



The Leading Open Source Backup Solution

Bacula[®] Developer's Guide

Kern Sibbald


May 2, 2023

This manual documents Bacula Community Edition 13.0.3 (02 May 2023)

Copyright © 1999-2023, Kern Sibbald

Bacula[®] is a registered trademark of Kern Sibbald.

This Bacula documentation by Kern Sibbald with contributions from many others, a complete list can be found in the License chapter. Creative Commons Attribution-ShareAlike 4.0 International License <http://creativecommons.org/licenses/by-sa/4.0/>







Contents

1	Bacula Developer Notes	1
1.1	Contributions	1
1.2	Patches	1
1.3	Copyrights	2
1.4	Copyright Assignment – Fiduciary License Agreement	2
1.5	The Development Cycle	3
1.5.1	Feature Requests	3
1.6	Bacula Code Submissions and Projects	5
1.7	Patches for Released Versions	5
1.8	Developing Bacula	6
1.9	Debugging	7
1.10	Using a Debugger	7
1.11	Memory Leaks	7
1.12	Special Files	8
1.13	When Implementing Incomplete Code	8
1.13.1	Bacula Source File Structure	8
1.14	Header Files	9
1.15	Programming Standards	9
1.16	Do Not Use	10
1.17	Avoid if Possible	10
1.18	Do Use Whenever Possible	10
1.19	Indenting Standards	11
1.20	Naming Convention	12
1.21	Locks and Threads	13
1.22	Tabbing	14



1.23	Don'ts	14
1.24	Message Classes	15
1.25	Debug Messages	15
1.25.1	Debug Tags	16
1.26	Error Messages	16
1.27	Job Messages	17
1.28	Queued Job Messages	17
1.29	Memory Messages	17
1.30	Bugs Database	18
2	Bacula Git Usage	19
2.1	Bacula Git repositories	19
2.2	Git Usage	19
2.2.1	Learning Git	20
2.3	Step by Step Modifying Bacula Code	21
2.3.1	More Details	23
2.4	Forcing Changes	24
3	Bacula FD Plugin API	27
3.1	Normal vs Command vs Options Plugins	27
3.2	Loading Plugins	28
3.3	loadPlugin	29
3.4	Plugin Entry Points	31
3.4.1	newPlugin(bpContext *ctx)	31
3.4.2	freePlugin(bpContext *ctx)	31
3.4.3	getPluginValue(bpContext *ctx, pVariable var, void *value)	32
3.4.4	setPluginValue(bpContext *ctx, pVariable var, void *value)	32
3.4.5	handlePluginEvent(bpContext *ctx, bEvent *event, void *value)	32
3.4.6	startBackupFile(bpContext *ctx, struct save_pkt *sp)	34
3.4.7	endBackupFile(bpContext *ctx)	35
3.4.8	startRestoreFile(bpContext *ctx, const char *cmd)	35
3.4.9	createFile(bpContext *ctx, struct restore_pkt *rp)	35
3.4.10	setFileAttributes(bpContext *ctx, struct restore_pkt *rp)	36
3.4.11	endRestoreFile(bpContext *ctx)	37



3.4.12	pluginIO(bpContext *ctx, struct io_pkt *io)	37
3.4.13	bool checkFile(bpContext *ctx, char *fname)	38
3.5	Bacula Plugin Entrypoints	38
3.5.1	bRC registerBaculaEvents(bpContext *ctx, ...)	38
3.5.2	bRC getBaculaValue(bpContext *ctx, bVariable var, void *value)	38
3.5.3	bRC setBaculaValue(bpContext *ctx, bVariable var, void *value)	39
3.5.4	bRC JobMessage(bpContext *ctx, const char *file, int line, int type, utime_t mtime, const char *fmt, ...)	39
3.5.5	bRC DebugMessage(bpContext *ctx, const char *file, int line, int level, const char *fmt, ...)	39
3.5.6	void baculaMalloc(bpContext *ctx, const char *file, int line, size_t size)	39
3.5.7	void baculaFree(bpContext *ctx, const char *file, int line, void *mem)	39
3.6	Building Bacula Plugins	39
3.7	Advanced Restore Options	40
3.8	Bacula Auth Plugin Documentation	41
3.8.1	Overview	41
3.8.2	Dictionary	41
3.8.3	Bacula DIR Plugin API	41
3.8.4	Bacula Pluggable Authentication Modules API Framework BPAM]	46
4	Platform Support	53
4.1	General	53
4.2	Requirements to become a Supported Platform	53
5	Daemon Protocol	55
5.1	General	55
5.2	Low Level Network Protocol	55
5.3	General Daemon Protocol	55
5.4	The Protocol Used Between the Director and the Storage Daemon	56
5.5	The Protocol Used Between the Director and the File Daemon	56
5.6	The Save Protocol Between the File Daemon and the Storage Daemon	57
5.6.1	Command and Control Information	57
5.6.2	Data Information	57
6	Director Services Daemon	59



7	File Services Daemon	61
7.1	Commands Received from the Director for a Backup	61
7.2	Commands Received from the Director for a Restore	62
8	Storage Daemon Design	63
8.1	SD Design Introduction	63
8.2	SD Development Outline	63
8.3	SD Connections and Sessions	64
8.3.1	SD Append Requests	64
8.3.2	SD Read Requests	65
8.4	SD Data Structures	65
9	Catalog Services	67
9.1	General	67
9.1.1	Filenames and Maximum Filename Length	67
9.1.2	Installing and Configuring MySQL	68
9.1.3	Installing and Configuring PostgreSQL	68
9.1.4	Internal Bacula Catalog	68
9.1.5	Database Table Design	68
9.2	Sequence of Creation of Records for a Save Job	68
9.3	Database Tables	69
10	Storage Media Output Format	81
10.1	General	81
10.2	Definitions	81
10.3	Storage Daemon File Output Format	83
10.4	Overall Format	83
10.5	Serialization	83
10.6	Block Header	83
10.7	Record Header	84
10.8	Version BB02 Block Header	85
10.9	Version 2 Record Header	85
10.10	Volume Label Format	85
10.11	Session Label	86



10.12	Overall Storage Format	87
10.13	Unix File Attributes	90
10.14	Old Depreciated Tape Format	91
11	Bacula Porting Notes	97
11.1	Porting Requirements	97
11.2	Steps to Take for Porting	98
11.1	General	101
11.1.1	Minimal Code in Console Program	101
11.1.2	GUI Interface is Difficult	101
11.2	Bvfs API	102
12	TLS	107
12.1	Introduction to TLS	107
12.2	New Configuration Directives	107
12.3	TLS API Implementation	108
12.3.1	Library Initialization and Cleanup	108
12.3.2	Manipulating TLS Contexts	108
12.3.3	Performing Post-Connection Verification	109
12.3.4	Manipulating TLS Connections	109
12.4	Bnet API Changes	110
12.4.1	Negotiating a TLS Connection	110
12.4.2	Manipulating Socket Blocking State	110
12.5	Authentication Negotiation	111
13	Bacula Regression Testing	113
13.1	Setting up Regression Testing	113
13.2	Running the Regression Script	113
13.2.1	Setting the Configuration Parameters	114
13.2.2	Building the Test Bacula	115
13.2.3	Setting up your SQL engine	115
13.2.4	Running the Disk Only Regression	115
13.2.5	Other Tests	117
13.2.6	If a Test Fails	117



13.3 Testing a Binary Installation	117
13.4 Running a Single Test	118
13.5 Writing a Regression Test	118
13.5.1 Running the Tests by Hand	118
13.5.2 Directory Structure	118
13.5.3 Adding a New Test	118
13.5.4 Adding a Unittest	119
13.5.5 Running a Test Under The Debugger	119
14 Bacula MD5 Algorithm	121
14.1 Command Line Message Digest Utility	121
14.1.1 Name	121
14.1.2 Synopsis	121
14.1.3 Description	121
14.1.4 Options	122
14.1.5 Files	122
14.1.6 Bugs	122
14.2 Download md5.zip (Zipped archive)	122
14.2.1 See Also	122
14.2.2 Exit Status	122
14.2.3 Copying	122
14.2.4 Acknowledgements	122
15 Bacula Memory Management	123
15.1 General	123
15.1.1 Statically Allocated Memory	123
15.1.2 Dynamically Allocated Memory	123
15.1.3 Pooled and Non-pooled Memory	124
16 TCP/IP Network Protocol	127
16.1 General	127
16.2 bnet and Threads	127
16.3 bnet_open	127
16.4 bnet_send	128



16.5	bnet_fsend	128
16.6	Additional Error information	128
16.7	bnet_recv	128
16.8	bnet_sig	128
16.9	bnet_strerror	129
16.10	bnet_close	129
16.11	Becoming a Server	129
16.12	Higher Level Conventions	129
17	Smart Memory Allocation	131
17.0.1	Installing SMARTALLOC	131
17.0.2	Squelching a SMARTALLOC	132
17.0.3	Living with Libraries	133
17.0.4	SMARTALLOC Details	134
17.0.5	When SMARTALLOC is Disabled	135
17.0.6	The alloc() Function	135
17.0.7	Overlays and Underhandedness	135
17.0.8	Test and Demonstration Program	136
17.0.9	Invitation to the Hack	136
17.1	Download smartall.zip (Zipped archive)	136
17.1.1	Copying	136
	Appendices	137
A	Acronyms	139
	Index	140





List of Figures

2.1	Git Edit Commit	24
3.1	BPAM General Workflow	50
11.1	Bat Brestore Panel	103
17.1	Smart Memory Allocation with Orphaned Buffer Detection	131





List of Tables

1.1	Error codes	17
9.1	Path table layout	69
9.2	File table layout	70
9.3	Job table layout	70
9.4	Job Types	71
9.5	Job Statuses	72
9.6	Filesets table layout	73
9.7	JobMedia table layout	73
9.8	Media table layout	73
9.9	Pool table layout	75
9.10	Client table layout	76
9.11	Storage table layout	76
9.12	Counter table layout	76
9.13	Jobhisto table layout	77
9.14	Log Table Layout	77
9.15	Location table layout	78
9.16	Locationlog table layout	78
9.17	Version table layout	78
9.18	Basefiles table layout	78
10.1	File Attributes	91
17.1	Systems functions	133





Chapter 1

Bacula Developer Notes

This document is intended mostly for developers and describes how you can contribute to the Bacula project and the the general framework of making Bacula source changes.

1.1 Contributions

Contributions to the Bacula project come in many forms: ideas, participation in helping people on the bacula-users email list, packaging Bacula binaries for the community, helping improve the documentation, and submitting code.

Contributions in the form of submissions for inclusion in the project are broken into two groups. The first are contributions that are aids and not essential to Bacula. In general, these will be scripts or will go into the [bacula/examples](#) directory. For these kinds of non-essential contributions there is no obligation to do a copyright assignment as described below. However, a copyright assignment would still be appreciated.

The second class of contributions are those which will be integrated with Bacula and become an essential part (code, scripts, documentation, ...) Within this class of contributions, there are two hurdles to surmount. One is getting your patch accepted, and two is dealing with copyright issues. The following text describes some of the requirements for such code.

1.2 Patches

Subject to the copyright assignment described below, your patches should be sent in [git format-patch](#) format relative to the current contents of the master branch of the Source Forge [git](#) repository. Please attach the output file or files generated by the [git format-patch](#) to the email rather than include them directory to avoid wrapping of the lines in the patch. Please be sure to use the Bacula indenting standard (see below) for source code. If you have checked out the source with [git](#), you can get a [diff](#) using.

```
| git pull  
| git format-patch -M
```

If you plan on doing significant development work over a period of time, after having your first patch reviewed and approved, you will be eligible for having developer [git](#) write access so that you can commit your changes directly to the [git](#) repository. To do so, you will need a userid on Source Forge.



1.3 Copyrights

To avoid future problems concerning changing licensing or copyrights, all code contributions more than a hand full of lines must be in the Public Domain or have the copyright transferred to the Free Software Foundation Europe e.V. with a Fiduciary Licence Agreement (FLA) as the case for all the current code.

Prior to November 2004, all the code was copyrighted by Kern Sibbald and John Walker. After November 2004, the code was copyrighted by Kern Sibbald, then on the 15th of November 2006, Kern transferred the copyright to the Free Software Foundation Europe e.V. In signing the FLA and transferring the copyright, you retain the right to use the code you have submitted as you want, and you ensure that Bacula will always remain Free and Open Source.

Your name should be clearly indicated as the author of the code, and you must be extremely careful not to violate any copyrights or patents or use other people's code without acknowledging it. The purpose of this requirement is to avoid future copyright, patent, or intellectual property problems. Please read the [LICENSE](#) agreement in the main Bacula source code directory. When you sign the Fiduciary Licence Agreement (FLA) and send it in, you are agreeing to the terms of that [LICENSE](#) file.

If you don't understand what we mean by future problems, please examine the difficulties Mozilla was having finding previous contributors at www.mozilla.org/MPL/missing.html. The other important issue is to avoid copyright, patent, or intellectual property violations as was (May 2003) claimed by SCO against IBM.

Although the copyright will be held by the Free Software Foundation Europe e.V., each developer is expected to indicate that he wrote and/or modified a particular module (or file) and any other sources. The copyright assignment may seem a bit unusual, but in reality, it is not. Most large projects require this.

If you have any doubts about this, please don't hesitate to ask. The objective is to assure the long term survival of the Bacula project.

Items not needing a copyright assignment are: most small changes, enhancements, or bug fixes of 5-10 lines of code, which amount to less than 20% of any particular file.

1.4 Copyright Assignment – Fiduciary License Agreement

Since this is not a commercial enterprise, and we prefer to believe in everyone's good faith, previously developers could assign the copyright by explicitly acknowledging that they do so in their first submission. This was sufficient if the developer is independent, or an employee of a not-for-profit organization or a university. However, in an effort to ensure that the Bacula code is really clean, beginning in August 2006, all previous and future developers with SVN write access will be asked to submit a copyright assignment (or Fiduciary Licence Agreement – FLA), which means you agree to the [LICENSE](#) in the main source directory. It also means that you receive back the right to use the code that you have submitted.

Any developer who wants to contribute and is employed by a company should either list the employer as the owner of the code, or get explicit permission from him to sign the copyright assignment. This is because in many countries, all work that an employee does whether on company time or in the employee's free time is considered to be Intellectual Property of the company. Obtaining official approval or an FLA from the company will avoid misunderstandings between the employee, the company, and the Bacula project. A good number of companies have already followed this procedure.

The Fiduciary Licence Agreement is posted on the Bacula web site at: www.bacula.org/en/FLA-bacula.en.pdf



The instructions for filling out this agreement are also at: www.bacula.org/?page=fsfe

It should be filled out, then sent to:

Kern Sibbald
Cotes-de-Montmoiret 9
1012 Lausanne
Switzerland

Please note that the above address is different from the officially registered office mentioned in the document. When you send in such a complete document, please notify me: <kern at sibbald dot com>, and please add your email address to the FLA so that I can contact you to confirm reception of the signed FLA.

1.5 The Development Cycle

As discussed on the email lists, the number of contributions are increasing significantly. We expect this positive trend will continue. As a consequence, we have modified how we do development, and instead of making a list of all the features that we will implement in the next version, each developer signs up for one (maybe two) projects at a time, and when they are complete, and the code is stable, we will release a new version. The release cycle will probably be roughly six months.

The difference is that with a shorter release cycle and fewer released feature, we will have more time to review the new code that is being contributed, and will be able to devote more time to a smaller number of projects (some prior versions had too many new features for us to handle correctly).

Future release schedules will be much the same, and the number of new features will also be much the same providing that the contributions continue to come – and they show no signs of let up :-)

1.5.1 Feature Requests

In addition, we have “formalized” the feature requests a bit.

Instead of me maintaining an informal list of everything I run into ([kernstodo](#)), we now maintain a “formal” list of projects. This means that all new feature requests, including those recently discussed on the email lists, must be formally submitted and approved.

Formal submission of feature requests will take two forms:

- 1 non-mandatory, but highly recommended is to discuss proposed new features on the mailing list.
- 2 Formal submission of an Feature Request in a special format. We'll give an example of this below, but you can also find it on the web site under Support → Feature Requests. Since it takes a bit of time to properly fill out a Feature Request form, you probably should check on the email list first.

Once the Feature Request is received by the keeper of the projects list, it will be sent to the Bacula project manager (Kern), and he will either accept it (90% of the time), send it back asking for clarification (10% of the time), send it to the email list asking for opinions, or reject it (very few cases).

If it is accepted, it will go in the [projects](#) file (a simple ASCII file) maintained in the main Bacula source directory.



Implementation of Feature Requests

Any qualified developer can sign up for a project. The project must have an entry in the projects file, and the developer's name will appear in the Status field.

How Feature Requests are accepted

Acceptance of Feature Requests depends on several things:

- 1 feedback from users. If it is negative, the Feature Request will probably not be accepted.
- 2 the difficulty of the project. A project that is so difficult that we cannot imagine finding someone to implement probably won't be accepted. Obviously if you know how to implement it, don't hesitate to put it in your Feature Request
- 3 whether or not the Feature Request fits within the current strategy of Bacula (for example an Feature Request that requests changing the tape to `tar` format probably would not be accepted, ...).

How Feature Requests are prioritized

Once an Feature Request is accepted, it needs to be implemented. If you can find a developer for it, or one signs up for implementing it, then the Feature Request becomes top priority (at least for that developer).

Between releases of Bacula, we will generally solicit Feature Request input for the next version, and by way of this email, we suggest that you send discuss and send in your Feature Requests for the next release. Please verify that the Feature Request is not in the current list (attached to this email).

Once users have had several weeks to submit Feature Requests, the keeper of the projects list will organize them, and request users to vote on them. This will allow fixing prioritizing the Feature Requests. Having a priority is one thing, but getting it implement is another thing – we are hoping that the Bacula community will take more responsibility for assuring the implementation of accepted Feature Requests.

Feature Request format:

```
===== Empty Feature Request form =====
Item n:  One line summary ...
Date:    Date submitted
Origin:  Name and email of originator.
Status:

What:    More detailed explanation ...

Why:     Why it is important ...

Notes:   Additional notes or features (omit if not used)
===== End Feature Request form =====
```

```
===== Example Completed Feature Request form =====
Item 1:  Implement a Migration job type that will move the job
        data from one device to another.
Origin:  Sponsored by Riege Software International GmbH. Contact:
        Daniel Holtkamp <holtkamp at riego dot com>
Date:    28 October 2005
Status:  Partially coded in 1.37 -- much more to do. Assigned to
        Kern.
```



What: The ability to copy, move, or archive data that is on a device to another device is very important.

Why: An ISP might want to backup to disk, but after 30 days migrate the data to tape backup and delete it from disk. Bacula should be able to handle this automatically. It needs to know what was put where, and when, and what to migrate -- it is a bit like retention periods. Doing so would allow space to be freed up for current backups while maintaining older data on tape drives.

Notes: Migration could be triggered by:
 Number of Jobs
 Number of Volumes
 Age of Jobs
 Highwater size (keep total size)
 Lowwater mark

=====

1.6 Bacula Code Submissions and Projects

Getting code implemented in Bacula works roughly as follows:

- Kern is the project manager, but prefers not to be a “gate keeper”. This means that the developers are expected to be self-motivated, and once they have experience submit directly to the [git](#) repositories. However, it is a good idea to have your patches reviewed prior to submitting, and it is a bad idea to submit monster patches because no one will be able to properly review them. See below for more details on this.
- There are growing numbers of contributions (very good).
- Some contributions come in the form of relatively small patches, which Kern reviews, integrates, documents, tests, and maintains.
- All Bacula developers take full responsibility for writing the code, posting as patches so that we can review it as time permits, integrating it at an appropriate time, responding to our requests for tweaking it (name changes, ...), document it in the code, document it in the manual (even though their mother tongue is not English), test it, develop and commit regression scripts, and answer in a timely fashion all bug reports – even occasionally accepting additional bugs :-)

This is a sustainable way of going forward with Bacula, and the direction that the project will be taking more and more. For example, in the past, we have had some very dedicated programmers who did major projects. However, some of these programmers due to outside obligations (job responsibilities change of job, school duties, ...) could not continue to maintain the code. In those cases, the code suffers from lack of maintenance, sometimes we patch it, sometimes not. In the end, if the code is not maintained, the code gets dropped from the project (there are two such contributions that are heading in that direction). When ever possible, we would like to avoid this, and ensure a continuation of the code and a sharing of the development, debugging, documentation, and maintenance responsibilities.

1.7 Patches for Released Versions

If you fix a bug in a released version, you should, unless it is an absolutely trivial bug, create and release a patch file for the bug. The procedure is as follows:

Fix the bug in the released branch and in the development master branch.

Make a patch file for the branch and add the branch patch to the patches directory in both the branch and the trunk. The name should be 2.2.4-xxx.patch where xxx is unique, in this case



it can be “restore”, e.g. `2.2.4-restore.patch`. Add to the top of the file a brief description and instructions for applying it – see for example `2.2.4-poll-mount.patch`. The best way to create the patch file is as follows:

```
(edit) 2.2.4-restore.patch
(input description)
(end edit)

git format-patch -M
mv 0001-xxx 2.2.4-restore.patch
```

check to make sure no extra junk got put into the patch file (i.e. it should have the patch for that bug only).

If there is not a bug report on the problem, create one, then add the patch to the bug report.

Then upload it to the 2.2.x release of bacula-patches.

So, end the end, the patch file is:

- Attached to the bug report
- In [Branch-2.2/bacula/patches/...](#)
- In the trunk
- Loaded on Source Forge bacula-patches 2.2.x release. When you add it, click on the check box to send an Email so that all the users that are monitoring SF patches get notified.

1.8 Developing Bacula

Typically the simplest way to develop Bacula is to open one `xterm` window pointing to the source directory you wish to update; a second `xterm` window at the top source directory level, and a third `xterm` window at the bacula directory [<top>/src/bacula](#). After making source changes in one of the directories, in the top source directory `xterm`, build the source, and start the daemons by entering:

```
| make
```

and

```
| ./startit
```

then in the enter:

```
| ./console
```

or

```
| ./gnome-console
```

to start the Console program. Enter any commands for testing. For example: `run kernsverify full`.

Note, the instructions here to use `./startit` are different from using a production system where the administrator starts Bacula by entering `./bacula start`. This difference allows a development version of **Bacula** to be run on a computer at the same time that a production



system is running. The `./startit` script starts **Bacula** using a different set of configuration files, and thus permits avoiding conflicts with any production system.

To make additional source changes, exit from the Console program, and in the top source directory, stop the daemons by entering:

```
| ./stopit
```

then repeat the process.

1.9 Debugging

Probably the first thing to do is to turn on debug output.

A good place to start is with a debug level of 20 as in `./startit -d20`. The `startit` command starts all the daemons with the same debug level. Alternatively, you can start the appropriate daemon with the debug level you want. If you really need more info, a debug level of 60 is not bad, and for just about everything a level of 200.

1.10 Using a Debugger

If you have a serious problem such as a segmentation fault, it can usually be found quickly using a good multiple thread debugger such as `gdb`. For example, suppose you get a segmentation violation in `bacula-dir`. You might use the following to find the problem:

<start the Storage and File daemons>

```
| cd dird
| gdb ./bacula-dir
| run -f -s -c ./dird.conf
```

<it dies with a segmentation fault> where The `-f` option is specified on the `run` command to inhibit `dird` from going into the background. You may also want to add the `-s` option to the `run` command to disable signals which can potentially interfere with the debugging.

As an alternative to using the debugger, each **Bacula** daemon has a built in back trace feature when a serious error is encountered. It calls the debugger on itself, produces a back trace, and emails the report to the developer. For more details on this, please see the chapter in the main Bacula manual entitled “What To Do When Bacula Crashes (Kaboom)”.

1.11 Memory Leaks

Because Bacula runs routinely and unattended on client and server machines, it may run for a long time. As a consequence, from the very beginning, Bacula uses `SmartAlloc` to ensure that there are no memory leaks. To make detection of memory leaks effective, all Bacula code that dynamically allocates memory **must** have a way to release it. In general when the memory is no longer needed, it should be immediately released, but in some cases, the memory will be held during the entire time that Bacula is executing. In that case, there **must** be a routine that can be called at termination time that releases the memory. In this way, we will be able to detect memory leaks. Be sure to immediately correct any and all memory leaks that are printed at the termination of the daemons.



1.12 Special Files

Kern uses files named 1, 2, ... 9 with any extension as scratch files. Thus any files with these names are subject to being rudely deleted at any time.

1.13 When Implementing Incomplete Code

Please identify all incomplete code with a comment that contains

```
| ***FIXME***
```

where there are three asterisks (*) before and after the word `FIXME` (in capitals) and no intervening spaces. This is important as it allows new programmers to easily recognize where things are partially implemented.

1.13.1 Bacula Source File Structure

The distribution generally comes as a `tar` file of the form `bacula.x.y.z.tar.gz` where x, y, and z are the version, release, and update numbers respectively.

Once you detar this file, you will have a directory structure as follows:

```
|
| Tar file:
| - depkgs
|   - mtx          (autochanger control program + tape drive info)
|   - sqlite       (SQLite database program)
|
| Tar file:
| - depkgs-win32
|   - pthreads     (Native win32 pthreads library -- dll)
|   - zlib         (Native win32 zlib library)
|   - wx           (wxWidgets source code)
|
| Project bacula:
| - bacula         (main source directory containing configuration
|                  and installation files)
|   - autoconf     (automatic configuration files, not normally used
|                  by users)
|   - intl         (programs used to translate)
|   - platforms    (OS specific installation files)
|     - redhat     (Red Hat installation)
|     - solaris    (Sun installation)
|     - freebsd    (FreeBSD installation)
|     - irix       (Irix installation -- not tested)
|     - unknown    (Default if system not identified)
|   - po           (translations of source strings)
|   - src          (source directory; contains global header files)
|     - plugins    (plugins for FD/SD/DIR)
|       - fd       (FileDaemon plugins)
|       - sd       (Storage Daemon plugins)
|       - dir      (Director Daemon plugins)
|     - cats       (SQL catalog database interface directory)
|     - console    (bacula user agent directory)
|     - dird       (Director daemon)
|     - filed      (Unix File daemon)
|     - findlib    (Unix file find library for File daemon)
|     - lib        (General Bacula library)
|     - stored     (Storage daemon)
|     - tools      (Various tool programs)
|     - win32      (Native Win32 File daemon)
|       - libwin32 (Win32 files to make bacula-fd be a service)
```



```

|- compat      (compatibility interface library)
|- filed       (links to src/filed)
|- findlib     (links to src/findlib)
|- lib         (links to src/lib)
|- console     (beginning of native console program)
|- wx-console  (wxWidget console Win32 specific parts)
|- win32-installer (Makensis installer for 32bit)
|- win64-installer (Makensis installer for 64bit)

Project regress:
|- regress     (Regression scripts)
|- bin         (temporary directory to hold Bacula installed binaries)
|- build       (temporary directory to hold Bacula source)
|- scripts     (scripts and .conf files)
|- tests       (test scripts)
|- tmp         (temporary directory for temp files)
|- working     (temporary working directory for Bacula daemons)

Project docs:
|- docs        (documentation directory)
|- developers  (Developer's guide)
|- home-page   (Bacula's home page source)
|- manual      (html document directory)
|- manual-fr   (French translation)
|- manual-de   (German translation)
|- techlogs    (Technical development notes);

Project gui:
|- gui         (Bacula GUI projects)
|- bacula-web  (Bacula web php management code)
|- bimagemgr   (Web application for burning CDRoms)

```

1.14 Header Files

Please carefully follow the scheme defined below as it permits in general only two header file includes per C file, and thus vastly simplifies programming. With a large complex project like Bacula, it isn't always easy to ensure that the right headers are invoked in the right order (there are a few kludges to make this happen – i.e. in a few include files because of the chicken and egg problem, certain references to typedefs had to be replaced with **void**).

Every file should include `bacula.h`. It pulls in just about everything, with very few exceptions. If you have system dependent ifdefing, please do it in `baconfig.h`. The version number and date are kept in `version.h`.

Each of the subdirectories (`console`, `cats`, `dird`, `filed`, `findlib`, `lib`, `stored`, ...) contains a single directory dependent include file generally the name of the directory, which should be included just after the include of `bacula.h`. This file (for example, for the `dird` directory, it is `dird.h`) contains either definitions of things generally needed in this directory, or it includes the appropriate header files. It always includes `protos.h`. See below.

Each subdirectory contains a header file named `protos.h`, which contains the prototypes for subroutines exported by files in that directory. `protos.h` is always included by the main directory dependent include file.

1.15 Programming Standards

For the most part, all code should be written in C unless there is a burning reason to use C++, and then only the simplest C++ constructs will be used. Note, Bacula is slowly evolving to use more and more C++.



Code should have some documentation – not a lot, but enough so that I can understand it. Look at the current code, and you will see that I document more than most, but am definitely not a fanatic.

We prefer simple linear code where possible. Gotos are strongly discouraged except for handling an error to either bail out or to retry some code, and such use of `gotos` can vastly simplify the program.

Remember this is a C program that is migrating to a **tiny** subset of C++, so be conservative in your use of C++ features.

1.16 Do Not Use

- STL – it is totally incomprehensible.

1.17 Avoid if Possible

- Using **`void *`** because this generally means that one must use casting, and in C++ casting is rather ugly. It is OK to use `void *` to pass structure address where the structure is not known to the routines accepting the packet (typically callback routines). However, declaring `void *buf` is a bad idea. Please use the correct types whenever possible.
- Using undefined storage specifications such as (`short`, `int`, `long`, `long long`, `size_t` ...). The problem with all these is that the number of bytes they allocate depends on the compiler and the system. Instead use Bacula's types (`int8_t`, `uint8_t`, `int32_t`, `uint32_t`, `int64_t`, and `uint64_t`). This guarantees that the variables are given exactly the size you want. Please try at all possible to avoid using `size_t`, `ssize_t` and the such. They are very system dependent. However, some system routines may need them, so their use is often unavoidable.
- Returning a malloc'ed buffer from a subroutine – someone will forget to release it.
- Heap allocation (`malloc`) unless needed – it is expensive. Use `POOL_MEM` instead.
- Templates – they can create portability problems.
- Fancy or tricky C or C++ code, unless you give a good explanation of why you used it.
- Too much inheritance – it can complicate the code, and make reading it difficult (unless you are in love with colons)

1.18 Do Use Whenever Possible

- Locking and unlocking within a single subroutine.
- A single point of exit from all subroutines. A `goto` is perfectly OK to use to get out early, but only to a label named `bail_out`, and possibly an `ok_out`. See current code examples.
- `malloc` and `free` within a single subroutine.
- Comments and global explanations on what your code or algorithm does.
- When committing a fix for a bug, make the comment of the following form:

| Reason for bug fix or other message. Fixes bug #1234

It is important to write the **bug #1234** like that because our program that automatically pulls messages from the `git` repository to make the `ChangeLog` looks for that pattern. Obviously the **1234** should be replaced with the number of the bug you actually fixed.

Providing the commit comment line has one of the following keywords (or phrases), it will be ignored:



```
tweak
typo
cleanup
bweb:
regress:
again
.gitignore
fix compilation
technotes
update version
update technotes
update kernstodo
update projects
update releasenotes
update version
update home
update release
update todo
update notes
update changelog
```

- Use the following keywords at the beginning of a `git` commit message

1.19 Indenting Standards

We find it very hard to read code indented 8 columns at a time. Even 4 at a time uses a lot of space, so we have adopted indenting 3 spaces at every level. Note, indentation is the visual appearance of the source on the page, while tabbing is replacing a series of up to 8 spaces from a tab character.

The closest set of parameters for the Linux **indent** program that will produce reasonably indented code are:

```
-nbad -bap -bbo -nbc -br -brs -c36 -cd36 -ncdb -ce -ci3 -cli0
-cp36 -d0 -di1 -ndj -nfc1 -nfca -hnl -i3 -ip0 -l85 -lp -npcs
-nprs -npsl -saf -sai -saw -nsob -nss -nbc -ncs -nbfd
```

You can put the above in your `.indent.pro` file, and then just invoke `indent` on your file. However, be warned. This does not produce perfect indenting, and it will mess up C++ class statements pretty badly.

Braces are required in all if statements (missing in some very old code). To avoid generating too many lines, the first brace appears on the first line (e.g. of an `if`), and the closing brace is on a line by itself. E.g.

```
if (abc) {
    some_code;
}
```

Just follow the convention in the code. For example we prefer non-indented cases.

```
switch (code) {
case 'A':
    do something
    break;
case 'B':
    again();
    break;
default:
    break;
}
```



Avoid using `//` style comments except for temporary code or turning off debug code. Standard C comments are preferred (this also keeps the code closer to C).

Attempt to keep all lines less than 85 characters long so that the whole line of code is readable at one time. This is not a rigid requirement.

Always put a brief description at the top of any new file created describing what it does and including your name and the date it was first written. Please don't forget any Copyrights and acknowledgments if it isn't 100% your code. Also, include the Bacula copyright notice that is in [src/c](#).

In general you should have two includes at the top of the an include for the particular directory the code is in, for includes are needed, but this should be rare.

