

Bacula Enterprise Edition For The Cloud

With Native Cloud Integration





Overview

Bacula Enterprise Edition not only connects to Public and Private Clouds (via S3, Azure, Oracle and Google interfaces), but also provides a vast array of tools to help the user be ahead regarding keeping Cloud-related costs right down to the bare minimum. Bacula Enterprise is designed from the ground up to provide features which specifically provide value and efficiency – and the efficiencies in cost management are one of the highest available in the industry. With Bacula, the user has an especially scalable and powerful backup and data recovery solution that can do cloud, virtual and physical all in one, single platform.

Unparalleled Flexibility Puts the User Back in Control of Costs

Bacula Enterprise Edition is known to combine the strength of an open source core with advanced enterprise-class functionality, together with world class support and training for its customers. The scalability of this product — enabled by its modular architecture and the great number of configuration options — allows big data centers to implement backup on the Petabyte scale and beyond, fully tailored to their specific needs.

The variety of options that Bacula Enterprise offers is now extended by its native cloud integration with S3, Google Cloud, Oracle Cloud or Azure as the target for its Storage Driver. There are little or no limitations on the platform's functionality: these new storage targets can be used with all of Bacula Enterprise's existing features right away. This means that features such as Bare Metal Recovery, Snapshots, Single File Recovery, Single Mailbox Recovery (Exchange), all the Scheduling features, advanced encryption, compression, comprehensive compatibility (such as backing up Microsoft databases and files) and full range of backup levels are all at the user's fingertips. The user is able to take backup and data recovery to a whole new level, with everything being orchestrated conveniently from one single interface.

Bacula Enterprise Edition's Cloud capabilities are simply just another part of its modular architecture, and the user will be able to back up virtual machines, databases and other applications — or just do classic file-based backups — and write all of it directly to a private or public cloud without the need to learn anything new. Of course, this results not only in ease of use but also convenience, with the ability to mix-and-match Bacula Enterprise Edition's GUI and command line interface controls, as the user pleases.

Use Case

But how does all this fit into the real world? Consider the below examples of two backup managers at two different companies who are tasked by both their CIO and CEO to implement cloud backup because it is "cheaper than sending our backups offsite".

The first does the obvious and starts using the cloud for the company's routine nightly backups. These backups aren't just for Disaster Recovery; they also serve the company's employees recovery requests for mistakenly deleted files and lost



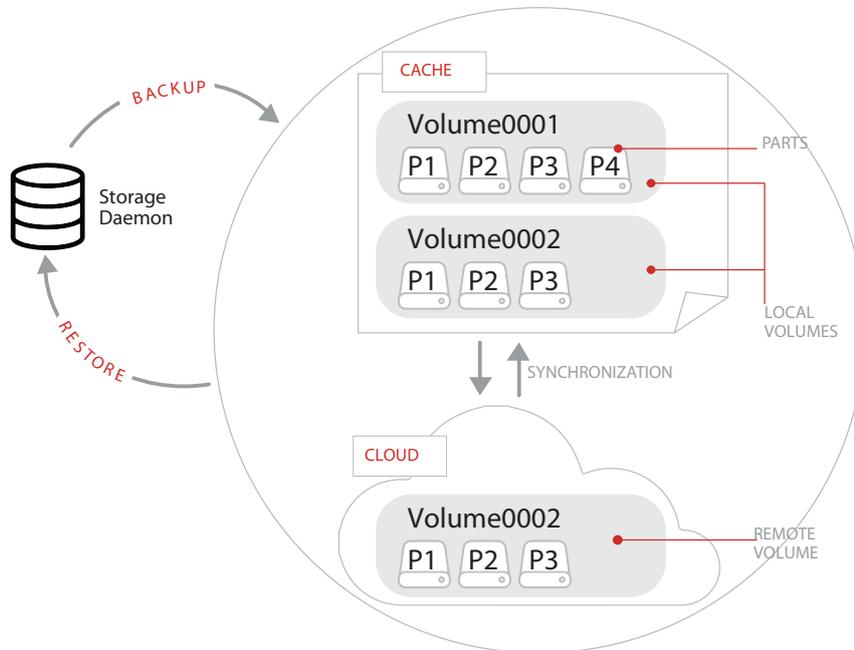
data. Since the backup manager doesn't see the monthly cloud bill, he doesn't see the costs for these recoveries start to rise (previously they were simply part of the sunk cost of a local backup infrastructure). Eventually, after 6 months, the CIO comes to him and asks why backup costs have been steadily going up. After all, the cloud was supposed to be cheaper! After hearing that restores from the cloud are much more expensive than backups, the boss makes a sweeping decision to go back to just using local backups. Of course by now, all of the company's backups for the past six months are in the cloud with no local copies, so doing this triggers a massive data transfer bill just to bring the data back on-site. The whole exercise has proven to be a very expensive mistake and the unfortunate employee gets blamed for the failed strategy.

