



PureConnect®

2018 R5

Generated:

12-November-2018

Content last updated:

12-January-2018

See [Change Log](#) for summary of changes.



CIC Virtualization

Technical Reference

Abstract

This technical reference content describes the best practices and strategies for successfully deploying PureConnect products on a virtualized platform. The guidelines apply to Microsoft Hyper-V and VMware servers.

For the latest version of this document, see the PureConnect Documentation Library at: <http://help.genesys.com/cic>.

For copyright and trademark information, see https://help.genesys.com/cic/desktop/copyright_and_trademark_information.htm.

Table of Contents

Table of Contents	2
Introduction	3
Characteristics of virtualization	4
Disk I/O	5
Locally attached storage	5
Sample Disk I/O Measurements	5
Storage Area Networks (SANs)	6
CPU resources and reservations	8
Network adapters	9
Multiple NICs	9
Hyper-V NIC configuration on HP servers	10
Virtual network adapter types	11
Emulated	11
Paravirtualized (VMWare)/Synthetic (Microsoft Hyper-V)	11
Recommendation	11
Memory allocation	12
Hardware recommendations for a host	13
Characteristics of PureConnect products	14
Supported for virtualization (Hyper-V and VMware)	14
Supported on VMware only	15
Products not supported for virtualization	16
Client application virtualization	16
Virtualizing Interaction Media Server in a production environment	17
Expectations for virtualized Interaction Media Server	17
Host recommendations for an Interaction Media Server virtual machine	18
Virtual machine requirements for Interaction Media Server	19
Required Interaction Media Server/CIC server configuration	19
Testing results of an Interaction Media Server guest in Hyper-V	20
Audio quality issues	20
Clock synchronization issues	20
Networking issues	22
Testing results of Interaction Media Server running natively on a Hyper-V host	25
Sizing and performance of a virtual CIC server guest VM	26
Sizing the Customer Interaction Center server	26
Small implementation	26
Medium implementation	26
Large implementation	27
Very large implementation	27
Using Customer Interaction Center with Interaction Process Automation (IPA)	27
Using Dialer-only or blended environments	28
Input/output operations per second (IOPs)	29
Dialer testing results with regard to IOPs	29
Interaction Dialer scenario #1	30
Interaction Dialer scenario #2	30
Hyper-V Live Migration and VMware VMotion	31
Data protection	33
Logging drive	34
Windows Disk Compression on Logging folders	34
Disk type and size	35
Importance of a testing environment	36
Good things to know	37
Running with guest enhancers	37
Identifying a virtual machine from the PureConnect logs	37
Hosts with mismatched CPUs, MAC addresses, and licensing	38
Common issue	38
Rare issue	38
Change Log	39

Introduction

The purpose of this document is to provide guidelines for running PureConnect products on a virtualized platform. This information is based on tested configurations that are proven to provide the best experience for each product. The examples in this document are provided as a courtesy and must not be taken as exact representations of any customer environment. The configurations presented might not be the best solution for larger, more complex environments. The intent of this document is to demonstrate some best practices and to provide a resource for engineers before implementation.

Important!

Commercially-reasonable efforts have been made to test PureConnect products in virtualized environments and to document general best practices. Since virtualization is a wide subject area with many factors that influence performance, Genesys does not guarantee that these best practices provide trouble-free operation in your contact center, data, or network environments; accordingly, the virtualization information in this document is provided on an “as is” basis and any reliance on it is solely at your risk.

For information about tested third-party virtualization solutions, see <http://testlab.genesys.com/ProductsPage.aspx?ProductType=33>.

If you have questions regarding virtualization of PureConnect products, contact your sales representative.

Characteristics of virtualization

Virtualization allows companies to take advantage of the same physical hardware to reduce costs and also provides a foundation for dynamic disaster recovery solutions. These aspects have benefits, but it is important to remember that the nature of virtualization leads to shared resources. *It is critical that you protect PureConnect products from performance issues caused by the following example problems:*

- Underperforming shared I/O
- Shared guest CPU starving
- Insufficient memory allocation
- Suboptimal network configuration
- Storage Area Network (SAN) transfers
- Provisioning new systems

Note: The PureConnect Customer Care is not responsible for the configuration or performance of your virtual environment.

Disk I/O

[Locally attached storage](#)

[Storage Area Networks \(SANs\)](#)

Locally attached storage

The host systems that house the guests have different options for physical storage. Locally-attached physical disks provide excellent performance and are relatively inexpensive. Our testing efforts had good success with HP DL 380G7s with 16 drives. On the following pages are some test results from benchmarking the I/O. You can see where the drive count is a significant factor in the performance. Locally attached storage is a good option for a smaller company that doesn't have a sufficient budget or doesn't want to make a strategic investment in a SAN.

The virtual guests that run on a host share the pool of drives. Over-subscribing the I/O leads to performance problems that affect all guests running on the host. We recommend that you measure the usage of the I/O with the following Microsoft Performance Monitor (Perfmon) counters:

- **Physical Disk > Avg. Disk Queue Length** This value should be no greater than 2 per physical disk in your array.
- **Physical Disk > % Idle time** Perfect would be 100%. If % **Idle time** is dropping below 30, the disk is starting to become over used.
- **Physical Disk > Data Transfers /sec** This value is what hardware vendors refer to as Input Output per second (IOPs).

Measuring the disk counters from the VM can be fairly accurate with Performance Monitor (Perfmon) although some documentation from Microsoft advises against this measurement. The disk counters seem to be a good indication of what is happening on the I/O for VMware guests as well. You can also use the host I/O counters in Hyper-V, but they appear as the aggregate of all the guests running on the machine. A **Performance** tab is provided for monitoring the host and the guests with similar counters on VMware vSphere.

Sample Disk I/O Measurements

Using SQLIO, we took measurements of Disk I/O on two different systems to compare their performance. SQLIO is a downloadable tool provided by Microsoft which can be used to determine the I/O capacity of a specified configuration. You can download SQLIO through the following URL address:

<http://www.microsoft.com/en-us/download/details.aspx?id=20163>

Configuration 1:

- Hardware - HP DL 380 G7 (RAID 10)
- Smart Array P410i controller 512 MB cache (25% read / 75% write)
- 14 x 500 GB SATA 3.0 GB p/s 7200 RPM
- Windows 2008 R2 SP1

Test Pattern	MB/s	IOPs
Random Write (worst case)	91	1400
Random Read*	120	2050
Sequential Write	540	8700
Sequential Read (Best Case)	545	8800

*Random Reads were sporadic. This number was an average, but it was as low as 16 MB/s and as high as 219 MB/s.

Configuration 2:

- Hardware - HP DL 380 G7 (RAID 10)
- Smart Array P410i controller 512 MB cache (25% read / 75% write)
- 8 x 500 GB SATA 3.0 GB p/s 7200 RPM
- Windows 2008 R2 SP1

Test Pattern	MB/s	IOPs
Random Write (worst case)	23	370
Random Read*	70	1000
Sequential Write	300	4800
Sequential Read (Best Case)**	185	2900

*Random Reads were sporadic. This number was an average, but it was as low as 11 MB/s and as high as 82 MB/s.

** Not a good number on Sequential Read with this small number of drives.

Configuration 3:

- Hardware - HP DL 360 G7 (RAID 10)
- Smart Array P812 controller 1 GB cache (25% read / 75% write)
- 14 x 300 GB SAS 6.0 GB p/s 10,000 RPM
- Windows 2008 R2 SP1

Test Pattern	MB/s	IOPs
Random Write	227	3628
Random Read	421	6730
Sequential Write	473	7564
Sequential Read	880	14079

Storage Area Networks (SANs)

Storage Area Networks (SANs) are more common in a medium to large company. SANs vary greatly in reliability, performance, and cost, and are beyond the scope of this document. SAN vendors commonly use the term *IOPs* (Input Output Per second) to describe the performance of the raw disks of their hardware. In a hardware sales process, customers want to know the following information:

- How many IOPs does my Customer Interaction Center server require?
- How many IOPs should my SQL Server be able to do?

Be careful using a vendor's claimed IOPs number for comparison purposes. Once you factor in Fibre Channel or iSCSI encapsulation, the virtualization layer, operating system, and application, it will be less than what is advertised with raw hardware. VMware does allow for a minimal reservation of IOPs per guest.

SANs offer much more flexibility and can be a good part of a disaster recovery plan. Depending on the vendor, there may be some excellent monitoring tools for making sure that the guests are not overwhelming the storage. Operations such as guest transfers, deployments, and cloning can potentially cause problems for other guests running and depending on the SAN. SANs can be more complicated than a local storage solution.

CPU resources and reservations

A virtual CPU reservation or a binding is a good way to give an important virtual machine guest guaranteed CPU time when other guests on the machine are busy. A virtual machine reservation is still not as good as a real CPU; however, if the host is using more CPU than is available, even a guest with a reservation or high shares is not entirely isolated.

For PureConnect products, we recommend that you use a minimum of two virtual CPUs. For example, virus applications can update virus definitions on one thread at real-time priority that dominates the CPU. On single-core machines, this kind of update causes the server to become unresponsive to the job it is trying to perform. It's easy to create a VM with only one virtual CPU, but it's rare to see any physical server hardware with one core and no hyper-threading.

PureConnect software is designed to take advantage of symmetric multiprocessing. Subsystems can dynamically adapt and use the processors for executing threads simultaneously so making more virtual processors available to the guest helps performance.

