



Technical report: SAP on Microsoft Windows and Microsoft SQL Server with IBM System Storage N series

Best practices

• • • • • • • • •

Document NS3585-0

May 22, 2008



Table of contents

Abstract	4
Introduction	4
Scope	4
Business challenges	4
IT challenges.....	4
N series solutions for SAP	5
Storage provisioning and management	5
Consolidation	5
Storage layout.....	10
Aggregate layout	10
FlexVol volume layout	11
Layout with MetroCluster.....	12
LUN layout.....	13
Recommended standard layout	14
FlexVol volume layout for large SAP systems with high I/O requirements	14
LUN layout for large SAP systems with high I/O requirements.....	15
SQL Server datafile layout	16
Sizing	18
Installation.....	19
General requirements.....	19
N series storage controller configuration.....	19
Storage migration.....	19
Overview of migration approaches.....	19
Migration on the operating-system level	20
Migration on the database level	20
Migration on the storage-system level	20
ReplicatorX.....	21
System management and maintenance	21
SAP system cloning	21
Business challenges.....	21
N series solution	22
N series solution for SAP system copies	22
Necessary steps at the source system	23
Necessary steps at the target system	23
Conclusion.....	24
SAP upgrade.....	24
Business challenges.....	24
N series solution	25
N series solution for SAP upgrades	26
Upgrading the development system	26
Upgrading the quality assurance system	26
Upgrading the productive system.....	27



Conclusion.....	28
Business continuance	29
Backup and recovery	29
Business challenges.....	29
N series solutions	30
N series SnapManager for SQL Server	30
SnapManager for SQL Server	30
Conventional backup to tape	32
SnapManager for SQL Server: Snapshot backups together with NDMP.....	32
SnapManager for SQL Server: Snapshot backups together with disk-to-disk backup and SnapVault	32
Restore from tape or SnapVault restore	33
SnapManager for SQL Server: restore using SnapRestore.....	33
Conclusion.....	33
ReplicatorX SAP-consistent landscape backup.....	35
High availability	36
Business challenges.....	36
N series solution.....	36
N series solution for SAP high availability	36
Conclusion.....	37
Disaster recovery	37
Business challenges.....	37
N series solutions	38
SnapMirror	38
MetroCluster	39
ReplicatorX.....	41
Conclusions	42
Archiving and compliance	42
Business challenges	42
Archiving.....	42
Compliance	42
N series solution.....	43
SnapLock	43
SnapLock compliance.....	44
SnapLock Enterprise	44
Conclusion.....	44
Trademarks and special notices.....	45



Abstract

This document provides customers and partners with general best-practices guidance for deploying IBM System Storage N series in support of SAP Business Suite solutions running in a Microsoft Windows environment using a Microsoft SQL Server database. IBM N series minimizes IT challenges around enterprise business process application deployment and management, allowing SAP users to maximize application benefits.

Introduction

Scope

This document is intended to provide customers and partners with best practices for deploying IBM® System Storage® N series storage systems in support of SAP™ Business Suite solutions running in a Microsoft® Windows® environment using a Microsoft SQL Server 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 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 enterprise resource planning (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 their 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 requirements. These



requirements 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 or as many as five or more instances supporting it. Deploying three SAP applications, for example ERP, CRM (for customer relationship management), and BW (for data warehousing), 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 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.

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

IBM N series offers highly available, scalable, and cost-effective storage consolidation solutions that incorporate a 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. N series enterprise-class storage solutions are proven interoperable across all platforms. Systems integrate easily into a complex enterprise and simultaneously support network attached storage (NAS), Fibre Channel (FC) storage area networks (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

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.

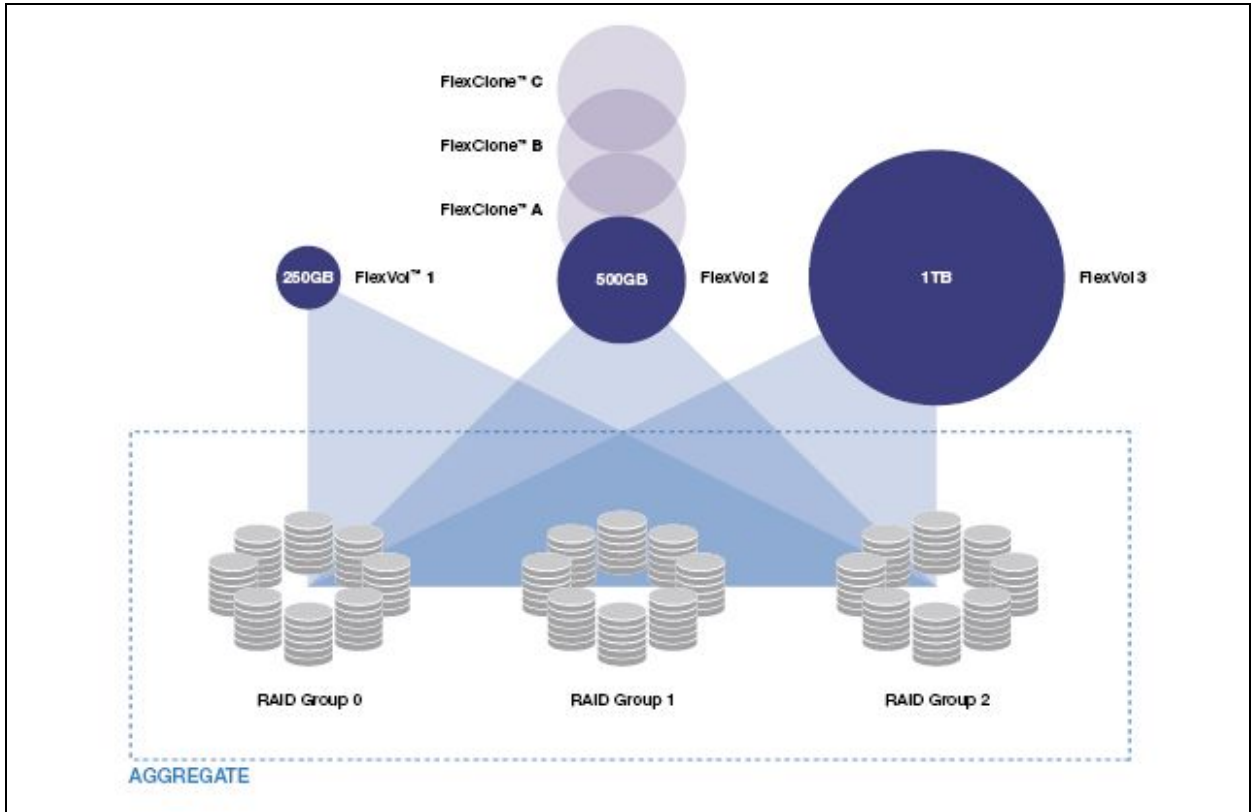


Figure 1) FlexVol technology.

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




Figure 2) FlexShare workload prioritization.

With IBM System Storage N series with SnapDrive® for Windows 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 NT file system (NTFS) 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 or FC access protocol. 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 application IBM System Storage N series with Snapshot™ backups. Snapshot backups are nondisruptive to applications and occur very quickly. Restoring data from a snapshot copy is nearly instantaneous.

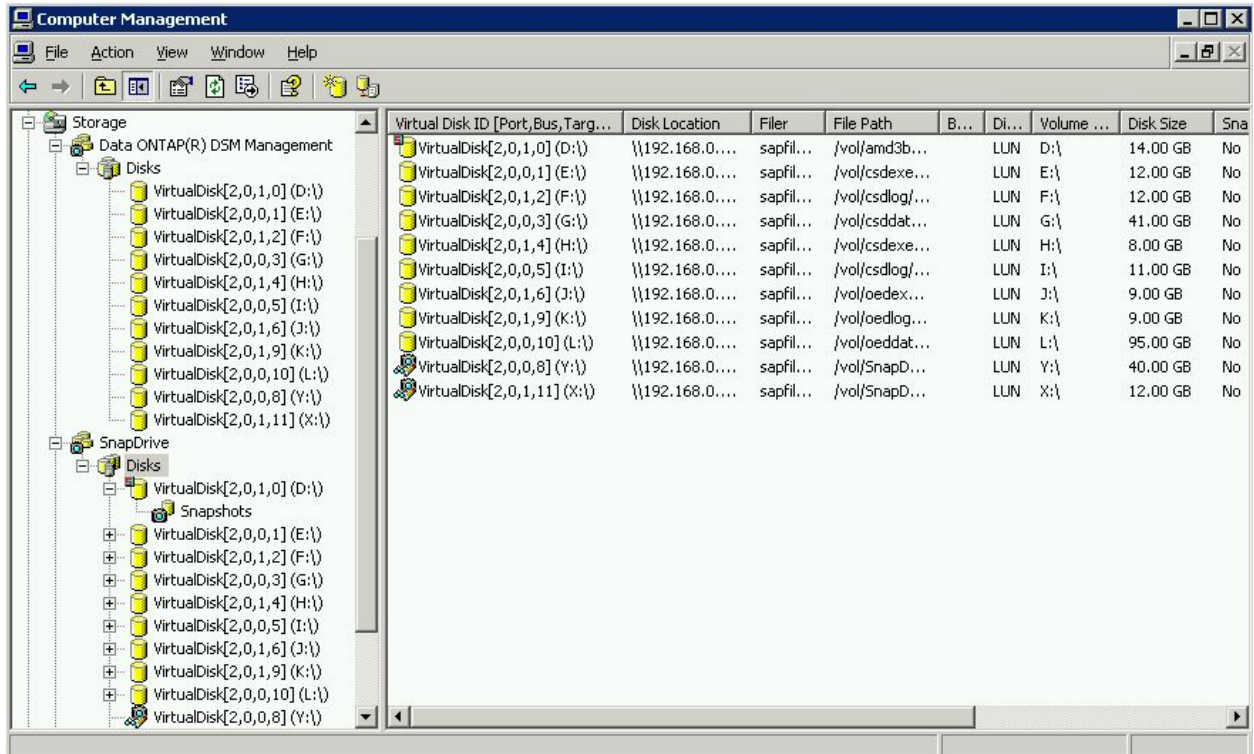


Figure 3) Storage management with SnapDrive and Data ONTAP.



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 connected using Gigabit Ethernet. The FC protocol facilitates storage access through a 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 N series storage systems. Thus a virtual disk can be created and accessed using the iSCSI or FC access protocol. Virtual disks are referred to as logical unit numbers (LUNs) when accessed over the iSCSI and FC protocols. It is possible to change the protocol from iSCSI to FC and vice versa without migrating the data.

The Data ONTAP DSM for Windows MPIO (DSM stands for device-specific module) enables you to manage multiple FC protocol paths between an N series storage system and a Windows host computer. If the active path fails, the DSM automatically switches to a passive path, maintaining the host's access to its storage. This product can be run as a standalone product or can be run with SnapDrive.

IBM System Storage N series with SnapManager® for SQL Server paves the way for database and storage administrators to simplify data management by utilizing the powerful capabilities of N series storage systems. SnapManager automates and simplifies the complex, manual, and time-consuming processes associated with the backup, restore, recovery, and cloning of SQL Server databases. It is integrated with Microsoft technology across both iSCSI and FC protocols to allow IT organizations to:

- Scale their storage infrastructure
- Meet increasingly stringent SLA commitments
- Improve the productivity of database and storage administrators across the enterprise.

SnapManager for SQL Server continues the drive to address the mission-critical needs of SAP Business Suite solutions.

