Technical report:

mySAP on UNIX and Oracle on IBM System Storage N series

Best Practices

Document NS3533-0

August 22, 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 when running in a UNIX and FCP environment and using an Oracle database.

Introduction

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 UNIX® fibre channel (FC) protocol environment using an Oracle database. Primary consideration is 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 are specific neither to any given SAP application nor to the size and scope of the 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 the 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 ERP landscapes based on the SAP NetWeaver™ technology platform. SAP 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 becomes possible; it becomes a requirement for corporations to keep pace with change.

IT Challenges

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

SAP has long encouraged customers to maintain separate development and quality assurance 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 sandbox and training. Driven by standard processes for development and testing within a corporation, it is also not uncommon to have multiple development instances as well as more than one system used for quality assurance, testing, or perhaps 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 these instances have differing performance, scalability, availability, and uptime profiles. These profiles can also fluctuate depending on the phases of a project implementation and whether the project is focused on an existing SAP implementation or a brand new one.

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

If the IT infrastructure supporting SAP applications is inflexible or is difficult to operate or manage, or if high cost of ownership barriers develop within IT, that can negatively affect the ability of business owners to deploy new and improved business 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 proven interoperable across all platforms. IBM N series fabric-attached storage (FAS) systems integrate easily into a complex enterprise and simultaneously support network attached storage (NAS), FC storage area network (SAN), and IP SAN (iSCSI) protocols.

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 means that operations can be simplified, utilization and efficiency can be increased, and changes can be applied more quickly and seamlessly. IBM N series storage solutions enable customers to add storage when and where they need it, without disruption and at the lowest incremental cost.

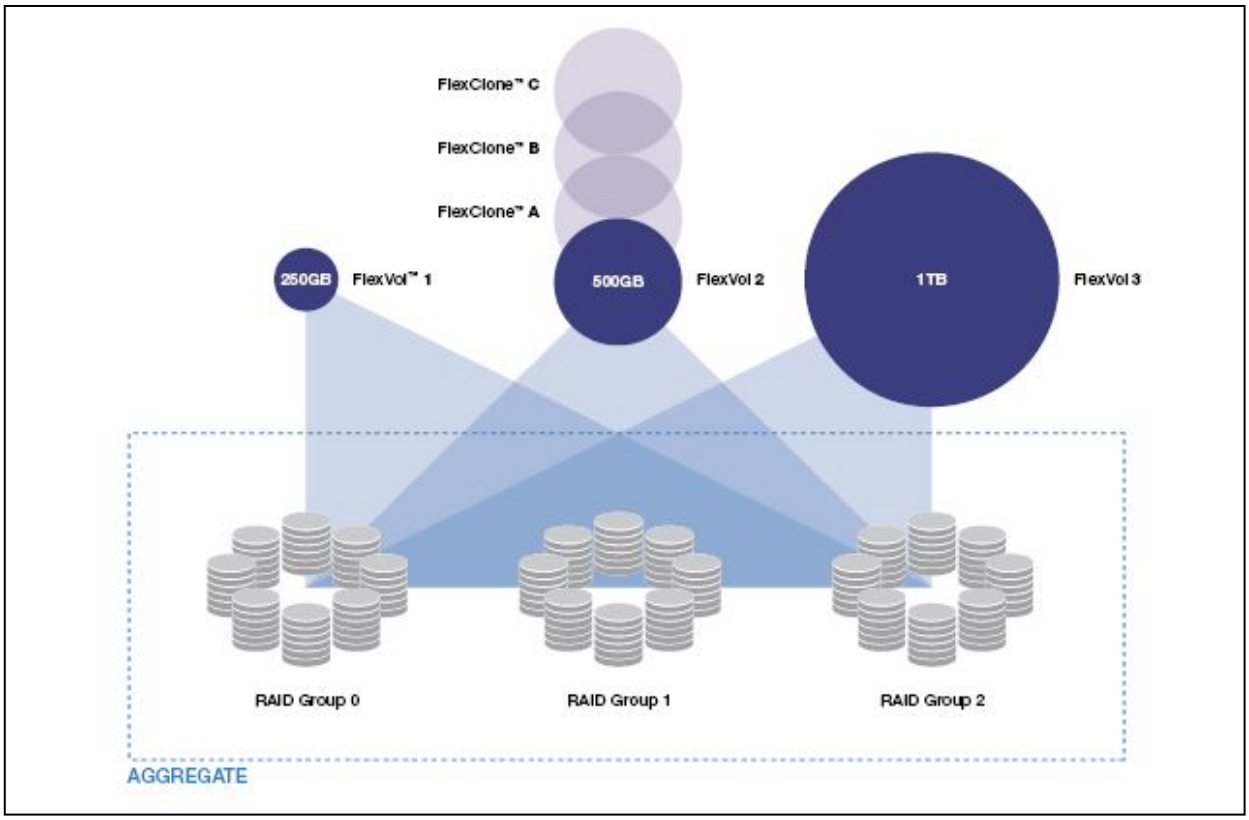

Figure 1. FlexVol technology.

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 volume 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. 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.

IBM System Storage N series with FlexShare™ gives administrators the ability to leverage existing infrastructure and increase processing utilization without sacrificing the performance required to meet critical business needs. It prioritizes processing resources for key services when the system is under heavy load. With the use of FlexShare, administrators can confidently consolidate different applications and data sets on a single storage system. FlexShare makes it possible for administrators to prioritize applications based on how critical they are to the business. For example, production SAP systems are configured with a higher priority than test and development systems.

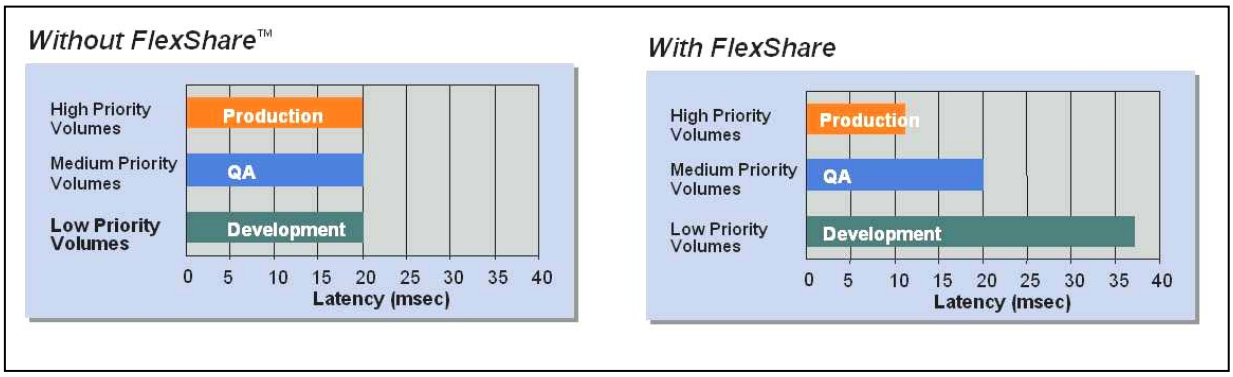

Figure 2. FlexShare.

FlexShare is an IBM System Storage N series with Data ONTAP® software feature that provides workload prioritization for a storage system.

With IBM System Storage N series with SnapDrive® for UNIX 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 the different UNIX Volume Managers and provides a layer of abstraction between application data and the physical storage associated with that data. SnapDrive for UNIX eliminates the need to maintain manual scripts normally used to back up and restore data to specific drives or mount points being used by various downstream applications and databases, without extensive downtime. It can also be used to easily add storage as needed, eliminating the need to preallocate large amounts of storage resources based only on forecasted demand. It also allows server and storage administrators to dynamically reallocate storage resources using the powerful FlexVol capabilities built into the Data ONTAP operating system.

Storage Layout

Aggregate Layout

It is recommended that you use a single aggregate per storage controller to store all data of all SAP systems. The use of a single large aggregate provides the performance benefits of all available disk spindles in the aggregate to every FlexVol volume in that aggregate. Adding a second aggregate is recommended only if the maximum capacity of the first aggregate is reached. In addition, it is recommended you use an aggregate other than the one used by the root volume.

The aggregates should be configured with IBM System Storage N series with RAID-DP™ (redundant array of inexpensive disks, double parity), which offers a high level of data protection. The reliability of RAID-DP is far better that of than RAID5 and very close to that of RAID1. Only if three disks within the same RAID group fail at the same time will data loss occur.

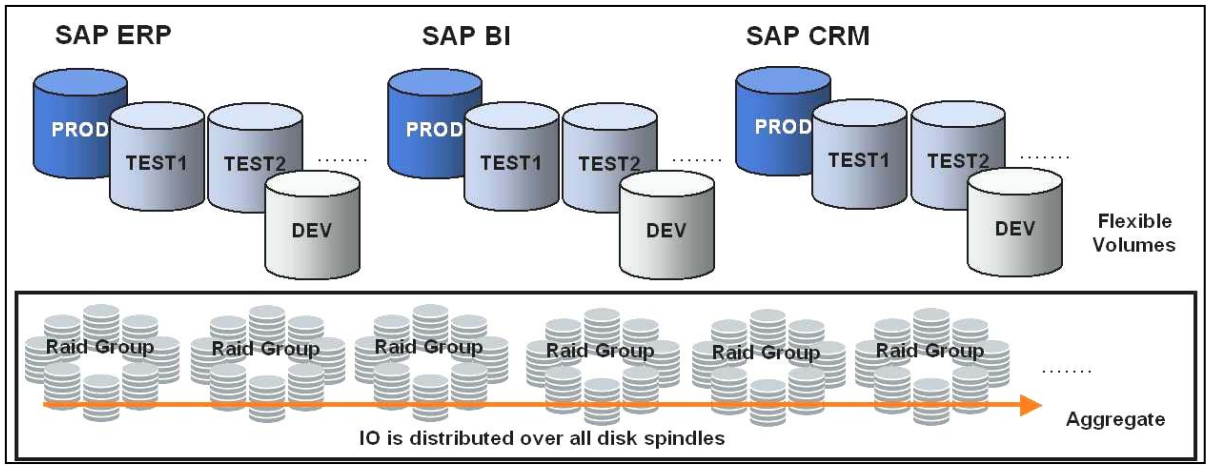

Figure 3. Aggregate layout.

The design of the physical disk layout is very simple because it is not done on a per-SAP-system basis. The aggregate is created as a physical storage resource pool, and storage resources are assigned on a logical, virtualized level with FlexVol volumes. The size of the FlexVol volumes can be easily increased or decreased during online operation without any reconfiguration of the underlying physical disk structure. This allows optimal utilization of the storage resources.

During normal operations the production systems need the highest performance and therefore the highest number of disk spindles compared to test and development systems. Based on the resource-sharing concepts with disk aggregates, the production systems will benefit from the disk spindles of the test and development systems, which are needed anyway because of capacity requirements.

