



NETAPP TECHNICAL REPORT

SAP ON UNIX AND SAP MAXDB WITH NFS AND NETAPP STORAGE

SAP Competency Center, NetApp
Werner Thesing, liveCache/MaxDB Technology, SAP AG
May 2008 | TR-3689-0708

EXECUTIVE SUMMARY

This document provides customers and partners with the best practices for deploying NetApp storage in support of SAP Business Suite solutions running in a UNIX and NFS environment using an SAP MaxDB database.

TABLE OF CONTENTS

1	INTRODUCTION	3
2	STORAGE PROVISIONING AND MANAGEMENT	5
2.1	CONSOLIDATION	5
2.2	STORAGE LAYOUT.....	7
2.3	SIZING	10
2.4	INSTALLATION.....	11
2.5	STORAGE MIGRATION.....	15
3	SYSTEM MANAGEMENT AND MAINTENANCE	16
3.1	SAP SYSTEM COPIES	16
3.2	SAP UPGRADES	20
4	BUSINESS CONTINUANCE.....	24
4.1	BACKUP AND RECOVERY	24
4.2	HIGH AVAILABILITY	28
4.3	DISASTER RECOVERY	29
5	ARCHIVING AND COMPLIANCE	31
6	CONCLUSION.....	33

1 INTRODUCTION

SCOPE

This document is intended to provide customers and partners with the best practices for deploying NetApp storage systems in support of SAP Business Suite solutions running in a UNIX or NFS environment using a MaxDB 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 based on the interaction of technical staff from NetApp, SAP, 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 profiles. These profiles can also fluctuate depending on the phases of a project implementation and whether the project is focused on an existing SAP implementation or a brand new one.

In summary, for each instance of SAP running in production, there can be as few as two and perhaps five or more instances supporting it. Deploying three SAP applications, such as 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 allows business owners and IT departments to work more efficiently and effectively toward the goal of improving enterprise business processes.

Storage consolidation with NetApp assures 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.

SAP MAXDB AND LIVECACHE

SAP MaxDB and liveCache are based on the same architecture. This technical report covers both instance types. Minor differences will be described below.

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, make sure of 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 automatically and enables you to create multiple flexible volumes on a large pool of disks (aggregate). This flexibility means that operations can be simplified, utilization and efficiency can be increased, and changes can be applied more quickly and seamlessly. NetApp storage solutions enable customers to add storage when and where they need it, without disruption and at the lowest incremental cost.

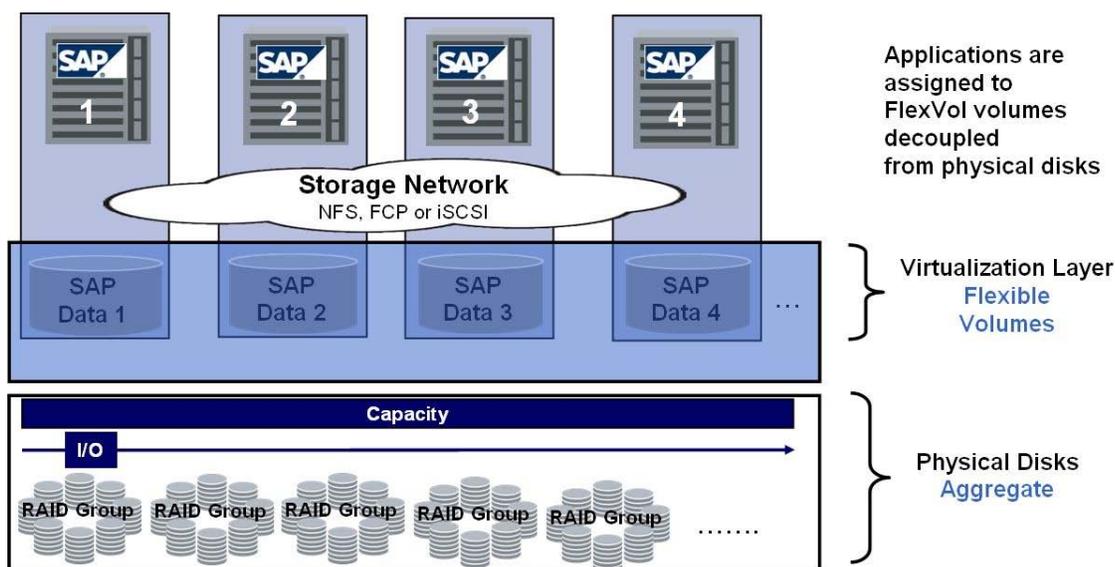


Figure 1) FlexVol technology.

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

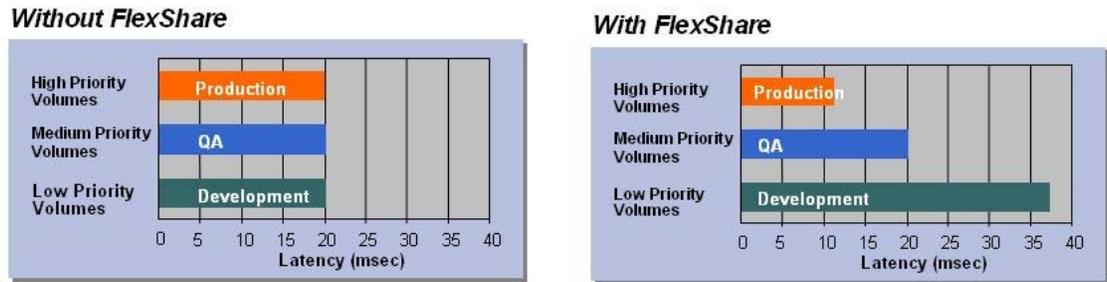


Figure 2) FlexShare.

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

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

The aggregates should be configured with RAID-DP®, which offers a high level of data protection. The reliability of RAID-DP is far better than that of RAID 5 and very close 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](#).

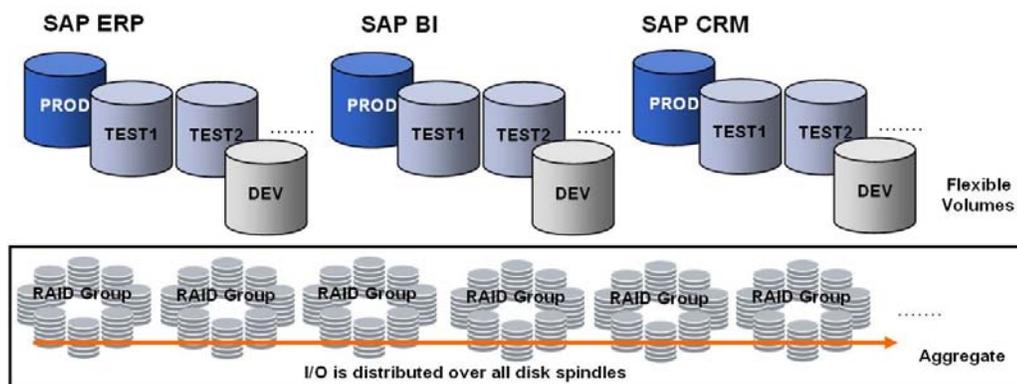


Figure 3) Aggregate layout.

