

Genesys VMWare

Deployment and Technical Considerations Guide

(Running Genesys SIP Routing and Voice Platform with VMware vMotion and HA on vSphere 5.5)

Note: This document is co-owned by Genesys and VMWare.



Running Genesys SIP Routing and Voice Platform with VMware vMotion and HA on vSphere 5.5

July 2014

DEPLOYMENT AND TECHNICAL CONSIDERATIONS GUIDE

Table of Contents

Introduction	3
Executive Summary.....	3
VMware and Genesys Overview.....	4
Genesys Suite Overview	4
VMware Virtual Infrastructure	5
Genesys Architecture and Deployment.....	6
Genesys Test Process and Results	7
vMotion	8
High Availability.....	14
Hardware and Software Configuration	18
Virtual Machine Configuration.....	19
Workload Used.....	19
Deployment Best Practices.....	20
Technical Support	22
Conclusions.....	22
Resources	23
Acknowledgements	23

Introduction

Virtualization is rapidly transforming the IT landscape and fundamentally changing the way that people compute. Today's powerful x86 computer hardware was originally designed to run only a single OS. This approach introduced a tight coupling between Operating System (OS) and hardware. Virtualization breaks that bond between the OS and hardware by making it possible to run multiple OS and applications on the same computer at the same time, increasing the utilization and flexibility of hardware.

This document provides direction to those interested in running vMotion and High Availability (HA) using the Genesys SIP Routing and Voice Portal components on VMware® vSphere™ 5.5. It provides basic information on the architecture of Genesys SIP Routing and Voice Portal, as well as the value of utilizing the VMware platform. Finally, the document outlines some best practices for utilizing the two product sets together in your datacenter.

Executive Summary

One of the largest challenges that customers have to contend with in any environment is how to deal with planned downtime. In the physical world servers always need to be powered down at some point, for various reasons: firmware upgrades, adding or removing hardware, or to upgrade to newer hardware. The same is true in the VMware virtual world.

VMware has pioneered a unique technology that minimizes planned downtime and allows workloads running on one physical ESX host server to be "hot" migrated (moved live) to another ESX host server, without any impact to end users. While this sounds like magic, VMware customers have been using vMotion technology successfully since it was first delivered in 2003.

VMware designed vMotion to handle three different use cases. It is the fundamental technology that enables each of these capabilities:

1. Minimizing planned downtime with operator managed (manual) vMotion.
2. Automated server load balancing with VMware® Distributed Resource Scheduler (DRS).
3. Power savings with the automated power management subsystem VMware® Distributed Power Management (DPM).

In case of Server and operating system failures the best way to minimize downtime is by using VMware High Availability (HA).

1. HA monitors virtual machines to detect operating system and hardware failures.
2. HA restarts virtual machines on other physical servers in the resource pool without manual intervention when server failure is detected.
3. HA protects applications from operating system failures by automatically restarting virtual machines when an operating system failure is detected.

You may want to know what impact VMware vMotion and HA will have on your applications when virtual machines are moved from one physical server to another or when there is server or operating system failure. The first question is around transactional integrity: will all your data be available after the migration completes? The second question is around understanding transaction rate and overall impact to SLAs: will your transactions per second and response times be negatively impacted by vMotion? This white paper provides data that will help answer these questions.

VMware and Genesys Overview

VMware (NYSE:VMW), the global leader in virtualization and cloud infrastructure, delivers customer-proven solutions that accelerate IT by reducing complexity and enabling more flexible, agile service delivery. VMware enables enterprises to adopt a cloud model that addresses their unique business challenges. VMware's approach accelerates the transition to cloud computing while preserving existing investments and improving security and control. With more than 500,000 customers and 55,000 partners, VMware solutions help organizations of all sizes to lower costs, increase business agility and ensure freedom of choice.

