White Paper

SAP Applications Built on NetApp

SAP Competency Center, NetApp
April 2015 | WP-7216

Abstract

This document provides an overview of solutions available for SAP on NetApp.

Starting with clustered Data ONTAP® as the basis for infrastructure solutions such as FlexPod®, this white paper lists important use cases, including backup and recovery, disaster recovery, and SAP system copies. In addition, we present an overview of integration into SAP Landscape Virtualization Manager.

TABLE OF CONTENTS

LIST OF FIGURES

1 Introduction

Companies 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. Typically, SAP business solution deployments are larger than a single production instance of SAP. Changing market conditions, restructuring activities, mergers, and acquisitions often result in the creation of new, complex SAP landscapes based on the SAP NetWeaver platform. New technologies such as SAP HANA and SAP S4Hana offer new business use cases, but they also create additional challenges at the infrastructure layer due to the changed behavior of these new technologies.

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 with current, up-to-date datasets as project schedules dictate 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 for most customers.

This document provides an overview of available NetApp® solutions for SAP environments.

1.1 Technology Challenges of an Expanding SAP Landscape

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

SAP recommends that customers maintain separate development and test instances for each production instance. In practice, the standard three-part SAP system (development, quality assurance, and production) is often expanded to include additional instances, such as sandbox and user training systems. It is also common practice to have multiple development instances and more than one system for quality assurance and testing. A final staging or performance system may also be needed prior to releasing SAP applications into production. Compound this list with the many different SAP applications or modules, such as enterprise resource planning (ERP), customer relation management (CRM), business warehouse (BW), supply chain management (SCM), supplier relation management (SRM), and the enterprise portal, and the number of systems requiring support can become very large, very quickly.

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 widely depending on the phase of a project and whether a project is focused around a new or existing SAP implementation. Projects typically rely on frequent refreshes of the nonproduction instances so that testing and training use the most up-to-date data. Frequent refreshes improve the quality and relevance of training and results to current business processes.

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

As SAP landscapes have expanded, the technology has also changed. SAP has evolved to take advantage of the latest technology, and new trends, such as in-memory database SAP HANA, have also emerged. In addition, 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 changing technology, IT organizations are unable to meet the business needs of the company.

1.2 NetApp Solutions for SAP

NetApp would like to minimize or eliminate many of the IT barriers associated with deploying new or improved business processes and applications. The combination of SAP solutions and a simplified and flexible NetApp clustered Data ONTAP infrastructure allows business owners and IT departments to improve enterprise business processes, regardless of which SAP technologies are deployed.

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

2 Storage Virtualization: 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 physical storage and network resources is the basis for scale-out and nondisruptive operations.

2.1 Hardware Support and Basic System Overview

As is shown in Figure 1, a clustered Data ONTAP system contains NetApp storage controllers. The basic building block is the high-availability (HA) pair, which is similar to the Data ONTAP 7G and 7-Mode environments used in the past. 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 nodes 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 by adding additional nodes of the same type or by adding 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 and more than 103PB of data storage.

Figure 1) Data ONTAP cluster overview.

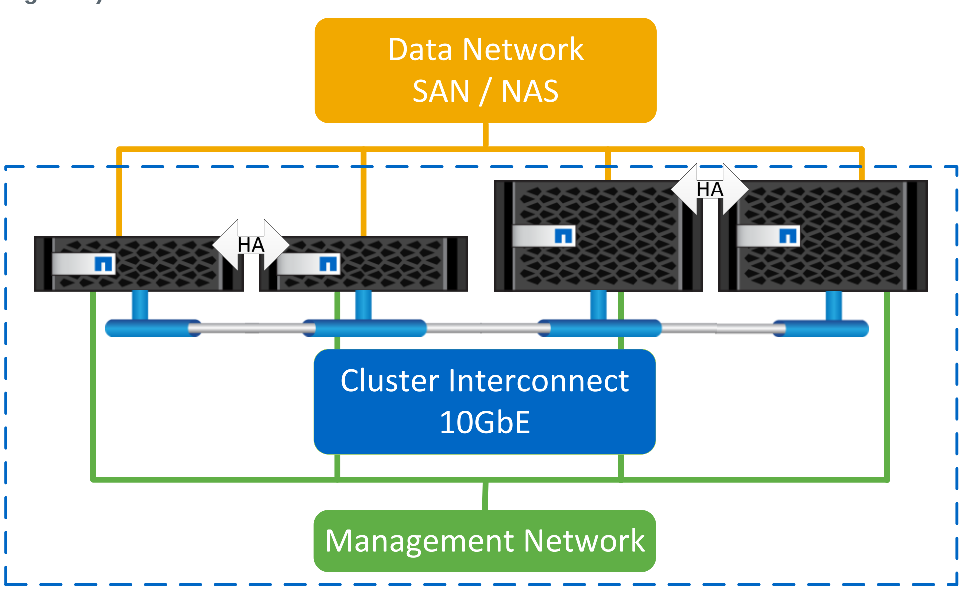

One of the key differentiators for a clustered Data ONTAP environment is that the storage nodes are combined into a cluster to form a shared pool of physical resources that are available to SAN hosts and NAS clients. This shared pool appears as a single system image for management purposes and provides a common point of management through either GUI or CLI tools for the entire cluster.

2.2 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 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.

Flash Cache™ cards can also be used to provide accelerated read performance for frequently accessed data. Starting with version 8.1.1, clustered Data ONTAP supports Flash Pool™ intelligent caching, which combines SSD with traditional hard drives for optimal performance and efficiency. The highly adaptable clustered Data ONTAP architecture is the key to delivering maximum on-demand flexibility for shared IT infrastructure and offers flexible options for performance, price, and capacity.

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.

2.3 Clustered Data ONTAP Networking

Figure 1 also shows the underlying network architecture of clustered Data ONTAP. Three networks are shown:

- **Cluster interconnect.** This private, dedicated, redundant network is used for communication between the cluster nodes and for DataMotion data migration within the cluster. Cluster-interconnect infrastructure for this network is provided with every clustered Data ONTAP configuration. This infrastructure takes the form of redundant, high-performance, high-throughput 10Gbs enterprise-class switch hardware in clusters of four or more nodes.

  Clusters of two nodes can optionally be configured without switches, with point-to-point connections used for the cluster interconnect. This configuration, available for the first time in clustered Data ONTAP 8.2, is known as a switchless cluster. This entry-level configuration provides all of the benefits of clustered Data ONTAP with a simpler infrastructure. Switchless clusters can be nondisruptively upgraded to include a switched cluster interconnect when the cluster grows beyond two nodes.
- **Management network.** All management traffic passes over this network. Management network switches can be included as part of a clustered Data ONTAP configuration, or customer-provided switches can be used. NetApp OnCommand® System Manager, OnCommand Unified Manager, and other NetApp applications are available for the management, configuration, and monitoring of clustered Data ONTAP systems. System Manager provides GUI management, including a number of easy-to-use wizards for common tasks. Unified Manager provides monitoring and alerts. A powerful CLI is included, and NetApp Manageability SDK is packaged and distributed in the Manage ONTAP® Software Developer's Kit.
- **Data networks.** These networks provide data access services over Ethernet to NAS clients or Fibre Channel to the SAN hosts. They are provided by the customer according to requirements and might also include connections to other clusters acting as volume replication targets for data protection.

2.4 Storage Efficiency and Data Protection

Storage efficiency built into clustered Data ONTAP offers substantial space savings, allowing more data to be stored at 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 up front, 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 may 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.

2.5 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.

2.6 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®, which is similar to 7G and 7-Mode. 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) that contain volumes and logical interfaces.

2.7 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.

Figure 2 shows a single SVM in a two-node cluster providing data services to SAN hosts and NAS clients. Each volume, shown by the orange circles, is provisioned within an aggregate on a cluster node, and the combination of all of the volumes constitutes the entire namespace or resource pool for LUNs. Volumes can be moved nondisruptively at any time from any aggregate to any other aggregate as required.

