



Technical Report

SAP on Windows and Microsoft SQL Server with NetApp Storage

SAP Competency Center, NetApp
May 2011 | TR-3585

ABSTRACT

This document provides customers and partners with the best practices for deploying NetApp® storage in support of SAP® Business Suite solutions running in a Windows® environment using a Microsoft® SQL Server® database.

TABLE OF CONTENTS

1	INTRODUCTION	4
1.1	SCOPE	4
1.2	BUSINESS CHALLENGES	4
1.3	IT CHALLENGES	4
1.4	NETAPP SOLUTIONS FOR SAP	5
2	STORAGE PROVISIONING AND MANAGEMENT	5
2.1	CONSOLIDATION	5
2.2	STORAGE LAYOUT	8
2.3	SIZING	15
2.4	INSTALLATION	16
2.5	STORAGE MIGRATION	17
3	SYSTEM MANAGEMENT AND MAINTENANCE	18
3.1	SAP SYSTEM CLONING	18
3.2	SAP UPGRADE	22
4	BUSINESS CONTINUANCE	25
4.1	BACKUP AND RECOVERY	25
4.2	HIGH AVAILABILITY	30
4.3	DISASTER RECOVERY	31
5	ARCHIVING AND COMPLIANCE	34
6	REFERENCES	36
	REVISION HISTORY	37

LIST OF TABLES

Table 1)	FlexVol volume layout	10
Table 2)	Number of LUNs, based on database size	12
Table 3)	FlexVol layout for large SAP systems	14
Table 5)	SAP system copy comparison	21

LIST OF FIGURES

Figure 1)	FlexVol technology	6
Figure 2)	FlexShare	6
Figure 3)	Storage management with SnapDrive and Data ONTAP DSM	7
Figure 4)	Database management with SnapManager for SQL	8
Figure 5)	Aggregate layout	9
Figure 6)	FlexVol volume layout	10

Figure 7) Storage layout with NetApp MetroCluster.	11
Figure 8) Storage layout with NetApp MetroCluster; only production systems are mirrored.	11
Figure 9) Standard storage layout.	13
Figure 10) FlexVol layout for large SAP systems.	14
Figure 11) LUN layout for large SAP systems.	15
Figure 12) SAP system cloning overview.	20
Figure 13) SAP upgrade overview.	22
Figure 14) SAP upgrade—development system.	23
Figure 15) SAP upgrade—quality assurance system.	24
Figure 16) SAP upgrade—productive system.	25
Figure 17) SAP backup and recovery—NetApp solution overview.	27
Figure 18) Comparison of different backup approaches.	28
Figure 19) Time needed for restore and recovery.	29
Figure 20) SAP testing cycle.	30
Figure 22) NetApp clustered storage system solution.	31
Figure 23) Disaster Recovery with SnapMirror.	32
Figure 24) MetroCluster over direct Fibre Channel switch connection.	33
Figure 25) MetroCluster over Fibre Channel and DWDM switch infrastructure.	33

1 INTRODUCTION

1.1 SCOPE

This document provides customers and partners with best practices for deploying NetApp storage systems in support of SAP Business Suite solutions running in a 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 NetApp and SAP products and was developed with input gathered from NetApp, SAP, Microsoft, and our customers.

1.2 BUSINESS CHALLENGES

Corporations deploying SAP solutions today are under great pressure to reduce TCO, accelerate 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.

1.3 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, like ERP, CRM, and 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 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.

1.4 NETAPP SOLUTIONS FOR SAP

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

2 STORAGE PROVISIONING AND MANAGEMENT

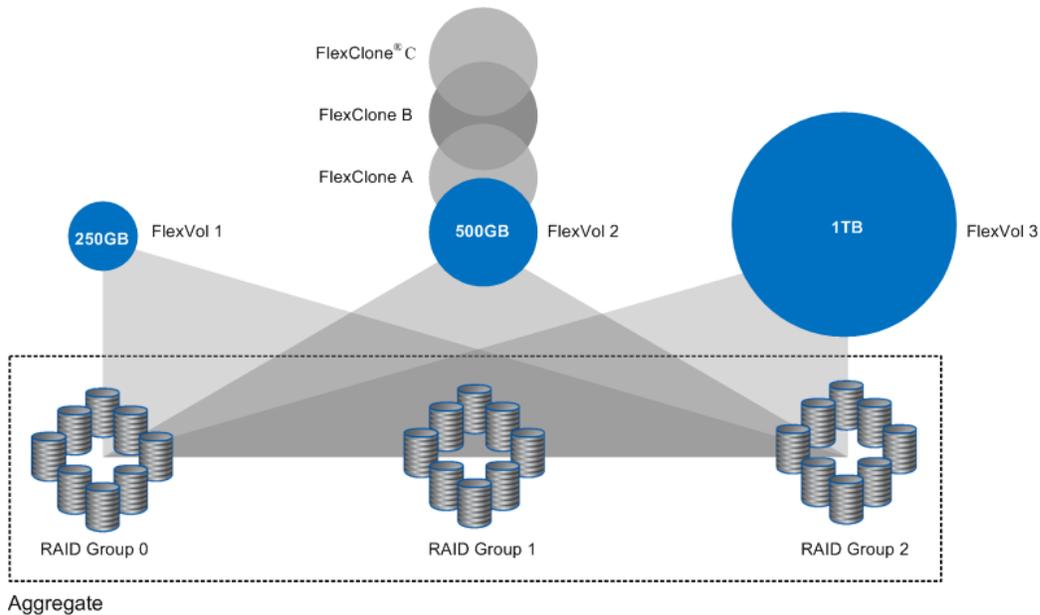
2.1 CONSOLIDATION

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 total cost of ownership (TCO).

NetApp offers highly available, scalable, and cost-effective storage consolidation solutions that incorporate the NetApp 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. NetApp enterprise-class storage solutions are proven interoperable across all platforms. NetApp fabric-attached storage (FAS) systems integrate easily into a complex enterprise and simultaneously support NAS, Fibre Channel SAN, and IP SAN (iSCSI) protocols.

NetApp [FlexVol](#)[®] 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. NetApp 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.

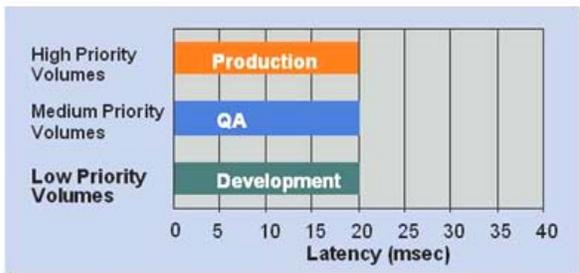


NetApp [FlexClone](#) enables true cloning—instant replication of data sets without requiring additional storage space at the time of creation. Each cloned volume is a transparent, virtual copy that can be used to test application patches, to run performance and data integrity tests, or to provide user-training 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.

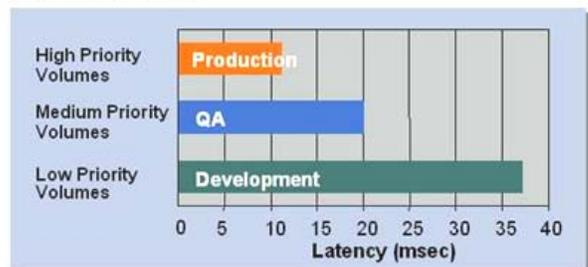
NetApp [FlexShare](#)® gives administrators the ability to leverage existing infrastructure and increase processing utilization without sacrificing the performance required to meet critical business needs. It prioritizes processing resources for key services when the system is under heavy load. With 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.

Without FlexShare



With FlexShare



FlexShare is a Data ONTAP® software feature that provides workload prioritization for a storage system.

