

SAP ON UNIX® AND DB2® WITH NFS AND NETWORK APPLIANCE™ STORAGE

BEST PRACTICES

SAP Competence Center, Network Appliance, Inc.
January 2008 | TR-3485

ABSTRACT

This document provides customers and partners with the best practices for deploying Network Appliance storage in support of mySAP™ Business Suite solutions running in a UNIX and NFS environment using a DB2 database.

TABLE OF CONTENTS

INTRODUCTION 3

STORAGE PROVISIONING AND MANAGEMENT 5

 CONSOLIDATION 5

 STORAGE LAYOUT 6

 SIZING 10

 INSTALLATION 11

 STORAGE MIGRATION 16

SYSTEM MANAGEMENT AND MAINTENANCE 17

 SAP SYSTEM CLONING 17

 SAP UPGRADE 21

BUSINESS CONTINUANCE 25

 BACKUP & RECOVERY 25

 HIGH AVAILABILITY 31

 DISASTER RECOVERY 32

ARCHIVING AND COMPLIANCE 35

SAP ADAPTIVE COMPUTING 38

ADDITIONAL REFERENCES 40

INTRODUCTION

Scope

This document is intended to provide customers and partners with the best practices for deploying Network Appliance storage systems in support of mySAP Business Suite solutions running in a UNIX and NFS environment using a DB2 database. Primary consideration has been 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 will be generic and specific to neither a specific SAP application nor size and scope of SAP implementation. This guide assumes a basic understanding of the technology and operation of NetApp and SAP products and was developed based upon the interaction of technical staff from Network Appliance, SAP, DB2, and our customers.

Business Challenges

Corporations deploying SAP solutions today are under great pressure to reduce TCO, accelerate ROI, and increase 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. 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 a requirement for corporations to keep pace with change.

IT Challenges

A typical production SAP landscape today 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 non-production 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 typical to have multiple development instances as well as more than one system used for quality assurance to be used for additional 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 each of these instances is going to have differing performance, scalability, availability, and uptime profiles. These profiles can also fluctuate depending on the phases of a project implementation and whether the project is focused on an existing SAP implementation or a brand-new one.

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

If the IT infrastructure supporting SAP applications is inflexible, difficult to operate or manage, or has high cost of ownership, barriers develop within IT that can negatively impact the ability of business owners to deploy new and improved business processes.

Network Appliance 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 ensures 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.

STORAGE PROVISIONING AND MANAGEMENT CONSOLIDATION

In today's rapidly changing business climate, enterprises demand cost-effective, flexible data management solutions that can handle the unpredictable and explosive growth of storage in heterogeneous environments. To enable global data management, ensure business continuity, satisfy regulatory and compliance standards, and improve resource utilization, a flexible and scalable storage network solution is required. The solution must also minimize complexity and reduce the total cost of ownership (TCO).

NetApp offers highly available, scalable, and cost-effective storage consolidation solutions that incorporate the NetApp unified storage platform and the feature-rich functionality of data and resource management software to deliver storage that improves enterprise productivity, performance, and profitability while providing investment protection and enhanced asset utilization. NetApp enterprise-class storage solutions are proven interoperable across all platforms. NetApp fabric attached storage (FAS) systems integrate easily into complex enterprises and simultaneously support NAS, Fibre Channel SAN, and IP SAN (iSCSI).

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 it is needed without disruption and at the lowest incremental cost.




Figure 1) FlexVol technology.

NetApp FlexClone™ technology enables true cloning—instant copying of datasets 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. NetApp FlexClone provides substantial space savings with minimal overhead. This means many more dataset variations can be managed—in less time and with less risk—to address and fuel the organization's business and development objectives.

STORAGE LAYOUT

In a cluster configuration the layout should be based on the most efficient usage of resources. Since both storage controllers are active parts of the cluster, the SAP systems should be distributed to both storage controllers in order to use the available CPU power and disk spindles equally.

It is therefore recommended to distribute the data of each SAP system among each single aggregate on each storage controller. In this case the DB2 log files and the archives will be stored on the aggregate on the second storage controllers while the database containers and the mirrored log files are stored in the aggregate on the first storage controller. Storing the database containers in a separate FlexVol volume is important to allow usage of Snapshot™ copies, SnapRestore®, FlexClone, and other Data ONTAP® features that work on the volume level.




Figure 2) Storage layout with NetApp cluster.

NetApp recommends using large aggregates to store all data of all SAP systems. These aggregates should be configured with RAID-DP™. Using large aggregates will provide the performance benefits of all available disk spindles in the aggregate to every FlexVol volume in the aggregate. In a cluster configuration each SAP system will use three FlexVol volumes. One FlexVol volume will be exclusively used for the DB2 database containers, a second FlexVol volume for the mirrored logs, and a third FlexVol volume will hold all the other SAP and DB2 file systems. Using this layout it is always possible to recover the database without data loss if one of the two aggregates is lost. Because DB2 log mirroring is not used by default, it has to be configured.

FlexVol saplog	FlexVol sapdata	FlexVol mirrlog
/usr/sap/<SID>	/db2/<SID>/sapdata1	/db2/<SID>/mirror_log
/usr/sap/trans	/db2/<SID>/sapdata2	
/sapmnt/<SID>		
/home/<SID>adm	/db2/<SID>/sapdataN	
/db2/<SID>		
/db2/<SID>/log_dir		
/db2/<SID>/log_archiv		
/db2/<SID>/log_retriev		
/db2/<SID>/saptemp1		
/db2/<SID>/db2dump		

The storage layout does not change if only one storage controller is used with two aggregates and three FlexVol volumes.




