a multi-copy device backup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

Sequential Autoloader

Select the "**Yes**" button if the selected tape drives are physically contained in a sequential autoloader or "*tape changer*". This may also be a *tape library* set to sequential mode. When you backup or restore using a 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 begins.

Random Library

Select the "**Yes**" to indicate the tape drives are physically 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 SBAdmin will issue the commands needed to return the tape to its original slot in the library, then grab and insert the next

tape in the magazine. 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, SBAdmin 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. Some pre-defined library definitions are available when you install the software. If your library is not in the list, you must add a new definition 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.

Directory (Disk) Devices

When you set the **Physical Device Type** to “*Directory/Disk*”, the device will be used to write the backups to disk on the server (or local disk if using *Workstation Edition*). You may then type one or more directory names where the backups will be written.

If using multiple directories, be sure each is in a different filesystem. Otherwise, when the first filesystem fills, the backup will attempt to write the next directory, which will be in the same full filesystem, and the backup will abort.

Directory Write Policy

The only *write policy* available for directory devices is *Sequential*. Backup data will be stored in files (generally one for each filesystem or partition being backed up) in the directory or directories specified.

If a single directory is specified, all backups will be written to that directory. If the filesystem where that directory exists fills up, the backup will abort.

If using multiple directories, be sure each is in a different filesystem. Otherwise, when the first filesystem fills, the backup will attempt to write the next directory, which will be in the same full filesystem, and the backup will abort.

If you specify more than one directory (each separated by a space), the backups will be written sequentially to the directories starting with the first directory in the list. If the filesystem containing the first directory fills, the backup will automatically continue by writing backup files to the next directory in the list, etc. If the filesystem containing the last directory fills, the backup will abort.

Separating Backups by Group and Client

This option does not apply to *Workstation Edition*.

There are some special notations that may be used when you specify the directories used within a directory-based device. These are:

%G This will be substituted by the *Group ID*
%C This will be substituted by the *Client Name*

Although you can use **%G** and **%C** within the name of the device itself, they only have special meaning when used to specify the directory the backups should be written to.

Although you may store client backups, and those from different groups in the same directory, you may still limit read permission to specific backups to the original client (owner) of the backup. This is accomplished using the Disk Backup Read Permissions in the Backup Profile settings.

The **%G** will be replaced by the *Group ID* of the group, thereby storing backups for different *groups* in different directories. It also prevents a client from one group from accessing backups of a different group. If you want all backups from different groups stored in the same directory, remove the %G notation.

Likewise, the **%C** notation will be replaced by the *client name*. Again, this will store each client's backups in its own directory, and will prevent a client from accessing another client's backups.

If you want backups for different groups to be accessible only by clients and servers within the same group, you can, for example, create a device containing the directory "**/backups/%G**". Using this device would create and place backups in a directory such as **/backups/19a955f3d19d49f8** (where *19a955f3d19d49f8* is the *group ID*).

If you want backups to be accessible only by the original client, you should either add the **%C** notation to the directory name, or create a new device for a directory containing the **%C** notation. Using the previous example, specifying "/backups/%G/%C" would create and place backups in a directory such as **/backups/19a955f3d19d49f8/mickey**", assuming the group is *19a955f3d19d49f8* and the client name is "*mickey*". In this case, no other client would be able to access this backup.

Sharing Backups to Directory Devices

At times you may want to create a backup that is to be shared across groups or clients. This is the case if you want to create, for example, a system backup to be used as a "golden" install images for installing other client systems (ie. cloning or provisioning). To do so, simply use a device whose directory does not contain the **%G** (group ID) or **%C** (clientname) notation. For example, if writing to a directory **"/backups/netinst**", all backups in that directory are accessible by all groups and clients.

Making a backup "accessible" to all groups or clients does not necessarily mean that it is also "readable" by all clients. By default, a backup is readable by all clients, but refer to the Disk Backup Read Permissions in the Backup Profile settings for how specific backups, even in a shared directory, may be limited.

Maximum Volume Size

A single backup "image file" or "volume" is created for each filesystem or raw partition (i.e. logical volume, slice, etc) that is included in the backup. Therefore, a filesystem containing 5 GB of data can end up in a single volume file up to 5 GB in size (assuming no compression). If the filesystem you are writing to does not support files of this size, or if you need to limit the size of a single file for any other reason, you can enter a number (in MB) in the **Maximum Volume Size** field.

If this number is non-zero, the backup will create a new image volume file on the disk any time a single image would exceed this value. Do not enter a small value here (minimum is 50 MB) for a large backup as this will unnecessarily create many volumes when one will often suffice.

System and Non-System Backups

You must select one or both of the options **Use for System Backups (network install images)** and/or **Use for Other (Non-System) Backups**. This indicate that this device option should (or should not) appear only when creating a backup job for a system or non-system backup, respectively.

This simply allows you to keep system backups, most commonly used to reinstall clients (network install images), and non-system (user data) backups separate, and prevents you from selecting the wrong device when configuring a backup job.

Default Directory Devices

For *Network Edition*, each time a server is added, three directory devices are automatically configured. This is mainly done to provide default backup options that illustrate the special naming convention for the physical directories and how that affects the location and sharing of the backup files. Those directories are:

- **/backups/data/%G/%C** Non-system backups placed in a private directory for each group and client. These backups are not shared among clients.
- **/backups/system/%G/%C** System backups placed in a private directory for each group and client. These backups are not shared among clients.
- **/backups/system/%G/shared** System backups placed in a directory for each group, not shared among groups. However, all client backups are in one directory that other clients in the same group can read (unless you specify host read-only permission for each backup). These backups are suitable for cloning one client system onto another or sharing of backup data.

You do not have to use any of these devices, and can remove them as long as they are not already assigned to a backup job. See Separating Backups by Group and Client above for more on the **%G** and **%C** notations.

Change a Backup Device

To change an existing device, select one of the following:

1. Configure→Backup Devices from the menu bar. Select the server and device from the drop-down lists at the top.
2. When **Clients, Servers and Devices** are shown in the Main Screen (Network Edition), click on a server, then the desired device, then press the **Change** button below.
3. Configure→Servers (*Network Edition*). From the server configuration screen, select a device from the **Device** listbox, then **Edit** button directly below.

The Device Configuration screen will appear. Make the desired changes and press the **Save** button to save and clear the options from the screen.

Remove a Backup Device

To remove an existing device, select one of the following:

1. Configure→Backup Devices from the menu bar. Select the server and device from the drop-down lists at the top, then press the **Remove** button at the bottom of the screen.
2. When **Clients, Servers and Devices** are shown in the Main Screen (Network Edition), click on a server, then the desired device, then press the **Change** button below. When the Configure Devices screen displays, press the **Remove** button at the bottom of the screen.
3. Configure→Servers (*Network Edition*). From the server configuration screen, select a device from the **Device** listbox, then **Edit** button directly below. When the Configure Devices screen displays, press the **Remove** button at the bottom of the screen.

9. Local System Backup Device

System Backups are primarily used for full system recovery of a client, or for cloning new systems from a backup of another client. This option allows a *client* or *workstation* to perform a full system backup to a device local to the client or workstation. No backup media server is required when performing a backup to a Local System Backup Device.

There are 3 types of Local System Backup Devices: Disk, NFS, and Tape. The type available depends on the license type:

- **Disk** - Network Edition and Workstation Edition
- **NFS** - Workstation Edition
- **Tape** - Network Edition and TSM Edition

Disk Local System Backup Device

A *Disk Local System Backup Device* (SBDIR) allows a *client* or *workstation* to perform a backup to local disk(s). The disk(s) may also be made bootable so that it may be used as system recovery boot media. At least one dedicated disk is required, but multiple disks may be combined when more backup space is needed, or configured individually for use in a backup rotation scheme.

The *Local System Backup Disk* option is available for the *Network Edition* and *Workstation Edition* license types and allows you to use a spare hard disk, portable/USB disk, or SAN-attached disks as full system recovery media. No server is required when a client is backed up to its own local media.

When using a *Network Edition* license, this option is configurable for each client and appears in the client configuration screen under the **Configure→Clients** option (see Configuring Clients). For *Workstation Edition* license, this is found under **Configure→Backup Devices→System Backup Disk/NFS**.

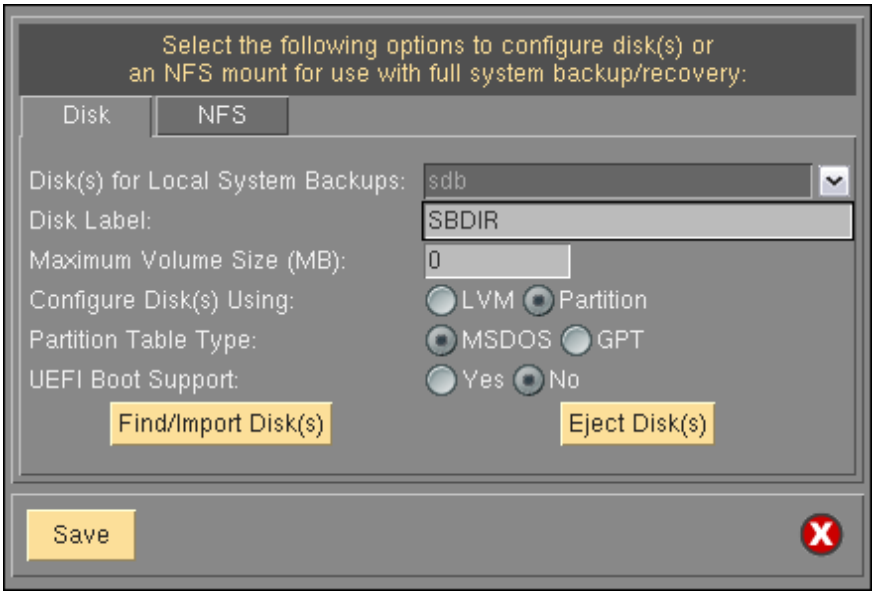

The above screen is for Workstation Edition. A screen with the same options will appear when configuring a client with Network Edition. The "Partition Table Type" and "UEFI Boot Support" are options that will only be seen with Linux systems.

In the **Disk(s) for Local System Backups** field, select the arrow to the right to display and select one or more disks to use. When selecting more than one disk Volume Groups or Zpools will be used to combine the disks into a single device. If disks are to be used in rotation they must be configured individually. 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 labeled IN_USE are disks marked removable and may be mounted by the operating system. They will be automatically un-mounted if used.

Disks containing partitions (*Linux*) or slices (*Solaris*) are considered available for a Local System Backup Device if they are not mounted or defined to be mounted.

Disks configured as LVM Physical Volumes (*Linux*/*AIX*) are considered available if they are either not assigned to a Volume Group or the Volume Group was exported. Although not mounted, if there is an entry in the /etc/fstab (*Linux*) or /etc/vfstab (*Solaris*), SBAdmin will consider the disk not available.

In the **Disk Label** field, enter something that may be used to identify the physical disk(s).

By default SBDIR will be used as the disk label. All disks may use this label and will still contain a unique identifier for telling devices apart.

In the **Maximum Volume Size** field, enter the size (in MB) the maximum file size of the backup image that will be created. Enter “0” for unlimited. See section Maximum Volume Size for more information.

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, and will depend on the operating system of the client:

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

- **Configure using ZFS** (Solaris) – For systems using ZFS, one or more disks may be used to create a ZFS pool and ZFS filesystem for used for system backup and recovery.

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

- **Partition Table Type** (Linux) – It is recommended that a GPT partition table be used when configuring a system backup disk larger than 2 TB. However, when using a GPT partition table, the disk will only be configured to boot from UEFI firmware and not BIOS. The default is MSDOS.

- **UEFI Boot Support** (Linux) – This option is only enabled for systems that have support for booting from UEFI firmware, and selecting “Yes” to this option will make the disk bootable from UEFI firmware. When using a MSDOS partition table, the disk will be made bootable for both UEFI and BIOS. When using a GPT partition table, the disk will be made bootable from UEFI only.

When you have finished your selections, press the **Save** button at the bottom of the screen. The messages indicating the progress of this configuration process will be displayed.

If you wish to physically remove the disk from the system, use the **Eject** button at the bottom of the screen. This will make the disk safe for removal. Ejecting will retain all information on the disk, and allow you to configure a second disk to be used in rotation.

If you have previously configured a *Local System Backup Disk* on another *Client* or system running *Workstation Edition*, you may import the configuration by physically attaching the disk and using the **Find/Import** button. The previously configured settings will populate the screen and backups may be run to the device.

If you no longer wish to use a disk as a *Local System Backup Disk*, expire all backups from the disk and use the **Unconfigure** button to remove all configurations. This process will remove all data from the device.

System Install Boot Disk

Upon creating a System backup disk, the disk will be made bootable and can be booted 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/BIOS/UEFI, 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 Edition*) to restore from if you do not want to use a backup on this disk.

Although this process will allow you to configure 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.

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

NFS Local System Backup Device

This option is available for the *Workstation Edition* license type only, and allows you to use a mounted NFS share as suitable media for a full system backup. The device will be referred to as **SBNFS** throughout the application. You may also perform system recovery from a NFS share.

Network Edition license supports using locally mounted NFS shares on a client as backup media with Shared NFS Servers.

To configure the NFS Local System Backup Device, the *workstation* must be able to NFS mount (read-write) a remote NFS share. At the time of configuring the device the NFS share may already be mounted. If not, the process will mount the share.

Only a single NFS Local System Backup Device may be configured. If you need to change or re-configure the device, you must first unconfigure it. This may be done by selecting the **Unconfigure** button.

NFS versions 2, 3, and 4 are supported.

To configure and manage the SBNFS device select **Configure→Backup Devices→System Backup Disk/NFS**

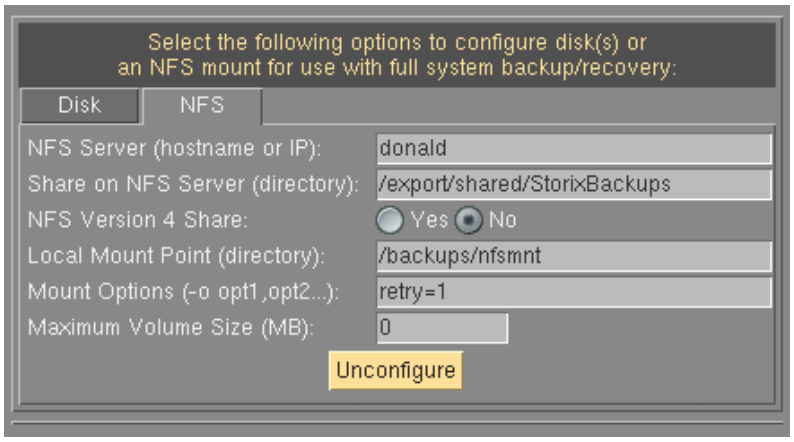

In the **NFS Server (hostname or IP)** field, enter the hostname or IP address of the NFS server that is exporting the share.

In the **Share on NFS Server (directory)** field, enter the directory path of the share on the NFS server.

The **NFS Version 4 Share** option allows you to indicate whether the NFS share that is to be mounted is a NFS v4 share. This is necessary because depending on the operating system type, special handling may be required for NFS v4.

In the **Local Mount Point (directory)** field, enter the directory to be used for storing local system backups. This directory will become the mount point for the NFS share.

In the **Mount Options (-o opt1,opt2…)** field, enter any options to the *mount* command required to mount the NFS share. The options should be comma separated. You should not include the "*-o*" in the options.

> **Using the directory specified here, a special device called "*SBNFS*" will be created. This option will be made available as a backup or restore device in other options.**

In the **Maximum Volume Size** field, enter the size (in MB) the maximum file size of the backup image that may be created. Enter "*0*" for unlimited. See section Maximum Volume Size for more information.

When you have finished your selections, press the **Save** button at the bottom of the screen. The messages indicating the progress of this configuration process will be displayed.

The configuration process will first check to see if the NFS share is currently mounted or currently defined on the *workstation*. If it is already mounted, you will be prompted as to whether you wish to remount it. If it is currently defined (found in */etc/fstab* (Linux), */etc/filesystems* (AIX), or */etc/vfstab* (Solaris)), you will be prompted as to whether you wish to have the entry edited.

10. 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. Therefore, it is 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 the software is initially installed, a set of pre-defined profiles are provided. 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.

> **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.**

When choosing a different backup type, the options on the screen will change to only display those that apply to the selected backup type. The following is an example of a **System Backup** profile when support for all UNIX client types is enabled:

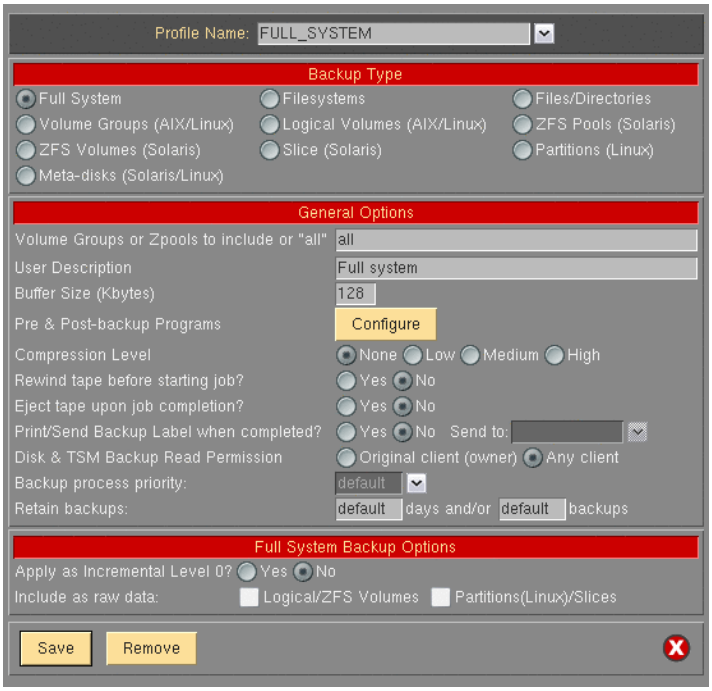

To add a new profile, enter a new profile name in the entry field at the top of the screen, then select the type of backup for this profile by pressing one of the buttons in the **Backup Type** section. A profile name may consist of any characters except a colon (:) or space (spaces will be changed to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

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 upon job completion***" 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.

Buffer Size

The buffer size represents the amount of data to accumulate in memory before writing that "buffer" to the backup device. In actuality, SBAdmin uses many buffers for best performance, but the amount of data written to any device at one time is set using the **Buffer Size** option.

Using a buffer size larger than the physical device can handle will result in an I/O error writing to the device, which generally varies by operating system or device driver. The default of 256K is adequate for most devices without exceeding their hardware limit. However, for best performance, especially when using high-speed tape drives, disk drives and RAID devices, you can increase this number. A value of 512K or 1024K is often best.

Note that, while not excessive, a larger buffer size will cause SBAdmin to use more memory during a backup. Also, if you use a large buffer size, you may see a backup slow down if the system is unable to write a large buffer to the device fast enough, or if there is limited memory on the system. For this reason, it's best to experiment with different buffer sizes until you find the best backup performance for your device.

Specifying the Data to Backup

The description of the first field in the **General Options** section 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 (or ZFS pools for Solaris) 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. Likewise, leaving the field empty on a ***Solaris*** system will exclude all filesystems and volumes contained in *ZFS pools*.

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 containing "/var /tmp /home", a client without a /tmp filesystem will only backup "/var" and "/home".

Disk & TSM Backup Read Permission

If you have the ***TSM Client Backup Feature*** installed, it is important to note the **Disk & TSM Backup Read Permission** field in the backup profile because this will determine if the backup of this client (node) will be readable by another client (node). When a backup is created, it is stored in a *TSM filespace* that corresponds to whether the backup should be private (owner-access only) or shared (any node can access).

For backups to a directory on a backup media server, setting this option will ensure that only the original client can read the backup. If you inadvertently set the permission incorrectly, you can later change the read permission of an existing backup by selecting the Change Read Permission of a Disk Backup.

If the backup data should not be accessible between different clients, be sure to set this option to "***Same client only***". Otherwise, select "***Any client***".

Note, however, that if you have created a backup that you wish to have installed onto different clients, such as used for system replication (cloning), you must make allow access to any client. Otherwise, only the original client can access it.

Backup Retention Policy

By default, all backups will be retained or overwritten according to Backup Retention Policy set in the Preferences section. However, certain backups or types of backups you may wish to retain for a longer period of time. For example, you may want to retain a full backup taken at year-end for several years, while a daily backup (replaced on a daily basis) may only need to be retained for a week. Rather than choosing to retain a backup for a certain number of days, you can also choose to retain a certain number of backups of the same job.

If you set a retention policy in the profile settings, it will override the policy set forth in the ***Preferences***. If you choose to customize a profile for a specific job, it will apply only to backups created by that job and will override the policy for both the main profile (here) and in the ***Preferences*** section.

For details on the use of the **Retain backups** field, refer to Backup Retention Policy set in the Preferences section.

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 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:

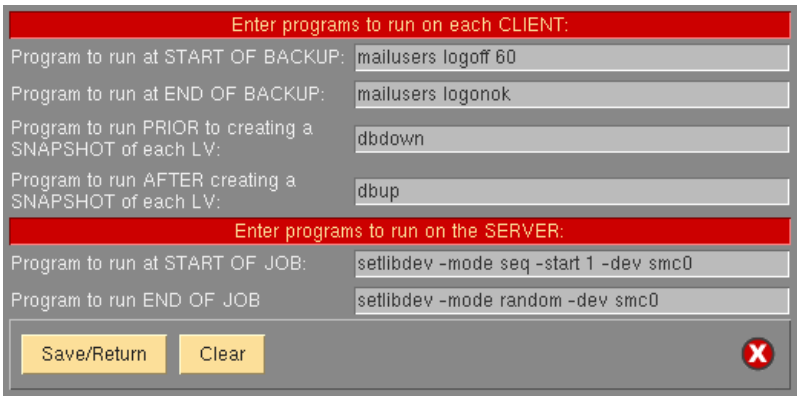

Note that the option for running "**Programs prior to or after creation of snapshot**" is only available with backups that are performing snapshots.

Using the web interface is slightly different. You will notice the option to configure pre and post backup programs directly from the main profile configuration screen:

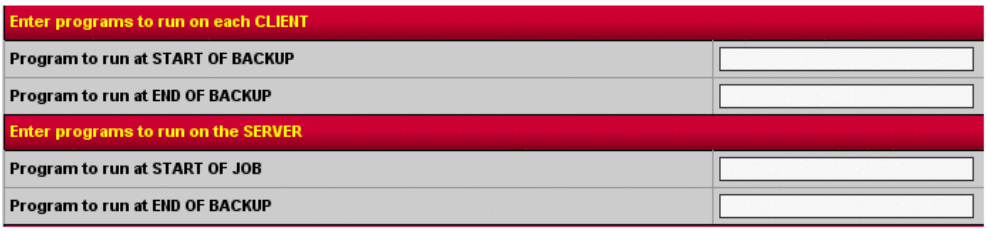

The pre-backup and post-backup programs will be executed with **ROOT USER** authority. Therefore, they must be placed in the ***DATADIR*/custom** directory by the root user on the client (where *DATADIR* is the directory you selected on each system when SBAdmin was installed - i.e. */storix*). The *custom* 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.

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 snapshot is created.

For Linux systems, snapshots are created for LVM logical volumes. For AIX systems, snapshots are created for JFS2 filesystems. For Solaris systems, snapshots are create for ZFS datasets.

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

The program names provided will be executed before the snapshot is created for *each snapshot source*. Therefore, the program must be intelligent enough to recognize the name of the device 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 will be created and the backup will continue.	The backup will continue normally.
Exit code 1	The snapshot will not be created and the backup will terminate.	The backup will terminate with an error.
Exit code 2	The snapshot will not be created, but the backup will continue using the active (online) data.	The backup will continue normally.
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

Server pre and post-backup programs are only available with Network Edition, but will be ignored for backup jobs sent to a *TSM Server*.

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 the server 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.

A post backup job program will be executed even if a client backup fails or another error occurs. 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_CLIENT	The name of the client
STX_SERVER	The name of the backup server
STX_DEVICE	The name of the device on the server
STX_JOBID	The Job ID
STX_BACKUPID	The Backup ID
STX_EXITCODE	The exit code of the backup command or job
STX_SNAPDEVNAME	The device name for which a snapshot is created.
STX_SNAPFSNAME	The filesystem name (mount point) of the snapshot device. This will show a dash "-" if the device is a logical volume without 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 one 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, 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.

Raw devices such as *logical volumes* (***AIX***) and other partitions (slices or *ZFS volumes* on ***Solaris***, meta-disks on ***Solaris*** & ***Linux***, etc) that do NOT contain mounted filesystems will always be backed up in their entirety if they have been written to since the last backup of a prior level. This assumes that you selected to “**Include as raw data**” this information in the backup profile. If included, we assume these raw devices will be included with the backups as well as the “files”.