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

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

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

FLEXVOL VOLUME LAYOUT

Each SAP system uses three FlexVol volumes:

- One volume for the database data files
- One volume for the log files and the saved log files
- One volume for the mirrored log files

The MaxDB data and mirrored log volumes are separated from the primary log volumes and the saved log files in the other aggregate. Therefore it is always possible to recover the database without data loss if one of the two aggregates is lost.

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

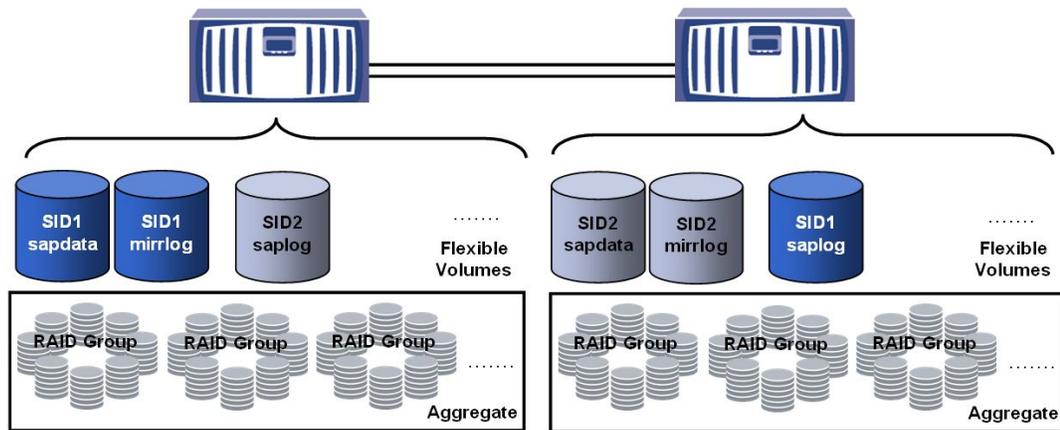


Figure 4) FlexVol volume layout.

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

Storage Controller1		Storage Controller2	
FlexVol sapdata	FlexVol mirrlogs	FlexVol saplog	
/sapdb/SID/sapdata	/sapdb/SID/sapmlog	/sapdb/SID/saplog	
		/sapdb/SID/saparch	
		/sapdb	
		/home/SIDadm	
		/usr/sap/trans	
		/sapmnt/SID	
		/usr/sap/SID	

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 MetroCluster and a normal cluster are the same.

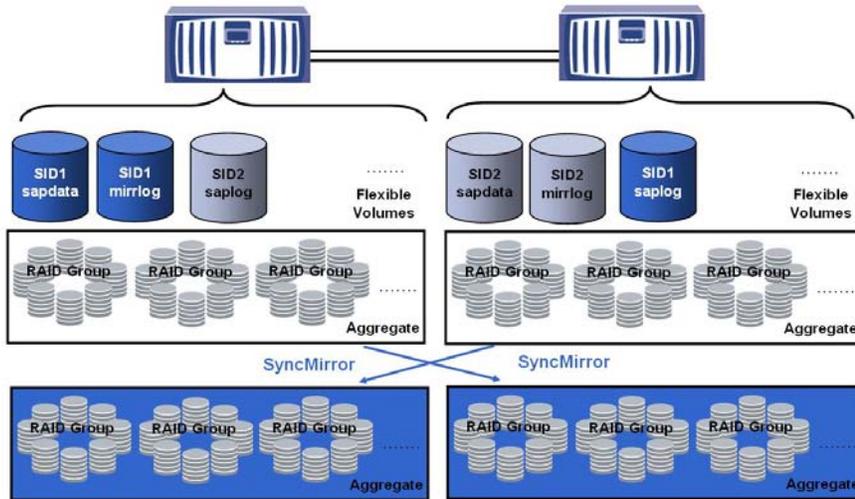


Figure 5) Storage layout with NetApp MetroCluster.

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.

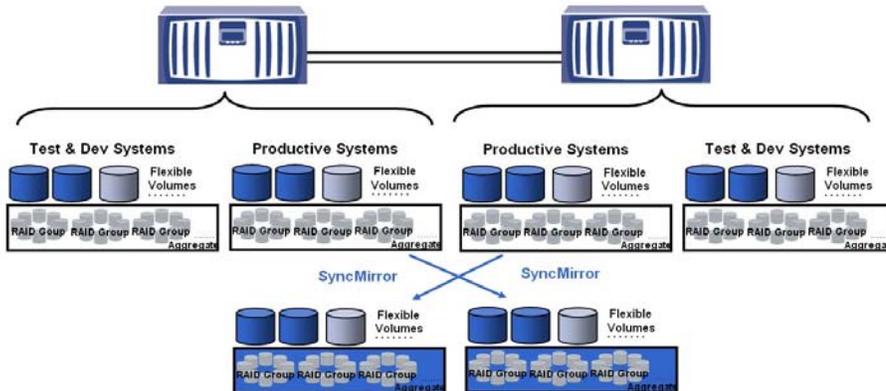


Figure 6) Storage layout with NetApp MetroCluster; only production systems are mirrored.

2.3 SIZING

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

- I/O requirements
- Capacity requirements
- Backup 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 MaxDB database analyzer or the SAP database performance monitor. For instance, iostat can be used if the measurement is done on the operating system level. Independent of which tools are used, it is very important that the measurement is done during peak loads on the SAP system. Especially when database tools are used for the measurement, a suitable time frame such as one hour must be chosen, 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 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 supporting both the I/O and capacity requirements.

2.4 INSTALLATION

This section describes the requirements and the configuration for installing an SAP Business Suite or SAP NetWeaver system with MaxDB on a UNIX server using the NFS protocol.

A common question that might arise in the SAP community would be “Is your solution SAP certified?” An important note is that there is no SAP certification program for storage. SAP supports UNIX with NFS as long as the database vendor supports or certifies the NFS storage system. SAP fully supports NetApp and NFS as a storage solution for the MaxDB database.

STORAGE NETWORK

A dedicated Gigabit Ethernet storage network needs to be used to attach the server(s) to the NetApp storage. This network should be used exclusively for the storage of traffic and not for any other services. Each server will therefore need a dedicated Gigabit Ethernet card to be connected to the NetApp storage.