With shared resources, it is always possible that there will be contention for available resources among systems. A stress test, which runs on a test system, might influence the response times of the productive systems because too many I/O resources might be used by the test systems. FlexShare can address this issue. FlexShare is a powerful tool that provides control-of-service for Data ONTAP storage systems. With FlexShare processing, resources can be prioritized on the FlexVol level. Productive systems are configured with a high priority compared to a medium or low priority for the test and development systems. The prioritization can be easily adapted during online operation.

FlexVol Volume Layout

Each SAP system uses three FlexVol volumes:

- One volume for the database data files
- One volume for the online redo log files and the archived log files
- One volume for the mirrored redo log files.

The Oracle data and mirrored log files are separated from the archived log files and the online redo logs stored in the other aggregate. Therefore it is always possible to recover the database without data loss if one of the two aggregates is lost.

Storing the database data files in a FlexVol volume separated from the redo logs is important to allow usage of IBM System Storage N series with Snapshot™ copies, IBM System Storage N series with SnapRestore®, and IBM N series with FlexClone and other Data ONTAP features that work on the volume level.

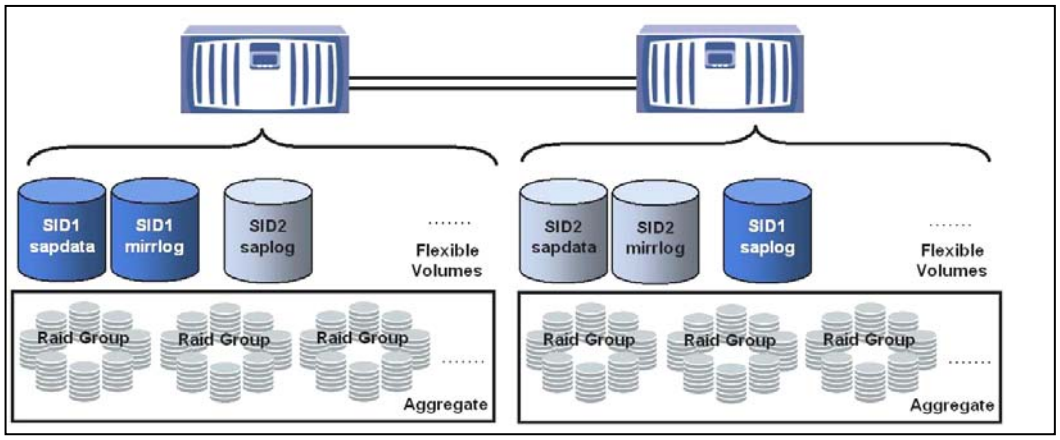

Figure 4. FlexVol volume layout.

Table 1 shows the distribution of file systems of a single SAP instance to the FlexVol volumes.

Storage Controller1		Storage Controller2	
FlexVol sapdata	**FlexVol mirrlogs**	**FlexVol saplog**	
Sapdata1	MirrlogA	OriglogA	
Sapdata2	MirrlogB	OriglogB	
Sapdata3		Oraarch	
Sapdata4		Oracle binaries	
Sapdata5		SAP binaries	
Sapdata6		Sapreorg	
Sapdata##		Sapbackup	

Table 1. FlexVol volume layout.

Layout with MetroCluster

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

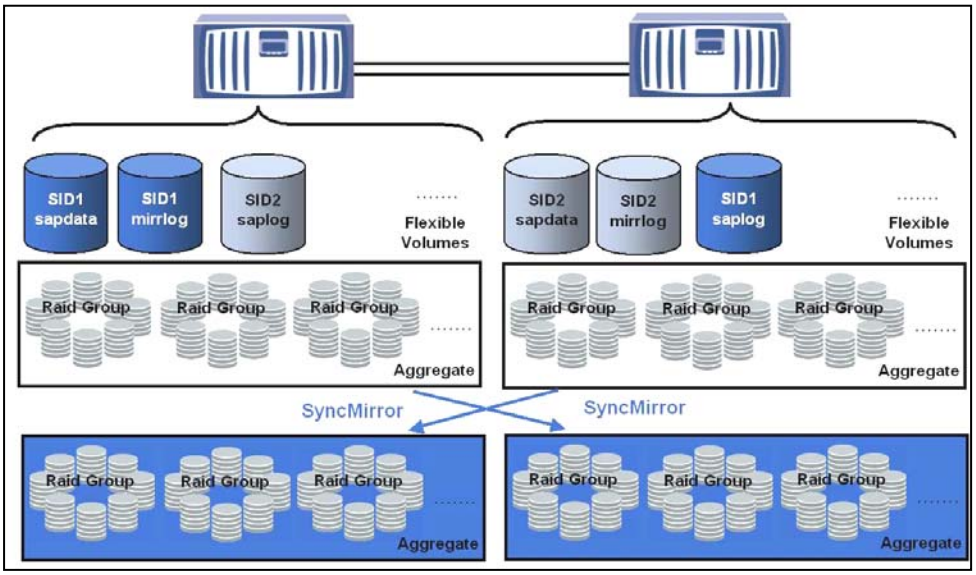

Figure 5. Storage layout with IBM N series with MetroCluster.

Additional aggregates are necessary only if parts of the landscape require synchronous mirroring. For example, the productive SAP systems require synchronous mirroring, but the test and development systems don't.

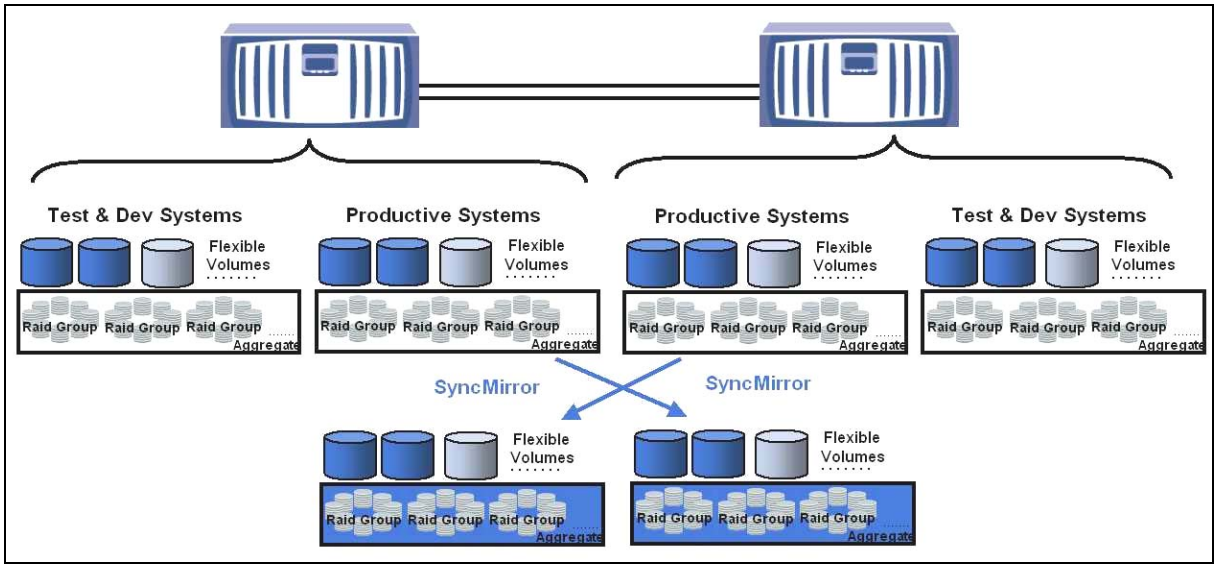

Figure 6. Storage layout with IBM N series with MetroCluster; only production systems are mirrored.

LUN and Volume Manager Layout

The following description assumes that a Volume Manager is being used on the UNIX host.

The size of the database determines the number and size of logical unit numbers (LUNs) required. The goal is to find a balance between the performance advantages of a large number of smaller LUNs and the ease of management that comes with a smaller number of large LUNs.

The following table gives some guidelines for a reasonable number of LUNs, based on the size of the database.

	Database Size	# of LUNs	Size of LUNs	# of LUNs	# of LUNs
Small	< 200GB	4-10	10GB –	1	1
Medium	200GB –	10-20	40GB –	2-4	2-4
Large	> 1TB	> 20	60GB –	>= 4	>=4

Table 2. Number of LUNs, based on database size.

Recommended Standard Layout

The following figure shows the LUN configuration for a small SAP system from the storage point of view. Four LUNs are configured in the sapdata FlexVol volume for the database data files. One LUN is configured for the mirrored online redo logs in the mirrlog FlexVol volume. Three LUNs in the saplog FlexVol volume are used to store the SAP and Oracle binaries, the online redo logs, and the Oracle archive logs.

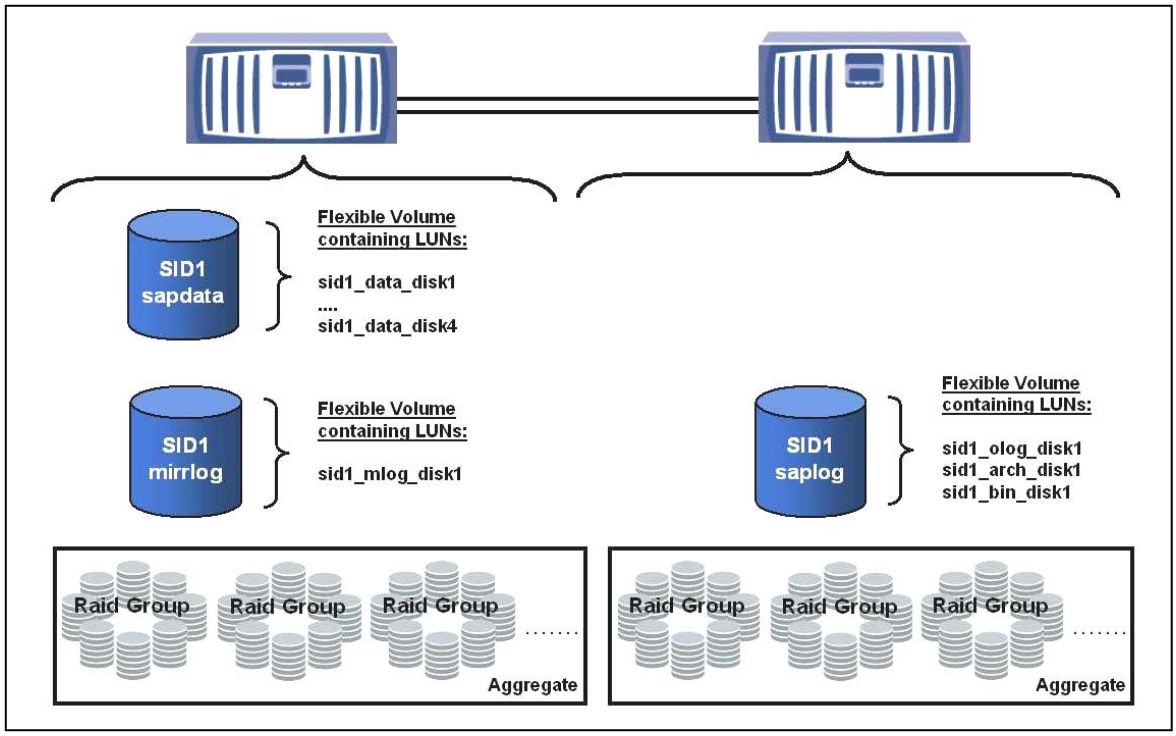

Figure 7. Standard LUN layout.

