

Backup Copy

Maria Zubareva
Team Leader EMEA&APAC

Outline

Backup copy concept

Synchronization interval

Limitations of backup copy

Backup copy phases

Compact and Health-check

Retention policy


Data Transport: Direct, Wan accelerators

Logs and troubleshooting

What's new in v9

Concept




3-2-1 rule




Concept

- Not a robo-copy:
 - Backup copy build its own chain (forever forward incremental method)
 - Specific blocks of selected VMs are copied:
Veeam extracts specific blocks of data from source backup job files, and sends them over to the BCJ storage
- Backup copy does not affect virtual infrastructure resources

Backup Copying Process



Inside the backup



OIB = [UUID_of_the_source_host][vm-ref]

Synchronization Intervals

- Time span in which the backup copy job must copy a VM restore point from the source backup repository to the target backup repository
- Minutely and Hourly vs Daily
- Backup Copy Window



Restore point selection

Time of restore point creation \geq current time – synchronization interval

- Example:

Synchronization interval = 24 hours;

Current time = **10/06/2016, 12:00**

Time of restore point creation = **09/06/2016, 13:00**

Veeam Backup & Replication will copy this new restore point, because:

09/06/2016, 13:00 \geq 10/06/2016, 12:00 – 24 hours

Backup Copy Limitations

- 1. Does not copy restore points from **the target backup repository**
- 1. Does not copy restore points from **imported backups**
- 1. Does not copy restore points that **have already been copied** by the same backup copy job to the target backup repository
- 1. Does not copy **corrupted restore points**
- 1. Does not copy restore points that are **locked by some tasks**: for example, a backup job creating a backup chain with the reverse incremental method or a restore process
- 1. Does not copy restore points if the **block size** of the restore point on the source backup repository **differs** from the block size of restore points on the target backup repository. *Restore point is located in backup file with different block size* message

Backup Copy Job Phases

- Idle state
- Synchronization process
- Transform operations
- Post-job activities

Compact Full Backup

- **Compacting a Full Backup File:**

Creates a new empty VBK file and copies to it all data blocks from the full backup file.


As a result, the full backup file gets defragmented, its size reduces and the speed of writing and reading to/from the file increases.

- **Compact full** backup option can be enabled only for the simple retention policy scheme.
- Target backup repository must have enough space to hold a backup file of the full backup size


Health-Check

- CRC check performed for restore points stored on the target repository
- Started at the beginning of the synchronization cycle before data transport
- Verifies only the most recent restore points for the VM(s)
- If a health check fails, Veeam Backup & Replication displays a warning in the job session report
- During the next synchronization interval, Veeam Backup & Replication transports valid data blocks from the source backup repository and stores them to the newly copied restore point on the target backup repository.

Retention Policy

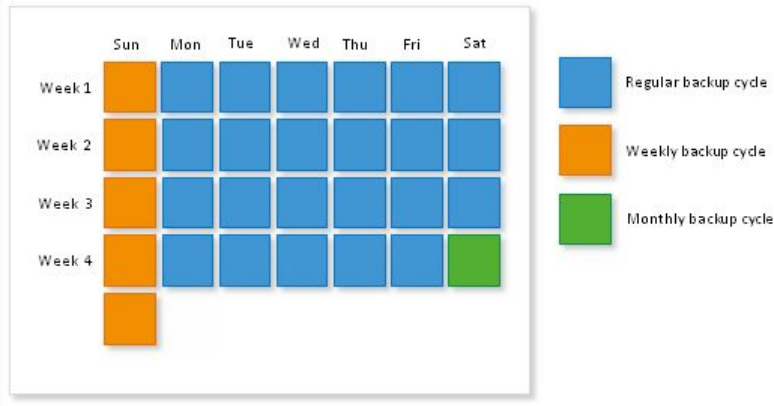


← simple


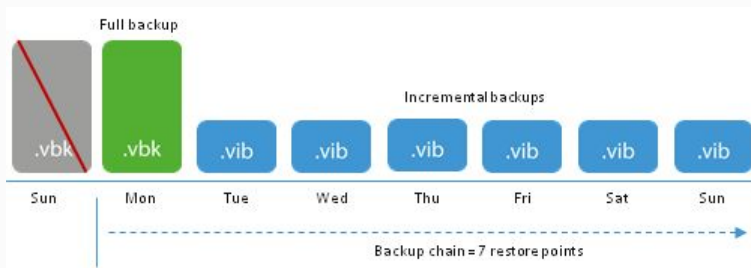
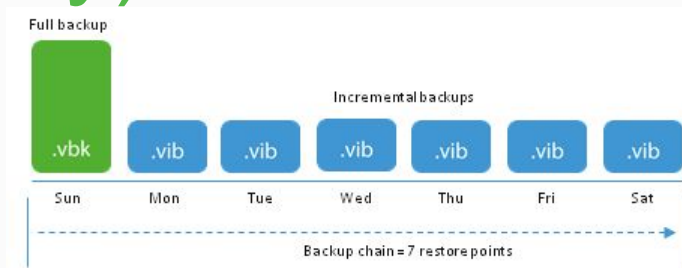


