

Bacula® - The Network Backup Solution

An Enterprise Class OpenSource Backup Tool

It comes by night and sucks the vital essence from your computers.

UniForum Chicago

Hemant Shah

E-mail: shahhe@gmail.com

Linked In Profile: <http://www.linkedin.com/in/shahhemant>

August 26, 2008

1

Presentation Goals

- bacula history
- bacula Features
- OS/Hardware supported by bacula
- bacula architecture
- bacula components
- Sample configuration files
- Future projects
- Backup strategy at LIDP
- Live demo

August 26, 2008

2

What is Bacula?

- It is a network backup solution
- It is opensource, most of code is licensed under GPL version 2
- Hosted at www.sourceforge.net
- Web site is www.bacula.org
- There is online documentation and a wiki at www.bacula.org
- Along with command line, there is a GTK GUI front-end (bat) and web interface (webacula)

August 26, 2008

3

Bacula History

- **Bacula** = Backup + Dracula
- January 2000 - Project started by Kern Sibbald
- 14 April 2002 -First release to Source Forge (version 1.16)
- 29 June 2006 -Release 1.38.11
- January 2007 -Release 2.0.0
- August 2007 - Release 2.2.0
- July 2008 - Release 2.4.2 (Current version)

August 26, 2008

4

Project Goals

- Backup any client from Palm to Mainframe
- Provide “Enterprise” features similar to commercial products
- Assure compatibility for 30 years
- Open source
- Scalable (upto 10,000 machines)

August 26, 2008

5

Features

- A central server and catalog with distributed backup
- All components communicate via the network.
- Internal scheduler for automatic and simultaneous job execution with priorities.
- Backup to single or multiple tape drives
- Backup to autochangers
- Backup to disk
- Backup to removable media (e.g. DVD, USB flash drives, etc.)
- Catalog database, supports multiple databases
- Track Jobs run, Volumes used, File locations, etc.

August 26, 2008

6

Features

- Very fast restores
- Interactive restore with many options
 - From most recent backup
 - Prior backup of time and date
 - List of files/directories
 - Restore by job id
 -
- Labeled volumes prevents accidental overwriting
 - Support for ANSI/IBM labels
- Machine independent data format - extensible
- Support for Unicode on win32 and UTF-8 on Unix
- Bare metal recovery (very complicated)

August 26, 2008

7

Features

- Supports bar code readers
- Backups can span multiple volumes
- Multiple backups on same volume
- You can run scripts
 - Before job
 - After job
 - After failed job

August 26, 2008

8

Security Features

- Daemon authorization with CRAM-MD5
- Director and storage daemon can be run non-root
- MD5, SHA1 signatures for each file
- CRC checksum for each volume block
- Restricted consoles and tray-monitors
- Communication encryption using TLS
- Data encryption using PKI
- Intrusion detection (Verify)

August 26, 2008

9

Supported Operating Systems

- Linux
- AIX
- HP-UX
- FreeBSD, OpenBSD
- Irix
- True64
- Mac OS X
- Windows (Win98/Me, WinNT/2K/XP, Vista) Client
- The bacula server binaries for Windows are available but not supported
- MS Exchange agent available (not developed by Bacula team)

August 26, 2008

10

Supported tape drives and autochangers

- Any tape drive that is supported by your OS will work with bacula
- Any autochanger that is supported by your OS will work with bacula
- List of tape drives and autochangers being used by bacula users is available on the web site

