# PureConnect Common Components Blueprint

Reference Architecture

| | |
|---|---|
| Authors | Adam Elkins |
| | Dave Gussin |
| | Pat O'Roark |
| | Reid Forrest |
| | Tom Bova |
| | Vikki Papesh |

| | |
|---|---|
| Version: | 1.0 |
| Status: | PUBLISH |

GENESYS™

# Table of Contents

1 Introduction .......... 1

1.1 Document Overview .......... 1

1.2 Intended Audience .......... 1

2 Definitions, Acronyms, and Document Standards .......... 2

2.1 Definitions .......... 2

2.2 Glossary .......... 2

2.3 Document Conventions .......... 3

3 Overall Architecture .......... 4

3.1 Logical Architecture Model .......... 5

3.2 Functional View .......... 6

3.2.1 Application Layer .......... 6

3.2.2 Media Processing Layer .......... 7

3.3 Standard Use Cases .......... 7

3.4 Component View .......... 9

3.4.1 Application Layer .......... 10

3.4.2 Media Processing Layer .......... 12

3.4.3 Client Applications .......... 15

3.4.4 3rd Party Components .......... 15

3.5 Limits and Constraints .......... 16

4 Deployment View .......... 17

4.1 Network Overview .......... 17

4.2 Genesys Deployment Options .......... 18

4.2.1 Deployment Option Overview .......... 18

4.2.2 High Availability .......... 21

4.2.3 Disaster Recovery .......... 47

4.3 Database .......... 48

5 Interaction View .......... 62

5.1 Agent Experience .......... 62

5.2 Call Flows .......... 63

5.3 External Interfaces ....65
5.4 Operational Management ....67
5.4.1 Network Management Systems ....67
5.4.2 Serviceability ....68
5.4.3 Usage ....68
6 Implementation View ....70
6.1 Solution Sizing Guidelines ....70
6.1.1 Sizing Assumptions ....70
6.1.2 Media Server Sizing ....71
6.1.3 Storage Sizing ....72
6.1.4 Database Sizing ....72
6.1.5 Network Sizing and Readiness ....73
6.1.6 Solutions Sizing - Dual Data Center ....73
6.2 Configuration Guidelines ....73
6.2.1 Configuration Recommendations ....74
6.3 Security ....74
6.3.1 Secure Connections ....74
6.3.2 VM and OS hardening ....75
6.3.3 PCI ....75
6.4 Localization and Internationalization ....75

## Table of Figures

Figure 1: Logical Architecture Model ....5
Figure 2: Common Component Detail ....9
Figure 3: Recording Interactions with Interaction Recorder Remote Content Service ....14
Figure 4: SIP Proxy Business Continuity with Local Trunking at Remote Site ....18
Figure 5: PureConnect Single Datacenter Model ....19
Figure 6: PureConnect Geo-Redundant Datacenter Model ....19
Figure 7: PureConnect Replicated Datacenter Model ....20
Figure 8: Disaster Recovery deployment (any delivery model) ....21

Figure 9: The Switchover State Machine ....26
Figure 10: How the Primary Monitor Determines the Status of the Network and the Connection ....29
Figure 11: Primary Monitor: Detect and Signal Connection Status....33
Figure 12: Primary Monitor: Process Module Monitor Notification ....35
Figure 13: Primary Monitor: Send Ping Request ....37
Figure 14: Primary Monitor: Ping Response Successful ....38
Figure 15: Primary Monitor: Schedule Next Ping ....39
Figure 16: Primary Monitor: Ping Response Timeout Ping Response Rejected....40
Figure 17: Primary Monitor: Ping Response Connection Loss....41
Figure 18: Reporting Database Growth Table ....48
Figure 19: Agent Desktop - Interaction Desktop ....62
Figure 20: Agent Desktop - Interaction Connect ....63
Figure 21: Standard Call Flow ....63
Figure 22: Call Flow Tools - Interaction Attendant....64

## Table of Tables

Table 1: Glossary....2
Table 2: Genesys Desktop Component List ....15
Table 3: 3rd Party Components ....16
Table 4: Database and Log Growth Chart....55
Table 5: External Interfaces ....67
Table 6: Sizing Inputs ....71
Table 7: Media Server Points....71
Table 8: Supported Languages....77

# Revision History

| Rev | Date Published | Author | Reason for Revision |
|---|---|---|---|
| 0.1 | | Adam Elkins, Dave Gussin,<br>Pat O'Roark, Reid Forrest, Tom Bova, Vikki Papesh | Initial draft |
| 1.0 | 2018/11/12 | | Published |

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS DOCUMENT ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENT ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THIS DOCUMENTATION IS BEING PROVIDED GRATUITOUSLY AND, THEREFORE GENESYS SHALL NOT BE LIABLE UNDER ANY THEORY FOR ANY DAMAGES SUFFERED BY LICENSEE OR ANY USER OF THE GENESYS DOCUMENTATION. UNDER NO CIRCUMSTANCES, INCLUDING NEGLIGENCE, SHALL GENESYS BE LIABLE FOR ANY INCIDENTAL, SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES ARISING OUT OF OR RELATING TO THIS DOCUMENTATION. SOME JURISDICTIONS DO NOT ALLOW THE LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES SO THIS LIMITATION MAY NOT APPLY TO YOU.

# 1 Introduction

The purpose of this Blueprint is to document the architecture for Genesys PureConnect. This document provides a prescriptive list of components (both Genesys and 3rd party) that should be included in the solution. It also provides guidance for implementing and deploying the solution including sizing and configuration as well as addressing several system concerns such as security, high availability, disaster recovery and serviceability.

## 1.1 Document Overview

The document contains the following sections:

- 2: Definitions and Acronyms
- 3: Overall Architecture
- 4: Deployment View
- 5: Interaction View
- 6: Implementation View

## 1.2 Intended Audience

Describing system and solution architectures can be difficult as there are multiple audiences each with different expectations. This document is intended for multiple audiences with various chapters being more interesting to some readers than others.

The *Overall Architecture* and *Deployment View* sections are likely meaningful to most audiences. However, the *Interaction View* and the *Implementation View* sections may be of more interest to those configuring the network and components.

# 2 Definitions, Acronyms, and Document Standards

## 2.1 Definitions

This document uses various abbreviations and acronyms that are commonly used in Genesys product documentation and the telecommunications and contact center industries. The following table defines terms that will be referenced subsequently in this document.

## 2.2 Glossary

| | |
|---|---|
| CTI | Computer-telephony integration, the adding of computer intelligence to monitoring and control of telephone calls |
| DB | Database |
| DBMS | Database Management System |
| DHCP | Dynamic Host Configuration Protocol |
| DN | Directory number |
| DNS | Domain Name System |
| FTP | File Transfer Protocol |
| GUI | Graphical User Interface |
| HA | High Availability |
| HTTP | Hypertext Transfer Protocol |
| IP | Internet Protocol |
| IVR | Interactive Voice Response |
| LAN | Local Area Network |
| LCA | Local Control Agent |
| OEM | Original Equipment Manufacturer |
| OS | Operating System |
| RDBMS | Relational Database Management System |
| SCS | Solution Control Server |
| SCXML | State Chart XML: State Machine Notation for Control Abstraction |
| SQL | Structured Query Language |
| VM | Virtual Machine |
| WAN | Wide Area Network |

*Table 1: Glossary*

## 2.3 Document Conventions

The following documentation and naming conventions are used throughout the document:

- `Code` and `configuration property names & values` will appear in `console font`.
- References to other documents are hyperlinked or italicized.

# 3 Overall Architecture

The PureConnect Common Components architecture covers standard elements which are shared by multiple blueprint architectures with the common components establishing the foundation that should be present in any architecture.

Genesys provides additional use case blueprints which cover add-on components required in separate documentation.

The appropriate PureConnect Blueprint Architecture should be consulted in addition to the Common Components to provide a complete understanding of the architecture.

Much of the guidance provided in this document is based on customer deployments and the lessons learned through implementing PureConnect within those real-world scenarios.

This Common Components architecture is focused on two areas:

1. Application Layer
2. Media Processing Layer.

***PureConnect Application Layer***– The PureConnect Application Layer includes the PureConnect Servers, Database and a Web Server for the Interaction Connect web-based client.

***PureConnect Media Processing Layer*** – The PureConnect Media processing layer includes Interaction Media Server, Remote Content Service and recording storage.

## 3.1 Logical Architecture Model

The following is a logical model of the PureConnect Common Components Architecture.

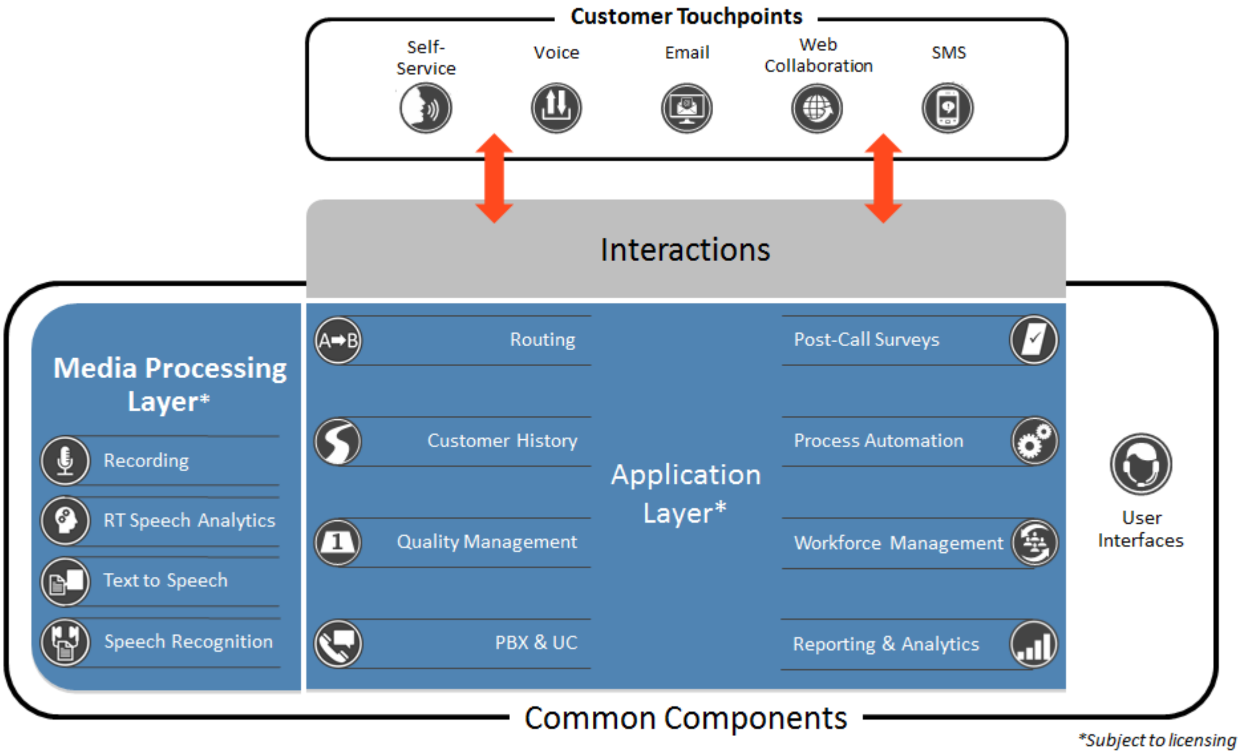

*Figure 1: Logical Architecture Model*

## 3.2 Functional View

The Common Components architecture delivers a consistent architecture that is intended to be leveraged across all solutions. Many customer deployments may start out with a single solution. The establishment of a common architecture for shared components allows the architecture to consider future customer needs both managing interaction types and the application of more advanced business strategies. Establishing a comprehensive common component architecture allows customers to easily integrate new Genesys components or evolve their customer experience management strategy, such as the sophistication of the routing that is used, without changes to the core foundation.

The functionality delivered can be broken down into the following areas:

- Applications
- Media Processing

### 3.2.1 Application Layer

The Applications provided by the common components deliver the following functionality, enabled via software licensing:

- Centralized configuration of all applications and resources in the PureConnect environment ranging from software configuration settings to resources such as agents and workgroups
- Intelligent routing of interactions (e.g. calls, chats, etc.) to available agents in Workgroup queues
- Access to contextual information about the current customer interaction, prior interactions and customer details to optimize the routing decisions
- Quality management tools for interaction playback and evaluation
- Built-in enterprise-grade PBX and unified communications tools such as internal chat
- Post-call survey tools capture the voice of the customer
- Workforce engagement tools for forecasting and scheduling agent activities
- Integration with external systems via a variety of methods including web services, DB, file/IO and TCP/IP

The Application Layer also provides reporting & analytics capabilities, which includes the following:

- Real-time dashboards for supervisors, team leads, and IT managers
- Canned historical reports, wizard-driven report creation, and the ability to upload custom Crystal-based reports
  - For interactions, queue statistics, agent and user activity, line and line groups, admin changes
- A well-defined, concise and documented data model enabling the historical data to be directly accessed and used by 3rd party analytics packages (PureConnect Data Dictionary)
- Intra-day metrics on contact center performance
- Visual alerting based upon user configured thresholds
- Integration with external sources through a variety of methods including web services, DB, file/IO and TCP/IP
- Report distribution capabilities to enable delivery of reports automatically via a variety of mediums (email, FTP, printer) and in various formats (PDF, Excel, CSV)

### 3.2.2 Media Processing Layer

The Media Processing provided by the common components enables:

- Recording, management and playback of audio interactions
- Real-time speech analytics via keyword spotting and assigning positive or negative values corresponding to spoken words or phrases
- Speaking out text and data to personalize call flows (TTS)
- Recognizing spoken words to direct call flows and provide self-service options (ASR)
- Audio conferencing with three or more parties
- Transcoding and encryption

## 3.3 Standard Use Cases

The Common Components Blueprint architecture provides a common set of capabilities which are used by all other Blueprint Architectures and required for all use cases. Details on the Standard Use Cases are available on Genie under Use Cases, as well as the Documentation portal.

Use cases which can be accomplished with components covered in the Common Components Blueprint alone are as follows. Click on links within the table to see a full description of use cases. Additional solution consulting and sizing will be necessary:

| ID | Title | Product Category | Subtitle |
|---|---|---|---|
| CE01 | Genesys Call Routing | Inbound | Route voice interactions to the best skilled resource |
| CE02 | Genesys Personalized Routing | Inbound | Apply personalized routing to voice interactions |
| CE07 | Genesys Customer Authentication | Self-Service and Automation | Identify and verify customers in your IVR |
| CE08 | Genesys Voice Payment | Self-Service and Automation | Capture payments in your IVR |
| CE09 | Genesys IVR Personalization | Self-Service and Automation | Increase self-service by personalizing your IVR |

| ID | Title | Product Category | Subtitle |
|---|---|---|---|
| CE16 | Genesys Email Routing | Digital | Route email interactions to the best skilled resource |
| CE18 | Genesys Chat Routing | Digital | Route chat interactions to the best skilled resource |
| CE22 | Genesys Digital Callback | Digital | Enable customers to request a callback from your website or app |
| CE29 | Genesys SMS Routing | Digital | Route SMS interactions to the best resource |
| CE32 | Genesys IVR MicroApps | Self-Service and Automation | Automate phone conversations with pre-built MicroApps and agent handover if needed |
| CE33 | Genesys Visual IVR | Self-Service and Automation | Extend engagement with customers from voice to online with Genesys MicroApps |
| EE07 | Genesys Voice Recording | Workforce Engagement | Record voice interactions |
| EE08 | Genesys Voice and Screen Recording | Workforce Engagement | Record voice and screen interactions |
| EE09 | Genesys Quality Management | Workforce Engagement | Improve employee performance with quality management |
| OP01 | Genesys Business Communications | Open Platform | Simplify contact center and business communications |

## 3.4 Component View

The Component View describes the higher-level modules that make up the solution. The following diagram depicts the components required as part of the solution.

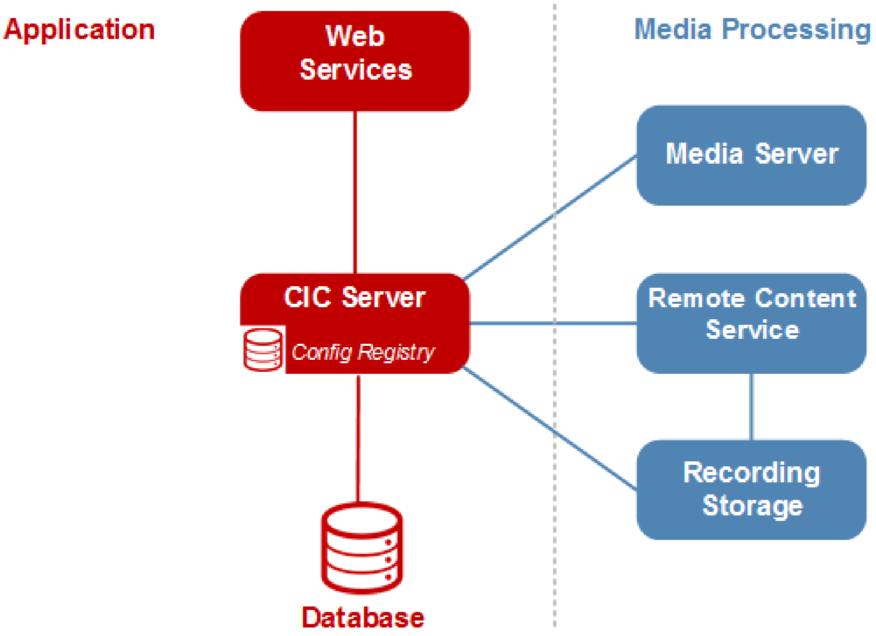

*Figure 2: Common Component Detail*

Note: The diagram is intended to provide a clear understanding of the different areas which are addressed by the Blueprint and not all components or connections are shown, e.g. the switchover CIC Server instance is not represented.

The diagram does show some detail such as the Remote Content Service. Deployment of the Remote Content Service is a common practice; however, CIC Server includes the Remote Content Service functionality natively (as a subsystem) so a separate instance may not be required for smaller deployments.

### 3.4.1 Application Layer

The following components are included within the logical Application layer:

- Customer Interaction Center (CIC) Server
- Web Services
- Database

Note: Web Services are utilized across several components however it may not be mandatory within your specific architecture.

**Customer Interaction Center (CIC) Server**

The CIC Server is a pure application server for the PureConnect platform and is the main conduit of communication between other components of the platform.

At its core, CIC Server is a routing engine. CIC supports multiple ways to route interactions (for example, calls and chat sessions) directly to any user or workgroup (queue). Automatic Call Distribution (ACD) is the intelligent routing of interactions to available agents in Workgroup queues. PureConnect uses assigned skill requirements to intelligently route incoming interactions to a qualified, available agent. All routing, whether direct or ACD based is native to the PureConnect solution and delivered via the CIC Server. Additional details on advanced routing can be found here.

CIC Server also provides many other features layered on top of the routing engine that are enabled via software licensing. For example, customer interaction history, quality management, workforce management, post-call surveys, PBX, unified communications, and more.

**Web Services**

PureConnect Web Services can serve two different roles: internal web-based applications (e.g. Interaction Connect, Data Extractor) and external web-based applications (e.g. web chat, web-based scheduled callbacks).

The following internal-facing web-based features are provided:

- **Interaction Connect** – web-based interaction manager offering call control functionality, chat features, and status management. No components are installed on the desktop.
- **Data Extractor** – web-based application that provides the ability to export CIC product data without requiring the user to have direct access to the IC database. Interaction Data Extractor generates an easily consumable, flat-format CSV file, which can be utilized directly or imported into a variety of powerful, 3rd-party Business Intelligence (BI) tools.
- **Interaction Optimizer Web** – agents can access their schedule and request time off through a dedicated web portal.
- **Interaction Process Automation Web Clients** – includes the IPA Work Item Client and the IPA Work Item Viewer. Those components are web-based, IPA-only applications for users who do not need the more full-featured Interaction Desktop. Users can pick up IPA work items, act on them, and then transfer them. They can also change their own statuses, activate or deactivate themselves in workgroups, and perform other tasks.

The following external-facing web-based features are provided via *Interaction Web Tools:*

- **Web Chat** – an agent can engage in text-based, one-on-one communication online with a website visitor or create a chat conference between multiple agents and a website visitor.
- **Intercom Chat** – an agent can chat online with fellow employees or a group of fellow employees.
- **Callback** – individuals can visit your website and request that an agent or member of a workgroup call them back. CIC routes the request to the appropriate agent or workgroup.
- **Response Management** – agents can use Response Management during a chat to send predefined messages, URLs, or files to a website visitor. An agent can also use Response Management for guidance in answering a website visitor's questions during a call made because of a callback request.

### Database

PureConnect uses industry-standard database connections to store call detail records, supporting Microsoft SQL Server (premise and cloud deployments) and Oracle databases (premise deployments only). The Genesys PureConnect Data Dictionary provides a concise and detailed description of the database tables Customer Interaction Center (CIC) uses to store and generate reports on historical data regarding interactions, queue statistics, agent and user activity, line and line groups, administrative changes, and other Interaction Administrator configuration information. In addition to documenting the structure and contents of CIC tables, a summary of historical data that CIC collects is available in this document.

PureConnect dashboards and reports are accessed primarily using IC Business Manager (Windows .NET desktop application). Several reporting and analytics tools are also offered via standard web browsers.

### 3.4.2 Media Processing Layer

The following components are included within the logical Media Processing layer:

- Media Server
- Remote Content Service
- Recording Storage

**Note:** The Remote Content Service is deployed for scale and/or to offload CIC Server from CPU-intensive processing for screen recording file management; it may not be mandatory within your specific architecture.

**Media Server**

Interaction Media Server is a required subsystem of Customer Interaction Center (CIC). Its main purpose is to handle audio streams in Internet Protocol communications. These audio streams include voice over IP (VoIP) active call connections, recordings, and the playing of prompts. Interaction Media Server does not issue or receive Session Initiation Protocol (SIP) commands, which control the establishment, management, and termination of telephone calls in the CIC environment. When Interaction Media Server must make a connection, change, or termination between two or more SIP calls, CIC supplies the necessary commands.

Interaction Media Servers are deployed in an N+1 configuration for redundancy where N represents the number of Interaction Media Servers required to service a system or location. The additional Interaction Media Server enables the system to continue functioning at its capacity without resource limitations when you must maintain or troubleshoot an Interaction Media Server.