In general (except for self-contained packages), prototypes should all be put in `protos.h` in each directory.

Always put space around assignment and comparison operators.

```
a = 1;
if (b >= 2) {
    cleanup();
}
```

but your can compress things in a **for** statement:

```
for (i=0; i < del.num_ids; i++) {
    ...
}
```

Don't overuse the inline if (`?:`). A full `if` is preferred, except in a print statement, e.g.:

```
if (ua->verbose && del.num_del != 0) {
    bsendmsg(ua, _("Pruned %d %s on Volume %s from catalog.\n"), del.num_del,
             del.num_del == 1 ? "Job" : "Jobs", mr->VolumeName);
}
```

Leave a certain amount of debug code (`Dmsg`) in code you submit, so that future problems can be identified. This is particularly true for complicated code likely to break. However, try to keep the debug code to a minimum to avoid bloating the program and above all to keep the code readable.

Please keep the same style in all new code you develop. If you include code previously written, you have the option of leaving it with the old indenting or re-indenting it. If the old code is indented with 8 spaces, then please re-indent it to Bacula standards.

If you are using `vim`, simply set your `tabstop` to 8 and your `shiftwidth` to 3.

1.20 Naming Convention

A fairly big amount of functions and variable are using the “snake_case” format, although you may find some classes that use the “CamelFormat”.

We strongly recommend to stick to the “snake_case” format to ease the reading of the code flow. Switching from one to the other is difficult and requires extra brain power.

Functions associated with a concept or a module and be prefixed with the name of the concept. For example, all Bacula Virtual FileSystem (bvfs) functions are prefixed with `bvfs_`.



Function and variable names must be readable, in general it is a bad idea to use striped down version of a function or a variable. In general you will find the following variables, functions or label in the code:

- `ret` - int/bool return code for a function
- `buf` - buffer
- `bail_out` - cleanup label for a function
- `ed1`, `ed2`, `ed3` - temporary buffer used to convert integers (50c long)
- `len` - length of a buffer
- `src` - source
- `dest` - destination
- `argc` - argument count
- `argk` - argument key (key=val)
- `argv` - argument value (key=val)
- `ua` - User Agent object
- `bvfs` - Bacula Virtual File System
- `lmgr` - Lock Manager
- `DCR` - Device Control Resource
- `ctx` - Context

1.21 Locks and Threads

Bacula has a builtin lock manager called `lmgr`. This lock manager is a wrapper for common pthread or mutex operation. The lock manager will overwrite the different POSIX thread functions via `src/lockmgr.h`.

It can detect deadlock situation during the run time. Found a deadlock !!!!

The lock manager can prevent deadlock and help developers to design with the mutex list (`lib/mutex_list.h`). If all mutex are acquired/released with some predefined order, deadlocks are not possible.

```
| ERROR: V out of order lock=%p %s:%i dumping lock
```

The lock manager will also dump all the mutex map during a backtrace. It can be analyzed easily to find the incorrect lock path. In the following example, we can see that one thread (0x7f67abe5f700) has requested to lock the same lock (0x6a8b40) two times from two different location.

```
| threadid=0x7f67abe5f700 max=4 current=1
|   lock=0x6a8b40 state=Granted priority=0 jobq.c:476
|   lock=0x6a8b40 state=Wanted priority=0 jobq.c:324
```

Using the thread id, we can locate in the backtrace what the thread was doing

```
| Thread 982 (Thread 0x7f67abe5f700 (LWP 6211)):
| #0 0x00007f6d1d8684ed in __lll_lock_wait () from /lib64/libpthread.so.0
| #1 0x00007f6d1d863dcb in _L_lock_883 () from /lib64/libpthread.so.0
| #2 0x00007f6d1d863c98 in pthread_mutex_lock () from /lib64/libpthread.so.0
```



```
#3 0x00007f6d1dcebdcb in lmgr_p (m=m@entry=0x6a8b40 <job_queue>) at lockmgr.c:105
#4 0x00007f6d1dced2ed in bthread_mutex_lock_p ("jobq.c", 324) at lockmgr.c:1026
#5 0x0000000000430d82 in jobq_remove (jcr=jcr@entry=0x7f6cfc542c8) at jobq.c:324
#6 0x000000000042ec9e in cancel_job (ua=ua@entry=0x7f6be807a7f8, jcr=jcr@entry=0x7f6cfc542c8,
    wait=wait@entry=60, cancel=cancel@entry=true) at job.c:746
#7 0x000000000042f10a in allow_duplicate_job (jcr=jcr@entry=0x7f6cfc361838) at job.c:1155
#8 0x00000000004317e5 in reschedule_job (je=0x7f6cfc48d8, jq=0x6a8b40, jcr=0x7f6cfc361838) at jobq.c:675
#9 jobq_server (arg=arg@entry=0x6a8b40 <job_queue>) at jobq.c:494
#10 0x00007f6d1dcece25 in lmgr_thread_launcher (x=0x7f6cfc369578) at lockmgr.c:1184
#11 0x00007f6d1d861dd5 in start_thread () from /lib64/libpthread.so.0
#12 0x00007f6d1c289ead in clone () from /lib64/libc.so.6
```

We can also have some information about the job 0x7f67abe5f700 looking the jcr dump list.

```
threadid=0x7f67abe5f700 JobId=936628 JobStatus=t jcr=0x7f6cfc361838
    name=XXX-SQL.2019-08-17_21.30.00_07
    use_count=1 killable=0
    JobType=B JobLevel=F
    sched_time=18-Aug-2019 03:13 start_time=18-Aug-2019 00:18
    end_time=18-Aug-2019 02:43 wait_time=01-Jan-1970 01:00
    db=0x7f69b0921e28 db_batch=0x7f6bb0033d38 batch_started=0
    wstore=0x7f6b246262e8 rstore=(nil) wjcr=(nil) client=0x7f6b240c5438
    reschedule_count=1 SD_msg_chan_started=0

BDB=0x7f69b0921e28 db_name=bacula db_user=bacula connected=true
    cmd="SELECT MediaId,VolumeName,VolJobs,..."
    RWLOCK=0x7f69b0921e40 w_active=0 w_wait=0
```

The lock manager can also handle threads, by default on Linux, it is not safe to call `pthread_kill` on a non existing thread. On linux, `pthread_t` is a pointer to a struct. As detached threads are released automatically, trying to kill an old thread will raise a segmentation fault. With the lockmanager, the replacement of the `pthread_kill` will check if the thread is registered in the lock manager before to kill it.

The lock manager also implements a ring of events. This list can be displayed in the backtrace file. It can be used to analyze the life of a mutex for example. (the backtrace is a fixed picture at a given time).

1.22 Tabbing

Tabbing (inserting the tab character in place of spaces) is as normal on all Unix systems – a tab is converted space up to the next column multiple of 8. My editor converts strings of spaces to tabs automatically – this results in significant compression of the files. Thus, you can remove tabs by replacing them with spaces if you wish. Please don't confuse tabbing (use of tab characters) with indenting (visual alignment of the code).

1.23 Don'ts

Please don't use:

```
strcpy()
strcat()
strncpy()
strncat();
sprintf()
snprintf()
```

They are system dependent and un-safe. These should be replaced by the Bacula safe equivalents:



```
char *bstrncpy(char *dest, char *source, int dest_size);
char *bstrncat(char *dest, char *source, int dest_size);
int bsnprintf(char *buf, int32_t buf_len, const char *fmt, ...);
int bvsprintf(char *str, int32_t size, const char *format, va_list ap);
```

See `src/lib/bsys.c` for more details on these routines.

Don't use the `%lld` or the `%q` printf format editing types to edit 64 bit integers – they are not portable. Instead, use `%s` with `edit_uint64()`. For example:

```
char buf[100];
uint64_t num = something;
char ed1[50];
bsnprintf(buf, sizeof(buf), "Num=%s\n", edit_uint64(num, ed1));
```

Note : `%lld` is now permitted in Bacula code – we have our own printf routines which handle it correctly. The `edit_uint64()` subroutine can still be used if you wish, but over time, most of that old style will be removed.

The edit buffer `ed1` must be at least 27 bytes long to avoid overflow. See `src/lib/edit.c` for more details. If you look at the code, don't start screaming that I use `lld`. I actually use subtle trick taught to me by John Walker. The `lld` that appears in the editing routine is actually `#define` to a what is needed on your OS (usually "lld" or "q") and is defined in `autoconf/configure.in` for each OS. C string concatenation causes the appropriate string to be concatenated to the "%".

Also please don't use the STL or Templates or any complicated C++ code.

1.24 Message Classes

Currently, there are five classes of messages: Debug, Error, Job, Memory, and Queued.

1.25 Debug Messages

Debug messages are designed to be turned on at a specified debug level and are always sent to `STDOUT`. There are designed to only be used in the development debug process. They are coded as:

```
DmsgN(level, message, arg1, ...)
```

where the `N` is a number indicating how many arguments are to be substituted into the message (i.e. it is a count of the number arguments you have in your message – generally the number of percent signs (%)). `level` is the debug level at which you wish the message to be printed. `message` is the debug message to be printed, and `arg1, ...` are the arguments to be substituted. Since not all compilers support `#defines` with `varargs`, you must explicitly specify how many arguments you have.

When the debug message is printed, it will automatically be prefixed by the name of the daemon which is running, the filename where the `Dmsg` is, and the line number within the file.

Some actual examples are:

```
Dmsg2(20, "MD5len=%d MD5=%s\n", strlen(buf), buf);
```



```
| Dmsg1(9, "Created client %s record\n", client->hdr.name);
```

1.25.1 Debug Tags

It is possible to use tags to organize debug messages. When using a tag, it is possible to hide/show an entire class of messages. The message tags are defined in `lib/messages.c/h`

```
...
#define DT_SCHEDULER (1<<23)          /* scheduler */
#define DT_PROTOCOL (1<<22)           /* protocol */
#define DT_DEDUP (1<<21)              /* BEEF deduplication */
#define DT_DDE (1<<20)                /* BEEF dedup engine */
#define DT_SNAPSHOT (1<<19)           /* Snapshot */
#define DT_RECORD (1<<18)             /* Record/block */
#define DT_CLOUD (1<<17)              /* cloud */
#define DT_ALL (0x7FFF0000)           /* all (up to debug_level 65635, 15 flags available) */

/* setdebug tag=all,-plugin */
static struct debugtags debug_tags[] = {
...
{ NT_("scheduler"), DT_SCHEDULER, _("Debug scheduler information")},
{ NT_("protocol"), DT_PROTOCOL, _("Debug protocol information")},
{ NT_("snapshot"), DT_SNAPSHOT, _("Debug snapshots")},
{ NT_("record"), DT_RECORD, _("Debug records")},
{ NT_("all"), DT_ALL, _("Debug all information")},
{ NULL, 0, NULL}
};
```

In the code, the tag can be used like in the following example:

```
| Dmsg0(DT_CLOUD|50, "This is a message about cloud\n");
```

The message will be displayed if the `ccloud` tag is set AND if the debug level is at least 50.

```
| setdebug level=50 tags=cloud storage
```

1.26 Error Messages

Error messages are messages that are related to the daemon as a whole rather than a particular job. For example, an out of memory condition may generate an error message. They should be very rarely needed. In general, you should be using Job and Job Queued messages (`Jmsg` and `Qmsg`). They are coded as:

`MsgN(error-code, level, message, arg1, ...)` As with debug messages, you must explicitly code the of arguments to be substituted in the message. Error-code indicates the severity or class of error, and it may be one of the following (See table 1.1 on the facing page:

Code	Meaning
M_ABORT	Causes the daemon to immediately abort. This should be used only in extreme cases. It attempts to produce a traceback.
M_ERROR_TERM	Causes the daemon to immediately terminate. This should be used only in extreme cases. It does not produce a traceback.
M_FATAL	Causes the daemon to terminate the current job, but the daemon keeps running.
M_ERROR	Reports the error. The daemon and the job continue running.

Continues on the following page



Code	Meaning
M_WARNING	Reports a warning message. The daemon and the job continue running.
M_INFO	Reports an informational message.

Table 1.1: Error codes

There are other error message classes, but they are in a state of being redesigned or deprecated, so please do not use them. Some actual examples are:

```
| Emsg1(M_ABORT, 0, "Cannot create message thread: %s\n",strerror(status));

| Emsg3(M_WARNING, 0, "Connect to File daemon %s at %s:%d failed. Retrying...\n",
|     client->hdr.name, client->address,client->port);

| Emsg3(M_FATAL, 0, "bdir<filed: bad response from Filed to %s command:
|     %d %s\n", cmd, n, strerror(errno));
```

1.27 Job Messages

Job messages are messages that pertain to a particular job such as a file that could not be saved, or the number of files and bytes that were saved. They Are coded as:

```
| Jmsg(jcr, M_FATAL, 0, "Text of message");
```

A Jmsg with M_FATAL will fail the job. The Jmsg() takes varargs so can have any number of arguments for substituted in a printf like format. Output from the Jmsg() will go to the Job report.

If the Jmsg is followed with a number such as Jmsg1(...), the number indicates the number of arguments to be substituted (varargs is not standard for #defines), and what is more important is that the file and line number will be prefixed to the message. This permits a sort of debug from user's output.

1.28 Queued Job Messages

Queued Job messages are similar to Jmsg()s except that the message is Queued rather than immediately dispatched. This is necessary within the network subroutines and in the message editing routines. This is to prevent recursive loops, and to ensure that messages can be delivered even in the event of a network error.

1.29 Memory Messages

Memory messages are messages that are edited into a memory buffer. Generally they are used in low level routines such as the low level device file `dev.c` in the Storage daemon or in the low level Catalog routines. These routines do not generally have access to the Job Control Record and so they return error messages reformatted in a memory buffer. Mmsg() is the way to do this.



1.30 Bugs Database

We have a bugs database which is at: bugs.bacula.org, and as a developer you will need to respond to bugs, perhaps bugs in general if you have time, otherwise just bugs that correspond to code that you wrote.

If you need to answer bugs, please be sure to ask the Project Manager (currently Kern) to give you Developer access to the bugs database. This allows you to modify statuses and close bugs.

The first thing is if you want to take over a bug, rather than just make a note, you should assign the bug to yourself. This helps other developers know that you are the principal person to deal with the bug. You can do so by going into the bug and clicking on the **Update Issue** button. Then you simply go to the **Assigned To** box and select your name from the drop down box. To actually update it you must click on the **Update Information** button a bit further down on the screen, but if you have other things to do such as add a Note, you might wait before clicking on the **Update Information** button.

Generally, we set the **Status** field to either acknowledged, confirmed, or feedback when we first start working on the bug. Feedback is set when we expect that the user should give us more information.

Normally, once you are reasonably sure that the bug is fixed, and a patch is made and attached to the bug report, and/or in the SVN, you can close the bug. If you want the user to test the patch, then leave the bug open, otherwise close it and set **Resolution** to **Fixed**. We generally close bug reports rather quickly, even without confirmation, especially if we have run tests and can see that for us the problem is fixed. However, in doing so, it avoids misunderstandings if you leave a note while you are closing the bug that says something to the following effect: We are closing this bug because . . . If for some reason, it does not fix your problem, please feel free to reopen it, or to open a new bug report describing the problem".

We do not recommend that you attempt to edit any of the bug notes that have been submitted, nor to delete them or make them private. In fact, if someone accidentally makes a bug note private, you should ask the reason and if at all possible (with his agreement) make the bug note public.

If the user has not properly filled in most of the important fields (platform, OS, Product Version, . . .) please do not hesitate to politely ask him. Also, if the bug report is a request for a new feature, please politely send the user to the Feature Request menu item on www.bacula.org. The same applies to a support request (we answer only bugs), you might give the user a tip, but please politely refer him to the manual and the Getting Support page of www.bacula.org.



Chapter 2

Bacula Git Usage

This chapter is intended to help you use the Git source code repositories to obtain, modify, and submit Bacula source code.

2.1 Bacula Git repositories

As of September 2009, the Bacula source code has been split into three Git repositories. One is a repository that holds the main Bacula source code with directories **bacula**, **gui**, and **regress**. The second repository contains the **docs** directory, and the third repository contains the **rescue** directory. All three repositories are hosted by UKFast.

Previously everything was in a single SVN repository. We have split the SVN repository into three because Git offers significant advantages for ease of managing and integrating developer's changes. However, one of the disadvantages of Git is that you must work with the full repository, while SVN allows you to checkout individual directories. If we put everything into a single Git repository it would be far bigger than most developers would want to checkout, so we have separated the docs and rescue into their own repositories, and moved only the parts that are most actively worked on by the developers (bacula, gui, and regress) to a the Git Bacula repository.

Each major version of Bacula uses a specific branch name rather than the default git "master" branch. For example, the version 11.0.1 code will be accessible via the branch Branch-11.0. In this section, the word "master" refers to the current active development branch.

Bacula developers must now have a certain knowledge of Git.

2.2 Git Usage

Please note that if you are familiar with SVN, Git is similar, (and better), but there can be a few surprising differences that can be very confusing (nothing worse than converting from CVS to SVN).

The main Bacula Git repo contains the subdirectories **bacula**, **gui**, and **regress**. With Git it is not possible to pull only a single directory, because of the hash code nature of Git, you must take all or nothing.

For developers, the most important thing to remember about Git and the bacula.org repository is not to "force" a **push** to the repository. Doing so, can possibly rewrite the Git repository history and cause a lot of problems for the project.

You can get a full copy of the Bacula Git repository with the following command:



```
git clone http://git.bacula.org/bacula.git bacula
git checkout -b Branch-11.0 origin/Branch-11.0
```

This will put a read-only copy into the directory **bacula** in your current directory, and **bacula** will contain the subdirectories: **bacula**, **gui**, and **regress**. Obviously you can use any name an not just **bacula**. In fact, once you have the repository in say **bacula**, you can copy the whole directory to another place and have a fully functional git repository.

The above command needs to be done only once. Thereafter, you can:

```
cd bacula
git pull      # refresh my repo with the latest code
```

As of August 2009, the size of the repository (**bacula** in the above example) will be approximately 55 Megabytes. However, if you build from source in this directory and do a lot of updates and regression testing, the directory could become several hundred megabytes.

2.2.1 Learning Git

If you want to learn more about Git, we recommend that you visit:

book.git-scm.com/.

Some of the differences between Git and SVN are:

- Your main Git directory is a full Git repository to which you can and must commit. In fact, we suggest you commit frequently.
- When you commit, the commit goes into your local Git database. You must use another command to write it to the bacula.org repository (see below).
- The local Git database is kept in the directory **.git** at the top level of the directory.
- All the important Git configuration information is kept in the file **.git/config** in ASCII format that is easy to manually edit.
- When you do a **commit** the changes are put in **.git** rather but not in the main bacula.org repository.
- You can push your changes to the external repository using the command **git push** providing you have write permission on the repository.
- We restrict developers just learning git to have read-only access until they feel comfortable with git before giving them write access.
- You can download all the current changes in the external repository and merge them into your **master** branch using the command **git pull**.
- The command **git add** is used to add a new file to the repository AND to tell Git that you want a file that has changed to be in the next commit. This has lots of advantages, because a **git commit** only commits those files that have been explicitly added. Note with SVN **add** is used only to add new files to the repo.
- You can add and commit all files modified in one command using **git commit -a**.
- This extra use of **add** allows you to make a number of changes then add only a few of the files and commit them, then add more files and commit them until you have committed everything. This has the advantage of allowing you to more easily group small changes and do individual commits on them. By keeping commits smaller, and separated into topics, it makes it much easier to later select certain commits for backporting.
- If you **git pull** from the main repository and make some changes, and before you do a **git push** someone else pushes changes to the Git repository, your changes will apply to an older version of the repository you will probably get an error message such as:



```
git push
To git@github.com:bacula/bacula.git
! [rejected]      Branch-11.0 -> Branch-11.0 (non-fast forward)
error: failed to push some refs to 'git@github.com:bacula/bacula.git'
```

which is Git's way of telling you that the main repository has changed and that if you push your changes, they will not be integrated properly. This is very similar to what happens when you do an "svn update" and get merge conflicts. As we have noted above, you should never ask Git to force the push. See below for an explanation of why.

- To integrate (merge) your changes properly, you should always do a **git pull** just prior to doing a **git push**.
- If Git is unable to merge your changes or finds a conflict it will tell you and you must do conflict resolution, which is much easier in Git than in SVN.
- Resolving conflicts is described below in the **github** section.

2.3 Step by Step Modifying Bacula Code

Suppose you want to download Bacula source code, build it, make a change, then submit your change to the Bacula developers. What would you do?

- Tell git who you are:

```
git config --global user.name "First-name Last-name"
git config --global user.email "email@address.com"
```

Where you put your real name and your email address. Since this is global, you only need to do it once on any given machine regardless of how many git repos you work with.

- Download the Source code:

```
git clone http://git.bacula.org/bacula.git bacula
```

- Configure and Build Bacula:

```
./configure (all-your-normal-options)
make
```

- Create a branch to work on:

```
cd bacula/bacula
git checkout -b bugfix Branch-11.0
```

- Edit, build, Test, ...

```
edit file jcr.h
make
test
```

Note: if you forget to create a working branch prior to making changes, and you make them on the development branch, this is no problem providing that you create the working branch before your first commit. So assuming that you have edited "master" instead of your bugfix branch, you can simply:

```
git checkout -b bugfix master
```

and a new bugfix branch will be created and checked out. You can then proceed to committing to your bugfix branch as described in the next step.

- commit your work:

```
git commit -am "Short comment on what I did"
```



- Possibly repeat the above two items
- Switch back to the master branch:

```
| git checkout master
```

- Pull the latest changes:

```
| git pull
```

- Get back on your bugfix branch:

```
| git checkout bugfix
```

- Merge your changes and correct any conflicts:

```
| git rebase master bugfix
```

- Fix any conflicts:

You will be notified if there are conflicts. The first thing to do is:

```
| git diff
```

This will produce a diff of only the files having a conflict. Fix each file in turn. When it is fixed, the diff for that file will go away.

For each file fixed, you must do the same as SVN, inform git with:

```
| git add (name-of-file-no-longer-in-conflict)
```

- When all files are fixed do:

```
| git rebase --continue
```

- If you find that it is impossible to reconcile the two branches or you made a mistake in correcting and adding files, before you enter the:

```
| git rebase --continue
```

you can instead enter:

```
| git rebase --abort
```

which will essentially cancel the the original git rebase and reset everything to the beginning with no changes to your bugfix branch.

- When you have completed the rebase and are ready to send a patch, do the following:

```
| git checkout bugfix  
| git format-patch -M master
```

Look at the files produced. They should be numbered 0001-xxx.patch where there is one file for each commit you did, number sequentially, and the xxx is what you put in the commit comment.

- If the patch files are good, send them by email to the developers as attachments.
- Then you can continue working on your code if you want, or start another branch with a new project.
- If you continue working on your bugfix branch, you should do a **git rebase master** from time to time, and when your changes are committed to the repo, you will be automatically synchronized. So that the next **git format-patch** will produce only the changes you made since the last format-patch you sent to the developers.



2.3.1 More Details

Normally, you will work by creating a branch of the development branch of your repository, make your modifications, then make sure it is up to date, and finally create format-patch patches or push it to the bacula.org. Assuming you call the Bacula repository **bacula**, you might use the following commands:

```
cd bacula
git checkout bacula
git pull
git checkout -b newbranch bacula
(edit, ...)
git add <file-edited>
git commit -m "<comment about commit>"
...
```

When you have completed working on your branch, you will do:

```
cd bacula
git checkout newbranch # ensure I am on my branch
git pull               # get latest source code
git rebase master       # merge my code
```

If you have completed your edits before anyone has modified the repository, the **git rebase master** will report that there was nothing to do. Otherwise, it will merge the changes that were made in the repository before your changes. If there are any conflicts, Git will tell you. Typically resolving conflicts with Git is relatively easy. You simply make a diff:

```
| git diff
```

Then edit each file that was listed in the **git diff** to remove the conflict, which will be indicated by lines of:

```
| other text
```

where **text** is what is in the Bacula repository, and **other text** is what you have changed.

Once you have eliminated the conflict, the **git diff** will show nothing, and you must do a:

```
| git add <file-with-conflicts-fixed>
```

Once you have fixed all the files with conflicts in the above manner, you enter:

```
| git rebase --continue
```

and your rebase will be complete.

If for some reason, before doing the `--continue`, you want to abort the rebase and return to what you had, you enter:

```
| git rebase --abort
```

Finally to make a set of patch files

```
| git format-patch -M master
```



When you see your changes have been integrated and pushed to the main repo, you can delete your branch with:

```
git checkout master  
git branch -D newbranch
```

2.4 Forcing Changes

If you want to understand why it is not a good idea to force a push to the repository, look at the following picture:

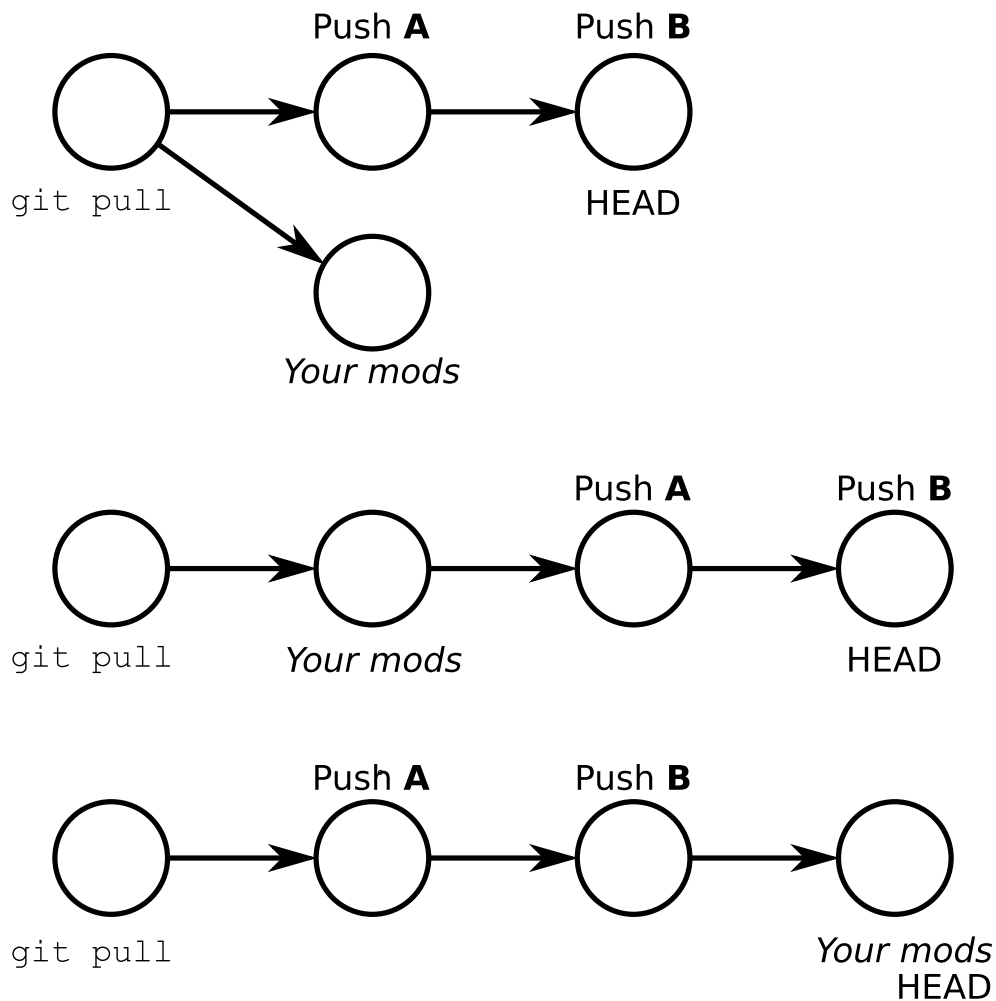


Figure 2.1: Git Edit Commit

The above graphic has three lines of circles. Each circle represents a commit, and time runs from the left to the right. The top line shows the repository just before you are going to do a push. Note the point at which you pulled is the circle on the left, your changes are represented by the circle labeled **Your mods**. It is shown below to indicate that the changes are only in your local repository. Finally, there are pushes A and B that came after the time at which you pulled.

If you were to force your changes into the repository, Git would place them immediately after the point at which you pulled them, so they would go before the pushes A and B. However, doing so would rewrite the history of the repository and make it very difficult for other users to



synchronize since they would have to somehow wedge their changes at some point before the current HEAD of the repository. This situation is shown by the second line of pushes.

What you really want to do is to put your changes after Push B (the current HEAD). This is shown in the third line of pushes. The best way to accomplish this is to work in a branch, pull the repository so you have your master equal to HEAD (in first line), then to rebase your branch on the current master and then commit it. The exact commands to accomplish this are shown in the next couple of sections.





Chapter 3

Bacula FD Plugin API

To write a Bacula plugin, you create a dynamic shared object program (or dll on Win32) with a particular name and two exported entry points, place it in the **Plugins Directory**, which is defined in the **bacula-fd.conf** file in the **Client** resource, and when the FD starts, it will load all the plugins that end with **-fd.so** (or **-fd.dll** on Win32) found in that directory.

3.1 Normal vs Command vs Options Plugins

In general, there are three ways that plugins are called. The first way, is when a particular event is detected in Bacula, it will transfer control to each plugin that is loaded in turn informing the plugin of the event. This is very similar to how a **RunScript** works, and the events are very similar. Once the plugin gets control, it can interact with Bacula by getting and setting Bacula variables. In this way, it behaves much like a RunScript. Currently very few Bacula variables are defined, but they will be implemented as the need arises, and it is very extensible.

We plan to have plugins register to receive events that they normally would not receive, such as an event for each file examined for backup or restore. This feature is not yet implemented.

The second type of plugin, which is more useful and fully implemented in the current version is what we call a command plugin. As with all plugins, it gets notified of important events as noted above (details described below), but in addition, this kind of plugin can accept a command line, which is a:

```
| Plugin = <command-string>
```

directive that is placed in the Include section of a FileSet and is very similar to the "File =" directive. When this Plugin directive is encountered by Bacula during backup, it passes the "command" part of the Plugin directive only to the plugin that is explicitly named in the first field of that command string. This allows that plugin to backup any file or files on the system that it wants. It can even create "virtual files" in the catalog that contain data to be restored but do not necessarily correspond to actual files on the filesystem.

The important features of the command plugin entry points are:

- It is triggered by a "Plugin =" directive in the FileSet
- Only a single plugin is called that is named on the "Plugin =" directive.
- The full command string after the "Plugin =" is passed to the plugin so that it can be told what to backup/restore.



The third type of plugin is the Options Plugin, this kind of plugin is useful to implement some custom filter on data. For example, you can implement a compression algorithm in a very simple way. Bacula will call this plugin for each file that is selected in a FileSet (according to Wild/Regex/Exclude/Include rules). As with all plugins, it gets notified of important events as noted above (details described below), but in addition, this kind of plugin can be placed in a Options group, which is a:

```
FileSet {
  Name = TestFS
  Include {
    Options {
      Compression = GZIP1
      Signature = MD5
      Wild = "*.txt"
      Plugin = <command-string>
    }
    File = /
  }
}
```

3.2 Loading Plugins

Once the File daemon loads the plugins, it asks the OS for the two entry points (loadPlugin and unloadPlugin) then calls the **loadPlugin** entry point (see below).

Bacula passes information to the plugin through this call and it gets back information that it needs to use the plugin. Later, Bacula will call particular functions that are defined by the **loadPlugin** interface.

When Bacula is finished with the plugin (when Bacula is going to exit), it will call the **unloadPlugin** entry point.

The two entry points are:

```
bRC loadPlugin(bInfo *lbinfo, bFuncs *lbfuncs, pInfo **pinfo, pFuncs **pfuncs)

and

bRC unloadPlugin()
```

both these external entry points to the shared object are defined as C entry points to avoid name mangling complications with C++. However, the shared object can actually be written in any language (preferably C or C++) providing that it follows C language calling conventions.

The definitions for **bRC** and the arguments are **src/typed/fd-plugins.h** and so this header file needs to be included in your plugin. It along with **src/lib/plugins.h** define basically the whole plugin interface. Within this header file, it includes the following files:

```
#include <sys/types.h>
#include "config.h"
#include "bc_types.h"
#include "lib/plugins.h"
#include <sys/stat.h>
```

Aside from the **bc_types.h** and **config.h** headers, the plugin definition uses the minimum code from Bacula. The **bc_type.h** file is required to ensure that the data type definitions in arguments correspond to the Bacula core code.

The return codes are defined as:



```
typedef enum {
    bRC_OK      = 0,          /* OK */
    bRC_Stop    = 1,          /* Stop calling other plugins */
    bRC_Error   = 2,          /* Some kind of error */
    bRC_More    = 3,          /* More files to backup */
    bRC_Term    = 4,          /* Unload me */
    bRC_Seen    = 5,          /* Return code from checkFiles */
    bRC_Core    = 6,          /* Let Bacula core handles this file */
    bRC_Skip    = 7,          /* Skip the proposed file */
} bRC;
```

At a future point in time, we hope to make the Bacula libbac.a into a shared object so that the plugin can use much more of Bacula's infrastructure, but for this first cut, we have tried to minimize the dependence on Bacula.

