Technical report:

SAP on Windows and Oracle with IBM System Storage N series

Best practices

Document NS3444-0

October 3, 2007

© Copyright IBM Corporation, 2007. All Rights Reserved.
All trademarks or registered trademarks mentioned herein are the property of their respective holders

Table of contents

Abstract

This document provides customers and partners with the best practices for deploying IBM System Storage N series in support of mySAP Business Suite solutions running in a Microsoft Windows environment using an Oracle database. Overall better storage availability and management as well as reliable and efficient business continuance and data archival capabilities are the result of a combined SAP and IBM N series solution.

Introduction

Scope

This document is intended to provide customers and partners with the best practices for deploying IBM® System Storage™ N series in support of mySAP™ Business Suite solutions running in a Microsoft® Windows® environment using an Oracle® database. Primary consideration has been given to addressing the common storage infrastructure design, operation, and management challenges faced by business and IT leaders deploying the latest generation of SAP solutions. Recommendations are generic and not specific to a particular SAP application or size and scope of SAP implementation. This guide assumes a basic understanding of the technology and operation of IBM N series and SAP products.

Business challenges

Corporations deploying SAP solutions today are under great pressure to reduce total cost of ownership (TCO), accelerate return on investment (ROI), and increase productivity and availability of their SAP landscapes through infrastructure simplification. Restructuring activities, mergers and acquisitions, and constantly changing market conditions often result in the creation of new Enterprise Resource Planning (ERP) landscapes based on the SAP NetWeaver™ technology platform. NetWeaver permits more flexible adoption and integration of new business processes and scenarios. Timely access to data and the ability to analyze it not only become possible but also become requirements for keeping pace with change.

IT challenges

A typical production SAP landscape consists of several different SAP systems. It is important to the successful operation and management of these production instances that the same careful attention is paid to the number of non-production instances that are required.

SAP has long encouraged its customers to maintain separate development and quality-assurance (QA) instances for each production instance. In practice, it is not uncommon for such a three-system landscape to be expanded to include separate systems supporting functions such as a technical development and training. Driven by standard processes for development and testing within a corporation, it is also typical to have many development instances — and for more than one system to be used for QA and for additional testing or perhaps as a final staging system prior to releasing applications into production.

Adding to the challenge of maintaining these databases and the servers needed to drive them is the fact that each of these instances is going to have different performance, scalability, availability, and uptime profiles. These profiles can also fluctuate depending on the phases of project implementation and whether the project is focused on an existing SAP implementation or on a brand new one.

In summary, for each instance of SAP running in production, there can be as few as two and as many as five instances, or even more, supporting it. Because each SAP application requires its own database instance, deploying three SAP applications, like R/3, customer relationship management (CRM), and business warehousing (BW), can easily result in IT departments having to account for 15 or more SAP instances in total. All these instances need to be backed up, copied or cloned to support test schedules or to create a reference instance for new projects, and factored into a disaster recovery (DR) plan.

If the IT infrastructure supporting SAP applications is inflexible, is difficult to operate or manage, or has high cost of ownership barriers, that can negatively affect the ability of business owners to deploy new and improved business processes.

IBM N series solutions for SAP

IBM N series minimizes or eliminates many of the IT barriers associated with deploying new or improved business processes and applications. The combination of SAP solutions based on the NetWeaver platform and a simplified, flexible IBM N series storage infrastructure allows enterprises to work more efficiently and effectively toward the goal of improving enterprise business processes.

Storage consolidation with IBM N series ensures high availability and performance of SAP data and applications to meet stringent service-level agreements (SLAs). IBM N series also helps reduce the administration and management costs associated with deploying business applications and processes.

Storage provisioning and management

Consolidation

In today's rapidly changing business climate, enterprises demand cost-effective, flexible data management solutions that can handle the unpredictable and explosive growth of storage in heterogeneous environments. To enable global data management, ensure business continuity, satisfy regulatory and compliance standards, and improve resource utilization, a flexible and scalable storage network solution is required. The solution must also minimize complexity and reduce TCO.

IBM N series offers highly available, scalable, and cost-effective storage consolidation solutions that incorporate the IBM N series unified storage platform and the feature-rich functionality of data and resource management software to deliver storage that improves enterprise productivity, performance, and profitability, while providing investment protection and enhanced asset utilization. IBM N series enterprise-class storage solutions are interoperable across all platforms. IBM N series fabric-attached storage (FAS) systems integrate easily into complex enterprises and simultaneously support Fibre Channel (FC) storage area network (SAN) and IP SAN (iSCSI).

IBM System Storage N series with FlexVol™ technology delivers true storage virtualization solutions that can lower overhead and capital expenses, reduce disruption and risk, and provide the flexibility to adapt quickly and easily to the dynamic needs of the enterprise. FlexVol technology pools storage resources automatically and enables you to create multiple flexible volumes on a large pool of disks (aggregate). This flexibility simplifies operations, increases utilization and efficiency, and applies changes more quickly and seamlessly. IBM N series storage solutions enable you to add storage when and where it is needed without disruption and at the lowest incremental cost.

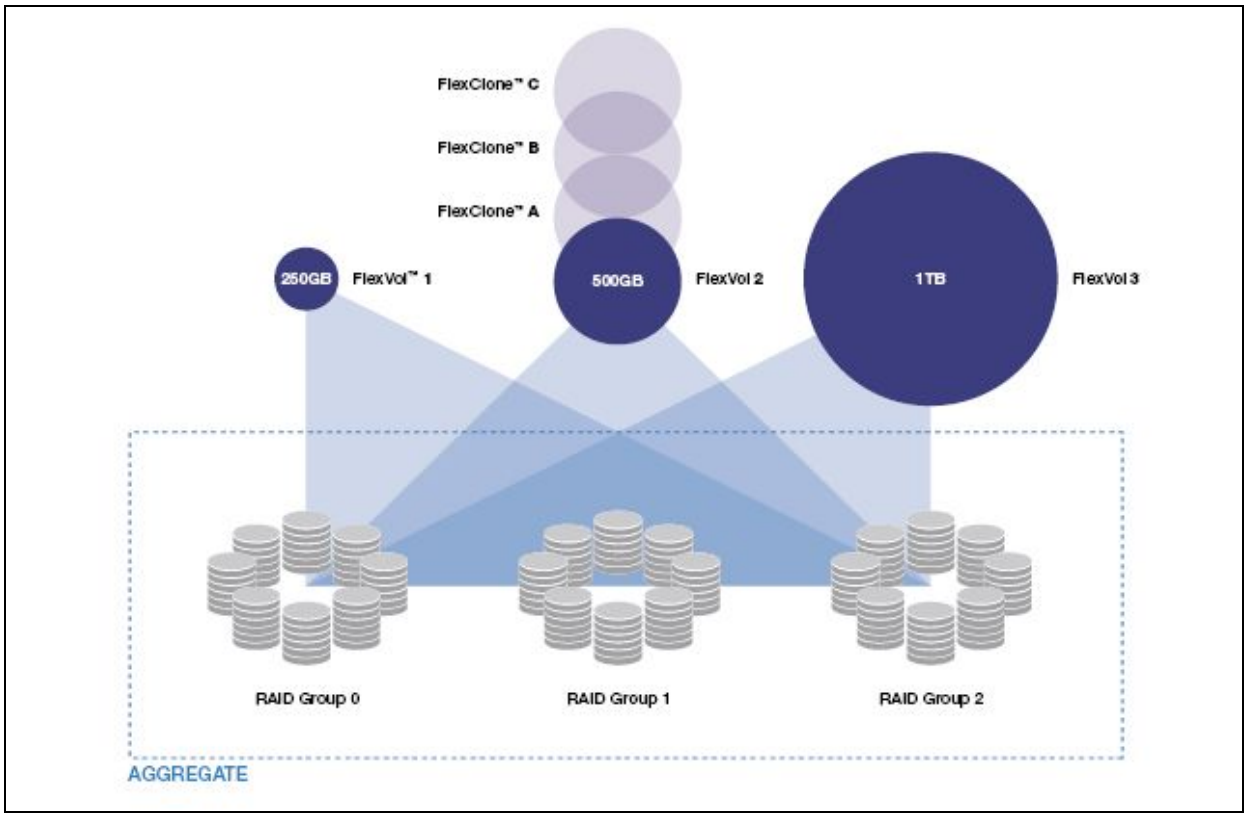

Figure 1. FlexVol technology.

IBM System Storage N Series with FlexClone™ technology enables true cloning—instant replication of data sets without requiring additional storage space at the time of creation. Each cloned logical unit number (LUN) is a transparent, virtual copy that can be used to test application patches, to run performance and data integrity tests, or to provide user-training environments with required copies of SAP components. IBM N series with FlexClone provides substantial space savings with minimal overhead. This means that many more data set variations can be managed—in less time and with less risk—to address and fuel the organization's business and development objectives.

With IBM System Storage N series with SnapDrive™ software, the cost and complexity of managing storage are reduced by enabling flexible and efficient utilization of storage resources to improve data and application availability. SnapDrive offers a rich set of capabilities to virtualize and enhance storage management for SAP environments. It is tightly integrated with NT file system (NTFS) technology and provides a layer of abstraction between application data and physical storage associated with that data. SnapDrive provides storage virtualization of volumes via the iSCSI and FC access protocols. IBM N series storage systems and SnapDrive software represent a complete data management solution for Windows applications. SnapDrive includes Windows device drivers and software that is used to manage IBM System Storage N series with Snapshot™ data backups. Snapshot backups are non-disruptive to applications and occur very quickly. Restoring data from a snapshot copy is nearly instantaneous.

The list below highlights some of the important benefits IBM N series storage systems offer to applications:

- Dynamic "on-the-fly" file system expansion, new disks are usable within seconds
- Patented, high-performance, low-latency file system with industry-leading reliability
- Robust yet easy-to-use data and storage management features and software
- Virtual disks created within a dynamic pool of storage that can be reallocated, scaled, and enlarged in real time, even while systems are accessing data
- Robust data integrity features such as advanced random array of inexpensive disks (RAID) functionality and built-in file system checksums that help protect against potential disk drive failures and disk errors.

SnapDrive is independent of the underlying storage access media and protocol. The iSCSI protocol provides storage access when the storage controller and host server are joined using Gigabit Ethernet. The FC protocol facilitates storage access through an FC host bus adapter (HBA) and SAN. The functionality and features intrinsic to SnapDrive are identical regardless of the underlying storage access protocol. This is because SnapDrive software uses either of the two methods to access virtual disks, which are created and stored on IBM N series storage systems. Thus a virtual disk can be created and accessed using the FC or iSCSI access protocols. Virtual disks are referred to as LUNs when accessed over the FC and iSCSI protocols. It is possible to change the protocol from FC to iSCSI and vice versa without migrating the data.

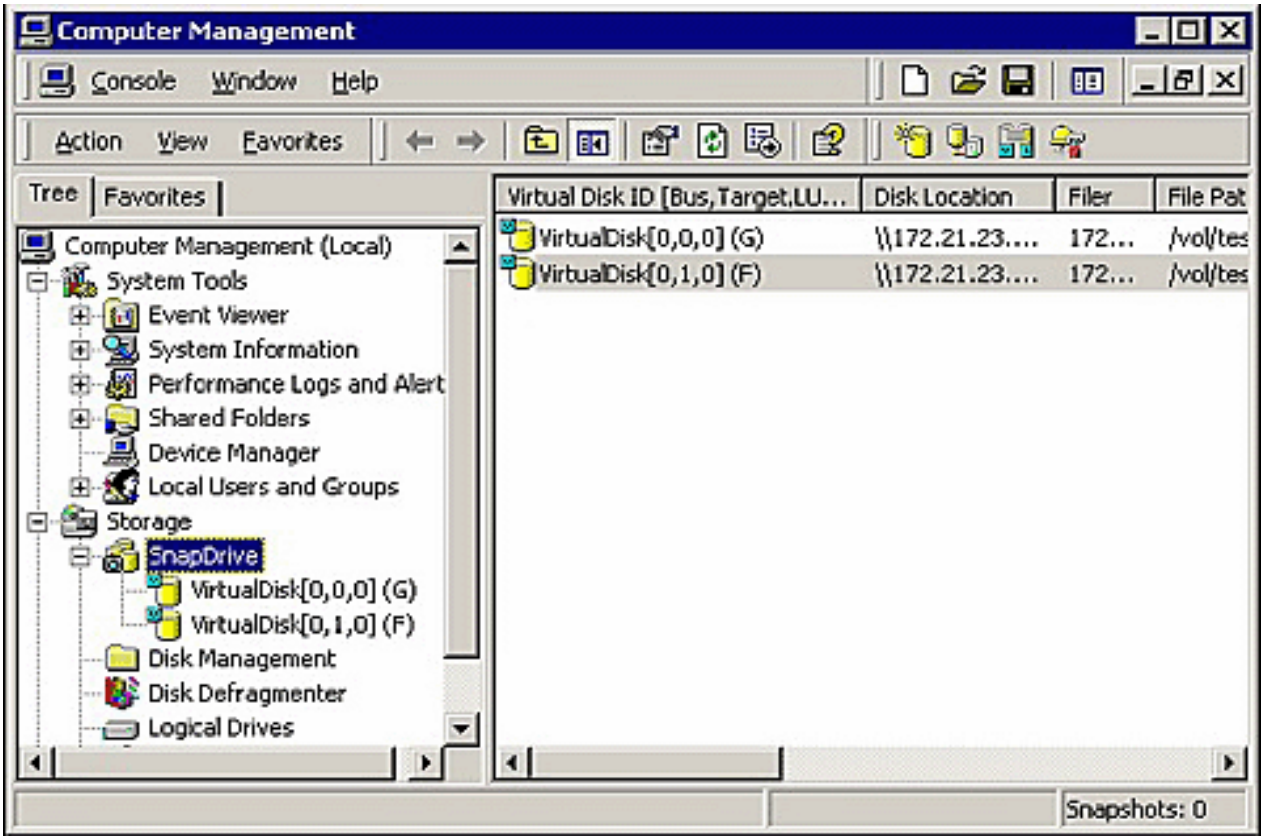

Figure 2. Storage management with SnapDrive.

