



Technical report:
**SAP, UNIX and Oracle on IBM System
Storage N series**

Backup and Recovery Best Practices

• • • • • • • • •

Document NS3520-0

August 22, 2007



Table of contents

Abstract	1
Introduction	1
Conventional backup to tape	4
Snapshot and SnapRestore.....	4
Conclusion.....	4
Tape backup with backup server	4
Tape backup with backup server and NDMP (serverless backup).....	4
Tape backup on backup server with Brbackup Split Mirror	5
Disk-to-disk backup with NearStore and SnapVault.....	5
Conclusion.....	5
Backup and recovery using Snapshot and SnapRestore	7
Disk-to-disk backup with SnapVault and NearStore	11
SAP Split Mirror backup for Oracle	14
Backup with backup server	16
NDMP backup with backup server	18
Database verification	20
Trademarks and special notices	21



Abstract

Downtime of SAP production and even development systems is critical since it will always cause a financial impact on business. Coupled with this reality are the challenges of SAP data protection, including growing database volumes and shrinking backup and recovery windows. This document describes the unique benefits and proven best-practice techniques of backing up and recovering SAP systems that run on IBM System Storage N series.

Introduction

Corporations today require their SAP applications to be available 24 hours per day, seven days per week, 365 days per year. They expect performance to be maintained, irrespective of the increasing data volumes, and the system to undergo routine tasks, for example, backups, without influencing the SAP system. Performing backups of the SAP databases is a critical task, as backups can have a significant performance impact on the production SAP system. Since backup windows are shrinking and the amount of data to be backed up is increasing, it is a complex task to define a point in time in which backups can be done with minimum influence on the business process. Downtime of SAP production and even development systems is critical, since it will always cause a financial impact on the business process. Thus, the time needed for a restore and recovery is of particular importance.

Summary of SAP backup and recovery challenges:

- **Performance impact on production SAP system.** Backups typically have a significant performance impact on the production SAP system because there is a high load on the database server, the storage system, and the storage network during backups.
- **Shrinking backup windows.** Since conventional backups have a significant performance impact on the production SAP system, backups can only be made during times with low dialog or batch activities on the SAP system. It becomes more and more difficult to define a backup window when the SAP system is used 24x7.
- **Rapid data growth.** Databases are growing. Rapid data growth together with shrinking backup windows results in ongoing investments in the backup infrastructure—more tape drives, new tape drive technology, faster storage networks, etc. Growing databases also result in more tape media or disk space for backups. Incremental backups can address these issues, but result in a very slow restore process, which usually is not acceptable.
- **Increasing cost of downtime, decreasing mean time to recover.** The mean time to recover is the time that is needed to recover from a database failure (logical or physical error). The MTTR cuts into two areas—time that is necessary to restore the database and time that is necessary to do the forward recovery of the database. The forward recovery time depends on the number of redo logs that need to be applied after a restore. Unplanned downtime of a SAP system will always cause a financial impact on the business process. A significant part of the unplanned downtime is the time that is needed to restore and recover the SAP system in the case of a database failure. The backup and recovery architecture has to be designed according to the maximum acceptable unplanned downtime.
- **Backup and recovery time included in SAP upgrade projects.** The project plan for an SAP upgrade always includes at least three backups of the SAP database. The time spent performing these backups cuts down the total available time for the upgrade process.



IBM® System Storage™ N series provides unique functions and tools to address these challenges.

IBM System Storage N series with SnapDrive® for UNIX® simplifies storage management and provisioning in SAP Fibre Channel (FC) storage environments. It integrates with IBM System Storage N series with Snapshot™ and IBM System Storage N series with SnapRestore® to simplify the process of taking error-free, host-consistent data snapshots.

IBM N series with Snapshot can create an online or offline database backup in seconds. The time needed to create a snapshot copy is independent of database size, since Snapshot does not move any data blocks. The use of Snapshot doesn't have any performance impact on the production SAP system, since the IBM N series with Snapshot implementation doesn't have to copy data blocks when the data in the active file system is changed. Therefore, creation of snapshot copies can be scheduled without considering dialog or batch activities. SAP and IBM N series users typically schedule several snapshot online backups during the day—for instance, every four hours.

Snapshot technology also provides key advantages for the restore and recovery operation. The IBM N series with SnapRestore functionality allows restoring the entire database or parts of the database to the point in time of the last snapshot copy. This restore process is done in a few minutes, regardless of the size of the database. Because several snapshot online backups have been created during the day, the time needed for the following recovery process is also dramatically reduced. Fewer logs need to be applied, because a restore can be done to an at most four-hour-old snapshot copy. The mean time to recover, which consists of the time needed for restore and recovery, is therefore reduced to several minutes, compared to several hours with conventional tape backups.

Snapshot backups are stored on the same disk system as the active online data. Therefore, it is recommended that snapshot backups be used as a supplement, not as a replacement for backups to a second location, whether backup to a secondary storage system or tape. Though backups to a second location are still necessary, there is only a slight probability that these backups will be needed for a restore and recovery. Most restore and recovery actions will be handled by using SnapRestore. Restores from a second location (disk or tape) are only necessary if the primary storage system holding the snapshot copies is damaged or there is the need to restore a backup that is no longer available from a snapshot copy—for instance, a two-week-old backup. Restores from a second location can be considered as a specific situation that will happen very seldom.

A backup and recovery solution using an IBM N series fabric-attached storage (FAS) system always consists of two parts:

- Backup and restore/recovery using SnapDrive for UNIX with Snapshot and SnapRestore
- Backup and restore to/from a second location, which can be disk or tape.

A backup to a second location is always based on snapshot copies created on the primary storage. Therefore, the data will be directly read from the primary storage system without generating load on the SAP database server. Several options to back up the data to a second location are possible.

Disk-to-disk backup using an IBM System Storage N series with NearStore® and IBM System Storage N series with SnapVault® software:

- The primary storage directly communicates with the secondary storage (NearStore) and sends the backup data to the destination. The IBM N series with SnapVault functionality offers significant advantages compared to tape backups. After an initial data transfer, in which all the data has to be transferred from the source to the destination, all following backups only copy the changed blocks to the secondary storage. Therefore, the load on the primary storage system and the time needed for a full backup are significantly reduced. Since SnapVault only stores the changed blocks at the destination, a full database backup needs significantly less disk space.

Backup to tape using third-party backup software:

- **Local area network (LAN) backup.** The backup server mounts the snapshot copy and writes the data to tape.
- **Network data management protocol (NDMP) backup (serverless backup).** The tape is directly connected to the FAS system. The data is written to tape using NDMP.
- **SAP Brbackup split mirror.** SAP Brbackup creates the snapshot copies on the FAS system, and the data is written to tape with backint.