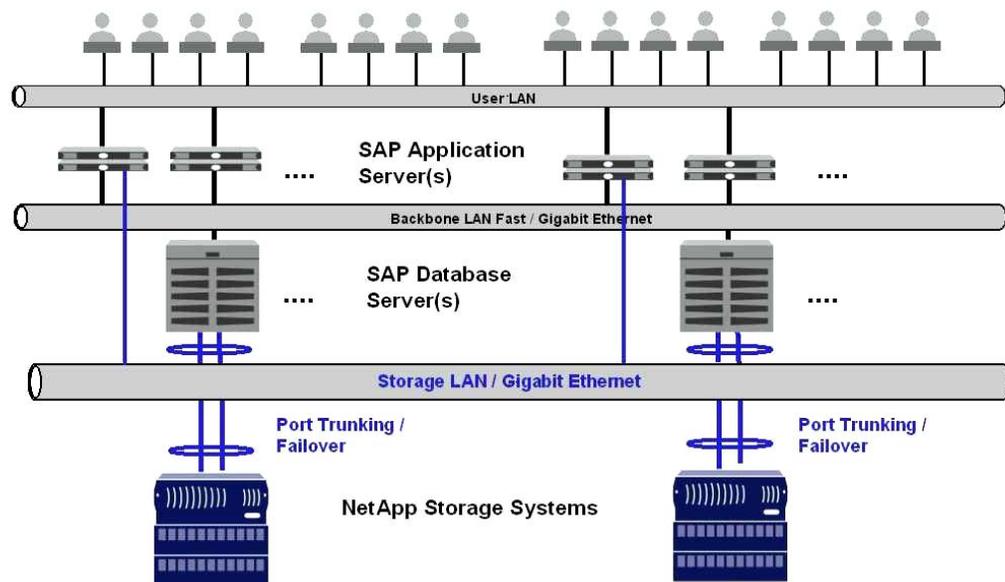


Figure 7) Infrastructure overview.

OPERATING SYSTEM CONFIGURATION

Correct NFS mount options are important to provide optimal performance and system stability. There are common mount options that are valid for all operating system platforms.

Common mount options: rw, bg, hard, nointr, rsize=32768, wsize=32768, vers=3, suid

Additional mount options for the different operating system platforms:

Linux®: <common>, tcp, timeo=600

Solaris™: <common>, proto=tcp, [forcdirectio or llock]

HP/UX: <common>, proto=tcp, timeo=600

AIX: <common>, proto=tcp, timeo=600, cio

More detailed information on mount options and NFS tuning can be found at [Mount options for databases on NetApp NFS](#).

Information on operating system specific tuning can be found at:

- Linux:
 - In the Linux 2.6 kernel, the NFS module introduces a new tunable parameter, sunrpc.tcp_slot_table_entries. This parameter controls the concurrent I/Os per mountpoint. The default value is 16, which is typically not sufficient. The maximum allowable setting of sunrpc.tcp_slot_table_entries is 128. This parameter should be set as described in the following technical reports:
 - [Linux \(RHEL 4\) 64-Bit Performance with NFS, iSCSI, and FCP](#)
 - [Using the Linux NFS Client with NetApp](#)
- Solaris:
 - [Database Performance with NAS](#)
- AIX:
 - [AIX Performance with NFS, iSCSI, and FCP](#)
- HP/UX:
 - [HP-UX 64-Bit NFS Performance with NetApp Storage](#)

NETAPP STORAGE CONTROLLER CONFIGURATION

During the SAP installation, the visibility of the Snapshot directory has to be switched off. Otherwise Sapinst will try to change permissions and ownership on Snapshot subdirectories. Since the Snapshot data is read only, Sapinst will fail, and the installation will abort. After the installation of the SAP system, the volume option can be switched on again.

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

Due to performance and security reasons the following options should be set:

```
filer> vol options vol_sapdata nvfail on
```

```
filer> vol options vol_sapdata no_atime_update on
```

FLEXVOL AND QTREE LAYOUT

For an SAP Business Suite or NetWeaver system installation on UNIX with MaxDB two or three FlexVol volumes need to be set up. Within these FlexVol volumes several qtrees will be configured.

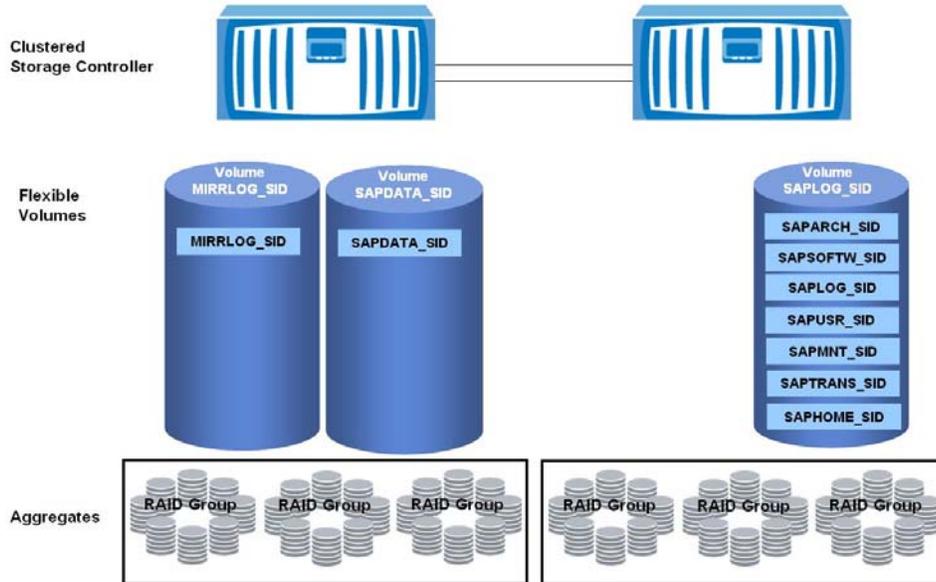


Figure 8) Installation with three FlexVol volumes in a storage system cluster.

SAP SYSTEM INSTALLATION

The following table shows the volume, qtree, and file system configuration with SID=MXD as an example.

Volume	Qtree	Mountpoint at SAP Server
sapdata_mxd	sapdata_mxd	/sapdb/MXD/sapdata
saplog_mxd	saplog_mxd	/sapdb/MXD/saplog
	saparch_mxd	/sapdb/MXD/saparch
	sapsoftw_mxd	/sapdb
	sapusr_mxd	/usr/sap/MXD
	sapmnt_mxd	/sapmnt/MXD
	saptrans_mxd	/usr/sap/trans
	saphome_mxd	/home/mxdadm
mirrog_mxd	mirrog_mxd	/sapdb/MXD/sapmlog

The necessary file systems for the SAP installation will be set up with the following steps:

1. Create directories /usr/sap/SID, /sapmnt/SID, /usr/sap/trans, /sapdb, and /home/sidadm.
2. Edit file system configuration file /etc/fstab to mount the corresponding file systems from the NetApp storage using the discussed mount options.
3. Mount the above file systems.
4. Create directory /sapdb/SID/sapdata and /sapdb/SID/saplog within the already mounted /sapdb file system.

5. Edit file system configuration file `/etc/fstab` to mount the `/sapdb/SID/sapdata` and `/sapdb/SID/saplog` file systems from the NetApp storage using the discussed mount options.