With NetApp [SnapDrive](#)® for Windows, 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 NTFS file system and provides a layer of abstraction between application data and physical storage associated with that data. SnapDrive provides storage

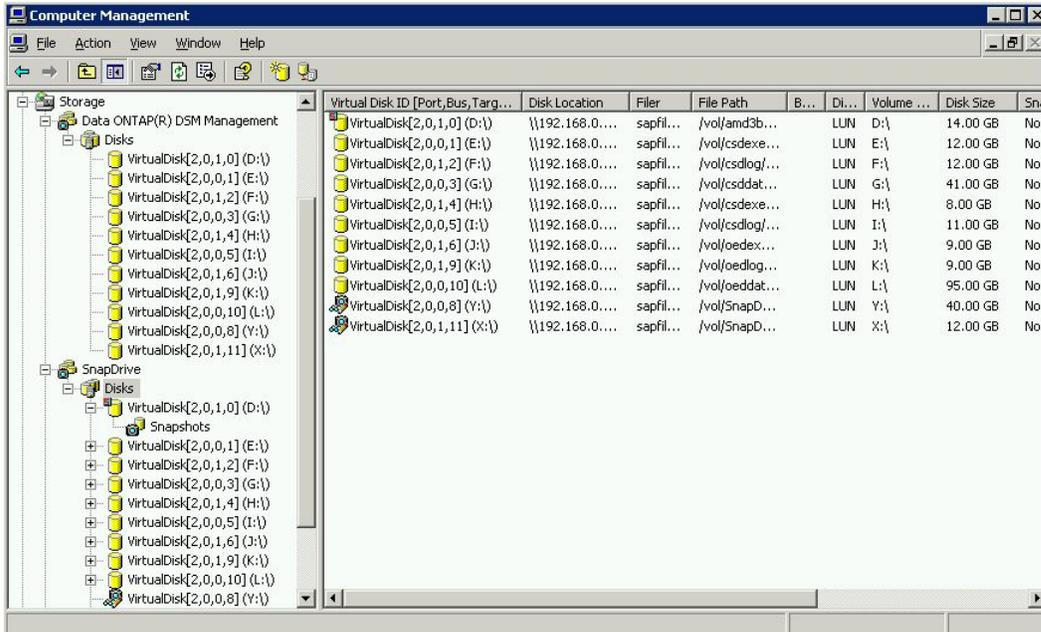
virtualization of volumes via the iSCSI or Fibre Channel (FCP) access protocol. NetApp storage systems and SnapDrive software represent a complete data management solution for Windows applications. SnapDrive includes Windows device drivers and software that is used to manage application Snapshot® backups. Snapshot backups are nondisruptive to applications and occur very quickly. Restoring data from a Snapshot copy is nearly instantaneous.

The list below highlights some of the important benefits NetApp storage systems offer to applications:

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

[SnapDrive](#) is independent of the underlying storage access media and protocol. The iSCSI protocol provides storage access when the storage controller and host server are connected using Gigabit Ethernet. The FCP protocol facilitates storage access through a Fibre Channel host bus adapter (HBA) and storage area network (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 NetApp storage systems. Thus a virtual disk can be created and accessed using the iSCSI or FCP access protocols. Virtual disks are referred to as logical unit numbers (LUNs) when accessed over the iSCSI and FCP protocols. It is possible to change the protocol from iSCSI to FCP and vice versa without migrating the data.

Figure 3) Storage management with SnapDrive and Data ONTAP DSM.



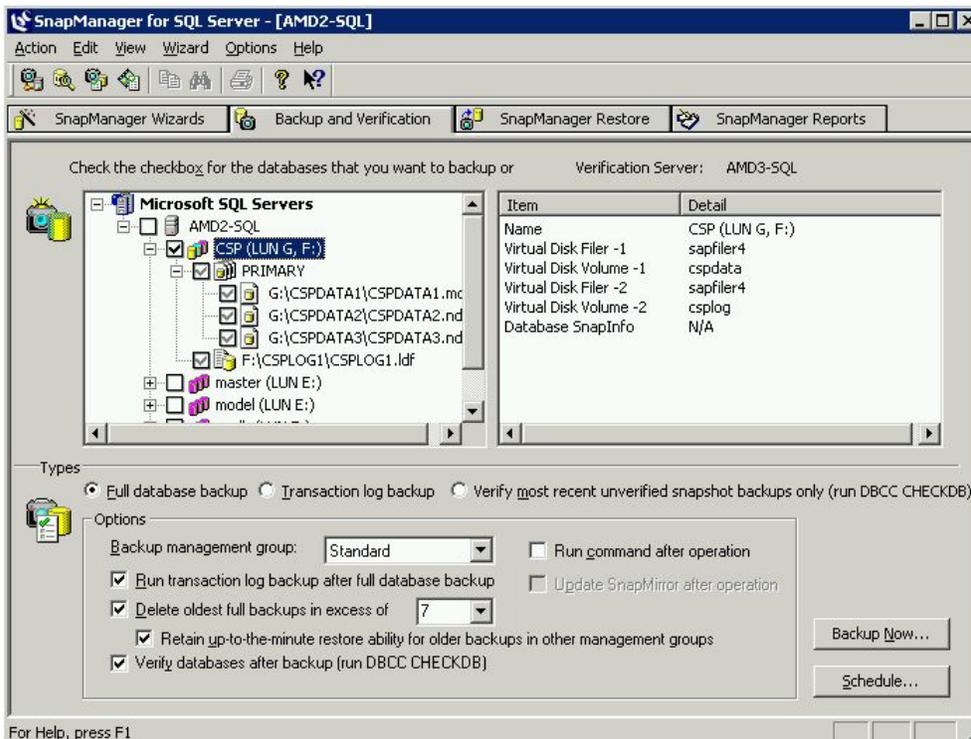
[SnapManager® for SQL Server](#) paves the way for database and storage administrators to simplify data management utilizing the powerful capabilities of NetApp storage systems.

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 Fibre Channel 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

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

Figure 4) Database management with SnapManager for SQL.



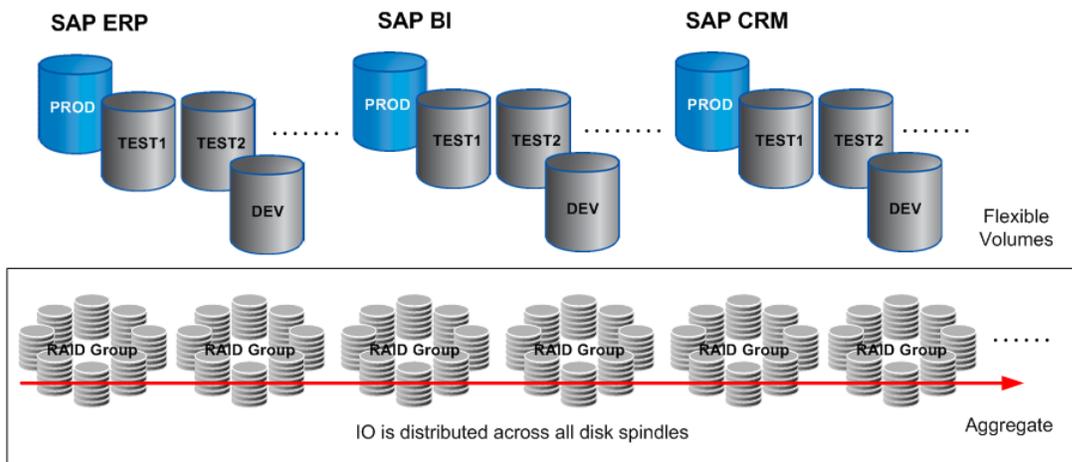
2.2 STORAGE LAYOUT

AGGREGATE LAYOUT

NetApp recommends using 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 RAID-DP®, which offers a high level of data protection. The reliability of RAID-DP is far greater than that 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 will benefit from the disk spindles of the test and development systems, which are needed anyway because of capacity requirements.

With shared resources, contention for available resources among systems is 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 adapted during online operation. For more information on FlexShare, see the [FlexShare Design and Implementation Guide](#).

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
- One volume for SnapManager for SQL Server SnapInfo

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®, 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.

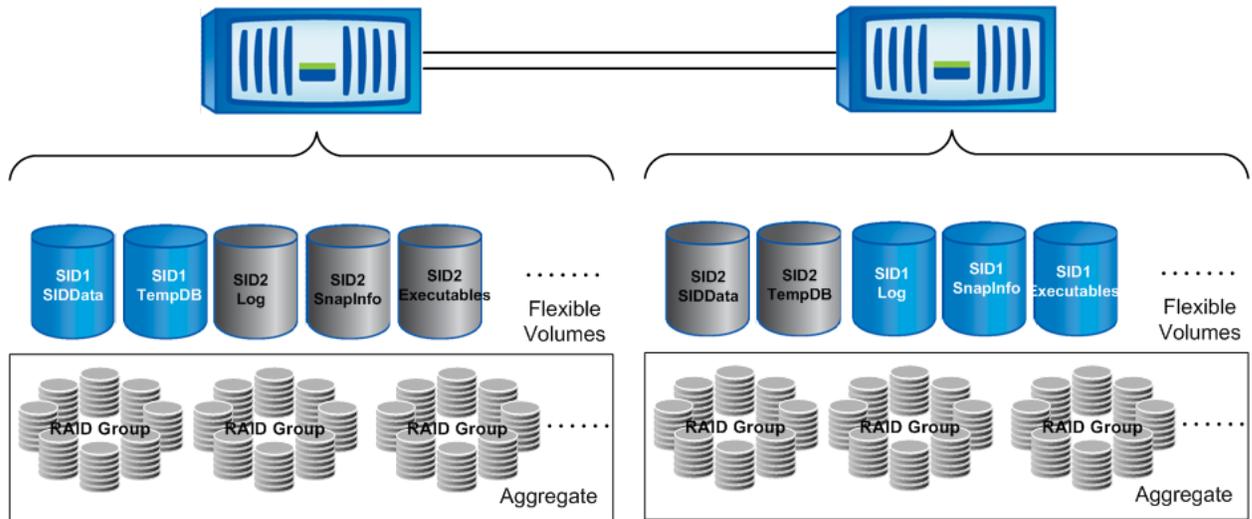


Table 1 shows the distribution of the file systems/LUNs of a single SAP instance to the FlexVol volumes.