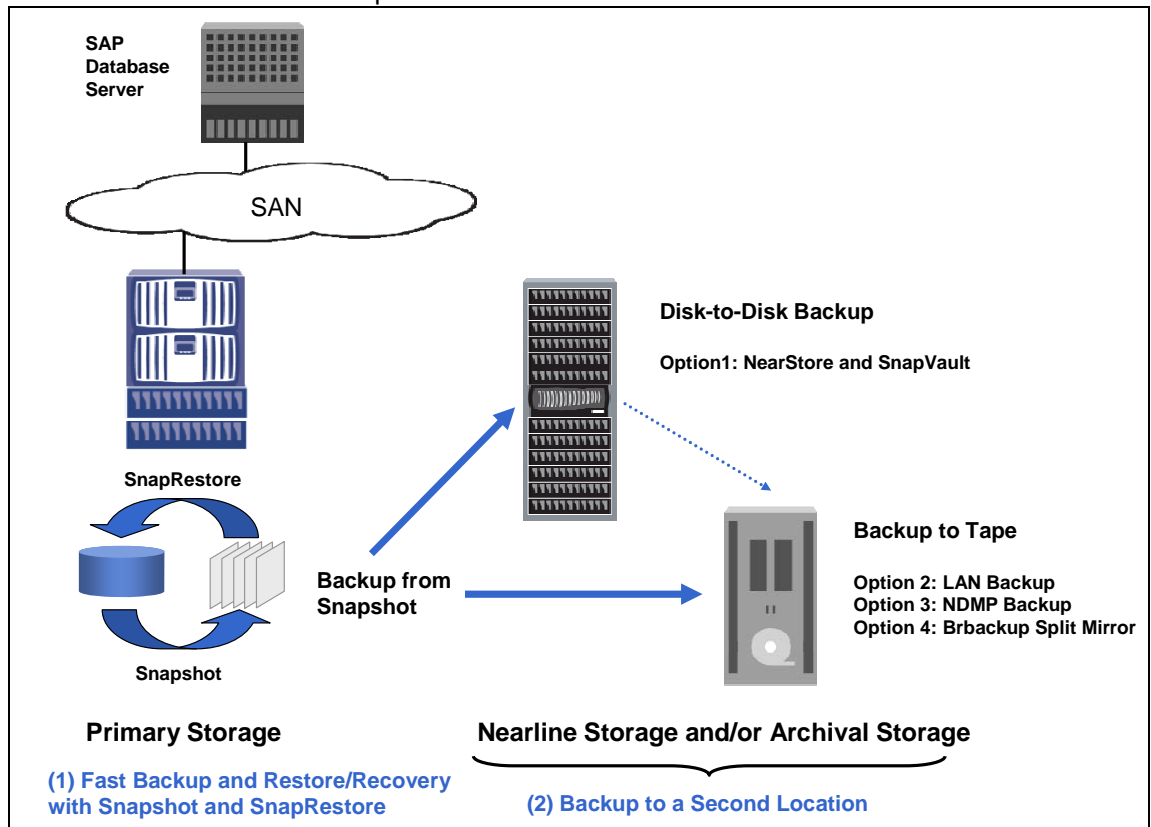


Figure 1. SAP backup and recovery—IBM N series solution overview.



Table 1 (further in the document) evaluates different backup and recovery concepts based on the most important attributes:

- Performance impact of the backup on the SAP production system: load that is generated by the backup on the database server, primary storage, and storage network
- Time that the database system will be in hot backup mode or offline
- Time needed for a full database backup
- Time needed for a full restore with recovery
- Space needed on tapes or disk for a full database backup
- Backup is usable for disaster recovery (DR)
- Integration into the SAP backup and recovery management tools.

Conventional backup to tape

A conventional backup to tape using SAP Brbackup with backint generates a significant load on the production SAP system: on the database server, the primary storage, and the storage network. Since this backup is not based on Snapshot, the database will be offline or in hot backup mode during the entire backup time. The time needed for the backup and for restore and recovery is high. Full backups will always need the full capacity on tape. The backup can be used for DR, and it is a fully integrated solution into SAP backup and recovery management.

Snapshot and SnapRestore

Snapshot backups do not generate any load on the database server, the primary storage, and the storage network. The database will be in hot backup mode or offline only for a few seconds. Using SnapRestore, the time needed for restore and recovery is very low. A full database backup based on Snapshot will only consume disk space for changed blocks. These backups can be automated using the operating system crontab scheduling facility or another third-party tool.

Conclusion

Snapshot and SnapRestore have significant advantages in all areas compared to a conventional tape backup. In a disaster scenario the primary storage system holding the snapshot backups might not be available anymore. Therefore, a backup to a second location must be accomplished as well.

All further backup concepts to a second location are based on snapshot copies created at the primary storage system. Therefore, with all concepts the database will only be offline or in hot backup mode for a short time.

Tape backup with backup server

When doing a backup with a separate backup server, which mounts the snapshot copy, there is no load on the database server. There is still a load on the primary storage system and storage network. The time needed for backup, restore and recovery may be significant. Full backups always need the full tape capacity. There is no integration into the SAP backup and recovery management.

Tape backup with backup server and NDMP (serverless backup)

When using NDMP with a backup server, the IBM N series FAS system writes the data directly to tape. Therefore, there will be no load on the storage network, but there will be a load on the primary storage. The time needed for the backup and for restore and recovery will be less, since the data is written directly to tape. Full backups will always need the full capacity on tape. There is no integration into the SAP backup and recovery management.



Tape backup on backup server with Brbackup Split Mirror

When executing Brbackup on a separate backup server, there will be no load on the database server, but there will still be a load on the primary storage and the storage network. The time needed for the backup and for restore and recovery may be significant. Full backups always need the full capacity on tape. The concept is integrated into the SAP backup and recovery management.

Disk-to-disk backup with NearStore and SnapVault

As SnapVault runs on the storage level, there is no load on the database server. SnapVault only transfers the changed blocks with each backup. Therefore, the load on primary storage and storage network is significantly reduced. For the same reason, the time needed to perform a full database backup is low. In addition, each full backup only stores the changed blocks at the destination. Therefore, the amount of disk space needed for a full backup is low compared to that needed for full tape backups. When executing a restore, it is always possible to do a full restore with no need to restore all incremental backups.

Conclusion

A disk-to-disk backup using SnapVault has significant advantages in all areas compared to all the above tape backup concepts.

	Backup Load On			Time Need			Space Needed for Backup	Backup Usable for Disaster Recovery	Integration within SAP
	Database Server	Primary Storage	Storage Network	DB Offline or Hot Backup Mode	For Backup	For Restore and Recovery			
Conventional Backup to Tape (e.g., Brbackup with backint)	High	High	High	High	High	High	High	Yes	Yes
Snapshot and SnapRestore	No load	No load	No load	Very low	Very low	Very low	Low	No	Partly
Backup to a Second Location Based on Snapshot									
Tape Backup with Backup Server	No load	High	High	Very low	High	High	High	Yes	No
Tape Backup with Backup Server and NDMP	No load	High	No load	Very low	Medium	Medium	High	Yes	No
Tape Backup on Backup Server with Brbackup Split Mirror	No load	High	High	Very low	High	High	High	Yes	Yes
Disk-to-Disk Backup with Near-line Storage and SnapVault	No load	Low	Low	Very low	Low	Medium	Low	Yes	Partly

Table 1).Comparison of different backup and recovery concepts.

The combination of Snapshot and SnapRestore with a disk-to-disk backup concept based on SnapVault offers significant improvement over conventional tape backups:

- Negligible impact of backups on the production SAP system
- Dramatically reduced mean time to recover
- Minimum disk space needed for database backups at the primary and the secondary storage systems (FAS system and NearStore system).

Database verification is an important part of a backup concept. Snapshot backups are well suited for performing a database verification using offline data files. Depending on the deployed backup concept, a database verification can run on a separate server with no load on the production database system. With the disk-to-disk backup concept, it is possible to run the verification on the secondary storage system, which means that the verification process is totally decoupled from the production SAP system.

The possibility of simply creating backups in seconds and being able to restore the SAP system to a point in time of any available snapshot copy is also very helpful in SAP test and development environments. Projects such as data import, SAP upgrades, or installation of support packages can be accelerated using fast backup and restore functionalities. During these projects, backups can be done at specific phases, and the system can be easily and quickly reset to a starting point in order to be able to repeat that phase.

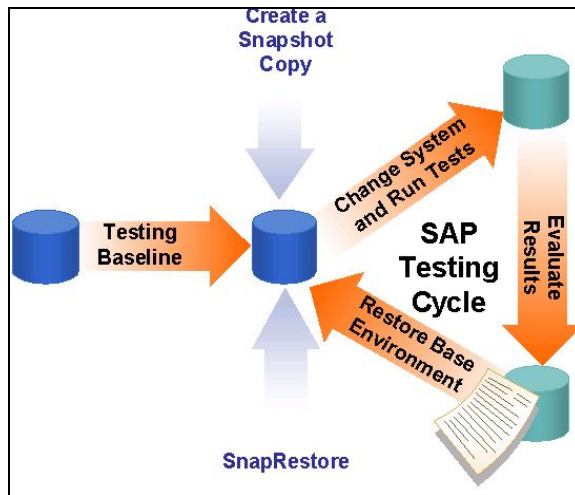


Figure 2. SAP testing cycle.

Performing an SAP upgrade or importing support packages and critical transports always influences the availability of the SAP system downtime. It is important that the time during which the system is not available is kept to a minimum and that the “old” status can always be restored. The specified system changes are usually first made in the development system in order to test the general functionality and procedures. In many cases, test systems must be upgraded several times, since problems can occur that can only be solved by restoring the system and restarting the upgrade. In this respect, Snapshot and SnapRestore technology save a considerable amount of time. A tape backup does not have to be made; a copy can be created instead. In the event of an error, the system can be quickly restored to its original status with SnapRestore and the upgrade repeated.

Time scheduling is extremely important when the production system is upgraded, since the system is unavailable at various stages in the upgrade. The time scheduling must also allow for restoring the system to its former release status. Depending on the size of the database and the time and effort required for the functional test and importing the transports for the modification adjustment, a 48-hour weekend may not be sufficient for the upgrade. Using Snapshot as a backup method and SnapRestore for restoring the system to its former release status ensures a higher level of flexibility with regard to scheduling. By creating several snapshot copies at certain stages of the upgrade, you can, to save time, restart the upgrade without having to fall back to the former release status.