Interaction Media Server handles the following media operations:

- **Call analysis** – Interaction Media Server detects if a person, an answering machine, or a voice mail system answers the call.
- **Recording calls** – Interaction Media Server records call conversations for the agent, the external party, or both parties. Interaction Media Server can record the audio from both participants on the call in one channel (mono) or two channels (stereo).
- **Securing call recordings** – When recording a call, Interaction Media Server encrypts and compresses the recording so that only approved users can listen to the calls.
- **Playing recordings** – When you play a call from Interaction Recorder Client, Interaction Media Server decompresses and streams the call recording.
- **Playing prompts** – When a caller enters your CIC system, an Interactive Voice Response system can present the caller with audio prompts to select digits that represent interest in speaking to a specific person or representative. Interaction Media Server streams these prompts into the call.
- **DTMF recognition** - Interaction Media Server recognizes Dual-Tone Multi Frequency (DTMF) tones, both in-band and with the RFC 2933 specification.
- **Playing on-hold music** – When an agent puts a call on hold, Interaction Media Server plays music to the other party while they wait.
- **Transcription and transcoding** – Interaction Media Server dynamically converts calls that use different codecs (G.711, G.729) or protocols (RTP, SRTP) so that all parties hear the audio streams.

- **Conferencing** – Interaction Media Server facilitates all audio communication for conference calls. For conference calls with fewer than 20 participants, the system hosts the call on a single Interaction Media Server. For a conference call with more than 20 participants, the system processes the call through multiple Interaction Media Servers.
- **Keyword spotting** – Interaction Media Server analyzes speech during telephone conversations and can recognize predefined keywords when spoken. For more information about this feature, see Interaction Analyzer Technical Reference.
- **Speech recognition** – Interaction Media Server analyzes speech provided through Interactive Voice Recognition, such as with Interaction Attendant. For information about this feature and provided grammars, see Interaction Speech Recognition Technical Reference and Interaction Attendant Help.
- **Text-to-Speech** - Interaction Media Server supports Speech API (SAPI), Media Resource Control Protocol (MRCP) version 2.0, and its own native Interaction Text to Speech (ITTS) product for playing of audio streams from TTS engines.

### Remote Content Service

The Interaction Recorder Remote Content Service (RCS) facilitates the retrieval and storage of both audio and screen recordings in your PureConnect environment. This capability offloads those actions from the
Interaction Recorder subsystem that resides on the Customer Interaction Center server thereby granting that server more processing and bandwidth resources for facilitating and handling interactions.

In PureConnect, there are three phases to making a recording:

1. Record, compress, and encrypt phase – All three of these actions are done simultaneously through Interaction Media Server or, in the case of screen recordings, on the agent workstation.
2. Storage phase – Moving the compressed and encrypted recording file to a location, such as a remote file server or SAN. **RCS is responsible for this phase.** While the call is recorded on the Media Server, the **Media Server does not do any file puts** once the recording is complete.
3. Database phase – Creating an entry in the PureConnect database so that the recording is cataloged and identified to reside in a specific storage location.

The following diagram displays the process of recording interactions with Interaction Recorder Remote Content Service:

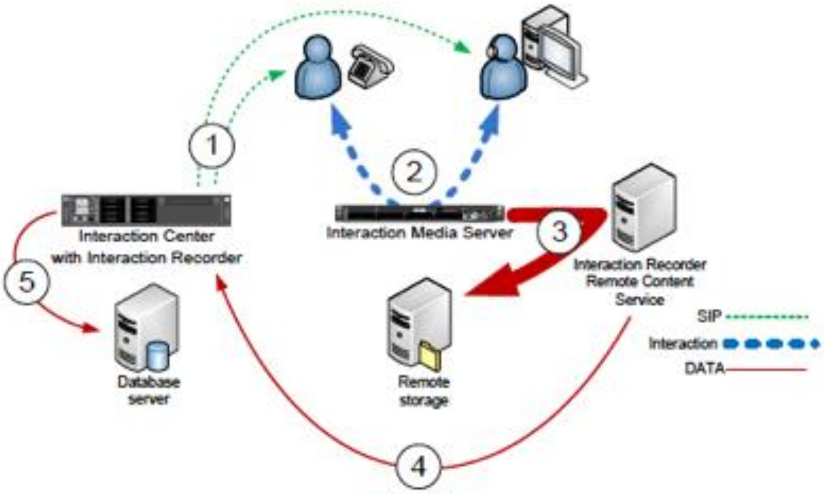

*Figure 3: Recording Interactions with Interaction Recorder Remote Content Service*

| Step | Description |
|---|---|
| 1 | PureConnect facilitates an interaction between a customer and an agent. |
| 2 | Interaction Media Server records, compresses, and encrypts the interaction |
| 3 | Interaction Recorder Remote Content Service moves and stores the recording either to itself or to a remote storage location |
| 4 | Interaction Recorder Remote Content Service notifies Interaction Recorder of the name and location of the recording. |
| 5 | Interaction Recorder writes an entry to the database that the name and location of the recording. |

Further information on Remote Content Service may be found in the PureConnect Technical Reference Library.

**Recording Storage:**

PureConnect audio and screen recordings must be stored in a location other than the PureConnect Server and Interaction Media Servers. This can be a SAN, NAS, storage server, etc. that must be accessible via a UNC (Universal Naming Convention) path such as \\server\shared_folder. PureConnect also supports storing recordings in AWS S3 buckets.

### 3.4.3 Client Applications

The following desktop applications or deployment services may also be deployed. These applications typically execute at the local desktop rather than as server applications.

| Category | Component | Notes |
|---|---|---|
| **Agent Application** | Interaction Desktop | Optional desktop agent application. Can be used in place of - or in addition to - the web-based agent interface (Interaction Connect) |
| **Supervisor/Manager Application** | Interaction Business Manager | Used by supervisors, managers and team-leads to operate, monitor, report on, and analyze contact center performance |
| **Script Development** | Interaction Attendant | Used for Auto-Attendant and IVR development |
| **System Administration Application** | Interaction Administrator | Used to administer most of the platform |
| **System Customization Application** | Interaction Designer | Used for advanced IVR development and system customization |

*Table 2: Genesys Desktop Component List*

### 3.4.4 3rd Party Components

The following table lists the recommended 3rd party components for the Common Components blueprint architecture.

| Category | Recommended | Other Supported Components |
|---|---|---|
| **Database** | Microsoft SQL Server | Oracle |
| **Web Server** | Microsoft IIS | NGINX, Apache |
| **File Server** | NAS/SAN | |
| **Virtualization** | VMWare | Hyper-V |
| **HTTP Load Balancer** | Customer preference | |

*Table 3: 3rd Party Components*

For detailed information about 3rd party products and support, see http://testlab.inin.com/

## 3.5 Limits and Constraints

There are components which have specific limitations on the operating system, database or deployment model. Key limitations are:

- The core components of CIC Servers, Interaction Media Servers and Remote Content Servers are supported on Windows Server OS only
- Interaction Media Servers are supported on physical hardware only--either on appliances sold by Genesys or on physical hardware provided by the customer in conjunction with Media Server software licenses sold by Genesys
- While PureConnect supports both MS SQL and Oracle databases, Oracle is supported with premise deployments only
- PureConnect core components including CIC Servers, Media Servers and RCS servers communicate with each other via a proprietary protocol called Notifier which is not supported over the Internet; these core components must communicate via private WAN or VPN
- Remote Content Server scalability constraints

# 4 Deployment View

## 4.1 Network Overview

PureConnect requires sufficient network bandwidth to deliver SIP and RTP/SRTP traffic in real time across the LAN and WAN. While load balancing NIC teaming is not supported, fault tolerant NIC teaming may be implemented to ensure a consistent connection across switches, and QoS tagging is implemented in accordance with RFC 2474. To facilitate proper configuration, a PureConnect QoS driver is installed during installation of server and client products. All switches and routers on the network must honor and pass the traffic with tagging intact. For more information, please consult the PureConnect Quality of Service Technical Reference document: https://help.genesys.com/cic/mergedProjects/wh_tr/desktop/pdfs/qos_tr.pdf

Use of dedicated voice VLANs is required to segregate voice traffic from broadcast traffic and to provide additional call security. All voice endpoints, including Interaction Media Servers and other servers in the CIC environment, should be placed in the voice VLAN. CIC clients, including those with soft phones, may remain in the data or default VLAN.

Interaction SIP Proxy is a program that enables a server to route SIP requests to a LAN, WAN, or PSTN. Additionally, Interaction SIP Proxy can host configuration files and firmware for IP telephones, facilitate locally-based routes, and provide consistent routing should a network connection or Customer Interaction Center server interruption occur between the remote site and the datacenter.

If the WAN link fails, SIP requests from the clients are routed through the Interaction SIP Proxy, and the audio path is sent directly to the gateway. This allows direct inbound and outbound calls at the remote office to continue in the event connectivity to CIC or primary trunking is down.

Interaction SIP Proxy uses regular expressions to evaluate SIP requests and route them to configured destinations. Multiple routes may be specified to provide greater fault tolerance, and the routing options include sequential, round robin, and randomization. Real time server status, visible through a web interface, displays the status of the connection to all configured destinations.

SIP proxy can act as a redirect server for SIP INVITE requests, and can serve as a registrar for SIP devices, eliminating the need for these devices to register directly with CIC and conserving WAN bandwidth. For continuity in a failover scenario, supported SIP endpoints can register to both Customer Interaction Center and SIP proxy. For more information on SIP Proxy, please consult the SIP Proxy Technical Reference: https://help.genesys.com/cic/mergedProjects/wh_ps/desktop/pdfs/sip_proxy_tr.pdf

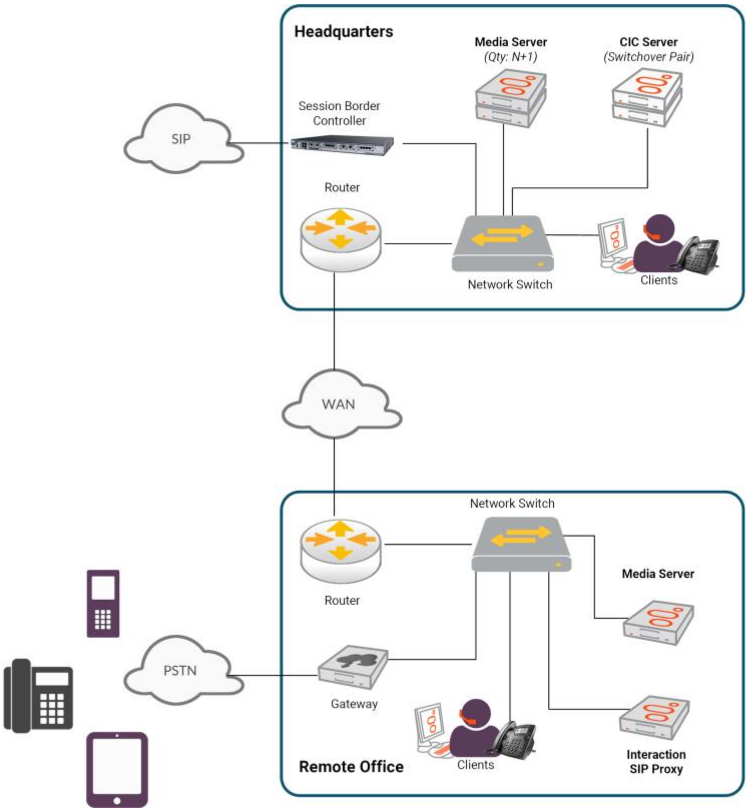

*Figure 4: SIP Proxy Business Continuity with Local Trunking at Remote Site*

## 4.2 Genesys Deployment Options

Genesys offers several supported PureConnect deployment models to meet varying business continuity scenarios. These include Central Datacenter Deployment, where active and standby CIC servers, media servers, recording storage, database, and trunking are located in a single site; Dual Data Center Deployment, where the CIC pair is split for geo-redundancy and trunking and media servers may be split or centralized; and a DR option available with any deployment model that consists of a cold standby server pair, media servers, and supporting resources.

### 4.2.1 Deployment Option Overview

PureConnect is deployed in an active-warm standby model. The following illustrations cover a few of the possible deployment models for PureConnect.

One deployment model of PureConnect is the Central Datacenter Deployment model, where the CIC switchover pair, Interaction Media Servers, telephony circuits and equipment, and other resources reside in a single physical location. This model allows for full interaction recovery in the event of a switchover.

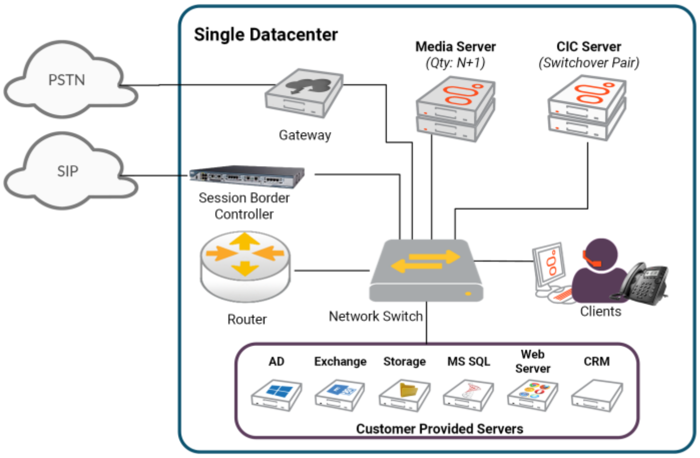

*Figure 5: PureConnect Single Datacenter Model*

Another deployment model of PureConnect separates the switchover pair across datacenters. Trunking remains centralized in the primary datacenter, and the secondary datacenter may or may not host redundant or additional telephony hardware and circuits.

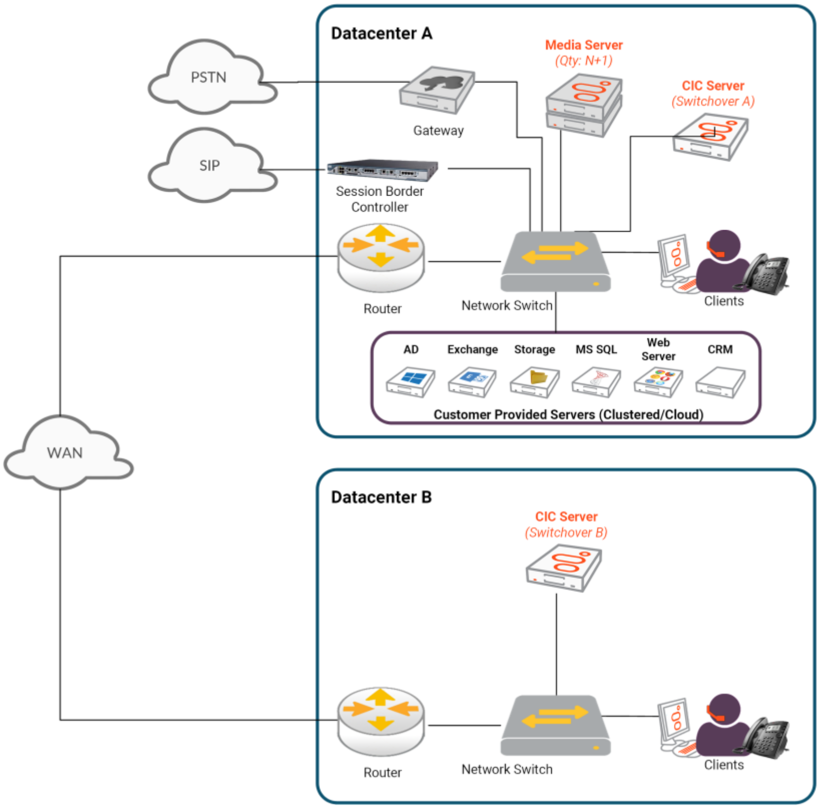

*Figure 6: PureConnect Geo-Redundant Datacenter Model*

A more robust geo-redundant deployment of PureConnect provides full redundancy of all CIC and supporting resources in two datacenters. Clustered database, Active Directory, Exchange, and Web servers are often combined with cloud CRM and SAN storage to provide always-on ancillary services. Telephony services are load balanced between sites or provide for circuit failover to reduce impact on calls. The PureConnect installation provides Interaction Media Servers and CIC server at each site to minimize impact of a datacenter outage.

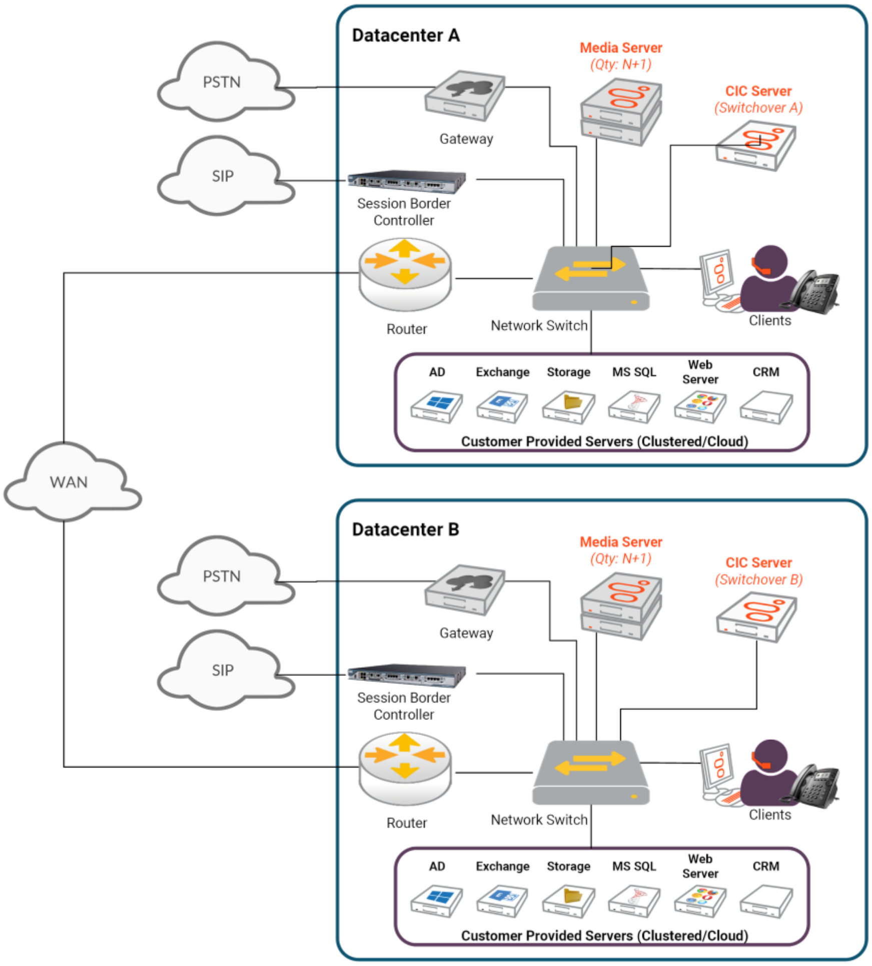

*Figure 7: PureConnect Replicated Datacenter Model*

One additional delivery model is available for assured business continuity in the event of a total outage of all the services in the above models. PureConnect Disaster Recovery (DR) provides a cold standby set of isolated PureConnect servers that can be manually, or script synchronized with production servers or kept in a known-good state for assured last-chance business continuity. Typical usage of this model provides essential communication and call handling, although it is capable of all services deployed in production.

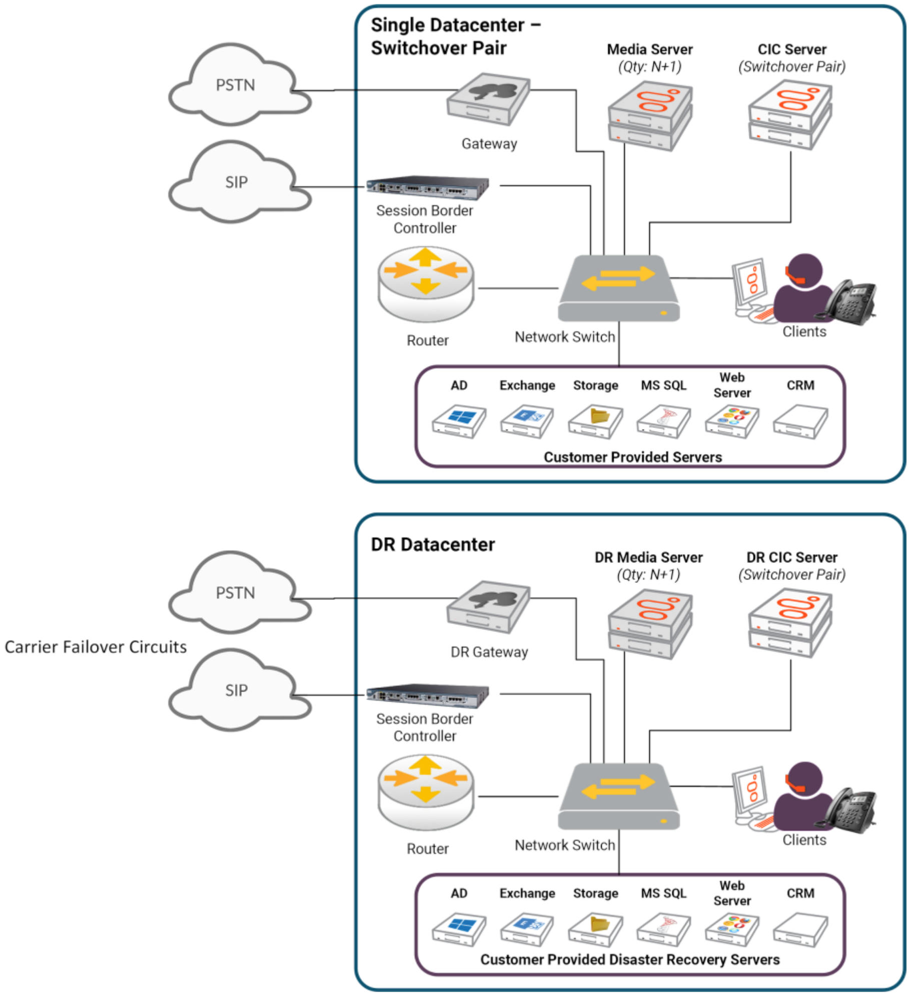

*Figure 8: Disaster Recovery deployment (any delivery model)*

### 4.2.2 High Availability

**Introduction**

Genesys recognizes the need to ensure that the PureConnect server functions in a highly reliable, fail-safe way to prevent unplanned down time. To accomplish this goal, you can configure 2 servers so that one performs the primary PureConnect functions and the other maintains a mirror image of the primary server. If the primary server fails or becomes disconnected, a **switchover** occurs, and the backup server takes over. In most cases, the switchover transition occurs quickly and does not disrupt communications between agents or users of the system. Replication occurs continuously between the two servers to ensure that in the event of a switchover, there is minimal data loss.