August 26, 2008

11

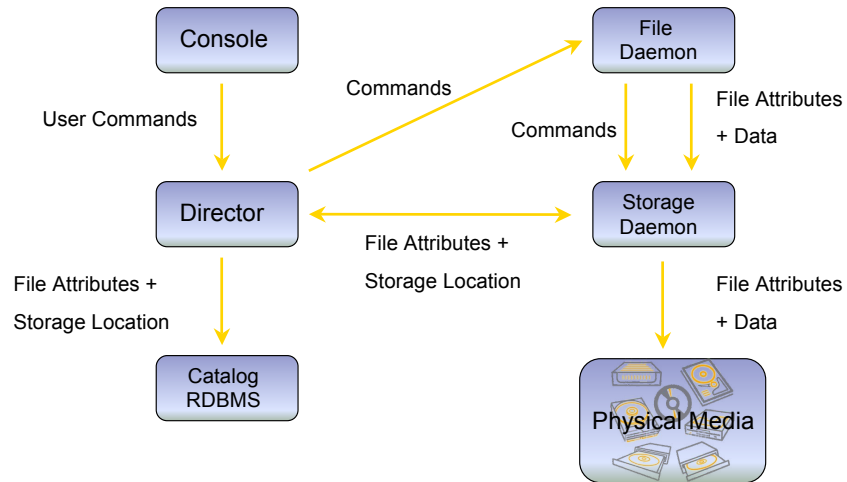
Other requirements

- One of these databases
 - MySQL 4.1 or higher
 - PostgreSQL 7.4 or higher
 - SQLite 2.8.16 or higher
- Gnome 2.x for gnome console
- GTK+ 2.x for GUI front end (bat)
- OS must support pthreads
- dvd+rw-tools with bacula patches applied, if you want to use DVDs as backup medium
- PHP 4 or above if you want to use webacula

August 26, 2008

12

Bacula Architecture



August 26, 2008

13

Bacula Components

- Director
- File Daemon (client)
- Storage Daemon
- Console
- Catalog

They could run on one machine or multiple machines

August 26, 2008

14

Director

- Supervise all operations
- Schedule jobs
- Start jobs
- Cancel jobs
- Maintain catalog
- Typically only one director is used
- It runs as a daemon
- Can be run non-root

August 26, 2008

15

File Daemon

- Installed on systems to be backed up
- Runs backup, restore, verify as requested by director
- Responsible for determining which files to backup
- Provides file attributes and data
- Communicates with director and storage daemon
- Needs access all files, so it must run as root or administrator
- One per client
- Runs as daemon on Unix and as service on Windows

August 26, 2008

16

Storage Daemon

- Read/writes data to physical media
- Accepts file attributes and data from file daemon and writes to physical media
- Reads file attributes and data from physical media and returns it to file daemon
- Sends file attributes and location to the director which sends it to catalog
- Runs as daemon
- Can be run non-root
- Can have more than one storage daemons
- Can have multiple devices per storage daemon
- Typically one storage daemon per director but with multiple devices

August 26, 2008

17

Console

- Allows administrator and users communicate with director
- Can be used for
 - Starting jobs
 - Canceling jobs
 - Review job output
 - Query/modify catalog
- Can accept input from stdin
- Restricted console can allow users to restore their own files
- Available as
 - Text-based (bconsole)
 - bat (GUI using QT4) for Unix, Linux
 - bwx-console (GUI using wxWidgets) for Linux, Unix and Win32
 - Web (bweb, webacula)

August 26, 2008

18

Catalog

- Only component not written by Bacula team
- Responsible for maintaining file indexes and volume databases
- Used to quickly locate and restore file(s)
- Can be used to generate reports of when and where files were backed up
- Supports three databases MySQL, PostgreSQL, and SQLite
- Old data can be automatically pruned by Director