Backup and recovery using Snapshot and SnapRestore

SnapDrive for UNIX simplifies storage management and provisioning in SAP FC storage environments. It integrates with IBM N series with Snapshot and SnapRestore to simplify the process of creating error-free, host-consistent data snapshot copies. SnapDrive for UNIX removes the need to maintain the manual scripts normally used to back up and restore data to specific drives or mountpoints without extensive downtime. And it can be used to easily add storage as needed, eliminating the need to preallocate too much storage. SnapDrive for UNIX is integral to all IBM N series and SAP FC solutions.

Snapshot can be used to create an online or an offline backup of an Oracle Database. With an online backup the Oracle Database has to be put in hot backup mode before the snapshot copy is created. With an offline backup the database will be shut down before the snapshot copy is created.

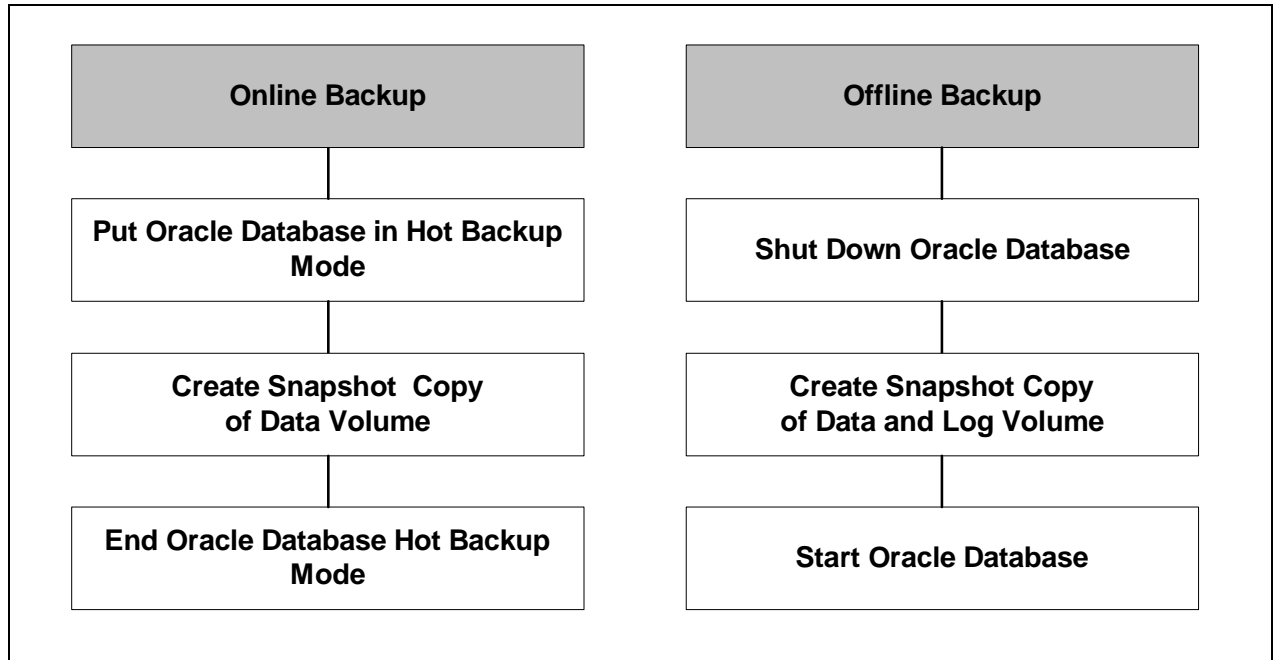


Figure 3. SAP database backups with Snapshot.

The creation of the snapshot copy is done in a few seconds. Therefore, the downtime with an offline backup and the time in hot backup mode with an online backup are very short. Snapshot does not have any performance impact on the running SAP system. There is no backup load on the database server or the storage system.

A conventional online backup to tape takes several hours and generates a significant load on the database server and on the storage system during this time.




Figure 4. Conventional backup to tape.

Snapshot backups take only a few seconds and do not generate any load on the database server or on the storage system. Snapshot backups will usually be scheduled more often, for instance, every four hours.




Figure 5. Snapshot backups.

A higher frequency of backups allows a more flexible restore strategy. Depending on the defined retention policy for the snapshot copies, there will always be a certain number of snapshot copies available that can be chosen for a restore. The number of redo logs that needs to be applied to do a forward recovery after a restore will be reduced with a higher backup frequency.

The amount of storage overhead associated with Snapshot technology is based on the number of blocks that are different between the snapshot copy and the active file system. As this size grows, the size of the snapshot copy's storage requirements increase as well. Therefore, the disk space needed for snapshots depends less on the number of snapshot copies than on the time frame in which a copy is stored. Recent snapshot copies are usually less onerous to store than older ones.

Based on the defined retention policy, there will be always a certain number of snapshot copies available that can be used for restores. Hourly, daily, and weekly snapshot copies can be created with different retention policies for each schedule.

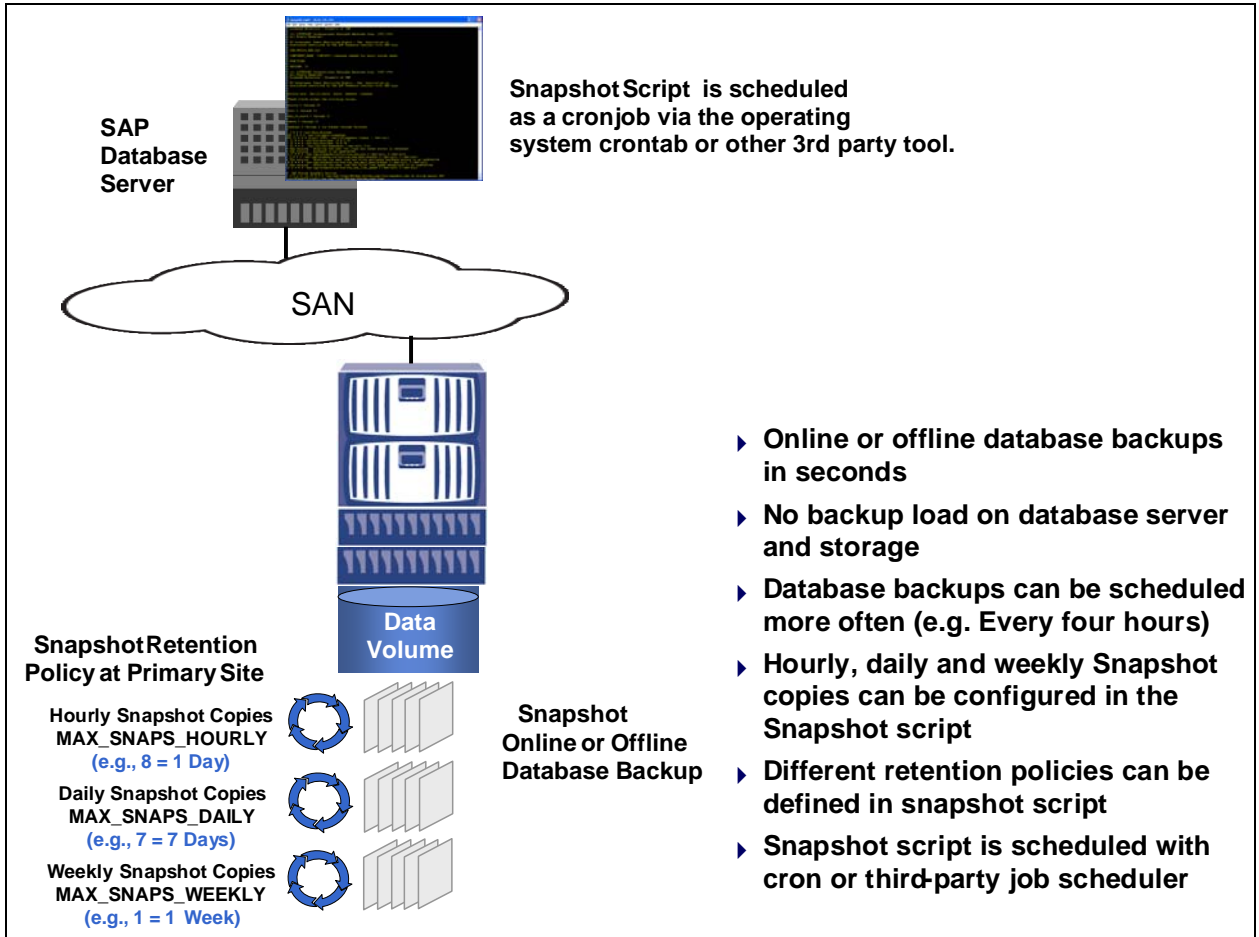


Figure 6. Snapshot schedule.

