



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

Kern Sibbald

November 22, 2005
This is the Bacula Developer's Guide

Contents

Bacula Developer Notes	4
General	4
Basic CVS Usage	6
Daemon Protocol	25
General	25
Low Level Network Protocol	25
General Daemon Protocol	25
Protocol Used Between the Director and the Storage Daemon	26
Protocol Used Between the Director and the File Daemon . .	27
Save Protocol Between the File Daemon and the Storage Daemon	28
Director Services Daemon	31
File Services Daemon	32
Commands Received from the Director for a Backup	33
Commands Received from the Director for a Restore	33
Storage Daemon Design	34
SD Design Introduction	34
SD Development Outline	34
SD Connections and Sessions	35
SD Data Structures	38

Catalog Services	40
General	40
Sequence of Creation of Records for a Save Job	42
Database Tables	43
Storage Media Output Format	56
General	56
Definitions	56
Storage Daemon File Output Format	59
Overall Format	59
Serialization	60
Block Header	60
Record Header	61
Version BB02 Block Header	63
Version 2 Record Header	63
Volume Label Format	63
Session Label	64
Overall Storage Format	65
Unix File Attributes	70
Old Depreciated Tape Format	72
Bacula Porting Notes	77
Porting Requirements	77
Steps to Take for Porting	78
Implementing a Bacula GUI Interface	81
General	81
0.1 TLS	83

TLS	83
0.1.1 Introduction to TLS	83
TLS Introduction	83
0.1.2 New Configuration Directives	83
New Configuration Directives	83
0.1.3 TLS API Implementation	85
TLS API Implementation	85
0.1.4 Bnet API Changes	87
Bnet API Changes	87
0.1.5 Authentication Negotiation	89
Authentication Negotiation	89
Bacula Regression Testing	90
General	90
Running the Regression Script	90
Writing a Regression Test	94
.	96
Command Line Message Digest Utility	96
Download md5.zip (Zipped archive)	98
Bacula Memory Management	100
General	100
TCP/IP Network Protocol	104
General	104
bnet and Threads	104
bnet_open	104
bnet_send	105

bnet_fsend	105
Additional Error information	105
bnet_recv	106
bnet_sig	106
bnet_strerror	107
bnet_close	107
Becoming a Server	107
Higher Level Conventions	107
Smart Memory Allocation With Orphaned Buffer Detection	109
Download smartall.zip (Zipped archive)	117

List of Figures

Smart Memory Allocation with Orphaned Buffer Detection	109
--	-----

List of Tables

Message Error Code Classes	22
Filename Table Layout	43
Path Table Layout	43
File Table Layout	44
Job Table Layout	45
Job Types	46
Job Statuses	46
File Sets Table Layout	47
JobMedia Table Layout	47
Media Table Layout	48
Pool Table Layout	49
Client Table Layout	50
Unsaved Files Table Layout	50
Counter Table Layout	51
Version Table Layout	51
Base Files Table Layout	51
File Attributes	71

Bacula Developer Notes

General

This document is intended mostly for developers and describes the the general framework of making Bacula source changes.

Contributions

Contributions from programmers 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 an examples or contributions directory.

The second class of contributions are those which will be integrated with Bacula and become an essential part. 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.

Patches

Subject to the copyright assignment described below, your patches should be sent in **diff -u** format relative to the current contents of the Source Forge CVS, which is the easiest for me to understand. If you have checked out the source with CVS, you can get a diff using:

```
cvs diff -u \lt{}change.patch
```

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 CVS access so that you can commit your changes directly to the CVS repository. To do so, you will need a userid on Source Forge.

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 assigned to Kern Sibbald as in the current

code. Note, prior to November 2004, the code was copyrighted by Kern Sibbald and John Walker.

Your name should be clearly indicated as the author of the code, and you must be extremely careful not to violate any copyrights or use other people's code without acknowledging it. The purpose of this requirement is to avoid future copyright, patent, or intellectual property problems. To understand on possible source of future problems, please examine the difficulties Mozilla is (was?) having finding previous contributors at <http://www.mozilla.org/MPL/missing.html>. The other important issue is to avoid copyright, patent, or intellectual property violations as are currently (May 2003) being claimed by SCO against IBM.

Although the copyright will be held by Kern, 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. In fact, the paperwork associated with making contributions to the Free Software Foundation, was for me unsurmountable.

If you have any doubts about this, please don't hesitate to ask. Our (John and my) track records with Autodesk are easily available; early programmers/founders/contributors and later employees had substantial shares of the company, and no one founder had a controlling part of the company. Even though Microsoft created many millionaires among early employees, the politics of Autodesk (during our time at the helm) is in stark contrast to Microsoft where the majority of the company is still tightly held among a few.

Items not needing a copyright assignment are: most small changes, enhancements, or bug fixes of 5-10 lines of code, and documentation.

Copyright Assignment

Since this is not a commercial enterprise, and I prefer to believe in everyone's good faith, developers can assign the copyright by explicitly acknowledging that they do so in their first submission. This is sufficient if the developer is independent, or an employee of a not-for-profit organization or a university. Any developer that wants to contribute and is employed by a company must get a copyright assignment from his employer. This is to avoid misunderstandings between the employee, the company, and the Bacula project.

Corporate Assignment Statement

The following statement must be filled out by the employer, signed, and mailed to my address (please ask me for my address and I will email it – I'd prefer not to include it here).

Copyright release and transfer statement.

<On company letter head>

To: Kern Sibbald

Concerning: Copyright release and transfer

<Company, Inc> is hereby agrees that <names-of-developers> and other employees of <Company, Inc> are authorized to develop code for and contribute code to the Bacula project for an undetermined period of time. The copyright as well as all other rights to and interests in such contributed code are hereby transferred to Kern Sibbald.

Signed in <City, Country> on <Date>:

<Name of Person>, <Position in Company, Inc>

This release/transfer statement must be sent to: Kern Sibbald Address-to-be-given

Basic CVS Usage

The Bacula CVS is kept at Source Forge. If you are not a developer, you only be able to access the public CVS, which runs about 6 hours behind the developer's CVS. If you are a developer, then you will have immediate access to "the" CVS and any changes ("commit") that you make will be immediately seen by other developers.

For developer access to the CVS, go to the Bacula page on Source Forge and click on the CVS menu item. Then follow the instructions for installing your SSH public key. It may take a few hours until the key is actually installed, and then you will have access to the CVS.

The Bacula CVS is divided into the following CVS "projects" or "modules".

bacula	(Bacula source code)
docs	(Bacula documentation)

gui	(Some of the GUI programs that do not use the Bacula core code)
rescue	(Bacula CDRom rescue code)
regress	(Bacula regression scripts)
ryol	(Roll your own Linux -- incomplete project)

Most of you will want either the Bacula source (bacula) and/or the documents (docs).

To get the source for a project, you must check it out ("checkout"), which you do usually once.

The first time you checkout the code for each project, you will need to tell the cvs program where the CVS repository is. You do so by doing:

```
export CVS_RSH=ssh
export CVSROOT=:ext:<nnnn>@cvs.bacula.sourceforge.net:/cvsroot/bacula
```

where you replace <nnnn> by your Source Forge user name.

Then do

```
cvs -z3 checkout -d <directory> bacula
```

where you replace <directory> by the name of the directory you want to contain the Bacula source code. If you want the docs, replace the word "bacula" on the above line by "docs" and be sure to put it in a different directory. The -z3 just tells CVS to use compression during the transmission, which makes things go faster.

