



Technical Report

SAP with Microsoft SQL Server on Windows Best Practices Using NetApp Clustered Data ONTAP and SnapCenter

Marco Schoen, NetApp
November 2015 | TR-4467

Abstract

This document provides customers and partners with best practices for deploying clustered NetApp® Data ONTAP® in support of SAP Business Suite solutions running in a Microsoft SQL Server on Windows environment.

TABLE OF CONTENTS

1	Introduction	4
1.1	Business Challenges Facing the SAP Customer	4
1.2	Technology Challenges of an Expanding SAP Landscape	4
1.3	NetApp Solutions for SAP	5
2	Storage Virtualization	5
2.1	Clustered Data ONTAP	5
2.2	Storage Management	8
3	Storage Configuration and Provisioning	8
3.1	Storage Layout	8
3.2	SAP System Installation	14
3.3	Microsoft iSCSI Software Initiator	14
3.4	Storage Migration	15
4	Business Continuance	16
4.1	Backup and Recovery	16
4.2	SAP Repair System	20
4.3	Disaster Recovery	21
5	System Management and Maintenance	23
5.1	SAP System Copy	23
5.2	SAP Testing Cycle	26
6	Conclusion	27
	References	27
	Best Practices Guides	27
	SAP	27

LIST OF TABLES

Table 1)	Recommended amount of SQL Server data files	10
Table 2)	FlexVol volume layout	11

LIST OF FIGURES

Figure 1)	Data ONTAP cluster overview	5
Figure 2)	LIF layout	9
Figure 3)	Aggregate layout	9

Figure 4) Minimal FlexVol volume layout.....	11
Figure 5) FlexVol volume layout with two data volumes.....	12
Figure 6) FlexVol volume layout.....	13
Figure 7) iSCSI initiator properties.....	15
Figure 8) Backup solution overview.....	18
Figure 9) Comparison of time required for different backup methods.....	19
Figure 10) Comparison of time needed for restore and recovery.....	19
Figure 11) SAP repair system.....	21
Figure 12) Disaster recovery with SnapMirror.....	22
Figure 13) Synchronous storage replication.....	23
Figure 14) Traditional SAP system copy.....	24
Figure 15) SAP system copy: NetApp approach.....	25
Figure 16) SAP system copy: standard approach.....	25
Figure 17) SAP system copy: NetApp approach.....	26
Figure 18) SAP testing cycle.....	26

1 Introduction

This document addresses the challenges of designing storage solutions to support SAP Business Suite products using a Microsoft SQL Server database. The primary focus is on the common storage infrastructure design, deployment, operation, and management challenges facing business and IT leaders who are using the latest generation of SAP solutions. Recommendations are generic; they are not specific to any given SAP application or 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. It was developed based on the interaction of technical staff from NetApp, SAP, Microsoft, and our customers.

1.1 Business Challenges Facing the SAP Customer

Corporations deploying SAP software today are under pressure to reduce cost, minimize risk, and control change by accelerating deployments and increasing the availability of their SAP landscapes. Changing market conditions, restructuring activities, and mergers and acquisitions often result in the creation of new SAP landscapes based on the SAP NetWeaver platform. Deployment of these business solutions is usually larger than a single production instance of SAP. Business process owners and project managers must coordinate with IT managers to optimize the scheduling and availability of systems to support rapid prototyping and development, frequent parallel testing or troubleshooting, and appropriate levels of end-user training. The ability to access these systems as project schedules dictate, with current datasets and without affecting production operations, often determines whether SAP projects are delivered on time and within budget. SAP systems are often used globally, resulting in 24/7 operation. Nondisruptive operation is therefore a key requirement.

1.2 Technology Challenges of an Expanding SAP Landscape

A typical SAP production landscape today consists of several different SAP systems. Just as important as the successful operation and management of these production instances are the many nonproduction instances that support them.

SAP recommends that customers maintain separate development and test instances for each production instance. In practice, standard SAP three-system (development, quality assurance, and production) landscapes often expand to include separate instances such as sandbox and user training systems. It is also common to have multiple development instances, as well as more than one system for quality assurance, testing, or perhaps a final staging system before releasing applications into production. Compound this with the many different SAP applications, such as ERP, CRM, BW, SCM, SRM, and Enterprise Portal, and the number of systems to support can become very large.

Adding to the challenge of maintaining these SAP systems is the fact that each of these instances has different performance and availability requirements. These requirements vary depending on the phase of the project and whether the project is focused on an existing SAP implementation or a new one. Projects rely on frequent refreshes of the nonproduction instances so that testing and training can occur with the current data.

As more test and training systems are required to accelerate test cycles by allowing parallel independent operation, the demand on the IT infrastructure increases. If the infrastructure that is supporting SAP systems and related applications is inflexible, expensive, and difficult to operate or manage, the ability of business owners to improve existing business processes and to deploy new ones might be restricted.

As SAP landscapes have expanded, the technology also has changed. SAP has evolved to take advantage of the latest technology trends. Database technologies such as SQL Server AlwaysOn functionality have introduced additional complexity into the database layer. Virtualization and cloud technologies have become predominant as corporations seek to leverage efficient computing methods to maximize their investment and reduce data center expenses. Without a storage infrastructure that can adapt to the needs of the changing technology, IT organizations would be unable to meet the business needs of the company.

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

2 Storage Virtualization

2.1 Clustered Data ONTAP

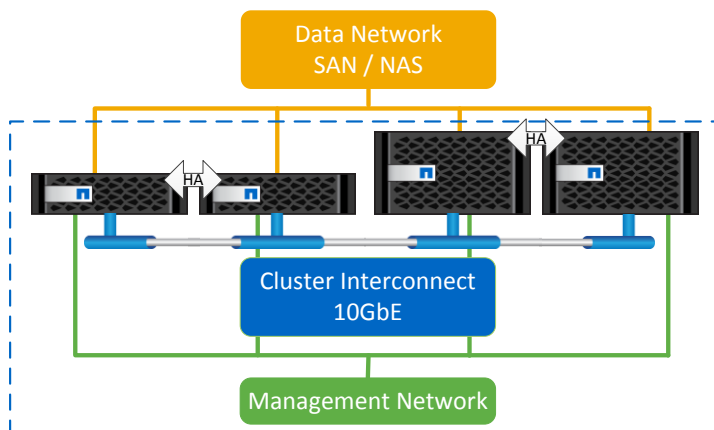
This section describes the architecture of clustered Data ONTAP, with an emphasis on the separation of physical resources and virtualized containers. Virtualization of storage and network physical resources is the basis for scale-out and nondisruptive operations.

Hardware Support and Basic System Overview

As shown in Figure 1, a clustered Data ONTAP system contains several storage controllers. The basic building block is the high-availability (HA) pair. An HA pair employs two identical nodes, or instances, of clustered Data ONTAP. Each node provides active data services and has redundant cable paths to the other node's disk storage. If either node is down for any reason, planned or unplanned, the HA partner can take over the failed node's storage and maintain access to the data. When the downed system rejoins the cluster, the partner node returns the storage resources.

The minimum cluster size is two matching nodes for an HA pair. Using NetApp nondisruptive technology refresh, a simple two-node, entry-level cluster can easily evolve into a much larger cluster with the addition of more nodes of the same type or more nodes of a more powerful controller model while the cluster is still online. At the time of writing, clusters with SAN protocols support up to eight nodes with midsize and high-end controllers. NAS-only clusters of high-end controllers scale up to 24 nodes.