All the snapshot copies are visible using the `snapdrive snap list` command. Each snapshot copy has a unique name, which is defined when the snapshot copy is created.

```
[root@ibmp650-rtp01]#: snapdrive snap list -vg vgsapdatana7
```

snap name	host	date	snapped
fas270c-rtp01:/vol/vol_sapdata_na7:sap_online_NA7_manual.0	ibmp650-rtp01	Aug 3 11:17	vgsapdatana7
fas270c-rtp01:/vol/vol_sapdata_na7:sap_online_NA7_syscopy_NA8	ibmp650-rtp01	Aug 3 10:10	vgsapdatana7
fas270c-rtp02:/vol/vol_sapdata_na7:sap_online_NA7_manual.0	ibmp650-rtp01	Aug 3 11:18	vgsapdatana7
fas270c-rtp02:/vol/vol_sapdata_na7:sap_online_NA7_syscopy_NA8	ibmp650-rtp01	Aug 3 10:10	vgsapdatana7

From the logical point of view, the content of a snapshot copy is an online or an offline database backup, including all files that would be normally written to tape. All these snapshot copies can be used as a source for restoring with SnapRestore.

Let's assume that the database of the SAP system needs to be restored and recovered after a database crash at 10:00 a.m. The database size is 350GB, and the tape infrastructure restores with 50MB per second (175GB per hour). Furthermore, the SAP system produces 10 redo logs per hour, and applying a redo log during recovery takes one minute per redo log.

With a conventional tape backup, the database will be restored from the last backup, which was created at 8:00 p.m. Restoring the 350GB database will take about two hours. For a point-in-time recovery, the redo logs from 8:00 p.m. (the time when the backup was created) to 10:00 a.m. (the time of the database crash) have to be applied: 14 hours with 10 redo logs per hour results in 140 redo logs that have to be applied. Applying 140 redo logs will take about 140 minutes. If the redo logs are not on disk anymore, additional time will be needed to restore the redo logs from tape.

A conventional tape restore and recovery will have a total downtime of about 4.5 hours.

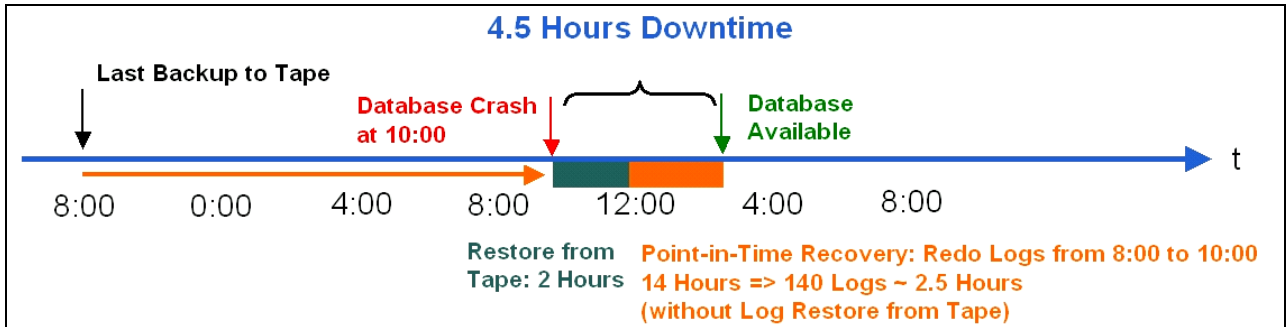


Figure 7. Conventional SAP database backups.

With IBM N series with Snapshot technology, the database will be restored from the last snapshot copy, which was created at 8:00 a.m. The database will be restored using SnapRestore. This restore is done in a few seconds. For a point-in-time recovery the redo logs from 8:00 a.m. (the time of the last snapshot copy) to 10:00 a.m. (the time of the database crash) have to be applied: two hours with 10 redo logs per hour results in 20 redo logs. Applying 20 redo logs will take about 20 minutes. Since only the redo logs of the last two hours are needed for the recovery, there is a good chance that the redo logs are still on disk and there is no need to restore them from tape.

A restore and recovery based on IBM N series with Snapshot and SnapRestore will have a total downtime of about 20 minutes.

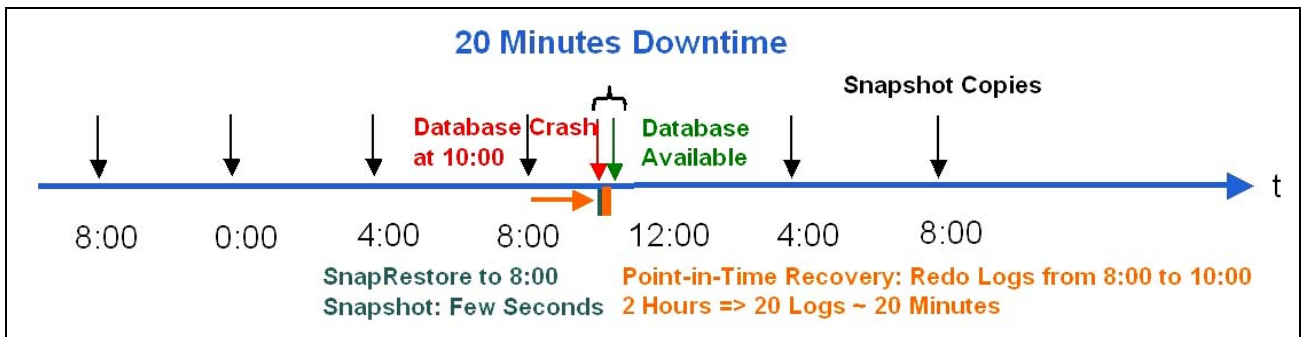


Figure 8. SAP database backups with Snapshot.

Backups based on Snapshot technology and the usage of SnapRestore dramatically reduce the mean time to recover and the downtime of the SAP system, from several hours to several minutes.

Disk-to-disk backup with SnapVault and NearStore

An alternative to traditional tape backups is to replace the tape architecture with ATA disk technologies and use tape capacity only for long-term archiving purposes. This cost-effective alternative is often referred to as nearline storage. NearStore technology provides a nearline storage solution to IBM N series systems. NearStore is a combination of the IBM System Storage N series with Data ONTAP® operating system with an array of ATA-based disk drives. This combination of inexpensive, ATA-based disks and the Data ONTAP operating system gives NearStore a number of advantages over traditional nearline storage.

SnapVault is a heterogeneous disk-to-disk data protection solution that is ideal for use with IBM N series with a NearStore feature. A SnapVault primary system corresponds to a backup client in the traditional backup architecture. The SnapVault secondary system is always a data storage system running Data ONTAP, such as a NearStore system. SnapVault software protects data residing on a SnapVault primary system, which will be an IBM N series FAS system.

There are two types of data transfers with a SnapVault backup:

- **Baseline initialization.** The baseline transfer will duplicate the entire source data set on the destination system. This baseline initialization has to be done only once. All subsequent transfers will be incremental updates.
- **Incremental updates.** Incremental updates are updates to an existing baseline initialization. An incremental update transfers the changed blocks between the baseline initialization or the previous Incremental update.

When comparing a conventional online tape backup with backups using Snapshot technology and SnapVault, we can see that a tape backup takes several hours and generates a significant load on the database server and on the storage system during this time.




Figure 9. Conventional tape backup.

With backups based on Snapshot and SnapVault, there is no backup load on the database server. The backup load on the primary storage system is significantly reduced, due to the fact that only changed blocks are transferred with each backup.

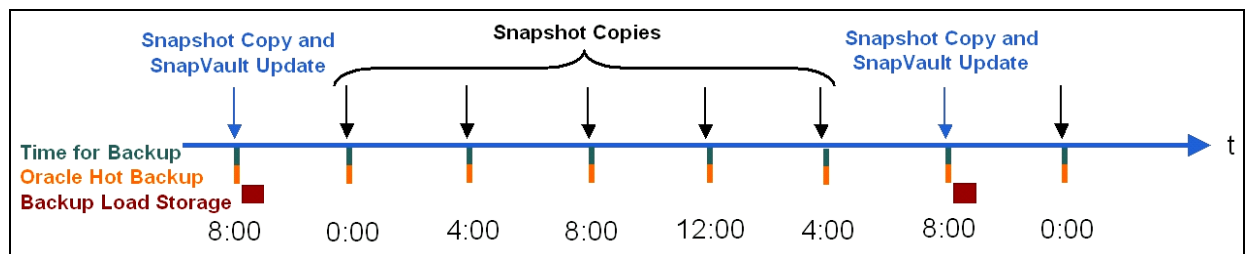


Figure 10. SnapVault backup.