August 26, 2008

19

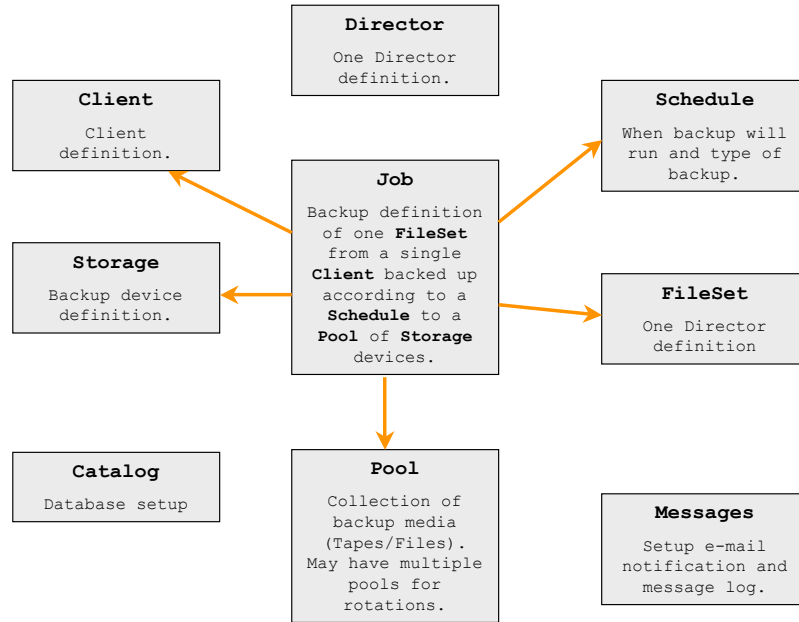
Bacula Configuration Files

- bacula-dir.conf - Director configuration
- bacula-sd.conf - Storage daemon configuration
- bacula-fd.conf - File daemon configuration file
- bconsole.conf - Console configuration file
- Other configuration files
 - bat.conf - bat configuration file
 - tray-monitor.conf - Tray monitor configuration file

August 26, 2008

20

Director config - bacula-dir.conf



August 26, 2008

21

Bacula Jobs

- Basic unifying structure
 - Name - unique name
 - Type - Backup, Restore, Admin,
 - Level - Full, Differential, Incremental
 - FileSet - Which files to backup/restore
 - Client - Which client to backup
 - Storage - Where to store the files
 - Pool - Which volume(s) (tapes, disk) to use
 - Schedule - When to do it

August 26, 2008

22

Sample bacula-dir.conf file

```
Director {
  Name = backupserv-dir
  DIRport = 9101
  QueryFile = "/usr/local/etc/query.sql"
  WorkingDirectory = "/backups/working"
  PidDirectory = "/var/run"
  Maximum Concurrent Jobs = 1
  Password = "SecretPassword"
}
```

August 26, 2008

23

Sample bacula-dir.conf file

```
Catalog {
  Name = MyCatalog
  QueryFile = "/usr/local/etc/query.sql"
  dbname = bacula
  user = bacula
  password = "SecretPassword"
}
```

August 26, 2008

24

Sample bacula-dir.conf file

```
Messages {
  Name = Standard
  mail = root@localhost = all, !skipped, !restore
  operator = root@localhost = mount
  console = all, !skipped
}
```

August 26, 2008

25

Sample bacula-dir.conf file

```
Client {
  Name = client1-fd
  Address = backupserv-dir
  FDPort = 9102
  Catalog = MyCatalog
  Password = "SecretPassword"
  File Retention = 60 days
  Job Retention = 1y
  AutoPrune = yes
}
```

August 26, 2008

26

Sample bacula-dir.conf file

```
Storage {  
  Name = DLTDrive  
  Address = backupserv-dir  
  SDPort = 9103  
  Password = "SecretPassword"  
  Device = "HP-920"  
  Media Type "LTO3"  
}
```

August 26, 2008

27

Sample bacula-dir.conf file

```
Pool {  
  Name = TapePool  
  Pool Type = Backup  
  Recycle = yes  
  AutoPrune = yes  
  Storage = DLTDrive  
}
```

August 26, 2008

28

Sample bacula-dir.conf file

```
Schedule {  
  Name = BackupCycle  
  Run = level=Full 1st fri at 18:00  
  Run = level=Differential 2nd-5th fri at 18:00  
  Run = level=Incremental sat-thu at 22:00  
}
```