Genesys Suite Overview

Genesys is the world's leading provider of customer service and contact software — with more than 4,000 customers in 80 countries. With more than 20 years of customer service innovation and experience, Genesys is uniquely positioned to help companies bring their people, insights and customer channels together to effectively drive today's customer conversation. Genesys software directs more than 100 million interactions every day, maximizing the value of customer engagement and differentiating the experience by driving personalization and multi-channel customer service – and extending customer service across the enterprise to optimize processes and the performance of customer-facing employees.

Figure 1: Genesys Suite Overview



VMware Virtual Infrastructure

VMware's leading virtualization solutions provide multiple benefits to IT administrators and users. VMware virtualization creates a layer of abstraction between the resources required by an application and operating system, and the underlying hardware that provides those resources. A summary of the value of this abstraction layer includes the following:

- **Consolidation:** VMware technology allows multiple application servers to be consolidated onto one physical server, with little or no decrease in overall performance.
- **Ease of Provisioning:** VMware virtualization encapsulates an application into an image that can be duplicated or moved, greatly reducing the cost of application provisioning and deployment.
- **Manageability:** Virtual machines may be moved from server to server with no downtime using VMware vMotion™, which simplifies common operations like hardware maintenance and reduces planned downtime.
- **Availability:** Unplanned downtime can be reduced and higher service levels can be provided to an application. VMware High Availability (HA) ensures that in the case of an unplanned hardware failure, any affected virtual machines are restarted on another host in a VMware cluster.

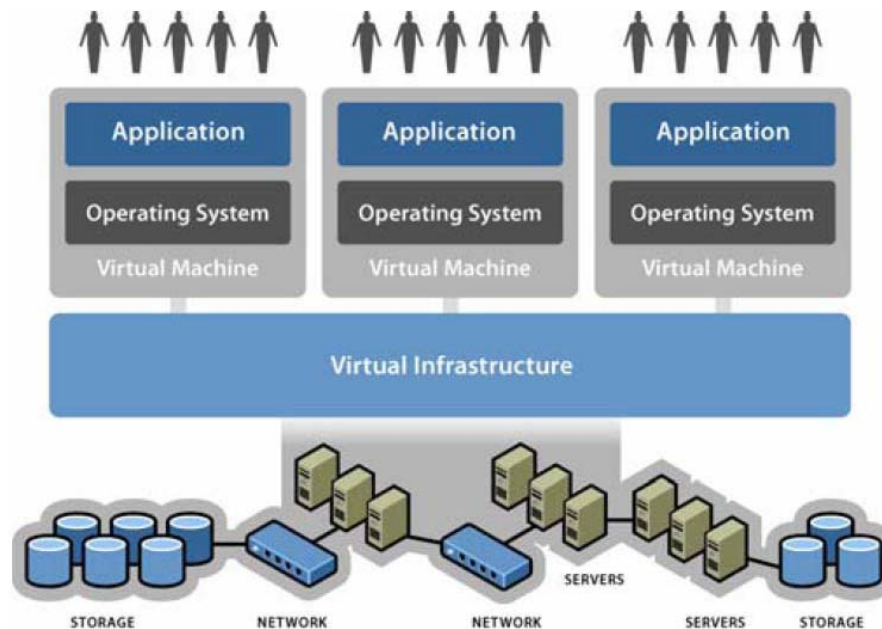


Figure 2: VMware vSphere Virtual Infrastructure

Genesys Architecture and Deployment

Genesys core components for SIP Server routing consist of Management Framework, SIP Server, Stat Server, Universal Routing Server (URS), Resource Manager and MCP. These components work together to form core architecture for routing all customer contacts to the appropriate agent or other resource.