For more information about CIC security considerations that may affect your switchover configuration, see *PureConnect Security Features Technical Reference* in the PureConnect Documentation Library at http://help.genesys.com/cic. http://help.genesys.com/cic.

For more information on configuring the recovery of call interactions, see *Call Recovery Feature Technical Reference* in the PureConnect Documentation Library.  in the PureConnect Documentation Library.

**CIC Switchover system processes and architecture**

A switchover environment uses a pair of identical PureConnect servers:

• The *active server* processes all PureConnect interactions, such as phone calls, email, faxes, web chats, and voice mail.

• The *backup server* is a mirror image of the active server, duplicating its hardware and software, including the current configuration of PureConnect. The backup server regularly *monitors* the active server. It validates specific PureConnect subsystems and looks for the appropriate return signal from an attempted call operation. The backup server also monitors and dynamically copies any changes to the configuration of the active server. These changes include new user entries, line configuration changes, handler updates, or Interaction Attendant profiles. All these changes keep the Directory Services (DS) tree on the backup server identical to the DS tree on the active server. The backup server also monitors and copies any changed files.

**Note:** This document uses the terms *active server* and *backup server* to signify the switchover process as occurring between two *states* as opposed to two physical servers.

The backup server starts the switchover process immediately when it detects that the active server is not responding to TS (Telephony System) pings.

**Synchronizing the registry on both servers**

This section describes how the registry is synchronized on the active and backup servers.

Most of the key configuration data for PureConnect is stored in the PureConnect server registry. The configuration data is dynamic and regularly updated using Interaction Administrator and Interaction Designer, serving as an interface to AdminServer and DSServer. The backup server must maintain an identical image of the IC-related registry keys on the active server to be able to take over processing for a failed active server.

Each time the Switchover system starts on the backup server, it establishes a Notifier connection to the active server. The Switchover system then runs a recursive compare-and-update algorithm to determine whether there are differences between the DSServer registry structures between both servers. If necessary, the Switchover system updates the backup server to synchronize both servers using DS notifications.

**Note:** Genesys recommends that you back up the registry regularly.

Typically, when the active server handles calls, the Switchover system on the backup server monitors changes on the active server by listening to change notifications from DSServer. Whenever a change to the configuration or PureConnect registry key occurs on the active server, the server broadcasts a DS notification and replicates the change to the backup server.

When a server in the switchover pair is started, it checks to see if its companion is an active server. If it cannot contact the other server or verify that it is running, it goes into fail-safe mode and starts as the active server.

**Note:** After every switchover a manual reboot is required on the failed server to re-enter a running state.

### Synchronizing PureConnect directories on both servers

When key files are added, changed, or removed from the active server, that change must be reflected on the backup server automatically. This ensures that the backup server synchronizes with the active server.

When the PureConnect Switchover service is started on the backup server, it automatically starts monitoring specific default directories on the active server and mirrors all file operations to the backup server.

You can refine the list of mirrored directories by setting one or more server parameters in Interaction Administrator. For more information, see *Switchover Server parameters*. .

### Default mirrored directories

By default, the backup server automatically mirrors the published handler directory *only once,* on startup. This directory is listed in the Handler Path server parameter. For example, \I3\IC\Handlers.

Any time a handler is published on the active server, the replication process also publishes it on the backup server.

If a switchover event occurs, the published, active handlers are identical on both servers. However, handler files in subdirectories under the Handlers directory are not mirrored on the backup server by default.

Audio files that are played in handlers using the Play Audio File tool are not replicated. When the handlers are published on the backup server, the new .ivp files are created. If the WhitePages.txt file is updated, it is automatically copied to the backup server.

If the backup server is unavailable when a change occurs to one of these directories, it replicates all aspects of the active server when it becomes available.

The other default mirrored directories are the Resource Path and I3Tables Path directories, which are continuously mirrored.

• The Resource Path server parameter defines the resources directory. For example, \I3\IC\Resources. This location is where system and user prompts reside along with the system white pages files.

• The I3Tables Path server parameter defines the i3tables directory (where Interaction Administrator data tables are stored). For example, \I3|IC\Server\I3Tables.

### Custom and initial mirrored directories

You can further define the list of mirrored directories by setting the CustomMirrorDir and InitialMirrorDir server parameters in Interaction Administrator on the active server.

### CustomMirrorDir

The CustomMirrorDir server parameter specifies one or more directories on the active server that are mirrored on the backup server. Any time a file is added, removed, or modified in one of these directories, the change is mirrored in the corresponding directory on the backup server.

The CustomMirrorDir server parameter includes the following directories by default:

- +D:\I3\IC\Resources
- +D:\I3\IC\HostTools
- +D:\I3\IC\Server\LRA
- +D:\I3\IC\ClientSettings
- +C:\I3\IC\Flows

- +C:\I3\IC\TFTPRoot
- +C:\I3\IC\Provision
- +C:\I3\IC\Certificates\LinesAuthority
- +C:\I3\IC\Certificates\Email
- +C:\I3\IC\Server\I3RxDocs
- +C:\I3\IC\Server\Reports

**Note: The plus sign (+) before the path indicates that all subdirectories of that directory are mirrored.**

### The ping process

The backup server regularly monitors the active server using the following ping processes:

- TS (Telephony Systems) ping
- IP (Interaction Processor) ping

### TS ping

During the TS ping process, the Switchover subsystem on the backup server sends a request through the Notifier subsystem on the active server to the TS subsystem on the active server. It then waits for a response. If the Switchover subsystem does not receive a response, it pauses for a specified amount of time (by default, 1 second) and then retries. If the second attempt is also unsuccessful, the backup server immediately starts the switchover procedure.

If you want to change either of the default timeout values, you can create and set the following server parameters in Interaction Administrator on the active server:

- Switchover TS Timeout: Specifies the number of seconds that the Switchover system waits from the time the ping is sent until it is marked as a TS failure. The value should be between 5 and 60 seconds. The default is 10 seconds.

- Switchover TS Failure Retry Delay: Specifies the number of seconds the Switchover system waits, after marking a TS failure, before sending the second ping. A second failure causes the system to switch.

Note: Set this value to greater than 0 seconds. The default is 1 second.

- Switchover Max TS Failures: Specifies the number of TS ping failures that the Switchover system on the backup server tolerates before starting a switchover.

Note: Set this value to greater than 0. The default value is 2. (Note that the error count is reset each time the Switchover system successfully receives a response from TS on the primary server.)

Frequently, a TS ping failure indicates that both the network connection and the Notifier connection have been lost, and that the TsServer is either backed up with processing requests or is unresponsive. Usually, the TsServer log is required to diagnose the root cause of the switchover event. However, you should also scan the system event logs and the Switchover log to verify that a lost network connection did not cause the ping to fail.

### Switchover system architecture

Whenever the Switchover system starts on the backup server, it first establishes a Notifier connection to the active server. It then runs a recursive compare-and-update algorithm to determine whether there are differences between the DSServer registry structures on both servers. Finally, it updates the backup server, if necessary, to synchronize the servers using DS notifications.

Note: Genesys recommends that you back up the registry as a part of performing regular server backups. You can also back up the server's configuration using Interaction Migrator. For more information, see *Data Backup Technical Reference* and *Interaction Migrator Technical Reference* in the PureConnect Documentation Library. http://help.genesys.com/cic. http://help.genesys.com/cic.

**Note: The active server name is replicated and used for both servers to optimize the synchronization process. The site names for both the active server and the backup server appear in the DS tree and in Interaction Administrator. These names stay the same.**

### Switchover system components

The following table shows the Switchover system components.

| Component | Description |
|---|---|
| SwitchoverU.exe | This is a CIC server subsystem that monitors the mirror server and signals the virtual switch to switch servers. |
| SwitchoverCtrlU.exe | This is the CIC client GUI module that monitors the state of the Switchover system and provides a manual Switchover command. |

### Switchover states

When the primary and backup servers are on the same release and running in auto-switch mode, this condition is called the Active state. When the backup server is running in manual switch only mode because it is on a different release than the primary server, it is in the Upgrade state. The Upgrade state has two secondary states:

- The Upgrade Higher state is when the backup server is on a newer release than the primary server.
- The Upgrade Lower state is when the backup server is on an older release than the primary server.

Both these Upgrade states provide limited replication. While in the Upgrade Higher or Upgrade Lower state, the backup server will not automatically become the active server if the primary server experiences a failure. However, a manual switch can be performed which will manually transition the backup server into the primary state and demote the old primary server into a failed state.

### Switchover Subsystem Components

The switchover subsystem consists of the following components:

- Switchover State Machine
    - Maintains state information

  - Processes events from several sources
- Switchover Main
  - Handles common tasks
  - Prepares for entering various switchover modes
- System Monitor
  - Loads and starts the primary monitor and the module monitor
- Remote Dispatchers
  - Conduits for establishing, monitoring, and using notifier connections to the primary
  - Posts events on transactions of the connections such as establishment and loss
  - Provides notifications that are integral to the switchover state machine's ability to detect connection status in real time
- Module Monitor
  - Sends routine pings to the module being monitored on the primary
- Primary Monitor
  - Is the central component for the backup
  - Sends routine pings to the notifier on the primary as a means of determining its status
  - Determines the status of the network connections and the primary server
  - Stops and restarts the module monitor when detecting the network status
  - Provides methods for determining network connection status to Switchover Main and the switchover state machine
  - Acts as the single provider for network and connection analysis

**The following diagram illustrates the Switchover Subsystem components on the primary and backup installations.**

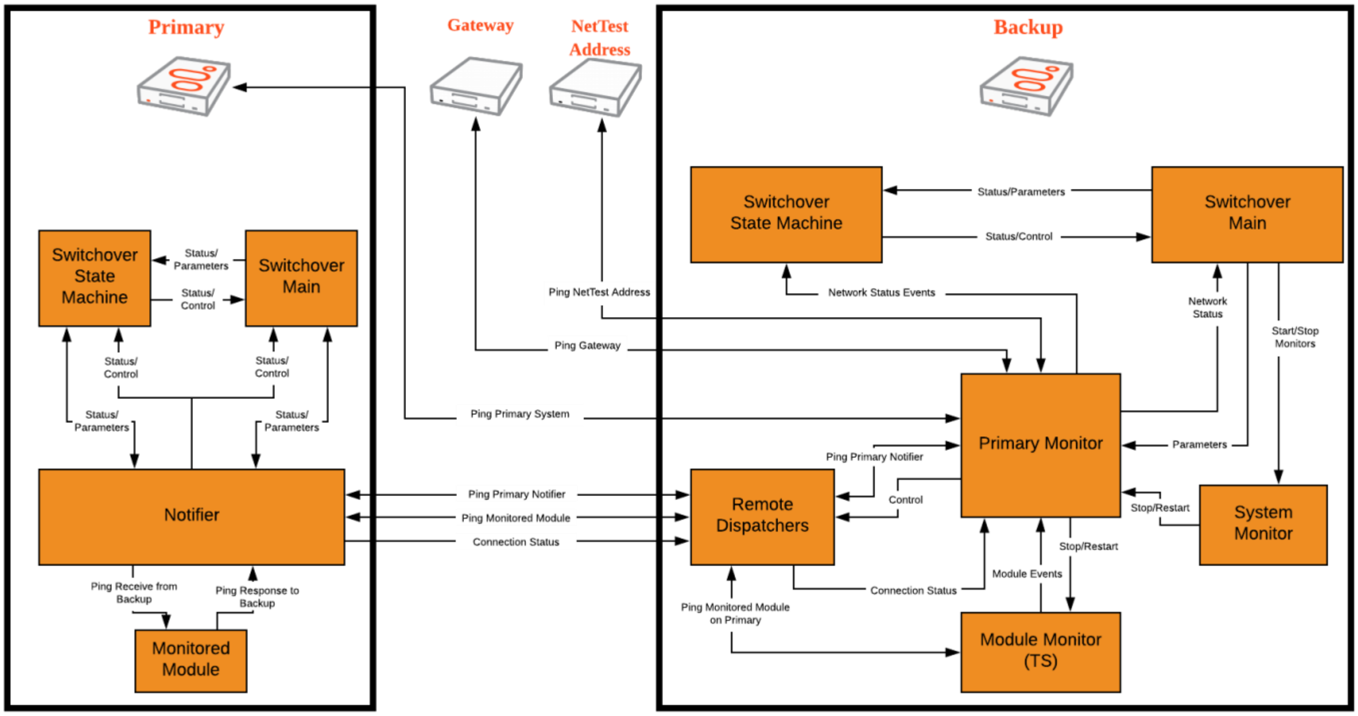

*Figure 9: The Switchover State Machine*

The switchover state machine provides the various modes of operation for switchover. There are several states that switchover routinely executes. The backup and reconnecting states are affected by the optimizations that were made in CIC 2015 R1 to improve connectivity management.

When switchover is running in the primary state, it operates as the primary. Likewise, when the switchover state machine runs the backup state, it operates as the backup.

When the switchover state machine runs as the backup, and connection or module issues are detected, the Switchover system transitions to the reconnecting state to take specific steps to reconnect to the primary. During reconnection, the Switchover system transitions between the backup and reconnecting state depending on the conditions encountered.

**Note: The transitions to the backup state during reconnection do not indicate that the backup server has returned to operating as the backup. Transitioning between the reconnecting and backup states is normal when the backup is attempting to re-establish communications and resume normal operations.**

The state machine on the backup will transition to the primary state and begin operations as the primary when any of the following occur:

1. The maximum number of reconnect attempts has been made.

2. The monitored module is determined to be down.

3. The primary is determined to be unreachable.

4. The reconnect interval has expired.

If reconnection succeeds or the monitored module is responding (in the case where reconnect was entered because the monitored module did not respond), the Switchover system will return to the backup state and performs a re-synchronization.

By default, the Switchover system will not attempt to determine if the primary is reachable when evaluating connection status. This is done to maintain legacy behavior that does not test whether the primary is reachable. You can set the following server parameters to specify the interval and the number times that the backup will try to reach the primary before deciding that it is unreachable.

- Switchover Unreachable Primary Ping Delay
- Switchover Unreachable Primary Ping Count

**System Monitors**

There are two types of monitors used by the backup:

- The primary monitor handles the bulk of the connection management and status detection.
- The module monitor maintains contact with the monitored module running on the primary.

Primary Monitor Purpose

The purpose of the primary monitor is to:

- Continually ping the notifier connections to the primary
- Process events from the module monitor
- Verify that the module monitor is running
- Provide diagnostics of the remote notifier connections
- Determine the network connection status of the backup
- Determine the network connection status to the primary

**Responsibilities**

The primary monitor continually monitors the main and auxiliary remote notifier connections to the primary by sending pings to the notifier on the primary. If any of the pings timeout or are otherwise not acknowledged, the primary monitor starts to diagnose the state of the network and connections. It posts events to the switchover state machine if there are issues with the connections or the network.

It also receives notifications from the module monitor that routinely pings the module being monitored on the primary. The module monitor will monitor the TsServer module or the IP module depending on the type of installation. The TsServer module is typically the module being monitored in a PureConnect Cloud environment.

**Note: In this document, the term *module* refers to either of the modules.**

The primary monitor tracks the time between events received from the module monitor. It does this to provide verification that the module monitor has not stopped for an unintended reason. When any event is received from the module monitor, the primary monitor resets the internal value that tracks the time of the last event. The state machine will poll the primary monitor at regular intervals to get the status of the module monitor. When the primary monitor receives this poll request from the state machine, it calculates the amount of time since the last module monitor event and determines if the duration is past the maximum allowed. The calculation of the maximum event interval is internal and based on the frequency of the pings sent by the module monitor to the primary. The primary monitor suspends tracking the intervals between module monitor events when reconnecting and resumes it when the connections are re-established.

If the module monitor indicates that the module is down, meaning that the module monitor exhausted its number of retries, the primary monitor will evaluate the network and connections.

When the primary monitor detects that there are issues with the network or connections, it will stop the module monitor to avoid receiving continual notifications from the module monitor and to prevent the module monitor from causing a switchover event before the primary monitor has been able to analyze the situation.

The primary monitor is only used on the backup.

**Startup**

The primary monitor is started by the system monitor (SysMonitor2), which passes the service parameters queried during switchovers initialization as a backup. The parameters include:

- Timeout, delays, and ping counts
- The name of the module being monitored
- The name of the primary
- A callback into the system monitor
- The list of NetTest addresses, if configured

During startup, the primary monitor queries the system to obtain the local gateway addresses of the backup. It uses standard Windows APIs to query information about all network adapters. Each adapter query returns the interface and gateway addresses. An adapter may have no gateway address, a single gateway address, or multiple gateway addresses. The primary monitor will store all gateway addresses returned and ping each one when determining network connectivity of the backup. Local host or empty addresses are excluded from the list.

The primary monitor will also attempt to get the address of the server that the primary is running on. It uses this address to test if the primary's server is reachable. If the address cannot be retrieved during startup, the primary monitor will try to get it when checking the network connectivity.

The following diagram shows the startup process.

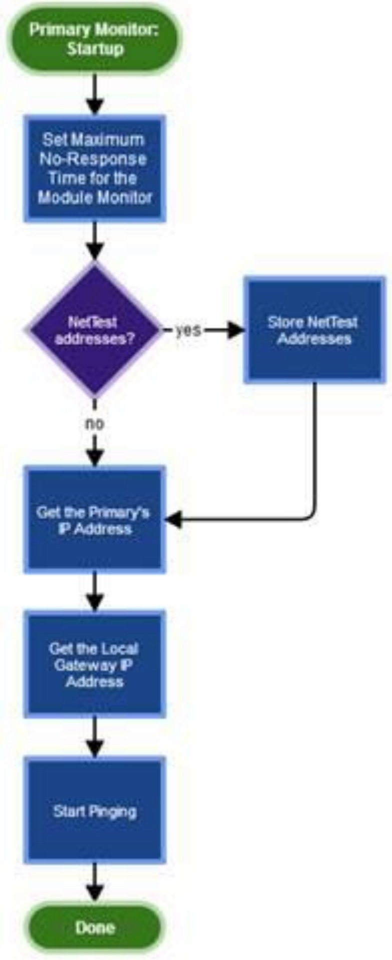
Primary Monitor: Startup
Set Maximum No-Response Time for the Module Monitor
NetTest addresses?
yes
Store NetTest Addresses
no
Get the Primary's IP Address
Get the Local Gateway IP Address
Start Pinging
Done

*Figure 10: How the Primary Monitor Determines the Status of the Network and the Connection*

When the primary monitor does not receive a ping reply from the notifier on the primary, it checks the following to detect the location of the problem:

- Status of the remote notifier connection
- Network status of the backup
- Network status of the primary

### Detecting Status of the Remote Notifier Connections

The primary monitor first examines both remote notifier connections. Determining the status of either connection is made by checking its connected state directly. Notifier on the backup maintains state information on each connection and the primary monitor (as well as any other switchover component) can check the status of any connection. The connected state value is used to evaluate if either connection is up or down.

If either connection is up, the primary monitor attempts to use them to check the status of notifier on the primary. If a closed connection is not re-established or if a connection was up and then going down, the primary monitor then examines the network.

### Detecting Network Connection Status of the Backup and Primary

To identify where there is a network issue, the primary monitor attempts the following steps:

1. First, the primary monitor pings NetTest addresses (if configured).

2. Next, the primary monitor pings the backup's local gateway (if enabled in the configuration).

3. Throughout the entire process, the primary monitor pings the server that the primary is running on (not any PureConnect products).

### Detecting Network Connection Status of the Backup and Primary

To identify where there is a network issue, the primary monitor attempts the following steps:

1. First, the primary monitor pings NetTest addresses (if configured).

2. Next, the primary monitor pings the backup's local gateway (if enabled in the configuration).

3. Throughout the entire process, the primary monitor pings the server that the primary is running on (not any CIC products).

The primary monitor pings the local gateway to determine if the backup is still connected to the network. There can be multiple gateways configured for each adapter on the backup. During startup, the primary monitor attempts to enumerate all the gateways for the adapters on the backup. The primary monitor stores all gateway addresses to use when pinging. Pinging the backup's gateway is optional; by default, it is disabled. To enable the gateway ping, set the **Switchover Disable Gateway Ping** server parameter to No or 0. To disable the gateway ping, set the parameter to Yes or 1. If you do not set the parameter, the gateway ping is disabled because some gateways block responses to pings (ICMP echoes) as a matter of security. If gateway pings are enabled and the primary monitor attempts to ping the primary's server but ping replies are blocked by the gateway, the primary monitor receives a false positive that the backup is not connected to the network. It is important to disable the primary monitor's gateway pinging if the gateway does not allow ping responses because there is no way for the primary monitor to know how the gateway is configured. However, gateway pinging is extremely valuable in determining the network connection state of the backup. If the backup's gateway responds to pings on the local network, it is strongly suggested that the backup be configured to enable gateway pinging. It is not unreasonable for the backup's gateway to respond to local pings since the backup should be on a private network.

The steps to check the network connections are taken in the following order based on the configuration of switchover on the backup. Once any step is taken, detection of the backup's network connection status is halted, and no further steps are taken.

The steps to check the network connections are taken in the following order based on the configuration of switchover on the backup. Once any step is taken, detection of the backup's network connection status is halted, and no further steps are taken. The steps are:

**Gateway pings are enabled**: The primary monitor determines the backup's network connection status based on the result of pinging each gateway address. The primary monitor stops pinging the gateway addresses once a response is received or all addresses have been pinged without a response. At this point, the primary monitor stops checking the network condition and sets the network status.

o Ping each gateway address until:

- A ping response is received from the gateway, the network connection status is set to Gateway Reached, and the backup's network connection is good.
- All gateway addresses have been pinged and no responses were received, the network connection status is set to Gateway Unreachable and the backup's network connection is not good.

o Network connection evaluation is stopped by the primary monitor.

**NetTest addresses are configured**: The primary monitor decides of the backup's network connection status based on the result of pinging each NetTest address. The primary monitor stops pinging the NetTest addresses once a response is received or all addresses have been pinged without a response. At this point, the primary monitor stops checking the network condition and sets the network status.

o Ping each NetTest address until:

- A ping response is received from a NetTest address, the network connection status is set to NetTest Reached, and the backup's network connection is good.
- All NetTest addresses have been pinged and no responses were received, the network connection status is set to NetTest Unreachable, and the backup's network connection is not good.

o Network connection evaluation is stopped by the primary monitor.

