SOLUTION
READINESS

SAP® Solutions on VMware® Business Continuance

Protecting Against Unplanned Downtime

VMware, Inc
3401 Hillview Ave
Palo Alto, CA 94304
www.vmware.com

Contents

1. Introduction

Business continuance (sometimes referred to as *business continuity*), describes the processes and procedures an organization puts in place to ensure that essential functions can continue in case of unplanned downtime. Unplanned downtime refers to an outage in system availability due to infrastructure failure (server, storage, network), or site disaster. SAP® products and solutions provide mission-critical business processes that need to be highly available even in a site disaster.

This paper describes three different high availability scenarios designed to protect the SAP single-points-of-failure. The architectures are based upon VMware® High Availability features (VMware Fault Tolerance and VMware High Availability) and third-party clustering software (Microsoft Cluster Server) in a virtualized environment. Factors that influence the final design are discussed. Finally, an SAP disaster recovery architecture based on VMware Site Recovery Manager is described. For background and more detail about these VMware functions and products see refer to the documents in the Resources section.

Though planning for business continuance of SAP implementations is part of a system-wide strategy, this document does not cover high availability features of network and storage. Consult the appropriate VMware and VMware partner guides for information on these topics.

2. SAP Distributed Architecture

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 has a scalable, fault-tolerant, multi-tier architecture, components of which can be protected either by horizontal scalability (e.g., NetWeaver application servers) or by cluster and switchover solutions that protect the single points of failure in the SAP architecture. The SAP single points of failure are identified and explained in Appendix A—they include the database, message and locking services. The latter two are included in constructs referred to as the Central Instance (CI) or ABAP SAP Central Services (ASCS).

3. Protection with VMware High Availability

VMware High Availability (HA) continuously monitors all VMware ESX™ hosts in a cluster and detects hardware failures. The VMware HA agent placed on each host maintains a heartbeat with the other hosts in the cluster using the service console network. Each server sends heartbeats to the other servers in the cluster at regular intervals. If any servers lose heartbeat, VMware HA initiates a failover action of restarting all affected virtual machines on other hosts.

Figure 1 depicts a typical scenario with the SAP database and Central Instance running in a single virtual machine with VMware HA applied, a configuration often found in existing installations. The table summarizes the features of this configuration.

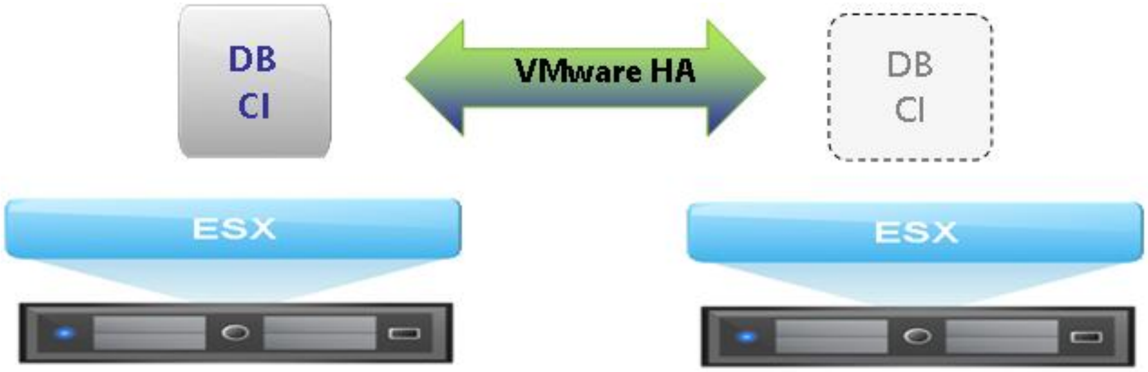

SAP Config VMware HA	High Availability Features	Comments
• 2-tier or 3-tier – app server VMs not shown • protected via VMware HA	• Protection against server failure • Auto restart of VMs • Startup scripts/service required to auto-start SAP/DB instances in guest OS • VMware HA easy to configure (VMware “out-of-the-box”)	• No monitoring of application • DB unavailable during failover • No enqueue and message services during CI failover

Figure 1. VMware HA Configuration for SAP

4. Clustering with MSCS in Virtual Machines

The Microsoft ® Cluster Server (MSCS) configuration is the standard switchover solution for SAP systems running on Windows platforms. In this setup, the SAP system is installed on two nodes of a cluster. Under normal operation, the SAP central services run on one node and the database runs on the other node of the cluster. If one of the nodes fails, the affected central service or database instance is automatically moved to the other node, preventing downtime. This configuration is documented in the SAP installation guide, NetWeaver 7.0 ABAP on Windows: MS SQL Server.