A brief description of Genesys core components:

- **Genesys Framework** is the foundation for all Genesys-based interaction management systems. Management Framework provides you with the following administration functions: Configuration, Access Control, Solution Control, Alarm Processing, Troubleshooting, and Fault Management..
- **SIP Server** is the Genesys software component that provides an interface between your telephony hardware and the rest of the Genesys software components in your enterprise. It translates and keeps track of events and requests that come from, and are sent to the telephony device. SIP Server is a TCP/IP-based server that can also act as a messaging interface between SIP Server clients. It is the critical point in allowing your Genesys solution to facilitate and track the contacts that flow through your enterprise.
- **Genesys Universal Routing** enables intelligent distribution of interactions throughout the enterprise, whether you have a single-tenant or a multi-tenant environment. Universal Routing can direct interactions from a wide variety of platforms, such as toll-free carrier switches, premise PBXs or ACDs, IVRs, IP PBXs, e-mail servers, web servers, and workflow servers. It can handle pure-voice, multimedia, and blended environments, enabling routing of each media type based on appropriate criteria. Routing strategies and business processes automate interaction routing to the most appropriate agent/resource based on factors such as the type of inquiry, the value of the customer, and the media channel.
- **Stat Server** tracks information about customer interaction networks (contact center, enterprise-wide, or multi-enterprise telephony and computer networks). It also converts the data accumulated for directory numbers (DNs), agents, agent groups, and non-telephony-specific object types, such as e-mail and chat sessions, into statistically useful information, and passes these calculations to other software applications that request data. For example, Stat Server sends data to Universal Routing Server (URS), because Stat Server reports on agent availability.
- **Resource Manager** provides high availability for media-processing resources by applying high-level business logic to maintain coordination of voice self service sessions. The Resource Manager controls access and routing to all resources in a GVP 8.5 deployment. The Resource Manager is the first element to process requests for services, and it interacts with the Configuration Server to determine the Interactive Voice Recognition (IVR) Profile, Voice Extensible Markup Language (VoiceXML), Call Control Extensible Markup Language (CCXML), Announcement, and Conference application, resource, and service profile required to deliver the service.
- **Media Control Platform** is the core component of GVP. The Media Control Platform executes the actual voice applications in the solution. In addition, it is used by other communication layer components, such as SIP Server, to provide broader customer service scenarios, such as agent interactions, and many other functions. The Media Control Platform (MCP) provides core functionality in a voice self service solution by controlling overall execution of voice applications and interacting with various media-processing resources.

The above components can be deployed in the customer's network by hosted service providers and across single or multiple sites. Deployment using virtualization is possible for all components, and it can provide cost savings and management advantages.




Figure 3: Genesys Architecture

Genesys Test Process and Results

In this exercise, VMware worked with Genesys to test the behavior of Genesys SIP Routing and Genesys Voice Platform with vMotion and High Availability. To make this testing relevant to real-world environments we used the Genesys Simulator which closely resembles industry standard benchmark. Given the mission-critical nature of the Routing System, they need to ensure that no calls were dropped during the VMware vMotion and whether all the services properly come up after VMware HA.

The primary objectives of testing was to check if vMotion and HA works resiliently on VMware virtual infrastructure. The joint tests were also designed to provide answers a common set of questions, and to derive and document a set of best practices that can be used by our customers.

Specifically, key areas under investigation were:

- Impact on calls per second during VMware vMotion/HA operation.
- Test whether the Primary component properly switches as a backup and processes all the calls after VMware HA.

vMotion

VMware vMotion technology moves running virtual machines from one physical server to another with no impact on end users. VMware vMotion keeps your IT environment up and running, giving you unprecedented flexibility and availability to meet the increasing demands of your business and end use.




Figure 4: VMware vMotion

20 Calls per second

In this test the Simulator was configured to generate load of 20 calls per second. The vMotion feature was tested with SIP Server, StatServer, URS, RM and MCP.