Question	Answer
Does having more virtual processors increase scheduling latency in VMWare?	The scheduling latency is low on VMware with virtual processors. To determine this low latency, Genesys used the Time Stamp Counter (TSC) and the Multimedia Timer (MMT). We did a comparison on the samples at a 1-ms frequency with the reported clock speed of the virtual processors from the perspective of the guest. We measured to see the jitter in MMT compared to real time. Under VMware, the TSC times almost match the exact MM timer intervals (measured from edge to edge, but not necessarily consecutive edges). We did extensive testing with such measurements, since we have also investigated the platform for handling RTP streaming. At the application level, the extra virtual CPUs give processes more threads to work with and determine how the thread pools are dynamically created.
Is a reservation handled the same way on Hyper-V as in VMWare?	No, the vendors have different approaches to reservations. In some cases, we ask you to reserve the core CPU MHz you are using. If the CIC server CPU is not available when it is required, you will miss calls.
Do Hyper-V and VMware use cores the same way?	<p>VMware does not count hyper-threaded (HT) cores in the reservation pool. Hyper-V does count hyper-threaded cores. When setting reservations, you cannot run as many virtual machine guests on a VMware host as you can on the same hardware with Hyper-V.</p> <p>For example: Let's say our host has this specification: E5620 @ 2.40 GHz with two physical sockets with four physical cores each (eight physical cores), with hyper-threading resulting in 16 logical CPUs.</p> <p>In this case, VMware only allows reservation of the physical cores, after subtracting the system overhead of running the hypervisor itself. So, in VMware, there would be approximately 17,420 MHz available for guests after it allocates CPU resources to running the hypervisor.</p> <p>$(2.40 \text{ GHz} \times 8 \text{ physical cores}) - \text{about } 1780\text{-MHz hypervisor overhead} = 17,420 \text{ MHz}$</p> <p>Hyper-V would make 38,400 Mhz available to guests with 100% reservations of all the virtual CPUs.</p> <p>$2.40 \text{ GHz} \times 16 \text{ logical CPUs} = 38,400 \text{ MHz}$</p>

Network adapters

With the simplest virtualization configurations, the host and guests may share a common NIC or even a teamed NIC. Shared NICs work well for small test environments but are not recommended for large production use. We have experienced excellent performance using a dedicated guest NIC with VMware when running network intensive operations. For Hyper-V, lab tests revealed a small performance gain of lower %DPC time.

[Multiple NICs](#)

[Hyper-V NIC configuration on HP servers](#)

[Virtual network adapter types](#)


Multiple NICs

PureConnect supports the use of multiple NICs, even in a virtual machine. However, using multiple NICs in a virtual machine greatly increases complexity. Consider the following factors to determine if you require multiple NICs in a virtual machine:

Software licensing	<p>PureConnect Licensing uses a combination of CPUID and MAC address. Changing NIC configurations that result in changes to the MAC address after applying a license requires that you acquire a new license and apply it.</p> <div style="background-color: #f0f0f0; padding: 5px;"><p>Important!</p><p>When creating a virtual machine for a CIC server, you must set a static MAC address value for the CIC server guest. PureConnect Licensing is based on a combination of the CPUID and the MAC address. For more information, see Hosts with mismatched CPUs, MAC addresses, and licensing.</p></div>
Switchover	<p>While not required, you can use a single NIC for data synchronization and status monitoring.</p>
VLANs	<p>If you use VLANs, you may want to dedicate separate NICs for different VLANs. This decision should include consideration of the bandwidth of your NICs, subnets, and protocols.</p>
Backups	<p>For large and frequently-changing databases, a separate, dedicated NIC for only backups can increase speed and not impact the bandwidth available on other NICs for communications and other data operations.</p>

Hyper-V NIC configuration on HP servers

If you don't have a good system in place, understanding how physical NICs map to the logical NICs can be confusing. On the physical host server, the ports are numbered 1-4 per four-port adapter. The problem is that Windows does not see the ports in this order and, at this point, it has nothing to do with Hyper-V. If you plug a cable into port 1, Windows may report it as Local Area Connection 3. As you plug in each cable, notice which adapter becomes live and rename it to something that makes sense to you. The following image is an example for teamed NICs:



Without naming these connections as you connect ports, all of the NICs become Local Area Connection and an arbitrary number. The Hewlett Packard Network Configuration Utility further complicates matters with its own naming system. There is a white paper available about the NIC configuration that seems outdated and no longer relevant:

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01663264/c01663264.pdf>

We experienced no problems with the order of the install with the latest drivers and Microsoft software updates during testing.

We use Network Fault Tolerance (NFT) teaming because time-critical, order-dependent data could be going through the guests. Use Virtual Network Manager and map your NIC teams or single NICs to a virtual NIC that the CIC guest uses.

Important!

Do not enable the **Allow management operating system to share this network adapter** check box. Use a dedicated Management NIC that is not shared with your production guest NICs.

Virtual network adapter types

Most hypervisors offer the following methods for defining network adapters for a virtual machine guest operating system: *emulated* and *paravirtualized/synthetic*.

Emulated

With an emulated network adapter, a hypervisor presents the definition of a legacy, real-world, PCI-based, network interface to the operating system of a guest virtual machine. Emulating an older network adapter, such as a basic 1Gb Intel NIC 82545E, helps ensure that the operating system has native device drivers that are compatible. Emulated network adapters work in nearly all situations and require no additional software.

Because a device driver is required for the emulated network adapter, the operating system of the guest virtual machine uses more CPU resources to process all network traffic. On a host with multiple guest virtual machines, those CPU resources are consumed more extensively.

It is important for you to understand that emulated virtual network adapters result in higher CPU resource utilization than network adapter drivers on a non-virtualized installation. For some PureConnect products that use significant CPU resources in normal operation, this increase in CPU utilization can result in errors and issues as CPU resources are consumed and exhausted.

Paravirtualized (VMWare)/Synthetic (Microsoft Hyper-V)

Paravirtualized or synthetic network adapters are a more-recent network solution for virtualization. For paravirtualized/synthetic network adapters, the hypervisor presents an abstraction of a network adapter as a generic PCI device to the operating system of the guest operating system.

The hypervisor provides software for installation on the operating system of the guest virtual machine. The software enables the operating system to use the abstraction as a network interface.

Generally, paravirtualized/synthetic network adapters are presented by the hypervisor to the operating system of the guest virtual machine as a 10Gb network interface. Because there is no specific device driver, the operating system of the guest virtual machine uses fewer CPU resources to process network traffic.

Recommendation

In general, we recommend that you use paravirtualized/synthetic virtual network adapters in virtual environments for PureConnect and its subsystems.

Memory allocation

Dynamic memory allocation and static memory assignments are possible for virtualized servers and clients. No PureConnect product dynamically extends to consume as much memory as it can get. Likewise, no PureConnect product installation program fails to install based on minimal memory requirements. Overall, it is best to give PureConnect products a static memory assignment with a 2 GB minimum for a server. There is valid argument that dynamic memory is more for client OS virtualization and that critical servers use reservations. There is always a risk that the memory needed will not be available without a minimal reservation of some type.

Hardware recommendations for a host

The following is a hardware recommendation for a virtualization platform host, running PureConnect software as a guest.

The hardware that we recommend for a host is greater than the minimum hardware requirements provided by the virtualization platform vendor to run the virtualization platform. These recommendations are based on extensive lab tests and experience with the PureConnect software. Consult your virtualization platform vendor for assistance in determining appropriate hardware specifications for your environment.

Component	Recommendation
Processor	<p>64-bit x86 CPUs with Intel VT or AMD-V technology enabled. Hardware Data Execution Prevention (DEP), specifically Intel XD bit (execute disable bit), or AMD NX bit (no execute bit), must be available and enabled.</p> <div style="border: 1px solid black; padding: 5px; background-color: #f0f0f0;"> <p>Caution!</p> <p>The number of CPU cores that you must assign to the hypervisor host to provide adequate processing for the guest virtual machines varies depending on many factors, such as hypervisor software, hardware specifications, the number of active guest VMs, the processing load of the guest VMs, the amount of network traffic, the need for immediate processing of communications, and others.</p> </div> <p>Minimum Processor Speed: 2.0Ghz or faster in Intel Nehalem class or Intel Xeon 5400, 7100/7300 or 7200/7400 @ 2.5Ghz or greater</p> <p>Recommended: Intel Nehalem 2.40 GHz or greater</p> <p>Examples: E5620 2.40Ghz, X5677 3.47GHz</p>
RAM	<p>Minimum 8 GB RAM for the host.</p> <p>Minimum 3 GB for each CIC guest unless otherwise specified in Supported for virtualization (Hyper-V and VMware).</p>
Available storage space	<ul style="list-style-type: none"> • 120 GB or greater for the host partition. • 300 GB of Fixed or Thick Provisioned storage space for each CIC guest. This storage space is required for CIC tracing and debugging dump files.
Network adapters	<p>Modern adapters have a speed of 1Gb/s or higher and support RSS and VMQ.</p> <div style="border: 1px solid black; padding: 5px; background-color: #f0f0f0;"> <p>Notes:</p> <p>Depending on the number of guests and network transmissions, 1Gb/s may not be sufficient.</p> <p>If you are using Windows Server 2012 R2 Hyper-V, we recommend that you enable the Receive Side Scaling (RSS) feature on your virtual NICs so that Hyper-V uses its vRSS feature to spread the processing of network activity across the multiple virtual CPU cores that are assigned to the virtual machine.</p> </div>


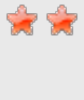
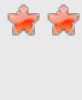





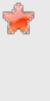
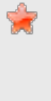
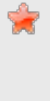

Characteristics of PureConnect products

PureConnect products are designed to work together, but they do not all have the same resource needs or characteristics. Some PureConnect products have higher CPU needs and some need more networking resources. The following tables provide insights into key virtualization resources that are critical for successful deployments.

Supported for virtualization (Hyper-V and VMware)

More stars indicate the intensity of the operation (Four stars = high intensity).