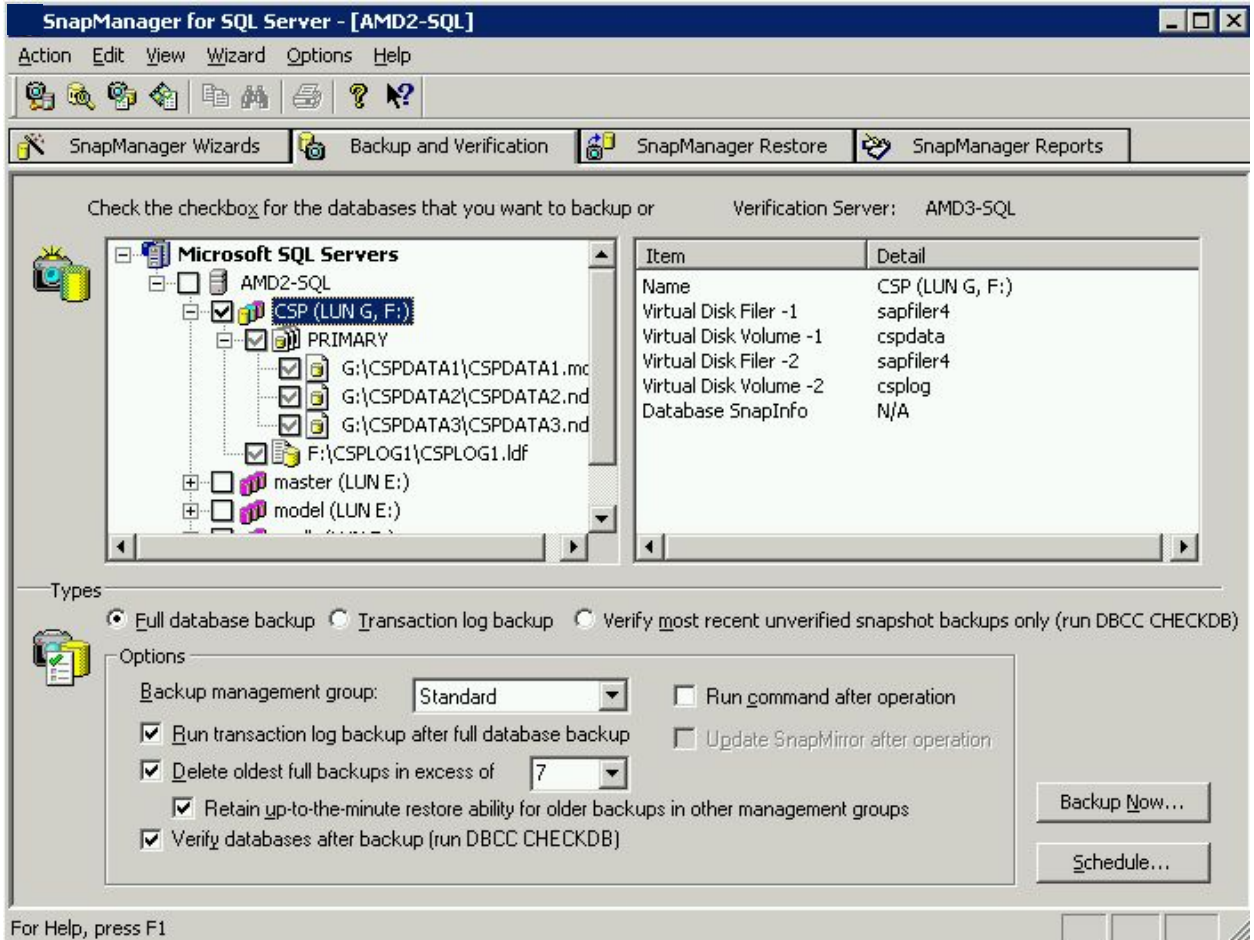


Figure 4) Database management with SnapManager for SQL.

The following list highlights some of the important benefits that 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

IBM N series works in combination with ReplicatorX™. ReplicatorX is enterprise-class replication software that works across platforms and over global distances. ReplicatorX software provides enterprise-wide data consistency and supports nondisruptive replication of production data to cost-efficient secondary storage. This enables comprehensive disaster recovery and block data migration. Plus, with N series FlexClone technology, you can also accelerate testing and development.

With ReplicatorX, you can continuously copy and protect SAP data—over any distance and across heterogeneous infrastructures—from multiple data centers to a target site over an IP network. ReplicatorX systems use lightweight host-based or SAN-resident client-side software that does not affect your SAP production system. At the target site, ReplicatorX software maintains an up-to-the-second replica of your

source data, which minimizes SAP data loss in the event of a disaster. ReplicatorX systems ensure data consistency among several SAP applications distributed across server and storage systems. Unlike conventional replication solutions, this unique capability eliminates complex and time-consuming data reconciliation and enables SAP landscape recovery in minutes.

ReplicatorX simplifies SAP data management and provides centralized visibility and control of all replication sessions. Administrators can quickly define policies, delegate access, and manage environments at any level, from a single volume to an entire geography. With ReplicatorX software, you can manage cloning operations through a single interface, with comprehensive and user-configurable historical reporting.

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 to every FlexVol volume in that aggregate. Adding a second aggregate is recommended only if the maximum capacity of the first aggregate is reached.

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

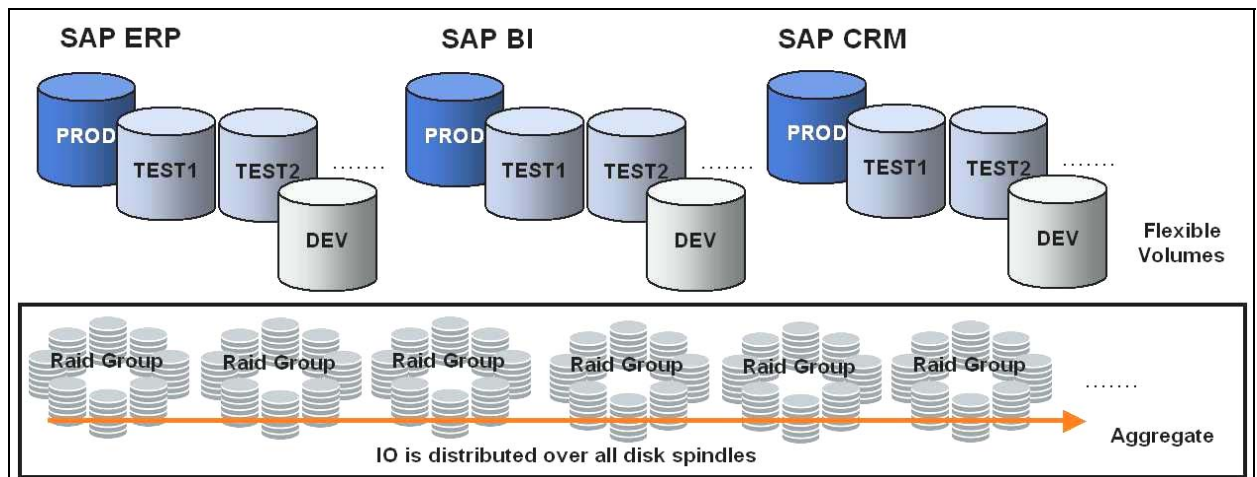


Figure 5) Aggregate layout.

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 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 benefit from the disk spindles of the test and development systems, which are necessary anyway because of capacity requirements.

With shared resources, contention for available resources among systems is always possible. 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 modified during online operation.

FlexVol volume layout

Each SAP system uses five FlexVol volumes:

- One volume for the database data files
- One volume for the database log files
- One volume for the executables, including SQL Server system databases
- One volume for SQL Server TempDB and the Quorum disk if Microsoft Cluster Server (MSCS) is used
- One volume for SnapManager for the SQL Server IBM System Storage N series with SnapInfo directory.

The SQL Server data files are separated into log files and log file backup (SnapInfo Dir). It is important to store the database data files in a FlexVol volume separated from the log files, to allow usage of snapshot copies, IBM System Storage N series with SnapRestore®, FlexClone, and other Data ONTAP features that work on the volume level. If log and data files are stored in the same FlexVol volume, using SnapRestore to restore the volume eliminates the ability to roll the database forward past the time of the snapshot copy, resulting in potential data loss.

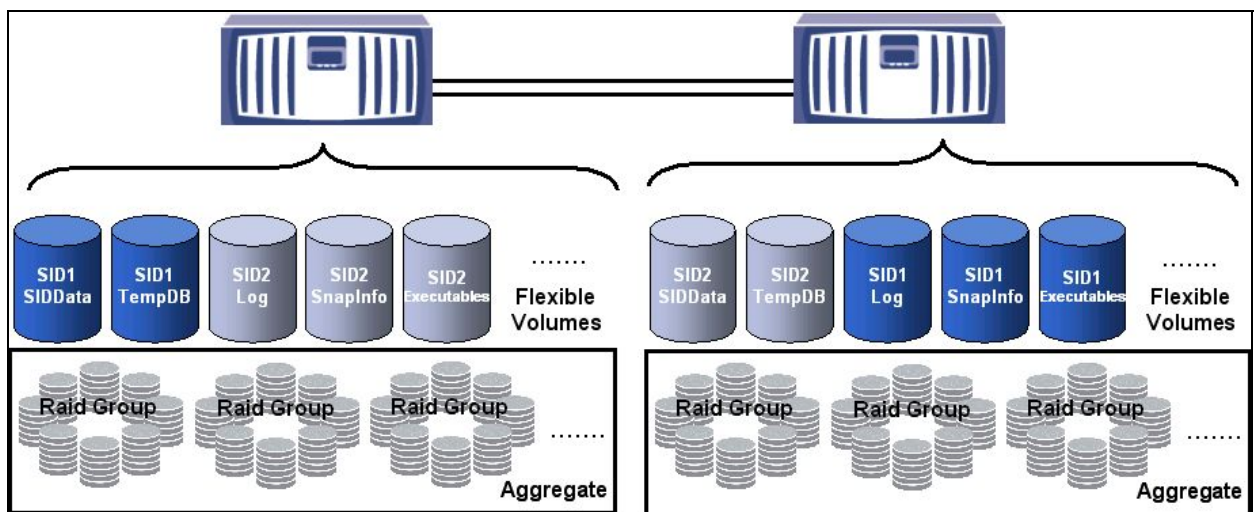


Figure 6) FlexVol volume layout.

The following table shows the distribution of the file systems and LUNs of a single SAP instance to the FlexVol volumes.

Storage Controller1		Storage Controller2		
Aggregate 0		Aggregate 1		
SIDData	TempDB	Executables	Log	SnapInfo
SIDData1	TempDB	SQL Executables	Log files	SnapInfo Dir
SIDData2	Quorum MSCS	System DBs		
SIDData3		SAP Executables		
SIDData#				

Table 1) FlexVol volume layout.

Layout with MetroCluster

IBM System Storage N series with MetroCluster and synchronous mirroring (using IBM System Storage N Series with SyncMirror®) work at 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 7) Storage layout with MetroCluster and SyncMirror.

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.