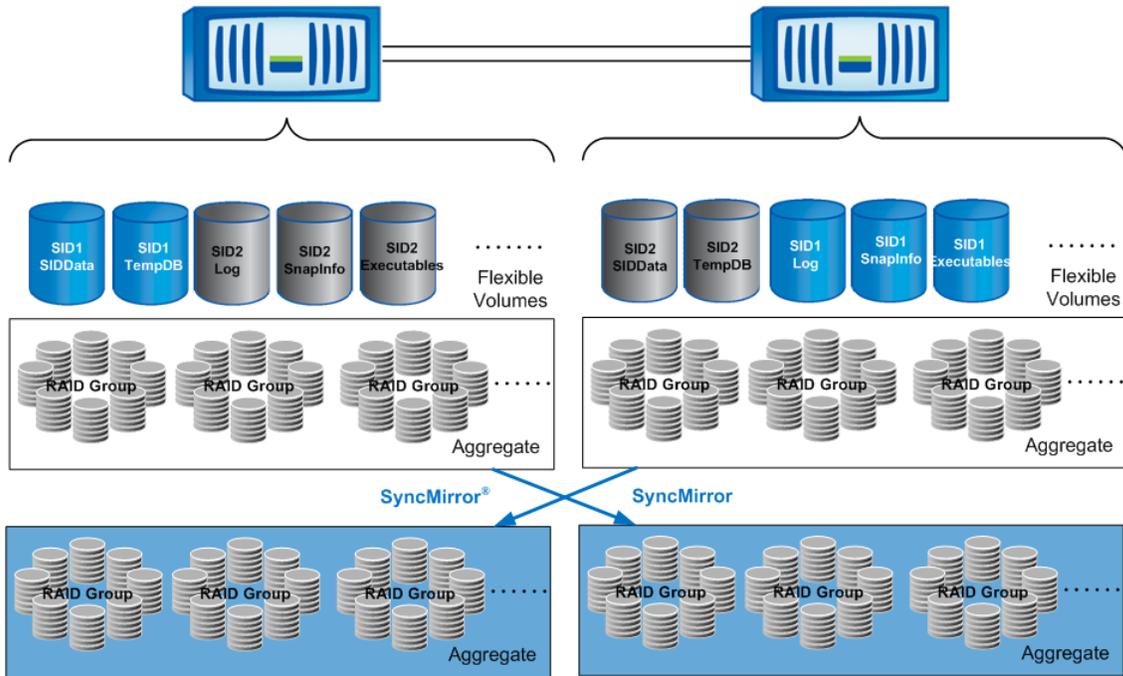
Table 1) FlexVol volume layout.

Storage Controller 1		Storage Controller 2		
Aggregate 0		Aggregate 1		
SIDData	TempDB	Executables	Log	SnapInfo
SIDData1	TempDB	SQL Executables	Logfile (s)	SnapInfo Dir
SIDData2	Quorum MSCS	System DBs		
SIDData3		SAP Executables		

LAYOUT WITH METROCLUSTER

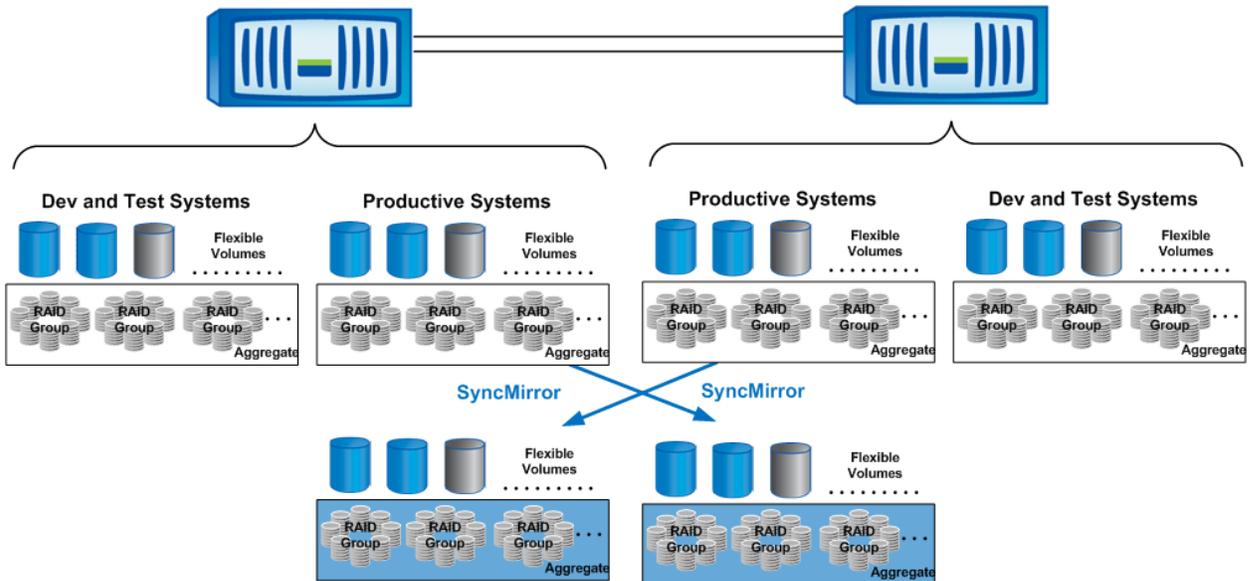
MetroCluster and synchronous mirroring work on the aggregate level. If all SAP systems are required to be mirrored synchronously, the layouts for a MetroCluster and a normal cluster are the same.

Figure 7) Storage layout with NetApp 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 8) Storage layout with NetApp MetroCluster; only production systems are mirrored.



For more information on MetroCluster, see section [4.3 Disaster Recovery](#).

LUN LAYOUT

The size of the database and the kind of SAP System determines 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. Table 2 gives some guidelines for a reasonable number of LUNs, based on the size of the database.

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

	Database Size	# of LUNs Data Files	Size of LUNs Data Files	# of LUNs Logfile(s)
Small System	< 200GB	1	100GB – 200GB	1
Medium System	200GB – 1TB	2-3	100GB – 350GB	1
Large System Or BI /BW Systems	> 1TB	> 3	200GB – 450GB	1

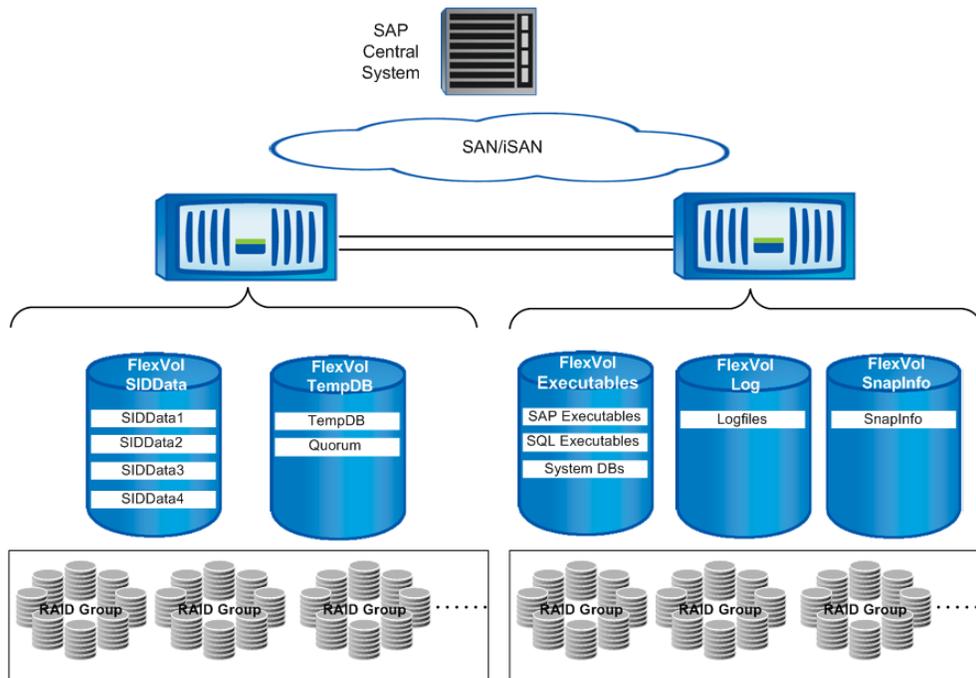
SQL SERVER DATA FILE LAYOUT

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, NetApp recommends storing each data file in its own LUN.

