



NETAPP TECHNICAL REPORT

SAP on Windows and Microsoft SQL Server with NetApp Storage and VMware

SAP Competency Center, NetApp
December 2008 | TR-3734

BEST PRACTICES

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 based on VMware® using a Microsoft® SQL Server® database.

TABLE OF CONTENTS

1	INTRODUCTION	3
2	STORAGE PROVISIONING AND MANAGEMENT	5
2.1	CONSOLIDATION	5
2.2	STORAGE LAYOUT	9
2.3	SIZING	14
2.4	INSTALLATION	15
2.5	STORAGE MIGRATION	16
3	SYSTEM MANAGEMENT AND MAINTENANCE	18
3.1	SAP HOMOGENEOUS SYSTEM COPIES	18
3.2	SAP UPGRADE	23
4	BUSINESS CONTINUANCE	27
4.1	BACKUP AND RECOVERY	27
4.2	HIGH AVAILABILITY	33
4.3	DISASTER RECOVERY	34
5	ARCHIVING AND COMPLIANCE: SAP ILM	38
6	CONCLUSION	40

1 INTRODUCTION

SCOPE

This document is intended to provide customers and partners with best practices for deploying NetApp storage systems in support of SAP Business Suite solutions running in a Windows environment based on VMware Virtual Infrastructure 3 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, VMware, and SAP products and was developed with input gathered from NetApp, VMware, SAP, Microsoft, and our customers.

BUSINESS CHALLENGES

Corporations deploying SAP solutions today are under great pressure to reduce total cost of ownership (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 become possible; they become requirements 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 its 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, and BI, 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.

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 in combination with server virtualization based on VMware's Virtual 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 the 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 and business continuity, satisfy regulatory and compliance standards, and improve resource utilization, a flexible and scalable storage network solution is required. The solution must also minimize complexity and reduce TCO.

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® 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. NetApp storage solutions enable customers to add storage when and where they need it, without disruption and at the lowest incremental cost.

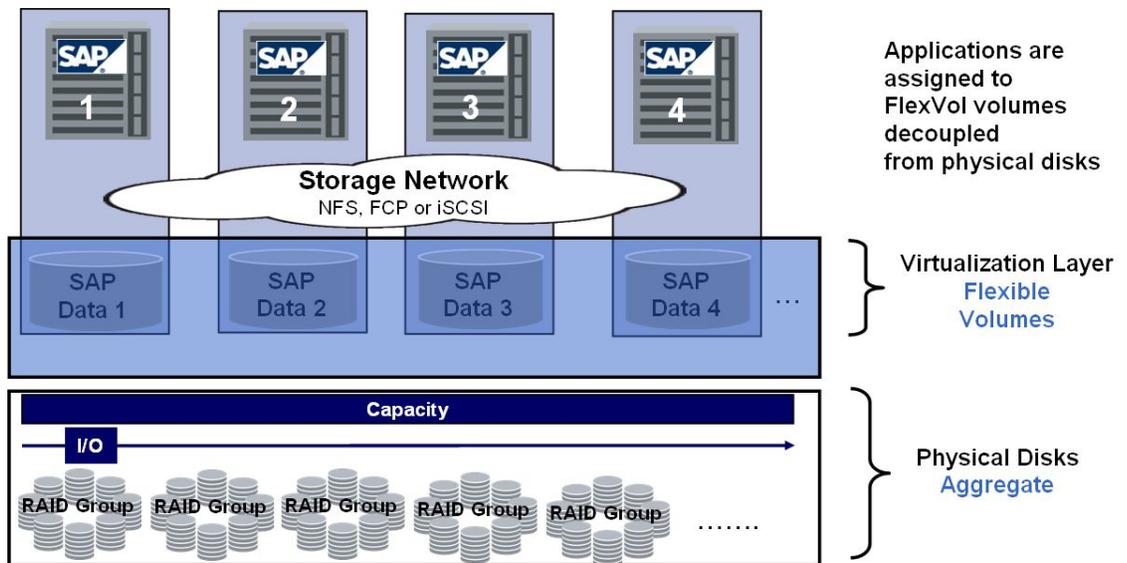


Figure 1) FlexVol technology.

NetApp FlexClone® 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.

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

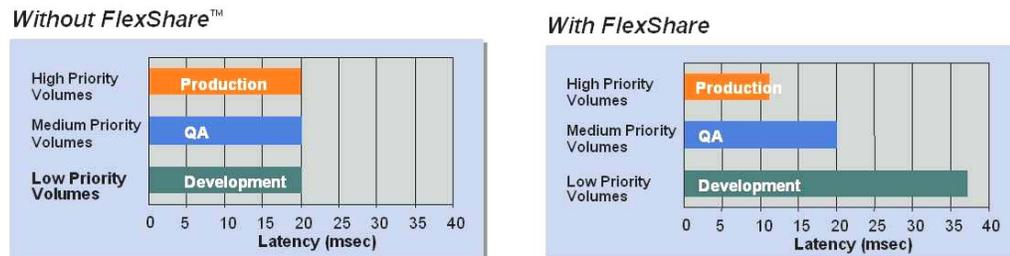


Figure 2) FlexShare workload prioritization.

SnapManager® for Virtual Infrastructure provides storage and virtual infrastructure administrators with an automated solution for data protection and recovery of virtual machines in a VMware ESX environment.

In addition to providing simplification and automation of the storage administrator tasks, SnapManager for Virtual Infrastructure empowers virtual infrastructure administrators to carry out virtual machine backup and recovery by themselves, within the policies set by the storage administrator.

Traditional physical and/or virtual server backups are CPU intensive and time consuming and use up too much extra storage. SnapManager for Virtual Infrastructure is removing these constraints by relying on NetApp Snapshot™ technology, which frees the servers for application-related workloads and provides granular restore capabilities at the virtual machine level.

Server virtualization, which contributes to the reduction of data center cost by increasing server asset utilization, has introduced a new set of challenges when it comes to data protection, business continuity, and provisioning.

The demand placed on traditional backup and DR in a virtualized environment can easily exceed the server and I/O capacity of the physical servers. Additionally, the mobility of virtual machines requires new levels of operational agility to manage the storage environment.

When managing multiple ESX servers, virtual infrastructure administrators can now rely on a seamless solution to centrally manage backups, recovery, and replication through an easy-to-use user interface, within the context and the framework of server virtualization with which they are familiar. Because NetApp Snapshot, SnapRestore®, and SnapMirror® are integrated in an easy-to-use user interface, the VMware administrator is now empowered with robust data protection capabilities that have now become VM aware through SnapManager for Virtual Infrastructure.

When managing the VMware ESX server environment and delivering the services necessary to support business needs, the VMware administrator can now rely on a seamless solution to protect and manage the virtual infrastructure environment:

- Automated data protection of data stores with granular recovery at the virtual machine level.
- Increased server utilization by eliminating interruptions and performance impact caused by the backup window. NetApp Snapshot and SnapRestore storage-based technology provides data protection and data recovery without any impact on the servers.
- Replication of data stores through NetApp SnapMirror, enabling DR.
- Easy-to-use user interface.

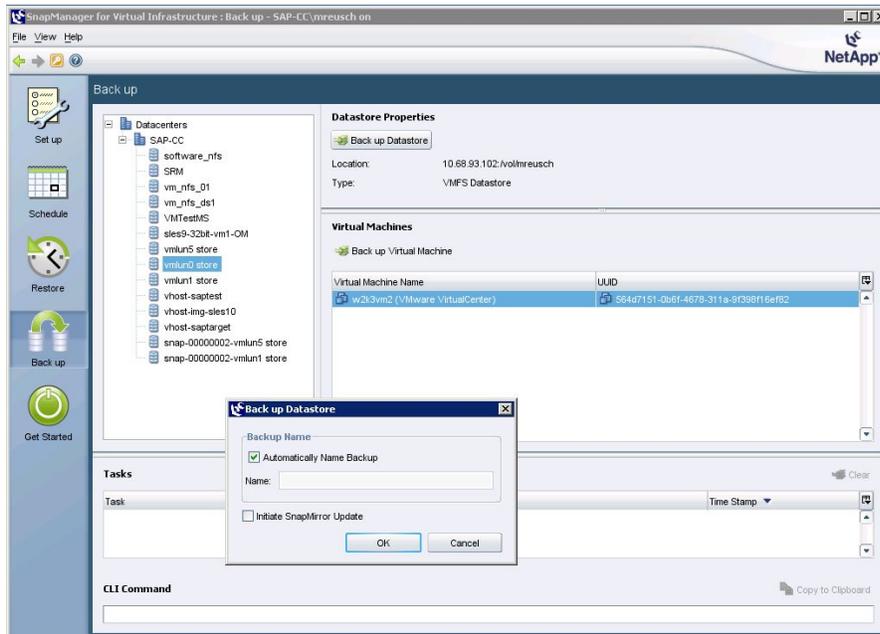


Figure 3) SnapManager VI backup.

With SnapDrive® for Windows software from NetApp, 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, also supporting virtual machines provided by VMware ESX Server. 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.