Storage layout

IBM N series recommends using a single aggregate to store all SAP system data. This aggregate should be configured with IBM System Storage N series with RAID-DP™ (RAID double parity). Using one large aggregate provides the performance of all available disk spindles in the aggregate. All FlexVol volumes configured within this aggregate can benefit from all disk spindles in the aggregate. Each SAP system uses two FlexVol volumes. One FlexVol volume is exclusively used for the Oracle data files. The other FlexVol volume holds all the other SAP and Oracle file systems. Storing the database data files in a separate FlexVol volume is important to allow usage of Snapshot, IBM System Storage N series with SnapRestore®, FlexClone, and other IBM System Storage N series with Data ONTAP® features that work on a volume level.

FlexVol saplog		FlexVol sapdata	
LUN for executables	\usr\sap\SID	***LUN for database files***	\oracle\SID\sapdata1
	\usr\sap\trans		\oracle\SID\sapdata2
	oracle\SID\<version>		…….
LUN for logs	\oracle\SID\origlogA		\oracle\SID\sapdataN
	\oracle\SID\origlogB		
	\oracle\SID\mirrlogA		
	\oracle\SID\mirrlogB		
	\oracle\SID\oraarch		
	\oracle\SID\sapreorg		
	\oracle\SID\sapbacku		
	\oracle\SID\saparch		

Table 1. FlexVol storage (A).

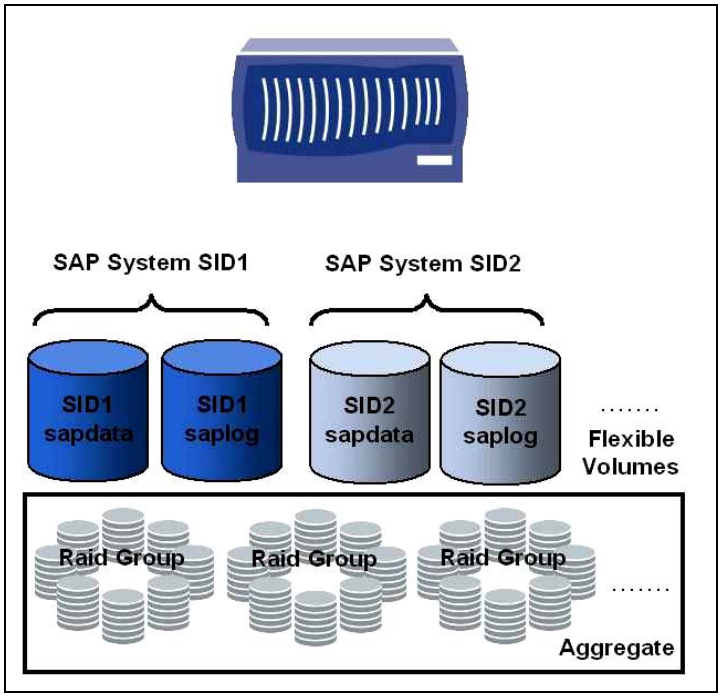

Figure 3. Storage layout with two FlexVol volumes.

RAID-DP offers a high level of data protection. The reliability of RAID-DP is far better than RAID5 and very close to RAID1. Data loss occurs only if three disks within the same raid group fail at the same time.

If this level of data protection is not sufficient, a second aggregate can be configured and the Oracle mirrored log files can be separated from the archived log files and the online redo logs stored in the second aggregate. Using this layout, it is always possible to recover the database without data loss if one of the two aggregates is lost. With this storage setup, it is recommended to distribute the “sapdata” volumes from different SAP systems equally to both aggregates. Of course, each aggregate has only half the disk spindles compared to the above setup and therefore this setup offers less efficient space and spindle utilization.

FlexVol saplog		FlexVol sapdata		FlexVol mirrlogs	
LUN for executables	\usr\sap\SID	***LUN for database files***	\oracle\SID\sapdata1	***LUN for mirrlogs***	\oracle\SID\mirrlogA
	\usr\sap\trans		\oracle\SID\sapdata2		\oracle\SID\mirrlogB
	oracle\SID\<version>		…….		
LUN for logs	\oracle\SID\origlogA		\oracle\SID\sapdataN		
	\oracle\SID\origlogB				
	\oracle\SID\oraarch				
	\oracle\SID\sapreorg				
	\oracle\SID\sapbackup				
	\oracle\SID\saparch				

Table 2. FlexVol storage (B).

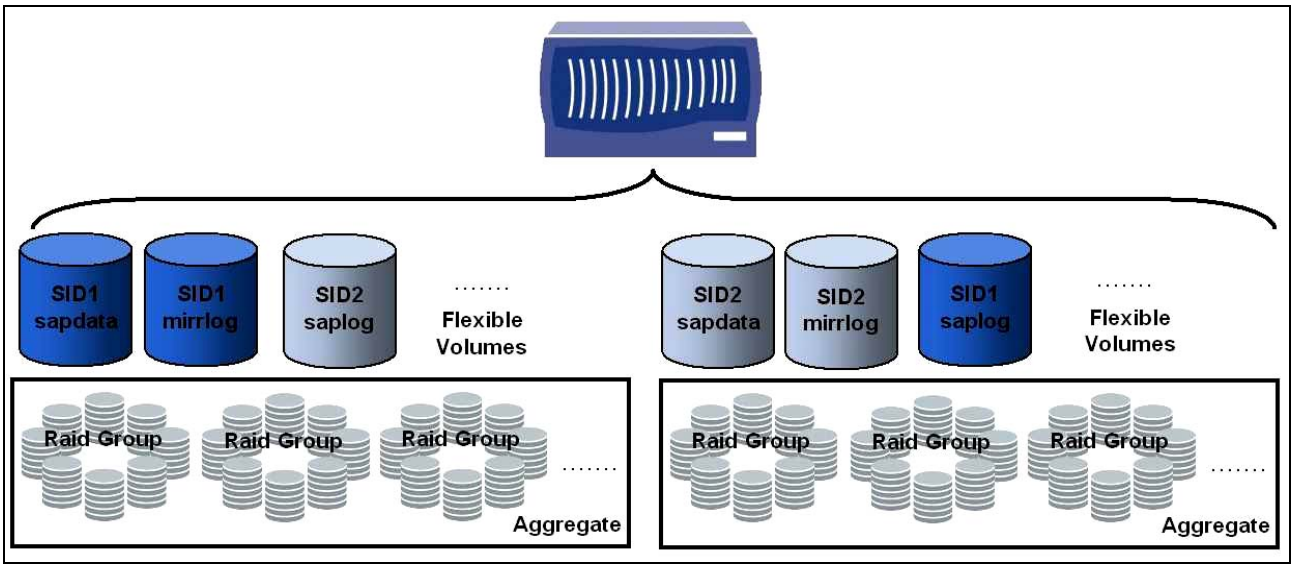

Figure 4. Storage layout with three FlexVol volumes.

In a cluster configuration, the layout should be based on the most efficient usage of resources. Because both storage controllers are active parts of the cluster, the SAP systems should be distributed to both storage controllers to use the available central processing unit (CPU) power and disk spindles equally.

It is therefore recommended to distribute the data of each SAP system among each single aggregate on each storage controller. In this case, the Oracle log files and the archives are stored in the aggregate on

the second storage controllers, while the data files and the mirrored log files are stored in the aggregate on the first storage controller.

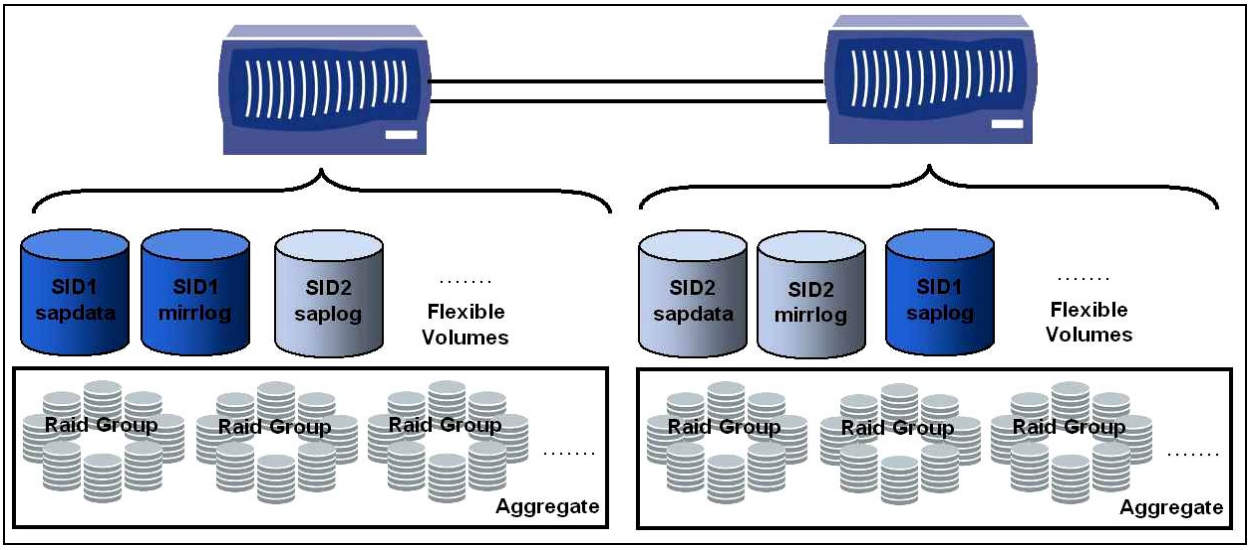

Figure 5. Storage layout with IBM N series cluster.

IBM System Storage N series with MetroCluster and synchronous mirroring work on the aggregate level. If all SAP systems must be mirrored synchronously, the layouts for a MetroCluster and a normal cluster are the same.

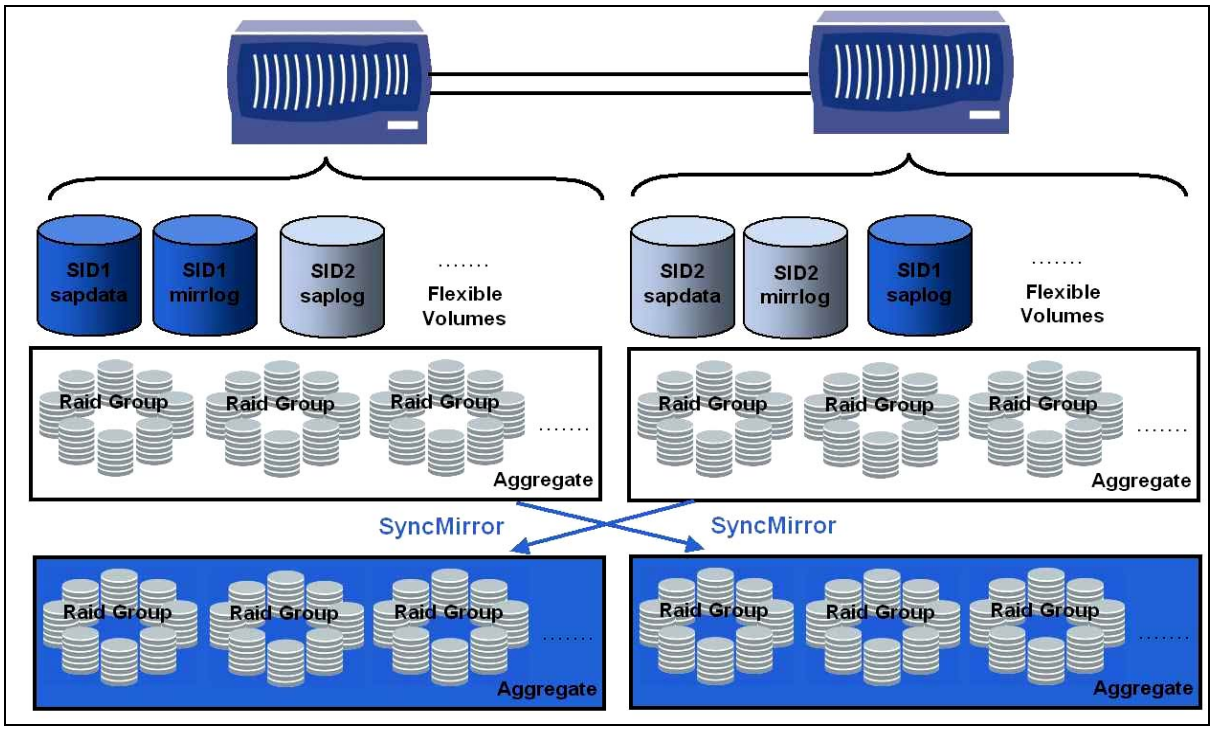

Figure 6. Storage layout with IBM N series MetroCluster.

Additional aggregates are necessary if only parts of the landscape require synchronous mirroring. For instance, production SAP systems require synchronous mirroring, but the test and development systems do not need to be mirrored.