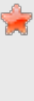

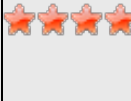




Application	CPU	I/O	Memory	Network	Notes
Customer Interaction Center server	☆☆☆ ☆☆☆	☆☆☆ ☆☆☆	☆☆☆ ☆☆☆	☆☆☆ ☆☆☆	For more information, see Sizing and performance of a virtual CIC server guest VM .
Interaction SIP Proxy	☆☆☆	☆☆	☆☆	☆☆☆ ☆☆☆	<ul style="list-style-type: none"> Up to 30 CPS (calls per second) 2 virtual processors minimum 2 GB RAM minimum 300 GB hard drive minimum[1] Dedicated NIC recommended
CIC Database SQL Server	☆☆☆	☆☆☆ ☆☆☆	☆☆☆ ☆☆☆	☆☆	<ul style="list-style-type: none"> 4 virtual processors 10 GB of RAM or more 8 CPS is around 250 IOPs with recording and reporting
Interaction Recorder Remote Content Service	☆☆	☆☆☆	☆☆	☆☆☆ ☆☆☆	<ul style="list-style-type: none"> 2 virtual processors 2 GB RAM minimum 300 GB hard drive minimum[1]
Off-server Session Manager (OSSM)	☆☆☆ ☆☆☆	☆☆☆	☆☆☆ ☆☆	☆☆☆ ☆☆☆	<ul style="list-style-type: none"> 4 virtual processors 4 GB RAM 4 cores 300 GB hard drive minimum[1] <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p>Note: A virtual OSSM with four virtual CPU cores supports a maximum of 500 sessions</p> </div>
WebServices	☆☆☆	☆☆	☆☆☆	☆☆☆	<ul style="list-style-type: none"> 4 virtual processors 4 GB RAM Up to 500 chats 300 GB hard drive minimum[1]
Voice XML	☆☆☆	☆☆	☆☆☆	☆☆☆	<ul style="list-style-type: none"> 2 virtual processors 2 GB RAM 300 GB hard drive minimum[1] Suitable for most deployments
Web Interaction Administrator	☆☆	☆☆	☆☆	☆☆	<ul style="list-style-type: none"> 2 virtual processors 2 GB RAM
Status Aggregator	☆☆☆ ☆☆☆	☆☆☆	☆☆☆ ☆☆☆	☆☆☆	<ul style="list-style-type: none"> 4 virtual processors 4 GB RAM 300 GB hard drive minimum[1]
RTMServer	☆☆☆	☆☆	☆☆	☆☆☆	<ul style="list-style-type: none"> 4 virtual processors 4 GB RAM 300 GB hard drive minimum[1]

Render Server					<ul style="list-style-type: none"> • 4 virtual processors • 2 GB RAM • 300 GB hard drive minimum[1]
Interaction Dialer CCS					<ul style="list-style-type: none"> • 2 virtual processors • 2 GB RAM • 300 GB hard drive minimum[1] • Suitable for most deployments
Scheduled Reports Server					<ul style="list-style-type: none"> • 2 virtual processors • 2 GB RAM • 300 GB hard drive minimum[1] • Suitable for most deployments

[1] For sites that may require increased tracing for troubleshooting purposes, it is highly possible that more than the minimum disk requirements will be needed. In some cases where increased tracing was performed, 600 GB was the increased disk size. Your individual requirements and results may vary.

Supported on VMware only

More stars indicate the intensity of the operation (Four stars = high intensity).


















Application	CPU	I/O	Memory	Network	Notes
ASR Server Nuance Recognizer					
TTS Server Nuance Vocalizer					

Note:

Genesys does not provide support or recommendations regarding the virtualization of third-party products. Consult the manufacturer-supplied documentation for each third-party product that you want to virtualize and integrate into the PureConnect environment. Genesys is not responsible for the configuration and performance of virtualized third-party products that you want to integrate into a PureConnect environment.

Products not supported for virtualization

More stars indicate the intensity of the operation (Four stars = high intensity).

Application	CPU	I/O	Memory	Network	Notes
Interaction Media Server	 	 			Important! You are allowed to virtualize Interaction Media Server for non-PureConnect Cloud environments but Genesys does not provide support for this specific usage. For sizing and configuration, see Virtualizing Interaction Media Server in a production environment .
IR Extreme Query Server		 	 		This application is built for performance.
Loquendo ASR	 				Highly sensitive to latency

Client application virtualization

You can virtualize PureConnect client applications, such as IC Business Manager and Interaction Desktop through desktop virtualization software, such as Citrix XenApp and VMWare Horizon. For more information about supported desktop virtualization solutions, see <http://testlab.genesys.com/ProductsPage.aspx?ProductType=27>.

Virtualizing Interaction Media Server in a production environment

Interaction Media Server processes communications in real-time as opposed to other systems that can perform data operations with minor, acceptable delays in processing. As such, Interaction Media Server is not well-suited and not recommended for virtualization as any processing delays and contention for resources can cause audio quality issues.

Some applications are better candidates for virtualization than others because of many variables, such as sensitivity to delays, direct interaction with hardware resources, and others. Not all virtualization platforms are equal as some solutions have different scaling capabilities, various features, and levels of separation between the guest and the host hardware platform.

Important!

Genesys allows installation and operation of Interaction Media Server in properly configured virtual machines in a non-PureConnect Cloud production environment. However, Genesys does not provide technical support for Interaction Media Server when run as a virtual machine. You are solely responsible for the configuration, performance, maintenance, and troubleshooting of Interaction Media Server as a virtual machine.

Expectations for virtualized Interaction Media Server

Virtualization of Interaction Media Server results in vastly-reduced performance and capacities as compared to installing and using Interaction Media Server on a host without a hypervisor.

For Interaction Media Server, media processing capacity is rated as points-per-media-engine (ppme). A media engine is a thread that runs on a CPU core that processes media, including audio, recording, speech analytics, transcription, transcoding, call analysis, and other features and operations.

For example, a recorded audio interaction using the TrueSpeech codec for both endpoints without any other features costs 1 point. If other features, such as speech analysis or transcoding, are used for that interaction, the point cost increases. When installed as the only software on a Windows Server operating system that is not virtualized, Interaction Media Server achieves 40 ppme. In tests we conducted, Interaction Media Server - when run as a virtual machine (VM) guest in a specific configuration with static reservations - can achieve, at most, 30 to 35 ppme. Other factors, such as the rated speed of the processors, network bandwidth, RAM speed, the hypervisor, and others, can further reduce the achievable ppme on the virtual CPU resources assigned to a VM guest.

Important!

VMWare vSphere does not allow usage of hyper-threaded CPU cores in virtual machines. As a result, a hypervisor host with two processor sockets at eight CPU cores per socket has 16 physical cores available under VMWare vSphere as opposed to 32 physical and hyper-threaded cores when used as a server running Hyper-V on Windows Server 2012 R2 or without a hypervisor.

As Interaction Media Server requires 100% vCPU reservation with high shares, you cannot assign all host CPUs to VM guests. The hypervisor host requires some CPU processing resources and reserves them.

One final consideration is that of the additional costs to achieve media processing capacities in a virtual environment. Because of the limitations of the hypervisor host and VM guest performance, you must purchase approximately three times the number of hardware servers for a virtual environment to match the media processing performance of a single, non-hypervisor installation of Interaction Media Server. Additionally, you would then have to purchase several more Interaction Media Server Software-Only licenses. For example, on a server with an Intel dual-socket E5-2690 running at 2.9 GHz, you have a total of 32 CPU cores (including hyper-threaded cores). You can install Interaction Media Server with a single Software-Only license on that server and achieve approximately 1280 points of media processing. To achieve a similar level of media processing power in a virtual environment with the same hardware specifications would require three additional hardware servers and nine additional Interaction Media Server Software-Only licenses.

Host recommendations for an Interaction Media Server virtual machine

Requirement area	Specifications
Virtualization software	<ul style="list-style-type: none"> VMWare vSphere 5.1 or higher VMWare vSphere ESXi 6.X Microsoft Hyper-V on Windows Server 2012 R2 <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p>Note: Microsoft Hyper-V on Windows Server 2008 R2 is not approved for virtualization of Interaction Media Server. For more information on Hyper-V 2008 R2 limitations with Interaction Media Server, see Testing results of an Interaction Media Server guest in Hyper-V.</p> </div>
Networking	<ul style="list-style-type: none"> NIC Two or more network interface cards with a bandwidth rating of 1Gb/s or higher. NIC configuration We recommend teaming NICs in a Fault Tolerance configuration with failover capabilities (two or more NICs available for communication but only one NIC is active at any given time) <p>Important! Do not use a load balancing NIC teaming configuration with an Interaction Media Server guest VM. You must use reservations (vSphere) or a combination of weight-based and minimum bandwidth rules (Hyper-V) to ensure that the Interaction Media Server VM guest communications are not delayed or interrupted by other guests or the host.</p> <ul style="list-style-type: none"> Virtual Network Switch The Virtual Network Switch must be dedicated to the Interaction Media Server guest VM. RSS/vRSS If available, you must enable Receive Side Scaling(RSS) on the physical NICs and its virtual counterparts (vRSS) to ensure that bandwidth usage is distributed among multiple network interfaces.
Central Processing Unit (CPU)	<ul style="list-style-type: none"> Intel Xeon E5-2690 or better, two sockets on each blade A speed rating of 2.9 GHz or better <p>Important! The host CPU load must not exceed 60% utilization. Failure to keep the host CPU load under 60% can result in problems, such as call audio issues or unanswered calls. Consult the documentation of your hypervisor for methods of determining CPU usage.</p>
Virtual drive source	<p>Reservation with High Shares</p> <p>A reservation of data storage resources does not protect the Interaction Media Server guest VM from interference by other guests. An oversubscribed host can cause poor performance of the Interaction Media Server guest VM and can result in the inability to answer calls.</p>
Clock synchronization	<p>For VMWare vSphere, we recommend configuring the Windows Server guest clock source (w32time) through the information contained at the following website: http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1318</p> <p>For Microsoft Hyper-V 2012 R2, we recommend using the w32tm command in the virtual machine instance to query its time synchronization sources and to synchronize the time to that of the host.</p>

Virtual machine requirements for Interaction Media Server

Requirement area	Specifications
Guest operating system	Microsoft Windows Server 2008 R2 Microsoft Windows Server 2012 R2
Virtual machine vCPUs	4 cores assigned to each Interaction Media Server guest Important! Interaction Media Server performance degrades if you assign more than 4 vCPUs to its guest VM. Important! Interaction Media Server requires 100% CPU reservation in the guest VM.
Networking	<ul style="list-style-type: none"> • Network - The Interaction Media Server guest VM requires a fully-dedicated network. • Network Interface Card (NIC) - The Interaction Media Server guest VM requires a fully-dedicated physical network interface. • Virtual Network Switch - The Interaction Media Server guest VM requires a dedicated Virtual Network Switch. • RSS/vRSS - Use RSS/vRSS to distribute network communications over multiple network interfaces. Important! Network traffic must not delay the transmission and receipt of RTP packets or you may experience audio issues.
Virtual RAM	2 GB
Bandwidth usage	The Network Sent utilization of the Interaction Media Server guest VM must not exceed 140 Mbps. Does the above statement still hold true? Consult the documentation of your hypervisor for methods of determining bandwidth usage.
Guest antivirus	Note: See http://testlab.genesys.com for the latest information on approved antivirus products for Interaction Media Server.