The above command should generate output that looks a bit like the following:

```
cvs checkout: Updating bacula
U bacula/Makefile
U bacula/bar.c
U bacula/foo.c
U bacula/main.c
...
```

Let's assume you used the name "bacula" for the directory, so your command was:

```
cvs -z3 checkout -d bacula bacula
```

When the command is done, simply do:

```
cd bacula
```

and then build Bacula. You will notice a lot of extra CVS directories in the source code. Don't ever change or delete them, or you will mess up your copy of the project. Also, do not rename or delete any of the files in your copy of the repository or you may have problems. Any files that you change will remain only on your system until you do a "commit", which is explained later.

Let's say you are working in the directory scripts. You would then do:

```
cd scripts  
(edit some files)
```

when you are happy with your changes, you can do the following:

```
cd bacula (to your top level directory)  
cvs diff -u >my-changes.patch
```

When the command is done, you can look in the file my-changes.patch and you will see all the changes you have made to your copy of the repository. Make sure that you understand all the changes that it reports before proceeding. If you modified files that you do not want to commit to the main repository, you can simply delete them from your local directory, and they will be restored from the repository with the "cvs update" that is shown below. Normally, you should not find changes to files that you do not want to commit, and if you find yourself in that position a lot, you are probably doing something wrong.

Let's assume that now you want to commit your changes to the main CVS repository.

First do:

```
cvs bacula  
cvs update
```

When you do this, it will pull any changes made by other developers into your local copy of the repository, and it will check for conflicts. If there are any, it will tell you, and you will need to resolve them. The problems of resolving conflicts are a bit more than this document can cover, but you can examine the files it claims have conflicts and look for <<<< or look in the .rej files that it creates. If you have problems, just ask on the developer's list.

Note, doing the above "cvs update" is not absolutely necessary. There are times when you may be working on code and you want to commit it, but you explicitly do not want to move up to the latest version of the code in the CVS. If that is the case, you can simply skip the "cvs update" and do the commit shown below. If the commit fails because of a conflict, it will tell you, and you must resolve the conflict before it will permit you to do the commit.

Once your local copy of the repository has been updated, you can now commit your changes:

```
cvs commit -m "Some comment about what you changed"
```

or if you really only want to commit a single file, you can do:

```
cvs commit -m "comment" scripts/file-I-edited
```

Note, if you have done a build in your directory, or you have added other new files, the commit will update only the files that are actually in the repository. For example, none of the object files are stored in the repository, so when you do a commit, those object files will simply be ignored.

If you want to add new files or remove files from the main CVS repository, and you are not experienced with CVS, please ask Kern to do it. If you follow the simple steps above, it is unlikely that you will do any damage to the repository, and if you do, it is always possible for us to recover, but it can be painful.

If you are only working in one subdirectory of say the bacula project, for example, the scripts directory, you can do your commit from that subdirectory, and only the changes in that directory and all its subdirectories will be committed. This can be helpful for translators. If you are doing a French translation, you will be working in docs/manual-fr, and if you are always cd'ed into that directory when doing your commits, your commit will effect only that directory. As long as you are careful only to change files that you want changed, you have little to worry about.

The Development Cycle

As I noted in the 1.38 ReleaseNotes, version 1.38 was different from prior versions because it had a lot more contributions. I expect that this trend will continue. As a consequence, I am going to modify how I normally do development, and instead of making a list of all the features that I will implement in the next version, I will personally sign up for one (maybe two) projects at a time, and when they are complete, I will release a new version.

The difference is that I 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 (1.38 had too many new features for me to handle correctly).

I expect that 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 :-)

Feature Requests:

In addition, I would like to "formalize" the feature requests a bit.

Instead of me maintaining an informal list of everything I run into (kernstodo), I would like to 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. I'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 I receive the Feature Request, I will either accept it, send it back asking for clarification, send it to the email list asking for opinions, or reject it.

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 I cannot imagine finding someone to implement probably won't be accepted.
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 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, I will generally solicit Feature Request input for the next version, and by way of this email, I 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, I 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 – I am 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 riege 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

=====

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` strip 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.

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.

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)”.

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.

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.

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.

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)
        |- cats     (SQL catalog database interface directory)
        |- console  (bacula user agent directory)
        |- dird     (Director daemon)
        |- filed    (Unix File daemon)
            |- win32 (Win32 files to make bacula-fd be a service)
        |- findlib  (Unix file find library for File daemon)
        |- gnome-console (GNOME version of console program)
        |- lib      (General Bacula library)
        |- stored   (Storage daemon)
        |- tconsole (Tcl/tk console program -- not yet working)
        |- testprogs (test programs -- normally only in Kern's tree)
        |- tools    (Various tool programs)
        |- win32    (Native Win32 File daemon)
            |- baculafd (Visual Studio project file)
            |- 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)
    |- wx-console      (wxWidgets console main source program)

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)

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 rescue:
|- rescue            (Bacula rescue CDROM)
  |- linux            (Linux rescue CDROM)
  |- cdrom           (Linux rescue CDROM code)
  ...
  |- solaris          (Solaris rescue -- incomplete)
  |- freebsd         (FreeBSD rescue -- incomplete)

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

```

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.

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.

I 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.

Do Not Use

- STL – it is totally incomprehensible.

Avoid if Possible

- Returning a malloc'ed buffer from a subroutine – someone will forget to release it.
- Using reference variables – it allows subroutines to create side effects.
- Heap allocation (malloc) unless needed – it is expensive.
- 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)

Do Use Whenever Possible

- Locking and unlocking within a single subroutine.
- Malloc and free within a single subroutine.
- Comments and global explanations on what your code or algorithm does.

Indenting Standards

I cannot stand code indented 8 columns at a time. This makes the code unreadable. Even 4 at a time uses a lot of space, so I 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 -nbfa
```

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. Originally I indented case clauses under a switch(), but now I 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 you 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.

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).

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));
```

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.

Message Classes

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

Debug Messages

Debug messages are designed to be turned on at a specified debug level and are always sent to `STDOUT`. They 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);
```

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:

`EmsgN(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:

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

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, "bdird<filed: bad response from Filed to %s
command: %d %s\n", cmd, n, strerror(errno));
```

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. `jbr;` 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.

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.

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.

Daemon Protocol

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.

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. I would not be appropriate for efficiency reasons to make every call to the BSOCK routines lock and unlock the packet.

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.

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 an 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).

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

```

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.