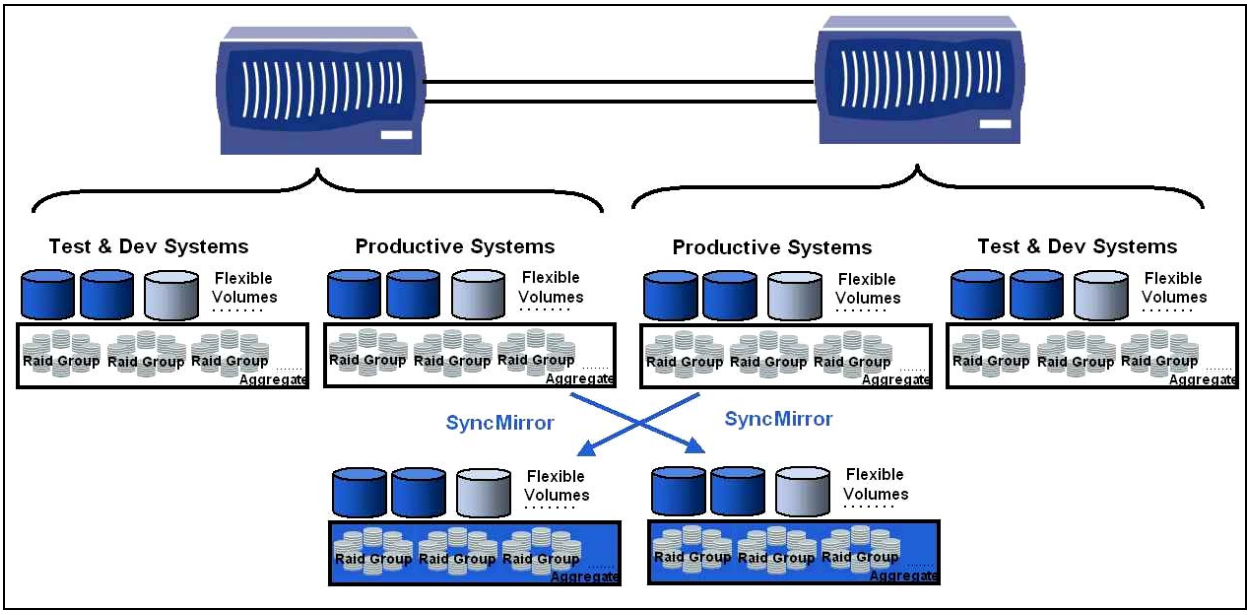

Figure 7. Storage layout with IBM N series with MetroCluster.

Sizing

This section gives an overview of the storage sizing for a SAP environment using IBM N series storage. The goal is to provide a basic understanding of what information is important to perform a storage sizing and how these requirements will influence the storage landscape.

Storage sizing for a SAP landscape is based on several conditions defined by customer requirements. All these requirements together will define the needed storage infrastructure:

- I/O requirements
- Capacity requirements
- Backup and recovery requirements (mean time to recover, backup window, retention policy)
- Cloning requirements (LUN clones or full copies)
- DR requirements (synchronous or asynchronous mirroring)
- High-availability requirements (storage system clustering).

Satisfying the I/O requirements is critical since overall SAP system performance is directly affected.

For existing SAP systems the I/O requirements need to be measured using database tools such as the Oracle STATSPACK utility, the SAP database performance monitor, or operating system tools such as the Windows performance monitor if the measurement is done on the operating system level. No matter which tools are used, it is very important that the measurement is done during peak loads of the SAP system. Especially when database tools are used for the measurement, a suitable time frame must be chosen, such as one hour, since these tools calculate an average value and the I/O sizing must be based on peak values.

For new SAP systems, where an I/O measurement is not possible, the SAP values for the systems, which are provided by the SAP Quick Sizer, can be used to estimate the I/O requirements. The storage sizing is much more accurate if I/O values are measured.

The load generated by asynchronous or synchronous mirroring should be added to the above I/O requirements. Also, the backup load must be added if the backup happens in a high activity phase of the system.

The type and amount of disk spindles and storage controllers are based on the I/O requirements.

To determine the needed capacity, the following information must be available:

- Size of each database
- Growth rate
- Number and retention policy of snapshot copies
- Number and durability of LUN clones
- Synchronous or asynchronous mirroring.

The type and amount of disks and the storage controller supporting the capacity are based on the capacity requirements.

The results of the I/O sizing and the capacity sizing are compared in a last step to define the right storage system for supporting both the I/O and capacity requirements.

Installation

This section describes the requirements and the configuration for installing a mySAP Business Suite or SAP NetWeaver system with Oracle database on a Windows Server or a Windows cluster (Microsoft Cluster Server, MSCS) using SnapDrive. Additional requirements and information can be found in the appropriate SnapDrive installation/administration guide and the corresponding SAP installation guide.

A common question asked by almost everyone in the SAP community is, “Is your solution SAP certified?” For Windows systems only, SAP commissions the German company AddOn to perform such certifications worldwide. They conduct quality and stability tests to analyze how hardware components behave in conjunction with the Microsoft platform, SAP products, and the database system. Upon successful completion, a formal certification is issued to the server partner. It is important to note that there is no SAP certification program for storage! The AddOn certification is based on the server systemboard, chipset, CPU, and I/O controller. Storage and some other standard components are deemed non-critical components under this program. SAP and AddOn make the hardware vendors of non-critical components responsible for the proper operation of their equipment.

General requirements

SnapDrive needs to be installed and configured on the server(s) as described in the appropriate SnapDrive installation/administration guide.

If iSCSI is used to attach the server(s) to the IBM N series storage, a dedicated storage local area network (LAN) needs to be set up. This network should be used exclusively for the storage traffic and not for any other services like user access or MSCS cluster interconnect. Each server therefore needs a dedicated Gigabit Ethernet network interface card (NIC) or iSCSI HBA to be connected to the IBM N series storage.

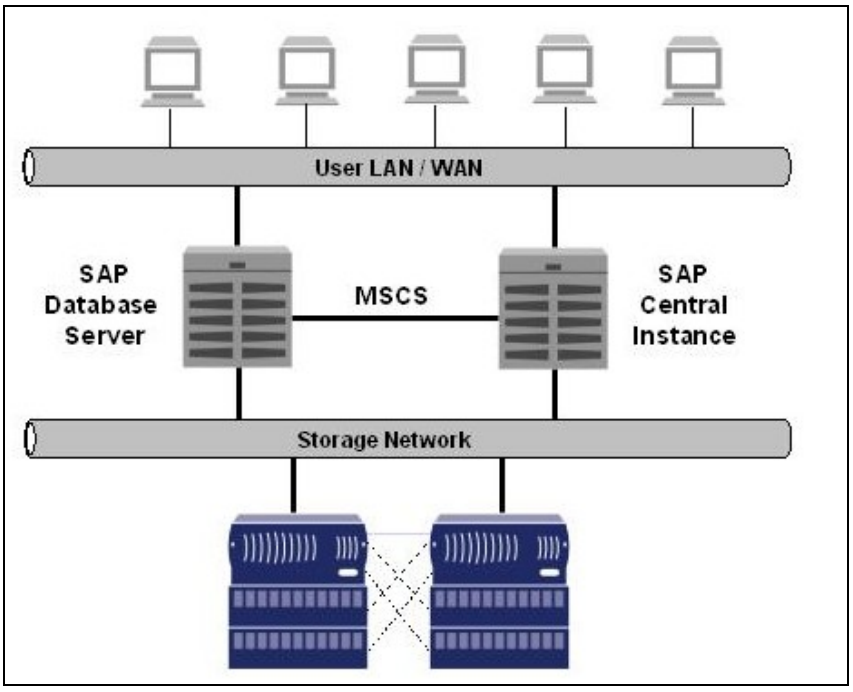

Figure 8. Infrastructure overview.

Installation in a non-MSCS environment

For a mySAP Business Suite or NetWeaver system installation on Windows or Oracle, at least three dedicated LUNs are needed:

- One LUN for the Oracle database files
- One LUN for the Oracle online logs, mirror logs, archive logs, database control files, and the directories `saptrace, sapreorg, sapcheck`, and `sapbackup`
- One LUN for the SAP and Oracle executables.

These LUNs are distributed to two FlexVol volumes in one single aggregate. The LUNs for the database files are stored in one FlexVol volume and the two other LUNs in the second FlexVol volume, as shown in Figure 9.

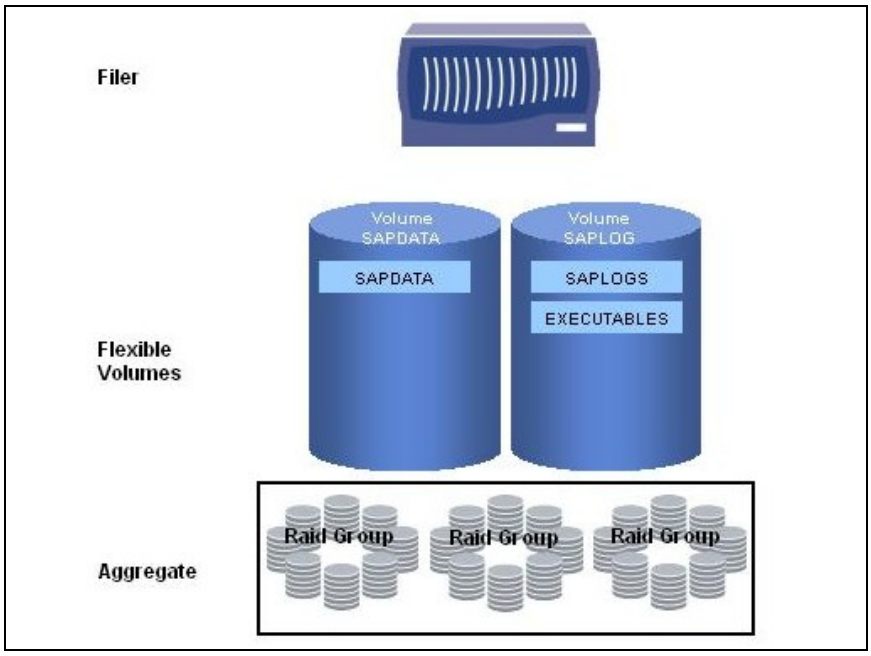

Figure 9. Installation with two FlexVol volumes.

If the level of data protection with a single aggregate is not sufficient, a second aggregate can be configured to separate the mirror logs from the redo logs and the archives in a second aggregate. With this configuration, the LUNs containing the mirror logs are placed in a separate FlexVol volume in the same aggregate as the data files. The archive logs and the redo logs are placed in a FlexVol volume in a separate aggregate.

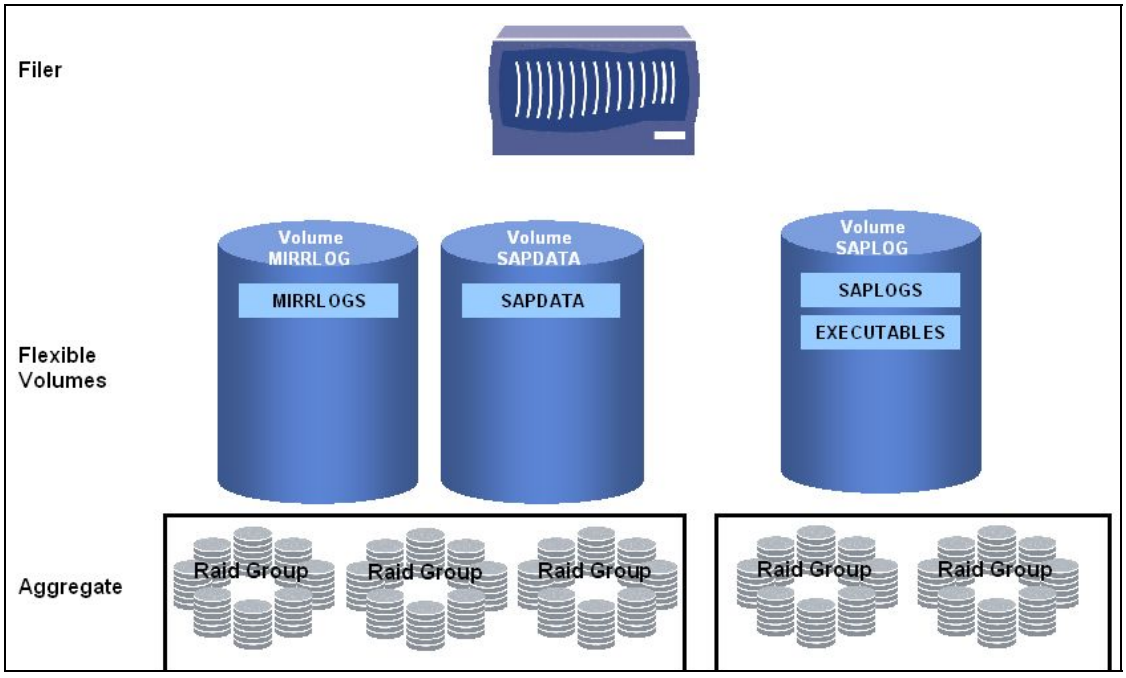

Figure 10. Installation with three FlexVol volumes.

If a clustered storage system is used, the configuration is identical to the above approach. One aggregate is configured on each storage controller and the LUNs are distributed as described above.