RECOMMENDED STANDARD LAYOUT

Figure 9 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 file(s) in the Log FlexVol volume. Two 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 and one LUN is configured in the SnapInfo FlexVol for the SMSSQL SnapInfo directory. The Quorum disk of a Microsoft Cluster should be stored in the SnapInfo FlexVol.

Figure 9) Standard storage layout.



FLEXVOL VOLUME LAYOUT FOR LARGE SAP SYSTEMS WITH HIGH I/O REQUIREMENTS

Systems with high I/O requirements should be distributed across 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 to make sure 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 file(s)
- One volume for the executables including SQL Server's system databases
- One volume for SQL Server's TempDB
- One volume for SnapManager for SQL Server SnapInfo directory and the Quorum Disk in case of MSCS

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

Figure 10) FlexVol layout for large SAP systems.

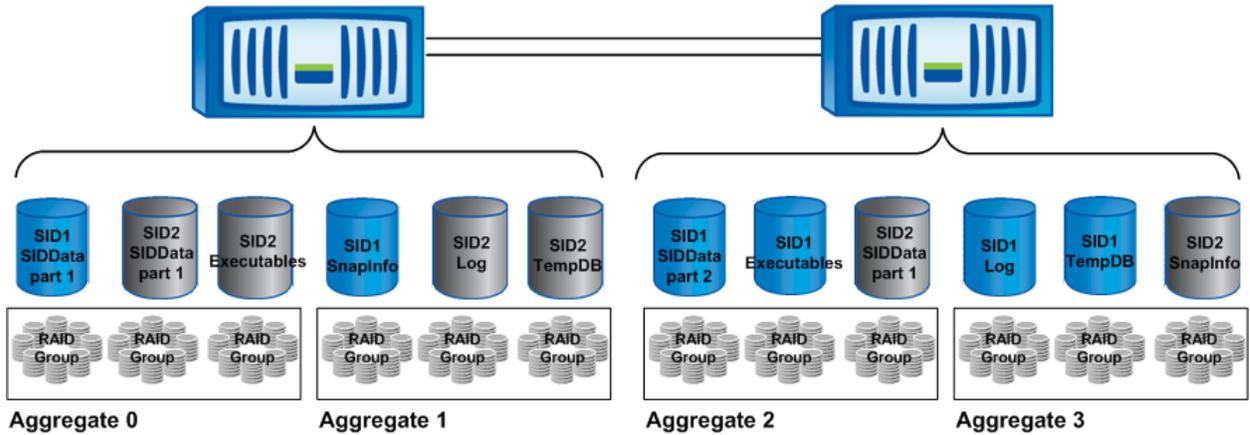


Table 3 shows the distribution of the file systems/LUNs of a single SAP instance to the FlexVol volumes. One LUN for each database file (x.mdf; x.ndf) is used, in this example 8.

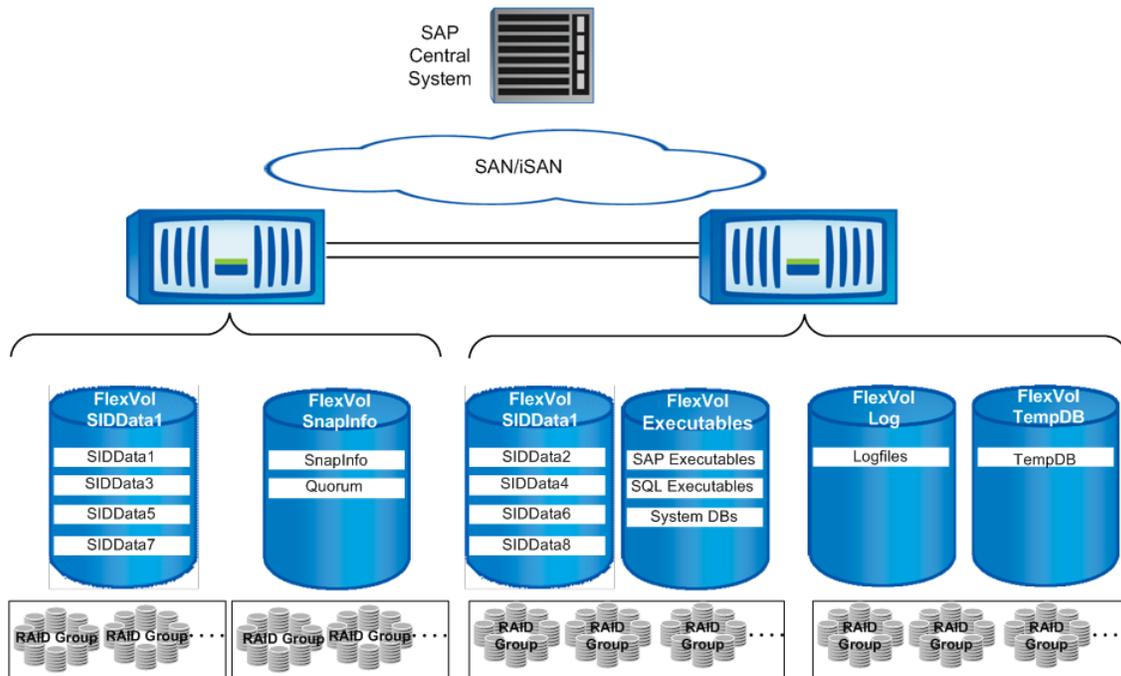
Table 3) FlexVol layout for large SAP systems.

Storage Controller 1		Storage Controller 2			
Aggregate 0	Aggregate 1	Aggregate 2		Aggregate 3	
SIDData	SnapInfo	SIDData	Executables	Log	TempDB
SIDData1	SnapInfo Dir	SIDData2	SAP Exe	Logfile (s)	TempDB
SIDData3	Quorum MSCS	SIDData4	System DBs		
SIDData5		SIDData6	SQL Exe		
SIDData7		SIDData8			

LUN LAYOUT FOR LARGE SAP SYSTEMS WITH HIGH I/O REQUIREMENTS

Figure 11 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, NetApp recommends mounting 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 SIDDATAx directories, in this example eight. Then create and mount the other seven LUNs to this newly created directories. For detailed information, refer to the SnapDrive for Windows Installation and Administration Guide. When the FCP protocol is used, NetApp recommends distributing the load of the data file LUNs to the available FCP HBAs or ports. With Data ONTAP DSM 4.0 or later, 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.

2.3 SIZING

This section gives an overview of the storage sizing for a SAP environment using NetApp 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.

NetApp can provide storage sizing to SAP customers, based on a sizing questionnaire that the customer fills out.

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. 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 1 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.

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 supporting both the I/O and capacity requirements.

2.4 INSTALLATION

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

GENERAL REQUIREMENTS

NetApp strongly recommends the use of SnapDrive for Windows, a NetApp host-based software product that simplifies storage management and provisioning in SAP Windows storage environments. It is required for SnapManager for SQL, which provides host-consistent database snapshots, SnapRestore and database cloning functionality. For information on SnapDrive system requirements, see <http://www.netapp.com/products/enterprise-software/manageability-software/server-suite/snapdrive-windows.html>

For additional recommendations on system configuration and tuning, see:

- [Best Practices Guide: Microsoft SQL Server](#)
- [SAN Booting with Windows Server Operating Systems](#)
- [iSCSI multipathing possibilities on Windows with Data Ontap](#)

NETAPP 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 or stored in qtrees to be able to use all desired software functionality like Qtree based SnapVault.

2.5 STORAGE MIGRATION

This section discusses different storage migration approaches. If a migration includes a change of the operating system or database system, the migration cannot 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 about 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 that needs to be migrated. In general, there are three approaches to storage migration of the SAP data:

- Migration on the operating system level
- Migration on the database level
- Migration on the storage system level
- NetApp SnapMirror for Open Systems

MIGRATION ON THE OPERATING SYSTEM LEVEL

In addition to the existing storage system, the NetApp storage system is connected to the database server. The NetApp 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 NetApp 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 FCP-to-FCP or an FCP-to-iSCSI 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. SnapManager for SQL Server offers a Wizard to perform such a migration.