August 26, 2008

29

Sample bacula-dir.conf file

```
FileSet {  
  Name = "Client 1 Full Set"  
  Include  
  {  
    Options  
    {  
      signature = MD5  
    }  
    File = /  
  }  
  Exclude  
  {  
    File = /tmp  
    File = /proc  
  }  
}
```

August 26, 2008

30

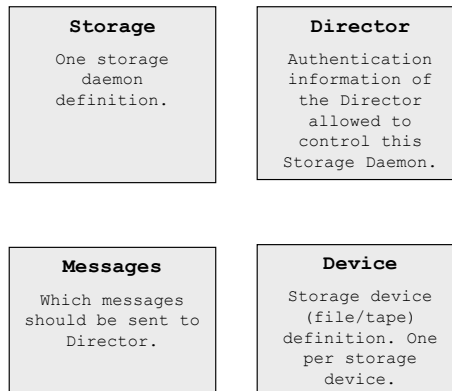
Sample bacula-dir.conf file

```
Job {
  Name = Client1-Backup
  Enabled = yes
  Type = Backup
  Level = Incremental # default level
  Client = Client1-fd
  File Set = "Client 1 Full Set"
  Schedule = BackupCycle
  Pool = TapePool
  Messages = Standard
}
```

August 26, 2008

31

Storage Daemon config - bacula-sd.conf



August 26, 2008

32

Sample bacula-sd.conf file

```
Storage {  
  Name = backupserv-dir  
  SDPort = 9103  
  WorkingDirectory = "/backups/working"  
  PID Directory = "/var/run"  
}
```

August 26, 2008

33

Sample bacula-sd.conf file

```
Director {  
  Name = backupserv-dir  
  Password = "SecretPassword"  
}
```

August 26, 2008

34

Sample bacula-sd.conf file

```
Messages {  
    Name = Standard  
    director = backupserv-dir = all  
}
```

August 26, 2008

35

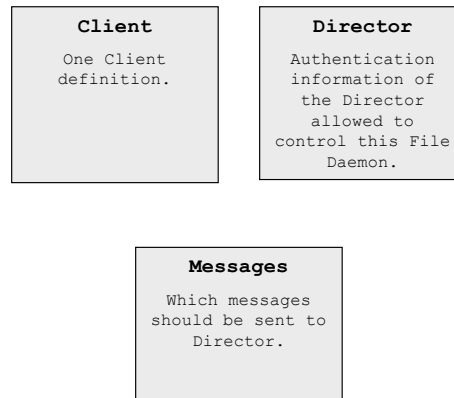
Sample bacula-sd.conf file

```
Storage {  
    Name = DLTDrive  
    Device = "HP-920"  
    Media Type "LTO3"  
    Drive Index = 0  
    Archive Device = /dev/nst0  
    Automatic Mount = yes  
    Always Open = no  
    Removable Media = yes  
    Random Access = no  
    Auto Changer = no  
}
```

August 26, 2008

36

File Daemon config - bacula-fd.conf



August 26, 2008

37

Sample bacula-fd.conf file

```
FileDaemon {
  Name = client1-fd
  FDport = 9102
  WorkingDirectory = "/usr/local/etc"
  PidDirectory "/var/run"
}
```

August 26, 2008

38

Sample bacula-fd.conf file

```
Director {  
  Name = backupserv-dir  
  Password = "SecretPassword"  
}
```

August 26, 2008

39

Sample bacula-fd.conf file

```
Messages {  
  Name = Standard  
  director = backupserv-dir = all, !skipped  
}
```

August 26, 2008

40

Console config - bconsole.conf

Director

Definition of the
director that
this Console can
connect to.