The second backup manager does his homework first, and realizes that the economics of the cloud are such that it really is better than traditional offsite backups for true disaster recovery — but not for day-to-day recoveries. He understands that the cloud eliminates the need for an expensive second data center or offsite archive company, but that local backups are still essential for smooth operation of the business and predictable costs. This company implements Bacula Enterprise as a single Backup and Recovery platform for physical, virtual and cloud backup - and enjoys the benefits of a flat monthly bill with practically no surprise costs, because their daily recoveries come from Bacula's local cache. However, a year later, the boss reads about a new cloud storage provider with even lower costs and orders that the company switch to their service. Because Bacula still has a local copy of all the backups, there is no data to be transferred expensively out of the old provider and into the new one. The old provider can be immediately cancelled and Bacula simply pointed at the new one, to begin uploading its second copy to the cloud in the background.

The Principles of Bacula Enterprise Edition's Cloud Architecture

Besides file and tape storage, Bacula now has a cloud storage backend. This storage backend for the Storage Daemon enables a Bacula user to store their backup data in cloud-backed object storage while still using all the features and functionality of Bacula.

Enterprise Edition - Native Cloud Integration



Native Cloud Integration With Bacula Enterprise Edition's Storage Daemon

Managing Objects Before Writing to Cloud

When sending data to the Cloud, Bacula Enterprise Edition's Storage Daemon (SD) behaves as it would for virtual or physical backup and restore tasks from both the operator's and File Daemon's perspective. Instead of writing continuous data streams to volumes, or reading volume data as needed, Bacula Enterprise Edition's cloud backend creates Volumes split into "parts" that are cached locally and synchronized to a cloud storage service in the background.

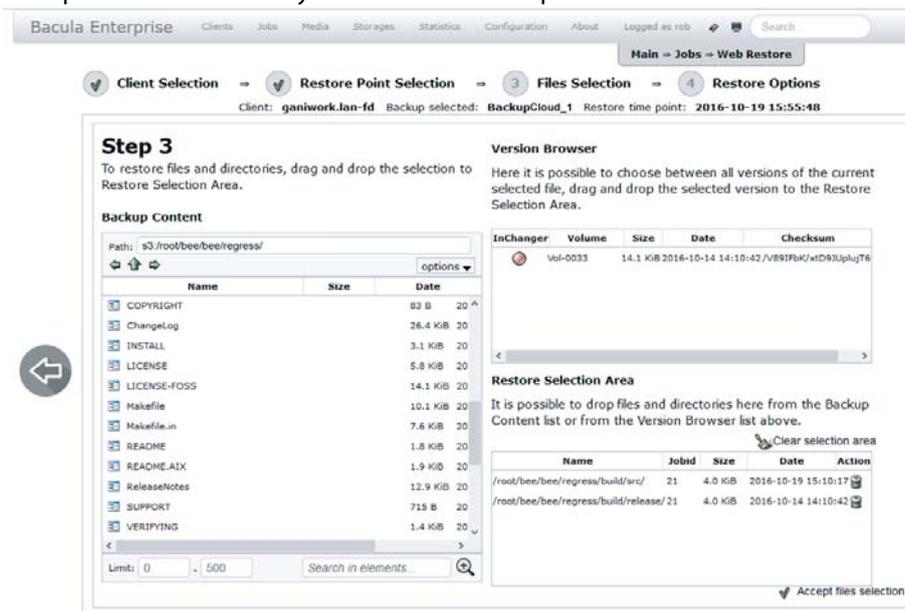
Unique "parts" approach (Bacula Cloud Volumes)

This special feature enables efficient upload/download differentiation of both Bacula disk volumes and new "Cloud Volumes". This directory contains all the file parts, and means that smaller 'parts' of data can be managed. The result is a finer-grain control of what data can be backed up and retrieved from the cloud, bringing the opportunity for yet further time and cost savings. Concurrent asynchronous upload and download of parts Bacula allows the user to exploit multiple connections to the cloud, uploading and/or downloading simultaneously to add extra flexibility and control to the user.