Required Interaction Media Server/CIC server configuration

Requirement area	Specifications
Codec (for recorded calls)	TrueSpeech
CPU mask	Ensure that Interaction Media Server is configured to use all vCPUs assigned to the guest VM. Use the Media Engine CPU Mask control in the Config-Parameters section of the Interaction Media Server web interface.

Testing results of an Interaction Media Server guest in Hyper-V

In lab tests, we observed a measurable difference in the ability of virtualization platforms to provide an adequately accurate clocking source required for real-time applications. When an Interaction Media Server is unable to obtain an accurate clock source, it inserts short periods of silence into the audio path and recordings. Users will perceive these silences as "pops". These artifacts also reduce the accuracy of call analysis and speech recognition.

Note:

Genesys does not advocate or support virtualization of Interaction Media Server on Microsoft Hyper-V on Microsoft Windows Server 2008 R2 or 2012 R2. Genesys does not provide technical support for virtualized instances of Interaction Media Server.

Interaction Media Server virtualization can be used in small development and testing environments using Microsoft Hyper-V on Windows Server. However, our tests consistently encountered poor audio quality and, in some cases, missing portions of audio recordings. Overall performance of Interaction Media Server in a Microsoft Hyper-V virtual machine on Windows Server 2008 R2 is unacceptable for this mission-critical product.


Subsequent tests indicate that Microsoft Hyper-V on Windows Server 2012 R2 has improved performance compared to Windows Server 2008 R2. However, we still recommend using a traditional server installation for Interaction Media Server instead of virtualization.

Note:

The following test results were done with Microsoft Hyper-V on Windows Server 2008 R2. Testing of Hyper-V on Windows Server 2012 R2 shows some improvement but the basic limitations of Interaction Media Server virtualization still apply.

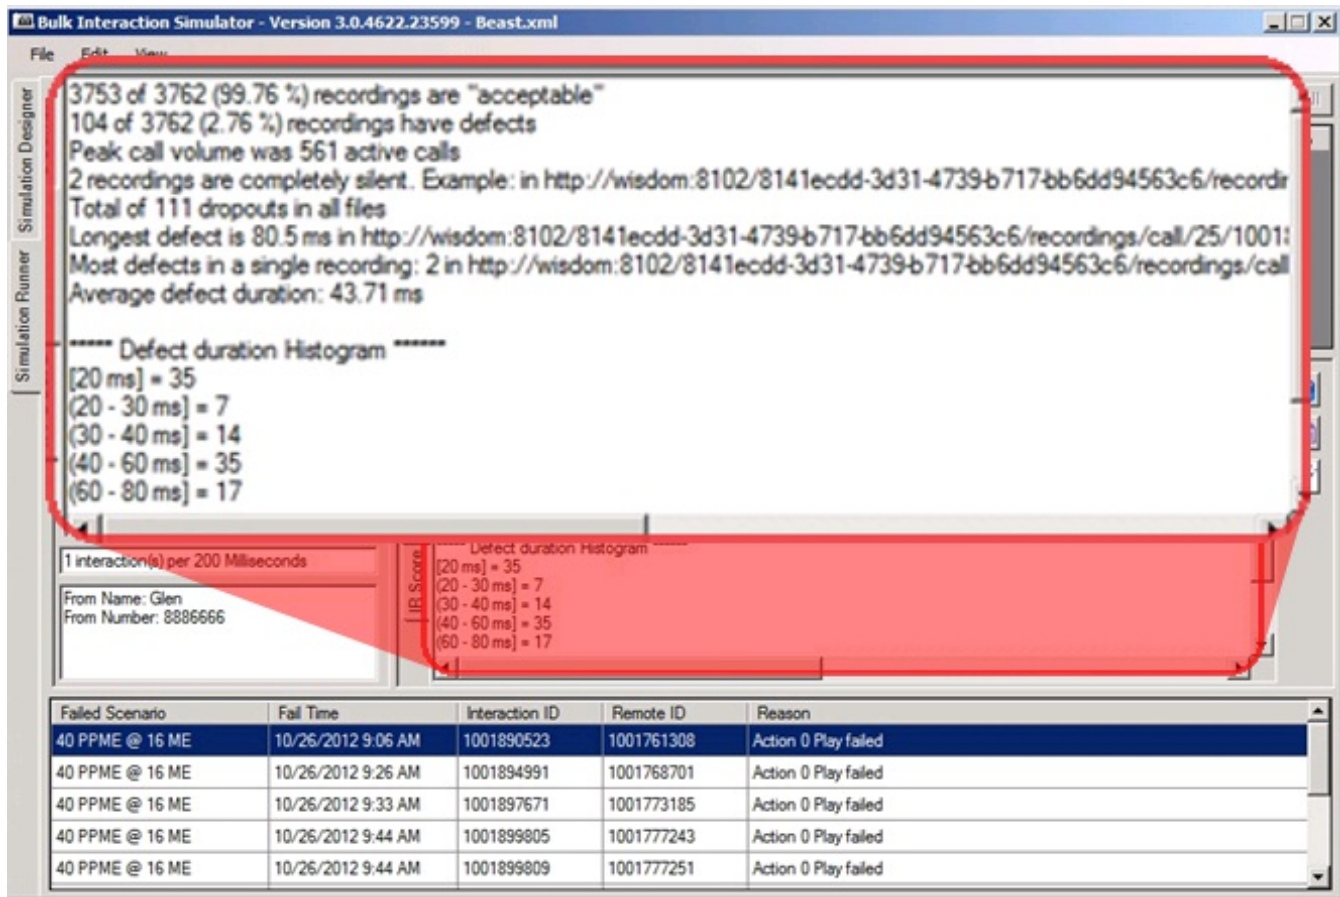
Audio quality issues

The following image shows a comparison of a recording with periods of silence, compared to the original audio file.

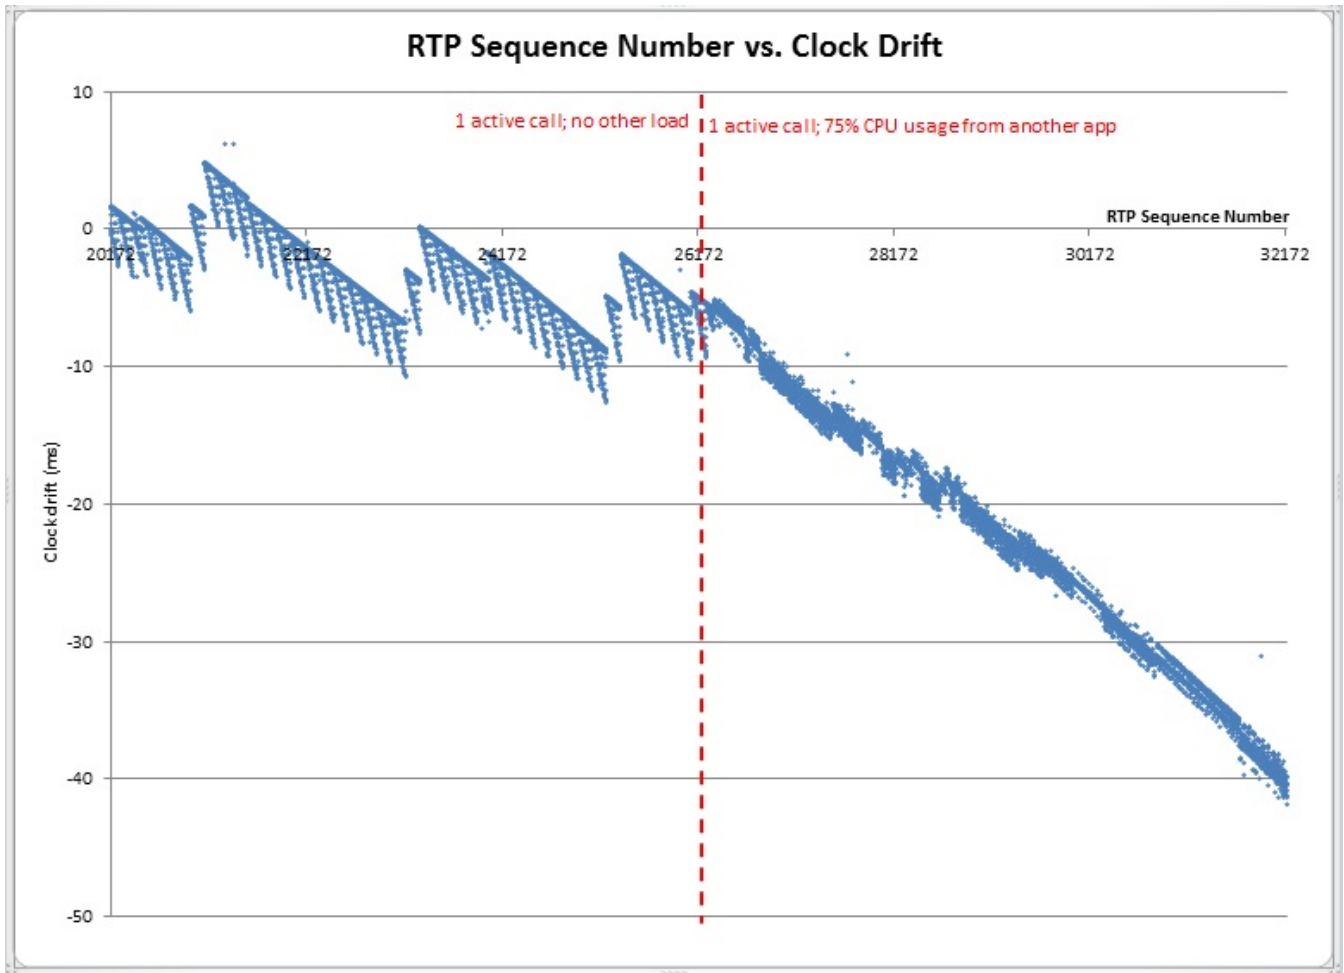


The Testing team uses tools that measure both segments of the audio stream to determine call quality. In the following example, our tools indicate that 2.76% of the audio has defects when inbound callers are conversing with agents. The highlighted area of the following image indicates the quality issues found during a test.