Figure 1) Data ONTAP cluster overview.



One of the key differentiators for a clustered Data ONTAP environment is that the storage nodes are combined into a cluster to form a shared pool of physical resources that are available to SAN hosts and

NAS clients. This shared pool appears as a single system image for management purposes and provides a common point of management for the entire cluster through either GUI or CLI tools.

Scalability

Clustered Data ONTAP supports different controller types within the same cluster, protecting initial hardware investments and providing the flexibility to adapt resources to meet dynamic business demands and workloads. Similarly, support for different disk types, including SAS, SATA, and solid-state disk (SSD), makes it possible to deploy integrated storage tiering for different data types, together with the transparent DataMotion™ data migration capabilities of clustered Data ONTAP.

Clustered Data ONTAP can scale both vertically and horizontally through the addition of nodes and storage to the cluster. This scalability, combined with proven, protocol-neutral storage efficiency, supports the most demanding workloads.

Storage Efficiency and Data Protection

Storage efficiency built into clustered Data ONTAP offers substantial space savings, allowing more data to be stored at a lower cost. Data protection provides replication services so that valuable data is backed up and recoverable:

- **Thin provisioning.** Volumes are created by using virtual sizing. Thin provisioning is the most efficient way to provision storage because storage is not preallocated, even though the clients see the total storage space assigned to them. In other words, when a volume or LUN is created with thin provisioning, no space on the storage system is used. The space remains unused until data is written to the LUN or the volume. At that time, only the space to store the data is used. Unused storage is shared across all volumes, and the volumes can grow and shrink on demand.
- **NetApp Snapshot® copies.** Automatically scheduled, point-in-time Snapshot copies take up no space and incur no performance overhead when created. Over time, Snapshot copies consume minimal storage space because only changes to the active file system are written. Individual files and directories can easily be recovered from any Snapshot copy, and the entire volume can be restored back to any Snapshot state in seconds.
- **NetApp FlexClone® volumes.** These near-zero-space, exact, writable, virtual copies of datasets offer rapid, space-efficient creation of additional data copies that are well suited for test and development environments.
- **Deduplication.** Duplicate data blocks are removed in primary and secondary storage. Only unique blocks are stored, which results in storage space and cost savings. Deduplication can run on a customizable schedule.
- **Compression.** Data blocks are compressed by replacing repeating patterns within a subset of a file. Compression is complementary with deduplication. Depending on the workload, compression only, deduplication only, or deduplication and compression together could provide the maximum space and cost savings.
- **NetApp SnapMirror® data replication software.** Asynchronous replication of volumes is supported, independent of the protocol, either within the cluster or to another clustered Data ONTAP system, for data protection and disaster recovery.
- **NetApp SnapVault® backup software.** Volumes can be copied for space-efficient, read-only, disk-to-disk backup, either within the cluster or to another clustered Data ONTAP system.
- **NetApp MetroCluster™ software.** Continuous data availability is supported beyond the data center or the cluster. MetroCluster is native within the NetApp Data ONTAP operating system. It provides a synchronous mirroring relationship between two distinct but identically configured two-node clusters up to 200km apart.

Cluster Virtualization and Multi-Tenancy Concepts

A cluster is composed of physical hardware, including storage controllers with attached disk shelves; network interface cards (NIC); and, optionally, Flash Cache™ cards. Together these components create a physical resource pool, which is virtualized as logical cluster resources to provide data access. Abstracting and virtualizing physical assets into logical resources provide flexibility and, potentially, multi-tenancy in clustered Data ONTAP. These processes also enable the DataMotion capabilities at the heart of nondisruptive operations.

Physical Cluster Components

Storage controllers, independent of the model, are considered equivalent in the cluster configuration, in that they are all presented and managed as cluster nodes. Clustered Data ONTAP is a symmetrical architecture, with all nodes performing the same data-serving function.

Individual disks are managed by defining them into aggregates. Groups of disks of a particular type are protected with NetApp RAID DP®. NICs and host bus adapters (HBAs) provide physical ports (Ethernet and FC) for connections to the management and data networks. The physical components of a system are visible to cluster administrators but not directly to the applications and hosts that use the cluster. The physical components provide a pool of shared resources from which the logical cluster resources are constructed. Applications and hosts only access data through storage virtual machines (SVMs), which contain volumes and logical interfaces.

Logical Cluster Components

The primary logical component of a cluster is the SVM; all client and host data access is through an SVM. Clustered Data ONTAP supports a minimum of one and up to hundreds of SVMs in a single cluster. Each SVM is configured for the client and host access protocols it supports in any combination of SAN and NAS. Each SVM contains at least one volume and at least one logical interface.

The administration of each SVM can optionally be delegated so that separate administrators are responsible for provisioning volumes and other SVM-specific operations. This is particularly appropriate for multi-tenant environments or when workload separation is desired. SVM-delegated administrators have visibility to only their specific SVM and have no knowledge of any other hosted SVM.

For NAS clients, the volumes in each SVM are joined together into a namespace for CIFS and Network File System (NFS) access. For SAN hosts, LUNs are defined within volumes and mapped to hosts.

The accessing hosts and clients connect to the SVM through a logical interface (LIF). LIFs present either an IP address (used by NAS clients and iSCSI hosts) or a worldwide port name (WWPN, for FC and FCoE access). Each LIF has a home port on a NIC or HBA. LIFs are used to virtualize the NIC and HBA ports rather than for mapping IP addresses or WWPNs directly to the physical ports, because there are almost always many more LIFs than physical ports in a cluster.

Each SVM requires its own dedicated set of LIFs, and up to 128 LIFs can be defined on any cluster node. A LIF defined for NAS access can be temporarily migrated to another port on the same or a different controller to preserve availability, rebalance client performance, or evacuate all resources on a controller for hardware lifecycle operations.

By virtualizing physical resources into the virtual server construct, Data ONTAP implements multi-tenancy and scale-out and allows a cluster to host many independent workloads and applications.

Storage QoS

Clustered Data ONTAP provides storage quality of service (QoS) policies for cluster objects. An entire SVM or a group of volumes or LUNs within an SVM can be dynamically assigned to a policy group, which specifies a throughput limit, defined in terms of IOPS or MBps. This can be used to reactively or proactively throttle workloads and prevent them from affecting other workloads. QoS policy groups can

also be used by service providers to prevent tenants from affecting each other as well as to avoid performance degradation of the existing tenants when a new tenant is deployed on the shared infrastructure.

For more information, refer to [NetApp Clustered Data ONTAP 8.3 and 8.2.x: An Introduction](#).

2.2 Storage Management

NetApp SnapCenter® data management software offers a rich set of capabilities to virtualize and enhance storage management for Microsoft Windows and for Microsoft SQL Server. The SnapCenter plug-in for Windows integrates tightly with the native file system and provides a layer of abstraction between application data and physical storage associated with that data.

Business does not have to stop every time the IT organization has to add more storage. With SnapCenter, adding, deleting, mapping, unmapping, and mirroring virtual disks can be done while systems are online. Capacity can be expanded with limited or no effect on application or system performance.