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 once again. This is the new backup from which all subsequent backups will be based. Any incremental backups performed prior to this level 0 will be considered 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, each weekday the *same* backup level will be performed. Since all data will be backed up since the last ***prior*** level, your previous day's 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 again. The size of the backup will grow each day until after the next Friday night backup is again 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. This will remove and replace all files in each filesystem.

b. Always restore full **Volume Groups** or **Filesystems** from incremental backups. If you choose to restore a *directory* from a Filesystem backup, all files will be restored from the backup, but changes will not be re-applied, such as re-removing files which had been removed prior to that incremental backup level.

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 than your level 1.

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. Any prior versions of the same backup level are considered obsolete.

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 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 options screen will then appear. Either enter the name of the profile in the field at the top of the screen or select the arrow button to the right of the entry field and select an existing profile from the list.

The profile settings will then be displayed. 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, enter or select the profile to remove, and then press the Remove button at the bottom of the screen.

11. Random Tape Libraries

A *tape library* is a device that contains one or more tape drives and is capable of moving tapes between the tape drives and various tape storage slots.

Most libraries can be configured either as a *sequential autoloader* or a *random library*. If it's set to sequential mode, you don't need to configure it as a random tape library in SBAdmin. But to use it was a sequential autoloader, you need to create a tape backup device and indicate that it is contained within an autoloader.

A tape library is a tape *changer*, not a backup device itself. Therefore, to use a tape library, you must create a tape backup device and specify that it is contained in a random tape library.

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.

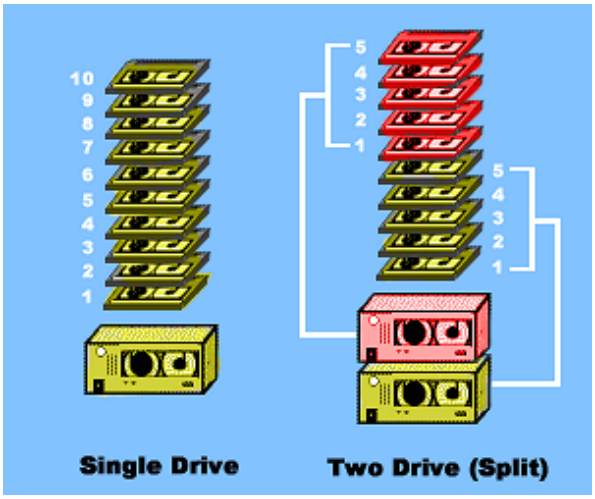

Multiple Drive Libraries

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 tape 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 Tape Write Policy for more information on multi-drive configurations.

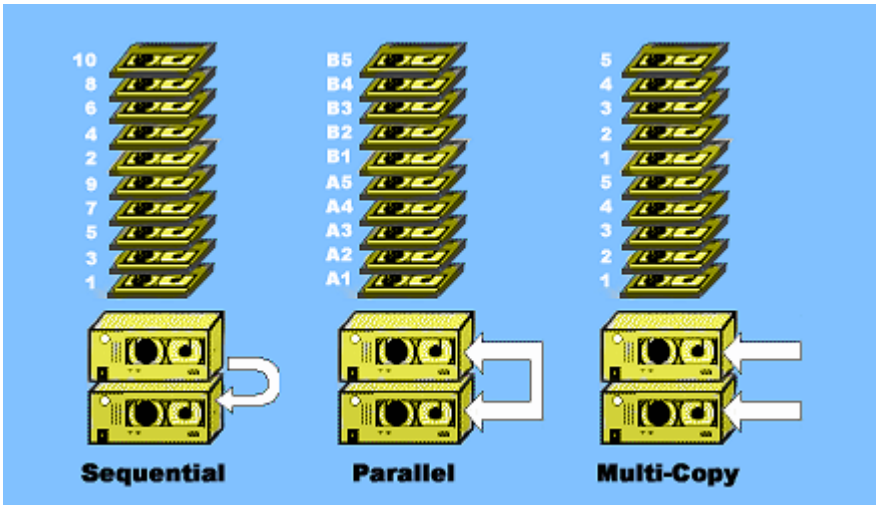

Using Multiple Drives in a Library Independently

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

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 screen similar to the following example will appear.

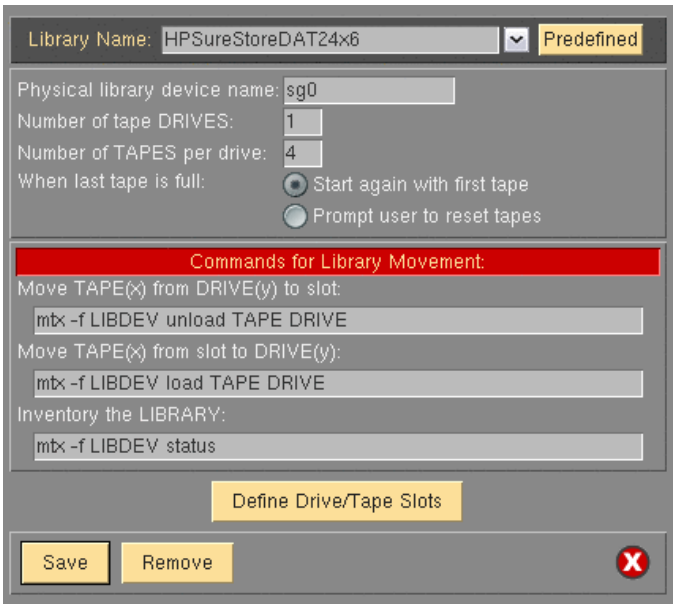

To create a new library definition, enter a unique name of the library in the **Library Name** field. You must then enter information in each of the remaining fields and define the library tape drives and tape slots.

If you want to add a new library from a list of predefined libraries, select the **Predefined** button and search for the library that best fits your environment. After making a selection, the name of the library is shown in the Library Name field and the configuration of that predefined library is displayed. 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/Solaris***) 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 an open-source utility available for Linux and Solaris systems 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:

DEVICE: This keyword will be replaced with the name of the backup device (the device your library is assigned to).

TAPEDEV: This keyword will be replaced with the physical tape drive 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 backup device is remote.

CLIENT: This keyword will be replaced with a client name when performing a backup of a client (only if *Network Edition* 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 DEVICE TAPEDEV SERVER /tmp/liblog
```

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

```
#!/bin/sh
action=$1
libdev=$2
drivenum=$3
tapenum=$4
devname=$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 $tapedevname ($devname)" >>$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.

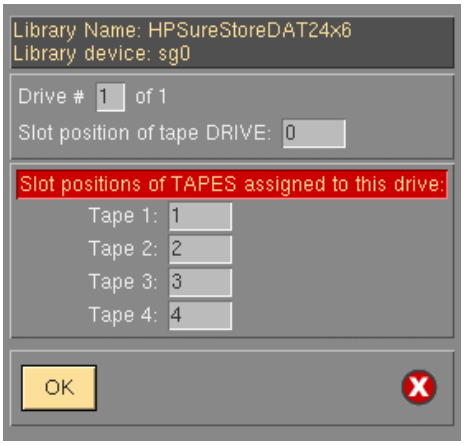

The slot positions referred to are the ***physical slot*** or ***element location*** that the library uses to reference the positions of tapes and drives. The location you provide will determine what tape position is assigned for each tape used with SBAdmin. In the above example, SBAdmin's *tape number*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

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.

12. 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.

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**:

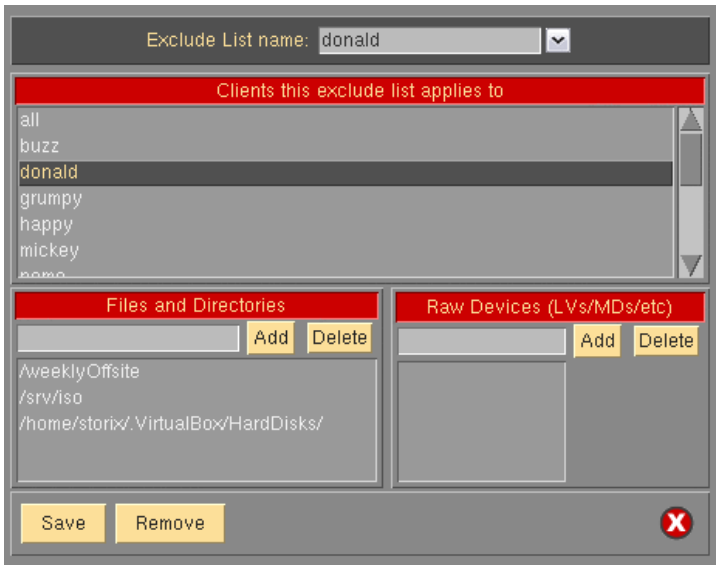

The *Clients* listbox will only appear when using Network Edition and TSM Edition.

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 using the arrow button to the left of the entry field. When doing so, the current settings for the selected exclude list, if any, are displayed.

In the first (**Clients**) 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. Note that, if this exclude list does not apply to a particular client, that exclude list will not appear as a selectable option when configuring a backup job. If, however, the exclude list is applied to a client, this does not automatically apply to all jobs. You must also select to use the exclude list (or lists) when configuring a backup job.

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***), *slices* and *ZFS volumes* (***Solaris***) or *meta-disk* (***Linux/Solaris***) 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 the device types that are supported for the various client operating system types enabled.

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** button at the bottom of the screen to save the entries and clear the entries. 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 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** 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, enter or select the exclude list at the top of the screen, then 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. You will be informed if the exclude list is assigned to any jobs before proceeding.

13. 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 a backup from the client using the **stbackup** command. The job information will identify the backup server, one or more clients to backup (when using *Network Edition*), the backup profile, and the device 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 *Network Edition*, 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.

> **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.**

Creating a Backup Job

To create a backup job, either

1. Select Configure→Backup Jobs from the menu bar, or
2. Press the **Add** button on the Main Screen when the Job Information is displayed.

The following **Configure Backup Job** screen will be displayed:

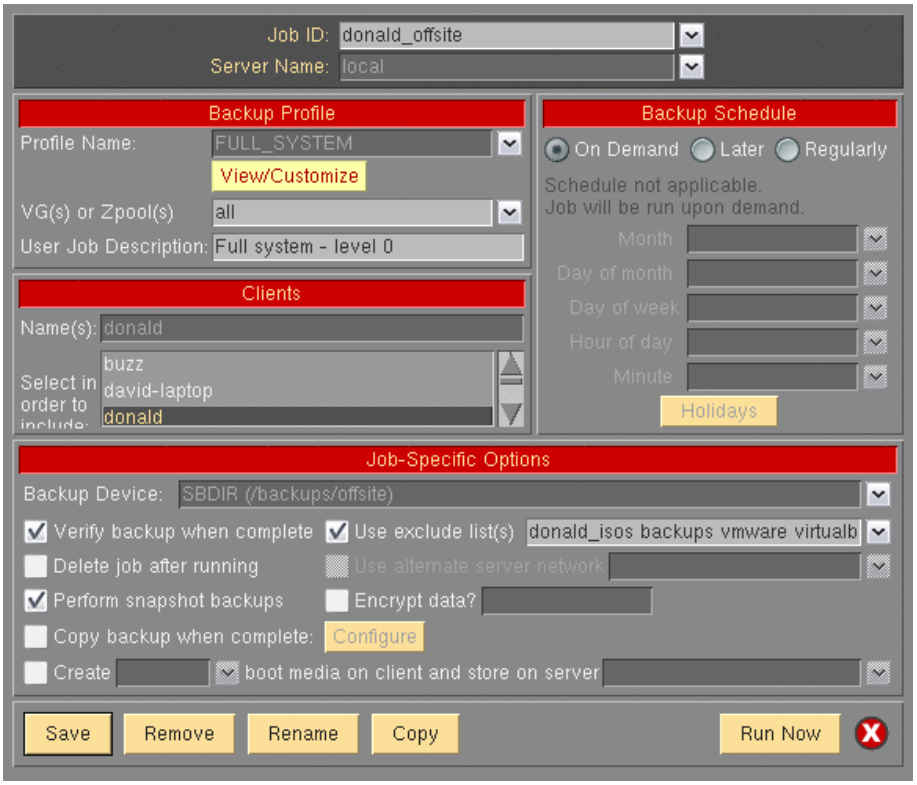

To create a new backup job, enter the **Backup Job ID** in the entry field at the top. The Job ID is used as unique identifier for this job, and may consist of any letters or numbers except for a colon (:) or space (spaces will be automatically replaced with an underscore).

Next, if you are using *Network Edition*, you must select a server using the arrow button next to the **Server Name** field.

When using *Network Edition*, you may also choose "**local (client tape/disk/nfs)**" for the **Server**. By selecting this option, you indicate that you want to backup to a disk (directory) a tape drive, or NFS mount attached to a client system, rather than a device configured on a server. When doing so, you may only select a *System Backup* profile type, a single *Client*, and the *Backup Device* will be either **SBDIR**, **SBTAPE**, or **SBNFS**. Refer to Creating a Local System Backup below for more information.

More details on the various entries which follow are described below. Remember, you may use QuickHelp anywhere on this screen for specific instructions or information on a specific option.

When all selections are complete, press the **Save** button at the bottom of the screen to save the profile and clear the current selections.

When using the *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.

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.

If you selected "local (disk/tape/nfs)" in the *Server* field, only *System Backup* profiles will be available.

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.

Selecting Clients to Backup

If using *Network Edition*, 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 until they appear in the desired order.

Only clients that apply to the selected **Backup Profile** will appear in the listbox. For instance, a *Partition* backup will only display ***Linux*** clients.

Also, if you have selected a ***TSM server***, then only clients that were configured as ***TSM nodes*** will be shown in the client list.

If you selected "**local (disk/tape/nfs)**" in the **Server** field, then you may only select one client from the list, since that client will be sending the backup to its own local media.

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)	**Meta-disk name(s)**
Filesystem mount point(s)	**Slice name(s)**
Files, Directories or @Flist	**Zpool name(s)**
Logical Volume(s)	**ZFS volume(s)**
Partition (PP) name(s)	

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.

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.

You may press the arrow button to the right of this field to list the options available for the selected backup type. Since it's undesirable to query every client (if several 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.

If performing a ***File/Directory*** backup, you also have the option of supplying the name of a file on the system containing the list of files to backup. Referred to as an **Include List**, the file must be a text file containing the name of all files and/or directories to include, beginning with a "/".

To use an include list, at the **Files, Directories or @Flist** field, enter an "@" followed by the pathname of the include list file (i.e. "@/home/Anthony/files-to-backup.txt")

If a directory appears in the include list, all files within that directory will be included. Therefore, it is important NOT to include both the name of a directory (i.e. /home) and its contents (/home/file) or the contents will be included more than once.

Selecting the Backup Media

The backup media option is not available when you have selected to backup to a *TSM server*.

Press the arrow button next to the **Backup device** field to select from a list of devices configured for this server. If the backup type from the profile is a *System Backup*, devices configured for ***System Backups*** will be shown. For all other backup types, the devices configured for other (non-system) backups will be shown. Refer to System and Non-System Backups in the Devices section for details.

For *Network Edition*, if you selected "***local (disk/tape/nfs)***" in the **Server** field, the client will be backing up to its own local media. The options that appear in the **Device** field will be **SBDIR *(disk local system backup device)***, **SBTAPE *(tape local system backup device)***, or **SBNFS *(NFS local system backup device)***, which must have been first configured for the client (refer to Tape for Local System Backups, Disk(s) for Local System Backups, and NFS Share for Local System Backups 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 network***: This option only appears for *Network Edition*, and is only enabled if one or more ***alternate IPs/hostnames*** were configured for the backup server. To set the alternate IP address or hostname for a server, refer to the server configuration. If an alternate network is available, select the check button to enable the field, then the arrow to the right to select from an available network.

  

  If you select an alternate network to use for the job, all clients must be able to contact the server using the IP address or hostname used to configure the alternate network. If one or more clients cannot contact the server on this network, the backup will automatically default to the primary network instead.

  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, there may be multiple networks available for reaching the server from the client, or you may want to offload the heavy backup data traffic onto a different network than other applications are using. The alternate network may use a different network adapter on the client, or may route through a different gateway to reach the server.

- ***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 check button, 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 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.

- ***Use 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 process 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 process 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.

- ***Create Boot Media:*** This option is only available with *System Backup* profiles. This option will create or update boot media for every client selected in the backup. Network or CDROM boot media can be created (or both), and be automatically transferred to the server selected. In the first field, select the type of boot media to create. In the second field you can select the server to store the boot media from the client, or select to keep the boot media on the client.

- ***Copy backup when completed:*** This option may be used only when the backup media is a directory,a and is not available for *Workstation Edition*. It allows you to configure the automatic copying of the completed backup to secondary media (any destination media type is supported) each time the backup job completes. This process may be performed manually for any completed backups as described in the chapter Copying Backups to Different Media. When selected, additional options appear for configuring the automatic copying, as described in the section Automatically Copying Backups below.

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 option, 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 Holidays button to specify certain days, contrary to your backup schedule, on which the backup should NOT run. Refer to Configuring Backup Holidays 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.

Creating a Local System Backup

Using *Network Edition*, it is possible to setup a backup job to perform a *System Backup* of a client to its own locally-attached backup media (Local System Backup Device). You must have previously configured a Tape for Local System Backup (*Network Edition* and *TSM Editioni*), Local System Backup Disks (*Network Edition and Workstation Edition*), or NFS Mount Local System Backup Device (*Workstation Edition*) before configuring a local System Backup

To create a **Local System Backup**, select Configure->Backup Jobs. Enter a name in the **Job ID** field, then choose "**local (tape/disk/nfs)"** from the **Server Name** field.

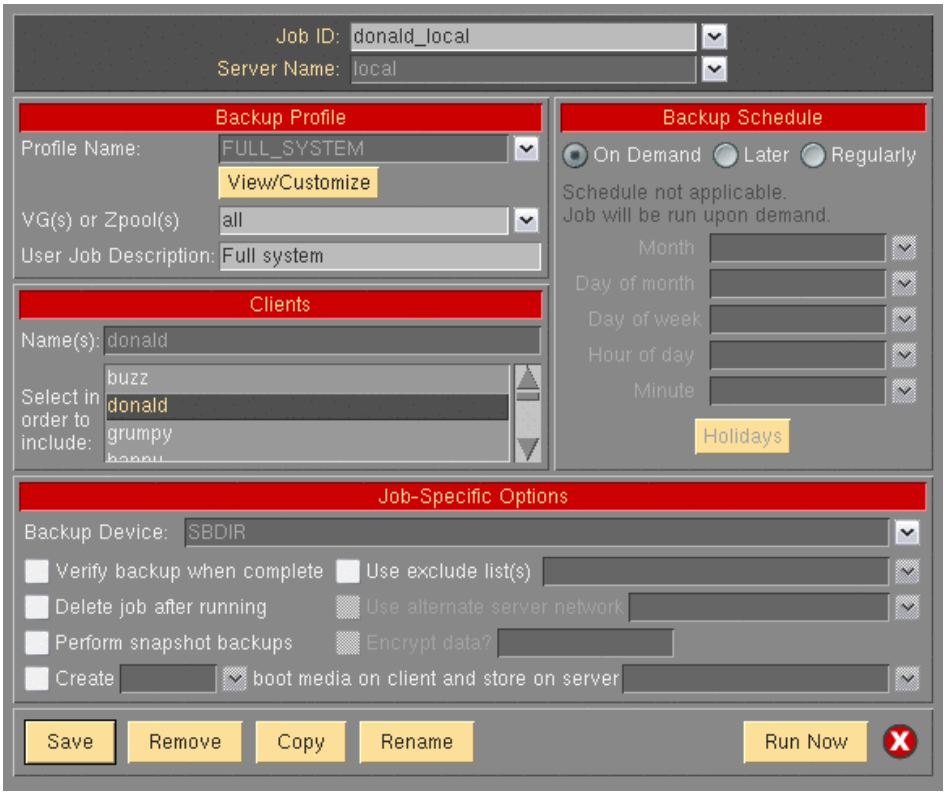

After choosing "**local (disk/tape/nfs)**" you will Configure the Backup Job as usual, with the following exceptions:

1. Only *System Backup* profiles will be available for selection.
2. You may select only one client since the backup will be sent from the client to its own local media.
3. Only the devices named **SBDIR** (backup to local directory/disk), **SBTAPE** (backup to local tape), **SBNFS** (backup to local NFS mount) will be available, and only if previously configured the device for the client or workstation.

Refer to Local System Backup Devices for more details.

Changing a Backup Job

To change information for an existing backup job, either:

1. Select Configure→Backup Jobs from the menu bar, then type or select the **Job ID** at the top of the screen.

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

The current job settings for the selected job will appear. Make all desired changes to the information on the screen, then press the **Save** button to save the changes clear the selections.

Copying a Backup Job

Use this option to copy the job configuration, thereby creating a new backup job with the same settings. To copy a configured backup job to a new job ID, either:

1. Select Configure→Backup Jobs from the menu bar, then type or select the **Job ID** to copy at the top of the screen.

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

The current job settings for the selected job will appear. Make all desired changes to the information on the screen and press the **Copy** button. You will be prompted with a new pop-up window for the job ID to copy the configuration to. After typing in the new job ID, press the **Copy** button and the job will be copied to the new job ID. Note that if you made changes to the options of the original job and did not save them, then these changes will only be applied to the new job ID.

Renaming a Backup Job

To rename a configured backup job to a new job id, either:

1. Select Configure→Backup Jobs from the menu bar, then type or select the **Job ID** to rename at the top of the screen.

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

The current job settings for the selected job will appear. Make all desired changes to the information on the screen and press the **Rename** button. You will be prompted with a new pop-up window for the job ID to rename the configuration as. After typing in the new job ID, press the **Rename** button and the job will be renamed as the new job ID. Note that if you made changes to the options of the original job and did not save them, then these changes will be applied to the new job ID.

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 type or select the **Job ID** at the top of the screen, then press the **Remove** button.

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

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

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 **Restart** 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 SBAdmin interface. To add a job to the queue, refer to the **stqueue command** in the ***Commands Reference Guide***.

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** in the ***Commands Reference Guide***). 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.

Automatically Copying Backups

This option is not available for *TSM Edition*, but may be used to copy to a TSM server when the optional *TSM Backup Feature* is installed. This option may only be used if the original backup is written to a directory device.

You may configure a backup job so that each time it completes, the new backup is automatically copied to another backup server or device. This is often referred to as a **Second-Stage Backup**, or (depending on the destination media) **D2D2T** (disk-to-disk-to-tape) or **D2D2D** (disk-to-disk-to-disk). Use this option if, for instance, a backup job is configured to backup one or more clients to a directory on a local network server, and you want to later copy those backups to another server or device (including tape, directory, **NFS** or **TSM** server).

This option allows the copying to take place automatically each time the job completes without any user intervention. You may still continue to manually copy backups from one server or media to another using the

Copying Backups to Different Media, but this requires you to manually select a backup that was previously created.

If configured, a new *copy backup job* is created after the backup job completes. This job is treated the same as other backup jobs and may be rescheduled or run on demand. However, once run, this copy job is permanently deleted, but will be re-created each time the original backup job completes. Refer to The Copy Backup Job below for more information.

To configure a job for creating an automatic copy:

1. Create a new job or select to change an existing job using the option Configure→Backup Jobs or Actions→Run a Backup Job from the menu bar.

2. You can also select an existing job when the Job Information is displayed on the Main Screen and press the **Change** button at the bottom of the screen.

3. With the Configure the Backup Job, select the checkbox for ***Copy backup when complete***, then press the **Configure** button to enter or change the copy backup settings.

The following screen will appear:

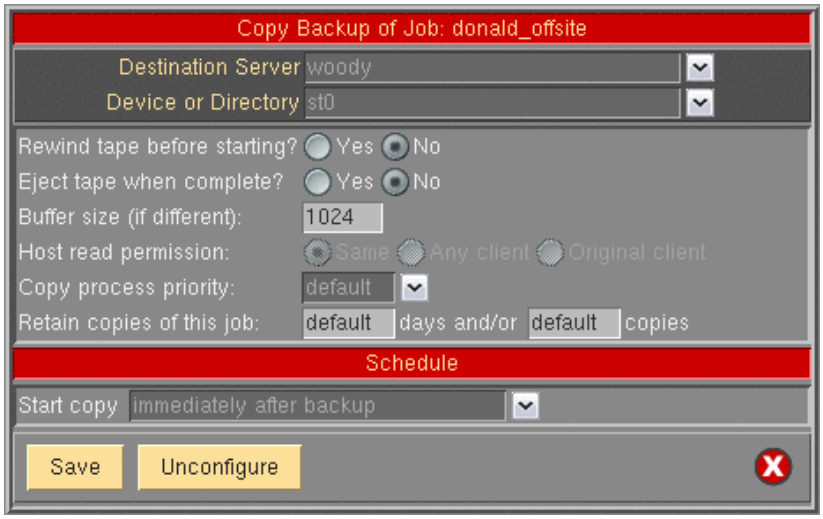

Select the ***Destination Server*** and ***Device or Directory*** at the top of the screen. The backup created by the job will be copied from the server and device defined in the job settings to the server and device selected here. Additional options are explained below:

Buffer Size

The buffer size represents the amount of data to accumulate in memory before writing that “buffer” to the *destination* device. In actuality, SBAdmin uses many buffers for best performance, but the amount of data written to any device at one time is set using the **Buffer Size** option. The buffer size setting has no affect on the *source* device, as the backup will be read using the buffer size at which it was written,

Using a buffer size larger than the physical device can handle will result in an I/O error writing to the device, which generally varies by operating system or device driver. The default of 128K is adequate for most devices without exceeding their hardware limit. However, for best performance, especially when using high-speed tape drives, disk drives and RAID devices, you can increase this number. A value of 512K or 1024K is often best.