When these steps are finished, the following file system structure will be available at the SAP server:

```
n200:sqdmxd 72> df
```

Filesystem	Mounted on
<code>sapfiler1-gig:/vol/sapdata_mxd/sapdata_mxd</code>	<code>/sapdb/MXD/sapdata</code>
<code>sapfiler1-gig:/vol/saplog_mxd/saplog_mxd</code>	<code>/sapdb/MXD/saplog</code>
<code>sapfiler1-gig:/vol/saplog_mxd/saparch_mxd</code>	<code>/sapdb/MXD/saparch</code>
<code>sapfiler1-gig:/vol/saplog_mxd/sapsoftw_mxd</code>	<code>/sapdb/MXD</code>
<code>sapfiler1-gig:/vol/saplog_mxd/sapusr_mxd</code>	<code>/usr/sap/MXD</code>
<code>sapfiler1-gig:/vol/saplog_mxd/sapmnt_mxd</code>	<code>/sapmnt/MXD</code>
<code>sapfiler1-gig:/vol/saplog_mxd/saptrans_mxd</code>	<code>/usr/sap/trans</code>
<code>sapfiler1-gig:/vol/saplog_mxd/saphome_mxd</code>	<code>/home/mxdadm</code>

```
n200:sqdx 75> ls -al /sapdb/MXD
```

```
dr-xr-xr-x  11 sdb  sdba   264 2008-03-27 14:34 db
drwxrwxr-x   2 sdb  sdba   168 2008-03-13 20:05 saparch
drwxr-x---   2 sdb  sdba   272 2008-04-23 14:17 sapdata
drwxr-x---   2 sdb  sdba    80 2008-04-05 17:27 saplog
```

The SAP installation tool Sapinst fully supports NFS mounts. The SAP installation can therefore be accomplished as described in the corresponding SAP Installation Guide.

2.5 STORAGE MIGRATION

In the following section different storage migration approaches are discussed.

OVERVIEW OF MIGRATION APPROACHES

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

- Migration on operating system level
- Migration on database level

MIGRATION ON OPERATING SYSTEM LEVEL

In addition to the existing storage system, the NetApp storage system will be connected to the database server. The NetApp storage system will be configured, and the file systems will be mounted to the server. Before the data migration is started, the database and the SAP system must be shut down. The data will then be copied using the server from the old storage system to the NetApp system. When all data is copied, the old storage system will be disconnected from the database server. If the file system structure remains the same, the database can be started immediately. If there is a change in the file system structure, the new structure needs to be configured within MaxDB.

The disadvantage of this approach is that the SAP system will not be available while the database files are copied. Depending on the database size, the downtime could be several hours.

MIGRATION ON DATABASE LEVEL

A database backup will be restored to the target system with an initialization of the target database. The UNIX system can transport a database backup directly using remote shell without writing the backup to any further medium. Please see SAP note [489615, MaxDB or liveCache System Copies via Remote Pipe](#), for further details.

The database parameters can be restored from the database backup. The database name of the target system can be different to the name of the source system.

To minimize the impact on the source SAP system, the restore can be done using a separate server connected to the NetApp storage. In addition, the saved log files will be continuously copied to the separate server. Before the final migration is started, the SAP database and the SAP system must be shut down. The NetApp storage will then be connected to the database server, and the file systems of the application server will be mounted to the server. The log area and the saved log files that have not been copied yet will now be copied from the old storage system to the NetApp storage. When all data is copied, the old storage system will be disconnected from the database server. Finally, a forward recovery of the database will be carried out.

This approach will reduce the downtime during the migration but will need an additional server during the migration process.

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 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 restore the data to the three target systems.

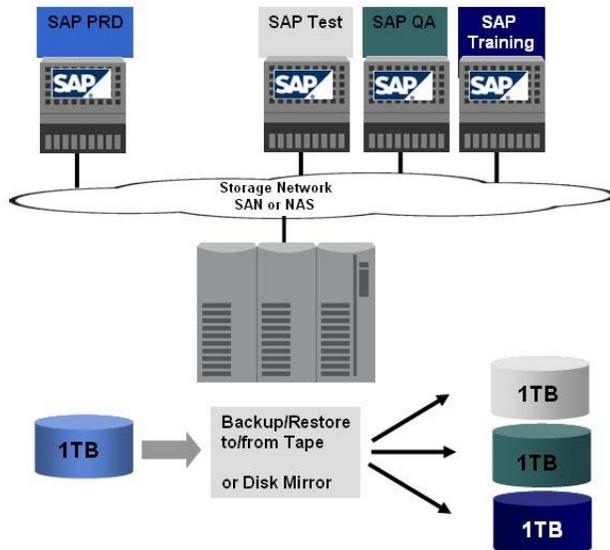


Figure 9) SAP system copy: standard approach.

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. 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 postprocessing tasks specific to SAP need to be performed.

All the above steps are fully automated and do not need any manual interaction. AN SAP system copy can be made in a few minutes using the NetApp solution. Figure 10 shows a high-level graphical representation of this infrastructure.

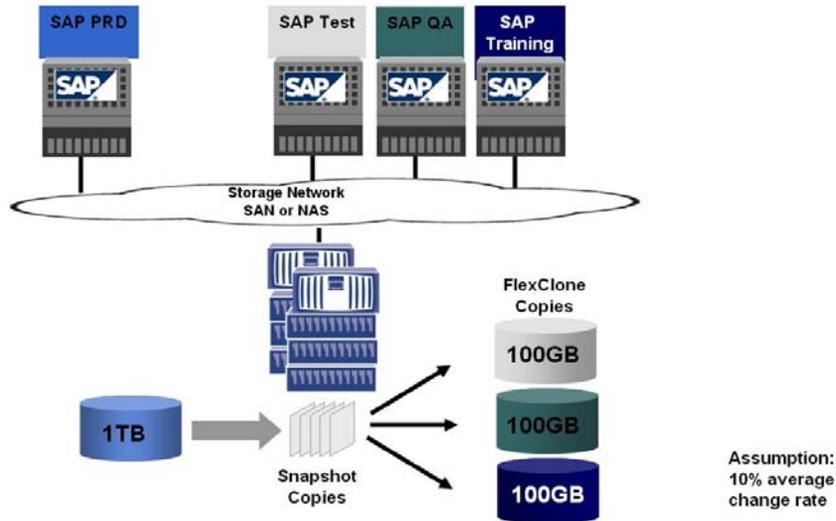


Figure 10) SAP system copy: NetApp approach.

TIME REQUIREMENTS

The time required to create an 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 postprocessing

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 11 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.

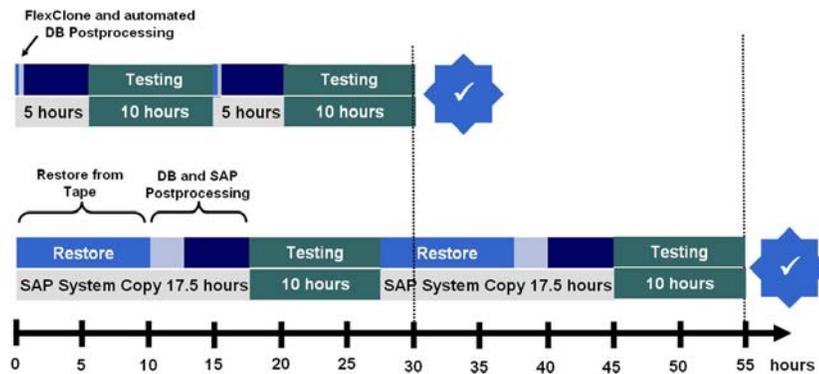


Figure 11) SAP system copy: time requirements.