Figure 3) Storage layout with three FlexVol volumes.

Because RAID-DP offers a high level of data protection, one aggregate can be used to store SAP systems on one NetApp storage controller. The reliability of RAID-DP is excellent; only if three disks within the same RAID group fail at the same time frame will data loss occur. More information on RAID-DP can be found at [NetApp Data Protection: Double parity RAID for enhanced data protection with RAID-DP](#).

One FlexVol volume will be exclusively used for the DB2 database containers. The other FlexVol volume will hold all the other SAP and DB2 file systems. The advantage of this solution is that all available disk spindles are usable for each SAP system, therefore providing more efficient space and spindle utilization. The usage of mirrored logs is optional for this approach.

FlexVol saplog	FlexVol sapdata
/usr/sap/<SID>	/db2/<SID>/sapdata1
/usr/sap/trans	/db2/<SID>/sapdata2
/sapmnt/<SID>	
/home/<SID>adm	/db2/<SID>/sapdataN
/db2/<SID>	
/db2/<SID>/log_dir	
/db2/<SID>/log_archiv	
/db2/<SID>/log_retriev	
/db2/<SID>/saptemp1	
/db2/<SID>/db2dump	
/db2/<SID>/mirror_log	




Figure 4) Storage layout with two FlexVol volumes.

The high availability / disaster recovery solution using MetroCluster with synchronous mirroring works on the aggregate level. If all SAP systems require being mirrored synchronously, the layout for a MetroCluster and a normal cluster is the same.




Figure 5) Storage layout with NetApp MetroCluster.

Additional aggregates are necessary if only parts of the landscape require synchronous mirroring. For instance only production SAP systems have the requirement of synchronous mirroring, but the test and development systems do not need to be mirrored.




Figure 6) Storage layout with NetApp MetroCluster and aggregates.

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 to perform a storage sizing and how these requirements will influence the storage landscape.

Storage sizing for a SAP landscape is based on several conditions defined by customer requirements. All these requirements together will 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 since overall SAP system performance will be directly affected.

For existing SAP systems, the I/O requirements need to be measured using database or operating system tools. For instance, the SAP database performance monitor can be used to obtain this data. A tool like iostat can be used if the measurement is done at the operating system level. Independent of which tools are used it is very important that the measurement be done during peak loads of the SAP system. Especially when database tools are used for the measurement, a suitable time frame must be chosen, e.g., one hour, since these tools calculate an average value and the I/O sizing must be based on peak values.

For new SAP systems in which 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. The storage sizing will be of course 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 above I/O requirements. Also, the backup load must be added if the backup happens during a high activity phase of the system.

Based on the I/O requirements, the type and amount of disk spindles and storage controllers will be determined.

In order to determine the required storage capacity, the following information must be available.

- Size of each database
- Growth rate
- Number and retention policy of Snapshot copies
- Number and duration of use of FlexClone volumes
- Synchronous or asynchronous mirroring

Based on the I/O requirements, the type and amount of disks and the storage controller supporting the capacity will be determined.

The results of the I/O sizing and the capacity sizing will be compared at the last step to define the right storage system supporting both the I/O and capacity requirements.

INSTALLATION

This section describes the requirements and the configuration for installing a mySAP Business Suite or SAP NetWeaver system with a DB2 database 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?” It is important to note 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.

IBM does certify storage vendors for use with DB2. Network Appliance storage is certified for use with DB2. This is a complete list of [verified network-attached storage systems for use with DB2 Universal Database Version 8](#).

GENERAL REQUIREMENTS

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 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.

Linux®: rw,bg,hard,nointr,rsize=32768,wsiz=32768,tcp,vers=3,suid,timeo=600

Solaris™: rw,bg,hard,nointr,rsize=32768,wsiz=32768,proto=tcp,vers=3,suid,[forcdirectio or llock]

AIX®, HP-UX® rw,bg,hard,nointr,rsize=32768,wsiz=32768,proto=tcp,vers=3,suid

More detailed information on mount options and NFS tuning can be found at:

- [Using the Linux NFS client with NetApp](#)
- [IBM DB2 UDB Enterprise Server Edition V8 for Unix: Integrating with a NetApp Storage System](#)
- [Database Layout with Data ONTAP 7G](#)