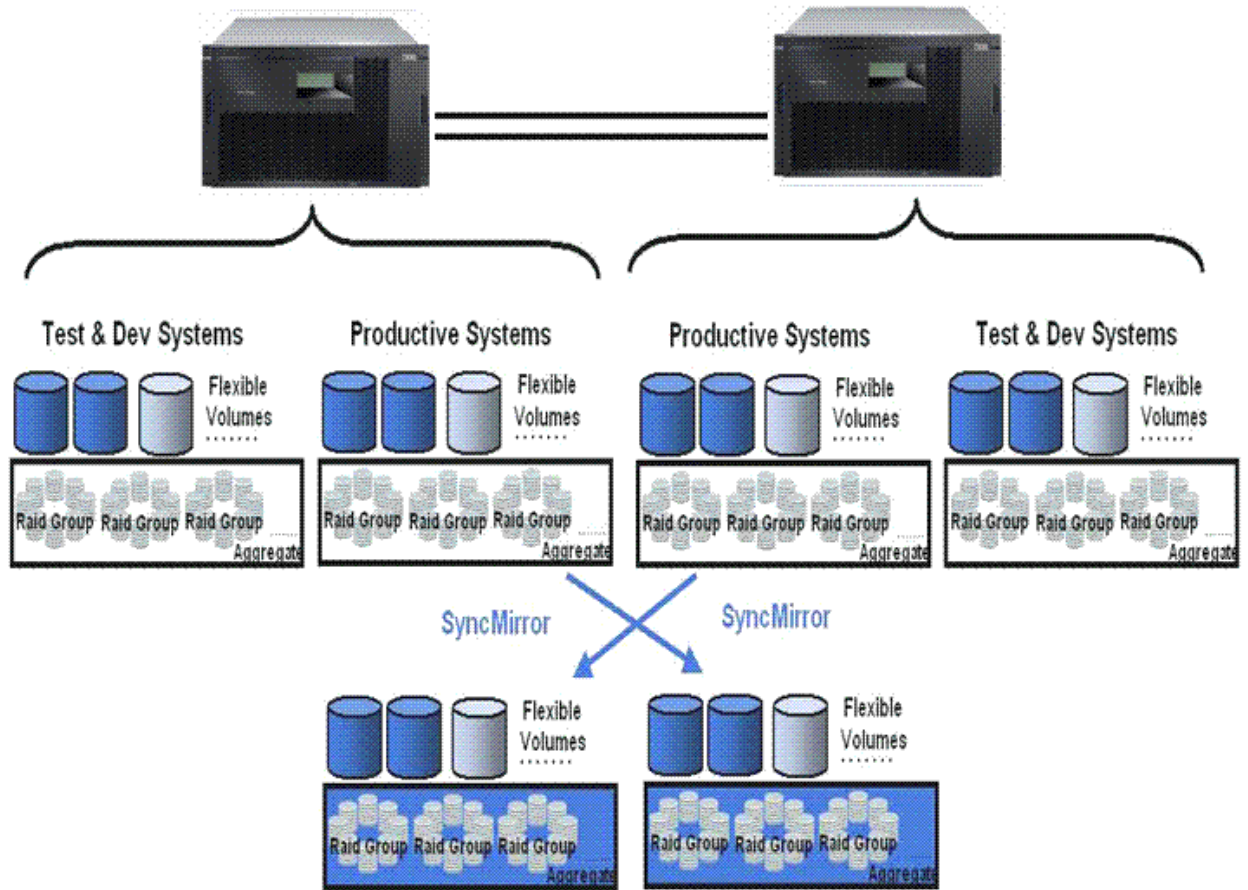


Figure 8) Storage layout with N series MetroCluster; only production systems are mirrored.

For more information on MetroCluster, see the document section discussing disaster recovery.

LUN layout

The size of the database and the type of SAP system determine the number and size of 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 data files	Size of LUNs data files	# of LUNs log files
Small System	< 200GB	1	100GB – 200GB	1
Medium System	200GB – 1TB	2-3	100GB – 350GB	1
Large System	> 1TB	> 3	200GB – 450GB	1

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. One LUN is configured in the SIDData FlexVol volume for the database data files and one LUN is configured for the log files in the Log FlexVol volume. Three LUNs in the Executables FlexVol volume are used to store the SAP and SQL Server binaries and SQL Servers system databases. Additionally, one LUN for the SQL Server TempDB is stored in the TempDB FlexVol volume and one LUN is configured in the SnapInfo FlexVol volume for the SMSSQL SnapInfo directory. The Quorum disk for a Microsoft Cluster should be stored in the TempDB FlexVol volume.




Figure 9) Standard storage layout.

FlexVol volume layout for large SAP systems with high I/O requirements

Systems with high input/output (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 prevents 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 six FlexVol volumes:

- Two volumes for the database data files
- One volume for the database log files
- One volume for the executables, including SQL Servers system databases
- One volume for SQL Servers TempDB
- One volume for SnapManager for the SQL Server SnapInfo directory and the Quorum Disk if MSCS is used.

The SQL Server data files are separated into the log files and log file backup (SnapInfo Dir). Each of these FlexVol volumes is stored in a separate aggregate. Therefore it is always possible to recover the database without data loss even if any of the aggregates is lost. Storing the database data files in a FlexVol volume separated from the log files is important to allow the use of snapshot copies, SnapRestore, FlexClone, and other Data ONTAP features that work on the volume level.

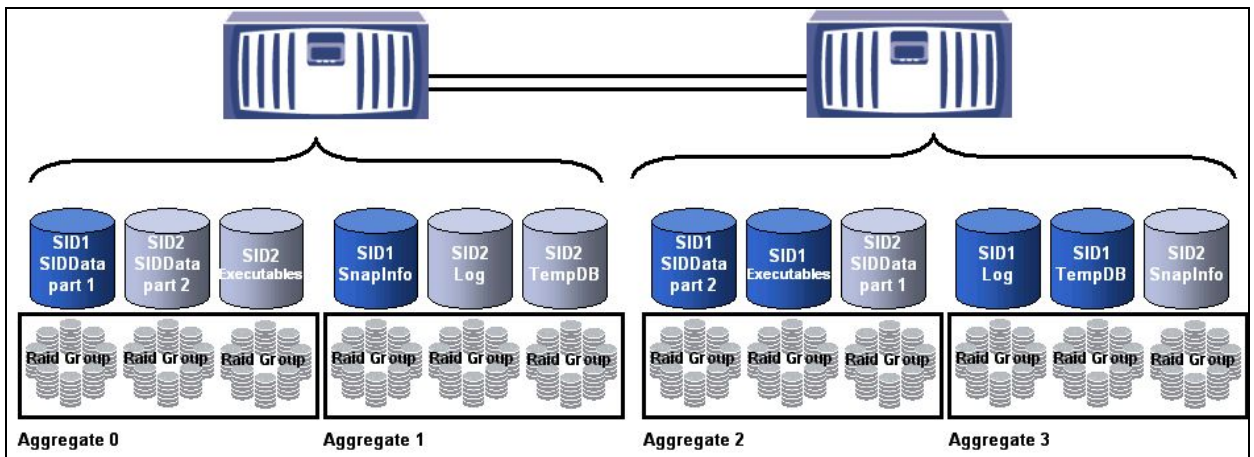


Figure 10) FlexVol layout for large SAP systems.

The following table shows the distribution of the file systems and LUNs of a single SAP instance to the FlexVol volumes. One LUN for each database file (x.mdf; x.ndf) is used, in this example a total of eight.

Storage Controller1		Storage Controller2			
Aggregate 0	Aggregate 1	Aggregate 2		Aggregate 3	
SIDData	SnapInfo	SIDData	Executables	Log	TempDB
SIDData1	SnapInfo Dir	SIDData2	SAP Exe	Log files	TempDB
SIDData3	Quorum MSCS	SIDData4	System DBs		
SIDData5		SIDData6	SQL Exe		
SIDData7		SIDData8			

Table 3) FlexVol layout for large SAP systems.

LUN 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. One LUN is configured for the log files in the log FlexVol volume. If possible, choose an independent I/O path to connect this LUN. In total, eight LUNs are configured for the database data files. Four LUNs are stored in the FlexVol volume SIDData on the first storage controller, and another four LUNs

are stored on the second storage controller. This approach allows load distribution to both storage controllers. To reduce the necessary number of drive letters, it is recommended that you mount the LUNs into an empty NTFS folder. To do so, create or connect the first LUN to a drive letter and create the necessary number of SIDDATx directories, in this example eight. Then create and mount the other seven LUNs to this newly created directory. For detailed information, refer to the appropriate SnapDrive installation and administration guide. When the FC protocol is used, it is recommended that you distribute the load of the data file LUNs to the available FC protocol HBAs or ports. With Data ONTAP DSM, the load-balancing option should be used.

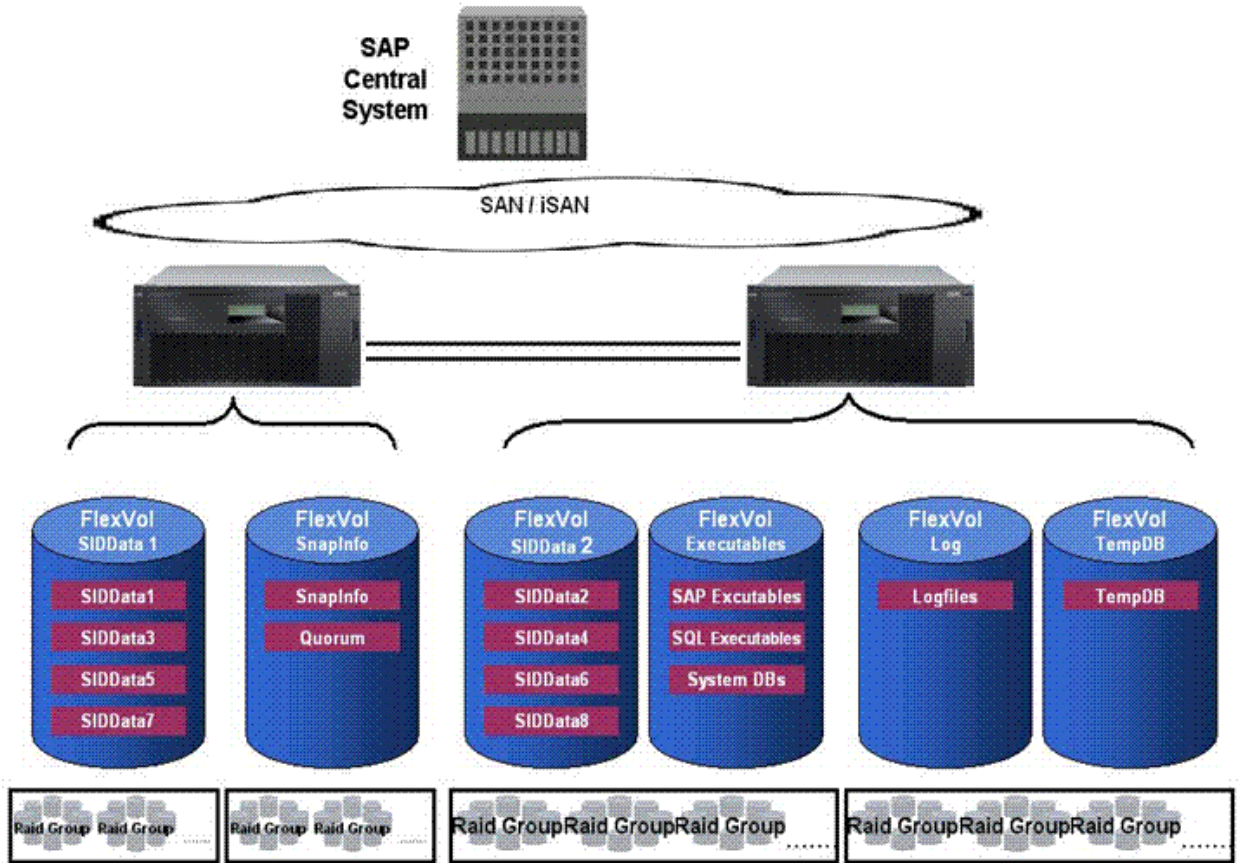


Figure 11) LUN layout for large SAP systems.

Three LUNs in the executables FlexVol volume are used to store the SAP and SQL Server binaries and SQL Servers system databases. Additionally, one LUN for the SQL Server TempDB is stored in the TempDB FlexVol volume and one LUN is configured in the SnapInfo FlexVol volume for the SMSSQL SnapInfo directory. The Quorum disk of a Microsoft Cluster should be stored in the SnapInfo FlexVol volume.