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>

```

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.

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 ...

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 Directory (this, of course, depends on how the Message resource is defined).

Commands Received from the Director for a Backup

To be written ...

Commands Received from the Director for a Restore

To be written ...

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.

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.

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.

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.

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.

SD Read Requests

Read open session = <JobId> <Volume-id> <start-file> <start-block> <end-file> <en
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 simultane-
ous 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

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

Catalog Services

General

This chapter is intended to be a technical discussion of the Catalog services and as such is not targeted at end users but rather at developers and system administrators that want or need to know more of the working details of Bacula.

The Bacula Catalog services consist of the programs that provide the SQL database engine for storage and retrieval of all information concerning files that were backed up and their locations on the storage media.

We have investigated the possibility of using the following SQL engines for Bacula: Beagle, mSQL, GNU SQL, PostgreSQL, SQLite, Oracle, and MySQL. Each presents certain problems with either licensing or maturity. At present, we have chosen for development purposes to use MySQL, PostgreSQL and SQLite. MySQL was chosen because it is fast, proven to be reliable, widely used, and actively being developed. MySQL is released under the GNU GPL license. PostgreSQL was chosen because it is a full-featured, very mature database, and because Dan Langille did the Bacula driver for it. PostgreSQL is distributed under the BSD license. SQLite was chosen because it is small, efficient, and can be directly embedded in Bacula thus requiring much less effort from the system administrator or person building Bacula. In our testing SQLite has performed very well, and for the functions that we use, it has never encountered any errors except that it does not appear to handle databases larger than 2GBytes. That said, we would not recommend it for serious production use.

The Bacula SQL code has been written in a manner that will allow it to be easily modified to support any of the current SQL database systems on the market (for example: mSQL, iODBC, unixODBC, Solid, OpenLink ODBC, EasySoft ODBC, InterBase, Oracle8, Oracle7, and DB2).

If you do not specify either `--with-mysql` or `--with-postgresql` or `--with-sqlite` on the `./configure` line, Bacula will use its minimalist internal database. This database is kept for build reasons but is no longer supported. Bacula requires one of the three databases (MySQL, PostgreSQL, or SQLite) to run.

Filenames and Maximum Filename Length

In general, either MySQL, PostgreSQL or SQLite permit storing arbitrary long path names and file names in the catalog database. In practice, there still may be one or two places in the Catalog interface code that restrict the maximum path length to 512 characters and the maximum file name length to 512 characters. These restrictions are believed to have been removed. Please note, these restrictions apply only to the Catalog database and thus to your ability to list online the files saved during any job. All information received and stored by the Storage daemon (normally on tape) allows and handles arbitrarily long path and filenames.

Installing and Configuring MySQL

For the details of installing and configuring MySQL, please see the Installing and Configuring MySQL chapter of this manual.

Installing and Configuring PostgreSQL

For the details of installing and configuring PostgreSQL, please see the Installing and Configuring PostgreSQL chapter of this manual.

Installing and Configuring SQLite

For the details of installing and configuring SQLite, please see the Installing and Configuring SQLite chapter of this manual.

Internal Bacula Catalog

Please see the Internal Bacula Database chapter of this manual for more details.

Database Table Design

All discussions that follow pertain to the MySQL database. The details for the PostgreSQL and SQLite databases are essentially

identical except for that all fields in the SQLite database are stored as ASCII text and some of the database creation statements are a bit different. The details of the internal Bacula catalog are not discussed here.

Because the Catalog database may contain very large amounts of data for large sites, we have made a modest attempt to normalize the data tables to reduce redundant information. While reducing the size of the database significantly, it does, unfortunately, add some complications to the structures.

In simple terms, the Catalog database must contain a record of all Jobs run by Bacula, and for each Job, it must maintain a list of all files saved, with their File Attributes (permissions, create date, ...), and the location and Media on which the file is stored. This is seemingly a simple task, but it represents a huge amount interlinked data. Note: the list of files and their attributes is not maintained when using the internal Bacula database. The data stored in the File records, which allows the user or administrator to obtain a list of all files backed up during a job, is by far the largest volume of information put into the Catalog database.

Although the Catalog database has been designed to handle backup data for multiple clients, some users may want to maintain multiple databases, one for each machine to be backed up. This reduces the risk of confusion of accidental restoring a file to the wrong machine as well as reducing the amount of data in a single database, thus increasing efficiency and reducing the impact of a lost or damaged database.

Sequence of Creation of Records for a Save Job

Start with StartDate, ClientName, Filename, Path, Attributes, MediaName, MediaCoordinates. (PartNumber, NumParts). In the steps below, “Create new” means to create a new record whether or not it is unique. “Create unique” means each record in the database should be unique. Thus, one must first search to see if the record exists, and only if not should a new one be created, otherwise the existing RecordId should be used.

1. Create new Job record with StartDate; save JobId
2. Create unique Media record; save MediaId

3. Create unique Client record; save ClientId
4. Create unique Filename record; save FilenameId
5. Create unique Path record; save PathId
6. Create unique Attribute record; save AttributeId store ClientId, FilenameId, PathId, and Attributes
7. Create new File record store JobId, AttributeId, MediaCoordinates, etc
8. Repeat steps 4 through 8 for each file
9. Create a JobMedia record; save MediaId
10. Update Job record filling in EndDate and other Job statistics

Database Tables

Filename		
Column Name	Data Type	Remark
FilenameId	integer	Primary Key
Name	Blob	Filename

The Filename table shown above contains the name of each file backed up with the path removed. If different directories or machines contain the same filename, only one copy will be saved in this table.

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

The Path table contains shown above the path or directory names of all directories on the system or systems. The filename and any MSDOS disk name are stripped off. As with the filename, only one copy of each directory name is kept regardless of how many machines or drives have the same directory. These path names should be stored in Unix path name format.

Some simple testing on a Linux file system indicates that separating the filename and the path may be more complication than is warranted by the space savings. For example, this system has a total of 89,097 files, 60,467 of which have unique filenames, and there are 4,374 unique paths.

Finding all those files and doing two stats() per file takes an average wall clock time of 1 min 35 seconds on a 400MHz machine running RedHat 6.1 Linux.

Finding all those files and putting them directly into a MySQL database with the path and filename defined as TEXT, which is variable length up to 65,535 characters takes 19 mins 31 seconds and creates a 27.6 MByte database.

Doing the same thing, but inserting them into Blob fields with the filename indexed on the first 30 characters and the path name indexed on the 255 (max) characters takes 5 mins 18 seconds and creates a 5.24 MB database. Rerunning the job (with the database already created) takes about 2 mins 50 seconds.

Running the same as the last one (Path and Filename Blob), but Filename indexed on the first 30 characters and the Path on the first 50 characters (linear search done there after) takes 5 mins on the average and creates a 3.4 MB database. Rerunning with the data already in the DB takes 3 mins 35 seconds.

Finally, saving only the full path name rather than splitting the path and the file, and indexing it on the first 50 characters takes 6 mins 43 seconds and creates a 7.35 MB database.

File		
Column Name	Data Type	Remark
FileId	integer	Primary Key
FileIndex	integer	The sequential file number in the Job
JobId	integer	Link to Job Record
PathId	integer	Link to Path Record
FilenameId	integer	Link to Filename Record
MarkId	integer	Used to mark files during Verify Jobs
LStat	tinyblob	File attributes in base64 encoding
MD5	tinyblob	MD5 signature in base64 encoding

The File table shown above contains one entry for each file backed

up by Bacula. Thus a file that is backed up multiple times (as is normal) will have multiple entries in the File table. This will probably be the table with the most number of records. Consequently, it is essential to keep the size of this record to an absolute minimum. At the same time, this table must contain all the information (or pointers to the information) about the file and where it is backed up. Since a file may be backed up many times without having changed, the path and filename are stored in separate tables.

This table contains by far the largest amount of information in the Catalog database, both from the stand point of number of records, and the stand point of total database size. As a consequence, the user must take care to periodically reduce the number of File records using the retention command in the Console program.

Job		
Column Name	Data Type	Remark
JobId	integer	Primary Key
Job	tinyblob	Unique Job Name
Name	tinyblob	Job Name
PurgedFiles	tinyint	Used by Bacula for purging/retention periods
Type	binary(1)	Job Type: Backup, Copy, Clone, Archive, Migration
Level	binary(1)	Job Level
ClientId	integer	Client index
JobStatus	binary(1)	Job Termination Status
SchedTime	datetime	Time/date when Job scheduled
StartTime	datetime	Time/date when Job started
EndTime	datetime	Time/date when Job ended
JobTDate	bigint	Start day in Unix format but 64 bits; used for Retention period.
VolSessionId	integer	Unique Volume Session ID
VolSessionTime	integer	Unique Volume Session Time
JobFiles	integer	Number of files saved in Job
JobBytes	bigint	Number of bytes saved in Job
JobErrors	integer	Number of errors during Job
JobMissingFiles	integer	Number of files not saved (not yet used)
PoolId	integer	Link to Pool Record
FileSetId	integer	Link to FileSet Record

PurgedFiles	tiny integer	Set when all File records purged
HasBase	tiny integer	Set when Base Job run

The Job table contains one record for each Job run by Bacula. Thus normally, there will be one per day per machine added to the database. Note, the JobId is used to index Job records in the database, and it often is shown to the user in the Console program. However, care must be taken with its use as it is not unique from database to database. For example, the user may have a database for Client data saved on machine Rufus and another database for Client data saved on machine Roxie. In this case, the two database will each have JobIds that match those in another database. For a unique reference to a Job, see Job below.

The Name field of the Job record corresponds to the Name resource record given in the Director's configuration file. Thus it is a generic name, and it will be normal to find many Jobs (or even all Jobs) with the same Name.

The Job field contains a combination of the Name and the schedule time of the Job by the Director. Thus for a given Director, even with multiple Catalog databases, the Job will contain a unique name that represents the Job.

For a given Storage daemon, the VolSessionId and VolSessionTime form a unique identification of the Job. This will be the case even if multiple Directors are using the same Storage daemon.

The Job Type (or simply Type) can have one of the following values:

Value	Meaning
B	Backup Job
V	Verify Job
R	Restore Job
C	Console program (not in database)
D	Admin Job
A	Archive Job (not implemented)

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

Value	Meaning
C	Created but not yet running
R	Running
B	Blocked
T	Terminated normally
E	Terminated in Error
e	Non-fatal error
f	Fatal error
D	Verify Differences
A	Canceled by the user
F	Waiting on the File daemon
S	Waiting on the Storage daemon
m	Waiting for a new Volume to be mounted
M	Waiting for a Mount
s	Waiting for Storage resource
j	Waiting for Job resource
c	Waiting for Client resource
d	Waiting for Maximum jobs
t	Waiting for Start Time
p	Waiting for higher priority job to finish

FileSet		
Column Name	Data Type	Remark
FileSetId	integer	Primary Key
FileSet	tinyblob	FileSet name
MD5	tinyblob	MD5 checksum of FileSet
CreateTime	datetime	Time and date Fileset created

The FileSet table contains one entry for each FileSet that is used. The MD5 signature is kept to ensure that if the user changes anything inside the FileSet, it will be detected and the new FileSet will be used. This is particularly important when doing an incremental update. If the user deletes a file or adds a file, we need to ensure that a Full backup is done prior to the next incremental.

JobMedia		
Column Name	Data Type	Remark
JobMediaId	integer	Primary Key

JobId	integer	Link to Job Record
MediaId	integer	Link to Media Record
FirstIndex	integer	The index (sequence number) of the first file written for this Job to the Media
LastIndex	integer	The index of the last file written for this Job to the Media
StartFile	integer	The physical media (tape) file number of the first block written for this Job
EndFile	integer	The physical media (tape) file number of the last block written for this Job
StartBlock	integer	The number of the first block written for this Job
EndBlock	integer	The number of the last block written for this Job
VolIndex	integer	The Volume use sequence number within the Job

The JobMedia table contains one entry for each volume written for the current Job. If the Job spans 3 tapes, there will be three JobMedia records, each containing the information to find all the files for the given JobId on the tape.

Media		
Column Name	Data Type	Remark
MediaId	integer	Primary Key
VolumeName	tinyblob	Volume name
Slot	integer	Autochanger Slot number or zero
PoolId	integer	Link to Pool Record
MediaType	tinyblob	The MediaType supplied by the user
FirstWritten	datetime	Time/date when first written
LastWritten	datetime	Time/date when last written
LabelDate	datetime	Time/date when tape labeled
VolJobs	integer	Number of jobs written to this media

VolFiles	integer	Number of files written to this media
VolBlocks	integer	Number of blocks written to this media
VolMounts	integer	Number of time media mounted
VolBytes	bigint	Number of bytes saved in Job
VolErrors	integer	Number of errors during Job
VolWrites	integer	Number of writes to media
MaxVolBytes	bigint	Maximum bytes to put on this media
VolCapacityBytes	bigint	Capacity estimate for this volume
VolStatus	enum	Status of media: Full, Archive, Append, Recycle, Read-Only, Disabled, Error, Busy
Recycle	tinyint	Whether or not Bacula can recycle the Volumes: Yes, No
VolRetention	bigint	64 bit seconds until expiration
VolUseDuration	bigint	64 bit seconds volume can be used
MaxVolJobs	integer	maximum jobs to put on Volume
MaxVolFiles	integer	maximum EOF marks to put on Volume

The Volume table (internally referred to as the Media table) contains one entry for each volume, that is each tape, cassette (8mm, DLT, DAT, ...), or file on which information is or was backed up. There is one Volume record created for each of the NumVols specified in the Pool resource record.

Pool		
Column Name	Data Type	Remark
PoolId	integer	Primary Key
Name	Tinyblob	Pool Name
NumVols	Integer	Number of Volumes in the Pool
MaxVols	Integer	Maximum Volumes in the Pool

UseOnce	tinyint	Use volume once
UseCatalog	tinyint	Set to use catalog
AcceptAnyVolume	tinyint	Accept any volume from Pool
VolRetention	bigint	64 bit seconds to retain volume
VolUseDuration	bigint	64 bit seconds volume can be used
MaxVolJobs	integer	max jobs on volume
MaxVolFiles	integer	max EOF marks to put on Volume
MaxVolBytes	bigint	max bytes to write on Volume
AutoPrune	tinyint	yes—no for autopruning
Recycle	tinyint	yes—no for allowing auto recycling of Volume
PoolType	enum	Backup, Copy, Cloned, Archive, Migration
LabelFormat	Tinyblob	Label format

The Pool table contains one entry for each media pool controlled by Bacula in this database. One media record exists for each of the NumVols contained in the Pool. The PoolType is a Bacula defined keyword. The MediaType is defined by the administrator, and corresponds to the MediaType specified in the Director's Storage definition record. The CurrentVol is the sequence number of the Media record for the current volume.

Client		
Column Name	Data Type	Remark
ClientId	integer	Primary Key
Name	TinyBlob	File Services Name
UName	TinyBlob	uname -a from Client (not yet used)
AutoPrune	tinyint	yes—no for autopruning
FileRetention	bigint	64 bit seconds to retain Files
JobRetention	bigint	64 bit seconds to retain Job

The Client table contains one entry for each machine backed up by Bacula in this database. Normally the Name is a fully qualified domain name.

UnsavedFiles		
Column Name	Data Type	Remark
UnsavedId	integer	Primary Key
JobId	integer	JobId corresponding to this record
PathId	integer	Id of path
FilenameId	integer	Id of filename

The UnsavedFiles table contains one entry for each file that was not saved. Note! This record is not yet implemented.

Counter		
Column Name	Data Type	Remark
Counter	tinyblob	Counter name
MinValue	integer	Start/Min value for counter
MaxValue	integer	Max value for counter
CurrentValue	integer	Current counter value
WrapCounter	tinyblob	Name of another counter

The Counter table contains one entry for each permanent counter defined by the user.

Version		
Column Name	Data Type	Remark
VersionId	integer	Primary Key

The Version table defines the Bacula database version number. Bacula checks this number before reading the database to ensure that it is compatible with the Bacula binary file.

BaseFiles		
Column Name	Data Type	Remark
BaseId	integer	Primary Key
BaseJobId	integer	JobId of Base Job
JobId	integer	Reference to Job
FileId	integer	Reference to File
FileIndex	integer	File Index number

The BaseFiles table contains all the File references for a particular JobId that point to a Base file – i.e. they were previously saved and hence were not saved in the current JobId but in BaseJobId under FileId. FileIndex is the index of the file, and is used for optimization of Restore jobs to prevent the need to read the FileId record when creating the in memory tree. This record is not yet implemented.

MySQL Table Definition

The commands used to create the MySQL tables are as follows:

```
USE bacula;
CREATE TABLE Filename (
  FilenameId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
  Name BLOB NOT NULL,
  PRIMARY KEY(FilenameId),
  INDEX (Name(30))
);
CREATE TABLE Path (
  PathId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
  Path BLOB NOT NULL,
  PRIMARY KEY(PathId),
  INDEX (Path(50))
);
CREATE TABLE File (
  FileId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
  FileIndex INTEGER UNSIGNED NOT NULL DEFAULT 0,
  JobId INTEGER UNSIGNED NOT NULL REFERENCES Job,
  PathId INTEGER UNSIGNED NOT NULL REFERENCES Path,
  FilenameId INTEGER UNSIGNED NOT NULL REFERENCES Filename,
  MarkId INTEGER UNSIGNED NOT NULL DEFAULT 0,
  LStat TINYBLOB NOT NULL,
  MD5 TINYBLOB NOT NULL,
  PRIMARY KEY(FileId),
  INDEX (JobId),
  INDEX (PathId),
  INDEX (FilenameId)
);
CREATE TABLE Job (
  JobId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
  Job TINYBLOB NOT NULL,
  Name TINYBLOB NOT NULL,
  Type BINARY(1) NOT NULL,
  Level BINARY(1) NOT NULL,
  ClientId INTEGER NOT NULL REFERENCES Client,
  JobStatus BINARY(1) NOT NULL,
```