Each incremental backup will be available as a snapshot copy at the SnapVault secondary system. Therefore, only blocks that have changed between two incremental updates will consume disk space. The “incremental backup forever” functionality of SnapVault saves a significant amount of backup media compared to a conventional tape backup solution.

Restores are still simple because each incremental backup is represented by a snapshot copy, which is a point-in-time view of the entire data set and can be restored in a single operation.

The Snapshot script is used to create snapshot copies on the primary storage system. In addition, the script supports SnapVault backups. For each schedule at the source—hourly, daily, weekly, and manual—SnapVault updates can be initiated. At the SnapVault destination, retention policies can be defined for each schedule. These retention policies are typically different from the retention policies at the source. For instance, daily snapshot copies at the source will be kept for seven days, while these snapshot copies will have a retention policy of four weeks at the destination.

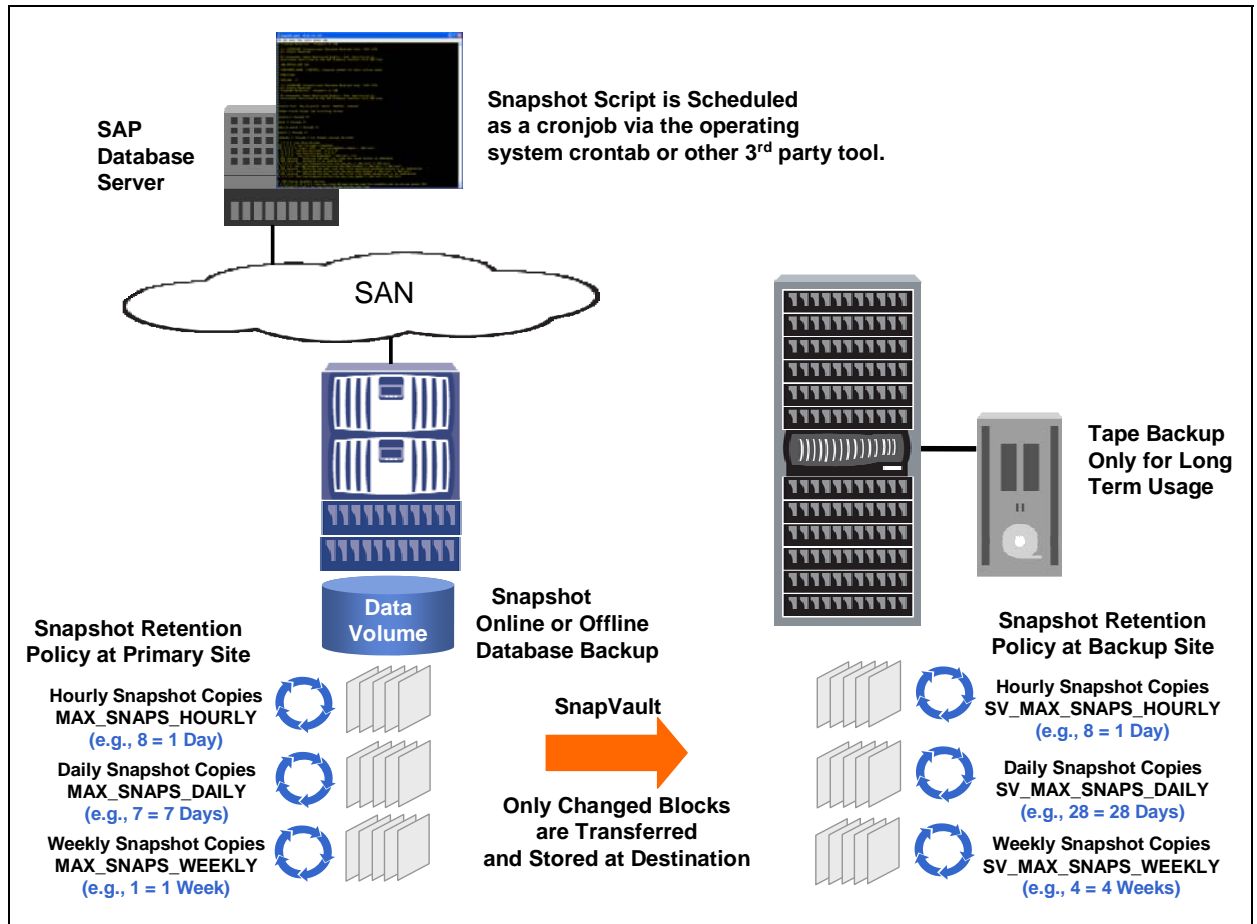


Figure 11. SnapVault backup overview.

Archive log backups can be accomplished using the available SAP tools `Brarchive` and `Brconnect`. `Brarchive` will be used to copy the archive logs to a second location, which will be a volume or qtree on the NearStore system (`brarchive -s`: save new archive logs). This volume or qtree will be mounted at the SAP database server. A second schedule of `Brarchive` will delete the archive logs on the primary systems (`brarchive -ds`: delete saved archive logs). A third schedule uses `Brconnect` to delete the archive logs at the destination (`brconnect -f` cleanup).

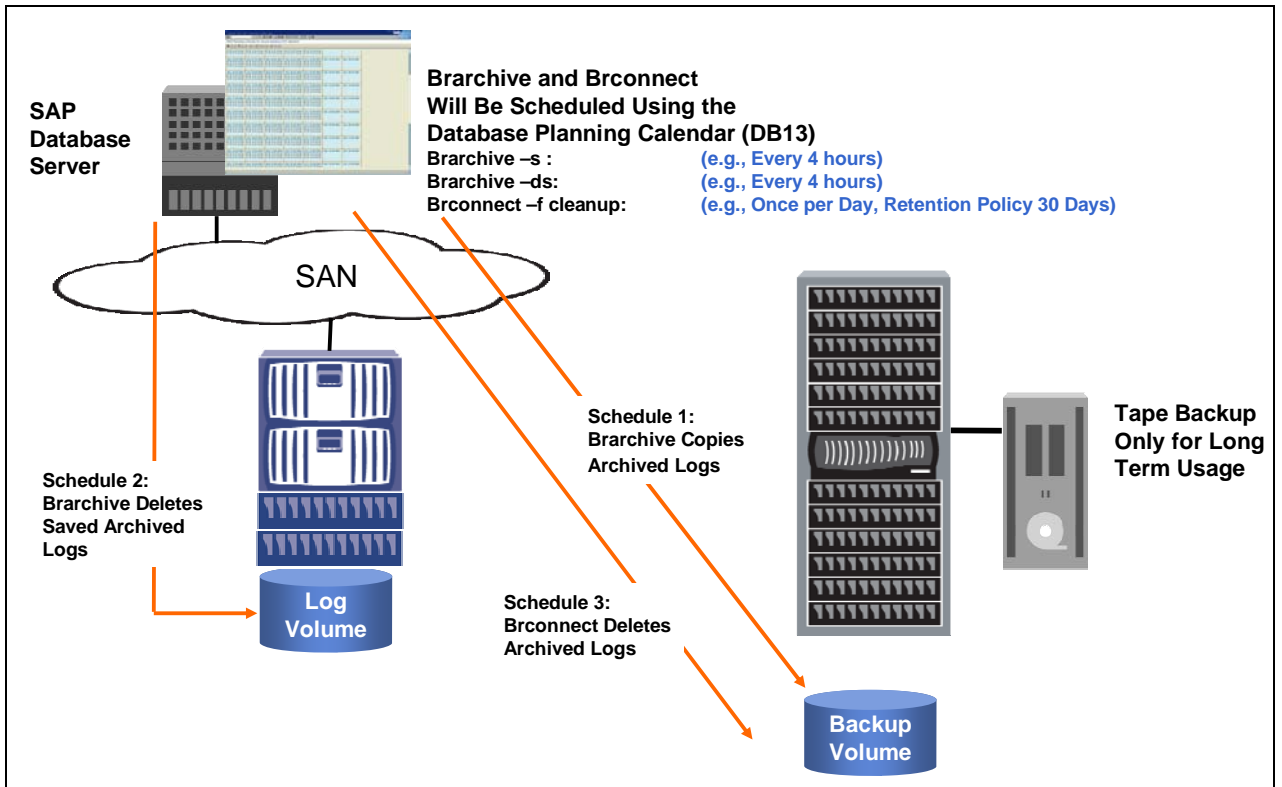


Figure 12. Archive log backup.