SQL Server datafile layout

The number of SQL Server data files for the SAP database should be equal to the number of CPU cores used in the database server. In general, the number of data files should not exceed 16. For



extremely large databases, up to 24 data files can be used. For large databases, it is recommended that you store each data file in its own LUN.

Sizing

This section gives an overview of the storage sizing for a SAP environment using 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 a 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 and recovery requirements (mean time to recover, backup window, retention policy)
- Cloning requirements (FlexClone copies or full copies)
- Disaster recovery requirements (synchronous or asynchronous mirroring)
- High-availability requirements (storage system clustering).

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

For existing SAP systems, the I/O requirements need to be measured using database or operating system tools. Database tools can be, for example, the SAP database performance monitor. For instance, Windows Performance Monitor can be used if the measurement is done on the operating system-level. Independently 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 values for the systems, which are provided by the SAP Quick Sizer (an internet-based tool from SAP), 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 is 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.

Based on the I/O requirements, the type and number of disk spindles and storage controllers is determined.

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 to support both the I/O and capacity requirements.



Installation

This section describes the requirements and the configuration for installing a SAP Business Suite or SAP NetWeaver system based on a Microsoft SQL Server database on a Windows server using the FC or iSCSI protocol.

General requirements

It is strongly recommended that you use SnapDrive for Windows, an N series host-based software product that simplifies storage management and provisioning in a SAP Windows storage environment. It is required for SnapManager for SQL, which provides host-consistent database snapshot copies, SnapRestore, and database cloning functionality.

N series storage controller configuration

Snapshot backups for database applications won't be consistent from the database point of view without the integration in the database management system. Therefore automatically scheduled snapshot copies on the storage level should be turned off on database volumes. SnapDrive for Windows automatically sets the correct options to the storage system volumes.

It is strongly recommended that all LUNs should be created and stored in qtrees to be able to use all desired software functionality, such as qtree-based IBM System Storage N series with SnapVault®.

Storage migration

This section discusses different storage migration approaches. If a migration includes a change of the operating system or database system, the migration can't be done solely at the storage level, and SAP migration tools must 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 of which migration approach fits best in a specific environment depends heavily on the acceptable downtime of the business application. Furthermore, the downtime depends on the amount of data to be migrated. In general, there are four approaches to storage migration of SAP data:

- Migration on the operating-system level
- Migration on the database level
- Migration on the storage-system level
- Use of ReplicatorX.



Migration on the operating-system level

In addition to the existing storage system, the N series storage system is connected to the database server. The 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 is then copied via the server from the old storage system to the 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 must be configured using SQL Server Management Studio.

A migration on the operating-system level can be done for an FC-to-FC protocol or an FC-to-iSCSI protocol migration. The disadvantage of this approach is that the SAP system is not available while the database files are being copied. Depending on the database size, the downtime could be several hours. SnapManager for SQL Server offers a wizard to perform this type of migration.

Migration on the database level

A SQL Server database backup is restored to the 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 N series storage. In addition, the log backups are continuously copied to the separate server. Before the final migration is started, the SAP system must be shut down and a log backup must be performed. The log backups that have not yet been copied are now copied from the old storage system to the N series storage system and a forward recovery of the database is carried out. Then detach the databases using SQL Server Management Studio, disconnect the old storage from the database server, and connect the N series storage system and the LUNs that are stored there.

Attach the database stored on the N series LUNs. It is possible to change the file system structure with this procedure.

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

Migration on the storage-system level

For data migration from a non-N series SAN to an N series SAN storage system, there is a fast, flexible SAN data migration service to suit individual customer requirements. Essentially, the service uses a separate, secondary storage system and host software, which hosts data-copying software. The secondary storage is attached to the existing SAN as well as to the N series storage system. The secondary storage and software must be configured to see the existing and new arrays, and then enabled to perform the migration; for example, security zones and domain reconfiguration. The secondary storage is configured to migrate the appropriate data sets and/or data files. The software automatically transfers the data from the source to the replacement (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 complete, the servers and storage infrastructure must be reconfigured to permit the servers to see the storage on the new N series storage controller. Only after all of the servers can access the new storage can the migration-enabling secondary storage 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 secondary storage described earlier for the migration process.

ReplicatorX

As noted previously, ReplicatorX is a replication and recovery solution for heterogeneous storage and server infrastructure. It offers the possibility to migrate seamless from any block-based storage to N series iSCSI or FC protocol storage and therefore overcomes the limitations of other storage-system-level migration tools.

The following table below 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 the whole copy process	None
Migration on the database level	FC-to-FC and FC-to-iSCSI	Medium during reconfigure and forward recovery	Server temporarily connected to N series storage
Migration on the storage-system level	FC-to-FC	Low during reconfigure	Secondary migration-enabling storagesystem
Using ReplicatorX	FC-to-FC and FC-to-iSCSI	Low during reconfigure	Server

Table 4) Summary of migration processes.

System management and maintenance

SAP system cloning

Business challenges

A typical SAP customer environment today consists of different SAP 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 a 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.

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 create a SAP system copy on demand is very quickly becoming more and 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 systems with current productive data must be set up quickly during a SAP upgrade project.
- Training systems must be set up or refreshed.

The traditional approach to create the system copies is not suitable to meet these demands.

SAP copies also consume a significant amount of storage. Because these copies are typically clones of the productive system, the amount of storage needed can be very large.

N series solution

The N series solution for SAP system cloning addresses these issues by providing an easy process to create a SAP system copy on demand, in a few minutes, without any impact to the source system. In addition, N series cloning functionality offered by SnapManager for SQL Server allows very efficient storage management by storing only data changes between the source and the clone.

N series solution for SAP system copies

SAP system copies are accomplished using SnapManager for SQL Server utilizing the N series FlexClone feature to clone SAP databases. A FlexClone copy is a writeable point-in-time image of an N series FlexVol volume. Based on a snapshot copy of the source FlexVol volume, the clone is created in a few seconds without interrupting the operation on the source system. FlexClone copies store only changed blocks between the source FlexVol volume and the FlexClone copy, and therefore significantly decrease the amount of disk space needed for SAP system copies. SAPINST (WebAS 7.0 or later) or the SAP STM tool (WebAS 6.x) is used to perform the SAP-specific task for a homogenous system copy.

The next figure 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 SAP database is created using SnapManager for SQL Server. This is done during online operation and has no performance impact on the source system. In case of a Java or ABAP/Java system, SAPINST must be executed to export the central instance to the Migration Export CD image before the backup is performed. If possible, store the image on a LUN that contains database files. After the restore described below is finished, the image is automatically available on the destination. On the source system, the previously created snapshot copy is used to perform a redirected restore to the target system. SnapManager for SQL Server creates a FlexClone image of the database LUNs and connects these LUNs to the target system. The database ID can be changed during the restore process.

On the target system, SAPINST must be executed to perform the SAP-specific tasks. In addition, SAP-specific post-processing tasks must be accomplished.

This solution can be used to refresh existing systems or to create new SAP systems and can be accomplished in a few minutes.

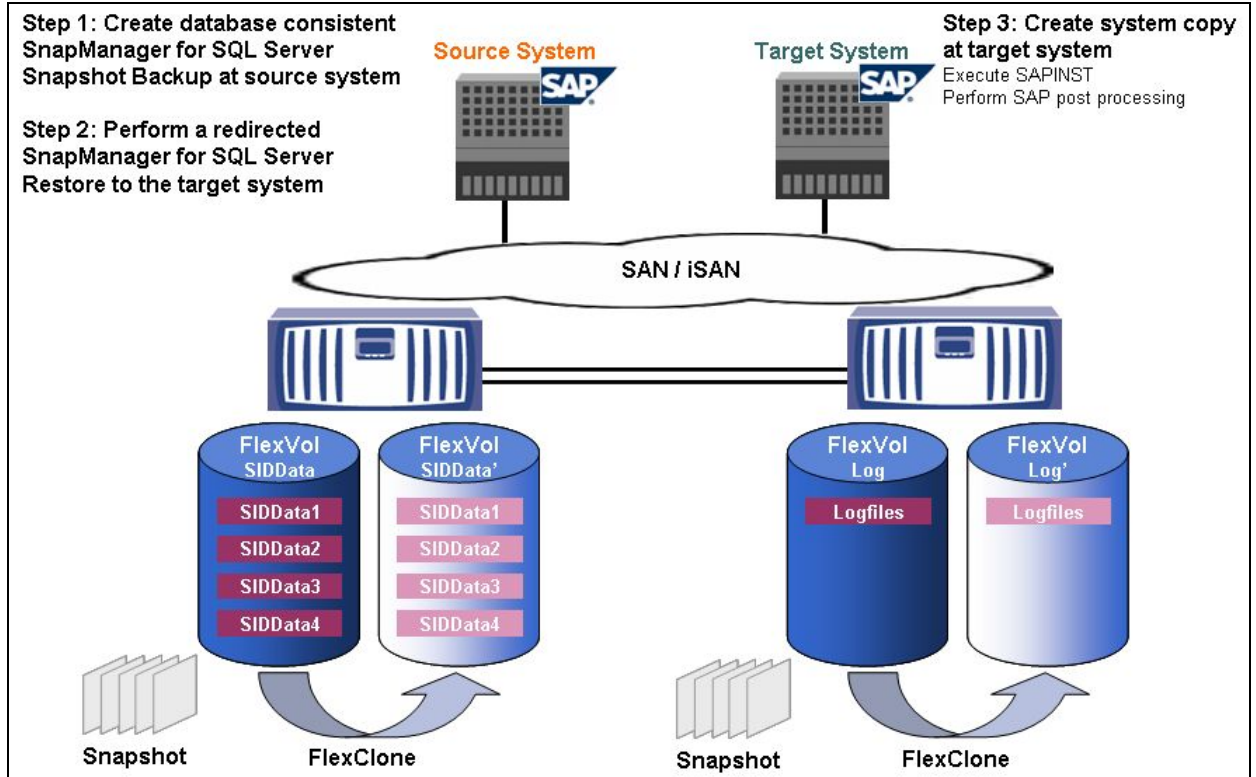


Figure 12) SAP system cloning overview.

Necessary steps at the source system

With the traditional approach, it is necessary to create a 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 can take several hours. The subsequent steps are typically carried out manually, consuming IT staff time.

With the N series approach, the backup is taken using N series Snapshot functionality. Creating a snapshot copy with SnapManager for SQL takes only a few seconds and has no performance impact on the source system. Therefore this step can be done at any time during online operation.

Necessary steps at the target system

If a new SAP test system needs to be set up, the Microsoft SQL Server must be installed once. This step is required with both approaches. With all subsequent refreshes of the system, this step is not necessary.

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. The following step to adapt the database to the new SID (system ID) is typically carried out manually.

With the N series approach, SnapManager for SQL Server creates a FlexClone image of the database based on the consistent snapshot database 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 step to adapt the database to the new SID is automatically done by SnapManager for SQL Server.

In both scenarios, SAPINST must be executed to perform the necessary SAP-specific tasks. Additionally, post-processing tasks must be performed.

Table 5 compares the traditional approach to the N series approach to perform a SAP system copy.