From the host point of view, two disk groups must be configured with the host Volume Manager:

- The Data Disk Group contains all LUNs for the database data files.
- The Log Disk Group contains all LUNs for the redo logs, archived logs, and binaries.

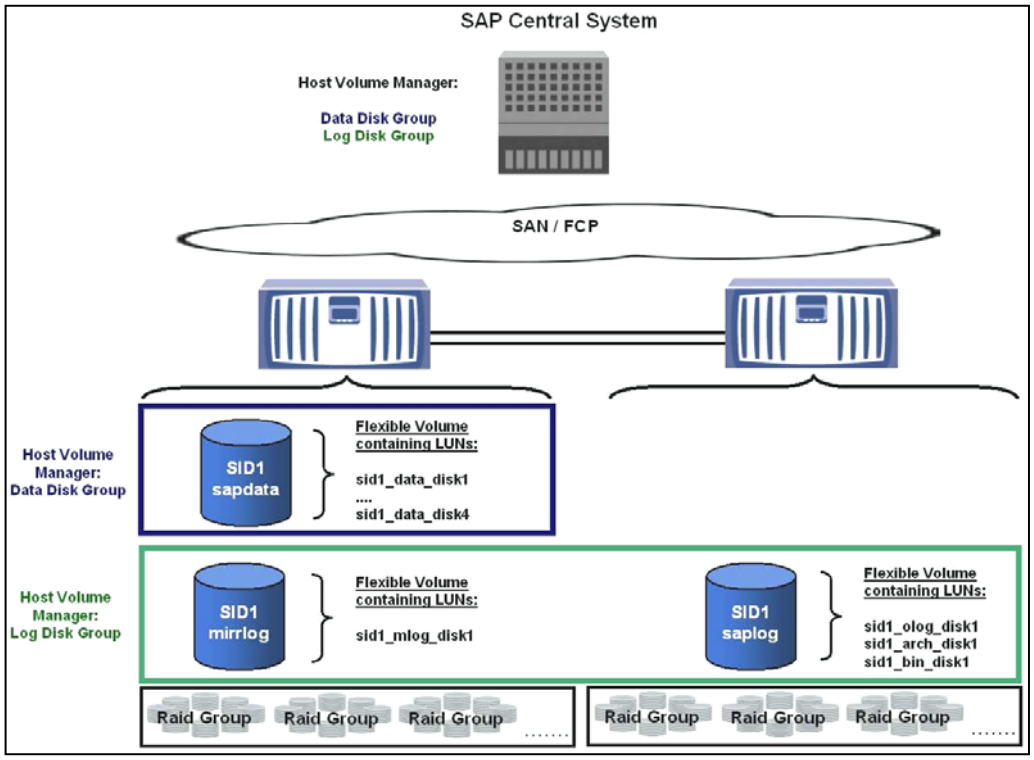

Figure 8. Standard LUN and Volume Manger layout.

The following table shows the logical volumes that need to be configured within the Volume Manager disk groups and the corresponding file systems at the host.

Data Disk Group		Log Disk Group	
LUN	**Logical Volume File System**	**LUN**	**Logical Volume File System**
sid1_data_disk1	/oracle/SID1/sapdata1	sid1_mlog_disk1	/oracle/SID1/mirrlogA
sid1_data_disk2	/oracle/SID1/sapdata2		/oracle/SID1/mirrlogB
sid1_data_disk3	/oracle/SID1/sapdata3	sid1_olog_disk1	/oracle/SID1/origlogA
sid1_data_disk4	/oracle/SID1/sapdata4		/oracle/SID1/origlogB
		sid1_arch_disk1	/oracle/SID1/orarch
		sid1_bin_disk1	/oracle/SID1
			/sapmnt/SID1
			/usr/sap/SID

Table 3. Logical volumes and file systems.

FlexVol Volume Layout For large SAP systems with High I/O Requirements

Systems with high I/O requirements should be distributed to both storage controllers. It can also be beneficial to distribute data from small production systems across both storage controllers to account for future growth. Taking this step during the initial installation will prevent costly downtime in the future as the production system's I/O requirements grow beyond the capacity of a single storage controller. With this configuration, a second aggregate per storage controller should be configured in order to guarantee that no data is lost if any aggregate is lost.

Each SAP system uses five FlexVol volumes:

- Two volumes for the database data files distributed to both storage controllers
- One volume for the online redo log files
- One volume for the archived log files
- One volume for the mirrored redo log files.

The Oracle data and mirrored log files are separated from the archived log files and the online redo logs stored in the other aggregates. Therefore it is always possible to recover the database without data loss if any of the aggregates is lost.

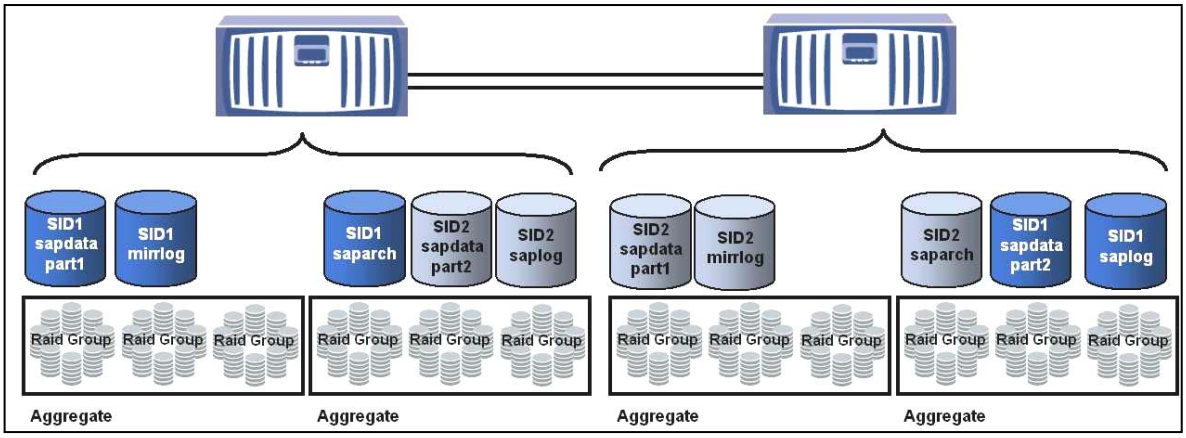

Figure 9. FlexVol layout for large SAP systems.

Table 4 shows the distribution of file systems of a single SAP instance to the FlexVol volumes.

Storage Controller1			Storage Controller2	
FlexVol sapdata	FlexVol saparch	FlexVol mirrlogs	FlexVol saplog	FlexVol sapdata
Sapdata1	Orarach	MirrlogA	OriglogA	Sapdata7
Sapdata2		MirrlogB	OriglogB	Sapdata8
Sapdata3			Oracle	Sapdata9
Sapdata4			SAP binaries	Sapdata10
Sapdata5			Sapreorg	Sapdata11
Sapdata6			Sapbackup	Sapdata12
Sapdata##				Sapdata##

Table 4. FlexVol layout for large SAP systems.

LUN and Volume Manager Layout For large SAP systems with High I/O Requirements

The following figure illustrates an example of the LUN configuration from the storage point of view. In total, 20 LUNs are configured for the database data files. Ten LUNs are stored in the FlexVol volume "sapdata" on the first storage controller, and another 10 LUNs are stored on the second storage controller. This approach allows load distribution to both storage controllers.

Two LUNs for the mirrored online redo logs are configured in the "mirrlog" FlexVol volume. Another two LUNs for the online redo logs are stored in FlexVol volume "saplog." The SAP and Oracle binaries are stored in two separate LUNs in the FlexVol volume "saplog." The archive logs are stored in the FlexVol volume "saparch."

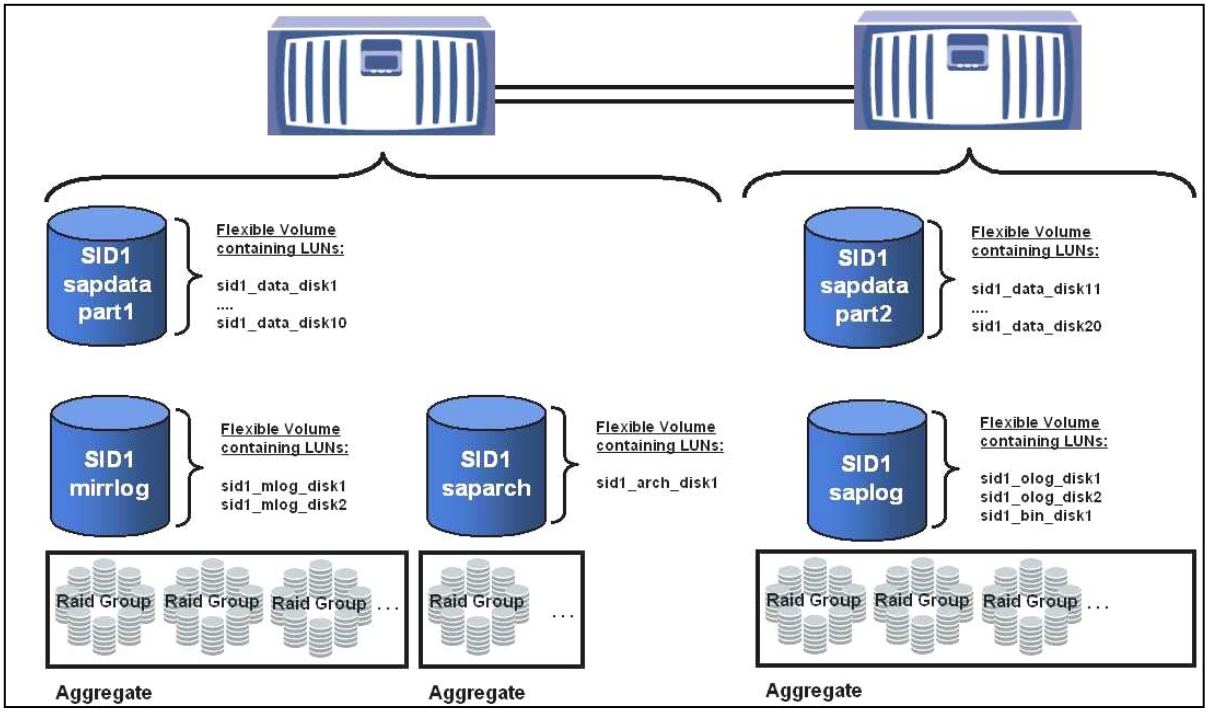

Figure 10. LUN layout for large SAP systems.

From the host point of view, two disk groups must be configured with the host Volume Manager:

- The Data Disk Group contains all LUNs for the database data files.
- The Log Disk Group contains all LUNs for the redo logs, archived logs, and binaries.