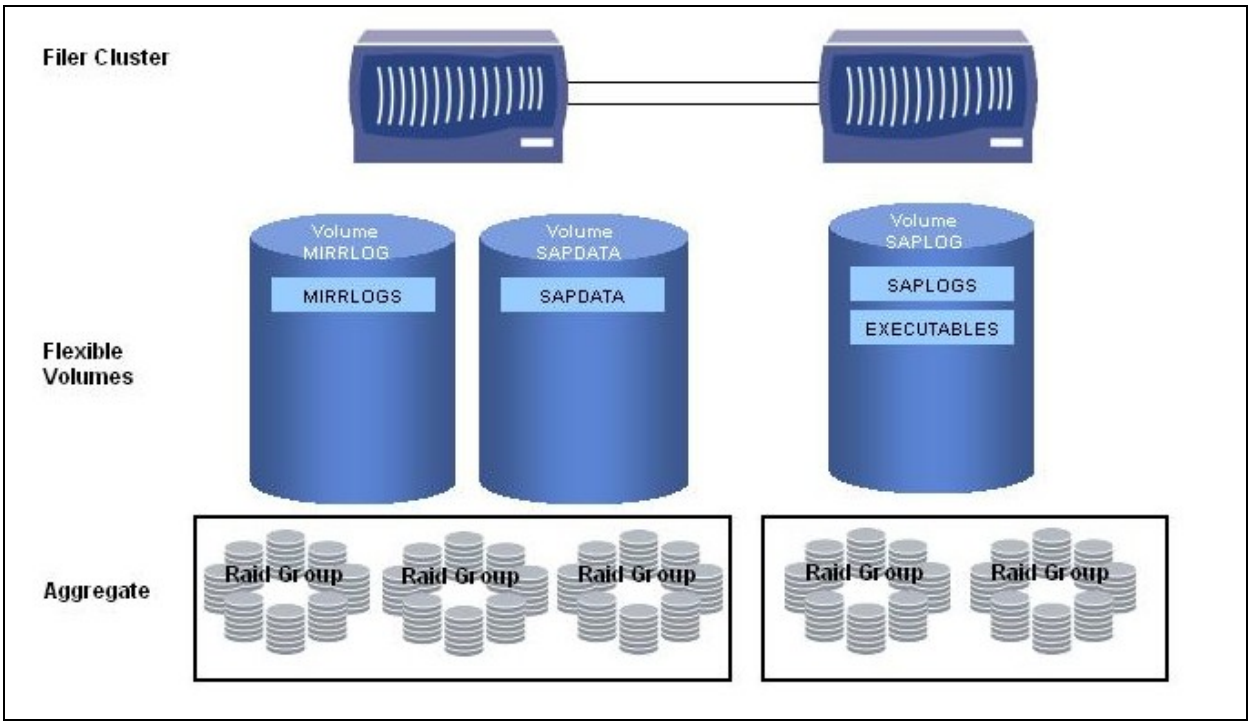

Figure 11) Installation with three FlexVol volumes in a storage system cluster.

The SAP installation is accomplished as described in the corresponding SAP installation guide.

Installation in an MSCS environment

For a mySAP Business Suite or NetWeaver system installation on Windows MSCS with Oracle, at least four shared LUNs are needed:

- One shared LUN for the Oracle database files
- One shared LUN for the Oracle online logs, mirror logs, archive logs, database control files, and the directories `saptrace`, `sapreorg`, `sapcheck`, and `sapbackup`
- One shared LUN for the SAP and Oracle executables
- One shared LUN for the Cluster Quorum disk.

IBM N series recommends using a clustered storage system when the application is clustered with an MSCS.

One aggregate is configured per storage controller. The mirror logs are separated from the redo logs and archives by storing them in different aggregates. The cluster quorum LUN is placed in the same FlexVol volume as the executables, the redo logs, and the archives.

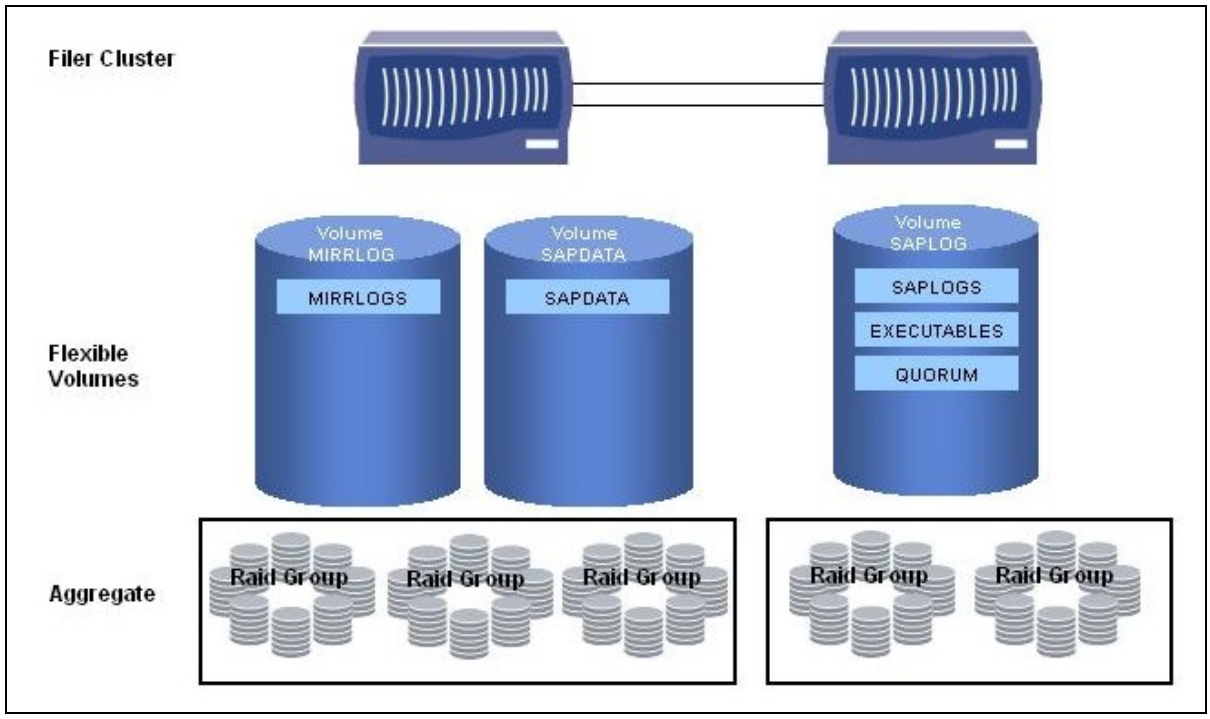

Figure 12. Installation in an MSCS environment.

Configure the cluster using the previously created LUNs as shared disks. The SAP installation is accomplished as described in the corresponding SAP installation guide.

Storage migration

In this section, different storage migration approaches are discussed. If a migration includes a change of the operating system or database system, the migration can't be done on the storage level, and the SAP migration tools need to be used. These tools export the data from the source environment and import the data into the target environment. The approach is therefore defined by SAP and is independent of the storage system used.

Overview of migration approaches

The decision about which migration approach fits best in a specific environment heavily depends on the acceptable downtime of the business application. Furthermore, the downtime depends on the amount of data that needs to be migrated. In general there are three approaches to storage migration of the SAP data:

- Migration on the operating system level
- Migration on the database level
- Migration on the storage level.

Migration on the operating system level

In addition to the existing storage system, the IBM N series storage system is connected to the database server. The IBM N series storage system is configured and the LUNs are mounted to the server. Before the data migration is started, the database and the SAP system must be shut down. The data will then be copied via the server from the old storage system to the IBM N series system. When all data is copied, the old storage system is disconnected from the database server. If the file system structure (drive letters) remains the same, the database can be started immediately. If there is a change in the file system structure, the new structure needs to be configured within Oracle by creating a new control file.

A migration on the operating system level can be done for an FC-to-FC or an FC-to-iSCSI protocol migration. The disadvantage of this approach is that the SAP system won't be available while the database files are copied. Depending on the database size, the downtime could be several hours.

Migration on the database level

An online or offline database backup is restored to the IBM N series storage system. To minimize the impact on the production SAP system, the restore can be done using a separate server connected to the IBM N series storage. In addition, the archive logs are continuously copied to the separate server. Before the final migration is started, the SAP database and the SAP system must be shut down. The IBM N series storage will then be connected to the database server and the LUNs are mounted to the server. The online logs, the control files, and the archive logs that have not yet been copied will now be copied from the old storage system to the IBM N series storage. When all data is copied, the old storage system is disconnected from the database server. If the file system structure (drive letters) remains the same, the database can be started immediately. If there is a change in the file system structure, the new structure needs to be configured within Oracle by creating a new control file. Finally, a forward recovery of the database is carried out.

A migration on the operating system level can be done for an FC-to-FC or an FC-to-iSCSI protocol migration. This approach will reduce downtime during the migration but will need an additional server during the migration process.

Migration on the storage level

IBM N series uses a fabric-attached appliance, which hosts data copying software. The appliance is attached to the existing SAN as well as to the IBM N series storage system. The appliance needs to be configured to see the existing and new arrays, and then must be enabled to perform the migration. The appliance is configured to migrate the appropriate data sets and/or data files. The software will automatically transfer the data from the source to the replacement (IBM N series) system. The transfer is done with no intervention from the hosts (servers) attached to the arrays. Impact on system performance is also minimal. Once the migration is done, the servers and storage infrastructure need to be reconfigured to permit the servers to see the storage on the new IBM N series storage controller. Only after all of the servers can access the new storage, can the migration appliance be pulled out.

A migration on the storage system level can be done only for an FC-to-FC protocol migration. This approach will minimize the downtime during the migration but will need the above-described appliance for the migration process.

The table below gives a summary of the migration processes.

	Supported Protocols	Downtime	Additional Hardware
Migration on the operating system level	FCP-to-FCP and FCP-to-iSCSI	**High** During reconfigure and whole copy process	None
Migration on the database system level	FCP-to-FCP and FCP-to-iSCSI	**Medium** During reconfigure and forward recovery	Server temporarily connected to IBM N series storage
Migration on the storage level	FCP-to-FCP	**Low** During reconfigure	Migration appliance

Table 3. Migration process summary.

System management and maintenance

SAP system cloning

Business challenges

A typical SAP customer environment consists of different mySAP Business Suite and SAP NetWeaver components. To be able to test application patches, run performance and data integrity tests, or provide user training environments, copies of SAP components are required. A typical SAP customer needs an average of 10 copies of different SAP components. These copies need to be refreshed, often on a weekly or monthly basis.

Performing a SAP system copy normally takes several days and negatively affects the production environment. Many manual steps are also performed, consuming valuable IT staff time.

The source database needs to be exported using SAP tools and imported at the target system, or an offline backup of the source database is restored at the target system. Depending on the database size, these steps have a significant impact on the application availability. It takes many hours to replicate a 1TB database from the source to the target system. Preparing the cloned system so that it can be used in the new environment takes several additional hours. This preparation is often done manually, which consumes SAP Basis administrator time.

The ability to quickly create a SAP system copy on demand is more important for these reasons:

- QA systems need to be refreshed on a weekly basis
- Additional test systems need to be set up quickly to perform specific integration tests
- Quick setup of test system with current production data is needed for an SAP upgrade project
- Training systems need to be set up or resynchronized.

The traditional approach to creating the system copies does not address these demands. SAP copies also consume a lot of storage, which needs to be provided. Since these copies are typically clones of the production system, the amount of storage needed can be huge.

IBM N series solution

The IBM N series solution for SAP system cloning addresses these issues by providing a fully automated process to create a SAP system copy on demand, in a few minutes, without any impact on the source production system. In addition, IBM N series cloning functionality makes managing storage very efficient by storing only data changes between the parent and the clone.

IBM N series solution for SAP system copies

SAP system copies are accomplished by cloning LUNs. A LUN clone is a writeable point-in-time image of an IBM N series LUN. A LUN clone is based on a snapshot copy of the source LUN and is created in a few seconds without interrupting the operation on the source system. LUN clones store only changed blocks between the source LUN and the clone LUN and therefore significantly decrease the amount of disk space needed for SAP system copies.

Figure 13 shows the basic concept of the system copy solution. Creating a SAP system copy consists of several steps on the source system and several steps on the destination system.

On the source system, a database-consistent snapshot copy of the Oracle data files is taken during online operation and has no performance impact on the source system. This step can therefore be carried out at any time, ignoring production online operation on the source system.

On the target system, this snapshot copy is the base for the LUN clone. The creation of the LUN clone takes only a few seconds. The LUN clone is then connected at the target system. The subsequent steps at the target system are the steps that are necessary to change the database and the SAP system ID (SID). In addition, SAP-specific post-processing tasks need to be done.

All these steps are fully automated. A SAP system copy can be accomplished in a few minutes using the IBM N series solution.