3.3 loadPlugin

As previously mentioned, the **loadPlugin** entry point in the plugin is called immediately after Bacula loads the plugin when the File daemon itself is first starting. This entry point is only called once during the execution of the File daemon. In calling the plugin, the first two arguments are information from Bacula that is passed to the plugin, and the last two arguments are information about the plugin that the plugin must return to Bacula. The call is:

```
| bRC loadPlugin(bInfo *lbinfo, bFuncs *lbfuncs, pInfo **pinfo, pFuncs **pfuncs)
```

and the arguments are:

lbinfo This is information about Bacula in general. Currently, the only value defined in the **lbinfo** structure is the version, which is the Bacula plugin interface version, currently defined as 1. The **size** is set to the byte size of the structure. The exact definition of the **lbinfo** structure as of this writing is:

```
typedef struct s_baculaInfo {
    uint32_t size;
    uint32_t version;
} bInfo;
```

lbfuncs The **bFuncs** structure defines the callback entry points within Bacula that the plugin can use register events, get Bacula values, set Bacula values, and send messages to the Job output or debug output.

The exact definition as of this writing is:

```
typedef struct s_baculaFuncs {
    uint32_t size;
    uint32_t version;
    bRC (*registerBaculaEvents)(bpContext *ctx, ...);
    bRC (*getBaculaValue)(bpContext *ctx, bVariable var, void *value);
    bRC (*setBaculaValue)(bpContext *ctx, bVariable var, void *value);
    bRC (*JobMessage)(bpContext *ctx, const char *file, int line,
        int type, utime_t mtime, const char *fmt, ...);
    bRC (*DebugMessage)(bpContext *ctx, const char *file, int line,
        int level, const char *fmt, ...);
    void *(*baculaMalloc)(bpContext *ctx, const char *file, int line,
        size_t size);
    void (*baculaFree)(bpContext *ctx, const char *file, int line, void *mem);
} bFuncs;
```

We will discuss these entry points and how to use them a bit later when describing the plugin code.

pInfo When the **loadPlugin** entry point is called, the plugin must initialize an information structure about the plugin and return a pointer to this structure to Bacula.

The exact definition as of this writing is:



```
typedef struct s_pluginInfo {
    uint32_t size;
    uint32_t version;
    const char *plugin_magic;
    const char *plugin_license;
    const char *plugin_author;
    const char *plugin_date;
    const char *plugin_version;
    const char *plugin_description;
} pInfo;
```

Where:

version is the current Bacula defined plugin interface version, currently set to 1. If the interface version differs from the current version of Bacula, the plugin will not be run (not yet implemented).

plugin_magic is a pointer to the text string `"*FDPluginData*"`, a sort of sanity check. If this value is not specified, the plugin will not be run (not yet implemented).

plugin_license is a pointer to a text string that describes the plugin license. Bacula will only accept compatible licenses (not yet implemented).

plugin_author is a pointer to the text name of the author of the program. This string can be anything but is generally the author's name.

plugin_date is the pointer text string containing the date of the plugin. This string can be anything but is generally some human readable form of the date.

plugin_version is a pointer to a text string containing the version of the plugin. The contents are determined by the plugin writer.

plugin_description is a pointer to a string describing what the plugin does. The contents are determined by the plugin writer.

The plnfo structure must be defined in static memory because Bacula does not copy it and may refer to the values at any time while the plugin is loaded. All values must be supplied or the plugin will not run (not yet implemented). All text strings must be either ASCII or UTF-8 strings that are terminated with a zero byte.

pFuncs When the loadPlugin entry point is called, the plugin must initialize an entry point structure about the plugin and return a pointer to this structure to Bacula. This structure contains pointer to each of the entry points that the plugin must provide for Bacula. When Bacula is actually running the plugin, it will call the defined entry points at particular times. All entry points must be defined.

The pFuncs structure must be defined in static memory because Bacula does not copy it and may refer to the values at any time while the plugin is loaded.

The exact definition as of this writing is:

```
typedef struct s_pluginFuncs {
    uint32_t size;
    uint32_t version;
    bRC (*newPlugin)(bpContext *ctx);
    bRC (*freePlugin)(bpContext *ctx);
    bRC (*getPluginValue)(bpContext *ctx, pVariable var, void *value);
    bRC (*setPluginValue)(bpContext *ctx, pVariable var, void *value);
    bRC (*handlePluginEvent)(bpContext *ctx, bEvent *event, void *value);
    bRC (*startBackupFile)(bpContext *ctx, struct save_pkt *sp);
    bRC (*endBackupFile)(bpContext *ctx);
    bRC (*startRestoreFile)(bpContext *ctx, const char *cmd);
    bRC (*endRestoreFile)(bpContext *ctx);
    bRC (*pluginIO)(bpContext *ctx, struct io_pkt *io);
    bRC (*createFile)(bpContext *ctx, struct restore_pkt *rp);
    bRC (*setFileAttributes)(bpContext *ctx, struct restore_pkt *rp);
    bRC (*checkFile)(bpContext *ctx, char *fname);
} pFuncs;
```

The details of the entry points will be presented in separate sections below.

Where:



size is the byte size of the structure.

version is the plugin interface version currently set to 3.

Sample code for loadPlugin:

```

bfuncs = lbfuns;           /* set Bacula funct pointers */
binfo  = lbinfo;
*pinfo = &pluginInfo;      /* return pointer to our info */
*pfuns = &pluginFuncs;     /* return pointer to our functions */

return bRC_OK;

```

where pluginInfo and pluginFuncs are statically defined structures. See bpipe-fd.c for details.

3.4 Plugin Entry Points

This section will describe each of the entry points (subroutines) within the plugin that the plugin must provide for Bacula, when they are called and their arguments. As noted above, pointers to these subroutines are passed back to Bacula in the pFuncs structure when Bacula calls the loadPlugin() externally defined entry point.

3.4.1 newPlugin(bpContext *ctx)

This is the entry point that Bacula will call when a new "instance" of the plugin is created. This typically happens at the beginning of a Job. If 10 Jobs are running simultaneously, there will be at least 10 instances of the plugin.

The bpContext structure will be passed to the plugin, and during this call, if the plugin needs to have its private working storage that is associated with the particular instance of the plugin, it should create it from the heap (malloc the memory) and store a pointer to its private working storage in the **pContext** variable. Note: since Bacula is a multi-threaded program, you must not keep any variable data in your plugin unless it is truly meant to apply globally to the whole plugin. In addition, you must be aware that except the first and last call to the plugin (loadPlugin and unloadPlugin) all the other calls will be made by threads that correspond to a Bacula job. The bpContext that will be passed for each thread will remain the same throughout the Job thus you can keep your private Job specific data in it (**bContext**).

```

typedef struct s_bpContext {
    void *pContext; /* Plugin private context */
    void *bContext; /* Bacula private context */
} bpContext;

```

This context pointer will be passed as the first argument to all the entry points that Bacula calls within the plugin. Needless to say, the plugin should not change the bContext variable, which is Bacula's private context pointer for this instance (Job) of this plugin.

3.4.2 freePlugin(bpContext *ctx)

This entry point is called when the this instance of the plugin is no longer needed (the Job is ending), and the plugin should release all memory it may have allocated for this particular instance (Job) i.e. the pContext. This is not the final termination of the plugin signaled by a call to **unloadPlugin**. Any other instances (Job) will continue to run, and the entry point **newPlugin** may be called again if other jobs start.



3.4.3 getPluginValue(bpContext *ctx, pVariable var, void *value)

Bacula will call this entry point to get a value from the plugin. This entry point is currently not called.

3.4.4 setPluginValue(bpContext *ctx, pVariable var, void *value)

Bacula will call this entry point to set a value in the plugin. This entry point is currently not called.

3.4.5 handlePluginEvent(bpContext *ctx, bEvent *event, void *value)

This entry point is called when Bacula encounters certain events (discussed below). This is, in fact, the main way that most plugins get control when a Job runs and how they know what is happening in the job. It can be likened to the **RunScript** feature that calls external programs and scripts. When the plugin is called, Bacula passes it the pointer to an event structure (bEvent), which currently has one item, the eventType:

```
typedef struct s_bEvent {
    uint32_t eventType;
} bEvent;
```

which defines what event has been triggered, and for each event, Bacula will pass a pointer to a value associated with that event. If no value is associated with a particular event, Bacula will pass a NULL pointer, so the plugin must be careful to always check value pointer prior to dereferencing it.

The current list of events are:

```
typedef enum {
    bEventJobStart                = 1,
    bEventJobEnd                  = 2,
    bEventStartBackupJob          = 3,
    bEventEndBackupJob            = 4,
    bEventStartRestoreJob         = 5,
    bEventEndRestoreJob           = 6,
    bEventStartVerifyJob          = 7,
    bEventEndVerifyJob            = 8,
    bEventBackupCommand           = 9,
    bEventRestoreCommand          = 10,
    bEventLevel                   = 11,
    bEventSince                   = 12,
    bEventCancelCommand           = 13, /* Executed by another thread */

    /* Just before bEventVssPrepareSnapshot */
    bEventVssBackupAddComponents  = 14,

    bEventVssRestoreLoadComponentMetadata = 15,
    bEventVssRestoreSetComponentsSelected = 16,
    bEventRestoreObject           = 17,
    bEventEndFileSet              = 18,
    bEventPluginCommand           = 19,
    bEventVssBeforeCloseRestore   = 21,

    /* Add drives to VSS snapshot
     * argument: char[27] drivelist
     * You need to add them without duplicates,
     * see fd_common.h add_drive() copy_drives() to get help
     */
    bEventVssPrepareSnapshot      = 22,
    bEventOptionPlugin            = 23,
    bEventHandleBackupFile        = 24 /* Used with Options Plugin */
}
```



```
} bEventType;
```

Most of the above are self-explanatory.

`bEventJobStart` is called whenever a Job starts. The value passed is a pointer to a string that contains: "Jobid=nnn Job=job-name". Where nnn will be replaced by the Jobid and job-name will be replaced by the Job name. The variable is temporary so if you need the values, you must copy them.

`bEventJobEnd` is called whenever a Job ends. No value is passed.

`bEventStartBackupJob` is called when a Backup Job begins. No value is passed.

`bEventEndBackupJob` is called when a Backup Job ends. No value is passed.

`bEventStartRestoreJob` is called when a Restore Job starts. No value is passed.

`bEventEndRestoreJob` is called when a Restore Job ends. No value is passed.

`bEventStartVerifyJob` is called when a Verify Job starts. No value is passed.

`bEventEndVerifyJob` is called when a Verify Job ends. No value is passed.

`bEventBackupCommand` is called prior to the `bEventStartBackupJob` and the plugin is passed the command string (everything after the equal sign in "Plugin =" as the value.

Note, if you intend to backup a file, this is an important first point to write code that copies the command string passed into your `pContext` area so that you will know that a backup is being performed and you will know the full contents of the "Plugin =" command (i.e. what to backup and what virtual filename the user wants to call it.

`bEventRestoreCommand` is called prior to the `bEventStartRestoreJob` and the plugin is passed the command string (everything after the equal sign in "Plugin =" as the value.

See the notes above concerning backup and the command string. This is the point at which Bacula passes you the original command string that was specified during the backup, so you will want to save it in your `pContext` area for later use when Bacula calls the plugin again.

`bEventLevel` is called when the level is set for a new Job. The value is a 32 bit integer stored in the `void*`, which represents the Job Level code.

`bEventSince` is called when the since time is set for a new Job. The value is a `time_t` time at which the last job was run.

`bEventCancelCommand` is called whenever the currently running Job is cancelled. Be warned that this event is sent by a different thread.

`bEventVssBackupAddComponents`

`bEventPluginCommand` is called for each `PluginCommand` present in the current `FileSet`. The event will be sent only on plugin specified in the command. The argument is the `Plugin-Command` (not valid after the call).

`bEventHandleBackupFile` is called for each file of a `FileSet` when using a `Options Plugin`. If the plugin returns `CF_OK`, it will be used for the backup, if it returns `CF_SKIP`, the file will be skipped. Anything else will backup the file with Bacula core functions.

During each of the above calls, the plugin receives either no specific value or only one value, which in some cases may not be sufficient. However, knowing the context of the event, the plugin can call back to the Bacula entry points it was passed during the **loadPlugin** call and get to a number of Bacula variables. (at the current time few Bacula variables are implemented, but it easily extended at a future time and as needs require).



3.4.6 startBackupFile(bpContext *ctx, struct save_pkt *sp)

This entry point is called only if your plugin is a command plugin, and it is called when Bacula encounters the "Plugin =" directive in the Include section of the FileSet. Called when beginning the backup of a file. Here Bacula provides you with a pointer to the **save_pkt** structure and you must fill in this packet with the "attribute" data of the file.

```

struct save_pkt {
    int32_t pkt_size;           /* size of this packet */
    char *fname;               /* Full path and filename */
    char *link;                /* Link name if any */
    struct stat statp;         /* System stat() packet for file */
    int32_t type;              /* FT_xx for this file */
    uint32_t flags;            /* Bacula internal flags */
    bool portable;            /* set if data format is portable */
    char *cmd;                 /* command */
    uint32_t delta_seq;        /* Delta sequence number */
    char *object_name;         /* Object name to create */
    char *object;              /* restore object data to save */
    int32_t object_len;        /* restore object length */
    int32_t index;            /* restore object index */
    int32_t pkt_end;          /* end packet sentinel */
};

```

The second argument is a pointer to the **save_pkt** structure for the file to be backed up. The plugin is responsible for filling in all the fields of the **save_pkt**. If you are backing up a real file, then generally, the **statp** structure can be filled in by doing a **stat** system call on the file.

If you are backing up a database or something that is more complex, you might want to create a virtual file. That is a file that does not actually exist on the filesystem, but represents say an object that you are backing up. In that case, you need to ensure that the **fname** string that you pass back is unique so that it does not conflict with a real file on the system, and you need to artificially create values in the **statp** packet.

Example programs such as **bpipe-fd.c** show how to set these fields. You must take care not to store pointers the stack in the pointer fields such as **fname** and **link**, because when you return from your function, your stack entries will be destroyed. The solution in that case is to **malloc()** and return the pointer to it. In order to not have memory leaks, you should store a pointer to all memory allocated in your **pContext** structure so that in subsequent calls or at termination, you can release it back to the system.

Once the backup has begun, Bacula will call your plugin at the **pluginIO** entry point to "read" the data to be backed up. Please see the **bpipe-fd.c** plugin for how to do I/O.

Example of filling in the **save_pkt** as used in **bpipe-fd.c**:

```

struct plugin_ctx *p_ctx = (struct plugin_ctx *)ctx->pContext;
time_t now = time(NULL);
sp->fname = p_ctx->fname;
sp->statp.st_mode = 0700 | S_IFREG;
sp->statp.st_ctime = now;
sp->statp.st_mtime = now;
sp->statp.st_atime = now;
sp->statp.st_size = -1;
sp->statp.st_blksize = 4096;
sp->statp.st_blocks = 1;
p_ctx->backup = true;
return bRC_OK;

```

Note: the filename to be created has already been created from the command string previously sent to the plugin and is in the plugin context (**p_ctx->fname**) and is a **malloc()**ed string. This example creates a regular file (**S_IFREG**), with various fields being created.

In general, the sequence of commands issued from Bacula to the plugin to do a backup while processing the "Plugin =" directive are:



- 1 generate a `bEventBackupCommand` event to the specified plugin and pass it the command string.
- 2 make a `startPluginBackup` call to the plugin, which fills in the data needed in `save_pkt` to save as the file attributes and to put on the Volume and in the catalog.
- 3 call Bacula's internal `save_file()` subroutine to save the specified file. The plugin will then be called at `pluginIO()` to "open" the file, and then to read the file data. Note, if you are dealing with a virtual file, the "open" operation is something the plugin does internally and it doesn't necessarily mean opening a file on the filesystem. For example in the case of the `bpipe-fd.c` program, it initiates a pipe to the requested program. Finally when the plugin signals to Bacula that all the data was read, Bacula will call the plugin with the "close" `pluginIO()` function.

3.4.7 `endBackupFile(bpContext *ctx)`

Called at the end of backing up a file for a command plugin. If the plugin's work is done, it should return `bRC_OK`. If the plugin wishes to create another file and back it up, then it must return `bRC_More` (not yet implemented). This is probably a good time to release any `malloc()`ed memory you used to pass back filenames.

3.4.8 `startRestoreFile(bpContext *ctx, const char *cmd)`

Called when the first record is read from the Volume that was previously written by the command plugin.

3.4.9 `createFile(bpContext *ctx, struct restore_pkt *rp)`

Called for a command plugin to create a file during a Restore job before restoring the data. This entry point is called before any I/O is done on the file. After this call, Bacula will call `pluginIO()` to open the file for write.

The data in the `restore_pkt` is passed to the plugin and is based on the data that was originally given by the plugin during the backup and the current user restore settings (e.g. where, `RegexWhere`, `replace`). This allows the plugin to first create a file (if necessary) so that the data can be transmitted to it. The next call to the plugin will be a `pluginIO` command with a request to open the file write-only.

This call must return one of the following values:

```
enum {
    CF_SKIP = 1,          /* skip file (not newer or something) */
    CF_ERROR,             /* error creating file */
    CF_EXTRACT,           /* file created, data to extract */
    CF_CREATED,           /* file created, no data to extract */
    CF_CORE               /* let bacula core handles the file creation */
};
```

in the `restore_pkt` value `create_status`. For a normal file, unless there is an error, you must return `CF_EXTRACT`.

```
struct restore_pkt {
    int32_t pkt_size;      /* size of this packet */
    int32_t stream;        /* attribute stream id */
    int32_t data_stream;   /* id of data stream to follow */
    int32_t type;          /* file type FT */
};
```



```
int32_t file_index;           /* file index */
int32_t LinkFI;               /* file index to data if hard link */
uid_t uid;                    /* userid */
struct stat statp;             /* decoded stat packet */
const char *attrEx;           /* extended attributes if any */
const char *ofname;           /* output filename */
const char *olname;           /* output link name */
const char *where;            /* where */
const char *RegexWhere;       /* regex where */
int replace;                   /* replace flag */
int create_status;             /* status from createFile() */
int32_t pkt_end;              /* end packet sentinel */

};
```

Typical code to create a regular file would be the following:

```
struct plugin_ctx *p_ctx = (struct plugin_ctx *)ctx->pContext;
time_t now = time(NULL);
sp->fname = p_ctx->fname; /* set the full path/filename I want to create */
sp->type = FT_REG;
sp->statp.st_mode = 0700 | S_IFREG;
sp->statp.st_ctime = now;
sp->statp.st_mtime = now;
sp->statp.st_atime = now;
sp->statp.st_size = -1;
sp->statp.st_blksize = 4096;
sp->statp.st_blocks = 1;
return bRC_OK;
```

This will create a virtual file. If you are creating a file that actually exists, you will most likely want to fill the statp packet using the stat() system call.

Creating a directory is similar, but requires a few extra steps:

```
struct plugin_ctx *p_ctx = (struct plugin_ctx *)ctx->pContext;
time_t now = time(NULL);
sp->fname = p_ctx->fname; /* set the full path I want to create */
sp->link = xxx; where xxx is p_ctx->fname with a trailing forward slash
sp->type = FT_DIREND
sp->statp.st_mode = 0700 | S_IFDIR;
sp->statp.st_ctime = now;
sp->statp.st_mtime = now;
sp->statp.st_atime = now;
sp->statp.st_size = -1;
sp->statp.st_blksize = 4096;
sp->statp.st_blocks = 1;
return bRC_OK;
```

The link field must be set with the full cononical path name, which always ends with a forward slash. If you do not terminate it with a forward slash, you will surely have problems later.

As with the example that creates a file, if you are backing up a real directory, you will want to do an stat() on the directory.

Note, if you want the directory permissions and times to be correctly restored, you must create the directory **after** all the file directories have been sent to Bacula. That allows the restore process to restore all the files in a directory using default directory options, then at the end, restore the directory permissions. If you do it the other way around, each time you restore a file, the OS will modify the time values for the directory entry.

3.4.10 setFileAttributes(bpContext *ctx, struct restore_pkt *rp)

This is call not yet implemented. Called for a command plugin.



See the definition of **restre_pkt** in the above section.

3.4.11 endRestoreFile(bpContext *ctx)

Called when a command plugin is done restoring a file.

3.4.12 pluginIO(bpContext *ctx, struct io_pkt *io)

Called to do the input (backup) or output (restore) of data from or to a file for a command plugin. These routines simulate the Unix read(), write(), open(), close(), and lseek() I/O calls, and the arguments are passed in the packet and the return values are also placed in the packet. In addition for Win32 systems the plugin must return two additional values (described below).

```
enum {
    IO_OPEN = 1,
    IO_READ = 2,
    IO_WRITE = 3,
    IO_CLOSE = 4,
    IO_SEEK = 5
};

struct io_pkt {
    int32_t pkt_size;           /* Size of this packet */
    int32_t func;               /* Function code */
    int32_t count;              /* read/write count */
    mode_t mode;                /* permissions for created files */
    int32_t flags;              /* Open flags */
    char *buf;                  /* read/write buffer */
    const char *fname;          /* open filename */
    int32_t status;              /* return status */
    int32_t io_errno;           /* errno code */
    int32_t lerror;             /* Win32 error code */
    int32_t whence;             /* lseek argument */
    boffset_t offset;           /* lseek argument */
    bool win32;                 /* Win32 GetLastError returned */
    int32_t pkt_end;            /* end packet sentinel */
};
```

The particular Unix function being simulated is indicated by the **func**, which will have one of the IO_OPEN, IO_READ, ... codes listed above. The status code that would be returned from a Unix call is returned in **status** for IO_OPEN, IO_CLOSE, IO_READ, and IO_WRITE. The return value for IO_SEEK is returned in **offset** which in general is a 64 bit value.

When there is an error on Unix systems, you must always set io_error, and on a Win32 system, you must always set win32, and the returned value from the OS call GetLastError() in lerror.

For all except IO_SEEK, **status** is the return result. In general it is a positive integer unless there is an error in which case it is -1.

The following describes each call and what you get and what you should return:

IO_OPEN You will be passed fname, mode, and flags. You must set on return: status, and if there is a Unix error io_errno must be set to the errno value, and if there is a Win32 error win32 and lerror.

IO_READ You will be passed: count, and buf (buffer of size count). You must set on return: status to the number of bytes read into the buffer (buf) or -1 on an error, and if there is a Unix error io_errno must be set to the errno value, and if there is a Win32 error, win32 and lerror must be set.



IO_WRITE You will be passed: `count`, and `buf` (buffer of size `count`). You must set on return: `status` to the number of bytes written from the buffer (`buf`) or `-1` on an error, and if there is a Unix error `io_errno` must be set to the `errno` value, and if there is a Win32 error, `win32` and `lerror` must be set.

IO_CLOSE Nothing will be passed to you. On return you must set `status` to `0` on success and `-1` on failure. If there is a Unix error `io_errno` must be set to the `errno` value, and if there is a Win32 error, `win32` and `lerror` must be set.

IO_LSEEK You will be passed: `offset`, and `whence`. `offset` is a 64 bit value and is the position to seek to relative to `whence`. `whence` is one of the following `SEEK_SET`, `SEEK_CUR`, or `SEEK_END` indicating to either to seek to an absolute position, relative to the current position or relative to the end of the file. You must pass back in `offset` the absolute location to which you seeked. If there is an error, `offset` should be set to `-1`. If there is a Unix error `io_errno` must be set to the `errno` value, and if there is a Win32 error, `win32` and `lerror` must be set.

Note: Bacula will call `IO_SEEK` only when writing a sparse file.

3.4.13 `bool checkFile(bpContext *ctx, char *fname)`

If this entry point is set, Bacula will call it after backing up all file data during an Accurate backup. It will be passed the full filename for each file that Bacula is proposing to mark as deleted. Only files previously backed up but not backed up in the current session will be marked to be deleted. If you return **false**, the file will be marked deleted. If you return **true** the file will not be marked deleted. This permits a plugin to ensure that previously saved virtual files or files controlled by your plugin that have not change (not backed up in the current job) are not marked to be deleted. This entry point will only be called during Accurate Incremental and Differential backup jobs.

3.5 Bacula Plugin Entrypoints

When Bacula calls one of your plugin entrypoints, you can call back to the entrypoints in Bacula that were supplied during the xxx plugin call to get or set information within Bacula.

3.5.1 `bRC registerBaculaEvents(bpContext *ctx, ...)`

This Bacula entrypoint will allow you to register to receive events that are not automatically passed to your plugin by default. This entrypoint currently is unimplemented.

3.5.2 `bRC getBaculaValue(bpContext *ctx, bVariable var, void *value)`

Calling this entrypoint, you can obtain specific values that are available in Bacula. The following Variables can be referenced:

- `bVarJobId` returns an `int`
- `bVarFDName` returns a `char *`
- `bVarLevel` returns an `int`
- `bVarClient` returns a `char *`
- `bVarJobName` returns a `char *`
- `bVarJobStatus` returns an `int`



- `bVarSinceTime` returns an `int` (`time_t`)
- `bVarAccurate` returns an `int`

3.5.3 `bRC setBaculaValue(bpContext *ctx, bVariable var, void *value)`

Calling this entrypoint allows you to set particular values in Bacula. The only variable that can currently be set is `bVarFileSeen` and the value passed is a `char *` that points to the full filename for a file that you are indicating has been seen and hence is not deleted.

3.5.4 `bRC JobMessage(bpContext *ctx, const char *file, int line, int type, utime_t mtime, const char *fmt, ...)`

This call permits you to put a message in the Job Report.

3.5.5 `bRC DebugMessage(bpContext *ctx, const char *file, int line, int level, const char *fmt, ...)`

This call permits you to print a debug message.

3.5.6 `void baculaMalloc(bpContext *ctx, const char *file, int line, size_t size)`

This call permits you to obtain memory from Bacula's memory allocator.

3.5.7 `void baculaFree(bpContext *ctx, const char *file, int line, void *mem)`

This call permits you to free memory obtained from Bacula's memory allocator.

3.6 Building Bacula Plugins

There is currently one sample program `example-plugin-fd.c` and one working plugin `bpipe-fd.c` that can be found in the Bacula `src/plugins/fd` directory. Both are built with the following:

```
cd <bacula-source>
./configure <your-options>
make
...
cd src/plugins/fd
make
make test
```

After building Bacula and changing into the `src/plugins/fd` directory, the `make` command will build the `bpipe-fd.so` plugin, which is a very useful and working program.

The `make test` command will build the `example-plugin-fd.so` plugin and a binary named `main`, which is build from the source code located in `src/fd/fd_plugins.c`.



If you execute `./main`, it will load and run the `example-plugin-fd` plugin simulating a small number of the calling sequences that Bacula uses in calling a real plugin. This allows you to do initial testing of your plugin prior to trying it with Bacula.

You can get a good idea of how to write your own plugin by first studying the `example-plugin-fd`, and actually running it. Then it can also be instructive to read the `bpipe-fd.c` code as it is a real plugin, which is still rather simple and small.

When actually writing your own plugin, you may use the `example-plugin-fd.c` code as a template for your code.

3.7 Advanced Restore Options

Some plugins can be configured at the restore time with “Plugin Restore Objects” stored in the catalog. It is possible to list these objects and display them for a list job jobids with the following command:

```
# /opt/bacula/bin/bconsole
*.bvfs_get_jobids client=127.0.0.1-fd
10,11,12
*llist pluginrestoreconf jobid=10,11,12
    jobid: 10
    restoreobjectid: 15
    objectname: RestoreOptions
    pluginname: bpipe:/@bpipe@/encrypt-bug.jpg:cat /tmp/encrypt-bug.jpg:cat >/tmp/encrypt-bug.jpg
    objecttype: 27

*list pluginrestoreconf jobid=10,11,12 id=15
# Plugin configuration file
# Version 1
OptPrompt="Restore command to use"
restore_command=@STR@
```

In this example, the `bpipe` plugin can be configured at restore time to overwrite the default restore command with a string (`@STR@`).

From a script it is possible to use the “Plugin Restore Option” menu at the restore step, or to submit a file with the appropriate configuration. For the `bpipe` plugin, the “`restore_command`” can be configured at the restore time. With the `bpipe` plugin, the “Plugin Restore Option” file would look like:

```
# cat /tmp/restore.opts
restore_command="cat > /tmp/a.jpg"
```

The file should be uploaded to the director during the restore session and used in the restore command via the `pluginrestoreconf` option. The file transfer and the restore command should be done in the same `bconsole` session.

```
# /opt/bacula/bin/bconsole
*@putfile akey /tmp/restore.opts
OK
*restore pluginrestoreconf="15:akey"
```

The “`pluginrestoreconf`” option of the restore command needs two parameters:

- `RestoreObjectId` displayed in the “`llist pluginrestoreconf`” command
- The Key of the file uploaded to the Director with the `@putfile` command

The “`RestoreObjectId`” is used to check the validity of the options provided by the user.



3.8 Bacula Auth Plugin Documentation

3.8.1 Overview

Large companies are using central systems to handle authentication and authorization data. It is very often based on LDAP databases. With one click, the access to all software in the company can be granted or revoked. Today, once a user can access the `bconsole.conf` file, Bacula doesn't require any other form of authentication, basically, any user that have access to `bconsole.conf` can interact with Bacula. It is mandatory to change the Console resource and reload the Director to disable a user. The idea would be to let the administrator the possibility to authenticate users with a central database. For that, this document propose to design an Authorization/Authentication mechanism based on a Director plugin. Once the Console is properly connected, the plugin would be able to authenticate a given user (password/user), and in a second time, would be also able to manage ACLs.

3.8.2 Dictionary

Authentication is the act of proving an assertion, such as the identity of a computer system user. In contrast with identification, the act of indicating a person or thing's identity, authentication is the process of verifying that identity. It might involve validating with personal password.

Authorization is the function of specifying access rights/privileges to resources, which is related to information security and computer security in general and to access control in particular. More formally, "to authorize" is to define an access policy.

Multi-factor authentication (or MFA) is an authentication method in which a computer user is granted access only after successfully presenting two or more pieces of evidence (or factors) to an authentication mechanism: knowledge (something the user and only the user knows), possession (something the user and only the user has), and inherence (something the user and only the user is).

Two-Factor Authentication (or 2FA) is a subset of MFA described above with only two different factors, i.e. something they know, something they have, or something they are.

3.8.3 Bacula DIR Plugin API

To write a Bacula plugin, you create a dynamic shared object program (or dll on Win32) with a particular name and two exported entry points, place it in the Plugins Directory, which is defined in the `bacula-dir.conf` file in the **Director** resource, and when the Director starts, it will load all the plugins that end with `-dir.so` (or `-dir.dll` on Win32) found in that directory.

Loading Plugins

Once the Director loads the plugins, it asks the dynamic loader for the two entry points (`loadPlugin` and `unloadPlugin`) then calls the `loadPlugin` entry point (see below). Bacula passes information to the plugin through this call and it gets back information that it needs to use the plugin. Later, Bacula will call particular functions that are defined by the `loadPlugin` interface. When Bacula is finished with the plugin (when Bacula is going to exit), it will call the `unloadPlugin` entry point.

The two entry points are:

```
bRC loadPlugin(bDirInfo *lbinfo, bDirFuncs *lbfuns, pDirInfo **pinfo, pDirFuncs **pfuns)
```

and



bRC unloadPlugin()

both these external entry points to the shared object are defined as C entry points to avoid name mangling complications with C++. However, the shared object can actually be written in any language (preferably C or C++) providing that it follows C language calling conventions.