SIP Server vMotion with 20 CPS




Figure 5: SIP Server vMotion with 20 CPS

RM vMotion with 20 CPS




Figure 6: RM vMotion with 20 CPS

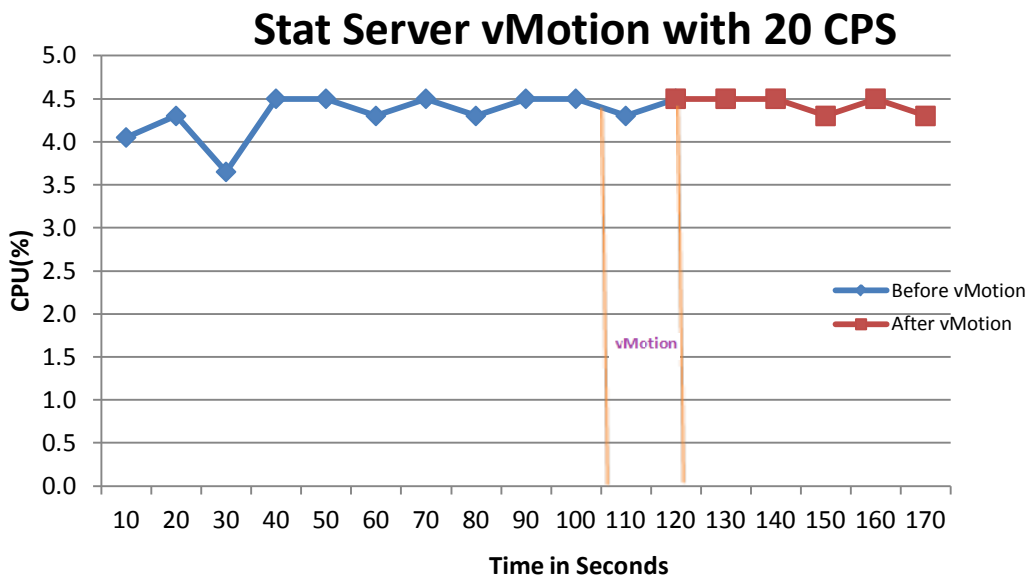


Figure 6: Stat Server vMotion with 20 CPS




Figure 7: URS vMotion with 20 CPS




Figure 8: MCP vMotion with 20 CPS

30 Calls per second

In this test the Simulator was configured to generate load of 30 calls per second. The vMotion feature was tested with SIP Server, StatServer, URS, RM and MCP..

