SAP on VMware Availability and Disaster Recovery Guide

TECHNICAL MARKETING DOCUMENTATION

V 1.0/AUGUST 2015/MOHAN POTHERI, VAS MITRA, ERIK RIEGER

vmware®

Table of Contents

Availability Considerations for Virtualizing SAP Applications

SAP provides a range of enterprise software applications and business solutions to manage and run the complete business of a company. These mission-critical systems require continuous availability. SAP NetWeaver–based products have a scalable, fault-tolerant, multitier architecture, components of which can be protected either by horizontal scalability—additional SAP NetWeaver Application Servers, for example—or by cluster and switchover solutions that protect the single points of failure in the SAP NetWeaver architecture. They include the database, message, and locking services. The latter two are included in constructs referred to as the central instance or SAP central services. The central instance is an older construct; in new releases, it is replaced with central services and the primary application server (PAS).

A correctly architected and highly available solution can provide applications such as those from SAP with acceptable operational uptime by countering the impact of unplanned downtime. Although downtime can also be planned for maintenance and patching activities, the unplanned outages have the greatest effect on production uptime. This document covers high-availability solutions of virtualized SAP solutions on VMware, to protect against unplanned downtime in the same data center, and disaster recovery (DR) for entire site failure. Unplanned downtime refers to an outage in system availability due to infrastructure or software failure—server, storage, network, software, or operating system (OS) crash, and so on—or an entire site disaster. SAP products and solutions provide mission-critical business processes that must be highly available even in the event of a site disaster.

Single Point of Failure (SPOF) Analysis

The following SPOFs exist in the SAP NetWeaver architecture:

- Database – Every application work process makes a private connection to the database at the start. If the connection is interrupted due to database instance failure, the work process attempts to set up a new connection and changes to "database reconnect" state until the database instance connection is reestablished. In-progress user sessions with database activity receive SQL error messages, but logged-in sessions are preserved on the application server.
- SAP Message Service – The Message Service is used to exchange and regulate messages between SAP instances in an SAP network. It manages functions such as determining which instance a user logs in to during client connect and scheduling of batch jobs on instances configured for batch.
- SAP Enqueue Service – The Enqueue Service manages the locking of business objects at the SAP transaction level. Locks are set in a lock table stored in the shared memory of the host on which the Enqueue Service runs. Failure of this service has a considerable effect on the system because all transactions that contain locks must be rolled back and any SAP updates being processed fail.

The following SAP architectural components are defined based upon the Message Service and Enqueue Service:

- Central instance (CI) – CI comprises Message Service and Enqueue Service in addition to other SAP work processes that enable execution of online and batch workloads. In later NetWeaver releases, the CI is replaced with SAP central services and the PAS.
- SAP central services – In newer versions of SAP, the message and enqueue processes are grouped into a standalone service. Separate central services exist for ABAP- and JAVA-based NetWeaver Application Servers. For ABAP variants, it is called ABAP SAP central services (ASCS). It is now recommended to install central services instead of the classical central instance.
- Primary application server – This is an SAP application server instance that is installed with central services in newer NetWeaver releases.

- Replicated Enqueue – This component consists of the standalone enqueue server and an enqueue replication server. The replicated enqueue server runs on another host and contains a replica of the lock table (replication table). If the standalone enqueue server fails, it must be restarted on the host on which the replication server is running, because this host contains the replication table in a shared memory segment. The restarted enqueue server uses this shared memory segment to generate the new lock table, after which this shared memory segment is deleted.

The isolation of the Message Service and Enqueue Service from the CI helps address the high-availability requirements of these SPOFs. The central services component is "lighter" than the CI and is much quicker to start up after a failure.

SAP Application Local-Site High-Availability Solutions and Configurations

There are multiple solutions for high availability. Many of these options can be combined to provide different levels of availability. We will look at all available solutions and discuss their configuration, pros and cons, and applicability to SAP application infrastructure.

VMware vSphere High Availability

VMware provides VMware vSphere® products with built-in and optional high-availability and DR solutions to protect a virtualized SAP HANA system at all levels.

The power behind vSphere high-availability and DR solutions is in how they are layered to protect against failures at every level of the data center—from individual components, such as network adapters and HBA card teaming, all the way up to the entire site. With VMware vSphere Replication™, the solutions provide protection against both planned and unplanned downtime.

Figure 1 shows the various solutions to protect against both component-level and complete site failures.

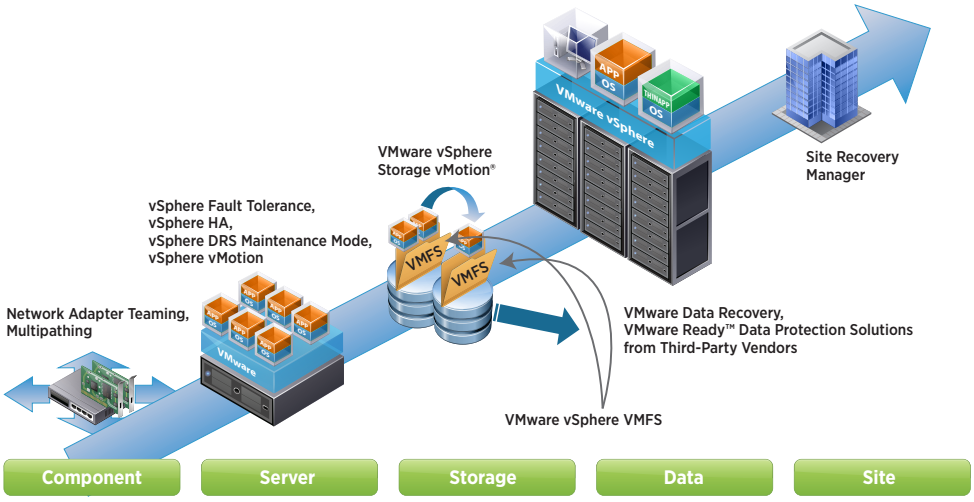

Figure 1. vSphere High-Availability and Disaster Recovery Solutions: Protection at Every Level

Many of the key features of virtualization, such as encapsulation and hardware independence, already offer inherent protection. Additional protection throughout the vSphere platform is provided to ensure that organizations can meet their availability requirements. The following features are provided:

- Protection against hardware failures
- Planned maintenance with zero downtime
- Protection against unplanned downtime and disasters

VMware vSphere High Availability (vSphere HA) delivers the availability required by most applications running in virtual machines (VMs), independent of the OS or applications running on them. vSphere HA provides uniform, cost-effective failover protection against hardware and OS outages within a virtualized IT environment through the following features:

- Monitors vSphere hosts and VMs to detect hardware and guest OS failures
- Restarts the VM on the same or another host in the vSphere cluster
- Restarts the VM, without any dependencies on the applications running on the VM or on other vSphere hosts in the cluster without manual intervention when a server outage is detected
- Reduces SAP application downtime by automatically restarting VMs upon detection of an OS failure

vSphere HA is the time-tested high-availability solution from VMware that is widely used for production environments. If configured with no SPOFs, it can protect the workloads from hardware failures. In the event of a hardware failure, the protected workload is automatically restarted in the remaining nodes of the cluster. There is an outage for the workload for the duration of the detection of the failure and the restart of the VM and the application. vSphere HA is very easy to set up and manage and is the simplest high-availability solution available for protecting virtual workloads.

Protecting SAP Applications with vSphere HA

VMware recommends leveraging vSphere HA and VMware vSphere Distributed Resource Scheduler™ (vSphere DRS) if licensing permits.

Prerequisites

- Verify that all the SAP application components and the configuration files reside on shared storage.
- Verify that the hosts are configured to access the shared storage, so VMs can be powered on by using different hosts in the cluster.
- Verify that hosts are configured to have access to the VM network.
- Verify that redundant management network connections are being used for vSphere HA.
- Verify that hosts are configured with at least two datastores, to provide redundancy for vSphere HA datastore heartbeating.
- Connect VMware vSphere Web Client to VMware vCenter Server™ using an account with cluster administrator permissions.

Procedure for New Cluster

1. In vSphere Web Client, browse to the data center where the cluster is to reside. Click **Create a Cluster**.
2. Complete the New Cluster Wizard. Do not turn on vSphere HA or vSphere DRS.
3. Click **OK** to close the wizard and create an empty cluster.
4. Based on the plan for resources and networking architecture of the cluster, use vSphere Web Client to add hosts to the cluster.

5. Browse to the cluster and enable vSphere HA.
   a. Click the **Manage** tab and click **Settings**.
   b. Select **vSphere HA** and click **Edit**.
   c. Select **Turn ON vSphere HA**.
6. Select **Host Monitoring**. This enables hosts in the cluster to exchange network heartbeats and enables vSphere HA to take action when it detects failures.
7. Choose a setting for **Virtual Machine Monitoring**. Select **VM Monitoring Only** to restart individual VMs if their heartbeats are not received within a set time. You can also select **VM and Application Monitoring** to enable application monitoring.
8. Click **OK**.

vSphere HA is now turned on and can help reduce downtime for all SAP application components during hardware failures. As part of the configuration, enable **Host Hardware Monitoring** with protection against storage connectivity loss. This helps in situations such as All Paths Down (APD) and Permanent Device Loss (PDL). In the event of a storage failure, the VM is restarted on a healthy host. For additional details, refer to the *vSphere 6.0 Availability Guide*.

vSphere HA and vSphere DRS Affinity Rules

vSphere HA failover must be specified for the following two types of rules:

- VM antiaffinity rules force specified VMs to remain apart during failover actions.
- VM host affinity rules place specified VMs on a particular host or a member of a defined group of hosts during failover actions.

When editing a vSphere DRS affinity rule, select the checkbox or checkboxes that enforce the preferred failover action for vSphere HA.

- vSphere HA is subject to VM antiaffinity rules during failover. If VMs with this rule are placed together, the failover is aborted.
- vSphere HA is subject to VM-to-host affinity rules during failover. vSphere HA attempts to place VMs with this rule on the specified hosts if possible.

vSphere DRS affinity rules help protect SAP application and database components by providing appropriate separation between the primary and standby servers and also between multiple application servers.

Multiple SAP application servers providing the same function must be separated from each other by using antiaffinity rules. An SAP application server and database server can be collocated on the same physical server by using affinity rules to optimize performance in certain situations.

Additional vSphere HA Considerations

- Enable and configure proper admission control for the cluster.
- Set the restart priority to **High** for the VM or VMs that are hosting the SAP applications and database.

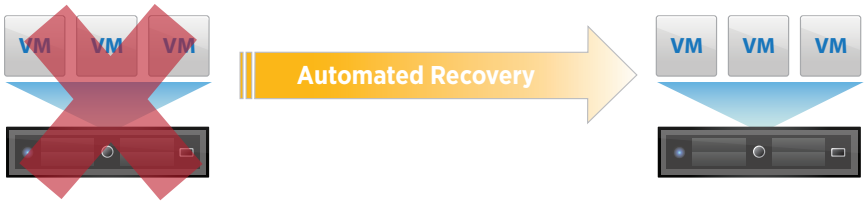

Figure 2. vSphere HA for Protection SAP Application VMs

vSPHERE HA KEY POINTS	CONSIDERATIONS
• Protection against server failure	• No monitoring of application
• Automatic restart of VMs	• Database instance and application instance unavailable during failover
• Start-up scripts and service required to autostart SAP database instances in guest OS	• Time to recover that includes time to boot guest OS and restart application
• Easy to configure: "VMware "out of the box"	• Database restart for crash-consistent recovery

Table 1. vSphere HA for SAP Protection

VMware vSphere Fault Tolerance

In the event of server failures, VMware vSphere Fault Tolerance (vSphere FT) provides continuous availability for applications with as many as four virtual CPUs. It does so by creating a live shadow instance of a VM that is always up to date with the primary VM. In the event of a hardware outage, vSphere FT automatically triggers failover, ensuring zero downtime and preventing data loss. Like vSphere HA, it protects against hardware failure but completely eliminates downtime with instantaneous cutover and recovery. After failover, vSphere FT automatically creates a new, secondary VM to deliver continuous protection for the application.

vSphere FT offers the following benefits:

- Protects mission-critical, high-performance applications regardless of OS
- Provides continuous availability, for zero downtime and zero data loss with infrastructure failures
- Delivers a fully automated response

When virtualizing an SAP application, technologies such as vSphere FT can help protect the SAP server from hardware failures. Compared to vSphere HA, vSphere FT can provide instantaneous protection, but the following limitations must be considered:

- The SAP system is limited to four vCPUs.
- vSphere FT protects against hardware failures but not against application failures.
- vSphere FT cannot reduce downtime for patching-related outages.
- vSphere FT has resource requirements that can create additional overhead.

