It comes in the night and sucks the essence from your computers.

Kern Sibbald

21 April 2005
This manual documents Bacula version 1.36.3

Contents

What is Bacula? 8

Who Needs Bacula? 8

Bacula Components or Services 8

Bacula Configuration 11

Conventions Used in this Document 12

Quick Start 13

Terminology 13

What Bacula is Not 17

Interactions Between the Bacula Services 18

Current State of Bacula 19

What is Implemented 19

Advantages of Bacula Over Other Backup Programs 21

Current Implementation Restrictions 22

Design Limitations or Restrictions 23

System Requirements 24

System Requirements 24

Supported Operating Systems 26

Supported Operating Systems 26

Supported Tape Drives 27

Supported Tape Drives 27
Unsupported Tape Drives 28
FreeBSD Users Be Aware!!! 28
Supported Autochangers 28
Getting Started with Bacula 29
Understanding Pools, Volumes and Labels 29
Setting Up Bacula Configuration Files 30
Testing your Configuration Files 33
Testing Bacula Compatibility with Your Tape Drive 34
Get Rid of the /lib/tls Directory 34
Running Bacula 35
Log Rotation 35
Disaster Recovery 35
Installing Bacula 36
General 36
Upgrading Bacula 36
Dependency Packages 37
Supported Operating Systems 38
Building Bacula from Source 38
What Database to Use? 42
Quick Start 43
Configure Options 43
Recommended Options for most Systems 44
RedHat 44
Solaris 45

FreeBSD 46
Win32 46
Windows Systems with CYGWIN Installed 46
Kern's Configure Script 47
Installing Bacula 47
Building a File Daemon or Client 48
Auto Starting the Daemons 48
Other Make Notes 49
Installing Tray Monitor 51
Modifying the Bacula Configuration Files 52
Critical Items to Implement Before Going Production 53
General 53
Critical Items 53
Recommended Items 54
Brief Tutorial 55
Before Running Bacula 55
Starting the Database 56
Starting the Daemons 56
Interacting with the Director to Query or Start Jobs 57
Running a Job 58
Restoring Your Files 64
Quitting the Console Program 67
Adding a Second Client 67
When The Tape Fills 69
Other Useful Console Commands 72

Debug Daemon Output 73
Have Patience When Starting the Daemons or Mounting Blank Tapes 73
Difficulties Connecting from the FD to the SD 74
Daemon Command Line Options 74
Creating a Pool 75
Labeling Your Volumes 76
Labeling Volumes with the Console Program 77
Customizing the Configuration Files 79
Resource Directive Format 80
Resource Types 84
Names, Passwords and Authorization 85
Detailed Information for each Daemon 86
Configuring the Director 87
Director Resource Types 87
Director Resource 88
Job Resource 91
JobDefs Resource 105
Schedule Resource 105
Technical Notes on Schedules 109
Client Resource 110
Storage Resource 112
Pool Resource 115
Catalog Resource 123
Messages Resource 125
Console Resource 125

Counter Resource 127
Example Director Configuration File 128
Client/File daemon Configuration 131
General 131
Client Resource 131
Director Resource 134
Message Resource 135
Example Client Configuration File 135
Storage Daemon Configuration 136
General 136
Storage Resource 136
Director Resource 139
Device Resource 139
Capabilities 148
Messages Resource 149
Sample Storage Daemon Configuration File 149
Messages Resource 152
Messages Resource 152
Console Configuration 158
General 158
Director Resource 158
ConsoleFont Resource 159
Console Resource 159
Console Commands 161
Sample Console Configuration File 161

Monitor Configuration 163
General 163
Monitor Resource 163
Director Resource 164
Client Resource 164
Storage Resource 165
Sample Monitor configuration file and related daemons' configuration records.166
Bacula Console 168
General 168
Configuration 168
Running the Console Program 168
Stopping the Console Program 169
Alphabetic List of Console Commands 170
Special dot Commands 184
Special At (@) Commands 184
Running the Console Program from a Shell Script 185
Adding Volumes to a Pool 186
Bacula Console Restore Command 188
General 188
Restore Command 188
Selecting Files by Filename 194
Command Line Arguments 195
Restoring Directory Attributes 197
Restoring on Windows 197
Restoring Files Can Be Slow 198

Problems Restoring Files 198
Example Restore Job Resource 199
File Selection Commands 200
Catalog Maintenance 203
Setting Retention Periods 203
Compacting Your MySQL Database 204
Repairing Your MySQL Database 205
Repairing Your PostgreSQL Database 205
Compacting Your PostgreSQL Database 206
Compacting Your SQLite Database 206
Migrating from SQLite to MySQL 207
Backing Up Your Bacula Database 207
Backing Up Third Party Databases 208
Database Size 209
Automatic Volume Recycling 211
Automatic Pruning 212
Prunning Directives 212
Recycling Algorithm 215
Recycle Status 216
Making Bacula Use a Single Tape 217
Daily, Weekly, Monthly Tape Usage Example 218
Automatic Pruning and Recycling Example 220
Manually Recycling Volumes 222
Basic Volume Management 224
Key Concepts and Resource Records 224

Example 229
Backing up to Multiple Disks 231
Considerations for Multiple Clients 232
Using Pools to Manage Volumes 236
Problem 236
Solution 236
Overall Design 237
Actual Conf Files 239
Backup Strategies 243
Simple One Tape Backup 243
Manually Changing Tapes 244
Daily Tape Rotation 245
Autochangers Support 252
Autochangers – General 252
Knowing What SCSI Devices You Have 253
Example Scripts 253
Slots 254
Multiple Devices 254
Device Configuration Records 255
Example Configuration File 257
Specifying Slots When Labeling 257
Dealing with Multiple Magazines 258
Simulating Barcodes in your Autochanger 259
Full Form of the Update Slots Command 260
FreeBSD Issues 261

Testing the Autochanger and Adapting Your mtx-changer Script 261
Using the Autochanger 263
Barcode Support 265
Bacula Autochanger Interface 265
Supported Autochangers 267
Supported Autochanger Models 267
Data Spooling 269
Data Spooling Directives 269
MAJOR WARNING !!! 270
Other Points 270
Bacula Frequently Asked Questions 272
Tips and Suggestions 284
Examples 284
Upgrading Bacula Versions 284
Getting Notified of Job Completion 284
Getting Email Notification to Work 286
Getting Notified that Bacula is Running 286
Maintaining a Valid Bootstrap File 288
Rejected Volumes After a Crash 290
Security Considerations 294
Creating Holiday Schedules 294
Automatic Labeling Using Your Autochanger 295
Backing Up Portables Using DHCP 296
Going on Vacation 296
How to Excude File on Windows Regardless of Case 297

Executing Scripts on a Remote Machine 297
Recycling All Your Volumes 298
Backing up ACLs on ext3 or XFS filesystems 299
Total Automation of Bacula Tape Handling 299
Running Concurrent Jobs 301
Volume Utility Tools 303
Specifying the Configuration File 303
Specifying a Device Name For a Tape 303
Specifying a Device Name For a File 303
Specifying Volumes 303
bls 304
bextract 308
bscan 310
bcopy 315
btape 315
Other Programs 318
bsmtp 318
dbcheck 319
testfind 322
bimagemgr 323
Testing Your Tape Drive With Bacula 327
Summary of Steps to Take to Get Your Tape Drive Working . 327
btape 329
Tips for Resolving Problems 332
Recovering Files Written to Tape With Fixed Block Sizes . . 341

Tape Blocking Modes 342
What To Do When Bacula Crashes (Kaboom) 343
Traceback 343
Testing The Traceback 344
Getting A Traceback On Other Systems 344
Manually Running Bacula Under The Debugger 345
Getting Debug Output from Bacula 346
Windows Version of Bacula 347
General 347
Disaster Recovery Using Bacula 361
General 361
Important Considerations 361
Steps to Take Before Disaster Strikes 361
Bare Metal Recovery on Linux with a Bacula Rescue CDROM 362
Requirements 363
Directories 363
Preparation for a Bare Metal Recovery 363
Creating a Bacula Rescue CDROM 364
Putting Two or More Systems on Your Rescue Disk 366
Restoring a Client System 368
Boot with your Bacula Rescue CDROM 369
Restoring a Server 373
Linux Problems or Bugs 374
FreeBSD Bare Metal Recovery 374
Solaris Bare Metal Recovery 376

Preparing Solaris Before a Disaster 376
Bugs and Other Considerations 377
Disaster Recovery of Win32 Systems 378
Resetting Directory and File Ownership and Permissions on Win32 Systems379
Alternate Disaster Recovery Suggestion for Win32 Systems . 379
Restoring to a Running System 380
Additional Resources 380
Using Bacula to Encrypt Communications to Clients 382
Communications Ports Used 382
Encryption 382
Picture 383
Certificates 383
Securing the Data Channel 384
Modification of bacula-dir.conf for the Data Channel 384
config Files for stunnel to Encrypt the Data Channel 385
Starting and Testing the Data Encryption 386
Encrypting the Control Channel 387
Modification of bacula-dir.conf for the Control Channel . . . 387
config Files for stunnel to Encrypt the Control Channel . . . 388
Starting and Testing the Control Channel 389
Using stunnel to Encrypt to a Second Client 389
Creating a Self-signed Certificate 390
Getting a CA Signed Certificate 391
Using ssh to Secure the Communications 391
Bacula Security Issues 392

Configuring and Testing TCP Wrappers with Bacula 393
Running as non-root 395
Dealing with Firewalls 397
Technical Details 397
Concrete Example 398
Using Bacula to Improve Computer Security 403
Details 403
Running the Verify 405
What To Do When Differences Are Found 406
Verify Configuration Example 408
Bacula® - RPM Packaging FAQ 410
Answers 410
Bootstrap File 413
File Format 413
Automatic Generation of Bootstrap Files 418
Final Example 418
Catalog Services 420
General 420
Sequence of Creation of Records for a Save Job 422
Database Tables 423
Installing and Configuring MySQL 435
Installing and Configuring MySQL – Phase I 435
Installing and Configuring MySQL – Phase II 436
Re-initializing the Catalog Database 437
Linking Bacula with MySQL 438

Installing and Configuring PostgreSQL 440
Installing and Configuring PostgreSQL – Phase I 440
Installing and Configuring PostgreSQL – Phase II 440
Re-initializing the Catalog Database 442
Converting from MySQL to PostgreSQL 442
Credits 444
Installing and Configuring SQLite 445
Installing and Configuring SQLite – Phase I 445
Installing and Configuring SQLite – Phase II 445
Linking Bacula with SQLite 446
Testing SQLite 446
Re-initializing the Catalog Database 446
Bacula internal database is no longer supported, please do not use it. 448
Internal Bacula Database 448
Disaster Recovery Using a Bacula Rescue Floppy 449
General 449
Important Considerations 449
Steps to Take Before Disaster Strikes 450
Bare Metal Floppy Recovery on Linux with a Bacula Floppy Rescue Disk 450
Restrictions 451
Directories 451
Preparation for a Bare Metal Recovery 451
Restoring Your Linux Client with a Floppy 457
Linux Problems or Bugs 463
tomsrtbt 464

Bacula Copyright, Trademark, and Licenses 466
GPL 466
LGPL 466
Public Domain 466
Trademark 466
Disclaimer 467
GNU General Public License 468
Table of Contents 468
GNU GENERAL PUBLIC LICENSE 468
Preamble 468
TERMS AND CONDITIONS 469
How to Apply These Terms to Your New Programs 474
GNU Lesser General Public License 476
Table of Contents 476
GNU LESSER GENERAL PUBLIC LICENSE 476
Preamble 476
TERMS AND CONDITIONS 479
How to Apply These Terms to Your New Libraries 486
Bacula Projects 488
Thanks 489
Copyrights and Trademarks 490
Bacula Bugs 491
Variable Expansion 492
General Functionality 492
Bacula Variables 492

Full Syntax 493

Semantics 495

Examples 495

List of Figures

Bacula Applications 8
Bacula Objects 11
Interactions between Bacula Services 18
Bacula Tray Monitor 31
Bacula Objects 79
Bacula CD Image Manager 324
Bacula CD Image Burn Progress Window 324
Bacula CD Image Burn Results 325
Win32 Client Setup Wizard 348
Win32 Component Selection Dialog 348
Win32 Directory Selection Dialog 348
Win32 Client Service Selection 349
Win32 Client Service Confirmation 349
Win32 Client Start 349
Win32 Client Setup Completed 349

List of Tables

Supported Tape Drives 27
Depedency Packages 37
Resource Types 84
Autochangers Known to Work with Bacula 267
WinNT/2K/XP Restore Portability Status 353
Filename Table Layout 423
Path Table Layout 423
File Table Layout 424
Job Table Layout 425
Job Types 426
Job Statuses 426
File Sets Table Layout 427
JobMedia Table Layout 427
Media Table Layout 428
Pool Table Layout 429
Client Table Layout 429
Unsaved Files Table Layout 430
Counter Table Layout 430
Version Table Layout 430

Base Files Table Layout 430

SQLite vs MySQL Database Comparison 448

What is Bacula?

Bacula is a set of computer programs that permit you (or the system administrator) to manage backup, recovery, and verification of computer data across a network of computers of different kinds. In technical terms, it is a network Client/Server based backup program. Bacula is relatively easy to use and efficient, while offering many advanced storage management features that make it easy to find and recover lost or damaged files. Due to its modular design, Bacula is scalable from small single computer systems to systems consisting of hundreds of computers located over a large network.

Who Needs Bacula?

If you are currently using a program such as **tar**, **dump**, or **bru** to backup your computer data, and you would like a network solution, more flexibility, or catalog services, Bacula will most likely provide the additional features you want. However, if you are new to Unix systems or do not have offsetting experience with a sophisticated backup package, we do not recommend using Bacula as it is much more difficult to setup and use than **tar** or **dump**.

If you are running **Amanda** and would like a backup program that can write to multiple volumes (i.e. is not limited by your tape drive capacity), Bacula can most likely fill your needs. In addition, quite a number of our users report that Bacula is simpler to setup and use than other equivalent programs.

If you are currently using a sophisticated commercial package such as Legato Networker. ARCserveIT, Arkeia, or PerfectBackup+, you may be interested in Bacula, which provides many of the same features, and is free software available under the GNU Version 2 software license.

Bacula Components or Services

Bacula is made up of the following five major components or services:

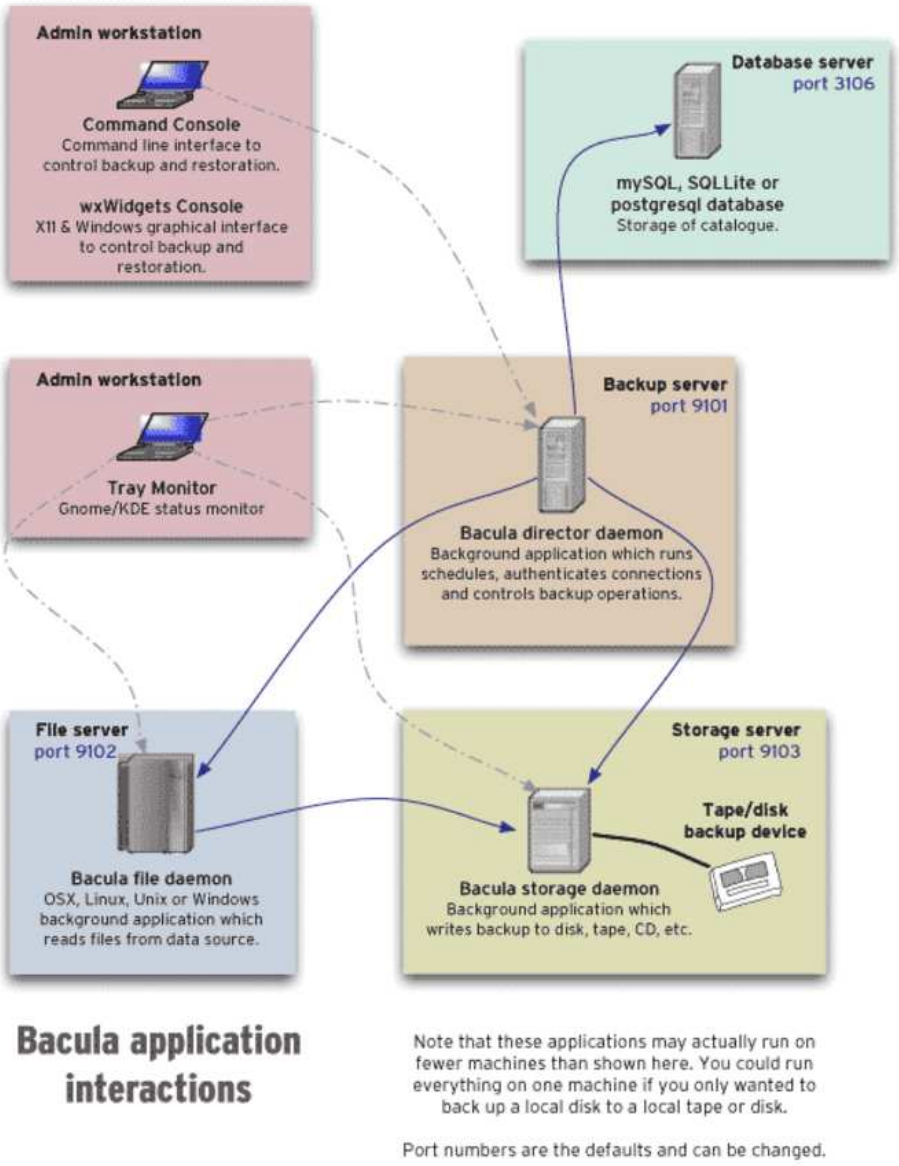

(thanks to Aristedes Maniatis for this graphic and the one below)

- **Bacula Director** service consists of the program that supervises all the backup, restore, verify and archive operations. The system administrator uses the Bacula Director to schedule backups and to recover files. For more details see the Director Services Daemon Design Document in the Bacula Developer's Guild. The Director runs as a daemon or a service (i.e. in the background).

- **Bacula Console** services is the program that allows the administrator or user to communicate with the **Bacula Director** (see above). Currently, the Bacula Console is available in three versions. The first and simplest is to run the Console program in a shell window (i.e. TTY interface). Most system administrators will find this completely adequate. The second version is a GNOME GUI interface that for the moment (23 November 2003) is far from complete, but quite functional as it has most the capabilities of the shell Console. The third version is a wxWidgets GUI with an interactive file restore. It also has most the capabilities of the shell console, allows command completion with tabulation, and gives you instant help about the command you are typing. For more details see the Bacula Console Design Document.

- **Bacula File** services (or Client program) is the software program that is installed on the machine to be backed up. It is specific to the operating system on which it runs and is responsible for providing the file attributes and data when requested by the Director. The File services are also responsible for the file system dependent part of restoring the file attributes and data during a recovery operation. For more details see the File Services Daemon Design Document in the Bacula Developer's Guide. This program runs as a daemon on the machine to be backed up, and in some of the documentation, the File daemon is referred to as the Client (for example in Bacula's configuration file). In addition to Unix/Linux File daemons, there is a Windows File daemon (normally distributed in binary format). The Windows File daemon runs on all currently known Windows versions (95, 98, Me, NT, 2000, XP).

- **Bacula Storage** services consist of the software programs that perform the storage and recovery of the file attributes and data to the physical backup media or volumes. In other words, the Storage daemon is responsible for reading and writing your tapes (or other storage media, e.g. files). For more details see the Storage Services Daemon Design Document in the Bacula Developer's Guild. The Storage services runs as a daemon on the machine that has the backup device (usually a tape drive).

- **Catalog** services are comprised of the software programs responsible for maintaining the file indexes and volume databases for all files backed up. The Catalog services permit the System Administrator or user to quickly locate and restore any desired file. The Catalog services sets Bacula apart from simple backup programs like tar and bru, because the catalog maintains a record of all Volumes used, all Jobs run, and all Files saved, permitting efficicient restoration and Volume management. Bacula currently supports three different

databases, MySQL, PostgreSQL, and SQLite, one of which must be chosen when building **Bacula**. There also exists an Internal database, but it is no longer supported.

The three SQL databases currently supported (MySQL, PostgreSQL or SQLite) provide quite a number of features, including rapid indexing, arbitrary queries, and security. Although we plan to support other major SQL databases, the current Bacula implementation interfaces only to MySQL, PostgreSQL and SQLite. For more details see the Catalog Services Design Document.

The RPMs for MySQL and PostgreSQL ship as part of the Linux RedHat release, or building it from the source is quite easy, see the Installing and Configuring MySQL chapter of this document for the details. For more information on MySQL, please see: www.mysql.com. Or see the Installing and Configuring PostgreSQL chapter of this document for the details. For more information on PostgreSQL, please see: www.postgresql.org.

Configuring and building SQLite is even easier. For the details of configuring SQLite, please see the Installing and Configuring SQLite chapter of this document.

- **Bacula Monitor** services is the program that allows the administrator or user to watch current status of **Bacula Directors**, **Bacula File Daemons** and **Bacula Storage Daemons** (see above). Currently, only a GTK+ version is available, which works with Gnome and KDE (or any window manager that supports the FreeDesktop.org system tray standard).

To perform a successful save or restore, the following four daemons must be configured and running: the Director daemon, the File daemon, the Storage daemon, and MySQL, PostgreSQL or SQLite.

Bacula Configuration

In order for Bacula to understand your system, what clients you want backed up, and how, you must create a number of configuration files containing resources (or objects). The following presents an overall picture of this:

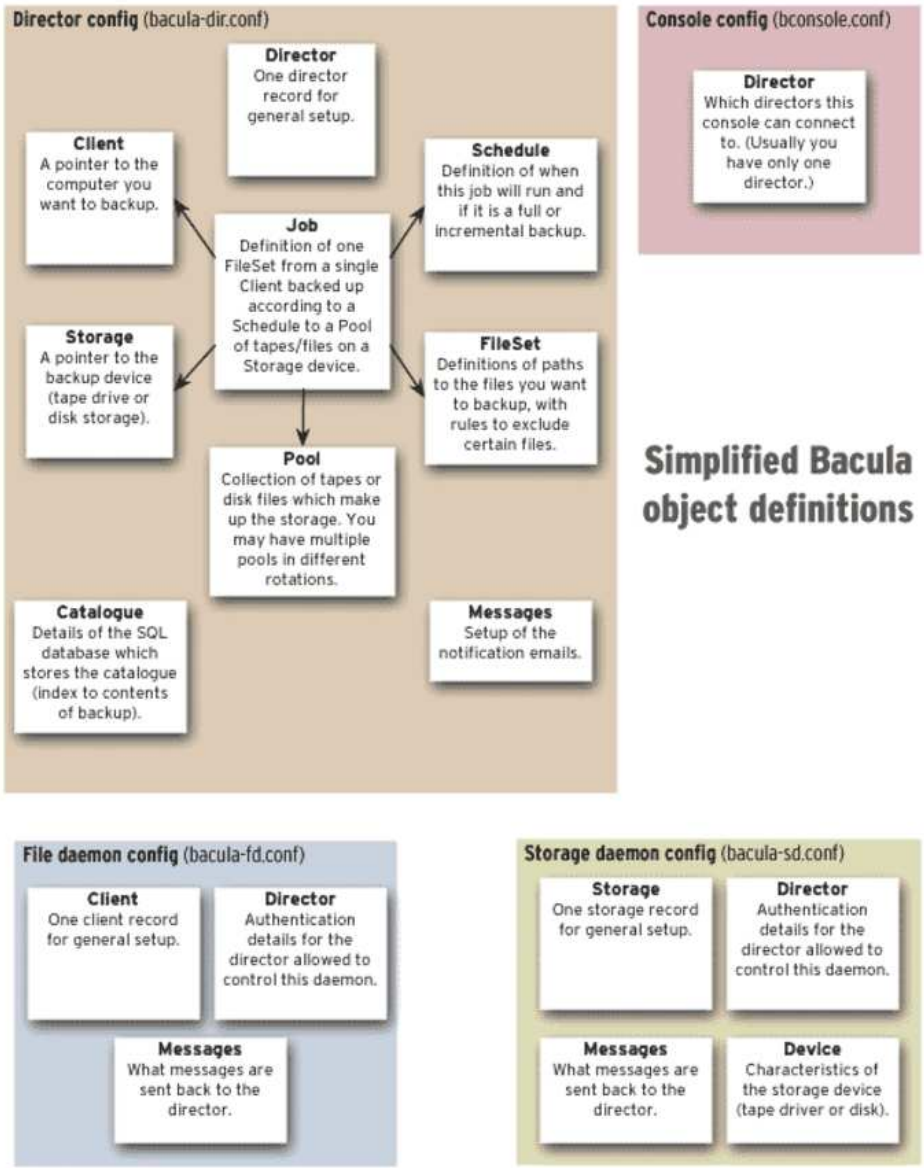

Conventions Used in this Document

Bacula is in a state of evolution, and as a consequence, this manual will not always agree with the code. If an item in this manual is preceded by an asterisk (*), it indicates that the particular feature is not implemented. If it is preceded by a plus sign (+), it indicates that the feature may be partially implemented.

If you are reading this manual as supplied in a released version of the software, the above paragraph holds true. If you are reading the online version of the manual, www.bacula.org/manual, please bear in mind that this version describes the current version in development (in the CVS) that may contain features not in the released version. Just the same, it generally lags behind the code a bit.

Quick Start

To get Bacula up and running quickly, we recommend that you first scan the Terminology section below, then quickly review the next chapter entitled The Current State of Bacula, then the Getting Started with Bacula, which will give you a quick overview of getting Bacula running. After which, you should proceed to the chapter on Installing Bacula, then How to Configure Bacula, and finally the chapter on Running Bacula.

Terminology

To facilitate communication about this project, we provide here the definitions of the terminology that we use.

Administrator The person or persons responsible for administrating the Bacula system.

Backup We use the term **Backup** to refer to a Bacula Job that saves files.

Bootstrap File The bootstrap file is an ASCII file containing a compact form of commands that allow Bacula or the stand-alone file extraction utility (**bextract**) to restore the contents of one or more Volumes, for example, the current state of a system just backed up. With a bootstrap file, Bacula can restore your system without a Catalog. You can create a bootstrap file from a Catalog to extract any file or files you wish.

Catalog The Catalog is used to store summary information about the Jobs, Clients, and Files that were backed up and on what Volume or Volumes. The information saved in the Catalog permits the administrator or user to determine what jobs were run, their status as well as the important characteristics of each file that was backed up. The Catalog is an online resource, but does not contain the data for the files backed up. Most of the information stored in the catalog is also stored on the backup volumes (i.e. tapes). Of course, the tapes

will also have a copy of the file in addition to the File Attributes (see below).

The catalog feature is one part of Bacula that distinguishes it from simple backup and archive programs such as **dump** and **tar**.

Client In Bacula's terminology, the word Client refers to the machine being backed up, and it is synonymous with the File services or File daemon, and quite often, we refer to it as the FD. A Client is defined in a configuration file resource.

Console The program that interfaces to the Director allowing the user or system administrator to control Bacula.

Daemon Unix terminology for a program that is always present in the background to carry out a designated task. On Windows systems, as well as some Linux systems, daemons are called **Services**.

Directive The term directive is used to refer to a statement or a record within a Resource in a configuration file that defines one specific thing. For example, the **Name** directive defines the name of the Resource.

Director The main Bacula server daemon that schedules and directs all Bacula operations. Occassionally, we refer to the Director as DIR.

Differential A backup that includes all files changed since the last Full save started. Note, other backup programs may define this differently.

File Attributes The File Attributes are all the information necessary about a file to identify it and all its properties such as size, creation date, modification date, permissions, etc. Normally, the attributes are handled entirely by Bacula so that the user never needs to be concerned about them. The attributes do not include the file's data.

File Daemon The daemon running on the client computer to be backed up. This is also referred to as the File services, and sometimes as the Client services or the FD.

FileSet A FileSet is a Resource contained in a configuration file that defines the files to be backed up. It consists of a list of included files or directories, a list of excluded files, and how the file is to be stored (compression, encryption, signatures). For more details, see the FileSet Resource definition in the Director chapter of this document.

Incremental A backup that includes all files changed since the last Full, Differential, or Incremental backup started. It is normally specified on the **Level** directive within the Job resource definition, or in a Schedule resourc.

Job A Bacula Job is a configuration resource that defines the work that Bacula must perform to backup or restore a particular Client. It consists of the **Type** (backup, restore, verify, etc), the **Level** (full, incremental,...), the **FileSet**, and **Storage** the files are to be backed up (Storage device, Media Pool). For more details, see the Job Resource definition in the Director chapter of this document.

Monitor The program that interfaces to the all the daemons allowing the user or system administrator to monitor Bacula status.

Resource A resource is a part of a configuration file that defines a specific unit of information that is available to Bacula. For example, the **Job** resource defines all the properties of a specific Job: name, schedule, Volume pool, backup type, backup level, ...

Restore A restore is a configuration resource that describes the operation of recovering a file (lost or damaged) from backup media. It is the inverse of a save, except that in most cases, a restore will normally have a small set of files to restore, while normally a Save backs up all the files on the system. Of course, after a disk crash, Bacula can be called upon to do a full Restore of all files that were on the system.

Schedule A Schedule is a configuration resource that defines when the Bacula Job will be scheduled for execution. To use the Schedule, the Job resource will refer to the name of the Schedule. For more details, see the Schedule Resource definition in the Director chapter of this document.

Service This is Windows terminology for a **daemon** – see above. It is now frequently used in Unix environments as well.

Storage Coordinates The information returned from the Storage Services that uniquely locates a file on a backup medium. It consists of two parts: one part pertains to each file saved, and the other part pertains to the whole Job. Normally, this information is saved in the Catalog so that the user doesn't need specific knowledge of the Storage Coordinates. The Storage Coordinates include the File Attributes (see above) plus the unique location of the information on the backup Volume.

Storage Daemon The Storage daemon, sometimes referred to as the SD, is the code that writes the attributes and data to a storage Volume (usually a tape or disk).

Session Normally refers to the internal conversation between the File daemon and the Storage daemon. The File daemon opens a **session**

with the Storage daemon to save a FileSet, or to restore it. A session has a one to one correspondence to a Bacula Job (see above).

Verify A verify is a job that compares the current file attributes to the attributes that have previously been stored in the Bacula Catalog. This feature can be used for detecting changes to critical system files similar to what **Tripwire** does. One of the major advantages of using Bacula to do this is that on the machine you want protected such as a server, you can run just the File daemon, and the Director, Storage daemon, and Catalog reside on a different machine. As a consequence, if your server is ever compromised, it is unlikely that your verification database will be tampered with.

Verify can also be used to check that the most recent Job data written to a Volume agrees with what is stored in the Catalog (i.e. it compares the file attributes), *or it can check the Volume contents against the original files on disk.

***Archive** An Archive operation is done after a Save, and it consists of removing the Volumes on which data is saved from active use. These Volumes are marked as Archived, and many no longer be used to save files. All the files contained on an Archived Volume are removed from the Catalog. NOT YET IMPLEMENTED.

***Update** An Update operation causes the files on the remote system to be updated to be the same as the host system. This is equivalent to an **rdist** capability. NOT YET IMPLEMENTED.

Retention Period There are various kinds of retention periods that Bacula recognizes. The most important are the **File** Retention Period, **Job** Retention Period, and the **Volume** Retention Period. Each of these retention periods applies to the time that specific records will be kept in the Catalog database. This should not be confused with the time that the data saved to a Volume is valid.

The File Retention Period determines the time that File records are kept in the catalog database. This period is important because the volume of the database File records by far use the most storage space in the database. As a consequence, you must ensure that regular "pruning" of the database file records is done. (See the Console **retention** command for more details on this subject).

The Job Retention Period is the length of time that Job records will be kept in the database. Note, all the File records are tied to the Job that saved those files. The File records can be purged leaving the Job records. In this case, information will be available about the jobs that ran, but not the details of the files that were backed up. Normally, when a Job record is purged, all its File records will also be purged.

The Volume Retention Period is the minimum of time that a Volume will be kept before it is reused. Bacula will normally never overwrite a Volume that contains the only backup copy of a file. Under ideal conditions, the Catalog would retain entries for all files backed up for all current Volumes. Once a Volume is overwritten, the files that were backed up on that Volume are automatically removed from the Catalog. However, if there is a very large pool of Volumes or a Volume is never overwritten, the Catalog database may become enormous. To keep the Catalog to a manageable size, the backup information should removed from the Catalog after the defined File Retention Period. Bacula provides the mechanisms for the catalog to be automatically pruned according to the retention periods defined.

Scan A Scan operation causes the contents of a Volume or a series of Volumes to be scanned. These Volumes with the information on which files they contain are restored to the Bacula Catalog. Once the information is restored to the Catalog, the files contained on those Volumes may be easily restored. This function is particularly useful if certain Volumes or Jobs have exceeded their retention period and have been pruned or purged from the Catalog. Scanning data from Volumes into the Catalog is done by using the **bscan** program. See the bscan section of the Bacula Utilities Chapter of this manual for more details.

Volume A Volume is an archive unit, normally a tape or a named disk file where Bacula stores the data from one or more backup jobs. All Bacula Volumes have a software label written to the Volume by Bacula so that it identify what Volume it is really reading. (Normally there should be no confusion with disk files, but with tapes, it is easy to mount the wrong one).

What Bacula is Not

Bacula is a backup, restore and verification program and is not a complete disaster recovery system in itself, but it can be a key part of one if you plan carefully and follow the instructions included in the Disaster Recovery Chapter of this manual.

With proper planning, as mentioned in the Disaster Recovery chapter **Bacula** can be a central component of your disaster recovery system. For example, if you have created an emergency boot disk, a Bacula Rescue disk to save the current partitioning information of your hard disk, and maintain a complete Bacula backup, it is possible to completely recover your system from “bare metal”.

If you have used the **WriteBootstrap** record in your job or some other means to save a valid bootstrap file, you will be able to use it to extract the necessary files (without using the catalog or manually searching for the files to restore).

Interactions Between the Bacula Services

The following block diagram shows the typical interactions between the Bacula Services for a backup job. Each block represents in general a separate process (normally a daemon). In general, the Director oversees the flow of information. It also maintains the Catalog.

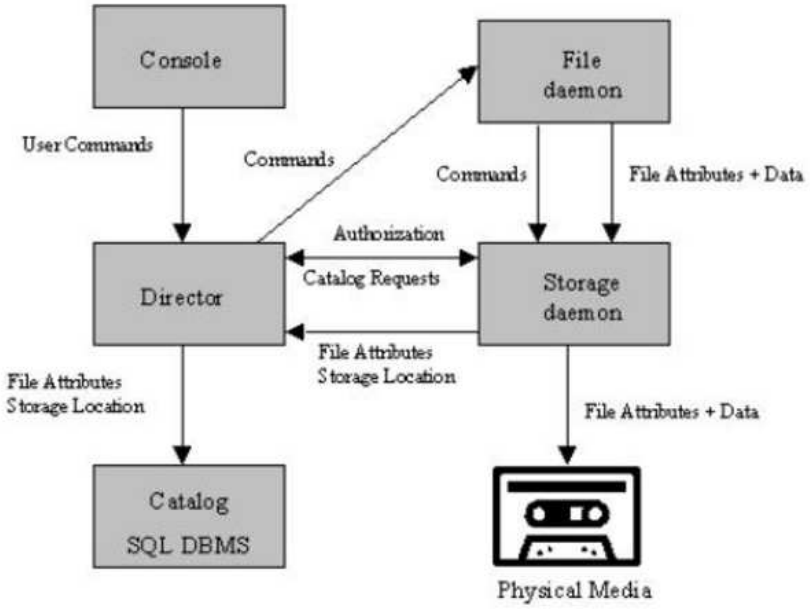

The Current State of Bacula

In other words, what is and what is not currently implemented and functional.

What is Implemented

- Network backup/restore with centralized Director.
- Internal scheduler for automatic Job execution.
- Scheduling of multiple Jobs at the same time.
- You may run one Job at a time or multiple simultaneous Jobs.
- Job sequencing using priorities.
- Restore of one or more files selected interactively either for the current backup or a backup prior to a specified time and date.
- Restore of a complete system starting from bare metal. This is mostly automated for Linux systems and partially automated for Solaris. See Disaster Recovery Using Bacula. This is also reported to work on Win2K/XP systems.
- Listing and Restoration of files using stand-alone **bls** and **bextract** tool programs. Among other things, this permits extraction of files when Bacula and/or the catalog are not available. Note, the recommended way to restore files is using the restore command in the Console. These programs are designed for use as a last resort.
- Ability to recreate the catalog database by scanning backup Volumes using the **bscan** program.
- Console interface to the Director allowing complete control. A shell, GNOME GUI and wxWidgets GUI versions of the Console program are available. Note, the GNOME GUI program currently offers very few additional features over the shell program.
- Verification of files previously cataloged, permitting a Tripwire like capability (system break-in detection).
- CRAM-MD5 password authentication between each component (daemon).
- A comprehensive and extensible configuration file for each daemon.

- Catalog database facility for remembering Volumes, Pools, Jobs, and Files backed up.
- Support for SQLite, PostgreSQL, and MySQL Catalog databases.
- User extensible queries to the SQLite, PostgreSQL and MySQL databases.
- Labeled Volumes, preventing accidental overwriting (at least by Bacula).
- Any number of Jobs and Clients can be backed up to a single Volume. That is, you can backup and restore Linux, Unix, Sun, and Windows machines to the same Volume.
- Multi-volume saves. When a Volume is full, **Bacula** automatically requests the next Volume and continues the backup.
- Pool and Volume library management providing Volume flexibility (e.g. monthly, weekly, daily Volume sets, Volume sets segregated by Client, ...).
- Machine independent Volume data format. Linux, Solaris, and Windows clients can all be backed up to the same Volume if desired.
- A flexible message handler including routing of messages from any daemon back to the Director and automatic email reporting.
- Multi-threaded implementation.
- Programmed to handle arbitrarily long filenames and messages.
- GZIP compression on a file by file basis done by the Client program if requested before network transit.
- Computation of MD5 or SHA1 signatures of the file data if requested.
- Saves and restores POSIX ACLs if enabled.
- Autochanger support using a simple shell interface that can interface to virtually any autoloader program. A script for **mtx** is provided.
- Support for autochanger barcodes – automatic tape labeling from barcodes.
- Automatic support for multiple autochanger magazines either using barcodes or by reading the tapes.
- Raw device backup/restore. Restore must be to the same device.

- All Volume blocks (approx 64K bytes) contain a data checksum.
- Access control lists for Consoles that permit restricting user access to only their data.
- Data spooling to disk during backup with subsequent write to tape from the spooled disk files. This prevents tape "shoe shine" during Incremental/Differential backups.
- Support for save/restore of files larger than 2GB.
- Support for 64 bit machines, e.g. amd64.
- Ability to encrypt communications between daemons using stunnel.

Advantages of Bacula Over Other Backup Programs

- Since there is a client for each machine, you can backup and restore clients of any type ensuring that all attributes of files are properly saved and restored.
- It is also possible to backup clients without any client software by using NFS or Samba. However, if possible, we recommend running a Client File daemon on each machine to be backed up.
- Bacula handles multi-volume backups.
- A full comprehensive SQL standard database of all files backed up. This permits online viewing of files saved on any particular Volume.
- Automatic pruning of the database (removal of old records) thus simplifying database administration.
- Any SQL database engine can be used making Bacula very flexible.
- The modular but integrated design makes Bacula very scalable.
- Since Bacula uses client file servers, any database or other application can be properly shutdown by Bacula using the native tools of the system, backed up, then restarted (all within a Bacula Job).
- Bacula has a built-in Job scheduler.
- The Volume format is documented and there are simple C programs to read/write it.
- Bacula uses well defined (registered) TCP/IP ports – no rpcs, no shared memory.

- Bacula installation and configuration is relatively simple compared to other comparable products.
- According to one user Bacula is as fast as the big major commercial application.
- According to another user Bacula is four times as fast as another commercial application, probably because that application stores its catalog information in a large number of individual files rather than an SQL database as Bacula does.
- Aside from a GUI administrative interface, Bacula has a comprehensive shell administrative interface, which allows the adminstrator to use tools such as ssh to administrate any part of Bacula from anywhere (even from home).
- Bacula has a Rescue CD for Linux systems with the following features:
  - You build it on your own system from scratch with one simple command: make – well, then make burn.
  - It uses your kernel
  - It captures your current disk parameters and builds scripts that allow you to automatically repartition a disk and format it to put it back to what you had before.
  - It has a script that will restart your networking (with the right IP address)
  - It has a script to automatically mount your hard disks.
  - It has a full Bacula FD statically linked
  - You can easily add additional data/programs, ... to the disk.

Current Implementation Restrictions

- It doesn't currently support ANSI and IBM tape labels.
- Typical of Microsoft, not all files can always be saved on WinNT, Win2K and WinXP when they are in use by another program. Anyone knowing the magic incantations please step forward. The files that are skipped seem to be in exclusive use by some other process, and don't appear to be too important.
- Unicode filenames (e.g. Chinese) cannot be saved or restored. This appears to be a problem only on Mac machines that are using remote mounted Windows volumes.

- If you have over 4 billion file entries stored in your database, the database FileId is likely to overflow. This is a monster database, but still possible. At some point, Bacula's FileId fields will be upgraded from 32 bits to 64 bits and this problem will go away. In the mean time, a good workaround is to use multiple databases.
- Files deleted after a Full save will be included in a restoration.
- Event handlers are not yet implemented (e.g. when Job terminates do this, ...)
- File System Modules (configurable routines for saving/restoring special files).
- Data encryption of the Volume contents.
- Bacula cannot automatically restore files for a single Job from two or more different storage devices or different media types. That is, if you use more than one storage device or media type to backup a single job, the restore process will require some manual intervention.
- There is no concept of a Pool of backup devices (i.e. if device /dev/nst0 is busy, use /dev/nst1, ...).

Design Limitations or Restrictions

- Names (resource names, Volume names, and such) defined in Bacula configuration files are limited to a fixed number of characters. Currently the limit is defined as 127 characters. Note, this does not apply to filenames, which may be arbitrarily long.

System Requirements

System Requirements

- **Bacula** has been compiled and run on Linux RedHat, FreeBSD, and Solaris systems.
- It requires GNU C++ version 2.95 or higher to compile. You can try with other compilers and older versions, but you are on your own. We have successfully compiled and used Bacula on RH8.0/RH9/RHEL 3.0 with GCC 3.2. Note, in general GNU C++ is a separate package (e.g. RPM) from GNU C, so you need them both loaded. On RedHat systems, the C++ compiler is part of the **gcc-c++** rpm package.
- There are certain third party packages that Bacula needs. Except for MySQL and PostgreSQL, they can all be found in the **depkgs** and **depkgs1** releases.
- If you want to build the Win32 binaries, you will need a Microsoft Visual C++ compiler (or Visual Studio). Although all components build (console has some warnings), only the File daemon has been tested.
- **Bacula** requires a good implementation of pthreads to work. This is not the case on some of the BSD systems.
- The source code has been written with portability in mind and is mostly POSIX compatible. Thus porting to any POSIX compatible operating system should be relatively easy.
- The GNOME Console program is developed and tested under GNOME 2.x. It also runs under GNOME 1.4 but this version is deprecated and thus no longer maintained.
- The wxWidgets Console program is developed and tested with the latest stable version of wxWidgets (2.4.2). It works fine with the Windows and GTK+-1.x version of wxWidgets, and should also works on other platforms supported by wxWidgets.
- The Tray Monitor program is developed for GTK+-2.x. It needs Gnome less or equal to 2.2, KDE greater or equal to 3.1 or any window manager supporting the FreeDesktop system tray standard.
- If you want to enable command line editing and history, you will need to have /usr/include/termcap.h and either the termcap or the ncurses library loaded (libtermcap-devel or ncurses-devel).

- If you want to use DVD as backup medium, you will need to download and install the dvd+rw-tools.

Supported Operating Systems

Supported Operating Systems

- Linux systems (built and tested on RedHat Enterprise Linux 3.0).
- If you have a recent Red Hat Linux system running the 2.4.x kernel and you have the directory **/lib/tls** installed on your system (normally by default), bacula will **NOT** run. This is the new pthreads library and it is defective. You must remove this directory prior to running Bacula, or you can simply change the name to **/lib/tls-broken**) then you must reboot your machine (one of the few times Linux must be rebooted). If you are not able to remove/rename /lib/tls, an alternative is to set the environment variable "LD_ASSUME_KERNEL=2.4.19" prior to executing Bacula. For this option, you do not need to reboot, and all programs other than Bacula will continue to use /lib/tls.

  The feedback that we have for 2.6 kernels is that the same problem exists. However, on 2.6 kernels, we would probably recommend using the environment variable override (LD_ASSUME_KERNEL=2.4.19) rather than removing /lib/tls.
- Most flavors of Linux (Gentoo, SuSE, Mandrake, Debian, ...).
- Solaris various versions.
- FreeBSD (tape driver supported in 1.30 – please see some **important** considerations in the Tape Modes on FreeBSD section of the Tape Testing chapter of this manual.)
- Windows (Win98/Me, WinNT/2K/XP) Client (File daemon) binaries.
- MacOS X/Darwin (see http://fink.sourceforge.net/ for obtaining the packages)
- OpenBSD Client (File daemon).
- Irix Client (File daemon).
- Tru64
- Bacula is said to work on other systems (AIX, BSDI, HPUX, ...) but we do not have first hand knowledge of these systems.
- RHat 7.2 AS2, AS3, AS4, Fedora Core 2, SuSE SLES 7,8,9 and Debian Woody and Sarge Linux on S/390 and Linux on zSeries.
- See the Porting chapter of the Bacula Developer's Guide for information on porting to other systems.

Supported Tape Drives

Supported Tape Drives

Even if your drive is on the list below, please check the Tape Testing Chapter of this manual for procedures that you can use to verify if your tape drive will work with Bacula. If your drive is in fixed block mode, it may appear to work with Bacula until you attempt to do a restore and Bacula wants to position the tape. You can be sure only by following the procedures suggested above and testing.

It is very difficult to supply a list of supported tape drives, or drives that are known to work with Bacula because of limited feedback (so if you use Bacula on a different drive, please let us know). Based on user feedback, the following drives are known to work with Bacula. A dash in a column means unknown:

OS	Man.	Media	Model
-	ADIC	DLT	Adic Scalar 100 DLT
-	ADIC	DLT	Adic Fastor 22 DLT
-	-	DDS	Compaq DDS 2,3,4
-	Exabyte	-	Exabyte drives less than 10 years
-	Exabyte	-	Exabyte VXA drives
-	HP	Travan 4	Colorado T4000S
-	HP	DLT	HP DLT drives
-	HP	LTO	HP LTO Ultrium drives
-	IBM	??	3480, 3480XL, 3490, 3490E, 3580 3590 drives
FreeBSD 4.10 RELEASE	HP	DAT	HP StorageWorks DAT72i
-	Overland	LTO	LoaderXpress LTO
-	Overland	-	Neo2000
-	OnStream	-	OnStream drives (see below)
-	Quantum	DLT	DLT-8000
Linux	Seagate	DDS-4	Scorpio 40
FreeBSD 4.9 STABLE	Seagate	DDS-4	STA2401LW
FreeBSD 5.2.1 pthreads patched RELEASE	Seagate	AIT-1	STA1701W
Linux	Sony	DDS-2,3,4	-
Linux	Tandberg	-	Tandbert MLR3
FreeBSD	Tandberg	-	Tandberg SLR6
Solaris	Tandberg	-	Tandberg SLR75

There is a list of supported autochangers in the Supported Autochangers chapter of this document, where you will find other tape drives that work with Bacula.

Unsupported Tape Drives

Previously OnStream IDE-SCSI tape drives did not work with Bacula. As of Bacula version 1.33 and the osst kernel driver version 0.9.14 or later, they now work. Please see the testing chapter as you must set a fixed block size.

QIC tapes are known to have a number of particularities (fixed block size, and one EOF rather than two to terminate the tape). As a consequence, you will need to take a lot of care in configuring them to make them work correctly with Bacula.

FreeBSD Users Be Aware!!!

Unless you have patched the pthreads library on most FreeBSD systems, you will lose data when Bacula spans tapes. This is because the unpatched pthreads library fails to return a warning status to Bacula that the end of the tape is near. Please see the Tape Testing Chapter of this manual for **important** information on how to configure your tape drive for compatibility with Bacula.

Supported Autochangers

For information on supported autochangers, please see the Autochangers Known to Work with Bacula section of the Supported Autochangers chapter of this manual.

Getting Started with Bacula

If you are like me, you want to get Bacula running immediately to get a feel for it, then later you want to go back and read about all the details. This chapter attempts to accomplish just that: get you going quickly without all the details. If you want to skip the section on Pools, Volumes and Labels, you can always come back to it, but please read to the end of this chapter, and in particular follow the instructions for testing your tape drive.

We assume that you have managed to build and install Bacula, if not, you might want to first look at the System Requirements then at the Compiling and Installing Bacula chapter of this manual.

Understanding Pools, Volumes and Labels

If you have been using a program such as **tar** to backup your system, Pools, Volumes, and labeling may be a bit confusing at first. A Volume is a single physical tape (or possibly a single file) on which Bacula will write your backup data. Pools group together Volumes so that a backup is not restricted to the length of a single Volume (tape). Consequently, rather than explicitly naming Volumes in your Job, you specify a Pool, and Bacula will select the next appendable Volume from the Pool and request you to mount it.

Although the basic Pool options are specified in the Director's Pool resource, the **real** Pool is maintained in the Bacula Catalog. It contains information taken from the Pool resource (bacula-dir.conf) as well as information on all the Volumes that have been added to the Pool. Adding Volumes to a Pool is usually done manually with the Console program using the **label** command.

For each Volume, Bacula maintains a fair amount of catalog information such as the first write date/time, the last write date/time, the number of files on the Volume, the number of bytes on the Volume, the number of Mounts, etc.

Before Bacula will read or write a Volume, the physical Volume must have a Bacula software label so that Bacula can be sure the correct Volume is mounted. This is usually done using the **label** command in the Console program.

The steps for creating a Pool, adding Volumes to it, and writing software labels to the Volumes, may seem tedious at first, but in fact, they are quite simple to do, and they allow you to use multiple Volumes (rather than

being limited to the size of a single tape). Pools also give you significant flexibility in your backup process. For example, you can have a "Daily" Pool of Volumes for Incremental backups and a "Weekly" Pool of Volumes for Full backups. By specifying the appropriate Pool in the daily and weekly backup Jobs, you thereby insure that no daily Job ever writes to a Volume in the Weekly Pool and vice versa, and Bacula will tell you what tape is needed and when.

For more on Pools, see the Pool Resource section of the Director Configuration chapter, or simply read on, and we will come back to this subject later.

Setting Up Bacula Configuration Files

After running the appropriate **./configure** command and doing a **make**, and a **make install**, if this is the first time you are running Bacula, you must create valid configuration files for the Director, the File daemon, the Storage daemon, and the Console programs. If you have followed our recommendations, default configuration files as well as the daemon binaries will be located in your installation directory. In any case, the binaries are found in the directory you specified on the **--sbindir** option to the **./configure** command, and the configuration files are found in the directory you specified on the **--sysconfdir** option.

When initially setting up Bacula you will need to invest a bit of time in modifying the default configuration files to suit your environment. This may entail starting and stopping Bacula a number of times until you get everything right. Please do not despair. Once you have created your configuration files, you will rarely need to change them nor will you stop and start Bacula very often. Most of the work will simply be in changing the tape when it is full.

Configuring the Console Program

The Console program is used by the administrator to interact with the Director and to manually start/stop Jobs or to obtain Job status information.

The Console configuration file is found in the directory specified on the **--sysconfdir** option that you specified on the **./configure** command and by default is named **console.conf**.

If you choose to build the GNOME console with the **--enable-gnome** option, you also find a default configuration file for it, named **gnome-console.conf**.

The same applies to the wxWidgets console, which is build with the **--enable-wx-console** option, and the name of the default configuration file is, in this case, **wx-console.conf**.

Normally, for first time users, no change is needed to these files. Reasonable defaults are set.

Configuring the Monitor Program

The Monitor program is typically an icon in the system tray. However, once the icon is expanded into a full window, the administrator or user can obtain status information about the Director or the backup status on the local workstation or any other Bacula daemon that is configured.

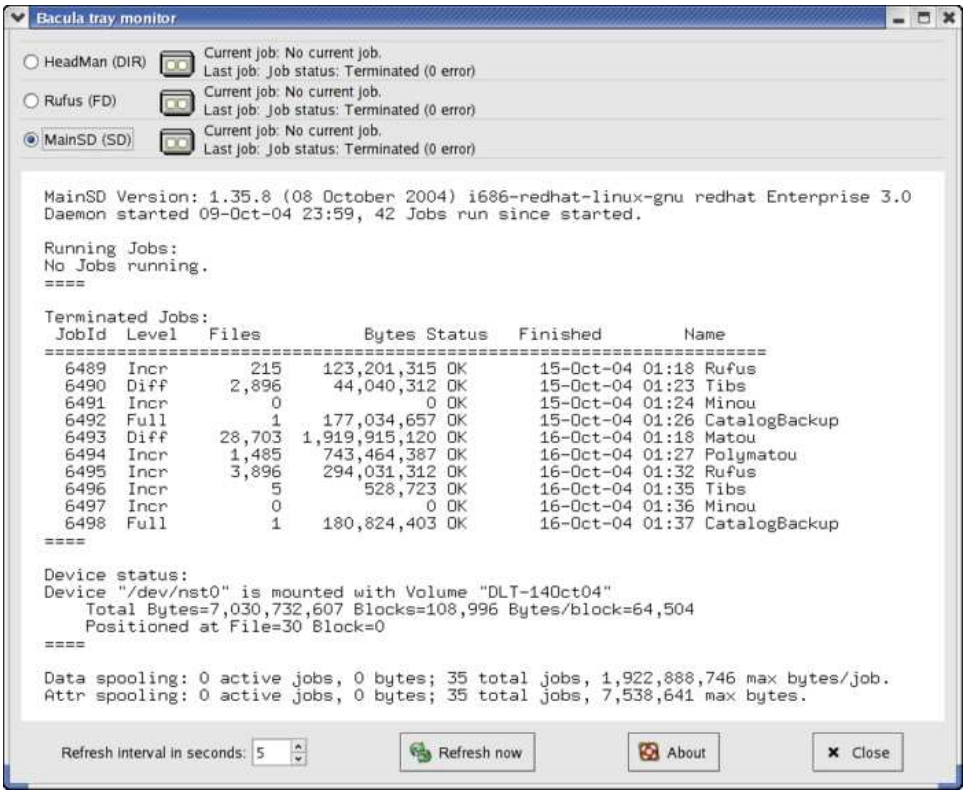

The image shows a tray-monitor configured for three daemons. By clicking on the radio buttons in the upper left corner of the image, you can see the status for each of the daemons. The image shows the status for the Storage

daemon (MainSD) that is currently selected.

The Monitor configuration file is found in the directory specified on the **--sysconfdir** option that you specified on the **./configure** command and by default is named **tray-monitor.conf**. Normally, for first time users, you just need to change the permission of this file to allow non-root users to run the Monitor, as this application must run as the same user as the graphical environment (don't forget allow non-root users to execute **bacula-tray-monitor**). This is not a security problem as long as you use the default settings.

Configuring the File daemon

The File daemon is a program that runs on each (Client) machine. At the request of the Director, finds the files to be backed up and sends them (their data) to the Storage daemon.

The File daemon configuration file is found in the directory specified on the **--sysconfdir** option that you specified on the **./configure** command. By default, the File daemon's configuration file is named **bacula-fd.conf**. Normally, for first time users, no change is needed to this file. Reasonable defaults are set. However, if you are going to back up more than one machine, you will need to install the File daemon with a unique configuration file on each machine to be backed up. The information about each File daemon must appear in the Director's configuration file.

Configuring the Director

The Director is the central control program for all the other daemons. It schedules and monitors all jobs to be backed up.

The Director configuration file is found in the directory specified on the **--sysconfdir** option that you specified on the **./configure** command. Normally the Director's configuration file is named **bacula-dir.conf**.

In general, the only change you must make is modify the FileSet resource so that the **Include** configuration directive contains at least one line with a valid name of a directory (or file) to be saved.

If you do not have a DLT tape drive, you will probably want to edit the Storage resource to contain names that are more representative of your actual storage device. You can always use the existing names as you are free to arbitrarily assign them, but they must agree with the corresponding

names in the Storage daemon's configuration file.

You may also want to change the email address for notification from the default **root** to your email address.

Finally, if you have multiple systems to be backed up, you will need a separate File daemon or Client specification for each system, specifying its name, address, and password. We have found that giving your daemons the same name as your system but post fixed with **-fd** helps a lot in debugging. That is, if your system name is **foobaz**, you would give the File daemon the name **foobaz-fd**. For the Director, you might use **foobaz-dir**, and for the storage daemon, you might use **foobaz-sd**.

Configuring the Storage daemon

The Storage daemon is responsible, at the Director's request, for accepting data from a File daemon and placing it on Storage media, or in the case of a restore request, to find the data and send it to the File daemon.

The Storage daemon's configuration file is found in the directory specified on the **--sysconfdir** option that you specified on the **./configure** command. By default, the Storage daemon's file is named **bacula-sd.conf**. Edit this file to contain the correct Archive device names for any tape devices that you have. If the configuration process properly detected your system, they will already be correctly set. These Storage resource name and Media Type must be the same as the corresponding ones in the Director's configuration file **bacula-dir.conf**. If you want to backup to a file instead of a tape, the Archive device must point to a directory in which the Volumes will be created as files when you label the Volume.

Testing your Configuration Files

You can test if your configuration file is syntactically correct by running the appropriate daemon with the **-t** option. The daemon will process the configuration file and print any error messages then terminate. For example, assuming you have installed your binaries and configuration files in the same directory.

```
cd <installation-directory>
./bacula-dir -t -c bacula-dir.conf
./bacula-fd -t -c bacula-fd.conf
./bacula-sd -t -c bacula-sd.conf
./bconsole -t -c bconsole.conf
```

```
./gnome-console -t -c gnome-console.conf
./wx-console -t -c wx-console.conf
su <normal user> -c "./bacula-tray-monitor -t -c tray-monitor.conf"
```

will test the configuration files of each of the main programs. If the configuration file is OK, the program will terminate without printing anything. Please note that, depending on the configure options you choose, some, or even all, of the three last commands will not be available on your system. If you have installed the binaries in traditional Unix locations rather than a single file, you will need to modify the above commands appropriately (no ./ in front of the command name, and a path in front of the conf file name).

Testing Bacula Compatibility with Your Tape Drive

Before spending a lot of time on Bacula only to find that it doesn't work with your tape drive, please read the btape – Testing Your Tape Drive chapter of this manual. If you have a modern standard SCSI tape drive on a Linux or Solaris, most likely it will work, but better test than be sorry. For FreeBSD (and probably other xBSD flavors), reading the above mentioned tape testing chapter is a must. Also, for FreeBSD, please see The FreeBSD Diary for a detailed description on how to make Bacula work on your system. In addition, users of FreeBSD prior to 4.9-STABLE dated Mon Dec 29 15:18:01 2003 UTC who plan to use tape devices, please see the file **platforms/freebsd/pthreads-fix.txt** in the main Bacula directory concerning important information concerning compatibility of Bacula and your system.

Get Rid of the /lib/tls Directory

The new pthreads library **/lib/tls** installed by default on recent Red Hat systems running kernel 2.4.x is defective. You must remove it or rename it then reboot your system before running Bacula otherwise after a week or so of running, Bacula will either block for long periods or deadlock entirely. The feedback that we have concerning 2.6 kernels is the same. However, on 2.6 systems, you may want to use the loader environment variable override rather than removing /lib/tls. Please see Supported Operating Systems for more information on this problem.

Running Bacula

Probably the most important part of running Bacula is being able to restore files. If you haven't tried recovering files at least once, when you actually have to do it, you will be under a lot more pressure, and prone to make errors, than if you had already tried it once.

To get a good idea how to use Bacula in a short time, we **strongly** recommend that you follow the example in the Running Bacula Chapter of this manual where you will get detailed instructions on how to run Bacula.

Log Rotation

If you use the default **bacula-dir.conf** or some variation of it, you will note that it logs all the Bacula output to a file. To avoid that this file grows without limit, we recommend that you copy the file **logrotate** from the **scripts/logrotate** to **/etc/logrotate.d/bacula**. This will cause the log file to be rotated once a month and kept for a maximum of 5 months. You may want to edit this file to change the default log rotation preferences.

Disaster Recovery

If you intend to use Bacula as a disaster recovery tool rather than simply a program to restore lost or damaged files, you will want to read the Disaster Recovery Using Bacula Chapter of this manual.

In any case, you are strongly urged to carefully test restoring some files that you have saved rather than wait until disaster strikes. This way, you will be prepared.

Installing Bacula

General

In general, you will need the Bacula source release, and if you want to run a Windows client, you will need the Bacula Windows binary release. However, Bacula needs certain third party packages (such as **SQLite**, **MySQL** to build properly depending on the options you specify. To simplify your task, we have combined a number of these packages into two **depkgs** releases (Dependency Packages). This can vastly simplify your life by providing you with all the necessary packages rather than requiring you to find them on the Web, load them, and install them.

Upgrading Bacula

If you are upgrading from one Bacula version to another, you should first carefully read the ReleaseNotes of all versions between your current version and the version to which you are upgrading. If the Bacula catalog database has been upgraded, you will either need to reinitialize your database starting from scratch, or save an ASCII copy of your database, then proceed to upgrade it. If there are several database upgrades between your version and the version to which you are upgradding, you will need to apply each database upgrade script. For your convenience, you can find all the old upgrade scripts in the **upgradedb** directory of the source code. You will need to edit the scripts to correspond to your system configuration. The final upgrade script, if any, will be in the **src/cats** directory as described in the ReleaseNotes.

If you are upgrading from one major version to another, you will need to replace all your components at the same time as generally the inter-daemon protocol will change. However, within any particular release (e.g. version 1.32.x) unless there is an oversight or bug, the daemon protocol will not change. If this is confusing, simply read the ReleaseNotes very carefully as they will note if all daemons must be upgraded at the same time.

Finally, please note that in general it is not necessary to do a **make uninstall** before doing an upgrade. In fact, if you do so, you will most likely delete all your conf files, which could be disasterous. For additional information on upgrading, please see the Upgrading Bacula Versions in the Tips chapter of this manual.

Dependency Packages

As discussed above, we have combined a number of third party packages that Bacula might need into the **depkgs** and **depkgs1** releases. You can, of course, get the latest packages from the original authors. The locations of where we obtained the packages are in the README file in each package. However, be aware that the packages in the depkgs files have been tested by us for compatibility with Bacula.

Typically, a dependency package will be named **depkgs-ddMMMyy.tar.gz** and **depkgs1-ddMMyy.tar.gz** where **dd** is the day we release it, **MMM** is the abbreviated month (e.g. Jan), and **yy** is the year. An actual example is: **depkgs-07Apr02.tar.gz**. To install and build this package (if needed), you do the following:

1. Create a **bacula** directory, into which you will place both the Bacula source as well as the dependency package.
2. Detar the **depkg** into the **bacula** directory.
3. cd bacula/depkgs
4. make

Although the exact composition of the dependency packages may change from time to time, the current makeup is the following:

3rd Party Package	depkgs	depkgs1	depkgs-win32
SQLite	X	-	-
mtx	X	-	-
readline	-	X	-
pthreads	-	-	X
zlib	-	-	X
wxWidgits	-	-	X

Note, some of these packages are quite large, so that building them can be a bit time consuming. The above instructions will build all the packages contained in the directory. However, when building Bacula, it will take only those pieces that it actually needs.

Alternatively, you can make just the packages that are needed. For example,

```
cd bacula/depkgs
```

```
make sqlite
```

will configure and build only the SQLite package.

You should build the packages that you will require in **depkgs** and/or **depkgs1** prior to configuring and building Bacula, since Bacula will need them during the build process.

Even if you do not use SQLite, you might find it worthwhile to build **mtx** because the **tapeinfo** program that comes with it can often provide you with valuable information about your SCSI tape drive (e.g. compression, min/max block sizes, ...).

The **depkgs-win32** package contains the source code for the pthreads and zlib libraries used by the native Win32 client program. It will only be needed if you intend to build the Win32 client from source.

Supported Operating Systems

Please see the Supported Operating Systems section of the QuickStart chapter of this manual.

Building Bacula from Source

The basic installation is rather simple.

1. Install and build any **depkgs** as noted above.
2. Configure and install MySQL or PostgreSQL (if desired). Installing and Configuring MySQL Phase I or Installing and Configuring PostgreSQL Phase I. If you are installing from rpms, and are using MySQL, please be sure to install **mysql-devel**, so that the MySQL header files are available while compiling Bacula. In addition, the MySQL client library **mysqlclient** requires the gzip compression library **libz.a** or **libz.so**. If you are using rpm packages, these libraries are in the **libz-devel** package. On Debian systems, you will need to load the **zlib1g-dev** package. If you are not using rpms or debs, you will need to find the appropriate package for your system. Note, if you already have a running MySQL or PostgreSQL on your system, you can skip this phase provided that you have built the thread safe libraries. And you have already installed the additional rpms noted above.

3. As an alternative to MySQL and PostgreSQL, configure and install SQLite, which is part of the **depkgs**. Installing and Configuring SQLite.

4. Detar the Bacula source code preferably into the **bacula** directory discussed above.

5. **cd** to the directory containing the source code.

6. ./configure (with appropriate options as described below)

7. Check the output of ./configure very carefully, especially the Install binaries and Install config files directories. If they are not correct, please rerun ./configure until they are. The output from ./configure is stored in **config.out** and can be re-displayed at any time without rerunning the ./configure by doing **cat config.out**.

8. If after running ./configure once, you decide to change options and re-run it, that is perfectly fine, but before re-running it, you should run:

   ```
   make distclean
   ```

   so that you are sure to start from scratch and not have a mixture of the two options. This is because ./configure caches much of the information. The **make distclean** is also critical if you move the source file from one machine to another. If the **make distclean** fails, just ignore it and continue on.

9. make

   If you get errors while linking in the Storage daemon directory (src/stored), it is probably because you have not loaded the static libraries on your system. I noticed this problem on a Solaris system. To correct it, make sure that you have not added **--enable-static-tools** to the **./configure** command.

10. make install

11. If you are new to Bacula, we **strongly** recommend that you skip the next step and use the default configuration files, then run the example program in the next chapter, then come back and modify your configuration files to suit your particular needs.

12. Customize the configuration files for each of the three daemons (Directory, File, Storage) and for the Console program. For the details of how to do this, please see Setting Up Bacula Configuration Files in

the Configuration chapter of this manual. We recommend that you start by modifying the default configuration files supplied, making the minimum changes necessary. Complete customization can be done after you have Bacula up and running. Please take care when modifying passwords, which were randomly generated, and the **Name**s as the passwords and names must agree between the configuration files for security reasons.

13. Create the Bacula MySQL database and tables (if using MySQL) Installing and Configuring MySQL Phase II or create the Bacula PostgreSQL database and tables Installing and Configuring PostgreSQL Phase II or alternatively if you are using SQLite Installing and Configuring SQLite Phase II.

14. Start Bacula (**./bacula start**) Note. the next chapter shows you how to do this in detail.

15. Interface with Bacula using the Console program

16. For the previous two items, please follow the instructions in the Running Bacula chapter of this manual, where you will run a simple backup and do a restore. Do this before you make heavy modifications to the configuration files so that you are sure that Bacula works and are familiar with it. After that changing the conf files will be easier.

17. If after installing Bacula, you decide to "move it", that is to install it in a different set of directories, proceed as follows:

```
make uninstall
make distclean
./configure (your-new-options)
make
make install
```

If all goes well, the **./configure** will correctly determine which operating system you are running and configure the source code appropriately. Currently, FreeBSD, Linux (RedHat), and Solaris are supported. MacOS X 10.3 is reported to work with the Client only as long as readline support is disabled.

If you install Bacula on more than one system, and they are identical, you can simply transfer the source tree to that other system and do a "make install". However, if there are differences in the libraries or OS versions, or you wish to install on a different OS, you should start from the original compress tar file. If you do transfer the source tree, and you have previously done a ./configure command, you MUST do:

```
make distclean
```

prior to doing your new ./configure. This is because the GNU autoconf tools cache the configuration, and if you re-use a configuration for a Linux machine on a Solaris, you can be sure your build will fail. To avoid this, as mentioned above, either start from the tar file, or do a "make distclean".

In general, you will probably want to supply a more complicated **configure** statement to ensure that the modules you want are built and that everything is placed into the correct directories.

For example, on RedHat, one could use the following:

```
CFLAGS="-g -Wall" \
  ./configure \
    --sbindir=$HOME/bacula/bin \
    --sysconfdir=$HOME/bacula/bin \
    --with-pid-dir=$HOME/bacula/bin/working \
    --with-subsys-dir=$HOME/bacula/bin/working \
    --with-mysql=$HOME/mysql \
    --with-working-dir=$HOME/bacula/bin/working \
    --with-dump-email=$USER
```

Note, the advantage of using the above configuration to start is that everything will be put into a single directory, which you can later delete once you have run the examples in the next chapter and learned how Bacula works. In addition, the above can be installed and run as non-root.

For the developer's convenience, I have added a **defaultconfig** script to the **examples** directory. This script contains the statements that you would normally use, and each developer/user may modify them to suit his needs. You should find additional useful examples in this directory as well.

The **--enable-conio** or **--enable-readline** options are useful because they provide a command line history and editing capability for the Console program. If you have included either option in the build, either the **termcap** or the **ncurses** package will be needed to link. On some systems, such as SuSE, the termcap library is not in the standard library directory. As a consequence, the option may be disabled or you may get an error message such as:

```
/usr/lib/gcc-lib/i586-suse-linux/3.3.1/.../ld:
cannot find -ltermcap
collect2: ld returned 1 exit status
```

while building the Bacula Console. In that case, you will need to set the **LDFLAGS** environment variable prior to building.

```
export LDFLAGS="-L/usr/lib/termcap"
```

The same library requirements apply if you wish to use the readline subroutines for command line editing and history or if you are using a MySQL library that requires encryption. If you need encryption, you can either export the appropriate additional library options as shown above or, alternatively, you can include them directly on the ./configure line as in:

```
LDFLAGS="-lssl -lcyrpto" \
   ./configure \
      <your-options>
```

On some systems such as Mandrake, readline tends to gobble up prompts, which makes it totally useless. If this happens to you, use the disable option, or if you are using version 1.33 and above try using **--enable-conio** to use a built-in readline replacement. You will still need the either termcap or ncurses library, but it is unlikely that the **conio** package will gobble up prompts.

readline is no longer supported after version 1.34. The code is still available and if users submit patches for it, I will be happy to apply them. However, due to the fact that each version of readline seems to be incompatible with previous versions, and that there are significant differences between systems, I can no longer afford to support it.

What Database to Use?

Before building Bacula you need to decide if you want to use SQLite, MySQL, or PostgreSQL. If you are not already running MySQL or PostgreSQL, we recommend that you start by using SQLite. This will greatly simplify the setup for you because SQLite is compiled into Bacula an requires no administration. It performs well and is suitable for small to medium sized installations (maximum 10-20 machines).

If you wish to use MySQL as the Bacula catalog, please see the Installing and Configuring MySQL chapter of this manual. You will need to install MySQL prior to continuing with the configuration of Bacula. MySQL is a high quality database that is very efficient and is suitable for any sized installation. It is slightly more complicated than SQLite to setup

and administer because it has a number of sophisticated features such as userids and passwords. It runs as a separate process, is truly professional and can manage a database of any size.

If you wish to use PostgreSQL as the Bacula catalog, please see the Installing and Configuring PostgreSQL chapter of this manual. You will need to install PostgreSQL prior to continuing with the configuration of Bacula. PostgreSQL is very similar to MySQL, though it tends to be slightly more SQL92 compliant and has many more advanced features such as transactions, stored procedures, and the such. It requires a certain knowledge to install and maintain. There are some important performance problems with PostgreSQL in Bacula versions prior to 1.35.5.

If you wish to use SQLite as the Bacula catalog, please see Installing and Configuring SQLite chapter of this manual.

Quick Start

There are a number of options and important considerations given below that you can skip for the moment if you have not had any problems building Bacula with a simplified configuration as shown above.

If you want to dive right into it, we recommend you skip to the next chapter, and run the example program. It will teach you a lot about Bacula and as an example can be installed into a single directory (for easy removal) and run as non-root. If you have any problems or when you want to do a real installation, come back to this chapter and read the details presented below.

Configure Options

The following command line options are available for **configure** to customize your installation.

> For more information on configuring and testing TCP wrappers, please see the Configuring and Testing TCP Wrappers section in the Security Capter.

Note, many other options are presented when you do a **./configure --help**, but they are not implemented.

Recommended Options for most Systems

For most systems, we recommend starting with the following options:

```
./configure \
  --enable-smartalloc \
  --sbindir=$HOME/bacula/bin \
  --sysconfdir=$HOME/bacula/bin \
  --with-pid-dir=$HOME/bacula/bin/working \
  --with-subsys-dir=$HOME/bacula/bin/working \
  --with-mysql=$HOME/mysql \
  --with-working-dir=$HOME/bacula/working
```

If you want to install Bacula in an installation directory rather than run it out of the build directory (as developers will do most of the time), you should also include the --sbindir and --sysconfdir options with appropriate paths. Neither are necessary if you do not use "make install" as is the case for most development work. The install process will create the sbindir and sysconfdir if they do not exist, but it will not automatically create the pid-dir, subsys-dir, or working-dir, so you must ensure that they exist before running Bacula for the first time. See below for an example of how Kern does it.

RedHat

Using SQLite:

```
CFLAGS="-g -Wall" ./configure \
  --sbindir=$HOME/bacula/bin \
  --sysconfdir=$HOME/bacula/bin \
  --enable-smartalloc \
  --with-sqlite=$HOME/bacula/depkgs/sqlite \
  --with-working-dir=$HOME/bacula/working \
  --with-pid-dir=$HOME/bacula/bin/working \
  --with-subsys-dir=$HOME/bacula/bin/working \
  --enable-gnome \
  --enable-conio
```

or

```
CFLAGS="-g -Wall" ./configure \
  --sbindir=$HOME/bacula/bin \
```

```
  --sysconfdir=$HOME/bacula/bin \
  --enable-smartalloc \
  --with-mysql=$HOME/mysql \
  --with-working-dir=$HOME/bacula/working
  --with-pid-dir=$HOME/bacula/bin/working \
  --with-subsys-dir=$HOME/bacula/bin/working
  --enable-gnome \
  --enable-conio
```

or finally, a completely traditional RedHat Linux install:

```
CFLAGS="-g -Wall" ./configure \
  --prefix=/usr \
  --sbindir=/usr/sbin \
  --sysconfdir=/etc/bacula \
  --with-scriptdir=/etc/bacula \
  --enable-smartalloc \
  --enable-gnome \
  --with-mysql \
  --with-working-dir=/var/bacula \
  --with-pid-dir=/var/run \
  --with-subsys-dir=/var/lock/subsys \
  --enable-conio
```

Note, Bacula assumes that /var/bacula, /var/run, and /var/loc/subsys exist so it will not automatically create them during the install process.

Solaris

To build Bacula from source, you will need the following installed on your system (they are not by default): libiconv, gcc 3.3.2, stdc++, libgcc (for stdc++ and gcc_s libraries), make 3.8 or later.

You will probably also need to: Add /usr/local/bin to PATH and Add /usr/ccs/bin to PATH for ar.

```
#!/bin/sh
CFLAGS="-g" ./configure \
  --sbindir=$HOME/bacula/bin \
  --sysconfdir=$HOME/bacula/bin \
  --with-mysql=$HOME/mysql \
  --enable-smartalloc \
  --with-pid-dir=$HOME/bacula/bin/working \
  --with-subsys-dir=$HOME/bacula/bin/working \
  --with-working-dir=$HOME/bacula/working
```

As mentioned above, the install process will create the sbindir and sysconfdir if they do not exist, but it will not automatically create the pid-dir, subsys-dir, or working-dir, so you must ensure that they exist before running Bacula for the first

FreeBSD

Please see: The FreeBSD Diary for a detailed description on how to make Bacula work on your system. In addition, users of FreeBSD prior to 4.9-STABLE dated Mon Dec 29 15:18:01 2003 UTC who plan to use tape devices, please see the Tape Testing Chapter of this manual for **important** information on how to configure your tape drive for compatibility with Bacula.

If you are using Bacula with MySQL, you should take care to compile MySQL with FreeBSD native threads rather than LinuxThreads, since Bacula is normal built with FreeBSD native threads rather than LinuxTreads. Mixing the two will probably not work.

Win32

To install the binary Win32 version of the File daemon please see the Win32 Installation Chapter in this document.

Windows Systems with CYGWIN Installed

As of version 1.34, Bacula no longer uses CYGWIN for the Win32 File daemon. However, it is still built under a CYGWIN build environment – though you can probably do it with VC Studio only. If you wish to build the Win32 File daemon from the source, you will need Microsoft C++ version 6.0 or greater. In Bacula prior to version 1.33, CYGWIN was used. Details for building it are in the README file of the src/win32 directory.

Note, although most parts of Bacula build on Windows systems, the only part that we have tested and used is the File daemon.

Finally, you should follow the installation instructions in the Win32 Installation section of this document.

Kern's Configure Script

The script that I use for building on my "production" Linux machines is:

```
#!/bin/sh
# This is Kern's configure script for Bacula
CFLAGS="-g -Wall" \
  ./configure \
    --sbindir=$HOME/bacula/bin \
    --sysconfdir=$HOME/bacula/bin \
    --enable-smartalloc \
    --enable-gnome \
    --with-pid-dir=$HOME/bacula/bin/working \
    --with-subsys-dir=$HOME/bacula/bin/working \
    --with-mysql=$HOME/mysql \
    --with-working-dir=$HOME/bacula/bin/working \
    --with-dump-email=$USER \
    --with-smtp-host=mail.your-site.com \
    --with-baseport=9101
exit 0
```

Note that I define the base port as 9101, which means that Bacula will use port 9101 for the Director console, port 9102 for the File daemons, and port 9103 for the Storage daemons. These ports should be available on all systems because they have been officially assigned to Bacula by IANA (Internet Assigned Numbers Authority). We strongly recommend that you use only these ports to prevent any conflicts with other programs. This is in fact the default if you do not specify a **--with-baseport** option.

You may also want to put the following entries in your **/etc/services** file as it will make viewing the connections made by Bacula easier to recognize (i.e. netstat -a):

```
bacula-dir      9101/tcp
bacula-fd       9102/tcp
bacula-sd       9103/tcp
```

Installing Bacula

Before setting up your configuration files, you will want to install Bacula in its final location. Simply enter:

```
make install
```

If you have previously installed Bacula, the old binaries will be overwritten, but the old configuration files will remain unchanged, and the "new" configuration files will be appended with a **.new**. Generally if you have previously installed and run Bacula you will want to discard or ignore the configuration files with the appended **.new**.

Building a File Daemon or Client

If you run the Director and the Storage daemon on one machine and you wish to back up another machine, you must have a copy of the File daemon for that machine. If the machine and the Operating System are identical, you can simply copy the Bacula File daemon binary file **bacula-fd** as well as its configuration file **bacula-fd.conf** then modify the name and password in the conf file to be unique. Be sure to make corresponding additions to the Director's configuration file (**bacula-dir.conf**).

If the architecture or the O/S level are different, you will need to build a File daemon on the Client machine. To do so, you can use the same **./configure** command as you did for your main program, starting either from a fresh copy of the source tree, or using **make distclean** before the **./configure**.

Since the File daemon does not access the Catalog database, you can remove the **--with-mysql** or **--with-sqlite** options, then add **--enable-client-only**. This will compile only the necessary libraries and the client programs and thus avoids the necessity of installing one or another of those database programs to build the File daemon. With the above option, you simply enter **make** and just the client will be built.

Auto Starting the Daemons

If you wish the daemons to be automatically started and stopped when your system is booted (a good idea), one more step is necessary. First, the ./configure process must recognize your system – that is it must be a supported platform and not **unknown**, then you must install the platform dependent files by doing:

```
(become root)
make install-autostart
```

Please note, that the auto-start feature is implemented only on systems that we officially support (currently, FreeBSD, RedHat Linux, and Solaris), and

has only been fully tested on RedHat Linux.

The **make install-autostart** will cause the appropriate startup scripts to be installed with the necessary symbolic links. On RedHat Linux systems, these scripts reside in **/etc/rc.d/init.d/bacula-dir /etc/rc.d/init.d/bacula-fd**, and **/etc/rc.d/init.d/bacula-sd**. However the exact location depends on what operating system you are using.

If you only wish to install the File daemon, you may do so with:

```
make install-autostart-fd
```

Other Make Notes

To simply build a new executable in any directory, enter:

```
make
```

To clean out all the objects and binaries (including the files named 1, 2, or 3, which Kern uses as temporary files), enter:

```
make clean
```

To really clean out everything for distribution, enter:

```
make distclean
```

note, this cleans out the Makefiles and is normally done from the top level directory to prepare for distribution of the source. To recover from this state, you must redo the **./configure** in the top level directory, since all the Makefiles will be deleted.

To add a new file in a subdirectory, edit the Makefile.in in that directory, then simply do a **make**. In most cases, the make will rebuild the Makefile from the new Makefile.in. In some case, you may need to issue the **make** a second time. In extreme cases, cd to the top level directory and enter: **make Makefiles**.

To add dependencies:

```
make depend
```

The **make depend** appends the header file dependencies for each of the object files to Makefile and Makefile.in. This command should be done in each directory where you change the dependencies. Normally, it only needs to be run when you add or delete source or header files. **make depend** is normally automatically invoked during the configuration process.

To install:

```
make install
```

This not normally done if you are developing Bacula, but is used if you are going to run it to backup your system.

After doing a **make install** the following files will be installed on your system (more or less). The exact files and location (directory) for each file depends on your **./configure** command (e.g. gnome-console and gnome-console.conf are not installed if you do not configure GNOME. Also, if you are using SQLite instead of mysql, some of the files will be different).

```
bacula
bacula-dir
bacula-dir.conf
bacula-fd
bacula-fd.conf
bacula-sd
bacula-sd.conf
bacula-tray-monitor
tray-monitor.conf
bextract
bls
bscan
btape
btraceback
btraceback.gdb
bconsole
bconsole.conf
create_mysql_database
dbcheck
delete_catalog_backup
drop_bacula_tables
drop_mysql_tables
fd
gnome-console
gnome-console.conf
make_bacula_tables
make_catalog_backup
make_mysql_tables
mtx-changer
query.sql
```

```
bsmtp
startmysql
stopmysql
wx-console
wx-console.conf
```

Installing Tray Monitor

The Tray Monitor is already installed if you used the **--enable-tray-monitor** configure option and ran **make install**.

As you don't run your graphical environment as root (if you do, you should change that bad habit), don't forget to allow your user to read **tray-monitor.conf**, and to execute **bacula-tray-monitor** (this is not a security issue).

Then log into your graphical environment (KDE, Gnome or something else), run **bacula-tray-monitor** as your user, and see if a cassette icon appear somewhere on the screen, usually on the task bar. If it doesn't, follow the instructions below related to your environment or window manager.

GNOME

System tray, or notification area if you use the GNOME terminology, has been supported in GNOME since version 2.2. To activate it, right-click on one of your panels, open the menu **Add to this Panel**, then **Utility** and finally click on **Notification Area**.

KDE

System tray has been supported in KDE since version 3.1. To activate it, right-click on one of your panels, open the menu **Add**, then **Applet** and finally click on **System Tray**.

Other window managers

Read the documentation to know if the Freedesktop system tray standard is supported by your window manager, and if applicable, how to activate it.

Modifying the Bacula Configuration Files

See the chapter Configuring Bacula in this manual for instructions on how to set Bacula configuration files.

Critical Items to Implement Before Going Production

General

We recommend you take your time before implementing a Bacula backup system since Bacula is a rather complex program, and if you make a mistake, you may suddenly find that you cannot restore the your files in case of a disaster. This is especially true if you have not previously used a major backup product.

If you follow the instructions in this chapter, you will have covered most of the major problems that can occur. It goes without saying that if ever you find that we have left out an important point, please point it out to us, so that we can document it to the benefit of everyone.

Critical Items

The following assumes that you have installed Bacula, you more or less understand it, you have at least worked through the tutorial or have equivalent experience, and that you have setup a basic production configuration. If you haven't done the above, please do so then come back here. The following is a sort of checklist that points you elsewhere in the manual with perhaps a brief explaination of why you should do it. The order is more or less the order you would use in setting up a production system (if you already are in production, use the checklist anyway).

- Test your tape drive with compatibility with Bacula by using the test command in the btape program.
- Better than doing the above is to walk through the nine steps in the Tape Testing chapter of the manual. It may take you a bit of time, but it will eliminate surprises.
- Make sure that /lib/tls is disabled. Bacula does not work with this library. See the second point under Supported Operating Systems.
- Do at least one restore of files. If you backup both Unix and Win32 systems, restore files from each system type. The Restoring Files chapter shows you how.
- Write a bootstrap file to a separate system for each backup job. The Write Bootstrap directive is described in the Director Configuration

chapter of the manual, and more details are available in the Bootstrap File chapter. Also, the default bacula-dir.conf comes with a Write Bootstrap directive defined. This allows you to recover the state of your system as of the last backup.

- Backup your catalog. An example of this is found in the default bacula-dir.conf file. The backup script is installed by default and should handle any database, though you may want to make your own local modifications.
- Write a bootstrap file for the catalog. An example of this is found in the default bacula-dir.conf file. This will allow you to quickly restore your catalog in the event it is wiped out – otherwise it is many excruciating hours of work.
- Make a Bacula Rescue CDROM! See the Disaster Recovery Using a Bacula Rescue CDROM chapter. It is trivial to make such a CDROM, and it can make system recovery in the event of a lost hard disk infinitely easier.

Recommended Items

Although these items may not be critical, they are recommended and will help you avoid problems.

- Read the Quick Start Guide to Bacula
- After installing and experimenting with Bacula, read and work carefully through the examples in the Tutorial chapter of this manual.
- Learn what each of the Bacula Utility Programs does.
- Set up reasonable retention periods so that your catalog does not grow to be too big. See the following three chapters:
  Recycling your Volumes,
  Basic Volume Management,
  Using Pools to Manage Volumes.
- Perform a bare metal recovery using the Bacula Rescue CDROM. See the Disaster Recovery Using a Bacula Rescue CDROM chapter.

A Brief Tutorial

This chapter will guide you through running Bacula. To do so, we assume you have installed Bacula, possibly in a single file as shown in the previous chapter, in which case, you can run Bacula as non-root for these tests. However, we assume that you have not changed the .conf files. If you have modified the .conf files, please go back and uninstall Bacula, then reinstall it, but do not make any changes. The examples in this chapter use the default configuration files, and will write the volumes to disk in your **/tmp** directory, in addition, the data backed up will be the source directory where you built Bacula. As a consequence, you can run all the Bacula daemons for these tests as non-root. Please note, in production, your File daemon(s) must run as root. See the Security chapter for more information on this subject.

The general flow of running Bacula is:

1. cd <install-directory>
2. Start the Database (if using MySQL or PostgreSQL)
3. Start the Daemons with **./bacula start**
4. Start the Console program to interact with the Director
5. Run a job
6. When the Volume fills, unmount the Volume, if it is a tape, label a new one, and continue running. In this chapter, we will write only to disk files so you won't need to worry about tapes for the moment.
7. Test recovering some files from the Volume just written to ensure the backup is good and that you know how to recover. Better test before disaster strikes
8. Add a second client.

Each of these steps is described in more detail below.

Before Running Bacula

Before running Bacula for the first time in production, we recommend that you run the **test** command in the **btape** program as described in the Utility Program Chapter of this manual. This will help ensure that Bacula

functions correctly with your tape drive. If you have a modern HP, Sony, or Quantum DDS or DLT tape drive running on Linux or Solaris, you can probably skip this test as Bacula is well tested with these drives and systems. For all other cases, you are **strongly** encouraged to run the test before continuing. **btape** also has a **fill** command that attempts to duplicate what Bacula does when filling a tape and writing on the next tape. You should consider trying this command as well, but be forewarned, it can take hours (about 4 hours on my drive) to fill a large capacity tape.

Starting the Database

If you are using MySQL or PostgreSQL as the Bacula database, you should start it before you attempt to run a job to avoid getting error messages from Bacula when it starts. The scripts **startmysql** and **stopmysql** are what I (Kern) use to start and stop my local MySQL. Note, if you are using SQLite, you will not want to use **startmysql** or **stopmysql**. If you are running this in production, you will probably want to find some way to automatically start MySQL or PostgreSQL after each system reboot.

If you are using SQLite (i.e. you specified the **--with-sqlite=xxx** option on the **./configure** command, you need do nothing. SQLite is automatically started by **Bacula**.

Starting the Daemons

To start the three daemons, from your installation directory, simply enter:

./bacula start This script starts the Storage daemon, the File daemon, and the Director daemon, which all normally run as daemons in the background. If you are using the autostart feature of Bacula, your daemons will either be automatically started on reboot, or you can control them individually with the files **bacula-dir**, **bacula-fd**, and **bacula-sd**, which are usually located in **/etc/init.d**, though the actual location is system dependent.

Note, on Windows, currently only the File daemon is ported, and it must be started differently. Please see the Windows Version of Bacula Chapter of this manual.

The rpm packages configure the daemons to run as user=root and group=bacula. The rpm installation also creates the group bacula if it does not exist on the system. Any users that you add to the group bacula will have access to files created by the daemons. To disable or alter this behavior

edit the daemon startup scripts:

- /etc/bacula/bacula
- /etc/init.d/bacula-dir
- /etc/init.d/bacula-sd
- /etc/init.d/bacula-fd

and then restart as noted above.

The installation chapter of this manual explains how you can install scripts that will automatically restart the daemons when the system starts.

Interacting with the Director to Query or Start Jobs

To communicate with the director and to query the state of Bacula or run jobs, from the top level directory, simply enter:

./bconsole

Note, on 1.32 versions and lower, the command name is **console** rather than bconsole. Alternatively to running the command line console, if you have GNOME installed and used the **--enable-gnome** on the configure command, you may use the GNOME Console program:

./gnome-console

For simplicity, here we will describe only the **./console** program. Most of what is described here applies equally well to **./gnome-console**.

The **./bconsole** runs the Bacula Console program, which connects to the Director daemon. Since Bacula is a network program, you can run the Console program anywhere on your network. Most frequently, however, one runs it on the same machine as the Director. Normally, the Console program will print something similar to the following:

```
[kern@polymatou bin]$ ./bconsole
Connecting to Director lpmatou:9101
1000 OK: HeadMan Version: 1.30 (28 April 2003)
*
```

the asterisk is the console command prompt.

Type **help** to see a list of available commands:

```
*help
  Command    Description
  =======    ===========
  add        add media to a pool
  autodisplay autodisplay [on/off] -- console messages
  automount  automount [on/off] -- after label
  cancel     cancel job=nnn -- cancel a job
  create     create DB Pool from resource
  delete     delete [pool=<pool-name> | media volume=<volume-name>]
  estimate   performs FileSet estimate debug=1 give full listing
  exit       exit = quit
  help       print this command
  label      label a tape
  list       list [pools | jobs | jobtotals | media <pool> |
             files jobid=<nn>]; from catalog
  llist      full or long list like list command
  messages   messages
  mount      mount <storage-name>
  prune      prune expired records from catalog
  purge      purge records from catalog
  query      query catalog
  quit       quit
  relabel    relabel a tape
  release    release <storage-name>
  restore    restore files
  run        run <job-name>
  setdebug   sets debug level
  show       show (resource records) [jobs | pools | ... | all]
  sqlquery   use SQL to query catalog
  status     status [storage | client]=<name>
  time       print current time
  unmount    unmount <storage-name>
  update     update Volume or Pool
  use        use catalog xxx
  var        does variable expansion
  version    print Director version
  wait       wait until no jobs are running
*
```

Details of the console program's commands are explained in the Console Chapter of this manual.

Running a Job

At this point, we assume you have done the following:

- Configured Bacula with **./configure --your-options**

- Built Bacula using **make**
- Installed Bacula using **make install**
- Have created your database with, for example, **./create_sqlite_database**
- Have created the Bacula database tables with, **./make_bacula_tables**
- Have possibly edited your **bacula-dir.conf** file to personalize it a bit. BE CAREFUL! if you change the Director's name or password, you will need to make similar modifications in the other .conf files. For the moment it is probably better to make no changes.
- You have started Bacula with **./bacula start**
- You have invoked the Console program with **./bconsole**

Furthermore, we assume for the moment you are using the default configuration files.

At this point, enter the following command:

```
show filesets
```

and you should get something similar to:

```
FileSet: name=Full Set
      Inc: /home/kern/bacula/bacula-1.30
      Exc: /proc
      Exc: /tmp
      Exc: /.journal
      Exc: /.fsck
FileSet: name=Catalog
      Inc: /home/kern/bacula/testbin/working/bacula.sql
```

This is a pre-defined **FileSet** that will backup the Bacula source directory. The actual directory names printed should correspond to your system configuration. For testing purposes, we have chosen a directory of moderate size (about 40 Megabytes) and complexity without being too big. The FileSet **Catalog** is used for backing up Bacula's catalog and is not of interest to us for the moment. The **Inc:** entries are the files or directories that will be included in the backup and the **Exc:** are those that will be excluded.

Now is the time to run your first backup job. We are going to backup your Bacula source directory to a File Volume in your **/tmp** directory just to show you how easy it is. Now enter:

```
status dir
```

and you should get the following output:

```
rufus-dir Version: 1.30 (28 April 2003)
Daemon started 28-Apr-2003 14:03, 0 Jobs run.
Console connected at 28-Apr-2003 14:03
No jobs are running.
Level          Type     Scheduled          Name
===============================================================
Incremental    Backup   29-Apr-2003 01:05  Client1
Full           Backup   29-Apr-2003 01:10  BackupCatalog
====
```

where the times and the Director's name will be different according to your setup. This shows that an Incremental job is scheduled to run for the Job **Client1** at 1:05am and that at 1:10, a **BackupCatalog** is scheduled to run. Note, you should probably change the name **Client1** to be the name of your machine, if not, when you add additional clients, it will be very confusing. For my real machine, I use **Rufus** rather than **Client1** as in this example.

