



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

Backup and Recovery

• • • • • • • • •

Document NS3366-0

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 over LAN	4
Tape Backup with Backup Server and NDMP (Serverless Backup)	4
Tape Backup on Database Server with brbackup Split Mirror.....	4
Tape Backup on Backup Server with brbackup Split Mirror	5
Disk-to-Disk Backup with NearStore and SnapMirror	5
Conclusion	5
Backup and Recovery Using Snapshot and SnapRestore	7
Disk-to-Disk Backup Using SnapMirror	11
SAP Split Mirror Backup for Oracle	12
LAN Backup with Backup Server	15
NDMP Backup with Backup Server	16
Database Verification	17
Trademarks and special notices	18



Abstract

SAP backup and recovery is increasingly critical, yet production systems must remain available 24x7. IBM System Storage N series—including in combination with SnapShot, SnapMirror, SnapRestore, and NearStore feature options—provides unique functionality to meet SAP backup and recovery challenges with efficiency and performance.

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 regardless of increasing data volumes, and the system to undergo routine maintenance, for example, for backups, without influencing availability. Performing backups of the SAP databases are a critical task, since backups can have a significant performance impact on the production SAP system. Since backup windows are shrinking and the amount of data that needs to be backed up is increasing, it is a complex task to define a point in time in which backups can be performed with minimal influence on business processes. Downtime of SAP production and even development systems is critical, since it will always have a financial impact. The time needed for a restore and recovery is thus of particular importance.

SAP backup and recovery presents several challenges:

- **Performance impact on the 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 be made only during times with low dialog or batch loads on the SAP system. It becomes more and more difficult to define an appropriate backup window when the SAP system is used 24x7.
- **Rapid data growth.** Database sizes 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 require more tape media or disk space for backups. Incremental backups can address these issues, but result in a very slow restore process that usually is not acceptable.
- **Increasing cost of downtime and decreasing mean time to recover.** The mean time to recover (MTTR) is the time necessary to recover from a database failure (due to logical or physical error). The MTTR cuts into two areas—the 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 an 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 needed to perform these backups will reduce the total available time for the upgrade process.

IBM® System Storage™ N series provides unique functionalities that address these challenges.



IBM System Storage 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 the size of the database, since Snapshot does not move any data blocks. The use of Snapshot technology 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. Snapshot can therefore be scheduled without considering dialog or batch activities. SAP/IBM N series customers typically schedule several Snapshot online backups during the day—for instance, every four hours.

Snapshot also provides key advantages for the restore and recovery operation. The IBM System Storage N series with SnapRestore® functionality allows restoring the entire database to the point in time of any available snapshot copy. This restore process is done in a few minutes, independently 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 a four-hour-old snapshot copy. The MTTR that 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, IBM N series recommends using snapshot backups as a supplement, not a replacement, for backups to a second location—tape backup or backup to disk. 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 IBM System Storage N series with SnapRestore. Restores from a second location (tape or disk) are necessary only if the primary storage system holding the snapshot copies is damaged or there is a need to restore a backup that is no longer available in a snapshot copy. Restores from a second location can be considered as a specific situation that will seldom occur.

A backup and recovery solution using IBM N series filers will always consists of two parts:

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

A backup to a second location will always be 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.

There are several possibilities for doing the backup to a second location.

Disk-to-disk backup using an IBM System Storage N series with NearStore® system (or filer) and IBM System Storage N series with SnapMirror® software:

- The primary storage directly communicates with the secondary storage (NearStore system or filer) and sends the backup data to the destination. IBM N series with SnapMirror functionality offers significant advantages compared to tape backups. After an initial data transfer, in which all the data is transferred from the source to the destination, all following backups copy only 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 SnapMirror stores only the changed blocks at the destination, a full database backup doesn't consume much disk space.

Backup to tape using third-party backup software:

- **Local area network (LAN) backup.** The backup server mounts the snapshot copy of the data logical unit number (LUN) as a writable snapshot copy and writes the data to tape.
- **Network data management protocol (NDMP) backup.** The tape is directly connected to the filer; the data is directly written to tape using NDMP.
- **SAP brbackup split mirror.** brbackup creates the snapshot copy, and the data is written to tape using backint.