```

SchedTime DATETIME NOT NULL,
StartTime DATETIME NOT NULL,
EndTime DATETIME NOT NULL,
JobTDate BIGINT UNSIGNED NOT NULL,
VolSessionId INTEGER UNSIGNED NOT NULL DEFAULT 0,
VolSessionTime INTEGER UNSIGNED NOT NULL DEFAULT 0,
JobFiles INTEGER UNSIGNED NOT NULL DEFAULT 0,
JobBytes BIGINT UNSIGNED NOT NULL,
JobErrors INTEGER UNSIGNED NOT NULL DEFAULT 0,
JobMissingFiles INTEGER UNSIGNED NOT NULL DEFAULT 0,
PoolId INTEGER UNSIGNED NOT NULL REFERENCES Pool,
FileSetId INTEGER UNSIGNED NOT NULL REFERENCES FileSet,
PurgedFiles TINYINT NOT NULL DEFAULT 0,
HasBase TINYINT NOT NULL DEFAULT 0,
PRIMARY KEY(JobId),
INDEX (Name(128))
);
CREATE TABLE FileSet (
FileSetId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
FileSet TINYBLOB NOT NULL,
MD5 TINYBLOB NOT NULL,
CreateTime DATETIME NOT NULL,
PRIMARY KEY(FileSetId)
);
CREATE TABLE JobMedia (
JobMediaId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
JobId INTEGER UNSIGNED NOT NULL REFERENCES Job,
MediaId INTEGER UNSIGNED NOT NULL REFERENCES Media,
FirstIndex INTEGER UNSIGNED NOT NULL DEFAULT 0,
LastIndex INTEGER UNSIGNED NOT NULL DEFAULT 0,
StartFile INTEGER UNSIGNED NOT NULL DEFAULT 0,
EndFile INTEGER UNSIGNED NOT NULL DEFAULT 0,
StartBlock INTEGER UNSIGNED NOT NULL DEFAULT 0,
EndBlock INTEGER UNSIGNED NOT NULL DEFAULT 0,
VolIndex INTEGER UNSIGNED NOT NULL DEFAULT 0,
PRIMARY KEY(JobMediaId),
INDEX (JobId, MediaId)
);
CREATE TABLE Media (
MediaId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
VolumeName TINYBLOB NOT NULL,
Slot INTEGER NOT NULL DEFAULT 0,
PoolId INTEGER UNSIGNED NOT NULL REFERENCES Pool,
MediaType TINYBLOB NOT NULL,
FirstWritten DATETIME NOT NULL,
LastWritten DATETIME NOT NULL,
LabelDate DATETIME NOT NULL,
VolJobs INTEGER UNSIGNED NOT NULL DEFAULT 0,
VolFiles INTEGER UNSIGNED NOT NULL DEFAULT 0,
VolBlocks INTEGER UNSIGNED NOT NULL DEFAULT 0,
VolMounts INTEGER UNSIGNED NOT NULL DEFAULT 0,
VolBytes BIGINT UNSIGNED NOT NULL DEFAULT 0,
VolErrors INTEGER UNSIGNED NOT NULL DEFAULT 0,
VolWrites INTEGER UNSIGNED NOT NULL DEFAULT 0,

```

```

VolCapacityBytes BIGINT UNSIGNED NOT NULL,
VolStatus ENUM('Full', 'Archive', 'Append', 'Recycle', 'Purged',
'Read-Only', 'Disabled', 'Error', 'Busy', 'Used', 'Cleaning') NOT NULL,
Recycle TINYINT NOT NULL DEFAULT 0,
VolRetention BIGINT UNSIGNED NOT NULL DEFAULT 0,
VolUseDuration BIGINT UNSIGNED NOT NULL DEFAULT 0,
MaxVolJobs INTEGER UNSIGNED NOT NULL DEFAULT 0,
MaxVolFiles INTEGER UNSIGNED NOT NULL DEFAULT 0,
MaxVolBytes BIGINT UNSIGNED NOT NULL DEFAULT 0,
InChanger TINYINT NOT NULL DEFAULT 0,
MediaAddressing TINYINT NOT NULL DEFAULT 0,
VolReadTime BIGINT UNSIGNED NOT NULL DEFAULT 0,
VolWriteTime BIGINT UNSIGNED NOT NULL DEFAULT 0,
PRIMARY KEY(MediaId),
INDEX (PoolId)
);
CREATE TABLE Pool (
PoolId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
Name TINYBLOB NOT NULL,
NumVols INTEGER UNSIGNED NOT NULL DEFAULT 0,
MaxVols INTEGER UNSIGNED NOT NULL DEFAULT 0,
UseOnce TINYINT NOT NULL,
UseCatalog TINYINT NOT NULL,
AcceptAnyVolume TINYINT DEFAULT 0,
VolRetention BIGINT UNSIGNED NOT NULL,
VolUseDuration BIGINT UNSIGNED NOT NULL,
MaxVolJobs INTEGER UNSIGNED NOT NULL DEFAULT 0,
MaxVolFiles INTEGER UNSIGNED NOT NULL DEFAULT 0,
MaxVolBytes BIGINT UNSIGNED NOT NULL,
AutoPrune TINYINT DEFAULT 0,
Recycle TINYINT DEFAULT 0,
PoolType ENUM('Backup', 'Copy', 'Cloned', 'Archive', 'Migration', 'Scratch') NOT NULL,
LabelFormat TINYBLOB,
Enabled TINYINT DEFAULT 1,
ScratchPoolId INTEGER UNSIGNED DEFAULT 0 REFERENCES Pool,
RecyclePoolId INTEGER UNSIGNED DEFAULT 0 REFERENCES Pool,
UNIQUE (Name(128)),
PRIMARY KEY (PoolId)
);
CREATE TABLE Client (
ClientId INTEGER UNSIGNED NOT NULL AUTO_INCREMENT,
Name TINYBLOB NOT NULL,
Uname TINYBLOB NOT NULL, /* full uname -a of client */
AutoPrune TINYINT DEFAULT 0,
FileRetention BIGINT UNSIGNED NOT NULL,
JobRetention BIGINT UNSIGNED NOT NULL,
UNIQUE (Name(128)),
PRIMARY KEY(ClientId)
);
CREATE TABLE BaseFiles (
BaseId INTEGER UNSIGNED AUTO_INCREMENT,
BaseJobId INTEGER UNSIGNED NOT NULL REFERENCES Job,
JobId INTEGER UNSIGNED NOT NULL REFERENCES Job,
FileId INTEGER UNSIGNED NOT NULL REFERENCES File,

```