Clock synchronization issues



We reported a known problem to Microsoft with an unreliable clock source being provided to the virtual machine guest. The following image displays unacceptable clock drift of one call in an Interaction Media Server over time.




Networking issues

Another potential problem with virtualization on Hyper-V is the inability of the host to provide a stable and reliable network layer to the guest. RTP packets that carry audio and voice over IP networks are extremely sensitive to latency. Advanced mechanisms in the Interaction Media Server can compensate for some loss and delay, but it has limits to the improvements that it makes. In our tests, Hyper-V on Windows Server 2008 R2 and Windows Server 2012 R2 could not provide an isolated and reliable network layer for hundreds of simultaneous RTP streams.

When an Interaction Media Server records calls, it temporarily stores recordings locally before moving them to an Interaction Recorder Remote Content Service (RCS) system or Interaction Recorder. The following image displays how jitter was introduced when Interaction Media Server transmits compressed RTP recordings. In lab tests, the network layer starts to fail when RTP approaches the 120 Mbps level. Dedicated NICs do not help to alleviate this problem on Hyper-V hosts.

Audio Issues Leaving the Host



The following image indicates that the sent (transmit) "leg" is usually the first to have quality problems. This problem is a result of the variability of the data transfer of recordings with RTP streams from Interaction Media Server to another network storage

location. The receive (sent) "leg" only has RTP streams coming into Interaction Media Server, but the problem will occur after a longer period of time. For this reason, it is important to evaluate both sides of the RTP steaming channels.

Wireshark: RTP Stream Analysis

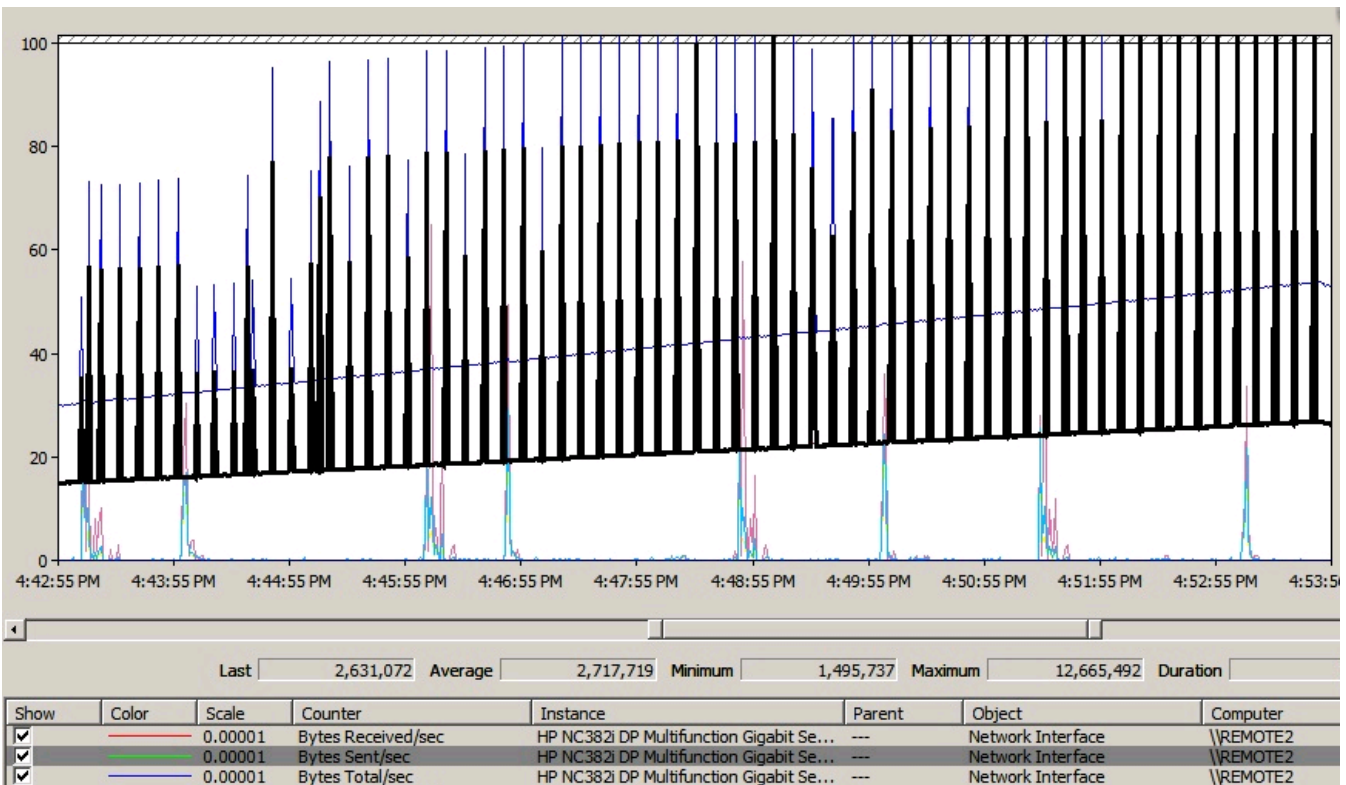
Forward Direction | Reversed Direction

Analysing stream from 174.20.0.22 port 19380 to 174.20.0.17 port 21218 SSRC = 0x1E26788C

Packet	Sequence	Delta(ms)	Filtered Jitter(ms)	Skew(ms)	IP BW(kbps)	Marker	Status
353	10450	20.00	0.39	1.19	81.60		[Ok]
354	10451	19.96	0.37	1.18	81.60		[Ok]
355	10452	20.02	0.35	1.16	81.60		[Ok]
356	10453	20.93	0.38	0.23	80.00		[Ok]
357	10454	19.06	0.42	1.17	81.60		[Ok]
358	10455	20.93	0.45	0.24	80.00		[Ok]
359	10456	19.18	0.47	1.06	81.60		[Ok]
360	10457	65.71	3.30	-44.66	76.80		[Ok]
361	10458	12.84	3.54	-37.49	78.40		[Ok]
362	10459	8.86	4.02	-26.35	78.40		[Ok]
363	10460	12.20	4.25	-18.55	80.00		[Ok]
364	10461	9.79	4.63	-8.34	80.00		[Ok]
365	10462	10.77	4.92	0.90	80.00		[Ok]
366	10463	20.77	4.66	0.13	80.00		[Ok]
367	10464	20.10	4.37	0.04	80.00		[Ok]
368	10465	18.89	4.17	1.14	81.60		[Ok]
369	10466	19.97	3.91	1.17	81.60		[Ok]
370	10467	19.92	3.67	1.25	81.60		[Ok]

Max delta = 65.72 ms at packet no. 360
 Max jitter = 6.27 ms. Mean jitter = 0.68 ms.
 Max skew = -44.66 ms.
 Total RTP packets = 944 (expected 944) Lost RTP packets = 0 (0.00%) Sequence errors = 0
 Duration 18.86 s (0 ms clock drift, corresponding to 8000 Hz (+0.00%))

Save payload... Save as CSV... Refresh Jump to Graph Player Next non-Ok Close



However, usage of virtual Receive Side Scaling (vRSS) in Hyper-V guests on Windows Server 2012 R2 significantly improves network performance for an Interaction Media Server Hyper-V guest on Windows Server 2012 R2.

Testing results of Interaction Media Server running natively on a Hyper-V host

We are often asked about the ability to run time-critical applications in parallel with a Microsoft Hyper-V host on a server. That is, can Interaction Media Server run on the Hyper-V host, not as a virtual machine, but installed and running on the host? The short answer is "don't do it." We strongly advise against using this configuration in any kind of environment.


The goal of the testing was to run Interaction Media Server in the host partition and run a number of common CIC servers as VM guests. We wanted to see how the guests could affect the host and if it was something that could be supported. The objective was to give Interaction Media Server access to more than 4 virtual CPUs (the maximum allowed in Hyper-V at that time) and a real-clock source with minimal latency.

Hyper-V is a Type 1 hypervisor. It uses a host or parent partition and guest partitions that run through the hypervisor. It might be helpful to think of a Type 1 hypervisor as the kernel layer for a physical machine.

The test used three guest VMs at 95% CPU utilization with reservations. Windows allows you to reserve up to the amount of logical processors seen included with hyper-threading. A test of this configuration used Interaction Media Server handling 500 encrypted TrueSpeech recordings at a rate of 6 calls per second (cps) and 1000 ACD agents.

The test results indicate that no host reservation for parallel processes is possible. Guest VMs do cause problems for the host as to its ability to provide suitable CPU resources to parallel processes. In this test, obvious audio quality issues when using an IVR were present. Pausing the guest VMs immediately resolves the problem.

The following image displays the network performance graphs for audio streams on an Interaction Media Server host with no guest VM contention and an Interaction Media Server with contention from the guest VMs.



The guest VM contention causes Interaction Media Server to miss numerous RTP packets, which results in bad MOS scores and audio gaps in recordings.