Delegated SVM administrators can provision volumes only in their own SVMs; these administrators have no visibility to any other SVM or even awareness that other SVMs exist. A delegated SVM administrator cannot perform volume moves around the cluster because this operation affects the capacity of aggregates shared by other SVMs. For this reason, only a cluster administrator can move volumes.

If SVM administration has been delegated, the cluster administrator must explicitly specify the aggregates available to the SVM administrator for provisioning volumes. This offers a mechanism whereby SVMs can provide different classes of service. For example, an SVM could be restricted to only using aggregates with SSD or SATA drives or only aggregates on a particular subset of controllers.

Figure 2) Cluster with a single SVM.

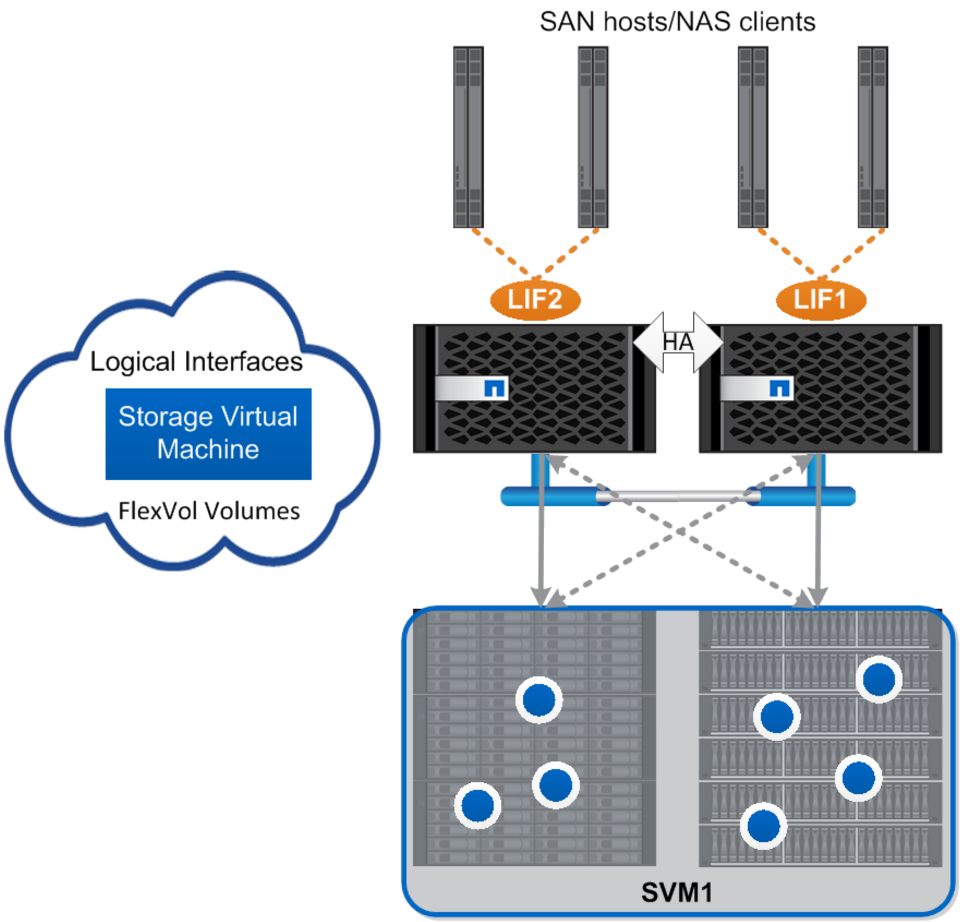

Figure 3 shows a more complex environment in which the cluster consists of four nodes with two SVMs providing data access. Each SVM consists of different volumes and LIFs for secure, compartmentalized access. Although the volumes and LIFs in each SVM share the same physical resources (network ports and storage aggregates), a host or client can only access the data in SVM1 through a LIF defined in that SVM. The same restriction is true for SVM2. Administrative controls make sure that a delegated administrator with access to SVM1 has visibility only to the logical resources assigned to that SVM, and an SVM2-delegated administrator similarly only sees SVM2 resources.

Figure 3) Cluster with multiple SVMs.

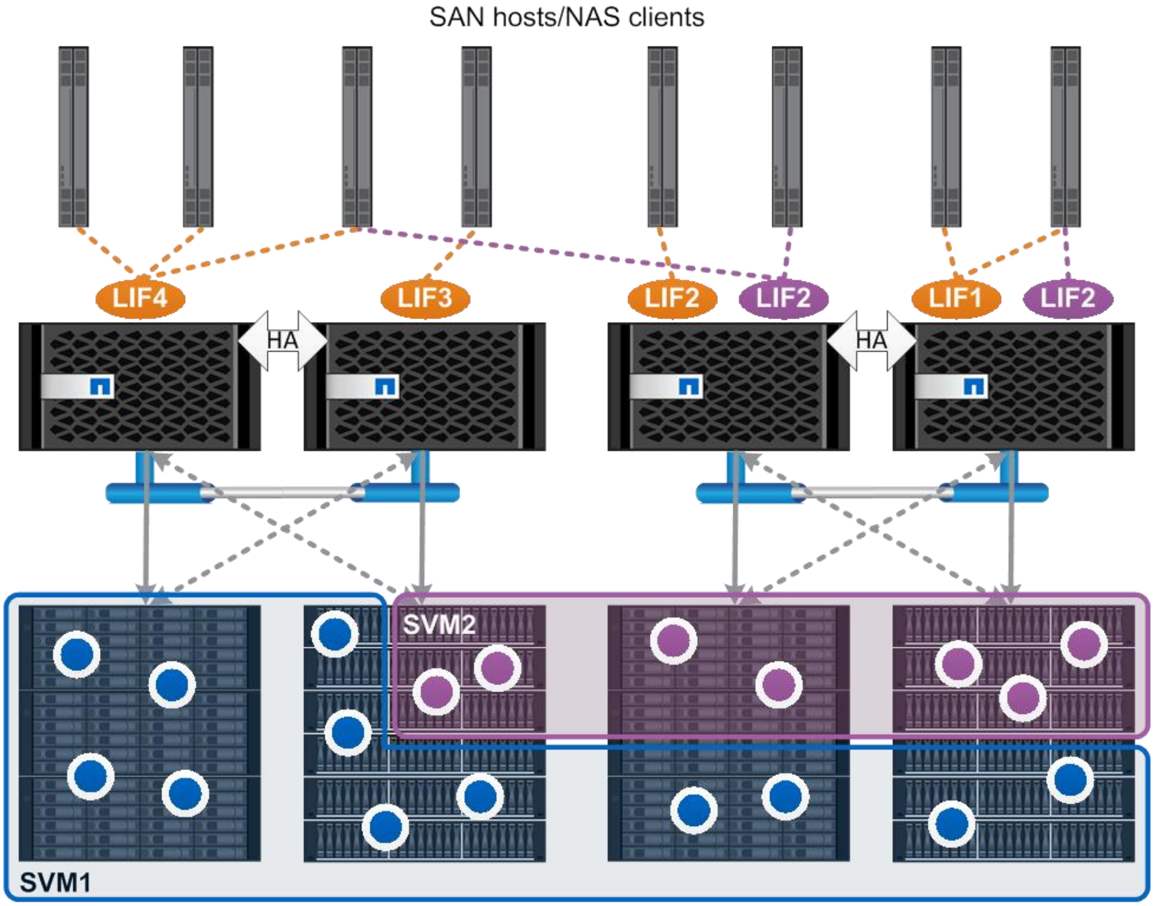

By virtualizing physical resources into the virtual server construct, Data ONTAP implements multi-tenancy and scale-out and allows a cluster to host many independent workloads and applications.

2.8 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.

3 Infrastructure Solutions: FlexPod

3.1 The Challenge: Disruptive Transitions from Inflexible Infrastructure Silos

IT departments of all sizes are increasingly challenged by data center complexity and inflexibility. Rapidly proliferating silos of server, storage, and networking resources combined with countless management tools and operational processes have led to crippling inefficiencies and costs.