**Ping the address of the primary's server**: The primary monitor stores the address of the primary's server during startup. If it was not able to get the primary server's address, it attempts to do so in this step. This is a standard ICMP echo sent to the IP address of the primary and is independent of any CIC application. This ping is used by the primary monitor to detect if the server running the primary is reachable.

o If the primary server's address is not stored and the primary monitor is able to get it, ping the primary's server address and if:

- A ping response *is received* from the primary's server address, the network connection status is set to Primary Reached and the backup's network connection is good. This indicates that there is an issue with the primary (CIC).
- A ping response *is not received* from the primary's server address, the network connection status is set to Primary Unreachable and the backup's network connection is not good.

This indicates one or both of the following conditions:

- There is a connection issue between the backup's network and the primary's network (when they're on separate networks such as in the case of a WAN configuration).
- The primary's server is down.

When this happens, the primary monitor cannot determine if any of the primary's components have failed and will do one of two things:

- Wait for a period or indefinitely (depending on the configuration parameters) for the primary to become reachable again.
- Immediately switchover to become the primary.

**Network Status versus Connection Status**

Network and connection status are determined by the primary monitor, as described previously. They are distinct and defined, in terms of switchover, as:

• **Network Status**: The status of the physical network connection for the server running switchover in the backup state. It is independent of any switchover or CIC applications.

• **Connection Status**: The status of both notifier connections between the backup and the primary. It is reflective of the network status since it depends on the physical network connection.

Whenever the primary monitor attempts to assess whether communications can occur between the backup and primary, the physical network connection is the mitigating factor; without the network connection, the notifier connections are unavailable. If the network connection is available, then the connection status provides further information about the state of the notifier connections. The connection status is one of the following:

- Both connections are up.
- The main connection is up (auxiliary connection is down).
- The auxiliary connection is up (main connection is down).
- Both connections are down.

Detection of the connection status includes:

- Checking the network connections
- Checking the main and auxiliary connections
- Determining if the primary's server is reachable
- Checking for no-response condition from the monitored module

monitor's response status is evaluated. The primary reachable status is examined if there is an issue with the network connections to determine if there's a general network problem or just the path to the primary's server. Depending on the results of all these checks, an event may be posted to the state machine indicating a condition that requires further reconnect processing. The diagram below illustrates the logic that determines the overall connection status.

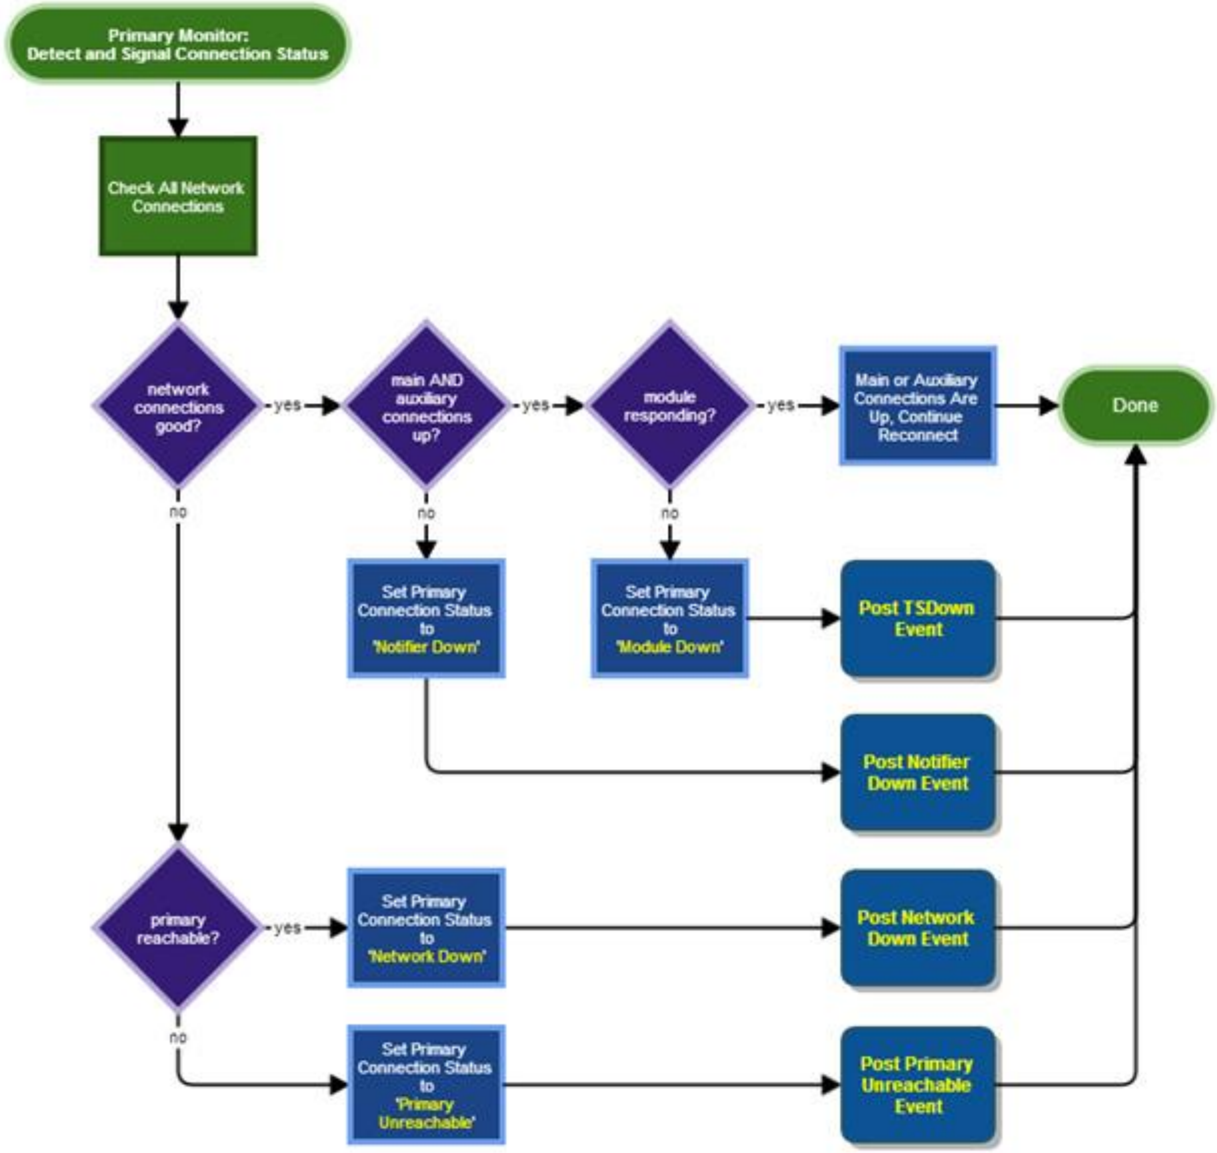
Primary Monitor:
Detect and Signal Connection Status
Check All Network Connections
network connections good?
yes
no
main AND auxiliary connections up?
yes
no
module responding?
yes
no
Main or Auxiliary Connections Are Up, Continue Reconnect
Done
Set Primary Connection Status to 'Notifier Down'
Set Primary Connection Status to 'Module Down'
Post TSDown Event
Post Notifier Down Event
primary reachable?
yes
no
Set Primary Connection Status to 'Network Down'
Post Network Down Event
Set Primary Connection Status to 'Primary Unreachable'
Post Primary Unreachable Event

*Figure 11: Primary Monitor: Detect and Signal Connection Status*

Note that the primary is reachable even if the network connections are not good. This is because the definition of a good network connection includes whether the primary's server is reachable. If it is not, the network connection's status is primary unreachable, and this is not a good condition.

This processing occurs when the primary monitor detects that it did not receive a ping from the notifier on the primary within the allotted amount of time or it detected a connection loss. When that happens, the primary monitor uses this logic to make a preliminary decision about what to do next. You can see in the backup state diagram how these events are handled. In the case of a primary unreachable determination, the state machine transitions to the reconnect state and tries to restore connections with the primary.

### Handling Module Monitor Notifications

The primary monitor receives all notifications from the monitor module. If a notification indicates an issue with the module, network, or connections, then the primary monitor posts an event to the state machine, if required. The primary monitor can then make one of the following determinations:

- **Network NOT OK:**
    - **Primary Reachable:** The server running the primary can be reached, so the remote notifier s down.
    - **Primary Unreachable:** The server running the primary cannot be pinged, so the primary is unreachable.
- **Network OK:**
    - **Both Connections Down**: The main and auxiliary remote connections are down, and the primary can be reached so the remote notifier is down.
    - **Both Connections Up:**
        - The Module is Down: The module cannot be pinged, so the module is down.
        - The Module is Up: The module monitor indicated that the module was down, but subsequent pings were answered so the module is up and there is no error.

The following flowchart illustrates this decision-making process.

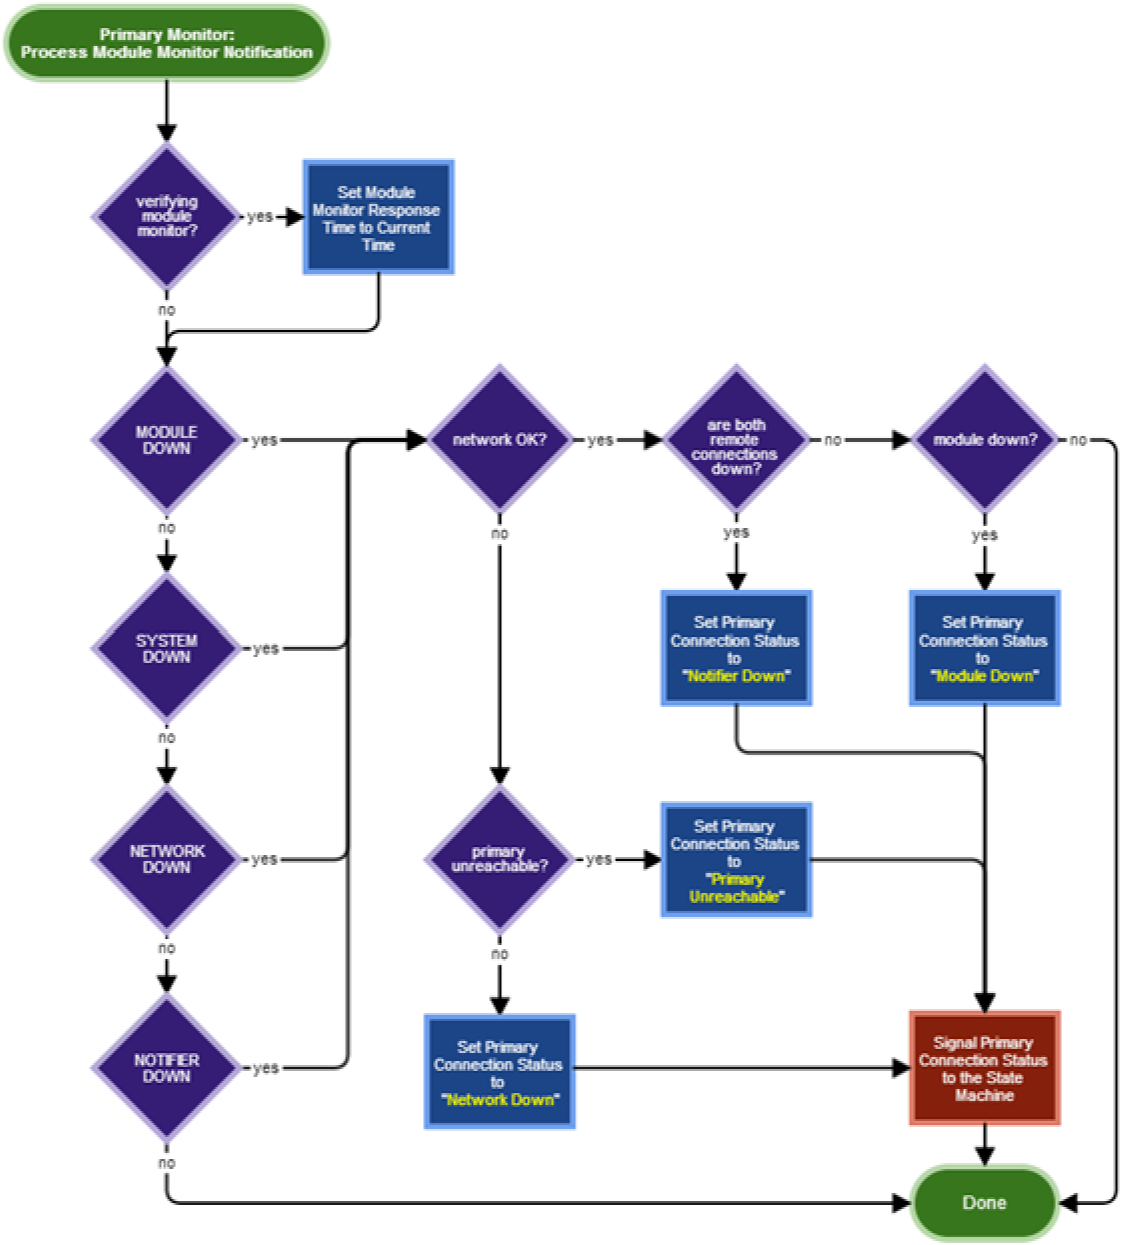
Primary Monitor:
Process Module Monitor Notification
verifying module monitor?
yes
Set Module Monitor Response Time to Current Time
no
MODULE DOWN
yes
no
SYSTEM DOWN
yes
no
NETWORK DOWN
yes
no
NOTIFIER DOWN
yes
no
network OK?
yes
no
are both remote connections down?
no
yes
module down?
no
yes
Set Primary Connection Status to "Notifier Down"
Set Primary Connection Status to "Module Down"
primary unreachable?
yes
no
Set Primary Connection Status to "Primary Unreachable"
Set Primary Connection Status to "Network Down"
Signal Primary Connection Status to the State Machine
Done

*Figure 12: Primary Monitor: Process Module Monitor Notification*

### Sending Ping Requests

The primary monitor sends ping requests to the notifier on the primary at the specified ping delay interval. The primary monitor schedules a callback whose timeout is the ping delay interval. During that callback, a request is sent to the remote notifier which acts like a network ping except that the response is from the remote notifier itself.

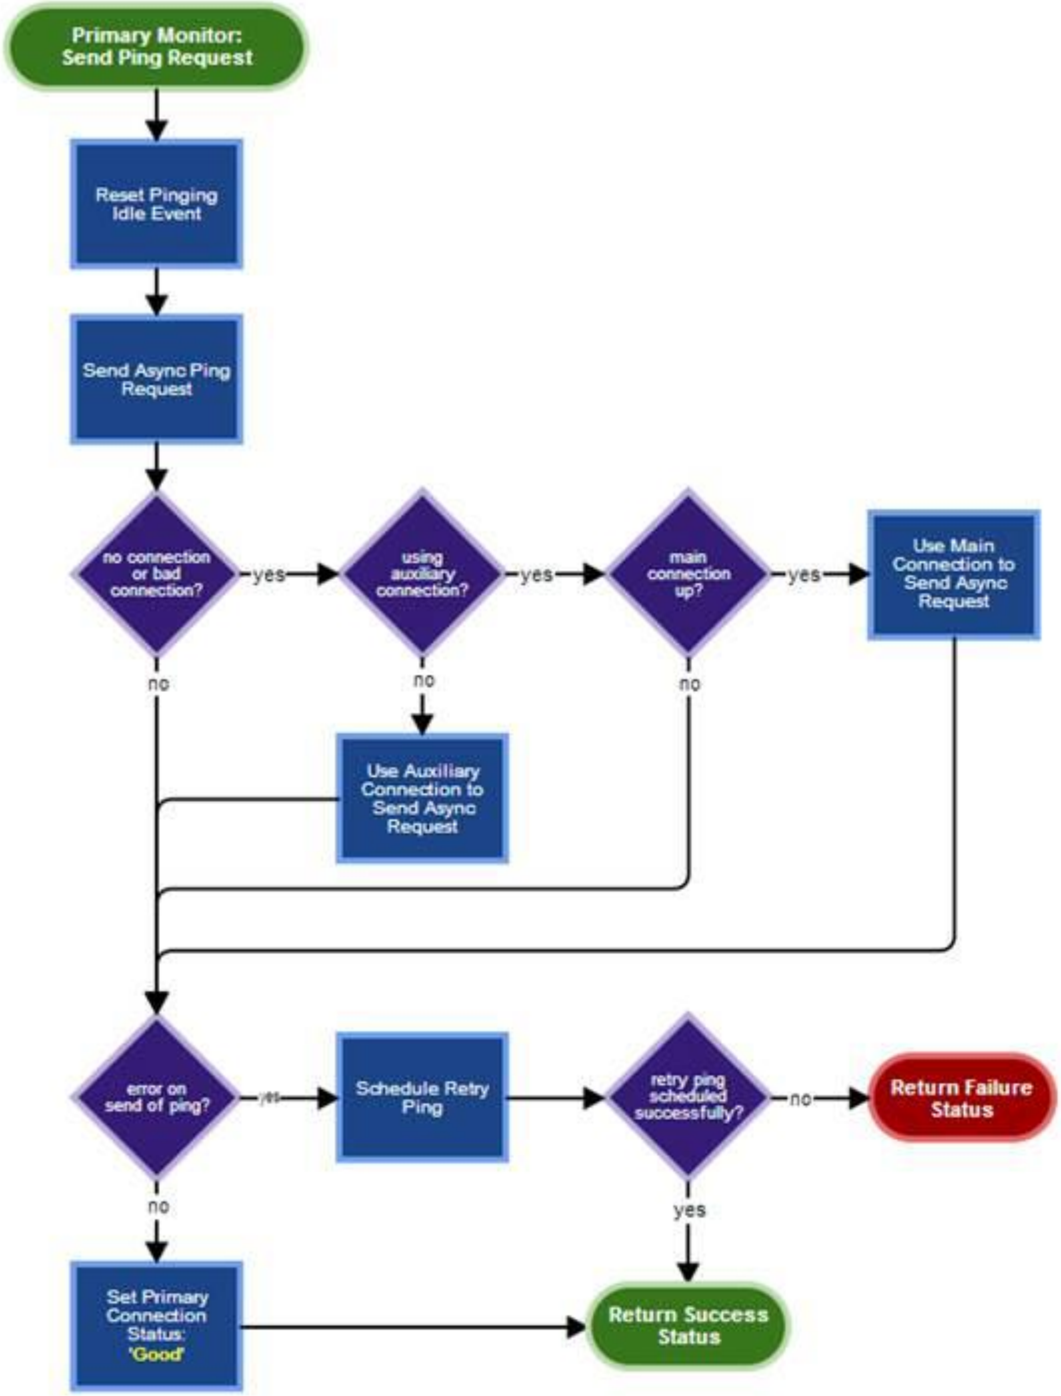
Primary Monitor: Send Ping Request
Reset Pinging Idle Event
Send Async Ping Request
no connection or bad connection?
yes
using auxiliary connection?
yes
main connection up?
yes
Use Main Connection to Send Async Request
no
no
no
Use Auxiliary Connection to Send Async Request
error on send of ping?
yes
Schedule Retry Ping
retry ping scheduled successfully?
no
Return Failure Status
no
yes
Set Primary Connection Status: 'Good'
Return Success Status

*Figure 13: Primary Monitor: Send Ping Request*

## Successful Ping Response

A successful ping response occurs when the notifier on the primary sends a response within the timeout specified in the ping request. An asynchronous ping request callback of the primary monitor is called when the response is received successfully and before the timeout. The retry count is reset when a successful ping response is retrieved, and the next ping is scheduled.

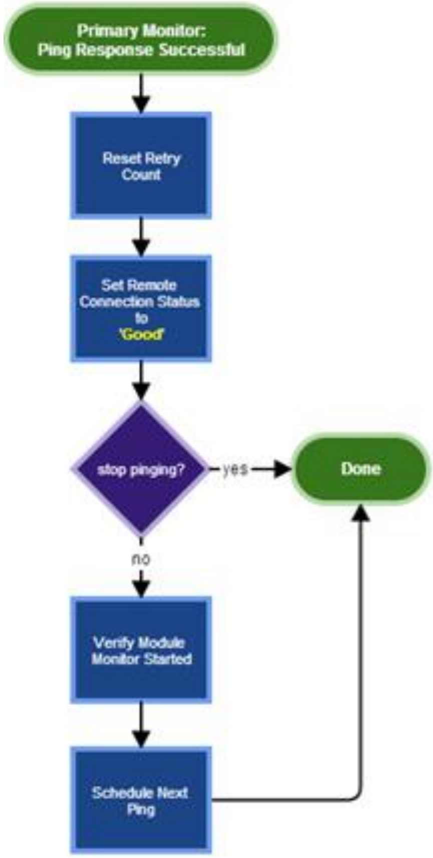

*Figure 14: Primary Monitor: Ping Response Successful*

After a ping has been received successfully, another ping request callback is scheduled to send the next ping after the required delay.

The primary monitor method that schedules the next ping also acts as a receiver for requests to asynchronously restart the module monitor. There are conditions where the primary monitor cannot synchronously restart the module monitor. In those cases, a request is scheduled, and this method will identify those requests and restart the module monitor. The method performs no other processing after the restart.