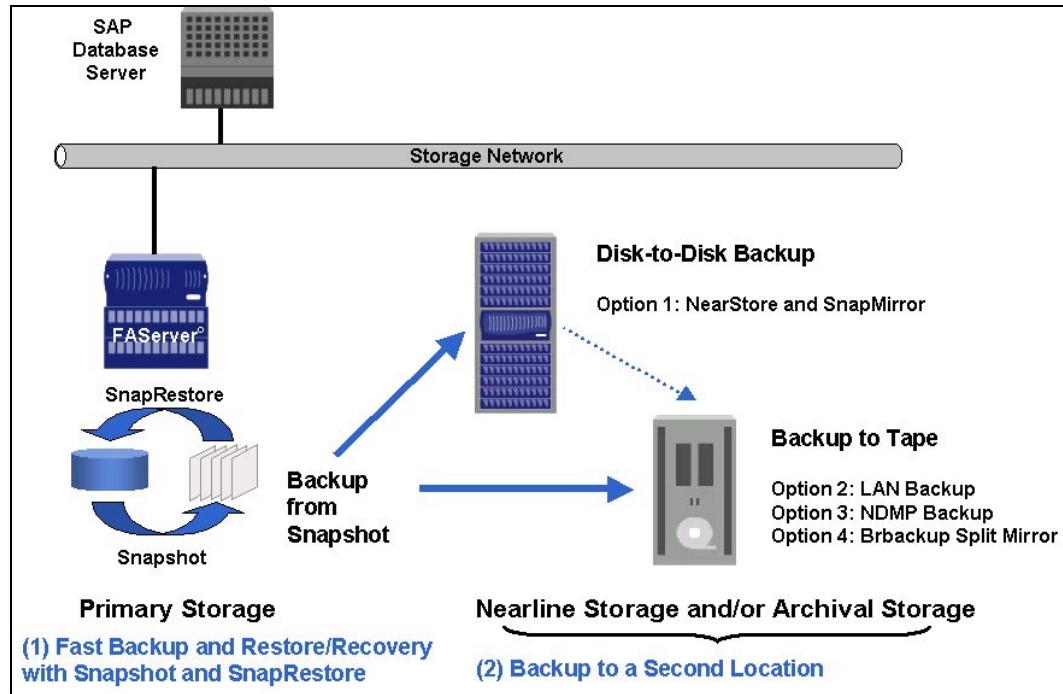


Figure 1. SAP backup and recovery—IBM N series solution overview.

Table 1 evaluates the 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 that 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, 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. The backup and restore speed is low. Full backups will always need the full capacity on tape. The backup can be used for DR, and it is fully integrated into SAP backup and recovery management.

Snapshot and SnapRestore

Snapshot and SnapRestore have significant advantages in all areas compared to conventional tape backups. Since snapshot backups are not a solution for DR, it is still necessary to do backups to a second location. A partial integration into SAP can be accomplished by using the SAP job scheduler to run the Snapshot script.

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 no longer be available. Therefore, a backup to a second location should be performed 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 over LAN

When using a separate backup server that mounts a writeable snapshot copy of the data LUN, there is no load on the database server. There is still a high load on the primary storage system and on the storage network. The backup and restore speed is low. Full backups will always need the full capacity on tape. The backup can be used for DR, and there is no integration into 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 filer 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 backup and restore speed will be higher compared to that of a backup over LAN.

Tape Backup on Database Server with `brbackup` Split Mirror

When executing `brbackup` on the database server, there is still a high load on the database server, primary storage, and the storage network. The time needed for the backup and for restore and recovery is high. Full backups will always need the full capacity on tape. The concept is integrated into 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 high load on primary storage and on the storage network. The time needed for the backup and for restore and recovery is high. Full backups will always need the full capacity on tape. The concept is integrated into SAP backup and recovery management.

Disk-to-Disk Backup with NearStore and SnapMirror

Since SnapMirror runs on the storage level, there will be no load on the database server. SnapMirror transfers only the changed blocks with each backup. The load on primary storage and the storage network is therefore significantly reduced. For the same reason, the time needed to perform a full database backup is low. In addition, each full backup stores only the changed blocks at the destination. Therefore, the amount of disk space that is needed for a full backup is small compared to that needed for full tape backups. When executing a restore, it is always possible to do a full restore without the need to restore all incremental backups.

Conclusion

A disk-to-disk backup using SnapMirror 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	Integration within SAP
	Database Server	Primary Storage	Storage Network	Database Offline or in Hot	For Backup	For Restore and			
Conventional Backup to Tape (e.g., <code>brbackup</code> with <code>backint</code>)	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 over LAN	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 DB Server with <code>brbackup</code> Split Mirror	High	High	High	Very low	High	High	High	Yes	Yes
Tape Backup on Backup Server with <code>brbackup</code> Split Mirror	No load	High	High	Very low	High	High	High	Yes	Yes
Disk-to-Disk Backup with Near-line Storage and SnapMirror	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 SnapMirror 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 secondary storage systems (filer and NearStore system).

Database verification is an important part of a backup concept. Snapshot backups are perfect for running a database verification using offline data files. Depending on the deployed backup concept, a database verification can be run on a separate server without creating any load on the production database 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.