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
```

Snapshot backups for database applications won't be consistent from the database point of view without shutting down the database or putting the DB2 database in write suspend mode. Therefore automatically scheduled Snapshot copies on storage level should be turned off on database volumes.

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

- More detailed information on the volume options can be found at [IBM DB2 UDB Enterprise Server Edition V8 for Unix: Integrating with a NetApp Storage System](#).

FLEXVOL AND QTREE LAYOUT

For a mySAP Business Suite or NetWeaver system installation on UNIX and DB2, three FlexVol volumes within two aggregates need to be set up. Within these FlexVol volumes several qtrees will be configured.

- Qtree sapdata_<sid> within the data volume
- Qtree mirrlog_<sid> within the mirrlog volume
- Qtrees saplog_<sid>, sapusr_<sid>, sapmnt_<sid>, saptrans_<sid> and saphome_<sid> within the log volume

The data volume and the mirrlog volumes are stored within the same aggregate and the log volume is stored in a second aggregate.




Figure 8) Installation with two FlexVol volumes.

If a clustered storage system is used, the configuration is identical to the above approach. One aggregate will be configured on each storage controller and the qtrees will be distributed as described above.




Figure 9) Installation with three FlexVol volumes in a storage system cluster.

If one aggregate at one storage controller should be used to ensure the most effective usage of disk spindles, only two FlexVol volumes and the required qtrees have to be created. The third FlexVol for the mirrlogs is not needed.




Figure 10) Installation with three FlexVol volumes.

SAP SYSTEM INSTALLATION

The following table shows the volume, qtree, and file system configuration with SID=R47 as an example. The additional mirrlog volume is only necessary if the decision is made to use two separate aggregates within a single storage controller or in a storage system cluster configuration.

Volume	Qtree	Mountpoint at SAP Server	Symbolic Links
sapdata_r47	sapdata_r47	/db2/R47_sapdata	ln -s /db2/R47_sapdata/sapdata1 /db2/R47/sapdata1 ln -s /db2/R47_sapdata/sapdata2 /db2/R47/sapdata2 ln -s /db2/R47_sapdata/sapdata4 /db2/R47/sapdata4
saplog_r47	saplog_r47	/db2	
saplog_r47	sapusr_r47	/usr/sap/R47	
saplog_r47	sapmnt_r47	/sapmnt/R47	
saplog_r47	saptrans_r47	/usr/sap/trans	
saplog_r47	saphome_r47	/home_r47adm	ln -s /home_r47adm /home/r47adm
mirrlog_r47	mirrlog_r47	/db2/R47/mirrlog_log	ln -s /db2/mirrlog_r47/mirrlog_r47 /db2/R47/mirrlog_log

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

- Create directories /usr/sap/<SID>, /sapmnt/<SID>, /usr/sap/trans, /db2, /home_<sid>adm
- Edit file system configuration file, i.e., /etc/fstab for Solaris or Linux or /etc/filesystems for AIX to mount the corresponding file systems from the NetApp storage using the discussed mount options
- Mount the above file systems
- Create directory /db2/<SID>_sapdata within the already mounted /db2 file system
- Create the symbolic links for the sapdata and the home directory as described in the above table

Installation with sapinst

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

STORAGE MIGRATION

In the following section, different storage migration approaches are discussed. If a migration includes a change of the operating system or database system, the migration cannot be done at the storage level. In this case, the SAP migration tools are required. These tools export the data from the source environment and import the data into the target environment.

OVERVIEW OF MIGRATION APPROACHES

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

- Migration at the operating system level
- Migration at the database level

Migration at the 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 on the server. Before the data migration is started, the database and the SAP system must be shut down. The data will then be copied via 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, for example, different mount points, the database can be updated using the db2 relocatedb utility to reflect the new path of database tablespace containers.

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 at the database level

A redirected restore of an online or offline database backup will be performed to the NetApp storage 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 archive logs 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 will be mounted to the server. The log files still on the old storage system will be copied to the NetApp storage. 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. Finally, a roll 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. Another approach is to use an offline backup for the redirected restore. However, this approach requires that SAP and DB2 be shut down during both the backup and restore operations. This amount of downtime is often difficult to schedule in production environments.

SYSTEM MANAGEMENT AND MAINTENANCE

SAP SYSTEM CLONING

BUSINESS CHALLENGES

A typical SAP customer environment today consists of multiple mySAP Business Suite and SAP NetWeaver components. In order to be able to test application patches, run performance and data integrity tests, or provide user training environments, one or more copies of SAP components is required. A typical SAP customer on average needs about 10 copies of different SAP components. These copies need to be refreshed often, on a weekly or monthly basis.

The creation of a SAP system copy normally takes several days and negatively impacts the production environment. In addition, a lot of manual steps are performed that consume valuable IT staff time.

The source database needs to be exported using SAP tools and imported at the target system or an offline backup of the source database will be restored at the target system. Depending on the database size, these steps have a significant impact on the application availability. It takes many hours to replicate a 1TB database from the source to the target system. Preparing the cloned system so that it can be used in the new environment takes several additional hours. This preparation is often done manually, consuming the SAP Basis administrator's time.

Being able to create a SAP system copy on demand very quickly becomes more and more important.

- Quality insurance systems need to be refreshed on a weekly basis
- Additional test systems need to be set up quickly to perform specific integration tests
- Quick setup of test system with current production data during a SAP upgrade project
- Setup or resync of training systems

The traditional approach to creating system copies is not suitable to address these demands.

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

NETAPP SOLUTION

The NetApp solution for SAP system cloning addresses these issues by providing a fully automated process to create a SAP system copy on demand, in a few minutes without any impact to the source production system. In addition, NetApp cloning functionality allows storage to be managed very efficiently by storing only data changes between the parent and the clone.

NETAPP SOLUTION FOR SAP SYSTEM COPIES

SAP system copies are accomplished using the NetApp FlexClone feature. A FlexClone copy is a writeable point-in-time image of a NetApp FlexVol volume. 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 image and therefore significantly decrease the amount of disk space needed for SAP system copies.

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

On the source system, a database-consistent Snapshot copy of the DB2 containers 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, ignoring operation on the source system.