The definitions for bRC and the arguments are `src/dird/dir_plugins.h` and so this header file needs to be included in your plugin. It along with `src/lib/plugins.h` define basically the whole plugin interface. Within this header file, it includes the following files:

```
#include <sys/types.h>
#include "config.h"
#include "bc_types.h"
#include "lib/plugins.h"
```

loadPlugin

As previously mentioned, the `loadPlugin` entry point in the plugin is called immediately after Bacula loads the plugin when the Director itself is first starting. This entry point is only called once during the execution of the Director. In calling the plugin, the first two arguments are information from Bacula that is passed to the plugin, and the last two arguments are information about the plugin that the plugin must return to Bacula. The call is:

```
bRC loadPlugin(bDirInfo *lbinfo, bDirFuncs *lbfuncs, pDirInfo **pinfo, pDirFuncs **pfuncs)
```

and the arguments are:

lbinfo This is information about Bacula in general. Currently, the only value defined in the `bInfo` structure is the version, which is the Bacula plugin interface version, currently defined as 1. The size is set to the byte size of the structure. The exact definition of the `bInfo` structure as of this writing is:

```
typedef struct s_dirbaculaInfo {
    uint32_t size;
    uint32_t version;
} bDirInfo;
```

lbfuncs The `bFuncs` structure defines the callback entry points within Bacula that the plugin can use register events, get Bacula values, set Bacula values, and send messages to the Job output or debug output. The exact definition as of this writing is:

```
typedef struct s_dirbaculaFuncs {
    uint32_t size;
    uint32_t version;
    bRC (*registerBaculaEvents)(bpContext *ctx, ...);
    bRC (*getBaculaValue)(bpContext *ctx, brDirVariable var, void *value);
    bRC (*setBaculaValue)(bpContext *ctx, bwDirVariable var, void *value);
    bRC (*JobMessage)(bpContext *ctx, const char *file, int line,
        int type, utime_t mtime, const char *fmt, ...);
    bRC (*DebugMessage)(bpContext *ctx, const char *file, int line,
        int level, const char *fmt, ...);
} bDirFuncs;
```

We will discuss these entry points and how to use them a bit later when describing the plugin code.



plInfo When the `loadPlugin` entry point is called, the plugin must initialize an information structure about the plugin and return a pointer to this structure to Bacula. The exact definition as of this writing is:

```
typedef struct s_dirpluginInfo {
    uint32_t size;
    uint32_t version;
    const char *plugin_magic;
    const char *plugin_license;
    const char *plugin_author;
    const char *plugin_date;
    const char *plugin_version;
    const char *plugin_description;
} pDirInfo;
```

Where:

- `version` is the current Bacula defined plugin interface version, currently set to 1. If the interface version differs from the current version of Bacula, the plugin will not be run (not yet implemented).
- `plugin_magic` is a pointer to the text string `DirPluginData` defined in `DIR_PLUGIN_MAGIC`, a sort of sanity check. If this value is not specified, the plugin will not be run (not yet implemented).
- `plugin_license` is a pointer to a text string that describes the plugin license. Bacula will only accept compatible licenses. The accepted licenses as of this writing are: "Bacula AGPLv3", "AGPLv3", "Bacula Systems(R) SA" and defined in `BPLUGIN_LICENSE` compile time definition.
- `plugin_author` is a pointer to the text name of the author of the plugin. This string can be anything but is generally the author's name.
- `plugin_date` is the pointer to a text string containing the date of the plugin. This string can be anything but is generally some human readable form of the date.
- `plugin_version` is a pointer to a text string containing the version of the plugin. The contents are determined by the plugin writer.
- `plugin_description` is a pointer to a text string describing what the plugin does. The contents are determined by the plugin writer.

The `plInfo` structure must be defined in static memory because Bacula does not copy it and may refer to the values at any time while the plugin is loaded. All values must be supplied or the plugin will not run (not yet implemented). All text strings must be either ASCII or UTF-8 strings that are terminated with a zero (`\nul`) byte.

pFuncs When the `loadPlugin` entry point is called, the plugin must initialize an entry point structure about the plugin and return a pointer to this structure to Bacula. This structure contains pointer to each of the entry points that the plugin must (or will) provide for Bacula. When Bacula is actually running the plugin, it will call the defined entry points at particular times. The `pFuncs` structure must be defined in static memory because Bacula does not copy it and may refer to the values at any time while the plugin is loaded.

The exact definition as of this writing is:

```
typedef struct s_dirpluginFuncs
{
    uint32_t size;
    uint32_t version;
    bRC (*newPlugin)(bpContext *ctx);
    bRC (*freePlugin)(bpContext *ctx);
    bRC (*getPluginValue)(bpContext *ctx, pDirVariable var, void *value);
    bRC (*setPluginValue)(bpContext *ctx, pDirVariable var, void *value);
    bRC (*handlePluginEvent)(bpContext *ctx, bDirEvent *event, void *value);
    bRC (*getPluginAuthenticationData)(bpContext *ctx, const char *param, void **data);
}
```



```
bRC (*getPluginAuthorizationData)(bpContext *ctx, const char *param, void **data);  
} pDirFuncs;
```

The details of the entry points will be presented in separate sections below. Where:

size is the byte size of the structure.

version is the plugin interface version currently set to 1 (defined in DIR_PLUGIN_INTERFACE_VERSION).

*

unloadPlugin

As previously mentioned, the unloadPlugin entry point in the plugin is called just before a Director finish execution. This entry point is responsible for releasing any resources allocated by plugin during loadPlugin.

Plugin Entry Points

This section will describe each of the entry points (subroutines) within the plugin that the plugin must provide for Bacula, when they are called and their arguments. As noted above, pointers to these subroutines are passed back to Bacula in the pFuncs structure when Bacula calls the loadPlugin() externally defined entry point.

newPlugin

This is the entry point that Bacula will call when a new instance of the plugin is created. This typically happens at the beginning of any Job or Console connection. If 10 Jobs are running simultaneously, there will be at least 10 instances of the plugin.

The bpContext structure will be passed to the plugin, and during this call, if the plugin needs to have its private working storage that is associated with the particular instance of the plugin, it should create it from the heap (malloc the memory) and store a pointer to its private working storage in the pContext variable.

Note: since Bacula is a multi-threaded program, you must not keep any variable data in your plugin unless it is truly meant to apply globally to the whole plugin. In addition, you must be aware that except the first and last call to the plugin (loadPlugin and unloadPlugin) all the other calls will be made by threads that correspond to a Bacula job. The bpContext that will be passed for each thread will remain the same throughout the Job thus you can keep your private Job specific data in it (bContext).

```
typedef struct s_bpContext {  
    void *pContext; /* Plugin private context */  
    void *bContext; /* Bacula private context */  
} bpContext;
```

This context pointer will be passed as the first argument to all the entry points that Bacula calls within the plugin. Needless to say, the plugin should not change the bContext variable, which is Bacula's private context pointer for this instance (Job) of this plugin.



freePlugin

This entry point is called when the this instance of the plugin is no longer needed (the Job/Console is ending), and the plugin should release all memory it may have allocated for this particular instance i.e. the pContext. This is not the final termination of the plugin signaled by a call to unloadPlugin. Any other instances (Job) will continue to run, and the entry point newPlugin may be called again if other jobs start.

getPluginValue

Bacula will call this entry point to get a value from the plugin. This entry point is currently not called.

setPluginValue

Bacula will call this entry point to set a value in the plugin. This entry point is currently not called.

handlePluginEvent

This entry point is called when Bacula encounters certain events. This is, in fact, the main way that most plugins get control when a Job/Console runs and how they know what is happening in the job. It can be likened to the RunScript feature that calls external programs and scripts. When the plugin is called, Bacula passes it the pointer to an event structure (bDirEvent), which currently has one item, the eventType:

```
typedef struct s_bDirEvent {
    uint32_t eventType;
} bDirEvent;
```

which defines what event has been triggered, and for each event, Bacula will pass a pointer to a value associated with that event. If no value is associated with a particular event, Bacula will pass a NULL pointer, so the plugin must be careful to always check value pointer prior to dereferencing it.

The current list of events are:

```
typedef enum {
    bDirEventJobStart           = 1,
    bDirEventJobEnd            = 2,
    bDirEventJobInit           = 3,
    bDirEventJobRun            = 4,
    bDirEventVolumePurged      = 5,
    bDirEventNewVolume         = 6,
    bDirEventNeedVolume        = 7,
    bDirEventVolumeFull        = 8,
    bDirEventRecycle           = 9,
    bDirEventGetScratch         = 10,
    bDirEventAuthenticationQuestion = 1000, // *value is a bDirAuthValue struct allocated
                                           // to get return value from
    bDirEventAuthenticationResponse = 1001, // *value is a char* to user response
    bDirEventAuthenticate      = 1002, // return BRC_OK when authenticate is succes
} bDirEventsType;
```



Most of the above are self-explanatory.

`bEventJobStart` is called whenever a Job starts.

`bEventJobEnd` is called whenever a Job ends. No value is passed.

`bDirEventJobInit` is called when

`bDirEventJobRun` is called when

`bDirEventVolumePurged` is called when

`ntAuthenticationQuestion` is called for authentication plugin (see description below - BPAM) whenever Director wants to get the next question in user (`bconsole`) authentication session. A **value** is a pointer to `bDirAuthValue`. Plugin should return a single `bDirAuthData` structure (described below) in this value.

`ntAuthenticationResponse` is called for authentication plugin to forward user response for the last handled authentication question. A value is a pointer to `bDirAuthValue`. Plugin will get a response in `value->response` (as a `char*`, a nul terminated string) and `value->seqdata` copied from authentication question. Plugin should return `bRC_OK` if Director can proceed to the next step.

`bDirEventAuthenticate` is called for authentication plugin whenever Director asks a plugin to final authenticate result where `bRC_OK` when successful and `bRC_Error` when error.

During each of the above calls, the plugin receives either no specific value or only one value, which in some cases may not be sufficient. However, knowing the context of the event, the plugin can call back to the Bacula entry points it was passed during the `loadPlugin` call and get to a number of Bacula variables. (at the current time few Bacula variables are implemented, but it easily extended at a future time and as needs require).

3.8.4 Bacula Pluggable Authentication Modules API Framework BPAM]

Starting from Bacula Enterprise 12.6 new user authentication API framework is introduced which allows to configure a different authentication mechanisms (user credentials verification) using a dedicated Director plugins and Console resource configuration. This is called **BPAM - Bacula Pluggable Authentication Modules**.

The new framework support standard **user/password** and **MFA** authentication schemes which are fully driven by external plugins. On the client side `bconsole` when noticed will perform user interaction to collect required credentials. Bacula will still support all previous authentication schemas including CRAM-MD5 and TLS. You can even configure **TLS Authentication** together with new BPAM authentication raising required security level. BPAM authentication is available for named Console resources only.

The BPAM framework extend a standard Director Plugin API architecture with the following plugin entry points:

```
bRC getPluginAuthenticationData(bpContext *ctx, const char *param, void **data);
bRC getPluginAuthorizationData(bpContext *ctx, const char *param, void **data);
```

and plugin events mentioned above:

```
...
bDirEventAuthenticationQuestion    = 1000,    // *value is a bDirAuthValue struct allocated b
// to get return value from
bDirEventAuthenticationResponse    = 1001,    // *value is a char* to user response
bDirEventAuthenticate              = 1002,    // return bRC_OK when authenticate is successfu
...
```



BPAM Plugin registration

BPAM assumes that any authentication or authorization workflows manages a very sensitive information (user credentials or permissions) which must be handled with extreme care, i.e. it should not be visible outside the selected plugin. This makes a clear break in general Bacula's plugin workflow design where every event is forwarded to every plugin until one of them raise event handling is done. This kind of event workflow handling can leads to unexpected user credentials data breach which is unacceptable.

Before a plugin will get authentication or authorization requests it has to register its services with `getPluginAuthenticationData()` or `getPluginAuthorizationData()` plugin entry points. Director will call this plugin's functions (if defined in `pDirFuncs` structure) on every new bconsole connection for selected plugin only when appropriate Console resource is configured (see below for more info).

```
Console {
    Name = "bpamauthconsole"
    Password = "xxx"

    # New directives
    Authentication Plugin = "<plugin>:<optional parameters>"
    Authorization Plugin = "<plugin>:<optional parameters>"      # not implemented yet!
    ...
}
```

bacula-dir.conf - Console resource configuration

The entry point `getPluginAuthenticationData()` is called when Director needs to forward authentication to selected `<plugin>` which is defined with `Authentication Plugin = ...` parameter. The name `<plugin>` has to match the filename of the Director plugin (without a `-dir.*` suffix). No other plugin will be bothered. The function takes the following parameters:

`param` is a nul terminated string defined on the right side of the parameter, including plugin name and optional plugin parameters.

`data` is a pointer to `bDirAuthenticationRegister` struct pointer (`bDirAuthenticationRegister**`) which should be filled as a return value for plugin data registration.

The `bDirAuthenticationRegister` structure is defined as:

```
typedef struct s_bDirAuthenticationRegister {
    const char * name;
    const char * welcome;
    const uint32_t num;
    const bDirAuthenticationData *data;
    const int32_t nsTTL;
} bDirAuthenticationRegister;
```

which together with `bDirAuthenticationData` structure will define a full authentication workflow managed by this plugin. Where parameters are:

`name` is a name of the authentication plugin and should match the `<plugin>` string in Console resource configuration.

`welcome` is an optional authentication welcome string which will be displayed before first authentication question.



num is a number of elements in authentication operation data table pointed by data below.

data the pointer to the authentication operation data (bDirAuthenticationData) table, the first element, with a number of elements defined in num above.

nsTTL currently unimplemented, for future usage.

The structure bDirAuthenticationData is defines as:

```
typedef struct s_bDirAuthenticationData {
    const bDirAuthenticationOperation operation;
    const char * question;
    const uint32_t seqdata;
} bDirAuthenticationData;
```

where:

operation a single authentication step (operation) which should be performed by Director + bconsole application. The list of possible operations include: display user the message, get login from user, get password from user, etc.

question the nul terminated string to display to the user during credentials collection.

seqdata 32 bit integer for full plugin use; this value will be added to the user response send to the plugin.

Both bDirAuthenticationRegister and bDirAuthenticationData structures could be a statically compiled data or assembled in runtime. Director will read data from it and never change. BPAM defines the following authentication operations:

```
typedef enum {
    bDirAuthenticationOperationPlugin,
    bDirAuthenticationOperationPluginAll,
    bDirAuthenticationOperationMessage,
    bDirAuthenticationOperationPlain,
    bDirAuthenticationOperationLogin = bDirAuthenticationOperationPlain,
    bDirAuthenticationOperationHidden,
    bDirAuthenticationOperationPassword = bDirAuthenticationOperationHidden,
    bDirAuthenticationOperationAuthenticate,
} bDirAuthenticationOperation;
```

bDirAuthenticationOperationPlugin Director will get the current operation from the plugin using bDirEventAuthenticationQuestion plugin event, then it should execute the operation returned.

bDirAuthenticationOperationPluginAll Director will get all authentication operations from the plugin with bDirEventAuthenticationQuestion plugin events loop.

bDirAuthenticationOperationMessage Director will ask bconsole to display a message pointed by bDirAuthenticationData.question and will proceed to the next operation in the list.

bDirAuthenticationOperationPlain (alias bDirAuthenticationOperationLogin) Director will ask bconsole to display a message pointed by bDirAuthenticationData.question, then get plain input visible to the user and proceed to the next operation in the list.

bDirAuthenticationOperationHidden (alias bDirAuthenticationOperationPassword) Director will ask bconsole to display a message pointed by bDirAuthenticationData.question, then get hidden (*) input invisible to the user, i.e. a password, and proceed to the next operation in the list.

bDirAuthenticationOperationAuthenticate tells Director to finish all user interaction and asks plugin for authentication using bDirEventAuthenticate plugin event.



Every authentication operation of `bDirAuthenticationOperationPlain` and `bDirAuthenticationOperationHidden` will return with an user response data which will be forwarded directly to plugin using `bDirEventAuthenticationResponse` plugin event. No other operations will do that. During this event a plugin is responsible to save user response as it won't be saved by Director and lost forever. These collected responses should be used for final user authenticate. The response is passed to plugin with `bDirAuthValue.response` variable and filled with `seqdata` value from current question operation.

As soon as the number of questions (**authenticate operations**) defined during registration (the `bDirAuthenticationRegister.num` value) comes to the end or one of the operation executed is `bDirAuthenticationOperationAuthenticate` then Director will stop asking a user more questions and will generate the `bDirEventAuthenticate` plugin event. During this event handling a plugin is responsible to verify all responses and perform user authentication.

There is a special operation in BPAM: `bDirAuthenticationOperationPluginAll` which tells Director to asks in runtime the plugin for every single operation to execute. This allow to build a dynamic authentication procedure with a variable number of questions and challenges. In this case a plugin is responsible to explicitly finish a workflow with `bDirAuthenticationOperationAuthenticate`.

You can find below some examples for above structures.

```
static bDirAuthenticationData testquestions0[] =
{
    // operation; question; data;
    {bDirAuthenticationOperationLogin, "Username:", 0},
    {bDirAuthenticationOperationPassword, "Password:", 1},
};

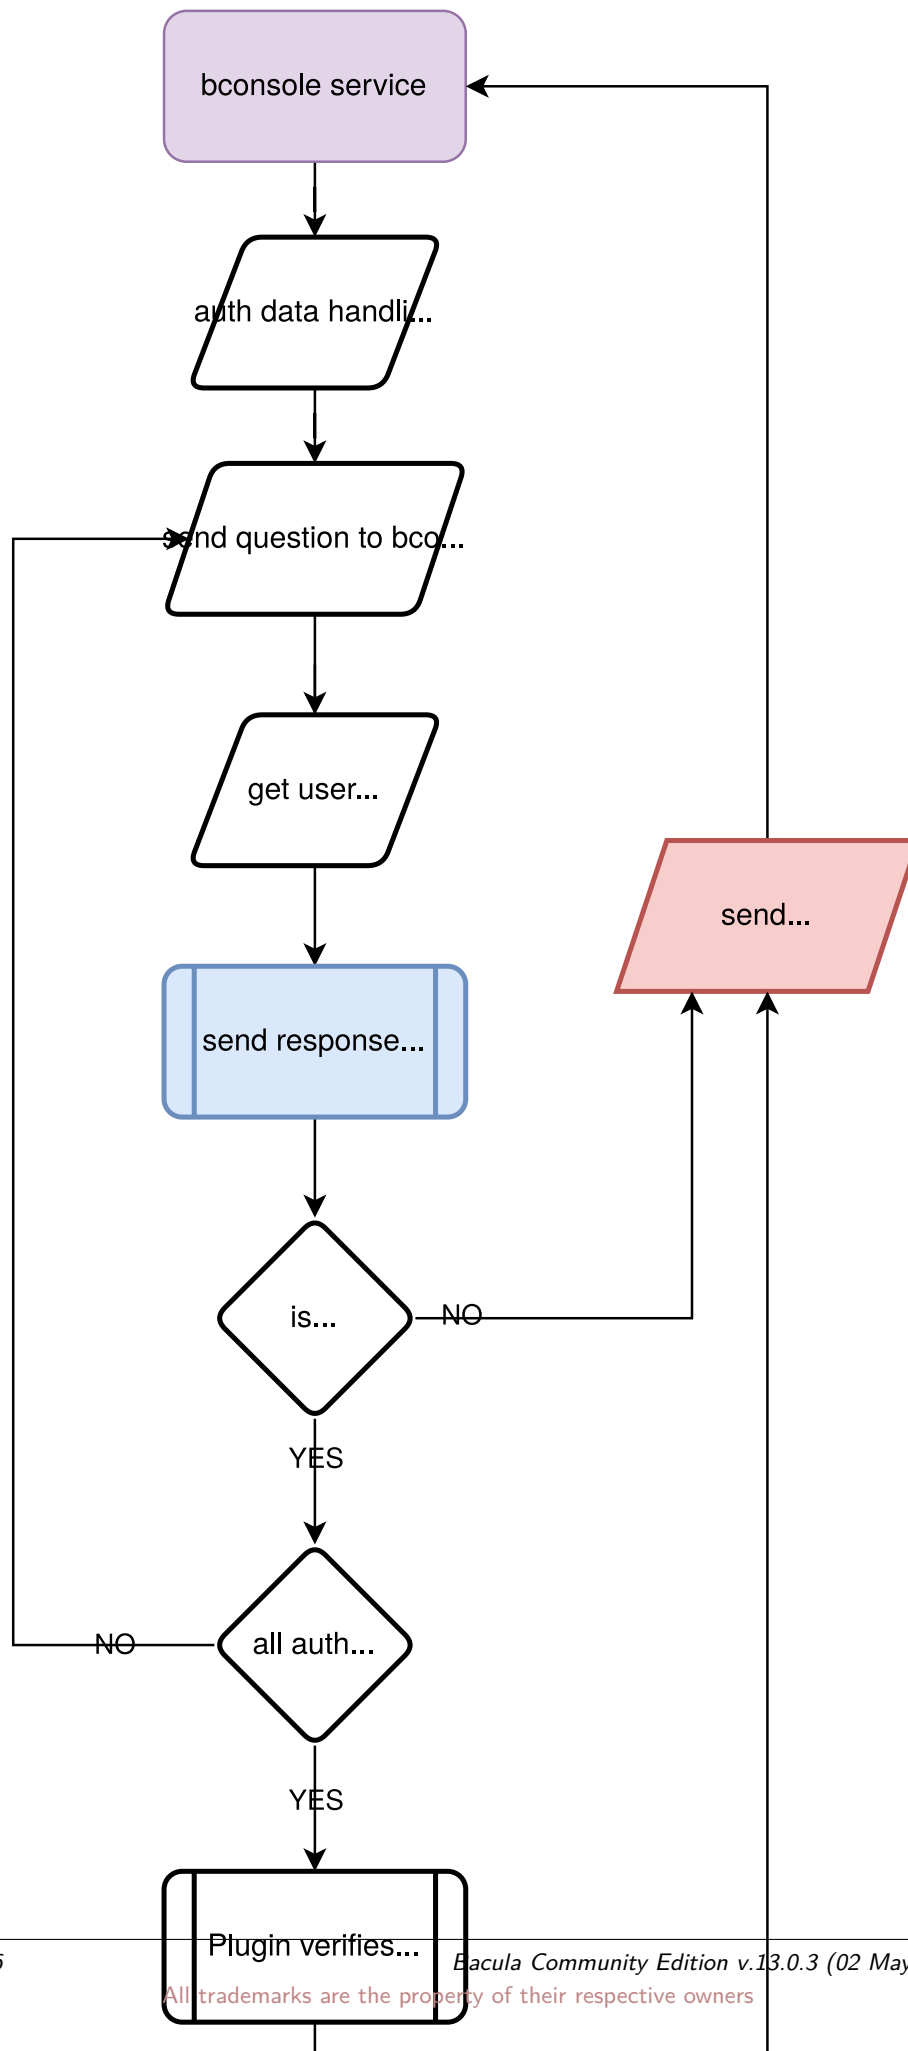
static bDirAuthenticationRegister testregister0 =
{
    .name = "PLUGIN_NAME",
    .welcome = "This is a test authplugin API Plugin. Use root/root to login.",
    .num = 2,
    .data = testquestions0,
    .nsTTL = 0,
};
```

Above you can find a simplest user/password authentication scheme with plain (visible) user input and hidden (invisible) password input.

```
static bDirAuthenticationData testquestions1[] =
{
    // operation; question; data;
    {bDirAuthenticationOperationLogin, "Username:", 0},
    {bDirAuthenticationOperationPlugin, NULL, 1},
    {bDirAuthenticationOperationPlain, "Response:", 2},
};

static bDirAuthenticationData testquestions1_msg =
{
    // operation; question; data;
    bDirAuthenticationOperationMessage, NULL, 0
};

static bDirAuthenticationRegister testregister1 =
{
    .name = "PLUGIN_NAME",
```





```
.welcome = "This is a test authplugin API Plugin. Use bacula username to login.",  
.num = 3,  
.data = testquestions1,  
.nsTTL = 0,  
};
```

Above you can find a simplest MFA (**challenge-response**) authentication scheme with out of band authentication verification. In this example a plugin will display a message prepared in runtime by plugin.





Chapter 4

Platform Support

4.1 General

This chapter describes the requirements for having a supported platform (Operating System). In general, Bacula is quite portable. It supports 32 and 64 bit architectures as well as bigendian and littleendian machines. For full support, the platform (Operating System) must implement POSIX Unix system calls. However, for File daemon support only, a small compatibility library can be written to support almost any architecture.

Currently Linux, FreeBSD, and Solaris are fully supported platforms, which means that the code has been tested on those machines and passes a full set of regression tests.

In addition, the Windows File daemon is supported on most versions of Windows, and finally, there are a number of other platforms where the File daemon (client) is known to run: NetBSD, OpenBSD, Mac OSX, SGI, ...

4.2 Requirements to become a Supported Platform

As mentioned above, in order to become a fully supported platform, it must support POSIX Unix system calls. In addition, the following requirements must be met:

- The principal developer (currently Kern) must have non-root ssh access to a test machine running the platform.
- The ideal requirements and minimum requirements for this machine are given below.
- There must be a defined platform champion who is normally a system administrator for the machine that is available. This person need not be a developer/programmer but must be familiar with system administration of the platform.
- There must be at least one person designated who will run regression tests prior to each release. Releases occur approximately once every 6 months, but can be more frequent. It takes at most a day's effort to setup the regression scripts in the beginning, and after that, they can either be run daily or on demand before a release. Running the regression scripts involves only one or two command line commands and is fully automated.
- Ideally there are one or more persons who will package each Bacula release.
- Ideally there are one or more developers who can respond to and fix platform specific bugs.

Ideal requirements for a test machine:

- The principal developer will have non-root ssh access to the test machine at all times.



- The principal developer will have a root password.
- The test machine will provide approximately 200 MB of disk space for continual use.
- The test machine will have approximately 500 MB of free disk space for temporary use.
- The test machine will run the most common version of the OS.
- The test machine will have an autochanger of DDS-4 technology or later having two or more tapes.
- The test machine will have MySQL and/or PostgreSQL database access for account "bacula" available.
- The test machine will have sftp access.
- The test machine will provide an smtp server.

Minimum requirements for a test machine:

- The principal developer will have non-root ssh access to the test machine when requested approximately once a month.
- The principal developer not have root access.
- The test machine will provide approximately 80 MB of disk space for continual use.
- The test machine will have approximately 300 MB of free disk space for temporary use.
- The test machine will run the the OS.
- The test machine will have a tape drive of DDS-4 technology or later that can be scheduled for access.
- The test machine will not have MySQL and/or PostgreSQL database access.
- The test machine will have no sftp access.
- The test machine will provide no email access.

Bare bones test machine requirements:

- The test machine is available only to a designated test person (your own machine).
- The designated test person runs the regression tests on demand.
- The test machine has a tape drive available.



Chapter 5

Daemon Protocol

5.1 General

This document describes the protocols used between the various daemons. As Bacula has developed, it has become quite out of date. The general idea still holds true, but the details of the fields for each command, and indeed the commands themselves have changed considerably.

It is intended to be a technical discussion of the general daemon protocols 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**.

5.2 Low Level Network Protocol

At the lowest level, the network protocol is handled by **BSOCK** packets which contain a lot of information about the status of the network connection: who is at the other end, etc. Each basic **Bacula** network read or write actually consists of two low level network read/writes. The first write always sends four bytes of data in machine independent byte order. If data is to follow, the first four bytes are a positive non-zero integer indicating the length of the data that follow in the subsequent write. If the four byte integer is zero or negative, it indicates a special request, a sort of network signaling capability. In this case, no data packet will follow. The low level BSOCK routines expect that only a single thread is accessing the socket at a time. It is advised that multiple threads do not read/write the same socket. If you must do this, you must provide some sort of locking mechanism. It would not be appropriate for efficiency reasons to make every call to the BSOCK routines lock and unlock the packet.

5.3 General Daemon Protocol

In general, all the daemons follow the following global rules. There may be exceptions depending on the specific case. Normally, one daemon will be sending commands to another daemon (specifically, the Director to the Storage daemon and the Director to the File daemon).

- Commands are always ASCII commands that are upper/lower case dependent as well as space sensitive.
- All binary data is converted into ASCII (either with printf statements or using base64 encoding).
- All responses to commands sent are always prefixed with a return numeric code where codes in the 1000's are reserved for the Director, the 2000's are reserved for the File



daemon, and the 3000's are reserved for the Storage daemon.

- Any response that is not prefixed with a numeric code is a command (or subcommand if you like) coming from the other end. For example, while the Director is corresponding with the Storage daemon, the Storage daemon can request Catalog services from the Director. This convention permits each side to send commands to the other daemon while simultaneously responding to commands.
- Any response that is of zero length, depending on the context, either terminates the data stream being sent or terminates command mode prior to closing the connection.
- Any response that is of negative length is a special sign that normally requires a response. For example, during data transfer from the File daemon to the Storage daemon, normally the File daemon sends continuously without intervening reads. However, periodically, the File daemon will send a packet of length -1 indicating that the current data stream is complete and that the Storage daemon should respond to the packet with an OK, ABORT JOB, PAUSE, etc. This permits the File daemon to efficiently send data while at the same time occasionally "polling" the Storage daemon for his status or any special requests. Currently, these negative lengths are specific to the daemon, but shortly, the range 0 to -999 will be standard daemon wide signals, while -1000 to -1999 will be for Director user, -2000 to -2999 for the File daemon, and -3000 to -3999 for the Storage daemon.

5.4 The Protocol Used Between the Director and the Storage Daemon

Before sending commands to the File daemon, the Director opens a Message channel with the Storage daemon, identifies itself and presents its password. If the password check is OK, the Storage daemon accepts the Director. The Director then passes the Storage daemon, the JobId to be run as well as the File daemon authorization (append, read all, or read for a specific session). The Storage daemon will then pass back to the Director a enabling key for this JobId that must be presented by the File daemon when opening the job. Until this process is complete, the Storage daemon is not available for use by File daemons.

```
SD: listens
DR: makes connection
DR: Hello <Director-name> calling <password>
SD: 3000 OK Hello
DR: JobId=nnn Allow=(append, read) Session=(*, SessionId)
    (Session not implemented yet)
SD: 3000 OK Job Authorization=<password>
DR: use device=<device-name> media_type=<media-type>
    pool_name=<pool-name> pool_type=<pool_type>
SD: 3000 OK use device
```

For the Director to be authorized, the <Director-name> and the <password> must match the values in one of the Storage daemon's Director resources (there may be several Directors that can access a single Storage daemon).

5.5 The Protocol Used Between the Director and the File Daemon

A typical conversation might look like the following:

```
FD: listens
DR: makes connection
DR: Hello <Director-name> calling <password>
```



```

FD: 2000 OK Hello
DR: JobId=nnn Authorization=<password>
FD: 2000 OK Job
DR: storage address = <Storage daemon address> port = <port-number>
    name = <DeviceName> mediatype = <MediaType>
FD: 2000 OK storage
DR: include
DR: <directory1>
DR: <directory2>
...
DR: Null packet
FD: 2000 OK include
DR: exclude
DR: <directory1>
DR: <directory2>
...
DR: Null packet
FD: 2000 OK exclude
DR: full
FD: 2000 OK full
DR: save
FD: 2000 OK save
FD: Attribute record for each file as sent to the
    Storage daemon (described above).
FD: Null packet
FD: <append close responses from Storage daemon>
    e.g.
    3000 OK Volumes = <number of volumes>
    3001 Volume = <volume-id> <start file> <start block>
        <end file> <end block> <volume session-id>
    3002 Volume data = <date/time of last write> <Number bytes written>
        <number errors>
    ... additional Volume / Volume data pairs for volumes 2 .. n
FD: Null packet
FD: close socket

```

5.6 The Save Protocol Between the File Daemon and the Storage Daemon

Once the Director has send a **save** command to the File daemon, the File daemon will contact the Storage daemon to begin the save.

In what follows: FD: refers to information set via the network from the File daemon to the Storage daemon, and SD: refers to information set from the Storage daemon to the File daemon.

5.6.1 Command and Control Information

Command and control information is exchanged in human readable ASCII commands.

```

FD: listens
SD: makes connection
FD: append open session = <JobId> [<password>]
SD: 3000 OK ticket = <number>
FD: append data <ticket-number>
SD: 3000 OK data address = <IPAddress> port = <port>

```

5.6.2 Data Information

The Data information consists of the file attributes and data to the Storage daemon. For the most part, the data information is sent one way: from the File daemon to the Storage daemon. This



allows the File daemon to transfer information as fast as possible without a lot of handshaking and network overhead.

However, from time to time, the File daemon needs to do a sort of checkpoint of the situation to ensure that everything is going well with the Storage daemon. To do so, the File daemon sends a packet with a negative length indicating that he wishes the Storage daemon to respond by sending a packet of information to the File daemon. The File daemon then waits to receive a packet from the Storage daemon before continuing.

All data sent are in binary format except for the header packet, which is in ASCII. There are two packet types used data transfer mode: a header packet, the contents of which are known to the Storage daemon, and a data packet, the contents of which are never examined by the Storage daemon.

The first data packet to the Storage daemon will be an ASCII header packet consisting of the following data.

<File-Index> <Stream-Id> <Info> where <File-Index> is a sequential number beginning from one that increments with each file (or directory) sent.

where <Stream-Id> will be 1 for the Attributes record and 2 for uncompressed File data. 3 is reserved for the MD5 signature for the file.

where <Info> transmit information about the Stream to the Storage Daemon. It is a character string field where each character has a meaning. The only character currently defined is 0 (zero), which is simply a place holder (a no op). In the future, there may be codes indicating compressed data, encrypted data, etc.

Immediately following the header packet, the Storage daemon will expect any number of data packets. The series of data packets is terminated by a zero length packet, which indicates to the Storage daemon that the next packet will be another header packet. As previously mentioned, a negative length packet is a request for the Storage daemon to temporarily enter command mode and send a reply to the File daemon. Thus an actual conversation might contain the following exchanges:

```
FD: <1 1 0> (header packet)
FD: <data packet containing file-attributes>
FD: Null packet
FD: <1 2 0>
FD: <multiple data packets containing the file data>
FD: Packet length = -1
SD: 3000 OK
FD: <2 1 0>
FD: <data packet containing file-attributes>
FD: Null packet
FD: <2 2 0>
FD: <multiple data packets containing the file data>
FD: Null packet
FD: Null packet
FD: append end session <ticket-number>
SD: 3000 OK end
FD: append close session <ticket-number>
SD: 3000 OK Volumes = <number of volumes>
SD: 3001 Volume = <volumeid> <start file> <start block>
    <end file> <end block> <volume session-id>
SD: 3002 Volume data = <date/time of last write> <Number bytes written>
    <number errors>
SD: ... additional Volume / Volume data pairs for
    volumes 2 .. n
FD: close socket
```

The information returned to the File daemon by the Storage daemon in response to the **append close session** is transmit in turn to the Director.