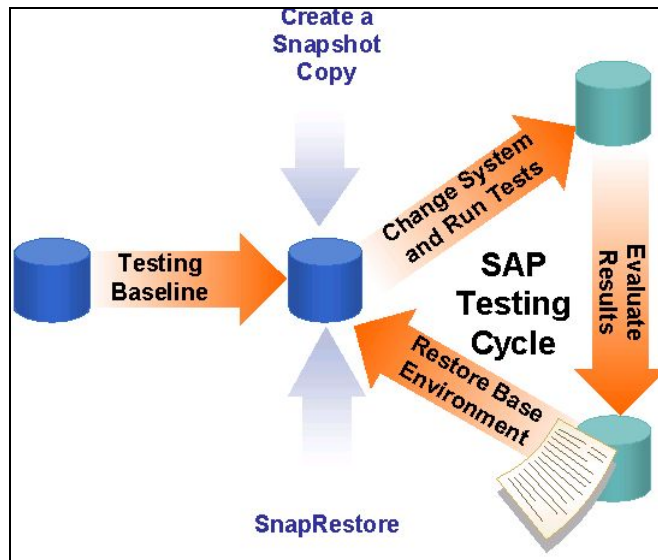


Figure 2. SAP testing cycle.

Carrying out an SAP upgrade or importing support packages and critical transports always involves SAP system downtime. It is important that this downtime be 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 be solved only by restoring the system and restarting the upgrade. In this respect, Snapshot and SnapRestore can save a considerable amount of time. A tape backup does not have to be made, since a snapshot 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 can be repeated.

Time management is extremely important when the production system is upgraded, since the system is not available at various stages during the upgrade. Scheduling must also include time 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, one normal

weekend may not be sufficient for the upgrade. Use of Snapshot as a backup method and SnapRestore for restoring the system to its former release status ensures a higher level of flexibility in scheduling. By creating several snapshot copies at certain stages of the upgrade, it is possible to restart the upgrade without having to revert to the former release status. This can save a lot of time.

Backup and Recovery Using Snapshot and SnapRestore

The key feature of the IBM N series approach to database backup is the use of Snapshot.

IBM System Storage N series with SnapDrive® is the management tool used to configure LUNs, create snapshot copies, and execute SnapRestore.

Snapshot can be used to create an online or offline backup of the Oracle® database. For an online backup, the Oracle database must be put in hot backup mode before the snapshot copy is created. For an offline backup, the database must be shut down before the snapshot copy is created.

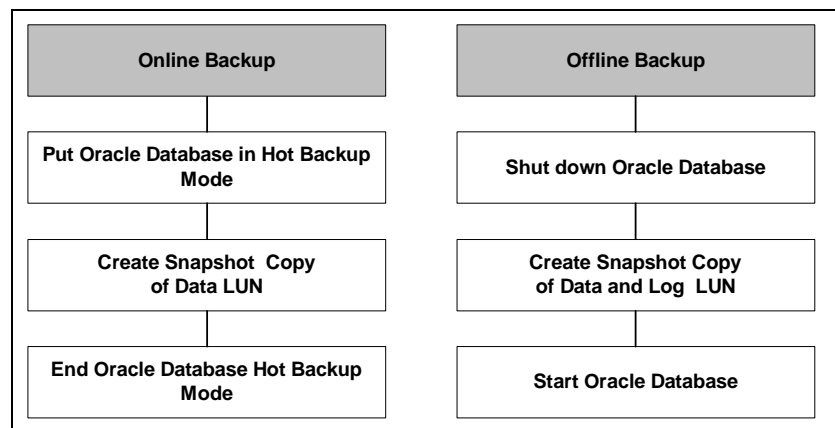


Figure 3. SAP database backups with Snapshot.

The creation of the snapshot copy takes only a few seconds — the downtime with an offline backup and the time in hot backup mode with an online backup are very short. Snapshot does not impact performance 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. The number of redo logs that must be applied to do a forward recovery after a restore is also reduced with a higher backup frequency. The amount of storage overhead associated with a snapshot copy is based upon the number of blocks that are different between the snapshot copy and the active file system. As this number increases, the size of the snapshot copy's storage requirements will increase as well. Therefore, the disk space needed for Snapshot depends less on the number of snapshot copies than on the time frame during which a snapshot copy is stored. Recent snapshot copies are generally less onerous to store than older ones.

Based on the defined retention policy, there will always be 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.