```

        FileIndex INTEGER UNSIGNED,
        PRIMARY KEY(BaseId)
    );
CREATE TABLE UnsavedFiles (
    UnsavedId INTEGER UNSIGNED AUTO_INCREMENT,
    JobId INTEGER UNSIGNED NOT NULL REFERENCES Job,
    PathId INTEGER UNSIGNED NOT NULL REFERENCES Path,
    FilenameId INTEGER UNSIGNED NOT NULL REFERENCES Filename,
    PRIMARY KEY (UnsavedId)
);
CREATE TABLE Version (
    VersionId INTEGER UNSIGNED NOT NULL
);
-- Initialize Version
INSERT INTO Version (VersionId) VALUES (7);
CREATE TABLE Counters (
    Counter TINYBLOB NOT NULL,
    MinValue INTEGER,
    MaxValue INTEGER,
    CurrentValue INTEGER,
    WrapCounter TINYBLOB NOT NULL,
    PRIMARY KEY (Counter(128))
);

```

Storage Media Output Format

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.

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 be 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 give 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.

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.

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.

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.

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.

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.

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.

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`.

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.

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.

Overall Storage Format

Current Bacula Tape Format

6 June 2001

Version BB02 added 28 September 2002

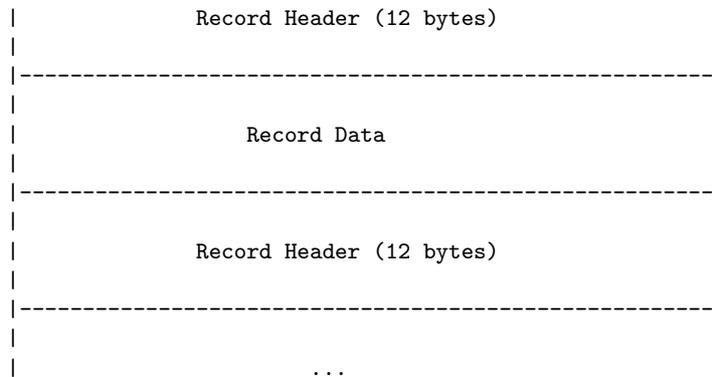
Version BB01 is the old deprecated format.

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)
|
|-----|
|

```



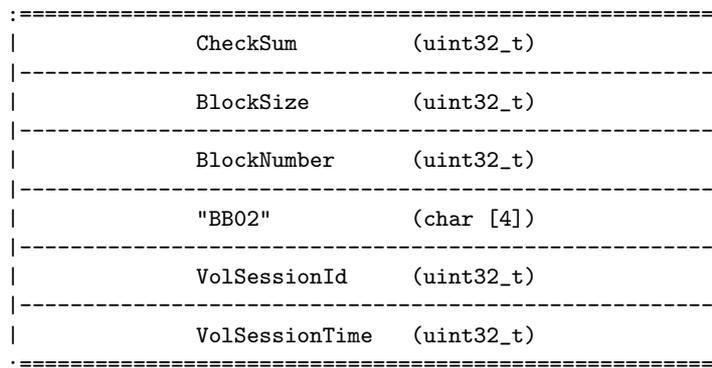
Block Header: the first item in each block. The format is shown below.

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



BB02: Serves to identify the block as a Bacula block and also servers 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.

```

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:

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

Old Depreciated 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:

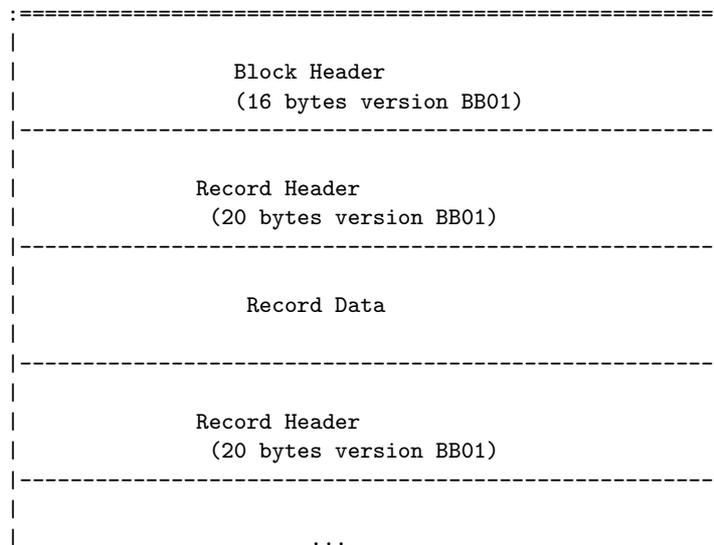
```
uint32_t VolSessionId; /* Unique ID for this session */
uint32_t VolSessionTime; /* Start time/date of session */
int32_t FileIndex;     /* File index supplied by File daemon */
int32_t Stream;        /* Stream number supplied by File daemon */
uint32_t DataSize;     /* size of following data record 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)

```

=====
|          CheckSum          (uint32_t)          |
|-----|
|          BlockSize         (uint32_t)          |
|-----|
|          BlockNumber       (uint32_t)          |
|-----|
|          "BB01"           (char [4])          |
|-----|
=====

```

BB01: Serves to identify the block as a Bacula block and also servers 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)

```

=====
|          VolSessionId      (uint32_t)          |
|-----|
|          VolSessionTime    (uint32_t)          |
|-----|
|          FileIndex         (int32_t)           |
|-----|
|          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.

```

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.

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 (Cygwin), 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.
- If you want to build the Win32 binaries, you will need the full Cygwin 1.5.5 release. Although all components build (console has some warnings), only the File daemon has been tested. Please note that if you attempt to build Bacula on any other version of Cygwin, particularly previous versions, you will be on your own.
- 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.

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 `pthread`s 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 Bacula GUI Interface

General

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

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 GNOME program. The major disadvantage is it increases the size of the code – however, compared to Networker the Bacula Director is really tiny.

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 `<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`

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-password;
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 GNOME GUI interface code in: `<bacula-source>/src/gnome-console/console.c`

0.1 TLS

Written by Landon Fuller

0.1.1 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.