Host Read Permission

For backups to a directory on a backup media server, setting this option will ensure that only the original client can read the backup. If you inadvertently set the permission incorrectly, you can later change the read permission of an existing backup by selecting the Change Read Permission of a Disk Backup.

If you have the ***TSM Client Backup Feature*** installed and are copying the backup to a TSM server, it is important to note the **Host Read Permission** field in the backup profile because this will determine if the backup of this client (node) will be readable by another client (node). When a backup is created, it is stored in a *TSM filespace* that corresponds to whether the backup should be private (owner-access only) or shared (any node can access).

If the backup data should not be accessible between different clients, be sure to set this option to "***Original Client***". Otherwise, select "***Any client***". If you want to retain the same permission setting of the original backup, as defined in the backup Profile, select "***Same***".

Note that if you have created a backup that you wish to have installed onto different clients, such as used for system replication (cloning), you must make allow access to any client. Otherwise, only the original client can access it.

Copy Process Priority

You can change the default CPU process priority of all jobs run from the admin system by selecting File→Preferences→General Preferences, and moving the slider in the **Backup Process Priority** section to the desired number.

The copy backup process will use the system default priority unless you set a different priority. SBAdmin represents this on a scale of 0 to 10, with 0 being lowest, 5 being normal (default), and 10 being highest. By default, this will bet set to 5, indicating that the normal operating system default priority should be used.

It is common to lower this value if you do not want the copy backup process to run as the same priority as other applications. This will usually cause the copy process to have less affect on the performance of other applications, but the copy could run a bit slower.

This is the same option as provided in the backup profile settings, but copying a backup is likely to have less effect on system performance than performing a backup.

If you want the copy process to run in the shortest possible time, with no regard to how it might adversely affect the performance of other applications, you can increase this value. Note that doing so will have little or no affect if there are no other applications or processes requiring CPU processing, but those that require the CPU could run slower.

Retention Policy

By default, all backup copies will be retained or overwritten according to Backup Retention Policy set in the Preferences section. However, certain backups or types of backups you may wish to retain for a longer or shorter period of time. In many cases, this feature is used to copy a temporary backup on the local network to longer-term storage, and therefore will want to set a different retention policy for this backup. The retention policy you chose here will apply to all copies of the original backup, but will not affect the retention period of the original backup.

To override the default retention policy, select the ***minimum number of days*** to retain backup copies created by this job, and/or the ***minimum number of copies*** to retain. Note that, by default,

Scheduling the Copy

The "***Start copy***" selection may be set to one of 3 options:

1. **Immediately after backup**: If selected, a new copy backup job will be created and placed in the Job Queue immediately after successful completion of the backup job.

2. **Hrs/mins after backup**: If selected, you must enter the number of hours and minutes after the backup job completes that you want the copy backup job to begin. You may want to use this option if you simply want to push the copy process into the future rather than run immediately. You must enter the hours (may be 0) and minutes, but limited to 23 hours to avoid conflicts with the same backup job running on a daily basis.

3. **Same or next day at (specified time)**: Use this option if you want to select a specific time of day to start the copy backup job. This is helpful to schedule the job to run at a time when it will have less effect on other network traffic. It is important to note that, *if the time of day provided must has already passed when the backup completes, it will be scheduled to run the next day*. Be careful, for example, if a daily backup job runs at 2 pm, and the copy is scheduled for 2:30, the backup may not complete until 2:45. In this case, the copy will start at 230 the next day, and will likely be running at the same time the daily backup job runs again.

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.

The Copy Backup Job

If a backup job is configured to Automatically Copy a Backup upon completion, a new job will be created as soon as the original backup job completes successfully. This new *copy job* will either be scheduled to run or immediately added to the Job Queue, depending on the job settings.

A copy job will only exist after the original backup job completes *successfully* and is automatically deleted after the copy completes.

The new copy job will be configured using the name of the original backup job with a "-COPY" extension. For example, a backup job called "*DailyClient1*" will produce a copy job called "*DailyClient1-COPY*"

The copy job will appear on the Main Screen when the Job Information is displayed. From there, you can select the **Remove** button to remove the copy job (this has no effect on the backup itself or on the original backup job), or the **Run** button to run the copy job immediately. Using the **Run** button assumes the job is not already running but is scheduled to run at a future time.

Note that you cannot change the settings of a copy job after it is created. The copy job is created according to the destination server, device, retention policies, etc, configured in the Copy Backup Settings screen for the original backup job. You may, however, remove the copy job or run it immediately. If you want to send the backup to a different server or device, you can do so manually using the option Copy a Backup Job to Different Media.

If the original backup job was set to copy immediately after completion, or when the time of day has arrived for the copy job to start, the copy job will be placed in the Job Queue. You can monitor or alter the progress of the copy job the same as any other backup job using the options on the Main Screen when the Queued/Running Jobs are displayed. This includes the buttons for **Remove** and **Hold Job** (if it has not yet started), **Restart** (if it has failed), or **Kill Job** (if it has already started). You can view the status and output messages, as well as the progress indicator of a running copy job by selecting the **Show Status/Output** button.

When the copy job has completed successfully, it will be removed from the queue and *the copy job will be deleted*. Of course, the next time the backup job runs and completes successfully, the same copy job will be created and scheduled (or immediately queued) to run again.

14. Holidays

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:

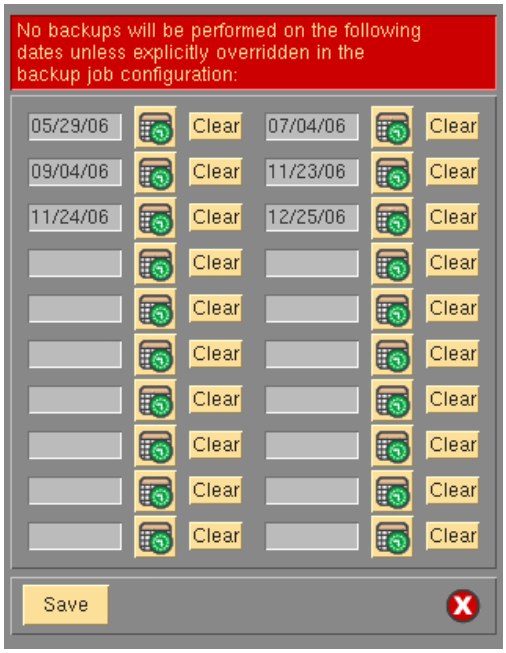

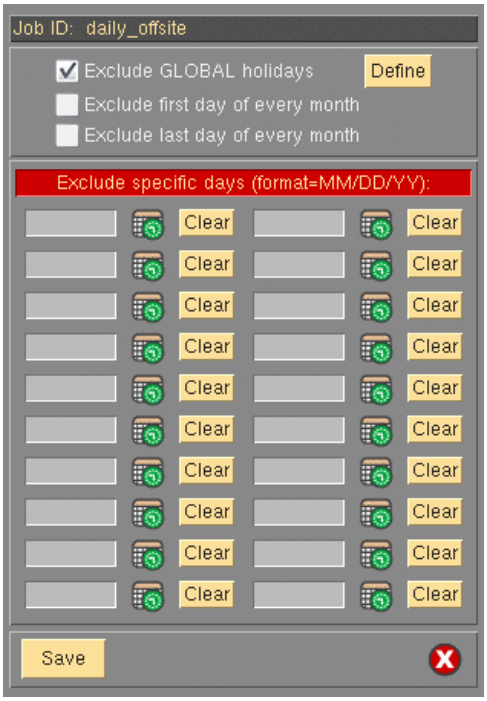

In the first screen, you may enter one or more dates on which ALL backup jobs will be excluded from running. In the second screen, you may enter additional dates in the date fields that will be excluded for this job in addition to those excluded on global holidays. Since it may be cumbersome to enter 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:

If you press a specific date on the calendar, that date will be automatically inserted into the date field and the 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.

When using the web interface the screen presented will look slightly different. To add a holiday, simply click on the date from the calendar. The date will be added to the list on the right. Any dates appearing on the right are automatically saved. To remove a holiday, simply click the **X** next to the date.

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 first 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.

15. Snapshot Backups

> **Snapshot feature is available only for *AIX* systems running AIX 6.1 and later, for Linux data contained in LVM logical volumes, and *Solaris* data contained in ZFS datasets.**

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

- For ***Linux***, an LVM *snapshot logical volume* is created for each logical volume to be backed up. Snapshots may be created for any logical volume, whether or not it contains a filesystem. 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.

- For ***AIX***, snapshots are available only for logical volumes containing *JFS2* (Extended JFS) filesystems. At ***AIX 6.1*** and later, either internal or external snapshots may be used (external snapshots created a temporary snapshot logical volume similar to Linux snapshots described above).

- For ***Solaris***, snapshots are available on ZFS filesystems and ZFS volumes. Snapshots may be created for any dataset, whether or not it contains a filesystem.

As the backup is performed, original data to be changed by another process is first copied to the snapshot LV (or internal snapshot area for AIX 6.1), 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.

Enabling Snapshot Backups

Snapshot backups are configured on each client for 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 or ZFS datasets 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.**

To enable snapshot backups, select the following from the menu bar:

Configure→Snapshot Backups

The configuration screen will appear as in the following sample:

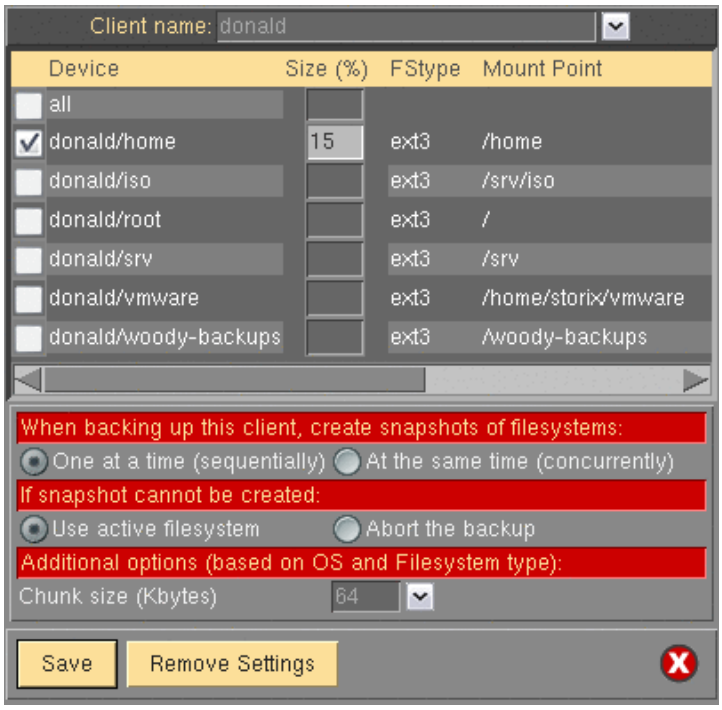

If using *Workstation Edition*, the **Client** entry box at the top does not appear. If using *Network Edition*, you must select the client for which to configure snapshots by pressing the arrow next to the **Client name** field. After doing so, a list of logical volumes or ZFS datasets on that system will be displayed in the second listbox.

You must either select the individual logical volumes or ZFS datasets for which snapshots may be created, or select "***all***". If "all" is selected, snapshots will be permitted for all supported snapshot devices.

> **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(%)**:

   For each logical volume (or for "all" if selected), indicate the size of the snapshot as a percentage of the original LV. The size needed 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.

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

   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, and 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 of 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 (*resynced*). 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*.

      Less disk space is required since only one snapshot LV is created at a time. For ***AIX 6.1*** using internal snapshots, no separate logical volume is created, so the added space does not apply. But, in any case, 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:**

   a. ***Use active filesystem.*** 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 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 be 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:

   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 expand the volume group, remove other unused logical volumes, or select to create smaller snapshots using the Size (%) 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.

16. Job Queues

When jobs are run, they 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 Queued/Running Jobs tab from the main screen:

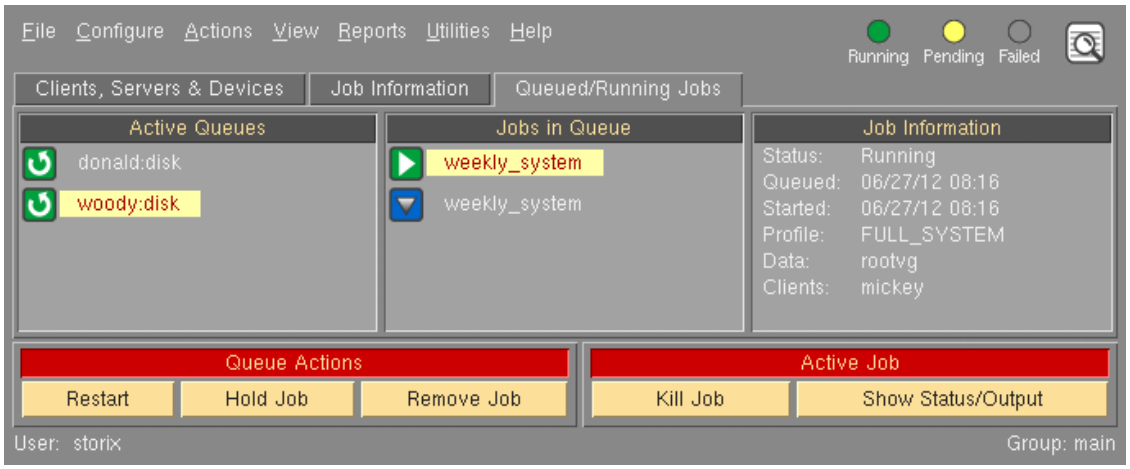

Active Queues

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 (if *Network Edition* used) the device name (for tape devices or local system backup devices) or “disk” (for directory devices). For *Workstation Edition*, only the device name or “disk” is used.

Backups to a directory are handled different than those to tape. Only one backup may be written to a tape drive at a time, but multiple backups can be written simultaneously to directories. It may not be desirable to have too many backups writing to a server’s disk or disks at one time, so the number of concurrent backups to disk on a server (or to a TSM server) can we limited using the option Concurrent Disk Backups in the General Preferences section.

There are 2 special device names, and therefore queue names, that are used for a *system backup* of a client to its local media (a Local System Backup Device). This applies when “***local (disk/tape/nfs)***” was selected as the *backup server* in the job settings with a *System Backup* profile and a single client to backup. The backup device, and therefore the queue name, will be “***client*:SBDIR**” (when backing up *client* to its system backup disk), “***client*:SBTAPE**” (when backing up to local tape), or “***client*:SBNFS**” (when backing up to local NFS mount). Refer to Local System Backup Devices for more details.

Jobs in Queue

To display the jobs within a queue, select the icon corresponding to the desired queue. When doing so, the selected queue will be highlighted 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 placed, and will be run, in the order they were added to the queue. To show a summary of the 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.

The status bar at the upper-right also shows the status of jobs in the queue. If the **Running** light is green, at least one backup job is currently running. The **Pending** light shows yellow if there are queued jobs that have not yet run because they are waiting for another job to complete or are waiting for the server to become available. The **Failed** (red) light illuminates if there is a failed job in the queue (usually preventing other jobs from running).

The action buttons below the list of queues or jobs 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


Job is waiting for a device to become available before starting.

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

By selecting a backup job on the Main Screen and then pressing the **Show Status/Output** button at the bottom of the screen, you can view the backup in progress, or information about a failed backup.

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 Show Status/Output 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:

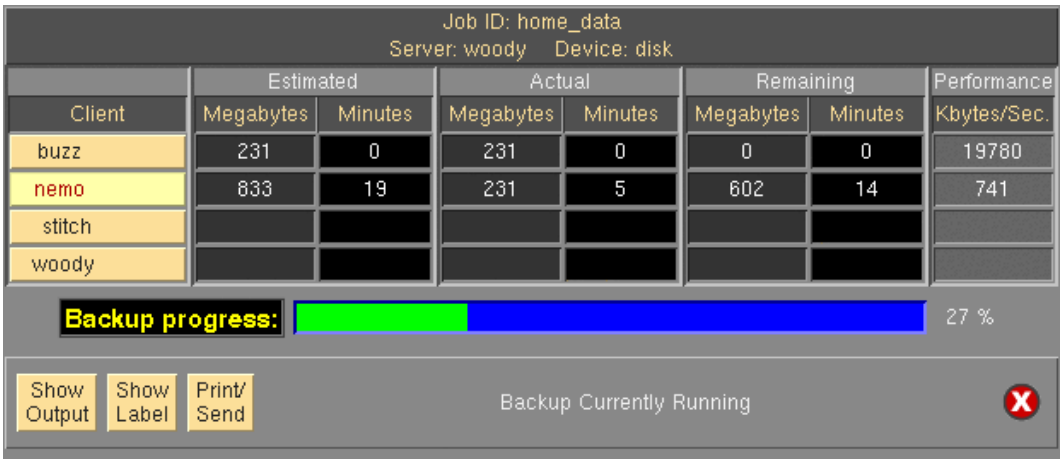

The **Job ID**, **Server** (if *Network Edition*) and **Device** are shown at the top of the screen. The middle section will contain a set of boxes for **each client** in the job. If using *Workstation Edition*, the **Client** column will not appear, and only one progress line will be shown. The corresponding 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 has started, select this button to find out why. The following is a sample output screen:

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**.

Using the ***Web Interface*** the Backup Status Screen and backup output messages are combined into a single screen as shown below:

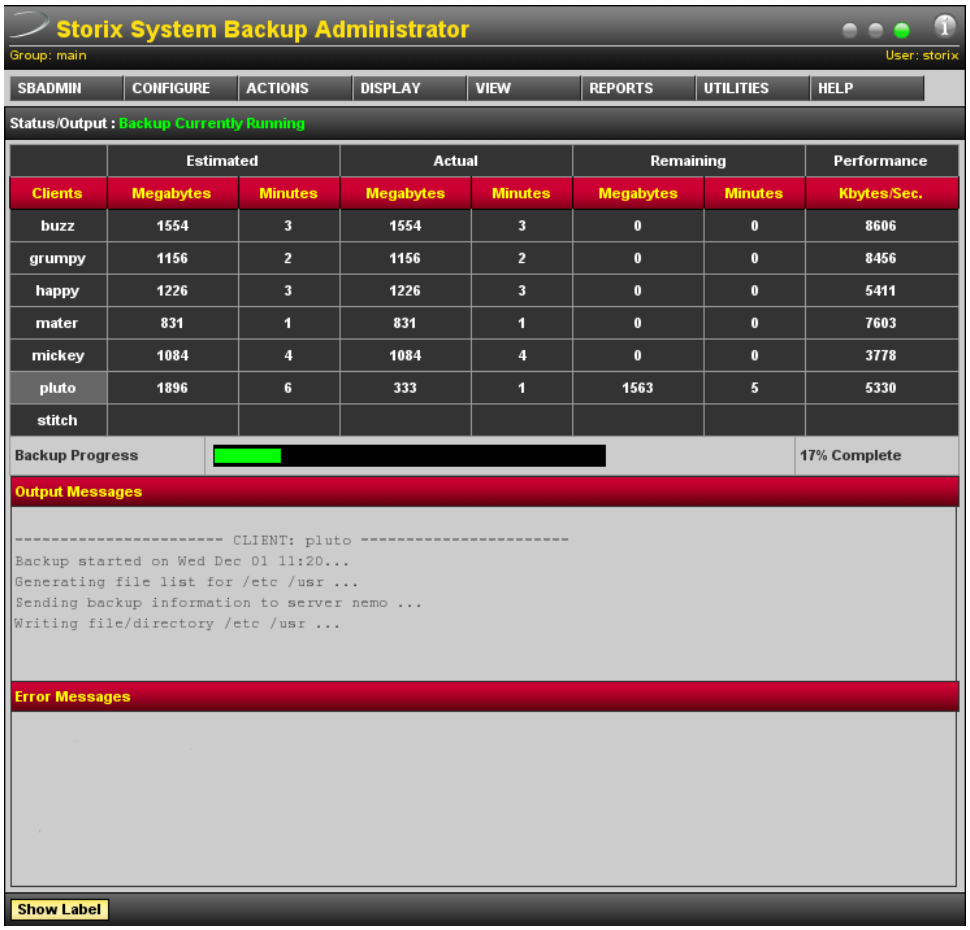

The Job Message Screen

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 Job Message screen will only appear if the SBAdmin user interface is running. If not, messages are logged and can be displayed the next time SBAdmin is started. Refer to the Backup Status Notification section for more information.**

The following is an example of this **job status message screen**:

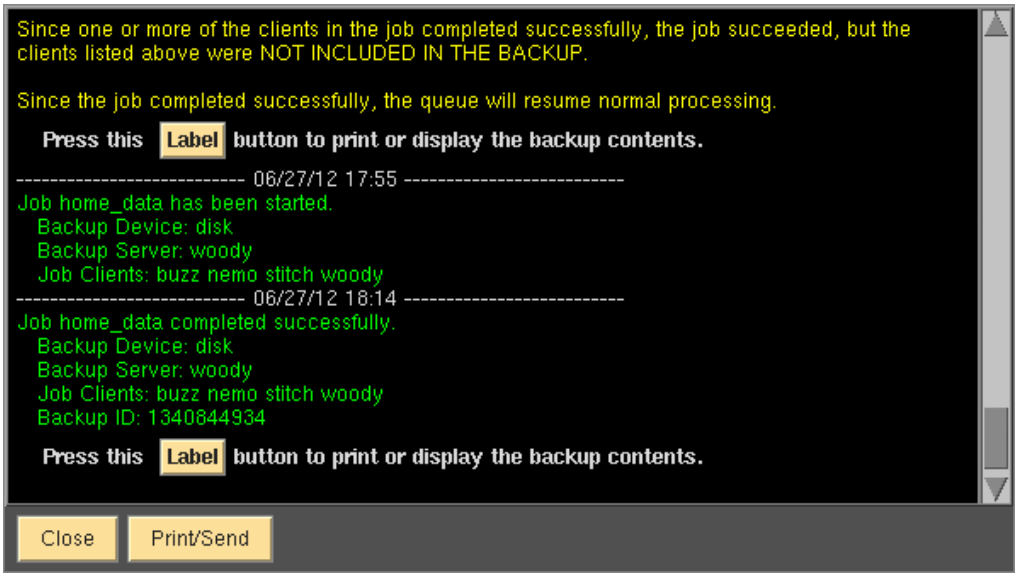

A scrollbar is provided to the right to scroll up and down the messages. This screen may only be displayed if the Backup Administrator application is running. When you **Close** this window, the messages will be removed from the log and cannot be displayed again.

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.

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 under the Queue Actions and Active Job headings 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.

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 **Restart** 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** 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.

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.

Show Status/Output

You can display the Backup Status Screen for the selected job by pressing the **Show Status/Output** button. Here you will see the progress indicator for the job. On that screen you can select a client to display, then press the **Show Output** button to display the Backup Output Display for the backup for the specified client.