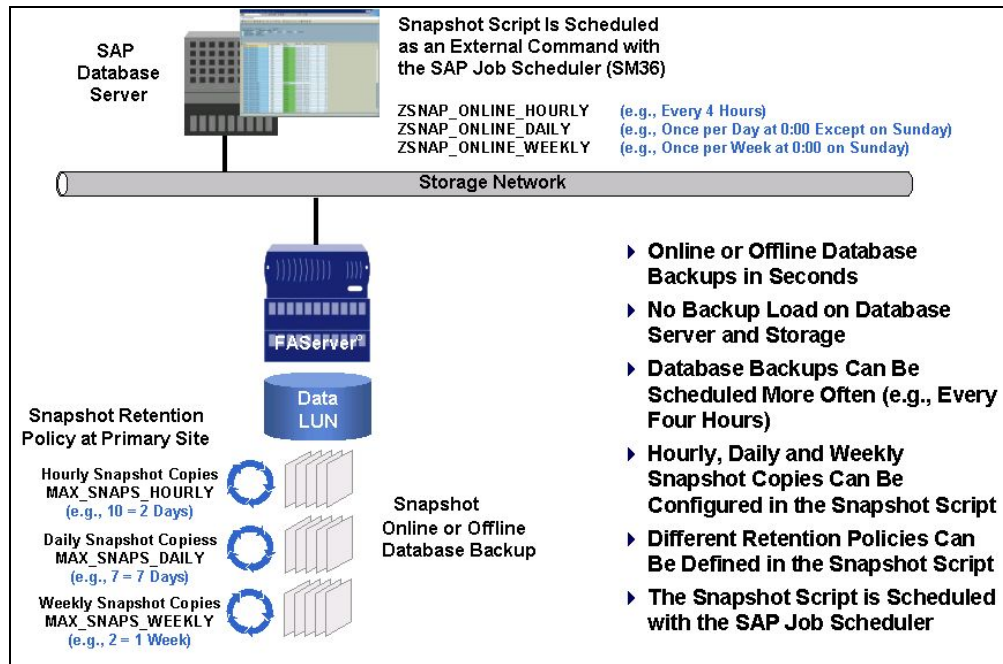


Figure 6. Snapshot backups overview.

All the snapshot copies for a LUN can be seen in the SnapDrive GUI. Each snapshot copy has a unique name, which is defined when the snapshot copy is created.

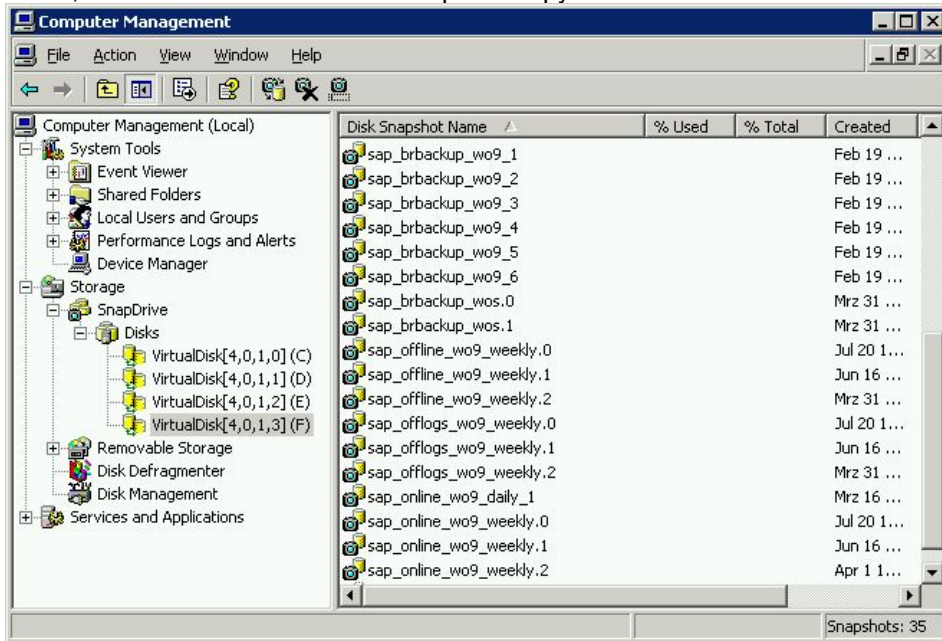


Figure 7. SnapDrive GUI.

Each snapshot copy contains a read-only copy of the data. This read-only copy can be used to create writable Snapshot LUNs, which can be mounted to any other server, e.g., for the creation of test systems.

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 normally be written to tape. These backups can be restored in few minutes by using the SnapRestore function integrated in SnapDrive.

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 at 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 no longer on disk, additional time will be needed to restore the redo logs from tape.

Note: Conventional tape restore and recovery has a total downtime of about 4.5 hours.




Figure 8. Conventional SAP database backups.

With IBM N series with Snapshot, 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 performed 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 from the last two hours are needed for the recovery, there is a good chance that the redo logs will still be on disk and there will be no need to restore them from tape.

Note: Restore and recovery based on IBM N series with Snapshot and SnapRestore has a total downtime of about 25 minutes.