Data Restore from the Cloud

Cached parts can be kept locally as long as required, or removed to free up local cache capacity as needed. During a restore, cached parts will be used, and — if needed — only the parts that are required will be retrieved from the cloud object storage. In most cases, this means a substantial reduction in cloud costs. This approach ensures that all kinds of data can be stored with the cloud backend — in particular, client encrypted or compressed backups and plugin-generated data are fully and transparently supported to avoid introducing any restrictions to existing backup scenarios.

Bacula System's BWeb is an extremely powerful GUI web application for Bacula Enterprise Edition. Below you can see an example of the wizard for data recovery.



Bacula Enterprises' easy to use BWeb GUI restores Cloud data in just 4 easy steps

Practical Considerations

Once configured, the cloud backed SD storage is usable in the same way as any other Bacula storage device. The flexible integration and configuration allows for batched or nearly-synchronous ongoing upload, and only downloads what is really necessary, thus enabling significant cost savings for cloud storage users.

With its configurable local cache, backup speeds can be as fast as backing up to local disk, and synchronization lag to the backend can be controlled. Since the amount of data that is stored to the cloud target is important, data compression is recommended in such configurations.

Security

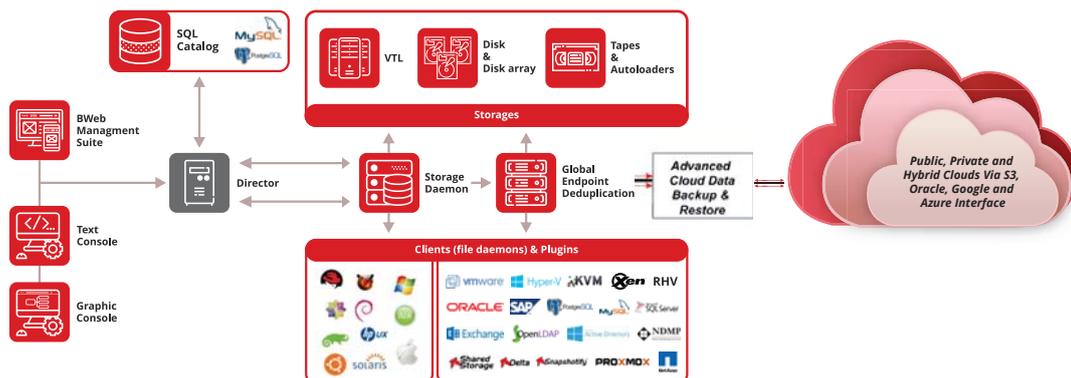
For some users, the cloud introduces security questions. It is important to note therefore, that Bacula Enterprise Edition is packed with state of the art security features, some for direct use with the Cloud, others for use in the physical (and VM) environment:

- ☒ Configurable Data (on Volume) encryption on a Client by Client basis
- ☒ Data encryption cipher (AES 128, AES192, AES256 or blowfish) and the digest algorithm
- ☒ Verification of files previously catalogued, permitting a Tripwire¹ like capability (system break-in detection)
- ☒ CRAM-MD5 password authentication between each component (daemon)
- ☒ Configurable TLS (SSL) communications encryption between each component
- ☒ Computation of MD5 or SHA1 signatures of the file data, if configured

Compatible Clouds

The cloud storage backend for the Storage Daemon uses the S3, S3-IA, Azure, Google Cloud and Oracle Cloud protocols, including https transport encryption for a wide range of public and private cloud services. As a direct result of the architecture design, a driver for Rados will quickly follow. Bacula Systems recommends using its data encryption when you chose a public cloud service provider.

The image below shows Bacula Enterprise Edition's architecture:



¹Tripwire is a trademark of Tripwire, Inc

For More Information

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

