GBS Backup

Backup & Deployment

Factsheet

KVM image backup (SnapLive)

Introduction

The Generic Backup Solution (GBS) offers a Linux based Backup Server that is capable of backing up virtual systems, non-virtual systems, SAN and network equipment.
It supports backup and (disaster) recovery for Linux, Solaris, HP-UX, OpenVMS and Windows Server.

For KVM we have complimented our RSYNC system backup with a KVM image backup mechanism. By taking a QEMU snapshot prior to the KVM image file backup, a consistent backup is guaranteed. The frozen image file is then transferred to the GBS backup server as a full backup. RSYNC will facilitate subsequent incremental system backups.

Backup Cycles

The GBS solution consists of a staged backup schedule:

- Backup preparation (Stage 0)
  Extracts database information in order to prepare databases for the backup.
  this requires local diskspace on the backup client.
- Overnight backup (Stage 1)
  Transfers all the backup data to the local backup server at nighttime.
  Consists of system and data backup mechanism using RSYNC protocol.
- Permanent backup (Stage 2)
  will transfer the overnight backup to a remote location, and will take care of historical backup retention. Maintains backup history: Daily weekly and monthly copies to offer a configurable retention time (3 months by default).
  Can be implemented with any vendor 3rd party backup application.

KVM image backup of virtual machines.

For Virtualized KVM deployments, GBS facilitates the backup of KVM virtual machines.
GBS has introduced a KVM image backup mechanism in stage 1, that will copy the KVM client's qcow2 image file to the backup server for disaster recovery purposes of the VM.

KVM hosts typically run multiple VM's each represented with a VM image file.
These VM image files (usually in /var/lib/libvirt/images) are copied to the GBS backup server as an initial "full" backup.
The KVM image backup is complementary to the regular overnight system backup using RSYNC.
We only need to copy the image file once, subsequent incremental backups are executed using the regular GBS system/data backup mechanisms. When a disaster recovery is needed, the VM image file is updated on the GBS backup server, and subsequently transferred to the KVM host to boot/restart the VM from.

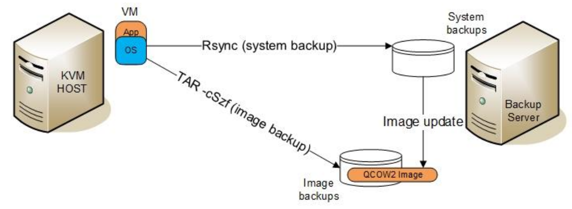

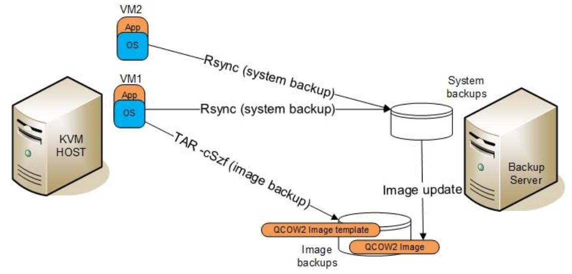

Our strategy is to create the full image backup on an incidental basis, and use the "system backup" for subsequent incremental backups.
That avoids the need to run a full image backup on a regular basis, avoiding massive data transfers associated with such a full image backup. Please note that we only need to collect the image backup in order to capture the layout and structure of the VM's internal disks. The actual file content will be sourced from the RSYNC system backup and copied into the image with the "image update".

SnapLive backup (QEMU 2.1)

In addition to the KVM image backup, GBS facilitates the creation of QEMU external live snapshots prior to the image backup in order to freeze the KVM image file into a consistent state, transforming the image backup into a "KVM SnapLive image backup".
In RHEL8 (qemu2.1 and higher), the VM image file can be copied as a hot backup using QEMU external snapshots, so the VM can keep running during the backup.
QEMU external snapshots are not supported in QEMU1.x, so for older Linux distributions (E.g RHEL7) the image backup needs to be created while the VM is shutdown. This only needs to be done incidentaly, as the image is only needed to capture the system disk layout.
For more detailed information, please read: https://wiki.libvirt.org/page/Live-disk-backup-with-active-blockcommit

Efficient live full disk backup in QEMU 4.2

In modern QEMU releases, QEMU 4.2 and higher, a new modern backup API is integrated and available.
This backup API omits the explicit use of the live external disk snapshot, and natively supports a "live full disk backup" with less hassle.
It also includes incremental backups natively.
QEMU 4.2 was included in RHEL 8.3/ CentOS 8.3 and higher.
GBS tooling will be modified and tested accordingly in the next GBS release to benefit from this QEMU/Libvirt functionality.
For more detailed information, please read: https://libvirt.org/kbase/live_full_disk_backup.html

Image Template backups:

Instead of making an image backup for each VM, it is sufficient to make a single "Template backup" for a single VM. This template backup can subsequently be used by similar VM's of the same "flavour", where the flavour is determined by the same disk size, LVM partition layout and filesystem layout. The image template can be taken from a different KVM host if needed.