VMware supports the use of MSCS software in virtual machines, and clustering virtual machines with MSCS can reduce hardware costs of traditional high availability clusters.

VMware vSphere™ is supported with MSCS. You can find technical details of MSCS and vSphere in the VMware technical guide, Setup for Failover Clustering and Microsoft Cluster Service.

The installation of SAP with MSCS by way of SAP install shield "sapinst" follows the same process as on a physical system. Each MSCS cluster node is a virtual machine and the resulting architecture is similar to that described in the SAP installation guide for Windows, as shown in Figure 2 ("REP ENQ" in the diagram stands for replicated enqueue server). The table outlines the features of this configuration.

The enqueue replication server contains a replica of the lock table (replication table) and behaves exactly the same way as in physical implementations. In normal operation the replication enqueue server is always active on the virtual machine where the ASCS is not running. If an enqueue server in an MSCS cluster with two nodes fails on the first MSCS node, the enqueue server fails over to the second MSCS node and starts there. It retrieves the data from the replication table on that node and writes it in its lock table. The enqueue replication server on the second MSCS node then becomes inactive. If the first MSCS node is available again, the enqueue replication server on the second MSCS node becomes active again.

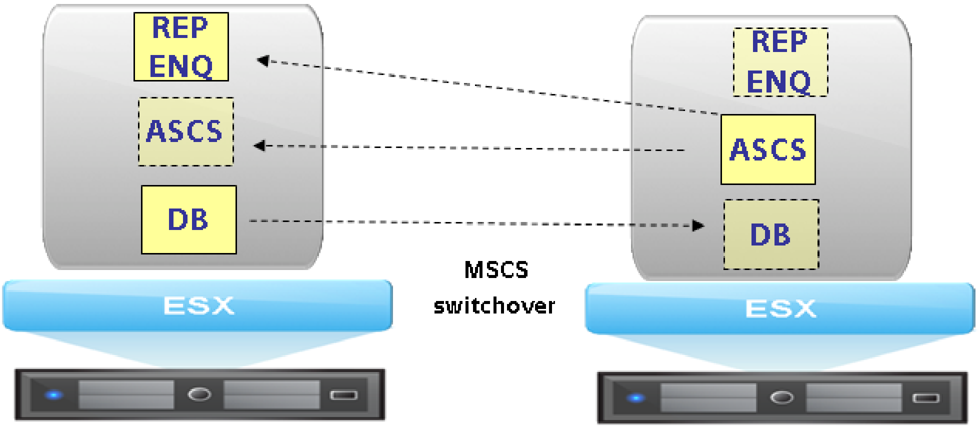

SAP Config MSCS in VMs	High Availability Features	Comments
• Assumes 3-tier - app server VMs not shown • Protected via MSCS agents for DB, ASCS, replicated enqueue	• Protection against server failure plus monitoring of DB and ASCS • Auto-restart of SAP services • Continuous availability of SAP locks due to replicated enqueue • No VM and guest OS boot required during failover • Planned downtime (e.g., guest OS patching) can be minimized by evacuating MSCS resources to the other node	• Unable to VMotion MSCS VMs • HA and DRS functionality must be disabled for individual MSCS VMs (ESX host and other VMs can still be part of HA/DRS cluster). For details please see the MSCS/ vSphere setup guide mentioned in the Resources section • DB and message service unavailable during failover • MSCS skills required, more complex setup

Figure 2. SAP and MSCS High Availability Configuration in Virtual Machines

5. SAP Central Services and VMware Fault Tolerance

Fault Tolerance (FT) relies on VMware vLockstep technology to establish and maintain an active secondary virtual machine that runs in virtual lockstep with the primary virtual machine. The secondary virtual machine resides on a different host and executes exactly the same sequence of virtual (guest) instructions as the primary virtual machine. The secondary observes the same inputs as the primary and is ready to take over at any time without any data loss or interruption of service should the primary fail. Both virtual machines are managed as a single unit, but run on different physical hosts. By allowing instantaneous failover between the two virtual machines, FT enables zero downtime for the application deployed within the virtual machine.