MAXDB PROCEDURE

The MaxDB architecture allows an easy procedure to create system copies using the NetApp FlexClone technology. No log files need to be applied. The procedure would be:

- 1) The database and SAP software needs to be installed once on the target system. The number and sizes of the database volumes in the target system are the same as those in the source system. This step is not needed for a refresh of the target system with the content of the source system.
- 2) Shut down the target system and unmount the file system containing the data volumes (for example, /sapdb/SID/sapdata).
- 3) Create a Snapshot copy of the NetApp volume containing the database data volumes of the source system.
- 4) Create a FlexClone copy based on the Snapshot copy created before and mount the new FlexVol volume at the target database server.
- 5) Start the target database into the ADMIN mode and execute the dbmcli command "db_execute clear log". Use the "clear log" command only for database copies. Never use this command for crash recovery.
- 6) Start the database into ONLINE mode and execute the dbmcli command "load_systab".
- 7) Rename the database user SAP<SourceSID> and SAP<SourceSID>DB to SAP<TargetSID> and SAP<TargetSID>DB using the sqlcli "rename user" command.
- 8) Do SAP postprocessing tasks.

Saved log files can be recovered after the "clear log". Nevertheless, if the database starts without recovered logs, then it goes online with a transaction-consistent point in time of the last database save point before the creation of the Snapshot copy and FlexClone volume. The SQL command "FORCE SAVEPOINT" lets the source database flush the modified data blocks from the cache to disk with a new save point.

More detailed information about this procedure can be found in [SAP note 371247](#).

CONCLUSION

A key requirement for successfully managing an SAP environment is the ability to create copies of production data for use in testing, development, training, and sandbox systems. NetApp Snapshot and FlexClone technologies allow the fast and efficient creation of SAP systems with MaxDB. The NetApp system copy solution significantly improves the process for creating SAP system copies by reducing the impact on the production system during the backup process. It also allows a complete system copy to be created in a few minutes with very little storage required.

3.2 SAP UPGRADES

An SAP upgrade is a typical example of a complex customer project within an 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 an SAP upgrade project at regular intervals.

Customers face several challenges in an 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 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.
- **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 an 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 an 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.

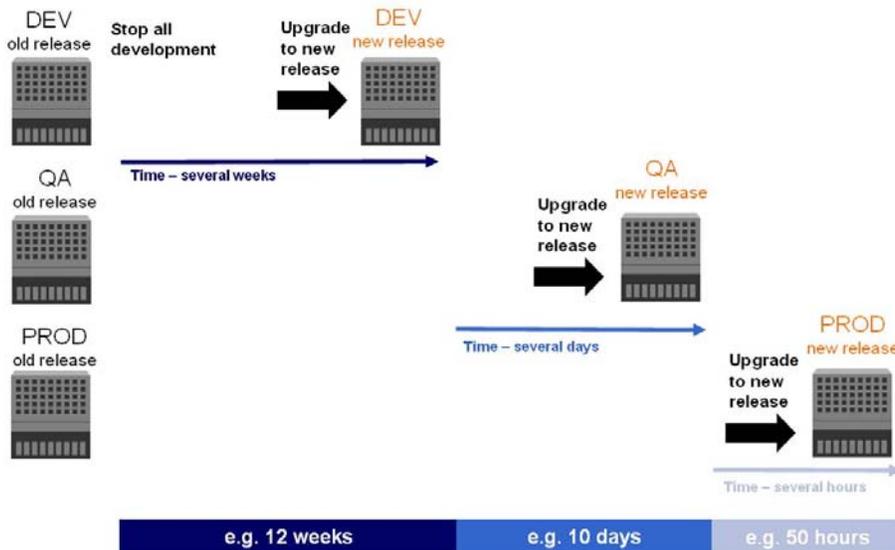


Figure 12) SAP upgrade: overview.

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.

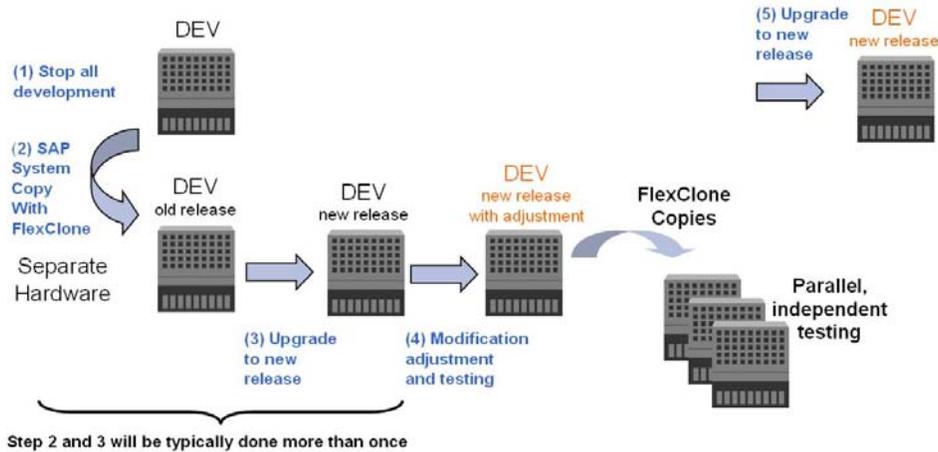


Figure 13) SAP upgrade: development system.

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.

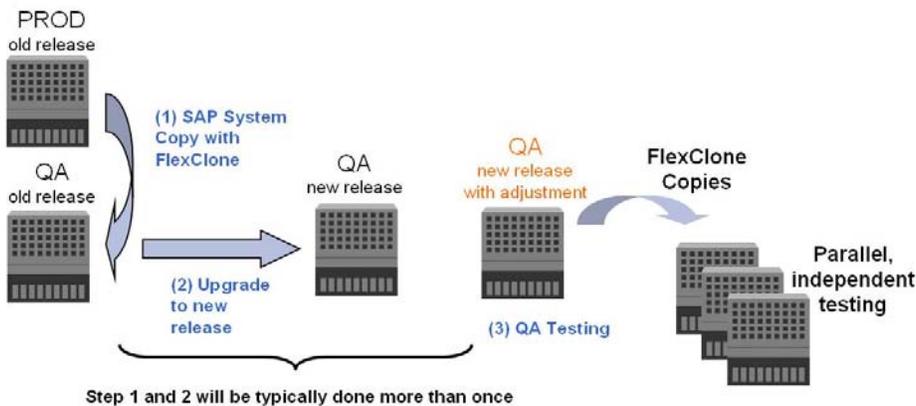


Figure 14) SAP upgrade: quality assurance system.

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.