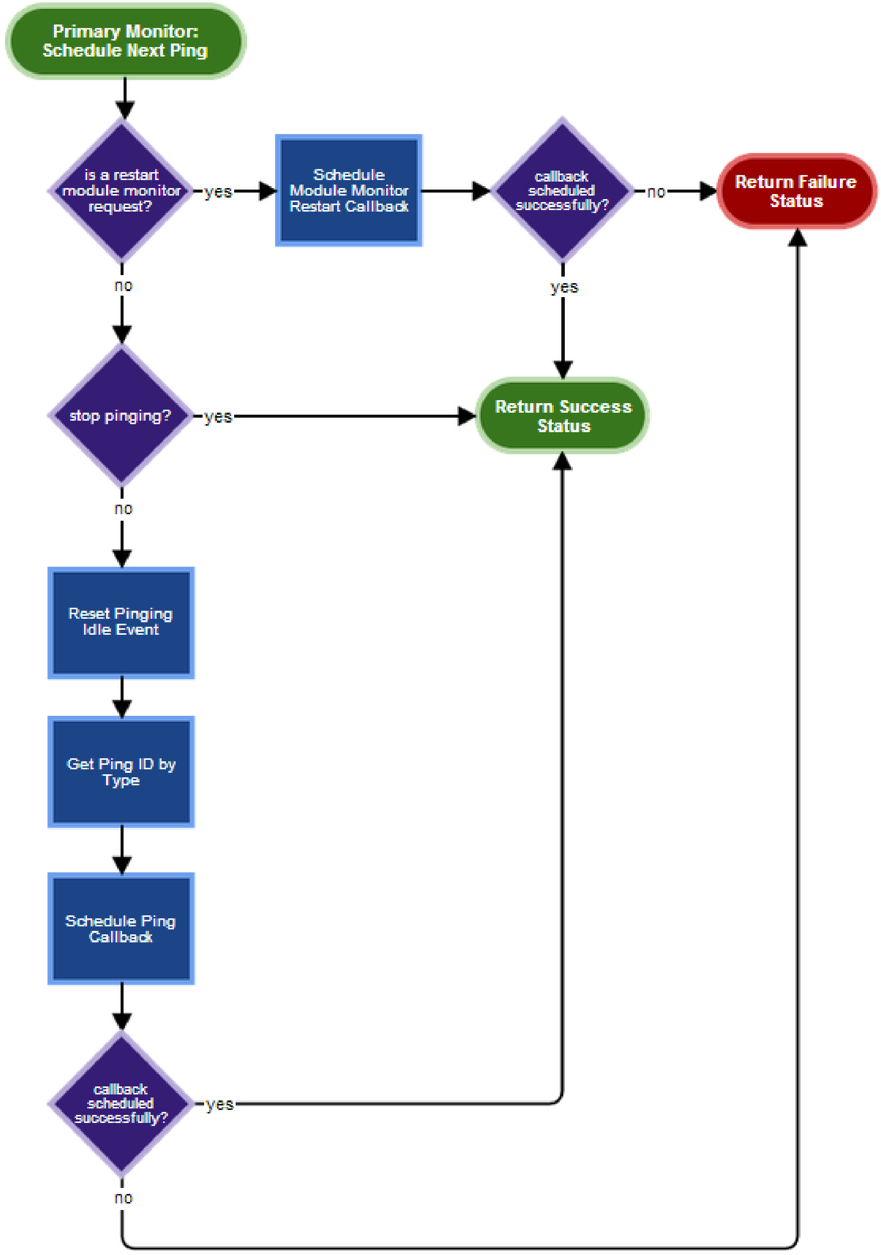
Primary Monitor: Schedule Next Ping
is a restart module monitor request?
yes
Schedule Module Monitor Restart Callback
callback scheduled successfully?
no
Return Failure Status
yes
no
stop pinging?
yes
Return Success Status
no
Reset Pinging Idle Event
Get Ping ID by Type
Schedule Ping Callback
callback scheduled successfully?
yes
no

*Figure 15: Primary Monitor: Schedule Next Ping*

### Non-Successful Ping Responses

Other responses are non-successful and appropriate action is taken if applicable. There are callbacks for each of the unsuccessful responses:

• **Timeout**: This response represents the potential loss of a remote notifier connection or an issue with the notifier on the primary. It triggers the retry logic that begins the process of sending retry pings.

• **Rejected**: Like the timeout response. The diagram below represents both the timeout and rejected response handling.

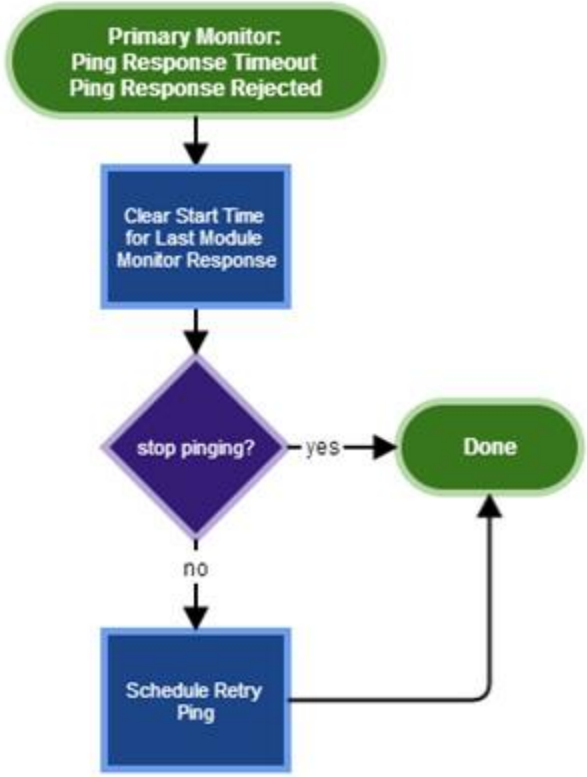

*Figure 16: Primary Monitor: Ping Response Timeout Ping Response Rejected*

**Connection Loss**: This response represents the potential loss of a connection to the notifier on the primary or an issue with the notifier on the primary.

o If the main connection is up, this response could indicate that the auxiliary connection went down or there was a momentary drop of the main connection, but it was recovered. In this case, a standard ping is sent.

o If the main connection is down, the primary monitor treats this as a potential reconnect condition and sends a retry ping.

The following diagram illustrates how this response is handled.

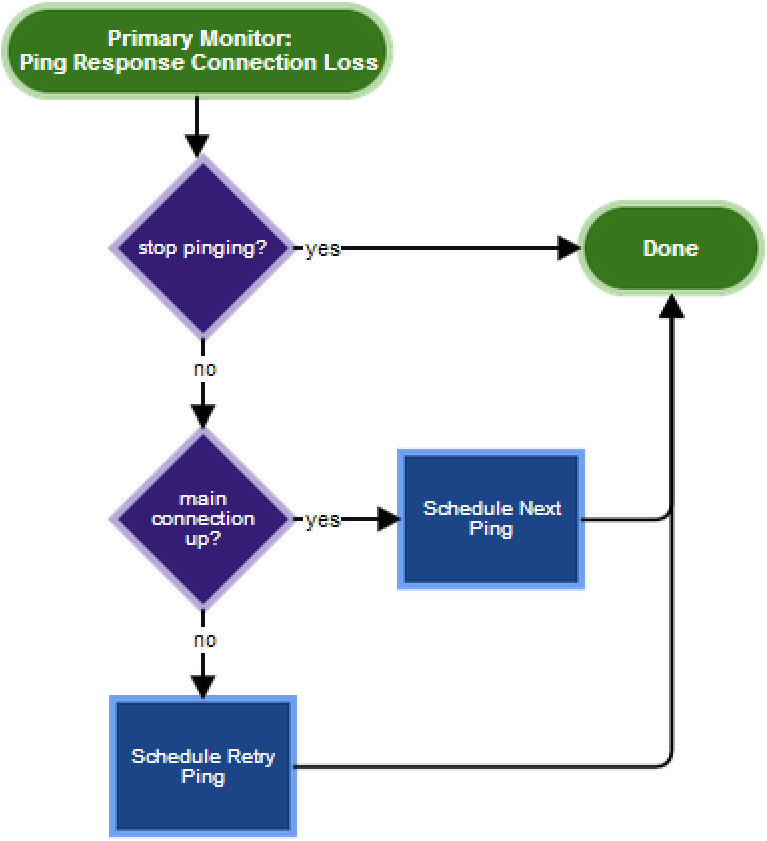

*Figure 17: Primary Monitor: Ping Response Connection Loss*

- **Cancelled**: No operations are performed upon receipt of this response because the primary monitor will cancel any outstanding pings when it is stopping (which generates this response). A retry ping is sent for an unsuccessful response. If the retry count has reached the maximum, no retry ping is sent and error processing starts to diagnose the state of the local network and remote connections to the primary.

### Recovery of ACD email, chat, and callback interactions during switchover

Starting in IC 4.0 SU 3, the Interaction Recovery Service subsystem supports recovery of email, chat, and callback interactions as described in the following sections. Before SU 3, those interactions, including any work in progress, were lost during a switchover. Email interactions were requeued as new interactions, chat users had to reconnect and be requeued to an agent, and callbacks were lost with no method for recovery.

The Interaction Recovery Service subsystem replicates the creation of interactions and processes them to take the appropriate action based on the interaction state. For example, interactions in a **c**onnected state on a user queue do not change and ACD automatically reprocesses interactions in an offering state on a workgroup queue.

Other subsystems use data from the Interaction Recovery Service subsystem to create and maintain mirrored interactions on the backup server.

**Enable the Interaction Recovery Service**

To enable the Interaction Recovery Service, open Interaction Administrator and then open the Server Parameters container. Add the parameter for each type of interaction that you want:

- Mail Interaction Recovery Enabled
- Chat Interaction Recovery Enabled
- Callback Interaction Recovery Enabled
- SMS Interaction Recovery Enabled

**Recovery of email interactions**

Previously following a switchover, email interactions were requeued as new interactions. Starting with IC 4.0 SU 3, the Interaction Recovery Service subsystem recovers email interactions with the following caveats:

- Email interactions generated by agents lose HTML formatting and convert to plain text. Any attachments are lost.
- Following the switchover, email interactions get new IDs.
- When a switchover occurs during the synchronization process, the interaction state could be lost.

The following table provides configuration information for the recovery of email interactions.

| | |
|---|---|
| **Enabled by default?** | No. Requires a server parameter. |
| **Server parameter** | **Mail Interaction Recovery Enabled**<br>This parameter must be set to 1. For more information, see *Optional Switchover Server parameters*. |
| **Requires a subsystem restart?** | No. When the backup server starts, or the parameter is enabled, the Interaction Recovery Service subsystem performs a full synchronization of email interactions with the active server. |

**Recovery of chat interactions**

Previously, following a switchover, chat users had to reconnect and wait for an agent. Starting with IC 4.0 SU 3, the Interaction Recovery Service subsystem recovers chat interactions with a transition that is almost seamless. Once switchover occurs, the WebProcessor subsystem prepares the mirrored interactions.

**Note the following:**

• Automatically generated status messages are lost after switchover. However, the switchover process replicates all the other texts that are exchanged during the chat session.

• Files that were transferred during the chat session are not available after switchover.

• The switchover behavior of the CIC clients is the same as it was in switchovers before IC 4.0 SU 3.

• All chat responses typed during the switchover are maintained. They are not lost.

The following table provides configuration information for the recovery of chat interactions.

| | |
|---|---|
| **Enabled by default?** | No. Requires a server parameter. |
| **Server parameter** | **Chat Interaction Recovery Enabled**<br>Set this parameter to 1. |
| **Requires a subsystem restart?** | Yes. Restart the WebProcessor subsystem on the active server and reboot the backup server. After the parameter is enabled and the backup server has restarted, the Interaction Recovery Service subsystem performs a full synchronization of chat interactions with the active server. |

### Recovery of callback interactions

Previously, following a switchover, callback interactions were lost and not recoverable. Starting with IC 4.0 SU 3, the Interaction Recovery Service subsystem recovers callback interactions. Note the following:

- • Recovery processing of callback interactions is like the recovery processing of chat interactions.
- • The callback window temporarily disappears while the CIC client tries to reconnect.

The following table provides configuration information for the recovery of email interactions.

| | |
|---|---|
| **Enabled by default?** | No. Requires a server parameter. |
| **Server parameter** | **Callback Interaction Recovery Enabled**<br>This parameter must be set to 1 |

| | |
|---|---|
| **Requires a subsystem restart?** | Yes. Restart the WebProcessor subsystem on the active server and reboot the backup server. After the parameter is enabled and the backup server has restarted, the Interaction Recovery Service subsystem performs a full synchronization of chat interactions with the active server. |

### Recovery of SMS interactions

Previously, following a switchover, SMS users had to reconnect and wait for an agent. Starting with CIC 2016 R4, the Interaction Recovery Service subsystem recovers SMS interactions with a transition that is almost seamless. Once switchover occurs, the WebProcessor subsystem prepares the mirrored interactions.

**Note the following:**

- Automatically generated status messages are lost after switchover. However, the switchover process replicates all the other texts that are exchanged during the SMS session.
- The switchover behavior of the CIC clients is the same as it was in switchovers before CIC 2016 R4.
- All SMS responses typed during the switchover are maintained. They are not lost.

The following table provides configuration information for the recovery of SMS interactions.

| | |
|---|---|
| **Enabled by default?** | No. Requires a server parameter. |
| **Server parameter** | **SMS Interaction Recovery Enabled**<br>Set this parameter to 1. |
| **Requires a subsystem restart?** | Yes. Restart the WebProcessor subsystem on the active server and reboot the backup server. After the parameter is enabled and the backup server has restarted, the Interaction Recovery Service subsystem performs a full synchronization of SMS interactions with the active server. |

### A note about the loss of duration information for interactions

When you enable the Mail Interaction Recovery Enabled, Chat Interaction Recovery Enabled, or Callback Interaction Recovery Enabled server parameter, every time an interaction is created on the active server, a corresponding shadow interaction is created on the backup server. If a switchover occurs, the interaction on the active server is automatically discarded. The shadow becomes the new active interaction. When you restore the switchover pair, the new backup automatically makes a new shadow interaction.

You may notice a discrepancy in the duration of the interaction if you look in IC Business Manager or at the Time in Queue value for the interaction.

The duration is the time that the interaction has existed, regardless of whether it is a shadow or active interaction. Because the second shadow was created after the first active and first shadow, the duration that the second shadow interaction has existed will be different than the original interaction and the first shadow. This is a known limitation of the switchover process.

**Example of how loss of duration appears**

1. Active1 and Shadow1 are created at 5 AM.

2. Switchover happens at 6 AM. Shadow1 becomes Active2.

3. Switchover pair is restored 7 AM. Shadow2 is created based on Active2.

4. Another switchover happens at 8 AM. Shadow2 becomes Active3.

5. Active3 duration is listed as 1 hour. Customer believes it should be 3 hours because the email originally came into the system 3 hours ago.

**Recovery of statistical data**

During normal processing, the primary server computes and caches statistical data which is saved into the database/PMQ periodically. The backup server does not receive any event, so it does not have any statistical data. As a result, when a switchover occurs the primary server and the backup server have different statistics.

To address this, you can enable the Interaction Recovery Service. When the Interaction Recovery Service runs, the Statserver on the backup server receives the same event notifications as the primary server. Because the same event notifications are stored on both servers, the statistical data on both servers is valid.

**Note: The Interaction Recovery Service supports only emails, chats, and callbacks. Therefore, statistical data can be generated only for these types of interactions. Statistical data for calls is not supported.**

**Location of log data**

During normal operations, the primary server (Server A) sends its log data to the database. The backup server (Server B) sends its log data to a CSV file. When a switchover occurs, Server B becomes the primary server, it flushes log data collected before switchover to CSV file, then sends new log data to the database.

The timespans of the database records reflect the switchover. For example, suppose the standard duration of a database log record is 1800 seconds (30 minutes). However, a switchover occurs 15 minutes into the logging process. Instead of a single log record for 1800 seconds, there would be 2 log records, each for 900 seconds.

The CSV log files are stored in I3\IC\CSVLogs. The CSV file is automatically overwritten every week. To keep a backup copy of the file, copy it to a different location.

If disk space is a concern, you can stop the backup server from sending its log data to the CSV log file. To do this, add the StatServer_DisableQPSLoggingOnBackup parameter and set its value to Yes. You must restart the Statserver on the backup server for this parameter to take effect.

**Note: When a switchover occurs, the log data from Server B is automatically sent to the CSV log file, regardless of the setting of the StatServer_DisableQPSLoggingOnBackup parameter. Also, if the Interaction Recover Service is not enabled, no data is logged to the CSV file because the Statserver does not receive any notifications.**

### Tracker Server logging

Beginning in SU 5, the Tracker Server actively monitors email interactions, chat interactions, and callback interactions on both the primary server and the backup server to capture the interaction data. This interaction data is logged in the Interaction Summary table and the Interaction Segment Detail table. Because the Tracker Server logging is active on both the primary server and the backup server, CIC captures the full history of these interaction types.

### The Tracker Server's processing on the primary server

On the primary server, interactions can be in one of the following states when the switchover happens.

- Pre-connected
- Connected
- Disconnected but waiting for the interaction to be de-allocated from the system.

During a *controlled* switchover, the Tracker Server on the primary server persists via PMQ (PMQ stands for Persistent Message Queuing, a delivery process which protects against the loss of database updates caused by network or hardware failures). the connected and disconnected interactions on the primary server. Pre-connected interactions are persisted on the primary server.

During an *uncontrolled* switchover, for example, when there is an operating system-level error, no interactions persist on the primary server.

### The Tracker Server's processing on the backup server

The Tracker Server actively processes all email interactions, chat interactions, and callback interactions on the backup server, regardless of their states. However, the Tracker Server does not log the data from the secondary server into the database. When the backup server becomes primary after a switchover, the history information captured about the non-call interactions is logged to the database.

### Identifying recovered interactions in the database

The primary server and the backup server share the InteractionIdKey in the database. When spanned interaction data is logged to the database, the sequence number column (seqno) is incremented to 1 in the following tables: Interaction Summary, Interaction Segment Detail, and Interaction Wrapup. Interaction records with the sequence number of 1 are the recovered interactions that have been shared between the primary server and the backup server.

### Switchover from the user's point of view

If a CIC server fails, control automatically and immediately switches to a backup server that has maintained a mirror-image of the primary CIC server's configuration.

The switchover time on a midsize system is typically less than 30 seconds and causes little disruption to the phone service.

Agents and users can observe the following effects when switchover occurs:

- Voice calls are switched automatically when a CIC client connects to the backup server.
- Client may temporarily gray out and become unresponsive when reconnecting on the new primary server.  Agent will regain control of the client once reconnected.  This may take up to 30 seconds.
- Email interactions in queues are requeued.
- Beginning with IC 4.0 SU 3, continuity of chats, emails, and callback objects is maintained.
- Beginning with IC 4.0 SU 4, continuity of calls is maintained. For complete details on the Call Recovery feature, see *Call Recovery Feature Technical Reference*.
- Emails connected to an agent remain connected.

**Note:** Emails that contain HTML are converted to plain text, and markup is lost if there is a switchover. Any attachment that has been added to a draft email is not synchronized and must be added again to the draft email after the switchover occurs.

- Chat interactions remain connected; the web user does not notice the switchover happening. Intercom chats remain connected and all the history information about the chat is maintained.
- If there is a file transfer during a switch, the transfer can fail. In that case, the user must reinitiate the transfer.
- Callback requests remain in the system and continue to be routed after the switchover has been completed.

For more detailed information about PureConnect High Availability features and configuration, see *Automated Switchover System Technical Reference* document in the PureConnect Documentation Library at http://help.genesys.com/cic. http://help.genesys.com/cic.

### 4.2.3 Disaster Recovery

At the foundation, PureConnect can be deployed at a single site with 100% redundancy for all services and features. This feature is accommodated via switchover services and maintains real-time replication of application configuration data. This configuration is an “Active/Passive” redundancy. In the event of a catastrophic failure of the active server, the backup (passive) server immediately begins providing service to the environment on the next inbound call.

Other design topologies can be enabled to extend business continuity (BC) in the event of a single site failure. This is accommodated easily by deploying the backup (passive) server of the switchover pair at a separate location. This configuration enables organizations to retain BC during periods where a single facility may be isolated from the WAN or completely down.

Many organization also deploy the core components within a data center that will not be impacted by office downtime, and in cases where a contact center location is inaccessible, agents are able to log in from a backup site or from home using either the standard or web-based client. The system supports calls being directed to any phone number, so agents can have queued calls terminate on their mobile or home phones.

With implementation of the options outlined above, PureConnect can meet the most demanding BC/DR requirements.

**Please refer to the previous section 4.2.2 for more detailed information relating to Automated Switchover and High availability.**

For more detailed information about PureConnect High Availability features and configuration, see the *Automated Switchover System Technical Reference* document in the PureConnect Documentation Library at http://help.genesys.com/cic. http://help.genesys.com/cic.

## 4.3 Database

**PureConnect supports Microsoft SQL and Oracle as standard databases:**

Microsoft SQL Server:

Genesys recommends a dedicated SQL Server computer for PureConnect database storage. If a SQL Server is already installed on the same network as the PureConnect Server, that server can be used for PureConnect database storage.

Database file size allocations

The following table shows the "rule of thumb" used to calculate initial data file size values for a Reporting, Tracker, and/or Recorder database. During a new installation, PureConnect database configuration in PureConnect Setup Assistant uses these default values if no values are entered.

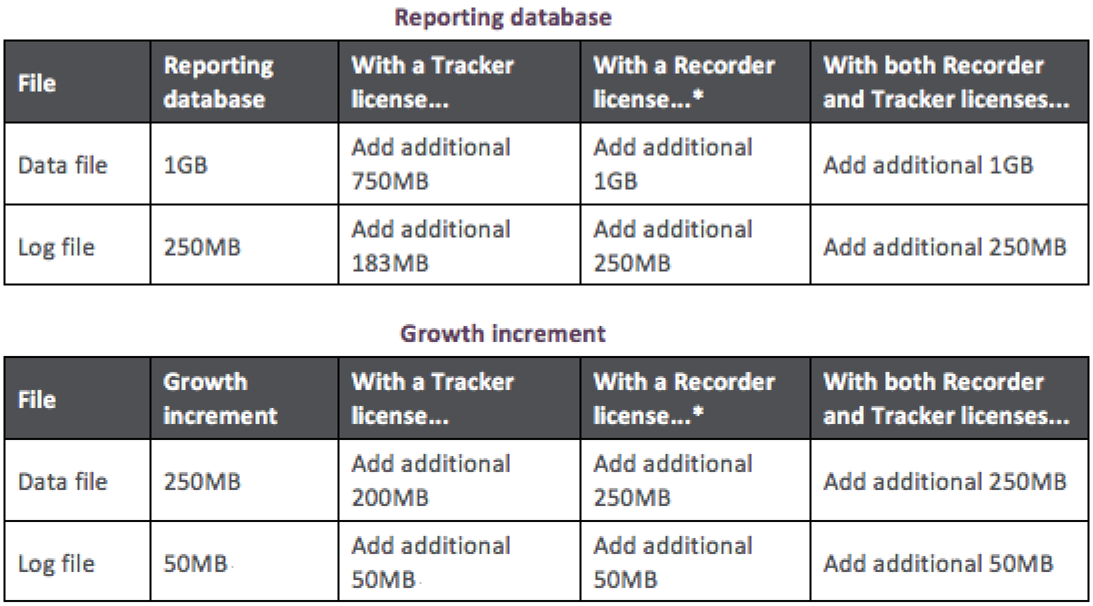

Reporting database

| File | Reporting database | With a Tracker license... | With a Recorder license...* | With both Recorder and Tracker licenses... |
|---|---|---|---|---|
| Data file | 1GB | Add additional 750MB | Add additional 1GB | Add additional 1GB |
| Log file | 250MB | Add additional 183MB | Add additional 250MB | Add additional 250MB |

Growth increment

| File | Growth increment | With a Tracker license... | With a Recorder license...* | With both Recorder and Tracker licenses... |
|---|---|---|---|---|
| Data file | 250MB | Add additional 200MB | Add additional 250MB | Add additional 250MB |
| Log file | 50MB | Add additional 50MB | Add additional 50MB | Add additional 50MB |

*Figure 18: Reporting Database Growth Table*

Oracle Server:

Genesys recommends using a dedicated Oracle server computer for PureConnect database storage. If an Oracle server is already installed on the same network as the PureConnect Server, that server can be used for PureConnect database storage.