Currently VMware FT supports only one virtual CPU and, as such, is a good candidate for the "lighter" central services component. A high availability configuration of the SAP database and ASCS is shown in Figure 3. The table summarizes the features of this setup. The database virtual machine is protected by VMware HA and the ASCS virtual machine by VMware FT.

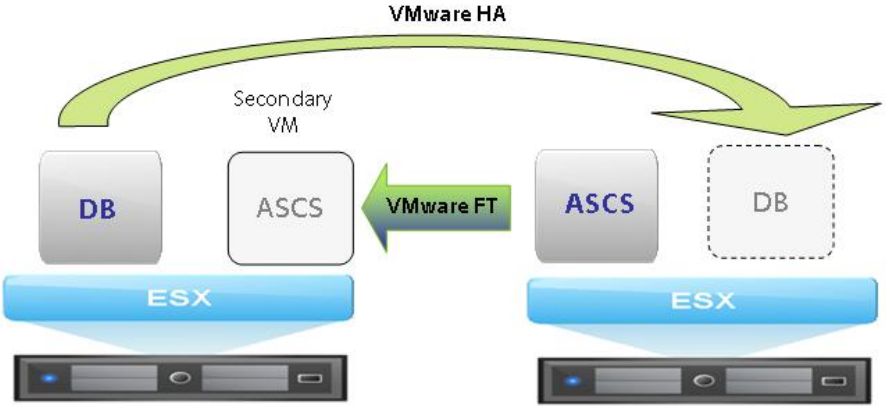

SAP CONFIG VMware FT for ASCS	High Availability Features	Comments
• Assumes 3-tier - app server VMs not shown • ASCS protected via VMware FT, DB protected via VMware HA	• Protection against server failure • Continuous availability of Central Services • Easy to configure – VMware “out-of-the-box” • DB still protected via VMware HA • New secondary ASCS VM automatically created after failover (assumes more ESX servers available) • 1 x vCPU VM for ASCS enough to support the smaller sized SAP environments (SAP Competency Center of server vendor can provide guidelines)	• No monitoring of application • VMware FT currently supports 1 x vCPU VM • DB unavailable during failover • Separate NIC/network recommended for FT logging traffic

Figure 3. SAP Central Services and VMware FT

The configuration shown in Figure 3 was installed in VMware labs and a small-scale functional test was conducted to verify continuous availability of central services during failover of the ASCS virtual machine protected via VMware FT. The results are described in Figure 4.

Test Setup	Results with failover
• 2 x ESX servers running vSphere • 1 x VM, 2 x vCPU, 8GB RAM running ECC 6.0: MSSQL database; Windows Server 2003 64-bit; dialog instance • 1 x VM, 1 x vCPU, 1GB RAM running ASCS protected by VMware FT 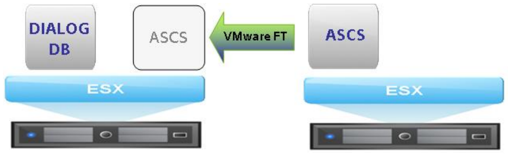	• 150 concurrent users < 0.5 sec response time (users generated by SD Benchmark Kit) • Successful completion of workload with no user or lock errors. • average CPU utilization of ASCS VM < 5% **Note**: This setup was not intended or tuned for benchmarking.

Figure 4. Lab Results – VMware FT Test with VM Running ASCS

The VMware hardware partner competency centers for SAP can provide further guidelines for determining the sizing of this distributed architecture. For technical details of VMware FT see the document *Protecting Mission-Critical Workloads with VMware Fault Tolerance*.

