GENESYS
AN ALCATEL-LUCENT COMPANY

Genesys 7

Hardware Sizing Guide

About Genesys

Genesys Telecommunications Laboratories, Inc., a subsidiary of Alcatel-Lucent, is 100% focused on software for contact centers. Genesys recognizes that better interactions drive better business and build company reputations. Customer service solutions from Genesys deliver on this promise for Global 2000 enterprises, government organizations, and telecommunications service providers across 80 countries, directing more than 100 million customer interactions every day. Sophisticated routing and reporting across voice, e-mail, and Web channels ensure that customers are quickly connected to the best available resource—the first time. Genesys offers solutions for customer service, help desks, order desks, collections, outbound telesales and service, and workforce management. Visit www.genesyslab.com for more information.

Each product has its own documentation for online viewing at the Genesys Technical Support website or on the Documentation Library DVD, which is available from Genesys upon request. For more information, contact your sales representative.

Notice

Although reasonable effort is made to ensure that the information in this document is complete and accurate at the time of release, Genesys Telecommunications Laboratories, Inc., cannot assume responsibility for any existing errors. Changes and/or corrections to the information contained in this document may be incorporated in future versions.

Your Responsibility for Your System's Security

You are responsible for the security of your system. Product administration to prevent unauthorized use is your responsibility. Your system administrator should read all documents provided with this product to fully understand the features available that reduce your risk of incurring charges for unlicensed use of Genesys products.

Trademarks

Genesys, the Genesys logo, and T-Server are registered trademarks of Genesys Telecommunications Laboratories, Inc. All other trademarks and trade names referred to in this document are the property of other companies. The Crystal monospace font is used by permission of Software Renovation Corporation, www.SoftwareRenovation.com.

Technical Support from VARs

If you have purchased support from a value-added reseller (VAR), please contact the VAR for technical support.

Technical Support from Genesys

If you have purchased support directly from Genesys, please contact Genesys Technical Support at the regional numbers provided on page 16. For complete contact information and procedures, refer to the ***Genesys Technical Support Guide***.

Ordering and Licensing Information

Complete information on ordering and licensing Genesys products can be found in the ***Genesys Licensing Guide.***

Released by

Genesys Telecommunications Laboratories, Inc. www.genesyslab.com

Document Version: 70g_sizing_03-2010_v7.0.008.49

Table of Contents

Preface

Welcome to the *Genesys 7 Hardware Sizing Guide*. This document provides you with system-level information about Genesys hardware sizing guidelines for the Genesys 7.x releases to date.

Note: This manual is continually updated when new Genesys products are released. For the most current version, please visit the Genesys Technical Support Website. To request the Documentation Library DVD, which you can order by e-mail from Genesys Order Management, contact `orderman@genesyslab.com.`

This Preface contains the following sections:

For information about related resources and about the conventions that are used in this document, see the supplementary material starting on page 327.

Note: This guide reflects the products currently available for purchase from Genesys. See Table 1 on page 14.

Intended Audience

The *Genesys 7 Hardware Sizing Guide* is primarily intended for those who make hardware and network LAN/WAN bandwidth purchasing recommendations. It assumes that you have a basic understanding of:

- Computer-telephony integration (CTI) concepts, processes, terminology, and applications.
- Network design and operation.
- Your own network configurations.

You should also be familiar with Genesys Framework architecture and functions.

Note: The Database Administrator must participate in all database sizing and configuration decisions.

Scope

This system-level guide has been prepared by Genesys. It is intended as a pre-sales tool for estimating the hardware required for installations.

The suggestions are based on field experience and are conservative estimates: a slight over-estimation of hardware requirements does not generate the production issues represented by under-estimating these requirements.

While the information in this guide can assist you in choosing hardware appropriate for your implementation, it is not intended to provide detailed information for every possible solution. These guidelines assume an “average” customer configuration of Genesys software deployed at a single site.

Unusual circumstances—such as a configuration with very complex routing strategies, multiple database lookups, or one distributed across multiple geographic locations—may require additional hardware or processors.

Be sure to consider the worst-case or peak scenarios when determining the appropriate configuration for your situation.

Once you have determined the hardware configuration that will accommodate your business needs, we recommend that you verify your proposed hardware solutions with Genesys Professional Services.

Recommendations

This guide shows the average complexity and call-flow scenario for small and medium contact centers, group-based routing, and queue routing. The recommendations are organized by contact center size, operating system, and Genesys solution.

Since hardware models change frequently, you will need to verify that the recommended models coincide with the current information from the hardware vendor.

This document covers four of the more common platforms and operating systems: IBM/AIX, Sun/Solaris, Microsoft Windows, and HP-UX. Other platforms may also be supported. Please check the *Genesys Supported Operating Environment Reference Manual* located on the Technical Support website at:

`http://genesyslab.com/support.`

The solution types, such as Framework and Enterprise Routing, have been factored out as much as possible so that you can assemble the elements for your specific needs.

Installation

It is necessary to use a **DVD ROM** when installing Genesys software.

Sizing Parameters

The parameters that have to be taken into consideration for hardware sizing are the following:

- Number of Agents
- Calls Per Hour
- Service Level
- Average Waiting Time
- Service Time

Note: Parameters cannot be changed independently.
For specific solutions and releases, particular parameters may have a greater impact on sizing. For example, Reporting statistics filters in 6.x have a high performance impact but not in 7.0.

Guide to Recommendations

Table 1 is a guide to recommendations and contains the following information:

- Contact center sizes.
- Corresponding number of agents.
- Maximum number of interactions the contact center should be receiving.
- Operating systems referred to in this document.

Note: Some Genesys solutions use different numbers to determine contact center size categories. In these cases, the different sizing criteria are provided in the applicable chapters in this document.

Table 1: Guide to Recommendations

Contact Center Size	Number of Agents	Average Interactions Per Second
Small Contact Center	<150	1.5
Medium Contact Center	150 to 1000	1.5-10
Large Contact Center	>1000, multi-site, distributed environment	
Operating Systems		
Sun Solaris		
Microsoft Windows		
IBM/AIX		
HP-UX		

Notes:

- The **Number of Agents** varies when using multi-channel routing. See MCR section of this document for specific MCR recommendations.
- The **Average Interactions/Second** numbers represent the average number of interactions per second that can be processed using reasonably complex routing strategies. With simpler routing strategies, the number of interactions per second that can be processed may be higher.
- In terms of **Contact Center Size**, this version of the guide is only intended to cover recommendations for small and medium installations. Large configurations may require special architectural design and should be done only in consultation with Genesys Professional Services.

This guide reflects products currently available for purchase from Genesys (both 6.5 and 7.x products), with the exception of Genesys Express, which is 4.x.

Multi-core/Processors

Genesys software applications are tested against Operating Systems and Databases only, and not specific hardware. Genesys does, however, indicate minimum sizing requirements, such as processor speed and RAM with regards to the various Operating Systems platforms it works on.

The majority of Genesys software has been designed as single-threaded applications and, as such, will run on a variety of hardware, provided the appropriate Operating System is installed and the minimum sizing requirements are met. Genesys is aware that there are new multi-core/processor platform architectures currently being designed and sold that are specifically created to handle new advanced multi-threaded software, such as Java web applications.

Although Genesys products will operate as designed on these new hardware platforms, it is important that you contact your hardware vendor and discuss the levels of performance that will be expected from single-threaded applications on their products. This question not only needs to be asked for Genesys single-threaded applications, but also for any other single-threaded software product you may have in your enterprise environment that is being considered to reside/run on these new architectures.

Genesys is aware of the new trends in hardware and has been monitoring the marketplace closely. Genesys has begun the process, as a result of this monitoring, of planning the required steps necessary to accommodate these changes in the marketplace.

Most hardware vendors who have these new multi-core/processor platform architectures do provide accurate information on their products about performance of various applications.

Making Comments on This Document

If you especially like or dislike anything about this document, feel free to e-mail your comments to Techpubs.webadmin@genesyslab.com.

You can comment on what you regard as specific errors or omissions, and on the accuracy, organization, subject matter, or completeness of this document. Please limit your comments to the scope of this document only and to the way in which the information is presented. Contact your Genesys Account Representative or Genesys Technical Support if you have suggestions about the product itself.

When you send us comments, you grant Genesys a nonexclusive right to use or distribute your comments in any way it believes appropriate, without incurring any obligation to you.

Contacting Genesys Technical Support

If you have purchased support directly from Genesys, contact Genesys Technical Support at the following regional numbers:

Region	Telephone	E-Mail
North and Latin America	+888-369-5555 (toll-free) +506-637-3900	support@genesyslab.com
Europe, Middle East, and Africa	+44-(0)-1276-45-7002	support@genesyslab.co.uk
Asia Pacific	+61-7-3368-6868	support@genesyslab.com.au
Malaysia	1-800-814-472 (toll-free) +61-7-3368-6868	
India	000-800-100-7136 (toll-free) +91-(022)-3918-0537	support@genesyslab.com.au
Japan	+81-3-6361-8950	support@genesyslab.co.jp
Before contacting technical support, refer to the *Genesys Technical Support Guide* for complete contact information and procedures.		

Chapter

1

Small Contact Centers

This chapter presents hardware guidelines and recommendations for small contact centers. Small contact centers generally have less than 150 agents and receive a maximum of 1.5 interactions per second (IPS).

The information in this chapter contains the following topics:

Note: Hardware architecture diagrams of sample configurations are in Chapter 6, "Hardware Architecture Diagrams," on page 41.

Recommendations

This section describes recommendations for small contact centers using Framework, Reporting, Routing, Outbound, Voice Callback, and others.

Table 2 on page 18 shows recommendations for small contact centers using Framework Management and Configuration Layers, T-Server, Historical Reporting, Enterprise Routing and/or Outbound Contact and Voice Callback on Sun Solaris, Microsoft Windows, IBM/AIX, and/or HP-UX operating systems.

All the solutions, including Stat Server and Reporting, may be located on one box or distributed as described in the following table.

Table 2: Framework, Reporting, Routing, Outbound, and Voice Callback

HP-UX Operating System	
OS	HP-UX 11i
Processor Type, Quantity, Speed	PA type processor: HP9000 pr 4410
Memory Size	2 GB RAM
Hard Disk Space	40 GB Hard Disk Drive
Ports	Networking Ports - Two TP Ethernet 1000/100/10 BASE-T cards that work in full duplex mode (one for Framework, one for T-Server)
IBM AIX Operating System	
OS	IBM AIX 5.3
Processor Type, Quantity, Speed	Four CPU System P5 510 type server with 4 GB RAM. Possible choices for CPU configuration could be: • Two dual CPU boxes Or • One single 4 CPU box (e.g., System p5 510Q)
Memory Size	4 GB RAM
Hard Disk Space	36 GB Hard Disk Drive
Ports	Networking Ports - Two TP Ethernet 1000/100/10 BASE-T cards that work in full duplex mode (one for Framework, one for T-Server).

Table 2: Framework, Reporting, Routing, Outbound, and Voice Callback (Continued)

Linux Operating System	
OS	Red Hat Linux Enterprise Server 3, 4, or 5
Processor Type, Quantity, Speed	2 CPU 2.4 MHz Intel Xeon Processor or higher.
Memory Size	4 GB Total RAM
Hard Disk Space	36 GB Hard Disk Drive
Ports	Networking Ports - Two TP Ethernet 1000/100/10 BASE-T cards that work in full duplex mode (one for Framework, one for T-Server).
Microsoft Windows Operating System	
OS	Microsoft Windows 2000/2003
Processor Type, Quantity, Speed	2 CPU 2.6 GHz dual core processor or higher. Possible choices for CPU configuration could be: • Two 1xCPU servers with 2 GB RAM Or: • One 2xCPU server with 4 GB RAM Genesys recommends using a server-class machine for all applications.
Memory Size	4 GB total
Hard Disk Space	80 GB per server
Ports	Networking Ports - 2 Ethernet 1000/100/10 BASE-T ports (one for Framework, one for T-Server)
Sun Solaris Operating System	
OS	Solaris 2.6-2.10, 32- or 64-bit
Processor Type, Quantity, Speed	4 CPU 1.5 GHz UltraSPARC IIIi. Possible choices for CPU configuration could be: • Four 1xCPU servers with 2 GB RAM (e.g., Sun Netra 240) Or • Two 2xCPU servers with 4 GB RAM (e.g. Sun Netra 240) Or • One 4xCPU server with 8 GB RAM (e.g., Sun Netra 440)
Memory Size	8 GB total

Table 2: Framework, Reporting, Routing, Outbound, and Voice Callback (Continued)

Hard Disk Space	146 GB per server
Ports	Networking Ports - 2 Ethernet 1000/100/10 BASE-T ports (one for Framework, one for T-Server)

Notes:

- If a POWER 4 is used, Stat Server may be run on the same box as the other applications. If a lesser machine is used, it might need to be located on a separate box depending on interaction volume. It is recommended that a separate Stat Server should be installed if Genesys Routing is deployed.
  The Database Server may be installed on the same box as the database; however, accurate database sizing information is essential. The number of HDDs and/or size of HDDs might need to be increased, depending on data storage requirements and workflow.
- **Important:** The Database Administrator needs to participate in all database sizing decisions.

Call Progress Detection Server

Table 3 shows recommendations for small contact centers using Call Progress Detection Server on a Microsoft Windows platform, if such a server is required in your configuration.

Table 3: Call Progress Detection Server

Microsoft Windows Operating System	
OS	Microsoft Windows 2000/2003
Processor Type, Quantity, Speed	One Intel Core 2 CPU, min. 2.6 GHz
Memory Size	1 GB RAM
Hard Disk Space	40 GB HDD
CPD Board	See the Supported Dialogic Boards Table in the *Genesys Supported Media Interfaces Reference Manual.*
Ports	Networking Ports - TP Ethernet 1000/100/10 BASE-T card that works in full duplex mode.

Note: Hardware configuration for systems where Call Progress Detection Server is used with Intel HMP software should comply with the requirements specified in Intel HMP documentation.

Multi-Channel Routing and Multimedia

Table 4 shows minimum requirements for small contact centers using Multi-Channel Routing (MCR) and Multimedia on a Microsoft Windows platform. Actual size depends on the type, number, size, and duration of the interactions your solution will be expected to handle.

Notes:

- The name of Multi-Channel Routing is changed to Multimedia beginning with release 7.2.
- For purposes of MCR and Multimedia only, a small contact center is defined as one with 75 or fewer agents.

Table 4: Multi-Channel Routing and Multimedia

Microsoft Windows Operating System	
OS	Microsoft Windows 2003 Server
Processor Type, Quantity, Speed	2xCPU 2.33 GHz processors or better
Memory Size	4 GB RAM
Hard Disk Space	80 GB available hard disk space [a]

a. Hard disk consumption depends not only on the size of the database (the amount of stored contact and interaction history), but also on the accumulated size of the log files (which you can regulate using options in Configuration Manager).

Genesys Agent Desktop

Table 5 shows recommendations for small contact centers using Genesys Agent Desktop on a Microsoft Windows platform.

- Genesys Agent Desktop Server is installed in standalone mode: one Tomcat 5.5 (Catalina) and Apache 2.2.4 in front.
- Java Server Pages should already be compiled using Java Development Kit 1.4.2_xx or 1.5.0_yy from Sun.

- Genesys Agent Desktop Server uses one Tomcat servlet per Genesys Agent Desktop container. The Genesys Agent Desktop Tomcat servlet container handles up to 100 agents.

Table 5: Genesys Agent Desktop

Microsoft Windows Operating System	
OS	Microsoft Windows 2000 Advanced Server Sp3, Microsoft Windows 2003 Advanced Server
Processor Type, Quantity, Speed	2xCPU Pentium IV-2.4 Ghz processor
Memory Size	1 GB SDRAM
Hard Disk Space	40 GB HDD (available hard disk space)
Ports	Networking Ports-TP Ethernet 100 MB full duplex mode
Server	Apache 2.2.4 Web Server

Genesys Express

Table 6 shows recommendations for small contact centers using Genesys Express, which runs Management and Configuration Layer, Framework, Routing, and, as an option, Contact Server, Email, and VTO. Express supports up to 150 agents but no more than 75 concurrently on a same host. For configuration of 75 to 150 agents, install Genesys Desktop on a separate host.

Table 6: Genesys Express

Microsoft Windows Operating System	
OS	Microsoft Windows 2000 Server/ Microsoft Windows 2003 Server
Processor Type, Quantity, Speed	CPU Minimum of Dual Intel Xeon Processors, 1.8 GHz, with 512 KB L2 cache per processor or the equivalent
Hard Drive	40 GB HDD (available hard disk space)
Ports	Networking Ports-TP Ethernet 100 MB full duplex mode
Memory Size	RAM: 2 GB
Cache Memory	512 KB L2

Notes:

- When Express CTI is used with an IVR, the customer and sales representative should use the appropriate sizing tools for IVR. Express CTI IVR can be deployed on the same host up to 24 ports. Routing and Reporting are not part of Express CTI. Only ports are part of Express CTI IVR Interface.
- Genesys Express 4.6 installation can span across several hosts. When Options (SIP Server, Call Concentrator, Outbound) are selected, contact your sales representative.

Genesys Agent Scripting

Table 7 shows the recommendations for small contact centers using Genesys Agent Scripting on a Microsoft Windows Platform:

Note: Genesys Agent Scripting is installed in standalone mode (one Tomcat and Apache or IIS web server.

Table 7: Genesys Agent Scripting

Microsoft Windows Operating System	
OS	Microsoft Windows 2000/2003/XP
Processor Type, Quantity, Speed	One Intel Core 2 CPU, min. 2.6 GHz
Memory Size	1 GB RAM
Hard Disk Space	40 GB HDD
Ports	Networking Ports-TP Ethernet 100 MB full duplex mode
Server	IIS 5.0, IIS 6.0, Apache/Tomcat 4.3.1, or Apache/Tomcat 5.0 Web Server running on Microsoft Windows 2000/2003 server

Note: The sizing information is for running scripts without integration with other products. If scripts are use or integrated with other products, then the sizing requirements for those products must be considered in addition to what is defined above.

Blue Pumpkin Integration

Table 8 shows the Blue Pumpkin Integration (7.1) hardware requirements for small contact centers (<500 agents).

Table 8: Blue Pumpkin Integration

WFM Data Aggregator	
OS	Microsoft Windows 2000, Microsoft Windows Server 2003
Processor Type, Quantity, Speed	Dual Intel Xeon 1 GHz with 1 GB L2 cache
Memory Size	1 GB RAM
Hard Disc Space	2 GB
Ports	Networking ports—TP Ethernet 1000/100 BASE-T card that works in full duplex mode

Client Workstation Requirements

Computers running the Statistics Configuration Utility, at minimum, should meet these specifications:

- Pentium III, 1 GHz or greater, with 1 GB RAM

Genesys Call Director Route

Table 9 shows the Genesys Call Director Route hardware requirements for small contact centers:

Table 9: Genesys Call Director Route

Microsoft Windows Operating System	
OS	Microsoft Windows 2000/2003 Server
Processor Type, Quantity, Speed	Dual Intel Xeon 1 GHz with 1 GB L2 cache
Memory Size	1 GB RAM
Hard Disk Space	40 GB HDD
Ports	Networking Ports—TP Ethernet 100 MB full duplex mode
Server	Tomcat 4.3.1 or Tomcat 5.0 Web Server

User Interface Hardware

User interfaces such as Configuration Management Environment (CME), Solution Control Interface (SCI), and CCPulse+, should use Microsoft Windows workstations with one 2.6 GHz CPU or better, and at least 1 GB of memory.

For Genesys Agent Desktop and Genesys Supervisor Desktop, a Microsoft Windows workstation with one 2.6GHz or better Intel Core 2 CPU, and at least 1GB of memory are recommended. If non-Genesys applications are being run on the same machine, 2GB of memory is recommended.

Genesys Agent and Genesys Supervisor Desktops also support Mac OS X workstations with one 1.5 GHz CPU or better, and at least 1 GB of memory.

Chapter

2 Medium Contact Centers

This chapter presents hardware guidelines and recommendations for medium contact centers. Medium contact centers are generally those that have from 150 to 1000 agents and receive a maximum of 12 interactions per second (IPS).

The information in this chapter contains the following topics:

Hardware architecture diagrams of sample configurations are in Chapter 6, "Hardware Architecture Diagrams," on page 41.

Note: You have to run License Server on one of the servers. This server requires a good network connection to other servers on which Genesys software is installed.

Framework, Reporting, Routing, Outbound, Voice Callback

Table 10 shows recommendations for medium contact centers using Framework Management and Configuration Layers, T-Server, Historical Reporting, Enterprise Routing, Outbound Contact, and/or Voice Callback on Sun Solaris, Microsoft Windows, IBM/AIX, and/or HP-UX operating systems.

Table 10: Framework, Reporting, Routing, Outbound, and VCB

Sun Solaris Operating System	
OS	Solaris 2.6-10, 32- or 64-bit
Processor Type, Quantity, Speed	4 CPU 1.5 GHz UltraSPARC IIIi. Possible choices for CPU configuration could be: • Four 1xCPU servers with 2 GB RAM (e.g., Sun Netra 240) Or: • Two 2xCPU servers with 4 GB RAM (e.g. Sun Netra 240) Or: • One 4xCPU server with 8 GB RAM (e.g., Sun Netra 440)
Memory Size	8 GB total
Hard Disk Space	146 GB per server
Ports	Networking Ports - 2 Ethernet 1000/100/10 BASE-T ports (one for Framework, one for T-Server)
Microsoft Windows Operating System	
OS	Microsoft Windows 2000 / 2003
Processor Type, Quantity, Speed	2 CPU 2.6 GHz dual core processor or higher. Possible choices for CPU configuration could be: • Two 1xCPU servers with 2 GB RAM Or: • One 2xCPU server with 4 GB RAM Genesys recommends using a server class machine for all applications.
Memory Size	4 GB Total
Hard Disk Space	80 GB per server
Ports	Networking Ports - 2 Ethernet 1000/100/10 BASE-T ports (one for Framework, one for T-Server)

Table 10: Framework, Reporting, Routing, Outbound, and VCB (Continued)

IBM AIX Operating System	
OS	IBM AIX
Processor Type, Quantity, Speed	4xCPU 1.0 GHz POWER 4processors or higher. Possible choices for CPU configuration amongst any combination of P-series 630 servers could be: • Four single CPU boxes (e.g., 7028-6C4-100C) Or: • Two 2xCPU (e.g., 7028-6C4-200C) Or: • One 4xCPU (e.g., 7028-6C4-400c)
Memory Size	4 GB RAM
Hard Drive	76 GB HDD
Networking Ports	Two TP Ethernet 1000/100/10 BASE-T cards that work in full duplex mode (one for Framework, one for T-Server)
Linux Operating System	
OS	Red Hat Linux Enterprise Server 3-5
Processor Type, Quantity, Speed	4 CPU 2.4 GHz Intel Xeon Processor or higher. Possible choices for CPU configuration could be: • Two 2xCPU 2 GB RAM Or: • One 4xCPU 4 GB RAM server
Memory Size	4 GB Total RAM
Hard Drive	76 GB HDD
Networking Ports	Two TP Ethernet 1000/100/10 BASE-T cards that work in full duplex mode (one for Framework, one for T-Server)
HP-UX Operating System	
OS	HP-UX 11i
Processor Type, Quantity, Speed	P7410 Server on 4xCPU PA8700+processor
Memory Size	4 GB RAM

Table 10: Framework, Reporting, Routing, Outbound, and VCB (Continued)

Hard Drive	76 GB HDD
Networking Ports	Two TP Ethernet 1000/100/10 BASE-T cards that work in full duplex mode (one for Framework, one for T-Server)

Notes:

- Stat Server may need to be located on a separate box depending on interaction volume. It is recommended that a separate Stat Server should be installed if Genesys Routing is deployed.
- The Database Server may be installed on the same box as the database; however accurate database sizing information is essential. The number of HDDs and/or size of HDDs may need to be increased depending on data storage requirements and workflow.
- You should have the DVD-ROM for Multi-Media (T-Server's) installation.
- **Important:** The Database Administrator needs to participate in all database sizing decisions.

Call Progress Detection Server

Table 11 shows recommendations for medium contact centers using Call Progress Detection Server on a Microsoft Windows platform.

Table 11: Call Progress Detection Server

Microsoft Windows Platform	
OS	Microsoft Windows 2000/2003
Processor Type, Quantity, Speed	One Intel Core 2 CPU, min. 2.6 GHz, min 1 GB RAM each
Memory Size	1 GB RAM
Hard Drive	40 GB HDD
CPD Board	See the Supported Dialogic Boards Table in the *Genesys Supported Media Interfaces Reference Manual.*
Networking Ports	10/100 TX PCI UTP Microsoft Windows

Note: Hardware configuration for systems where Call Progress Detection Server is used with Intel HMP software should comply with the requirements specified in Intel HMP documentation.

Multi-Channel Routing and Multimedia

Table 12 on page 33 shows the minimum requirements for medium contact centers using Multi-Channel Routing (MCR) and Multimedia on a Microsoft Windows platform.

Notes:

- The name of Multi-Channel Routing is changed to Multimedia beginning with release 7.2.
- For purposes of MCR and Multimedia only, a medium contact center is defined as one with no more than 300 agents doing simultaneous chat processing and no more than 500 agents for e-mail processing, with an e-mail volume of no more than five interactions per second.

Component Distribution

Genesys recommends that you distribute MCR/Multimedia and associated components among several host machines, as follows:

- **Server**: Multimedia (previously MCR) core components: Interaction Server (with separate installation of DB Server), Classification Server, Universal Contact Server (UCS), E-mail Server Java*, Chat Server*
- **Server**: Multimedia web components
- **Database**: Universal Contact Server database and Interaction Server cache
- **Server**: Kana Response Live Server (part of Genesys web collaboration, requires a separate application container from Web API Server)
- **Desktop**: Interaction Routing Designer (IRD), Knowledge Manager, UCS Manager

* Based on the load and nature of contact center media (email, chat or blended) you might need to deploy these components on separate machines.

Note: The corporate mail server should also be deployed on a separate computer.

Sample Architecture

The sample architecture in Figure 1 shows all the components required for a Multimedia solution. Some key points about these requirements that you should be aware of:

- The number of servers and distribution of components complies with the recommendations in "Component Distribution".
- Adjust the number of servers and component distribution as necessary, according to the sizing information in the tables later in this section.
- If necessary, substitute any of the Microsoft Windows Server 2003 hosts in this architecture for a UNIX host (AIX, Solaris, Linux), thus allowing for a mixed approach.

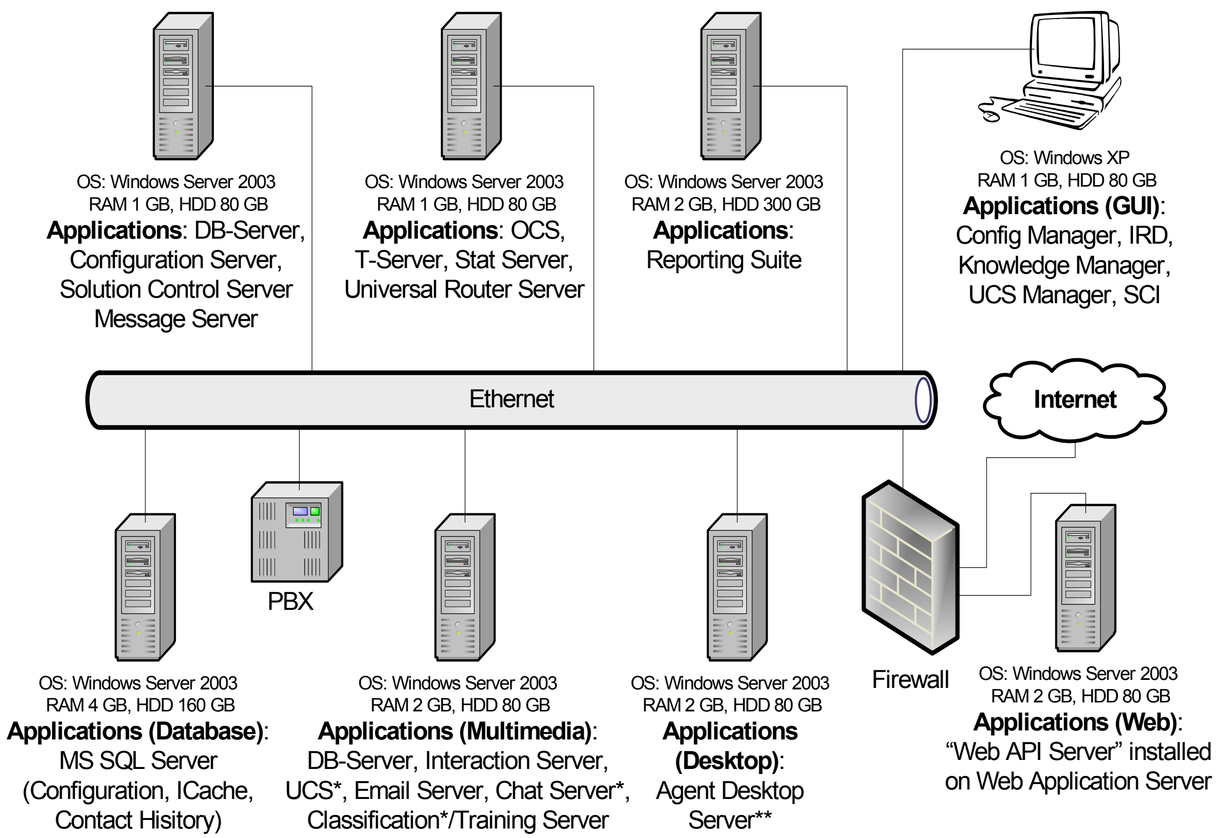

Figure 1: Sample Multimedia Architecture

* Depending on your solution, some of these servers might need to be installed on a separate machine. For example, in an E-mail solution, both Classification Server and UCS must be installed on separate machines.

** A computer with the listed specifications can support approximately 150-200 agents. For more agents, you must distribute the load over additional machines. For further explanation, see "Genesys Agent Desktop" on page 36.

Minimum Hardware Sizes

Table 12 contains separate requirements for the MCR/Multimedia components themselves and for the associated databases. You should consider these numbers as minimum recommendations only. Actual size depends on the type, number, size, and duration of the interactions your solution will be expected to handle. To help determine your size requirements, see the guidelines for memory and processing loads in "Interaction Processing Loads" on page 34.

Table 12: Multi-Channel Routing and Multimedia

Microsoft Windows Platform	
MCR/Multimedia Component	
OS	Microsoft Windows 2003 Server
Processor Type, Quantity, Speed	6xCPU 2.33 GHz processors or better. A possible CPU configuration could be: Distribution between three 2xCPU boxes, each with 1.5 GB of memory.
Memory Size	4 GB RAM
Hard Drive	160 GB available hard disk space **Note**: Hard disk requirements depend not only on the projected size of your database but also on the accumulated size of your log files (which you can regulate via options).
Database	
OS	Microsoft Windows Server 2003
Processor Type, Quantity, Speed	2xCPU 2.33 GHz processors or better
Memory Size	4 GB RAM
Hard Drive	Depends on contact center policy regarding the preservation of log files.

Note: Hard disk space requirements for the MCR/Multimedia databases vary greatly depending on how long the contact center saves its contact and interaction history data. It is therefore difficult to provide simple recommendations about hard disk space for these databases. However, the *Multimedia 7.6 Deployment Guide* does provide suggested formulas that you can use to calculate database sizing estimates. See the "Deployment Planning" chapter of that guide for details.

Interaction Processing Loads

See the following tables for information about:

- "Processing Loads for E-mail Interactions"
- "Processing Loads for Chat Interactions"

The numbers in these tables are derived from hardware configurations tested in Genesys laboratories—different configurations yield different results, so consider these values as guidelines to help plan the basic layout of your deployment. For example, in a blended solution (E-mail and Chat), depending on the capability of the various systems on your network and the volume of interactions you anticipate, you might consider setting up E-mail and Chat functionality on separate machines.

Processing Loads for E-mail Interactions

Table 13 shows both CPU and memory consumption in a contact center servicing 340 agents logged into a single instance of Genesys Agent Desktop. Components were deployed on computers with at least two 2.33 GHz Xeon-powered processors.

Table 13: Processing Loads for E-mail Interactions

Multimedia Components	Average CPU usage	Maximum CPU usage	Memory footprint
E-mail Server Java	4%	8%	less than 100 MB
Interaction Server	18%	57%	less than 24 MB
Universal Contact Server	22%	74%	less than 100 MB

Table 13: Processing Loads for E-mail Interactions (Continued)

Classification Server	22%	74%	less than 100 MB
Databases	**Average CPU usage**	**Maximum CPU usage**	**Memory footprint**
Interaction Server Database (ICache) on MS SQL Server	8%	11%	1.0 GB
Universal Contact Server Database on MS SQL Server	15%	50%	980 MB

Processing Loads for Chat Interactions

Table 14 shows both CPU and memory consumption in a contact center with 300 logged in agents, handling simultaneous chat sessions.

Table 14: Processing Loads for Chat Interactions

Multimedia Components	Average CPU usage	Memory footprint
Chat Server	44%	Total less than 50 MB
Web API Server	70% Apache usage: 6%	400 MB (Maximum 640 MB) Apache usage: 47 MB
Interaction Server (and DB Server)	Total 2%	Total less than 100 MB
Universal Contact Server	1%	65 MB
Databases	**Average CPU usage**	**Memory footprint**
Interaction Server Database (ICache) on MS SQL Server	1%	1.7 GB
Universal Contact Server Database on MS SQL Server	2%	

Sample Log File Sizes in a Chat Solution

The values in Table 15 give you a sample of how much hard disk space you will need for the log files that the various Multimedia components generate when running a Chat solution at a medium output level for 30 days. Use these numbers as guidelines only—actual size requirements vary, depending on the number and length of the chat sessions that your solution serves.

Note: Lab testing of Genesys components used a total of 300 simultaneous chat sessions, each with a duration of 5 minutes, creating approximately two and half million interactions over a one-month period.

Table 15: Log File Sizes in a Chat Solution

Multimedia Component	Average consumption per chat session	Total monthly HDD space consumption
Chat Server	30 K [a]	About 100 GB
Interaction Server	100 - 150 K [b]	260 - 390 GB
Universal Contact Server	25 K	100 GB
Classification Server	5K	100 GB
Training Server	Negligible consumption [c]	
Knowledge Manager	This GUI application has no impact on log size.	

a. Based on a 5-minute session containing 10 messages of 1.1 K each.

b. Very much dependent on the complexity of the interaction workflow (ie. strategies)

c. Since the server is typically run only once a week or month.

Genesys Agent Desktop

Table 16 shows recommendations for medium contact centers using Genesys Agent Desktop Server on a Microsoft Windows platform.

- Genesys Agent Desktop Server is installed in either standalone or in load balancing mode (several Tomcats and Apache in front), as follows:

  a. Standalone mode: Genesys Agent Desktop Server can support a maximum of 400 agents with Apache 2.2.4 with one Tomcat 5.5 (Catalina) with a maximum of 1.8 EMS or 8 CPS. See Table 16.

  Or:

  b. Load balancing mode: Use network load balancing hardware or software (for example: Microsoft Network Load Balancing for W2K Advanced Server) for configuration of more than 400 agents. Desirable configuration is one network node (2x2.4 MHz 2 GB memory box) with one Web Server and one Genesys Agent Desktop Server per each 400 agents.

c. Java Server Pages should already be compiled using Java Development Kit 1.4.2_xx or 1.5.0_yy from Sun.

Table 16: Genesys Agent Desktop

Microsoft Windows Operating System	
OS	Microsoft Windows 2000 Advanced Server SP3/Microsoft Windows 2003 Advanced Server
Processor Type, Quantity, Speed	2xPentium Xeon 2.4 GHz
Memory Size	2 GB SDRAM
Hard Drive	80 GB HDD
Ports	Networking Ports-TP Ethernet 100 MB full duplex mode
Web Server	Apache 2.2.4 Web Server

Genesys Desktop .NET Server

Genesys Desktop .NET Server is no longer available. From release 7.2, all .NET Server capabilities have been delivered in Genesys Integration Server (GIS). See Chapter 13, “Genesys Integration Server,” page 251 for details.

Genesys Agent Scripting

Table 17 shows the recommendations for medium contact centers using Genesys Agent Scripting on a Microsoft Windows Platform:

Note: Genesys Agent Scripting is installed in standalone mode (one Tomcat and Apache or IIS web server.

Table 17: Genesys Agent Scripting

Microsoft Windows Operating System	
OS	Microsoft Windows 2000/2003/XP
Processor Type, Quantity, Speed	2.6 GHz Processor or higher
Memory Size	1 GB RAM
Hard Disk Space	40 GB HDD

Table 17: Genesys Agent Scripting (Continued)

Microsoft Windows Operating System	
Ports	Networking Ports-TP Ethernet 100 MB full duplex mode
Server	IIS 5.0, IIS 6.0, Apache/Tomcat 4.3.1, or Apache/Tomcat 5.0 Web Server running on Microsoft Windows 2000/2003 server

Note: The sizing information is for running scripts without integration with other products. If scripts are use or integrated with other products, then the sizing requirements for those products must be considered in addition to what is defined above.

Blue Pumpkin Integration

Table 18 shows the Blue Pumpkin Integration (7.1) hardware requirements for medium contact centers (500-1000 agents).

Table 18: Blue Pumpkin Integration

WFM Data Aggregator	
OS	Microsoft Windows 2000, Microsoft Windows Server 2003
Processor Type, Quantity, Speed	Dual Intel Xeon 1 GHz with 1 GB L2 cache
Memory Size	1 GB RAM
Hard Disc Space	40 GB (Contact Genesys Technical Support website at `http://genesyslab.com/support` for specific requirements since 40 GB is just a baseline.)
Ports	2 TP Ethernet 1000/100 BASE-T cards in full duplex mode

Client Workstation Requirements

Computers running the Statistics Configuration Utility, at minimum, should meet these specifications:

- Pentium III, 1 GHz or greater, with 1 GB RAM

Genesys Call Director Route

Table 19 shows the Genesys Call Director Route hardware requirements for medium contact centers:

Table 19: Genesys Call Director Route

Microsoft Windows Operating System	
OS	Microsoft Windows 2000/2003 Server
Processor Type, Quantity, Speed	1.8 GHz Processor or higher
Memory Size	1 GB RAM
Hard Disk Space	40 GB HDD
Ports	Networking Ports—TP Ethernet 100 MB full duplex mode
Server	Tomcat 4.3.1 or Tomcat 5.0 Web Server

User Interface Hardware

User interfaces such as Configuration Management Environment (CME), Solution Control Interface (SCI), and CCPulse+, should use Microsoft Windows workstations with one 2.6 GHz CPU or better, and at least 1 GB of memory.

For Genesys Agent Desktop and Genesys Supervisor Desktop, a Microsoft Windows workstation with one 2.6GHz or better Intel Core 2 CPU, and at least 1GB of memory are recommended. If non-Genesys applications are being run on the same machine, 2GB of memory is recommended.

Genesys Agent and Genesys Supervisor Desktops also support Mac OS X workstations with one 1.5 GHz CPU or better, and at least 1 GB of memory.

Chapter

3 Hardware Architecture Diagrams

This chapter presents hardware architecture diagrams prepared by Genesys. It includes generic sizing information and diagrams for small and medium contact centers running Microsoft Windows, Sun Solaris, IBM/AIX, and HP-UX operating systems.

The information in this chapter contains the following topics:

Note: Genesys did not size box for SQL Server engine.

Small Configurations

The diagrams in this section are for small configurations on Microsoft Microsoft Windows, Sun Solaris, and IBM/AIX platforms.

Microsoft Windows Platform

Figure 2 shows a diagram of a sample small configuration on a Microsoft Windows Platform.

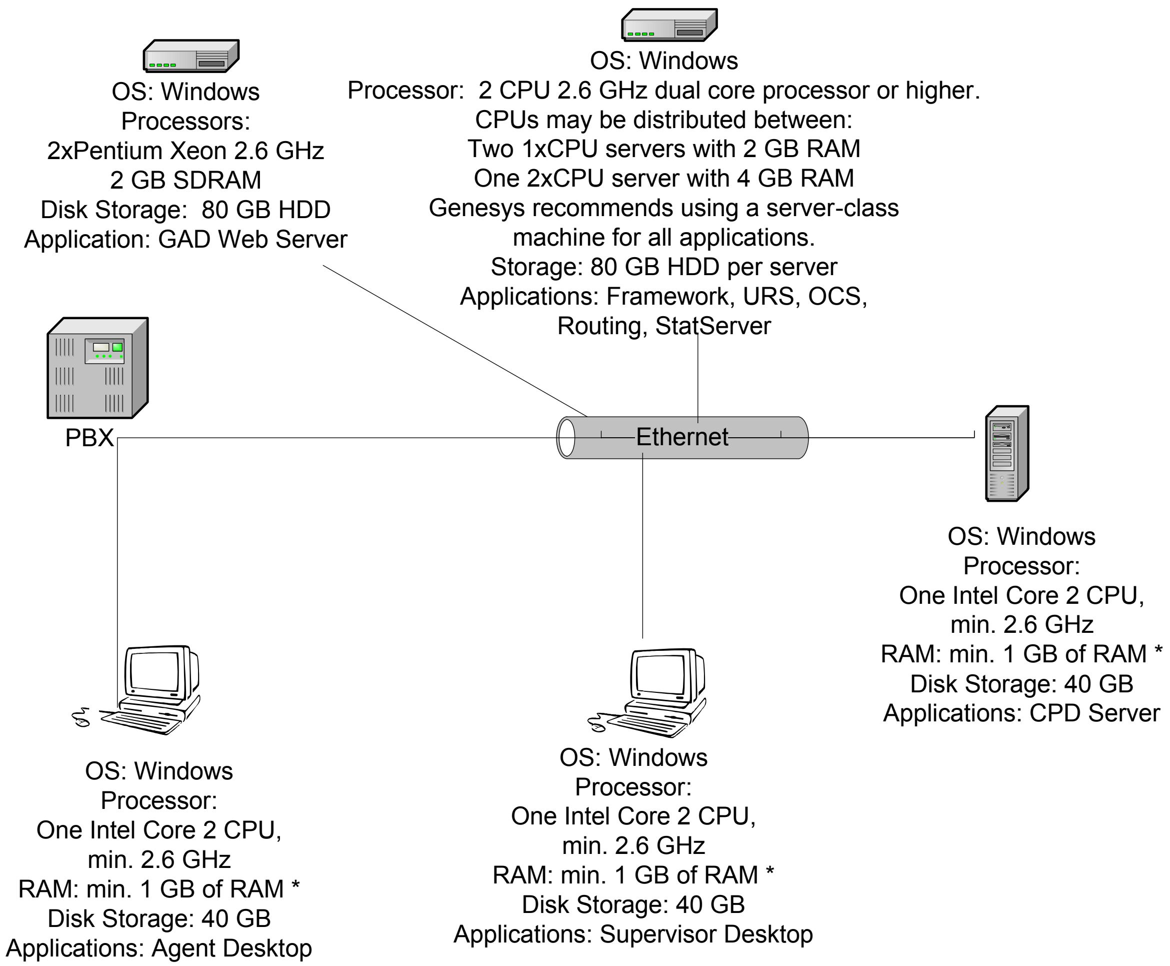

* Note: 2 GB RAM preferable if non-Genesys applications are being run.

Figure 2: Sample Small Configuration on a Microsoft Windows Platform

Sun Solaris Platform

Figure 3 shows a diagram of a sample small configuration on a Sun Solaris platform.

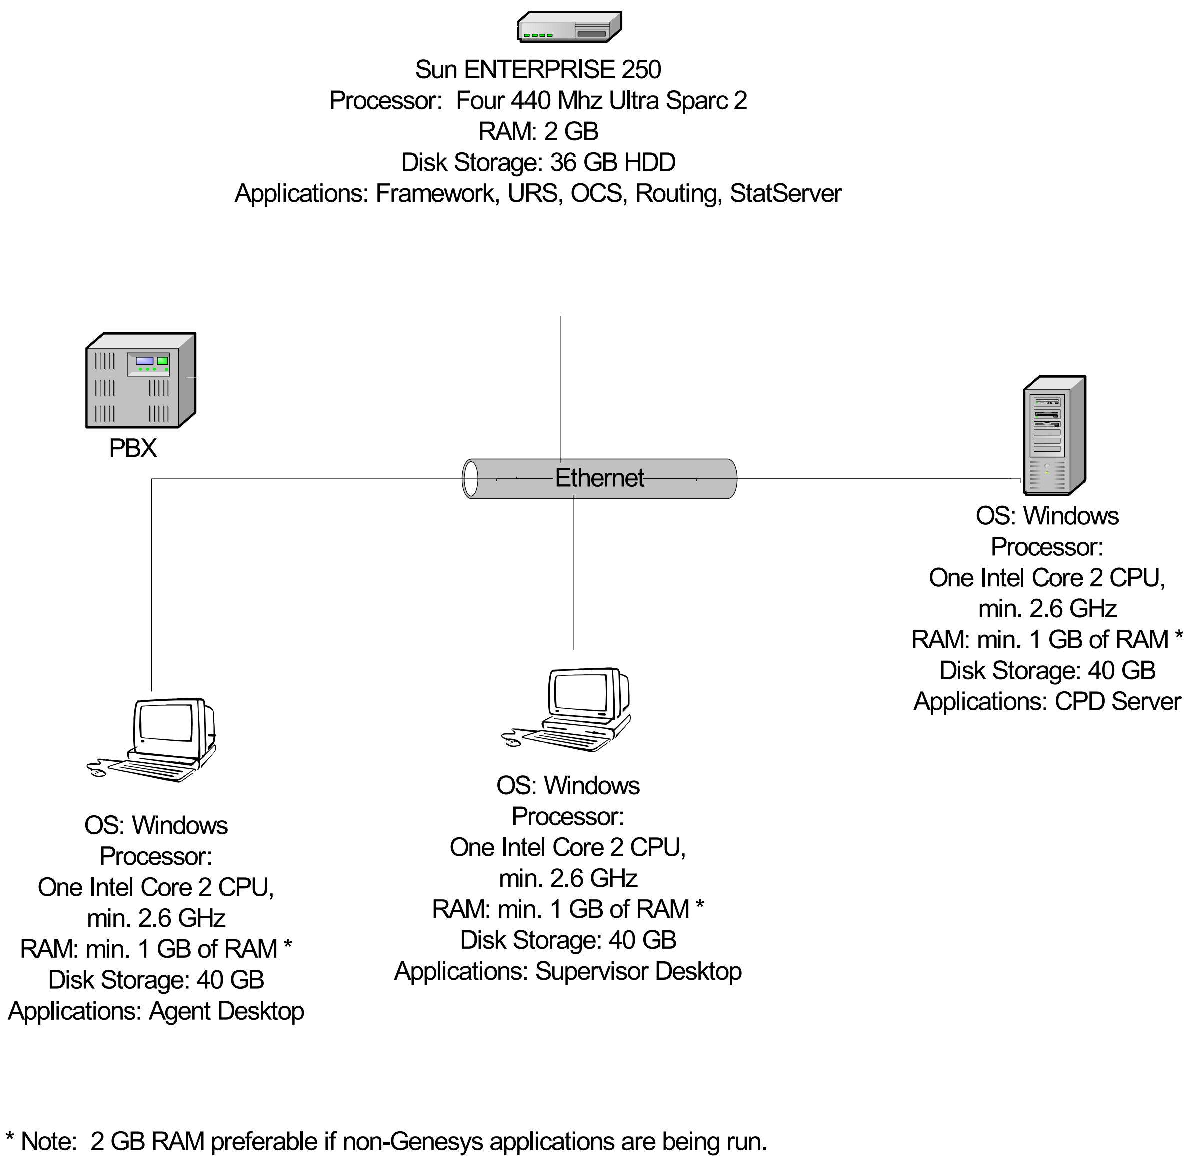

Figure 3: Sample Small Configuration on a Sun Solaris Platform

IBM/AIX Platform

Figure 4 shows a diagram of a sample small configuration on an IBM/AIX platform.

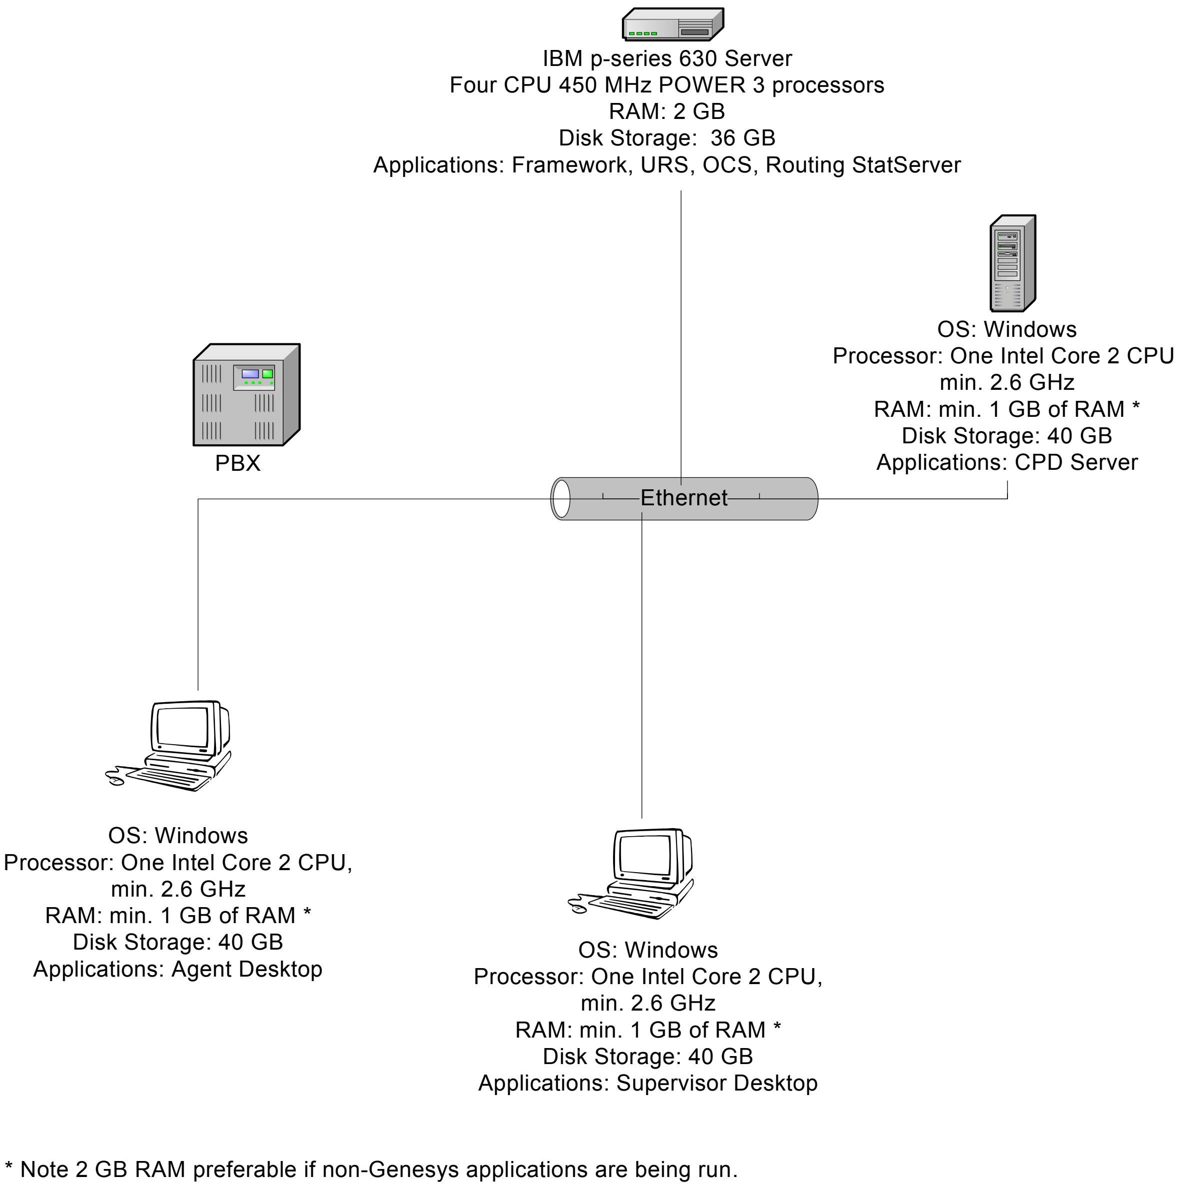

Figure 4: Sample Small Configuration on an IBM/AIX Platform

HP-UX Platform

Figure 5 shows a diagram of a sample small configuration on an HP-UX platform.

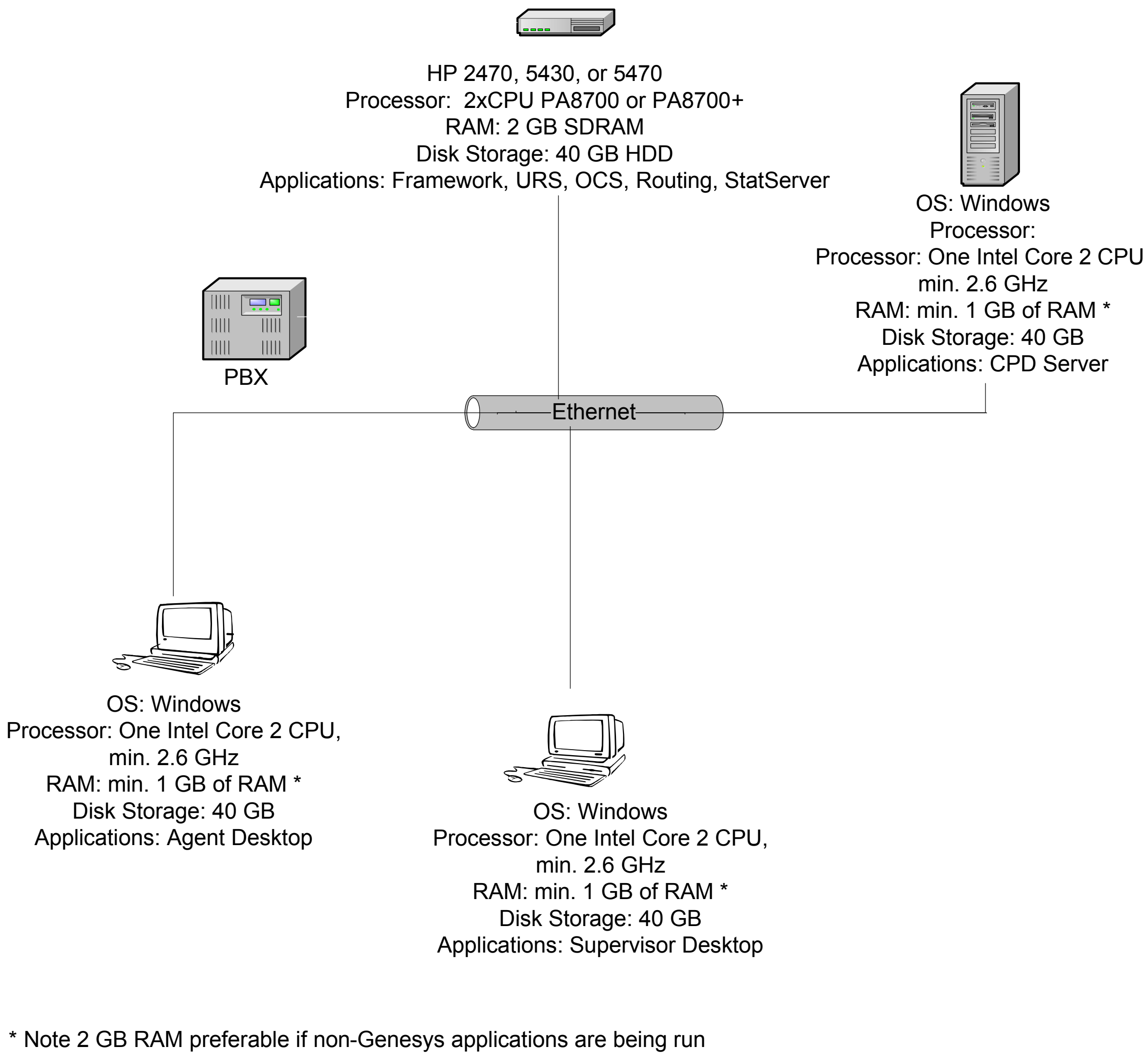

Figure 5: Sample Small Configuration on an HP-UX Platform

Medium Configurations

The diagrams in this section are for medium configurations on Microsoft Windows, Sun Solaris, and IBM/AIX platforms.

Microsoft Windows Platform

Figure 6 shows a diagram of a sample medium configuration on a Microsoft Windows platform.

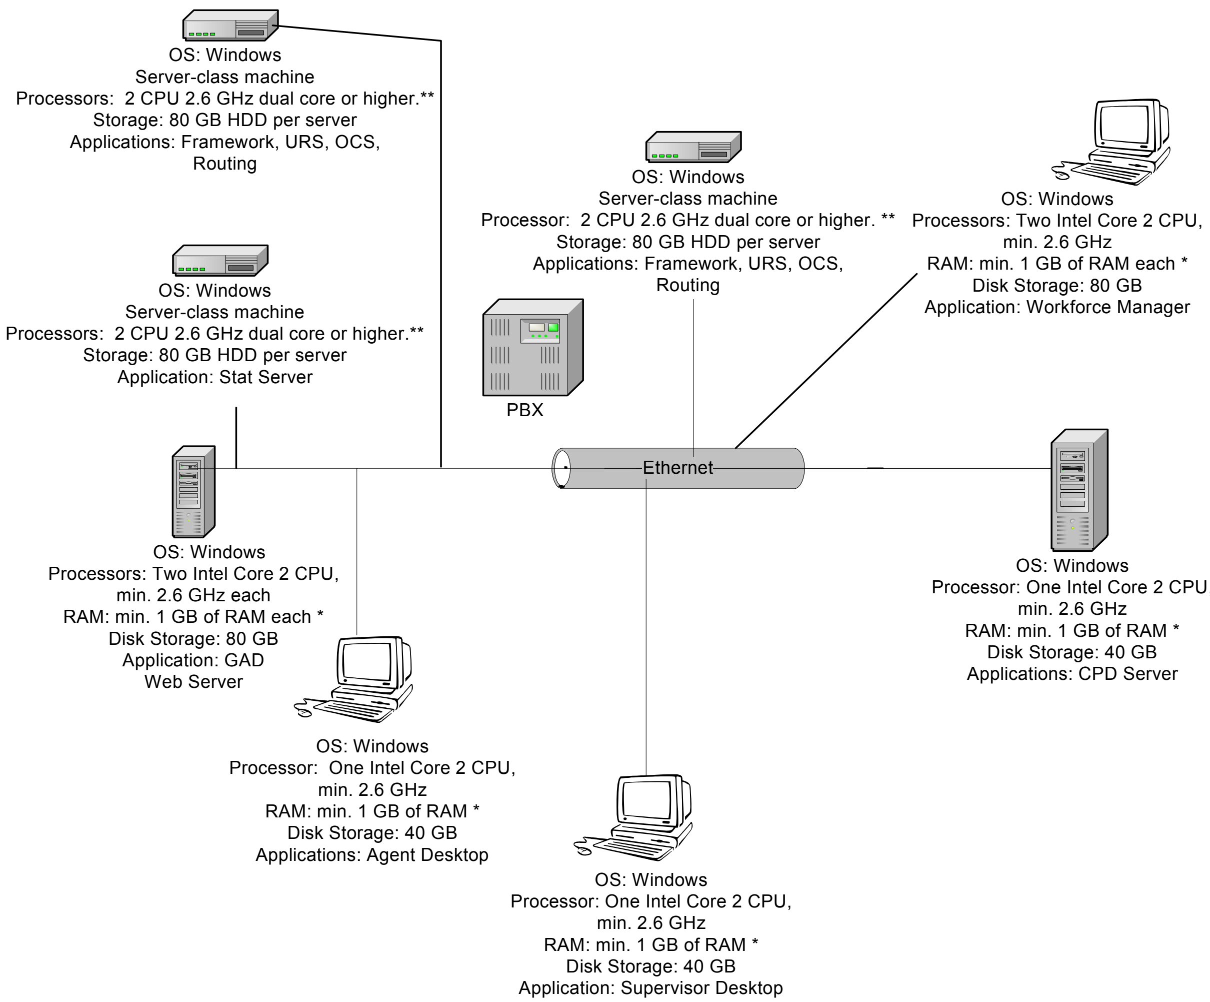

* 2 GB RAM preferable if non-Genesys applications are being run.
** Possible choices for CPU configuration could be: Two 1xCPU servers with 2 GB RAM, or, One 2xCPU server with 4 GB RAM. Genesys recommends using a server-class machine for all applications.

Figure 6: Sample Medium Configuration on a Microsoft Windows Platform

Sun Solaris Platform

Figure 7 shows a diagram of a sample medium configuration on a Sun Solaris platform.

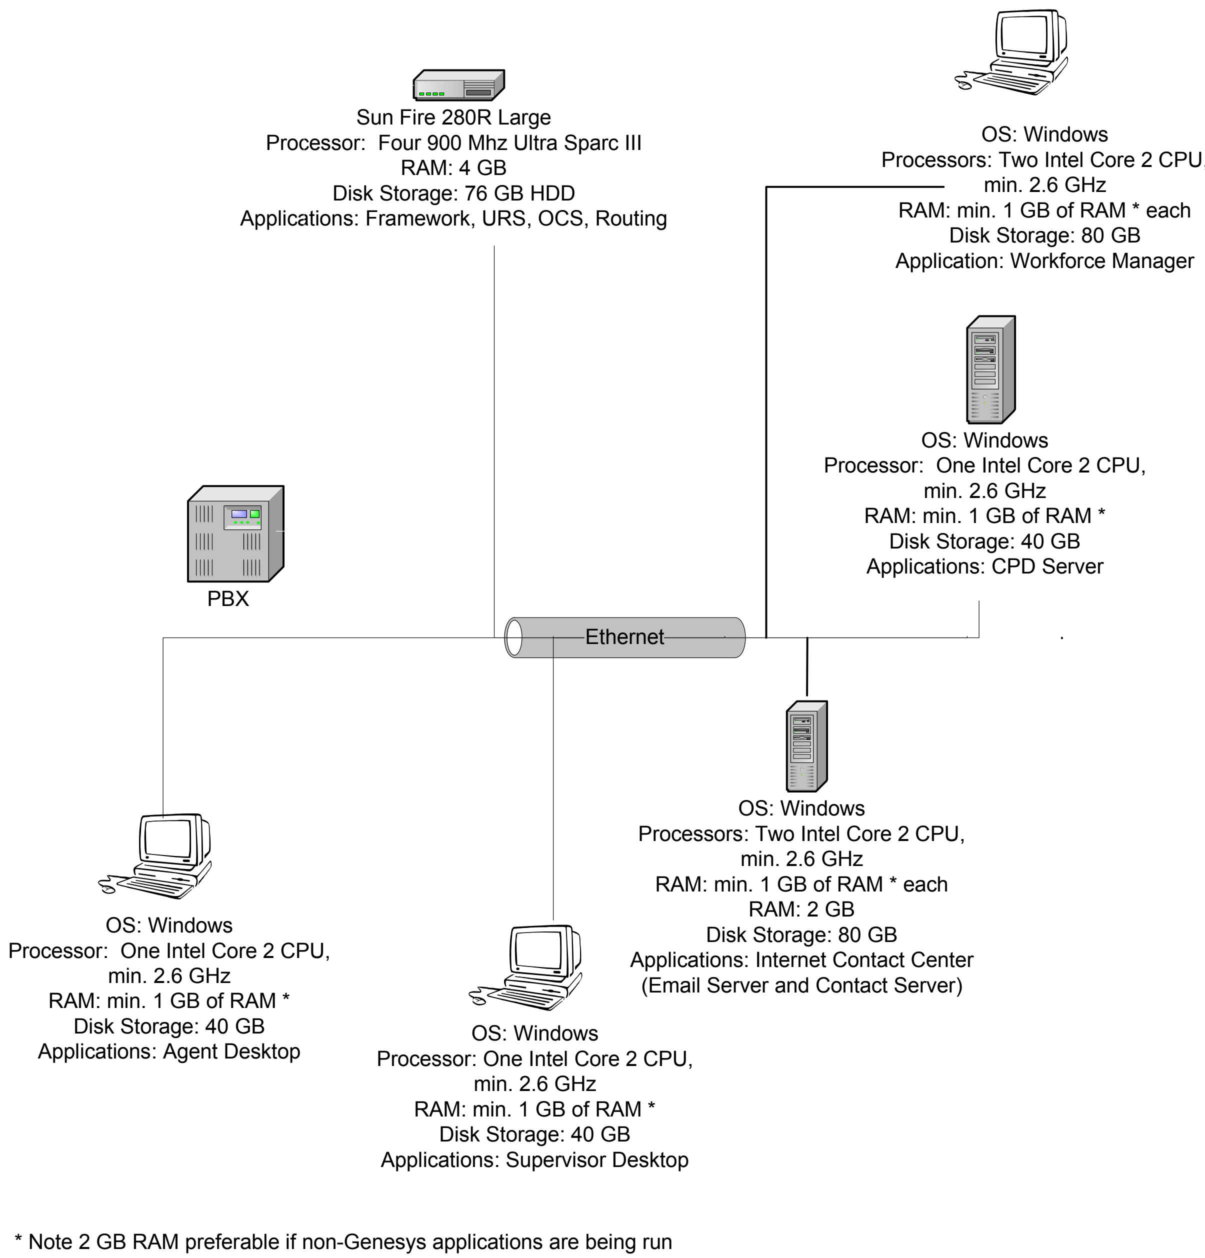

Figure 7: Sample Medium Configuration on a Sun Solaris Platform

Note: Box for DB Engine and corporate Mail Server was not counted.

IBM/AIX Platform

Figure 8 shows a diagram of a sample medium configuration on an IBM/AIX platform.

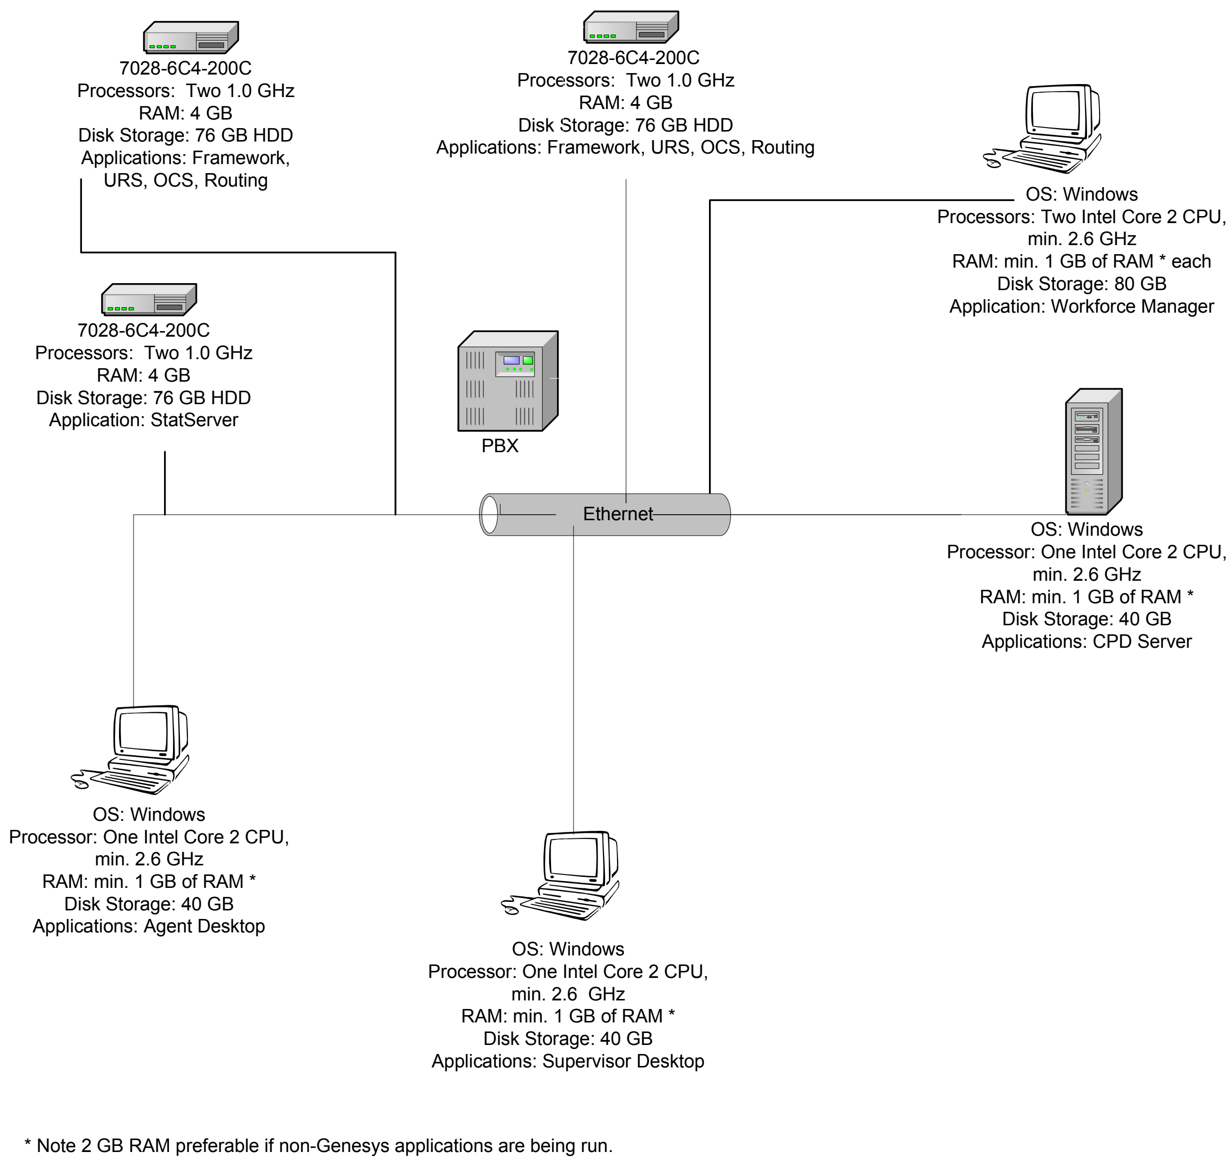

Figure 8: Sample Medium Configuration on an IBM/AIX Platform

Note: Box for DB Engine and corporate Mail Server was not counted.

HP-UX Platform

Figure 9 shows a diagram of a sample medium configuration on an HP-UX platform.

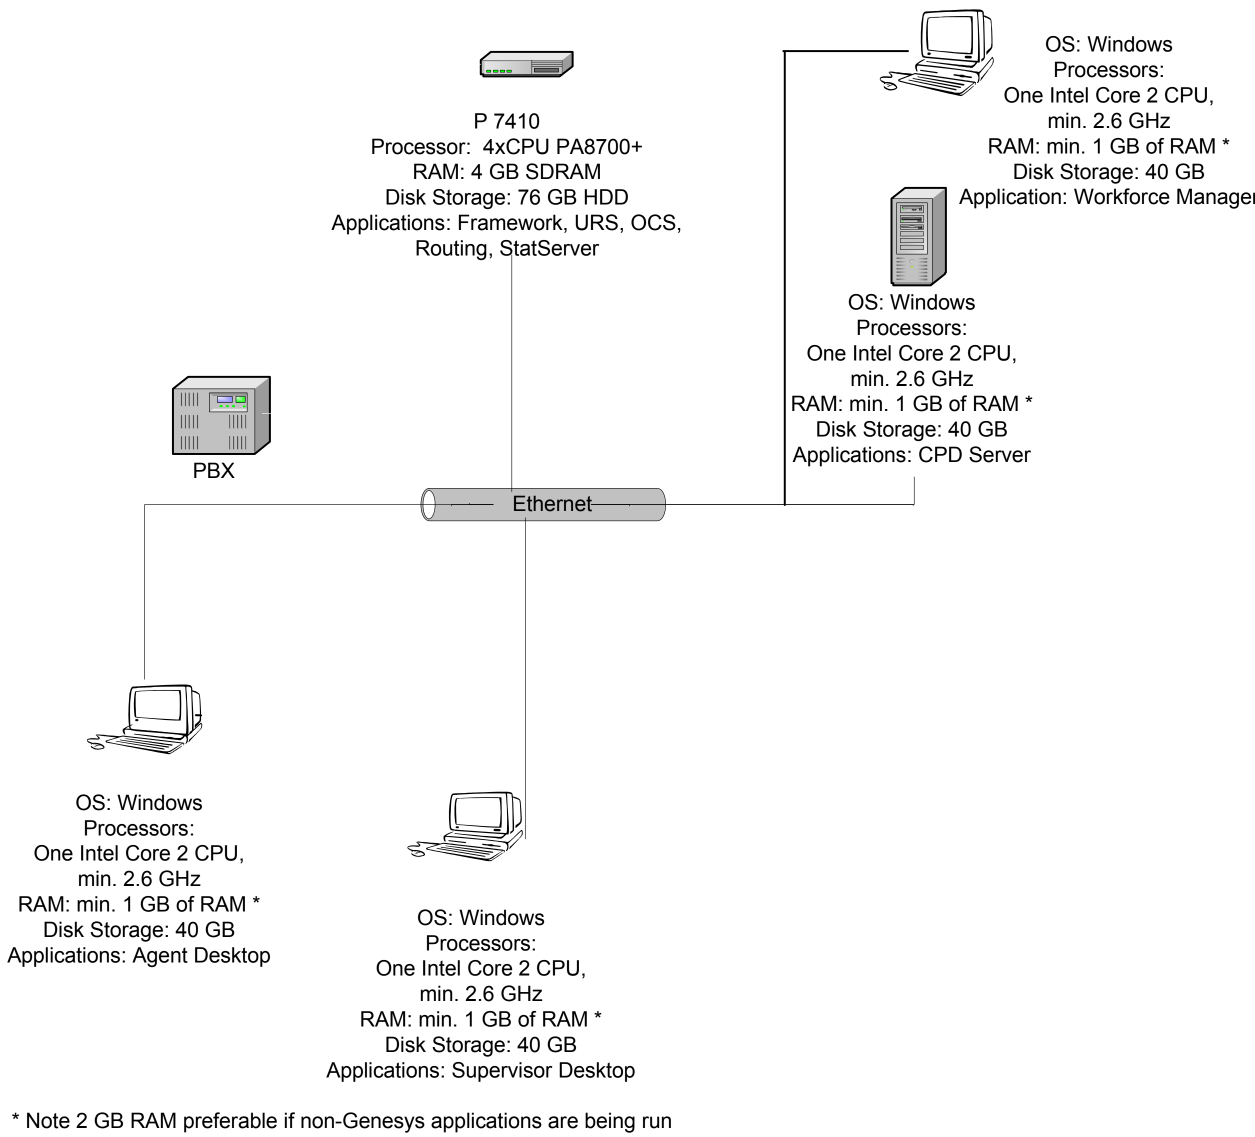

Figure 9: Sample Medium Configuration on an HP-UX Platform

Note: Box for DB Engine and corporate Mail Server was not counted.

Chapter

4 Genesys Interaction Concentrator

Interaction Concentrator collects and stores detailed data about the interactions and resources in customer interaction networks that use Genesys Framework (contact center, enterprise-wide, or multi-enterprise telephony and computer networks).

Interaction Concentrator consists of two components:

- Interaction Concentrator (ICON) server
- Interaction Database (IDB)

This chapter describes the hardware architecture for the ICON server and IDB components, providing examples of architectures for single-site and multi-site deployments. It also describes the factors that affect Interaction Concentrator performance, and lists sample performance measurements for reference platforms for both Windows and UNIX in single-site and multi-site deployments. This chapter also describes testing that was conducted in a large-scale environment to examine the performance and scalability of ICON 7.6.*x*.

Note: Before proceeding, review the product overview chapter in the *Interaction Concentrator 7.6 Deployment Guide,* to familiarize yourself with the product architecture.

This chapter contains the following sections:

Note: Because of their variability and complexity, large contact centers require special planning. If you want to plan an Interaction Concentrator deployment for a large contact center, contact Genesys Professional Services.

Hardware Architectures

This section provides examples of architectures for single-site and multi-site Interaction Concentrator deployments, and also provides hardware guidelines for the Interaction Concentrator (ICON) server and Interaction Database (IDB) components.

The Interaction Concentrator architecture is flexible enough to store reporting data for a contact center environment of practically any size. The hardware architecture that you select for your Interaction Concentrator deployment depends primarily on:

- The size of your contact center, in terms of the number of daily interactions and the number of agents. This determines the requirements for your Genesys Framework Configuration Layer and Media Layer components, which are the major sources of data for Interaction Concentrator.
- The type of data that you need to collect. In a multi-site environment with multiple IDBs, considerations include whether and how your downstream reporting application will merge inter-site interactions.

  You can balance the database-writing load by distributing some ICON roles among ICON instances. For more information, see the section about ICON roles in the product overview chapter in the *Interaction Concentrator 7.6 Deployment Guide.*

For information about the hardware requirements for the Framework Configuration Layer, Management Layer, and Media Layer components, see Chapter 4, "Small Contact Centers," on page 17 and Chapter 5, "Medium Contact Centers," on page 27.

For information about the hardware requirements for the ICON server, see "ICON Server" on page 56.

For information about sizing IDB in order to calculate the hardware requirements for the database, see "Interaction Database" on page 57.

Deployment Scenarios

This section provides examples of the following basic deployment types:

- A single-site deployment, with a single ICON instance writing to a single IDB instance
- A multi-site deployment, with a single ICON instance writing to a single IDB instance for the entire contact center (see page 53)
- A multi-site deployment, with multiple ICON instances writing to a single IDB instance (see page 54)
- A multi-site deployment, with multiple ICON instances writing to multiple IDB instances (see page 55)

For more information about the supported Interaction Concentrator deployments, see the section about deployment scenarios in the product overview chapter in the *Interaction Concentrator 7.6 Deployment Guide.*

Diagram Conventions

To maintain focus, the diagrams in this section omit the following components that are required for additional, optional Interaction Concentrator functionality:

- For reporting on virtual queue usage, a Universal Routing Server (URS), release 7.2.000.11 or later, connected to T-Server.
- For reporting on outbound activities, an Outbound Contact Server (OCS), release 7.2 or later, connected to an ICON instance.