Chapter 6

Director Services Daemon

This chapter is intended to be a technical discussion of the Director 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 Director** services consist of the program that supervises all the backup and restore operations.

To be written . . .





Chapter 7

File Services Daemon

Please note, this section is somewhat out of date as the code has evolved significantly. The basic idea has not changed though.

This chapter is intended to be a technical discussion of the File daemon 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 File Services** consist of the programs that run on the system to be backed up and provide the interface between the Host File system and Bacula – in particular, the Director and the Storage services.

When time comes for a backup, the Director gets in touch with the File daemon on the client machine and hands it a set of “marching orders” which, if written in English, might be something like the following:

OK, **File daemon**, it's time for your daily incremental backup. I want you to get in touch with the Storage daemon on host archive.mysite.com and perform the following save operations with the designated options. You'll note that I've attached include and exclude lists and patterns you should apply when backing up the file system. As this is an incremental backup, you should save only files modified since the time you started your last backup which, as you may recall, was 2000-11-19-06:43:38. Please let me know when you're done and how it went. Thank you.

So, having been handed everything it needs to decide what to dump and where to store it, the File daemon doesn't need to have any further contact with the Director until the backup is complete providing there are no errors. If there are errors, the error messages will be delivered immediately to the Director. While the backup is proceeding, the File daemon will send the file coordinates and data for each file being backed up to the Storage daemon, which will in turn pass the file coordinates to the Director to put in the catalog.

During a **Verify** of the catalog, the situation is different, since the File daemon will have an exchange with the Director for each file, and will not contact the Storage daemon.

A **Restore** operation will be very similar to the **Backup** except that during the **Restore** the Storage daemon will not send storage coordinates to the Director since the Director presumably already has them. On the other hand, any error messages from either the Storage daemon or File daemon will normally be sent directly to the Director (this, of course, depends on how the Message resource is defined).

7.1 Commands Received from the Director for a Backup

To be written. . .



7.2 Commands Received from the Director for a Restore

To be written ...



Chapter 8

Storage Daemon Design

This chapter is intended to be a technical discussion of the Storage daemon 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**.

This document is somewhat out of date.

8.1 SD Design Introduction

The Bacula Storage daemon provides storage resources to a Bacula installation. An individual Storage daemon is associated with a physical permanent storage device (for example, a tape drive, CD writer, tape changer or jukebox, etc.), and may employ auxiliary storage resources (such as space on a hard disk file system) to increase performance and/or optimize use of the permanent storage medium.

Any number of storage daemons may be run on a given machine; each associated with an individual storage device connected to it, and BACULA operations may employ storage daemons on any number of hosts connected by a network, local or remote. The ability to employ remote storage daemons (with appropriate security measures) permits automatic off-site backup, possibly to publicly available backup repositories.

8.2 SD Development Outline

In order to provide a high performance backup and restore solution that scales to very large capacity devices and networks, the storage daemon must be able to extract as much performance from the storage device and network with which it interacts. In order to accomplish this, storage daemons will eventually have to sacrifice simplicity and painless portability in favor of techniques which improve performance. My goal in designing the storage daemon protocol and developing the initial prototype storage daemon is to provide for these additions in the future, while implementing an initial storage daemon which is very simple and portable to almost any POSIX-like environment. This original storage daemon (and its evolved descendants) can serve as a portable solution for non-demanding backup requirements (such as single servers of modest size, individual machines, or small local networks), while serving as the starting point for development of higher performance configurable derivatives which use techniques such as POSIX threads, shared memory, asynchronous I/O, buffering to high-speed intermediate media, and support for tape changers and jukeboxes.



8.3 SD Connections and Sessions

A client connects to a storage server by initiating a conventional TCP connection. The storage server accepts the connection unless its maximum number of connections has been reached or the specified host is not granted access to the storage server. Once a connection has been opened, the client may make any number of Query requests, and/or initiate (if permitted), one or more Append sessions (which transmit data to be stored by the storage daemon) and/or Read sessions (which retrieve data from the storage daemon).

Most requests and replies sent across the connection are simple ASCII strings, with status replies prefixed by a four digit status code for easier parsing. Binary data appear in blocks stored and retrieved from the storage. Any request may result in a single-line status reply of "3201 Notification pending", which indicates the client must send a "Query notification" request to retrieve one or more notifications posted to it. Once the notifications have been returned, the client may then resubmit the request which resulted in the 3201 status.

The following descriptions omit common error codes, yet to be defined, which can occur from most or many requests due to events like media errors, restarting of the storage daemon, etc. These details will be filled in, along with a comprehensive list of status codes along with which requests can produce them in an update to this document.

8.3.1 SD Append Requests

append open session = <JobId> [<Password>] A data append session is opened with the Job ID given by *JobId* with client password (if required) given by *Password*. If the session is successfully opened, a status of 3000 OK is returned with a "ticket = *number*" reply used to identify subsequent messages in the session. If too many sessions are open, or a conflicting session (for example, a read in progress when simultaneous read and append sessions are not permitted), a status of "3502 Volume busy" is returned. If no volume is mounted, or the volume mounted cannot be appended to, a status of "3503 Volume not mounted" is returned.

append data = <ticket-number> If the append data is accepted, a status of 3000 OK data address = <IPAddress> port = <port> is returned, where the *IPAddress* and *port* specify the IP address and port number of the data channel. Error status codes are 3504 Invalid ticket number and 3505 Session aborted, the latter of which indicates the entire append session has failed due to a daemon or media error.

Once the File daemon has established the connection to the data channel opened by the Storage daemon, it will transfer a header packet followed by any number of data packets. The header packet is of the form:

```
<file-index> <stream-id> <info>
```

The details are specified in the [Daemon Protocol](#) section of this document.

*append abort session = <ticket-number> The open append session with ticket *ticket-number* is aborted; any blocks not yet written to permanent media are discarded. Subsequent attempts to append data to the session will receive an error status of 3505 Session aborted.

append end session = <ticket-number> The open append session with ticket *ticket-number* is marked complete; no further blocks may be appended. The storage daemon will give priority to saving any buffered blocks from this session to permanent media as soon as possible.

append close session = <ticket-number> The append session with ticket *ticket* is closed. This message does not receive an 3000 OK reply until all of the content of the session are stored on permanent media, at which time said reply is given, followed by a list of volumes, from first to last, which contain blocks from the session, along with the first and last file and



block on each containing session data and the volume session key identifying data from that session in lines with the following format:

Volume = <Volume-id> <start-file> <start-block> <end-file> <end-block> <volume-session-id> where *Volume-id* is the volume label, *start-file* and *start-block* are the file and block containing the first data from that session on the volume, *end-file* and *end-block* are the file and block with the last data from the session on the volume and *volume-session-id* is the volume session ID for blocks from the session stored on that volume.

8.3.2 SD Read Requests

Read open session = <JobId> <Volume-id> <start-file> <start-block> <end-file> <end-block> <volume-session-id> <password> where *Volume-id* is the volume label, *start-file* and *start-block* are the file and block containing the first data from that session on the volume, *end-file* and *end-block* are the file and block with the last data from the session on the volume and *volume-session-id* is the volume session ID for blocks from the session stored on that volume.

If the session is successfully opened, a status of

3100 OK Ticket = *number*

is returned with a reply used to identify subsequent messages in the session. If too many sessions are open, or a conflicting session (for example, an append in progress when simultaneous read and append sessions are not permitted), a status of "3502 Volume busy" is returned. If no volume is mounted, or the volume mounted cannot be appended to, a status of "3503 Volume not mounted" is returned. If no block with the given volume session ID and the correct client ID number appears in the given first file and block for the volume, a status of "3505 Session not found" is returned.

Read data = <Ticket> > <Block> The specified Block of data from open read session with the specified Ticket number is returned, with a status of 3000 OK followed by a "Length = size" line giving the length in bytes of the block data which immediately follows. Blocks must be retrieved in ascending order, but blocks may be skipped. If a block number greater than the largest stored on the volume is requested, a status of "3201 End of volume" is returned. If a block number greater than the largest in the file is requested, a status of "3401 End of file" is returned.

Read close session = <Ticket> The read session with Ticket number is closed. A read session may be closed at any time; you needn't read all its blocks before closing it.

by John Walker January 30th, MM

8.4 SD Data Structures

In the Storage daemon, there is a Device resource (i.e. from conf file) that describes each physical device. When the physical device is used it is controlled by the DEVICE structure (defined in [dev.h](#)), and typically referred to as dev in the C++ code. Anyone writing or reading a physical device must ultimately get a lock on the DEVICE structure – this controls the device. However, multiple Jobs (defined by a JCR structure [src/jcr.h](#)) can be writing a physical DEVICE at the same time (of course they are sequenced by locking the DEVICE structure). There are a lot of job dependent "device" variables that may be different for each Job such as spooling (one job may spool and another may not, and when a job is spooling, it must have an i/o packet open, each job has its own record and block structures, ...), so there is a device control record or DCR that is the primary way of interfacing to the physical device. The DCR contains all the job specific data as well as a pointer to the Device resource (DEVRES structure) and the physical DEVICE structure.



Now if a job is writing to two devices (it could be writing two separate streams to the same device), it must have two DCRs. Today, the code only permits one. This won't be hard to change, but it is new code.

Today three jobs (threads), two physical devices each job writes to only one device:

```
Job1 -> DCR1 -> DEVICE1
Job2 -> DCR2 -> DEVICE1
Job3 -> DCR3 -> DEVICE2
```

To be implemented three jobs, three physical devices, but job1 is writing simultaneously to three devices:

```
Job1 -> DCR1 -> DEVICE1
      -> DCR4 -> DEVICE2
      -> DCR5 -> DEVICE3
Job2 -> DCR2 -> DEVICE1
Job3 -> DCR3 -> DEVICE2

Job = job control record
DCR = Job control data for a specific device
DEVICE = Device only control data
```




Chapter 9

Catalog Services

9.1 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, Oracle, and MySQL. Each presents certain problems with either licensing or maturity. At present, we have chosen for development purposes to use MySQL and PostgreSQL. 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.

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-postgresq` 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.

9.1.1 Filenames and Maximum Filename Length

In general, either MySQL or PostgreSQL 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.



9.1.2 Installing and Configuring MySQL

For the details of installing and configuring MySQL, please see the **Installing and Configuring MySQL** chapter (chapter 50 page 555) of the Bacula Community Edition Main manual.

9.1.3 Installing and Configuring PostgreSQL

For the details of installing and configuring PostgreSQL, please see the **Installing and Configuring PostgreSQL** chapter (chapter 51 page 561) of the Bacula Community Edition Main manual.

9.1.4 Internal Bacula Catalog

Please see the **Internal Bacula Database** chapter (chapter ?? page ??) of the Bacula Community Edition Miscellaneous Guide for more details.

9.1.5 Database Table Design

All discussions that follow pertain to the MySQL database. The details for the PostgreSQL databases are quite similar.

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.

9.2 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

9.3 Database Tables

Table 9.1: Path table layout

Column Name	Data Type	Remark
PathId	integer	Primary Key
Path	Blob	Full Path

The **Path table** contains 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.

**Table 9.2:** File table layout

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
Filename	blob	Filename Record
MarkId	integer	Used to mark files during Verify Jobs
LStat	tinyblob	File attributes in base64 encoding
MD5	tinyblob	MD5/SHA1 signature in base64 encoding
DeltaSeq	integer	Delta Sequence number

The **File table** 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 is 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.

Table 9.3: Job table layout

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
RealEndTime	datetime	Time/date when original Job ended
JobTDate	bigint	Start day in Unix format but 64 bits; used for Retention period.
VolSessionId	integer	Unique Volume Session ID

Continues on the following page



Column Name	Data Type	Remark
VolSessionTime	integer	Unique Volume Session Time
JobFiles	integer	Number of files saved in Job
JobBytes	bigint	Number of bytes saved in Job
ReadBytes	bigint	Number of bytes read 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
PrioJobId	integer	Link to prior Job Record when migrated
PurgedFiles	tiny integer	Set when all File records purged
HasBase	tiny integer	Set when Base Job run
Reviewed	tiny integer	Set when the error is acknowledged
Comment	tinyblob	Comment about this Job
PrioJob	tinyblob	Prior Job name when migrated

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:

Table 9.4: Job Types

Value	Meaning
B	Backup Job
M	Migrated Job
V	Verify Job
R	Restore Job
U	Console program (not in database)
I	Internal or system Job

Continues on the following page



Value	Meaning
D	Admin Job
A	Archive Job (not implemented)
C	Copy of a Job
c	Copy Job
g	Migration Job

Note, the Job Type values in table 9.4 on the previous page noted above are not kept in an SQL table.

The JobStatus field specifies how the job terminated, and can be one of the following:

Table 9.5: Job Statuses

Value	Meaning
C	Created but not yet running
R	Running
B	Blocked
T	Terminated normally
W	Terminated normally with warnings
E	Terminated in Error
e	Non-fatal error
f	Fatal error
D	Verify Differences
A	Canceled by the user
I	Incomplete Job
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	Waiting for Maximum jobs
t	Waiting for Start Time
p	Waiting for higher priority job to finish
i	Doing batch insert file records
a	SD despooling attributes
l	Doing data despooling
L	Committing data (last despool)

**Table 9.6:** Filesets table layout

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.

Table 9.7: JobMedia table layout

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
VollIndex	integer	The Volume use sequence number within the Job

The **JobMedia** table contains one entry at the following: start of the job, start of each new tape file, start of each new tape, end of the job. Since by default, a new tape file is written every 2GB, in general, you will have more than 2 JobMedia records per Job. The number can be varied by changing the "Maximum File Size" specified in the Device resource. This record allows Bacula to efficiently position close to (within 2GB) any given file in a backup. For restoring a full Job, these records are not very important, but if you want to retrieve a single file that was written near the end of a 100GB backup, the JobMedia records can speed it up by orders of magnitude by permitting forward spacing files and blocks rather than reading the whole 100GB backup.

Table 9.8: Media table layout

Column Name	Data Type	Remark
MediaId	integer	Primary Key
VolumeName	tinyblob	Volume name

Continues on the following page



Column Name	Data Type	Remark
Slot	integer	Autochanger Slot number or zero
PoolId	integer	Link to Pool Record
MediaType	tinyblob	The MediaType supplied by the user
MediaTypeId	integer	The MediaTypeId
LabelType	tinyint	The type of label on the Volume
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
VolParts	integer	The number of parts for a Volume (DVD)
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
Enabled	tinyint	Whether or not Volume can be written
Recycle	tinyint	Whether or not Bacula can recycle the Volumes: <yes no>
ActionOnPurge	tinyint	What happens to a Volume after purging
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
InChanger	tinyint	Whether or not Volume in autochanger
StorageId	integer	Storage record ID
DeviceId	integer	Device record ID
MediaAddressing	integer	Method of addressing media
VolReadTime	bigint	Time Reading Volume
VolWriteTime	bigint	Time Writing Volume
EndFile	integer	End File number of Volume
EndBlock	integer	End block number of Volume
LocationId	integer	Location record ID
RecycleCount	integer	Number of times recycled

Continues on the following page



Column Name	Data Type	Remark
InitialWrite	datetime	When Volume first written
ScratchPoolId	integer	Id of Scratch Pool
RecyclePoolId	integer	Pool ID where to recycle Volume
Comment	blob	User text field

The [Volume table](#)¹ 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.

Table 9.9: Pool table layout

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
MaxPoolBytes	bigint	max bytes to write on the Pool
AutoPrune	tinyint	<yes no> for autopruning
Recycle	tinyint	<yes no> for allowing auto recycling of Volume
ActionOnPurge	tinyint	Default Volume ActionOnPurge
PoolType	enum	Backup, Copy, Cloned, Archive, Migration
LabelType	tinyint	Type of label ANSI / Bacula
LabelFormat	Tinyblob	Label format
Enabled	tinyint	Whether or not Volume can be written
ScratchPoolId	integer	Id of Scratch Pool
RecyclePoolId	integer	Pool ID where to recycle Volume
NextPoolId	integer	Pool ID of next Pool
MigrationHighBytes	bigint	High water mark for migration
MigrationLowBytes	bigint	Low water mark for migration
MigrationTime	bigint	Time before migration

¹Internally referred to as the Media table



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.

Table 9.10: Client table layout

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.

Table 9.11: Storage table layout

Column Name	Data Type	Remark
StorageId	integer	Unique Id
Name	tinyblob	Resource name of Storage device
AutoChanger	tinyint	Set if it is an autochanger

The **Storage table** contains one entry for each Storage used.

Table 9.12: Counter table layout

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.

**Table 9.13:** Jobhisto table layout

Column Name	Data Type	Remark
JobId	integer	Primary Key
Job	tinyblob	Unique Job Name
Name	tinyblob	Job Name
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
RealEndTime	datetime	Time/date when original 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
ReadBytes	bigint	Number of bytes read 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
PrioJobId	integer	Link to prior Job Record when migrated
PurgedFiles	tiny integer	Set when all File records purged
HasBase	tiny integer	Set when Base Job run
Reviewed	tiny integer	Set when the error is acknowledged
Comment	tinyblob	Comment about this Job

The **JobHisto** table is the same as the Job table, but it keeps long term statistics (i.e. it is not pruned with the Job).

Table 9.14: Log Table Layout

Column Name	Data Type	Remark
LogIdId	integer	Primary Key
JobId	integer	Points to Job record
Time	datetime	Time/date log record created

Continues on the following page



Column Name	Data Type	Remark
LogText	blob	Log text

The Log table contains a log of all Job output.

Table 9.15: Location table layout

Column Name	Data Type	Remark
LocationId	integer	Primary Key
Location	tinyblob	Text defining location
Cost	integer	Relative cost of obtaining Volume
Enabled	tinyint	Whether or not Volume is enabled

The Location table defines where a Volume is physically.

Table 9.16: Locationlog table layout

Column Name	Data Type	Remark
locLogId	integer	Primary Key
Date	datetime	Time/date log record created
MediaId	integer	Points to Media record
LocationId	integer	Points to Location record
NewVolStatus	integer	enum: Full, Archive, Append, Recycle, Purged Read-only, Disabled, Error, Busy, Used, Cleaning
Enabled	tinyint	Whether or not Volume is enabled

The Location Log table contains a log of all Job output.

Table 9.17: Version table layout

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.

Table 9.18: Basefiles table layout

Column Name	Data Type	Remark
BaseId	integer	Primary Key
BaseJobId	integer	JobId of Base Job

Continues on the following page



Column Name	Data Type	Remark
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.





Chapter 10

Storage Media Output Format

10.1 General

This document describes the media format written by the Storage daemon. The Storage daemon reads and writes in units of blocks. Blocks contain records. Each block has a block header followed by records, and each record has a record header followed by record data.

This chapter is intended to be a technical discussion of the Media Format 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**.

10.2 Definitions

Block A block represents the primitive unit of information that the Storage daemon reads and writes to a physical device. Normally, for a tape device, it will be the same as a tape block. The Storage daemon always reads and writes blocks. A block consists of block header information followed by records. Clients of the Storage daemon (the File daemon) normally never see blocks. However, some of the Storage tools (bls, bscan, bextract, ...) may use block header information. In older Bacula tape versions, a block could contain records (see record definition below) from multiple jobs. However, all blocks currently written by Bacula are block level BB02, and a given block contains records for only a single job. Different jobs simply have their own private blocks that are intermingled with the other blocks from other jobs on the Volume (previously the records were intermingled within the blocks). Having only records from a single job in any given block permitted moving the VolumeSessionId and VolumeSessionTime (see below) from each record heading to the Block header. This has two advantages: 1. a block can be quickly rejected based on the contents of the header without reading all the records. 2. because there is on the average more than one record per block, less data is written to the Volume for each job.

Record A record consists of a Record Header, which is managed by the Storage daemon and Record Data, which is the data received from the Client. A record is the primitive unit of information sent to and from the Storage daemon by the Client (File daemon) programs. The details are described below.

JobId A number assigned by the Director daemon for a particular job. This number will be unique for that particular Director (Catalog). The daemons use this number to keep track of individual jobs. Within the Storage daemon, the JobId may not be unique if several Directors are accessing the Storage daemon simultaneously.

Session A Session is a concept used in the Storage daemon corresponds one to one to a Job with the exception that each session is uniquely identified within the Storage daemon by



- a unique SessionId/SessionTime pair (see below).
- VolSessionId** A unique number assigned by the Storage daemon to a particular session (Job) it is having with a File daemon. This number by itself is not unique to the given Volume, but with the VolSessionTime, it is unique.
- VolSessionTime** A unique number assigned by the Storage daemon to a particular Storage daemon execution. It is actually the Unix time_t value of when the Storage daemon began execution cast to a 32 bit unsigned integer. The combination of the **VolSessionId** and the **VolSessionTime** for a given Storage daemon is guaranteed to be unique for each Job (or session).
- FileIndex** A sequential number beginning at one assigned by the File daemon to the files within a job that are sent to the Storage daemon for backup. The Storage daemon ensures that this number is greater than zero and sequential. Note, the Storage daemon uses negative FileIndexes to flag Session Start and End Labels as well as End of Volume Labels. Thus, the combination of VolSessionId, VolSessionTime, and FileIndex uniquely identifies the records for a single file written to a Volume.
- Stream** While writing the information for any particular file to the Volume, there can be any number of distinct pieces of information about that file, e.g. the attributes, the file data, ... The Stream indicates what piece of data it is, and it is an arbitrary number assigned by the File daemon to the parts (Unix attributes, Win32 attributes, data, compressed data, ...) of a file that are sent to the Storage daemon. The Storage daemon has no knowledge of the details of a Stream; it simply represents a numbered stream of bytes. The data for a given stream may be passed to the Storage daemon in single record, or in multiple records.
- Block Header** A block header consists of a block identification ("BB02"), a block length in bytes (typically 64,512) a checksum, and sequential block number. Each block starts with a Block Header and is followed by Records. Current block headers also contain the VolSessionId and VolSessionTime for the records written to that block.
- Record Header** A record header contains the Volume Session Id, the Volume Session Time, the FileIndex, the Stream, and the size of the data record which follows. The Record Header is always immediately followed by a Data Record if the size given in the Header is greater than zero. Note, for Block headers of level BB02 (version 1.27 and later), the Record header as written to tape does not contain the Volume Session Id and the Volume Session Time as these two fields are stored in the BB02 Block header. The in-memory record header does have those fields for convenience.
- Data Record** A data record consists of a binary stream of bytes and is always preceded by a Record Header. The details of the meaning of the binary stream of bytes are unknown to the Storage daemon, but the Client programs (File daemon) defines and thus knows the details of each record type.
- Volume Label** A label placed by the Storage daemon at the beginning of each storage volume. It contains general information about the volume. It is written in Record format. The Storage daemon manages Volume Labels, and if the client wants, he may also read them.
- Begin Session Label** The Begin Session Label is a special record placed by the Storage daemon on the storage medium as the first record of an append session job with a File daemon. This record is useful for finding the beginning of a particular session (Job), since no records with the same VolSessionId and VolSessionTime will precede this record. This record is not normally visible outside of the Storage daemon. The Begin Session Label is similar to the Volume Label except that it contains additional information pertaining to the Session.
- End Session Label** The End Session Label is a special record placed by the Storage daemon on the storage medium as the last record of an append session job with a File daemon. The End Session Record is distinguished by a FileIndex with a value of minus two (-2). This record is useful for detecting the end of a particular session since no records with the same VolSessionId and VolSessionTime will follow this record. This record is not normally visible outside of the Storage daemon. The End Session Label is similar to the Volume Label except that it contains additional information pertaining to the Session.



10.3 Storage Daemon File Output Format

The file storage and tape storage formats are identical except that tape records are by default blocked into blocks of 64,512 bytes, except for the last block, which is the actual number of bytes written rounded up to a multiple of 1024 whereas the last record of file storage is not rounded up. The default block size of 64,512 bytes may be overridden by the user (some older tape drives only support block sizes of 32K). Each Session written to tape is terminated with an End of File mark (this will be removed later). Sessions written to file are simply appended to the end of the file.

10.4 Overall Format

A Bacula output file consists of Blocks of data. Each block contains a block header followed by records. Each record consists of a record header followed by the record data. The first record on a tape will always be the Volume Label Record.

No Record Header will be split across Bacula blocks. However, Record Data may be split across any number of Bacula blocks. Obviously this will not be the case for the Volume Label which will always be smaller than the Bacula Block size.

To simplify reading tapes, the Start of Session (SOS) and End of Session (EOS) records are never split across blocks. If this is about to happen, Bacula will write a short block before writing the session record (actually, the SOS record should always be the first record in a block, excepting perhaps the Volume label).

Due to hardware limitations, the last block written to the tape may not be fully written. If your drive permits backspace record, Bacula will backup over the last record written on the tape, re-read it and verify that it was correctly written.

When a new tape is mounted Bacula will write the full contents of the partially written block to the new tape ensuring that there is no loss of data. When reading a tape, Bacula will discard any block that is not totally written, thus ensuring that there is no duplication of data. In addition, since Bacula blocks are sequentially numbered within a Job, it is easy to ensure that no block is missing or duplicated.

10.5 Serialization

All Block Headers, Record Headers, and Label Records are written using Bacula's serialization routines. These routines guarantee that the data is written to the output volume in a machine independent format.

10.6 Block Header

The format of the Block Header (version 1.27 and later) is:

```
uint32_t CheckSum;           /* Block check sum */
uint32_t BlockSize;         /* Block byte size including the header */
uint32_t BlockNumber;       /* Block number */
char ID[4] = "BB02";        /* Identification and block level */
uint32_t VolSessionId;      /* Session Id for Job */
uint32_t VolSessionTime;    /* Session Time for Job */
```

The Block header is a fixed length and fixed format and is followed by Record Headers and Record Data. The CheckSum field is a 32 bit checksum of the block data and the block header but not



including the CheckSum field. The Block Header is always immediately followed by a Record Header. If the tape is damaged, a Bacula utility will be able to recover as much information as possible from the tape by recovering blocks which are valid. The Block header is written using the Bacula serialization routines and thus is guaranteed to be in machine independent format. See below for version 2 of the block header.

10.7 Record Header

Each binary data record is preceded by a Record Header. The Record Header is fixed length and fixed format, whereas the binary data record is of variable length. The Record Header is written using the Bacula serialization routines and thus is guaranteed to be in machine independent format.

The format of the Record Header (version 1.27 or later) is:

```
int32_t FileIndex; /* File index supplied by File daemon */
int32_t Stream;    /* Stream number supplied by File daemon */
uint32_t DataSize; /* size of following data record in bytes */
```

This record is followed by the binary Stream data of DataSize bytes, followed by another Record Header record and the binary stream data. For the definitive definition of this record, see record.h in the src/stored directory.

Additional notes on the above:

The **VolSessionId** is a unique sequential number that is assigned by the Storage Daemon to a particular Job. This number is sequential since the start of execution of the daemon.

The **VolSessionTime** is the time/date that the current execution of the Storage Daemon started. It assures that the combination of VolSessionId and VolSessionTime is unique for every jobs written to the tape, even if there was a machine crash between two writes.

The **FileIndex** is a sequential file number within a job. The Storage daemon requires this index to be greater than zero and sequential. Note, however, that the File daemon may send multiple Streams for the same FileIndex. In addition, the Storage daemon uses negative FileIndices to hold the Begin Session Label, the End Session Label, and the End of Volume Label.

The **Stream** is defined by the File daemon and is used to identify separate parts of the data saved for each file (Unix attributes, Win32 attributes, file data, compressed file data, sparse file data, ...). The Storage Daemon has no idea of what a Stream is or what it contains except that the Stream is required to be a positive integer. Negative Stream numbers are used internally by the Storage daemon to indicate that the record is a continuation of the previous record (the previous record would not entirely fit in the block).

For Start Session and End Session Labels (where the FileIndex is negative), the Storage daemon uses the Stream field to contain the JobId. The current stream definitions are:

```
#define STREAM_UNIX_ATTRIBUTES 1 /* Generic Unix attributes */
#define STREAM_FILE_DATA 2 /* Standard uncompressed data */
#define STREAM_MD5_SIGNATURE 3 /* MD5 signature for the file */
#define STREAM_GZIP_DATA 4 /* GZip compressed file data */
/* Extended Unix attributes with Win32 Extended data. Deprecated. */
#define STREAM_UNIX_ATTRIBUTES_EX 5 /* Extended Unix attr for Win32 EX */
#define STREAM_SPARSE_DATA 6 /* Sparse data stream */
#define STREAM_SPARSE_GZIP_DATA 7
#define STREAM_PROGRAM_NAMES 8 /* program names for program data */
#define STREAM_PROGRAM_DATA 9 /* Data needing program */
#define STREAM_SHA1_SIGNATURE 10 /* SHA1 signature for the file */
#define STREAM_WIN32_DATA 11 /* Win32 BackupRead data */
#define STREAM_WIN32_GZIP_DATA 12 /* Gzipped Win32 BackupRead data */
#define STREAM_MACOS_FORK_DATA 13 /* Mac resource fork */
```



```
#define STREAM_HFSPLUS_ATTRIBUTES 14 /* Mac OS extra attributes */
#define STREAM_UNIX_ATTRIBUTES_ACCESS_ACL 15 /* Standard ACL attributes on UNIX */
#define STREAM_UNIX_ATTRIBUTES_DEFAULT_ACL 16 /* Default ACL attributes on UNIX */
```

The **DataSize** is the size in bytes of the binary data record that follows the Session Record header. The Storage Daemon has no idea of the actual contents of the binary data record. For standard Unix files, the data record typically contains the file attributes or the file data. For a sparse file the first 64 bits of the file data contains the storage address for the data block.

The Record Header is never split across two blocks. If there is not enough room in a block for the full Record Header, the block is padded to the end with zeros and the Record Header begins in the next block. The data record, on the other hand, may be split across multiple blocks and even multiple physical volumes. When a data record is split, the second (and possibly subsequent) piece of the data is preceded by a new Record Header. Thus each piece of data is always immediately preceded by a Record Header. When reading a record, if Bacula finds only part of the data in the first record, it will automatically read the next record and concatenate the data record to form a full data record.

10.8 Version BB02 Block Header

Each session or Job has its own private block. As a consequence, the SessionId and SessionTime are written once in each Block Header and not in the Record Header. So, the second and current version of the Block Header BB02 is:

```
uint32_t CheckSum; /* Block check sum */
uint32_t BlockSize; /* Block byte size including the header */
uint32_t BlockNumber; /* Block number */
char ID[4] = "BB02"; /* Identification and block level */
uint32_t VolSessionId; /* Applies to all records */
uint32_t VolSessionTime; /* contained in this block */
```

As with the previous version, the BB02 Block header is a fixed length and fixed format and is followed by Record Headers and Record Data. The CheckSum field is a 32 bit CRC checksum of the block data and the block header but not including the CheckSum field. The Block Header is always immediately followed by a Record Header. If the tape is damaged, a Bacula utility will be able to recover as much information as possible from the tape by recovering blocks which are valid. The Block header is written using the Bacula serialization routines and thus is guaranteed to be in machine independent format.

10.9 Version 2 Record Header

Version 2 Record Header is written to the medium when using Version BB02 Block Headers. The memory representation of the record is identical to the old BB01 Record Header, but on the storage medium, the first two fields, namely VolSessionId and VolSessionTime are not written. The Block Header is filled with these values when the First user record is written (i.e. non label record) so that when the block is written, it will have the current and unique VolSessionId and VolSessionTime. On reading each record from the Block, the VolSessionId and VolSessionTime is filled in the Record Header from the Block Header.

10.10 Volume Label Format

Tape volume labels are created by the Storage daemon in response to a **label** command given to the Console program, or alternatively by the **btape** program. created. Each volume is labeled



with the following information using the Bacula serialization routines, which guarantee machine byte order independence.

For Bacula versions 1.27 and later, the Volume Label Format is:

```
char Id[32];           /* Bacula 1.0 Immortal\n */
uint32_t VerNum;       /* Label version number */
/* VerNum 11 and greater Bacula 1.27 and later */
btime_t label_btime;  /* Time/date tape labeled */
btime_t write_btime;  /* Time/date tape first written */
/* The following are 0 in VerNum 11 and greater */
float64_t write_date;  /* Date this label written */
float64_t write_time;  /* Time this label written */
char VolName[128];     /* Volume name */
char PrevVolName[128]; /* Previous Volume Name */
char PoolName[128];    /* Pool name */
char PoolType[128];    /* Pool type */
char MediaType[128];   /* Type of this media */
char HostName[128];    /* Host name of writing computer */
char LabelProg[32];    /* Label program name */
char ProgVersion[32];  /* Program version */
char ProgDate[32];     /* Program build date/time */
```

Note, the LabelType (Volume Label, Volume PreLabel, Session Start Label, ...) is stored in the record FileIndex field of the Record Header and does not appear in the data part of the record.