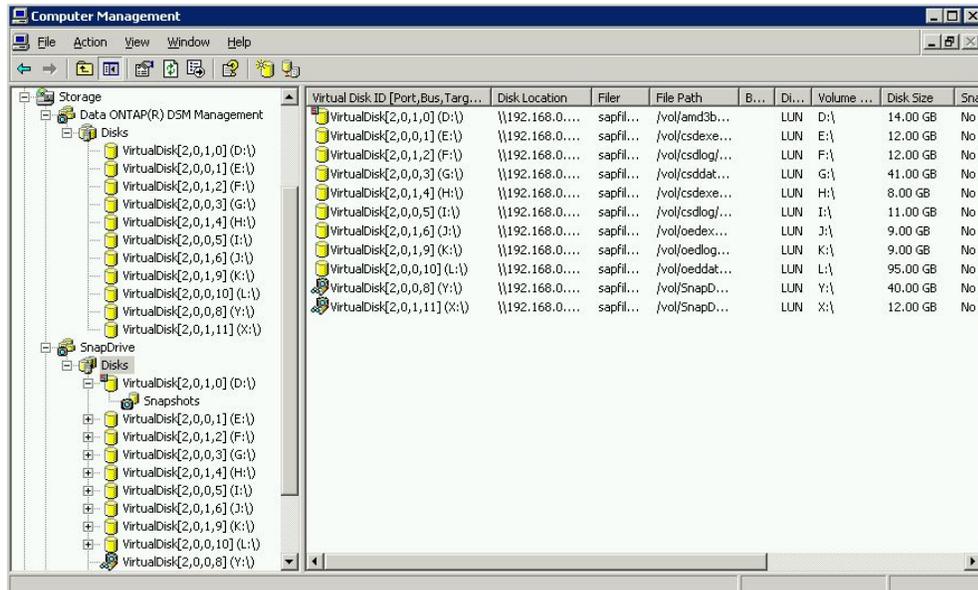


Figure 4) Storage management with SnapDrive.

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 Fibre Channel 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 protocol. 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.

SnapManager for SQL Server paves the way for database and storage administrators to simplify data management by 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.

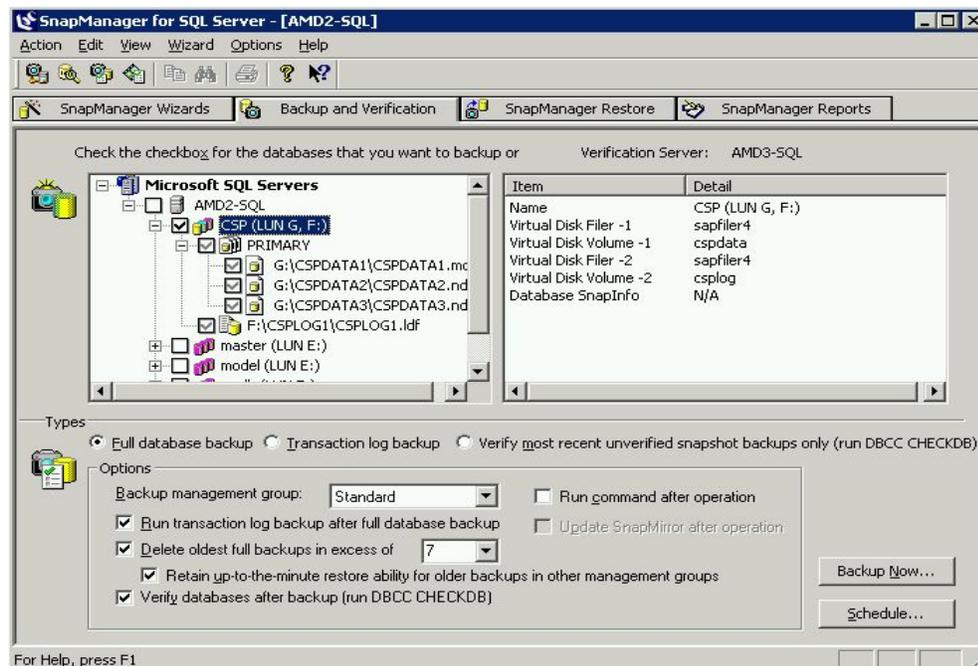


Figure 5) Database management with SnapManager for SQL Server.

The following list highlights some of the important benefits that 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.

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 of RAID 5 and is comparable to that of RAID 1. Only if three disks within the same RAID group fail at the same time will data loss occur. For information on RAID-DP, see [NetApp Data Protection: Double-Parity RAID for Enhanced Data Protection with RAID-DP](#).

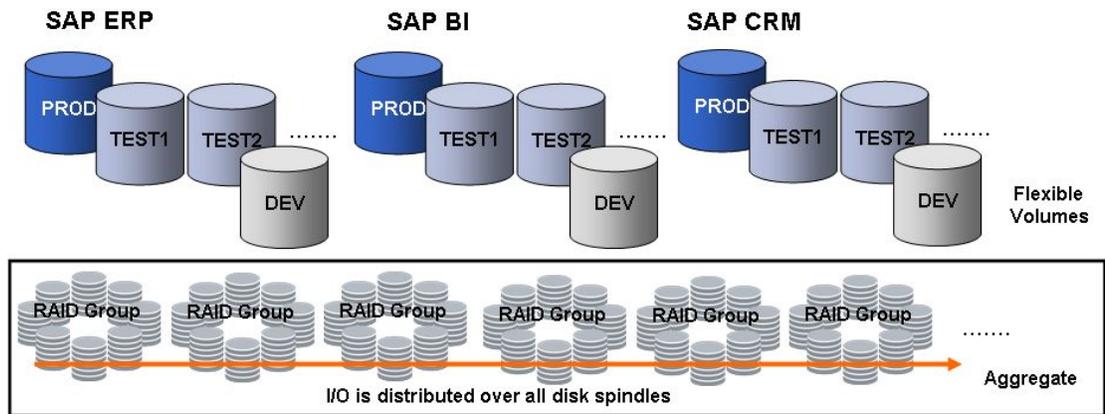


Figure 6) Aggregate layout.

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. For more information on FlexShare, see the [FlexShare Design and Implementation Guide](#).

VIRTUAL MACHINE LAYOUT: OPERATING SYSTEM AND SWAP

The data stores for the virtual machines, operating system and swap, should be stored separately in FlexVol volumes, independent of the SAP application and database files. These FlexVol volumes can be connected using the NFS, iSCSI, or FCP protocol. In case VMware Site Recovery Manager with SnapMirror will be used as the disaster recovery solution, iSCSI and FCP can be used. It is recommended to use the Virtual Machine File System (VMFS) for data stores connected via iSCSI or FC. This will enable the usage of

SnapManager for Virtual Infrastructure for backing up and restoring the operating system. In addition VMware Virtual Infrastructure functionality like template creation, cloning, and deploying of new virtual machines out of these templates are available. This is especially useful if a new virtual machine for an additional SAP QA or training system is needed. For further information how to provision NetApp storage for virtual machines, see the [NetApp and VMware Virtual Infrastructure 3: Storage Best Practices Guide](#)

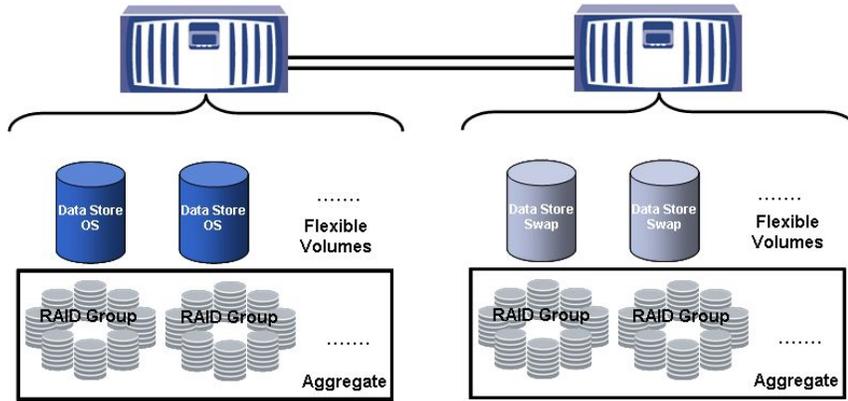


Figure 7) FlexVol volume: operating system.

FLEXVOL VOLUME LAYOUT: SAP SYSTEMS

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 the SnapManager for SQL Server 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, 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.

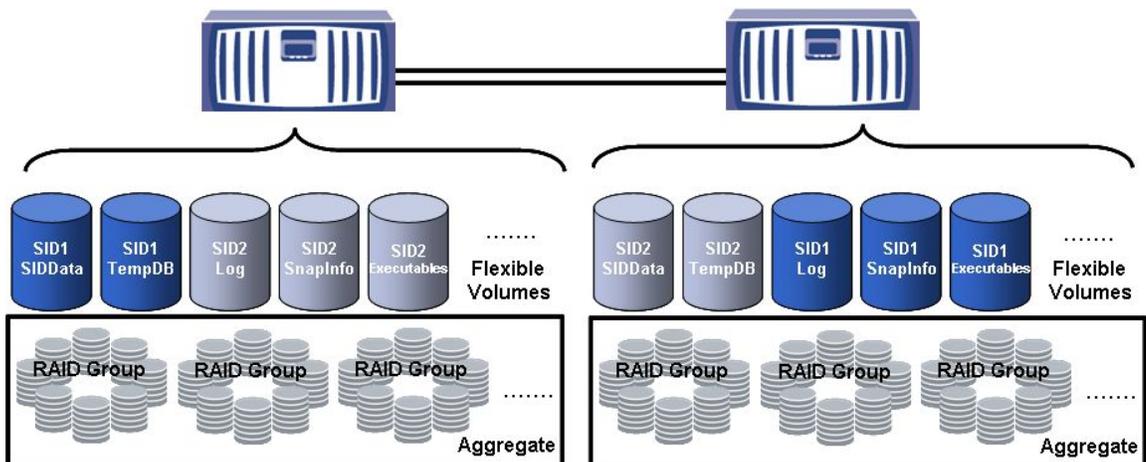


Figure 8) FlexVol volume layout: SAP system.