Apparently the Hyper-V host operating system has no protection from the guests and may even be worse when sharing the hypervisor with guests that use reservations. The following webpage contains a diagram showing what the host is sharing. However, it also contains a statement that the host has direct access to the hardware:

[http://msdn.microsoft.com/en-us/library/cc768520\(v=bts.10\).aspx](http://msdn.microsoft.com/en-us/library/cc768520(v=bts.10).aspx)

Important!

Genesys does not support any PureConnect products in the host partition of a Hyper-V server.

Sizing and performance of a virtual CIC server guest VM

[Sizing the Customer Interaction Center server](#)

[Small implementation](#)

[Medium implementation](#)

[Large implementation](#)

[Very large implementation](#)

[Using Customer Interaction Center with Interaction Process Automation \(IPA\)](#)

[Using Dialer-only or blended environments](#)

Sizing the Customer Interaction Center server

The following information is based on inbound contact center environments only. The small and medium virtual implementations have been evaluated in lab tests with satisfactory performance when using the following configurations and recommendations.

The following information is with regard to the virtual machine guest for the CIC server. Depending on other virtual machines running on the host, the host will require a hardware configuration to adequately satisfy its virtual machine needs. For more information, see [Hardware recommendations for a host](#). Consult your virtualization platform vendor for assistance in determining appropriate hardware specifications for your environment.

Important!

We tested virtualized CIC servers up to 8 calls per second with 1000 ACD agents while recording all calls. These tests had tracing set at a level of 80. As the Large and Very large implementation specifications are theoretical and not tested, you may or may not observe a slightly higher call processing rate.

Small implementation

For up to 300 ACD agents and 30 supervisors:

- Two virtual processors with 100% reservation highly recommended
- 4 GB RAM minimum
- 300 GB hard drive minimum^[1]

[1] For sites that may require increased tracing for troubleshooting purposes, it is highly possible that more than the minimum disk requirement will be needed. In some cases where increased tracing was performed, 600 GB was the increased disk size. Your individual requirements and results may vary.

Note:

The small implementation requirements may not consistently provide a processing rate of 8 calls per second. The call rate depends on many variables in your virtualized environment and the specific tasks that the CIC server is processing.

Medium implementation

For 300 - 1000 ACD agents and 30 100 supervisors:

- Four virtual processors with 100% reservation highly recommended
- 6 GB RAM minimum
- 300 GB hard drive minimum [1]
- For VMware, a dedicated NIC is recommended

[1] For sites that may require increased tracing for troubleshooting purposes, it is highly possible that more than the minimum disk requirement will be needed. In some cases where increased tracing was performed, 600 GB was the increased disk size. Your individual requirements and results may vary.

Large implementation

Important!

Much of the large-scale testing data and experience that we accumulated has been with physical hardware.

The following specifications are from our physical servers. We have no reason to believe that the hardware specifications for our physical servers would be lessened in a virtual environment. If anything, the virtual server requirements are greater than the physical hardware, due to the nature of any virtual environment.

For 1000 - 2500 ACD agents and 100 – 250 supervisors:

- Eight virtual processors with 100% reservation required
- 12 GB RAM minimum
- 300 GB hard drive minimum
- For VMware, a dedicated NIC is recommended

Very large implementation

Important!

Much of the large-scale testing data and experience that we accumulated has been with physical hardware.

The following specifications are from our physical servers. We have no reason to believe that the hardware specifications for our physical servers would be lessened in a virtual environment. If anything, the virtual server requirements are greater than the physical hardware, due to the nature of any virtual environment.

For 2500 - 4000 ACD agents and 250-400 supervisors:

- 16 virtual processors minimum with 100% reservation required
- 12 GB RAM minimum
- 300 GB hard drive minimum^[2]
- For VMware, a dedicated NIC is recommended

[2]For sites that may require increased tracing for troubleshooting purposes, it is highly possible that more than the minimum disk requirement will be needed. In some cases where increased tracing was performed, 600 GB was the increased disk size. Your individual requirements and results may vary.

Using Customer Interaction Center with Interaction Process Automation (IPA)

Interaction Process Automation (IPA) has higher memory requirements for the CIC server. Testing has shown that using a smaller-sized, relatively simple process (around six variables), with 0 - 50,000 processes in the system requires approximately 2 GB of additional memory in the CIC server supporting the contact center. For a similarly sized process, running up 50,000 - 100,000 processes, adding an additional 4 GB of memory is required.

For Customer Interaction Center 4.0 and later, changes were made to allow support for processes to use up to 4 GB of memory. Since IPA processes can remain on the system for weeks or months at a time, it is possible for ProcessAutomationServerU.exe to use up to 4 GB of memory. When creating an IPA process, we recommend that you use variable scope when creating variables. Process level variables are similar to global variables and they remain resident in memory for the life of the process. If a variable does not need to retain information for the life of the process, create variables with State, Tasks, or work item scopes.

It is important to note that the number of variables defined in a process affects the maximum number of active processes that can successfully run on a CIC server.

The following URL address provides information on Windows memory limits:

<http://msdn.microsoft.com/en-us/library/windows/desktop/aa366778%28v=vs.85%29.aspx>

Using Dialer-only or blended environments

To date, our testing in a virtual environment has been with either inbound call center or inbound Dialer scenarios. We have not done extensive testing with the combination of the two scenarios. From a general application capacity standpoint, however, the outbound dialer call rate yields to inbound calls.

For Interaction Dialer (outbound), testing focused on the Central Campaign Server (CCS) as a virtual machine.

A stand-alone CCS has the following minimum configuration requirements:

- Two virtual processors with 100% reservation highly recommended
- 2 GB RAM minimum
- 300 GB hard drive minimum
- No more than 50 call per second (CPS) on average, with one outbound campaign, 10% connect rate, 2 million record call list, and 500 agents

A CCS and database server combination has the following minimum requirements:

- Four virtual processors with 100% reservation highly recommended
- 12 GB RAM minimum
- 300 GB hard drive minimum for the CCS and database server application, plus appropriate hard drive space for the database.

For optimum database server performance and database size planning for the database, there are a number of Technical Reference documents available in the [Documentation Library](#) with specific guidelines. Database server documentation from vendors is also available at the websites for those vendors.

Input/output operations per second (IOPs)

IOPs is a storage device industry term that refers to the number of input/output operations per second on a storage device. IOP is a measure of how fast a system can accept or produce a given dataset or file. Hardware vendors use this term in their hardware specifications. Most PureConnect customers want to know "how many IOPs is the CIC server?" or "How does CIC affect the database server, in IOPs?" or "Does my hardware have enough IOPs?" Be careful comparing anything with a vendor's claimed IOPs number. Once you factor in fibre or iSCSI, virtualization, operating system, and application, disk IOPs are not going to equal the advertised values of bare hardware.

IOPs can also be an indication of workload generated by an application, like a CIC server or database server. Looking at IOPs from this point of view, one can make an assessment of the IOPs workload that the CIC server brings to an environment.

The application workload is an artifact of several things including file access and file access frequency, how an application uses cache versus disk reads and writes, the overall type of processing, and so on. The CIC server is a data intensive application. The storage system must be fast enough to keep the CIC server from waiting and delaying processing. Since the CIC server is event driven and data intensive, a high IOPs requirement met with a slow storage system will result in potentially substandard and inconsistent application performance.

In Windows Performance Monitor (Perfmon), the counter **PhysicalDisk>Disk Transfer/sec** is what the Testing department studies to learn about the disk load that the CIC server is putting on the database server.

We have collected some Performance Monitor (Perfmon) data from SQL Server, during some high-load tests:

CIC scenario

- 4000 contact center users and 400 supervisors on CIC server
- 20 workgroups
- Workgroup membership: One workgroup membership per user; 200 users per workgroup
- Approximately 3-4min call duration
- 20CPS, with recordings
- Alerts were configured for supervisors and clients

The data was gathered during a test run over 14 hrs.

Perfmon Counter for the Datastore drive	Avg	Max
PhysicalDisk>Disk Transfers/sec	507	28,767
PhysicalDisk>Disk Reads/sec	71	3,320
PhysicalDisk>Disk Writes/sec	436	28,740

Pay attention to the **Maximum** values but consider them differently from the **Average** values. For example, a database query that is run during the test would make the **Disk Read/sec Maximum** higher than if the additional queries are not run.

Dialer testing results with regard to IOPs

We also collected Perfmon data during some Interaction Dialer testing:

Interaction Dialer scenario #1

- 40cps with about a 10% connect rate, recorded
- 5000 calls maximum
- One campaign
- Controlled database access (few database queries during this testing)

Perfmon Counter	Operating System		SQL Server Data Files		SQL Server Log Files		SQL Server Transaction Log Files	
	C:\Avg	C:\Max	D:\Avg	D:\Max	E:\Avg	E:\Max	F:\Avg	F:\Max
PhysicalDisk> Disk Transfers/sec	127.514	833.674	81.54	3183.79	41.545	386.037	1.151	28.154
PhysicalDisk> Disk Reads/sec	18.45	452.579	64.069	3174.05	0.001	0.266	0.191	13.343
PhysicalDisk> Disk Writes/sec	109.063	797.663	17.479	452.32	41.544	386.037	0.96	14.811

Interaction Dialer scenario #2

- 40cps with about a 10% connect rate, recorded
- 5000 calls maximum
- 18 campaigns
- Controlled database access (few database queries during this testing)

Perfmon Counter	Operating System		SQL Server Data Files		SQL Server Log Files		SQL Server Transaction Log Files	
	C:\Avg	C:\Max	D:\Avg	D:\Max	E:\Avg	E:\Max	F:\Avg	F:\Max
PhysicalDisk> Disk Transfers/sec	175.21	754.677	2224.305	5476.231	184.164	510.797	4.88	95.97
PhysicalDisk> Disk Reads/sec	85.292	466.683	2084.351	5418.996	0.041	2.668	2.304	45.951
PhysicalDisk> Disk Writes/sec	89.918	549.142	139.954	1503.727	184.123	510.797	2.577	50.019


Pay attention to the **Maximum** values but consider them differently from the **Average** values. For example: A database query during the test would make the **Disk Read/sec Maximum** higher than if the additional queries are not run.