Savvy organizations understand the financial and operational benefits of moving from infrastructure silos to a converged, virtualized environment. However, many organizations hesitate to make the transition because of potential short-term business disruptions and long-term architectural inflexibility that can impede scalability and responsiveness to business changes.

3.2 The Solution: Prevalidated Data Center Solutions to Simplify Transition to the Cloud

Enterprises and service providers alike need a tested, cost-effective data-center solution that:

- Supports both virtualized and nonvirtualized resources
- Can be easily implemented in large IT environments
- Can be managed within existing infrastructures
- Can scale without disruption to meet future cloud computing objectives

FlexPod, jointly developed by NetApp and Cisco, is a flexible infrastructure platform composed of presized storage, networking, and server components. FlexPod is designed to ease your IT transformation from virtualization to cloud computing with maximum efficiency and minimal risk.

FlexPod differs from other virtualization solutions by providing:

- Integrated, validated technologies from industry leaders and top-tier software partners
- A single platform built from unified compute, fabric, and storage technologies that lets you scale to large-scale data centers without architectural changes
- Centralized, simplified management of infrastructure resources, including end-to-end automation
- A choice of validated FlexPod management solutions from trusted partners leveraging our open APIs
- A flexible cooperative support model that resolves issues rapidly and covers new and legacy products

Investment Protection with Standardized, Flexible IT

Together, NetApp and Cisco offer a data center platform that is ready for virtualized environments today and yet is flexible enough to grow at your own pace to a fully private cloud. The FlexPod future-proof unified architecture runs multiple workloads on all protocols and fits right into your current infrastructure, leveraging existing resources and minimizing or eliminating technology replacement costs.

FlexPod components are integrated and standardized to help you achieve timely, repeatable, consistent deployments that eliminate guesswork and minimize risk. As a result, you can better predict the exact power, floor space, usable capacity, performance, and cost of each FlexPod deployment.

Scalability for Any Cloud Solution

FlexPod can be scaled up or down and then duplicated in modular fashion to fit your capacity needs. For example, large enterprises or service providers with mature IT processes and rapid growth expectations can deploy and scale out one or more FlexPod configurations to:

- Transition to a converged infrastructure with many applications
- Improve agility to meet growth and business initiatives
- Lower the cost per user without sacrificing scalability
- Reduce operational skill requirements, process complexity, and costs
- Deploy dedicated software as a service, virtual desktop infrastructure, and business-critical applications
- Securely separate multi-tenant environments for virtualized workloads alongside nonvirtualized ones

Business and Disaster Recovery

FlexPod can be configured with integrated data protection software to provide fast recovery from system, site, and regional outages for business continuity. The combination of NetApp MetroCluster™ and SnapMirror with Cisco Unified Computing System (Cisco UCS) Manager and Wide Area Application Services (WAAS) offers automated monitoring and failover as well as cost-effective replication to a secondary site for continuous protection against unplanned downtime. Our solution also lets you move virtual servers and storage resources and data nondisruptively across hardware to eliminate planned downtime.

Secure Multi-Tenancy and Secure Separation

FlexPod leverages the Cisco security architectures for the enterprise (SAFE) architecture and NetApp technology to deliver the industry's only secure multi-tenancy (SMT) architecture. SMT enables each tenant, an application, business unit, or customer, to be securely isolated within the FlexPod environment. SMT provides the data separation and service-level guarantees offered by application silos, while delivering the efficiencies of a converged, virtualized infrastructure. FlexPod also securely separates nonvirtualized workloads across the data center. SMT has undergone and passed several security compliance tests, including PCI compliance validation, FISMA certification and ICSA labs audit. For more details, refer to Secure Multi-Tenancy in Clustered Data ONTAP: Overview and Design Considerations.

Best-in-Class Components for Enhanced Data Center Efficiency

FlexPod includes the following components in a standardized configuration that supports a typical set of mixed-enterprise application workloads. This integrated solution can significantly reduce your capital and operating expenses through end-to-end virtualization and higher efficiencies at each layer.

The Cisco UCS is a data center platform that is designed to eliminate time-consuming manual configuration, reduce total cost of ownership, and increase business agility. Cisco UCS combines compute, networking, storage access, and virtualization into a scalable, modular system that is easily managed as a single entity by Cisco UCS Manager. This dramatically simplified architecture greatly reduces the number of devices that must be purchased, configured, managed, and secured. Service profile templates enable automatic, policy-based hardware configuration and deployment for large, stateless computing environments. The highly efficient Cisco UCS extended-memory technology also reduces memory requirements by up to 60%.

Cisco Nexus data center switches use award-winning unified fabric technology to identify and consolidate all network traffic onto a single simplified, cost-effective architecture. Cisco Nexus switches offer "zero-touch" installation, automatic configuration, and excellent scalability, including in-service upgrades. A single point of policy management also increases efficiency, availability, and security.

NetApp FAS systems reduce cost and complexity for virtualized infrastructures by meeting all of your storage requirements with a single, highly scalable solution. The NetApp unified storage platform supports all protocols so that you no longer must purchase separate systems to accommodate different storage needs. You can slash capacity use by up to 50% with built-in deduplication and thin provisioning as well as space-efficient backup and cloning. NetApp systems also enhance operational efficiency with automated storage management, data protection, and security and optimize performance with 10GbE or FCoE.

Choice of Management Solutions

NetApp and Cisco work with trusted partners to provide you with a choice of management solutions. The FlexPod architecture enables this flexibility by providing APIs at each layer, so FlexPod is easily integrated with a broad range of software solutions for end-to-end management. Validated FlexPod management solutions have been tested in NetApp and Cisco labs to verify that they deliver essential functionality. Together with our partners, we can provide automation and orchestration, monitoring and analytics, and configuration management, among other capabilities.

Rapid Resolution to Minimize Disruption

FlexPod includes a simplified, flexible, cooperative support model that is coordinated between NetApp, Cisco, and VMware. The joint support model offers global, 24/7 support with streamlined responses from technical experts, covers new and legacy products from each company for maximum flexibility, and also includes a cooperative support lab designed to replicate and rapidly resolve complex customer issues, leveraging onsite experts and state-of-the-art equipment from all three vendors.

NetApp and Cisco extend world-class technical support to a broad range of FlexPod customer environments by partnering with other hypervisor, application, and management vendors through the Technical Support Alliance Network (TSANet).

Proven Partnership

Industry leaders NetApp and Cisco have a powerful global presence and have been working together on a shared, virtualized data center vision since 2003. We are transforming the data center with unified architecture and virtualization technologies that create highly efficient virtualization and cloud computing solutions built on jointly validated reference architectures. Together, we have helped thousands of our mutual customers to increase efficiency and agility and reduce total cost of ownership.

Open Delivery Ecosystem

You can choose from a broad network of world-class solution-delivery partners to implement FlexPod. These partners understand your business requirements, and they are trained and certified on NetApp, Cisco, and complementary technologies to deliver a complete cloud solution that fits your business needs.

For more information, refer to FlexPod Datacenter with VMware vSphere 5.5 Update 1 and NetApp All-Flash FAS: NVA Deployment; FlexPod Datacenter Solutions Guide; and FlexPod Datacenter with Microsoft Private Cloud.

4 Business Continuance

4.1 Backup and Recovery

Corporations today require their SAP applications to be available 24 hours a day and 7 days a week. Consistent performance levels are expected, independent 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 effect on the production SAP system. Because backup windows are shrinking and the amount of data that must be backed up is increasing, it is difficult to define a point in time when backups can be performed with minimal effect on the business process. The time needed to restore and recover SAP systems is of particular concern because downtime for SAP production and nonproduction systems must be minimized.

This list summarizes SAP backup and recovery challenges:

- **Performance effects on production SAP systems.** Backups create a heavy load on the database server, the storage system, and the storage network and thus have a significant performance effect on the production SAP system.
- **Shrinking backup windows.** Conventional backups have a significant performance effect on the production SAP system. Indeed, 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 always active.

