

WHITE PAPER

SAP Solutions on VMware® vSphere™: High Availability



Table of Contents

| | |
|-----------------------------------------------------------------------|----------|
| Introduction..... | 1 |
| vSphere Overview | 1 |
| VMware Fault Tolerance | 1 |
| Background on High Availability for SAP Solutions..... | 2 |
| VMware HA | 3 |
| Clustering SAP Solutions in Virtual Machines with MSCS | 4 |
| SAP Central Services and VMware Fault Tolerance..... | 5 |
| SAP Database/Central Instance and VMware Fault Tolerance | 7 |
| Summary | 8 |
| Further Reading..... | 9 |

Introduction

VMware® has forged a strong, collaborative relationship with SAP AG to maximize the benefits for joint customers. VMware® Infrastructure 3, shipping since 2006, has proven to be a flexible and cost-effective industry-standard hardware virtualization suite for SAP solution implementations, helping customers to manage a smarter SAP application-based landscape. VMware® vSphere™ 4 is the new, renamed version of VMware Infrastructure, the cloud operating system from VMware. In this release, VMware vSphere provides enhancements and new features that greatly benefit deployments of SAP solutions on VMware virtualization software.

This document describes and compares different high availability scenarios available for SAP solutions running on vSphere. The new VMware® Fault Tolerance (FT) feature provides further protection to critical SAP application services with zero downtime, enhancing the high availability provided by the VMware® High Availability (HA) feature introduced in earlier releases. SAP solutions deployments are excellent candidates to take advantage of Fault Tolerance because they are mission critical, requiring maximum uptime and protection against single points of failure such as the database and SAP central services.

This document covers four high availability scenarios, each providing different degrees of fault tolerance and uptime for the SAP single points of failure (SPOF):

- VMware HA
- Use of Microsoft clustering software for SAP in virtual machines (scope here is restricted to one SAP system in a two-node cluster)
- VMware FT to protect SAP central services
- VMware FT to protect SAP database and central instance

The first two scenarios are solutions available with VMware Infrastructure 3 and the second two are based on the new VMware FT functionality in vSphere.

vSphere Overview

The latest major version of VMware Infrastructure, called VMware vSphere, is the industry's first cloud operating system and transforms datacenters into simplified environments to enable the next generation of flexible, reliable IT services. A cloud operating system is a new category of software that is specifically designed to holistically manage large collections of infrastructure – CPUs, storage, networking – as a seamless, flexible and dynamic operating environment.

You can find an overview of the new and enhanced technical features in vSphere here:

<http://www.vmware.com/products/vsphere/mid-size-and-enterprise-business/features.html>.

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.

vLockstep technology delivers architectural guarantees that the states of the primary and secondary virtual machines are identical at any point in the execution of instructions. vLockstep accomplishes this feat by having the primary and the secondary virtual machines execute identical sequences of x86 instructions.

For technical details of VMware FT see the document *Protecting Mission-Critical Workloads with VMware Fault Tolerance* and the VMworld breakout session presentation *VMworld Europe 2009 Session TA24, Fault Tolerance Performance Preview* in the further reading section.

Background on High Availability for SAP Solutions

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 scope of discussion and examples covered here are based on the following system considerations:

- Solutions built on an ABAP-based NetWeaver 04s system (but the concepts can be applied to ABAP+JAVA and JAVA based NetWeaver architectures).
- SAP clustering software from Microsoft (Microsoft Cluster Service - MSCS).
- Although high availability for SAP solutions should be part of a system-wide strategy, this document does not include high availability features for switches, routers, network, and storage within its scope. Consult the appropriate VMware and VMware partner guides for information on these topics.

The following single points of failure are identified 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 will 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 would 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, since 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 hence is much quicker to start up after a failure.

The next section describes four high availability scenarios with varying degrees of fault tolerance. The examples below consider the SAP SPOFs (database, CI and Central Services) - the SAP application instances are assumed to run in separate virtual machines and are not shown in the architectures below.

VMware HA

VMware 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 the 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 |
|--------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> 2-tier or 3-tier - app server VMs not shown protected via VMware HA | <ul style="list-style-type: none"> 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") | <ul style="list-style-type: none"> No monitoring of application DB unavailable during failover No enqueue and message services during CI failover |

Figure 1. SAP VMware HA Configuration

Clustering SAP Solutions in Virtual Machines with MSCS

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" (available from <http://service.sap.com/instguides>).


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.

vSphere is supported with MSCS (see "What's New in VMware vSphere 4.0" at http://vmware.com/support/vsphere4/doc/vsp_40_new_feat.html). You can find technical details of MSCS and vSphere in the following technical guide, http://vmware.com/pdf/vsphere4/r40/vsp_40_mscs.pdf.

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 below ("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). 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 |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> Assumes 3-tier - app server VMs not shown protected via MSCS agents for DB, ASCS, replicated enqueue | <ul style="list-style-type: none"> Protection against server failure + monitoring of DB and ASCS Auto restart of SAP services Continuous availability of SAP locks due to replicated enqueue No guest OS boot required during failover | <ul style="list-style-type: none"> Unable to VMotion VMs VMs cannot be part of VMware HA/DRS cluster For other technical caveats please see the MSCS/ vSphere setup guide mentioned in the "Further reading" section DB unavailable during failover Message Service unavailable during failover MSCS skills required / more complex setup |

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

SAP Central Services and VMware Fault Tolerance

VMware FT provides continuous availability for applications in the event of server failures, by creating a live shadow instance of a virtual machine that is in virtual lockstep with the primary instance. By allowing instantaneous failover between the two instances in the event of hardware failure, FT enables zero downtime for the application.