The following table shows the distribution of the file systems and 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	Log files	SnapInfo Dir
SIDData2	Quorum MSCS	System DBs		
SIDData3		SAP Executables		
SIDData#				

LAYOUT WITH METROCLUSTER

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

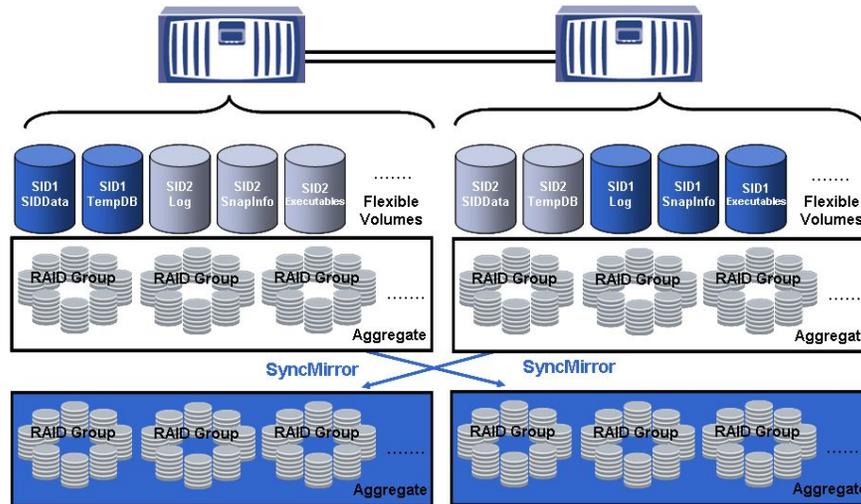


Figure 9) Storage layout with NetApp MetroCluster.

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.

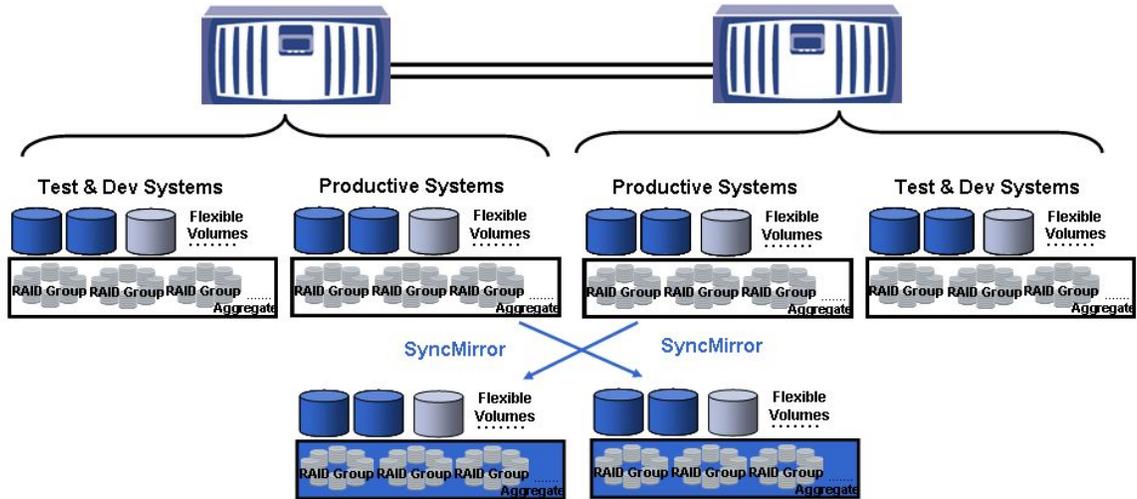


Figure 10) Storage layout with NetApp MetroCluster; only production systems are mirrored.

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

LUN Layout

The size of the database and the 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.

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

	Database Size	Number of LUN Data Files	Size of LUN Data Files	Number of LUN Log Files
Small System	< 200GB	1	100GB – 200GB	1
Medium System	200GB –	2-3	100GB – 350GB	1
Large System	> 1TB	> 3	200GB – 450GB	1

RECOMMENDED STANDARD LAYOUT

The following figure shows the LUN configuration for one 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 Server 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.

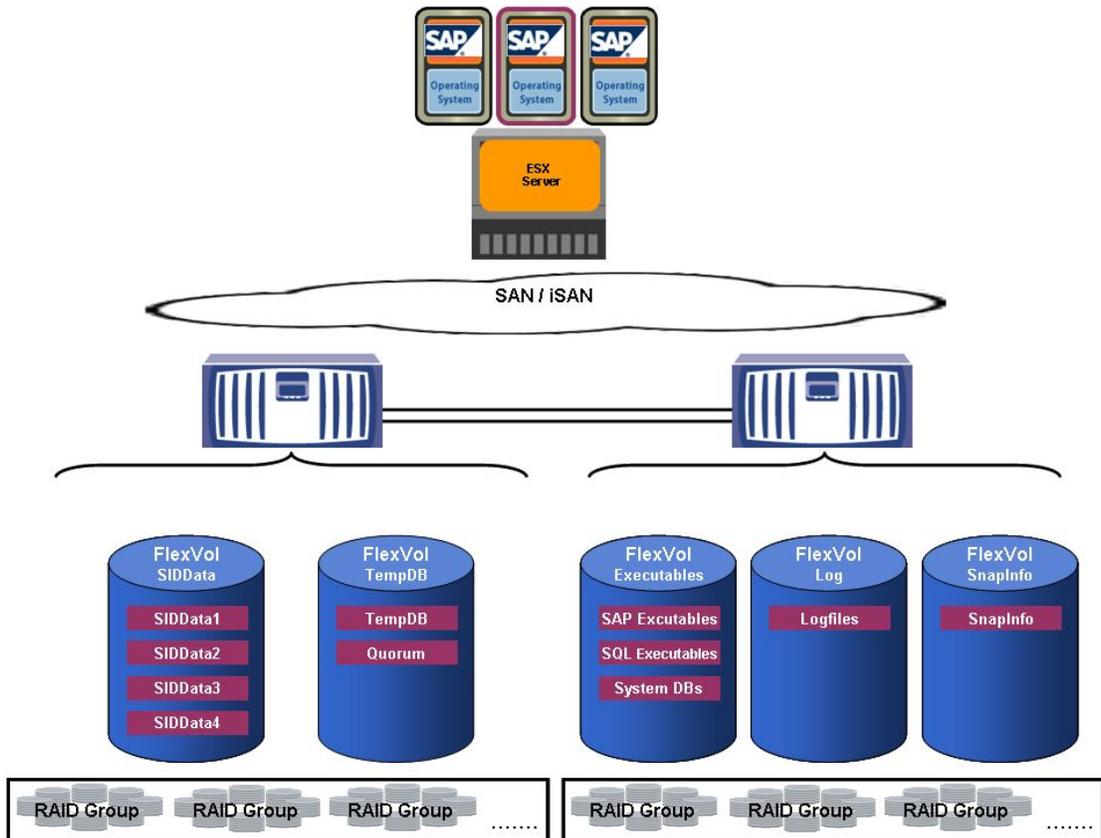


Figure 11) Standard storage layout for one SAP system.

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.

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. 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 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 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 are 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.

2.4 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 running as guest on a VMware ESX server using the FCP or iSCSI protocol. If the FCP protocol is used, raw device mapping (RDM) will be used to provide the LUNs to the operating system. In case of iSCSI the Microsoft iSCSI software initiator installed within the Windows guest operating system has to be used to provision the required LUNs.

GENERAL REQUIREMENTS

NetApp strongly recommends the use of SnapDrive for Windows, a NetApp 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. For information on SnapDrive system requirements, see www.netapp.com/us/products/management-software/snapdrive-windows.html.

For additional recommendations on system configuration and tuning, see:

- [Best Practices Guide: Microsoft SQL Server](#)
- [NetApp and VMware Virtual Infrastructure 3](#)

NETAPP STORAGE CONTROLLER CONFIGURATION

Snapshot backups for database applications won't be consistent from the database point of view without 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.

All LUNs should be created and connected to the Windows operating system using SnapDrive for Windows.

SnapDrive for Windows will set all required storage parameters such as FlexVol options, partition offset, and so on and manages the raw device mapping through the ESX Server.

It is strongly recommended that all LUNs should be stored in qtrees to be able to use all desired software functionality, such as 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 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 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 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

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 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 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 must 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 SQL Server Management Studio, disconnect the old storage from the database server, and connect the NetApp storage system and the LUNs that are 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 third-party SAN to a NetApp SAN storage system, NetApp has partnered with the best 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 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 complete, 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 www.netapp.com/us/services/professional/implementation.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. Table 3 summarizes the migration processes.

Table 3) Summary of migration processes.