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

All the above steps can be fully automated and do not need any manual interaction. A SAP system copy can be accomplished in a few minutes using the NetApp solution.




Figure 11) SAP system cloning overview.

The table below compares the traditional approach to the NetApp approach to perform a SAP system copy.

Necessary steps at the source system:

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

With the NetApp approach, the backup will be taken using Snapshot functionality. Creating a Snapshot copy will take only a few seconds and has no performance impact on the source system. Therefore this step can be scheduled at any time during online operation. The creation of the Snapshot copy and all the subsequent steps can be fully automated.

Necessary steps at the target system:

If a new SAP test system needs to be set up, the SAP and DB2 software needs to be installed once. This step needs to be carried out with both approaches. With all subsequent refreshes of this system, this step is not necessary.

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

With the NetApp approach a FlexClone image will be created based on the consistent Snapshot database backup that was created at the source system. The creation of the FlexClone image will take only a few seconds and can be scheduled at any time. The following steps to adapt the file system and the database to the new SID and the SAP post-processing tasks can be fully automated.

Traditional Approach	NetApp Approach	NetApp Advantages
Necessary steps at the source system		
- Database backup... - Online backup - Copy active logs to new location OR - Shut down system - Offline backup - Restart system	- Database backup... - Write suspend - Snapshot - Copy active logs to shared location - Write resume OR - Shut down system - Snapshot - Restart system	- Traditional approach has a significant impact on performance during online backup - No impact on operation with NetApp online Snapshot solution; can be scheduled at any time - NetApp offline Snapshot takes only minutes - Difficult to schedule hours of downtime for offline backup with traditional approach
Necessary steps at the target system		
Install SAP system (if not existing yet)	Install SAP system (if not existing yet)	Same approach

• Offline backup... • Redirected restore of offline backup • Run db2relocatedb • SAP-specific post-processing tasks OR • Online backup... • Redirected restore of online backup • Run db2relocatedb • Roll forward recovery • SAP-specific post-processing tasks	• Offline Snapshot... • Create FlexClone image based on offline Snapshot backup; mount the FlexClone image at the target system • Run db2relocatedb • SAP-specific post-processing tasks OR • Online Snapshot... • Create FlexClone image based on online Snapshot backup; mount the FlexClone image at the target system • Adapt directory names to new SID • Run db2inidb with the RELOCATE USING option • Roll forward recovery • SAP-specific post-processing tasks	• Restore takes several hours compared to several seconds with FlexClone technology • Manual process vs. fully automatable process with NetApp solution
--	--	--

Conclusion: The NetApp system copy solution significantly improves the process to create SAP system copies.

- A system copy can be accomplished within several minutes compared to several days with the traditional approach.
- System copies can be scheduled at any time since there is no impact on the online operation of the source system (performance, backup infrastructure).
- Snapshot and FlexClone functionality reduces the time necessary to copy the data from the source system to the target system from several hours to several seconds.
- All storage, operating system, database, and SAP-specific tasks can be fully automated, therefore minimizing the interaction of IT staff.
- Snapshot and FlexClone functionality significantly reduces the necessary disk space for a SAP system copy by storing only data changes between the source and the target system.

SAP UPGRADE

BUSINESS CHALLENGES

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

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

Business processes are influenced during the upgrade project time period, since all development needs to be stopped and SAP support packages can't be imported. Therefore it is very important to minimize the overall time for the upgrade project.

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

During a SAP upgrade project SAP Basis administrators need to create several system copies to run the upgrade with current data from the development or production SAP system. The creation of a SAP system copy normally takes several days and negatively impacts the production environment. In addition a lot of manual steps are performed that consume valuable IT staff time.




Figure 12) SAP upgrade overview.

NETAPP SOLUTION

The NetApp solution for SAP upgrades addresses the issues above by providing a solution for creating fully automated SAP system copies in a few minutes. The NetApp backup and recovery solution helps to minimize the downtime during all upgrade phases with the capability to create database backups and restore databases in seconds. The NetApp solutions help to minimize the risk, reduce the downtime, and consume less IT staff time during a SAP upgrade project.

NETAPP SOLUTION FOR SAP UPGRADES

Upgrading the development system

The upgrade of the development system usually will be carried out on a copy of the current development system running on separate hardware. During the upgrade process of the copy of the development system the functionality of the upgrade is tested in the specific customer environment. In almost all cases the upgrade of the development system will be carried out more than once in order to define the necessary actions in all specific upgrade phases.

The setup of the separate SAP system will be done based on a system copy of the original development system. This system copy can be carried out using the NetApp system copy solution described in the section "NetApp Solution for SAP System Copies." Reducing the time needed to create this system copy is critical specifically because the copy is typically carried out several times. During the upgrade process and during the modification adjustment Snapshot backups are very helpful because they allow resetting the system to any Snapshot copy and restarting the upgrade phase.




Figure 13) SAP upgrade—development system.

Upgrading the quality assurance system

The upgrade of the quality assurance system is done using a fresh system copy of the production SAP system. One important result of the upgrade of the quality assurance system is the run time of the upgrade with real production data. The NetApp SAP system copy solution allows for efficient refreshing of the quality assurance system. Reducing the necessary time to create this copy is also critical when upgrading the quality assurance system because the copy is typically done more than once. Snapshot backups are helpful during the upgrade process and before the modification adjustments are imported. These Snapshot copies allow restoring the system to any specific Snapshot copy, allowing restart to an upgrade phase or restart of the import.




Figure 14) SAP upgrade—quality assurance system.

Upgrading the production system

Scheduling is extremely important when the production system is going to be upgraded, since the system is not available at various stages during the upgrade. Scheduling must also allow time for restoring the system to its former release status. Depending on the size of the database and the time and effort required for the functional test and importing the transports for the modification adjustment, one normal weekend may not be sufficient for the upgrade.