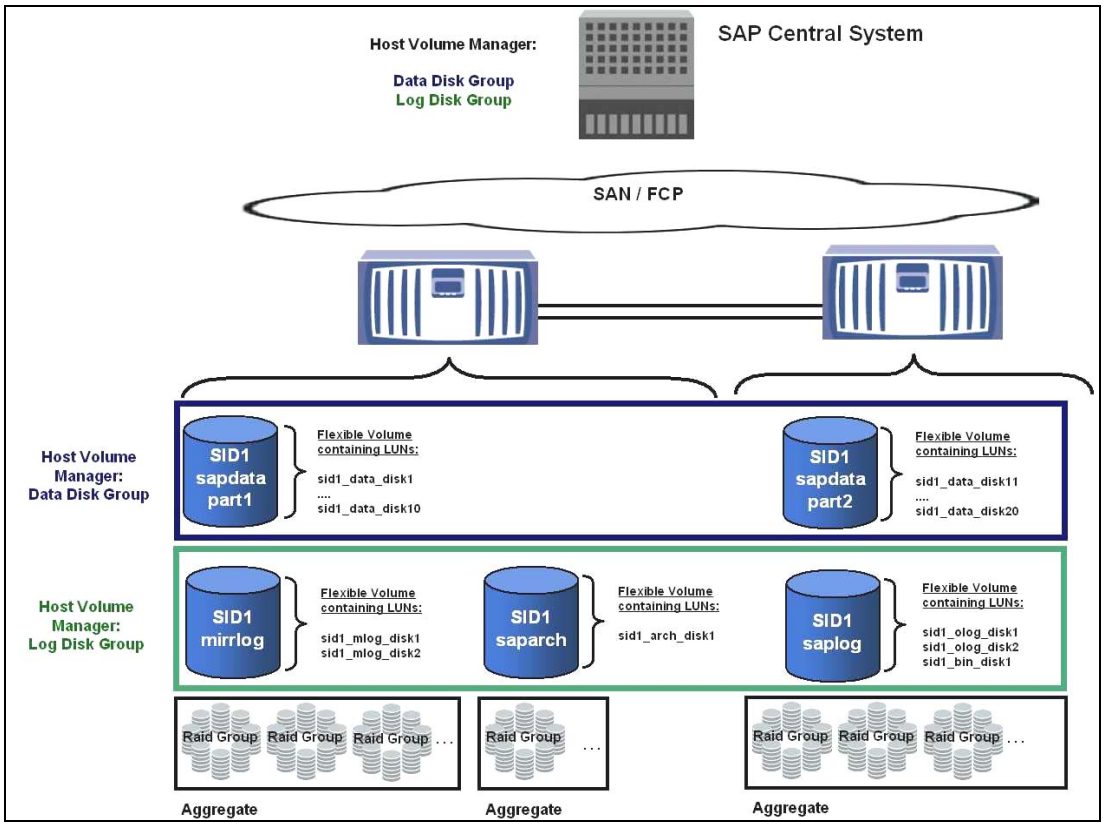

Figure 11. Volume Manager layout for large SAP systems.

The following table shows the logical volumes that must be configured in the Volume Manager disk groups and the corresponding file systems at the host. Each logical volume for the database data files must be configured as a concatenated volume that consists of one LUN from each storage controller. In this approach, the load distribution to both storage controllers is done for each sapdata file system.

Data Disk Group		Log Disk Group	
LUN	Logical Volume File System	LUN	Logical Volume File System
sid1_data_disk1 sid1_data_disk11	/oracle/SID1/sapdata1	sid1_mlog_disk1 sid1_mlog_disk2	/oracle/SID1/mirrlogA
sid1_data_disk2 sid1_data_disk12	/oracle/SID1/sapdata2		/oracle/SID1/mirrlogB
sid1_data_disk3 sid1_data_disk13	/oracle/SID1/sapdata3	sid1_olog_disk1 sid1_olog_disk2	/oracle/SID1/origlogA
			/oracle/SID1/origlogB
		sid1_arch_disk1	/oracle/SID1/orarch
		sid1_bin_disk1	/oracle/SID1
			/sapmnt/SID1
			/usr/sap/SID

Table 5. Logical volumes and file systems for large SAP systems.

Sizing

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

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

- I/O requirements
- Capacity requirements
- Backup/recovery requirements (mean time to recover, backup window, retention policy)
- Cloning requirements (FlexClone copies or full copies)
- DR requirements (synchronous or asynchronous mirroring)
- High-availability requirements (storage system clustering)
- Satisfying I/O requirements is critical, directly affecting overall SAP system performance.

For existing SAP systems, the I/O requirements need to be measured using database or operating system tools. Database tools can be, for example, Oracle Statspack or the SAP database performance monitor. For instance, iostat can be used if the measurement is done on the operating system level. Independent of which tools are used, it is very important that the measurement is done during peak loads on the SAP system. Especially when database tools are used for the measurement, a suitable time frame must be chosen, such as one hour, because 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 SAPS values for the systems, which are provided by the SAP Quick Sizer, can be used to estimate the I/O requirements. Of course, the storage sizing is much more accurate if I/O values are measured.

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

The I/O requirements determine the type and number of disk spindles and storage controllers.

In order 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 FlexClone volumes
- Synchronous or asynchronous mirroring.

Based on the capacity requirements, the type and number of disks and the storage controller supporting the capacity are determined.

The results of the I/O sizing and the capacity sizing are compared in a final step to define the right storage system 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 an Oracle database on a UNIX server with the FC protocol.

General Requirements

It is recommended that you use SnapDrive for UNIX, an IBM N series host-based software product that simplifies storage management and provisioning in SAP FC storage environments. It integrates with IBM N series with Snapshot and SnapRestore to simplify the process of taking error-free, host-consistent data snapshots.

IBM N series Storage Controller Configuration

Snapshot backups for database applications won't be consistent from a database point of view without shutting down the database or putting the Oracle database in hot backup mode. Thus, you must turn off automatically scheduled snapshot copies on the storage level on database volumes:

```
filer> vol options <volname> nosnap on
```

For performance and security reasons, the following options should be set:

filer> vol options vol_sapdata nvfail on
filer> vol options vol_sapdata no_atime_update on

Storage Migration

This section discusses various storage migration approaches. Migrations that include a change of operating system or database system can't be done solely at the storage level; you must use SAP migration tools that 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 depends heavily on the acceptable downtime of the application. Furthermore, downtime depends on the amount of data 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 Volume Manager level
- Migration on the database level
- Migration on the storage system 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 starts, the database and SAP system must be shut down. The data is then 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 are changes in the file system structure, the new structure must be configured in 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 Volume Manager level

If a Volume Manager is used on the host, a migration can be done using a host-based mirror. The IBM N series storage system must be configured and attached to the host. All data that needs to be migrated must be mirrored to the IBM N series storage by adding an additional plex to the disk/volume group. The synchronizing of the new mirror should be scheduled during a low-activity period, because there is a high load on the server during the synchronization. When the synchronization is finished, all new data is synchronously mirrored to the IBM N series storage system. Before the storage system switch can be done the SAP and database system must be shut down. The mirror must be broken and the Volume Manager must be reconfigured so that the plex stored at the old storage system is not used any more.

A migration on the Volume Manager level can be done for an FC-to-FC protocol migration. This approach requires downtime during the connection of the IBM N series storage to the database host and during the final switch to the new storage system. The data migration is done during online operation, but has a performance impact on the database host.

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 productive 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 is then connected to the database server and the LUNs are mounted to the server. The online logs, control files, and archive logs that have not yet been copied are now 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 remains the same, the database can be started immediately. If there is a change in the file system structure, the new structure must be configured within Oracle by creating a new control file. Finally, a forward recovery of the database is carried out.

A database-level migration can be done for an FC-to-FC or a FC-to-iSCSI protocol migration. This approach reduces migration downtime but requires another server during the migration process.

Migration on the storage system level

For data migration from a non-IBM N series SAN to an IBM N series SAN storage system, there is a fast, flexible SAN data migration service to suit individual user requirements. Essentially, the service uses a FAS appliance and host software, 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 and software must be configured to see the existing and new arrays, and then must be enabled to perform the migration; for example, security zones and domain reconfiguration. The appliance is configured to

migrate the appropriate data sets and/or data files. The software automatically transfers 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. When the migration is finished, the servers and storage infrastructure must 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 minimizes the downtime during the migration, but needs the appliance described earlier for the migration process. Table 6 summarizes the migration processes.

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

Table 6. Summary of different migration 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 and flexible IBM N series storage infrastructure allows business owners and IT departments to work more efficiently and effectively toward the goal of improving enterprise business processes.

Storage consolidation with IBM N series ensures the high availability and performance of SAP data and applications so that stringent service-level agreements (SLAs) are met. In addition, IBM N series helps reduce the administration and management costs associated with deploying these new business applications and processes.

System Management and Maintenance

SAP System Cloning

Business Challenges

A typical SAP customer environment today 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 about 10 copies of different SAP components. These copies must be refreshed, often on a weekly or monthly basis.

The creation of an SAP system copy normally takes several days and negatively affects the production environment. Lots of manual steps also consume valuable IT staff time.

The source database must be exported using SAP tools and imported at the target system, or an offline backup of the source database will be 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, consuming SAP basis administrators' time.

Being able to quickly create an SAP system copy on demand is becoming more important:

- Quality insurance systems must be refreshed on a weekly basis.
- Additional test systems must be set up quickly to perform specific integration tests.
- Test system with current productive data must be set up quickly for SAP upgrade projects.
- Training systems must be set up or resynced.
- The traditional approach to create system copies does not address these demands.

SAP copies also consume large amounts of storage, which must be provided. Because these copies are typically clones of a productive 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 an SAP system copy on demand, in a few minutes, without any impact to the source productive system. In addition, IBM N series cloning functionality allows very efficient storage management by storing only data changes between the parent and the clone.

IBM N series Solution for SAP System Copies

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

Figure 12 shows the basic concept of the system copy solution. Creating an 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 created. This is done during online operation and has no performance impact on the source system.

On the target system, this snapshot copy is the base for the LUN clone image. The creation of the LUN clone image takes only a few seconds. The LUN clone image 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 postprocessing tasks need to be accomplished.

All of these steps can be fully automated and do not need any manual interaction. A SAP system copy can be accomplished in a few minutes using the IBM N series solution.

Table 7 (next page) compares the traditional approach to the IBM N series approach to perform an SAP system copy.