Now enter:

```
status client
```

and you should get something like:

```
The defined Client resources are:
     1: rufus-fd
Item 1 selected automatically.
Connecting to Client rufus-fd at rufus:8102
rufus-fd Version: 1.30 (28 April 2003)
Daemon started 28-Apr-2003 14:03, 0 Jobs run.
Director connected at: 28-Apr-2003 14:14
No jobs running.
====
```

In this case, the client is named **rufus-fd** your name will be different, but the line beginning with **rufus-fd Version ...** is printed by your File daemon, so we are now sure it is up and running.

Finally do the same for your Storage daemon with:

```
status storage
```

and you should get:

```
The defined Storage resources are:
     1: File
Item 1 selected automatically.
Connecting to Storage daemon File at rufus:8103
rufus-sd Version: 1.30 (28 April 2003)
Daemon started 28-Apr-2003 14:03, 0 Jobs run.
Device /tmp is not open.
No jobs running.
====
```

You will notice that the default Storage daemon device is named **File** and that it will use device **/tmp**, which is not currently open.

Now, let's actually run a job with:

```
run
```

you should get the following output:

```
Using default Catalog name=MyCatalog DB=bacula
A job name must be specified.
The defined Job resources are:
     1: Client1
     2: BackupCatalog
     3: RestoreFiles
Select Job resource (1-3):
```

Here, Bacula has listed the three different Jobs that you can run, and you should choose number **1** and type enter, at which point you will get:

```
Run Backup job
JobName:  Client1
FileSet:  Full Set
Level:    Incremental
Client:   rufus-fd
Storage:  File
Pool:     Default
When:     2003-04-28 14:18:57
OK to run? (yes/mod/no):
```

At this point, take some time to look carefully at what is printed and understand it. It is asking you if it is OK to run a job named **Client1** with FileSet **Full Set** (we listed above) as an Incremental job on your Client (your client name will be different), and to use Storage **File** and Pool **Default**, and finally, it wants to run it now (the current time should be displayed by your console).

Here we have the choice to run (**yes**), to modify one or more of the above parameters (**mod**), or to not run the job (**no**). Please enter **yes**, at which point you should immediately get the command prompt (an asterisk). If you wait a few seconds, then enter the command **messages** you will get back something like:

```
28-Apr-2003 14:22 rufus-dir: Last FULL backup time not found. Doing
                  FULL backup.
28-Apr-2003 14:22 rufus-dir: Start Backup JobId 1,
                  Job=Client1.2003-04-28_14.22.33
28-Apr-2003 14:22 rufus-sd: Job Client1.2003-04-28_14.22.33 waiting.
                  Cannot find any appendable volumes.
Please use the "label"  command to create a new Volume for:
    Storage:      FileStorage
    Media type:   File
    Pool:         Default
```

The first message, indicates that no previous Full backup was done, so Bacula is upgrading our Incremental job to a Full backup (this is normal). The second message indicates that the job started with JobId 1., and the third message tells us that Bacula cannot find any Volumes in the Pool for writing the output. This is normal because we have not yet created (labeled) any Volumes. Bacula indicates to you all the details of the volume it needs.

At this point, the job is blocked waiting for a Volume. You can check this if you want by doing a **status dir**. In order to continue, we must create a Volume that Bacula can write on. We do so with:

```
label
```

and Bacula will print:

```
The defined Storage resources are:
     1: File
Item 1 selected automatically.
Enter new Volume name:
```

at which point, you should enter some name beginning with a letter and containing only letters and numbers (period, hyphen, and underscore) are also permitted. For example, enter **TestVolume001**, and you should get back:

```
Defined Pools:
     1: Default
Item 1 selected automatically.
Connecting to Storage daemon File at rufus:8103 ...
Sending label command for Volume "TestVolume001" Slot 0 ...
3000 OK label. Volume=TestVolume001 Device=/tmp
Catalog record for Volume "TestVolume002", Slot 0  successfully created.
Requesting mount FileStorage ...
3001 OK mount. Device=/tmp
```

Finally, enter **messages** and you should get something like:

```
28-Apr-2003 14:30 rufus-sd: Wrote label to prelabeled Volume
   "TestVolume001" on device /tmp
28-Apr-2003 14:30 rufus-dir: Bacula 1.30 (28Apr03): 28-Apr-2003 14:30
JobId:                  1
Job:                    Client1.2003-04-28_14.22.33
FileSet:                Full Set
Backup Level:           Full
Client:                 rufus-fd
Start time:             28-Apr-2003 14:22
End time:               28-Apr-2003 14:30
Files Written:          1,444
Bytes Written:          38,988,877
Rate:                   81.2 KB/s
Software Compression:   None
Volume names(s):        TestVolume001
Volume Session Id:      1
Volume Session Time:    1051531381
Last Volume Bytes:      39,072,359
FD termination status:  OK
SD termination status:  OK
Termination:            Backup OK
28-Apr-2003 14:30 rufus-dir: Begin pruning Jobs.
28-Apr-2003 14:30 rufus-dir: No Jobs found to prune.
28-Apr-2003 14:30 rufus-dir: Begin pruning Files.
28-Apr-2003 14:30 rufus-dir: No Files found to prune.
28-Apr-2003 14:30 rufus-dir: End auto prune.
```

If you don't see the output immediately, you can keep entering **messages** until the job terminates, or you can enter, **autodisplay on** and your messages will automatically be displayed as soon as they are ready.

If you do an **ls -l** of your **/tmp** directory, you will see that you have the following item:

```
-rw-r-----    1 kern     kern     39072153 Apr 28 14:30 TestVolume001
```

This is the file Volume that you just wrote and it contains all the data of the job just run. If you run additional jobs, they will be appended to this Volume unless you specify otherwise.

You might ask yourself if you have to label all the Volumes that Bacula is going to use. The answer for disk Volumes, like the one we used, is no. It is possible to have Bacula automatically label volumes. For tape Volumes, you will most likely have to label each of the Volumes you want to use.

If you would like to stop here, you can simply enter **quit** in the Console program, and you can stop Bacula with **./bacula stop**. To clean up, simply delete the file **/tmp/TestVolume001**, and you should also re-initialize your database using:

```
./drop_bacula_tables
./make_bacula_tables
```

Please note that this will erase all information about the previous jobs that have run, and that you might want to do it now while testing but that normally you will not want to re-initialize your database.

If you would like to try restoring the files that you just backed up, read the following section.

Restoring Your Files

If you have run the default configuration and the save of the Bacula source code as demonstrated above, you can restore the backed up files in the Console program by entering:

```
restore all
```

where you will get:

```
First you select one or more JobIds that contain files
to be restored. You will be presented several methods
of specifying the JobIds. Then you will be allowed to
select which files from those JobIds are to be restored.
To select the JobIds, you have the following choices:
     1: List last 20 Jobs run
```

```
     2: List Jobs where a given File is saved
     3: Enter list of comma separated JobIds to select
     4: Enter SQL list command
     5: Select the most recent backup for a client
     6: Select backup for a client before a specified time
     7: Enter a list of files to restore
     8: Enter a list of files to restore before a specified time
     9: Cancel
Select item:  (1-9):
```

As you can see, there are a number of options, but for the current demonstration, please enter **5** to do a restore of the last backup you did, and you will get the following output:

```
Defined Clients:
     1: rufus-fd
Item 1 selected automatically.
The defined FileSet resources are:
     1: 1  Full Set  2003-04-28 14:22:33
Item 1 selected automatically.
+-------+-------+----------+---------------------+---------------+
| JobId | Level | JobFiles | StartTime           | VolumeName    |
+-------+-------+----------+---------------------+---------------+
| 1     | F     | 1444     | 2003-04-28 14:22:33 | TestVolume002 |
+-------+-------+----------+---------------------+---------------+
You have selected the following JobId: 1
Building directory tree for JobId 1 ...
1 Job inserted into the tree and marked for extraction.
The defined Storage resources are:
     1: File
Item 1 selected automatically.
You are now entering file selection mode where you add and
remove files to be restored. All files are initially added.
Enter "done" to leave this mode.
cwd is: /
$
```

where I have truncated the listing on the right side to make it more readable. As you can see by starting at the top of the listing, Bacula knows what client you have, and since there was only one, it selected it automatically, likewise for the FileSet. Then Bacula produced a listing containing all the jobs that form the current backup, in this case, there is only one, and the Storage daemon was also automatically chosen. Bacula then took all the files that were in Job number 1 and entered them into a **directory tree** (a sort of in memory representation of your filesystem). At this point, you can use the **cd** and **ls** ro **dir** commands to walk up and down the directory tree and view what files will be restored. For example, if I enter **cd /home/kern/bacula/bacula-1.30** and then enter **dir** I will get a listing

of all the files in the Bacula source directory. On your system, the path will be somewhat different. For more information on this, please refer to the Restore Command Chapter of this manual for more details.

To exit this mode, simply enter:

```
done
```

and you will get the following output:

```
Bootstrap records written to
   /home/kern/bacula/testbin/working/restore.bsr
The restore job will require the following Volumes:

   TestVolume001
1444 files selected to restore.
Run Restore job
JobName:    RestoreFiles
Bootstrap:  /home/kern/bacula/testbin/working/restore.bsr
Where:      /tmp/bacula-restores
Replace:    always
FileSet:    Full Set
Client:     rufus-fd
Storage:    File
JobId:      *None*
When:       2003-04-28 14:53:54
OK to run? (yes/mod/no):
```

If you answer **yes** your files will be restored to **/tmp/bacula-restores**. If you want to restore the files to their original locations, you must use the **mod** option and explicitly set **Where:** to nothing (or to /). We recommend you go ahead and answer **yes** and after a brief moment, enter **messages**, at which point you should get a listing of all the files that were restored as well as a summary of the job that looks similar to this:

```
28-Apr-2003 14:56 rufus-dir: Bacula 1.30 (28Apr03): 28-Apr-2003 14:56
JobId:                  2
Job:                    RestoreFiles.2003-04-28_14.56.06
Client:                 rufus-fd
Start time:             28-Apr-2003 14:56
End time:               28-Apr-2003 14:56
Files Restored:         1,444
Bytes Restored:         38,816,381
Rate:                   9704.1 KB/s
FD termination status:  OK
Termination:            Restore OK
28-Apr-2003 14:56 rufus-dir: Begin pruning Jobs.
```

```
28-Apr-2003 14:56 rufus-dir: No Jobs found to prune.
28-Apr-2003 14:56 rufus-dir: Begin pruning Files.
28-Apr-2003 14:56 rufus-dir: No Files found to prune.
28-Apr-2003 14:56 rufus-dir: End auto prune.
```

After exiting the Console program, you can examine the files in **/tmp/bacula-restores**, which will contain a small directory tree with all the files. Be sure to clean up at the end with:

```
rm -rf /tmp/bacula-restore
```

Quitting the Console Program

Simply enter the command **quit**.

Adding a Second Client

If you have gotten the example shown above to work on your system, you may be ready to add a second Client (File daemon). That is you have a second machine that you would like backed up. The only part you need installed on the other machine is the binary **bacula-fd** (or **bacula-fd.exe** for Windows) and its configuration file **bacula-fd.conf**. You can start with the same **bacula-fd.conf** file that you are currently using and make one minor modification to it to create the conf file for your second client. Change the File daemon name from whatever was configured, **rufus-fd** in the example above, but your system will have a different name. The best is to change it to the name of your second machine. For example:

```
...
#
# "Global" File daemon configuration specifications
#
FileDaemon {                          # this is me
  Name = rufus-fd
  FDport = 9102                  # where we listen for the director
  WorkingDirectory = /home/kern/bacula/working
  Pid Directory = /var/run
}
...
```

would become:

```
...
#
# "Global" File daemon configuration specifications
#
FileDaemon {                            # this is me
  Name = matou-fd
  FDport = 9102                  # where we listen for the director
  WorkingDirectory = /home/kern/bacula/working
  Pid Directory = /var/run
}
...
```

where I show just a portion of the file and have changed **rufus-fd** to **matou-fd**. The names you use are your choice. For the moment, I recommend you change nothing else. Later, you will want to change the password.

Now you should install that change on your second machine. Then you need to make some additions to your Director's configuration file to define the new File daemon or Client. Starting from our original example which should be installed on your system, you should add the following lines (essentially copies of the existing data but with the names changed) to your Director's configuration file **bacula-dir.conf**.

```
#
# Define the main nightly save backup job
#   By default, this job will back up to disk in /tmp
Job {
  Name = "Matou"
  Type = Backup
  Client = matou-fd
  FileSet = "Full Set"
  Schedule = "WeeklyCycle"
  Storage = File
  Messages = Standard
  Pool = Default
  Write Bootstrap = "/home/kern/bacula/working/matou.bsr"
}
# Client (File Services) to backup
Client {
  Name = matou-fd
  Address = matou
  FDPort = 9102
  Catalog = MyCatalog
  Password = "xxxxx"                # password for
  File Retention = 30d              # 30 days
  Job Retention = 180d              # six months
  AutoPrune = yes                   # Prune expired Jobs/Files
}
```

Then make sure that the Address parameter in the Storage resource is set to the fully qualified domain name and not to something like "localhost". The address specified is sent to the File daemon (client) and it must be a fully qualified domain name. If you pass something like "localhost" it will not resolve correctly and will result in a time out when the File daemon fails to connect to the Storage daemon.

That is all that is necessary. I copied the existing resource to create a second Job (Matou) to backup the second client (matou-fd). It has the name **Matou**, the Client is named **matou-fd**, and the bootstrap file name is changed, but everything else is the same. This means that Matou will be backed up on the same schedule using the same set of tapes. You may want to change that later, but for now, let's keep it simple.

The second change was to add a new Client resource that defines **matou-fd** and has the correct address **matou**, but in real life, you may need a fully qualified machine address or an IP address. I also kept the password the same (shown as xxxxx for the example).

At this point, if you stop Bacula and restart it, and start the Client on the other machine, everything will be ready, and the prompts that you saw above will now include the second machine.

To make this a real production installation, you will possibly want to use different Pool, or a different schedule. It is up to you to customize. In any case, you should change the password in both the Director's file and the Client's file for additional security.

For some important tips on changing names and passwords, and a diagram of what names and passwords must match, please see Authorization Errors in the FAQ chapter of this manual.

When The Tape Fills

If you have scheduled your job, typically nightly, there will come a time when the tape fills up and **Bacula** cannot continue. In this case, Bacula will send you a message similar to the following:

```
rufus-sd: block.c:337 === Write error errno=28: ERR=No space left
          on device
```

This indicates that Bacula got a write error because the tape is full. Bacula will then search the Pool specified for your Job looking for an appendable

volume. In the best of all cases, you will have properly set your Retention Periods and you will have all your tapes marked to be Recycled, and **Bacula** will automatically recycle the tapes in your pool requesting and overwriting old Volumes. For more information on recycling, please see the Recycling chapter of this manual. If you find that your Volumes were not properly recycled (usually because of a configuration error), please see the Manually Recycling Volumes section of the Recycling chapter.

If like me, you have a very large set of Volumes and you label them with the date the Volume was first writing, or you have not set up your Retention periods, Bacula will not find a tape in the pool, and it will send you a message similar to the following:

```
rufus-sd: Job kernsave.2002-09-19.10:50:48 waiting. Cannot find any
          appendable volumes.
Please use the "label"  command to create a new Volume for:
    Storage:      SDT-10000
    Media type:   DDS-4
    Pool:         Default
```

Until you create a new Volume, this message will be repeated an hour later, then two hours later, and so on doubling the interval each time up to a maximum interval of 1 day.

The obvious question at this point is: What do I do now?

The answer is simple: first, using the Console program, close the tape drive using the **unmount** command. If you only have a single drive, it will be automatically selected, otherwise, make sure you release the one specified on the message (in this case **STD-10000**).

Next, you remove the tape from the drive and insert a new blank tape. Note, on some older tape drives, you may need to write an end of file mark (**mt -f /dev/nst0 weof**) to prevent the drive from running away when Bacula attempts to read the label.

Finally, you use the **label** command in the Console to write a label to the new Volume. The **label** command will contact the Storage daemon to write the software label, if it is successful, it will add the new Volume to the Pool, then issue a **mount** command to the Storage daemon. See the previous sections of this chapter for more details on labeling tapes.

The result is that Bacula will continue the previous Job writing the backup to the new Volume.

If you have a Pool of volumes and Bacula is cycling through them, instead of the above message "Cannot find any appendable volumes.", Bacula may ask you to mount a specific volume. In that case, you should attempt to do just that. If you do not have the volume any more (for any of a number of reasons), you can simply mount another volume from the same Pool, providing it is appendable, and Bacula will use it. You can use the **list volumes** command in the console program to determine which volumes are appendable and which are not.

If like me, you have your Volume retention periods set correctly, but you have no more free Volumes, you can relabel and reuse a Volume as follows:

- Do a **list volumes** in the Console and select the oldest Volume for relabeling.
- If you have setup your Retention periods correctly, the Volume should have VolStatus **Purged**.
- If the VolStatus is not set to Purged, you will need to purge the database of Jobs that are written on that Volume. Do so by using the command **purge jobs volume** in the Console. If you have multiple Pools, you will be prompted for the Pool then enter the VolumeName (or MediaId) when requested.
- Then simply use the **relabel** command to relabel the Volume.

To manually relabel the Volume use the following additional steps:

- To delete the Volume from the catalog use the **delete volume** command in the Console and select the VolumeName (or MediaId) to be deleted.
- Use the **unmount** command in the Console to unmount the old tape.
- Physically relabel the old Volume that you deleted so that it can be reused.
- Insert the old Volume in the tape drive.
- From a command line do: **mt -f /dev/st0 rewind** and **mt -f /dev/st0 weof**, where you need to use the proper tape drive name for your system in place of **/dev/st0**.
- Use the **label** command in the Console to write a new Bacula label on your tape.

- Use the **mount** command in the Console if it is not automatically done, so that Bacula starts using your newly labeled tape.

Other Useful Console Commands

status dir Print a status of all running jobs and jobs scheduled in the next 24 hours.

status The console program will prompt you to select a daemon type, then will request the daemon's status.

status jobid=nn Print a status of JobId nn if it is running. The Storage daemon is contacted and requested to print a current status of the job as well.

list pools List the pools defined in the Catalog (normally only Default is used).

list media Lists all the media defined in the Catalog.

list jobs Lists all jobs in the Catalog that have run.

list jobid=nn Lists JobId nn from the Catalog.

list jobtotals Lists totals for all jobs in the Catalog.

list files jobid=nn List the files that were saved for JobId nn.

list jobmedia List the media information for each Job run.

messages Prints any messages that have been directed to the console.

unmount storage=storage-name Unmounts the drive associated with the storage device with the name **storage-name** if the drive is not currently being used. This command is used if you wish Bacula to free the drive so that you can use it to label a tape.

mount storage=storage-name Causes the drive associated with the storage device to be mounted again. When Bacula reaches the end of a volume and requests you to mount a new volume, you must issue this command after you have placed the new volume in the drive. In effect, it is the signal needed by Bacula to know to start reading or writing the new volume.

quit Exit or quit the console program.

Most of the commands given above, with the exception of **list**, will prompt you for the necessary arguments if you simply enter the command name.

Debug Daemon Output

If you want debug output from the daemons as they are running, start the daemons from the install directory as follows:

```
./bacula start -d20
```

To stop the three daemons, enter the following from the install directory:

```
./bacula stop
```

The execution of **bacula stop** may complain about pids not found. This is OK, especially if one of the daemons has died, which is very rare.

To do a full system save, each File daemon must be running as root so that it will have permission to access all the files. None of the other daemons require root privileges. However, the Storage daemon must be able to open the tape drives. On many systems, only root can access the tape drives. Either run the Storage daemon as root, or change the permissions on the tape devices to permit non-root access. MySQL and PostgreSQL can be installed and run with any userid; root privilege is not necessary.

Have Patience When Starting the Daemons or Mounting Blank Tapes

When you start the Bacula daemons, the Storage daemon attempts to open all defined storage devices and verify the currently mounted Volume (if configured). Until all the storage devices are verified, the Storage daemon will not accept connections from the Console program. If a tape was previously used, it will be rewound, and on some devices this can take several minutes. As a consequence, you may need to have a bit of patience when first contacting the Storage daemon after starting the daemons. If you can see your tape drive, once the lights stop flashing, the drive will be ready to be used.

The same considerations apply if you have just mounted a blank tape in a drive such as an HP DLT. It can take a minute or two before the drive properly recognizes that the tape is blank. If you attempt to **mount** the tape with the Console program during this recognition period, it is quite possible that you will hang your SCSI driver (at least on my RedHat Linux system). As a consequence, you are again urged to have patience when inserting blank tapes. Let the device settle down before attempting to access it.

Difficulties Connecting from the FD to the SD

If you are having difficulties getting one or more of your File daemons to connect to the Storage daemon, it is most likely because you have not used a fully qualified Internet address on the **Address** directive in the Director's Storage resource. That is the resolver on the File daemon's machine (not on the Director's) must be able to resolve the name you supply into an IP address. An example of an address that is guaranteed not to work: **localhost**. An example that may work: **megalon**. An example that is more likely to work: **magalon.mydomain.com**. On Win32 if you don't have a good resolver (often true on older Win98 systems), you might try using an IP address in place of a name.

If your address is correct, then make sure that no other program is using the port 9103 on the Storage daemon's machine. The Bacula port number are authorized by IANA, and should not be used by other programs, but apparently some HP printers do use these port numbers. A **netstat -a** on the Storage daemon's machine can determine who is using the 9103 port (used for FD to SD communications in Bacula).

Daemon Command Line Options

Each of the three daemons (Director, File, Storage) accepts a small set of options on the command line. In general, each of the daemons as well as the Console program accepts the following options:

-c <**file**> Define the file to use as a configuration file. The default is the daemon name followed by **.conf** i.e. **bacula-dir.conf** for the Director, **bacula-fd.conf** for the File daemon, and **bacula-sd** for the Storage daemon.

-d nn Set the debug level to **nn**. Higher levels of debug cause more information to displayed on STDOUT concerning what the daemon is doing.

-f Run the daemon in the foreground. This option is needed to run the daemon under the debugger.

-s Do not trap signals. This option is needed to run the daemon under the debugger.

-t Read the configuration file and print any error messages, then immediately exit. Useful for syntax testing of new configuration files.

-v Be more verbose or more complete in printing error and informational messages. Recommended.

-? Print the version and list of options.

The Director has the following additional Director specific option:

-r <job> Run the named job immediately. This is for debugging and should not be used.

The File daemon has the following File daemon specific option:

-i Assume that the daemon is called from **inetd** or **xinetd**. In this case, the daemon assumes that a connection has already been made and that it is passed as STDIN. After the connection terminates the daemon will exit.

The Storage daemon has no Storage daemon specific options.

The Console program has no console specific options.

Creating a Pool

Creating the Pool is automatically done when **Bacula** starts, so if you understand Pools, you can skip to the next section.

When you run a job, one of the things that Bacula must know is what Volumes to use to backup the FileSet. Instead of specifying a Volume (tape) directly, you specify which Pool of Volumes you want Bacula to consult when it wants a tape for writing backups. Bacula will select the first available Volume from the Pool that is appropriate for the Storage device you have specified for the Job being run. When a volume has filled up with data, **Bacula** will change its VolStatus from **Append** to **Full**, and then **Bacula** will use the next volume and so on. If no appendable Volume exists in the Pool, the Director will attempt to recycle an old Volume, if there are still no appendable Volumes available, **Bacula** will send a message requesting the operator to create an appropriate Volume.

Bacula keeps track of the Pool name, the volumes contained in the Pool, and a number of attributes of each of those Volumes.

When Bacula starts, it ensures that all Pool resource definitions have been recorded in the catalog. You can verify this by entering:

```
list pools
```

to the console program, which should print something like the following:

```
*list pools
Using default Catalog name=MySQL DB=bacula
+--------+---------+---------+---------+----------+-------------+
| PoolId | Name    | NumVols | MaxVols | PoolType | LabelFormat |
+--------+---------+---------+---------+----------+-------------+
| 1      | Default | 3       | 0       | Backup   | *           |
| 2      | File    | 12      | 12      | Backup   | File        |
+--------+---------+---------+---------+----------+-------------+
*
```

If you attempt to create the same Pool name a second time, **Bacula** will print:

```
Error: Pool Default already exists.
Once created, you may use the {\bf update} command to
modify many of the values in the Pool record.
```

Labeling Your Volumes

Bacula requires that each Volume contain a software label. There are several strategies for labeling volumes. The one I use is to label them as they are needed by **Bacula** using the console program. That is when Bacula needs a new Volume, and it does not find one in the catalog, it will send me an email message requesting that I add Volumes to the Pool. I then use the **label** command in the Console program to label a new Volume and to define it in the Pool database, after which Bacula will begin writing on the new Volume. Alternatively, I can use the Console **relabel** command to relabel a Volume that is no longer used providing it has VolStatus **Purged**.

Another strategy is to label a set of volumes at the start, then use them as **Bacula** requests them. This is most often done if you are cycling through a set of tapes, for example using an autochanger. For more details on recycling, please see the Automatic Volume Recycling chapter of this manual.

If you run a Bacula job, and you have no labeled tapes in the Pool, Bacula will inform you, and you can create them "on-the-fly" so to speak. In my case, I label my tapes with the date, for example: **DLT-18April02**. See below for the details of using the **label** command.

Labeling Volumes with the Console Program

Labeling volumes is normally done by using the console program.

1. ./bconsole
2. label

If Bacula complains that you cannot label the tape because it is already labeled, simply **unmount** the tape using the **unmount** command in the console, then physically mount a blank tape and re-issue the **label** command.

Since the physical storage media is different for each device, the **label** command will provide you with a list of the defined Storage resources such as the following:

```
The defined Storage resources are:
     1: File
     2: 8mmDrive
     3: DLTDrive
     4: SDT-10000
Select Storage resource (1-4):
```

At this point, you should have a blank tape in the drive corresponding to the Storage resource that you select.

It will then ask you for the Volume name.

```
Enter new Volume name:
```

If Bacula complains:

```
Media record for Volume xxxx already exists.
```

It means that the volume name **xxxx** that you entered already exists in the Media database. You can list all the defined Media (Volumes) with the **list media** command. Note, the LastWritten column has been truncated for proper printing.

```
+--------------+---------+--------+----------------+-----/~/-+------------+-----+
```

VolumeName	MediaTyp	VolStat	VolBytes	LastWri	VolReten	Recy
DLTVol0002	DLT8000	Purged	56,128,042,217	2001-10	31,536,000	0
DLT-07Oct2001	DLT8000	Full	56,172,030,586	2001-11	31,536,000	0
DLT-08Nov2001	DLT8000	Full	55,691,684,216	2001-12	31,536,000	0
DLT-01Dec2001	DLT8000	Full	55,162,215,866	2001-12	31,536,000	0
DLT-28Dec2001	DLT8000	Full	57,888,007,042	2002-01	31,536,000	0
DLT-20Jan2002	DLT8000	Full	57,003,507,308	2002-02	31,536,000	0
DLT-16Feb2002	DLT8000	Full	55,772,630,824	2002-03	31,536,000	0
DLT-12Mar2002	DLT8000	Full	50,666,320,453	1970-01	31,536,000	0
DLT-27Mar2002	DLT8000	Full	57,592,952,309	2002-04	31,536,000	0
DLT-15Apr2002	DLT8000	Full	57,190,864,185	2002-05	31,536,000	0
DLT-04May2002	DLT8000	Full	60,486,677,724	2002-05	31,536,000	0
DLT-26May02	DLT8000	Append	1,336,699,620	2002-05	31,536,000	1

Once Bacula has verified that the volume does not already exist, it will then prompt you for the name of the Pool in which the Volume (tape) to be created. If there is only one Pool (Default), it will be automatically selected.

If the tape is successfully labeled, a media record will also be created in the Pool. That is the Volume name and all its other attributes will appear when you list the Pool. In addition, that Volume will be available for backup if the MediaType matches what is requested by the Storage daemon.

When you labeled the tape, you answered very few questions about it – principally the Volume name, and perhaps the Slot. However, a Volume record in the catalog database (internally known as a Media record) contains quite a few attributes. Most of these attributes will be filled in from the default values that were defined in the Pool (i.e. the Pool holds most of the default attributes used when creating a Volume).

It is also possible to add media to the pool without physically labeling the Volumes. This can be done with the **add** command. For more information, please see the Console Chapter of this manual.

Customizing the Configuration Files

When each of the Bacula programs starts, it reads a configuration file specified on the command line or the default **bacula-dir.conf**, **bacula-fd.conf**, **bacula-sd.conf**, or **console.conf** for the Director daemon, the File daemon, the Storage daemon, and the Console program respectively.

Each service (Director, Client, Storage, Console) has its own configuration file containing a set of Resource definitions. These resources are very similar from one service to another, but may contain different directives (records) depending on the service. For example, in the Director's resource file, the **Director** resource defines the name of the Director, a number of global Director parameters and his password. In the File daemon configuration file, the **Director** resource specifies which Directors are permitted to use the File daemon.

Before running Bacula for the first time, you must customize the configuration files for each daemon. Default configuration files will have been created by the installation process, but you will need to modify them to correspond to your system. An overall view of the resources can be seen in the following:

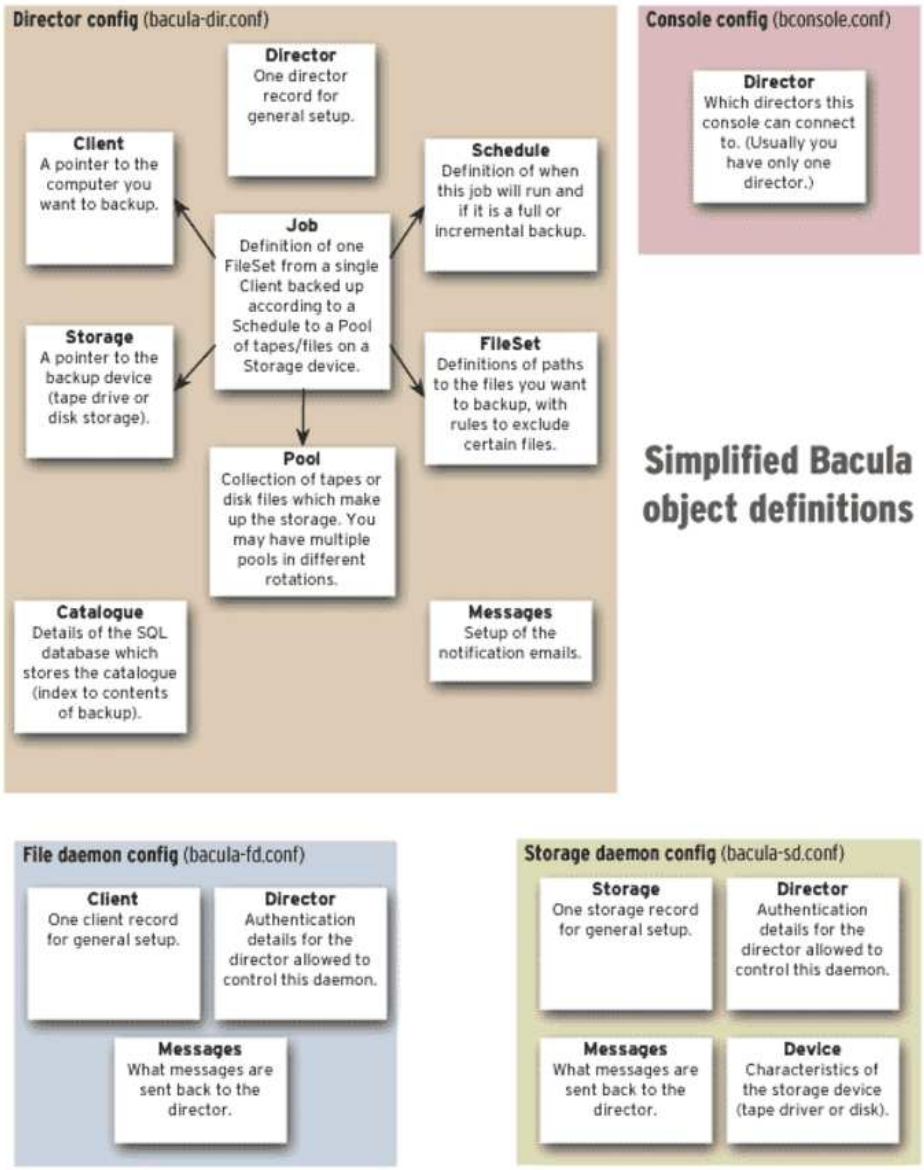

(thanks to Aristedes Maniatis for the above graphic)

Resource Directive Format

Although, you won't need to know the details of all the directives a basic knowledge of Bacula resource directives is essential. Each directive contained within the resource (within the braces) is composed of a keyword followed by an equal sign (=) followed by one or more values. The keywords must be

one of the known Bacula resource record keywords, and it may be composed of upper or lower case characters and spaces.

Each resource definition MUST contain a Name directive, and may optionally contain a Description directive (or record). The Name directive is used to uniquely identify the resource. The Description directive is (will be) used during display of the Resource to provide easier human recognition. For example:

```
Director {
  Name = "MyDir"
  Description = "Main Bacula Director"
  WorkingDirectory = "$HOME/bacula/bin/working"
}
```

Defines the Director resource with the name "MyDir" and a working directory $HOME/bacula/bin/working. In general, if you want spaces in a name to the right of the first equal sign (=), you must enclose that name within double quotes. Otherwise quotes are not generally necessary because once defined, quoted strings and unquoted strings are all equal.

Comments

When reading the configuration file, blank lines are ignored and everything after a hash sign (#) until the end of the line is taken to be a comment. A semicolon (;) is a logical end of line, and anything after the semicolon is considered as the next statement. If a statement appears on a line by itself, a semicolon is not necessary to terminate it, so generally in the examples in this manual, you will not see many semicolons.

Upper and Lower Case and Spaces

Case (upper/lower) and spaces are totally ignored in the resource directive keywords (the part before the equal sign).

Within the keyword (i.e. before the equal sign), spaces are not significant. Thus the keywords: **name**, **Name**, and **N a m e** are all identical.

Spaces after the equal sign and before the first character of the value are ignored.

In general, spaces within a value are significant (not ignored), and if the value is a name, you must enclose the name in double quotes for the spaces

to be accepted. Names may contain up to 127 characters. Currently, a name may contain any ASCII character. Within a quoted string, any character following a backslash (\) is taken as itself (handy for inserting blackslashes and double quotes (“).

Please note, however, that Bacula resource names as well as certain other names (e.g. Volume names) will in the future be severely limited to permit only letters (including ISO accented letters), numbers, and a few special characters (space, underscore, ...). All other characters and punctuation will be illegal.

Including other Configuration Files

If you wish to break your configuration file into smaller pieces, you can do so by including other files using the syntax @**filename** where **filename** is the full path and filename of another file. The @filename specification can be given anywhere a primitive token would appear.

Recognized Primitive Data Types

When parsing the resource directives, Bacula classifies the data according to the types listed below. The first time you read this, it may appear a bit overwhelming, but in reality, it is all pretty logical and straight forward.

name A keyword or name consisting of alpha numeric characters, including the hyphen, underscore, and dollar characters. The first character of a **name** must be a letter. A name has a maximum length currently set to 127 bytes. Typically keywords appear on the left side of an equal (i.e. they are Bacula keywords – i.e. Resource names or directive names). Keywords may not be quoted.

name-string A name-string is similar to a name, except that the name may be quoted and can thus contain additional characters including spaces. Name strings are limited to 127 characters in length. Name strings are typically used on the right side of an equal (i.e. they are values to be associated with a keyword.

string A quoted string containing virtually any character including spaces, or a non-quoted string. A string may be of any length. Strings are typically values that correspond to filenames, directories, or system command names. A backslash (\) turns the next character into itself,

so to include a double quote in a string, you precede the double quote with a backslash. Likewise to include a backslash.

directory A directory is either a quoted or non-quoted string. A directory will be passed to your standard shell for expansion when it is scanned. Thus constructs such as **$HOME** are interpreted to be their correct values.

password This is a Bacula password and it is stored internally in MD5 hashed format.

integer A 32 bit integer value. It may be positive or negative.

positive integer A 32 bit positive integer value.

long integer A 64 bit integer value. Typically these are values such as bytes that can exceed 4 billion and thus require a 64 bit value.

yes—no Either a **yes** or a **no**.

size A size specified as bytes. Typically, this is a floating point scientific input format followed by an optional modifier. The floating point input is stored as a 64 bit integer value. If a modifier is present, it must immediately follow the value with no intervening spaces. The following modifiers are permitted:

k 1,024 (kilobytes)

kb 1,000 (kilobytes)

m 1,048,576 (megabytes)

mb 1,000,000 (megabytes)

g 1,073,741,824 (gigabytes)

gb 1,000,000,000 (gigabytes)

time A time or duration specified in seconds. The time is stored internally as a 64 bit integer value, but it is specified in two parts: a number part and a modifier part. The number can be an integer or a floating point number. If it is entered in floating point notation, it will be rounded to the nearest integer. The modifer is mandatory and follows the number part, either with or without intervening spaces. The following modifiers are permitted:

seconds seconds

minutes minutes (60 seconds)

hours hours (3600 seconds)

days days (3600*24 seconds)

weeks weeks (3600*24*7 seconds)

months months (3600*24*30 seconds)

quarters quarters (3600*24*91 seconds)

years years (3600*24*365 seconds)

Any abbreviation of these modifiers is also permitted (i.e. **seconds** may be specified as **sec** or **s**. A specification of **m** will be taken as months.

The specification of a time my have as may number/modifier parts as you wish. For example:

```
1 week 2 days 3 hours 10 mins
1 month 2 days 30 sec
```

are valid date specifications (beginning with version 1.35.1).

Note! in Bacula version 1.31 and below, the modifier was optional. It is now manditory.

Resource Types

The following table lists all current Bacula resource types. It shows what resources must be defined for each service (daemon). The default configuration files will already contain at least one example of each permitted resource, so you need not worry about creating all these kinds of resources from scratch.

Resource	Director	Client	Storage	Console
Catalog	Yes	No	No	No
Client	Yes	Yes	No	No
Console	Yes	No	No	Yes
Device	No	No	Yes	No
Director	Yes	Yes	Yes	Yes
FileSet	Yes	No	No	No
Job	Yes	No	No	No
JobDefs	Yes	No	No	No
Message	Yes	Yes	Yes	No
Pool	Yes	No	No	No
Schedule	Yes	No	No	No
Storage	Yes	No	Yes	No

Names, Passwords and Authorization

In order for one daemon to contact another daemon, it must authorize itself with a password. In most cases, the password corresponds to a particular name, so both the name and the password must match to be authorized.

The default configuration files are automatically defined for correct authorization with random passwords. If you add to or modify these files, you will need to take care to keep them consistent.

Here is sort of a picture of what names/passwords in which files/Resources must match up:

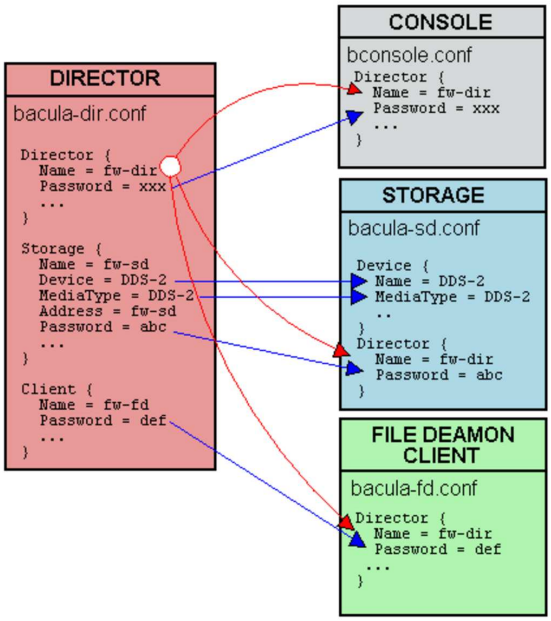

In the left column, you will find the Director, Storage, and Client resources, with their names and passwords – these are all in **bacula-dir.conf**. In the right column are where the corresponding values should be found in the Console, Storage daemon (SD), and File daemon (FD) configuration files.

Please note that the Address, **fd-sd**, that appears in the Storage resource of the Director, preceded with and asterisk in the above example, is passed to the File daemon in symbolic form. The File daemon then resolves it to an IP address. For this reason, you must use either an IP address or a fully qualified name. A name such as **localhost**, not being a fully qualified name, will resolve in the File daemon to the localhost of the File daemon, which is most likely not what is desired. The password used for the File daemon to

authorize with the Storage daemon is a temporary password unique to each Job created by the daemons and is not specified in any .conf file.

Detailed Information for each Daemon

The details of each Resource and the directives permitted therein are described in the following chapters.

The following configuration files must be defined:

- Console – to define the resources for the Console program (user interface to the Director). It defines which Directors are available so that you may interact with them.
- Director – to define the resources necessary for the Director. You define all the Clients and Storage daemons that you use in this configuration file.
- Client – to define the resources for each client to be backed up. That is, you will have a separate Client resource file on each machine that runs a File daemon.
- Storage – to define the resources to be used by each Storage daemon. Normally, you will have a single Storage daemon that controls your tape drive or tape drives. However, if you have tape drives on several machines, you will have at least one Storage daemon per machine.

Configuring the Director

Of all the configuration files needed to run **Bacula**, the Director's is the most complicated, and the one that you will need to modify the most often as you add clients or modify the FileSets.

For a general discussion of configuration file and resources including the data types recognized by **Bacula**. Please see the Configuration chapter of this manual.

Director Resource Types

Director resource type may be one of the following:

Job, JobDefs, Client, Storage, Catalog, Schedule, FileSet, Pool, Director, or Messages. We present them here in the most logical order for defining them:

- Director – to define the Director's name and its access password used for authenticating the Console program. Only a single Director resource definition may appear in the Director's configuration file. If you have either **/dev/random** or **bc** on your machine, Bacula will generate a random password during the configuration process, otherwise it will be left blank.
- Job – to define the backup/restore Jobs and to tie together the Client, FileSet and Schedule resources to be used for each Job.
- JobDefs – optional resource for providing defaults for Job resources.
- Schedule – to define when a Job is to be automatically run by **Bacula's** internal scheduler.
- FileSet – to define the set of files to be backed up for each Client.
- Client – to define what Client is to be backed up.
- Storage – to define on what physical device the Volumes should be mounted.
- Pool – to define what the pool of Volumes that can be used for a particular Job.
- Catalog – to define in what database to keep the list of files and the Volume names where they are backed up.

- Messages – to define where error and information messages are to be sent or logged.

The Director Resource

The Director resource defines the attributes of the Directors running on the network. In the current implementation, there is only a single Director resource, but the final design will contain multiple Directors to maintain index and media database redundancy.

Director Start of the Director resource. One and only one director resource must be supplied.

Name = <name> The director name used by the system administrator. This directive is required.

Description = <text> The text field contains a description of the Director that will be displayed in the graphical user interface. This directive is optional.

Password = <UA-password> Specifies the password that must be supplied for the default Bacula Console to be authorized. The same password must appear in the **Director** resource of the Console configuration file. For added security, the password is never actually passed across the network but rather a challenge response hash code created with the password. This directive is required. If you have either **/dev/random bc** on your machine, Bacula will generate a random password during the configuration process, otherwise it will be left blank and you must manually supply it.

Messages = <Messages-resource-name> The messages resource specifies where to deliver Director messages that are not associated with a specific Job. Most messages are specific to a job and will be directed to the Messages resource specified by the job. However, there are a few messages that can occur when no job is running. This directive is required.

Working Directory = <Directory> This directive is mandatory and specifies a directory in which the Director may put its status files. This directory should be used only by Bacula but may be shared by other Bacula daemons. Standard shell expansion of the **Directory** is done when the configuration file is read so that values such as **$HOME** will be properly expanded. This directive is required.

Pid Directory = <Directory> This directive is mandatory and specifies a directory in which the Director may put its process Id file files. The process Id file is used to shutdown Bacula and to prevent multiple copies of Bacula from running simultaneously. Standard shell expansion of the **Directory** is done when the configuration file is read so that values such as **$HOME** will be properly expanded.

Typically on Linux systems, you will set this to: **/var/run**. If you are not installing Bacula in the system directories, you can use the **Working Directory** as defined above. This directive is required.

QueryFile = <Path> This directive is mandatory and specifies a directory and file in which the Director can find the canned SQL statements for the **Query** command of the Console. Standard shell expansion of the **Path** is done when the configuration file is read so that values such as **$HOME** will be properly expanded. This directive is required.

Maximum Concurrent Jobs = <number> where <number> is the maximum number of total Director Jobs that should run concurrently. The default is set to 1, but you may set it to a larger number.

Please note that the Volume format becomes much more complicated with multiple simultaneous jobs, consequently, restores can take much longer if Bacula must sort through interleaved volume blocks from multiple simultaneous jobs. This can be avoided by having each simultaneously running job write to a different volume or by using data spooling, which will first spool the data to disk simultaneously, then write each spool file to the volume in sequence.

There may also still be some cases where directives such as **Maximum Volume Jobs** are not properly synchronized with multiple simultaneous jobs (subtle timing issues can arise), so careful testing is recommended.

At the current time, there is no configuration parameter set or limit the number console connections. A maximum of five simultaneous console connections are permitted.

For more details on getting concurrent jobs to run, please see Running Concurrent Jobs in the Tips chapter of this manual.

FD Connect Timeout = <time> where **time** is the time that the Director should continue attempting to contact the File daemon to start a job, and after which the Director will cancel the job. The default is 30 minutes.

SD Connect Timeout = <time> where **time** is the time that the Director should continue attempting to contact the Storage daemon

to start a job, and after which the Director will cancel the job. The default is 30 minutes.

DirAddresses = <IP-address-specification> Specify the ports and addresses on which the Director daemon will listen for Bacula Console connections. Probably the simplest way to explain is to show an example:

```
DirAddresses  = { ip = {
       addr = 1.2.3.4; port = 1205; }
   ipv4 = {
       addr = 1.2.3.4; port = http; }
   ipv6 = {
       addr = 1.2.3.4;
       port = 1205;
   }
   ip = {
       addr = 1.2.3.4
       port = 1205
   }
   ip = {
       addr = 1.2.3.4
   }
   ip = {
       addr = 201:220:222::2
   }
   ip = {
       addr = bluedot.thun.net
   }
}
```

where ip, ip4, ip6, addr, and port are all keywords. Note, that the address can be specified as either a dotted quadruple, or IPv6 colon notation, or as a symbolic name (only in the ip specification). Also, port can be specified as a number or as the mnemonic value from the /etc/services file. If a port is not specified, the default will be used. If an ip section is specified, the resolution can be made either by IPv4 or IPv6. If ip4 is specified, then only IPv4 resolutions will be permitted, and likewise with ip6.

DIRport = <port-number> Specify the port (a positive integer) on which the Director daemon will listen for Bacula Console connections. This same port number must be specified in the Director resource of the Console configuration file. The default is 9101, so normally this directive need not be specified. This directive is not needed if you specify DirAddresses.

DirAddress = <IP-Address> This directive is optional, but if it is specified, it will cause the Director server (for the Console program)

to bind to the specified **IP-Address**, which is either a domain name or an IP address specified as a dotted quadruple in string or quoted string format. If this directive is not specified, the Director will bind to any available address (the default). Note, unlike the DirAddresses specification noted above, this directive only permits a single address to be specified. This directive is not needed if you specify a DirAddresses (not plural).

The following is an example of a valid Director resource definition:

```
Director {
  Name = HeadMan
  WorkingDirectory = "$HOME/bacula/bin/working"
  Password = UA_password
  PidDirectory = "$HOME/bacula/bin/working"
  QueryFile = "$HOME/bacula/bin/query.sql"
  Messages = Standard
}
```

The Job Resource

The Job resource defines a Job (Backup, Restore, ...) that Bacula must perform. Each Job resource definition contains the names of the Clients and their FileSets to backup or restore, the Schedule for the Job, where the data are to be stored, and what media Pool can be used. In effect, each Job resource must specify What, Where, How, and When or FileSet, Storage, Backup/Restore/Level, and Schedule respectively.

Only a single type (**Backup**, **Restore**, ...) can be specified for any job. If you want to backup multiple FileSets on the same Client or multiple Clients, you must define a Job for each one.

Job Start of the Job resource. At least one Job resource is required.

Name = <name> The Job name. This name can be specified on the **Run** command in the console program to start a job. If the name contains spaces, it must be specified between quotes. It is generally a good idea to give your job the same name as the Client that it will backup. This permits easy identification of jobs.

When the job actually runs, the unique Job Name will consist of the name you specify here followed by the date and time the job was scheduled for execution. This directive is required.

Type = <job-type> The **Type** directive specifies the Job type, which may be one of the following: **Backup**, **Restore**, **Verify**, or **Admin**. This directive is required. Within a particular Job Type, there are also Levels as discussed in the next item.

Backup Run a backup Job. Normally you will have at least one Backup job for each client you want to save. Normally, unless you turn off cataloging, most all the important statistics and data concerning files backed up will be placed in the catalog.

Restore Run a restore Job. Normally, you will specify only one Restore job which acts as a sort of prototype that you will modify using the console program in order to perform restores. Although certain basic information from a Restore job is saved in the catalog, it is very minimal compared to the information stored for a Backup job – for example, no File database entries are generated since no Files are saved.

Verify Run a verify Job. In general, **verify** jobs permit you to compare the contents of the catalog to the file system, or to what was backed up. In addition, to verifying that a tape that was written can be read, you can also use **verify** as a sort of tripwire intrusion detection.

Admin Run a admin Job. An **Admin** job can be used to periodically run catalog pruning, if you do not want to do it at the end of each **Backup** Job. Although an Admin job is recorded in the catalog, very little data is saved.

Level = <job-level> The Level directive specifies the default Job level to be run. Each different Job Type (Backup, Restore, ...) has a different set of Levels that can be specified. The Level is normally overridden by a different value that is specified in the **Schedule** resource. This directive is not required, but must be specified either by a **Level** directive or as a override specified in the **Schedule** resource.

For a **Backup** Job, the Level may be one of the following:

Full is all files in the FileSet whether or not they have changed.

Incremental is all files that have changed since the last successful backup of the specified FileSet. If the Director cannot find a previous Full backup then the job will be upgraded into a Full backup. When the Director looks for a "suitable" backup record in the catalog database, it looks for a previous Job with:

- The same Job name.
- The same Client name.

- The same FileSet (any change to the definition of the FileSet such as adding or deleting a file in the Include or Exclude sections constitutes a different FileSet.
- The Job was a Full, Differential, or Incremental backup.
- The Job terminated normally (i.e. did not fail or was not canceled).

If all the above conditions do not hold, the Director will upgrade the Incremental to a Full save. Otherwise, the Incremental backup will be performed as requested.

The File daemon (Client) decides which files to backup for an Incremental backup by comparing start time of the prior Job (Full, Differential, or Incremental) against the time each file was last "modified" (st_mtime) and the time its attributes were last "changed"(st_ctime). If the file was modified or its attributes changed on or after this start time, it will then be backed up.

Please note that some virus scanning software may change st_ctime while doing the scan. For exaple, if the the virus scanning program attempts to reset the access time (st_atime), which Bacula does not use, it will cause st_ctime to change and hence Bacula will backup the file during an Incremental or Differential backup. In the case of Sophos virus scanning, you can prevent it from resetting the access time (st_atime) and hence changing st_ctime by using the **--no-reset-atime** option. For other software, please see their manual.

When Bacula does an Incremental backup, all modified files that are still on the system are backed up. However, any file that has been deleted since the last Full backup remains in the Bacula catalog, which means that if between a Full save and the time you do a restore, some files are deleted, those deleted files will also be restored. The deleted files will no longer appear in the catalog after doing another Full save. However, to remove deleted files from the catalog during a Incremental backup is quite a time consuming process and not currently implemented in Bacula.

Differential is all files that have changed since the last successful Full backup of the specified FileSet. If the Director cannot find a previous Full backup or a suitable Full backup, then the Differential job will be upgraded into a Full backup. When the Director looks for a "suitable" Full backup record in the catalog database, it looks for a previous Job with:

- The same Job name.
- The same Client name.

- The same FileSet (any change to the definition of the FileSet such as adding or deleting a file in the Include or Exclude sections constitutes a different FileSet.
- The Job was a FULL backup.
- The Job terminated normally (i.e. did not fail or was not canceled).

If all the above conditions do not hold, the Director will upgrade the Differential to a Full save. Otherwise, the Differential backup will be performed as requested.

The File daemon (Client) decides which files to backup for a differential backup by comparing the start time of the prior Full backup Job against the time each file was last "modified" (st_mtime) and the time its attributes were last "changed" (st_ctime). If the file was modified or its attributs were changed on or after this start time, it will then be backed up. The start time used is displayed after the **Since** on the Job report. In rare cases, using the start time of the prior backup may cause some files to be backed up twice, but it ensures that no change is missed. As with the Incremental option, you shouldensure that the clocks on your server and client are synchronized or as close as possible to avoid the possibility of a file being skipped. Note, on versions 1.33 or greater Bacula automatically makes the necessary adjstments to the time between the server and the client so that the times Bacula uses are synchronized.

When Bacula does an Differential backup, all modified files that are still on the system are backed up. However, any file that has been deleted since the last Full backup remains in the Bacula catalog, which means that if between a Full save and the time you do a restore, some files are deleted, those deleted files will also be restored. The deleted files will no longer appear in the catalog after doing another Full save. However, to remove deleted files from the catalog during a Differential backup is quite a time consuming process and not currently implemented in Bacula.

For a **Restore** Job, no level need be specified.

For a **Verify** Job, the Level may be one of the following:

InitCatalog does a scan of the specified **FileSet** and stores the file attributes in the Catalog database. Since no file data is saved, you might ask why you would want to do this. It turns out to be a very simple and easy way to have a **Tripwire** like feature using **Bacula**. In other words, it allows you to save the state of a set of files defined by the **FileSet** and later check to see

if those files have been modified or deleted and if any new files have been added. This can be used to detect system intrusion. Typically you would specify a **FileSet** that contains the set of system files that should not change (e.g. /sbin, /boot, /lib, /bin, ...). Normally, you run the **InitCatalog** level verify one time when your system is first setup, and then once again after each modification (upgrade) to your system. Thereafter, when your want to check the state of your system files, you use a **Verify level = Catalog**. This compares the results of your **InitCatalog** with the current state of the files.

Catalog Compares the current state of the files against the state previously saved during an **InitCatalog**. Any discrepancies are reported. The items reported are determined by the **verify** options specified on the **Include** directive in the specified **FileSet** (see the **FileSet** resource below for more details). Typically this command will be run once a day (or night) to check for any changes to your system files.

Please note! If you run two Verify Catalog jobs on the same client at the same time, the results will certainly be incorrect. This is because Verify Catalog modifies the Catalog database while running in order to track new files.

VolumeToCatalog This level causes Bacula to read the file attribute data written to the Volume from the last Job. The file attribute data are compared to the values saved in the Catalog database and any differences are reported. This is similar to the **Catalog** level except that instead of comparing the disk file attributes to the catalog database, the attribute data written to the Volume is read and compared to the catalog database. Although the attribute data including the signatures (MD5 or SHA1) are compared the actual file data is not compared (it is not in the catalog).

Please note! If you run two Verify VolumeToCatalog jobs on the same client at the same time, the results will certainly be incorrect. This is because the Verify VolumeToCatalog modifies the Catalog database while running.

DiskToCatalog This level causes Bacula to read the files as they currently are on disk, and to compare the current file attributes with the attributes saved in the catalog from the last backup for the job specified on the **VerifyJob** directive. This level differs from the **Catalog** level described above by the fact that it compare not against a previous Verify job but against a previous backup. When you run this level, you must supply the verify options on your Include statements. Those options determine

what attribute fields are compared.

This command can be very useful if you have disk problems because it will compare the current state of your disk against the last successful backup, which may be several jobs.

Note, the current implementation (1.32c) does not identify files that have been deleted.

Verify Job = <Job-Resource-Name> If you run a verify job without this directive, the last job run will be compared with the catalog, which means that you must immediately follow a backup by a verify command. If you specify a **Verify Job** Bacula will find the last job with that name that ran. This permits you to run all your backups, then run Verify jobs on those that you wish to be verified (most often a **VolumeToCatalog**) so that the tape just written is re-read.

JobDefs = <JobDefs-Resource-Name> If a JobDefs-Resource-Name is specified, all the values contained in the named JobDefs resource will be used as the defaults for the current Job. Any value that you explicitly define in the current Job resource, will override any defaults specified in the JobDefs resource. The use of this directive permits writing much more compact Job resources where the bulk of the directives are defined in one or more JobDefs. This is particularly useful if you have many similar Jobs but with minor variations such as different Clients. A simple example of the use of JobDefs is provided in the default bacula-dir.conf file.

Bootstrap = <bootstrap-file> The Bootstrap directive specifies a bootstrap file that, if provided, will be used during **Restore** Jobs and is ignored in other Job types. The **bootstrap** file contains the list of tapes to be used in a restore Job as well as which files are to be restored. Specification of this directive is optional, and if specified, it is used only for a restore job. In addition, when running a Restore job from the console, this value can be changed.

If you use the **Restore** command in the Console program, to start a restore job, the **bootstrap** file will be created automatically from the files you select to be restored.

For additional details of the **bootstrap** file, please see Restoring Files with the Bootstrap File chapter of this manual.

Write Bootstrap = <bootstrap-file-specification> The **writebootstrap** directive specifies a file name where Bacula will write a **bootstrap** file for each Backup job run. Thus this directive applies only to Backup Jobs. If the Backup job is a Full save, Bacula will erase any current contents of the specified file before

writing the bootstrap records. If the Job is an Incremental save, Bacula will append the current bootstrap record to the end of the file.

Using this feature, permits you to constantly have a bootstrap file that can recover the current state of your system. Normally, the file specified should be a mounted drive on another machine, so that if your hard disk is lost, you will immediately have a bootstrap record available. Alternatively, you should copy the bootstrap file to another machine after it is updated.

If the **bootstrap-file-specification** begins with a vertical bar (—), Bacula will use the specification as the name of a program to which it will pipe the bootstrap record. It could for example be a shell script that emails you the bootstrap record.

For more details on using this file, please see the chapter entitled The Bootstrap File of this manual.

Client = <client-resource-name> The Client directive specifies the Client (File daemon) that will be used in the current Job. Only a single Client may be specified in any one Job. The Client runs on the machine to be backed up, and sends the requested files to the Storage daemon for backup, or receives them when restoring. For additional details, see the Client Resource section of this chapter. This directive is required.

FileSet = <FileSet-resource-name> The FileSet directive specifies the FileSet that will be used in the current Job. The FileSet specifies which directories (or files) are to be backed up, and what options to use (e.g. compression, ...). Only a single FileSet resource may be specified in any one Job. For additional details, see the FileSet Resource section of this chapter. This directive is required.

Messages = <messages-resource-name> The Messages directive defines what Messages resource should be used for this job, and thus how and where the various messages are to be delivered. For example, you can direct some messages to a log file, and others can be sent by email. For additional details, see the Messages Resource Chapter of this manual. This directive is required.

Pool = <pool-resource-name> The Pool directive defines the pool of Volumes where your data can be backed up. Many Bacula installations will use only the **Default** pool. However, if you want to specify a different set of Volumes for different Clients or different Jobs, you will probably want to use Pools. For additional details, see the Pool Resource section of this chapter. This resource is required.

Full Backup Pool = <pool-resource-name> The *Full Backup Pool* specifies a Pool to be used for Full backups. It will override any Pool specification during a Full backup. This resource is optional.

Differential Backup Pool = <pool-resource-name> The *Differential Backup Pool* specifies a Pool to be used for Differential backups. It will override any Pool specification during a Differentia backup. This resource is optional.

Incremental Backup Pool = <pool-resource-name> The *Incremental Backup Pool* specifies a Pool to be used for Incremental backups. It will override any Pool specification during a Incremental backup. This resource is optional.

Schedule = <schedule-name> The Schedule directive defines what schedule is to be used for the Job. The schedule determines when the Job will be automatically started and what Job level (i.e. Full, Incremental, ...) is to be run. This directive is optional, and if left out, the Job can only be started manually. For additional details, see the Schedule Resource Chapter of this manual. If a Schedule resource is specified, the job will be run according to the schedule specified. If no Schedule resource is specified for the Job, the job must be manually started using the Console program. Although you may specify only a single Schedule resource for any one job, the Schedule resource may contain multiple **Run** directives, which allow you to run the Job at many different times, and each **run** directive permits overriding the default Job Level Pool, Storage, and Messages resources. This gives considerable flexibility in what can be done with a single Job.

Storage = <storage-resource-name> The Storage directive defines the name of the storage services where you want to backup the FileSet data. For additional details, see the Storage Resource Chapter of this manual. This directive is required.

Max Start Delay = <time> The time specifies maximum delay between the scheduled time and the actual start time for the Job. For example, a job can be scheduled to run at 1:00am, but because other jobs are running, it may wait to run. If the delay is set to 3600 (one hour) and the job has not begun to run by 2:00am, the job will be canceled. This can be useful, for example, to prevent jobs from running during day time hours. The default is 0 which indicates no limit.

Max Run Time = <time> The time specifies maximum allowed time that a job may run, counted from the when the job starts (**not**

necessarily the same as when the job was scheduled). This directive is implemented only in version 1.33 and later.

Max Wait Time = <time> The time specifies maximum allowed time that a job may block waiting for a resource (such as waiting for a tape to be mounted, or waiting for the storage or file daemons to perform their duties), counted from the when the job starts (**not** necessarily the same as when the job was scheduled). This directive is implemented only in version 1.33 and later. Note, the implementation is not yet complete, so this directive does not yet work correctly.

Prune Jobs = <yes—no> Normally, pruning of Jobs from the Catalog is specified on a Client by Client basis in the Client resource with the **AutoPrune** directive. If this directive is specified (not normally) and the value is **yes**, it will override the value specified in the Client resource. The default is **no**.

Prune Files = <yes—no> Normally, pruning of Files from the Catalog is specified on a Client by Client basis in the Client resource with the **AutoPrune** directive. If this directive is specified (not normally) and the value is **yes**, it will override the value specified in the Client resource. The default is **no**.

Prune Volumes = <yes—no> Normally, pruning of Volumes from the Catalog is specified on a Client by Client basis in the Client resource with the **AutoPrune** directive. If this directive is specified (not normally) and the value is **yes**, it will override the value specified in the Client resource. The default is **no**.

Run Before Job = <command> The specified **command** is run as an external program prior to running the current Job. Any output sent by the job to standard output will be included in the Bacula job report. The command string must be a valid program name or name of a shell script. This directive is not required, but if it is defined, and if the exit code of the program run is non-zero, the current Bacula job will be canceled. In addition, the command string is parsed then feed to the execvp() function, which means that the path will be searched to execute your specified command, but there is no shell interpretation, as a consequence, if you complicated commands or want any shell features such as redirection or piping, you must call a shell script and do it inside that script.

Before submitting the specified command to the operating system, Bacula performs character substitution of the following characters:

```
%% = %
```

```
%c = Client's name
%d = Director's name
%i = JobId
%e = Job Exit Status
%j = Unique Job name
%l = Job Level
%n = Job name
%t = Job type
%v = Volume name
```

The Job Exit Status code %e edits the following values:

- OK
- Error
- Fatal Error
- Canceled
- Differences
- Unknown term code

Thus if you edit it on a command line, you will need to enclose it within some sort of quotes.

Bacula checks the exit status of the RunBeforeJob program. If it is non-zero, the job will be error terminated. Lutz Kittler has pointed out that this can be a simple way to modify your schedules during a holiday. For example, suppose that you normally do Full backups on Fridays, but Thursday and Friday are holidays. To avoid having to change tapes between Thursday and Friday when no one is in the office, you can create a RunBeforeJob that returns a non-zero status on Thursday and zero on all other days. That way, the Thursday job will not run, and on Friday the tape you insert on Wednesday before leaving will be used.

Run After Job = <command> The specified **command** is run as an external program after the current job terminates. This directive is not required. The command string must be a valid program name or name of a shell script. If the exit code of the program run is non-zero, the current Bacula job will terminate in error. Before submitting the specified command to the operating system, Bacula performs character substitution as described above for the **Run Before Job** directive.

An example of the use of this command is given in the Tips Chapter of this manual. As of version 1.30, Bacula checks the exit status of the RunAfter program. If it is non-zero, the job will be terminated in error.

Client Run Before Job = <command> This command is the same as **Run Before Job** except that it is run on the client machine. The same restrictions apply to Unix systems as noted above for the **Run Before Job**. In addition, for a Windows client on version 1.33 and above, please take careful note that you must ensure a correct path to your script, and the script or program can be a .com, .exe or a .bat file. However, if you specify a path, you must also specify the full extension. Unix like commands will not work unless you have installed and properly configured Cygwin in addition to and separately from Bacula.

Special Windows Considerations The command can be anything that cmd.exe or command.com will recognize as a executable file. Specifiying the executable's extention is optional, unless there is an ambiguity. (i.e. ls.bat, ls.exe)

The System %Path% will be searched for the command. (under the envrionment variable dialog you have have both System Environment and User Environment, we believe that only the System environment will be available to bacual-fd, if it is running as a service.)

System environment varaible can be called out using the %var% syntax and used as either part of the command name or arguments.

When specifiying a full path to an executable if the path or executable name contains whitespace or special characters they will need to be quoted. Arguments containing whitespace or special characters will also have to be quoted.

```
ClientRunBeforeJob = "\"C:/Program Files/Software
     Vendor/Executable\" /arg1 /arg2 \"foo bar\""
```

The special characters &()[]{}^=;!'+,‘~ will need to be quoted if part of a filename or argument.

If someone is logged in a blank "command" window running the commands will be present during the execution of the command.

Some Suggestions from Phil Stracchino for running on Win32 machines with the native Win32 File daemon:

1. You might want the ClientRunBeforeJob directive to specify a .bat file which runs the actual client-side commands, rather than trying to run (for example) regedit /e directly.
2. The batch file should explicitly 'exit 0' on successful completion.
3. The path to the batch file should be specified in Unix form:
   ClientRunBeforeJob = "c:/bacula/bin/systemstate.bat"

rather than DOS/Windows form:
ClientRunBeforeJob = "c:\bacula\bin\systemstate.bat"
INCORRECT

Client Run After Job = <command> This command is the same as **Run After Job** except that it is run on the client machine. Note, please see the notes above in **Client Run Before Job** concerning Windows clients.

Rerun Failed Levels = <yes—no> If this directive is set to **yes** (default no), and Bacula detects that a previous job at a higher level (i.e. Full or Differential) has failed, the current job level will be upgraded to the higher level. This is particularly useful for Laptops where they may often be unreachable, and if a prior Full save has failed, you wish the very next backup to be a Full save rather than whatever level it is started as.

Spool Data = <yes—no> If this directive is set to **yes** (default no), the Storage daemon will be requested to spool the data for this Job to disk rather than write it directly to tape. Once all the data arrives or the spool file maximum sizes are reached, the data will be despooled and written to tape. When this directive is set to yes, the Spool Attributes is also automatically set to yes. Spooling data prevents tape shoe-shine (start and stop) during Incremental saves. This option should not be used if you are writing to a disk file.

Spool Attributes = <yes—no> The default is set to **no**, which means that the File attributes are sent by the Storage daemon to the Director as they are stored on tape. However, if you want to avoid the possibility that database updates will slow down writing to the tape, you may want to set the value to **yes**, in which case the Storage daemon will buffer the File attributes and Storage coordinates to a temporary file in the Working Directory, then when writing the Job data to the tape is completed, the attributes and storage coordinates will be sent to the Director. The default is **no**.

Where = <directory> This directive applies only to a Restore job and specifies a prefix to the directory name of all files being restored. This permits files to be restored in a different location from which they were saved. If **Where** is not specified or is set to backslash (/), the files will be restored to their original location. By default, we have set **Where** in the example configuration files to be **/tmp/bacula-restores**. This is to prevent accidental overwriting of your files.

Replace = <replace-option> This directive applies only to a Restore job and specifies what happens when Bacula wants to restore a file

or directory that already exists. You have the following options for **replace-option**:

always when the file to be restored already exists, it is deleted then replaced by the copy backed up.

ifnewer if the backed up file (on tape) is newer than the existing file, the existing file is deleted and replaced by the back up.

ifolder if the backed up file (on tape) is older than the existing file, the existing file is deleted and replaced by the back up.

never if the backed up file already exists, Bacula skips restoring this file.

Prefix Links=<yes—no> If a **Where** path prefix is specified for a recovery job, apply it to absolute links as well. The default is **No**. When set to **Yes** then while restoring files to an alternate directory, any absolute soft links will also be modified to point to the new alternate directory. Normally this is what is desired – i.e. everything is self consistent. However, if you wish to later move the files to their original locations, all files linked with absolute names will be broken.

Maximum Concurrent Jobs = <number> where <number> is the maximum number of Jobs from the current Job resource that can run concurrently. Note, this directive limits only Jobs with the same name as the resource in which it appears. Any other restrictions on the maximum concurrent jobs such as in the Director, Client, or Storage resources will also apply in addition to the limit specified here. The default is set to 1, but you may set it to a larger number. We strongly recommend that you read the WARNING documented under Maximum Concurrent Jobs in the Director's resource.

Reschedule On Error = <yes—no> If this directive is enabled, and the job terminates in error, the job will be rescheduled as determined by the **Reschedule Interval** and **Reschedule Times** directives. If you cancel the job, it will not be rescheduled. The default is **no** (i.e. the job will not be rescheduled).

This specification can be useful for portables, laptops, or other machines that are not always connected to the network or switched on.

Reschedule Interval = <time-specification> If you have specified **Reschedule On Error = yes** and the job terminates in error, it will be rescheduled after the interval of time specified by **time-specification**. See the time specification formats in the Configure chapter for details of time specifications. If no interval is specified, the job will not be rescheduled on error.

Reschedule Times = <count> This directive specifies the maximum number of times to reschedule the job. If it is set to zero (the default) the job will be rescheduled an indefinite number of times.

Priority = <number> This directive permits you to control the order in which your jobs run by specifying a positive non-zero number. The higher the number, the lower the job priority. Assuming you are not running concurrent jobs, all queued jobs of priority 1 will run before queued jobs of priority 2 and so on, regardless of the original scheduling order.

The priority only affects waiting jobs that are queued to run, not jobs that are already running. If one or more jobs of priority 2 are already running, and a new job is scheduled with priority 1, the currently running priority 2 jobs must complete before the priority 1 job is run.

The default priority is 10.

If you want to run concurrent jobs, which is not recommended, you should keep these points in mind:

- To run concurrent jobs, you must set Maximum Concurrent Jobs = 2 in 5 or 6 distinct places: in bacula-dir.conf in the Director, the Job, the Client, the Storage resources; in bacula-fd in the FileDaemon (or Client) resource, and in bacula-sd.conf in the Storage resource. If any one is missing, it will throttle the jobs to one at a time.
- Bacula concurrently runs jobs of only one priority at a time. It will not simultaneously run a priority 1 and a priority 2 job.
- If Bacula is running a priority 2 job and a new priority 1 job is scheduled, it will wait until the running priority 2 job terminates even if the Maximum Concurrent Jobs settings would otherwise allow two jobs to run simultaneously.
- Suppose that bacula is running a priority 2 job and new priority 1 job is scheduled and queued waiting for the running priority 2 job to terminate. If you then start a second priority 2 job, the waiting priority 1 job will prevent the new priority 2 job from running concurrently with the running priority 2 job. That is: as long as there is a higher priority job waiting to run, no new lower priority jobs will start even if the Maximum Concurrent Jobs settings would normally allow them to run. This ensures that higher priority jobs will be run as soon as possible.

If you have several jobs of different priority, it is best not to start them at exactly the same time, because Bacula must examine them one at a time. If by chance Bacula treats a lower priority first, then it will

run before your high priority jobs. To avoid this, start any higher priority a few seconds before lower ones. This insures that Bacula will examine the jobs in the correct order, and that your priority scheme will be respected.

The following is an example of a valid Job resource definition:

```
Job {
  Name = "Minou"
  Type = Backup
  Level = Incremental                 # default
  Client = Minou
  FileSet="Minou Full Set"
  Storage = DLTDrive
  Pool = Default
  Schedule = "MinouWeeklyCycle"
  Messages = Standard
}
```

The JobDefs Resource

The JobDefs resource permits all the same directives that can appear in a Job resource. However, a JobDefs resource does not create a Job, rather it can be referenced within a Job to provide defaults for that Job. This permits you to concisely define several nearly identical Jobs, each one referencing a JobDefs resource which contains the defaults. Only the changes from the defaults need be mentioned in each Job.

The Schedule Resource

The Schedule resource provides a means of automatically scheduling a Job as well as the ability to override the default Level, Pool, Storage and Messages resources. If a Schedule resource is not referenced in a Job, the Job may only be run manually. In general, you specify an action to be taken and when.

Schedule Start of the Schedule directives. No **Schedule** resource is required, but you will need at least one if you want Jobs to be automatically started.

Name = <name> The name of the schedule being defined. The Name directive is required.

Run = <Job-overrides> <Date-time-specification> The Run directive defines when a Job is to be run, and what overrides if any to apply. You may specify multiple **run** directives within a **Schedule** resource. If you do, they will all be applied (i.e. multiple schedules). If you have two **Run** directives that start at the same time, two Jobs will start at the same time (well, within one second of each other).

The **Job-overrides** permit overriding the Level, the Storage, the Messages, and the Pool specifications provided in the Job resource. In addition, the FullPool, the IncrementalPool, and the DifferentialPool specifications permit overriding the Pool specification according to what backup Job Level is in effect.

By the use of overrides, you may customize a particular Job. For example, you may specify a Messages override for your Incremental backups that outputs messages to a log file, but for your weekly or monthly Full backups, you may send the output by email by using a different Messages override.

Job-overrides are specified as: **keyword=value** where the keyword is Level, Storage, Messages, Pool, FullPool, DifferentialPool, or IncrementalPool, and the **value** is as defined on the respective directive formats for the Job resource. You may specify multiple **Job-overrides** on one **Run** directive by separating them with one or more spaces or by separating them with a trailing comma. For example:

Level=Full is all files in the FileSet whether or not they have changed.

Level=Incremental is all files that have changed since the last backup.

Pool=Weekly specifies to use the Pool named **Weekly**.

Storage=DLT_Drive specifies to use **DLT_Drive** for the storage device.

Messages=Verbose specifies to use the **Verbose** message resource for the Job.

FullPool=Full specifies to use the Pool named **Full** if the job is a full backup, or is upgraded from another type to a full backup.

DifferentialPool=Differential specifies to use the Pool named **Differential** if the job is a differential backup.

IncrementalPool=Incremental specifies to use the Pool named **Incremental** if the job is an incremental backup.

SpoolData=yes—no tells Bacula to request the Storage daemon to spool data to a disk file before putting it on tape.

WritePartAfterJob=yes—no tells Bacula to request the Storage daemon to write the current part file to the device when the job is finished (see Write Part After Job directive in the Job resource). Please note, this directive is implemented only in version 1.37 and later.

Date-time-specification determines when the Job is to be run. The specification is a repetition, and as a default Bacula is set to run a job at the beginning of the hour of every hour of every day of every week of every month of every year. This is not normally what you want, so you must specify or limit when you want the job to run. Any specification given is assumed to be repetitive in nature and will serve to override or limit the default repetition. This is done by specifing masks or times for the hour, day of the month, day of the week, week of the month, week of the year, and month when you want the job to run. By specifying one or more of the above, you can define a schedule to repeat at almost any frequency you want.

Basically, you must supply a **month**, **day**, **hour**, and **minute** the Job is to be run. Of these four items to be specified, **day** is special in that you may either specify a day of the month such as 1, 2, ... 31, or you may specify a day of the week such as Monday, Tuesday, ... Sunday. Finally, you may also specify a week qualifier to restrict the schedule to the first, second, third, fourth, or fifth week of the month.

For example, if you specify only a day of the week, such as **Tuesday** the Job will be run every hour of every Tuesday of every Month. That is the **month** and **hour** remain set to the defaults of every month and all hours.

Note, by default with no other specification, your job will run at the beginning of every hour. If you wish your job to run more than once in any given hour, you will need to specify multiple **run** specifications each with a different minute.

The date/time to run the Job can be specified in the following way in pseudo-BNF:

```
<void-keyword>    = on
<at-keyword>      = at
<week-keyword>    = 1st | 2nd | 3rd | 4th | 5th | first |
                    second | third | forth | fifth
<wday-keyword>    = sun | mon | tue | wed | thu | fri | sat |
                    sunday | monday | tuesday | wednesday |
                    thursday | friday
<week-of-year-keyword> = w00 | w01 | ... w52 | w53
<month-keyword>   = jan | feb | mar | apr | may | jun | jul |
                    aug | sep | oct | nov | dec | january |
                    february | ... | december
```

```
<daily-keyword>   = daily
<weekly-keyword>  = weekly
<monthly-keyword> = monthly
<hourly-keyword>  = hourly
<digit>           = 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0
<number>          = <digit> | <digit><number>
<12hour>          = 0 | 1 | 2 | ... 12
<hour>            = 0 | 1 | 2 | ... 23
<minute>          = 0 | 1 | 2 | ... 59
<day>             = 1 | 2 | ... 31
<time>            = <hour>:<minute> |
                    <12hour>:<minute>am |
                    <12hour>:<minute>pm
<time-spec>       = <at-keyword> <time> |
                    <hourly-keyword>
<date-keyword>    = <void-keyword>  <weekly-keyword>
<day-range>       = <day>-<day>
<month-range>     = <month-keyword>-<month-keyword>
<wday-range>      = <wday-keyword>-<wday-keyword>
<range>           = <day-range> | <month-range> |
                          <wday-range>
<date>            = <date-keyword> | <day> | <range>
<date-spec>       = <date> | <date-spec>
<day-spec>        = <day> | <wday-keyword> |
                    <day-range> | <wday-range> |
                    <daily-keyword>
<day-spec>        = <day> | <wday-keyword> |
                    <week-keyword> <wday-keyword>
<month-spec>      = <month-keyword> | <month-range> |
                    <monthly-keyword>
<date-time-spec>  = <month-spec> <day-spec> <time-spec>
```

Note, the Week of Year specification wnn follows the ISO standard definition of the week of the year, where Week 1 is the week in which the first Thursday of the year occurs, or alternatively, the week which contains the 4th of January. Weeks are numbered w01 to w53. w00 for Bacula is the week that precedes the first ISO week (i.e. has the first few days of the year if any occur before Thursday). w00 is not defined by the ISO specification. A week starts with Monday and ends with Sunday.

An example schedule resource that is named **WeeklyCycle** and runs a job with level full each Sunday at 1:05am and an incremental job Monday through Saturday at 1:05am is:

```
Schedule {
  Name = "WeeklyCycle"
  Run = Level=Full sun at 1:05
  Run = Level=Incremental mon-sat at 1:05
}
```

An example of a possible monthly cycle is as follows:

```
Schedule {
  Name = "MonthlyCycle"
  Run = Level=Full Pool=Monthly 1st sun at 1:05
  Run = Level=Differential 2nd-5th sun at 1:05
  Run = Level=Incremental Pool=Daily mon-sat at 1:05
}
```

The first of every month:

```
Schedule {
  Name = "First"
  Run = Level=Full on 1 at 1:05
  Run = Level=Incremental on 2-31 at 1:05
}
```

Every 10 minutes:

```
Schedule {
  Name = "TenMinutes"
  Run = Level=Full hourly at 0:05
  Run = Level=Full hourly at 0:15
  Run = Level=Full hourly at 0:25
  Run = Level=Full hourly at 0:35
  Run = Level=Full hourly at 0:45
  Run = Level=Full hourly at 0:55
}
```

Technical Notes on Schedules

Internally Bacula keeps a schedule as a bit mask. There are six masks and a minute field to each schedule. The masks are hour, day of the month (mday), month, day of the week (wday), week of the month (wom), and week of the year (woy). The schedule is initialized to have the bits of each of these masks set, which means that at the beginning of every hour, the job will run. When you specify a month for the first time, the mask will be cleared and the bit corresponding to your selected month will be selected. If you specify a second month, the bit corresponding to it will also be added to the mask. Thus when Bacula checks the masks to see if the bits are set corresponding to the current time, your job will run only in the two months you have set. Likewise, if you set a time (hour), the hour mask will be cleared, and the hour you specify will be set in the bit mask and the minutes will be stored in the minute field.

For any schedule you have defined, you can see how these bits are set by doing a **show schedules** command in the Console program. Please note that the bit mask is zero based, and Sunday is the first day of the week (bit zero).

The Client Resource

The Client resource defines the attributes of the Clients that are served by this Director; that is the machines that are to be backed up. You will need one Client resource definition for each machine to be backed up.

Client (or FileDaemon) Start of the Client directives.

Name = <name> The client name which will be used in the Job resource directive or in the console run command. This directive is required.

Address = <address> Where the address is a host name, a fully qualified domain name, or a network address in dotted quad notation for a Bacula File server daemon. This directive is required.

FD Port = <port-number> Where the port is a port number at which the Bacula File server daemon can be contacted. The default is 9102.

Catalog = <Catalog-resource-name> This specifies the name of the catalog resource to be used for this Client. This directive is required.

Password = <password> This is the password to be used when establishing a connection with the File services, so the Client configuration file on the machine to be backed up must have the same password defined for this Director. This directive is required. If you have either **/dev/random bc** on your machine, Bacula will generate a random password during the configuration process, otherwise it will be left blank.

File Retention = <time-period-specification> The File Retention directive defines the length of time that Bacula will keep File records in the Catalog database. When this time period expires, and if **AutoPrune** is set to **yes** Bacula will prune (remove) File records that are older than the specified File Retention period. Note, this affects only records in the catalog database. It does not effect your archive backups.

File records may actually be retained for a shorter period than you specify on this directive if you specify either a shorter **Job Retention** or shorter **Volume Retention** period. The shortest retention period

of the three takes precedence. The time may be expressed in seconds, minutes, hours, days, weeks, months, quarters, or years. See the Configuration chapter of this manual for additional details of time specification.

The default is 60 days.

Job Retention = <time-period-specification> The Job Retention directive defines the length of time that Bacula will keep Job records in the Catalog database. When this time period expires, and if **AutoPrune** is set to **yes** Bacula will prune (remove) Job records that are older than the specified File Retention period. As with the other retention periods, this affects only records in the catalog and not data in your archive backup.

If a Job record is selected for pruning, all associated File and JobMedia records will also be pruned regardless of the File Retention period set. As a consequence, you normally will set the File retention period to be less than the Job retention period. The Job retention period can actually be less than the value you specify here if you set the **Volume Retention** directive in the Pool resource to a smaller duration. This is because the Job retention period and the Volume retention period are independently applied, so the smaller of the two takes precedence.

The Job retention period is specified as seconds, minutes, hours, days, weeks, months, quarters, or years. See the Configuration chapter of this manual for additional details of time specification.

The default is 180 days.

AutoPrune = <yes—no> If AutoPrune is set to **yes** (default), Bacula (version 1.20 or greater) will automatically apply the File retention period and the Job retention period for the Client at the end of the Job. If you set **AutoPrune = no**, pruning will not be done, and your Catalog will grow in size each time you run a Job. Pruning affects only information in the catalog and not data stored in the backup archives (on Volumes).

Maximum Concurrent Jobs = <number> where <number> is the maximum number of Jobs with the current Client that can run concurrently. Note, this directive limits only Jobs for Clients with the same name as the resource in which it appears. Any other restrictions on the maximum concurrent jobs such as in the Director, Job, or Storage resources will also apply in addition to any limit specified here. The default is set to 1, but you may set it to a larger number. We strongly recommend that you read the WARNING documented under Maximum Concurrent Jobs in the Director's resource.

***Priority = <number>** The number specifies the priority of this client relative to other clients that the Director is processing simultaneously. The priority can range from 1 to 1000. The clients are ordered such that the smaller number priorities are performed first (not currently implemented).

The following is an example of a valid Client resource definition:

```
Client {
  Name = Minimatou
  Address = minimatou
  Catalog = MySQL
  Password = very_good
}
```

The Storage Resource

The Storage resource defines which Storage daemons are available for use by the Director.

Storage Start of the Storage resources. At least one storage resource must be specified.

Name = <name> The name of the storage resource. This name appears on the Storage directive specified in the Job directive and is required.

Address = <address> Where the address is a host name, a **fully qualified domain name**, or an **IP address**. Please note that the <address> as specified here will be transmitted to the File daemon who will then use it to contact the Storage daemon. Hence, it is **not**, a good idea to use **localhost** as the name but rather a fully qualified machine name or an IP address. This directive is required.

SD Port = <port> Where port is the port to use to contact the storage daemon for information and to start jobs. This same port number must appear in the Storage resource of the Storage daemon's configuration file. The default is 9103.

Password = <password> This is the password to be used when establishing a connection with the Storage services. This same password also must appear in the Director resource of the Storage daemon's configuration file. This directive is required. If you have either **/dev/random bc** on your machine, Bacula will generate a random password during the configuration process, otherwise it will be left blank.

Device = <device-name> This directive specifies the name of the device to be used to for the storage. This name is not the physical device name, but the logical device name as defined on the **Name** directive contained in the **Device** resource definition of the **Storage daemon** configuration file. You can specify any name you would like (even the device name if you prefer) up to a maximum of 127 characters in length. The physical device name associated with this device is specified in the **Storage daemon** configuration file (as **Archive Device**). Please take care not to define two different Storage resource directives in the Director that point to the same Device in the Storage daemon. Doing so may cause the Storage daemon to block (or hang) attempting to open the same device that is already open. This directive is required.

Media Type = <MediaType> This directive specifies the Media Type to be used to store the data. This is an arbitrary string of characters up to 127 maximum that you define. It can be anything you want. However, it is best to make it descriptive of the storage media (e.g. File, DAT, "HP DLT8000", 8mm, ...). In addition, it is essential that you make the **Media Type** specification unique for each storage media type. If you have two DDS-4 drives that have incompatible formats, or if you have a DDS-4 drive and a DDS-4 autochanger, you almost certainly should specify different **Media Types**. During a restore, assuming a **DDS-4** Media Type is associated with the Job, Bacula can decide to use any Storage daemon that support Media Type **DDS-4** and on any drive supports it. If you want to tie Bacula to using a single Storage daemon or drive, you must specify a unique Media Type for that drive. This is an important point that should be carefully understood. You can find more on this subject in the Basic Volume Management chapter of this manual.

The **MediaType** specified here, **must** correspond to the **Media Type** specified in the **Device** resource of the **Storage daemon** configuration file. This directive is required, and it is used by the Director and the Storage daemon to ensure that a Volume automatically selected from the Pool corresponds to the physical device. If a Storage daemon handles multiple devices (e.g. will write to various file Volumes on different partitions), this directive allows you to specify exactly which device.

As mentioned above, the value specified in the Director's Storage resource must agree with the value specified in the Device resource in the **Storage daemon's** configuration file. It is also an additional check so that you don't try to write data for a DLT onto an 8mm device.

Autochanger = <yes—no> If you specify **yes** for this command (the default is **no**), when you use the **label** command or the **add** command to create a new Volume, **Bacula** will also request the Autochanger Slot number. This simplifies creating database entries for Volumes in an autochanger. If you forget to specify the Slot, the autochanger will not be used. However, you may modify the Slot associated with a Volume at any time by using the **update volume** command in the console program. When **autochanger** is enabled, the algorithm used by Bacula to search for available volumes will be modified to consider only Volumes that are known to be in the autochanger's magazine. If no **in changer** volume is found, Bacula will attempt recycling, pruning, ..., and if still no volume is found, Bacula will search for any volume whether or not in the magazine. By privileging in changer volumes, this procedure minimizes operator intervention. The default is **no**.

For the autochanger to be used, you must also specify **Autochanger = yes** in the Device Resource in the Storage daemon's configuration file as well as other important Storage daemon configuration information. Please consult the Using Autochangers manual of this chapter for the details of using autochangers.

Maximum Concurrent Jobs = <number> where <number> is the maximum number of Jobs with the current Storage resource that can run concurrently. Note, this directive limits only Jobs for Jobs using this Storage daemon. Any other restrictions on the maximum concurrent jobs such as in the Director, Job, or Client resources will also apply in addition to any limit specified here. The default is set to 1, but you may set it to a larger number. We strongly recommend that you read the WARNING documented under Maximum Concurrent Jobs in the Director's resource.

While it is possible to set the Director's, Job's, or Client's maximum concurrent jobs greater than one, you should take great care in setting the Storage daemon's greater than one. By keeping this directive set to one, you will avoid having two jobs simultaneously write to the same Volume. Although this is supported, it is not currently recommended.

The following is an example of a valid Storage resource definition:

```
# Definition of tape storage device
Storage {
  Name = DLTDrive
  Address = lpmatou
  Password = storage_password # password for Storage daemon
  Device = "HP DLT 80"    # same as Device in Storage daemon
```

```
  Media Type = DLT8000    # same as MediaType in Storage daemon
}
```

The Pool Resource

The Pool resource defines the set of storage Volumes (tapes or files) to be used by Bacula to write the data. By configuring different Pools, you can determine which set of Volumes (media) receives the backup data. This permits, for example, to store all full backup data on one set of Volumes and all incremental backups on another set of Volumes. Alternatively, you could assign a different set of Volumes to each machine that you backup. This is most easily done by defining multiple Pools.

Another important aspect of a Pool is that it contains the default attributes (Maximum Jobs, Retention Period, Recycle flag, ...) that will be given to a Volume when it is created. This avoids the need for you to answer a large number of questions when labeling a new Volume. Each of these attributes can later be changed on a Volume by Volume basis using the **update** command in the console program. Note that you must explicitly specify which Pool Bacula is to use with each Job. Bacula will not automatically search for the correct Pool.

Most often in Bacula installations all backups for all machines (Clients) go to a single set of Volumes. In this case, you will probably only use the **Default** Pool. If your backup strategy calls for you to mount a different tape each day, you will probably want to define a separate Pool for each day. For more information on this subject, please see the Backup Strategies chapter of this manual.

To use a Pool, there are three distinct steps. First the Pool must be defined in the Director's configuration file. Then the Pool must be written to the Catalog database. This is done automatically by the Director each time that it starts, or alternatively can be done using the **create** command in the console program. Finally, if you change the Pool definition in the Director's configuration file and restart Bacula, the pool will be updated alternatively you can use the **update pool** console command to refresh the database image. It is this database image rather than the Director's resource image that is used for the default Volume attributes. Note, for the pool to be automatically created or updated, it must be explicitly referenced by a Job resource.

Next the physical media must be labeled. The labeling can either be done with the **label** command in the **console** program or using the **btape** program. The preferred method is to use the **label** command in the **console**

program.

Finally, you must add Volume names (and their attributes) to the Pool. For Volumes to be used by Bacula they must be of the same **Media Type** as the archive device specified for the job (i.e. if you are going to back up to a DLT device, the Pool must have DLT volumes defined since 8mm volumes cannot be mounted on a DLT drive). The **Media Type** has particular importance if you are backing up to files. When running a Job, you must explicitly specify which Pool to use. Bacula will then automatically select the next Volume to use from the Pool, but it will ensure that the **Media Type** of any Volume selected from the Pool is identical to that required by the Storage resource you have specified for the Job.

If you use the **label** command in the console program to label the Volumes, they will automatically be added to the Pool, so this last step is not normally required.

It is also possible to add Volumes to the database without explicitly labeling the physical volume. This is done with the **add** console command.

As previously mentioned, each time Bacula starts, it scans all the Pools associated with each Catalog, and if the database record does not already exist, it will be created from the Pool Resource definition. **Bacula** probably should do an **update pool** if you change the Pool definition, but currently, you must do this manually using the **update pool** command in the Console program.

The Pool Resource defined in the Director's configuration file (bacula-dir.conf) may contain the following directives:

Pool Start of the Pool resource. There must be at least one Pool resource defined.

Name = <name> The name of the pool. For most applications, you will use the default pool name **Default**. This directive is required.

Number of Volumes = <number> This directive specifies the number of volumes (tapes or files) contained in the pool. Normally, it is defined and updated automatically by the Bacula catalog handling routines.

Maximum Volumes = <number> This directive specifies the maximum number of volumes (tapes or files) contained in the pool. This directive is optional, if omitted or set to zero, any number of volumes will be permitted. In general, this directive is useful for Autochangers where there is a fixed number of Volumes, or for File

storage where you wish to ensure that the backups made to disk files do not become too numerous or consume too much space.

Pool Type = <type> This directive defines the pool type, which corresponds to the type of Job being run. It is required and may be one of the following:

Backup

*Archive

*Cloned

*Migration

*Copy

*Save

Use Volume Once = <yes—no> This directive if set to **yes** specifies that each volume is to be used only once. This is most useful when the Media is a file and you want a new file for each backup that is done. The default is **no** (i.e. use volume any number of times). This directive will most likely be phased out (deprecated), so you are recommended to use **Maximum Volume Jobs = 1** instead.

Please note that the value defined by this directive in the bacula-dir.conf file is the default value used when a Volume is created. Once the volume is created, changing the value in the bacula-dir.conf file will not change what is stored for the Volume. To change the value for an existing Volume you must use the **update** command in the Console.

Maximum Volume Jobs = <positive-integer> This directive specifies the maximum number of Jobs that can be written to the Volume. If you specify zero (the default), there is no limit. Otherwise, when the number of Jobs backed up to the Volume equals **positive-integer** the Volume will be marked **Used**. When the Volume is marked **Used** it can no longer be used for appending Jobs, much like the **Full** status but it can be recycled if recycling is enabled. By setting **MaximumVolumeJobs** to one, you get the same effect as setting **UseVolumeOnce = yes**.

Please note that the value defined by this directive in the bacula-dir.conf file is the default value used when a Volume is created. Once the volume is created, changing the value in the bacula-dir.conf file will not change what is stored for the Volume. To change the value for an existing Volume you must use the **update** command in the Console.

Maximum Volume Files = <positive-integer> This directive specifies the maximum number of files that can be written to

the Volume. If you specify zero (the default), there is no limit. Otherwise, when the number of files written to the Volume equals **positive-integer** the Volume will be marked **Used**. When the Volume is marked **Used** it can no longer be used for appending Jobs, much like the **Full** status but it can be recycled if recycling is enabled. This value is checked and the **Used** status is set only at the end of a job that writes to the particular volume.

Please note that the value defined by this directive in the bacula-dir.conf file is the default value used when a Volume is created. Once the volume is created, changing the value in the bacula-dir.conf file will not change what is stored for the Volume. To change the value for an existing Volume you must use the **update** command in the Console.