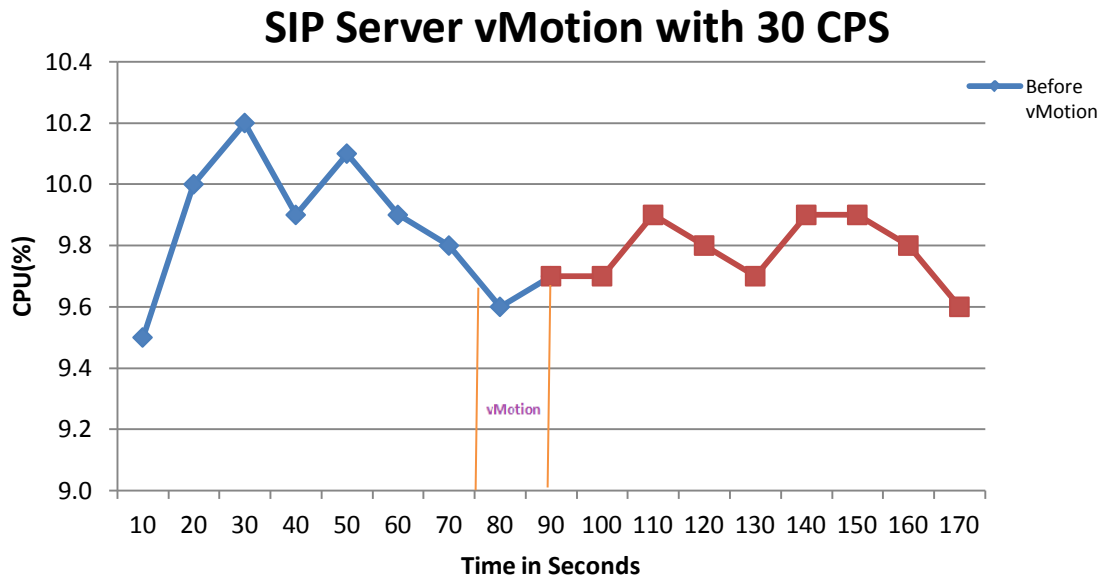


Figure 9: RM vMotion with 30 CPS




Figure 10: RM vMotion with 30 CPS




Figure 11: Stat Server vMotion with 30 CPS




Figure 12: URS vMotion with 30 CPS

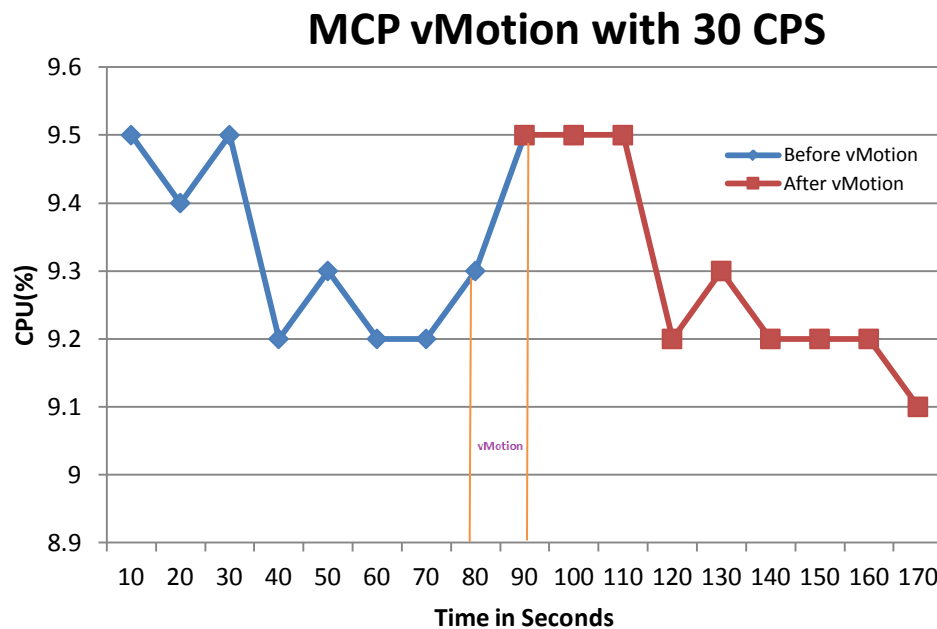


Figure 13: MCP vMotion with 30 CPS

Observations and Highlights

1. Average vMotion time for all tests was approximately between 10 to 15 seconds.
2. No calls were dropped or abandoned.
3. There was no impact to agent login sessions.

High Availability

vSphere High Availability (HA) delivers the availability needed by many applications running in virtual machines, independent of the operating system and application running in it. HA provides uniform, cost-effective failover protection against hardware and operating system failures within your virtualized IT environment.




Figure 14: vSphere High Availability (HA)

In this scenario, SIP Server, StatServer, RM, URS and MCP were enabled for VMware HA. After a host failure, the backup nodes on the other host in the cluster automatically become primary nodes and re-establish connection with different components in the System. The scenario was tested as follows:

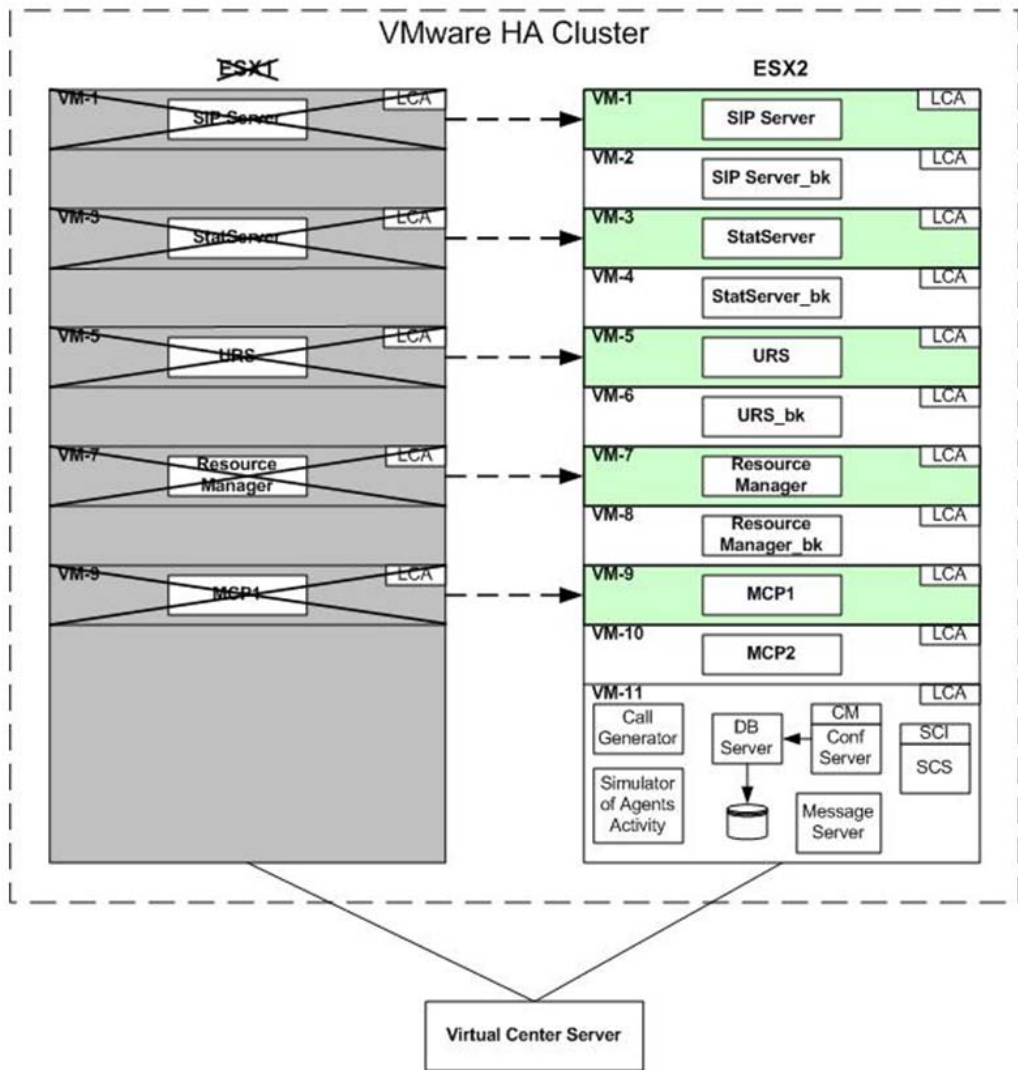


Figure 15: Affirmed HA Architecture




Figure 16: VMware HA with 10 CPS

Observation and Highlights

1. HA is enabled for the SIP Server, StatServer, URS, RM and MCP VM's which are running in primary mode.
2. Another host in the cluster has all the backup components of all the above 5 VM's.
3. In a HA process the host with all the Primary VM's is rebooted and it was observed that all the backup components on the other host resume processing with minimal interruption. During this transition it was observed that we had lost 139 calls.
4. As part of VMware HA, all the primary components reboot on the another host in the cluster

Hardware and Software Configuration

The following section provides details on the hardware and software used in the testing.

The following table describes the configuration of ESX host servers and storage in the Genesys test configurations.

Table 1: ESX Host Hardware

HARDWARE	CONFIGURATION
HP Proliant DL380p G8	Three HP Proliant DL380p G8 servers. Each server is equipped with: <ul style="list-style-type: none"> • Intel(R) Xeon(R) CPU E5-2680 @2.70 GHz • Sockets: 2 • Number of Cores per Socket: 8 • Hyperthreading: Enabled • 132 GB RAM • 4 X 1 GB NICS per ESX host
Storage	EMC VNX 5600: <ul style="list-style-type: none"> • 13 disks, 15K RPM • 8 Gbps Fibre Channel

Installed Software

Table 2 lists the software used for the Genesys solution running on VMware virtual infrastructure.

