



Backup Concept and Hardware Sizing

Bacula Enterprise Edition



**Bacula
Systems
White
Paper**

Version 2.0, February 16, 2017
Copyright (C) 2008-2017, Bacula Systems S.A.
All rights reserved.

www.baculasystems.com

Contents

1	Introduction	2
2	General hardware recommendations for Bacula Enterprise Edition	3
2.1	Bacula Director (DIR)	3
2.2	Bacula Catalog	3
2.3	Bacula Storage Daemon (SD)	4
2.4	Bacula File Daemons (FD)	4
2.5	Hardware requirements for enhanced security	4
3	Example configuration for a Bronze subscription	5



Introduction

Bacula Enterprise Edition is the leading open core enterprise backup platform. From its inception it has been designed to scale from a small number of servers to many thousands without requiring major investments in hardware or software. Due to its modern, single platform, scalable, modular and open design, Bacula Enterprise Edition has a very light footprint and places relatively few demands on hardware. Its performance is easily tuned following recommended guidelines and best practices available from Bacula Systems Support or through our Consulting Team. In this document, the hardware requirements for each of the components of Bacula Enterprise Edition are described individually to allow customers to assess the hardware requirements specific to their environments.



General hardware recommendations for Bacula Enterprise Edition

2.1 Bacula Director (DIR)

The Bacula Director itself does not require many resources. It is a very lightweight process that only does some scheduling. A default server with minimal hardware requirements will suffice.

2.2 Bacula Catalog

If you plan to run the Catalog on the same machine where the DIR resides (which is done quite often), the server machine needs to be what you would buy for a good database server: a lot of RAM and fast disk I/O, i.e. SSD for the database (RAID-1).

You should run PostgreSQL. This database scales and performs best from our experience. When you have high availability and performance demands, you should consider running a PostgreSQL cluster of several machines.

The bacula-dir process does not need many resources (see section 2.1), so it can easily run on the same machine where you have the Catalog database. A well equipped database server will be capable of running the bacula-dir service too.

If you know the size that the Bacula Catalog Database will have in the end, you can design your SSD RAID-1 accordingly. If you don't know for sure, you should make use of LVM, because you will be able to resize (i.e. increase) the database partition later. While it is true that the additional LVM layer will cost you a little bit of performance, a modern Linux distribution on a fast processor will be able to handle that, and the benefits of LVM outweigh the (small) reduction in performance.

The biggest table in the Bacula Catalog database is the File table. Each file backed up will produce an entry of approximately 250 Bytes. If you know the number of files in all of your file sets and the rate/number with which files change, you can multiply that and you will end up with an estimate for the overall size of the Bacula Catalog. Multiply the number you get with a factor of 1.5 to account for the rest of the Catalog tables/entries and for future data growth. Your retention times are also

important here, because Bacula can (and should) be configured to prune Catalog entries once the retention time for files has passed. This will keep your Catalog lean.

2.3 Bacula Storage Daemon (SD)

A Storage Daemon machine mostly needs CPU power and a good connection to the storage backend (FC or iSCSI) and to the network that connects to the clients (10 GE). High bandwidth is of paramount importance for all networks. For the file systems on the storage itself we recommend XFS file system on top of LVM and of course a sufficient RAID level (RAID-6). Memory is not that important unless Global Endpoint Deduplication (GED) is being implemented. If you plan on using our GED solution, memory requirements for an SD will increase significantly, and we have a simple calculation to help you size the memory in your SDs appropriately.

2.4 Bacula File Daemons (FD)

The agent that gets installed on the backup client machines is a very lightweight process with a small CPU and memory footprint. On Windows systems it runs as a Windows Service in the background.

2.5 Hardware requirements for enhanced security

If you plan to do TLS on any of the Bacula components you will need additional CPU power to encrypt the channel.

Data encryption will put additional CPU load on the FD only. Multiple CPU sockets and a lot of cores help here. If you plan to use Bacula Accurate Backup, the client should have as much RAM as possible because the comparison for each file to be backed up will be memory intensive.



Example configuration for a Bronze subscription

We will present example configurations for one DIR, two SDs, 200 clients (FDs) and 20TB of data. In the table below you will find the minimum configuration and the optimal configuration that we recommend for your Bacula Enterprise Edition environment.

The average size of a Full backup for a client in this example is only 100GB, but this does not mean that all clients are the same. There could be big file servers with several TB of data, and other clients that only have a few GB. We have made a few assumptions to arrive at the numbers below:

- The number of all files in all Full backups for all clients is 100 million.
- The backup schedule is one Full each month on a weekend and one Incremental each workday (Mon-Fri).
- The size of all Incremental backups for a given client accumulated for a month equals the size of its Full.
- Restore should be possible with daily granularity for the last two months.



(Illustrative material only)

Minimal configuration	
DIR	SD
CPU	2 sockets, 8 cores
RAM	8 GB RAM
Storage	100 TB (RAID 6)
Management Network	1 GE
Backup network	10 GE
Storage network	2-4x 1 GE or 10 GE
Optimal configuration	
DIR	SD
CPU	2 sockets, 8 cores
RAM	16 GB RAM
Expansion ports	minimum one free and usable PCIe x8 slot
Storage	200 GB for system & logs (RAID1) 100 TB for Bacula Volumes (RAID 6, LVM) allow for SAN-attached volume store
Management Network	1 GE
Backup network	10 GE
Storage network	10 GE (NFS, iSCSI) 8 Gb FC (FC SAN) No. of ports depends on desired throughput and zoning

For More Information

For more information on Bacula Enterprise Edition, or any part of the broad Bacula Systems services portfolio, visit www.baculasystems.com.

Headquarters

Bacula Systems SA
Rue Galilée 5
CH-1400 Yverdon-les-Bains
Switzerland
Phone: +41 21 641 6080
Fax: +41 21 641 6081

USA

Western Region

100 100th Ave SE. #3
Bellevue, WA 98004

Eastern Region & Canada

269 Carmita Ave
Rutherford, NJ 07070

Toll Free: +1 800 256 0192

Rev V. 2.0
Author(s): DOG