Because vSphere FT is suitable for workloads with a maximum of four vCPUs and 64GB of memory, it can be used for SAP components such as ASCS, which is critical and a SPOF.

Figure 3. vSphere Fault Tolerance for ASCS Protection

vSphere FT can be turned on through vSphere Web Client.

When vSphere FT is turned on, vCenter Server resets the VM's memory limit and sets the memory reservation to the memory size of the VM. While vSphere FT remains on, memory reservation, size, and limit, as well as number of vCPUs and shares, cannot be changed. In addition, disks for the VM cannot be added or removed. Disk usage is twice the normal amount: two copies of the data are stored.

Connect vSphere Web Client to vCenter Server by using an account with cluster administrator permissions.

Prerequisites for Using vSphere FT

All hosts with vSphere FT enabled require a dedicated VMkernel interface.

The option to turn on vSphere FT is unavailable (dimmed) if any of these conditions apply:

- The VM resides on a host that does not have a license for the feature.
- The VM resides on a host that is in maintenance mode or standby mode.
- The VM is disconnected or orphaned—that is, its VMX file cannot be accessed.
- The user does not have permission to turn the feature on.

It is required that the hosts have a dedicated 10GBps network interface for vSphere FT logging traffic.

Configuration Procedure

1. In vSphere Web Client, browse to the VM for which vSphere FT is to be turned on.
2. Right-click the VM representing the ASCS server and select **Fault Tolerance** > **Turn On Fault Tolerance**.
3. Click **Yes**.
4. Select a datastore on which to place the secondary VM configuration files. Then click **Next**.
5. Select a host on which to place the secondary VM. Then click **Next**.
6. Review the selections and then click **Finish**.

The specified ASCS VM is designated as a primary VM; a secondary VM is established on another host. The primary ASCS VM is now fault tolerant.

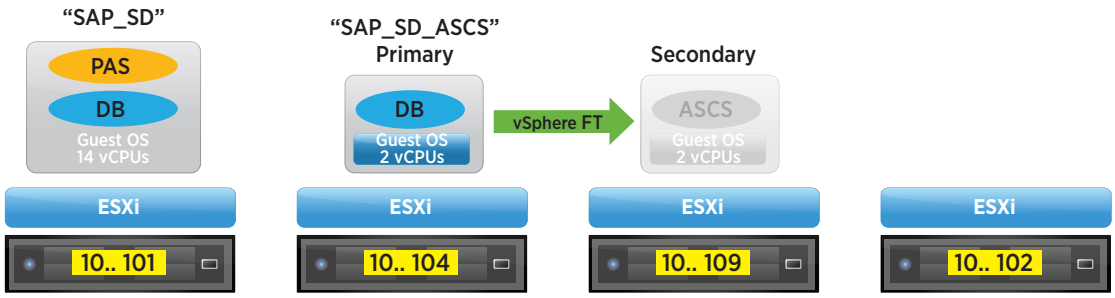

Figure 4. Lab Configuration: VMware FT Test with Virtual Machine Running ASCS on vSphere 6.0

The configuration shown in Figure 4 was installed in VMware Labs. A small-scale functional test was conducted to verify continuous availability of central services during failover of the ASCS VM protected via vSphere FT. The results are shown in Table 2.

TEST SETUP	RESULTS WITH FAILOVER
• Four VMware ESXi™ hosts running vSphere 6.0	• 1,000 concurrent users: < 0.5 sec response time (users generated by SAP workload kit)
• One VM, 14 vCPUs, running ECC 6.0: PAS and database instance (together in one VM for lab test only)	• Successful completion of workload with no user or lock errors
• One VM, two vCPUs, 8GB RAM running ASCS protected by vSphere FT	• Average CPU utilization of ASCS VM: < 19 percent
• *NOTE: This setup was not intended or tuned for benchmarking.*	• New secondary VM automatically created after failover

Table 2. Lab Results: vSphere FT Test with Virtual Machine Running ASCS on vSphere 6.0

When deploying SAP central services standalone in a VM, note the following:

- Linux-based guest OS is supported by SAP; there are no caveats.
- For Microsoft Windows–based guest OS, see SAP note 1609304: *Installing a Standalone ASCS Instance*. To obtain support on Windows for a standalone deployment, follow these guidelines:
  - Use a "SAPinst" that allows installation of standalone central services (available from NetWeaver Application Server 7.3 but also possible with some earlier versions).
  - Attend to RFC destinations that point to the virtual host name of the central services by maintaining RFC group destinations or implementing a standalone gateway.
  - In case of an upgrade, choose the correct upgrade tools. For advice, open an SAP message under support component BC-UPG.
  - For clarification, open an SAP ticket under support component BC-OP-NT-ESX before proceeding with an installation.

vSPHERE FT FOR ASCS KEY POINTS	CONSIDERATIONS
• Assumes three-tier; application server VMs not shown	• No monitoring of application
• ASCS protected via vSphere FT; Database instance protected via vSphere HA	• Current vSphere FT support for four-vCPU VMs; fits 99.9 percent of all SAP ASCS systems
• Protection against server failure	• Separate network adapter and network recommended for vSphere FT logging traffic
• Continuous availability of central services	
• Easy to configure: VMware "out of the box"	
• Database instance still protected via vSphere HA	
• New secondary ASCS VM automatically created after failover; assumes more ESXi hosts available	

Table 3. vSphere FT for SAP Protection

Third-Party Tools Leveraging the VMware Application Monitoring API

Third-party tools such as Symantec ApplicationHA, SUSE Linux Enterprise High Availability Extension (SUSE SLES HAE), and Red Hat Cluster Suite provide monitoring capabilities for applications running inside VMs. vSphere HA has an API that enables third-party vendors to develop agents that can monitor the health of an application running within the guest OS and inform vSphere HA when a problem is detected. These VMware partners have developed an agent for providing application detection within a vSphere cluster.

Agents exist for monitoring SAP SPOFs. The application agent runs a utility to verify the status of the instance—central instance or database, for example. The agent detects application failure if the monitoring routine reports an improper function of the instance processes. When this application failure occurs, the availability agent for SAP tries to restart the instance. If it further fails, a VM reboot is triggered.

The key points of this solution are summarized in Table 4.

THIRD-PARTY vSPHERE HA AGENTS KEY POINTS	CONSIDERATIONS
• Agents build upon vSphere HA to enable application-level detection to improve application availability.	• Recovery time depends on the time it takes to restart the service or processes.
• They do not impede the functionality of VMware® products such as vSphere DRS and VMware vSphere vMotion®.	• If vSphere HA is invoked, downtime is incurred for the amount of time it takes to boot the guest OS and start the application.
• Application monitoring and management are through a single pane of glass using vCenter Server plug-in.	• Agents are not supported for use with vSphere FT protected VMs.
• Application and dependency detection enables graceful start-up and shutdown.	
• Agent setup is less complex than for clustering.	

Table 4. Third-Party vSphere HA Agents for SAP Protection

Virtual Machine Guest Operating System Clustering Solutions

It is common in physical setups to install SAP with third-party cluster solutions. Such solutions can also be deployed with SAP running on VMs. A typical clustering setup includes disks that are shared between nodes. A shared disk is required as a quorum disk. In a cluster of VMs across physical hosts, the shared disk must be either Fibre Channel (FC) SAN, FCoE, or iSCSI storage protocol. A private heartbeat network is required between the nodes.

A list of clustering solutions certified with SAP is available at http://scn.sap.com/docs/DOC-31701. If third-party clustering solutions are being run on VMware technologies for an SAP deployment and there are any installation issues or support questions, follow these guidelines:

- For Windows Server Failover Clustering (WSFC) and Oracle RAC, contact VMware Global Support Services (GSS).
- For the following, first contact the respective cluster vendor. If the vendor proves that there is a VMware issue, they will contact VMware Support and VMware will work with the vendor and customer to resolve the issue.
  - Red Hat Enterprise Linux High Availability – "Virtualization Support for High Availability in Red Hat Enterprise Linux 5, 6, and 7": https://access.redhat.com/articles/29440
  - SUSE SLES HAE – "SUSE Alliance Partners: SUSE + VMware": https://www.suse.com/partners/alliance-partners/vmware/
  - Veritas Cluster Server – VMware Knowledge Base article 2046035: "Symantec Storage Foundation High Availability 6.x in Guest Storage, High Availability, and Disaster Recovery Support (Partner Verified and Supported)": http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2046035
  - NEC EXPRESSCLUSTER X – "High Availability for SAP by EXPRESSCLUSTER": http://www.nec.com/en/global/prod/expresscluster/solution/sap/index.html
  - SteelEye LifeKeeper –
    - *SIOS LifeKeeper Single Server Protection for Linux*: http://us.sios.com/products/Lifekeeper/
    - *Adding Application Protection in Virtualized SAP Environments in vSphere*: http://www.cc-dresden.de/fileadmin/Downloads/PDF/Praesentationen/SAP_TechEd_2013.pdf

Figure 5 shows two VMs that run clustering software at the OS level with application monitoring and remediation. The VMs share a private heartbeat and a public network connection backed by shared storage.

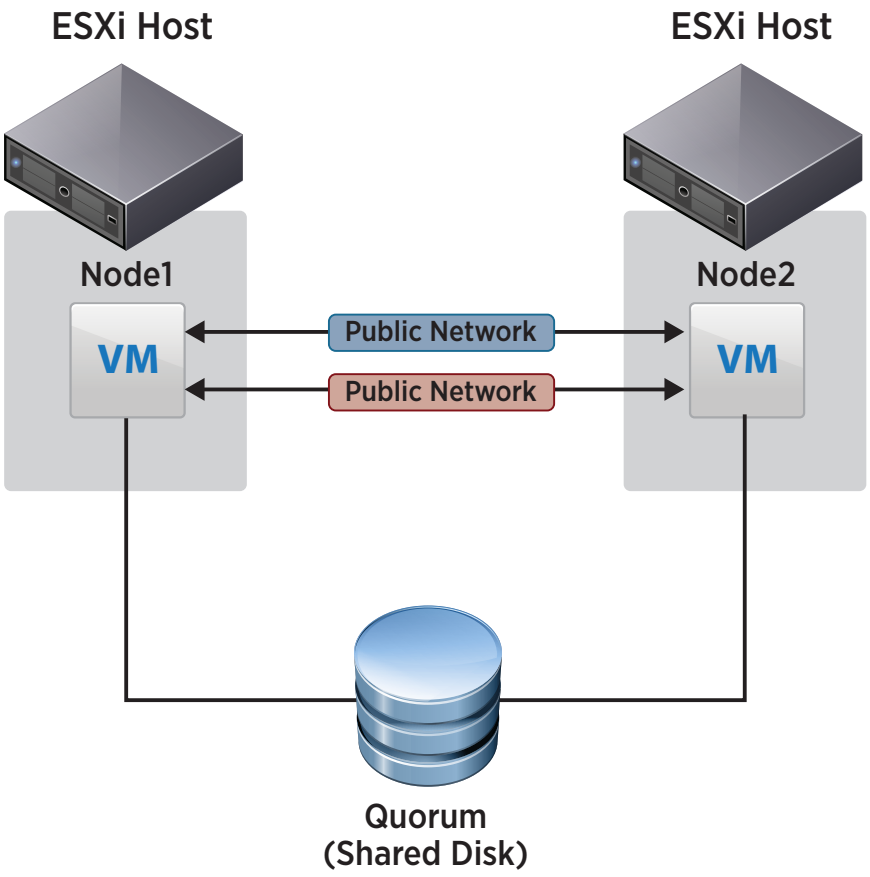

Figure 5. In-Guest Clustering Setup (Source: https://vikernel.files.wordpress.com/2014/04/cib_diagram.jpg)

This solution uses a primary and a standby VM for the SAP application being protected. The cluster framework monitors the health of the application resources. If a configured application instance or associated services become unavailable, the cluster services fail over the services to the standby node. Planned downtimes relating to patching the OS are protected by this solution because the application can be failed over to the standby VM during OS patching and the downtime can be minimized.