- **Rapid data growth.** Rapid data growth, together with shrinking backup windows, necessitate an ongoing investment in the backup infrastructure, including more tape drives, new tape-drive technology, and faster storage networks. Growing databases also need more tape media or disk space for backups. Incremental backups can address these issues, but they typically result in a very slow restore process, which is usually not acceptable.
- **Increasing cost of downtime.** Unplanned downtime for an SAP system always affects a business financially. A significant part of unplanned downtime is caused by restore and recovery processes for an SAP system following 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. Go/no-go decisions are based on the amount of time required to restore and recover the database from a previously created backup. 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 effect on the production SAP system because the Snapshot 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 course of a day (for example, 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 to our customers for restore and recovery operations. The NetApp SnapRestore® functionality allows the restoration of the entire database to the point in time when any available Snapshot copy was created. In combination with database redo or transactional logs, a recovery of the database to any point in time with an available backup is possible. The restore process can be completed in a few minutes, independent of the size of the database.

Because several online Snapshot backups were created during the course of the day, the time needed for the recovery process is also dramatically reduced. Because a restore can be performed with a Snapshot copy that is at most 8 hours old, fewer transaction logs must be applied. The mean time to recover (the time needed for restore and recovery) is therefore reduced to several minutes instead of several hours, as is the case 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, it is unlikely that these backups are 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 one must restore a backup that is no longer available from a Snapshot copy (for example, a two-week-old backup).

Figure 4) Backup solution overview.

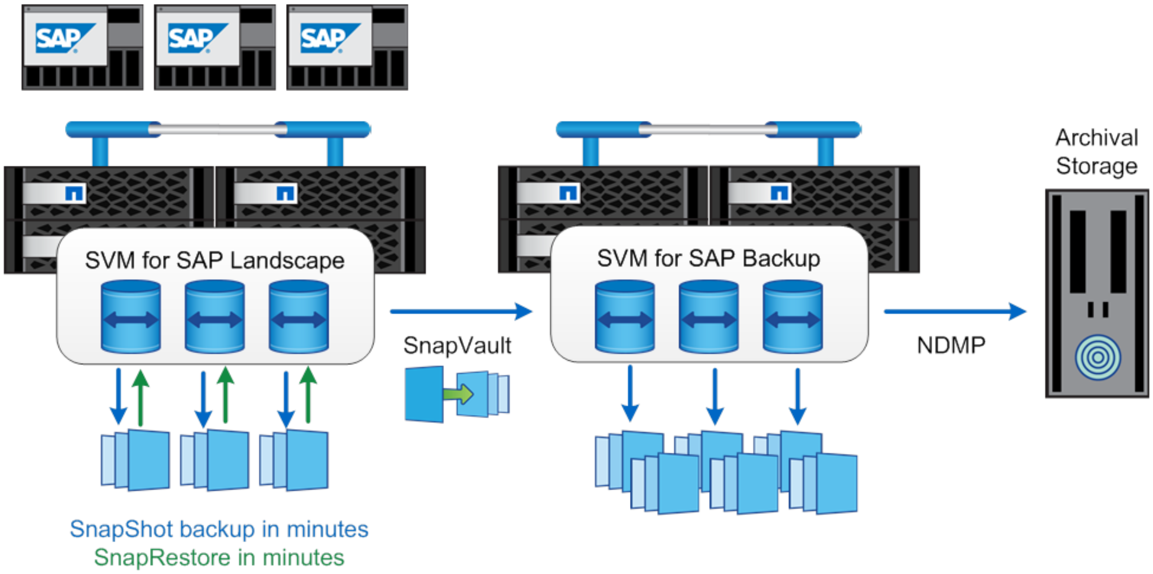

As is shown in Figure 4, a backup and recovery solution using a NetApp storage system should always consists of two parts:

- Backup and restore by using Snapshot and SnapRestore technologies
- Backup to and restore 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 for backing up data to a secondary location:

- **Disk-to-disk backup using SnapVault software.** The primary SVM communicates directly with the secondary SVM 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 must 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 if traditional database systems are used. 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. For example, this could be a monthly backup that is kept for a year. In this case the tape infrastructure can be connected directly to the secondary SVM, 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, and data is written to tape using NDMP.

Figure 5 compares the different backup approaches with regard to the performance effect of a backup and the time in which the database must be in hot backup mode or offline.

Figure 5) Comparison of time required for different backup methods.

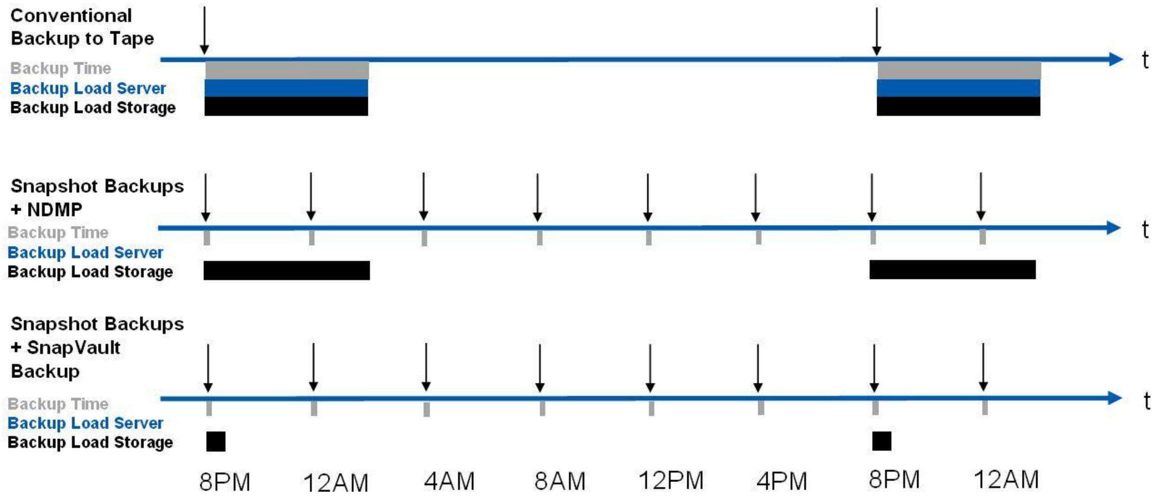

Snapshot Backups Together with NDMP Backups

Snapshot backups do not generate any load on the database server or the primary storage system. A full database backup based on a Snapshot copy consumes disk space only for changed blocks. Snapshot 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. In the case in which a tape backup is required, the NDMP backup can be performed at a secondary storage system to a SnapVault destination, as is shown in Figure 4. This approach would remove the heavy load of the NDMP backup from the primary storage system.

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 when compared with a full tape backup.

Figure 6 compares the time required to perform restore and recovery using different methods.

Figure 6) Comparison of time needed for restore and recovery.

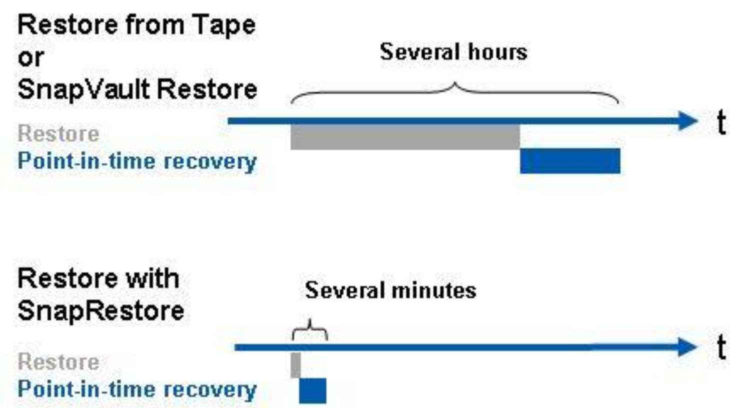

Restore from Tape or SnapVault Restore

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, some number of transaction logs must 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. Therefore, the forward recovery time is much faster, because fewer transaction logs must 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 copies and SnapRestore with disk-to-disk backup based on SnapVault offers significant improvements over conventional tape backups:

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

Database Verification

Independent database verification is an important part of any company's backup strategy. Snapshot backups are excellent for running database consistency checks. NetApp SnapManager® software 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.