6. SAP High Availability Options and Uptime Discussion

The previous sections described three high availability scenarios based on VMware HA/FT and MSCS in virtual machines—the following discusses and summarizes their features.

Both MSCS (with virtual machines) and VMware HA/ FT provide protection against hardware failure. The essential differences between these scenarios are:

- Complexity of setup –VMware HA and FT are simple to set up and configure compared to deploying clustering software such as MSCS. For example, zero-downtime protection against hardware failure for Central Services is possible with VMware FT without the complexity of configuring replicated enqueue in a clustered environment.
- Cost – Clustering is often a costly solution because it requires dedicated standby servers, expensive OS licenses, and advanced technical skills on-site. VMware HA is significantly more cost-effective because it doesn’t require dedicated standby servers and is much simpler to configure.
- Microsoft supports the use of MSCS with VMware virtual machines on ESX. However, this support only applies when the virtual machines are not migrated with VMotion or protected with VMware HA or Fault Tolerance. VMware customers have successfully deployed MSCS with the use of VMotion and VMware HA, but this is not an officially supported configuration.
- Monitoring of the SAP application software – MSCS monitors the health of the database and central services instance, whereas VMware HA and Fault Tolerance only protect against hardware failures and in some cases against OS failures.
- Reduction in planned downtime – the guest OS of clustered virtual machines can be patched with minimal downtime to the SAP services. Resources from a clustered node can be moved to the other node to allow for OS patching.

The following considerations can impact the high availability design:

- Customers past experience in running SAP applications – how often has only the database or central services component failed at the application level (and hardware was not the cause of failure)?
- Some customers may not prefer automatic restart of the application in the event of an application error only.
- While clustering solutions have proven to provide higher uptime for SAP installations in the last 10+ years in physical environments, there is now a substantial list of SAP on VMware production deployments where SAP customers have achieved their required uptime SLAs with VMware HA (and vMotion) used as a substitute for MSCS.
- Though it may be generally considered that clustering software provides the higher degree of uptime measured as “four nines,” this level of uptime cannot realistically be achieved with clustering software alone as overall system availability also depends on redundancy designed into the other parts of the infrastructure (network, power, storage, etc.).
- What is the business cost/uptime trade-off? Clustering is the more expensive solution, and VMware is the more cost-efficient solution. Therefore, the business’ willingness to incur additional costs for increased availability is a key consideration.
- What is a customer’s Service Level Agreement (SLA) with respect to uptime/downtime, or how much downtime is a business willing to tolerate?

The following table shows the translation from a given availability percentage to the corresponding amount of time a system would be unavailable per year.

Table 1. Availability Percentages Calculated on a Yearly Basis

Availability %	Downtime per Year
99	3.65 days
99.9 (three nines)	8.76 hours
99.99 (four nines)	52.6 minutes
99.999 (five nines)	5.26 minutes

Depending on duration of manual procedures for a customer to recover from potential SAP application failures (i.e., corruption of the database and/or central services), it may be possible to achieve “three nines” availability with VMware HA/FT. Clustering software typically can provide the “four nines” degree of availability. Both these cases are anecdotal approximations and assume other parts of the infrastructure have adequate redundancy designed into them. A more accurate assessment is possible with customer-specific data as planned downtime (for patch updates, SAP transports, SAP upgrades, etc) needs to be factored into the overall availability calculation. For example, a major SAP release upgrade may impact the final availability percentage.

The final design choice between using or not using clustering software in a VMware virtualized deployment depends on how much realistic downtime a business can tolerate and the cost they are willing to invest in the extra resources and skills to install and operate clustering software—it is a tradeoff.

7. VMware vCenter Site Recovery Manager

VMware vCenter™ Site Recovery Manager (SRM) 4.0 provides business continuity and disaster recovery protection for virtual environments.

Disaster recovery testing comprises a logistical plan for how an organization will recover and restore partially or completely interrupted critical function(s) within a predetermined time after a disaster or extended disruption. Common wisdom states that any disaster recovery plan is only as good as the last (successful) test. Disaster recovery efforts can fail because the IT team neglects to test often, the result being an insurance policy that does not pay off when disaster hits.