SUSE SLES HAE or Red Hat Cluster Suite can be leveraged for protecting SAP applications deployed on Linux. WSFC is one method of guest OS clustering that can be used for SAP applications running on Windows.

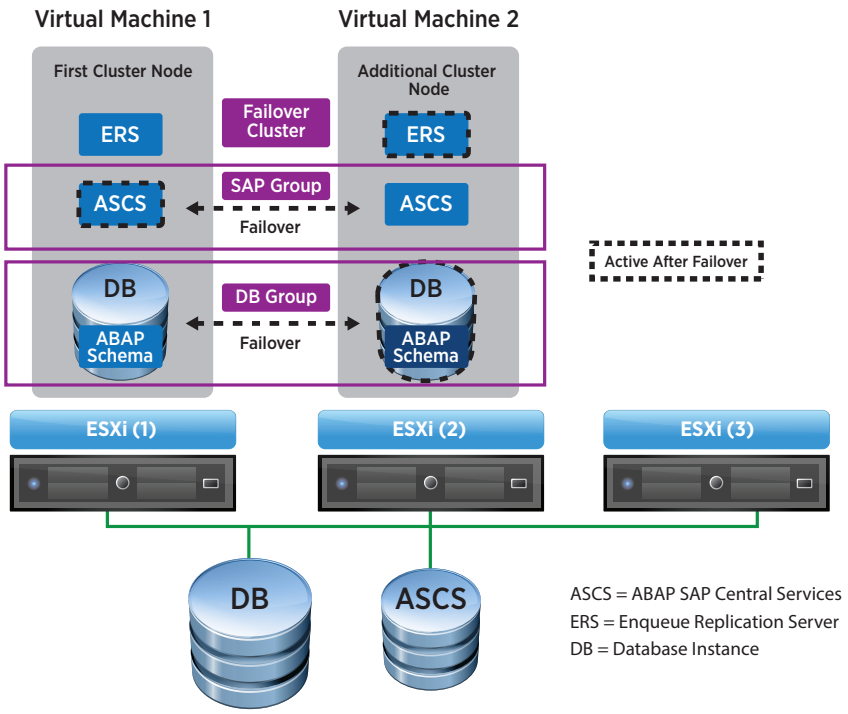

Figure 6. SAP with WSFC (Source: SAP Windows Installation Guide)

Figure 6 shows a two-node WSFC cluster on a three-node ESXi cluster. If ESXi host 1 fails, the database instance fails over to VM 2 on ESXi host 2. Meanwhile, vSphere HA restarts the failed VM 1 on ESXi host 3. VM 1 then rejoins the WSFC cluster.

IN-GUEST CLUSTERING KEY POINTS	CONSIDERATIONS
• Layers clustering technology on top of the vSphere cluster	• Recovery time that depends on the time it takes to restart the service or processes
• Can impede the functionality of VMware products such as vSphere DRS and vSphere vMotion	• Can provide reduced downtime during OS and application maintenance
• Application monitoring and management through independent agents and management	• Not supported for use with vSphere FT protected VMs
• Application and dependency detection that enables graceful start-up and shutdown	
• Most complex setup	

Table 5. Guest Failover Clustering for SAP Protection

SAP Database Vendor Replication Considerations

The solutions referenced in the previous section can be used to protect SAP databases in a crash-consistent manner. Specialized high-availability and replication solutions are available for the most common SAP databases for data-consistent protection. This involves a mirrored copy of the database that requires duplicate storage—that is, a nonshared storage solution. Table 6 shows the solutions that can be layered on top of VMware products to maintain data consistency.

These solutions help minimize recovery point objective (RPO) and recovery time objective (RTO). The data consistency of the replicated database enables faster restart times of the replica, which can help lower the RTO.

DATABASE	TECHNOLOGY
Oracle Database	Data Guard
Microsoft SQL Server	AlwaysOn Availability Groups (SQL AAG)
SAP HANA	SAP HANA System Replication

Table 6. Database Availability Technologies

These solutions can be used as follows:

1. As a local high-availability solution, within the same data center, for mission-critical SAP databases that require fast recovery in case of primary database failure; duplicate storage removes SAN as a SPOF; focus of this section is on this use case
2. As a DR solution where the database vendor replication technology is used instead of a SAN-based replication solution

Oracle Data Guard

Oracle Data Guard protects Oracle databases by enabling management, monitoring, and automation software to create and maintain one or more standby databases and providing high availability for mission-critical applications. Data Guard is included as part of Oracle Database Enterprise Edition.

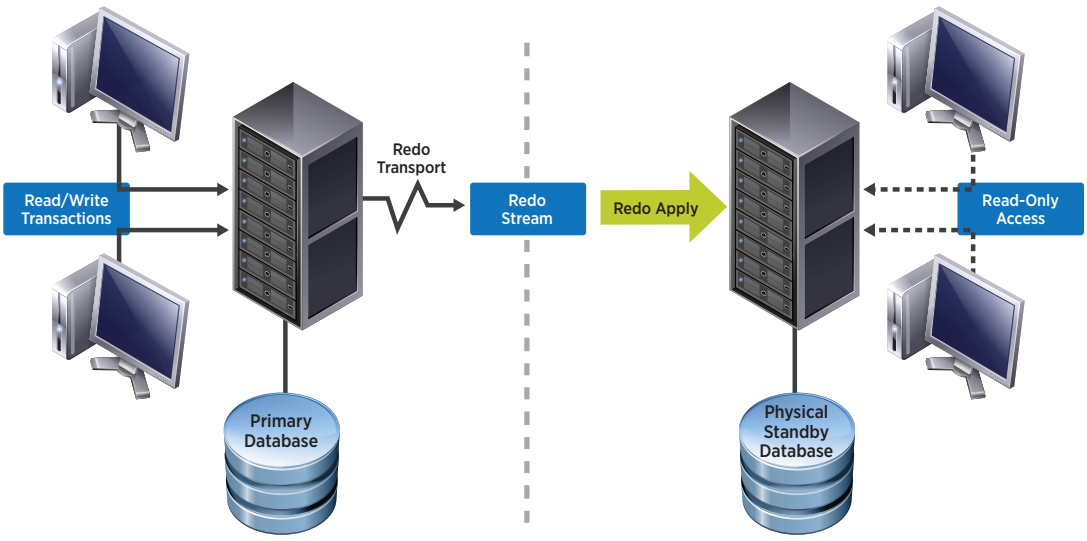

Figure 7. Oracle Database Protection with Data Guard

Microsoft SQL AlwaysOn Availability Groups (AAG)

The AlwaysOn Availability Groups (AAG) feature is a high-availability solution that provides a good alternative to database mirroring. AAG was introduced in Microsoft SQL Server 2012 and can be used to maximize the availability of business-critical databases such as those used for SAP. It can be combined with WSFC to automate failover to replica on failover.

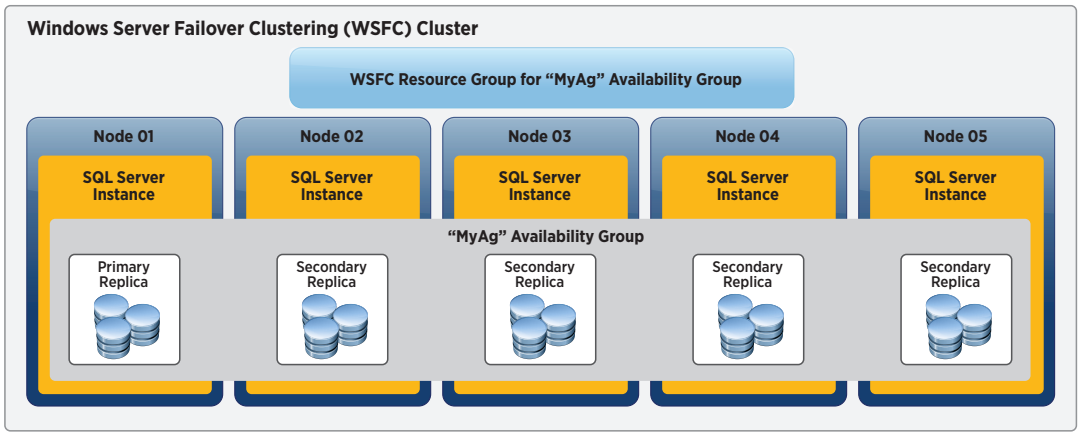

Figure 8. Microsoft SQL Server with AAG

The following diagram shows an installation of SAP on SQL Server with AAG. For details, see http://blogs.msdn.com/b/saponsqlserver/archive/2011/11/17/microsoft-s-sap-deployment-and-sql-server-2012.aspx.

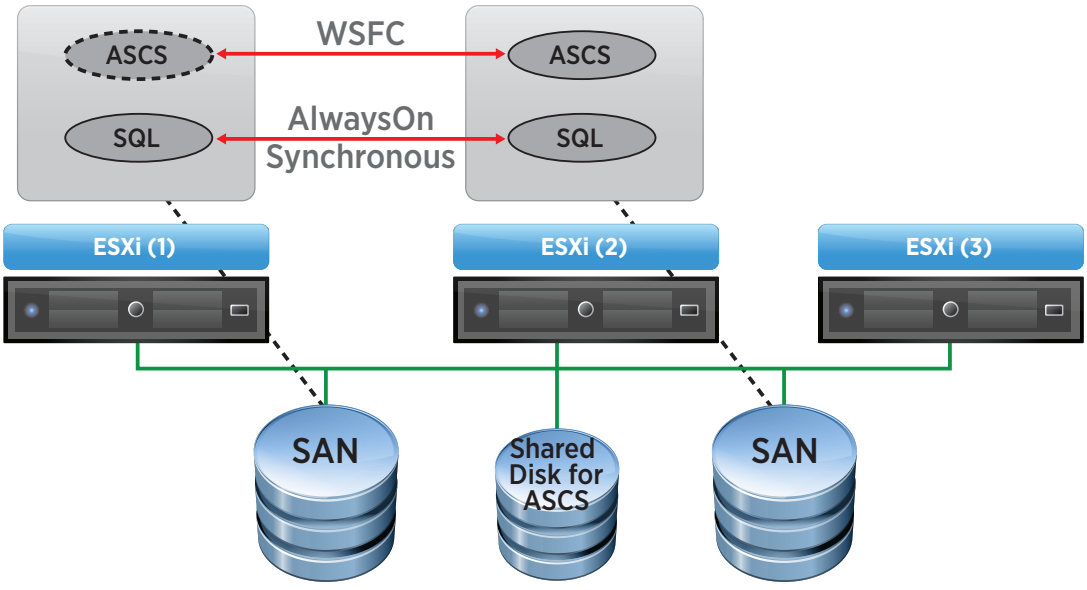

Figure 9. SAP HA with WSFC for ASCS and AlwaysOn for SQL

In the architecture shown in Figure 9, the SAP central services component is protected with shared-disk WSFC configuration. Meanwhile, the SQL Server instance is deployed on nonshared storage with AAG.

SAP HANA Availability

Using vSphere HA to protect an SAP HANA system is the easiest and most cost-efficient way to protect a virtualized SAP HANA system against OS and hardware failures, without the dependency on any external components such as DNS servers or features that the SAP HANA Storage Connector API provides. In the event of a failover, the entire SAP HANA VM is restarted on the same host; if the hardware of the initial host is defective, the entire SAP HANA VM is restarted on another host in the vSphere cluster. Because all virtualized disks—such as OS, data and log VMDKs, and in-guest mounted NFS volumes—are failed over to a new ESXi host as well, no specific storage tasks or cluster solutions are required.

