

Backup Method - Virtual Full Backup

Backup Method - Virtual Full Backup


The Virtual Full Backup method was pioneered by R1Soft and overcomes limitations of the other known backup methods.

The chief drawbacks to all of the other backup methods that Virtual Full Backups overcome are:

- The need for a periodic Full Backup to be done e.g. every week. The weekly Full Backup that other methods require is extremely time consuming and intensive on critical server Disk I/O resources.
- Incremental Backup methods make restoring very complex as each incremental must be replayed to return files to a particular point in time.
- Incremental and Differential Backup methods provide no long term archiving or convenient rotation of old backups. For example you can not point and click Delete an Incremental backup without breaking other backup sets and losing data.
- The other backup methods are based on tape and do not properly take advantage of disk based storage. For example a Delta or unchanged file is stored in every Incremental backup set multiple times. This makes sense when your storage media is tape and it needs to be duplicated in each tape. Many Disk Based backups use Virtual Tape Library (VTL) software to make a file appear like a large tape. If the storage medium is Disk Based there is no need to store a Delta or unchanged file more than once! Virtual Full Backups overcome this limitation.

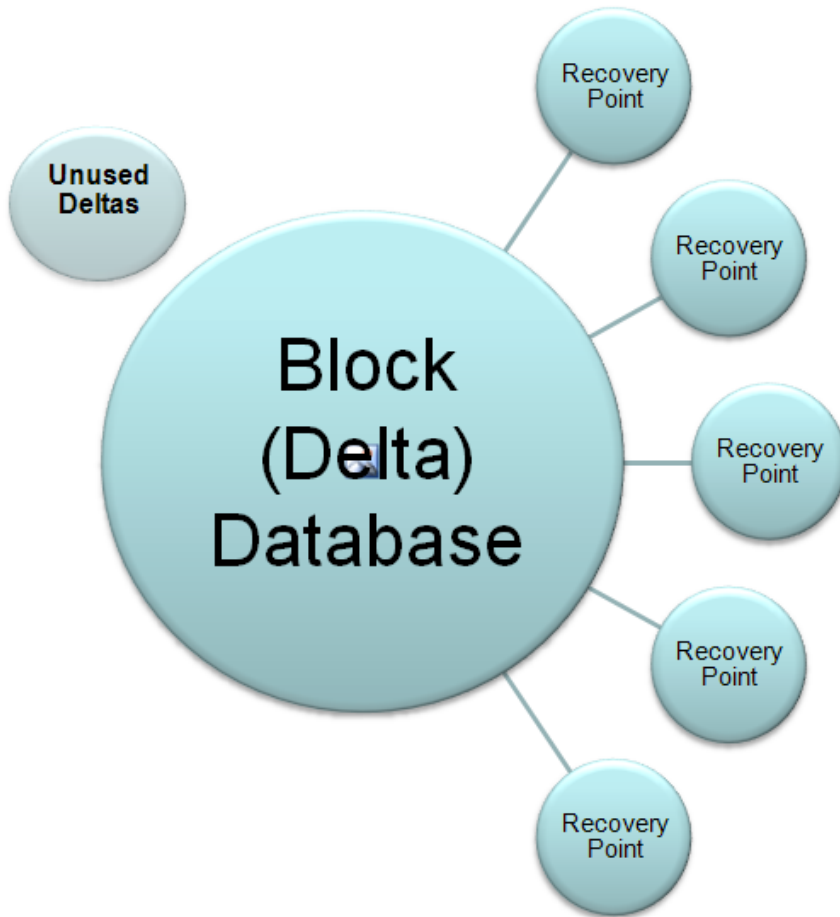
Initial Replica

The Virtual Full Method involves one initial Full backup called an Initial Replica. It is called an Initial Replica because it is performed only one time ever as long as the storage medium is not replaced. For example if the Virtual Full Backups are stored on a NAS device at \\myNAS\backups then there is only one Virtual Full Backup until a new storage location is used. This is very different from Incremental or Differential backups where the full backup is performed periodically for example weekly.



Disk Safe a Database for Block Deltas


The Virtual Full Backup method requires that a database be used to store block deltas. R1Soft calls this a Disk Safe®. While in theory an SQL type database could be used in practice they do not deal with large amounts of binary block efficiently so custom tailored data stores are used. The important function of the block deltas database is to map what versions of blocks are allocated and used by particular recovery points. This facilitates the efficient retrieval of blocks belonging to a recovery point as well as the process of merging. Merging involves deleting a recovery point and deltas specific to that recovery point not needed by other recovery points that are not being deleted.