Maximum Volume Bytes = <size> This directive specifies the maximum number of bytes that can be written to the Volume. If you specify zero (the default), there is no limit except the physical size of the Volume. Otherwise, when the number of bytes written to the Volume equals **size** the Volume will be marked **Used**. When the Volume is marked **Used** it can no longer be used for appending Jobs, much like the **Full** status but it can be recycled if recycling is enabled. This value is checked and the **Used** status set while the job is writing to the particular volume.

Please note that the value defined by this directive in the bacula-dir.conf file is the default value used when a Volume is created. Once the volume is created, changing the value in the bacula-dir.conf file will not change what is stored for the Volume. To change the value for an existing Volume you must use the **update** command in the Console.

Volume Use Duration = <time-period-specification> The Volume Use Duration directive defines the time period that the Volume can be written beginning from the time of first data write to the Volume. If the time-period specified is zero (the default), the Volume can be written indefinitely. Otherwise, when the time period from the first write to the volume (the first Job written) exceeds the time-period-specification, the Volume will be marked **Used**, which means that no more Jobs can be appended to the Volume, but it may be recycled if recycling is enabled.

You might use this directive, for example, if you have a Volume used for Incremental backups, and Volumes used for Weekly Full backups. Once the Full backup is done, you will want to use a different Incremental Volume. This can be accomplished by setting the Volume Use Duration for the Incremental Volume to six days. I.e. it will be used for the 6 days following a Full save, then a different Incremental volume will be used.

This value is checked and the **Used** status is set only at the end of a job that writes to the particular volume, which means that even though the use duration may have expired, the catalog entry will not be updated until the next job that uses this volume is run.

Please note that the value defined by this directive in the bacula-dir.conf file is the default value used when a Volume is created. Once the volume is created, changing the value in the bacula-dir.conf file will not change what is stored for the Volume. To change the value for an existing Volume you must use the **update** command in the Console.

Catalog Files = <yes—no> This directive defines whether or not you want the names of the files that were saved to be put into the catalog. The default is **yes**. The advantage of specifying **Catalog Files = No** is that you will have a significantly smaller Catalog database. The disadvantage is that you will not be able to produce a Catalog listing of the files backed up for each Job (this is often called Browsing). Also, without the File entries in the catalog, you will not be able to use the Console **restore** command nor any other command that references File entries.

AutoPrune = <yes—no> If AutoPrune is set to **yes** (default), Bacula (version 1.20 or greater) will automatically apply the Volume Retention period when new Volume is needed and no appendable Volumes exist in the Pool. Volume pruning causes expired Jobs (older than the **Volume Retention** period) to be deleted from the Catalog and permits possible recycling of the Volume.

Volume Retention = <time-period-specification> The Volume Retention directive defines the length of time that **Bacula** will keep Job records associated with the Volume in the Catalog database. When this time period expires, and if **AutoPrune** is set to **yes** Bacula will prune (remove) Job records that are older than the specified Volume Retention period. All File records associated with pruned Jobs are also pruned. The time may be specified as seconds, minutes, hours, days, weeks, months, quarters, or years. The **Volume Retention** applied independently to the **Job Retention** and the **File Retention** periods defined in the Client resource. This means that the shorter period is the one that applies. Note, that when the **Volume Retention** period has been reached, it will prune both the Job and the File records.

The default is 365 days. Note, this directive sets the default value for each Volume entry in the Catalog when the Volume is created. The value in the catalog may be later individually changed for each Volume using the Console program.

By defining multiple Pools with different Volume Retention periods, you may effectively have a set of tapes that is recycled weekly, another Pool of tapes that is recycled monthly and so on. However, one must keep in mind that if your **Volume Retention** period is too short, it may prune the last valid Full backup, and hence until the next Full backup is done, you will not have a complete backup of your system, and in addition, the next Incremental or Differential backup will be promoted to a Full backup. As a consequence, the minimum **Volume Retention** period should be at twice the interval of your Full backups. This means that if you do a Full backup once a month, the minimum Volume retention period should be two months.

Please note that the value defined by this directive in the bacula-dir.conf file is the default value used when a Volume is created. Once the volume is created, changing the value in the bacula-dir.conf file will not change what is stored for the Volume. To change the value for an existing Volume you must use the **update** command in the Console.

Recycle = <yes—no> This directive specifies the default for recycling Purged Volumes. If it is set to **yes** and Bacula needs a volume but finds none that are appendable, it will search for Purged Volumes (i.e. volumes with all the Jobs and Files expired and thus deleted from the Catalog). If the Volume is recycled, all previous data written to that Volume will be overwritten.

Please note that the value defined by this directive in the bacula-dir.conf file is the default value used when a Volume is created. Once the volume is created, changing the value in the bacula-dir.conf file will not change what is stored for the Volume. To change the value for an existing Volume you must use the **update** command in the Console.

Recycle Oldest Volume = <yes—no> This directive instructs the Director to search for the oldest used Volume in the Pool when another Volume is requested by the Storage daemon and none are available. The catalog is then **pruned** respecting the retention periods of all Files and Jobs written to this Volume. If all Jobs are pruned (i.e. the volume is Purged), then the Volume is recycled and will be used as the next Volume to be written. This directive respects any Job, File, or Volume retention periods that you may have specified, and as such it is **much** better to use this directive than the Purge Oldest Volume.

This directive can be useful if you have a fixed number of Volumes in the Pool and you want to cycle through them and you have specified the correct retention periods. However, if you use this directive and have only one Volume in the Pool, you will immediately recycle your Volume if you fill it and Bacula needs another one. Thus your backup will be totally invalid. Please use this directive with care.

Recycle Current Volume = <yes—no> If Bacula needs a new Volume, this directive instructs Bacula to Prune the volume respecting the Job and File retention periods. If all Jobs are pruned (i.e. the volume is Purged), then the Volume is recycled and will be used as the next Volume to be written. This directive respects any Job, File, or Volume retention periods that you may have specified, and thus it is **much** better to use it rather than the Purge Oldest Volume directive.

This directive can be useful if you have: a fixed number of Volumes in the Pool, you want to cycle through them, and you have specified retention periods that prune Volumes before you have cycled through the Volume in the Pool. However, if you use this directive and have only one Volume in the Pool, you will immediately recycle your Volume if you fill it and Bacula needs another one. Thus your backup will be totally invalid. Please use this directive with care.

Purge Oldest Volume = <yes—no> This directive instructs the Director to search for the oldest used Volume in the Pool when another Volume is requested by the Storage daemon and none are available. The catalog is then **purged** irrespective of retention periods of all Files and Jobs written to this Volume. The Volume is then recycled and will be used as the next Volume to be written. This directive overrides any Job, File, or Volume retention periods that you may have specified.

This directive can be useful if you have a fixed number of Volumes in the Pool and you want to cycle through them and when all Volumes are full, but you don't want to worry about setting proper retention periods. However, by using this option you risk losing valuable data.

Please be aware that Purge Oldest Volume disregards all retention periods. If you have only a single Volume defined and you turn this variable on, that Volume will always be immediately overwritten when it fills! So at a minimum, ensure that you have a decent number of Volumes in your Pool before running any jobs. If you want retention periods to apply do not use this directive. To specify a retention period, use the **Volume Retention** directive (see above).

We **highly** recommend against using this directive, because it is sure that some day, Bacula will recycle a Volume that contains current data.

Accept Any Volume = <yes—no> This directive specifies whether or not any volume from the Pool may be used for backup. The default is **yes** as of version 1.27 and later. If it is **no** then only the first writable volume in the Pool will be accepted for writing backup data,

thus Bacula will fill each Volume sequentially in turn before using any other appendable volume in the Pool. If this is **no** and you mount a volume out of order, Bacula will not accept it. If this is **yes** any appendable volume from the pool mounted will be accepted.

If your tape backup procedure dictates that you manually mount the next volume, you will almost certainly want to be sure this directive is turned on.

If you are going on vacation and you think the current volume may not have enough room on it, you can simply label a new tape and leave it in the drive, and assuming that **Accept Any Volume** is **yes** Bacula will begin writing on it. When you return from vacation, simply remount the last tape, and Bacula will continue writing on it until it is full. Then you can remount your vacation tape and Bacula will fill it in turn.

Cleaning Prefix = <string> This directive defines a prefix string, which if it matches the beginning of a Volume name during labeling of a Volume, the Volume will be defined with the VolStatus set to **Cleaning** and thus Bacula will never attempt to use this tape. This is primarily for use with autochangers that accept barcodes where the convention is that barcodes beginning with **CLN** are treated as cleaning tapes.

Label Format = <format> This directive specifies the format of the labels contained in this pool. The format directive is used as a sort of template to create new Volume names during automatic Volume labeling.

The **format** should be specified in double quotes, and consists of letters, numbers and the special characters hyphen (**-**), underscore (**_**), colon (**:**), and period (**.**), which are the legal characters for a Volume name. The **format** should be enclosed in double quotes (").

In addition, the format may contain a number of variable expansion characters which will be expanded by a complex algorithm allowing you to create Volume names of many different formats. In all cases, the expansion process must resolve to the set of characters noted above that are legal Volume names. Generally, these variable expansion characters begin with a dollar sign (**$**) or a left bracket ([). If you specify variable expansion characters, you should always enclose the format with double quote characters (“). For more details on variable expansion, please see the Variable Expansion Chapter of this manual.

If no variable expansion characters are found in the string, the Volume name will be formed from the **format** string appended with the number of volumes in the pool plus one, which will be edited as four

digits with leading zeros. For example, with a **Label Format = "File-"**, the first volumes will be named **File-0001**, **File-0002**, ...

With the exception of Job specific variables, you can test your **LabelFormat** by using the var command the Console Chapter of this manual.

In almost all cases, you should enclose the format specification (part after the equal sign) in double quotes. Please note that this directive is deprecated and is replaced in version 1.37 and greater with a Python script for creating volume names.

In order for a Pool to be used during a Backup Job, the Pool must have at least one Volume associated with it. Volumes are created for a Pool using the **label** or the **add** commands in the **Bacula Console**, program. In addition to adding Volumes to the Pool (i.e. putting the Volume names in the Catalog database), the physical Volume must be labeled with valid Bacula software volume label before **Bacula** will accept the Volume. This will be automatically done if you use the **label** command. Bacula can automatically label Volumes if instructed to do so, but this feature is not yet fully implemented.

The following is an example of a valid Pool resource definition:

```
Pool {
  Name = Default
  Pool Type = Backup
}
```

The Catalog Resource

The Catalog Resource defines what catalog to use for the current job. Currently, Bacula can only handle a single database server (SQLite, MySQL, built-in) that is defined when configuring **Bacula**. However, there may be as many Catalogs (databases) defined as you wish. For example, you may want each Client to have its own Catalog database, or you may want backup jobs to use one database and verify or restore jobs to use another database.

Catalog Start of the Catalog resource. At least one Catalog resource must be defined.

Name = <name> The name of the Catalog. No necessary relation to the database server name. This name will be specified in the Client

resource directive indicating that all catalog data for that Client is maintained in this Catalog. This directive is required.

password = <password> This specifies the password to use when logging into the database. This directive is required.

DB Name = <name> This specifies the name of the database. If you use multiple catalogs (databases), you specify which one here. If you are using an external database server rather than the internal one, you must specify a name that is known to the server (i.e. you explicitly created the Bacula tables using this name. This directive is required.

user = <user> This specifies what user name to use to log into the database. This directive is required.

DB Socket = <socket-name> This is the name of a socket to use on the local host to connect to the database. This directive is used only by MySQL and is ignored by SQLite. Normally, if neither **DB Socket** or **DB Address** are specified, MySQL will use the default socket.

DB Address = <address> This is the host address of the database server. Normally, you would specify this instead of **DB Socket** if the database server is on another machine. In that case, you will also specify **DB Port**. This directive is used only by MySQL and is ignored by SQLite if provided. This directive is optional.

DB Port = <port> This defines the port to be used in conjunction with **DB Address** to access the database if it is on another machine. This directive is used only by MySQL and is ignored by SQLite if provided. This directive is optional.

The following is an example of a valid Catalog resource definition:

```
Catalog
{
  Name = SQLite
  dbname = bacula;
  user = bacula;
  password = ""                       # no password = no security
}
```

or for a Catalog on another machine:

```
Catalog
{
```

```
  Name = MySQL
  dbname = bacula
  user = bacula
  password = ""
  DB Address = remote.acme.com
  DB Port = 1234
}
```

The Messages Resource

For the details of the Messages Resource, please see the Messages Resource Chapter of this manual.

The Console Resource

As of Bacula version 1.33 and higher, there are three different kinds of consoles, which the administrator or user can use to interact with the Director. These three kinds of consoles comprise three different security levels.

- The first console type is an **anonymous** or **default** console, which has full privileges. There is no console resource necessary for this type since the password is specified in the Director's resource and consequently such consoles do not have an name as defined on a **Name =** directive. This is the kind of console that was initially implemented in versions prior to 1.33 and remains valid. Typically you would use it only for administrators.
- The second type of console, and new to version 1.33 and higher is a "named" console defined within a Console resource in both the Director's configuration file and in the Console's configuration file. Both the names and the passwords in these two entries must match much as is the case for Client programs.

  This second type of console begins with absolutely no privileges except those explicitly specified in the Director's Console resource. Thus you can have multiple Consoles with different names and passwords, sort of like multiple users, each with different privileges. As a default, these consoles can do absolutely nothing – no commands what so ever. You give them privileges or rather access to commands and resources by specifying access control lists in the Director's Console resource. The ACLs are specified by a directive followed by a list of access names. Examples of this are shown below.

- The third type of console is similar to the above mentioned one in that it requires a Console resource definition in both the Director and the Console. In addition, if the console name, provided on the **Name =** directive, is the same as a Client name, that console is permitted to use the **SetIP** command to change the Address directive in the Director's client resource to the IP address of the Console. This permits portables or other machines using DHCP (non-fixed IP addresses) to "notify" the Director of their current IP address.

The Console resource is optional and need not be specified. The following directives are permited within the Director's configuration resource:

Name = <name> The name of the console. This name must match the name specified in the Console's configuration resource (much as is the case with Client definitions).

Password = <password> Specifies the password that must be supplied for a named Bacula Console to be authorized. The same password must appear in the **Console** resource of the Console configuration file. For added security, the password is never actually passed across the network but rather a challenge response hash code created with the password. This directive is required. If you have either **/dev/random** **bc** on your machine, Bacula will generate a random password during the configuration process, otherwise it will be left blank.

JobACL = <name-list> This directive is used to specify a list of Job resource names that can be accessed by the console. Without this directive, the console cannot access any of the Director's Job resources. Multiple Job resource names may be specified by separating them with commas, and/or by specifying multiple JobACL directives. For example, the directive may be specified as:

```
JobACL = kernsave, "Backup client 1", "Backup client 2"
JobACL = "RestoreFiles"
```

With the above specification, the console can access the Director's resources for the four jobs named on the JobACL directives, but for no others.

ClientACL = <name-list> This directive is used to specify a list of Client resource names that can be accessed by the console.

StorageACL = <name-list> This directive is used to specify a list of Storage resource names that can be accessed by the console.

ScheduleACL = <name-list> This directive is used to specify a list of Schedule resource names that can be accessed by the console.

PoolACL = <name-list> This directive is used to specify a list of Pool resource names that can be accessed by the console.

FileSetACL = <name-list> This directive is used to specify a list of FileSet resource names that can be accessed by the console.

CatalogACL = <name-list> This directive is used to specify a list of Catalog resource names that can be accessed by the console.

CommandACL = <name-list> This directive is used to specify a list of of console commands that can be executed by the console.

Aside from Director resource names and console command names, the special keyword ***all*** can be specified in any of the above access control lists. When this keyword is present, any resource or command name (which ever is appropriate) will be accepted. For an example configuration file, please see the Console Configuration chapter of this manual.

The Counter Resource

The Counter Resource defines a counter variable that can be accessed by variable expansion used for creating Volume labels with the **LabelFormat** directive. See the LabelFormat directive in this chapter for more details.

Counter Start of the Counter resource. Counter directives are optional.

Name = <name> The name of the Counter. This is the name you will use in the variable expansion to reference the counter value.

Minimum = <integer> This specifies the minimum value that the counter can have. It also becomes the default. If not supplied, zero is assumed.

Maximum = <integer> This is the maximum value value that the counter can have. If not specified or set to zero, the counter can have a maximum value of 2,147,483,648 (2 to the 31 power). When the counter is incremented past this value, it is reset to the Minimum.

***WrapCounter = <counter-name>** If this value is specified, when the counter is incremented past the maximum and thus reset to the minimum, the counter specified on the **WrapCounter** is incremented. (This is not currently implemented).

Catalog = <catalog-name> If this directive is specified, the counter and its values will be saved in the specified catalog. If this directive is not present, the counter will be redefined each time that Bacula is started.

Example Director Configuration File

An example Director configuration file might be the following:

```
#
# Default Bacula Director Configuration file
#
#  The only thing that MUST be changed is to add one or more
#   file or directory names in the Include directive of the
#   FileSet resource.
#
#  For Bacula release 1.15 (5 March 2002) -- redhat
#
#  You might also want to change the default email address
#   from root to your address.  See the "mail" and "operator"
#   directives in the Messages resource.
#
Director {                            # define myself
  Name = rufus-dir
  QueryFile = "/home/kern/bacula/bin/query.sql"
  WorkingDirectory = "/home/kern/bacula/bin/working"
  PidDirectory = "/home/kern/bacula/bin/working"
  Password = "XkSfzu/Cf/wX4L8Zh4G4/yhCbpLcz3YVdmVoQvU3EyF/"
}
# Define the backup Job
Job {
  Name = "NightlySave"
  Type = Backup
  Level = Incremental                 # default
  Client=rufus-fd
  FileSet="Full Set"
  Schedule = "WeeklyCycle"
  Storage = DLTDrive
  Messages = Standard
  Pool = Default
}
Job {
  Name = "Restore"
  Type = Restore
  Client=rufus-fd
  FileSet="Full Set"
  Where = /tmp/bacula-restores
  Storage = DLTDrive
  Messages = Standard
  Pool = Default
}
```

```
# List of files to be backed up
FileSet {
  Name = "Full Set"
  Include {
    Options { signature=SHA1 }
#
#  Put your list of files here, one per line or include an
#    external list with:
#
#    @file-name
#
#  Note: / backs up everything
  File = /
  }
  Exclude { }
}
# When to do the backups
Schedule {
  Name = "WeeklyCycle"
  Run = Full sun at 1:05
  Run = Incremental mon-sat at 1:05
}
# Client (File Services) to backup
Client {
  Name = rufus-fd
  Address = rufus
  Catalog = MyCatalog
  Password = "MQk6lVinz4GG2hdIZk1dsKE/LxMZGo6znMHiD7t7vzF+"
  File Retention = 60d      # sixty day file retention
  Job Retention = 1y        # 1 year Job retention
  AutoPrune = yes           # Auto apply retention periods
}
# Definition of DLT tape storage device
Storage {
  Name = DLTDrive
  Address = rufus
  Password = "jMeWZvfikUHvt3kzKVVPpQ0ccmV6emPnF2cPYFdhLApQ"
  Device = "HP DLT 80"      # same as Device in Storage daemon
  Media Type = DLT8000      # same as MediaType in Storage daemon
}
# Definition of DDS tape storage device
Storage {
  Name = SDT-10000
  Address = rufus
  Password = "jMeWZvfikUHvt3kzKVVPpQ0ccmV6emPnF2cPYFdhLApQ"
  Device = SDT-10000        # same as Device in Storage daemon
  Media Type = DDS-4        # same as MediaType in Storage daemon
}
# Definition of 8mm tape storage device
Storage {
  Name = "8mmDrive"
  Address = rufus
  Password = "jMeWZvfikUHvt3kzKVVPpQ0ccmV6emPnF2cPYFdhLApQ"
```

```
  Device = "Exabyte 8mm"
  MediaType = "8mm"
}
# Definition of file storage device
Storage {
  Name = File
  Address = rufus
  Password = "jMeWZvfikUHvt3kzKVVPpQ0ccmV6emPnF2cPYFdhLApQ"
  Device = FileStorage
  Media Type = File
}
# Generic catalog service
Catalog {
  Name = MyCatalog
  dbname = bacula; user = bacula; password = ""
}
# Reasonable message delivery -- send most everything to
#   the email address and to the console
Messages {
  Name = Standard
  mail = root@localhost = all, !skipped, !terminate
  operator = root@localhost = mount
  console = all, !skipped, !saved
}

# Default pool definition
Pool {
  Name = Default
  Pool Type = Backup
  AutoPrune = yes
  Recycle = yes
}
#
# Restricted console used by tray-monitor to get the status of the director
#
Console {
  Name = Monitor
  Password = "GN0uRo7PTUmlMbqrJ2Gr1p0fk0HQJTxwnFyE4WSST3MWZseR"
  CommandACL = status, .status
}
```

Client/File daemon Configuration

General

The Client (or File Daemon) Configuration is one of the simpler ones to specify. Generally, other than changing the Client name so that error messages are easily identified, you will not need to modify the default Client configuration file.

For a general discussion of configuration file and resources including the data types recognized by **Bacula**, please see the Configuration chapter of this manual. The following Client Resource definitions must be defined:

- Client – to define what Clients are to be backed up.
- Director – to define the Director's name and its access password.
- Messages – to define where error and information messages are to be sent.

The Client Resource

The Client Resource (or FileDaemon) resource defines the name of the Client (as used by the Director) as well as the port on which the Client listens for Director connections.

Client (or FileDaemon) Start of the Client records. There must be one and only one Client resource in the configuration file, since it defines the properties of the current client program.

Name = <name> The client name that must be used by the Director when connecting. Generally, it is a good idea to use a name related to the machine so that error messages can be easily identified if you have multiple Clients. This record is required.

Working Directory = <Directory> This directive is mandatory and specifies a directory in which the File daemon may put its status files. This directory should be used only by **Bacula**, but may be shared by other Bacula daemons. This record is required

Pid Directory = <Directory> This directive is mandatory and specifies a directory in which the Director may put its process Id file files. The process Id file is used to shutdown Bacula and to prevent

multiple copies of Bacula from running simultaneously. This record is required. Standard shell expansion of the **Directory** is done when the configuration file is read so that values such as **$HOME** will be properly expanded.

Typically on Linux systems, you will set this to: **/var/run**. If you are not installing Bacula in the system directories, you can use the **Working Directory** as defined above.

Heartbeat Interval = <time-interval> This record defines an interval of time. For each heartbeat that the File daemon receives from the Storage daemon, it will forward it to the Director. In addition, if no heartbeat has been received from the Storage daemon and thus forwarded the File daemon will send a heartbeat signal to the Director and to the Storage daemon to keep the channels active. The default interval is zero which disables the heartbeat. This feature is particularly useful if you have a router such as 3Com that does not follow Internet standards and times out an inactive connection after a short duration.

Maximum Concurrent Jobs = <number> where <number> is the maximum number of Jobs that should run concurrently. The default is set to 2, but you may set it to a larger number. Each contact from the Director (e.g. status request, job start request) is considered as a Job, so if you want to be able to do a **status** request in the console at the same time as a Job is running, you will need to set this value greater than 1.

FDAddresses = <IP-address-specification> Specify the ports and addresses on which the Director daemon will listen for Bacula Console connections. Probably the simplest way to explain is to show an example:

```
FDAddresses  = { ip = {
        addr = 1.2.3.4; port = 1205; }
    ipv4 = {
        addr = 1.2.3.4; port = http; }
    ipv6 = {
        addr = 1.2.3.4;
        port = 1205;
    }
    ip = {
        addr = 1.2.3.4
        port = 1205
    }
    ip = {
        addr = 1.2.3.4
    }
```

```
    ip = {
        addr = 201:220:222::2
    }
    ip = {
        addr = bluedot.thun.net
    }
}
```

where ip, ip4, ip6, addr, and port are all keywords. Note, that the address can be specified as either a dotted quadruple, or IPv6 colon notation, or as a symbolic name (only in the ip specification). Also, port can be specified as a number or as the mnemonic value from the /etc/services file. If a port is not specified, the default will be used. If an ip section is specified, the resolution can be made either by IPv4 or IPv6. If ip4 is specified, then only IPv4 resolutions will be permitted, and likewise with ip6.

FDPort = <port-number> This specifies the port number on which the Client listens for Director connections. It must agree with the FDPort specified in the Client resource of the Director's configuration file. The default is 9102.

FDAddress = <IP-Address> This record is optional, and if it is specified, it will cause the File daemon server (for Director connections) to bind to the specified **IP-Address**, which is either a domain name or an IP address specified as a dotted quadruple. If this record is not specified, the File daemon will bind to any available address (the default).

SDConnectTimeout = <time-interval> This record defines an interval of time that the File daemon will try to connect to the Storage daemon. The default is 30 minutes. If no connection is made in the specified time interval, the File daemon cancels the Job.

Maximum Network Buffer Size = <bytes> where <bytes> specifies the initial network buffer size to use with the File daemon. This size will be adjusted down if it is too large until it is accepted by the OS. Please use care in setting this value since if it is too large, it will be trimmed by 512 bytes until the OS is happy, which may require a large number of system calls. The default value is 32,768 bytes.

The following is an example of a valid Client resource definition:

```
Client {                              # this is me
  Name = rufus-fd
  WorkingDirectory = $HOME/bacula/bin/working
```

```
  Pid Directory = $HOME/bacula/bin/working
}
```

The Director Resource

The Director resource defines the name and password of the Directors that are permitted to contact this Client.

Director Start of the Director records. There may be any number of Director resources in the Client configuration file. Each one specifies a Director that is allowed to connect to this Client.

Name = <name> The name of the Director that may contact this Client. This name must be the same as the name specified on the Director resource in the Director's configuration file. This record is required.

Password = <password> Specifies the password that must be supplied for a Director to be authorized. This password must be the same as the password specified in the Client resource in the Director's configuration file. This record is required.

Monitor = <yes—no> If Monitor is set to **no** (default), this director will have full access to this Client. If Monitor is set to **yes**, this director will only be able to fetch the current status of this Client.

Please note that if this director is being used by a Monitor, we highly recommend to set this directive to **yes** to avoid serious security problems.

Thus multiple Directors may be authorized to use this Client's services. Each Director will have a different name, and normally a different password as well.

The following is an example of a valid Director resource definition:

```
#
# List Directors who are permitted to contact the File daemon
#
Director {
  Name = HeadMan
  Password = very_good                # password HeadMan must supply
}
Director {
  Name = Worker
  Password = not_as_good
  Monitor = Yes
}
```

The Message Resource

Please see the Messages Resource Chapter of this manual for the details of the Messages Resource.

There must be at least one Message resource in the Client configuration file.

Example Client Configuration File

An example File Daemon configuration file might be the following:

```
#
# Default  Bacula File Daemon Configuration file
#
#  For Bacula release 1.35.2 (16 August 2004) -- gentoo 1.4.16
#
# There is not much to change here except perhaps to
#   set the Director's name and File daemon's name
#   to something more appropriate for your site.
#
#
# List Directors who are permitted to contact this File daemon
#
Director {
  Name = rufus-dir
  Password = "/LqPRkX++saVyQE7w7mmiFg/qxYc1kufww6FEyY/47jU"
}
#
# Restricted Director, used by tray-monitor to get the
#   status of the file daemon
#
Director {
  Name = rufus-mon
  Password = "FYpq4yyI1y562EMS35bAOJOQCOM2L3t5cZObxT3XQxgxppTn"
  Monitor = yes
}
#
# "Global" File daemon configuration specifications
#
FileDaemon {                          # this is me
  Name = rufus-fd
  WorkingDirectory = $HOME/bacula/bin/working
  Pid Directory = $HOME/bacula/bin/working
}
# Send all messages except skipped files back to Director
Messages {
  Name = Standard
  director = rufus-dir = all, !skipped
}
```

Storage Daemon Configuration

General

The Storage Daemon configuration file has relatively few resource definitions. However, due to the great variation in backup media and system capabilities, the storage daemon must be highly configurable. As a consequence, there are quite a large number of directives in the Device Resource definition that allow you to define all the characteristics of your Storage device (normally a tape drive). Fortunately, with modern storage devices, the defaults are sufficient, and very few directives are actually needed.

Examples of **Device** resource directives that are known to work for a number of common tape drives can be found in the <**bacula-src**>**/examples/devices** directory, and most will also be listed here.

For a general discussion of configuration file and resources including the data types recognized by **Bacula**, please see the Configuration chapter of this manual. The following Storage Resource definitions must be defined:

- Storage – to define the name of the Storage daemon.
- Director – to define the Director's name and his access password.
- Device – to define the characteristics of your storage device (tape drive).
- Messages – to define where error and information messages are to be sent.

Storage Resource

In general, the properties specified under the Storage resource define global properties of the Storage daemon. Each Storage daemon configuration file must have one and only one Storage resource definition.

Name = <**Storage-Daemon-Name**> Specifies the Name of the Storage daemon. This directive is required.

Working Directory = <**Directory**> This directive is mandatory and specifies a directory in which the Storage daemon may put its status

files. This directory should be used only by **Bacula**, but may be shared by other Bacula daemons. This directive is required

Pid Directory = <Directory> This directive is mandatory and specifies a directory in which the Director may put its process Id file files. The process Id file is used to shutdown Bacula and to prevent multiple copies of Bacula from running simultaneously. This directive is required. Standard shell expansion of the **Directory** is done when the configuration file is read so that values such as **$HOME** will be properly expanded.

Typically on Linux systems, you will set this to: **/var/run**. If you are not installing Bacula in the system directories, you can use the **Working Directory** as defined above.

Heartbeat Interval = <time-interval> This directive defines an interval of time. When the Storage daemon is waiting for the operator to mount a tape, each time interval, it will send a heartbeat signal to the File daemon. The default interval is zero which disables the heartbeat. This feature is particularly useful if you have a router such as 3Com that does not follow Internet standards and times out an inactive connection after a short duration.

Maximum Concurrent Jobs = <number> where <number> is the maximum number of Jobs that should run concurrently. The default is set to 10, but you may set it to a larger number. Each contact from the Director (e.g. status request, job start request) is considered as a Job, so if you want to be able to do a **status** request in the console at the same time as a Job is running, you will need to set this value greater than 1. To run simultaneous Jobs, you will need to set a number of other directives in the Director's configuration file. Which ones you set depend on what you want, but you will almost certainly need to set the **Maximum Concurrent Jobs** in the Storage resource in the Director's configuration file and possibly those in the Job and Client resources.

SDAddresses = <IP-address-specification> Specify the ports and addresses on which the Storage daemon will listen for Director connections. Normally, the default is sufficient and you do not need to specify this directive. Probably the simplest way to explain how this directive works is to show an example:

```
SDAddresses  = { ip = {
       addr = 1.2.3.4; port = 1205; }
   ipv4 = {
       addr = 1.2.3.4; port = http; }
   ipv6 = {
```

```
            addr = 1.2.3.4;
            port = 1205;
        }
        ip = {
            addr = 1.2.3.4
            port = 1205
        }
        ip = {
            addr = 1.2.3.4
        }
        ip = {
            addr = 201:220:222::2
        }
        ip = {
            addr = bluedot.thun.net
        }
    }
```

where ip, ip4, ip6, addr, and port are all keywords. Note, that the address can be specified as either a dotted quadruple, or IPv6 colon notation, or as a symbolic name (only in the ip specification). Also, port can be specified as a number or as the mnemonic value from the /etc/services file. If a port is not specified, the default will be used. If an ip section is specified, the resolution can be made either by IPv4 or IPv6. If ip4 is specified, then only IPv4 resolutions will be permitted, and likewise with ip6.

Using this directive, you can replace both the SDPort and SDAddress directives shown below.

SDPort = <port-number> Specifies port number on which the Storage daemon listens for Director connections. The default is 9103.

SDAddress = <IP-Address> This directive is optional, and if it is specified, it will cause the Storage daemon server (for Director and File daemon connections) to bind to the specified **IP-Address**, which is either a domain name or an IP address specified as a dotted quadruple. If this directive is not specified, the Storage daemon will bind to any available address (the default).

The following is a typical Storage daemon Storage definition.

```
#
# "Global" Storage daemon configuration specifications appear
# under the Storage resource.
#
Storage {
  Name = "Storage daemon"
```

```
  Address = localhost
  WorkingDirectory = "~/bacula/working"
  Pid    Directory = "~/bacula/working"
}
```

Director Resource

The Director resource specifies the Name of the Director which is permitted to use the services of the Storage daemon. There may be multiple Director resources. The Director Name and Password must match the corresponding values in the Director's configuration file.

Name = <Director-Name> Specifies the Name of the Director allowed to connect to the Storage daemon. This directive is required.

Password = <Director-password> Specifies the password that must be supplied by the above named Director. This directive is required.

Monitor = <yes—no> If Monitor is set to **no** (default), this director will have full access to this Storage daemon. If Monitor is set to **yes**, this director will only be able to fetch the current status of this Storage daemon.

Please note that if this director is being used by a Monitor, we highly recommend to set this directive to **yes** to avoid serious security problems.

The following is an example of a valid Director resource definition:

```
Director {
  Name = MainDirector
  Password = my_secret_password
}
```

Device Resource

The Device Resource specifies the details of each device (normally a tape drive) that can be used by the Storage daemon. There may be multiple Device resources for a single Storage daemon. In general, the properties specified within the Device resource are specific to the Device.

Name = *Device-Name* Specifies the Name that the Director will use when asking to backup or restore to or from to this device. This is the logical Device name, and may be any string up to 127 characters in length. It is generally a good idea to make it correspond to the English name of the backup device. The physical name of the device is specified on the **Archive Device** directive described below. The name you specify here is also used in your Director's conf file on the Device directive in its Storage resource.

Archive Device = *name-string* The specified **name-string** gives the system file name of the storage device managed by this storage daemon. This will usually be the device file name of a removable storage device (tape drive), for example "**/dev/nst0**" or "**/dev/rmt/0mbn**". For a DVD-writer, it will be for example **/dev/hdc**. It may also be a directory name if you are archiving to disk storage. In this case, you must supply the full absolute path to the directory. When specifying a tape device, it is preferable that the "non-rewind" variant of the device file name be given. In addition, on systems such as Sun, which have multiple tape access methods, you must be sure to specify to use Berkeley I/O conventions with the device. The **b** in the Solaris (Sun) archive specification **/dev/rmt/0mbn** is what is needed in this case. Bacula does not support SysV tape drive behavior.

As noted above, normally the Archive Device is the name of a tape drive, but you may also specify an absolute path to an existing directory. If the Device is a directory Bacula will write to file storage in the specified directory, and the filename used will be the Volume name as specified in the Catalog. If you want to write into more than one directory (i.e. to spread the load to different disk drives), you will need to define two Device resources, each containing an Archive Device with a different directory.

In addition to a tape device name or a directory name, Bacula will accept the name of a FIFO. A FIFO is a special kind of file that connects two programs via kernel memory. If a FIFO device is specified for a backup operation, you must have a program that reads what Bacula writes into the FIFO. When the Storage daemon starts the job, it will wait for **MaximumOpenWait** seconds for the read program to start reading, and then time it out and terminate the job. As a consequence, it is best to start the read program at the beginning of the job perhaps with the **RunBeforeJob** directive. For this kind of device, you never want to specify **AlwaysOpen**, because you want the Storage daemon to open it only when a job starts, so you must explicitly set it to **No**. Since a FIFO is a one way device, Bacula will

not attempt to read a label of a FIFO device, but will simply write on it. To create a FIFO Volume in the catalog, use the **add** command rather than then **label** command to avoid attempting to write a label.

During a restore operation, if the Archive Device is a FIFO, Bacula will attempt to read from the FIFO, so you must have an external program that writes into the FIFO. Bacula will wait **MaximumOpenWait** seconds for the program to begin writing and will then time it out and terminate the job. As noted above, you may use the **RunBeforeJob** to start the writer program at the beginning of the job.

The Archive Device directive is required.

Media Type = *name-string* The specified **name-string** names the type of media supported by this device, for example, "DLT7000". Media type names are arbitrary in that you set it to anything you want, but must be known to the volume database to keep track of which storage daemons can read which volumes. The same **name-string** must appear in the appropriate Storage resource definition in the Director's configuration file.

Even though the names you assign are arbitrary (i.e. you choose the name you want), you should take care in specifying them because the Media Type is used to determine which storage device Bacula will select during restore. Thus you should probably use the same Media Type specification for all drives where the Media can be freely interchanged. This is not generally an issue if you have a single Storage daemon, but it is with multiple Storage daemons, especially if they have incompatible media.

For example, if you specify a Media Type of "DDS-4" then during the restore, Bacula will be able to choose any Storage Daemon that handles "DDS-4". If you have an autochanger, you might want to name the Media Type in a way that is unique to the autochanger, unless you wish to possibly use the Volumes in other drives. You should also ensure to have unique Media Type names if the Media is not compatible between drives. This specification is required for all devices.

Autochanger = *Yes—No* If **Yes**, this device is an automatic tape changer, and you should also specify a **Changer Device** as well as a **Changer Command**. If **No** (default), the volume must be manually changed. You might also want to add an identical directive to the Storage resource in the Director's configuration file so that when labeling tapes you are prompted for the slot.

Changer Device = *name-string* The specified **name-string** gives the system file name of the autochanger device name that corresponds

to the **Archive Device** specified. This device name is specified if you have an autochanger or if you want to use the **Alert Command** (see below). Normally you will specify the **generic SCSI** device name in this directive. For example, on Linux systems, for archive device **/dev/nst0**, This directive is optional. See the Using Autochangers chapter of this manual for more details of using this and the following autochanger directives.

Changer Command = *name-string* The **name-string** specifies an external program to be called that will automatically change volumes as required by **Bacula**. Most frequently, you will specify the Bacula supplied **mtx-changer** script as follows:

```
Changer Command = "/path/mtx-changer %c %o %S %a %d"
```

and you will install the **mtx** on your system (found in the **depkgs** release). An example of this command is in the default bacula-sd.conf file. For more details on the substitution characters that may be specified to configure your autochanger please see the Autochangers chapter of this manual. For FreeBSD users, you might want to see one of the several **chio** scripts in **examples/autochangers**.

Alert Command = *name-string* The **name-string** specifies an external program to be called at the completion of each Job after the device is released. The purpose of this command is to check for Tape Alerts, which are present when something is wrong with your tape drive (at least for most modern tape drives). The same substitution characters that may be specified in the Changer Command may also be used in this string. For more information, please see the Autochangers chapter of this manual.

Note, it is not necessary to have an autochanger to use this command. The example below uses the **tapeinfo** program that comes with the **mtx** package, but it can be used on any tape drive. However, you will need to specify a **Changer Device** directive in your Device resource (see above) so that the generic SCSI device name can be edited into the command (with the %c).

An example of the use of this command to print Tape Alerts in the Job report is:

```
Alert Command = "sh -c 'tapeinfo -f %c | grep TapeAlert'"
```

and an example output when there is a problem could be:

```
bacula-sd  Alert: TapeAlert[32]: Interface: Problem with SCSI interface
                  between tape drive and initiator.
```

Drive Index = *number* The **Drive Index** that you specify is passed to the **mtx-changer** script and is thus passed to the **mtx** program. By default, the Drive Index is zero, so if you have only one drive in your autochanger, everything will work normally. However, if you have multiple drives, you may specify two Bacula Device resources. The first will either set Drive Index to zero, or leave it unspecified, and the second Device Resource should contain a Drive Index set to 1. This will then permit you to use two or more drives in your autochanger. However, you must ensure that Bacula does not request the same Volume on both drives at the same time. You may also need to modify the mtx-changer script to do locking so that two jobs don't attempt to use the autochanger at the same time. An example script can be found in **examples/autochangers/locking-mtx-changer**.

Maximum Changer Wait = *time* This directive specifies the maximum time for Bacula to wait for an autochanger to change the volume. If this time is exceeded, Bacula will invalidate the Volume slot number stored in the catalog and try again. If no additional changer volumes exist, Bacula will ask the operator to intervene. The default time out is 5 minutes.

Always Open = *Yes—No* If **Yes** (default), Bacula will always keep the device open unless specifically **unmounted** by the Console program. This permits Bacula to ensure that the tape drive is always available. If you set **AlwaysOpen** to **no Bacula** will only open the drive when necessary, and at the end of the Job if no other Jobs are using the drive, it will be freed. To minimize unnecessary operator intervention, it is highly recommended that **Always Open = yes**. This also ensures that the drive is available when Bacula needs it.

If you have **Always Open = yes** (recommended) and you want to use the drive for something else, simply use the **unmount** command in the Console program to release the drive. However, don't forget to remount the drive with **mount** when the drive is available or the next Bacula job will block.

For File storage, this directive is ignored. For a FIFO storage device, you must set this to **No**.

Please note that if you set this directive to **No** Bacula will release the tape drive between each job, and thus the next job will rewind the tape and position it to the end of the data. This can be a very time consuming operation.

Volume Poll Interval = *time* If the time specified on this directive is non-zero, after asking the operator to mount a new volume Bacula will periodically poll (or read) the drive at the specified interval to see if a new volume has been mounted. If the time interval is zero (the default), no polling will occur. This directive can be useful if you want to avoid operator intervention via the console. Instead, the operator can simply remove the old volume and insert the requested one, and Bacula on the next poll will recognize the new tape and continue. Please be aware that if you set this interval too small, you may excessively wear your tape drive if the old tape remains in the drive, since Bacula will read it on each poll. This can be avoided by ejecting the tape using the **Offline On Unmount** and the **Close on Poll** directives.

Close on Poll= *Yes—No* If **Yes**, Bacula close the device (equivalent to an unmount except no mount is required) and reopen it at each poll. Normally this is not too useful unless you have the **Offline on Unmount** directive set, in which case the drive will be taken offline preventing wear on the tape during any future polling. Once the operator inserts a new tape, Bacula will recognize the drive on the next poll and automatically continue with the backup.

Maximum Open Wait = *time* This directive specifies the maximum amount of time that Bacula will wait for a device that is busy. The default is 5 minutes. If the device cannot be obtained, the current Job will be terminated in error. Bacula will re-attempt to open the drive the next time a Job starts that needs the the drive.

Removable media = *Yes—No* If **Yes**, this device supports removable media (for example, tapes or CDs). If **No**, media cannot be removed (for example, an intermediate backup area on a hard disk).

Random access = *Yes—No* If **Yes**, the archive device is assumed to be a random access medium which supports the **lseek** (or **lseek64** if Largefile is enabled during configuration) facility.

Minimum block size = *size-in-bytes* On most modern tape drives, you will not need to specify this directive, and if you do so, it will be to make Bacula use fixed block sizes. This statement applies only to non-random access devices (e.g. tape drives). Blocks written by the storage daemon to a non-random archive device will never be smaller than the given **size-in-bytes**. The Storage daemon will attempt to efficiently fill blocks with data received from active sessions but will, if necessary, add padding to a block to achieve the required minimum size.

To force the block size to be fixed, as is the case for some non-random access devices (tape drives), set the **Minimum block size** and the **Maximum block size** to the same value (zero included). The default is that both the minimum and maximum block size are zero and the default block size is 64,512 bytes. If you wish the block size to be fixed and different from the default, specify the same value for both **Minimum block size** and **Maximum block size**.

For example, suppose you want a fixed block size of 100K bytes, then you would specify:

```
Minimum block size = 100K
Maximum block size = 100K
```

Please note that if you specify a fixed block size as shown above, the tape drive must either be in variable block size mode, or if it is in fixed block size mode, the block size (generally defined by **mt**) **must** be identical to the size specified in Bacula – otherwise when you attempt to re-read your Volumes, you will get an error.

If you want the block size to be variable but with a 64K minimum and 200K maximum (and default as well), you would specify:

```
Minimum block size = 64K
Maximum blocksize = 200K
```

Maximum block size = *size-in-bytes* On most modern tape drives, you will not need to specify this directive. If you do so, it will most likely be to use fixed block sizes (see Minimum block size above). The Storage daemon will aways attempt to write blocks of the specified **size-in-bytes** to the archive device. As a consequence, this statement specifies both the default block size and the maximum block size. The size written never exceed the given **size-in-bytes**. If adding data to a block would cause it to exceed the given maximum size, the block will be written to the archive device, and the new data will begin a new block.

If no value is specified or zero is specified, the Storage daemon will use a default block size of 64,512 bytes (126 * 512).

Hardware End of Medium = *Yes—No* If **No**, the archive device is not required to support end of medium ioctl request, and the storage daemon will use the forward space file function to find the end of

the recorded data. If **Yes**, the archive device must support the `ioctl MTEOM` call, which will position the tape to the end of the recorded data. In addition, your SCSI driver must keep track of the file number on the tape and report it back correctly by the **MTIOCGET** ioctl. Note, some SCSI drivers will correctly forward space to the end of the recorded data, but they do not keep track of the file number. On Linux machines, the SCSI driver has a **fast-eod** option, which if set will cause the driver to lose track of the file number. You should ensure that this option is always turned off using the **mt** program.

Default setting for Hardware End of Medium is **Yes**. This function is used before appending to a tape to ensure that no previously written data is lost. We recommend if you have a non standard or unusual tape drive that you use the **btape** program to test your drive to see whether or not it supports this function. All modern (after 1998) tape drives support this feature.

If you set Hardware End of Medium = no, you should also set **Fast Forward Space File = no**. If you do not, Bacula will most likely be unable to correctly find the end of data on the tape.

Fast Forward Space File = *Yes—No* If **No**, the archive device is not required to support keeping track of the file number (**MTIOCGET** ioctl) during forward space file. If **Yes**, the archive device must support the `ioctl MTFSF` call, which virtually all drivers support, but in addition, your SCSI driver must keep track of the file number on the tape and report it back correctly by the **MTIOCGET** ioctl. Note, some SCSI drivers will correctly forward space, but they do not keep track of the file number or more seriously, they do not report end of meduim.

Default setting for Fast Forward Space File is **Yes**. If you disable Hardware End of Medium, most likely you should also disable Fast Forward Space file. The **test** command in the program **btape** will test this feature and advise you if it should be turned off.

BSF at EOM = *Yes—No* If **No**, the default, no special action is taken by Bacula with the End of Medium (end of tape) is reached because the tape will be positioned after the last EOF tape mark, and Bacula can append to the tape as desired. However, on some systems, such as FreeBSD, when Bacula reads the End of Medium (end of tape), the tape will be positioned after the second EOF tape mark (two successive EOF marks indicated End of Medium). If Bacula appends from that point, all the appended data will be lost. The solution for such systems is to specify **BSF at EOM** which causes Bacula to backspace over the second EOF mark. Determination of whether or not you need this directive is done using the **test** command in the **btape** program.

TWO EOF = *Yes—No* If **Yes**, Bacula will write two end of file marks when terminating a tape – i.e. after the last job or at the end of the medium. If **No**, the default, Bacula will only write one end of file to terminate the tape.

Backward Space Record = *Yes—No* If *Yes*, the archive device supports the `MTBSR ioctl` to backspace records. If *No*, this call is not used and the device must be rewound and advanced forward to the desired position. Default is **Yes** for non random-access devices.

Backward Space File = *Yes—No* If *Yes*, the archive device supports the **MTBSF** and **MTBSF ioctl**s to backspace over an end of file mark and to the start of a file. If *No*, these calls are not used and the device must be rewound and advanced forward to the desired position. Default is **Yes** for non random-access devices.

Forward Space Record = *Yes—No* If *Yes*, the archive device must support the **MTFSR ioctl** to forward space over records. If **No**, data must be read in order to advance the position on the device. Default is **Yes** for non random-access devices.

Forward Space File = *Yes—No* If **Yes**, the archive device must support the `MTFSF ioctl` to forward space by file marks. If *No*, data must be read to advance the position on the device. Default is **Yes** for non random-access devices.

Offline On Unmount = *Yes—No* The default for this directive is **No**. If **Yes** the archive device must support the `MTOFFL ioctl` to rewind and take the volume offline. In this case, Bacula will issue the offline (eject) request before closing the device during the **unmount** command. If **No** Bacula will not attempt to offline the device before unmounting it. After an offline is issued, the cassette will be ejected thus **requiring operator intervention** to continue, and on some systems require an explicit load command to be issued (**mt -f /dev/xxx load**) before the system will recognize the tape. If you are using an autochanger, some devices require an offline to be issued prior to changing the volume. However, most devices do not and may get very confused.

Maximum Volume Size = *size* No more than **size** bytes will be written onto a given volume on the archive device. This directive is used mainly in testing Bacula to simulate a small Volume. It can also be useful if you wish to limit the size of a File Volume to say less than 2GB of data. In some rare cases of really antiquated tape drives that do not properly indicate when the end of a tape is reached during writing (though I have read about such drives, I have never personally encountered one). Please note, this directive is deprecated (being

phased out) in favor of the **Maximum Volume Bytes** defined in the Director's configuration file.

Maximum File Size = *size* No more than **size** bytes will be written into a given logical file on the volume. Once this size is reached, an end of file mark is written on the volume and subsequent data are written into the next file. Breaking long sequences of data blocks with file marks permits quicker positioning to the start of a given stream of data and can improve recovery from read errors on the volume. The default is one Gigabyte.

Block Positioning = *yes—no* This directive is not normally used (and has not yet been tested). It will tell Bacula not to use block positioning when it is reading tapes. This can cause Bacula to be **extremely** slow when restoring files. You might use this directive if you wrote your tapes with Bacula in variable block mode (the default), but your drive was in fixed block mode. If it then works as I hope, Bacula will be able to re-read your tapes.

Maximum Network Buffer Size = *bytes* where *bytes* specifies the initial network buffer size to use with the File daemon. This size will be adjusted down if it is too large until it is accepted by the OS. Please use care in setting this value since if it is too large, it will be trimmed by 512 bytes until the OS is happy, which may require a large number of system calls. The default value is 32,768 bytes.

Maximum Spool Size = *bytes* where the bytes specify the maximum spool size for all jobs that are running. The default is no limit.

Maximum Job Spool Size = *bytes* where the bytes specify the maximum spool size for any one job that is running. The default is no limit. This directive is implemented only in version 1.37 and later.

Spool Directory = *directory* specifies the name of the directory to be used to store the spool files for this device. This directory is also used to store temporary part files when writing to a device that requires mount (DVD). The default is to use the working directory.

Capabilities

Label media = *Yes—No* If **Yes**, permits this device to automatically label blank media without an explicit operator command. It does so by using an internal algorithm as defined on the Label Format record in each Pool resource. If this is **No** as by default, Bacula will label tapes only by specific operator command (**label** in the Console) or

when the tape has been recycled. The automatic labeling feature is most useful when writing to disk rather than tape volumes.

Automatic mount = *Yes—No* If **Yes** (the default), permits the daemon to examine the device to determine if it contains a Bacula labeled volume. This is done initially when the daemon is started, and then at the beginning of each job. This directive is particularly important if you have set **Always Open = no** because it permits Bacula to attempt to read the device before asking the system operator to mount a tape.

Messages Resource

For a description of the Messages Resource, please see the Messages Resource Chapter of this manual.

Sample Storage Daemon Configuration File

A example Storage Daemon configuration file might be the following:

```
#
# Default Bacula Storage Daemon Configuration file
#
#  For Bacula release 1.35.2 (16 August 2004) -- gentoo 1.4.16
#
# You may need to change the name of your tape drive
#   on the "Archive Device" directive in the Device
#   resource.  If you change the Name and/or the
#   "Media Type" in the Device resource, please ensure
#   that bacula-dir.conf has corresponding changes.
#
Storage {                             # definition of myself
  Name = rufus-sd
  Address = rufus
  WorkingDirectory = "$HOME/bacula/bin/working"
  Pid Directory = "$HOME/bacula/bin/working"
  Maximum Concurrent Jobs = 20
}
#
# List Directors who are permitted to contact Storage daemon
#
Director {
  Name = rufus-dir
  Password = "ZF9Ctf5PQoWCPkmR3s4atCBOusUPg+vWWyIo2VS5ti6k"
}
#
```

```
# Restricted Director, used by tray-monitor to get the
#   status of the storage daemon
#
Director {
  Name = rufus-mon
  Password = "9usxgc307dMbe7jbD16v0PXlhD64UVasIDD0DH2WAujcDsc6"
  Monitor = yes
}
#
# Devices supported by this Storage daemon
# To connect, the Director's bacula-dir.conf must have the
#  same Name and MediaType.
#
Device {
  Name = "HP DLT 80"
  Media Type = DLT8000
  Archive Device = /dev/nst0
  AutomaticMount = yes;                 # when device opened, read it
  AlwaysOpen = yes;
  RemovableMedia = yes;
}
#Device {
#  Name = SDT-7000                     #
#  Media Type = DDS-2
#  Archive Device = /dev/nst0
#  AutomaticMount = yes;               # when device opened, read it
#  AlwaysOpen = yes;
#  RemovableMedia = yes;
#}
#Device {
#  Name = Floppy
#  Media Type = Floppy
#  Archive Device = /mnt/floppy
#  RemovableMedia = yes;
#  Random Access = Yes;
#  AutomaticMount = yes;               # when device opened, read it
#  AlwaysOpen = no;
#}
#Device {
#  Name = FileStorage
#  Media Type = File
#  Archive Device = /tmp
#  LabelMedia = yes;                   # lets Bacula label unlabeled media
#  Random Access = Yes;
#  AutomaticMount = yes;               # when device opened, read it
#  RemovableMedia = no;
#  AlwaysOpen = no;
#}
#Device {
#  Name = "NEC ND-1300A"
#  Media Type = DVD
#  Archive Device = /dev/hda
#  LabelMedia = yes;                   # lets Bacula label unlabeled media
#  Random Access = Yes;
```

```
#  AutomaticMount = yes;               # when device opened, read it
#  RemovableMedia = yes;
#  AlwaysOpen = no;
#  MaximumPartSize = 800M;
#  RequiresMount = yes;
#  SpoolDirectory = /tmp/backup;
#}
#
# A very old Exabyte with no end of media detection
#
#Device {
#  Name = "Exabyte 8mm"
#  Media Type = "8mm"
#  Archive Device = /dev/nst0
#  Hardware end of medium = No;
#  AutomaticMount = yes;               # when device opened, read it
#  AlwaysOpen = Yes;
#  RemovableMedia = yes;
#}
#
# Send all messages to the Director,
# mount messages also are sent to the email address
#
Messages {
  Name = Standard
  director = rufus-dir = all
  operator = root = mount
}
```

Messages Resource

The Messages Resource

The Messages resource defines how messages are to be handled and destinations to which they should be sent.

Even though each daemon has a full message handler, within the File daemon and the Storage daemon, you will normally choose to send all the appropriate messages back to the Director. This permits all the messages associated with a single Job to be combined in the Director and sent as a single email message to the user, or logged together in a single file.

Each message that Bacula generates (i.e. that each daemon generates) has an associated type such as INFO, WARNING, ERROR, FATAL, etc. Using the message resource, you can specify which message types you wish to see and where they should be sent. In addition, a message may be sent to multiple destinations. For example, you may want all error messages both logged as well as sent to you in an email. By defining multiple messages resources, you can have different message handling for each type of Job (e.g. Full backups versus Incremental backups).

In general, messages are attached to a Job and are included in the Job report. There are some rare cases, where this is not possible, e.g. when no job is running, or if a communications error occurs between a daemon and the director. In those cases, the message may remain in the system, and should be flushed at the end of the next Job. However, since such messages are not attached to a Job, any that are mailed will be sent to **/usr/lib/sendmail**. On some systems, such as FreeBSD, if your sendmail is in a different place, you may want to link it to the the above location.

The records contained in a Messages resource consist of a **destination** specification followed by a list of **message-types** in the format:

destination = message-type1, message-type2, message-type3, ...

or for those destinations that need and address specification (e.g. email):

destination = address = message-type1, message-type2, message-type3, ...

Where **destination** is one of a predefined set of keywords that define where the message is to be sent (**stdout**, **file**, ...), **message-type** is

one of a predefined set of keywords that define the type of message generated by **Bacula** (**ERROR**, **WARNING**, **FATAL**, ...), and **address** varies according to the **destination** keyword, but is typically and email address or a filename.

The following are the list of the possible record definitions that can be used in a message resource.

Messages Start of the Messages records.

Name = <name> The name of the Messages resource. The name you specify here will be used to tie this Messages resource to a Job and/or to the daemon.

MailCommand = <command> In the absence of this resource, Bacula will send all mail using the following command:

mail -s "Bacula Message" <recipients>

In many cases, depending on your machine, this command may not work. Using the **MailCommand**, you can specify exactly how to send the mail. During the processing of the **command**, normally specified as a quoted string, the following substitutions will be used:

- %% = %
- %c = Client's name
- %d = Director's name
- %e = Job Exit code (OK, Error, ...)
- %i = Job Id
- %j = Unique Job name
- %l = Job level
- %n = Job name
- %r = Recipients
- %t = Job type (e.g. Backup, ...)

The following is the command I (Kern) use. Note, the whole command should appear on a single line in the configuration file rather than split as is done here for presentation:

mailcommand = "/home/kern/bacula/bin/bsmtp -h mail.whitehouse.com -f \"\(Bacula\) %r\" -s \"Bacula: %t %e of %c %l\" %r"

Note, the **bsmtp** program is provided as part of **Bacula**. For additional details, please see the bsmtp – Customizing Your Email Messages section of the Bacula Utility Programs chapter of this manual. Please test any **mailcommand** that you use to ensure that your bsmtp gateway accepts the addressing form that you use. Certain program such as Exim can be very selective as to what forms are permitted particularly in the from part.

OperatorCommand = <command> This resource specification is similar to the **MailCommand** except that it is used for Operator messages. The substitutions performed for the **MailCommand** are also done for this command. Normally, you will set this command to the same value as specified for the **MailCommand**.

Debug = <debug-level> This sets the debug message level to the debug level, which is an integer. Higher debug levels cause more debug information to be produced. You are requested not to use this record since it will be deprecated.

<destination> = <message-type1>, <message-type2&gt, ...
Where **destination** may be one of the following:

stdout Send the message to standard output.

stderr Send the message to standard error.

console Send the message to the console (Bacula Console). These messages are held until the console program connects to the Director.

<destination> = <address> = <message-type1>, <message-type2&gt, ...

Where **address** depends on the **destination**, which may be one of the following:

director Send the message to the Director whose name is given in the **address** field. Note, in the current implementation, the Director Name is ignored, and the message is sent to the Director that started the Job.

file Send the message to the filename given in the **address** field. If the file already exists, it will be overwritten.

append Append the message to the filename given in the **address** field. If the file already exists, it will be appended to. If the file does not exist, it will be created.

syslog Send the message to the system log (syslog) using the facility specified in the **address** field. Note, for the moment, the **address** field is ignored and the message is always sent to the **LOG_ERR** facility.

mail Send the message to the email addresses that are given as a comma separated list in the **address** field. Mail messages are grouped together during a job and then sent as a single email message when the job terminates. The advantage of this destination is that you are notified about every Job that runs. However, if you backup 5 or 10 machines every night, the volume of email messages can be important. Some users use filter programs such as **procmail** to automatically file this email based on the Job termination code (see **mailcommand**).

mail on error Send the message to the email addresses that are given as a comma separated list in the **address** field if the Job terminates with an error condition. MailOnError messages are grouped together during a job and then sent as a single email message when the job terminates. This destination differs from the **mail** destination in that if the Job terminates normally, the message is totally discarded (for this destination). If the Job terminates in error, it is emailed. By using other destinations such as **append** you can ensure that even if the Job terminates normally, the output information is saved.

operator Send the message to the email addresses that are specified as a comma separated list in the **address** field. This is similar to **mail** above, except that each message is sent as received. Thus there is one email per message. This is most useful for **mount** messages (see below).

For any destination, the **message-type** field is a comma separated list of the following types or classes of messages:

info General information messages.

warning Warning messages. Generally this is some unusual condition but not expected to be serious.

error Non-fatal error messages. The job continues running. Any error message should be investigated as it means that something went wrong.

fatal Fatal error messages. Fatal errors cause the job to terminate.

terminate Message generated when the daemon shuts down.

saved Files saved normally.

notsaved Files not saved because of some error. Usually because the file cannot be accessed (i.e. it does not exist or is not mounted).

skipped Files that were skipped because of a user supplied option such as an incremental backup or a file that matches an exclusion pattern. This is not considered an error condition such as the files listed for the **notsaved** type because the configuration file explicitly requests these types of files to be skipped. For example, any unchanged file during an incremental backup, or any subdirectory if the no recursion option is specified.

mount Volume mount or intervention requests from the Storage daemon. These requests require a specific operator intervention for the job to continue.

restored The **ls** style listing generated for each file restored is sent to this message class.

all All message types.

***security** Security info/warning messages (not currently implemented).

The following is an example of a valid Messages resource definition, where all messages except files explicitly skipped or daemon termination messages are sent by email to enforcement@sec.com. In addition all mount messages are sent to the operator (i.e. emailed to enforcement@sec.com). Finally all messages other than explicitly skipped files and files saved are sent to the console:

```
Messages {
  Name = Standard
  mail = enforcement@sec.com = all, !skipped, !terminate
  operator = enforcement@sec.com = mount
  console = all, !skipped, !saved
}
```

With the exception of the email address (changed to avoid junk mail from robot's), Kern's Director's Messages resource is as follows. Note, the **mailcommand** and **operatorcommand** are on a single line – they had to be split for this manual:

```
Messages {
  Name = Standard
  mailcommand = "bacula/bin/bsmtp -h mail.whitehouse.com \
    -f \"\(Bacula\) %r\" -s \"Bacula: %t %e of %c %l\" %r"
  operatorcommand = "bacula/bin/bsmtp -h mail.whitehouse.com \
```

```
    -f \"\(Bacula\) %r\" -s \"Bacula: Intervention needed \
        for %j\" %r"
  MailOnError = security@whitehouse.com = all, !skipped, \
                !terminate
  append = "bacula/bin/log" = all, !skipped, !terminate
  operator = security@whitehouse.com = mount
  console = all, !skipped, !saved
}
```

Console Configuration

General

The Console configuration file is the simplest of all the configuration files, and in general, you should not need to change it except for the password. It simply contains the information necessary to contact the Director or Directors.

For a general discussion of configuration file and resources including the data types recognized by **Bacula**, please see the Configuration chapter of this manual.

The following Console Resource definition must be defined:

- Director – to define the Director's name and his access password. Note, you may define more than one Director resource in the Console configuration file. If you do so, the Console program will ask you which one you want to use.

The Director Resource

The Director resource defines the attributes of the Director running on the network. You may have multiple Director resource specifications in a single Console configuration file. If you have more than one, you will be prompted to choose one when you start the **Console** program.

Director Start of the Director records.

Name = <name> The director name used to select among different Directors, otherwise, this name is not used.

DIRPort = <port-number> Specify the port to use to connect to the Director. This value will most likely already be set to the value you specified on the **--with-base-port** option of the **./configure** command. This port must be identical to the **DIRport** specified in the **Director** resource of the Director's configuration file. The default is 9101 so this record is not normally specified.

Address = <address> Where the address is a host name, a fully qualified domain name, or a network address used to connect to the Director.

Password = <password> Where the password is the password needed for the Director to accept the Console connection. This password must be identical to the **Password** specified in the **Director** resource of the Director's configuration file. This record is required.

An actual example might be:

```
Director {
  Name = HeadMan
  address = rufus.cats.com
  password = xyz1erploit
}
```

The ConsoleFont Resource

The ConsoleFont resource is available only in the GNOME version of the console. It permits you to define the font that you want used to display in the main listing window.

ConsoleFont Start of the ConsoleFont records.

Name = <name> The name of the font.

Font = <X-Window Font Specification> The string value given here defines the desired font. It is specified in the standard cryptic X Window format. For example, the default specification is:

```
Font = "-misc-fixed-medium-r-normal-*-*-130-*-*-c-*-iso8859-1"
```

Thanks to Phil Stracchino for providing the code for this feature.

An actual example might be:

```
ConsoleFont {
  Name = Default
Font = "-misc-fixed-medium-r-normal-*-*-130-*-*-c-*-iso8859-1"
}
```

The Console Resource

As of Bacula version 1.33 and higher, there are three different kinds of consoles, which the administrator or user can use to interact with the

Director. These three kinds of consoles comprise three different security levels.

- The first console type is an **anonymous** or **default** console, which has full privileges. There is no console resource necessary for this type since the password is specified in the Director resource. This is the kind of console that was initially implemented in versions prior to 1.33 and remains valid. Typically you would use it only for administrators.
- The second type of console, and new to version 1.33 and higher is a "named" console defined within a Console resource in both the Director's configuration file and in the Console's configuration file. Both the names and the passwords in these two entries must match much as is the case for Client programs.

  This second type of console begins with absolutely no privileges except those explicitly specified in the Director's Console resource. Thus you can have multiple Consoles with different names and passwords, sort of like multiple users, each with different privileges. As a default, these consoles can do absolutely nothing – no commands what so ever. You give them privileges or rather access to commands and resources by specifying access control lists in the Director's Console resource. Note, if you are specifying such a console, you will want to put a null password in the Director resource.
- The third type of console is similar to the above mentioned one in that it requires a Console resource definition in both the Director and the Console. In addition, if the console name, provided on the **Name** = directive, is the same as a Client name, the user of that console is permitted to use the **SetIP** command to change the Address directive in the Director's client resource to the IP address of the Console. This permits portables or other machines using DHCP (non-fixed IP addresses) to "notify" the Director of their current IP address.

The Console resource is optional and need not be specified. However, if it is specified, you can use ACLs (Access Control Lists) in the Director's configuration file to restrict the particular console (or user) to see only information pertaining to his jobs or client machine.

The following configuration files were supplied by Phil Stracchino. For example, if we define the following in the user's bconsole.conf file (or perhaps the wx-console.conf file):

```
Director {
```

```
   Name = MyDirector
   DIRport = 9101
   Address = myserver
   Password = "XXXXXXXXXX"    # no, really.  this is not obfuscation.
 }

 Console {
   Name = restricted-user
   Password = "UntrustedUser"
 }
```

Where the Password in the Director section is deliberately incorrect, and the Console resource is given a name, in this case **restricted-client**. Then in the Director's bacula-dir.conf file (not directly accessible by the user), we define:

```
Console {
  Name = restricted-user
  Password = "UntrustedUser"
  JobACL = "Restricted Client Save"
  ClientACL = restricted-client
  StorageACL = main-storage
  ScheduleACL = *all*
  PoolACL = *all*
  FileSetACL = "Restricted Client's FileSet"
  CatalogACL = DefaultCatalog
  CommandACL = run
}
```

the user logging into the Director from his Console will get logged in as **restricted-client**, and he will only be able to see or access a Job with the name **Restricted Client Save** a Client with the name **restricted-client**, a Storage device **main-storage**, any Schedule or Pool, a FileSet named **Restricted Client's File**, a Catalog named **DefaultCatalog**, and the only command he can use in the Console is the **run** command. In other words, this user is rather limited in what he can see and do with Bacula.

Console Commands

For more details on running the console and its commands, please see the Bacula Console chapter of this manual.

Sample Console Configuration File

A example Console configuration file might be the following:

```
#
# Bacula Console Configuration File
#
Director {
  Name = HeadMan
  address = "my_machine.my_domain.com"
  Password = Console_password
}
```

Monitor Configuration

General

The Monitor configuration file is a stripped down version of the Director configuration file, mixed with a Console configuration file. It simply contains the information necessary to contact Directors, Clients, and Storage daemons you want to monitor.

For a general discussion of configuration file and resources including the data types recognized by **Bacula**, please see the Configuration chapter of this manual.

The following Monitor Resource definition must be defined:

- Monitor – to define the Monitor's name used to connect to all the daemons and the password used to connect to the Directors. Note, you must not define more than one Monitor resource in the Monitor configuration file.
- At least one Client, Storage or Director resource, to define the daemons to monitor.

The Monitor Resource

The Monitor resource defines the attributes of the Monitor running on the network. The parameters you define here must be configured as a Director resource in Clients and Storages configuration files, and as a Console resource in Directors configuration files.

Monitor Start of the Monitor records.

Name = <name> Specify the Director name used to connect to Client and Storage, and the Console name used to connect to Director. This record is required.

Password = <password> Where the password is the password needed for Directors to accept the Console connection. This password must be identical to the **Password** specified in the **Console** resource of the Director's configuration file. This record is required if you wish to monitor Directors.

Refresh Interval = <time> Specifies the time to wait between status requests to each daemon. It can't be set to less than 1 second, or more than 10 minutes, and the default value is 5 seconds.

The Director Resource

The Director resource defines the attributes of the Directors that are monitored by this Monitor.

As you are not permitted to define a Password in this resource, to avoid obtaining full Director privileges, you must create a Console resource in the Director's configuration file, using the Console Name and Password defined in the Monitor resource. To avoid security problems, you should configure this Console resource to allow access to no others daemon, and permit the use of only two commands: **status** and **.status** (see below for an example).

You may have multiple Director resource specifications in a single Monitor configuration file.

Director Start of the Director records.

Name = <name> The Director name used to identify the Director in the list of monitored daemons. It is not required to be the same as defined in the Director's configuration file. This record is required.

DIRPort = <port-number> Specify the port to use to connect to the Director. This value will most likely already be set to the value you specified on the **--with-base-port** option of the **./configure** command. This port must be identical to the **DIRport** specified in the **Director** resource of the Director's configuration file. The default is 9101 so this record is not normally specified.

Address = <address> Where the address is a host name, a fully qualified domain name, or a network address used to connect to the Director. This record is required.

The Client Resource

The Client resource defines the attributes of the Clients that are monitored by this Monitor.

You must create a Director resource in the Client's configuration file, using the Director Name defined in the Monitor resource. To avoid security

problems, you should set the **Monitor** directive to **Yes** in this Director resource.

You may have multiple Director resource specifications in a single Monitor configuration file.

Client (or FileDaemon) Start of the Client records.

Name = <name> The Client name used to identify the Director in the list of monitored daemons. It is not required to be the same as defined in the Client's configuration file. This record is required.

Address = <address> Where the address is a host name, a fully qualified domain name, or a network address in dotted quad notation for a Bacula File daemon. This record is required.

FD Port = <port-number> Where the port is a port number at which the Bacula File daemon can be contacted. The default is 9102.

Password = <password> This is the password to be used when establishing a connection with the File services, so the Client configuration file on the machine to be backed up must have the same password defined for this Director. This record is required.

The Storage Resource

The Storage resource defines the attributes of the Storages that are monitored by this Monitor.

You must create a Director resource in the Storage's configuration file, using the Director Name defined in the Monitor resource. To avoid security problems, you should set the **Monitor** directive to **Yes** in this Director resource.

You may have multiple Director resource specifications in a single Monitor configuration file.

Storage Start of the Storage records.

Name = <name> The Storage name used to identify the Director in the list of monitored daemons. It is not required to be the same as defined in the Storage's configuration file. This record is required.

Address = <address> Where the address is a host name, a fully qualified domain name, or a network address in dotted quad notation for a Bacula Storage daemon. This record is required.

SD Port = <port> Where port is the port to use to contact the storage daemon for information and to start jobs. This same port number must appear in the Storage resource of the Storage daemon's configuration file. The default is 9103.

Password = <password> This is the password to be used when establishing a connection with the Storage services. This same password also must appear in the Director resource of the Storage daemon's configuration file. This record is required.

Sample Monitor configuration file and related daemons' configuration records.

A example Monitor configuration file might be the following:

```
#
# Bacula Tray Monitor Configuration File
#
Monitor {
  Name = rufus-mon        # password for Directors
  Password = "GN0uRo7PTUmlMbqrJ2Gr1p0fk0HQJTxwnFyE4WSST3MWZseR"
  RefreshInterval = 10 seconds
}

Client {
  Name = rufus-fd
  Address = rufus
  FDPort = 9102           # password for FileDaemon
  Password = "FYpq4yyI1y562EMS35bA0J0QC0M2L3t5cZObxT3XQxgxppTn"
}
Storage {
  Name = rufus-sd
  Address = rufus
  SDPort = 9103           # password for StorageDaemon
  Password = "9usxgc307dMbe7jbD16v0PXlhD64UVasIDD0DH2WAujcDsc6"
}
Director {
  Name = rufus-dir
  DIRport = 9101
  address = rufus
}
```

Sample File daemon's Director record.

Click here to see the full example.

```
#
# Restricted Director, used by tray-monitor to get the
#   status of the file daemon
#
Director {
  Name = rufus-mon
  Password = "FYpq4yyI1y562EMS35bA0J0QC0M2L3t5cZObxT3XQxgxppTn"
  Monitor = yes
}
```

Sample Storage daemon's Director record.

Click here to see the full example.

```
#
# Restricted Director, used by tray-monitor to get the
#   status of the storage daemon
#
Director {
  Name = rufus-mon
  Password = "9usxgc307dMbe7jbD16v0PXlhD64UVasIDD0DH2WAujcDsc6"
  Monitor = yes
}
```

Sample Director's Console record.

Click here to see the full example.

```
#
# Restricted console used by tray-monitor to get the status of the director
#
Console {
  Name = Monitor
  Password = "GN0uRo7PTUmlMbqrJ2Gr1p0fk0HQJTxwnFyE4WSST3MWZseR"
  CommandACL = status, .status
}
```

Bacula Console

General

The **Bacula Console** (sometimes called the User Agent) is a program that allows the user or the System Administrator, to interact with the Bacula Director daemon while the daemon is running.

The current Bacula Console comes in two versions: a shell interface (TTY style), and a GNOME GUI interface. Both permit the administrator or authorized users to interact with Bacula. You can determine the status of a particular job, examine the contents of the Catalog as well as perform certain tape manipulations with the Console program.

In addition, there is a wx-console built with wxWidgets that allows a graphic restore of files. As of version 1.34.1 it is in an early stage of development, but it quite useful.

Since the Console program interacts with the Director by the network, your Console and Director programs do not necessarily need to run on the same machine.

In fact, a certain minimal knowledge of the Console program is needed in order for Bacula to be able to write on more than one tape, because when Bacula requests a new tape, it waits until the user, via the Console program, indicates that the new tape is mounted.

Configuration

When the Console starts, it reads a standard Bacula configuration file named **bconsole.conf** or **gnome-console.conf** in the case of the GNOME Console version. This file allows default configuration of the Console, and at the current time, the only Resource Record defined is the Director resource, which gives the Console the name and address of the Director. For more information on configuration of the Console program, please see the Console Configuration File Chapter of this document.

Running the Console Program

After launching the Console program (bconsole), it will prompt you for the next command with an asterisk (*). (Note, in the GNOME version, the prompt is not present; you simply enter the commands you want in

the command text box at the bottom of the screen.) Generally, for all commands, you can simply enter the command name and the Console program will prompt you for the necessary arguments. Alternatively, in most cases, you may enter the command followed by arguments. The general format is:

```
<command> <keyword1>[=<argument1>] <keyword2>[=<argument2>] ...
```

where **command** is one of the commands listed below; **keyword** is one of the keywords listed below (usually followed by an argument); and **argument** is the value. The command may be abbreviated to the shortest unique form. If two commands have the same starting letters, the one that will be selected is the one that appears first in the **help** listing. If you want the second command, simply spell out the full command. None of the keywords following the command may be abbreviated.

For example:

```
list files jobid=23
```

will list all files saved for JobId 23. Or:

```
show pools
```

will display all the Pool resource records.

Stopping the Console Program

Normally, you simply enter **quit** or **exit** and the Console program will terminate. However, it waits until the Director acknowledges the command. If the Director is already doing a lengthy command (e.g. prune), it may take some time. If you want to immediately terminate the Console program, enter the **.quit** command.

There is currently no way to interrupt a Console command once issued (i.e. ctl-C does not work). However, if you are at a prompt that is asking you to select one of several possibilities and you would like to abort the command, you can enter a period (**.**), and in most cases, you will either be returned to the main command prompt or if appropriate the previous prompt (in the case of nested prompts). In a few places such as where it is asking for a Volume name, the period will be taken to be the Volume name. In that case, you will most likely be able to cancel at the next prompt.

Alphabetic List of Console Commands

The following commands are currently implemented:

add [pool=<pool-name> storage=<storage> jobid=<JobId>]
This command is used to add Volumes to an existing Pool. The Volume names entered are placed in the Catalog and thus become available for backup operations. Normally, the **label** command is used rather than this command because the **label** command labels the physical media (tape) and does the equivalent of the **add** command. This command affects only the Catalog and not the physical media (data on Volumes). The physical media must exist and be labeled before use (usually with the **label** command). This command can, however, be useful if you wish to add a number of Volumes to the Pool that will be physically labeled at a later time. It can also be useful if you are importing a tape from another site. Please see the **label** command below for the list of legal characters in a Volume name.

autodisplay on/off This command accepts **on** or **off** as an argument, and turns auto-display of messages on or off respectively. The default for the console program is **off**, which means that you will be notified when there are console messages pending, but they will not automatically be displayed. The default for the gnome-console program is **on**, which means that messages will be displayed when they are received (usually within 5 seconds of them being generated).

When autodisplay is turned off, you must explicitly retrieve the messages with the **messages** command. When autodisplay is turned on, the messages will be displayed on the console as they are received.

automount on/off This command accepts **on** or **off** as the argument, and turns auto-mounting of the tape after a **label** command on or off respectively. The default is **on**. If **automount** is turned off, you must explicitly **mount** the tape after a label command to use it.

cancel [jobid=<number> job=<job-name>] This command is used to cancel a job and accepts **jobid=nnn** or **job=xxx** as an argument where nnn is replaced by the JobId and xxx is replaced by the job name. If you do not specify a keyword, the Console program will prompt you with the names of all the active jobs allowing you to choose one.

Once a Job is marked to be canceled, it may take a bit of time (generally within a minute) before it actually terminates, depending on what operations it is doing.

create [pool=<pool-name>] This command is used to create a Pool record in the database using the Pool resource record defined in the Director's configuration file. So in a sense, this command simply transfers the information from the Pool resource in the configuration file into the Catalog. Normally this command is done automatically for you when the Director starts providing the Pool is referenced within a Job resource. If you use this command on an existing Pool, it will automatically update the Catalog to have the same information as the Pool resource. After creating a Pool, you will most likely use the **label** command to label one or more volumes and add their names to the Media database.

When starting a Job, if Bacula determines that there is no Pool record in the database, but there is a Pool resource of the appropriate name, it will create it for you. If you want the Pool record to appear in the database immediately, simply use this command to force it to be created.

delete [volume=<vol-name> pool=<pool-name> job jobid=<id>] The delete command is used to delete a Volume, Pool or Job record from the Catalog as well as all associated Volume records that were created. This command operates only on the Catalog database and has no effect on the actual data written to a Volume. This command can be dangerous and we strongly recommend that you do not use it unless you know what you are doing.

If the keyword **Volume** appears on the command line, the named Volume will be deleted from the catalog, if the keyword **Pool** appears on the command line, a Pool will be deleted, and if the keyword **Job** appears on the command line, a Job and all its associated records (File and JobMedia) will be deleted from the catalog. The full form of this command is:

delete pool=<pool-name>

or

delete volume=<volume-name> pool=<pool-name> or

delete JobId=<job-id> JobId=<job-id2> ... or

delete Job JobId=n,m,o-r,t ...

The first form deletes a Pool record from the catalog database. The second form deletes a Volume record from the specified pool in the catalog database. The third form delete the specified Job record from the catalog database. The last form deletes JobId records for JobIds n,m,o,p, q,r, and t. When each one of the n,m,... is, of course, a number.

estimate Using this command, you can get an idea how many files will be backed up, or if you are unsure about your Include statements in your FileSet, you can test them without doing an actual backup. The default is to assume a Full backup. However, you can override this by specifying a **level=Incremental** or **level=Differential** on the command line. A Job name must be specified or you will be prompted for one, and optionally a Client and FileSet may be specified on the command line. It then contacts the client which computes the number of files and bytes that would be backed up. Please note that this is an estimated calculated from the number of blocks in the file rather than by reading the actual bytes. As such, the estimated backup size will generally be larger than an actual backup.

Optionally you may specify the keyword **listing** in which case, all the files to be backed up will be listed. Note, it could take quite some time to display them if the backup is large. The full form is:

estimate job=<job-name> listing client=<client-name> fileset=<fileset-name> level=<level-name>

Specification of the **job** is sufficient, but you can also override the client, fileset and/or level by specifying them on the estimate command line.

As an example, you might do:

```
@output /tmp/listing
estimate job=NightlySave listing level=Incremental
@output
```

which will do a full listing of all files to be backed up for the Job **NightlySave** during an Incremental save and put it in the file **/tmp/listing**.

help This command displays the list of commands available.

label This command is used to label physical volumes. The full form of this command is:

label storage=<storage-name> volume=<volume-name> slot=<slot>

If you leave out any part, you will be prompted for it. The media type is automatically taken from the Storage resource definition that you supply. Once the necessary information is obtained, the Console program contacts the specified Storage daemon and requests that the tape be labeled. If the tape labeling is successful, the Console program will create a Volume record in the appropriate Pool.

The Volume name is restricted to letters, numbers, and the special characters hyphen (**-**), underscore (**_**), colon (**:**), and period (**.**). All other characters including a space are illegal. This restriction is to ensure good readability of Volume names to reduce operator errors.

Please note, when labeling a blank tape, Bacula will get read I/O error when it attempts to ensure that the tape is already labeled. If you wish to avoid getting these messages, please write and EOF mark on your tape before attempting to label it:

```
mt rewind
mt weof
```

The label command can fail for a number of reasons:

1. The Volume name you specify is already in the Volume database.
2. The Storage daemon has a tape already mounted on the device, in which case you must **unmount** the device, insert a blank tape, then do the **label** command.
3. The tape in the device is already a Bacula labeled tape. (Bacula will never relabel a Bacula labeled tape unless it is recycled and you use the **relabel** command).
4. There is no tape in the drive.

There are two ways to relabel a volume that already has a Bacula label. The brute force method is to write an end of file mark on the tape using the system **mt** program, something like the following:

```
mt -f /dev/st0 rewind
mt -f /dev/st0 weof
```

Then you use the **label** command to add a new label. However, this could leave traces of the old volume in the catalog.

The preferable method to relabel a tape is to first **purge** the volume, either automatically, or explicitly with the **purge** command, then use the **relabel** command described below.

If your autochanger has barcode labels, you can label all the Volumes in your autochanger one after another by using the **label barcodes** command. For each tape in the changer containing a barcode, Bacula will mount the tape and then label it with the same name as the barcode. An appropriate Media record will also be created in the catalog. Any barcode that begins with the same characters as specified on the "CleaningPrefix=xxx" command, will be treated as a cleaning tape, and will not be labeled. For example with:

```
Pool {
  Name ...
  Cleaning Prefix = "CLN"
}
```

Any slot containing a barcode of CLNxxxx will be treated as a cleaning tape and will not be mounted. Note, the full form of the command is:

```
update storage=xxx pool=yyy slots=1-5,10 barcodes
```

list The list command lists the requested contents of the Catalog. The various fields of each record are listed on a single line. If there The various forms of the list command are:

list jobs

list jobid=<id>

list job=<job-name>

list jobmedia

list jobmedia jobid=<id>

list jobmedia job=<job-name>

list files jobid=<id>

list files job=<job-name>

list pools

list clients

list jobtotals

list volumes

list volumes jobid=<id>

list volumes pool=<pool-name>

list volumes job=<job-name>

list volume=<volume-name> list nextvolume job=<job-name>

list nextvol job=<job-name>

What most of the above commands do should be more or less obvious. In general if you do not specify all the command line arguments, the command will prompt you for what is needed.

The **list nextvol** command will print the Volume name to be used by the specified job. You should be aware that exactly what Volume will be used depends on a lot of factors including the time and what

a prior job will do. It may fill a tape that is not full when you issue this command. As a consequence, this command will give you a good estimate of what Volume will be used but not a definitive answer. In addition, this command may have certain side effect because it runs through the same algorithm as a job, which means it may automatically purge or recycle a Volume.

If you wish to add specialized commands that list the contents of the catalog, you can do so by adding them to the **query.sql** file. However, this takes some knowledge of programming SQL. Please see the **query** command below for additional information. See below for listing the full contents of a catalog record with the **llist** command.

As an example, the command **list pools** might produce the following output:

```
+------+---------+---------+---------+----------+-------------+
| PoId | Name    | NumVols | MaxVols | PoolType | LabelFormat |
+------+---------+---------+---------+----------+-------------+
|    1 | Default |       0 |       0 | Backup   | *           |
|    2 | Recycle |       0 |       8 | Backup   | File        |
+------+---------+---------+---------+----------+-------------+
```

As mentioned above, the **list** command lists what is in the database. Some things are put into the database immediately when Bacula starts up, but in general, most things are put in only when they are first used, which is the case for a Client as with Job records, etc.

Bacula should create a client record in the database the first time you run a job for that client. Doing a **status** will not cause a database record to be created. The client database record will be created whether or not job fails, but it must at least start. When the Client is actually contacted, additional info from the client will be added to the client record (a "uname -a" output).

If you want to see what Client resources you have available in your conf file, you use the Console command **show clients**.

llist The llist or "long list" command takes all the same arguments that the list command described above does. The difference is that the llist command list the full contents of each database record selected. It does so by listing the various fields of the record vertically, with one field per line. It is possible to produce a very large number of output lines with this command.

If instead of the **list pools** as in the example above, you enter **llist pools** you might get the following output:

```
          PoolId: 1
            Name: Default
         NumVols: 0
         MaxVols: 0
         UseOnce: 0
      UseCatalog: 1
AcceptAnyVolume: 1
   VolRetention: 1,296,000
 VolUseDuration: 86,400
     MaxVolJobs: 0
    MaxVolBytes: 0
      AutoPrune: 0
        Recycle: 1
       PoolType: Backup
    LabelFormat: *
         PoolId: 2
           Name: Recycle
        NumVols: 0
        MaxVols: 8
        UseOnce: 0
     UseCatalog: 1
AcceptAnyVolume: 1
   VolRetention: 3,600
 VolUseDuration: 3,600
     MaxVolJobs: 1
    MaxVolBytes: 0
      AutoPrune: 0
        Recycle: 1
       PoolType: Backup
    LabelFormat: File
```

messages This command causes any pending console messages to be immediately displayed.

mount The mount command is used to get Bacula to read a volume on a physical device. It is a way to tell Bacula that you have mounted a tape and that Bacula should examine the tape. This command is used only after there was no Volume in a drive and Bacula requests you to mount a new Volume or when you have specifically unmounted a Volume with the **unmount** console command, which causes Bacula to close the drive. If you have an autoloader, the mount command will not cause Bacula to operate the autoloader. The various forms of the mount command are:

mount storage=<storage-name>

mount [jobid=<id> — job=<job-name>]

If you have specified **Automatic Mount = yes** in the Storage daemon's Device resource, under most circumstances, Bacula

will automatically access the Volume unless you have explicitly **unmount**ed it in the Console program.

prune The Prune command allows you to safely remove expired database records from Jobs and Volumes. This command works only on the Catalog database and does not affect data written to Volumes. In all cases, the Prune command applies a retention period to the specified records. You can Prune expired File entries from Job records; you can Prune expired Job records from the database, and you can Prune both expired Job and File records from specified Volumes.

prune files—jobs—volume client=<client-name> volume=<volume-name>

For a Volume to be pruned, the **VolStatus** must be Full, Used, or Append, otherwise the pruning will not take place.

purge The Purge command will delete associated Catalog database records from Jobs and Volumes without considering the retention period. **Purge** works only on the Catalog database and does not affect data written to Volumes. This command can be dangerous because you can delete catalog records associated with current backups of files, and we recommend that you do not use it unless you know what you are doing. The permitted forms of **purge** are: purge files jobid=<jobid>—job=<job-name>—client=<client-name>

purge jobs client=<client-name> (of all jobs)

purge volume—volume=<vol-name> (of all jobs)

For the **purge** command to work on Volume Catalog database records the **VolStatus** must be Append, Full, Used, or Error.

The actual data written to the Volume will be unaffected by this command.

relabel This command is used to label physical volumes. The full form of this command is:

relabel storage=<storage-name> volume=<newvolume-name> name=<old-volume-name>

If you leave out any part, you will be prompted for it. In order for the Volume (old-volume-name) to be relabeled, it must be in the catalog, and the volume status must be marked **Purged** or **Recycle**. This happens automatically as a result of applying retention periods, or you may explicitly purge the volume using the **purge** command.

Once the volume is physically relabeled, the old data written on the Volume is lost and cannot be recovered.

release This command is used to cause the Storage daemon to rewind (release) the current tape in the drive, and to re-read the Volume label the next time the tape is used.

release storage=<storage-name>

After a release command, the device is still kept open by Bacula (unless Always Open is set to No in the Storage Daemon's configuration) so it cannot be used by another program. However, with some tape drives, the operator can remove the current tape and to insert a different one, and when the next Job starts, Bacula will know to re-read the tape label to find out what tape is mounted. If you want to be able to use the drive with another program (e.g. **mt**), you must use the **unmount** command to cause Bacula to completely release (close) the device.

restore The restore command allows you to select one or more Jobs (JobIds) to be restored using various methods. Once the JobIds are selected, the File records for those Jobs are placed in an internal Bacula directory tree, and the restore enters a file selection mode that allows you to interactively walk up and down the file tree selecting individual files to be restored. This mode is somewhat similar to the standard Unix **restore** program's interactive file selection mode.

restore storage=<storage-name> client=<client-name> where=<path> pool=<pool-name> fileset=<fileset-name> select current all done

Where **current**, if specified, tells the restore command to automatically select a restore to the most current backup. If not specified, you will be prompted. The **all** specification tells the restore command to restore all files. If it is not specified, you will be prompted for the files to restore. For details of the **restore** command, please see the Restore Chapter of this manual.

run This command allows you to schedule jobs to be run immediately. The full form of the command is:

run job=<job-name> client=<client-name> fileset=<FileSet-name> level=<level-keyword> storage=<storage-name> where=<directory-prefix> when=<universal-time-specification> yes

Any information that is needed but not specified will be listed for selection, and before starting the job, you will be prompted to accept, reject, or modify the parameters of the job to be run, unless you have specified **yes**, in which case the job will be immediately sent to the scheduler.

On my system, when I enter a run command, I get the following prompt:

```
A job name must be specified.
The defined Job resources are:
     1: Matou
     2: Polymatou
     3: Rufus
     4: Minimatou
     5: Minou
     6: PmatouVerify
     7: MatouVerify
     8: RufusVerify
     9: Watchdog
Select Job resource (1-9):
```

If I then select number 5, I am prompted with:

```
Run Backup job
JobName:  Minou
FileSet:  Minou Full Set
Level:    Incremental
Client:   Minou
Storage:  DLTDrive
Pool:     Default
When:     2003-04-23 17:08:18
OK to run? (yes/mod/no):
```

If I now enter **yes**, the Job will be run. If I enter **mod**, I will be presented with the following prompt.

```
Parameters to modify:
     1: Level
     2: Storage
     3: Job
     4: FileSet
     5: Client
     6: When
     7: Pool
Select parameter to modify (1-7):
```

If you wish to start a job at a later time, you can do so by setting the When time. Use the **mod** option and select **When** (no. 6). Then enter the desired start time in YYYY-MM-DD HH:MM:SS format.

setdebug This command is used to set the debug level in each daemon. The form of this command is:

setdebug level=nn [trace=0/1 client=<client-name> — dir — director — storage=<storage-name> — all]

If trace=1 is set, then the tracing will be enabled, and the daemon where the setdebug applies will be placed in trace mode, and all debug output will go to the file **bacula.trace** in the current directory of the daemon. Normally, tracing is used only for Win32 clients where the debug output cannot be written to a terminal or redirected to a file. When tracing, each debug output message is appended to the trace file. You must explicitly delete the file when you are done.

show The show command will list the Director's resource records as defined in the Director's configuration file (normally **bacula-dir.conf**). This command is used mainly for debugging purposes by developers. The following keywords are accepted on the show command line: directors, clients, counters, jobs, storages, catalogs, schedules, filesets, groups, pools, messages, all, help. Please don't confuse this command with the **list**, which displays the contents of the catalog.

sqlquery The sqlquery command puts the Console program into SQL query mode where each line you enter is concatenated to the previous line until a semicolon (;) is seen. The semicolon terminates the command, which is then passed directly to the SQL database engine. When the output from the SQL engine is displayed, the formation of a new SQL command begins. To terminate SQL query mode and return to the Console command prompt, you enter a period (.) in column 1.

Using this command, you can query the SQL catalog database directly. Note you should really know what you are doing otherwise you could damage the catalog database. See the **query** command below for simpler and safer way of entering SQL queries.

Depending on what database engine you are using (MySQL or SQLite), you will have somewhat different SQL commands available. For more detailed information, please refer to the MySQL or SQLite documentation.

status This command will display the status of the next jobs that are scheduled during the next twenty-four hours as well as the status of currently running jobs. The full form of this command is:

status [all — dir=<dir-name> — director — client=<client-name> — storage=<storage-name>]

If you do a **status dir**, the console will list any currently running jobs, a summary of all jobs scheduled to be run in the next 24 hours, and a listing of the last 10 terminated jobs with their statuses. The scheduled jobs summary will include the Volume name to be used. You should be aware of two things: 1. to obtain the volume name, the code goes through the same code that will be used when the job runs, which means that it may prune or recycle a Volume; 2. The Volume

listed is only a best guess. The Volume actually used may be different because of the time difference (more durations may expire when the job runs) and another job could completely fill the Volume requiring a new one.

In the Running Jobs listing, you may find the following types of information:

```
2507 Catalog MatouVerify.2004-03-13_05.05.02 is waiting execution
5349 Full    CatalogBackup.2004-03-13_01.10.00 is waiting for higher
             priority jobs to finish
5348 Differe Minou.2004-03-13_01.05.09 is waiting on max Storage jobs
5343 Full    Rufus.2004-03-13_01.05.04 is running
```

Looking at the above listing from bottom to top, obviously JobId 5343 (Rufus) is running. JobId 5348 (Minou) is waiting for JobId 5343 to finish because it is using the Storage resource, hence the "waiting on max Storage jobs". JobId 5349 has a lower priority than all the other jobs so it is waiting for higher priority jobs to finish, and finally, JobId 2508 (MatouVerify) is waiting because only one job can run at a time, hence it is simply "waiting execution&quot.¡/dd¿

unmount This command causes the indicated Bacula Storage daemon to unmount the specified device. The forms of the command are the same as the mount command:

unmount storage=<storage-name>

unmount [jobid=<id> — job=<job-name>]

update This command will update catalog for either a specific Pool record, a Volume record, or the Slots in an autochanger with barcode capability. In the case of updating a Pool record, the new information will be automatically taken from the corresponding Director's configuration resource record. It can be used to increase the maximum number of volumes permitted or to set a maximum number of volumes. The following main keywords may be specified:

media, volume, pool, slots

In the case of updating a Volume, you will be prompted for which value you wish to change. The following Volume parameters may be changed:

```
Volume Status
Volume Retention Period
Volume Use Duration
Maximum Volume Jobs
```

```
Maximum Volume Files
Maximum Volume Bytes
Recycle Flag
Slot
InChanger Flag
Pool
Volume Files
Volume from Pool
All Volumes from Pool
```

For slots **update slots**, Bacula will obtain a list of slots and their barcodes from the Storage daemon, and for each barcode found, it will automatically update the slot in the catalog Media record to correspond to the new value. This is very useful if you have moved cassettes in the magazine, or if you have removed the magazine and inserted a different one. As the slot of each Volume is updated, the InChanger flag for that Volume will also be set, and any other Volumes in the Pool will have their InChanger flag turned off. This permits Bacula to know what magazine (tape holder) is currently in the autochanger.