10.11 Session Label

The Session Label is written at the beginning and end of each session as well as the last record on the physical medium. It has the following binary format:

```
char Id[32];           /* Bacula Immortal ... */
uint32_t VerNum;       /* Label version number */
uint32_t JobId;        /* Job id */
uint32_t VolumeIndex;  /* sequence no of vol */
/* Prior to VerNum 11 */
float64_t write_date;  /* Date this label written */
/* VerNum 11 and greater */
btime_t write_btime;  /* time/date record written */
/* The following is zero VerNum 11 and greater */
float64_t write_time;  /* Time this label written */
char PoolName[128];    /* Pool name */
char PoolType[128];    /* Pool type */
char JobName[128];     /* base Job name */
char ClientName[128];  /* Added in VerNum 10 */
char Job[128];         /* Unique Job name */
char FileSetName[128]; /* FileSet name */
uint32_t JobType;
uint32_t JobLevel;
```

In addition, the EOS label contains:

```
/* The remainder are part of EOS label only */
uint32_t JobFiles;
uint64_t JobBytes;
uint32_t start_block;
uint32_t end_block;
uint32_t start_file;
uint32_t end_file;
uint32_t JobErrors;
```

In addition, for VerNum greater than 10, the EOS label contains (in addition to the above):

```
uint32_t JobStatus      /* Job termination code */
```



: Note, the LabelType (Volume Label, Volume PreLabel, Session Start Label, ...) is stored in the record FileIndex field and does not appear in the data part of the record. Also, the Stream field of the Record Header contains the JobId. This permits quick filtering without actually reading all the session data in many cases.

10.12 Overall Storage Format

```

Current Bacula Tape Format
    6 June 2001
    Version BB02 added 28 September 2002
    Version BB01 is the old deprecated format.
A Bacula tape is composed of tape Blocks. Each block
has a Block header followed by the block data. Block
Data consists of Records. Records consist of Record
Headers followed by Record Data.
:=====:
|
|      Block Header (24 bytes)
|
|-----|
|
|      Record Header (12 bytes)
|
|-----|
|
|      Record Data
|
|-----|
|
|      Record Header (12 bytes)
|
|-----|
|
|      ...
|
Block Header: the first item in each block. The format is
shown below.
Partial Data block: occurs if the data from a previous
block spills over to this block (the normal case except
for the first block on a tape). However, this partial
data block is always preceded by a record header.
Record Header: identifies the Volume Session, the Stream
and the following Record Data size. See below for format.
Record data: arbitrary binary data.
    Block Header Format BB02
:=====:
|      CheckSum      (uint32_t)
|-----|
|      BlockSize     (uint32_t)
|-----|
|      BlockNumber   (uint32_t)
|-----|
|      "BB02"        (char [4])
|-----|
|      VolSessionId  (uint32_t)
|-----|
|      VolSessionTime (uint32_t)
|-----|
:=====:
BB02: Serves to identify the block as a
Bacula block and also serves as a block format identifier
should we ever need to change the format.
BlockSize: is the size in bytes of the block. When reading
back a block, if the BlockSize does not agree with the
actual size read, Bacula discards the block.
Checksum: a checksum for the Block.
BlockNumber: is the sequential block number on the tape.
VolSessionId: a unique sequential number that is assigned
by the Storage Daemon to a particular Job.
This number is sequential since the start
of execution of the daemon.
VolSessionTime: the time/date that the current execution
of the Storage Daemon started. It assures

```



that the combination of VolSessionId and VolSessionTime is unique for all jobs written to the tape, even if there was a machine crash between two writes.

Record Header Format BB02

=====		
	FileIndex	(int32_t)

	Stream	(int32_t)

	DataSize	(uint32_t)
	=====	

FileIndex: a sequential file number within a job. The Storage daemon enforces this index to be greater than zero and sequential. Note, however, that the File daemon may send multiple Streams for the same FileIndex. The Storage Daemon uses negative FileIndices to identify Session Start and End labels as well as the End of Volume labels.

Stream: defined by the File daemon and is intended to be used to identify separate parts of the data saved for each file (attributes, file data, ...). The Storage Daemon has no idea of what a Stream is or what it contains.

DataSize: the size in bytes of the binary data record that follows the Session Record header. The Storage Daemon has no idea of the actual contents of the binary data record. For standard Unix files, the data record typically contains the file attributes or the file data. For a sparse file the first 64 bits of the data contains the storage address for the data block.

Volume Label

=====		
	Id	(32 bytes)

	VerNum	(uint32_t)

	label_date	(float64_t)
	label_btime	(btime_t VerNum 11

	label_time	(float64_t)
	write_btime	(btime_t VerNum 11

	write_date	(float64_t)
	0	(float64_t) VerNum 11

	write_time	(float64_t)
	0	(float64_t) VerNum 11

	VolName	(128 bytes)

	PrevVolName	(128 bytes)

	PoolName	(128 bytes)

	PoolType	(128 bytes)

	MediaType	(128 bytes)

	HostName	(128 bytes)

	LabelProg	(32 bytes)

	ProgVersion	(32 bytes)

	ProgDate	(32 bytes)

	=====	

Id: 32 byte Bacula identifier "Bacula 1.0 immortal\n"
(old version also recognized:)

Id: 32 byte Bacula identifier "Bacula 0.9 mortal\n"



```

LabelType (Saved in the FileIndex of the Header record).
    PRE_LABEL -1    Volume label on unwritten tape
    VOL_LABEL -2    Volume label after tape written
    EOM_LABEL -3    Label at EOM (not currently implemented)
    SOS_LABEL -4    Start of Session label (format given below)
    EOS_LABEL -5    End of Session label (format given below)
VerNum: 11
label_date: Julian day tape labeled
label_time: Julian time tape labeled
write_date: Julian date tape first used (data written)
write_time: Julian time tape first used (data written)
VolName: "Physical" Volume name
PrevVolName: The VolName of the previous tape (if this tape is
              a continuation of the previous one).
PoolName: Pool Name
PoolType: Pool Type
MediaType: Media Type
HostName: Name of host that is first writing the tape
LabelProg: Name of the program that labeled the tape
ProgVersion: Version of the label program
ProgDate: Date Label program built
          Session Label
:=====:
|          Id          (32 bytes)      |
|-----|
|          VerNum      (uint32_t)      |
|-----|
|          JobId        (uint32_t)      |
|-----|
|          write_btime  (btime_t)  VerNum 11 |
|-----|
|          0            (float64_t) VerNum 11 |
|-----|
|          PoolName     (128 bytes)     |
|-----|
|          PoolType     (128 bytes)     |
|-----|
|          JobName      (128 bytes)     |
|-----|
|          ClientName   (128 bytes)     |
|-----|
|          Job          (128 bytes)     |
|-----|
|          FileSetName   (128 bytes)     |
|-----|
|          JobType       (uint32_t)     |
|-----|
|          JobLevel     (uint32_t)     |
|-----|
|          FileSetMD5    (50 bytes)  VerNum 11 |
|-----|
|          Additional fields in End Of Session Label
|-----|
|          JobFiles      (uint32_t)     |
|-----|
|          JobBytes      (uint32_t)     |
|-----|
|          start_block   (uint32_t)     |
|-----|
|          end_block     (uint32_t)     |
|-----|
|          start_file    (uint32_t)     |
|-----|
|          end_file      (uint32_t)     |
|-----|
|          JobErrors     (uint32_t)     |
|-----|
|          JobStatus     (uint32_t) VerNum 11 |
|-----|
:=====:
* => fields deprecated
Id: 32 byte Bacula Identifier "Bacula 1.0 immortal\n"
LabelType (in FileIndex field of Header):
    EOM_LABEL -3    Label at EOM
    SOS_LABEL -4    Start of Session label
    EOS_LABEL -5    End of Session label

```



```

VerNum: 11
JobId: JobId
write_btime: Bacula time/date this tape record written
write_date: Julian date tape this record written - deprecated
write_time: Julian time tape this record written - deprecated.
PoolName: Pool Name
PoolType: Pool Type
MediaType: Media Type
ClientName: Name of File daemon or Client writing this session
           Not used for EOM_LABEL.

```

10.13 Unix File Attributes

The Unix File Attributes packet consists of the following:

<File-Index> <Type> <Filename>@<File-Attributes>@<Link> @<Extended-Attributes@> where

@ represents a byte containing a binary zero.

FileIndex is the sequential file index starting from one assigned by the File daemon.

Type is one of the following:

```

#define FT_LNKSAVED  1  /* hard link to file already saved */
#define FT_REGE      2  /* Regular file but empty */
#define FT_REG       3  /* Regular file */
#define FT_LNK       4  /* Soft Link */
#define FT_DIR       5  /* Directory */
#define FT_SPEC      6  /* Special file -- chr, blk, fifo, sock */
#define FT_NOACCESS  7  /* Not able to access */
#define FT_NOFOLLOW  8  /* Could not follow link */
#define FT_NOSTAT    9  /* Could not stat file */
#define FT_NOCHG     10 /* Incremental option, file not changed */
#define FT_DIRNOCHG  11 /* Incremental option, directory not changed */
#define FT_ISARCH    12 /* Trying to save archive file */
#define FT_NORECURSE 13 /* No recursion into directory */
#define FT_NOFSCHG   14 /* Different file system, prohibited */
#define FT_NOOPEN    15 /* Could not open directory */
#define FT_RAW       16 /* Raw block device */
#define FT_FIFO      17 /* Raw fifo device */

```

Filename is the fully qualified filename.

File-Attributes consists of the 13 fields of the stat() buffer in ASCII base64 format separated by spaces. These fields and their meanings are shown below. This stat() packet is in Unix format, and MUST be provided (constructed) for ALL systems.

Link when the FT code is FT_LNK or FT_LNKSAVED, the item in question is a Unix link, and this field contains the fully qualified link name. When the FT code is not FT_LNK or FT_LNKSAVED, this field is null.

Extended-Attributes The exact format of this field is operating system dependent. It contains additional or extended attributes of a system dependent nature. Currently, this field is used only on WIN32 systems where it contains a ASCII base64 representation of the WIN32_FILE_ATTRIBUTE_DATA structure as defined by Windows. The fields in the base64 representation of this structure are like the File-Attributes separated by spaces.

The File-attributes consist of the following:



Table 10.1: File Attributes

Field No.	Stat Name	Unix	Win98 / NT	MacOS
1	st_dev	Device number of filesystem	Drive number	vRefNum
2	st_ino	Inode number	Always 0	fileID / dirID
3	st_mode	File mode	File mode	777 dirs / apps; 666 docs; 444 locked docs
4	st_nlink	Number of links to the file	Number of link (only on NTFS)	Always 1
5	st_uid	Owner ID	Always 0	Always 0
6	st_gid	Group ID	Always 0	Always 0
7	st_rdev	Device ID for special files	Drive No.	Always 0
8	st_size	File size in bytes	File size in bytes	Data fork file size in bytes
9	st_blksize	Preferred block size	Always 0	Preferred block size
10	st_blocks	Number of blocks allocated	Always 0	Number of blocks allocated
11	st_atime	Last access time since epoch	Last access time since epoch	Last access time -66 years
12	st_mtime	Last modify time since epoch	Last modify time since epoch	Last access time -66 years
13	st_ctime	Inode change time since epoch	File create time since epoch	File create time -66 years

10.14 Old Deprecated Tape Format

The format of the Block Header (version 1.26 and earlier) is:

```
uint32_t CheckSum;      /* Block check sum */
uint32_t BlockSize;     /* Block byte size including the header */
uint32_t BlockNumber;   /* Block number */
char ID[4] = "BB01";    /* Identification and block level */
```

The format of the Record Header (version 1.26 or earlier) is:

```
uint32_t VolSessionId;  /* Unique ID for this session */
uint32_t VolSessionTime; /* Start time/date of session */
int32_t FileIndex;      /* File index supplied by File daemon */
```




```
int32_t Stream;          /* Stream number supplied by File daemon */
uint32_t DataSize;       /* size of following data record in bytes */
```

Current Bacula Tape Format 6 June 2001

Version BB01 is the old deprecated format.

A Bacula tape is composed of tape Blocks. Each block has a Block header followed by the block data. Block Data consists of Records. Records consist of Record Headers followed by Record Data.




Block Header: the first item in each block. The format is shown below.

Partial Data block: occurs if the data from a previous block spills over to this block (the normal case except for the first block on a tape). However, this partial data block is always preceded by a record header.

Record Header: identifies the Volume Session, the Stream and the following Record Data size. See below for format.

Record data: arbitrary binary data.

Block Header Format BB01 (deprecated)



BB01: Serves to identify the block as a Bacula block and also serves as a block format identifier should we ever need to change the format.

BlockSize: is the size in bytes of the block. When reading back a block, if the BlockSize does not agree with the actual size read, Bacula discards the block.


Checksum: a checksum for the Block.

BlockNumber: is the sequential block number on the tape.

VolSessionId: a unique sequential number that is assigned by the Storage Daemon to a particular Job. This number is sequential since the start of execution of the daemon.

VolSessionTime: the time/date that the current execution of the Storage Daemon started. It assures that the combination of VolSessionId and VolSessionTime is unique for all jobs written to the tape, even if there was a machine crash between two writes.

Record Header Format BB01 (deprecated)





```

|          Stream          (int32_t)          |
|-----|
|          DataSize        (uint32_t)         |
|=====|
VolSessionId: a unique sequential number that is assigned
               by the Storage Daemon to a particular Job.
               This number is sequential since the start
               of execution of the daemon.
VolSessionTime: the time/date that the current execution
                 of the Storage Daemon started. It assures
                 that the combination of VolSessionId and
                 VolSessionTime is unique for all jobs
                 written to the tape, even if there was a
                 machine crash between two writes.
FileIndex: a sequential file number within a job. The
            Storage daemon enforces this index to be
            greater than zero and sequential. Note,
            however, that the File daemon may send
            multiple Streams for the same FileIndex.
            The Storage Daemon uses negative FileIndices
            to identify Session Start and End labels
            as well as the End of Volume labels.
Stream: defined by the File daemon and is intended to be
         used to identify separate parts of the data
         saved for each file (attributes, file data,
         ...). The Storage Daemon has no idea of
         what a Stream is or what it contains.
DataSize: the size in bytes of the binary data record
           that follows the Session Record header.
           The Storage Daemon has no idea of the
           actual contents of the binary data record.
           For standard Unix files, the data record
           typically contains the file attributes or
           the file data. For a sparse file
           the first 64 bits of the data contains
           the storage address for the data block.
           Volume Label
|=====|
|          Id              (32 bytes)          |
|-----|
|          VerNum          (uint32_t)          |
|-----|
|          label_date      (float64_t)         |
|-----|
|          label_time      (float64_t)         |
|-----|
|          write_date      (float64_t)         |
|-----|
|          write_time      (float64_t)         |
|-----|
|          VolName         (128 bytes)         |
|-----|
|          PrevVolName     (128 bytes)         |
|-----|
|          PoolName        (128 bytes)         |
|-----|
|          PoolType        (128 bytes)         |
|-----|
|          MediaType       (128 bytes)         |
|-----|
|          HostName        (128 bytes)         |
|-----|
|          LabelProg       (32 bytes)          |
|-----|
|          ProgVersion     (32 bytes)          |
|-----|
|          ProgDate        (32 bytes)          |
|-----|
|=====|

Id: 32 byte Bacula identifier "Bacula 1.0 immortal\n"
(old version also recognized:)
Id: 32 byte Bacula identifier "Bacula 0.9 mortal\n"
LabelType (Saved in the FileIndex of the Header record).
PRE_LABEL -1 Volume label on unwritten tape

```



VOL_LABEL -2 Volume label after tape written
EOM_LABEL -3 Label at EOM (not currently implemented)
SOS_LABEL -4 Start of Session label (format given below)
EOS_LABEL -5 End of Session label (format given below)

label_date: Julian day tape labeled

label_time: Julian time tape labeled

write_date: Julian date tape first used (data written)

write_time: Julian time tape first used (data written)

VolName: "Physical" Volume name

PrevVolName: The VolName of the previous tape (if this tape is
a continuation of the previous one).

PoolName: Pool Name

PoolType: Pool Type

MediaType: Media Type

HostName: Name of host that is first writing the tape

LabelProg: Name of the program that labeled the tape

ProgVersion: Version of the label program

ProgDate: Date Label program built

Session Label

```
=====
|          Id          (32 bytes)          |
|-----|
|         VerNum       (uint32_t)         |
|-----|
|         JobId        (uint32_t)         |
|-----|
| *write_date         (float64_t) VerNum 10 |
|-----|
| *write_time         (float64_t) VerNum 10 |
|-----|
|        PoolName      (128 bytes)        |
|-----|
|        PoolType      (128 bytes)        |
|-----|
|        JobName       (128 bytes)        |
|-----|
|       ClientName     (128 bytes)        |
|-----|
|         Job          (128 bytes)        |
|-----|
|      FileSetName     (128 bytes)        |
|-----|
|        JobType       (uint32_t)         |
|-----|
|       JobLevel       (uint32_t)         |
|-----|
|      FileSetMD5      (50 bytes)  VerNum 11 |
|-----|
| Additional fields in End Of Session Label |
|-----|
|       JobFiles       (uint32_t)         |
|-----|
|      JobBytes        (uint32_t)         |
|-----|
|     start_block      (uint32_t)         |
|-----|
|     end_block        (uint32_t)         |
|-----|
|     start_file       (uint32_t)         |
|-----|
|     end_file         (uint32_t)         |
|-----|
|      JobErrors       (uint32_t)         |
|-----|
|     JobStatus        (uint32_t) VerNum 11 |
|-----|
=====
```

* => fields deprecated

Id: 32 byte Bacula Identifier "Bacula 1.0 immortal\n"

LabelType (in FileIndex field of Header):

EOM_LABEL -3 Label at EOM

SOS_LABEL -4 Start of Session label

EOS_LABEL -5 End of Session label

VerNum: 11

JobId: JobId

write_btime: Bacula time/date this tape record written



```
write_date: Julian date tape this record written - deprecated
write_time: Julian time tape this record written - deprecated.
PoolName: Pool Name
PoolType: Pool Type
MediaType: Media Type
ClientName: Name of File daemon or Client writing this session
            Not used for EOM_LABEL.
```





Chapter 11

Bacula Porting Notes

This document is intended mostly for developers who wish to port Bacula to a system that is not **officially** supported.

It is hoped that Bacula clients will eventually run on every imaginable system that needs backing up (perhaps even a Palm). It is also hoped that the Bacula Directory and Storage daemons will run on every system capable of supporting them.

11.1 Porting Requirements

In General, the following holds true:

- **Bacula** has been compiled and run on Linux RedHat, FreeBSD, and Solaris systems.
- In addition, clients exist on Win32, and Irix
- It requires GNU C++ to compile. You can try with other compilers, but you are on your own. The Irix client is built with the Irix compiler, but, in general, you will need GNU.
- Your compiler must provide support for 64 bit signed and unsigned integers.
- You will need a recent copy of the **autoconf** tools loaded on your system (version 2.13 or later). The **autoconf** tools are used to build the configuration program, but are not part of the Bacula source distribution.
- There are certain third party packages that Bacula needs. Except for MySQL, they can all be found in the **depkgs** and **depkgs1** releases.
- To build the Win32 binaries, we use Microsoft VC++ standard 2003. Please see the instructions in bacula-source/src/win32/README.win32 for more details. If you want to use VC++ Express, please see README.vc8. Our build is done under the most recent version of Cygwin, but Cygwin is not used in the Bacula binaries that are produced. Unfortunately, we do not have the resources to help you build your own version of the Win32 FD, so you are pretty much on your own. You can ask the bacula-devel list for help, but please don't expect much.
- **Bacula** requires a good implementation of pthreads to work.
- 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.



11.2 Steps to Take for Porting

- The first step is to ensure that you have version 2.13 or later of the **autoconf** tools loaded. You can skip this step, but making changes to the configuration program will be difficult or impossible.
- The run a **./configure** command in the main source directory and examine the output. It should look something like the following:

```
Configuration on Mon Oct 28 11:42:27 CET 2002:
Host: i686-pc-linux-gnu -- redhat 7.3
Bacula version: 1.27 (26 October 2002)
Source code location: .
Install binaries: /sbin
Install config files: /etc/bacula
C Compiler: gcc
C++ Compiler: c++
Compiler flags: -g -O2
Linker flags:
Libraries: -lpthread
Statically Linked Tools: no
Database found: no
Database type: Internal
Database lib:
Job Output Email: root@localhost
Traceback Email: root@localhost
SMTP Host Address: localhost
Director Port 9101
File daemon Port 9102
Storage daemon Port 9103
Working directory /etc/bacula/working
SQL binaries Directory
Large file support: yes
readline support: yes
cweb support: yes /home/kern/bacula/depkgs/cweb
TCP Wrappers support: no
ZLIB support: yes
enable-smartalloc: yes
enable-gnome: no
gmp support: yes
```

The details depend on your system. The first thing to check is that it properly identified your host on the **Host:** line. The first part (added in version 1.27) is the GNU four part identification of your system. The part after the **--** is your system and the system version. Generally, if your system is not yet supported, you must correct these.

- If the **./configure** does not function properly, you must determine the cause and fix it. Generally, it will be because some required system routine is not available on your machine.
- To correct problems with detection of your system type or with routines and libraries, you must edit the file **<bacula-src>/autoconf/configure.in**. This is the “source” from which **configure** is built. In general, most of the changes for your system will be made in **autoconf/aclocal.m4** in the routine **BA_CHECK_OPSYS** or in the routine **BA_CHECK_OPSYS_DISTNAME**. I have already added the necessary code for most systems, but if yours shows up as **unknown** you will need to make changes. Then as mentioned above, you will need to set a number of system dependent items in **configure.in** in the **case** statement at approximately line 1050 (depending on the Bacula release).
- The items to in the case statement that corresponds to your system are the following:
 - **DISTVER** – set to the version of your operating system. Typically some form of **uname** obtains it.
 - **TAPEDRIVE** – the default tape drive. Not too important as the user can set it as an option.
 - **PSCMD** – set to the **ps** command that will provide the PID in the first field and the program name in the second field. If this is not set properly, the **bacula stop** script will most likely not be able to stop Bacula in all cases.



- `hostname` – command to return the base host name (non-qualified) of your system. This is generally the machine name. Not too important as the user can correct this in his configuration file.
 - `CFLAGS` – set any special compiler flags needed. Many systems need a special flag to make pthreads work. See cygwin for an example.
 - `LDFLAGS` – set any special loader flags. See cygwin for an example.
 - `PTHREAD_LIB` – set for any special pthreads flags needed during linking. See freebsd as an example.
 - `lld` – set so that a “long long int” will be properly edited in a `printf()` call.
 - `llu` – set so that a “long long unsigned” will be properly edited in a `printf()` call.
 - `PFILES` – set to add any files that you may define in your platform subdirectory. These files are used for installation of automatic system startup of Bacula daemons.
- To rebuild a new version of **configure** from a changed **autoconf/configure.in** you enter **make configure** in the top level Bacula source directory. You must have done a `./configure` prior to trying to rebuild the configure script or it will get into an infinite loop.
 - If the **make configure** gets into an infinite loop, `ctl-c` it, then do `./configure` (no options are necessary) and retry the **make configure**, which should now work.
 - To rebuild **configure** you will need to have **autoconf** version 2.57-3 or higher loaded. Older versions of autoconf will complain about unknown or bad options, and won't work.
 - After you have a working **configure** script, you may need to make a few system dependent changes to the way Bacula works. Generally, these are done in **src/baconfig.h**. You can find a few examples of system dependent changes toward the end of this file. For example, on Irix systems, there is no definition for **socklen_t**, so it is made in this file. If your system has structure alignment requirements, check the definition of **BALIGN** in this file. Currently, all Bacula allocated memory is aligned on a **double** boundary.
 - If you are having problems with Bacula's type definitions, you might look at **src/bc_types.h** where all the types such as **uint32_t**, **uint64_t**, etc. that Bacula uses are defined.





Implementing a GUI Interface

11.1 General

This document is intended mostly for developers who wish to develop a new GUI interface to Bacula.

11.1.1 Minimal Code in Console Program

Until now, I have kept all the Catalog code in the Directory (with the exception of [dbcheck](#) and [bscan](#)). This is because at some point I would like to add user level security and access. If we have code spread everywhere such as in a GUI this will be more difficult. The other advantage is that any code you add to the Director is automatically available to both the tty console program and the WX program. The major disadvantage is it increases the size of the code – however, compared to Networker the Bacula Director is really tiny.

11.1.2 GUI Interface is Difficult

Interfacing to an interactive program such as Bacula can be very difficult because the interfacing program must interpret all the prompts that may come. This can be next to impossible. There are a number of ways that Bacula is designed to facilitate this:

- The Bacula network protocol is packet based, and thus pieces of information sent can be ASCII or binary.
- The packet interface permits knowing where the end of a list is.
- The packet interface permits special “signals” to be passed rather than data.
- The Director has a number of commands that are non-interactive. They all begin with a period, and provide things such as the list of all Jobs, list of all Clients, list of all Pools, list of all Storage, ... Thus the GUI interface can get to virtually all information that the Director has in a deterministic way. See [\bbracket{bacula-source}/src/dird/ua/_dotcmds.c](#) for more details on this.
- Most console commands allow all the arguments to be specified on the command line: e.g. `run job=NightlyBackup level=Full` command

One of the first things to overcome is to be able to establish a conversation with the Director. Although you can write all your own code, it is probably easier to use the Bacula subroutines. The following code is used by the Console program to begin a conversation.

```
static BSOCK *UA_sock = NULL;
static JCR *jcr;
...
    read-your-config-getting-address-and-pasword;
```



```
UA_sock = bnet_connect(NULL, 5, 15, "Director daemon", dir->address,
                       NULL, dir->DIRport, 0);

if (UA_sock == NULL) {
    terminate_console(0);
    return 1;
}
jcr.dir_bsock = UA_sock;
if (!authenticate_director(&jcr, dir)) {
    fprintf(stderr, "ERR=%s", UA_sock->msg);
    terminate_console(0);
    return 1;
}
read_and_process_input(stdin, UA_sock);
if (UA_sock) {
    bnet_sig(UA_sock, BNET_TERMINATE); /* send EOF */
    bnet_close(UA_sock);
}
exit 0;
```

Then the `read_and_process_input` routine looks like the following:

```
get-input-to-send-to-the-Director;
bnet_fsend(UA_sock, "%s", input);
stat = bnet_recv(UA_sock);
process-output-from-the-Director;
```

For a GUI program things will be a bit more complicated. Basically in the very inner loop, you will need to check and see if any output is available on the `UA_sock`. For an example, please take a look at the WX GUI interface code in: `<bacula-source/src/wx-console`

11.2 Bvfs API

To help developers of restore GUI interfaces, we have added new *dot commands* that permit browsing the catalog in a very simple way.

Bat has now a bRestore panel that uses Bvfs to display files and directories.

The Bvfs module works correctly with BaseJobs, Copy and Migration jobs.

This project was funded by Bacula Systems.

General notes

- All fields are separated by a tab
- You can specify `limit=` and `offset=` to quickly obtain partial listings big directories.
- All operations (except cache creation) are designed to run very fast.
- The cache creation depends on the number of directories, and since Bvfs shares information across jobs, the first creation can take some time.
- All fields returned are separated by a tab.
- Due to potential filename encoding problems, we recommend to always use pathid in queries.

Get dependent jobs from a given JobId

Bvfs allows you to query the catalog against any combination of jobs. You can combine all Jobs and all FileSet for a Client in a single session.




Figure 11.1: Bat Brestore Panel

To get all JobId needed to restore a particular job, you can use the `.bvfs_get_jobids` command.

```
| .bvfs_get_jobids jobid=num [all]
```

```
| .bvfs_get_jobids jobid=10
| 1,2,5,10
| .bvfs_get_jobids jobid=10 all
| 1,2,3,5,10
```

In this example, a normal restore will need to use JobIds 1,2,5,10 to compute a complete restore of the system.

With the `all` option, the Director will use all defined FileSet for this client.

Delete a Catalog File Record

In some cases, it is possible to delete a specific File record with the `.bvfs_delete_fileid` command.

```
| .bvfs_delete_fileid jobid=num fileid=num
```

Generating Bvfs cache

The `.bvfs_update` command computes the directory cache for jobs specified in argument, or for all jobs if unspecified.

```
| .bvfs_update [jobid=numlist]
```



```
.bvfs_update jobid=1,2,3
```

```
.bvfs_versions client=filedaemon pathid=num filenameid=num jobid=1
PathId FilenameId FileId JobId LStat Md5 VolName Inchanger
PathId FilenameId FileId JobId LStat Md5 VolName Inchanger
...
```

```
.bvfs_versions client=localhost-fd pathid=1 fnid=47 jobid=1
1 47 52 12 gD HRid IGk D Po Po A P BAA I A /uPgWaxMgKZlnMti7LChyA Vol1 1
```

```
.bvfs_lsdirs pathid=num path=/apath jobid=numlist limit=num offset=num
PathId  FilenameId  FileId  JobId  LStat  Path
PathId  FilenameId  FileId  JobId  LStat  Path
PathId  FilenameId  FileId  JobId  LStat  Path
...
```

```
.bvfs_lsdirs pathid=4 jobid=1,11,12
```

4	0	0	0	A A A A A A A A A A A A A A A A	.
5	0	0	0	A A A A A A A A A A A A A A A A	..
3	0	0	0	A A A A A A A A A A A A A A A A	regress/

```
.bvfs_lsdirs pathid=3 jobid=1,11,12
```

3	0	0	0	A A A A A A A A A A A A A A A A	.
4	0	0	0	A A A A A A A A A A A A A A A A	..
2	0	0	0	A A A A A A A A A A A A A A A A	tmp/

All trademarks are the property of their respective owners



```
.bvfs_lsfiles pathid=num path=/apath jobid=numlist limit=num offset=num
PathId FilenameId FileId JobId LStat Path
PathId FilenameId FileId JobId LStat Path
PathId FilenameId FileId JobId LStat Path
...
```

You need to pathid or path. Using path="" will list "/" on Unix and all drives on Windows. If FilenameId is 0, the record listed is a directory.

```
.bvfs_lsfiles pathid=4 jobid=1,11,12
4 0 0 0 A A A A A A A A A A A A .
5 0 0 0 A A A A A A A A A A A A ..
1 0 0 0 A A A A A A A A A A A A regress/
```

In this example, to list files present in regress/, you can use

```
.bvfs_lsfiles pathid=1 jobid=1,11,12
1 47 52 12 gD HRid IGk BAA I BMqcPH BMqcPE BMqe+t A titi
1 49 53 12 gD HRid IGk BAA I BMqe/K BMqcPE BMqe+t B toto
1 48 54 12 gD HRie IGk BAA I BMqcPH BMqcPE BMqe+3 A tutu
1 45 55 12 gD HRid IGk BAA I BMqe/K BMqcPE BMqe+t B fichero1.txt
1 46 56 12 gD HRie IGk BAA I BMqe/K BMqcPE BMqe+3 D fichero2.txt
```

In this example, we list all files present in a given job

```
.bvfs_lsfiles allfiles jobid=1
83 1 1 1 m BDPE IGk B Po Bk A t1 BAA I BiQWDy BiQWDy BiQWDy A A G /tmp/regress/build/cor
83 2 2 1 m BDO4 IHT B Po Bk A BI5j BAA JI BiQWEB BiQWDk BiQWDk A A G /tmp/regress/build/lib
83 3 3 1 m BDMV IGk B Po Bk A CsI BAA Y BiQWEB BiQWDi BiQWDi A A G /tmp/regress/build/Mal
83 4 4 1 m BDMI IHT B Po Bk A aaw BAA DY BiQWDk BiQWDh BiQWDh A A G /tmp/regress/build/cor
83 5 5 1 m BDIX IGk B Po Bk A XZj BAA DA BiQWdf BiQWDy BiQWDy A A G /tmp/regress/build/cor
```

Restore set of files

Bvfs allows you to create a SQL table that contains files (as well as plugin objects) that you want to restore. This table can be provided to a restore command with the file option.

```
.bvfs_restore fileid=numlist dirid=numlist path=b2num objectid=numlist
OK
restore file=?b2num ...
```

To include a directory (with dirid), Bvfs needs to run a query to select all files. This query could be time consuming.

The path argument represents the name of the table that Bvfs will store results. The format of this table is b2[0-9]+. (Should start by b2 and followed by digits).