Disaster recovery testing is often difficult because it is usually very disruptive, expensive in terms of resources and extremely complex. By leveraging virtualization, SRM addresses this problem while making planning and testing simpler to execute.

Using SRM, two sites are involved—a protected site and a recovery site. SRM leverages array-based replication between a protected site and a recovery site to copy virtual machines.

Recovery Point Objective (RPO) and Recovery Time Objective (RTO) are the two most important performance metrics IT administrators need to keep in mind while designing and executing a disaster recovery plan. RPO is addressed by the storage provider who provides certified storage replication adapters that integrate with SRM to enable a fully automated test or real recovery. RTO is addressed by the SRM recovery plans that automate the startup sequence of multiple virtual machines that comprise a virtualized SAP landscape and automate network connectivity at the remote site.

7.1 Site Recovery Manager Architecture

An SAP landscape can consist of a considerable number of separate systems to host the multiple SAP products, each with separate production and non-production systems. In the production environment, multiple SAP systems typically interface to a myriad of third-party bolt-on applications. In addition, the multi-tier architecture of Netweaver may result in separate tiers of application and database servers. Hence, a fully virtualized SAP environment results in numerous virtual machines with data interfaces/flows between these virtual machines. Such a volume of virtual machines can be managed with the workflow features of SRM that process the correct sequence and order of recovery of virtual machines after a site failure.

Figure 5 shows the architecture of a deployment of a virtualized SAP landscape with SRM. In this example, production SAP systems are replicated from the protected to a recovery site. Each site hosts a separate storage array. Customer-specific business requirements determine if non-production systems also need to be replicated and protected against site failure. Other non-production SAP systems are hosted at the protected site. The SAP landscape is logically depicted here by three SAP systems for simplicity, each of which is connected via interfaces to demonstrate that business processes can traverse separate systems (the actual landscape would have more SAP systems and third-party bolt-on applications).