If you do not have barcodes, you can accomplish the same thing in version 1.33 and later by using the **update slots scan** command. The **scan** keyword tells Bacula to physically mount each tape and to read its VolumeName.

For Pool **update pool**, Bacula will move the Volume record from its existing poole to the pool specified.

For **Volume from Pool** and **All Volumes from Pool**, the following values are updated from the Pool record: Recycle, VolRetention, VolUseDuration, MaxVolJobs, MaxVolFiles, and MaxVolBytes.

The full form of the update command with all command line arguments is:

```
update volume=xxx pool=yyy slots volstatus=xxx VolRetention=ddd
  VolUse=ddd MaxVolJobs=nnn MaxVolBytes=nnn Recycle=yes|no
  slot=nnn
```

use This command allows you to specify which Catalog database to use. Normally, you will be using only one database so this will be done automatically. In the case that you are using more than one database, you can use this command to switch from one to another.

use <database-name>

var This command takes a string or quoted string and does variable expansion on it the same way variable expansion is done on the **LabelFormat** string. Thus, for the most part, you can test your LabelFormat strings. The difference between the **var** command and the actual LabelFormat process is that during the var command, no job is running so "dummy" values are used in place of Job specific variables. Generally, however, you will get a good idea of what is going to happen in the real case.

version The command prints the Director's version.

quit This command terminates the console program. The console program sends the **quit** request to the Director and waits for acknowledgment. If the Director is busy doing a previous command for you that has not terminated, it may take some time. You may quit immediately by issuing the **.quit** command (i.e. quit preceded by a period).

query This command reads a predefined SQL query from the query file (the name and location of the query file is defined with the QueryFile resource record in the Director's configuration file). You are prompted to select a query from the file, and possibly enter one or more parameters, then the command is submitted to the Catalog database SQL engine.

The following queries are currently available (version 1.24):

```
Available queries:
  1: List Job totals:
  2: List where a file is saved:
  3: List where the most recent copies of a file are saved:
  4: List total files/bytes by Job:
  5: List total files/bytes by Volume:
  6: List last 20 Full Backups for a Client:
  7: List Volumes used by selected JobId:
  8: List Volumes to Restore All Files:
  9: List where a File is saved:
Choose a query (1-9):
```

exit This command terminates the console program.

wait The wait command causes the Director to pause until there are no jobs running. This command is useful in a batch situation such as regression testing where you wish to start a job and wait until that job completes before continuing.

Special dot Commands

There is a list of commands that are prefixed with a period (.). These commands are intended to be used either by batch programs or graphical user interface front-ends. They are not normally used by interactive users. Once GUI development begins, this list will be considerably expanded. The following is the list of dot commands:

```
.die         cause the Director to segment fault (for debugging)
.jobs        list all job names
.filesets    list all fileset names
.clients     list all client names
.msgs        return any queued messages
.quit        quit
.exit        quit
```

Special At (@) Commands

Normally, all commands entered to the Console program are immediately forwarded to the Director, which may be on another machine, to be executed. However, there is a small list of **at** commands, all beginning with a at character (@), that will not be sent to the Director, but rather interpreted by the Console program directly. Note, these commands are implemented only in the tty console program and not in the GNOME Console. These commands are:

@input <filename> Read and execute the commands contained in the file specified.

@output <filename> w/a Send all following output to the filename specified either overwriting the file (w) or appending to the file (a). To redirect the output to the terminal, simply enter **@output** without a filename specification. WARNING: be careful not to overwrite a valid file. A typical example during a regression test might be:

```
@output /dev/null
commands ...
@output
```

@tee <filename> w/a Send all subsequent output to both the specified file and the terminal. It is turned off by specifying **@tee** or **@output** without a filename.

@sleep <**seconds**> Sleep the specified number of seconds.

@time Print the current time and date.

@version Print the console's version.

@quit quit

@exit quit

@# anything Comment

Running the Console Program from a Shell Script

You can automate many Console tasks by running the console program from a shell script. For example, if you have created a file containing the following commands:

```
./bconsole -c ./bconsole.conf <<END_OF_DATA
unmount storage=DDS-4
quit
END_OF_DATA
```

when that file is executed, it will unmount the current DDS-4 storage device. You might want to run this command during a Job by using the **RunBeforeJob** or **RunAfterJob** records.

It is also possible to run the Console program from file input where the file contains the commands as follows:

```
./bconsole -c ./bconsole.conf <filename
```

where the file named **filename** contains any set of console commands.

As a real example, the following script is part of the Bacula regression tests. It labels a volume (a disk volume), runs a backup, then does a restore of the files saved.

```
bin/bconsole -c bin/bconsole.conf <<END_OF_DATA
@output /dev/null
messages
@output /tmp/log1.out
label volume=TestVolume001
```

```
run job=Client1 yes
wait
messages
@#
@# now do a restore
@#
@output /tmp/log2.out
restore current all
yes
wait
messages
@output
quit
END_OF_DATA
```

The output from the backup is directed to /tmp/log1.out and the output from the restore is directed to /tmp/log2.out. To ensure that the backup and restore ran correctly, the output files are checked with:

```
grep "^Termination: *Backup OK" /tmp/log1.out
backupstat=$?
grep "^Termination: *Restore OK" /tmp/log2.out
restorestat=$?
```

Adding Volumes to a Pool

If you have used the **label** command to label a Volume, it will be automatically added to the Pool, and you will not need to add any media to the pool.

Alternatively, you may choose to add a number of Volumes to the pool without labeling them. At a later time when the Volume is requested by **Bacula** you will need to label it.

Before adding a volume, you must know the following information:

1. The name of the Pool (normally "Default")
2. The Media Type as specified in the Storage Resource in the Director's configuration file (e.g. "DLT8000")
3. The number and names of the Volumes you wish to create.

For example, to add media to a Pool, you would issue the following commands to the console program:

```
*add
Enter name of Pool to add Volumes to: Default
Enter the Media Type: DLT8000
Enter number of Media volumes to create. Max=1000: 10
Enter base volume name: Save
Enter the starting number: 1
10 Volumes created in pool Default
*
```

To see what you have added, enter:

```
*list media pool=Default
+-------+----------+---------+---------+-------+------------------+
| MedId | VolumeNa | MediaTyp| VolStat | Bytes | LastWritten      |
+-------+----------+---------+---------+-------+------------------+
|    11 | Save0001 | DLT8000 | Append  |     0 | 0000-00-00 00:00 |
|    12 | Save0002 | DLT8000 | Append  |     0 | 0000-00-00 00:00 |
|    13 | Save0003 | DLT8000 | Append  |     0 | 0000-00-00 00:00 |
|    14 | Save0004 | DLT8000 | Append  |     0 | 0000-00-00 00:00 |
|    15 | Save0005 | DLT8000 | Append  |     0 | 0000-00-00 00:00 |
|    16 | Save0006 | DLT8000 | Append  |     0 | 0000-00-00 00:00 |
|    17 | Save0007 | DLT8000 | Append  |     0 | 0000-00-00 00:00 |
|    18 | Save0008 | DLT8000 | Append  |     0 | 0000-00-00 00:00 |
|    19 | Save0009 | DLT8000 | Append  |     0 | 0000-00-00 00:00 |
|    20 | Save0010 | DLT8000 | Append  |     0 | 0000-00-00 00:00 |
+-------+----------+---------+---------+-------+------------------+
*
```

Notice that the console program automatically appended a number to the base Volume name that you specify (Save in this case). If you don't want it to append a number, you can simply answer 0 (zero) to the question "Enter number of Media volumes to create. Max=1000:", and in this case, it will create a single Volume with the exact name you specify.

The Bacula Console Restore Command

General

Below, we will discuss restoring files with the Console **Restore** command, which is the recommended way of doing it. However, there is a standalone program named **bextract**, which also permits restoring files. For more information on this program, please see the Bacula Utility Programs chapter of this manual. You will also want to look at the **bls** program in the same chapter, which allows you to list the contents of your Volumes. Finally, if you have an old Volume that is no longer in the catalog, you can restore the catalog entries using the program named **bscan**, documented in the same Bacula Utility Programs chapter.

In general, to restore a file or a set of files, you must run a **restore** job. That is a job with **Type = Restore**. As a consequence, you will need a predefined **restore** job in your **bacula-dir.conf** (Director's config) file. The exact parameters (Client, FileSet, ...) that you define are not important as you can either modify them manually before running the job or if you use the **restore** command, explained below, they will be automatically set for you.

Since Bacula is a network backup program, you must be aware that when you restore files, it is up to you to ensure that you or Bacula have selected the correct Client and the correct hard disk location for restoring those files. **Bacula** will quite willingly backup client A, and restore it by sending the files to a different directory on client B. Normally, you will want to avoid this, but assuming the operating systems are not too different in their file structures, this should work perfectly well, if so desired.

The Restore Command

Since Bacula maintains a catalog of your files and on which Volumes (disk or tape), they are stored, it can do most of the bookkeeping work, allowing you simply to specify what kind of restore you want (current, before a particular date), and what files to restore. Bacula will then do the rest.

This is accomplished using the **restore** command in the Console. First you select the kind of restore you want, then Bacula Once the JobIds are selected, the File records for those Jobs are placed in an internal Bacula directory tree, and the restore enters a file selection mode that allows you to interactively walk up and down the file tree selecting individual files to

be restored. This mode is somewhat similar to the standard Unix **restore** program's interactive file selection mode.

Within the Console program, after entering the **restore** command, you are presented with the following selection prompt:

```
First you select one or more JobIds that contain files
to be restored. You will be presented several methods
of specifying the JobIds. Then you will be allowed to
select which files from those JobIds are to be restored.
To select the JobIds, you have the following choices:
     1: List last 20 Jobs run
     2: List Jobs where a given File is saved
     3: Enter list of JobIds to select
     4: Enter SQL list command
     5: Select the most recent backup for a client
     6: Select backup for a client before a specified time
     7: Enter a list of files to restore
     8: Enter a list of files to restore before a specified time
     9: Cancel
Select item:  (1-9):
```

- Item 1 will list the last 20 jobs run. If you find the Job you want, you can then select item 3 and enter its JobId(s).
- Item 2 will list all the Jobs where a specified file is saved. If you find the Job you want, you can then select item 3 and enter the JobId.
- Item 3 allows you the enter a list of comma separated JobIds whose files will be put into the directory tree.
- Item 4 allows you to enter any arbitrary SQL command. This is probably the most primitive way of finding the desired JobIds, but at the same time, the most flexible. Once you have found the JobId(s), you can select item 3 and enter them.
- Item 5 will automatically select the most recent Full backup and all subsequent incremental and differential backups for a specified Client. These are the Jobs and Files which if reloaded will restore your system to the most current saved state. It automatically enters the JobIds found into the directory tree. This is probably the most convenient of all the above options to use if you wish to restore a selected Client to its most recent state.
- Item 6 allows you to specify a date and time then Bacula will automatically select the most recent Full backup and all subsequent incremental and differential backups that started before the specified date and time.

- Item 7 allows you to specify one or more filenames (complete path required) to be restored. Each filename is entered one at a time or if you prefix a filename with the less-than symbol (<) Bacula will read that file and assume it is a list of filenames to be restored. The filename entry mode is terminated by entering a blank line.
- Item 8 allows you to specify a date and time before entering the filenames. See Item 7 above for more details.
- Item 9 allows you to cancel the restore command.

As an example, suppose that we select item 5 (restore to most recent state). It will then ask for the desired Client, which on my system, will print all the Clients found in the database as follows:

```
Defined clients:
     1: Rufus
     2: Matou
     3: Polymatou
     4: Minimatou
     5: Minou
     6: MatouVerify
     7: PmatouVerify
     8: RufusVerify
     9: Watchdog
Select Client (File daemon) resource (1-9):
```

You will probably have fare fewer Clients than this example, and if you have only one Client, it will be automatically selected, but in this case, I enter **Rufus** to select the Client. Then Bacula needs to know what FileSet is to be restored, so it prompts with:

```
The defined FileSet resources are:
     1: Full Set
     2: Kerns Files
Select FileSet resource (1-2):
```

I choose item 1, which is my full backup. Normally, you will only have a single FileSet for each Job, and if your machines are similar (all Linux) you may only have one FileSet for all your Clients.

At this point, **Bacula** has all the information it needs to find the most recent set of backups. It will then query the database, which may take a bit of time, and it will come up with something like the following. Note, some of the columns are truncated here for presentation:

```
+-------+------+----------+-------------+-------------+------+-------+------------+
| JobId | Levl | JobFiles | StartTime   | VolumeName  | File | SesId | VolSesTime |
+-------+------+----------+-------------+-------------+------+-------+------------+
| 1,792 | F    |  128,374 | 08-03 01:58 | DLT-19Jul02 |   67 |    18 | 1028042998 |
| 1,792 | F    |  128,374 | 08-03 01:58 | DLT-04Aug02 |    0 |    18 | 1028042998 |
| 1,797 | I    |      254 | 08-04 13:53 | DLT-04Aug02 |    5 |    23 | 1028042998 |
| 1,798 | I    |       15 | 08-05 01:05 | DLT-04Aug02 |    6 |    24 | 1028042998 |
+-------+------+----------+-------------+-------------+------+-------+------------+
You have selected the following JobId: 1792,1792,1797
Building directory tree for JobId 1792 ...
Building directory tree for JobId 1797 ...
Building directory tree for JobId 1798 ...
cwd is: /
$
```

Depending on the number of **JobFiles** for each JobId, the **Building directory tree ...“** can take a bit of time.

In our example, Bacula found four Jobs that comprise the most recent backup of the specified Client and FileSet. Two of the Jobs have the same JobId because that Job wrote on two different Volumes. The third Job was an incremental backup to the previous Full backup, and it only saved 254 Files compared to 128,374 for the Full backup. The fourth Job was also an incremental backup that saved 15 files.

Next Bacula entered those Jobs into the directory tree, with no files marked to be restored as a default, tells you how many files are in the tree, and tells you what the current working directory (**cwd**) is /. Finally, Bacula prompts with the dollar sign ($) to indicate that you may enter commands to move around the directory tree and to select files.

Instead of choosing item 5 on the first menu (Select the most recent backup for a client), if we had chosen item 3 (Enter list of JobIds to select) and we had entered the JobIds **1792,1797,1798** we would have arrived at the same point.

One point to note if you are manually entering JobIds is that you must enter them in the order they were run (generally in increasing JobId order). If you enter them out of order and the same file was saved in two or more of the Jobs, you may end up with an old version of that file (i.e. not the most recent).

While in file selection mode, you can enter **help** or a question mark (?) to produce a summary of the available commands:

```
  Command    Description
  =======    ===========
```

```
cd          change current directory
count       count marked files in and below the cd
dir         list current directory
done        leave file selection mode
estimate    estimate restore size
exit        exit = done
find        find files -- wildcards allowed
help        print help
ls          list current directory -- wildcards allowed
lsmark      list the marked files in and below the cd
mark        mark file to be restored
markdir     mark directory entry to be restored -- nonrecursive
pwd         print current working directory
unmark      unmark file to be restored
unmarkdir   unmark directory -- no recursion
quit        quit
?           print help
```

As a default no files have been selected for restore. If you want to restore everything, at this point, you should enter **mark ***, and then **done** and **Bacula** will write the bootstrap records to a file and request your approval to start a restore job.

If you do not enter the above mentioned **mark *** command, you will start with an empty slate. Now you can simply start looking at the tree and **mark** particular files or directories if you want restored. It is easy to make a mistake in specifying a file to mark or unmark, and Bacula's error handling is not perfect, so please check your work by using the **ls** or **dir** commands to see what files are actually selected. Any selected file has its name preceded by an asterisk.

To check what is marked or not marked, enter the **count** command, which displays:

```
128401 total files. 128401 marked to be restored.
```

Each of the above commands will be described in more detail in the next section. We continue with the above example, having accepted to restore all files as Bacula set by default. On entering the **done** command, Bacula prints:

```
Bootstrap records written to /home/kern/bacula/working/restore.bsr
The restore job will require the following Volumes:
```

```
   DLT-19Jul02
   DLT-04Aug02
128401 files selected to restore.
Run Restore job
JobName:    kernsrestore
Bootstrap:  /home/kern/bacula/working/restore.bsr
Where:      /tmp/bacula-restores
Replace:    always
FileSet:    Kerns Files
Client:     Rufus
Storage:    SDT-10000
JobId:      *None*
OK to run? (yes/mod/no):
```

Please examine each of the items very carefully to make sure that they are correct. In particular, look at **Where**, which tells you where in the directory structure the files will be restored, and **Client**, which tells you which client will receive the files. These items will not always be completed with the correct values depending on which of the restore options you chose.

The above assumes that you have defined a **Restore** Job resource in your Director's configuration file. Normally, you will only need one Restore Job resource definition because by its nature, restoring is a manual operation, and using the Console interface, you will be able to modify the Restore Job to do what you want.

An example Restore Job resource definition is given below.

Returning to the above example, you should verify that the Client name is correct before running the Job. However, you may want to modify some of the parameters of the restore job. For example, in addition to checking the Client it is wise to check that the Storage device chosen by Bacula is indeed correct. Although the **FileSet** is shown, it will be ignored in restore. The restore will choose the files to be restored either by reading the **Bootstrap** file, or if not specified, it will restore all files associated with the specified backup **JobId** (i.e. the JobId of the Job that originally backed up the files).

Finally before running the job, please note that the default location for restoring files is **not** their original locations, rather the directory **/tmp/bacula-restores**. You can change this default by modifying your **bacula-dir.conf** file, or you can modify it using the **mod** option. If you want to restore the files to their original location, you must have **Where** set to nothing or to the root, i.e. /.

If you now enter **yes**, Bacula will run the restore Job. The Storage daemon will first request Volume **DLT-19Jul02** and after the appropriate files have

been restored from that volume, it will request Volume **DLT-04Aug02**.

Selecting Files by Filename

If you have a small number of files to restore, and you know the filenames, you can either put the list of filenames in a file to be read by Bacula, or you can enter the names one at a time. The filenames must include the full path and filename. No wild cards are used.

To enter the files, after the **restore**, you select item number 7 from the prompt list:

```
To select the JobIds, you have the following choices:
     1: List last 20 Jobs run
     2: List Jobs where a given File is saved
     3: Enter list of JobIds to select
     4: Enter SQL list command
     5: Select the most recent backup for a client
     6: Select backup for a client before a specified time
     7: Enter a list of files to restore
     8: Enter a list of files to restore before a specified time
     9: Cancel
Select item:  (1-9): 7
```

which then prompts you with for the client name:

```
Defined Clients:
     1: Timmy
     2: Tibs
     3: Rufus
Select the Client (1-3): 3
```

Of course, your client list will be different, and if you have only one client, it will be automatically selected. And finally, Bacula requests you to enter a filename:

```
Enter filename:
```

At this point, you can enter the full path and filename

```
Enter filename: /home/kern/bacula/k/Makefile.in
Enter filename:
```

as you can see, it took the filename. If Bacula cannot find a copy of the file, it prints the following:

```
Enter filename: junk filename
No database record found for: junk filename
Enter filename:
```

If you want Bacula to read the filenames from a file, you simply precede the filename with a less-than symbol (<). When you have entered all the filenames, you enter a blank line, and Bacula will write the bootstrap file, tell you what tapes will be used, and propose a Restore job to be run:

```
Enter filename:
Automatically selected Storage: DDS-4
Bootstrap records written to /home/kern/bacula/working/restore.bsr
The restore job will require the following Volumes:

   test1
1 file selected to restore.
Run Restore job
JobName:    kernsrestore
Bootstrap:  /home/kern/bacula/working/restore.bsr
Where:      /tmp/bacula-restores
Replace:    always
FileSet:    Kerns Files
Client:     Rufus
Storage:    DDS-4
When:       2003-09-11 10:20:53
Priority:   10
OK to run? (yes/mod/no):
```

It is possible to automate the selection by file by putting your list of files in say **/tmp/file-list**, then using the following command:

```
restore client=Rufus file=</tmp/file-list
```

If in modifying the parameters for the Run Restore job, you find that Bacula asks you to enter a Job number, this is because you have no yet specified either a Job number or a Bootstrap file. Simply entering zero will allow you to continue and to select another option to be modified.

Command Line Arguments

If all the above sounds complicated, you will probably agree that it really isn't after trying it a few times. It is possible to do everything that was

shown above, with the exception of selecting the FileSet, by using command line arguments with a single command by entering:

```
restore client=Rufus select current all done yes
```

The **client=Rufus** specification will automatically select Rufus as the client, the **current** tells Bacula that you want to restore the system to the most current state possible, and the **yes** suppresses the final **yes/mod/no** prompt and simply runs the restore.

The full list of possible command line arguments are:

- **all** – select all Files to be restored.
- **select** – use the tree selection method.
- **done** – do not prompt the user in tree mode.
- **current** – automatically select the most current set of backups for the specified client.
- **client=xxxx** – select the specified client.
- **jobid=nnn** – specify a JobId or comma separated list of JobIds to be restored.
- **before=YYYY-MM-DD HH:MM:SS** – specify a date and time to which the system should be restored. Only Jobs started before the specified date/time will be selected, and as is the case for **current** Bacula will automatically find the most recent prior Full save and all Differential and Incremental saves run before the date you specify. Note, this command is not too user friendly in that you must specify the date/time exactly as shown.
- **file=filename** – specify a filename to be restored. You must specify the full path and filename. Prefixing the entry with a less-than sign (<) will cause Bacula to assume that the filename is on your system and contains a list of files to be restored. Bacula will thus read the list from that file. Multiple file=xxx specifications may be specified on the command line.
- **jobid=nnn** – specify a JobId to be restored.
- **pool=pool-name** – specify a Pool name to be used for selection of Volumes when specifying options 5 and 6 (restore current system, and restore current system before given date). This permits you to have

several Pools, possibly one offsite, and to select the Pool to be used for restoring.

- **yes** – automatically run the restore without prompting for modifications (most useful in batch scripts).

Restoring Directory Attributes

Depending how you do the restore, you may or may not get the directory entries back to their original state. Here are a few of the problems you can encounter, and for same machine restores, how to avoid them.

- You backed up on one machine and are restoring to another that is either a different OS or doesn't have the same users/groups defined. Bacula does the best it can in these situations.
- You are restoring into a directory that is already created and has file creation restrictions. Bacula tries to reset everything but without walking up the full chain of directories and modifying them all during the restore, which Bacula does and will not do, getting permissions back correctly in this situation depends to a large extent on your OS.
- You selected one or more files in a directory, but did not select the directory entry to be restored. In that case, if the directory is not on disk Bacula simply creates the directory with some default attributes which may not be the same as the original. If you do not select a directory and all its contents to be restored, you can still select items within the directory to be restored by individually marking those files, but in that case, you should individually use the "markdir" command to select all higher level directory entries (one at a time) to be restored if you want the directory entries properly restored.

Restoring on Windows

If you are restoring on WinNT/2K/XP systems, Bacula will restore the files with the original ownerships and permissions as would be expected. This is also true if you are restoring those files to an alternate directory (using the Where option in restore). However, if the alternate directory does not already exist, the Bacula File daemon (Client) will create it, and since the File daemon runs under the SYSTEM account, the directory will be created with SYSTEM ownership and permissions. In this case, you may have problems accessing the newly restored files.

To avoid this problem, you can create the alternate directory before doing the restore. Bacula will not change the ownership and permissions of the directory if it is already created as long as it is not one of the directories being restored (i.e. written to tape).

Restoring Files Can Be Slow

Restoring files is generally **much** slower than backing it up for several reasons. The first is that during a backup the tape is normally already positioned and Bacula need only write. On the other hand, because restoring files is done so rarely, Bacula keeps only the he start file and block on the tape for the whole job rather than on a file by file basis which would use quite a lot of space in the catalog.

Bacula versions 1.31a and older would seek to the first file on the first tape, then sequentially search the tape for the specified files. If you were doing a full restore, this is OK, but if you want to restore one or two files, the process could be quite long.

This deficiency has been corrected in version 1.32. The consequence is that Bacula will forward space to the correct file mark on the tape for the Job, then forward space to the correct block, and finally sequentially read each record until it gets to the correct one(s) for the file or files you want to restore. Once the desired files are restored, Bacula will stop reading the tape. For restoring a small number of files, version 1.32 and greater are hundreds of times faster than previous versions.

Finally, instead of just reading a file for backup, during the restore, Bacula must create the file, and the operating system must allocate disk space for the file as Bacula is restoring it.

For all the above reasons the restore process is generally much slower than backing up.

Problems Restoring Files

The most frequent problems users have restoring files are error messages such as:

```
04-Jan 00:33 z217-sd: RestoreFiles.2005-01-04_00.31.04 Error:
block.c:868 Volume data error at 20:0! Short block of 512 bytes on
device /dev/tape discarded.
```

or

```
04-Jan 00:33 z217-sd: RestoreFiles.2005-01-04_00.31.04 Error:
block.c:264 Volume data error at 20:0! Wanted ID: "BB02", got ".".
Buffer discarded.
```

Both these kinds of messages indicate that you were probably running your tape drive in fixed block mode rather than variable block mode. Fixed block mode will work with any program that reads tapes sequentially such as tar, but Bacula repositions the tape on a block basis when restoring files because this will speed up the restore by orders of magnitude when only a few files are restore. There are several ways that you can attempt to recover from this unfortunate situation.

Try the following things each separately, and reset your Device resource to what it is now after each individual test:

1. Set "Block Positioning = no“ in your Device resource and try the restore. This is a new directive and untested.
2. Set "Minimum Block Size = 512“ and "Maximum Block Size = 512“ and try the restore. Again send me the full job report output. If you are able to determine the block size your drive was previously using, you should try that size if 512 does not work.
3. Try editing the restore.bsr file at the Run xxx yes/mod/no prompt before starting the restore job and remove all the VolBlock statements. These are what causes Bacula to reposition the tape, and where problems occur if you have a fixed block size set for your drive. The VolFile commands also cause repositioning, but this will work regardless of the block size.
4. Use bextract to extract the files you want – it reads the Volume sequentially if you use the include list feature, or if you use a .bsr file, but remove all the VolBlock statements after the .bsr file is created (at the Run yes/mod/no) prompt but before you start the restore.

Example Restore Job Resource

```
Job {
  Name = "RestoreFiles"
  Type = Restore
  Client = Any-client
  FileSet = "Any-FileSet"
```

```
  Storage = Any-storage
  Where = /tmp/bacula-restores
  Messages = Standard
  Pool = Default
}
```

If **Where** is not specified, the default location for restoring files will be their original locations.

File Selection Commands

After you have selected the Jobs to be restored and Bacula has created the in-memory directory tree, you will enter file selection mode as indicated by the dollar sign (**$**) prompt. While in this mode, you may use the commands listed above. The basic idea is to move up and down the in memory directory structure with the **cd** command much as you normally do on the system. Once you are in a directory, you may select the files that you want restored. As a default no files are marked to be restored. If you wish to start with all files, simply enter: **cd /** and **mark ***. Otherwise proceed to select the files you wish to restore by marking them with the **mark** command. The available commands are:

cd The **cd** command changes the current directory to the argument specified. It operates much like the Unix **cd** command. Wildcard specifications are not permitted.

Note, on Windows systems, the various drives (c:, d:, ...) are treated like a directory within the file tree while in the file selection mode. As a consequence, you must do a **cd c:** or possibly in some cases a **cd C:** (note upper case) to get down to the first directory.

dir The **dir** command is similar to the **ls** command, except that it prints it in long format (all details). This command can be a bit slower than the **ls** command because it must access the catalog database for the detailed information for each file.

estimate The **estimate** command prints a summary of the total files in the tree, how many are marked to be restored, and an estimate of the number of bytes to be restored. This can be useful if you are short on disk space on the machine where the files will be restored.

find The **find** command accepts one or more arguments and displays all files in the tree that match that argument. The argument may have wildcards. It is somewhat similar to the Unix command **find / -name arg**.

ls The **ls** command produces a listing of all the files contained in the current directory much like the Unix **ls** command. You may specify an argument containing wildcards, in which case only those files will be listed. Any file that is marked to be restored will have its name preceded by an asterisk (*). Directory names will be terminated with a forward slash (/) to distinguish them from filenames.

lsmark The **lsmark** command is the same as the **ls** except that it will print only those files marked for extraction. The other distinction is that it will recursively descend into any directory selected.

mark The **mark** command allows you to mark files to be restored. It takes a single argument which is the filename or directory name in the current directory to be marked for extraction. The argument may be a wildcard specification, in which case all files that match in the current directory are marked to be restored. If the argument matches a directory rather than a file, then the directory and all files contained in that directory (recursively) are marked to be restored. Any marked file will have its name preceded with an asterisk (*) in the output produced by the **ls** or **dir** commands. Note, supplying a full path on the mark command does not work as expected to select a file or directory in the current directory. Also, the **mark** command works on the current and lower directories but does not touch higher level directories.

After executing the **mark** command, it will print a brief summary:

```
No files marked.
```

If no files were marked, or:

```
nn files marked.
```

if some files are marked.

unmark The **unmark** is identical to the **mark** command, except that it unmarks the specified file or files so that they will not be restored. Note: the **unmark** command works from the current directory, so it does not unmark any files at a higher level. First do a **cd /** before the **unmark *** command if you want to unmark everything.

pwd The **pwd** command prints the current working directory. It accepts no arguments.

count The **count** command prints the total files in the directory tree and the number of files marked to be restored.

done This command terminates file selection mode.

exit This command terminates file selection mode (the same as done).

quit This command terminates the file selection and does not run the restore job.

help This command prints a summary of the commands available.

? This command is the same as the **help** command.

Catalog Maintenance

Without proper setup and maintenance, your Catalog may continue to grow indefinitely as you run Jobs and backup Files. How fast the size of your Catalog grows depends on the number of Jobs you run and how many files they backup. By deleting records within the database, you can make space available for the new records that will be added during the next Job. By constantly deleting old expired records (dates older than the Retention period), your database size will remain constant.

If you started with the default configuration files, they already contain reasonable defaults for a small number of machines (less that 5), so if you fall into that case, catalog maintenance will not be urgent if you have a few hundred megabytes of disk space free. Whatever the case may be, some knowledge of retention periods will be useful.

Setting Retention Periods

Bacula uses three Retention periods: the **File Retention** period, the **Job Retention** period, and the **Volume Retention** period. Of these three, the File Retention period is by far the most important in determining how large your database will become.

The **File Retention** and the **Job Retention** are specified in each Client resource as is shown below. The **Volume Retention** period is specified in the Pool resource, and the details are given in the next chapter of this manual.

File Retention = <time-period-specification> The File Retention record defines the length of time that Bacula will keep File records in the Catalog database. When this time period expires, and if **AutoPrune** is set to **yes**, Bacula will prune (remove) File records that are older than the specified File Retention period. The pruning will occur at the end of a backup Job for the given Client. Note that the Client database record contains a copy of the File and Job retention periods, but Bacula uses the current values found in the Director's Client resource to do the pruning.

Retention periods are specified in seconds, but as a convenience, there are a number of modifiers that permit easy specification in terms of minutes, hours, days, weeks, months, quarters, or years on the record. See the Configuration chapter of this manual for additional details of modifier specification.

The default is 60 days.

Job Retention = <time-period-specification> The Job Retention record defines the length of time that **Bacula** will keep Job records in the Catalog database. When this time period expires, and if **AutoPrune** is set to **yes** Bacula will prune (remove) Job records that are older than the specified File Retention period. Note, if a Job record is selected for pruning, all associated File and JobMedia records will also be pruned regardless of the File Retention period set. As a consequence, you normally will set the File retention period to be less than the Job retention period.

The retention period is specified in seconds, but as a convenience, there are a number of modifiers that permit easy specification in terms of minutes, hours, days, weeks, months, quarters, or years. See the Configuration chapter of this manual for additional details of modifier specification.

The default is 180 days.

AutoPrune = <yes/no> If AutoPrune is set to **yes** (default), Bacula will automatically apply the File retention period and the Job retention period for the Client at the end of the Job.

If you turn this off by setting it to **no**, your Catalog will grow each time you run a Job.

Compacting Your MySQL Database

Over time, as noted above, your database will tend to grow. I've noticed that even though Bacula regularly prunes files, **MySQL** does not effectively use the space, and instead continues growing. To avoid this, from time to time, you must compact your database. Normally, large commercial database such as Oracle have commands that will compact a database to reclaim wasted file space. MySQL has the **OPTIMIZE TABLE** command that you can use, and SQLite version 2.8.4 and greater has the **VACUUM** command. We leave it to you to explore the utility of the **OPTIMIZE TABLE** command in MySQL.

All database programs have some means of writing the database out in ASCII format and then reloading it. Doing so will re-create the database from scratch producing a compacted result, so below, we show you how you can do this for both MySQL and SQLite.

For a **MySQL** database, you could write the Bacula database as an ASCII file (bacula.sql) then reload it by doing the following:

```
mysqldump -f --opt bacula > bacula.sql
mysql bacula < bacula.sql
rm -f bacula.sql
```

Depending on the size of your database, this will take more or less time and a fair amount of disk space. For example, if I cd to the location of the MySQL Bacula database (typically /opt/mysql/var or something similar) and enter:

```
du bacula
```

I get **620,644** which means there are that many blocks containing 1024 bytes each or approximately 635 MB of data. After doing the **msqldump**, I had a bacula.sql file that had **174,356** blocks, and after doing the **mysql** command to recreate the database, I ended up with a total of **210,464** blocks rather than the original **629,644**. In other words, the compressed version of the database took approximately one third of the space of the database that had been in use for about a year.

As a consequence, I suggest you monitor the size of your database and from time to time (once every 6 months or year), compress it.

Repairing Your MySQL Database

If you find that you are getting errors writing to your MySQL database, or Bacula hangs each time it tries to access the database, you should consider running MySQL's database check and repair routines. The program you need to run depends on the type of database indexing you are using. If you are using the default, you will probably want to use **myisamchk**. For more details on how to do this, please consult the MySQL document at: http://www.mysql.com/doc/en/Repair.html.

If the errors you are getting are simply SQL warnings, then you might try running dbcheck before (or possibly after) using the MySQL database repair program. It can clean up many of the orphanned record problems, and certain other inconsistencies in the Bacula database.

Repairing Your PostgreSQL Database

The same considerations apply that are indicated above for MySQL. That is, consult the PostgreSQL documents for how to repair the database, and also

consider using Bacula's dbcheck program if the conditions are reasonable for using (see above).

Compacting Your PostgreSQL Database

Over time, as noted above, your database will tend to grow. I've noticed that even though Bacula regularly prunes files, PostgreSQL has a **VACUUM** command that will compact your database for you. Alternatively you may want to use the **vacuumdb** command, which can be run from a cron job.

All database programs have some means of writing the database out in ASCII format and then reloading it. Doing so will re-create the database from scratch producing a compacted result, so below, we show you how you can do this for PostgreSQL.

For a **PostgreSQL** database, you could write the Bacula database as an ASCII file (bacula.sql) then reload it by doing the following:

```
pg_dump bacula > bacula.sql
cat bacula.sql | psql bacula
rm -f bacula.sql
```

Depending on the size of your database, this will take more or less time and a fair amount of disk space. For example, you can **cd** to the location of the Bacula database (typically /usr/local/pgsql/data or possible /var/lib/pgsql/data) and check the size.

Compacting Your SQLite Database

First please read the previous section that explains why it is necessary to compress a database. SQLite version 2.8.4 and greater have the **Vacuum** command for compacting the database.

```
cd {\bf working-directory}
echo 'vacuum;' | sqlite bacula.db
```

As an alternative, you can use the following commands, adapted to your system:

```
cd {\bf working-directory}
```

```
echo '.dump' | sqlite bacula.db > bacula.sql
rm -f bacula.db
sqlite bacula.db < bacula.sql
rm -f bacula.sql
```

Where **working-directory** is the directory that you specified in the Director's configuration file. Note, in the case of SQLite, it is necessary to completely delete (rm) the old database before creating a new compressed version.

Migrating from SQLite to MySQL

You may begin using Bacula with SQLite then later find that you want to switch to MySQL for any of a number of reasons: SQLite tends to use more disk than MySQL, SQLite apparently does not handle database sizes greater than 2GBytes, ... Several users have done so by first producing an ASCII "dump" of the SQLite database, then creating the MySQL tables with the **create_mysql_tables** script that comes with Bacula, and finally feeding the SQLite dump into MySQL using the **-f** command line option to continue past the errors that are generated by the DDL statements that SQLite's dump creates. Of course, you could edit the dump and remove the offending statements. Otherwise, MySQL accepts the SQL produced by SQLite.

Backing Up Your Bacula Database

If ever the machine on which you Bacula database crashes, and you need to restore from backup tapes, one of your first priorities will probably be to recover the database. Although Bacula will happily backup your catalog database if it is specified in the FileSet, this is not a very good way to do it because the database will be saved while Bacula is modifying it. Thus the database may be in and instable state. Worse yet, you will backup the database before all the Bacula updates have been applied.

To resolve these problems, you need backup the database after all the backup jobs have been run. In addition, you will want to make a copy while Bacula is not modifying it. To do so, you can use two scripts provided in the release **make_catalog_backup** and **delete_catalog_backup**. These files will be automatically generated along with all the other Bacula scripts. The first script will make an ASCII copy of your Bacula database into **bacula.sql** in the working directory you specified on your configuration, and the second will delete the **bacula.sql** file.

The basic sequence of events to make this work correctly is as follows:

- Run all your nightly backups
- After running your nightly backups, run a Catalog backup Job
- The Catalog backup job must be scheduled after your last nightly backup
- You use **RunBeforeJob** to create the ASCII backup file and **RunAfterJob** to clean up

Assuming that you start all your nightly backup jobs at 1:05 am (and that they run one after another), you can do the catalog backup with the following additional Director configuration statements:

```
# Backup the catalog database (after the nightly save)
Job {
  Name = "BackupCatalog"
  Type = Backup
  Client=rufus-fd
  FileSet="Catalog"
  Schedule = "WeeklyCycleAfterBackup"
  Storage = DLTDrive
  Messages = Standard
  Pool = Default
  RunBeforeJob = "/home/kern/bacula/bin/make_catalog_backup"
  RunAfterJob  = "/home/kern/bacula/bin/delete_catalog_backup"
}
# This schedule does the catalog. It starts after the WeeklyCycle
Schedule {
  Name = "WeeklyCycleAfterBackup
  Run = Full sun-sat at 1:10
}
# This is the backup of the catalog
FileSet {
  Name = "Catalog"
  Include = signature=MD5 {
     @working_directory@/bacula.sql
  }
}
```

Backing Up Third Party Databases

If you are running a database in production mode on your machine, Bacula will happily backup the files, but if the database is in use while Bacula is reading it, you may back it up in an unstable state.

The best solution is to shutdown your database before backing it up, or use some tool specific to your database to make a valid live copy perhaps by dumping the database in ASCII format. I am not a database expert, so I cannot provide you advice on how to do this, but if you are unsure about how to backup your database, you might try visiting the Backup Central site, which has been renamed Storage Mountain (www.backupcentral.com). In particular, their Free Backup and Recovery Software page has links to scripts that show you how to shutdown and backup most major databases.

Database Size

As mentioned above, if you do not do automatic pruning, your Catalog will grow each time you run a Job. Normally, you should decide how long you want File records to be maintained in the Catalog and set the **File Retention** period to that time. Then you can either wait and see how big your Catalog gets or make a calculation assuming approximately 154 bytes for each File saved and knowing the number of Files that are saved during each backup and the number of Clients you backup.

For example, suppose you do a backup of two systems, each with 100,000 files. Suppose further that you do a Full backup weekly and an Incremental every day, and that the Incremental backup typically saves 4,000 files. The size of your database after a month can roughly be calculated as:

```
   Size = 154 * No. Systems * (100,000 * 4 + 10,000 * 26)
```

where we have assumed 4 weeks in a month and 26 incremental backups per month. This would give the following:

```
   Size = 154 * 2 * (100,000 * 4 + 10,000 * 26)
or
   Size = 308 * (400,000 + 260,000)
or
   Size = 203,280,000 bytes
```

So for the above two systems, we should expect to have a database size of approximately 200 Megabytes. Of course, this will vary according to how many files are actually backed up.

Below are some statistics for a MySQL database containing Job records for five Clients beginning September 2001 through May 2002 (8.5 months) and File records for the last 80 days. (Older File records have been pruned). For

these systems, only the user files and system files that change are backed up. The core part of the system is assumed to be easily reloaded from the RedHat rpms.

In the list below, the files (corresponding to Bacula Tables) with the extension .MYD contain the data records whereas files with the extension .MYI contain indexes.

You will note that the File records (containing the file attributes) make up the large bulk of the number of records as well as the space used (459 Mega Bytes including the indexes). As a consequence, the most important Retention period will be the **File Retention** period. A quick calculation shows that for each File that is saved, the database grows by approximately 150 bytes.

Size in Bytes	Records	File
168	5	Client.MYD
3,072		Client.MYI
344,394,684	3,080,191	File.MYD
115,280,896		File.MYI
2,590,316	106,902	Filename.MYD
3,026,944		Filename.MYI
184	4	FileSet.MYD
2,048		FileSet.MYI
49,062	1,326	JobMedia.MYD
30,720		JobMedia.MYI
141,752	1,378	Job.MYD
13,312		Job.MYI
1,004	11	Media.MYD
3,072		Media.MYI
1,299,512	22,233	Path.MYD
581,632		Path.MYI
36	1	Pool.MYD
3,072		Pool.MYI
5	1	Version.MYD
1,024		Version.MYI

This database has a total size of approximately 450 Megabytes.

If we were using SQLite, the determination of the total database size would be much easier since it is a single file, but we would have less insight to the size of the individual tables as we have in this case.

Note, SQLite databases may be as much as 50% larger than MySQL databases due to the fact that all data is stored as ASCII strings. That is even binary integers are stored as ASCII strings, and this seems to increase the space needed.

Automatic Volume Recycling

Normally, Bacula will write on a volume, and once the tape is written, it can append to the volume, but it will never overwrite the data thus destroying it. When we speak of **recycling** volumes, we mean that **Bacula** can write over the previous contents of a volume. Thus all previous data will be lost.

You may not want Bacula to automatically recycle (reuse) tapes. This requires a large number of tapes, and in such a case, it is possible to manually recycle tapes. For more on manual recycling, see the section entitled Manually Recycling Volumes below in this chapter.

Most people prefer to have a Pool of tapes that are used for daily backups and recycled once a week, another Pool of tapes that are used for Full backups once a week and recycled monthly, and finally a Pool of tapes that are used once a month and recycled after a year or two. With a scheme like this, your pool of tapes remains constant.

By properly defining your Volume Pools with appropriate Retention periods, Bacula can manage the recycling (such as defined above) automatically.

Automatic recycling of Volumes is controlled by three records in the **Pool** resource definition in the Director's configuration file. These three records are:

- **AutoPrune = yes**
- **VolumeRetention = <time>**
- **Recycle = yes**

Automatic recycling of Volumes is performed by Bacula only when it wants a new Volume and no appendable Volumes are available in the Pool. It will then search the Pool for any Volumes with the **Recycle** flag set and whose Volume Status is **Full**. At that point, the recycling occurs in two steps. The first is that the Catalog for a Volume must be purged of all Jobs and Files contained on that Volume, and the second step is the actual recycling of the Volume. The Volume will be purged if the VolumeRetention period has expired. When a Volume is marked as Purged, it means that no Catalog records reference that Volume, and the Volume can be recycled. Until recycling actually occurs, the Volume data remains intact. If no Volumes can be found for recycline for any of the reasons stated above, Bacula will request operator intervention (i.e. it will ask you to label a new volume).

A key point mentioned above that can be a source of frustration is that Bacula will only recycle purged Volumes if there is no other appendable Volume available, otherwise, it will always write to an appendable Volume before recycling even if there are Volume marked as Purged. This preserves your data as long as possible. So, if you wish to "force" Bacula to use a purged Volume, you must first ensure that no other Volume in the Pool is marked **Append**. If necessary, you can manually set a volume to **Full**. The reason for this is that Bacula wants to preserve the data on your old tapes (even though purged from the catalog) as long as absolutely possible before overwriting it.

Automatic Pruning

As Bacula writes files to tape, it keeps a list of files, jobs, and volumes in a database called the catalog. Among other things, the database helps Bacula to decide which files to back up in an incremental or differential backup, and helps you locate files on past backups when you want to restore something. However, the catalog will grow larger and larger as time goes on, and eventually it can become unacceptably large.

Bacula's process for removing entries from the catalog is called Pruning. The default is Automatic Pruning, which means that once an entry reaches a certain age (e.g. 30 days old) it is removed from the catalog. Once a job has been pruned, you can still restore it from the backup tape, but one additional step is required: scanning the volume with bscan. The alternative to Automatic Pruning is Manual Pruning, in which you explicitly tell Bacula to erase the catalog entries for a volume. You'd usually do this when you want to reuse a Bacula volume, because there's no point in keeping a list of files that USED TO BE on a tape. Or, if the catalog is starting to get too big, you could prune the oldest jobs to save space. Manual pruning is done with the prune command in the console. (thanks to Bryce Denney for the above explanation).

Prunning Directives

There are three pruning durations. All apply to catalog database records and not to the actual data in a Volume. The pruning (or retention) durations are for: Volumes (Media records), Jobs (Job records), and Files (File records). The durations inter-depend a bit because if Bacula prunes a Volume, it automatically removes all the Job records, and all the File records. Also when a Job record is pruned, all the File records for that Job are also pruned (deleted) from the catalog.

Having the File records in the database means that you can examine all the files backed up for a particular Job. They take the most space in the catalog (probably 90-95% of the total). When the File records are pruned, the Job records can remain, and you can still examine what Jobs ran, but not the details of the Files backed up. In addition, without the File records, you cannot use the Console restore command to restore the files.

When a Job record is pruned, the Volume (Media record) for that Job can still remain in the database, and if you do a "list volumes", you will see the volume information, but the Job records (and its File records) will no longer be available.

In each case, pruning removes information about where older files are, but it also prevents the catalog from growing to be too large. You choose the retention periods in function of how many files you are backing up and the time periods you want to keep those records online, and the size of the database.You can always re-insert the records (with 98% of the original data) by using "bscan" to scan in a whole Volume or any part of the volume that you want.

By setting **AutoPrune** to **yes** you will permit **Bacula** to automatically prune all Volumes in the Pool when a Job needs another Volume. Volume pruning means removing records from the catalog. It does not shrink the size of the Volume or effect the Volume data until the Volume gets overwritten. When a Job requests another volume and there are no Volumes with Volume Status **Append** available, Bacula will begin volume pruning. This means that all Jobs that are older than the **VolumeRetention** period will be pruned from every Volume that has Volume Status **Full** or **Used** and has Recycle set to **yes**. Pruning consists of deleting the corresponding Job, File, and JobMedia records from the catalog database. No change to the physical data on the Volume occurs during the pruning process. When all files are pruned from a Volume (i.e. no records in the catalog), the Volume will be marked as **Purged** implying that no Jobs remain on the volume. The Pool records that control the pruning are described below.

AutoPrune = <yes—no> If AutoPrune is set to **yes** (default), Bacula (version 1.20 or greater) will automatically apply the Volume retention period when running a Job and it needs a new Volume but no appendable volumes are available. At that point, Bacula will prune all Volumes that can be pruned (i.e. AutoPrune set) in an attempt to find a usable volume. If during the autoprune, all files are pruned from the Volume, it will be marked with VolStatus **Purged**. The default is **yes**.

Volume Retention = <time-period-specification> The Volume

Retention record defines the length of time that Bacula will guarantee that the Volume is not reused counting from the time the last job stored on the Volume terminated.

When this time period expires, and if **AutoPrune** is set to **yes**, and a new Volume is needed, but no appendable Volume is available, Bacula will prune (remove) Job records that are older than the specified Volume Retention period.

The Volume Retention period takes precedence over any Job Retention period you have specified in the Client resource. It should also be noted, that the Volume Retention period is obtained by reading the Catalog Database Media record rather than the Pool resource record. This means that if you change the VolumeRetention in the Pool resource record, you must ensure that the corresponding change is made in the catalog by using the **update pool** command. Doing so will insure that any new Volumes will be created with the changed Volume Retention period. Any existing Volumes will have their own copy of the Volume Retention period that can only be changed on a Volume by Volume basis using the **update volume** command.

When all file catalog entries are removed from the volume, its VolStatus is set to **Purged**. The files remain physically on the Volume until the volume is overwritten.

Retention periods are specified in seconds, minutes, hours, days, weeks, months, quarters, or years on the record. See the Configuration chapter of this manual for additional details of time specification.

The default is 1 year.

Recycle = <yes—no> This statement tells Bacula whether or not the particular Volume can be recycled (i.e. rewritten). If Recycle is set to **no** (the default), then even if Bacula prunes all the Jobs on the volume and it is marked **Purged**, it will not consider the tape for recycling. If Recycle is set to **yes** and all Jobs have been pruned, the volume status will be set to **Purged** and the volume may then be reused when another volume is needed. If the volume is reused, it is relabeled with the same Volume Name, however all previous data will be lost.

Note, it is also possible to "force" pruning of all Volumes in the Pool associated with a Job by adding **Prune Files = yes** to the Job resource.

Recycling Algorithm

After all Volumes of a Pool have been pruned (as mentioned above, this happens when a Job needs a new Volume and no appendable Volumes are available), Bacula will look for the oldest Volume that is Purged (all Jobs and Files expired), and if the **Recycle** flag is on (Recycle=yes) for that Volume, Bacula will relabel it and write new data on it.

The full recycling algorithm that Bacula uses when it needs a new Volume is:

- Search the Pool for a Volume with VolStatus=Append (if there is more than one, the Volume with the oldest date last written is chosen. If two have the same date then the one with the lowest MediaId is chosen).
- Search the Pool for a Volume with VolStatus=Recycle (if there is more than one, the Volume with the oldest date last written is chosen. If two have the same date then the one with the lowest MediaId is chosen).
- Prune volumes applying Volume retention period (Volumes with VolStatus Full, Used, or Append are pruned).
- Search the Pool for a Volume with VolStatus=Purged
- Attempt to create a new Volume if automatic labeling enabled
- Prune the oldest Volume if RecycleOldestVolume=yes (the Volume with the oldest LastWritten date and VolStatus equal to Full, Recycle, Purged, Used, or Append is chosen). This record ensures that all retention periods are properly respected.
- Purge the oldest Volume if PurgeOldestVolume=yes (the Volume with the oldest LastWritten date and VolStatus equal to Full, Recycle, Purged, Used, or Append is chosen). We strongly recommend against the use of **PurgeOldestVolume** as it can quite easily lead to loss of current backup data.
- Give up and ask operator.

The above occurs when Bacula has finished writing a Volume or when no Volume is present in the drive.

On the other hand, if you have inserted a different Volume after the last job, and Bacula recognizes the Volume as valid, it will request authorization from the Director to use this Volume. In this case, if you have set **Recycle**

Current Volume = yes and the Volume is marked as Used or Full, Bacula will prune the volume and if all jobs were removed during the pruning (respecting the retention periods), the Volume will be recycled and used. For this to work, you must have **Accept Any Volume = yes** in the Pool. The recycling algorithm in this case is:

- If the VolStatus is **Append** or **Recycle** and **Accept Any Volume** is set, the volume will be used.
- If **Recycle Current Volume** is set and the volume is marked **Full** or **Used**, Bacula will prune the volume (applying the retention period). If all Jobs are pruned from the volume, it will be recycled.

This permits users to manually change the Volume every day and load tapes in an order different from what is in the catalog, and if the volume does not contain a current copy of your backup data, it will be used.

Recycle Status

Each Volume inherits the Recycle status (yes or no) from the Pool resource record when the Media record is created (normally when the Volume is labeled). This Recycle status is stored in the Media record of the Catalog. Using the the Console program, you may subsequently change the Recycle status for each Volume. For example in the following output from **list volumes**:

```
+----------+-------+--------+---------+------------+--------+-----+
| VolumeNa | Media | VolSta | VolByte | LastWritte | VolRet | Rec |
+----------+-------+--------+---------+------------+--------+-----+
| File0001 | File  | Full   | 4190055 | 2002-05-25 | 14400  | 1   |
| File0002 | File  | Full   | 1896460 | 2002-05-26 | 14400  | 1   |
| File0003 | File  | Full   | 1896460 | 2002-05-26 | 14400  | 1   |
| File0004 | File  | Full   | 1896460 | 2002-05-26 | 14400  | 1   |
| File0005 | File  | Full   | 1896460 | 2002-05-26 | 14400  | 1   |
| File0006 | File  | Full   | 1896460 | 2002-05-26 | 14400  | 1   |
| File0007 | File  | Purged | 1896466 | 2002-05-26 | 14400  | 1   |
+----------+-------+--------+---------+------------+--------+-----+
```

all the volumes are marked as recyclable, and the last Volume, **File0007** has been purged, so it may be immediately recycled. The other volumes are all marked recyclable and when their Volume Retention period (14400 seconds or 4 hours) expires, they will be eligible for pruning, and possible recycling. Even though Volume **File0007** has been purged, all the data on the Volume

is still recoverable. A purged Volume simply means that there are no entries in the Catalog. Even if the Volume Status is changed to **Recycle**, the data on the Volume will be recoverable. The data is lost only when the Volume is re-labeled and re-written.

To modify Volume **File0001** so that it cannot be recycled, you use the **update volume pool=File** command in the console program, or simply **update** and Bacula will prompt you for the information.

```
+----------+------+-------+---------+-------------+-------+-----+
| VolumeNa | Media| VolSta| VolByte | LastWritten | VolRet| Rec |
+----------+------+-------+---------+-------------+-------+-----+
| File0001 | File | Full  | 4190055 | 2002-05-25  | 14400 | 0   |
| File0002 | File | Full  | 1897236 | 2002-05-26  | 14400 | 1   |
| File0003 | File | Full  | 1896460 | 2002-05-26  | 14400 | 1   |
| File0004 | File | Full  | 1896460 | 2002-05-26  | 14400 | 1   |
| File0005 | File | Full  | 1896460 | 2002-05-26  | 14400 | 1   |
| File0006 | File | Full  | 1896460 | 2002-05-26  | 14400 | 1   |
| File0007 | File | Purged| 1896466 | 2002-05-26  | 14400 | 1   |
+----------+------+-------+---------+-------------+-------+-----+
```

In this case, **File0001** will never be automatically recycled. The same effect can be achieved by setting the Volume Status to Read-Only.

Making Bacula Use a Single Tape

Most people will want Bacula to fill a tape and when it is full, a new tape will be mounted, and so on. However, as an extreme example, it is possible for Bacula to write on a single tape, and every night to rewrite it. To get this to work, you must do two things: first, set the VolumeRetention to less than your save period (one day), and the second item is to make Bacula mark the tape as full after using it once. This is done using **UseVolumeOnce = yes**. If this latter record is not used and the tape is not full after the first time it is written, Bacula will simply append to the tape and eventually request another volume. Using the tape only once, forces the tape to be marked **Full** after each use, and the next time **Bacula** runs, it will recycle the tape.

An example Pool resource that does this is:

```
Pool {
  Name = DDS-4
  Use Volume Once = yes
  Pool Type = Backup
  AutoPrune = yes
```

```
  VolumeRetention = 12h # expire after 12 hours
  Recycle = yes
}
```

A Daily, Weekly, Monthly Tape Usage Example

This example is meant to show you how one could define a fixed set of volumes that Bacula will rotate through on a regular schedule. There are an infinite number of such schemes, all of which have various advantages and disadvantages.

We start with the following assumptions:

- A single tape has more than enough capacity to do a full save.
- There are 10 tapes that are used on a daily basis for incremental backups. They are prelabeled Daily1 ... Daily10.
- There are 4 tapes that are used on a weekly basis for full backups. They are labeled Week1 ... Week4.
- There are 12 tapes that are used on a monthly basis for full backups. They are numbered Month1 ... Month12
- A full backup is done every Saturday evening (tape inserted Friday evening before leaving work).
- No backups are done over the weekend (this is easy to change).
- The first Friday of each month, a Monthly tape is used for the Full backup.
- Incremental backups are done Monday - Friday (actually Tue-Fri mornings).

We start the system by doing a Full save to one of the weekly volumes or one of the monthly volumes. The next morning, we remove the tape and insert a Daily tape. Friday evening, we remove the Daily tape and insert the next tape in the Weekly series. Monday, we remove the Weekly tape and re-insert the Daily tape. On the first Friday of the next month, we insert the next Monthly tape in the series rather than a Weekly tape, then continue. When a Daily tape finally fills up, **Bacula** will request the next one in the series, and the next day when you notice the email message, you will mount it and **Bacula** will finish the unfinished incremental backup.

What does this give? Well, at any point, you will have a the last complete Full save plus several Incremental saves. For any given file your want to recover (or your whole system), you will have a copy of that file every day for at least the last 14 days. For older versions, you will have at least 3 and probably 4 Friday full saves of that file, and going back further, you will have a copy of that file made on the beginning of the month for at least a year.

So you have copies of any file (or your whole system) for at least a year, but as you go back in time, the time between copies increases from daily to weekly to monthly.

What would the Bacula configuration look like to implement such a scheme?

```
Schedule {
  Name = "NightlySave"
  Run = Level=Full Pool=Monthly 1st sat at 03:05
  Run = Level=Full Pool=Weekly 2nd-5th sat at 03:05
  Run = Level=Incremental Pool=Daily tue-fri at 03:05
}
Job {
  Name = "NightlySave"
  Type = Backup
  Level = Full
  Client = LocalMachine
  FileSet = "File Set"
  Messages = Standard
  Storage = DDS-4
  Pool = Daily
  Schedule = "NightlySave"
}
# Definition of file storage device
Storage {
  Name = DDS-4
  Address = localhost
  SDPort = 9103
  Password = XXXXXXXXXXXX
  Device = FileStorage
  Media Type = 8mm
}
FileSet {
  Name = "File Set"
  Include = signature=MD5 {
    ffffffffffffffff
  }
  Exclude = { *.o }
}
Pool {
  Name = Daily
  Pool Type = Backup
  AutoPrune = yes
```

```
  VolumeRetention = 10d   # recycle in 10 days
  Maximum Volumes = 10
  Recycle = yes
}
Pool {
  Name = Weekly
  Use Volume Once = yes
  Pool Type = Backup
  AutoPrune = yes
  VolumeRetention = 30d  # recycle in 30 days (default)
  Recycle = yes
}
Pool {
  Name = Monthly
  Use Volume Once = yes
  Pool Type = Backup
  AutoPrune = yes
  VolumeRetention = 365d  # recycle in 1 year
  Recycle = yes
}
```

Automatic Pruning and Recycling Example

Perhaps the best way to understand the various resource records that come into play during automatic pruning and recycling is to run a Job that goes through the whole cycle. If you add the following resources to your Director's configuration file:

```
Schedule {
  Name = "30 minute cycle"
  Run = Level=Full Pool=File Messages=Standard Storage=File
        hourly at 0:05
  Run = Level=Full Pool=File Messages=Standard Storage=File
        hourly at 0:35
}
Job {
  Name = "Filetest"
  Type = Backup
  Level = Full
  Client=XXXXXXXXXX
  FileSet="Test Files"
  Messages = Standard
  Storage = File
  Pool = File
  Schedule = "30 minute cycle"
}
# Definition of file storage device
Storage {
  Name = File
  Address = XXXXXXXXXX
```

```
  SDPort = 9103
  Password = XXXXXXXXXXXX
  Device = FileStorage
  Media Type = File
}
FileSet {
  Name = "Test Files"
  Include = signature=MD5 {
    ffffffffffffffff
  }
  Exclude = { *.o }
}
Pool {
  Name = File
  Use Volume Once = yes
  Pool Type = Backup
  LabelFormat = "File"
  AutoPrune = yes
  VolumeRetention = 4h
  Maximum Volumes = 12
  Recycle = yes
}
```

Where you will need to replace the **ffffffff**'s by the appropriate files to be saved for your configuration. For the FileSet Include, choose a directory that has one or two megabytes maximum since there will probably be approximately 8 copies of the directory that **Bacula** will cycle through.

In addition, you will need to add the following to your Storage daemon's configuration file:

```
Device {
  Name = FileStorage
  Media Type = File
  Archive Device = /tmp
  LabelMedia = yes;
  Random Access = Yes;
  AutomaticMount = yes;
  RemovableMedia = no;
  AlwaysOpen = no;
}
```

With the above resources, Bacula will start a Job every half hour that saves a copy of the directory you chose to /tmp/File0001 ... /tmp/File0012. After 4 hours, Bacula will start recycling the backup Volumes (/tmp/File0001 ...). You should see this happening in the output produced. Bacula will automatically create the Volumes (Files) the first time it uses them.

To turn it off, either delete all the resources you've added, or simply comment out the **Schedule** record in the **Job** resource.

Manually Recycling Volumes

Although automatic recycling of Volumes is implemented in version 1.20 and later (see the Automatic Recycling of Volumes chapter of this manual), you may want to manually force reuse (recycling) of a Volume.

Assuming that you want to keep the Volume name, but you simply want to write new data on the tape, the steps to take are:

- Use the **update volume** command in the Console to ensure that the **Recycle** field is set to **1**
- Use the **purge jobs volume** command in the Console to mark the Volume as **Purged**. Check by using **list volumes**.

Once the Volume is marked Purged, it will be recycled the next time a Volume is needed.

If you wish to reuse the tape by giving it a new name, follow the following steps:

- Use the **purge jobs volume** command in the Console to mark the Volume as **Purged**. Check by using **list volumes**.
- In Bacula version 1.30 or greater, use the Console **relabel** command to relabel the Volume.

Please note that the relabel command applies only to tape Volumes.

For Bacula versions prior to 1.30 or to manually relabel the Volume, use the instructions below:

- Use the **delete volume** command in the Console to delete the Volume from the Catalog.
- If the a different tape is mounted, use the **unmount** command, remove the tape, and insert the tape to be renamed.
- Write an EOF mark in the tape using the following commands:

```
mt -f /dev/nst0 rewind
mt -f /dev/nst0 weof
```

where you replace **/dev/nst0** with the appropriate device name on your system.

- Use the **label** command to write a new label to the tape and to enter it in the catalog.

Please be aware that the **delete** command can be dangerous. Once it is done, to recover the File records, you must either restore your database as it was before the **delete** command, or use the **bscan** utility program to scan the tape and recreate the database entries.

Basic Volume Management

This chapter presents most all the features needed to do Volume management. Most of the concepts apply equally well to both tape and disk Volumes. However, the chapter was originally written to explain backing up to disk, so you will see it is slanted in that direction, but that all the directives presented here apply equally well whether your volume is disk or tape.

If you have a lot of hard disk storage or you absolutely must have your backups run within a small time window, you may want to direct Bacula to backup to disk Volumes rather than tape Volumes. This chapter is intended to give you some of the options that are available to you so that you can manage either disk or tape volumes.

Key Concepts and Resource Records

Getting Bacula to write to disk rather than tape in the simplest case is rather easy. In the Storage daemon's configuration file, you simply define an **Archive Device** to be a directory. For example, if you want your disk backups to go into the directory **/home/bacula/backups**, you could use the following:

```
Device {
  Name = FileBackup
  Media Type = File
  Archive Device = /home/bacula/backups
  Random Access = Yes;
  AutomaticMount = yes;
  RemovableMedia = no;
  AlwaysOpen = no;
}
```

Assuming you have the appropriate **Storage** resource in your Director's configuration file that references the above Device resource,

```
Storage {
  Name = FileStorage
  Address = ...
  Password = ...
  Device = FileBackup
  Media Type = File
}
```

Bacula will then write the archive to the file **/home/bacula/backups/<volume-name>** where <volume-name> is the volume name of a Volume defined in the Pool. For example, if you have labeled a Volume named **Vol001**, Bacula will write to the file **/home/bacula/backups/Vol001**. Although you can later move the archive file to another directory, you should not rename it or it will become unreadable by Bacula. This is because each archive has the filename as part of the internal label, and the internal label must agree with the system filename before Bacula will use it.

Although this is quite simple, there are a number of problems, the first is that unless you specify otherwise, Bacula will always write to the same volume until you run out of disk space.

Pool Options to Limit the Volume Usage

Some of the options you have, all of which are specified in the Pool record, are:

- To write each Volume only once (i.e. one Job per Volume or file in this case), use:

  UseVolumeOnce = yes.

- To write nnn Jobs to each Volume, use:

  Maximum Volume Jobs = nnn.

- To limit the maximum size of each Volume, use:

  Maximum Volume Bytes = mmmm.

- To limit the use time (i.e. write the Volume for a maximum of 5 days), use:

  Volume Use Duration = ttt.

Note that although you probably would not want to limit the number of bytes on a tape as you would on a disk Volume, the other options can be very useful in limiting the time Bacula will use a particular Volume (be it tape or disk). For example, the above directives can allow you to ensure that you rotate through a set of daily Volumes if you wish.

As mentioned above, each of those directives are specified in the Pool or Pools that you use for your Volumes. In the case of **Maximum Volume Job**, **Maximum Volume Bytes**, and **Volume Use Duration**, you can

actually specify the desired value on a Volume by Volume basis. The value specified in the Pool record becomes the default when labeling new Volumes. As an example of the use of one of the above, suppose your Pool resource contains:

```
Pool {
  Name = File
  Pool Type = Backup
  Volume Use Duration = 23h
}
```

then if you run a backup once a day (every 24 hours), Bacula will use a new Volume each backup because each Volume it writes can only be used for 23 hours after the first write.

Automatic Volume Labeling

Use of the above records brings up another problem – that of labeling your Volumes. For automated disk backup, you can either manually label each of your Volumes, or you can have Bacula automatically label new Volumes when they are needed. While, the automatic Volume labeling in version 1.30 and prior is a bit simplistic, but it does allow for automation, the features added in version 1.31 permit automatic creation of a wide variety of labels including information from environment variables and special Bacula Counter variables.

Please note that automatic Volume can also be used with tapes, but it is not nearly so practical since the tapes must be pre-mounted. This requires some user interaction. Automatic labeling from templates does NOT work with autochangers since Bacula will not access unknown slots. There are several methods of labeling all volumes in an autochanger magazine. For more information on this, please see the Autochanger chapter of this manual.

Automatic Volume labeling is enabled by making a change to both the Pool resource (Director) and to the Device resource (Storage daemon) shown above. In the case of the Pool resource, you must provide Bacula with a label format that it will use to create new names. In the simplest form, the label format is simply the Volume name, to which Bacula will append a four digit number. This number starts at 0001 and is incremented for each Volume the pool contains. Thus if you modify your Pool resource to be:

```
Pool {
  Name = File
```

```
  Pool Type = Backup
  Volume Use Duration = 23h
  LabelFormat = "Vol"
}
```

Bacula will create Volume names Vol0001, Vol0002, and so on when new Volumes are needed. Much more complex and elaborate labels can be created using variable expansion defined in the Variable Expansion chapter of this manual.

The second change that is necessary to make automatic labeling work is to give the Storage daemon permission to automatically label Volumes. Do so by adding **LabelMedia = yes** to the Device resource as follows:

```
Device {
  Name = File
  Media Type = File
  Archive Device = /home/bacula/backups
  Random Access = Yes;
  AutomaticMount = yes;
  RemovableMedia = no;
  AlwaysOpen = no;
  LabelMedia = yes
}
```

You can find more details of the **Label Format** Pool record in Label Format description of the Pool resource records.

Restricting the Number of Volumes and Recycling

Automatic labeling discussed above brings up the problem of Volume management. With the above scheme, a new Volume will be created every day. If you have not specified Retention periods, your Catalog will continue to fill keeping track of all the files Bacula has backed up, and this procedure will create one new archive file (Volume) every day.

The tools Bacula gives you to help automatically manage these problems are the following:

1. Catalog file record retention periods, the File Retention = ttt record in the Client resource.
2. Catalog job record retention periods, the Job Retention = ttt record in the Client resource.

3. The AutoPrune = yes record in the Client resource to permit application of the above two retention periods.
4. The Volume Retention = ttt record in the Pool resource.
5. The AutoPrune = yes record in the Pool resource to permit application of the Volume retention period.
6. The Recycle = yes record in the Pool resource to permit automatic recycling of Volumes whose Volume retention period has expired.
7. The Recycle Oldest Volume = yes record in the Pool resource tells Bacula to Prune the oldest volume in the Pool, and if all files were pruned to recyle this volume and use it.
8. The Recycle Current Volume = yes record in the Pool resource tells Bacula to Prune the currently mounted volume in the Pool, and if all files were pruned to recyle this volume and use it.
9. The Purge Oldest Volume = yes record in the Pool resource permits a forced recycling of the oldest Volume when a new one is needed. **N.B. This record ignores retention periods! We highly recommend not to use this record, but instead use Recycle Oldest Volume**
10. The Maximum Volumes = nnn record in the Pool resource to limit the number of Volumes that can be created.

The first three records (File Retention, Job Retention, and AutoPrune) determine the amount of time that Job and File records will remain in your Catalog, and they are discussed in detail in the Automatic Volume Recycling chapter of this manual.

Volume Retention, AutoPrune, and Recycle determine how long Bacula will keep your Volumes before reusing them, and they are also discussed in detail in the

Automatic Volume Recycling chapter of this manual.

The Maximum Volumes record can also be used in conjunction with the Volume Retention period to limit the total number of archive Volumes (files) that Bacula will create. By setting an appropriate Volume Retention period, a Volume will be purged just before it is needed and thus Bacula can cycle through a fixed set of Volumes. Cycling through a fixed set of Volumes can also be done by setting **Recycle Oldest Volume = yes** or **Recycle Current Volume = yes**. In this case, when Bacula needs a new Volume, it will prune the specified volume.

An Example

The following example is not very practical, but can be used to demonstrate the proof of concept in a relatively short period of time. The example consists of a single client that is backed up to a set of 12 archive files (Volumes). Each Volume is used (written) only once, and there are four Full saves done every hour (so the whole thing cycles around after three hours).

The Director's configuration file is as follows:

```
Director {
  Name = my-dir
  QueryFile = "~/bacula/bin/query.sql"
  PidDirectory = "~/bacula/working"
  WorkingDirectory = "~/bacula/working"
  Password = dir_password
}
Schedule {
  Name = "FourPerHour"
  Run = Level=Full Pool=Recycle Storage=File hourly at 0:05
  Run = Level=Full Pool=Recycle Storage=File hourly at 0:20
  Run = Level=Full Pool=Recycle Storage=File hourly at 0:35
  Run = Level=Full Pool=Recycle Storage=File hourly at 0:50
}
Job {
  Name = "RecycleExample"
  Type = Backup
  Level = Full
  Client = Rufus
  FileSet= "Example FileSet"
  Messages = Standard
  Storage = FileStorage
  Pool = Recycle
  Schedule = FourPerHour
}
FileSet {
  Name = "Example FileSet"
  Include = compression=GZIP signature=SHA1 {
    /home/kern/bacula/bin
  }
}
Client {
  Name = Rufus
  Address = rufus
  Catalog = BackupDB
  Password = client_password
}
Storage {
  Name = FileStorage
  Address = rufus
```

```
  Password = local_storage_password
  Device = RecycleDir
  Media Type = File
}
Catalog {
  Name = BackupDB
  dbname = bacula; user = bacula; password = ""
}
Messages {
  Name = Standard
  ...
}
Pool {
  Name = Recycle
  Use Volume Once = yes
  Pool Type = Backup
  LabelFormat = "Vol"
  AutoPrune = yes
  VolumeRetention = 2h
  Maximum Volumes = 12
  Recycle = yes
}
```

and the Storage daemon's configuration file is:

```
Storage {
  Name = my-sd
  WorkingDirectory = "~/bacula/working"
  Pid Directory = "~/bacula/working"
  MaximumConcurrentJobs = 10
}
Director {
  Name = my-dir
  Password = local_storage_password
}
Device {
  Name = RecycleDir
  Media Type = File
  Archive Device = /home/bacula/backups
  LabelMedia = yes;
  Random Access = Yes;
  AutomaticMount = yes;
  RemovableMedia = no;
  AlwaysOpen = no;
}
Messages {
  Name = Standard
  director = my-dir = all
}
```

In this example, the Jobs will be backed up to directory **/home/bacula/backups** with Volume names Vol0001, Vol0002, ...

Vol0012. Every backup Job will write a new volume cycling through the volume numbers, and two hours after a job has started, the volume will be pruned **Volume Retention = 2h**.

With a little bit of work, you can change the above example into a weekly or monthly cycle (take care about the amount of archive disk space used).

Backing up to Multiple Disks

Bacula can, of course, use multiple disks, but in general, each disk must be a separate Device specification in the Storage daemon's conf file, and you must then select what clients to backup to each disk.

The situation is a bit more complicated if you want to treat two disk drives logically as a single drive, which Bacula does not directly support. However, it is possible to back up your data to multiple disks as if they were a single drive by linking the Volumes from the first disk to the second disk.

For example, assume that you have two disks named **/disk1** and **/disk2¿**. If you then create a standard Storage daemon Device resource for backing up to the first disk, it will look like the following:

```
Device {
  Name = client1
  Media Type = File
  Archive Device = /disk1
  LabelMedia = yes;
  Random Access = Yes;
  AutomaticMount = yes;
  RemovableMedia = no;
  AlwaysOpen = no;
}
```

Since there is no way to get the above Device resource to reference both **/disk1** and **/disk2** we do it by pre-creating Volumes on /disk2 with the following:

```
ln -s /disk2/Disk2-vol001 /disk1/Disk2-vol001
ln -s /disk2/Disk2-vol002 /disk1/Disk2-vol002
ln -s /disk2/Disk2-vol003 /disk1/Disk2-vol003
...
```

At this point, you can label the Volumes as Volume **Disk2-vol001**, **Disk2-vol002**, ... and Bacula will use them as if they were on /disk1 but actually

write the data to /disk2. The only minor inconvenience with this method is that you must explicitly name the disks and cannot use automatic labeling unless you arrange to have the labels exactly match the links you have created.

Considerations for Multiple Clients

If we take the above example and add a second Client, here are a few considerations:

- Although the second client can write to the same set of Volumes, you will probably want to write to a different set.
- You can write to a different set of Volumes by defining a second Pool, which has a different name and a different **LabelFormat**.
- If you wish the Volumes for the second client to go into a different directory (perhaps even on a different filesystem to spread the load), you would do so by defining a second Device resource in the Storage daemon. The **Name** must be different, and the **Archive Device** could be different. To ensure that Volumes are never mixed from one pool to another, you might also define a different MediaType (e.g. **File1**).

In this example, we have two clients, each with a different Pool and a different number of archive files retained. They also write to different directories with different Volume labeling.

The Director's configuration file is as follows:

```
Director {
  Name = my-dir
  QueryFile = "~/bacula/bin/query.sql"
  PidDirectory = "~/bacula/working"
  WorkingDirectory = "~/bacula/working"
  Password = dir_password
}
# Basic weekly schedule
Schedule {
  Name = "WeeklySchedule"
  Run = Level=Full fri at 1:30
  Run = Level=Incremental sat-thu at 1:30
}
FileSet {
  Name = "Example FileSet"
```

```
  Include = compression=GZIP signature=SHA1 {
    /home/kern/bacula/bin
  }
}
Job {
  Name = "Backup-client1"
  Type = Backup
  Level = Full
  Client = client1
  FileSet= "Example FileSet"
  Messages = Standard
  Storage = File1
  Pool = client1
  Schedule = "WeeklySchedule"
}
Job {
  Name = "Backup-client2"
  Type = Backup
  Level = Full
  Client = client2
  FileSet= "Example FileSet"
  Messages = Standard
  Storage = File2
  Pool = client2
  Schedule = "WeeklySchedule"
}
Client {
  Name = client1
  Address = client1
  Catalog = BackupDB
  Password = client1_password
  File Retention = 7d
}
Client {
  Name = client2
  Address = client2
  Catalog = BackupDB
  Password = client2_password
}
# Two Storage definitions permits different directories
Storage {
  Name = File1
  Address = rufus
  Password = local_storage_password
  Device = client1
  Media Type = File
}
Storage {
  Name = File2
  Address = rufus
  Password = local_storage_password
  Device = client2
  Media Type = File
}
```

```
Catalog {
  Name = BackupDB
  dbname = bacula; user = bacula; password = ""
}
Messages {
  Name = Standard
  ...
}
# Two pools permits different cycling periods and Volume names
# Cycle through 15 Volumes (two weeks)
Pool {
  Name = client1
  Use Volume Once = yes
  Pool Type = Backup
  LabelFormat = "Client1-"
  AutoPrune = yes
  VolumeRetention = 13d
  Maximum Volumes = 15
  Recycle = yes
}
# Cycle through 8 Volumes (1 week)
Pool {
  Name = client2
  Use Volume Once = yes
  Pool Type = Backup
  LabelFormat = "Client2-"
  AutoPrune = yes
  VolumeRetention = 6d
  Maximum Volumes = 8
  Recycle = yes
}
```

and the Storage daemon's configuration file is:

```
Storage {
  Name = my-sd
  WorkingDirectory = "~/bacula/working"
  Pid Directory = "~/bacula/working"
  MaximumConcurrentJobs = 10
}
Director {
  Name = my-dir
  Password = local_storage_password
}
# Archive directory for Client1
Device {
  Name = client1
  Media Type = File
  Archive Device = /home/bacula/client1
  LabelMedia = yes;
  Random Access = Yes;
  AutomaticMount = yes;
```

```
  RemovableMedia = no;
  AlwaysOpen = no;
}
# Archive directory for Client2
Device {
  Name = client2
  Media Type = File
  Archive Device = /home/bacula/client2
  LabelMedia = yes;
  Random Access = Yes;
  AutomaticMount = yes;
  RemovableMedia = no;
  AlwaysOpen = no;
}
Messages {
  Name = Standard
  director = my-dir = all
}
```

Using Pools to Manage Volumes

If you manage 5 or 10 machines and have a nice tape backup, you don't need Pools, and you may wonder what they are good for. In this chapter, you will see that Pools can help you optimize disk storage space. The same techniques can be applied to a shop that has multiple tape drives, or that wants to mount various different Volumes to meet their needs.

The rest of this chapter will give an example involving backup to disk Volumes, but most of the information applies equally well for tape Volumes.

The Problem

A site that I administer (a charitable organization) had a tape DDS-3 tape drive that was failing. The exact reason for the failure is still unknown. Worse yet, their full backup size is about 15GB whereas the capacity of their broken DDS-3 was at best 8GB (rated 6/12). A new DDS-4 tape drive and the necessary cassettes was more expensive than their budget could handle.

The Solution

They want to maintain 6 months of backup data, and be able to access the old files on a daily basis for a week, a weekly basis for a month, then monthly for 6 months. In addition, and offsite capability was not needed (well perhaps it really is, but it was never used). Their daily changes amount to about 300MB on the average, or about 2GB per week.

As a consequence, the total volume of data they need to keep to meet their needs is about 100GB (15GB x 6 + 2GB x 5 + 0.3 x 7) = 102.1GB.

The chosen solution was to buy a 120GB hard disk for next to nothing – far less than 1/10th the price of a tape drive and the cassettes to handle the same amount of data, and to have Bacula write to disk files.

The rest of this chapter will explain how to setup Bacula so that it would automatically manage a set of disk files with the minimum intervention on my part. The system has been running since 22 January 2004 until today (08 April 2004) with no intervention. Since we have not yet crossed the six month boundary, we still lack some data to be sure the system performs as desired.

Overall Design

Getting Bacula to write to disk rather than tape in the simplest case is rather easy, and is documented in the previous chapter. In addition, all the directives discussed here are explained in that chapter. We'll leave it to you to look at the details there. If you haven't read it and are not familiar with Pools, you probably should at least read it once quickly for the ideas before continuing here.

One needs to consider about what happens if we have only a single large Bacula Volume defined on our hard disk. Everything works fine until the Volume fills, then Bacula will ask you to mount a new Volume. This same problem applies to the use of tape Volumes if your tape fills. Being a hard disk and the only one you have, this will be a bit of a problem. It should be obvious that it is better to use a number of smaller Volumes and arrange for Bacula to automatically recycle them so that the disk storage space can be reused. The other problem with a single Volume, is that at the current time (1.34.0) Bacula does not seek within a disk Volume, so restoring a single file can take more time than one would expect.

As mentioned, the solution is to have multiple Volumes, or files on the disk. To do so, we need to limit the use and thus the size of a single Volume, by time, by number of jobs, or by size. Any of these would work, but we chose to limit the use of a single Volume by putting a single job in each Volume with the exception of Volumes containing Incremental backup where there will be 6 jobs (a week's worth of data) per volume. The details of this will be discussed shortly.

The next problem to resolve is recycling of Volumes. As you noted from above, the requirements are to be able to restore monthly for 6 months, weekly for a month, and daily for a week. So to simplify things, why not do a Full save once a month, a Differential save once a week, and Incremental saves daily. Now since each of these different kinds of saves needs to remain valid for differing periods, the simplest way to do this (and possibly the only) is to have a separate Pool for each backup type.

The decision was to use three Pools: one for Full saves, one for Differential saves, and one for Incremental saves, and each would have a different number of volumes and a different Retention period to accomplish the requirements.

Full Pool

Putting a single Full backup on each Volume, will require six Full save Volumes, and a retention period of six months. The Pool needed to do that is:

```
Pool {
  Name = Full-Pool
  Pool Type = Backup
  Recycle = yes
  AutoPrune = yes
  Volume Retention = 6 months
  Accept Any Volume = yes
  Maximum Volume Jobs = 1
  Label Format = Full-
  Maximum Volumes = 6
}
```

Since these are disk Volumes, no space is lost by having separate Volumes for each backup (done once a month in this case). The items to note are the retention period of six months (i.e. they are recycled after 6 months), that there is one job per volume (Maximum Volume Jobs = 1), the volumes will be labeled Full-0001, ... Full-0006 automatically. One could have labeled these manual from the start, but why not use the features of Bacula.

Differential Pool

For the Differential backup Pool, we choose a retention period of a bit longer than a month and ensure that there is at least one Volume for each of the maximum of five weeks in a month. So the following works:

```
Pool {
  Name = Diff-Pool
  Pool Type = Backup
  Recycle = yes
  AutoPrune = yes
  Volume Retention = 40 days
  Accept Any Volume = yes
  Maximum Volume Jobs = 1
  Label Format = Diff-
  Maximum Volumes = 6
}
```

As you can see, the Differential Pool can grow to a maximum of six volumes, and the Volumes are retained 40 days and there after can be recycled. Finally

there is one job per volume. This, of course, could be tightened up a lot, but the expense here is a few GB which is not too serious.

Incremental Pool

Finally, here is the resource for the Incremental Pool:

```
Pool {
  Name = Inc-Pool
  Pool Type = Backup
  Recycle = yes
  AutoPrune = yes
  Volume Retention = 20 days
  Accept Any Volume = yes
  Maximum Volume Jobs = 6
  Label Format = Inc-
  Maximum Volumes = 5
}
```

We keep the data for 20 days rather than just a week as the needs require. To reduce the proliferation of volume names, we keep a week's worth of data (6 incremental backups) in each Volume. In practice, the retention period should be set to just a bit more than a week and keep only two or three volumes instead of five. Again, the lost is very little and as the system reaches the full steady state, we can adjust these values so that the total disk usage doesn't exceed the disk capacity.

The Actual Conf Files

The following example shows you the actual files used, with only a few minor modifications to simplify things.

The Director's configuration file is as follows:

```
Director {          # define myself
  Name = bacula-dir
  DIRport = 9101
  QueryFile = "/home/bacula/bin/query.sql"
  WorkingDirectory = "/home/bacula/working"
  PidDirectory = "/home/bacula/working"
  Maximum Concurrent Jobs = 1
  Password = " "
  Messages = Standard
}
```

```
#   By default, this job will back up to disk in /tmp
Job {
  Name = client
  Type = Backup
  Client = client-fd
  FileSet = "Full Set"
  Schedule = "WeeklyCycle"
  Storage = File
  Messages = Standard
  Pool = Default
  Full Backup Pool = Full-Pool
  Incremental Backup Pool = Inc-Pool
  Differential Backup Pool = Diff-Pool
  Write Bootstrap = "/home/bacula/working/client.bsr"
  Priority = 10
}
# List of files to be backed up
FileSet {
  Name = "Full Set"
  Include = signature=SHA1 compression=GZIP9 {
    /
    /usr
    /home
  }
  Exclude = {
     /proc /tmp /.journal /.fsck
  }
}
Schedule {
  Name = "WeeklyCycle"
  Run = Full 1st sun at 1:05
  Run = Differential 2nd-5th sun at 1:05
  Run = Incremental mon-sat at 1:05
}
Client {
  Name = client-fd
  Address = client
  FDPort = 9102
  Catalog = MyCatalog
  Password = " "
  AutoPrune = yes      # Prune expired Jobs/Files
  Job Retention = 6 months
  File Retention = 60 days
}
Storage {
  Name = File
  Address = localhost
  SDPort = 9103
  Password = " "
  Device = FileStorage
  Media Type = File
}
Catalog {
  Name = MyCatalog
```

```
  dbname = bacula; user = bacula; password = ""
}
Pool {
  Name = Full-Pool
  Pool Type = Backup
  Recycle = yes           # automatically recycle Volumes
  AutoPrune = yes         # Prune expired volumes
  Volume Retention = 6 months
  Accept Any Volume = yes # write on any volume in the pool
  Maximum Volume Jobs = 1
  Label Format = Full-
  Maximum Volumes = 6
}
Pool {
  Name = Inc-Pool
  Pool Type = Backup
  Recycle = yes           # automatically recycle Volumes
  AutoPrune = yes         # Prune expired volumes
  Volume Retention = 20 days
  Accept Any Volume = yes
  Maximum Volume Jobs = 6
  Label Format = Inc-
  Maximum Volumes = 5
}
Pool {
  Name = Diff-Pool
  Pool Type = Backup
  Recycle = yes
  AutoPrune = yes
  Volume Retention = 40 days
  Accept Any Volume = yes
  Maximum Volume Jobs = 1
  Label Format = Diff-
  Maximum Volumes = 6
}
Messages {
  Name = Standard
  mailcommand = "bsmtp -h mail.domain.com -f \"\(Bacula\) %r\"
      -s \"Bacula: %t %e of %c %l\" %r"
  operatorcommand = "bsmtp -h mail.domain.com -f \"\(Bacula\) %r\"
      -s \"Bacula: Intervention needed for %j\" %r"
  mail = root@domain.com = all, !skipped
  operator = root@domain.com = mount
  console = all, !skipped, !saved
  append = "/home/bacula/bin/log" = all, !skipped
}
```

and the Storage daemon's configuration file is:

```
Storage {               # definition of myself
  Name = bacula-sd
  SDPort = 9103       # Director's port
```

```
  WorkingDirectory = "/home/bacula/working"
  Pid Directory = "/home/bacula/working"
}
Director {
  Name = bacula-dir
  Password = " "
}
Device {
  Name = FileStorage
  Media Type = File
  Archive Device = /files/bacula
  LabelMedia = yes;    # lets Bacula label unlabeled media
  Random Access = Yes;
  AutomaticMount = yes;   # when device opened, read it
  RemovableMedia = no;
  AlwaysOpen = no;
}
Messages {
  Name = Standard
  director = bacula-dir = all
}
```

Backup Strategies

Although Recycling and Backing Up to Disk Volume have been discussed in previous chapters, this chapter is meant to give you an overall view of possible backup strategies and to explain their advantages and disadvantages.

Simple One Tape Backup

Probably the simplest strategy is to back everything up to a single tape and insert a new (or recycled) tape when it fills and Bacula requests a new one.

Advantages

- The operator intervenes only when a tape change is needed. (once a month at my site).
- There is little chance of operator error because the tape is not changed daily.
- A minimum number of tapes will be needed for a full restore. Typically the best case will be one tape and worst two.
- You can easily arrange for the Full backup to occur a different night of the month for each system, thus load balancing and shortening the backup time.

Disadvantages

- If your site burns down, you will lose your current backups, and in my case about a month of data.
- After a tape fills and you have put in a blank tape, the backup will continue, and this will generally happen during working hours.

Practical Details

This system is very simple. When the tape fills and Bacula requests a new tape, you **unmount** the tape from the Console program, insert a new tape

and **label** it. In most cases after the label, Bacula will automatically mount the tape and resume the backup. Otherwise, you simply **mount** the tape.

Using this strategy, one typically does a Full backup once a week following by daily Incremental backups. To minimize the amount of data written to the tape, once can do (as I do) a Full backup once a month on the first Sunday of the month, a Differential backup on the 2nd-5th Sunday of the month, and incremental backups the rest of the week.

Manually Changing Tapes

If you use the strategy presented above, Bacula will ask you to change the tape, and you will **unmount** it and the remount it when you have inserted the new tape.

If you do not wish to interact with Bacula to change each tape, there are several ways to get Bacula to release the tape:

- In your Storage daemon's Device resource, set **AlwaysOpen = no** In this case, Bacula will release the tape after every job. If you run several jobs, the tape will be rewound and repositioned to the end at the beginning of every job. This is not very efficient, but does let you change the tape whenever you want.
- Use a **RunAfterJob** statement to run a script after your last job. This could also be an **Admin** job that runs after all your backup jobs. The script could be something like:

```
#!/bin/sh
/full-path/console -c /full-path/console.conf <<END_OF_DATA
release storage=your-storage-name
END_OF_DATA
```

  In this example, you would have **AlwaysOpen=yes**, but the **release** command would tell Bacula to rewind the tape and on the next job assume the tape has changed. This strategy may not work on some systems, or on autochangers because Bacula will still keep the drive open.

- The final strategy is the similar to the previous case except that you would use the unmount command to force Bacula to release the drive. Then you would eject the tape, and remount it as follows:

```
#!/bin/sh
/full-path/console -c /full-path/console.conf <\&ltEND_OF_DATA
unmount storage=your-storage-name
END_OF_DATA
# the following is a shell command
mt eject
/full-path/console -c /full-path/console.conf <<END_OF_DATA
mount storage=your-storage-name
END_OF_DATA
```

Daily Tape Rotation

This scheme is quite different from the one mentioned above in that a Full backup is done to a different tape every day of the week. Generally, the backup will cycle continuously through 5 or 6 tapes each week. Variations are to use a different tape each Friday, and possibly at the beginning of the month. Thus if backups are done Monday through Friday only, you need only 5 tapes, and by having two Friday tapes, you need a total of 6 tapes. Many sites run this way, or using modifications of it based on two week cycles or longer.

Advantages

- All the data is stored on a single tape, so recoveries are simple and faster.
- Assuming the previous day's tape is taken offsite each day, a maximum of one days data will be lost if the site burns down.

Disadvantages

- The tape must be changed every day requiring a lot of operator intervention.
- More errors will occur because of human mistakes.
- If the wrong tape is inadvertently mounted, the Backup for that day will not occur exposing the system to data loss.
- There is much more movement of the tape each day (rewinds) leading to shorter tape drive life time.

- Initial setup of Bacula to run in this mode is more complicated than the Single tape system described above.
- Depending on the number of systems you have and their data capacity, it may not be possible to do a Full backup every night for time reasons or reasons of tape capacity.

Practical Details

The simplest way to "force" Bacula to use a different tape each day is to define a different Pool for each day of the the week a backup is done. In addition, you will need to specify appropriate Job and File retention periods so that Bacula will relabel and overwrite the tape each week rather than appending to it. Nic Bellamy has supplied an actual working model of this which we include here.

What is important is to create a different Pool for each day of the week, and on the **run** statement in the Schedule, to specify which Pool is to be used. He has one Schedule that accomplishes this, and a second Schedule that does the same thing for the Catalog backup run each day after the main backup (Priorities were not available when this script was written). In addition, he uses a **Max Start Delay** of 22 hours so that if the wrong tape is premounted by the operator, the job will be automatically canceled, and the backup cycle will re-synchronize the next day. He has named his Friday Pool **WeeklyPool** because in that Pool, he wishes to have several tapes to be able to restore to a time older than one week.

And finally, in his Storage daemon's Device resource, he has **Automatic Mount = yes** and **Always Open = No**. This is necessary for the tape ejection to work in his **end_of_backup.sh** script below.