By leveraging VM restart polices and affinity and antiaffinity rules, it is possible to provide high-availability protection for independent SAP and non–SAP HANA VMs on a single cluster. The only requirement is that the failover host has enough remaining resources to start the failed VM and that SAP HANA services are configured to start automatically after an OS reboot.

In an SAP HANA production environment, resource commitments are required and must be enforced to ensure optimal performance. In the case of an SAP HANA VM hardware event, the VM is restarted on a host with sufficient resources remaining. Another method is to configure the noncritical VM without resource commitments to intentionally enable overcommitment of resources. In the event of a failover, the SAP HANA VM requests all required resources from the ESXi host; the noncritical VM receives only unclaimed resources remaining on the ESXi host.

Figure 10 shows two failure situations to which vSphere HA can react:

- Event 1 shows a hardware defect, affecting two VMs running on a failed host.
- Event 2 affects only the VM itself.

As event 1 occurs, two arrows show that "SAP HANA 1 worker" VM is moved to the host that runs beside SAP HANA master 2 "ANY workload" VM. "SAP HANA 2 worker" is moved to host 4 and restarts the SAP HANA worker 2 VM on this host.

Event 2 shows a VM OS failure. This VM is moved to host 4, because an "ANY workload" VM runs on this host. By using affinity and antiaffinity rules—and the possibility to shut down or reconfigure a no-resource commitment for the "ANY workload" VM—it is possible to provide a high-availability environment that does not waste any resources.

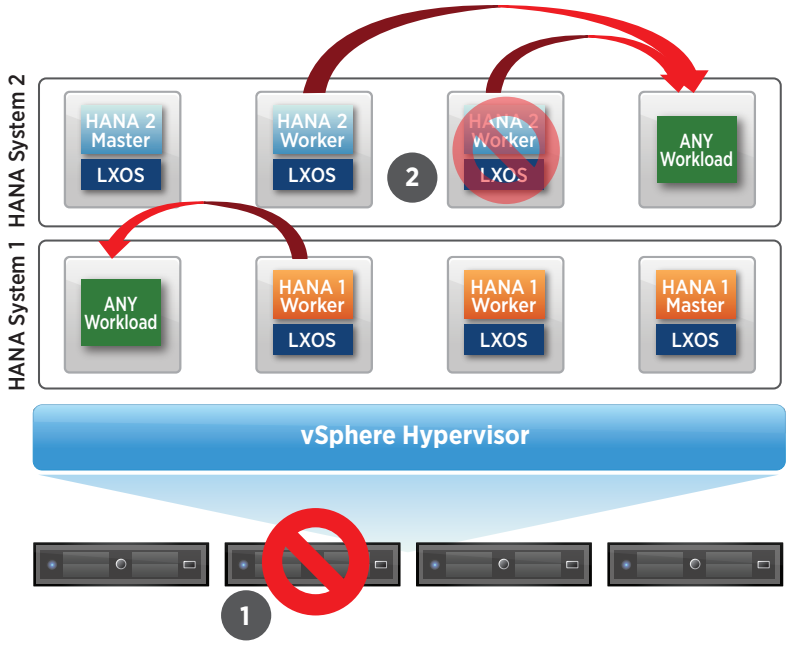

Figure 10. vSphere HA Protected SAP HANA Systems

During the failover, it is not possible to access the tables maintained by the affected SAP HANA worker or master VM. Depending on the last written savepoint and log, SAP HANA initiates a crash recovery after the SAP HANA processes have started. After the crash recovery, the SAP HANA system is accessible and usable as it was before the fault (RPO=0).

NOTE: Logical or file corruption failures are not recoverable.

SAP HANA Auto-Restart Feature with vSphere HA

SAP HANA provides a service auto-restart and watchdog function to automatically detect and restart stopped HANA processes.

All configured SAP HANA services—index server, name server, and others—are restarted by the SAP HANA service auto-restart function, which automatically detects failures, whether a software failure or an intentionally stopped SAP HANA process by an administrator. When detected, these stopped SAP HANA processes restart automatically. The in-memory loaded SAP HANA data not affected by a process failure is not reloaded. After the process has restarted, only the affected SAP HANA service process data is reloaded into memory, and SAP HANA resumes its function.

vSphere HA combined with the SAP HANA auto-restart feature provides a very solid high-availability solution to minimize failures due to power, hardware, OS, or SAP HANA application faults. Combining these two solutions provides a high-availability solution with better than 99.9 percent effectiveness.

In a view of the protected components of a virtualized SAP HANA system, Figure 11 shows an invocation of vSphere HA and the SAP HANA auto-restart feature.

Additionally, SAP HANA application failures—such as process crashes or OS failures—can also be protected via third-party solutions, such as that offered by Symantec. A benefit of such a third-party solution is that it can also detect OS failures and can leverage the vSphere HA API to initiate a VM restart or failover; the SAP HANA watchdog can only monitor and restart the SAP HANA processes. See the third-party vendor product pages for support availability for SAP HANA scale-out configurations.

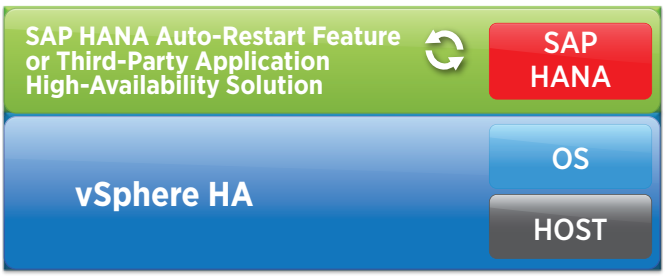

Figure 11. vSphere HA Protected Components

Using vSphere HA and the SAP HANA auto-restart feature provides the following:

- The ability of vSphere HA to protect against OS and host failures without specific configuration requirements
- Protection against SAP HANA application process failures, improving application uptime to more than 99.9 percent
- Third-party VMware virtualized SAP HANA solutions that can extend the monitoring beyond SAP HANA processes

SAP HANA Systems Protected with Host Auto-Failover to a Standby VM

Protecting a VMware virtualized SAP HANA system with vSphere HA already provides very high system availability. Customers who want to use the host auto-failover solution to fail over to a standby VM can do this with selected storage vendors that offer SAP HANA Storage Connector API or STONITH implementations validated for the vSphere virtualized environment.

A benefit of using a standby VM to protect an SAP HANA system is having an operational VM with a preinstalled OS and SAP HANA configuration that is available and online. This VM can be used for specific maintenance tasks only available with a standby VM, such as OS or SAP HANA application binary patching, where the active SAP HANA system configuration can be moved over to the standby node that already has a patched OS and SAP HANA binaries available. This is not possible in a vSphere HA environment because only one VM per SAP HANA worker exists at any time. Nevertheless, before performing such maintenance tasks, it is possible to take a snapshot of the VM to do or undo in the case of a binary problem, and any patches of configuration changes are restored within seconds. After the maintenance, such snapshots must be deleted; otherwise, the I/O performance of this VM will be negatively impacted. An organization must decide which solution is best. Virtualizing SAP HANA with vSphere does not limit the customer's freedom of choice.

Host Auto-Failover

When an SAP HANA has a host auto-failover—also called SAP HANA standby host solution—the active (worker) VMs are monitored:

1. The SAP HANA name server also acts as the cluster manager that regularly monitors all SAP HANA worker VMs.
2. If a failing SAP HANA instance is detected, the cluster manager determines which standby system is used for the failover and ensures that the standby VM take over the role of the failing VM and start its database instance using the persisted data and log files of the failed instance. The SAP HANA storage connector implementation of the storage vendor is executed, ensuring that the failed SAP HANA instance no longer has write access to the log and data files stored in VMDK disks. This prevents data corruption and is done via a STONITH script or SCSI-3 persistent reservations and the SAP HANA Storage Connector API.
3. When the SAP HANA Storage Connector returns successfully, the disks are accessible only by the standby VM.
4. The data and log files are loaded into memory while SAP HANA automatically starts the required crash recovery steps to ensure data consistency for this instance.

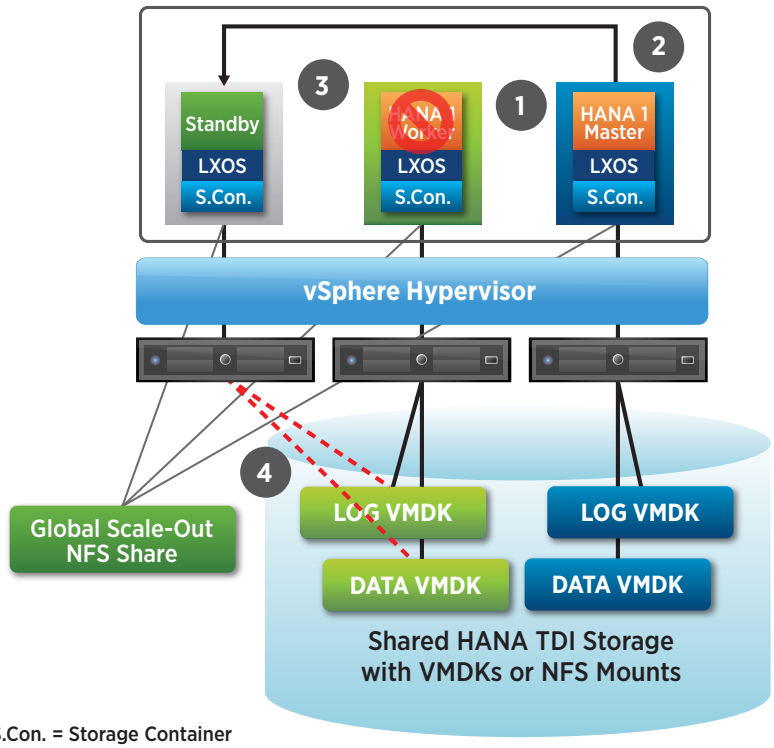

Figure 12. SAP HANA Host Auto-Failover Step

SAP HANA Disaster Recovery Solutions with VMware vSphere

SAP HANA supports replication technologies used for DR solutions based on vSphere. SAP HANA System Replication provides a very robust solution to replicate the SAP HANA database content to a secondary disaster site; this storage-based system replication can be used as well.

When using SAP HANA System Replication, the same number of SAP HANA VMs must exist at the DR site. These VMs must be configured and installed similarly to a natively running SAP HANA system with System Replication enabled.

SAP HANA System Replication provides various modes for system replication:

- Synchronous
- Synchronous in-memory
- Asynchronous

Depending on requirements, the DR VMs can consume more or fewer resources on the DR vSphere cluster. For instance, selecting the synchronous in-memory mode consumes the same amount of RAM as the primary systems. This mode is required only if the customer requests the shortest recovery time. In most customer scenarios, using synchronous data replication is sufficient. SAP states that by only replicating the data, about 10 percent of the system resources are required, enabling up to 90 percent of the resources to continue to be used by other systems such as test or QA systems.

Figure 13 shows the scenario of a codeployed configuration in which all SAP HANA scale-out VMs are replicated to a DR site per VM level.