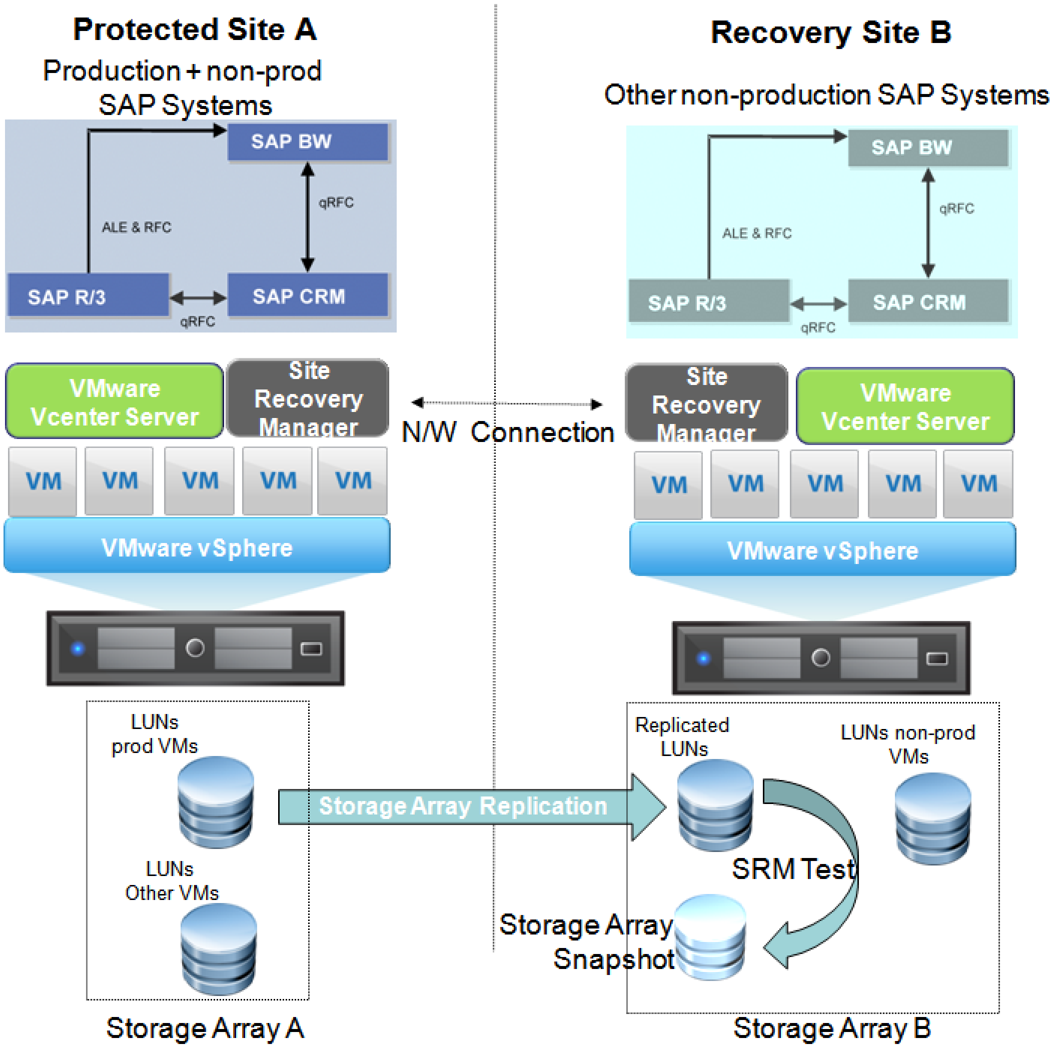

Figure 5. Example Deployment of SAP Landscape with Site Recovery Manager

Overview of the architecture:

- ESX hosts at the recovery site run some non-production systems to maximize resource usage; these servers do not need to be idle. Another scenario (not shown here) based on two-way storage array replication is feasible with SRM whereby production systems can be split between the two sites and each can be acting as a failover to the other.
- A SRM server is installed both at the protected and recovery site. Both sites are managed by their own vCenter Server. The SRM Server operates as an extension to the vCenter Server and the SRM user interface installs as a vSphere client plug-in.
- A certified storage array vendor is required that has an adapter which integrates with Site Recovery Manager. See the Resources section for a list of certified storage products. Storage array replication needs to be correctly installed and configured for Site Recovery Manager to operate. This should follow the same process as the physical environments. Site Recovery Manger automatically detects the replicated LUNs that contain virtual machines.
- The protected and recovery sites should be connected by a reliable IP network. Storage arrays might have additional network requirements for replication. The SRM servers at both sites communicate with each other during normal operations.
- On the protected site, production virtual machines are replicated via storage array replication. On the recovery site (storage array B), the replicated LUNs are not visible to the ESX hosts.

7.2 Executing Recovery Plans

- Protection groups are created on the protected site. A protection group is a collection of virtual machines that all use the same set of replicated LUNs and failover together.
- Recovery plans are created at the recovery site and are created from the protection groups. The recovery plan is essentially an automated runbook that consists of a set of steps that control what happens during a failover.
    - The recovery plan determines the order of production virtual machine startup during a failover and also can suspend non-production virtual machines already running at the recovery site. Enough server resources are required at the recovery site to run the production systems, as well as any non-production systems that are also needed to run per business requirements (otherwise, non-production systems can be suspended by the recovery plan).
    - Call outs to custom scripts can be included in the recovery plan for customer specific requirements.
    - The SAP application can be configured to auto start after a guest OS boot within the virtual machine.
    - The execution of the recovery plans enable customers to achieve faster RTO.

A recovery plan can be executed in either of two modes:

- Actual failover – array replication is halted and the replicated LUNs on the recovery site are enabled for read and write capabilities, and SRM initiates the power up of the virtual machines in the recovery site according to the startup order in the recovery plan. SRM does not automatically detect a site disaster—recovery has to be manually started via the SRM user interface at the recovery site.
- Test failover – the replicated LUNs on the recovery site still remain unavailable to the ESX hosts. They are copied using storage array snapshot functionality and these copied snapshot LUNs are presented to the ESX hosts. The snapshot is reasonably quick as data is not duplicated (this is part of the storage array feature). The production virtual machines are started according to the recovery plan and can then be user tested. After testing is complete, a manual step to continue via the SRM user interface stops the production virtual machines and removes the storage array snapshot. Meanwhile, any suspended non-production virtual machines are started again on the protected site. During this test cycle the replicated LUNs are still being refreshed per the storage array replication schedule, and production systems continue to function normally on the protected site.