SAP Split Mirror backup for Oracle

The SAP backup tool `Brbackup` supports the backup of the database using mirrored disks. This mirror is split before the backup is made. Using this technique, it is possible to do an online or an offline backup. The time the database is offline or in hot backup mode is very short. The backup can be done on a separate backup server where `Brbackup` is implemented. Also, the verification options can be used to ensure the physical consistence of the database.

The usage of a separate backup server minimizes the performance impact of the backup or the database verification on the production database server. The backup of the database can be made using all backup methods that are supported by `Brbackup` (local tape, `backint`, etc.).

Using IBM N series with Snapshot technology offers several advantages in relation to normal techniques in the SAP split mirror environment:

- No double disk capacity is necessary for the mirror
- No complex configuration for the disk mirrors is necessary
- No file-system synchronization is necessary before the mirror is split
- Split command consists of simply creating a snapshot copy
- No file-system check is necessary after the mirror is created
- No need to synchronize the mirror disks after the backup.

When using a separate backup server, the backup consists of the following steps—executing `Brbackup` with the `-t offline_split` option (offline backup) or the `-t online_split` option (online backup):

1. `Brbackup` stops the database (offline backup) or sets the appropriate table spaces in hot backup mode (online backup). These commands are executed via SQL*Net.
2. `Brbackup` creates a snapshot copy.
3. `Brbackup` starts the database (offline backup) or ends the hot backup mode for the appropriate table spaces (online backup). These commands are executed via SQL*Net.
4. The database is backed up from the snapshot copy with or without database verification, or a database verification is done without a preceding backup.

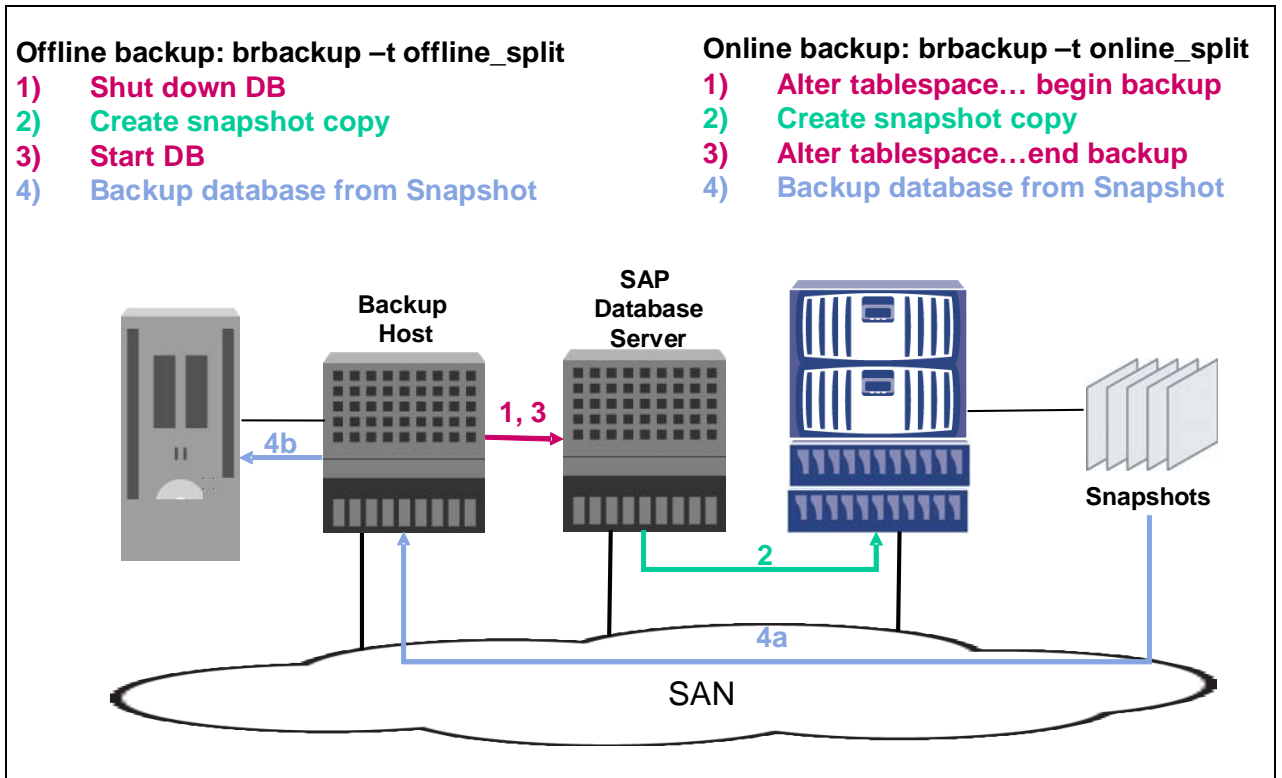


Figure 13. Split mirror backup with separate backup server.

The advantage of this approach is the short time that the database is offline or in hot backup. Additionally, there will be no load on the database server during backups and during database verifications. Therefore, the system can operate as usual, which means no performance impacts for batch jobs or dialog activities while the backup or the database verification is done.

Backup with backup server

The backup using a backup server is based on the snapshot copies created with the Snapshot script. The backup server will mount the data volume and will copy the data to tape. This schedule is independent of the snapshot schedule and has to be configured at the backup server.

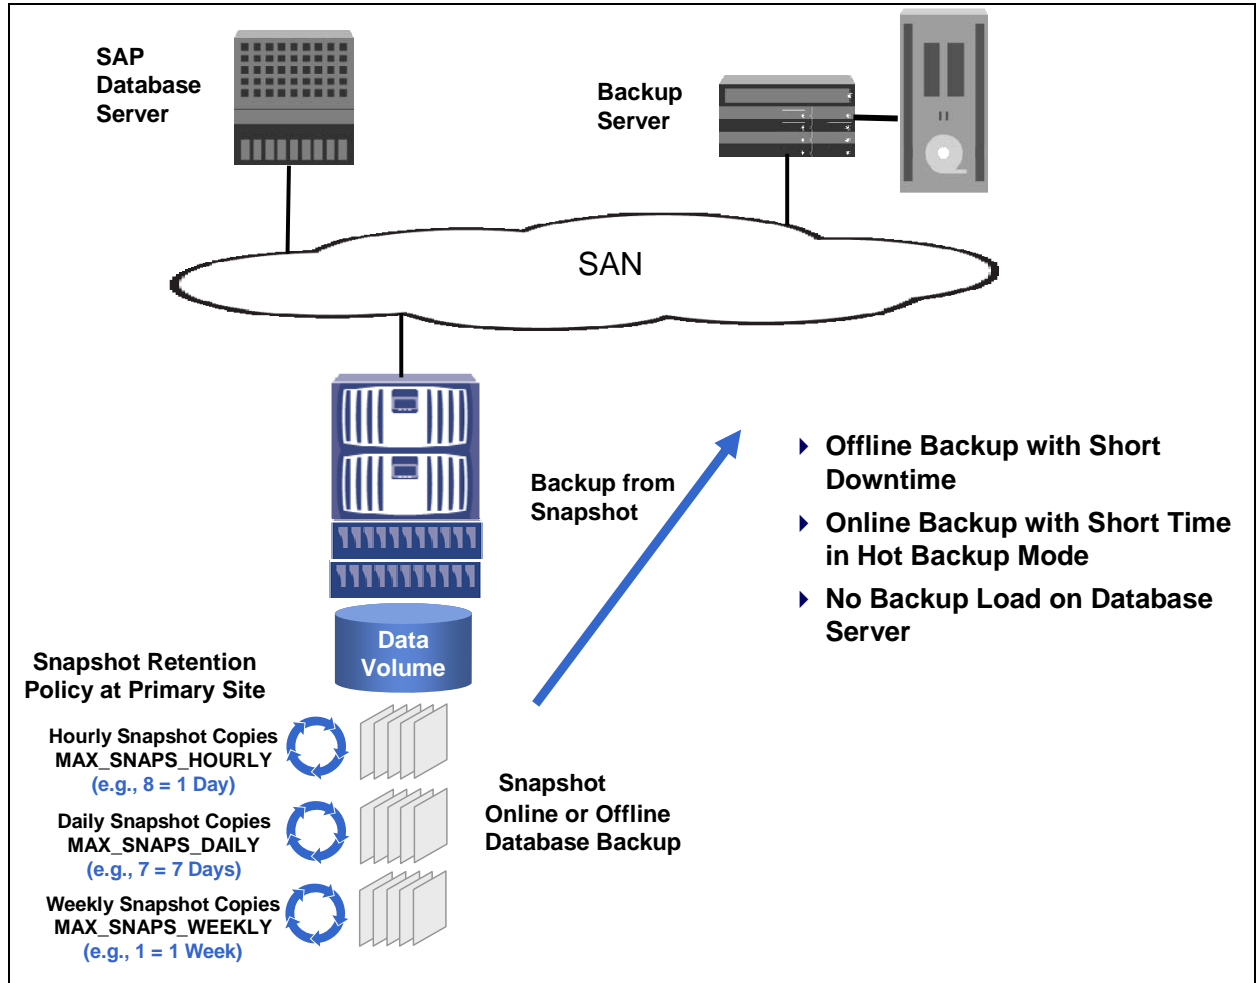


Figure 14. Backup with backup server.