## IC Database Configuration and Maintenance for SQL Server

### *CIC Database Configuration Technical Summary*

### I/O Configuration

The configuration of the I/O system is extremely important for the performance of the database. CIC can generate thousands of records per second, and all those records must be written to the database as quickly as possible. A poorly configured I/O system can cause a queuing problem, meaning that the database is not able to keep up with all the records CIC is generating. Some recommended web sites are listed below. (Several sites state that they are for SQL Server 2005, but many of the recommendations are valid for SQL Server 2008 R2.)

- Physical Database Storage Design
- SQL Server I/O Best Practices
- Formatting drives for maximum SQL Server performance

Basic principles when considering the drives to use include:

1. When using traditional hard drives, get the highest RPM available, which is 15,000 RPM.
2. Do not get the largest drives available. It is much better to have many multiple smaller drives than a few large drives. Having a larger number of drives allows the I/O to be spread out over more devices. Drives smaller than 100 GB are ideal, though they may be hard to find. Stay away from drives that are 500 GB and larger.
3. If the storage system is a Storage Area Network (SAN), it is vital to work with the SAN administrators to set up the SAN for database performance. There are many articles available on the web that describe best practices. Here are some very useful websites that discuss using a SAN with SQL Server:
    - SAN Storage Best Practices for SQL Server
    - SAN Multipathing Part 1: What are Paths?
    - SAN Multipathing Part 2: What Multipathing Does
4. Solid State Drives (SDD) can achieve remarkable I/O rates, though they are very expensive. Configure the server with a minimum of four (4) physically separate data storage arrays, and allocate them as follows:
    a. Array1 - Operating system.
        i. Use RAID 1 or RAID 10.
    b. Array2 - SQL (user) database data files.
        i. Use RAID 10.
    c. Array3 - SQL (user) database transaction log files.
        i. Use RAID 1 or RAID 10.
    d. Array4 - SQL tempdb data file and transaction log file.
        i. Use RAID 0, RAID 1, or RAID 10.
5. RAID 5 is not recommended, because of the high number of writes that PureConnect generates. For databases that have extremely high volume and through put, consider placing indexes and tables onto separate arrays. In addition, potentially large tables such as the tracker tables IntxSegment and Intx_Participant may benefit from having their own filegroups.

### Database Size

1. Configure the tempdb data file between 500 and 1000mb. If Interaction Tracker and/or Interaction Recorder are used, add additional 500 to 800mb for each of them. The two links below have good

recommendations for setting up temp DB. Generally, SQL Server should have one tempdb file for each CPU core, e.g. a machine with two dual core CPUs would have four tempdb files. It is important that each file is sized the same and the files are preallocated, so that tempdb does not have to expand frequently. For machines with many cores, a maximum of eight tempdb files may be configured.

   a. Optimizing tempdb Performance
   b. Properly Sizing the SQL Server TempDB Database

2. Configure the tempdb transaction log file between 300 and 500mb. If you have Interaction Tracker and/or Interaction Recorder licenses as well, add additional 250 to 500mb for each of them.
3. Use the database space planning spreadsheet, IC_DB_SpacePlanning.xls (available in the CIC .iso file and in the CIC Documentation Library in the CIC Resource Center) to calculate, and then configure, the size of the IC database. The computation should include the number of years of data planned to be kept in the database. Allow sufficient and constant free space in both the database and transaction log for rebuilding indexes. (For a 2300mb database, 1000 to 1400mb of free space in both the database and transaction log are needed to rebuild indexes that are 50% fragmented.)
4. Allow both the database and transaction logs to grow automatically (as a safeguard) but use a suitably sized growth increment.

### Database Options Configuration

Leave the defaults in place, except for ‘optimize for ad hoc workloads’. sp_configure 'show advanced options',1; reconfigure; go sp_configure 'optimize for ad hoc workloads',1; reconfigure; go

### Database Maintenance

Perform database integrity checks and/or index reorganization in a scheduled job that is separate from the daily backup maintenance plan. Regular index reorganization is vital for peak system performance. Use sys.dm_db_index_physical_stats to identify which indexes need to be rebuilt or reorganized. We do not recommend setting up a regular job to rebuild all indexes; only those indexes that have reached a predetermined threshold of both size and fragmentation should be rebuilt or reorganized.

See Microsoft’s online documentation about rebuilding indexes.

### Hardware

#### CPU

When selecting a CPU for the server, select one with a large L2 cache. This is especially important for multiple-processor servers. Select at least a 1MB L2 cache for one or two CPUs. Four or more CPUs should have at a least 2MB L2 cache in each CPU. The greater the L2 cache, the greater the server's CPU performance because it reduces the amount of wait time experienced by the CPU when reading and writing data to main memory.

Simple, single table queries and updates, along with query joins on small tables take minimal CPU processing power. On the other hand, large joins, aggregations, and sorting of large result sets use a high level of CPU processing power. Keep this in mind when choosing the hardware configuration for SQL Server.

### Memory

In most cases, the more physical RAM SQL Server has the greater SQL Server's performance. If possible, purchase enough RAM to hold the largest database table in memory. If such a purchase is not possible during the initial setup, leave room for adding more RAM later. We strongly recommend running the 64-bit version of SQL Server with a minimum of 4 GB RAM. Systems with high performance demands should be prepared to use up to 128 GB RAM or more. To take advantage of SQL Server's ability to run parallel queries, plan on investing in more RAM. Parallel queries use much more RAM than non-parallel queries.

### AWE

If SQL Server is running on a 32-bit Windows box, consider using Address Windowing Extensions (AWE) to increase the amount of memory available to SQL Server. Normally, 32-bit CPUs can only support up to 4GB of RAM because of limited address space. SQL Server supports AWE to bypass this limitation.

Note: AWE will be removed in the next version of SQL Server. SQL Server 2008 R2 is the last version that will support AWE. We strongly recommend using the 64-bit version instead of the 32-bit version of SQL Server.

AWE support is not turned on automatically. To enable AWE support, change the "awe enabled" advanced option from 0 to 1.

### Drives

Avoid locating read-intensive and write-intensive activity on the same drive or array. For example, do not locate an OLTP and an OLAP database or heavily random and sequential activity on the same physical device. Whenever a drive or array must change back and forth between activities, efficiency is lost. NTFS-formatted partitions should not exceed 80% of their capacity. For example, a 20GB drive should never hold more than 16GB. NTFS needs room to work, and when capacity exceeds 80%, NTFS becomes less efficient and I/O suffers. Consider creating a system alert to indicate when an array exceeds 80% of capacity so that immediate action can be taken to correct the problem.

### I/O and RAID

Use hardware-based RAID rather than software-based RAID because the latter can't offload the work to a separate processor, making it much slower than a hardware-based RAID solution.

Do not store the operating system, application software, or databases on single disk drives because they do not afford any fault tolerance. Instead, always choose a RAID array made up of three or more physical drives that offers fault tolerance. Common fault tolerant RAID configurations include RAID Level 1 (mirroring or duplexing) and RAID Level 10 (also called 1+0, which includes both striping without parity and mirroring). Non-fault tolerant RAID configurations include RAID 0 which is simple disk striping. RAID 0 offers excellent performance and can be used for the TEMP tablespace. Each of these RAID levels offers different performance levels. Ideally, if the budget allows, chose RAID Level 10, which offers both high-speed and fault tolerance.

### I/O controller

Select the best I/O controller possible. Top-notch controllers offload much of the I/O work onto its own local CPU, freeing up CPU time on the server to do other tasks. For the ultimate in I/O controllers, consider a fiber channel connection instead of a SCSI connection. The controller should have the largest amount of cache RAM possible, with a minimum of 128mb of cache RAM. Generally, the greater the RAM cache on the controller, the higher the performance of the overall I/O, because data can be read ahead and stored in the cache, even if the data is not currently requested by Oracle. The data Oracle wants next from the array will likely be in the cache, speeding up data access.

Do not put DAT, DLT, CD-ROM, scanners, or other non-hard disk devices on the same I/O controllers that connect to the hard disk arrays. In addition, do not put hard disks on the same I/O controller if they have different speeds. Putting devices with different speeds on the same I/O controller slows the faster devices. Always put slower devices on their own I/O controller. For maximum I/O throughput, assign each type of major I/O activity (database, log files, tempdb, etc.) to its own separate RAID controller and dedicated RAID array.

### OLTP vs. OLAP

If the budget does not allow for the ideal number of disk controllers and hard disks to maximize the server's I/O performance, remember that optimal OLTP I/O is achieved by increasing disk reads and writes. The best way to do this is to add more hard disks to the array(s) that hold the database files and/or transaction logs. Adding more disks helps OLTP-based applications more than increasing the number or speed of disk controllers would, because OLTP-based applications tend to be limited by the number of transfer operations (read/writes) rather than bandwidth. However, for OLAP-based applications, adding more and faster disk controllers to the array is generally a better way to boost I/O than increasing the number of disk drives, because OLAP applications tend to be more limited by bandwidth than by read/write operations. Adding faster or more disk controllers increases the bandwidth and helps to remove any bottlenecks.

### I/O Testing

There are several utilities available to test I/O subsystem performance. We recommend SQLIO.

### Networking

If SQL Server is not connected to a switch (as recommended for best performance), try the following suggestions for boosting network performance. If Ethernet is running, reduce the number of computers on the same segment as SQL Server, and if possible, reduce the length of the physical segment SQL Server is connected to. Do not cascade hubs. If Token-Ring is running, reduce the physical length of the ring and increase the number of active stations on the ring, which reduces token travel time.

### Network Protocols

For best performance, SQL Server should be running on a dedicated server. Limit the number of network protocols installed on the server, because unnecessary network protocols increase overhead on the server and send out unnecessary network traffic. For the best overall performance, only install TCP/IP on the server.

### Routers

While not always possible (especially for WANs and Internet connections), try to avoid a router between SQL Server clients and SQL Server. In particular, avoid routers between two or more SQL Servers that need to communicate with each other. Routers are often a bottleneck for network traffic and can affect SQL Server client/server performance. If SQL Server must communicate over a router, ensure that the router has been properly tuned for maximum performance.

### Network Cards

SQL Server should have a minimum of one 100Mbs network card, and perhaps two. Two cards can be used to increase network throughput and to offer redundancy. In addition, the network card(s) should be connected to full-duplex switched ports for best performance. Be sure that the network card(s) in the server are set to the same duplex level (half or full) and speed as the switched port they are connected to (assuming they are connected to a switch and not a hub). If there is a mismatch, the server may still be able to connect to the network, but network performance can be significantly impaired. Do not rely on network cards or switches that are supposed to auto-sense duplex or speed settings, because they often do not work correctly. Manually set the duplex and speed for the card from the operating system, and if necessary, manually make the same changes to the switch.

Windows allows network cards to save energy by going to sleep when they are not used. If any network card on a production server has a power management feature, ensure that the power savings feature is off. Otherwise, unexpected results, such as a network card that fails to wake up, or intermittent performance problems, may occur.

Check to see if the network card has a power management feature by viewing the Properties sheet for the network card's driver. View the Power Management tab on the Properties sheet, to verify the settings.

### Software

The network libraries chosen during SQL Server installation can affect the speed of communications between the server and its clients. Of the three key network libraries, TCP/IP is the fastest and Multi-Protocol is the slowest. Due to the speed advantage, use TCP/IP on both the servers and clients. Do not install unused network libraries on the server, because they will contribute unnecessary overhead.

### Service Packs

We recommend staying current with database service packs and maintenance. In most cases, you will want to install the latest SQL Server service packs.

### Services

Do not install unnecessary SQL Server services, such as Microsoft Search, OLAP, or English Query, as they only add additional overhead to your server. If applications do not use the Microsoft Distributed Transaction Coordinator (MS DTC), turn this service off by setting it to manual using the Services icon in the Control Panel. Leaving this service on adds unnecessary overhead to your server.

### File Structure and distribution

Place the database files (.mdf, .ndf) and transaction log files (.ldf) for all production databases on separate arrays to isolate potentially conflicting reads and writes. This means that the server will have at least two physical RAID arrays, one to store the database files, and a separate one to store the transaction log files. The operating system can be stored on a mirrored set of drives. Use separate storage areas for tempdb and the operating system. Since tempdb is rebuilt each time SQL Server is started, this storage does not have to be fault tolerant (i.e., you could use RAID 0). The physical location for the master, msdb, and model databases is not as critical as user databases because they are not used excessively in production environments.

### MDF

For database files (.mdf), the best performance is gained by storing them on RAID 10 arrays. Each RAID array should have as many physical disks in the array as the controller will support, with the fastest RPM available. This allows reads and writes to be performed simultaneously on each physical drive in the array, significantly boosting disk I/O.

### LDF

For database log files (.ldf), the best performance is often gained by storing them on a RAID 1 (mirrored or duplexed) array. This assumes that there is only a single log file on the RAID 1 array. If there is only a single log file on the RAID 1 array, the file can be written to sequentially, speeding up log writes. But if there are multiple log files (from multiple databases) sharing the same RAID 1 array, then there is little or no advantage of using a RAID 1 array. This is because although writing to a log is done sequentially, multiple log files on the same array means that the array will no longer be able to write sequentially, but will have to write randomly, negating much of the benefits of a RAID 1 array.

### Note

You can put each database log on its own separate RAID 1 array. Another option is to put the log on a RAID 10 array. While this is expensive, it will provide optimum performance.

### Multiple file Strategy

If your database is very large and very busy, multiple files can be used to increase performance. One example of using multiple files would be a single table with 10 million rows that is heavily queried. If the table is in a single file, such as a single database file, then SQL Server would only use one thread to perform a sequential read of the rows in the table. But if the table were divided into three physical files (all part of the same filegroup), then SQL Server would use three threads (one per physical file) to sequentially read the table, which potentially could be much faster. In addition, if each file were on its own separate disk or disk array, the performance would even be greater.

Essentially, the more separate physical files that a large table is divided into, the greater the potential performance. Of course, there is a point where the additional threads aren't of much use when you max out the server's I/O. But up until you do max out the I/O, additional threads (and files) should increase performance.

### Tempdb

If SQL Server's tempdb database is heavily used by your application(s), then locate it on an array of its own (such as RAID 0, RAID 1 or RAID 10). If SQL Server has multi-gigabyte databases, create a new file for tempdb that is at least 1 GB in size. Having multiple tembdb files can also improve performance, taking care to make sure that each file is the same size.

### SQL Server Database Settings

### AutoGrowth

Every time a database file or transaction log grows automatically, it takes up a little extra CPU and I/O time. Minimize how often automatic growth occurs by sizing the database and transaction logs as accurately as possible to their "final" size. This recommendation is particularly important for transaction logs, because the more often that SQL Server must increase the size of a transaction log, the moretransaction log virtual files that must be created and maintained by SQL Server. A transaction virtual file is used by SQL Server to internally divide and manage the physical transaction log file.

In SQL Server, database and log files can be set to grow automatically. The default growth amount is 10%. This automatic growth number may or may not be ideal. If the database is growing automatically often (such as daily or several times a week), change the growth percentage to a larger number, such as 20% or 30%. Each time the database must be increased, SQL Server will suffer a small performance hit. By increasing the amount, the database grows each time, the less often it will have to grow.

If your database is very large, 10GB or larger, you may want to use a fixed growth amount instead of a percentage growth amount. This is because a percentage growth amount can be large on a large database. For example, a 10% growth rate on a 10GB database means that when the database grows, it will increase by 1GB. If the percentage growth is too large, change the settings to use a fixed growth size.

Size the database properly, to ensure that the database is not subject to frequent growth. In addition, set the growth increment to a reasonably sufficient value, to prevent the database from growing in numerous, small increments. In this configuration, the Auto grow feature is used as a safeguard only to prevent the database from stopping if it unexpectedly runs out of space. The table below shows typical growth increments for databases of the indicated size.

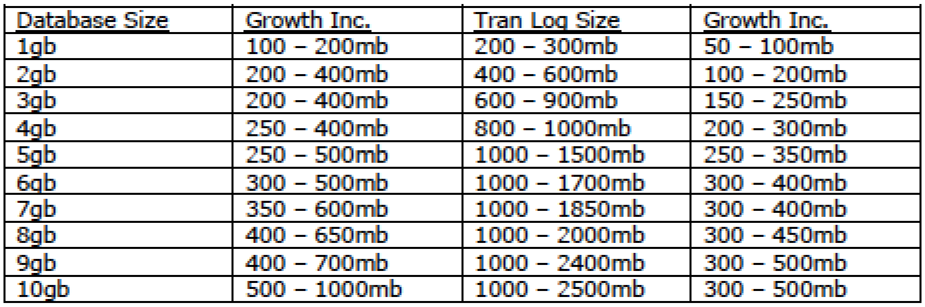

| Database Size | Growth Inc. | Tran Log Size | Growth Inc. |
|---|---|---|---|
| 1gb | 100 – 200mb | 200 – 300mb | 50 – 100mb |
| 2gb | 200 – 400mb | 400 – 600mb | 100 – 200mb |
| 3gb | 200 – 400mb | 600 – 900mb | 150 – 250mb |
| 4gb | 250 – 400mb | 800 – 1000mb | 200 – 300mb |
| 5gb | 250 – 500mb | 1000 – 1500mb | 250 – 350mb |
| 6gb | 300 – 500mb | 1000 – 1700mb | 300 – 400mb |
| 7gb | 350 – 600mb | 1000 – 1850mb | 300 – 400mb |
| 8gb | 400 – 650mb | 1000 – 2000mb | 300 – 450mb |
| 9gb | 400 – 700mb | 1000 – 2400mb | 300 – 500mb |
| 10gb | 500 – 1000mb | 1000 – 2500mb | 300 – 500mb |

*Table 4: Database and Log Growth Chart*

Additionally, the use of a file defragmentation tool is recommended. Be aware that these tools are very resource intensive, so you should only run them when the server is not servicing production requests.

**ODBC**

Do not use ODBC connection pooling and temporary stored procedures at the same time, or SQL Server will experience a performance hit. When a DSN is used to make a connection from your application to SQL Server, the MDAC driver, by default, converts any dynamic Transact-SQL from the application to temporary stored procedures in SQL Server. The theory behind this is that if the application resends the same Transact-SQL to SQL Server more than once, then it will save the SQL Server overhead of additional parsing and compilation. The recommended configuration consists of turning the convert T-SQL to temporary stored procedure feature off. This feature is configurable from the ODBC Database Wizard when creating or modifying a DSN.

Connection pooling is another option that can be configured using the ODBC Database Wizard when creating or modifying a DSN. It is also on by default, and it pools database connections from the application, which allows connections to be reused, which in turn reduces the overhead of making and breaking database connections. The recommended configuration consists of turning the connection pooling feature on.

Note that pooling improves performance more than temporary stored procedures. The problem is that if both options are on, which is often the case in DSNs, SQL Server can take a performance hit. Here's what can happen. When dynamic Transact-SQL is converted into a temporary stored procedure by the MDAC driver, the temporary stored procedure is stored in the tempdb database. When connection pooling is not enabled, and the connection between the client application and SQL Server is ended, any temporary stored procedures created during the connection are deleted. But, when connection pooling is enabled, things work differently. When a database connection is ended by the client application, it is not ended at SQL Server. SQL Server still thinks the connection is still open, even though the client application does not. This means the temporary stored procedures created during the connection are not deleted. With a busy client application that often starts and stops database connections, the tempdb database can fill up with temporary stored procedures, putting unnecessary overhead on SQL Server.

### Database Maintenance Plans

Create a Database Maintenance Plan to maintain databases. Database integrity options include "Check database integrity" and "Include indexes", which test the data and index page allocations in the databases for any errors. These tests are resource intensive and impair the performance of the server during the tests. These tests should only be run during off hours.

The database maintenance plan screen also has an option called "Perform these tests before backing up the database or transaction log". This is not a good choice to make from a performance standpoint. If chosen, every time the Maintenance Plan is used to perform a database or transaction log backup--without exception--the integrity tests are automatically run. If the databases are large, and/or if frequent database or transaction log backups occur, running these tests this often can significantly degrade the server's performance. The recommended configuration consists of creating separate scheduled SQL tasks (in SQL Agent) to perform database/log backups and database integrity checks. It is imperative to run backups, DBCC CHECKDB, and index rebuilds or reorgs regularly! It is equally imperative that indexes rebuild, and reorgs are done selectively. All indexes should not be rebuilt or reorged during the same job!

### Third-party (Application) Access to Database

Be wary of allowing users to directly access the databases (especially OLTP databases) with third-party database access tools such as Microsoft Excel or Access. Many of these tools can wreak havoc with database performance. Here are some reasons why:

- Often these users aren't experienced with these tools and create overly complex queries that eat up server resources. At the other extreme, their queries may not be complex enough (such as lacking effective WHERE clauses) and return thousands, if not millions, of unnecessary rows of data.
- This reporting activity can often lock rows, pages, or tables, creating user contention for data and reducing database performance.
- These tools are often file-based. This means that even if an effective query is written, the query is not performed at the server. Instead, the entire table (or multiple tables in the case of joins) must be returned to the client software where the query is performed. This leads to excessive server activity and can play havoc on your network.

If users must be allowed access to the data, limit hits on the production OLTP databases by pointing them to a "reporting" server that is replicated or in the form of a datamart or data warehouse For tools to help performance tune the operating system, see Sysinternals. The site has tools to defrag the server's swap file, among many others. And best of all, most are free.

For more detailed information see *CIC Database Configuration and Maintenance For SQL Server Technical Reference* document in the PureConnect Documentation Library at http://help.genesys.com/cic. document in the PureConnect Documentation Library at http://help.genesys.com/cic.

### IC Database Configuration and Maintenance for Oracle

### CIC Database Configuration Technical Summary

### I/O Configuration

The configuration of the I/O system is extremely important for the performance of the database. IC can generate thousands of records per second, and all those records must be written to the database as quickly as possible. A poorly configured I/O system can cause a queuing problem, meaning that the database is not able to keep up with all the records IC is generating.