The recovery plan test simulates an actual recovery as it performs the same sequence of actions to recover the production SAP systems. It has the advantage of being run as frequently as required and demonstrates an enormous business benefit of being able to test a disaster recovery plan on demand to satisfy any auditing requirements.

7.3 Network Customization

Typically there are separate networks at the protected and recovery sites. While each network should be connected via routers, the subnet and IP address will differ between the locations. Therefore, when performing site failover, IT administrators can be faced with the following challenges on the recovery site:

- Network properties of the production virtual machines need to be customized according to the network specification of the recovery site.
- Domain Name Server (DNS) records pertaining to these virtual machines need to be updated.

After failover to a disparate network, the network properties of virtual machines such as IP addresses, Gateway and DNS domain all need to change to return to a functional state. SRM addresses this at the recovery site via the following features:

- Customization Specification Manager – this allows administrators to create a custom network specification for each production virtual machine that is replicated from the protected site. Network properties (IP address, gateway, etc) can be assigned to the virtual machine so that when it starts up in a recovery plan it will function correctly on the recovery site network. The hostname of the guest OS in the virtual machine needs to remain the same so as not to impact the SAP application (SAP instance files once installed have the hostname of the OS in various configuration and startup files, but IP address is not hard coded in the files).
- After changing the IP addresses of virtual machines, DNS records of the virtual machines need updating.

These features are covered in detail in the document, Automating Network Setting Changes and DNS Updates on Recovery Site Using VMware vCenter Site Recovery Manager.

7.4 Storage Array Replication

As previously mentioned, the SRM solution requires storage array tools to replicate the LUNs from the protected to the recovery site. Storage array replication needs to be installed and configured in the same manner as in physical environments, and administrators should follow guidelines from their storage vendor. Similarly, SAP database LUN layout on the storage array should follow the same recommendations as for physical environments. The major storage array vendors have SAP practices that have developed best practice guidelines for LUN layouts of SAP databases and how they should be replicated between separate sites in a disaster recovery scenario. The same guidelines should be followed with SRM. For example:

- Best practice for I/O performance requires production database virtual machines not to be shared with other virtual machines.
- Where applicable, some storage array vendors may prefer the use of RDMs as they are compatible with their disaster recovery tools. In these cases the virtual machine guest OS drive ("root" or "C:\") would be VMFS format and the database datafiles would be RDM-based.

The RPO objective is managed by the storage array replication schedule. The frequency of replication and subsequent cost with respect to bandwidth requirements over a long distance is managed by the storage vendor specifications and is balanced against the business requirements. Two broad replication methods are available from storage vendors that impact RPO, and in both cases SRM does not manage the consistency of the SAP database during replication (quiescing of the database). This is addressed by the storage vendor technology or by separate procedures:

- Synchronous replication – guarantees zero data loss, where a write either completes on both sides or not at all. The storage vendor technology typically guarantees consistency of the database that is spread across multiple LUNs.
- Asynchronous replication – write is considered complete as soon as local storage acknowledges it. The remote storage is not guaranteed to have the current copy of data. A potential scenario to guarantee database consistency in this situation involves putting the database into online backup mode before replicating. On a separate, more frequent schedule, replicate the database log files. Database recovery then involves starting the database and applying logs to roll forward the database. Such a process may be created manually or be part of tools/products from the storage vendor.

8. Conclusion

SAP software solutions enable a variety of mission-critical business functions such as sales order entry, manufacturing and accounting that depend on the availability of IT services. The consequence of a failure to meet the business demands can be costly and require an investment in infrastructure that is designed for high availability to protect against failures within the datacenter as well as against events that may cause a site disaster. Such failures result in unplanned downtime.