Single-Site Deployment—One ICON, One IDB

Figure 10 depicts a sample hardware architecture for a small-size contact center in which a single ICON instance writes to a single IDB instance. The ICON instance is connected to a single T-Server.

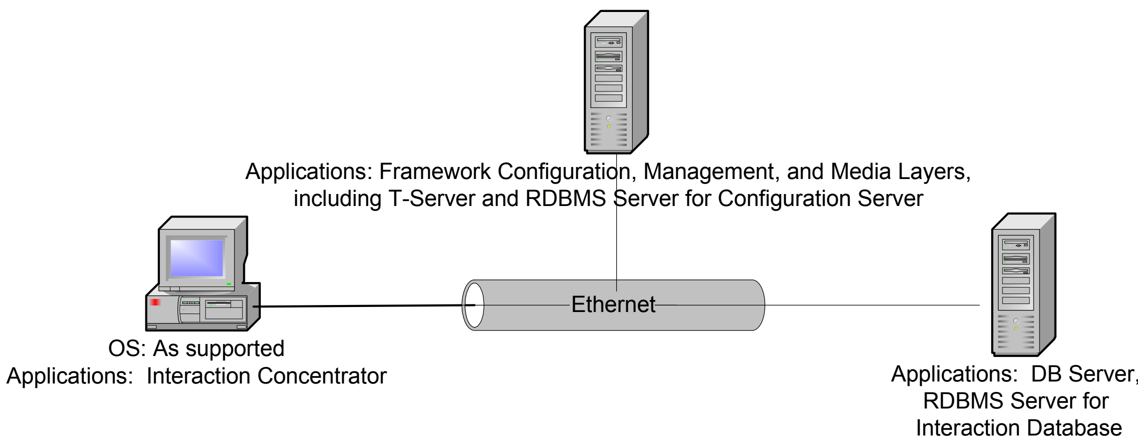

Figure 10: Single-Site Deployment with One ICON and One IDB

Multi-Site Deployment—One ICON, One IDB

Figure 11 depicts a sample hardware architecture for a medium-size contact center in which a single ICON instance, writing to a single IDB instance, serves all sites in the contact center. The ICON instance is connected to multiple T-Servers.

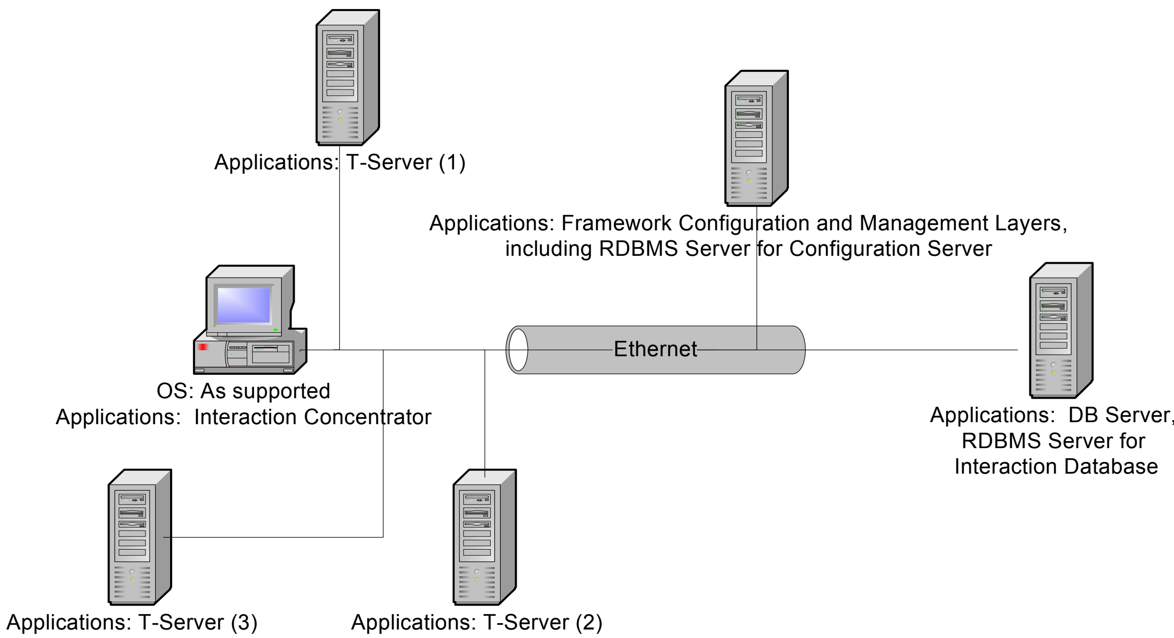

Figure 11: Multi-Site Deployment with One ICON and One IDB

Multi-Site Deployment—Multiple ICONs, One IDB

Figure 12 depicts a sample hardware architecture for a medium-size contact center in which there are multiple ICON instances, all of which write to the same IDB instance. Each ICON instance is connected to a single T-Server.

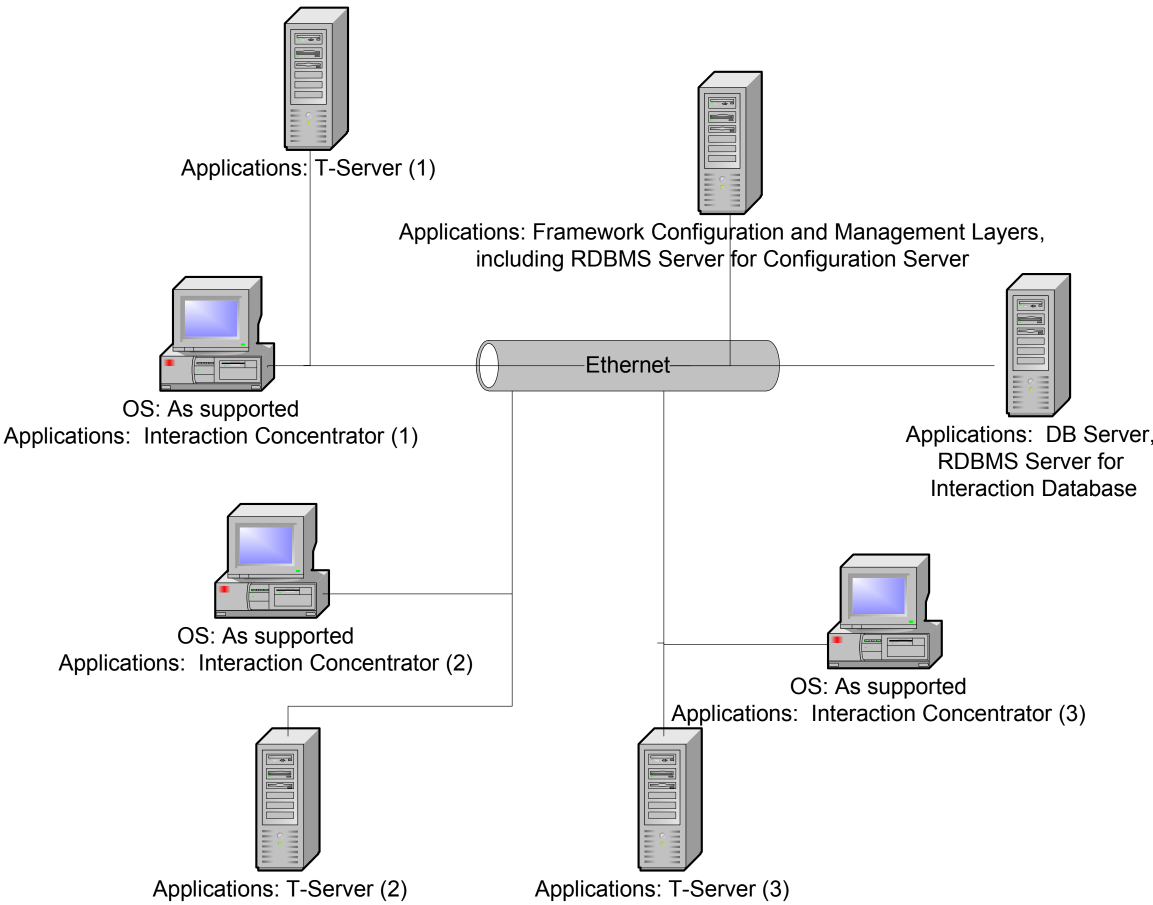

Figure 12: Multi-Site Deployment with Multiple ICONs and One IDB

Note: Genesys recommends using the multiple ICON, single IDB deployment only when there is a compelling reason to do so (for example, to use multiple network T-Servers in load-balancing mode). Genesys test results show that the performance of IDB is the limiting factor for overall Interaction Concentrator performance. Compared with the multiple ICON, multiple IDB deployment (see Figure 13), the multiple ICON, single IDB deployment requires more RDBMS resources and is more likely to encounter database conflicts that adversely impact IDB performance.

Multi-Site Deployment—Multiple ICONs, Multiple IDBs

Figure 13 depicts a sample hardware architecture for a medium-size or large-size contact center in which there are multiple ICON instances, each of which writes to its own IDB instance.

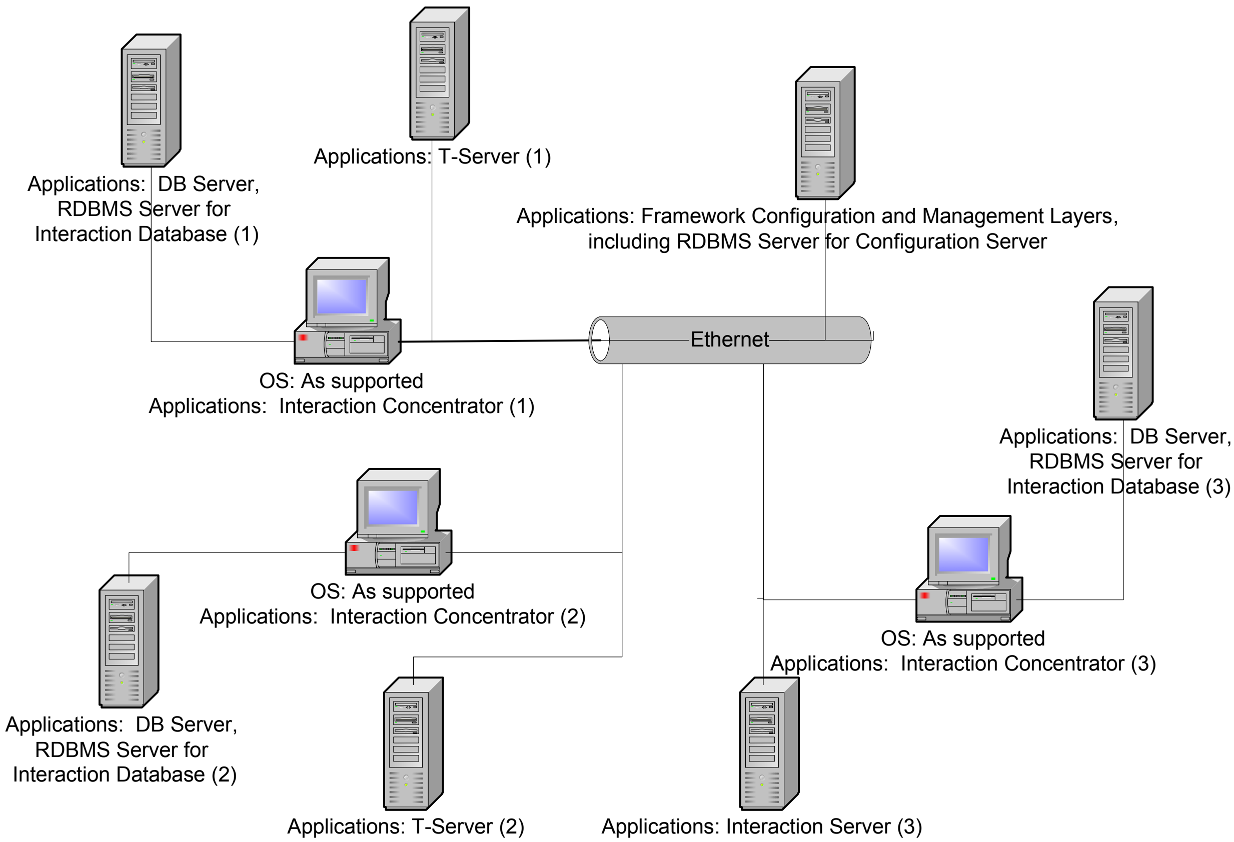

Figure 13: Multi-Site Deployment with Multiple ICONs and Multiple IDBs

Hardware Guidelines

This section describes hardware guidelines for the ICON server and IDB components. For information about the hardware requirements for the Genesys Framework components, see Chapter 4 on page 17 and Chapter 5 on page 27.

ICON Server

The memory requirements for the ICON server application, including the persistent queue, depend on the specific deployment and contact center characteristics.

The following factors affect the amount of memory and processing capacity that ICON requires:

- The number of interactions, which is a function of both the overall contact center size and the complexity of interaction flows
- The ICON configuration—for example, the roles that the ICON application has been configured to perform, or configuration settings that affect persistent queue and database-writing operations
- The requirements for merging interaction data

Because requirements are so deployment-specific, it is not possible to provide precise hardware specifications for the ICON server. However, Genesys has conducted performance tests for Interaction Concentrator on reference platforms for both Windows and UNIX. The test hardware, in combination with the reported usage results, can serve as a baseline that you can use when calculating your own requirements.

Interaction Database

The hardware requirements for IDB depend on the estimated size of your IDB instance or instances.

Genesys provides an interactive tool to help you estimate the required size of your IDB. This tool, the *Interaction Concentrator 7.6 Database Size Estimator,* is a Microsoft Office Excel spreadsheet that is available on the Genesys Technical Support website.

The spreadsheet uses relational database management system (RDBMS)–specific information and user input about general flow characteristics to provide an estimate of table, index, and total database size by day, month, and year, assuming that IDB is not periodically purged. Tooltips in the spreadsheet itself explain the information that you are required to enter.

Note: The *Interaction Concentrator 7.6 Database Size Estimator* provides a reasonably generous estimate, based on average projected activity. The estimator does not factor in growth associated with increased business.

Database Configuration and Optimization

Tablespace configuration and database tuning can significantly affect the performance and stability of Interaction Concentrator and your downstream reporting application.

Reducing I/O Contention

Any system configuration or database strategies to optimize database input and output (I/O) will significantly improve Interaction Concentrator performance. In particular, storing RDBMS logs, indexes, and table data on different disk drives reduces I/O contention. Genesys strongly recommends that you equip your RDBMS host with a fiber array or with a disk subsystem that contains multiple SCSI disk drives.

Recommendations for High Call Volumes

For environments with high call volumes, Genesys strongly recommends a multi-spindle disk subsystem, preferably with an advanced controller with write-back cache. Genesys also recommends locating the database log and temporary tablespace on disks that are separate from the disks where tables and indexes are stored.

In addition, for high call volumes, Genesys recommends that you configure IDB with a large buffer cache (hundreds of megabytes, if not gigabytes).

Database Settings For information about the database settings that were used for the Genesys 7.5 performance tests, see "Database Settings" on page 84.

Interaction Concentrator Performance

The primary factors affecting Interaction Concentrator performance are the following:

- The size of IDB
- The type of RDBMS
- RDBMS settings
- For multi-site deployments writing to a centralized IDB, the frequency with which the merge procedure is run
- Tuning of IDB
- The speed of the network connections between components
- The amount of business data attached to interactions

The following sections discuss the significance of each of these factors in the context of the specific ICON release.

Interaction Concentrator 7.6.1 Performance

Genesys performance testing of ICON 7.6.1 focused on validating the following performance-based requirements of ICON 7.6.1 in a large-scale testing environment:

- Handling a large number of active interactions
- Filtering out unnecessary Multimedia interaction data
- Purging Multimedia interaction data

Two separate test environments were utilized—a voice-specific deployment and an open media (e-mail) deployment. The following sections describe the sample ICON 7.6.1 environments, call flows, tests, results, and recommendations.

Voice Test Environment

The voice test environment was organized as a single contact center under a single tenant that used a multi-tenant Configuration Server (see Figure 14). The contact center had four T-Servers that represented four separate sites. The routing targets consisted of 5,000 inbound agents per T-Server, for a total of 20,000 agents. During the tests, the logging level was set to `Standard` for all processes.

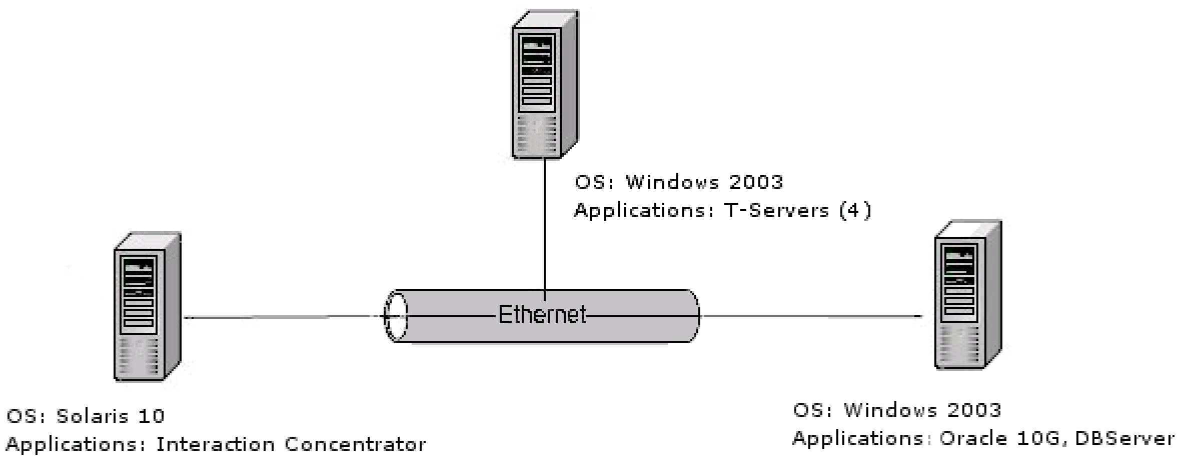

Figure 14: ICON 7.6.1 Voice Test Environment

Hardware and Software

Table 20 describes the hardware and software that hosted the ICON 7.6.1 application in the voice-specific test environment.

Table 20: Hardware and Software in the ICON 7.6.1 Voice Test Environment

Application	Processor	Memory	Application Software Version	
			On Windows 2003 Server[a]	On Solaris 10[b]
ICON	Quad dual core SPARC64 IV 2.15 GHz	32 GB RAM	n/a	See Table 21
RDBMS (Oracle 10)	E5410 dual quad core Intel Xeon 2.33 GHz	8 GB RAM	10.2.0.4	n/a
DB Server (ICON)	E5410 dual quad core Intel Xeon 2.33 GHz	8 GB RAM	7.6.000.08	n/a
Stat Server	X5355 dual quad core Intel Xeon 2.66 GHz	8 GB RAM	7.6.100.12	n/a
URS	X5355 dual quad core Intel Xeon 2.66 GHz	8 GB RAM	7.6.100.04	n/a

a. Windows Server 2003 Enterprise Edition SP2 32-bit

b. Solaris 10 (Sun OS 5.10) 64-bit 11/06

Table 21: ICON Releases Tested

Test	ICON Release
ICON 7.6 and 7.6.1 Comparison Test	7.6.000.16 and 7.6.100.09
Voice Purge Test	7.6.100.09

Call Flow and Configuration

Figure 15 shows the call flow and configuration that were used in the voice-specific tests.

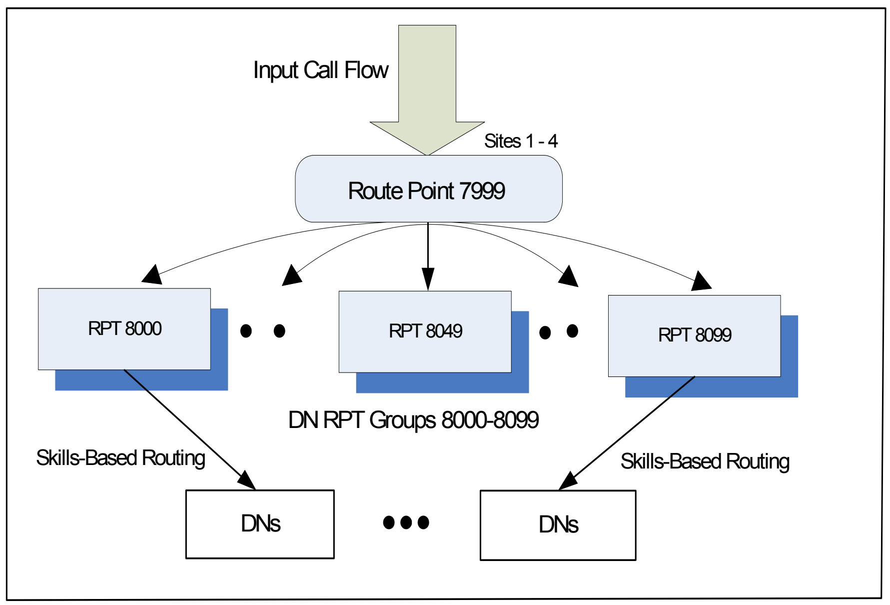

Figure 15: Voice-Specific Call Flow for ICON 7.6.1 Tests

The call flow is as follows:

- Inbound calls arrive at a route point, where they are routed in turn to a DN group of 100 route points.
- At route point two, Configuration Server options are assigned to routing-strategy variables for reference at various points in the routing strategy.
- One KB of data is attached to each call.
- A skills-based routing strategy routes the call to the most suitable agent (across all four sites), where the call is handled for 120 seconds.

Tests and Results

The following tests were run in the voice-specific environment:

- ICON 7.6 and 7.6.1 Comparison Test (Voice) (see page 61)

- Voice Purge Test (see page 61)

The details of these tests and the obtained results are discussed in the following sections.

ICON 7.6 and 7.6.1 Comparison Test (Voice)

Genesys monitored ICON 7.6 and ICON 7.6.1 over a period of two hours to compare the performance of each version. The tests were conducted under the following conditions:

- Call rate of 34 calls/second
- Duration of test: 2 hours
- Twenty thousand (20,000) agents configured
- ICON application running on Solaris 10

Test Results

Performance results are consistent between the 7.6 and 7.6.1 releases of ICON. Table 22 provides the observed results.

Table 22: ICON 7.6 and 7.6.1 Comparison Test (Voice) Results

ICON Release	ICON		Oracle[a]	
	Maximum CPU (%)[b]	Maximum RAM (MB)	Maximum CPU (%)[b]	Maximum RAM (MB)
7.6	32	756	102	2,421
7.6.1	32	768	106	2,423

a. Multi-threaded application

b. Peak usage of one core/thread (%)

Voice Purge Test

Genesys tested the new `gsysPurge76` stored procedure in the voice-specific environment over a period of three days while under a constant load. The test was conducted under the following conditions:

- Call rate of 17 calls/second
- Twenty thousand (20,000) agents configured
- Memory-management options enabled: `om-memory-optimization` set to `true`; `om-max-in-memory` set to `100` (default)
- Non-partitioned database

Test Results

The purge procedure was executed once every 24 hours for two days. It was able to clear approximately 204 million rows from IDB tables in 150 minutes. Although ICON did fall behind processing calls and attached-data queues while IDB tables were purged, it recovered quickly after the purge procedure completed. Table 23 displays the observed results.

Table 23: Voice Purge Test Results

Application	Maximum CPU (%)[a]	Maximum RAM (MB)
ICON	36	686
Oracle (RDBMS)[b]	175	2,440

a. Peak usage of one core/thread (%)

b. Multi-threaded application

Open Media Test Environment

The open media test environment was organized as a single contact center under a single tenant that used a multi-tenant Configuration Server. A single ICON application on a Windows platform, writing to an Oracle database on Microsoft Windows Server 2003, was configured to collect open media (e-mail) interactions from Interaction Server (see Figure 16). During the tests, the logging level was set to `Standard` for all processes.

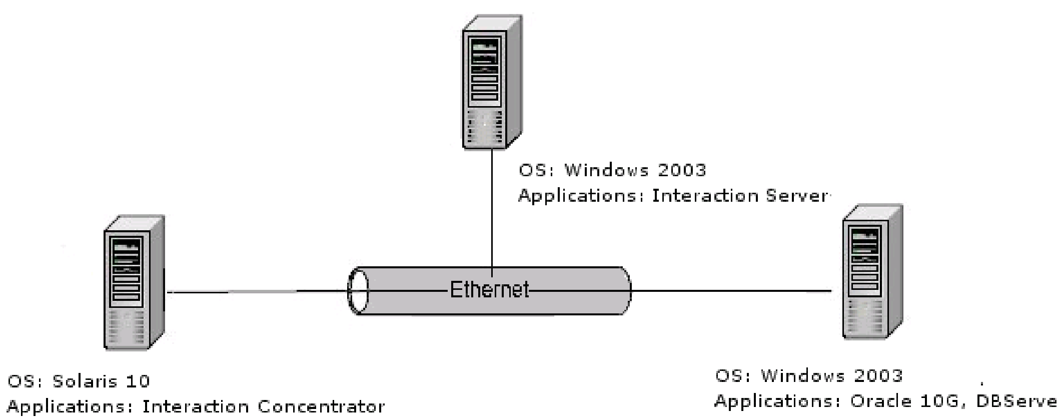

Figure 16: ICON 7.6.1 Open Media Test Environment

Hardware and Software

Table 24 describes the hardware that hosted the ICON application and other components in the 7.6.1 open media environment.

Table 24: Hardware and Software in ICON 7.6.1 Open Media Environment

Application	Processor	Memory	Application Software Version	
			On Windows 2003 Server[a]	On Solaris 10[b]
ICON	Quad dual core SPARC64 IV 2.15 GHz	32 GB RAM	n/a	See Table 25
Interaction Server	X5355 dual quad core Intel Xeon 2.66 GHz	8 GB RAM	7.6.100.24	n/a
Interaction Proxies	Dual UltraSPARC-IIIi 1.0 GHz	2 GB RAM	n/a	7.6.100.04
RDBMS (Oracle 10)	E5410 dual quad core Intel Xeon 2.33 GHz	8 GB RAM	10.2.0.4	n/a
DB Server (ICON)	E5410 dual quad core Intel Xeon 2.33 GHz	8 GB RAM	7.6.000.08	n/a
DB Server (Interaction Server)	E5410 dual quad core Intel Xeon 2.33 GHz	8 GB RAM	7.6.000.08	n/a
Stat Server	X5355 dual quad core Intel Xeon 2.66 GHz	8 GB RAM	7.6.100.12	n/a
URS	X5355 dual quad core Intel Xeon 2.66 GHz	8 GB RAM	7.6.100.04	n/a

a. Windows Server 2003 Enterprise Edition SP2 32-bit

b. Solaris 10 (Sun OS 5.10) 64-bit 11/06

Table 25: ICON Releases Tested

Test	ICON Release
Open Media Backlog Test	7.6.100.09 (with 400,000 e-mails backlogged) 7.6.100.05 (with 0, 1 million, and 10 million e-mails backlogged)
Data Filtering Test	7.6.100.05

Table 25: ICON Releases Tested (Continued)

Test	ICON Release
ICON 7.6 and 7.6.1 Comparison Test (Open Media)	7.6.000.16 and 7.6.100.09
Multimedia Purge Test1	7.6.100.05
Multimedia Purge Test2	7.6.100.05

Configuration and Interaction Flow

In the 7.6.1 open media model, the interaction flow is as follows:

- An interaction arrives at the entry queue and is routed to a target queue based on the business strategy.
- Five (5) KB of data, consisting of 100 key-value pairs of data, is attached.
- The interaction is directed by a skills-based routing strategy to one of the 2,000 agents who are logged in to two separate Interaction Server proxies (total of 4,000 agents).
- The agent handles the interaction for 5 minutes (300 seconds), which simulates the approximate amount of time that is required to respond to an e-mail interaction.
- The interaction is routed to an archive queue.

Figure 17 shows the call flow and configuration used in the open media tests.

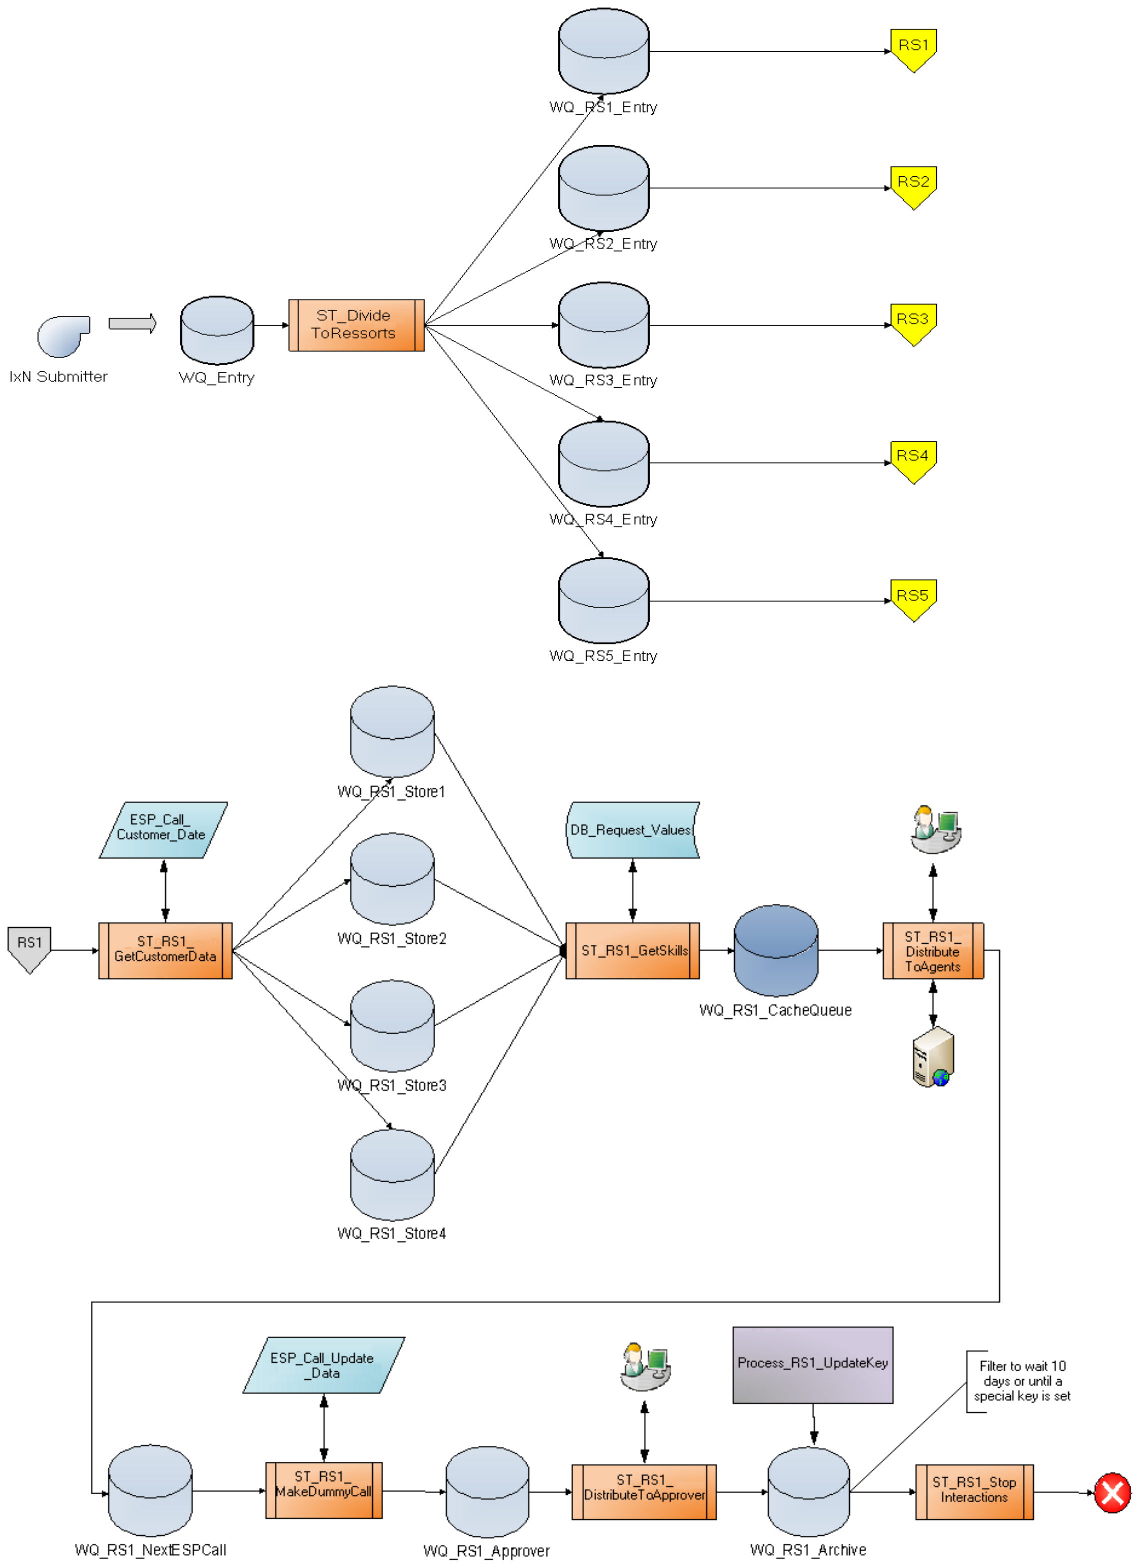

Figure 17: Open Media Call Flow for ICON 7.6.1 Tests

Tests and Results

The following tests were run in the open media environment:

- Open Media Backlog Test (see page 66)
- Data Filtering Test (see page 67)
- ICON 7.6 and 7.6.1 Comparison Test (Open Media) (see page 68)
- Multimedia Purge Test 1 (see page 69)
- Multimedia Purge Test 2 (see page 70)

The details of these tests and the obtained results are discussed in the following sections.

Open Media Backlog Test

In this test, Genesys varied the number of e-mail interactions that were backlogged to determine the effect on memory consumption. It was conducted under the following conditions:

- Four thousand (4,000 agents) logged in to queue
- Duration of test: 2 hours
- Memory-management options enabled: `om-memory-optimization` set to `true`; `om-max-in-memory` set to `100` (default); `om-memory-clean` set to `1` on the interaction queues that are receiving the interactions.
- Four-hundred thousand (400,000) processed interactions in IDB
- Number of e-mails backlogged: 0; 400,000; 1 million; and 10 million
- ICON application running on Solaris 10

Test Results

Table 26 displays the following results:

- Oracle CPU gradually increased with the backlog of e-mails until it stabilized after approximately 6 hours.
- ICON CPU remained constant at 48 percent, as the number of backlogged e-mails increased.

Table 26: Open Media Backlog Test Results

Application	ICON		Oracle[a]	
# E-mails Backlogged	Maximum CPU (%)[b]	Maximum RAM (GB)	Maximum CPU (%)[b]	Maximum RAM (MB)
0	48	538	113	2,302
400,000	48	559	185	2,421

Table 26: Open Media Backlog Test Results (Continued)

Application	ICON		Oracle[a]	
# E-mails Backlogged	Maximum CPU (%)[b]	Maximum RAM (GB)	Maximum CPU (%)[b]	Maximum RAM (MB)
1 million	48	563	182	2,300
10 million	48	550	152	2,307

a. Multi-threaded application

b. Peak usage of one core/thread (%)

Data-Filtering Test

In this test, Genesys monitored the performance of ICON as the data-filtering configuration option, `udata-history-terminated`, was set first to `0` (default) and then to `1`.

Note: By setting the `udata-history-terminated` configuration option to `1` on the ICON `Application` object, ICON will not store the call-termination value of UserData keys to IDB. For more information on this feature, see the *Interaction Concentrator 7.6 User's Guide.*

The test was conducted under the following conditions:

- Four thousand (4,000) agents logged into queue
- Duration of test: 2 hours
- No backlogged interactions
- Four-hundred thousand (400,000) processed interactions in IDB
- Memory management options enabled

Test Results

Table 27 displays the following results:

- There was no impact on ICON's CPU when filtering was enabled compared to when filtering was disabled.
- Fifty-five (55) percent fewer attached-data keys were recorded in the `G_USERDATA_HISTORY` table, when filtering was enabled.

Table 27: Data-Filtering Test Results

Application	ICON		Oracle[a]	
Filtering	Maximum CPU (%)[b]	Maximum RAM (GB)	Maximum CPU (%)[b]	Maximum RAM (MB)
OFF	48	538	113	2,302
ON	48	535	104	2,426

a. Multi-threaded application

b. Peak usage of one core/thread (%)

ICON 7.6 and 7.6.1 Comparison Test (Open Media)

Genesys monitored ICON 7.6. and 7.6.1 over a period of 2 hours to measure any changes in performance between releases. The test was conducted under the following conditions:

- Four thousand (4,000) agents configured
- Interaction rate of 14 e-mails/second
- Duration of test: 2 hours
- Two-hundred twenty thousand (220,000) previously processed interactions in IDB
- Memory-management options enabled

Test Results

Table 28 displays the following results:

- ICON 7.6 memory increased as expected throughout the test, reaching a maximum of 4.3 GB.
- Memory utilization in ICON 7.6.1 was much better than in ICON 7.6. CPU in ICON 7.6 reached 80 percent of the available CPU, or 67 percent more than the ICON 7.6.1 results.
- Oracle CPU utilization in ICON 7.6 was almost 300 percent higher than the ICON 7.6.1 results.

Table 28: ICON 7.6 and 7.6.1 Comparison Test (Open Media) Results

ICON Release	ICON		Oracle[a]	
	Maximum CPU (%)[b]	Maximum RAM (MB)	Maximum CPU (%)[b]	Maximum RAM (MB)
7.6	80	4,315	334	2,422
7.6.1	48	538	113	2,302

a. Multi-threaded application

b. Peak usage of one core/thread (%)

Multimedia Purge Test 1

Genesys tested the new `gsysPurge76` stored procedure in the open media environment over a period of three days while under a constant load. No sizing parameters were set. The test was conducted under the following conditions:

- Interaction rate of 14 e-mails/second
- Four thousand (4,000) agents configured
- Memory-management options enabled
- Test duration: 3 days under constant load
- Non-partitioned database
- ICON application running on Solaris 10

Test Results

Testing produced the following results:

- In two main executions, the purge procedure cleared approximately 361 million and 363 million rows from IDB tables in 293 minutes and 333 minutes, respectively.
- The execution of the purge procedure did not affect ICON CPU or memory consumption (see Table 29).
- Although ICON did fall behind processing calls and attached-data queues while IDB tables were purged, it recovered quickly after the purge procedure completed.

Table 29: Multimedia Purge Test 1 Results

Application	Maximum CPU (%)[a]	Maximum RAM (GB)
ICON	48	580
Oracle[b]	300	2480

a. Peak usage of one core/thread (%)

b. Multi-threaded application

Multimedia Purge Test 2

Genesys repeated the multimedia purge test—this time, setting an IDB sizing parameter (`rowspertransaction`) to specify the maximum size of one database transaction. This test was conducted under the following conditions:

- Interaction rate of 14 e-mails/second
- Four thousand (4,000) agents configured
- Memory-management options set to `true`
- `rowspertransaction` parameter set to `10 million`
- Duration of test: 3 days under constant load
- Non-partitioned database
- ICON application running on Solaris 10

Test Results

Testing produced the following results:

- In two purge procedure executions, 360.8 million and 360.91 million rows were cleared from IDB tables in 303 minutes and 323 minutes, respectively.
- Although ICON did fall behind processing calls and attached-data queues while IDB tables were purged, it recovered quickly after the purge procedure completed.

Table 30 provides the observed results.

Table 30: Multimedia Purge Test 2 Results

Application	Maximum CPU (%)[a]	Maximum RAM (GB)
ICON	48	575
Oracle[b]	265	2464

a. Peak usage of one core/thread (%)

b. Multi-threaded application

ICON 7.6.1 Performance Conclusions

Based on the results of performance testing that was conducted in a large-scale environment, Genesys has the following observations and conclusions:

- ICON 7.6.1 demonstrates a significant improvement over ICON 7.6.
- The new memory-management options are effective in controlling ICON memory consumption, in that ICON can continue to operate under load conditions while having a large number of interactions backlogged.
- ICON 7.6.1 demonstrates a significant (approximately 67 percent) improvement over ICON 7.6 in handling open media (e-mail) interactions under load conditions. There is also a significant (295 percent) reduction to the Oracle database load.
- ICON can operate successfully under load conditions while the new `gsysPurge76` procedure executes in a timely manner to clear IDB tables. Although ICON does fall behind with database insertion while IDB tables are purged, it recovers quickly after the purge procedure has completed.

Interaction Concentrator 7.6 Performance

Genesys has conducted a number of tests to examine the performance and scalability of ICON 7.6 in a large-scale environment. Two separate test environments were utilized—a voice-specific deployment and an open media (e-mail) deployment.

The following sections describe the sample 7.6 environments, call flows, tests, results, and recommendations.

Voice-Specific Environment

In the voice-specific test environment, two separate ICON applications were configured to collect parallel calls on two separate platforms, as shown in Figure 18.

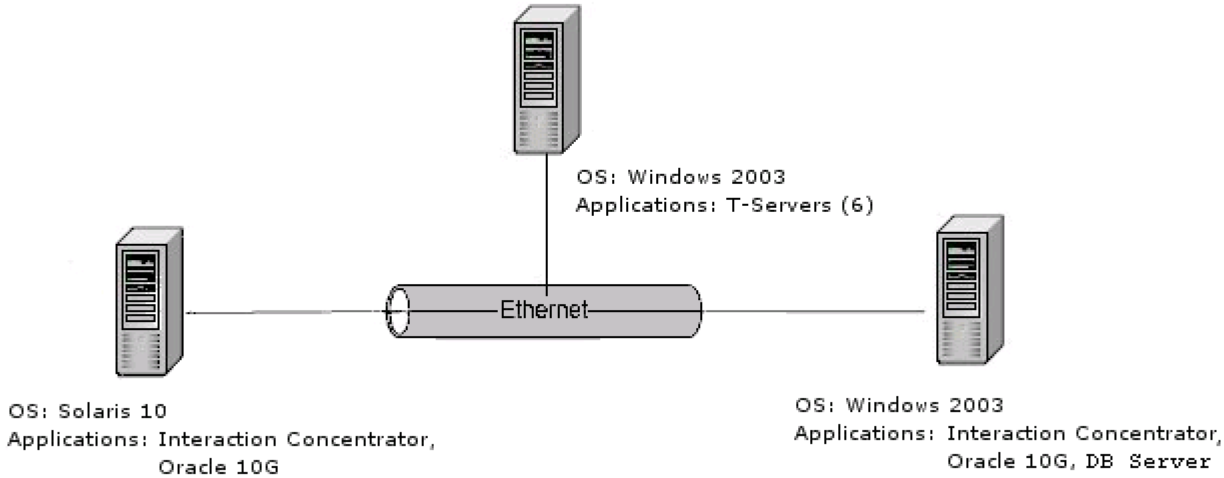

Figure 18: ICON 7.6 Voice-Specific Test Environment

Hardware and Software

Table 31 describes the hardware that hosted the ICON 7.6 application, persistent queue, and other components in the voice-specific test environment.

Table 31: Hardware and Software in Voice Test Environment

Application	Processor	Memory	Application Software Version	
			On Windows 2003 Server[a]	On Solaris 10[b]
ICON	Sunfire Dual Ultra-Sparc IIIi 1.504 GHz	2GB RAM	See Table 32	See Table 32
RDBMS (Oracle 10)	Dual dual-core Xeon 5160 3.0 GHz	4 GB RAM	10.2.0.4.0	n/a
DB Server	Dual dual-core Xeon 5160 3.0 GHz	4 GB RAM	7.6.000.08	n/a
Stat Server	Dual Xeon 2.8 GHz	1 GB RAM	7.6.000.18	n/a
URS	Dual Xeon 2.8 GHz	1 GB RAM	7.6.001.06	n/a
T-Server (Avaya G3)	Dual Xeon 3.06 GHz	2 GB RAM	7.6.000.04	n/a

a. Windows Server 2003 Enterprise Edition SP2 32-bit

b. Solaris 10 (Sun OS 5.10) 64-bit 11/06

Table 32: ICON Releases Tested

Test	ICON Release
Comparison Test	7.5.000.22 and 7.6.000.14
Attached Data	7.6.000.16
Vary Call Rate	7.6.000.14
Data Filtering	7.6.000.14
ICON Recovery	7.6.000.16

Call Flow and Configuration

Figure 19 shows the sample call flow and configuration for the scenarios that were used in the voice-specific tests.

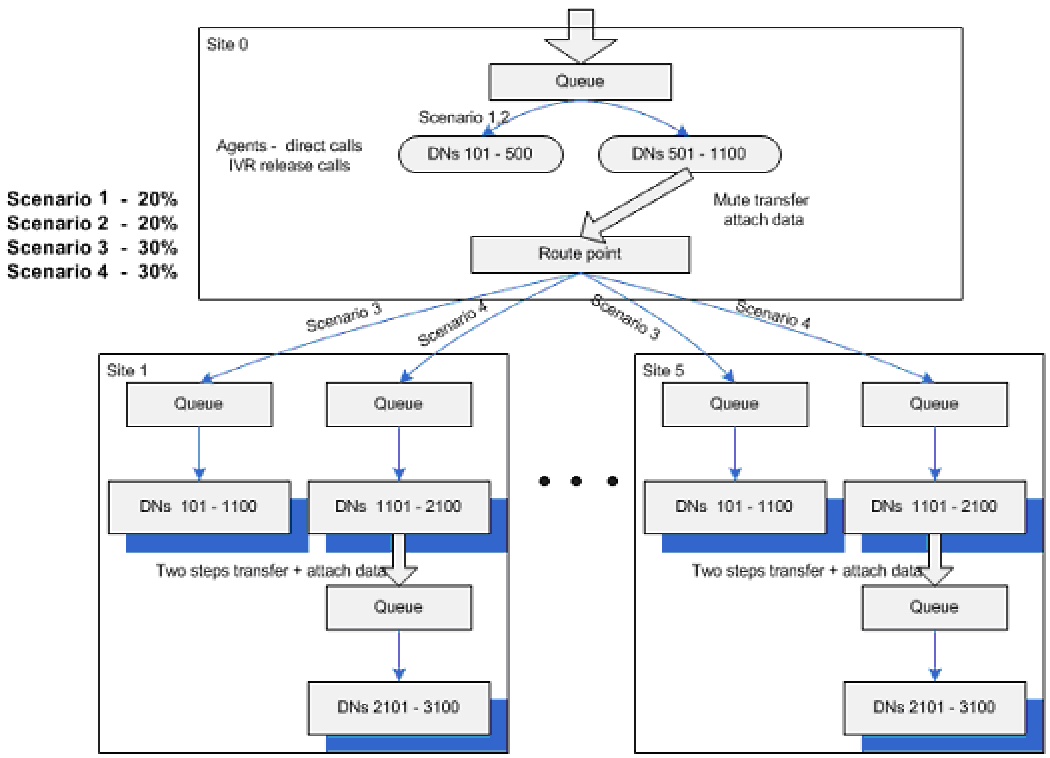

Figure 19: Sample Call Flow and Configuration (Voice Only)

The following four scenarios were tested with this configuration:

Scenario 1 Calls routed directly to an agent and released (20 percent of total calls)

Scenario 2 Calls routed to an IVR and released (20 percent of total calls)

Scenario 3 Calls routed to an IVR, transferred to a routing point, transferred to an agent, and then released (30 percent of total calls)

Scenario 4 Call routed to an IVR, transferred to a routing point, transferred to an agent, transferred and routed to a second agent, and then released (30 percent of total calls)

Tests and Results

The following tests were run in the voice-specific environment:

- ICON 7.5 and 7.6 Comparison Test (see page 74)
- Attached-Data Varied Test (see page 75)
- Call-Rate (Voice) Varied Test (see page 75)
- ICON Recovery (DB Server Disconnect) Test (see page 77)
- Data-Filtering Test (see page 78)

The details of these tests and the obtained results are discussed in the following sections.

ICON 7.5 and 7.6 Comparison Test

Genesys monitored ICON 7.5 and 7.6 over an extended period of time to ensure that no performance degradation occurred between releases. The test was conducted under the following conditions:

- Merge procedure run at 5-minute intervals
- Ten (10) predefined key-value pairs (KVPs) of data attached to every call.
- Five (5) days' endurance run
- Call rate of 14 calls/second (cps)

Test Results

Performance results are consistent between the 7.5 and 7.6 releases of ICON. Table 33 provides the observed results. Note that all values have been rounded up and are approximate.

Table 33: ICON 7.5 and 7.6 Comparison Test Results

Application	Average CPU (%)	Maximum CPU (%)	Maximum RAM (MB)	Average RAM (MB)
Solaris Operating System				
ICON 7.6	33	42	1,235	1,235
ICON 7.5	33	42	1,021	1,021

Table 33: ICON 7.5 and 7.6 Comparison Test Results

Application	Average CPU (%)	Maximum CPU (%)	Maximum RAM (MB)	Average RAM (MB)
Windows Operating System				
ICON 7.6	22	~60-80	712	708
ICON 7.5	19	~70	600	600

Attached-Data Varied Test

In this test, Genesys varied the number of key-value pairs that were attached to each call to observe the effect on CPU and RAM. The test was conducted under the following conditions:

- Number of attached KVPs varied from 10 to 50
- Average of 255 bytes of data attached to KVPs
- Merge procedure run at 5-minute intervals
- Duration of test: 2 hours
- Five (5) days of data prior to start

Test Results

Table 34 shows that increasing the number of KVPs had a linear affect on ICON storage. The CPU increased quickly with the number of KVPs, but had no effect on the merge time, which took an average time of 4 seconds. Note that all values have been rounded up and are approximate.

Table 34: Attached-Data Varied Test Results

# KVPs Attached/call	Average CPU (%)	Maximum CPU (%)	Maximum RAM (MB)	Average RAM (MB)	Merge Time (s)
10	33	40	887	880	3.6
30	40	48	1,017	1,002	4
50	46	58	998	984	4

Call-Rate (Voice) Varied Test

In this test, Genesys varied the call rate to observe the effect on the CPU and the merge time. The test was conducted under the following conditions:

- Call volume: Variable, starting at 7 calls/second, incrementing to 30 calls/second

- Merge procedure run at 5-minute intervals
- Duration of test: 2 hours
- Five (5) days of data prior to start

Test Results

Table 35 shows that the following results were produced, as the call rate increased:

- The average RAM was fairly constant; approximately 890 MB for call rates of 10 and 14calls/second, and a slightly higher 950 MB for call rates of 10, 15, 20, 25, 28, and 30 calls/second. Note that all values have been rounded up and are approximate.
- The duration of the merge procedure increased from 2 to 11 seconds (see Figure 20).

Table 35: Call Rate Varied Test Results

Call Rate (cps)	Average CPU (%)	Maximum CPU (%)	Maximum RAM (MB)	Average RAM (MB)
7	16	20	887	887
10	23	28	948	948
14	33	40	887	880
15	35	42	948	948
20	48	56	948	948
25	60	70	948	948
28	68	80	948	945
30	73	84	961	958

Figure 20 shows the linear relationship between the duration of the merge procedure and the call rate.

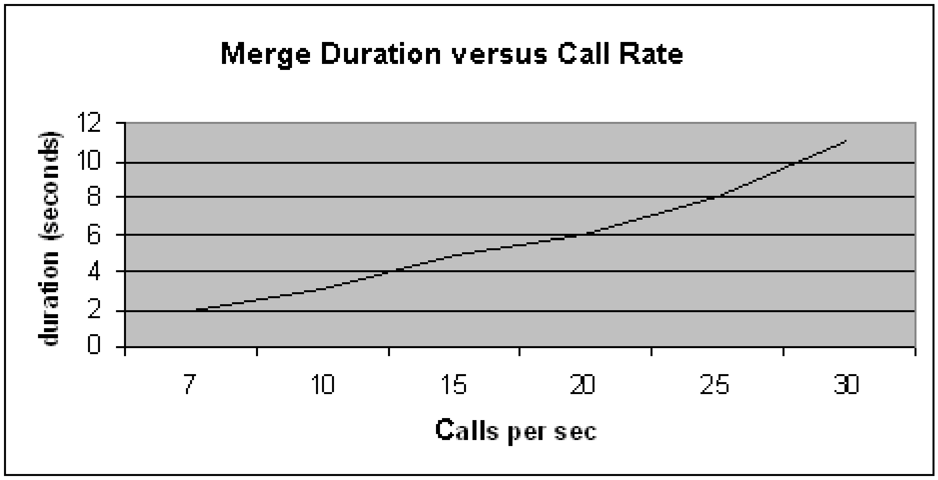

Figure 20: ICON Merge Duration Versus Call Rate

ICON Recovery (DB Server Disconnect) Test

Genesys has conducted a recovery test in which ICON was disconnected from DB Server for a period of time and then reconnected to determine how long it would take for ICON to recover (clear the persistent queue). The test was conducted under the following conditions:

- Call flow of 14 calls/second
- Merge procedure run at 5-minute intervals
- Duration of disconnection from DB Server: 0 minutes and 120 minutes

Test Results

After a disconnect from DB server of 2 hours, it took ICON 72 minutes to clear the persistent queue.

During the period of disconnect, testing produced the following results:

- The average CPU increased to 50 percent.
- RAM increased slightly from an average of 880 MB to 962 MB (see Table 36).
- Merge time increased to an average of 20 seconds (peak of 69 seconds). Typically, the merge time is about 4 seconds (see Figure 21).
- ICON took approximately 60 percent of the time that it was disconnected from DB Server to recover (clear the persistent queue, or PQ).

Note that all values in Table 36 have been rounded up and are approximate.

Table 36: ICON Recovery Test Results

DB Server Disconnect Time (minutes)	Time to Clear PQ (minutes)	Average CPU (%)	Maximum CPU (%)	Maximum RAM (MB)	Average RAM (MB)
0	N/A	33	40	887	880
120	72	50	64	962	962

Figure 21 shows the backlog of transactions in the persistent queue while ICON was disconnected from DB Server. After its reconnection to DB Server, ICON cleared the backlog within approximately 72 minutes.

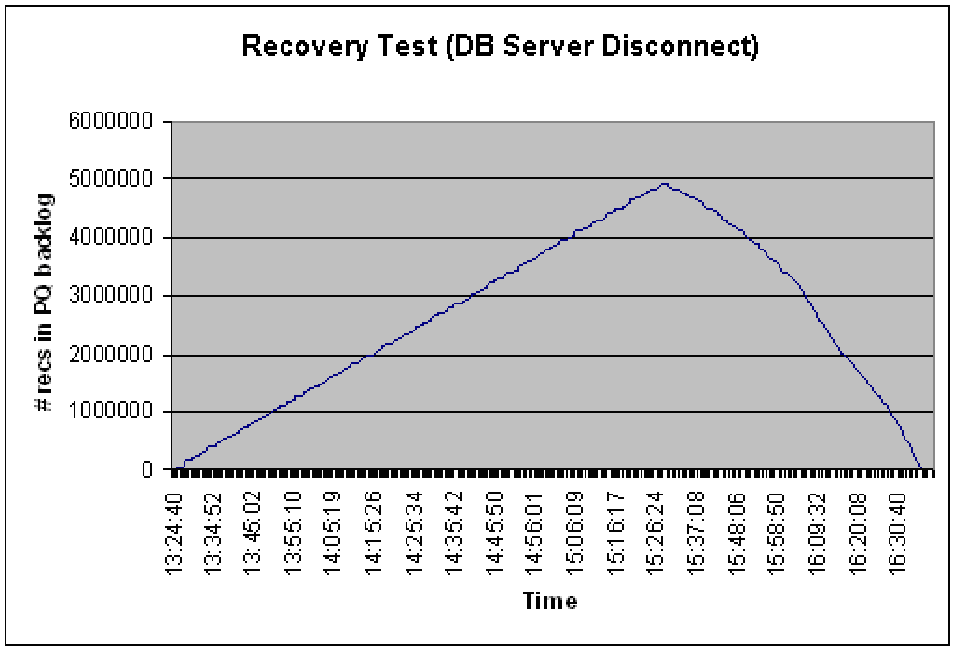

Figure 21: Records in ICON's Persistent Queue During and After DB Server Disconnect

Data-Filtering Test

In this test, the performance of ICON was measured when data filter options were set to filter out (exclude) data from IDB storage. The tests were conducted under the following conditions:

- Filters set to: `none`, three (`call-metrics`, `call-history`, `ir-history`), and `all`
- Merge procedure run at 5-minute intervals
- Call flow of 14 calls/second
- Duration of test: 60 minutes and 130 minutes

Note: For more information about ICON's data-filtering feature, see the *Interaction Concentrator 7.6 User's Guide.*

Test Results

There was no impact on the CPU when three data filters were applied versus none. When all of the data filters were applied, however, the maximum CPU decreased by 25 percent, as shown in Table 37. Increasing the number of data filters also decreased the rate of data storage (see Figure 22).

Table 37: Data Filtering Test Results

Data Filters Set	Average CPU (%)	Maximum CPU (%)	Maximum RAM (MB)	Average RAM (MB)
None	33	40	887	880
Three (`call-metrics`, `call-history`, `ir-history`)	32	40	924	913
All	24	30	752	747

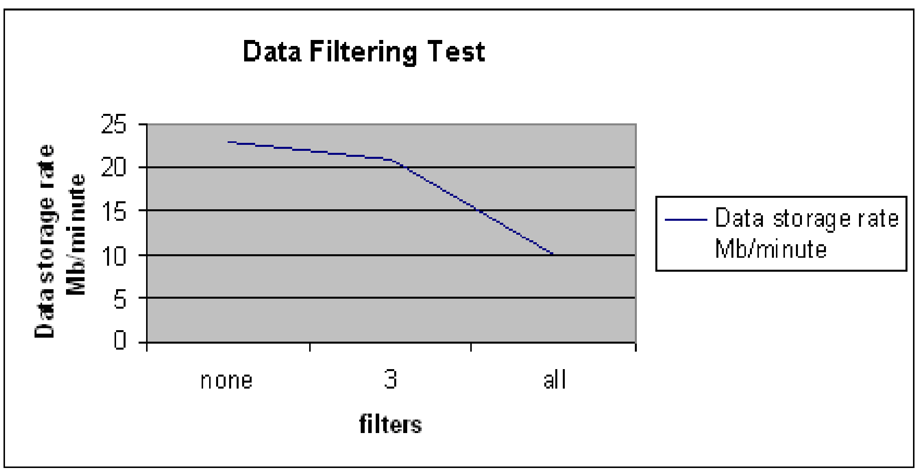

Figure 22: Rate of Data Storage Versus Number of Data Filters

Open Media Environment

In the open media test environment, a single ICON application—running on a Windows platform, and writing to an Oracle database on a Windows Server 2003—was configured to collect open media (e-mail) interactions from Interaction Server (see Figure 23).

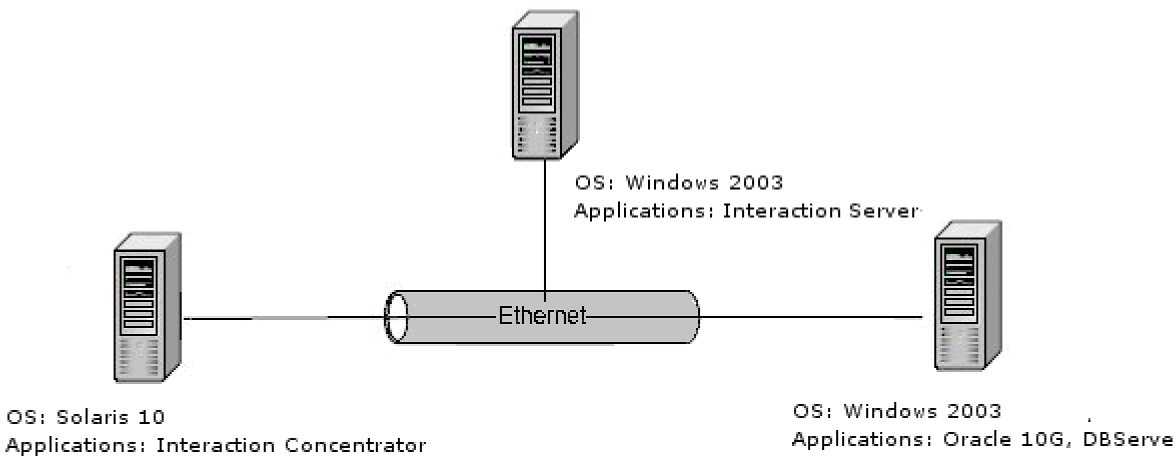

Figure 23: ICON 7.6 Open Media Test Environment

Hardware and Software

Table 38 describes the hardware that hosted the ICON application, persistent queue, and other components in the open media environment.

Table 38: Hardware and Software in ICON Open Media Environment

Application	Operating System	Processor	Memory	Software Version
ICON	Solaris	T5220 1.2 GHz 8 Core, 64 thread	32 GB RAM	7.6.000.16 7.6.100.00
RDBMS (Oracle 10)	Windows	Dual dual-core Xeon 5160 3.0 MHz	4 GB RAM	10.2.0.4.0
DB Server	Windows	Dual dual-core Xeon 5160 3.0 MHz	4 GB RAM	7.6.000.08
Stat Server	Windows	Dual dual-core Xeon 5160 3.0 MHz	4 GB RAM	7.6.000.18
URS	Windows	Dual dual-core Xeon 5160 3.0 MHz	4 GB RAM	7.6.001.06
Interaction Server and Proxies	Windows	Dual dual-core Xeon 5160 3.0 MHz	4 GB RAM	7.6.000.10

Call Flow and Configuration

In the open media model, the interaction flow is as follows:

1. The e-mail arrives at the first queue, where 10 KVPs of data are attached by the routing strategy.
2. The e-mail moves through the next two queues each of which has a strategy loaded onto it.
3. The e-mail arrives at an agent, where it is handled for 120 seconds.
4. The interaction moves through a final queue and routing strategy, before it is terminated.

Tests and Results

The following test was run in the open media environment:

- Rate of E-ails Varied Test (see page 81)

The following sections provide a detailed description of this test and a discussion of the obtained results.

Rate of E-Mails Varied Test

In this test, Genesys varied the number of e-mail interactions that were submitted during 2 hours of testing and with 1 day of previously processed data in IDB (approximately 1.2 million e-mails). The test was conducted under the following conditions:

- Data attached as KVPs to each e-mail at the first queue
- Forty-five hundred (4500) agents and four (4) queues configured
- One-hundred twenty (120) seconds handling time
- Rate of submitted e-mails: 7 e-mails/second and 14 e-mails/second

Test Results

Table 39 shows that, at a rate of 14 e-mails/second, the CPU averaged 93 percent and reached a maximum of 99 percent. As the rate of submitted e-mails increased, ICON's CPU increased linearly (see Figure 24).

Table 39: Rate of E-Mails Varied Test Results

Submitted Rate (e-mails/sec)	Average CPU (%)	Maximum CPU (%)
7	50	53
14	92	99

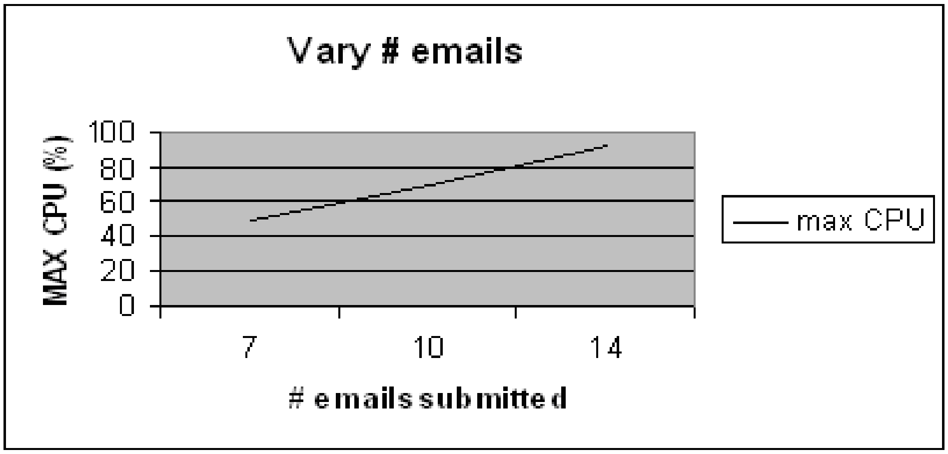

Figure 24: E-Mails Submitted Versus CPU

ICON 7.6 Performance Conclusions

Based on the results of performance testing that was conducted in a large-scale environment, Genesys has the following conclusions and recommendations:

- There is no degradation in performance between ICON releases 7.5 and 7.6 on either Solaris or Windows operating systems.
- The call rate and the number of key-value pairs (attached data) are significant factors that affect ICON's scalability.
- ICON recovers quickly from DB Server disconnects—typically, recovering in a shorter time interval than the original outage.
- Filtering of data did not have a significant impact on RAM and CPU, when a small number of filters (three) was set. The reduction was more significant— 43 percent reduction in RAM and 25 percent reduction in CPU— when all filters were set.

Interaction Concentrator 7.5 Performance

Genesys has run several 7.5 performance tests for selected combinations of the following variables:

- Operating system (OS)—Windows (Microsoft Windows Server 2003 Enterprise), UNIX (Solaris 10)
- Relational database management system (RDBMS)—Microsoft SQL Server 2000, Oracle 10
- Call volume—starting at 10 calls/second, incrementing by 5 calls/second

The following sections describe the sample ICON 7.5 environment, call flows, tests, results, and recommendations.

Sample Environments

This subsection describes the Genesys sample 7.5 environments and conditions:

- Hardware specifications for the ICON server (see page 83)
- Hardware specifications for the DB Server, RDBMS server, and IDB (see page 84)
- Database settings (see page 84)
- Environment configuration and call flows (see page 85)

Hardware Specifications

Table 40 describes the hardware that hosted the ICON 7.5 application and persistent queue.

Table 40: Hardware Specifications—ICON 7.5 Application

Windows	
Operating System	Microsoft Windows Server 2003 Enterprise
Processor	2 x Intel Xeon 3.2 GHz/2 MB
Memory	4 GB RAM
Disk Storage	2 HDD x 80 GB SATA
UNIX	
Operating System	Solaris 10
Processor	2 x Sun UltraSPARC-IIIi 1.0 GHz/1 MB
Memory	2 GB RAM
Disk Storage	2 x 36GB Ultra160 SCSI drives

Table 41 describes the hardware that hosted the DB Server, RDBMS Server, and IDB.

Table 41: Hardware Specifications—DB Server, RDBMS Server, and IDB

Windows	
Operating System	Microsoft Windows Server 2003 Enterprise
RDBMS	Microsoft SQL Server 2000
Processor	2 x Intel Xeon 3.2 GHz/2 MB
Memory	4 GB RAM
Disk Storage	2 HDD x 80 GB SATA
UNIX	
Operating System	Solaris 10
RDBMS	Oracle 10
Processor	SUN SPARC Enterprise 9000, 16 GB memory, 4 x CPU quad core 2.8 GHZ [a]
Memory	8 GB Real RAM
Disk Storage	2 Gbit Sun StorEdge 3510 Fiber Channel Array

a. Each core supports four hardware threads of execution. The hardware threads are scheduled on the core's processing unit in round-robin order. A different software thread can run on each hardware thread, so 32 software threads can run in parallel on a single T1 processor.

Database Settings

Table 42 summarizes the Microsoft SQL database settings that were used in the Genesys sample 7.5 configurations.

Table 42: Microsoft SQL Database Settings

Setting	Value
Memory	2497 MB
Auto update statistics	On
Torn page detection	On

Table 42: Microsoft SQL Database Settings (Continued)

Setting	Value
Auto create statistics	On
Model	Simple

Table 43 summarizes the Oracle database settings that were used in the Genesys sample 7.5 configurations.

Table 43: Oracle Database Settings

Setting	Value
Memory	2,543 MB
Database block size (`db_block_size`)	8 KB
File system I/O options (`filesystemio_options`)*	directIO
*The Solaris file system was mounted with the option `forcedirectio`.	

Environment Configurations and Call Flows

The Genesys 7.5 performance results were obtained under the following conditions:

- Genesys used the following site configurations and call-flow scenarios:
  - Single site, with simple call flow (see page 85)
  - Single site, with complex call flow (see page 86)
  - Multi-site (single ICON, single IDB), with complex call flow (see page 88)

  The call flows include several Genesys-defined attached-data items, and eight user-defined attached data key-value pairs (approximately 200 bytes).
- The ICON `Application` used the default configuration settings.
- The logging level was set to `Standard`.
- All tests started with an empty database.
- In the tests that used the multi-site scenario with complex call flow, the merge procedure was run at 5-minute intervals.
- The ICON Performance Counter pages were queried frequently to obtain the reported measurements.

Single Site (Simple Call Flow)

Figure 25 shows the configuration and call flow for the simple call-flow scenario in a single-site deployment. The routing strategy has nine target queues, and each queue has 100 logged-in agents.

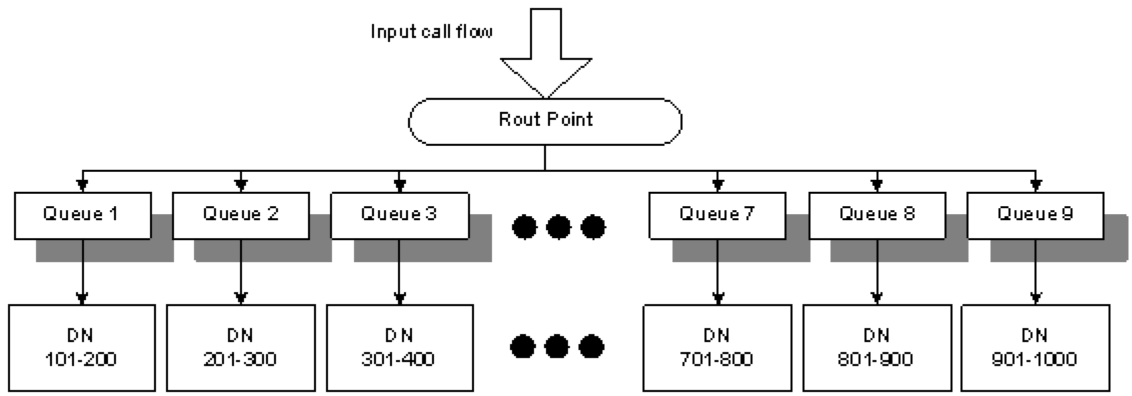

Figure 25: Single Site, Simple Call Flow

The call flow is as follows:

1. A new inbound call arrives at the Routing Point.
2. URS sends the call to one of the nine target queues.
3. The call is directed to one of the agents who are logged in to the queue.
4. The agent answers the call, handles the interaction, and then releases the call. The agent talk time is 40 seconds.

Single Site (Complex Call Flow)

Figure 26 shows the configuration and call flow for the complex call-flow scenario in a single-site deployment. There is an input queue, which has 400 logged-in DNs. The routing strategy has six target queues, and each queue has 400 logged-in DNs. There are 600 agents logged in to two additional queues.

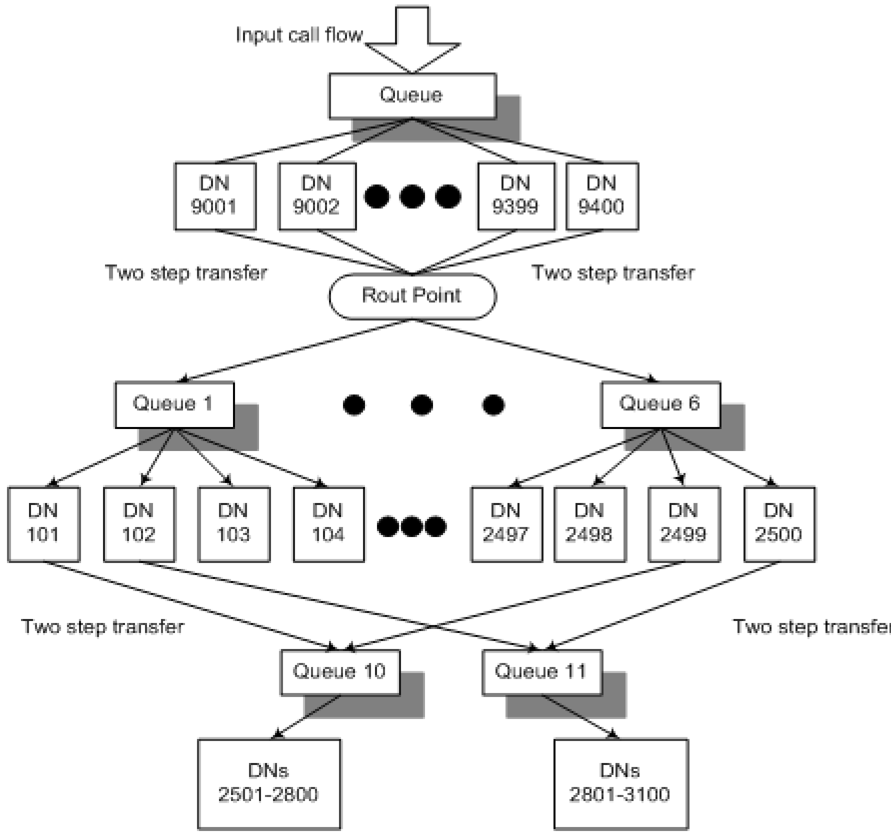

Figure 26: Single Site, Complex Call Flow

The call flow is as follows:

1. A new inbound call arrives at the input queue.
2. The call is directed to 1 of the 400 DNs that are logged in to the queue.
3. An agent simulator on these DNs attaches some business data and performs a two-step transfer to send the call to the Routing Point.
4. URS sends the call to one of the six target queues.
5. The call is directed to one of the agents who are logged in to the queue.
6. The call flow scenario splits:
   - For 90 percent of the agents logged in to the target queues, the agent answers the call, handles the interaction, and then releases the call. The agent talk time is 30 seconds.
   - For the remaining 10 percent of the agents logged in to the target queues, the agent answers the call, attaches some additional data, and then makes a two-step transfer to one of two additional queues, where there are a total of 600 logged in agents. The talk time of the agent in the last queue is 15 seconds, before the agent releases the call.

Multi-Site (Complex Call Flow)

Figure 27 shows the configuration and call flow for the complex call flow scenario in a multi-site deployment. There is an input queue, which has 900 logged-in DNs. The routing strategy has one target queue at each of five sites, and each queue has 900 logged-in DNs. Each site also has 1,100 agents logged in to each of two additional queues.

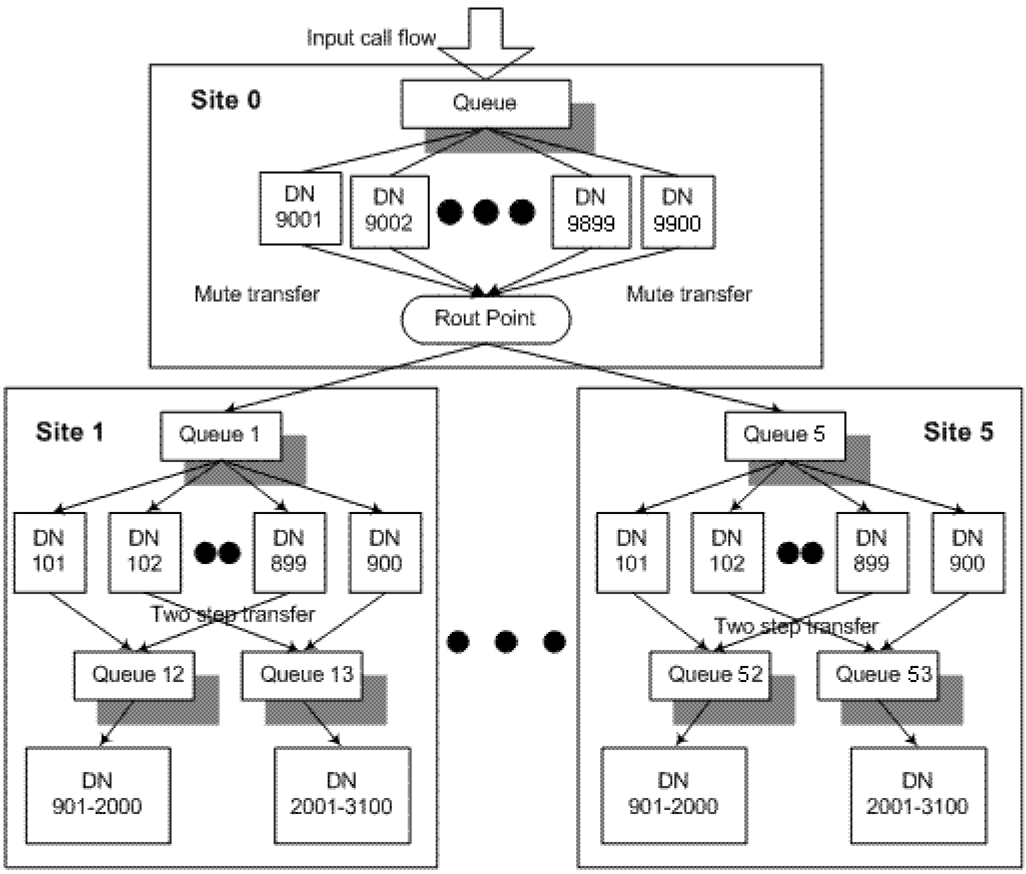

Figure 27: Multi-Site, Complex Call Flow

The call flow is as follows:

1. A new inbound call arrives at the input queue.
2. The call is directed to 1 of the 900 DNs that are logged in to the queue.
3. An agent simulator on these DNs attaches some business data and performs a mute transfer to send the call to the Routing Point.
4. URS sends the call to one of the five target queues.
5. The call is directed to one of the agents who are logged in to the queue.
6. The agent answers the call, attaches some additional data, and then makes a two-step transfer to one of two additional queues.
7. The last agent talk time is 40 seconds, before the agent releases the call.