Figure 9. SAP database backups with Snapshot.

Backups based on Snapshot and 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 Using SnapMirror

An alternative to traditional tape backups is replacement of 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 near-line storage. The NearStore technology was created to provide a near-line storage solution. 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 near-line storage.

IBM N series offers with SnapMirror a solution that enables asynchronous mirroring between different storage systems using a transmission control protocol (TCP)/IP network connection. Snapshot backups will be transferred automatically from the primary storage system (filer) to a secondary storage system (NearStore system or filer) using SnapMirror. SnapDrive always starts an update of a SnapMirror source volume to the destination volume after a snapshot copy is created. Therefore, the disk-to-disk backup solution using SnapMirror is an extension of the IBM N series with Snapshot backup and recovery solution described in Chapter 2: [Backup and Recovery Using Snapshot and SnapRestore](#). Only a SnapMirror relation between primary volume and secondary volume must be set up additionally.

After an initial transfer, SnapMirror transfers only changed physical blocks to the destination between two updates (incremental forever). Therefore, the load on the primary storage system is much lower, and less storage space is needed to store all backups compared to a traditional full backup to tape.

When comparing a conventional online tape backup with backups using Snapshot and SnapMirror, 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 10. Conventional tape backup.

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




Figure 11. SnapMirror backup.

Also, all existing snapshot copies of the primary site will be accessible on the destination site. Since snapshot copies are logical points of view of complete volumes, a restore of a snapshot copy from the destination to the source is always a complete full restore. It is not necessary to restore the first backup

(initial transfer) and then all others afterward, as must be done using conventional incremental backup. Therefore, the IBM N series disk-to-disk backup solution with SnapMirror combines the advantages of traditional full and incremental backups, while avoiding the traditionally associated disadvantages.

The disk-to-disk backup solution using SnapMirror is based on the snapshot backup solution described previously. When a snapshot backup is done, SnapDrive automatically triggers a SnapMirror update. Therefore, only a SnapMirror relationship must be configured, and that can be done during normal operations.

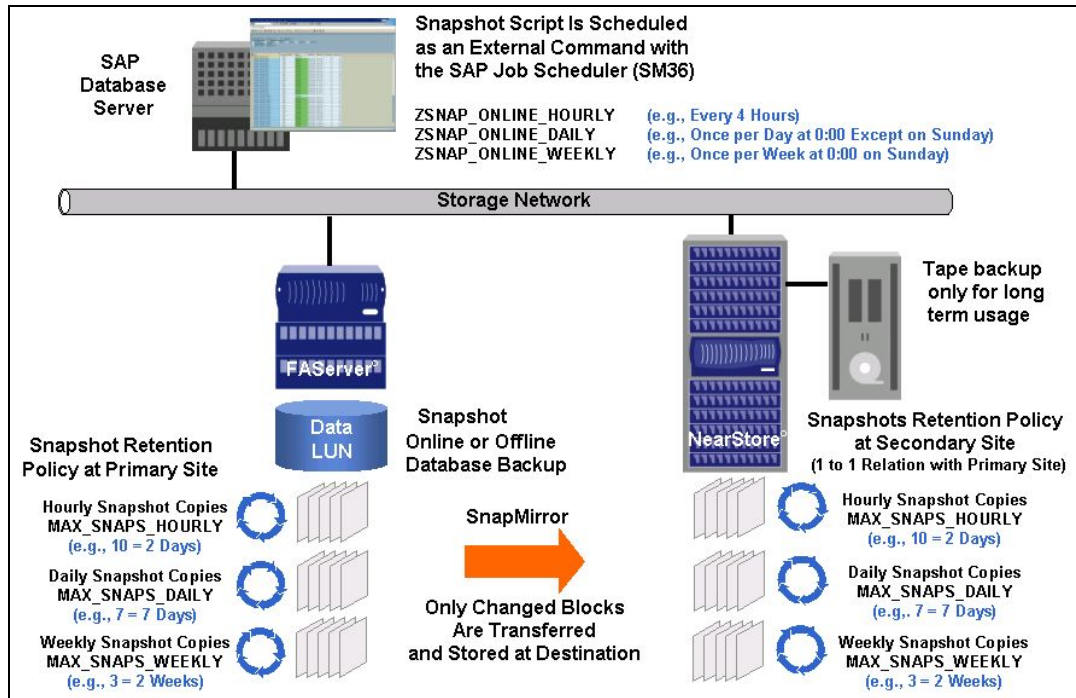


Figure 12. SnapMirror backup overview.

SAP Split Mirror Backup for Oracle

The SAP backup tool `brbackup` supports 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. This functionality enables the user to back up large databases or systems that do not allow large backup windows. The backup can be done either at the database server or with 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 influence 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.). An archive log backup will be performed as usual by executing `brarchive` at the backup server.