17. 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 Edition*) 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 job containing multiple client backups:

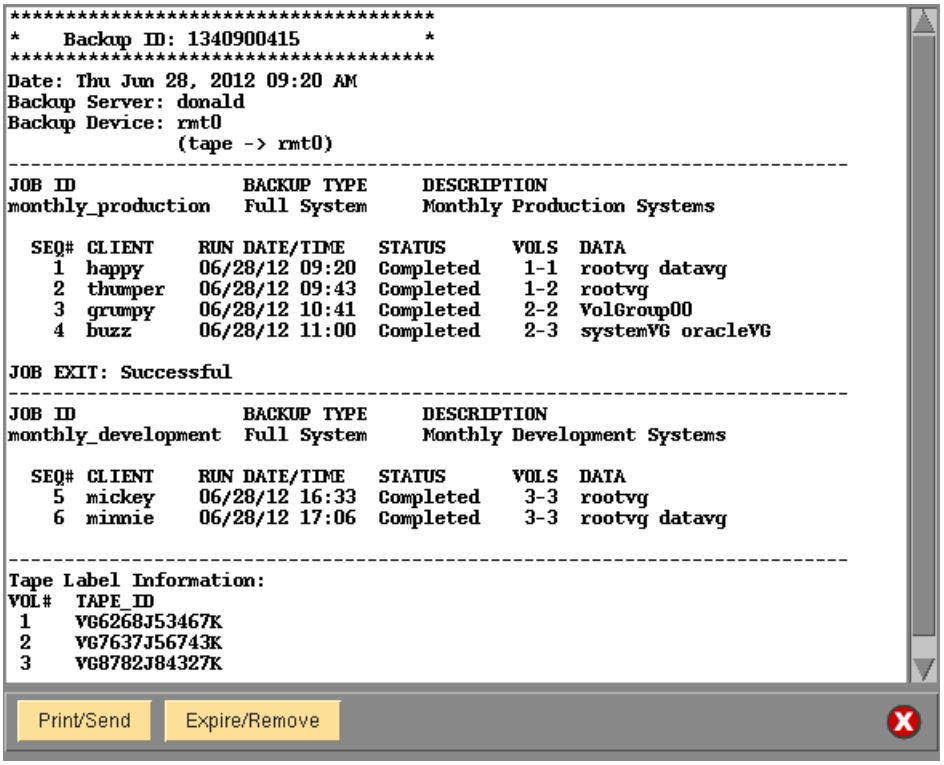

The **Backup ID** appears at the top. This ID is a unique identifier generated automatically for each label and is also stored on the backup media itself. This way, it is possible to read the Backup ID from the backup media and reference its contents in the label information. Also at the top of the label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

The **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 to 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 with SBAdmin 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 **Save** 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.

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:

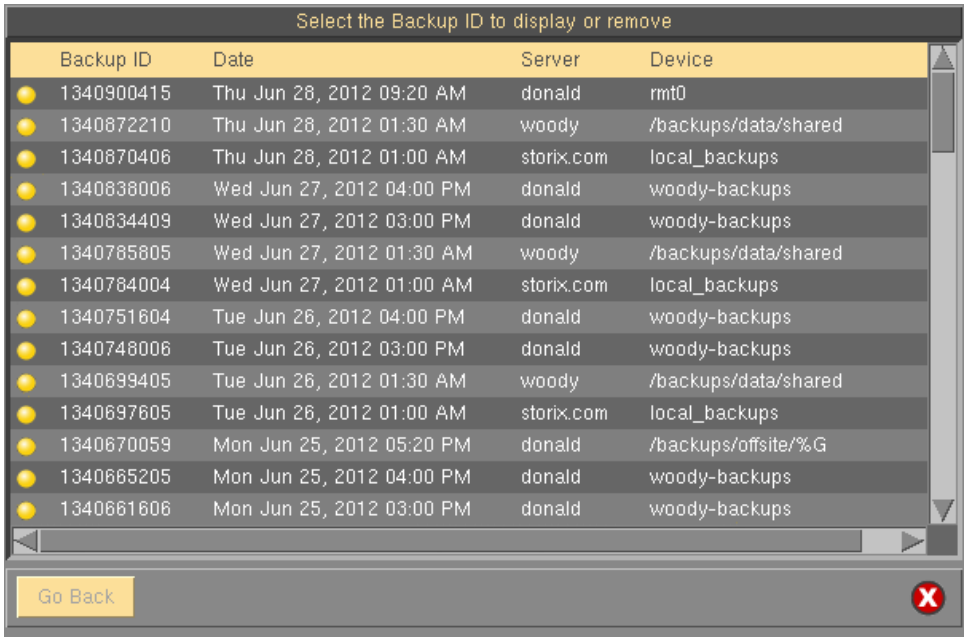

This list could become very lengthy if there are a lot of labels on file. To display the detailed label information, click on the button to the left of the desired Backup ID.

View by Disk Label ID

To display a backup label in which a *Local system backup disk* was used, select View→Backup Labels→By Disk Label ID from the menu bar. A list of disk labels currently associated with backup labels is displayed. Multiple disks may have the same backup label and can be distinguished by unique identifier as shown in the following example:

Select the Disk Label by selecting the button to the left. After doing so a list of backups on the device will be displayed as shown in the image below:

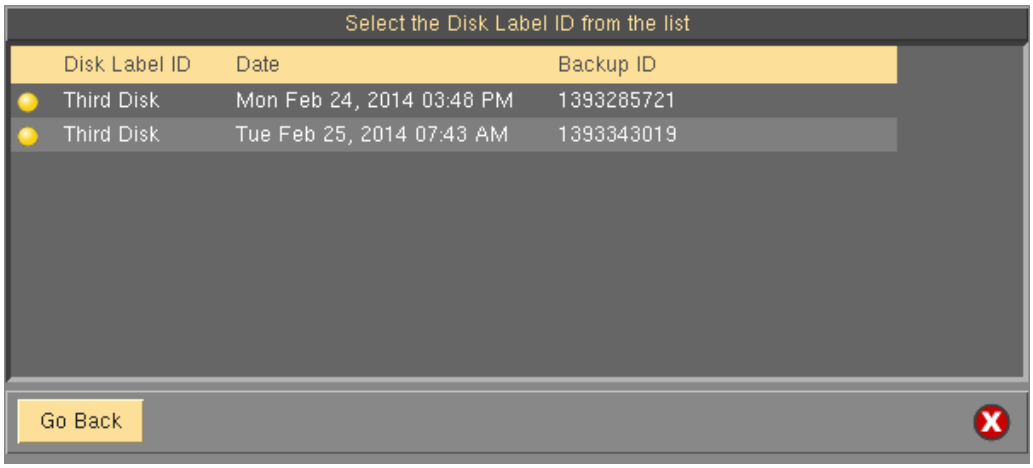

To view the backup label, select the button on the line with the corresponding backup id.

View by Tape Label ID

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:

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

View by Server

This option is not available when using Workstation Edition.

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

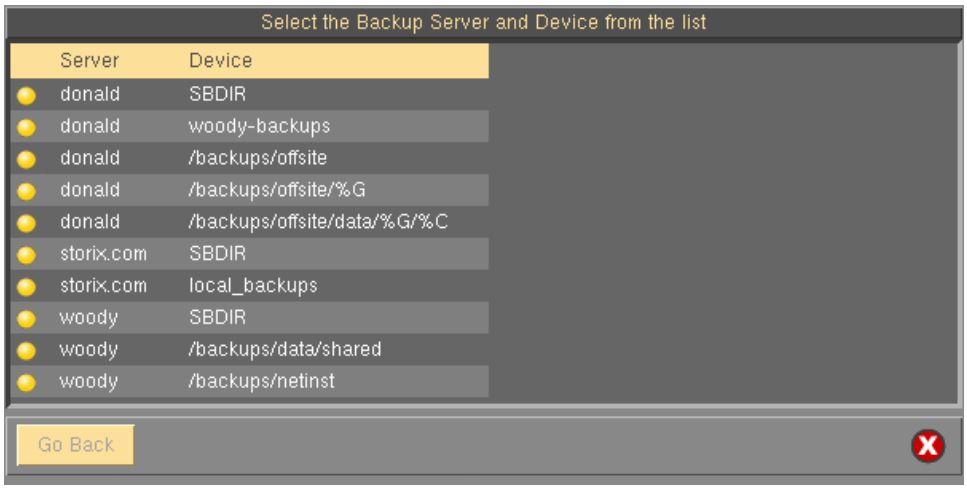

Select the server and device from the list by clicking on the button to the left. Once you do so, a list of backups for the selected server and device is displayed, as shown below:

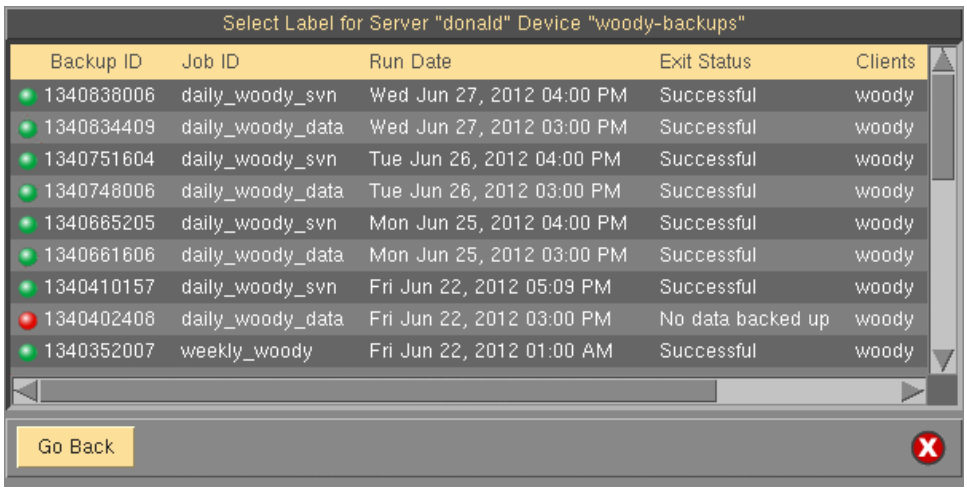

Note that the color of the button represents whether the backup was successful (green) or failed (red). A yellow button indicates that the job completed successfully, but with warning messages. The Job ID, date and time, and a list of clients on the media are displayed for each label in the list. To 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:

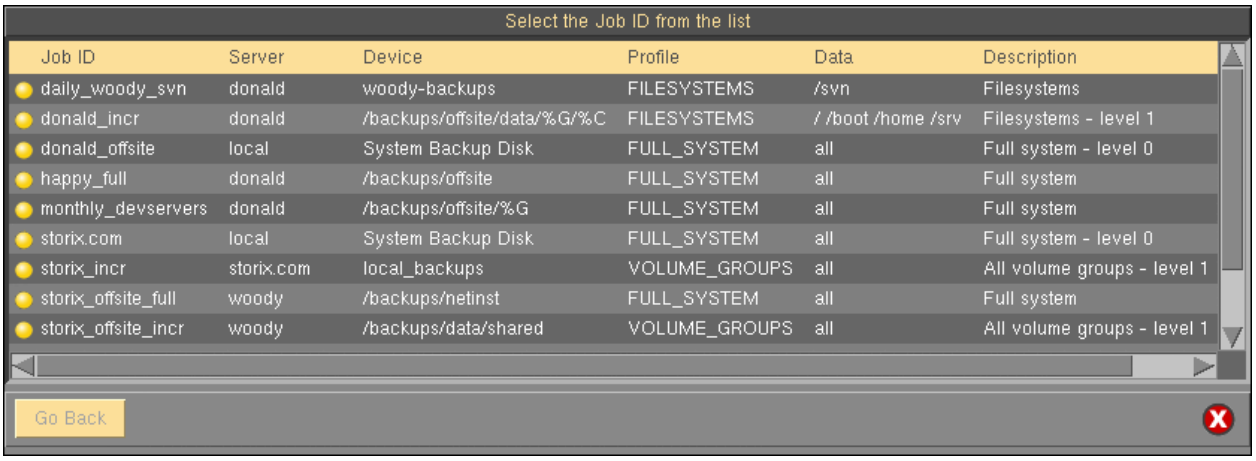

Select the desired job. An additional list will display, showing the dates the job has been run:

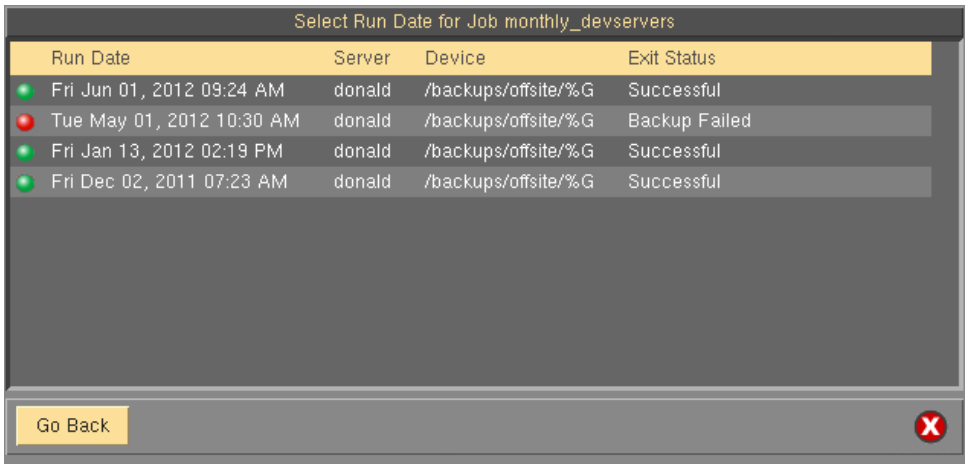

Note that the color of the button represents whether the backup was successful (green) or failed (red). The Job ID, date and time, and a list of clients on the media are displayed for each label in the list. To display a label, select a specific ***run date*** from the list. The label will be displayed (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 not available when using Workstation Edition.**

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 if 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:

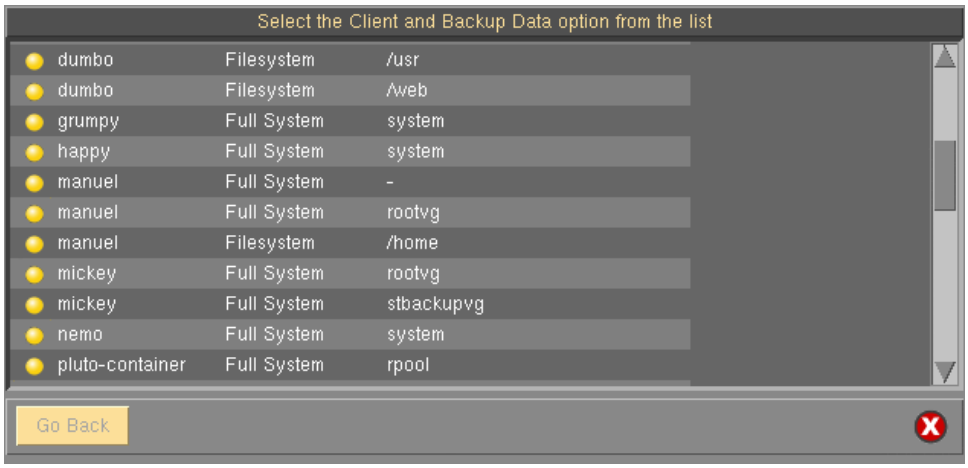

Select the button next to the client and backup type you wish to display. An additional list of specific backup dates for the selected client and backup data will be shown:

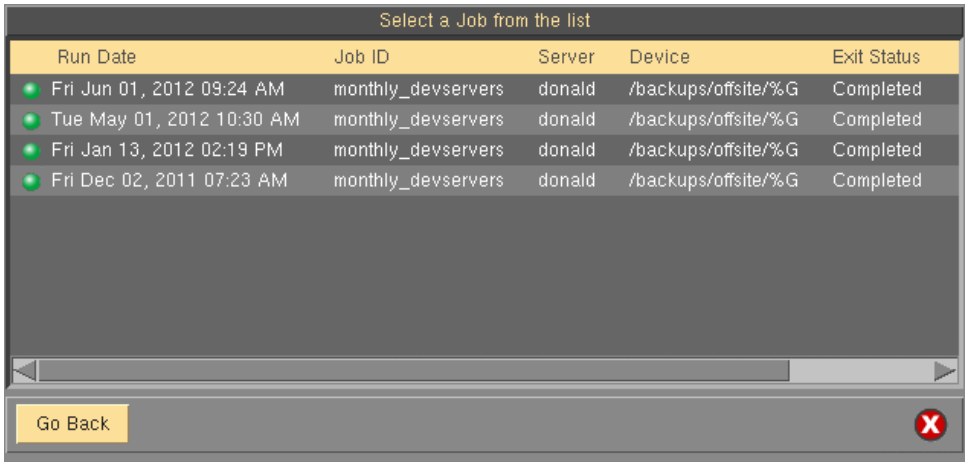

Note that the color of the button represents whether the backup was successful (green) or failed (red). The Job ID, date and time, and a list of clients on the media are displayed for each label in the list. To display a label, select a specific ***run date*** from the list. The label will be displayed (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 Server (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 Server (*Network Edition*) or View→Backup Labels→Read from Media (*Workstation Edition*). A window similar to the following will display:

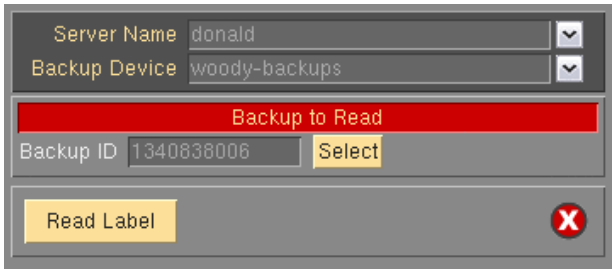

2. Use the arrow button to select a **Server Name** (if *Network Edition*).

3. Use the arrow button to select the **Backup Device**.

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

5. Select the **Read Label** button.

The tape or disk image 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 or TSM 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 removed 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.

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*** determines when a prior backup can be automatically expired.

In either case, the retention policy only applies when writing a new backup using the same backup job. When this occurs, the policy may allow:

1. No backup to be expired, removed or overwritten without manual expiration.
2. Any backup to be automatically expired and overwritten
3. Expiration of backups that are older than a certain number of days
4. Expiration of a backup only if there are a minimum number of un-expired backups of the same job remaining.

18. 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 may appear by selecting one of the following:

3. Select the **Show Status/Output** button on the Job Queues Display
4. Select the **Status/Output** button on the Job Information Display
5. Selecting View→Backup Status/Output from the menu bar (as described throughout this section)

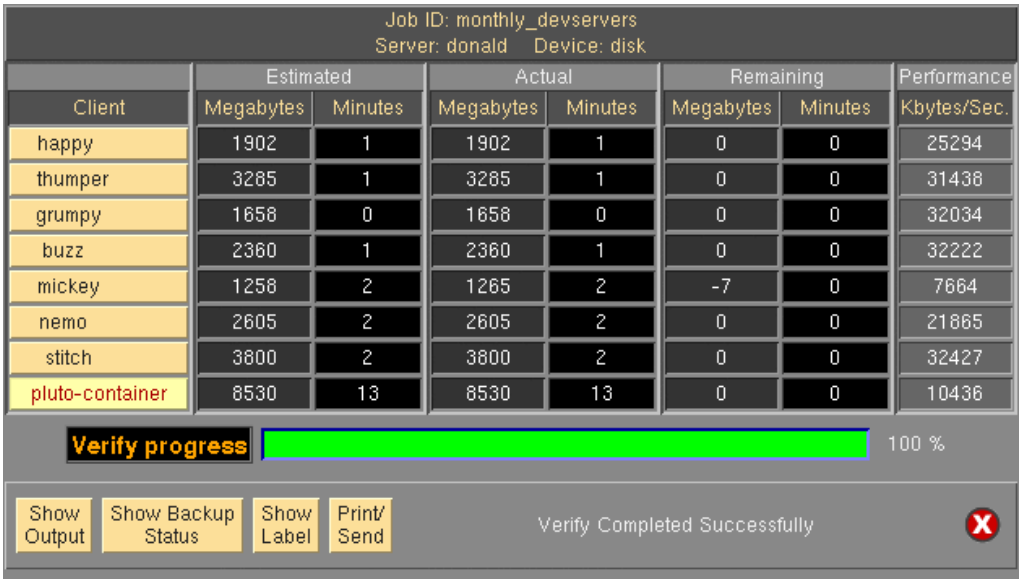

The following is the Backup Output Display which is displayed when the **Show Output** button is pressed on the Backup Status Report screen above:

The desired job status for completed (or failed) jobs may be obtained in each of the following ways:

View by Server

> **This option is not available when using Workstation Edition.**

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

Select the server and device 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 are 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 are 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

This option is not available when using Workstation Edition.

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 are 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.

19. 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 a backup upon completion within the Backup Job configuration, the backup data was verified at that time, and it is generally not necessary to verify again. However, if you did 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. A window similar to the following will appear:

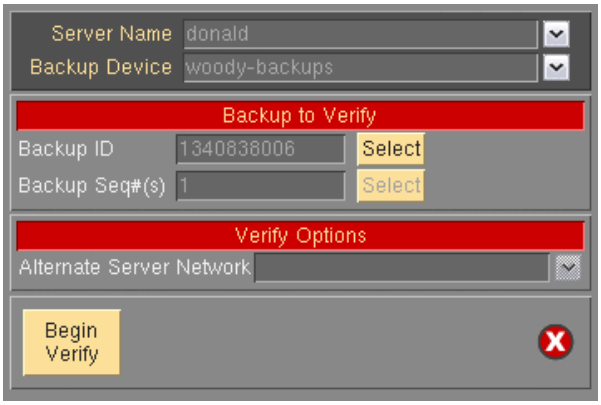

2. If using a *Network Edition* license, select the **Server Name**. When using *Network Edition*, you can also choose "**local (client tape/disk)**" for the **Server**. By selecting this option, you indicate that you want to read the backup information from a disk (directory) or a tape drive attached to a client system, rather than a device configured on a server.

3. If you selected "**local (client tape/disk)**" for the server, a new **Client Name** field will appear, where you must select the client where the backup exists.

   or

4. If you selected a Backup Media Server, select the **Backup Device**.

   or

5. If you selected a TSM Server, select the **Original Client (owner)**. In this case, any client that was part of the backup job to be verified will be sufficient, since we need to list the Backup IDs which the client is part of.

6. Press the **Select** button next to the **Backup ID** field to read the backup media for a list of backup labels.

7. If the selected device is tape, the tape will be automatically read and the backup label ID will be displayed **Backup ID** field. Press the **Begin Verify** button to continue.

8. 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:

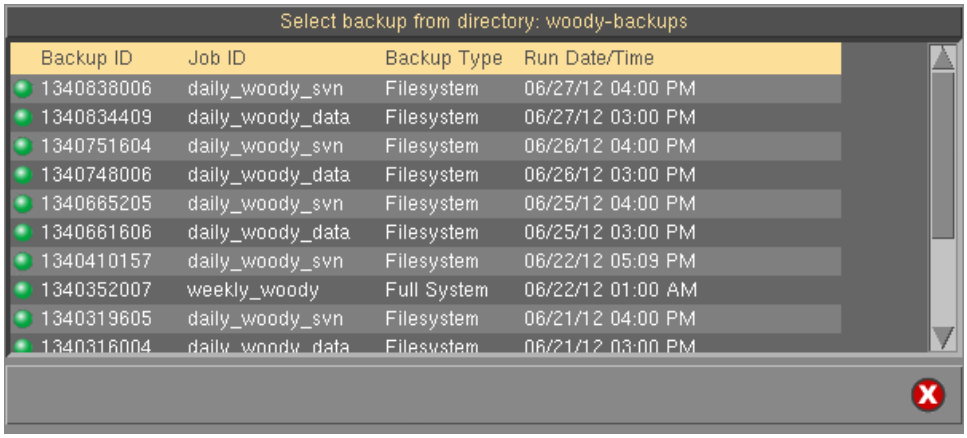

Select the specific backup job to verify by clicking on the button to the left of the desired job.

9. Next, if there are multiple backup on the media, another screen will appear with a list of backups to select from. 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:

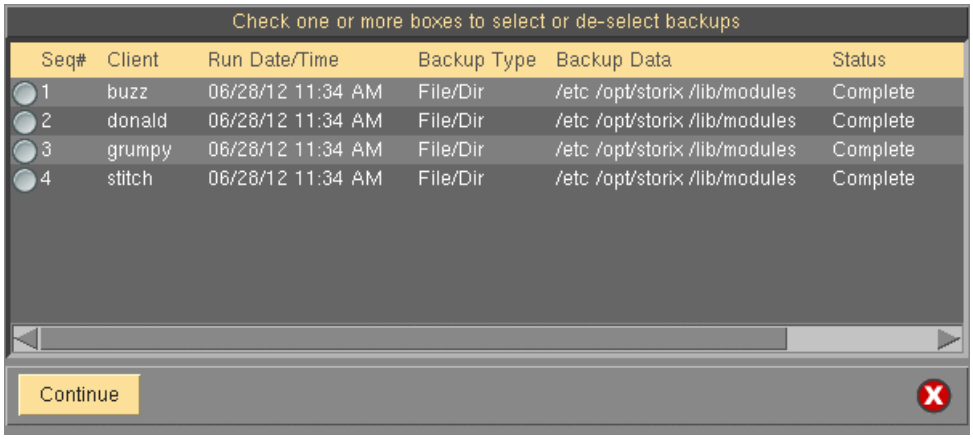

You may select any one or more backups to verify by clicking on the button to the left of the desired selection and a checkmark will appear in the button. If you wish to de-select an option, simply click the button again and the checkmark will disappear. When all selections have been made, click the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

The selected **Backup ID** and **Backup Seq#(s)** will appear on the ***Verify Data from a Backup*** screen. From this screen, press the **Begin Verify** button to start reading the backups.

Using an Alternate Network to Verify from the Server

When using *Network Edition*, it may at time be desirable to have the client read 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 verify using the alternate network.