For example, his bacula-dir.conf file looks like the following:

```
# /etc/bacula/bacula-dir.conf
#
# Bacula Director Configuration file
#
Director {
  Name = ServerName
  DIRport = 9101
  QueryFile = "/etc/bacula/query.sql"
  WorkingDirectory = "/var/lib/bacula"
  PidDirectory = "/var/run"
  SubSysDirectory = "/var/lock/subsys"
  Maximum Concurrent Jobs = 1
```

```
  Password = "console-pass"
  Messages = Standard
}
#
# Define the main nightly save backup job
#
Job {
  Name = "NightlySave"
  Type = Backup
  Client = ServerName
  FileSet = "Full Set"
  Schedule = "WeeklyCycle"
  Storage = Tape
  Messages = Standard
  Pool = Default
  Write Bootstrap = "/var/lib/bacula/NightlySave.bsr"
  Max Start Delay = 22h
}
# Backup the catalog database (after the nightly save)
Job {
  Name = "BackupCatalog"
  Type = Backup
  Client = ServerName
  FileSet = "Catalog"
  Schedule = "WeeklyCycleAfterBackup"
  Storage = Tape
  Messages = Standard
  Pool = Default
 # This creates an ASCII copy of the catalog
  RunBeforeJob = "/usr/lib/bacula/make_catalog_backup -u bacula"
 # This deletes the copy of the catalog, and ejects the tape
  RunAfterJob  = "/etc/bacula/end_of_backup.sh"
  Write Bootstrap = "/var/lib/bacula/BackupCatalog.bsr"
  Max Start Delay = 22h
}
# Standard Restore template, changed by Console program
Job {
  Name = "RestoreFiles"
  Type = Restore
  Client = ServerName
  FileSet = "Full Set"
  Storage = Tape
  Messages = Standard
  Pool = Default
  Where = /tmp/bacula-restores
}
# List of files to be backed up
FileSet {
  Name = "Full Set"
  Include = signature=MD5 {
    /
    /data
  }
  Exclude = { /proc /tmp /.journal }
```

```
}
#
# When to do the backups
#
Schedule {
  Name = "WeeklyCycle"
  Run = Level=Full Pool=MondayPool Monday at 8:00pm
  Run = Level=Full Pool=TuesdayPool Tuesday at 8:00pm
  Run = Level=Full Pool=WednesdayPool Wednesday at 8:00pm
  Run = Level=Full Pool=ThursdayPool Thursday at 8:00pm
  Run = Level=Full Pool=WeeklyPool Friday at 8:00pm
}
# This does the catalog. It starts after the WeeklyCycle
Schedule {
  Name = "WeeklyCycleAfterBackup"
  Run = Level=Full Pool=MondayPool Monday at 8:15pm
  Run = Level=Full Pool=TuesdayPool Tuesday at 8:15pm
  Run = Level=Full Pool=WednesdayPool Wednesday at 8:15pm
  Run = Level=Full Pool=ThursdayPool Thursday at 8:15pm
  Run = Level=Full Pool=WeeklyPool Friday at 8:15pm
}
# This is the backup of the catalog
FileSet {
  Name = "Catalog"
  Include = signature=MD5 {
     /var/lib/bacula/bacula.sql
  }
}
# Client (File Services) to backup
Client {
  Name = ServerName
  Address = dionysus
  FDPort = 9102
  Catalog = MyCatalog
  Password = "client-pass"
  File Retention = 30d
  Job Retention = 30d
  AutoPrune = yes
}
# Definition of file storage device
Storage {
  Name = Tape
  Address = dionysus
  SDPort = 9103
  Password = "storage-pass"
  Device = Tandberg
  Media Type = MLR1
}
# Generic catalog service
Catalog {
  Name = MyCatalog
  dbname = bacula; user = bacula; password = ""
}
# Reasonable message delivery -- send almost all to email address
```

```
#  and to the console
Messages {
  Name = Standard
  mailcommand = "/usr/sbin/bsmtp -h localhost -f \"\(Bacula\) %r\"
     -s \"Bacula: %t %e of %c %l\" %r"
  operatorcommand = "/usr/sbin/bsmtp -h localhost -f \"\(Bacula\) %r\"
     -s \"Bacula: Intervention needed for %j\" %r"
  mail = root@localhost = all, !skipped
  operator = root@localhost = mount
  console = all, !skipped, !saved
  append = "/var/lib/bacula/log" = all, !skipped
}

# Pool definitions
#
# Default Pool for jobs, but will hold no actual volumes
Pool {
  Name = Default
  Pool Type = Backup
}
Pool {
  Name = MondayPool
  Pool Type = Backup
  Recycle = yes
  AutoPrune = yes
  Volume Retention = 6d
  Accept Any Volume = yes
  Maximum Volume Jobs = 2
}
Pool {
  Name = TuesdayPool
  Pool Type = Backup
  Recycle = yes
  AutoPrune = yes
  Volume Retention = 6d
  Accept Any Volume = yes
  Maximum Volume Jobs = 2
}
Pool {
  Name = WednesdayPool
  Pool Type = Backup
  Recycle = yes
  AutoPrune = yes
  Volume Retention = 6d
  Accept Any Volume = yes
  Maximum Volume Jobs = 2
}
Pool {
  Name = ThursdayPool
  Pool Type = Backup
  Recycle = yes
  AutoPrune = yes
  Volume Retention = 6d
  Accept Any Volume = yes
```

```
  Maximum Volume Jobs = 2
}
Pool {
  Name = WeeklyPool
  Pool Type = Backup
  Recycle = yes
  AutoPrune = yes
  Volume Retention = 12d
  Accept Any Volume = yes
  Maximum Volume Jobs = 2
}
# EOF
```

Note, the mailcommand and operatorcommand should be on a single line each. They were split to preserve the proper page width. In order to get Bacula to release the tape after the nightly backup, he uses a **RunAfterJob** script that deletes the ASCII copy of the database back and then rewinds and ejects the tape. The following is a copy of **end_of_backup.sh**

```
#! /bin/sh
/usr/lib/bacula/delete_catalog_backup
mt rewind
mt eject
exit 0
```

Finally, if you list his Volumes, you get something like the following:

```
*list media
Using default Catalog name=MyCatalog DB=bacula
Pool: WeeklyPool
+-----+-----------+-------+--------+-----------+-----------------+-------+------+
| MeId| VolumeName| MedTyp| VolStat| VolBytes  | LastWritten     | VolRet| Recyc|
+-----+-----------+-------+--------+-----------+-----------------+-------+------+
| 5   | Friday_1  | MLR1  | Used   | 2157171998| 2003-07-11 20:20| 103680| 1    |
| 6   | Friday_2  | MLR1  | Append | 0         | 0               | 103680| 1    |
+-----+-----------+-------+--------+-----------+-----------------+-------+------+
Pool: MondayPool
+-----+-----------+-------+--------+-----------+-----------------+-------+------+
| MeId| VolumeName| MedTyp| VolStat| VolBytes  | LastWritten     | VolRet| Recyc|
+-----+-----------+-------+--------+-----------+-----------------+-------+------+
| 2   | Monday    | MLR1  | Used   | 2260942092| 2003-07-14 20:20| 518400| 1    |
+-----+-----------+-------+--------+-----------+-----------------+-------+------+
Pool: TuesdayPool
+-----+-----------+-------+--------+-----------+-----------------+-------+------+
| MeId| VolumeName| MedTyp| VolStat| VolBytes  | LastWritten     | VolRet| Recyc|
+-----+-----------+-------+--------+-----------+-----------------+-------+------+
| 3   | Tuesday   | MLR1  | Used   | 2268180300| 2003-07-15 20:20| 518400| 1    |
+-----+-----------+-------+--------+-----------+-----------------+-------+------+
Pool: WednesdayPool
```

```
+-----+-----------+-------+--------+-----------+-----------------+-------+------+
| MeId| VolumeName| MedTyp| VolStat| VolBytes  | LastWritten     | VolRet| Recyc|
+-----+-----------+-------+--------+-----------+-----------------+-------+------+
| 4   | Wednesday | MLR1  | Used   | 2138871127| 2003-07-09 20:2 | 518400| 1    |
+-----+-----------+-------+--------+-----------+-----------------+-------+------+
Pool: ThursdayPool
+-----+-----------+-------+--------+-----------+-----------------+-------+------+
| MeId| VolumeName| MedTyp| VolStat| VolBytes  | LastWritten     | VolRet| Recyc|
+-----+-----------+-------+--------+-----------+-----------------+-------+------+
| 1   | Thursday  | MLR1  | Used   | 2146276461| 2003-07-10 20:50| 518400| 1    |
+-----+-----------+-------+--------+-----------+-----------------+-------+------+
Pool: Default
No results to list.
```

Note, I have truncated a number of the columns so that the information fits on the width of a page.

Autochangers Support

Autochangers – General

Beginning with version 1.23, Bacula provides autochanger support for reading and writing tapes. In order to work with an autochanger, Bacula requires three things, each of which is explained in more detail after this list:

- A script that actually controls the autochanger according to commands sent by Bacula. We furnish such a script that works with **mtx** found in the **depkgs** distribution. This script works only with single drive autochangers.
- That each Volume (tape) to be used must be defined in the Catalog and have a Slot number assigned to it so that Bacula knows where the Volume is in the autochanger. This is generally done with the **label** command. See below for more details.
- Modifications to your Storage daemon's Device configuration resource to identify that the device is a changer, as well as a few other parameters.
- Optionally, you can modify your Storage resource definition in the Director's configuration file so that you are automatically prompted for the Slot when labeling a Volume.

Bacula uses its own **mtx-changer** script to interface with a program that actually does the tape changing. Thus in principle, **mtx-changer** can be adapted to function with any autochanger program. The current version of **mtx-changer** works with the **mtx** program. However, FreeBSD users have provided a script in the **examples** directory that allows Bacula to use the **chio** program.

As of version 1.30 and later, Bacula supports autochangers with barcode readers. This support includes two new Console commands: **label barcodes** and **update slots**. For more details on these commands, see the "Barcode Support" section below.

Current Bacula autochanger support does not include cleaning, stackers, or silos. However, under certain conditions, you may be able to make Bacula work with stackers (gravity feed and such). Bacula supports only single drive autochangers. Bacula does have code to operate multi-drive autochangers. However, the implementation is only partial. See below for more details.

In principle, if **mtx** will operate your changer correctly, then it is just a question of adapting the **mtx-changer** script (or selecting one already adapted) for proper interfacing. You can find a list of autochangers supported by **mtx** at the following link: http://mtx.badtux.net/compatibility.php. The home page for the **mtx** project can be found at: http://mtx.badtux.net/.

If you are having troubles, please use the **auto** command in the **btape** program to test the functioning of your autochanger with Bacula. When Bacula is running, please remember that for many distributions (e.g. FreeBSD, Debian, ...) the Storage daemon runs as **bacula.tape** rather than **root.root**, so you will need to ensure that the Storage daemon has sufficient permissions to access the autochanger.

Knowing What SCSI Devices You Have

Under Linux, you can

```
cat /proc/scsi/scsi
```

to see what SCSI devices you have available. You can also:

```
cat /proc/scsi/sg/device_hdr /proc/scsi/sg/devices
```

to find out how to specify their control address (**/dev/sg0** for the first, **/dev/sg1** for the second, ...) on the **Changer Device =** Bacula directive.

Under FreeBSD, you can use:

```
camcontrol devlist
```

To list the SCSI devices as well as the **/dev/passn** that you will use on the Bacula **Changer Device =** directive.

Example Scripts

Please read the sections below so that you understand how autochangers work with Bacula. Although we supply a default **mtx-changer** script, your autochanger may require some additional changes. If you want to

see examples of configuration files and scripts, please look in the <**bacula-src**>**/examples/devices** directory where you will find an example **HP-autoloader.conf** Bacula Device resource, and several **mtx-changer** scripts that have been modified to work with different autochangers.

Slots

To properly address autochangers, Bacula must know which Volume is in each **slot** of the autochanger. Slots are where the changer cartridges reside when not loaded into the drive. Bacula numbers these slots from one to the number of cartridges contained in the autochanger.

Bacula will not automatically use a Volume in your autochanger unless it is labeled and the slot number is stored in the catalog and the Volume is marked as InChanger. For each Volume in your changer, you will, using the Console program, assign a slot. This information is kept in **Bacula's** catalog database along with the other data for the volume. If no slot is given, or the slot is set to zero, Bacula will not attempt to use the autochanger even if all the necessary configuration records are present. In addition, the console **mount** command does not cause Bacula to operate the autochanger, it only tells Bacula to read any tape that may be in the drive.

You can check if the Slot number and InChanger flag are set by doing a:

```
list Volumes
```

in the Console program.

Multiple Devices

Some autochangers have more than one read/write device (drive). The current implementation has limited support for multiple devices by using the **Drive Index** directive in the Device resource of the Storage daemon's configuration file. Drive numbers or the Device Index are numbered beginning at zero, which is the default. To use the second Drive in an autochanger, you need to define a second Device resource and set the Drive Index to one for that device. In general, the second device will have the same **Changer Device** (control channel) as the first drive, but a different **Archive Device**.

The current implementation of Bacula does not coordinate between the two drives, so you must make sure that Bacula doesn't attempt to mount the

same Volume on both drives at the same time. There are a number of ways to do this. One was is to use different pools for each drive.

Worse than the above, the **mtx** program apparently does not prevent two accesses to the same control device at the same time, which means that if Bacula happens to attempt to call the mtx-changer script for two drives simultaneously, something will break.

A user supplied modified version of the mtx-changer script, which does locking to avoid this problem can be found in **examples/autochangers/locking-mtx-changer**. If you are using multiple drives, you will probably want to modify this script to work for you.

Device Configuration Records

Configuration of autochangers within Bacula is done in the Device resource of the Storage daemon. Four records: **Autochanger**, **Changer Device**, **Changer Command**, and **Maximum Changer Wait** control how Bacula uses the autochanger.

These four records, permitted in **Device** resources, are described in detail below:

Autochanger = *Yes—No* The **Autochanger** record specifies that the current device is or is not an autochanger. The default is **no**.

Changer Device = <**device-name**> In addition to the Archive Device name, you must specify a **Changer Device** name. This is because most autochangers are controlled through a different device than is used for reading and writing the cartridges. For example, on Linux, one normally uses the generic SCSI interface for controlling the autochanger, but the standard SCSI interface for reading and writing the tapes. On Linux, for the **Archive Device = /dev/nst0**, you would typically have **Changer Device = /dev/sg0**. Note, some of the more advanced autochangers will locate the changer device on **/dev/sg1**. Such devices typically have several drives and a large number of tapes.

On FreeBSD systems, the changer device will typically be on **/dev/pass0** through **/dev/passn**.

On Solaris, the changer device will typically be some file under **/dev/rdsk**.

Changer Command = <command> This record is used to specify the external program to call and what arguments to pass to it. The command is assumed to be a standard program or shell script that can be executed by the operating system. This command is invoked each time that Bacula wishes to manipulate the autochanger. The following substitutions are made in the **command** before it is sent to the operating system for execution:

```
%% = %
%a = archive device name
%c = changer device name
%d = changer drive index base 0
%f = Client's name
%j = Job name
%o = command  (loaded, load, or unload)
%s = Slot base 0
%S = Slot base 1
%v = Volume name
```

An actual example for using **mtx** with the **mtx-changer** script (part of the Bacula distribution) is:

```
Changer Command = "/etc/bacula/mtx-changer %c %o %S %a %d"
```

Where you will need to adapt the **/etc/bacula** to be the actual path on your system where the mtx-changer script resides. Details of the three commands currently used by Bacula (loaded, load, unload) as well as the output expected by Bacula are give in the **Bacula Autochanger Interface** section below.

Maximum Changer Wait = <time> This record is used to define the maximum amount of time that Bacula will wait for an autoloader to respond to a command (e.g. load). The default is set to 120 seconds. If you have a slow autoloader you may want to set it longer.

If the autoloader program fails to respond in this time, it will be killed and Bacula will request operator intervention.

Drive Index = <number> This record allows you to tell Bacula to use the second or subsequent drive in an autochanger with multiple drives. Since the drives are numbered from zero, the second drive is defined by

```
Device Index = 1
```

To use the second drive, you need a second Device resource definition in the Bacula configuration file. See the Multiple Drive section above in this chapter for more information.

An Example Configuration File

The following **Device** resource implements an autochanger:

```
Device {
  Name = "Autochanger"
  Media Type = DDS-4
  Archive Device = /dev/nst0    # Normal archive device
  Changer Device = /dev/sg0     # Generic SCSI device name
  Changer Command = "/etc/bacula/mtx-changer %c %o %S %a %d"
  Autochanger = yes
  LabelMedia = no;
  AutomaticMount = yes;
  AlwaysOpen = yes;
  Mount Anonymous Volumes = no;
}
```

where you will adapt the **Archive Device**, the **Changer Device**, and the path to the **Changer Command** to correspond to the values used on your system.

The above **Device** resource will work equally well for any standard tape drive (with device name **/dev/nst0**) since the extra autochanger commands will not be used unless a **slot** has been specified in the catalog record for the Volume to be used. See below for more details on the **slot**.

Specifying Slots When Labeling

If you add an **Autochanger = yes** record to the Storage resource in your Director's configuration file, the Bacula Console will automatically prompt you for the slot number and whether or not the Volume is in the changer when you **add** or **label** tapes for that Storage device. You must also set **Autochanger = yes** in the Device resource as we have described above in order for the autochanger to be used. Please see the Storage Resource in the Director's chapter and the Device Resource in the Storage daemon chapter for more details on these records.

Thus all stages of dealing with tapes can be totally automated. It is also

possible to set or change the Slot using the **update** command in the Console and selecting **Volume Parameters** to update.

Even though all the above configuration statements are specified and correct, Bacula will attempt to access the autochanger only if a **slot** is non-zero in the catalog Volume record (with the Volume name).

p¿If your autochanger has barcode labels, you can label all the Volumes in your autochanger one after another by using the **label barcodes** command. For each tape in the changer containing a barcode, Bacula will mount the tape and then label it with the same name as the barcode. An appropriate Media record will also be created in the catalog. Any barcode that begins with the same characters as specified on the "CleaningPrefix=xxx" command, will be treated as a cleaning tape, and will not be labeled. For example with:

```
Pool {
  Name ...
  Cleaning Prefix = "CLN"
}
```

Any slot containing a barcode of CLNxxxx will be treated as a cleaning tape and will not be mounted.

Dealing with Multiple Magazines

If you have several magazines or if you insert or remove cartridges from a magazine, you will need to notify Bacula of this. By doing so, Bacula will as a preference, use Volumes that it knows to be in the autochanger before accessing Volumes that are not in the autochanger. This prevents unneeded operator intervention.

If your autochanger has barcodes (machine readable tape labels), the task of informing Bacula is simple. Every time, you change a magazine, or add or remove a cartridge from the magazine, simply do

```
update slots
```

in the Console program. This will cause Bacula to request the autochanger to return the current Volume names in the magazine. This will be done without actually accessing or reading the Volumes because the barcode reader does this during inventory when the autochanger is first turned on.

Bacula will ensure that any Volumes that are currently marked as being in the magazine are marked as no longer in the magazine, and the new list of Volumes will be marked as being in the magazine. In addition, the Slot numbers of the Volumes will be corrected in Bacula's catalog if they are incorrect (added or moved).

If you do not have a barcode reader on your autochanger, you have several alternatives.

1. You can manually set the Slot and InChanger flag using the **update volume** command in the Console (quite painful).
2. You can issue a

   ```
   update slots scan
   ```

   command that will cause Bacula to read the label on each of the cartridges in the magazine in turn and update the information (Slot, InChanger flag) in the catalog. This is quite effective but does take time to load each cartridge into the drive in turn and read the Volume label.
3. You can modify the mtx-changer script so that it simulates an autochanger with barcodes. See below for more details.

Simulating Barcodes in your Autochanger

You can simulate barcodes in your autochanger by making the **mtx-changer** script return the same information that an autochanger with barcodes would do. This is done by commenting out the one and only line in the **list)** case, which is:

```
  ${MTX} -f $ctl status | grep " *Storage Element [0-9]*:.*Full" | awk "{print \$3 \$4}" | sed "
```

at approximately line 99 by putting a # in column one of that line, or by simply deleting it. Then in its place add a new line that prints the contents of a file. For example:

```
cat /etc/bacula/changer.volumes
```

Be sure to include a full path to the file, which can have any name. The contents of the file must be of the following format:

```
1:Volume1
2:Volume2
3:Volume3
...
```

Where the 1, 2, 3 are the slot numbers and Volume1, Volume2, ... are the Volume names in those slots. You can have multiple files that represent the Volumes in different magazines, and when you change magazines, simply copy the contents of the correct file into your **/etc/bacula/changer.volumes** file. There is no need to stop and start Bacula when you change magazines, simply put the correct data in the file, then run the **update slots** command, and your autochanger will appear to Bacula to be an autochanger with barcodes.

The Full Form of the Update Slots Command

If you change only one cartridge in the magazine, you may not want to scan all Volumes, so the **update slots** command (as well as the **update slots scan** command) has the additional form:

```
update slots=n1,n2,n3-n4, ...
```

where the keyword **scan** can be appended or not. The n1,n2, ... represent Slot numbers to be updated and the form n3-n4 represents a range of Slot numbers to be updates (e.g. 4-7 will update Slots 4,5,6, and 7).

This form is particularly useful if you want to do a scan (time expensive) and restrict the update to one or two slots.

For example, the command:

```
update slots=1,6 scan
```

will cause Bacula to load the Volume in Slot 1, read its Volume label and update the Catalog. It will do the same for the Volume in Slot 6. The command:

```
update slots=1-3,6
```

will read the barcoded Volume names for slots 1,2,3 and 6 and make the appropriate updates in the Catalog. If you don't have a barcode reader or have not modified the mtx-changer script as described above, the above command will not find any Volume names so will do nothing.

FreeBSD Issues

If you are having problems on FreeBSD when Bacula tries to select a tape, and the message is **Device not configured**, this is because FreeBSD has made the tape device **/dev/nsa1** disappear when there is no tape mounted in the autochanger slot. As a consequence, Bacula is unable to open the device. The solution to the problem is to make sure that some tape is loaded into the tape drive before starting Bacula. This problem is corrected in Bacula versions 1.32f-5 and later.

Please see the Tape Testing chapter of this manual for **important** information concerning your tape drive before doing the autochanger testing.

Testing the Autochanger and Adapting Your mtx-changer Script

Before attempting to use the autochanger with Bacula, it is preferable to "hand-test" that the changer works. To do so, we suggest you do the following commands (assuming that the **mtx-changer** script is installed in **/etc/bacula/mtx-changer**):

Make sure Bacula is not running.

/etc/bacula/mtx-changer /dev/sg0 list 0 /dev/nst0 0 This command should print:

```
1:
2:
3:
...
```

or one number per line for each slot that is occupied in your changer, and the number should be terminated by a colon (**:**). If your changer has barcodes, the barcode will follow the colon. If an error message is printed, you must resolve the problem (e.g. try a different SCSI

control device name if **/dev/sg0** is incorrect. For example, on FreeBSD systems, the autochanger SCSI control device is generally **/dev/pass2**.

/etc/bacula/mtx-changer /dev/sg0 slots 0 /dev/nst0 0 This command should return the number of slots in your autochanger.

/etc/bacula/mtx-changer /dev/sg0 unload If a tape is loaded, this should cause it to be unloaded.

/etc/bacula/mtx-changer /dev/sg0 load 3 /dev/nst0 0 Assuming you have a tape in slot 3, it will be loaded into the read slot (0).

/etc/bacula/mtx-changer /dev/sg0 loaded 0 /dev/nst0 0 It should print "3"

/etc/bacula/mtx-changer /dev/sg0 unload

Once all the above commands work correctly, assuming that you have the right **Changer Command** in your configuration, Bacula should be able to operate the changer. The only remaining area of problems will be if your autoloader needs some time to get the tape loaded after issuing the command. After the **mtx-changer** script returns, Bacula will immediately rewind and read the tape. If Bacula gets rewind I/O errors after a tape change, you will probably need to insert a **sleep 20** after the **mtx** command, but be careful to exit the script with a zero status by adding **exit 0** after any additional commands you add to the script. This is because Bacula checks the return status of the script, which should be zero if all went well.

You can test whether or not you need a **sleep** by putting the following commands into a file and running it as a script:

```
#!/bin/sh
/etc/bacula/mtx-changer /dev/sg0 unload
/etc/bacula/mtx-changer /dev/sg0 load 3
mt -f /dev/st0 rewind
mt -f /dev/st0 weof
```

If the above script runs, you probably have no timing problems. If it does not run, start by putting a **sleep 30** or possibly a **sleep 60** in the the script just after the mtx-changer load command. If that works, then you should move the sleep into the actual **mtx-changer** script so that it will be effective when Bacula runs.

A second problem that comes up with a small number of autochangers is that they need to have the cartridge ejected before it can be removed. If this is the case, the **load 3** will never succeed regardless of how long you wait. If this seems to be your problem, you can insert an eject just after the unload so that the script looks like:

```
#!/bin/sh
/etc/bacula/mtx-changer /dev/sg0 unload
mt -f /dev/st0 offline
/etc/bacula/mtx-changer /dev/sg0 load 3
mt -f /dev/st0 rewind
mt -f /dev/st0 weof
```

Obviously, if you need the **offline** command, you should move it into the mtx-changer script ensuring that you save the status of the **mtx** command or always force an **exit 0** from the script, because Bacula checks the return status of the script.

As noted earlier, there are several scripts in **<bacula-source>/examples/devices** that implement the above features, so they may be a help to you in getting your script to work.

If Bacula complains "Rewind error on /dev/nst0. ERR=Input/output error." you most likely need more sleep time in your **mtx-changer** before returning to Bacula after a load command has been completed.

Using the Autochanger

Let's assume that you have properly defined the necessary Storage daemon Device records, and you have added the **Autochanger = yes** record to the Storage resource in your Director's configuration file.

Now you fill your autochanger with say six blank tapes.

What do you do to make Bacula access those tapes?

One strategy is to prelabel each of the tapes. Do so by starting Bacula, then with the Console program, enter the **label** command:

```
./console
Connecting to Director rufus:8101
1000 OK: rufus-dir Version: 1.26 (4 October 2002)
*{\bf label}
```

it will then print something like:

```
Using default Catalog name=BackupDB DB=bacula
The defined Storage resources are:
     1: Autochanger
     2: File
Select Storage resource (1-2): {\bf 1}
```

I select the autochanger (1), and it prints:

```
Enter new Volume name: {\bf TestVolume1}
Enter slot (0 for none): {\bf 1}
```

where I entered **TestVolume1** for the tape name, and slot **1** for the slot. It then asks:

```
Defined Pools:
     1: Default
     2: File
Select the Pool (1-2): {\bf 1}
```

I select the Default pool. This will be automatically done if you only have a single pool, then Bacula will proceed to unload any loaded volume, load the volume in slot 1 and label it. In this example, nothing was in the drive, so it printed:

```
Connecting to Storage daemon Autochanger at localhost:9103 ...
Sending label command ...
3903 Issuing autochanger "load slot 1" command.
3000 OK label. Volume=TestVolume1 Device=/dev/nst0
Media record for Volume=TestVolume1 successfully created.
Requesting mount Autochanger ...
3001 Device /dev/nst0 is mounted with Volume TestVolume1
You have messages.
*
```

You may then proceed to label the other volumes. The messages will change slightly because Bacula will unload the volume (just labeled TestVolume1) before loading the next volume to be labeled.

Once all your Volumes are labeled, Bacula will automatically load them as they are needed.

To "see" how you have labeled your Volumes, simply enter the **list volumes** command from the Console program, which should print something like the following:

```
*{\bf list volumes}
Using default Catalog name=BackupDB DB=bacula
Defined Pools:
     1: Default
     2: File
Select the Pool (1-2): {\bf 1}
+-------+----------+--------+---------+-------+--------+----------+-------+------+
| MedId | VolName  | MedTyp | VolStat | Bites | LstWrt | VolReten | Recyc | Slot |
+-------+----------+--------+---------+-------+--------+----------+-------+------+
| 1     | TestVol1 | DDS-4  | Append  | 0     | 0      | 30672000 | 0     | 1    |
| 2     | TestVol2 | DDS-4  | Append  | 0     | 0      | 30672000 | 0     | 2    |
| 3     | TestVol3 | DDS-4  | Append  | 0     | 0      | 30672000 | 0     | 3    |
| ...                                                                          |
+-------+----------+--------+---------+-------+--------+----------+-------+------+
```

Barcode Support

Bacula provides barcode support with two Console commands, **label barcodes** and **update slots**.

The **label barcodes** will cause Bacula to read the barcodes of all the cassettes that are currently installed in the magazine (cassette holder) using the **mtx-changer list** command. Each cassette is mounted in turn and labeled with the same Volume name as the barcode.

The **update slots** command will first obtain the list of cassettes and their barcodes from **mtx-changer**. Then it will find each volume in turn in the catalog database corresponding to to the barcodes and set its Slot to correspond to the value just read. If the Volume is not in the catalog, then nothing will be done. This command is useful for synchronizing Bacula with the current magazine in case you have changed magazines or in case you have moved cassettes from one slot to another.

The **Cleaning Prefix** statement can be used in the Pool resource to define a Volume name prefix, which if it matches that of the Volume (barcode) will cause that Volume to be marked with a VolStatus of **Cleaning**. This will prevent Bacula from attempting to write on the Volume.

Bacula Autochanger Interface

Bacula calls the autochanger script that you specify on the **Changer Device** statement. Normally this script will be the **mtx-changer** script that we can provide, but it can in fact be any program. The only requirements are that the "commands" that Bacula uses are **loaded**, **load**,

unload, **list**, and (**slots**. In addition, each of those commands must return the information in the precise format as specified below:

```
- Currently the changer commands used are:
    loaded -- returns number of the slot that is loaded in
              the drive or 0 if the drive is empty.
    load   -- loads a specified slot (note, some autochangers
              require a 30 second pause after this command) into
              the drive.
    unload -- unloads the device (returns cassette to its slot).
    list   -- returns one line for each cassette in the autochanger
              in the format <slot>:<barcode>. Where
              the {\bf slot} is the non-zero integer representing
              the slot number, and {\bf barcode} is the barcode
              associated with the cassette if it exists and if you
              autoloader supports barcodes. Otherwise the barcode
              field is blank.
    slots  -- returns total number of slots in the autochanger.
```

Bacula checks the exit status of the program called, and if it is zero, the data is accepted. If the exit status is non-zero, Bacula ignores any information returned and treats the drive as if it is not an autochanger.

Supported Autochangers

Supported Autochanger Models

I hesitate to call these "supported" autochangers because the only autochanger that I have in my possession and am able to test is the HP SureStore DAT40X6. All the other autochangers have been reported to work by Bacula users. Note, in the Capacity/Slot column below, I quote the Compressed capacity per tape (or Slot).

OS	Man.	Media	Model	Slots	Cap/Slot
Linux	Adic	LTO-1/2, SDLT 320	Adic Scalar 24	24	100GB
Linux	Adic	LTO-2	Adic FastStor 2, Sun Storedge L8	8	200GB
-	CA-VM	??	Tape	??	??
-	Dell	LTO-2	PowerVault 132T/136T	-	100GB
-	DFSMS	??	VM RMM	-	??
z/VM	IBM	??	IBM Tape Manager	-	??
z/VM	IBM	??	native tape	-	??
Linux	Exabyte	VXA2	VXA PacketLoader 1x10 2U	10	80/160GB
Linux Gentoo 1.4	Exabyte	AIT-2	215A	15 (2 drives)	50GB
Linux	HP	DDS-4	SureStore DAT-40X6	6	40GB
Linux	HP	Ultrium-2/LTO	MSL 6000/ 60030/ 5052	28	200/400GB
-	HP	DLT	A4853 DLT	30	40/70GB
Linux	HP (Compaq)	DLT VI	Compaq TL-895	96+4 import export	35/70GB

SuSE 9.0	IBM	LTO	IBM 3581 Ultrium Tape Loader	7	200/400GB
-	Overland	LTO	Overland LoaderXpress LTO	10-19	100GB
-	Overland	LTO	Overland Neo2000 LTO	26-30	100GB
-	Quantum	??	Super Loader	??	??
FreeBSD 4.9	QUALSTAR TLS-4210 (Qualstar)	AIT1: 36GB, AIT2: 50GB all uncomp	QUALSTAR TLS-4210	12	AIT1: 36GB, AIT2: 50GB all uncomp
Linux	Skydata	DLT	ATL-L200	8	40/80
-	Sony	DDS-4	TSL-11000	8	40GB
Linux	Sony	AIT-2	LIB-304(SDX-500C)	?	200GB
FreeBSD 4.9-STABLE	Sony	AIT-1	TSL-SA300C	4	45/70GB
-	Storagetek	DLT	Timberwolf DLT	6	40/70
-	Storagetek	??	ACSLS	??	??
Solaris	Sun	4mm DLT	Sun Desktop Archive Python 29279	4	20GB
Linux	Tandberg	DLT VI	VS 640	8?	35/70GB
Linux 2.6.x	Tandberg Data	SLR100	SLR100 Autoloader	8	50/100GB

Data Spooling

Bacula allows you to specify that you want the Storage daemon to initially write your data to disk and then subsequently to tape. This serves several important purposes.

- It can take a long time for data to come in from the File daemon during an Incremental backup. If it is directly written to tape, the tape will start and stop or shoe-shine as it is often called causing tape wear. By first writing the data to disk, then writing it to tape, the tape can be kept in continual motion.
- While the spooled data is being written to the tape, the despooling process has exclusive use of the tape. This means that you can spool multiple simultaneous jobs to disk, then have them very efficiently despooled one at a time without having the data blocks from several jobs intermingled, thus substantially improving the time needed to restore files.
- Writing to a tape can be slow. By first spooling your data to disk, you can often reduce the time the File daemon is running on a system, thus reducing downtime, and/or interference with users.

Data spooling is exactly that "spooling". It is not a way to first write a "backup" to a disk file and then to a tape. When the backup spooled to disk, it is not complete and cannot be restored until it is written to tape. In a future version, Bacula will support writing a backup to disk then later **Migrating** or **Copying** it to a tape.

The remainder of this chapter explains the various directives that you can use in the spooling process.

Data Spooling Directives

The following directives can be used to control data spooling.

- To turn data spooling on/off at the Job level in the Job resource in the Director's conf file (default **no**).

  SpoolData = yes—no

- To override the Job specification in a Schedule Run directive in the Director's conf file.

SpoolData = yes—no

- To limit the maximum total size of the spooled data for a particular device. Specified in the Device resource of the Storage daemon's conf file (default unlimited).

  Maximum Spool Size = size

  Where size is a the maximum spool size for all jobs specified in bytes.

- To limit the maximum total size of the spooled data for a particular device for a single job. Specified in the Device Resource of the Storage daemon's conf file (default unlimited).

  Maximum Job Spool Size = size

  Where size is the maximum spool file size for a single job specified in bytes.

- To specify the spool directory for a particular device. Specified in the Device Resource of the Storage daemon's conf file (default, the working directory).

  Spool Directory = directory

!!! MAJOR WARNING !!!

Please be very careful to exclude the spool directory from any backup, otherwise, your job will write enormous amounts of data to the Volume, and most probably terminate in error. This is because in attempting to backup the spool file, the backup data will be written a second time to the spool file, and so on ad infinum.

Another advice is to always specify the maximum spool size so that your disk doesn't completely fill up. In principle, data spooling will properly detect a full disk, and despool data allowing the job to continue. However, attribute spooling is not so kind to the user. If the disk on which attributes are being spooled fills, the job will be canceled.

Other Points

- When data spooling is enabled, Bacula automatically turns on attribute spooling. In other words, it also spools the catalog entries to disk. This is done so that in case the job fails, there will be no catalog entries pointing to non-existent tape backups.

- Attribute despooling is done at the end of the job, as a consequence, after Bacula stops writing the data to the tape, there may be a pause while the attributes are sent to the Directory and entered into the catalog before the job terminates.
- Attribute spool files are always placed in the working directory.
- When Bacula begins despooling data spooled to disk, it takes exclusive use of the tape. This has the major advantage that in running multiple simultaneous jobs at the same time, the blocks of several jobs will not be intermingled.
- It probably does not make a lot of sense to enable data spooling if you are writing to disk files.
- It is probably best to provide as large a spool file as possible to avoid repeatedly spooling/despooling. Also, while a job is despooling to tape, the File daemon must wait (i.e. spooling stops for the job while it is despooling).
- If you are running multiple simultaneous jobs, Bacula will continue spooling other jobs while one is despooling to tape, provided there is sufficient spool file space.

Bacula Frequently Asked Questions

See the bugs section of this document for a list of known bugs and solutions.

What is Bacula? **Bacula** is a network backup and restore program.

Does Bacula support Windows? Yes, Bacula compiles and runs on Windows machines (Win98, WinMe, WinXP, WinNT, and Win2000). We provide a binary version of the Client (bacula-fd), but have not tested the Director nor the Storage daemon. Note, Win95 is no longer supported because it doesn't have the GetFileAttributesExA API call.

What language is Bacula written in? It is written in C++, but it is mostly C code using only a limited set of the C++ extensions over C. Thus Bacula is completely compiled using the C++ compiler. There are several modules, including the Win32 interface that are written using the object oriented C++ features. Over time, we are slowly adding a larger subset of C++.

On what machines does Bacula run? **Bacula** builds and executes on RedHat Linux (versions RH7.1-RHEL 3.0, SuSE, Gentoo, Debian, Mandrake, ...), FreeBSD, Solaris, Alpha, SGI (client), NetBSD, OpenBSD, Mac OS X (client), and Win32 (client).

Bacula has been my only backup tool for over four years backing up 5 machines nightly (3 Linux boxes running RedHat, a WinXP machine, and a WinNT machine).

Is Bacula Stable? Yes, it is remarkably stable, but remember, there are still a lot of unimplemented or partially implemented features. With a program of this size (100,000+ lines of C++ code not including the SQL programs) there are bound to be bugs. The current test environment (a twisted pair local network and a HP DLT backup tape) is rather ideal, so additional testing on other sites is necessary. The File daemon has never crashed – running months at a time with no intervention. The Storage daemon is remarkably stable with most of the problems arising during labeling or switching tapes. Storage daemon crashes are rare. The Director, given the multitude of functions it fulfills is also relatively stable. In a production environment, it rarely if ever crashes. Of the three daemons, the Director is the most prone to having problems. It frequently runs several months with no problems.

There are a number of reasons for this stability.

1. The program was largely written by one person to date (Kern).

2. The program constantly is checking the chain of allocated memory buffers to ensure that no overruns have occurred.

3. All memory leaks (orphaned buffers) are reported each time the program terminates.

4. Any signal (segmentation fault, ...) generates a traceback that is emailed to the developer. This permits quick resolution of bugs even if they only show up rarely in a production system.

5. There is a reasonably comprehensive set of regression tests that avoids re-creating the most common errors in new versions of Bacula.

I'm Getting Authorization Errors. What is Going On? For security reasons, Bacula requires that both the File daemon and the Storage daemon know the name of the Director as well as his password. As a consequence, if you change the Director's name or password, you must make the corresponding change in the Storage daemon and in the File daemon configuration files.

During the authorization process, the Storage daemon and File daemon also require that the Director authenticate itself, so both ends require the other to have the correct name and password.

If you have edited the conf files and modified any name or any password, and you are getting authentication errors, then your best bet is to go back to the original conf files generated by the Bacula installation process. Make only the absolutely necessary modifications to these files – e.g. add the correct email address. Then follow the instructions in the Running Bacula chapter of this manual. You will run a backup to disk and a restore. Only when that works, should you begin customization of the conf files.

Another reason that you can get authentication errors is if you are running Multiple Concurrent Jobs in the Director, but you have not set them in the File daemon or the Storage daemon. Once you reach their limit, they will reject the connection producing authentication (or connection) errors.

If you are having problems connecting to a Windows machine that previously worked, you might try restarting the Bacula service since Windows frequently encounters networking connection problems.

Here is sort of a picture of what names/passwords in which files/Resources must match up:

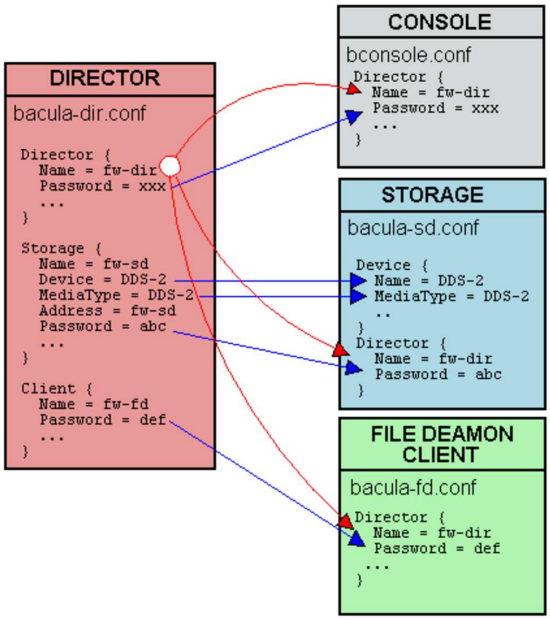

In the left column, you will find the Director, Storage, and Client resources, with their names and passwords – these are all in **bacula-dir.conf**. In the right column are where the corresponding values should be found in the Console, Storage daemon (SD), and File daemon (FD) configuration files.

Bacula Runs Fine but Cannot Access a Client on a Different Machine. Why?

There are several reasons why Bacula could not contact a client on a different machine. They are:

- It is a Windows Client, and the client died because of an improper configuration file. Check that the Bacula icon is in the system tray and the the menu items work. If the client has died, the icon will disappear only when you move the mouse over the icon.
- The Client address or port is incorrect or not resolved by DNS. See if you can ping the client machine using the same address as in the Client record.
- You have a firewall, and it is blocking traffic on port 9102 between the Director's machine and the Clients machine (or on port 9103 between the Client and the Storage daemon machines).
- Your password or names are not correct in both the Director and the Client machine. Try configuring everything identical to how

you run the client on the same machine as the Director, but just change the Address. If that works, make the other changes one step at a time until it works.

My Catalog is Full of Test Runs, How Can I Start Over? If you are using MySQL do the following:

```
cd <bacula-source>/src/cats
./drop_mysql_tables
./make_mysql_tables
```

If you are using SQLite, do the following:

```
Delete bacula.db from your working directory.
cd <bacula-source>/src/cats
./drop_sqlite_tables
./make_sqlite_tables
```

Then write an EOF on each tape you used with **Bacula** using:

```
mt -f /dev/st0 rewind
mt -f /dev/st0 weof
```

where you need to adjust the device name for your system.

I Run a Restore Job and Bacula Hangs. What do I do? On Bacula version 1.25 and prior, it expects you to have the correct tape mounted prior to a restore. On Bacula version 1.26 and higher, it will ask you for the tape, and if the wrong one it mounted, it will inform you.

If you have previously done an **unmount** command, all Storage daemon sessions (jobs) will be completely blocked from using the drive unmounted, so be sure to do a **mount** after your unmount. If in doubt, do a second **mount**, it won't cause any harm.

I Cannot Get My Windows Client to Start Automatically? You are probably having one of two problems: either the Client is dying due to an incorrect configuration file, or you didn't do the Installation commands necessary to install it as a Windows Service.

For the first problem, see the next FAQ question. For the second problem, please review the Windows Installation instructions in this manual.

My Windows Client Immediately Dies When I Start It The most common problem is either that the configuration file is not where it expects it to be, or that there is an error in the configuration file. You must have the configuration file in **c:\bacula\bin\bacula-fd.conf**.

To **see** what is going on when the File daemon starts on Windows, do the following:

```
Start a DOS shell Window.
cd c:\bacula\bin
bacula-fd -d100 -c c:\bacula\bin\bacula-fd.conf
```

This will cause the FD to write a file **bacula.trace** in the current directory, which you can examine and determine the problem.

When I Start the Console, the Error Messages Fly By. How can I see them? Either use a shell window with a scroll bar, or use the gnome-console. In any case, you probably should be logging all output to a file, and then you can simply view the file using an editor or the **less** program. To log all output, I have the following in my Director's Message resource definition:

```
append = "/home/kern/bacula/bin/log" = all, !skipped
```

Obviously you will want to change the filename to be appropriate for your system.

I didn't realize that the backups were not working on my Windows Client. What should I You should be sending yourself an email message for each job. This will avoid the possibility of not knowing about a failed backup. To do so put something like:

```
Mail = yourname@yourdomain = all, !skipped
```

in your Director's message resource. You should then receive one email for each Job that ran. When you are comfortable with what is going on (it took me 9 months), you might change that to:

```
MailOnError = yourname@yourdomain = all, !skipped
```

then you only get email messages when a Job errors as is the case for your Windows machine.

You should also be logging the Director's messages, please see the previous FAQ for how to do so.

All my Jobs are scheduled for the same time. Will this cause problems? No, not at all. Bacula will schedule all the Jobs at the same time, but will run them one after another unless you have increased the number of simultaneous jobs in the configuration files for the Director, the File daemon, and the Storage daemon. The appropriate configuration record is **Maximum Concurrent Jobs = nn**. At the current time, we recommend that you leave this set to **1** for the Director.

Can Bacula Backup My System To Files instead of Tape? Yes, in principle, Bacula can backup to any storage medium as long as you have correctly defined that medium in the Storage daemon's Device resource. For an example of how to backup to files, please see the Pruning Example in the Recycling chapter of this manual. Also, there is a whole chapter devoted to Backing Up to Disk.

Can Bacula Backup and Restore Files Greater than 2 Giga bytes in Size? If your operating system permits it, and you are running Bacula version 1.26 or later, the answer is yes. To the best of our knowledge all client system supported by Bacula can handle files larger than 2 Giga bytes.

I Started A Job then Decided I Really Did Not Want to Run It. Is there a better Yes, you normally should use the Console command **cancel** to cancel a Job that is either scheduled or running. If the Job is scheduled, it will be marked for cancellation and will be canceled when it is scheduled to start. If it is running, it will normally terminate after a few minutes. If the Job is waiting on a tape mount, you may need to do a **mount** command before it will be canceled.

Why have You Trademarked the Name Bacula®? We have trademarked the name Bacula to ensure that all media written by any program named Bacula will always be compatible. Anyone may use the name Bacula, even in a derivative product as long as it remains totally compatible in all respects with the program defined here.

Why is Your Online Document for Version 1.35 of Bacula when the Currently Rel As Bacula is being developed, the document is also being enhanced, more often than not it has clarifications of existing features that can be very useful to our users, so we publish the very latest document. Fortunately it is rare that there are confusions with new features.

If you want to read a document that pertains only to a specific version, please use the one distributed in the source code.

How Can I Be Sure that Bacula Really Saves and Restores All Files? It is really quite simple, but took me awhile to figure out

how to "prove" it. First make a Bacula Rescue disk, see the Disaster Recovery Using Bacula of this manual. Second, you run a full backup of all your files on all partitions. Third, you run an Verify InitCatalog Job on the same FileSet, which effectively makes a record of all the files on your system. Fourth, you run a Verify Catalog job and assure yourself that nothing has changed (well, between an InitCatalog and Catalog one doesn't expect anything). Then do the unthinkable, write zeros on your MBR (master boot record) wiping out your hard disk. Now, restore your whole system using your Bacula Rescue disk and the Full backup you made, and finally re-run the Verify Catalog job. You will see that with the exception of the directory modification and access dates and the files changed during the boot, your system is identical to what it was before you wiped your hard disk.

I did a Full backup last week, but now in running an Incremental, Bacula says it did not fi

Before doing an Incremental or a Differential backup, Bacula checks to see if there was a prior Full backup of the same Job that terminated successfully. If so, it uses the date that full backup started as the time for comparing if files have changed. If Bacula does not find a successfully full backup, it proceeds to do one. Perhaps you canceled the full backup, or it terminated in error. In such cases, the full backup will not be successful. You can check by entering **list jobs** and look to see if there is a prior Job with the same Name that has Level F and JobStatus T (normal termination).

Another reason why Bacula may not find a suitable Full backup is that every time you change the FileSet, Bacula will require a new Full backup. This is necessary to ensure that all files are properly backed up in the case where you have added more files to the FileSet. Beginning with version 1.31, the FileSets are also dated when they are created, and this date is displayed with the name when you are listing or selecting a FileSet. For more on backup levels see below.

How Can You Claim to Handle Unlimited Path and Filename Lengths when All Other Pro

Most of those other programs have been around for a long time, in fact since the beginning of Unix, which means that they were designed for rather small fixed length path and filename lengths. Over the years, these restrictions have been relaxed allowing longer names. Bacula on the other hand was designed in 2000, and so from the start, Path and Filenames have been keep in buffers that start at 256 bytes in length but can grow as needed to handle any length. Most of the work is carried out by lower level routines making the coding rather easy.

What Is the Really Unique Feature of Bacula? Well, it is hard to come up with unique features when backup programs for Unix machines have been around since the 1960s. That said, I believe that Bacula is the first and only program to use a standard SQL interface to its catalog database. Although this adds a bit of complexity and possibly overhead, it provides an amazingly rich set of features that are easy to program and enhance. The current code has barely scratched the surface in this regard (version 1.31).

The second feature, which gives a lot of power and flexibility to Bacula is the Bootstrap record definition.

The third unique feature, which is currently (1.30) unimplemented, and thus can be called vaporware :-), is Base level saves. When implemented, this will enormously reduce tape usage.

If I Do Run Multiple Simultaneous Jobs, How Can I Force One Particular Job to

Yes, you can set Priorities on your jobs so that they run in the order you specify. Please see: the Priority record in the Job resource.

] The most common problem is that you have not specified a fully qualified email address and your bsmtp server is rejecting the mail. The next most common problem is that your bsmtp server doesn't like the syntax on the From part of the message. For more details on this and other problems, please see the Getting Email Notification to Work section of the Tips chapter of this manual. The section Getting Notified of Job Completion of the Tips chapter may also be useful. For more information on the **bsmtp** mail program, please see bsmtp in the Volume Utility Tools chapter of this manual.

I Change Recycling, Retention Periods, or File Sizes in my Pool Resource and the

The different variables associated with a Pool are defined in the Pool Resource, but are actually read by Bacula from the Catalog database. On Bacula versions prior to 1.30, after changing your Pool Resource, you must manually update the corresponding values in the Catalog by using the **update pool** command in the Console program. In Bacula version 1.30, Bacula does this for you automatically every time it starts.

When Bacula creates a Media record (Volume), it uses many default values from the Pool record. If you subsequently change the Pool record, the new values will be used as a default for the next Volume that is created, but if you want the new values to apply to existing Volumes, you must manually update the Volume Catalog entry using the **update volume** command in the Console program.

I Have Configured Compression On, But None of My Files Are Compressed. Why?

There are two kinds of compression. One is tape compression. This is done by the tape drive hardware, and you either enable or disable it with system tools such as **mt**. This compression works independently of Bacula.

Bacula also has compression code, which is normally used only when backing up to file Volumes. There are two conditions for this "software" to be enabled.

1. You must have the zip development libraries loaded on your system when building Bacula and Bacula must find this library, normally **/usr/lib/libz.a**. On RedHat systems, this library is provided by the **zlib-devel** rpm.
   If the library is found by Bacula during the **./configure** it will be in dicated on the **config.out** line by:

   ```
   ZLIB support:  yes
   ```

2. You must add the **compression=gzip** option on your Include statement in the Director's configuration file.

Bacula is Asking for a New Tape After 2 GB of Data but My Tape holds 33 GB. Why?

There are several reasons why Bacula will request a new tape.

- There is an I/O error on the tape. Bacula prints an error message and requests a new tape. Bacula does not attempt to continue writing after an I/O error.
- Bacula encounters and end of medium on the tape. This is not always distinguishable from an I/O error.
- You have specifically set some size limitation on the tape. For example the **Maximum Volume Bytes** or **Maximum Volume Files** in the Director's Pool resource, or **Maximum Volume Size** in the Storage daemon's Device resource.

Bacula is Not Doing the Right Thing When I Request an Incremental Backup. Why?

As explained in one of the previous questions, Bacula will automatically upgrade an Incremental or Differential job to a Full backup if it cannot find a prior Full backup or a suitable Full backup. For the gory details on how/when Bacula decides to upgrade levels please see the Level record in the Director's configuration chapter of this manual.

If after reading the above mentioned section, you believe that Bacula is not correctly handling the level (Differential/Incremental), please send us the following information for analysis:

- Your Director's configuration file.
- The output from **list jobs** covering the period where you are having the problem.
- The Job report output from the prior Full save (not critical).
- An **llist jobid=nnn** where nnn is the JobId of the prior Full save.
- The Job report output from the save that is doing the wrong thing (not critical).
- An **llist jobid=nnn** where nnn is the JobId of the job that was not correct.
- An explanation of what job went wrong and why you think it did.

The above information can allow us to analyze what happened, without it, there is not much we can do.

I am Backing Up an Offsite Machine with an Unreliable Connection. The Director

Bacula was written on the assumption that it will have a good TCP/IP connection between all the daemons. As a consequence, the current Bacula doesn't deal with faulty connection very well. This situation is slowly being corrected over time.

There are several things you can do to improve the situation.

- Upgrade to version 1.32 and use the new SDConnectTimeout record. For example, set:

```
SD Connect Timeout = 5 min
```

  in the FileDaemon resource.
- Run these kinds of jobs after all other jobs.

When I ssh into a machine and start Bacula then attempt to exit, ssh hangs forev

This happens because Bacula leaves stdin, stdout, and stderr open for debug purposes. To avoid it, the simplest thing to do is to redirect the output of those files to **/dev/null** or another file in your startup script (the RedHat autostart scripts do this automatically). For example, you start the Director with:

```
bacula-dir -c bacula-dir.conf ... 0>\&1 2>\&1 >/dev/null
```

and likewise for the other daemons.

I'm confused by the different Retention periods: File Retention, Job Retention, Volume Re

Yes, this certainly can be confusing. The basic reason for so many is to allow flexibility. The File records take quite a lot of space in the catalog, so they are typically records you want to remove rather quickly. The Job records, take very little space, and they can be useful even without the File records to see what Jobs actually ran and when. One must understand that if the File records are removed from the catalog, you cannot use the **restore** command to restore an individual file since Bacula no longer knows where it is. However, as long as the Volume Retention period has not expired, the data will still be on the tape, and can be recovered from the tape.

For example, I keep a 30 day retention period for my Files to keep my catalog from getting too big, but I keep my tapes for a minimum of one year, just in case.

Why Does Bacula Ignore the MaxVolumeSize Set in my Pool?

The MaxVolumeSize that Bacula uses comes from the Media record, so most likely you changed your Pool, which is used as the default for creating Media records, **after** you created your Volume. Check what is in the Media record by doing:

```
llist Volume=xxx
```

If it doesn't have the right value, you can use:

```
update Volume=xxx
```

to change it.

In connecting to my Client, I get "ERR:Connection Refused. Packet Size too big from File

This is typically a communications error resulting from one of the following:

- Old versions of Bacula, usually a Win32 client, where two threads were using the same I/O packet. Fixed in more recent versions. Please upgrade.
- Some other program such as an HP Printer using the same port (9102 in this case).

If it is neither of the above, please submit a bug report at bugs.bacula.org.

Another solution might be to run the daemon with the debug option by:

```
Start a DOS shell Window.
cd c:\bacula\bin
bacula-fd -d100 -c c:\bacula\bin\bacula-fd.conf
```

This will cause the FD to write a file **bacula.trace** in the current directory, which you can examine and determine the problem.

Tips and Suggestions

Examples

There are a number of example scripts for various things that can be found in the **example** subdirectory and its subdirectories of the Bacula source distribution.

Upgrading Bacula Versions

The first thing to do before upgrading from one version to another is to ensure that don't overwrite or delete your production (current) version of Bacula until you have tested that the new version works.

If you have installed Bacula into a single directory, this is simple: simply make a copy of your Bacula directory.

If you have done a more typical Unix installation where the binaries are placed in one directory and the configuration files are placed in another, then the simplest way is to configure your new Bacula to go into a single file. Alternatively, make copies of all your binaries and especially your conf files.

Whatever your situation may be (one of the two just described), you should probably start with the **defaultconf** script that can be found in the **examples** subdirectory. Copy this script to the main Bacula directory, modify it as necessary (there should not need to be many modifications), configure Bacula, build it, install it, then stop your production Bacula, copy all the ***.conf** files from your production Bacula directory to the test Bacula directory, start the test version, and run a few test backups. If all seems good, then you can proceed to install the new Bacula in place of or possibly over the old Bacula.

When installing a new Bacula you need not worry about losing the changes you made to your configuration files as the installation process will not overwrite them providing that you do not do a **make uninstall**.

Getting Notified of Job Completion

One of the first things you should do is to ensure that you are being properly notified of the status of each Job run by Bacula, or at a minimum of each Job that terminates with an error.

Until you are completely comfortable with **Bacula**, we recommend that you send an email to yourself for each Job that is run. This is most easily accomplished by adding an email notification address in the **Messages** resource of your Director's configuration file. An email is automatically configured in the default configuration files, but you must ensure that the default **root** address is replaced by your email address.

For examples of how I (Kern) configure my system, please take a look at the **.conf** files found in the **examples** sub-directory. We recommend the following configuration (where you change the paths and email address to correspond to your setup). Note, the **mailcommand** and **operatorcommand** should be on a single line. They were split here for presentation:

```
Messages {
  Name = Standard
  mailcommand = "/home/bacula/bin/bsmtp -h localhost
                -f \"\(Bacula\) %r\"
                -s \"Bacula: %t %e of %c %l\" %r"
  operatorcommand = "/home/bacula/bin/bsmtp -h localhost
                -f \"\(Bacula\) %r\"
                -s \"Bacula: Intervention needed for %j\" %r"
  Mail = your-email-address = all, !skipped, !terminate
  append = "/home/bacula/bin/log" = all, !skipped, !terminate
  operator = your-email-address = mount
  console = all, !skipped, !saved
}
```

You will need to ensure that the **/home/bacula/bin** path on the **mailcommand** and the **operatorcommand** lines points to your **Bacula** binary directory where the **bsmtp** program will be installed. You will also want to ensure that the **your-email-address** is replaced by your email address, and finally, you will also need to ensure that the **/home/bacula/bin/log** points to the file where you want to log all messages.

With the above Messages resource, you will be notified by email of every Job that ran, all the output will be appended to the **log** file you specify, all output will be directed to the console program, and all mount messages will be emailed to you. Note, some messages will be sent to multiple destinations.

The form of the mailcommand is a bit complicated, but it allows you to distinguish whether the Job terminated in error or terminated normally. Please see the Mail Command section of the Messages Resource chapter of this manual for the details of the substitution characters used above.

Once you are totally comfortable with Bacula as I am, or if you have a large

number of nightly Jobs as I do (eight), you will probably want to change the **Mail** command to **Mail On Error** which will generate an email message only if the Job terminates in error. If the Job terminates normally, no email message will be sent, but the output will still be appended to the log file as well as sent to the Console program.

Getting Email Notification to Work

The section above describes how to get email notification of job status. Occasionally, however, users have problems receiving any email at all. In that case, the things to check are the following:

- Ensure that you have a valid email address specified on your **Mail** record in the Director's Messages resource. The email address should be fully qualified. Simply using **root** generally will not work, rather you should use **root@localhost** or better yet your full domain.
- Ensure that you do not have a **Mail** record in the Storage daemon's or File daemon's configuration files. The only record you should have is **director**:

```
director = director-name = all
```

- If all else fails, try replacing the **mailcommand** with

```
mailcommand = "mail -s test your@domain.com"
```

- Once the above is working, assuming you want to use **bsmtp**, submit the desired bsmtp command by hand and ensure that the email is delivered, then put that command into **Bacula**. Small differences in things such as the parenthesis around the word Bacula can make a big difference to some bsmtp programs. For example, you might start simply by using:

```
mailcommand = "/home/bacula/bin/bsmtp -f \"root@localhost\" %r"
```

Getting Notified that Bacula is Running

If like me, you have setup Bacula so that email is sent only when a Job has errors, as described in the previous section of this chapter, inevitably, one

day, something will go wrong and **Bacula** can stall. This could be because Bacula crashes, which is vary rare, or more likely the network has caused **Bacula** to **hang** for some unknown reason.

To avoid this, you can use the **RunAfterJob** command in the Job resource to schedule a Job nightly, or weekly that simply emails you a message saying that Bacula is still running. For example, I have setup the following Job in my Director's configuration file:

```
Schedule {
  Name = "Watchdog"
  Run = Level=Full sun-sat at 6:05
}
Job {
  Name = "Watchdog"
  Type = Admin
  Client=Watchdog
  FileSet="Verify Set"
  Messages = Standard
  Storage = DLTDrive
  Pool = Default
  Schedule = "Watchdog"
  RunAfterJob = "/home/kern/bacula/bin/watchdog %c %d"
}
Client {
  Name = Watchdog
  Address = rufus
  FDPort = 9102
  Catalog = Verify
  Password = ""
  File Retention = 1day
  Job Retention = 1 month
  AutoPrune = yes
}
```

Where I established a schedule to run the Job nightly. The Job itself is type **Admin** which means that it doesn't actually do anything, and I've defined a FileSet, Pool, Storage, and Client, all of which are not really used (and probably don't need to be specified). The key aspect of this Job is the command:

```
  RunAfterJob = "/home/kern/bacula/bin/watchdog %c %d"
```

which runs my "watchdog" script. As an example, I have added the Job codes %c and %d which will cause the Client name and the Director's name to be passed to the script. For example, if the Client's name is **Watchdog** and the Director's name is **main-dir** then referencing $1 in the script would

get **Watchdog** and referencing $2 would get **main-dir**. In this case, having the script know the Client and Director's name is not really useful, but in other situations it may be.

You can put anything in the watchdog scrip. In my case, I like to monitor the size of my catalog to be sure that **Bacula** is really pruning it. The following is my watchdog script:

```
#!/bin/sh
cd /home/kern/mysql/var/bacula
du . * |
/home/kern/bacula/bin/bsmtp  \
   -f "\(Bacula\) abuse@whitehouse.com" -h mail.yyyy.com \
   -s "Bacula running" abuse@whitehouse.com
```

If you just wish to send yourself a message, you can do it with:

```
#!/bin/sh
cd /home/kern/mysql/var/bacula
/home/kern/bacula/bin/bsmtp  \
   -f "\(Bacula\) abuse@whitehouse.com" -h mail.yyyy.com \
   -s "Bacula running" abuse@whitehouse.com <<END-OF-DATA
Bacula is still running!!!
END-OF-DATA
```

Maintaining a Valid Bootstrap File

By using a WriteBootstrap record in each of your Director's Job resources, you can constantly maintain a bootstrap file that will enable you to recover the state of your system as of the last backup without having the Bacula catalog. This permits you to more easily recover from a disaster that destroys your Bacula catalog.

When a Job resource has a **WriteBootstrap** record, Bacula will maintain the designated file (normally on another system but mounted by NSF) with up to date information necessary to restore your system. For example, in my Director's configuration file, I have the following record:

```
 Write Bootstrap = "/mnt/deuter/files/backup/client-name.bsr"
```

where I replace **client-name** by the actual name of the client that is being backed up. Thus, Bacula automatically maintains one file for each of my clients. The necessary bootstrap information is appended to this file during

each **Incremental** backup, and the file is totally rewritten during each **Full** backup.

Note, one major disadvantage of writing to an NFS mounted volume as I do is that if the other machine goes down, the OS will wait forever on the fopen() call that Bacula makes. As a consequence, Bacula will completely stall until the machine exporting the NSF mounts comes back up. The solution to this problem was provided by Andrew Hilborne, and consists of using the **soft** option instead of the **hard** option when mounting the NFS volume, which is typically done in **/etc/fstab**/. The NFS documentation explains these options in detail.

If you are starting off in the middle of a cycle (i.e. with Incremental backups) rather than at the beginning (with a Full backup), the **bootstrap** file will not be immediately valid as it must always have the information from a Full backup as the first record. If you wish to synchronize your bootstrap file immediately, you can do so by running a **restore** command for the client and selecting a full restore, but when the restore command asks for confirmation to run the restore Job, you simply reply no, then copy the bootstrap file that was written to the location specified on the **Write Bootstrap** record. The restore bootstrap file can be found in **restore.bsr** in the working directory that you defined. In the example given below for the client **rufus**, my input is shown in bold. Note, the JobId output has been partially truncated to fit on the page here:

```
(in the Console program)
*{\bf restore}
First you select one or more JobIds that contain files
to be restored. You will then be presented several methods
of specifying the JobIds. Then you will be allowed to
select which files from those JobIds are to be restored.
To select the JobIds, you have the following choices:
     1: List last 20 Jobs run
     2: List Jobs where a given File is saved
     3: Enter list of JobIds to select
     4: Enter SQL list command
     5: Select the most recent backup for a client
     6: Cancel
Select item:  (1-6): {\bf 5}
The defined Client resources are:
     1: Minimatou
     2: Rufus
     3: Timmy
Select Client (File daemon) resource (1-3): {\bf 2}
The defined FileSet resources are:
     1: Kerns Files
Item 1 selected automatically.
+-------+------+-------+---------+---------+------+-------+------------+
```

```
| JobId | Levl | Files | StrtTim | VolName | File | SesId | VolSesTime |
+-------+------+-------+---------+---------+------+-------+------------+
| 2     | F    | 84    |  ...    | test1   | 0    | 1     | 1035645259 |
+-------+------+-------+---------+---------+------+-------+------------+
You have selected the following JobId: 2
Building directory tree for JobId 2 ...
The defined Storage resources are:
     1: File
Item 1 selected automatically.
You are now entering file selection mode where you add and
remove files to be restored. All files are initially added.
Enter "done" to leave this mode.
cwd is: /
$ {\bf done}
84 files selected to restore.
Run Restore job
JobName:    kernsrestore
Bootstrap:  /home/kern/bacula/working/restore.bsr
Where:      /tmp/bacula-restores
FileSet:    Kerns Files
Client:     Rufus
Storage:    File
JobId:      *None*
OK to run? (yes/mod/no): {\bf no}
{\bf quit}
(in a shell window)
{\bf cp ../working/restore.bsr /mnt/deuter/files/backup/rufus.bsr}
```

Rejected Volumes After a Crash

Bacula keeps a count of the number of files on each Volume in its Catalog database so that before appending to a tape, it can verify that the number of files are correct, and thus prevent overwriting valid data. If the Director or the Storage daemon crashes before the job has completed, the tape will contain one more file than is noted in the Catalog, and the next time you attempt to use the same Volume, Bacula will reject it due to a mismatch between the physical tape (Volume) and the catalog.

The easiest solution to this problem is to label a new tape and start fresh. If you wish to continue appending to the current tape, you can do so by using the **update** command in the console program to change the **Volume Files** entry in the catalog. A typical sequence of events would go like the following:

```
- Bacula crashes
- You restart Bacula
```

Bacula then prints:

```
17-Jan-2003 16:45 rufus-dir: Start Backup JobId 13,
                  Job=kernsave.2003-01-17_16.45.46
17-Jan-2003 16:45 rufus-sd: Volume test01 previously written,
                  moving to end of data.
17-Jan-2003 16:46 rufus-sd: kernsave.2003-01-17_16.45.46 Error:
                  I cannot write on this volume because:
                  The number of files mismatch! Volume=11 Catalog=10
17-Jan-2003 16:46 rufus-sd: Job kernsave.2003-01-17_16.45.46 waiting.
                   Cannot find any appendable volumes.
Please use the "label"  command to create a new Volume for:
    Storage:      SDT-10000
    Media type:   DDS-4
    Pool:         Default
```

(note, lines wrapped for presentation) The key here is the line that reads:

```
  The number of files mismatch! Volume=11 Catalog=10
```

It says that Bacula found eleven files on the volume, but that the catalog says there should be ten. When you see this, you can be reasonably sure that the SD was interrupted while writing before it had a chance to update the catalog. As a consequence, you can just modify the catalog count to eleven, and even if the catalog contains references to files saved in file 11, everything will be OK and nothing will be lost. Note that if the SD had written several file marks to the volume, the difference between the Volume cound and the Catalog count could be larger than one, but this is unusual.

If on the other hand the catalog is marked as having more files than Bacula found on the tape, you need to consider the possible negative consequences of modifying the catalog. Please see below for a more complete discussion of this.

Continuing with the example of **Volume = 11 Catalog = 10**, to enable to Bacula to append to the tape, you do the following:

```
{\bf update}
Update choice:
     1: Volume parameters
     2: Pool from resource
     3: Slots from autochanger
Choose catalog item to update (1-3): {\bf 1}
Defined Pools:
     1: Default
     2: File
Select the Pool (1-2):
+-------+---------+---------+---------+-----------+------+----------+------+-----+
| MedId | VolName | MedTyp | VolStat | VolBytes  | Last | VolReten | Recy | Slt |
```

```
+-------+---------+--------+---------+-----------+------+----------+------+-----+
| 1     | test01  | DDS-4  | Error   | 352427156 | ...  | 31536000 | 1    | 0   |
+-------+---------+--------+---------+-----------+------+----------+------+-----+
Enter MediaId or Volume name: {\bf 1}
```

(note table output truncated for presentation) First, you chose to update the Volume parameters by entering a **1**. In the volume listing that follows, notice how the VolStatus is **Error**. We will correct that after changing the Volume Files. Continuing, you respond 1,

```
Updating Volume "test01"
Parameters to modify:
     1: Volume Status
     2: Volume Retention Period
     3: Volume Use Duration
     4: Maximum Volume Jobs
     5: Maximum Volume Files
     6: Maximum Volume Bytes
     7: Recycle Flag
     8: Slot
     9: Volume Files
    10: Pool
    11: Done
Select parameter to modify (1-11): {\bf 9}
Warning changing Volume Files can result
in loss of data on your Volume
Current Volume Files is: 10
Enter new number of Files for Volume: {\bf 11}
New Volume Files is: 11
Updating Volume "test01"
Parameters to modify:
     1: Volume Status
     2: Volume Retention Period
     3: Volume Use Duration
     4: Maximum Volume Jobs
     5: Maximum Volume Files
     6: Maximum Volume Bytes
     7: Recycle Flag
     8: Slot
     9: Volume Files
    10: Pool
    11: Done
Select parameter to modify (1-10): {\bf 1}
```

Here, you have selected **9** in order to update the Volume Files, then you changed it from **10** to **11**, and you now answer **1** to change the Volume Status.

```
Current Volume status is: Error
```

```
Possible Values are:
     1: Append
     2: Archive
     3: Disabled
     4: Full
     5: Used
     6: Read-Only
Choose new Volume Status (1-6): {\bf 1}
New Volume status is: Append
Updating Volume "test01"
Parameters to modify:
     1: Volume Status
     2: Volume Retention Period
     3: Volume Use Duration
     4: Maximum Volume Jobs
     5: Maximum Volume Files
     6: Maximum Volume Bytes
     7: Recycle Flag
     8: Slot
     9: Volume Files
    10: Pool
    11: Done
Select parameter to modify (1-11): {\bf 11}
Selection done.
```

At this point, you have changed the Volume Files from **10** to **11** to account for the last file that was written but not updated in the database, and you changed the Volume Status back to **Append**.

This was a lot of words to describe something quite simple.

The **Volume Files** option exists only in version 1.29 and later, and you should be careful using it. Generally, if you set the value to that which Bacula said is on the tape, you will be OK, especially if the value is one more than what is in the catalog.

Now lets consider the case:

```
  The number of files mismatch! Volume=10 Catalog=12
```

Here the Bacula found fewer files on the volume than what is marked in the catalog. Now, in this case, you should hesitate lot before modifying the count in the catalog, because if you force the catalog from 12 to 10, Bacula will start writing after the file 10 on the tape, possibly overwriting valid data, and if you ever try to restore any of the files that the catalog has marked as saved on Files 11 and 12, all chaos will break out. In this case, you will probably be better off using a new tape. In fact, you might want

to see what files the catalog claims are actually stored on that Volume, and back them up to another tape and recycle this tape.

Security Considerations

Only the File daemon needs to run with root permission (so that it can access all files). As a consequence, you may run your Director, Storage daemon, and MySQL or PostgreSQL database server as non-root processes. Version 1.30 has the **-u** and the **-g** options that allow you to specify a userid and groupid on the command line to be used after Bacula starts.

As of version 1.33, thanks to Dan Langille, it is easier to configure the Bacula Director and Storage daemon to run as non-root.

You should protect the Bacula port addresses (normally 9101, 9102, and 9103) from outside access by a firewall or other means of protection to prevent unauthorized use of your daemons.

You should ensure that the configuration files are not world readable since they contain passwords that allow access to the daemons. Anyone who can access the Director using a console program can restore any file from a backup Volume.

You should protect your Catalog database. If you are using SQLite, make sure that the working directory is readable only by root (or your Bacula userid), and ensure that **bacula.db** has permissions **-rw-r--r--** (i.e. 640) or more strict. If you are using MySQL or PostgreSQL, please note that the Bacula setup procedure leaves the database open to anyone. At a minimum, you should assign the user **bacula** a userid and add it to your Director's configuration file in the appropriate Catalog resource.

Creating Holiday Schedules

If you normally change tapes every day or at least every Friday, but Thursday is a holiday, you can use a trick proposed by Lutz Kittler to ensure that no job runs on Thursday so that you can insert Friday's tape and be sure it will be used on Friday. To do so, define a **RunJobBefore** script that normally returns zero, so that the Bacula job will normally continue. You can then modify the script to return non-zero on any day when you do not want Bacula to run the job.

Automatic Labeling Using Your Autochanger

If you have an autochanger but it does not support barcodes, using a "trick" you can make Bacula automatically label all the volumes in your autochanger's magazine.

First create a file containing one line for each slot in your autochanger that has a tape to be labeled. The line will contain the slot number a colon (:) then the Volume name you want to use. For example, create a file named **volume-list**, which contains:

```
1:Volume001
2:TestVolume02
5:LastVolume
```

The records do not need to be in any order and you don't need to mention all the slots. Normally, you will have a consistent set of Volume names and a sequential set of numbers for each slot you want labeled. In the example above, I've left out slots 3 and 4 just as an example. Now, modify your **mtx-changer** script and comment out all the lines in the **list)** case by putting a # in column 1. Then add the following two lines:

```
  cat <absolute-path>/volume-list
  exit 0
```

so that the whole case looks like:

```
  list)
#
# commented out lines
   cat <absolute-path>/volume-list
   exit 0
   ;;
```

where you replace <absolute-path> with the full path to the volume-list file. Then using the console, you enter the following command:

```
   label barcodes
```

and Bacula will proceed to mount the autochanger Volumes in the list and label them with the Volume names you have supplied. Bacula will think

that the list was provided by the autochanger barcodes, but in reality, it was you who supplied the <barcodes>.

If it seems to work, when it finishes, enter:

```
list volumes
```

and you should see all the volumes nicely created.

Backing Up Portables Using DHCP

You may want to backup laptops or portables that are not always connected to the network. If you are using DHCP to assign an IP address to those machines when they connect, you will need to use the Dynamic Update capability of DNS to assign a name to those machines that can be used in the Address field of the Client resource in the Director's conf file.

Going on Vacation

At some point, you may want to be absent for a week or two and you want to make sure Bacula has enough tape left so that the backups will complete. You start by doing a **list volumes** in the Console program:

```
{\bf list volumes}

Using default Catalog name=BackupDB DB=bacula
Pool: Default
+---------+--------------+-----------+-----------+----------------+-
| MediaId | VolumeName   | MediaType | VolStatus |       VolBytes |
+---------+--------------+-----------+-----------+----------------+-
|      23 | DLT-30Nov02  | DLT8000   | Full      | 54,739,278,128 |
|      24 | DLT-21Dec02  | DLT8000   | Full      | 56,331,524,629 |
|      25 | DLT-11Jan03  | DLT8000   | Full      | 67,863,514,895 |
|      26 | DLT-02Feb03  | DLT8000   | Full      | 63,439,314,216 |
|      27 | DLT-03Mar03  | DLT8000   | Full      | 66,022,754,598 |
|      28 | DLT-04Apr03  | DLT8000   | Full      | 60,792,559,924 |
|      29 | DLT-28Apr03  | DLT8000   | Full      | 62,072,494,063 |
|      30 | DLT-17May03  | DLT8000   | Full      | 65,901,767,839 |
|      31 | DLT-07Jun03  | DLT8000   | Used      | 56,558,490,015 |
|      32 | DLT-28Jun03  | DLT8000   | Full      | 64,274,871,265 |
|      33 | DLT-19Jul03  | DLT8000   | Full      | 64,648,749,480 |
|      34 | DLT-08Aug03  | DLT8000   | Full      | 64,293,941,255 |
|      35 | DLT-24Aug03  | DLT8000   | Append    |  9,999,216,782 |
+---------+--------------+-----------+-----------+----------------+
```

Note, I have truncated the output for presentation purposes. What is significant for is that I can see that my current tape has almost 10 Gbytes of data, and that the average amount of data I get on my tapes is about 60 Gbytes. So if I go on vacation now, I don't need to worry about tape capacity (at least not for short absences).

Equally significant is the fact that I did go on vacation the 28th of June 2003, and when I did the **list volumes** command, my current tape at that time, DLT-07Jun03 MediaId 31, had 56.5 Gbytes written. I could see that the tape would fill shortly. Consequently, I manually marked it as **Used** and replaced it with a fresh tape that I labeled as DLT-28Jun03, thus assuring myself that the backups would all complete without my intervention.

How to Excude File on Windows Regardless of Case

This tip was submitted by Marc Brueckner who wasn't sure of the case of some of his files on Win32, which is case insensitive. The problem is that Bacula thinks that **/UNIMPORTANT FILES** is different from **/Unimportant Files**. Marc was aware that the file exclusion permits wild-cards. So, he specified:

```
"/[Uu][Nn][Ii][Mm][Pp][Oo][Rr][Tt][Aa][Nn][Tt]  [Ff][Ii][Ll][Ee][Ss]"
```

As a consequence, the above exclude works for files of any case.

Please note that this works only in Bacula Exclude statement and not in Include.

Executing Scripts on a Remote Machine

This tip also comes from Marc Brueckner. (Note, this tip is probably outdated by the addition of **ClientRunBeforJob** and **ClientRunAfterJob** Job records, but the technique still could be useful.) First I thought the "Run Before Job" statement in the Job-resource is for executing a script on the remote machine(the machine to be backed up). It could be usefull to execute scripts on the remote machine e.g. for stopping databases or other services while doing the backup. (Of cause I have to start the services again when the backup has finished) I found the following solution: Bacula could execute scrips on the remote machine by using ssh. The authentication is done automatically using a private key. First You have to generate a keypair. I ve done this by:

```
ssh-keygen -b 4096 -t dsa -f Bacula_key
```

This statement may take a little time to run. It creates a public/private key pair with no pass phrase. You could save the keys in /etc/bacula. Now you have two new files : Bacula_key which contains the private key and Bacula_key.pub which contains the public key.

Now you have to append the Bacula_key.pub file to the file authorized_keys in the \root\.ssh directory of the remote machine. Then you have to add (or uncomment) the line

```
AuthorizedKeysFile            %h/.ssh/authorized_keys
```

to the sshd_config file on the remote machine. Where the %h stands for the home-directory of the user (root in this case).

Assuming that your sshd is already running on the remote machine, you can now enter the folloing on the machine where Bacula runs:

```
ssh -i Bacula_key  -l root "ls -la"
```

This should execute the “ls -la” command on the remote machine.

Now you could add lines like the following to your Director’s conf file:

```
...
Run Before Job = ssh -i /etc/bacula/Bacula_key 192.168.1.1 \
                 "/etc/init.d/database stop"
Run After Job = ssh -i /etc/bacula/Bacula_key 192.168.1.1 \
                 "/etc/init.d/database start"
...
```

Even though Bacula version 1.32 has a ClientRunBeforeJob, the ssh method still could be useful for updating all the Bacula clients on several remote machines in a single script.

Recycling All Your Volumes

This tip comes from Phil Stracchino.

If you decide to blow away your catalog and start over, the simplest way to re-add all your prelabelled tapes with the minimum of fuss (provided

you don't care about the data on the tapes) is to add the tape labels using the console **add** command, then go into the catalog and manually set the VolStatus of every tape to **Recycle**.

The SQL command to do this is very simple:

```
update Media set VolStatus = "Recycle";
```

Bacula will then ignore the data already stored on the tapes and just re-use each tape without further objection.

Backing up ACLs on ext3 or XFS filesystems

This tip comes from Volker Sauer.

Note, this tip was given prior to implementation of ACLs in Bacula (version 1.34.5). It is left here because dumping/displaying ACLs can still be useful in testing/verifying that Bacula is backing up and restoring your ACLs properly. Please see the aclsupport FileSet option in the configuration chapter of this manual.

For example, you could dump the ACLs to a file with a script similar to the following:

```
#!/bin/sh
BACKUP_DIRS="/foo /bar"
STORE_ACL=/root/acl-backup
umask 077
for i in $BACKUP_DIRS; do
 cd $i /usr/bin/getfacl -R --skip-base .>$STORE_ACL/${i//\//_}
done
```

Then use Bacula to backup **/root/acl-backup**.

The ACLs could be restored using Bacula to the **/root/acl-backup** file, then restored to your system using:

```
setfacl --restore/root/acl-backup
```

Total Automation of Bacula Tape Handling

This tip was provided by Alexander Kuehn.

Bacula is a really nice backup program except that the manual tape changing requires user interaction with the bacula console.

Fortunately I can fix this. NOTE!!! This suggestion applies for people who do *NOT* have tape autochangers and must change tapes manually.!!!!!

Bacula supports a variety of tape changers through the use of mtx-changer scripts/programs. This highly flexible approach allowed me to create this shell script which does the following: Whenever a new tape is required it sends a mail to the operator to insert the new tape. Then it waits until a tape has been inserted, sends a mail again to say thank you and let's bacula continue it's backup. So you can schedule and run backups without ever having to log on or see the console. To make the whole thing work you need to create a Device resource which looks something like this ("Archive Device", "Maximum Changer Wait", "Media Type" and "Label media" may have different values):

```
Device {
   Name=DDS3
   Archive Device = # use yours not mine! ;)/dev/nsa0
   Changer Device = # not really required/dev/nsa0
   Changer Command = "# use this (maybe change the path)!
         /usr/local/bin/mtx-changer %o %a %S"
   Maximum Changer Wait = 3d          # 3 days in seconds
   AutomaticMount = yes;              # mount on start
   AlwaysOpen = yes;                  # keep device locked
   Media Type = DDS3                  # it's just a name
   RemovableMedia = yes;              #
   Offline On Unmount = Yes;          # keep this too
   Label media = Yes;                 #
}
```

As the script has to emulate the complete wisdom of a mtx-changer it has an internal "database" where which tape is stored, you can see this at that line:

```
labels="VOL-0001 VOL-0002 VOL-0003 VOL-0004 VOL-0005 VOL-0006
VOL-0007 VOL-0008 VOL-0009 VOL-0010 VOL-0011 VOL-0012"
```

The above should be all on one line, and it effectivly tells Bacula that volume "VOL-0001" is located in slot 1 (which is our lowest slot), that volume "VOL-0002" is located in slot 2 and so on.. The script also maintains a logfile (/var/log/mtx.log) where you can monitor its operation.

Running Concurrent Jobs

Bacula can run multiple concurrent jobs, but the default configuration files are not set to do so. Using the **Maximum Concurrent Jobs** directive, you have a lot of control over how many jobs can run at the same time, and which jobs can run simultaneously. The downside is that it can be a bit tricky to set it up for the first time as you need to set the concurrency in at least five different places.

The Director, the File daemon, and the Storage daemon each have a **Maximum Concurrent Jobs** directive that determines overall number of concurrent jobs the daemon will run. The default is one for the Director and ten for both the File daemon and the Storage daemon, so assuming you will not be running more than ten concurrent jobs, the only changes that are needed are in the Director's conf file (bacula-dir.conf).

Within the Director's configuration file, **Maximum Concurrent Jobs** can be set in the Direct, Job, Client, and Storage resources. Each one must be set properly, according to your needs, otherwise your jobs may be run one at a time.

For example, if you want two different jobs to run simultaneously backing up the same Client to the same Storage device, they will run concurrentl only if you have set **Maximum Concurrent Jobs** greater than one in the Director resource, the Client resource, and the Storage resource in bacula-dir.conf.

We recommend that you carefully test your multiple concurrent backup including doing thorough restore testing before you put it into production.

Below is a super stripped down bacula-dir.conf file showing you the four places where the the file has been modified to allow the same job **NightlySave** to run up to four times concurrently. The change to the Job resource is not necessary if you want different Jobs to run at the same time, which is the normal case.

```
#
# Bacula Director Configuration file -- bacula-dir.conf
#
Director {
  Name = rufus-dir
  Maximum Concurrent Jobs = 4
  ...
}
Job {
  Name = "NightlySave"
```

```
  Maximum Concurrent Jobs = 4
  Client = rufus-fd
  Storage = File
  ...
}
Client {
  Name = rufus-fd
  Maximum Concurrent Jobs = 4
  ...
}
Storage {
  Name = File
  Maximum Concurrent Jobs = 4
  ...
}
```

Volume Utility Tools

This document describes the utility programs written to aid Bacula users and developers in dealing with Volumes external to Bacula.

Specifying the Configuration File

Starting with version 1.27, each of the following programs requires a valid Storage daemon configuration file (actually, the only part of the configuration file that these programs need is the **Device** resource definitions). This permits the programs to find the configuration parameters for your archive device (generally a tape drive). By default, they read **bacula-sd.conf** in the current directory, but you may specify a different configuration file using the **-c** option.

Specifying a Device Name For a Tape

Each of these programs require a **device-name** where the Volume can be found. In the case of a tape, this is the physical device name such as **/dev/nst0** or **/dev/rmt/0ubn** depending on your system. For the program to work, it must find the identical name in the Device resource of the configuration file. See below for specifying Volume names.

Specifying a Device Name For a File

If you are attempting to read or write an archive file rather than a tape, the **device-name** should be the full path to the archive location including the filename. The filename (last part of the specification) will be stripped and used as the Volume name, and the path (first part before the filename) must have the same entry in the configuration file. So, the path is equivalent to the archive device name, and the filename is equivalent to the volume name.

Specifying Volumes

In general, you must specify the Volume name to each of the programs below (with the exception of **btape**). The best method to do so is to specify a **bootstrap** file on the command line with the **-b** option. As part of the bootstrap file, you will then specify the Volume name or Volume names if more than one volume is needed. For example, suppose you want to

read tapes **tape1** and **tape2**. First construct a **bootstrap** file named say, **list.bsr** which contains:

```
Volume=test1|test2
```

where each Volume is separated by a vertical bar. Then simply use:

```
./bls -b list.bsr /dev/nst0
```

In the case of Bacula Volumes that are on files, you may simply append volumes as follows:

```
./bls /tmp/test1\|test2
```

where the backslash (\) was necessary as a shell escape to permit entering the vertical bar (—).

And finally, if you feel that specifying a Volume name is a bit complicated with a bootstrap file, you can use the **-V** option (on all programs except **bcopy**) to specify one or more Volume names separated by the vertical bar (—). For example,

```
./bls -V Vol001 /dev/nst0
```

You may also specify an asterisk (*) to indicate that the program should accept any volume. For example:

```
./bls -V* /dev/nst0
```

bls

bls can be used to do an **ls** type listing of a **Bacula** tape or file. It is called:

```
Usage: bls [-d debug_level] <device-name>
       -b <file>       specify a bootstrap file
       -c <file>       specify a configuration file
       -d <level>       specify a debug level
       -e <file>       exclude list
       -i <file>       include list
       -j              list jobs
```

```
       -k              list blocks
       -L              list tape label
    (none of above)    list saved files
       -p              proceed inspite of I/O errors
       -t              use default tape device
       -v              be verbose
       -V              specify Volume names (separated by |)
       -?              print this message
```

For example, to list the contents of a tape:

```
./bls -V Volume-name /dev/nst0
```

Or to list the contents of a file:

```
./bls /tmp/Volume-name
or
./bls -V Volume-name /tmp
```

Note that, in the case of a file, the Volume name becomes the filename, so in the above example, you will replace the **xxx** with the name of the volume (file) you wrote.

Normally if no options are specified, **bls** will produce the equivalent output to the **ls -l** command for each file on the tape. Using other options listed above, it is possible to display only the Job records, only the tape blocks, etc. For example:

```
./bls /tmp/File002
bls: butil.c:148 Using device: /tmp
drwxrwxr-x   3 k  k  4096 02-10-19 21:08  /home/kern/bacula/k/src/dird/
drwxrwxr-x   2 k  k  4096 02-10-10 18:59  /home/kern/bacula/k/src/dird/CVS/
-rw-rw-r--   1 k  k    54 02-07-06 18:02  /home/kern/bacula/k/src/dird/CVS/Root
-rw-rw-r--   1 k  k    16 02-07-06 18:02  /home/kern/bacula/k/src/dird/CVS/Repository
-rw-rw-r--   1 k  k  1783 02-10-10 18:59  /home/kern/bacula/k/src/dird/CVS/Entries
-rw-rw-r--   1 k  k 97506 02-10-18 21:07  /home/kern/bacula/k/src/dird/Makefile
-rw-r--r--   1 k  k  3513 02-10-18 21:02  /home/kern/bacula/k/src/dird/Makefile.in
-rw-rw-r--   1 k  k  4669 02-07-06 18:02  /home/kern/bacula/k/src/dird/README-config
-rw-r--r--   1 k  k  4391 02-09-14 16:51  /home/kern/bacula/k/src/dird/authenticate.c
-rw-r--r--   1 k  k  3609 02-07-07 16:41  /home/kern/bacula/k/src/dird/autoprune.c
-rw-rw-r--   1 k  k  4418 02-10-18 21:03  /home/kern/bacula/k/src/dird/bacula-dir.conf
...
-rw-rw-r--   1 k  k    83 02-08-31 19:19  /home/kern/bacula/k/src/dird/.cvsignore
bls: Got EOF on device /tmp
84 files found.
```

Listing Jobs

If you are listing a Volume to determine what Jobs to restore, normally the **-j** option provides you with most of what you will need as long as you don't have multiple clients. For example,

```
./bls -j /tmp/test1
Volume Record: SessId=2 SessTime=1033762386 JobId=0 DataLen=144
Begin Session Record: SessId=2 SessTime=1033762386 JobId=1 Level=F Type=B
End Session Record: SessId=2 SessTime=1033762386 JobId=1 Level=F Type=B
Begin Session Record: SessId=3 SessTime=1033762386 JobId=2 Level=I Type=B
End Session Record: SessId=3 SessTime=1033762386 JobId=2 Level=I Type=B
Begin Session Record: SessId=4 SessTime=1033762386 JobId=3 Level=I Type=B
End Session Record: SessId=4 SessTime=1033762386 JobId=3 Level=I Type=B
bls: Got EOF on device /tmp
```

shows a full save followed by two incremental saves.

Adding the **-v** option will display virtually all information that is available for each record:

Listing Blocks

Normally, except for debugging purposes, you will not need to list Bacula blocks (the "primitive" unit of Bacula data on the Volume). However, you can do so with:

```
./bls -k /tmp/File002
bls: butil.c:148 Using device: /tmp
Block: 1 size=64512
Block: 2 size=64512
...
Block: 65 size=64512
Block: 66 size=19195
bls: Got EOF on device /tmp
End of File on device
```

By adding the **-v** option, you can get more information, which can be useful in knowing what sessions were written to the volume:

```
./bls -k -v /tmp/File002
Volume Label:
Id                : Bacula 0.9 mortal
VerNo             : 10
```

```
VolName           : File002
PrevVolName       :
VolFile           : 0
LabelType         : VOL_LABEL
LabelSize         : 147
PoolName          : Default
MediaType         : File
PoolType          : Backup
HostName          :
Date label written: 2002-10-19 at 21:16
Block: 1 blen=64512 First rec FI=VOL_LABEL SessId=1 SessTim=1035062102 Strm=0 rlen=147
Block: 2 blen=64512 First rec FI=6 SessId=1 SessTim=1035062102 Strm=DATA rlen=4087
Block: 3 blen=64512 First rec FI=12 SessId=1 SessTim=1035062102 Strm=DATA rlen=5902
Block: 4 blen=64512 First rec FI=19 SessId=1 SessTim=1035062102 Strm=DATA rlen=28382
...
Block: 65 blen=64512 First rec FI=83 SessId=1 SessTim=1035062102 Strm=DATA rlen=1873
Block: 66 blen=19195 First rec FI=83 SessId=1 SessTim=1035062102 Strm=DATA rlen=2973
bls: Got EOF on device /tmp
End of File on device
```

Armed with the SessionId and the SessionTime, you can extract just about anything.

If you want to know even more, add a second **-v** to the command line to get a dump of every record in every block.

```
./bls -k -v -v /tmp/File002
bls: block.c:79 Dump block  80f8ad0: size=64512 BlkNum=1
               Hdrcksum=b1bdfd6d cksum=b1bdfd6d
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=VOL_LABEL Strm=0 len=147 p=80f8b40
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=SOS_LABEL Strm=-7 len=122 p=80f8be7
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=1 Strm=UATTR len=86 p=80f8c75
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=2 Strm=UATTR len=90 p=80f8cdf
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=3 Strm=UATTR len=92 p=80f8d4d
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=3 Strm=DATA len=54 p=80f8dbd
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=3 Strm=MD5 len=16 p=80f8e07
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=4 Strm=UATTR len=98 p=80f8e2b
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=4 Strm=DATA len=16 p=80f8ea1
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=4 Strm=MD5 len=16 p=80f8ec5
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=5 Strm=UATTR len=96 p=80f8ee9
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=5 Strm=DATA len=1783 p=80f8f5d
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=5 Strm=MD5 len=16 p=80f9668
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=6 Strm=UATTR len=95 p=80f968c
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=6 Strm=DATA len=32768 p=80f96ff
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=6 Strm=DATA len=32768 p=8101713
bls: block.c:79 Dump block  80f8ad0: size=64512 BlkNum=2
               Hdrcksum=9acc1e7f cksum=9acc1e7f
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=6 Strm=contDATA len=4087 p=80f8b40
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=6 Strm=DATA len=31970 p=80f9b4b
bls: block.c:92    Rec: VId=1 VT=1035062102 FI=6 Strm=MD5 len=16 p=8101841
...
```

bextract

Normally, you will restore files by running a **Restore** Job from the **Console** program. However, **bextract** can be used to extract a single file or a list of files from a Bacula tape or file. In fact, **bextract** can be a useful tool to restore files to an empty system assuming you are able to boot, you have statically linked **bextract** and you have an appropriate **bootstrap** file.

It is called:

```
Usage: bextract [-d debug_level] <device-name> <directory-to-store-files>
       -b <file>       specify a bootstrap file
       -dnn            set debug level to nn
       -e <file>       exclude list
       -i <file>       include list
       -p              proceed inspite of I/O errors
       -V              specify Volume names (separated by |)
       -?              print this message
```

where **device-name** is the Archive Device (raw device name or full filename) of the device to be read, and **directory-to-store-files** is a path prefix to prepend to all the files restored.

NOTE: On Windows systems, if you specify a prefix of say d:/tmp, any file that would have been restored to **c:/My Documents** will be restored to **d:/tmp/My Documents**. That is, the original drive specification will be stripped. If no prefix is specified, the file will be restored to the original drive.

Extracting with Include or Exclude Lists

Using the **-e** option, you can specify a file containing a list of files to be excluded. Wildcards can be used in the exclusion list. This option will normally be used in conjunction with the **-i** option (see below). Both the **-e** and the **-i** options may be specified at the same time as the **-b** option. The bootstrap filters will be applied first, then the include list, then the exclude list.