Note: 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
- Simple configuration for the mirrored disks using SnapDrive
- No file-system check necessary after the mirror is split
- Easy mount of the mirror disks after the mirror is split
- No resynchronization of the mirror after the backup.

When the backup is carried out on the database server, the backup consists of the following steps:

1. Executing `brbackup` with the `-t offline_split` option (offline backup) or the `-t online_split` option (online backup).
2. `brbackup` stops the database (offline backup) or sets the appropriate table spaces in hot backup mode (online backup).
3. The database server creates snapshot copies and connects these as writable Snapshot LUNs.
4. `brbackup` starts the database (offline backup) or ends the hot backup mode on the appropriate table spaces (online backup).
5. The database is backed up from the mirrored disks.
6. Verification of the database can be done additionally or without a backup.
7. The writable Snapshot LUNs will be disconnected.

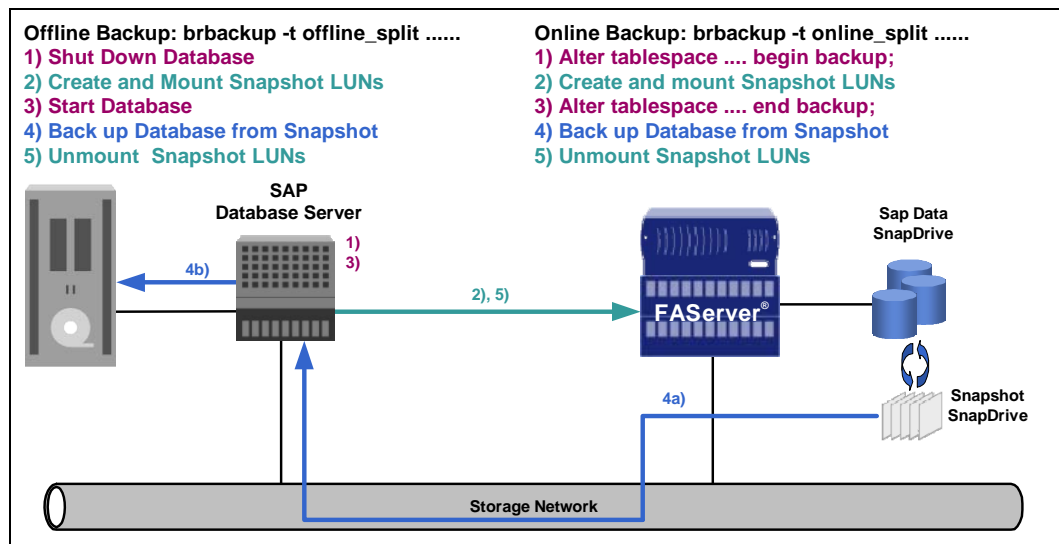


Figure 13. Split mirror backup at database server.

The advantage of this approach is the very short time that the database is offline or in hot backup mode. The load on the database server is the same as that for a standard backup.

When using a separate backup server, the backup consists of the following steps:

1. brbackup is executed with the `-t offline_split` option (offline backup) or the `-t online_split` option (online backup).
2. brbackup stops the database (offline backup) or sets the appropriate table spaces in hot backup mode (online backup).
3. These commands are carried out over SQL*Net.
4. The backup server creates snapshot copies and connects these as writable Snapshot LUNs.
5. brbackup starts the database (offline backup) or ends the hot backup mode on the appropriate table spaces (online backup). These commands are carried out over SQL*Net.
6. The database is backed up from the mirrored disks.
7. Verification of the database can be done additionally or without a backup.
8. The writable Snapshot LUNs will be disconnected.