Depending on the database used, different NetApp products provide the functionality described previously. For example, NetApp SnapManager for MS SQL Server is used with Microsoft SQL Server, while the NetApp Snap Creator® framework is used for SAP HANA.

For further information, refer to [SAP HANA Backup and Recovery Using Snap Creator;]() [SAP with Microsoft SQL Server on Windows: Best Practices Using NetApp Clustered Data ONTAP](); and [SAP with Oracle on UNIX and NFS with NetApp Clustered Data ONTAP]().

4.2 SAP Repair System

Companies are increasingly confronting logical errors within complex SAP environments, in which 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 just before the logical error occurred. This approach has the following disadvantages:

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

Therefore, SAP customers are looking for a more efficient and flexible solution to address logical errors. NetApp Snapshot and FlexClone technologies provide solutions that allow for the recovery from logical errors without restore and recovery of the affected system.

Figure 7) SAP repair system.

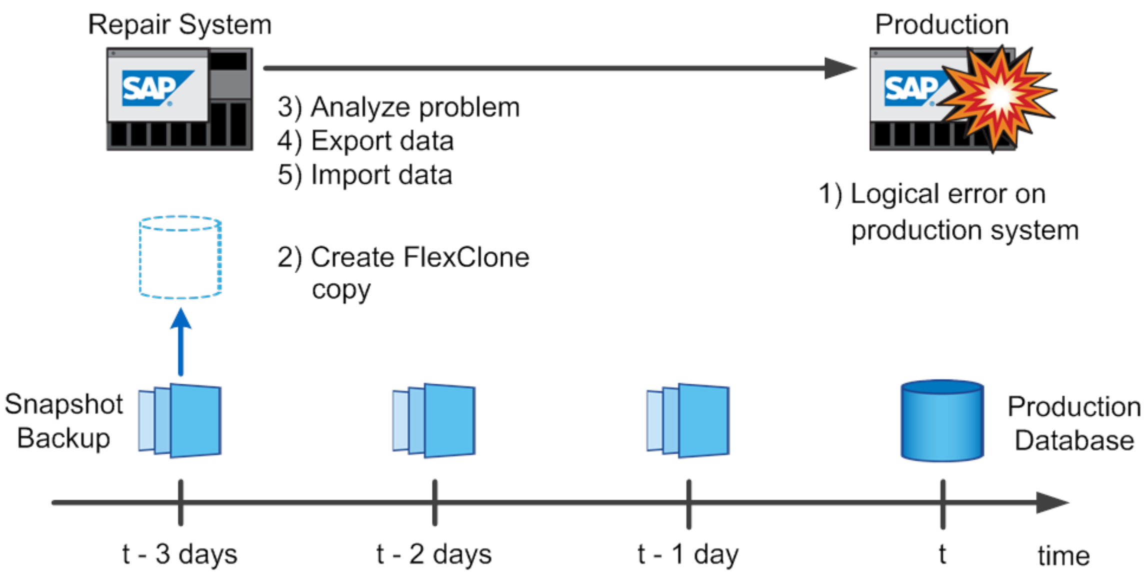

Figure 7 shows the general process of creating and using the repair system.

1. A logical error is discovered on the production system. Depending on the type of logical error, the production system might be shut down or, alternatively, kept online so that some but not all business processes are affected.
2. Several Snapshot backups of the production system are available, and any of them can be used to create a 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 effect 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 that 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 logical errors.

Depending on the database used, different NetApp products provide the functionality described previously. For example, NetApp SnapManager for MS SQL Server is used with Microsoft SQL Server, while the NetApp Snap Creator framework is used for SAP HANA.

For further information, refer to SAP HANA Backup and Recovery Using Snap Creator; SAP with Microsoft SQL Server on Windows - Best Practices Using NetApp Clustered Data ONTAP; and SAP with Oracle on UNIX and NFS with NetApp Clustered Data ONTAP.

4.3 Disaster Recovery

Almost all organizations recognize the importance of having a business continuance plan in place to deal with a disaster. The costs of not having such a plan—loss of productivity, revenue, and customer loyalty and possibly even business failure—make such a plan mandatory for minimum downtime and rapid recovery from a disaster.

A disaster recovery solution based on SnapMirror fulfills and exceeds all business requirements needed for even global SAP systems. By replicating data at high speeds over a LAN or WAN, SnapMirror software provides high data availability and fast recovery. SnapMirror technology mirrors data to one or more SVMs. It updates mirrored data to keep it current and available for disaster recovery, tape backup, read-only data distribution, testing, online data migration, and more.

SnapMirror performs an initial transfer to initialize the disaster recovery site. Incremental changes are then passed to the disaster recovery site asynchronously. The SnapMirror disaster recovery solution is based on NetApp backup and recovery; Snapshot backups are mirrored to the disaster recovery site. Additionally, the volumes where the log files and the log-file backups are stored are mirrored with SnapMirror. The frequency of SnapMirror updates for the log files and backups determines the quantity of data lost in a disaster.

Figure 8) Disaster recovery with SnapMirror.

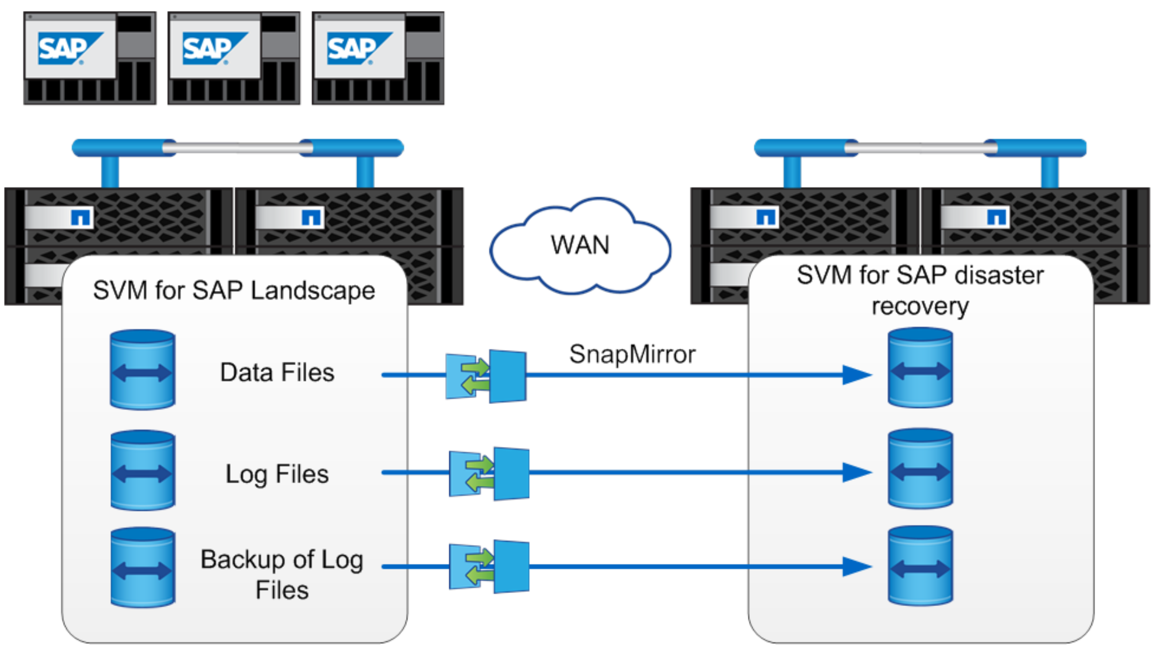

For further information, refer to SAP HANA Disaster Recovery with Asynchronous Storage Replication Using Snap Creator and SnapMirror; SAP with Microsoft SQL Server on Windows: Best Practices Using NetApp Clustered Data ONTAP; and SAP with Oracle on UNIX and NFS with NetApp Clustered Data ONTAP.

5 System Management and Maintenance

5.1 SAP System Copy

Business Challenges

A typical SAP customer environment consists of different component for SAP Suite on HANA and SAP NetWeaver. 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 requires approximately 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 for NetWeaver-based systems 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 to restore the data to the three target systems.

Figure 9) Traditional SAP system copy.