Table 2: Software Installed for Genesys Testing on VMware Virtual Infrastructure

INSTALLED SOFTWARE	
VMware	VMware vSphere ESX 5.5.0 Build 1623387
Microsoft	Windows 2008 R2 Enterprise 64 bit <ul style="list-style-type: none"> • SIP Server • Stat Server • URS • Resource Manager • MCP1 and MCP2 • Simulator

Virtual Machine Configuration

The following table describes the configuration of virtual machines running on ESX host servers in the Genesys test configurations.

Table 3: Virtual Machine Configuration

VIRTUAL MACHINE	HARDWARE CONFIGURATION
11 virtual machine instances TServer TServer_bk Stats Server StatServer_bk URS URS_bk RM RM_bk MCP1 MCP2 Simulators	Each configured with:- <ul style="list-style-type: none"> • 4 vCPUs • 4 GB memory • Ethernet card (1 Gbps Network) • Disk storage 1 x 40 GB • vmdk only

Workload Used

To accomplish the testing, Genesys tools to simulate loads were used:

- CCAS/Autoserv: Simulates customer Inbound callers
- CCAS/Autoserv: Simulates Agent activity

Deployment Best Practices

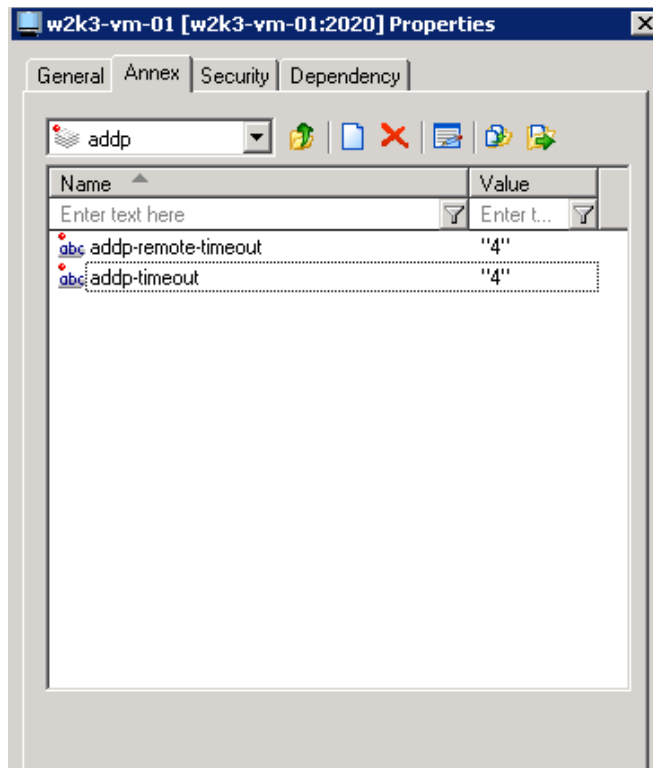
Here are some of the best practices derived from configuration and testing of the Genesys application in a VMware virtual infrastructure environment:

1. When using Intel-based systems, ensure that the BIOS settings enable VT and EPT options on all ESX hosts. Hyperthreading should be enabled.
2. Configure the VMs and the VMkernel on a separate NIC on a separate vSwitch.
3. Use VMware paravirtualized vmxnet3 adapter for better network throughput.
4. Install VMware tools on the virtual machines. The VMware Tools package provides support required for shared folders and for drag and drop operations. Other tools in the package support synchronization of time in the guest operating system and the host, automatic grabbing and releasing of the mouse cursor, copying and pasting between guest and host, and improved mouse performance in some guest operating systems.
5. Consider using server-class network interface cards (NICs) for the best performance and configure paravirtualized vmxnet3 adapters for better network throughput.
6. Virtual machines dependent on each other should be configured on same ESX host and same vSwitch. This will reduce traffic flow through the wire and avoid unnecessary CPU and network overhead. An example would be having SQLSERVER and one of the APPS on the same ESX host.
7. Disconnect or disable unused or unnecessary physical hardware devices, such as:
 - COM ports
 - LPT ports
 - USB controllers
 - Floppy drives
 - Optical drives (that is, CD or DVD drives)