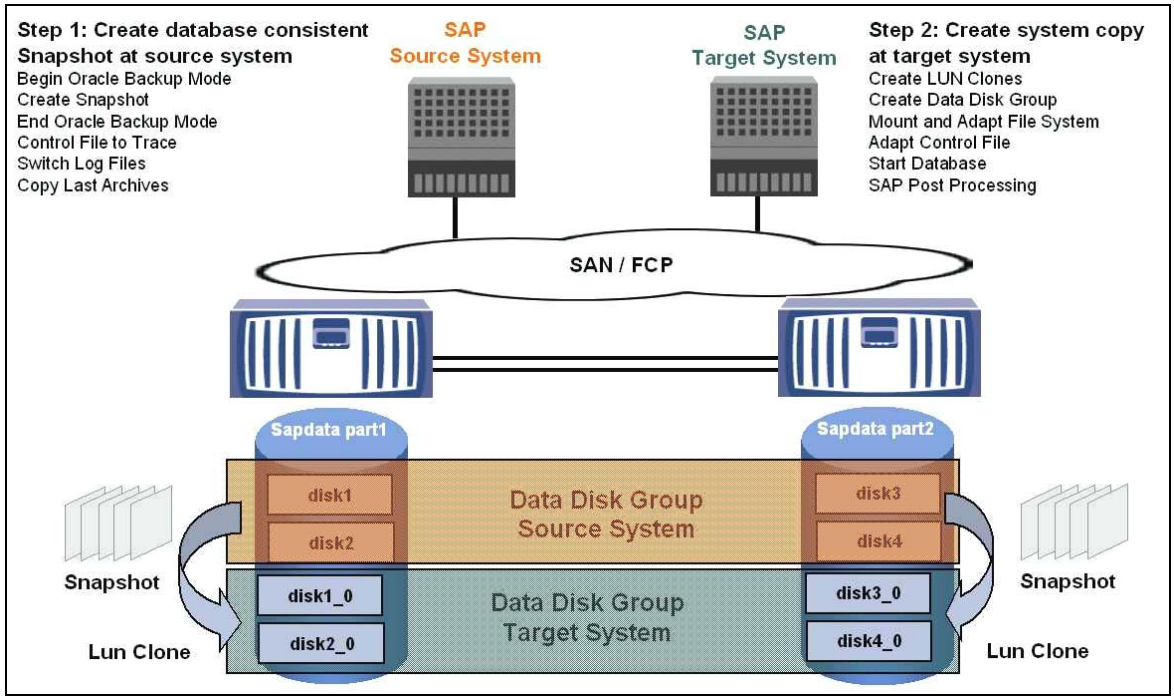

Figure 12. SAP system cloning overview.

Necessary steps at the source system:

With the traditional approach, it is necessary to create an online or offline backup of the source database. A backup typically has a significant performance impact on the source system and therefore cannot be scheduled during working hours. Depending on the database size, the backup of the database takes several hours. The subsequent steps are typically carried out manually, consuming IT staff time.

With the IBM N series approach, the backup is taken using Snapshot functionality. Creating a snapshot copy takes 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:

To set up a new SAP test system, the SAP and Oracle software must be installed once. This is required with both approaches. With subsequent system refreshes, this step is unnecessary.

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

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

Traditional Approach	IBM N Series Approach	Advantages
Necessary steps at the source system		
• Offline or online backup • Control file to trace • Switch log files and copy archived logs to shared location (with online backup)	• Snapshot backup during online operation • Control file to trace • Switch log files and copy archived logs to shared location	• Significant impact on operation, difficult to schedule with the traditional solution. No impact on operation with IBM N series solution. Can be scheduled at any time. • Manual process vs. fully automated process with IBM N series solution. • Manual process vs. fully automated process with IBM N series solution.
Necessary steps at the target system		
• Install SAP system (if not existing yet) • Restore offline backup from source system • Adapt directory names to new SID • Create new control file with new SID based on control file trace from source system • Adjust Oracle security • SAP-specific post processing tasks	• Install SAP system (if not existing yet) • Create LUN clone image based on snapshot backup; mount the LUN clone image at the target system • Adapt directory names to new SID • Create new control file with new SID based on control file trace from source system • Adjust Oracle security • SAP-specific post processing tasks	• Same approach. • Restore takes several hours vs. several seconds with LUN clone technology. Fully automated process with IBM N series solution. • Manual process vs. fully automated process with IBM N series solution. • Manual process vs. fully automated process with IBM N series solution. • Manual process vs. fully automated process with IBM N series solution. • Manual process vs. fully automated process with IBM N series solution.

Table 7).SAP system copy comparison.

Conclusions: The IBM N series system copy solution significantly improves the process to create SAP system copies:

- A system copy can be accomplished in several minutes, compared to several days with the traditional approach.
- System copies can be scheduled at any time because there is no impact on the online operation of the source system.
- Snapshot and LUN clone functionality reduce 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 and frequently repeatable, minimizing the involvement of IT staff.
- Snapshot and LUN clone functionality significantly reduce the necessary disk space for an SAP system copy by storing only data changes between the source and the target system.

SAP Upgrade

Business Challenges

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

Upgrading to a new SAP release is a challenging project that consumes large amounts of employee time, including IT staff time.

Business processes are affected during the upgrade project time period, because 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 productive SAP system. Every hour that can be saved while running the upgrade of the productive system is important. Database backups consume extensive time. Optimizing backup and restore functionality is therefore very critical.

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

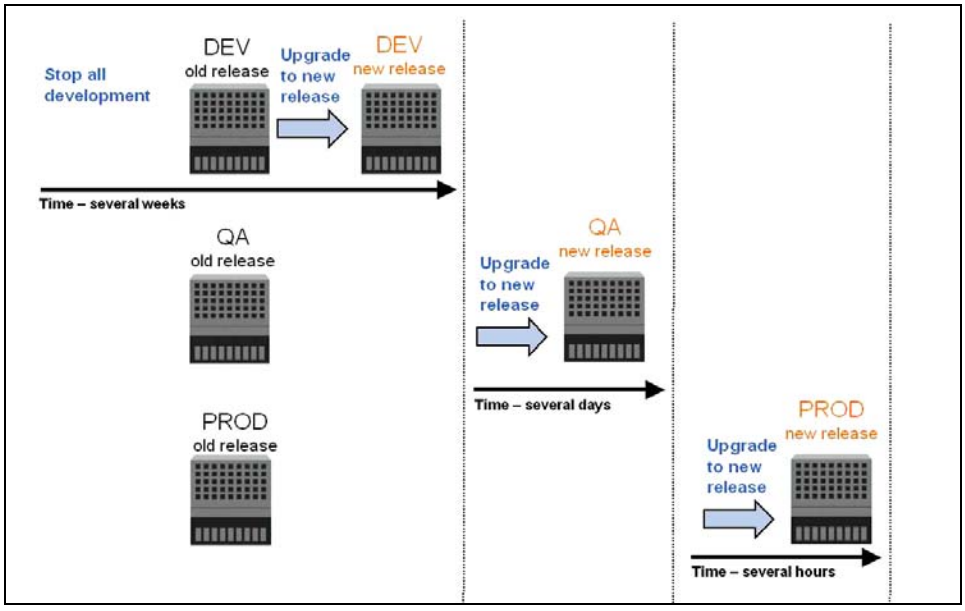

Figure 13. SAP upgrade overview.

IBM N series Solution

The IBM N series solution for SAP upgrades addresses the issues just described by providing a solution for creating fully automated SAP system copies in a few minutes. The IBM N series backup and recovery solution helps to 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 to minimize the risk, reducing the downtime and reducing overload of IT staff resources during an 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 copy of the current development system running on separate hardware. During the upgrade process of the copy 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 system is carried out more than once in order to define the necessary actions in all upgrade phases.

The setup of the separate SAP system is done based on a system copy of the original development system. This system copy can be carried out using the IBM N series system copy solution described in “IBM N series Solution for SAP” section of the document. Reducing the time needed to create this system copy is critical because the copy is typically carried out several times. During the upgrade process and during the modification adjustment, snapshot backups are very helpful, allowing the system to be reset to any snapshot copy and restarting the upgrade phase.

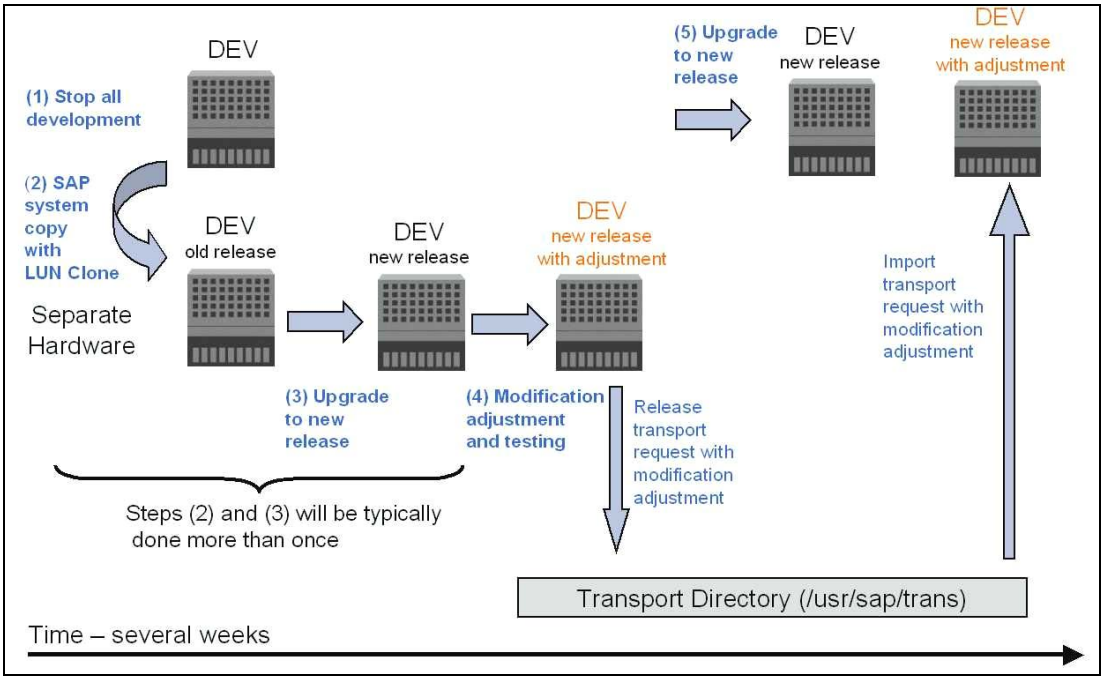

Figure 14. SAP upgrade—development system.

Upgrading the Quality Assurance System

The quality assurance system is upgraded using a fresh system copy of the productive SAP system. One important result of the upgrade of the quality assurance system is the run time of the upgrade with real production data. The IBM N series and SAP system copy solution allows efficient refreshing of the quality assurance system. Reducing the time necessary to create this copy is also critical when upgrading the quality assurance 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, allowing restart to an upgrade phase or restarting the import.