The upgrade of the production system includes at least three backups of the database. The first backup needs to be done directly before the upgrade is started. After the upgrade is done, 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 needs to be restored to the old release level.

Having Snapshot copies as a backup method and SnapRestore for restoring the system to its former release status ensures a higher level of flexibility with regard to scheduling. Normal tape backups will need several hours, a fact that needs to be considered when planning the upgrade schedule. The backup time is reduced to several minutes when using Snapshot and SnapRestore features.




Figure 15) SAP upgrade—production system.

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

- The NetApp system copy solution allows quick refreshing of the separate development and the quality assurance systems, reducing the time necessary from several days to several minutes.
- Using the NetApp backup and recovery solution, Snapshot backups of the database can be created that allow a database to be restored to any specific Snapshot copy in a few seconds and a restart of any upgrade phase.
- The NetApp backup and recovery solution significantly reduces the total upgrade time of the production system, providing a higher level of flexibility with regard to scheduling the SAP upgrade.

BUSINESS CONTINUANCE

BACKUP & RECOVERY

BUSINESS CHALLENGES

Corporations today require their SAP applications to be available 24 hours per day, seven days per week, 365 days per year. They expect performance to be maintained, irrespective of the increasing data volumes, and the system to undergo routine tasks, for example, backups, without influencing the SAP system. Performing backups of SAP databases is a critical task, as backups can have a significant performance impact on the production SAP system. Since 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 in which backups can be performed with minimum influence on the business process. Reducing downtime of SAP production and even development systems is critical, since it will always cause a financial impact on the business process. Thus the time needed for a restore and recovery is of particular importance.

To summarize SAP backup and recovery challenges:

- **Performance impact on production SAP system.** Backups typically have a significant performance impact on the production SAP system because there is a high load on the database server, the storage system, and the storage network during backups.
- **Shrinking backup windows.** Since conventional backups have a significant performance impact on the production SAP system, backups can only be made during times with low dialog or batch activities on the SAP system. It becomes more and more difficult to define a backup window when the SAP system is used 24x7.
- **Rapid data growth.** Databases are growing. Rapid data growth together with shrinking backup windows results in ongoing investments in the backup infrastructure—more tape drives, new tape drive technology, faster storage networks, etc. Growing databases also result in more tape media or disk space for backups. Incremental backups can address these issues, but result in a very slow restore process, which usually is not acceptable.
- **Increasing cost of downtime, decreasing mean time to recover (MTTR).** The mean time to recover is the time that is needed to recover from a database failure (logical or physical error). The MTTR cuts into two areas—time that is necessary to restore the database and time that is necessary to do the forward recovery of the database. The forward recovery time depends on the number of transaction logs that need to be applied after a restore. Unplanned downtime of a SAP system will always cause a financial impact on the business process. A significant part of the unplanned downtime is the time that is needed to restore and recover the SAP system in the case of a database failure. The backup and recovery architecture has to be designed according to the maximum acceptable unplanned downtime.
- **Backup and recovery time included in SAP upgrade projects.** The project plan for a SAP upgrade always includes at least three backups of the SAP database. The time needed to perform these backups will cut down the total available time for the upgrade process.

NETAPP SOLUTION

Network Appliance provides unique functionalities that address the above challenges.

Network Appliance Snapshot 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, since Snapshot does not move any data blocks. The use of Snapshot doesn't have any performance impact on the production SAP system, since the NetApp Snapshot implementation doesn't have to copy data blocks when the data in the active file system is changed. Therefore, creation of Snapshot copies can be scheduled without any concerns for affecting the

performance of dialog or batch activities. SAP and NetApp customers typically schedule several Snapshot online backups during the day—for instance, every four hours.

Snapshot also provides key advantages for the restore and recovery operation. The Network Appliance SnapRestore functionality allows restoring the entire database or parts of the database to the point in time of any available Snapshot copy. This restore process is done in a few minutes, independent of the size of the database. Because several Snapshot online backups have been created during the day, the time needed for the following recovery process is also dramatically reduced. Fewer logs need to be applied, because a restore can be done to an at most four-hour-old Snapshot copy. 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, Network Appliance recommends using Snapshot backups as a supplement, not a replacement for backups to a second location, whether backup to disk or tape. Though backups to a second 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 will be handled by using SnapRestore. Restores from a second location (disk or tape) are only necessary if the primary storage system holding the Snapshot copies is damaged or there is the need to restore a backup that is no longer available from a Snapshot copy—for instance, a two-week-old backup. Restores from a second location can be considered as a specific situation that will happen very rarely.

A backup and recovery solution using a Network Appliance storage system will always consist of two parts:

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

A backup to a second location will always be based on Snapshot copies created on the primary storage. Therefore, the data will be directly read 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 Network Appliance NearStore® system or other storage system and SnapVault® software:

- The primary storage directly communicates with the secondary storage (NearStore or other storage system) and sends the backup data to the destination. The Network Appliance 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 only copy 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. Since SnapVault only stores the changed blocks at the destination, a full database backup needs significantly less disk space.

Backup to tape using third-party backup software:

- **LAN backup.** The backup server mounts the Snapshot directory and writes the data to tape.
- **NDMP backup (serverless backup).** The tape is directly connected to the primary storage system. The data is written to tape using NDMP.