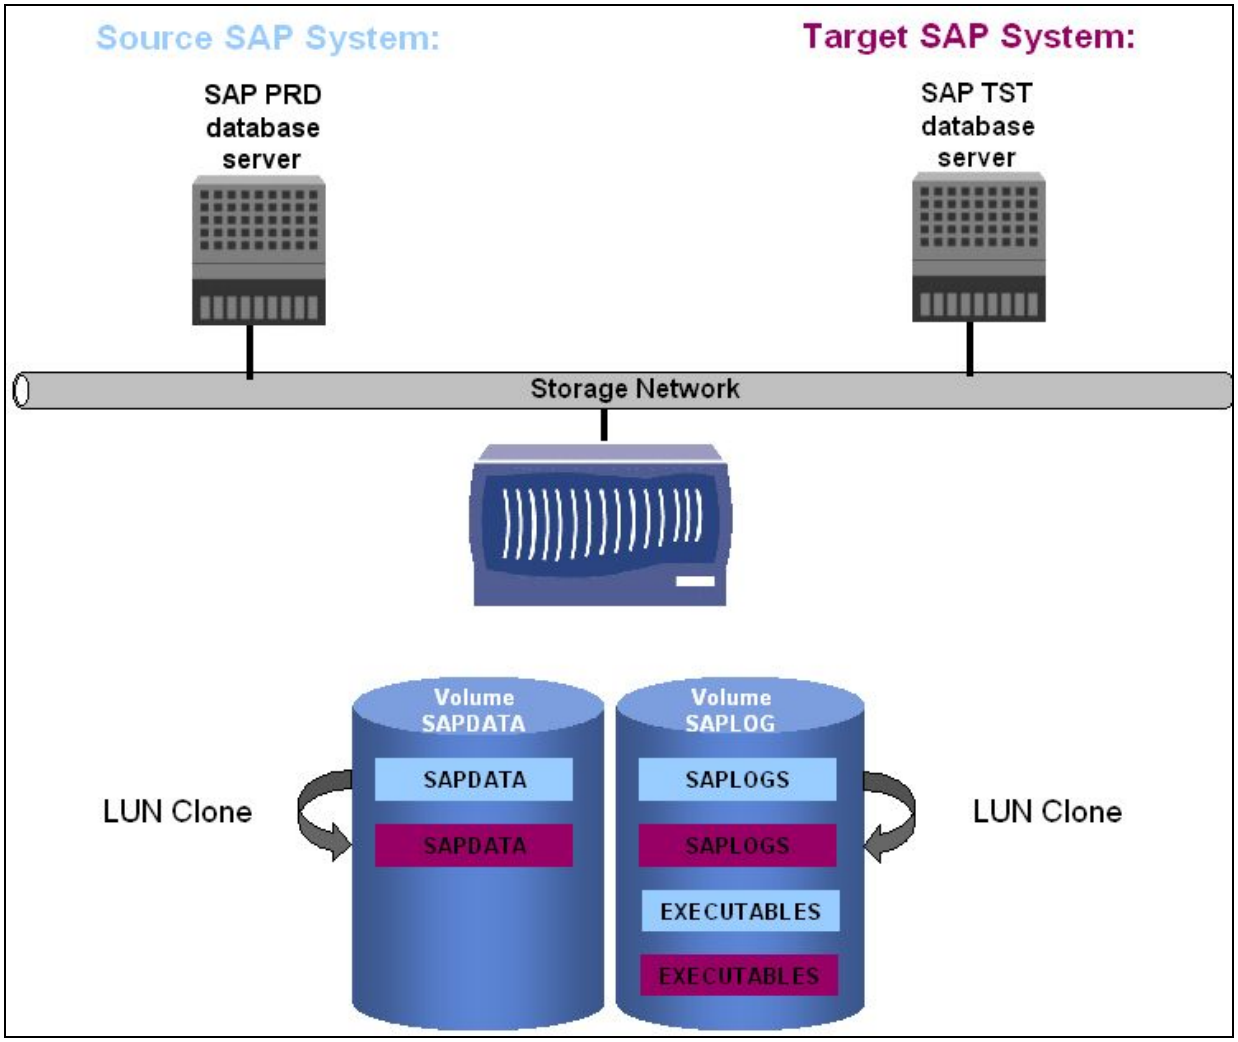

Figure 13. SAP system cloning overview.

The table below provides a comparison of the traditional approach compared to the IBM N series approach to perform an SAP system copy.

Necessary steps at the source system

The traditional approach requires the creation of an online or offline backup of the source database. Typically, a backup significantly impacts performance on the source system and, thus, is difficult to schedule. Depending on the database size, the backup of the database will take several hours. The subsequent steps are typically carried out manually, consuming IT staff time.

With the IBM N series approach, the backup is done using Snapshot functionality. Creating a snapshot copy will take only a few seconds and has no performance impact on the source system. Therefore this step can be scheduled at any time during online operation. The creation of the snapshot copy and all the subsequent steps are fully automated.

Necessary Steps at the Target System

If a new SAP test system needs to be set up, the SAP and Oracle software need to be installed once. This step needs to be carried out with both approaches. With all subsequent refreshes of this system, this step is not necessary any more.

With the traditional approach, the next step is to restore the offline or online backup from the source system. Depending on the database size, this step may take several hours. Scheduling the restore might also be difficult, since the restore blocks the backup infrastructure. The following steps to adapt the file system and the database to the new SID and the SAP post-processing tasks are typically carried out manually, consuming IT staff time.

With the IBM N series approach, a LUN clone is created based on the consistent Snapshot technology database backup that was done at the source system. The creation of the LUN clone takes only a few seconds and can be scheduled at any time. The following steps to adapt the file system and the database to the new SID and the SAP post-processing tasks are fully automated.

Traditional Approach	IBM N series Approach	Advantages
Necessary Steps at the Source System		
• Offline or online backup	• Snapshot backup during online operation	• No impact on production operation with IBM N series solution; schedule any time
• Control file to trace	• Control file to trace	• Fully automated process with IBM N series
• Switch log files and copy archived logs to shared location (with online backup)	• Switch log files and create a LUN clone of the LUN where the archive logs are stored	• Fully automated process with IBM N series
Necessary Steps at the Target System		
• Install SAP system (if not existing yet)	• Install SAP system (if not existing yet)	• Same approach
• Restore offline backup from source system	• Create LUN clone based on snapshot backup; connect LUN clone at target system	• Restore takes only seconds with LUN clone technology; fully automated process with IBM N series
• Adapt directory names to new SID	• Adapt directory names to new SID	• Fully automated process with IBM N series
• Create new control file with new SID based on control file trace from source system	• Create new control file with new SID based on control file trace from source system	• Fully automated process with IBM N series
• Adjust Oracle security	• Adjust Oracle security	
• SAP-specific postprocessing tasks	• SAP-specific postprocessing tasks	• Fully automated process with IBM N series
		• Fully automated process with IBM N series

Table 4. SAP system copy approaches.

Conclusion

- The IBM N series system copy solution significantly improves the process to create SAP system copies.
- A system copy can be accomplished within several minutes, compared to several days with the traditional approach.
- System copies can be scheduled at any time, since there is no impact on the online operation of the source system (performance, backup infrastructure).
- Snapshot and LUN clone functionality reduces the time necessary to copy the data from the source system to the target system from several hours to several seconds.
- All storage, operating system, database, and SAP-specific tasks are fully automated, minimizing the interaction of IT staff.
- Snapshot and LUN clone functionality significantly reduces the necessary disk space for a SAP system copy by storing only data changes between the source and the target system.

SAP upgrade

Business challenges

Existing SAP customers face pressure to upgrade to new SAP solutions because new technology and functionality are needed or because the existing release runs out of maintenance.

Upgrading to a new SAP release is a challenging project, which consumes a lot of resources, including IT staff.

Business processes are influenced during the upgrade project time period, since all development needs to be stopped and SAP support packages can't be imported. Therefore it is very important to minimize the overall time for the upgrade project.

In complex environments with large databases, a normal two-day weekend might not be sufficient to run the upgrade of the production SAP system. Every hour that can be saved while running the upgrade is important. Database backups consume a lot of time. Optimizing backup and restore functionality is therefore very critical.

During a SAP upgrade project, SAP basis administrators need to create several system copies to run the upgrade with current data from the development or production SAP system. The creation of a SAP system copy normally takes several days and negatively affects the production environment. In addition, a lot of steps are performed manually, which consumes valuable IT staff time.

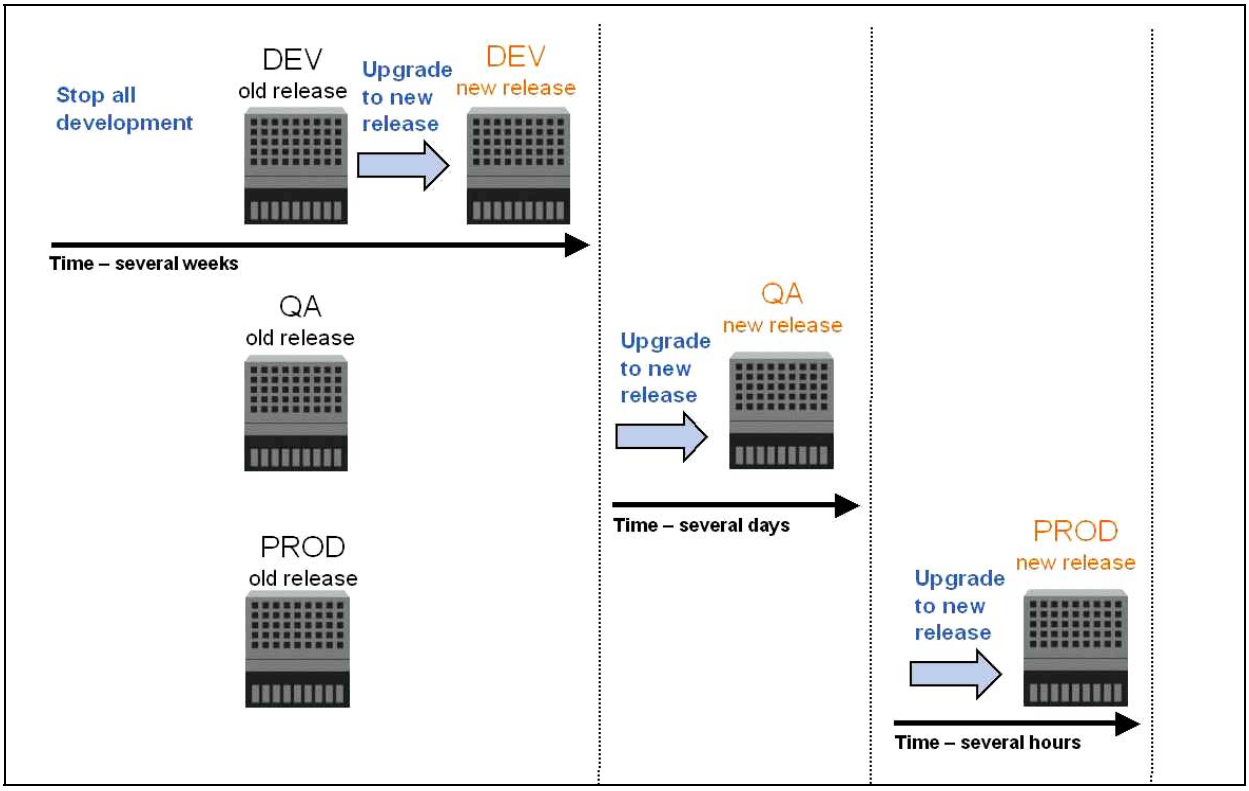

Figure 14. SAP upgrade overview.

IBM N series solution

The IBM N series solution for SAP upgrades addresses the issues above by providing a solution for creating fully automated SAP system copies in a few minutes. The IBM N series backup and recovery solution helps minimize the downtime during all upgrade phases, with the capability to create database backups and restore databases in seconds. The IBM N series solutions help minimize the risk by reducing the downtime and overload of IT staff resources during a SAP upgrade project.

IBM N series solution for SAP upgrades

Upgrading the development system

The upgrade of the development system is usually carried out on a separate development system running on separate hardware. During the upgrade process of the development system, the functionality of the upgrade is tested in the specific customer environment. In almost all cases, the upgrade of the development is carried out more than once to be able to define the necessary actions in all specific upgrade phases.

The setup of the separate SAP system is based on a system copy of the original development system. This system copy can be made using the IBM N series system copy solution. Reducing the time needed to create this system copy is critical, specifically because the copy is typically done several times. During the upgrade process and during the modification adjustment, snapshot backups are very helpful because they allow resetting the system to any snapshot copy and restarting the upgrade phase

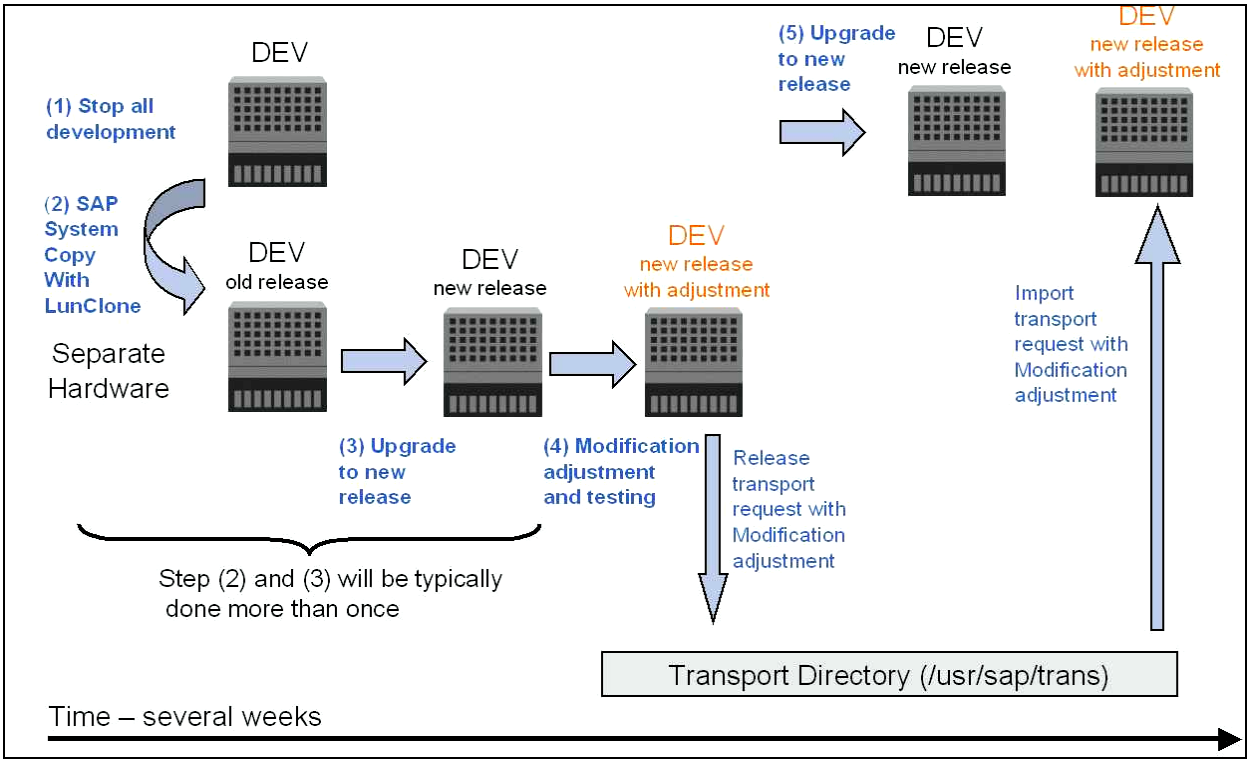

Figure 15. SAP upgrade–development system.