Example:

```
.bvfs_restore fileid=1,2,3,4 jobid=10 path=b20001 objectid=1,2,3
OK
```

Cleanup after Restore

To drop the table used by the restore command, you can use the `.bvfs_cleanup` command.

```
.bvfs_cleanup path=b20001
```



Clearing the BVFS Cache

To clear the BVFS cache, you can use the `.bvfs_clear_cache` command command.

```
.bvfs_clear_cache yes
OK
```

Decode the LStat Attribute

```
*.bvfs_decode_lstat lstat="A A IHA A A A A BAA B BXTDon BXTDon BXTDon A A C"
st_nlink=0
st_mode=33216
st_uid=0
st_gid=0
st_size=0
st_blocks=1
st_ino=0
st_ctime=1464613415
st_mtime=1464613415
st_atime=1464613415
st_dev=0
LinkFI=0
```




Chapter 12

TLS

Written by Landon Fuller

12.1 Introduction to TLS

This patch includes all the back-end code necessary to add complete TLS data encryption support to Bacula. In addition, support for TLS in Console/Director communications has been added as a proof of concept. Adding support for the remaining daemons will be straight-forward. Supported features of this patchset include:

- Client/Server TLS Requirement Negotiation
- TLSv1 Connections with Server and Client Certificate Validation
- Forward Secrecy Support via Diffie-Hellman Ephemeral Keying

This document will refer to both “server” and “client” contexts. These terms refer to the accepting and initiating peer, respectively.

Diffie-Hellman anonymous ciphers are not supported by this patchset. The use of DH anonymous ciphers increases the code complexity and places explicit trust upon the two-way Cram-MD5 implementation. Cram-MD5 is subject to known plaintext attacks, and is should be considered considerably less secure than PKI certificate-based authentication.

Appropriate autoconf macros have been added to detect and use OpenSSL. Two additional preprocessor defines have been added: `HAVE_TLS` and `HAVE_OPENSSL`. All changes not specific to OpenSSL rely on `HAVE_TLS`. OpenSSL-specific code is constrained to `src/lib/tls.c` to facilitate the support of alternative TLS implementations.

12.2 New Configuration Directives

Additional configuration directives have been added to both the Console and Director resources. These new directives are defined as follows:

- `TLS Enable` (*yes/no*) Enable TLS support.
- `TLS Require` (*yes/no*) Require TLS connections.
- `TLS Certificate` (*path*) Path to PEM encoded TLS certificate. Used as either a client or server certificate.



- **TLS Key** (*path*) Path to PEM encoded TLS private key. Must correspond with the TLS certificate.
- **TLS Verify Peer** (*yes/no*) Verify peer certificate. Instructs server to request and verify the client's x509 certificate. Any client certificate signed by a known-CA will be accepted unless the TLS Allowed CN configuration directive is used. Not valid in a client context.
- **TLS Allowed CN** (*string list*) Common name attribute of allowed peer certificates. If directive is specified, all client certificates will be verified against this list. This directive may be specified more than once. Not valid in a client context.
- **TLS CA Certificate File** (*path*) Path to PEM encoded TLS CA certificate(s). Multiple certificates are permitted in the file. One of *TLS CA Certificate File* or *TLS CA Certificate Dir* are required in a server context if *TLS Verify Peer* is also specified, and are always required in a client context.
- **TLS CA Certificate Dir** (*path*) Path to TLS CA certificate directory. In the current implementation, certificates must be stored PEM encoded with OpenSSL-compatible hashes. One of *TLS CA Certificate File* or *TLS CA Certificate Dir* are required in a server context if *TLS Verify Peer* is also specified, and are always required in a client context.
- **TLS DH File** (*path*) Path to PEM encoded Diffie-Hellman parameter file. If this directive is specified, DH ephemeral keying will be enabled, allowing for forward secrecy of communications. This directive is only valid within a server context. To generate the parameter file, you may use openssl:

```
| openssl dhparam -out dh1024.pem -5 1024
```

12.3 TLS API Implementation

To facilitate the use of additional TLS libraries, all OpenSSL-specific code has been implemented within *src/lib/tls.c*. In turn, a generic TLS API is exported.

12.3.1 Library Initialization and Cleanup

```
| int init_tls (void);
```

Performs TLS library initialization, including seeding of the PRNG. PRNG seeding has not yet been implemented for win32.

```
| int cleanup_tls (void);
```

Performs TLS library cleanup.

12.3.2 Manipulating TLS Contexts

```
| TLS_CONTEXT *new_tls_context (const char *ca_certfile,  
|                               const char *ca_cermdir, const char *certfile,  
|                               const char *keyfile, const char *dhfile, bool verify_peer);
```

Allocates and initializes a new opaque *TLS_CONTEXT* structure. The *TLS_CONTEXT* structure maintains default TLS settings from which *TLS_CONNECTION* structures are instantiated. In the future the *TLS_CONTEXT* structure may be used to maintain the TLS session cache. *ca_certfile* and *ca_cermdir* arguments are used to initialize the CA verification stores. The *certfile* and *keyfile* arguments are used to initialize the local certificate and private key. If *dhfile* is non-NULL, it is used to initialize Diffie-Hellman ephemeral keying. If *verify_peer* is *true*, client certificate validation is enabled.



```
| void free_tls_context (TLS_CONTEXT *ctx);
```

Deallocates a previously allocated `TLS_CONTEXT` structure.

12.3.3 Performing Post-Connection Verification

```
| bool tls_postconnect_verify_host (TLS_CONNECTION *tls, const char *host);
```

Performs post-connection verification of the peer-supplied x509 certificate. Checks whether the `subjectAltName` and `commonName` attributes match the supplied `host` string. Returns `true` if there is a match, `false` otherwise.

```
| bool tls_postconnect_verify_cn (TLS_CONNECTION *tls, alist *verify_list);
```

Performs post-connection verification of the peer-supplied x509 certificate. Checks whether the `commonName` attribute matches any strings supplied via the `verify_list` parameter. Returns `true` if there is a match, `false` otherwise.

12.3.4 Manipulating TLS Connections

```
| TLS_CONNECTION *new_tls_connection (TLS_CONTEXT *ctx, int fd);
```

Allocates and initializes a new `TLS_CONNECTION` structure with context `ctx` and file descriptor `fd`.

```
| void free_tls_connection (TLS_CONNECTION *tls);
```

Deallocates memory associated with the `tls` structure.

```
| bool tls_bsock_connect (BSOCK *bsock);
```

Negotiates a TLS client connection via `bsock`. Returns `true` if successful, `false` otherwise. Will fail if there is a TLS protocol error or an invalid certificate is presented

```
| bool tls_bsock_accept (BSOCK *bsock);
```

Accepts a TLS client connection via `bsock`. Returns `true` if successful, `false` otherwise. Will fail if there is a TLS protocol error or an invalid certificate is presented.

```
| bool tls_bsock_shutdown (BSOCK *bsock);
```

Issues a blocking TLS shutdown request to the peer via `bsock`. This function may not wait for the peer's reply.

```
| int tls_bsock_writen (BSOCK *bsock, char *ptr, int32_t nbytes);
```

Writes `nbytes` from `ptr` via the `TLS_CONNECTION` associated with `bsock`. Due to OpenSSL's handling of `EINTR`, `bsock` is set non-blocking at the start of the function, and restored to its original blocking state before the function returns. Less than `nbytes` may be written if an error occurs. The actual number of bytes written will be returned.

```
| int tls_bsock_readn (BSOCK *bsock, char *ptr, int32_t nbytes);
```

Reads `nbytes` from the `TLS_CONNECTION` associated with `bsock` and stores the result in `ptr`. Due to OpenSSL's handling of `EINTR`, `bsock` is set non-blocking at the start of the function, and restored to its original blocking state before the function returns. Less than `nbytes` may be read if an error occurs. The actual number of bytes read will be returned.



12.4 Bnet API Changes

A minimal number of changes were required in the Bnet socket API. The BSOCK structure was expanded to include an associated TLS_CONNECTION structure, as well as a flag to designate the current blocking state of the socket. The blocking state flag is required for win32, where it does not appear possible to discern the current blocking state of a socket.

12.4.1 Negotiating a TLS Connection

bnet_tls_server() and *bnet_tls_client()* were both implemented using the new TLS API as follows:

```
| int bnet_tls_client(TLS_CONTEXT *ctx, BSOCK * bsock);
```

Negotiates a TLS session via *bsock* using the settings from *ctx*. Returns 1 if successful, 0 otherwise.

```
| int bnet_tls_server(TLS_CONTEXT *ctx, BSOCK * bsock, alist *verify_list);
```

Accepts a TLS client session via *bsock* using the settings from *ctx*. If *verify_list* is non-NULL, it is passed to *tls_postconnect_verify_cn()* for client certificate verification.

12.4.2 Manipulating Socket Blocking State

Three functions were added for manipulating the blocking state of a socket on both Win32 and Unix-like systems. The Win32 code was written according to the MSDN documentation, but has not been tested.

These functions are prototyped as follows:

```
| int bnet_set_nonblocking (BSOCK *bsock);
```

Enables non-blocking I/O on the socket associated with *bsock*. Returns a copy of the socket flags prior to modification.

```
| int bnet_set_blocking (BSOCK *bsock);
```

Enables blocking I/O on the socket associated with *bsock*. Returns a copy of the socket flags prior to modification.

```
| void bnet_restore_blocking (BSOCK *bsock, int flags);
```

Restores blocking or non-blocking IO setting on the socket associated with *bsock*. The *flags* argument must be the return value of either *bnet_set_blocking()* or *bnet_restore_blocking()*.



12.5 Authentication Negotiation

Backwards compatibility with the existing SSL negotiation hooks implemented in `src/lib/cram-md5.c` have been maintained. The `cram_md5_get_auth()` function has been modified to accept an integer pointer argument, `tls_remote_need`. The TLS requirement advertised by the remote host is returned via this pointer.

After exchanging cram-md5 authentication and TLS requirements, both the client and server independently decide whether to continue:

```
if (!cram_md5_get_auth(dir, password, &tls_remote_need) ||
    !cram_md5_auth(dir, password, tls_local_need)) {
    [snip]
    /* Verify that the remote host is willing to meet our TLS requirements */
    if (tls_remote_need < tls_local_need && tls_local_need != BNET_TLS_OK &&
        tls_remote_need != BNET_TLS_OK) {
        sendit_("Authorization problem:"
            " Remote server did not advertise required TLS support.\n");
        auth_success = false;
        goto auth_done;
    }

    /* Verify that we are willing to meet the remote host's requirements */
    if (tls_remote_need > tls_local_need && tls_local_need != BNET_TLS_OK &&
        tls_remote_need != BNET_TLS_OK) {
        sendit_("Authorization problem:"
            " Remote server requires TLS.\n");
        auth_success = false;
        goto auth_done;
    }
}
```





Chapter 13

Bacula Regression Testing

13.1 Setting up Regression Testing

This document is intended mostly for developers who wish to ensure that their changes to Bacula don't introduce bugs in the base code. However, you don't need to be a developer to run the regression scripts, and we recommend them before putting your system into production, and before each upgrade, especially if you build from source code. They are simply shell scripts that drive Bacula through bconsole and then typically compare the input and output with `diff`.

You can find the existing regression scripts in the Bacula developer's [git](#) repository on SourceForge. We strongly recommend that you **clone** the repository because afterwards, you can easily get pull the updates that have been made.

To get started, we recommend that you create a directory named `bacula`, under which you will put the current source code and the current set of regression scripts. Below, we will describe how to set this up.

The top level directory that we call `bacula` can be named anything you want. Note, all the standard regression scripts run as non-root and can be run on the same machine as a production Bacula system (the developers run it this way).

To create the directory structure for the current trunk and to clone the repository, do the following (note, we assume you are working in your home directory in a non-root account):

```
| git clone http://git.bacula.org/bacula.git bacula
```

This will create the directory `bacula` and populate it with three directories: `bacula`, `gui`, and `regress`. `bacula` contains the Bacula source code; `gui` contains certain GUI programs that you will not need, and `regress` contains all the regression scripts. The above should be needed only once. Thereafter to update to the latest code, you do:

```
| cd bacula
| git pull
```

There are two different aspects of regression testing that this document will discuss: 1. Running the Regression Script, 2. Writing a Regression test.

13.2 Running the Regression Script

There are a number of different tests that may be run, such as: the standard set that uses disk Volumes and runs under any userid; a small set of tests that write to tape; another set of tests where you must be root to run them. Normally, I run all my tests as non-root and very rarely run the root tests. The tests vary in length, and



running the full tests including disk based testing, tape based testing, autochanger based testing, and multiple drive autochanger based testing can take 3 or 4 hours.

13.2.1 Setting the Configuration Parameters

There is nothing you need to change in the source directory.

To begin:

```
| cd bacula/regress
```

The very first time you are going to run the regression scripts, you will need to create a custom config file for your system. We suggest that you start by:

```
| cp prototype.conf config
```

Then you can edit the `config` file directly.

```
# Where to get the source to be tested
BACULA_SOURCE="${HOME}/bacula/bacula"

# Where to send email  !!!!! Change me !!!!!!!
EMAIL=your-name@your-domain.com
SMTP_HOST="localhost"

TAPE_DRIVE="/dev/nst0"
# if you don't have an autochanger set AUTOCHANGER to /dev/null
AUTOCHANGER="/dev/sg0"
# For two drive tests -- set to /dev/null if you do not have it
TAPE_DRIVE1="/dev/null"

# This must be the path to the autochanger including its name
AUTOCHANGER_PATH="/usr/sbin/mtx"

# Set what backend to use "postgresql" "mysql" or "sqlite3"
DBTYPE="postgresql"

# Set your database here
#WHICHDB="--with-${DBTYPE}=${SQLITE3_DIR}"
WHICHDB="--with-${DBTYPE}"

# Set this to "--with-tcp-wrappers" or "--without-tcp-wrappers"
TCPWRAPPERS="--with-tcp-wrappers"

# Set this to "" to disable OpenSSL support, "--with-openssl=yes"
# to enable it, or provide the path to the OpenSSL installation,
# eg "--with-openssl=/usr/local"
OPENSSL="--with-openssl"

# You may put your real host name here, but localhost or 127.0.0.1
# is valid also and it has the advantage that it works on a
# non-networked machine
HOST="localhost"
```

`BACULA_SOURCE` should be the full path to the Bacula source code that you wish to test. It will be loaded, configured, compiled, and installed with the "make setup" command, which needs to be done only once each time you change the source code.

`EMAIL` should be your email address. Please remember to change this or I will get a flood of unwanted messages. You may or may not want to see these emails. In my case, I don't need them so I direct it to the bit bucket.

`SMTP_HOST` defines where your SMTP server is.

`TAPE_DRIVE` is the full path to your tape drive. The base set of regression tests do not use a tape, so this is only important if you want to run the full tests. Set this to `/dev/null` if you do not have a tape drive.

`TAPE_DRIVE1` is the full path to your second tape drive, if have one. The base set of regression tests do not use a tape, so this is only important if you want to run the full two drive tests. Set this to `/dev/null` if you do not have a second tape drive.

`AUTOCHANGER` is the name of your autochanger control device. Set this to `/dev/null` if you do not have an autochanger.



AUTOCHANGER_PATH is the full path including the program name for your autochanger program (normally `mtx`). Leave the default value if you do not have one.

TCPWRAPPERS defines whether or not you want the `./configure` to be performed with `tcpwrappers` enabled.

OPENSSL used to enable/disable SSL support for Bacula communications and data encryption.

HOST is the hostname that it will use when building the scripts. The Bacula daemons will be named `<HOST>-dir`, `<HOST>-fd`, ... It is also the name of the HOST machine that to connect to the daemons by the network. Hence the name should either be your real hostname (with an appropriate DNS or `/etc/hosts` entry) or `localhost` as it is in the default file.

bin is the binary location.

scripts is the bacula scripts location (where we could find database creation script, autochanger handler, etc.)

13.2.2 Building the Test Bacula

Once the above variables are set, you can build the setup by entering:

```
| make setup
```

This will setup the regression testing and you should not need to do this again unless you want to change the database or other regression configuration parameters.

13.2.3 Setting up your SQL engine

If you are using SQLite or SQLite 3, there is nothing more to do; you can simply run the tests as described in the next section.

If you are using MySQL or PostgreSQL, you will need to establish an account with your database engine for the user name `regress` and you will need to manually create a database named `regress` that can be used by user name `regress`, which means you will have to give the user `regress` sufficient permissions to use the database named `regress`. There is no password on the `regress` account.

You have probably already done this procedure for the user name and database named `bacula`. If not, the manual describes roughly how to do it, and the scripts in `bacula/regress/build/src/cats` named `create_mysql_database`, `create_postgresql_database`, `grant_mysql_privileges`, and `grant_postgresql_privileges` may be of a help to you.

Generally, to do the above, you will need to run under `root` to be able to create databases and modify permissions within MySQL and PostgreSQL.

It is possible to configure MySQL access for database accounts that require a password to be supplied. This can be done by creating a `~/my.cnf` file which supplies the credentials by default to the MySQL commandline utilities.

```
| [client]
| host      = localhost
| user      = regress
| password  = asecret
```

A similar technique can be used PostgreSQL regression testing where the database is configured to require a password. The `~/pgpass` file should contain a line with the database connection properties.

```
| hostname:port:database:username:password
```

13.2.4 Running the Disk Only Regression

The simplest way to copy the source code, configure it, compile it, link it, and run the tests is to use a helper script:

```
| ./do_disk
```

This will run the base set of tests using disk Volumes. If you are testing on a non-Linux machine several of the tests may not be run. In any case, as we add new tests, the number will vary. It will take about 1 hour and you don't need to be `root` to run these tests (I run under my regular userid). The result should be something similar to:



```
Test results
===== auto-label-test OK 12:31:33 =====
===== backup-bacula-test OK 12:32:32 =====
===== bextract-test OK 12:33:27 =====
===== bscan-test OK 12:34:47 =====
===== bsr-opt-test OK 12:35:46 =====
===== compressed-test OK 12:36:52 =====
===== compressed-encrypt-test OK 12:38:18 =====
===== concurrent-jobs-test OK 12:39:49 =====
===== data-encrypt-test OK 12:41:11 =====
===== encrypt-bug-test OK 12:42:00 =====
===== fifo-test OK 12:43:46 =====
===== backup-bacula-fifo OK 12:44:54 =====
===== differential-test OK 12:45:36 =====
===== four-concurrent-jobs-test OK 12:47:39 =====
===== four-jobs-test OK 12:49:22 =====
===== incremental-test OK 12:50:38 =====
===== query-test OK 12:51:37 =====
===== recycle-test OK 12:53:52 =====
===== restore2-by-file-test OK 12:54:53 =====
===== restore-by-file-test OK 12:55:40 =====
===== restore-disk-seek-test OK 12:56:29 =====
===== six-vol-test OK 12:57:44 =====
===== span-vol-test OK 12:58:52 =====
===== sparse-compressed-test OK 13:00:00 =====
===== sparse-test OK 13:01:04 =====
===== two-jobs-test OK 13:02:39 =====
===== two-vol-test OK 13:03:49 =====
===== verify-vol-test OK 13:04:56 =====
===== weird-files2-test OK 13:05:47 =====
===== weird-files-test OK 13:06:33 =====
===== migration-job-test OK 13:08:15 =====
===== migration-jobspan-test OK 13:09:33 =====
===== migration-volume-test OK 13:10:48 =====
===== migration-time-test OK 13:12:59 =====
===== hardlink-test OK 13:13:50 =====
===== two-pool-test OK 13:18:17 =====
===== fast-two-pool-test OK 13:24:02 =====
===== two-volume-test OK 13:25:06 =====
===== incremental-2disk OK 13:25:57 =====
===== 2drive-incremental-2disk OK 13:26:53 =====
===== scratch-pool-test OK 13:28:01 =====
Total time = 0:57:55 or 3475 secs
```

and the working tape tests are run with

```
| make full_test
```

```
Test results
===== Bacula tape test OK =====
===== Small File Size test OK =====
===== restore-by-file-tape test OK =====
===== incremental-tape test OK =====
===== four-concurrent-jobs-tape OK =====
===== four-jobs-tape OK =====
```

Each separate test is self contained in that it initializes to run Bacula from scratch (i.e. newly created database). It will also kill any Bacula session that is currently running. In addition, it uses ports 8101, 8102, and 8103 so that it does not interfere with a production system.

Alternatively, you can do the `./do_disk` work by hand with:

```
| make setup
```

The above will then copy the source code within the regression tree (in directory `regress/build`), configure it, and build it. There should be no errors. If there are, please correct them before continuing. From this point on, as long as you don't change the Bacula source code, you should not need to repeat any of the above steps. If you pull down a new version of the source code, simply run `make setup` again.

Once Bacula is built, you can run the basic disk only non-root regression test by entering:



```
| make test
```

13.2.5 Other Tests

There are a number of other tests that can be run as well. All the tests are a simply shell script keep in the `regress` directory. For example the `make test` simply executes `./all-non-root-tests`. The other tests, which are invoked by directly running the script are:

`all_non-root-tests` All non-tape tests not requiring root. This is the standard set of tests, that in general, backup some data, then restore it, and finally compares the restored data with the original data.

`all-root-tests` All non-tape tests requiring root permission. These are a relatively small number of tests that require running as root. The amount of data backed up can be quite large. For example, one test backs up `/usr`, another backs up `/etc`. One or more of these tests reports an error – I'll fix it one day.

`all-non-root-tape-tests` All tape test not requiring root. There are currently three tests, all run without being root, and backup to a tape. The first two tests use one volume, and the third test requires an autochanger, and uses two volumes. If you don't have an autochanger, then this script will probably produce an error.

`all-tape-and-file-tests` All tape and file tests not requiring root. This includes just about everything, and I don't run it very often.

13.2.6 If a Test Fails

If you one or more tests fail, the line output will be similar to:

```
| !!!!! concurrent-jobs-test failed!!! !!!!!
```

If you want to determine why the test failed, you will need to rerun the script with the debug output turned on. You do so by defining the environment variable `REGRESS_DEBUG` with commands such as:

```
| REGRESS_DEBUG=1
| export REGRESS_DEBUG
```

Then from the `regress` directory (all regression scripts assume that you have `regress` as the current directory), enter:

```
| tests/test-name
```

where test-name should be the name of a test script – for example: `tests/backup-bacula-test`.

13.3 Testing a Binary Installation

If you have installed your Bacula from a binary release such as (rpms or debs), you can still run regression tests on it. First, make sure that your regression `config` file uses the same catalog backend as your installed binaries. Then define the variables `bin` and `scripts` variables in your config file.

Example:

```
| bin=/opt/bacula/bin
| scripts=/opt/bacula/scripts
```

The `./scripts/prepare-other-loc` will tweak the regress scripts to use your binary location. You will need to run it manually once before you run any regression tests.

```
| $ ./scripts/prepare-other-loc
| $ ./tests/backup-bacula-test
| ...
```

All regression scripts must be run by hand or by calling the test scripts. These are principally scripts that begin with `all_...` such as `all_disk_tests`, `./all_test`, ... None of the `./do_disk`, `./do_all`, `./nightly...` scripts will work.

If you want to switch back to running the regression scripts from source, first remove the `bin` and `scripts` variables from your `config` file and rerun the `make setup` step.



13.4 Running a Single Test

If you wish to run a single test, you can simply:

```
cd regress
tests/<name-of-test>
```

or, if the source code has been updated, you would do:

```
cd bacula
git pull
cd regress
make setup
tests/backup-to-null
```

13.5 Writing a Regression Test

Any developer, who implements a major new feature, should write a regression test that exercises and validates the new feature. Each regression test is a complete test by itself. It terminates any running Bacula, initializes the database, starts Bacula, then runs the test by using the console program.

13.5.1 Running the Tests by Hand

You can run any individual test by hand by cd'ing to the `regress` directory and entering:

```
tests/<test-name>
```

13.5.2 Directory Structure

The directory structure of the regression tests is:

```
regress          - Makefile, scripts to start tests
|----- scripts - Scripts and conf files
|-----tests    - All test scripts are here
|
|-----         -- All directories below this point are used
|                  for testing, but are created from the
|                  above directories and are removed with
|                  "make distclean"
|
|----- bin      - This is the install directory for
|                  Bacula to be used testing
|----- build    - Where the Bacula source build tree is
|----- tmp      - Most temp files go here
|----- working  - Bacula working directory
|----- weird-files - Weird files used in two of the tests.
```

13.5.3 Adding a New Test

If you want to write a new regression test, it is best to start with one of the existing test scripts, and modify it to do the new test.

When adding a new test, be extremely careful about adding anything to any of the daemons' configuration files. The reason is that it may change the prompts that are sent to the console. For example, adding a Pool means that the current scripts, which assume that Bacula automatically selects a Pool, will now be presented with a new prompt, so the test will fail. If you need to enhance the configuration files, consider making your own versions.

Once written, if the test can be executed automatically, you can do a separated commit to schedule the new test via CDASH by editing the `regress/DartTestFile.txt.in`



13.5.4 Adding a Unittest

For low level operations, it is important to write unittests. They are usually directly included at the end of the file that has the code to test. It is possible to create a new file in `regress/src` or `bacula/src/tools`.

The C code should be compiled with the appropriate Makefile rule (see `alist_test`) for example.

```
mytest_test: Makefile libbac.la mytest.c unittests.o
    $(CXX) -DTEST_PROGRAM $(DEFS) $(DEBUG) -c $(CPPFLAGS) -I$(srcdir) -I$(basedir) $(DINCLUDE) $(CFLAGS) m
    $(LIBTOOL_LINK) $(CXX) $(LDLFLAGS) -L. -o $@ mytest_test.o unittests.o $(DLIB) -lbac -lm $(LIBS) $(OPEN
    $(LIBTOOL_INSTALL) $(INSTALL_PROGRAM) $@ $(DESTDIR)$$(sbindir)/
```

The Bacula Unittest framework is a simple set of functions that helps to write a unittest. It has functions to

- Test variables
- Log messages
- Configure/Unconfigure Bacula tools such as mempool, lock manager
- Do some high level filesystem functions (`mkdir`, `rmdir`, `stat`)
- Ease some string operations

The following code explains how to use the Bacula Unittest framework:

```
int function_to_test(int parameter1, int parameter2)
{
    /* some interesting code */
    return 1;
}

/* In this case, the test procedure is at the end of the file,
 * it can be also in a separated file
 */
#ifdef TEST_PROGRAM
#include "lib/unittests.h"

int main(int argc, char **argv)
{
    Unittests tests("app_test");
    tests.set_nb_tests(4); /* Plan the number of expected tests */
    tests.configure(TEST_QUIET|TEST_PRINT_LOCAL); /* Configure */

    log("Test a the function_to_test()");
    is(function_to_test(1,2), 1, "Test function_to_test(1,2)");
    ok(function_to_test(1,2) == 1, "Test function_to_test(1,2)");
    isnt(function_to_test(3,4), 2, "Test function_to_test(3,4)");
    nok(function_to_test(3,4) == 2, "Test function_to_test(3,4)");

    return report();
}
#endif /* TEST_PROGRAM */
```

The current options in the unittest library are:

`TEST_QUIET` Display only test in error

`TEST_PRINT_LOCAL` Display the local variables after an error

These options can be set via an environment variable (`UNITTEST_TEST_QUIET`, `UNITTEST_TEST_PRINT_LOCAL`).

13.5.5 Running a Test Under The Debugger

You can run a test under the debugger (actually run a Bacula daemon under the debugger) by first setting the environment variable `REGRESS_WAIT` with commands such as:

```
REGRESS_WAIT=1
export REGRESS_WAIT
```

Then executing the script. When the script prints the following line:



| Start Bacula under debugger and enter anything when ready ...

You start the Bacula component you want to run under the debugger in a different shell window. For example:

```
| cd .../regress/bin  
| gdb bacula-sd  
| (possibly set breakpoints, ...)  
| run -s -f
```

Then enter any character in the window with the above message. An error message will appear saying that the daemon you are debugging is already running, which is the case. You can simply ignore the error message.



Chapter 14

Bacula MD5 Algorithm

14.1 Command Line Message Digest Utility

This page describes **md5**, a command line utility usable on either Unix or MS-DOS/Windows, which generates and verifies message digests (digital signatures) using the MD5 algorithm. This program can be useful when developing shell scripts or Perl programs for software installation, file comparison, and detection of file corruption and tampering.

14.1.1 Name

md5 - generate / check MD5 message digest

14.1.2 Synopsis

```
md5 [ -csignature ] [ -u ] [ -dinput_text | infile ] [ outfile ]
```

14.1.3 Description

A *message digest* is a compact digital signature for an arbitrarily long stream of binary data. An ideal message digest algorithm would never generate the same signature for two different sets of input, but achieving such theoretical perfection would require a message digest as long as the input file. Practical message digest algorithms compromise in favour of a digital signature of modest size created with an algorithm designed to make preparation of input text with a given signature computationally infeasible. Message digest algorithms have much in common with techniques used in encryption, but to a different end; verification that data have not been altered since the signature was published.

Many older programs requiring digital signatures employed 16 or 32 bit *cyclical redundancy codes* (CRC) originally developed to verify correct transmission in data communication protocols, but these short codes, while adequate to detect the kind of transmission errors for which they were intended, are insufficiently secure for applications such as electronic commerce and verification of security related software distributions.

The most commonly used present-day message digest algorithm is the 128 bit MD5 algorithm, developed by Ron Rivest of the [MIT Laboratory for Computer Science](#) and [RSA Data Security, Inc.](#) The algorithm, with a reference implementation, was published as Internet [RFC 1321](#) in April 1992, and was placed into the public domain at that time. Message digest algorithms such as MD5 are not deemed "encryption technology" and are not subject to the export controls some governments impose on other data security products. (Obviously, the responsibility for obeying the laws in the jurisdiction in which you reside is entirely your own, but many common Web and Mail utilities use MD5, and I am unaware of any restrictions on their distribution and use.)

The MD5 algorithm has been implemented in numerous computer languages including C, [Perl](#), and [Java](#); if you're writing a program in such a language, track down a suitable subroutine and incorporate it into your program. The program described on this page is a *command line* implementation of MD5, intended for use in shell scripts and Perl programs (it is much faster than computing an MD5 signature directly in Perl). This **md5** program was originally developed as part of a suite of tools intended to monitor large collections of files (for example, the contents of a Web site) to detect corruption of files and inadvertent (or perhaps malicious) changes. That



task is now best accomplished with more comprehensive packages such as [Tripwire](#), but the command line **md5** component continues to prove useful for verifying correct delivery and installation of software packages, comparing the contents of two different systems, and checking for changes in specific files.

14.1.4 Options

- csignature** Computes the signature of the specified *infile* or the string supplied by the **-d** option and compares it against the specified *signature*. If the two signatures match, the exit status will be zero, otherwise the exit status will be 1. No signature is written to *outfile* or standard output; only the exit status is set. The signature to be checked must be specified as 32 hexadecimal digits.
- dinput_text** A signature is computed for the given *input_text* (which must be quoted if it contains white space characters) instead of input from *infile* or standard input. If input is specified with the **-d** option, no *infile* should be specified.
- u** Print how-to-call information.

14.1.5 Files

If no *infile* or **-d** option is specified or *infile* is a single "-", **md5** reads from standard input; if no *outfile* is given, or *outfile* is a single "-", output is sent to standard output. Input and output are processed strictly serially; consequently **md5** may be used in pipelines.

14.1.6 Bugs

The mechanism used to set standard input to binary mode may be specific to Microsoft C; if you rebuild the DOS/Windows version of the program from source using another compiler, be sure to verify binary files work properly when read via redirection or a pipe.

This program has not been tested on a machine on which `int` and/or `long` are longer than 32 bits.

14.2 Download md5.zip (Zipped archive)

The program is provided as [md5.zip](#), a [Zipped](#) archive containing an ready-to-run Win32 command-line executable program, `md5.exe` (compiled using Microsoft Visual C++ 5.0), and in source code form along with a `Makefile` to build the program under Unix.

14.2.1 See Also

`sum(1)`

14.2.2 Exit Status

md5 returns status 0 if processing was completed without errors, 1 if the **-c** option was specified and the given signature does not match that of the input, and 2 if processing could not be performed at all due, for example, to a nonexistent input file.

14.2.3 Copying

This software is in the public domain. Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, without any conditions or restrictions. This software is provided "as is" without express or implied warranty.

14.2.4 Acknowledgements

The MD5 algorithm was developed by Ron Rivest. The public domain C language implementation used in this program was written by Colin Plumb in 1993. *by John Walker January 6th, MIM*



Chapter 15

Bacula Memory Management

15.1 General

This document describes the memory management routines that are used in Bacula and is meant to be a technical discussion for developers rather than part of the user manual.

Since Bacula may be called upon to handle filenames of varying and more or less arbitrary length, special attention needs to be used in the code to ensure that memory buffers are sufficiently large. There are four possibilities for memory usage within **Bacula**. Each will be described in turn. They are:

- Statically allocated memory.
- Dynamically allocated memory using `malloc()` and `free()`.
- Non-pooled memory.
- Pooled memory.

15.1.1 Statically Allocated Memory

Statically allocated memory is of the form:

```
| char buffer[MAXSTRING];
```

The use of this kind of memory is discouraged except when you are 100% sure that the strings to be used will be of a fixed length. One example of where this is appropriate is for **Bacula** resource names, which are currently limited to 127 characters (`MAX_NAME_LENGTH`). Although this maximum size may change, particularly to accommodate Unicode, it will remain a relatively small value.