August 26, 2008

41

Sample bconsole.conf file

```
Director {  
  Name = "backupserv-dir"  
  DIRport = 9101  
  Address = backupserv  
  Password = "SecretPassword"  
}
```

August 26, 2008

42

Including Files

- You can include other configuration files
 - @/full/path/to/file
- You can use output of a command
 - Execute command on director
 - “| command”
 - Execute command on client
 - “\| command”

August 26, 2008

43

List of current resource types

Resources	Director	Client	Storage	Console
Autochanger	No	No	Yes	No
Catalog	Yes	No	No	No
Client	Yes	Yes	No	No
Console	Yes	No	No	Yes
Device	No	No	Yes	No
Director	Yes	Yes	Yes	Yes
FileSet	Yes	No	No	No
Job	Yes	No	No	No
JobDefs	Yes	No	No	No
Messages	Yes	Yes	Yes	No
Pool	Yes	No	No	No
Schedule	Yes	No	No	No
Storage	Yes	No	Yes	No

August 26, 2008

44

Future Projects

- Accurate restoration of renamed/deleted files
- Merge multiple backups (Synthetic backup or consolidation)
- Add plugins to the FileSet include statements
- Database driver (LIBDBI) to support more SQL engines
- Better job scheduling conflict resolution
- Certificate based authentication

August 26, 2008

45

Volume Utilities Tools

- bls - "ls" type listing of bacula tape or file
- bextract - extract file(s) from a bacula tape or file (restore does not work)
- bscan - used to re-create bacula catalog database from bacula tape or file
- bcopy - copy bacula one archive to to another (e.g. from file to tape). Information about new volume is not stored in database
- btape - test tape drives
- bsmtplib - SMTP client, runs on Unix and Windows
- dbcheck - check/repair bacula catalog database
- bregex - allows you to test regular expressions
- bwild - allows you to test wildcard expressions
- testfind - file listing using same search engine that is used in include statement

August 26, 2008

46

Resources

- Presentation by Kern Sibbald at BSDCan, May 2008
- Bacula users guide
 - <http://www.bacula.org>
- Bacula wiki
 - <http://wiki.bacula.org>
- Bacula users mailing list
 - bacula-users@lists.sourceforge.net
- Bacula developers mailing list
 - bacula-devel@lists.sourceforge.net
- Bug reports
 - <http://bugs.bacula.org>
- SVN at source forge

August 26, 2008

47

Backup Strategy at LIDP

- Hardware/Software
 - Generic Intel PC with Dual Core Xeon at 2.00 Ghz
 - 4GB memory
 - 3ware 9650SE-8LPML RAID controller
 - 7 x 500GB drives in RAID 6
 - 1 500GB hot spare drive
 - 2 Gigabit interfaces bonded together
 - HP-920 LTO-3 Tape Drive
 - 2TB file system for backup files
 - 100GB file system for bacula catalog
 - Server running 64-bit Fedora 9
 - Server running bacula 2.4.2
 - Clients running bacula 2.2.8
 - PostgreSQL 8.3.3-2 for catalog database

August 26, 2008

48

Backup Strategy at LIDP

- 30 Clients
 - AIX 5.3
 - Red Hat Linux (RHEL 3, RHL 7.3, 8, Fedora 7, 8, 9)
 - Gentoo Linux
 - HP-UX 11
 - Windows 2000 server

August 26, 2008

49

Backup Strategy at LIDP

- Full backup size - 565 GB (uncompressed)
- Full backups on 1st Friday of the month
 - First run full backup to tape (20 hours)
 - Then run full backup to disk (13 hours)
- Incremental backups to disk on other days
- Slow I/O on several old machines
- Backup of millions of small files (source code)
- Retain tapes for 2 years
- Retain disk backups for 2 months

August 26, 2008

50

Resources defined at LIDP

```
Schedule
{
  Name = FullBackupToTape
  Run = Level=Full Storage=DLT_Drive SpoolData=yes 1st fri at 20:00
}

Schedule
{
  Name = BackupToDisk
  Run = Level=Full          1st fri at 21:00
  Run = Level=Incremental  sun-thu at 22:00
  Run = Level=Incremental  2nd-5th fri at 22:00
  Run = Level=Incremental  2nd-5th sat at 22:00
}
```