Likewise, and probably more importantly, with the **-i** option, you can specify a file that contains a list (one file per line) of files and directories to include to be restored. The list must contain the full filename with the path. If you specify a path name only, all files and subdirectories of that path will

be restored. If you specify a line containing only the filename (e.g. **myfile.txt**) it probably will not be extracted because you have not specified the full path.

For example, if the file **include-list** contains:

```
/home/kern/bacula
/usr/local/bin
```

Then the command:

```
./bextract -i include-list -V Volume /dev/nst0 /tmp
```

will restore from the Bacula archive **/dev/nst0** all files and directories in the backup from **/home/kern/bacula** and from **/usr/local/bin**. The restored files will be placed in a file of the original name under the directory **/tmp** (i.e. /tmp/home/kern/bacula/... and /tmp/usr/local/bin/...).

Extracting With a Bootstrap File

The **-b** option is used to specify a **bootstrap** file containing the information needed to restore precisely the files you want. Specifying a **bootstrap** file is optional but recommended because it gives you the most control over which files will be restored. For more details on the **bootstrap** file, please see Restoring Files with the Bootstrap File chapter of this document. Note, you may also use a bootstrap file produced by the **restore** command. For example:

```
./bextract -b bootstrap-file /dev/nst0 /tmp
```

The bootstrap file allows detailed specification of what files you want restored (extracted). You may specify a bootstrap file and include and/or exclude files at the same time. The bootstrap conditions will first be applied, and then each file record seen will be compared to the include and exclude lists.

Extracting From Multiple Volumes

If you wish to extract files that span several Volumes, you can specify the Volume names in the bootstrap file or you may specify the Volume names

on the command line by separating them with a vertical bar. See the section above under the **bls** program entitled **Listing Multiple Volumes** for more information. The same techniques apply equally well to the **bextract** program.

bscan

The **bscan** program can be used to re-create a database (catalog) from the backup information written to one or more Volumes. This is normally needed only if one or more Volumes have been pruned or purged from your catalog so that the records on the Volume are no longer in the catalog.

With some care, it can also be used to synchronize your existing catalog with a Volume. Since **bscan** modifies your catalog, we strongly recommend that you do a simple ASCII backup of your database before running **bscan** just to be sure. See Compacting Your Database.

bscan can also be useful in a disaster recovery situation, after the loss of a hard disk, if you do not have a valid **bootstrap** file for reloading your system, or if a Volume has been recycled but not overwritten, you can use **bscan** to re-create your database, which can then be used to **restore** your system or a file to its previous state.

It is called:

```
Usage: bscan [options] <bacula-archive>
       -b bootstrap   specify a bootstrap file
       -c <file>      specify configuration file
       -d <nn>        set debug level to nn
       -m             update media info in database
       -n <name>      specify the database name (default bacula)
       -u <user>      specify database user name (default bacula)
       -P <password>  specify database password (default none)
       -h <host>      specify database host (default NULL)
       -p             proceed inspite of I/O errors
       -r             list records
       -s             synchronize or store in database
       -v             verbose
       -V <Volumes>   specify Volume names (separated by |)
       -w <dir>       specify working directory (default from conf file)
       -?             print this message
```

If you are using MySQL or PostgreSQL, there is no need to supply a working directory since in that case, bscan knows where the databases are. However, if you have provided security on your database, you may need to supply

either the database name (**-b** option), the user name (**-u** option), and/or the password (**-p**) options.

As an example, let's suppose that you did a backup to Volume "Vol001" and that sometime later all record of that Volume was pruned or purged from the database. By using **bscan** you can recreate the catalog entries for that Volume and then use the **restore** command in the Console to restore whatever you want. A command something like:

```
bscan -c bacula-sd.conf -v -V Vol001 /dev/nst0
```

will give you a give you an idea of what is going to happen without changing your catalog. Of course, you may need to change the path to the Storage daemon's conf file, the Volume name, and your tape (or disk) device name. This command must read the entire tape, so if it has a lot of data, it may take a long time, and thus you might want to immediately use the command listed below. Note, if you are writing to a disk file, replace the device name with the path to the directory that contains the Volume. This must correspond to the Archive Device in the conf file.

Then to actually write or store the records in the catalog, add the **-s** option as follows:

```
bscan -s -m -c bacula-sd.conf -v -V Vol001 /dev/nst0
```

When writing to the database, if bscan finds existing records, it will generally either update them if something is wrong or leave them alone. Thus if the Volume you are scanning is all or partially in the catalog already, no harm will be done to that existing data. Any missing data will simply be added.

If you have multiple tapes, you can scan them with:

```
bscan -s -m -c bacula-sd.conf -v -V Vol001\|Vol002\|Vol003 /dev/nst0
```

You should, where ever possible try to specify the tapes in the order they are written. However, bscan can handle scanning tapes that are not sequential. Any incomplete records at the end of the tape will simply be ignored in that case.

Note, the restoration process using bscan is not identical to the original creation of the catalog data. This is because certain non-essential data such as volume reads, volume mounts, etc is not stored on the Volume, and thus

is not restored by bscan. The results of bscanning are, however, perfectly valid, and will permit restoration of any or all the files in the catalog using the normal Bacula console commands.

Using bscan to Compare a Volume to an existing Catalog

If you wish to compare the contents of a Volume to an existing catalog without changing the catalog, you can safely do so if and only if you do **not** specify either the **-m** or the **-s** options. However, at this time (Bacula version 1.26), the comparison routines are not as good or as thorough as they should be, so we don't particularly recommend this mode other than for testing.

Using bscan to Recreate a Catalog from a Volume

This is the mode for which **bscan** is most useful. You can either **bscan** into a freshly created catalog, or directly into your existing catalog (after having made an ASCII copy as described above). Normally, you should start with a freshly created catalog that contains no data.

Starting with a single Volume named **TestVolume1**, you run a command such as:

```
./bscan -V TestVolume1 -v -s -m -c bacula-sd.conf /dev/nst0
```

If there is more than one volume, simply append it to the first one separating it with a vertical bar. You may need to precede the vertical bar with a forward slash escape the shell – e.g. **TestVolume1\—TestVolume2** . The **-v** option was added for verbose output (this can be omitted if desired). The **-s** option that tells **bscan** to store information in the database. The physical device name **/dev/nst0** is specified after all the options.

For example, after having done a full backup of a directory, then two incrementals, I reinitialized the SQLite database as described above, and using the bootstrap.bsr file noted above, I entered the following command:

```
./bscan -b bootstrap.bsr -v -s -c bacula-sd.conf /dev/nst0
```

which produced the following output:

```
bscan: bscan.c:182 Using Database: bacula, User: bacula
bscan: bscan.c:673 Created Pool record for Pool: Default
bscan: bscan.c:271 Pool type "Backup" is OK.
bscan: bscan.c:632 Created Media record for Volume: TestVolume1
bscan: bscan.c:298 Media type "DDS-4" is OK.
bscan: bscan.c:307 VOL_LABEL: OK for Volume: TestVolume1
bscan: bscan.c:693 Created Client record for Client: Rufus
bscan: bscan.c:769 Created new JobId=1 record for original JobId=2
bscan: bscan.c:717 Created FileSet record "Kerns Files"
bscan: bscan.c:819 Updated Job termination record for new JobId=1
bscan: bscan.c:905 Created JobMedia record JobId 1, MediaId 1
bscan: Got EOF on device /dev/nst0
bscan: bscan.c:693 Created Client record for Client: Rufus
bscan: bscan.c:769 Created new JobId=2 record for original JobId=3
bscan: bscan.c:708 Fileset "Kerns Files" already exists.
bscan: bscan.c:819 Updated Job termination record for new JobId=2
bscan: bscan.c:905 Created JobMedia record JobId 2, MediaId 1
bscan: Got EOF on device /dev/nst0
bscan: bscan.c:693 Created Client record for Client: Rufus
bscan: bscan.c:769 Created new JobId=3 record for original JobId=4
bscan: bscan.c:708 Fileset "Kerns Files" already exists.
bscan: bscan.c:819 Updated Job termination record for new JobId=3
bscan: bscan.c:905 Created JobMedia record JobId 3, MediaId 1
bscan: Got EOF on device /dev/nst0
bscan: bscan.c:652 Updated Media record at end of Volume: TestVolume1
bscan: bscan.c:428 End of Volume. VolFiles=3 VolBlocks=57 VolBytes=10,027,437
```

The key points to note are that **bscan** prints a line when each major record is created. Due to the volume of output, it does not print a line for each file record unless you supply the **-v** option twice or more on the command line.

In the case of a Job record, the new JobId will not normally be the same as the original Jobid. For example, for the first JobId above, the new JobId is 1, but the original JobId is 2. This is nothing to be concerned about as it is the normal nature of databases. **bscan** will keep everything straight.

Although **bscan** claims that it created a Client record for Client: Rufus three times, it was actually only created the first time. This is normal.

You will also notice that it read an end of file after each Job (Got EOF on device ...). Finally the last line gives the total statistics for the bscan.

If you had added a second **-v** option to the command line, Bacula would have been even more verbose, dumping virtually all the details of each Job record it encountered.

Now if you start Bacula and enter a **list jobs** command to the console program, you will get:

```
+-------+----------+------------------+------+-----+----------+----------+---------+
| JobId | Name     | StartTime        | Type | Lvl | JobFiles | JobBytes | JobStat |
+-------+----------+------------------+------+-----+----------+----------+---------+
| 1     | kernsave | 2002-10-07 14:59 | B    | F   | 84       | 4180207  | T       |
| 2     | kernsave | 2002-10-07 15:00 | B    | I   | 15       | 2170314  | T       |
| 3     | kernsave | 2002-10-07 15:01 | B    | I   | 33       | 3662184  | T       |
+-------+----------+------------------+------+-----+----------+----------+---------+
```

which corresponds virtually identically with what the database contained before it was re-initialized and restored with bscan. All the Jobs and Files found on the tape are restored including most of the Media record. The Volume (Media) records restored will be marked as **Full** so that they cannot be rewritten without operator intervention.

It should be noted that **bscan** cannot restore a database to the exact condition it was in previously because a lot of the less important information contained in the database is not saved to the tape. Nevertheless, the reconstruction is sufficiently complete, that you can run **restore** against it and get valid results.

Using bscan to Correct the Volume File Count

If the Storage daemon crashes during a backup Job, the catalog will no be properly updated for the Volume being used at the time of the crash. This means that the Storage daemon will have written say 20 files on the tape, but the catalog record for the Volume indicates only 19 files.

Bacula refuses to write on a tape that contains a different number of files from what is in the catalog. To correct this situation, you may run a **bscan** with the **-m** option (but **without** the **-s** option) to update only the final Media record for the Volumes read.

After bscan

If you use **bscan** to enter the contents of the Volume into an existing catalog, you should be aware that the records you entered may be immediately pruned during the next job particularly if the Volume is very old or had been previously purged. To avoid this, after running **bscan**, you can manually set the volume status (VolStatus) to **Read-Only** by using the **update** command in the catalog. This will allow you to restore from the volume without having it immediately purged. When you have restored and backed up the data, you can reset the VolStatus to **Used** and the Volume will be purged from the catalog.

bcopy

The **bcopy** program can be used to copy one **Bacula** archive file to another. For example, you may copy a tape to a file, a file to a tape, a file to a file, or a tape to a tape. For tape to tape, you will need two tape drives. (a later version is planned that will buffer it to disk). In the process of making the copy, no record of the information written to the new Volume is stored in the catalog. This means that the new Volume, though it contains valid backup data, cannot be accessed directly from existing catalog entries. If you wish to be able to use the Volume with the Console restore command, for example, you must first bscan the new Volume into the catalog.

bcopy Command Options

```
Usage: bcopy [-d debug_level] <input-archive> <output-archive>
       -b bootstrap      specify a bootstrap file
       -c <file>         specify configuration file
       -dnn              set debug level to nn
       -i                specify input Volume names (separated by |)
       -o                specify output Volume names (separated by |)
       -p                proceed inspite of I/O errors
       -v                verbose
       -w dir            specify working directory (default /tmp)
       -?                print this message
```

By using a **bootstrap** file, you can copy parts of a Bacula archive file to another archive.

One of the objectives of this program is to be able to recover as much data as possible from a damaged tape. However, the current version does not yet have this feature.

As this is a new program, any feedback on its use would be appreciated. In addition, I only have a single tape drive, so I have never been able to test this program with two tape drives.

btape

This program permits a number of elementary tape operations via a tty command interface. The **test** command, described below, can be very useful for testing older tape drive compatibility problems. Aside from initial testing of tape drive compatibility with **Bacula**, **btape** will be mostly used by developers writing new tape drivers.

btape can be dangerous to use with existing **Bacula** tapes because it will relabel a tape or write on the tape if so requested regardless that the tape may contain valuable data, so please be careful and use it only on blank tapes.

To work properly, **btape** needs to read the Storage daemon's configuration file. As a default, it will look for **bacula-sd.conf** in the current directory. If your configuration file is elsewhere, please use the **-c** option to specify where.

The physical device name must be specified on the command line, and that this same device name must be present in the Storage daemon's configuration file read by **btape**

```
Usage: btape [-c config_file] [-d debug_level] [device_name]
       -c <file>   set configuration file to file
       -dnn        set debug level to nn
       -s          turn off signals
       -t          open the default tape device
       -?          print this message.
```

Using btape to Verify your Tape Drive

An important reason for this program is to ensure that a Storage daemon configuration file is defined so that Bacula will correctly read and write tapes.

It is highly recommended that you run the **test** command before running your first Bacula job to ensure that the parameters you have defined for your storage device (tape drive) will permit **Bacula** to function properly. You only need to mount a blank tape, enter the command, and the output should be reasonably self explanatory. Please see the Tape Testing Chapter of this manual for the details.

btape Commands

The full list of commands are:

```
  Command    Description
  =======    ===========
  bsf        backspace file
  bsr        backspace record
  cap        list device capabilities
```

```
clear      clear tape errors
eod        go to end of Bacula data for append
test       General test Bacula tape functions
eom        go to the physical end of medium
fill       fill tape, write onto second volume
unfill     read filled tape
fsf        forward space a file
fsr        forward space a record
help       print this command
label      write a Bacula label to the tape
load       load a tape
quit       quit btape
rd         read tape
readlabel  read and print the Bacula tape label
rectest    test record handling functions
rewind     rewind the tape
scan       read tape block by block to EOT and report
status     print tape status
test       test a tape for compatibility with Bacula
weof       write an EOF on the tape
wr         write a single record of 2048 bytes
```

The most useful commands are:

- test – test writing records and EOF marks and reading them back.
- fill – completely fill a volume with records, then write a few records on a second volume, and finally, both volumes will be read back. Please be aware that the data written will be quite similar every record, so you might want to turn compression off. One user found that the fill command wrote 750Gb to a tape that can hold 35Gb – so you can see that the hardware compression really worked well!
- readlabel – read and dump the label on a Bacula tape.
- cap – list the device capabilities as defined in the configuration file and as perceived by the Storage daemon.

The **readlabel** command can be used to display the details of a Bacula tape label. This can be useful if the physical tape label was lost or damaged.

In the event that you want to relabel a **Bacula**, you can simply use the **label** command which will write over any existing label. However, please note for labeling tapes, we recommend that you use the **label** command in the **Console** program since it will never overwrite a valid Bacula tape.

Other Programs

The following programs are general utility programs and in general do not need a configuration file nor a device name.

bsmtp

bsmtp is a simple mail transport program that permits more flexibility than the standard mail programs typically found on Unix systems. It can even be used on Windows machines.

It is called:

```
Usage: bsmtp [-f from] [-h mailhost] [-s subject] [-c copy] [recipient ...]
       -c          set the Cc: field
       -dnn        set debug level to nn
       -f          set the From: field
       -h          use mailhost:port as the bsmtp server
       -s          set the Subject: field
       -?          print this message.
```

If the **-f** option is not specified, **bsmtp** will use your userid. If the option is not specified **bsmtp** will use the value in the environment variable **bsmtpSERVER** or if there is none **localhost**. By default port 25 is used.

recipients is a space separated list of email recipients.

The body of the email message is read from standard input.

An example of the use of **bsmtp** would be to put the following statement in the **Messages** resource of your **bacula-dir.conf** file. Note, these commands should appear on a single line each.

```
mailcommand = "/home/bacula/bin/bsmtp -h mail.domain.com -f \"\(Bacula\) %r\"
              -s \"Bacula: %t %e of %c %l\" %r"
operatorcommand = "/home/bacula/bin/bsmtp -h mail.domain.com -f \"\(Bacula\) %r\"
                  -s \"Bacula: Intervention needed for %j\" %r"
```

Where you replace **/home/bacula/bin** with the path to your **Bacula** binary directory, and you replace **mail.domain.com** with the fully qualified name of your bsmtp (email) server, which normally listens on port 25. For more details on the substitution characters (e.g. %r) used in the above line, please see the documentation of the MailCommand in the Messages Resource chapter of this manual.

It is HIGHLY recommended that you test one or two cases by hand to make sure that the **mailhost** that you specified is correct and that it will accept your email requests. Since **bsmtp** always uses a TCP connection rather than writing in the spool file, you may find that your **from** address is being rejected because it does not contain a valid domain, or because your message is caught in your spam filtering rules. Generally, you should specify a fully qualified domain name in the **from** field, and depending on whether your bsmtp gateway is Exim or Sendmail, you may need to modify the syntax of the from part of the message. Please test.

When running **bsmtp** by hand, you will need to terminate the message by entering a ctl-d in column 1 of the last line.

dbcheck

dbcheck is a simple program that will search for inconsistencies in your database, and optionally fix them. The **dbcheck** program can be found in the <**bacula-source**>**/src/tools** directory of the source distribution. Though it is built with the make process, it is not normally "installed".

It is called:

```
Usage: dbcheck [-c config] [-C catalog name] [-d debug_level]     []
       -b              batch mode
       -C              catalog name in the director conf file
       -c              director conf filename
       -dnn            set debug level to nn
       -f              fix inconsistencies
       -v              verbose
       -?              print this message
```

If the **-c** option is given with the Director's conf file, there is no need to enter any of the command line arguments, in particular the working directory as dbcheck will read them from the file.

If the **-f** option is specified, **dbcheck** will repair (**fix**) the inconsistencies it finds. Otherwise, it will report only.

If the **-b** option is specified, **dbcheck** will run in batch mode, and it will proceed to examine and fix (if -f is set) all programmed inconsistency checks. If the **-b** option is not specified, **dbcheck** will enter interactive mode and prompt with the following:

```
Hello, this is the database check/correct program.
```

```
Please select the function you want to perform.
     1) Toggle modify database flag
     2) Toggle verbose flag
     3) Repair bad Filename records
     4) Repair bad Path records
     5) Eliminate duplicate Filename records
     6) Eliminate duplicate Path records
     7) Eliminate orphaned Jobmedia records
     8) Eliminate orphaned File records
     9) Eliminate orphaned Path records
    10) Eliminate orphaned Filename records
    11) Eliminate orphaned FileSet records
    12) Eliminate orphaned Client records
    13) Eliminate orphaned Job records
    14) Eliminate all Admin records
    15) Eliminate all Restore records
    16) All (3-15)
    17) Quit
Select function number:
```

By entering 1 or 2, you can toggle the modify database flag (-f option) and the verbose flag (-v). It can be helpful and reassuring to turn off the modify database flag, then select one or more of the consistency checks (items 3 through 9) to see what will be done, then toggle the modify flag on and re-run the check.

The inconsistencies examined are the following:

- Duplicate filename records. This can happen if you accidentally run two copies of Bacula at the same time, and they are both adding filenames simultaneously. It is a rare occurrence, but will create an inconsistent database. If this is the case, you will receive error messages during Jobs warning of duplicate database records. If you are not getting these error messages, there is no reason to run this check.

- Repair bad Filename records. This checkes and corrects filenames that have a trailing slash. They should not.

- Repair bad Path records. This checks and corrects path names that do not have a trailing slash. They should.

- Duplicate path records. This can happen if you accidentally run two copies of Bacula at the same time, and they are both adding filenames simultaneously. It is a rare occurrence, but will create an inconsistent database. See the item above for why this occurs and how you know it is happening.

- Orphaned JobMedia records. This happens when a Job record is deleted (perhaps by a user issued SQL statement), but the corresponding JobMedia record (one for each Volume used in the Job) was not deleted. Normally, this should not happen, and even if it does, these records generally do not take much space in your database. However, by running this check, you can eliminate any such orphans.

- Orphaned File records. This happens when a Job record is deleted (perhaps by a user issued SQL statement), but the corresponding File record (one for each Volume used in the Job) was not deleted. Note, searching for these records can be **very** time consuming (i.e. it may take hours) for a large database. Normally this should not happen as Bacula takes care to prevent it. Just the same, this check can remove any orphaned File records. It is recommended that you run this once a year since orphaned File records can take a large amount of space in your database.

- Orphaned Path records. This condition happens any time a directory is deleted from your system and all associated Job records have been purged. During standard purging (or pruning) of Job records, Bacula does not check for orphaned Path records. As a consequence, over a period of time, old unused Path records will tend to accumulate and use space in your database. This check will eliminate them. It is strongly recommended that you run this check at least once a year.

- Orphaned Filename records. This condition happens any time a file is deleted from your system and all associated Job records have been purged. This can happen quite frequently as there are quite a large number of files that are created and then deleted. In addition, if you do a system update or delete an entire directory, there can be a very large number of Filename records that remain in the catalog but are no longer used.

  During standard purging (or pruning) of Job records, Bacula does not check for orphaned Filename records. As a consequence, over a period of time, old unused Filename records will accumulate and use space in your database. This check will eliminate them. It is strongly recommended that you run this check at least once a year, and for large database (more than 200 Megabytes), it is probably better to run this once every 6 months.

- Orphaned Client records. These records can remain in the database long after you have removed a client.

- Orphaned Job records. If no client is defined for a job or you do not run a job for a long time, you can accumulate old job records. This

option allow you to remove jobs that are not attached to any client (and thus useless).

- All Admin records. This command will remove all Admin records, regardless of their age.
- All Restore records. This command will remove all Restore records, regardless of their age.

testfind

testfind permits listing of files using the same search engine that is used for the **Include** resource in Job resources. Note, much of the functionality of this program (listing of files to be included) is present in the estimate command in the Console program.

The original use of testfind was to ensure that Bacula's file search engine was correct and to print some statistics on file name and path length. However, you may find it useful to see what bacula would do with a given **Include** resource. The **testfind** program can be found in the <**bacula-source**>**/src/tools** directory of the source distribution. Though it is built with the make process, it is not normally "installed".

It is called:

```
Usage: testfind [-d debug_level] [-] [pattern1 ...]
       -a          print extended attributes (Win32 debug)
       -dnn        set debug level to nn
       -           read pattern(s) from stdin
       -?          print this message.
Patterns are used for file inclusion -- normally directories.
Debug level>= 1 prints each file found.
Debug level>= 10 prints path/file for catalog.
Errors are always printed.
Files/paths truncated is a number with len> 255.
Truncation is only in the catalog.
```

Where a pattern is any filename specification that is valid within an **Include** resource definition. If none is specified, / (the root directory) is assumed. For example:

```
./testfind /bin
```

Would print the following:

```
Dir: /bin
Reg: /bin/bash
Lnk: /bin/bash2 -> bash
Lnk: /bin/sh -> bash
Reg: /bin/cpio
Reg: /bin/ed
Lnk: /bin/red -> ed
Reg: /bin/chgrp
...
Reg: /bin/ipcalc
Reg: /bin/usleep
Reg: /bin/aumix-minimal
Reg: /bin/mt
Lnka: /bin/gawk-3.1.0 -> /bin/gawk
Reg: /bin/pgawk
Total files    : 85
Max file length: 13
Max path length: 5
Files truncated: 0
Paths truncated: 0
```

Even though **testfind** uses the same search engine as **Bacula**, each directory to be listed, must be entered as a separate command line entry or entered one line at a time to standard input if the **-** option was specified.

Specifying a debug level of one (i.e. **-d1**) on the command line will cause **testfind** to print the raw filenames without showing the Bacula internal file type, or the link (if any). Debug levels of 10 or greater cause the filename and the path to be separated using the same algorithm that is used when putting filenames into the Catalog database.

bimagemgr

bimagemgr is a utility for those who backup to disk volumes in order to commit them to CDR disk, rather than tapes. It is a web based interface written in perl, used to monitor when a volume file needs to be burned to disk. It requires:

- A web server running on the bacula server
- A CD recorder installed and configured on the bacula server
- The cdrtools package installed on the bacula server.
- perl, perl-DBI module, and either DBD-MySQL or DBD-PostgreSQL modules

SQLite databases and DVD burning are not supported by **bimagemgr** at this time, but both planned for future releases.

bimagemgr installation

Please see the README file in the bimagemgr directory of the distribution for instructions.

bimagemgr usage

Calling the program in your web browser, e.g. `http://localhost/cgi-bin/bimagemgr.pl` will produce a display as shown below in Figure 1. The program will query the bacula database and display all volume files with the date last written and the date last burned to disk. If a volume needs to be burned (last written is newer than last burn date) a "Burn" button will be displayed in the right most column.

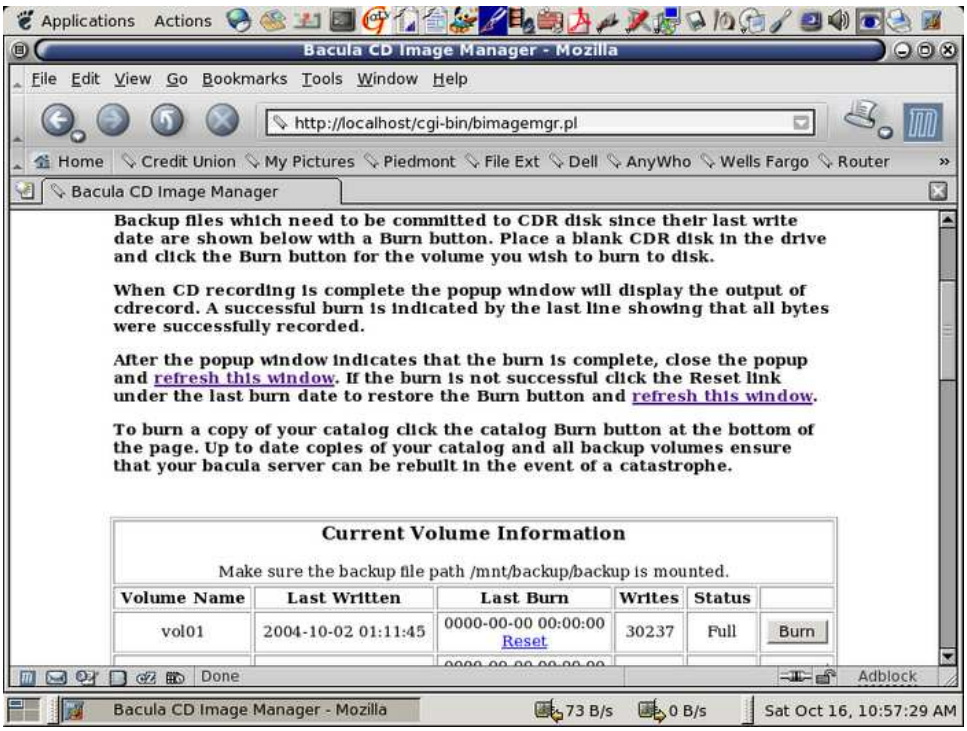

Figure 1

Place a blank CDR disk in your recorder and click a "Burn" button. This will cause a pop up window as shown in Figure 2 to display the burn progress.

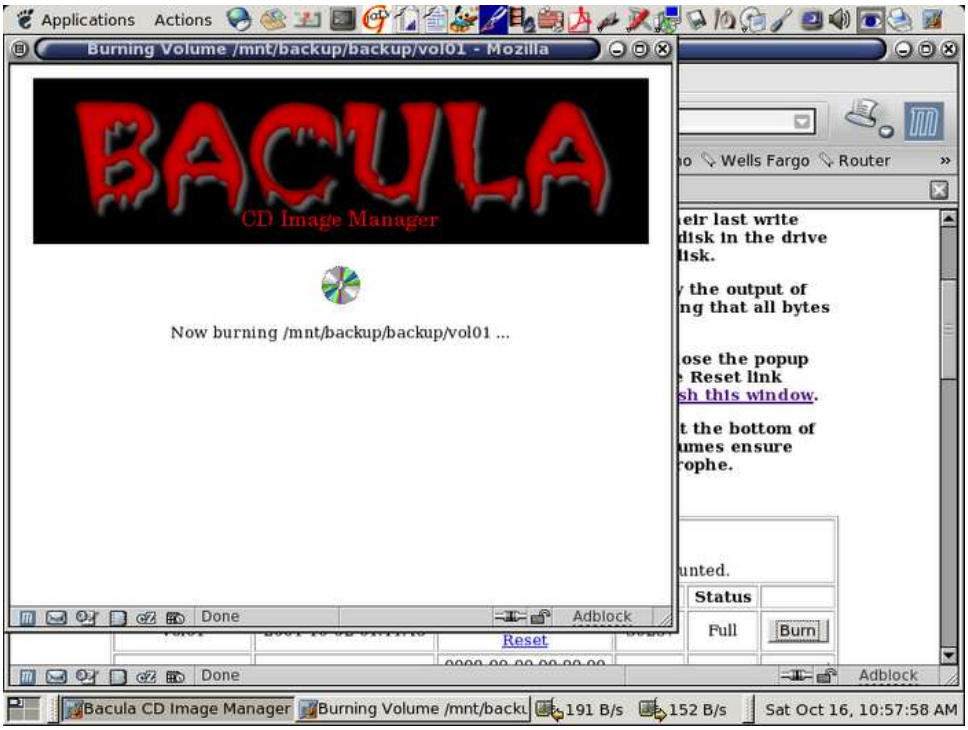

Figure 2

When the burn finishes the pop up window will display the results of cdrecord as shown in Figure 3. Close the pop up window and refresh the main window. The last burn date will be updated and the "Burn" button for that volume will disappear. Should you have a failed burn you can reset the last burn date of that volume by clicking it's "Reset" link.

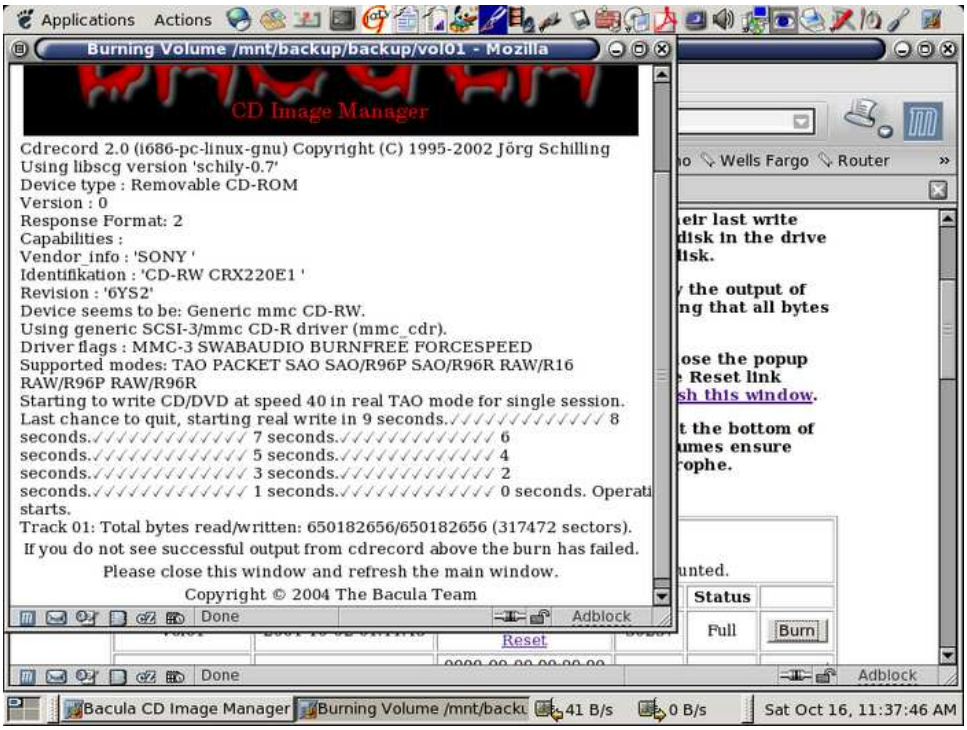

Figure 3

In the bottom row of the main display window are two more buttons labeled "Burn Catalog" and "Blank CDRW". "Burn Catalog" will place a copy of your bacula catalog on a disk. If you use CDRW disks rather than CDR then "Blank CDRW" allows you to erase the disk before re-burning it. Regularly committing your backup volume files and your catalog to disk with **bimagemgr** insures that you can rebuild easily in the event of some disaster on the bacula server itself.

Testing Your Tape Drive With Bacula

This chapter is concerned with testing and configuring your tape drive to make sure that it will work properly with Bacula using the **btape** program.

Summary of Steps to Take to Get Your Tape Drive Working

In general, you should follow the following steps to get your tape drive to work with Bacula. Start with a tape mounted in your drive. If you have an autochanger, load a tape into the drive. We use **/dev/nst0** as the tape drive name, you will need to adapt it according to your system.

Do not proceed to the next item until you have succeeded with the previous one.

1. Use tar to write to, then read from your drive:

```
mt -f /dev/nst0 rewind
tar cvf /dev/nst0 .
mt -f /dev/nst0 rewind
tar tvf /dev/nst0
```

2. Make sure you have a valid and correct Device resource corresponding to your drive. For Linux users, generally, the default one works. For FreeBSD users, there are two possible Device configurations (see below).

3. Run the btape **test** command:

```
./btape -c bacula-sd.conf /dev/nst0
test
```

   It isn't necessary to run the autochanger part of the test at this time, but do not go past this point until the basic test succeeds.

4. Run the btape **fill** command, preferrably with two volumes. This can take a long time. If you have an autochanger and it is configured, Bacula will automatically use it. If you do not have it configured, you can manual issue the appopriate **mtx** command, or press the autochanger buttons to change the tape when requested to do so.

5. FreeBSD users, run the **tapetest** program, and make sure your system is patched if necessary. See below for more details.

6. Run Bacula, and backup a reasonably small directory, say 60 Megabytes. Do three successive backups of this directory.

7. Stop Bacula, then restart it. Do another full backup of the same directory. Then stop and restart Bacula.

8. Do a restore of the directory backed up, by entering the following restore command, being careful to restore it to an alternate location:

```
restore select all done
yes
```

   Do a **diff** on the restored directory to ensure it is identical to the original directory.

9. If you have an autochanger, you should now go back to the btape program and run the autochanger test:

```
./btape -c bacula-sd.conf /dev/nst0
auto
```

   Adjust your autochanger as necessary to ensure that it works correctly. See the Autochanger chapter of this manual for a complete discussion of testing your autochanger.

If you have reached this point, you stand a good chance of having everything work. If you get into trouble at any point, **carefully** read the documentation given below. If you cannot get past some point, ask the **bacula-users** email list, but specify which of the steps you have successfully completed. In particular, you may want to look at the Tips for Resolving Problems section below.

Specifying the Configuration File

Starting with version 1.27, each of the tape utility programs including the **btape** program requires a valid Storage daemon configuration file (actually, the only part of the configuration file that **btape** needs is the **Device** resource definitions). This permits **btape** to find the configuration parameters for your archive device (generally a tape drive). Without those

parameters, the testing and utility programs do not no how to properly read and write your drive. By default, they use **bacula-sd.conf** in the current directory, but you may specify a different configuration file using the **-c** option.

Specifying a Device Name For a Tape

btape device-name where the Volume can be found. In the case of a tape, this is the physical device name such as **/dev/nst0** or **/dev/rmt/0ubn** depending on your system that you specify on the Archive Device directive. For the program to work, it must find the identical name in the Device resource of the configuration file. If the name is not found in the list of phsical names, the utility program will compare the name you entered to the Device names (rather than the Archive device names). See below for specifying Volume names.

Specifying a Device Name For a File

If you are attempting to read or write an archive file rather than a tape, the **device-name** should be the full path to the archive location including the filename. The filename (last part of the specification) will be stripped and used as the Volume name, and the path (first part before the filename) must have the same entry in the configuration file. So, the path is equivalent to the archive device name, and the filename is equivalent to the volume name.

btape

This program permits a number of elementary tape operations via a tty command interface. The **test** command, described below, can be very useful for testing tape drive compatibility problems. Aside from initial testing of tape drive compatibility with **Bacula**, **btape** will be mostly used by developers writing new tape drivers.

btape can be dangerous to use with existing **Bacula** tapes because it will relabel a tape or write on the tape if so requested regardless of whether or not the tape contains valuable data, so please be careful and use it only on blank tapes.

To work properly, **btape** needs to read the Storage daemon's configuration file. As a default, it will look for **bacula-sd.conf** in the current directory.

If your configuration file is elsewhere, please use the **-c** option to specify where.

The physical device name or the Device resource name must be specified on the command line, and that this same device name must be present in the Storage daemon's configuration file read by **btape**

```
Usage: btape [options] device_name
       -b <file>   specify bootstrap file
       -c <file>   set configuration file to file
       -d <nn>     set debug level to nn
       -s          turn off signals
       -v          be verbose
       -?          print this message.
```

Using btape to Verify your Tape Drive

An important reason for this program is to ensure that a Storage daemon configuration file is defined so that Bacula will correctly read and write tapes.

It is highly recommended that you run the **test** command before running your first Bacula job to ensure that the parameters you have defined for your storage device (tape drive) will permit **Bacula** to function properly. You only need to mount a blank tape, enter the command, and the output should be reasonably self explanatory. For example:

```
(ensure that Bacula is not running)
./btape -c /usr/bin/bacula/bacula-sd.conf /dev/nst0
```

The output will be:

```
Tape block granularity is 1024 bytes.
btape: btape.c:376 Using device: /dev/nst0
*
```

Enter the test command:

```
test
```

The output produced should be something similar to the following: I've cut the listing short because it is frequently updated to have new tests.

```
=== Append files test ===
This test is essential to Bacula.
I'm going to write one record  in file 0,
                   two records in file 1,
             and three records in file 2
btape: btape.c:387 Rewound /dev/nst0
btape: btape.c:855 Wrote one record of 64412 bytes.
btape: btape.c:857 Wrote block to device.
btape: btape.c:410 Wrote EOF to /dev/nst0
btape: btape.c:855 Wrote one record of 64412 bytes.
btape: btape.c:857 Wrote block to device.
btape: btape.c:855 Wrote one record of 64412 bytes.
btape: btape.c:857 Wrote block to device.
btape: btape.c:410 Wrote EOF to /dev/nst0
btape: btape.c:855 Wrote one record of 64412 bytes.
btape: btape.c:857 Wrote block to device.
btape: btape.c:855 Wrote one record of 64412 bytes.
btape: btape.c:857 Wrote block to device.
btape: btape.c:855 Wrote one record of 64412 bytes.
btape: btape.c:857 Wrote block to device.
btape: btape.c:410 Wrote EOF to /dev/nst0
btape: btape.c:387 Rewound /dev/nst0
btape: btape.c:693 Now moving to end of media.
btape: btape.c:427 Moved to end of media
We should be in file 3. I am at file 3. This is correct!
Now the important part, I am going to attempt to append to the tape.
...
=== End Append files test ===
```

If you do not successfully complete the above test, please resolve the problem(s) before attempting to use **Bacula**. Depending on your tape drive, the test may recommend that you add certain records to your configuration. We strongly recommend that you do so and then re-run the above test to insure it works the first time.

Some of the suggestions it provides for resolving the problems may or may not be useful. If at all possible avoid using fixed blocking. If the test suddenly starts to print a long series of:

```
Got EOF on tape.
Got EOF on tape.
...
```

then almost certainly, you are running your drive in fixed block mode rather than variable block mode. Please see below for help on resolving that.

For FreeBSD users, please see the notes below for doing further testing of your tape drive.

Linux SCSI Tricks

You can find out what SCSI devices you have by doing:

```
cat /proc/scsi/scsi
```

For example, I get the following:

```
Attached devices:
Host: scsi2 Channel: 00 Id: 01 Lun: 00
  Vendor: HP       Model: C5713A           Rev: H107
  Type:   Sequential-Access                ANSI SCSI revision: 02
Host: scsi2 Channel: 00 Id: 04 Lun: 00
  Vendor: SONY     Model: SDT-10000        Rev: 0110
  Type:   Sequential-Access                ANSI SCSI revision: 02
```

If you want to remove the SDT-10000 device, you can do so as root with:

```
echo "scsi remove-single-device 2 0 4 0">/proc/scsi/scsi
```

and you can put add it back with:

```
echo "scsi add-single-device 2 0 4 0">/proc/scsi/scsi
```

where the 2 0 4 0 are the Host, Channel, Id, and Lun as seen on the output from **cat /proc/scsi/scsi**. Note, the Channel must be specified as numeric.

Tips for Resolving Problems

Bacula Saves But Cannot Restore Files

If you are getting error messages such as:

```
Volume data error at 0:1! Wanted block-id: "BB02", got "". Buffer discarded
```

It is very likely that Bacula has tried to do block positioning and ended up at an invalid block. This can happen if your tape drive is in fixed block mode while Bacula's default is variable blocks. Note that in such cases, Bacula is

perfectly able to write to your Volumes (tapes), but cannot position to read them.

There are two possible solutions.

1. The first and best is to always ensure that your drive is in variable block mode. Note, it can switch back to fixed block mode on a reboot or if another program uses the drive. So on such systems you need to modify the Bacula startup files to explicitly set:

   ```
   mt -f /dev/nst0 defblksize 0
   ```

   or whatever is appropriate on your system.

2. The second possibility, especially, if Bacula wrote while the drive was in fixed block mode, is to turn off block positioning in Bacula. This is done by adding:

   ```
   Block Positioning = no
   ```

   to the Device resource. This is not the recommended procedure because it can enormously slow down recovery of files, but it may help where all else fails. This directive is available in version 1.35.5 or later (and not yet tested).

Bacula Cannot Open the Device

If you get an error message such as:

```
dev open failed: dev.c:265 stored: unable to open
device /dev/nst0:> ERR=No such device or address
```

the first time you run a job, it is most likely due to the fact that you specified the incorrect device name on your **Archive Device**.

If Bacula works fine with your drive, then all off a sudden you get error messages similar to the one shown above, it is quite possible that your driver module is being removed because the kernel deems it idle. This is done via **crontab** with the use of **rmmod -a**. To fix the problem, you can remove this entry from **crontab**, or you can manually **modprob** your driver module (or add it to the local startup script). Thanks to Alan Brown for this tip.

Incorrect File Number

When Bacula moves to the end of the medium, it normally uses the **ioctl(MTEOM)** function. Then Bacula uses the **ioctl(MTIOCGET)** function to retrieve the current file position from the **mt_fileno** field. Some SCSI tape drivers will use a fast means of seeking to the end of the medium and in doing so, they will not know the current file position and hence return a **-1**. As a consequence, if you get **"This is NOT correct!"** in the positioning tests, this may be the cause. You must correct this condition in order for Bacula to work.

There are two possible solutions to the above problem of incorrect file number:

- Figure out how to configure your SCSI driver to keep track of the file position during the MTEOM request. This is the preferred solution.
- Modify the **Device** resource of your **bacula-sd.conf** file to include:

  ```
  Hardware End of File = no
  ```

  This will cause Bacula to use the MTFSF request to seek to the end of the medium, and Bacula will keep track of the file number itself.

Incorrect Number of Blocks or Positioning Errors during btape Testing

Bacula's preferred method of working with tape drives (sequential devices) is to run in variable block mode, and this is what is set by default. You should first ensure that your tape drive is set for variable block mode (see below).

If your tape drive is in fixed block mode and you have told Bacula to use different fixed block sizes or variable block sizes (default), you will get errors when Bacula attempts to forward space to the correct block (the kernel driver's idea of tape blocks will not correspond to Bacula's).

All modern tape drives support variable tape blocks, but some older drives (in particular the QIC drives) as well as the ATAPI ide-scsi driver run only in fixed block mode. The Travan tape drives also apparently must run in fixed block mode (to be confirmed).

Even in variable block mode, with the exception of the first record on the second or subsequent volume of a multi-volume backup, Bacula will write blocks of a fixed size. However, in reading a tape, Bacula will assume that for each read request, exactly one block from the tape will be transferred. This the most common way that tape drives work and is well supported by **Bacula**.

Drives that run in fixed block mode can cause serious problems for Bacula if the drive's block size does not correspond exactly to **Bacula's** block size. In fixed block size mode, drivers may transmit a partial block or multiple blocks for a single read request. From **Bacula's** point of view, this destroys the concept of tape blocks. It is much better to run in variable block mode, and almost all modern drives (the OnStream is an exception) run in variable block mode. In order for Bacula to run in fixed block mode, you must include the following records in the Storage daemon's Device resource definition:

```
Minimum Block Size = nnn
Maximum Block Size = nnn
```

where **nnn** must be the same for both records and must be identical to the driver's fixed block size.

We recommend that you avoid this configuration if at all possible by using variable block sizes.

If you must run with fixed size blocks, make sure they are not 512 bytes. This is too small and the overhead that Bacula has with each record will become excessive. If at all possible set any fixed block size to something like 64,512 bytes or possibly 32,768 if 64,512 is too large for your drive. See below for the details on checking and setting the default drive block size.

To recover files from tapes written in fixed block mode, see below.

Ensuring that the Tape Modes Are Properly Set – Linux Only

If you have a modern SCSI tape drive and you are having problems with the **test** command as noted above, it may be that some program has set one or more of the your SCSI driver's options to non-default values. For example, if your driver is set to work in SysV manner, Bacula will not work correctly because it expects BSD behavior. To reset your tape drive to the default values, you can try the following, but **ONLY** if you have a SCSI tape drive on a **Linux** system:

```
become super user
mt -f /dev/nst0 rewind
mt -f /dev/nst0 stoptions buffer-writes async-writes read-ahead
```

The above commands will clear all options and then set those specified. None of the specified options are required by Bacula, but a number of other options such as SysV behavior must not be set. Bacula does not support SysV tape behavior. On systems other than Linux, you will need to consult your **mt** man pages or documentation to figure out how to do the same thing. This should not really be necessary though – for example, on both Linux and Solaris systems, the default tape driver options are compatible with Bacula.

You may also want to ensure that no prior program has set the default block size, as happened to one user, by explicitly turning it off with:

```
mt -f /dev/nst0 defblksize 0
```

If you would like to know what stoptions you have set before making any of the changes noted above, you can now view them on Linux systems, thanks to a tip provided by Willem Riede. Do the following:

```
become super user
mt -f /dev/nst0 stsetoptions 0
grep st0 /var/log/messages
```

and you will get output that looks something like the following:

```
kernel: st0: Mode 0 options: buffer writes: 1, async writes: 1, read ahead: 1
kernel: st0:    can bsr: 0, two FMs: 0, fast mteom: 0, auto lock: 0,
kernel: st0:    defs for wr: 0, no block limits: 0, partitions: 0, s2 log: 0
kernel: st0:    sysv: 0 nowait: 0
```

Note, I have chopped off the beginning of the line with the date and machine name for presentation purposes.

Some people find that the above settings only last until the next reboot, so please check this otherwise you may have unexpected problems.

Beginning with Bacula version 1.35.8, if Bacula detects that you are running in variable block mode, it will attempt to set your drive appropriately. All OSes permit setting variable block mode, but some OSes do not permit setting the other modes that Bacula needs to function properly.

Checking and Setting Tape Hardware Compression and Blocking Size

As far as I can tell, there is no way with the **mt** program to check if your tape hardware compression is turned on or off. You can, however, turn it on by using (on Linux):

```
become super user
mt -f /dev/nst0 defcompression 1
```

and of course, if you use a zero instead of the one at the end, you will turn it off.

You may also want to ensure that no prior program has set the default block size, as happened to one user, by explicitly turning it off with:

```
mt -f /dev/nst0 defblksize 0
```

If you have built the **mtx** program in the **depkgs** package, you can use tapeinfo to get quite a bit of information about your tape drive even if it is not an autochanger. This program is called using the SCSI control device. On Linux for tape drive /dev/nst0, this is usually /dev/sg0, while on FreeBSD for /dev/nsa0, the control device is often /dev/pass2. For example on my DDS-4 drive (/dev/nst0), I get the following:

```
tapeinfo -f /dev/sg0
Product Type: Tape Drive
Vendor ID: 'HP      '
Product ID: 'C5713A          '
Revision: 'H107'
Attached Changer: No
MinBlock:1
MaxBlock:16777215
SCSI ID: 5
SCSI LUN: 0
Ready: yes
BufferedMode: yes
Medium Type: Not Loaded
Density Code: 0x26
BlockSize: 0
```

where the **DataCompEnabled: yes** means that tape hardware compression is turned on. You can see it turn on and off (yes—no) by using the **mt** commands given above. Also, this output will tell you if the

BlockSize is non-zero and hence set for a particular block size. Bacula is not likely to work in such a situation because it will normally attempt to write blocks of 64,512 bytes, except the last block of the job which will generally be shorter. The first thing to try is setting the default block size to zero using the **mt -f /dev/nst0 defblksize 0** command as shown above. On FreeBSD, this would be something like: **mt -f /dev/nsa0 blocksize 0**.

If your tape drive requires fixed block sizes (very unusual), you can use the following records:

```
Minimum Block Size = nnn
Maximum Block Size = nnn
```

in your Storage daemon's Device resource to force Bacula to write fixed size blocks (where you sent nnn to be the same for both of the above records). This should be done only if your drive does not support variable block sizes, or you have some other strong reasons for using fixed block sizes. As mentioned above, a small fixed block size of 512 or 1024 bytes will be very inefficient. Try to set any fixed block size to something like 64,512 bytes or larger if your drive will support it.

Also, note that the **Medium Type** field of the output of **tapeinfo** reports **Not Loaded**, which is not correct. As a consequence, you should ignore that field as well as the **Attached Changer** field.

To recover files from tapes written in fixed block mode, see below.

Tape Modes on FreeBSD

On most FreeBSD systems such as 4.9 and most tape drives, Bacula should run with:

```
mt \  -f \  /dev/nsa0 \  seteotmodel \  2
mt \  -f \  /dev/nsa0 \  blocksize \  0
mt \  -f \ /dev/nsa0 \  comp \  enable
```

You might want to put those commands in a startup script to make sure your tape driver is properly initialized before running Bacula.

Then according to what the **btape test** command returns, you will probably need to set the following (see below for an alternative):

```
Hardware End of Medium = no
BSF at EOM = yes
Backward Space Record = no
Backward Space File = no
Fast Forward Space File = no
TWO EOF = yes
```

Then be sure to run some append tests with Bacula where you start and stop Bacula between appending to the tape, or use **btape** version 1.35.1 or greater, which includes simulation of stopping/restarting Bacula.

Please see the file **platforms/freebsd/pthreads-fix.txt** in the main Bacula directory concerning **important** information concerning compatibility of Bacula and your system. A much more optimal Device configuration is shown below, but does not work with all tape drives. Please test carefully before putting either into production.

Note, for FreeBSD 4.10-RELEASE, using a Sony TSL11000 L100 DDS4 w/Autochanger set to variable block size and DCLZ compression, Brian McDonald reports that to get Bacula to append correctly between Bacula executions, the correct values to use are:

```
mt \  -f \  /dev/nsa0 \  seteotmodel \  1
mt \  -f \  /dev/nsa0 \  blocksize \  0
mt \  -f \ /dev/nsa0 \  comp \  enable
```

and

```
Hardware End of Medium = no
BSF at EOM = no
Backward Space Record = no
Backward Space File = no
Fast Forward Space File = yes
TWO EOF = no
```

This has been confirmed by several other people using different hardware. This configuration is the preferred one because it uses one EOF and no backspacing at the end of the tape, which works much more efficiently and reliably with modern tape drives.

Determining What Tape Drives and Autochangers You Have on FreeBSD

On FreeBSD, you can do a **camcontrol devlist** as root to determine what drives and autochangers you have. For example,

```
undef# camcontrol devlist
    at scbus0 target 2 lun 0 (pass0,sa0)
    at scbus0 target 4 lun 0 (pass1,sa1)
    at scbus0 target 4 lun 1 (pass2)
```

from the above, you can determine that there is a tape drive on **/dev/sa0** and another on **/dev/sa1** in addition since there is a second line for the drive on **/dev/sa1**, you know can assume that it is the control device for the autochanger (i.e. **/dev/pass2**). It is also the control device name to use when invoking the tapeinfo program. E.g.

```
tapeinfo -f /dev/pass2
```

Using the OnStream driver on Linux Systems

Bacula version 1.33 (not 1.32x) is now working and ready for testing with the OnStream kernel osst driver version 0.9.14 or above. Osst is available from: http://sourceforge.net/projects/osst/.

To make Bacula work you must first load the new driver then, as root, do:

```
  mt -f /dev/nosst0 defblksize 32768
```

Also you must add the following to your Device resource in your Storage daemon's conf file:

```
 Minimum Block Size = 32768
 Maximum Block Size = 32768
```

Here is a Device specification provided by Michel Meyers that is known to work:

```
Device {
  Name = "Onstream DI-30"
  Media Type = "ADR-30"
  Archive Device = /dev/nosst0
  Minimum Block Size = 32768
  Maximum Block Size = 32768
  Hardware End of Medium = yes
  BSF at EOM = no
  Backward Space File = yes
```

```
  Fast Forward Space File = yes
  Two EOF = no
  AutomaticMount = yes
  AlwaysOpen = yes
  Removable Media = yes
}
```

Using btape to Simulate Bacula Filling a Tape

Because there are often problems with certain tape drives or systems when end of tape conditions occur, **btape** has a special command **fill** that causes it to write random data to a tape until the tape fills. It then writes at least one more Bacula block to a second tape. Finally, it reads back both tapes to ensure that the data has been written in a way that Bacula can recover it. Note, there is also a single tape option as noted below, which you should use rather than the two tape test. See below for more details.

This can be an extremely time consuming process (here is is about 6 hours) to fill a full tape. Note, that btape writes random data to the tape when it is filling it. This has two consequences: 1. it takes a bit longer to generate the data, especially on slow CPUs. 2. the total amount of data is approximately the real physical capacity of your tape, regardless of whether or not the tape drive compression is on or off. This is because random data does not compress very much.

To begin this test, you enter the **fill** command and follow the instructions. There are two options: the simple single tape option and the multiple tape option. Please use only the simple single tape option because the multiple tape option still doesn't work totally correctly. If the single tape option does not succeed, you should correct the problem before using Bacula.

Recovering Files Written to Tape With Fixed Block Sizes

If you have been previously running your tape drive in fixed block mode (default 512) and Bacula with variable blocks (default), then in version 1.32f-x and 1.34 and above, Bacula will fail to recover files because it does block spacing, and because the block sizes don't agree between your tape drive and Bacula it will not work.

The long term solution is to run your drive in variable block mode as described above. However, if you have written tapes using fixed block sizes, this can be a bit of a pain. The solution to the problem is: while you are

doing a restore command using a tape written in fixed block size, ensure that your drive is set to the fixed block size used while the tape was written. Then when doing the **restore** command in the Console program, do not answer the prompt **yes/mod/no**. Instead, edit the bootstrap file (the location is listed in the prompt) using any ASCII editor. Remove all **VolBlock** lines in the file. When the file is re-written, answer the question, and Bacula will run without using block positioning, and it should recover your files.

Tape Blocking Modes

SCSI tapes may either be written in **variable** or **fixed** block sizes. Newer drives support both modes, but some drives such as the QIC devices always use fixed block sizes. Bacula attempts to fill and write complete blocks (default 65K), so that in normal mode (variable block size), Bacula will always write blocks of the same size except the last block of a Job. If Bacula is configured to write fixed block sizes, it will pad the last block of the Job to the correct size. Bacula expects variable tape block size drives to behave as follows: Each write to the drive results in a single record being written to the tape. Each read returns a single record. If you request less byte than are in the record, only those number of bytes will be returned, but the entire logical record will have been read (the next read will retrieve the next record). Thus data from a single write is always returned in a single read, and sequentially written records are returned by sequential reads.

Bacula expects fixed block size tape drives to behave as follows: If a write length is greater than the physical block size of the drive, the write will be written as two blocks each of the fixed physical size. This a single write may become multiple physical records on the tape. (This is not a good situation). According to the documentation, one may never write an amount of data that is not the exact multiple of the blocksize (it is not specified if an error occurs or if the the last record is padded). When reading, it is my understanding that each read request reads one physical record from the tape. Due to the complications of fixed block size tape drives, you should avoid them if possible with Bacula, or you must be ABSOLUTELY certain that you use fixed block sizes within Bacula that correspond to the physical block size of the tape drive. This will ensure that Bacula has a one to one correspondence between what it writes and the physical record on the tape.

Please note that Bacula will not function correctly if it writes a block and that block is split into two or more physical records on the tape. Bacula assumes that each write causes a single record to be written, and that it can sequentially recover each of the blocks it has written by using the same number of sequential reads as it had written.

What To Do When Bacula Crashes (Kaboom)

If you are running on a Linux system, and you have a set of working configuration files, it is very unlikely that **Bacula** will crash. As with all software, however, it is inevitable that someday, it may crash, particularly if you are running on another operating system or using a new or unusual feature.

This chapter explains what you should do if one of the three **Bacula** daemons (Director, File, Storage) crashes.

Traceback

Each of the three Bacula daemons has a built-in exception handler which, in case of an error, will attempt to produce a traceback. If successful the traceback will be emailed to you.

For this to work, you need to ensure that a few things are setup correctly on your system:

1. You must have an installed copy of **gdb** (the GNU debugger), and it must be on **Bacula's** path.
2. The Bacula installed script file **btraceback** must be in the same directory as the daemon which dies, and it must be marked as executable.
3. The script file **btraceback.gdb** must have the correct path to it specified in the **btraceback** file.
4. You must have a **mail** program which is on **Bacula's** path.

If all the above conditions are met, the daemon that crashes will produce a traceback report and email it to you. If the above conditions are not true, you can either run the debugger by hand as described below, or you may be able to correct the problems by editing the **btraceback** file. I recommend not spending too much time on trying to get the traceback to work as it can be very difficult.

The changes that might needed are to add a correct path to the **gdb** program, correct the path to the **btraceback.gdb** file, change the **mail** program or its path, or change your email address. The key line in the **btraceback** file is:

```
gdb -quiet -batch -x /home/kern/bacula/bin/btraceback.gdb \
     $1 $2 2>\&1 | mail -s "Bacula traceback" your-address@xxx.com
```

Since each daemon has the same traceback code, a single btraceback file is sufficient if you are running more than one daemon on a machine.

Testing The Traceback

To "manually" test the traceback feature, you simply start **Bacula** then obtain the **PID** of the main daemon thread (there are multiple threads). Unfortunately, the output had to be split to fit on this page:

```
[kern@rufus kern]$ ps fax --columns 132 | grep bacula-dir
 2103 ?        S      0:00 /home/kern/bacula/k/src/dird/bacula-dir -c
                                 /home/kern/bacula/k/src/dird/dird.conf
 2104 ?        S      0:00  \_ /home/kern/bacula/k/src/dird/bacula-dir -c
                                 /home/kern/bacula/k/src/dird/dird.conf
 2106 ?        S      0:00      \_ /home/kern/bacula/k/src/dird/bacula-dir -c
                                 /home/kern/bacula/k/src/dird/dird.conf
 2105 ?        S      0:00      \_ /home/kern/bacula/k/src/dird/bacula-dir -c
                                 /home/kern/bacula/k/src/dird/dird.conf
```

which in this case is 2103. Then while Bacula is running, you call the program giving it the path to the Bacula executable and the **PID**. In this case, it is:

```
./btraceback /home/kern/bacula/k/src/dird 2103
```

It should produce an email showing you the current state of the daemon (in this case the Director), and then exit leaving **Bacula** running as if nothing happened. If this is not the case, you will need to correct the problem by modifying the **btraceback** script.

Typical problems might be that **gdb** is not on the default path. Fix this by specifying the full path to it in the **btraceback** file. Another common problem is that the **mail** program doesn't work or is not on the default path. On some systems, it is preferable to use **Mail** rather than **mail**.

Getting A Traceback On Other Systems

It should be possible to produce a similar traceback on systems other than Linux, either using **gdb** or some other debugger. Solaris with **gdb**

loaded works quite fine. On other systems, you will need to modify the **btraceback** program to invoke the correct debugger, and possibly correct the **btraceback.gdb** script to have appropriate commands for your debugger. If anyone succeeds in making this work with another debugger, please send us a copy of what you modified.

Manually Running Bacula Under The Debugger

If for some reason you cannot get the automatic traceback, or if you want to interactively examine the variable contents after a crash, you can run Bacula under the debugger. Assuming you want to run the Storage daemon under the debugger (the technique is the same for the other daemons, only the name changes), you would do the following:

1. Start the Director and the File daemon. If the Storage daemon also starts, you will need to find its PID as shown above (ps fax — grep bacula-sd) and kill it with a command like the following:

   ```
   kill -15 PID
   ```

   where you replace **PID** by the actual value.

2. At this point, the Director and the File daemon should be running but the Storage daemon should not.

3. cd to the directory containing the Storage daemon

4. Start the Storage daemon under the debugger:

   ```
   gdb ./bacula-sd
   ```

5. Run the Storage daemon:

   ```
   run -s -f -c ./bacula-sd.conf
   ```

   You may replace the **./bacula-sd.conf** with the full path to the Storage daemon's configuration file.

6. At this point, Bacula will be fully operational.

7. In another shell command window, start the Console program and do what is necessary to cause Bacula to die.

8. When Bacula crashes, the **gdb** shell window will become active and **gdb** will show you the error that occurred.

9. To get a general traceback of all threads, issue the following command:

```
thread apply all bt
```

   After that you can issue any debugging command.

Getting Debug Output from Bacula

Each of the daemons normally has debug compiled into the program, but disabled. There are two ways to enable the debug output. One is to add the **-d nnn** option on the command line when starting the debugger. The **nnn** is the debug level, and generally anything between 50 and 200 is reasonable. The higher the number, the more output is produced. The output is written to standard output.

The second way of getting debug output is to dynamically turn it on using the Console using the **setdebug** command. The full syntax of the command is:

```
setdebug level=nnn client=client-name storage=storage-name dir
```

If none of the options are given, the command will prompt you. You can selectively turn on/off debugging in any or all the daemons (i.e. it is not necessary to specify all the components of the above command).

The Windows Version of Bacula

General

At the current time only the File daemon or Client program has been tested on Windows. As a consequence, when we speak of the Windows version of Bacula below, we are referring to the File daemon only.

The Windows version of the Bacula File daemon has been tested on Win98, WinMe, WinNT, and Win2000 systems. We have coded to support Win95, but no longer have a system for testing. The Windows version of Bacula is a native Win32 port, but there are very few source code changes, which means that the Windows version is for the most part running code that has long proved stable on Unix systems. When running, it is perfectly integrated with Windows and displays its icon in the system icon tray, and provides a system tray menu to obtain additional information on how Bacula is running (status and events dialog boxes). If so desired, it can also be stopped by using the system tray menu, though this should normally never be necessary.

Once installed Bacula normally runs as a system service. This means that it is immediately started by the operating system when the system is booted, and runs in the background even if there is no user logged into the system.

Win32 Installation

Normally, you will install the Windows version of Bacula from the binaries. This install is standard Windows .exe that runs an install wizard using the NSIS Free Software installer, so if you have already installed Windows software, it should be very familiar to you.

If you have a previous version Cygwin of Bacula (1.32 or lower) installed, you should stop the service, uninstall it, and remove the directory possibly saving your bacula-fd.conf file for use with the new version you will install. The new native version of Bacula has far fewer files than the old Cygwin version.

Finally, proceed with the installation.

- Simply double click on the **winbacula-1.xx.0.exe** NSIS install icon. The actual name of the icon will vary from one release version to another.

winbacula-1.xx.0.exe

- Once launched, the installer wizard will ask you if you want to install Bacula.

- If you proceed, you will be asked to select the components to be installed. You may install the Bacula program (Bacula File Service) and or the documentation. Both will be installed in sub-directories of the install location that you choose later. The components dialog looks like the following:

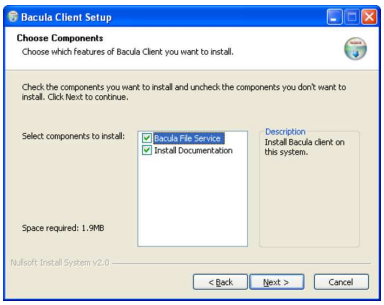

- Next you will be asked to select an installation directory.

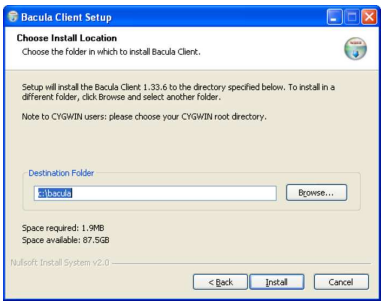

- If you are installing for the first time, you will be asked if you want to edit the bacula-fd.conf file, and if you respond with yes, it will be opened in notepad.

- Then the installer will ask if wish to install Bacula as a service. You should always choose to do so:

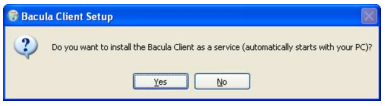

- If everything goes well, you will receive the following confirmation:

- Then you will be asked if you wish to start the service. If you respond with yes, any running Bacula will be shutdown and the new one started. You may see a DOS box momentarily appear on the screen as the service is started. It should disappear in a second or two:

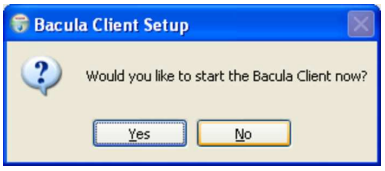

- Finally, the finish dialog will appear:

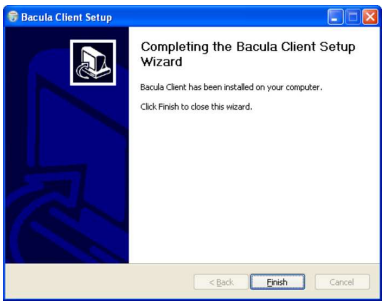

That should complete the installation process. When the Bacula File Server is ready to serve files, an icon  representing a cassette (or tape) will appear in the system tray 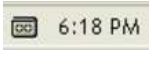

; right click on it and a

menu will appear. 

The **Events** item is currently unimplemented, by selecting the **Status** item, you can verify whether any jobs are running or not.

When the Bacula File Server begins saving files, the color of the holes in the cassette icon will change from white to green  , and if there is an error, the holes in the cassette icon will change to red  .

If you are using remote desktop connections between your windows boxes, be warned that that tray icon does not always appear. It will always be visible when you log into the console, but the remote desktop may not display it.

Post Win32 Installation

After installing Bacula and before running it, you should check the contents of **c:\bacula\bin\bacula-fd.conf** to ensure that it corresponds to your configuration.

Uninstalling Bacula on Win32

Once Bacula has been installed, it can be uninstalled using the standard Windows Add/Remove Programs dialog found on the Control panel.

Dealing with Win32 Problems

The most likely source of problems is authentication when the Director attempts to connect to the File daemon that you installed. This can occur if the names and the passwords defined in the File daemon's configuration file **c:\bacula\bin\bacula-fd.conf** on the Windows machine do not match with the names and the passwords in the Director's configuration file **bacula-dir.conf** located on your Unix/Linux server.

More specifically, the password found in the **Client** resource in the Director's configuration file must be the same as the password in the **Director** resource of the File daemon's configuration file. In addition, the name of the

Director resource in the File daemon's configuration file must be the same as the name in the **Director** resource of the Director's configuration file.

It is a bit hard to explain in words, but if you understand that a Director normally has multiple Clients and a Client (or File daemon) may permit access by multiple Directors, you can see that the names and the passwords on both sides must match for proper authentication.

One user had serious problems with the configuration file until he realized that the Unix end of line conventions were used and Bacula wanted them in Windows format. This has not been confirmed though.

Running Unix like programs on Windows machines is a bit frustrating because the Windows command line shell (DOS Window) is rather primitive. As a consequence, it is not generally possible to see the debug information and certain error messages that Bacula prints. With a bit of work, however, it is possible. When everything else fails and you want to **see** what is going on, try the following:

```
Start a DOS shell Window.
cd c:\bacula\bin
bacula-fd -t >out
type out
```

The **-t** option will cause Bacula to read the configuration file, print any error messages and then exit. the > redirects the output to the file named **out**, which you can list with the **type** command.

If something is going wrong later, or you want to run **Bacula** with a debug option, you might try starting it as:

```
bacula-fd -d 100 >out
```

In this case, Bacula will run until you explicitly stop it, which will give you a chance to connect to it from your Unix/Linux server. In later versions of Bacula (1.34 on, I think), when you start the File daemon in debug mode it can write the output to a trace file **bacula.trace** in the current directory. To enable this, before running a job, use the console, and enter:

```
trace on
```

then run the job, and once you have terminated the File daemon, you will find the debug output in the **bacula.trace** file.

In addition, you should look in the System Applications log on the Control Panel to find any Windows errors that Bacula got during the startup process.

Finally, due to the above problems, when you turn on debugging, and specify trace=1 on a setdebug command in the Console, Bacula will write the debug information to the file **bacula.trace** in the directory from which Bacula is executing.

Windows Compatibility Considerations

If any applications are running during the backup and they have files opened exclusively, Bacula will not be able to backup those files, so be sure you close your applications (or tell your users to close their applications) before the backup. Most Microsoft applications do not open files exclusively so that they can be backed up. However, you will need to experiment. In any case, if Bacula cannot open the file, it will print an error message, so you will always know which files were not backed up.

During backup, Bacula doesn't know about the system registry, so you will either need to write it out to an ASCII file using **regedit /e** or use a program specifically designed to make a copy or backup the registry.

In Bacula version 1.31 and later, we use Windows backup API calls by default. Typical of Windows, programming these special BackupRead and BackupWrite calls is a real nightmare of complications. The end result gives some distinct advantages and some disadvantages.

First, the advantages are that on WinNT/2K/XP systems, the security and ownership information is now backed up. In addition, with the exception of files in exclusive use by another program (a major disaster for backup programs on Windows), Bacula can now access all system files. This means that when you restore files, the security and ownership information will be restored on WinNT/2K/XP along with the data.

The disadvantage of the Windows backup API calls is that it produces non-portable backups. That is files and their data that are backed up on WinNT using the native API calls (BackupRead/BackupWrite) cannot be restored on Win95/98/Me or Unix systems. In principle, a file backed up on WinNT can be restored on WinXP, but this remains to be seen in practice (not yet tested). In addition, the stand-alone tools such as **bls** and **bextract** cannot be used to retrieve the data for those files because those tools are not available on Windows. All restores must use the Bacula **restore** command. This restriction is mentioned for completeness, but in practice should not create any problems.

As a default, Bacula backs up Windows systems using the Windows API calls. If you want to backup data on a WinNT/2K/XP system and restore it on a Unix/Win95/98/Me system, we have provided a special **portable** option that backups the data in a portable fashion by using portable API calls. See the portable option on the Include statement in a FileSet resource in the Director's configuration chapter for the details on setting this option. However, using the portable option means you may have permissions problems accessing files, and none of the security and ownership information will be backed up or restored. The file data can, however, be restored on any system.

You should always be able to restore any file backed up on Unix or Win95/98/Me to any other system. On some systems, such as WinNT/2K/XP, you may have to reset the ownership of such restored files. Any file backed up on WinNT/2K/XP should in principle be able to be restored to a similar system (i.e. WinNT/2K/XP), however, I am unsure of the consequences if the owner information and accounts are not identical on both systems. Bacula will not let you restore files backed up on WinNT/2K/XP to any other system (i.e. Unix Win95/98/Me) if you have used the defaults.

Finally, if you specify the **portable=yes** option on the files you back up. Bacula will be able to restore them on any other system. However, any WinNT/2K/XP specific security and ownership information will be lost.

The following matrix will give you an idea of what you can expect. Thanks to Marc Brueckner for doing the tests:

+

Backup OS	**Restore OS**	**Results**
WinMe	WinMe	Works
WinMe	WinNT	Works (SYSTEM permissions)
WinMe	WinXP	Works (SYSTEM permissions)
WinMe	Linux	Works (SYSTEM permissions)
WinXP	WinXP	Works
WinXP	WinNT	Works (all files OK, but got "The data is invalid" message)
WinXP	WinMe	Error: Win32 data stream not supported.
WinXP	WinMe	Works if **Portable=yes** specified during backup.
WinXP	Linux	Error: Win32 data stream not supported.

WinXP	Linux	Works if **Portable=yes** specified during backup.
WinNT	WinNT	Works
WinNT	WinXP	Works
WinNT	WinMe	Error: Win32 data stream not supported.
WinNT	WinMe	Works if **Portable=yes** specified during backup.
WinNT	Linux	Error: Win32 data stream not supported.
WinNT	Linux	Works if **Portable=yes** specified during backup.
Linux	Linux	Works
Linux	WinNT	Works (SYSTEM permissions)
Linux	WinMe	Works
Linux	WinXP	Works (SYSTEM permissions)

Windows Firewalls

If you turn on the firewalling feature on Windows (default in WinXP SR2), you are likely to find that the Bacula ports are blocked and you cannot communicated to the other daemons. This can be deactivated through the **Security Notification** dialog, which is apparently somewhere in the **Security Center**. I don't have this on my computer, so I cannot give the exact details.

The command:

```
netsh firewall set opmode disable
```

is purported to disable the firewall, but this command is not accepted on my WinXP Home machine.

Windows Port Usage

If you want to see if the File daemon has properly opened the port and is listening, you can enter the following command in a shell window:

```
netstat -an | findstr 910[123]
```

Windows Disaster Recovery

We don't currently have a good solution for disaster recovery on Windows as we do on Linux. The main piece lacking is a Windows boot floppy or a Windows boot CD. Microsoft releases a Windows Pre-installation Environment (**WinPE**) that could possibly work, but we have not investigated it. This means that until someone figures out the correct procedure, you must restore the OS from the installation disks, then you can load a Bacula client and restore files. Please don't count on using **bextract** to extract files from your backup tapes during a disaster recovery unless you have backed up those files using the **portable** option. **bextract** does not run on Windows, and the normal way Bacula saves files using the Windows API prevents the files from being restored on a Unix machine. Once you have an operational Windows OS loaded, you can run the File daemon and restore your user files.

Please see Disaster Recovery of Win32 Systems for the latest suggestion, which looks very promising.

It looks like Bart PE Builder, which creates a Windows PE (Pre-installation Environment) Boot-CD, may be just what is needed to build a complete disaster recovery system for Win32. This distribution can be found at http://www.nu2.nu/pebuilder/ .

Windows Ownership and Permissions Problems

If you restore files backed up from WinNT/XP/2K to an alternate directory, Bacula may need to create some higher level directories that were not saved (or restored). In this case, the File daemon will create them under the SYSTEM account because that is the account that Bacula runs under as a service. As of version 1.32f-3, Bacula creates these files with full access permission. However, there may be cases where you have problems accessing those files even if you run as administrator. In principle, Microsoft supplies you with the way to cease the ownership of those files and thus change the permissions. However, a much better solution to working with and changing Win32 permissions is the program **SetACL**, which can be found at http://setacl.sourceforge.net/ .

Manually resetting the Permissions

The following solution was provided by Dan Langille <dan at langille in the dot org domain>. The steps are performed using Windows 2000 Server but they should apply to most Win32 platforms. The procedure outlines how to deal with a problem which arises when a restore creates a top-level new directory. In this example, "top-level" means something like **c:\src**, not **c:\tmp\src** where **c:\tmp** already exists. If a restore job specifies / as the **Where:** value, this problem will arise.

The problem appears as a directory which cannot be browsed with Windows Explorer. The symptoms include the following message when you try to click on that directory:

If you encounter this message, the following steps will change the permissions to allow full access.

1. right click on the top level directory (in this example, **c:/src**) and select **Properties**.

2. click on the Security tab.

3. If the following message appears, you can ignore it, and click on **OK**.

You should see something like this:

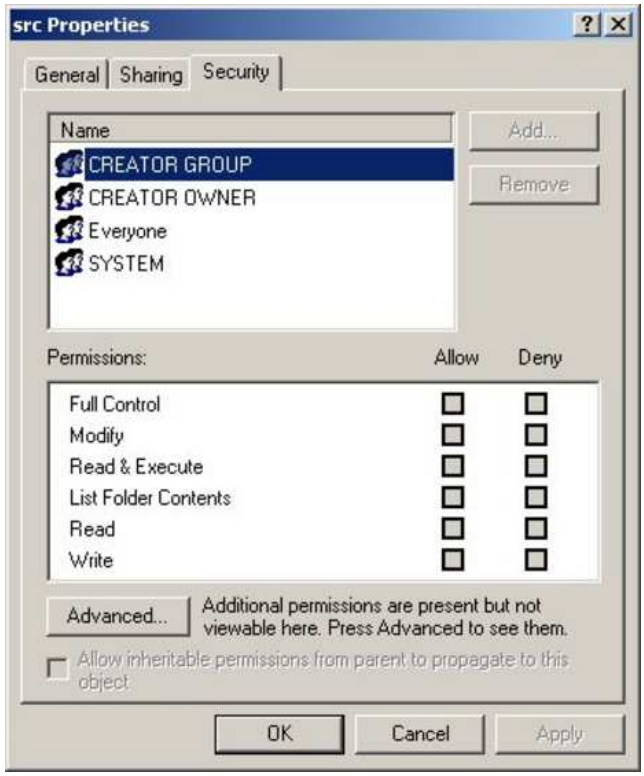

4. click on Advanced

5. click on the Owner tab

6. Change the owner to something other than the current owner (which is **SYSTEM** in this example as shown below).

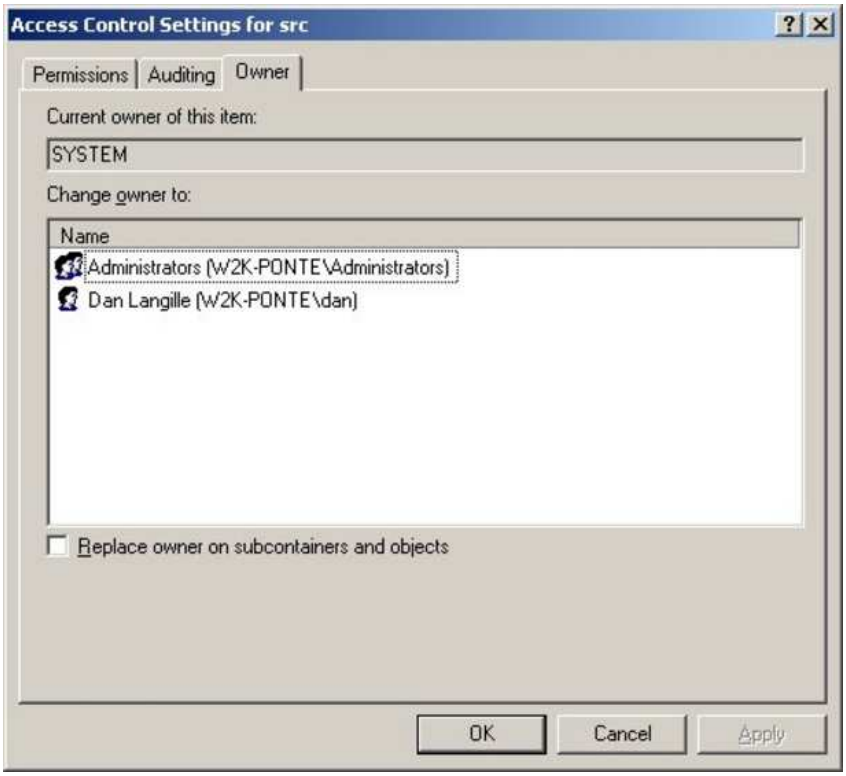

7. ensure the "Replace owner on subcontainers and objects" box is checked

8. click on OK

9. When the message "You do not have permission to read the contents of directory "c:\src\basis. Do you wish to replace the directory permissions with permissions granting you Full Control¿', click on Yes.

10. Click on OK to close the Properties tab

With the above procedure, you should now have full control over your restored directory.

Backing Up the WinNT/XP/2K System State

A suggestion by Damian Coutts using Microsoft's NTBackup utility in conjunction with Bacula should permit a full restore of any damaged system files on Win2K/XP. His suggestion is to do an NTBackup of the critical system state prior to running a Bacula backup with the following command:

```
ntbackup backup systemstate /F c:\systemstate.bkf
```

The **backup** is the command, the **systemstate** says to backup only the system state and not all the user files, and the **/F c:\systemstate.bkf** specifies where to write the state file. this file must then be saved and restored by Bacula.

To restore the system state, you first reload a base operating system if the OS is damaged, otherwise, this is not necessary, then you would use Bacula to restore all the damaged or lost user's files and to recover the **c:\systemstate.bkf** file. Finally if there are any damaged or missing system files or registry problems, you run **NTBackup** and **catalogue** the system statefile, and then select it for restore. The documentation says you can't run a command line restore of the systemstate.

To the best of my knowledge, this has not yet been tested. If you test it, please report your results to the Bacula email list.

Windows Considerations for Filename Specifications

Please see the Director's Configuration chapter of this manual for important considerations on how to specify Windows paths in Bacula FileSet Include and Exclude directives.

Command Line Options Specific to the Bacula Windows File Daemon (Client)

These options are not normally seen or used by the user, and are documented here only for information purposes. At the current time, to change the default options, you must either manually run **Bacula** or you must manually edit the system registry and modify the appropriate entries.

In order to avoid option clashes between the options necessary for **Bacula** to run on Windows and the standard Bacula options, all Windows specific

options are signaled with a forward slash character (/), while as usual, the standard Bacula options are signaled with a minus (-), or a minus minus (--). All the standard Bacula options can be used on the Windows version. In addition, the following Windows only options are implemented:

/servicehelper Run the service helper application (don't use this it is deprecated.).

/service Start Bacula as a service

/run Run the Bacula application

/install Install Bacula as a service in the system registry

/remove Uninstall Bacula from the system registry

/about Show the Bacula about dialogue box

/status Show the Bacula status dialogue box

/events Show the Bacula events dialogue box (not yet implemented)

/kill Stop any running **Bacula**

/help Show the Bacula help dialogue box

It is important to note that under normal circumstances the user should never need to use these options as they are normally handled by the system automatically once Bacula is installed. However, you may note these options in some of the .bat files that have been created for your use.

Shutting down Windows Systems

Some users like to shutdown their windows machines after a backup using a Client Run After Job directive. If you want to do something similar, you might take the shutdown program from the apcupsd project or one from the Sysinternals project.

Disaster Recovery Using Bacula

General

When disaster strikes, you must have a plan, and you must have prepared in advance otherwise the work of recovering your system and your files will be considerably greater. For example, if you have not previously saved the partitioning information for your hard disk, how can you properly rebuild it if the disk must be replaced?

Unfortunately, many of the steps one must take before and immediately after a disaster are very operating system dependent. As a consequence, this chapter will discuss in detail disaster recovery (also called Bare Metal Recovery) for **Linux** and **Solaris**. For Solaris, the procedures are still quite manual. For FreeBSD the same procedures may be used but they are not yet developed. For Win32, no luck. Apparently an "emergency boot" disk allowing access to the full system API without interference does not exist.

Important Considerations

Here are a few important considerations concerning disaster recovery that you should take into account before a disaster strikes.

- If the building which houses your computers burns down or is otherwise destroyed, do you have off-site backup data?
- Disaster recovery is much easier if you have several machines. If you have a single machine, how will you handle unforeseen events if your only machine is down?
- Do you want to protect your whole system and use Bacula to recover everything? or do you want to try to restore your system from the original installation disks and apply any other updates and only restore user files?

Steps to Take Before Disaster Strikes

- Create a Bacula Rescue CDROM for each of your Linux systems. Note, it is possible to create one CDROM by copying the bacula-hostname directory from each machine to the machine where you will be burning the CDROM.

- Ensure that you always have a valid bootstrap file for your backup that is saved to an alternate machine. This will permit you to easily do a full restore of your system.

- If possible copy your catalog nightly to an alternate machine. If you have a valid bootstrap file, this is not necessary, but can be very useful if you do not want to reload everything. .

- Ensure that you always have a valid bootstrap file for your catalog backup that is saved to an alternate machine. This will permit you to restore your catalog more easily if needed.

- Test using the Bacula Rescue CDROM before you are forced to use it in an emergency situation.

Bare Metal Recovery on Linux with a Bacula Rescue CDROM

The remainder of this section concerns recovering a **Linux** computer, and parts of it relate to the Red Hat version of Linux. The **Solaris** procedures can be found below under the Solaris Bare Metal Recovery section of this chapter.

If you wish to use a floppy for restoration, please see the chapter Bare Metal Floppy Recovery on Linux with a Bacula Floppy Rescue Disk, but be aware that the Bacula floppy disk is deprecated and replaced by the CDROM rescue described in this chapter.

A so called "Bare Metal" recovery is one where you start with an empty hard disk and you restore your machine. There are also cases where you may lose a file or a directory and want it restored. Please see the previous chapter for more details for those cases.

Bare Metal Recovery assumes that you have the following items for your system:

- A Bacula Rescue CDROM containing a copy of your OS and a copy of your hard disk information, as well as a statically linked version of the Bacula File daemon. This chapter describes how to build such a CDROM.

- A full Bacula backup of your system possibly including Incremental or Differential backups since the last Full backup

- A second system running the Bacula Director, the Catalog, and the Storage daemon. (this is not an absolute requirement, but how to get around it is not yet documented here)

Requirements

In addition, to the above assumptions, the following conditions or restrictions apply:

- Linux only – tested only on Red Hat, but should work on other Linuxes
- The scripts handle only SCSI and IDE disks
- All partitions will be recreated, but only **ext2**, **ext3**, **rfs** and **swap** partitions will be reformatted. Any other partitions such as Windows FAT partitions will not be formatted by the scripts, but you can do it by hand
- You are using either **lilo** or **grub** as a boot loader, and you know which one (not automatically detected)
- The partitioning and reformating scripts will *should* work with RAID devices, but probably not with other "complicated" disk partitioning/formating schemes. They also should work with Reiser filesystems. Please check them carefully. You will probably need to edit the scripts by hand to make them work.
- You will need mkisofs (might be part of cdrtools, but is a separate rpm on my system); cdrecord or some other tool for burning the CDROM.

Directories

To build the Bacula Rescue CDROM, you will find the necessary scripts in **rescue/linux/cdrom** subdirectory of the Bacula source code. If you installed the bacula-rescue rpm package the scripts will be found in the **/etc/bacula/rescue/cdrom** directory.

Preparation for a Bare Metal Recovery

Before you can do a Bare Metal recovery, you must create a Bacula Rescue CDROM, which will contain everything you need to begin recovery. This

assumes that you will have your Directory and Storage daemon running on a different machine. If you want to recover a machine where the Director and/or the database were previously running things will be much more complicated.

Creating a Bacula Rescue CDROM

The primary goals of the Bacula rescue CD are:

- NOT to be a general or universal recovery disk.
- to capture and setup a restore environment for a single system running as a Client.
- to capture the current state of the hard disks on your system, so that they can be easily restored from pre-generated scripts. Note, this is not done by any other rescue CDROM, as far as I am aware.
- to create and save a statically linked copy of your current Bacula FD. Thus you need no packages or other software to be installed before using this CDROM and the Bacula File daemon on it.
- to be relatively easy to create. In most cases you simply type **make all** in the **rescue/linux/cdrom** directory, then burn the ISO image created. In contrast, if you have looked at any of the documentation on how to remaster a CD or how to roll your own, your head will spin (at least mine did).
- to be easy for you to add any additional files, binaries, or libraries to the CD.
- to build and work on any (or almost any) Linux flavor or release.
- you might ask why I don't use Knoppix or some other preprepared recovery disk, especially since Knoppix is very kind and provides the Bacula FD on their disk. The answer is that: I am more comfortable having my Linux boot up in rescue mode rather than another flavor. In addition, the Bacula rescue CDROM contains a complete snapshot of your disk partitioning, which is not the case with any other rescue disk. If your harddisk dies, do you remember all the partitions you had and how big they are? I don't, and without that information, you have little hope of reformatting your harddisk and rebuilding your system.

One of the main of the advantages of a Bacula Rescue CDROM is that it contains a bootable copy of your system, so you should be familiar with it.

You should probably make a new rescue CDROM each time you make any major updates to your kernel, and every time you upgrade a major version of Bacula.

The whole process with the exception of burning the CDROM is done with the following commands:

```
(Build a working version of Bacula in the
 bacula-source directory)
cd <bacula-source>
./configure (your options)
make
cd <bacula-source>/rescue/linux/cdrom
su (become root)
make all
```

For users of the bacula-rescue rpm the static bacula-fd has already been built and placed in **/etc/bacula/rescue/cdrom/bin/** along with a symbolic link to your **/etc/bacula/bacula-fd.conf** file. Rpm users only need to do the second step:

```
cd /etc/bacula/rescue/cdrom
su (become root)
make all
```

At this point, if the scripts are successful, they should have done the following things:

- Made a copy of your kernel and its essential files.
- Copied a number of binary files from your system.
- Copied all the necessary shared libraries to run the above binary files.
- Made a statically-linked version of your File daemon and copied it into the CDROM build area.
- Made an ISO image and left it in **bootcd.iso**

Once this is accomplished, you need only burn it into a CDROM. This can be done directly from the makefile with:

```
make burn
```

However, you may need to modify the Makefile to properly specify your CD burner as the detection process is complicated especially if you have two CDROMs or do not have **cdrecord** loaded on your system. Users of the rescue rpm package should definitely examine the Makefile since it was configured on the host used to produce the rpm package. If you find that the **make burn** does not work for you, try doing a:

```
make scan
```

and use the output of that to modify the Makefile accordingly.

The "make all" that you did above actually does the equivalent to the following:

```
make kernel
make binaries
make bacula
make iso
```

If you wish, you can modify what you put on the CDROM and redo any part of the make that you wish. For example, if you want to add a new directory, you might do the first three makes, then add a new directory to the CDROM, and finally do a "make iso". Please see the README file in the **rescue/linux/cdrom** or **/etc/bacula/rescue/cdrom**directory for instructions on changing the contents of the CDROM.

At the current time, the size of the CDROM is about 50MB (compressed to about 20MB), so there is quite a bit more room for additional program. Keep in mind that when this CDROM is booted, *everything* is in memory, so the total size cannot exceed your memory size, and even then you will need some reserve memory for running programs, ...

Putting Two or More Systems on Your Rescue Disk

You can put multiple systems on the same rescue CD if you wish. This is because the information that is specific to your OS will be stored in the **/bacula-hostname** directory, where **hostname** is the name of the host on which you are building the CD. Suppose for example, you have two systems. One named **client1** and one named **client2**. Assume also that your CD burner is on client1, and that is the machine we start on, and that we can ssh into client2 and also client2's disks are mounted on client1.

```
ssh client2
cd <bacula-source>
./configure (your options)
make
cd rescue/linux/cdrom
su
(enter root password)
make bacula
exit
exit
```

Again, for rpm package users the above command set would be:

```
ssh client2
cd /etc/bacula/rescue/cdrom
su
(enter root password)
make bacula
exit
exit
```

Thus we have just built a Bacula rescue directory on client2. Now, on client1, we copy the appropriate directory to two places (explained below), then build an ISO and burn it:

```
cd <bacula-source>
./configure (your options)
make
cd rescue/linux/cdrom
su
(enter root password)
c=/mnt/client2/home/user/bacula/rescue/linux/cdrom
cp -a $c/roottree/bacula-client2 roottree
cp -a $c/roottree/bacula-client2 cdtree
make all
make burn
exit
```

And with the rpm package:

```
cd /etc/bacula/rescue/cdrom
su
(enter root password)
c=/mnt/client2/etc/bacula/rescue/cdrom
cp -a $c/roottree/bacula-client2 roottree
cp -a $c/roottree/bacula-client2 cdtree
make all
make burn
exit
```

In summary, with the above commands, we first build a Bacula directory on client2 in roottree/bacula-client2, then we copied the bacula-client2 directory into the client1's roottree so it is available in memory after booting, and we also copied it into the cdtree so it will also be on the CD as a separate directory and thus can be read without booting the CDROM. Then we made and burned the CDROM for client1, which of course, contains the client2 data.

Restoring a Client System

Now, let's assume that your hard disk has just died and that you have replaced it with an new identical drive. In addition, we assume that you have:

1. A recent Bacula backup (Full plus Incrementals)
2. A Bacula Rescue CDROM.
3. Your Bacula Director, Catalog, and Storage daemon running on another machine on your local network.

This is a relatively simple case, and later in this chapter, as time permits, we will discuss how you might recover from a situation where the machine that crashes is your main Bacula server (i.e. has the Director, the Catalog, and the Storage daemon).

You will take the following steps to get your system back up and running:

1. Boot with your Bacula Rescue CDROM.
2. Start the Network (local network)
3. Re-partition your hard disk(s) as it was before
4. Re-format your partitions
5. Restore the Bacula File daemon (static version)
6. Perform a Bacula restore of all your files
7. Re-install your boot loader
8. Reboot

Now for the details ...

Boot with your Bacula Rescue CDROM

When the CDROM boots, you will be presented with a script that looks like:

```
      Welcome to the Bacula Rescue Disk 1.1.0
To proceed, press the <ENTER> key or type "linux <runlevel>"

   linux 1     -> shell
   linux 2     -> login  (default if ENTER pressed)
   linux 3     -> network started and login (network not working yet)
   linux debug -> print debug during boot then login
```

Normally, at this point, you simply press ENTER. However, you may supply options for the boot if you wish.

Once it has booted, you will be requested to login something like:

```
Welcome to the Bacula Rescue CDROM
2.4.21-15.0.4.EL #1 Wed Aug 4 03:08:03 EDT 2004
Please login using root and your root password ...
RescueCD login:
```

Note, you must enter the root password for the system on which you loaded the kernel or on which you did the build of the CDROM. Once you are logged in, your will be in the home directory for **root**, and you can proceed to examine your system.

The complete Bacula rescue part of the CD will be in the directory: **/bacula-hostname**, where hostname is replaced by the name of the host machine on which you did the build for the CDROM. This naming procedure allows you to put multiple restore environments for each of your machines on a single CDROM if you so wish to do. Please see the README document in the **rescue/linux/cdrom** directory for more information on adding to the CDROM.

Start the Network: At this point, you should bring up your network. Normally, this is quite simple and requires just a few commands. Please cd into the /bacula-hostname directory before continuing. To simplify your task, we have created a script that should work in most cases by typing:

```
cd /bacula-hostname
./start_network
```

You can test it by pinging another machine, or pinging your broken machine machine from another machine. Do not proceed until your network is up.