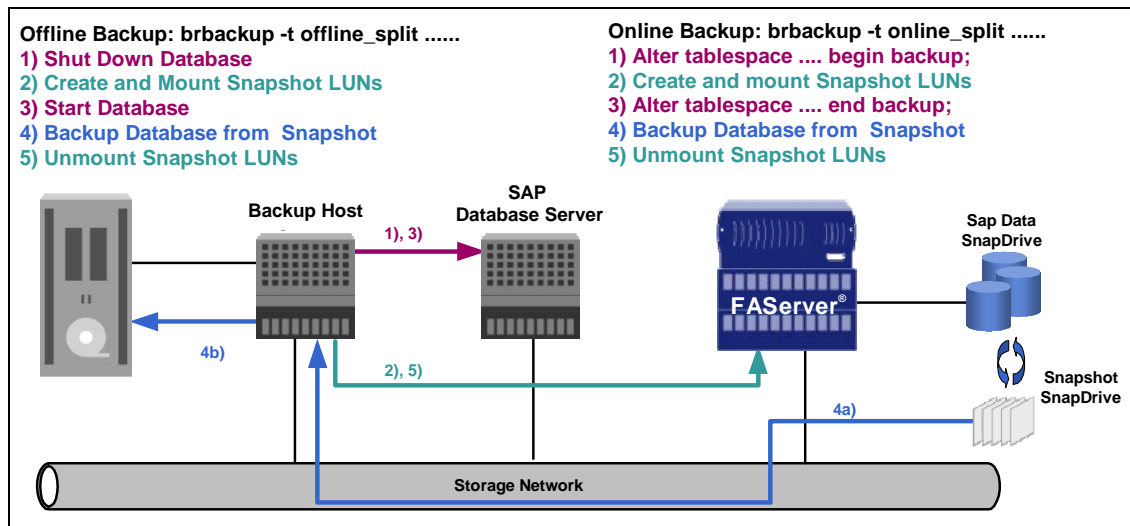


Figure 14. Split mirror backup with separate backup server.

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

LAN Backup with Backup Server

The LAN backup using a backup server is based on Snapshot, as described in Section 2: [Backup and Recovery Using Snapshot and SnapRestore](#). The backup server will mount the snapshot copy of the data LUN as a writable snapshot copy.

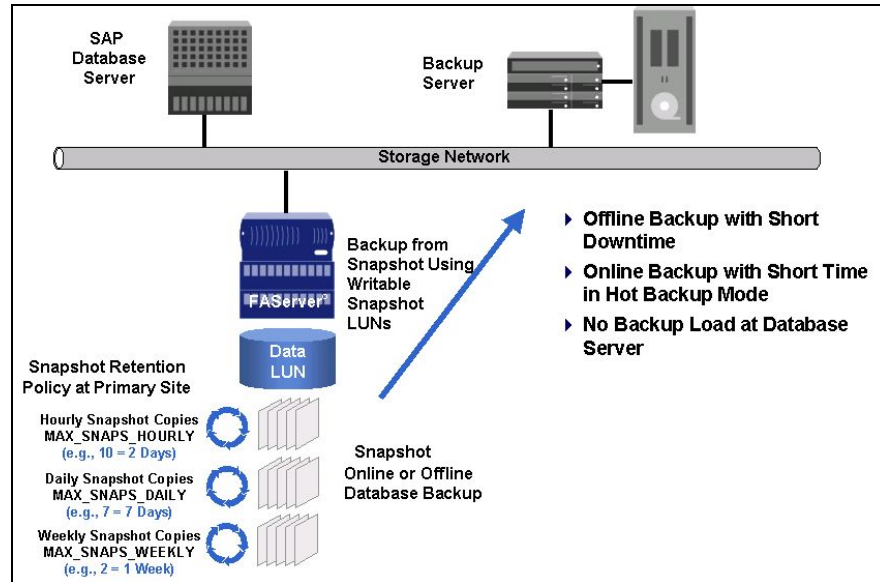


Figure 15. LAN backup with backup server.

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 LUN on the filer (`brarchive -s`: save new archive logs). This LUN will be connected to the backup server and will be shared on 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).

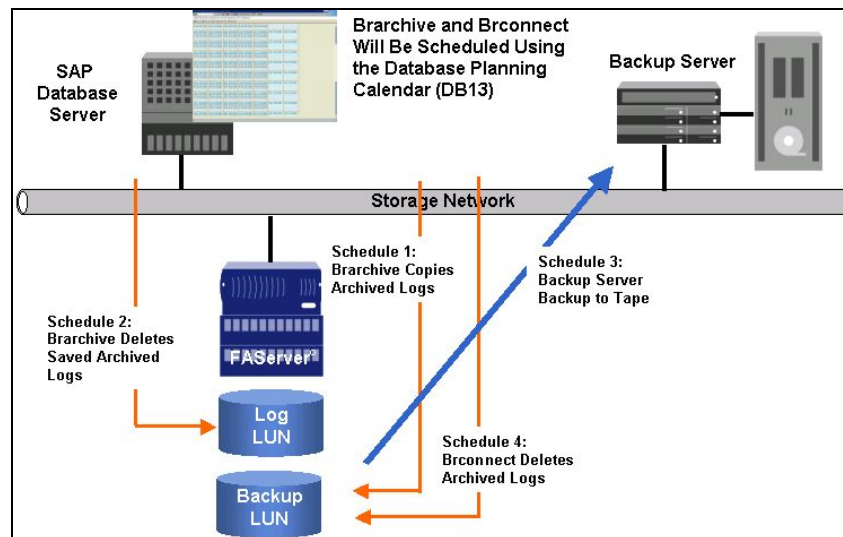


Figure 16. Archive log backup.