KVM Image Backup basic operation

The KVM image backup consists of three parts:

- *gbs-kvmimg-get* operation will search for the given VM instance name in the backup data of the KVM host.
  - it will retrieve all the necessary metadata for the VM from the backup of the KVM host.
  - it will determine which image files make up the VM's image, and will retrieve these KVM image files for backup purpose. Subsequently, a system backup (with RSYNC) will be made on a daily basis.
- *gbs-kvmimg-update* operation will then allow to update the KVM image backup with the latest data taken from the overnight system backup. The updated KVM image can then be used for disaster recovery of the VM.
- KVM live snapshot with block commit.
  Prepares VM images to be frozen (become consistent) for backup purpose, to allow for online backup of KVM images.
  gbs-kvmsnap-create operation triggers the creation of an external QEMU snapshot prior to the KVM image backup.
  gbs-kvmsnap-purge operation deletes any outstanding QEMU snapshots that were created by the gbs-kvmsnap-create.
  The snapshot operations are only supported on qemu-kvm 2.1 and newer (RHEL8).
  For older versions of RedHat/CentOS (7) the VM needs to be shutdown for the KVM image backup.

Backup of the KVM host

In addition to the KVM image backups, the KVM hosts are also included in the backup schedule.
A system backup is created for each KVM host, complemented with a ReaR backup for disaster recovery purpose.
ReaR rescue captures the system's metadata, and is capable of re-creating the system disk from scratch in case of a disaster.

VDO data deduplication and compression

GBS deployment makes use of deduplication mechanisms:

- Rsnapshot for the permanent backup uses "hardlink" deduplication.
  it only copies a file if the content has changed since the previous day.
  if a file did not change, only a hardlink is created to the pervious day/file.
  Typically, the permanent backup contains 20 backups, requiring only 2,5 times the original size.
- VDO is used to reduce backup storage requirements even further.
  The backup repository will contain multiple copies of the same file, as OS system files generally are identical on multiple backup clients. VDO is able to reduce the physical diskspace by a factor of 3-4 times (e.g storing 15TB of data on a 4TB physical disk).

Geographic Redundancy

A geographic redundant backup server can be offered as an additional option in GBS, if the customer has two or more sites Permanent backups are stored on a remote backup server in order to provide geo-redundancy.
The geo-redundant GBS will store the backups locally and on the remote site:

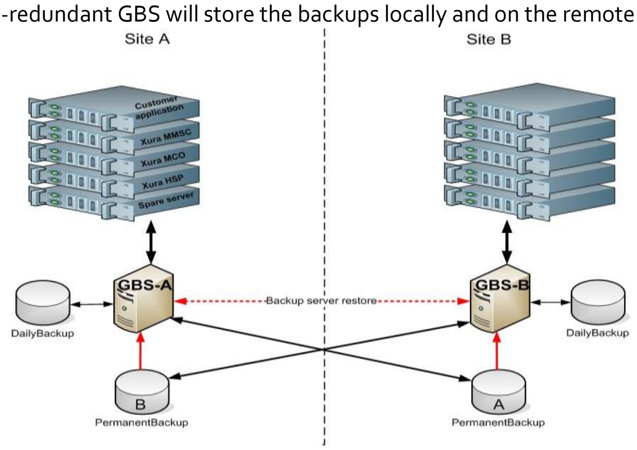

Geo redundancy includes the following features:

- Site A GBS server can be restored from site B GBS server.
- Site B GBS server can be restored from site A GBS server.
- The whole site B can be restored from site A backup.
- The whole site A can be restored from site B backup.
- Disaster recovery fully supported from remote backup.

The remote storage ("Permanent backup") will be mounted on the GBS server via NFSv4 (Network File System).
Precondition for a working geo-redundancy is a good network connection between the sites.

Proven technology, existing mechanism

GBS originates from an existing telecom environment. It has a proven track record of backup and restore capabilities. The existing disaster recovery, image backup and data backup mechanisms are versatile mechanisms that work for the virtual environment just as well as for the original bare metal environment.
GBS was already adapted for VMware/vSphere and now it has been tested successfully for OpenStack and KVM.

Site disaster recovery plan

The GBS backup solution includes a disaster recovery service.
A disaster recovery plan identifies the critical components in the customers telecom/enterprise solution, and defines a step by step approach on how to restore things in case of a disaster.
In addition to restoring individual computer systems, a site disaster recovery planning is also included, to describe the whole recovery procedure in case the whole site needs to be rebuilt from remote backup.

More GBS info

For more info on the GBS backup products and solution, please visit our website: www.gbsbackup.com
Or please feel free to contact our Technical Director:

Roel Otten
info@gbsbackup.com