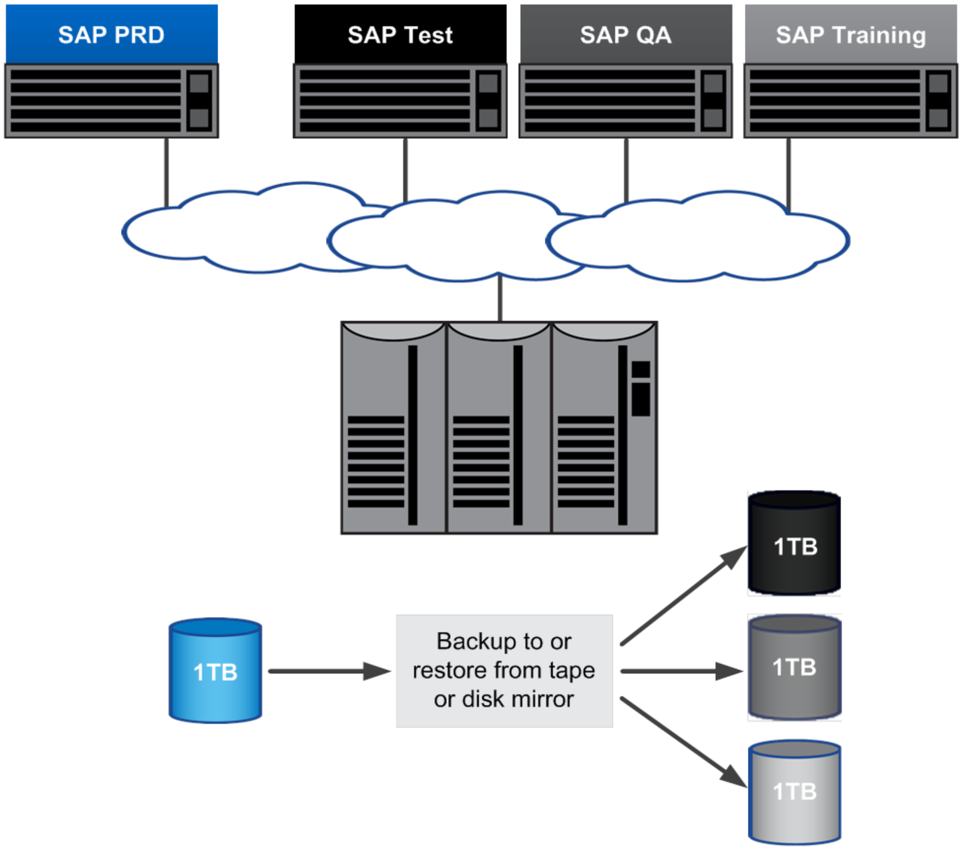

In contrast, when using NetApp FlexClone technology to create SAP system copies, only a fraction of the original storage space is required. FlexClone technology uses Snapshot copies, which are created in a few seconds without interrupting the operation on the source system, to perform SAP system copies.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

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.

The FlexClone copy can be created in the same storage system or in a secondary storage system. The secondary storage system can already be in place and used as a disk-to-disk backup device or a disaster recovery solution. The backup or disaster recovery replication images can be accessed for reading and writing with FlexClone technology. Existing backup or disaster recovery images can be 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 10) SAP system copy: NetApp approach.

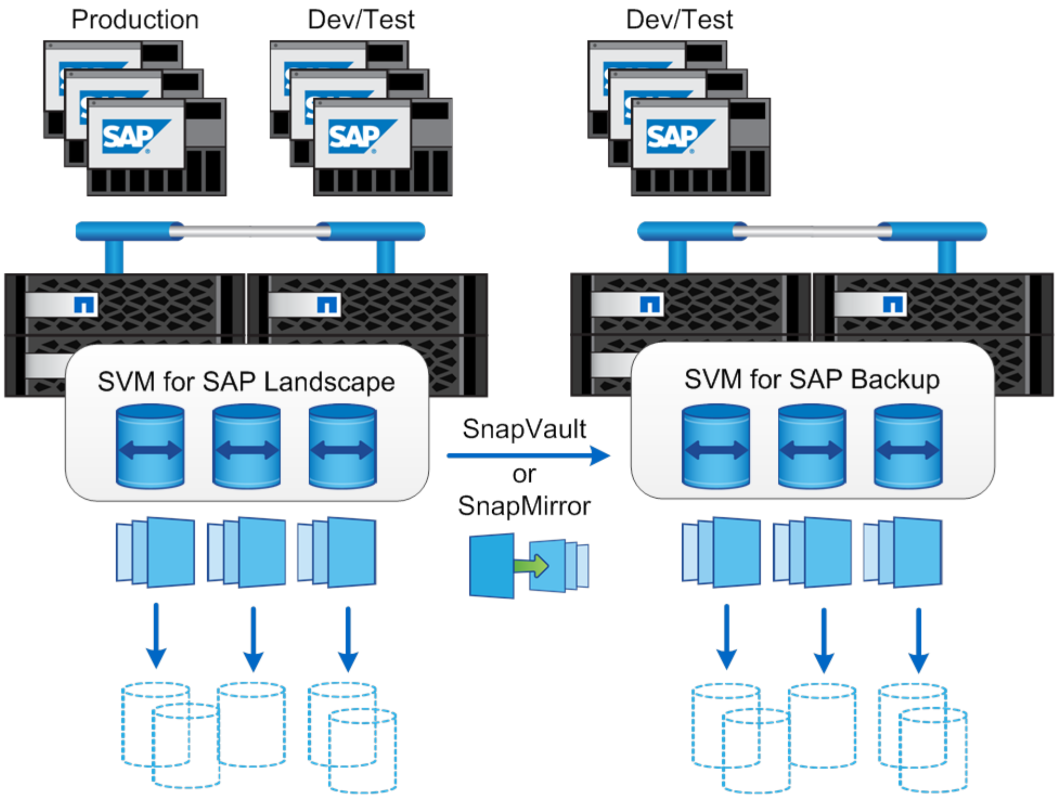

Time Requirements

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

- The time to restore the backup to the target system.
- The time to perform OS and database-specific postprocessing.
- The time to perform SAP application postprocessing. 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 effect on the source system during backup. Because of the large number of logs that must 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 a loss of productivity.

Figure 11 and Figure 12 provide examples depicting the difference between the amount of time spent creating an SAP system copy with NetApp storage versus the time spent using a conventional approach.

Figure 11) SAP system copy: timeline for standard approach.

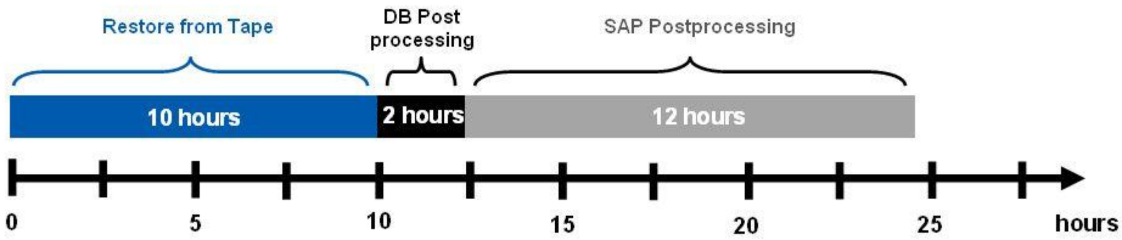

All steps up to the point when the SAP system can be started on the target host can be accomplished in a few minutes in the NetApp solution versus several hours in the standard approach. In each case, SAP postprocessing must be done as an additional step.

Figure 12) SAP system copy: timeline for NetApp approach.