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

3 SYSTEM MANAGEMENT AND MAINTENANCE

Testing properly before implementing changes in the production environment is critical. Risk can be reduced by improving the quality of testing and training.

Test teams need to be able to test using real business data. For example, if the data provided for testing is six months old, the results of the test could be significantly different from those that occur when the changes are migrated to the production environment.

Test teams can also improve the quality of testing by getting their own test environment, which will allow them to do independent parallel testing without interference from other test teams.

The same applies to training environments, which should be up-to-date images of the production or development environment. For example, if training is performed on a system that is not synchronized with the production environment, users will not be familiar with the business processes when they are asked to perform them in the live production system.

Another way to reduce risk is to make sure that there is a current backup of the production environment. When you execute changes in the production environment, several backups need to be created both before and after the changes are made. Being able to make fast backups and being able to restore and recover as fast as possible will allow you to do more testing before going live.

With conventional environments, all approaches to reducing risk result in additional costs and additional time requirements. Providing real business data for testing or training environments is a complex, time-consuming process. In addition, testing and training environments consume a significant amount of disk space, which adds additional costs and complexity. Therefore testing and training are typically done with outdated production data, and the number of available test systems is reduced to a minimum.

Making backups of large databases takes a significant amount of time, which is often not available. Also, the time for restoring the production environment needs to be considered when implementing production changes. Additional testing is therefore reduced to a minimum because of time limitations.

3.1 SAP HOMOGENEOUS SYSTEM COPIES

CAPACITY REQUIREMENTS

When creating SAP system copies with most storage architectures, space must be allocated to accommodate the entire size of the source database. This can drastically increase the amount of storage required to support a single production SAP instance.

During a typical project a 1TB SAP production system will be copied to a quality assurance system, a test system, and a training system. With conventional storage architectures, this requires an additional 3TB of storage. Furthermore, it requires a significant amount of time to first back up the source system and then to restore the data to the three target systems.

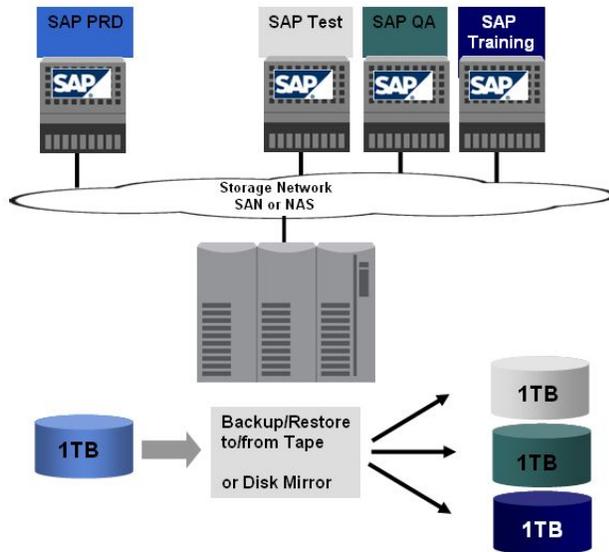


Figure 12) SAP system copy: standard approach.

In contrast, when using NetApp FlexClone technology to create SAP system copies, only a fraction of the storage space is required. NetApp FlexClone technology uses Snapshot copies to create copies of SAP systems. Because the data is not copied but referenced in place, the amount of storage required is limited to only data that is changed at the source or the target system. As a result, the storage requirements for a system copy in a NetApp storage environment can be kept to around 10% of the size of the source database.

SAP system copies are made using the NetApp FlexClone feature integrated into SnapManager for SQL Server. A FlexClone copy is based on a Snapshot copy of the source FlexVol volume and is created in a few seconds without interrupting the operation on the source system. FlexClone copies store only changed blocks between the source FlexVol volume and the FlexClone copy and therefore significantly decrease the disk space needed for SAP system copies.

On the source system a database-consistent Snapshot copy of the data files will be created. This is done during online operation and has no performance impact on the source system. This step can therefore be carried out at any time.

On the target system this Snapshot copy will be the base for the FlexClone copy. The creation of the FlexClone copy only takes a few seconds. It is then connected at the target system. The subsequent steps at the target system are necessary to change the database and the SAP SID. In addition post processing tasks specific to SAP need to be performed.

All the above steps are fully automated and do not need any manual interaction. A SAP system copy can be made in a few minutes using the NetApp solution. Figure 13 shows a high-level graphical representation of this infrastructure.

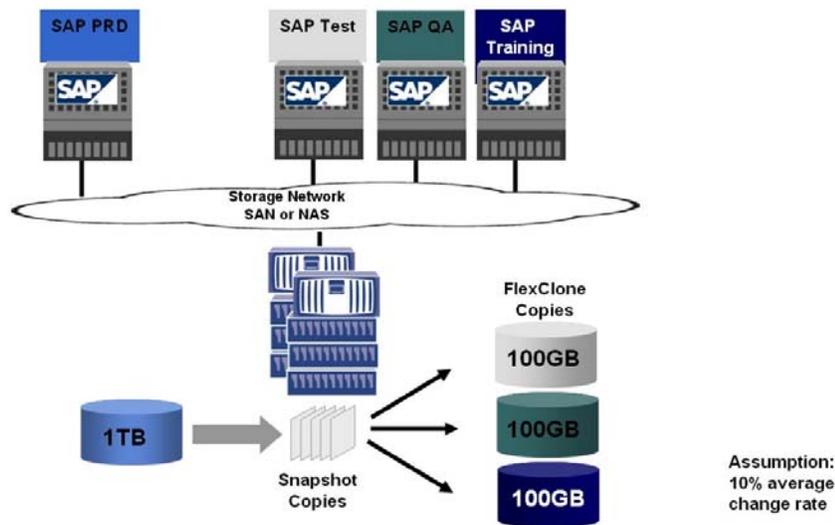


Figure 13) SAP system copy: NetApp approach.

TIME REQUIREMENTS

The time required to create a SAP system copy can be subdivided into four parts:

- Time to create a backup of the source system
- Time to restore the backup to the target system
- Time for initializing the log area of the new database and recover the backup
- Time to perform SAP application post processing

In a conventional system copy process, the data is backed up to tape and then restored, which takes a great deal of time. If an online backup is used, there is no downtime for the source system; however, there might be a performance impact on the source system during the backup. Also, the time required to recover the database and make it consistent is greatly increased—possibly adding hours to the system copy process—because of the large number of logs that need to be applied. If an offline backup is used, then the source system is shut down, resulting in a loss of productivity.

In contrast, a backup made on NetApp storage takes just a few seconds. As a result with MaxDB no logs need to be applied to the database.

For both processes, the SAP postprocessing time is the same.

Figure 14 shows the difference between the amounts of time that are spent testing with NetApp storage versus the time spent testing using a conventional approach. The NetApp approach yields more than twice the amount of testing as the conventional approach yields in the same period of time. The ability to run more test cycles allows customers to correct more issues before going live, reducing risk and increasing end-user satisfaction and productivity.

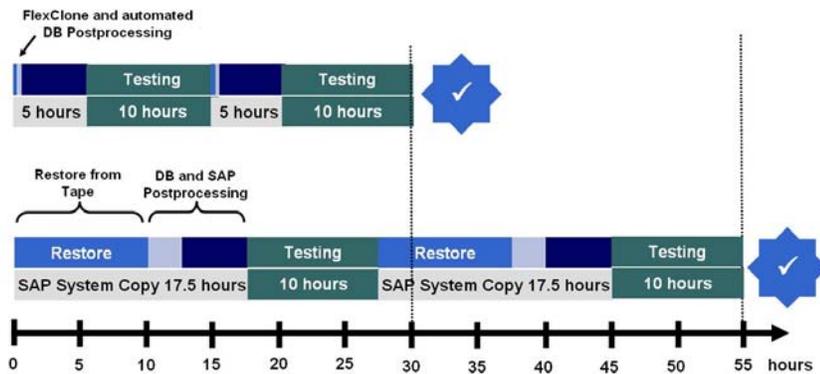


Figure 14) SAP system copy: time requirements.

NETAPP SOLUTION FOR SAP SYSTEM COPIES

SAP system copies are accomplished using SnapManager for SQL Server utilizing the NetApp FlexClone feature to clone SAP databases. A FlexClone copy is a writable point-in-time image of a NetApp 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 tasks specific to SAP for a homogeneous system copy.

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 and Java system, SAPINST must be executed to export the deployment manager configuration and to save the Java applets stored in the file system 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 tasks specific to SAP. In addition, post processing tasks specific to SAP 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. Also it is possible to clone a SAP system which runs on a physical server to a virtual machine and vice versa. For more details please, refer to the [SAP on Windows and Microsoft SQL Server System Copy Implementation Guide](#).

VIRTUAL MACHINE PROVISIONING

NetApp Snapshot and FlexClone technologies offer the possibility to create a SAP system copy in a very short amount of time and without the need to provide the full amount of disk space for the copy. This offers the possibility to create additional SAP systems for parallel testing of developments, upgrades, and other projects or training of users.

Of course, additional servers are needed for additional SAP systems; in an environment based on VMware new virtual machines are needed. VMware Virtual Infrastructure offers template creation of virtual machines (VMs), which can be cloned to new virtual machines. This cloning functionality integrates the Microsoft Sysprep tool for automated Windows operating system setup. For details, refer to the VMware Virtual Infrastructure documentation.