NDMP Backup with Backup Server

The NDMP backup using a backup server is based on Snapshot, as described in Section 2: [Backup and Recovery Using Snapshot and SnapRestore](#). The backup server will schedule NDMP backups on the filer, and the data (the whole LUN) will be copied from the Snapshot directory to tape.

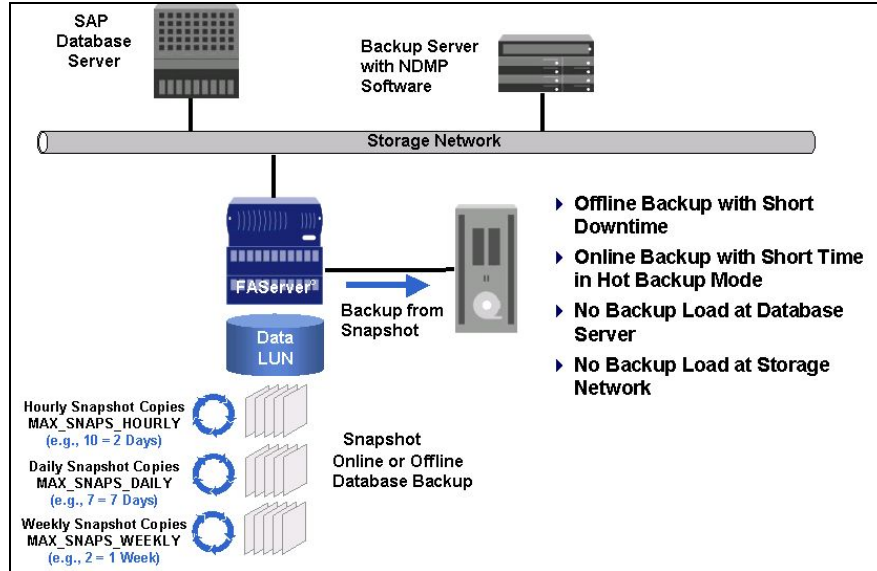


Figure 17. NDMP backup with backup server.

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 filer (`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.

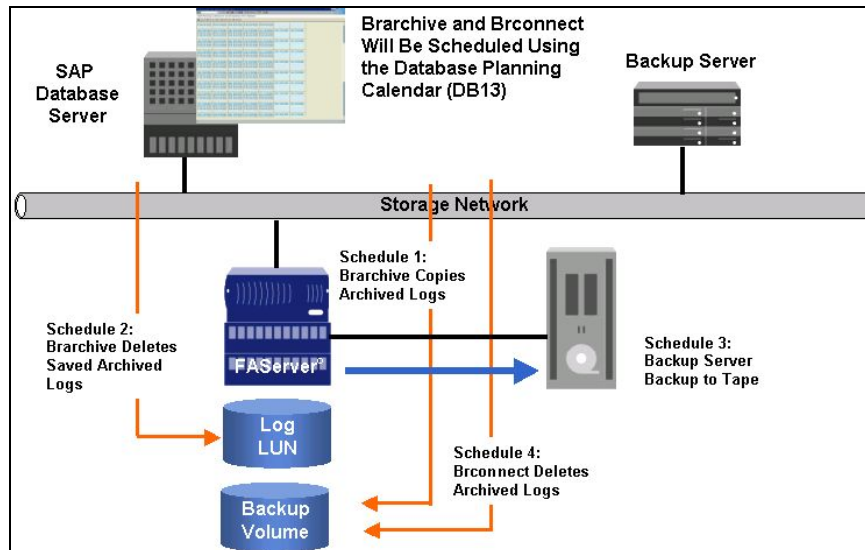


Figure 18. Archive log backup.



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 performs 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 prevents the verification from running during online operation.

The advantage of the `dbverify` command is 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 affect on the production SAP system. Using offline files also has the advantage that a verification will be always correct. If `dbverify` is run 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 using the `dbverify` command 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 v8)
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/Import	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.

	Database Verification With	Load on Database Server
Conventional Backup to Tape (e.g., <code>brbackup</code> with <code>backint</code>)	<code>brbackup</code> on DB server	High
Database Verification Based on Snapshot		
Tape Backup with Backup Server over LAN	<code>brbackup</code> on any separate server	No load
Tape Backup with Backup Server and NDMP	<code>brbackup</code> on any separate server	No load
Tape Backup on DB Server with <code>brbackup</code> Split Mirror	<code>brbackup</code> on DB server	High
Tape Backup on Backup Server with <code>brbackup</code> Split Mirror	<code>brbackup</code> on backup server	No load
Disk-to-Disk Backup with Near-line Storage and SnapMirror	<code>brbackup</code> on any separate server	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, Snapshot, SnapRestore, Data ONTAP, SnapDrive, SnapMirror and NearStore are trademarks or registered trademarks of Network Appliance, Inc., in the U.S. 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.