August 26, 2008

51

Resources defined at LIDP

```
#####
# Full backup to tape pool definition
#####
Pool
{
  Name = FullBackupTapePool
  Pool Type = Backup
  Recycle = no
  AutoPrune = yes
  Volume Retention = 2 years
  Maximum Volumes = 24
  Storage = DLT_Drive
  Label Format = "FullBackup-${Year}-${Month:p/2/0/r}-${Day:p/2/0/r}"
}
```

August 26, 2008

52

Resources defined at LIDP

```
Pool
{
  Name = lidp4-FullBackupDiskPool
  Pool Type = Backup
  AutoPrune = yes
  Maximum Volume Jobs = 1
  Use Volume Once = yes
  Recycle = no
  Volume Retention = 2 months
  Storage = lidp4-File
  Label Format = "Full-lidp4-${Year}-${Month:p/2/0/r}-${Day:p/2/0/r}-
    ${Hour:p/2/0/r}:${Minute:p/2/0/r}:${Second:p/2/0/r}"
}
```

August 26, 2008

53

Resources defined at LIDP

```
Pool
{
  Name = lidp4-IncrBackupDiskPool
  Pool Type = Backup
  AutoPrune = yes
  Maximum Volume Jobs = 1
  Use Volume Once = yes
  Recycle = no
  Volume Retention = 2 months
  Storage = lidp4-File
  Label Format = "Incr-lidp4-${Year}-${Month:p/2/0/r}-${Day:p/2/0/r}-
    ${Hour:p/2/0/r}:${Minute:p/2/0/r}:${Second:p/2/0/r}"
}
```

August 26, 2008

54

Resources defined at LIDP

```
Storage
{
  Name = lidp4-File
  Address = lidp11
  SDPort = 9103
  Password = "xxxxxxxxxxxxx"
  Device = lidp4-File
  Media Type = File
  Maximum Concurrent Jobs = 5
}
```

August 26, 2008

55

Resources defined at LIDP

```
Job
{
  Name = lidp4-FullBackupToTape
  Enabled = yes
  Type = Backup
  Client = lidp4-fd
  FileSet = "lidp4 Tape set"
  Schedule = FullBackupToTape
  Messages = Standard
  Priority = 5
  Pool = FullBackupTapePool
  Write Bootstrap = "/backups/bootstrap/lidp4-BackupToTape.bsr"
}
```

August 26, 2008

56

Resources defined at LIDP

```
Job
{
  Name = lidp4-BackupToDisk
  Enabled = yes
  Type = Backup
  Client = lidp4-fd
  FileSet = "lidp4 Disk set"
  Schedule = BackupToDisk
  Messages = Standard
  Priority = 10
  Pool = Default
  Full Backup Pool = lidp4-FullBackupDiskPool
  Incremental Backup Pool = lidp4-IncrBackupDiskPool
  Write Bootstrap = "/backups/bootstrap/lidp4-BackupToDisk.bsr"
}
```

August 26, 2008

57

Resources defined at LIDP

```
Fileset
{
  Name = "lidp4 Disk set"
  Ignore Fileset Changes = yes
  Include
  {
    Options
    {
      signature = MD5
      oneofs=yes
    }
    Options
    {
      WildDir = "*/.snapshot"
      WildFile = "*/core"
      WildDir = "*/cob"
      WildDir = "*/dbcob"
      WildDir = "*/tmp"
      exclude=yes
    }
  }
}
```

August 26, 2008

58

Resources defined at LIDP

```
File = "\\usr/local/adm/CreateBaculaIncludeList"
Options
{
    compression = gzip
}
}
Exclude
{
    Options
    {
        signature = MD5
    }
    File = /proc
    File = /tmp
    File = /dev
}
}
```

August 26, 2008

59

Live Demo

August 26, 2008

60