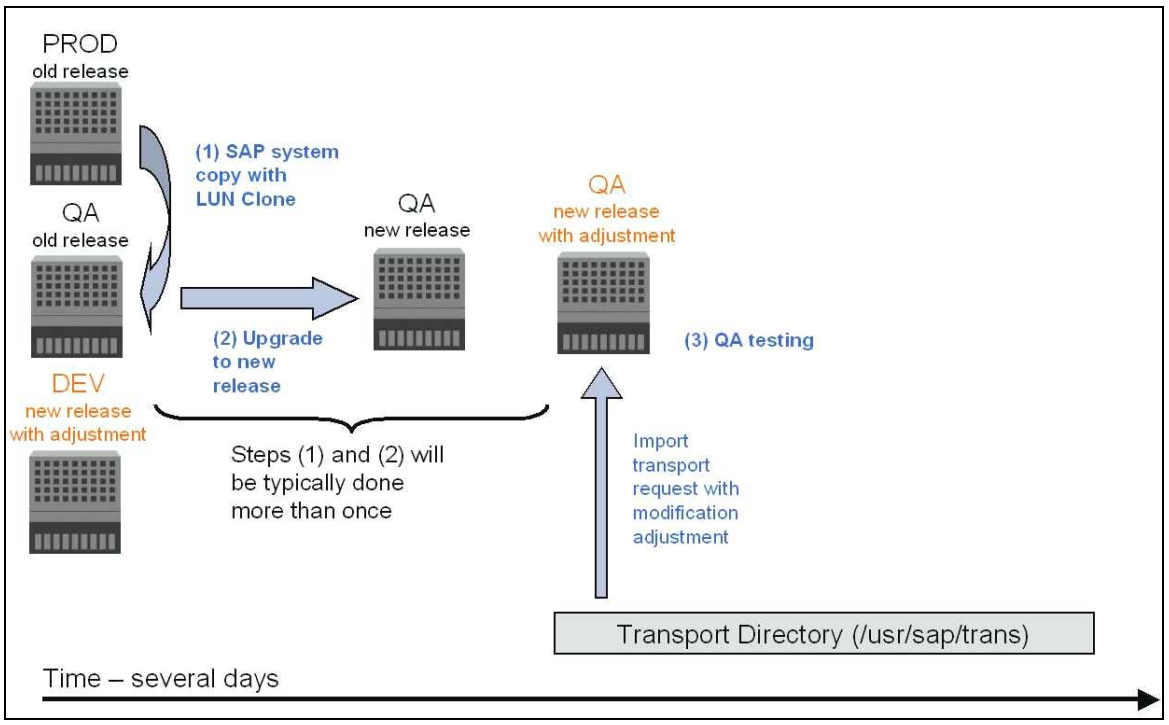

Figure 15. SAP upgrade—quality assurance system.

Upgrading the Productive System

Scheduling is extremely important when the production system is going to be upgraded, because 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 two-day weekend may not be sufficient for the upgrade.

The upgrade of the productive system includes at least three backups of the database. The first backup must be done immediately before the upgrade is started. After the upgrade is finished, 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 must be restored to the old release level.

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

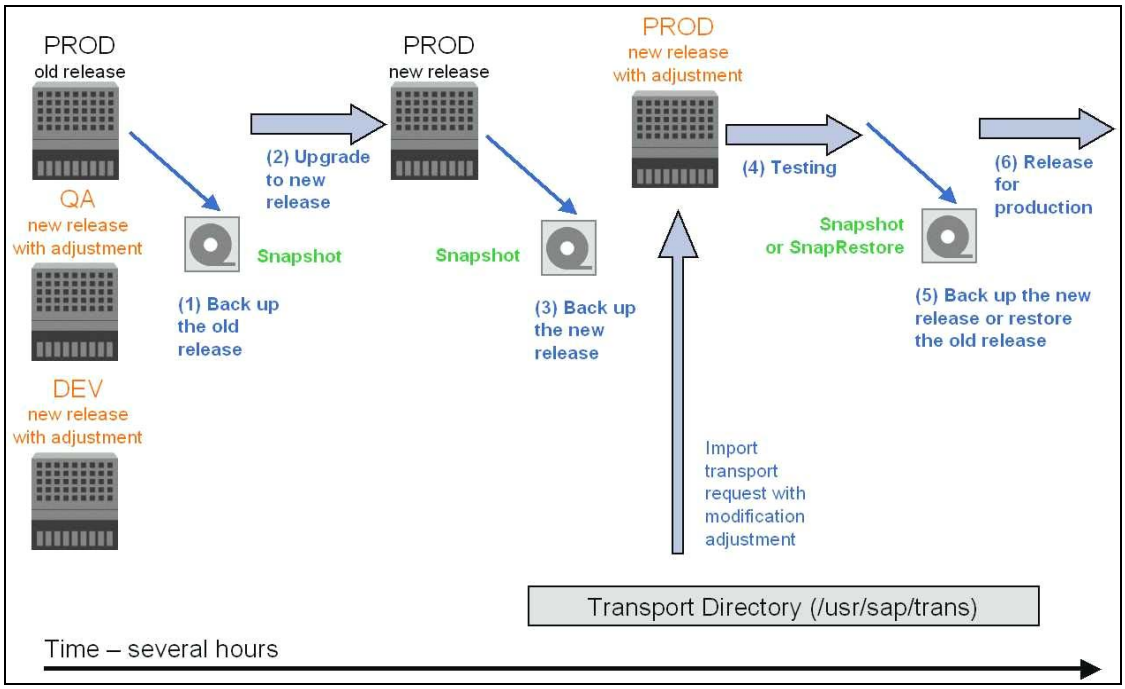

Figure 16. SAP upgrade—productive system.

Conclusions: The IBM N series system copy and backup and recovery solutions significantly improve the SAP upgrade process:

- The IBM N series system copy solution allows quickly refreshing the separate development and the quality assurance system, reducing the time required from several days to several minutes.
- Using the IBM N series backup and recovery solution, snapshot backups of the database can be created, allowing a database to be restored to any specific snapshot copy in a few seconds and restarting any upgrade phase.
- The IBM N series backup and recovery solution significantly reduces the total upgrade time of the productive system, providing a higher level of flexibility for 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, 365 days a year. Consistent levels of performance are expected, regardless of increasing data volumes and routine maintenance tasks, such as system backups. Performing backups of the SAP databases is a critical task, and can have a significant performance impact on the productive 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 point in time at which backups can be performed with minimum impact on the business process. Downtime of SAP productive and even development systems is always a concern. Thus the time needed to restore and recover SAP systems is of particular importance.

Summary of SAP backup and recovery challenges:

- **Performance impact on productive SAP system.** Backups typically have a significant performance impact on the productive SAP system because there is a high load on the database server, the storage system, and the storage network during backups.
- **Shrinking backup windows.** Because conventional backups have a significant performance impact on the productive SAP system, backups can be made only during times with low dialog or batch activities on the SAP system. It becomes more and more difficult to define a backup window when the SAP system is used 24x7.
- **Rapid data growth.** Databases are growing. Rapid data growth together with shrinking backup windows results in ongoing investments in the backup infrastructure—more tape drives, new tape drive technology, faster storage networks, etc. Growing databases also result in more tape media or disk space for backups. Incremental backups can address these issues, but result in a very slow restore process, which is usually 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—the time that is needed to restore the database and the time that is needed 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 always causes 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 must 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 cuts down the total available time for the upgrade process.

IBM N series Solution

IBM N series provides unique storage solutions to address the challenges just described.

IBM N series with Snapshot technology 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, because Snapshot does not move any data blocks. The use of Snapshot technology doesn't have any performance impact on the productive SAP system, because the IBM N series with Snapshot implementation doesn't have to copy data blocks when the data in the active file system is changed. Therefore, creation of snapshot copies can be scheduled without having to consider peak dialog or batch activities periods. SAP and IBM N series users typically schedule several snapshot online backups during the day—for instance, every 4 hours.

Snapshot copies also provide key advantages for the restore and recovery operation. The IBM N series with SnapRestore functionality allows restoring the entire database or parts of the database to the point in time of any available snapshot copy. This restore process is done in a few minutes, independent of the size of the database. Because several snapshot online backups have been created during the day, the time needed for the recovery process is also dramatically reduced. Fewer logs need to be applied, because a restore can be done to a snapshot copy that is at most 4 hours old. The mean time to recover, which consists of the time needed for restore and recovery, is therefore reduced to several minutes, compared to several hours with conventional tape backups.

Snapshot backups are stored on the same disk system as the active online data. Therefore, it is recommended that you use snapshot backups as a supplement, not a replacement for backups to a secondary location, such as disk or tape. Although backups to a secondary 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 are handled by using SnapRestore. Restores from a secondary location (disk or tape) are necessary only if the primary storage system holding the snapshot copies is damaged or if it is necessary to restore a backup that is no longer available from a snapshot copy—for instance, a 2-week-old backup.

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

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

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

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

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

Backup to tape using third-party backup software:

- Network data management protocol (NDMP) backup (serverless backup): The tape is connected directly to the primary storage system. The data is written to tape using NDMP.
- SAP Brbackup split mirror: SAP Brbackup creates the snapshot copies on the storage system, and the data is written to tape with backint.

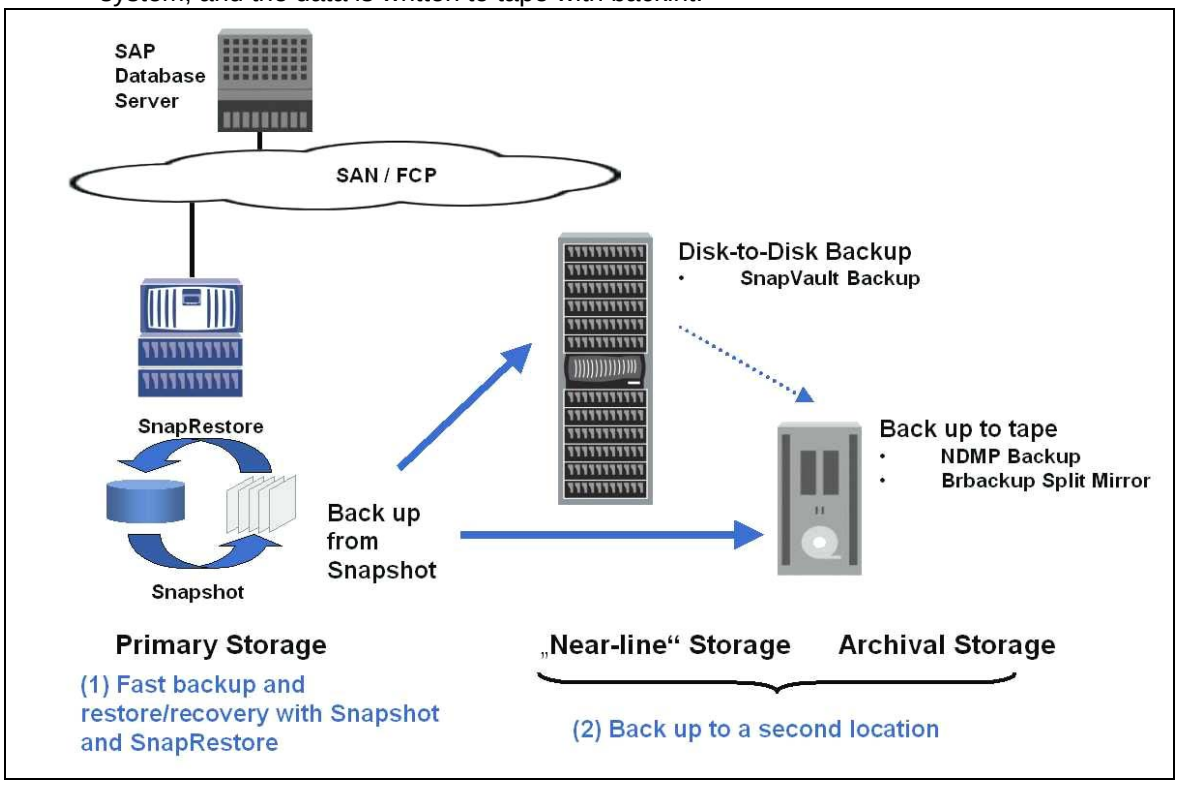

Figure 17. SAP backup and recovery—IBM N series solution overview.