Usually the base VMs with the operating system are nearly identical from the data point of view. Therefore NetApp [deduplication](#) software eliminates redundant data objects and references just the original object, and an immediate benefit is obtained through storage space efficiencies. The result is twofold:

- Cost benefit: Reduced initial storage acquisition cost or longer intervals between storage capacity upgrades
- Management benefit: The ability to store “more” data per storage unit or retain online data for longer periods of time

3.2 SAP UPGRADE

BUSINESS CHALLENGES

A SAP upgrade is a typical example of a complex customer project within a SAP landscape. Similar to upgrade projects, the described issues also occur during Unicode conversions, when applying support packages, and during custom developments. The procedures to overcome those issues are similar, too.

Running a new version of the SAP applications is often a prerequisite in order to innovate business processes. Over time, SAP versions will also reach the end of their maintenance period. Therefore SAP customers need to go through a SAP upgrade project at regular intervals.

Customers face several challenges in a SAP upgrade project:

- **Costs.** The SAP upgrade project consumes large amounts of employee time, especially IT staff time. Therefore project time needs to be reduced.
- **Delayed innovation.** Business processes are affected during the upgrade project time 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.
- **Risk.** An upgrade causes change, which introduces risk that the business processes might not work as expected after the upgrade. The risk needs to be reduced by increasing the frequency and improving the quality of testing and training.
- **Production system downtime.** During the upgrade of the production system, the system will not be available. Production system downtime has to be minimized.

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

During a SAP upgrade project, SAP basis administrators need to create several system copies to test the upgrade process with current data from the development or production SAP system. The creation of a SAP system copy usually takes several days and might negatively affect the production environment. In addition many manual steps must be performed, consuming valuable IT staff time.

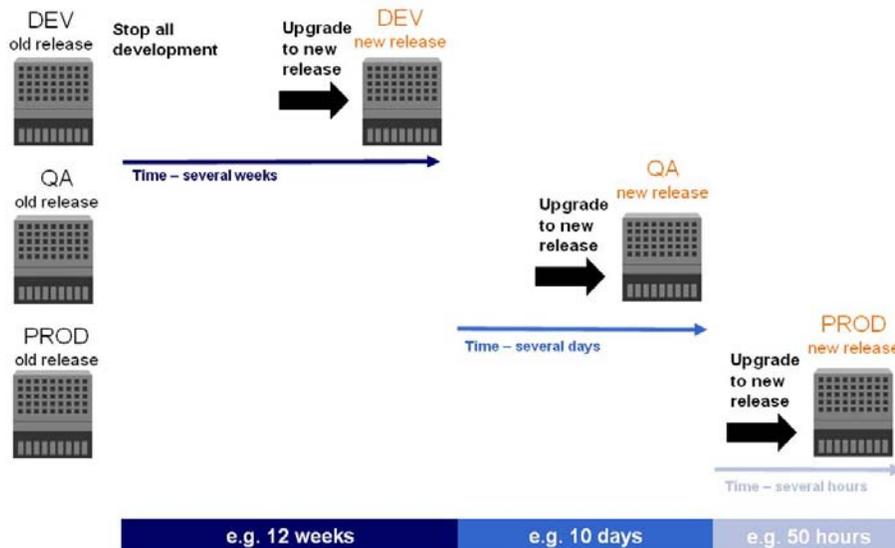


Figure 15) SAP upgrade: overview.

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, 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 for 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 provided using the NetApp system copy solution. Using this solution will significantly reduce the time and resources needed for the system copy. Reducing the time is critical because in most cases the copy is created 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 to restart the upgrade phase.

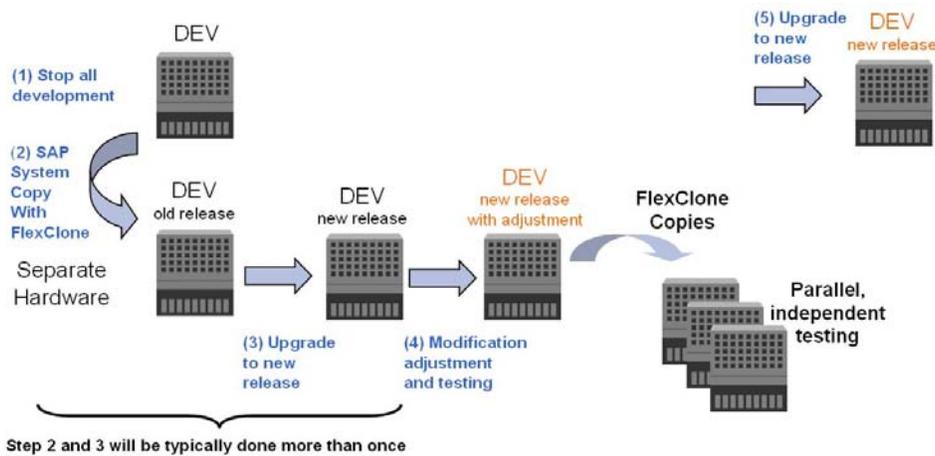


Figure 16) SAP upgrade: development system.

UPGRADING THE QUALITY ASSURANCE SYSTEM

The quality assurance system is upgraded using a fresh system copy of the production SAP system. One important result of this upgrade is testing the upgrade with production data. The NetApp 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 usually made more than once to support multiple tests of the upgrade process and also multiple test cycles. 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 you to restart an upgrade phase or restart the import.

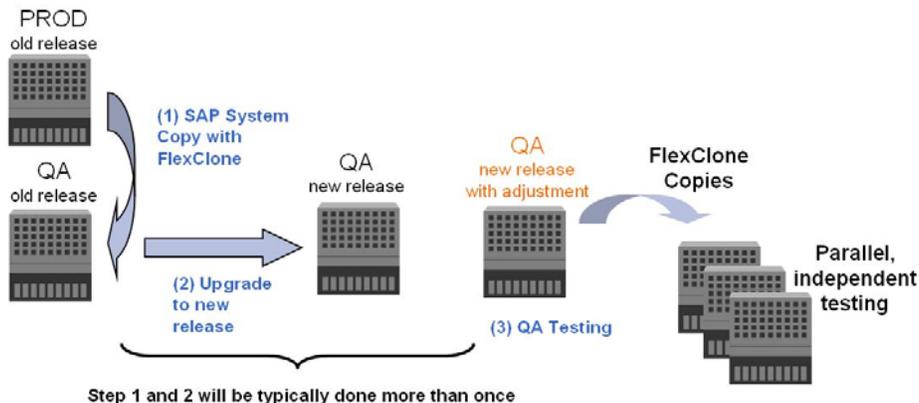


Figure 17) SAP upgrade: quality assurance system.

UPGRADING THE PRODUCTION SYSTEM

Scheduling is extremely important when upgrading the production system, because the system is not available at various stages during the upgrade. The schedule has to provide time to restore 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 for importing the transports for the modification adjustment, a 48-hour weekend might not be enough time to complete the upgrade.

The production system upgrade 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 previous release level.

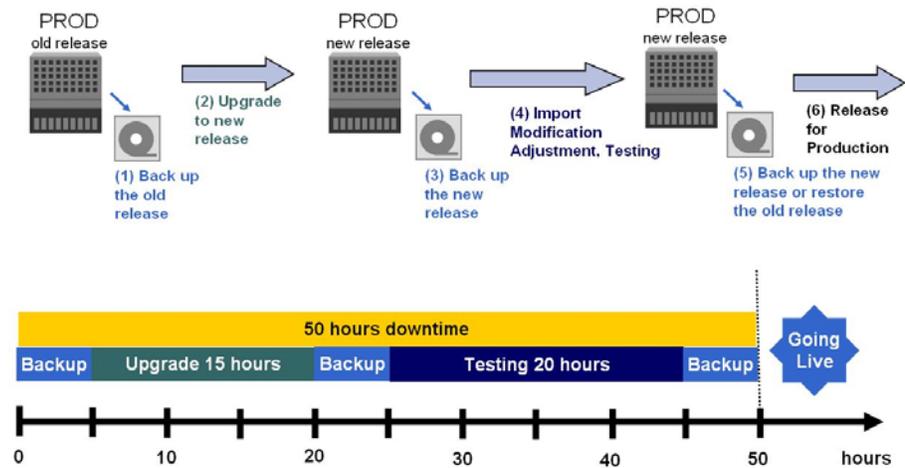


Figure 18) SAP upgrade: production system.

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

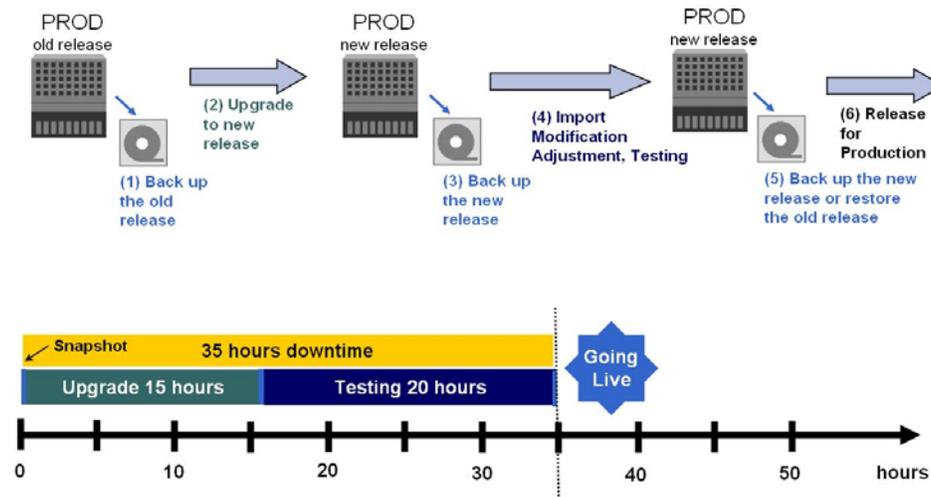


Figure 19) SAP upgrade: production system with NetApp. [[NOTE: Figures 19 and 20 have the same title. Please change one of them.]]