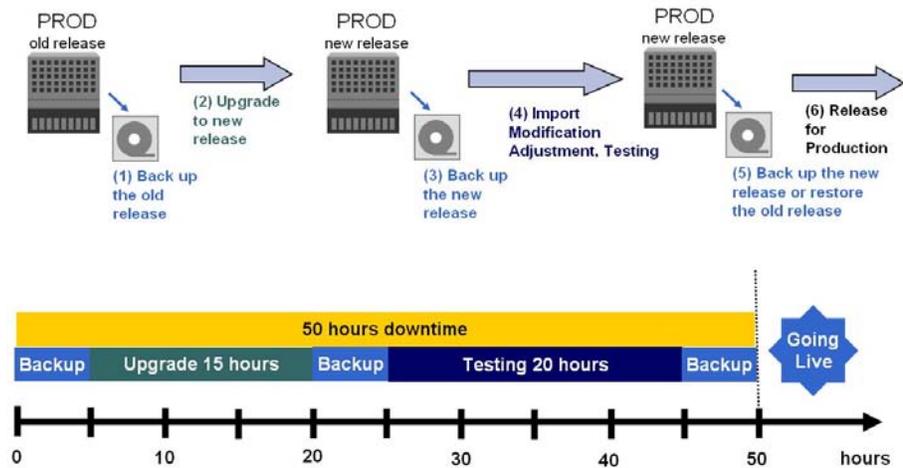


Figure 15) SAP upgrade: production system.

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.

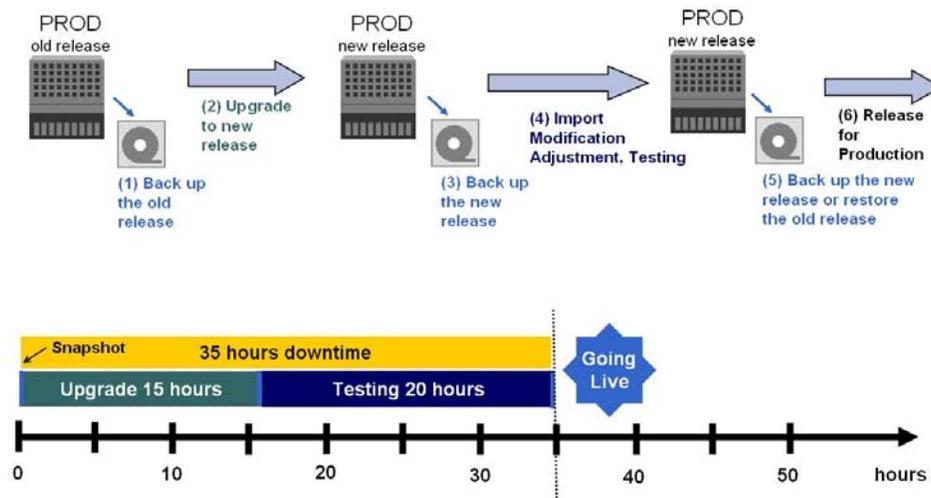


Figure 16) SAP upgrade: production system with NetApp.

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.

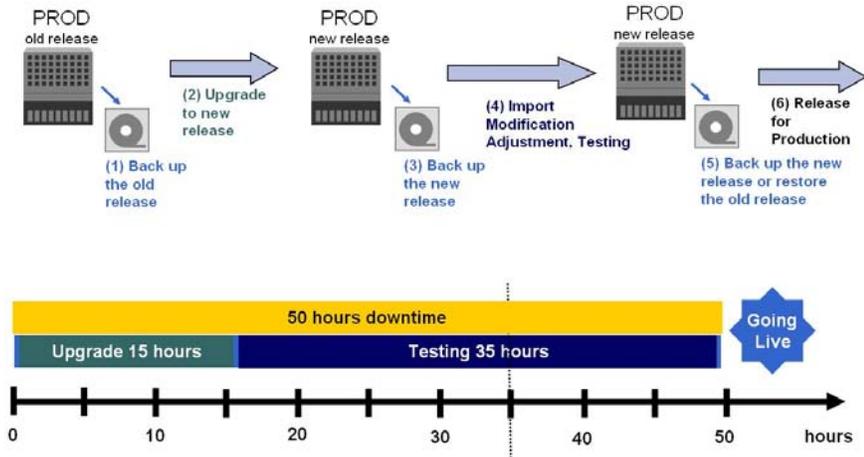


Figure 17) SAP upgrade: production system with NetApp.

4 BUSINESS CONTINUANCE

4.1 BACKUP AND RECOVERY

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 and 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 an 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 an SAP upgrade always includes at least three backups of the SAP database. The time needed to perform these backups cuts down the total available time for the upgrade process. 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 Snapshot technology can create an online or offline database backup in seconds. The time needed to create a Snapshot copy is independent of the size of the database, because a Snapshot does not move any data blocks. The use of Snapshot technology has no performance impact on the production SAP system because the NetApp Snapshot implementation does not copy data blocks when the Snapshot copy is created or when data in the active file system is changed. Therefore, creation of Snapshot copies can be scheduled without having to consider peak dialog or batch activity periods. SAP and NetApp customers typically schedule several online Snapshot backups during the day, for instance, every four hours.

Snapshot copies also provide key advantages for the restore and recovery operation. The NetApp SnapRestore functionality allows restoration of the entire database or parts of the database to the point in time when any available Snapshot copy was created. This restore process is done in a few minutes, independent of the size of the database. Because several online Snapshot backups were created during the day, the time needed for the recovery process is also dramatically reduced. Because a restore can be done using a Snapshot copy that is at most four hours old, fewer transaction logs need to be applied. The mean

time to recover, which is the time needed for restoration 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 restoration and recovery. Most restore and recovery actions are handled by using SnapRestore. Restores from a secondary location 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 two-week-old backup.

A backup and recovery solution using a NetApp storage system always consists of two parts:

1. Backup and restoration/recovery using Snapshot and SnapRestore
2. Backup and restoration to/from a secondary location

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

- **Disk-to-disk backup using a NetApp near-line or primary storage system 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 has to be transferred from the source to the destination, all following backups copy only the changed blocks to the secondary storage. Therefore, the load on the primary storage system and the time needed for a full backup are significantly reduced. Because SnapVault stores only the changed blocks at the destination, a full database backup requires significantly less disk space.
- **Backup to tape using third-party backup software such as NDMP backup (serverless backup).** The tape is connected directly to the primary storage system. The data is written to tape using NDMP.

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

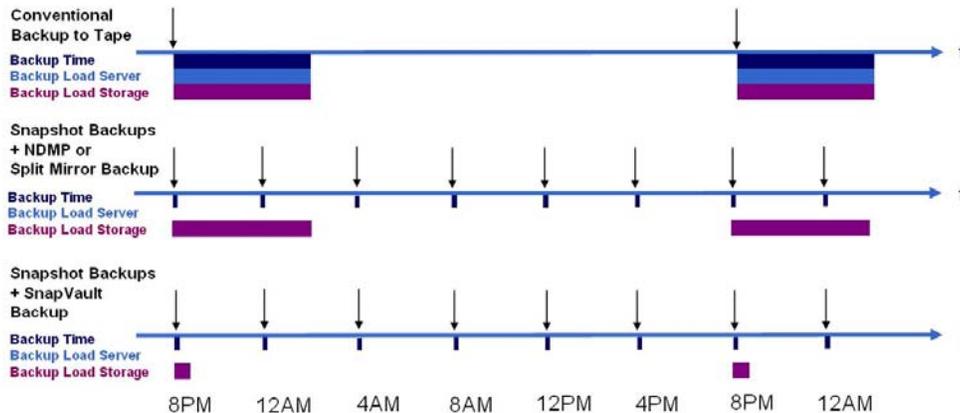


Figure 18) Comparison of time required for different backup methods.