ICON 7.5 Performance Conclusions

The Genesys performance results report Interaction Concentrator 7.5 runtime performance, related to the database-writing side of Interaction Concentrator activity. Call simulation and performance measurement started after ICON initialization was complete.

Table 44 summarizes the application processing usage and related recommendations for maximum call flows, on the basis of the performance results. The maximum call flows represent the average number of interactions per second that ICON can write to IDB without unsustainable persistent queue backlogs being generated.

For more information about the call-flow scenarios that produced these performance results, see "Environment Configurations and Call Flows" on page 85.

Table 44: Maximum Call Flows and Application Usage, by Scenario

Scenario	Maximum Call Flow (calls/sec)	Application Processing Usage		
		ICON CPU (%)	ICON Memory (MB)	RDBMS CPU (%)
OS: Windows Server 2003 Enterprise RDBMS: Microsoft SQL Server 2000				
Single-Site, Simple Call Flow	75	46	431	42.9
Single-Site, Complex Call Flow	60	57.9	421	70.5
Multi-Site, Complex Call Flow*	25	40.9	437	58.6
OS: Solaris 10 RDBMS: Oracle 10 * **				
Single-Site, Simple Call Flow	30**	38	725	4
Single-Site, Complex Call Flow	15**	40	757	4
Multi-Site, Complex Call Flow*	32**	42	854	52

*The merge procedure was run at 5-minute intervals.

**The hardware configuration of the database host had an adverse impact on ICON performance.

***Single side data was obtained using T2000 class of Servers while data for complex call flow was obtained using l M9000 class Server for DB hosting.

Endurance Test

In addition to the incremental call-flow performance tests, Genesys subjected Interaction Concentrator 7.5 to an endurance test under the following conditions:

- OS: Windows Server 2003 Enterprise
- RDBMS: Microsoft SQL Server 2000
- Call-flow scenario: Multi-Site (Complex Call Flow) (for more information, see page 88)
- Call speed: 25 calls/second
- Merge-procedure frequency: At 5-minute call intervals
- Duration of test: Approximately 72 hours

Endurance Test Results

Table 45 summarizes the endurance test results. All values have been rounded up and are approximate.

Table 45: Endurance Test Results

Measurement	Value
ICON CPU usage (%)	35
ICON memory usage (MB)	612
RDBMS CPU usage (%)	52
RDBMS memory usage (MB)	1,750
DB Server multiserver CPU usage (%)	0.5
DB Server multiserver memory usage (MB)	12
DB Server client 1 CPU usage (%)	9
DB Server client 1 memory usage (MB)	50
DB Server client 2 CPU usage (%)	6
DB Server client 2 memory usage (MB)	32
Merge time (sec)	30–40
Total number of calls	6,210,000
Total number of ICON transactions	372,887,000

Recovery Test

In addition to the incremental call-flow performance tests, Genesys subjected Interaction Concentrator 7.5 to a recovery test, following disconnection from DB Server. The test was conducted under the following conditions:

- OS: Windows Server 2003 Enterprise.
- RDBMS: Microsoft SQL Server 2000.
- Call flow scenario: Multi-Site (Complex Call Flow) (for more information, see page 88).
- Call speed: 10 calls/second. Call flow was not stopped when DB Server reconnected.
- Duration of disconnection from DB Server: 15 minutes.

Recovery Test Results

While ICON was disconnected from DB Server, transactions backlogged in the persistent queue. After reconnection to DB Server, ICON cleared the backlog within approximately 17 minutes.

On the basis of the recovery test results, Genesys suggests that the recovery time can be estimated as follows:

```
Recovery time = Time_without_DBServer * Input_call_flow /
(Max_call_flow – Input_call_flow)
```

Chapter

5 Genesys Info Mart Solution

The information in this chapter applies to the 7.2, 7.5, and 7.6 releases of Genesys Info Mart. Genesys Info Mart release 7.2 uses Call Concentrator as the primary source of data. Genesys Info Mart releases 7.5 and 7.6 use Interaction Concentrator as the primary source of data.

Genesys Info Mart produces a data mart containing several star schemas that you can use for contact center historical reporting. Before proceeding, review the "Architecture" chapter in the *Genesys Info Mart Deployment Guide* for your release of Genesys Info Mart to familiarize yourself with the product architecture.

This chapter consists of the following sections:

- Genesys Info Mart 7.2 Database Size Estimation, page 179

Refer to the sections that are applicable to your release of Genesys Info Mart.

Release 7.6 Guidelines

The following sections provide hardware sizing guidelines for environments with Genesys Info Mart 7.6:

- "Hardware Architectures in Release 7.6" on page 95, provides architecture samples for Genesys Info Mart release 7.6, for small and medium-size contact centers.
- "Genesys Info Mart 7.6 ETL Runtime Performance" on page 99, lists the factors that affect Genesys Info Mart 7.6 performance. It also provides general configuration details that apply to both sample configurations.
- "Release 7.6 ETL Performance for Sample Configuration 1" on page 106, lists specific performance measurements for Genesys Info Mart 7.6 as deployed in a larger contact center environment on both Oracle and Microsoft SQL Server database platforms.
- "Release 7.6 ETL Performance for Sample Configuration 2" on page 118, lists specific performance measurements for Genesys Info Mart 7.6 as deployed in a smaller contact center environment on a Microsoft SQL Server database platform.
- "Genesys Info Mart 7.6 Performance Tuning" on page 124, provides a series of guidelines to help you tune your database environment for superior performance.
- "Genesys Info Mart 7.6 Database Size Estimation" on page 127, shows you how to estimate the size of the Staging Area and the Genesys Info Mart databases.

The information is organized according to the size of the contact center as a function of the daily number of interactions and the number of agents. Table 46 provides the definition of the contact center size, as applicable to release 7.6 of Genesys Info Mart.

Table 46: Contact Center Size for Release 7.6

Size Category	Number of Agents	Daily Number of Interactions	Maximum Calls Per Second
Small	Fewer than 150	Fewer than 50,000	1.5
Medium	Fewer than 1,000	Fewer than 300,000	12
Large	More than 1,000	More than 300,000	> 12

Note: Large contact centers require special planning because of their variability and complexity. Contact Genesys Technical Support if you want to plan a Genesys Info Mart deployment for a large contact center.

Hardware Architectures in Release 7.6

The examples in this section are organized by the size of your contact center and your choice of operating system.

The hardware architecture you select for your Genesys Info Mart 7.6 deployment depends primarily on:

- The size of your contact center in relation to the number of daily interactions and the number of agents.
- The time available for Genesys Info Mart to extract, transform, and load data.
- Your choice of operating system.

To minimize the time it takes for Genesys Info Mart 7.6 to extract, transform, and load data, Genesys strongly recommends that you run the Genesys Info Mart Server and the Staging Area relational database management system (RDBMS) server on different machines.

Notes:

- Hardware recommendations for the Interaction Concentrator, Stat Server RDBMS, and Genesys Info Mart RDBMS servers are outside the scope of this guide.
- The type of hardware you deploy depends on whether this database is used directly for reporting (and the number of concurrent users), or as a secondary staging area to hold the data until it can be uploaded to a data warehouse.

Small Contact Center—Windows—Release 7.6

Figure 28 depicts a sample hardware architecture of Genesys Info Mart release 7.6 for a small-size contact center on a Microsoft Windows operating system.

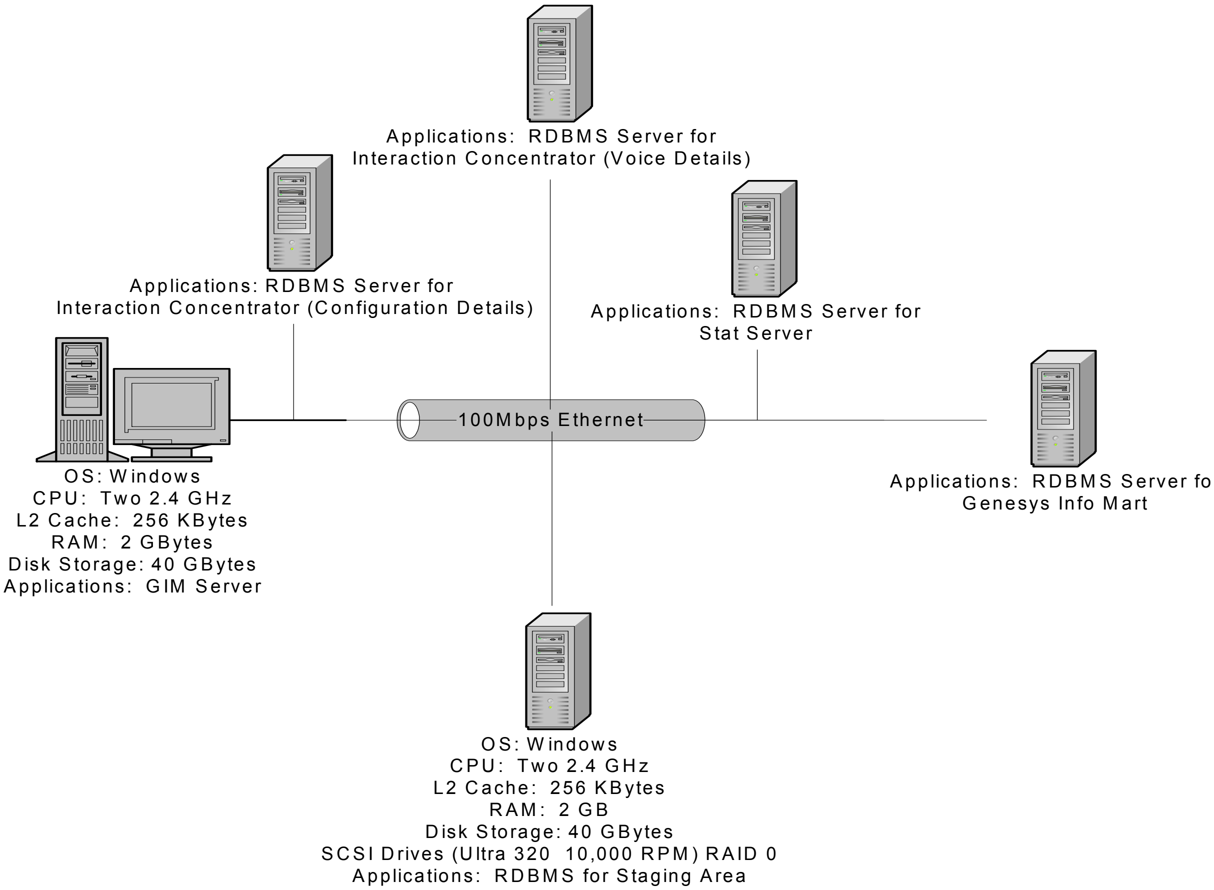

Figure 28: Small-Size Contact Center—Windows—Release 7.6

Note: Genesys strongly recommends that you use a Staging Area RDBMS disk subsystem that contains multiple SCSI disk drives. Storing RDBMS logs, indexes, and table data on different disk drives improves extraction, transformation, and loading (ETL) performance by reducing I/O contention.

Medium Contact Center—Windows—Release 7.6

Figure 29 depicts a sample hardware architecture of Genesys Info Mart release 7.6 for a medium-size contact center on a Microsoft Windows operating system.

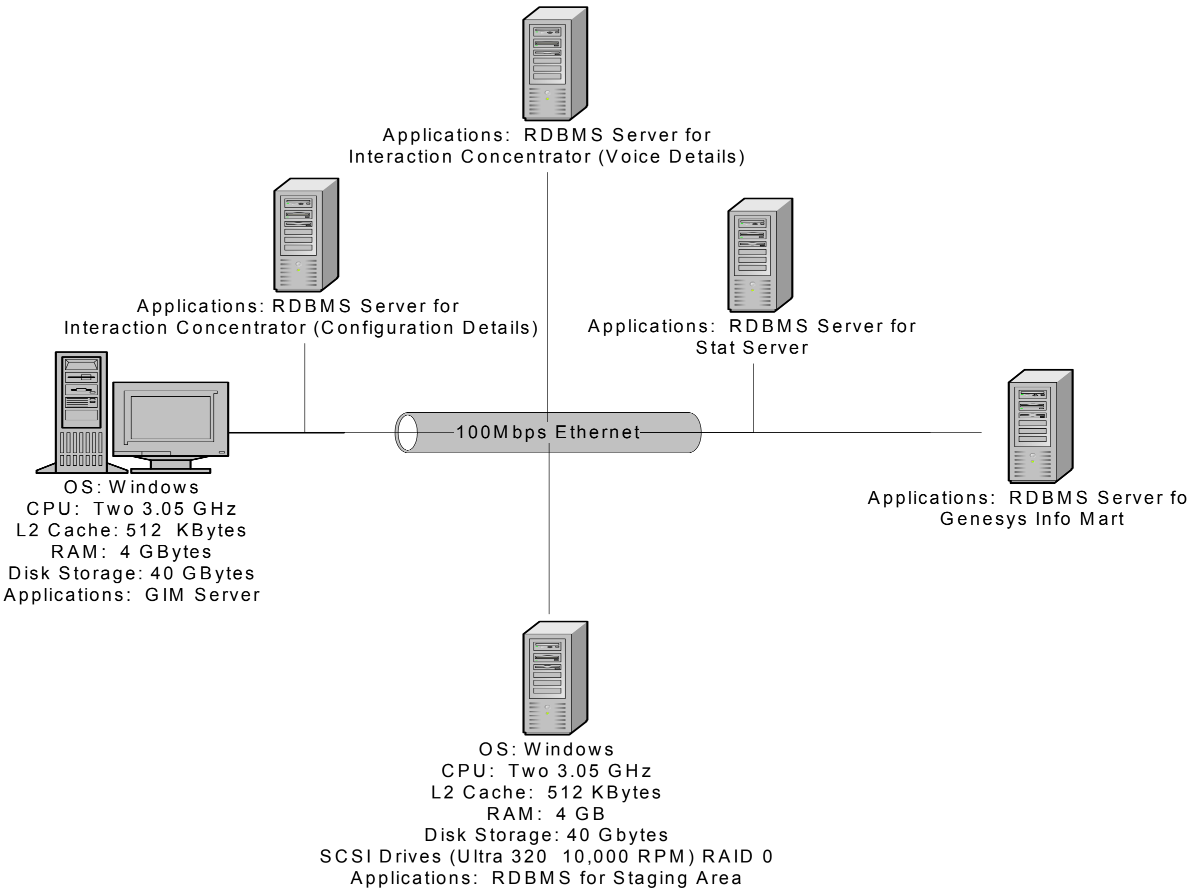

Figure 29: Medium-Size Contact Center—Windows—Release 7.6

Note: Genesys strongly recommends that you use a Staging Area RDBMS disk subsystem that contains multiple SCSI disk drives. Storing RDBMS logs, indexes, and table data on different disk drives improves extraction, transformation, and loading (ETL) performance by reducing I/O contention.

Small Contact Center—UNIX (Solaris)—Release 7.6

Figure 30 depicts a sample hardware architecture of Genesys Info Mart release 7.6 for a small-size contact center on a UNIX (Solaris) operating system.

To obtain the best Genesys Info Mart Server performance on UNIX (Solaris), Genesys recommends selecting a machine that has the most powerful processors, rather than one with multiple slower processors. For example, a machine with two 990-MHz processors performs better than one with four 440-MHz processors.

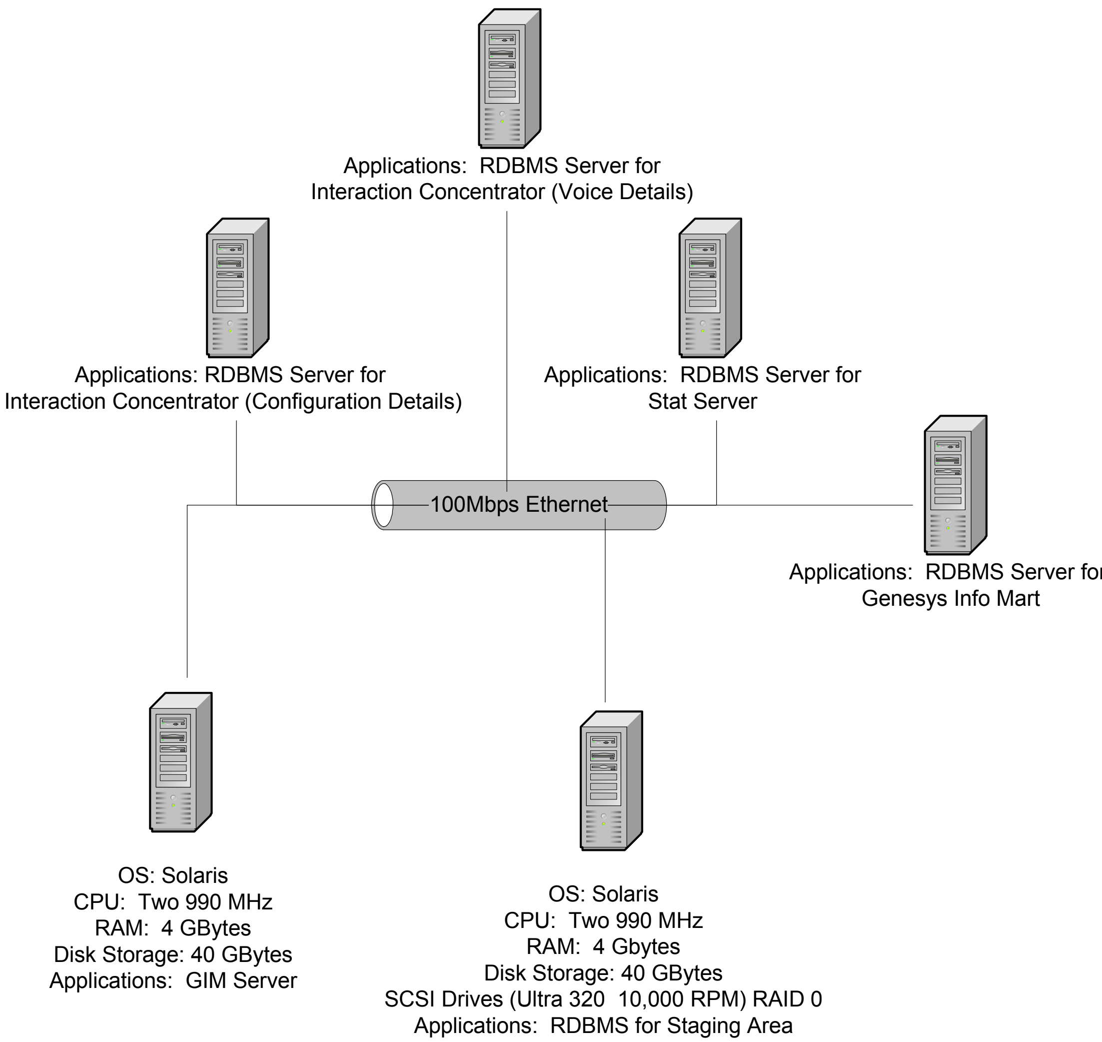

Figure 30: Small-Size Contact Center—Solaris—Release 7.6

Note: Genesys strongly recommends that you use a Staging Area RDBMS disk subsystem that contains multiple SCSI disk drives. Storing RDBMS logs, indexes, and table data on different disk drives improves extraction, transformation, and loading (ETL) performance by reducing I/O contention.

Genesys Info Mart 7.6 ETL Runtime Performance

This section provides information on the following topics:

About Genesys Info Mart 7.6 Performance Testing

To test Genesys Info Mart 7.6 ETL performance, the ETL cycle was run on actively populated source data for the following test contact center sizes:

- Sample Configuration 1 handling up to 1,000,000 calls a day, over a period of three consecutive days
- Sample Configuration 2 handling up to 260,000 calls a day, over a period of seven consecutive days

Testing was performed on two database platforms:

- Oracle 10g for Sample Configuration 1
- Microsoft SQL Server 2005 for Sample Configuration 2

Genesys Info Mart 7.6 was tested with Interaction Concentrator 7.6.

General Configuration Details

Testing for both Sample Configurations 1 and 2 used identical call flows and Genesys Info Mart configuration, but different call volumes. The call rate was constant for a full 24 hours each day, with no scheduling for peak or off-peak hours.

For further configuration details that apply generally to all of the release 7.6 tests, see the following sections:

- "Release 7.6 Performance Testing Configuration" on page 100
- "Scheduling and Configuration of the ETL Jobs" on page 104

Performance Results

For some performance trends that applied generally for all of the release 7.6 tests, see "General Performance Results for Release 7.6" on page 106.

You can find more specific performance results—including a series of graphs that show the duration of the various ETL jobs—in the following sections:

- “Release 7.6 ETL Performance for Sample Configuration 1” on page 106
- “Release 7.6 ETL Performance for Sample Configuration 2” on page 118

Release 7.6 Performance Testing Configuration

Table 47 provides configuration details that pertain to both sample configurations.

Table 47: Genesys Info Mart 7.6 Performance Testing Configuration Details

Parameter	Descriptions
Call flow	Testing consisted of four call flow scenarios: • Call is routed directly to an agent, and is then released (20 percent) • Call is routed to an Interactive Voice Response (IVR), and is then released (20 percent) • Call is routed to an IVR, and is then transferred to a Routing Point, and then to an agent, and is finally released (30 percent) • Call is routed to an IVR, is transferred to a Routing Point, and then to an agent; after which the call is transferred and is routed to a second agent, and is finally released (30 percent)
Attached data	20 predefined key-value pairs were stored in various Genesys Info Mart fact and dimension tables, as were five user data string dimensions and five user data facts.
Agents	The testing configuration defined about 15,000 agents. The average talk time for a call handled by an agent was 120 seconds.

Genesys Info Mart Application

Table 48 lists the Genesys Info Mart release 7.6 application options that were used for testing. It includes only those options that affected performance.

Table 48: Genesys Info Mart 7.6 Application Configuration Options

Section	Option
gim-aggregates-tenant	maximum-aggregation-level = MONTH populate-agent-state-aggregates = FALSE populate-agent-state-interval-aggregates = TRUE populate-ixn-agent-aggregates = TRUE populate-ixn-agent-interval-aggregates = TRUE populate-ixn-service-type-aggregates = TRUE populate-queue-aggregates = TRUE populate-skill-combination-aggregates = FALSE populate-skill-demand-aggregates = FALSE
gim-etl	aggregate-time-range-limit = 3 aggregate-time-range-units = DAYS days-to-keep-stg-ha-ir-ids = 3 default-ivr-to-self-service = FALSE extract-data-time-range-limit = 1 extract-data-time-range-units = HOURS extract-ha-voice-agent-activity = TRUE ha-agent-all-connections-required = FALSE ha-cfg-all-connections-required = TRUE ha-ir-extract-comparison-timeout = 0 ir-merge-interval = 3 limit-extract-data = TRUE load-transaction-size = 10000 max-session-duration-in-hours = 24 max-wrap-delay = 2 memory-threshold = 95 populate-chat-ixns = FALSE populate-chat-resource-activity = FALSE populate-detailed-ixn-subtype = FALSE populate-dt-chat-resource-activity = FALSE populate-dt-email-resource-activity = FALSE populate-dt-open-media-resource-activity = FALSE populate-dt-voice-resource-activity = FALSE populate-email-ixns = FALSE populate-email-resource-activity = FALSE populate-ocs-ixns = TRUE populate-open-media-ixns = FALSE

Table 48: Genesys Info Mart 7.6 Application Configuration Options (Continued)

Section	Option
gim-etl (continued)	populate-open-media-resource-activity = FALSE populate-sm-chat-resource-activity = FALSE populate-sm-email-resource-activity = FALSE populate-sm-open-media-resource-activity = FALSE populate-sm-voice-resource-activity = TRUE populate-voice-resource-states-for-queues = TRUE sub-hour-level-aggregation = 30
gim-etl-tenant	days-to-keep-dt-resource-activity-facts = 30 days-to-keep-gim-facts = 400
gim-transformation	show-abandoned-detail = TRUE show-conference-detail = TRUE transformation-buffer-size = 5 voice-init-resp-duration = TIME_TO_FIRST_AGENT
optional-tables	populate-acd-queue-facts = TRUE populate-dt-dnd-facts = FALSE populate-dt-resource-state-facts = FALSE populate-dt-resource-state-reason-facts = FALSE populate-gvp-var-facts = FALSE populate-interaction-resource-facts = TRUE populate-interaction-resource-state-facts = TRUE populate-place-group-facts = TRUE populate-resource-group-facts = TRUE populate-resource-session-facts = FALSE populate-resource-skill-facts = TRUE populate-resource-state-facts = TRUE populate-resource-state-reason-facts = TRUE populate-sm-resource-session-facts = TRUE populate-sm-resource-state-facts = TRUE populate-sm-resource-state-reason-facts = TRUE populate-virtual-queue-facts = TRUE

Table 48: Genesys Info Mart 7.6 Application Configuration Options (Continued)

Section	Option
schedule	aggregate-start-time = `02:30` etl-end-time = `01:00` etl-start-time = `04:00` load-recent-start-time = `01:30` load-start-time = `02:00` maintain-start-time = `03:00` max-concurrent-extract-jobs = `10` migration-duration-in-hours = `1` migration-start-time = `04:00` populate-intraday-aggregates = `TRUE` run-aggregates = `TRUE` run-load-recent-with-extract-and-transform = `TRUE` run-maintain = `TRUE` run-migration = `FALSE` run-scheduler = `TRUE` etl-frequency = `30`

Java Memory

The Java memory setting for the Genesys Info Mart Server process in the `gim_server.bat` file was set to:

`-Xmx1400m -Xms1400m -Xrs -Xss512k`

Interaction Concentrator Application

Table 49 lists the Interaction Concentrator release 7.6 application options that were used for testing with Sample Configuration 1. It includes only the data-filtering options of the Interaction Concentrator that was dedicated to handling voice details.

Table 49: Interaction Concentrator 7.6 Application Data-Filtering Configuration Options

Section	Option
filter-data	acd-party-history = 1 acd-party-metrics = 1 call-history = 1 call-metrics = 1 gls-metrics = 1 ir-history = 1 observer-party = 1 udata-history-terminated = 1

Scheduling and Configuration of the ETL Jobs

Sample Configurations 1 and 2 used the same scheduling and configuration of the intraday and daily ETL jobs.

Scheduling

The intraday ETL cycle was scheduled with an `etl-frequency` of 30 minutes (1,800 seconds), which targeted 48 cycles per day. Because the call rate remained constant for a full 24 hours a day, no allowances were made for peak or off-peak hours.

The intraday jobs were scheduled to suspend their execution between 1:00 AM and 4:00 AM, to allow for the execution of the daily ETL jobs. The daily jobs were scheduled to start running between 1:30 AM and 2:30 AM.

About the Extraction Jobs

The intraday ETL cycle included data extraction jobs for configuration details, voice interaction details, and voice agent activity details.

- An instance of `Job_ExtractICON` for role `ICON_CFG` was used to extract configuration details:
  - In Sample Configuration 1, the job extracted the configuration details from a highly available (HA) pair of Interaction Concentrators that were dedicated to handling configuration details.
  - In Sample Configuration 2, the job extracted the configuration details from a non-HA Interaction Concentrator.
- An instance of `Job_ExtractICON` for role `ICON_CORE` was used to extract voice details, including voice interaction details and voice agent activity details:
  - In Sample Configuration 1, the job extracted the voice interaction details and voice agent activity details from an HA pair of Interaction Concentrators that were dedicated to handling voice details.
  - In Sample Configuration 2, the job extracted the voice details from a non-HA Interaction Concentrator.

About the Daily ETL Jobs

The daily ETL jobs included the following:

- `Job_LoadGIM`—This job moves facts from the intraday fact tables (where they were loaded by `Job_LoadRecent` as part of the intraday ETL cycle) to the historical fact tables.
- `Job_AggregateGIM`—This job calculates or recalculates the historical aggregates based on:
  - Data that changes since the last load of the historical fact tables.

- New settings for configuration options that control aggregation.

Note: The results for Job_MaintainGIM were not measured because the testing did not generate more than ten days worth of data, so there were no fact or aggregate tables old enough to be eligible for deletion.

How Job Durations Were Calculated

The performance results graphs throughout this chapter show the durations for the various intraday and daily ETL jobs.

To measure the duration of the extraction phase of the intraday ETL cycle, only the longest running job was used—JobExtractICON for role ICON_CORE.

The overall duration of the intraday ETL cycle can be determined using either method:

- By combining the runtimes of the following intraday jobs: Job_ExtractICON for role ICON_CORE + Job_TranformGIM + Job_LoadRecent, as it was done in tests for Sample Configuration 1.
- By subtracting the start time of Job_ExtractICON for role ICON_CFG from the end time of the last intraday job, Job_LoadRecent: (end time of Load Recent) - (start time of Extract ICON CFG), as it was done in tests for Sample Configuration 2.

Factors Affecting Release 7.6 ETL Performance

The primary factors affecting extraction, transformation, and loading (ETL) performance in release 7.6 are as follows:

- The number of daily interactions in your contact center.
- The number of agents in your contact center.
- The complexity of your interaction flows.
- The amount of business data attached to interactions.
- The presence of HA data sources in your deployment.
- The hardware on which the Genesys Info Mart Server is running, primarily the CPU speed and available memory.
- The hardware on which the Staging Area and Info Mart RDBMS server(s) are running, primarily the CPU speed, disk speed, and available memory.
- The tuning of the Staging Area and Info Mart databases.
- The speed of the network connections between components.
- The hardware on which the Interaction Concentrator RDBMS server is running; primarily the CPU speed, disk speed, and available memory.
- The amount of historical data retained in the Interaction Concentrator database.

General Performance Results for Release 7.6

Because similar configurations were used for both Sample Configurations 1 and 2, certain similar trends were noticed in the results.

Daily Spikes in the Results Graphs

In the performance graphs throughout this chapter, notice the daily spikes in the intraday ETL cycle execution time. In these tests, the intraday ETL cycle was suspended for three hours each day to allow the daily ETL jobs to run. These regular spikes show the time it took for the resumed intraday ETL cycle to process the backlog of source data that had accumulated during those three hours on most days.

Gradual Increase in Average Duration from Day to Day

In these performance tests, no provision was made for long-term maintenance of the Genesys Info Mart database. This resulted in a gradual increase in the average duration of the intraday ETL cycles from one day to the next. To keep this increase within an acceptable range for your environment, Genesys recommends that you periodically perform index maintenance on, and purge old data from, the Interaction Database tables containing voice details. For more information about both purging and index maintenance, see "Genesys Info Mart 7.6 Performance Tuning" on page 124.

Release 7.6 ETL Performance for Sample Configuration 1

Genesys Info Mart 7.6 performance testing used the time that it took to execute the ETL jobs as its primary measurement of performance. In these tests, the intraday and daily ETL jobs were reviewed separately. Testing was performed on the Oracle 10g database platform.

Configuration Details

For configuration details specific to Sample Configuration 1, see:

- "Call Volume for Sample Configuration 1" on page 107.
- "Hardware Architecture—Sample Configuration 1" on page 107.
- "Oracle Database Settings" on page 107.

Performance Results

For performance results, see:

- "Volume of Generated Data—Sample Configuration 1" on page 108.
- "Release 7.6 Performance Trends for Oracle 10g" on page 111.

Call Volume for Sample Configuration 1

To test Genesys Info Mart 7.6 ETL performance in Sample Configuration 1, calls were generated at a constant rate of 12 calls per second (cps) for three days. This meant that the total number of calls generated each day was 1,036,800.

Hardware Architecture—Sample Configuration 1

Figure 31 shows the distribution of software components across four servers used in the testing of Sample Configuration 1. Two pairs of HA Interaction Concentrators were used to generate redundant sets of configuration details and voice details, respectively.

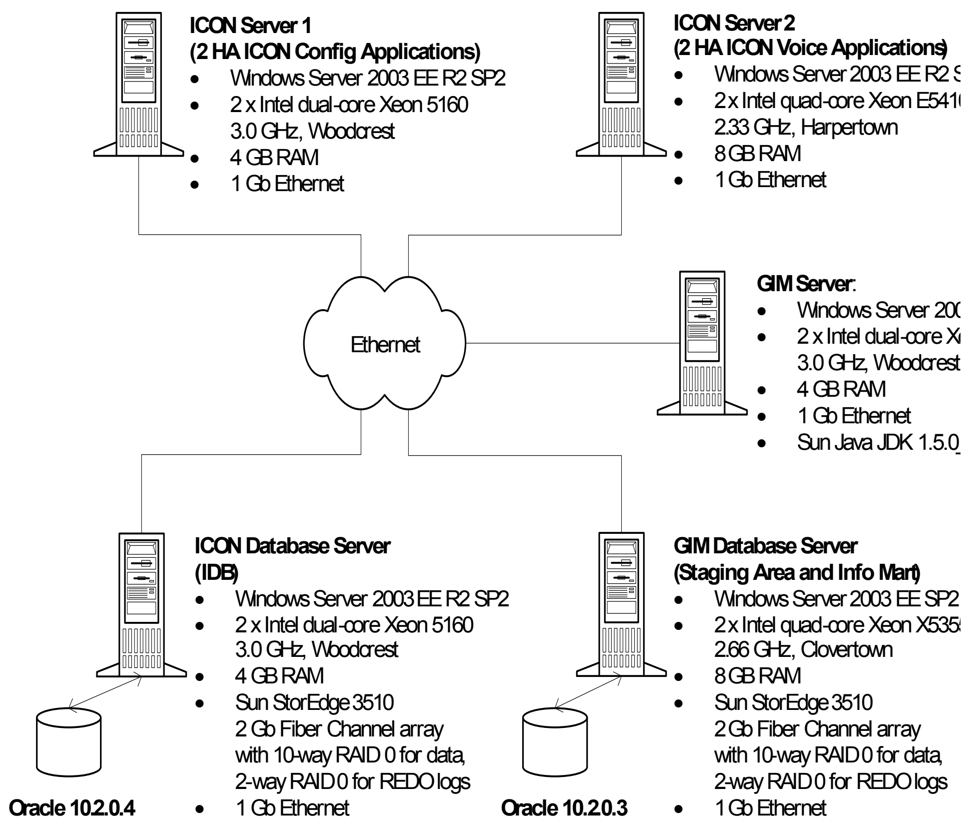

Figure 31: Hardware Architecture—Sample Configuration 1

Oracle Database Settings

Table 50 lists the non-default Oracle settings used for the:

- Database hosting the Interaction Databases (IDBs).
- Database hosting the Staging Area and Info Mart databases.

Table 50: Oracle Database Settings

Database	Settings
IDBs	job_queue_processes = 10 sga_max_size = 2504M shared_pool_reserved_size = 40M undo_retention = 500 db_block_size = 16384 open_cursors = 300 pga_aggregate_target = 800M processes = 150 sga_target = 1600M
Staging Area and Info Mart	job_queue_processes = 10 sga_max_size = 2152M undo_retention = 500 db_block_size = 16384 open_cursors = 600 pga_aggregate_target = 500M processes = 150 sga_target = 2152M

Volume of Generated Data—Sample Configuration 1

Tables 51–53 provide details about the volume of data extracted and loaded during Sample Configuration 1 testing.

Table 51 lists the number of rows extracted from the IDB that stored voice details during a representative 30-minute intraday ETL cycle—at the Sample Configuration 1 rate of 12 cps. The two columns represent:

- An intraday ETL cycle during which there was no backlog of data and the intraday ETL jobs processed the data from the current call generation. With the ETL cycle frequency set to 30 minutes, this column represents roughly 30 minutes worth of Interaction Concentrator (ICON) data.
- An intraday ETL cycle when the intraday ETL jobs had to process the backlog of data (that accumulated during the three-hour pause in the intraday ETL cycle) before resuming the current call data processing. With the `extract-data-time-range-limit` option set to 1 hour, this column represents roughly one hour of ICON data.

Table 51: Data Extracted from IDB Per Intraday ETL Cycle—12 Calls/Second

IDB Table Name	Rows Extracted	
	No Backlog	With Backlog
G_AGENT_STATE_HISTORY	340091	676985
G_AGENT_STATE_RC	43743	80853
G_CALL	65116	127676
G_IR	38891	78272
G_IS_LINK_HISTORY	34800	70044
G_LOGIN_SESSION	13768	16863
G_PARTY	165048	311809
G_PARTY_HISTORY	238384	432206
G_PARTY_STAT	82298	148676
G_ROUTE_RESULT	17547	35317
G_USERDATA_HISTORY	820712	1633499
GX_SESSION_ENDPOINT	27582	34402

Table 52 lists the number of rows loaded to Genesys Info Mart during a representative 30-minute intraday ETL cycle—at the Sample Configuration 1 rate of 12 cps. The two columns represent:

- An intraday ETL cycle during which there was no backlog of data and the intraday ETL jobs processed the data from the current call generation.
- An intraday ETL cycle when the intraday ETL jobs had to process the backlog of data (that accumulated during the three-hour pause in the intraday ETL cycle) before resuming the current call data processing.

Table 52: Data Loaded Per Intraday ETL Cycle—12 Calls/Second

Genesys Info Mart Table Name	Rows Loaded	
	No Backlog	With Backlog
R_INTERACTION_FACT	21491	42349
R_INTERACTION_RESOURCE_FACT	47717	94019

Table 52: Data Loaded Per Intraday ETL Cycle—12 Calls/Second (Continued)

Genesys Info Mart Table Name	Rows Loaded	
	No Backlog	With Backlog
R_INTERACTION_SEGMENT_FACT	147495	290613
R_IXN_RESOURCE_STATE_FACT	174145	343686
R_MEDIATION_SEGMENT_FACT	47569	93730
R_SM_RES_SESSION_FACT	8807	18556
R_SM_RES_STATE_FACT	109113	244680
R_SM_RES_STATE_REASON_FACT	61374	140254
R_VOICE_IXN_FACT_EXT	21491	42349
R_VOICE_RES_FACT_EXT	47717	94019
R_VOICE_SEG_FACT_EXT	147495	290613

Table 53 lists the average total data volume that the daily ETL jobs loaded to Genesys Info Mart—at the Sample Configuration 1 rate of 12 cps.

Table 53: Data Loaded Per Day—12 Calls/Second

Genesys Info Mart Table Name	Rows Loaded
INTERACTION_FACT	1040781
INTERACTION_RESOURCE_FACT	2310590
INTERACTION_SEGMENT_FACT	7141985
IXN_RESOURCE_STATE_FACT	8440883
MEDIATION_SEGMENT_FACT	2303512
SM_RES_SESSION_FACT	420730
SM_RES_STATE_FACT	5335569
SM_RES_STATE_REASON_FACT	3048507
VOICE_IXN_FACT_EXT	1040781
VOICE_RES_FACT_EXT	2310590
VOICE_SEG_FACT_EXT	7141985

Note: Tables 51–53 do not include the numbers of Genesys Info Mart dimension and configuration fact table rows extracted and loaded during each cycle. This is because the expected data volume for new and changed rows is expected to be very low compared to the number of interaction and resource fact table rows.

Release 7.6 Performance Trends for Oracle 10g

This section includes general notes on the release 7.6 testing results, as well as a number of graphs showing the results, in the following subsections:

Notes On Oracle Testing

In reviewing the results of the Oracle performance tests, note how nightly suspension of the intraday ETL jobs affected the results:

- On the morning when there was no backlog of data, the intraday ETL jobs extracted about 30 minutes of ICON data per cycle from the current call generation.
- On the mornings when the backlog of data accumulated during the three-hour pause in the intraday ETL cycle, the intraday ETL jobs had to process the backlog of data before resuming the current call data processing. The backlog ETL cycles processed about an hour of ICON data per cycle. It took up to six ETL cycles (three hours) to process the backlog of data.

Note: To improve the Genesys Info Mart performance, you can adjust the value of the `extract-data-time-range-limit` option, depending on the current backlog and call rate.

Total Job Duration by Cycle—Oracle

Figure 32 depicts the total execution time for all intraday jobs throughout the testing period. The regular spikes are the result of the three-hour pause in the intraday ETL cycle to allow for the daily ETL jobs—the spikes show the subsequent "catch-up" period that the intraday ETL cycle required to process the backlog of data.

Notice the marked duration increase around cycle 28, which resulted from slow performance of interval-based aggregation queries. Refer to the *Genesys Info Mart 7.6 Release Notes* for recommendations on locking database statistics for the intraday fact tables (on Oracle 10g), or re-running `ANALYZE` type statistics (on Oracle 9i).

Also note that it took an average of six cycles—up to three hours—to process each three-hour backlog of data, implying that the system can sustain significantly larger spikes of activity.

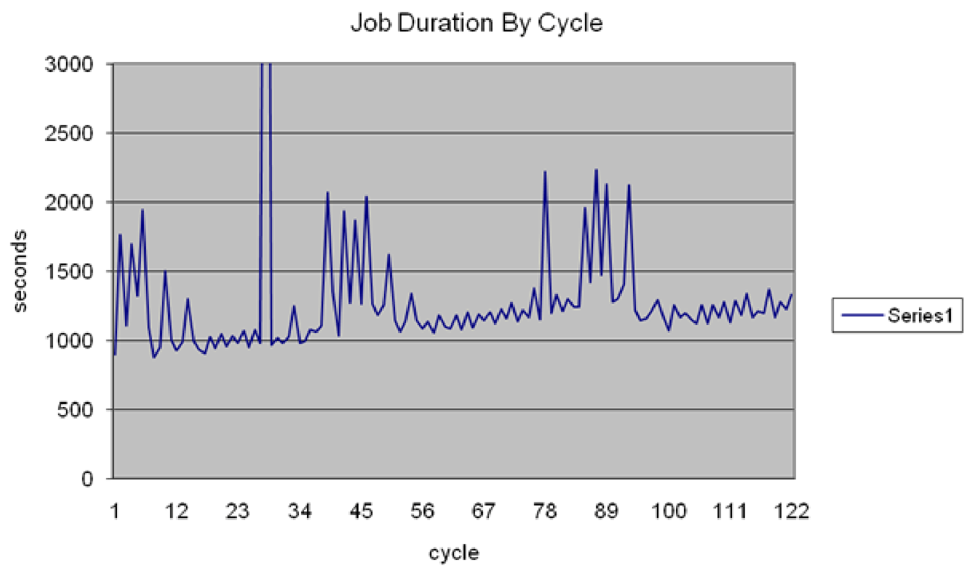

Figure 32: Total Job Duration by Cycle—Oracle

Average Intraday ETL Cycle Duration—Oracle

Figure 33 depicts the daily average duration of the intraday ETL cycles. You can see evidence of a gradual increase in the duration from day to day. This increase is caused by the accumulation of historical data in the Interaction Database, specifically the voice details extracted by `Job_ExtractICON` for role `ICON_CORE`. To keep the duration within an acceptable range for your environment, Genesys recommends that you periodically perform index maintenance on, and purge old data from, the Interaction Database tables that contain voice details.

Also note that for the configured 30-minute intraday ETL cycle, fewer than 23 minutes were used, indicating that the environment could sustain higher call rates.

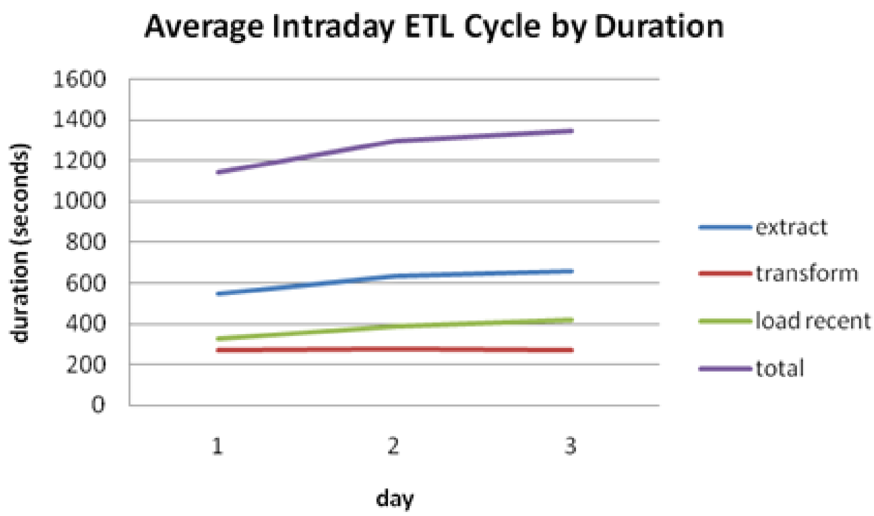

Figure 33: Average Intraday ETL Cycle Duration—Oracle

Average Intraday Extraction Cycle Duration—Oracle

Figure 34 depicts the performance trends for the individual extraction jobs, showing daily averages for each. The three-day test did not indicate any increase in the duration of `Job_ExtractICON` for the role of `ICON_CFG`.

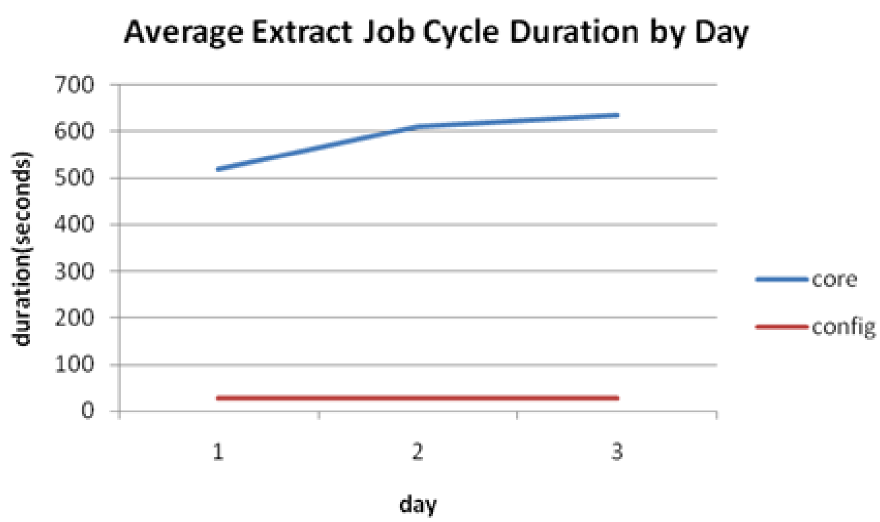

Figure 34: Average Intraday Extraction Cycle Duration—Oracle

Job_ExtractICON for Role ICON_CORE (All Call Details)

Figures 35 shows that `Job_ExtractICON` (as the longest-running extraction job) for the role of `ICON_CORE` determined the duration of the extraction phase of the intraday ETL cycle. You can see a slight, gradual increase in the duration of this job over time, despite the constant call volume. This increase, which would be more noticeable over a longer time period, is directly related to the overall size of the Interaction Database.

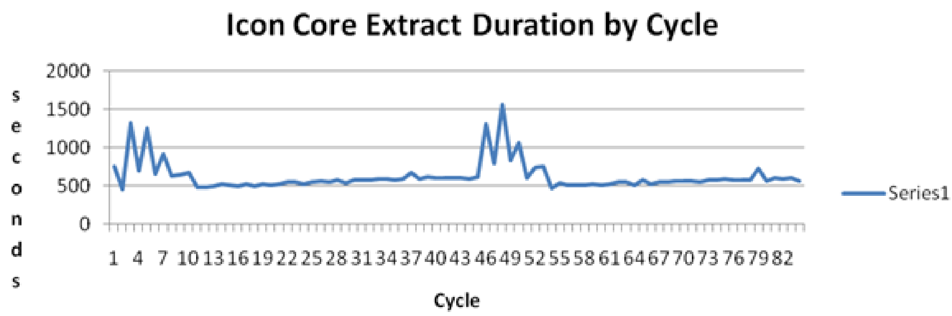

Figure 35: Job_ExtractICON for Role ICON_CORE—Oracle

Note: Performance testing did not include any purging or index maintenance on the Interaction Database. Periodic purging or index maintenance of the Interaction Database will prevent the intraday ETL cycle duration from increasing to an unacceptable level.

Transform Duration

Figure 36 shows the duration for `JOB_TransformGIM`. You can see that the transform duration increases during the day, but it remains constant between day two and day three.

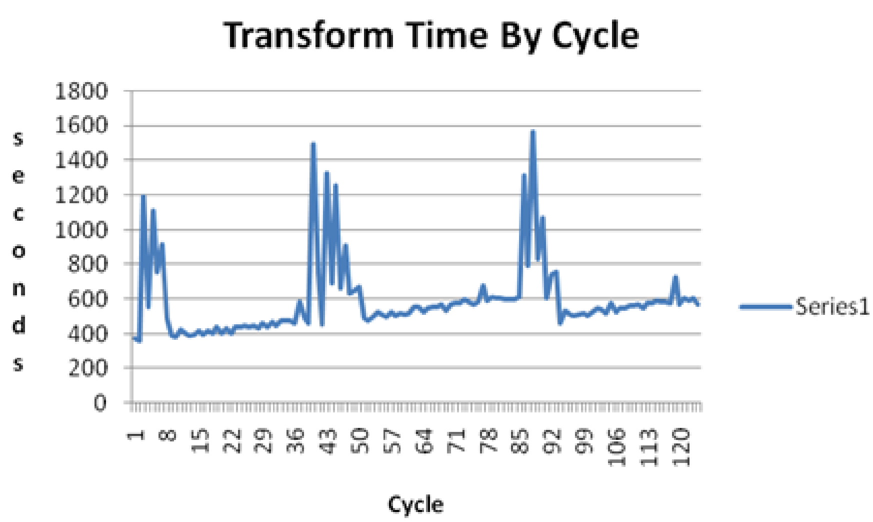

Figure 36: Job_TransformGIM—Oracle

Load Recent Duration

Figure 37 shows the duration of the Job_LoadRecent phase of the intraday ETL cycle throughout the testing period. The reasons for the large spikes around cycles 78 and 99 are unknown.

Despite the constant call rate, a slight linear increase in the job duration is visible over time. By the end of the testing period, the average duration of Job_LoadRecent had not increased beyond 450 seconds. The increase in the job duration results from growth of the Info Mart database size. To mitigate the growth, consider using Job_MaintainGIM to purge the Info Mart database as necessary.

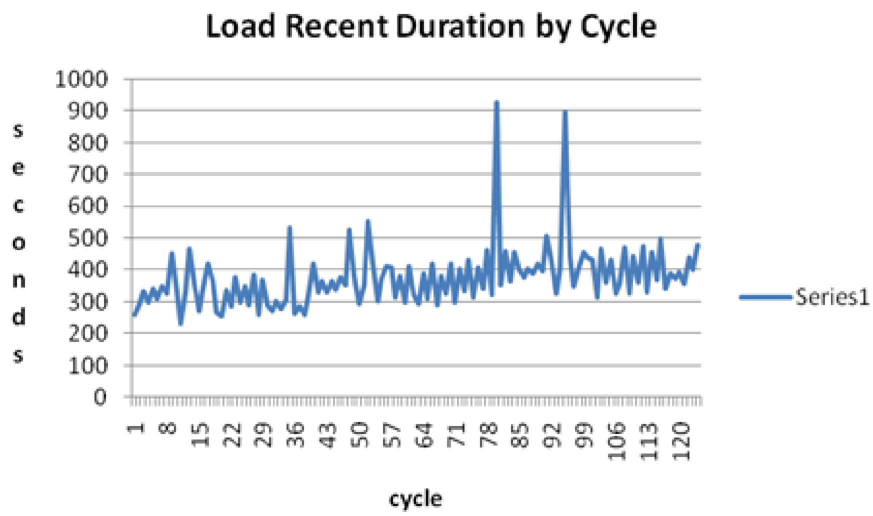

Figure 37: Load Recent Duration by Cycle—Oracle

Daily Job Executions

No optimization for target database growth was performed before or during testing.

Genesys recommends that you consult a database administrator (DBA) knowledgeable in data warehouses to analyze the projected growth of your Genesys Info Mart database—according to a size estimation spreadsheet—and that you plan the storage, partitioning, and any other optimization measures accordingly.

Figure 38 shows the performance result information for Job_LoadGIM. This job's behavior is directly affected by the amount of records being loaded into the Info Mart database. The job's duration also depends on the amount of data accumulated in the Info Mart database. No increase in the job duration was shown in the test because the load remained constant from day to day and because the fact tables were not pre-loaded with a large amount of data before the test was started. In a contact center environment, you may see an increase in duration of `Job_LoadGIM` over time as well as during the days with high call volume.

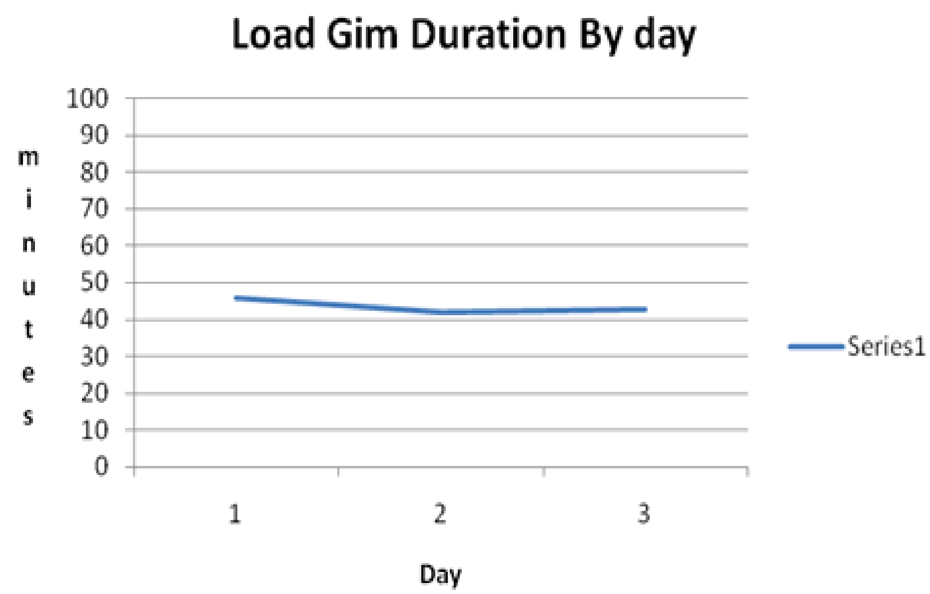

Figure 38: Load Genesys Info Mart

Figure 39 shows the performance results for Aggregate Genesys Info Mart Duration. The job duration did not change materially from day to day during the three-day test. The test did not demonstrate any significant increase in the job duration because of the short time frame. In a longer term, you should expect an increase in the aggregate duration because of the growth in size of the Info Mart database.

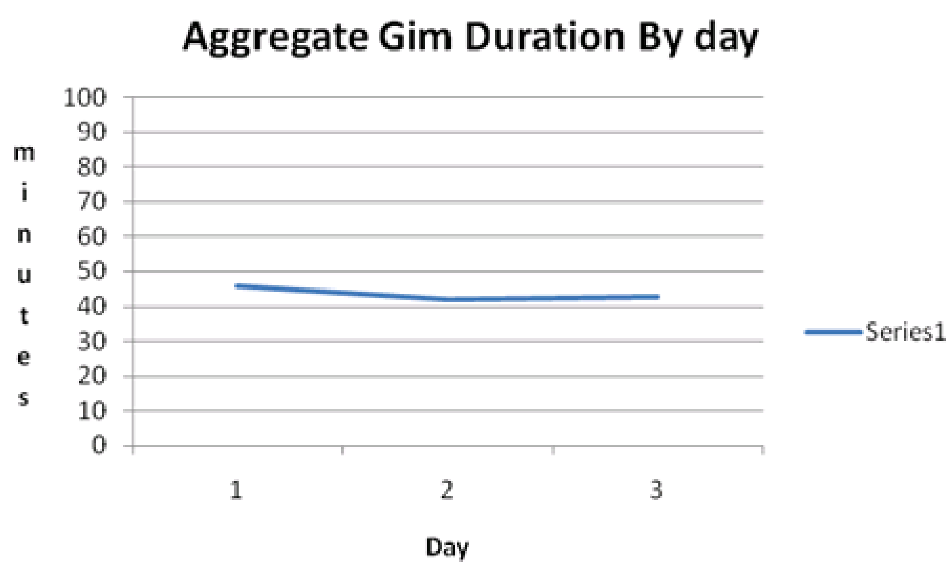

Figure 39: Aggregate Genesys Info Mart Duration

Data Consumption

Table 54 shows the amount of hard disk space that Genesys Info Mart 7.6 ETL performance testing consumed. (No purging was performed on any of the test databases.)

Table 54: Hard Disk Space Consumption for Release 7.6 ETL Performance Testing

Database	Space Consumption on Oracle
IDB	3.1 million recorded calls used 676 GB of database space (including HA storage of ICON Voice details and Configuration details), so daily consumption was about 225 GB/day. The space was used in an IDB voice schema as follows: • Tables 69 GB • Indexes 46 GB • Total space for tables and indexes 115 GB The space was used in an IDB configuration schema as follows: • Tables 156 GB • Indexes 67 GB • Total space for tables and indexes 223 GB
Genesys Info Mart	3.1 million recorded calls used 54.9 GB of database space, so daily consumption was about 18.3 GB/day. The space was used in the Staging Area schema as follows: • Tables ~1 GB • Indexes ~1.3 GB • Total space for tables and indexes ~2.3 GB The space was used in the Info Mart schema as follows: • Tables ~33.7 GB • Indexes 18.9 GB • Total space for tables and indexes ~52.6 GB

Release 7.6 ETL Performance for Sample Configuration 2

Genesys Info Mart 7.6 ETL performance testing for Sample Configuration 2 was conducted using the same call flows and Genesys Info Mart configuration that were used for the Sample Configuration 1 tests, but at a much lower call volume of 4.5 cps, without ICON data filtering, and without HA. Testing was performed for the Microsoft SQL Server 2005 database platform.

Configuration Details

For more configuration details specific to Sample Configuration 2, see:

- “Call Volume for Sample Configuration 2” on page 118.
- “Hardware Architecture—Sample Configuration 2” on page 119.

Performance Results

For performance results on Microsoft SQL Server 2005, see:

- “Release 7.6 Performance Trends for Microsoft SQL Server 2005” on page 119.
- “Data Consumption” on page 123.

Call Volume for Sample Configuration 2

To test Genesys Info Mart 7.6 ETL performance in Sample Configuration 2, calls were generated at a constant rate of 4.5 cps for 7 days. This meant that:

- At a rate of 4.5 cps, 8,100 calls were generated during every 30-minute cycle.
- The total number of call generated each day was 388,800.

Hardware Architecture—Sample Configuration 2

Testing on the Sample Configuration 2 used a distribution of software components across several servers, as shown in Figure 40.

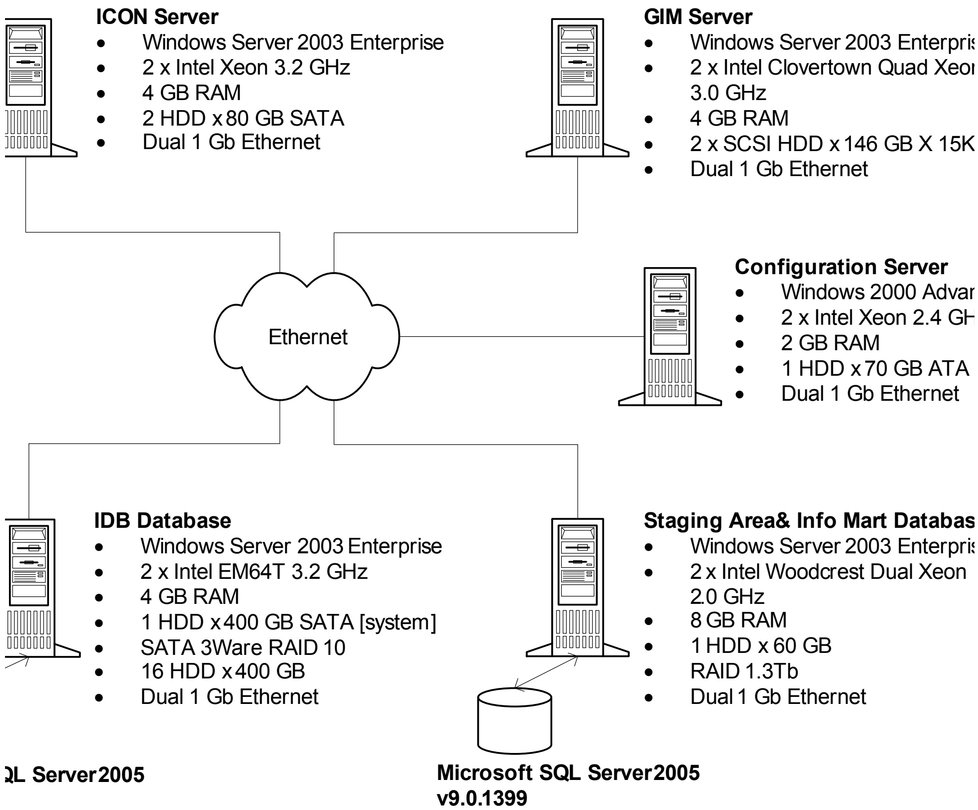

Figure 40: Hardware Architecture—Sample Configuration 2

Microsoft SQL Database Settings

Default Microsoft SQL Server 2005 settings were used for the:

- Database hosting the IDB.
- Database hosting the Staging Area and Info Mart databases.

Release 7.6 Performance Trends for Microsoft SQL Server 2005

This section includes general notes on the release 7.6 testing results, as well as a number of graphs showing the results, in the following subsections:

- "Notes on Microsoft SQL Server Testing" on page 120
- "Total Job Duration By Date—Microsoft SQL Server" on page 120
- "Extract Duration By Date—Microsoft SQL Server" on page 120

- “Transform Genesys Info Mart Duration By Date—Microsoft SQL Server” on page 121
- “Daily Jobs—Microsoft SQL Server” on page 122

Notes on Microsoft SQL Server Testing

- In Figures 42 through 44, the time axis uses the logarithmic scale for better visual presentation.
- In Figures 41 through 43, the large peaks in runtime are “catch-up” periods for the intraday jobs after the three-hour nightly window for the running of the daily jobs.

Total Job Duration By Date—Microsoft SQL Server

Figure 41 shows that the overall runtime of the combined intraday jobs has increased throughout the test. The total job duration in this test was calculated as follows: `(end time of Load Recent) - (start time of Extract ICON CFG)`.

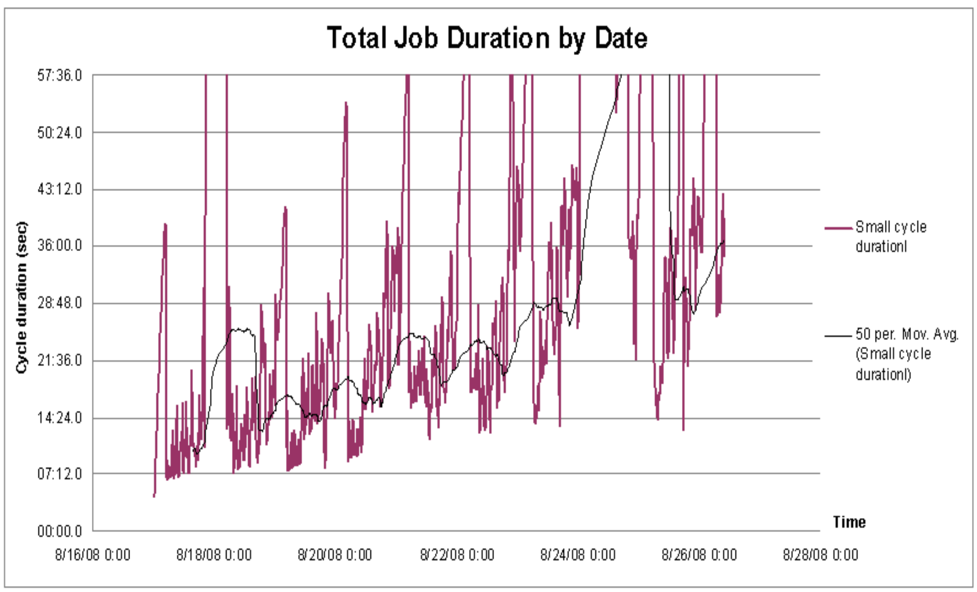

Figure 41: Total Job Duration by Date

Extract Duration By Date—Microsoft SQL Server

Figure 42 illustrates no increase in the duration of `Job_ExtractICON` over time, under the constant call volume. On certain days, the job took approximately ten times longer to complete the data extraction. Index maintenance on certain key IDB tables, which was applied when extract performance slowed, accounts for the dramatic restoration of performance throughout the test.

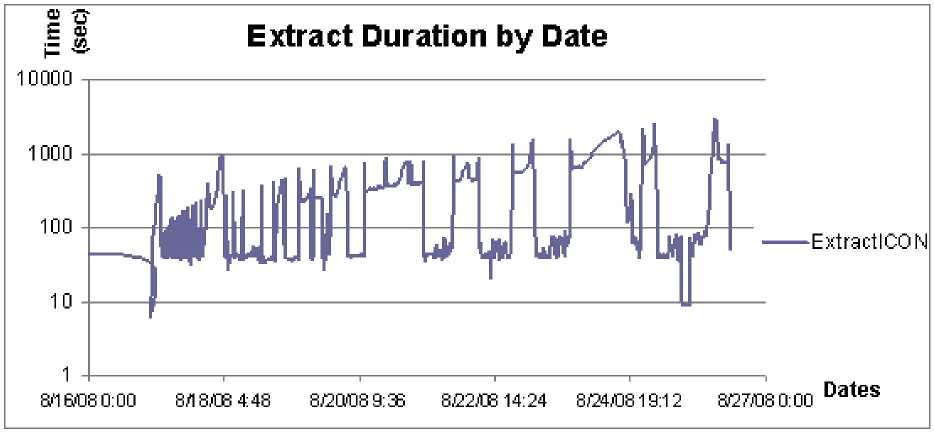

Figure 42: Extract Duration by Date

Transform Genesys Info Mart Duration By Date—Microsoft SQL Server

Figure 43 shows that there was no significant increase in the duration of `Job_TransformGIM`. The spikes appear in the graph because of nightly backlog processing.

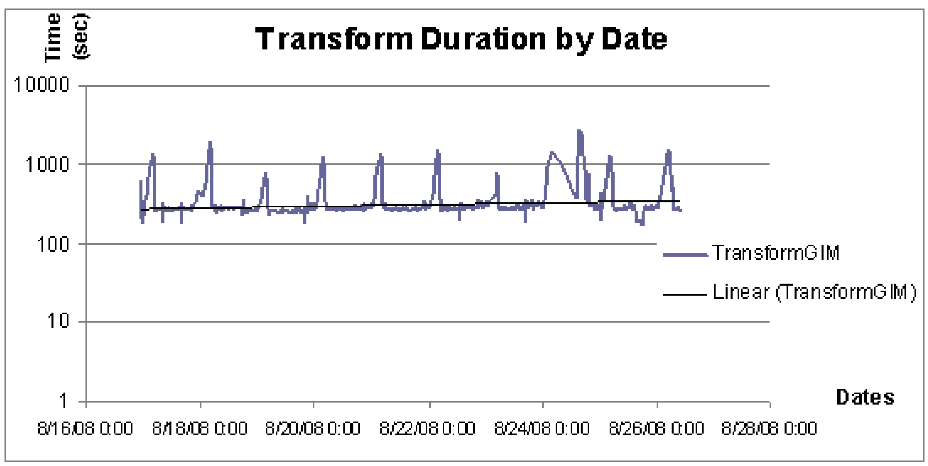

Figure 43: Transform Duration by Date

Load Recent Duration By Date—Microsoft SQL Server

Figure 44 shows a linear increase in `Job_LoadRecent` duration over the testing period. The increase results from the growth of the Info Mart database. The large spike on 8/25/08 could be attributed to some environmental factors.

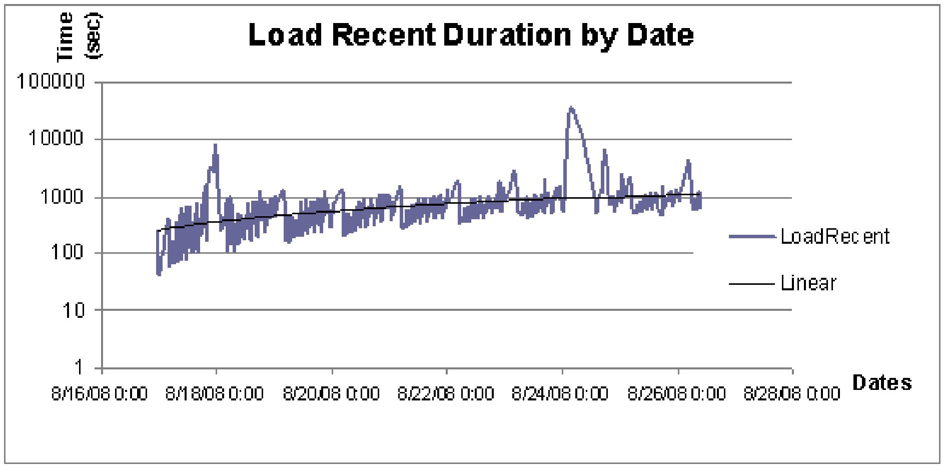

Figure 44: Load Recent Duration by Date

Daily Jobs—Microsoft SQL Server

Figure 45 illustrates that the execution time of `Job_LoadGIM` did not show significant changes during the test—compare the job durations on day 3 and day 7.

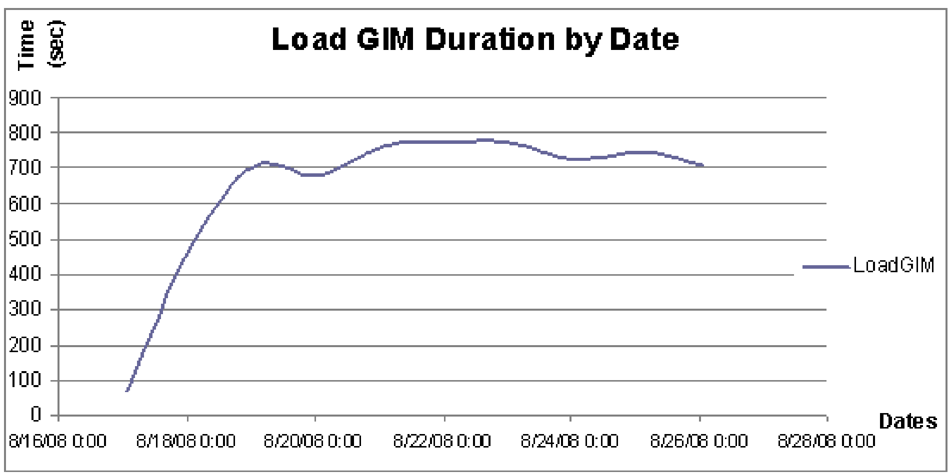

Figure 45: Load GIM Duration by Date

Figure 46 illustrates that `Job_AggregateGIM` increased in duration linearly as the amount of aggregated data increased as the test progressed.

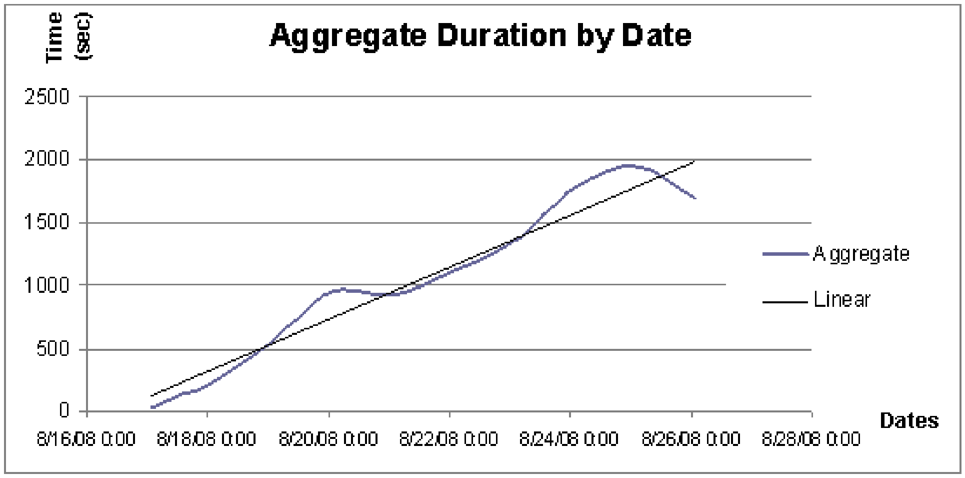

Figure 46: Aggregate GIM Duration by Date

Data Consumption

Table 55 shows the amount of hard disk space that Genesys Info Mart 7.6 ETL performance testing consumed.

Table 55: Hard Disk Space Consumption for Release 7.6 ETL Performance Testing

	On Microsoft SQL Server
ICON	Total database space consumption was about 64.5 GB for IDB.
Genesys Info Mart	Total database space consumption was about 31.3 GB for the Info Mart database.

Genesys Info Mart 7.6 Performance Tuning

Genesys recommends that you tune Genesys Info Mart 7.6 performance according to the guidelines in Table 56.

Table 56: Release 7.6 Performance Tuning Guidelines

Guideline	Description
Purge the IDB.	To maintain ETL job performance over time, Genesys recommends that you periodically purge the Interaction Concentrator database (IDB). The frequency required depends on the hardware and data volume. Current Call Concentrator (CCON) customers have procedures for manually purging the CCON database. Because ICON generates several times more data for the same call volume, it needs to be purged more frequently than CCON. For more information about procedures to purge the Interaction Concentrator database, see the *Genesys Info Mart 7.6 Deployment Guide.*
Defragment indexes.	If you find the performance of `Job_ExtractICON` degrading over time, check for index fragmentation on the IDB tables. In particular, look at the following five indexes: • `IDX_G_IR_ROOTIRID` on table `G_IR` • `IDX_G_CALL_ROOTIRID` on table `G_CALL` • `GIM_IDX_PAR_CALLID` on table `G_PARTY` • `IDX_G_PARTY_H_PID` on table `G_PARTY_HISTORY` • `GIM_IDX_UH_CALLID` on table `G_USERDATA_HISTORY` Consider defragmenting these indexes by using either of these methods: • For Oracle, use `COALESCE` • For Microsoft SQL Server, rebuild the index with a different fill factor (20–50 percent).
Filter ICON data.	Use the ICON data filtering options listed in Table 49 on page 103 because, with these options, ICON writes the data only to those tables and row types that Genesys Info Mart actually uses. This will reduce the amount of data that is stored in IDB and improve the performance of Info Mart extraction, in addition to ICON performance improvements.

Table 56: Release 7.6 Performance Tuning Guidelines (Continued)

Guideline	Description
Provision for large database volumes.	If you want to retain significant volumes in the database, you must make provisions at the database level to manage the Genesys Info Mart database. A volume of one million calls per day quickly leads to unacceptable performance from both the Load and Aggregate jobs, unless proper measures are put in place. Genesys recommends that you consult a Database Administrator (DBA) knowledgeable in data warehouses to analyze the projected growth of your Genesys Info Mart database—according to a size estimation spreadsheet—and that you plan the storage, partitioning, and any other optimization measures accordingly.
Use caution when adding indexes.	Genesys Info Mart load time is also affected by the number of indexes on the large fact tables. Therefore, use caution when adding indexes. In the performance tests, one table in particular took much longer than the others to load: `INTERACTION_SEGMENT_FACT`, which had eight indexes.
Partition the `Interaction_Segment_Fact` table by date.	Although the last two days of execution time for `Job_LoadGIM` were flat for the Oracle test, we can expect the general trend of increasing execution time to continue. The factor most influencing this time was the rebuilding of indexes on `INTERACTION_SEGMENT_FACT`. Because all the indexes were global, they all had to be rebuilt for the entire table with every run of `Job_LoadGIM`. Partitioning this table by `STD_TENANT_DATE_KEY` and making the indexes local would significantly flatten this performance trend. Each index would be rebuilt only for the data contained in that partition, rather than for the entire table.
Keep databases on separate servers.	These performance tests used a single database instance for both the Staging Area database and the Genesys Info Mart database. This topology did not significantly impact the test results, because only `Job_LoadRecent` would use both databases at the same time. In an actual customer deployment, you should locate the Staging Area database and the Genesys Info Mart database on separate database servers, so that the user reporting queries do not impact the intraday ETL jobs.
Optimize the storage subsystem.	Genesys recommends that you take special note of the storage subsystem of the database servers used in testing, because earlier tests indicated that sub-optimal database storage can easily become a critical performance bottleneck.

Table 56: Release 7.6 Performance Tuning Guidelines (Continued)

Guideline	Description
Use asynchronous merge.	The Genesys Info Mart application option `ir-merge-interval` specifies the time interval, in minutes, that the Genesys Info Mart is to periodically run the ICON IR Merge stored procedure. The stored procedure merges completed calls between switches monitored by the same ICON application. If you set the value to `0`, `Job_ExtractICON` calls this IR Merge procedure, and the extract queries are not executed until the IR Merge procedure completes. The test results included here used an `ir-merge-interval` of `0`. If you set the value greater than `0`, then the Genesys Info Mart Server calls the IR Merge procedure asynchronously for all the ETL jobs. This decreases the execution time of `Job_ExtractICON` and increases the though put of the ETL process. The recommended value of `ir-merge-interval` is `5`.
Allow sufficient database resources.	Be generous with the amount of RAM that you allocate to the database server buffer caches. Work with your DBA to tune the IDB, Staging Area and Genesys Info Mart database systems with production data volumes. These systems have different operational characteristics and need to be tuned differently.
Provide adequate space for transaction logs.	By default, the daily `Job_LoadGIM` moves data from an intraday table to the matching historical table using transactions that are sized at one million rows each. Depending on database platform and daily data volume, the transaction logs for such a large amount of data can consume gigabytes of disk space. Make sure that your environment allows for adequate space to handle these transaction logs. Use the `load-transaction-size` option that is configured at the Genesys Info Mart application level to adjust the transaction size.