Basic principles when considering the drives to use include:

1. When using traditional hard drives, get the highest RPM available, which is 15,000 RPM.
2. Do not get the largest drives available. It is much better to have many multiple smaller drives than a few large drives. Having a larger number of drives allows the I/O to be spread out over more devices. Drives smaller than 100 GB are ideal, though they may be hard to find. Stay away from drives that are 500 GB and larger.
3. If the storage system is a Storage Area Network (SAN), it is vital to work with the SAN administrators to set up the SAN for database performance. In addition to the previous link, Tips for Oracle database design on SAN has a short list of items to consider.
   Solid State Drives (SSD) can achieve remarkable I/O rates, though they are very expensive.
   Oracle Automatic Storage Management may be considered.

Configure the server with a minimum of four (4) physically separate data storage arrays, and allocate them as follows:

   a. Array1 - Operating system.
      i. Use RAID 1 or RAID 10.
   b. Array2 – Application tablespaces (one for tables and one for indexes are required).
      i. Use RAID 10.
   c. Array3 – REDO tablespace.
      i. Use RAID 1 or RAID 10.
   d. Array4 – TEMP tablespace.
      i. Use RAID 0, RAID 1, or RAID 10.e

4. RAID 5 is not recommended because of the high number of writes that IC generates. For databases with extremely high volume and through put, consider placing indexes and tables onto separate arrays. In addition, potentially large tables such as the tracker tables IntxSegment and Intx_Participant may benefit from having their own tablespaces.

### Database Size

1. Use the database space planning spreadsheet, IC_DB_SpacePLanning.xls (available in the CIC .iso file and in the CIC Documentation Library in the CIC Resource Center) to calculate, and then configure, the size of the CIC database. The computation should include the number of years of data planned to be kept in the database. Allow sufficient and constant free space in the application, TEMP and REDO tablespaces.

2. Allow tablespaces to grow automatically (as a safeguard) but use a suitably sized growth increment.

### Database Maintenance

Refer to the Oracle Database Administrator's Guide regarding database maintenance and tuning.

### Hardware

### CPU

When selecting a CPU for the server, select one with a large L2 cache. This is especially important for multiple-processor servers. Select at least a 1MB L2 cache for one or two CPUs. Four or more CPUs should have at a least 2MB L2 cache in each CPU. The greater the L2 cache, the greater the server's CPU performance because it reduces the amount of wait time experienced by the CPU when reading and writing data to main memory. Simple, single table queries and updates, along with query joins on small tables take minimal CPU processing power. On the other hand, large joins, aggregations, and sorting of large result sets uses a high level of CPU processing power. Keep this in mind when choosing the hardware configuration for the Oracle database server.

### Memory

In most cases, the more physical RAM Oracle has the greater Oracle's performance. If possible, purchase enough RAM to hold the largest database table in memory. If such a purchase is not possible during the initial setup, leave room for adding more RAM later. We strongly recommend running the 64-bit version of Oracle with a minimum of 4 GB RAM. Systems with high performance demands should be prepared to use up to 128 GB RAM or more.  To take advantage of Oracle's ability to run parallel queries, plan on investing in more RAM. Parallel queries use much more RAM than non-parallel queries.

### AWE

If Oracle is running on a 32-bit Windows box, consider using Address Windowing Extensions (AWE) to increase the amount of memory available to Oracle. Normally, 32-bit CPUs can only support up to 4GB of RAM because of limited address space. Oracle supports AWE to bypass this limitation and allows up to 64GB of RAM to be addressed.

### Drives

Avoid locating read-intensive and write-intensive activity on the same drive or array. For example, do not locate an OLTP and an OLAP database or heavily random and sequential activity on the same physical device. Whenever a drive or array must change back and forth between activities, efficiency is lost. NTFS-formatted partitions should not exceed 80% of their capacity. For example, a 20GB drive should never hold more than 16GB. NTFS needs room to work, and when capacity exceed 80%, NTFS become less efficient and I/O suffers. Consider creating a system alert to indicate when an array exceeds 80% of capacity so that immediate action can be taken to correct the problem.

### I/O and RAID

Use hardware-based RAID rather than software-based RAID because the latter can't offload the work to a separate processor, making it much slower than a hardware-based RAID solution. Do not store the operating system, application software, or databases on single disk drives because they do not afford any fault tolerance. Instead, always choose a RAID array made up of three or more physical drives that offers fault tolerance. Common fault tolerant RAID configurations include RAID Level 1 (mirroring or duplexing) and RAID Level 10 (also called 1+0, which includes both striping without parity and mirroring). Non-fault tolerant RAID configurations include RAID 0 which is simple disk striping. RAID 0 offers excellent performance and can be used for the TEMP tablespace. Each of these RAID levels offers different performance levels. Ideally, if the budget allows, chose RAID Level 10, which offers both high-speed and fault tolerance.

### I/O controller

Select the best I/O controller possible. Top-notch controllers offload much of the I/O work onto its own local CPU, freeing up CPU time on the server to do other tasks. For the ultimate in I/O controllers, consider a fiber channel connection instead of a SCSI connection. The controller should have the largest amount of cache RAM possible, with a minimum of 128mb of cache RAM. Generally, the greater the RAM cache on the controller, the higher the performance of the overall I/O, because data can be read ahead and stored in the cache, even if the data is not currently requested by Oracle. The data Oracle wants next from the array will likely be in the cache, speeding up data access.

Do not put DAT, DLT, CD-ROM, scanners, or other non-hard disk devices on the same I/O controllers that connect to the hard disk arrays. In addition, do not put hard disks on the same I/O controller if they have different speeds. Putting devices with different speeds on the same I/O controller slows the faster devices. Always put slower devices on their own I/O controller.

For maximum I/O throughput, assign each type of major I/O activity (database, REDO tablespaces, TEMP tablespaces, etc.) to its own separate RAID controller and dedicated RAID array.

### OLTP vs. OLAP

If the budget does not allow for the ideal number of disk controllers and hard disks to maximize the server's I/O performance, remember that optimal OLTP I/O is achieved by increasing disk reads and writes. The best way to do this is to add more hard disks to the array(s) that hold the database files and/or transaction logs. Adding more disks helps OLTP-based applications more than increasing the number or speed of disk controllers would, because OLTP-based applications tend to be limited by the number of transfer operations (read/writes) rather than bandwidth.

However, for OLAP-based applications, adding more and faster disk controllers to the array is generally a better way to boost I/O than increasing the number of disk drives, because OLAP applications tend to be more limited by bandwidth than by read/write operations. Adding faster or more disk controllers increases the bandwidth and helps to remove any bottlenecks.

### I/O Testing

There are several utilities available to test I/O subsystem performance. We recommend Oracle’s Orion tool.

### Networking

If the Oracle database server is not connected to a switch (as recommended for best performance), try the following suggestions for boosting network performance. If Ethernet is running, reduce the number of computers on the same segment as Oracle, and if possible, reduce the length of the physical segment Oracle is connected to. Do not cascade hubs. If Token-Ring is running, reduce the physical length of the ring and increase the number of active stations on the ring, which reduces token travel time.

### Network Protocols

For best performance, Oracle should be running on a dedicated server. Limit the number of network protocols installed on the server, because unnecessary network protocols increase overhead on the server and send out unnecessary network traffic. For the best overall performance, only install TCP/IP on the server.

### Routers

While not always possible (especially for WANs and Internet connections), try to avoid a router between Oracle clients and Oracle. In particular, avoid routers between two or more Oracle servers that need to communicate with each other. Routers are often a bottleneck for network traffic and can affect Oracle client/server performance. If Oracle must communicate over a router, ensure that the router has been properly tuned for maximum performance.

### Network Cards

The Oracle database server should have a minimum of one 100Mbs network card, and perhaps two. Two cards can be used to increase network throughput and to offer redundancy. In addition, the network card(s) should be connected to fullduplex switched ports for best performance.

Be sure that the network card(s) in the server are set to the same duplex level (half or full) and speed as the switched port they are connected to (assuming they are connected to a switch and not a hub). If there is a mismatch, the server may still be able to connect to the network, but network performance can be significantly impaired. Do not rely on network cards or switches that are supposed to auto-sense duplex or speed settings, because they often do not work correctly. Manually set the duplex and speed for the card from the operating system, and if necessary, manually make the same changes to the switch.

Windows allows network cards to save energy by going to sleep when they are not used. If any network card on a production server has a power management feature, ensure that the power savings feature is off. Otherwise, unexpected results, such as a network card that fails to wake up, or intermittent performance problems, may occur.  Check to see if the network card has a power management feature by viewing the Properties sheet for the network card's driver. View the Power Management tab on the Properties sheet, to verify the settings.

### Software

The network libraries chosen during Oracle installation can affect the speed of communications between the server and its clients. Of the three key network libraries, TCP/IP is the fastest and Multi-Protocol is the slowest. Due to the speed advantage, use TCP/IP on both the servers and clients. Do not install unused network libraries on the server, because they will contribute unnecessary overhead.

### Critical Patch Updates

We recommend staying current with Oracle Critical Patch Updates (CPUs).

### Services

Do not install any unnecessary Oracle services.

### File Structure and distribution

Place the database files and REDO files for all production databases on separate arrays to isolate potentially conflicting reads and writes. This means that the server will have at least two physical RAID arrays, one to store the database files, and a separate one to store the REDO files. The operating system can be stored on a mirrored set of drives. Use separate storage areas for TEMP tablespaces and the operating system.

### Third-party (Application) Access to Database

Be wary of allowing users to directly access the databases (especially OLTP databases) with third-party database access tools, such as Microsoft Excel or Access. Many of these tools can wreak havoc with database performance. Here are some reasons why:

- Often these users aren't experienced with these tools and create overly complex queries that eat up server resources. At the other extreme, their queries may not be complex enough (such as lacking effective WHERE clauses) and return thousands, if not millions, of unnecessary rows of data.
- This reporting activity can often lock rows, pages, or tables, creating user contention for data and reducing database performance.
- These tools are often file-based. This means that even if an effective query is written, the query is not performed at the server. Instead, the entire table (or multiple tables in the case of joins) must be returned to the client software where the query is performed. This leads to excessive server activity and can play havoc on your network.

If users must be allowed access to the data, limit hits on the production OLTP databases by pointing them to a "reporting" server that is replicated or in the form of a datamart or data warehouse.

For more detailed information see *CIC Database Configuration and Maintenance For ORACLE Server Technical Reference* document in the PureConnect Documentation Library at http://help.genesys.com/cic. http://help.genesys.com/cic.

# 5 Interaction View

## 5.1 Agent Experience

Interaction Desktop is the interface used to manage interactions and indicate availability. It is an extremely intuitive graphical application that uses single/double-click and drag and drop functionality to perform most operations. Interaction Desktop allows access to several interaction control buttons and tabs. It is important to note that access to all buttons, tabs, and configuration options are controlled by rights and permissions defined by the system administrator. Interaction Desktop is an installable Windows application.

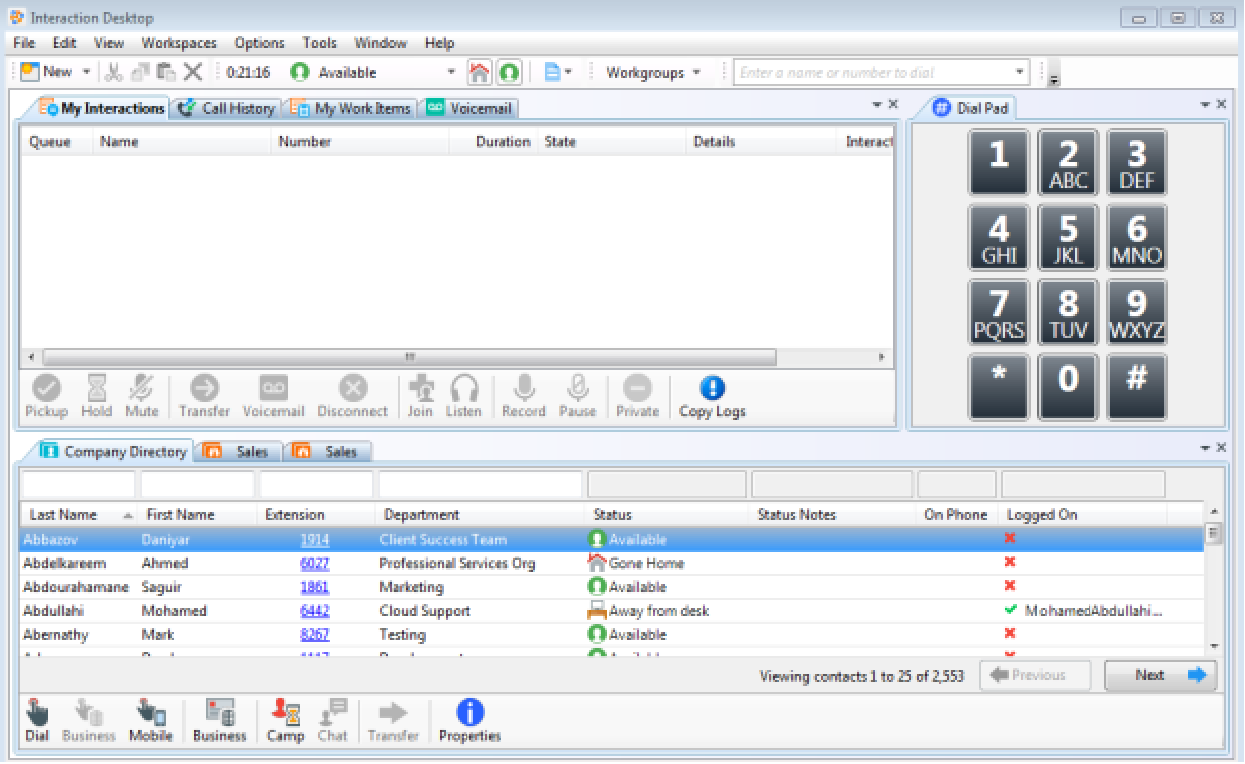

| Last Name | First Name | Extension | Department | Status | Status Notes | On Phone | Logged On |
|---|---|---|---|---|---|---|---|
| Abbazov | Daniyar | 1914 | Client Success Team | Available | | | ✗ |
| Abdelkareem | Ahmed | 6027 | Professional Services Org | Gone Home | | | ✗ |
| Abdourahamane | Saguir | 1861 | Marketing | Available | | | ✗ |
| Abdullahi | Mohamed | 6442 | Cloud Support | Away from desk | | | ✓ MohamedAbdullahi... |
| Abernathy | Mark | 8267 | Testing | Available | | | ✗ |

Viewing contacts 1 to 25 of 2,553

*Figure 19: Agent Desktop - Interaction Desktop*

Interaction Connect is the web-based agent interface. It allows for a zero-install deployment to agents and is the preferred user interface. Going forward, enhancements will only be made to Interaction Connect rather than to Interaction Desktop. Supervisory features are also quickly being added to Interaction Connect.

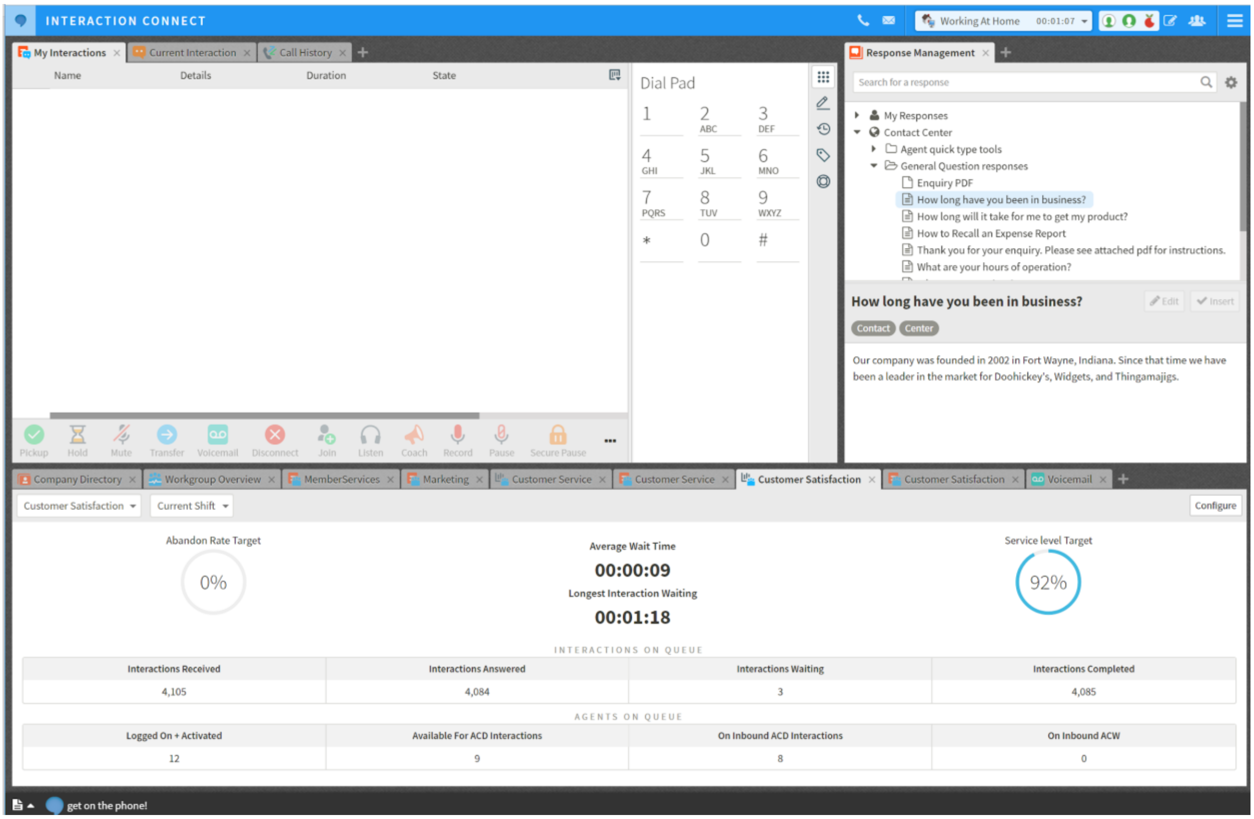

**How long have you been in business?**

Our company was founded in 2002 in Fort Wayne, Indiana. Since that time we have been a leader in the market for Doohickey's, Widgets, and Thingamajigs.

Abandon Rate Target: 0%

Average Wait Time: 00:00:09

Longest Interaction Waiting: 00:01:18

Service level Target: 92%

INTERACTIONS ON QUEUE

| Interactions Received | Interactions Answered | Interactions Waiting | Interactions Completed |
|---|---|---|---|
| 4,105 | 4,084 | 3 | 4,085 |

AGENTS ON QUEUE

| Logged On + Activated | Available For ACD Interactions | On Inbound ACD Interactions | On Inbound ACW |
|---|---|---|---|
| 12 | 9 | 8 | 0 |

*Figure 20: Agent Desktop - Interaction Connect*

## 5.2 Call Flows

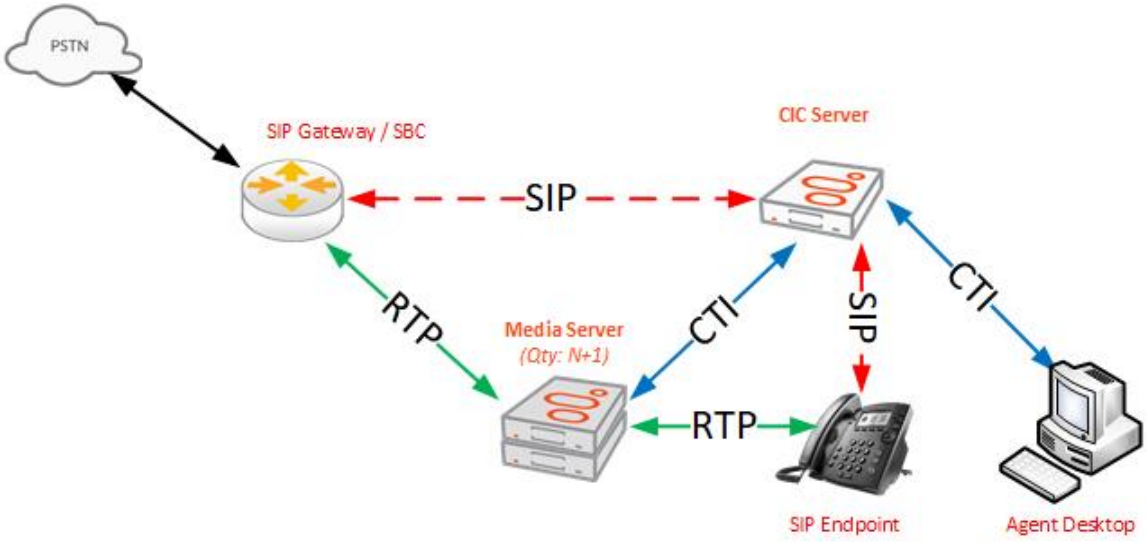

*Figure 21: Standard Call Flow*

## Call Flow Tools

Interaction Attendant is an easy-to-use graphical console that can be used to configure interactive voice response (IVR) behavior. Simply select the options to associate with each number on the telephone keypad. Options the user can select from include playing a message, routing to a queue, sending a fax, opening a submenu, and many others. Users can choose to start a handler to incorporate unlimited custom functionality.

Interaction Attendant includes the ability to schedule the availability of different menus at different times. For example, after hours and on holidays the user might only want to offer a subset of options.