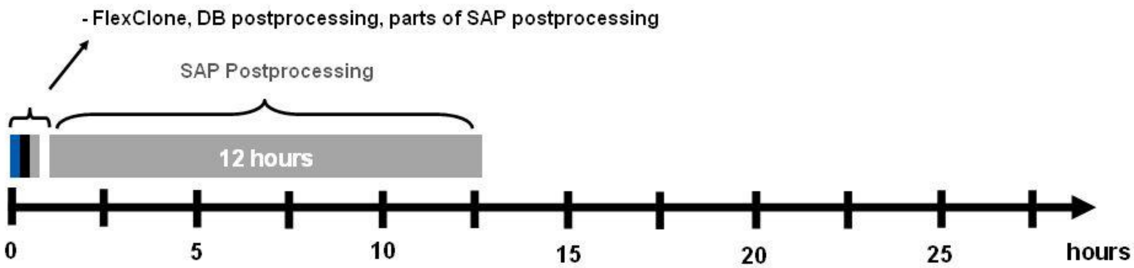

A key requirement for successfully managing an SAP environment is the ability to create copies of production data for use in testing, quality assurance, and training. NetApp Snapshot and FlexClone technologies allow for the fast, space-efficient creation of SAP system copies.

Depending on the database used, different NetApp products provide the functionality described previously. For example, NetApp SnapManager for MS SQL Server is used with Microsoft SQL Server, while the NetApp Snap Creator framework is used for SAP HANA.

For further information, refer to SAP HANA Backup and Recovery Using Snap Creator; SAP with Microsoft SQL Server on Windows: Best Practices Using NetApp Clustered Data ONTAP; and SAP with Oracle on UNIX and NFS with NetApp Clustered Data ONTAP.

5.2 SAP Testing Cycle

Easily creating backups in seconds and restoring the SAP system to the point in time of any available Snapshot copy are 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 the repeat of that phase.

Figure 13) SAP testing cycle.

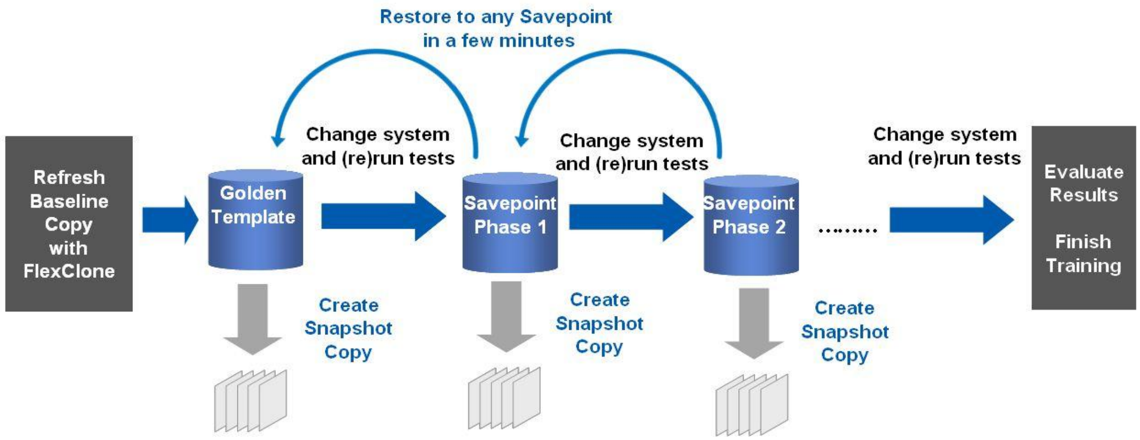

Carrying out an SAP upgrade or importing support packages and critical transports always involves SAP system downtime. It is important to keep this downtime to a minimum and make sure that the previous status can be restored. The specified system changes are usually made in the development system first to test 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 of 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. On the other hand, NetApp SnapManager software offers Snapshot copies as the backup method and SnapRestore to quickly restore the system to its former release status. This allows a higher level of scheduling flexibility. By creating several Snapshot copies at certain stages during the upgrade, one can restart the upgrade without having to revert to the former release status.

5.3 SAP Landscape Virtualization Manager

SAP Landscape Virtualization Manager (LVM) has a long history of enabling SAP system administrators to automate SAP system operations, including end-to-end, SAP-system clone, copy, and refresh operations. SAP LVM allows infrastructure providers such as NetApp to integrate their products so that customers can use these added functions in the SAP LVM management GUI.

For many years, NetApp has offered the NetApp Storage Services Connector (SSC), which allows LVM to directly access NetApp technologies such as NetApp FlexClone and SnapMirror to minimize the storage required for SAP systems and shorten the time required to provide SAP system clones (without system identifier (SID) change) and copies (with SID change).

By leveraging the rich set of extensibility features in SAP LVM, NetApp customers can directly integrate storage-based backups as a customer-defined system management function or use the NetApp HANA plugin to automate snapshot-based HANA backups as a mass operations directly out of the LVM.

These features are accessible for customers when running their own on-premises data center or their own private clouds. They are also available in a hybrid cloud solution that integrates public cloud providers such as Amazon Web Services into the overall data-center concept. SAP LVM, together with NetApp SSC, bridges the gap between on-premises and cloud resources by defining clear data ownership and providing tools for moving systems seamlessly between clouds.

System Architecture

From an administrator's perspective, SAP LVM is the central tool that communicates with the storage and virtualization layers. Apart from creating SAP system copies and clones, LVM acts as a management component that provides extensive monitoring capabilities and controls for interaction with SAP systems (start, stop, relocate, and so on). For more information about LVM, refer to SAP Landscape Virtualization Management.

Architecture Overview

SAP LVM can be used to manage SAP systems running on any kind of infrastructure that supports running SAP applications. This includes:

- Standard physical servers in an on-premise data center
- Cloudlike infrastructure managed through converged systems such as FlexPod
- Virtual environments such as VMware, Hyper-V, and Linux KVM
- Cloud infrastructures such as Amazon Web Services or Microsoft Azure

Figure 14 shows a typical on-premises data center setup. SAP LVM can integrate any SAP system, including classical NetWeaver-based SAP systems as well as SAP HANA running on supported operating systems such as Windows, Linux, and UNIX.

Figure 14) SAP LVM system landscape.

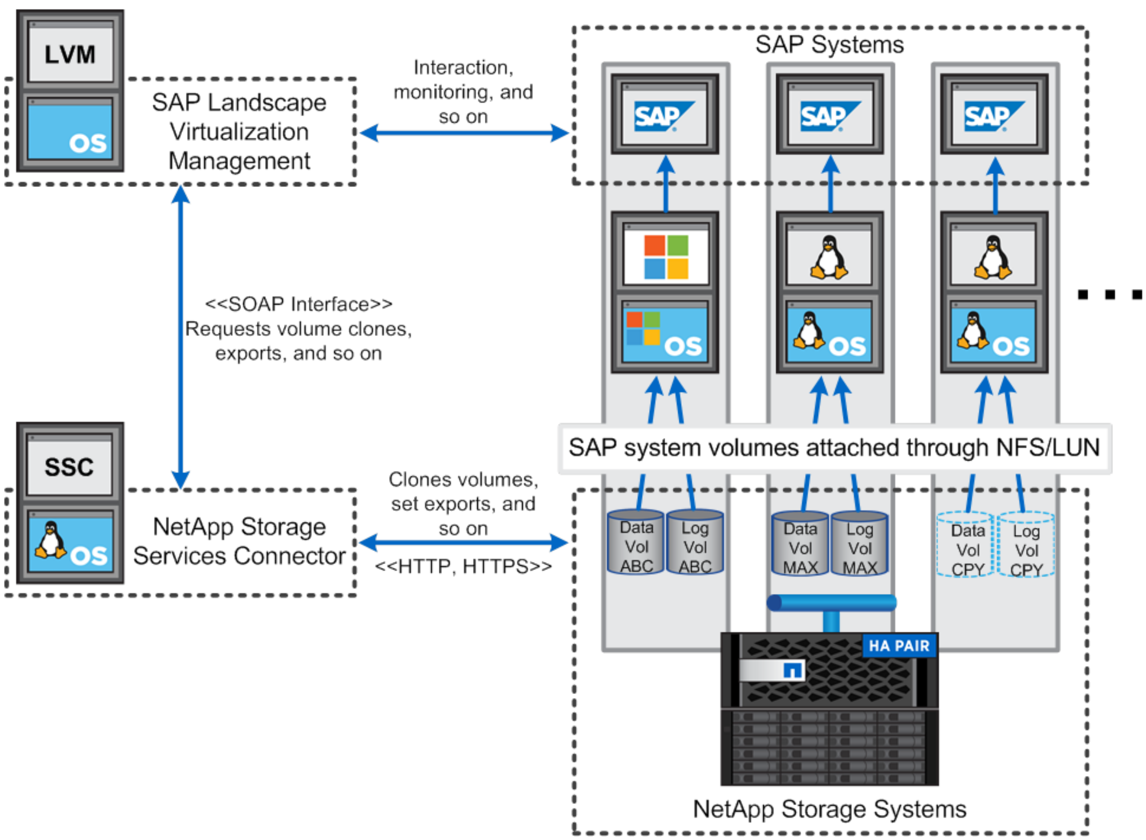

NetApp Storage Services Connector

The NetApp SSC is the central technical component that enables communication between SAP LVM and NetApp storage controllers. SSC implements the SAP LVM storage adpater API and allows SAP administrators to leverage NetApp storage efficiency tools when using SAP workflows to create volume clones or supply exports. Communication between SAP LVM and SSC is established through a SOAP interface. It uses well-defined XML objects to transmit messages, thereby providing transparency and interoperability.