Hyper-V Live Migration and VMware VMotion


Redundancy is built into the full suite of PureConnect products. The fact that redundancy is already available gives the PureConnect software application layer awareness.

For a standard CIC server pair running with a Primary and a Backup, the application state is not replicated, nor would it be desirable. Problems such as memory corruption, viruses, or application faults would be replicated as well.

In the diagram below, the Primary Server (IC1) is a guest virtual machine on Host A and the Backup Server (IC2) is a guest virtual machine on Host B.



If Power is lost to Host A, the VM software detects the failure and tries to restore the servers running on Host A. In the meantime, the server on Host B has taken over the call processing. The state information is lost just as it would be with a real server.



It is not recommended to move VMs handling load in production. In lab testing, the timing can be sensitive on machines handling multiple interactions per second. Some real-time data can be potentially lost. It is for this reason that in a VMware environment - we strongly recommend assigning VMotion traffic to a dedicated NIC. Live Migration traffic should also run on a separate NIC away from the guests and management traffic.

It is a best practice to run redundant PureConnect products on separate physical hosts to take advantage of the hardware redundancy they afford. If a host restarts (even in a clustered environment), all guests on that host are restarted. Running similar guests on the same host does not provide the flexibility for services driver updates and software updates to a host.

Important!

Test your resiliency strategy before going into production. Ensure that the guests can fail from different hosts in the cluster without causing a licensing issue.

VMware High Availability is a feature that provides multiple hosts in a cluster with the capability to recognize when a vSphere host failure has occurred and to respond to the resultant failure of guest virtual machines. Any guest virtual machines that were on the failed host are powered off (like a power failure on physical servers). This may be fine for redundant services like Off-Host Session Manager, but an unscheduled restart of a CIC server during business hours is typically unacceptable. To mitigate this risk present in any virtualization platform, we recommend the use of Switchover systems to maintain the best fault tolerance available.

Many times an administrator does not know what host is running a guest in a large VM environment. Consider this scenario when running multiple non-redundant applications on a single server even with a cluster. When using VMware Distributed Resource Scheduler (DRS), we recommend creating rules to prevent VMotion from moving both servers to the same host, rack, and circuit where possible.

VMware introduced Fault Tolerance (FT) in vSphere 4.0. This feature creates a secondary copy of the virtual machine on a different host in the farm and replicates all instruction to this secondary machine. If the host of the FT enabled virtual machine fails, the secondary continues on as the host and a new clone is created. Fault tolerance can avoid a restart of a virtual machine if the physical host fails. However, it does not replace application level failover as the FT-created clone cannot be managed independently of the FT-enabled virtual machine (that is, a restart affects both). There are several requirements that must be met, including a limit to one virtual CPU (vCPU). If the virtual machine running PureConnect products requires more than one vCPU to meet the needs of your environment, this feature will not be available to it.

The Testing department has not yet performed an analysis of the FT feature.

Important!

We highly recommend that you first test the Fault Tolerance feature in vSphere to ensure that it is meeting your expectations before moving it forward to a production environment.

Data protection

There are a number of software solutions available for backing up VMs. Backup Exec, Microsoft Data Protection Manager, Acronis, and ARCserve are a few third-party solutions. SANs and the native features of the virtualization vendor may also provide some solutions. Bare metal recovery is simplified with virtualization and works well. Many solutions have the ability to throttle the bandwidth and set hours for restriction. Most servers no longer require downtime for a complete backup.

The most important thing to validate is that the solution you chose does not disrupt the performance of the VM guest if scheduled downtime is unacceptable for your environment. For example, in a Switchover pair the Secondary server switches and becomes the Primary if the backup software causes any network disruptions or significant delay to the guest. Using a tool like the SIP Bulk Caller can help test the readiness of a backup solution before it is used in production.

Most vendors tell you that a snapshot is not a sufficient backup. A snapshot in Hyper-V is not the same as VMware. If you delete a snapshot in Hyper-V, don't restart the machine and then continue to run for an extended time. The time it takes to merge the differencing disk becomes exponentially longer. The next time the VM is restarted, such as for a software update or server maintenance, the differencing disk will be merged with the production VHD. During that time, the VM is not available until the merge completes. In lab tests, we have also observed a performance hit when using a snapshot in Hyper-V.

Important!

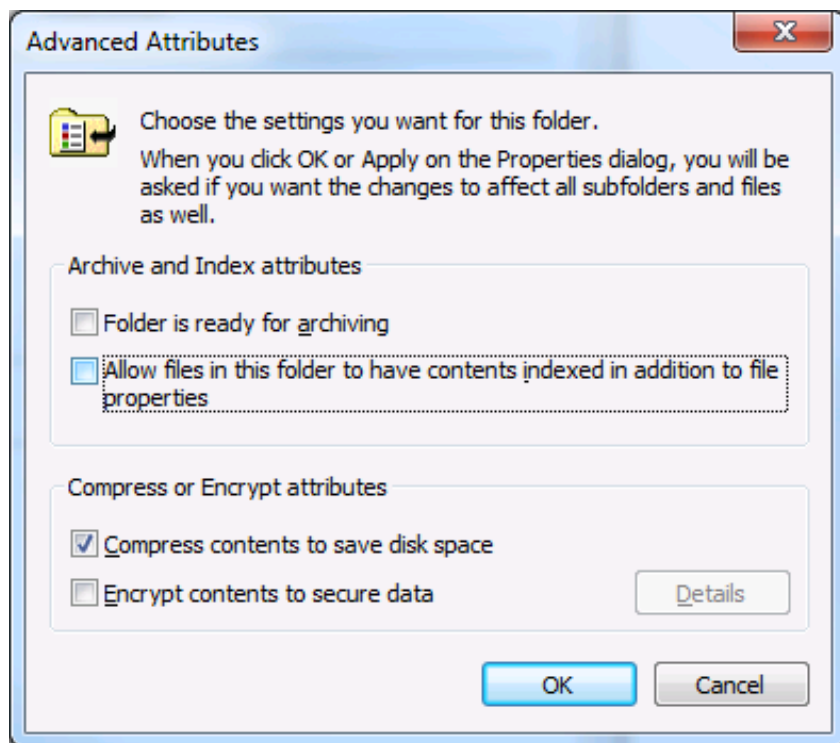
We recommend not using snapshots as a backup of your configuration as they cause a Switchover of the CIC server.

Logging drive

PureConnect uses a feature called Application Buffering, which buffers tracing data before it is written to trace log. Application Buffering uses a default size of 64 KB. This behavior has an override with an environment variable. The `ININ_TRACE_BUFFER_SIZE` environment variable allows buffer size to be adjusted from 0 (off) through 512. You specify size in kilobytes, so 512 would be 512 KB.

Windows Disk Compression on Logging folders

Consider using Windows Disk Compression on your Logging folder. This compression is only supported on the trace logs and not the entire drive or PureConnect program files.



It is possible to put tracing on a separate disk for other reasons. Many times, the cost of SAN disk space is much more expensive than another disk solution. PureConnect has been designed to trace less at default and verbose tracing levels thereby using less disk space.

Note:

For sites that may require increased tracing for troubleshooting purposes, it is highly possible that more than the minimum disk requirement will be needed. In some cases where increased tracing was performed, 600GB was the increased disk size. Your individual requirements and results may vary.

Another consideration is if you want to back up the tracing data. Depending on the backup strategy and solution chosen, it may not be possible to back up the guest without also backing up the logging drive. For example, VMs that are backed up at the host level with data position measurement (DPM) get all the drives so that you cannot choose what to exclude. You could back up the machine at the guest level to get more granularity.

Disk type and size

Hard disks are recommended over dynamic disks. Hard disks give some performance benefit but take longer to initialize. You must properly scope the drive space and a hard disk guarantees the PureConnect products have the space they need.

- In Hyper-V, it is possible to oversubscribe storage with dynamic VHDS. As a best practice, use fixed VHDS, especially on a SAN.
- In VMware, virtual disks are created as VMDK files and can use thick or thin provisioning. Thick provisioning reserves all of the capacity up front. Thin provisioning dynamically grows and exists to allow increased storage density with an eye toward adding storage later as it is consumed. Both types require zeroing of the storage before data can be written. From a performance perspective, VMware has shown that the difference between thick and thin provisioned volumes is negligible both before and after zeroing (http://www.vmware.com/pdf/vsp_4_thinprov_perf.pdf).

Importance of a testing environment

We strongly recommend that you create a robust testing environment to ensure that you can simulate the load of a production environment.

Careful planning, design, and implementation cannot account for operational variables. Prove the concepts on a smaller scale before trying to deploy them. Discuss changes and document them.