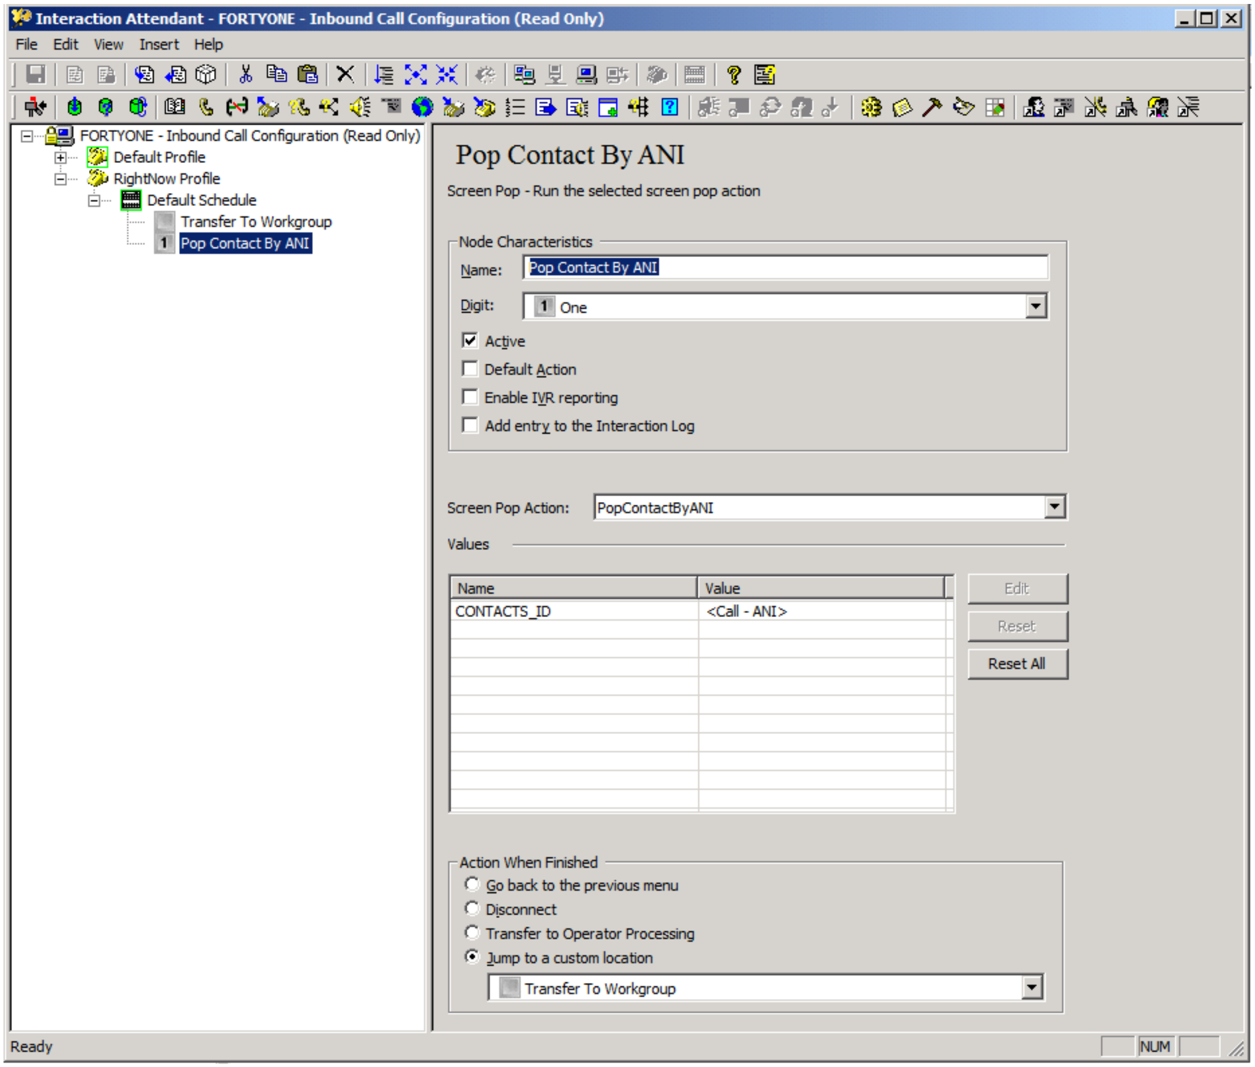
Interaction Attendant - FORTYONE - Inbound Call Configuration (Read Only)
File
Edit
View
Insert
Help
Edit
Reset
Reset All
Ready
NUM

FORTYONE - Inbound Call Configuration (Read Only)
- Default Profile
- RightNow Profile
  - Default Schedule
    - Transfer To Workgroup
    - 1 Pop Contact By ANI

Pop Contact By ANI

Screen Pop - Run the selected screen pop action

Node Characteristics

Name: Pop Contact By ANI

Digit: 1 One

☑ Active

☐ Default Action

☐ Enable IVR reporting

☐ Add entry to the Interaction Log

Screen Pop Action: PopContactByANI

Values

| Name | Value |
|---|---|
| CONTACTS_ID | <Call - ANI> |

Action When Finished

○ Go back to the previous menu

○ Disconnect

○ Transfer to Operator Processing

● Jump to a custom location

Transfer To Workgroup

*Figure 22: Call Flow Tools - Interaction Attendant*

## 5.3 External Interfaces

Describes the external interface for the solution.

List the external interfaces and relevant tasks for integration

| Interface | Protocol | Solution Components | Integration Tasks | Description |
|---|---|---|---|---|
| Media Gateway/ Session Border Controller (SBC) | SIP and RTP<br>SIP/TLS and SRTP | CIC Server,<br>Media Server,<br>SIP Proxy,<br>SIP Endpoints | Add the necessary bandwidth to the network<br>Provision the network infrastructure (e.g. DNS) for the new traffic<br>Provision the MG and SBC appropriately for the integration<br>Configure the codec list for supported codecs. | This interface is used to handle ingress and egress voice traffic from the network.<br>Note that the SfB Mediation Server may be the key component integrated into the SBC in which case the integration will be with the Mediation Server and not the SBC. |
| Reporting Databases (Relational Database/RDBMS) | TCP/SQL | CIC Server,<br>Central Campaign Server | Provision the network infrastructure (e.g. DNS) for the new traffic<br>Run the database scripts (.sql)<br>Provision appropriate user access to required database tables | This interface is used to store historical reporting data. It is also used by Interaction Dialer campaigns. |

| Interface | Protocol | Solution Components | Integration Tasks | Description |
|---|---|---|---|---|
| Corporate Backend Servers | HTTPS (REST or SOAP), RDBMS access methods (optional) | CIC Server,<br>Interaction Scripter | Provision the network infrastructure (e.g. DNS) for the new traffic<br>Create and provision the security information (certificates, etc.) | This interface is used to get data from the corporate systems to make decisions in solution (agent desktop, routing strategy, etc.). It can also be used to perform certain business actions. |
| Enterprise Authentication Service (Optional) | Active Directory,<br>ADFS,<br>SAML 2.0 | CIC Server | Provision the network infrastructure (e.g. DNS) for the new traffic<br>Create and provision the security information (certificates, etc.) | This interface is used to perform authentication of users using the solution. Genesys provides the option to have user passwords authenticated by an external authentication service or authentication can be managed by Genesys. |
| Corporate Network Management System (Optional) | SNMP | Genesys SNMP Master Agent | Provision the network infrastructure (e.g. DNS) for the new traffic<br>Create and provision the security information (certificates, etc.) | This interface is used to integrate the solution with the Corporate network management system. |

| Interface | Protocol | Solution Components | Integration Tasks | Description |
|---|---|---|---|---|
| Domain Name Servers | DNS | SIP Proxy, SIP Server, Workspace, Interactive Insights, SIP Endpoints, etc. (potentially any Genesys application). | Provision the DNS records along with appropriate weightings | This interface is used by the clients to perform the name/IP address translation. For specific cold standby components, the DNS entries will be manually modified to redirect traffic in the event of a site failure. |

*Table 5: External Interfaces*

## 5.4 Operational Management

Operational management of the PureConnect Solution is enhanced by the consolidated nature of the solution. Most management tools are contained on the PureConnect Server itself which improves ease of use and provides management tools that dramatically improve the performance of any business. Additional considerations for operational management include network operations, serviceability, and other management considerations, all of which are covered in the following section.

### 5.4.1 Network Management Systems

If the customer does have a Network Management System (NMS), then Genesys components need to be integrated into their NMS. This is typically done by setting up the ININ SNMP Service to send SNMP events (traps) and info to their NMS.

Examples of supportable NMS includes HP OpenView and OpenNMS (an open source NMS - http://www.opennms.org/).

PureConnect products and the PureConnect SNMP service support SNMPv2c and SNMPv3. The PureConnect SNMP service also supports SNMP agents that use the SNMPv1, SNMPv2, and SNMPv3 standards.

Documentation regarding SNMP and the PureConnect components that support it can be found at https://help.genesys.com/cic/mergedProjects/wh_tr/desktop/pdfs/snmp_tr.pdf.

### 5.4.2 Serviceability

Serviceability relates to the ability of technical support to identify issues and defects within the system. Most of this relates to the ability to retrieve logs and configuration information and pass them back to technical support.

Setting up logical logging locations is a best practice that can reduce the time to send logs to support. Configuring 3rd party components to log into the same location is ideal as well. Establishing a "log" directory in the root of the disk structure and logging there is recommended. On CIC servers it is best practice to place trace logs on a separate drive or partition. Trace logs can become quite large on a busy system if trace levels are turned up for troubleshooting. It is not recommended to continuously run at high trace levels during operation, but rather to temporarily increase log levels during troubleshooting. Typical locations for trace logs are:

```
D:\I3\IC\Logs
C:\Windows\Temp\inin_tracing
```

Guidelines for setting up trace logging are contained in https://help.genesys.com/cic/mergedProjects/wh_iandc/desktop/ic_installation_and_configuration_guide.htm.

Log Retrieval Assistant (LRA) is a feature that allows support organizations such as Genesys and certified partners to configure logging and retrieve logs from CIC servers at specific times. Log Retrieval Assistant is installed automatically during the CIC server install. For more information, see the Log Retrieval Assistant Customer Site Technical Reference in the PureConnect Documentation Library.

### 5.4.3 Usage

Starting with PureConnect 2018 R2 Patch 2, new on-premises customers have the option to use subscription billing. Usage data for subscribers is collected in encrypted files and then uploaded, either automatically or manually, for billing. Billing data for subscribers appears in MyCloud. Subscription customers must have a subscription agreement with Genesys.

The usage files, located in the i3\ic\work\usage directory, can be automatically uploaded daily to the Genesys PureConnect Amazon S3 location. Subscription customers must configure firewall rules to allow sending of these files on port 443. The Amazon S3 location is **us-east-2**. Some enterprises block egress TCP 443, so you need to Verify Access for the Automated Upload to make sure your site can access the AWS billing service and no firewalls block the connection.

Your organization can manually upload the usage data monthly from the Genesys Customer Care Portal. If you have access, click the **Subscription Usage Upload** tile in the portal dashboard to go to the upload page where you can upload your encrypted usage data files.

Your organization must upload the usage data for the required billing period before the 4th of each month. The billing period is the 28th through the 27th. For example, if the billing period is April 28th through May 27th, then all usage files must be uploaded from the portal by June 4th. Failure to do so may result in disablement of the system, as well as financial penalties. If you choose to manually upload the usage files, you are responsible for deleting the files from the usage directory.

For automated uploads, the license file must include the **I3_FEATURE_USAGE_TRACKING_UPLOAD** license key. For manual uploads, the license management site loads the license file with the technical license key **I3_FEATURE_USAGE_TRACKING_LOCAL**, which prevents the automatic upload of usage data. If the license file does not include either license key, the system does not collect usage data. For more information about licensing, see the PureConnect Licensing Technical Reference.

The amount of disk space needed for the usage files varies depending on the number of logins per user per day. For most customers, the file will be less than 3 MB per day, while the largest customers may see files as large as 10 MB per day. PureConnect creates one usage file each day for each PureConnect server and creates additional usage files when a server is restarted.

# 6 Implementation View

The Implementation View describes details such as sizing, security and configuration of the solution based on the previous deployment and interaction views.

## 6.1 Solution Sizing Guidelines

Solution sizing is important across scalable components of PureConnect. Sizing examples and relevant calculators are provided below for Media Servers, Storage, Database, and Network Sizing.

### 6.1.1 Sizing Assumptions

The following assumptions are made regarding the sizing of this solution.

| **Input Assumptions** | |
|---|---|
| Agents | 2,000 |
| Agent utilization | 80% |
| Call qualification time | 60s |
| Queue time | 120s |
| Talk time | 180s |
| Transfer Rate | 10% |
| Conference Rate | 10% |
| Percentage of Queued Calls | 50% |
| Log retention | Debug 1 week |
| Reporting History | 2 years |

| | |
|---|---|
| Non-aggregated Reporting History | N/A |
| ***Calculated worst case values*** | |
| Concurrent active calls (IVR + Queue + Agent) | 1662 |
| Peak CAPS | 6 |
| Busy hour calls | 21600 |
| Emails/Day | 17,280 |
| Chat – Peak interaction rate (in sec) | 0.5 |
| Chats/Day | 8640 |

*Table 6: Sizing Inputs*

### 6.1.2 Media Server Sizing

Media Server sizing is based on total points required, which are assigned based on session types. Session types may require 1 or ½ a session depending on the session type. For example, a call session requires 1 point, and to add transcription adds ½ a point, which would total 1 ½ points for a call with transcription. The Media Server Estimation Spreadsheet should be used to calculate these points, and a full listing of Media Server point requirements can be seen in this document. Once Media Server points are calculated, the spreadsheet will also indicate the server count required. This will be calculated for Small, Medium, and Large Server hardware, with points per server indicated in the table below. Server Counts and N+1 Configuration are indicated in the Media Server Estimation Spreadsheet.

| **Media Server Platform** | **Points Per Server** |
|---|---|
| Small Appliance (G9 DL-60) | 240 |
| Medium Appliance (G9 DL-360) | 480 |
| Large Appliance (G9 DL-360) | 960 |

*Table 7: Media Server Points*

### 6.1.3 Storage Sizing

Voice recording storage sizing should be based on:

- Average voice recording size: bitrate times the average file length in seconds divided by 8
  - 16 kbps for mp3 (stereo)
  - 8 kbps for mp3 (mono)
- Average recording size times the retention period

Screen recording storage – refer to the Screen Recording Calculator.

### 6.1.4 Database Sizing

Database sizing is based on the following inputs:

- Data retention period in months
- Average calls per hour
- Days of week in operation
- Hours per day of operations
- Average shift length
- Number of agents (users)
- Number of workgroups
- Average workgroups per agent
- Average agent status changes / day
- Average number of faxes sent/received by an agent per day
- Lines / circuits
- Line report groups

Calculator information:

- **Calculator name: db_spaceplanning.xls**
- Example Calculation (Please use calculator for more precise estimates):
  - 12 Months Retention
  - 5 Days a week of operations
  - 12 Hour Days
  - 8 Hour Shifts
  - 2000 Agents
  - 50 Workgroups
  - 10 Status Changes per day
  - 1 Fax per agent per day
  - 14 Lines
  - 14 Line groups
  - Calculated Database Size: ~36 GB.

For additional details on the PureConnect Database, please refer to the Installation and Configuration Guide, CIC Database Configuration and Maintenance for SQL Server Technical Reference, and CIC Database Configuration and Maintenance for Oracle Technical Reference.

### 6.1.5 Network Sizing and Readiness

For PureConnect there are two main considerations for network bandwidth:

- Application Traffic: calculation of bandwidth sized per application instance.
  - Interaction Desktop Client
  - Interaction Connect Client
  - ICBM Client
  - Off-host Session Manager to CIC communication
  - Text to Speech session (TTS)
  - Interaction Recorder file transfer
  - SIP Call Control Traffic
  - Media Server Sessions
  - Numbers of utterances (Analyzer)
  - 3rd party WFM RTA Integration

- Call Traffic: calculation of bandwidth sized per call instance.
  - G.729 Calls
  - G.711 Calls
  - T.38 and TTS
  - Conference Calls

Use "Bandwidth Calculator" to perform all Application and Call Traffic sizing per site." to perform all Application and Call Traffic sizing per site.

### 6.1.6 Solutions Sizing - Dual Data Center

PureConnect Server Sizing considerations should include the following:

- PureConnect Server: Runs in Switchover Mode which should be sized with switchover calculator
- Media Servers: Sized N+1 based on local trunking requirements (Use Media Server Calculator)
- Off Server Session Managers.
- Central Campaign Servers (Dialer).

## 6.2 Configuration Guidelines

Configuration guidelines which should be followed include an installation checklist which can be used to perform a new CIC installation. Full configuration guidelines and recommendations can be found in the Installation and Configuration Guide. This document describes how to install and configure a new PureConnect installation. It includes PureConnect CIC Server, Interaction Media Server, client workstation, database server, mail server, add-on server requirements, and post-installation procedures.

### 6.2.1 Configuration Recommendations

While it is difficult to offer best practice guidance to cover every situation, we offer the following content linked, which is part of a growing body of knowledge. Thus, the following list of best practices is not comprehensive, and Genesys SC and PS resources should be used as necessary to cover all design and deployment best practices on a case by case basis:

Interaction Dialer

Interaction Attendant

CIC Web applications configuration

Interaction Web Portal and Interaction Marquee

PureConnect products on a virtualized platform

Migration from CIC and Interaction Dialer 2.4 or 3.0 to current PureConnect releases

Testing results and supported best practices can also be found at Testlab.

## 6.3 Security

Protecting the customer's infrastructure should be imperative for any solution deployment. Many customers have their own security procedures that our solution needs to conform to. The following are additional documents explaining requirements and recommendations for the customer to maintain a secure environment:

PureConnect Security Features Technical Reference. This document covers basic security features employed in PureConnect, including its use of SSL, TLS, and SRTP protocols along with public key cryptography and certificates to enhance application security.

PureConnect Security Precautions Technical Reference includes information on general security practices, a "Top 10" list of security topics, and details about how to make the CIC server more secure.

Secure Input Technical Reference describes how to install and configure Genesys support for secure input of sensitive or confidential data such as credit card numbers. describes how to install and configure Genesys support for secure input of sensitive or confidential data such as credit card numbers.

Genesys PureConnect Documentation Library is the location where these documents are maintained.

### 6.3.1 Secure Connections

PureConnect Port Maps and Data Flow Diagrams Technical Reference provides diagrams illustrating the default network ports and protocols for data flow between workstations, devices, and PureConnect products, as well as existing devices on the customer network environment and PureConnect systems. These diagrams can also be provided to customers to assist in network security considerations.

### 6.3.2 VM and OS hardening

Operating Systems are often pre-configured for ease of use and development and not necessarily security. If the O/S is being installed or is part of a set of VMs being delivered, that O/S should be hardened to ensure that typical security holes are addressed.

PureConnect Installation and Configuration Guide provides Windows security-related settings, including Windows permissions, Windows Firewall settings, User Account Control (UAC) settings, and DCOM settings which should be followed.

PureConnect Virtualization Technical Reference provides guidelines for running PureConnect products on a virtualized platform. This information is based on tested configurations that are proven to provide the best experience for each product.

### 6.3.3 PCI

The Payment Card Industry (PCI) Security Standards Council defines the PCI standard as follows:

A worldwide information security standard ... created to help organizations that process card payments prevent credit card fraud through increased controls around data and its exposure to compromise. The current version of the standard (1.2) specifies 12 requirements for compliance, organized into six logically related groups, called "control objectives."

http://www.pcisecuritystandards.org

Secure Input, which is a component of the PCI standard, separates and encrypts, or obfuscates, submitted, personal data to protect it from theft or misuse.

PureConnect supports Secure Input through separation and data protection. PureConnect separates confidential customer information input into the system from other parties in the interaction to prevent disclosure of confidential data to agents, coaches, monitors, etc. In addition to separating user input, PureConnect protects the data from access by other users by providing a secured, untraced session, and minimal access to sensitive information in memory, only by authorized processes which are necessary to complete the interaction, in a way that minimizes risk of accidental disclosure.

Additional information on PureConnect Secure Input features can be found in the Secure Input Technical Reference.

## 6.4 Localization and Internationalization

PureConnect is localized into more than 21 languages. Additional details are available in the Language Pack Technical Reference. The following PureConnect Language Packs are available. Please note, localization scope can be different based on the language:

| Language | Localization scope |
|---|---|

| | |
|---|---|
| Arabic | Prompts, IC User Apps |
| Chinese (Simplified) | Prompts, IC User Apps, IC Business Manager Apps including Reports, IC Server Manager Apps, Interaction Client Web Edition, Interaction Connect |
| Chinese (Traditional) | Prompts, IC User Apps, IC Business Manager Apps including Reports, IC Server Manager Apps, Interaction Client Web Edition, Interaction Connect |
| Danish | Prompts, IC User Apps, Interaction Client Web Edition, Interaction Connect |
| Dutch | Prompts, IC User Apps, IC Business Manager Apps, Interaction Client Web Edition, Interaction Connect |
| English (Australia) | Prompts |
| English (New Zealand) | Prompts |
| English (United Kingdom) | Prompts |
| French | Prompts, IC User Apps, IC Business Manager Apps including Reports, IC Server Manager Apps, Interaction Client Web Edition, Interaction Connect |
| French (Canada) | Prompts, IC User Apps, IC Business Manager Apps, IC Server Manager Apps, Interaction Client Web Edition, Interaction Connect |
| German | Prompts, IC User Apps, IC Business Manager Apps including Reports, IC Server Manager Apps, Interaction Web Portal*, Interaction Client Web Edition, Interaction Connect |
| Hebrew | Prompts, IC User Apps |
| Italian | Prompts, IC User Apps, IC Business Manager Apps, Interaction Client Web Edition, Interaction Connect |

| | |
|---|---|
| Japanese | Prompts, IC User Apps, IC Business Manager Apps including Reports, IC Server Manager Apps, Interaction Client Web Edition, Interaction Client Mobile Web Edition, Interaction Connect |
| Korean | Prompts, IC User Apps, IC Business Manager Apps, Interaction Client Web Edition, Interaction Connect |
| Norwegian | Prompts, IC User Apps, IC Business Manager Apps, Interaction Client Web Edition, Interaction Connect |
| Polish | Prompts, IC User Apps, IC Business Manager Apps including Reports, IC Server Manager Apps, Interaction Client Web Edition, Interaction Connect |
| Portuguese (Brazil) | Prompts, IC User Apps, IC Business Manager Apps including Reports, IC Server Manager Apps, Interaction Client Web Edition, Interaction Connect |
| Russian | Prompts, IC User Apps, Interaction Client Web Edition, Interaction Connect |
| Serbian (Latin) | Prompts, IC User Apps, Interaction Client Web Edition, Interaction Connect |
| Spanish (Latin America) | Prompts, IC User Apps, IC Business Manager Apps, IC Server Manager Apps, Interaction Client Web Edition, Interaction Connect |
| Spanish (Spain) | Prompts, IC User Apps, IC Business Manager Apps including Reports, IC Server Manager Apps, Interaction Client Web Edition, Interaction Connect |
| Swedish | Prompts, IC User Apps, IC Business Manager Apps, Interaction Client Web Edition, Interaction Connect |
| Turkish | Prompts, IC User Apps, IC Business Manager Apps, Interaction Client Web Edition, Interaction Connect |

*Table 8: Supported Languages*