Traditional Approach	N series Approach	Advantages
Necessary steps at the source system		
<ul style="list-style-type: none"> Offline or online backup 	<ul style="list-style-type: none"> Snapshot backup during online operation using SnapManager 	<ul style="list-style-type: none"> No impact on operation with N series solution. Can be scheduled at any time.
Necessary steps at the target system		
<ul style="list-style-type: none"> Install MS SQL Server (if not yet existing) Restore backup from source system Change database name SAP-specific post-processing tasks 	<ul style="list-style-type: none"> Install MS SQL Server (if not yet existing) Create a FlexClone image based on a snapshot backup; mount the LUNs at the target system Change database name SAP-specific post-processing tasks 	<ul style="list-style-type: none"> Same approach. Restore takes several seconds with FlexClone technology used by SnapManager for SQL Server Fully integrated in SnapManager for SQL Same approach

Table 5) SAP system copy comparison.

Conclusion

The 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 performed at any time because there is no impact to the online operation of the source system.
- Snapshot and FlexClone functionality offered by SnapManager for SQL Server reduces the time necessary to copy the data from the source system to the target system from several hours to several seconds.
- Snapshot and FlexClone 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 the pressure of upgrading 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, because all development must be stopped and SAP support packages cannot 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 upgrading the productive system is important. Database backups consume extensive time. Optimizing backup and restore functionality is therefore very critical.

During a SAP upgrade project, SAP basis administrators must create several system copies to run the upgrade with current data from the development or productive SAP system. The creation of a SAP system copy normally takes several days and negatively affects the productive environment.

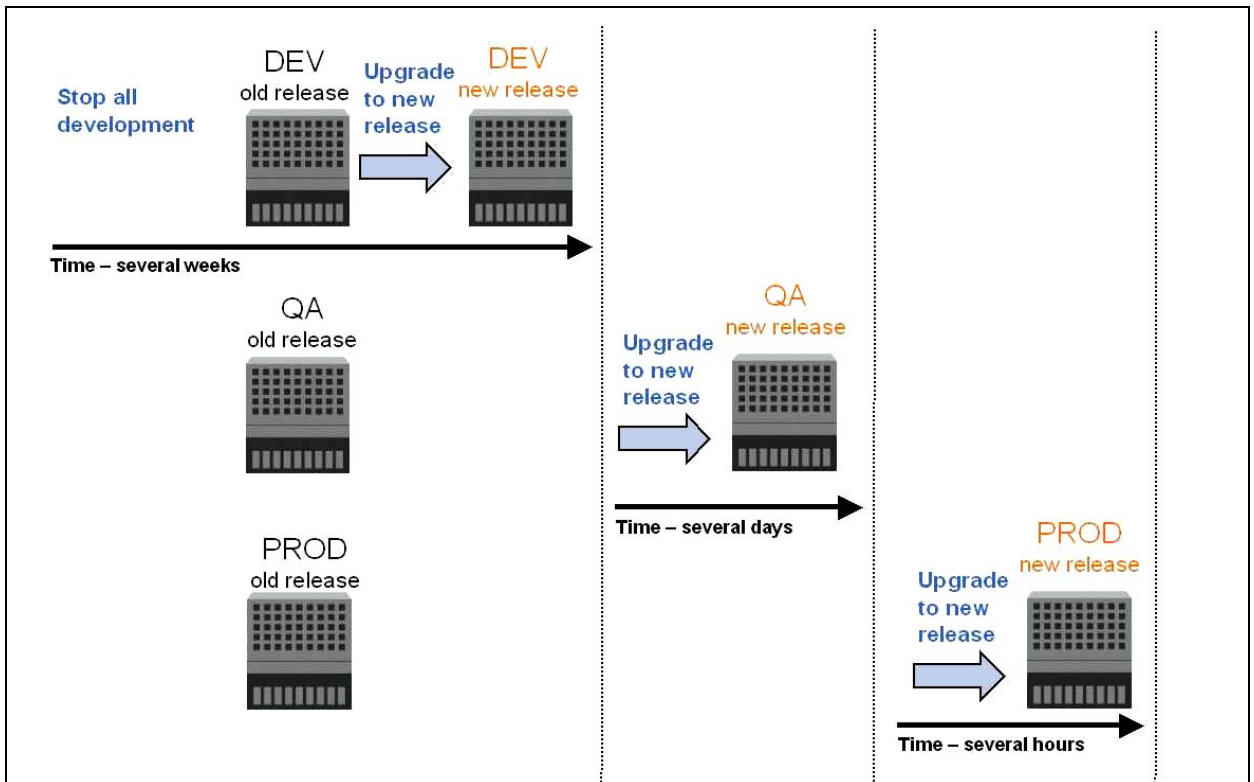


Figure 13) SAP upgrade overview.

N series solution

The 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 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 N series solutions help to minimize the risk, reducing the downtime and reducing overload of IT staff resources during a SAP upgrade project. Because of this, the N series solution offers the possibility to increase the amount and the parallelization of testing cycles during an upgrade project using the same or less time and resources compared to the traditional approach. This reduces dramatically the risk of an unexpected problem occurring during the upgrade of the production system and accelerates the upgrade project overall.

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 N series system copy solution described in the earlier section on the N series solution for SAP system copies.

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.

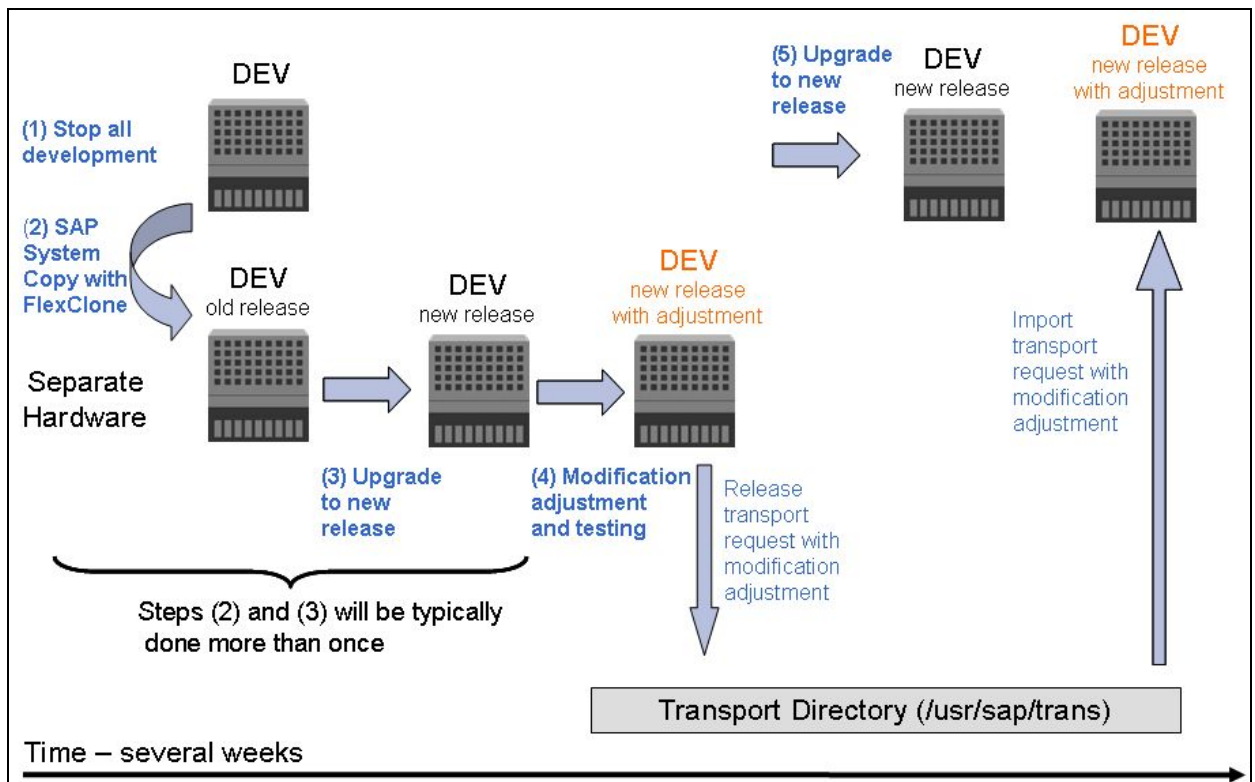


Figure 14) SAP upgrade—development system.

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 N series 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 restoration of the system to any specific snapshot copy, allowing a restart of an upgrade phase or restarting the import.

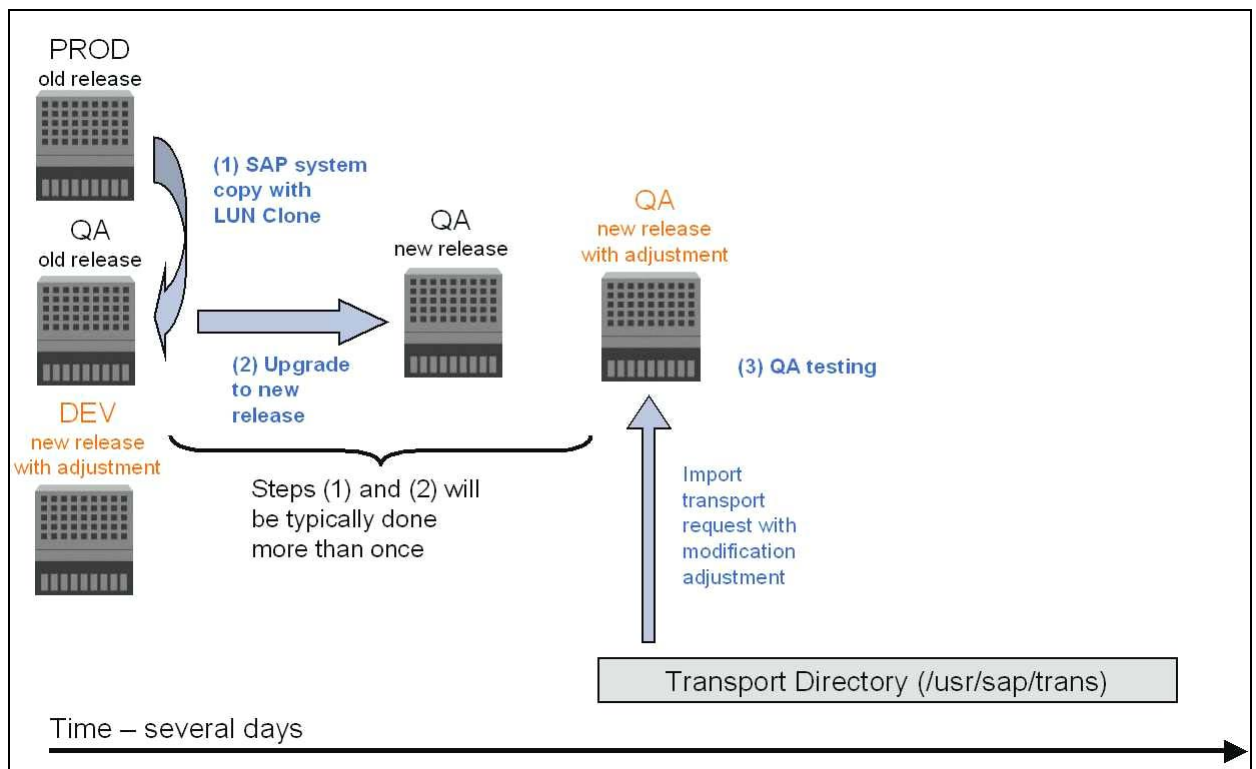


Figure 15) SAP upgrade—quality assurance system.

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. The schedule 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.

SnapManager for SQL offers snapshot copies as a backup method and SnapRestore for restoring the system to its former release status, ensuring a higher level of scheduling flexibility. Normal tape backups take several hours, which must be considered when planning the upgrade schedule. This time is reduced to several minutes when using SnapManager for SQL Server. The N series solution offers additional flexibility during this upgrade. If an unexpected error occurs, more time is available to fix the issue before the system must be reverted to the previous phase or restored to the old release, resulting in rescheduling the upgrade.