MIGRATION ON THE DATABASE LEVEL

A SQL Server database backup is restored to the NetApp storage system. To minimize the impact on the productive SAP system, the restore can be done using a separate server connected to the NetApp 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 has to be performed. The log backups that have not yet been copied are now copied from the old storage system to the NetApp storage system and a forward recovery of the database is carried out. Then, detach the databases using the SQL Server Management Studio, disconnect the old storage from the database server and connect the NetApp storage system and the LUNs stored there.

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

A migration on the database level can be done for an FCP-to-FCP or an FCP-to-iSCSI 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-NetApp SAN to a NetApp SAN storage system, NetApp has partnered with the best-in-class SAN data migration companies to provide a fast, flexible SAN data migration service to suit individual customer requirements. Essentially, the service uses a fabric-attached appliance and host software, which hosts data-copying software. The appliance is attached to the existing SAN as well as to the NetApp 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 (NetApp) 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 NetApp storage controller. Only after all of the servers can access the new storage can the migration appliance be pulled out. For more details, see SAN Data Migration Services at <http://www.netapp.com/support/consultingedge.html>.

A migration on the storage-system level can be done only for an FCP-to-FCP migration. This approach minimizes the downtime during the migration, but needs the appliance described earlier for the migration process.

3 SYSTEM MANAGEMENT AND MAINTENANCE

3.1 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 very quickly is 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 system 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 address these demands.

SAP copies also consume a significant amount of storage, which needs to be provided. Because these copies are typically clones of the productive system, the amount of storage needed can be huge.

NETAPP SOLUTION

The NetApp 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 productive system. In addition, NetApp cloning functionality offered by SnapManager for SQL Server allows very efficient storage management by storing only data changes between the source and the clone.

NETAPP SOLUTION FOR SAP SYSTEM COPIES

SAP system copies are accomplished using SnapManager for SQL Server for cloning the SAP database, utilizing the NetApp FlexClone feature. A FlexClone copy is a writeable point-in-time image of a NetApp FlexVol volume 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. FlexClone copies store only changed blocks between the source FlexVol volume and the FlexClone, and therefore significantly decrease the amount of disk space needed for SAP system copies. SAPINST (WebAS 7.0 or higher) or the SAP STM tool (WebAS 6.x) is used to perform the SAP specific task for a homogenous system copy

Figure 12 shows the basic concept of the system copy solution. Creating a SAP system copy consists of several steps on the source system and several steps on the destination system.

On the source system, a database-consistent Snapshot copy of the 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.

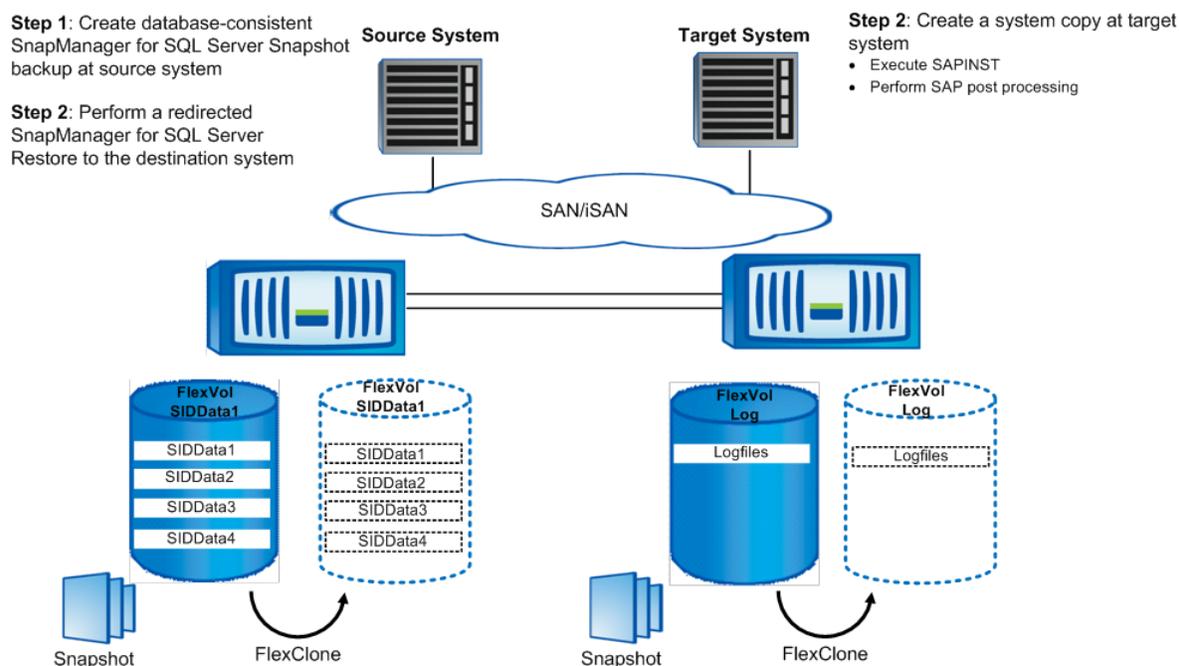


Table 5 compares the traditional approach to the NetApp approach to perform a SAP system copy.

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 of the database takes several hours. The subsequent steps are typically carried out manually, consuming IT staff time.

With the NetApp approach, the backup is taken using 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 is typically carried out manually.

With the NetApp approach SnapManger for SQL Server creates a FlexClone 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 database to the new SID is automatically done by SnapManager for SQL Server.

In both scenarios SAPINST has to be executed to perform the necessary SAP specific tasks. Additionally post-processing tasks have to be done.

Table 4) SAP system copy comparison.

Traditional Approach	NetApp 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 NetApp 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 existing yet) Restore backup from source system Change of database name SAP-specific post processing tasks 	<ul style="list-style-type: none"> Install MS SQL Server (if not existing yet) Create FlexClone image based on Snapshot backup; mount the LUNs at the target system Change of 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 Manual process

Conclusions

The NetApp 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 as there is no impact on the online operation of the source system.
- Snapshot and FlexClone functionality offered by SnapManager for SQL Server reduce 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 reduce the necessary disk space for a SAP system copy by storing only data changes between the source and the target system.

3.2 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.

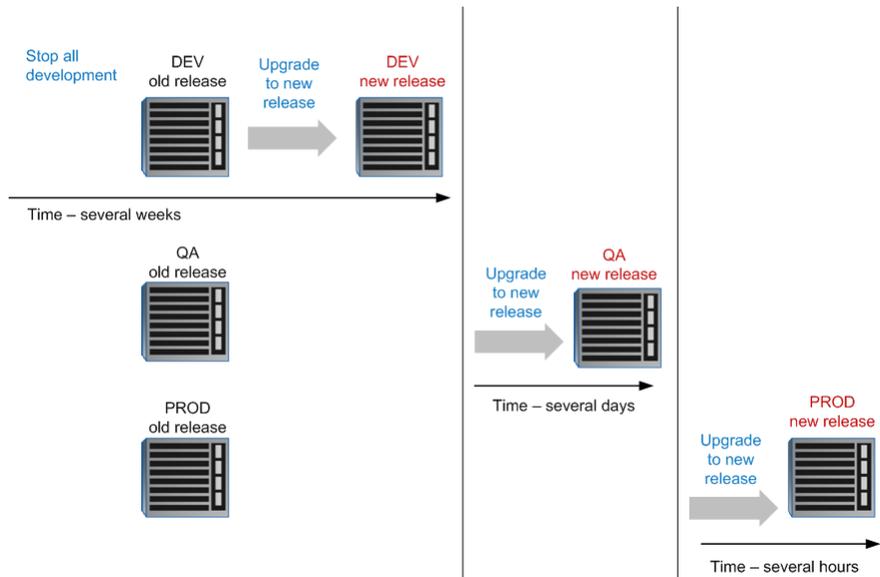
Upgrading to a new SAP release is a challenging project that consumes large amounts of employee time, including IT staff time.

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 a 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 a SAP system copy normally takes several days and negatively affects the productive environment.