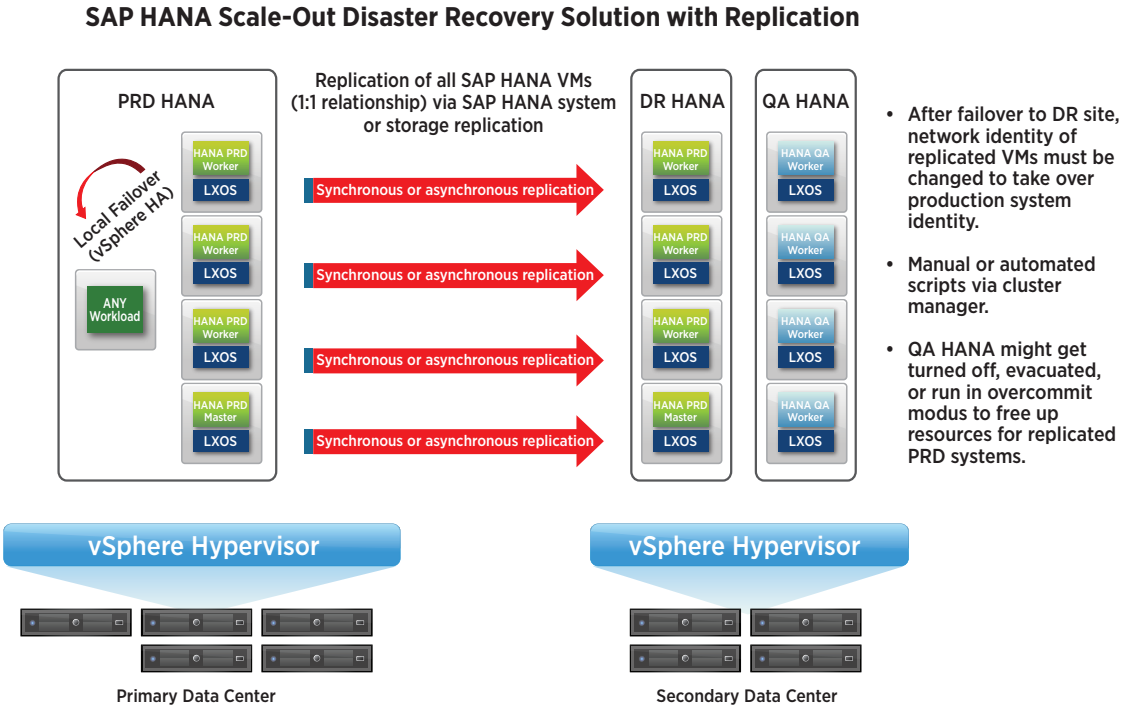

Figure 13. One-to-One VM Replication via SAP HANA System or Storage Replication

In this scenario, resource overcommitments are used to enable the codeployment of such an environment. By using resource pools and resource shares, it is possible to provide the required resources to the DR SAP HANA scale-out VMs. The codeployed system, with fewer resource shares, experiences performance degradation after the DR systems are used following a site failover. Evacuate these VMs to other available vSphere systems to free up all resources for the DR SAP HANA VMs. This is another option to running the two systems in parallel—with resource limitations—on the same platform.

To switch the network identity (IP redirect) of the replicated systems from the DR configuration to the production configuration, system replication via storage—or the SAP HANA replication solution—requires additional steps after a site failover. This can be done manually or via automated tools such as HP Serviceguard, SUSE HA Cluster, SAP Landscape Virtualization Manager (SAP LVM), or other cluster managers. The configuration of such a solution in a virtualized environment is similar to that of natively running systems. Contact the respective storage vendor to discuss a cluster manager solution supported by their storage solution.

Summary: SAP HANA High-Availability and Disaster Recovery Solutions with VMware vSphere

Table 7 provides an overview of the various SAP HANA high-availability and DR solutions available when running on vSphere.

The complexity and cost level can increase with certain solutions, from vSphere HA combined with the SAP HANA Auto-Restart feature, to a completely replicated SAP HANA system, to a second DR site.

Tools such as VMware vCenter™ Site Recovery Manager™ can help reduce the complexity of a DR solution.

HIGH-AVAILABILITY SOLUTION: QUALITIES	vSPHERE HA	vSPHERE HA + SAP HANA AUTO-RESTART FEATURE	SAP HANA HOST AUTO-FAILOVER (STANDBY VM)	SAP HANA SYSTEM REPLICATION	STORAGE SYSTEM REPLICATION	STORAGE SYSTEM REPLICATION + SITE RECOVERY MANAGER
Scenario Description	VMware standard high-availability solution. vSphere HA restarts or fails over VM to another host in case of a detected OS or hardware failure.	Standard vSphere HA is combined with SAP HANA auto-restart watchdog running inside a VM to monitor SAP HANA application status and trigger an SAP HANA process restart. OS and hardware failures are handled by vSphere HA.	SAP HANA standby VM automatically takes over the role of another SAP HANA VM in case of a detected failure.	Data replication between primary and secondary sites by leveraging SAP HANA System Replication functionality; there is no automated failover process. Third-party cluster solutions can be used to automate site failover.	Data replication between primary and secondary site by leveraging storage system replication functionality. No automated failover process. Third-party cluster solution can be used to automate site failover.	Site Recovery Manager automates the storage replication of the virtualized SAP HANA servers to another site.
OS Failures	Yes	Yes	Yes	Yes	Yes	Yes
Hardware Failures	Yes	Yes	Yes	Yes	Yes	Yes
Application Failures	No	Yes	Yes	Yes	Yes	Yes
IP Redirect/ DNS Update	Not necessary	Not necessary	Not necessary	Yes, manual or third-party software, such as cluster manager solution, necessary	Not necessary because complete VMs with all data disks are replicated	Not necessary because complete VMs with all data disks are replicated
RTO	Medium (crash recovery of database)	Medium (crash recovery of database)	Short to medium (only if in-memory data loading)	Shortest to medium RTO, depending on IP redirect solution	Long and manual restart (crash recovery of database)	Medium (crash recovery of database); automated restart
RPO	0	0	0	0 (synchrony) > 0 (asynchrony)	0 (synchrony) > 0 (asynchrony	0 (synchrony) > 0 (asynchrony
Performance Ramp	Minutes to hours (bootstrap + loading)	Minutes to hours (bootstrap + loading)	Minutes to hours (bootstrap + loading)	Seconds to minutes if synchronous replication into memory is selected, depending on the selected modus and IP redirect solution.	Hours to days (bootstrap + loading)	Hours (bootstrap + loading)
Cost	Included	Included	Included	High	High	High
Complexity	Low	Low	Medium	High	High	High

Table 7. SAP HANA Scale-Out High-Availability and Disaster Recovery Solutions with VMware vSphere

SAP HANA has its own availability that leverages in-guest clustering solutions—Linux HA cluster software, for example. Most of the protection that this is designed to do is for hardware failures. Because vSphere HA can protect against hardware failures in a VMware environment, it can be leveraged for SAP HANA.

The solution for application-level protection called SAP HANA System Replication can be leveraged even in a VMware environment. By combining SAP HANA System Replication with vSphere HA, high levels of availability can be achieved for SAP HANA environments.

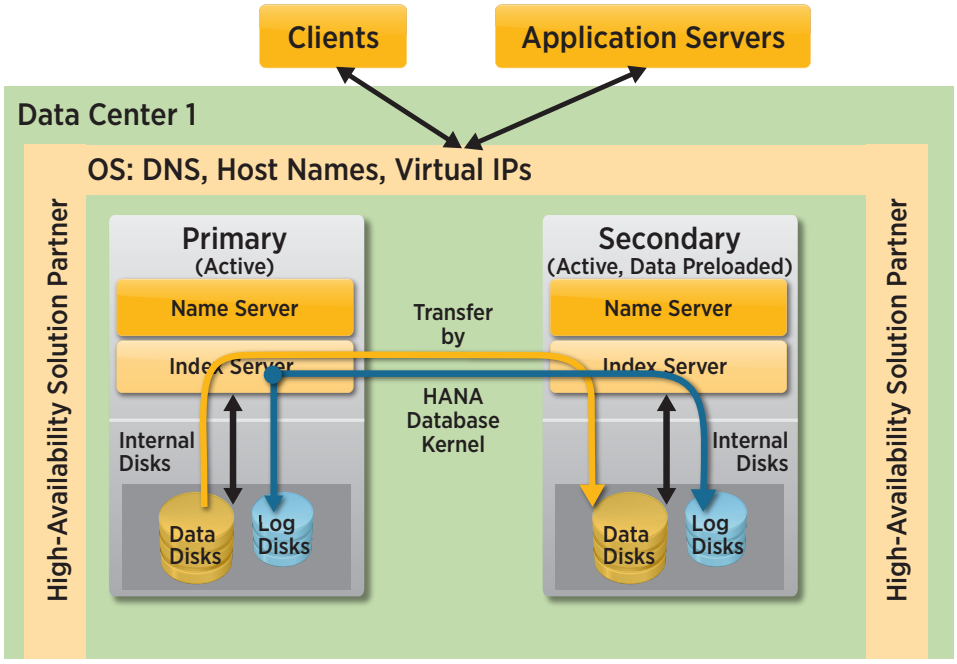

Figure 14. SAP HANA System Replication. Source: http://scn.sap.com/docs/DOC-60341

Automatic failover is available via SAP partner solutions:

- SUSE – https://www.suse.com/communities/conversations/saphanasr-helps-automate-sap-hana-system-replication-sles-sap-applications/
- Red Hat – https://access.redhat.com/articles/1466063

DATABASE REPLICATION KEY POINTS	CONSIDERATIONS
• Acts at the database level, ensuring data consistency	• Possibility of built-in automation for cutover
• Does not impede the functionality of VMware products such as vSphere DRS and vSphere vMotion	• Duplicate storage required in local high-availability setups and must be on separate SAN to prevent a SAN SPOF
• Unique solutions based on database vendor	
• Potential additional licensing cost for use of features	
• SAN not a SPOF in local high-availability setups: database replication configured in synchronous mode	

Table 8. Data-Consistent Protection for SAP

vSphere vMotion for Improved Availability

vSphere live migration, vSphere vMotion, enables moving an entire running VM from one physical server to another, without downtime. The VM retains its network identity and connections, ensuring a seamless migration process. Transfer the VM's active memory and precise execution state over a high-speed network, enabling the VM to switch from running on the source vSphere host to the destination vSphere host. This entire process takes less than 2 seconds on a Gigabit Ethernet network. SAP infrastructures can leverage vSphere vMotion to accomplish the following tasks:

- Automatically optimize VMs within resource pools
- Perform hardware maintenance without scheduling downtime or disrupting business operations
- Move VMs away from failing or underperforming servers

The previously mentioned third-party cluster and high-availability solutions such as WSFC, Oracle RAC, and SUSE HA Cluster are compatible with vSphere vMotion if certain prerequisites are met.

Summary: Comparison of High-Availability Options

SAP can leverage all the previously discussed solutions for high availability. The solutions actually used depend on customer requirements. Table 9 shows the various solutions, their features, and their cost and complexity ratings. The more solutions that are used, the more expensive and complex they are. The choice should be dictated by minimum customer requirements for SAP availability.

HIGH-AVAILABILITY SOLUTION	VM RESTARTING	APPLICATION MONITORING	FAILOVER TIME	DATABASE CONSISTENCY	HARDWARE MAINTENANCE DOWNTIME	OS MAINTENANCE DOWNTIME	COST AND COMPLEXITY
vSphere HA	YES	YES	MED	NO	NONE	HIGH	LOW
vSphere HA + vSphere FT	YES	YES	NONE	NO	NONE	HIGH	MED
vSphere HA + Third-Party Monitoring	YES	YES	MED	NO	NONE	HIGH	MED
Guest Failover Cluster	NO	YES	LOW	NO	MED	LOW	HIGH
vSphere HA + Guest Failover Cluster	YES	YES	LOW	NO	NONE	LOW	HIGH
vSphere HA + Guest Failover Cluster + Database Availability + Consistency	YES	YES	LOW	YES	NONE	LOW	HIGH

Table 9. Availability Solutions for SAP Solutions and Their Characteristics

Disaster Recovery for Virtualized SAP

Due to their nature, business-critical applications (BCAs) must have robust DR. In legacy physical server infrastructures, protecting BCAs against disasters is a monumental task for the following reasons:

1. Identical dedicated hardware is needed in the recovery site, which is expensive and underutilized.
2. Recovery site applications, OSs, and hardware must be kept up to date.
3. Any testing requires downtime of the production environment.

With the advent of virtualization and the concept of encapsulation of VMs, replicating entire business-critical workloads has been greatly simplified. VMs are represented as a set of files that can easily be replicated to the recovery site and reinstantiated on different hardware.

SAP applications can effectively be protected using VMware solutions such as Site Recovery Manager. Site Recovery Manager leverages the unique aspects of VMs in combination with replication management and workflows to automate DR for SAP applications. The following DR use cases are of most benefit to virtualized SAP environments.