Table 56: Release 7.6 Performance Tuning Guidelines (Continued)

Guideline	Description
Limit extraction data.	Use the Genesys Info Mart application options `limit-extract-data` and `extract-data-time-range-limit` to set the maximum amount of data that Genesys Info Mart will extract during a single ETL cycle. For more information about this data throttling process, see the Genesys Info Mart Deployment Guide for details.
Tune Genesys Info Mart Server memory.	For optimal performance, use the largest transformation buffer that the RAM on your machine can support. This buffer holds the objects that `Job_TransformGIM` will process on a single pass. To tune Genesys Info Mart Server memory: 1. Set the Java heap size parameter in the `gim_server` startup script to either 1500 MB for a 32-bit operating system or 2000 MB for a 64-bit operating system, provided that these amounts of RAM are available on the machine. 2. Set the Genesys Info Mart application option `transformation-buffer-size` to `5`. 3. Restart the server and begin running jobs. 4. If you receive a `Java Out of Memory` exception, set the option to `4`. 5. Repeat step 4 (decreasing the value of the option by one) until you no longer get a memory exception.

Genesys Info Mart 7.6 Database Size Estimation

Genesys Info Mart 7.6 reads data from the Interaction Concentrator, Stat Server, and GVP VAR databases, and writes data to its Staging Area and Genesys Info Mart databases.

The information in this section helps you estimate the size of the Staging Area and Info Mart data, depending on the size of your contact center and the length of time you want to retain data in the Genesys Info Mart.

The variables that have the greatest effect on the size of the Staging Area and Genesys Info Mart databases are the:

- Number of daily interactions in your contact center.
- Number of agents in your contact center.
- Complexity of your interaction flows.
- Amount of business data attached to interactions.

Notes:

- The Interaction Concentrator, Stat Server, and GVP VAR data size estimates are outside the scope of this guide.
- The data sizes are only estimates. Be sure to factor in extra space to accommodate variations in average data lengths.
- The Staging Area size estimate allows for up to three days of extracted data in case a network outage, database outage, or some other error temporarily prevents Genesys Info Mart from transforming or loading data.
- The Info Mart size estimate includes default indexes. Be sure to provide for additional storage for any indexes you add to the Info Mart to enhance query performance.

Genesys Info Mart provides the Genesys Info Mart Database Size Estimator, a spreadsheet in Microsoft Excel, to help you estimate the data size of the Staging Area and Info Mart databases.

- The spreadsheet requires that you fill in the basic information about the resources and daily interaction volumes for your contact center.
- You should also include the number of days you intend to store data in the Info Mart database.
- The spreadsheet then calculates the estimated initial size and daily delta size for the Staging Area and Info Mart databases.

The Genesys Info Mart 7.6 Database Size Estimator spreadsheet (`GIMSize_760.xls`) is available on the:

- Genesys Technical Support Website at `http://genesyslab.com/support`.
- Genesys Documentation Library DVD, which you can order by e-mail from Genesys Order Management at `orderman@genesyslab.com.`

Release 7.5 Guidelines

The following sections provide hardware sizing guidelines for environments with Genesys Info Mart 7.5:

- "Hardware Architectures in Release 7.5" on page 129, provides architecture samples for Genesys Info Mart release 7.5, for small and medium-size contact centers.
- "Genesys Info Mart 7.5 ETL Runtime Performance" on page 134, lists the factors that affect Genesys Info Mart 7.5 performance. It also provides general configuration details that apply to both sample configurations.

- "Release 7.5 ETL Performance for Sample Configuration 1" on page 139, lists specific performance measurements for Genesys Info Mart 7.5 as deployed in a larger contact center environment on both Oracle and Microsoft SQL Server database platforms.
- "Release 7.5 ETL Performance for Sample Configuration 2" on page 156, lists specific performance measurements for Genesys Info Mart 7.5 as deployed in a smaller contact center environment on a Microsoft SQL Server database platform.
- "Genesys Info Mart 7.5 Performance Tuning" on page 164, provides a series of guidelines to help you tune your database environment for superior performance.
- "Genesys Info Mart 7.5 Database Size Estimation" on page 167, shows you how to estimate the size of the Staging Area and the Genesys Info Mart databases.

The information is organized according to the size of the contact center as a function of the daily number of interactions and the number of agents. Table 57 provides the definition of the contact center size, as applicable to release 7.5 of Genesys Info Mart.

Table 57: Contact Center Size for Release 7.5

Size Category	Number of Agents	Daily Number of Interactions	Maximum Calls Per Second
Small	Fewer than 150	Fewer than 50,000	1.5
Medium	Fewer than 1,000	Fewer than 300,000	12
Large	More than 1,000	More than 300,000	> 12

Note: Large contact centers require special planning because of their variability and complexity. Contact Genesys Technical Support if you want to plan a Genesys Info Mart deployment for a large contact center.

Hardware Architectures in Release 7.5

The examples in this section are organized by the size of your contact center and your choice of operating system.

The hardware architecture you select for your Genesys Info Mart 7.5 deployment depends primarily on:

- The size of your contact center in relation to the number of daily interactions and the number of agents.

- The time available for Genesys Info Mart to extract, transform, and load data.
- Your choice of operating system.

To minimize the time it takes for Genesys Info Mart 7.5 to extract, transform, and load data, Genesys strongly recommends that you run the Genesys Info Mart Server and the Staging Area relational database management system (RDBMS) server on different machines.

Notes:

- Hardware recommendations for the Interaction Concentrator, Stat Server RDBMS, and Genesys Info Mart RDBMS servers are outside the scope of this guide.
- The type of hardware you deploy depends on whether this database is used directly for reporting (and the number of concurrent users), or as a secondary staging area to hold the data until it can be uploaded to a data warehouse.

Small Contact Center—Windows—Release 7.5

Figure 47 depicts a sample hardware architecture of Genesys Info Mart release 7.5 for a small-size contact center on a Windows operating system.

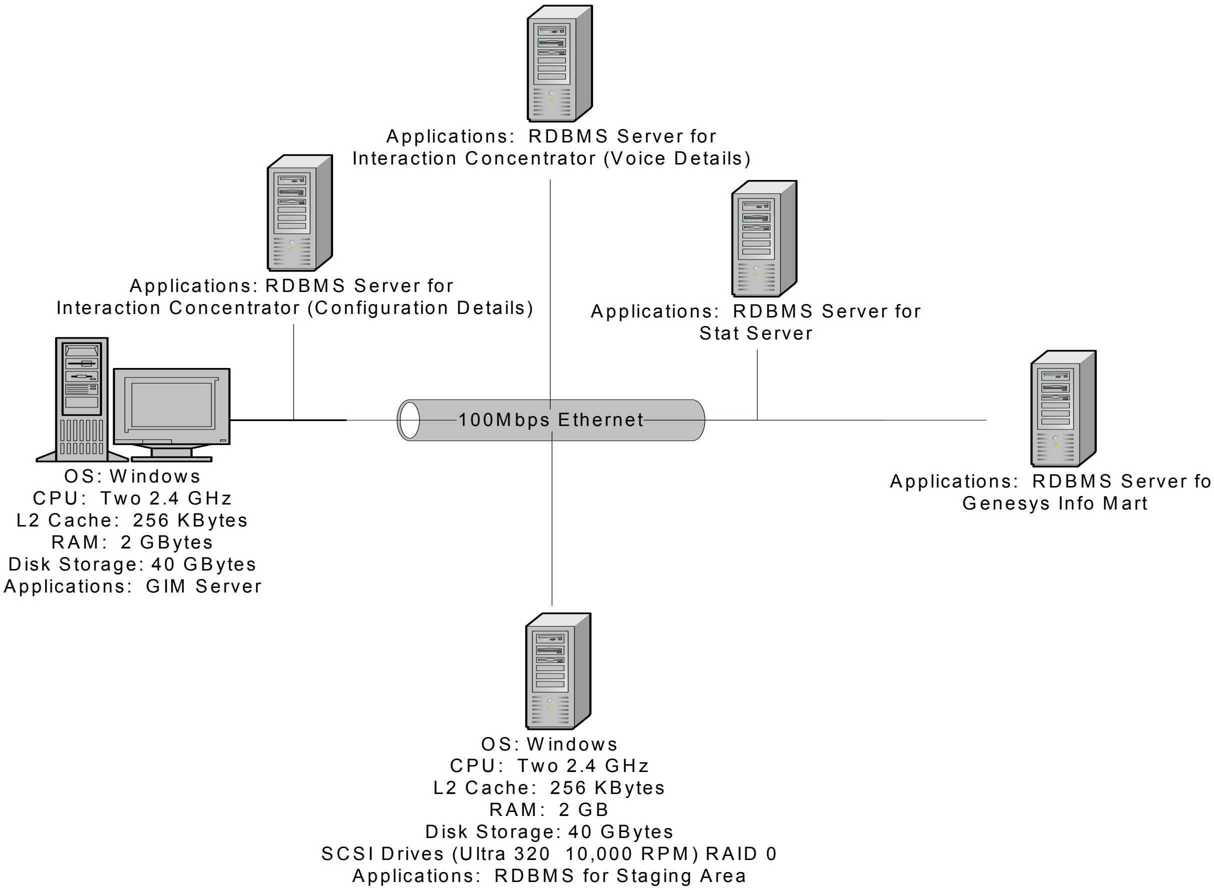

Figure 47: Small-Size Contact Center—Windows—Release 7.5

Note: Genesys strongly recommends that you use a Staging Area RDBMS disk subsystem that contains multiple SCSI disk drives. Storing RDBMS logs, indexes, and table data on different disk drives improves extraction, transformation, and loading (ETL) performance by reducing I/O contention.

Medium Contact Center—Windows—Release 7.5

Figure 48 depicts a sample hardware architecture of Genesys Info Mart release 7.5 for a medium-size contact center on a Windows operating system.

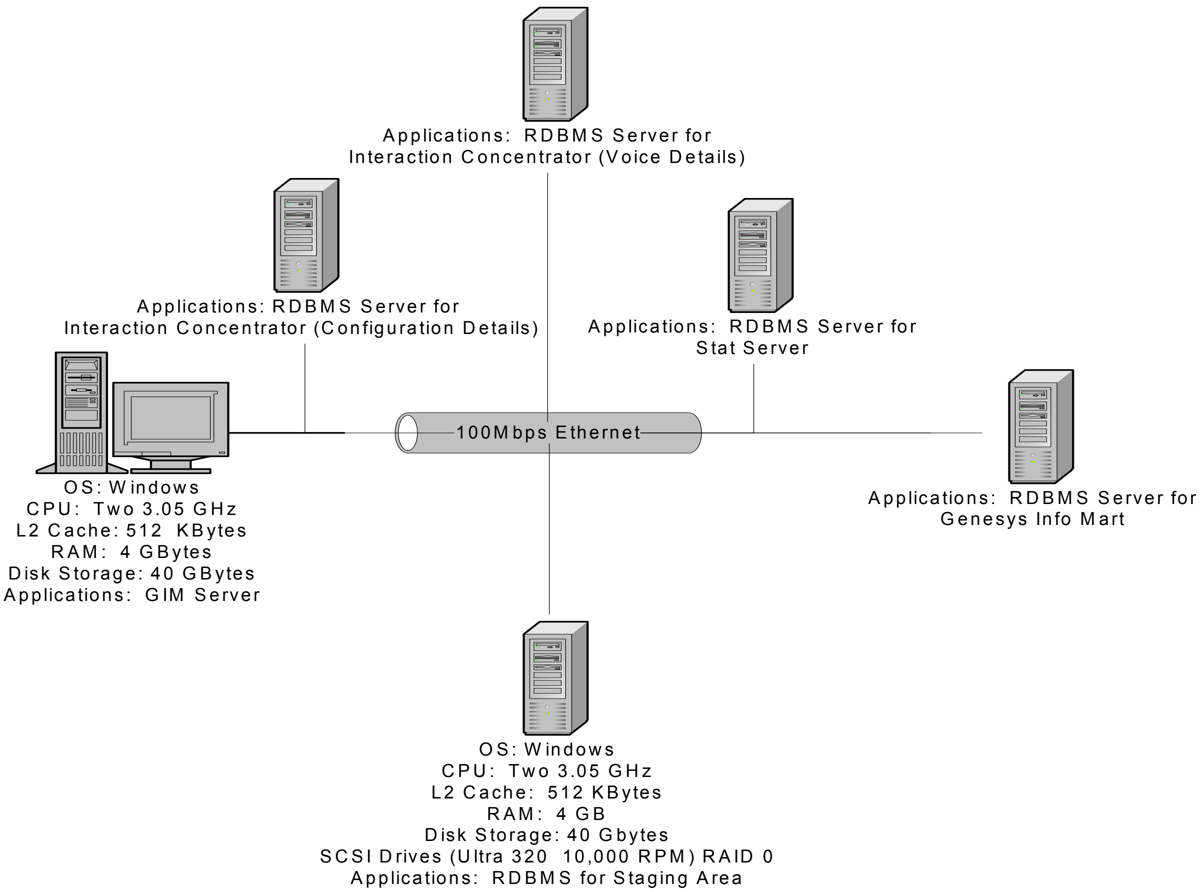

Figure 48: Medium-Size Contact Center—Windows—Release 7.5

Note: Genesys strongly recommends that you use a Staging Area RDBMS disk subsystem that contains multiple SCSI disk drives. Storing RDBMS logs, indexes, and table data on different disk drives improves extraction, transformation, and loading (ETL) performance by reducing I/O contention.

Small Contact Center—UNIX (Solaris)—Release 7.5

Figure 49 depicts a sample hardware architecture of Genesys Info Mart release 7.5 for a small-size contact center on a UNIX (Solaris) operating system.

To obtain the best Genesys Info Mart Server performance on UNIX (Solaris), Genesys recommends selecting a machine that has the most powerful processors, rather than one with multiple slower processors. For example, a machine with two 990-MHz processors performs better than one with four 440-MHz processors.

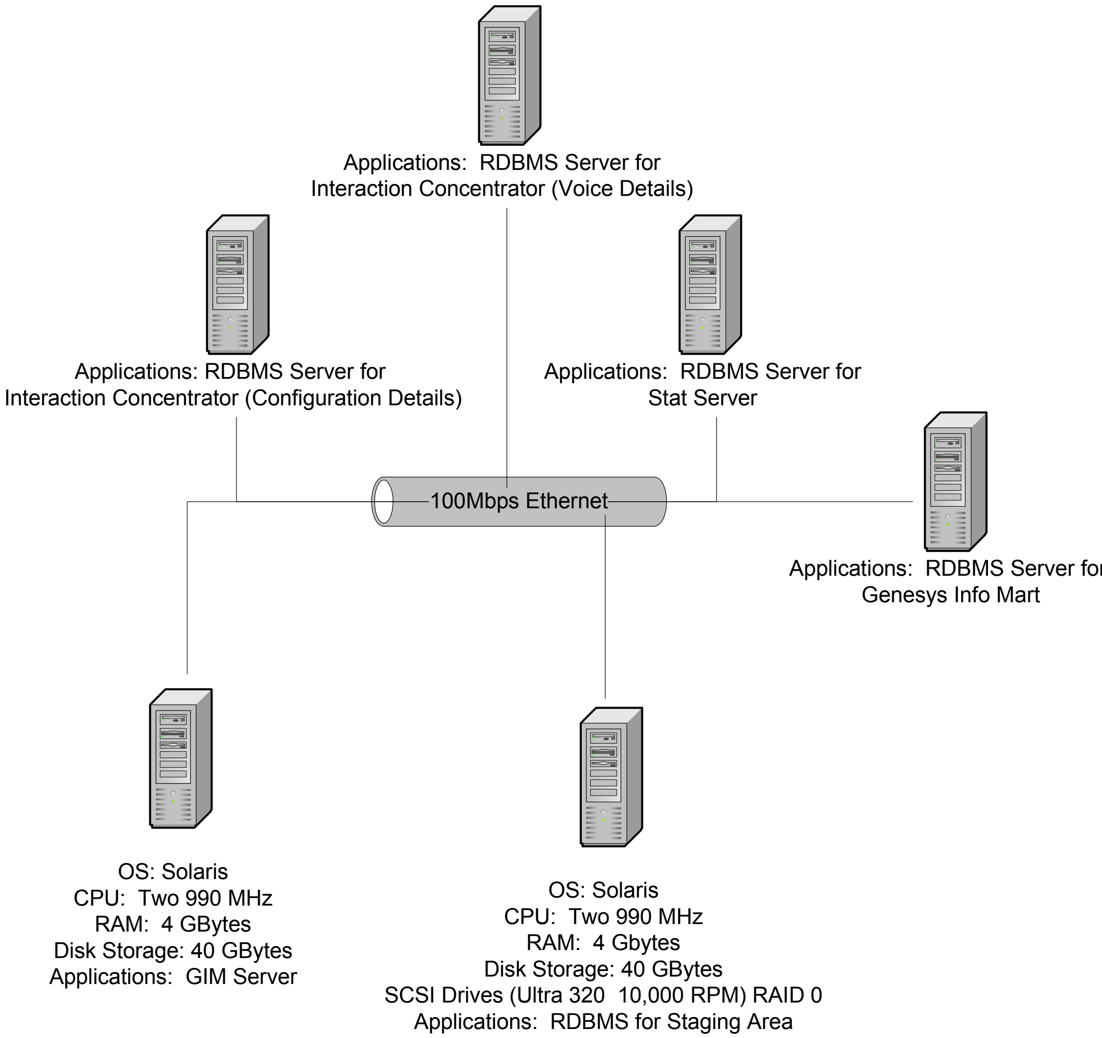

Figure 49: Small-Size Contact Center—Solaris—Release 7.5

Note: Genesys strongly recommends that you use a Staging Area RDBMS disk subsystem that contains multiple SCSI disk drives. Storing RDBMS logs, indexes, and table data on different disk drives improves extraction, transformation, and loading (ETL) performance by reducing I/O contention.

Genesys Info Mart 7.5 ETL Runtime Performance

This section provides information on the following topics:

About Genesys Info Mart 7.5 Performance Testing

To test Genesys Info Mart 7.5 ETL performance, the ETL cycle was run on actively populated source data over a number of days for the following test contact center sizes:

- Sample Configuration 1 handling up to 1,000,000 calls a day
- Sample Configuration 2 handling up to 260,000 calls a day

Testing was performed on two database platforms: Oracle 10g and Microsoft SQL Server 2005 for the Sample Configuration 1, and Microsoft SQL Server 2005 only for Sample Configuration 2.

Genesys Info Mart 7.5 was tested with Interaction Concentrator 7.5.

General Configuration Details

Testing for both Sample Configurations 1 and 2 used identical call flows and Genesys Info Mart configuration, but different call volumes. The call rate was constant for a full 24 hours each day, with no scheduling for peak or off-peak hours.

For further configuration details that apply generally to all of the release 7.5 tests, see the following sections:

- "Release 7.5 Performance Testing Configuration" on page 135
- "Scheduling and Configuration of the ETL Jobs" on page 136

Performance Results

For some performance trends that applied generally for all of the release 7.5 tests, see "General Performance Results for Release 7.5" on page 138.

You can find more specific performance results—including a series of graphs that show the duration of the various ETL jobs—in the following sections:

- "Release 7.5 ETL Performance for Sample Configuration 1" on page 139

- “Release 7.5 ETL Performance for Sample Configuration 2” on page 156

Release 7.5 Performance Testing Configuration

Table 58 provides configuration details that pertain to both sample configurations.

Table 58: Genesys Info Mart 7.5 Performance Testing Configuration Details

Parameter	Descriptions
Call flow	Testing consisted of four call flow scenarios: • Call is routed directly to an agent, and is then released (20 percent) • Call is routed to an Interactive Voice Response (IVR), and is then released (20 percent) • Call is routed to an IVR, and is then transferred to a Routing Point, and then to an agent, and is finally released (30 percent) • Call is routed to an IVR, is transferred to a Routing Point, and then to an agent; after which the call is transferred and is routed to a second agent, and is finally released (30 percent)
Attached data	Ten predefined key-value pairs were stored in various Genesys Info Mart fact and dimension tables, as were five user data string dimensions and five user data facts.
Agents	The testing configuration defined about 6,000 agents. Three shifts, of 2,000 agents each, handled the call volume. The average talk time for the calls handled by an agent was 120 seconds.

Genesys Info Mart Application

Table 59 lists the Genesys Info Mart release 7.5 application options that were used for testing. It includes only those options that affected performance.

Table 59: Genesys Info Mart 7.5 Application Configuration Options

Section	Option
gim-aggregates-tenant	maximum-aggregation-level = MONTH populate-agent-state-aggregates = TRUE populate-skill-combination-aggregates = TRUE populate-skill-demand-aggregates = TRUE
gim-etl	extract-data-time-range-limit = 4 extract-data-time-range-units = HOURS limit-extract-data = TRUE
gim-transformation	transformation-buffer-size = 5
optional-tables	populate-resource-session-facts = TRUE populate-resource-state-facts = TRUE populate-resource-state-reason-facts = TRUE
schedule	aggregate-start-time = 01:45 etl-end-time = 01:00 etl-frequency = 15 etl-start-time = 04:00 load-recent-start-time = 22:15 load-start-time = 01:30 maintain-start-time = 02:00 max-concurrent-extract-jobs = 10 populate-intraday-aggregates = FALSE run-aggregates = TRUE run-load-recent-with-extract-and-transform = TRUE run-scheduler = TRUE ir-merge-interval = 0

Java Memory The Java memory setting for the Genesys Info Mart Server process in the `gim_server.bat` file was set to `-Xmx1500m`.

Scheduling and Configuration of the ETL Jobs

Sample Configurations 1 and 2 used the same scheduling and configuration of the intraday and daily ETL jobs.

Scheduling

The intraday ETL cycle was scheduled with an `etl-frequency` of 15 minutes (900 seconds), which allowed for 80 cycles per day. Because the call rate remained constant for a full 24 hours a day, no allowances were made for peak or off-peak hours.

The daily jobs were scheduled to start running between 1:30 AM and 2:00 AM.

About the Extraction Jobs

The intraday ETL cycle included data extraction jobs for configuration details, voice interaction details, and voice agent details. These details were extracted from both Interaction Concentrator and Stat Server databases using the following jobs:

- Separate instances of `Job_ExtractICON` for roles `ICON_CFG` and `ICON_CORE` were used to extract configuration details and voice interaction details, respectively.
- `JobExtractSS` was used to extract voice agent details from Stat Server database.

About the Daily ETL Jobs

The daily ETL jobs included the following:

- `Job_LoadGIM` - This job moves facts from the intraday fact tables (where they were loaded by `Job_LoadRecent` as part of the intraday ETL cycle) to the historical fact tables
- `Job_AggregateGIM` - This job calculates aggregates based on the newly loaded historical fact table data.

Note: The results for `Job_MaintainGIM` were not measured because the testing did not generate more than ten days worth of data, so there were no fact tables old enough to be eligible for deletion.

How Job Durations Were Calculated

The performance results graphs throughout this chapter show the durations for the various intraday and daily ETL jobs.

To measure the duration of the extraction phase of the intraday ETL cycle, only the longest running job was used—`JobExtractICON` for role `ICON_CORE`.

The overall duration of the intraday ETL cycle was determined by combining the runtimes of the following intraday jobs: `Job_ExtractICON` for role `ICON_CORE` + `Job_TranformGIM` + `Job_LoadRecent`.

Factors Affecting Release 7.5 ETL Performance

The primary factors affecting extraction, transformation, and loading (ETL) performance in release 7.5 are as follows:

- The number of daily interactions in your contact center.
- The number of agents in your contact center.
- The complexity of your interaction flows.
- The amount of business data attached to interactions.
- The hardware on which the Genesys Info Mart Server is running, primarily the CPU speed and available memory.
- The hardware on which the Staging Area RDBMS server is running, primarily the CPU speed, disk speed, and available memory.
- The tuning of the Staging Area database
- The speed of the network connections between components.
- The hardware on which the Interaction Concentrator RDBMS server is running; primarily the CPU speed, disk speed, and available memory.
- The amount of historical data retained in the Interaction Concentrator database.

General Performance Results for Release 7.5

Because similar configurations were used for both Sample Configurations 1 and 2, certain similar trends were noticed in the results.

Daily Spikes in the Results Graphs

In the performance graphs throughout this chapter, notice the daily spikes in the intraday ETL cycle execution time. In these tests, the intraday ETL cycle was suspended for four hours each day to allow the daily ETL jobs to run. These regular spikes show the time it took for the resumed intraday ETL cycle to process the backlog of source data that had accumulated during those four hours.

Gradual Increase in Average Duration from Day to Day

In these performance tests, no provision was made for long-term maintenance of the Genesys Info Mart database. This resulted in a gradual increase in the average duration of the intraday ETL cycles from one day to the next. To keep this increase within an acceptable range for your environment, Genesys recommends that you periodically perform index maintenance on, and purge old data from, the Interaction Database tables containing voice details. For more information about both purging and index maintenance, see “Genesys Info Mart 7.5 Performance Tuning” on page 164.

Release 7.5 ETL Performance for Sample Configuration 1

Genesys Info Mart 7.5 performance testing used the time that it took to execute the ETL jobs as its primary measurement of performance. In these tests, the intraday and daily ETL jobs were reviewed separately. Testing was performed on two database platforms: Oracle 10g and Microsoft SQL Server 2005.

Configuration Details

For configuration details specific to Sample Configuration 1, see:

- "Call Volume for Sample Configuration 1" on page 139
- "Hardware Architecture—Sample Configuration 1" on page 140

Performance Results

For performance results on the two database platforms, see:

- "Volume of Generated Data—Sample Configuration 1" on page 140
- "Release 7.5 Performance Trends for Oracle 10g" on page 142
- "Release 7.5 Performance Trends for Microsoft SQL Server 2005" on page 159
- "Aggregate Duration by Day" on page 155

Call Volume for Sample Configuration 1

To test Genesys Info Mart 7.5 ETL performance in Sample Configuration 1, calls were generated at a constant rate of 12 calls per second (cps) for seven days, and then raised to 14 cps for two and a half more days. This meant that:

- At the rate of 12 cps, 10,800 calls were generated during each 15-minute cycle.
- The total number of calls generated each day was 1,036,800.

Hardware Architecture—Sample Configuration 1

Figure 50 shows the distribution of software components across four servers used in the testing of Sample Configuration 1.

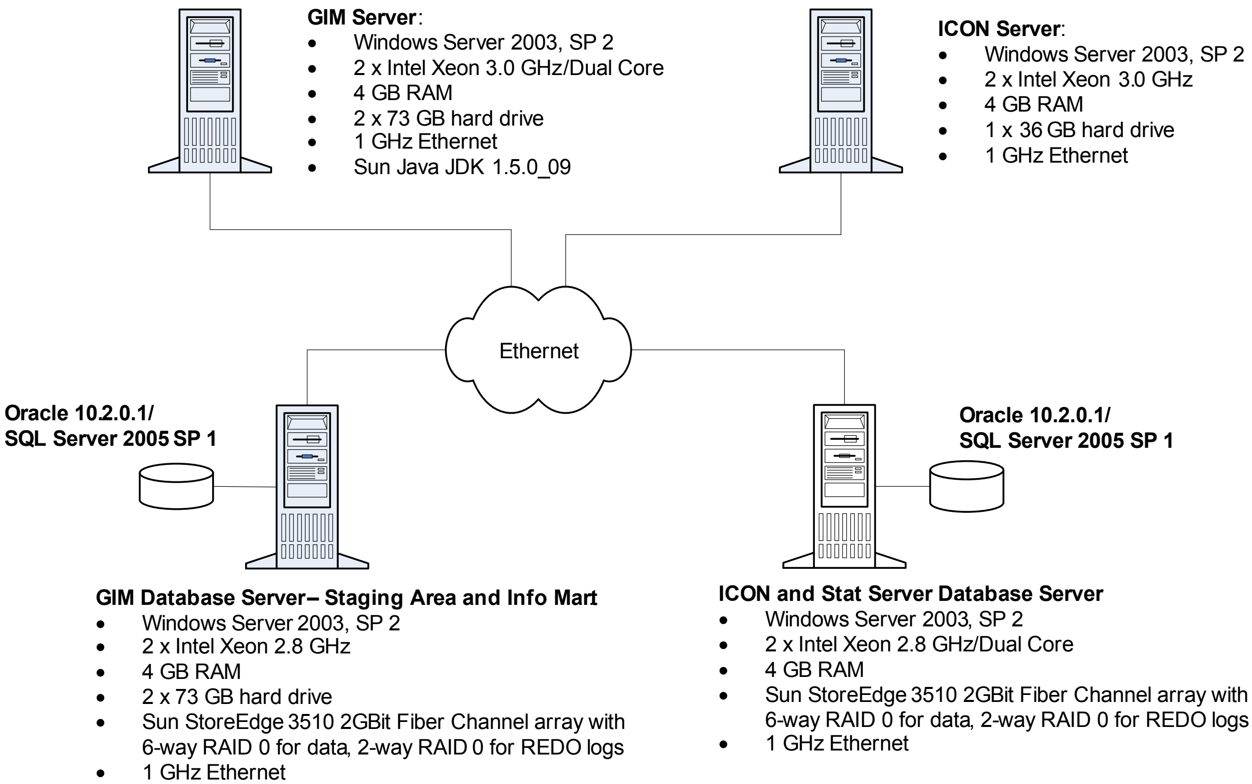

Figure 50: Hardware Architecture—Sample Configuration 1

Volume of Generated Data—Sample Configuration 1

Tables 60–62 provide details about the volume of data extracted and loaded during Sample Configuration 1 testing:

Table 60 lists the number of rows extracted from the Stat Server database and the IDB during each 15-minute intraday ETL cycle—at the Sample Configuration 1 rate of 12 cps.

Table 60: Data Extracted Per Intraday ETL Cycle—12 Calls/Second

Source	Source Table Name	Rows Extracted
Stat Server	STATUS_TABLE	127,300
	VOICE_REASONS	26,600

Table 60: Data Extracted Per Intraday ETL Cycle—12 Calls/Second (Continued)

Source	Source Table Name	Rows Extracted
ICON	G_CALL	32,200
	G_IR	19,300
	G_IS_LINK_HISTORY	17,300
	G_PARTY	81,600
	G_PARTY_HISTORY	199,600
	G_PARTY_STAT	81,600
	G_ROUTE_RESULT	8,700
	G_USERDATA_HISTORY	405,800

Table 61 lists the number of rows loaded to Genesys Info Mart during each 15-minute intraday ETL cycle—at the Sample Configuration 1 rate of 12 cps.

Table 61: Data Loaded Per Intraday ETL Cycle—12 Calls/Second

Genesys Info Mart Table Name	Rows Loaded
R_INTERACTION_FACT	10,800
R_INTERACTION_SEGMENT_FACT	74,300
R_RESOURCE_STATE_FACT	127,600
R_RESOURCE_STATE_REASON_FACT	26,600
R_VOICE_IXN_FACT_EXT	10,800
R_VOICE_SEG_FACT_EXT	74,300

Table 62 lists the average total data volume that the daily ETL jobs loaded to Genesys Info Mart—at the Sample Configuration 1 rate of 12 cps.

Table 62: Data Loaded Per Day—12 Calls/Second

Genesys Info Mart Table Name	Rows Loaded
INTERACTION_FACT	1,032,000
INTERACTION_SEGMENT_FACT	7,100,000
RESOURCE_STATE_FACT	12,125,000
RESOURCE_STATE_REASON_FACT	2,522,000
VOICE_IXN_FACT_EXT	1,032,000
VOICE_SEG_FACT_EXT	7,100,000

Note: Tables 60–62 do not include the numbers of Genesys Info Mart dimension and configuration fact table rows extracted and loaded during each cycle. This is because the expected data volume for new and changed rows is expected to be very low compared to the number of interaction and resource fact table rows.

Release 7.5 Performance Trends for Oracle 10g

This section includes general notes on the release 7.5 testing results, as well as a number of graphs showing the results, in the following subsections:

- "Notes On Oracle Testing" on page 142
- "Total Job Duration by Cycle—Oracle" on page 143
- "Average Intraday ETL Cycle Duration—Oracle" on page 144
- "Average Intraday Extraction Cycle Duration—Oracle" on page 144
- "Job_ExtractICON for Role ICON_CORE (All Call Details)" on page 145
- "Transform Duration" on page 147
- "Load Recent Duration" on page 147
- "Daily Job Executions" on page 148

Notes On Oracle Testing

In reviewing the results of the Oracle performance tests, note how the following mid-test changes to the testing parameters affected the results:

- Before cycle 221, testing did not include the collection of voice agent details from the Stat Server database. However, after cycle 221, testing began to include these Stat Server details—at this point in the graphs, you can see a marked increase in the intraday ETL cycle time.
- At cycle 547, the call rate was increased from 12 to 14 cps, which accounts for a second marked increase in the intraday ETL cycle time.

Total Job Duration by Cycle—Oracle

Figure 51 depicts the total execution time for all intraday jobs throughout the testing period. The regular spikes are the result of the four-hour pause in the intraday ETL cycle to allow for the daily ETL jobs—the spikes show the subsequent “catch-up” period that the intraday ETL cycle required to process the backlog of data.

Notice the marked increases in these spikes at two stages:

- After the “No Stat Server” stage of testing (shown as the blue portion of the graph), when testing began to include details from the Stat Server database.
- When the call rate changed from 12 cps (shown as the gray portion) to 14 cps (shown as the pink portion).

Also note that it took an average of two cycles—totaling between 60 and 90 minutes—to process each four-hour backlog of data, implying that the system can sustain significantly larger spikes of activity.

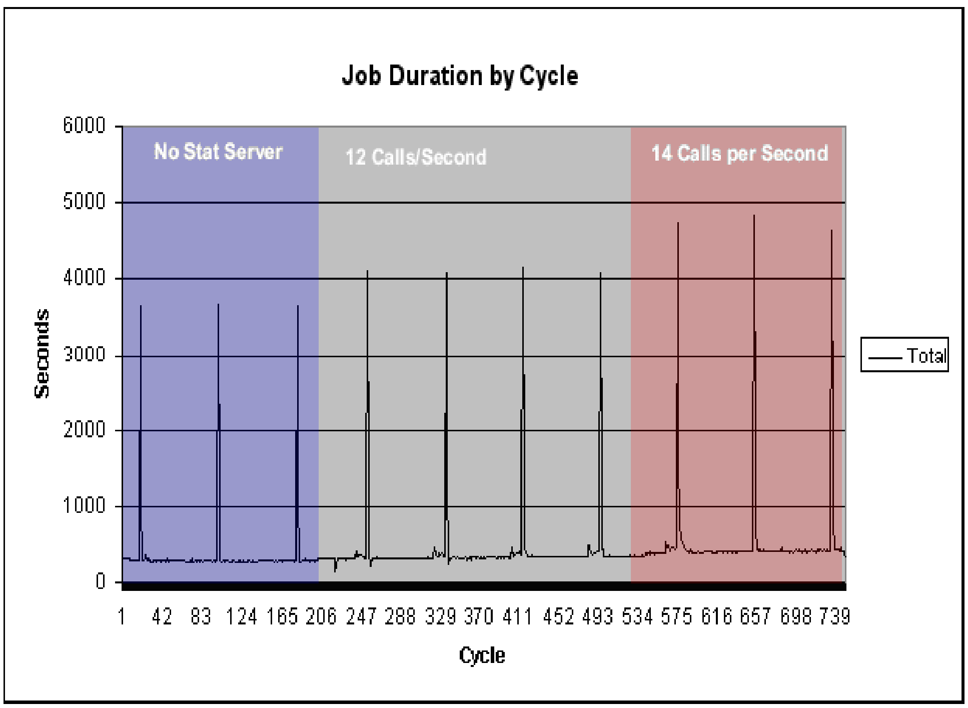

Figure 51: Total Job Duration by Cycle—Oracle

Average Intraday ETL Cycle Duration—Oracle

Figure 52 depicts the daily average duration of the intraday ETL cycles. You can see evidence of a gradual increase in the duration from day to day. This increase is caused by the accumulation of historical data in the Interaction Database, specifically the voice details extracted by `Job_ExtractICON` for role `ICON_CORE`. To keep the duration within an acceptable range for your environment, Genesys recommends that you periodically perform index maintenance on, and purge old data from, the Interaction Database tables that contain voice details.

Also note that for the configured 15-minute intraday ETL cycle, fewer than 10 minutes were used, indicating that the environment could sustain higher call rates.

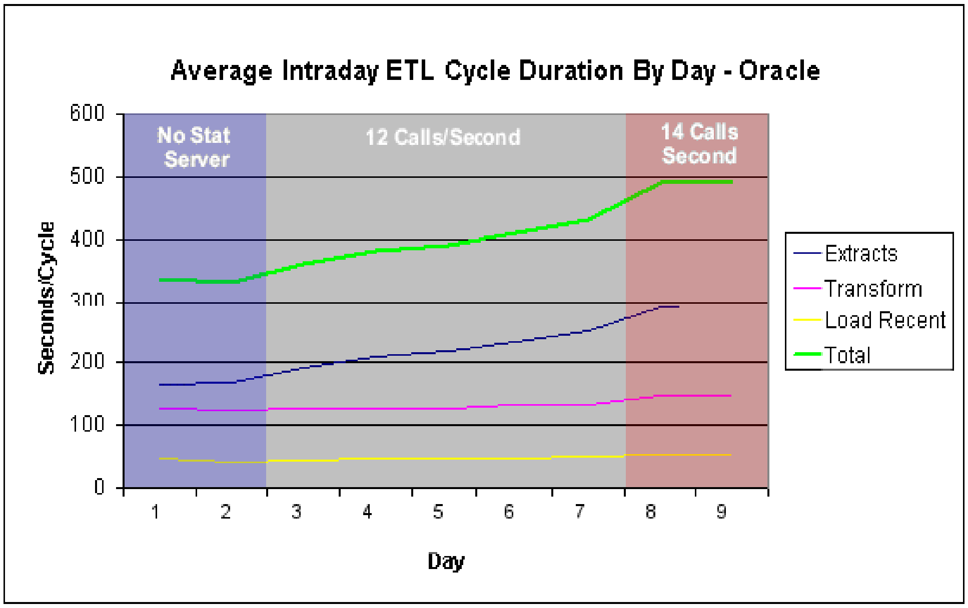

Figure 52: Average Intraday ETL Cycle Duration—Oracle

Average Intraday Extraction Cycle Duration—Oracle

Figure 53 depicts the performance trends for the individual extraction jobs, showing daily averages for each. The duration of `Job_ExtractICON` for role `ICON_CFG` grew as a side effect of the increasing load on the Interaction Database Server imposed by `Job_ExtractICON` for role `ICON_CORE`.

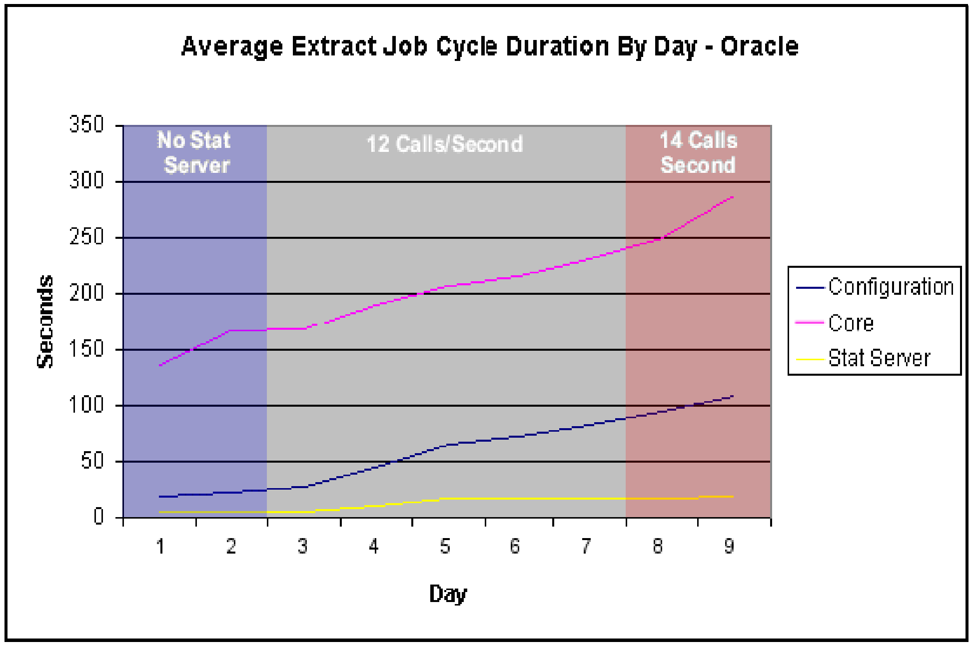

Figure 53: Average Intraday Extraction Cycle Duration—Oracle

Job_ExtractICON for Role ICON_CORE (All Call Details)

Figure 54 shows that `Job_ExtractICON` (as the longest-running extraction job) for role `ICON_CORE`, determined the duration of the extraction phase of the intraday ETL cycle. You can see a gradual increase in the duration of this job over time, despite the constant call volume. This increase may be related to the overall size of the Interaction Concentrator database, or to fragmentation of the ICON indexes on the tables from which the data was extracted.

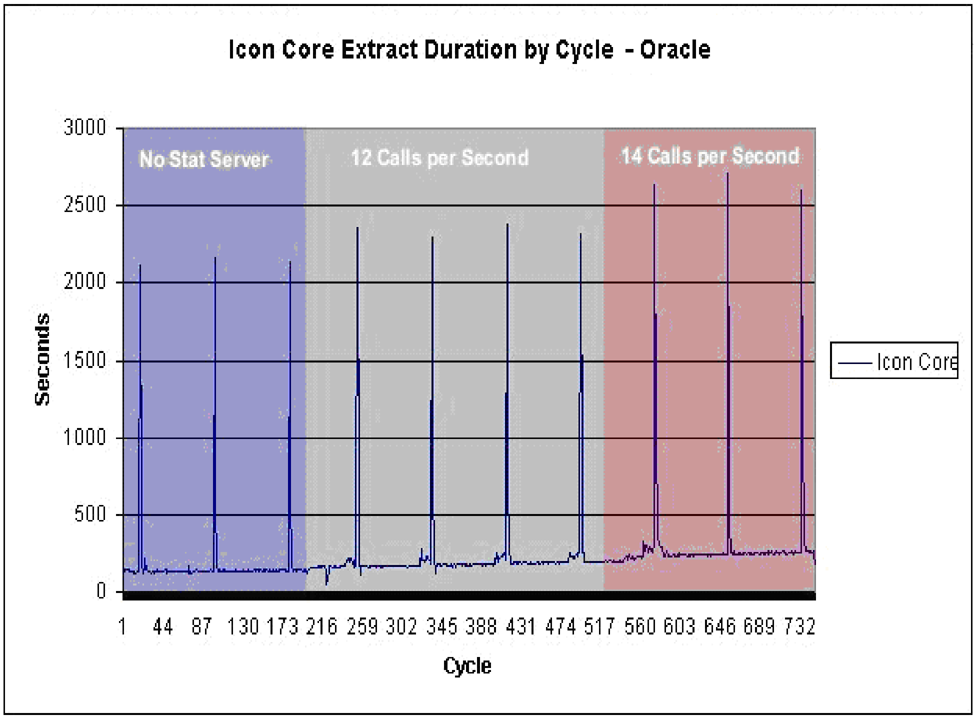

Figure 54: Job_ExtractICON for Role ICON_CORE (All Call Details)

You can see this gradual increase more clearly in the edited graph (Figure 54) in which most of the daily "catch-up" spikes were removed.

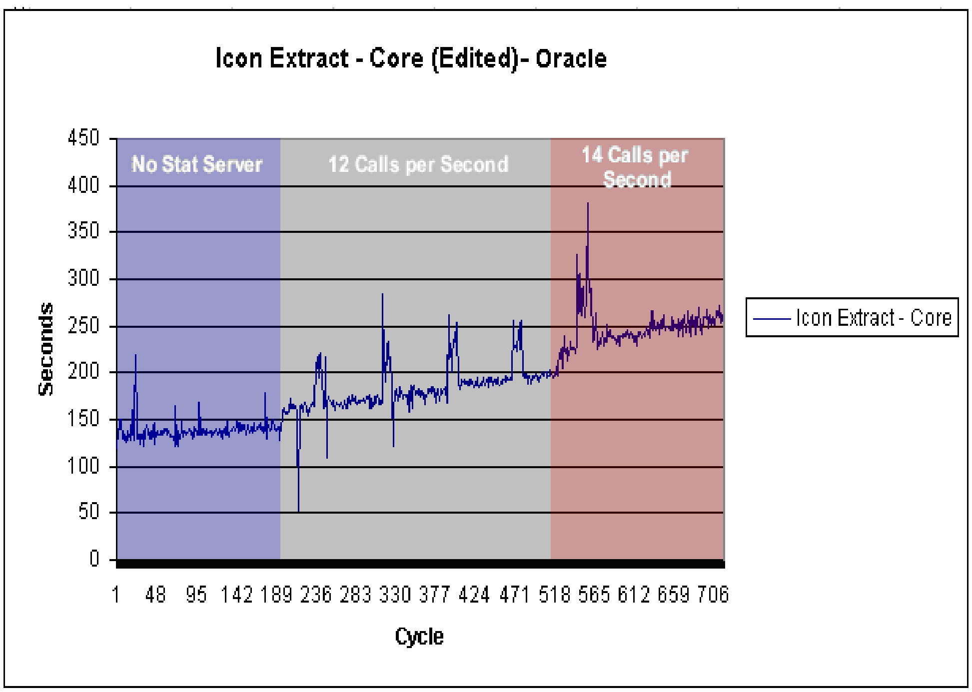

Figure 55: Job_ExtractICON for Role ICON_CORE (Edited)

Note: Performance testing did not include any purging or index maintenance on the Interaction Concentrator database. Periodic purging or index maintenance of the Interaction Concentrator database will prevent the intraday ETL cycle duration from increasing to an unacceptable level.

Transform Duration

Figure 56 shows that the duration for `JOB_TransformGIM` remained basically flat, except for two isolated increases attributable to the addition of Stat Server source data (cycle 221) and the increased call rate (cycle 547).

This graph shows the limited impact of Stat Server data and increased call volume. By the end of the test, the average duration for `JOB_TransformGIM` had increased to 110 seconds.

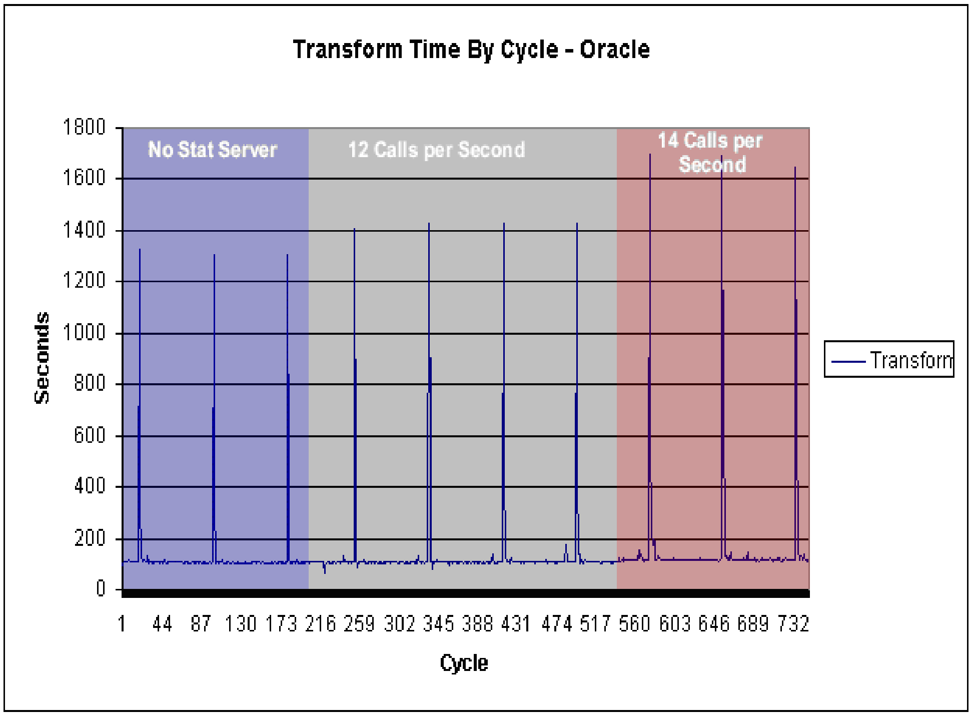

Figure 56: Transform Time by Cycle—Oracle

Load Recent Duration

Figure 57 shows the duration of the `Job_LoadRecent` phase of the intraday ETL cycle throughout the testing period. Apart from increases due to the additional load from Stat Server (cycle 221), and the increase in call rate (cycle 547), the duration remained flat. By the end of the testing period, the average duration of `Job_LoadRecent` had not increased beyond 50 seconds.

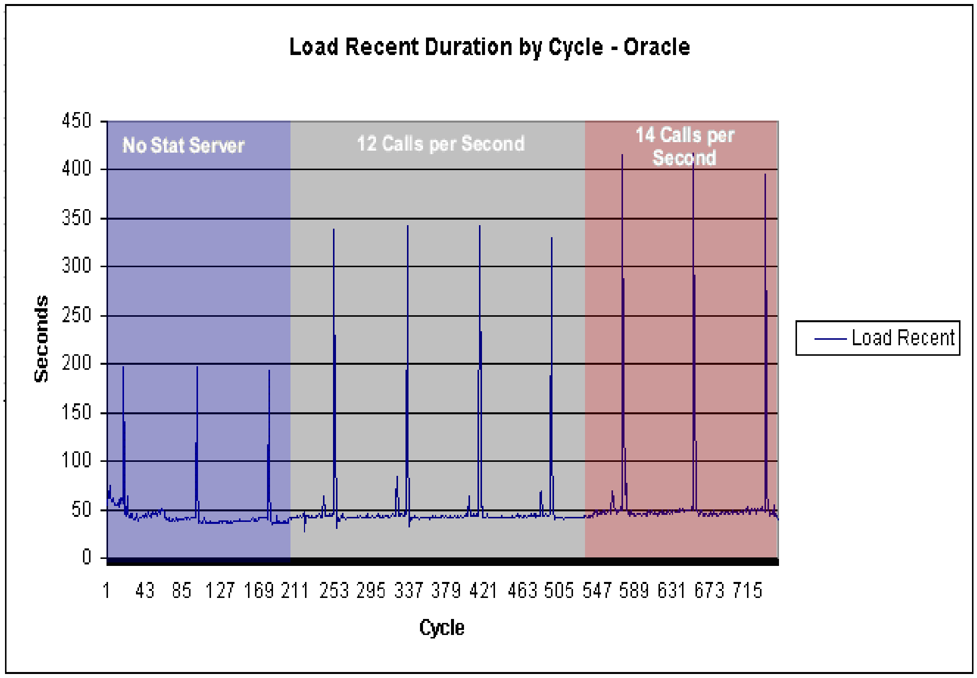

Figure 57: Load Recent Duration by Cycle—Oracle

Daily Job Executions

No optimization for target database growth was performed before or during testing. The performance results, shown in Figures 58 and 59 demonstrate that, for these call volumes, an adequate environment requires such database optimization.

Genesys recommends that you consult a database administrator (DBA) knowledgeable in data warehouses to analyze the projected growth of your Genesys Info Mart database—according to a size estimation spreadsheet—and that you plan the storage, partitioning, and any other optimization measures accordingly.

Figure 58 shows the performance result information for Load Genesys Info Mart.

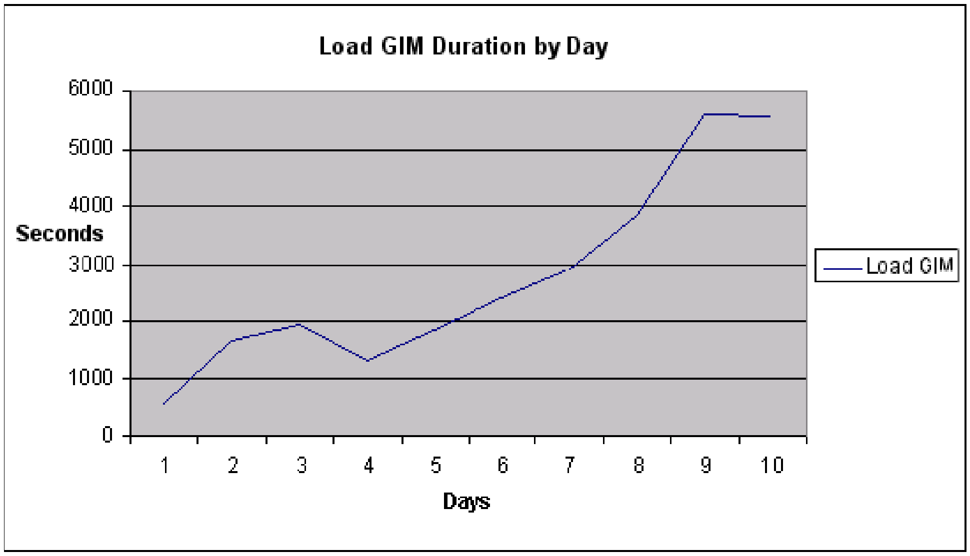

Figure 58: Load Genesys Info Mart

Figure 59 shows the performance results for Aggregate Genesys Info Mart Duration.

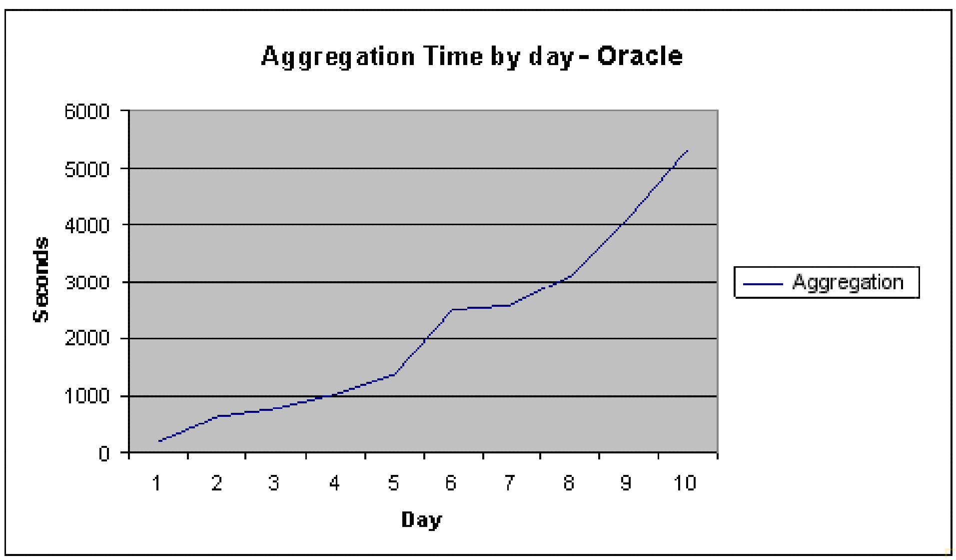

Figure 59: Aggregate Genesys Info Mart Duration

Release 7.5 Performance Trends for Microsoft SQL Server 2005

This section includes general notes on the release 7.5 testing results as well as a number of graphs showing the results themselves, in the following subsections:

- "Notes On Microsoft SQL Server Testing" on page 150
- "Total Job Duration by Cycle—Microsoft SQL Server" on page 150
- "Extract Duration by Cycle—Microsoft SQL Server" on page 152
- "Extract Duration By Cycle" on page 152
- "Daily Jobs—Microsoft SQL Server" on page 154

Notes On Microsoft SQL Server Testing

- Extract job runtimes were erratic during the first 30 cycles. At cycle 30, the Interaction Concentrator database server was reconfigured with `AUTO_UPDATE_STATISTICS_ASYNC ON`. This smoothed out the extract runtimes.
- At cycle 400, the Daily Aggregation job ran long due to outdated statistics on the `INTERACTION_FACT` and `INTERACTION_SEGMENT_FACT` tables. This caused the intraday jobs to wait eight hours instead of the usual four hours. The intraday jobs then ran longer than normal for several cycles.

Note: In the generally-available 7.5.003.05 release of Genesys Info Mart, a modification was made to correct this problem of stale statistics impacting the aggregation.

- At cycle 430, the call generators were down for environmental reasons. Normal processing resumed at cycle 460.

Total Job Duration by Cycle—Microsoft SQL Server

Figure 60 shows that aside from the "catch-up" periods, the overall runtime of the combined intraday jobs—`Extract ICON Core` + `TransformGIM` + `LoadRecent`—remained consistent throughout the test, at about 400 seconds.

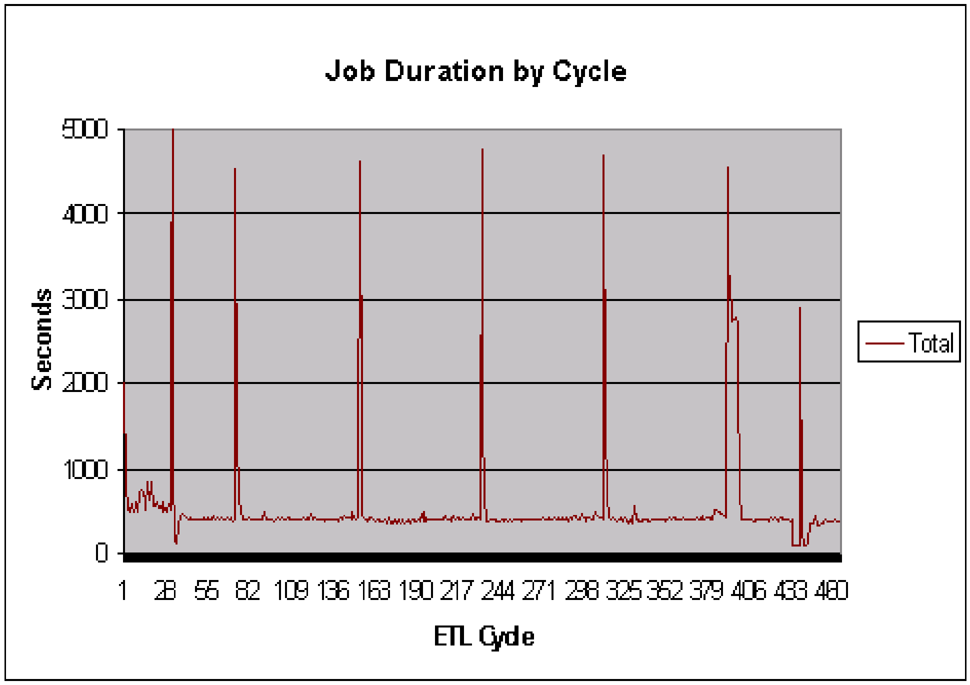

Figure 60: Job Duration By Cycle

Extract Duration by Cycle—Microsoft SQL Server

Figure 61 shows that the ICON Core job extract showed no signs of degrading for the duration of this test. At this point, no purging or other maintenance on the Interaction Concentrator database is needed.

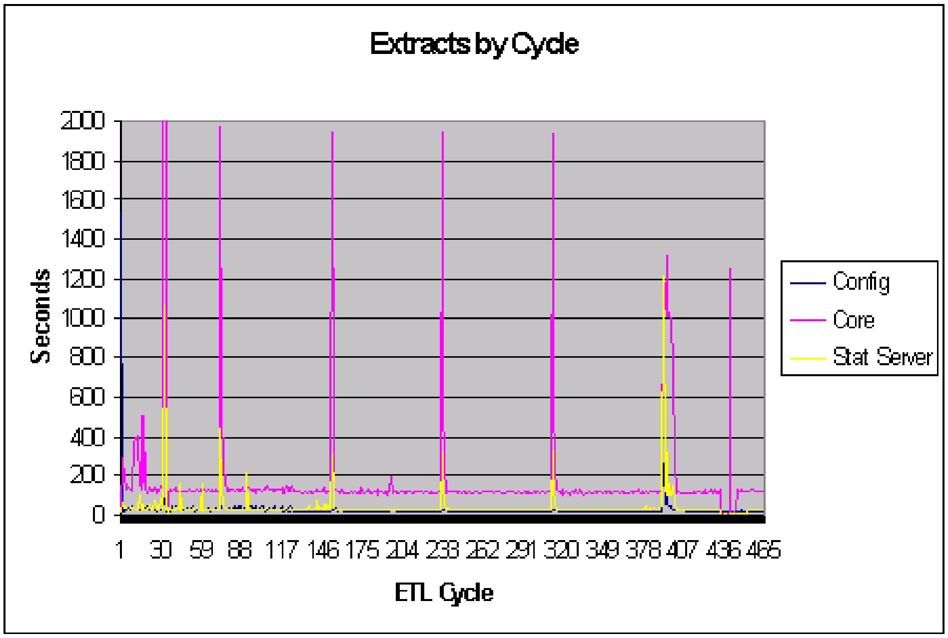

Figure 61: Extract Duration By Cycle

Transform and Load Recent Duration by Cycle—Microsoft SQL Server

Figure 62 shows that after the environment was tuned, the execution time for the Transform and LoadRecent jobs remained flat.

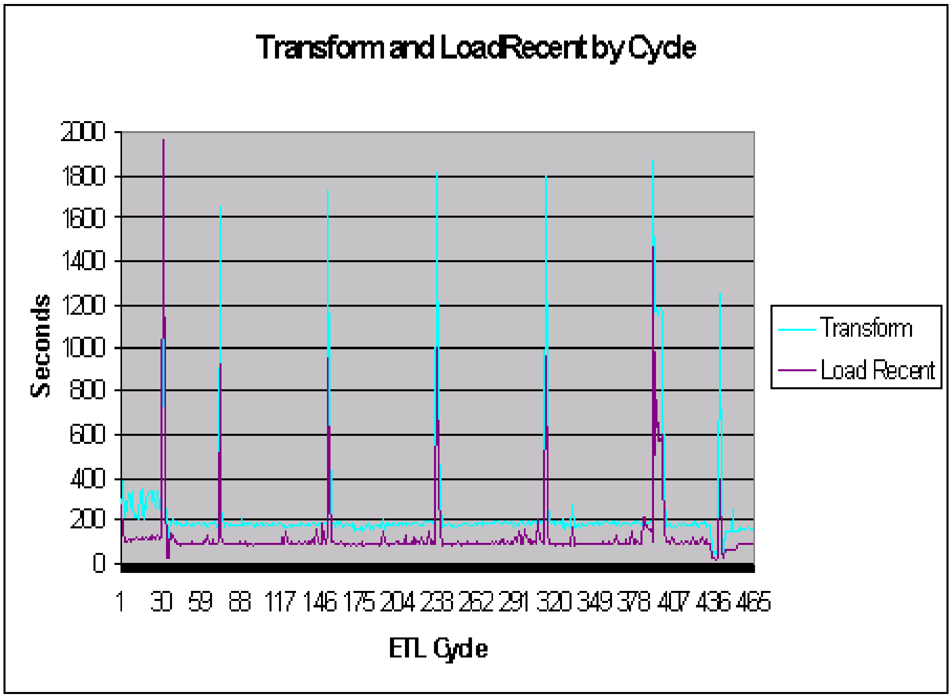

Figure 62: Transform and Load Recent by Cycle

Daily Jobs—Microsoft SQL Server

Figure 63 shows the slow increase in execution time over the course of the test.

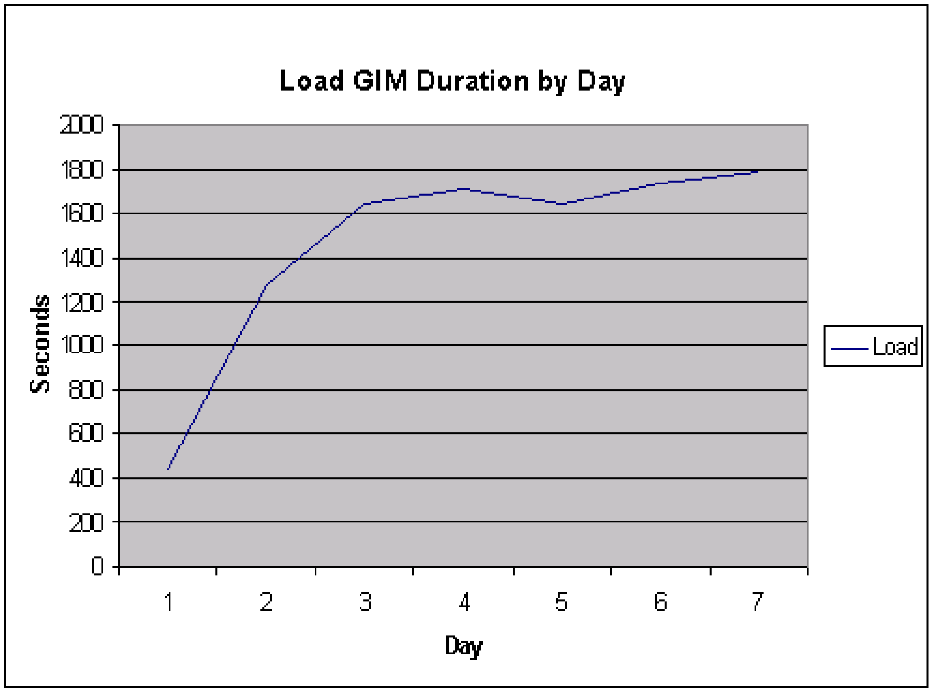

Figure 63: Load Duration by Day

Figure 64 shows the rapid increase in the length of the Aggregation job. Genesys recommends that you take steps to address this issue at the database level. Given that the aggregation job deals primarily with only the most recently loaded day of data, one solution involves partitioning the main fact tables (especially `INTERACTION_FACT` and `INTERACTION_SEGMENT_FACT`) by day, using the `STD_TENANT_DATE_KEY` column. This should improve the performance of the aggregation queries and flatten the execution time curve.

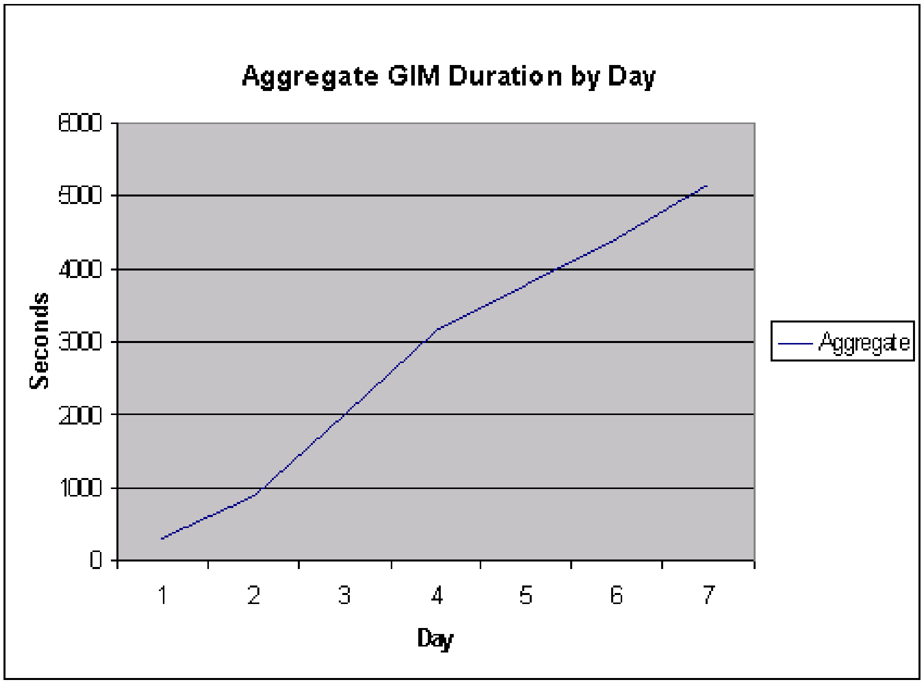

Figure 64: Aggregate Duration by Day

Data Consumption

Table 63 shows the amount of hard disk space that Genesys Info Mart 7.5 ETL performance testing consumed.

Table 63: Hard Disk Space Consumption for Release 7.5 ETL Performance Testing

	On Oracle	On Microsoft SQL Server
ICON and Stat Server	10 million recorded calls used 350 GB of database space (including configuration data), so daily consumption was about 35 KB/call, or 35 GB/day.	Daily consumption was 38 GB/day for ICON and 1.4 GB/day for Stat Server.
Genesys Info Mart	10 million recorded calls used 100 GB of database space, so daily consumption was about 10 GB/day	Daily consumption was about 12 GB/day.

Release 7.5 ETL Performance for Sample Configuration 2

Genesys Info Mart 7.5 ETL performance testing for Sample Configuration 2 was conducted using the same call flows and Genesys Info Mart configuration as were used for the Sample Configuration 1 tests, but at a much lower call volume of 3 calls per second (cps). Testing was performed for Microsoft SQL Server 2005 only.

Configuration Details

For more configuration details specific to Sample Configuration 2, see:

- “Call Volume for Sample Configuration 2” on page 156
- “Hardware Architecture—Sample Configuration 2” on page 157

Performance Results

For performance results on Microsoft SQL Server 2005, see:

- “Volume of Generated Data for Sample Configuration 2” on page 157
- “Release 7.5 Performance Trends for Microsoft SQL Server 2005” on page 159
- “Data Consumption” on page 163

Call Volume for Sample Configuration 2

To test Genesys Info Mart 7.5 ETL performance in Sample Configuration 2, calls were generated at a constant rate of 3 cps for 15 days. This meant that:

- At a rate of 3 cps, 2,700 calls were generated during every 15-minute cycle.
- The total number of call generated each day was 259,200.

Hardware Architecture—Sample Configuration 2

Testing on the Sample Configuration 2 used a distribution of software components across several servers, as shown in Figure 65.

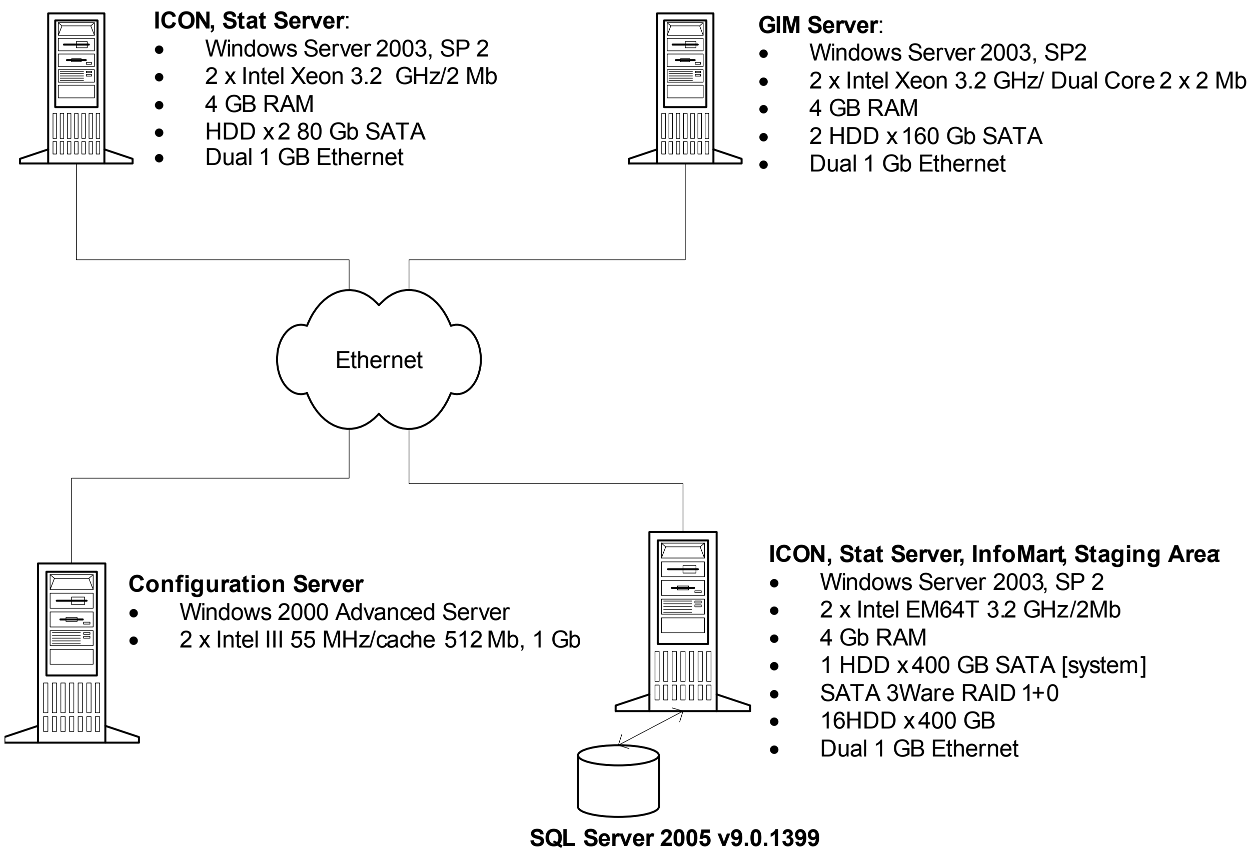

Figure 65: Hardware Architecture—Sample Configuration 2

Volume of Generated Data for Sample Configuration 2

Tables 64–66 provide details about the volume of data generated during testing for Sample Configuration 2.

Table 64 lists the number of rows extracted for Stat Server and ICON during each 15-minute cycle—at the Sample Configuration 2 rate of 3 cps.

Table 64: Data Volume Per Intraday ETL Cycle—3 Calls/Second

Source System	Source Table Name	Rows Extracted
Stat Server	STATUS_TABLE	31,800
	VOICE_REASONS	6,700

Table 64: Data Volume Per Intraday ETL Cycle—3 Calls/Second (Continued)

Source System	Source Table Name	Rows Extracted
ICON	G_CALL	8,000
	G_IR	4,800
	G_IS_LINK_HISTORY	4,300
	G_PARTY	20,400
	G_PARTY_HISTORY	49,900
	G_PARTY_STAT	20,400
	G_ROUTE_RESULT	2,200
	G_USERDATA_HISTORY	101,400

Table 65 lists the number of rows loaded to Genesys Info Mart during each 15-minute cycle—at the Sample Configuration 2 rate of 3 cps.

Table 65: Data Loaded Per Intraday ETL Cycle—3 Calls/Second

Genesys Info Mart Table Name	Rows Loaded
R_INTERACTION_FACT	2,700
R_INTERACTION_SEGMENT_FACT	18,600
R_RESOURCE_STATE_FACT	31,900
R_RESOURCE_STATE_REASON_FACT	6,700
R_VOICE_IXN_FACT_EXT	2,700
R_VOICE_SEG_FACT_EXT	18,600

Table 66 lists the average total data volume for a daily load at 3 cps.

Table 66: Data Loaded Per Day—3 Calls/Second

Genesys Info Mart Table Name	Rows Loaded
INTERACTION_FACT	258,000
INTERACTION_SEGMENT_FACT	1,775,000

Table 66: Data Loaded Per Day—3 Calls/Second (Continued)

Genesys Info Mart Table Name	Rows Loaded
RESOURCE_STATE_FACT	3,031,000
RESOURCE_STATE_REASON_FACT	630,000
VOICE_IXN_FACT_EXT	258,000
VOICE_SEG_FACT_EXT	1,775,000

Note: Tables 64–66 do not include the number of dimension and configuration table rows in Genesys Info Mart that are extracted and loaded during each cycle. This is because the expected data volume for new and changed rows is anticipated to be very low compared to the number of interaction and resource fact table rows.

Release 7.5 Performance Trends for Microsoft SQL Server 2005

This section includes general notes on the release 7.5 testing results as well as a number of graphs showing the results, in the following subsections:

- "Notes On Microsoft SQL Server Testing" on page 150
- "Total Job Duration by Cycle—Microsoft SQL Server" on page 150
- "Extract Duration by Cycle—Microsoft SQL Server" on page 152
- "Transform and Load Recent Duration by Cycle—Microsoft SQL Server" on page 153
- "Daily Jobs—Microsoft SQL Server" on page 154

Notes on Microsoft SQL Server Testing

- Figures 66 and 67 show that around cycle 200, the performance of the Extract ICON Core job began to decrease steadily. Various attempts were made to improve the performance around cycles 600–680, but with inconsistent effect.

  At cycle 687, several indexes in ICON appeared severely fragmented. Five indexes were then rebuilt with a Fill Factor of 20 percent. This dramatically improved both the Extract ICON execution and the performance of the ICON process itself.

  Note: The cause of the initial performance degradation was attributed to index fragmentation of the Interaction Concentrator database.

The Interaction Concentrator indexes were:

- IDX_G_IR_ROOTIRID on table G_IR.
- IDX_G_CALL_ROOTIRID on table G_CALL.
- GIM_IDX_PAR_CALLID on table G_PARTY.
- IDX_G_PARTY_H_PID on table G_PARTY_HISTORY.
- GIM_IDX_UH_CALLID on table G_USERDATA_HISTORY.

- In Figures 66 and 68, the large peaks in runtime are "catch-up" periods for the Intraday jobs after the four-hour nightly window for the running of the daily jobs. In Figures 66 and 67, the small peaks between the large ones were eliminated at around cycle 900, when the Interaction Concentrator database server was reconfigured with AUTO_UPDATE_STATISTICS_ASYNC ON.
- Job_AggregateGIM was not run for this test.