Upgrading the QA system

The upgrade of the QA system is done using a fresh system copy of the production SAP system. One important result of the upgrade of the quality system is the run-time of the upgrade with real production data. The IBM N series SAP system copy solution makes it easy to refresh the QA system. Reducing the necessary time to create this copy is also critical when upgrading the QA system, because the copy is typically done more than once. Snapshot backups are helpful during the upgrade process and before the modification adjustments are imported. These snapshot copies allow restoring the system to any specific snapshot copy and restarting an upgrade phase or the import.

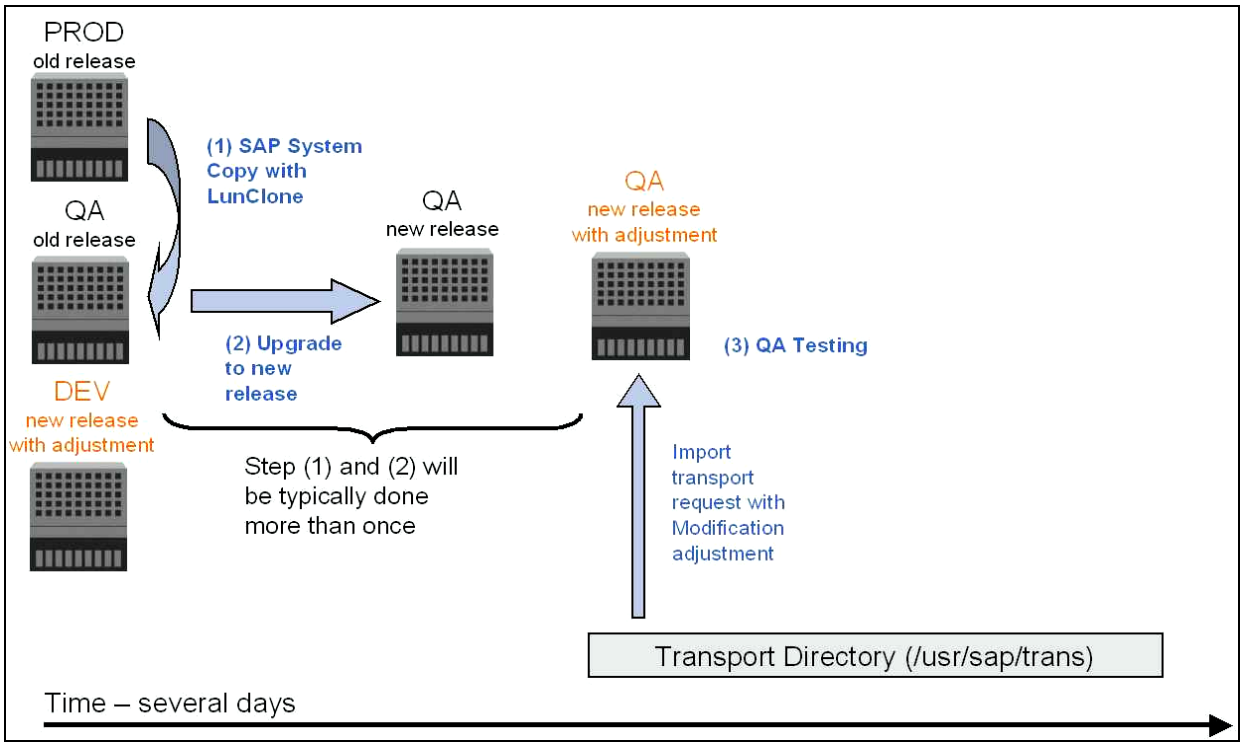

Figure 16. SAP upgrade–QA system.

Upgrading the production system

Scheduling is extremely important when the production system is going to be upgraded, since the system is not available at various stages during the upgrade. Scheduling must also allow 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.

Production system upgrade includes at least three backups of the database. The first backup needs to be done directly before the upgrade is started. After the upgrade is done, a second backup is required before the modification adjustments are imported. After importing the adjustments and finishing the functionality tests, a third backup is required. If functionality testing fails, the system needs to be restored to the old release level.

Using snapshot copies as a backup method and SnapRestore for restoring the system to its former release status increases flexibility with regard to scheduling. Normal tape backups will need several hours, which must be considered when planning the upgrade schedule. This time is reduced to several minutes when using Snapshot and SnapRestore features.

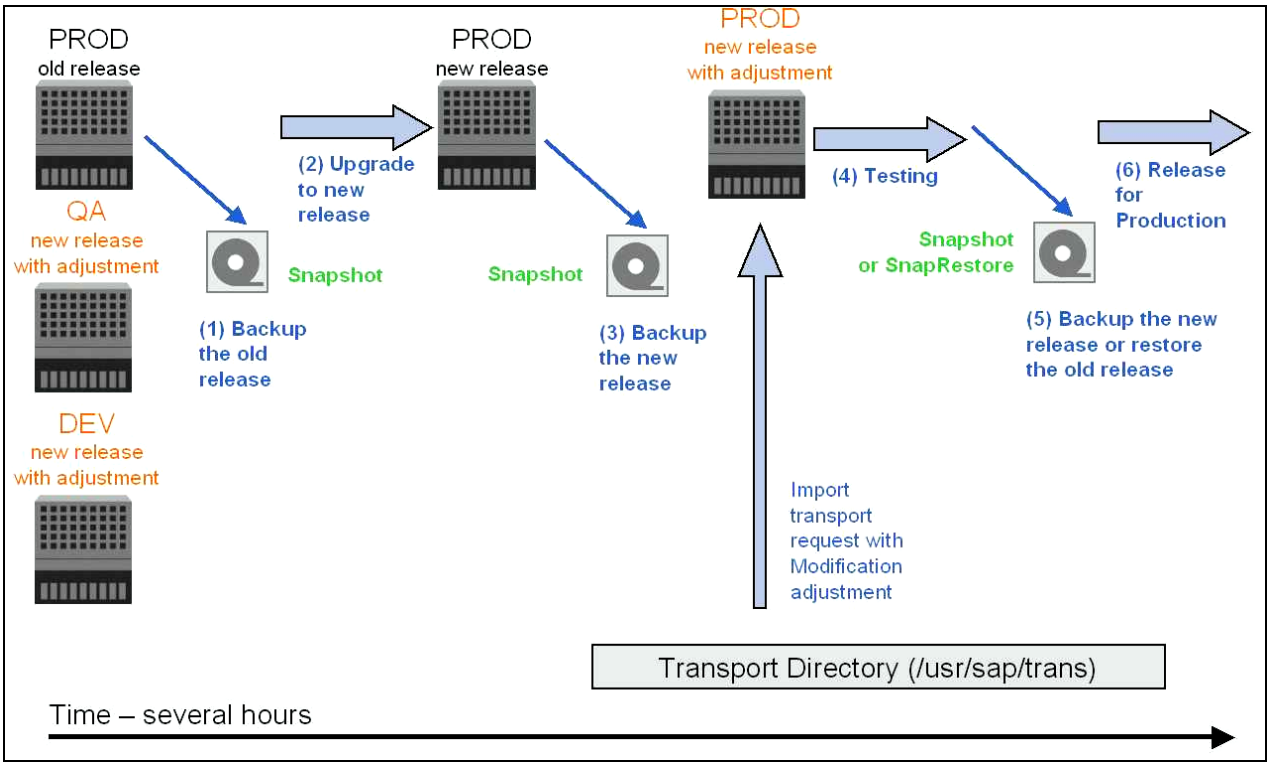

Figure 17. SAP upgrade–production system.

Conclusion

- The IBM N series system copy and backup and recovery solution significantly improves the SAP upgrade process.
- The IBM N series system copy solution makes refreshing the separate development and QA systems faster, reducing the time needed for backup and recovery from several days to several minutes.
- Using the IBM N series backup and recovery solution backs up the database with Snapshot, allowing the database to be restored to any specific snapshot copy in a few seconds and allowing immediate restart of any upgrade phase.
- The IBM N series backup and recovery solution significantly reduces the total upgrade time of the production system, providing more flexibility with regard to scheduling the SAP upgrade.

Business continuance

Backup and recovery

Business challenges

Corporations today require their SAP applications to be available 24 hours a day, 7 days a week. They expect performance to be maintained without respect to increasing data volumes, and they expect the system to undergo routine tasks such as backups without influencing the SAP system. Performing backups of the SAP databases is a critical task, because backups can have a significant performance impact on the production SAP system. Because 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 time during which backups can be performed with minimum influence on the business process. Downtime of SAP production and even development systems is critical, since it always has a financial impact on the business process. Thus the time needed for backup and recovery is of particular importance.

Summary of SAP backup and recovery challenges:

- **Performance impact on production SAP systems.** 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 a 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 result in more tape media or disk space for backups. Incremental backups can address these issues, but they 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 (MTTR) 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 always has 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 a SAP upgrade always includes at least three backups of the SAP database. The time needed to perform these backups will cut down the total time available for the upgrade process.

IBM N series solution

IBM N series provides unique functionalities that address the above challenges.

IBM N series with Snapshot creates 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 usage 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, a snapshot copy can 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 N series with SnapRestore functionality allows restoring of 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. Due to the fact that 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 from 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 are needed for a restore and recovery. Most restore and recovery actions are handled by using 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 happen.

A backup and recovery solution using IBM N series storage system always consists of two parts:

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

A backup to a second location is always based on snapshot copies created on primary storage. Thus, data is directly read from the primary storage system without generating load on the SAP database server.

There are several possible ways to do the backup to a second location:

- Disk-to-disk backup using a IBM System Storage N series with NearStore® feature or FAS system and IBM System Storage N series with SnapMirror® software

The primary storage directly communicates with the secondary storage and sends the backup data to the destination. SnapMirror 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 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

- **LAN backup.** The backup server mounts the snapshot copy of the data 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 storage controller; 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`.

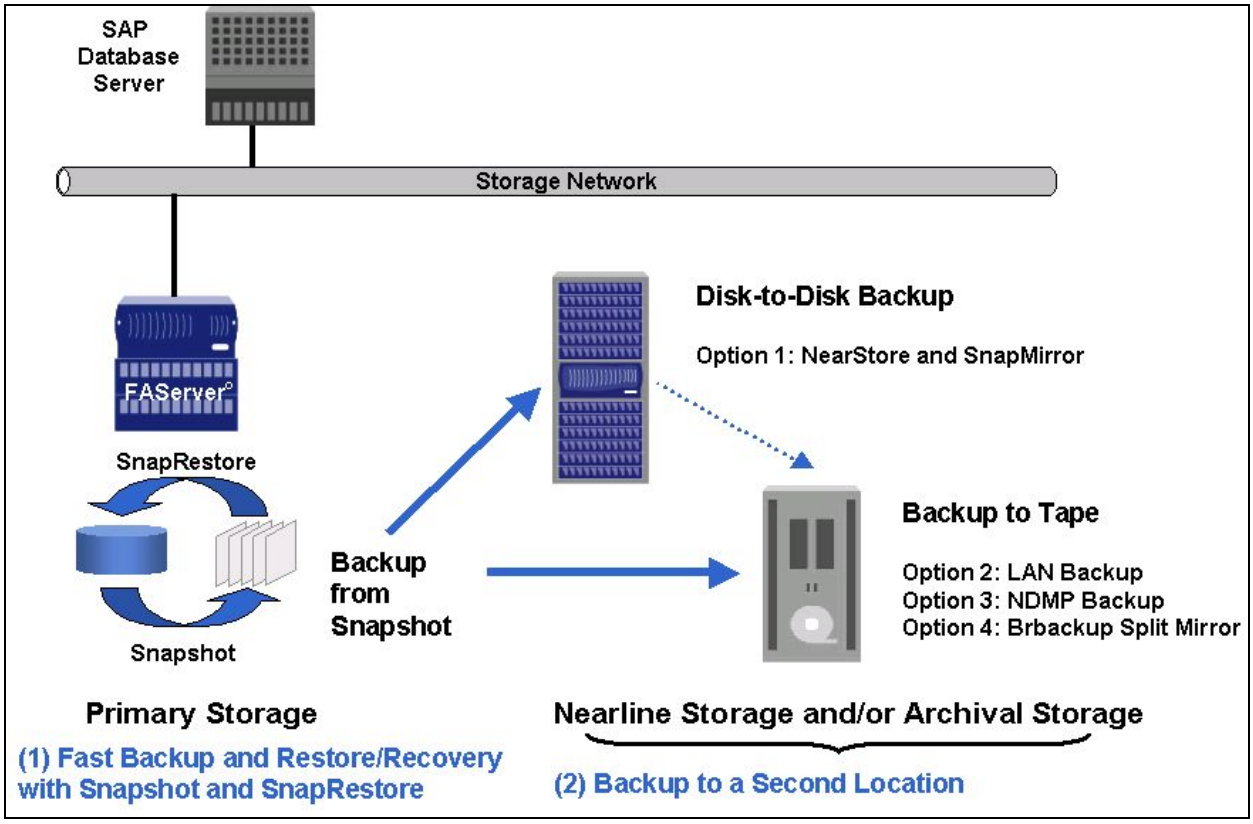

Figure 18. SAP backup and recovery solution overview.