SNAPSHOT BACKUPS TOGETHER WITH NDMP OR SPLIT MIRROR BACKUPS

Snapshot backups do not generate any load on the database server or the primary storage system. The Snapshot generation doesn't need any synchronization with the online MaxDB database. A full database backup based on Snapshot 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 restoration process and reduces the number of logs that must be applied during forward recovery. In

addition, a full NDMP backup to tape or a split mirror backup to tape is scheduled once a day. This backup still creates a heavy load on the primary storage system and takes the same amount of time as conventional tape backup.

SNAPSHOT BACKUPS TOGETHER WITH DISK-TO-DISK BACKUP AND SNAPVAULT

Snapshot backups are used here 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 small compared to full tape backups.

Figure 19 compares the time required to perform restoration and recovery.

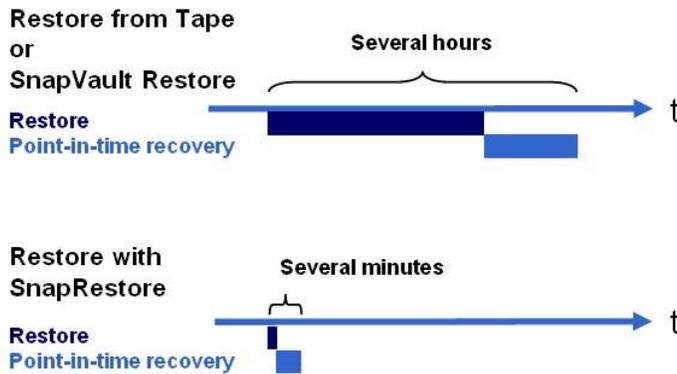


Figure 19) Comparison of time needed for restoration and recovery using NetApp solutions.

RESTORE FROM TAPE OR SNAPVAULT RESTORE

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

RESTORE WITH SNAPRESTORE

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

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

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 production SAP system
- Dramatically reduced mean time to recovery
- Minimum disk space needed for database backups at the primary and the secondary storage systems

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.

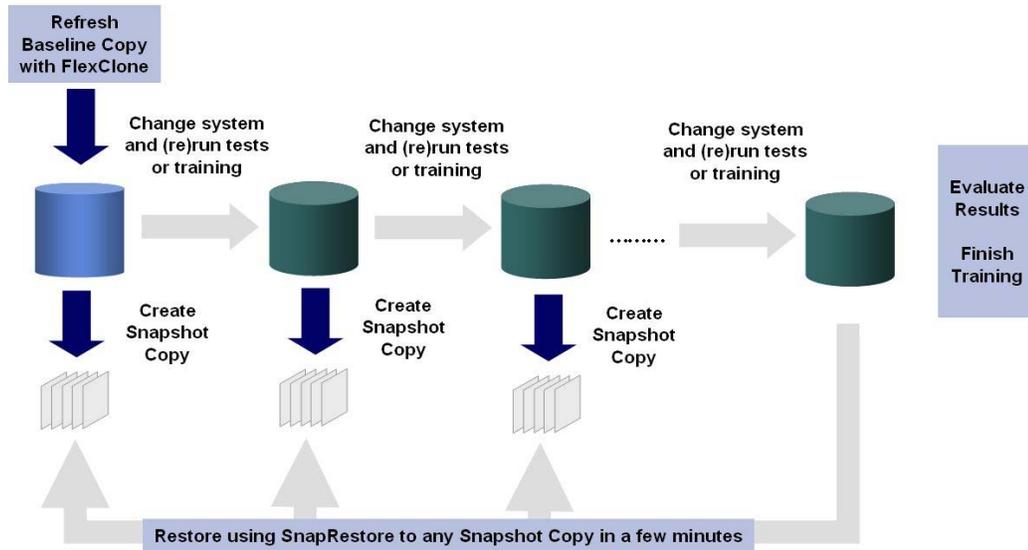


Figure 20) SAP testing cycle.

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.

4.2 HIGH AVAILABILITY

BUSINESS CHALLENGES

Productive SAP systems are business-critical applications that require 24x7 availability. Meeting this requirement requires an infrastructure without 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 assures 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 using a cluster interconnect. This connection is redundant and is used to exchange cluster heartbeats and synchronize the NVRAM on both storage controllers. The disk shelves of the cluster partner are connected to the second storage controller using a second Fibre Channel loop. If the first storage controller fails, the second storage controller handles its disk shelves. The MAC and IP addresses and the WWPN of the first storage controller are also adopted. Since the NVRAM is mirrored on both storage controllers using 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.

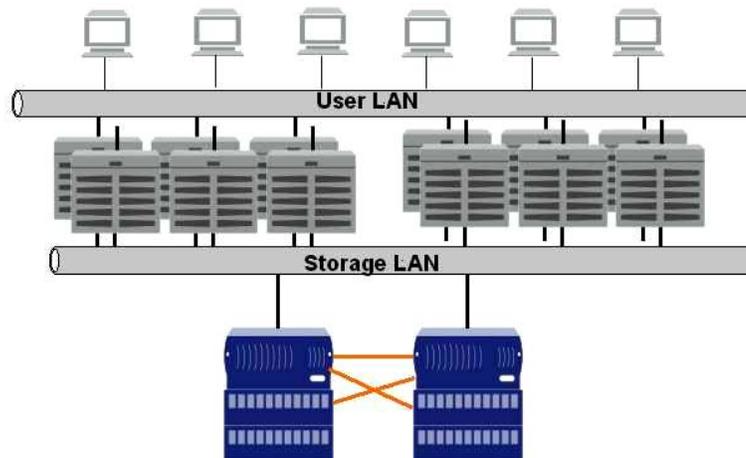


Figure 21) NetApp clustered storage system solution.

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

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

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 costs of not having one—lost productivity, revenue, and customer loyalty and possibly even business failure—make it mandatory to have a plan that assures an absolute minimum of downtime and rapid recovery from a disaster, with no or minimal 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 your selection of a disaster recovery solution that utilizes one or many NetApp products.

SNAPMIRROR

NetApp SnapMirror® software 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 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 Level 0 transfer to initialize the disaster recovery site. After the initial transfer, incremental changes are then 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. Selective Snapshot backups will be mirrored to the disaster recovery site. Additionally, the qtree where the saved log files are stored has to be mirrored using SnapMirror. NetApp recommends a frequent SnapMirror update of the log files, for example, every 10 minutes, to assure a minimum of data loss.