Total Job Duration By Cycle—Microsoft SQL Server

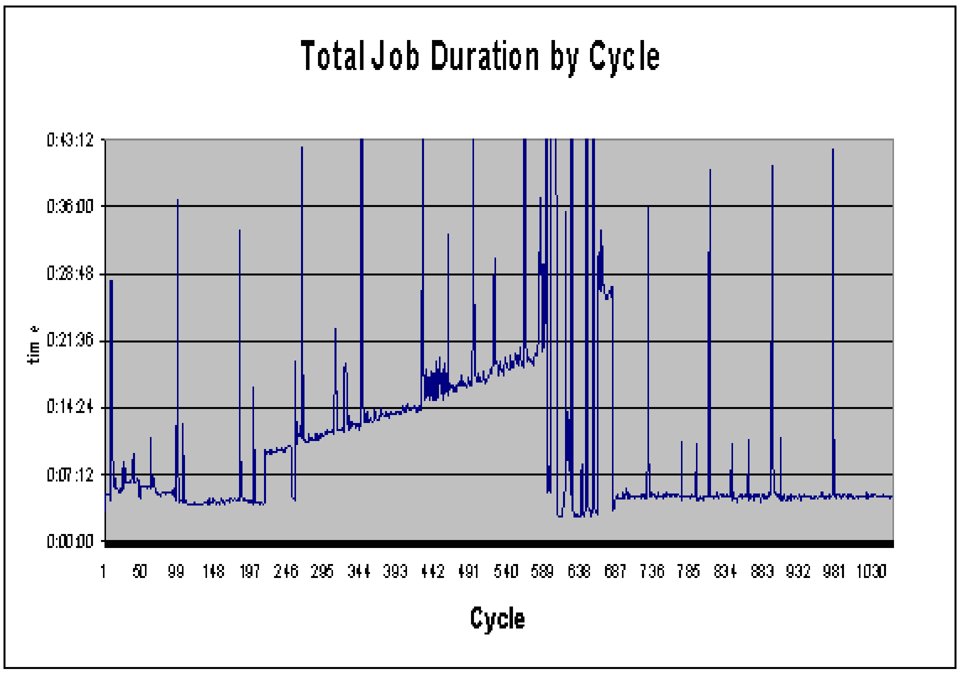

Figure 66: Total Job Duration by Cycle

Extract Duration By Cycle—Microsoft SQL Server

Figure 67 illustrates the slowdown caused by index fragmentation of the Interaction Concentrator database (see "Notes on Microsoft SQL Server Testing" on page 159).

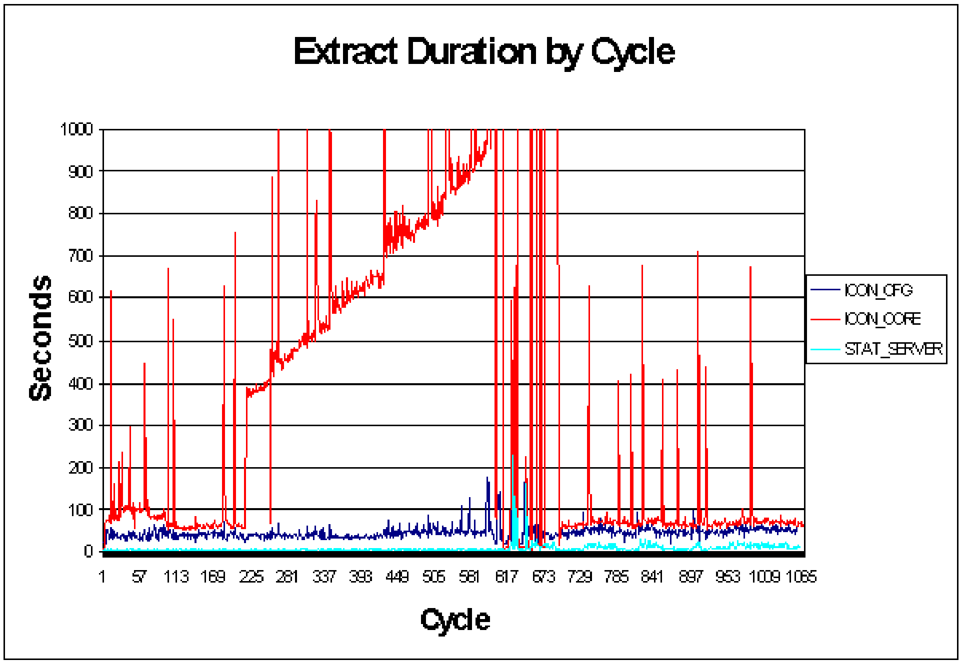

Figure 67: Extract Duration by Cycle

Transform and Load Recent Duration By Cycle—Microsoft SQL Server

Figure 68 shows that around cycles 660–680, the execution times for `Job_TransformGIM` and `Job_LoadRecent` became erratic due to problems with `Job_ExtractICON` (see "Notes on Microsoft SQL Server Testing" on page 159). When the identified problems were resolved, the performance became consistent.

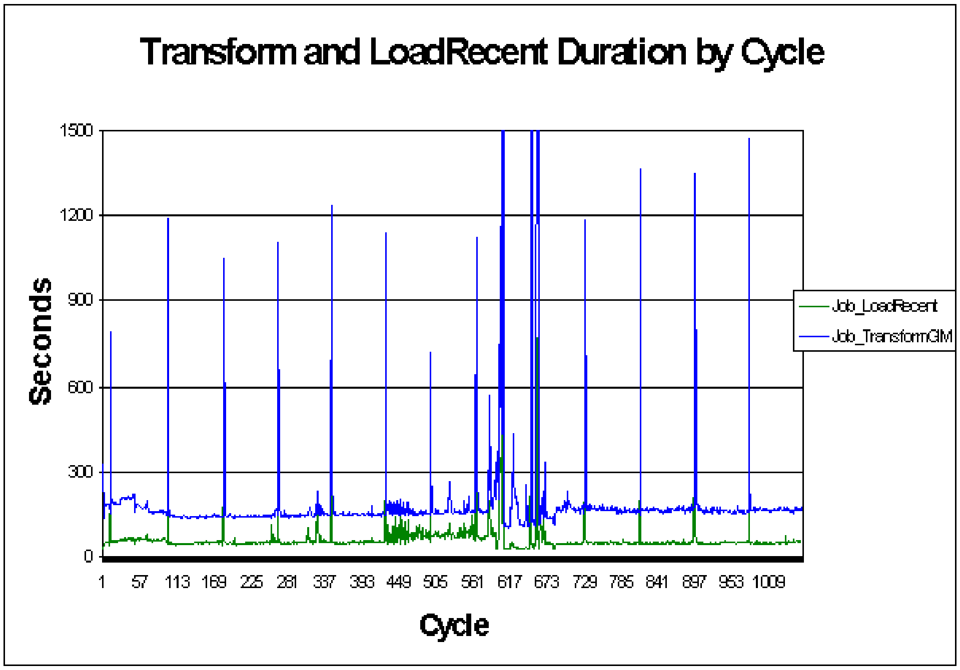

Figure 68: Transform and LoadRecent Duration by Cycle

Daily Jobs—Microsoft SQL Server

Figure 69 shows that the execution time of `Job_LoadGIM` decreased when the Info Mart database server was reconfigured with `AUTO_UPDATE_STATISTICS_ASYNC ON` after day nine.

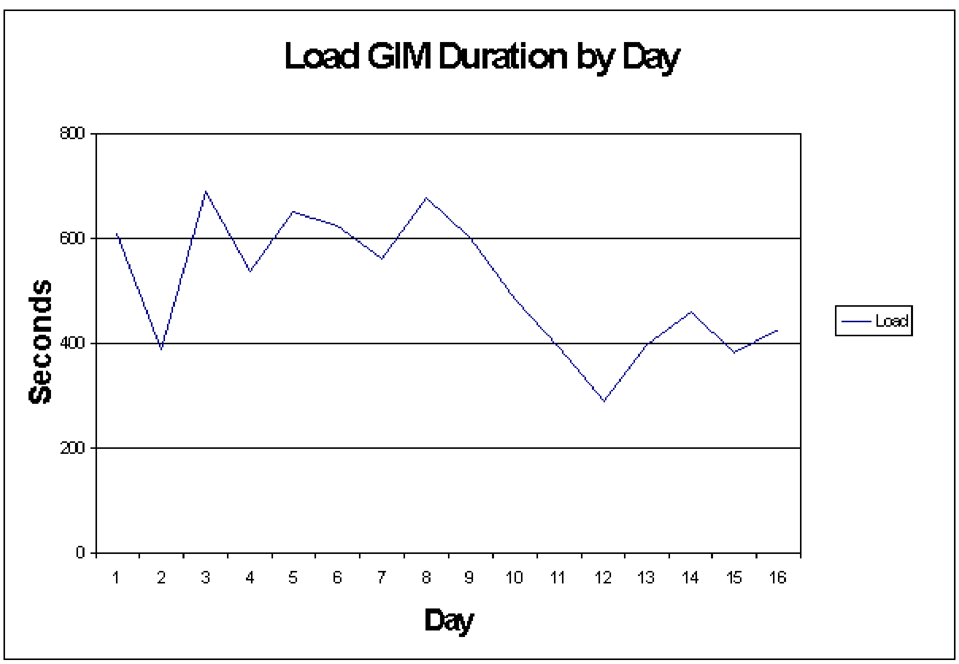

Figure 69: Load GIM Duration by Day

Data Consumption

Table 67 shows the amount of hard disk space that Genesys Info Mart 7.5 ETL performance testing consumed.

Table 67: Hard Disk Space Consumption for Release 7.5 ETL Performance Testing

	On SQL Server
ICON and Stat Server	Daily consumption was 9.5 GB/day for ICON, and 350 MB/day for Stat Server.
Genesys Info Mart	Daily consumption was 3 GB/day.

Genesys Info Mart 7.5 Performance Tuning

Genesys recommends that you tune Genesys Info Mart 7.5 performance according to the guidelines in Table 68.

Table 68: Release 7.5 Performance Tuning Guidelines

Guideline	Description
Purge the IDB.	To maintain ETL job performance over time, Genesys recommends that you periodically purge the Interaction Concentrator database (IDB). The frequency required depends on the hardware and data volume. Current Call Concentrator (CCON) customers have procedures for manually purging the CCON database. Because ICON generates several times more data for the same call volume, it needs to be purged more frequently than CCON. For more information about procedures to purge the Interaction Concentrator database, see the *Genesys Info Mart 7.5 Deployment Guide.*
Defragment indexes.	If you find the performance of `Job_ExtractICON` degrading over time, check for index fragmentation on the IDB tables. In particular, look at the following five indexes: • `IDX_G_IR_ROOTIRID` on table `G_IR` • `IDX_G_CALL_ROOTIRID` on table `G_CALL` • `GIM_IDX_PAR_CALLID` on table `G_PARTY` • `IDX_G_PARTY_H_PID` on table `G_PARTY_HISTORY` • `GIM_IDX_UH_CALLID` on table `G_USERDATA_HISTORY` Consider defragmenting these indexes by using either of these methods: • For Oracle, use `COALESCE` • For Microsoft SQL Server, rebuild the index with a different fill factor (20–50 percent).
Filter ICON data.	If you are using Genesys Info Mart 7.5 in an environment with Interaction Concentrator 7.6, set the ICON data filtering options listed below; with these options, ICON writes the data only to those tables and row types that Genesys Info Mart actually uses. This will reduce the amount of data that is stored in IDB and improve the performance of Info Mart extraction, in addition to ICON performance improvements. The following set of ICON data filtering options work with Genesys Info Mart 7.5: [filter-data] call-history = 1 call-metrics = 1 gls-metrics = 1 ir-history = 1 udata-history-terminated = 1

Table 68: Release 7.5 Performance Tuning Guidelines (Continued)

Guideline	Description
Provision for large database volumes.	If you want to retain significant volumes in the database, you must make provisions at the database level to manage the Genesys Info Mart database. A volume of one million calls per day quickly leads to unacceptable performance from both the Load and Aggregate jobs, unless proper measures are put in place. Genesys recommends that you consult a Database Administrator (DBA) knowledgeable in data warehouses to analyze the projected growth of your Genesys Info Mart database—according to a size estimation spreadsheet—and that you plan the storage, partitioning, and any other optimization measures accordingly.
Use caution when adding indexes.	Genesys Info Mart load time is also affected by the number of indexes on the large fact tables. Therefore, use caution when adding indexes. In the performance tests, one table in particular took much longer than the others to load: `INTERACTION_SEGMENT_FACT`, which had eight indexes.
Partition the `Interaction_Segment_Fact` table by date.	Although the last two days of execution time for `Job_LoadGIM` were flat for the Oracle test, we can expect the general trend of increasing execution time to continue. The factor most influencing this time was the rebuilding of indexes on `INTERACTION_SEGMENT_FACT`. Because all the indexes were global, they all had to be rebuilt for the entire table with every run of `Job_LoadGIM`. Partitioning this table by `STD_TENANT_DATE_KEY` and making the indexes local would significantly flatten this performance trend. Each index would be rebuilt only for the data contained in that partition, rather than for the entire table.
Keep databases on separate servers.	These performance tests used a single database instance for both the Staging Area database and the Genesys Info Mart database. This topology did not significantly impact the test results, because only `Job_LoadRecent` would use both databases at the same time. In an actual customer deployment, you should locate the Staging Area database and the Genesys Info Mart database on separate database servers, so that the user reporting queries do not impact the intraday ETL jobs.
Optimize the storage subsystem.	Genesys recommends that you take special note of the storage subsystem of the database servers used in testing, because earlier tests indicated that sub-optimal database storage can easily become a critical performance bottleneck.

Table 68: Release 7.5 Performance Tuning Guidelines (Continued)

Guideline	Description
Use asynchronous merge.	The Genesys Info Mart application option `ir-merge-interval` specifies the time interval, in minutes, that the Genesys Info Mart is to periodically run the ICON IR Merge stored procedure. The stored procedure merges completed calls between switches monitored by the same ICON application. If you set the value to `0`, `Job_ExtractICON` calls this IR Merge procedure, and the extract queries are not executed until the IR Merge procedure completes. The test results included here used an `ir-merge-interval` of `0`. If you set the value greater than `0`, then the Genesys Info Mart Server calls the IR Merge procedure asynchronously for all the ETL jobs. This decreases the execution time of `Job_ExtractICON` and increases the though put of the ETL process. The recommended value of `ir-merge-interval` is `5`.
Allow sufficient database resources.	Be generous with the amount of RAM that you allocate to the database server buffer caches. Work with your DBA to tune the IDB, Staging Area and Genesys Info Mart database systems with production data volumes. These systems have different operational characteristics and need to be tuned differently.
Provide adequate space for transaction logs.	The daily `Job_LoadGIM` uses a single transaction to move a day's worth of data from an intraday table to the matching historical table. Depending on database platform and daily data volume, the transaction logs for such a large amount of data can consume gigabytes of disk space. Make sure that your environment allows for adequate space to handle these transaction logs.
Limit extraction data.	Use the Genesys Info Mart application options `limit-extract-data` and `extract-data-time-range-limit` to set the maximum amount of data that Genesys Info Mart will extract during a single ETL cycle. For more information about this data throttling process, see the Genesys Info Mart Deployment Guide for details.
Tune Genesys Info Mart Server memory.	For optimal performance, use the largest transformation buffer that the RAM on your machine can support. This buffer holds the objects that `Job_TransformGIM` will process on a single pass. To tune Genesys Info Mart Server memory: 1. Set the Java heap size parameter in the `gim_server` startup script to either 1500 MB for a 32-bit operating system or 2000 MB for a 64-bit operating system, provided that these amounts of RAM are available on the machine. 2. Set the Genesys Info Mart application option `transformation-buffer-size` to `5`. 3. Restart the server and begin running jobs. 4. If you receive a `Java Out of Memory` exception, set the option to `4`. 5. Repeat step 4 (decreasing the value of the option by one) until you no longer get a memory exception.

Genesys Info Mart 7.5 Database Size Estimation

Genesys Info Mart 7.5 reads data from the Interaction Concentrator, Stat Server, and GVP VAR databases, and writes data to its Staging Area and Genesys Info Mart databases.

The information in this section helps you estimate the size of the Staging Area and Info Mart data, depending on the size of your contact center and the length of time you want to retain data in the Genesys Info Mart.

The variables that have the greatest effect on the size of the Staging Area and Genesys Info Mart databases are the:

- Number of daily interactions in your contact center.
- Number of agents in your contact center.
- Complexity of your interaction flows.
- Amount of business data attached to interactions.

Notes:

- The Interaction Concentrator, Stat Server, and GVP VAR data size estimates are outside the scope of this guide.
- The data sizes are only estimates. Be sure to factor in extra space to accommodate variations in average data lengths.
- The Staging Area size estimate allows for up to three days of extracted data in case a network outage, database outage, or some other error temporarily prevents Genesys Info Mart from transforming or loading data.
- The Info Mart size estimate includes default indexes. Be sure to provide for additional storage for any indexes you add to the Info Mart to enhance query performance.

Genesys Info Mart provides the Genesys Info Mart Database Size Estimator, a spreadsheet in Microsoft Excel, to help you estimate the data size of the Staging Area and Info Mart databases.

- The spreadsheet requires that you fill in the basic information about the resources and daily interaction volumes for your contact center.
- You should also include the number of days you intend to store data in the Info Mart database.
- The spreadsheet then calculates the estimated initial size and daily delta size for the Staging Area and Info Mart databases.

The Genesys Info Mart 7.5 Database Size Estimator spreadsheet (`GIMSize_750.xls`) is available on the:

- Genesys Technical Support Website at `http://genesyslab.com/support`.
- Genesys Documentation Library DVD, which you can order by e-mail from Genesys Order Management at `orderman@genesyslab.com`.

Release 7.2 Guidelines

The following sections provide hardware sizing guidelines for environments with Genesys Info Mart 7.5:

- “Hardware Architectures in Release 7.2” on page 169, provides architecture samples for Genesys Info Mart release 7.2, for small and medium-size contact centers.
- “Genesys Info Mart 7.2 ETL Runtime Performance” on page 173, lists the factors that affect Genesys Info Mart 7.2 performance. It also lists sample performance measurements for small, medium, and large contact centers.
- “Release 7.2 ETL Performance with the Windows 2000/2003 Operating System” on page 175, lists specific performance measurements for Genesys Info Mart 7.2 as deployed on Windows 2000/2003 with different RDBMS options.
- “Genesys Info Mart 7.2 Database Size Estimation” on page 179, shows you how to estimate the size of the Staging Area and the Genesys Info Mart databases for release 7.2.

The information is organized based on the size of your contact center, as a function of the daily number of interactions and the number of agents. Table 46 provides the definition of the contact center size, as applicable to release 7.2 of Genesys Info Mart.

Table 69: Contact Center Size for Release 7.2

Size Category	Daily Number of Interactions	Number of Agents
Small	Fewer than 50,000	Fewer than 1,000
Medium	Fewer than 100,000	Fewer than 2,000
Large	More than 100,000	More than 2,000

Note: Large contact centers require special planning because of their variability and complexity. Contact Genesys Technical Support if you want to plan a Genesys Info Mart deployment for a large contact center.

Hardware Architectures in Release 7.2

The examples in this section are organized by the size of your contact center and your choice of operating system.

The hardware architecture you select for your Genesys Info Mart 7.2 deployment depends primarily on:

- The size of your contact center in relation to the number of daily interactions and the number of agents.
- The time available for Genesys Info Mart to extract, transform, and load data.
- Your choice of operating system.

To minimize the time it takes for Genesys Info Mart 7.2 to extract, transform, and load data, Genesys strongly recommends that you run the Data Integrator Job Server and the Staging Area database server in different machines.

Notes:

- Hardware recommendations for the Configuration Server, Call Concentrator, Interaction Concentrator, Stat Server, and Genesys Info Mart RDBMS servers are outside the scope of this guide.
- The type of hardware you deploy depends on whether this database is used directly for reporting (and the number of concurrent users), or as a secondary staging area to hold the data until it can be uploaded to a data warehouse.

Small Contact Center—Windows—Release 7.2

Figure 70 depicts hardware architecture of Genesys Info Mart release 7.2 for a small-size contact center on a Windows operating system.

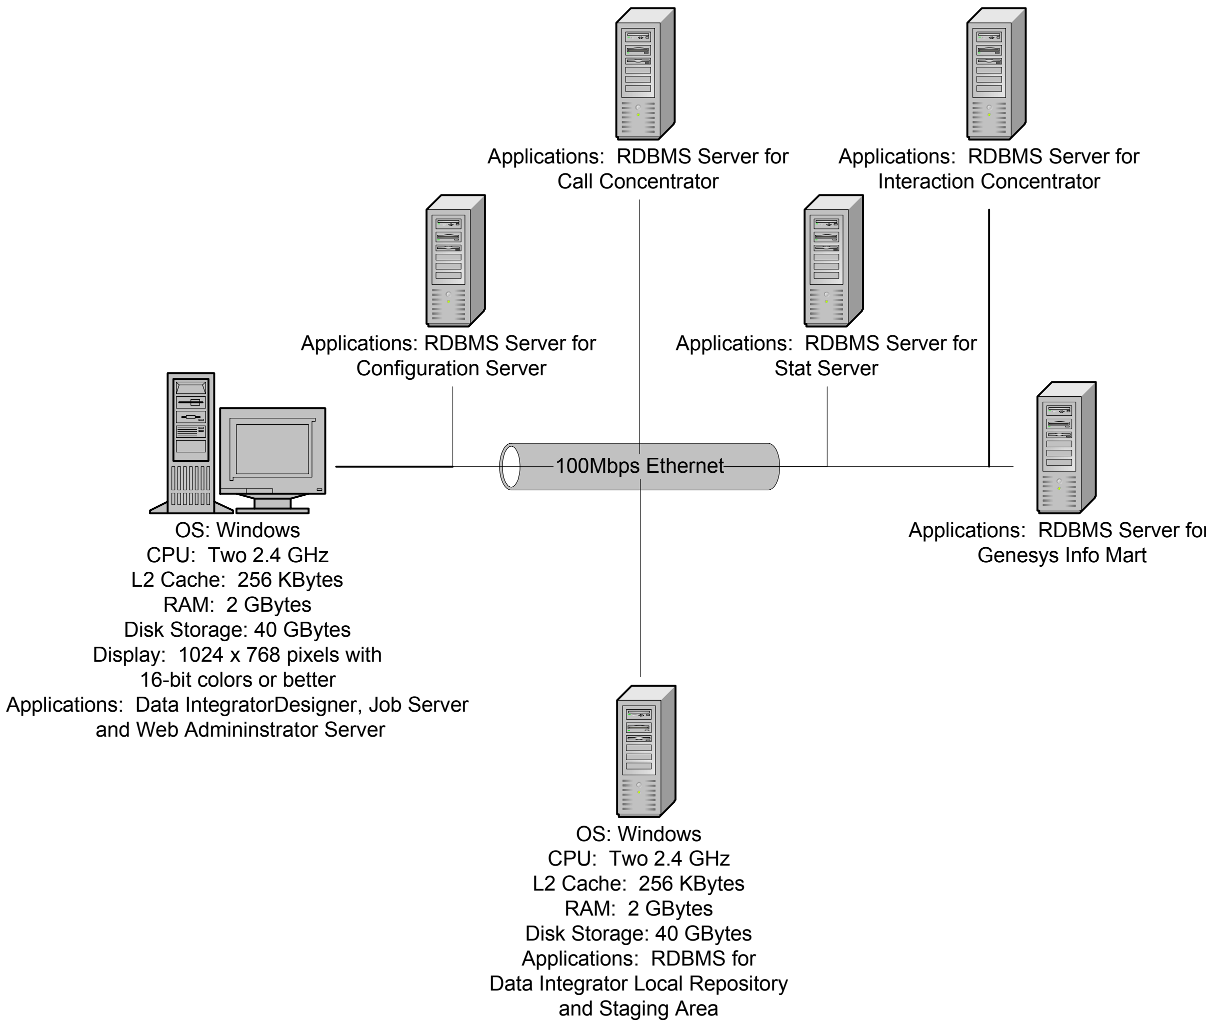

Figure 70: Small-Size Contact Center—Windows—Release 7.2

Medium Contact Center—Windows—Release 7.2

Figure 71 depicts hardware architecture of Genesys Info Mart release 7.2 for a medium-size contact center on a Windows operating system.

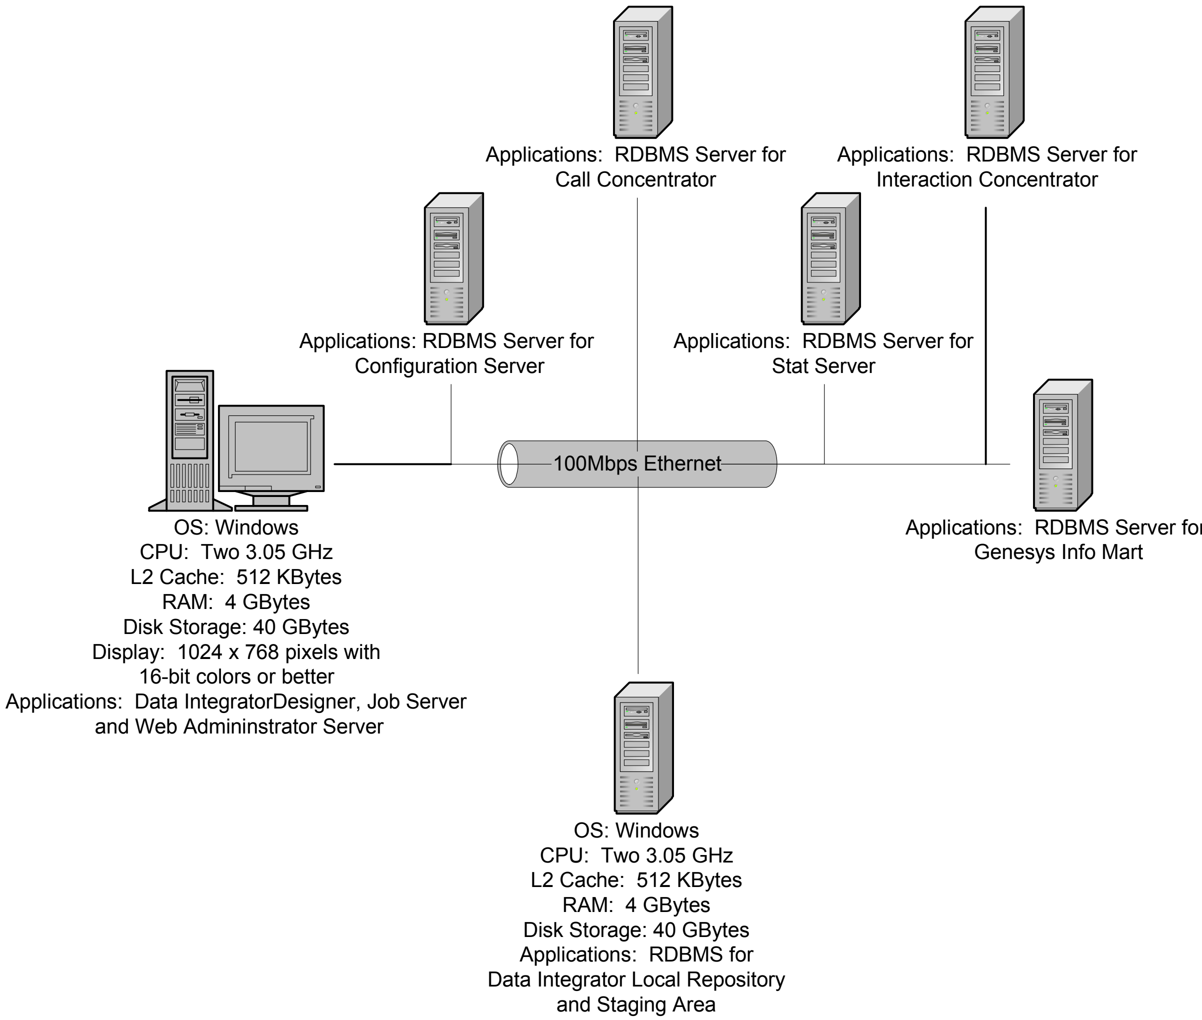

Figure 71: Medium-Size Contact Center—Windows—Release 7.2

Note: Genesys strongly recommends that you use a Staging Area RDBMS disk subsystem that contains multiple SCSI disk drives. Storing RDBMS logs, indexes and table data on different disk drives improves Extraction, Transformation, and Loading (ETL) performance by reducing I/O contention.

Small or Medium Contact Center—UNIX—Release 7.2

Figure 72 depicts hardware architecture of Genesys Info Mart release 7.2 for a small- or medium-size contact center on a UNIX operating system. In this architecture, Data Integrator Designer and Job Server run on different machines, since Designer runs on Windows only.

To obtain the best Data Integrator Job Server performance on UNIX, Genesys recommends selecting a machine that has the most powerful processors, rather than one with multiple slower processors. For example, a machine with two 990 MHz processors performs better than one with four 440 MHz processors.

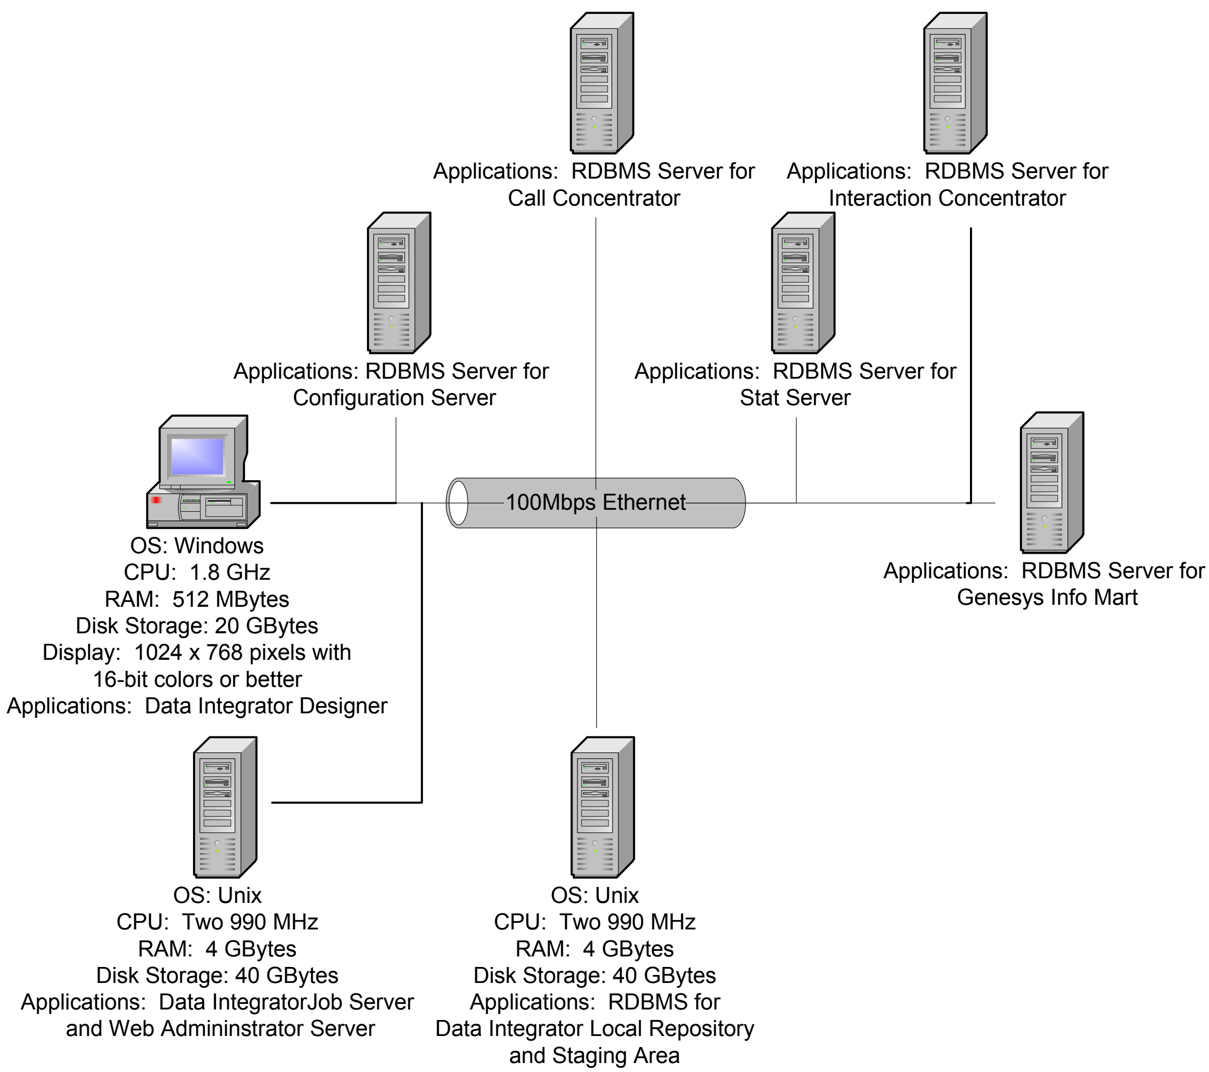

Figure 72: Small- or Medium-Size Contact Center—UNIX—Release 7.2

Note: Genesys strongly recommends that you use a Staging Area RDBMS disk subsystem that contains multiple SCSI disk drives. Storing RDBMS logs, indexes and table data on different disk drives improves Extraction, Transformation, and Loading (ETL) performance by reducing I/O contention.

Genesys Info Mart 7.2 ETL Runtime Performance

The primary factors affecting Extraction, Transformation, and Loading (ETL) performance in release 7.2 are as follows:

- The number of daily interactions in your contact center.
- The number of agents in your contact center.
- The complexity of your interaction flows.
- The amount of business data attached to interactions.
- The hardware on which the Data Integrator Job Server is running, primarily the CPU speed and available memory.
- The hardware on which the Staging Area RDBMS server is running, primarily the CPU speed, disk speed, and available memory.
- Tuning of the Staging Area database.
- The speed of the network connections between components.

Genesys ran several performance tests for select combinations of the following variables:

- Operating system (Windows 2000/2003)
- RDBMS (Oracle, MS SQL Server, DB2)
- Call volume (25,000, 50,000, 100,000, 250,000)

Call volumes between 25,000 and 50,000 calls per day are representative of small contact centers, while call volumes up to 100,000 calls per day are representative of medium-sized contact centers. Call volumes greater than 100,000 calls per day are representative of large contact centers.

The performance tests use the following call flow that yields six interaction segments per interaction:

1. A new inbound call arrives at an ACD queue.
2. The ACD queue diverts the call to an IVR port. The IVR application collects digits and attaches data representing the customer ID and the requested service type. The IVR application mute-transfers the call to a routing point.

3. The routing strategy loaded at the routing point uses the customer ID to determine the customer segment. Then it uses the requested service type and customer segment to select a target agent group, and attaches data representing the selected target. Finally, the routing strategy routes the call to an agent in the target agent group.
4. The agent answers the call, handles the interaction, attaches data representing a business result and a case ID, and finally releases the call.

The call flows include several Genesys-defined attached data items, but do not include any user-defined attached data.

The performance measurements focus on the time it takes to extract, transform, load, and aggregate data from Call Concentrator and Stat Server data sources, because processing from these sources takes the most time. The source data does not include:

- contact center configuration changes extracted from Configuration Server.
- user-defined attached data extracted from Call Concentrator.
- voice reason codes extracted from Stat Server.
- Outbound Contact Solution data extracted from Interaction Concentrator.

You can expect the ETL time to increase if source data is present for the subject areas listed above. This is especially true for Outbound Contact Solution data, for which the data volume is relatively high.

Table 70 shows the number of source rows extracted for the call flow that yields six interaction segments per interaction.

Table 70: Total Rows Extracted (Six Interaction Segments per Interaction)

Source Database	Source Table	Number of Rows Extracted			
		25,000 Calls	50,000 Calls	100,000 Calls	250,000 Calls
Call Concentrator	SCDR	77,748	155,829	310,570	780,995
	EVREFEX	75,243	150,559	300,571	752,239
	AREC	300,020	600,031	1,200,040	3,000,133
Stat Server	STATUS	124,095	195,780	334,066	937,710
	LOGIN	5,400	5,400	5,400	5,400
Total		582,506	1,107,599	2,150,647	5,476,477

Note: The number of rows extracted per day from the Configuration Server database is not included because the data volume from this data source is expected to be very low compared to the data volume from the Call Concentrator and Stat Server databases.

Table 71 shows the number of Genesys Info Mart fact table rows loaded for the call flow that yields six interaction segments per interaction.

Table 71: Total Number of Rows Loaded (Six Interaction Segments Per Interaction)

Genesys Info Mart Tables	Number of Rows Loaded			
	25,000 Calls	50,000 Calls	100,000 Calls	250,000 Calls
INTERACTION_SEGMENT_FACT	150,000	300,000	600,000	1,499,520
VOICE_SEG_FACT_EXT	150,000	300,000	600,000	1,499,520
INTERACTION_FACT	25,000	50,000	100,000	249,921
VOICE_IXN_FACT_EXT	25,000	50,000	100,000	249,921
RESOURCE_SESSION_FACT	2,700	2,700	2,700	2,700
RESOURCE_STATE_FACT	96,398	168,911	317,621	909,578
Total	449,098	871,611	1,720,321	4,411,160

Note: The numbers of Genesys Info Mart 7.2 dimension table rows, and configuration fact table rows, loaded per day are not included. This is because the expected data volume for new and changed rows is expected to be very low compared to the number of interaction and resource fact table rows.

Release 7.2 ETL Performance with the Windows 2000/2003 Operating System

This section presents Extraction, Transformation, and Loading (ETL) performance measurements for Genesys Info Mart 7.2 as deployed on Windows 2000/2003 with different RDBMS options. The hardware architecture used for all these tests is shown in Figure 73.

In this figure, the Staging Area RDBMS disk subsystem contains six SCSI disk drives in two RAID 0 configurations. The RDBMS data files are striped across one four drive RAID 0 array. The RDBMS log files are striped across another two drive RAID 0 array. For the Data Integrator Job Server, the `Global_DOP` option was set to `8` to fully utilize both CPUs. The `SkipValidationAtStartup` option was set to `1` to reduce the time that it takes to load the ETL jobs. The `Enable_Statistics` option was set to `FALSE` to reduce job overhead.

For more information on setting these options, see the section on “Data Integrator Performance Tuning,” in the *Genesys Info Mart 7.2 Deployment Guide*.

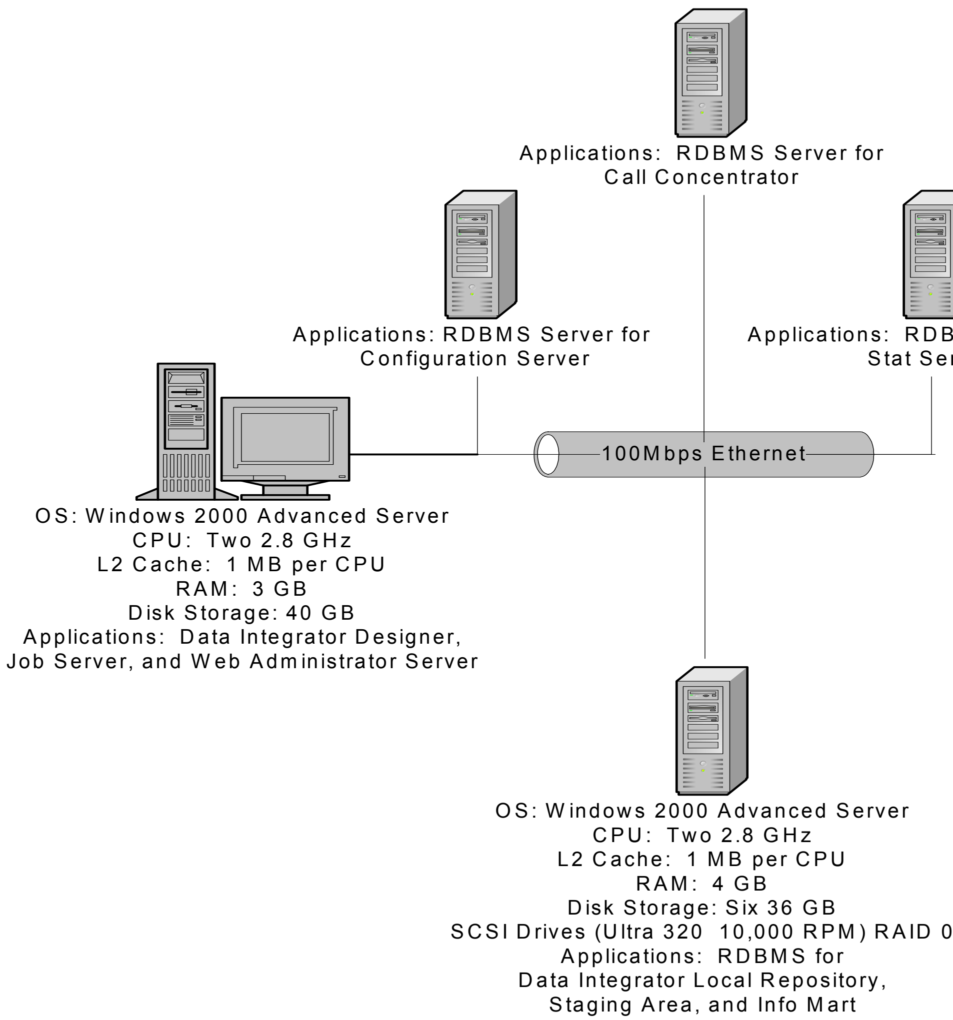

Figure 73: Release 7.2 Hardware Architecture for Windows 2000/2003 Performance Tests

Note: The best performance measurements were obtained with hyper-threading disabled. Genesys recommends that you disable hyper-threading on the Job Server machine.

Release 7.2 Performance with Windows 2000/2003 and Oracle

The hardware architecture used to measure the release 7.2 performance is depicted in Figure 73 on page 176. See also the adjacent Note about disabling hyperthreading.

Table 72 shows the time to extract Call Concentrator and Stat Server data for the different call volumes, with 6 segments per interaction.

Table 72: Extraction Time (Minutes): Oracle, 6 Segments per Interaction

Number of Calls	JOB_ExtractCCON	JOB_ExtractSS
25,000	1	1
50,000	2	1
100,000	3	2
250,000	8	3

Table 73 shows the time to transform and load Call Concentrator and Stat Server data for the different call volumes, with 6 segments per interaction.

Table 73: Transformation and Loading Time (Minutes): Oracle, 6 Segments per Interaction

Number of Calls	JOB_ TransformGIM	JOB_Load Recent	JOB_ LoadRecent with intraday aggregation	JOB_Load GIM	JOB_ Aggregate GIM
25,000	25	3	4	2	1
50,000	43	4	5	2	2
100,000	79	4	7	3	3
250,000	205	6	17	5	6

Release 7.2 Performance with Windows 2000/2003 and Microsoft SQL Server

The release 7.2 hardware architecture used to measure the performance is depicted in Figure 73 on page 176. See also the adjacent Note about disabling hyperthreading.

Table 74 shows the time to extract Call Concentrator and Stat Server data for the different call volumes, with 6 segments per interaction.

Table 74: Extraction Time (Minutes): Microsoft SQL Server, 6 Segments per Interaction

Number of Calls	JOB_ExtractCCON	JOB_ExtractSS
25,000	1	1
50,000	2	1
100,000	3	1
250,000	6	2

Table 75 shows the time to transform and load Call Concentrator and Stat Server data for the different call volumes, with 6 segments per interaction.

Table 75: Transformation and Loading Time (Minutes): Microsoft SQL Server, 6 Segments per Interaction

Number of Calls	JOB_ TransformGIM	JOB_ LoadRecent	JOB_ LoadRecent with intraday aggregation	JOB_LoadGIM	JOB_ AggregateGIM
25,000	31	3	4	1	1
50,000	54	4	6	2	1
100,000	101	6	9	2	2
250,000	254	12	16	5	3

Release 7.2 Performance with Windows 2000/2003 and DB2

The release 7.2 hardware architecture used to measure the performance is depicted in Figure 73 on page 176. See also the adjacent Note about disabling hyperthreading.

Table 76 shows the time to extract Call Concentrator and Stat Server data for the different call volumes, with 6 segments per interaction.

Table 76: Extraction Time (Minutes): DB2, 6 Segments per Interaction

Number of Calls	JOB_ExtractCCON	JOB_ExtractSS
25,000	1	1
50,000	2	1

Table 76: Extraction Time (Minutes): DB2, 6 Segments per Interaction (Continued)

Number of Calls	JOB_ExtractCCON	JOB_ExtractSS
25,000	1	1
100,000	3	1
250,000	6	2

Table 77 shows the time to transform and load Call Concentrator and Stat Server data for the different call volumes, with 6 segments per interaction.

Table 77: Transformation and Loading Time (Minutes): DB2, 6 Segments per Interaction

Number of Calls	JOB_ TransformGIM	JOB_ LoadRecent	JOB_ LoadRecent with intraday aggregation	JOB_LoadGIM	JOB_ AggregateGIM
25,000	20	3	4	2	1
50,000	34	4	5	2	2
100,000	62	5	7	3	2
250,000	156	10	14	5	8

Genesys Info Mart 7.2 Database Size Estimation

Genesys Info Mart 7.2 reads data from Genesys Configuration Server, Call Concentrator, Interaction Concentrator, and Stat Server databases, and writes data to its Staging Area and Info Mart databases.

The information in this section helps you estimate the size of the Staging Area and Info Mart data, depending on the size of your contact center and the length of time you want to retain data in the Genesys Info Mart. The estimates include only raw data: they do not include database overhead.

The variables that have the greatest effect on the size of the Staging Area and Genesys Info Mart databases are the:

- Number of daily interactions in your contact center.
- Number of agents in your contact center.
- Complexity of your interaction flows.
- Amount of business data attached to interactions.

Notes:

- The Configuration Server, Call Concentrator, Interaction Concentrator, and Stat Server data size estimates are outside the scope of this guide.
- The data sizes are only estimates. Be sure to factor in extra space to accommodate variations in average data lengths and database overhead.
- The Staging Area size estimate allows for up to three days of extracted data in case a network outage, database outage, or some other error temporarily prevents Genesys Info Mart from transforming or loading data.
- The Info Mart size estimate includes default indexes. Be sure to provide for additional storage for any indexes you add to the Info Mart to enhance query performance.

Genesys Info Mart provides the Genesys Info Mart Database Size Estimator, a spreadsheet in Microsoft Excel, to help you estimate the raw data size of the Staging Area and Info Mart databases.

- The spreadsheet requires that you fill in the basic information about the resources and daily interaction volumes for your contact center.
- You should also include the number of days you intend to store data in the Info Mart database.
- The spreadsheet then calculates the estimated initial size and daily delta size for the Staging Area and Info Mart databases.

The Genesys Info Mart 7.2 Database Size Estimator spreadsheet (`GIMSize_720.xls`) is available on the:

- Genesys Technical Support Website at `http://genesyslab.com/support`.
- Genesys Documentation Library CD, which you can order by e-mail from Genesys Order Management at `orderman@genesyslab.com`.

Chapter

6 Genesys Interactive Insights Performance

In the 7.6 release, Genesys Interactive Insights introduces 17 reports that you can run to summarize the inbound, call-related data stored in your Info Mart 7.6 database. Genesys has conducted a number of laboratory tests on the Interactive Insights reports to assess how long it takes to create or open a report instance under varying conditions such as different database size, different RDBMSs, different numbers of object selections for the hourly and daily reports, and different number of generated pages. The results of these tests can help you determine how Interactive Insights might perform in your own environment.

This chapter describes the environment in which the tests were run and summarizes the results in the following sections:

Environment

Hardware/Software Used

All tests were run with BusinessObjects Enterprise XI 3.0 and Interactive Insights 7.6, operating on an Intel Xeon box with the following specifications:

- 2 x Intel Xeon 3.2 GHz/2 MB
- 4 GB RAM
- 2 HDD x 80 GB SATA

- Dual GigEthernet

The tests were also run with Genesys Info Mart (GIM) 7.6 and either Oracle 10g or Microsoft SQL Server 2005, on an Intel Clovertown Quad Xeon box with the following specifications:

- 2 x Intel Clovertown Quad Xeon X5365 3.0 GHz/L2=2x4 MB
- 4 GB RAM
- 16 x SCSI HDD x 146 GB x 15K RPM
- Dual GigEthernet

The operating system running on both boxes was Microsoft Windows Server 2003 Enterprise edition. Except where noted, GIM ETL processes were scheduled to run in such a way that they did not compete for system resources with the running of Interactive Insights reports.

Call Flow Model

The call flow scenario for these tests simulated 4 to 6 inbound calls per second entering queue `81001` with 40 percent of them directly routed to agent DNs or IVR ports that were configured as handling resources. The remaining 60 percent were routed through routing point `22002` to one of four ACD queues before distribution to agent DNs, as shown in Figure 74. Additional inbound call flows entered other queues at a low rate, to simulate abandoned and other interactions.

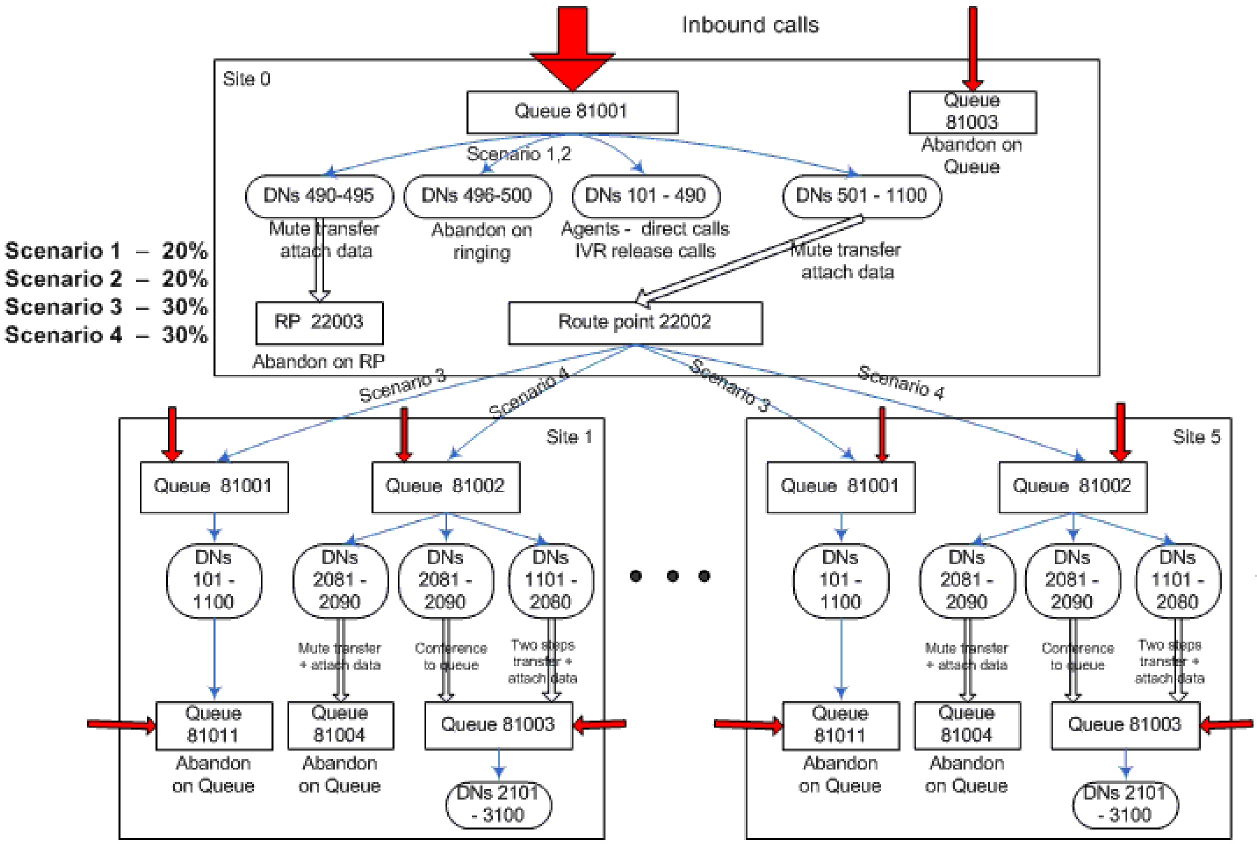

Figure 74: Simulated Call Flows

Attached Data Description

User data was affiliated with each interaction to associate a customer segment, service type and subtypes, business result, and/or service objective, and to apply one of ten string user data keys. This configuration enables data for the `Call Volume` reports to be populated.

How Time Was Measured

Opening a report consists of refreshing the report's data to re-query and re-retrieve data from the Info Mart, as opposed to displaying the results of an already generated report instance for which data was stored in the report's cube. (Refer to Business Objects documentation for information about the content of a report cube.) The time required to open a report begins when the `Run` button is clicked in the `User Prompt Input` area of a report, and continues until the report's results are displayed on-screen. Web Intelligence's internal timer provided the official measurement; the results of this timer are displayed the next time the query is run, in the `Running Query` message box shown in Figure 75. The level of precision available for this measurement is seconds.

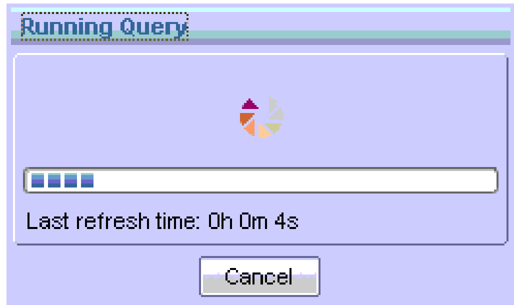

Figure 75: Web Intelligence Running Query Reflects Last Refresh Time

For scheduled reports, two time measurements were taken for the performance results gathered—that of report creation and that of report display. The time spent displaying a report was gathered in the same manner as described above. The time spent creating a report was gathered by viewing the details of the report's success history within Web Intelligence. The information about start and end times is recorded on the `HistoryStatus` page shown in Figure 76.

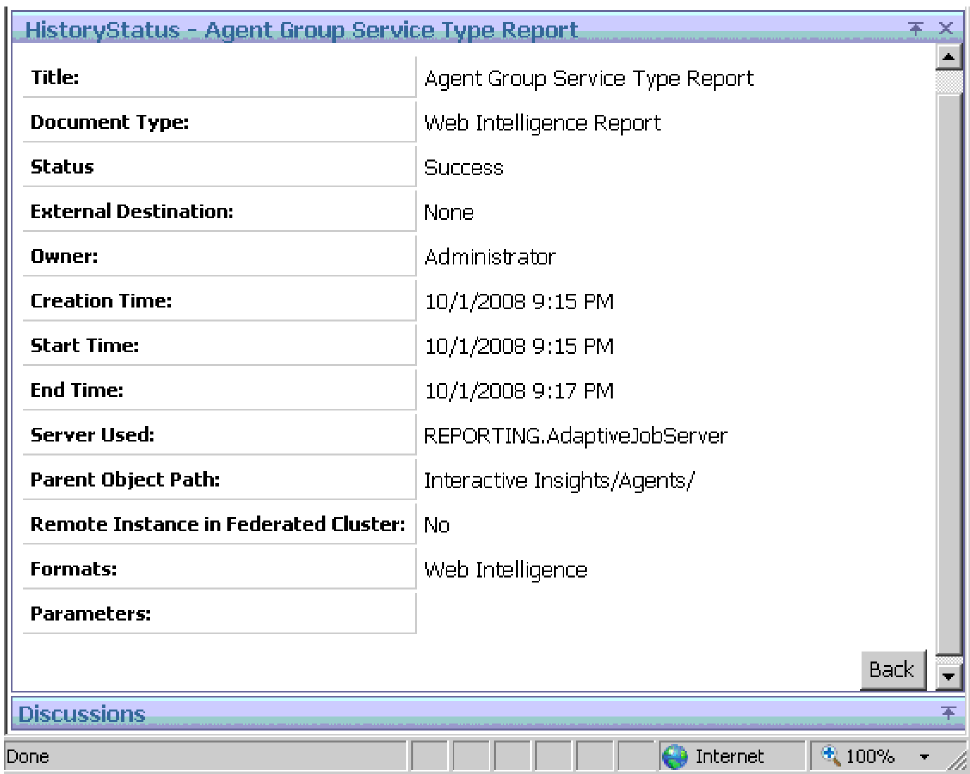

Figure 76: HistoryStatus Page Shows Minutes As Unit of Precision

The level of precision available for report creation measurements is minutes; therefore, a span of time is provided in the results. In Figure 76, for example, the performance results in this chapter would report two to three minutes for the time required to create the `Agent Group Service Type Report`.

Configuration Objects Used

The objects selected for the Agent reports were three different agent groups: AG1, AG2, and AG3, comprised of 300, 500, and 800 agents, respectively. The objects selected for the Call Volume reports were three service types or service subtypes: T1, T2, and T3. Finally, different report runs for the queue-type reports used selections of:

- One queue object, queue 81003—abbreviated 1Q.
- Two queue objects, queues 81003 and 81004— abbreviated 2Q.
- Three queue objects, queues 81003, 81004, and 81011—abbreviated 3Q.

Report Performance for a Microsoft SQL Info Mart

Tables 78 through 49 show the amount of time, in second(s), that was required to open each Interactive Insight report given a Genesys Info Mart (GIM) 7.6 source on a Microsoft SQL Server 2005 RDBMS. Database size was 61 gigabytes, the data amassed over 18 days, and none other than the requisite processes were simultaneously competing for system resources while the reports were running. No GIM 7.6 processes were simultaneously running.

Table 78 shows the time it took to run and open the Interactive Insights reports, which provided hourly breakdowns of the results for the 18-day range over which data was collected.

Table 78: Results of Hourly Reports on Microsoft SQL

Agent Reports		**Time to Open (s)**		
		AG1	**AG2**	**AG3**
Agent ACW Report	Time	2	2	3
	Pages	37	61	98
Agent Group Inbound Call Handling Report	Time	3	3	4
	Pages	2	2	2
Agent Inbound Call Handling VQ Report	Time	5	3	5
	Pages	603	690	690
Agent Inbound Utilization Report	Time	5	3	3
	Pages	61	101	161
Agent Interval Based Report	Time	5	5	7
	Pages	101	167	267
Agent Not Ready Reason Code Report	Time	60	60	60
	Pages	43	72	115
Agent Not Ready Report	Time	5	5	8
	Pages	155	233	276
Agent Service Report	Time	6	6	6
	Pages	3	3	3
Daily Agent Login-Logout Report	Time	2	2	2
	Pages	185	305	485
Daily Agent State Detail Report	Time	13	13	13
	Pages	57	89	124
Business Result Reports		**T1**	**T1,T2**	**T1-T3**
Call Volume Service Subtype Report	Time	3	5	5
	Pages	2	3	4
Call Volume Service Report	Time	3	3	3
	Pages	1	1	1
Queue Reports		**1Q**	**2Q**	**3Q**
Abandon Delay Report	Time	4	4	5
	Pages	1	1	2
Inbound Voice Traffic Group Report	Time	3	4	4
	Pages	1	1	1
Inbound Voice Traffic Report	Time	3	3	3
	Pages	1	1	1
Queue-Virtual Queue Summary Report	Time	3	3	3
	Pages	1	1	1
Speed of Answer Report	Time	3	3	4
	Pages	1	1	1

Table 79 provides these results in daily breakdowns of the same data over the same date range.

Table 79: Results of Daily Reports on Microsoft SQL

Agent Reports		**Time to Open (s)**		
		AG1	**AG2**	**AG3**
Agent Group Inbound Call Handling Report	Time	60	60	65
	Pages	5	5	5
Agent Inbound Call Handling VQ Report	Time	56	56	60
	Pages	162	269	424
Agent Inbound Utilization Report	Time	120	120	125
	Pages	281	462	741
Agent Not Ready Reason Code Report	Time	62	62	65
	Pages	18	18	21
Business Result Reports		**T1**	**T1-T2**	
Call Volume Service Subtype Report	Time	352	680	
	Pages	3	5	
Call Volume Service Report	Time	356	690	
	Pages	2	3	
Queue Reports		**1Q**	**2Q**	**3Q**
Abandon Delay Report	Time	50	236	415
	Pages	2	2	3
Inbound Voice Traffic Group Report	Time	190	320	410
	Pages	1	1	1
Inbound Voice Traffic Report	Time	315	575	
	Pages	2	2	
Queue-Virtual Queue Summary Report	Time	220	480	
	Pages	2	3	
Speed of Answer Report	Time	220	580	
	Pages	2	3	

Table 80 shows the time it took to create report instances and to open the reports for two scheduled reports.

Table 80: Results of Scheduled Hourly Reports on Microsoft SQL

Agent Reports for AG3	**Time to Create Report**	**Time to Open**
Agent Not Ready Reason Code Report	1– 2 minutes	2 seconds
Daily Agent Login-Logout Report	1– 2 minutes	3 seconds

Report Performance for a Microsoft SQL Info Mart

Tables 78 through 83 show the amount of time, in second(s), that was required to open each Interactive Insight report given a Genesys Info Mart (GIM) 7.6 source on a Microsoft SQL Server 2005 RDBMS. Database size was 61 gigabytes, the data amassed over 18 days, and none other than the requisite processes were simultaneously competing for system resources while the reports were running. No GIM 7.6 processes were simultaneously running.

Table 78 shows the time it took to run and open the Interactive Insights reports, which provided hourly breakdowns of the results for the 18-day range over which data was collected. Table 79 provides these results in daily breakdowns of the same data over the same date range. Table 83 shows the time it took to create report instances and to open the reports for two scheduled reports.

Table 81: Results of Hourly Reports on Microsoft SQL

Agent Reports		**Time to Open (s)**		
		AG1	**AG2**	**AG3**
Agent ACW Report	Time	2	2	3
	Pages	37	61	98
Agent Group Inbound Call Handling Report	Time	3	3	4
	Pages	2	2	2
Agent Inbound Call Handling VQ Report	Time	5	3	5
	Pages	603	690	690
Agent Inbound Utilization Report	Time	5	3	3
	Pages	61	101	161
Agent Interval Based Report	Time	5	5	7
	Pages	101	167	267
Agent Not Ready Reason Code Report	Time	60	60	60
	Pages	43	72	115
Agent Not Ready Report	Time	5	5	8
	Pages	155	233	276
Agent Service Report	Time	6	6	6
	Pages	3	3	3
Daily Agent Login-Logout Report	Time	2	2	2
	Pages	185	305	485
Daily Agent State Detail Report	Time	13	13	13
	Pages	57	89	124
Business Result Reports		**T1**	**T1,T2**	**T1-T3**
Call Volume Service Subtype Report	Time	3	5	5
	Pages	2	3	4
Call Volume Service Report	Time	3	3	3
	Pages	1	1	1

Table 81: Results of Hourly Reports on Microsoft SQL (Continued)

Queue Reports		1Q	2Q	3Q
Abandon Delay Report	Time	4	4	5
	Pages	1	1	2
Inbound Voice Traffic Group Report	Time	3	4	4
	Pages	1	1	1
Inbound Voice Traffic Report	Time	3	3	3
	Pages	1	1	1
Queue-Virtual Queue Summary Report	Time	3	3	3
	Pages	1	1	1
Speed of Answer Report	Time	3	3	4
	Pages	1	1	1

Table 82: Results of Daily Reports on Microsoft SQL

Agent Reports		Time to Open (s)		
		AG1	AG2	AG3
Agent Group Inbound Call Handling Report	Time	60	60	65
	Pages	5	5	5
Agent Inbound Call Handling VQ Report	Time	56	56	60
	Pages	162	269	424
Agent Inbound Utilization Report	Time	120	120	125
	Pages	281	462	741
Agent Not Ready Reason Code Report	Time	62	62	65
	Pages	18	18	21
Business Result Reports		**T1**	**T1-T2**	
Call Volume Service Subtype Report	Time	352	680	
	Pages	3	5	
Call Volume Service Report	Time	356	690	
	Pages	2	3	
Queue Reports		**1Q**	**2Q**	**3Q**
Abandon Delay Report	Time	50	236	415
	Pages	2	2	3
Inbound Voice Traffic Group Report	Time	190	320	410
	Pages	1	1	1
Inbound Voice Traffic Report	Time	315	575	
	Pages	2	2	
Queue-Virtual Queue Summary Report	Time	220	480	
	Pages	2	3	
Speed of Answer Report	Time	220	580	
	Pages	2	3	

Table 83: Results of Scheduled Hourly Reports on Microsoft SQL

Agent Reports for AG3	Time to Create Report	Time to Open
Agent Not Ready Reason Code Report	1– 2 minutes	2 seconds
Daily Agent Login-Logout Report	1– 2 minutes	3 seconds

Report Performance for an Oracle Info Mart

Two sets of tests were run and results gathered, to measure the time required to open reports given a relatively small-sized Info Mart (30 gigabytes) and a relatively medium-sized one (125 gigabytes) on an Oracle 10g RDBMS. For the smaller Info Mart, data amassed over a 7-day period with call flow volume of 4 to 6 calls per second. For the larger Info Mart, data amassed over a simulated 365-day period with the same call flow volume. The results are provided in "Results for a Small-Sized Oracle Info Mart" and "Results for a Medium-Sized Oracle Info Mart" on page 193 sections below.

Results for a Small-Sized Oracle Info Mart

The objects selected for the reports were the same as those selected for the Microsoft SQL tests on page 185. Table 84 shows an hourly breakdown of the time it took to run and open reports for the 7-day range over which data was collected. Table 85 shows a daily breakdown of the same data. Table 86 shows the time it took for a scheduled process to create a report instance for the `Agent Not Ready Reason Code Report` and open the report. Finally, Table 87 shows results, comparable to those in Table 84, when GIM ETL processes were running and competing for system resources.

Table 84: Results of Hourly Report Performance for a Small-Sized Oracle Info Mart

Agent Reports		**Time to Open (s)**		
		AG1	**AG2**	**AG3**
Agent ACW Report	Time	3	3	3
	Pages	133	218	347
Agent Group Inbound Call Handling Report	Time	4	4	4
	Pages	1	1	2
Agent Inbound Call Handling VQ Report	Time	4	4	4
	Pages	222	222	610
Agent Inbound Utilization Report	Time	5	5	5
	Pages	14	54	114
Agent Interval Based Report	Time	4	4	5
	Pages	198	327	521
Agent Not Ready Reason Code Report	Time	36	36	36
	Pages	38	63	101
Agent Not Ready Report	Time	3	3	4
	Pages	130	214	341
Agent Service Type Report	Time	6	6	6
	Pages	3	3	3
Daily Agent Login-Logout Report	Time	3	3	3
	Pages	110	167	301
Daily Agent State Detail Report	Time	3	3	3
	Pages	148	249	400
Business Result Reports		**T1**	**T1,T2**	**T1-T3**
Call Volume Service Subtype Report	Time	3	3	3
	Pages	2	3	4
Call Volume Service Type Report	Time	3	3	3
	Pages	1	1	1
Queues Reports		**1Q**	**2Q**	**3Q**
Abandon Delay Report	Time	3	3	3
	Pages	1	1	1
Inbound Voice Traffic Group Report	Time	3	3	3
	Pages	1	1	1
Inbound Voice Traffic Report	Time	3	3	3
	Pages	1	1	1
Queue-Virtual Queue Summary Report	Time	3	3	3
	Pages	1	1	1
Speed of Answer Report	Time	3	3	3
	Pages	1	1	1

Table 85: Results of Daily Report Performance for a Small-Sized Oracle Info Mart

Agent Reports		Time to Open (s)		
		AG1	**AG2**	**AG3**
Agent Group Inbound Call Handling Report	Time	18	18	22
	Pages	3	3	4
Agent Inbound Call Handling VQ Report	Time	5	5	5
	Pages	102	177	279
Agent Inbound Utilization Report	Time	10	10	14
	Pages	157	277	457
Agent Service Type Report	Time	8	8	8
	Pages	12	12	14
Business Result Reports		**1T**	**2T**	**3T**
Call Volume Service Subtype Report	Time	2	2	2
	Pages	4	8	11
Call Volume Service Report	Time	2	2	2
	Pages	1	2	2
Queue Reports		**1Q**	**2Q**	**3Q**
Abandon Delay Report	Time	2	2	2
	Pages	1	2	2
Inbound Voice Traffic Group Report	Time	2	2	3
	Pages	1	1	1
Inbound Voice Traffic Report	Time	3	3	3
	Pages	1	2	2
Queue-Virtual Queue Summary Report	Time	3	3	3
	Pages	1	2	2
Speed of Answer Report	Time	3	3	3
	Pages	1	1	1

Table 86: Results of Scheduled Hourly Reports on Oracle

Agent Reports for AG3	Time to Create	Time to Open
Agent Not Ready Reason Code Report	< =1 minute	2 seconds

Table 87: Results of Scheduled Hourly Reports with GIM Running

Agent Reports for AG1	Time to Open (s)	
Agent ACW Report	Time	8–12
	Pages	134
Daily Agent State Detail Report	Time	22–28
	Pages	161

Table 87: Results of Scheduled Hourly Reports with GIM Running (Continued)

Business Result Reports (All Service Types/Subtypes)	Time to Open (s)	
Call Volume Service Type Report	Time	2–4
	Pages	2
Queues Reports (All Queues)		
Inbound Voice Traffic Group Report	Time	2
	Pages	1
Queue-Virtual Queue Summary Report	Time	2–3
	Pages	2

Results for a Medium-Sized Oracle Info Mart

The objects selected for the reports were the same as those selected for the Microsoft SQL tests on page 185. Table 88 shows an hourly breakdown of the time it took to run and open reports for the simulated 300-day range over which data was collected with no GIM processes simultaneously running. Table 89 shows the same results in a daily breakdown. Table 90 shows the time to create and open one scheduled report, `Agent Not Ready Reason Code Report`. Finally, Table 91 shows the results when six users concurrent running reports. This table repeats the corresponding tests and results that were run in Table 88, to enable comparison between single and concurrent report runs.

Table 88: Results of Hourly Report Performance for a Medium-Sized Oracle Info Mart

Agent Reports		**Time to open (s)**		
		AG1	**AG2**	**AG3**
Agent ACW Report	Time	146	157	166
	Pages	124	244	327
Agent Group Inbound Call Handling Report	Time	12	15	16
	Pages	2	2	2
Agent Inbound Call Handling VQ Report	Time	45	52	74
	Pages	217	217	603
Agent Inbound Utilization Report	Time	101	111	117
	Pages	61	101	161
Agent Interval Based Report	Time	290	330	330
	Pages	186	365	490
Agent Not Ready Reason Code Report	Time	690	745	770
	Pages	61	101	161
Agent Not Ready Report	Time	146	158	170
	Pages	124	243	326
Agent Service Type Report	Time	56	64	73
	Pages	13	18	20
Daily Agent State Detail Report	Time	112	115	113
	Pages	91	151	200
Business Result Reports		**T1**	**T1,T2**	**T1-T3**
Call Volume Service Subtype Report	Time	2	2	2
	Pages	2	4	4
Call Volume Service Type Report	Time	2	2	2
	Pages	1	1	2
Queue Reports		**1Q**	**2Q**	**3Q**
Abandon Delay Report	Time	2	2	2
	Pages	1	1	1
Inbound Voice Traffic Group Report	Time	2	2	2
	Pages	1	1	1
Inbound Voice Traffic Report	Time	2	2	2
	Pages	1	1	1
Queue-Virtual Queue Summary Report	Time	2	2	2
	Pages	1	1	1
Speed of Answer Report	Time	2	2	2
	Pages	1	1	1

Table 89: Results of Daily Report Performance for a Medium-Sized Oracle Info Mart

Agent Reports		**Time to open (s)**		
		AG1	**AG2**	**AG3**
Agent Group Inbound Call Handling Report	Time	33	34	35
	Pages	2	2	2
Agent Inbound Call Handling VQ Report	Time	14	14	27
	Pages	262	262	728
Agent Inbound Utilization Report	Time	112	113	114
	Pages	76	126	201
Agent Service Type Report	Time	55	63	67
	Pages	13	18	20
Business Result Reports		**T1**	**T1,T2**	**T1-T3**
Call Volume Service Subtype Report	Time	2	2	2
	Pages	2	4	6
Call Volume Service Report	Time	2	2	2
	Pages	1	1	2
Queues Reports		**1Q**	**2Q**	**3Q**
Abandon Delay Report	Time	2	2	2
	Pages	1	1	1
Inbound Voice Traffic Group Report	Time	2	2	2
	Pages	1	1	1
Inbound Voice Traffic Report	Time	2	2	2
	Pages	1	1	1
Queue-Virtual Queue Summary Report	Time	2	2	2
	Pages	1	1	1
Speed of Answer Report	Time	2	2	2
	Pages	1	1	1

Table 90: Scheduled Daily Report on Oracle

Agent Report for AG3		**Time to Create**	**Time to Open**
Agent Not Ready Reason Code Report	Time	<= 12 minutes	2 seconds

Note: The daily results of the single report runs in Table 91 were generated based on data from a different day or span of days than the results reported in Table 88.

Table 91: Comparison of Daily Results for Concurrent and Noncurrent Report Users on Oracle

	Pages	Time to Open (s)	
		Single	Concurrent
Agent Reports for AG1			
Agent ACW Report	124	157	166–173
Agent Group Inbound Call Handling Report	2	7	7–9
Agent Inbound Call Handling VQ Report	42	45	101–107
Agent Inbound Utilization Report	61	81	88–94
Agent Interval Based Report	186	278	316–340
Agent Not Ready Reason Code Report	61	660	726–752
Agent Not Ready Report	124	140	179–182
Agent Service Type Report	3	20	22–23
Daily Agent State Detail Report	91	109	116–120
Business Result Reports for T1			
Call Volume Service Subtype Report	2	2	2
Call Volume Service Type Report	1	2	2
Queue Reports 1Q			
Abandon Delay Report	1	2	2
Inbound Voice Traffic Group Report	1	2	2
Inbound Voice Traffic Report	1	2	2
Queue/Virtual Queue Summary Report	2	1	1–2
Speed of Answer Report	2	2	1–2

Conclusions

As expected, the number of objects selected affects the time required to create and open a report. It generally takes longer to create a report than it does to open it. The larger the selection pool (the number of generated pages), the longer the report creation and open time. Reports take longer to open with concurrent usage than nonconcurrent (single) usage.

Chapter

7

Workforce Management

This chapter presents the following Workforce Management information:

Software Co-location Recommendations

Genesys makes the following recommendations:

- Co-locate the WFM Server and the WFM database on the same LAN segment.
- Co-locate the WFM Data Aggregator (and its backup, if any) on the same LAN segment as the WFM Server and the WFM database.

Note: In instances when a backup Data Aggregator must be connected to the WFM database over a WAN, performance will be sensitive to network bandwidth and latency. Insufficient network bandwidth or excessive network latency may result in the following behaviors:

- Excessive memory consumption on the Data Aggregator server, due to the server taking an extended amount of time to write data to the WFM database.
- WFM Web displaying outdated schedule state data in the real-time adherence view.

Small Contact Centers

Table 92 on page 198 and Table 93 on page 198 show the Workforce Management hardware requirements for small contact centers.

Table 92: Small Contact Centers (<500 Agents)—Server 1

Tomcat/WebSphere, WFM Web, WFM Server, WFM Data Aggregator	
Processor Type, Quantity, Speed	Dual Intel Xeon 2.4 GHz with 512MB L2 cache
Memory Size	4 GB RAM
Hard Disc Space	40 GB (Contact Genesys Technical Support website at http://genesyslab.com/support for specific requirements since 40 GB is just a baseline.)
Ports	Networking ports—TP Ethernet 1000/100 BASE-T card that works in full duplex mode
CD ROM	32 X

Table 93: Small Contact Centers (<500 Agents)—Server 2

WFM Builder, WFM Daemon	
Processor Type, Quantity, Speed	Single Intel P4 2.4 GHz
Memory Size	2 GB RAM
Hard Disc Space	40 GB (Contact Genesys Technical Support website at http://genesyslab.com/support for specific requirements since 40 GB is just a baseline.)
Ports	Networking ports—TP Ethernet 1000/100/10 BASE-T card that works in full duplex mode
CD ROM	32 x

Notes:

- Starting with WFM 7.2, there is no longer a separate WFM Server component for building reports. The report generating functionality is now integrated into the WFM Web component.
- When WFM Builder and the instance of WFM Web dedicated to report generation are deployed on the same machine, in some cases usage may overlap and cause delays for pending user requests. If this may be an issue, deploy WFM Builder and the WFM Web dedicated to report generation on separate machines.
- WFM Builder hardware requirements are dependent on the size of the schedules you are generating. See Figure 77 on page 203 for a complete chart of WFM Builder RAM requirements.

Client Workstation Requirements

Computers running the WFM client applications WFM Configuration Utility, WFM Database Utility, WFM Web for Supervisors, and WFM Web for Agents: at minimum, should meet these specifications:

- WFM Configuration Utility, WFM Database Utility, or WFM Web for Supervisors: Pentium IV, 2 GHz or greater (or its equivalent), with 512MB RAM.
- WFM Web for Agents: Pentium IV, 2 GHz or greater (or its equivalent), with 256MB RAM.

Medium Contact Centers

Table 94 shows the Workforce Management hardware requirement for medium contact centers.

Table 94: Medium Contact Centers (500-1000 Agents)—Server 1 (Sheet 1 of 2)

Tomcat/WebSphere & WFM Web	
Processor Type, Quantity, Speed	Dual Intel Xeon 2.4 GHz with 512 MB L2 cache
Memory Size	2 GB RAM
Hard Disc Space	40 GB (Contact Genesys Technical Support website at http://genesyslab.com/support for specific requirements since 40 GB is just a baseline.)

Table 94: Medium Contact Centers (500-1000 Agents)—Server 1 (Sheet 2 of 2) (Continued)

Tomcat/WebSphere & WFM Web	
Ports	2 TP Ethernet 1000/100 BASE-T cards in full duplex mode
CD ROM	32 x

Sizing for WFM Web

Note: Starting with WFM 7.2, there is no longer a separate WFM server component for building reports. The report-generating functionality is now integrated into the WFM Web component. Because report generation can be CPU-intensive, some contact centers may want to maintain a dedicated server for this process.
For the same reason, contact centers may not want the same instance of WFM Web used to generate reports to also serve up all agent and supervisor-facing user interfaces. Therefore, an instance of WFM Web may be deployed and dedicated to report-building only. Instructions for how to configure this may be found in the current *Workforce Management Administrator's Guide*.

Performance test results show that you can calculate the number of requests one WFM Web machine can support using the formula in the following section. The relationship between the number of users and the number of requests depends on what actions the users are performing. Most actions require at least two requests.

Notes:

- The formula calculates CPU usage. It assumes adequate memory resources for all users. If your memory resources are limited, this also affects the number of sessions you can have open simultaneously.
- This formula assumes that the WFM Web component is running on a dedicated machine, which houses only a supported version of Tomcat and is equivalent to a Wintel Web server, dual 3GHz processor with 2GB RAM. If you are running WFM Web in a WebSphere servlet container, contact Genesys Technical Support for sizing recommendations.
- This formula applies to an average load over a period of time. A server may be able to support higher numbers of users during short-term peak usage.