Adaptive Computing Enabled

Adaptive computing is a technology SAP has developed to separate the OS from the SAP installation, making it possible to move an installed SAP system from one host to another. This function enables an administrator to:

- Change the host type (CPU or memory)
- Patch the OS without disturbing the running SAP system
- Move from a virtualized host to a physical host or vice versa

All of these functions can be performed without changing the SAP configuration or reinstalling the system. These options dramatically improve the manageability of an SAP system and enable customers to easily move from a physical to a virtual environment or even from on-premises to cloudlike infrastructure.

Figure 15) Adaptive design principles.

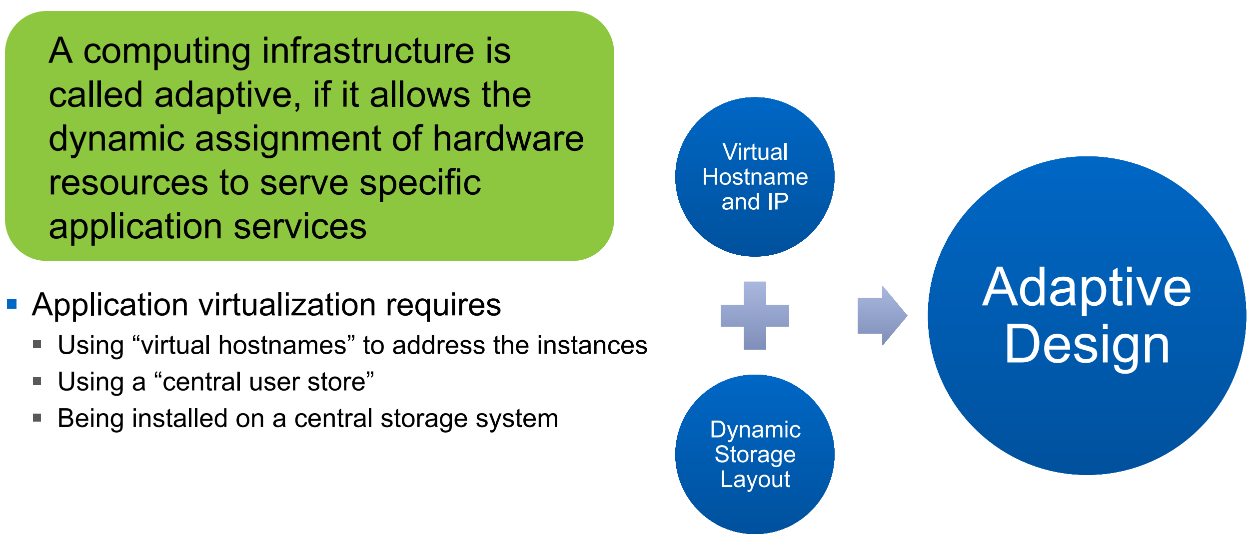

The difference between an adaptive computing environment and a standard installation can be found in the OS changes that occur during an SAP installation. When an SAP system is installed on a host, not only are users, groups, services, and file systems created, but many configuration options, such as the host name and IP, are stored in configuration files and SAP’s own database. If an SAP system is not installed according to adaptive design principles, it becomes very difficult to move the system from the original installation host to another more powerful host or to change the host for OS maintenance and patching purposes.

An adaptive-enabled installation has the following characteristics:

- Centralized user management such as Lightweight Directory Access Protocol (LDAP) or Microsoft Active Directory
- Central storage instead of local disks. In other words, all SAP specifics must be separated from the host file system.

- Use of virtual host names to address the instances. Instead of installing a SAP system onto a host using a physical host name and IP address, every SAP service has its own set of virtual hostnames and IP addresses.

If an SAP application is relocated from one host to another, SAP storage is moved (unmounted and mounted) to the new host, and the virtual IP addresses are also relocated. SAP client traffic still connects to the relocated virtual host names and IP addresses, and, from a networking perspective, the relocated SAP system remains unchanged.

While it is possible to install an SAP system without following adaptive design principles, this option should be considered.

SAP LVM Extensibility

Over the years, SAP has included many options for extending the built-in functionality of SAP LVM with customer-specific functions, and SAP even allows customers to add additional options to the LVM GUI. The following list shows some of the available extensibility options:

- Custom provisioning
- Custom operations and hooks
- Custom services
- External interfaces
- Custom links and tabs
- Custom notifications
- Custom validation

NetApp customers can implement the following features:

- Snapshot copy-based SAP backups triggered directly by SAP LVM
- Custom provisioning
- Improvements based on NetApp Workflow Automation (WFA) to further automate SAP management processes

Many of the additional functions, once implemented, can be used within the standard automation and mass operations of SAP LVM. For example, backup integration can be used to plan and schedule backups for all SAP systems using NetApp Snapshot software and have all logs propagated and available within SAP LVM.

For further information, refer to Integrating NetApp Storage with SAP NetWeaver Landscape Virtualization Management.

6 References

Best Practices Guides

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

http://www.netapp.com/us/media/tr-4160.pdf

Secure Multi- Tenancy in Clustered Data ONTAP: Overview and Design Considerations

http://www.netapp.com/us/media/tr-4160.pdf

FlexPod Datacenter with VMware vSphere 5.5 Update 1 and NetApp All-Flash FAS: NVA Deployment

http://www.netapp.com/us/media/nva-0011-netapp-all-flash.pdf

FlexPod Datacenter Solutions Guide

http://www.netapp.com/us/media/tr-3884.pdf

FlexPod Datacenter with Microsoft Private Cloud

http://www.netapp.com/us/media/nva-0010-flexpod-datacenter.pdf

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

http://www.netapp.com/us/media/tr-4251.pdf

SAP with Oracle on UNIX and NFS with NetApp Clustered Data ONTAP

http://www.netapp.com/us/media/sap-with-oracle-unix-nfs-with-netapp-clustered-data-ontap.pdf

SAP HANA on NetApp FAS Systems with Fibre Channel Protocol: Configuration Guide

http://www.netapp.com/us/media/tr-4384.pdf

SAP HANA on NetApp FAS Systems with NFS: Configuration Guide

http://www.netapp.com/us/media/tr-4290.pdf

SAP HANA on VMware vSphere and NetApp FAS Systems: Reference Architecture

http://www.netapp.com/us/media/tr-4338.pdf

SAP HANA Backup and Recovery Using Snap Creator

http://www.netapp.com/us/media/tr-4313.pdf

SAP HANA Disaster Recovery with Asynchronous Storage Replication Using Snap Creator and SnapMirror

http://www.netapp.com/us/media/tr-4279.pdf

Integrating NetApp Storage with SAP NetWeaver Landscape Virtualization Management

http://www.netapp.com/us/media/tr-4018.pdf

SAP

SAP Landscape Virtualization Management

http://www.sap.com/pc/tech/cloud/software/virtualization/index.html

Global SAP Homepage

http://go.sap.com/index.html

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.