If an alternate network IP Address or hostname was defined for the server you are restoring from, an addition option will be available on the verify options screen above, "**Alternate Server Network**". If you want to use the alternate network to perform the verify process, select a server hostname or IP address from the list. 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 Adding Alternate Networks in the server configuration.

Displaying the Status and Output of the Verify

When the verify process will begins, 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.

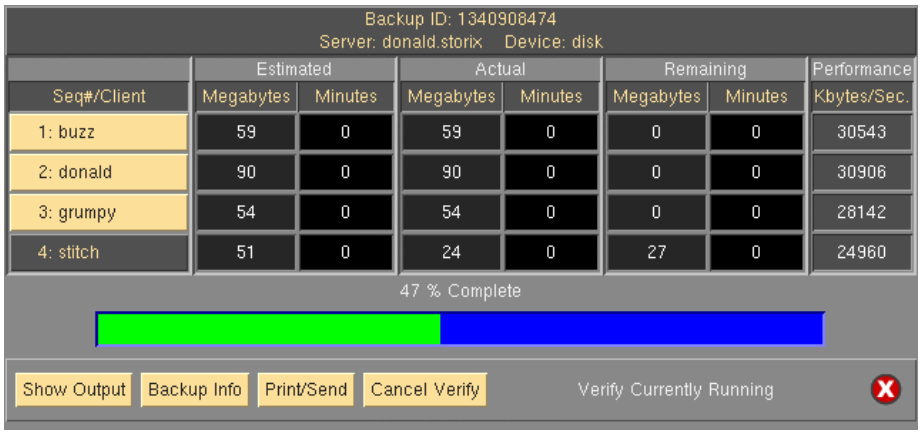

Note that this screen may not be closed as long as the verify process is running. It will remain on the screen after the verify process completes until it is closed by the user. Once the screen is closed, the verify status and output messages may not be redisplayed.

To view the output of the verification of a particular backup, first select the **Client** button or **Seq#** button (if not using *Network Edition*), 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:

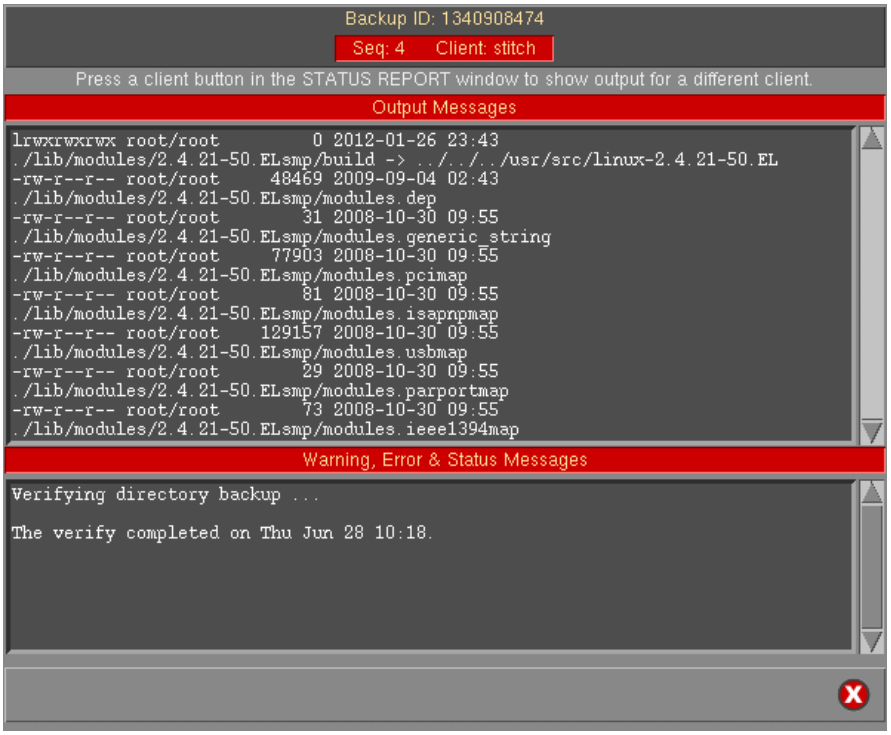

If the verify is of any backup type other than a raw device backup (such as a logical volume or meta-disk), a list of files on the backup will be displayed as each file is read. This screen may be closed and redisplayed at any time, even after the verify completes, as long as the Verify Status Report screen has not been closed.

In addition to the job output, summary information for the selected backup may be displayed by selecting the **Client** button or **Seq#** button (if not using *Network Edition*), then pressing the **Backup Info** button. A screen similar to the following example will appear:

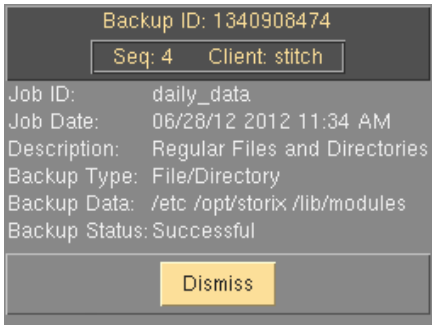

Simply press the **Dismiss** button to close this window.

20. 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* and *Solaris* configurations, allowing creation of devices onto other devices (i.e. meta-disks on logical volumes or ZFS volumes on slices, etc), this feature is not available for Linux or Solaris 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. 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 exists. 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 Xwindows installed. If Xwindows cannot be found on the client, an appropriate message will be displayed and you may not continue. You must either install Xwindows on the client or rebuild the volume groups, logical volumes or filesystems manually on the client. This option is not available using the Web Interface.

Recreate Volume Groups

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

To recreate a volume group, perform the following steps:

1. Select Actions➔Recreate Volume Groups from the menu bar. A screen similar to the following will display:

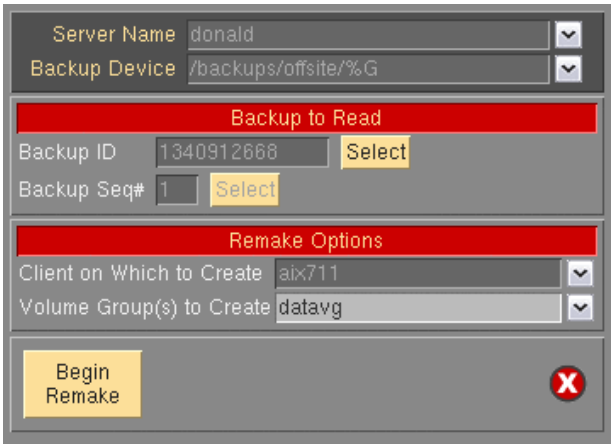

1. If using *Network Edition* license, select the **Server Name**.

2. If using a *Network Edition* license, select the **Server Name**. When using *Network Edition*, you can also choose "**local (client tape/disk)**" for the **Server**. By selecting this option, you indicate that you want to read the backup information from a disk (directory) or a tape drive attached to a client system, rather than a device configured on a server.

3. If you selected a Backup Media Server, select the **Backup Device**.

   or

4. If you selected "**local (client tape/disk)**" for the server, a new **Client Name** field will appear, where you must select the client where the backup exists.

   or

5. If you selected a TSM Server, select the **Original Client (owner**) of the backup.

6. If the selected device is tape, the tape will be automatically read and the backup label ID will be displayed **Backup ID** field.

7. If you selected to recreate from 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:

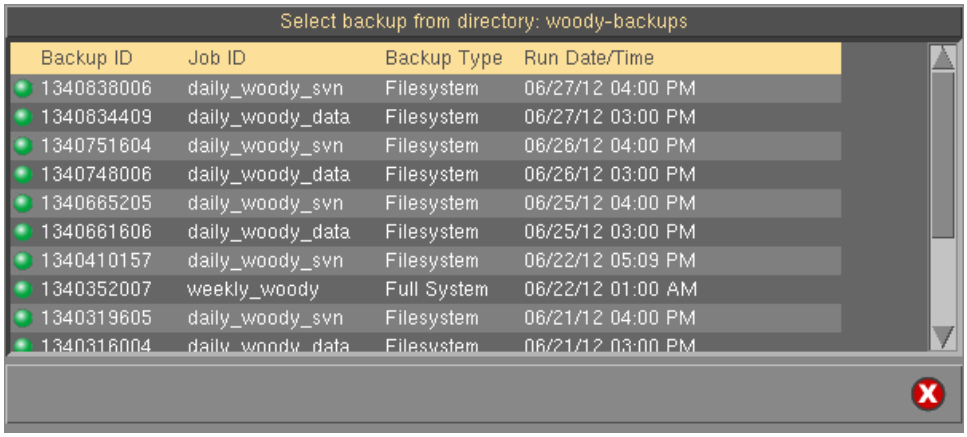

   Select the specific backup job to read by clicking on the button to the left of the desired job.

8. Next, if there are multiple backup on the media, another screen will appear with a list of backups to select from. 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:

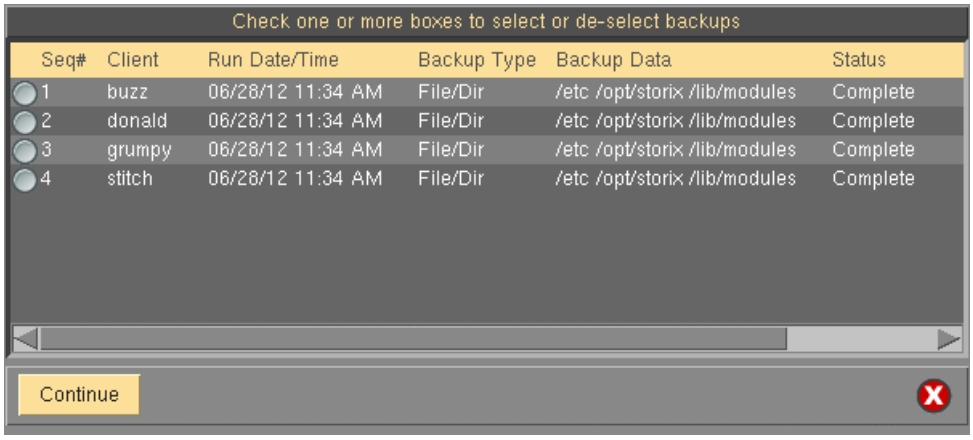

You may select the backup from which to recreate the VG by clicking on the button to the left of the desired selection and a checkmark will appear in the button. Only one selection may be made. If you select a different option, the checkmark will be removed from the previous selection. After making your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

9. You will be returned to the previous window, where you must select each of the following:

   Client on which to create - If using *Network Edition*, this 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.

   Volume Group(s) to create – Use this option 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 this field. You must select at least one volume group to continue.

10. When all selections have been made, press the **Begin Remake** 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:

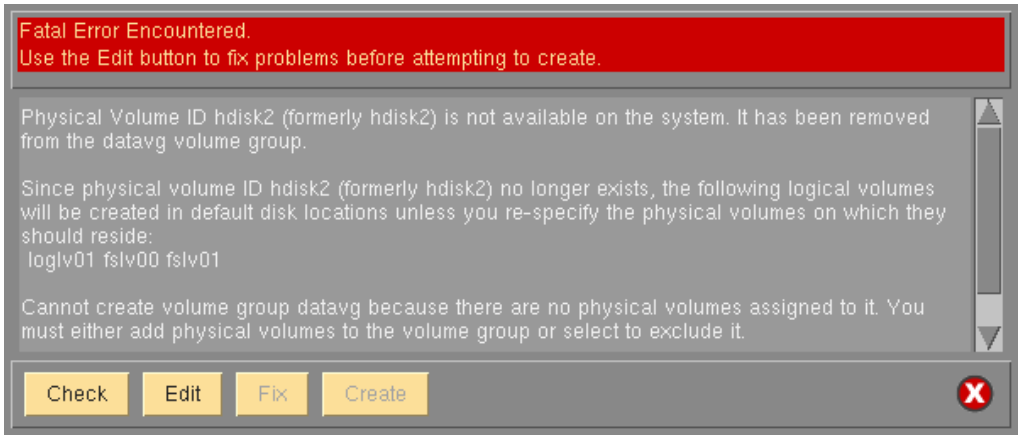

If there are changes required to make the selected volume group fit onto the current system, the **Edit** and **Fix** buttons will become available. If there are no problems found, the **Create** button will be available.

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:

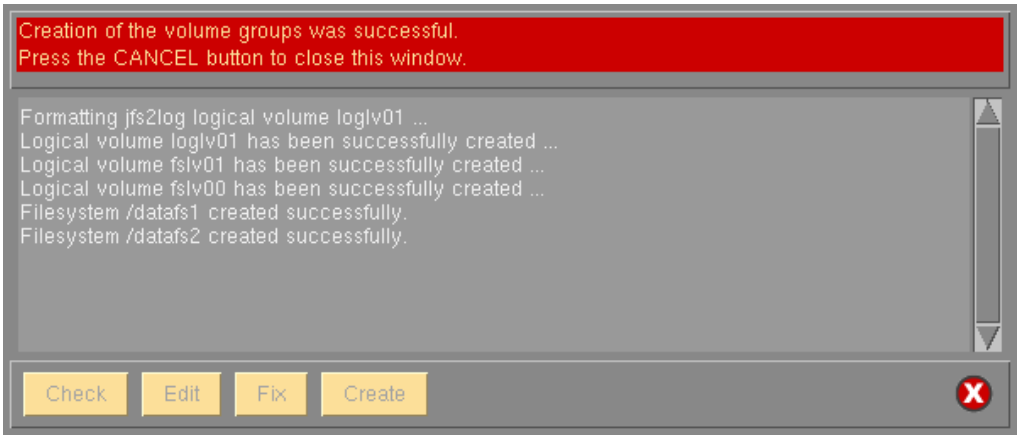

Recreate Logical Volumes or Filesystems

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. A screen similar to the following will display:

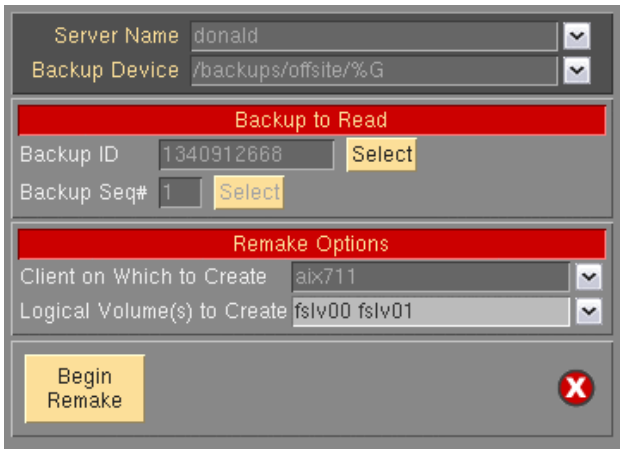

8. If using a *Network Edition* license, select the **Server Name**. When using *Network Edition*, you can also choose "**local (client tape/disk)**" for the **Server**. By selecting this option, you indicate that you want to read the backup information from a disk (directory) or a tape drive attached to a client system, rather than a device configured on a server.

9. If you selected a Backup Media Server, select the **Backup Device**.

   or

10. If you selected "**local (client tape/disk)**" for the server, a new **Client Name** field will appear, where you must select the client where the backup exists.

    or

11. If you selected a TSM Server, select the **Original Client (owner)** of the backup.

12. If the selected device is tape, the tape will be automatically read and the backup label ID will be displayed **Backup ID** field.

13. If you selected to recreate from 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:

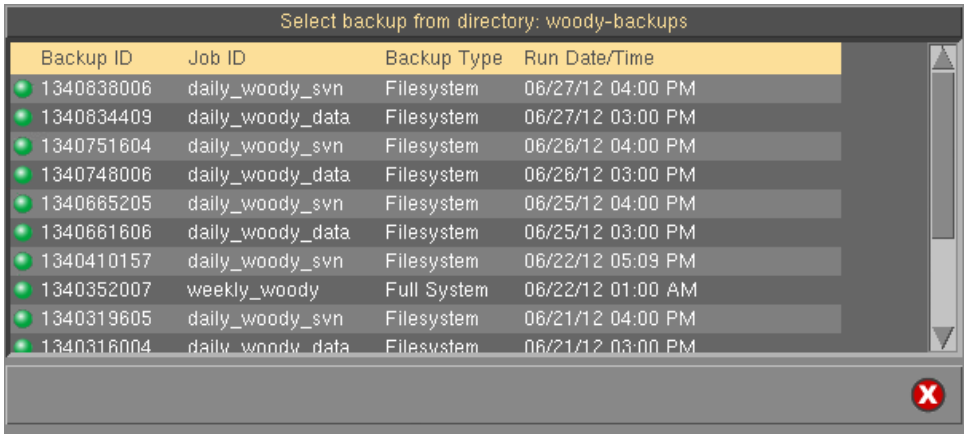

14. Select the specific backup job to read by clicking on the button to the left of the desired job.

15. Next, if there are multiple backup on the media, another screen will appear with a list of backups to select from. 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:

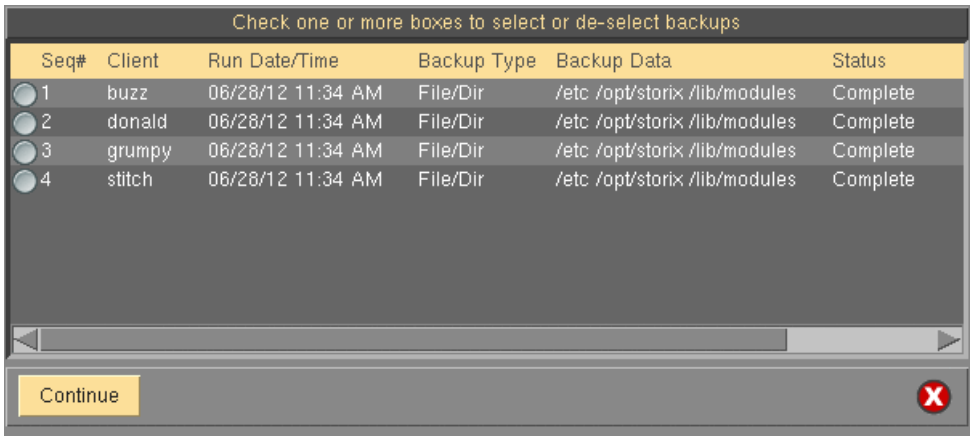

You may select the backup from which to recreate the LV or filesystem by clicking on the button to the left of the desired selection and a checkmark will appear in the button. Only one selection may be made. If you select a different option, the checkmark will be removed from the previous selection.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

10. You will be returned to the previous window, where you must select each of the following:

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

    Logical Volume(s) to create - To list and select the logical volume(s) (and filesystems) that are defined on the backup and select one or more to create from the list, select the arrow next to this field. You must select at least one logical volume to continue.

11. 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:

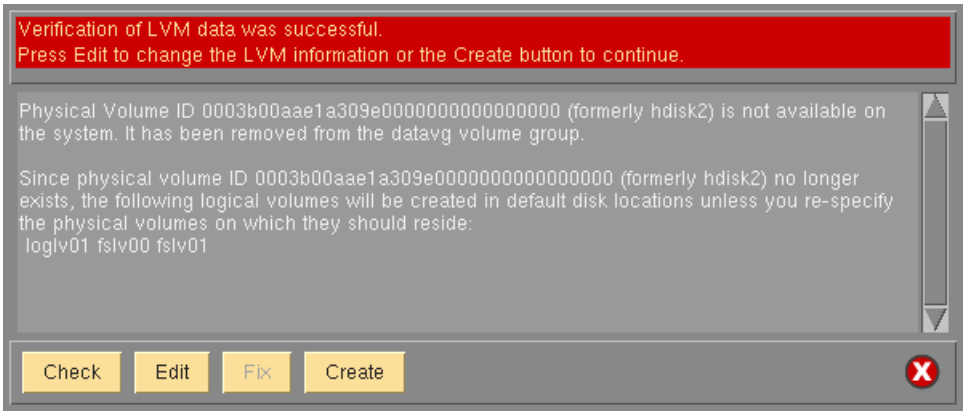

If there are changes required to make the selected volume group fit onto the current system, the **Edit** and **Fix** buttons will become available. If there are no problems found, the **Create** button will be available.

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:

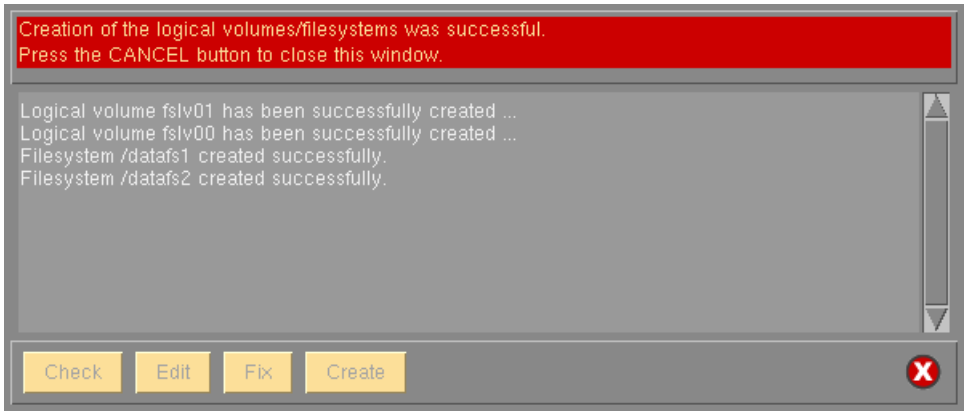

21. Restore Data from a Backup

Data may be restored from any backup server to any client using SBAdmin. Using a *Network Edition* 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* or *Solaris ZFS pools*, each of which may contain *raw 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. A window similar to the following will appear:

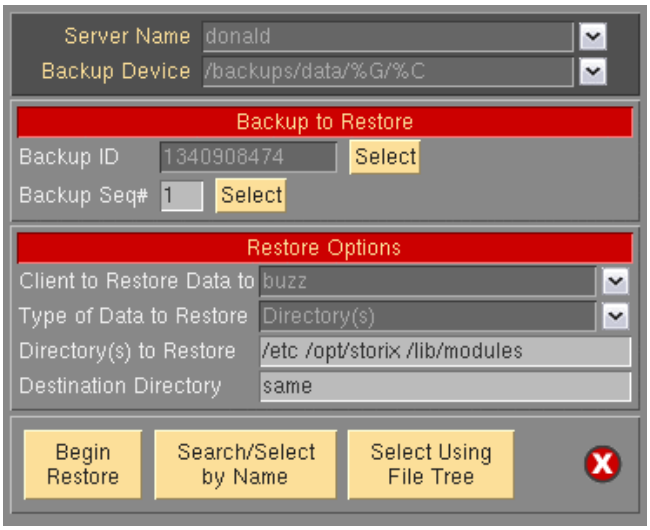

2. If using a *Network Edition* license, select the **Server Name**. When using *Network Edition*, you can also choose "**local (client tape/disk)**" for the **Server**. By selecting this option, you indicate that you want to read the backup information from a disk (directory) or a tape drive attached to a client system, rather than a device configured on a server.

3. If you selected a Backup Media Server, select the **Backup Device**.

   or

4. If you selected "**local (client tape/disk)**" for the server, a new **Client Name** field will appear, where you must select the client where the backup exists.

   or

5. If you selected a TSM Server, select the **Original Client (owner)** of the backup.

6. If the selected device is tape, the tape will be automatically read and the backup label ID will be displayed **Backup ID** field. Press the **Begin Restore** button to continue.

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

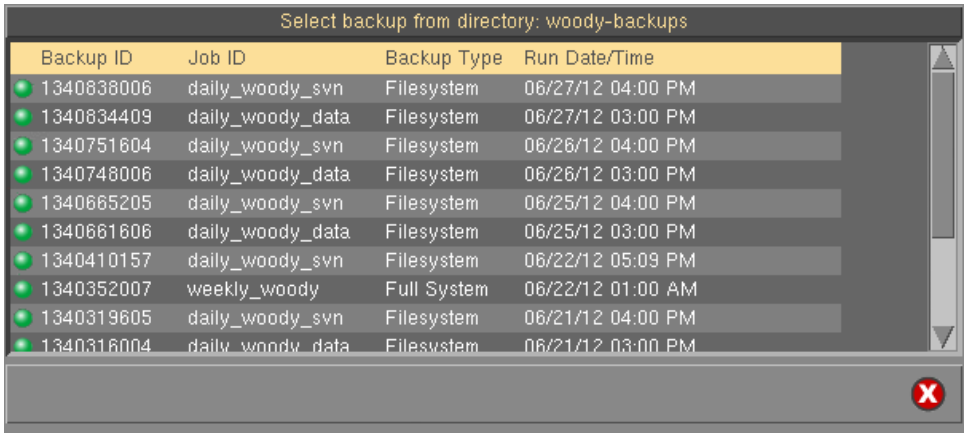

Select the specific backup job to restore from by clicking on the button to the left of the desired job.