0.1.2 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
```

0.1.3 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.

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.

Manipulating TLS Contexts

```
TLS_CONTEXT *new_tls_context (const char *ca_certfile,  
                             const char *ca_certdir, 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_certdir* 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);
```

Deallocated a previously allocated *TLS_CONTEXT* structure.

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.

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 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.

0.1.4 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.

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.

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()*.

0.1.5 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;
}
```

Bacula Regression Testing

General

This document is intended mostly for developers who wish to ensure that their changes to Bacula don't introduce bugs in the base code.

You can find the existing regression script in the Bacula CVS on the SourceForge CVS in the project tree named regress.

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

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. To date, each subset of tests runs no more than about 15 minutes.

Setting the Configuration Parameters

Once you have the regression directory loaded, you will first need to create a custom xxx.conf file for your system. You can either edit prototype.conf directly or copy it to a new file and edit it. To see a real example of a configuration file, look at kern.conf. The variables you need to modify are:

```
# Where to get the source to be tested
# BACULA_SOURCE="${HOME}/bacula/branch-1.36.2"
BACULA_SOURCE="${HOME}/bacula/k"

# Where to send email   !!!!! Change me !!!!!!!
EMAIL=kern@sibbald.com
SMTP_HOST="localhost"

# Full "default" path where to find sqlite (no quotes!)
#SQLITE_DIR=${HOME}/bacula/depkgs/sqlite3
SQLITE_DIR=${HOME}/bacula/depkgs/sqlite
```

```

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/nst1"

# This must be the path to the autochanger including its name
AUTOCHANGER_PATH="/usr/sbin/mtx"

# Set your database here
#SQLITE_DIR=${HOME}/bacula/depkgs/sqlite3
WHICHDB?="--with-sqlite=${SQLITE_DIR}"
#WHICHDB="--with-mysql=${HOME}/mysql"

# Set this to "--with-tcp-wrappers" or "--without-tcp-wrappers"
TCPWRAPPERS="--with-tcp-wrappers"

```

- **BACULA_SOURCE** should be the full path to the Bacula source code that you wish to test.
- **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.
- **SQLITE_DIR** should be the full path to the sqlite package, must be build before running a Bacula regression, if you are using SQLite. This variable is ignored if you are using MySQL or PostgreSQL. To use PostgreSQL, edit the Makefile and change (or add) **WHICHDB?**="--with-postgresql". For MySQL use "**WHICHDB?**="--with-mysql".
- **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.
- **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.
- **AUTOCHANGER** is the name of your autochanger device. Set this to /dev/null if you do not have one.

- `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.

Building the Test Bacula

Once the above variables are set, you can build Bacula by entering:

```
./configure xxx.conf  
make setup
```

Where `xxx.conf` is the name of the conf file containing your system parameters. This will build a Makefile from `Makefile.in`, 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.

Running the Disk Only Regression

Once Bacula is built, you can run the basic disk only non-root regression test by entering:

```
make test
```

This will run the base set of tests using disk Volumes, currently (19 Dec 2003), there are current 18 separate tests that run. If you are testing on a non-Linux machine two of the tests will not be run. In any case, as we add new tests, the number will vary. It will take about 5 or 10 minutes if you have a fast (2 GHz) machine, 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

```
==== Backup Bacula Test OK ====  
==== Verify Volume Test OK ====  
==== sparse-test OK ====  
==== compressed-test OK =====
```

```

===== sparse-compressed-test OK =====
===== Weird files test OK =====
===== two-jobs-test OK =====
===== two-vol-test OK =====
===== six-vol-test OK =====
===== bscan-test OK =====
===== Weird files2 test OK =====
===== concurrent-jobs-test OK =====
===== four-concurrent-jobs-test OK =====
===== bsr-opt-test OK =====
===== bextract-test OK =====
===== recycle-test OK =====
===== span-vol-test OK =====
===== restore-by-file-test OK =====
===== restore2-by-file-test OK =====
===== four-jobs-test OK =====
===== incremental-test OK =====

```

and the working tape tests are:

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.

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 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.

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 modify the script so that it prints. Do so by finding the file in regress/tests that corresponds to the name printed. For example, the script for the above error message is in: regress/tests/concurrent-jobs-test.

In order to see the output produced by Bacula, you need only change the lines that start with @output to @tee, then rerun the test by hand. it is very important to start the test from the regress directory.

To modify the test mentioned above so that you can see the output, change every occurrence of @output to @tee in the file. In rare cases you might need to remove the 2>&1 >/dev/null from the end of the bacula, bconsole, or diff lines, but this is rare.

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.

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>
```

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.
```

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.

Bacula MD5 Algorithm

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.

Name

`md5` - generate / check MD5 message digest

Synopsis

```
md5 [ -csignature ] [ -u ] [ -dinput_text — infile ] [ outfile ]
```

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.

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.

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.

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.

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.

See Also

`sum(1)`

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.

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.

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*

Bacula Memory Management

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.

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.

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.

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

TCP/IP Network Protocol

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.

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.

`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.

`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.

`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`.

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).

bnet_recv

To read a packet, one uses the subroutine:

int bnet_recv(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_recv()** **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_recv()** is a blocking read.

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

`bnet_strerror`

Returns a formatted string corresponding to the last error that occurred.

`bnet_close`

The connection with the server remains open until closed by the subroutine:

```
void bnet_close(BSOCK *sock)
```

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.

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_rcv()`. Thus the normal exchange between the server (S) and the client (C) are:

S: wait for connection	C: attempt connection
S: accept connection	C: <code>bnet_send()</code> send request
S: <code>bnet_rcv()</code> wait for request	
S: act on request	
S: <code>bnet_send()</code> send ack	C: <code>bnet_rcv()</code> 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.

Smartalloc

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.