The SnapCenter plug-in for Windows integrates Snapshot technology to capture near-instantaneous, point-in-time images of application and user data. It also gives access to Snapshot copies by mounting them as virtual disks. These virtual disks can be used for routine administrative tasks such as online backup, testing new applications, or populating data marts with limited or no downtime of business-critical information. Restoring data can be done in minutes with SnapRestore® and SnapCenter.

SnapCenter makes management simple and intuitive in Windows environments by allowing administration through the centralized web GUI or PowerShell commands. Interactive wizards and easy-to-use interfaces guide administrators through all management tasks and create automatic schedules of operations.

The SnapCenter plug-in for SQL Server paves the way for database and storage administrators to simplify data management by using the powerful capabilities of NetApp storage systems.

SnapCenter automates and simplifies the complex, time-consuming manual processes associated with the backup, restore, recovery, and cloning of SQL Server databases. It is integrated with Microsoft technology iSCSI and FC protocols to allow IT organizations to:

- Scale their storage infrastructure
- Meet increasingly stringent SLA commitments
- Improve the productivity of database and storage administrators across the enterprise
- Support federated backup of several independent SQL Server databases

3 Storage Configuration and Provisioning

3.1 Storage Layout

Storage Virtual Machine


The use of one storage virtual machine (SVM) is recommended for accessing all SAP systems within one SAP landscape. This offers the possibility of creating clones of the database by using the SnapCenter plug-in for SQL Server: for example, for performing SAP system copies. The aggregate hosting the source database must be assigned to the SVM to allow SnapCenter to create clones. Execute the following command to do so:

```
> vservers modify -vservers <vservers name> -aggr-list <list of used aggregates>
```


As shown in Figure 2, two data LIFs for the desired protocol (iSCSI, FCoE, or FCP) need to be configured on each physical controller and on the corresponding HA partner where data is stored and where it is planned to store data in the future. In the case of iSCSI, one data LIF is suitable if the data LIF is assigned to an interface group. The LIFs on one controller should be assigned to separate physical ports that are connected to different switches. This makes sure that the connection path contains no single point of failure, and it also increases the available bandwidth. In addition, it avoids nonoptimal data paths by using the cluster interconnect.

Note: Depending on throughput requirements, additional data LIFs might be required.

Figure 2) LIF layout.

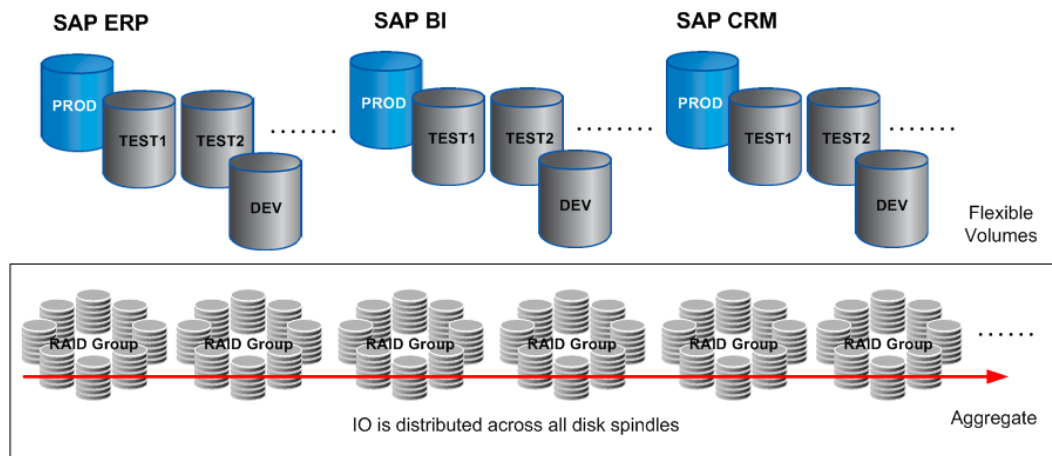


Aggregate Layout

As shown in Figure 3, 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 10. Data loss occurs only if three disks within the same RAID group fail at the same time.

Figure 3) Aggregate layout.



The design of the physical disk layout is very simple because it is not done on a per-SAP system basis. The aggregate is created as a physical storage resource pool, and storage resources are assigned on a

logical, virtualized level with FlexVol volumes. The size of the volumes can be easily increased or decreased during online operation without any reconfiguration of the underlying physical disk structure. This allows optimal utilization of the storage resources.

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

With shared resources, 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 production systems because too many I/O resources might be used by the test systems. Quality of service (QoS) can address this issue. QoS provides control of service for clustered Data ONTAP systems. With QoS, resources can be prioritized on the FlexVol volume level. Production 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.

SQL Server Data File Layout

The amount of SQL Server data files for the SAP database depends on the number of CPU threads/vCPUs used for the database server. Table 1 shows the amount of data files for some t-shirt sizes.

Table 1) Recommended amount of SQL Server data files.

T-Shirt Sizes	Amount of CPU Threads/vCPUs	Amount of Data Files
Small	4–8	4
Medium	8–32	8–16
Large	32–64	16–32
Extra large	64–128	32–64
Extra extra large	>128	64

LUN Layout

The size of the database and the kind of SAP system determine the number and size of LUNs required. The goal is to find a balance between the performance advantages of a large number of smaller LUNs and the ease of management that comes with a smaller number of larger LUNs. NetApp recommends storing two SQL Server data files in one LUN.

Minimal FlexVol Volume Layout

As shown in Figure 4, each SAP system uses five FlexVol volumes:

- One volume for the database data files
- One volume for the database log file or files
- One volume for the executables, including SQL Server system databases
- One volume for SQL Server TempDB
- One volume for the SQL Server log backup directory (and the quorum disk, in the case of Microsoft failover cluster)

The SQL Server data files are separated from the log files and log file backups (log backup directory). It is important to store the database data files in a FlexVol volume separated from the log files to allow use of Snapshot copies, SnapRestore, FlexClone, and other Data ONTAP features that work on the volume level. If log files and data files are stored in the same FlexVol volume, using SnapRestore to restore the

volume eliminates the ability to roll the database forward past the time of the Snapshot copy, resulting in potential data loss.

Figure 4) Minimal FlexVol volume layout.

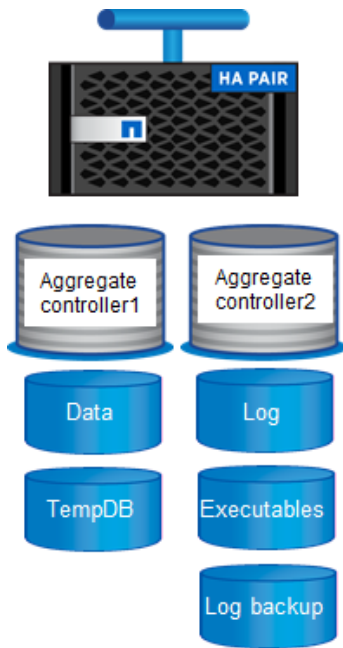


Table 2 shows the distribution of the file systems/LUNs of a single SAP instance to the FlexVol volumes. Three LUNs are configured in the SIDData FlexVol volume for the database data files, and one LUN is configured for the log files in the log FlexVol volume. Two LUNs in the executables FlexVol volume are used to store the SAP and SQL Server binaries and SQL Server system databases. Additionally, one LUN for the SQL Server TempDB is stored in the TempDB FlexVol volume, and one LUN is configured in the SnapInfo FlexVol volume for the SMS SQL SnapInfo directory. The quorum disk of a Microsoft failover cluster should be stored in the SnapInfo FlexVol volume.