8. Next, if there are multiple backup on the media, another screen will appear with a list of backups to select from. 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:

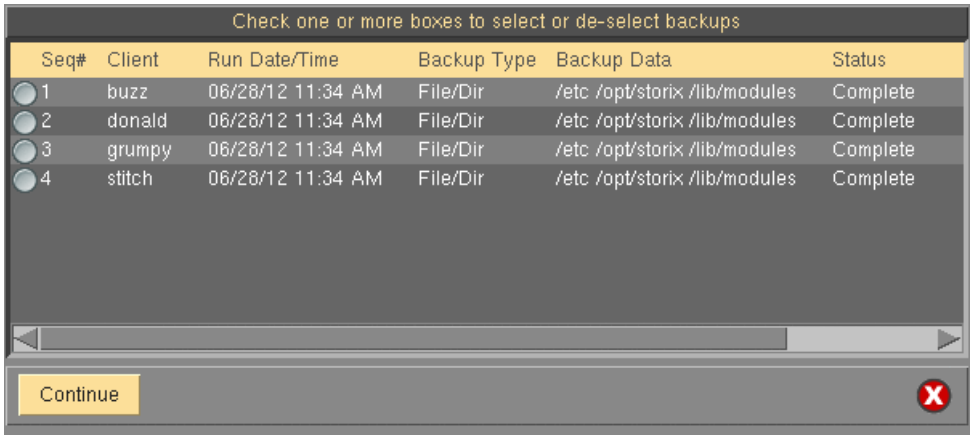

9. You may select only one backup to restore from by clicking on the button to the left of the desired selection and a checkmark will appear in the button. If you wish to de-select an option, simply click the button again and the checkmark will disappear. When all selections have been made, click the **Continue** button at the bottom of the screen.

   The **Backup ID** and **Backup Seq#** selected will be displayed in the fields in the previous window.

Selecting Restore Options

After selecting the backup media to restore from, you will be returned to the previous window, where you may enter or select the following restore options:

The following options may differ depending on the type of backup to restore from and the type of data to be restored.

1. ***Client to restore data to***: This option does not appear on *Workstation Edition* 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, and new buttons will appear at the bottom of the screen from which you may select items to restore from a file tree or search the list for specific patterns. Refer to Selecting Data to Restore below 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 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. ***Alternate Server network***: This option appears only on *Network Edition* and only if the selected backup server has one or more Alternate Networks defined. You may select 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 (Linux/AIX) Zpools (Solaris) Filesystems Directories Regular Files Logical Volumes ZFS Volumes Meta-disks (Linux) Partitions (Linux)	Volume Group Zpools Filesystem, Directory Directory Directory Logical Volume XFS Volume Meta-disk Partition
Volume Group (Linux/AIX)	Volume Groups Filesystems Directories Regular Files Logical Volumes	Volume Group Filesystem, Directory Directory Directory Logical Volume
Zpool (Solaris)	Zpool Filesystems Directories Regular Files ZFS Volume	Zpool Filesystem, Directory Directory Directory ZFS Volume
Filesystem	Filesystems Directories Regular Files	Filesystem, Directory Directory Directory
Directory	Directories Regular Files	Directory Directory
Logical Volume (Linux/AIX)	Logical Volume	Logical Volume
Meta-disk (Linux/Solaris)	Meta-disk (Linux or Solaris)	Meta-disk (Linux or Solaris)
Partition (Linux)	Partition (Linux)	Partition (Linux)
Slice (Solaris)	Slice	Slice
ZFS Volume (Solaris)	ZFS Volume	ZFS Volume

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 *Volume Groups, Filesystems, Logical Volumes*, *Meta-disks* (Linux/Solaris), *Partitions* (Linux), Slices (Solaris) 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 using the graphical interface:

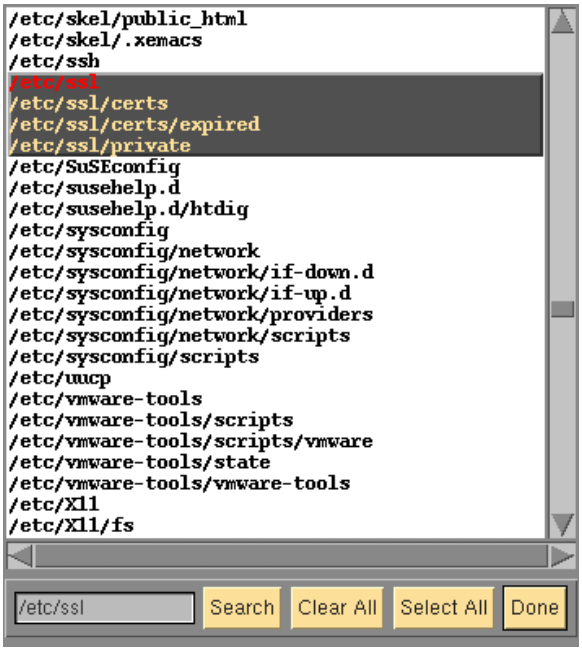

From this window, you may:

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. A 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.

When using the web interface it is recommended to use a search term to limit the **Search/Select by Name** button. Due to limitations with web browsers, SBAdmin will only display up to 1000 files in the search result. Listing more files could cause the browser to crash.

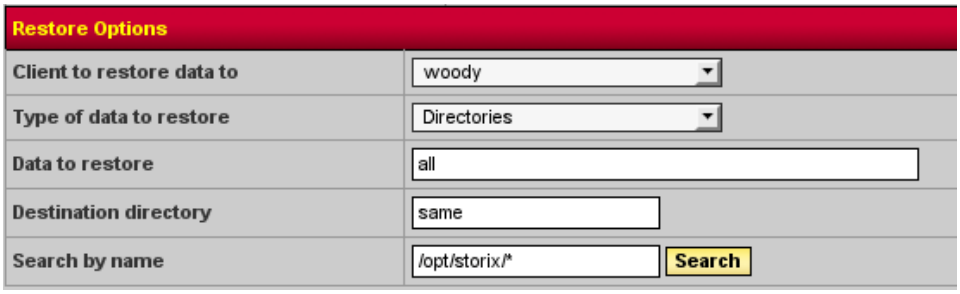

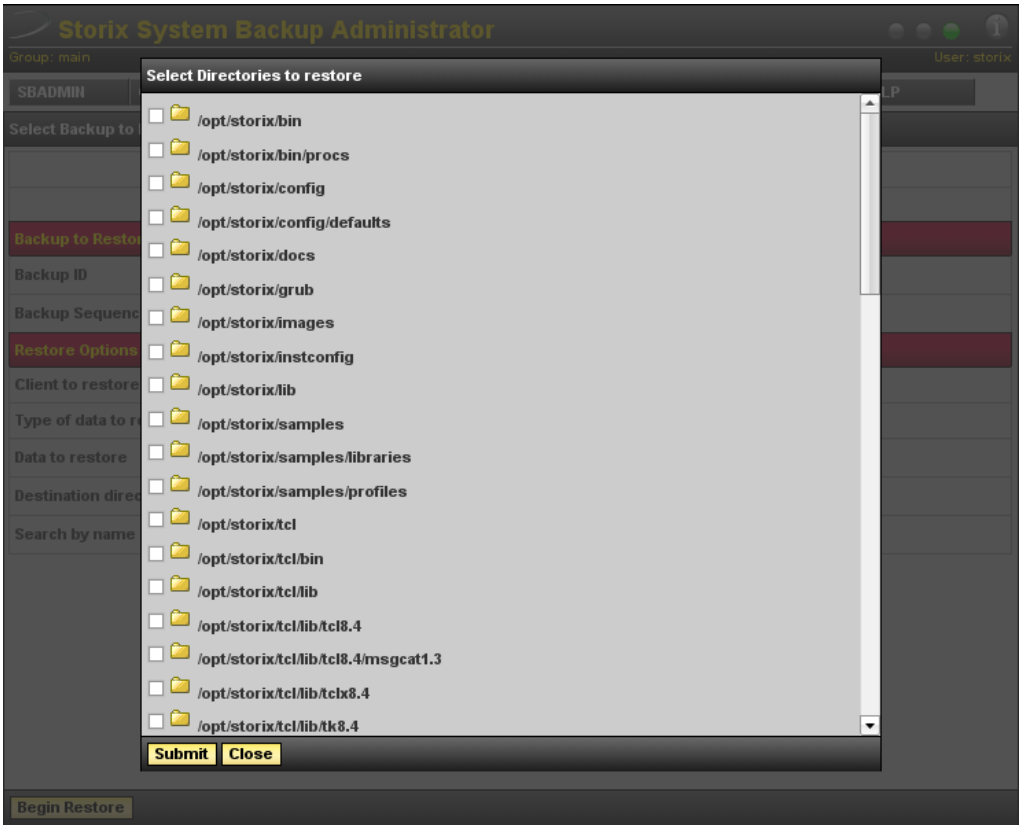

From this popup window you may select one or more files to restore. After selecting the **Submit** button you will be returned to the main screen and can begin the restore.

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).

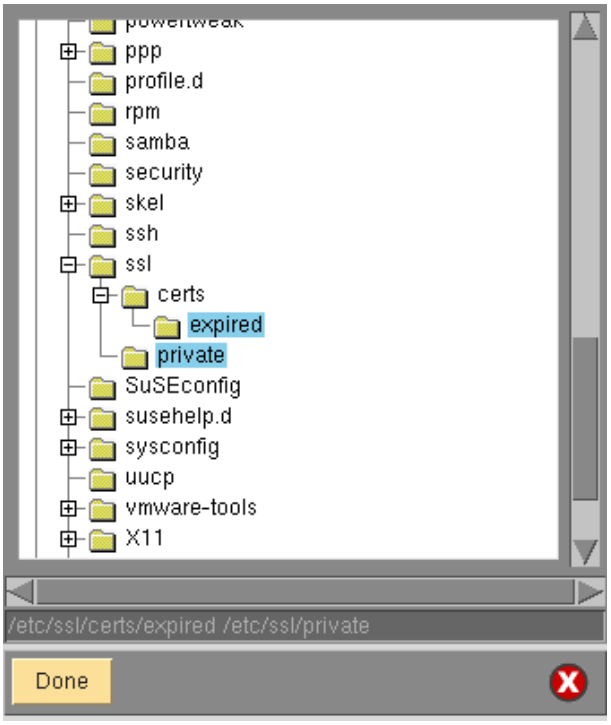

From this window, you may:

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 that contain similar names that you want to restore without having to select each and every file. 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
```

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

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

but this will NOT result in files such as:

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

These are not found 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
```

This 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** 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, ZFS Pool* 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, ZFS Pool* 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 *Network Edition*, it may at times 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.

If an alternate network IP Address or hostname was defined for the server you are restoring from, an addition option will be available on the restore options screen above, "**Alternate Server Network**". If you want to use the alternate network to perform the restore, select a server hostname or IP address from the list. 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 Adding Alternate Networks in 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.

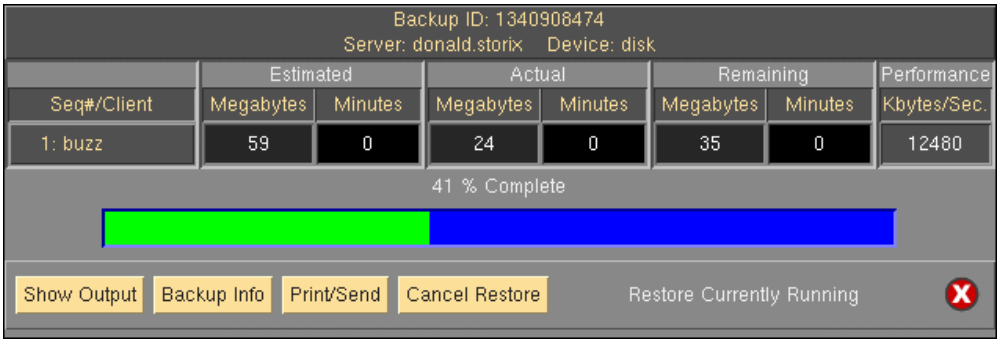

Note that this screen may not be closed as long as the restore is running. It must remain on the screen after the restore completes, after which time it may be closed by pressing the cancel button. Once the screen is closed, the restore status and output messages may not be redisplayed.

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:

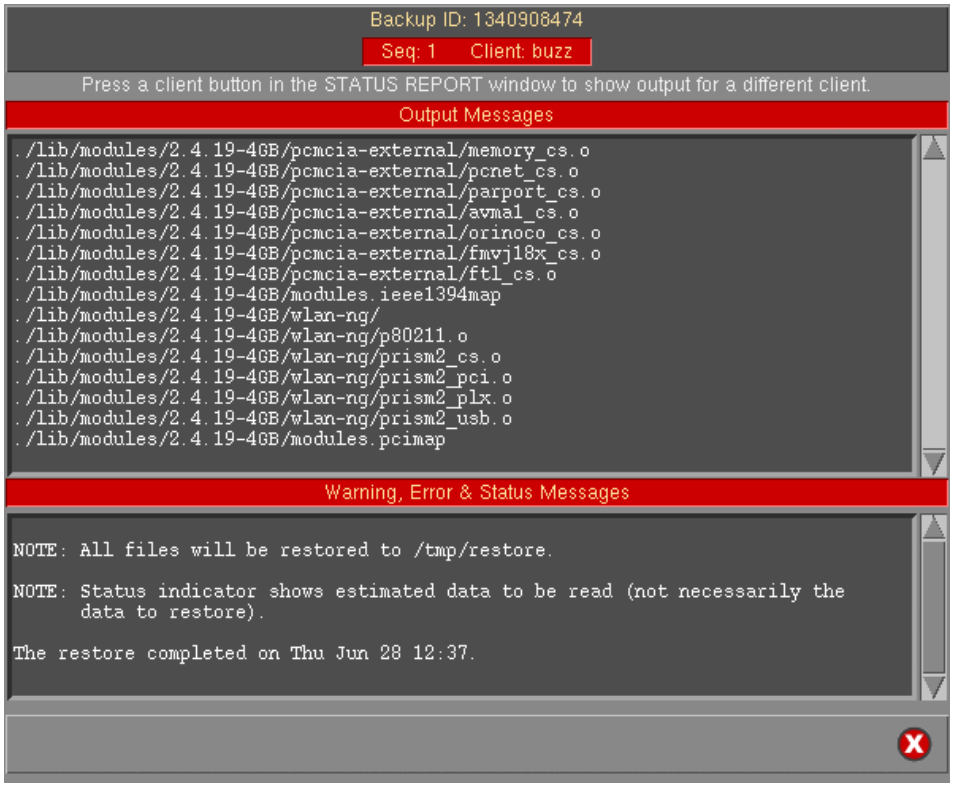

If the restore is of any backup type containing filesystem data, the files will be listed on the screen as they are restored. For **Logical Volume (*AIX/Linux*)**, **Partition** (***Linux***), **Meta-disk** (***Linux/Solaris***), **Slices** or **ZFS Volumes** (Solaris), only one message is displayed 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.

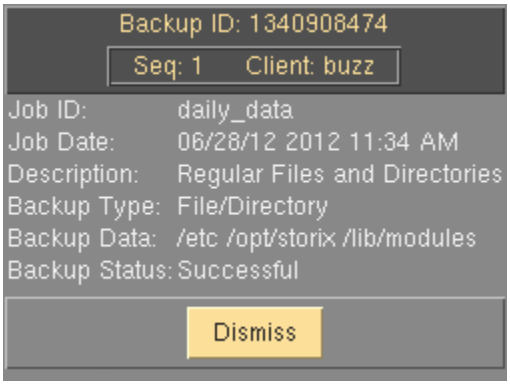

Simply press the **Dismiss** button to close this window.

22. Copying Backups to Different Media

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 take 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:

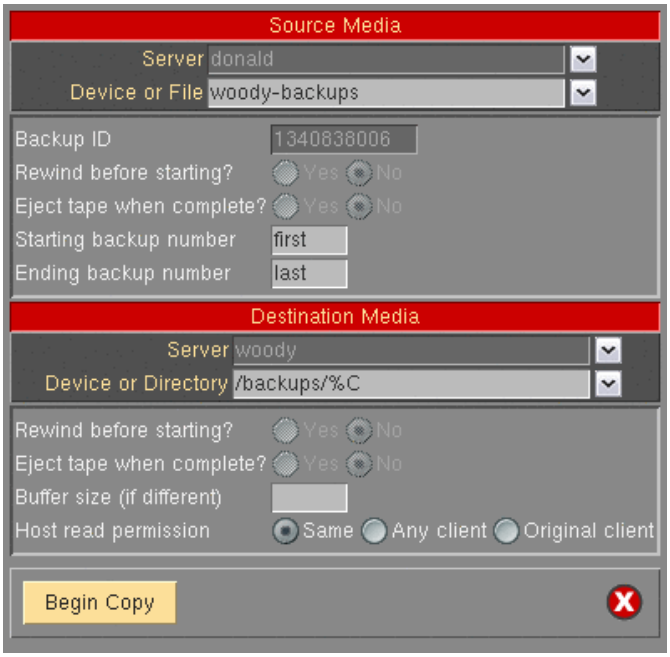

This screen is broken into two sections, one for the source backup and one for the destination media. If using *Network Edition*, you must specify a source and destination server. Otherwise, these fields do not appear.

Source Media

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

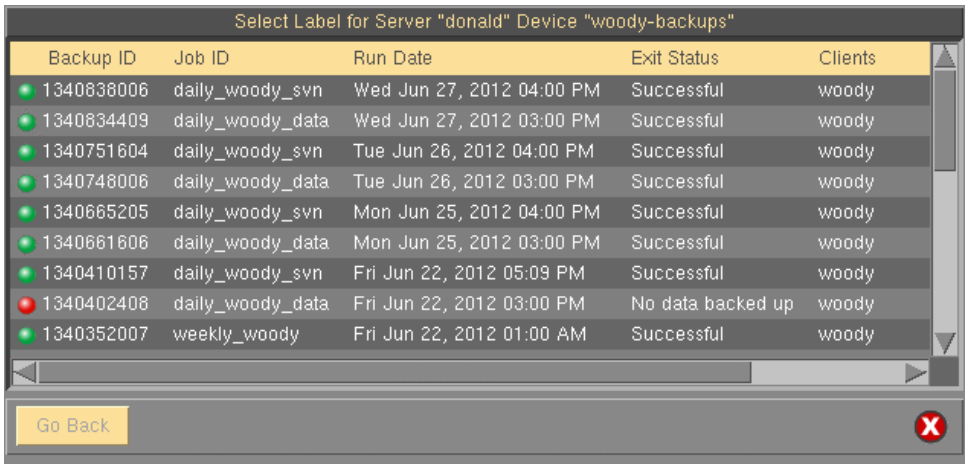

If the backup contained multiple backup sequence numbers, you may select the starting and ending backup sequence numbers to copy in the **Starting backup number** and **Ending backup number** fields. This is useful if, for instance, you created a backup of multiple clients but want to copy only one client backup in the list to tape. Another example 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.

> **You will not be able to use a device configured as a random tape library as a source device. This is because SBAdmin 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 a tape device not configured as a library 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 a tape device to a directory device, from one tape device to another, from directory to tape device, or from one directory device to another. If using *Network Edition*, the selected backup may be copied from any server to any other server.

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 end of the last backup written to it). 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 128K buffer size may be adequate when you wrote your original disk backup file, but when copying to a high-speed tape drive or disk, a higher buffer size (i.e. 256K to 1024K) 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. Refer to Buffer Size in the profile settings for more information.

If using *Network Edition*, 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 using *Workstation Edition*, 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:

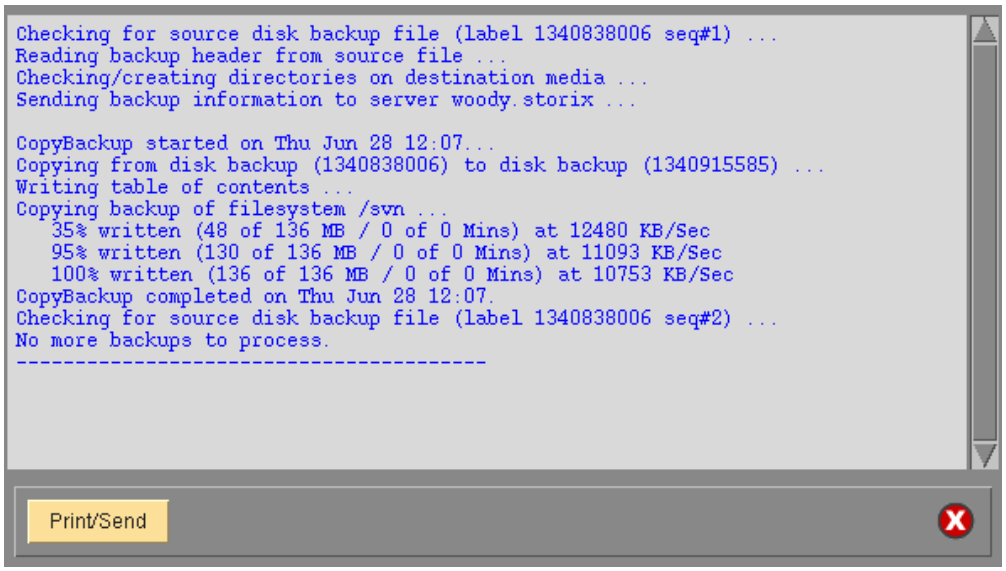

Canceling the Operation

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.

23. 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:

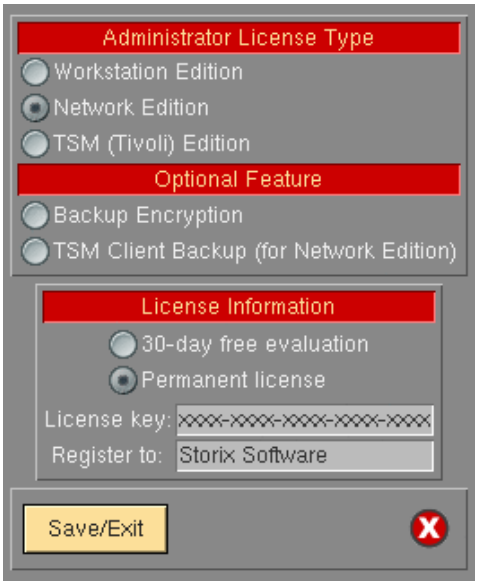

When using the web interface:

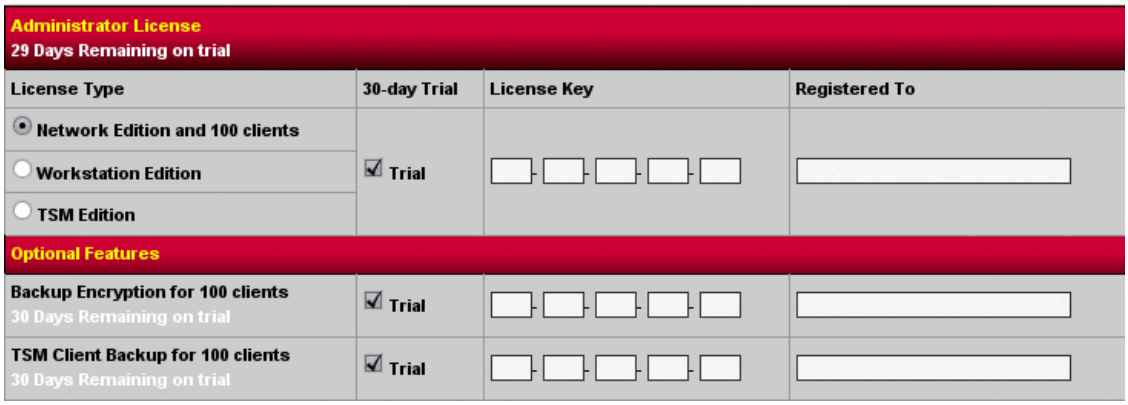

It may be desirable to change the license for a number of reasons:

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 install a permanent (purchased) license key.

3. You purchased a Network Edition for 10 clients, and now want to add support for another 6 clients.

4. You wish to install a new optional feature.

Administrator License

When the software is initially installed, you indicated the type of administrator you would be installing. This was ***Workstation***, ***Network,* or *TSM Edition***. You also entered a license key, which matched the administrator license type, and also indicated the number of clients supported by a Network or TSM Edition 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. **TSM Client Backup Feature**: This optional feature may be added to *Network Edition* to allow any SBAdmin backup to be performed to a TSM server. This feature will support any number of TSM servers. You must also configure each of your clients with TSM node information, but you may only configure the number of TSM nodes up to the number of clients the license supports.

   Note that this is not the same as ***TSM Edition*** (administrator license). The TSM Edition provides only TSM client and server options and only System Backup options. Adding ***TSM Support*** to the Network Edition provides all of the options of both ***Network Edition*** and a ***TSM Edition***.

When entering this function, your current *Administrator License* information is displayed. In the **License Type** section, you may select a different *Administrator License* Type (***Workstation, Network*** or ***TSM***), or you can select an *optional feature* (Encryption, TSM). When you select a different option in the **License Type** section, the current license information for that feature is displayed.

Any Administrator license or optional feature may be installed for a 30-day, one-time, trial period. 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 program (***sbadmin***) must be terminated and restarted for the changes to appear.