Partition Your Hard Disk(s): Assuming that your hard disk crashed and needs repartitioning, proceed with:

```
./partition.hda
```

If you have multiple disks, do the same for each of them. For SCSI disks, the repartition script will be named: **partition.sda**. If the script complains about the disk being in use, simply go back and redo the **df** command and **umount** commands until you no longer have your hard disk mounted. Note, in many cases, if your hard disk was seriously damaged or a new one installed, it will not automatically be mounted. If it is mounted, it is because the emergency kernel found one or more possibly valid partitions.

If for some reason this procedure does not work, you can use the information in **partition.hda** to re-partition your disks by hand using **fdisk**.

Format Your Hard Disk(s): If you have repartitioned your hard disk, you must format it appropriately. The formatting script will put back swap partitions, normal Unix partitions (ext2) and journaled partitions (ext3) as well as Reiser partitions (rei). Do so by entering for each disk:

```
./format.hda
```

The format script will ask you if you want a block check done. We recommend to answer yes, but realize that for very large disks this can take hours.

Mount the Newly Formatted Disks: Once the disks are partitioned and formatted, you can remount them with the **mount_drives** script. All your drives must be mounted for Bacula to be able to access them. Run the script as follows:

```
./mount_drives
df
```

The **df** command will tell you if the drives are mounted. If not, re-run the script again. It isn't always easy to figure out and create the mount points

and the mounts in the proper order, so repeating the **./mount_drives** command will not cause any harm and will most likely work the second time. If not, correct it by hand before continuing.

Restore and Start the File Daemon: If you have booted with a Bacula Rescue CDROM, your statically linked Bacula File daemon and the bacula-fd.conf file with be in the /bacula-hostname/bin directory. Make sure **bacula-fd** and **bacula-fd.conf** are both there.

Edit the Bacula configuration file, create the working/pid/subsys directory if you haven't already done so above, and start Bacula. Before starting Bacula, you will need to move it and bacula-fd.conf from /bacula-hostname/bin, to the /mnt/disk/tmp directory so that it will be on your hard disk. Then start it with the following command:

```
chroot /mnt/disk /tmp/bacula-fd -c /tmp/bacula-fd.conf
```

The above command starts the Bacula File daemon with your the proper root disk location (i.e. **/mnt/disk/tmp**. If Bacula does not start correct the problem and start it. You can check if it is running by entering:

```
ps fax
```

You can kill Bacula by entering:

```
kill -TERM <pid>
```

where **pid** is the first number printed in front of the first occurrence of **bacula-fd** in the **ps fax** command.

Now, you should be able to use another computer with Bacula installed to check the status by entering:

```
status client=xxxx
```

into the Console program, where xxxx is the name of the client you are restoring.

One common problem is that your **bacula-dir.conf** may contain machine addresses that are not properly resolved on the stripped down system to

be restored because it is not running DNS. This is particularly true for the address in the Storage resource of the Director, which may be very well resolved on the Director's machine, but not on the machine being restored and running the File daemon. In that case, be prepared to edit **bacula-dir.conf** to replace the name of the Storage daemon's domain name with its IP address.

Restore Your Files: On the computer that is running the Director, you now run a **restore** command and select the files to be restored (normally everything), but before starting the restore, there is one final change you must make using the **mod** option. You must change the **Where** directory to be the root by using the **mod** option just before running the job and selecting **Where**. Set it to:

```
/
```

then run the restore.

You might be tempted to avoid using **chroot** and running Bacula directly and then using a **Where** to specify a destination of **/mnt/disk**. This is possible, however, the current version of Bacula always restores files to the new location, and thus any soft links that have been specified with absolute paths will end up with **/mnt/disk** prefixed to them. In general this is not fatal to getting your system running, but be aware that you will have to fix these links if you do not use **chroot**.

Final Step: At this point, the restore should have finished with no errors, and all your files will be restored. One last task remains and that is to write a new boot sector so that your machine will boot. For **lilo**, you enter the following command:

```
./run_lilo
```

If you are using grub instead of lilo, you must enter the following:

```
./run_grub
```

Note, I've had quite a number of problems with **grub** because it is rather complicated and not designed to install easily under a simplified system. So,

if you experience errors or end up unexpectedly in a **chroot** shell, simply exit back to the normal shell and type in the appropriate commands from the **run_grub** script by hand until you get it to install. When you run the run_grub script, it will print the commands that you should manually enter if that is necessary.

Reboot: First unmount all your hard disks, otherwise they will not be cleanly shutdown, then reboot your machine by entering **exit** until you get to the main prompt then enter **ctl-d**. Once back to the main CDROM prompt, you will need to turn the power off then back on to your machine to get it to reboot.

If everything went well, you should now be back up and running. If not, re-insert the emergency boot CDROM, boot, and figure out what is wrong.

Restoring a Server

Above, we considered how to recover a client machine where a valid Bacula server was running on another machine. However, what happens if your server goes down and you no longer have a running Director, Catalog, or Storage daemon? There are several solutions:

1. Bring up static versions of your Director, Catalog, and Storage daemon.
2. Move your server to another machine.

The first option, is very difficult because it requires you to have created a static version of the Director and the Storage daemon as well as the Catalog. If the Catalog uses MySQL or PostgreSQL, this may or may not be possible. In addition, to loading all these programs on a bare system (quite possible), you will need to make sure you have a valid driver for your tape drive.

The second suggestion is probably a much simpler solution, and one I have done myself. To do so, you might want to consider the following steps:

- If you are using MySQL or PostgreSQL, configure, build and install it from source (or user rpms) on your new system.
- Load the Bacula source code onto your new system, configure, install it, and create the Bacula database.

- If you have a valid saved Bootstrap file as created for your damaged machine with WriteBootstrap, use it to restore the files to the damaged machine, where you have loaded a static Bacula File daemon using the Bacula Rescue disk). This is done by using the restore command and at the yes/mod/no prompt, selecting **mod** then specifying the path to the bootstrap file.

- If you have the Bootstrap file, you should now be back up and running, if you do not have a Bootstrap file, continue with the suggestions below.

- Using **bscan** scan the last set of backup tapes into your MySQL, PostgreSQL or SQLite database.

- Start Bacula, and using the Console **restore** command, restore the last valid copy of the Bacula database and the the Bacula configuration files.

- Move the database to the correct location.

- Start the database, and restart Bacula. Then use the Console **restore** command, restore all the files on the damaged machine, where you have loaded a Bacula File daemon using the Bacula Rescue disk.

Linux Problems or Bugs

Since every flavor and every release of Linux is different, there are likely to be some small difficulties with the scripts, so please be prepared to edit them in a minimal environment. A rudimentary knowledge of **vi** is very useful. Also, these scripts do not do everything. You will need to reformat Windows partitions by hand, for example.

Getting the boot loader back can be a problem if you are using **grub** because it is so complicated. If all else fails, reboot your system from your floppy but using the restored disk image, then proceed to a reinstallation of grub (looking at the run-grub script can help). By contrast, lilo is a piece of cake.

FreeBSD Bare Metal Recovery

The same basic techniques described above also apply to FreeBSD. Although we don't yet have a fully automated procedure, Alex Torres Molina has

provided us with the following instructions with a few additions from Jesse Guardiani and Dan Languille:

1. Boot with the FreeBSD installation disk
2. Go to Custom, Partition and create your slices and go to Label and create the particions that you want. Apply changes.
3. Go to Fixit to start a emergency console.
4. Create devs ad0 .. .. if don't exist under /mnt2/dev (in my situation) with MAKEDEV. The device or devices you create depend on what hard drives you have. ad0 is your first ATA drive. da0 would by your first SCSI drive. Under OS version 5 and greater, your device files are most likely automatically created for you.
5. mkdir /mnt/disk this is the root of the new disk
6. mount /mnt2/dev/ad0s1a /mnt/disk mount /mnt2/dev/ad0s1c /mnt/disk/var mount /mnt2/dev/ad0s1d /mnt/disk/usr ..... The same hard drive isssues as above apply here too. Note, under OS version 5 or higher, your disk devices may be in /dev not /mnt2/dev.
7. Network configuraion (ifconfig xl0 ip/mask + route add default ip-gateway)
8. mkdir /mnt/disk/tmp
9. cd /mnt/disk/tmp
10. Copy bacula-fd and bacula-fd.conf to this path
11. If you need to use sftp to copy files then you must do this: ln -s /mnt2/usr/bin /usr/bin
12. chmod u+x bacula-fd
13. Modify bacula-fd.conf to fit this machine
14. Copy /bin/sh to /mnt/disk, neccesary for chroot
15. Don't forget to put your bacula-dir's IP address and domain name in /mnt/disk/etc/hosts if it's not on a public net. Otherwise the FD on the machine you are restoring to won't be able to contact the SD and DIR on the remote machine.
16. mkdir -p /mnt/disk/var/db/bacula

17. chroot /mnt/disk /tmp/bacula-fd -c /tmp/bacula-fd.conf to start bacula-fd
18. Now you can go to bacula-dir and restore the job with the entire contents of the broken server.
19. You must create /proc

Solaris Bare Metal Recovery

The same basic techniques described above apply to Solaris:

- the same restrictions as those given for Linux apply
- you will need to create a Bacula Rescue disk

However, during the recovery phase, the boot and disk preparation procedures are different:

- there is no need to create an emergency boot disk since it is an integrated part of the Solaris boot.
- you must partition and format your hard disk by hand following manual procedures as described in W. Curtis Preston's book "Unix Backup & Recovery"

Once the disk is partitioned, formatted and mounted, you can continue with bringing up the network and reloading Bacula.

Preparing Solaris Before a Disaster

As mentioned above, before a disaster strikes, you should prepare the information needed in the case of problems. To do so, in the **rescue/solaris** subdirectory enter:

```
su
./getdiskinfo
./make_rescue_disk
```

The **getdiskinfo** script will, as in the case of Linux described above, create a subdirectory **diskinfo** containing the output from several system utilities. In addition, it will contain the output from the **SysAudit** program as described in Curtis Preston's book. This file **diskinfo/sysaudit.bsi** will contain the disk partitioning information that will allow you to manually follow the procedures in the "Unix Backup & Recovery" book to repartition and format your hard disk. In addition, the **getdiskinfo** script will create a **start_network** script.

Once you have your your disks repartitioned and formatted, do the following:

- Start Your Network with the **start_network** script
- Restore the Bacula File daemon as documented above
- Perform a Bacula restore of all your files using the same commands as described above for Linux
- Re-install your boot loader using the instructions outlined in the "Unix Backup & Recovery" book using installboot

Bugs and Other Considerations

Directory Modification and Access Times are Modified on pre-1.30 Baculas : When a pre-1.30 version of Bacula restores a directory, it first must create the directory, then it populates the directory with its files and subdirectories. The act of creating the files and subdirectories updates both the modification and access times associated with the directory itself. As a consequence, all modification and access times of all directories will be updated to the time of the restore.

This has been corrected in Bacula version 1.30 and later. The directory modification and access times is reset to the value saved in the backup after all the files and subdirectories have been restored. This has been tested and verified on normal restore operations, but not verified during a bare metal recovery.

Strange Bootstrap Files: If any of you look closely at the bootstrap file that is produced and used for the restore (I sure do), you will probably notice that the FileIndex item does not include all the files saved to the tape. This is because in some instances there are duplicates (especially in the case of an Incremental save), and in such circumstances, **Bacula** restores only the last of multiple copies of a file or directory.

Disaster Recovery of Win32 Systems

Due to open system files, and registry problems, Bacula cannot save and restore a complete Win2K/XP/NT environment.

A suggestion by Damian Coutts using Microsoft's NTBackup utility in conjunction with Bacula should permit a Full bare metal restore of Win2K/XP (and possibly NT systems). His suggestion is to do an NTBackup of the critical system state prior to running a Bacula backup with the following command:

```
ntbackup backup systemstate /F c:\systemstate.bkf
```

The **backup** is the command, the **systemstate** says to backup only the system state and not all the user files, and the **/F c:\systemstate.bkf** specifies where to write the state file. this file must then be saved and restored by Bacula.

To restore the system state, you first reload a base operating system, then you would use Bacula to restore all the users files and to recover the **c:\systemstate.bkf** file, and finally, run **NTBackup** and **catalogue** the system statefile, and then select it for restore. The documentation says you can't run a command line restore of the systemstate.

This procedure has been confirmed to work by Ludovic Strappazon – many thanks!

A new tool is provided in the form of a bacula plugin for the BartPE rescue CD. BartPE is a self-contained WindowsXP boot CD which you can make using the PeBuilder tools available at http://www.nu2.nu/pebuilder/ and a valid Windows XP SP1 CDROM. The plugin is provided as a zip archive. Unzip the file and copy the bacula directory into the plugin directory of your BartPE installation. Edit the configuration files to suit your installation and build your CD according to the instructions at Bart's site. This will permit you to boot from the cd, configure and start networking, start the bacula file client and access your director with the console program. The programs menu on the booted CD contains entries to install the file client service, start the file client service, and start the WX-Console. You can also open a command line window and CD Programs\Bacula and run the command line console bconsole.

Resetting Directory and File Ownership and Permissions on Win32 Systems

Bacula versions after 1.31 should properly restore ownership and permissions on all WinNT/XP/2K systems. If you do experience problems, generally in restores to alternate directories because higher level directories were not backed up by Bacula, you can correct any problems with the **SetACL** available under the GPL license at: http://sourceforge.net/projects/setacl/.

Alternate Disaster Recovery Suggestion for Win32 Systems

Ludovic Strappazon has suggested an interesting way to backup and restore complete Win32 partitions. Simply boot your Win32 system with a Linux Rescue disk as described above for Linux, install a statically linked Bacula, and backup any of the raw partitions you want. Then to restore the system, you simply restore the raw partition or partitions. Here is the email that Ludovic recently sent on that subject:

```
I've just finished testing my brand new cd LFS/Bacula
with a raw Bacula backup and restore of my portable.
I can't resist sending you the results: look at the rates !!!
hunt-dir: Start Backup JobId 100, Job=HuntBackup.2003-04-17_12.58.26
hunt-dir: Bacula 1.30 (14Apr03): 17-Apr-2003 13:14
JobId:                  100
Job:                    HuntBackup.2003-04-17_12.58.26
FileSet:                RawPartition
Backup Level:           Full
Client:                 sauvegarde-fd
Start time:             17-Apr-2003 12:58
End time:               17-Apr-2003 13:14
Files Written:          1
Bytes Written:          10,058,586,272
Rate:                   10734.9 KB/s
Software Compression:   None
Volume names(s):        000103
Volume Session Id:      2
Volume Session Time:    1050576790
Last Volume Bytes:      10,080,883,520
FD termination status:  OK
SD termination status:  OK
Termination:            Backup OK
hunt-dir: Begin pruning Jobs.
hunt-dir: No Jobs found to prune.
hunt-dir: Begin pruning Files.
hunt-dir: No Files found to prune.
hunt-dir: End auto prune.
hunt-dir: Start Restore Job RestoreFilesHunt.2003-04-17_13.21.44
hunt-sd: Forward spacing to file 1.
```

```
hunt-dir: Bacula 1.30 (14Apr03): 17-Apr-2003 13:54
JobId:                  101
Job:                    RestoreFilesHunt.2003-04-17_13.21.44
Client:                 sauvegarde-fd
Start time:             17-Apr-2003 13:21
End time:               17-Apr-2003 13:54
Files Restored:         1
Bytes Restored:         10,056,130,560
Rate:                   5073.7 KB/s
FD termination status:  OK
Termination:            Restore OK
hunt-dir: Begin pruning Jobs.
hunt-dir: No Jobs found to prune.
hunt-dir: Begin pruning Files.
hunt-dir: No Files found to prune.
hunt-dir: End auto prune.
```

Restoring to a Running System

If for some reason you want to do a Full restore to a system that has a working kernel, you will need to take care not to overwrite the following files:

```
/etc/grub.conf
/etc/X11/Conf
/etc/fstab
/etc/mtab
/lib/modules
/usr/modules
/usr/X11R6
/etc/modules.conf
```

Additional Resources

Many thanks to Charles Curley who wrote Linux Complete Backup and Recovery HOWTO for the The Linux Documentation Project. This is an excellent document on how to do Bare Metal Recovery on Linux systems, and it was this document that made me realize that Bacula could do the same thing.

You can find quite a few additional resources, both commercial and free at Storage Mountain, formerly known as Backup Central.

And finally, the O'Reilly book, "Unix Backup & Recovery" by W. Curtis

Preston covers virtually every backup and recovery topic including bare metal recovery for a large range of Unix systems.

Using Bacula to Encrypt Communications to Clients

At the current time, Bacula does not have built-in communications encryption. However, without too much effort, it is possible to encrypt the communications between any of the daemons. This chapter will show you how to use **stunnel** to encrypt communications to your client programs. We assume the Director and the Storage daemon are running on one machine that will be called **server** and the Client or File daemon is running on a different machine called **client**. Although the details may be slightly different, the same principles apply whether you are encrypting between Unix, Linux, or Win32 machines. This example was developed between two Linux machines running stunnel version 4.04-4 on a Red Hat Enterprise 3.0 system.

Communications Ports Used

First, you must know that with the standard Bacula configuration, the Director will contact the File daemon on port 9102. The File daemon then contacts the Storage daemon using the address and port parameters supplied by the Director. The standard port used will be 9103. This in the typical server/client view of the world, the File daemon is a server to the Director (i.e. listens for the Director to contact it), and the Storage daemon is a server to the File daemon.

Encryption

The encryption is accomplished between the Director and the File daemon by using an stunnel on the Director's machine (server) to encrypt the data and to contact a stunnel on the File daemon's machine (client), which decrypts the data and passes it to the client.

Between the File daemon and the Storage daemon, we use an stunnel on the File daemon's machine to encrypt the data and another stunnel on the Storage daemon's machine to decrypt the data.

As a consequence, there are actually four copies of stunnel running, two on server and two on client. This may sound a bit complicated, but it really isn't. To accomplish this, we will need to construct four separate conf files for stunnel, and we will need to make some minor modifications to the Director's conf file. None of the other conf files need to be changed.

A Picture

Since pictures usually help a lot, here is an overview of what we will be doing. Don't worry about all the details of the port numbers and such for the moment.

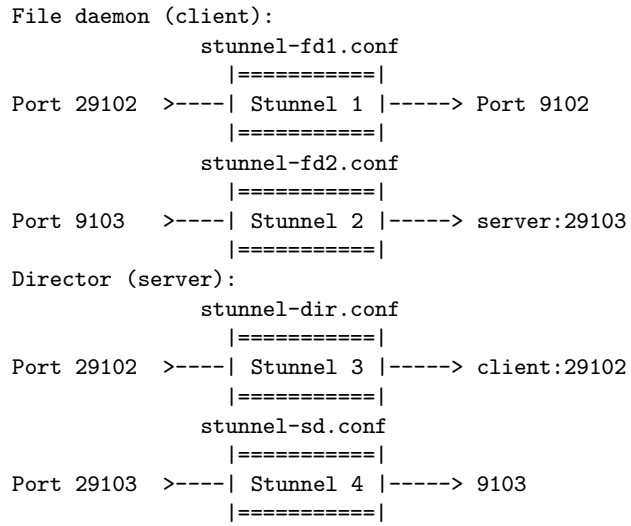

Certificates

In order for stunnel to function as a server, which it does in our diagram for Stunnel 1 and Stunnel 4, you must have a certificate and the key. It is possible to keep the two in separate files, but normally, you keep them in one single .pem file. You may create this certificate yourself in which case, it will be self-signed, or you may have it signed by a CA.

If you want your clients to verify that the server is in fact valid (Stunnel 2 and Stunnel 3), you will need to have the server certificates signed by a CA (Certificate Authority), and you will need to have the CA's public certificate (contains the CA's public key).

Having a CA signed certificate is **highly** recommended if you are using your client across the Internet, otherwise you are exposed to the man in the middle attack and hence loss of your data.

See below for how to create a self-signed certificate.

Securing the Data Channel

To simplify things a bit, let's for the moment consider only the data channel. That is the connection between the File daemon and the Storage daemon, which takes place on port 9103. In fact, in a minimalist solution, this is the only connection needs to be encrypted, because it is the one that transports your data. The connection between the Director and the File daemon is simply a control channel used to start the job and get the job status.

Normally the File daemon will contact the Storage daemon on port 9103 (supplied by the Director), so we need a stunnel that listens on port 9103 on the File daemon's machine, encrypts the data and sends it to the Storage daemon. This is depicted by Stunnel 2 above. Note that this stunnel is listening on port 9103 and sending to server:29103. We use port 29103 on the server because if we sent the data to port 9103, it would go directly to the Storage daemon, which doesn't understand encrypted data. On the server machine, we run Stunnel 4, which listens on port 29103, decrypts the data and sends it to the Storage daemon, which is listening on port 9103.

Modification of bacula-dir.conf for the Data Channel

The Storage resource of the bacula-dir.conf normally looks something like the following:

```
Storage {
  Name = File
  Address = server
  SDPort = 9103
  Password = storage_password
  Device = File
  Media Type = File
}
```

Notice that this is running on the server machine, and it points the File daemon back to server:9103, which is where our Storage daemon is listening. We modify this to be:

```
Storage {
  Name = File
  Address = localhost
  SDPort = 9103
  Password = storage_password
  Device = File
  Media Type = File
}
```

This causes the File daemon to send the data to the stunnel running on localhost (the client machine). We could have used client as the address as well.

config Files for stunnel to Encrypt the Data Channel

In the diagram above, we see above Stunnel 2 that we stunnel-fd2.conf on client. A pretty much minimal config file would look like the following:

```
client = yes
[29103]
accept = localhost:9103
connect = server:29103
```

The above config file does encrypt the data but it does not require a certificate, so it is subject to the man in the middle attack. The file I actually used, stunnel-fd2.conf, looked like this:

```
#
# Stunnel conf for Bacula client -> SD
#
pid = /home/kern/bacula/bin/working/stunnel.pid
#
# A cert is not mandatory here. If verify=2, a
#  cert signed by a CA must be specified, and
#  either CAfile or CApath must point to the CA's
#  cert
#
cert = /home/kern/stunnel/stunnel.pem
CAfile = /home/kern/ssl/cacert.pem
verify = 2
client = yes
# debug = 7
# foreground = yes
[29103]
accept = localhost:9103
connect = server:29103
```

You will notice that I specified a pid file location because I ran stunnel under my own userid so I could not use the default, which requires root permission. I also specified a certificate that I have as well as verify level 2 so that the certificate is required and verified, and I must supply the location of the CA (Certificate Authority) certificate so that the stunnel certificate can be verified. Finally, you will see that there are two lines commented out, which when enabled, produce a lot of nice debug info in the command window.

If you do not have a signed certificate (stunnel.pem), you need to delete the cert, CAfile, and verify lines.

Note that the stunnel.pem, is actually a private key and a certificate in a single file. These two can be kept and specified individually, but keeping them in one file is more convenient.

The config file, stunnel-sd.conf, needed for Stunnel 4 on the server machine is:

```
#
# Bacula stunnel conf for Storage daemon
#
pid = /home/kern/bacula/bin/working/stunnel.pid
#
# A cert is mandatory here, it may be self signed
#  If it is self signed, the client may not use
#  verify
#
cert   = /home/kern/stunnel/stunnel.pem
client = no
# debug = 7
# foreground = yes
[29103]
accept = 29103
connect = 9103
```

Starting and Testing the Data Encryption

It will most likely be the simplest to implement the Data Channel encryption in the following order:

- Setup and run Bacula backing up some data on your client machine without encryption.
- Stop Bacula.
- Modify the Storage resource in the Director's conf file.
- Start Bacula
- Start stunnel on the server with:

  ```
  stunnel stunnel-sd.conf
  ```

- Start stunnel on the client with:

```
stunnel stunnel-fd2.conf
```

- Run a job.
- If it doesn't work, turn debug on in both stunnel conf files, restart the stunnels, rerun the job, repeat until it works.

Encrypting the Control Channel

The Job control channel is between the Director and the File daemon, and as mentioned above, it is not really necessary to encrypt, but it is good practice to encrypt it as well. The two stunnels that are used in this case will be Stunnel 1 and Stunnel 3 in the diagram above. Stunnel 3 on the server might normally listen on port 9102, but if you have a local File daemon, this will not work, so we make it listen on port 29102. It then sends the data to client:29102. Again we use port 29102 so that the stunnel on the client machine can decrypt the data before passing it on to port 9102 where the File daemon is listening.

Modification of bacula-dir.conf for the Control Channel

We need to modify the standard Client resource, which would normally look something like:

```
Client {
  Name = client-fd
  Address = client
  FDPort = 9102
  Catalog = BackupDB
  Password = "xxx"
}
```

to be:

```
Client {
  Name = client-fd
  Address = localhost
  FDPort = 29102
  Catalog = BackupDB
  Password = "xxx"
}
```

This will cause the Director to send the control information to localhost:29102 instead of directly to the client.

config Files for stunnel to Encrypt the Control Channel

The stunnel config file, stunnel-dir.conf, for the Director's machine would look like the following:

```
#
# Bacula stunnel conf for the Directory to contact a client
#
pid = /home/kern/bacula/bin/working/stunnel.pid
#
# A cert is not mandatory here. If verify=2, a
#  cert signed by a CA must be specified, and
#  either CAfile or CApath must point to the CA's
#  cert
#
cert   = /home/kern/stunnel/stunnel.pem
CAfile = /home/kern/ssl/cacert.pem
verify = 2
client = yes
# debug = 7
# foreground = yes
[29102]
accept = localhost:29102
connect = client:29102
```

and the config file, stunnel-fd1.conf, needed to run stunnel on the Client would be:

```
#
# Bacula stunnel conf for the Directory to contact a client
#
pid = /home/kern/bacula/bin/working/stunnel.pid
#
# A cert is not mandatory here. If verify=2, a
#  cert signed by a CA must be specified, and
#  either CAfile or CApath must point to the CA's
#  cert
#
cert   = /home/kern/stunnel/stunnel.pem
CAfile = /home/kern/ssl/cacert.pem
verify = 2
client = yes
# debug = 7
# foreground = yes
[29102]
```

```
accept = localhost:29102
connect = client:29102
```

Starting and Testing the Control Channel

It will most likely be the simplest to implement the Control Channel encryption in the following order:

- Stop Bacula.
- Modify the Client resource in the Director's conf file.
- Start Bacula
- Start stunnel on the server with:

```
stunnel stunnel-dir.conf
```

- Start stunnel on the client with:

```
stunnel stunnel-fd1.conf
```

- Run a job.
- If it doesn't work, turn debug on in both stunnel conf files, restart the stunnels, rerun the job, repeat until it works.

Using stunnel to Encrypt to a Second Client

On the client machine, you can just duplicate the setup that you have on the first client file for file and it should work fine.

In the bacula-dir.conf file, you will want to create a second client pretty much identical to how you did for the first one, but the port number must be unique. We previously used:

```
Client {
  Name = client-fd
  Address = localhost
  FDPort = 29102
  Catalog = BackupDB
  Password = "xxx"
}
```

so for the second client, we will, of course, have a different name, and we will also need a different port. Remember that we used port 29103 for the Storage daemon, so for the second client, we can use port 29104, and the Client resource would look like:

```
Client {
  Name = client2-fd
  Address = localhost
  FDPort = 29104
  Catalog = BackupDB
  Password = "yyy"
}
```

Now, fortunately, we do not need a third stunnel to on the Director's machine, we can just add the new port to the config file, stunnel-dir.conf, to make:

```
#
# Bacula stunnel conf for the Directory to contact a client
#
pid = /home/kern/bacula/bin/working/stunnel.pid
#
# A cert is not mandatory here. If verify=2, a
#  cert signed by a CA must be specified, and
#  either CAfile or CApath must point to the CA's
#  cert
#
cert   = /home/kern/stunnel/stunnel.pem
CAfile = /home/kern/ssl/cacert.pem
verify = 2
client = yes
# debug = 7
# foreground = yes
[29102]
accept = localhost:29102
connect = client:29102
[29104]
accept = localhost:29102
connect = client2:29102
```

There are no changes necessary to the Storage daemon or the other stunnel so that this new client can talk to our Storage daemon.

Creating a Self-signed Certificate

You may create a self-signed certificate for use with stunnel that will permit you to make it function, but will now allow certificate validation. The

.pem file containing both the certificate and the key can be made with the following, which I put in a file named **makepem**:

```
#!/bin/sh
#
# Simple shell script to make a .pem file that can be used
# with stunnel and Bacula
#
OPENSSL=openssl
   umask 77
   PEM1=`/bin/mktemp openssl.XXXXXX`
   PEM2=`/bin/mktemp openssl.XXXXXX`
   ${OPENSSL} req -newkey rsa:1024 -keyout $PEM1 -nodes \
       -x509 -days 365 -out $PEM2
   cat $PEM1 > stunnel.pem
   echo ""   >>stunnel.pem
   cat $PEM2 >>stunnel.pem
   rm $PEM1 $PEM2
```

The above script will ask you a number of questions. You may simply answer each of them by entering a return, or if you wish you may enter your own data.

Getting a CA Signed Certificate

The process of getting a certificate that is signed by a CA is quite a bit more complicated. You can purchase one from quite a number of PKI vendors, but that is not at all necessary for use with Bacula. To get a CA signed certificate, you will either need to find a friend that has setup his own CA or to become a CA yourself, and thus you can sign all your own certificates. The book OpenSSL by John Viega, Matt Mesier & Pravir Chandra from O'Reilly explains how to do it, or you can read the documentation provided in the Open-source PKI Book project at Source Forge: http://ospkibook.sourceforge.net/docs/OSPKI-2.4.7/OSPKI-html/ospki-book.htm. Note, this link may change.

Using ssh to Secure the Communications

Please see the script **ssh-tunnel.sh** in the **examples** directory. It was contributed by Stephan Holl.

Bacula Security Issues

- Security means being able to restore your files, so read the Critical Items Chapter of this manual.
- The Clients (**bacula-fd**) must run as root to be able to access all the system files.
- It is not necessary to run the Director as root.
- It is not necessary to run the Storage daemon as root, but you must ensure that it can open the tape drives, which are often restricted to root access by default. In addition, if you do not run the Storage daemon as root, it will not be able to automatically set your tape drive parameters on most OSes since these functions, unfortunately require root access.
- You should restrict access to the Bacula configuration files, so that the passwords are not world-readable. The **Bacula** daemons are password protected using CRAM-MD5 (i.e. the password is not sent across the network). This will ensure that not everyone can access the daemons. It is a reasonably good protection, but can be cracked by experts.
- If you are using the recommended ports 9101, 9102, and 9103, you will probably want to protect these ports from external access using a firewall and/or using tcp wrappers (**etc/hosts.allow**).
- Currently all data that is sent across the network is unencrypted. As a consequence, unless you use **ssh** or **stunnel** for port forwarding, it is not recommended to do a backup across an insecure network (e.g. the Internet). In a future version, we plan to have **ssl** encryption built-in.
- You should ensure that the Bacula working directories are readable and writable only by the Bacula daemons.
- If you are using **MySQL** it is not necessary for it to run with **root** permission.
- The default Bacula **grant-mysql-permissions** script grants all permissions to use the MySQL database without a password. If you want security, please tighten this up!
- Don't forget that Bacula is a network program, so anyone anywhere on the network with the console program and the Director's password can access Bacula and the backed up data.

- You can restrict what IP addresses Bacula will bind to by using the appropriate **DirAddress**, **FDAddress**, or **SDAddress** records in the respective daemon configuration files.

Configuring and Testing TCP Wrappers with Bacula

TCP Wrappers are implemented if you turn them on when configuring (**./configure --with-libwrap**). With this code enabled, you may control who may access your daemons. This control is done by modifying the file: **/etc/hosts.allow**. The program name that **Bacula** uses when applying these access restrictions is the name you specify in the daemon configuration file. You must not use the **twist** option in your **/etc/hosts.allow** or it will terminate the Bacula daemon when a connection is refused.

Dan Languille has provided the following information on configuring and testing TCP wrappers with Bacula.

If you read hosts_options(5), you will see an option called twist. This option replaces the current process by an instance of the specified shell command. Typically, something like this is used:

```
ALL : ALL \
 : severity auth.info \
 : twist /bin/echo "You are not welcome to use %d from %h."
```

The libwrap code tries to avoid **twist** if it runs in a resident process, but that test will not protect the first hosts_access() call. This will result in the prcess (e.g. bacula-fd, bacula-sd, bacula-dir) being terminated if the first connection to their port results in the twist option being invoked. The potential, and I strees potential, exists for an attacker to prevent the daemons from running. This situation is eliminated if your /etc/hosts.allow file contains an appropriate ruleset. The following example is sufficent:

```
undef-fd : localhost : allow
undef-sd : localhost : allow
undef-dir : localhost : allow
undef-fd : ALL : deny
undef-sd : ALL : deny
undef-dir : ALL : deny
```

You must adjust the daemon names to those found in the respective daemon configuration files. In these examples, the Director is undef- dir, the Storage

Daemon is undef-sd, and the File Daemon is undef-fd. Adjust to suit your situation. The above example rules assume that the SD, FD, and DIR all reside on the same box. If you have a remote FD client, then the following ruleset on the remote client will suffice:

```
undef-fd : director.example.org : allow
undef-fd : ALL : deny
```

where director.example.org is the host which will be contacting the client (ie. the box on which the Bacula Director daemon runs). The use of "ALL : deny" ensures that the twist option (if present) is not invoked. To properly test your configuration, start the daemon(s), then attempt to connect from an IP address which should be able to connect. You should see something like this:

```
$ telnet undef 9103
Trying 192.168.0.56...
Connected to undef.example.org.
Escape character is '^]'.
Connection closed by foreign host.
$
```

This is the correct response. If you see this:

```
$ telnet undef 9103
Trying 192.168.0.56...
Connected to undef.example.org.
Escape character is '^]'.
You are not welcome to use undef-sd from xeon.example.org.
Connection closed by foreign host.
$
```

then twist has been invoked and your configuration is not correct and you need to add the deny statement. It is important to note that your testing must include restarting the daemons after each connection attempt. You can also tcpdchk(8) and tcpdmatch(8) to validate your /etc/hosts.allow rules. Here is a simple test using tcpdmatch:

```
$ tcpdmatch undef-dir xeon.example.org
warning: undef-dir: no such process name in /etc/inetd.conf
client: hostname xeon.example.org
client: address 192.168.0.18
server: process undef-dir
matched: /etc/hosts.allow line 40
option: allow
access: granted
```

If you are running Bacula as a standalone daemon, the warning above can be safely ignored. Here is an example which indicates that your rules are missing a deny statement and the twist option has been invoked.

```
$ tcpdmatch undef-dir 10.0.0.1
warning: undef-dir: no such process name in /etc/inetd.conf
client: address 10.0.0.1
server: process undef-dir
matched: /etc/hosts.allow line 91
option: severity auth.info
option: twist /bin/echo "You are not welcome to use
  undef-dir from 10.0.0.1."
access: delegated
```

Running as non-root

Security advice from Dan Languille:

It is a good idea to run daemons with the lowest possible privileges. In other words, if you can, don't run applications as root which do not have to be root. The Storage Daemon and the Director Daemon do not need to be root. The File Daemon needs to be root in order to access all files on your system. In order to run as non-root, you need to create a user and a group. Choosing `bacula` as both the user name and the group name sounds like a good idea to me.

The FreeBSD port creates this user and group for you (actually, as I write this, the port doesn't do that, but it soon will). Here is what those entries looked like on my FreeBSD laptop:

```
bacula:*:1002:1002::0:0:Bacul Daemon:/var/db/bacula:/sbin/nologin
```

I used vipw to create this entry. I selected a User ID and Group ID of 1002 as they were unused on my system.

I also created a group in /etc/group:

```
bacula:*:1002:
```

The bacula user (as opposed to the Bacula daemon) will have a home directory of `/var/db/bacula` which is the default location for the Bacula database.

Now that you have both a bacula user and a bacula group, you can secure the bacula home directory by issuing this command:

```
chown -R bacula:bacula /var/db/bacula/
```

This ensures that only the bacula user can access this directory. It also means that if we run the Director and the Storage daemon as bacula, those daemons also have restricted access. This would not be the case if they were running as root.

It is important to note that the storage daemon actually needs to be in the operator group for normal access to tape drives etc (at least on a FreeBSD system, that's how things are set up by default) Such devices are normally chown root:operator. It is easier and less error prone to make Bacula a member of that group than it is to play around with system permissions.

Starting the Bacula daemons

To start the bacula daemons on a FreeBSD system, issue the following command:

```
/usr/local/etc/rc.d/bacula.sh start
```

To confirm they are all running:

```
$ ps auwx | grep bacula
root\ 63416\ 0.0\ 0.3\ 2040 1172\ ??\ Ss\ 4:09PM 0:00.01
    /usr/local/sbin/bacula-sd -v -c /usr/local/etc/bacula-sd.conf
root\ 63418\ 0.0\ 0.3\ 1856 1036\ ??\ Ss\ 4:09PM 0:00.00
    /usr/local/sbin/bacula-fd -v -c /usr/local/etc/bacula-fd.conf
root\ 63422\ 0.0\ 0.4\ 2360 1440\ ??\ Ss\ 4:09PM 0:00.00
    /usr/local/sbin/bacula-dir -v -c /usr/local/etc/bacula-dir.conf
```

Dealing with Firewalls

If you have a firewall or a DMZ installed on your computer, you may experience difficulties contacting one or more of the Clients to back them up. This is especially true if you are trying to backup a Client across the Internet.

Technical Details

If you are attempting to do this, the sequence of network events in Bacula to do a backup are the following:

```
Console -> DIR:9101
DIR     -> SD:9103
DIR     -> FD:9102
FD      -> SD:9103
```

Where it should be obvious that DIR represents the Director, FD the File daemon or client, and SD the Storage daemon. The numbers that follow those names are the standard ports used by Bacula, and the -> represents the left side making a connection to the right side (i.e. the right side is the "server" or is listening on the specified port), and the left side is the "client" who initiates the conversation.

Note, port 9103 serves both the Director and the File daemon, each having its own independent connection.

If you are running **iptables**, you might add something like:

```
-A FW-1-INPUT -m state --state NEW -m tcp -p tcp --dport 9101:9103 -j ACCEPT
```

on your server, and

```
-A FW-1-INPUT -m state --state NEW -m tcp -p tcp --dport 9102 -j ACCEPT
```

on your client. In both cases, I assume that the machine is allowed to initiate connections on any port. If not, you will need to allow outgoing connections on ports 9102 and 9103 on your server and 9103 on your client. Thanks to Raymond Norton for this tip.

A Concrete Example

Jesse Guardiani's solution for his network for this problem, in his own words, is:

My bacula server is on the 192.168.1.0/24 network at IP address 192.168.1.52. For the sake of discussion we will refer to this network as the 'internal' network because it connects to the internet through a NAT'd firewall. We will call the network on the public (internet) side of the NAT'd firewall the 'external' network. Also, for the sake of discussion we will call my bacula server:

```
server.int.mydomain.tld
```

when a fully qualified domain name is required, or simply:

```
server
```

if a hostname is adequate. We will call the various bacula daemons running on the server.int.mydomain.tld machine:

```
server-fd
server-sd
server-dir
```

In addition, I have two clients that I want to back up with Bacula. The first client is on the internal network. Its fully qualified domain name is:

```
private1.int.mydomain.tld
```

And its hostname is:

```
private1
```

This machine is a client and therefore runs just one bacula daemon:

```
private1-fd
```

The second client is on the external network. Its fully qualified domain name is:

```
public1.mydomain.tld
```

And its hostname is:

```
public1
```

This machine also runs just one bacula daemon:

```
public1-fd
```

Finally, I have a NAT firewall/gateway with two network interfaces. The first interface is on the internal network and serves as a gateway to the internet for all the machines attached to the internal network (For example, server.int.mydomain.tld and private1.int.mydomain.tld). The second interface is on the external (internet) network. The external interface has been assigned the name:

```
firewall.mydomain.tld
```

Remember:

```
*.int.mydomain.tld = internal network
    *.mydomain.tld = external network
```

The Bacula Configuration Files for the Above

server-sd manages a 4 tape AIT autoloader. All of my backups are written to server-sd. I have just *one* Device resource in my server-sd.conf file:

```
Device {
  Name = "autochanger1";
  Media Type = AIT-1;
  Archive Device = /dev/nrsa1;
  Changer Device = /dev/ch0;
  Changer Command = "/usr/local/sbin/chio-bacula %c %o %S %a";
  Label Media = yes;
  AutoChanger = yes;
  AutomaticMount = yes;               # when device opened, read it
  AlwaysOpen = yes;
    Hardware End of Medium = No
    Fast Forward Space File = No
    BSF at EOM = yes
}
```

(note, please see the Tape Testing chapter of this manual for important FreeBSD information.) However, I have *two* Storage resources in my server-dir.conf file:

```
Storage {
  Name = "autochanger1-int"    # Storage device for backing up
  Address = server.int.mydomain.tld
  SDPort = 9103
  Password = "mysecretpassword"
  Device = "autochanger1"
  Media Type = AIT-1
  Autochanger = yes
}
Storage {
  Name = "autochanger1-ext"    # Storage device for backing up
  Address = firewall.mydomain.tld
  SDPort = 9103
  Password = "mysecretpassword"
  Device = "autochanger1"
  Media Type = AIT-1
  Autochanger = yes
}
```

Note that BOTH of the above server-dir.conf Storage resources use the same 'autochanger1' Device resource from server-sd.conf.

My backup jobs run consecutively, one after the other, so only one of the above Storage resources is being used by Bacula file daemons at any given time. I don't know if this would cause problems at a site that runs more than one backup in parallel to a single tape device.

In addition to the above, I have two Client resources defined in server-dir.conf:

```
Client {
  Name = private1-fd
  Address = private1.int.mydomain.tld
  FDPort = 9102
  Catalog = MyCatalog
  Password = "mysecretpassword"       # password for FileDaemon
}
Client {
  Name = public1-fd
  Address = public1.mydomain.tld
  FDPort = 9102
  Catalog = MyCatalog
  Password = "mysecretpassword"       # password for FileDaemon
}
```

And finally, to tie it all together, I have two Job resources defined in server-dir.conf:

```
Job {
  Name = "Private1-Backup"
  Type = Backup
  Client = private1-fd
  FileSet = "Private1"
  Schedule = "WeeklyCycle"
  Storage = "autochanger1-int"
  Messages = Standard
  Pool = "Weekly"
  Write Bootstrap = "/var/db/bacula/Private1-Backup.bsr"
  Priority = 12
}
Job {
  Name = "Public1-Backup"
  Type = Backup
  Client = public1-fd
  FileSet = "Public1"
  Schedule = "WeeklyCycle"
  Storage = "autochanger1-ext"
  Messages = Standard
  Pool = "Weekly"
  Write Bootstrap = "/var/db/bacula/Public1-Backup.bsr"
  Priority = 13
}
```

It is important to notice that because the 'Private1-Backup' Job is intended to back up a machine on the internal network it uses the 'autochanger1-int' Storage resource. On the other hand, the 'Public1-Backup' Job is intended to back up a machine on the external network, so it uses the 'autochanger1-ext' Storage resource.

I have left the Pool, Catalog, Messages, FileSet, Schedule, and Director resources out of the above server-dir.conf examples because they are not pertinent to the discussion.

How Does It Work?

If I want to run a backup of private1.int.mydomain.tld and store that backup using server-sd then my understanding of the order of events is this:

1. I execute my Bacula 'console' command on server.int.mydomain.tld.
2. console connects to server-dir.

3. I tell console to 'run' backup Job 'Private1-Backup'.
4. console relays this command to server-dir.
5. server-dir connects to private1-fd at private1.int.mydomain.tld:9102
6. server-dir tells private1-fd to start sending the files defined in the 'Private1-Backup' Job's FileSet resource to the Storage resource 'autochanger1-int', which we have defined in server-dir.conf as having the address:port of server.int.mydomain.tld:9103.
7. private1-fd connects to server.int.mydomain.tld:9103 and begins sending files.

Alternatively, if I want to run a backup of public1.mydomain.tld and store that backup using server-sd then my understanding of the order of events is this:

1. I execute my Bacula 'console' command on server.int.mydomain.tld.
2. console connects to server-dir.
3. I tell console to 'run' backup Job 'Public1-Backup'.
4. console relays this command to server-dir.
5. server-dir connects, through the NAT'd firewall, to public1-fd at public1.mydomain.tld:9102
6. server-dir tells public1-fd to start sending the files defined in the 'Public1-Backup' Job's FileSet resource to the Storage resource 'autochanger1-ext', which we have defined in server-dir.conf as having the address:port of firewall.mydomain.tld:9103.
7. public1-fd connects to firewall.mydomain.tld:9103 and begins sending files.

Important Note

In order for the above 'Public1-Backup' Job to succeed, firewall.mydomain.tld:9103 MUST be forwarded using the firewall's configuration software to server.int.mydomain.tld:9103. Some firewalls call this 'Server Publication'. Others may call it 'Port Forwarding'.

Using Bacula to Improve Computer Security

Since Bacula maintains a catalog of files, their attributes, and either SHA1 or MD5 signatures, it can be an ideal tool for improving computer security. This is done by making a snapshot of your system files with a **Verify** Job and then checking the current state of your system against the snapshot, on a regular basis (e.g. nightly).

The first step is to set up a **Verify** Job and to run it with:

```
Level = InitCatalog
```

The **InitCatalog** level tells **Bacula** simply get the information on the specified files and to put it into the catalog. That is your database is initialized and no comparison is done. The **InitCatalog** is normally run one time manually.

Thereafter, you will run a Verify Job on a daily (or whatever) basis with:

```
Level = Catalog
```

The **Level = Catalog** level tells Bacula to compare the current state of the files on the Client to the last **InitCatalog** that is stored in the catalog and to report any differences. See the example below for the format of the output.

You decide what files you want to form your "snapshot" by specifying them in a **FileSet** resource, and normally, they will be system files that do not change, or that only certain features change.

Then you decide what attributes of each file you want compared by specifying comparison options on the **Include** statements that you use in the **FileSet** resource of your **Catalog** Jobs.

The Details

In the discussion that follows, we will make reference to the Verify Configuration Example that is included below in the **A Verify Configuration Example** section. You might want to look it over now to get an idea of what it does.

The main elements consist of adding a schedule, which will normally be run daily, or perhaps more often. This is provided by the **VerifyCycle** Schedule, which runs at 5:05 in the morning every day.

Then you must define a Job, much as is done below. We recommend that the Job name contain the name of your machine as well as the word **Verify** or **Check**. In our example, we named it **MatouVerify**. This will permit you to easily identify your job when running it from the Console.

You will notice that most records of the Job are quite standard, but that the **FileSet** resource contains **verify=pins1** option in addition to the standard **signature=SHA1** option. If you don't want SHA1 signature comparison, and we cannot imagine why not, you can drop the **signature=SHA1** and none will be computed nor stored in the catalog. Or alternatively, you can use **verify=pins5** and **signature=MD5**, which will use the MD5 hash algorithm. The MD5 hash computes faster than SHA1, but is cryptographically less secure.

The **verify=pins1** is ignored during the **InitCatalog** Job, but is used during the subsequent **Catalog** Jobs to specify what attributes of the files should be compared to those found in the catalog. **pins1** is a reasonable set to begin with, but you may want to look at the details these and other options. They can be found in the FileSet Resource section of this manual. Briefly, however, the **p** of the **pins1** tells Verify to compare the permissions bits, the **i** is to compare inodes, the **n** causes comparison of the number of links, the **s** compares the file size, and the **1** compares the SHA1 checksums (this requires the **signature=SHA1** option to have been set also).

You must also specify the **Client** and the **Catalog** resources for your Verify job, but you probably already have them created for your client and do not need to recreate them, they are included in the example below for completeness.

As mentioned above, you will need to have a **FileSet** resource for the Verify job, which will have the additional **verify=pins1** option. You will want to take some care in defining the list of files to be included in your **FileSet**. Basically, you will want to include all system (or other) files that should not change on your system. If you select files, such as log files or mail files, which are constantly changing, your automatic Verify job will be constantly finding differences. The objective in forming the FileSet is to choose all unchanging important system files. Then if any of those files has changed, you will be notified, and you can determine if it changed because you loaded a new package, or because someone has broken into your computer and modified your files. The example below shows a list of files that I use on my RedHat 7.3 system. Since I didn't spend a lot of time working on it, it probably

is missing a few important files (if you find one, please send it to me). On the other hand, as long as I don't load any new packages, none of these files change during normal operation of the system.

Running the Verify

The first thing you will want to do is to run an **InitCatalog** level Verify Job. This will initialize the catalog to contain the file information that will later be used as a basis for comparisons with the actual file system, thus allowing you to detect any changes (and possible intrusions into your system).

The easiest way to run the **InitCatalog** is manually with the console program by simply entering **run**. You will be presented with a list of Jobs that can be run, and you will choose the one that corresponds to your Verify Job, **MatouVerify** in this example.

```
The defined Job resources are:
     1: MatouVerify
     2: kernsrestore
     3: Filetest
     4: kernsave
Select Job resource (1-4): 1
```

Next, the console program will show you the basic parameters of the Job and ask you:

```
Run Verify job
JobName:  MatouVerify
FileSet:  Verify Set
Level:    Catalog
Client:   MatouVerify
Storage:  DLTDrive
OK to run? (yes/mod/no): mod
```

Here, you want to respond **mod** to modify the parameters because the Level is by default set to **Catalog** and we want to run an **InitCatalog** Job. After responding **mod**, the console will ask:

```
Parameters to modify:
     1: Job
     2: Level
     3: FileSet
     4: Client
     5: Storage
Select parameter to modify (1-5): 2
```

you should select number 2 to modify the **Level**, and it will display:

```
Levels:
     1: Initialize Catalog
     2: Verify from Catalog
     3: Verify Volume
     4: Verify Volume Data
Select level (1-4): 1
```

Choose item 1, and you will see the final display:

```
Run Verify job
JobName:  MatouVerify
FileSet:  Verify Set
Level:    Initcatalog
Client:   MatouVerify
Storage:  DLTDrive
OK to run? (yes/mod/no): yes
```

at which point you respond **yes**, and the Job will begin.

There after the Job will automatically start according to the schedule you have defined. If you wish to immediately verify it, you can simply run a Verify **Catalog** which will be the default. No differences should be found.

What To Do When Differences Are Found

If you have setup your messages correctly, you should be notified if there are any differences and exactly what they are. For example, below is the email received after doing an update of OpenSSH:

```
HeadMan: Start Verify JobId 83 Job=RufusVerify.2002-06-25.21:41:05
HeadMan: Verifying against Init JobId 70 run 2002-06-21 18:58:51
HeadMan: File: /etc/pam.d/sshd
HeadMan:       st_ino   differ. Cat: 4674b File: 46765
HeadMan: File: /etc/rc.d/init.d/sshd
HeadMan:       st_ino   differ. Cat: 56230 File: 56231
HeadMan: File: /etc/ssh/ssh_config
HeadMan:       st_ino   differ. Cat: 81317 File: 8131b
HeadMan:       st_size  differ. Cat: 1202 File: 1297
HeadMan:       SHA1 differs.
HeadMan: File: /etc/ssh/sshd_config
HeadMan:       st_ino   differ. Cat: 81398 File: 81325
HeadMan:       st_size  differ. Cat: 1182 File: 1579
HeadMan:       SHA1 differs.
```

```
HeadMan: File: /etc/ssh/ssh_config.rpmnew
HeadMan:       st_ino   differ. Cat: 812dd File: 812b3
HeadMan:       st_size  differ. Cat: 1167 File: 1114
HeadMan:       SHA1 differs.
HeadMan: File: /etc/ssh/sshd_config.rpmnew
HeadMan:       st_ino   differ. Cat: 81397 File: 812dd
HeadMan:       st_size  differ. Cat: 2528 File: 2407
HeadMan:       SHA1 differs.
HeadMan: File: /etc/ssh/moduli
HeadMan:       st_ino   differ. Cat: 812b3 File: 812ab
HeadMan: File: /usr/bin/scp
HeadMan:       st_ino   differ. Cat: 5e07e File: 5e343
HeadMan:       st_size  differ. Cat: 26728 File: 26952
HeadMan:       SHA1 differs.
HeadMan: File: /usr/bin/ssh-keygen
HeadMan:       st_ino   differ. Cat: 5df1d File: 5e07e
HeadMan:       st_size  differ. Cat: 80488 File: 84648
HeadMan:       SHA1 differs.
HeadMan: File: /usr/bin/sftp
HeadMan:       st_ino   differ. Cat: 5e2e8 File: 5df1d
HeadMan:       st_size  differ. Cat: 46952 File: 46984
HeadMan:       SHA1 differs.
HeadMan: File: /usr/bin/slogin
HeadMan:       st_ino   differ. Cat: 5e359 File: 5e2e8
HeadMan: File: /usr/bin/ssh
HeadMan:       st_mode  differ. Cat: 89ed File: 81ed
HeadMan:       st_ino   differ. Cat: 5e35a File: 5e359
HeadMan:       st_size  differ. Cat: 219932 File: 234440
HeadMan:       SHA1 differs.
HeadMan: File: /usr/bin/ssh-add
HeadMan:       st_ino   differ. Cat: 5e35b File: 5e35a
HeadMan:       st_size  differ. Cat: 76328 File: 81448
HeadMan:       SHA1 differs.
HeadMan: File: /usr/bin/ssh-agent
HeadMan:       st_ino   differ. Cat: 5e35c File: 5e35b
HeadMan:       st_size  differ. Cat: 43208 File: 47368
HeadMan:       SHA1 differs.
HeadMan: File: /usr/bin/ssh-keyscan
HeadMan:       st_ino   differ. Cat: 5e35d File: 5e96a
HeadMan:       st_size  differ. Cat: 139272 File: 151560
HeadMan:       SHA1 differs.
HeadMan: 25-Jun-2002 21:41
JobId:                  83
Job:                    RufusVerify.2002-06-25.21:41:05
FileSet:                Verify Set
Verify Level:           Catalog
Client:                 RufusVerify
Start time:             25-Jun-2002 21:41
End time:               25-Jun-2002 21:41
Files Examined:         4,258
Termination:            Verify Differences
```

At this point, it was obvious that these files were modified during installation of the RPMs. If you want to be super safe, you should run a **Verify Level=Catalog** immediately before installing new software to verify that there are no differences, then run a **Verify Level=InitCatalog** immediately after the installation.

To keep the above email from being sent every night when the Verify Job runs, we simply re-run the Verify Job setting the level to **InitCatalog** (as we did above in the very beginning). This will re-establish the current state of the system as your new basis for future comparisons. Take care that you don't do an **InitCatalog** after someone has placed a Trojan horse on your system!

If you have included in your **FileSet** a file that is changed by the normal operation of your system, you will get false matches, and you will need to modify the **FileSet** to exclude that file (or not to Include it), and then re-run the **InitCatalog**.

The FileSet that is show below is what I use on my RedHat 7.3 system. With a bit more thought, you can probably add quite a number of additional files that should be monitored.

A Verify Configuration Example

```
Schedule {
  Name = "VerifyCycle"
  Run = Level=Catalog sun-sat at 5:05
}
Job {
  Name = "MatouVerify"
  Type = Verify
  Level = Catalog                     # default level
  Client = MatouVerify
  FileSet = "Verify Set"
  Messages = Standard
  Storage = DLTDrive
  Pool = Default
  Schedule = "VerifyCycle"
}
#
# The list of files in this FileSet should be carefully
# chosen. This is a good starting point.
#
FileSet {
  Name = "Verify Set"
  Include = verify=pins1 signature=SHA1 {
     /boot
     /bin
```

```
    /sbin
    /usr/bin
    /lib
    /root/.ssh
    /home/kern/.ssh
    /var/named
    /etc/sysconfig
    /etc/ssh
    /etc/security
    /etc/exports
    /etc/rc.d/init.d
    /etc/sendmail.cf
    /etc/sysctl.conf
    /etc/services
    /etc/xinetd.d
    /etc/hosts.allow
    /etc/hosts.deny
    /etc/hosts
    /etc/modules.conf
    /etc/named.conf
    /etc/pam.d
    /etc/resolv.conf
  }
  Exclude = { }
}
Client {
  Name = MatouVerify
  Address = lmatou
  Catalog = Bacula
  Password = ""
  File Retention = 80d                # 80 days
  Job Retention = 1y                  # one year
  AutoPrune = yes                     # Prune expired Jobs/Files
}
Catalog {
  Name = Bacula
  dbname = verify; user = bacula; password = ""
}
```

Bacula® - RPM Packaging FAQ

1. How do I build Bacula for platform xxx?
2. How do I control which database support gets built?
3. What other defines are used?
4. I'm getting errors about not having permission when I try to build the packages. Do I need to be r
5. I'm building my own rpms but on all platforms and compiles I get an unresolved dependancy for so

Answers

1. **How do I build Bacula for platform xxx?** The bacula spec file contains defines to build for several platforms: RedHat 7.x (rh7), RedHat 8.0 (rh8), RedHat 9 (rh9), Fedora Core (fc1), Whitebox Enterprise Linux (RHEL) 3.0 (wb3), Mandrake 10.x (mdk) and SuSE 9.x (su9). The package build is controlled by a mandatory define set at the beginning of the file. These defines basically just control the dependancy information that gets coded into the finished rpm package. The platform define may be edited in the spec file directly (by default all defines are set to 0 or "not set"). For example, to build the RedHat 7.x package find the line in the spec file which reads

   ```
   %define rh7 0
   ```

   and edit it to read

   ```
   %define rh7 1
   ```

   Alternately you may pass the define on the command line when calling rpmbuild:

   ```
   rpmbuild -ba --define "build_rh7 1" bacula.spec
   rpmbuild --rebuild --define build_rh7 1" bacula-x.x.x-x.src.rpm
   ```

2. **How do I control which database support gets built?** Another mandatory build define controls which database support is compiled, one of build_sqlite, build_mysql or build_postgresql. To get the MySQL package and support either set the

```
%define mysql 0
```

to

```
%define mysql 1
```

in the spec file directly or pass it to rpmbuild on the command line:

```
rpmbuild -ba --define "build_rh7 1" --define "build_mysql 1" bacula.spec
```

3. **What other defines are used?** Two other building defines of note are the depkgs_version and tomsrtbt identifiers. These two defines are set with each release and must match the version of those sources that are being used to build the packages. You would not ordinarily need to edit these.

4. **I'm getting errors about not having permission when I try to build the packages. Do I need to be root?** No, you do not need to be root and, in fact, it is better practice to build rpm packages as a non-root user. Bacula packages are designed to be built by a regular user but you must make a few changes on your system to do this. If you are building on your own system then the simplest method is to add write permissions for all to the build directory (/usr/src/redhat/). To accomplish this execute the following command as root:

   ```
   chmod -R 777 /usr/src/redhat
   ```

   If you are working on a shared system where you can not use the method above then you need to recreate the /usr/src/redhat directory tree with all of it's subdirectories inside your home directory. Then create a file named `.rpmmacros` in your home directory (or edit the file if it already exists) and add the following line:

   ```
   %_topdir /home/myuser/redhat
   ```

5. **I'm building my own rpms but on all platforms and compiles I get an unresolved dependancy for something called /usr/afsws/bin/pagsh.** This is a shell from the OpenAFS (Andrew File System). If you are seeing this then you chose to include the

docs/examples directory in your package. One of the example scripts in this directory is a pagsh script. Rpmbuild, when scanning for dependancies, looks at the shebang line of all packaged scripts in addition to checking shared libraries. To avoid this do not package the examples directory.

The Bootstrap File

The information in this chapter is provided so that you may either create your own bootstrap files, or so that you can edit a bootstrap file produced by **Bacula**. However, normally the bootstrap file will be automatically created for you during the restore command in the Console program, or by using a Write Bootstrap record in your Backup Jobs, and thus you will never need to know the details of this file.

The **bootstrap** file contains ASCII information that permits precise specification of what files should be restored. It is a relatively compact form of specifying the information, is human readable, and can be edited with any text editor.

File Format

The general format of a **bootstrap** file is:

<**keyword**>= <**value**>

Where each **keyword** and the **value** specify which files to restore. More precisely the **keyword** and their **values** serve to limit which files will be restored and thus act as a filter. The absence of a keyword means that all records will be accepted.

Blank lines and lines beginning with a pound sign (#) in the bootstrap file are ignored.

There are keywords which permit filtering by Volume, Client, Job, FileIndex, Session Id, Session Time, ...

The more keywords that are specified, the more selective the specification of which files to restore will be. In fact, each keyword is **AND**ed with other keywords that may be present.

For example,

```
Volume = Test-001
VolSessionId = 1
VolSessionTime = 108927638
```

directs the Storage daemon (or the **bextract** program) to restore only those files on Volume Test-001 **AND** having VolumeSessionId equal to one **AND** having VolumeSession time equal to 108927638.

The full set of permitted keywords presented in the order in which they are matched against the Volume records are:

Volume The value field specifies what Volume the following commands apply to. Each Volume specification becomes the current Volume, to which all the following commands apply until a new current Volume (if any) is specified. If the Volume name contains spaces, it should be enclosed in quotes.

Count The value is the total number of files that will be restored for this Volume. This allows the Storage daemon to know when to stop reading the Volume.

VolFile The value is a file number, a list of file numbers, or a range of file numbers numbers to match on the current Volume. The file number represents the physical file on the Volume where the data is stored. For a tape volume, this record is used to position to the correct starting file, and once the tape is past the last specified file, reading will stop.

VolBlock The value is a block number, a list of block numbers, or a range of block numbers numbers to match on the current Volume. The block number represents the physical block on the Volume where the data is stored. This record is currently not used.

VolSessionTime The value specifies a Volume Session Time to be matched from the current volume.

VolSessionId The value specifies a VolSessionId, a list of volume session ids, or a range of volume session ids to be matched from the current Volume. Each VolSessionId and VolSessionTime pair corresponds to a unique Job that is backed up on the Volume.

JobId The value specifies a JobId, list of JobIds, or range of JobIds to be selected from the current Volume. Note, the JobId may not be unique if you have multiple Directors, or if you have reinitialized your database. The JobId filter works only if you do not run multiple simultaneous jobs.

Job The value specifies a Job name or list of Job names to be matched on the current Volume. The Job corresponds to a unique VolSessionId and VolSessionTime pair. However, the Job is perhaps a bit more readable by humans. Standard regular expressions (wildcards) may be used to match Job names. The Job filter works only if you do not run multiple simultaneous jobs.

Client The value specifies a Client name or list of Clients to will be matched on the current Volume. Standard regular expressions (wildcards) may be used to match Client names. The Client filter works only if you do not run multiple simultaneous jobs.

FileIndex The value specifies a FileIndex, list of FileIndexes, or range of FileIndexes to be selected from the current Volume. Each file (data) stored on a Volume within a Session has a unique FileIndex. For each Session, the first file written is assigned FileIndex equal to one and incremented for each file backed up.

This for a given Volume, the triple VolSessionId, VolSessionTime, and FileIndex uniquely identifies a file stored on the Volume. Multiple copies of the same file may be stored on the same Volume, but for each file, the triple VolSessionId, VolSessionTime, and FileIndex will be unique. This triple is stored in the Catalog database for each file.

Slot The value specifies the autochanger slot. There may be only a single **Slot** specification for each Volume.

Stream The value specifies a Stream, a list of Streams, or a range of Streams to be selected from the current Volume. Unless you really know what you are doing (the internals of **Bacula**, you should avoid this specification.

***JobType** Not yet implemented.

***JobLevel** Not yet implemented.

The **Volume** record is a bit special in that it must be the first record. The other keyword records may appear in any order and any number following a Volume record.

Multiple Volume records may be specified in the same bootstrap file, but each one starts a new set of filter criteria for the Volume.

In processing the bootstrap file within the current Volume, each filter specified by a keyword is **AND**ed with the next. Thus,

```
Volume = Test-01
Client = "My machine"
FileIndex = 1
```

will match records on Volume **Test-01 AND** Client records for **My machine AND** FileIndex equal to **one**.

Multiple occurrences of the same record are **OR**ed together. Thus,

```
Volume = Test-01
Client = "My machine"
Client = "Backup machine"
FileIndex = 1
```

will match records on Volume **Test-01 AND** (Client records for **My machine OR Backup machine**) **AND** FileIndex equal to **one**.

For integer values, you may supply a range or a list, and for all other values except Volumes, you may specify a list. A list is equivalent to multiple records of the same keyword. For example,

```
Volume = Test-01
Client = "My machine", "Backup machine"
FileIndex = 1-20, 35
```

will match records on Volume **Test-01 AND (**Client records for **My machine OR Backup machine) AND (**FileIndex 1 **OR** 2 **OR** 3 ... **OR** 20 **OR** 35**)**.

As previously mentioned above, there may be multiple Volume records in the same bootstrap file. Each new Volume definition begins a new set of filter conditions that apply to that Volume and will be **OR**ed with any other Volume definitions.

As an example, suppose we query for the current set of tapes to restore all files on Client **Rufus** using the **query** command in the console program:

```
Using default Catalog name=MySQL DB=bacula
*query
Available queries:
     1: List Job totals:
     2: List where a file is saved:
     3: List where the most recent copies of a file are saved:
     4: List total files/bytes by Job:
     5: List total files/bytes by Volume:
     6: List last 10 Full Backups for a Client:
     7: List Volumes used by selected JobId:
     8: List Volumes to Restore All Files:
Choose a query (1-8): 8
Enter Client Name: Rufus
+-------+------------------+------------+-----------+----------+------------+
| JobId | StartTime        | VolumeName | StartFile | VolSesId | VolSesTime |
+-------+------------------+------------+-----------+----------+------------+
| 154   | 2002-05-30 12:08 | test-02    | 0         | 1        | 1022753312 |
| 202   | 2002-06-15 10:16 | test-02    | 0         | 2        | 1024128917 |
| 203   | 2002-06-15 11:12 | test-02    | 3         | 1        | 1024132350 |
```

```
| 204   | 2002-06-18 08:11 | test-02    | 4         | 1        | 1024380678 |
+-------+------------------+------------+-----------+----------+------------+
```

The output shows us that there are four Jobs that must be restored. The first one is a Full backup, and the following three are all Incremental backups.

The following bootstrap file will restore those files:

```
Volume=test-02
VolSessionId=1
VolSessionTime=1022753312
Volume=test-02
VolSessionId=2
VolSessionTime=1024128917
Volume=test-02
VolSessionId=1
VolSessionTime=1024132350
Volume=test-02
VolSessionId=1
VolSessionTime=1024380678
```

As a final example, assume that the initial Full save spanned two Volumes. The output from **query** might look like:

```
+-------+------------------+------------+-----------+----------+------------+
| JobId | StartTime        | VolumeName | StartFile | VolSesId | VolSesTime |
+-------+------------------+------------+-----------+----------+------------+
| 242   | 2002-06-25 16:50 | File0003   | 0         | 1        | 1025016612 |
| 242   | 2002-06-25 16:50 | File0004   | 0         | 1        | 1025016612 |
| 243   | 2002-06-25 16:52 | File0005   | 0         | 2        | 1025016612 |
| 246   | 2002-06-25 19:19 | File0006   | 0         | 2        | 1025025494 |
+-------+------------------+------------+-----------+----------+------------+
```

and the following bootstrap file would restore those files:

```
Volume=File0003
VolSessionId=1
VolSessionTime=1025016612
Volume=File0004
VolSessionId=1
VolSessionTime=1025016612
Volume=File0005
VolSessionId=2
VolSessionTime=1025016612
Volume=File0006
VolSessionId=2
VolSessionTime=1025025494
```

Automatic Generation of Bootstrap Files

One thing that is probably worth knowing: the bootstrap files that are generated automatically at the end of the job are not as optimized as those generated by the restore command. This is because the ones created at the end of the file, contain all files written to the Volume for that job. As a consequence, all the files saved to an Incremental or Differential job will be restored first by the Full save, then by any Incremental or Differential saves.

When the bootstrap file is generated for the restore command, only one copy (the most recent) of each file is restored.

So if you have spare cycles on your machine, you could optimize the bootstrap files by doing the following:

```
./console
restore client=xxx select all
no
quit
Backup bootstrap file.
```

The above will not work if you have multiple FileSets because that will be an extra prompt. However, the **restore client=xxx select all** builds the in-memory tree, selecting everything and creates the bootstrap file.

The **no** answers the **Do you want to run this (yes/mod/no)** question.

A Final Example

If you want to extract or copy a single Job, you can do it by selecting by JobId (code not tested) or better yet, if you know the VolSessionTime and the VolSessionId (printed on Job report and in Catalog), specifying this is by far the best. Using the VolSessionTime and VolSessionId is the way Bacula does restores. A bsr file might look like the following:

```
Volume="Vol001"
VolSessionId=10
VolSessionTime=1080847820
```

If you know how many files are backed up (on the job report), you can enormously speed up the selection by adding (let's assume there are 157 files):

```
FileIndex=1-157
Count=157
```

Finally, if you know the File number where the Job starts, you can also cause bcopy to forward space to the right file without reading every record:

```
VolFile=20
```

There is nothing magic or complicated about a BSR file. Parsing it and properly applying it within Bacula *is* magic, but you don't need to worry about that.

If you want to see a *real* bsr file, simply fire up the **restore** command in the console program, select something, then answer no when it prompts to run the job. Then look at the file **restore.bsr** in your working directory.

Catalog Services

General

This chapter is intended to be a technical discussion of the Catalog services and as such is not targeted at end users but rather at developers and system administrators that want or need to know more of the working details of **Bacula**.

The **Bacula Catalog** services consist of the programs that provide the SQL database engine for storage and retrieval of all information concerning files that were backed up and their locations on the storage media.

We have investigated the possibility of using the following SQL engines for Bacula: Beagle, mSQL, GNU SQL, PostgreSQL, SQLite, Oracle, and MySQL. Each presents certain problems with either licensing or maturity. At present, we have chosen for development purposes to use MySQL, PostgreSQL and SQLite. MySQL was chosen because it is fast, proven to be reliable, widely used, and actively being developed. MySQL is released under the GNU GPL license. PostgreSQL was chosen because it is a full-featured, very mature database, and because Dan Langille did the Bacula driver for it. PostgreSQL is distributed under the BSD license. SQLite was chosen because it is small, efficient, and can be directly embedded in **Bacula** thus requiring much less effort from the system administrator or person building **Bacula**. In our testing SQLite has performed very well, and for the functions that we use, it has never encountered any errors except that it does not appear to handle databases larger than 2GBytes.

The Bacula SQL code has been written in a manner that will allow it to be easily modified to support any of the current SQL database systems on the market (for example: mSQL, iODBC, unixODBC, Solid, OpenLink ODBC, EasySoft ODBC, InterBase, Oracle8, Oracle7, and DB2).

If you do not specify either **--with-mysql** or **--with-postgresql** or **--with-sqlite** on the ./configure line, Bacula will use its minimalist internal database. This database is kept for build reasons but is no longer supported. Bacula **requires** one of the three databases (MySQL, PostgreSQL, or SQLite) to run.

Filenames and Maximum Filename Length

In general, either MySQL, PostgreSQL or SQLite permit storing arbitrary long path names and file names in the catalog database. In practice, there

still may be one or two places in the Catalog interface code that restrict the maximum path length to 512 characters and the maximum file name length to 512 characters. These restrictions are believed to have been removed. Please note, these restrictions apply only to the Catalog database and thus to your ability to list online the files saved during any job. All information received and stored by the Storage daemon (normally on tape) allows and handles arbitrarily long path and filenames.

Installing and Configuring MySQL

For the details of installing and configuring MySQL, please see the Installing and Configuring MySQL chapter of this manual.

Installing and Configuring PostgreSQL

For the details of installing and configuring PostgreSQL, please see the Installing and Configuring PostgreSQL chapter of this manual.

Installing and Configuring SQLite

For the details of installing and configuring SQLite, please see the Installing and Configuring SQLite chapter of this manual.

Internal Bacula Catalog

Please see the Internal Bacula Database chapter of this manual for more details.

Database Table Design

All discussions that follow pertain to the MySQL database. The details for the PostgreSQL and SQLite databases are essentially identical except for that all fields in the SQLite database are stored as ASCII text and some of the database creation statements are a bit different. The details of the internal Bacula catalog are not discussed here.

Because the Catalog database may contain very large amounts of data for large sites, we have made a modest attempt to normalize the data

tables to reduce redundant information. While reducing the size of the database significantly, it does, unfortunately, add some complications to the structures.

In simple terms, the Catalog database must contain a record of all Jobs run by Bacula, and for each Job, it must maintain a list of all files saved, with their File Attributes (permissions, create date, ...), and the location and Media on which the file is stored. This is seemingly a simple task, but it represents a huge amount interlinked data. Note: the list of files and their attributes is not maintained when using the internal Bacula database. The data stored in the File records, which allows the user or administrator to obtain a list of all files backed up during a job, is by far the largest volume of information put into the Catalog database.

Although the Catalog database has been designed to handle backup data for multiple clients, some users may want to maintain multiple databases, one for each machine to be backed up. This reduces the risk of confusion of accidental restoring a file to the wrong machine as well as reducing the amount of data in a single database, thus increasing efficiency and reducing the impact of a lost or damaged database.

Sequence of Creation of Records for a Save Job

Start with StartDate, ClientName, Filename, Path, Attributes, MediaName, MediaCoordinates. (PartNumber, NumParts). In the steps below, "Create new" means to create a new record whether or not it is unique. "Create unique" means each record in the database should be unique. Thus, one must first search to see if the record exists, and only if not should a new one be created, otherwise the existing RecordId should be used.

1. Create new Job record with StartDate; save JobId
2. Create unique Media record; save MediaId
3. Create unique Client record; save ClientId
4. Create unique Filename record; save FilenameId
5. Create unique Path record; save PathId
6. Create unique Attribute record; save AttributeId store ClientId, FilenameId, PathId, and Attributes
7. Create new File record store JobId, AttributeId, MediaCoordinates, etc

8. Repeat steps 4 through 8 for each file
9. Create a JobMedia record; save MediaId
10. Update Job record filling in EndDate and other Job statistics

Database Tables

Filename		
Column Name	**Data Type**	**Remark**
FilenameId	integer	Primary Key
Name	Blob	Filename

The **Filename** table shown above contains the name of each file backed up with the path removed. If different directories or machines contain the same filename, only one copy will be saved in this table.

Path		
Column Name	**Data Type**	**Remark**
PathId	integer	Primary Key
Path	Blob	Full Path

The **Path** table contains shown above the path or directory names of all directories on the system or systems. The filename and any MSDOS disk name are stripped off. As with the filename, only one copy of each directory name is kept regardless of how many machines or drives have the same directory. These path names should be stored in Unix path name format.

Some simple testing on a Linux file system indicates that separating the filename and the path may be more complication than is warranted by the space savings. For example, this system has a total of 89,097 files, 60,467 of which have unique filenames, and there are 4,374 unique paths.

Finding all those files and doing two stats() per file takes an average wall clock time of 1 min 35 seconds on a 400MHz machine running RedHat 6.1 Linux.

Finding all those files and putting them directly into a MySQL database with the path and filename defined as TEXT, which is variable length up to 65,535 characters takes 19 mins 31 seconds and creates a 27.6 MByte database.

Doing the same thing, but inserting them into Blob fields with the filename indexed on the first 30 characters and the path name indexed on the 255 (max) characters takes 5 mins 18 seconds and creates a 5.24 MB database. Rerunning the job (with the database already created) takes about 2 mins 50 seconds.

Running the same as the last one (Path and Filename Blob), but Filename indexed on the first 30 characters and the Path on the first 50 characters (linear search done there after) takes 5 mins on the average and creates a 3.4 MB database. Rerunning with the data already in the DB takes 3 mins 35 seconds.

Finally, saving only the full path name rather than splitting the path and the file, and indexing it on the first 50 characters takes 6 mins 43 seconds and creates a 7.35 MB database.

File		
Column Name	**Data Type**	**Remark**
FileId	integer	Primary Key
FileIndex	integer	The sequential file number in the Job
JobId	integer	Link to Job Record
PathId	integer	Link to Path Record
FilenameId	integer	Link to Filename Record
MarkId	integer	Used to mark files during Verify Jobs
LStat	tinyblob	File attributes in base64 encoding
MD5	tinyblob	MD5 signature in base64 encoding

The **File** table shown above contains one entry for each file backed up by Bacula. Thus a file that is backed up multiple times (as is normal) will have multiple entries in the File table. This will probably be the table with the most number of records. Consequently, it is essential to keep the size of this record to an absolute minimum. At the same time, this table must contain all the information (or pointers to the information) about the file and where it is backed up. Since a file may be backed up many times without having changed, the path and filename are stored in separate tables.

This table contains by far the largest amount of information in the Catalog database, both from the stand point of number of records, and the stand point of total database size. As a consequence, the user must take care to periodically reduce the number of File records using the **retention** command in the Console program.

Job		
Column Name	**Data Type**	**Remark**
JobId	integer	Primary Key
Job	tinyblob	Unique Job Name
Name	tinyblob	Job Name
PurgedFiles	tinyint	Used by Bacula for purging/retention periods
Type	binary(1)	Job Type: Backup, Copy, Clone, Archive, Migration
Level	binary(1)	Job Level
ClientId	integer	Client index
JobStatus	binary(1)	Job Termination Status
SchedTime	datetime	Time/date when Job scheduled
StartTime	datetime	Time/date when Job started
EndTime	datetime	Time/date when Job ended
JobTDate	bigint	Start day in Unix format but 64 bits; used for Retention period.
VolSessionId	integer	Unique Volume Session ID
VolSessionTime	integer	Unique Volume Session Time
JobFiles	integer	Number of files saved in Job
JobBytes	bigint	Number of bytes saved in Job
JobErrors	integer	Number of errors during Job
JobMissingFiles	integer	Number of files not saved (not yet used)
PoolId	integer	Link to Pool Record
FileSetId	integer	Link to FileSet Record
PurgedFiles	tiny integer	Set when all File records purged
HasBase	tiny integer	Set when Base Job run

The **Job** table contains one record for each Job run by Bacula. Thus normally, there will be one per day per machine added to the database. Note, the JobId is used to index Job records in the database, and it often is shown to the user in the Console program. However, care must be taken with its use as it is not unique from database to database. For example, the user may have a database for Client data saved on machine Rufus and another database for Client data saved on machine Roxie. In this case, the two database will each have JobIds that match those in another database. For a unique reference to a Job, see Job below.

The Name field of the Job record corresponds to the Name resource record given in the Director's configuration file. Thus it is a generic name, and it

will be normal to find many Jobs (or even all Jobs) with the same Name.

The Job field contains a combination of the Name and the schedule time of the Job by the Director. Thus for a given Director, even with multiple Catalog databases, the Job will contain a unique name that represents the Job.

For a given Storage daemon, the VolSessionId and VolSessionTime form a unique identification of the Job. This will be the case even if multiple Directors are using the same Storage daemon.

The Job Type (or simply Type) can have one of the following values:

Value	**Meaning**
B	Backup Job
V	Verify Job
R	Restore Job
C	Console program (not in database)
D	Admin Job
A	Archive Job (not implemented)

The JobStatus field specifies how the job terminated, and can be one of the following:

Value	**Meaning**
C	Created but not yet running
R	Running
B	Blocked
T	Terminated normally
E	Terminated in Error
e	Non-fatal error
f	Fatal error
D	Verify Differences
A	Canceled by the user
F	Waiting on the File daemon
S	Waiting on the Storage daemon
m	Waiting for a new Volume to be mounted
M	Waiting for a Mount
s	Waiting for Storage resource
j	Waiting for Job resource
c	Waiting for Client resource
d	Wating for Maximum jobs
t	Waiting for Start Time

p	Waiting for higher priority job to finish

FileSet		
Column Name	**Data Type**	**Remark**
FileSetId	integer	Primary Key
FileSet	tinyblob	FileSet name
MD5	tinyblob	MD5 checksum of FileSet
CreateTime	datetime	Time and date Fileset created

The **FileSet** table contains one entry for each FileSet that is used. The MD5 signature is kept to ensure that if the user changes anything inside the FileSet, it will be detected and the new FileSet will be used. This is particularly important when doing an incremental update. If the user deletes a file or adds a file, we need to ensure that a Full backup is done prior to the next incremental.

JobMedia		
Column Name	**Data Type**	**Remark**
JobMediaId	integer	Primary Key
JobId	integer	Link to Job Record
MediaId	integer	Link to Media Record
FirstIndex	integer	The index (sequence number) of the first file written for this Job to the Media
LastIndex	integer	The index of the last file written for this Job to the Media
StartFile	integer	The physical media (tape) file number of the first block written for this Job
EndFile	integer	The physical media (tape) file number of the last block written for this Job
StartBlock	integer	The number of the first block written for this Job
EndBlock	integer	The number of the last block written for this Job
VolIndex	integer	The Volume use sequence number within the Job

The **JobMedia** table contains one entry for each volume written for the current Job. If the Job spans 3 tapes, there will be three JobMedia records, each containing the information to find all the files for the given JobId on the tape.

Media		
Column Name	**Data Type**	**Remark**
MediaId	integer	Primary Key
VolumeName	tinyblob	Volume name
Slot	integer	Autochanger Slot number or zero
PoolId	integer	Link to Pool Record
MediaType	tinyblob	The MediaType supplied by the user
FirstWritten	datetime	Time/date when first written
LastWritten	datetime	Time/date when last written
LabelDate	datetime	Time/date when tape labeled
VolJobs	integer	Number of jobs written to this media
VolFiles	integer	Number of files written to this media
VolBlocks	integer	Number of blocks written to this media
VolMounts	integer	Number of time media mounted
VolBytes	bigint	Number of bytes saved in Job
VolErrors	integer	Number of errors during Job
VolWrites	integer	Number of writes to media
MaxVolBytes	bigint	Maximum bytes to put on this media
VolCapacityBytes	bigint	Capacity estimate for this volume
VolStatus	enum	Status of media: Full, Archive, Append, Recycle, Read-Only, Disabled, Error, Busy
Recycle	tinyint	Whether or not Bacula can recycle the Volumes: Yes, No
VolRetention	bigint	64 bit seconds until expiration
VolUseDuration	bigint	64 bit seconds volume can be used
MaxVolJobs	integer	maximum jobs to put on Volume
MaxVolFiles	integer	maximume EOF marks to put on Volume

The **Volume** table (internally referred to as the Media table) contains one

entry for each volume, that is each tape, cassette (8mm, DLT, DAT, ...), or file on which information is or was backed up. There is one Volume record created for each of the NumVols specified in the Pool resource record.

Pool		
Column Name	**Data Type**	**Remark**
PoolId	integer	Primary Key
Name	Tinyblob	Pool Name
NumVols	Integer	Number of Volumes in the Pool
MaxVols	Integer	Maximum Volumes in the Pool
UseOnce	tinyint	Use volume once
UseCatalog	tinyint	Set to use catalog
AcceptAnyVolume	tinyint	Accept any volume from Pool
VolRetention	bigint	64 bit seconds to retain volume
VolUseDuration	bigint	64 bit seconds volume can be used
MaxVolJobs	integer	max jobs on volume
MaxVolFiles	integer	max EOF marks to put on Volume
MaxVolBytes	bigint	max bytes to write on Volume
AutoPrune	tinyint	yes—no for autopruning
Recycle	tinyint	yes—no for allowing auto recycling of Volume
PoolType	enum	Backup, Copy, Cloned, Archive, Migration
LabelFormat	Tinyblob	Label format

The **Pool** table contains one entry for each media pool controlled by Bacula in this database. One media record exists for each of the NumVols contained in the Pool. The PoolType is a Bacula defined keyword. The MediaType is defined by the administrator, and corresponds to the MediaType specified in the Director's Storage definition record. The CurrentVol is the sequence number of the Media record for the current volume.

Client		
Column Name	**Data Type**	**Remark**
ClientId	integer	Primary Key
Name	TinyBlob	File Services Name
UName	TinyBlob	uname -a from Client (not yet used)
AutoPrune	tinyint	yes—no for autopruning
FileRetention	bigint	64 bit seconds to retain Files

JobRetention	bigint	64 bit seconds to retain Job

The **Client** table contains one entry for each machine backed up by Bacula in this database. Normally the Name is a fully qualified domain name.

UnsavedFiles		
Column Name	**Data Type**	**Remark**
UnsavedId	integer	Primary Key
JobId	integer	JobId corresponding to this record
PathId	integer	Id of path
FilenameId	integer	Id of filename

The **UnsavedFiles** table contains one entry for each file that was not saved. Note! This record is not yet implemented.

Counter		
Column Name	**Data Type**	**Remark**
Counter	tinyblob	Counter name
MinValue	integer	Start/Min value for counter
MaxValue	integer	Max value for counter
CurrentValue	integer	Current counter value
WrapCounter	tinyblob	Name of another counter

The **Counter** table contains one entry for each permanent counter defined by the user.

Version		
Column Name	**Data Type**	**Remark**
VersionId	integer	Primary Key

The **Version** table defines the Bacula database version number. Bacula checks this number before reading the database to ensure that it is compatible with the Bacula binary file.

BaseFiles		
Column Name	**Data Type**	**Remark**
BaseId	integer	Primary Key
BaseJobId	integer	JobId of Base Job
JobId	integer	Reference to Job
FileId	integer	Reference to File
FileIndex	integer	File Index number

The **BaseFiles** table contains all the File references for a particular JobId that point to a Base file – i.e. they were previously saved and hence were not saved in the current JobId but in BaseJobId under FileId. FileIndex is the index of the file, and is used for optimization of Restore jobs to prevent the need to read the FileId record when creating the in memory tree. This record is not yet implemented.

MySQL Table Definition

The commands used to create the MySQL tables are as follows:

```
USE bacula;
CREATE TABLE Filename (
  FilenameId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
  Name BLOB NOT NULL,
  PRIMARY KEY(FilenameId),
  INDEX (Name(30))
  );
CREATE TABLE Path (
   PathId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
   Path BLOB NOT NULL,
   PRIMARY KEY(PathId),
   INDEX (Path(50))
   );
CREATE TABLE File (
   FileId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
   FileIndex INTEGER UNSIGNED NOT NULL DEFAULT 0,
   JobId INTEGER UNSIGNED NOT NULL REFERENCES Job,
   PathId INTEGER UNSIGNED NOT NULL REFERENCES Path,
   FilenameId INTEGER UNSIGNED NOT NULL REFERENCES Filename,
   MarkId INTEGER UNSIGNED NOT NULL DEFAULT 0,
   LStat TINYBLOB NOT NULL,
   MD5 TINYBLOB NOT NULL,
   PRIMARY KEY(FileId),
   INDEX (JobId),
```

```
   INDEX (PathId),
   INDEX (FilenameId)
   );
CREATE TABLE Job (
   JobId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
   Job TINYBLOB NOT NULL,
   Name TINYBLOB NOT NULL,
   Type BINARY(1) NOT NULL,
   Level BINARY(1) NOT NULL,
   ClientId INTEGER NOT NULL REFERENCES Client,
   JobStatus BINARY(1) NOT NULL,
   SchedTime DATETIME NOT NULL,
   StartTime DATETIME NOT NULL,
   EndTime DATETIME NOT NULL,
   JobTDate BIGINT UNSIGNED NOT NULL,
   VolSessionId INTEGER UNSIGNED NOT NULL DEFAULT 0,
   VolSessionTime INTEGER UNSIGNED NOT NULL DEFAULT 0,
   JobFiles INTEGER UNSIGNED NOT NULL DEFAULT 0,
   JobBytes BIGINT UNSIGNED NOT NULL,
   JobErrors INTEGER UNSIGNED NOT NULL DEFAULT 0,
   JobMissingFiles INTEGER UNSIGNED NOT NULL DEFAULT 0,
   PoolId INTEGER UNSIGNED NOT NULL REFERENCES Pool,
   FileSetId INTEGER UNSIGNED NOT NULL REFERENCES FileSet,
   PurgedFiles TINYINT NOT NULL DEFAULT 0,
   HasBase TINYINT NOT NULL DEFAULT 0,
   PRIMARY KEY(JobId),
   INDEX (Name(128))
   );
CREATE TABLE FileSet (
   FileSetId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
   FileSet TINYBLOB NOT NULL,
   MD5 TINYBLOB NOT NULL,
   CreateTime DATETIME NOT NULL,
   PRIMARY KEY(FileSetId)
   );
CREATE TABLE JobMedia (
   JobMediaId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
   JobId INTEGER UNSIGNED NOT NULL REFERENCES Job,
   MediaId INTEGER UNSIGNED NOT NULL REFERENCES Media,
   FirstIndex INTEGER UNSIGNED NOT NULL DEFAULT 0,
   LastIndex INTEGER UNSIGNED NOT NULL DEFAULT 0,
   StartFile INTEGER UNSIGNED NOT NULL DEFAULT 0,
   EndFile INTEGER UNSIGNED NOT NULL DEFAULT 0,
   StartBlock INTEGER UNSIGNED NOT NULL DEFAULT 0,
   EndBlock INTEGER UNSIGNED NOT NULL DEFAULT 0,
   VolIndex INTEGER UNSIGNED NOT NULL DEFAULT 0,
   PRIMARY KEY(JobMediaId),
   INDEX (JobId, MediaId)
   );
CREATE TABLE Media (
   MediaId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
   VolumeName TINYBLOB NOT NULL,
   Slot INTEGER NOT NULL DEFAULT 0,
   PoolId INTEGER UNSIGNED NOT NULL REFERENCES Pool,
```

```
   MediaType TINYBLOB NOT NULL,
   FirstWritten DATETIME NOT NULL,
   LastWritten DATETIME NOT NULL,
   LabelDate DATETIME NOT NULL,
   VolJobs INTEGER UNSIGNED NOT NULL DEFAULT 0,
   VolFiles INTEGER UNSIGNED NOT NULL DEFAULT 0,
   VolBlocks INTEGER UNSIGNED NOT NULL DEFAULT 0,
   VolMounts INTEGER UNSIGNED NOT NULL DEFAULT 0,
   VolBytes BIGINT UNSIGNED NOT NULL DEFAULT 0,
   VolErrors INTEGER UNSIGNED NOT NULL DEFAULT 0,
   VolWrites INTEGER UNSIGNED NOT NULL DEFAULT 0,
   VolCapacityBytes BIGINT UNSIGNED NOT NULL,
   VolStatus ENUM('Full', 'Archive', 'Append', 'Recycle', 'Purged',
    'Read-Only', 'Disabled', 'Error', 'Busy', 'Used', 'Cleaning') NOT NULL,
   Recycle TINYINT NOT NULL DEFAULT 0,
   VolRetention BIGINT UNSIGNED NOT NULL DEFAULT 0,
   VolUseDuration BIGINT UNSIGNED NOT NULL DEFAULT 0,
   MaxVolJobs INTEGER UNSIGNED NOT NULL DEFAULT 0,
   MaxVolFiles INTEGER UNSIGNED NOT NULL DEFAULT 0,
   MaxVolBytes BIGINT UNSIGNED NOT NULL DEFAULT 0,
   InChanger TINYINT NOT NULL DEFAULT 0,
   MediaAddressing TINYINT NOT NULL DEFAULT 0,
   VolReadTime BIGINT UNSIGNED NOT NULL DEFAULT 0,
   VolWriteTime BIGINT UNSIGNED NOT NULL DEFAULT 0,
   PRIMARY KEY(MediaId),
   INDEX (PoolId)
   );
CREATE TABLE Pool (
   PoolId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
   Name TINYBLOB NOT NULL,
   NumVols INTEGER UNSIGNED NOT NULL DEFAULT 0,
   MaxVols INTEGER UNSIGNED NOT NULL DEFAULT 0,
   UseOnce TINYINT NOT NULL,
   UseCatalog TINYINT NOT NULL,
   AcceptAnyVolume TINYINT DEFAULT 0,
   VolRetention BIGINT UNSIGNED NOT NULL,
   VolUseDuration BIGINT UNSIGNED NOT NULL,
   MaxVolJobs INTEGER UNSIGNED NOT NULL DEFAULT 0,
   MaxVolFiles INTEGER UNSIGNED NOT NULL DEFAULT 0,
   MaxVolBytes BIGINT UNSIGNED NOT NULL,
   AutoPrune TINYINT DEFAULT 0,
   Recycle TINYINT DEFAULT 0,
   PoolType ENUM('Backup', 'Copy', 'Cloned', 'Archive', 'Migration', 'Scratch') NOT NULL,
   LabelFormat TINYBLOB,
   Enabled TINYINT DEFAULT 1,
   ScratchPoolId INTEGER UNSIGNED DEFAULT 0 REFERENCES Pool,
   RecyclePoolId INTEGER UNSIGNED DEFAULT 0 REFERENCES Pool,
   UNIQUE (Name(128)),
   PRIMARY KEY (PoolId)
   );
CREATE TABLE Client (
   ClientId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
   Name TINYBLOB NOT NULL,
   Uname TINYBLOB NOT NULL,       /* full uname -a of client */
```

```
   AutoPrune TINYINT DEFAULT 0,
   FileRetention BIGINT UNSIGNED NOT NULL,
   JobRetention  BIGINT UNSIGNED NOT NULL,
   UNIQUE (Name(128)),
   PRIMARY KEY(ClientId)
   );
CREATE TABLE BaseFiles (
   BaseId INTEGER UNSIGNED AUTO_INCREMENT,
   BaseJobId INTEGER UNSIGNED NOT NULL REFERENCES Job,
   JobId INTEGER UNSIGNED NOT NULL REFERENCES Job,
   FileId INTEGER UNSIGNED NOT NULL REFERENCES File,
   FileIndex INTEGER UNSIGNED,
   PRIMARY KEY(BaseId)
   );
CREATE TABLE UnsavedFiles (
   UnsavedId INTEGER UNSIGNED AUTO_INCREMENT,
   JobId INTEGER UNSIGNED NOT NULL REFERENCES Job,
   PathId INTEGER UNSIGNED NOT NULL REFERENCES Path,
   FilenameId INTEGER UNSIGNED NOT NULL REFERENCES Filename,
   PRIMARY KEY (UnsavedId)
   );
CREATE TABLE Version (
   VersionId INTEGER UNSIGNED NOT NULL
   );
-- Initialize Version
INSERT INTO Version (VersionId) VALUES (7);
CREATE TABLE Counters (
   Counter TINYBLOB NOT NULL,
   MinValue INTEGER,
   MaxValue INTEGER,
   CurrentValue INTEGER,
   WrapCounter TINYBLOB NOT NULL,
   PRIMARY KEY (Counter(128))
   );
```

Installing and Configuring MySQL

Installing and Configuring MySQL – Phase I

If you use the ./configure --with-mysql=mysql-directory statement for configuring **Bacula**, you will need MySQL version 3.23.33 or later installed in the **mysql-directory** (we are currently using 3.23.56). If MySQL is installed in the standard system location, you need only enter **--with-mysql** since the configure program will search all the standard locations. If you install MySQL in your home directory or some other non-standard directory, you will need to provide the full path to it.