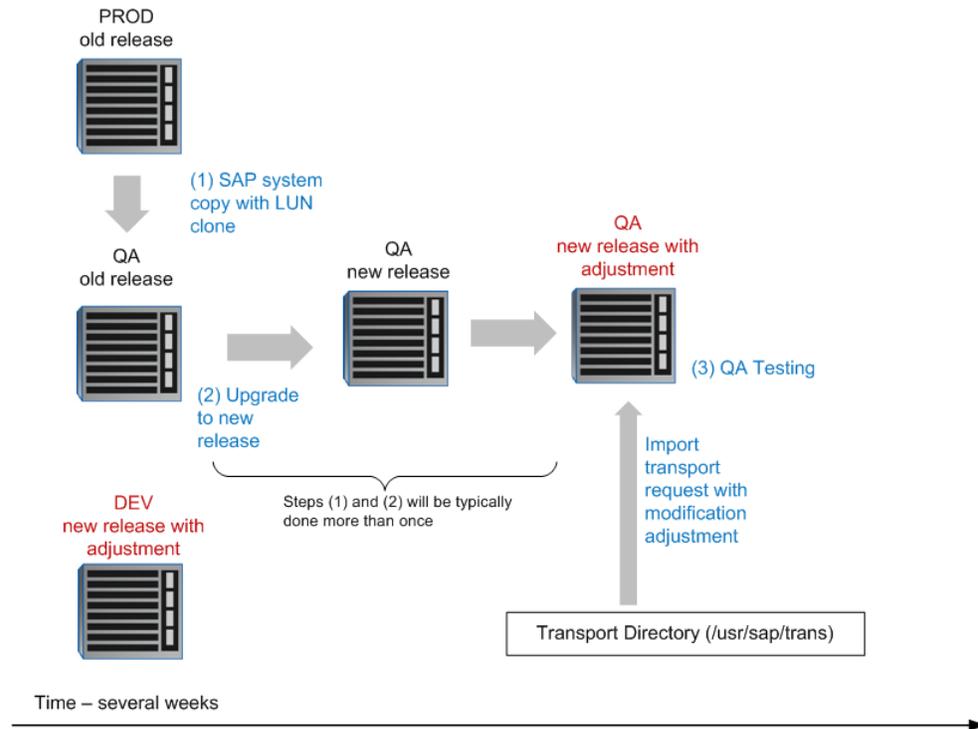
Figure 13) SAP upgrade overview.



NETAPP SOLUTION

The NetApp 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 NetApp 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 NetApp solutions help to minimize the risk, reducing the downtime and reducing overload of IT staff resources during a SAP upgrade project.

Figure 15) SAP upgrade—quality assurance system.



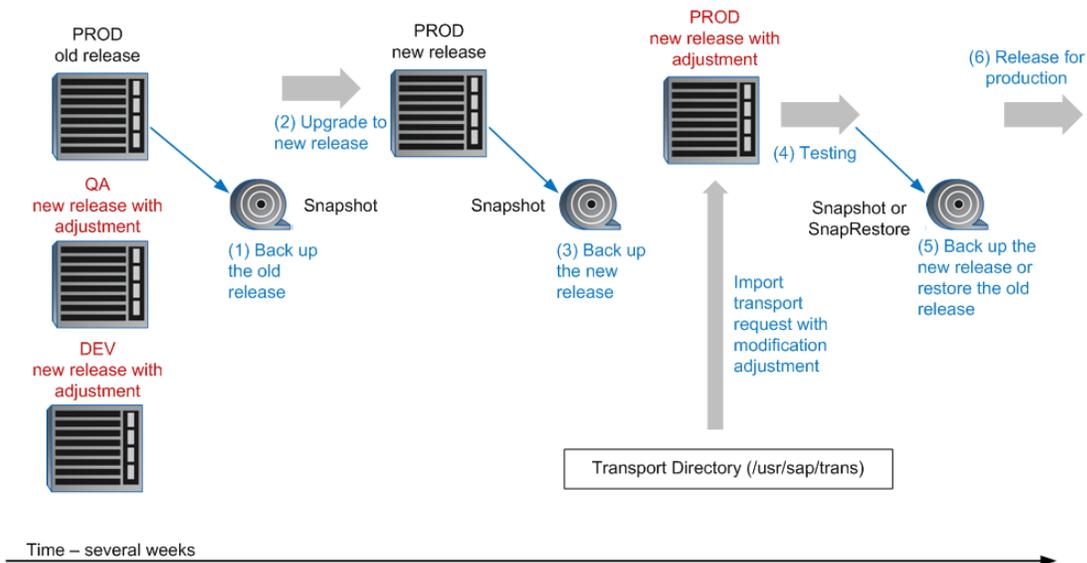
Upgrading the Productive System

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.

SnapManger for SQL offers Snapshot copies as a backup method and SnapRestore for restoring the system to its former release status to provide 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 SnapManager for SQL Server.

Figure 16) SAP upgrade—productive system.



Conclusions

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

- The NetApp 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 NetApp 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.
- NetApp SnapManager for SQL Server significantly reduces the total upgrade time of the productive system, providing a higher level of flexibility for scheduling the SAP upgrade.

4 BUSINESS CONTINUANCE

4.1 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 a 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 a 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.

NETAPP SOLUTION

NetApp provides unique storage solutions to address the challenges just described.

NetApp [SnapManager for SQL Server](#) can create SQL Server database backups in seconds using Snapshot technology. 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 NetApp 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 NetApp customers typically schedule several Snapshot online backups during the day—for instance, every 4 hours.

SnapManager for SQL Server also provides key advantages for the restore and recovery operation. The NetApp [SnapRestore](#)[®] functionality integrated into SnapManager for SQL Server 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 backups have 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 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, NetApp recommends using 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 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 a NetApp SnapManager for SQL Server (SMSQL) always consists of two parts:

1. Backup and restore/recovery using Snapshot and SnapRestore through SMSQL
2. 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 a NetApp [NearStore](#)® system or other storage system and [SnapVault](#)® software:

- The primary storage communicates directly with the secondary storage (NearStore or other storage system) and sends the backup data to the destination. The NetApp 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. SnapManager for SQL Server automatically transfers the changed blocks during a backup.

Backup to tape using third-party backup software:

- NDMP 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—NetApp solution overview.

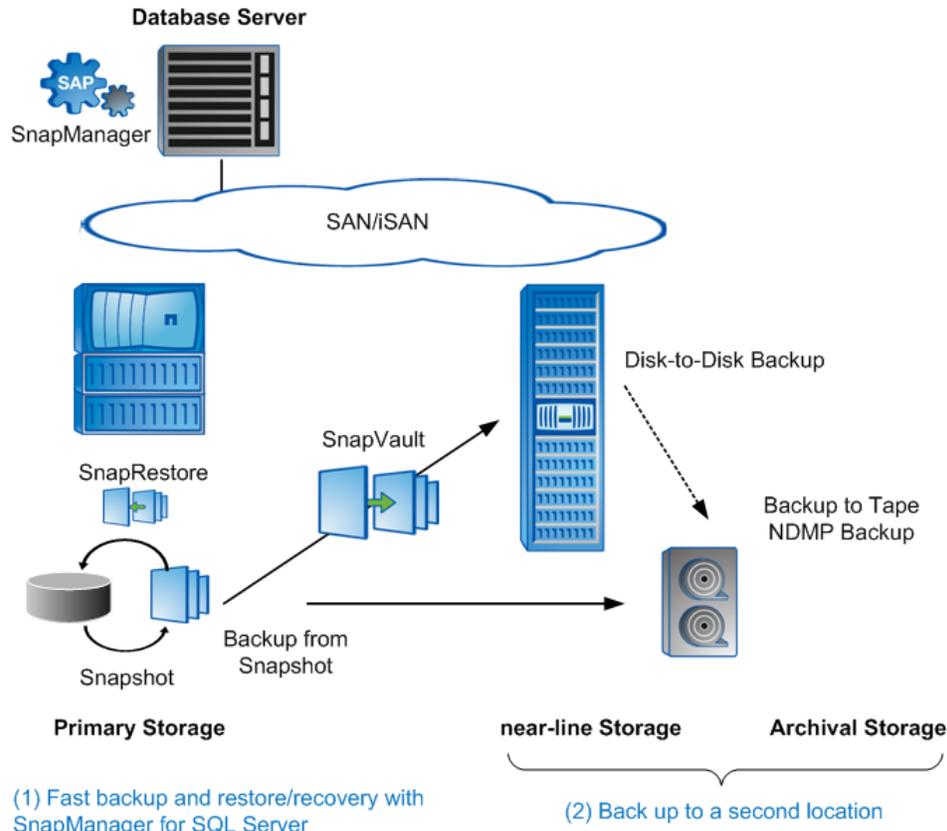
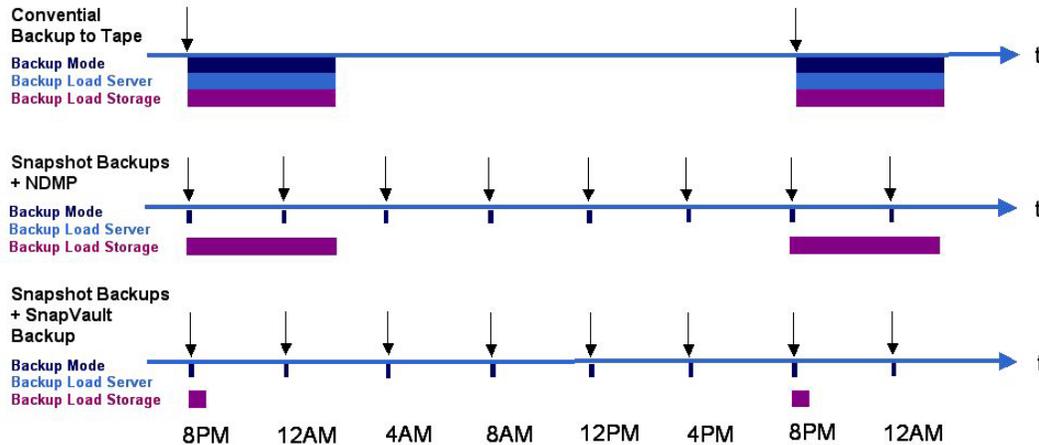