Reducing the time needed for backup and restore will allow minimizing the upgrade downtime of the production SAP system. One option is to go live earlier with the upgraded release.

The other option is to use more time for testing the upgraded release before going live. More testing reduces the risk and provides more time to fix any issues that arise.

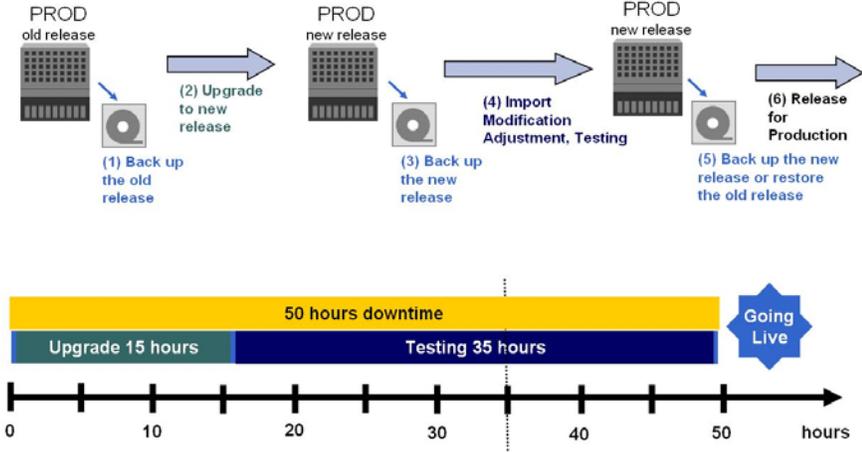


Figure 20) SAP upgrade: production system with NetApp.

4 BUSINESS CONTINUANCE

4.1 BACKUP AND RECOVERY

BUSINESS CHALLENGES

Corporations today require their SAP applications to be available 24 hours a day, seven days a week. Consistent levels of performance are expected, regardless of increasing data volumes and routine maintenance tasks such as system backups. Performing backups of SAP databases is a critical task and can have a significant performance impact on the production SAP system. Because backup windows are shrinking and the amount of data that needs to be backed up is increasing, it is a complex task to define a point in time during which backups can be performed with minimum impact on the business process. The time needed to restore and recover SAP systems is of particular concern because the downtime of SAP production and nonproduction systems must be minimized.

The following summarize SAP backup and recovery challenges:

- Performance impact on production SAP systems. Backups typically have a significant performance impact on the production SAP system because there is a heavy 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 production SAP system, backups can be made only during times when there is little dialog or a small number of batch activities taking place 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 result in ongoing investments in the backup infrastructure: more tape drives, new tape drive technology, faster storage networks, and so on. 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. 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 the case of a database failure. The backup and recovery architecture must be designed based on the maximum acceptable RTO.
- 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. The go/no-go decision is based on the amount of time required to restore and recover the database from the backup that was created previously. The option to restore very quickly allows more time to solve problems with the upgrade rather than to restore the backup.

NETAPP SOLUTION

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

NetApp [SnapManager for SQL Server](#) is able to 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 activity periods. SAP and NetApp customers typically schedule several Snapshot online backups during the day—for instance, every four 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 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 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 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 NetApp SnapManager for SQL Server (SMSQL) always consists of two parts:

- Backup and restore/recovery using Snapshot and SnapRestore through SMSQL
- 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 storage system with nearline storage (SATA disks) and SnapVault software: The primary storage communicates directly with the secondary storage 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 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 automatically starts the SnapVault transfer of 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.

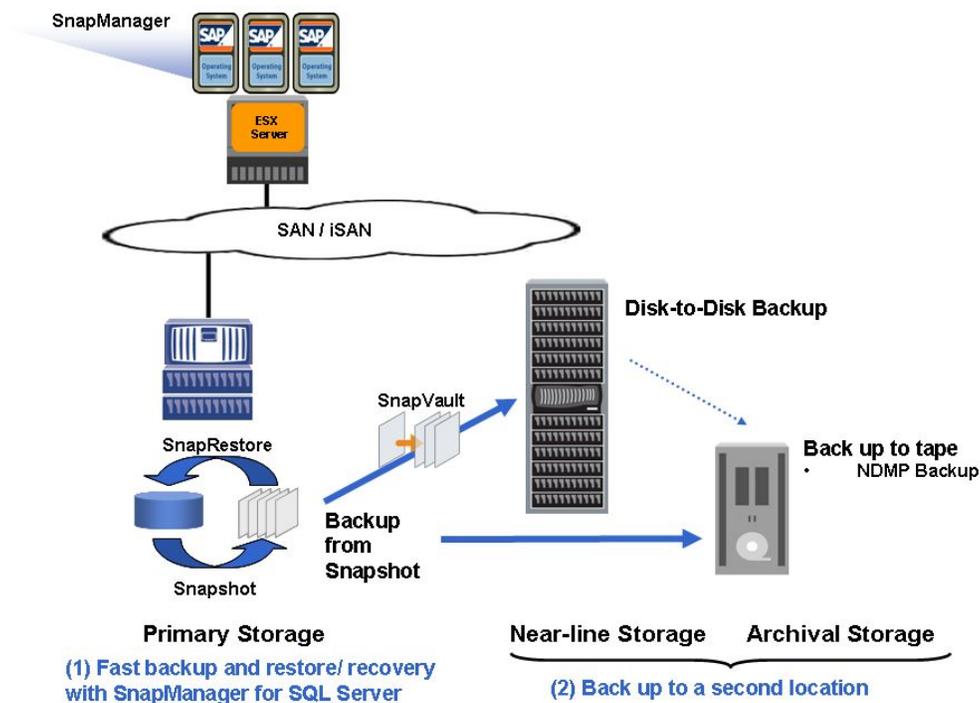


Figure 21) SAP backup and recovery: NetApp solution overview.

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).

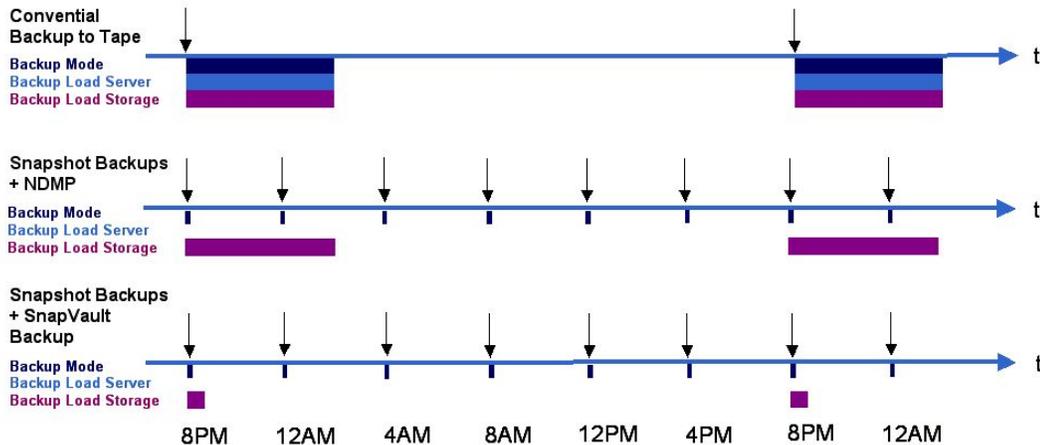


Figure 22) Comparison of different backup approaches.

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.

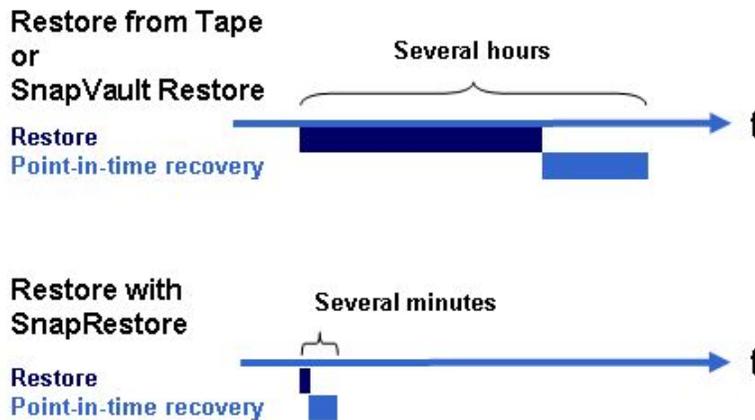


Figure 23) Time needed for restore and recovery.

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 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 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.