Archive log backups can be accomplished using the available SAP tools `Brarchive` and `Brconnect`. `Brarchive` will be used to copy the archive logs to a second location, which will be an additional volume on the FAS system (`brarchive -s`: save new archive logs). This volume or qtree will be mounted at the SAP database server and at the backup server. A second schedule of `Brarchive` will delete the archive logs (`brarchive -ds`: delete saved archive logs). The backup server will copy the archive logs to tape. A fourth schedule uses `Brconnect` to delete the archive logs at the destination (`brconnect -f` cleanup).

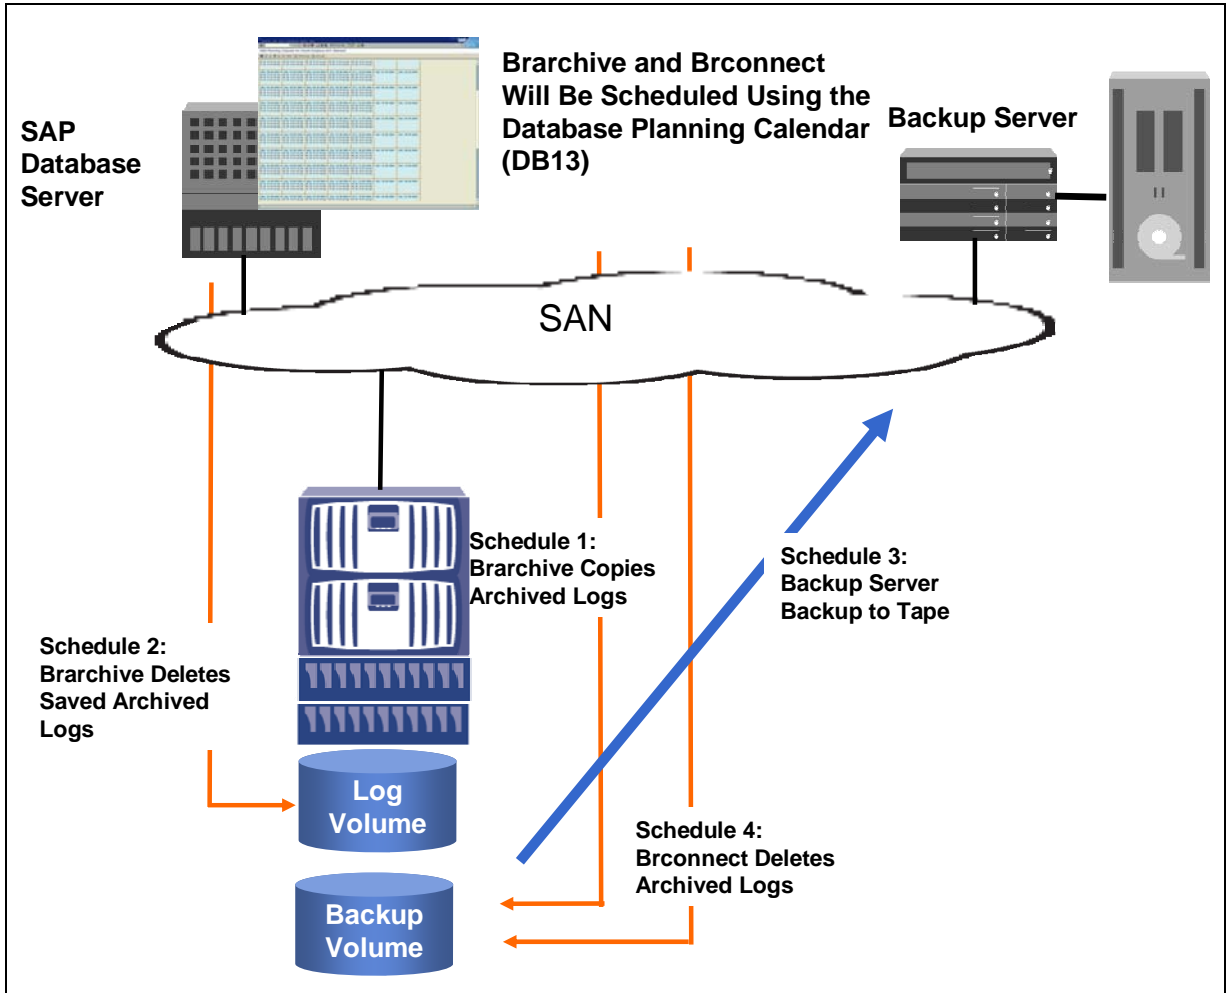


Figure 15. Archive log backup.

NDMP backup with backup server

The NDMP backup using a backup server is based on the snapshot copies created with the Snapshot script. The backup server will schedule NDMP backups on the FAS system, and the data will be copied from the data volume to tape. This schedule is independent of the snapshot schedule and has to be configured at the backup server.

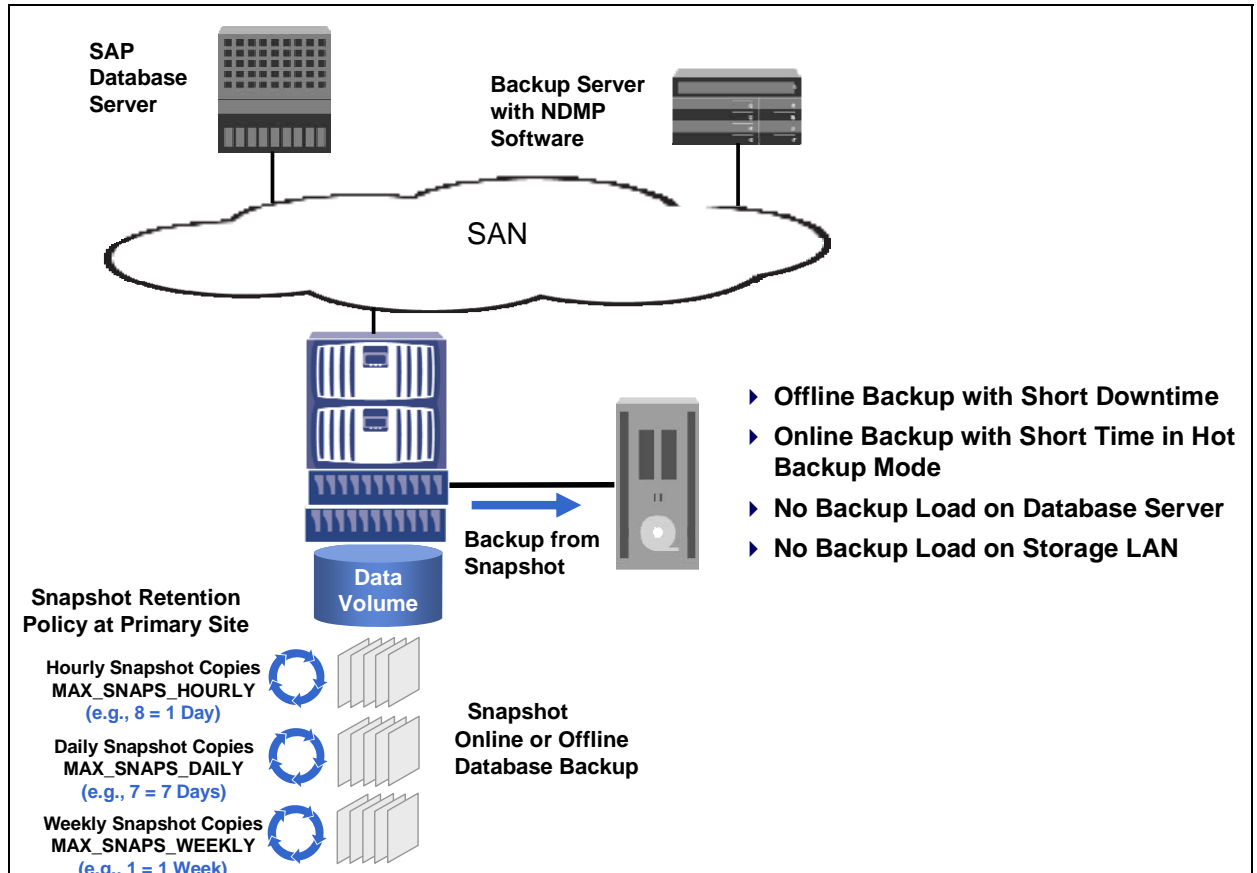


Figure 16. NDMP backup with backup server.