Which method is best depends on a number of factors, including:

- 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 is 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 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 technology, the database is offline or in hot backup mode during the whole backup time. The backup and restore speed is slow. Full backups always need the full capacity on tape. The backup can be used for DR purposes, and it is a fully integrated solution 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 DR solution, 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

Again, Snapshot and SnapRestore have significant advantages in all areas compared to a conventional tape backup. However, 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 done 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 doing a backup with a separate backup server that mounts a writable 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 slow. Full backups will always need the full capacity on tape. The backup can be used for DR, but 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 storage controller writes the data directly to tape. Therefore, there is no load on the storage network, but there is a load on the primary storage. The backup and restore speed is higher than 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, the 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 the SAP backup and recovery management.

Tape backup on backup server with Brbackup Split Mirror

When executing `Brbackup` on a separate backup server, there is no load on the database server, but there will still be a high load on the primary storage and the storage network. The time needed for the backup and for restore and recovery is high. 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 SnapMirror

Because SnapMirror runs on the storage level, there is be no load on the database server. SnapMirror transfers only the changed blocks with each backup. This significantly reduces the load on the primary storage and the storage network. For the same reason, little time is needed to perform a full database backup. In addition, each full backup stores only the changed blocks at the destination. Therefore, the amount of disk space needed for a full backup is very low 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.

	Backup Load On			Time Needed					
	Database Server	Primary Storage	Storage Network	Database Offline or in Hot	For Backup	For Restore and	Space Needed for Backup	Backup Usable for Disaster	Integration within SAP
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 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 Brbackup Split Mirror	High	High	High	Very low	High	High	High	Yes	Yes
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 Nearline 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 MTTR
- Minimum disk space needed for database backups at the primary and secondary storage systems (FAS and NearStore systems).

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, 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, and 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 to repeat that phase.

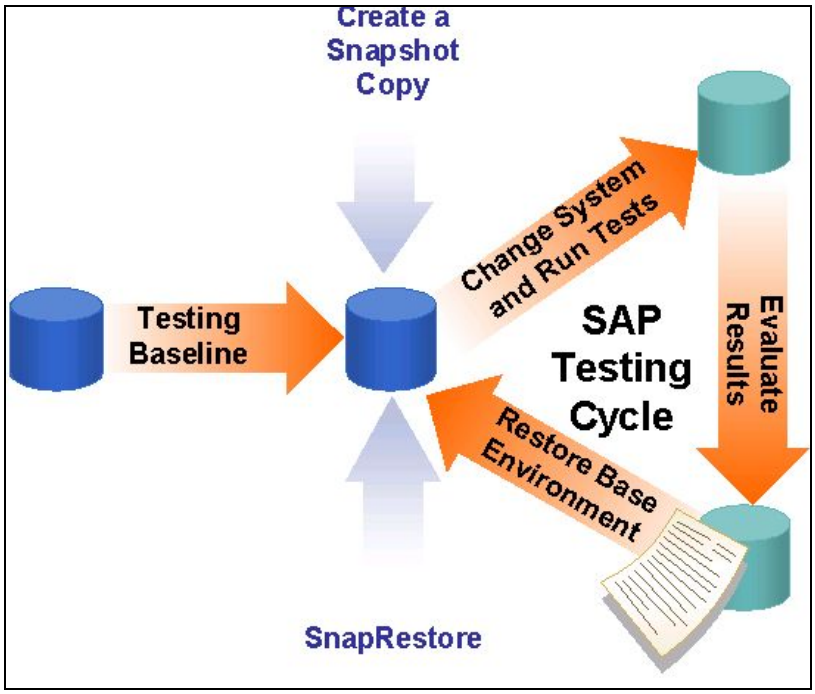

Figure 19. SAP testing cycle.

Carrying out a 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 made in the development system first 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 can save a considerable amount of time. A tape backup does not have to be made; 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. Snapshot as a backup method and SnapRestore for restoring the system to its former release status increase flexibility with regard to scheduling. By creating several snapshot copies at certain stages during 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.

High availability

Business challenges

Production SAP systems are business-critical applications that require 24x7 availability. Meeting these requirements requires an infrastructure without a single point of failure. SAP systems have two single points of failure that require a high-availability solution. The database server and central instance must be available.

IBM N series solution

IBM N series Clustered Failover delivers a robust and highly available data service for business-critical environments. Installed on a pair of IBM N series storage controllers, IBM N series Clustered Failover ensures data availability by transferring the data service of an unavailable storage controller to the other storage controller in the cluster.

IBM N series solution for SAP high availability

Figure 20 shows a sample clustered failover configuration. A cluster can be created with two storage controllers by connecting the storage controllers with a cluster interconnect. This connection is redundant and is used to exchange cluster heartbeats and synchronize the non-volatile random access memory (NVRAM) on both storage controllers. The disk shelves of the cluster partner are connected to the second storage controller through a second FC loop. If the first storage controller fails, the second storage controller handles its disk shelves. The MAC and IP addresses and the world-wide port name (WWPN) of the first storage controller are also adopted. Since the NVRAM is mirrored on both storage controllers through the cluster interconnect, no data is lost.

Because both storage controllers can be active in a cluster configuration, it is possible to use a single cluster to provide high availability for both the central instance and the database server. It is also possible to support other systems on the cluster.

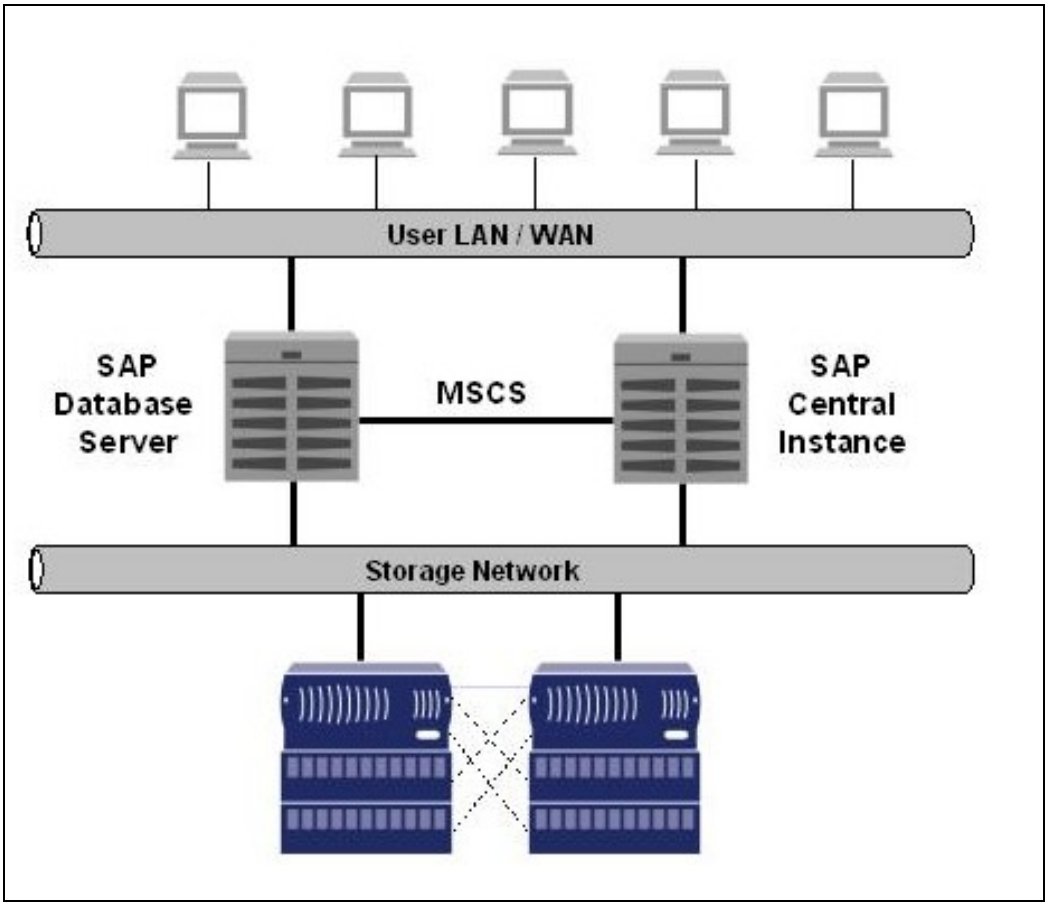

Figure 20. IBM N series clustered storage system solution.

Conclusion

- The IBM N series Clustered Failover technology provides an extremely robust high-availability solution.
- A cluster has an availability level of nearly 100%.
- Both storage controllers in the cluster can be used actively, providing high availability for both the database server and central instance.
- A clustered storage system is recommended if server clustering is used for the application.

DR (disaster recovery)

Business challenges

Organizations recognize the importance of a bulletproof business continuance plan in place to deal with a disaster. The cost of not having one—lost productivity, revenue, and customer loyalty, and possibly even business failure—makes it mandatory to have a plan that ensures an absolute minimum of downtime and rapid recovery from a disaster, with little or no loss of data. IBM N series offers several solutions that can be configured to meet your corporation's specific recovery point objective (RPO) and recovery time objective (RTO). Working with your corporation's business users to determine the acceptable values for RPO and RTO will guide you in your selection of a DR solution that uses one or more IBM N series products.

IBM N Series solution

SnapMirror

IBM N series with SnapMirror software delivers the DR solution that today's global SAP systems need. By replicating data at high speeds over a LAN or wide area network (WAN), SnapMirror software provides the highest possible data availability and fastest recovery.

SnapMirror technology mirrors data to one or more storage controllers. It updates the mirrored data to keep it current and available for DR, tape backup, read-only data distribution, testing, online data migration, and more.

SnapMirror performs an initial Level 0 transfer to initialize the DR site. After the initial transfer, incremental changes are asynchronously passed to the recovery site. The amount of data lost in the event of a disaster depends on the frequency of the incremental asynchronous transfers. The SnapMirror DR solution is based on the IBM N series backup and recovery solution. Every snapshot backup is mirrored to the DR site. SnapDrive automatically performs a SnapMirror update when a snapshot copy is taken. The LUN where the archive logs are stored also has to be mirrored using SnapMirror. SnapDrive offers the possibility to start a SnapMirror update by executing one command. IBM N series recommends a frequent SnapMirror update of the archive logs, e.g., every 10 minutes, to ensure a minimum of data loss.

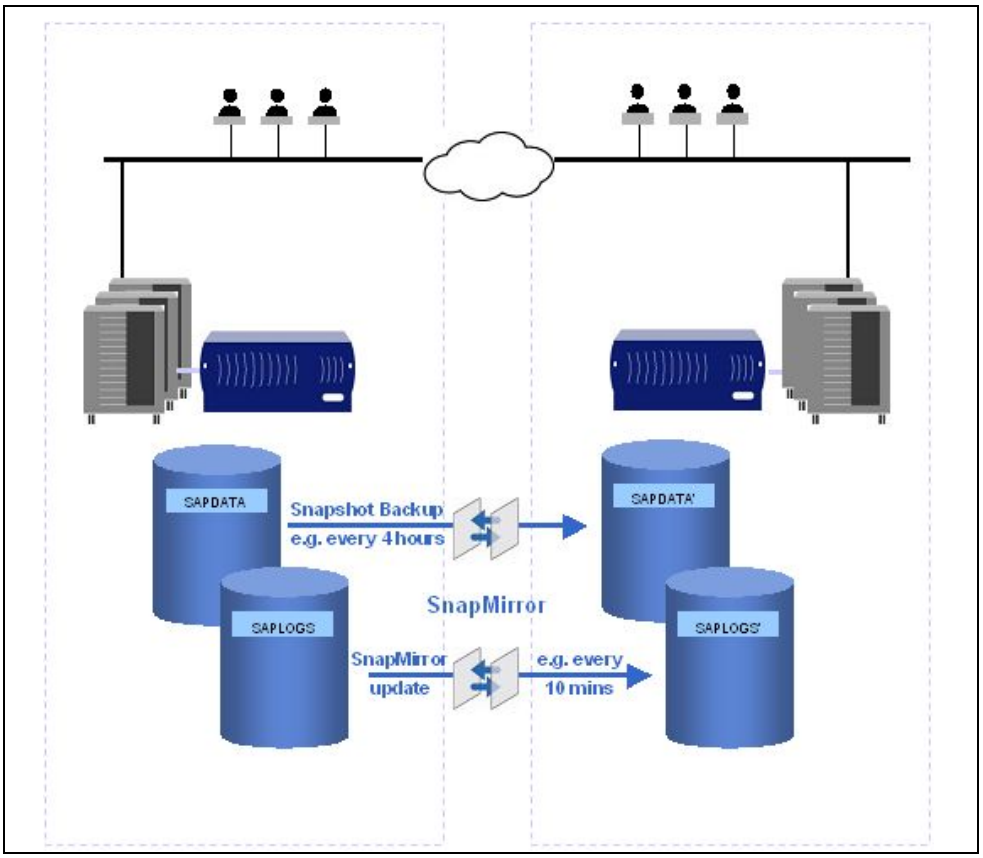

Figure 21. DR with SnapMirror.

MetroCluster

IBM N series with MetroCluster is an integrated high-availability and business continuance solution that provides DR with no data loss. MetroCluster extends failover capability from within a data center to a site located many miles away. It also replicates data from the primary site to the remote site to ensure that the data there is completely current. The combination of failover and data replication ensures that you can recover from disaster—with no loss of data—in minutes rather than hours or days. MetroCluster is much like IBM N series Clustered Failover but with the added benefit of DR. Clustered Failover creates a cluster of IBM N series storage appliances in one location with access to both sets of disks. MetroCluster extends this cluster configuration to locations up to 30 kilometers away. Because there is no physical connection to the cluster appliance's disk in case of a site failure, MetroCluster requires the use of IBM System Storage N series with SyncMirror® to ensure that both storage controllers in the cluster have copies of the other storage controller's data.