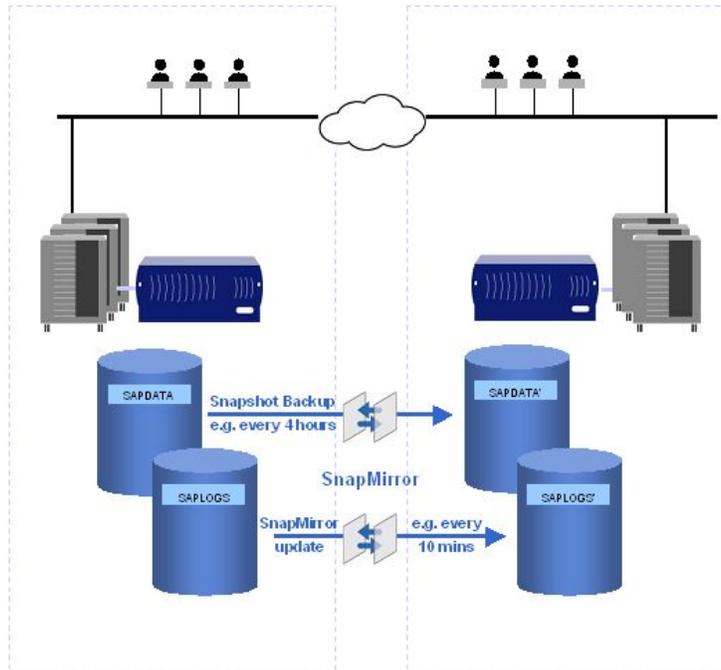


Figure 22) Disaster recovery with SnapMirror.

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 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.

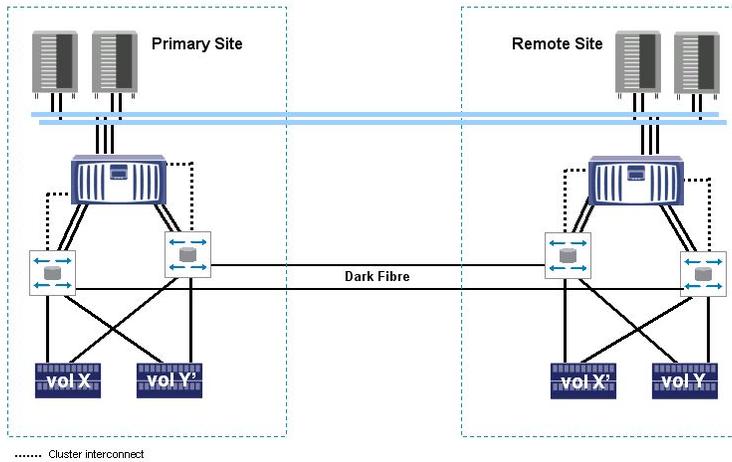


Figure 23) MetroCluster over direct Fibre Channel switch connection.

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

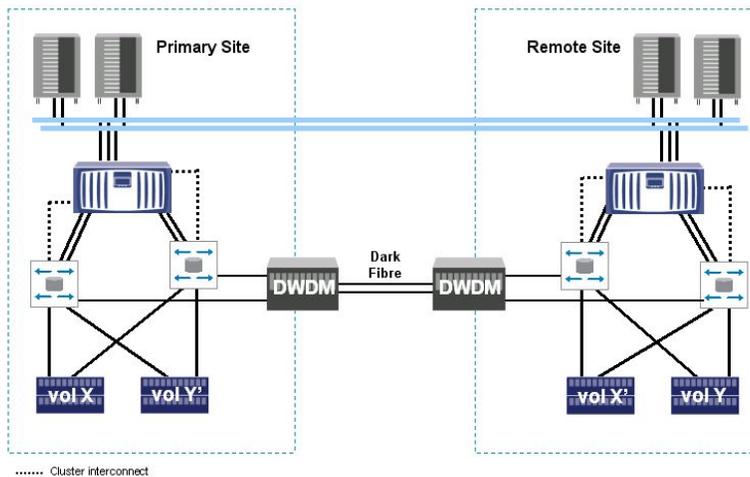


Figure 24) MetroCluster over Fibre Channel and DWDM switch infrastructure.

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 maintain performance and availability while reducing management overhead.

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. Current write once, read many (WORM) technologies such as 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

Besides system size and performance issues, SAP customers have to keep up with increasingly numerous industry regulations that introduced financial penalties for failing to comply with retention, indexing, auditing, privacy, and reporting requirements. These regulations affect 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 purchasing new storage hardware and software.

Historically, most regulated data has been stored on optical devices, tape, paper, and/or microfiche/microfilm. According to the Enterprise Storage Group (ESG), about 10% of regulated data is stored on disk today. Disk has not often been utilized because of a number of factors, which include cost and the lack of necessity to retrieve information quickly. However, ESG estimates that, moving forward, disk will be the fastest-growing area in the storage of regulated data.

NETAPP SOLUTION

SAP data archiving is based on the archive development kit (ADK), in which archiving objects are used to remove data that is no longer needed for business processes in online databases and to store it in such a way that it is still accessible. The purpose of XML-based data archiving is the same as that of ADK-based archiving. The key differences are universally accepted and widely used standards: the XML format is used to archive business objects, HTTP as a communication service, and Web-Based Distributed Authoring and Versioning (WebDAV) as a general protocol to connect storage systems.

The ADK is the software layer that encapsulates the technical aspects of data archiving programs. The ADK provides an application programming interface used by SAP internally and also used by customers and partners 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 using workflow.

WebDAV is a set of extensions to HTTP 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 using WebDAV as the data archive interface is certified by SAP.

Once archive files have been created, the data marked for archiving 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 near-line storage 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 archiving storage subsystem for SAP landscapes. The near-line storage leverages NetApp Data ONTAP 7G technology and takes full advantage of value-added software from NetApp such as SnapLock. NetApp RAID-DP technology enables near-line storage

systems to tolerate single disk failures with no data loss. If additional capacity or performance is required for any reason, NetApp fabric-attached storage systems can be substituted for near-line storage in the SAP storage landscape.

SnapLock is the NetApp implementation of high-performance disk-based magnetic WORM storage. SnapLock provides secure, storage-enforced data retention functionality using open file protocols such as CIFS and NFS while leveraging existing NetApp technologies to the greatest degree possible. This implementation also includes significant efforts to harden Data ONTAP, and its administration interfaces used by SnapLock are hardened and certified to protect data in stringently regulated environments. Even the storage administrator is not considered to be a trusted 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 flexibly regulated environments.

SnapLock provides special-purpose volumes in which files can be stored and committed to a nonerasable, nonrewritable state, either forever or for a designated retention period. SnapLock allows this retention to be performed at the granular level of individual files through standard open file protocols such as CIFS and NFS. The retention period of these files is enforced by Data ONTAP, which acts as a gatekeeper. All access to physical media is strictly controlled.

SnapLock is based on 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 a 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 storage system.

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

Conclusion: NetApp provides a flexible, scalable, and secure solution for SAP compliance and data archiving:

- 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.
- Data can be managed and backed up using the customer's current products and strategies.
- The solution can incorporate existing storage from NetApp or other vendors.
- Better ROI and lower TCO are achieved through increased availability, enhanced system performance, lower administration overhead, and increased staff productivity.
- Compliance and archived data can be easily accessed 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 becomes increasingly complex. The NetApp solutions for SAP bring together technologies that simplify and accelerate this process and align with the SAP application lifecycle.

The NetApp solutions for SAP 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.