Table 2) FlexVol volume layout.

Storage Controller 1		Storage Controller 2		
Aggregate 0		Aggregate 1		
Data	TempDB	Executables	Log	Log Backup
SIDData1	TempDB	SQL Server executables	Log files	Log backup dir
SIDData2		System DBs		Quorum failover cluster
SIDData3		SAP executables		

FlexVol Volume Layout with Distributed Data Files

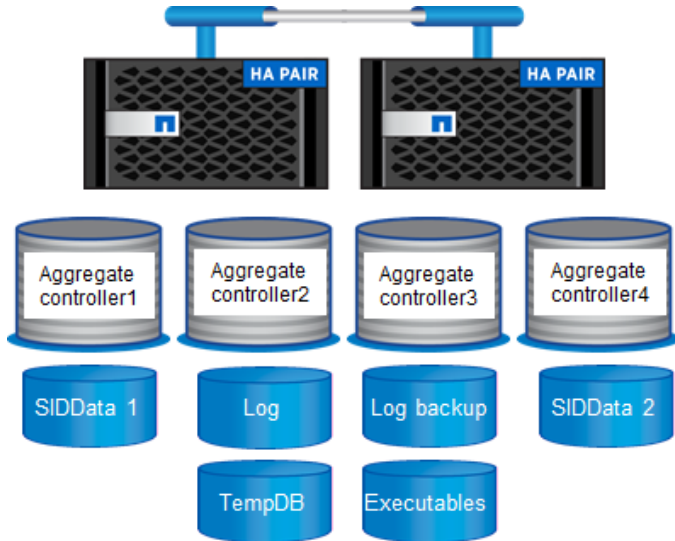
As shown in Figure 5, each SAP system uses six FlexVol volumes:

- Two volumes for the database data files
- One volume for the database log file or files
- One volume for the executables, including SQL Server system databases

- One volume for SQL Server TempDB
- One volume for the SQL Server log backup directory (and the quorum disk, in the case of failover cluster)

Compared to the minimal FlexVol volume layout, the LUNs for the data files are distributed to different volumes. This makes it possible to increase the throughput beyond the borders of one physical storage controller.

Figure 5) FlexVol volume layout with two data volumes.



The distribution of LUNs to volumes is similar to the minimal FlexVol, but the LUNs of the data files are distributed to two different volumes. Figure 5 shows that in addition, the volumes for the executables and SQL Server log backup directory are stored on a separate aggregate. This is optional.


FlexVol Volume Layout for Highest Flexibility

For the easiest, most granular, and most seamless migration of database files to another storage controller—either to redistribute the load or to use a new storage controller—use one volume for each LUN.

As shown in Figure 6, an SAP system with 16 SQL Server data files uses 12 FlexVol volumes:

- Eight volumes for the database data files (Keep in mind that two data files are recommended per LUN.)
- One volume for the database log file or files (If several log files are required, each should be stored in its own LUN.)
- One volume for the executables, including SQL Server system databases
- One volume for SQL Server TempDB
- One volume for the SQL Server log backup directory (and the quorum disk, in the case of failover cluster)

Figure 6) FlexVol volume layout.



The distribution of LUNs to volumes is similar to the previous layouts, but the LUNs of the data files are distributed to eight different volumes. Figure 6 shows that, in addition, the volumes for the executables and SQL Server log backup directory are stored on a separate aggregate. This is optional.

Sizing

This section is an overview of storage sizing for an SAP environment using NetApp storage. It provides a basic understanding of what kind of information is important in performing storage sizing and how these requirements influence the storage landscape.

NetApp can provide storage sizings to SAP customers, based on a sizing questionnaire that the customer completes. Customers should work with their NetApp account team to request and perform a sizing.

Storage sizing for an SAP landscape is based on several conditions that are defined by customer requirements. Together, these requirements define the required storage infrastructure:

- Throughput 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
- High-availability requirements

For existing SAP systems, the throughput load is measured with database or operating system tools. Independent of which tools are used, it is important to do the measurement during peak loads on the SAP system. When database tools are used for the measurement, a suitable time frame, such as one hour or less, must be chosen, because these tools calculate an average value, and the throughput sizing must be based on peak values.

For new SAP systems, where a throughput measurement is not possible, the SAP Application Performance Standard (SAPS) values for the systems, which are provided by the SAP Quick Sizer, can be used to estimate the throughput requirements. The storage sizing is much more accurate when real throughput values are measured. SAPS-based sizing should be done only if no other data is available. In the case of a SAPS-based sizing, it is important to perform additional monitoring to review the sizing at different points in the project to be sure that the original SAPS-based sizing is accurate.

The type and number of disk spindles and storage controllers are determined based on the throughput requirements.

To determine the required capacity, the following information must be available:

- Size of each database
- Growth rate
- Daily change rate
- Number and retention policy of Snapshot copies
- Number and durability of FlexClone volumes
- Synchronous or asynchronous mirroring

The type and number of disks and the storage controller that support the capacity are determined based on the capacity requirements.

The results of the throughput sizing and the capacity sizing are compared in a final step to define the right storage system to support both the throughput and capacity requirements.

3.2 SAP System Installation

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

With VMware, either raw device mapping (RDM) or direct iSCSI out of the virtual machine (VM) is recommended to provide the LUNs to the operating system. With Microsoft Hyper-V, either direct iSCSI access out of the VM or pass-through disks are recommended to provide LUNs to the virtual machine.

General Requirements

NetApp recommends the use of SnapCenter plug-in for Windows, a NetApp software product that simplifies storage management and provisioning in an SAP and Windows storage environment. It is required for SnapCenter plug-in for SQL Server, which provides host-consistent database Snapshot copies, SnapRestore, and database cloning functionality. For information about SnapCenter and SnapCenter plug-ins, see [SnapCenter Overview](#).

Creating LUNs and connecting them to the operating system, Snapshot backups for database applications are not consistent from the database point of view without integration in the database management system. Therefore, automatically scheduled Snapshot copies on the storage level should be turned off on database volumes. SnapCenter plug-in for Windows automatically sets the correct options on the storage system volumes.