Merging (Deleting) Recovery Points

In a Virtual Full Backup each recovery point is comprised of many deltas. A deltas represents the state of a disk volume block at a particular point in time and each delta can be used by one or more recovery point. Since the relationship between Deltas and Recovery Points is maintained by the Disk Safe it is possible to delete any unwanted Recovery Point.

1. The process of deleting a recovery point is actually a merge between the unwanted recovery point and closest recovery point still existing that created since the unwanted recovery point was taken.




2. The two recovery points are compared. Recall that not all parts of the Disk are used all the time. And as files are allocated and created some portions of the Disk which were used become unused available to store data.

As the two recovery points are compared two types of deltas can be discarded:

A) deltas associated with blocks that are allocated in recovery point N (the unwanted recovery point) and Not allocated in Recovery Point N+1.

B) deltas that are allocated in both recovery points that are not duplicated. Or deltas that changed from N to N+1 so that the corresponding deltas in N are no longer needed and can be discarded.



Advanced Recovery Point Archiving Rules


Merging recovery points allows for archiving rules. For example a backup policy or schedule can be defined to support the following example configurations:

- Synchronize Hourly and Keep only the last X Recovery Points
- Keep the Last X Hourly Recovery Points, Y Daily Recovery Points and Z Monthly Recovery Points

Synchronization Process

After the Initial Replica the process is very different from the other methods. Every time a backup job is run a Synchronization is performed. During the synchronization process a point-in-time snapshot of the Disk Volume is created and Deltas are computed based on the last completed synchronization. The Deltas are block level and are usually computed at a level below the file system very close to the raw disk structure. These deltas represent the low level changes to the disk volume made since the last recovery point.

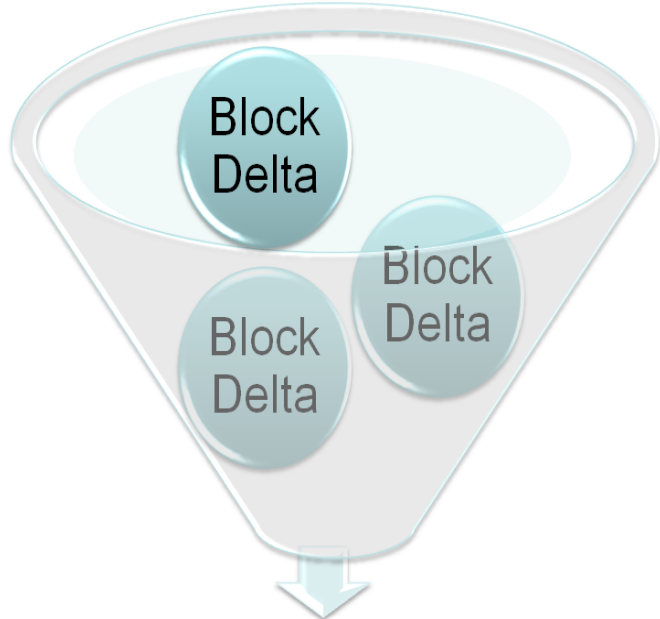
A Recovery Point appears like a Full backup to the user. That is only in appearance as each recovery point only contains block level Deltas or changes since the last Synchronization. The method for computing deltas can vary in implementation. R1Soft uses a near-Continuous Delta method described here [Computing Deltas - near-Continuous \(CDP\)](#). Deltas could potentially be computed by comparing all disk volume blocks every synchronization however this would be undesirable due to the time needed to do the comparison and heavy impact on disk I/O.



Restore with Virtual Full Backup Method

Restores are easy with Virtual Full Backups. The Disk Safe makes each Virtual Full Backup appear as a complete Full backup even though it is comprised of Deltas taken across many or even all recovery points. Each Recovery Point is then usable as a complete Volume or Disk image. R1Soft has special software capable of reading raw Disk blocks into a usable file system like NTFS on Windows or ext3 and reiserfs on Linux so that files can be browsed and restored seamlessly as if the files are being seen like they appeared on the live server at the point-in-time the synchronization was performed.

Behind the scenes the work of making a recovery point looks like the image below where Deltas are read from the Disk Safe to comprise a complete Full backup.



State of Volume at Synchronization (Recovery Point)

Virtual Full Backup Disk I/O and Storage Impact

The impact of the Virtual Full Backups on performance is as little as possible since only the deltas between synchronizations are read from the live server Disk(s). And unlike the other methods Deltas are only ever stored one time in the backup storage medium.