Disconnecting or disabling devices will help free up interrupt resources. For example, traditionally, some devices, such as USB controllers, operate on a polling scheme that consumes extra CPU resources. Some PCI devices reserve blocks of memory, making that memory unavailable to ESXi. HA Failover Architecture.

8. To reduce the number of the abandoned/unprocessed calls during Genesys HA recovery, ADDP parameters should be set for 4 seconds on the host object in CME, as shown below:

Figure 16: ADDP Parameters



9. Genesys HA and VMware HA are complementary, and may be used together, provided that anti-affinity rules are followed (see #10 below).

- Genesys HA should always be used to provide application level protection, and ensure that calls and contacts in progress are not interrupted. (DRS anti-affinity required.)
- VMware HA adds an additional layer of hardware failure protection. If there is a physical server failure, VMware HA will restart any failed "Genesys resiliency HA VM's impacted by the server failure" on different servers automatically. Without VMware HA, the "Genesys resiliency HA VM's" would remain on the failed server, powered down, disabling Genesys HA's functionality.
- The combination of VMware HA, anti-affinity rules and Genesys HA allow Genesys HA to be brought back to full protection within a few minutes of a failed physical server that is hosting Genesys HA components, while at the same time facilitating the repair / replacement of the failed server.

10. To respect anti-affinity rules when restarting virtual machines after a failure follow the instructions as per the below KB article:-

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2033250

Technical Support

For technical support on virtualization issues, existing Genesys customers should use the following contacts:

North America	support@genesys.com
Europe, Middle East, and Africa	support@genesys.co.uk
Asia Pacific, Malaysia, India	support@genesys.com.au
Japan	support@genesys.co.jp

For general questions on Genesys virtualization support, contact virtualizationsupport@genesys.com

Conclusions

Overall, testing results show that, running VMware vMotion and HA with Genesys SIP Routing and Genesys Voice Platform performs well. Furthermore, it has potential to reduce cost, increase service levels and simplify the manageability of applications. It is critical that new virtualization server capabilities are used, since this will significantly increase performance of Genesys. It is equally important that you follow hardware recommendations and best practices as mentioned in this guide.

Resources

Customers can find more information about VMware and Genesys products using the links listed below.

VMware Resources

- VMware official website:
<http://www.vmware.com/>
- VMware Infrastructure 5 product Web site:
http://www.vmware.com/products/data_center.html
- VMware download Web site:
<https://www.vmware.com/download/>
- VMware support Web site:
<http://www.vmware.com/vmtrn/>
- VMware Performance Tuning Paper:
http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.5.pdf
- System Compatibility Guide for a complete list of compatible hardware:
http://www.vmware.com/resources/compatibility/pdf/vi_systems_guide.pdf
- Storage/SAN Compatibility Guide for a complete list of compatible storage devices:
http://www.vmware.com/resources/compatibility/pdf/vi_san_guide.pdf
- I/O Compatibility Guide for a complete list of compatible networking devices:
<http://www.vmware.com/resources/compatibility/search.php>

Genesys Resources

- Genesys Web site:
www.Genesys.com

Acknowledgements

The following individuals contributed to the creation of this deployment guide:

- Balayya Kamanboina, ISV Validation Engineer, VMware
- Leonid Mitnik, Senior Software Engineer, Quality Assurance, Genesys
- Leonid Ryklansky, Senior Manager, Quality Assurance, Genesys