Good things to know

[Running with guest enhancers](#)

[Identifying a virtual machine from the PureConnect logs](#)

[Common issue](#)

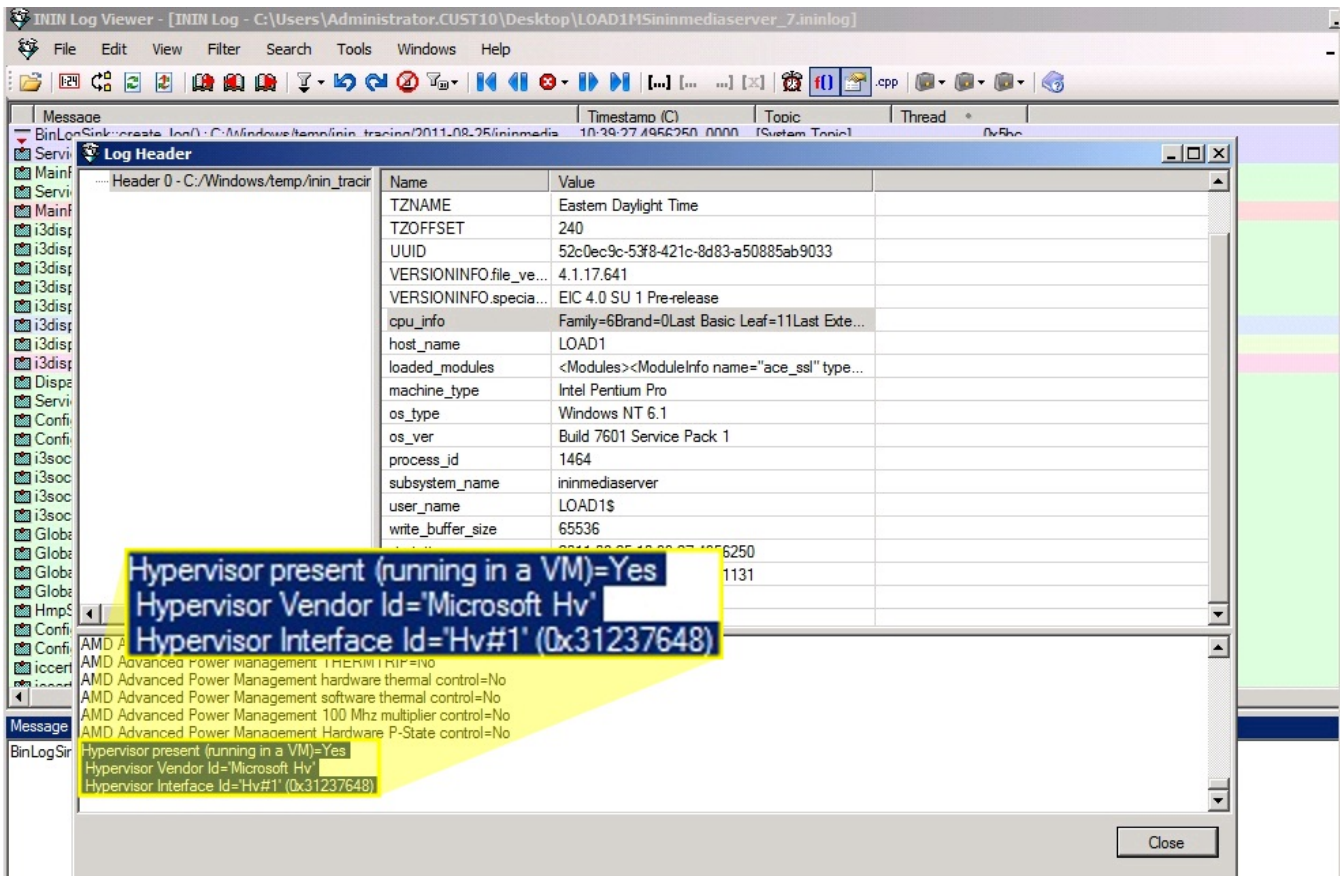
[Rare issue](#)

Running with guest enhancers

We test with guest enhancers enabled. It is easy to forget guest enhancers. It is important that you run your guest OS with the vendor recommended enhancer. VMware provides VMware Tools to improve the graphics, synchronize the time, and improve overall performance. Running a guest on Windows Server 2008 R2 already has the Integrated Services installed. If you are migrating Windows Server 2008 R2 Hyper-V guests to a Windows Server 2012 R2 Hyper-V host, be sure to update the Integrated Services install within the guest. The application event log will contain warnings until you complete the Integrated Services update.

Identifying a virtual machine from the PureConnect logs

You can determine when a machine is running on a virtual machine by viewing the log header of any log files:



Hosts with mismatched CPUs, MAC addresses, and licensing

Common issue

When a license is generated for the CIC server, it requires a CPUID and the MAC address to generate a HOSTID. To ensure that the license will continue to function when a VM host becomes part of a cluster, the CPUID and MAC address must be managed.

The first component of the HOSTID, the CPUID, can be generalized so that guests can move from host to host in a cluster and maintain the same CPUID. If the CPU types are the same for all hosts in the cluster, the general CPUID is usually not problematic. VMware and Hyper-V both provide some mechanisms to accommodate a consistent CPUID when the hosts are not identical.

Important!

Contact the virtualization software vendor if you need assistance in providing a consistent CPUID to guests in a cluster.

Genesys is not responsible for the configuration and performance of your virtualization environment.

Warning!

The second component of the HOSTID, the MAC address, must also be identical for the CIC server guest when run on a different host. Any changes to a virtual environment after you have generated the CIC server license will likely impact the license. For this reason, you must configure the virtual machine guest to have a static MAC address value.

Once you verify that your CPUID and MAC address do not change when the CIC server guest uses different hosts in the cluster, only then should you generate your CIC server license.

The following examples provide common scenarios that result in license issues when the CPUID or MAC address is not maintained across hosts:

Example #1: A customer added a host that was running the CIC server guest in a new host cluster, not realizing it would cause the CPUID to change.

Example #2: The guest running the CIC server guest does a VMotion or Live Migration onto a new server with a different dynamic MAC address pool. The customer had not configured a static MAC address because the MAC address had never changed in the previous environment.

Example #3: The host supporting a CIC server guest failed. The guest was restored on a new host. Even though clustering was not involved, the different CPU platform of the new host machine generated a different CPUID value that invalidated the CIC server license.

Important!

Always test your resiliency strategy before putting CIC guest VMs into production. Ensure that the guests can move to different hosts in the cluster without causing a licensing issue.

Rare issue

A host should be running with the same version of sockets if it has more than one. It is possible to buy a server with one socket and then later upgrade the machine with another socket. The additional socket can cause a situation where the CIC license check fails because the stepping ID, for example, can be different even though the sockets appear to be the same model and clock speed.

The CPUID instruction used for processor identification is different from the CPU with which the license was originally generated. Migration of a virtual machine that is not part of a cluster to another host causes a licensing failure because the CPUID used for the HOSTID is different.

Change Log

Date	Changes
20-May-2011	Initial Draft
06-June-2011	Got some feedback from CaaS
09-June-2011	Input from Testing Team
01-September-2011	More Updates based off Testing Results
07-September-2011	Changes for Licensing, and Log Header info
19-September-2011	Renamed tables for "Supported" and "Not Supported" for virtualization
04-October-2011	Copyediting
31-October-2011	Added new section "Running with Guest Enhancers".
06-November-2011	Added Nuance ASR to "Not Supported" table
16-January-2012	<ul style="list-style-type: none"> • Added new section "Using a Virtualized Interaction Media Server in Development Environment". • Updated copyright information.
09-March-2012	Added ASR Server Nuance Recognizer as "Supported on VMware only" and changed the table description.
20-April-2012	Added information about VMware HA and Switchover in the "Hyper-V Live Migration and VMware vMotion" section.
20-November-2012	Major Rewrite to accommodate Virtual Media Server and Large sized VMs
07-December-2012	Light edit (Acrolinx).
14-December-2012	<ul style="list-style-type: none"> • Created new section "Virtualized Interaction Media Server in a Development and Testing Environment" to clarify that virtual Interaction Media Server is supported only in these environments. • Provided additional information on IOPs and sizing and performance for virtual IC server guests.
18-December-2012	<ul style="list-style-type: none"> • Moved Interaction Media Server table row from "Supported on VMware Only" table to "Supported for Virtualization (Hyper-V and VMware)" table." • Moved section "Licensing a Virtualized Interaction Media Server in a Dev and Testing Environment" after the Hyper-V section since it applied to any virtualized Media Server, and edited heading level. • Added qualification statements to section "Virtual Interaction Media Server Guest in a Hyper-V Environment."
19-July-2013	Added content on virtualization support for Interaction Media Server.
02-May-2014	<ul style="list-style-type: none"> • Updated copyright and trademark statements. • Updated content for using Interaction Media Server in a virtual machine. • Added third-party ASR product information. • Added content advising against usage of snapshots as a means of configuration backup. • Added content to clarify IC server licensing with regards to virtualization and redundancy. • Applied general editing. • Added statements regarding responsibility of the virtualization environment. • Added e-mail address for virtualization inquiries.
14-July-2014	Corrected content regarding performance expectations for Interaction Media Server in a virtual environment.
08-August-2014	<ul style="list-style-type: none"> • Updated documentation to reflect changes required in the transition from version 4.0 SU# to CIC 2015 R1, such as updates to product version numbers, system requirements, installation procedures, references to Interactive Intelligence Product Information site URLs, and copyright and trademark information. • Added Windows Server 2012 R2 as a supported operating system • Removed content regarding Windows Server 2003 as it is no longer supported

02-March-2015	<ul style="list-style-type: none"> • Updated "Copyright and Trademark Information" page • Updated content to reflect changes in corporate policy removing support for virtualization of Interaction Media Server • Minor edits
12-March-2015	Removed "Resources" section that contained URL addresses to obsolete reference material for HP servers and Microsoft Hyper-V
13-October-2015	<ul style="list-style-type: none"> • Updated cover page for corporate branding • Updated "Copyright and trademark information" page • Minor edits
02-February-2016	Added content describing the differences between emulated and paravirtualized/synthetic network adapters available in hypervisors.
03-March-2016	Corrected erroneous e-mail address in "Change Log"
27-May-2016	Updated minimum recommended hardware requirements for the hypervisor host and guest virtual machines in Hardware recommendations for a host .
05-January-2017	<ul style="list-style-type: none"> • Updated Copyright and Trademark Information page • Added ESXi version 6.X as a validated hypervisor in Virtual machine requirements for Interaction Media Server • Added content and links to testlab.genesys.com for the latest tested virtualization products and versions
24-March-2017	<ul style="list-style-type: none"> • Updated instances of organization name • Removed e-mail address for submission of hardware specifications • Minor rebranding
28-April-2017	<ul style="list-style-type: none"> • Updated documentation to reflect the removal of Interaction Client Web Edition.
24-August-2017	<ul style="list-style-type: none"> • Updates to important notes in "Products not supported for virtualization" section and "Virtualizing Interaction Media server in a production environment" section to clarify that virtualization of Interaction Media Server is only for on-premises installations of PureConnect, and does not apply to PureConnect Cloud. • Rebranding. Updated title page, copyright, terminology.
12-January-2018	Conversion to HTML.