General Preferences

By selecting File→Preference→General Preferences, a screen similar to the following will be displayed, from which you may set or change numerous attributes defining the behavior of the SBAdmin application:

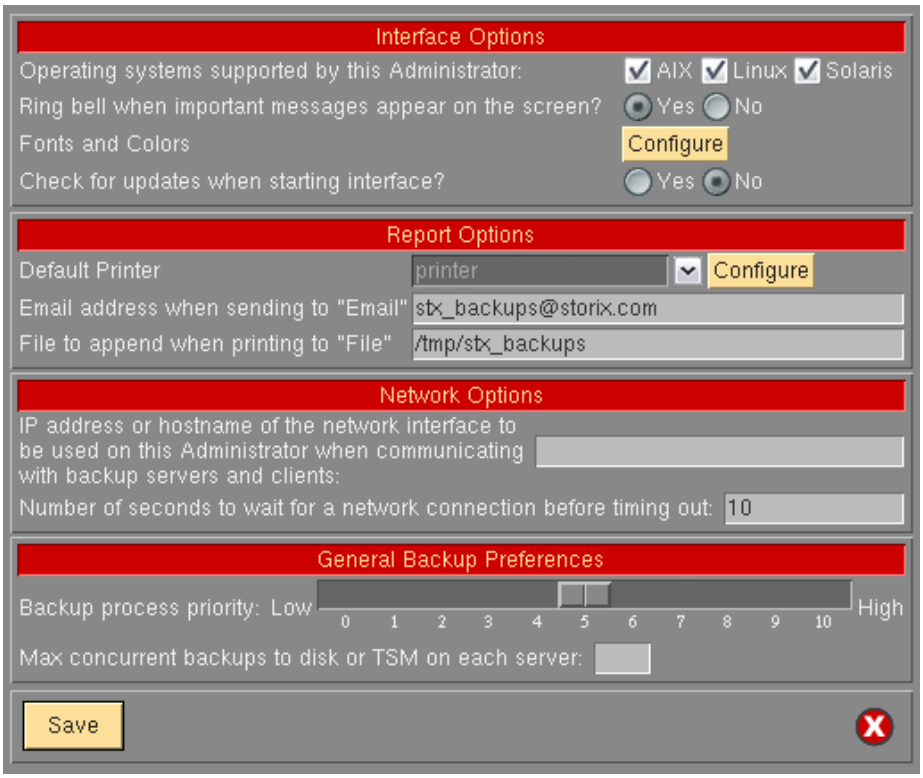

The **Interface Options** define the features and appearance of the SBAdmin GUI interface:

Operating Systems Support

> **This option will not appear if using Workstation Edition.**

By default, when the Backup Administrator software is first installed, only support for the operating system running on 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*** or ***Solaris*** clients, then you will want to add into the application those options that are applicable to Linux or Solaris systems also. You may later turn off support for client operating systems that will not be managed by the *Network Administrator*.

For example, only ***Solaris*** systems support *ZFS*. Therefore, by turning off Solaris support, no ZFS options (i.e. Zpool or ZFS Volume backups) will appear anywhere in the interface.

To enable or disable support for a particular operating system, simply select or de-select the button next to the corresponding operating system type. When finished, press the **Save** button.

Sound On/Off

The option “Ring bell when important messages appear on the screen” allows 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. You may then select **Yes** or **No** indicating whether or not the bell should ring.

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 File→Preference→General Preferences, press the **Configure** Button next to **Fonts & Colors**. The following screen will appear:

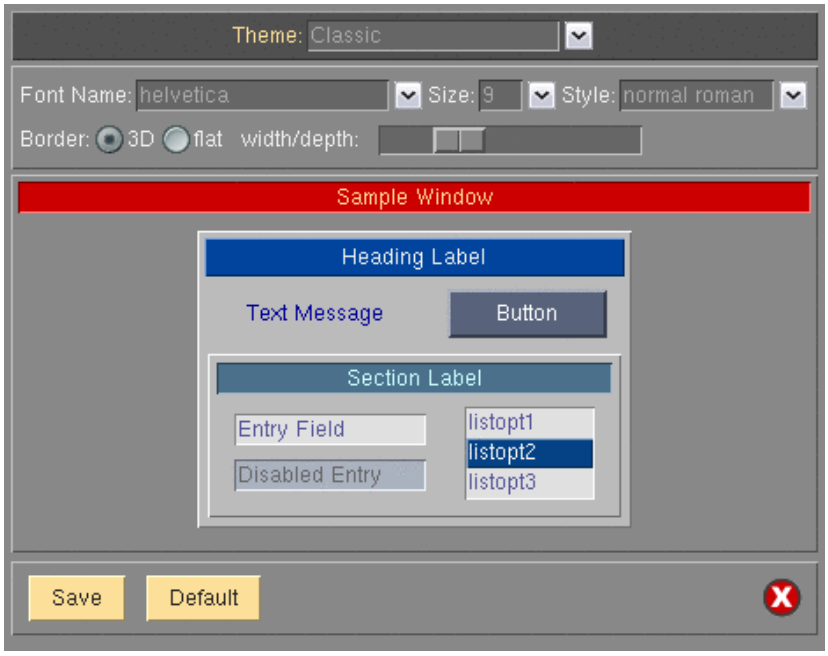

This screen will always be displayed using the ***Storix*** color theme, even when another theme has been applied to other screens. To display a different color theme in the **Sample** section of the display, select one of the following from the **Theme** drop-down list:

Aquamarine

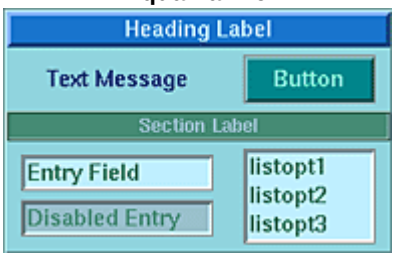

Carnival

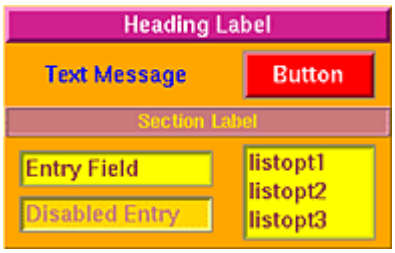

GolfCourse

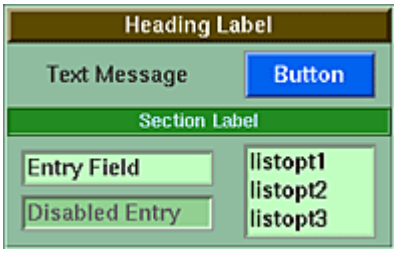

MardiGras

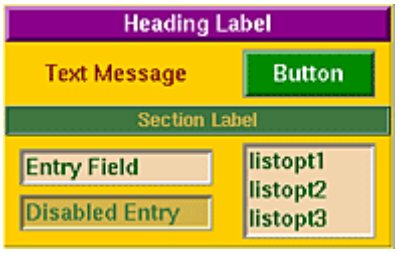

PowderPuff

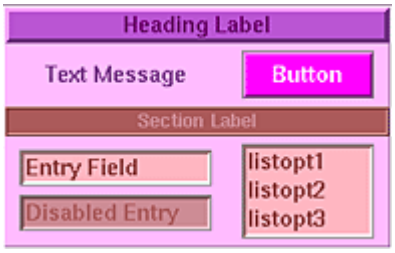

SantaFe

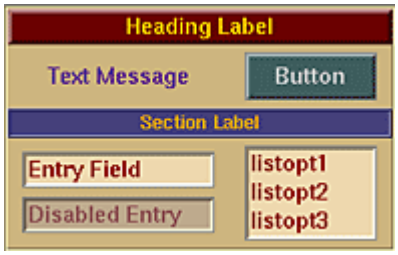

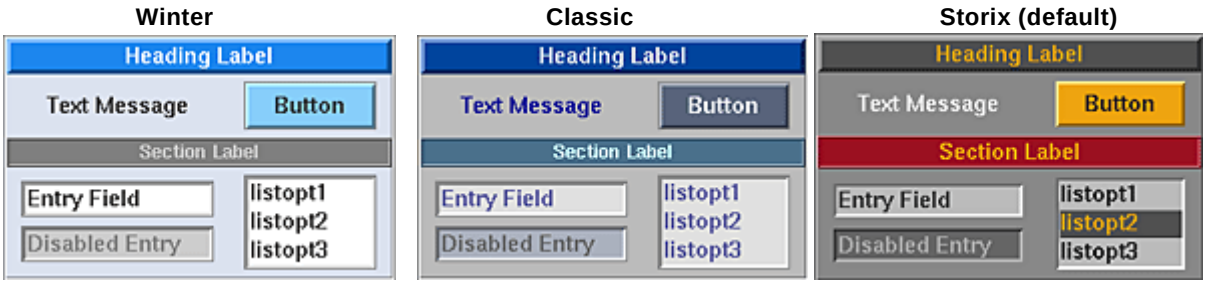

To change the **Font Name**, **Size** or **Style**, click the arrow button next to the desired selection, then select an option from the list. The sample box will be changed to show your selections.

To change the **3D Effect** and the **Width/Size of Borders**, select the option or adjust the values to the desired settings. 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 (*Classic*) font and colors, press the **Default** button or 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.

Check for Updates

When launching the Administrator, the application will contact the Storix website in order to compare the currently running version of SBAdmin with the latest version available. If you are not running the most current version, a note will be displayed at the bottom of the interface. You may prevent this communication from occurring by selecting "*No*" to this option.

Network Options

Network Interface

By default, the *admin system* will use the network adapter associated with the default *hostname* of the system when communicating with clients. If the system has multiple network adapters, you may choose a different network, which will only be used for communicating between this admin system and the clients. To do so, select File→Preferences→General Preferences from the menu bar, go to the **Network Options** section, and enter the IP address or hostname in the "**IP address or hostname of the network interface**…" field. If you want the system to go back to using the default adapter (according to the primary hostname of the system), simply remove the entry from this field.

The network adapter selected will be used to pass information between the admin system and clients, 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 TSM servers.

Important! After setting this value, some or all of the clients servers may show as unavailable when the Clients, Servers and Devices are displayed on the Main Screen. If this should occur, it means that the client does not have the admin system defined using the alternate network adapter. To resolve this problem, edit the ***/storix*/config/system_admin_hosts** file on the client (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

For SBAdmin to perform any operations on a client, 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 within 10 seconds, it is assumed that the client 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→General Preferences from the menu bar, then enter the new number of seconds in the in the field labeled "**Number of seconds to wait for a network connection before timing out**".

When increasing the value, it will take longer to determine that a system is unavailable, so some processes used by the interface (most notably when displaying client, servers and devices from the Main Screen) may take longer to update if a client cannot be contacted. It is not advisable to increase this value to more than 30 seconds, depending on the total number of clients configured.

> **Changing this option will affect only the timeout value when the Administrator contacts the clients. It has no affect on the timeout TSM uses for the clients to contact the TSM 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.**

Report Options

Reports and backup 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.

To edit configure or change these options, select File→Preference→General Preferences, and add or change the following options in the **Report Options** section.

Default Printer (*AIX*)

It is assumed that AIX systems will already have AIX printer *queues* 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/Solaris*)

Linux and Solaris 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 or Solaris 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 select the default printer used for reports and notifications, select from the pull down menu next to the **Default Printer** entry field, and select a pre-configured printer from the list. If you've not yet configured the printers within SBAdmin, press **Configure** button, and a screen will appear like the example below:

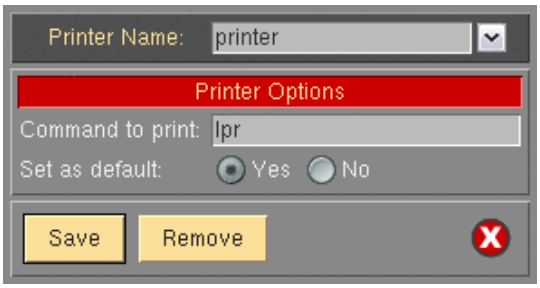

To define a new printer, enter the name of your printer in the box at the top of the screen. To change an existing printer definition, either type its name or select it using the arrow button to the right of the entry field.

If configuring a new printer, the name you enter may be any name you choose, not necessarily the name of the printer queue as defined to Linux or Solaris. 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 **Save** 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. SBAdmin uses either the mail or mailx command on the system to send emails. If you are not receiving emails, verify that the mail or mailx command is located on the Admin system and that you can successfully send emails from the command line.

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 automatically be created. Any parent directories of the file must already exist.

Network Options

These options are not available when using Workstation Edition.

Network Interface

By default, the *admin system* will use the network adapter associated with the default *hostname* of the system when communicating with backup servers and clients. If the system has multiple network adapters, you may choose a different network, which will only be used for communicating between this admin system and the backup servers or clients. To do so, select File→Preferences→General Preferences from the menu bar, go to the **Network Options** section, and enter the IP address or hostname in the "**IP address or hostname of the network interface**…" field. If you want the system to

go back to using the default adapter (according to the primary hostname of the system), simply remove any entry from this field.

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 even if the admin system is the backup server. To configure the network interface used for sending backup data, refer to the alternate hostname 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 Clients, Servers and Devices 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, edit the ***/storix/*config/system_admin_hosts** 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

For SBAdmin to perform any operations on a client or server, from querying its 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→General Preferences from the menu bar, then enter the new number of seconds in the in the field labeled "**Number of seconds to wait for a network connection before timing out**".

When increasing the value, it will take longer to determine that a system is unavailable, so some processes used by the interface (most notably when displaying client, servers and devices from the Main Screen) may take longer to update if a client or server cannot be contacted. It is not advisable to increase this value to more than 30 seconds, depending on the total number of clients and servers configured.

Changing this option will affect only the timeout value when the 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.

Backup Process Priority

You can change the default CPU process priority of all backup jobs run from the admin system by selecting File→Preferences→General Preferences, and moving the slider in the **Backup Process Priority** section to the desired number.

All backups will use the system default priority unless you set a different priority. The process is represented by the operating system using a scale of highest to lowest. You can view a process priority by typing "`ps -ef`" and referring to the NI column. For ***Linux***, the number will be –19 (highest) to 20 (lowest), with a default of 0. For ***AIX*** and ***Solaris***, the number will be from 0 (highest) to 39 (lowest), with a default of 20.

To make this simpler, SBAdmin represents this on a scale of 0 to 10, with 0 being lowest, 5 being normal (default), and 10 being highest. By default, this will bet set to 5, indicating that the normal operating system default priority should be used.

It is common to lower this value if you do not want your backups to run as the same priority as other applications. This will usually cause the backup process to have less affect on the performance of other applications, but the backup could run a bit slower. You might consider setting this value to 3 to give backup processes a lower priority, but not so low that the backup takes too much longer.

If you want the backup process to run in the shortest possible time, with no regard to how it might adversely affect the performance of other applications, you can increase this value. Note that doing so will have little or no affect if there are no other applications or processes requiring CPU processing, but those that require the CPU could run slower. In this case, you might consider a setting of 7 to give priority to the backup without bringing all other applications to a halt.

Setting this option to any value other than 5 will apply to all backup jobs unless you change the backup job process priority within the job profile settings.

Concurrent Backups

For backups to a disk (directory) on a Backup Media Server, multiple backups may be writing to the same device concurrently. For TSM, SBAdmin has no control over where the backup data is actually stored on the TSM server. Therefore, all backups to the TSM server may be run concurrently, even if they are writing to the same tape library or disk drives.

This may pose a problem on the server for a number of reasons:

1. This may cause degraded I/O performance on the server, affecting other applications.

2. A single large backup to a server's disk could cause many smaller backups to take much longer since they must share the same I/O.

3. TSM may simultaneously write multiple backups to the same tape (interspersing the backup data), but all of the simultaneous backups will be limited by the speed of that single tape drive.

It would make more sense in any case to limit the number of concurrent backups to a manageable number. To limit the number of disk (or TSM) backups which may be written concurrently to the same server, select File→Preferences→General Preferences, and enter a number in the **Limit concurrent backups *to [a disk or the TSM server]* to:** field.

By leaving this field blank (or 0), there will be no limit to the number of concurrent backups.

Auto-Terminate Stalled Backups

You may specify the number of minutes to wait before considering a backup as stalled and have the backup job automatically killed. This preference does not apply to backups written to tape devices because the timeout could be related to waiting for a new tape when reaching end of media. The default is to wait forever and not automatically kill the backup job.

A backup is considered to have stalled once all backup pre-processing has completed and backup data begins to write, but no progress has been made for the length of the defined timeout. If the timeout has been reached, notification will be sent based on your notification settings, and a kill signal will be sent to the backup job.

To define stalled backups timeout select File→Preferences→General Preferences and enter the number of minutes in the **Minutes to wait before killing a stalled backup to a non-tape device** field.

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 backup 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 the Backup Retention Policy in the Backup Profile settings for more information on overriding the global backup retention policy.

To change the global backup retention policy, select File→Preferences→Backup Retention Policies from the menu bar. The following screen will appear:

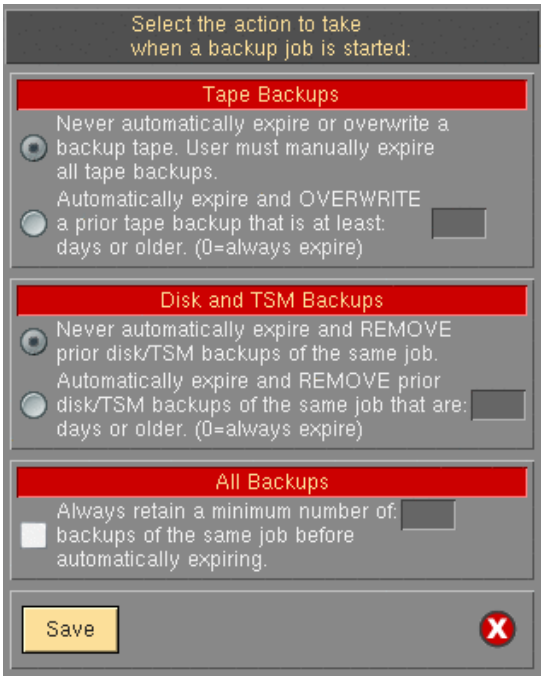

Tape Backups

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. ***Never automatically expire or overwrite a backup tape. User must manually expire all tape backups.***

   Select this option if 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. ***Automatically expire and OVERWRITE a prior tape backup that is at least: __ days or older.***

   Select this option if the backup will proceed ONLY if the backup to be overwritten is at least the specified number of days old. You can ***always allow expiration and overwriting of a prior backup tape*** by setting this value to 0. If the backup currently on the tape 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.

These global retention policies may be overridden for specific jobs by changing the Retain option in the job settings.

Disk or TSM Backups

Backups written to disk (or a *TSM server*) are never overwritten by a new backup since each backup has a different filename (constructed using the backup Label ID, job ID, etc). A new copy of a backup will be written to the server each time a job is re-run. This option is used to free-up space on the server by automatically removing old backups when new backups are started. Two options are available for backups to disk (directories):

1. ***Never automatically expire or remove prior disk backups of the same job.***

   Select this option if no backups on disk should ever be automatically expired when writing new backups. This will require that the user manually expire and remove each prior backup that is no longer needed. The user must take the steps to expire the older backups, else they will remain on disk indefinitely, using storage space which may be needed by newer backups. Refer to Expiring a Backup for the steps to manually expire a backup.

2. ***Automatically expire and REMOVE prior [disk/TSM] backups of the same job that are ___ days or older.***

   By selecting this option, older backups of the same job that is being run will be automatically expired and removed from server if they are over the specified number of days old. If no backups are needed for more than 90 days, selecting this option will ensure that backups over 90 days old are removed and replaced by a newer backup created using the same backup job.

   If you enter “0” in this field, you are indicating that ***all prior backups of the same job will be removed and expired*** before the new backup is created.

These global retention policies may be overridden for specific jobs by changing the Retain option in the job settings.

Number of Backups to Retain

In addition to, or in replacement of, the number of days that a backup must be retained, you may also specify the number of backups that must be retained for each job. To do so, select the checkbox next to the label ***Always retain a minimum number of ___ backups of the same job before automatically expiring***, and enter a number in the corresponding field.

By entering a number in this field, you will ensure that a minimum number of backups are retained for each job before a prior backup may be automatically expired. If, however, you set your **Tape** or **Disk** retention policy to never allow a backup to be automatically expired, this entry will have no affect.

If you have set the **Tape** and **Disk** (or TSM**)** retention options to always expire backups, then this setting alone will ensure that you always retain a minimum number of backups of each job, regardless of how old they are.

When specifying the Number of Backups to Retain, keep in mind that this is the number of backups that will ALWAYS exist on the system. If you specify to keep a minimum of 3 copies, there will be a total of 4 backups after running the backup job. This is because we expire only the 4th backup at the start of the backup to ensure that if the backup were to fail, you would still have your minimum of 3 backups.

Press the **Save** button to save your selections and close this window, or press the **Cancel** button to cancel changes.

Backup Status Notifications

Because scheduled backups may be running even when the *Backup Administrator* 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 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 Notifications on the menu bar. The following screen will appear:

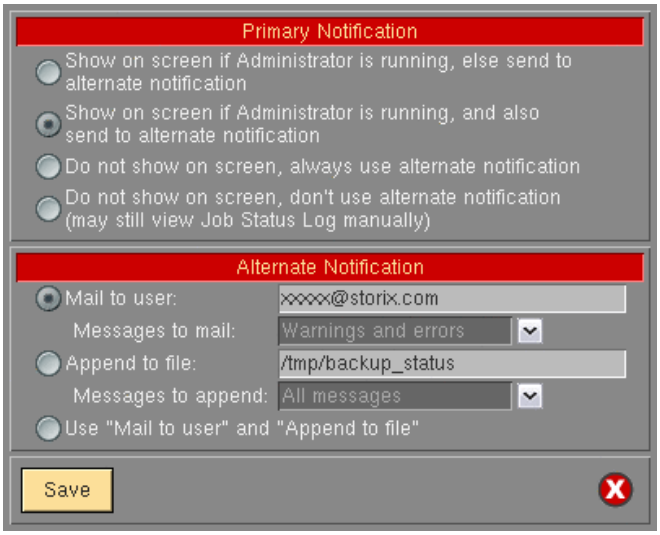

Primary Notification

One of the options in this section must be selected to determine where backup status messages should be reported:

1. ***Show on Screen if Administrator is running, else send to alternate notification.***

   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. ***Show on screen if Administrator is running, and also send to alternate notification***

   If you want messages always reported on the screen (when Administrator is running) and also sent using the alternate notification method, select this option.

3. ***Do not show on screen, always use alternate notification***

   Select this 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.

4. ***Do not show on screen, do not use alternate notification***

   If selected, the messages will not pop-up on the screen and will not be sent to the alternate notification. The messages are still logged for displaying on the screen at a later time. You will be asked if you want to display the messages whenever SBAdmin is started, when it is exited (in which case the log is cleared), and you may also press the **View Log** button on the main screen to display the log at any time.

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.

1. Mail to user

By default, messages will be sent to the *roo*t user's mail when the alternate notification is used. If you want a different user to receive the mail messages, select this option, 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"). You may want to limit the types of messages that are emailed to the user. To do so, select the arrow-button next to the Messages to mail field, and select either **All messages**, **Warnings and errors** or **Errors only**. A description of which messages fall under each category is shown below.