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
```

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.

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 “actually*xxx*” is furnished which provides direct access to the underlying system function, as follows:

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("scrdump.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()`.

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.

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.

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.

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.

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.

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.

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.

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

Index

- csignature , 97
- dinput_text , 97
- Download smartall.zip
(Zipped archive) ,
117
- Acknowledgements , 99
- Adding a New Test, 95
- Additional Error information,
105
- all-non-root-tape-tests, 94
- all-root-tests, 94
- all-tape-and-file-tests, 94
- all_non-root-tests, 93
- Alloc() Function , 115
- ALSO
SEE , 98
- API Changes
Bnet, 87
- API Implimentation
TLS, 85
- Archive
Download smartall.zip
Zipped , 117
Download md5.zip Zipped
, 98
- Attributes
Unix File, 70
- Authentication Negotiation,
89
- Backup
Commands Received from
the Director for a , 33
- Bacula
 - Building the Test, 92
 - Bacula Memory Manage-
ment, 100
 - Bacula Porting Notes, 77
 - Bacula Regression Testing, 90
 - Becoming a Server, 107
 - Begin Session Label, 58
 - Block, 56
 - Block Header, 58, 60
 - Blocking State
 - Socket
Manipulating, 88
 - Bnet and Threads, 104
 - Bnet API Changes, 87
 - Bnet_close, 107
 - Bnet_fsend, 105
 - Bnet_open, 104
 - Bnet_recv, 106
 - Bnet_send, 105
 - Bnet_sig, 106
 - Bnet_strerror, 107
 - Bugs , 98
 - Building the Test Bacula, 92
 - Catalog
 - Internal Bacula , 41
 - Catalog Services , 40
 - Classes
 - Message, 22
 - Command and Control Infor-
mation , 28
 - Command Line Message Di-
gest Utility , 96
 - Commands Received from the
Director for a Backup

- , 33
- Commands Received from the Director for a Restore , 33
- Connections
 - Manipulating TLS, 86
- Contexts
 - Manipulating TLS, 85
- Conventions
 - Higher Level, 107
- Copying , 99, 118
- Daemon
 - Director Services , 31
 - File Services , 32
 - Protocol Used Between the Director and the File , 27
 - Protocol Used Between the Director and the Storage , 26
 - Save Protocol Between the File Daemon and the Storage , 28
- Daemon Protocol , 25
- Data Information , 29
- Data Record, 58
- Database Table Design , 41
- Database Tables , 43
- DataSize, 62
- Debug Messages, 22
- Definition
 - MySQL Table , 52
- Definitions, 56
- Description , 96
- Design
 - Database Table , 41
 - Storage Daemon , 34
- Details
 - SMARTALLOC , 113
- Detection
 - Smart Memory Allocation With Orphaned Buffer , 109
- Difficult
 - GUI Interface is , 81
- Directives
 - TLS Configuration, 83
- Director Services Daemon , 31
- Directory Structure, 95
- Disabled
 - When SMARTALLOC is , 115
- Don'ts, 21
- Download md5.zip (Zipped archive) , 98
- Dynamically Allocated Memory, 100
- End Session Label, 59
- Error Messages, 22
- Exit Status , 98
- Extended-Attributes, 70
- Fails
 - If a Test, 94
- File Services Daemon , 32
- File-Attributes, 70
- FileIndex, 57, 61, 70
- Filename, 70
- Filenames and Maximum Filename Length , 41
- Files , 98
- Format
 - Old Depreciated Tape, 72
 - Overall, 59
 - Overall Storage, 65
 - Storage Daemon File Output, 59
 - Storage Media Output, 56
 - Volume Label, 63
- Function
 - alloc , 115
- General, 56, 90, 100, 104
- General , 25, 40, 81
- General Daemon Protocol , 25

- GUI Interface is Difficult , 81
- Hack
 - Invitation to the , 117
- Hand
 - Running the Tests by, 95
- Header
 - Block, 60
 - Record, 61
 - Version 2 Record, 63
 - Version BB02 Block, 63
- Higher Level Conventions, 107
- If a Test Fails, 94
- Implementing a Bacula GUI Interface , 81
- Indenting Standards, 18
- Information
 - Additional Error, 105
 - Command and Control , 28
 - Data , 29
- Initialization and Cleanup
 - Library, 85
- Installing and Configuring MySQL , 41
- Installing and Configuring PostgreSQL , 41
- Installing and Configuring SQLite , 41
- Installing SMARTALLOC , 110
- Interface
 - Implementing a Bacula GUI , 81
- Internal Bacula Catalog , 41
- Introduction
 - SD Design , 34
 - TLS, 83
- Invitation to the Hack , 117
- Job
 - Sequence of Creation of Records for a Save , 42
- Job Messages, 23
- JobId, 57
- Label
 - Session, 64
- Length
 - Filenames and Maximum Filename , 41
- Libraries
 - Living with , 111
- Library Initialization and Cleanup, 85
- Link, 70
- Living with Libraries , 111
- Low Level Network Protocol , 25
- Management
 - Bacula Memory, 100
- Manipulating Socket Blocking State, 88
- Memory
 - Dynamically Allocated, 100
 - Pooled and Non-pooled, 101
 - Statically Allocated, 100
- Memory Messages, 24
- Message Classes, 22
- Messages
 - Debug, 22
 - Error, 22
 - Job, 23, 24
 - Memory, 24
- Minimal Code in Console Program , 81
- MySQL
 - Installing and Configuring , 41
- MySQL Table Definition , 52
- Name, 96
- Negotiating a TLS Connection, 87
- Negotiation

- TLS Authentication, 89
- Notes
 - Bacula Porting, 77
- Old Depreciated Tape Format, 72
- Options , 97
- Other Tests, 93
- Outline
 - SD Development , 34
- Overall Format, 59
- Overall Storage Format, 65
- Overlays and Underhandedness , 116
- Parameters
 - Setting the Configuration, 90
- Pooled and Non-pooled Memory, 101
- Porting
 - Steps to Take for, 78
- Porting Requirements, 77
- PostgreSQL
 - Installing and Configuring , 41
- Program
 - Minimal Code in Console , 81
 - Test and Demonstration , 117
- Protocol
 - Daemon , 25
 - General Daemon , 25
 - Low Level Network , 25
 - TCP/IP Network, 104
- Protocol Used Between the Director and the File Daemon , 27
- Protocol Used Between the Director and the Storage Daemon , 26
- Queued Job Messages, 24
- Record, 56
- Record Header, 58, 61
- Regression
 - Running the Disk Only, 92
- Requests
 - SD Append , 35
 - SD Read , 37
- Requirements
 - Porting, 77
- Restore
 - Commands Received from the Director for a , 33
- Running the Disk Only Regression, 92
- Running the Regression Script, 90
- Running the Tests by Hand, 95
- Save Protocol Between the File Daemon and the Storage Daemon , 28
- Script
 - Running the Regression, 90
- SD Append Requests , 35
- SD Connections and Sessions , 35
- SD Data Structures, 38
- SD Design Introduction , 34
- SD Development Outline , 34
- SD Read Requests , 37
- See Also , 98
- Sequence of Creation of Records for a Save Job , 42
- Serialization, 60
- Server
 - Becoming a, 107
- Services
 - Catalog , 40
- Session, 57
- Session Label, 64

- Sessions
 - SD Connections and , 35
- Setting the Configuration Parameters, 90
- Smart Memory Allocation With Orphaned Buffer Detection , 109
- SMARTALLOC
 - Installing , 110
 - Squelching a , 111
- SMARTALLOC Details , 113
- Socket Blocking State
 - Manipulating, 88
- SPAN class , 35–37
- SQLite
 - Installing and Configuring , 41
- Squelching a SMARTALLOC , 111
- Standards
 - Indenting, 18
- Statically Allocated Memory, 100
- Status
 - Exit , 98
- Steps to Take for Porting, 78
- Storage Daemon Design , 34
- Storage Daemon File Output Format, 59
- Storage Media Output Format, 56
- Stream, 57, 61
- Structure
 - Directory, 95
- Synopsis , 96
- Tabbing, 20
- Tables
 - Database , 43
- TCP/IP Network Protocol, 104
- Test
 - Adding a New, 95
 - Writing a Regression, 94
- Test and Demonstration Program , 117
- Testing
 - Bacula Regression, 90
- Tests
 - Other, 93
- Threads
 - bnet and, 104
- TLS, 83
- TLS API Implimentation, 85
- TLS Configuration Directives, 83
- TLS Connection
 - Negotiating, 87
- TLS Connection Manipulation, 86
- TLS Context Manipulation, 85
- TLS Introduction, 83
- TLS Post-Connection Verification, 86
- Type, 70
- Underhandedness
 - Overlays and , 116
- Unix File Attributes, 70
- Utility
 - Command Line Message Digest , 96
- Verification
 - TLS Post-Connection, 86
- Version 2 Record Header, 63
- Version BB02 Block Header, 63
- VolSessionId, 57, 61
- VolSessionTime, 57, 61
- Volume Label, 58
- Volume Label Format, 63
- When SMARTALLOC is Disabled , 115
- Writing a Regression Test, 94