Formula

The formula to determine how many clients (either agents or supervisors) can be supported on a single one of these servers is:

```
Total_number_of_clients = (4500 -
(Number_of_simultaneous_adherence_views * 60 / refresh_time)) /
Average_number_of_requests_of_other_clients
```

- `refresh_time`—refresh time in seconds of the real-time agent-adherence view.
- `60`—number of requests per minute for one real-time agent-adherence view assuming a refresh rate of 2 seconds.
- `4500 = 75 * 60` —a number, empirically derived from tests, showing that a Wintel Web server, dual 3GHz processor with 2GB RAM supports about 75 real-time agent-adherence views.
- `Number_of_simultaneous_adherence_views` = the number of constantly running real-time agent adherence views. Note that this can be greater than the number of users. If 1 supervisor runs 2 views, it counts as 2 views, not 1.
- `Average_number_of_requests_of_other_clients` = the average number of all types of requests per minute for all users (except real-time agent adherence views) during the day. Here, a request refers to any time a user submits something from a Web page, navigates to a new view, gets data from the server, and so on.

Note: A more powerful computer would be able to support more than 75 views, a less powerful one, fewer. However, because this number is an estimate, based on test results, it should be used as a guideline rather than an absolute figure.

Example

- Assume that on average, each user, whether agent or supervisor, makes one nonreal-time agent-adherence request every 3 minutes during the day.
- Assume there are 30 supervisors, each running 1 real-time agent adherence view with a refresh rate of 2 seconds.

```
Total_number_of_clients = (4500 - (30 * 60)/2) / 0.333 = 10,811
```

Thus, in this type of environment, a dedicated Web server meeting the specifications described should be able to handle an average of 10, 811 users per day. During peak times, this number would be lower. For example, during peak times, users may access the system 150% as much as the average level. In this case, the number of supported users drops to 5,405. See Table 95 on page 202 and Table 96 on page 202.

The default value set in the refresh_time option is 2 seconds. If you increase this refresh rate, you can increase the number of real-time adherence users. For example, if you increase the refresh rate to 5 seconds, you increase the potential number of users by about 100 percent.

Table 95: Medium Contact Centers (500-1000 Agents)—Server 2

WFM Server, WFM Data Aggregator, WFM Daemon	
Processor Type, Quantity, Speed	Dual Intel Xeon 2.4 GHz with 512 MB L2 cache
Memory Size	2 GB RAM
Hard Disc Space	40 GB (Contact Genesys Technical Support website at http://genesyslab.com/support for specific requirements since 40 GB is just a baseline.)
Ports	2 TP Ethernet 1000/100 BASE-T cards in full duplex mode
CD ROM	32 x

Table 96: Medium Contact Centers (500-1000 Agents)—Server 3

WFM Builder	
Processor Type, Quantity, Speed	Dual Intel Xeon 2.4 GHz with 512 MB L2 cache
Memory Size	2 GB RAM
Hard Disc Space	40 GB (Contact Genesys Technical Support website at http://genesyslab.com/support for specific requirements since 40 GB is just a baseline.)
Ports	2 TP Ethernet 1000/100 BASE-T cards in full duplex mode
CD ROM	32 x

Note: WFM Builder hardware requirements are dependent on the size of the schedules you are generating. See Figure 77 on page 203 for a complete chart of WFM Builder RAM usage.

Client Workstation Requirements

Computers running the WFM client applications WFM Configuration Utility, WFM Database Utility, WFM Web for Supervisors, and WFM Web for Agents: at minimum, should meet these specifications:

- WFM Configuration Utility, WFM Database Utility, or WFM Web for Supervisors: Pentium IV, 2 GHz or greater (or its equivalent), with 512MB RAM.
- WFM Web for Agents: Pentium IV, 2 GHz or greater (or its equivalent), with 256MB RAM.

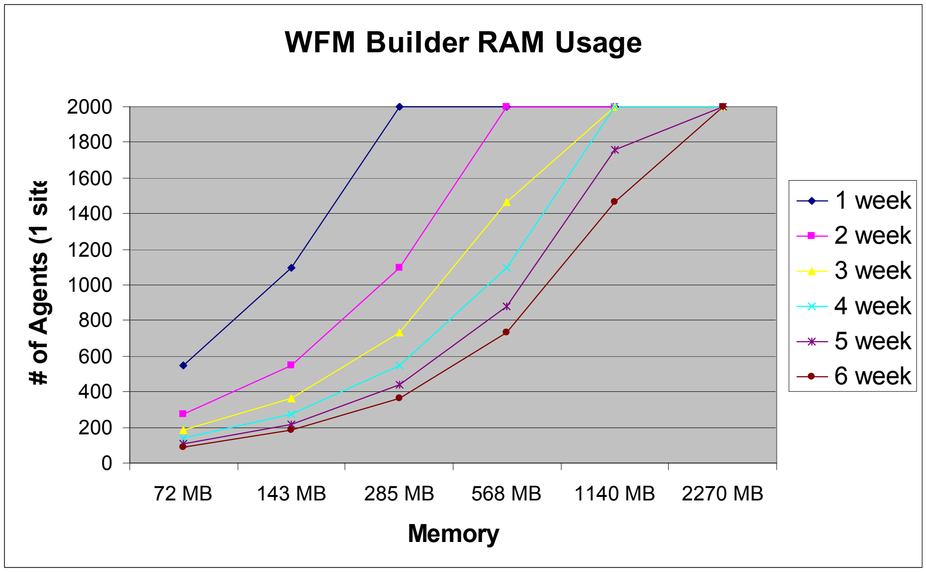

Figure 77: WFM Builder Memory Requirements

Chapter

8 Genesys Voice Platform 7.x

This chapter describes the hardware and capacity information for Genesys Voice Platform (GVP) 7.2, 7.5, and 7.6 releases. It also describes the performance criteria, machine setup, application profile, and typical call volume.

This chapter contains the following sections:

Capacity and Traffic Terminology

This section describes the capacity and traffic terminology for GVP.

Criteria

GVP has several capacity criteria. This section discusses three major performance measures, which are:

- Call Arrival Per Second
- Call Duration
- Caller Perceived Latency

Call Arrival Per Second

The Call Arrival Per Second (CAPS) is the indicator that shows the system traffic. Ten (10) CAPS means that there are 10 calls presented to GVP every second, which indicates busy traffic.

The CAPS is very similar to the legacy engineering term for telephony traffic called *Busy Hour Call Attempt,* or *Centum Call Seconds* (CCS). If the traffic has been known using the term CCS, then CAPS = CCS/36.

Call Duration

The Call Duration (CD) defines how long the call stays in the GVP system.

With the CAPS information and the CD, you can calculate the Erlang distribution for how many ports that are required for handling such traffic.

Further, if you assume an Erlang B traffic pattern, then the following is also used:

Erlang = CAPS x CD

The Erlang is the number of ports that are required to handle the traffic given CAPS and CD.

Caller Perceived Latency (CPL)

The CPL is defined as follows:

Assuming that the time the caller input ends (speech or DTMF key) is T1, and the time the next prompt starts is T2, then CPL = T2 – T1 (see Figure 78).

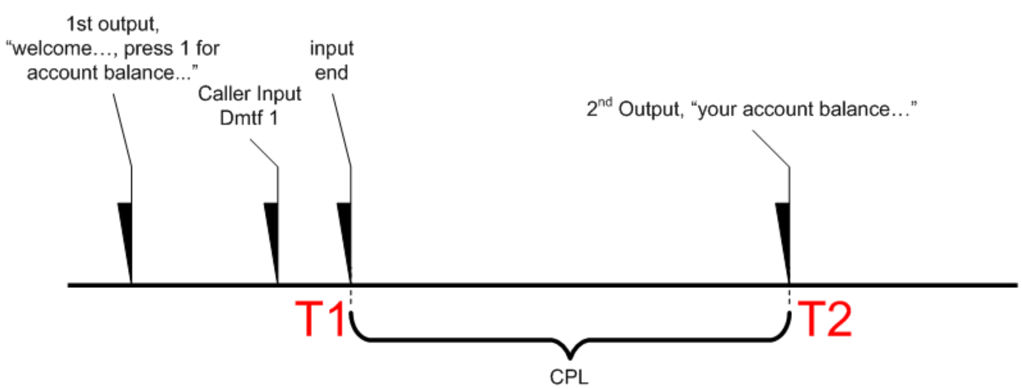

Figure 78: Caller Perceived Latency Example

The following factors have a direct impact on the CPL:

- Recognition engine.
- End of speech/DTMF timeout.
- Application pages, prompts, and grammar caching and fetching mechanism.
- The size of the application pages themselves.

- Call traffic, including call arrival rate and call duration.
- Speech density: percentage of time during the call that speech or DTMF detection is on and that the caller can give input.
- Speech recognition grammar size, and how often they are used in the application.
- Back end operation: how long it is needed to get information from the business layer (database query, and so on) and then present the result back to the caller.

Call Duration Increase

Instead of measuring individual CPLs within the application, sometimes the increase in the call duration is measured.

For a single call, the assumption is that the normal call duration (hang up by application) is `D1`. When the load increases, the call duration is expected to increase due to an increase of the CPLs within the application. Assume that for `x` simultaneous calls in the system, the measured call duration is `Dx`. The capacity measurement goal is:

`Dx/D1 = 120%`

When the 120% threshold is reached, the call volume `x` stops.

The CPL and call duration threshold will be used interchangeably for the following capacity results.

Error Rate

The error rate is the measurement that indicates how many calls did not finish the call flow during the performance load test. If a total of 1000 calls were placed, and five calls did not finish because of various reasons, the error rate is:

`5/1000=0.5%`

It is preferred to find a performance load that has the minimum error rate, which is preferred to be less than 0.1%.

GVP 7.2

This section provides the hardware sizing requirements for Genesys Voice Platform (GVP) 7.2.

Hardware and Operating Systems

Table 97 describes the hardware and operating system recommendations. This table shows the baseline requirement for a Intel processor based machine.

Table 97: Hardware and Operating System

Hardware Specification	Recommendation	Comments
Microsoft Windows		
CPU	Dual Xeon 3.6GHz or higher	Slower CPU such as Dual 2.4 or 3.0 is supported.
Memory	2 GB or more	
Network	GigaBit Ethernet	100MBit is also supported.
Storage	RAID 1 disk storage with at least 36 GB	
Operating System	Microsoft Windows 2003 Server, Standard Edition (32-bit)	Microsoft Windows 2000 Server is supported.
Solaris		
CPU	Sun Fire V240 with sparcv9 dual 1002 MHz processors or higher	
Memory	2 GB or more	
Network	GigaBit Ethernet	100MBit is also supported
Storage	RAID 1 disk storage with at least 36 GB	
OS	Solaris 9	

System Capacity

This section describes the capacity recommendations for GVP.

Criteria

The system capacity is defined as how many ports the GVP system can support that maximizes the machine's CPU and memory without significantly impacting the CPL, and without errors.

An increasing volume of call traffic that comes into the GVP system will incur the usage of additional ports, hence increasing the CPU and memory usage on the GVP system, and possibly delaying the CPL and generating errors.

The Genesys performance results show that the capacity has an average CPU usage of less than 75 percent, and all of the results have the CPL of less than three (3) seconds. These results show that memory is not an issue for capacity in the performance, and the CPU is always the impeding factor in determining the number of ports supported on a single GVP Server.

It is also important to note that performance is largely dependent on the application itself, its complexity, and its speech density. Genesys tests performance using applications of average complexity and speech density.

Note: The Media Resource Control Protocol (MRCP) server that supports the ASR application is always on a different machine from the GVP machine.

The recommended hardware in Table 97 was used during the capacity tests, which means that Dual 3.6GHz machines were used in the tests. Even though GVP supports Dual 2.4GHz and 3.0GHz machines, the performance number will be lower if lower speed machines are used.

TDM/Dialogic Capacity Result

The results shown in Table 98 are based on Genesys performance test results.

Table 98: TDM/Dialogic Capacity Result (Microsoft Windows OS only)

Application Type	Maximum CAPS	Ports	Comment
DTMF only	5	288 (12 T1)	CAPS=5 is the maximum that GVP can support
Simple ASR (digits collection)	5	192 (8 T1)	Might need multiple MRCP servers
Large Grammar ASR (40,000 items grammar)	5	144 (6 T1)	Might need multiple MRCP servers

VoIP Capacity Result

The results shown in Table 99 are based on Genesys performance test results.

Table 99: VoIP Capacity Result

Application Type	Maximum CAPS	Ports	Comment
Microsoft Windows			
DTMF only	5	300	CAPS=5 is the maximum that GVP can support

Table 99: VoIP Capacity Result (Continued)

Application Type	Maximum CAPS	Ports	Comment
Simple ASR	5	220	Might need multiple MRCP servers
Large Grammar ASR	5	220	Might need multiple MRCP servers
Solaris			
DTMF only	1	70	CAPS=1 is the maximum that GVP can support
Simple ASR	1	60	Might need multiple MRCP Servers
Large Grammar ASR	1	60	Might need multiple MRCP Servers

VoIP Capacity Result for Dual-Core Machine

The information in this section is provided to help you to understand the normal behavior of GVP on a typical dual-core dual machine.

Hardware used in testing (dual-core dual CPU)

- IBM x3550, Dual-Core Intel Xeon 5230 2.00GHz/1333MHz 2x2MB L2 Cache
- Two Processors
- 1GB Memory
- 15K RPM SAS HDD

Application: DTMF only, phone card refill application with different call duration

- 18 VoiceXML pages (including one application root page)
- 8 user DTMF input (maximum sequence length of 16 digits [credit card number])
- Call duration is approximately 73 seconds
- DTMF only, no ASR/TTS resource usage
- 6 VoiceXML pages that contain 13 simple java script functions (moderately ECMA script intensive)
- 77 unique audio files, 262 `<audio>` tags (high audio file-based prompt and file fetching intensive)

Application: ASR, tipping guide application

- 6 user inputs, including number, money amount, an item menu, and yes/no
- 3 mandatory voice inputs, while the other 3 are optional between DTMF or voice
- Moderate ECMA script inside the page
- Base call duration is approximately 65 seconds

The results shown in Table 100 are based on Genesys performance test results.

Table 100: VoIP Capacity on Dual-Core Dual-Processor Machine

Approximate DTMF Application Duration (seconds)	Maximum CAPS	Tested Ports	Comment
27	18.3	500	CPU ~87%
40	12.3	500	CPU ~80%
85	5.9	500	CPU ~77%
Approximate ASR Application Duration (seconds)	**Maximum CAPS**	**Tested Ports**	**Comment**
65	3	200	CPU ~42% Error rate > 0.1% if CAPS goes higher

Capacity versus Call Duration

When the CAPS is greater than five, GVP still supports an application running on the platform; however, the port density might change.

In order to obtain the capacity relationship between call duration and CAPS, Genesys ran tests with the DTMF-only performance result shown in the preceding example application. Genesys added or shortened the call durations to evaluate the maximum CAPS, and then ran the same performance test to check the maximum CAPS.

Figure 79 highlights the duration summary of the application.

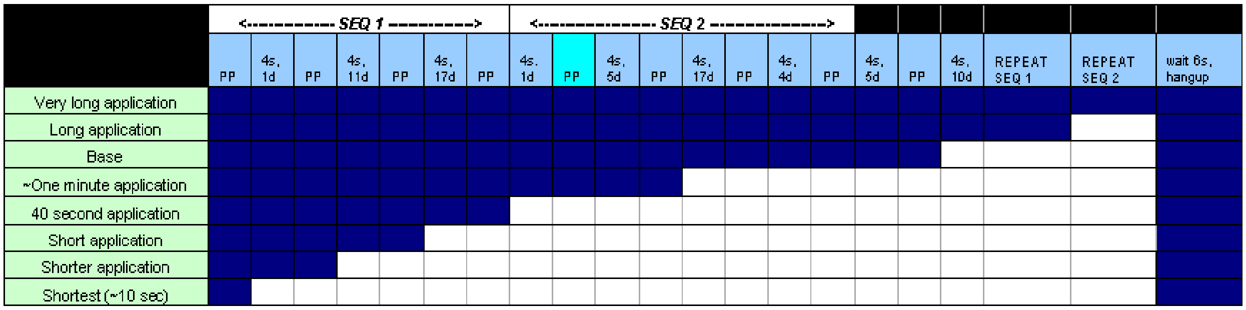

Figure 79: Application Duration Summary

Notes: `PP` = play prompt.

`Xs,Yd` = wait `X` seconds for DTMF input with length of `Y`. For example, `4s,5d` means wait four seconds for the input of five DTMF keys.

Figure 80 shows the result.

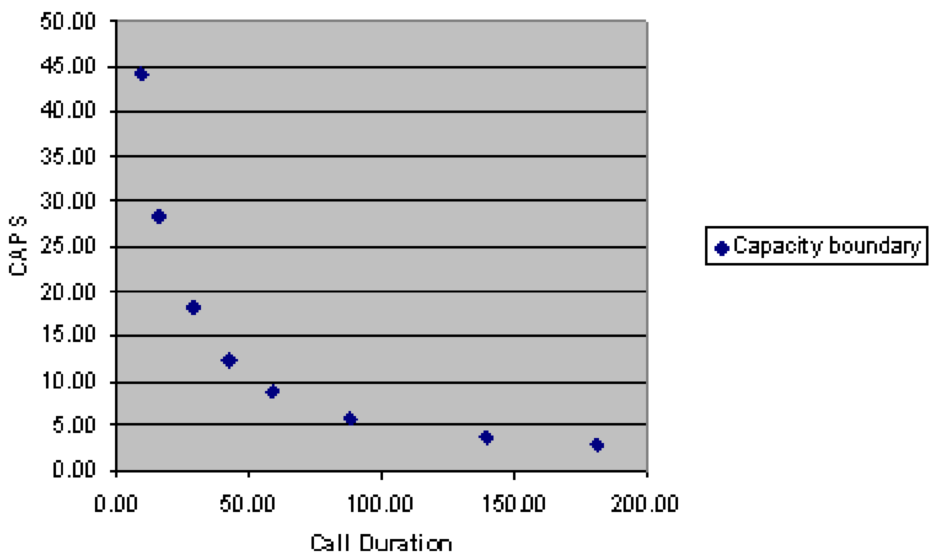

Figure 80: CAPS vs. Call Duration (Combining Peak for Each Call Flow)

You can quickly estimate the capacity by CAPS x CD.

EMS/EMPS Capacity Result

There are two types of components in the Element Management System solution, real-time components and non-real-time components. The real-time components are the software components that are part of live IVR and CTI calls. These components must respond fast enough so that there is no delay to the end caller. All other components are non-real-time components.

Real-Time Components

- Policy Manager
- IVR Server Client
- IP Call Manager
- Dispenser

Table 101 shows the capacity based on the hardware recommended in Table 97.

Table 101: EMS Real-Time Components Capacity

EMS Real-Time Components	Maximum CAPS	Capacity
Microsoft Windows		
Policy Manager	80	4000 ports, 50 customers
IVR Server Client	50	2000 ports, 50 customers
IP Call Manager	25	4000 ports/sessions
Dispenser	N/A	4000 ports
Solaris		
Policy Manager	16	1000 ports, 50 customers
IVR Server Client	16	1000 ports, 50 customers
IP Call Manager	8	2500 ports/sessions
Dispenser	N/A	4000 ports

Non-Real-Time Components

- Element Management Provisioning System/LDAP
- EventC (database)
- Bandwidth Manager
- Reporter

For the non-real-time components, the capacity is not related to ports. It is either measured on calls per day or HTTP requests per seconds. Table 102 shows the capacity of each component. This assumes that each component is running on a separate machine with the hardware specification as described in Table 97.

Table 102: EMS Non-Real-Time Components Capacity (Microsoft Windows OS only)

EMS Non-Real-Time Components	Capacity
EMPS/LDAP	200 DID per application, others unlimited
EventC	200,000 calls/hour (3 server solution)
Bandwidth Manager	Unlimited
Reporter	Unlimited, 50 simultaneous sessions/logins

Data Network Sizing

The centralized EMS real-time components require sufficient data network connection to the VCS/IPCS in different locations. Insufficient data network bandwidth causes delays between EMS real-time components and the VCS/IPCS.

The general rule for bandwidth between EMS and VCS is that every 2300 ports (100 ISDN T1) require a T1 data network connection from the VCS to the EMS. For a remote VCS site that has 4600 ports, 2 T1 data network bandwidth are required.

For a local area network within VCS/IPCS, Genesys strongly recommends that Gigabit Network be used to ensure enough bandwidth.

Deployment Examples

Table 103 and Table 104 are deployment examples for a solution that supports a 1000, 2500, or 5000 port system. The assumptions are:

- Microsoft Windows Operating System
- DTMF only applications.
- Single or multiple locations for VCS or IPCS.
- Robbed Bit T1 protocol for TDM.
- A cache server, which is required one each per site, is not included in the example.

Table 103: Deployment Examples for Genesys Voice Platform: Network Edition

	TDM 1000	TDM 2500	TDM 5000	VoIP 1000	VoIP 2500	VoIP 5000
	Number of Machines					
VCS or IPCS	4 (12 T1 each)	9	18	4	9	17
Policy Manager (EMS Runtime)	1	1	2	1	1	2
IVR Server Client (EMS Runtime)	1	2	3	1	2	3
IP Call Manager (EMS Runtime)	n/a	n/a	n/a	1	1	2
EMS Runtime Redundancy	2	3	5	3	4	7
EMPS	1	1	1	1	1	1
EventC	1	2	3	1	2	3
Reporting (Login Server, Call Status Monitor, Reporter)	1	2	2	1	2	2
EventC Redundancy	1	1	1	1	1	1
Total	**12**	**21**	**35**	**14**	**23**	**38**

Table 104: Deployment Examples for Genesys Voice Platform: Enterprise Edition

	TDM 1000	TDM 2500	TDM 5000	VoIP 1000	VoIP 2500	VoIP 5000
	Number of Machines					
VCS or IPCS	4 (12 T1 each)	9	18	4	9	17
Call Manager (optional)	n/a	n/a	n/a	1	1	2

Table 104: Deployment Examples for Genesys Voice Platform: Enterprise Edition (Continued)

	TDM 1000	TDM 2500	TDM 5000	VoIP 1000	VoIP 2500	VoIP 5000
	Number of Machines					
VPM	1	1	1	1	1	1
Total	**5**	**10**	**19**	**6**	**11**	**20**

GVP 7.5 and 7.6

This section provides the hardware sizing requirements for Genesys Voice Platform (GVP) 7.5 and 7.6.

Hardware and Operating Systems

Table 105 describes the hardware and operating system recommendations.

Table 105: Hardware and Operating System

Hardware Specification	Recommendation	Comments
Microsoft Windows		
CPU	Dual Xeon 3.6GHz or higher Dual core Dual 2.0GHz or higher	Slower CPU such as Dual 2.4 or 3.0 is supported.
Memory	2 GB or more	VCS with Dialogic has a 4 GB limit.
Network	GigaBit Ethernet	100MBit is also supported.
Storage	RAID 1 disk storage with at least 36 GB	No comment.
Operating System	Microsoft Windows 2003 Enterprise Server SP1	Microsoft Windows 2000 is also supported on GVP 7.5. Microsoft Windows 2000 is not supported on GVP 7.6.

Table 105: Hardware and Operating System (Continued)

Hardware Specification	Recommendation	Comments
Solaris		
CPU	Sun Fire T1000 with 1.0 GHz UltraSPARC T1 processor, 6 cores, 24 CoolThreads (GVP 7.6) Sun Fire V240 with UltraSPARC IIIi Dual 1.5 GHz processors or higher	Sun Fire V490 has also been validated.
Memory	2 GB or more	No comment.
Network	GigaBit Ethernet	100MBit is also supported
Storage	RAID 1 disk storage with at least 36 GB with 15K RPM	No comment.
OS	Solaris 10	Solaris 9 is also supported on GVP 7.5. Solaris 9 is not supported on GVP 7.6.

System Capacity

This section describes the capacity recommendations for GVP.

Criteria

The system capacity is defined as how many ports the GVP system can support that maximizes the machine's CPU and memory without significantly impacting the CPL, and with a low or zero error rate.

An increasing volume of call traffic that comes into the GVP system will incur the usage of additional ports, hence increasing the CPU and memory usage on the GVP system, and which will then increase CPL and error rates.

By configuring GVP with the recommended memory requirements, CPU use for memory management can be ignored.

The Genesys performance results show that the capacity has an average CPU usage of less than 75 percent, all of the results have a CPL of less than three (3) seconds, and the call duration ratio Dx/D1 is less than 120 percent, with an error rate of less than 0.1 percent. These results show that memory is not an issue for capacity in the performance, and the CPU is always the impeding factor in determining the number of ports supported on a single GVP Server.

Genesys recommends that the Media Resource Control Protocol (MRCP) server that supports the ASR and TTS application reside on a different machine from the GVP VCS machine.

The recommended hardware shown in Table 105 has been used for the capacity test, which means that Dual 3.6GHz Xeon servers were used in the test. Even though GVP supports Dual 2.4GHz and 3.0GHz machines, the performance number will be lower if lower clocked machines are used.

Capacity and Voice Application Complexity

Genesys uses several voice applications to demonstrate the GVP capacity for different voice application profiles. By no means have these voice applications covered the wide range of the voice applications that are deployed in a customer's lab and production. When you plan to capacity, there are other factors that have direct impact on how many ports GVP can support. These factors are:

- Size of the VoiceXML pages—Large pages tend to have more VoiceXML variable declarations and more ECMAScripts on the page, which causes the GVP VoiceXML browser to use more CPU cycles.
- Dynamic or static VoiceXML pages—If a voice application has the majority of its pages dynamically created, there will be less opportunity for GVP to cache the compiled `.xml` pages in GVP 7.5 or 7.6, while applications that are dominated by static pages will see some significant improvement for CPL and port density.
- Application logic on the server-side or on the GVP-side—Depending on how the voice applications are being developed, if the majority of the application logic is coded in the VoiceXML pages, it will impact the capacity because GVP will spend more CPU time handling this logic.

TDM/Dialogic Capacity Result

The results shown in Table 106 are based on Genesys performance test results.

Table 106: TDM/Dialogic Capacity Result (Microsoft Windows OS only, Dual Xeon)

Application Type	Maximum CAPS	Tested Ports	Comment
Simple application, DTMF only	5	288 (12 T1)	The limit is with the Dell machine used in the performance test. Dell machines do not have enough electric power in the backplane bus to support four DM/V Dialogic boards.
Simple application, ASR (digits collection)	5	288 (12 T1)	The limit is with the Dell machine used in the performance test. Dell machines do not have enough electric power in the backplane bus to support four DM/V Dialogic boards.
Simple application, Large Grammar ASR (40,000 grammar items)	5	192 (8 T1)	No comment.
Very complex application with ASR	0.4	48	Dynamic application pages.

VoIP Capacity Result

The results shown in Table 107 are based on Genesys performance test results.

Note: The performance test results for the SunFire T1000 Server apply to GVP 7.6.

Table 107: VoIP Capacity Result

Application Type	Maximum CAPS	Tested Ports	Comment
Microsoft Windows Dual Xeon			
Simple application, DTMF only	5	250	No comment
Simple application, ASR	5	200	No comment
Simple application, Large Grammar ASR	5	200	No comment

Table 107: VoIP Capacity Result (Continued)

Application Type	Maximum CAPS	Tested Ports	Comment
Typical TTS only	5	200	No comment
OSDM core application	3	95	No comment
Very complex application with ASR	0.4	48	No comment
Typical TTS only	5	200	No comment
Solaris V240			
Simple application, DTMF only	1.5	100	No comment
Simple application, ASR	1.5	80	No comment
Simple application, Large Grammar ASR	1.5	80	No comment
OSDM core application	1.3	40	No comment
Typical TTS only	1.5	80	No comment
Solaris V490			
Simple application, DTMF only	5	320	All V490 results require special configuration on the machine.
Simple application, ASR	5	250	All V490 results require special configuration on the machine.
Simple application, Large Grammar ASR	5	250	All V490 results require special configuration on the machine.
Typical TTS only	3.5	230	All V490 results require special configuration on the machine.
Solaris T1000 (GVP 7.6)			
Simple application, DTMF only	6	390	All T1000 results require special configuration on the machine.
Simple application, ASR	4	240	All T1000 results require special configuration on the machine.
Simple application, Large Grammar ASR	4	240	All T1000 results require special configuration on the machine.

VoIP Capacity Result for Dual-Core Machine

The information in this section is provided to help you to understand the normal behavior of GVP on a typical dual-core dual-processor machine.

Hardware used in testing (dual-core dual-processor CPU)

- IBM x3550, Dual-Core Intel Xeon 5230 2.00GHz/1333MHz 2x2MB L2 Cache
- Two Processors
- 1GB Memory
- 15K RPM SAS HDD

Application: DTMF only, phone card refill application with different call duration

- 18 VoiceXML pages (including one application root page)
- 8 user DTMF input (maximum sequence length of 16 digits [credit card number])
- Call duration is 73 seconds
- DTMF only, no ASR/TTS resource usage
- 6 VoiceXML pages out of a total of 18 that contain 13 simple java script functions (moderately ECMA script intensive)
- 77 unique audio files, 262 `<audio>` tags

Application: ASR, tipping guide application

- 6 user inputs, including number, money amount, an item menu, and yes/no
- 3 mandatory voice inputs, while the other 3 are optional between DTMF and voice
- Moderate ECMA script inside the page
- Base call duration is approximately 65 seconds

The results shown in Table 108 are based on Genesys performance test results.

Table 108: VoIP Capacity on Dual-Core Dual-Processor Machine

Approximate DTMF Application Duration (seconds)	Maximum CAPS	Tested Ports	Comment
27	17.8	500	CPU ~97%
40	12.1	500	CPU ~89%
85	5.8	500	CPU ~86%

Table 108: VoIP Capacity on Dual-Core Dual-Processor Machine (Continued)

Approximate DTMF Application Duration (seconds)	Maximum CAPS	Tested Ports	Comment
Approximate ASR Application Duration (seconds)	**Maximum CAPS**	**Tested Ports**	**Comment**
69	3	220	CPU ~42% Call duration > 120% of base if CAPS goes higher

Capacity versus Call Duration

When the CAPS is greater than five, GVP still supports an application running on the platform; however, the port density might change.

In order to obtain the capacity relationship between call duration and CAPS, Genesys ran tests with the DTMF-only performance result shown above. Genesys added or shortened the call durations to evaluate the maximum CAPS.

Figure 81 highlights the duration summary of the application.

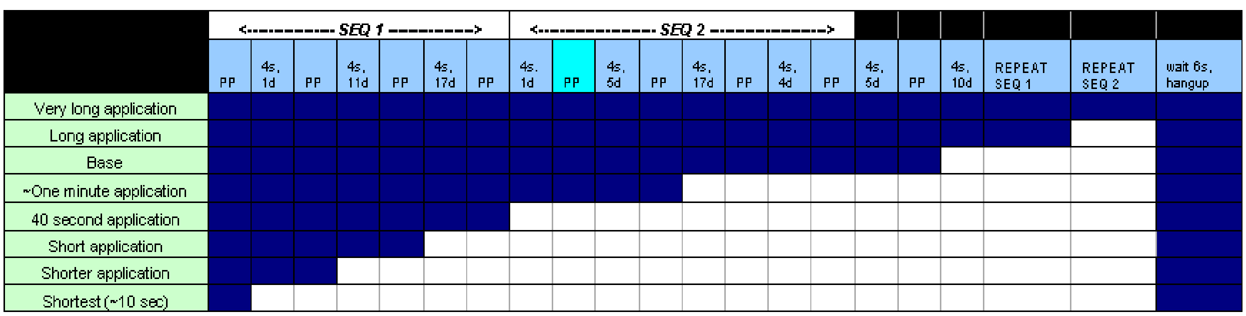

Figure 81: Application Duration Summary

Notes: `PP` = play prompt.

`Xs,Yd` = wait `X` seconds for DTMF input with length of `Y`. For example, `4s,5d` means wait four seconds for the input of five DTMF keys.

Figure 82 shows the result.

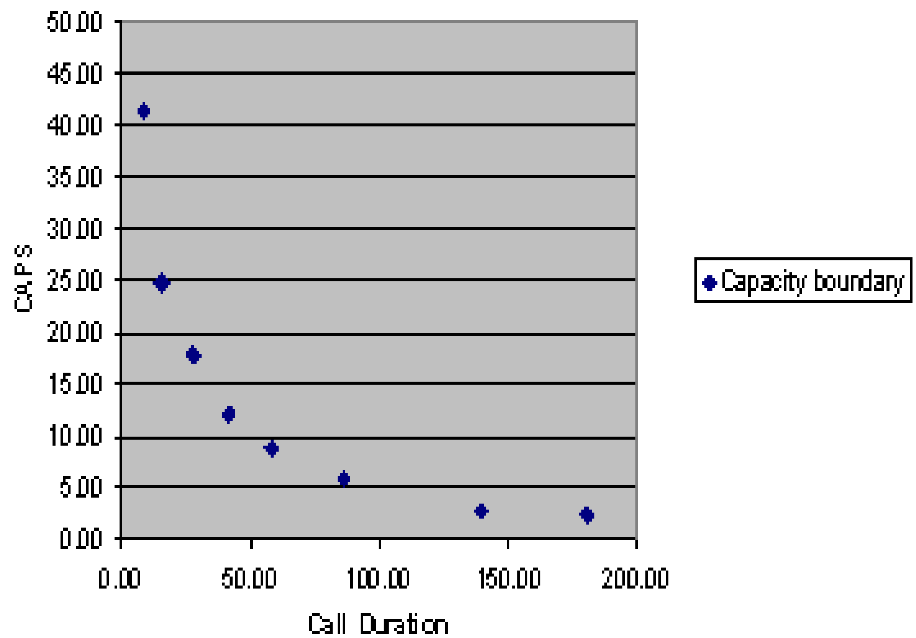

Figure 82: CAPS vs. Call Duration (Combining Peak for Each Call Flow)

You can quickly estimate the capacity by CAPS x CD.

GVP Performance Improvement in VoiceXML Browser

GVP 7.5 and 7.6 introduces the compiled XML page caching technology within the GVP VoiceXML browser. For static VoiceXML pages, instead of compiling the page and creating an action object for the elements for each occurrence, the compiled pages are saved in memory for future use. When the same page is received by the GVP VoiceXML browser, the VoiceXML browser will not recompile the page, but will use the existing compiled page in memory. This reduces the CPU usage of the VoiceXML browser, and improves the overall system efficiency and port density.

This improvement mainly benefits a voice application that has a majority of static pages.

Figure 83 shows the VoiceXML browser usage of CPU when the CAPS increases (consequently the concurrent ports) in the GVP system. The GVP 7.5nc example indicates a DTMF telephone book application with instrumentation, so every page has a dynamic comment, causing the page to not be cacheable (nc). The GVP7.5c is a standard DTMF telephone book application, where all VoiceXML pages are static, so it is cacheable (c).

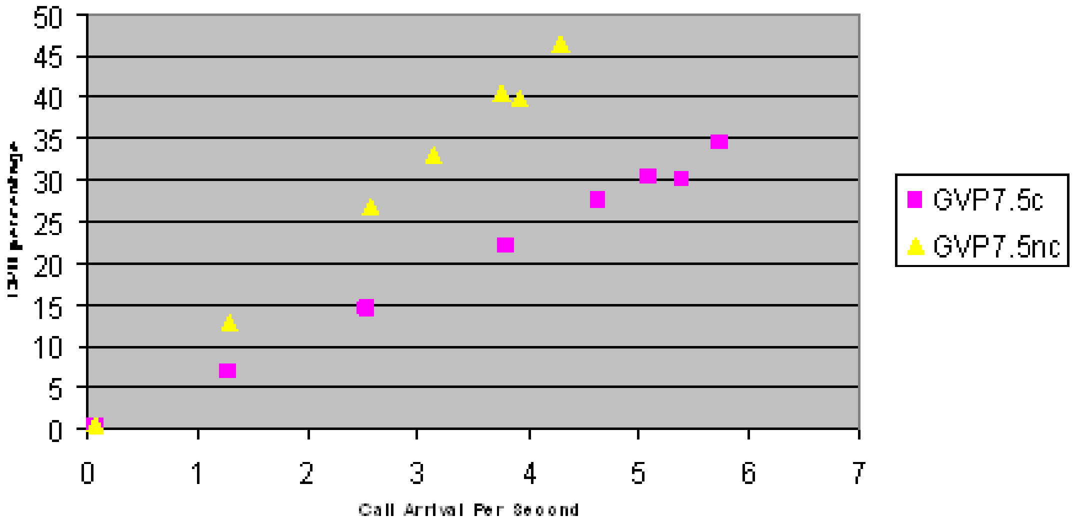

Figure 83: VoiceXML Browser (PopGateway) CPU Usage vs. CAPS

A typical application will fit between the above two graph lines. GVP7.5c represents the best case of using compiled page caching technology introduced in GVP 7.5 and 7.6, while GVP7.5nc represents the worst case.

EMS/EMPS Capacity Result

There are two types of components in the Element Management System solution, run time components and non-run time components. The run time components are the software components that are part of ongoing IVR and CTI call handling. These components must respond quickly enough so that there is no delay to the end caller. All other components are non-run time components.

Run Time Components

- Policy Manager
- IVR Server Client
- IP Call Manager
- Dispenser

In GVP 7.5 and 7.6, the EMS and EMPS components can be combined and placed on one physical server. Table 109 shows the capacity based on the hardware recommended in Table 105.

Table 109: EMS Run Time Components Capacity

EMS Run Time Components	Maximum CAPS	Capacity
Microsoft Windows Dual Xeon		
Policy Manager	150	8000 ports, 60 customers
IVR Server Client	150	5000 ports, 60 customers
IP Call Manager	45	750–6000 ports/sessions (depends on call duration)
Dispenser	Not applicable.	4000 applications
Solaris		
Policy Manager	50	1500 ports, 60 customers
IVR Server Client	60	2000 ports, 60 customers
IP Call Manager	8	1000 ports/sessions
Dispenser	Not applicable.	4000 applications

Non-Run Time Components

- Element Management Provisioning System/LDAP
- EventC (database)
- Bandwidth Manager
- Reporter

For the non-run time components, the capacity is not related to ports. It is either measured on calls per day or HTTP requests per seconds. Table 110 shows the capacity of each component. This assumes that each component is running on a separate machine with the hardware specification as described in Table 105.

Table 110: EMS Non-Run Time Components Capacity

EMS Non-Run Time Components	Capacity
Microsoft Windows Dual Xeon	
EMPS/LDAP	200 DID per application, others unlimited
EventC	126,000 call/hour (one server solution)

Table 110: EMS Non-Run Time Components Capacity (Continued)

EMS Non-Run Time Components	Capacity
Bandwidth Manager	Unlimited
Reporter	Unlimited
Solaris	
EMPS/LDAP	Not applicable.
EventC	35,000 calls/hour (one server solution)
Bandwidth Manager	CAPS=2, recording lost < 1.5%
Reporter	Unlimited

Because GVP 7.5 and 7.6 allows logical software components to reside on the same physical machine, the capacity information in Table 111 shows the GVP solution under different configurations.

Table 111: GVP on One Machine and Two Machine Configuration (Microsoft Windows Only, Dual Xeon)

EMS Non Run Time Components with EMS Run Time Components	Maximum CAPS	Maximum Ports Tested	Comment
EMS+EMPS+VCS+ASR +TTS (all in one TDM)	4 4.14 3	288 216 168	Simple DTMF Simple ASR Simple ASR, large grammar
EMS+EMPS+IPCS+ASR +TTS (all in one IP)	3.7 2 1.58	245 130 90	Simple DTMF Simple ASR Simple ASR, large grammar
EMS+EMPS+VCS	4	288	Simple DTMF **Note:** The limit is with the Dell machine used in the performance test. Dell machines do not have enough power to support four DM/V Dialogic boards.
EMS+EMPS+IPCS	3.7	245	Simple DTMF

Table 111: GVP on One Machine and Two Machine Configuration (Microsoft Windows Only, Dual Xeon) (Continued)

EMS Non Run Time Components with EMS Run Time Components	Maximum CAPS	Maximum Ports Tested	Comment
EMS+EMPS on first machine; VCS on second machine	4	288	Simple DTMF **Note:** The limit is with the Dell machine used in the performance test. Dell machines do not have enough power to support four DM/V Dialogic boards.
EMS+EMPS on first machine; IPCS on second machine	3.7	245	Simple DTMF EMS+EMPS only use 1.34% CPU

Data Network Sizing

The centralized EMS run time components require sufficient data network connection to the VCS/IPCS in different locations. Insufficient data network bandwidth causes delays between EMS run time components and the VCS/IPCS.

The general rule for bandwidth between EMS and VCS is that every 2300 ports (100 ISDN T1) require a T1 data network connection from the VCS to the EMS. For a remote VCS site that has 4600 ports, two T1 data network bandwidth are required.

For a local area network within VCS/IPCS, Genesys highly recommends that Gigabit Network be used to ensure enough bandwidth.

High Performance Configuration

This section applies to GVP 7.2, GVP 7.5, and GVP 7.6 releases. For the GVP 7.6 release, only the Popgateway section and PageCollector section in Table 112 are required. The other sections of the table (McuXML and McsXML/asr/mrcp) are not applicable; you do not need to enter them.

In order for the IPCS to support more than 200 ports on a single machine, the parameter changes shown in Table 112 are required.

Table 112: High Performance Parameter Configuration

Node/Parameter	Default	High Performance
Popgateway		
MaxPoolThreads	200	500
MaxTPThreads	400	500
McuXML		
MaxTPThreads	100	200
RtpSendThreads	#cpu (2 if dual CPU)	64
RtpRcvThreads	#cpu	64
RtpMinPortNum	10000	10000
RtpMaxPortNum	12000	16000
McuXml/asr/mrcp		
Rtpportrange	30500–32000	30500–40000
TTS Template	25000–30000	40500–50000
Tts_genericproc	25000–30000	40500–50000
PageCollector		
MaxPoolThreads	200	400
Registry on Microsoft Windows HKLM\Software\CallNet\CnInet Settings		
MaxThreadPool	64	256

Chapter

9 Genesys Voice Platform 8.1

This chapter specifies the recommended hardware and operating systems for Genesys Voice Platform (GVP) 8.1 and provides information about capacity and performance testing.

This chapter contains the following sections:

Hardware and Operating Systems

With the exception of the operating systems, the recommended specifications in Table 113 are the same for Windows and Linux.

Table 113: Recommended Hardware and Operating Systems

Hardware	Recommended Specification
Central Processing Unit (CPU)	Dual Quad Core Xeon, 2.66 GHz or Higher (Xeon with Core 2 technology is recommended for optimal performance.)
Memory	4 GB or more
Network	GigaBit Ethernet (100 MegaBit is supported.)
Storage	SAS disk storage with at least 36 GB (15k rpm is recommended for maximum performance.)
Operating Systems (OS)	Windows Server 2003, R2 Enterprise Edition (Windows Server 2003, Standard Edition is also supported.)
	Red Hat Enterprise Linux 4.0, Update 6

Capacity and Performance Planning

Use the information in this section to determine the required capacity of your GVP servers, based on anticipated traffic characteristics or by running tests on an existing system. You can then use the information in the following sections to determine the appropriate hardware configuration for your GVP 8.1 environment.

Capacity and Traffic Terminology

System capacity is defined as the maximum number of ports (Port Density) or maximum call-arrival rate (Call Arrival Per Second) at which GVP can maximize the use of hardware resources while maintaining all of the criteria within a predefined threshold.

This section defines the terms used in capacity and performance management and provides the formulas used to calculate it.

Capacity Metrics and Formulas

Two units of measure are used for capacity planning—Port Density (PD) and Call Arrival Per Second (CAPS).

Port Density

PD is defined as the maximum number of simultaneous calls being served by GVP at any given time, or the number of ports that are required to handle the call traffic.

Use the following formula to calculate Port Density:

```
PD = CAPS x Avg(CD)
```

(See also, "Call Duration".)

Call Arrival Per Second

CAPS is used as a watermark for traffic within the system. For example, 10 CAPS means that GVP is receiving 10 calls every second, which is considered *busy traffic*.

CAPS is similar to Busy Hour Call Attempts (BHCA) or Centum Call Seconds (CCS), which is the legacy engineering term for telephony traffic.

Use the following formula to calculate CAPS in terms of CCS:

```
CAPS = CCS/36
```

Performance Metrics and Formulas

Three units of measure are used to assess performance—Call Duration (CD), Call Setup Latency (CSL), Caller Perceived Latency (CPL), and Call Passrate (CP).

Call Duration

CD is defined as the length of time that a call stays within GVP. CD and CAPS are used to calculate the PD required to handle call traffic. Sometimes, the increase in the total CD is measured rather than the individual CPLs within the application.

For a single call, the assumption is that the normal call duration—which ends when the application terminates the call—is CD1. When the load increases on the system, the call duration is expected to increase due to increased latencies within the application.

Use the CDx metric as shown in the following example to obtain a capacity measurement goal:

Assume *x* represents the number of simultaneous calls within GVP.

- The capacity measurement goal for Average CDx is

  `Avg(CDx)/CD1 <= 110%`
- The capacity measurement goal for the 95th percentile CDx is

  `95%(CDx)/CD1 <= 120%`

In the example, when the threshold (110% or 120%, whichever calculation is used) is reached, the call volume *x* is considered peak capacity for this criteria. The CPL and CD thresholds are used interchangeably in this capacity result.

Call Setup Latency

CSL is defined as the delay between the initial SIP `INVITE` message and the first audible RTP packet sent from GVP. The dialogue in Figure 84 is an example of a typical SIP call flow.

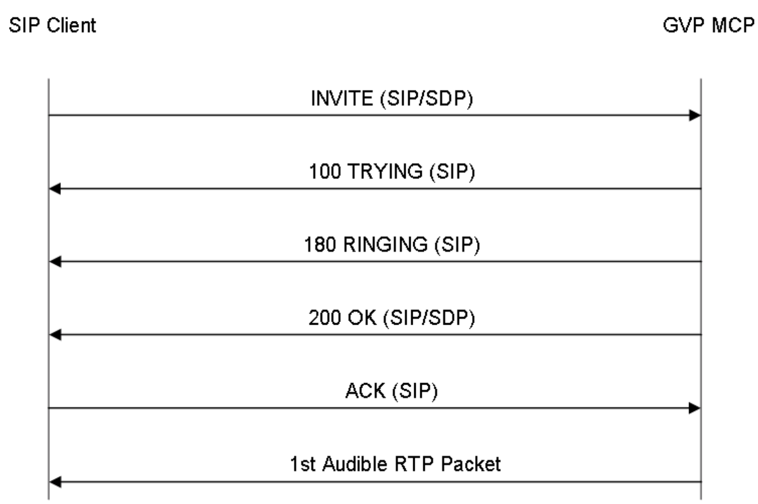

Figure 84: Typical SIP Call Flow

CSL consists of the following requests and responses:

- User SIP `INVITE` request received > SIP `200 OK` response sent.
- SIP `200 OK` response sent > User SIP `ACK` request sent.
- User SIP `ACK` request sent > First audible media response sent.

Caller Perceived Latency

CPL (or Response Time Latency [RTL]) is defined as the time between the last user input (speech or DTMF) and the next prompt. In Figure 85, the time between T1 and T2 is the CPL period.

Calculate CPL by using the formula, `CPL = T2 - T1`.

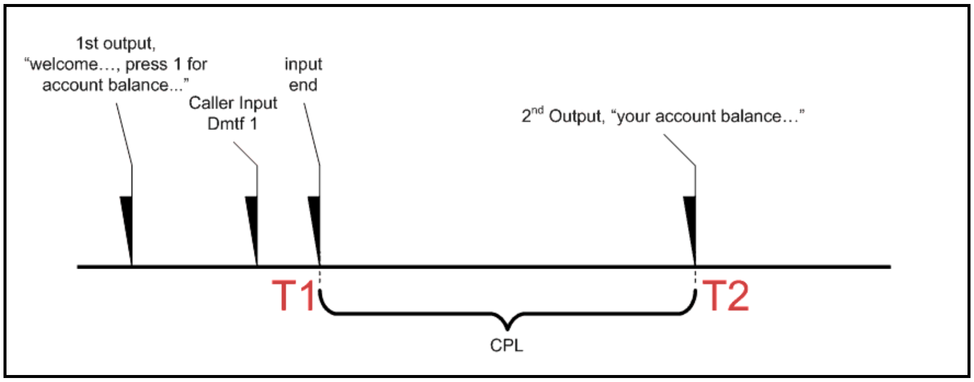

Figure 85: Caller Perceived Latency

CPL is impacted by the following factors:

- Recognition engines.
- End of speech or a DTMF time out.
- Application pages and prompts.
- Grammars caching and fetching mechanisms.
- The size of application pages.
- Call traffic, including call arrival rate and call duration.
- Speech density—during a call, the percentage of time that speech or DTMF detection is on and the caller can provide input.
- The size of speech recognition grammars and the how often they are used in an application.
- Back end operation—the length of time required to obtain information from the business layer (such as, database queries and CRM findings) and return the results to the caller.

Call Passrate

CP is defined as the number of calls that finished the predefined call flow successfully during a performance load test.

Use the following formula to calculate the CP:

> Assume the number of calls is 1000 and 5 calls did not finish the predefined call flow successfully.
>
> `(1000-5)/1000 = 99.5%`

The capacity measurement formulas are:

ASR/TTS-dependent application	Passrate >= 99.95%	Error rate <= 0.05%
DTMF-only application	Passrate >= 99.99%	Error rate <= 0.01%

Similar to the formulas for CD, *x* is considered the peak capacity for this criteria when the threshold is reached at call volume *x*.

Capacity Test Cases

When measuring peak capacity on a single GVP machine, CPU usage is usually the determining factor—memory has not been an issue in most test cases. Therefore, the sample test results in this section concentrates on CPU usage and other criteria.

In addition, the Media Resource Control Protocol (MRCP) server supporting Automatic Speech Recognition (ASR) applications, must not share a host with a GVP server. You can use multiple MRCP servers for a particular test, however, it is important that the MRCP resources do not cause a bottleneck during testing.

This section contains the following capacity test cases:

- Call Setup Latency Test Case, page 238
- Caller Perceived Latency Test Case, page 238
- Cachable VoiceXML Content Test Cases, page 239

VoIP Capacity Summary

The complexity of VoiceXML applications impacts capacity testing, therefore, the Genesys QA performance testing results in this section are from test cases in which two different VoiceXML applications are used (one being more complex than the other).

The application profiles are described in Table 114. The call flow duration for each application profile is for a single call or CD1 (see “Call Duration” on page 230).

Table 114: Test Application Profiles

Profile Name	Type	Details
`VoiceXML_App1`	A simple DTMF-only application designed to refill calling cards.	• Total number of digits (DTMF input only) = 52, including: ◦ Toll free number from the back of the card. ◦ Refill card PIN number. ◦ Refill dollar amount. ◦ Credit card number. ◦ Credit card expiration date. ◦ Zip Code of caller.

Table 114: Test Application Profiles (Continued)

Profile Name	Type	Details
`VoiceXML_App1` (continued)	A simple DTMF-only application designed to refill calling cards. (continued)	• Number of VoiceXML pages = 18 • VoiceXML complexity = low • Number of audio prompts = 9 • Number of audio files used in prompts (no TTS) = 107 • ECMA script complexity = moderate • Number of VoiceXML pages = 6 • Number of Java script help functions in each VoiceXML page = 13 • Call flow duration = 74 seconds

`VoiceXML_App2`	A complex ASR application designed for injury insurance coverage.	• Speech input, including:
		  ◦ Type of request
		  ◦ ID card number
		  ◦ Confirmation
		  ◦ Relationship with insurance plan holder
		  ◦ Date of birth confirmation
		• Number of VoiceXML pages = 10
		• VoiceXML complexity (~ 1 MB of content) = 10
		• Number of audio prompts = 7
		• Number of audio files used in prompts (3 with TTS) = 29
		• ECMA script complexity = high
		• Call flow duration = 85 seconds

Note: Some capacity tests in this section were performed on systems with hardware specifications other than those recommended in Table 1 on page 5. Major differences in test results can occur depending on the model and number of CPUs used. In Table 115, the hardware used for each capacity test is specified.

Table 115: GVP VoIP Capacity Testing

Application Type	Hardware	Peak CAPS	Peak Ports	Comments
Windows				
`VoiceXML_App1`	2x Core 2 Quad Xeon x5355 2.66 GHz	17.5	1300	
`VoiceXML_App1`	2x Core 2 Quad Xeon x5160 3.00 GHz	9.5	700	
`VoiceXML_App1`	1x Core 2 Quad Xeon x5160 3.00 GHz	6	450	
`VoiceXML_App2`	2x Core 2 Quad Xeon x5355 2.66 GHz	5.5	500	
`VoiceXML_App2`	2x Core 2 Quad Xeon x5160 3.00 GHz	2.6	250	
Audio transcoding (G711u <-> AMR)	n/a	1/4 non-transcoding case	1/4 non-transcoding case	In a worst case scenario, transcoding causes a 3/4 drop in peak capacity (or CAPS) compared to the case in which there is no transcoding.
Netann announcement (-3 seconds audio)	2x Core 2 Quad Xeon x5355 2.66 GHz	120	500	Preferred, meaning the highest capacity while maintaining optimal user experience, with call setup and call tear down latency < 1sec (500ms each).
		166	1100	Peak, meaning the absolutely highest capacity the system can sustain regardless of the user experience (long delays).

Table 115: GVP VoIP Capacity Testing (Continued)

Application Type	Hardware	Peak CAPS	Peak Ports	Comments
Audio conferencing	1x Core 2 Quad Xeon x5160, 3.00 GHz CPU	n/a	500 (total participants)	The Media Control Platform supports up to 500 total conference participants regardless of the number of conferences, participants per conference, or number of speakers per conference.
Simple CTI application (Media Control Platform performance with GVPi)	2x Core 2 Quad Xeon x5355, 2.66 GHz PU	25 (per Media Control Platform)	500 (per Media Control Platform)	Simple prompts playing VoiceXML applications with CTI information query and store functionality.
Simple CTI application (overall system performance with multiple Media Control Platforms)	2x Core 2 Quad Xeon x5355, 2.66 GHz CPU	40 (overall system)	900 (overall system)	Simple prompts playing VoiceXML applications with CTI information query and store functionality (behind-the-switch configuration.)
SIP call (Resource Manager performance)	1x Core 2 Quad Xeon x5160, 3.00 GHz CPU	500	Any number	Occurs regardless of the port density or the type of calls routed. Multiple Media Control Platforms required to achieve the peak CAPS. Reporting Server is disabled, however, with Reporting Server enabled, a peak CAPS bottleneck would be caused by the Reporting Server. See SIP call (Reporting Server with Microsoft SQL Server) in this table.

Table 115: GVP VoIP Capacity Testing (Continued)

Application Type	Hardware	Peak CAPS	Peak Ports	Comments
SIP call (Reporting Server with Microsoft SQL Server)	Reporting Server: 2x Core 2 Quad Xeon x5355 2.66 GHz CPU Microsoft SQL Server: 2x Core 2 Quad Xeon x5355 2.66 GHz CPU	120	Any number	Occurs regardless of the port density or the type of calls processed. Resource Manager and Media Control Platform log information to the Reporting Server using default settings. Increased reporting and logging can reduce Reporting Server capacity. Microsoft SQL database installed on Windows 2003 Server with the database residing on a 15k rpm HDD eSAS SAN.
SIP call (Reporting Server with Oracle Database Server)	Reporting Server: 2x Core 2 Quad Xeon x5355, 2.66 GHz CPU Oracle Server: 2x Core 2 Quad Xeon x5355, 2.66 GHz CPU	60	Any number	Occurs regardless of the port density or the type of calls processed. Resource Manager and Media Control Platform log information to the Reporting Server using default settings. Increased reporting and logging can reduce Reporting Server capacity. Oracle Database Server is installed on Windows 2003 with the database residing on a 15k rpm HDD eSAS SAN.
Linux				
`VoiceXML_App1`	2x Core 2 Quad Xeon x5355, 2.66 GHz CPU	8	600 (tested)	A single all-in-one GVP system (with a database server, Reporting Server, Media Control Platform, Fetching Module, Resource Manager, and all-in-one SIP Server).
`VoiceXML_App2`	2x Core 2 Quad Xeon x5355, 2.66 GHz CPU	5	400 (peak)	Peak ports are over 400 and cause the call duration to increases beyond the threshold.

In the following test cases, maximum capacity was achieved within the constraints of specific thresholds. However, the system was also tested beyond the recommended capacity to determine the extent of performance degradation.

Call Setup Latency Test Case

The test case in Figure 86 uses the `VoiceXML_App1` profile (in Table 114 on page 233) to show how the CSL increases as the PD increases. The rate at which the CSL increases is relatively constant until the system reaches a bottleneck—for example, when the system load is beyond peak capacity.

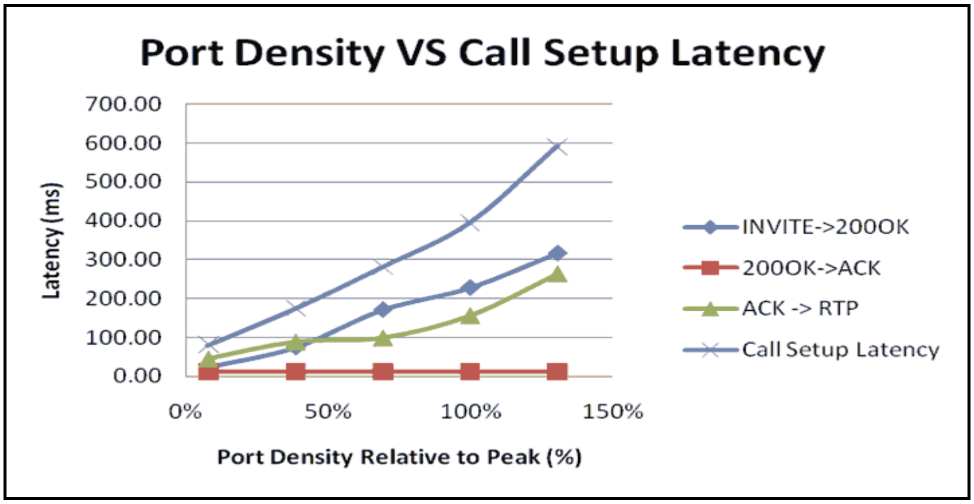

Figure 86: PD Versus CSL

Caller Perceived Latency Test Case

The graph in Figure 87 shows the DTMF response-to-audio-prompt latency at various port densities (relative to the peak capacity indicated in Table 115 on page 235). Notice that the TTS prompts produce ~300 ms more latency than the audio file prompts. This is due to the beginning silence played by the TTS engine.

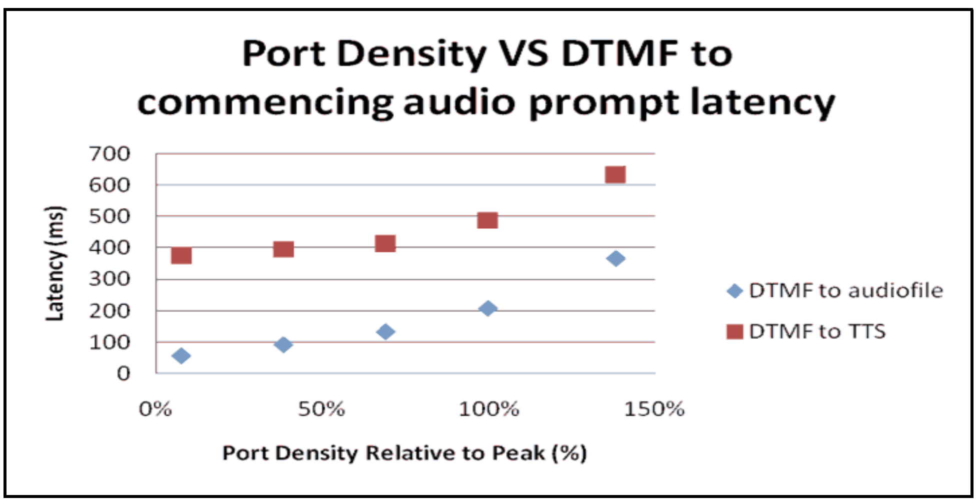

Figure 87: PD Versus DTMF

When there is speech input, additional latency is usually caused by the ASR engine. In Figure 88, the latency result is from 1000 words of grammar using the Nuance OSR3 MRCP version 1 (MRCPv1) engine. The result can vary, depending on the type of MRCP engine used, the type of speech grammar used, and the load on the speech engine.

The performance results in Figure 88 were obtained from isolated ASR engines supporting the same number of recognition sessions at all Media Control Platform port densities; the MRCP engines did not cause a bottleneck. Therefore, depending on the load on the Media Control Platform, it can add as much as ~100 ms of latency.

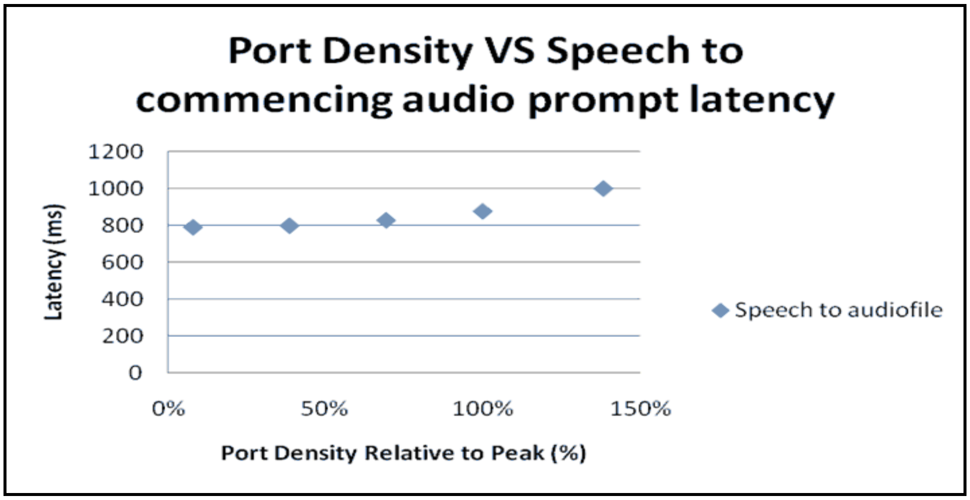

Figure 88: PD Versus Speech

Cachable VoiceXML Content Test Cases

GVP can cache internal, compiled VoiceXML objects. Caching VoiceXML objects saves a significant amount of compilation time, which results in less CPU consumption. The `VoiceXML_App1` application is used for the test case in Figure 89 and is based on the peak capacity indicated in Table 115 on page 235.

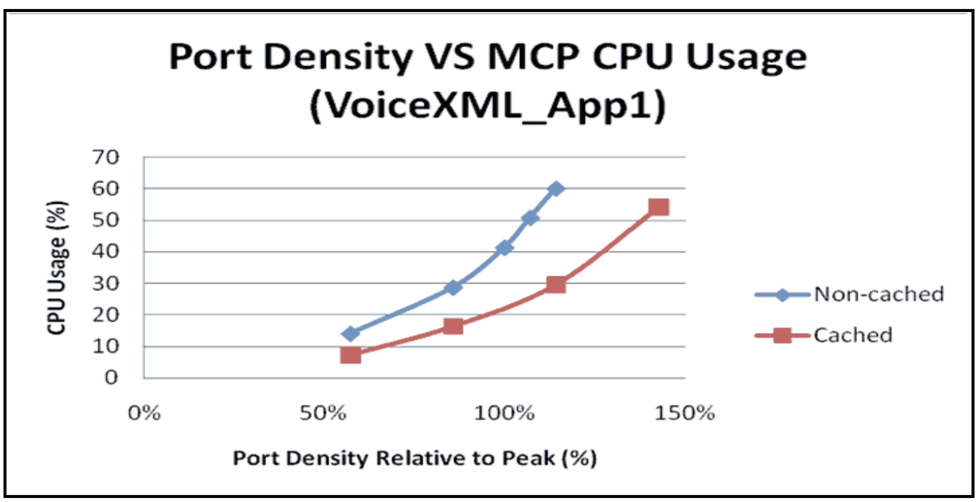

Figure 89: PD Versus CPU (VoiceXML_App1)

The more complex the VoiceXML content, the greater the benefit of having cachable content. The test case in Figure 90 is similar to the one in Figure 89 except the more complex `VoiceXML_App2` application is used (see Table 114 on page 233).

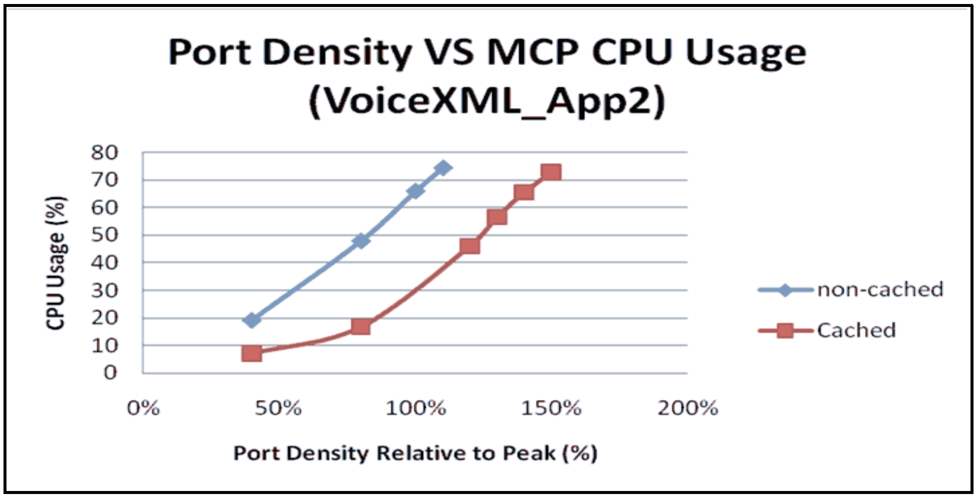

Figure 90: PD Versus CPU (VoiceXML_App2)

In Figures 89 and 90 the processing of cachable and non-cachable content are compared with the Media Control Platform using the same level of CPU consumption for both applications. The following results clearly show the benefits of using cachable content:

- CPU consumption—Media Control Platform at peak capacity.
  - 15% less consumption than non-cached content using `VoiceXML_App1`.
  - ~30% less consumption than non-cached content using `VoiceXML_App2`.
- Port density—CPU consumption at same level for both applications:
  - ~30-35% greater than non-cached content using `VoiceXML_App1`,
  - ~50% greater than non-cached content using `VoiceXML_App2`.

Performance Test Cases

The Genesys QA performance testing samples in this section include application and component test cases:

Application Test Cases

The following application test cases are described in this section:

- NetAnn Announcement Test Case, page 240
- Transcoding Test Case, page 242
- Conference Test Case, page 243

NetAnn Announcement Test Case

When the Media Control Platform acts as a media announcement server, high call rates can be sustained. It can sustain up to 166 CAPS (~1100 ports) for a typical audio playback of 3 seconds, however, call setup and tear down latency increases.

The graph in Figure 91 shows call durations at various CAPS. When CAPS reaches its peak (166 CAPS), the setup and tear down latency can reach 3.2 seconds. Optimally, call setup and tear down latency should be maintained at <1sec (or 500 ms each) with CAPS at 120 (with 500 ports).

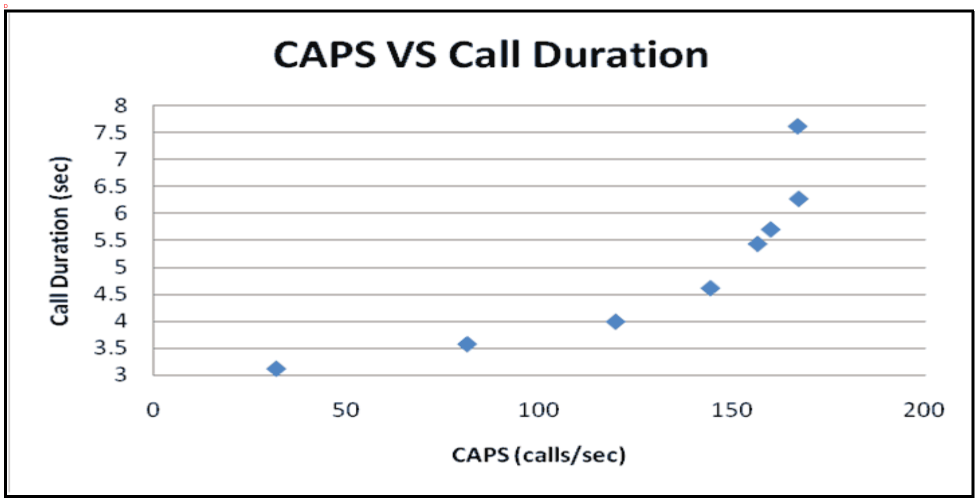

Figure 91: CAPS Versus CD (Announcement)

In Figure 92, as the call duration increases with higher port density, the additional call setup and tear down latency prevents the CAPS from scaling linearly in relation to the port density.

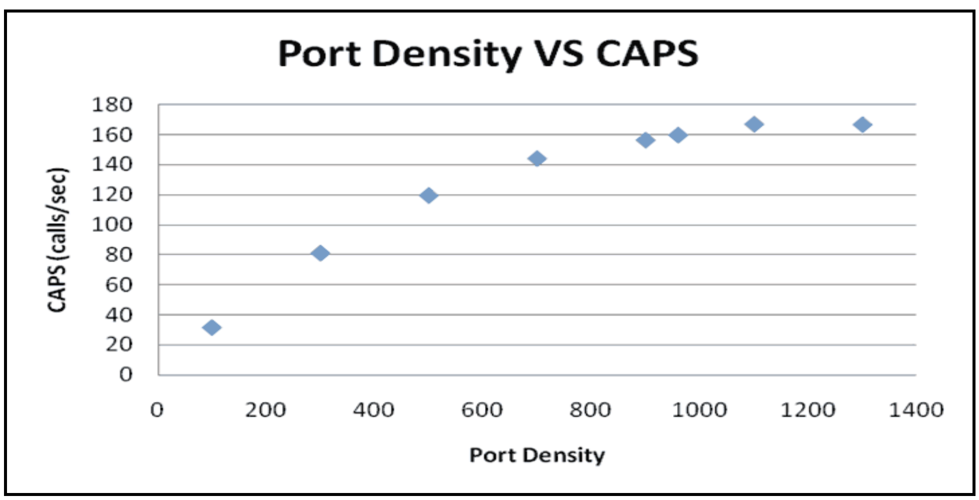

Figure 92: PD Versus CAPS (Announcement)

In Figure 93, a bottleneck is caused by the media streaming. Shorter audio announcements increase the time spent on call setup and tear down and, although the load on the system decreases, shorter audio prompts cause the peak CAPS to increase.

The graph in Figure 93 depicts a use case where a 1 second audio announcement drives the peak CAPS to ~235. Optimally, in this use case, call setup and tear down latency should be maintained at <1sec and CAPS at 200 (with ~500 ports).

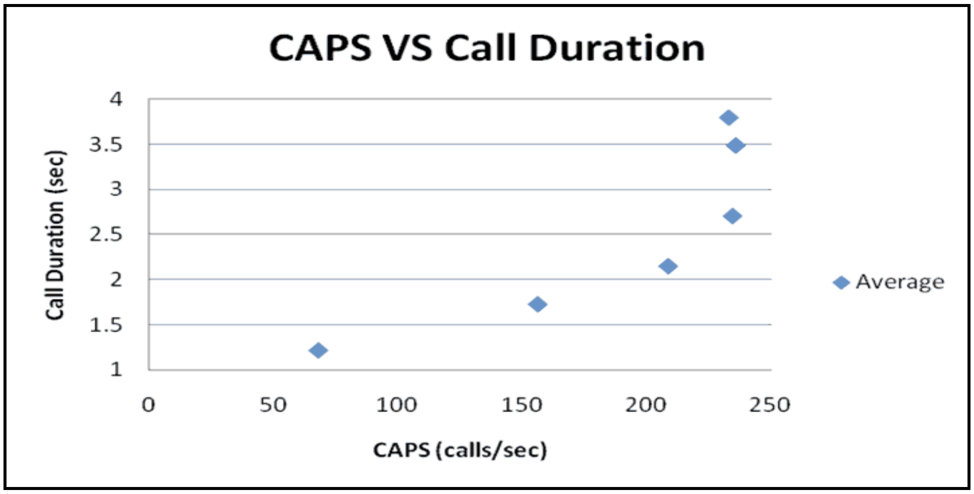

Figure 93: CAPS Versus CD (Average)

Transcoding Test Case

The Media Control Platform can transcode various media codecs in real time. The impact on performance from transcoding overhead varies, depending on the codec the Media Control Platform is transcoding to or from. Other variables that contribute to transcoding overhead are the number of audio prompts played by GVP and the amount of customer input received.

The worst case scenario occurs when the Media Control Platform is constantly transcoding between two codecs during the entire call. (Most VoiceXML applications require minimal encoding [G711u to AMR]). In Figure 94, the least amount of transcoding overhead is between G711u and G711a codecs, where the peak capacity drops by ~25%.

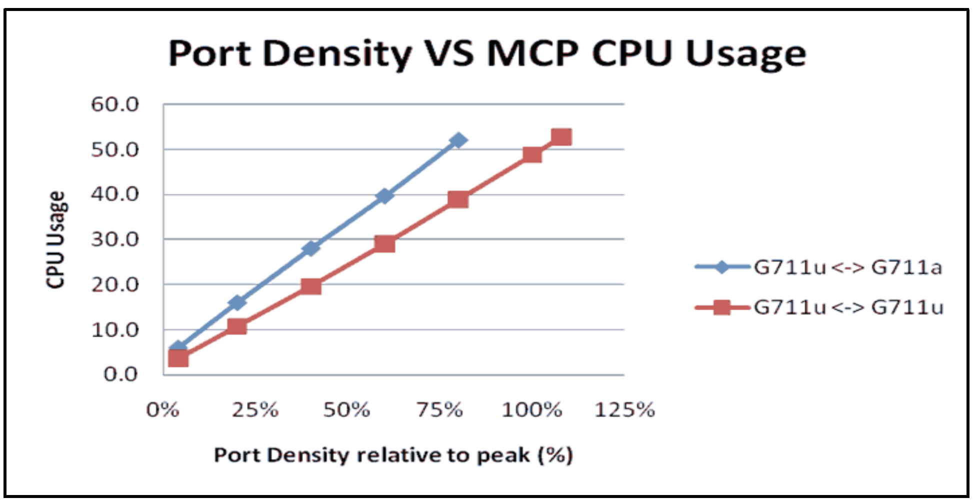

Figure 94: PD Versus CPU Usage (G711u and G711a)

The graph in Figure 95 illustrates the impact of transcoding overhead. It is greater between the G711u and AMR codecs, where the peak capacity drops by ~75%.

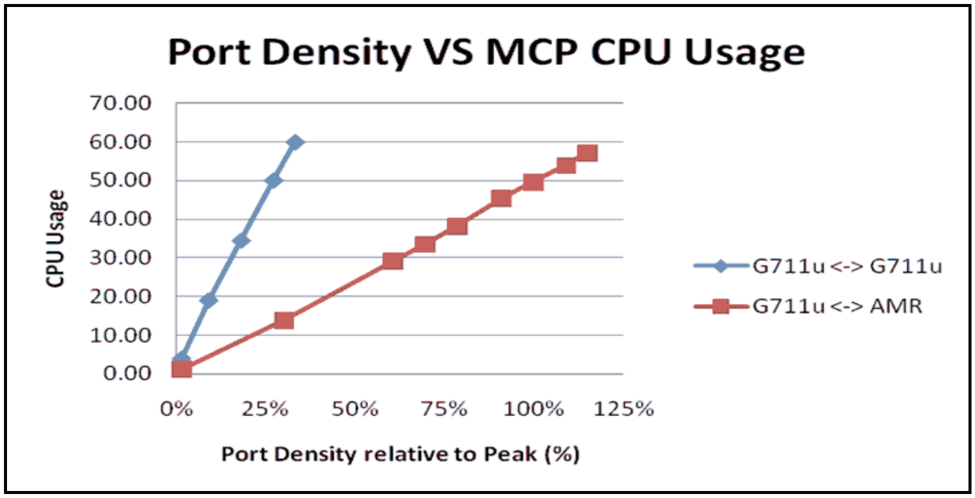

Figure 95: PD Versus CPU Usage (G711u and AMR)

As mentioned previously, the transcoding test cases in this section depict the worst case scenario involving constant transcoding between two codecs. However, in a call flow scenario where the audio stream was suppressed or silent, and the application was waiting for user input 50% of the time, transcoding overhead would be reduced by 50%.

Conference Test Case

In the following conference use cases, three variables affect Media Control Platform performance:

- The number of simultaneous conferences.
- The number of participants per conference.
- The number of speaking participants.

As the graph in Figure 96 illustrates, the variable impacting performance the most is the total number of participant hosted by GVP (the number of conferences multiplied by the number of participants per conference).

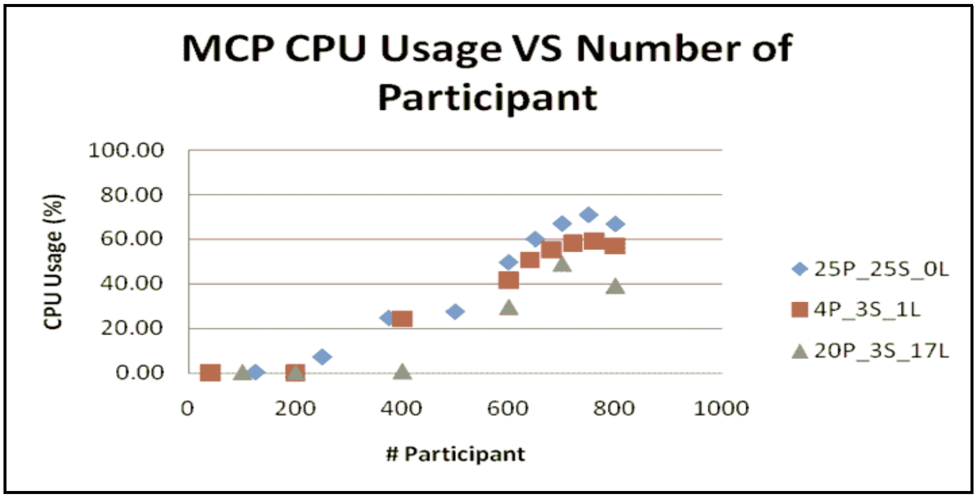

Figure 96: CPU Usage Versus Total Number of Participants

The symbols and legend in the graph in Figure 96 are explained below:

- 25P_25S_0L = 25 participants per conference (25 active speakers + 0 listen only)
- 4P_3S_1L = 4 participants per conference (3 active speakers + 1 listen only)
- 20P_3S_17L = 20 participants per conference (3 active speakers + 17 listen only)

Overall, the CPU usage increases with a higher percentage of actively speaking participants, however, regardless of the conference configuration, the system bottleneck occurs when the total number of participants reaches ~600 (on a 2x Xeon5160 @3.0GHz server).