Figure 16) SAP backup and recovery—NetApp solution overview.

Table 1 evaluates the different backup and recovery concepts based on the most important attributes:

- Performance impact of the backup on the SAP production system: load that is generated by the backup on the database server, primary storage, and storage network
- Time that the database system will be in a “write suspended” state or offline
- Time needed for a full database backup
- Time needed for a full restore with recovery
- Space needed on tapes or disk for a full database backup
- Backup is usable for disaster recovery
- Integration into the SAP backup and recovery management tools

Conventional Backup to Tape

A conventional backup to tape using SAP Brbackup with backint generates a significant load on the production SAP system: on the database server, the primary storage, and the storage network. Since this backup is not based on Snapshot, the database will be offline or in hot backup mode during all the backup time. The time needed for the backup and for restore and recovery is high. Full backups will always need the full capacity on tape. The backup can be used for disaster recovery, and it is a fully integrated solution into SAP backup and recovery management.

Snapshot and SnapRestore

Snapshot backups do not generate any load on the database server, the primary storage, or the storage network. The database will be in a “write suspended” state or offline only for a few seconds. Using SnapRestore, the time needed for restore and recovery is very short. A full database backup based on Snapshot will only consume disk space for changed blocks. A partial integration into SAP can be accomplished by using the SAP job scheduler to run the Snapshot script.

Conclusion

Snapshot and SnapRestore have significant advantages in all areas compared to a conventional tape backup. In a disaster scenario the primary storage system holding the Snapshot backups might not be available any more. Therefore, a backup to a second location should be accomplished as well.

All further backup concepts to a second location are based on Snapshot copies created at the primary storage system. Therefore, with all concepts the database will only be offline or in hot backup mode for a short time.

Tape Backup with Backup Server over LAN

When doing a backup with a separate backup server that mounts the Snapshot file system, there is no load on the database server. There is still a high load on the primary storage system and on the storage network. The time needed for the backup and for restore and recovery is high. Full backups will always need the full capacity on tape. There is no integration into the SAP backup and recovery management.

Tape Backup with Backup Server and NDMP (Serverless Backup)

When using NDMP with a backup server, the NetApp storage system writes the data directly to tape. The time needed for the backup and for restore and recovery will be less, since the data is directly written to tape. Full backups will always need the full capacity on tape. There is no integration into the SAP backup and recovery management.

Disk-to-Disk Backup with NearStore and SnapVault

As SnapVault runs on the storage level, there will be no load on the database server. SnapVault only transfers the changed blocks with each backup. Therefore, the load on the primary storage and the storage network is significantly reduced. For the same reason the time needed to perform a full database backup is low. In addition, each full backup only stores the changed blocks at the destination. Therefore, the amount of disk space that is needed for a full backup is very low compared to that needed for full tape backups. When executing a restore, it is always possible to do a full restore without the need to restore all incremental backups.

Conclusion

A disk-to-disk backup using SnapVault has significant advantages in all areas compared to all the above tape backup concepts.

	Backup Load On			Time Need			Space Needed for Backup	Backup Usable for Disaster Recovery	Integration within SAP
	Database Server	Primary Storage	Storage Network	DB Offline or Hot Backup Mode	For Backup	For Restore and Recovery			
Conventional Backup to Tape	High	High	High	High	High	High	High	Yes	Yes
Snapshot and SnapRestore	No load	No load	No load	Very low	Very low	Very low	Low	No	Partly
Backup to a Second Location Based on Snapshot									
Tape Backup with Backup Server over LAN	No load	High	High	Very low	High	High	High	Yes	No
Tape Backup with Backup Server and NDMP	No load	High	No load	Very low	Medium	Medium	High	Yes	No
Disk-to-Disk Backup with Nearline Storage and SnapVault	No load	Low	Low	Very low	Low	Medium	Low	Yes	Partly

Table 1) Comparison of different backup and recovery concepts.

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 recover
- Minimum disk space needed for database backups at the primary and the secondary storage systems (filer and NearStore system)

The possibility of simply creating backups in seconds and being able to restore the SAP system to a point in time of any available Snapshot copy is also very helpful in SAP test and development environments. Projects such as data import, SAP upgrades, or installation of support packages can be accelerated using fast backup and restore functionalities. During these projects, backups can be done at specific phases, and the system can be easily and quickly reset to a starting point in order to be able to repeat that phase.




Figure 17) SAP testing cycle.

Carrying out a SAP upgrade or importing support packages and critical transports always involves SAP system downtime. It is important that this downtime be kept to a minimum and that the “old” status can always be restored. The specified system changes are usually first made in the development system in order to test the general functionality and procedures. In many cases, test systems must be upgraded several times, since problems can occur that can only be solved by restoring the system and restarting the upgrade. In this respect, Snapshot and SnapRestore can save a considerable amount of time. A tape backup does not have to be made; a Snapshot copy can be created instead. In the event of an error, the system can be quickly restored to its original status with SnapRestore, and the upgrade can be repeated.

Time management is extremely important when the production system is upgraded, since the system is not available at various stages during the upgrade. Scheduling must also include time for restoring the system to its former release status. Depending on the size of the database and the time and effort required for the functional test and importing the transports for the modification adjustment, one normal weekend may not be sufficient for the upgrade. Snapshot as a backup method and SnapRestore for restoring the system to its former release status ensure a higher level of flexibility with regard to scheduling. By creating several Snapshot copies at certain stages during the upgrade, it is possible to restart the upgrade without having to revert to the former release status. This can save a lot of time.

HIGH AVAILABILITY

BUSINESS CHALLENGES

Production SAP systems are business-critical applications that require 24x7 availability. Meeting these requirements 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 the SAP central instance.

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 ensures 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 can be created with two storage controllers by connecting the storage controllers via a cluster interconnect. This connection is redundant and is used to exchange cluster heartbeats and 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 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 via the cluster interconnect, no data is lost.