2. **Append to file**

   If, rather than sending mail, you want messages to be appended to a text file on the *admin system*, select this option, then 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
   Job Clients:   mickey minnie goofy
------------------------------------------------------------
SBA JOBERR: 000003
June 17 16:53:44 PDT 1999
Job 000003 cannot be written to the tape.
   Backup Device: rmt1
   Backup Server: spiderman
   Job clients:   mickey minnie goofy
   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.
------------------------------------------------------------
SBA JOBSTART: 000003
June 17 16:58:21 PDT 1999
Job 000003 has been started.
   Backup Device: rmt1
   Backup Server: spiderman
   Job Clients:   mickey minnie goofy
------------------------------------------------------------
SBA JOBOK: 000003
June 17 17:45:12 PDT 1999
Job 000003 completed successfully.
   Backup Device: rmt1
   Backup Server: spiderman
   Job Clients:   mickey minnie goofy
   Backup ID: 929665124
```

For easier identification of important messages, all messages in this file contain a header indicating the message type. These include:

TAG	MSG TYPE	DESCRIPTION
ERROR	error	general error (error message)
INFO	note	general info (all messages)
VOLCHG	warning	tape volume change requested (all messages)
JOBSTART	note	a job has started (all messages)
JOBOK	note	a job completed successfully
JOBWARN	warning	a job completed successfully with a warning message
JOBERR	error	a job terminated with an error

If you set the Messages to append value to “**All messages**”, then you will see all messages in the file. If you set it to “**Errors and warnings**”, then all “note” messages will be omitted. Likewise, if you set it to “**Errors only**”, then only the messages of type “error” will be shown.

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

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 Notification 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:

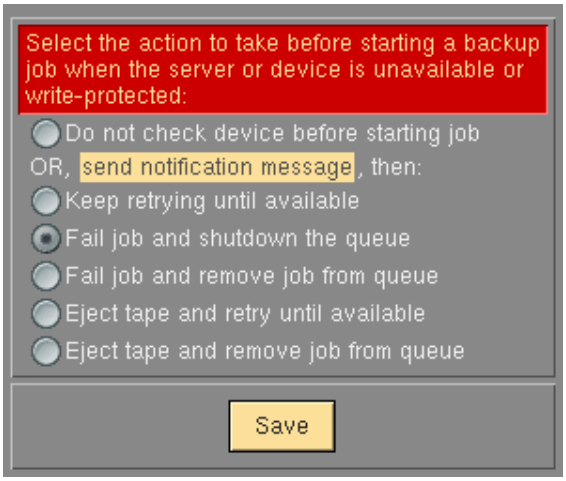

By default, the backup server and device availability is checked before a backup job is started. The first option, **Do not check device before starting job**, allows a backup to begin without first checking the availability. This may be preferable, for instance, when you have a pre-backup program which 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 **Save** and this new error handling behavior will be applied for all devices and servers.

24. Calendar

When a backup job is scheduled via the job configuration screen, an entry is created in the systems crontab configuration. All scheduled backup jobs are started using the cron daemon. You can view the backup schedule for a particular month, week, or day in a calendar view by selecting View from the menu bar and selecting Calendar.

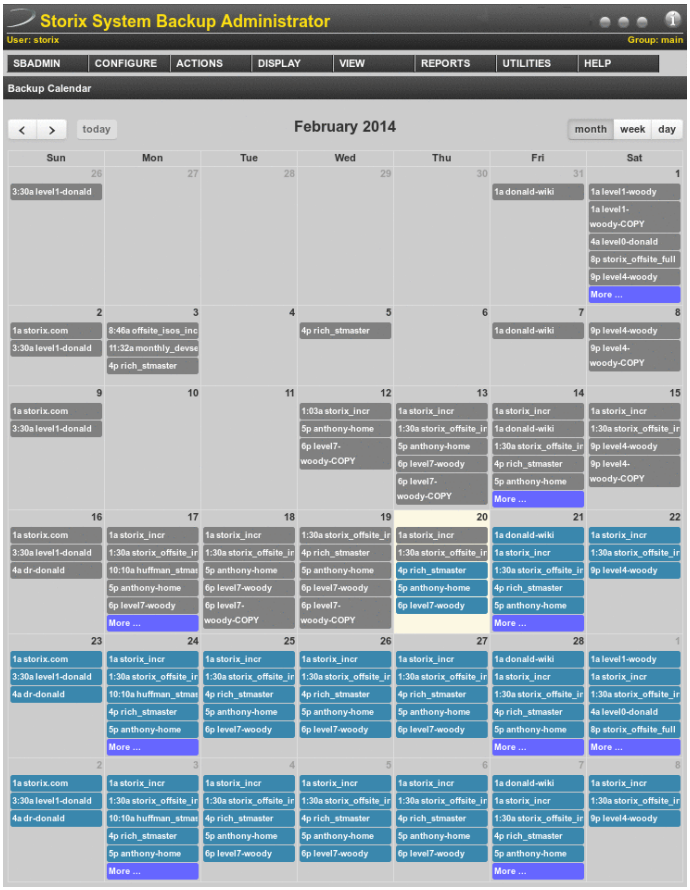

This feature is only available using the Web Interface

The default view for the calendar is the **Monthly** view of scheduled jobs for that month. Jobs that are scheduled to run in the future will be displayed with a teal colored event box. When you select the event box for a backup job that has not run yet, you will be redirected to the Job Configuration screen for the selected job.

Backup jobs that have already run will be displayed based on the backup label information and will have a grey event box. If you select an event job for a previously run backup job, you will be redirected to the Job Status/Output screen for that job. **If the backup label has been removed, backup information will not appear in the calendar.**

Only 5 events are displayed for each day when viewing by **Month**. If there are more events, there will be a blue event box titled "More ..." that when selected, will change to view to the selected **Day** view. The **Week** and **Day** view do not limit the number of events and the page will scroll if there are more events than will fit within the browser window.

Jobs cannot be scheduled using the Calendar; this page is only used to view the backup job schedule.

25. 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:

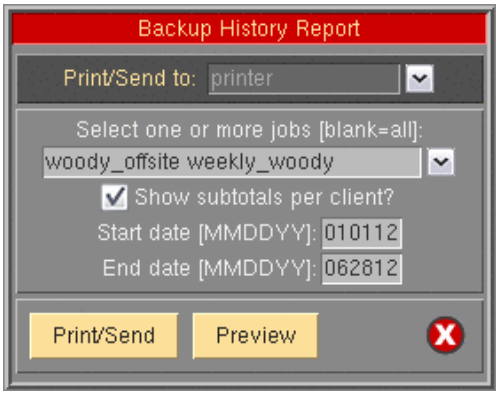

The above example is used when print a "**Backup Job History**" report. The option at the top, **Print/Send to**, and the **Print/Send** and **Preview** buttons at the bottom of the screen are provided for all report options. The other options will vary for each report option selected.

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):

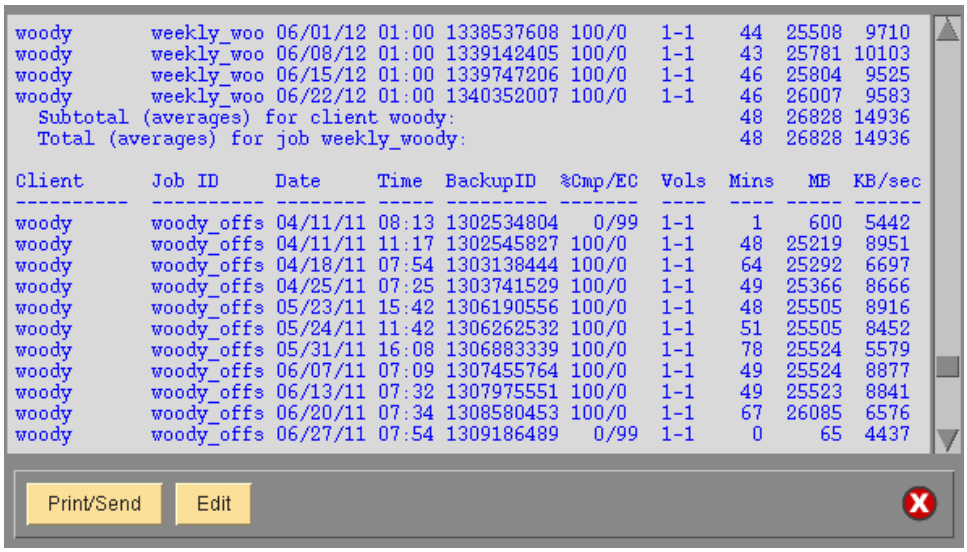

```
woody       weekly_woo 06/01/12 01:00 1338537608 100/0   1-1    44  25508  9710
woody       weekly_woo 06/08/12 01:00 1339142405 100/0   1-1    43  25781 10103
woody       weekly_woo 06/15/12 01:00 1339747206 100/0   1-1    46  25804  9525
woody       weekly_woo 06/22/12 01:00 1340352007 100/0   1-1    46  26007  9583
  Subtotal (averages) for client woody:                             48  26828 14936
  Total (averages) for job weekly_woody:                            48  26828 14936

Client      Job ID     Date     Time  BackupID  %Cmp/EC  Vols  Mins   MB  KB/sec
----------  ---------- -------- ----- --------- -------  ----  ---- ----- ------
woody       woody_offs 04/11/11 08:13 1302534804   0/99  1-1     1   600  5442
woody       woody_offs 04/11/11 11:17 1302545827 100/0   1-1    48 25219  8951
woody       woody_offs 04/18/11 07:54 1303138444 100/0   1-1    64 25292  6697
woody       woody_offs 04/25/11 07:25 1303741529 100/0   1-1    49 25366  8666
woody       woody_offs 05/23/11 15:42 1306190556 100/0   1-1    48 25505  8916
woody       woody_offs 05/24/11 11:42 1306262532 100/0   1-1    51 25505  8452
woody       woody_offs 05/31/11 16:08 1306883339 100/0   1-1    78 25524  5579
woody       woody_offs 06/07/11 07:09 1307455764 100/0   1-1    49 25524  8877
woody       woody_offs 06/13/11 07:32 1307975551 100/0   1-1    49 25523  8841
woody       woody_offs 06/20/11 07:34 1308580453 100/0   1-1    67 26085  6576
woody       woody_offs 06/27/11 07:54 1309186489   0/99  1-1     0    65  4437
```


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 not available when using Workstation Edition.

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 device details**: Check this box if you want to also show a list of the devices configured for each server listed.

Devices

Select Reports→Devices to view or print the device information for one or all servers. Refer to the main Reports section above for details on the Print and Preview options. When selecting this option, an additional option is provided:

- **Server**: To show devices for a single server, select the server in this field, otherwise leave blank for all servers.

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**: 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.

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 running the report in the order of client or job ID.

When using Workstation 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.

Restore History

To print a Restore History Report showing details regarding the restore, select Reports→Restore History from the main menu bar. A further option is provided for running the report in the order of job or backup ID.

Restore history was introduced in SBAdmin v8.1. Therefore, only restores performed with SBAdmin v8.1 will be reported.

Running the report by Job ID: Select Reports→Restore 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.

Running the report by Backup ID: Select Reports→Restore History→By Backup ID. You may select one or more Backup IDs for which to print the report or, by not specifying any backup IDs, the report will be reported for all backup IDs.

You may also select starting and ending dates for this report. If provided, the report will only include restores that occurred within that date range.

Backup Expiration Report

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.

Network Install Clients

This option is not available when using Workstation Edition.

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. It 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.

26. 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/Manage Boot Media

System Installation media is bootable media which may be used to boot the system to the **SBAdmin System Installation** process. To create boot media, select the following:

Utilities→Create/Manage Boot Media

Numerous options are available for creating system boot media, depending on the operating system and system type:

- **CDROM image** – A CDROM image is an *ISO9660 format filesystem image*, 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, on Solaris use the "**cdrw**" 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.

- **Tape** – Bootable tapes may be made created if the hardware platform supports booting from tape. At this time, only the ***IBM System p and System i*** type systems support booting from tape, and SBAdmin will support creation of bootable tapes for both ***AIX*** and ***Linux*** on this platform. When a **System Backup** is written to the beginning of a tape, the tape is automatically made bootable for the client's system type (if the hardware supports it). For AIX systems, you may also specify in the backup profile the type of system for which to create the boot tape (i.e. CHRP, RS6K, etc). Refer to the Backup Profiles for additional information on the ***platform type*** for bootable tapes.

- **Hard Disks** – Any hard disk, internal, external, or portable, may be made bootable after configuring the hard disk as a ***Local System Backup Disk***. This option is available when configuring **Clients** (if using Network Edition) or **Backup Devices and Directories** (for *Workstation Edition*) as described 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 the need for any other boot media.

- **Network** – 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 systems (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***.

> **For *Linux,* the network boot images are created and copied to the boot server, but some *bootloader* configuration must be manually performed by the user. This is automated on *AIX and Solaris* systems, but is more difficult for *Linux* due to the number of different boot loaders and configuration file formats that are available.**

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 System Recovery Guide*** for details.

Remote Installation Manager (RIM)

After a client system is booted from the *SBAdmin System Installation Media*, the installation process may be performed remotely from any compatible "**ssh**" client program. SBAdmin includes an *ssh client* for this purpose, allowing you to display, manage and perform the system installation of the client directly from the SBAdmin user interface on the *admin system*.

Before you can access the installation process on the client, the client must have started RIM, providing password required to log on. RIM and the associated password may have been setup and started within the client installation menus after booting from the media, or may have been configured when the boot media was created.

> **Only a user on a remote system (including this admin system) using a compatible *ssh client* program and the previously configured password can access the client's installation process.**

Once the client has enabled access, select the menu option Utilities→Start Remote Installation Manager. You will be asked for the client IP address and password. A new window will be displayed with the ***System Installation Menus***. You may then change the settings or perform the installation for the client.

Details of the configuring and enabling *Remote Install Manager* access is described in detail in the ***SBAdmin System Recovery Guide***.

Write a Tape Label ID to a Tape

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 contain 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.

To 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:

You must press the arrow keys next to each entry field to list and select the backup **Server Name** (if *Network Edition* used) and the **Device Name** in which the tape is inserted.

When using *Network Edition*, you can also choose "**local (client tape)**" for the **Server**. By selecting this option, you indicate that you want to write to a tape drive attached to a client system rather than a device configured on a server. When doing so, a new option "**Client Name**" will appear where you must select the client where the tape drive is attached.

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. This is because 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 Tape Operations

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:

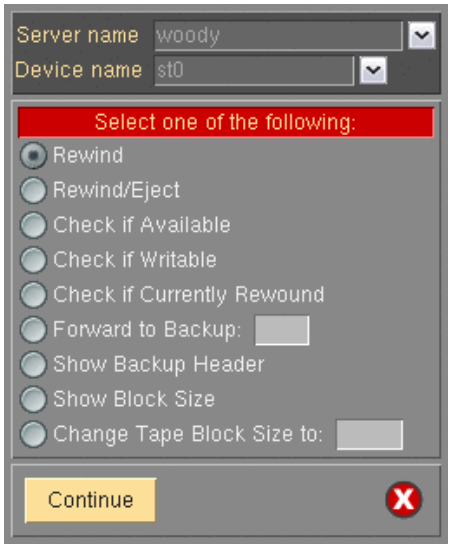

You must press the arrow keys next to each entry field to list and select the backup **Server Name** (if *Network Edition* used) and the **Device** in which the tape is inserted.

When using *Network Edition*, you can also choose "**local (client tape)**" for the **Server**. By selecting this option, you indicate that you want to write to a tape drive attached to a client system rather than a device configured on a server. When doing so, a new option "**Client Name**" will appear where you must select the client where the tape drive is attached.

Next, 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. 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 the drive was reconfigured, you will need to set the block size to 0 again before you can read a SBAdmin backup.

Perform Tape Library Operations

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 menu bar. After doing so, the following will be displayed:

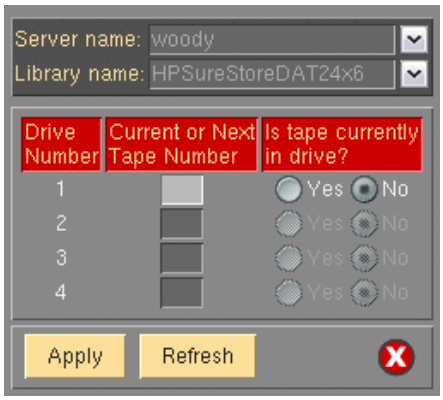

Select the **Library Name** by pressing the arrow button to the right of the entry field. Entry boxes will become available for the number of drives that are configured for this library. Next to the **Drive Number**, enter the **Next Tape to Use**, which must correspond to a tape position *for that drive*.

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:

Select the server name and the device name by pressing the button to the right of each entry field. Only devices 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:

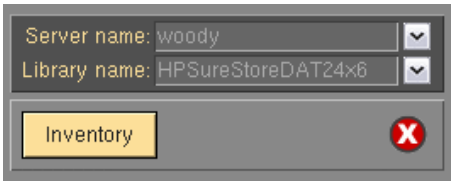

Select the server name (if *Network Edition* is used) and the library name by pressing the arrow buttons to the right of each field, then press the **Inventory** button at the bottom of the screen. After doing so, a new window will appear showing the command to execute and the output of that command, such as:

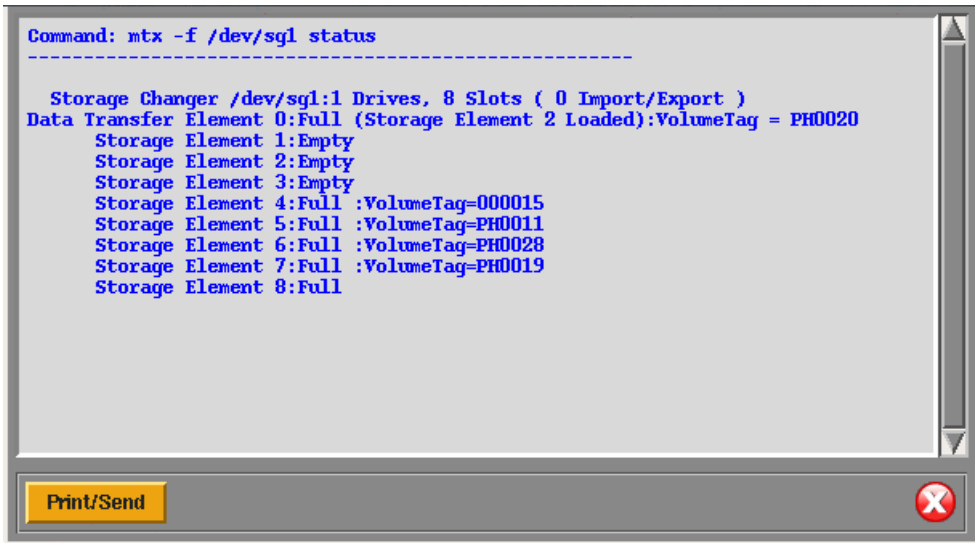

Any error messages that occur will also be displayed in that window. Simply press the **Close** button when done.

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.

> **This option will allow you to select disk backup files, although this is typically only done with tapes. When a backup label is expired for a disk backup, the actual disk backup is normally removed as well. The option would be necessary if you were, for instance, to manually move disk backup files from one server to another.**

To rebuild the backup label, select Utilities→Rebuild (unexpire) a Backup Label from the menu bar. A screen similar to the following will appear:

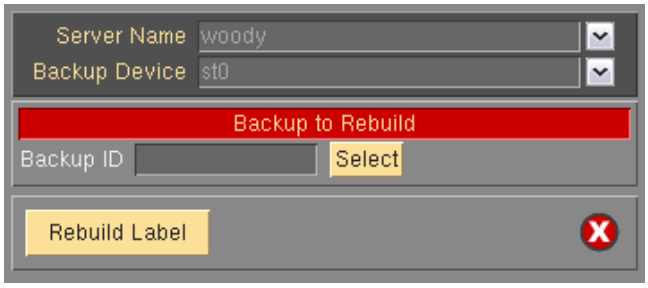

Use the arrow button to select a **Server Name** from the drop-down list. After doing so, you can select the **Backup Device** in the next field.

When using *Network Edition*, you can also choose "**local (client tape/disk)**" for the **Server**. By selecting this option, you indicate that you want to read the backup information from a disk (directory) or a tape drive attached to a client system, rather than a device configured on a server. After doing so, a new **Client Name** field will appear where you must select the client where the backup exists.

Then press the **Select** button to read the backup media to obtain the ***Backup ID***, which will be displayed in the **Backup ID** field.

If you have read the appropriate backup, just press the **Rebuild Label** button to begin reading the media and rebuilding the backup label. As this occurs, a screen similar to the following will appear:

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.

Read Error Settings

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 Settings from the menu bar. The following screen will appear:

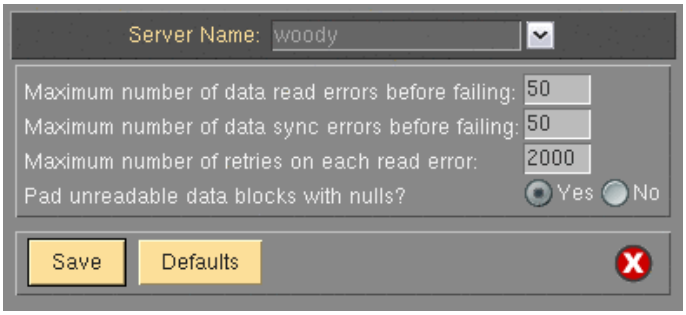

The following is an explanation of each field:

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.

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.

> **It is very important for the data to remain in the correct alignment when restoring raw device backups (logical/ZFS volumes, slices or meta-disks). 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, Zpool, Filesystem and File/Directory backups use an underlying restore command that is capable of re-synchronizing 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 of 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.

Change Access Permission of a Disk Backup

> **This option is not available when using Workstation Edition.**

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 original client that wrote the backup may read the contents, or any client with access to the directory on 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. A screen similar to the following will appear:

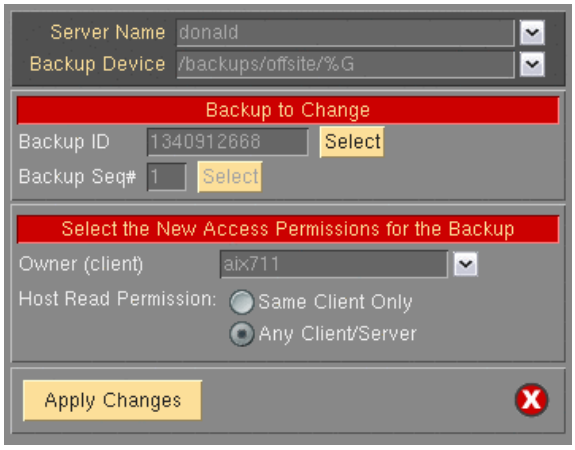

You will be asked to select the **Server name** and the **Backup Device** to which the desired backup was written. The directory or directories specified by the device will be searched and a list of backups will be displayed. Some of the backups may be part of the same *Backup ID* if there were multiple clients included in the same

backup job. In this case, you must select the specific client backup (sequence number) to change from the list. After doing so, the **Backup ID** and the **Backup Seq#** will be displayed.

You may change the following options:

- **Owner (client)** – This is the name of the client from which the backup originated. 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.

- **Host read permission** – By setting this to **Any client/server**, then any client may read from this backup if they have permission to read the directory on the server. If the option is set to **Same client only**, then only the owner, as shown in the Owner field, may read from this backup.

After you have made your selections, press the **Apply Changes** button to save your changes and exit this function.

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 Edition* license communicates via the **Transmission Control Protocol/Internet Protocol (TCP/IP).** This communication is handled through two different ports, the **Data port** and the **Status port**. 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 or CDs will attempt to communicate through the old port numbers if installing from a remote server. It is advised to create your boot media/images after changing your port numbers.

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 network ports. In addition, only specific commands can be run remotely.

All remote commands are executed using the "**strexec**" executable, which may be executed only by the root user on the system.

All attempts to run remote commands are checked on the remote system for authenticity as follows:

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 & Solaris or */usr/lpp/storix* for AIX).

1. The IP address of the sender is checked to see if it is a valid SBAdmin administrator system. Valid admin systems are specified in the **$STXPATH/config/admin_servers** file when SBAdmin is installed onto a client or server. If the caller is an administrator system, no further hostname or IP address checking is performed.

2. The *groupid* of the caller is checked that it is a member of the same group. Or, if a client is calling a server, the **$STXPATH/config/serverinfo** file is checked to see if the caller is a member of an enabled group. This file is created by the SBAdmin administrator system and copied to each server when changes are made to the server's access permissions.

3. The IP address of the sender is checked to see if it is a valid client (if calling a server). The **$STXPATH/config/serveraccess_*groupid*** file determines the permitted hosts. This file is created by the SBAdmin administrator system and copied to each server when clients are added or removed from the server's group.

4. 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.

5. 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.

6. 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. Commands in this file may be designated as only available to a calling administrator, client, server owner, server group, etc.

7. 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.

8. When executing a command to read or write backup media (i.e."**stio**" or "**sttape**"), the device specification may be:

    a. A configured SBAdmin device name for the server. If it's a directory-based device, the base name of the file must also be specified in a separate option (since the directory-device refers only to the pathname).

    b. A tape drive name only if it is also physical drive configured within an SBAdmin device for the server.

    c. The full pathname to a backup image file created by SBAdmin, and only if the pathname to the file is also configured as a directory within an SBAdmin device for the server.

Note that all of the above configuration files and directories may only be written by the root user on the system.

Remote Installation Manager

The ***Remote Installation Manager (RIM)*** provides a remote system anywhere on the network to connect to the system installation process of a client. This access is provided using a secure (**ssh**) connection. Only one remote ssh program may connect to the client at a time. Since this access is only available when this option has

been configured, a password has been set, and the client is booted to the system installation process, there is little security risk, but it is worth noting that the remote user will have access to all system installation process options and installation media available to the client system.

RIM access is only provided after booting from SBAdmin system installation boot media and either:

a. Enabling RIM access within the system installation menus from the client

or

b. Pre-configuring automatic enabling of RIM access when the installation media is configured.

In either case, a password is selected that the remote ssh program must use to connect to the client's installation process.

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 or the information therein is ever sent over the network.

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 in the ***Commands Reference Guide*** for details.

Storix Support

Should you encounter a problem using ***SBAdmin*** or have any questions, numerous support options are available. Select **Help→Storix Support** from the menu bar 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

C

D

E

F

V

W

Z