ACCELERATING TEST AND TRAINING CYCLES

When the amount of time required to perform system copies and database restores is reduced, the amount of time left to test is increased. This increase in time means more test cycles can be performed, and any problems encountered can be corrected quickly in the source system. Because the NetApp Snapshot copy process takes only a few seconds and does not impact the source system, a new source image for the test systems can be created quickly and easily. The parallel tests can then be rerun against the adjusted system. This process can be repeated until all tests are successful or the testing period ends. The end result is that more and higher-quality test cycles can be performed in less time.

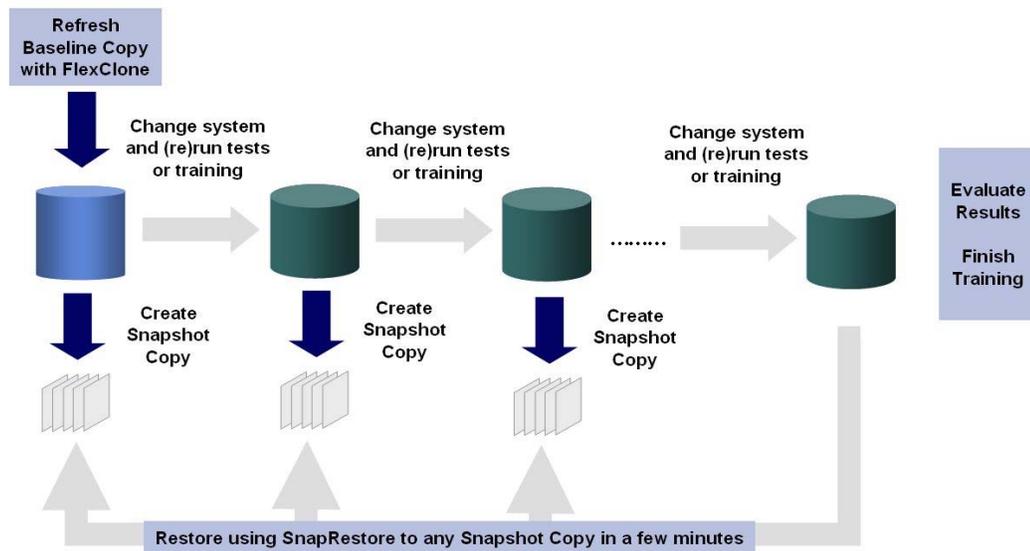


Figure 24) SAP testing cycle.

The same concepts also apply to training systems. Because you can rapidly restore system backups from a NetApp Snapshot copy using SnapRestore, training system refreshes take only minutes. This allows more training courses to be run in a shorter period of time because days are not required to restore the training system to its original state.

OPERATING SYSTEM BACKUP USING SNAPMANAGER FOR VIRTUAL INFRASTRUCTURE

NetApp SnapManager for Virtual Infrastructure (SMVI) offers the same advantages for backing up and restoring of virtual machine images (operating system) as SnapManager for SQL Server does for the SAP databases. SMVI is integrated into VMware Virtual Infrastructure and VMotion® aware. It supports backup and restore at data store and VM level and offers verification of the backups.

4.2 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.

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

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 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. To make sure 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.

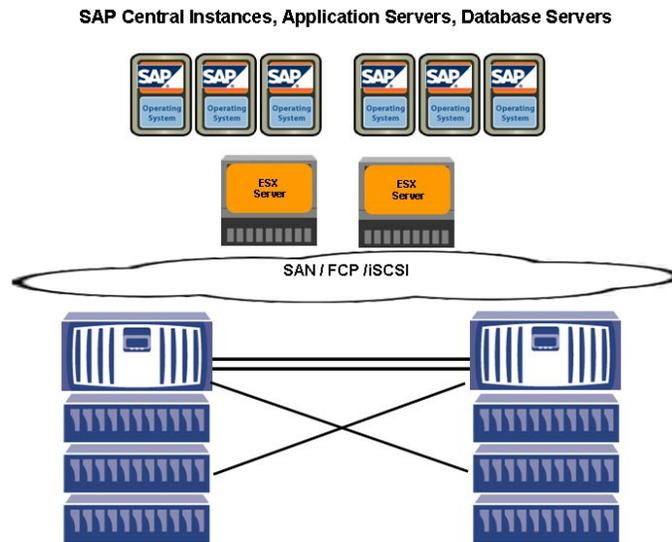


Figure 25) NetApp clustered storage system solution.

Conclusion: The NetApp clustered failover technology provides an extremely robust high-availability solution.

- A cluster has an availability level of 99.99+%.
- Both storage controllers in the cluster can be used actively, providing high availability for both the database server and the central instance.
- A clustered storage system is recommended if server clustering (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 makes sure of 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](#) software delivers a 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 the fastest recovery. SnapMirror is integrated with SnapManager for SQL Server, SnapManager for Virtual Infrastructure, and VMware Site Recovery Manager.

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 NetApp backup and recovery solution SnapManager for SQL Server and SnapManager for Virtual Infrastructure. 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, enabling a minimum of data loss. This can be achieved by scheduling SnapManager log backups every 10 minutes.

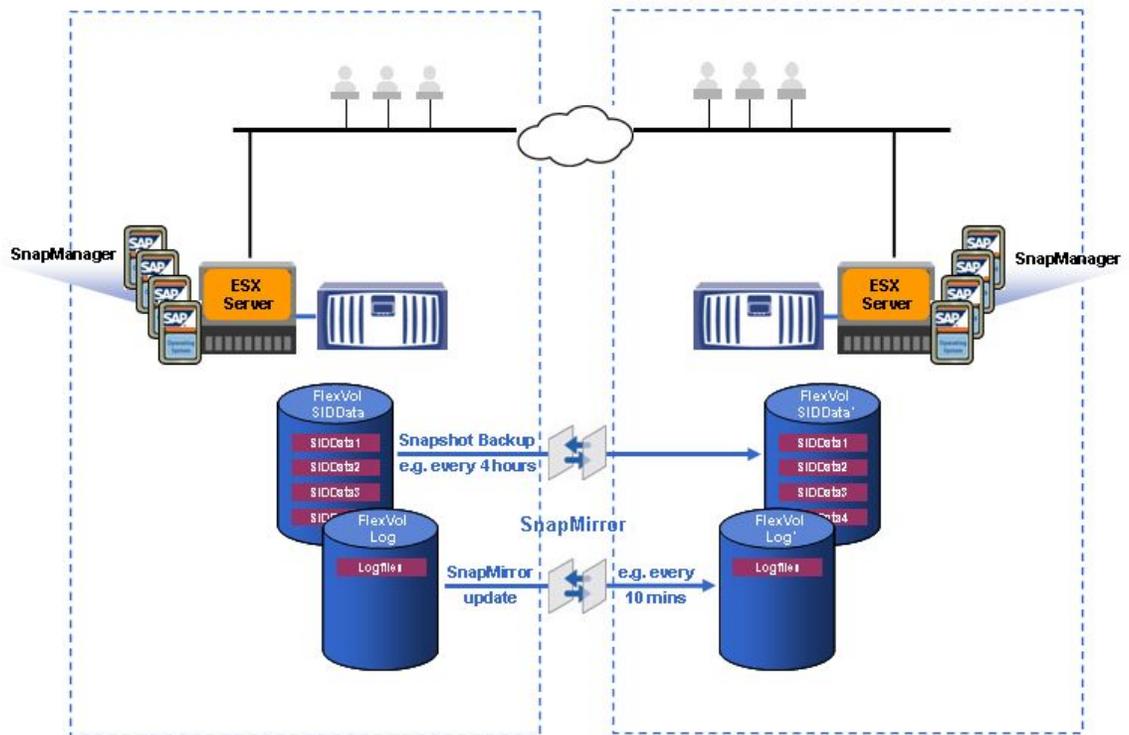


Figure 26) Disaster recovery with SnapMirror.

VMware Site Recovery Manager and SnapMirror

Perhaps the most difficult and time-consuming part of DR failover in a VMware environment is the execution of the steps necessary to connect, inventory, reconfigure, and power up virtual machines at the DR site. VMware has elegantly solved these problems with the introduction of Site Recovery Manager. Site Recovery Manager, or SRM, enables two separate VMware environments, the primary and the DR (or paired) sites, to communicate with each other. Virtual machines can be quickly and easily collected into protection groups at the primary site that share common resources and can be recovered together. These protection groups are configured into recovery plans at the DR site.

Upon the execution of a DR plan SRM will:

- Quiesce and break the NetApp SnapMirror relationships.
- Map the LUNs to existing igroups.
- Trigger the DR ESX hosts to rescan and detect the storage.
- If desired, power off virtual machines, such as test/dev instances, at the DR site, freeing compute resources.
- Reconfigure the virtual machines as defined for the network at the DR site.
- Power on the virtual machines in the order defined in the recovery plan.

By adding SRM into the NetApp and VMware environment, DR testing can be completed in a matter of minutes or just a few hours.

Upon the execution of a DR test SRM will:

- Create FlexClone volumes of the FlexVol volumes on the DR storage appliance.
- Map the LUNs contained within these FlexVol volumes to existing igroups.

- Trigger the DR ESX hosts to rescan and detect the storage.
- Connect the virtual machine network adapters to a private test bubble network.
- Reconfigure the virtual machines as defined for the network at the DR site.
- Power on the virtual machines in the order defined in the recovery plan.

For detailed information for SRM with SnapMirror, refer to [VMware Site Recovery Manager in a NetApp Environment](#).