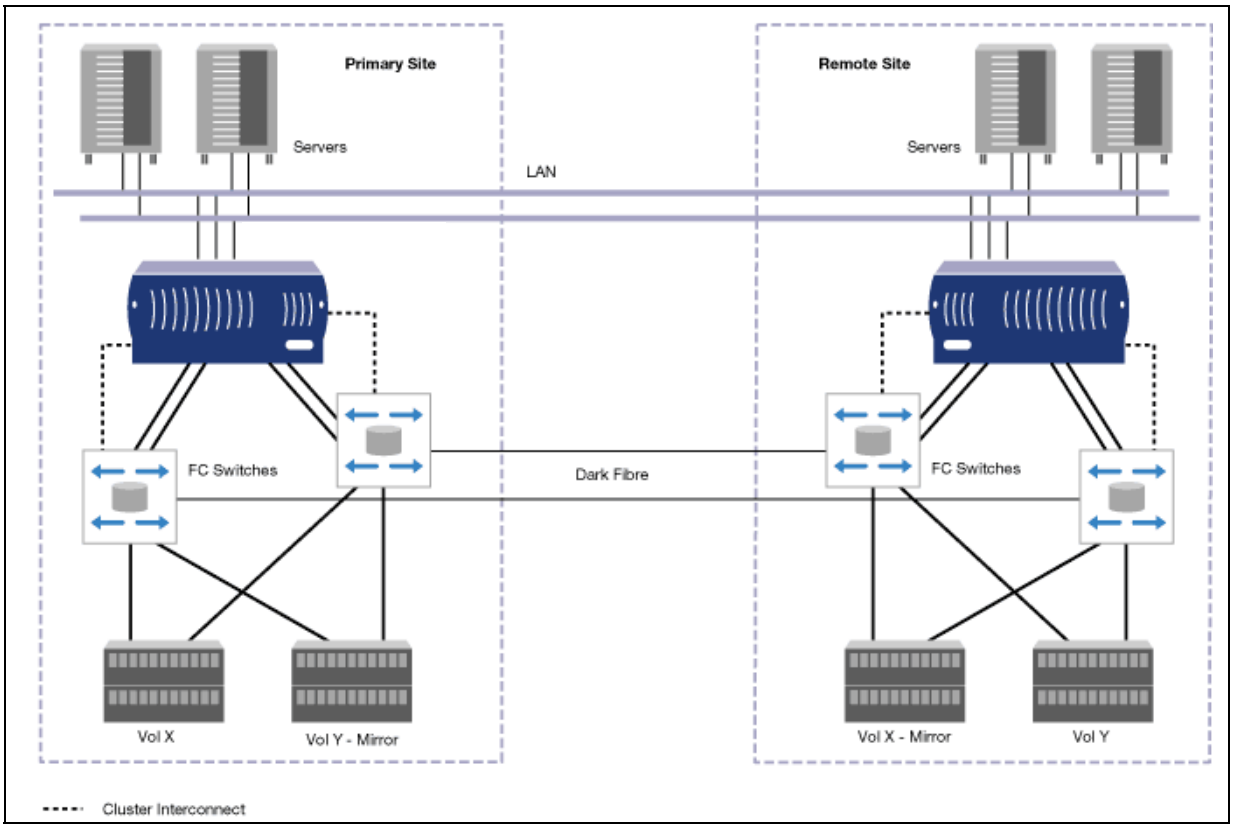

Figure 22. MetroCluster over direct FC switch connection.

This solution provides high availability and disaster protection in a campus environment.

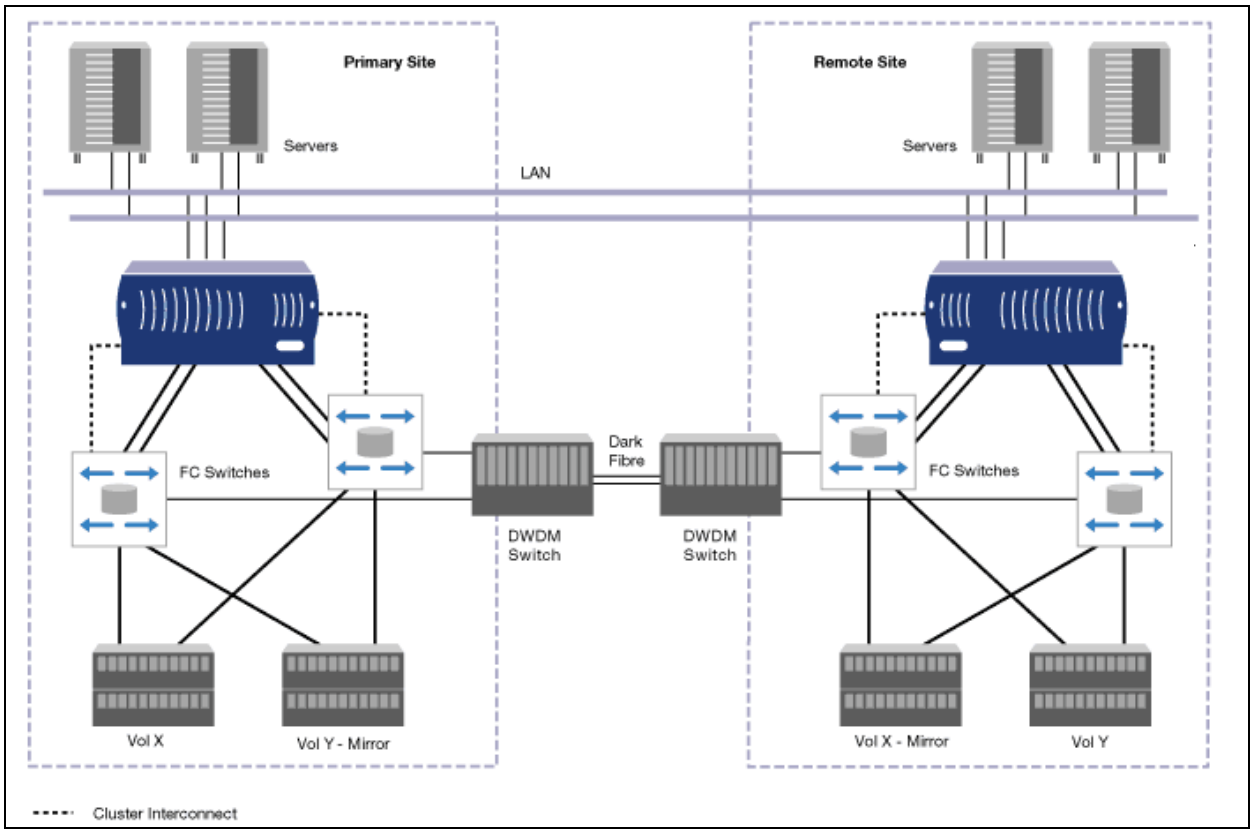

Figure 23. MetroCluster over FC and DWDM switch infrastructure.

This solution connects distant sites in metropolitan areas

Conclusion

- IBM N series has multiple DR solutions to support different business and financial requirements.
- SnapMirror provides an efficient and cost-effective DR solution.
- MetroCluster enables DR in a high-availability cluster configuration with no data loss.

Archiving and compliance

Business challenges

Archiving

The long-term accumulation of data in the SAP database can ultimately affect the performance and availability of SAP applications. To keep your SAP systems and applications running at peak efficiency, it is vital to implement a data archiving process to enhance availability while reducing performance and management overhead.

Simply deleting this data is often not an option, because read access to individual data objects may still be required. For this reason, the data must be relocated from the database in such a way that it is secure and can still be accessed when the need arises.

Choosing media type and platform for archival storage requires companies to conform to not just one, but many content retention mandates. IT organizations have to respond by analyzing the business requirement and then choosing the proper solution based on factors such as time to data, risk, storage scalability, compatibility, and TCO. Current WORM (write once, read many) technologies like WORM optical disk and WORM tape do not provide sufficiently rapid access, high reliability, or low TCO. What organizations need is a solution that easily and inexpensively integrates archived storage with corporate applications and enables them to comply with existing retention and security regulations for reference data.

Compliance

In addition to managing system size and performance, SAP customers are keenly aware of increasing industry regulations that have introduced significant financial penalties for failing to comply with retention, indexing, auditing, privacy, and reporting requirements. These regulations span practically all public companies and industry sectors. Nearly every major corporation must put a regulatory compliance solution in place or face the risk of being exposed to litigation and fines. In most cases this solution will require the purchase of new storage subsystem hardware and software.

Historically, most regulated data has been stored on optical devices, tape, paper, or microfiche/microfilm. According to the Enterprise Storage Group (ESG), only about 10% of regulated data is stored on disk today. Disk storage has seldom been used due to a number of factors that include cost and the lack of necessity to retrieve information quickly. However, ESG estimates that disk will be the fastest growing area for the storage of regulated data.

IBM N series solution

SAP data archiving is based on the Archive Development Kit (ADK), where the archiving objects are used to remove from the database data that is no longer needed in online business processes and to store it in such a way that it is still accessible in the future. The purpose of XML-based data archiving is the same as ADK-based archiving. The key difference is that it is based on universally accepted and widely used standards: XML is used to archive business objects, hyper text transfer protocol (HTTP) as a communication service, and Web-based Distributed Authoring and Versioning (WebDAV) as a general protocol for connecting storage systems.

The ADK is the software layer that encapsulates the technical aspects of data archiving programs. ADK provides an application programming interface, also used by SAP, that customers and partners can use to develop their own archiving solutions. ArchiveLink is an interface as well as a service for facilitating the process-driven management of business documents. Business-related documents can be linked to and retrieved from application objects via workflow.

WebDAV is a set of extensions to the HTTP protocol that allow users to collaboratively edit and manage files on remote Web servers. The major features of the protocol include locking, metadata management, and namespace manipulation.

Once archive files have been created, the data marked for archive can be deleted from the source system. The archived data can then be transferred directly from the primary storage system to an external content or archive server. IBM N series solutions for SAP archiving such as NearStore and IBM System Storage N series with SnapLock® work hand-in-hand with technologies from SAP and their archiving partners. The result of effective SAP archiving is better-performing applications that cost less to operate and manage.

IBM N series with NearStore is the preferred compliance and archive storage subsystem for SAP landscapes. The NearStore product family uses IBM N series Data ONTAP 7G technology and takes full advantage of value-added software from IBM N series such as SnapLock. NearStore scales from 7 to 504TB by using more economical ATA disk technology. With nearly 100% field-measurable uptime, IBM N series RAID-DP technology enables NearStore systems to tolerate single disk failures with no data loss. Should additional capacity or performance be required for any reason, IBM N series FAS systems can be substituted for IBM N series with NearStore within the SAP storage landscape.

IBM N series with SnapLock is the IBM implementation of high-performance, disk-based magnetic WORM storage. SnapLock provides secure, storage-enforced data retention functionality through open file protocols such as common internet file system (CIFS) and network file system (NFS) while using existing IBM N series technologies to the greatest degree possible. This implementation also includes significant efforts in hardening Data ONTAP and its administrative interfaces to the degree that SnapLock can be deployed for protecting data in regulatory environments so strict that even the storage administrator is considered an untrusted party. An example of such an environment is the broker-dealer market regulated by SEC 240.17a-4. Alternate configurations of SnapLock can be deployed for unregulated or more flexible regulated environments.

SnapLock provides special-purpose volumes in which files can be stored and committed to a nonerasable, nonrewritable state either forever or for a designated retention period. SnapLock allows this retention to be performed at the granularity of individual files through standard open file protocols such as CIFS and NFS. The retention of these files is enforced by Data ONTAP, which controls all access to the physical media and acts as the gate through which all file protocol or administrative access to the data must pass.

SnapLock is based on open file protocol interfaces and does not require the use of any kind of proprietary application programming interface (API). You can perform all SnapLock-specific operations, such as setting file retention periods and committing files to WORM state, through regular file system operations available on all clients. Applications use the regular programmatic library interfaces they would use for file operations on any other kind of storage system.

SnapLock is available in two versions. One or the other of these versions can be implemented within Data ONTAP, but not both.

SAP customers who have chosen compliance and archiving solutions from iXOS, such as iXOS-eCONserver, or from FileNet, such as the P8 platform, can take full advantage of the products' integration with SAP and IBM N series with SnapLock.

SnapLock Compliance enables organizations to satisfy strict records-retention regulations such as SEC Rule 17a-4 (broker-dealers), HIPAA (healthcare), Sarbanes-Oxley (public companies), 21CFR Part 11 (life sciences), and DOD 5015.2 (government). Only an act of willful destruction, such as physically removing disks from an IBM N series system, can result in record deletion or alteration prior to the specified retention date.

SnapLock Enterprise enables adherence to rigorous organizational best practices through functionality similar to that of SnapLock Compliance, but allows administrators to delete entire SnapLock Enterprise volumes. Under no circumstances is it possible for any SnapLock Enterprise user or administrator to delete or modify individual SnapLock Enterprise WORM records or to undermine SnapLock Compliance WORM volumes. SnapLock is supported on IBM N series FAS and NearStore platforms.

Conclusion

- IBM N series provides a flexible, scalable, and secure solution for SAP compliance and data archiving needs.
- SnapLock enables locking of some files without forcing WORM behavior for all data.
- There is no risk of software vendor lock-in. IBM N series works well with existing document and content management packages such as iXOS-eCONserver and the FileNet P8 platform.
- Data can be managed and backed up using the customer's current products and strategies.
- The solution can incorporate existing IBM N series or other vendor's storage.
- The solution improves ROI and lowers TCO through increased availability, enhanced system performance, lower administration overhead, and increased staff productivity.
- Compliance and archived data remains easily accessible on NearStore—a more cost-effective alternative for archiving SAP data than adding database storage or processing power.

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, RAID-DP, FlexClone, FlexVol, SnapLock, SnapRestore, SnapMirror, SyncMirror, NearStore and SnapDrive are trademarks or registered trademarks of Network Appliance, Inc., in the U.S. and other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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.