Architectural scenarios were described showing how the SAP SPOFs can be protected against hardware failure with VMware HA and FT, or with MSCS clustering software in virtual machines. While MSCS requires a more complex setup it provides the highest degree of monitoring by additionally checking the health of the SAP application software. An anecdotal assessment estimates that it may be possible to achieve “three nines” availability with VMware HA and FT, assuming adequate manual procedures are in place to recover from software corruption of the SAP SPOFs. MSCS configurations typically can achieve “four nines” availability, but this is also anecdotal. Note that both these high availability specifications require redundancy designed into other parts of the infrastructure (e.g., network, storage and power).

Designing a highly available SAP system on VMware vSphere requires a tradeoff between the level of downtime that can be tolerated (which in turn has a business cost), and the complexity of the setup which has a cost with respect to skills and IT resources. So organizations need to determine their realistic requirements for availability.

The architectural deployment of an SAP landscape with VMware Site Recovery Manager was described which provides an automated disaster recovery and testing solution for SAP landscapes. Site Recovery Manager enables on demand and frequent testing of disaster recovery plans while having no impact to the production systems. This can help to satisfy internal audits and business compliance requirements. Site Recovery Manager can help to achieve the disaster recovery RPO and RTO priorities of organizations running SAP applications. RTO is addressed by recovery plans that automate the sequence of virtual machine recovery at the remote site, including network re-configuration. RPO is managed by the storage array that controls the frequency of replication to the remote site and manages the consistency of data. A successful SRM deployment requires a solid partner approach between the customer, VMware and the storage array vendor.

9. Resources

SAP Note 821904 - Separating SCS instances for ABAP and J2EE

SAP Note 524816 - Standalone enqueue server

SAP note 175047 - Causes for FI document number gaps

VMware HA: Concepts and Best Practices
http://www.vmware.com/files/pdf/VMwareHA_twp.pdf

Protecting Mission-Critical Workloads with VMware Fault Tolerance:
http://www.vmware.com/files/pdf/resources/ft_virtualization_wp.pdf

Setup for Failover Clustering and Microsoft Cluster Service (Update 1 ESX 4.0

ESXi 4.0 vCenter Server 4.0)
http://www.vmware.com/pdf/vsphere4/r40_u1/vsp_40_u1_mscs.pdf

Installation Guide: NetWeaver 7.0 ABAP on Windows: MS SQL Server
http://service.sap.com/instguides

VMware Site Recovery Manager
http://www.vmware.com/files/pdf/srm_datasheet.pdf

VMware vCenter Site Recovery Manager Administration Guide
http://www.vmware.com/pdf/srm_10_admin.pdf

VMware Site Recovery Manager Storage Partners
http://www.vmware.com/pdf/srm_storage_partners.pdf

Automating Network Setting Changes and DNS Updates on Recovery Site Using VMware vCenter Site Recovery Manager
http://communities.vmware.com/docs/DOC-11516

Appendix A – SAP Single Points of Failure

The following single points of failure exist in the SAP architecture:

- Database – Every ABAP work process makes a private connection to the database at the start, and 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 comes back up. User sessions in the middle of database activity receive SQL error messages, but their logged on sessions are preserved on the application server.
- SAP Message Service – The SAP Message Service is used to exchange and regulate messages between SAP instances in a SAP network. It manages functions such as determining which instance a user logs onto during client connect and scheduling of batch jobs on instances configured for batch.
- SAP Enqueue Service – The enqueue service manages 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 the transactions that contain locks have to be rolled back and any SAP updates being processed would fail (and potentially, depending on business requirements, would have to be manually reapplied via SAP transaction SM13 once the enqueue service is back up. See SAP note 175047 - Causes for FI document number gaps).

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

- Central Instance (CI) – comprises message and enqueue services in addition to other SAP work processes that allow execution of online and batch workloads. The CI is a SPOF as it includes message and enqueue.
- Central Services – in newer versions of SAP the message and enqueue processes have been separated from the CI and 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).
- 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 and enqueue service from the CI helps to 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.