Automated DR failover

- Initiate recovery plan execution from the vSphere Web Client with a single click of a button
- Halt replication and promote replicated VMs for fastest possible recovery
- Execute user-defined scripts and pauses during recovery

Planned migration and disaster avoidance

- Graceful shutdown of protected VMs at the original site
- Replication synchronization of protected VMs prior to migration, to avoid data loss
- Restart of protected VMs in an application-consistent state

Seamless workflow automation with centralized recovery plans

- Create and manage recovery plans directly from the next-generation user interface of vSphere Web Client
- Predefine the boot sequence of VMs for automated recovery
- Reconfigure IP addresses upon failover at the subnet or individual address level

Test SAP Disaster Recovery Setup

A test environment was set up to create a DR environment for SAP. The environment consists of two sites: the primary site and the recovery site. The primary site is located in Wenatchee, Washington; the recovery site is located in Cambridge, Massachusetts. A test SAP environment was implemented in the primary site. In addition, infrastructure components such as domain controllers and Site Recovery Manager servers are also deployed. The environment and the VMs representing SAP and other applications are shown in Figure 15.

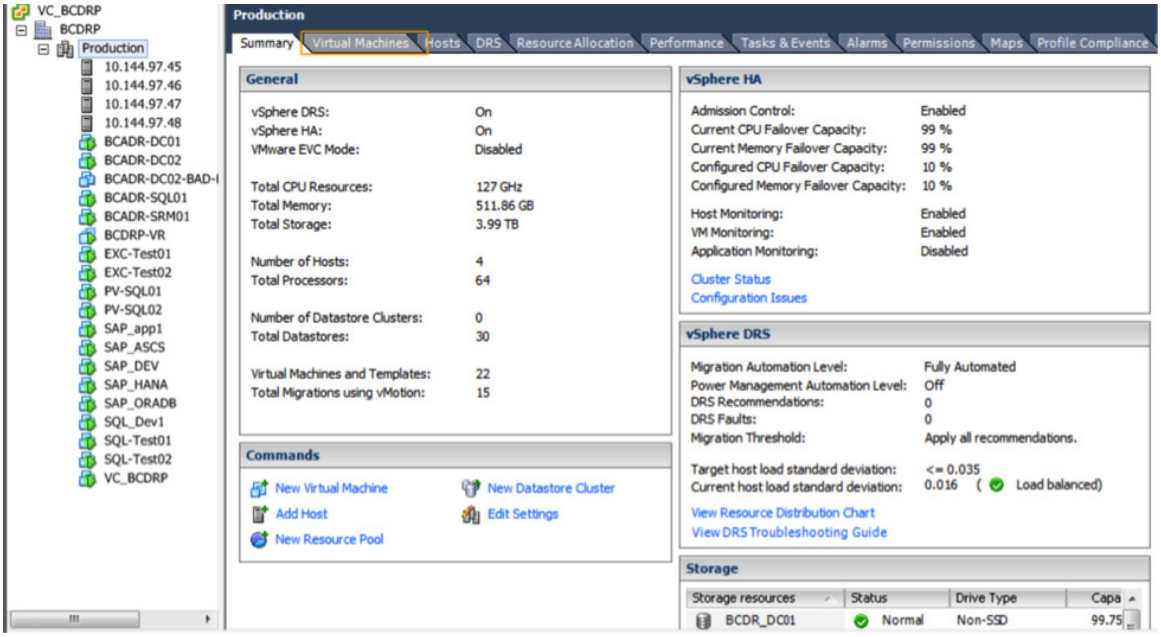

Figure 15. Primary Site with Infrastructure and BCA Components

The recovery site has three vSphere servers along with some local servers relating to infrastructure such as domain controllers, Site Recovery Manager servers, and other local applications. The remote site has placeholder VMs created by Site Recovery Manager for all protected VMs from the primary site.

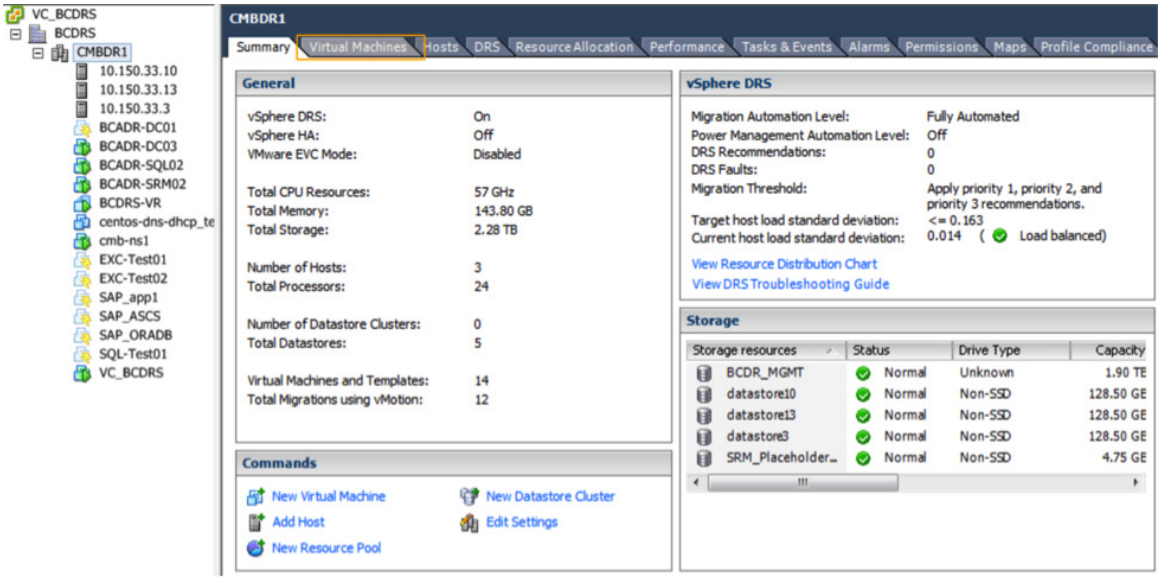

Figure 16. Recovery Site with Site Recovery Manager and Infrastructure Components

Replication

Site Recovery Manager

Site Recovery Manager can help reduce the complexity of a system replication DR solution by automating the complex DR steps on any level. It is designed for DR of a complete site or data center failure. It supports both unidirectional and bidirectional failover. It also supports a "shared recovery site," enabling organizations to fail over multiple protected sites into a single, shared recovery site. This site can, for example, also be a VMware vCloud® Air™ provided cloud data center, the VMware cloud service offering.

The following are the key elements that compose a Site Recovery Manager deployment for SAP:

- Site Recovery Manager is designed for virtual-to-virtual DR and requires a vCenter Server management server instance at each site. These two instances are independent, each managing its own site. Site Recovery Manager enables them to detect the VMs they must recover if a disaster occurs.
- Site Recovery Manager updates, manages, and executes DR plans. It is managed via a vCenter Server plug-in.
- Site Recovery Manager relies on storage vendors' array-based replication Fibre Channel or NFS storage that supports replication at the block level to replicate SAP HANA data and log files to the DR site. It communicates with the replication via storage replication adapters that the storage vendor offers and that have been certified for Site Recovery Manager.
- vSphere Replication has no such restrictions on use of storage type or adapters and can be used for non–performance-critical or static VMs, such as infrastructure services or SAP application servers with an RPO of 15 minutes or longer.

Figure 17 shows an example SAP landscape protected by Site Recovery Manager. The VMs running on the primary site contain all required infrastructure and SAP components such as LDAP, SAP HANA database, and SAP application servers, as in an SAP Business Suite implementation. The VMs can be replicated, depending on RPO needs, via vSphere, SAP HANA, or storage replication. vSphere Replication can be used with VMs that tolerate an RPO of 15 minutes or longer.

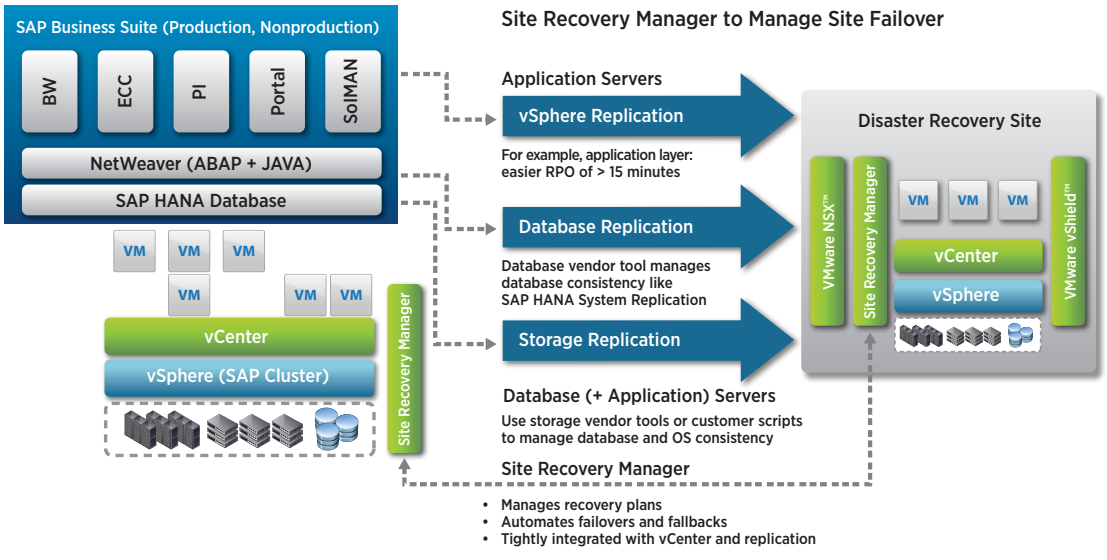

Figure 17. Site Recovery Manager Protected SAP Landscape

The following is a summary of the benefits of using Site Recovery Manager for managing the DR process for SAP landscapes:

- Reduced cost of DR by up to 50 percent
- Application-agnostic protection that eliminates the need for application-specific point solutions
- Support for vSphere Replication and array-based replication that offers choices and options for synchronous replication with zero data loss
- Centralized management of recovery plans directly from vSphere Web Client that replaces manual runbooks
- Automated protection through self-service, policy based provisioning via VMware vRealize™ Automation™
- Highly predictable recovery objectives via frequent, nondisruptive testing of recovery plans
- Reliably reduced RTOs via automated orchestration of site failover and failback with a single click
- Planned migration workflows that enable disaster avoidance and data center mobility

The backbone for DR is replication of the protected workloads from the primary to the recovery site. There are various types of replication.

Synchronous replication is used in active–active environments with zero RPO requirements. The scope of synchronous replication is within metro areas because latencies are required to be less than 10ms. Every write in the primary site is acknowledged only when it has been written to both sites. This solution is typically very expensive and is used by organizations that have zero RPO as a requirement. EMC vPLEX and NetApp MetroClusters are examples of stretched clusters that leverage sync replication techniques.

Asynchronous replication is used for the majority of DR deployments. The RPO for asynchronous replication can range from a few minutes to hours, depending on customer requirements. This replication is usually constrained by the bandwidth between the primary and recovery sites. Site Recovery Manager is typically used with asynchronous replication. The following two types of replication are used in Site Recovery Manager deployments.

- Storage replication – This is array-based replication provided by the storage vendor. It requires the same type and vendor—EMC SRDF or NetApp Snapmirror, for example—of the storage solution on both the primary and recovery sites. These solutions are very mature, have been used over the past few decades, and provide granular features and robust recovery mechanisms. Storage replication with Site Recovery Manager has a storage replication adapter (SRA) that is provided by the storage vendor. Site Recovery Manager communicates with the storage array via this adapter and uses it as a proxy for replication via the array.
- vSphere Replication – vSphere Replication is a VMware solution that can replicate at the individual VM level. The storage backing the VMs can be of any type, including local storage. The primary and recovery sites can have different types of storage. vSphere Replication operates at the VM level and can replicate all the VM disks or a chosen subset of disks. It occurs at the VMware kernel level, with any changes to storage captured and replicated. The RPO for vSphere Replication can range from 15 minutes to 24 hours, based on customer requirements. vSphere Replication requires appliances that are deployed by Site Recovery Manager. These appliances coordinate the replication of changes between the primary and the recovery site.

Both of these replication methodologies use crash-consistent recovery. The application has not been quiesced, so the recovery is similar to that of a machine after a power outage. There is little probability of data corruption for database-type workloads. In a later section in which we discuss individual BCAs, we will look at how application-level replication can be used to protect against the risk of data corruption.