15.1.2 Dynamically Allocated Memory

Dynamically allocated memory is obtained using the standard `malloc()` routines. As in:

```
| char *buf;  
| buf = malloc(256);
```

This kind of memory can be released with:

```
| free(buf);
```

It is recommended to use this kind of memory only when you are sure that you know the memory size needed and the memory will be used for short periods of time – that is it would not be appropriate to use statically allocated memory. An example might be to obtain a large memory



buffer for reading and writing files. When **SmartAlloc** is enabled, the memory obtained by `malloc()` will automatically be checked for buffer overwrite (overflow) during the `free()` call, and all `malloc`'ed memory that is not released prior to termination of the program will be reported as Orphaned memory.

15.1.3 Pooled and Non-pooled Memory

In order to facility the handling of arbitrary length filenames and to efficiently handle a high volume of dynamic memory usage, we have implemented routines between the C code and the `malloc` routines. The first is called "Pooled" memory, and is memory, which once allocated and then released, is not returned to the system memory pool, but rather retained in a Bacula memory pool. The next request to acquire pooled memory will return any free memory block. In addition, each memory block has its current size associated with the block allowing for easy checking if the buffer is of sufficient size. This kind of memory would normally be used in high volume situations (lots of `malloc`(s) and `free`(s)) where the buffer length may have to frequently change to adapt to varying filename lengths.

The non-pooled memory is handled by routines similar to those used for pooled memory, allowing for easy size checking. However, non-pooled memory is returned to the system rather than being saved in the Bacula pool. This kind of memory would normally be used in low volume situations (few `malloc`(s) and `free`(s)), but where the size of the buffer might have to be adjusted frequently.

Types of Memory Pool: Currently there are three memory pool types:

- `PM_NOPOOL` – non-pooled memory.
- `PM_FNAME` – a filename pool.
- `PM_MESSAGE` – a message buffer pool.
- `PM_EMMSG` – error message buffer pool.

Getting Memory: To get memory, one uses:

```
| void *get_pool_memory(pool);
```

where **pool** is one of the above mentioned pool names. The size of the memory returned will be determined by the system to be most appropriate for the application.

If you wish non-pooled memory, you may alternatively call:

```
| void *get_memory(size_t size);
```

The buffer length will be set to the size specified, and it will be assigned to the `PM_NOPOOL` pool (no pooling).

Releasing Memory: To free memory acquired by either of the above two calls, use:

```
| void free_pool_memory(void *buffer);
```

where `buffer` is the memory buffer returned when the memory was acquired. If the memory was originally allocated as type `PM_NOPOOL`, it will be released to the system, otherwise, it will be placed on the appropriate Bacula memory pool free chain to be used in a subsequent call for memory from that pool.



Determining the Memory Size: To determine the memory buffer size, use:

```
| size_t sizeof_pool_memory(void *buffer);
```

Resizing Pool Memory: To resize pool memory, use:

```
| void *realloc_pool_memory(void *buffer);
```

The buffer will be reallocated, and the contents of the original buffer will be preserved, but the address of the buffer may change.

Automatic Size Adjustment: To have the system check and if necessary adjust the size of your pooled memory buffer, use:

```
| void *check_pool_memory_size(void *buffer, size_t new-size);
```

where **new-size** is the buffer length needed. Note, if the buffer is already equal to or larger than **new-size** no buffer size change will occur. However, if a buffer size change is needed, the original contents of the buffer will be preserved, but the buffer address may change. Many of the low level Bacula subroutines expect to be passed a pool memory buffer and use this call to ensure the buffer they use is sufficiently large.

Releasing All Pooled Memory: In order to avoid orphaned buffer error messages when terminating the program, use:

```
| void close_memory_pool();
```

to free all unused memory retained in the Bacula memory pool. Note, any memory not returned to the pool via `free_pool_memory()` will not be released by this call.

Pooled Memory Statistics: For debugging purposes and performance tuning, the following call will print the current memory pool statistics:

```
| void print_memory_pool_stats();
```

an example output is:

Pool	Maxsize	Maxused	Inuse
0	256	0	0
1	256	1	0
2	256	1	0





Chapter 16

TCP/IP Network Protocol

16.1 General

This document describes the TCP/IP protocol used by Bacula to communicate between the various daemons and services. The definitive definition of the protocol can be found in `src/lib/bsock.h`, `src/lib/bnet.c` and `src/lib/bnet_server.c`.

Bacula's network protocol is basically a "packet oriented" protocol built on a standard TCP/IP streams. At the lowest level all packet transfers are done with `read()` and `write()` requests on system sockets. Pipes are not used as they are considered unreliable for large serial data transfers between various hosts.

Using the routines described below (`bnet_open`, `bnet_write`, `bnet_recv`, and `bnet_close`) guarantees that the number of bytes you write into the socket will be received as a single record on the other end regardless of how many low level `write()` and `read()` calls are needed. All data transferred are considered to be binary data.

16.2 bnet and Threads

These `bnet` routines work fine in a threaded environment. However, they assume that there is only one reader or writer on the socket at any time. It is highly recommended that only a single thread access any BSOCK packet. The exception to this rule is when the socket is first opened and it is waiting for a job to start. The wait in the Storage daemon is done in one thread and then passed to another thread for subsequent handling.

If you envision having two threads using the same BSOCK, think twice, then you must implement some locking mechanism. However, it probably would not be appropriate to put locks inside the `bnet` subroutines for efficiency reasons.

16.3 bnet_open

To establish a connection to a server, use the subroutine:

BSOCK `*bnet_open(void *jcr, char *host, char *service, int port, int *fatal)` `bnet_open()`, if successful, returns the Bacula sock descriptor pointer to be used in subsequent `bnet_send()` and `bnet_read()` requests. If not successful, `bnet_open()` returns a NULL. If `fatal` is set on return, it means that a fatal error occurred and that you should not repeatedly call `bnet_open()`. Any error message will generally be sent to the JCR.



16.4 bnet_send

To send a packet, one uses the subroutine:

`int bnet_send(BSOCK *sock)` This routine is equivalent to a `write()` except that it handles the low level details. The data to be sent is expected to be in `sock->msg` and be `sock->msglen` bytes. To send a packet, `bnet_send()` first writes four bytes in network byte order than indicate the size of the following data packet. It returns:

- Returns 0 on failure
- Returns 1 on success

In the case of a failure, an error message will be sent to the JCR contained within the `bsock` packet.

16.5 bnet_fsend

This form uses:

`int bnet_fsend(BSOCK *sock, char *format, ...)` and it allows you to send a formatted messages somewhat like `fprintf()`. The return status is the same as `bnet_send`.

16.6 Additional Error information

For additional error information, you can call `is_bnet_error(BSOCK *bsock)` which will return 0 if there is no error or non-zero if there is an error on the last transmission. The `is_bnet_stop(BSOCK *bsock)` function will return 0 if there no errors and you can continue sending. It will return non-zero if there are errors or the line is closed (no more transmissions should be sent).

16.7 bnet_rcv

To read a packet, one uses the subroutine:

`int bnet_rcv(BSOCK *sock)` This routine is similar to a `read()` except that it handles the low level details. `bnet_read()` first reads packet length that follows as four bytes in network byte order. The data is read into `sock->msg` and is `sock->msglen` bytes. If the `sock->msg` is not large enough, `bnet_rcv()` `realloc()` the buffer. It will return an error (-2) if `maxbytes` is less than the record size sent. It returns:

- * Returns number of bytes read
- * Returns 0 on end of file
- * Returns -1 on hard end of file (i.e. network connection close)
- * Returns -2 on error

It should be noted that `bnet_rcv()` is a blocking read.

16.8 bnet_sig

To send a "signal" from one daemon to another, one uses the subroutine:



int bnet_sig(BSOCK *sock, SIGNAL) where SIGNAL is one of the following:

- 1 BNET_EOF - deprecated use BNET_EOD
- 2 BNET_EOD - End of data stream, new data may follow
- 3 BNET_EOD_POLL - End of data and poll all in one
- 4 BNET_STATUS - Request full status
- 5 BNET_TERMINATE - Conversation terminated, doing close()
- 6 BNET_POLL - Poll request, I'm hanging on a read
- 7 BNET_HEARTBEAT - Heartbeat Response requested
- 8 BNET_HB_RESPONSE - Only response permitted to HB
- 9 BNET_PROMPT - Prompt for UA

16.9 bnet_strerror

Returns a formatted string corresponding to the last error that occurred.

16.10 bnet_close

The connection with the server remains open until closed by the subroutine:

```
void bnet_close(BSOCK *sock)
```

16.11 Becoming a Server

The bnet_open() and bnet_close() routines described above are used on the client side to establish a connection and terminate a connection with the server. To become a server (i.e. wait for a connection from a client), use the routine **bnet_thread_server**. The calling sequence is a bit complicated, please refer to the code in bnet_server.c and the code at the beginning of each daemon as examples of how to call it.

16.12 Higher Level Conventions

Within Bacula, we have established the convention that any time a single record is passed, it is sent with bnet_send() and read with bnet_recv(). Thus the normal exchange between the server (S) and the client (C) are:

S: wait for connection	C: attempt connection
S: accept connection	C: bnet_send() send request
S: bnet_recv() wait for request	
S: act on request	
S: bnet_send() send ack	C: bnet_recv() wait for ack



Thus a single command is sent, acted upon by the server, and then acknowledged.

In certain cases, such as the transfer of the data for a file, all the information or data cannot be sent in a single packet. In this case, the convention is that the client will send a command to the server, who knows that more than one packet will be returned. In this case, the server will enter a loop:

```
while ((n=bnet_rcv(bsock)) > 0) {  
    act on request  
}  
if (n < 0)  
    error
```

The client will perform the following:

```
bnet_send(bsock);  
bnet_send(bsock);  
...  
bnet_sig(bsock, BNET_EOD);
```

Thus the client will send multiple packets and signal to the server when all the packets have been sent by sending a zero length record.



Chapter 17

Smart Memory Allocation

Figure 17.1: Smart Memory Allocation with Orphaned Buffer Detection

Few things are as embarrassing as a program that leaks, yet few errors are so easy to commit or as difficult to track down in a large, complicated program as failure to release allocated memory. SMARTALLOC replaces the standard C library memory allocation functions with versions which keep track of buffer allocations and releases and report all orphaned buffers at the end of program execution. By including this package in your program during development and testing, you can identify code that loses buffers right when it's added and most easily fixed, rather than as part of a crisis debugging push when the problem is identified much later in the testing cycle (or even worse, when the code is in the hands of a customer). When program testing is complete, simply recompiling with different flags removes SMARTALLOC from your program, permitting it to run without speed or storage penalties.

In addition to detecting orphaned buffers, SMARTALLOC also helps to find other common problems in management of dynamic storage including storing before the start or beyond the end of an allocated buffer, referencing data through a pointer to a previously released buffer, attempting to release a buffer twice or releasing storage not obtained from the allocator, and assuming the initial contents of storage allocated by functions that do not guarantee a known value. SMARTALLOC's checking does not usually add a large amount of overhead to a program (except for programs which use `realloc()` extensively; see below). SMARTALLOC focuses on proper storage management rather than internal consistency of the heap as checked by the `malloc_debug` facility available on some systems. SMARTALLOC does not conflict with `malloc_debug` and both may be used together, if you wish. SMARTALLOC makes no assumptions regarding the internal structure of the heap and thus should be compatible with any C language implementation of the standard memory allocation functions.

17.0.1 Installing SMARTALLOC

SMARTALLOC is provided as a Zipped archive, [smartall.zip](#); see the download instructions below.

To install SMARTALLOC in your program, simply add the statement:

to every C program file which calls any of the memory allocation functions (`malloc`, `calloc`, `free`, etc.). SMARTALLOC must be used for all memory allocation with a program, so include



file for your entire program, if you have such a thing. Next, define the symbol SMARTALLOC in the compilation before the inclusion of smartall.h. I usually do this by having my Makefile add the “-DSMARTALLOC” option to the C compiler for non-production builds. You can define the symbol manually, if you prefer, by adding the statement:

```
#define SMARTALLOC
```

At the point where your program is all done and ready to relinquish control to the operating system, add the call:

```
sm_dump(datadump);
```

where *datadump* specifies whether the contents of orphaned buffers are to be dumped in addition printing to their size and place of allocation. The data are dumped only if *datadump* is nonzero, so most programs will normally use “sm_dump(0);”. If a mysterious orphaned buffer appears that can't be identified from the information this prints about it, replace the statement with “sm_dump(1);”. Usually the dump of the buffer's data will furnish the additional clues you need to excavate and extirpate the elusive error that left the buffer allocated.

Finally, add the files “smartall.h” and “smartall.c” from this release to your source directory, make dependencies, and linker input. You needn't make inclusion of smartall.c in your link optional; if compiled with SMARTALLOC not defined it generates no code, so you may always include it knowing it will waste no storage in production builds. Now when you run your program, if it leaves any buffers around when it's done, each will be reported by sm_dump() on stderr as follows:

```
| Orphaned buffer:      120 bytes allocated at line 50 of gutshot.c
```

17.0.2 Squelching a SMARTALLOC

Usually, when you first install SMARTALLOC in an existing program you'll find it nattering about lots of orphaned buffers. Some of these turn out to be legitimate errors, but some are storage allocated during program initialisation that, while dynamically allocated, is logically static storage not intended to be released. Of course, you can get rid of the complaints about these buffers by adding code to release them, but by doing so you're adding unnecessary complexity and code size to your program just to silence the nattering of a SMARTALLOC, so an escape hatch is provided to eliminate the need to release these buffers.

Normally all storage allocated with the functions malloc(), calloc(), and realloc() is monitored by SMARTALLOC. If you make the function call:

```
|      sm_static(1);
```

you declare that subsequent storage allocated by malloc(), calloc(), and realloc() should not be considered orphaned if found to be allocated when sm_dump() is called. I use a call on “sm_static(1);” before I allocate things like program configuration tables so I don't have to add code to release them at end of program time. After allocating unmonitored data this way, be sure to add a call to:

```
|      sm_static(0);
```

to resume normal monitoring of buffer allocations. Buffers allocated while sm_static(1) is in effect are not checked for having been orphaned but all the other safeguards provided by SMARTALLOC remain in effect. You may release such buffers, if you like; but you don't have to.



17.0.3 Living with Libraries

Some library functions for which source code is unavailable may gratuitously allocate and return buffers that contain their results, or require you to pass them buffers which they subsequently release. If you have source code for the library, by far the best approach is to simply install SMARTALLOC in it, particularly since this kind of ill-structured dynamic storage management is the source of so many storage leaks. Without source code, however, there's no option but to provide a way to bypass SMARTALLOC for the buffers the library allocates and/or releases with the standard system functions.

For each function `xxx` redefined by SMARTALLOC, a corresponding routine named “`actuallyxxx`” is furnished which provides direct access to the underlying system function, as follows:

Table 17.1: Systems functions

Standard function	Direct access function
<code>malloc(<i>size</i>)</code>	<code>actuallymalloc(<i>size</i>)</code>
<code>calloc(<i>nelem</i>, <i>elsize</i>)</code>	<code>actuallycalloc(<i>nelem</i>, <i>elsize</i>)</code>
<code>realloc(<i>ptr</i>, <i>size</i>)</code>	<code>actuallyrealloc(<i>ptr</i>, <i>size</i>)</code>
<code>free(<i>ptr</i>)</code>	<code>actuallyfree(<i>ptr</i>)</code>

For example, suppose there exists a system library function named “`getimage()`” which reads a raster image file and returns the address of a buffer containing it. Since the library routine allocates the image directly with `malloc()`, you can't use SMARTALLOC's `free()`, as that call expects information placed in the buffer by SMARTALLOC's special version of `malloc()`, and hence would report an error. To release the buffer you should call `actuallyfree()`, as in this code fragment:

```
struct image *ibuf = getimage("ratpack.img");
display_on_screen(ibuf);
actuallyfree(ibuf);
```

Conversely, suppose we are to call a library function, “`putimage()`”, which writes an image buffer into a file and then releases the buffer with `free()`. Since the system `free()` is being called, we can't pass a buffer allocated by SMARTALLOC's allocation routines, as it contains special information that the system `free()` doesn't expect to be there. The following code uses `actuallymalloc()` to obtain the buffer passed to such a routine.

```
struct image *obuf =
    (struct image *) actuallymalloc(sizeof(struct image));
dump_screen_to_image(obuf);
putimage("screndump.img", obuf); /* putimage() releases obuf */
```

It's unlikely you'll need any of the “actually” calls except under very odd circumstances (in four products and three years, I've only needed them once), but they're there for the rare occasions that demand them. Don't use them to subvert the error checking of SMARTALLOC; if you want to disable orphaned buffer detection, use the `sm_static(1)` mechanism described above. That way you don't forfeit all the other advantages of SMARTALLOC as you do when using `actuallymalloc()` and `actuallyfree()`.



17.0.4 SMARTALLOC Details

When you include “smartall.h” and define SMARTALLOC, the following standard system library functions are redefined with the #define mechanism to call corresponding functions within smartall.c instead. (For details of the redefinitions, please refer to smartall.h.)

```
void *malloc(size_t size)
void *calloc(size_t nelem, size_t elsize)
void *realloc(void *ptr, size_t size)
void free(void *ptr)
void cfree(void *ptr)
```

cfree() is a historical artifact identical to free().

In addition to allocating storage in the same way as the standard library functions, the SMARTALLOC versions expand the buffers they allocate to include information that identifies where each buffer was allocated and to chain all allocated buffers together. When a buffer is released, it is removed from the allocated buffer chain. A call on sm_dump() is able, by scanning the chain of allocated buffers, to find all orphaned buffers. Buffers allocated while sm_static(1) is in effect are specially flagged so that, despite appearing on the allocated buffer chain, sm_dump() will not deem them orphans.

When a buffer is allocated by malloc() or expanded with realloc(), all bytes of newly allocated storage are set to the hexadecimal value 0x55 (alternating one and zero bits). Note that for realloc() this applies only to the bytes added at the end of buffer; the original contents of the buffer are not modified. Initializing allocated storage to a distinctive nonzero pattern is intended to catch code that erroneously assumes newly allocated buffers are cleared to zero; in fact their contents are random. The calloc() function, defined as returning a buffer cleared to zero, continues to zero its buffers under SMARTALLOC.

Buffers obtained with the SMARTALLOC functions contain a special sentinel byte at the end of the user data area. This byte is set to a special key value based upon the buffer's memory address. When the buffer is released, the key is tested and if it has been overwritten an assertion in the free function will fail. This catches incorrect program code that stores beyond the storage allocated for the buffer. At free() time the queue links are also validated and an assertion failure will occur if the program has destroyed them by storing before the start of the allocated storage.

In addition, when a buffer is released with free(), its contents are immediately destroyed by overwriting them with the hexadecimal pattern 0xAA (alternating bits, the one's complement of the initial value pattern). This will usually trip up code that keeps a pointer to a buffer that's been freed and later attempts to reference data within the released buffer. Incredibly, this is *legal* in the standard Unix memory allocation package, which permits programs to free() buffers, then raise them from the grave with realloc(). Such program “logic” should be fixed, not accommodated, and SMARTALLOC brooks no such Lazarus buffer“ nonsense.

Some C libraries allow a zero size argument in calls to malloc(). Since this is far more likely to indicate a program error than a defensible programming stratagem, SMARTALLOC disallows it with an assertion.

When the standard library realloc() function is called to expand a buffer, it attempts to expand the buffer in place if possible, moving it only if necessary. Because SMARTALLOC must place its own private storage in the buffer and also to aid in error detection, its version of realloc() always moves and copies the buffer except in the trivial case where the size of the buffer is not being changed. By forcing the buffer to move on every call and destroying the contents of the old buffer when it is released, SMARTALLOC traps programs which keep pointers into a buffer across a call on realloc() which may move it. This strategy may prove very costly to programs which make extensive use of realloc(). If this proves to be a problem, such programs may wish to use actuallymalloc(), actuallyrealloc(), and actuallyfree() for such frequently-adjusted buffers, trading error detection for performance. Although not specified in the System V Interface Definition, many C library implementations of realloc()



permit an old buffer argument of NULL, causing `realloc()` to allocate a new buffer. The SMARTALLOC version permits this.

17.0.5 When SMARTALLOC is Disabled

When SMARTALLOC is disabled by compiling a program with the symbol SMARTALLOC not defined, calls on the functions otherwise redefined by SMARTALLOC go directly to the system functions. In addition, compile-time definitions translate calls on the "actually..." functions into the corresponding library calls; "actuallymalloc(100)", for example, compiles into "malloc(100)". The two special SMARTALLOC functions, `sm_dump()` and `sm_static()`, are defined to generate no code (hence the null statement). Finally, if SMARTALLOC is not defined, compilation of the file `smartall.c` generates no code or data at all, effectively removing it from the program even if named in the link instructions.

Thus, except for unusual circumstances, a program that works with SMARTALLOC defined for testing should require no changes when built without it for production release.

17.0.6 The `alloc()` Function

Many programs I've worked on use very few direct calls to `malloc()`, using the identically declared `alloc()` function instead. Alloc detects out-of-memory conditions and aborts, removing the need for error checking on every call of `malloc()` (and the temptation to skip checking for out-of-memory).

As a convenience, SMARTALLOC supplies a compatible version of `alloc()` in the file `alloc.c`, with its definition in the file `alloc.h`. This version of `alloc()` is sensitive to the definition of SMARTALLOC and cooperates with SMARTALLOC's orphaned buffer detection. In addition, when SMARTALLOC is defined and `alloc()` detects an out of memory condition, it takes advantage of the SMARTALLOC diagnostic information to identify the file and line number of the call on `alloc()` that failed.

17.0.7 Overlays and Underhandedness

String constants in the C language are considered to be static arrays of characters accessed through a pointer constant. The arrays are potentially writable even though their pointer is a constant. SMARTALLOC uses the compile-time definition `./smartall.wml` to obtain the name of the file in which a call on buffer allocation was performed. Rather than reserve space in a buffer to save this information, SMARTALLOC simply stores the pointer to the compiled-in text of the file name. This works fine as long as the program does not overlay its data among modules. If data are overlaid, the area of memory which contained the file name at the time it was saved in the buffer may contain something else entirely when `sm_dump()` gets around to using the pointer to edit the file name which allocated the buffer.

If you want to use SMARTALLOC in a program with overlaid data, you'll have to modify `smartall.c` to either copy the file name to a fixed-length field added to the `abufhead` structure, or else allocate storage with `malloc()`, copy the file name there, and set the `abfname` pointer to that buffer, then remember to release the buffer in `sm_free`. Either of these approaches are wasteful of storage and time, and should be considered only if there is no alternative. Since most initial debugging is done in non-overlaid environments, the restrictions on SMARTALLOC with data overlaying may never prove a problem. Note that conventional overlaying of code, by far the most common form of overlaying, poses no problems for SMARTALLOC; you need only be concerned if you're using exotic tools for data overlaying on MS-DOS or other address-space-challenged systems.

Since a C language "constant" string can actually be written into, most C compilers generate a unique copy of each string used in a module, even if the same constant string appears many



times. In modules that contain many calls on allocation functions, this results in substantial wasted storage for the strings that identify the file name. If your compiler permits optimization of multiple occurrences of constant strings, enabling this mode will eliminate the overhead for these strings. Of course, it's up to you to make sure choosing this compiler mode won't wreak havoc on some other part of your program.

17.0.8 Test and Demonstration Program

A test and demonstration program, `smtest.c`, is supplied with SMARTALLOC. You can build this program with the Makefile included. Please refer to the comments in `smtest.c` and the Makefile for information on this program. If you're attempting to use SMARTALLOC on a new machine or with a new compiler or operating system, it's a wise first step to check it out with `smtest` first.

17.0.9 Invitation to the Hack

SMARTALLOC is not intended to be a panacea for storage management problems, nor is it universally applicable or effective; it's another weapon in the arsenal of the defensive professional programmer attempting to create reliable products. It represents the current state of evolution of expedient debug code which has been used in several commercial software products which have, collectively, sold more than third of a million copies in the retail market, and can be expected to continue to develop through time as it is applied to ever more demanding projects.

The version of SMARTALLOC here has been tested on a Sun SPARCStation, Silicon Graphics Indigo2, and on MS-DOS using both Borland and Microsoft C. Moving from compiler to compiler requires the usual small changes to resolve disputes about prototyping of functions, whether the type returned by buffer allocation is `char *` or `void *`, and so forth, but following those changes it works in a variety of environments. I hope you'll find SMARTALLOC as useful for your projects as I've found it in mine.

17.1 Download smartall.zip (Zipped archive)

SMARTALLOC is provided as [smartall.zip](#), a [Zipped](#) archive containing source code, documentation, and a Makefile to build the software under Unix.

17.1.1 Copying

SMARTALLOC is in the public domain. Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, without any conditions or restrictions. This software is provided "as is" without express or implied warranty.

by John Walker October 30th, 1998



Appendices





Appendix A

Acronyms

ANSI	American National Standards Institute
ASCII	American Standard Code for Information Interchange
DNS	Domain Name Service
EOF	End Of File
FLA	Fiduciary Licence Agreement
GNU	Gnu is not Unix
GPL	Gnu General Public License
GUI	Graphical User Interface
MD	Message Digest
POSIX	Portable Operating System Interface
SHA	Secure Hash Algorithm
SMTP	Simple Mail Transfer Protocol
SQL	Structured Query Language
SSL	Secure Socket Layer
SVN	Subversion



Index

Symbols

-csignature	122
-dinput_text	122
GUI Interface is Difficult	101
Bacula	
Building the Test	115
Bacula Regression Testing	113
MySQL	
Installing and Configuring	68
PostgreSQL	
Installing and Configuring	68
Download smartall.zip (Zipped archive)	136

A

Acknowledgements	122
Adding a New Test	118
Adding a Unittest	119
Additional Error information	128
all-non-root-tape-tests	117
all-root-tests	117
all-tape-and-file-tests	117
all_non-root-tests	117
Alloc() Function	135
ALSO	
SEE	122
API Changes	
Bnet	110
API Implimentation	
TLS	108
Archive	
Download smartall.zip Zipped	136
Download md5.zip Zipped	122
Attributes	
Unix File	90
Authentication Negotiation	111

B

Backup	
Commands Received from the Director for a	61
Bacula Memory Management	123
Bacula Porting Notes	97
Becoming a Server	129
Begin Session Label	82
Block	81
Block Header	82, 83
Blocking State	
Socket	
Manipulating	110
Bnet and Threads	127
Bnet API Changes	110
Bnet_close	129
Bnet_fsend	128
Bnet_open	127
Bnet_recv	128
Bnet_send	128
Bnet_sig	128
Bnet_strerror	129



Bugs	122
Bugs Database.....	18
Building the Test Bacula	115

C

Catalog	
Internal Bacula	68
Catalog Services	67
Classes	
Message	15
Command and Control Information	57
Command Line Message Digest Utility	121
Commands Received from the Director for a Backup	61
Commands Received from the Director for a Restore	62
Connections	
Manipulating TLS.....	109
Contexts	
Manipulating TLS.....	108
Conventions	
Higher Level	129
Copying	122, 136

D

Daemon	
Director Services	59
File Services	61
Protocol Used Between the Director and the File	56
Protocol Used Between the Director and the Storage	56
Save Protocol Between the File Daemon and the Storage	57
Daemon Protocol	55
Data Information	57
Data Record	82
Database	
Bugs.....	18
Database Table Design	68
Database Tables	69
DataSize	85
Debug Messages	15
Debugger	119
Definitions.....	81
Description	121
Design	
Database Table	68
Storage Daemon	63
Details	
SMARTALLOC	134
Detection	
Smart Memory Allocation With Orphaned Buffer	131
Difficult	
GUI Interface is	101
Directives	
TLS Configuration	107
Director Services Daemon	59
Directory Structure	118
Disabled	
When SMARTALLOC is	135
Do Use Whenever Possible	10
Don'ts	14
Download md5.zip (Zipped archive)	122
Dynamically Allocated Memory.....	123



E	
End Session Label	82
Error Messages	16
Exit Status	122
Extended-Attributes	90
F	
Fails	
If a Test	117
File Services Daemon	61
File-Attributes	90
FileIndex	82, 84, 90
Filename	90
Filenames and Maximum Filename Length	67
Files	122
Format	
Old Depreciated Tape	91
Overall	83
Overall Storage	87
Storage Daemon File Output	83
Storage Media Output	81
Volume Label	85
Function	
alloc	135
G	
General	67, 81, 101, 123, 127
General	53, 55
General Daemon Protocol	55
Git	19
Repo	19
Git Usage	19
H	
Hack	
Invitation to the	136
Hand	
Running the Tests by	118
Header	
Block	83
Record	84
Version 2 Record	85
Version BB02 Block	85
Higher Level Conventions	129
I	
If a Test Fails	117
Implementing a Bacula GUI Interface	101
Indenting Standards	11
Information	
Additional Error	128
Command and Control	57
Data	57
Initialization and Cleanup	
Library	108
Installing and Configuring MySQL	68
Installing and Configuring PostgreSQL	68
Installing SMARTALLOC	131
Interface	
Implementing a Bacula GUI	101



Internal Bacula Catalog	68
Introduction	
SD Design	63
TLS	107
Invitation to the Hack	136
J	
Job	
Sequence of Creation of Records for a Save	68
Job Messages	17
JobId	81
L	
Label	
Session	86
Learning Git	20
Length	
Filenames and Maximum Filename	67
Libraries	
Living with	133
Library Initialization and Cleanup	108
Link	90
Living with Libraries	133
Low Level Network Protocol	55
M	
Management	
Bacula Memory	123
Manipulating Socket Blocking State	110
Memory	
Dynamically Allocated	123
Pooled and Non-pooled	124
Statically Allocated	123
Memory Messages	17
Message Classes	15
Messages	
Debug	15
Error	16
Job	17
Memory	17
Queued Job	17
Minimal Code in Console Program	101
N	
Name	121
Negotiating a TLS Connection	110
Negotiation	
TLS Authentication	111
Notes	
Bacula Porting	97
O	
Old Depreciated Tape Format	91
Options	122
Other Tests	117
Outline	
SD Development	63
Overall Format	83
Overall Storage Format	87
Overlays and Underhandedness	135

**P**

Parameters	
Setting the Configuration	114
Platform Requirements	53
Platform Support	53
Pooled and Non-pooled Memory	124
Porting	
Steps to Take for	98
Porting Requirements	97
Possible	
Do Use Whenever	10
Program	
Minimal Code in Console	101
Test and Demonstration	136
Protocol	
Daemon	55
General Daemon	55
Low Level Network	55
TCP/IP Network	127
Protocol Used Between the Director and the File Daemon	56
Protocol Used Between the Director and the Storage Daemon	56

Q

Queued Job Messages	17
---------------------------	----

R

Record	81
Record Header	82, 84
Regression	
Running the Disk Only	115
Requests	
SD Append	64
SD Read	65
Requirements	
Platform	53
Porting	97
Restore	
Commands Received from the Director for a	62
Running a Single Test	118
Running the Disk Only Regression	115
Running the Regression Script	113
Running the Tests by Hand	118

S

Save Protocol Between the File Daemon and the Storage Daemon	57
Script	
Running the Regression	113
SD Append Requests	64
SD Connections and Sessions	64
SD Data Structures	65
SD Design Introduction	63
SD Development Outline	63
SD Read Requests	65
See Also	122
Sequence of Creation of Records for a Save Job	68
Serialization	83
Server	
Becoming a	129
Services	
Catalog	67



Session	81
Session Label	86
Sessions	
SD Connections and	64
Setting the Configuration Parameters	114
Setting up Regression Testing	113
Setting up your SQL engine	115
Smart Memory Allocation With Orphaned Buffer Detection	131
SMARTALLOC	
Installing	131
Squelching a	132
SMARTALLOC Details	134
Socket Blocking State	
Manipulating	110
SPAN class	64, 65
Squelching a SMARTALLOC	132
Standards	
Indenting	11
Statically Allocated Memory	123
Status	
Exit	122
Steps to Take for Porting	98
Storage Daemon Design	63
Storage Daemon File Output Format	83
Storage Media Output Format	81
Stream	82, 84
Structure	
Directory	118
Support	
Platform	53
Synopsis	121

T

Tabbing	14
Tables	
Database	69
TCP/IP Network Protocol	127
Test	
Adding a New	118, 119
Testing a Binary Installation	117
Writing a Regression	118
Test and Demonstration Program	136
Testing	
Bacula Regression	113
Tests	
Other	117
Threads	
bnet and	127
TLS	107
TLS API Implimentation	108
TLS Configuration Directives	107
TLS Connection	
Negotiating	110
TLS Connection Manipulation	109
TLS Context Manipulation	108
TLS Introduction	107
TLS Post-Connection Verification	109
Type	90

**U**

Underhandedness	
Overlays and	135
Unix File Attributes	90
Utility	
Command Line Message Digest	121

V

Verification	
TLS Post-Connection	109
Version 2 Record Header	85
Version BB02 Block Header	85
VolSessionId	82, 84
VolSessionTime	82, 84
Volume Label	82
Volume Label Format	85

W

When SMARTALLOC is Disabled	135
Writing a Regression Test	118