Installing and Configuring MySQL is not difficult but can be confusing the first time. As a consequence, below, we list the steps that we used to install it on our machines. Please note that our configuration leaves MySQL without any user passwords. This may be an undesirable situation if you have other users on your system.

Please note that as of Bacula version 1.31, the thread safe version of the MySQL client library is used, and hence you must add the **--enable-thread-safe-client** option to the **./configure** as shown below:

1. Download MySQL source code from www.mysql.com/downloads
2. Detar it with something like:

   tar xvfz mysql-filename

   Note, the above command requires GNU tar. If you do not have GNU tar, a command such as:

   zcat mysql-filename — tar xvf -

   will probably accomplish the same thing.
3. cd **mysql-source-directory**

   where you replace **mysql-source-directory** with the directory name where you put the MySQL source code.
4. ./configure --enable-thread-safe-client --prefix=mysql-directory

   where you replace **mysql-directory** with the directory name where you want to install mysql. Normally for system wide use this is /usr/local/mysql. In my case, I use ~kern/mysql.
5. make

   This takes a bit of time.

6. make install

   This will put all the necessary binaries, libraries and support files into the **mysql-directory** that you specified above.

7. ./scripts/mysql_install_db

   This will create the necessary MySQL databases for controlling user access. Note, this script can also be found in the **bin** directory in the installation directory

The MySQL client library **mysqlclient** requires the gzip compression library **libz.a** or **libz.so**. If you are using rpm packages, these libraries are in the **libz-devel** package. On Debian systems, you will need to load the **zlib1g-dev** package. If you are not using rpms or debs, you will need to find the appropriate package for your system.

At this point, you should return to completing the installation of **Bacula**. Later after Bacula is installed, come back to this chapter to complete the installation. Please note, the installation files used in the second phase of the MySQL installation are created during the Bacula Installation.

Installing and Configuring MySQL – Phase II

At this point, you should have built and installed MySQL, or already have a running MySQL, and you should have configured, built and installed **Bacula**. If not, please complete these items before proceeding.

Please note that the **./configure** used to build **Bacula** will need to include **--with-mysql=mysql-directory**, where **mysql-directory** is the directory name that you specified on the ./configure command for configuring MySQL. This is needed so that Bacula can find the necessary include headers and library files for interfacing to MySQL.

Bacula will install scripts for manipulating the database (create, delete, make tables etc) into the main installation directory. These files will be of the form *_bacula_* (e.g. create_bacula_database). These files are also available in the <bacula-src>/src/cats directory after running ./configure. If you inspect create_bacula_database, you will see that it calls create_mysql_database. The *_bacula_* files are provided for convenience. It doesn't matter what database you have chosen; create_bacula_database will always create your database.

Now you will create the Bacula MySQL database and the tables that Bacula uses.

1. Start **mysql**. You might want to use the **startmysql** script provided in the Bacula release.

2. cd <install-directory>

   This directory contains the Bacula catalog interface routines.

3. ./grant_mysql_privileges

   This script creates unrestricted access rights for **kern**, **kelvin**, and **bacula**. You may want to modify it to suit your situation. Please note that none of these userids including root are password protected.

4. ./create_mysql_database

   This script creates the MySQL **bacula** database. The databases you create as well as the access databases will be located in <install-dir>/var/ in a subdirectory with the name of the database, where <install-dir> is the directory name that you specified on the **--prefix** option. This can be important to know if you want to make a special backup of the Bacula database or to check its size.

5. ./make_mysql_tables

   This script creates the MySQL tables used by **Bacula**.

Each of the three scripts (grant_mysql_privileges, create_mysql_database and make_mysql_tables) allows the addition of a command line argument. This can be useful for specifying the user and or password. For example, you might need to add **-u root** to the command line to have sufficient privilege to create the Bacula tables.

To take a closer look at the access privileges that you have setup with the above, you can do:

```
mysql-directory/bin/mysql -u root mysql
select * from user;
```

Re-initializing the Catalog Database

After you have done some initial testing with **Bacula**, you will probably want to re-initialize the catalog database and throw away all the test Jobs that you ran. To do so, you can do the following:

```
cd <install-directory>
./drop_mysql_tables
./make_mysql_tables
```

Please note that all information in the database will be lost and you will be starting from scratch. If you have written on any Volumes, you must write and end of file mark on the volume so that Bacula can reuse it. Do so with:

```
(stop Bacula or unmount the drive)
mt -f /dev/nst0 rewind
mt -f /dev/nst0 weof
```

Where you should replace **/dev/nst0** with the appropriate tape drive device name for your machine.

Linking Bacula with MySQL

After configuring Bacula with

./configure --enable-thread-safe-client --prefix=<mysql-directory> where <mysql-directory> is in my case **/home/kern/mysql**, you may have to configure the loader so that it can find the MySQL shared libraries. If you have previously followed this procedure and later add the **--enable-thread-safe-client** options, you will need to rerun the **ldconfig** program shown below. If you put MySQL in a standard place such as **/usr/lib** or **/usr/local/lib** this will not be necessary, but in my case it is. The description that follows is Linux specific. For other operating systems, please consult your manuals on how to do the same thing:

First edit: **/etc/ld.so.conf** and add a new line to the end of the file with the name of the mysql-directory. In my case, it is:

/home/kern/mysql/lib/mysql then rebuild the loader's cache with:

/sbin/ldconfig If you upgrade to a new version of **MySQL**, the shared library names will probably changes, and you must re-run the **/sbin/ldconfig** command so that the runtime loader can find them.

Alternatively, your system my have a loader environment variable that can be set. For example, on a Solaris system where I do not have root permission, I use:

LD_LIBRARY_PATH=/home/kern/mysql/lib/mysql

Finally, if you have encryption enabled in MySQL, you may need to add **-lssl -lcrypto** to the link. In that case, you can either export the appropriate LDFLAGS definition, or alternatively, you can include them directly on the ./configure line as in:

```
LDFLAGS="-lssl -lcyrpto" \
   ./configure \
      <your-options>
```

Installing and Configuring PostgreSQL

Installing and Configuring PostgreSQL – Phase I

If you use the **./configure --with-postgresql=PostgreSQL-Directory** statement for configuring **Bacula**, you will need PostgreSQL version 7.3 or later installed. NOTE! PostgreSQL versions earlier than 7.3 do not work with Bacula. If PostgreSQL is installed in the standard system location, you need only enter **--with-postgresql** since the configure program will search all the standard locations. If you install PostgreSQL in your home directory or some other non-standard directory, you will need to provide the full path with the **--with-postgresql** option.

Installing and configuring PostgreSQL is not difficult but can be confusing the first time. If you prefer, you may want to use a package provided by your chosen operating system. Binary packages are available on most PostgreSQL mirrors.

If you prefer to install from source, we recommend following the instructions found in the PostgreSQL documentation.

If you are using FreeBSD, this FreeBSD Diary article will be useful. Even if you are not using FreeBSD, the article will contain useful configuration and setup information.

After installing PostgreSQL, you should return to completing the installation of **Bacula**. Later, after Bacula is installed, come back to this chapter to complete the installation. Please note, the installation files used in the second phase of the PostgreSQL installation are created during the Bacula Installation.

Installing and Configuring PostgreSQL – Phase II

At this point, you should have built and installed PostgreSQL, or already have a running PostgreSQL, and you should have configured, built and installed **Bacula**. If not, please complete these items before proceeding.

Please note that the **./configure** used to build **Bacula** will need to include **--with-postgresql=PostgreSQL-directory**, where **PostgreSQL-directory** is the directory name that you specified on the ./configure command for configuring PostgreSQL (if you didn't specify a directory or PostgreSQL is installed in a default location, you do not need to specify the directory). This is needed so that Bacula can find the necessary include

headers and library files for interfacing to PostgreSQL.

Bacula will install scripts for manipulating the database (create, delete, make tables etc) into the main installation directory. These files will be of the form *_bacula_* (e.g. create_bacula_database). These files are also available in the <bacula-src>/src/cats directory after running ./configure. If you inspect create_bacula_database, you will see that it calls create_postgresql_database. The *_bacula_* files are provided for convenience. It doesn't matter what database you have chosen; create_bacula_database will always create your database.

Now you will create the Bacula PostgreSQL database and the tables that Bacula uses. These instructions assume that you already have PostgreSQL running. You will need to perform these steps as a user that is able to create new databases. This can be the PostgreSQL user (on most systems, this is the pgsql user).

1. cd <install-directory>

   This directory contains the Bacula catalog interface routines.

2. ./create_bacula_database

   This script creates the PostgreSQL **bacula** database.

3. ./make_bacula_tables

   This script creates the PostgreSQL tables used by **Bacula**.

4. ./grant_bacula_privileges

   This script creates the database user **bacula** with restricted access rights. You may want to modify it to suit your situation. Please note that this database is not password protected.

Each of the three scripts (create_bacula_database, make_bacula_tables, and grant_bacula_privileges) allows the addition of a command line argument. This can be useful for specifying the user name. For example, you might need to add **-h hostname** to the command line to specify a remote database server.

To take a closer look at the access privileges that you have setup with the above, you can do:

```
PostgreSQL-directory/bin/psql --command \\dp bacula
```

Re-initializing the Catalog Database

After you have done some initial testing with **Bacula**, you will probably want to re-initialize the catalog database and throw away all the test Jobs that you ran. To do so, you can do the following:

```
cd <install-directory>
./drop_bacula_tables
./make_bacula_tables
./grant_bacula_privileges
```

Please note that all information in the database will be lost and you will be starting from scratch. If you have written on any Volumes, you must write and end of file mark on the volume so that Bacula can reuse it. Do so with:

```
(stop Bacula or unmount the drive)
mt -f /dev/nst0 rewind
mt -f /dev/nst0 weof
```

Where you should replace **/dev/nst0** with the appropriate tape drive device name for your machine.

Converting from MySQL to PostgreSQL

The conversion procedure presented here was worked out by Norm Dressler <ndressler at dinmar dot com>

This process was tested using the following software versions:

- Linux Mandrake 10/Kernel 2.4.22-10 SMP
- Mysql Ver 12.21 Distrib 4.0.15, for mandrake-linux-gnu (i586)
- PostgreSQL 7.3.4
- Bacula 1.34.5

WARNING: Always as a precaution, take a complete backup of your databases before proceeding with this process!

1. Shutdown bacula (cd /etc/bacula;./bacula stop)

2. Run the following command to dump your Mysql database:

```
mysqldump -f -t -n >bacula-backup.dmp>
```

3. Make a backup of your /etc/bacula directory (but leave the original in place).

4. Go to your Bacula source directory and rebuild it to include PostgreSQL support rather then Mysql support. Check the config.log file for your original configure command and replace enable-mysql with enable-postgresql.

5. Recompile Bacula with a make and if everything compiles completely, perform a make install.

6. Shutdown Mysql.

7. Start PostgreSQL on your system.

8. Create a bacula user in Postgres with createuser command. Depending on your Postgres install, you may have to SU to the user who has privileges to create a user.

9. Verify your pg_hba.conf file contains sufficient permissions to allow bacula to access the server. Mine has the following since it's on a secure network:

```
local all all trust

host all all 127.0.0.1 255.255.255.255 trust

NOTE: you should restart your postgres server if you
      made changes
```

10. Change into the /etc/bacula directory and prepare the database and tables with the following commands:

```
./create_postgresql_database

./make_postgresql_tables

./grant_postgresql_privileges
```

11. Verify you have access to the database:

```
psql -Ubacula bacula
```

You should not get any errors.

12. Load your database from the Mysql database dump with:

```
psql -Ubacula bacula <bacula-backup.dmp>
```

13. Reseqence your tables with the following commands:

```
psql -Ubacula bacula

SELECT SETVAL('basefiles_baseid_seq', (SELECT
MAX(baseid) FROM basefiles));
SELECT SETVAL('client_clientid_seq', (SELECT
MAX(clientid) FROM client));
SELECT SETVAL('file_fileid_seq', (SELECT MAX(fileid)
FROM file));
SELECT SETVAL('filename_filenameid_seq', (SELECT
MAX(filenameid) FROM filename));

SELECT SETVAL('fileset_filesetid_seq', (SELECT
MAX(filesetid) FROM fileset));

SELECT SETVAL('job_jobid_seq', (SELECT MAX(jobid) FROM job));
SELECT SETVAL('jobmedia_jobmediaid_seq', (SELECT
MAX(jobmediaid) FROM jobmedia));
SELECT SETVAL('media_mediaid_seq', (SELECT MAX(mediaid) FROM media));
SELECT SETVAL('path_pathid_seq', (SELECT MAX(pathid) FROM path));

SELECT SETVAL('pool_poolid_seq', (SELECT MAX(poolid) FROM pool));
```

14. At this point, start up Bacula, verify your volume library and perform a test backup to make sure everything is working properly.

Credits

Many thanks to Dan Languille for writing the PostgreSQL driver. This will surely become the most popular database that Bacula supports.

Installing and Configuring SQLite

Installing and Configuring SQLite – Phase I

If you use the **./configure --with-sqlite** statement for configuring **Bacula**, you will need SQLite version 2.2.3 or later installed. Our standard location (for the moment) for SQLite is in the dependency package **depkgs/sqlite-2.2.3**. Please note that the version will be updated as new versions are available and tested.

Installing and Configuring is quite easy.

1. Download the Bacula dependency packages
2. Detar it with something like:

   tar xvfz depkgs.tar.gz

   Note, the above command requires GNU tar. If you do not have GNU tar, a command such as:

   zcat depkgs.tar.gz — tar xvf -

   will probably accomplish the same thing.
3. **cd depkgs**
4. **make sqlite**

At this point, you should return to completing the installation of **Bacula**.

Please note that the **./configure** used to build **Bacula** will need to include **--with-sqlite**.

Installing and Configuring SQLite – Phase II

This phase is done **after** you have run the **./configure** command to configure **Bacula**.

Bacula will install scripts for manipulating the database (create, delete, make tables etc) into the main installation directory. These files will be of the form *_bacula_* (e.g. create_bacula_database). These files are also available in the <bacula-src>/src/cats directory after running ./configure. If you inspect create_bacula_database, you will see that it calls create_sqlite_database. The *_bacula_* files are provided for convenience. It

doesn't matter what database you have chosen; create_bacula_database will always create your database.

At this point, you can create the SQLite database and tables:

1. cd <install-directory>

   This directory contains the Bacula catalog interface routines.

2. ./make_sqlite_tables

   This script creates the SQLite database as well as the tables used by **Bacula**. This script will be automatically setup by the **./configure** program to create a database named **bacula.db** in **Bacula's** working directory.

Linking Bacula with SQLite

If you have followed the above steps, this will all happen automatically and the SQLite libraries will be linked into **Bacula**.

Testing SQLite

As of this date (20 March 2002), we have much less "production" experience using SQLite than using MySQL. That said, we should note that SQLite has performed flawlessly for us in all our testing.

Re-initializing the Catalog Database

After you have done some initial testing with **Bacula**, you will probably want to re-initialize the catalog database and throw away all the test Jobs that you ran. To do so, you can do the following:

```
cd <install-directory>
./drop_sqlite_tables
./make_sqlite_tables
```

Please note that all information in the database will be lost and you will be starting from scratch. If you have written on any Volumes, you must write and end of file mark on the volume so that Bacula can reuse it. Do so with:

```
(stop Bacula or unmount the drive)
mt -f /dev/nst0 rewind
mt -f /dev/nst0 weof
```

Where you should replace **/dev/nst0** with the appropriate tape drive device name for your machine.

The Bacula internal database is no longer supported, please do not use it.

Internal Bacula Database

Previously it was intended to be used primarily by Bacula developers for testing; although SQLite is also a good choice for this. We do not recommend its use in general.

This database is simplistic in that it consists entirely of Bacula's internal structures appended sequentially to a file. Consequently, it is in most cases inappropriate for sites with many clients or systems with large numbers of files, or long-term production environments.

Below, you will find a table comparing the features available with SQLite and MySQL and with the internal Bacula database. At the current time, you cannot dynamically switch from one to the other, but must rebuild the Bacula source code. If you wish to experiment with both, it is possible to build both versions of Bacula and install them into separate directories.

Feature	SQLite or MySQL	Bacula
Job Record	Yes	Yes
Media Record	Yes	Yes
FileName Record	Yes	No
File Record	Yes	No
FileSet Record	Yes	Yes
Pool Record	Yes	Yes
Client Record	Yes	Yes
JobMedia Record	Yes	Yes
List Job Records	Yes	Yes
List Media Records	Yes	Yes
List Pool Records	Yes	Yes
List JobMedia Records	Yes	Yes
Delete Pool Record	Yes	Yes
Delete Media Record	Yes	Yes
Update Pool Record	Yes	Yes
Implement Verify	Yes	No
MD5 Signatures	Yes	No

In addition, since there is no SQL available, the Console commands: **sqlquery**, **query**, **retention**, and any other command that directly uses SQL are not available with the Internal database.

Disaster Recovery Using a Bacula Rescue Floppy

General

Please note that the Bacula Rescue Floppy is now deprecated and is and is replaced by the Bacula Rescue CDROM described in another chapter of this manual.

When disaster strikes, you must have a plan, and you must have prepared in advance otherwise the work of recovering your system and your files will be considerably greater. For example, if you have not previously saved the partitioning information for your hard disk, how can you properly rebuild it if the disk must be replaced?

Unfortunately, many of the steps one must take before and immediately after a disaster are very operating system dependent. As a consequence, this chapter will discuss in detail disaster recovery (also called Bare Metal Recovery) for **Linux** and **Solaris**. For Solaris, the procedures are still quite manual. For FreeBSD the same procedures may be used but they are not yet developed. For Win32, no luck. Apparently an "emergency boot" disk allowing access to the full system API without interference does not exist.

Important Considerations

Here are a few important considerations concerning disaster recovery that you should take into account before a disaster strikes.

- If the building which houses your computers burns down or is otherwise destroyed, do you have off-site backup data?
- Disaster recovery is much easier if you have several machines. If you have a single machine, how will you handle unforeseen events if your only machine is down?
- Do you want to protect your whole system and use Bacula to recover everything? or do you want to try to restore your system from the original installation disks and apply any other updates and only restore user files?

Steps to Take Before Disaster Strikes

- Create a Bacula Rescue floppy for each of your Linux systems.
- Ensure that you always have a valid bootstrap file that is saved to an alternate machine.
- If possible copy your catalog nightly to an alternate machine. If you have a valid bootstrap file, this is not necessary, but can be very useful if you do not want to reload everything. .
- Test using the Bacula Rescue floppy before you are forced to use it in an emergency situation.

Bare Metal Floppy Recovery on Linux with a Bacula Floppy Rescue Disk

Since floppies are being used less and less, the Bacula Floppy rescue disk is deprecated, which means that it is no longer really supported. For those of you who have or need floppy rescue, we include the recovery instructions here for your reference.

The remainder of this section concerns recovering a **Linux** computer using a floppy, and parts of it relate to the Red Hat version of Linux.

A so called "Bare Metal" recovery is one where you start with an empty hard disk and you restore your machine. There are also cases where you may lose a file or a directory and want it restored. Please see the previous chapter for more details for those cases.

Bare Metal Recovery assumes that you have the following four items for your system:

- An emergency boot disk allowing you to boot without a hard disk
- A Bacula Rescue floppy containing your disk information and a number of helpful scripts (described below) including a statically linked version of the Bacula File daemon
- A full Bacula backup of your system possibly including Incremental or Differential backups since the last Full backup
- A second system running the Bacula Director, the Catalog, and the Storage daemon. (this is not an absolute requirement, but how to get around it is not yet documented here)

Restrictions

In addition, to the above assumptions, the following conditions or restrictions apply:

- Linux only – tested only on Red Hat, but should work on other Linuxes
- The scripts handle only SCSI and IDE disks
- All partitions will be recreated, but only **ext2**, **ext3**, and **swap** partitions will be reformatted. Any other partitions such as Windows FAT partitions will not be formatted by the scripts, but you can do it by hand
- You are using either **lilo** or **grub** as a boot loader, and you know which one (not automatically detected)
- The partitioning and reformating scripts will *should* work with RAID devices, but probably not with other "complicated" disk partitioning/formating schemes. Please check them carefully. You will probably need to edit the scripts by hand to make them work.

Directories

If you are building a self-contained Bacula Rescue CDROM, you will find the necessary scripts in **rescule/linux/cdrom** subdirectory of the Bacula source code.

If you wish to build the Bacula Rescue floppy disk, the scripts discussed below can be found in the **rescue/linux/floppy** subdirectory of the Bacula source code.

Preparation for a Bare Metal Recovery

There are two things you should do immediately on all (Linux) systems for which you wish to do a bare metal recovery:

1. Create a system emergency boot disk or alternatively a system installation boot floppy. This step can be skipped if you have an Installation CDROM and your machine will boot from CDROM (most modern computers will).

2. Create a Bacula Rescue floppy, which captures the current working state of your computer and creates scripts to restore it. In addition, it creates a statically linked version of the Bacula File daemon (Client) program, which is key to successfully restoring from scratch.

Creating an Emergency Boot Disk

Here you have several choices:

- Create a tomsrtbt disk (any Linux)
- Create an emergency boot disk (any Linux I think)
- Create a Red Hat Installation disk (Red Hat specific)
- Others

tomsrtbt: If you have created a Bacula Rescue CDROM, you can skip this section.

If you *must* use a boot floppy, my preference is to create and use a **tomsrtbt** emergency boot disk because it gives you a very clean Linux environment (with a 2.2 kernel) and the most utilities. See http://www.toms.net/rb/ for more details on this. It is very easy to do and well worth the effort. However, I recommend that you create both especially if you have non-standard hardware. You may find that **tomsrtbt** will not work with your network driver (he surely has one, but you must explicitly put it on the disk), whereas the Linux rescue is more likely to work.

Emergency Boot Disk: If you have created a Bacula Rescue CDROM, you can skip this section.

To create a standard Linux emergency boot disk you must first know the name of the kernel, which you can find with:

```
ls -l /boot
```

and looking on the **vmlinux-...** line or alternative do an

```
uname -a
```

then become root and with a blank floppy in the drive, enter the following command:

```
mkbootdisk --device /dev/fd0 2.4.18-18
```

where you replace "2.4.18-18" by your system name.

This disk can then be booted and you will be in an environment with a number of important tools available. Some disadvantages of this environment as opposed to **tomsrtbt** are that you must enter **linux rescue** at the boot prompt or the boot will fail without a hard disk; it requires a disk boot image or a CDROM to be mounted, if the CDROM is released, you will loose a large number of the tools.

Red Hat Installation Disk: If you have created a Bacula Rescue CDROM, you can skip this section.

Specific to Red Hat Linux, is to create an Installation floppy, which can also be used as an emergency boot disk. The advantage of this method is that it works in conjunction with the installation CDROM and hence during the first part of restoring the system, you have a much larger number of tools available (on the CDROM). This can be extremely useful if you are not sure what really happened and you need to examine your system in detail.

To make a Red Hat Linux installation disk, do the following:

```
mount the Installation CDROM (/mnt/cdrom)
cd /mnt/cdrom/images
dd if=boot.img of=/dev/fd0 bs=1440k
```

Now that you have either an emergency boot disk or an installation floppy, you will be able to reboot your system in the absence of your hard disk or with a damaged hard disk. This method has the same disadvantages compared to **tomsrtbt** disk as mentioned above for the Emergency Boot Disk.

Creating a Bacula Rescue Disk

If you have created a Bacula Rescue CDROM, this step will be automatically done for you.

Simply having a boot disk is not sufficient to re-create things as they were. To solve this problem, we will create a Bacula Rescue disk. Everything that will be written to this disk will first be placed into the **<bacula-src>/rescue/linux** directory.

The first step is while your system is up and running normally, you use a Bacula script called **getdiskinfo** to capture certain important information about your hard disk configuration (partitioning, formatting, mount points, ...). **getdiskinfo** will also create a number of scripts using the information found that can be used in an emergency to repartition your disks, reformat them, and restore a statically linked version of the Bacula file daemon so that your disk can be restored from within a minimal boot environment.

The first step is to run **getdiskinfo** as follows:

```
su
cd <bacula-src>/rescue/linux
./getdiskinfo
```

getdiskinfo works for either IDE or SCSI drives and recognizes both ext2 and ext3 file systems. If you wish to restore other file systems, you will need to modify the code. This script can be run multiple times, but really only needs to be run once unless you change your hard disk configuration.

Assuming you have a single hard disk on device /dev/hda, **getdiskinfo** will create the following files:

partition.hda This file contains the shell commands to repartition your hard disk drive /dev/hda to the current state. If you have additional drives (e.g. /dev/hdc), you will find one of these files for each drive. DO NOT EXECUTE THIS SCRIPT UNLESS YOU WANT YOUR HARD DISK REPARTITIONED

format.hda This file contains the shell commands that will format each of the partitions on your hard drive. It knows about ext2, ext3, and swap partitions. All other partitions, you must manually format. It is recommended that any Microsoft partitions be partitioned with Microsoft's **format** command rather than using Unix tools. DO NOT EXECUTE THIS SCRIPT UNLESS YOU WANT YOUR HARD DISK REFORMATTED

mount_drives This script will mount all ext2 and ext3 drives that were previously mounted. They will be mounted on

/mnt/drive/. This is used just before running the statically linked Bacula so that it can access your drives for the restore.

restore_bacula This script will restore the File daemon from the Bacula Rescue disk. Building the Bacula Rescue disk will be described later. This will provide your emergency boot environment with a Bacula file daemon. Note, this is a special statically linked version of the file daemon (i.e. it does not need or use shared libraries).

start_network This script will start your network using the simplest possible commands. You will need to verify that the IP address used in this script is correct. In addition, if you have several ethernet cards, you may need to make other modifications to this script.

sfdisk This is the program that will repartition your hard disk, and it is normally found in **/sbin/sfdisk**. It is placed in this directory so that it will be included on the rescue disk as it is not normally available with all emergency boot environments.

sfdisk.gz This is the version of sfdisk that works with **tomsrtbt**. The standard sfdisk described above will not run under tomsrtbt.

The **getdiskinfo** program (actually a shell script) will also create a subdirectory named **diskinfo**, which contains the following files:

```
df.bsi
disks.bsi
fstab.bsi
ifconfig.bsi
mount.bsi
mount.ext2.bsi
mount.ext3.bsi
mtab.bsi
route.bsi
sfdisk.disks.bsi
sfdisk.hda.bsi
sfdisk.make.hda.bsi
```

Each of these files contains some important piece of information (sometimes redundant) about your hard disk setup or your network. Normally, you will not need this information, but it will be written to the Bacula Rescue disk just in case. Since it is normally not used, we will leave it to you to examine those files at your leisure.

Building a Static File Daemon: If you have created a Bacula Rescue CDROM, this step will be automatically done for you.

The second of the three steps in creating your Bacula Rescue disk is to build a static version of the File daemon. Do so by either configuring Bacula as

follows or by allowing the **make_rescue_disk** script described below make it for you:

```
cd <bacula-src>
./configure <normal-options>
make
cd src/filed
make static-bacula-fd
strip static-bacula-fd
cp static-bacula-fd ../../rescue/linux/bacula-fd
cp bacula-fd.conf ../../rescue/linux
```

Note, above, we built **static-bacula-fd** and changed its name to **bacula-fd** when copying it to the rescue/linux directory.

Finally, in <bacula-src>/rescue/linux, ensure that the WorkingDirectory and PIDDirectory both point to reasonable locations on a stripped down system. If you are using **tomsrtbt** you will also want to replace machine names with IP addresses since there is no resolver running. With the Linux Rescue disk, network address mapping seems to work. Don't forget that at the time this version of the Bacula File daemon runs, your file system will not be restored. In my bacula-fd.conf, I use **/var/working**.

Writing the Bacula Rescue Floppy: When you have everything you need (output of getdiskinfo, Bacula File daemon, ...), you create your rescue floppy by putting a blank tape into your floppy disk drive and entering:

```
su
./make_rescue_disk
```

This script will reformat the floppy and write everything in the current directory and all files in the **diskinfo** directory to the floppy. If you supply the appropriate command line options, it will also build a static version of the Bacula file daemon and copy it along with the configuration file to the disk. Also using a command line option, you can make it write a compressed tar file containing all the files whose names are in **backup.etc.list** to the floppy. The list as provided contains names of files in **/etc** that you might need in a disaster situation. It is not needed, but in some cases such as a complex network setup, you may find it useful.

Options for make_rescue_disk: The following command line options are available for the make_rescue_disk script:

```
Usage: make_rescue_disk
  -h, --help             print this message
  --make-static-bacula   make static File daemon and add to diskette
  --copy-static-bacula   copy static File daemon to diskette
  --copy-etc-files       copy files in etc list to diskette
```

Briefly the options are:

Please examine the contents of the rescue floppy to ensure that it has everything you want and need. If not modify the scripts as necessary and re-run it until it is correct. Now that you have both a system boot floppy and a Bacula Rescue floppy, assuming you have a full backup of your system made by Bacula, you are ready to handle nearly any kind of emergency restoration situation. **Restoring Your Linux Client with a Floppy**

Now, let's assume that your hard disk has just died and that you have replaced it with an new identical drive. In addition, we assume that you have:A recent Bacula backup (Full plus Incrementals) An emergency boot floppy (preferably **tomsrtbt**) A Bacula Rescue Floppy Disk Your Bacula Director, Catalog, and Storage daemon running on another machine on your local network.

This is a relatively simple case, and later in this chapter, as time permits, we will discuss how you might recover from a situation where the machine that crashes is your main Bacula server (i.e. has the Director, the Catalog, and the Storage daemon).

You will take the following steps to get your system back up and running:

1. Boot with your Emergency Floppy
2. Mount your Bacula Rescue floppy
3. Start the Network (local network)
4. Re-partition your hard disk(s) as it was before
5. Re-format your partitions
6. Restore the Bacula File daemon (static version)
7. Perform a Bacula restore of all your files
8. Re-install your boot loader
9. Reboot

Now for the details ...

Boot with your Emergency Floppy

First you will boot with your emergency floppy. If you use the Installation floppy described above, when you get to the boot prompt:

```
boot:
```

you enter **linux rescue**.

If you are booting from **tomsrtbt** simply enter the default responses.

When your machine finishes booting, you should be at the command prompt possibly with your hard disk mounted on **/mount/sysimage** (Linux emergency only). To see what is actually mounted, use:

```
df
```

Mount your Bacula Rescue Floppy: Make sure that the mount point **/mnt/floppy** exists. If not, enter:

```
mkdir -p /mnt/floppy
```

the mount your **Bacula Rescue** disk and cd to it with:

```
mount /dev/fd0 /mnt/floppy
cd /mnt/floppy
```

To simplify running the scripts make sure the current directory is on your path by:

```
PATH=$PATH:.
```

Start the Network: At this point, you should bring up your network. Normally, this is quite simple and requires just a few commands. If you have booted from your Bacula Rescue CDROM, please cd into the /bacula-hostname directory before continuing. To simplify your task, we have created a script that should work in most cases by typing:

```
./start_network
```

You can test it by pinging another machine, or pinging your broken machine machine from another machine. Do not proceed until your network is up.

Unmount Your Hard Disk (if mounted): When you are sure you want to repartition your disk, normally, if your disk was damaged or if you are using **tomsrtbt** your hard disk will not be mounted. However, if it is you must first unmount it so that it is not in use. Do so by entering **df** and then enter the correct commands to unmount the disks. For example:

```
umount /mnt/sysimage/boot
umount /mnt/sysimage/usr
umount /mnt/sysimage/proc
umount /mnt/sysimage/
```

where you explicitly unmount (**umount**) each sysimage partition and finally, the last one being the root. Do another **df** command to be sure you successfully unmount all the sysimage partitions.

This is necessary because **sfdisk** will refuse to partition a disk that is currently mounted. As mentioned, this should never be necessary with **tomsrtbt**.

Partition Your Hard Disk(s): If you are using **tomsrtbt**, you will need to do the following steps to get the correct sfdisk:

```
rm -f sfdisk
bzip2 -d sfdisk.bz2
```

Do not do the above steps if you are using a standard Linux boot disk or the Bacula Rescue CDROM.

Then proceed with partitioning your hard disk by:

```
./partition.hda
```

If you have multiple disks, do the same for each of them. For SCSI disks, the repartition script will be named: **partition.sda**. If the script complains about the disk being in use, simply go back and redo the **df** command and **umount** commands until you no longer have your hard disk mounted. Note, in many cases, if your hard disk was seriously damaged or a new one installed, it will not automatically be mounted. If it is mounted, it is because the emergency kernel found one or more possibly valid partitions.

If for some reason this procedure does not work, you can use the information in **partition.hda** to re-partition your disks by hand using **fdisk**.

Format Your Hard Disk(s): After partitioning your disk, you must format it appropriately. The formatting script will put back swap partitions, normal Unix partitions (ext2) and journaled partitions (ext3). Do so by entering for each disk:

```
./format.hda
```

The format script will ask you if you want a block check done. We recommend to answer yes, but realize that for very large disks this can take hours.

Mount the Newly Formatted Disks: Once the disks are partitioned and formatted, you can remount them with the **mount_drives** script. All your drives must be mounted for Bacula to be able to access them. Run the script as follows:

```
./mount_drives
df
```

The **df** will tell you if the drives are mounted. If not, re-run the script again. It isn't always easy to figure out and create the mount points and the mounts in the proper order, so repeating the **./mount_drives** command will not cause any harm and will most likely work the second time. If not, correct it by hand before continuing.

Unmount the CDROM: Next, if you are using the Red Hat installation disk, unmount the CDROM drive by doing:

```
umount /mnt/cdrom
```

This is not necessary if you are running **tomsrtbt**. In doing this, I find it is always busy, and I haven't figured out how to unmount it (Linux boot only).

Restore and Start the File Daemon: If you have booted with a Bacula Rescue CDROM, your statically linked Bacula File daemon and the bacula-fd.conf file with be in the /bacula-hostname/bin directory. Please skip the following paragraph and continue with editing the Bacula configuration file.

If you have not used a Bacula Rescue CDROM, now change (cd) to some directory where you want to put the image of the Bacula File daemon. I use the tmp directory my hard disk (mounted as

/mnt/disk/tmp) because it is easy. Then install into the current directory Bacula by running the **restore_bacula** script from the floppy drive. For example:

```
cd /mnt/disk
mkdir -p /mnt/disk/tmp
mkdir -p /mnt/disk/tmp/working
/mnt/floppy/restore_bacula
ls -l
```

Make sure **bacula-fd** and **bacula-fd.conf** are both there.

Edit the Bacula configuration file, create the working/pid/subsys directory if you haven't already done so above, and start Bacula by entering:

```
chroot /mnt/disk /tmp/bacula-fd -c /tmp/bacula-fd.conf
```

The above command starts the Bacula File daemon with your the proper root disk location (i.e. **/mnt/disk/tmp**. If Bacula does not start correct the problem and start it. You can check if it is running by entering:

```
ps fax
```

You can kill Bacula by entering:

```
kill -TERM <pid>
```

where **pid** is the first number printed in front of the first occurrence of **bacula-fd** in the **ps fax** command.

Now, you should be able to use another computer with Bacula installed to check the status by entering:

```
status client=xxxx
```

into the Console program, where xxxx is the name of the client you are restoring.

One common problem is that your **bacula-dir.conf** may contain machine addresses that are not properly resolved on the stripped down system to be restored because it is not running DNS. This is particularly true for the address in the Storage resource of the Director, which may be very well resolved on the Director's machine, but not on the machine being restored and running the File daemon. In that case, be prepared to edit **bacula-dir.conf** to replace the name of the Storage daemon's domain name with its IP address.

Restoring using the RedHat Installation Disk: Suppose your system was damaged for one reason or another, so that the hard disk and the partitioning and much of the filesystems are intact, but you want to do a full restore. If you have booted into your system with the RedHat Installation Disk by specifying **linux rescue** at the **boot:** prompt, you will find yourself in a shell command with your disks already mounted (if it was possible) in **/mnt/sysimage**. In this case, you can do much like you did above to restore your system:

```
cd /mnt/sysimage/tmp
mkdir -p /mnt/sysimage/tmp/working
/mnt/floppy/restore_bacula
ls -l
```

Make sure that **bacula-fd** and **bacula-fd.conf** are both in the current directory and that the directory names in the **bacula-fd.conf** correctly point to the appropriate directories. Then start Bacula with:

```
chroot /mnt/sysimage /tmp/bacula-fd -c /tmp/bacula-fd.conf
```

Restore Your Files: On the computer that is running the Director, you now run a **restore** command and select the files to be restored (normally everything), but before starting the restore, there is one final change you must make using the **mod** option. You must change the **Where** directory to be the root by using the **mod** option just before running the job and selecting **Where**. Set it to:

```
/
```

then run the restore.

You might be tempted to avoid using **chroot** and running Bacula directly and then using a **Where** to specify a destination of **/mnt/disk**. This is possible, however, the current version of Bacula always restores files to the new location, and thus any soft links that have been specified with absolute paths will end up with **/mnt/disk** prefixed to them. In general this is not fatal to getting your system running, but be aware that you will have to fix these links if you do not use **chroot**.

Final Step: At this point, the restore should have finished with no errors, and all your files will be restored. One last task remains and that is to write a new boot sector so that your machine will boot. For **lilo**, you enter the following command:

```
run_lilo
```

If you are using grub instead of lilo, you must enter the following:

```
run_grub
```

Note, I've had quite a number of problems with **grub** because it is rather complicated and not designed to install easily under a simplified system. So, if you experience errors or end up unexpectedly in a **chroot** shell, simply exit back to the normal shell and type in the appropriate commands from the **run_grub** script by hand until you get it to install.

Reboot: Reboot your machine by entering **exit** until you get to the main prompt then enter **ctl-d**.

If everything went well, you should now be back up and running. If not, re-insert the emergency boot floppy, boot, and figure out what is wrong.

At this point, you will probably want to remove the temporary copy of Bacula that you installed. Do so with:

```
rm -f /bacula-fd /bacula-fd.conf
rm -rf /working
```

Linux Problems or Bugs

Since every flavor and every release of Linux is different, there are likely to be some small difficulties with the scripts, so please be prepared to edit them in a minimal environment. A rudimentary knowledge of **vi** is very useful. Also, these scripts do not do everything. You will need to reformat Windows partitions by hand, for example.

Getting the boot loader back can be a problem if you are using **grub** because it is so complicated. If all else fails, reboot your system from your floppy but using the restored disk image, then proceed to a reinstallation of grub (looking at the run-grub script can help). By contrast, lilo is a piece of cake.

Bugs

When performing the bare metal recovery using the Red Hat emergency boot disk (actually the installation boot disk), I was never able to release the cdrom, and when the system came up **/mnt/cdrom** was soft linked to **/mnt/disk/dev/hdd**, which is not correct. I fixed this in each case by deleting and simply remaking it with **mkdir -p /mnt/cdrom**.

tomsrtbt

This is a single floppy (1.722Meg) that really has A LOT of software. For example, by default (version 2.0.103) you get:

AHA152X AHA1542 AIC7XXX BUSLOGIC DAC960 DEC_ELCP(TULIP) EATA EEXPRESS/PRO/PRO100 EL2 EL3 EXT2 EXT3 FAT FD IDE-CD/DISK/TAPE IMM INITRD ISO9660 JOLIET LOOP MATH_EMULATION MINIX MSDOS NCR53C8XX NE2000 NFS NTFS PARPORT PCINE2K PCNET32 PLIP PPA RTL8139 SD SERIAL/_CONSOLE SLIP SMC_ULTRA SR ST VFAT VID_SELECT VORTEX WD80x3 .exrc 3c589_cs agetty ash badblocks basename boot.b buildit.s busybox bz2bzImage bzip2 cardmgr cardmgr.pid cat chain.b chattr chgrp chmod chown chroot clear clone.s cmp common config cp cpio cs cut date dd dd-lfs debugfs ddate df dhcpcd-- dirname dmesg domainname ds du dumpe2fs e2fsck echo egrep elvis ex false fdflush fdformat fdisk filesize find findsuper fmt fstab grep group gunzip gzip halt head hexdump hexedit host.conf hostname hosts httpd i82365 ifconfig ile init inittab insmod install.s issue kernel key.lst kill killall killall5 ld ld-linux length less libc libcom_err libe2p libext2fs libtermcap libuuid lilo lilo.conf ln loadkmap login ls lsattr lsmod lua luasocket man map md5sum miterm mkdir mkdosfs mke2fs mkfifo mkfs.minix mknod mkswap more more.help mount mt mtab mv nc necho network networks nmclan_cs nslookup passwd pax pcmcia_core pcnet_cs pidof ping poweroff printf profile protocols ps pwd rc.0 rc.S rc.custom rc.custom.gz rc.pcmcia reboot rescuept reset resolv.conf rm rmdir rmmod route rsh rshd script sed serial serial_cs services setserial settings.s sh shared slattach sleep sln sort split stab strings swapoff swapon sync tail tar tcic tee telnet telnetd termcap test tomshexd tomsrtbt.FAQ touch traceroute true tune2fs umount undeb-- unpack.s unrpm-- update utmp vi vi.help view watch wc wget which xargs xirc2ps_cs yecho yes zcat

In addition, at Tom's Web Site, you can find a lot of additional kernel drivers and other software (such as **sdisk**, which is used by Bacula.

Building his floppy is a piece of cake. Simply download his .tar.gz file then:

```
- detar the .tar.gz archive
- become root
- cd to the tomsrtbt-<version> directory
- load a blank floppy with no bad sectors
- ./install.s
```

Bacula Copyright, Trademark, and Licenses

There are a number of different licenses that are used in Bacula.

GPL

The vast bulk of the code is released under a modified version of the GNU General Public License version 2. The modifications (actually additions) are described in the source file LICENSE, and their purpose is not to alter the essential qualities of the GPL but to permit more freedom in linking certain third party software supposedly non-GPL compatable, provide termination for Patent (and IP) actions, clarify contributors IP and Copyright claims and non-infringment intentions. The details and governing text are in the file LICENSE in the main source directory.

Most of this code is copyrighted: Copyright ©2000-2004 Kern Sibbald and John Walker. or Copyright ©2000-2005 Kern Sibbald

Portions may be copyrighted by other people (ATT, the Free Software Foundation, ...). Generally these portions are released under a non-modified GPL 2 license.

LGPL

Some of the Bacula library source code is released under the GNU Lesser General Public License. This permits third parties to use these parts of our code in their proprietary programs to interface to Bacula.

Public Domain

Some of the Bacula code has been released to the public domain. E.g. md5.c, SQLite.

Trademark

Bacula®is a registered trademark of Kern Sibbald and John Walker.

We have trademarked the Bacula name to ensure that any variant of Bacula will be exactly compatible with the program that we have released. The use of the name Bacula is restricted to software systems that agree exactly with the program presented here.

Disclaimer

GNU General Public License

image of a Philosophical GNU

- What to do if you see a possible GPL violation
- Translations of the GPL

Table of Contents

- GNU GENERAL PUBLIC LICENSE
  - Preamble
  - TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION
  - How to Apply These Terms to Your New Programs

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These

restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program

(independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- **a)** You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
- **b)** You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
- **c)** If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of

this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- **a)** Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- **b)** Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- **c)** Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to

copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such

claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING,

BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
{\it one line to give the program's name and an idea of what it does.}
Copyright (C) {\it yyyy}  {\it name of author}
This program is free software; you can redistribute it and/or
modify it under the terms of the GNU General Public License
as published by the Free Software Foundation; either version 2
of the License, or (at your option) any later version.
This program is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
GNU General Public License for more details.
You should have received a copy of the GNU General Public License
along with this program; if not, write to the Free Software
```

```
Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA  02111-1307, USA.
```

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) {\it year} {\it name of author}
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details
type ‘show w’.  This is free software, and you are welcome
to redistribute it under certain conditions; type ‘show c’
for details.
```

The hypothetical commands ‘`show w`’ and ‘`show c`’ should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than ‘`show w`’ and ‘`show c`’; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a “copyright disclaimer” for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright
interest in the program ‘Gnomovision’
(which makes passes at compilers) written
by James Hacker.
{\it signature of Ty Coon}, 1 April 1989
Ty Coon, President of Vice
```

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License. Return to GNU’s home page.

FSF & GNU inquiries & questions to gnu@gnu.org. Other ways to contact the FSF.

Comments on these web pages to webmasters@www.gnu.org, send other questions to gnu@gnu.org.

Copyright notice above. Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111, USA

Updated: 3 Jan 2000 rms

GNU Lesser General Public License

image of a Philosophical GNU [English — Japanese]

- Why you shouldn't use the Lesser GPL for your next library
- What to do if you see a possible LGPL violation
- Translations of the LGPL
- The GNU Lesser General Public License as a text file
- The GNU Lesser General Public License as a Texinfo file

This GNU Lesser General Public License counts as the successor of the GNU Library General Public License. For an explanation of why this change was necessary, read the Why you shouldn't use the Lesser GPL for your next library article.

Table of Contents

- GNU LESSER GENERAL PUBLIC LICENSE
  - Preamble
  - TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION
  - How to Apply These Terms to Your New Libraries

GNU LESSER GENERAL PUBLIC LICENSE

Version 2.1, February 1999

```
Copyright (C) 1991, 1999 Free Software Foundation, Inc.
59 Temple Place, Suite 330, Boston, MA  02111-1307  USA
Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.
[This is the first released version of the Lesser GPL.  It also counts
 as the successor of the GNU Library Public License, version 2, hence
 the version number 2.1.]
```

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the

Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this

license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

TERMS AND CONDITIONS

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy

and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- **a)** The modified work must itself be a software library.
- **b)** You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- **c)** You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
- **d)** If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

  (For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

  These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

  Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

  In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so

that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables

containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

- **a)** Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)
- **b)** Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.
- **c)** Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- **d)** If distribution of the work is made by offering access to copy

from a designated place, offer equivalent access to copy the above specified materials from the same place.

- **e)** Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

- **a)** Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.
- **b)** Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are

prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries,

so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL

DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
{\it one line to give the library's name and an idea of what it does.}
Copyright (C) {\it year}  {\it name of author}
This library is free software; you can redistribute it and/or
modify it under the terms of the GNU Lesser General Public
License as published by the Free Software Foundation; either
version 2.1 of the License, or (at your option) any later version.
This library is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU
Lesser General Public License for more details.
You should have received a copy of the GNU Lesser General Public
License along with this library; if not, write to the Free Software
Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA  02111-1307  USA
```

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in
the library 'Frob' (a library for tweaking knobs) written
by James Random Hacker.
{\it signature of Ty Coon}, 1 April 1990
Ty Coon, President of Vice
```

That's all there is to it! Return to GNU's home page.

FSF & GNU inquiries & questions to gnu@gnu.org. Other ways to contact the FSF.

Comments on these web pages to webmasters@www.gnu.org, send other questions to gnu@gnu.org.

Updated: 27 Nov 2000 paulv

Bacula Projects

Please see the projects page on the web site at: www.bacula.org.projects.html, or see the **projects** file in the main source directory. For a current list of tasks you can see **kernstodo** in the main source directory.

Thanks

Thanks to Richard Stallman for starting the Free Software movement and for bringing us gcc and all the other GNU tools as well as the GPL license.

Thanks to Linus Torvalds for bring us Linux.

Thanks to all the Free Software programmers. Without being able to peek at your code, and in some cases, take parts of it, this project would have been much more difficult.

Thanks to John Walker for suggesting this project, giving it a name, contributing software he has written, and for his programming efforts on Bacula as well as having acted as a constant sounding board and source of ideas.

Thanks to the apcupsd project where I started my Free Software efforts, and from which I was able to borrow some ideas and code that I had written.

Special thanks to D. Scott Barninger for writing the bacula RPM spec file, building all the RPM files and loading them onto Source Forge. This has been a tremendous help.

Many thanks to Karl Cunningham for converting the manual from html format to LaTeX. It was a major effort flawlessly done that will benefit the Bacula users for many years to come. Thanks Karl.

Thanks to Dan Langille for the **incredible** amount of testing he did on FreeBSD. His perseverance is truly remarkable. Thanks also for the many contributions he has made to improve Bacula (pthreads patch for FreeBSD, improved start/stop script and addition of Bacula userid and group, stunnel, ...), his continuing support of Bacula users. He also wrote the PostgreSQL driver for Bacula and has been a big help in correcting the SQL.

Thanks to Phil Stracchino for writing the gnome-console **ConsoleFont** configuration command, all the suggestions he has made, and his continuing suppport of Bacula users.

Thanks to multiple other Bacula Packagers who make and release packages for different platforms for Bacula.

Thanks to Christopher Hull for developing the native Win32 Bacula emulation code and for contributing it to the Bacula project.

Thanks to Nicolas Boichat for writing wx-console and the bacula-tray-monitor. These are very nice GUI additions to Bacula.

Thanks to Nic Bellamy for providing the bacula-dir.conf file that he uses to implement daily tape rotation using multiple Pools.

Thanks to Johan Decock for providing numerous corrections to the manual.

Thanks to all the Bacula users, especially those of you who have contributed ideas, bug reports, patches, and new features.

The original variable expansion code used in the LabelFormat comes from the Open Source Software Project (www.ossp.org). It has been adapted and extended for use in Bacula.

For all those who I have left out, please send me a reminder, and in any case, thanks for your contribution.

Copyrights and Trademarks

Certain words and/or products are Copyrighted or Trademarked such as Windows (by Microsoft). Since they are numerous, and we are not necessarily aware of the details of each, we don't try to list them here. However, we acknowledge all such Copyrights and Trademarks, and if any copyright or trademark holder wishes a specific acknowledgment, notify us, and we will be happy to add it where appropriate.

Bacula Bugs

Well fortunately there are not too many bugs, but thanks to Dan Langille, we have a bugs database where bugs are reported. Generally, when a bug is fixed, a patch for the currently release version will be attached to the bug report.

The directory **patches** in the current CVS always contains a list of the patches that have been created for the previously released version of Bacula. In addition, the file **patches-version-number** in the **patches** directory contains a summary of each of the patches.

A "raw" list of the current task list and known issues can be found in **kernstodo** in the main Bacula source directory.

Variable Expansion

Please note that as of version 1.37, the Variable Expansion is deprecated and replaced by Python scripting (not yet documented).

Variable expansion is somewhat similar to Unix shell variable expansion. Currently (version 1.31), it is used only in format labels, but in the future, it will most likely be used in more places.

General Functionality

This is basically a string expansion capability that permits referencing variables, indexing arrays, conditional replacement of variables, case conversion, substring selection, regular expression matching and replacement, character class replacement, padding strings, repeated expansion in a user controlled loop, support of arithmetic expressions in the loop start, step and end conditions, and recursive expansion.

When using varaiable expansion characters in a Volume Label Format record, the format should always be enclosed in double quotes (**"**).

For example, **${HOME}** will be replaced by your home directory as defined in the environment. If you have defined the variable **xxx** to be **Test**, then the reference **${xxx:p/7/Y/r}** will right pad the contents of **xxx** to a length of seven characters filling with the character **Y** giving **YYYTest**.

Bacula Variables

Within Bacula, there are three main classes of variables with some minor variations within the classes. The classes are:

Counters Counters are defined by the **Counter** resources in the Director's conf file. The counter can either be a temporary counter that lasts for the duration of Bacula's execution, or it can be a variable that is stored in the catalog, and thus retains its value from one Bacula execution to another. Counter variables may be incremented by postfixing a plus sign (+ after the variable name).

Internal Variables Internal variables are read-only, and may be related to the current job (i.e. Job name), or may be special variables such as the date and time. The following variables are available:

Year – the full year

Month – the current month 1-12
Day – the day of the month 1-31
Hour – the hour 0-24
Minute – the current minute 0-59
Second – the current second 0-59
WeekDay – the current day of the week 0-6 with 0 being Sunday
Job – the job name
Dir – the Director's name
Level – the Job Level
Type – the Job type
JobId – the JobId
JobName – the unique job name composed of Job and date
Storage – the Storage daemon's name
Client – the Client's name
NumVols – the current number of Volumes in the Pool
Pool – the Pool name
Catalog – the Catalog name
MediaType – the Media Type

Environment Variables Environment variables are read-only, and must be defined in the environment prior to executing Bacula. Environment variables may be either scalar or an array, where the elements of the array are referenced by subscripting the variable name (e.g. **${Months[3]}**). Environment variable arrays are defined by separating the elements with a vertical bar (—), thus **set Months="Jan—Feb—Mar—Apr—..."** defines an environment variable named **Month** that will be treated as an array, and the reference **${Months[3]}** will yield **Mar**. The elements of the array can have differing lengths.

Full Syntax

Since the syntax is quite extensive, below, you will find the pseudo BNF. The special characters have the following meaning:

```
::=     definition
( )     grouping if the parens are not quoted
|       separates alternatives
'/'     literal / (or any other character)
CAPS    a character or character sequence
*       preceding item can be repeated zero or more times
?       preceding item can appear zero or one time
+       preceding item must appear one or more times
```

And the pseudo BNF describing the syntax is:

```
input       ::= ( TEXT
                | variable
                | INDEX_OPEN input INDEX_CLOSE (loop_limits)?
                )*
variable    ::= DELIM_INIT (name|expression)
name        ::= (NAME_CHARS)+
expression  ::= DELIM_OPEN
                (name|variable)+
                (INDEX_OPEN num_exp INDEX_CLOSE)?
                (':' command)*
                DELIM_CLOSE
command     ::= '-' (TEXT_EXP|variable)+
              | '+' (TEXT_EXP|variable)+
              | 'o' NUMBER ('-'|',') (NUMBER)?
              | '#'
              | '*' (TEXT_EXP|variable)+
              | 's' '/' (TEXT_PATTERN)+
                    '/' (variable|TEXT_SUBST)*
                    '/' ('m'|'g'|'i'|'t')*
              | 'y' '/' (variable|TEXT_SUBST)+
                    '/' (variable|TEXT_SUBST)*
                    '/'
              | 'p' '/' NUMBER
                    '/' (variable|TEXT_SUBST)*
                    '/' ('r'|'l'|'c')
              | '%' (name|variable)+
                    ('(' (TEXT_ARGS)? ')')?
              | 'l'
              | 'u'
num_exp     ::= operand
              | operand ('+'|'-'|'*'|'/'|'%') num_exp
operand     ::= ('+'|'-')? NUMBER
              | INDEX_MARK
              | '(' num_exp ')'
              | variable
loop_limits ::= DELIM_OPEN
                (num_exp)? ',' (num_exp)? (',' (num_exp)?)?
                DELIM_CLOSE
NUMBER      ::= ('0'|...|'9')+
TEXT_PATTERN::= (^('/'))+
TEXT_SUBST  ::= (^(DELIM_INIT|'/'))+
TEXT_ARGS   ::= (^(DELIM_INIT|')'))+
TEXT_EXP    ::= (^(DELIM_INIT|DELIM_CLOSE|':'|'+'))+
TEXT        ::= (^(DELIM_INIT|INDEX_OPEN|INDEX_CLOSE))+
DELIM_INIT  ::= '$'
DELIM_OPEN  ::= '{'
DELIM_CLOSE ::= '}'
INDEX_OPEN  ::= '['
INDEX_CLOSE ::= ']'
INDEX_MARK  ::= '#'
NAME_CHARS  ::= 'a'|...|'z'|'A'|...|'Z'|'0'|...|'9'
```

Semantics

The items listed in **command** above, which always follow a colon (**:**) have the following meanings:

```
-    perform substitution if variable is empty
+    perform substitution if variable is not empty
o    cut out substring of the variable value
#    length of the variable value
*    substitute empty string if the variable value is not empty,
     otherwise substitute the trailing parameter
s    regular expression search and replace. The trailing
     options are: m = multiline, i = case insensitive,
                  g = global,    t = plain text (no regexp)
y    transpose characters from class A to class B
p    pad variable to l = left, r = right or c = center,
     with second value.
%    special function call (none implemented)
l    lower case the variable value
u    upper case the variable value
```

The **loop_limits** are start, step, and end values.

A counter variable name followed immediately by a plus (+) will cause the counter to be incremented by one.

Examples

To create an ISO date:

```
DLT-${Year}-${Month:p/2/0/r}-${Day:p/2/0/r}
```

on 20 June 2003 would give **DLT-2003-06-20**

If you set the environment variable **mon** to

```
January|February|March|April|May|...
File-${mon[${Month}]}/${Day}/${Year}
```

on the first of March would give **File-March/1/2003**

Director Index

*Priority , 112
*WrapCounter , 127

a name , 83, 96
Accept Any Volume , 121
Address , 110, 112, 158
Admin , 92
always , 103
append , 154
Autochanger , 113
AutoPrune , 111, 119, 204

Backup , 92
Bootstrap , 96

Catalog , 95, 110, 123, 128
Catalog Files , 119
CatalogACL , 127
Cleaning Prefix , 122
Client , 97
Client (or FileDaemon) , 110
Client Run After Job , 102
Client Run Before Job , 101
ClientACL , 126
CommandACL , 127
count , 202
Counter , 127
Counters , 492

days , 84
DB Address , 124
DB Name , 124
DB Port , 124
DB Socket , 124
Description , 88
destination , 152

Device , 113
Differential , 93
Differential Backup Pool , 98
DifferentialPool , 106
dir , 200
DirAddress , 90
DirAddresses , 90
Director , 88
director , 154
directory , 83
DIRPort , 158
DIRport , 90
DiskToCatalog , 95
done , 202

Environment Variables , 493
estimate , 200

FD Connect Timeout , 89
FD Port , 110
file , 154
File Retention , 110, 203
FileSet , 97
FileSetACL , 127
find , 200
Full , 92
Full Backup Pool , 98
FullPool , 106

Heartbeat Interval , 132
hours , 83

ifnewer , 103
ifolder , 103
Incremental , 92
Incremental Backup Pool , 98

IncrementalPool , 106
InitCatalog , 94
integer , 83
Internal Variables , 492

Job , 91
Job Retention , 111, 204
JobACL , 126
JobDefs , 96

Label Format , 122
Level , 92, 106
long integer , 83

MailCommand , 153
mark , 201
Max Run Time , 98
Max Start Delay , 98
Max Wait Time , 99
Maximum , 127
Maximum Concurrent Jobs , 89, 103, 111, 114, 132
Maximum Volume Bytes , 118
Maximum Volume Files , 117
Maximum Volume Jobs , 117
Maximum Volumes , 116
Media Type , 113
Messages , 88, 97, 106, 153
Minimum , 127
minutes , 83
months , 84
mount , 156

Name , 88, 91, 105, 110, 112, 116, 123, 126, 127, 153
never , 103
Number of Volumes , 116

Password , 88, 110, 112, 126, 159
password , 83, 124
Pid Directory , 89, 131
Pool , 97, 106, 116
Pool Type , 117
PoolACL , 127
positive integer , 83
Prefix Links , 103
Priority , 104
Prune Files , 99
Prune Jobs , 99
Prune Volumes , 99
Purge Oldest Volume , 121
pwd , 201

quarters , 84
QueryFile , 89

Recycle , 120
Recycle Current Volume , 121
Recycle Oldest Volume , 120
Replace , 102
Rerun Failed Levels , 102
Reschedule Interval , 103
Reschedule On Error , 103
Reschedule Times , 104
Restore , 92
restored , 156
Run , 106
Run After Job , 100
Run Before Job , 99

Schedule , 98, 105
ScheduleACL , 127
SD Connect Timeout , 89
SD Port , 112
seconds , 83
setdebug , 179
sfdisk.gz , 455
Spool Attributes , 102
Spool Data , 102
SpoolData , 106
sqlquery , 180
status , 180
Storage , 98, 106, 112
StorageACL , 126

Type , 92

unmark , 201
Use Volume Once , 117
user , 124

Verify , 92

Verify Job , 96
Volume Retention , 119
Volume Use Duration , 118
VolumeToCatalog , 95

weeks , 84
Where , 102
Working Directory , 88
WritePartAfterJob , 107

years , 84
yes, 83

File Daemon Index

*Archive , 16
*JobLevel , 415
*JobType , 415
*Update , 16
*security , 156
-r <job> , 75
/about , 360
/events , 360
/help , 360
/install , 360
/kill , 360
/remove , 360
/run , 360
/service , 360
/servicehelper , 360
/status , 360
<destination> , 154

a name , 14, 15
Address , 164, 165
Administrator , 13
all , 156

Backup , 13
Bootstrap File , 13

Catalog , 13
Client , 14, 415
Client (or FileDaemon) , 131, 165
Console , 14
Count , 414

Daemon , 14
Debug , 154
Differential , 14
Directive , 14
Director , 14, 134, 164
DIRPort , 164

error , 155
exit , 202

fatal , 155
FD Port , 165
File Attributes , 14
File Daemon , 14
FileIndex , 415
format.hda , 454

help , 202

Incremental , 14
info , 155

Job , 414
JobId , 414

lsmark , 201

mail , 155
mail on error , 155
Monitor , 15, 134, 163
mount_drives , 454

Name , 131, 134, 163–165
name , 82
name-string , 82
notsaved , 156

operator , 155
OperatorCommand , 154

partition.hda , 454
Password , 134, 163, 165

quit , 202

Recycle , 214
Refresh Interval , 164
Resource , 15
Restore , 15
restore_bacula , 455
Retention Period , 16

saved , 155
Schedule , 15
SD Port , 166
Service , 15
sfdisk , 455
skipped , 156
Slot , 415
start_network , 455
stderr , 154
stdout , 154
Storage , 165
Storage Coordinates , 15
Storage Daemon , 15
Stream , 415
string , 82
syslog , 155

terminate , 155

VolBlock , 414
VolFile , 414
VolSessionId , 414
VolSessionTime , 414
Volume , 414

warning , 155
Working Directory , 131

Storage Daemon Index

-c <file> , 74
-d nn , 74
/etc/bacula/mtx-changer /dev/sg0 list , 261
/etc/bacula/mtx-changer /dev/sg0 load , 261
/etc/bacula/mtx-changer /dev/sg0 loaded , 261
/etc/bacula/mtx-changer /dev/sg0 slots , 261
/etc/bacula/mtx-changer /dev/sg0 unload , 262
/etc/bacula/mtx-changer /dev/sg0 unloaded , 262

Alert Command , 142
Always Open , 143
Archive Device , 140
Autochanger , 141, 255
Automatic mount , 149

Backward Space File , 147
Backward Space Record , 147
Block Positioning , 148
BSF at EOM , 146

Changer Command , 142, 256
Changer Device , 141, 255
Close on Poll , 144

Drive Index , 143, 256

Fast Forward Space File , 146
Forward Space File , 147
Forward Space Record , 147

Hardware End of Medium , 145
Heartbeat Interval , 137

Label media , 148

Make sure Bacula is not running. , 261
Maximum block size , 145
Maximum Changer Wait , 143, 256
Maximum Concurrent Jobs , 137
Maximum File Size , 148
Maximum Job Spool Size , 148
Maximum Network Buffer Size , 148
Maximum Open Wait , 144
Maximum Spool Size , 148
Maximum Volume Size , 147
Media Type , 141
Minimum block size , 144
Monitor , 139
mount storage , 72

Name , 136, 139, 140

Offline On Unmount , 147

Password , 139, 166
Pid Directory , 137

quit , 72

Random access , 144
Removable media , 144

Scan , 17
SDAddress , 138
SDAddresses , 137
SDPort , 138
Session , 15
Spool Directory , 148

TWO EOF , 147

Verify , 16
Volume , 17
Volume Poll Interval , 144

Working Directory , 136

Console Index

<destination> , 154

add [pool , 170
anything , 185
autodisplay on/off , 170
automount on/off , 170
AutoPrune , 213

cancel [jobid , 170
console , 154
ConsoleFont , 159
create [pool , 171

delete , 171
Director , 158

estimate , 172
exit , 183

FDAddress , 133
FDAddresses , 132
FDPort , 133
Font , 159

help , 172

label , 172
list , 174
list files jobid , 72
list jobid , 72
list jobmedia , 72
list jobs , 72
list jobtotals , 72
list media , 72
list pools , 72
llist , 175

Maximum Network Buffer Size , 133
messages , 72, 176
mount , 176

Name , 158, 159

prune , 177
purge , 177

query , 183
quit , 183

relabel , 177
release , 178
restore , 178
run , 178

SDConnectTimeout , 133
show , 180
status , 72
status dir , 72
status jobid , 72

unmount , 181
unmount storage , 72
update , 181
use , 182

var name , 183
version , 183
Volume Retention , 213

wait , 183

General Index

MAJOR WARNING , 270

Above
    Bacula Configuration Files for the , 399
Actual Conf Files , 239
Adding a Second Client , 67
Adding Volumes to a Pool , 186
Additional Resources , 380
Advantages , 243, 245
Advantages of Bacula Over Other Backup Programs , 21
After bscan , 314
Algorithm
    Recycling , 215
All my Jobs are scheduled for the same time. Will this cause problems?
, 277
Alphabetic List of Console Commands , 170
Alternate Disaster Recovery Suggestion for Win32 Systems , 379
Answers , 410
Arguments
    Command Line , 195
Attributes
    Restoring Directory , 197
Authorization
    Names Passwords and , 85
Auto Starting the Daemons , 48
Autochanger
    Automatic Labeling Using Your , 295
    Simulating Barcodes in your , 259
    Using the , 263
Autochangers
    Supported , 28, 267
Autochangers – General , 252
Autochangers Support , 252
Automatic Generation of Bootstrap Files , 418
Automatic Labeling Using Your Autochanger , 295
Automatic Pruning , 212
Automatic Pruning and Recycling Example , 22
Automatic Volume Labeling , 226
Automatic Volume Recycling , 211
Aware
    FreeBSD Users Be , 28

Backing up ACLs on ext3 or XFS filesystems , 2
Backing Up Portables Using DHCP , 296
Backing Up the WinNT/XP/2K System State ,
Backing Up Third Party Databases , 208
Backing up to Multiple Disks , 231
Backing Up Your Bacula Database , 207
Backup
    Simple One Tape , 243
Backup Strategies , 243
Bacula
    Before Running , 55
    Disaster Recovery Using , 361
    Installing , 36, 47
    Running , 35
    Upgrading , 36
    What is , 8
    Who Needs , 8
Bacula Autochanger Interface , 265
Bacula Bugs , 491
Bacula Cannot Open the Device , 333
Bacula Components or Services , 8
Bacula Configuration , 11
Bacula Configuration Files for the Above , 399
Bacula Console , 168
Bacula Console Restore Command , 188
Bacula Copyright, Trademark, and Licenses , 46
Bacula Frequently Asked Questions , 272

Bacula internal database is no longer supported, please convert , 448
Bacula is Asking for a New Tape After 2 GB of Data but My Tape holds 33 GB. Why? , 280
Bacula is Not Doing the Right Thing When I Request a Full Backup. Why? , 280
Bacula Projects , 488
Bacula Runs Fine but Cannot Access a Client on a Different Machine. Why? , 274
Bacula Saves But Cannot Restore Files , 332
Bacula Security Issues , 392
Bacula Variables , 492
Bacula® - RPM Packaging FAQ , 410
Barcode Support , 265
Bare Metal Floppy Recovery on Linux with a Bacula Floppy Rescue Disk , 450
Bare Metal Recovery on Linux with a Bacula Rescue CDROM , 362
Basic Volume Management , 224
Bcopy , 315
Bcopy Command Options , 315
Before Running Bacula , 55
Bextract , 308
bimagemgr
    Installation , 324
    Usage , 324
Bimagemgr , 323
bimagemgr Installation , 324
bimagemgr Usage , 324
bls
    Listing Blocks , 306
    Listing Jobs , 306
bls , 304
Boot with your Bacula Rescue CDROM , 369
Boot with your Emergency Floppy , 458
Bootstrap File , 413
Brief Tutorial , 55
Bscan
    After , 314
bscan , 310
Bsmtp , 318
Btape , 315, 329
Btape Commands , 316
Bugs
    Bacula , 491
    Linux Problems or , 374, 463
    Bugs , 464
    Bugs and Other Considerations , 377
Building a File Daemon or Client , 48
Building Bacula from Source , 38

Can Bacula Backup and Restore Files Greater than 2 Gi
Size? , 277
Can Bacula Backup My System To Files instead of Tape?
Capabilities , 148
Case
    How to Excude File on Windows Regardless of , 297
Catalog
    Internal Bacula , 421
    Using bscan to Compare a Volume to an existing , 31
Catalog Maintenance , 203
Catalog Resource , 123
Catalog Services , 420
CDROM
    Bare Metal Recovery on Linux with a Bacula Rescue
    Boot with your Bacula Rescue , 369
    Creating a Bacula Rescue , 364
Certificate
    Creating a Self-signed , 390
    Getting a CA Signed , 391
Certificates , 383
Channel
    config Files for stunnel to Encrypt the Control , 388
    config Files for stunnel to Encrypt the Data , 385
    Encrypting the Control , 387
    Modification of bacula-dir.conf for the Control , 387
    Modification of bacula-dir.conf for the Data , 384
    Securing the Data , 384
    Starting and Testing the Control , 389
Checking and Setting Tape Hardware Compression and Bl
, 337
Client
    Adding a Second , 67
    Building a File Daemon or , 48
    Command Line Options Specific to the Bacula Wi
Daemon , 359
    Using stunnel to Encrypt to a Second , 389
Client Resource , 110, 131, 164
Client/File daemon Configuration , 131

Clients
Considerations for Multiple , 232
Using Bacula to Encrypt Communications between , 382
Command
Bacula Console Restore , 188
Full Form of the Update Slots , 260
Restore , 188
Command Line Arguments , 195
Command Line Options Specific to the Bacula Console (Client) , 359
Commands
Alphabetic List of Console , 170
btape , 316
Console , 161
File Selection , 200
Other Useful Console , 72
, 184
Special dot , 184
Comments , 81
Communications
Using ssh to Secure the , 391
Communications Ports Used , 382
Compacting Your MySQL Database , 204
Compacting Your PostgreSQL Database , 206
Compacting Your SQLite Database , 206
Completion
Getting Notified of Job , 284
Concrete Example , 398
CONDITIONS
TERMS AND , 469, 479
Config Files for stunnel to Encrypt the Control Channel , 388
Config Files for stunnel to Encrypt the Data Channel , 385
Configuration
Bacula , 11
Client/File daemon , 131
Console , 158
Monitor , 163
Storage Daemon , 136
Configuration , 168
Configure Options , 43
Configuring and Testing TCP Wrappers with Bacula , 393
Configuring the Console Program , 30
Configuring the Director , 32, 87
Configuring the File daemon , 32
Configuring the Monitor Program , 31
Configuring the Storage daemon , 33
Consoles , 382
Bugs and Other , 377
Important , 361, 449
Security , 294
Windows Compatibility , 352
Considerations for Multiple Clients , 232
Contains Windows File Daemon
Bacula , 168
Console Commands , 161
Console Configuration , 158
Console Resource , 125, 159
ConsoleFont Resource , 159
Contents
Table of , 468, 476
Conventions Used in this Document , 12
Converting from MySQL to PostgreSQL , 442
Copyrights and Trademarks , 490
Count
Using bscan to Correct the Volume File , 31
Counter Resource , 127
Crash
Rejected Volumes After a , 290
Creating a Bacula Rescue CDROM , 364
Creating a Bacula Rescue Disk , 453
Creating a Pool , 75
Creating a Self-signed Certificate , 390
Creating an Emergency Boot Disk , 452
Creating Holiday Schedules , 294
Credits , 444
Critical Items , 53
Critical Items to Implement Before Going Produ
Current Implementation Restrictions , 22
Current State of Bacula , 19
Customizing the Configuration Files , 79

Daemon
Configuring the File , 32
Configuring the Storage , 33
Daemon Information for each , 86
Daemon Command Line Options , 74
Daemons
Auto Starting the , 48
Data Encryption , 393

Starting the , 56
Daily Tape Rotation , 245
Daily, Weekly, Monthly Tape Usage Example, 218
Data Spooling , 269
Data Spooling Directives , 269
Database
Backing Up Your Bacula , 207
Compacting Your MySQL , 204
Compacting Your PostgreSQL , 206
Compacting Your SQLite , 206
Internal Bacula , 448
Re-initializing the Catalog , 437, 442, 446
Repairing Your MySQL , 205
Repairing Your PostgreSQL , 205
Starting the , 56
Database Size , 209
Database Table Design , 421
Database Tables , 423
Databases
Backing Up Third Party , 208
Dbcheck , 319
Dealing with Firewalls , 397
Dealing with Multiple Magazines , 258
Dealing with Win32 Problems , 350
Debug Daemon Output , 73
Debugger
Manually Running Bacula Under The , 345
Definition
MySQL Table , 431
Dependency Packages , 37
Design
Database Table , 421
Overall , 237
Design Limitations or Restrictions , 23
Detailed Information for each Daemon , 86
Details
Practical , 243, 246
Technical , 397
Details , 403
Determining What Tape Drives and Autochangers You Have on FreeBSD , 339
Device
Bacula Cannot Open the , 333
Device Configuration Records , 255
Device Resource , 139
Devices
Multiple , 254
devices
SCSI, 253
DHCP
Backing Up Portables Using , 296
Differential Pool , 238
Difficulties Connecting from the FD to the SD , 74
Directives
Data Spooling , 269
Pruning , 212
Director
Configuring the , 32, 87
Director Resource , 88, 134, 139, 158, 164
Director Resource Types , 87
Directories , 363, 451
Directory
Get Rid of the /lib/tls , 34
Disadvantages , 243, 245
Disaster
Preparing Solaris Before a , 376
Disaster Recovery , 35
Disaster Recovery of Win32 Systems , 378
Disaster Recovery Using a Bacula Rescue Floppy , 449
Disaster Recovery Using Bacula , 361
Disclaimer , 467
Disk
Bare Metal Floppy Recovery on Linux with a Bac
Rescue , 450
Creating a Bacula Rescue , 453
Creating an Emergency Boot , 452
Putting Two or More Systems on Your Rescue , 366
Disks
Backing up to Multiple , 231
Document
Conventions Used in this , 12
Does Bacula support Windows? , 272
Domain
Publicly , 466
Drive
Testing Bacula Compatibility with Your Tape , 34
Using btape to Verify your Tape , 316, 330
Drives

Supported Tape , 27
Unsupported Tape , 28

Encrypting the Control Channel , 387
Encryption
Starting and Testing the Data , 386
Encryption , 382
Ensuring that the Tape Modes Are Properly Set – Linux Only , 335
Example
Automatic Pruning and Recycling , 220
Concrete , 398
Daily Weekly Monthly Tape Usage , 218
Final , 418
Verify Configuration , 408
Example , 229
Example Client Configuration File , 135
Example Configuration File , 257
Example Director Configuration File , 128
Example Restore Job Resource , 199
Example Scripts , 253
Examples , 284, 495
Executing Scripts on a Remote Machine , 297
Expansion
Variable , 492
Extracting From Multiple Volumes , 309
Extracting With a Bootstrap File , 309
Extracting with Include or Exclude Lists , 308

FAQ
Bacula® - RPM Packaging , 410
File
Bootstrap , 413
Example Client Configuration , 135
Example Configuration , 257
Example Director Configuration , 128
Extracting With a Bootstrap , 309
Maintaining a Valid Bootstrap , 288
Sample Console Configuration , 161
Sample Storage Daemon Configuration , 149
Specifying a Device Name For a , 303, 329
Specifying the Configuration , 328
File Format , 413
File Selection Commands , 200
Filename
Selecting Files by , 194
Filenames and Maximum Filename Length , 420
Files
Actual Conf , 239
Automatic Generation of Bootstrap , 418
Bacula Saves But Cannot Restore , 332
Customizing the Configuration , 79
Including other Configuration , 82
Modifying the Bacula Configuration , 52
Problems Restoring , 198
Restoring Your , 64
Setting Up Bacula Configuration , 30
Testing your Configuration , 33
Filesystems
Backing up ACLs on ext3 or XFS , 299
Fills
When The Tape , 69
Final Example , 418
Firewalls
Dealing with , 397
Windows , 354
Floppy
Boot with your Emergency , 458
Disaster Recovery Using a Bacula Rescue ,
Restoring Your Linux Client with a , 457
Format
File , 413
Resource Directive , 80
Found
What To Do When Differences Are , 406
FreeBSD
Determining What Tape Drives and Autoch
339
Tape Modes on , 338
FreeBSD , 46
FreeBSD Bare Metal Recovery , 374
FreeBSD Issues , 261
FreeBSD Users Be Aware , 28
Full Form of the Update Slots Command , 260
Full Pool , 238
Full Syntax , 493
Functionality
General , 492

General
Autochangers – , 252
General , 36, 53, 131, 136, 158, 163, 168, 188, 347, 361, 420, 449
General Functionality , 492
Get Rid of the /lib/tls Directory , 34
Getting a CA Signed Certificate , 391
Getting A Traceback On Other Systems , 354
Getting Debug Output from Bacula , 346
Getting Email Notification to Work , 286
Getting Notified of Job Completion , 284
Getting Notified that Bacula is Running , 286
Getting Started with Bacula , 29
GNOME , 51
GNU GENERAL PUBLIC LICENSE , 468
GNU General Public License , 468
GNU LESSER GENERAL PUBLIC LICENSE , 476
GNU Lesser General Public License , 476
Going on Vacation , 296
GPL , 466

Handling
Total Automation of Bacula Tape , 209
Have
Knowing What SCSI Devices You , 253
Have Patience When Starting the Daemons or Mounting Blank Tapes , 73
How Can I Be Sure that Bacula Really Saves and Restores All Files? , 277
How Can You Claim to Handle Unlimited Filename Lengths when All Other Programs Have Fixed Limits? , 278
How Does It Work? , 401
How to Apply These Terms to Your New Libraries , 486
How to Apply These Terms to Your New Programs , 474
How to Excude File on Windows Regardless of Case , 297

I am Backing Up an Offsite Machine with Intermittent Connection. The Director Waits Forever for the Client to Contact It. What Can I Do. , 281
I Am Not Getting Email Notification, What Can I Do? , 279
I Cannot Get My Windows Client to Start Automatically , 275
I Change Recycling, Retention Periods, or File Sizes in my Pool Resource and they Still Dont Work. , 279
I did a Full backup last week, but now in running an Incremental, Bacula says it did not find a FULL backup. Why? , 278
I didn't realize that the backups were not working on m
Client, and now I have lots of data. What can I do? , 276
I Have Configured Compression On, But None of My Compressed. Why? , 279
I Run a Restore Job and Bacula Hangs. What do I do? ,
I Started A Job then Decided I Really Did Not Want to there a better way than ./bacula stop to stop it? , 277
I'm confused by the different Retention periods: File Ret Retention, Volume Retention. Why are there so many? ,
I'm Getting Authorization Errors. What is Going On? , 2
If I Do Run Multiple Simultaneous Jobs, How Can I Particular Job to Run After Another Job? , 279
Implemented
What is , 19
Important Considerations , 361, 449
Important Note , 402
In connecting to my Client, I get htmlQuoteERR:Connecti Packet Size too big from File daemon:192.168.1.4:9102 Why? , 282
Including other Configuration Files , 82
Incorrect File Number , 334
Incorrect Number of Blocks or Positioning Errors during bt , 334
Incremental Backup , 289
Installation , 347
Installed
Windows Systems with CYGWIN , 46
Installing and Configuring MySQL , 421, 435
Installing and Configuring MySQL – Phase I , 435
Installing and Configuring MySQL – Phase II , 436
Installing and Configuring PostgreSQL , 421, 440
Installing and Configuring PostgreSQL – Phase I , 440
Installing and Configuring PostgreSQL – Phase II , 440
Installing and Configuring SQLite , 421, 445
Installing and Configuring SQLite – Phase I , 445
Installing and Configuring SQLite – Phase II , 445
Installing Bacula , 36, 47
Installing Bat , 279
Interacting with the Director to Query or Start Jobs , 57
Interactions between the Bacula Services , 18
Interface
Bacula Administration , 265
Internal Bacula Database , 421

Internal Bacula Database , 448
Is Bacula Stable? , 272
Issues
    Bacula Security , 392
    FreeBSD , 261
Items
    Critical , 53
    Recommended , 54

Job
    Running a , 58
    Sequence of Creation of Records for a Save , 422
Job Resource , 91
JobDefs Resource , 105
Jobs
    Interacting with the Director to Query or Start , 57
    Running Concurrent , 301

Kaboom
    What To Do When Bacula Crashes , 343
KDE , 51
Kern's Configure Script , 47
Key Concepts and Resource Records , 224
Knowing What SCSI Devices You Have , 253

Labeling
    Automatic Volume , 226
    Specifying Slots When , 257
Labeling Volumes with the Console Program , 77
Labeling Your Volumes , 76
Labels
    Understanding Pools Volumes and , 29
Length
    Filenames and Maximum Filename , 420
LGPL , 466
Libraries
    How to Apply These Terms to Your New , 488
LICENSE
    GNU GENERAL PUBLIC , 468
    GNU LESSER GENERAL PUBLIC , 476
License
    GNU General Public , 468
    GNU Lesser General Public , 476
Licenses
    Bacula Copyright Trademark and , 466
Linking Bacula with MySQL , 438
Linking Bacula with SQLite , 446
Linux Problems or Bugs , 374, 463
Linux SCSI Tricks , 332
Listing Blocks with bls , 306
Listing Jobs with bls , 306
Lists
    Extracting with Include or Exclude , 308
Log Rotation , 35

Machine
    Executing Scripts on a Remote , 297
Magazines
    Dealing with Multiple , 258
Maintaining a Valid Bootstrap File , 288
Maintenance
    Catalog , 203
Making Bacula Use a Single Tape , 217
Management
    Basic Volume , 224
Managers
    Other window , 51
Manually Changing Tapes , 244
Manually Recycling Volumes , 222
Manually resetting the Permissions , 356
Manually Running Bacula Under The Debugger
Message Resource , 135
Messages Resource , 125, 149, 152
Migrating from SQLite to MySQL , 207
Models
    Supported Autochanger , 267
Modes
    Tape Blocking , 342
Modification of bacula-dir.conf for the Control
Modification of bacula-dir.conf for the Data Cha
Modifying the Bacula Configuration Files , 52
Monitor
    Installing Tray , 51
Monitor Configuration , 163
Monitor Resource , 163
Multiple Devices , 254
My Catalog is Full of Test Runs, How Can I Sta
My Windows Client Immediately Dies When I S
MySQL
    Configure , 436

Installing and Configuring , 421, 435
Linking Bacula with , 438
Migrating from SQLite to , 207
MySQL Table Definition , 431

Names, Passwords and Authorization , 85
Not
What Bacula is , 17
Note
Important , 402
Notes
Other Make , 49
Number
Incorrect File , 334

On what machines does Bacula run? , 272
Only
Ensuring that the Tape Modes Are Properly Set – Linux , 335
Options
bcopy Command , 315
Configure , 43
Daemon Command Line , 74
Other Make Notes , 49
Other Points , 270
Other Programs , 318
Other Useful Console Commands , 72
Other window managers , 51
Output
Debug Daemon , 73
Overall Design , 237

Packages
Dependency , 37
Periods
Setting Retention , 203
Permissions
Manually resetting the , 356
Phase I
Installing and Configuring MySQL – Phase I , 435
Installing and Configuring PostgreSQL – Phase I , 440
Installing and Configuring SQLite – Phase I , 445
Phase II
Installing and Configuring MySQL – Phase II , 436
Installing and Configuring PostgreSQL – Phase II , 441
Installing and Configuring SQLite – Phase II , 445
Picture , 383
Points
Other , 270
Pool
Adding Volumes to a , 186
Creating a , 75
Differential , 238
Full , 238
Incremental , 239
Pool Options to Limit the Volume Usage , 225
Pool Resource , 115
Post Win32 Installation , 350
PostgreSQL
Converting from MySQL to , 442
Installing and Configuring , 421, 440
Practical Details , 243, 246
Preamble , 468 , 476
Preparation for a Bare Metal Recovery , 363, 451
Preparing Solaris Before a Disaster , 376
Problem , 236
Problems
Tips for Resolving , 332
Windows Ownership and Permissions , 355
Problems Restoring Files , 198
Production
Critical Items to Implement Before Going , 53
Program
Configuring the Console , 30
Configuring the Monitor , 31
Labeling Volumes with the Console , 77
Quitting the Console , 67
Running the Console , 168
Stopping the Console , 169
Programs
Advantages of Bacula Over Other Backup , 21
How to Apply These Terms to Your New , 474
Other , 318
Projects , 435, 440, 445
Pruning
Automatic , 212
Pruning Directives , 212
Public Domain , 466
Putting Two or More Systems on Your Rescue Disk , 366

Questions
Bacula Frequently Asked , 272
Quick Start , 13, 43
Quitting the Console Program , 67

Re-initializing the Catalog Database , 437, 442, 446
Recognized Primitive Data Types , 82
Recommended Items , 54
Recommended Options for most Systems , 44
Record
Sample Director's Console , 167
Sample File daemon's Director , 166
Sample Storage daemon's Director , 167
Records
Device Configuration , 255
Key Concepts and Resource , 224
Sample Monitor configuration file and related daemons' configuration , 166
Recovering Files Written to Tape With Fixed Block Sizes , 341
Recovery
Disaster , 35
FreeBSD Bare Metal , 374
Preparation for a Bare Metal , 363, 451
Solaris Bare Metal , 376
Windows Disaster , 355
Recycle Status , 216
Recycling
Automatic Volume , 211
Restricting the Number of Volumes and , 227
Recycling Algorithm , 215
Recycling All Your Volumes , 298
RedHat , 44
Rejected Volumes After a Crash , 290
Repairing Your MySQL Database , 205
Repairing Your PostgreSQL Database , 205
Requirements
System , 24
Requirements , 363
Resetting Directory and File Ownership and Permissions on Win32 Systems , 379
Resource
Catalog , 123
Client , 110, 131, 164
Console , 125, 159
ConsoleFont , 159
Counter , 127
Device , 139
Director , 88, 134, 139, 158, 164
Example Restore Job , 199
Job , 91
JobDefs , 105
Message , 135
Messages , 125, 149, 152
Monitor , 163
Pool , 115
Schedule , 105
Storage , 112, 136, 165
Resource Directive Format , 80
Resource Types , 84
Resources
Additional , 380
Restore Command , 188
Restoring a Client System , 368
Restoring a Server , 373
Restoring Directory Attributes , 197
Restoring Files Can Be Slow , 198
Restoring on Windows , 197
Restoring to a Running System , 380
Restoring Your Files , 64
Restoring Your Linux Client with a Floppy , 457
Restricting the Number of Volumes and Recycli
Restrictions
Current Implementation , 22
Design Limitations or , 23
Restrictions , 451
Rotation
Daily Tape , 245
Log , 35
Running
Getting Notified that Bacula is , 286
Running a Job , 58
Running as non-root , 395
Running Bacula , 35
Running Concurrent Jobs , 301
Running the Console Program , 168
Running the Console Program from a Shell Scri
Running the Verify , 405

Sample Console Configuration File , 161
Sample Director's Console record. , 167
Sample File daemon's Director record. , 166
Sample Monitor configuration file and related daemons' configuration records. , 166
Sample Storage Daemon Configuration File , 149
Sample Storage daemon's Director record , 167
Schedule Resource , 105
Schedules
    Creating Holiday , 294
    Technical Notes on , 109
Script
    Kern's Configure , 47
    Running the Console Program from a Shell , 185
    Testing the Autochanger and Adapting mtx-changer , 261
Scripts
    Example , 253
SCSI devices, 253
SD
    Difficulties Connecting from the FD to the , 74
Securing the Data Channel , 384
Security
    Using Bacula to Improve Computer , 403
Security Considerations , 294
Selecting Files by Filename , 194
Semantics , 495
Sequence of Creation of Records for a Save Job , 422
Server
    Restoring a , 373
Services
    Bacula Components or , 8
    Catalog , 420
    Interactions Between the Bacula , 18
Setting Retention Periods , 203
Setting Up Bacula Configuration Files , 30
Shutting down Windows Systems , 360
Simple One Tape Backup , 243
Simulating Barcodes in your Autochanger , 259
Size
    Checking and Setting Tape Hardware Compression , 337
    Database , 209
Sizes
    Recovering Files Written to Tape With Fixed Block , 341
Slots , 254
Slow
    Restoring Files Can Be , 198
Solaris , 45
Solaris Bare Metal Recovery , 376
Solution , 236
Source
    Building Bacula from , 38
Spaces
    Upper and Lower Case and , 81
) Commands , 184
Special dot Commands , 184
Specifications
    Windows Considerations for Filename , 359
Specifying a Device Name For a File , 303, 329
Specifying a Device Name For a Tape , 303, 329
Specifying Slots When Labeling , 257
Specifying the Configuration File , 303, 328
Specifying Volumes , 303
Spooling , 74
    Data , 269
SQLite
    Installing and Configuring , 421, 445
    Linking Bacula with , 446
    Testing , 446
Start
    Quick , 42, 43
Starting and Testing the Control Channel , 389
Starting and Testing the Data Encryption , 386
Starting the Daemons , 56
Starting the Database , 56
State
    Backing Up the WinNT/XP/2K System , 359
Status
    Recycle , 216
Steps to Take Before Disaster Strikes , 361, 450
Stopping the Console Program , 169
Storage Daemon Configuration , 136
Storage Resource , 112, 136, 165
Strategies
    Backup , 243
Strikes
    Steps to Take Before Disaster , 361, 450
Suggested Block , 341

Tips and , 284
Summary of Steps to Take to Get Your Tape Drive Working , 327
Support
Autochangers , 252
Barcode , 265
Supported Autochanger Models , 267
Supported Autochangers , 28, 267
Supported Operating Systems , 26, 38
Supported Tape Drives , 27
Syntax
Full , 493
System
Restoring a Client , 368
Restoring to a Running , 380
System Requirements , 24
Systems
Alternate Disaster Recovery Suggestion for Win32 , 379
Disaster Recovery of Win32 , 378
Getting A Traceback On Other , 344
Recommended Options for most , 44
Resetting Directory and File Ownership and Permissions on Win32 , 379
Shutting down Windows , 360
Supported Operating , 26, 38
Using the OnStream driver on Linux , 340

Table of Contents , 468, 476
Tables
Database , 423
Tape
Making Bacula Use a Single , 217
Specifying a Device Name For a , 303, 329
Using btape to Simulate Bacula Filling a , 341
Tape Blocking Modes , 342
Tape Modes on FreeBSD , 338
Tapes
Have Patience When Starting the Daemons or Mounting Blank , 73
Manually Changing , 244
Technical Details , 397
Technical Notes on Schedules , 109
Terminology , 13
TERMS AND CONDITIONS , 469, 479
Testfind , 322
Testing
Incorrect Number of Blocks or Positioning , 327, 334
Testing Bacula Compatibility with Your Tape D
Testing SQLite , 446
Testing the Autochanger and Adapting Your mtx
Testing The Traceback , 344
Testing your Configuration Files , 33
Testing Your Tape Drive With Bacula , 327
Thanks , 489
Tips and Suggestions , 284
Tips for Resolving Problems , 332
Tomsrtbt , 464
Tools
Volume Utility , 303
Total Automation of Bacula Tape Handling , 29
Traceback
Testing The , 344
Traceback , 343
Trademark , 466
Trademarks
Copyrights and , 490
Tricks
Linux SCSI , 332
Tutorial
Brief , 55
Types
Director Resource , 87
Recognized Primitive Data , 82
Resource , 84

Understanding Pools, Volumes and Labels , 29
Uninstalling Bacula on Win32 , 350
Unsupported Tape Drives , 28
Upgrading Bacula , 36
Upgrading Bacula Versions , 284
Upper and Lower Case and Spaces , 81
Usage
Pool Options to Limit the Volume , 225
Windows Port , 354
Use
What Database to , 42
Use it

Bacula internal database is no longer supported , 448
Used
Communications Ports , 382
Using Bacula to Encrypt Communications to Clients , 382
Using Bacula to Improve Computer Security , 403
Using bscan to Compare a Volume to an existing Catalog , 312
Using bscan to Correct the Volume File Count , 314
Using bscan to Recreate a Catalog from a Volume , 312
Using btape to Simulate Bacula Filling a Tape , 345
Using btape to Verify your Tape Drive , 316, 330
Using Pools to Manage Volumes , 236
Using ssh to Secure the Communications , 391
Using stunnel to Encrypt to a Second Client , 389
Using the Autochanger , 263
Using the OnStream driver on Linux Systems , 340

Vacation
Going on , 296
Variable Expansion , 492
Variables
Bacula , 492
Verify
Running the , 405
Verify Configuration Example , 408
Versions
Upgrading Bacula , 284
Volume
Using bscan to Recreate a Catalog from a , 312
Volume Utility Tools , 303
Volumes
Extracting From Multiple , 309
Labeling Your , 76
Manually Recycling , 222
Recycling All Your , 298
Specifying , 303
Using Pools to Manage , 236

WARNING
MAJOR , 270
What Bacula is Not , 17
What Database to Use? , 42
What is Bacula? , 8, 272
What is Implemented , 19
What Is the Really Unique Feature of Bacula? , 278
What language is Bacula written in? , 272
What To Do When Bacula Crashes (Kaboom) , 343
What To Do When Differences Are Found , 406
When I boot my machine and start Bacula then attempt backups, I get: ... Connection refused. , 281
When I Start the Console, the Error Messages Fly By. , 276
Where, 314
Where Is My Tape , 342
When Tape Fills , 69
Who Needs Bacula? , 8
Why Does Bacula Ignore the MaxVolumeSize Set in my P
Why have You Trademarked the Name Bacula®? , 277
Why is Your Online Document for Version 1.35 of Bacul
Why Release Version is 1.34? , 277
Win32
Dealing with Problems , 350
Installation , 347
Post Installation , 350
Uninstalling Bacula , 350
Win32 , 46
Windows
Restoring on , 197
Windows Compatibility Considerations , 352
Windows Considerations for Filename Specifications , 359
Windows Disaster Recovery , 355
Windows Firewalls , 354
Windows Ownership and Permissions Problems , 355
Windows Port Usage , 354
Windows Systems with CYGWIN Installed , 46
Windows Version of Bacula , 347
Work
Getting Email Notification to , 286
How Does It , 401
Working
Summary of Steps to Take to Get Your Tape Drive ,