Archive log backups can be accomplished using the available SAP tools `Brarchive` and `Brconnect`. `Brarchive` will be used to copy the archive logs to a second location, which will be an additional volume on the FAS system (`brarchive -s`: save new archive logs). A second schedule of `Brarchive` will delete the archive logs (`brarchive -ds`: delete saved archive logs). The backup server will copy the archive logs to tape. A fourth schedule uses `Brconnect` to delete the archive logs at the destination (`brconnect -f cleanup`).

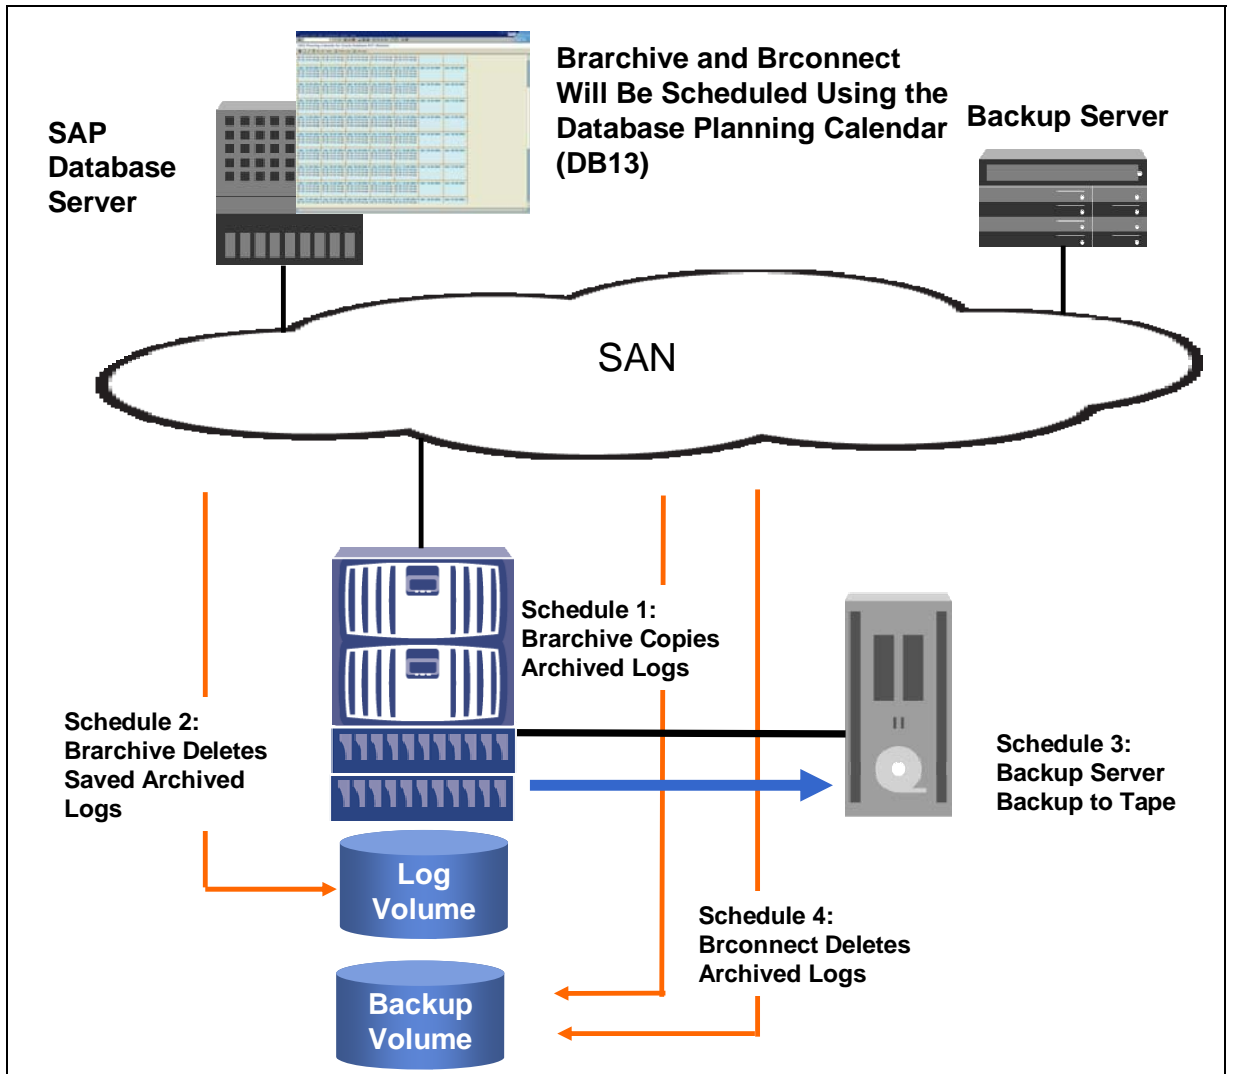


Figure 17. Archive log backup.

Database verification

Database verification is an important part of a backup concept. SAP recommends doing a consistency check of the database at least once per week (see also SAP OSS notes 540463 and 23345).

SAP recommends using the analyze command to do a database verification, if possible. The analyze command does the most checks but has the disadvantage of loading blocks into the SGA, which will negatively influence the buffer quality. In addition, up to Oracle V8, tables are locked during the analyze run, which will not allow the verification to run during online operation.

The advantages of dbverify are that blocks are not loaded into SGA during the verification and that it can run using offline files. If dbverify is used with offline files, there will be no influence on the production SAP system. Using offline files has the other advantage that a verification will be always correct. When running dbverify with online data files, which means that the database is online during the verification, dbverify can report a block corruption even if the block is not corrupted.

If the analyze command is not a choice for verifying the whole database, SAP recommends checking only the system table space with the analyze command and doing a dbverify for all the other table spaces.

	Run Time	Checks	Influence on Production System
Analyze Table, Verify Structure Cascade	Long	Tables Indexes Cross references -----	Blocks are loaded into SGA Tables are locked (up to Oracle 8)
Dbverify	Short	Tables Indexes ----- Blank db blocks	Blocks are <i>not</i> loaded into SGA Can be run on offline files (e.g., restored data files)
Export	Medium	Tables ----- ----- -----	Blocks are loaded into SGA

Table 2. Comparison of different database verification methods.

Snapshot backups are well suited for executing dbverify using offline data files. A consistency check based on `Brbackup` with the verify-only option can be accomplished on a separate server. This can be done independently of the chosen backup process.

When using a separate backup server that runs `Brbackup`, there will be no load on the database server during the verification. With the disk-to-disk backup concept it is possible to run the verification on the secondary storage system, which means that there will be no load on the primary storage system and therefore no influence of the verification on the production SAP system.

	Database verification with	Load on database server	Load on primary storage
Conventional backup to tape (e.g. <code>Brbackup</code> with <code>backint</code>)	<code>Brbackup</code> on DB server	High	High
Database Verification Based on Snapshot			
Tape backup with backup server over LAN	<code>Brbackup</code> on any separate server	No load	High
Tape backup with backup server and NDMP	<code>Brbackup</code> on any separate server	No load	High
Tape backup on backup server with <code>Brbackup Split Mirror</code>	<code>Brbackup</code> on backup server	No load	High
Disk-to-disk backup with Near-Line storage and <code>SnapVault</code>	<code>Brbackup</code> on any separate server	No load	No load

Table 3. Database verification based on Snapshot.



Trademarks and special notices

© International Business Machines 1994-2007. IBM, the IBM logo, System Storage, and other referenced IBM products and services are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. All rights reserved

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Network Appliance, the Network Appliance logo, Data ONTAP, NearStore, SnapDrive, SnapRestore, SnapVault and Snapshot are trademarks or registered trademarks of Network Appliance, Inc., in the U.S. and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.