Figure 18 compares the different backup approaches with regard to the performance impact of a backup and the time that the database must be in backup mode.

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, the database is in backup mode 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 is in backup mode 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 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 in backup mode while running the backup to tape.

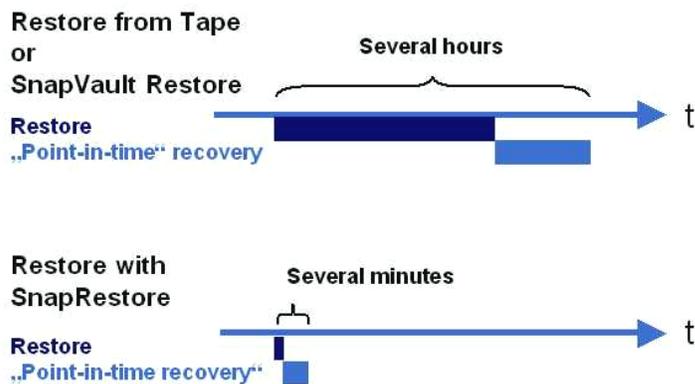
SnapManager for SQL Server: Snapshot Backups Together with Disk-to-Disk Backup and SnapVault

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 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.

Figure 19 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 amount of log backups need to 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 taken with a higher frequency, such as every 4 hours, so the forward recovery is much faster, because fewer amount of log backups need to be applied.

If Snapshot backups are used in combination with tape or SnapVault backups, most restore cases 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.

Conclusions

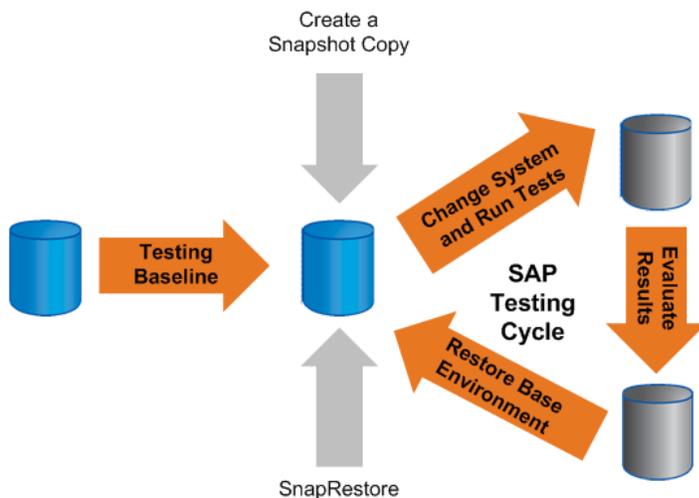
The combination of Snapshot and SnapRestore with a disk-to-disk backup concept based on SnapVault offers significant improvement over conventional tape backups:

- 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 consistency check (DBCC). SnapManager offers the possibility to run a database consistency check on a separate server automatically or manually after a SnapManager for SQL 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 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 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, 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 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. SnapManager for SQL Server offers Snapshot as a backup method, and SnapRestore for restoring the system to its former release status. This allows a higher level of flexibility with regard to scheduling. By creating several Snapshot copies at certain stages during the upgrade, it is possible to restart the upgrade without having to revert to the former release status.

4.2 HIGH AVAILABILITY

BUSINESS CHALLENGES

Productive SAP systems are business-critical applications that require 24 x7 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.

NETAPP SOLUTION

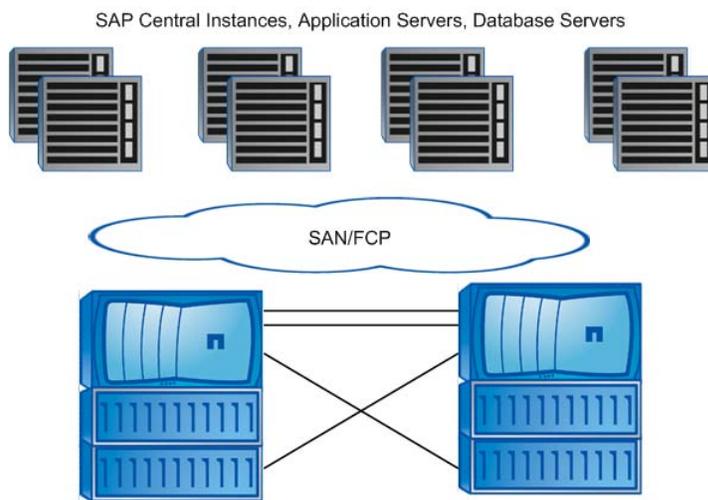
NetApp clustered failover delivers a robust and highly available data service for business-critical environments. Installed on a pair of NetApp storage controllers, NetApp clustered failover provides data availability by transferring the data service of an unavailable storage controller to the other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

NETAPP SOLUTION FOR SAP HIGH AVAILABILITY

Figure 21 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 NVRAM on both storage controllers. The disk shelves of the cluster partner are connected to the second storage controller via a second Fibre Channel 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. 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) NetApp clustered storage system solution.



Conclusions

The NetApp clustered failover 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 (MSCS) is used for the application.

4.3 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 provides an absolute minimum of downtime and rapid recovery from a disaster, with minimal or no loss of data. NetApp 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 NetApp products.

NETAPP SOLUTIONS

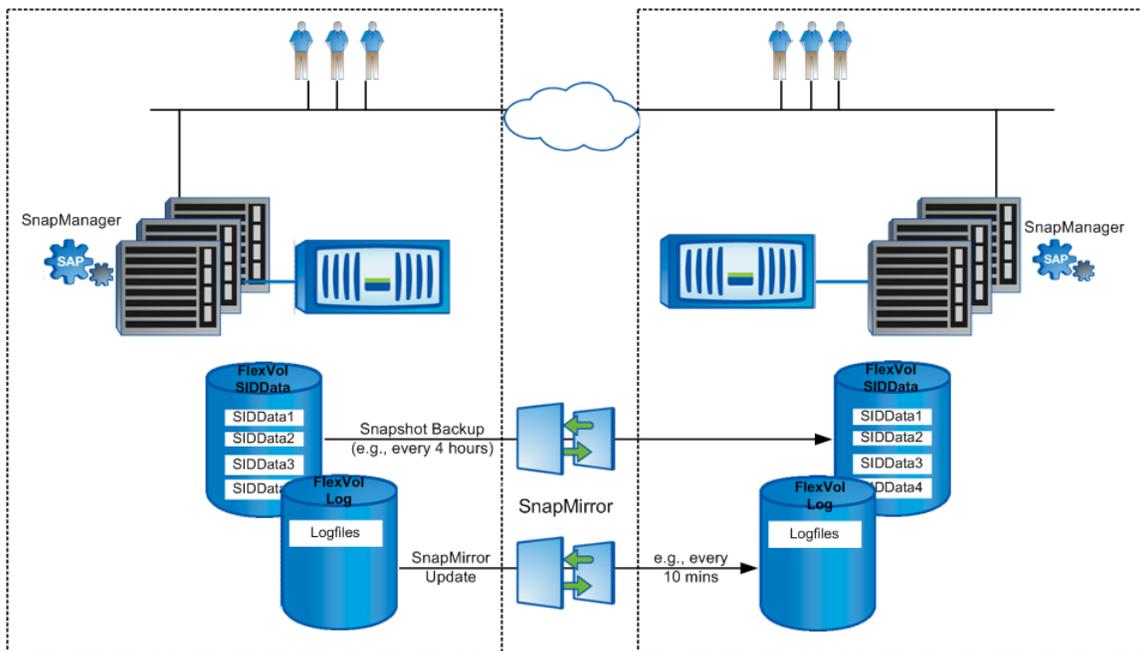
SnapMirror

NetApp [SnapMirror](#) delivers the disaster recovery 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 is integrated within SnapManager for SQL Server.