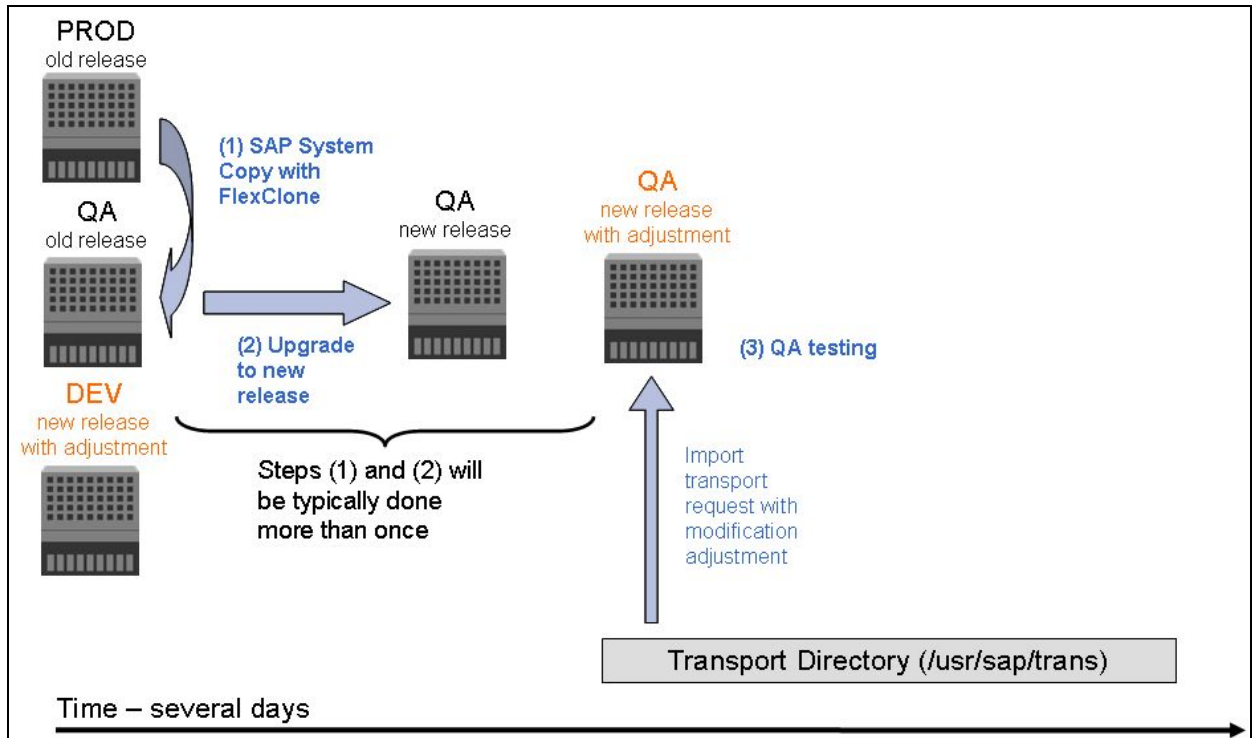


Figure 16) SAP upgrade—productive system.

Conclusion

The N series system copy and backup and recovery solutions significantly improve the SAP upgrade process.

- The N series system copy solution allows quick refreshing of the separate development and quality assurance systems, reducing the time required from several days to several minutes.
- With the same or less time and resources, more parallel testing cycles are possible, reducing the overall risk of the upgrade projects.
- Using N series SnapManager for SQL Server, 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.
- N series SnapManager for SQL Server significantly reduces the total upgrade time of the productive system, providing a higher level of flexibility for scheduling and executing 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 systems 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 only be taken 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 a SAP system always causes a financial impact on the business. A significant part of the unplanned downtime is the time that is needed to restore and recover the SAP system in 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 a SAP upgrade always includes at least three backups of the SAP database. The time needed to perform these backups cuts down the total time available for the upgrade process.



N series solutions

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

N series SnapManager for SQL Server

N series SnapManager for SQL Server is able to create SQL Server database backups in seconds using N series Snapshot technology. The time needed to create a snapshot copy is independent of the size of the database because N series Snapshot does not move any data blocks. The use of N series Snapshot technology doesn't have any performance impact on the productive SAP system, because the N series Snapshot implementation doesn't have to copy data blocks when the data in the active file system is changed. Therefore SnapManager for SQL Server backups can be scheduled without having to consider peak dialog or batch activities periods. SAP and N series customers typically schedule several snapshot online backups during the day—for instance, every four hours.

SnapManager for SQL Server

SnapManager for SQL Server also provides key advantages for the restore and recovery operation. The N series SnapRestore functionality integrated into SnapManager for SQL Server allows restoration of 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. As a result of several snapshot backups having been created during the day, the time needed for the recovery process is also dramatically reduced. Fewer log backups need to be applied, because a restore can be done to a snapshot copy that is at most four 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 SnapManager for SQL Server utilizing SnapRestore. Restores from a secondary location (disk or tape) are necessary only if the primary storage system that is storing 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 two-week-old backup.

A backup and recovery solution using a N series SnapManager for SQL Server always consists of two parts:

- Backup and restore/recovery using N series Snapshot and SnapRestore through SnapManager for SQL Server
- 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 by SnapManager for SQL Server. 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® system or other storage system and SnapVault software: The primary storage communicates directly with the secondary storage (N series with NearStore option or other storage system) and sends the backup data to the destination. The N series SnapVault functionality offers significant advantages compared to tape backups. After an initial data transfer, in which all the data is transferred from the source to the destination, all subsequent 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. SnapManager for SQL Server automatically starts the SnapVault transfer of changed blocks during a backup.

Backup to tape using third-party backup software: For NDMP (network data management protocol) backup (serverless backup), the tape is connected directly to the primary storage system. The data is written to tape using NDMP.

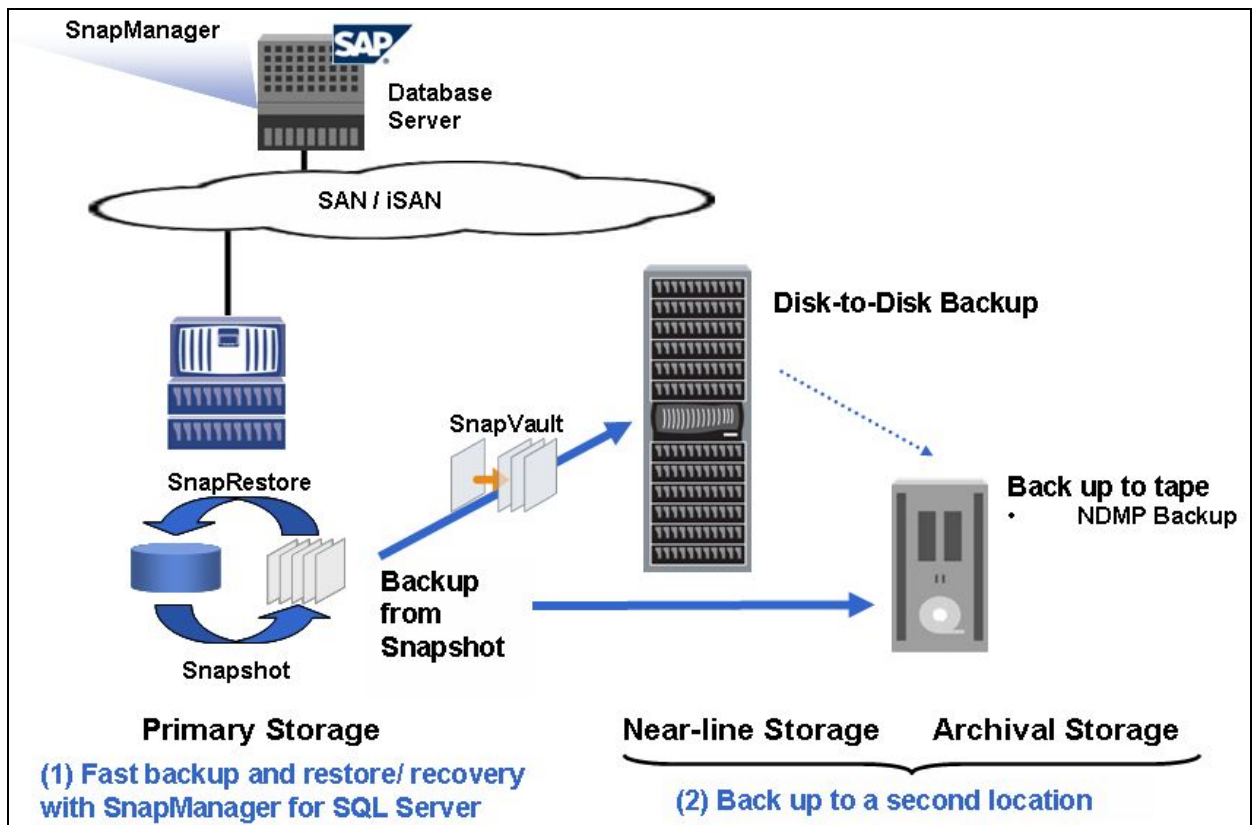


Figure 17) SAP backup and recovery—N series solution overview.

The following figure compares the different backup approaches with regard to the performance impact of a backup and the time that the database needs to perform the backup (backup mode).




Figure 18) Comparison of different backup approaches.

Conventional backup to tape

A conventional backup to tape generates a significant load on the productive database server and the primary storage system. Because this backup is not based on snapshot copies, the database is affected through the backup during the whole backup time. A full backup to tape is typically scheduled once a day.

SnapManager for SQL Server: Snapshot backups together with NDMP

Snapshot backups do not generate any load on the database server or the primary storage system. The database time needed for the backup is only a few seconds. A full database backup based on snapshot copies consumes disk space only for changed blocks. Snapshot backups are typically scheduled more often, for example, every four 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 is scheduled once a day. This backup still creates a high load on 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 is not affected while running the backup to tape.

SnapManager for SQL Server: 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 significantly reduced. 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 to perform a restore. Because the backup frequency is typically once a day, a certain amount of log backups must be applied after the restore is finished.

SnapManager for SQL Server: restore using 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 generally taken more frequently, such as every four hours, so the forward recovery is much faster because fewer log backups need to be applied.

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

Conclusion

The combination of N series 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 N series with NearStore option system).

Database verification is an important part of database maintenance. Snapshot backups are perfect for running a database consistency check (DBCC). SnapManager provides the capability to run a DBCC on a separate server automatically or manually after a SnapManager for SQL Server backup without creating any load on the productive database system.

The ability to easily create backups in seconds and to restore the SAP system to a point in time based on any available snapshot copy is also very useful 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 repeat that phase.