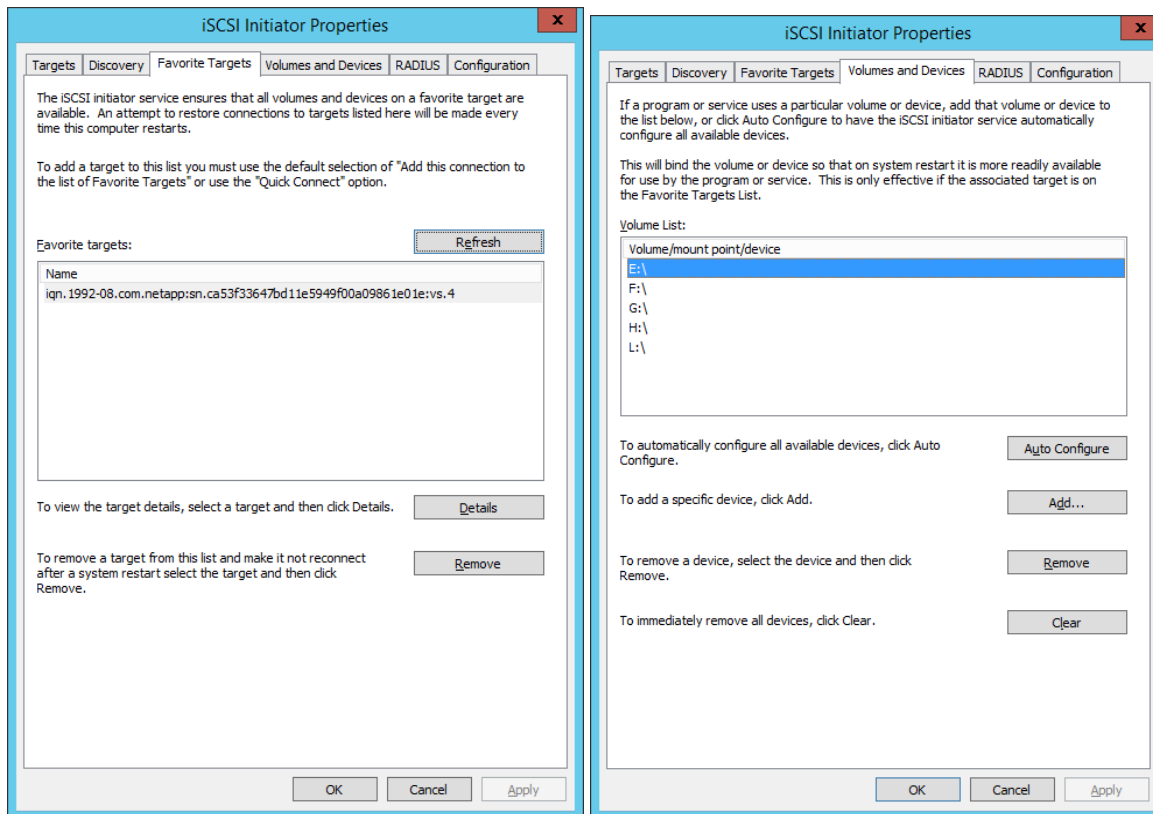
All LUNs described in section 3.1, “Storage Layout,” are created and connected to the Windows operating system by using SnapCenter plug-in for Windows.

SnapCenter plug-in for Windows sets all required storage parameters such as FlexVol options and partition offset. It also manages the device mapping through the hypervisor, VMware ESX server, or Microsoft Hyper-V.

3.3 Microsoft iSCSI Software Initiator

If iSCSI LUNs are connected using the Microsoft iSCSI initiator, make sure that the SVM is listed as the favorite target and the LUNs are listed in the volume list of the Microsoft iSCSI software initiator, as shown in Figure 7.

Figure 7) iSCSI initiator properties.



If this is not the case, follow the instructions mentioned at the iSCSI initiator property pages. Otherwise the LUNs might be not available after a reboot of the server.

3.4 Storage Migration

This section discusses different storage migration approaches. If a migration includes a change of the operating system or database system, it cannot be done solely at the storage level, and SAP migration tools must be used. These tools export the data from the source environment and import it into the target environment. The approach is therefore defined by SAP and is independent of the storage system used.

Overview of Migration Approaches

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

- Migration on the operating system level
- Migration on the database level

Migration on the Operating System Level

In addition to the existing storage system, the NetApp storage system is connected to the database server. The NetApp storage system is configured, and the LUNs are mounted to the server. Before the data migration is started, the database and the SAP system must be shut down. The data is then copied using the server from the old storage system to the NetApp system. When all data has been copied, the old storage system is disconnected from the database server. If the file system structure (drive letters)

remains the same, the database can be started immediately. If there is a change in the file system structure, the new structure must be configured by using SQL Server Management Studio.

A migration on the operating-system level can be done for an FCP-to-FCP or an FCP-to-iSCSI migration. The disadvantage of this approach is that the SAP system is not available while the database files are being copied. Depending on the database size, the downtime could be several hours.

Migration on the Database Level

A SQL Server database backup is restored to the NetApp storage system. To minimize the impact on the production SAP system, the restore can be done by using a separate server connected to the NetApp storage. In addition, the log backups are continuously copied to the separate server. Before the final migration is started, the SAP system must be shut down, and a log backup has to be performed. The log backups that have not yet been copied are now copied from the old storage system to the NetApp storage system, and a forward recovery of the database is carried out. Then, the databases are detached by using the SQL Server Management Studio, the old storage is disconnected from the database server, and the NetApp storage system and the LUNs stored there are connected.

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

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

4 Business Continuation

4.1 Backup and Recovery

Enterprise requires their SAP applications to be available 24 hours a day, 7 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 that 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 difficult to define a point in time when backups can be performed with minimal 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 list summarizes 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.** Conventional backups have a significant performance impact on the production SAP system; backups can be made only during times when few dialog or batch activities are taking place on the SAP system. It becomes more difficult to define a backup window when the SAP system is used 24/7.
- **Rapid data growth.** 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. Growing databases also result in more tape media or disk space for backups. Incremental backups can address these issues, but they result in a very slow restore process, which is usually not acceptable.
- **Increasing cost of downtime.** Unplanned downtime of an SAP system always has a financial impact on the business. A significant part of the unplanned downtime is the time that is needed to restore

and recover the SAP system in case of a failure. The backup and recovery architecture must be designed based on an acceptable recovery time objective (RTO).

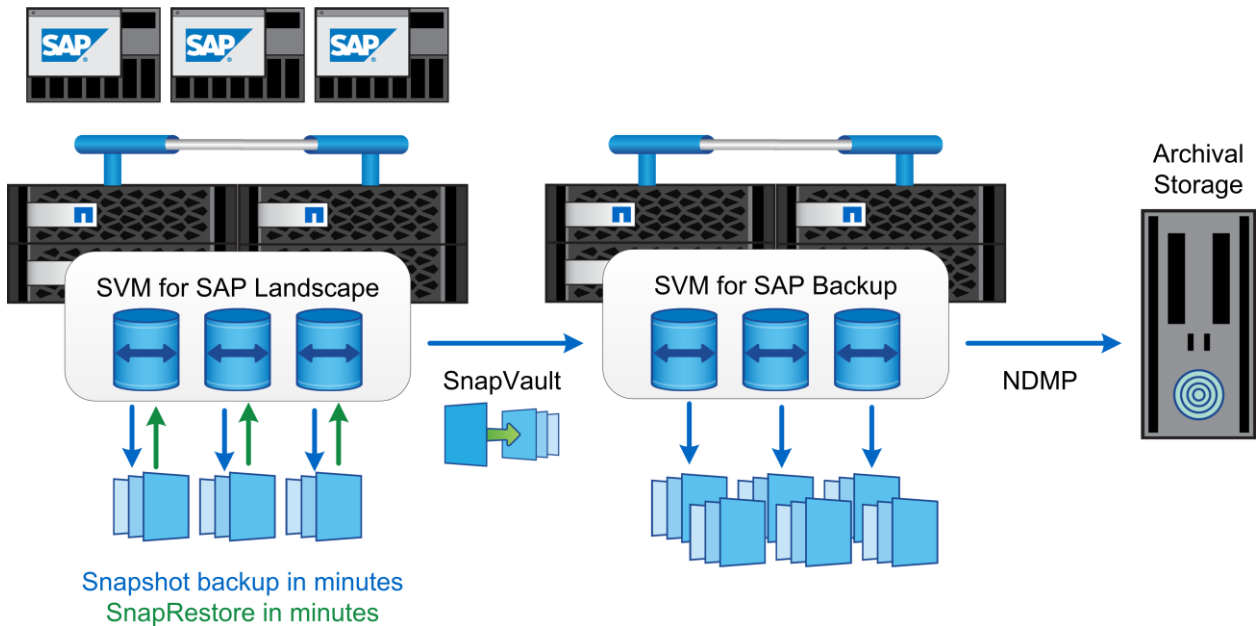
- **Backup and recovery time included in SAP upgrade projects.** The project plan for an SAP upgrade always includes at least three backups of the SAP database. The time needed to perform these backups cuts down the total available time for the upgrade process. The go/no-go decision is based on the amount 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.