The following figure compares the different backup approaches with regard to the performance impact of a backup and the time that the database must be in hot backup mode or offline.

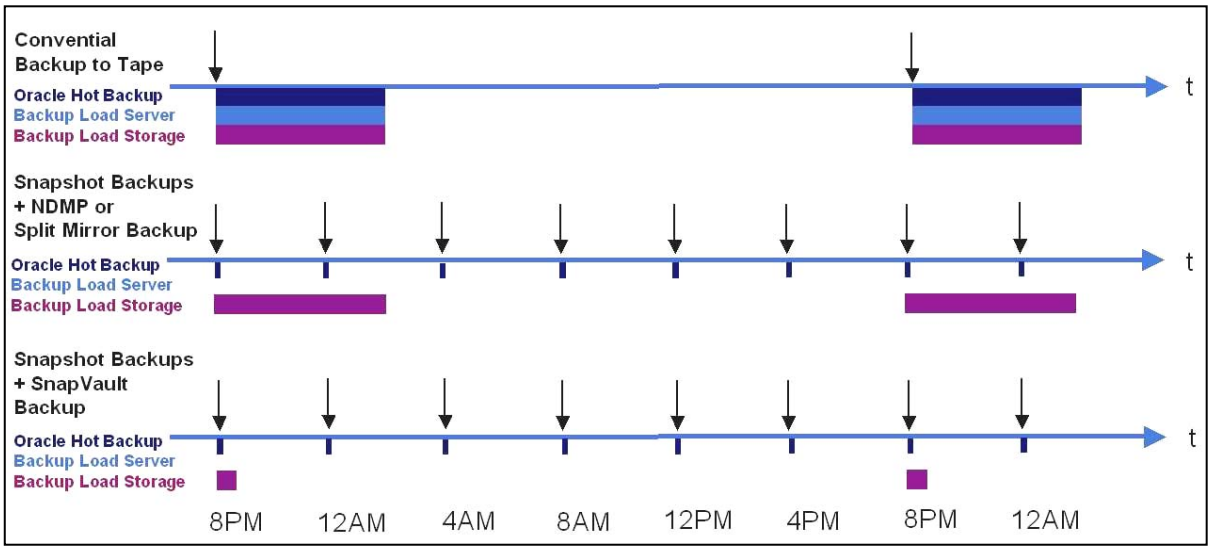

Figure 18. Comparison of different backup approaches.

Conventional Backup to Tape

A conventional backup to tape using SAP Brbackup with backint generates a significant load on the productive database server and the primary storage system. Because this backup is not based on Snapshot, the database is offline or in hot backup mode during the whole backup time. A full backup to tape is typically scheduled once a day.

Snapshot Backups Together with NDMP or Split Mirror Backups

Snapshot backups do not generate any load on the database server or the primary storage system. The database is in hot backup mode or offline for only a few seconds. A full database backup based on Snapshot consumes disk space only for changed blocks. Snapshot backups are typically scheduled more often, for example, every 4 hours. A higher backup frequency allows a more flexible restore process and reduces the number of logs that must be applied during forward recovery. In addition, a full NDMP backup to tape or a split mirror backup to tape is scheduled once a day. This backup still creates a high load at the primary storage system and takes the same amount of time as the conventional tape backup. Because the backup is taken from a snapshot copy, the database doesn't need to be offline or in hot backup mode while running the backup to tape.

Snapshot Backups Together with Disk-to-Disk Backup and SnapVault

Snapshot backups are used in the same way as described in the previous section.

Because SnapVault runs at the storage level, there is no load on the database server. SnapVault transfers only the changed blocks with each backup. Therefore, the load on the primary storage is significantly reduced. For the same reason, the time needed to perform a full database backup is short. 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 very low compared to full tape backups.

The following figure compares the time required to do a restore and recovery.

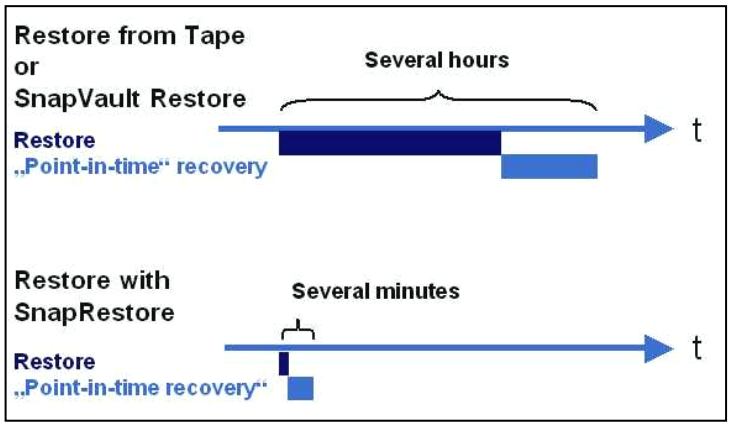

Figure 19. Time needed for restore and recovery.

Restore from Tape or SnapVault Restore

The time that is needed to restore the database from tape or disk depends on the size of the database and the tape or disk infrastructure that is used. In either case, several hours are required for performing a restore. Because the backup frequency is typically one backup a day, a certain number of redo logs need to be applied after the restore is finished.

Restore with SnapRestore

The time needed to restore the database with SnapRestore is independent of the size of the database. A SnapRestore process is always finished in a few seconds. Snapshot backups are taken with a higher frequency, such as every 4 hours, so the forward recovery is much faster, because fewer redo logs need to be applied.

If snapshot backups are used in combination with tape or SnapVault backups, most restore cases are handled with SnapRestore. A restore from tape or disk is necessary only if a snapshot copy is no longer available.

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

- Negligible impact of backups on the productive SAP system
- Dramatically reduced mean time to recover
- Minimum disk space needed for database backups at the primary and the secondary storage systems (primary storage system 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, database verification can be run on a separate server without creating any load on the productive 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 in order to be able to repeat that phase.

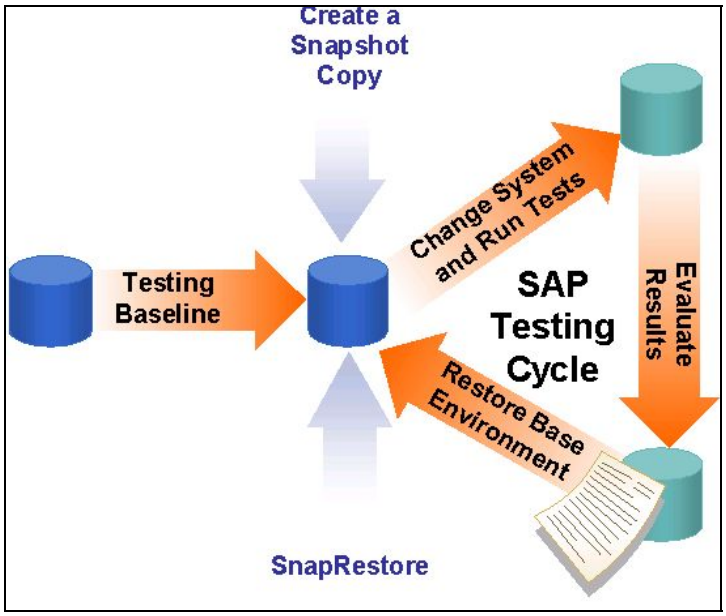

Figure 20. SAP testing cycle.

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 previous 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, because 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, because 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, ensure a higher level of flexibility with regard to scheduling. By creating several snapshot copies at certain stages during the upgrade, the upgrade can be restarted without reverting to the former release status.

High Availability

Business Challenges

Productive SAP systems are business-critical applications that require 24x7 availability. Meeting these requirements requires an infrastructure that does not have any 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 with Clustered Failover delivers a robust and highly available data service for business-critical environments. Installed on a pair of IBM N series storage controllers, 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

The following figure shows a sample clustered failover configuration. A cluster can be created with two storage controllers by connecting the storage controllers via a cluster interconnect. This connection is redundant and is used to exchange cluster heartbeats and to 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 via a second FC loop. If the first storage controller fails, the second storage controller is able to access its partner's disk shelves. The MAC and IP addresses and the worldwide port name (WWPN) of the first storage controller are also adopted. Because the NVRAM is mirrored on both storage controllers via 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.

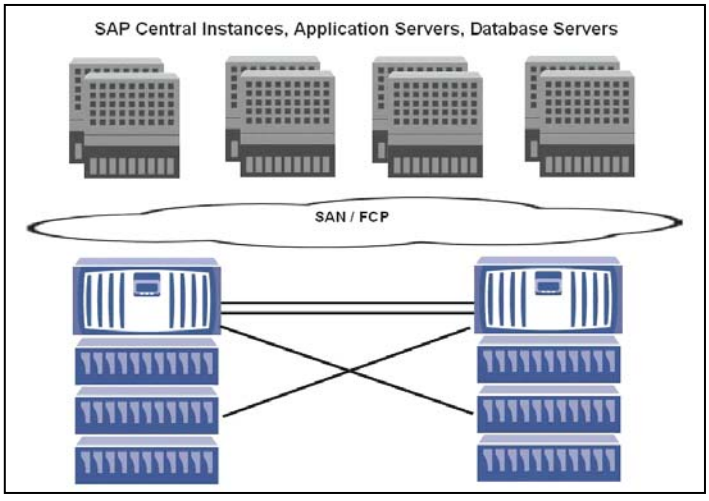

Figure 21. Clustered storage system solution.

Conclusions: The Clustered Failover technology provides an extremely robust high-availability solution.

- A cluster has an availability level of 99.99+%
- Both storage controllers in the cluster can be used actively, providing high availability for both the database server and the central instance.
- A clustered storage system is recommended if server clustering is used for the application.

Disaster Recovery

Business Challenges

Organizations recognize the importance of having 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 minimal 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 selecting a DR solution that utilizes one or many IBM N series products.

IBM N series Solutions

SnapMirror

IBM System Storage 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 a 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 transfer to initialize the DR site. After the initial transfer, incremental changes are passed to the DR site asynchronously. The amount of data lost in the event of a disaster depends on the frequency of the incremental asynchronous transfers. The SnapMirror disaster recovery solution is based on the IBM N series backup and recovery solution. Selective snapshot backups are mirrored to the DR site. Additionally, the qtree where the archive logs are stored must be mirrored using SnapMirror. It is recommended that a frequent SnapMirror update of the archive logs be taken, such as every 10 minutes, to ensure a minimum of data loss.