Component Test Cases

The following component test cases are described in this section:

- Media Control Platform Test Case, page 244
- CTI Connector Test Case, page 245
- Resource Manager Test Case, page 246
- Reporting Server Test Case, page 247

Media Control Platform Test Case

Depending on the use case, performance results differ when GVP is running on Linux or Windows—overall performance tends to be better on Windows. For example, in Figure 97, at peak capacity, GVP performs ~30% better on Windows.

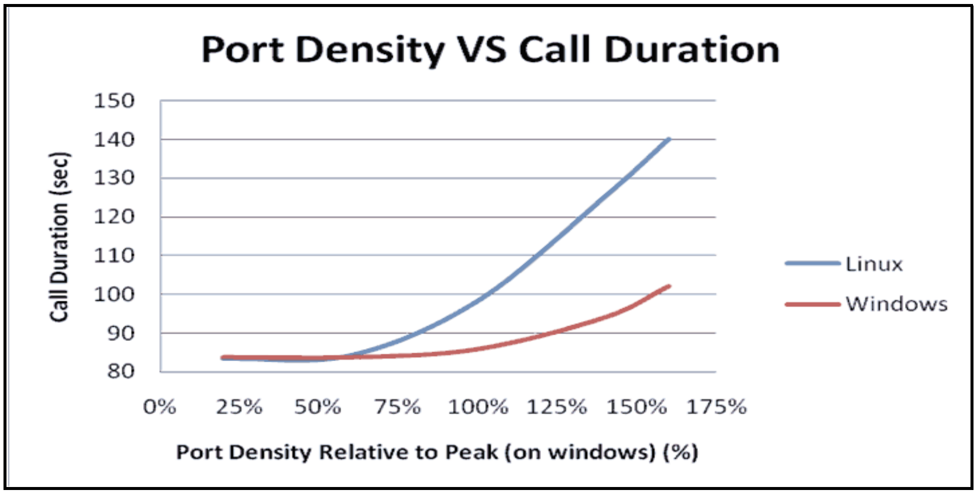

Figure 97: PD Versus CD (Windows and Linux)

CTI Connector Test Case

When the Media Control Platform is configured to use GVPi and a simple CTI information querying application is running, a port density of 450 ports (20 CAPS) is sustainable. Beyond this density, excessive disk input/output (I/O) forces the overall call passrate below the 99.95% criteria. See Figure 98.

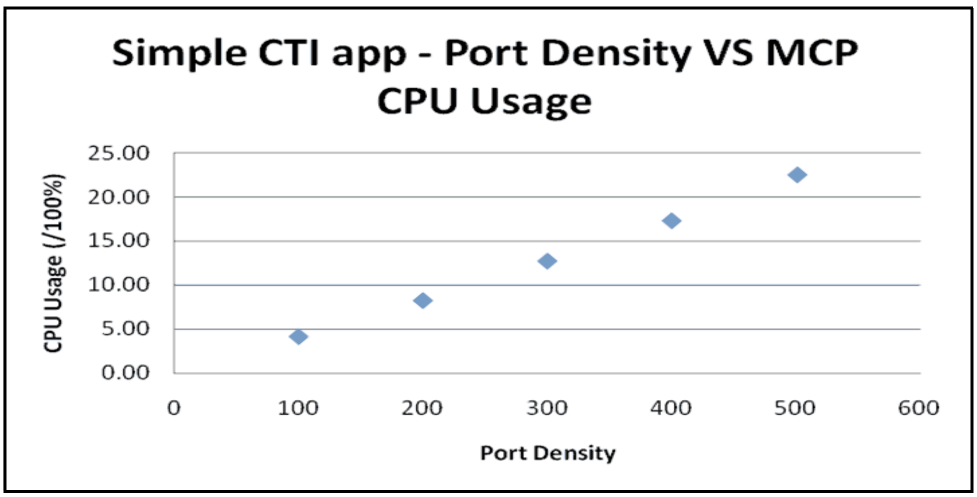

Figure 98: PD Versus CPU Usage (CTI Connector)

At the system level (with multiple Media Control Platforms), a peak of 55 CAPS can be achieved with a PD of 1300. Beyond the peak CAPS, the call passrate drops below the 99.95% threshold. See Figure 99.

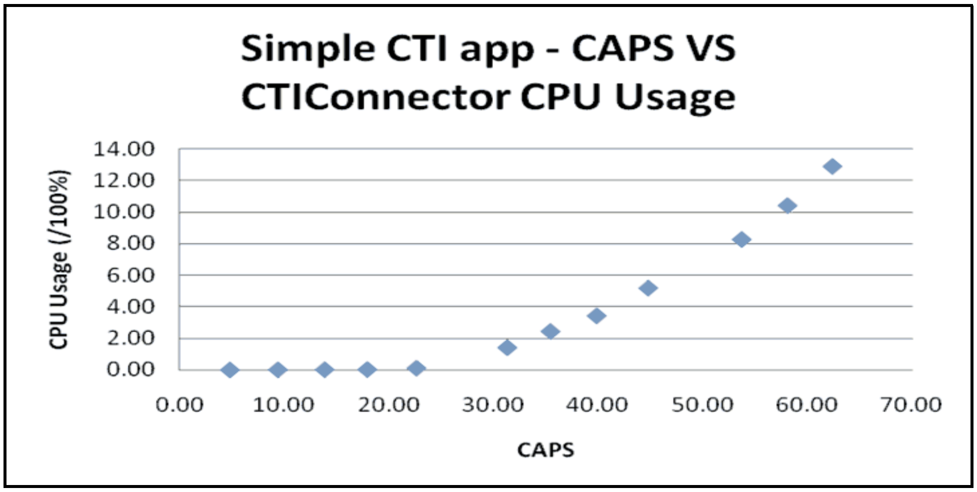

Figure 99: CAPS versus CPU Usage (CTI Connector)

Resource Manager Test Case

Resource Manager performance is measured in terms of CAPS. Performance is not affected by the number of simultaneous calls held by Resource Manager. Resource Manager performs most efficiently when multiple Media Control Platforms are used.

The affect on Resource Manager performance differs, depending on the type of call being processed (for example, conference versus announcement calls), but generally, a peak of 500 CAPS can be sustained regardless of the call type. The same is true even if the Resource Manager is configured for High Availability.

CPU consumption on the Resource Manager is very low. The 500 CAPS limit mentioned previously is due to the use of virtual memory, which exceeds the 2 GB limit when Resource Manager is running consistently beyond 500 CAPS.

The graph in Figure 100 depicts Resource Manager CPU usage.

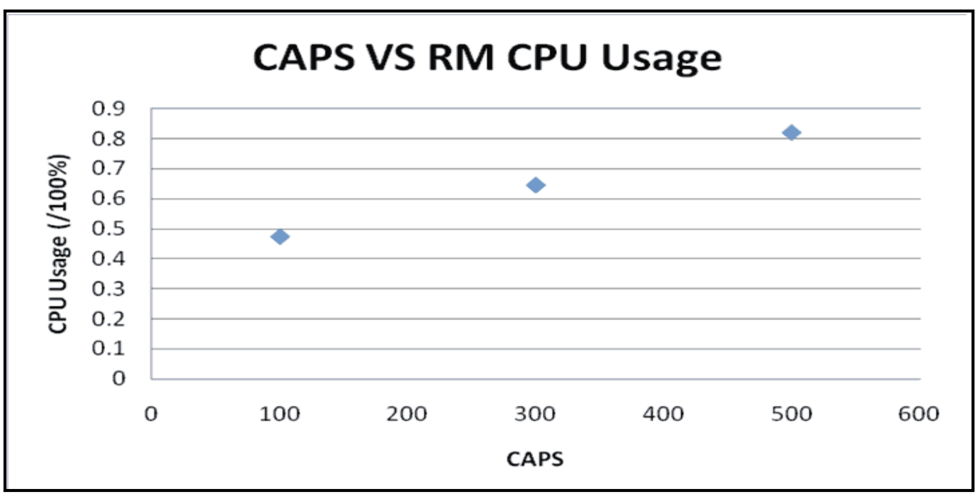

Figure 100: CAPS Versus CPU Usage (Resource Manager)

In Figure 101, Resource Manager uses only <1% of the CPU cycles at 500 CAPS.

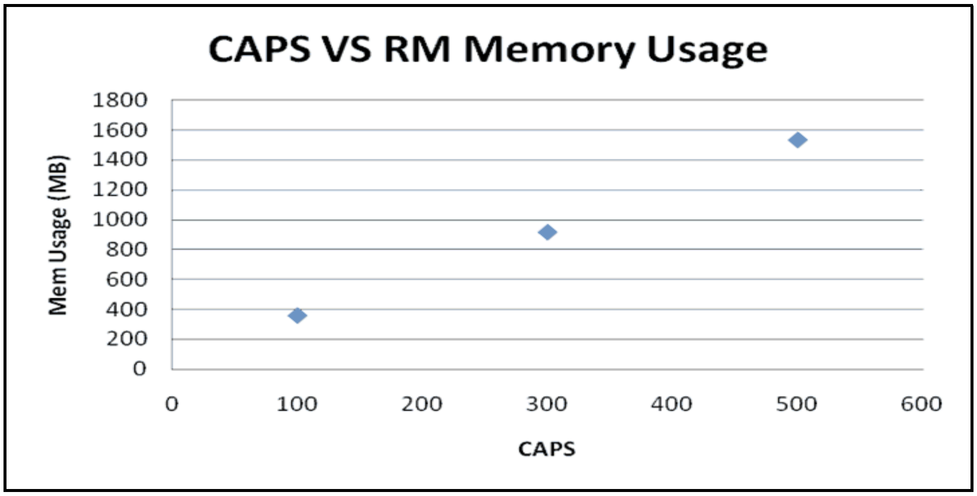

Figure 101: CAPS Versus Memory Usage (Resource Manager)

It is possible for Resource Manager to sustain 600 CAPS, however, the memory usage almost reaches the 2 GB limit and a 500 CAPS peak capacity is more appropriate.

When configured with the Reporting Server, Resource Manager sustained 500 to 600 CAPS, but the Reporting Server performance invariably caused a bottleneck to occur. See "Reporting Server Test Case".

Reporting Server Test Case

Like the Resource Manager, Reporting Server performance is measured in terms of CAPS. The number of simultaneous calls being processed by GVP does not affect performance and there are no known performance bottlenecks with the Reporting Server software, however, performance can be affected by the database setup.

A use case was conducted on a Microsoft SQL database server with the Resource Manager and the Media Control Platform streaming information to the Reporting Server for each call with a default log level. The result was a peak capacity of 120 CAPS.

A second use case was conducted on an Oracle database server using the same hardware and system configuration as in the first use case. The result was a peak capacity of 60 CAPS.

Note: 15 disk eSAS SAN and a Dual Quad Core Xeon, 2.66 GHz CPU were used on both the Microsoft Sequel Server (MSSQL) and Oracle Database Servers. The Reporting Server was installed on a Dual Quad Core Xeon server with a 2.66 GHz CPU, separate from the database server

Regardless of the database server being used, approximately 824 bytes of data is stored for each call that is accepted by GVP (through the Resource Manager to the Media Control Platform).

The following data throughput statistics are typically stored in the database:

VAR_CDRS	> 49 bytes
MCP_CDR	> 171 bytes
RM_CDR	> 169 bytes
EVENT_LOGS	> 435 bytes

~19 transactions per call are reported by Microsoft SQL Server.

~50 transactions per call are reported by Oracle Database Server.

Performance and Scalability Comparisons

GVP 8.x contains many performance enhancements not included in GVP 7.6 or VoiceGenie (VG) 7.2.

Performance Comparisons

In this section, performance and scalability comparisons are made between GVP 8.x and previous releases using the application profiles in Table 114 on page 233.

Tested with VoiceXML_App1

Peak capacity of GVP 8.x:

- with NGI
  - ~50% higher than VG 7.2.
  - ~90% higher than GVP 7.6.
- with GVPi
  - equivalent to GVP 7.6.

Comparing GVP 8.x and GVP 7.6 (with GVPi) to GVP 8.1:

- 8.1 uses significantly fewer CPU cycles (relatively 30%).
- 8.1 uses less memory (relatively 30%).
- In 8.1, the peak capacity is identical to previous releases (using identical temp file management mechanisms), as the bottleneck is due to disk IO.

Tested with VoiceXML_App2

Peak capacity of GVP 8.x:

- with NGI
  - ~66% higher than VG 7.2.
  - ~100% higher than GVP 7.6.
- with GVPi
  - equivalent to GVP 7.6.

In the use case with GVPi, the peak capacity for GVP 8.x is identical to GVP 7.6 (using identical temp file management mechanisms), because the bottleneck is due to disk IO.

Scalability Comparisons

For applications that are CPU-dependent (or applications in which bottlenecks occur due to CPU cycles) GVP 8.x can use additional CPU cycles and cores. Use case results showed that peak port densities scaled upward linearly relative to an increase in CPU clock speed.

Table 116 contains examples of peak capacity when `VoiceXML_App1` is used:

Table 116: Peak Capacity—VoiceXML_App1

Processor	Total Clock Speed	Peak Port Density
2x Core 2 Quad, 2.66 GHz	21.28 GHz	1300
2x Core 2 Dual, 3.00 GHz	12 GHz	700
1x Core 2 Dual, 3.00 GHz	6 GHz	400

Figure 102 is a graphical depiction of the peak port density in Table 116.

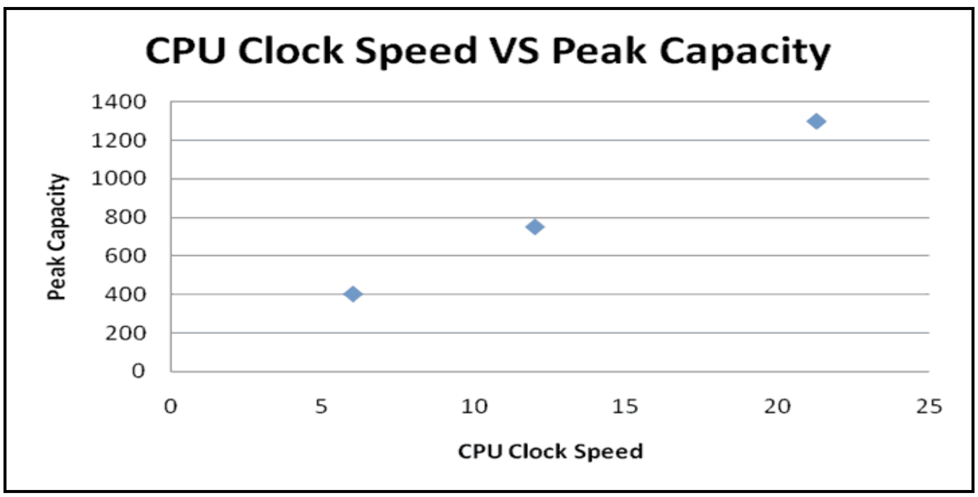

Figure 102: CPU Clock Speed Versus Peak Capacity

To increase the total clock speed by 100%, the peak capacity would have to increase by ~90 to 100%, assuming:

- the type of CPUs are the same as the ones in Table 116.
- the `VoiceXML_App1` application is used.
- the overall system bottleneck CPU cycles remain the same.

High Performance Configuration

The Media Control Platform can support more than 400 ports on a single host, however, some configuration changes are required. Use Genesys Administrator to configure the Media Control Platform for high performance by modifying the options and default values in Table 117, and configure the Windows Registry on the Media Control Platform to support either the NGI, GVPi, or both.

Table 117: High Performance Configuration for the Media Control Platform

Section	Option/Key	Default Value	High Performance Value
Media Control Platform with NGI			
mcp	`maxmediathreads`	32	16
vxmli	`max_num_documents`	2000	6000
Windows Registry key: `HKEY_LOCAL_MACHINE\CurrentControlSet\Services\Tcpip\Parameteres\TimedWaitDelay`		None	Type = DWORD Value = 30 or 1e (hex)
Media Control Platform with GVPi			
mcp	`maxmediathreads`	32	32
PageCollector	`maxpoolthreads`	512	>= Port Density
PopGateway1	`maxpoolthreads`	512	>= Port Density
Windows Registry key: `HKEY_LOCAL_MACHINE\Software\CallNet\CnlnetSettings\MaxThreadPool`		None	Type = DWORD Value >= Port Density /2

Chapter

10 Genesys Integration Server

This chapter presents hardware sizing recommendations for selected Genesys Integration Server (GIS, formerly known as Genesys Interface Server) 7.x services. It covers the following topics:

Overall Recommendations

To match or exceed the performance standards reported in this chapter, Genesys recommends that you use hardware that is at least as powerful as that in the referenced environment.

To maximize scalability, Genesys recommends that you run multiple instances of GIS on separate computers (to the extent that network bandwidth and latency are not constraints).

GIS Interaction Service

This section provides general recommendations about hardware sizing for the GIS voice interaction service.

Recommended Environment for Voice Interaction Service

Table 118 shows the referenced hardware and software environment.

Table 118: Referenced Hardware and Software Environment

Microsoft Windows 2000/2003	
Hardware Specification	**Recommendation**
Processor Quantity, Type, Speed, and Cache	Dual Intel Xeon 2.4 GHz with 1 GB L2 cache
Memory Size	4 GB
Operating System	Microsoft Windows 2000 Advanced Server, Microsoft Windows 2003 Advanced Server
Application Server	Tomcat 4

Environment Limitations

Number of agent desktops per GIS instance: up to 500

Recommended browser HTTP mode: HTTP 1.1

Note: These are general recommendations. Genesys cannot anticipate all possible production environments. If you have a very complex configuration, please consult Genesys Professional Services for specific sizing recommendations.

Interaction Service: Performance Test Statistics

Genesys conducted tests to define the maximum performance of GIS 7.5 with both SOAP and GSAP connectors working with phone calls.

Test Environment

Table 119 details the test environment used to arrive at the recommended limits for deploying the GIS Interaction Service.

Note: One basis for understanding the optimized results of all the tests in this section is the use of a dedicated 1-Gigabit network in the lab environment.

Table 119: Interaction Service Test Environment Components

Component	Host	Hardware	Operating System
Configuration Server 7.5 T-Server G3 7.5 Switch Simulator 6.5.300.02	Host_01	2x Opteron Dual Core; 2.2 GHz/cache 512 KB; 4 GB memory; dedicated 1-Gigabit Ethernet	Microsoft Windows 2003 Server
GIS 7.5	Host_04 Host_05		Microsoft Windows 2003 Server; (J2SDK), Standard Ed., 1.5.0_09
Universal Contact Server (UCS) 7.5 Interaction Server 7.5	Host_02		Microsoft Windows 2003 Server
Microsoft SQL Server 2000 (UCS Database)	Host_03		Microsoft Windows 2003 Server

Table 120 lists the agent simulators used in the test environment.

Table 120: Interaction Service Simulators Used

Component	Host	Hardware	Operating System
.NET-based agent simulators	Host_06 Host_07 Host_08 Host_09 Host_10	2x Opteron Dual Core; 2.2 GHz/cache 512 KB; 4 GB memory; dedicated 1-Gigabit Ethernet	Microsoft Windows 2000 Advanced Server

Test Environment Topology

Figure 103 on page 254 illustrates the test environment's specific topology.

Note: `Host_05` is used to run a second instance of GIS for the purpose of High Availability, but High Availability is not covered in these GIS details.

Also note that GIS is not connected to the Multimedia servers for this testing.

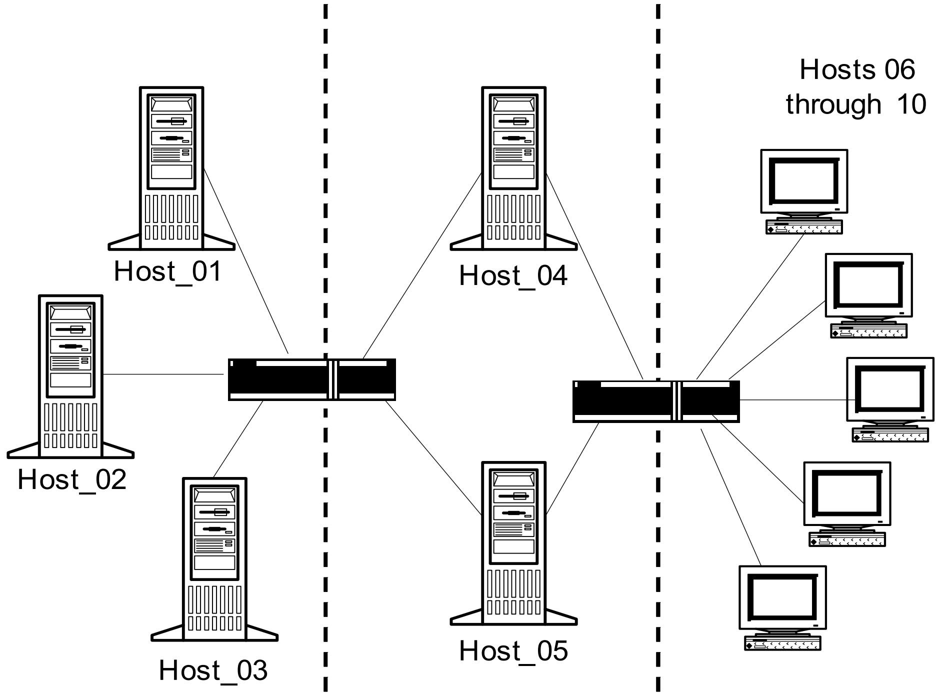

Figure 103: Topology of the Interaction Service Test Environment

Test Tools

The agent simulators were built with Genesys Integration Server 7.5 Proxies/Libraries.

Test Procedure

The performance test was executed for three differently sized contact centers, 100, 200, and 500 agents.

In each case the agents worked according to the following scenario:

1. Log in to queue 8001.
2. Become `Ready`.
3. Wait for an incoming call.
4. Answer the incoming call.
5. Become `Not Ready`.
6. Wait 10 seconds before releasing the call.
7. Wait two seconds.
8. Become `Ready`.
9. Wait for the next call.
10. Repeat these steps.

The test lab generated calls with a switch simulator for queue 8001 at a rate of up to 80 calls per second. The maximum calls-per-second rate was used to measure test runs with 100, 200, and 500 agents who were placed in push notification mode. In addition to measuring the time it took to answer calls, this operation defined system reactivity.

Test Results

This section details the test results for both GIS:SOAP and GIS:GSAP. Table 121 shows the values for the main performance parameters that were measured for GIS:SOAP in the process of testing. In all cases clients were set to work in push notification mode.

Table 121: Performance Test Results for GIS:SOAP

Number of Agents in Test	CPU Usage Percent	MBytes Per Second Over the Network Interface	Calls Per Second	Time to Answer Call (in msecs)[a]
100	3	0.45	7.3	10
200	12	0.9	15.3	10
500	26	2.1	36	10

a. This measurement was done on the server side through internal counters.

The following system settings were applied to be able to reach these numbers for GIS:SOAP:

- JVM 1.5.0_09 tuning: `-Xmx512M -Xss128K`
- Tomcat tuning: `MaxThreads = 1000`

Table 122 shows the results for GIS:GSAP under the same conditions as applied to GIS:SOAP.

Table 122: Performance Test Results for GIS:GSAP

Number of Agents in Test	CPU Usage Percent	MBytes Per Second Over the Network Interface	Calls Per Second	Time to Answer Call (in msecs)
100	4	0.1	7.4	10
200	7	0.3	15.7	10
500	12	0.7	38.2	10

Note: Genesys performed the same tests with GIS 7.2. The main differences are that GIS 7.5 consumes about 10% less CPU than 7.2, and that, though the numbers are not shown here, at 1000 agents, 7.2 experiences an `OutOfMemory` error.

Interaction Service Test Conclusions

The GIS:SOAP connector can support 500 agents in push notification mode with a call flow rate equal to 65 calls per second, with moderate CPU consumption (given the software and hardware specifications described here for the test environment).

The GIS:GSAP connector can support 500 agents in push notification mode with a call flow rate equal to 72 calls per second, with low CPU consumption (given the same software and hardware specifications for the test environment).

In the end, the primary bottleneck is the ability to simulate more agents. It is possible that the performance of this solution might improve with more agents, but testing for that is difficult.

Based on the results here, we can conclude that:

- with machines and a network that can handle the necessary load, GIS:SOAP can support up to 500 agents at a comparable call rate to that supported by GIS:GSAP.
- with new mutt-core architectures suitable for Java applications (and their inherent multi-threaded natures) we see a significant (approximately 50%) improvement over GIS 7.2 numbers (which resulted from tests on single-core processors).

GIS Configuration Options

This section displays the GIS options used for this performance testing.

```
In GIS application:
[ail-services]
agent-wait-status-stable=100
cache-lease-time=30
interaction-voice-create-new-timetolive=30
timeout=30

[core-services]
events-adapter-soap-commons-sender-enabled=true
events-adapter-soap-commons-sender-http11-mode=true
events-adapter-soap-commons-sender-max-connections-host=101
events-adapter-soap-commons-sender-max-total-connections=500
events-adapter-soap-invokeonewaymode=false
```

```
events-adapter-soap-notification-failure-nbretries=3
events-adapter-soap-notification-polling=true
events-buffered-period=100
events-buffered-size=201
events-keep-alive-time=5
events-maximum-pool-size=100
events-minimum-pool-size=10
events-polling-one-by-one=false
events-subscriber-timetolive-polling=10
events-thread-priority=7
jmx-html-adapter-port=9002

[GSAP.General]
display_nb_threads=true
ha-count-servers=2
ha-server-no=1
server.backlog=1000
server.clientsocket.enablekeepalive=true
server.clientsocket.enabletcpnodelay=true
server.clientsocket.sendbufsize=524288
server.clientsocket.sotimeout=10000
server.handlers.max=2000
server.handlers.recomended=500
server.process.interval=30
use-cluster-ha=false

[GSAP.KeepAlive]
handler-daemon-audit-interval=5000
handler-daemon-inactivity-interval=1000
ping-daemon-audit-interval=3000
ping-daemon-inactivity-interval=500
ping-times-threshold=3
session-timeout-threshold=3600000

[license]
gis_configurationservice-block-size=1
gis_interactionservice-block-size=13
gis_statservice-block-size=1
license-file=7260@<License_Host>

[log]
all=../logs/gis_all.log,stdout
Buffering=false
debug=..\logs\gis_debug.log
expire=20
ha-settings=dump-cache-short
lb-console=error
lb-file=error, gis-gsap-lb.log, 20, 3
lb-msgsrv=off
MessageFile=
prop-console=error
```

```
prop-file=error,gis-gsap.log, 20, 3
prop-msgsrv=off
segment=100000
standard=../logs/gis_std.log,stdout
trace=
verbose=error

[SessionService]
sessionTimeout=3600

[StatService]
error_check=true
restriction_time=30
scopeStatEvents=15
```

The following are the options configured for the AIL-based applications connected to GIS:

```
[dn-at-switch]
enabled=false

[kworker]
auto-markdone=true
easy-newcall=true
reroute=false
reroute-location=<location>

[license]
attempts-interval=5
attempts-max=10
license-file=7260@<License_Host>

[loading]
srl-on-demand=false

[log]
console=error
file=warn,ail_gsap_,10M

[multimedia]
chat-addp-protocol=false
chat-busy-threshold=1
collaboration-mode=push
collaboration-workbin=desktop-collaboration-email-workbin
email-busy-threshold=1
email-default-queue=email-default-queue
email-drafts-workbin=draft
email-outbound-queue=email-outbound-queue
email-quote-char=>
email-quote-header=On <date>, <contact> wrote:
```

```
email-trsf-ext-queue=
qa-review-skill-name=

[server]
type=ail

[settings]
ucsapi-max-active-rmi-client=30
ucsapi-max-timeout-rmi-client=30

[signature]
include-agent-name=true
line=
prefix=--

[voice]
database=all
dms-last-digits=-1
enable-all-routing-events=false
enable-attached-data-for-transfer=true
enable-interaction-id-tracking=true
enable-possible-changed-event=true
```

GIS Statistics Service—Microsoft Windows 2000/2003

This section provides general recommendations about hardware sizing for the GIS Statistics Service on the Microsoft Windows 2000/2003 Advanced Server platforms.

Recommended Environment for Statistics Service

Table 123 shows the referenced hardware and software environment.

Table 123: Referenced Hardware and Software Environment

Microsoft Windows 2000/2003	
Hardware Specification	**Recommendation**
Processor Quantity, Type, Speed, and Cache	Dual Intel Xeon 2.4 GHz with 1 GB L2 cache
Memory Size	4 GB

Table 123: Referenced Hardware and Software Environment (Continued)

Microsoft Windows 2000/2003	
Hardware Specification	**Recommendation**
Operating System	Microsoft Windows 2000 Advanced Server, Microsoft Windows 2003 Advanced Server
Application Server	Tomcat 4

Environment Limitations

Real-Time Reporting Mode

Minimum statistics update interval: 2 seconds

Number of connected clients: up to 15

Total statistics update per second, per GIS instance: up to 300

Maximum statistics update per second, per client: 63

Recommended browser HTTP mode: HTTP 1.1

Historical Reporting Mode

Statistics update interval: 15 minutes

Number of connected clients: 2

Total statistics updates per interval, per GIS instance: up to 52,000

Recommended browser HTTP mode: HTTP 1.1

Note: These are general recommendations. Genesys cannot anticipate all possible production environments. If you have a very complex configuration, please consult Genesys Professional Services for specific sizing recommendations.

Statistics Service for Microsoft Windows: Performance Test Statistics

Genesys conducted tests to define how many messages per second GIS 7.2 could handle, and how many agents it could support if the agent applications worked in `unsolicited` notification mode.

Test Environment

Table 124 details the test environment used to arrive at the recommended limits for using the GIS Statistics Service in a Microsoft Windows deployment.

Table 124: Statistics Service for Microsoft Windows Test Environment Components

Component	Host	Hardware	Operating System
Genesys Integration Server, 7.2.000.15, with SOAP connector	Host_01	2 x Intel Xeon 2.4 GHz/cache 512k; 2 GB memory; 2 Gigabit Ethernet	Microsoft Windows 2000, SP4
.NET-based client simulators	Host_02		Microsoft Windows 2000, SP4
.NET-based client simulators	Host_03		Microsoft Windows Server 2003 Enterprise edition, SP1
Stat Server 7.2.000.10	Host_01		Microsoft Windows 2000, SP4

Test Environment Topology

Figure 104 illustrates the GIS Statistics Service test environment's specific topology.

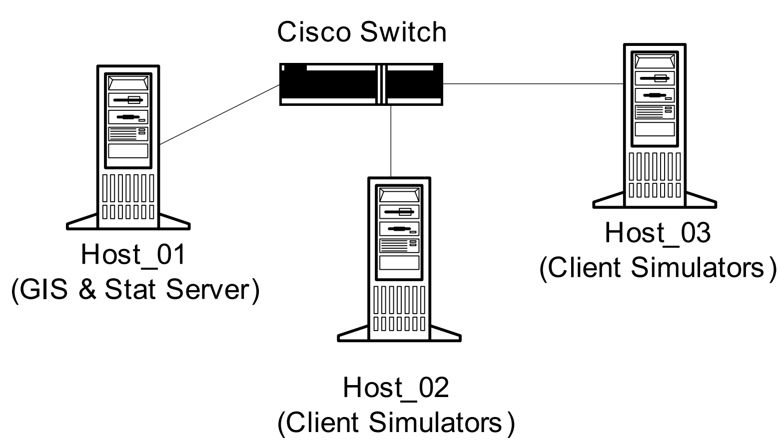

Figure 104: Topology for Statistics Service for Microsoft Windows Test Environment

Statistics Service for Microsoft Windows Test Conclusions

In the first tested configuration, with GIS using its Statistics Service in unsolicited notification mode, it was able to support stable work for 40 clients when each of them subscribed to 12.5 notifications per second (25 notifications every 2 seconds).

A second test revealed that a maximum of 63 notifications per second can be delivered to a client, allowing GIS to support work for one client that is processing 63 notifications per second. But another test demonstrated that when 10 clients were connected, GIS was able to send them no more than 32 messages per second.

One instance of GIS is able to deliver 52,000 statistics in a 15-minute subscription interval.

Note: These results are for .NET clients working in HTTP 1.1 mode only. Results for others client implementations may be differ.

Modules.conf Values

The following are the contents of `modules.conf` for the tests conducted in this section:

```
<?xml version="1.0" encoding="utf-8"?>
<modules>
   <module name="logging"
class="com.genesyslab.gis.modules.logging.LoggingModule">
   <file>logging.conf</file>
   </module>
   <module name="license"
class="com.genesyslab.gis.modules.license.LicenseModule"/>
   <module name="lca"
class="com.genesyslab.gis.modules.lca.LcaModule"/>
   <module name="notification"
class="com.genesyslab.gis.modules.notification.NotificationModule">
      <maxThreads>500</maxThreads>
      <minThreads>10</minThreads>
      <priority>7</priority>
      <http_version>HTTP 1.1</http_version>
   </module>

</modules>
```

GIS Statistics Service—Solaris 9

This section provides general recommendations about hardware sizing for the GIS Statistics Service on the Solaris 9 platform with WebSphere 5.1.

Recommended Environment for Statistics Service

Table 125 shows the referenced hardware and software environment.

Table 125: Referenced Hardware and Software Environment

Solaris 9	
Hardware Specification	**Recommendation**
Processor Quantity, Type, and Speed	2 CPU 1.5 GHz UltraSPARC IIIi. (e.g. Sun Netra 240)
Memory Size	4 GB
Operating System	Solaris 9
Application Server	WebSphere 5.1

Environment Limitations

Real-Time Reporting Mode

Minimum statistics update interval: 2 seconds

Number of connected clients: up to 14

Total statistics update per second, per GIS instance: up to 200

Maximum statistics update per second, per client: 90

Recommended browser HTTP mode: HTTP 1.1

Historical Reporting Mode

Statistics update interval: 15 minutes

Number of connected clients: 1

Total statistics updates per interval, per GIS instance: up to 25,000

Recommended browser HTTP mode: HTTP 1.1

Note: These are general recommendations. Genesys cannot anticipate all possible production environments. If you have a very complex configuration, please consult Genesys Professional Services for specific sizing recommendations.

Statistics Service for UNIX: Performance Test Statistics

Genesys conducted tests to define how many messages per second GIS 7.2 could handle, and how many agents it could support if the agent applications worked in `unsolicited` notification mode.

Test Environment

Table 126 details the test environment used to arrive at the recommended limits for using the GIS Statistics Service in a UNIX (Solaris 9) deployment.

Table 126: Statistics Service for Solaris 9 Test Environment Components

Component	Host	Hardware	Operating System
Genesys Interface Server 7.0.101.27 (name changed with 7.2.x release)	Solaris_Host	Sun Fire V240: 2 x UltraSPARC III 1000 MHz, 2 GB memory, 1 Gigabit Ethernet	Solaris 9 WebSphere 5.1
Component	**Host**	**Hardware**	**Operating System**
.NET-based client simulators	Host_02	2 x Intel Xeon 2.4 GHz/cache 512k; 2 Gigabit Ethernet	Microsoft Windows 2000 Advanced Server
.NET-based client simulators	Host_03		
Stat Server 7.0.000.21	Host_04		

Test Environment Topology

Figure 105 illustrates the GIS Statistics Service test environment's specific topology.

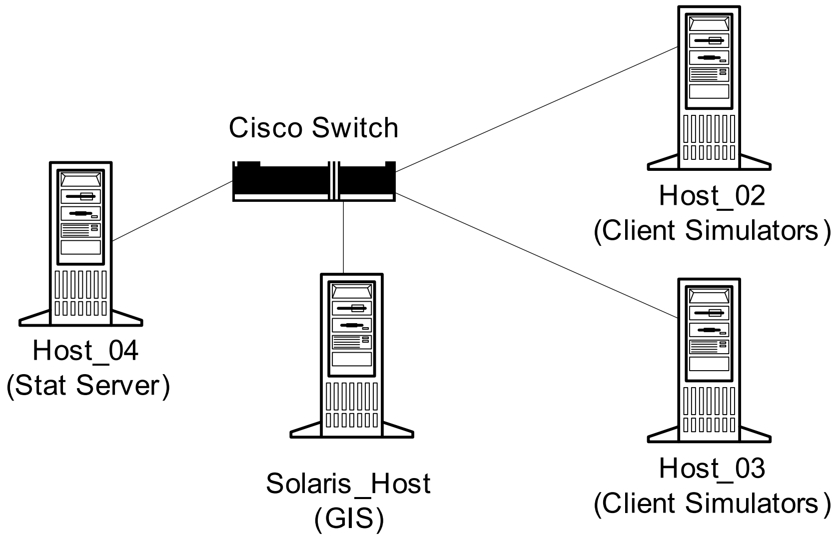

Figure 105: Topology for Statistics Service for Solaris Test Environment

Statistics Service for Solaris Test Conclusions

One instance of GIS is able to deliver 52,000 statistics in a 15-minute subscription interval.

In the first tested configuration, with GIS using its Statistics Service in `unsolicited` notification mode, it was able to support stable work for 14 clients when each of them subscribed to 12.5 notifications per second (25 notifications every 2 seconds). GIS delivered messages at a rate of about 175 per second.

A second test revealed that a maximum of 90 notifications per second can be delivered to a given client.

Other tests showed that GIS is able to process about 180 messages per second.

GIS can deliver 25,000 statistics in a 15-minute interval.

Modules.conf Values

The following are the contents of `modules.conf` for the tests conducted in this section:

```
<?xml version="1.0" encoding="utf-8"?>
<modules>
   <module name="logging"
class="com.genesyslab.gis.modules.logging.LoggingModule">
   <file>logging.conf</file>
   </module>
   <module name="license"
class="com.genesyslab.gis.modules.license.LicenseModule"/>
```

```
   <module name="lca"
class="com.genesyslab.gis.modules.lca.LcaModule"/>
   <module name="notification"
class="com.genesyslab.gis.modules.notification.NotificationModule">
      <maxThreads>500</maxThreads>
      <minThreads>100</minThreads>
      <priority>7</priority>
      <http_version>HTTP 1.1</http_version>
   </module>

</modules>
```

Note: The parameter `error_check` in the Configuration Layer had a value of `false` for these tests.

Chapter

11 SIP Server Solution

This chapter provides hardware sizing guidelines for deploying the SIP Server solution releases 7.6 and 8.0. It contains the following sections:

Introduction

This chapter provides hardware sizing guidelines and the basic information required for deploying and capacity planning of a Voice over Internet Protocol (VOIP) solution based on SIP Server releases 7.6 and 8.0. It covers the following topics:

- Required hardware resources (the number of physical host computers)
- Number of applications (SIP Servers and Stream Managers)
- Recommended network architecture
- Application placement across the host computers
- Expected system loads (CPU, memory, and traffic per network interface)

This chapter also provides benchmarks for the following basic scenarios:

- SIP Inbound Call
- Transfer
- Recording
- Conference
- Consultation Call
- Internal Call (using third-party call control [3pcc])

Running a particular scenario involves multiple system components. For a reference system architecture, see “Reference Architecture” on page 268.

For the benchmark data obtained for the reference scenarios, see “Scenarios and Benchmarks” on page 270.

This chapter also provides performance data for SIP Server. SIP Server performance is measured for variable call rates. Other important parameters are the number of concurrent sessions and the complexity of the scenarios. The results presented in this chapter are based on lab benchmarks obtained on a 32-bit version of Microsoft® Windows® Server 2003 Service Pack 2 (SP2), and an Intel® CPU. For details of the hardware platform that was used, see “Hardware Details” on page 290.

To match or exceed the performance levels described in this chapter, Genesys recommends that you use hardware that is at least as powerful as that in the lab environment.

Reference Architecture

Figure 106 depicts the generic architecture of a VOIP solution that was deployed for benchmarking.

Incoming SIP calls that are to be balanced across multiple instances of SIP Server are processed by the Network SIP Server. Media processing (RTP traffic) is handled by multiple Stream Managers. Typically, Stream Managers support treatments (for example, Music On Hold), conferences, and call recording. Media gateways provide media conversion between the TDM (Time-Division Multiplexing) and VOIP domains. The media gateway is a third-party component.

High-availability (HA) capability of the VOIP solution is achieved through the introduction of redundant backup components (shown as shadow boxes in Figure 106). The backup components should be deployed on separate physical hosts, different from the hosts on which the primary components are running. This chapter does not discuss the impact of HA.

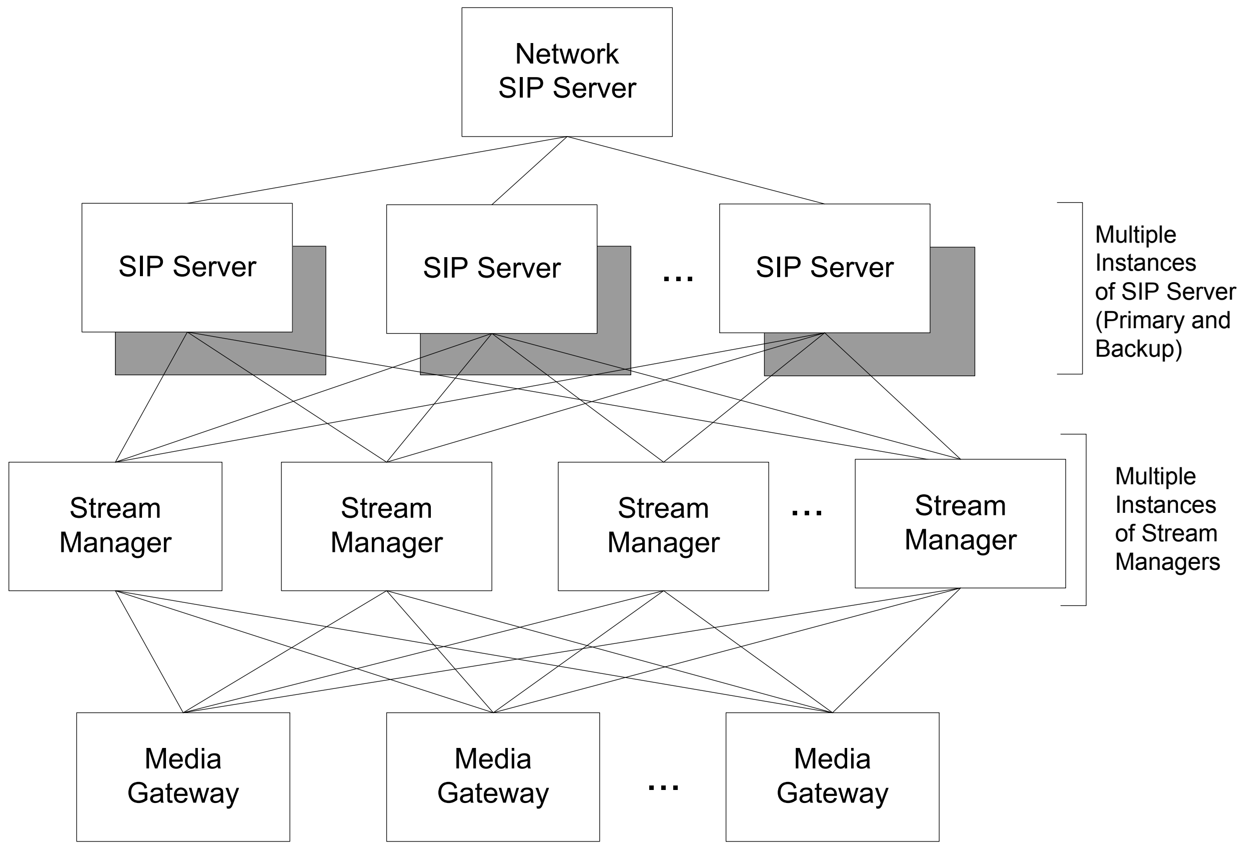

Figure 106: Reference Architecture for VOIP Solution

For proper sizing, this architecture assumes the following constraints:

- No more than one server application per physical CPU (core) is deployed.

  For example, a dual Quad-Core Intel Xeon® host computer runs no more than eight server applications, either SIP Servers or Stream Managers.
- SIP Server and Stream Manager must be placed on separate physical host computers.
- The total number of open TCP/UDP ports per host should not exceed 10,000.
- The default socket buffer size is set to at least 64 KB.
- To prevent voice quality degradation, the amount of the network traffic (the sum of the sent/received bytes) per network interface should be limited to 100 Mbit/sec, and to 250 Mbit/sec per network segment.
- TCP/IP transport is used for SIP communication between SIP Server and Stream Manager(s).

The quality of the network interface cards (NIC) and NIC drivers is important for achieving optimal performance. The hardware buffer size of the NIC should be at least 64 KB. For example, Genesys recommends using the following NICs on Windows platforms:

- Broadcom BCM5708C NetXtreme II GigE
- Intel E1000 (Intel PRO/1000 Family)

Genesys recommends deploying the latest drivers that are available from the NIC vendor. The default operating system (OS) drivers may not be optimal.

Scenarios and Benchmarks

This section discusses benchmarks for the following reference scenarios:

- SIP Inbound Call (first-party call control [1pcc])
- Transfer
- Recording
- Conference
- Consultation Call
- Internal Call (using 3pcc)

These scenarios are considered to be typical for a VOIP solution. For each scenario, this chapter provides the dependencies between the call rate (measured in calls per second [calls/sec]) and CPU usage. These dependencies were evaluated for several SIP Server operating conditions, with the major parameter being the number of concurrent sessions. SIP Server performance was evaluated for the following set of values for this parameter:

- 500 concurrent sessions
- 1,000 concurrent sessions
- 1,500 concurrent sessions

Table 127 indicates the typical number of SIP dialogs, SIP messages and T-Library messages for each of the reference scenarios. Subsequent sections describe how to estimate the CPU usage, memory usage and network traffic associated with each scenario, and provide detailed message flows for each scenario.

Table 127: Reference Call Flow Details

Scenario Type	Number of SIP Dialogs	Number of SIP Messages	Number of T-Library Requests/Events
SIP Inbound Call	2	14	20
Transfer	3	24	48
Recording	4	36	23

Table 127: Reference Call Flow Details (Continued)

Scenario Type	Number of SIP Dialogs	Number of SIP Messages	Number of T-Library Requests/Events
Conference	6	43	44
Consultation Call	3	39	56
Internal Call (using 3pcc)	3	25	32

CPU Usage Estimation

Based on the lab measurements, CPU usage for a given number of concurrent sessions appears to be a linear function of call rate. Due to the effective implementation of a caching mechanism in SIP Server, the dependency between CPU usage and the number of concurrent sessions is relatively small. The following is the recommended procedure for evaluating SIP Server performance when running mixed call flows.

Procedure: Estimating CPU Usage

Purpose: To estimate the required number of SIP Server instances and CPU loads, given an arbitrary mix of reference scenarios.

Start of procedure

1. Express the specific call flows in terms of the reference scenarios.

   For the details of the scenarios, see "Reference Call Flows" on page 279.
2. For each scenario, specify the average call rate (in calls/sec) and average call duration (in seconds). Calculate the average number of concurrent calls per scenario by using this formula:

   $$\mathtt{ConcurrentCalls}_i = \mathtt{CallRate}_i * \mathtt{CallDuration}_i$$

3. Calculate the total number of concurrent calls for all scenarios:

   $$TotalConcurrentCalls = \Sigma\ ConcurrentCalls_i$$

   Check that the calculated value for `TotalConcurrentCalls` does not exceed 2,000. If it does, you must use multiple instances of SIP Server, and a Network SIP Server for load balancing, to support your requirements for call volume.

   To estimate the number of SIP Server instances required, use this formula:

   `nSIPServers = Max Integer (TotalConcurrentCalls / 2,000)`

   As shown in the formula, the result must be rounded to the next highest integer value. The number of SIP Server instances per host must not exceed the number of CPU cores. At maximum load, each SIP Server requires up to 300 MB of resident memory, so the latter constraint is bound by the amount of memory that is available to applications on a 32-bit version of Windows Server 2003 (default user space of 2 GB, assuming 4 GB of physical RAM).

4. If you are using more than one SIP Server, calculate the effective call rate per instance:

   $$EffectiveCallRate_i = CallRate_i\ /nSIPServers$$

5. For each scenario and effective call rate, estimate CPU usage by using an approximation formula for the corresponding call flow. Table 128 contains these formulas, which take form $CPU_i = K_i * CallRate_i$, where the index ($_i$) corresponds to a specific call flow (scenario).

Table 128: Calibration Table for CPU Usage (for Intel Xeon CPU, 2.33 GHz)

Scenario Type	Formula for CPU Usage CPU(%) vs. Call Rate (Calls/Sec)	Applicable Call Rate (Calls/Sec)
SIP Inbound Call	CPU = 0.69 * EffectiveCallRate	0–90
Transfer	CPU = 1.42 * EffectiveCallRate	0–40
Recording	CPU = 2.25 * EffectiveCallRate	0–30
Conference	CPU = 2.52 * EffectiveCallRate	0–25
Consultation Call	CPU = 2.33 * EffectiveCallRate	0–35
Internal Call (using 3pcc)	CPU = 1.05 * EffectiveCallRate	0–50

Notes:
- The typical margin of error is 10 percent.
- For simplicity, the calibration table (Table 128 on page 272) assumes the maximum number of concurrent sessions and a CPU clock frequency of 2.33 GHz. For Intel Xeon CPUs with an L2 cache size above 2 MB, performance of the application scales linearly with the CPU clock frequency.

If the value of parameter `EffectiveCallRate` is out of the applicable range indicated in Table 128, you must deploy additional instances of SIP Server and recalculate `EffectiveCallRate` as follows:

a. Using the calibration table (Table 128), calculate the total CPU usage by summing individual values:

$\mathtt{TotalCPU} = \sum \mathtt{CPU}_i$

b. If a Network SIP Server is present in the configuration, use the following formula:

$\mathtt{TotalCPU} = 1.20 * \sum \mathtt{CPU}_i$

This formula takes into account the additional overhead associated with the additional messaging that the Network SIP Server introduces. If `TotalCPU` exceeds 70 percent, you must deploy additional instances of SIP Server. Increment the number of SIP Servers (parameter `nSIPServers`), recalculate `EffectiveCallRate`, and then use the formulas in Table 128 to find the CPU usage, repeating this process until `TotalCPU` drops below 70 percent.

The preceding estimates are valid when all SIP Server clients share a single T-Library connection to the server (for example, Genesys Agent Desktop is deployed). If clients use individual multiple connections to SIP Server, a connection correction factor (CCF) must be applied to `TotalCPU`, as shown in Table 129. The resulting CPU usage is calculated as follows:

```
CorrectedTotalCPU = CCF * TotalCPU
```

This new parameter, `CorrectedTotalCPU`, must be used instead of parameter `TotalCPU` in the calculations for estimating the number of the SIP Server instances (`nSIPServers`); however, such iteration is not usually required.

Table 129: Correction Factor for Multiple Connections

Number of Client Connections	CCF
1–100	1.00
100–500	1.02
500–000	1.05
1000–2000	1.10

End of procedure

CPU Clock Scaling

The results for CPU load estimation are given for the maximum number of concurrent sessions and a CPU clock frequency of 2.33 GHz. For Intel Xeon CPUs with an L2 cache size above 2 MB, performance of the application scales linearly with the clock frequency. For example, if you are upgrading from Intel Xeon model L5410 (Quad Core, 2x6 MB cache, 2.33 GHz, 1333 MHz FSB (Front-Side Bus) to Xeon model X5460 (Quad Core 2x6 MB cache, 3.16 GHz, 1333 MHz FSB), you should expect a performance boost of about 35 percent (3.16 GHz/2.33 GHz). In this example, the scaling coefficients in Table 128 on page 272 must be *reduced* by 35 percent accordingly.

Memory Usage

Based on lab measurements, a single instance of SIP Server uses up to 300 MB of resident memory on a 32-bit Windows Server 2003 operating system. This value assumes the maximum number of concurrent sessions and the maximum call rate. Less memory is used under a lighter load.

Network Traffic Estimation

Table 130 provides formulas for estimating the amount of incoming `(RX)` and outgoing `(TX)` traffic per SIP Server instance and specific scenario. When using Table 130, you must use parameter $\texttt{EffectiveCallRate}_i$ for each scenario (denoted by index $_i$), as calculated in Step 4 on page 272.

Table 130: Calibration Table for Network Traffic

Scenario Type	Formula for Incoming Traffic (RX): KB Received vs. EffectiveCallRate (Calls/Sec)	Formula for Outgoing Traffic (TX): KB Sent vs. EffectiveCallRate (Calls/Sec)	Applicable Call Rate (Calls/Sec)
SIP Inbound Call	RX = 8.9 * EffectiveCallRate	TX = 69.4 * EffectiveCallRate	0–90
Transfer	RX = 13.1 * EffectiveCallRate	TX = 86.4 * EffectiveCallRate	0–40
Recording	RX = 14.9 * EffectiveCallRate	TX = 73.9 * EffectiveCallRate	0–30
Conference	RX = 23.7 * EffectiveCallRate	TX = 116.7 * EffectiveCallRate	0–25
Consultation Call	RX = 19.0 * EffectiveCallRate	TX = 107.9 * EffectiveCallRate	0–35
Internal Call (using 3pcc)	RX = 11.5 * EffectiveCallRate	TX = 58.9 * EffectiveCallRate	0–50

Note: The typical margin of error is 5 percent.

The resulting network traffic is the sum of individual components (calculated per scenario):

$$TotalRX = \Sigma RX_i$$

$$TotalTX = \Sigma TX_i$$

Examples

This section provides several examples to show how the described approach works in practice.

Simple Example

This example illustrates how to use the proposed methods to estimate CPU usage and network traffic when no Network SIP Server is deployed. Typically, this means that the aggregated call rate and the number of concurrent sessions are relatively low.

For the purpose of this example, the system specifications are as follows:

- The system supports three basic scenarios, which are reasonably close to the following reference call flows:
  - SIP Inbound Call
  - Transfer
  - Recording

- The common call rates and average call durations are:
  - For the SIP Inbound Call scenario: 2 calls/sec, 180 sec
  - For the Transfer scenario: 1 call/sec, 240 sec
  - For the Recording scenario: 2 calls/sec, 300 sec
- The CPU type is Quad Core Intel Xeon E5405, 2x6 MB cache, 2.0 GHz, 1333 MHz FSB.
- The contact center has 1,500 agents, and each agent application has an individual T-Library connection to a SIP Server.

CPU Usage Estimation

Using the procedure "Estimating CPU Usage" on page 271, the calculations are as follows:

Step 1: Express call flows in terms of reference scenarios.

This step is not required, because the input data provides this information.

Step 2: Calculate the number of concurrent calls per scenario:

a. Per SIP Inbound Call scenario:

$ConcurrentCalls_1 = 2.0 \text{ calls/sec} * 180 \text{ sec} = 360 \text{ calls}$

b. Per Transfer scenario:

$ConcurrentCalls_2 = 1.0 \text{ call/sec} * 240 \text{ sec} = 240 \text{ calls}$

c. Per Recording scenario:

$ConcurrentCalls_3 = 2.0 \text{ calls/sec} * 300 \text{ sec} = 600 \text{ calls}$

Step 3: Calculate the maximum number of concurrent calls in the system:

$TotalConcurrentCalls = 360 + 240 + 600 = 1,200$

In this step, we simply check that the number of concurrent calls is below 2,000, which means that one instance of SIP Server is sufficient (and no Network SIP Server is required).

Step 4: Calculate the effective call rate per scenario:

Because we are dealing with a single SIP Server instance, the effective call rates are the same as the original call rates:

a. Per SIP Inbound Call scenario:

$EffectiveCallRate_1 = 2.0 \text{ calls/sec}$

b. Per Transfer scenario:

$EffectiveCallRate_2 = 1.0 \text{ call/sec}$

c. Per Recording scenario:

$EffectiveCallRate_3 = 2.0 \text{ calls/sec}$

Step 5: Estimate CPU usage per scenario:

Using the calibration table for CPU usage (Table 128 on page 272), we get the following estimates for CPU usage per scenario:

a. Per SIP Inbound Call scenario:

$CPU_1 = 0.69 \ (\% * \text{sec/call}) * 2.0 \text{ calls/sec} = 1.38\%$

b. Per Transfer scenario:

`CPU`$_2$ `= 1.42 (% * sec/call) * 1.0 call/sec = 1.42%`

c. Per Recording scenario:

`CPU`$_3$ `= 2.25 (% * sec/call)* 2.0 calls/sec = 5.5%`

So, `TotalCPU = 1.38 + 1.42 + 5.5 = 8.3%`

Because we have multiple client connections to SIP Server (a maximum of 1,500), we must factor the CPU usage according to Table 129 on page 274:

```
CorrectedTotalCPU = 1.10 * 8.3 = 9.13%
```

However, the calibration table for CPU usage assumes an Intel CPU running at 2.33 GHz, and we are using a 2.0 GHz CPU. Therefore, we must scale the result for the lower clock frequency:

```
CPU = 9.13 * (2.33/2.0) = 10.6%
```

Network Traffic Estimation

To estimate the amount of network traffic, use the calibration table for network traffic (Table 130 on page 275), summing the values for individual scenarios to get the following results (for SIP traffic):

- For received `(RX)` traffic:

```
RX = 8.9 * 2.0 + 13.1 * 1.0 + 14.9 * 2.0 = 60.7 (KB/sec)
```

- For transmitted `(TX)` traffic:

```
TX = 69.4 * 2.0 + 86.4 * 1.0 + 73.9 * 2.0 = 373.0 (KB/sec)
```

Complex Example

This example shows how to use the proposed procedure to estimate CPU usage and network traffic when a Network SIP Server is deployed. This means that the aggregated call rate and the number of concurrent sessions are relatively high.

For the purpose of this example, the system specifications are as follows:

- The system supports three basic scenarios, which are reasonably close to the following reference call flows:
  - SIP Inbound Call
  - Transfer
  - Consultation Call
- The common call rates and average call durations are:
  - For the SIP Inbound Call scenario: 20 calls/sec, 180 sec
  - For the Transfer scenario: 10 calls/sec, 240 sec
  - For the Consultation Call scenario: 2 calls/sec, 300 sec
- The CPU type is Quad Core Intel Xeon E5430, 2x6 MB cache, 2.66 GHz, 1333 MHz FSB.

- The contact center has 8,000 agents, with individual T-Library client connections aggregated by means of Genesys Agent Desktop.

CPU Usage Estimation

Using the procedure “Estimating CPU Usage” on page 271, the calculations are as follows:

Step 1: Express call flows in terms of reference scenarios.

This step is not required, because the input data provides this information.

Step 2: Calculate the number of concurrent calls per scenario:

a. Per SIP Inbound Call scenario:

$\texttt{ConcurrentCalls}_1$ `= 20 calls/sec * 180 sec = 3,600 calls`

b. Per Transfer scenario:

$\texttt{ConcurrentCalls}_2$ `= 10 calls/sec * 240 sec = 2,400 calls`

c. Per Consultation Call scenario:

$\texttt{ConcurrentCalls}_3$ `= 2.0 calls/sec * 300 sec = 600 calls`

Step 3: Calculate the maximum number of concurrent calls in the system:

`TotalConcurrentCalls = 3,600 + 2,400 + 600 = 6,600 calls`

To keep the number of sessions for a single instance of SIP Server, we would need four SIP Servers and a Network SIP Server:

`nSIPServers = Max Integer (6,600 / 2,000) = 4`

The average number of concurrent sessions per instance is:

`(6,600 / 4) = 1,650.`

Step 4: Calculate the effective call rate per scenario:

Because we are using multiple SIP Servers, we must calculate the effective call rate per instance (assuming even call distribution across instances):

a. Per SIP Inbound Call scenario:

$\texttt{EffectiveCallRate}_1$ `= 20.0/4 = 5.0 calls/sec`

b. Per Transfer scenario:

$\texttt{EffectiveCallRate}_2$ `= 10.0/4 = 2.5 calls/sec`

c. Per Consultation Call scenario:

$\texttt{EffectiveCallRate}_3$ `= 2.0/4 = 0.5 call/sec`

Step 5: Estimate CPU usage per scenario:

Using the calibration table for CPU usage (Table 128 on page 272), we get the following estimates for CPU usage per scenario:

a. Per SIP Inbound Call scenario:

$\texttt{CPU}_1$ `= 0.69 (% * sec/call) * 5.0 calls/sec = 3.45%`

b. Per Transfer scenario:

$\texttt{CPU}_2$ `= 1.42 (% * sec/call) * 2.5 calls/sec = 3.55%`

c. Per Consultation Call scenario:

CPU_3 = 2.33 (% * sec/call)* 0.5 call/sec = 1.17%

So, TotalCPU = 3.45 + 3.55 + 1.17 = 8.17%

Because we are using a Network SIP Server for call distribution, we must factor the CPU usage according to Table 129 on page 274:

CorrectedTotalCPU = 1.20 * 8.17 = 9.8%

However, the calibration table for CPU usage assumes an Intel CPU running at 2.33 GHz, and we are using a 2.66 GHz CPU. Therefore, we must scale the result for the higher clock frequency:

CPU = 9.8 * (2.33/2.66) = 8.6%

Network Traffic Estimation

To estimate the amount of network traffic (per instance), use the calibration table for network traffic (Table 130 on page 275), summing the values for individual scenarios to get the following results (for SIP traffic per SIP Server instance):

- For received (RX) traffic:

  RX = 8.9 * 20.0 + 13.1 * 10.0 + 19.0 * 0.5 = 318.5 (KB/sec)

- For transmitted (TX) traffic:

  TX = 69.4 * 0.0 + 86.4 * 10.0 + 107.0 * 0.5 = 2,305 (KB/sec)

Note: The amount of traffic via host network interfaces depends on the particular placement of SIP Server instances.

Reference Call Flows

This section provides reference call flows and details of SIP messages that were used for benchmarking of the following call scenarios:

- "SIP Inbound Call Scenario" on page 280
- "Transfer Scenario" on page 281
- "Recording Scenario" on page 283
- "Conference Call Scenario" on page 285
- "Consultation Call Scenario" on page 287
- "Internal Call Scenario" on page 289

SIP Inbound Call Scenario

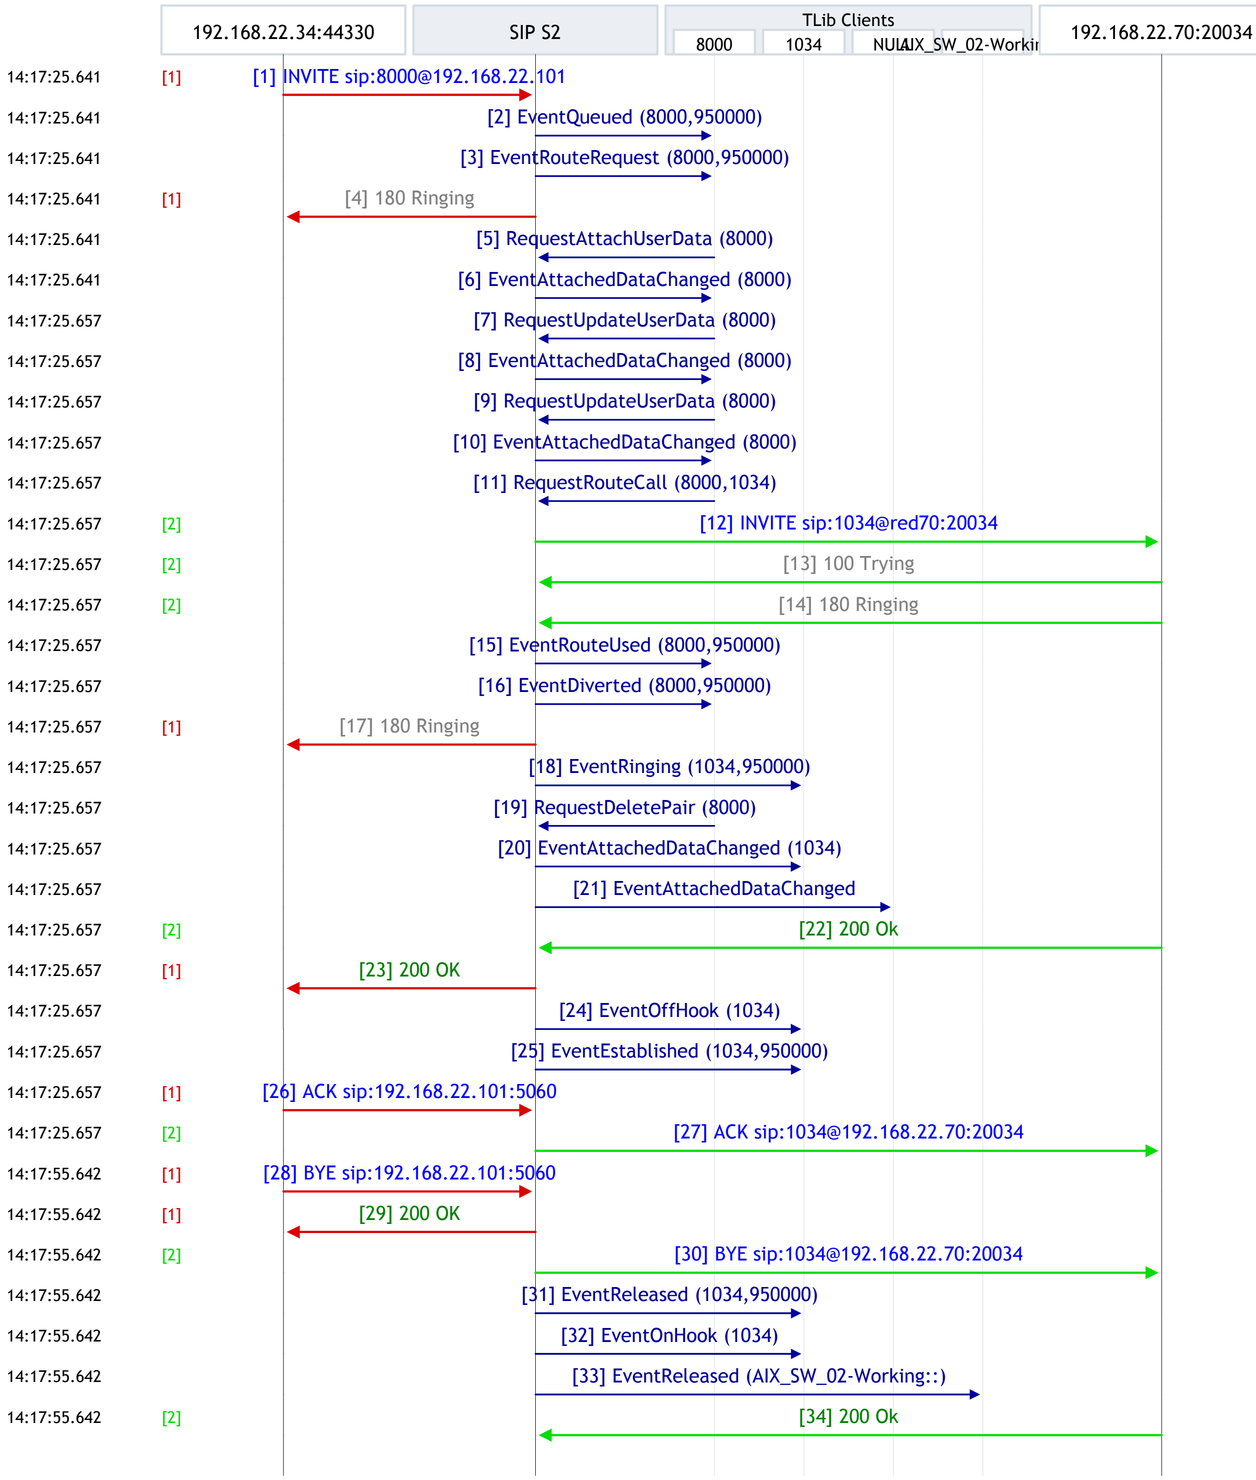

Transfer Scenario

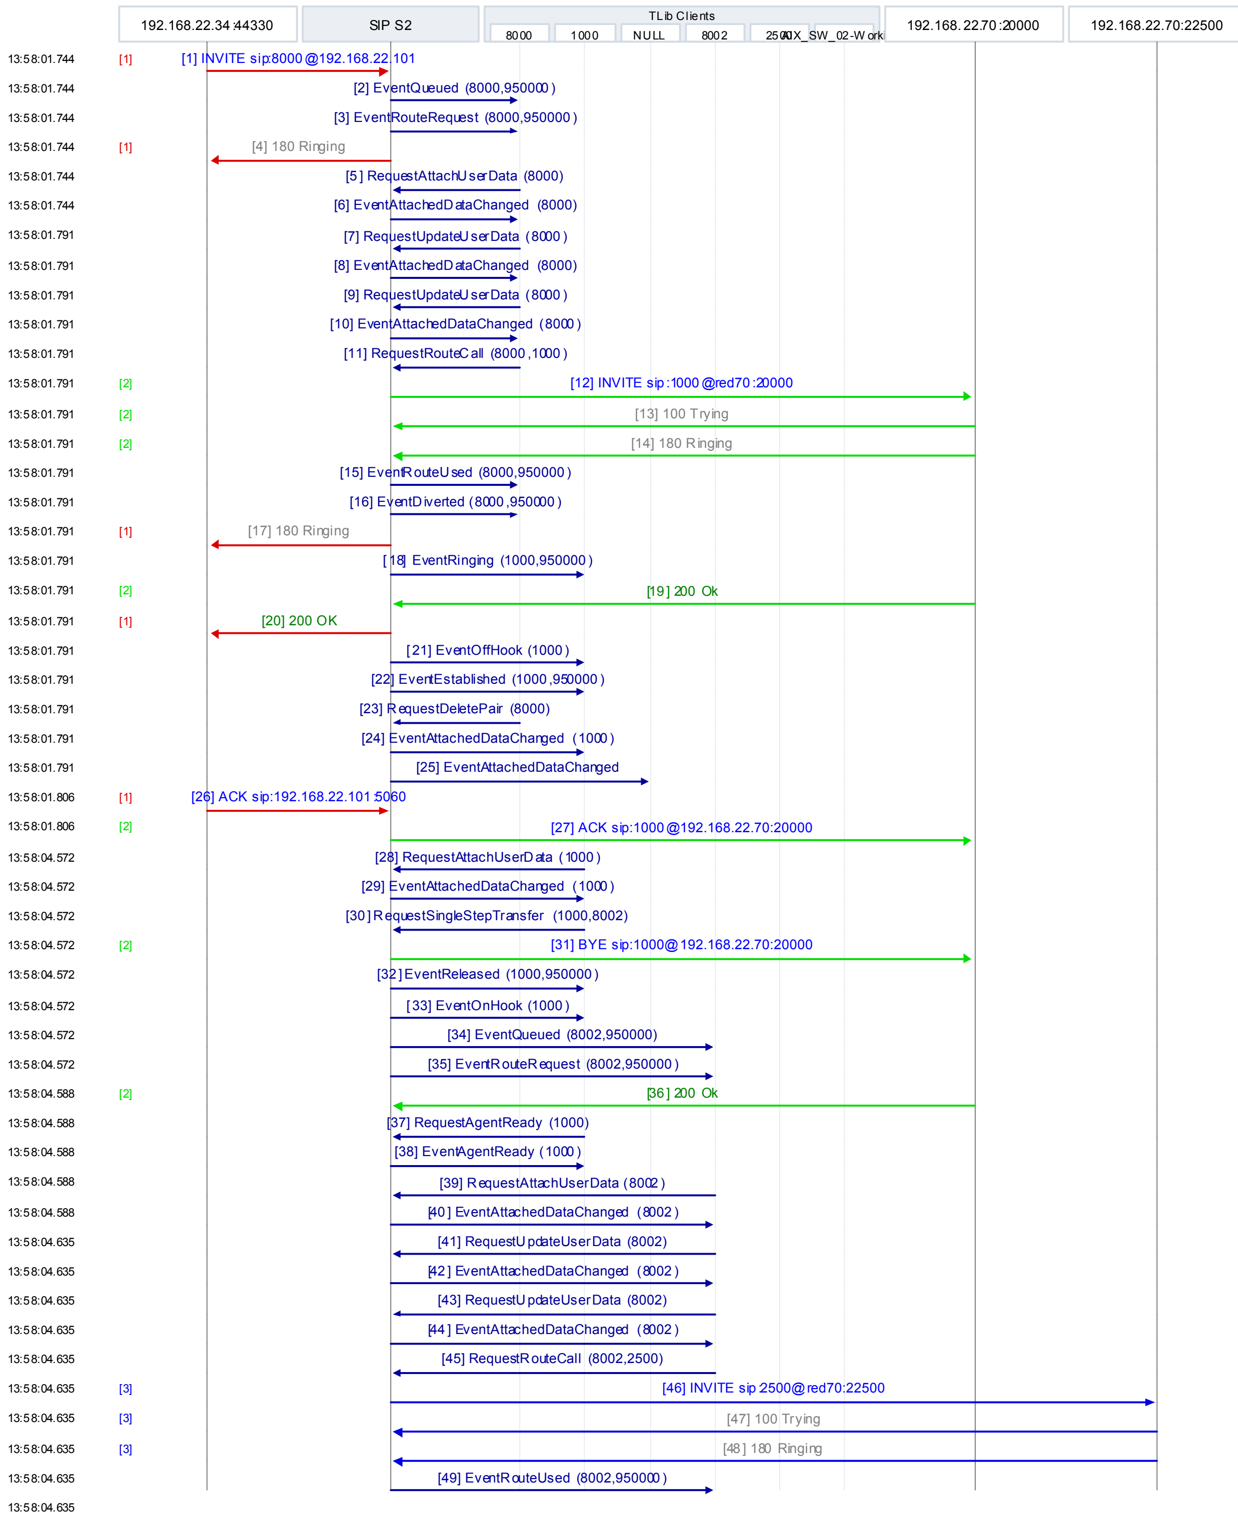

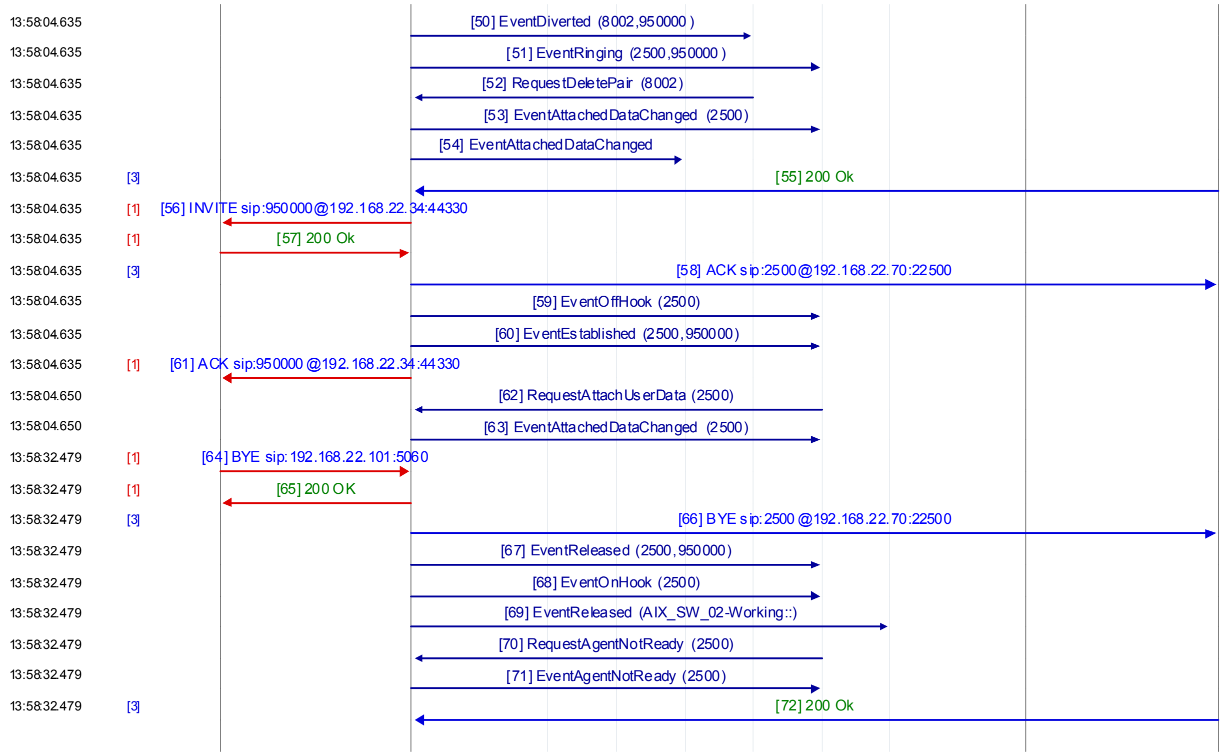
13:58:04.635 [50] EventDiverted (8002,950000)
13:58:04.635 [51] EventRinging (2500,950000)
13:58:04.635 [52] RequestDeletePair (8002)
13:58:04.635 [53] EventAttachedDataChanged (2500)
13:58:04.635 [54] EventAttachedDataChanged
13:58:04.635 [3] [55] 200 Ok
13:58:04.635 [1] [56] INVITE sip:950000@192.168.22.34:44330
13:58:04.635 [1] [57] 200 Ok
13:58:04.635 [3] [58] ACK sip:2500@192.168.22.70:22500
13:58:04.635 [59] EventOffHook (2500)
13:58:04.635 [60] EventEstablished (2500,950000)
13:58:04.635 [1] [61] ACK sip:950000@192.168.22.34:44330
13:58:04.650 [62] RequestAttachUserData (2500)
13:58:04.650 [63] EventAttachedDataChanged (2500)
13:58:32.479 [1] [64] BYE sip:192.168.22.101:5060
13:58:32.479 [1] [65] 200 OK
13:58:32.479 [3] [66] BYE sip:2500@192.168.22.70:22500
13:58:32.479 [67] EventReleased (2500,950000)
13:58:32.479 [68] EventOnHook (2500)
13:58:32.479 [69] EventReleased (AIX_SW_02-Working::)
13:58:32.479 [70] RequestAgentNotReady (2500)
13:58:32.479 [71] EventAgentNotReady (2500)
13:58:32.479 [3] [72] 200 Ok