NetApp Snapshot technology can create an online or offline database backup in minutes. The time needed to create a Snapshot copy is independent of the size of the database, because a Snapshot copy does not move any data blocks. The use of Snapshot technology has no performance impact on the production SAP system because the Snapshot technology implementation does not copy data blocks when the copy is created or when data in the active file system is changed. Therefore, the 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 4 hours. These Snapshot backups are typically kept for 3 to 5 days on the primary storage system.

Snapshot copies also offer key advantages for restore and recovery operations. The NetApp SnapRestore functionality allows restoring 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 by using a Snapshot copy that is at most eight hours old, fewer transaction logs need to be applied. The mean time to recover (the time needed for restore and recovery) is therefore reduced to several minutes, compared to several hours with conventional tape backups.

Snapshot backups are stored on the same disk system as the active online data. Therefore NetApp recommends using Snapshot backups as a supplement, not as 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 restore and recovery. Most restore and recovery actions are handled by using SnapRestore on the primary storage system. 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.

Figure 8) Backup solution overview.



As shown in Figure 8, a backup and recovery solution using a NetApp storage system always consists of two parts:


- Backup and restore using Snapshot and SnapRestore technologies
- Backup and restore to and 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. There are two options to back up the data to a secondary location:

- **Disk-to-disk backup using SnapVault software.** The primary storage virtual machine communicates directly with the secondary storage virtual machine 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 of the data has to be transferred from the source to the destination, all subsequent backups copy only the changed blocks to the secondary storage. The typical block change rate for an SAP system is between 2% and 5% per day. 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. Backing up data to tape as a long-term backup might still be required. This could be, for example, a monthly backup that is kept for a year. In this case the tape infrastructure can be directly connected to the secondary storage virtual machine, and the data is written to tape using NDMP.
- **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 9 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 9) Comparison of time required for different backup methods.



Snapshot Backups Together with NDMP Backups


Snapshot copy backups do not generate any load on the database server or the primary storage system. A full database backup based on Snapshot copies consumes disk space only for changed blocks. Snapshot copy backups are typically scheduled frequently: for example, every 4 hours. Frequent backups allow a more flexible restore process and reduce the number of logs that must be applied during forward recovery. In addition, a full NDMP backup to tape is scheduled once a day. This backup creates a heavy load on the primary storage system and takes the same amount of time as a 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 subsection. Because SnapVault runs at the storage level, there is no load on the database server. SnapVault transfers only the changed blocks with each backup, so the load on the primary storage system 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 that for a full tape backup.

Figure 10 compares the time required to perform restore and recovery.

Figure 10) Comparison of time needed for restore and recovery.



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 to perform 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 minutes. Snapshot backups are created frequently, such as every 4 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.

The combination of Snapshot and SnapRestore with disk-to-disk backup based on SnapVault offers the following significant improvements over conventional tape backups:

- Negligible impact of backups on the production SAP system
- Heavily reduced RTO
- Minimum disk space needed for database backups on the primary and the secondary storage systems

Database Verification

Database verification is an important part of a backup concept. Snapshot backups are excellent for running database consistency checks. SnapCenter plug-in for SQL Server offers the possibility of running a database consistency check on a separate server automatically or manually after a backup without creating any load on the production database system.

4.2 SAP Repair System

More and more companies are facing the challenge of addressing logical errors in a complex SAP environment, where several SAP systems exchange data with each other.

A logical error can be addressed by restoring the system using the last backup and doing a forward recovery up to the point before the logical error occurred. This approach has the following disadvantages:

- Downtime for the analysis of when the logical error occurred and for the restore and recovery process
- Data loss, because the system is recovered to a point in time in the past
- Inconsistency between the system that got restored and recovered to a point in time in the past and the other systems that exchange data with that system

Therefore, SAP customers are looking for a more efficient and flexible solution to address logical errors. The NetApp Snapshot and FlexClone technologies help to provide a solution that allows recovery from logical errors without the need to restore and recover the affected system.

Figure 11) SAP repair system.

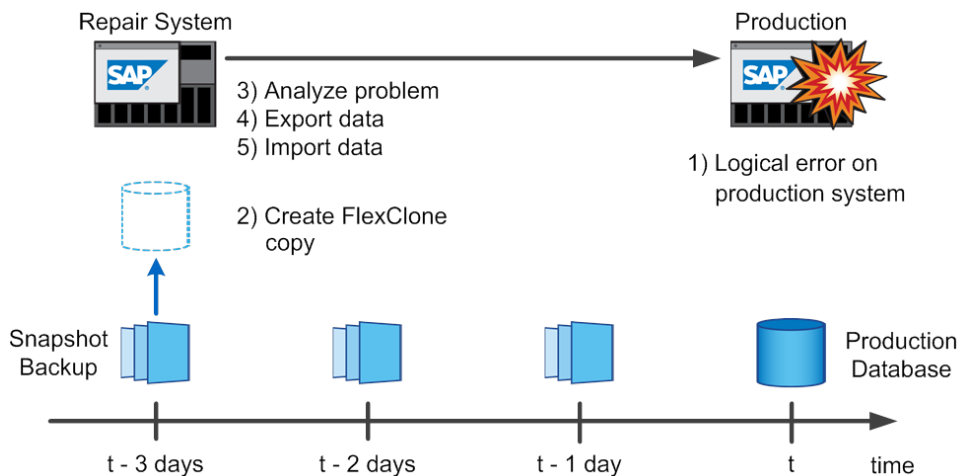


Figure 11 shows the general process of creating and using the repair system.

1. A logical error is discovered on the production system. Depending on the kind of logical error, the decision can be made to shut down the production system or to keep it online so that only parts of the business process are affected.
2. Several Snapshot backups of the production system are available, and any of them can be used to create an SAP system copy of the production system. The SAP system copy is created by using a FlexClone copy of the Snapshot copy.
3. The repair system is used to analyze the problem.
4. The appropriate data is exported from the repair system.
5. The data is imported into the production system.

In this example, there is little or no impact on the production system, no data loss, and no inconsistency in the SAP landscape.

The described scenario is quite simple, and it is obvious that not all logical errors can be solved this easily. However, the repair system approach also helps in more complex scenarios because there is greater flexibility and there are more options to analyze and to recover from the logical error.