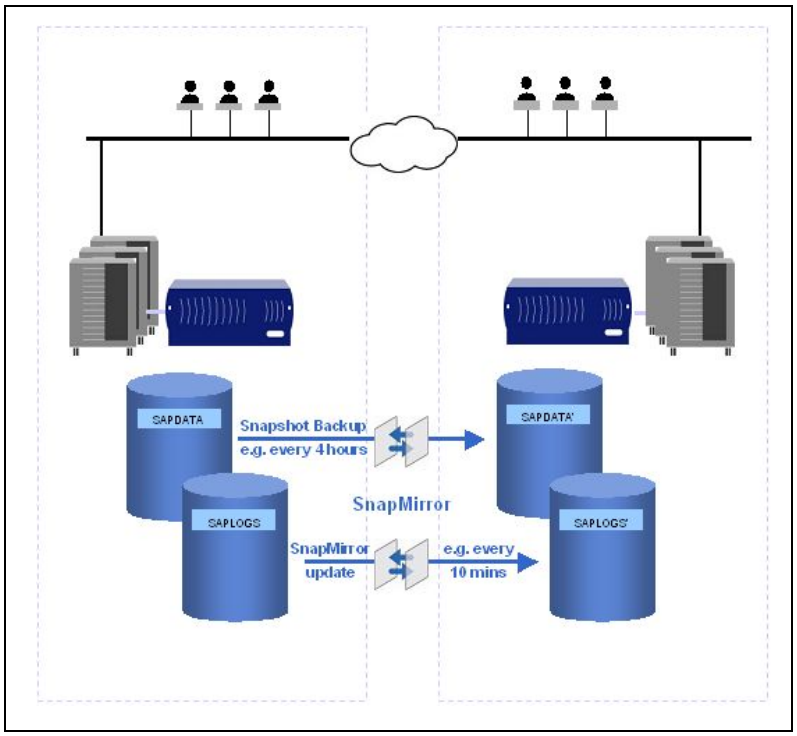

Figure 22. DR with SnapMirror.

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 data there is completely current. The combination of failover and data replication ensures that you can recover from a disaster—with no loss of data—in minutes rather than hours or days.

MetroCluster is much like 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 remote locations up to 30km. 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.

If the distance between the nodes is 500 meters or less, then a Stretch MetroCluster configuration can be used with OM3 cabling. If the distance is greater than 500m, then a Fabric MetroCluster must be chosen, making FC switches necessary.

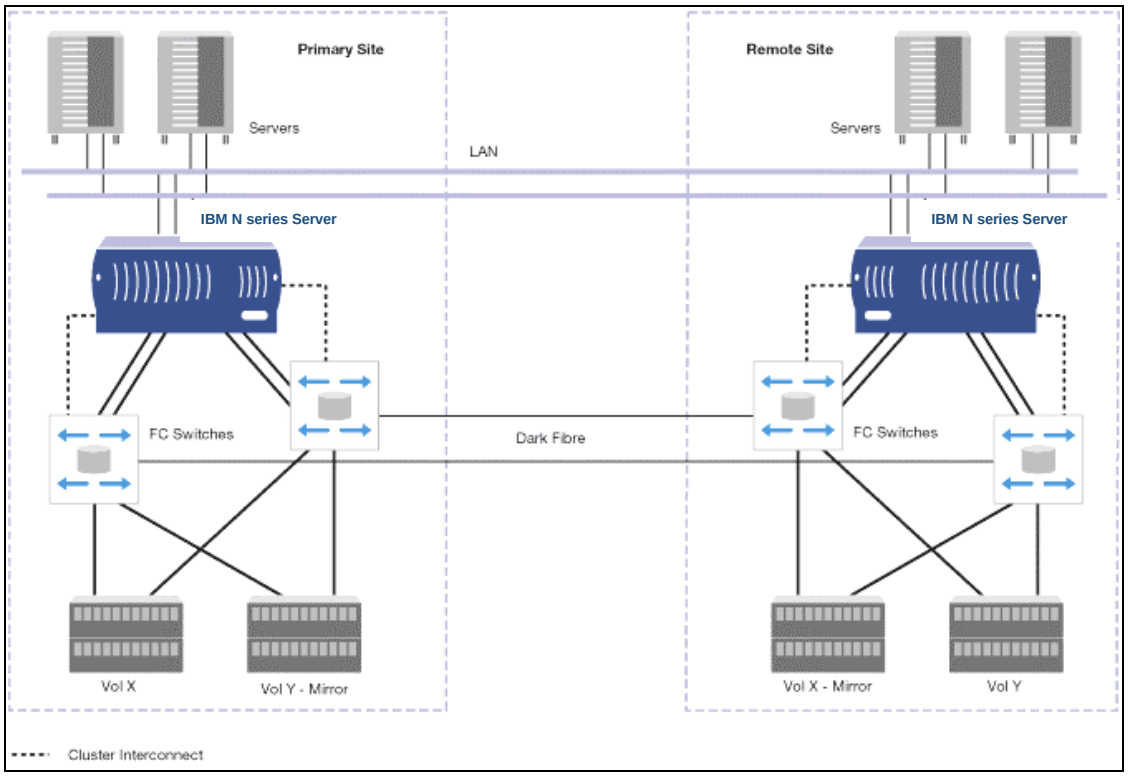

Figure 23. MetroCluster over direct FC switch connection.

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

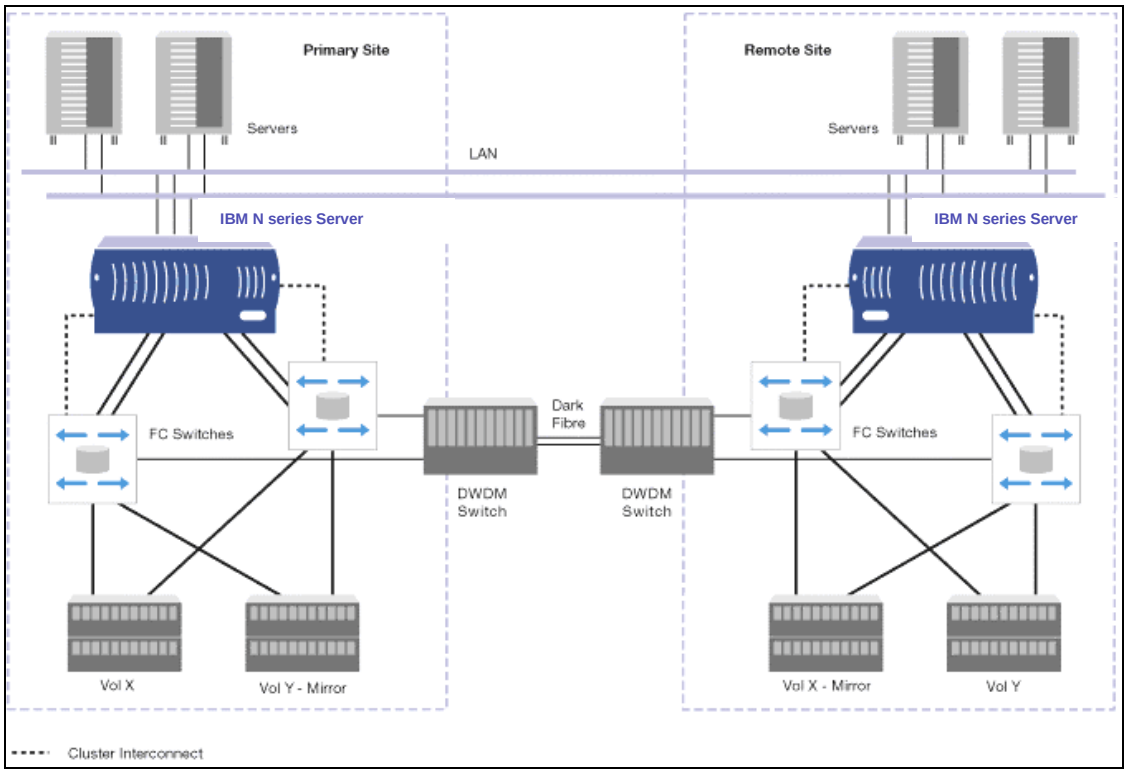

Figure 24. MetroCluster over FC and DWDM switch infrastructure.

This solution connects distant sites in metropolitan areas.

Conclusions:

- 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 the media type and platform for archival storage requires companies to conform to not just one but many content retention mandates. IT organizations must 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 almost 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 requires the purchase of new storage subsystem hardware and software.

Historically, most regulated data has been stored on optical devices, tape, paper, and/or microfiche/microfilm. According to Enterprise Storage Group (ESG), about 10% of regulated data is stored on disk today. Disk has not often been utilized, due to a number of factors that include cost and the lack of necessity to retrieve information quickly. However, ESG estimates that moving forward, 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), in which the archiving objects are used to remove data that is no longer needed in online business processes from the database 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 that of ADK-based archiving. The key difference is that it is based on universally accepted and widely used standards: XML format is used to archive business objects, hyper text transfer protocols (HTTPs) as a communication service, and 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 stands for Web-Based Distributed Authoring and Versioning. It is a set of extensions to the HTTP protocol that allows users to collaboratively edit and manage files on remote Web servers. The major features of the protocol are locking, metadata management, and namespace manipulation. IBM N series storage is certified by SAP using the WebDAV for Data Archiving interface. More information can be found at Partner Information Center. Once archive files have been created, the data marked for archive can be deleted from the source system. The archiving 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 IBM N series with NearStore feature 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 feature is the preferred compliance and archive storage subsystem for SAP landscapes. The NearStore product family leverages IBM N series technology introduced with Data ONTAP 7.1 and takes full advantage of value-added software from IBM N series such as SnapLock. An IBM N series with the NearStore feature scales from 7 to 504TB using economical ATA disk technology. With over 99.995% field-measurable uptime, IBM N series RAID-DP technology enables IBM N series with NearStore feature to tolerate single disk failures with no data loss.

IBM N series with SnapLock provides high-performance disk-based magnetic WORM storage. SnapLock provides secure, storage-enforced data retention functionality via open file protocols such as common internet file system (CIFS) and network file system (NFS) while leveraging 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 non-erasable, non-rewritable state, either forever or for a designated retention period. SnapLock allows this retention to be performed at the granularity of individual files through the 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 gatekeeper through which all file protocol and administrative access to the data must pass.

SnapLock is based on the 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 that are available on all clients. Applications use regular programmatic library interfaces used 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 in 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 their P8 platform, can take full advantage of the integration of these products 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 the IBM N series.

Conclusions: 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.
- Solution can incorporate existing IBM N series or other vendor's storage.
- Better ROI and lower TCO are achieved through increased availability, enhanced system performance, lower administration overhead, and increased staff productivity.
- Compliance and archived data remain easily accessible on the NearStore feature—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. FlexClone, FlexShare, SnapDrive, Snapshot, SnapRestore, SnapVault, NearStore, SnapMirror, RAID-DP, SyncMirror, SnapLock and FlexVol are trademarks or registered trademarks of Network Appliance, Inc., in the U.S. and other countries.

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

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

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

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

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

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