Recording Scenario

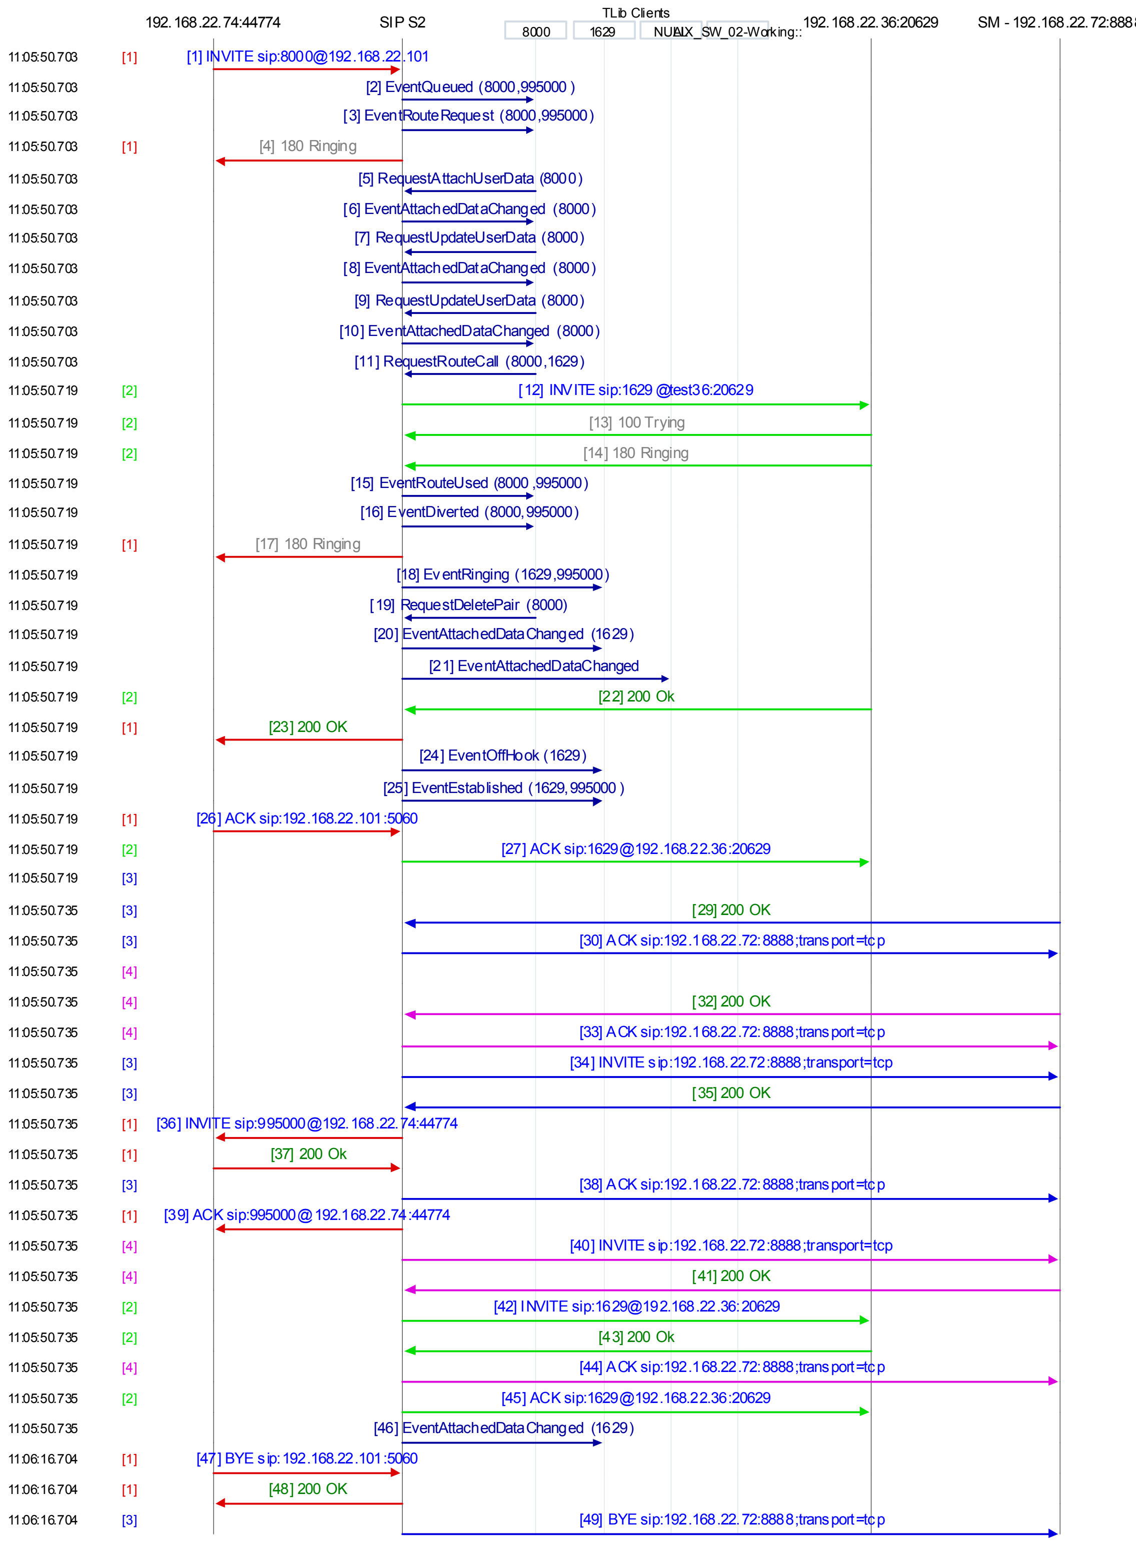

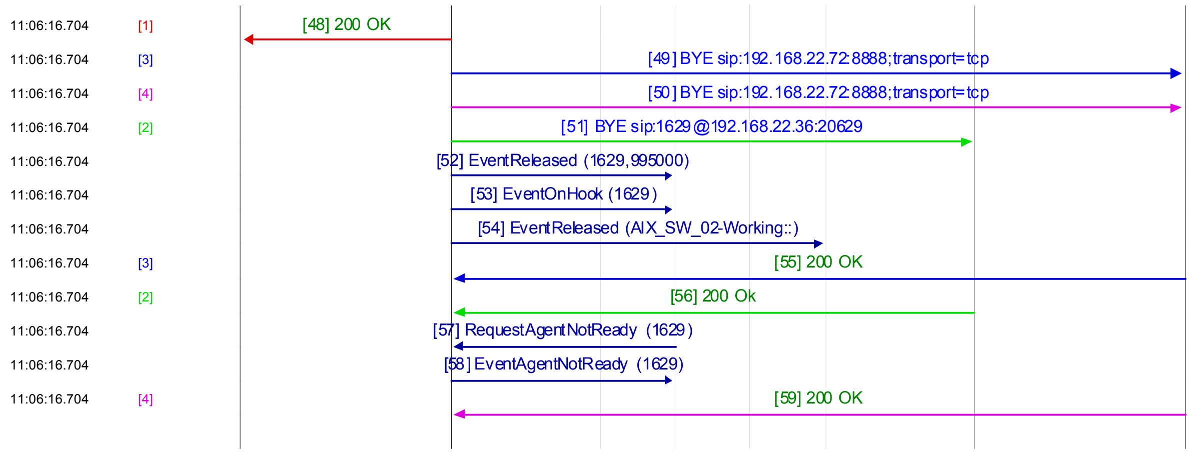
11:06:16.704 [1] [48] 200 OK
11:06:16.704 [3] [49] BYE sip:192.168.22.72:8888;transport=tcp
11:06:16.704 [4] [50] BYE sip:192.168.22.72:8888;transport=tcp
11:06:16.704 [2] [51] BYE sip:1629@192.168.22.36:20629
11:06:16.704 [52] EventReleased (1629,995000)
11:06:16.704 [53] EventOnHook (1629)
11:06:16.704 [54] EventReleased (AIX_SW_02-Working::)
11:06:16.704 [3] [55] 200 OK
11:06:16.704 [2] [56] 200 Ok
11:06:16.704 [57] RequestAgentNotReady (1629)
11:06:16.704 [58] EventAgentNotReady (1629)
11:06:16.704 [4] [59] 200 OK

Conference Call Scenario

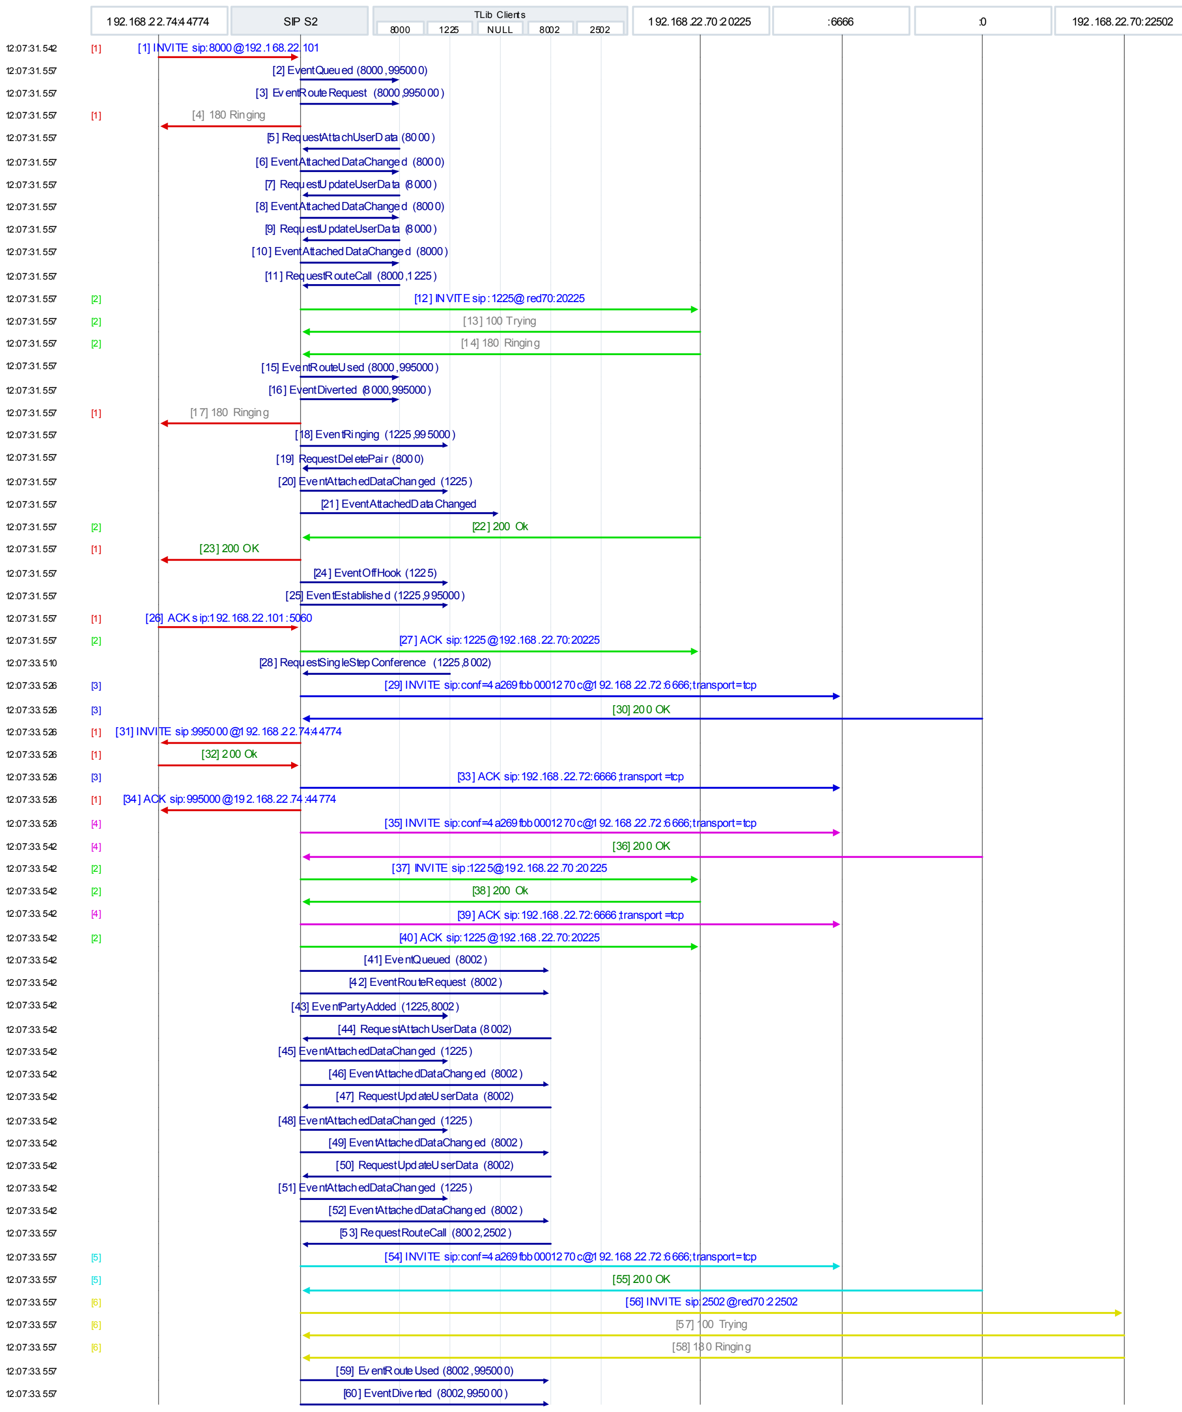

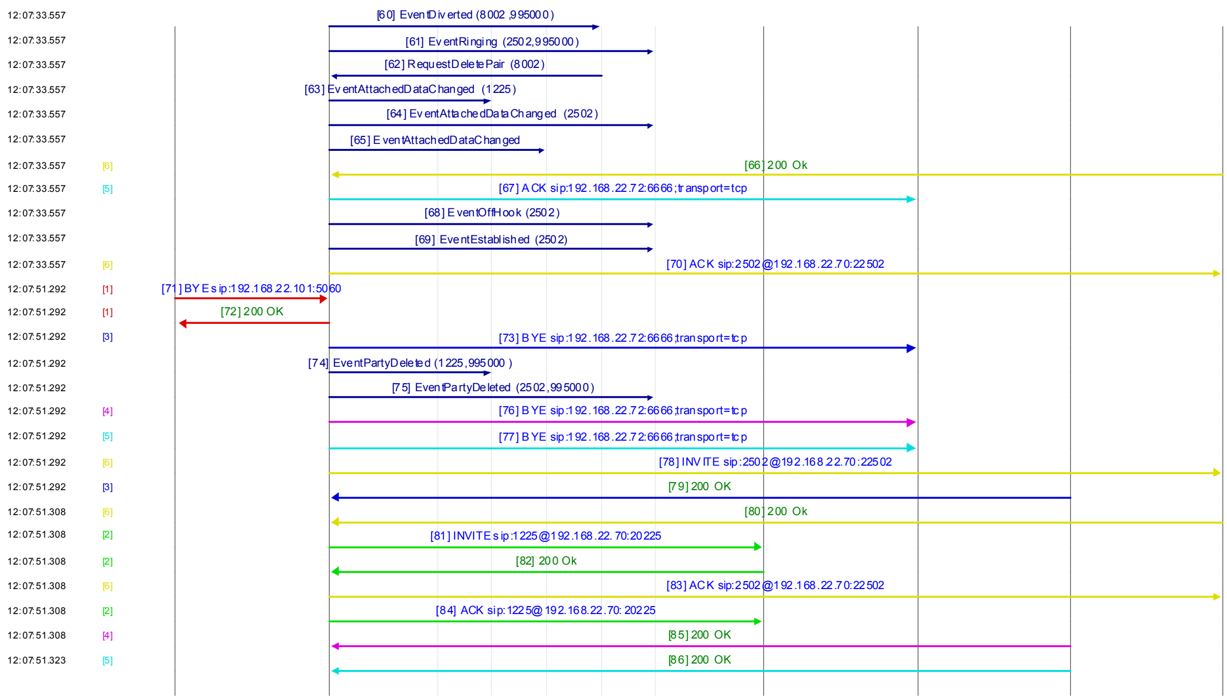
12:07:33.557 [60] EventDiverted (8002,995000)
12:07:33.557 [61] EventRinging (2502,995000)
12:07:33.557 [62] RequestDeletePair (8002)
12:07:33.557 [63] EventAttachedDataChanged (1225)
12:07:33.557 [64] EventAttachedDataChanged (2502)
12:07:33.557 [65] EventAttachedDataChanged
12:07:33.557 [6] [66] 200 Ok
12:07:33.557 [5] [67] ACK sip:192.168.22.72:6666;transport=tcp
12:07:33.557 [68] EventOffHook (2502)
12:07:33.557 [69] EventEstablished (2502)
12:07:33.557 [6] [70] ACK sip:2502@192.168.22.70:22502
12:07:51.292 [1] [71] BYE sip:192.168.22.101:5060
12:07:51.292 [1] [72] 200 OK
12:07:51.292 [3] [73] BYE sip:192.168.22.72:6666;transport=tcp
12:07:51.292 [74] EventPartyDeleted (1225,995000)
12:07:51.292 [75] EventPartyDeleted (2502,995000)
12:07:51.292 [4] [76] BYE sip:192.168.22.72:6666;transport=tcp
12:07:51.292 [5] [77] BYE sip:192.168.22.72:6666;transport=tcp
12:07:51.292 [6] [78] INVITE sip:2502@192.168.22.70:22502
12:07:51.292 [3] [79] 200 OK
12:07:51.308 [6] [80] 200 Ok
12:07:51.308 [2] [81] INVITE sip:1225@192.168.22.70:20225
12:07:51.308 [2] [82] 200 Ok
12:07:51.308 [6] [83] ACK sip:2502@192.168.22.70:22502
12:07:51.308 [2] [84] ACK sip:1225@192.168.22.70:20225
12:07:51.308 [4] [85] 200 OK
12:07:51.323 [5] [86] 200 OK

Consultation Call Scenario

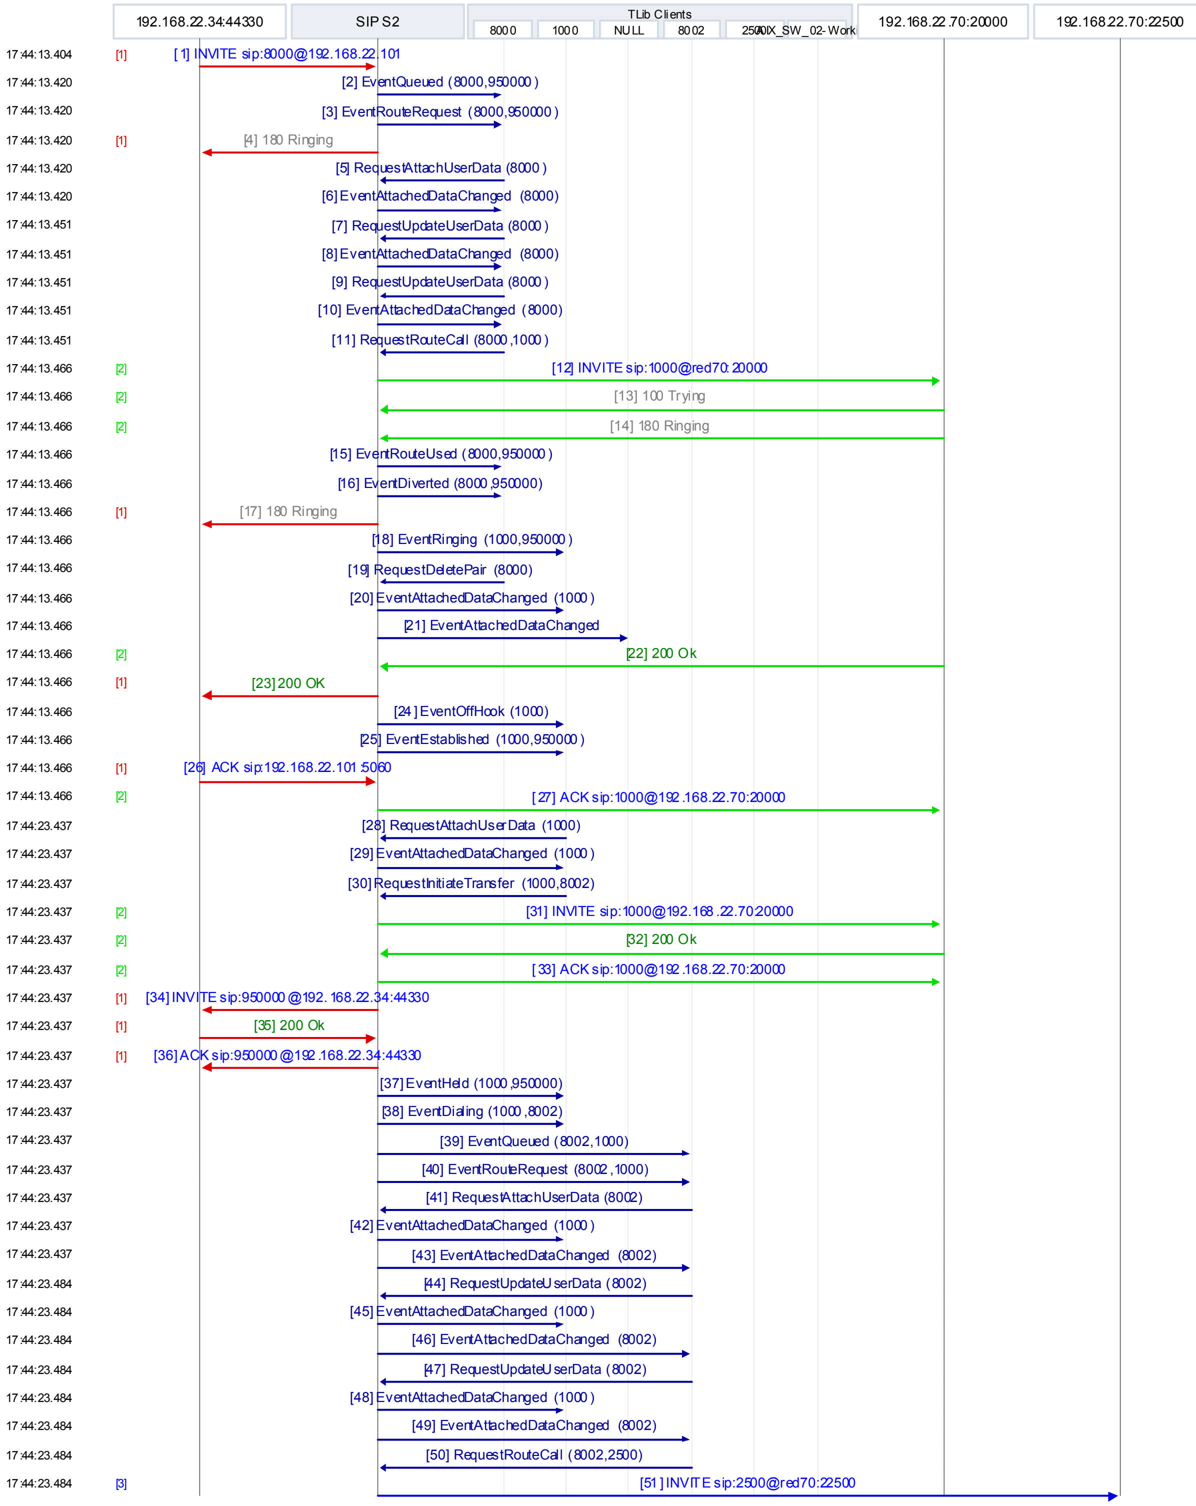

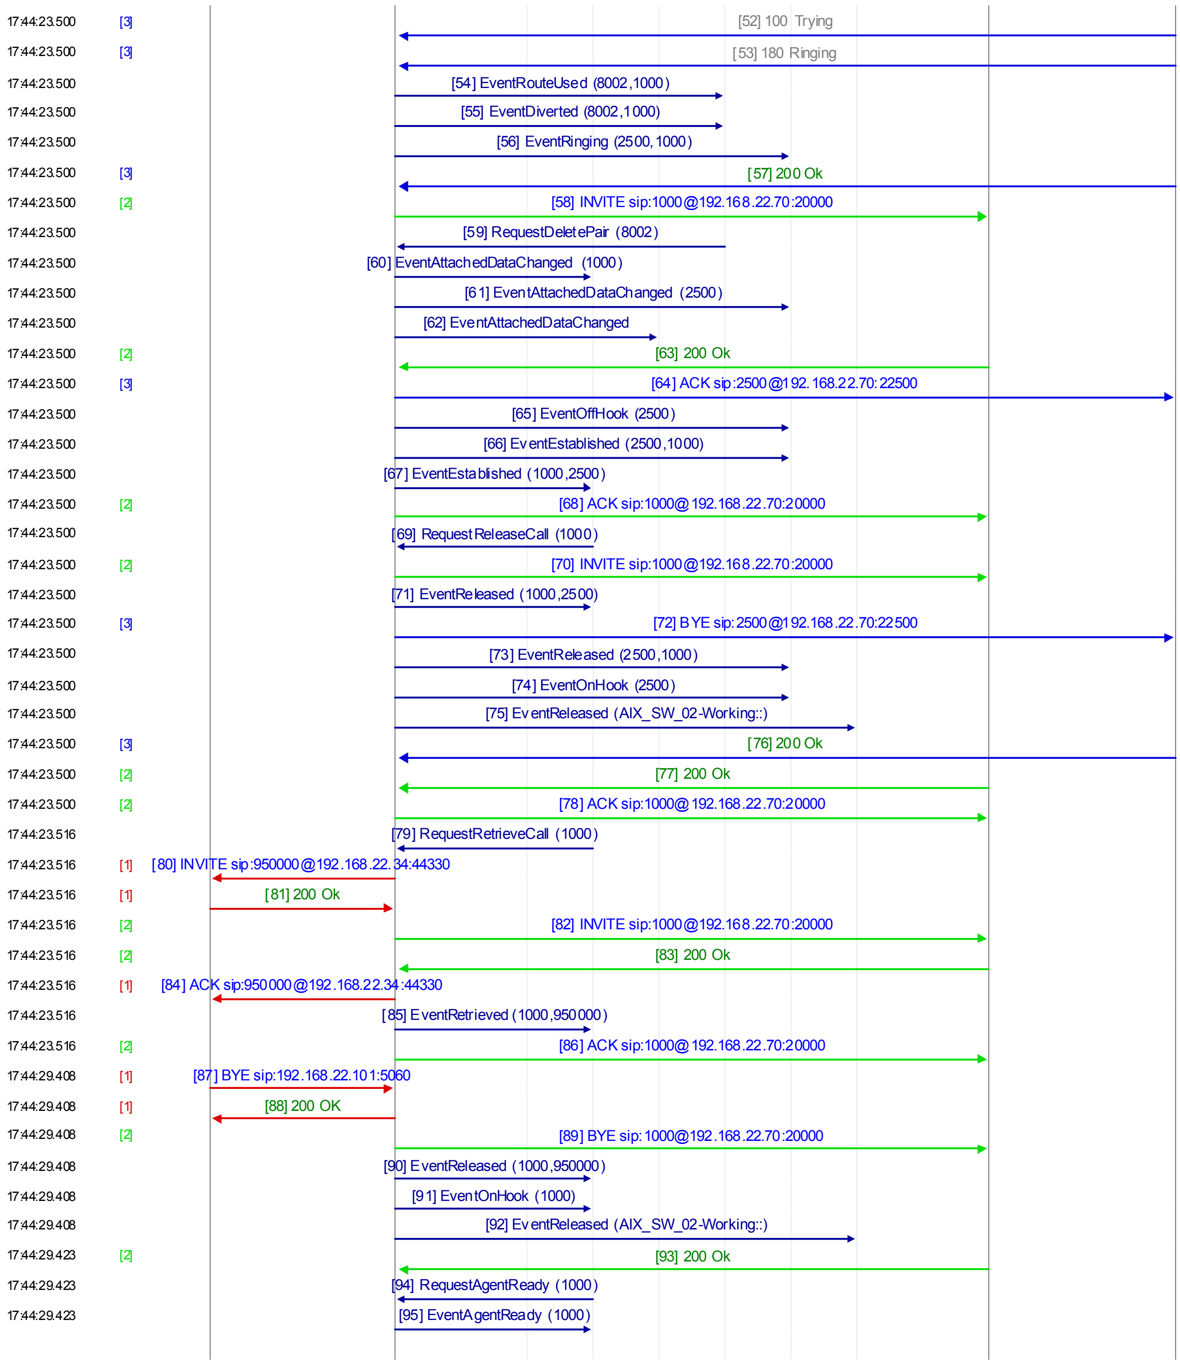
17:44:23.500 [3] [52] 100 Trying
17:44:23.500 [3] [53] 180 Ringing
17:44:23.500 [54] EventRouteUsed (8002,1000)
17:44:23.500 [55] EventDiverted (8002,1000)
17:44:23.500 [56] EventRinging (2500,1000)
17:44:23.500 [3] [57] 200 Ok
17:44:23.500 [2] [58] INVITE sip:1000@192.168.22.70:20000
17:44:23.500 [59] RequestDeletePair (8002)
17:44:23.500 [60] EventAttachedDataChanged (1000)
17:44:23.500 [61] EventAttachedDataChanged (2500)
17:44:23.500 [62] EventAttachedDataChanged
17:44:23.500 [2] [63] 200 Ok
17:44:23.500 [3] [64] ACK sip:2500@192.168.22.70:22500
17:44:23.500 [65] EventOffHook (2500)
17:44:23.500 [66] EventEstablished (2500,1000)
17:44:23.500 [67] EventEstablished (1000,2500)
17:44:23.500 [2] [68] ACK sip:1000@192.168.22.70:20000
17:44:23.500 [69] RequestReleaseCall (1000)
17:44:23.500 [2] [70] INVITE sip:1000@192.168.22.70:20000
17:44:23.500 [71] EventReleased (1000,2500)
17:44:23.500 [3] [72] BYE sip:2500@192.168.22.70:22500
17:44:23.500 [73] EventReleased (2500,1000)
17:44:23.500 [74] EventOnHook (2500)
17:44:23.500 [75] EventReleased (AIX_SW_02-Working::)
17:44:23.500 [3] [76] 200 Ok
17:44:23.500 [2] [77] 200 Ok
17:44:23.500 [2] [78] ACK sip:1000@192.168.22.70:20000
17:44:23.516 [79] RequestRetrieveCall (1000)
17:44:23.516 [1] [80] INVITE sip:950000@192.168.22.34:44330
17:44:23.516 [1] [81] 200 Ok
17:44:23.516 [2] [82] INVITE sip:1000@192.168.22.70:20000
17:44:23.516 [2] [83] 200 Ok
17:44:23.516 [1] [84] ACK sip:950000@192.168.22.34:44330
17:44:23.516 [85] EventRetrieved (1000,950000)
17:44:23.516 [2] [86] ACK sip:1000@192.168.22.70:20000
17:44:29.408 [1] [87] BYE sip:192.168.22.101:5060
17:44:29.408 [1] [88] 200 OK
17:44:29.408 [2] [89] BYE sip:1000@192.168.22.70:20000
17:44:29.408 [90] EventReleased (1000,950000)
17:44:29.408 [91] EventOnHook (1000)
17:44:29.408 [92] EventReleased (AIX_SW_02-Working::)
17:44:29.423 [2] [93] 200 Ok
17:44:29.423 [94] RequestAgentReady (1000)
17:44:29.423 [95] EventAgentReady (1000)

Internal Call Scenario

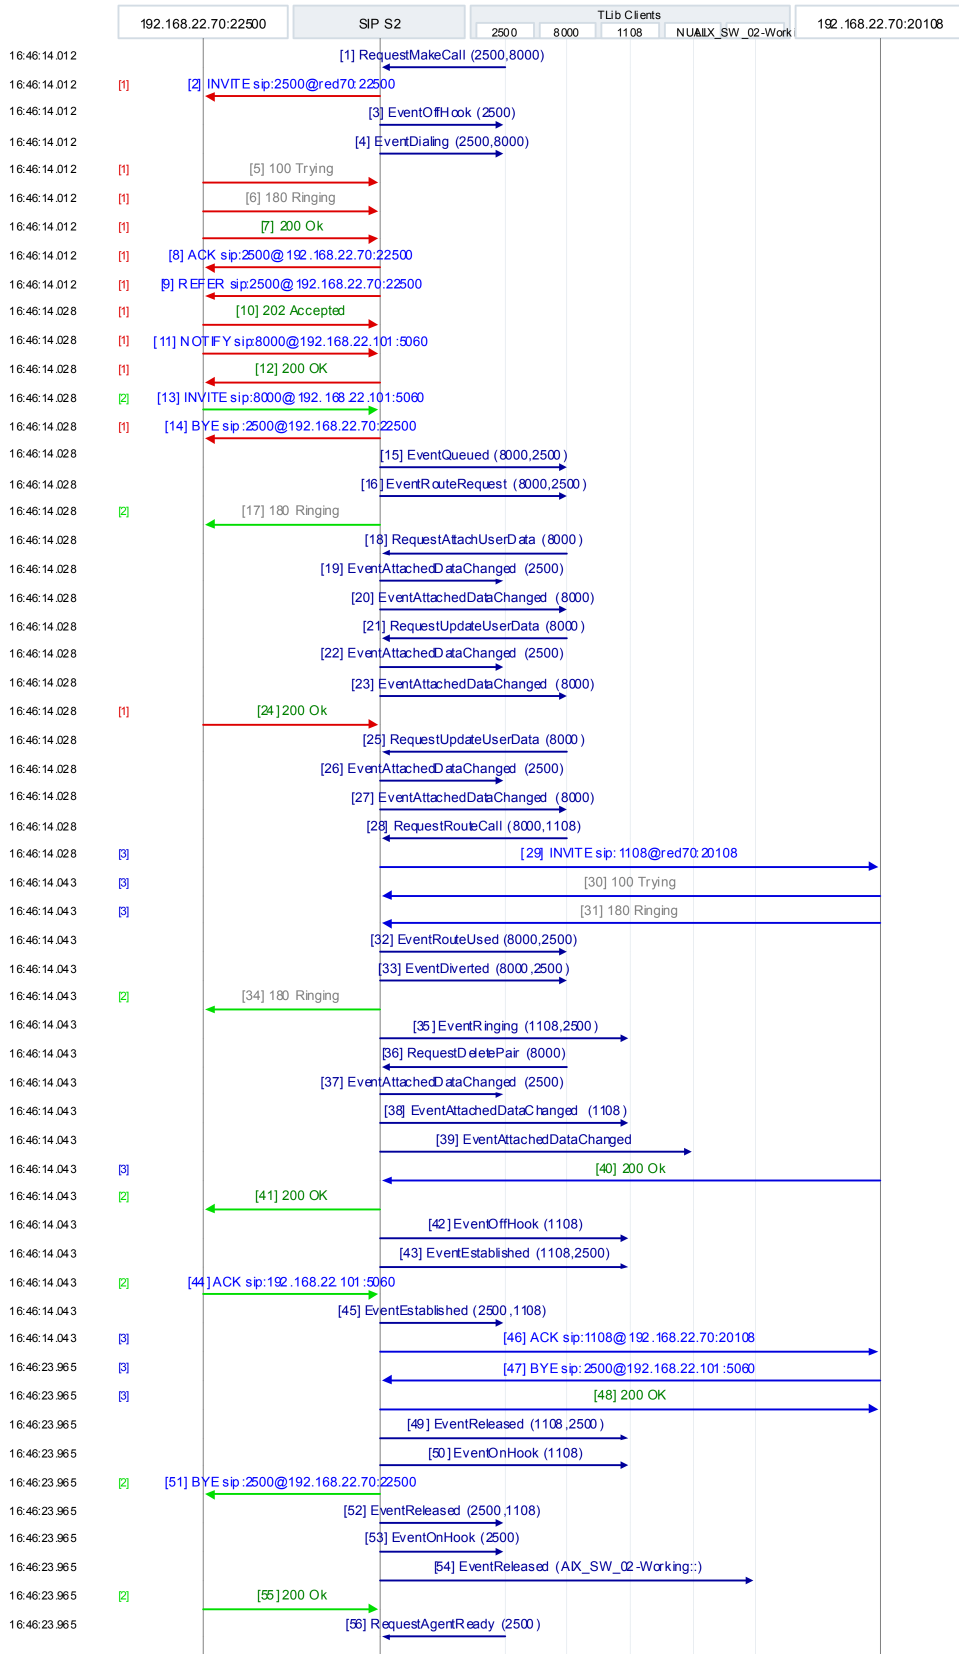

Hardware Details

Table 131 provides details of the hardware platform that was used for benchmarks on Windows OS.

Table 131: Hardware Details

Feature	Details	Notes
Operating System	Microsoft Windows Server 2003 Enterprise Edition, 32-bit SP2	Default socket buffer size adjusted to 64 KB
Hardware Platform	Dell 2950	
CPU	Dual Quad Core Intel Xeon L5410	Total number of cores: 8 L2 cache size: 2x6 MB
RAM	8 GB	2x4 GB, 667 MHz (only 4 GB of RAM is used by OS)
HDD	Two (2) 73 GB 15K RPM Serial-Attach SCSI 3 Gbps 3.5-in	Controller Type: Integrated SAS/SATA, PERC 6/I, no RAID configuration
NIC	Two (2) Broadcom BCM5708C NetXtreme II GigE	1 GB Full Duplex, TCP Offload Engine (TOE) disabled

Default Socket Buffer Size

For optimal SIP Server and Stream Manager performance, the default socket buffer size (for both TCP and UDP sockets) must be increased to at least 64 KB. The default buffer size is OS-specific, and is typically less than 64 KB. For example, on Windows Server 2003 with SP2, the default settings could be 8,096 bytes.

The methods for adjusting and checking the default buffer size are OS-specific. For more information, see the latest version of the *SIP Server Deployment Guide*.

Chapter

12 Agent Interaction Layer

This chapter covers the actual performance test results done on Agent Interaction Layer (AIL) 7.5. It presents the following topic:

AIL Performance Test Statistics

Genesys conducted tests to measure the performance of AIL when connected to the Genesys Framework. These tests used voice-only, e-mail-only, and blended scenarios. All tests used AIL 7.5 (which is built with Platform SDK libraries).

Test Environment

Table 132 details the test environment used to arrive at the recommended limits for deploying the AIL-based applications.

Note: One basis for understanding the optimized results of all the tests in this section is the use of a dedicated 1-Gigabit network in the lab environment.

Table 132: AIL Test Environment Components

Component	Host	Hardware	Operating System
Configuration Server 7.5 T-Server G3 7.5 Switch Simulator 6.5.300.02	Host_01	2x Xeon 3.4 GHz/cache 512 KB; 2 GB memory; dedicated 1-Gigabit Ethernet	Microsoft Windows 2003 Server

Table 132: AIL Test Environment Components (Continued)

Component	Host	Hardware	Operating System
AIL 7.5	Host_04	2x Opteron Dual Core; 2.2 GHz/cache 512 KB; 4 GB memory; dedicated 1-Gigabit Ethernet	Microsoft Windows 2003 Server; (J2SDK), Standard Ed., 1.4.2_08
Universal Contact Server (UCS) 7.5 Interaction Server 7.5	Host_02	2x Xeon 3.4 GHz/cache 512 KB; 2 GB memory; dedicated 1-Gigabit Ethernet	Microsoft Windows 2003 Server
Microsoft SQL Server 2000 (UCS Database)	Host_03	Intel Xeon 3.0 GHz /cache 512 KB; 2 GB memory; dedicated 1-Gigabit Ethernet	Microsoft Windows 2003 Server

Test Tools

The tests used the following steps according to contact center role.

Behavior of Voice Agents

Upon receiving a call, an agent:

1. Answers the call.
2. Waits for 10 seconds.
3. Sets the agent's status to `NotReady`.
4. Releases the call.
5. Waits for 2 seconds.
6. Becomes `Ready`.
7. Repeats these steps.

Behavior of E-Mail Agents

Upon receiving an e-mail, an agent:

1. Answers the incoming e-mail.
2. Replies with an outgoing e-mail.
3. Waits for 10 seconds.
4. Sets the agent's status to `NotReady`.
5. Sends the outgoing e-mail.
6. Waits for 2 seconds.

7. Becomes `Ready`.
8. Repeats these steps.

Test Procedure

The general scenario for the tests proceeded as follows:

1. Launched the AIL-based application.
2. Sequentially logged all agents (for both voice and e-mail).
3. Initiated voice traffic and/or sent e-mails.
4. Measured the call/interaction rates.

Voice Traffic

The tests used the Genesys 6.5 Switch Simulator and its accompanying test tools to generate voice traffic. The call rate was set to a level above the theoretical maximum so as to minimize delays between any two calls. (Calls that could not be distributed simply were not.)

E-Mail Traffic

The tests used a customized e-mail spammer to send e-mails. 10 times more e-mails were sent than there were agents in the test scenarios.

Note: The tests looked at the performance of AIL, but also at various Genesys Multimedia components, including E-mail Server Java, Interaction Server, and Universal Contact Server. The tests used a corporate e-mail server.

In some tests, it was clear that E-mail Server Java did not process incoming e-mail quickly enough to keep all agents busy. Those tests were re-done. (E-mail Server Java may have been impacted by the processing of outgoing e-mails.

Test Results

Test results here are grouped according to different criteria.

Startup and Login

It takes AIL about 8.8 seconds to start (which includes optimization that the 7.5 release performs at that time).

Voice-Only Test Results

Table 133 displays the results of voice-only testing. This test did not save voice calls to the UCS database. The theoritical maximum call-rate values depend on the number of agents logged in and the time spent for each call (plus its wrap-up time). The difference between this theoretical maximum and the actual number represents the cost of processing by AIL.

Table 133: Results of Voice-Only Test

Number of Agents	AIL Release Number	Call Rate Achieved (calls/second)	Theoretical Maximum Call Rate (calls/second)
100	7.5.018.00	7.7	8.3
200		15.3	16.6
500		38.4	41.6

Voice-Only (but with writing to Database) Test Results

Table 134 shows the results of a test that did save the voice calls to the UCS database.

Table 134: Results of Voice with Database Test

Number of Agents	AIL Release Number	Call Rate Achieved (calls/second)	Theoretical Maximum Call Rate (calls/second)
100	7.5.018.00	7.66	8.3
200		14.2	16.6
500		23.5	41.6

Saving voice calls to the UCS database significantly impacted the call rate.

E-Mail–Only Test Results

Table 135 shows the results of the e-mail–only test.

Table 135: Results of the E-Mail–Only Test

Number of Agents	AIL Release Number	E-Mail Rate Achieved (e-mails/second)	Theoretical Maximum E-Mail Rate (e-mails/second)
100	7.5.018.00	5.8	8.3
200		11.1	16.6
500		13.6	41.6

This test did not reveal full AIL performance results. A full test would have also shown the E-mail Server Java statistics to help understand the comparison to voice-only traffic.

Blended-Processing Test Results

Table 136 shows the results of the blended-processing test, in which phone calls were saved to the UCS database.

Table 136: Results of the Blended-Processing Test

Number of Agents	AIL Release	Call Rate Achieved (calls/second)	E-Mail Rate Achieved (e-mails/ second)	Theoretical Maximum Call/E-Mail Rate (calls or e-mails/second)
100	7.5.018.00	7.3	5.7	8.3
200		13.7	10.5	16.6
500		24.6	15.4	41.6

The blended-processing test showed that the rate for voice interactions was not significantly affected by the accompanying distribution of e-mails. Unexpectedly, more e-mails were distributed for the 500-agent test than for smaller contact-center scenarios.

Chapter

13 LivePerson Adapter

The LivePerson Adapter (the Adapter) has been developed for small-to-medium sized Contact Centers. This chapter presents the hardware requirements and scalability information. The sections include:

Small Contact Centers

This section provides hardware requirements for small Contact Centers (less than 500 agents) that run on the following operation systems:

- Windows (see Table 139)
- Linux (see Table 138)
- Solaris (see Table 139)

Windows

Table 137 provides the Adapter hardware requirements for a small Contact Center on a Microsoft Windows Server.

Table 137: Small Contact Center—Microsoft Windows Server

Microsoft Windows Server	
Processor Type, Quantity, Speed	2xCPU 2.33 GHz processors or better
Memory Size	4 GB RAM

Table 137: Small Contact Center—Microsoft Windows Server (Continued)

Microsoft Windows Server	
Hard Disk Space	20 GB available hard disk space Note: Hard disk requirements depend not only on the projected size of your database but also on the accumulated size of your log files (which you can regulate using options). (Contact Genesys Technical Support website at http://genesyslab.com/support for specific requirements since 20 GB is just a baseline.)
CD ROM	32 X

Linux

Table 138 provides the Adapter hardware requirements for a small Contact Center on Linux.

Table 138: Small Contact Center—Linux

Linux	
Processor Type, Quantity, Speed	2xCPU 2.33 GHz processors or better
Memory Size	4 GB RAM
Hard Disk Space	20 GB available hard disk space Note: Hard disk requirements depend not only on the projected size of your database but also on the accumulated size of your log files (which you can regulate via options). (Contact Genesys Technical Support website at http://genesyslab.com/support for specific requirements since 20 GB is just a baseline.)
CD ROM	32 X

Solaris

Table 139 provides the Adapter hardware requirements for a small Contact Center on Solaris.

Table 139: Small Contact Center—Solaris

Solaris	
Processor Type, Quantity, Speed	2 x Sun UltraSPARC-IIIi 1.0 GHz
Memory Size	4 GB RAM
Hard Disk Space	20 GB available hard disk space Note: Hard disk requirements depend not only on the projected size of your database but also on the accumulated size of your log files (which you can regulate using options).
CD ROM	32 X

Medium Contact Centers

This section provides hardware requirements for medium-sized Contact Centers (500 to 1000 agents) that run on the following operation systems:

- Windows (see Table 140)
- Linux (see Table 141)
- Solaris (see Table 142)

Windows

Table 140 provides the Adapter hardware requirements for a medium-sized Contact Center on a Microsoft Windows Server.

Table 140: Medium Contact Center—Microsoft Windows Server

Microsoft Windows Server	
Processor Type, Quantity, Speed	4xCPU 2.33 GHz processors or better
Memory Size	6 GB RAM

Table 140: Medium Contact Center—Microsoft Windows Server (Continued)

Microsoft Windows Server	
Hard Disk Space	20 GB available hard disk space Note: Hard disk requirements depend not only on the projected size of your database but also on the accumulated size of your log files (which you can regulate using options). (Contact Genesys Technical Support website at http://genesyslab.com/support for specific requirements since 20 GB is just a baseline.)
CD ROM	32 X

Linux

Table 141 provides the Adapter hardware requirements for a medium-sized Contact Center on Linux.

Table 141: Medium Contact Center—Linux

Linux	
Processor Type, Quantity, Speed	4xCPU 2.33 GHz processors or better
Memory Size	6 GB RAM
Hard Disk Space	20 GB available hard disk space Note: Hard disk requirements depend not only on the projected size of your database but also on the accumulated size of your log files (which you can regulate via options). (Contact Genesys Technical Support website at http://genesyslab.com/support for specific requirements since 20 GB is just a baseline.)
CD ROM	32 X

Solaris

Table 142 provides the Adapter hardware requirements for a medium-sized Contact Center on Solaris.

Table 142: Medium Contact Center—Solaris

Solaris	
Processor Type, Quantity, Speed	Quad dual core SPARC64 IV 2.15 GHz
Memory Size	6 GB RAM
Hard Disk Space	20 GB available hard disk space Note: Hard disk requirements depend not only on the projected size of your database but also on the accumulated size of your log files (which you can regulate using options).
CD ROM	32 X

Scalability (Windows, Linux, and Solaris)

The Adapter on Microsoft Windows and Unix-based platforms provides the following scalability for small-to-medium sized Contact Centers:

- 50,000 proactive interactions per day (all chat, all callback, or a combination of chat and callback)
- 100 simultaneous chat sessions (maximum)
- 125 simultaneous callback sessions (maximum)
- 250 agents monitored (maximum)
- 18 messages per second on LP link (peak rate)

Chapter

14 Generated System Traffic

This chapter presents information on the traffic generated by the system. It covers the following topic:

Traffic Generated by Components and Solutions

This section presents information on traffic generated by components and inbound solutions.

Traffic Generated by Framework 7.x Components

Table 143 provides basic data about link traffic among various Framework components. Use this information to help you determine the optimal component location on the network.

Table 143: Traffic Generated by Framework 7.x Components

Component 1	Component 2	Elements Determining Total Message Traffic	Primary Transaction Types	Average Message Length	Number of Messages per Transaction	Total Traffic Volume	Timelines s of Message Delivery
Config Server	DB Server ICON	frequency of configuration changes	config changes	1 KB	one per change	very low	not critical
Config Server	Any client application	frequency of configuration changes	regular config updates	1 KB	one per update	low	not critical
Any client application	Message Server[2]	type of client application (see the log events documentation)	log events	0.5 KB	one per event	low (standard log output level)	not critical
Message Server	DB Server	number of message server clients	log records	0.5 KB	one per event	medium	not critical
Message Server	Message Server - Solution Control Server	• number of configured alarm conditions • overall system stability	alarms	0.5 KB	two per alarm condition	low	not critical

Table 143: Traffic Generated by Framework 7.x Components (Continued)

Component 1	Component 2	Elements Determining Total Message Traffic	Primary Transaction Types	Average Message Length	Number of Messages per Transaction	Total Traffic Volume	Timeliness of Message Delivery
Solution Control Server	Solution Control Interface	• number of solutions • number of components per solution • number of computers serving the installation • number of configured alarm conditions • overall system stability	• solution control requests • changes in solution and application statuses • alarm information	0.5 KB	up to 20	low	not critical
DB Server for Log Database	Solution Control Interface	frequency of log view changes	log record blocks	50 KB	a block of 100 log records per single view change	medium	not critical
Solution Control Server	Local Control Agent	number of applications located on the same computer	• application startup • shutdown • redundancy arbitration commands	0.25 KB	one per control operation	very low	not critical

Traffic Generated by Inbound Solutions

Table 144 provides basic data about traffic among various inbound solutions with Reporting and IP Telephony Components included. The results shown in this table are based on Genesys performance test results.

Table 144: Traffic Generated by 7.5 Inbound Solution with Reporting Components

Component1	Component2	Protocol	Estimated bi-directional traffic	Criticality	Comments
T-Server	ICON	TLIB	9 Kbyte/call	Very high	Traffic was measured on generic average call flow.
T-Server	StatServer	TLIB	9 Kbyte/call	Very high	
T-Server	URS	TLIB	3.5 Kbyte/call	Very high	For one routing attempt.
URS	StatServer	StatLib	0.1 Kbyte/call	Very high	Could depend on strategy complexity.
StatServer	SS client	StatLib	0.1 Kbyte per statistic per update	High	
Interaction Concentrator (ICON) server	DBServer	DBLib	9 Kbyte/call	Very high	Traffic was measured for MS SQL on generic average call flow.
Genesys Info Mart	DataBase	JDBC	7 Kbyte/call	Very high	Aggregation and data transformation for supporting generic call flow.

Table 144: Traffic Generated by 7.5 Inbound Solution with Reporting Components (Continued)

Component1	Component2	Protocol	Estimated bi-directional traffic	Criticality	Comments
SIPServer	SIP Components	SIP	Simple incoming call: 18.5 Kbyte/call	Very high	Significantly depends on call flow and network conditions. If network connection is poor, messages could be re sent according SIP protocol.
CTI Link	T-Server	Link vendor appropriate	2 Kbyte /call	Critical	Given number is an estimation only and could vary for different switch vendors and configurations. Traffic was measured on generic average call flow.
Stream Manager	Stream component	RTP	• G711 codec: 22 Kbyte/sec per call leg • G729 codec: 7.67 Kbyte/sec per call leg • GSM codec: 6 Kbyte/sec per call leg	Critical	Examples of stream components are: Media Gateways, Genesys Voice Platform, SIP phones etc.
Genesys Integration Server (GIS)	Voice client	GSAP	12 Kbyte/call	Very high	Depends on call flow

Table 144: Traffic Generated by 7.5 Inbound Solution with Reporting Components (Continued)

Component1	Component2	Protocol	Estimated bi-directional traffic	Criticality	Comments
GIS	Voice client	SOAP	27 Kbyte/call	Very high	Depends on call flow
GIS	Stat client	SOAP	2.36 Kbyte per update per sec	Very high	Depends on call flow

Generic Call Flow

The generic call flow is as follows:

--> Queue --> 400 DNs -- ttransfer, attach ~110bytes --> Route Point --> 9 queues(100 DNs each) -->

--> 90% agents talk = 10 sec, release

--> 10% agents --ttransfer --> 2 queues(50 DNs each) -->

--> talk = 10 sec, release

Inbound Solution with Reporting Components Traffic Flow

Figure 107 shows the Inbound Solution with reporting components traffic flow.

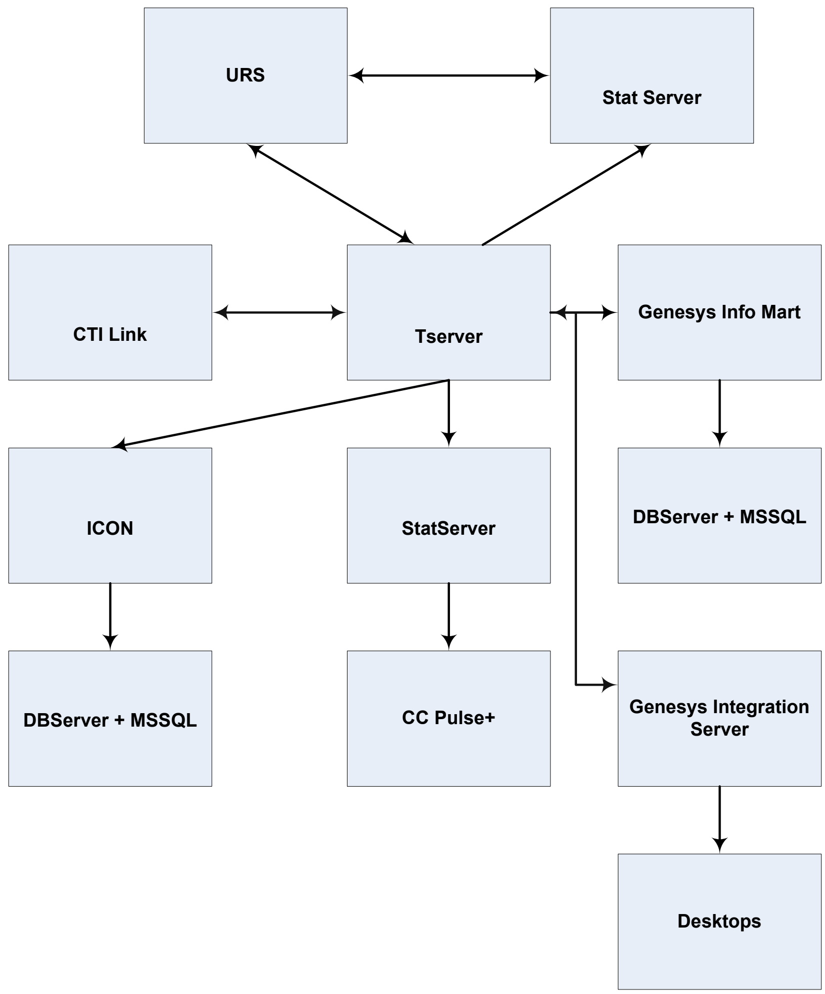

Figure 107: Inbound Solution with Reporting Components Traffic Flow

Chapter

15 Database Sizing Estimation

This chapter presents tools and information for estimating the size of databases for Genesys components. It provides you with a reference to a convenient Microsoft Excel dynamic spreadsheet document that includes database sizing information for various products. In the spreadsheet, you can enter configuration data to automatically calculate an estimation of the database size for your environment.

In addition, this chapter provides references to other sections in this guide and other Genesys deployment guides that provide further information about database sizing estimation. It is recommended that you review the information in these additional resources before you begin any spreadsheet calculation. You might find information about alternative methods of estimation that could be more appropriate for your database.

When using a spreadsheet, keep in mind that it is a tool meant to help you determine an approximate estimate for the required disk memory volume. It is not providing the actual size of the database.

Some of the references provided in this chapter may contain more information on how the estimated size is relevant to the actual size.

This chapter contains:

Database Sizing Estimator Spreadsheets

This section provides information on the Genesys database sizing estimator spreadsheets.

The *Genesys Database Sizing Estimator Spreadsheets* document is a Microsoft Excel file that provides tabs to access sizing spreadsheets for various products, including the following:

- Configuration Database
- Genesys Voice Platform (GVP)
- Multimedia
- Outbound Contact (OCS)
- Data Mart and Operational Data Storage (ODS) databases for Historical Reporting (CCPulse+ and CCAnalyzer)
- Interaction Concentrator (ICON)
- Workforce Management (WFM)

The *Genesys Database Sizing Estimator 7.0* document for the 7.6 release (760_DBSizing_WorkSheets.xls), is a system-level document, and is available on:

- Genesys Technical Support Website at `http://genesyslab.com/support`. From the top menu, click Genesys Knowledge Base; Browse. Select the “system level documents by release" tab; System-Level Documents - 7.6: *Database Sizing Estimator 7.0.*
- Genesys Documentation Library DVD: see the System-Level Documents page. You can order this DVD by e-mail from Genesys Order Management at `orderman@genesyslab.com`.

The *Genesys Database Sizing Estimator 7.0* document for the 7.5 release (750_DBSizing_WorkSheets.xls), is a system-level document, and is available on:

- Genesys Technical Support Website at `http://genesyslab.com/support`. From the top menu, click Genesys Knowledge Base; Browse. Select the “system level documents by release" tab; System-Level Documents - 7.5: *Database Sizing Estimator 7.0.*
- Genesys Documentation Library DVD: see the System-Level Documents page. You can order this DVD by e-mail from Genesys Order Management at `orderman@genesyslab.com`.

Reference Information

For more information, refer to the resources listed below. The Genesys product documentation is available on:

- Genesys Technical Support Website at `http://genesyslab.com/support`.
- Genesys Documentation Library DVD, which you can order by e-mail from Genesys Order Management at `orderman@genesyslab.com`.

Resources

- Configuration Database: see the *Management Framework 7.6 Deployment Guide.*
- Data Mart and Operational Data Storage (ODS): see the *Reporting 7.6 Deployment Guide*.
- Genesys Info Mart: see "Genesys Info Mart Solution" on page 93 in this guide.
- Genesys Voice Platform: see "Genesys Voice Platform 7.x" on page 205 in this guide.
- Multimedia: see the *Multimedia 7.6 Deployment Guide,* and "Multi-Channel Routing and Multimedia" on page 21 in this guide.
- Operational Data Storage (ODS): see the Reporting 7.6 Deployment Guide.
- Outbound Contact: see the Outbound Contact 7.6 Deployment Guide, and"Framework, Reporting, Routing, Outbound, Voice Callback" on page 28 in this guide.

Appendix

A Performance Improvements

This section discusses the performance improvements for Genesys 7.x release and covers the following topics:

Configuration Server Performance Improvements

Availability, Scalability, and Performance Improvements

- Startup time reduced up to 50%
- Enabling technology
  - New startup algorithm
    - Simultaneous initialization and data reading
  - Logical errors detection algorithm
    - Automatic correction or report of logical discrepancies

Note: Such errors no longer prevent Configuration Server Startup.

Logging Improvements

With the new options available in Genesys 7, the log performance can be up to five times better than in Genesys 6, depending on platform.

Log Test Description

Environment

- Two hosts are used in the test environment:
  - Dell Workstation PWS530, with 2 Xenon™ CPU 1.8GHz
  - Dell OptiPlex, Intel Pentium III CPU, 600MHz
- The network between workstations is 100 MB.
- Maximum disk write capacity on first workstation is about 12-13 MB/s.
- For all tests, the same configuration is used.

See Figure 108.

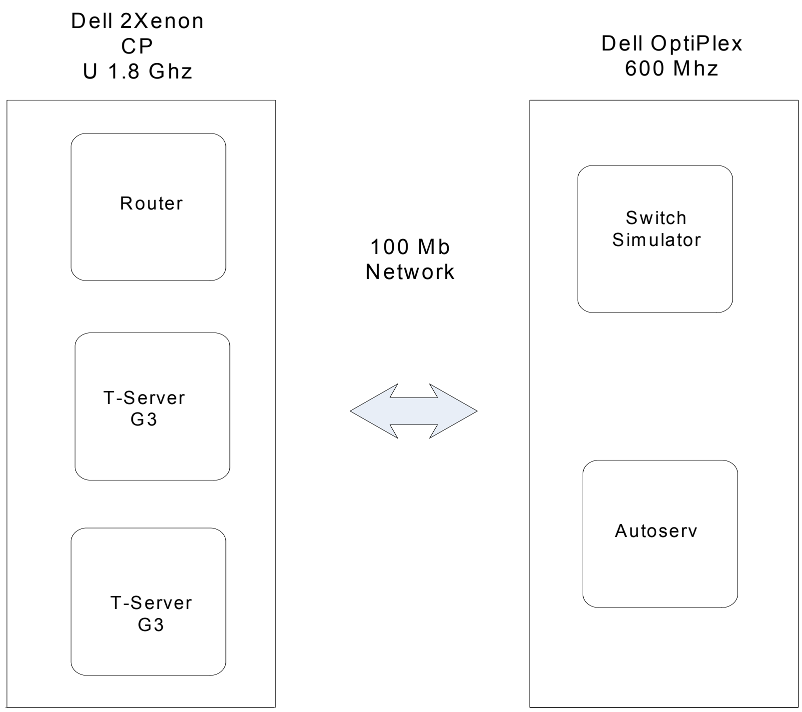

Figure 108: Dell 2 Xenon CPU 1.8 Ghz

Solution Performance Test

1. Two software environments are used: 6.5 and 7.0
   a. Environment 6.5:
      - Router 6.5.110.0
      - Stat server 6.5.101.06
      - T-server G3 6.5.306.01
   b. Environment 7.0:
      - Router 7.0.000.05
      - Stat server 7.0.000.11
      - T-server G3 7.0.011.00

2. Log level is set to the maximum for all applications and is not changed during any of the tests.
3. Log buffering is set to the value "true" for 6.5 applications.
4. All tests are started with 3 calls per second, and the number of calls are increased during test. Results are omitted because the test goal is to find upper limit for call volume.
5. Test is considered successful if there are no abandoned calls during a ten minute interval and are failed or stopped otherwise.
6. Number of calls, CPU, and disk usage are used for test interpretations.

See Table 145 for test results.

Table 145: Test Results

NN	Application	Output	Calls
1	T-Server 6.5	File	7
	Router 6.5	File	
	Stat Server 6.5	File	
2	T-Server 7.0	File	24
	Router 7.0	File	
	Stat Server 7.0	File	
3	T-Server 7.0	Memory map	36
	Router 7.0	Memory map	
	Stat Server 7.0	Memory map	

Results

The test shows significant performance improvement for Genesys 7.0 applications when compared to the 6.5 release. This is due to improved applications design and logging subsystem redesign: the relation is approximately 50-50.

According to the test results, this hardware configuration is recommended for call centers that have an average volume of about 12-18 calls per second and in cases where full logging is necessary. This volume is in middle of calculated results, which allows us to guarantee that the environment will be stable and that there will be almost double reserve for peak loading.

The memory-mapped log output is recommended for troubleshooting cases in a working environment as an output with minimum impact for performance; this output is also recommended for call centers with maximum activity.

If greater call volume is needed, use CPU with greater frequency and setup memory-mapped log output.

Stat Server Performance Improvements 7.0 vs. 6.5

Availability, Scalability, and Performance Improvements

- Ability to apply more statistical filters in a Stat Server instance because of improved filtering mechanism.
- Reduced start-up time, which increases real time statistics availability.
- Increased maximum amount of objects for which Stat Server can calculate statistics:
  - Example of usage: the same statistics can be ordered for more groups of agents.

Stat Server Performance

Improvements 7.0 vs. 6.5

All stat Server tests performed on Micron Netframe 2101 2 CPU Pentium 2 400 Mhz, 512 MB of RAM, Win2000

Performance is measured for 5, 10, and 20 seconds update interval for statistics. See Figure 109

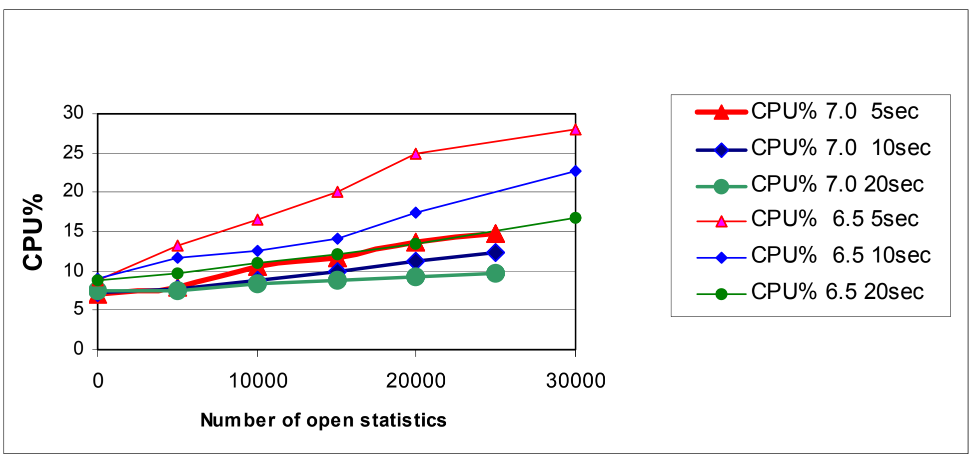

Figure 109: Stat Server Performance Improvements 7.0 vs. 6.5

Skill-Based Routing Performance Enhancement 7.0

- Performance for skill-based routing increased 4-5 times compared to release 6.x.

Note: Routing target is *skill expression*.

- For example:
  v. 6.5 on NT: 16 calls/sec takes 65.4% CPU
  v. 7.0 on NT: 16 calls/sec takes 13.4% CPU
- Hardware: Micron Netframe 2101 2 CPU Pentium

Note: Maximum number of calls per second will be determined later.

- Endurance run 120 hours at 24 calls/sec.

Note: Performance for ACD Queue (agent/virtual group) routing same as in 6.5.

See Figure 110 for Router CPU on List900 strategy.

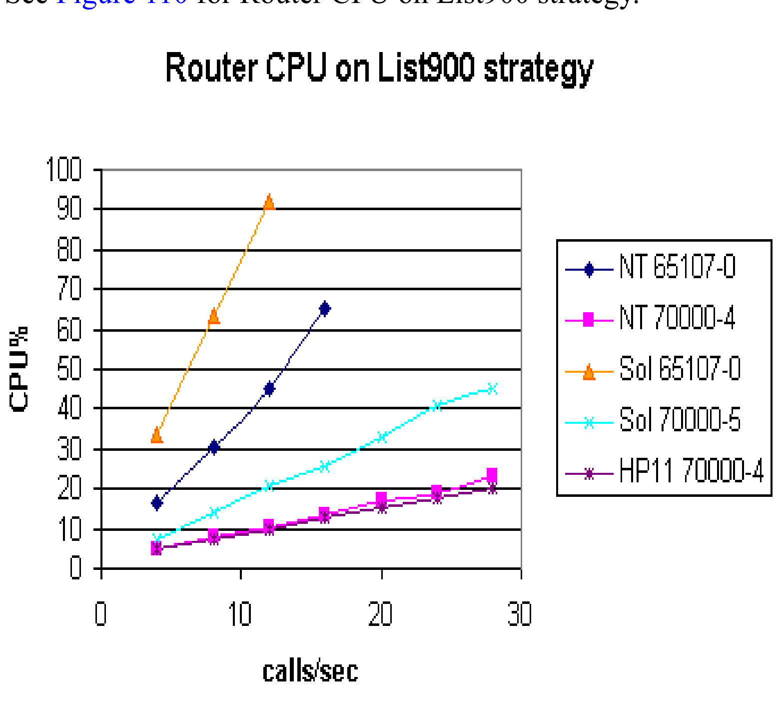

Figure 110: Router CPU on List900 Strategy

Appendix

B Log Database 7.0 Sizing Examples

This appendix contains:

Examples

Size of Log Record

Average size of the single Log Record (record in table G_LOG_MESSAGES), considering index space, is 500 bytes.

Log Record Attributes

Some of the Log Records (in particular, Log Records of level INTERACTION) would have several Attributes attached.

- Each Attribute attached to the Log Record is stored in Log Database as a record in table G_LOG_ATTRS.
- Average size of the single Log Record Attribute (record in table G_LOG_ATTRS) considering index space, is 300 bytes.

Log database space

Total amount of space, allocated by the log database, depends on many factors, in particular:

- How many applications are running.
- Whether or not Network log output is enabled for the applications.
- Log Messages of what level that are sent to the network log output.
- Call volume.

Consider the following simple environment as an example:

- DB Server
- Configuration Server
- Message Server
- Solution Control Server
- T-Server
- URS
- Stat Server

Case 1

Only STANDARD-level log messages are sent to the network log output

Average amount of Standard log messages produced by application could be estimated as approximately 100 per day. In our sample environment, 700 log messages would be written to the database each day. Most of the Standard log messages do not contain attached attributes. Therefore, the size of the log database needed to store log messages produced in 1 day could be estimated as ~340K. In this situation, the amount of space needed to store log records produced within 1 week is ~2.5 MB.

Note that this estimation assumes there is no error condition in our environment. If there is an error condition (such as one of the applications is not started, and other applications are repeatedly trying to connect to it), more log messages could be produced. Since error conditions do not usually persist for a long period of time, the temporary high volume of the log messages could be disregarded in our calculation.

Case 2

INTERACTION-level log messages are sent to the network log output

Interaction-level log messages contain information about ongoing interactions, and on the average has 5 attached attributes each. Therefore, in cases where Interaction-level log messages are sent to the network log output, the size of the log database would depend upon the call volume.

Not all applications produce Interaction log messages. For our sample environment, only T-Server, URS, and Stat Server produce such log messages.

The amount of Interaction log messages per call depends significantly upon the call scenario and the applications.

1. For example, let us assume that:
   - T-Server produces on the average 10 Interaction-level log messages per Interaction.
   - URS produces on the average 5 Interaction-level log messages per Interaction.
   - Stat Server produces on the average 5 Interaction-level log messages per Interaction.
2. Let us assume that:
   - The average call volume is 5 calls per second.
3. Therefore:

   (10 + 5 + 5)*5 = 100 Interaction log messages will be produced each second,

   with

   100 * 5 = 500 attached attributes records
4. These log records would require:

   100*500 + 500*300 = 195Kb of log database space per second

   Which for one hour is:

   195 * 60 * 60 = ~690 MB of log database space.

Note: The Interaction log messages are targeted for **test environments** in order to tune the interaction-related messaging between applications. It is not recommended to send Interaction-level log messages to the Network log output in the production environment.

Case 3

TRACE-level log messages are sent to the network log output

Applications produce even more Trace log messages than Interaction or Standard log messages. Therefore, it is also not recommended to send Trace log messages to the Network log output in the production environment.

Log Length Dependence for 7.0

- Call Flow—8 calls/sec, 5000 calls made.
- Attach data ~ 108 bytes
- Each application generates snapshot log file ~ 128 Kb always.

Table 146 specifies 7.0 applications:

Table 146: 7.0 Applications

Application	All		Trace		Standard		None		All/None Rating
	CPU%	Log (MB)	CPU%	Log (MB)	CPU%	Log (MB)	CPU%	Log (MB)	
T-Server	32.0	318.0	15.1	26.200	13.5	0.001	13.3	0.001	2.4
T-Server backup	21.0	269.0	6.3	0.001	5.8	0.001	6.0	0.001	3.5
HA-proxy	9.5	65.4	6.0	0.001	5.5	0.001	5.8	0.001	1.6
Ha-proxy backup	3.0	31.4	1.2	0.002	1.2	0.001	1.2	0.001	2.5
Router	4.7	51.3	2.8	0.776	2.8	0.001	1.8	0.001	2.6
Stat Server	7.3	62.0	4.0	8.977	3.2	0.001	3.4	0.001	2.1

Table 147 specifies 6.1 applications.

Table 147: 6.1 Applications

Application	All		Trace		Standard		None		All/None Rating
	CPU%	Log (MB)	CPU%	Log (MB)	CPU%	Log (MB)	CPU%	Log (MB)	
T-Server	26	200	7	0	7	0	7.5	0	3.5
T-Server backup	17.5	165	3.2	0	3.2	0	3.5	0	5.0
HA-proxy	13.5	61	5.2	0	5.2	0	3.5	0	3.9

Table 147: 6.1 Applications (Continued)

Application	All		Trace		Standard		None		All/None Rating
	CPU%	Log (MB)	CPU%	Log (MB)	CPU%	Log (MB)	CPU%	Log (MB)	
Ha-proxy backup	1.7	17	0.4	0	0.4	0	0.5	0	3.4
Router	17	120	11	1.9	11	0	3.3	0	5.2
Stat Server	9.2	80	6.2	0	6	0	3.7	0	2.5

Supplements

Related Documentation Resources

The following resources provide additional information that is relevant to this software. Consult these additional resources as necessary.

Genesys

- *Genesys Technical Publications Glossary,* which ships on the Genesys Documentation Library DVD and which provides a comprehensive list of the Genesys and computer-telephony integration (CTI) terminology and acronyms used in this document.
- *Genesys Migration Guide*, which ships on the Genesys Documentation Library DVD, and which provides documented migration strategies for Genesys product releases. Contact Genesys Technical Support for more information.
- Release Notes and Product Advisories for this product, which are available on the Genesys Technical Support website at `http://genesyslab.com/support`.

Information about supported hardware and third-party software is available on the Genesys Technical Support website in the following documents:

- *Genesys Supported Operating Environment Reference Manual*
- *Genesys Supported Media Interfaces Reference Manual*

Consult these additional resources as necessary:

- *Genesys 7 Interoperability Guide,* which provides information on the compatibility of Genesys products with various Configuration Layer Environments; Interoperability of Reporting Templates and Solutions; and Gplus Adapters Interoperability.
- *Genesys Licensing Guide,* which introduces you to the concepts, terminology, and procedures relevant to the Genesys licensing system.
- *Genesys Database Sizing Estimator 7.6 Worksheets,* which provides a range of expected database sizes for various Genesys products.

For additional system-wide planning tools and information, see the release-specific listings of System Level Documents on the Genesys Technical Support website, accessible from the `system level documents by release` tab in the Knowledge Base `Browse Documents` Section.

Genesys product documentation is available on the:

- Genesys Technical Support website at `http://genesyslab.com/support`.
- Genesys Documentation Library DVD, which you can order by e-mail from Genesys Order Management at `orderman@genesyslab.com`.

Document Conventions

This document uses certain stylistic and typographical conventions—introduced here—that serve as shorthands for particular kinds of information.

Document Version Number

A version number appears at the bottom of the inside front cover of this document. Version numbers change as new information is added to this document. Here is a sample version number:

80fr_ref_06-2008_v8.0.001.00

You will need this number when you are talking with Genesys Technical Support about this product.

Screen Captures Used in This Document

Screen captures from the product graphical user interface (GUI), as used in this document, may sometimes contain minor spelling, capitalization, or grammatical errors. The text accompanying and explaining the screen captures corrects such errors *except* when such a correction would prevent you from installing, configuring, or successfully using the product. For example, if the name of an option contains a usage error, the name would be presented exactly as it appears in the product GUI; the error would not be corrected in any accompanying text.

Type Styles

Table 148 describes and illustrates the type conventions that are used in this document.

Table 148: Type Styles

Type Style	Used For	Examples
Italic	• Document titles • Emphasis • Definitions of (or first references to) unfamiliar terms • Mathematical variables Also used to indicate placeholder text within code samples or commands, in the special case where angle brackets are a required part of the syntax (see the note about angle brackets on page 329).	Please consult the *Genesys Migration Guide* for more information. Do *not* use this value for this option. A *customary and usual* practice is one that is widely accepted and used within a particular industry or profession. The formula, $x + 1 = 7$ where x stands for . . .
Monospace font (Looks like `teletype` or `typewriter text`)	All programming identifiers and GUI elements. This convention includes: • The *names* of directories, files, folders, configuration objects, paths, scripts, dialog boxes, options, fields, text and list boxes, operational modes, all buttons (including radio buttons), check boxes, commands, tabs, CTI events, and error messages. • The values of options. • Logical arguments and command syntax. • Code samples. Also used for any text that users must manually enter during a configuration or installation procedure, or on a command line.	Select the `Show variables on screen` check box. In the `Operand` text box, enter your formula. Click `OK` to exit the `Properties` dialog box. T-Server distributes the error messages in `EventError` events. If you select `true` for the `inbound-bsns-calls` option, all established inbound calls on a local agent are considered business calls. Enter `exit` on the command line.
Square brackets ([ ])	A particular parameter or value that is optional within a logical argument, a command, or some programming syntax. That is, the presence of the parameter or value is not required to resolve the argument, command, or block of code. The user decides whether to include this optional information.	`smcp_server -host [/flags]`
Angle brackets (< >)	A placeholder for a value that the user must specify. This might be a DN or a port number specific to your enterprise. **Note:** In some cases, angle brackets are required characters in code syntax (for example, in XML schemas). In these cases, italic text is used for placeholder values.	`smcp_server -host <confighost>`

Index