Replication Configuration

vSphere Replication is leveraged in this DR deployment exercise and is configured to individually replicate each VM. Figure 18 shows the replication configuration for the Oracle VM. The selected RPO is 15 minutes, and the target location in the DR site is specified.

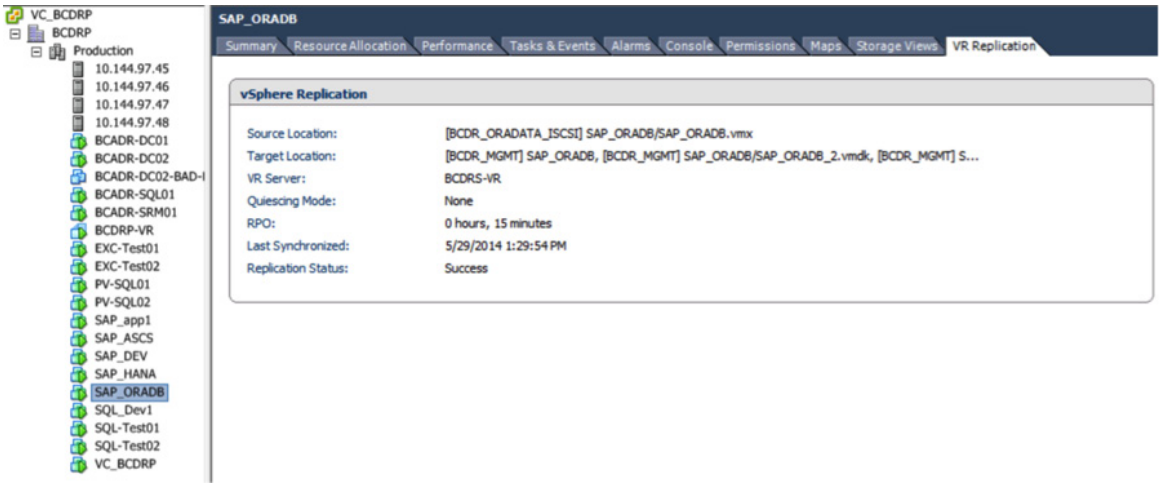

Figure 18. vSphere Replication Configuration for the SAP Database Oracle Backend

Similarly, replication is set up for all VMs belonging to applications that must be protected. As part of the vSphere Replication process, a relationship is established between the vSphere Replication servers in the primary and the recovery sites. An initial synchronization is run for all VMs and their disks. After it has completed, the replication occurs periodically, based on the RPO setting. The status of all replicated VMs in the recovery site can be observed through the Site Recovery Manager interface, as is shown in Figure 19.

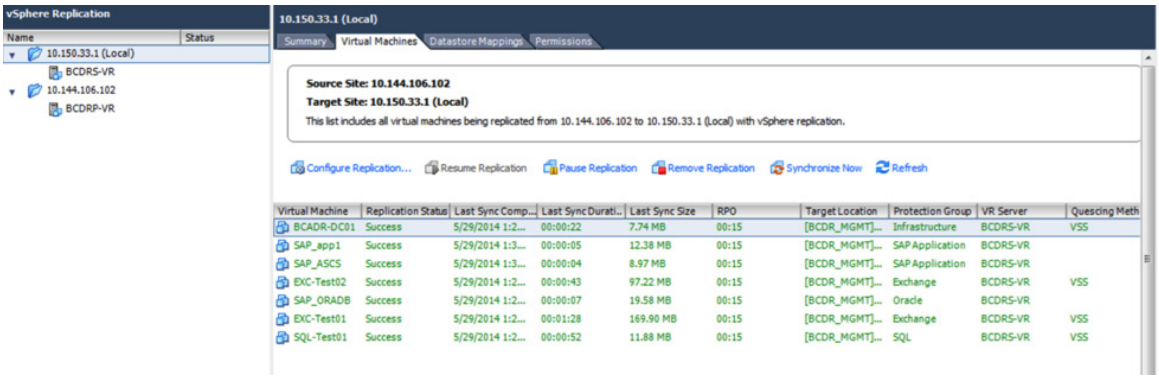

Figure 19. vSphere Replication Status

Protection Groups

A protection group is a group of VMs that fail over together to the recovery site. It contains VMs whose data has been replicated by array-based replication or by VR. It also typically contains VMs that are related in some way. The following are two examples:

- VMs for an SAP application module (application servers, central services, database server)
- VMs whose disk files are part of the same datastore group

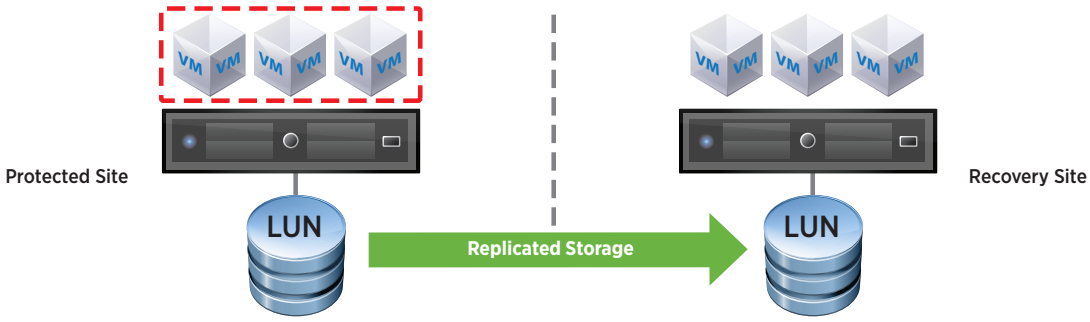

Figure 20. Protection Groups

In our use case, we have various protection groups, as is shown in Figure 21 along with the composition of the SAP application protection group.

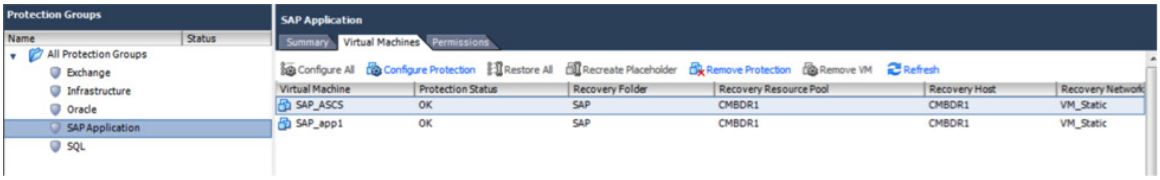

Figure 21. SAP Application Protection Group

As is shown, individual protection groups for SAP application and Oracle Database were created.

Recovery Plans

Protection groups are the building blocks of recovery plans. A protection group can be included in multiple recovery plans. A recovery plan is a series of steps executed to recover VMs in a specified sequence and priority.

Figure 22. Recovery Plans for Various Applications and for All

Each SAP application module can have its own recovery plan and can be recovered independently of others. Most applications have dependencies on certain infrastructure components such as Microsoft Active Directory and DNS. The infrastructure protection group must be included in most application recovery plans to ensure that the application is usable after recovery. A recovery plan for the entire site ("All") can include all applications. The recovery plan provides the capability to prioritize various applications to create a specified order for the recovery process.

The recovery plan for the SAP application is shown in Figure 23. The SAP application requires Oracle Database in addition to the application servers. Therefore, the two protection groups are included in the plan.

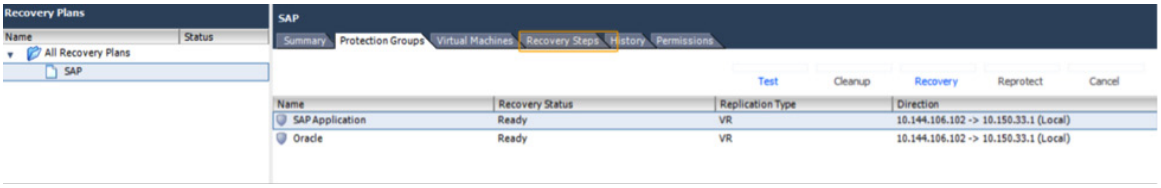

Figure 23. Recovery Plan for the SAP Application

IP Addressing at the Recovery Site

The option exists for the recovery site to use the same IP address as that of the primary site or a completely different one. This is dependent on the customer infrastructure.

- Some customers leverage capabilities such as stretched VLANs or relocatable VLANs to have the same IP address in the primary site and the recovery site.
- It is quite common to have a completely different set of IP addresses at the recovery site for the VMs. Site Recovery Manager recovery plans provide the capability to automatically re-IP the VMs before recovery.

Figure 24 shows an example IP mapping for the SAP Oracle VM at the recovery site.

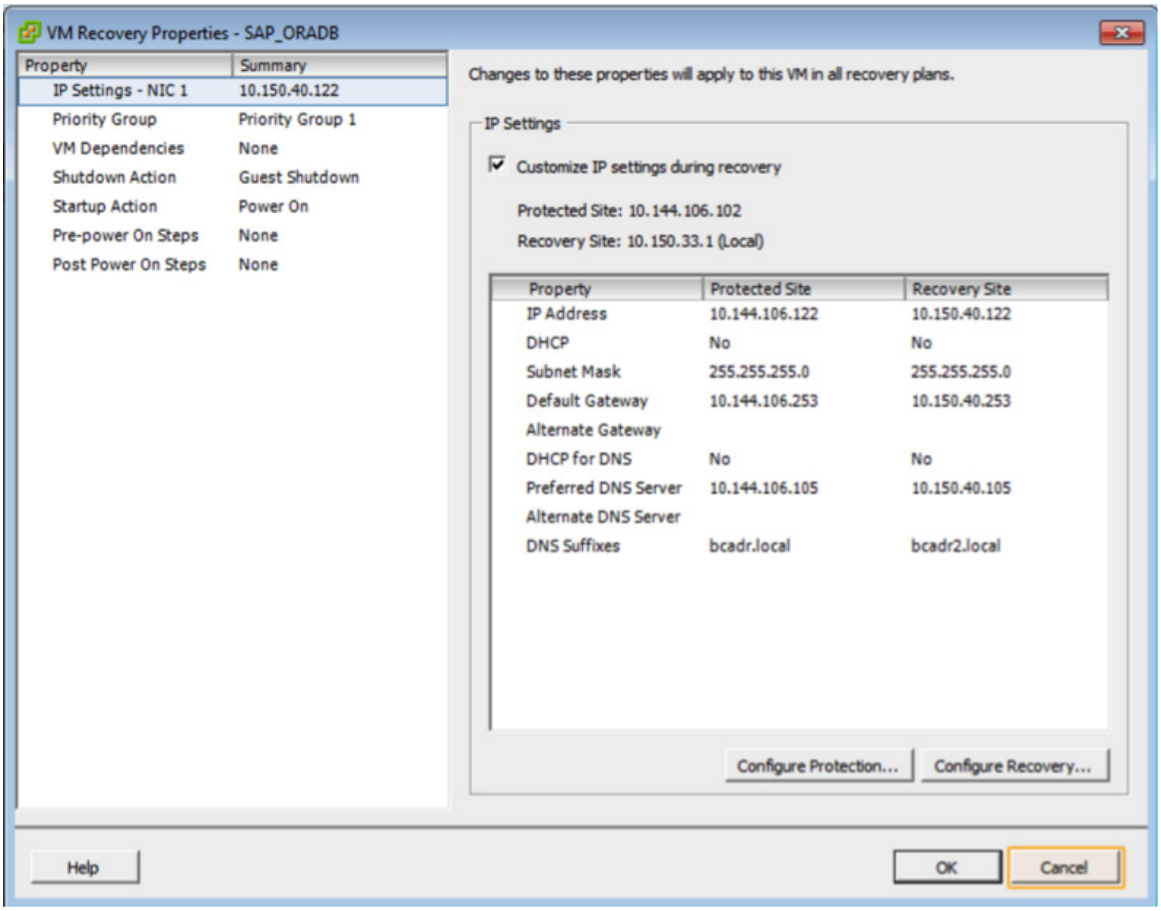

Figure 24. IP Customization During Recovery

Running an SAP Disaster Recovery Test with Site Recovery Manager

One of the major advantages of Site Recovery Manager is the capability to test DR of BCAs with no impact on the running production applications. This is a big differentiator for Site Recovery Manager as compared to DR for legacy physical environments. Site Recovery Manager provides the ability to test recovery in an isolated environment without any downtime and helps tune the recovery and learn from the tests. When an actual disaster occurs, there are very few surprises and the RTO is truly optimized. The processes for testing and actual recovery are identical. The only difference is that when a DR actually occurs with the primary site's being down, a test recovery takes place in isolation with the primary site's being up.