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 below. 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 |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> Assumes 3-tier - app server VMs not shown ASCS protected via VMware FT, DB protected via VMware HA | <ul style="list-style-type: none"> 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) | <ul style="list-style-type: none"> No monitoring of application VMware FT limited to 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 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 |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> 2 x ESX servers running pre-release build of vSphere 1 x VM, 2 x vCPU, 8GB RAM running ECC 6.0: MSSQL database; Windows Server 2003 64bit; dialog instance 1 x VM, 1 x vCPU, 1GB RAM running ASCS protected by VMware FT | <ul style="list-style-type: none"> 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: the 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.

SAP Database/Central Instance and VMware Fault Tolerance

While VMware FT currently supports one virtual CPU it is capable of protecting smaller complete SAP systems installed in a uni-processor virtual machine. The configuration and summary of features are shown in Figure 5.



| SAP CONFIG VMware FT for DB+CI | High Availability Features | Comments |
|--------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> 2-tier or 3-tier - app server VMs not shown DB+CI protected via VMware FT | <ul style="list-style-type: none"> Protection against server failure Continuous availability of CI + DB Easy to configure – VMware “out-of-the-box” New secondary VM automatically created after failover (assumes more ESX servers available) | <ul style="list-style-type: none"> No monitoring of application VMware FT supports 1 x vCPU – smallest sizing building block for production SAP VMs is typically 2 x vCPU Separate NIC/network recommended for FT logging traffic |

Figure 5. SAP DB + CI and VMware FT

The configuration above was tested as part of the LinuxLab@SAP Workshop conducted at SAP in Walldorf in the first quarter of 2009. The purpose of the workshop was to verify new virtualization technologies with SAP software. The test setup and results are described in Figure 6.

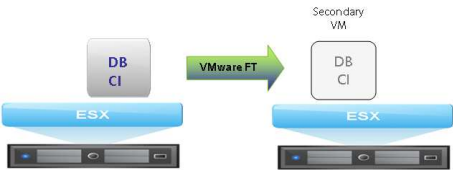
| Test Setup | Results with failover |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> • 2 x ESX servers running pre-release build of vSphere • 1 x VM, 1 vCPU, 8GB • SAP Linux Certification Suite (SLCS) running SAP 2005 EhP4 Unicode Instance and the Application Load Driver. Guest OS = Red Hat Enterprise Linux 5.1  | <ul style="list-style-type: none"> • 190 concurrent users • < 1 sec average response time • successful completion of workload with no user or application errors • Note: the setup was not intended or tuned for benchmarking. |

Figure 6. LinuxLab@SAP Workshop Results - VMware FT Test with VM running DB + CI

Summary

The four high availability scenarios discussed in this document provide varying degrees of fault tolerance and uptime to the single points of failure in the SAP architecture. These scenarios are summarized in the table below.


| SAP HA Scenario | Setup - Relative Complexity | Comments |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VMware HA: <ul style="list-style-type: none"> • 2-tier / 3-tier • CI + DB protected by VMware HA |  | <ul style="list-style-type: none"> • Protects against h/w failure • VMware "out-of-the-box" – easy to configure • DB + CI unavailable during failover |
| VMware FT: <ul style="list-style-type: none"> • 3-tier • CI or ASCS + DB protected by VMware FT | | <ul style="list-style-type: none"> • VMware "out-of-the-box" – easy to configure • Continuous availability of DB, message and enqueue services • VMware FT currently supports 1 x vCPU VM - SAP VMs typically sized at 2 x vCPU and above • Separate network for FT logging recommended |
| VMware FT + VMware HA: <ul style="list-style-type: none"> • 3-tier • ASCS in separate VM protected by VMware FT • DB protected by VMware HA | | <ul style="list-style-type: none"> • Protects against h/w failure • VMware "out-of-the-box" – easy to configure • Continuous availability for ASCS (enqueue + message service) • DB unavailable during failover • VMware FT currently supports 1 x vCPU VM – server vendor SAP competency center can provide sizing guidelines • Separate network for FT logging recommended |
| MSCS in virtual machines: <ul style="list-style-type: none"> • 3-tier • MSCS Agents for SAP service switchover (DB, ASCS, Replicated Enqueue) | | <ul style="list-style-type: none"> • Protects against h/w and s/w failure of DB and ASCS. • Continuous availability for enqueue locks only • DB unavailable during failover • Message Service unavailable during failover • More complex to configure • VMotion restriction for MSCS clustered VMs |

Figure 7. Summary of SAP on vSphere High Availability Scenarios

As clustering with MSCS includes monitoring of the SAP application in addition to hardware, it inherently provides the highest degree of monitoring but involves a more complex implementation compared to the VMware technology only solutions. Meanwhile customers currently running SAP solutions on VMware software are successfully utilizing VMware HA with VMotion to satisfy the high availability requirements of mission critical SAP solutions-based environments. The VMware FT functionality further addresses the high availability requirements of SAP solutions, especially in the area of protecting SPOFs like SAP central services. All of the high availability scenarios described here are viable but customer business and sizing requirements and SLAs will determine which option is best.

Further Reading

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

VMworld Europe 2009 Session TA24, Fault Tolerance Performance Preview

<http://www.vmworld.com/docs/DOC-2930>

(Requires VMworld.com logon credentials)

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 (ESX 4.0
ESXi 4.0 vCenter Server 4.0)*

http://vmware.com/pdf/vsphere4/r40/vsp_40_mscs.pdf

Installation Guide: NetWeaver 7.0 ABAP on Windows: MS SQL Server

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



VMware, Inc 3401 Hillview Ave Palo Alto, CA 94304 USA

Tel 650-427-5000 Fax 650-427-5001 www.vmware.com

© 2009 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156, 6,795,966, 6,880,022, 6,961,941, 6,961,806, 6,944,699, 7,069,413; 7,082,598 and 7,089,377; patents pending.