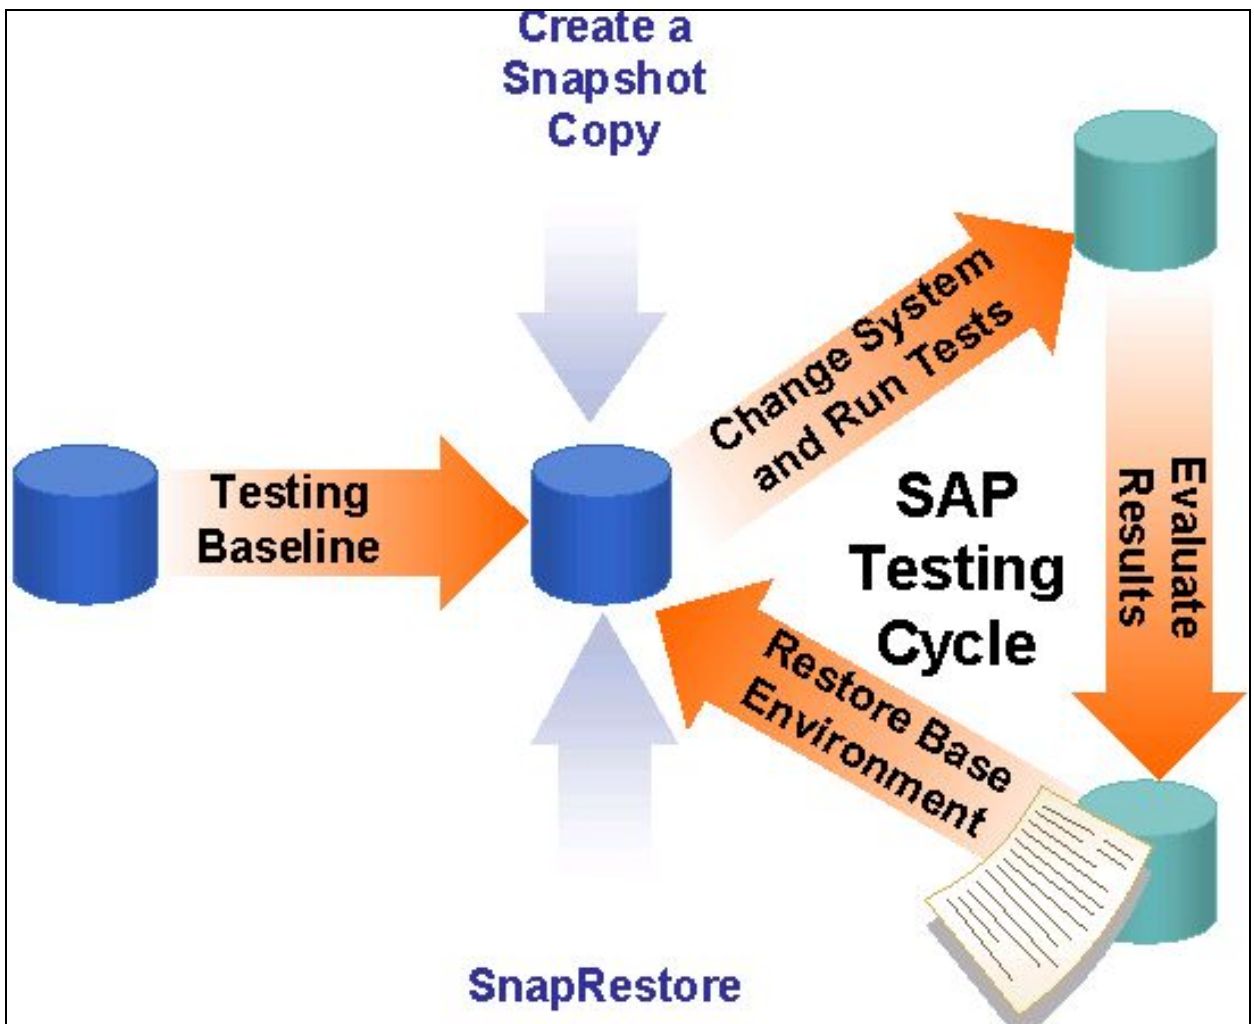


Figure 20) SAP testing cycle.

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 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 be



solved only by restoring the system and restarting the upgrade. In this respect, SnapManager for SQL Server 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, 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 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. SnapManager for SQL Server provides N series Snapshot technology as a backup method, and SnapRestore for restoring the system to its former release status. This provides a higher level of flexibility for 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.

ReplicatorX SAP-consistent landscape backup

Today's SAP Business Suit and SAP NetWeaver solutions consist of several components, which are tightly integrated by data transfers. A business process is processed in several different systems. To offer data consistency across the entire landscape, all SAP system must be backed up at exactly the same point in time.

As already noted, ReplicatorX is a replication and recovery solution with the unique functionality of guaranteeing data consistency between several systems. Based on a global clock, every I/O gets a time stamp, which allows the creation of a consistent image of the complete SAP landscape, also known as SAP Federated Backup. FlexClone functionally is integrated with in ReplicatorX, which allows cloning a complete SAP landscape, for example for QA and test purposes or disaster recovery tests. The replication and consistency group settings are easy to configure with the ReplicatorX graphical user interface (GUI).

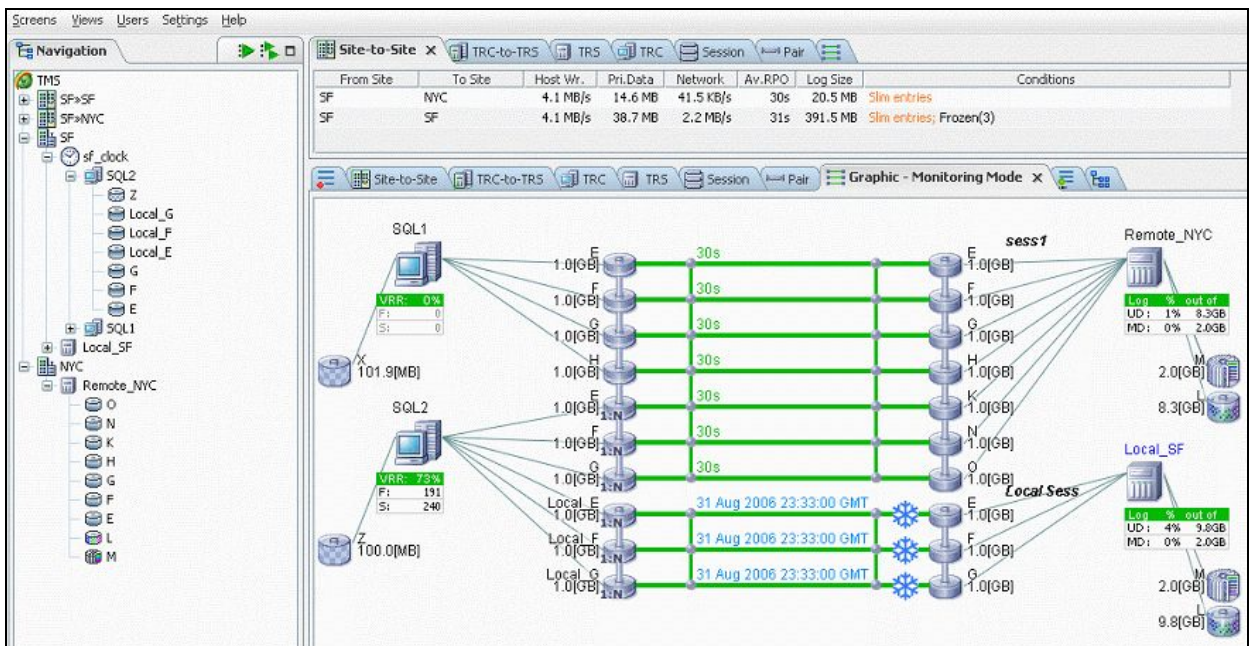


Figure 21) ReplicatorX GUI.

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.

N series solution

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

N series solution for SAP high availability

The following figure shows a sample clustered failover configuration. A cluster is created with two storage controllers by connecting the storage controllers via a cluster interconnect. This connection is fully redundant and is used to exchange cluster heartbeats and to synchronize the 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 WWPN of the first storage controller are also adopted. To ensure that no data is lost, the NVRAM is mirrored on both storage controllers via the cluster interconnect.

As a result of both storage controllers being 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.

SAP Central Instances, Application Servers, Database Servers




Figure 22) N series clustered storage system solution.

Conclusion

The N series clustered failover technology provides an extremely robust high-availability solution.

- A cluster has an availability level of 99.999+%.
- 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 (MSCS) 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 disaster recovery solution that utilizes one or many N series products.

N series solutions

SnapMirror

IBM System Storage N series with SnapMirror® software delivers a disaster recovery solution that today's global SAP systems need. By replicating data at high speeds over a local area network (LAN) or a wide area network (WAN), SnapMirror software provides the highest possible data availability and the fastest recovery. SnapMirror is integrated with SnapManager for SQL Server.

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

SnapMirror performs an initial transfer to initialize the disaster recovery site. After the initial transfer, incremental changes are passed to the disaster recovery 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 N series backup and recovery solution SnapManager for SQL Server. Selective snapshot backups are mirrored to the disaster recovery site. Additionally, the LUN where the log backups are stored (SnapManager for SQL Server SnapInfo Directory) must be mirrored using SnapMirror. It is recommended that you use a frequent SnapMirror update of the archive logs, such as every 10 minutes, enabling a minimum of data loss. This can be achieved by scheduling SnapManager log backups every 10 minutes.

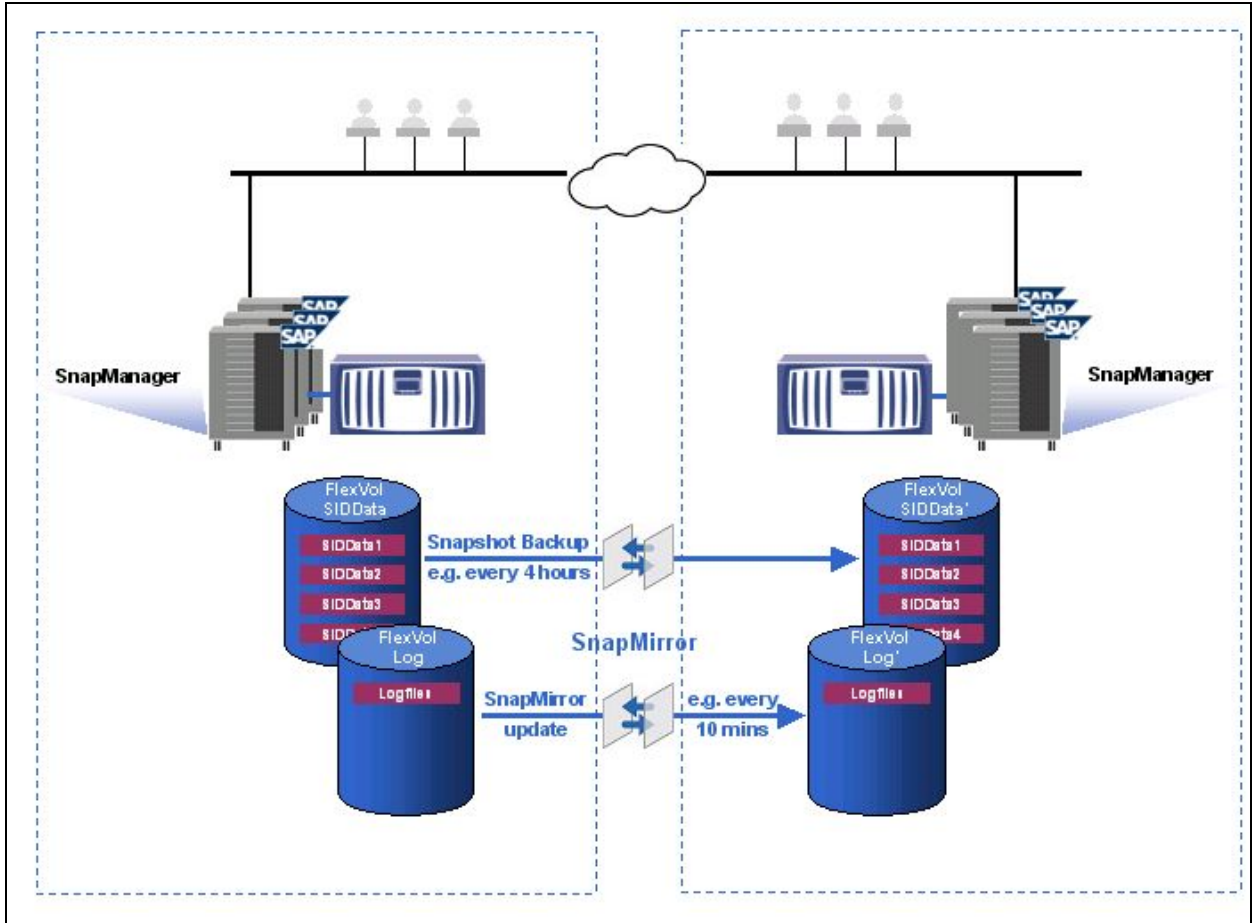


Figure 23) Disaster recovery with SnapMirror.