The environment that was set up with SAP ERP modules was tested with Site Recovery Manager. DR tests were executed by running recovery plans that had been set up for the various applications. Figure 25 shows the testing process for the SAP recovery plan. For an actual recovery, the recovery use case is run instead of the test use case. There are multiple steps in the recovery. Some are done in parallel; others are done based on priority. This provides optimization of the recovery time.

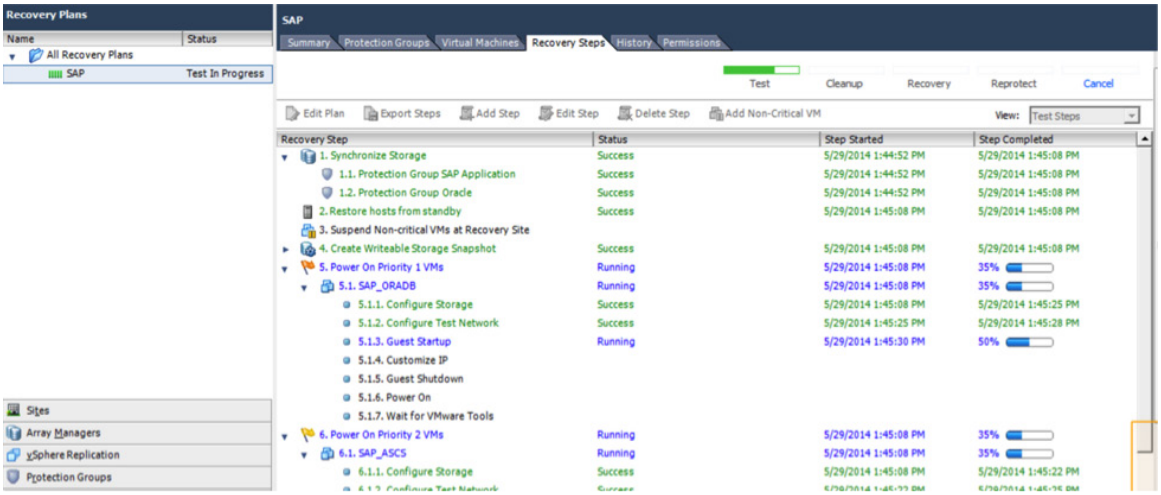

Figure 25. SAP Recovery with Site Recovery Manager

When the recovery is complete, reports are available that provide a view of all the steps and their duration. These history reports are invaluable for the tuning of the recovery process.

After testing, Site Recovery Manager makes it seamless to clean up the test environment through the use of the Cleanup function. The temporary storage, VM, and networking used for the tests are cleaned up in just a few minutes.

For an actual disaster and recovery, a reprotect process—rather than a cleanup—is initiated when the primary site infrastructure is available. The reprotect process reverses the direction of replication and the roles of the primary and recovery sites.

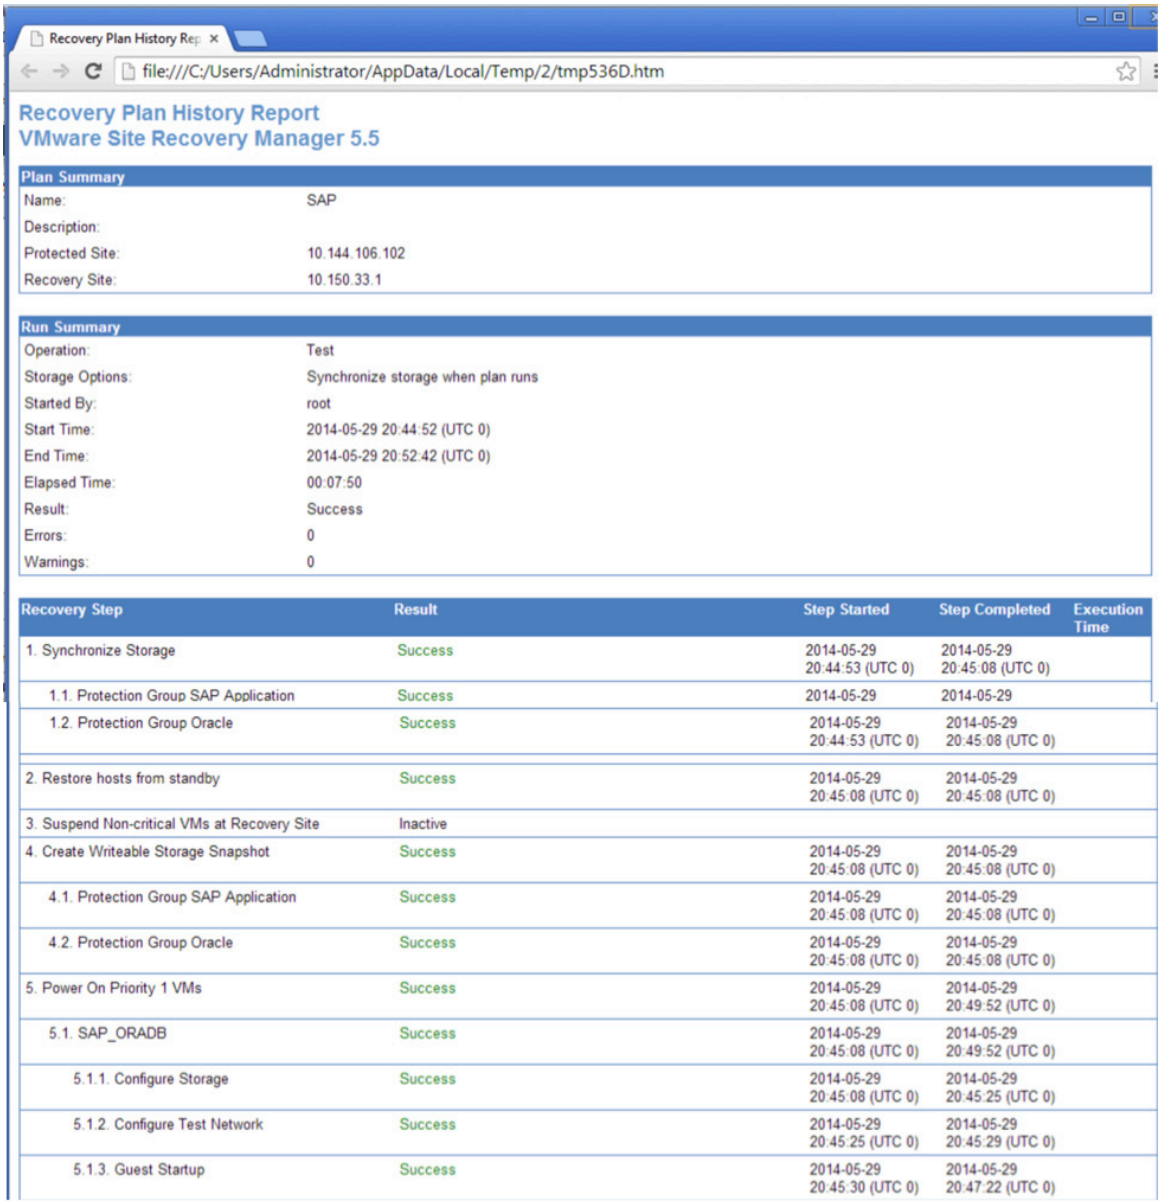

Recovery Plan History Rep ×

file:///C:/Users/Administrator/AppData/Local/Temp/2/tmp536D.htm

Recovery Plan History Report
VMware Site Recovery Manager 5.5

Plan Summary

Name:	SAP
Description:	
Protected Site:	10.144.106.102
Recovery Site:	10.150.33.1

Run Summary

Operation:	Test
Storage Options:	Synchronize storage when plan runs
Started By:	root
Start Time:	2014-05-29 20:44:52 (UTC 0)
End Time:	2014-05-29 20:52:42 (UTC 0)
Elapsed Time:	00:07:50
Result:	Success
Errors:	0
Warnings:	0

Recovery Step	Result	Step Started	Step Completed	Execution Time
1. Synchronize Storage	Success	2014-05-29 20:44:53 (UTC 0)	2014-05-29 20:45:08 (UTC 0)	
1.1. Protection Group SAP Application	Success	2014-05-29	2014-05-29	
1.2. Protection Group Oracle	Success	2014-05-29 20:44:53 (UTC 0)	2014-05-29 20:45:08 (UTC 0)	
2. Restore hosts from standby	Success	2014-05-29 20:45:08 (UTC 0)	2014-05-29 20:45:08 (UTC 0)	
3. Suspend Non-critical VMs at Recovery Site	Inactive			
4. Create Writeable Storage Snapshot	Success	2014-05-29 20:45:08 (UTC 0)	2014-05-29 20:45:08 (UTC 0)	
4.1. Protection Group SAP Application	Success	2014-05-29 20:45:08 (UTC 0)	2014-05-29 20:45:08 (UTC 0)	
4.2. Protection Group Oracle	Success	2014-05-29 20:45:08 (UTC 0)	2014-05-29 20:45:08 (UTC 0)	
5. Power On Priority 1 VMs	Success	2014-05-29 20:45:08 (UTC 0)	2014-05-29 20:49:52 (UTC 0)	
5.1. SAP_ORADB	Success	2014-05-29 20:45:08 (UTC 0)	2014-05-29 20:49:52 (UTC 0)	
5.1.1. Configure Storage	Success	2014-05-29 20:45:08 (UTC 0)	2014-05-29 20:45:25 (UTC 0)	
5.1.2. Configure Test Network	Success	2014-05-29 20:45:25 (UTC 0)	2014-05-29 20:45:29 (UTC 0)	
5.1.3. Guest Startup	Success	2014-05-29 20:45:30 (UTC 0)	2014-05-29 20:47:22 (UTC 0)	

Figure 26. History Report for SAP Recovery

Conclusion

Based on customer requirements, multiple high-availability options and deployment modes are available for virtualized SAP environments. VMware vSphere High Availability, VMware vSphere Fault Tolerance, guest failover clustering, and database availability solutions can be leveraged for SAP availability in local and multisite configurations. VMware vCenter Site Recovery Manager radically enhances the capability of businesses to protect business-critical applications from disasters. In addition, Site Recovery Manager enables customers to migrate applications temporarily or permanently between data centers, with minimal downtime and impact to business. The offline testing, workflow, and runbook capabilities provided by Site Recovery Manager are indispensable for businesses seeking to minimize their RTO and RPO for their business-critical applications.

References

1. *vSphere 6.0 Availability Guide*
2. Overview of AlwaysOn Availability Groups
3. Introduction to Oracle Data Guard
4. SAP HANA Replication
5. Site Recovery Manager Features

About the Authors

Mohan Potheri is currently a senior solutions architect at VMware, focusing on the virtualization of business-critical applications. He has more than 20 years of experience in IT infrastructure with VMware virtualization, enterprise UNIX, and business-critical applications. Mohan is a CISSP and VMware Certified Design Expert (VCDX#98). He holds master's degrees in electrical engineering and business administration from the University of Houston. Follow Mohan @ITVista on Twitter and on the VMware vSphere Blog.

Vas Mitra is an SAP solutions architect in the VMware Alliances organization. He has been working with VMware partners to develop SAP environments on VMware solutions and best practice guidelines for deploying SAP software on VMware infrastructure. He also works with the VMware sales organization on presales architecture design of customer implementations of SAP software on VMware ESXi. Prior to joining VMware, Vas worked on numerous SAP projects since 1993 as an SAP presales architect, basis administrator, and ABAP programmer.

Acknowledgment

Our sincere thanks to Duncan Epping for his detailed review and insight.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com