Because both storage controllers can be active in a cluster configuration, it is possible to use a single cluster to provide high availability for both the central instance and the database server. It is also possible to support other systems on the cluster.




Figure 18) 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 and central instance
- A clustered storage system is recommended if server clustering is used for the application

DISASTER RECOVERY

BUSINESS CHALLENGES

Organizations recognize the importance of having a bulletproof business continuance plan in place to deal with a disaster. The cost of not having one—lost productivity, revenue, and customer loyalty and possibly even business failure—makes it mandatory to have a plan that ensures an absolute minimum of downtime and rapid recovery from a disaster, with 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.

NETAPP SOLUTION

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 based on the SnapMirror configuration. 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 archive logs are stored has to be mirrored using SnapMirror. NetApp recommends a frequent SnapMirror update of the archive logs, e.g., every 10 minutes, to ensure a minimum of data loss.




Figure 19) 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 ensure that data there is completely current. The combination of failover and data replication ensures 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 disks. MetroCluster extends this cluster configuration to remote locations up to 30km. Because there is no physical connection to the cluster appliance's disk in case of a site failure, MetroCluster requires the use of SyncMirror® to insure that both storage controllers in the cluster have copies of the other storage controller's data.




Figure 20) MetroCluster over direct Fibre Channel switch connection.

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




Figure 21) MetroCluster over Fibre Channel and DWDM switch infrastructure.

This solution connects distant sites in metropolitan areas.

Conclusion:

- 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

ARCHIVING AND COMPLIANCE

BUSINESS CHALLENGES

Archiving

The long-term accumulation of data in the SAP database can ultimately affect 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, as 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 media type and platform for archival storage requires companies to conform to not just one but many content retention mandates. IT organizations have to respond by analyzing the business requirement and then choosing the proper solution based on factors such as time to data, risk, storage scalability, compatibility, and TCO. Current WORM (write once, read many) technologies like WORM optical disk and WORM tape do not provide sufficiently rapid access, high reliability, or low total cost of ownership (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 practically all public companies and industry sectors. Nearly every major corporation will put a regulatory compliance solution in place or face the risk of being exposed to litigation and fines. In most cases this solution is going to require 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), disk today represents about 10% of where regulated data is stored. Disk has not often been utilized due to a number of factors that include cost and the lack of necessity to retrieve information quickly. However, ESG estimates the moving forward disk will be the fastest growing area for the storage of regulated data.

NETAPP SOLUTION

SAP data archiving is based on the Archive Development Kit (ADK) in which archiving objects are used to remove data no longer needed in online business processes from the database and to store it in such a way that it is accessible in the future. The purpose of XML-based data archiving is the same as ADK-based archiving. The key difference is that it is based on universally accepted and widely used standards: the XML format is used to archive business objects, HTTP(s) as a communication service, and WebDAV as a general protocol for connecting storage systems.

The Archive Development Kit is the software layer that encapsulates the technical aspects of data archiving programs. ADK provides an application programming interface, also used by SAP, that customers and partners can use to develop their own archiving solutions. ArchiveLink is an interface as well as a service for facilitating the process-driven management of business documents. Business-related documents can be linked to and retrieved from application objects via workflow.

WebDAV stands for "Web-based Distributed Authoring and Versioning." It is a set of extensions to the HTTP protocol that allows users to collaboratively edit and manage files on remote Web servers. The major features of the protocol include locking, metadata management, and namespace manipulation.

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 NearStore and SnapLock®, work hand in hand with technologies from SAP and their archiving partners. The result of effective SAP archiving is better performing applications that cost less to operate and manage.

NetApp NearStore is the preferred compliance and archive storage subsystem for SAP landscapes. The NearStore product family leverages Network Appliance Data ONTAP 7G technology and takes full advantage of value-added software from Network Appliance, such as SnapLock. NearStore scales from 8TB to 96TB using more economical disk technology. With over 99.995% field-measurable uptime, NetApp RAID-DP technology enables NearStore systems to tolerate double disk failures with no data loss. Should additional capacity or performance be required for any reason, NetApp fabric-attached storage (FAS) systems may be substituted for NearStore within the SAP storage landscape.

SnapLock is the Network Appliance implementation of high performance disk-based magnetic WORM storage. SnapLock provides secure, storage-enforced data retention functionality via open file protocols such as CIFS and NFS while leveraging existing NetApp technologies to the greatest degree possible. This implementation also includes significant efforts in hardening Data ONTAP and its administrative interfaces to the degree that SnapLock can be deployed for protecting data in regulatory environments so strict that even the storage administrator is considered an untrusted party. An example of such an environment is the broker/dealer market regulated by SEC 240.17a-4. Alternate configurations of SnapLock can be deployed for unregulated or more flexible regulated environments.

SnapLock provides special purpose volumes in which files can be stored and committed to a non-erasable, non-rewritable state either forever or for a designated retention period. SnapLock allows this retention to be performed at the granularity of individual files through 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 or administrative access to the data must pass.

SnapLock is based on the open file protocol interfaces and does not require the use of any kind of proprietary API. You can perform all SnapLock-specific operations, such as setting file retention periods and committing files to WORM state, through regular file system operations available on all clients. Applications can use the regular programmatic library interfaces they would use for file operations on any other kind of storage system.

SnapLock is available in two versions. One or the other may be implemented within Data ONTAP but not both.

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