MetroCluster

N series MetroCluster is an integrated high-availability and business continuance solution that provides disaster recovery 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, that the data there is completely up to date. The combination of failover and data replication offers the possibility that you can recover from a disaster—with no loss of data—in minutes rather than hours or days.

MetroCluster is similar to N series clustered failover, but with the added benefit of disaster recovery. Clustered failover creates a cluster of N series systems 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 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 300 meters or less, then a Stretch MetroCluster configuration can be used with OM3 cabling. If the distance is greater than 300m, then a Fabric MetroCluster must be chosen, requiring FC switches as part of the solution.

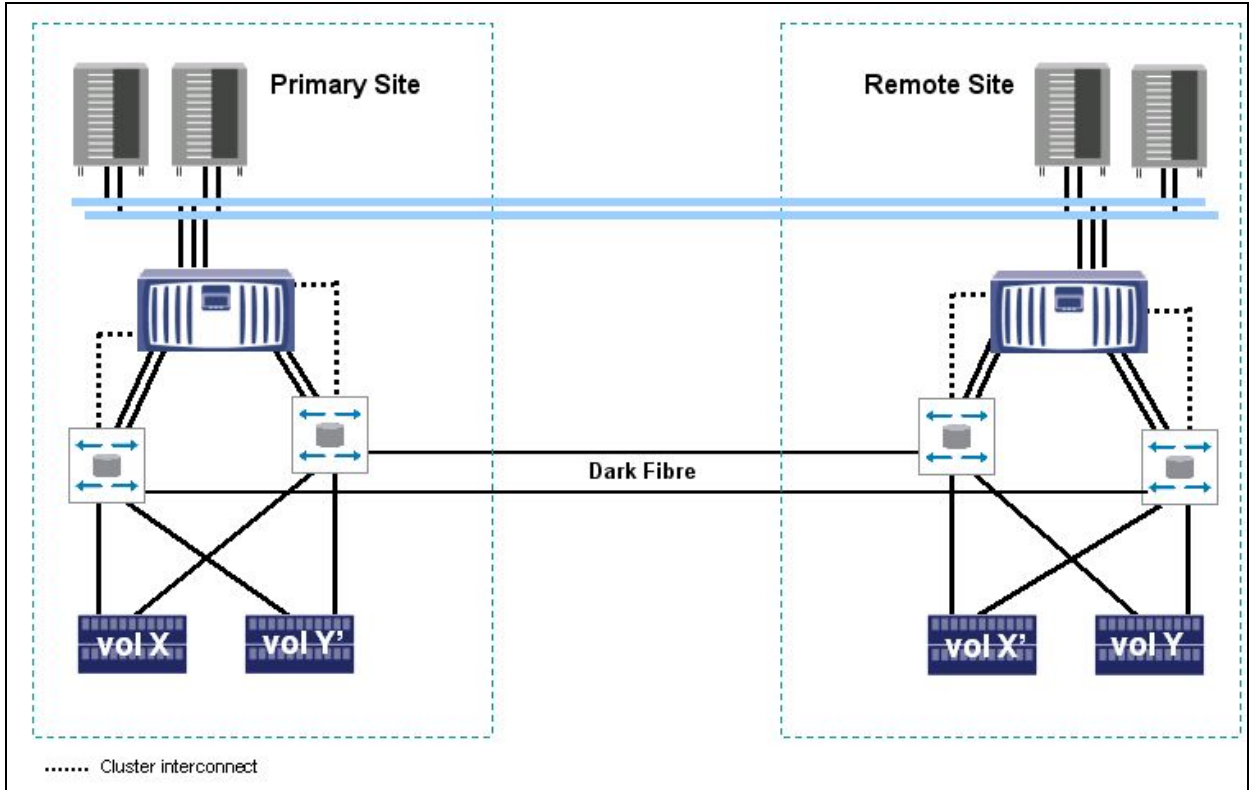


Figure 24) MetroCluster over direct FC switch connection.

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

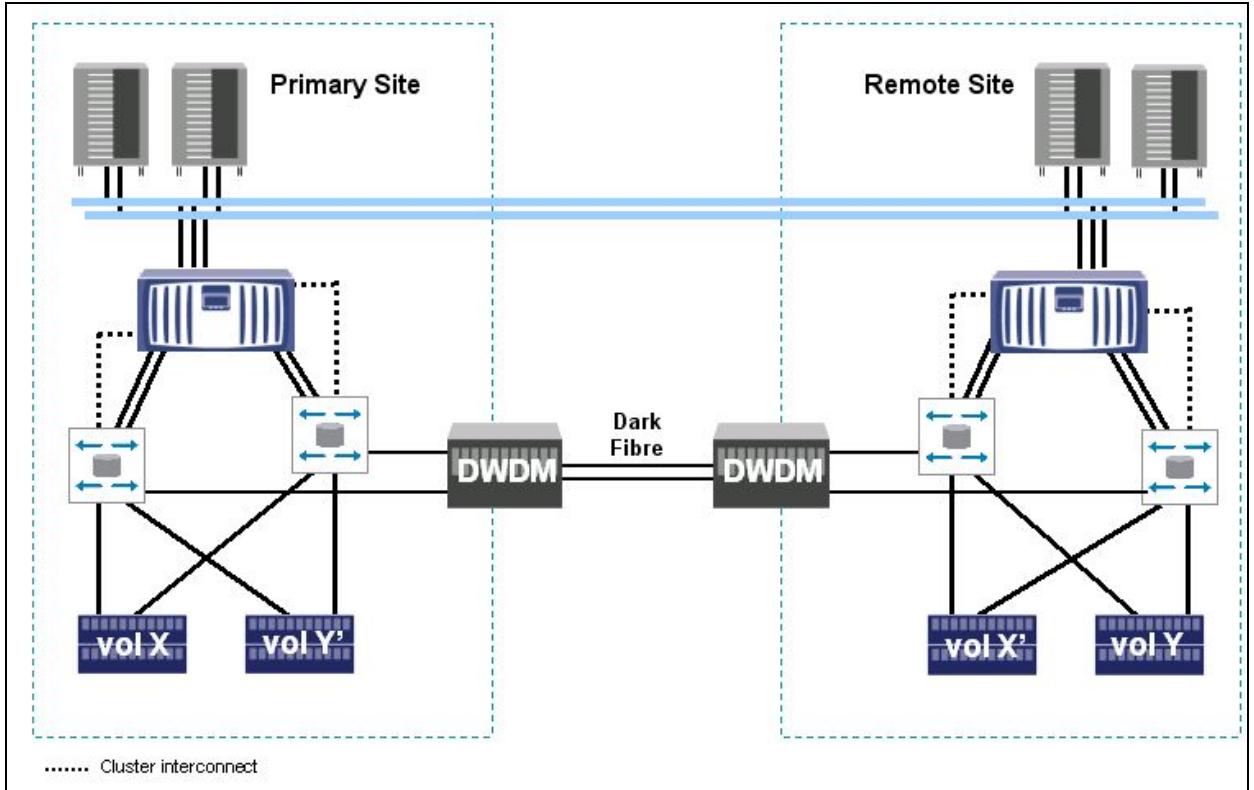


Figure 25) MetroCluster over FC and DWDM switch infrastructure.

This solution connects distant sites in metropolitan areas.

ReplicatorX

As already noted, ReplicatorX is enterprise-class replication software that works across platforms and over global distances. ReplicatorX software ensures enterprise-wide data consistency and supports nondisruptive replication of production data to cost-efficient secondary storage. This enables comprehensive disaster recovery and block data migration. Together with N series FlexClone technology, you can also accelerate testing and development.

With ReplicatorX, you can continuously copy and protect SAP data—over any distance and across heterogeneous infrastructures—from multiple data centers to a target site over an IP network. ReplicatorX systems use lightweight host-based or SAN-resident client-side software that does not affect your SAP production system. At the target site, ReplicatorX software maintains an up-to-the-second replica of your source data, which minimizes SAP data loss in the event of a disaster. ReplicatorX systems ensure data consistency among several SAP applications distributed across server and storage systems. Unlike conventional replication solutions, this unique capability eliminates complex and time-consuming data reconciliation and enables SAP landscape recovery in minutes. ReplicatorX simplifies SAP data management and provides centralized visibility and control of all replication sessions. Administrators can quickly define policies, delegate access, and manage environments at any level, from a single volume to an entire geography. With ReplicatorX software, you can manage cloning operations through a single interface, with comprehensive and user-configurable historical reporting.



Conclusions

- IBM N series has multiple disaster recovery solutions to support different business and financial requirements.
- SnapMirror provides an efficient and cost-effective disaster recovery solution.
- MetroCluster enables disaster recovery in a high-availability cluster configuration with no data loss.
- ReplicatorX enables consistent SAP landscape disaster recovery in heterogeneous environments.

Archiving and compliance

Business challenges

Archiving

The long-term accumulation of data in the SAP database can eventually 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 access data, risk, storage scalability, compatibility, and TCO. Current WORM (write once, read many) technologies like optical disk and 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 an Enterprise Storage Group (ESG) estimate, about 10% of regulated data is stored on disk today. To date, disks have not often been utilized, due to a number of factors, including 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.



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 for 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, HTTP(s) as a communication service, and WebDAV (web-based distributed authoring and versioning) 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 allows users to collaboratively edit and manage files on remote Web servers. The major features of this protocol are: 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 archiving data can then be transferred directly from the primary storage system to an external content or archive server. N series solutions for SAP archiving such as N series systems 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.

N series SATA storage is the preferred compliance and archive storage subsystem for SAP landscapes. N series with NearStore option leverages N series Data ONTAP 7.1 and later technology and takes full advantage of value-added software such as SnapLock. With over 99.995% field-measurable uptime, N series RAID-DP technology enables N series systems to tolerate double disk failures with no data loss.

SnapLock

SnapLock is the N series implementation of 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 N series technologies. This implementation also includes significant efforts in securing Data ONTAP and its administrative interfaces to the degree that SnapLock can be deployed for protecting data in regulatory environments. 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. These volumes can be locked either for a designated retention time or permanently. 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 plug in (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 can use the regular programmatic library interfaces that they would use for file operations on any other kind of storage system.

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 N series SnapLock.

SnapLock compliance

SnapoLock compliance enables organizations to satisfy strict records-retention regulations, such as SEC Rule 17a-4 (broker-dealers), HIPAA (health care), 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 N series system, can result in record deletion or alteration prior to the specified retention date.

SnapLock Enterprise

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 all N series primary and N series with NearStore option platforms except the IBM System Storage N series Gateway product line.

Conclusion

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. N series works well with existing document and content management packages.
- Data can be managed and backed up using the customer's current products and strategies.
- Solution can incorporate existing 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 N series with NearStore option , a more cost-effective alternative for archiving SAP data than adding database storage or processing power.



Trademarks and special notices

© International Business Machines 1994-2008. 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.

Data ONTAP, FlexClone, FlexShare, FlexVol, MetroCluster, NearStore, Network Appliance, the Network Appliance logo, RAID-DP, ReplicatorX, SnapDrive, SnapInfo, SnapLock, SnapManager, SnapMirror, SnapRestore, Snapshot, SnapVault SyncMirror 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.

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

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.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

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.