4.3 Disaster Recovery

Organizations recognize the importance of having a business continuance plan in place to deal with a disaster. The costs of not having one—loss of productivity, revenue, customer loyalty, and possibly even business failure—make it mandatory to have a plan that enables minimum downtime and rapid recovery from a disaster.

A disaster recovery solution based on SnapMirror and MetroCluster fulfills and exceeds all business requirements for running even global SAP systems. By replicating data at high speeds over a LAN or a WAN, SnapMirror software provides data high availability and fast recovery.

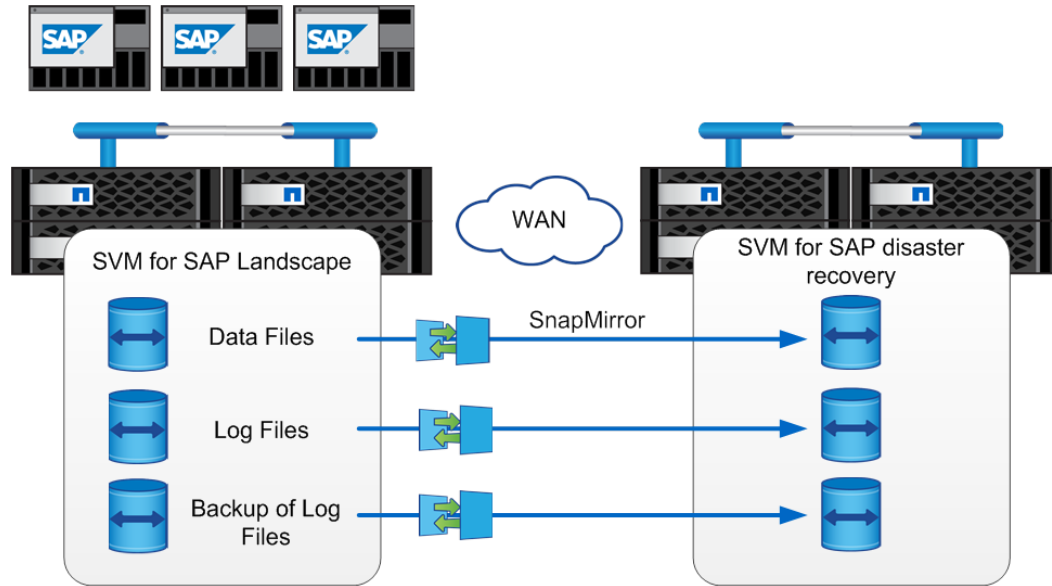
Asynchronous Replication

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

SnapMirror performs an initial transfer to initialize the disaster recovery site. Incremental changes are then passed to the disaster recovery site asynchronously. The SnapMirror disaster recovery solution is

based on the NetApp backup and recovery solution whereby Snapshot copy backups are mirrored to the disaster recovery site. Additionally, the volumes where the log files and the log file backups are stored are mirrored by using SnapMirror. The frequency of SnapMirror updates of the log files and log backups determines the amount of data lost in the event of a disaster. Figure 12 shows disaster recovery with SnapMirror.


Figure 12) Disaster recovery with SnapMirror.



Synchronous Replication

The synchronous disaster recovery solution for SAP on SQL Server is based on NetApp MetroCluster software. Figure 13 shows a high-level overview of the solution. The storage cluster at each site provides local high availability and is used for production workloads. The data on each site is synchronously replicated to the other location and is available in case of disaster failover.

Figure 13) Synchronous storage replication.



5 System Management and Maintenance

5.1 SAP System Copy

Business Challenges

A typical SAP customer environment consists of different SAP Business Suite and SAP NetWeaver components. Copies of SAP components are required to test application patches, run performance and data integrity tests, or provide user training environments. A typical SAP customer needs about 10 copies of different SAP components. These copies must be refreshed frequently, often on a weekly or monthly basis.

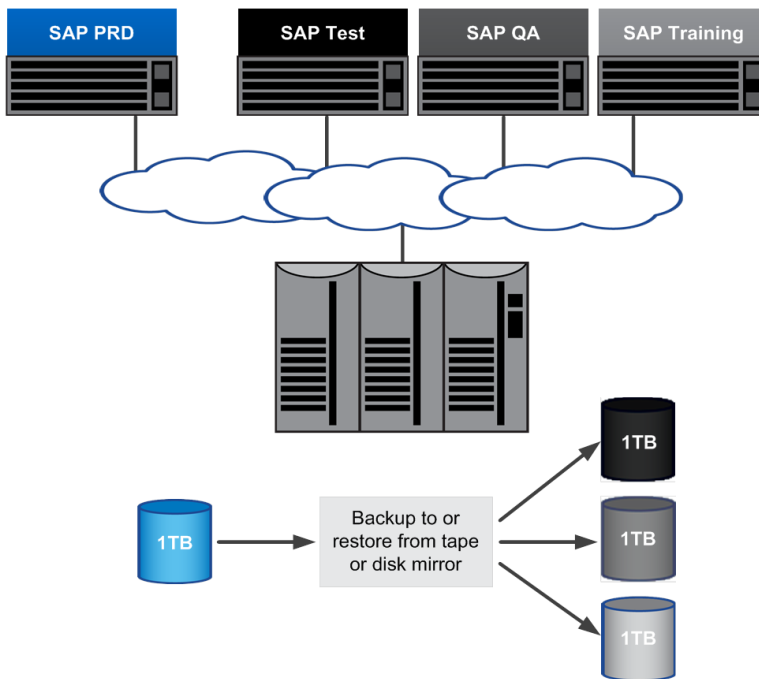
Rapid and space-efficient provisioning of test systems allows SAP customers to run more test or project systems and to refresh the systems more often. This enables project teams to reduce project cycles by running parallel testing and improves the quality of testing and training with more actual data from production.

Capacity Requirements

When creating SAP system copies with most storage architectures, space must be allocated to accommodate the entire 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 is copied to a quality assurance (QA) 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 14 shows the traditional SAP system copy architecture.

Figure 14) Traditional SAP system copy.



In contrast, when using NetApp FlexClone technology to create SAP system copies, only a fraction of the storage space is required. FlexClone technology uses Snapshot copies, which are created in a few seconds without interrupting the operation on the source system, to perform SAP system copies. Because the data is not copied, but is referenced in place, the amount of storage required is limited to data that is changed at the source and the target system, and therefore, the disk space needed for SAP system copies is significantly decreased.