SnapMirror technology mirrors data to one or more storage controllers. It updates the mirrored data to keep it current and 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 NetApp 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. NetApp recommends a frequent SnapMirror update of the archive logs, such as every 10 minutes, for minimum of data loss. This can be achieved by schedule SnapManager log backups every 10 minutes.

Figure 22) Disaster Recovery with SnapMirror.



MetroCluster

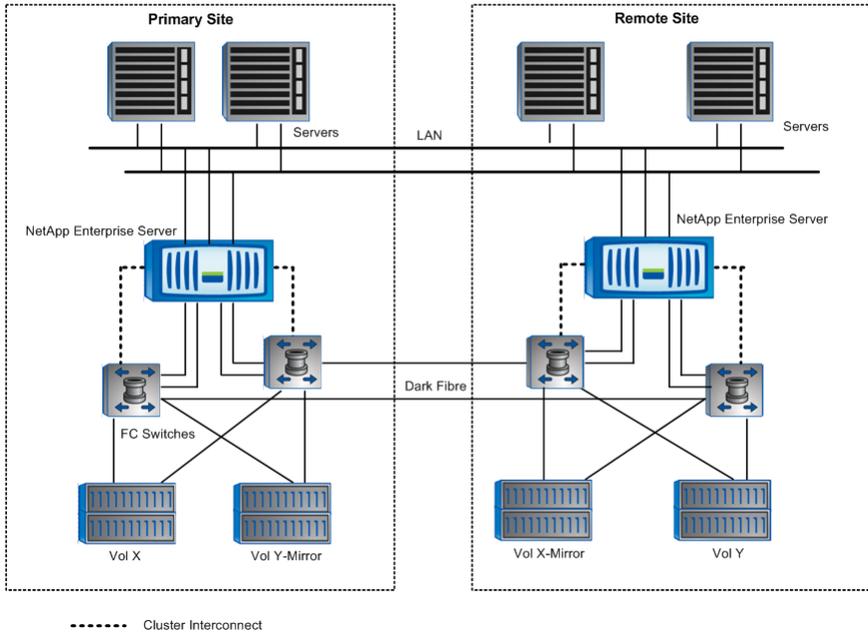
NetApp [MetroCluster](#) is an integrated high-availability and business continuance solution that provides disaster recovery with no data loss. MetroCluster extends failover capability from within a data center to a site located many miles away. It also replicates data from the primary site to the remote site to make sure that data there is completely current. The combination of failover and data replication makes sure that you can recover from a disaster—with no loss of data—in minutes rather than hours or days.

MetroCluster is much like NetApp Clustered Failover, but with the added benefit of disaster recovery. Clustered Failover creates a cluster of NetApp 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 SyncMirror to make sure that both storage controllers in the cluster have copies of the other storage controller's data.

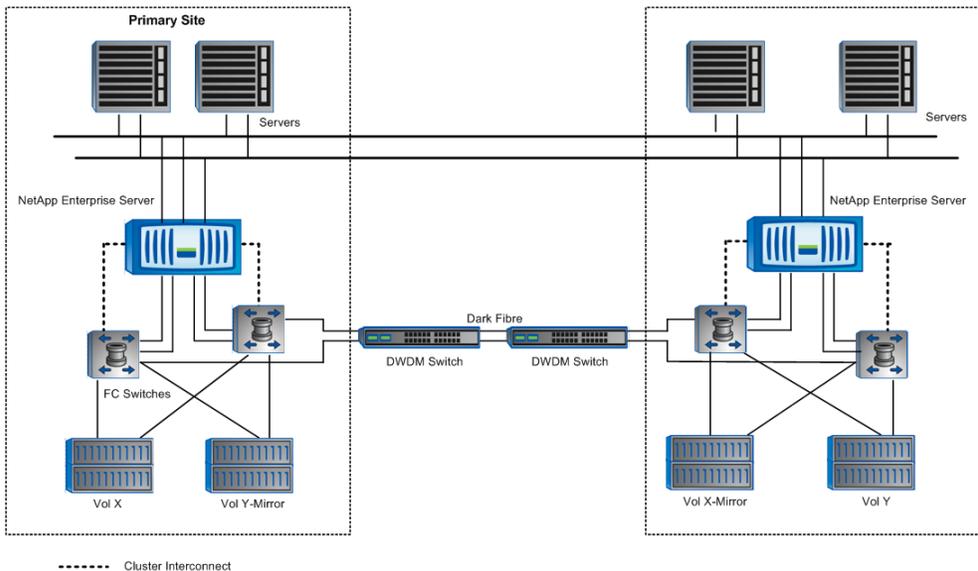
If the distance between the nodes is 500 meters or less, a stretch MetroCluster configuration can be used with OM3 cabling. If the distance is greater than 500m, a Fabric MetroCluster must be chosen, making Fibre Channel switches necessary.

Figure 23) MetroCluster over direct Fibre Channel switch connection.



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

Figure 24) MetroCluster over Fibre Channel and DWDM switch infrastructure.



This solution connects distant sites in metropolitan areas.

Conclusions

- NetApp 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.

5 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 total cost of ownership (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.

NETAPP 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, HTTP(s) as a communication service, and WebDAV as a general protocol for connecting storage systems.

The Archive Development Kit 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. NetApp 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. NetApp solutions for SAP archiving such as NearStore and [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.

NetApp FAS systems with SATA disks are the preferred compliance and archive storage subsystem for SAP landscapes. The product family leverages NetApp Data ONTAP 7G technology and takes full advantage of value-added software from NetApp such as SnapLock. With over 99.995% field-measurable uptime, NetApp RAID-DP technology enables FAS systems to tolerate single disk failures with no data loss. [SnapLock](#) is the NetApp implementation of high-performance disk-based magnetic WORM storage. SnapLock provides secure, storage-enforced data retention functionality via open file protocols such as CIFS and NFS while leveraging existing NetApp 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 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 they would use for file operations on any other kind of storage system.

SnapLock is available in two versions. One or the other of these versions can be implemented 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 NetApp 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 a NetApp 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 all NetApp FAS and NearStore platforms except the F7xx, F810, F820, F840, and V-Series product line.

Conclusions

NetApp 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. NetApp 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 NetApp 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 NearStore—a more cost-effective alternative for archiving SAP data than adding database storage or processing power.

6 REFERENCES

SQL Server Backup and Recovery

TR-3323—SnapManager for Microsoft SQL Server 2000 Best Practices

<http://www.netapp.com/library/tr/3323.pdf>

Data ONTAP 7G: FlexVol and FlexClone

TR-3348—Block Management with Data ONTAP 7G: FlexVol, FlexClone, and Space Guarantees

<http://media.netapp.com/documents/tr-3348.pdf>

[A thorough introduction to FlexClone volumes](#)

TR-3347—FlexShare Design and Implementation Guide

<http://media.netapp.com/documents/tr-3347.pdf>

Databases and Operating System Tuning

TR-3373—Data ONTAP 7G: The ideal platform for Database Applications

<http://media.netapp.com/documents/tr-3373.pdf>

TR-3431—Snapmanager 5.0 for SQL Server: Best Practices Guide

<http://media.netapp.com/documents/tr-3431.pdf>

TR3363—Microsoft SQL Server 2000: Sizing and Capacity Planning Guidelines

<http://www.netapp.com/library/tr/3363.pdf>

TR-3376—[SAN Booting with Windows Server Operating Systems](#)

<http://www.netapp.com/library/tr/3376.pdf>

TR- 3441—Windows Multipathing Options with Data ONTAP: FCP and iSCSI

[iSCSI multipathing possibilities on Windows with Data Ontap](#)

Miscellaneous

Mercer Study: Total Cost Comparison: Storage Solutions in Enterprise SAP Environments

<http://media.netapp.com/documents/ar1038.pdf>

SAP

SAP Adaptive Computing

<http://service.sap.com/adaptive>

Global SAP Homepage
<http://www.sap.com/index.epx>

SAP Service Marketplace
<http://service.sap.com/>

SAP Developer Network
<http://sdn.sap.com/>

SAP Help Portal
<http://help.sap.com/>

REVISION HISTORY

Date	Name	Description
04/25/2006	Marco Schoen	Creation

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.