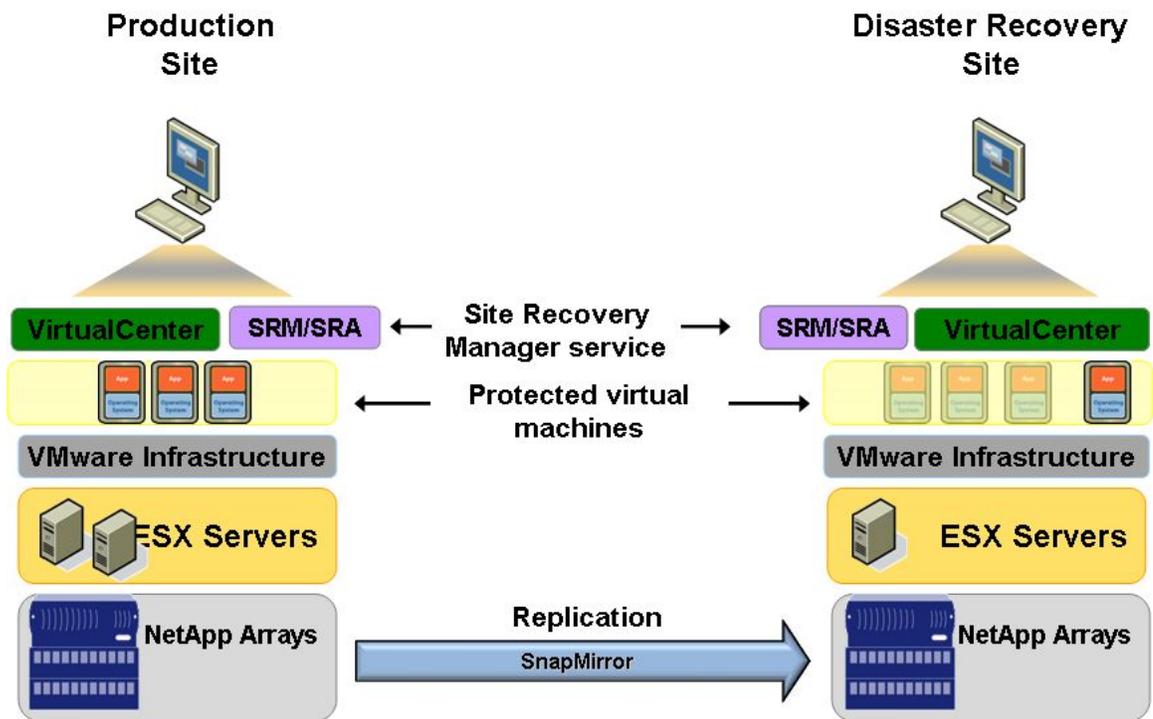


Figure 27) VMware Site Recovery Manager with SnapMirror.

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 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 disk. MetroCluster extends this cluster configuration to remote locations up to 30 kilometers. 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.

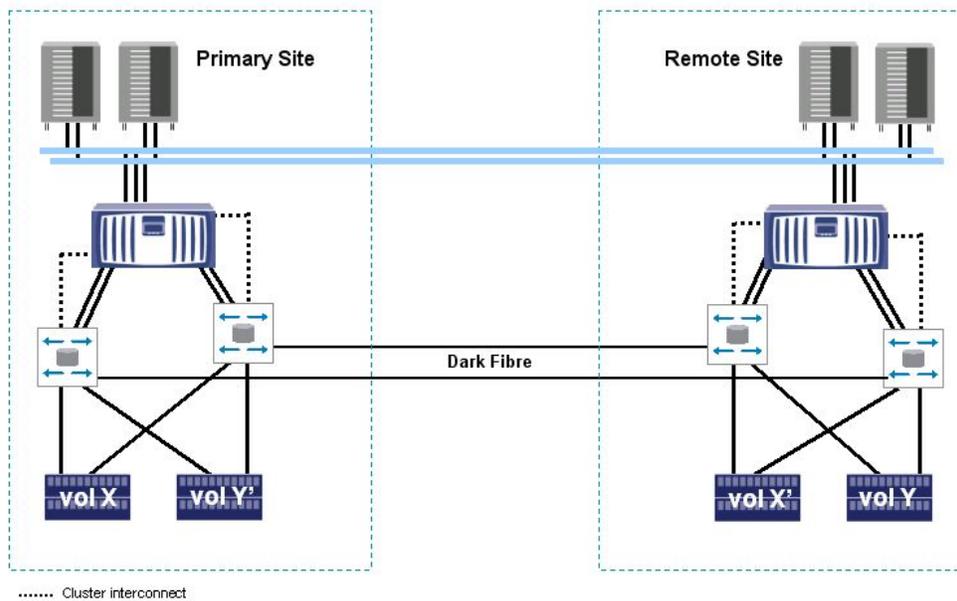


Figure 28) MetroCluster over direct Fibre Channel switch connection.

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

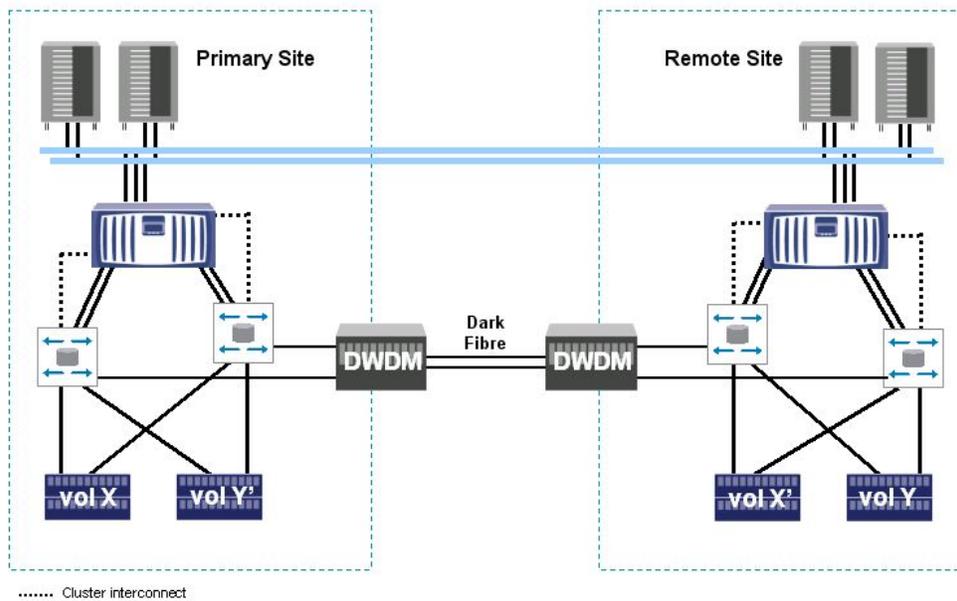


Figure 29) MetroCluster over Fibre Channel and DWDM switch infrastructure.

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: SAP ILM

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 total cost of ownership (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 Enterprise Storage Group (ESG), 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.

NETAPP SOLUTION

NetApp offers solutions for SAP ADK/ArchiveLink and SAP Information Lifecycle Management (ILM) which consists of three solutions: data archiving using the WebDAV protocol, Information Retention Manager, and Retention Warehouse.

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, used by SAP ILM, 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 for SAP ILM - WebDAV Storage Interface 2.0 (BC-ILM 2.0) in combination with PBS ContentLink. More information can be found at the SAP [Partner Information Center](#).

Once archive files have been created, the data marked for archive can be deleted from the source system. The archiving data will then be transferred directly from the primary storage system to the archive server. The SAP Retention Manager will set and manage the retention of the archived data based on legal requirements, including legal hold management. For system decommissioning of legacy systems SAP offers the Retention Warehouse. Detailed information can be obtained from [SAP](#).

NetApp solutions for SAP archiving such as NetApp FAS systems 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 near-line storage is the preferred compliance and archive storage subsystem for SAP landscapes. The FAS 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 NetApp systems to tolerate double 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. 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 API. You can perform all operations specific to SnapLock, 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 OpenText, FileNet, PBS, or other software vendors 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 (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 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 platforms.

Conclusion: NetApp provides a flexible, scalable, and secure solution for SAP ILM, 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 OpenText, FileNet, PBS, and other software vendors.
- 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 near-line storage, a more cost-effective alternative for archiving SAP data than adding database storage or processing power.

6 CONCLUSION

As SAP landscapes grow to support more and more business-critical applications, the job of maintaining those landscapes, including virtual machine management, becomes increasingly complex. The NetApp solutions for SAP on VMware bring together technologies that simplify and accelerate this process and align with the SAP application lifecycle.

The NetApp solutions for SAP on VMware accelerate upgrades and changes; enable fast SAP system copies; and provide simplified, economical, and highly available disk-based archiving. NetApp solutions help enterprises to reduce cost and complexity, minimize risk, and control change in their SAP environments.



www.netapp.com

© 2008 NetApp. All rights reserved. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, Data ONTAP, FlexClone, FlexShare, FlexVol, NearStore, RAID-DP, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapRestore, Snapshot, SnapVault, and SyncMirror are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Microsoft, Windows, and SQL Server are registered trademarks of Microsoft Corporation. Java is a trademark of Sun Microsystems, Inc. SAP and NetWeaver are registered trademarks of SAP AG. VMware and VMotion are registered trademarks of VMware, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.