SnapLock Compliance enables organizations to satisfy strict records-retention regulations such as SEC Rule 17a-4 (broker-dealers), HIPAA (healthcare), Sarbanes-Oxley (public companies), 21CFR Part 11 (life sciences), and DOD 5015.2 (government). Only an act of willful destruction, such as physically removing disks from a NetApp system, can result in record deletion or alteration prior to the specified retention date.

SnapLock Enterprise enables adherence to rigorous organizational best practices through functionality similar to that of SnapLock Compliance, but allows administrators to delete entire SnapLock Enterprise volumes. Under no circumstances is it possible for any SnapLock Enterprise user or administrator to delete or modify individual SnapLock Enterprise WORM records or undermine SnapLock Compliance WORM volumes. SnapLock is supported on all NetApp FAS and NearStore platforms except the F7xx, F810, F820, F840, and V-Series.

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

- SnapLock enables locking of some files without forcing WORM behavior for all data
- No risk of software vendor lock-in; NetApp works well with existing document and content management packages such as iXOS-eCONserver and the FileNet P8 platform
- Can be managed and backed up using customer's current products and strategies
- Solution can incorporate existing Network Appliance or other vendors' storage
- Better ROI and lower TCO through increased availability, enhanced system performance, lower administration overhead, and increased staff productivity
- Compliance and archived data remains easily accessible on NearStore, a more cost-effective alternative for archiving SAP data than adding database storage or processing power

SAP ADAPTIVE COMPUTING

BUSINESS CHALLENGES

SAP system landscapes have become increasingly complex. Along with R/3 and BW systems, customers increasingly are installing CRM, SEM, Portal, and other SAP applications. Each of these applications requires database server, central instances, and possibly application servers.

These systems may also have complex usage profiles that require additional computing resources at different times of the day, month, or even year. In order to support the worst-case scenario, IT organizations are forced to invest large amounts of capital into computing resources that may go underutilized or unused for long times. While it is possible to redeploy these resources, the effort to do so often outweighs the cost because of outages and time required for employees to make the changes.

To address this challenge, SAP has developed an adaptive computing solution that allows companies to virtualize their computing, network, and storage resources. This allows them to shift the IT infrastructure to where it is needed most at any particular time. This makes the infrastructure flexible enough to meet changing business requirements while managing the total cost of a complex SAP solution. Network Appliance has embraced this technology and offers solutions to support a SAP adaptive computing environment.




Figure 22) SAP adaptive computing overview.

In a traditional SAP environment, each server and storage system is tied to a particular SAP application. With adaptive computing, the servers and storage become flexible and can be used to support any application in the landscape that is part of the adaptive environment. The SAP Adaptive Computing Controller (ACC) allows any SAP service to start on any available compute node in the compute node pool. Compute power can therefore be easily allocated to different SAP applications depending on the workload requirements of the specific application. For instance, compute nodes that are normally used for QA, tests, or training can be allocated to any production application supporting batch runs during the night or at month end for salary calculation within an HR system.

Another benefit of the adaptive landscape is the ability to "park" inactive systems. Companies with large SAP environments may have test and development systems that are inactive for long periods. In a traditional landscape, the system is restricted to a dedicated server. As a result, resources may be unused for long periods. Adaptive computing allows that server to be deployed to other systems while the development or test system is dormant. This allows the customer to support more systems with fewer servers.

NETAPP SOLUTION FOR SAP ADAPTIVE COMPUTING

NetApp storage systems provide the storage layer for the SAP adaptive computing environment. The SAP Adaptive Computing Controller allows any SAP service to start on any available compute node in the compute node pool. Therefore any compute node needs to have access to any data belonging to the specific SAP service.




Figure 23) Adaptive computing on NetApp with NFS.

NetApp storage solutions for SAP adaptive computing give customers the option to assign IT resources where needed. When using the NFS protocol to attach the compute nodes to the storage system, only standard hardware and software components are used.

- Standard Gigabit Ethernet NICs and NFS, TCP/IP protocol on server level
- Standard Gigabit Ethernet switches on network level
- Standard NFS and TCP/IP protocol on storage level

The NFS protocol is available within all UNIX operating systems and there is no need to implement any additional software layers. The NFS file system provides all functionalities of a shared file system that are necessary in an adaptive computing environment.

With the Linux operating system the compute nodes can be booted with a single shared operating system image on the NetApp storage system. With this approach there is no need for any OS deployment software and the local disks in the compute nodes are only used for paging. Additional compute nodes can be added to the pool without installing any operating system. New compute nodes can therefore be integrated into the landscape in several minutes.

With NetApp storage and SAP adaptive computing software, it is possible to meet the needs of a changing SAP landscape at a moment's notice.

Conclusion: NetApp technology in a SAP adaptive computing environment improves TCO.

- Reduced complexity
- Improved flexibility and scalability
- Improved infrastructure utilization

ADDITIONAL REFERENCES

SAP Backup & Recovery

Data ONTAP 7G: FlexVol and FlexClone

[Block Management with Data ONTAP 7G: FlexVol, FlexClone and space guarantees](#)

[A thorough introduction to FlexClone volumes](#)

RAID-DP

[NetApp Data Protection: Double parity RAID for enhanced data protection with RAID-DP](#)

Databases and Operating System Tuning for NFS

[Data ONTAP 7G: The ideal platform for database applications: FlexVol and FlexClone for database workloads](#)

[Database Layout with Data ONTAP 7G](#)

[Using the Linux NFS client with NetApp](#)

DB2 on NetApp Storage

[IBM® DB2 UDB Enterprise Server Edition V8 for UNIX®: Backup and Recovery Using a NetApp Storage System](#)

SAP

[SAP Adaptive Computing](#)

[Global SAP Homepage](#)

[SAP Service Marketplace](#)

[SAP Developer Network](#)

[SAP Help Portal](#)

Network Appliance, Inc.



© 2006 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, the Network Appliance logo, Data ONTAP, NearStore, SnapLock, SnapMirror, SnapRestore, SnapVault, and SyncMirror are registered trademarks and Network Appliance, FlexClone, FlexVol, RAID-DP, and Snapshot are trademarks of Network Appliance, Inc. in the U.S. and other countries. Solaris is a trademark of Sun Microsystems, Inc. Linux is a registered trademark of Linus Torvalds. DB2 and IBM are registered trademarks of IBM Corporation. UNIX is a registered trademark of The Open Group. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.