GFS Retention Policy

- Grandfather-Father-Son retention policy is a backup rotation scheme intended for long-term archiving




GFS (Weekly)



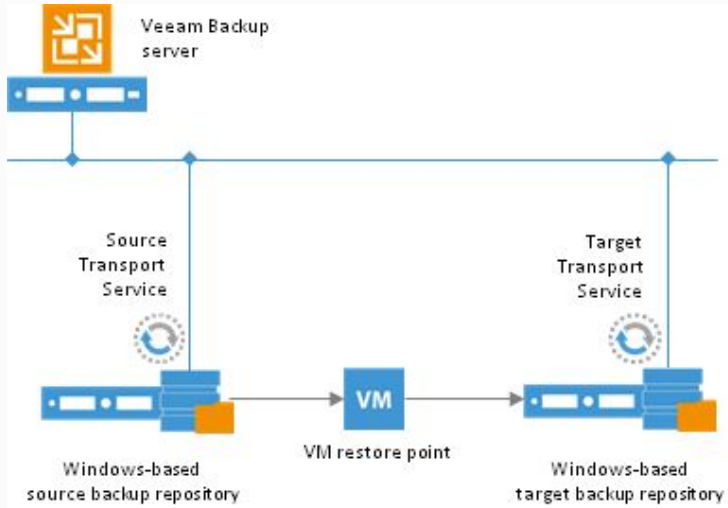
GFS (Weekly)

- The newly created full backup file remains in the backup chain and is used as a starting point for incremental backups created by the regular backup cycle

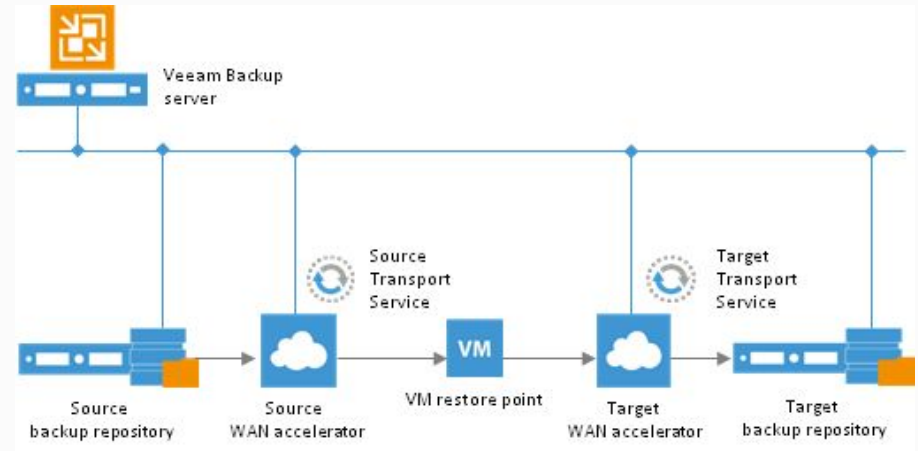


Data Transport Path

- Direct transport path



- Through built-in WAN accelerators




WAN Acceleration

- Global Data Deduplication
 - When you first run a remote job, Veeam Backup & Replication analyzes data blocks going over WAN
 - With every new cycle of a remote job, Veeam Backup & Replication uses the data redundancy algorithm to find duplicate data blocks in copied files. Veeam Backup & Replication analyzes data blocks in files on the source side and compares them with those that have been previously transferred over the WAN. If an identical data block is found, Veeam Backup & Replication deduplicates it.
- Veeam Backup & Replication deduplicates data blocks within one VM disk and in restore points for one VM only
- Deduplication between VM disks and restore points of different VM is performed indirectly, via the global cache.

WAN Acceleration

- VeeamWAN folder on the **source** WAN accelerator stores files with digests required for global deduplication
- VeeamWAN folder on the **target** WAN accelerator stores global cache data.



What's new in v9

- Active Full Backup: doesn't respect the usual sync interval
 - using extended sync interval for as long as it's needed
- Parallel processing of tasks:
 - VM's, but not disks
 - Several disks in one VM will be still processed sequentially

What's new in v9

- Read the entire restore point from the source backup instead of synthesizing it from increments (Default option for deduplicating storages):
https://helpcenter.veeam.com/backup/vsphere/backup_copy_gfs_active_full.html
- Backup Copy job can be used with the repository option “Use per VM backup files” to divide or gather the source backup

Thank you!