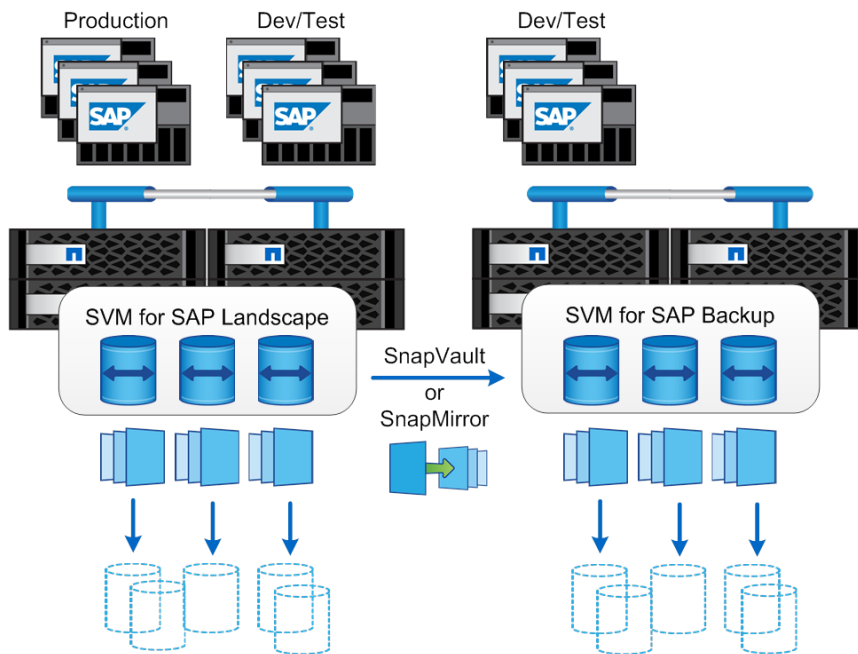
As a result, the capacity requirements for a system copy in a NetApp storage environment depend on the refresh cycle of the target systems. The longer test systems are kept, the more block changes happen from the source and the target system, and the more storage space is required. Storage requirements also depend on the number of copies that are made from the same source. Of course, more copies of the same source system result in higher storage savings.

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

The FlexClone copy can be created at the same storage system or at a secondary storage system.

The secondary storage system can already be in place and used as a disk-to-disk backup device or as a disaster recovery solution. The backup or disaster recovery replication images can be accessed for reading and writing by using FlexClone technology. Existing backup or disaster recovery images are used for test environments, turning expenses into assets. As a side effect, the backup and recovery or disaster recovery solution is tested without any additional effort and without any interruption. Figure 15 shows SAP system copy using the NetApp approach.

Figure 15) SAP system copy: NetApp approach.



Time Requirements

The time required to create an SAP system copy can be divided into three parts:


- Time to restore the backup to the target system
- Time to perform OS and database-specific postprocessing
- Time to perform SAP application postprocessing

Note: The SAP postprocessing depends on the customer's SAP environment. Some customers can finish the postprocessing in a few hours, while others need several days to accomplish this task.

In a conventional system-copy process, the data is backed up to tape and then restored, which takes a long 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. Because of the large number of logs that need to be applied, the time required to recover the database and make it consistent is greatly increased, possibly adding hours to the system copy process. If an offline backup is used, the source system is shut down, resulting in loss of productivity.


Figure 16 and Figure 17 show an example describing the difference between the amount of time spent creating an SAP system copy using a conventional approach versus with NetApp storage.

Figure 16) SAP system copy: standard approach.



All steps up to the point when the SAP system can be started on the target host can be accomplished in a few minutes by using the NetApp solution, compared to several hours with the standard approach. But in both cases, the SAP postprocessing must be done as an additional step.

Figure 17) SAP system copy: NetApp approach.

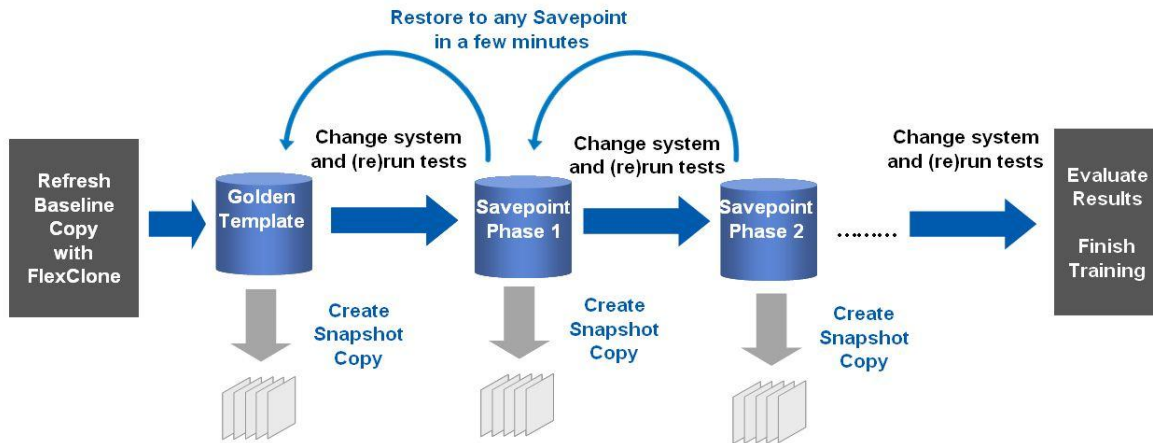


A key requirement to successfully managing an SAP environment is the ability to create copies of production data for use in testing, quality assurance, or training. NetApp Snapshot and FlexClone technologies allow the fast and space-efficient creation of SAP system copies.

5.2 SAP Testing Cycle

The possibility of easily creating backups in seconds and being able to restore the SAP system to a point in time of any available Snapshot copy is very helpful in SAP development and test environments. Projects such as data import, SAP upgrades, and installation of support packages can be accelerated by using fast backup and restore functionalities. During these projects, backups can be done at specific phases, and the system can easily and quickly be reset to a starting point, enabling that step to be repeated. Figure 18 shows the SAP testing cycle.

Figure 18) SAP testing cycle.



Carrying out an SAP upgrade or importing support packages and critical transports always involves SAP system downtime. It is important to keep this downtime to a minimum and to make sure that the previous system status can always be restored. The specified system changes are usually made in the development system first, to test the general functionality and procedures. In many cases, test systems must be upgraded several times, because problems can occur that can only be solved by restoring the system and restarting the upgrade. In this respect, NetApp Snapshot copies and FlexClone functionality can save a considerable amount of time. There is no need for a tape backup; a Snapshot copy can be

created instead. In the event of an error, the system can be quickly restored to its original status, and the upgrade can be repeated.

Time management is extremely important when the production system is upgraded, because the system is not available at various stages during the upgrade. Scheduling must also include time for restoring the system to its former release status. Depending on the size of the database and the time and effort required for the functional test and importing the transports for the modification adjustment, one weekend might not be sufficient for the upgrade. Alternatively, NetApp SnapCenter plug-in for SQL Server software offers Snapshot technology as a backup method and SnapRestore technology for restoring the system to its former release status quickly. This allows a higher level of flexibility with regard to scheduling. By creating several Snapshot copies at certain stages during the upgrade, it is possible to restart the upgrade without having to revert to the former release status.

6 Conclusion

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 clustered NetApp Data ONTAP infrastructure together with NetApp SnapCenter allows business owners and IT departments to work more efficiently and effectively toward the goal of improving enterprise business processes.

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

References

Best Practices Guides

Best Practices for Scalable SAN in Clustered Data ONTAP 8.3.1

<http://www.netapp.com/us/media/tr-4080.pdf>

NetApp Clustered Data ONTAP 8.3.x and 8.2.x: An Introduction

<http://www.netapp.com/us/media/tr-3982.pdf>

SAP

Global SAP Homepage

<http://www.sap.com/index.epx>

SAP Service Marketplace

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

SAP Developer Network

<http://sdn.sap.com/>

SAP Help Portal

<http://help.sap.com/>

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2015 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the Web at <http://www.netapp.com/us/legal/netapptmlist.aspx>. TR-4467-1115