

297-2183-106

Nortel Networks Symposium Call Center Server

Planning and Engineering Guide

Product release 5.0

Standard 2.0

June 2004

NORTEL
NETWORKS™

Nortel Networks Symposium Call Center Server Planning and Engineering Guide

Publication number: 297-2183-106
Product release: 5.0
Document release: Standard 2.0
Date: June 2004

Copyright © 2004 Nortel Networks, All Rights Reserved

Information is subject to change without notice. Nortel Networks reserves the right to make changes in design or components as progress in engineering and manufacturing may warrant.

The process of transmitting data and call messaging between the Meridian 1 or DMS/MSL-100 switch and Symposium Call Center Server is proprietary to Nortel Networks. Any other use of the data and the transmission process is a violation of the user license unless specifically authorized in writing by Nortel Networks prior to such use. Violations of the license by alternative usage of any portion of this process or the related hardware constitutes grounds for an immediate termination of the license and Nortel Networks reserves the right to seek all allowable remedies for such breach.

*Nortel Networks, the Nortel Networks logo, the Globemark, CallPilot, Contivity, DMS, IVR, Meridian, Meridian 1, Meridian Mail, Meridian SL, Optivity, Succession, and Symposium are trademarks of Nortel Networks.

INTEL, INTEL XEON, and PENTIUM are trademarks of Intel Corporation.

MICROSOFT, MICROSOFT ACCESS, WINDOWS, WINDOWS NT, and WINDOWS XP are trademarks of Microsoft Corporation.

REPLICATION SERVER and SYBASE are trademarks of Sybase, Inc.

PCANYWHERE is a trademark of Symantec Corporation.

Publication history

June 2004

The Standard 2.0 version of the *Nortel Networks Symposium Call Center Server Planning and Engineering Guide*, Release 5.0, is released. The CapTool user documentation has been removed from this version and placed in a separate *CapTool User's Guide*.

April 2004

The Standard 1.0 version of the *Nortel Networks Symposium Call Center Server Planning and Engineering Guide*, Release 5.0, is released.

Contents

1	Introduction	11
	Overview	12
	What's new in Release 5.0	13
	Engineering a contact center	17
	Skills you need	18
	Related documents	19
2	Contact center architecture	25
	Overview	26
	Meridian 1/Succession 1000	27
	DMS	34
3	Engineering Symposium Call Center Server	37
	Overview	38
	Hardware configurations	39
	Call load	42
	MLS and HDX performance impact	46
	Guidelines to minimize capacity requirements	52
	Capacity estimation	55
4	Engineering the client	59
	Overview	60
	Section A: Symposium Web Client	63
	Architecture	64
	Requirements	65
	Application server performance	66
	Client performance	72
	Minimizing CPU load	73
	Section B: Classic Client	75
	Classic Client configuration	76

5	Engineering the switch	77
	Meridian 1/Succession 1000	78
	DMS	84
6	Engineering the network	91
	CLAN requirements	92
	ELAN requirements	94
	ELAN connection to CLAN	99
	WAN requirements (Meridian 1/Succession 1000 only)	101
	Recommended network equipment	103
7	Engineering the voice processing system (Meridian 1/ Succession 1000 only)	105
	Overview	106
	Port usage	107
	ACCESS requirements	109
	Symposium Voice Services on CallPilot requirements	110
	Symposium Voice Services on Meridian Mail requirements	111
8	Setting up remote support with a VPN	117
	Overview	118
	Guidelines for the remote support VPN at the customer's premises	119
	VPN configurations	120
A	Product limits	125
	Product limits	126
B	Standard call models	133
	Inbound call models	134
C	IP Multicast Networking	139
	Overview	140
	Multicast sending and receiving	141
	Implementing IP multicasting for Symposium Call Center Server	151

D	Calculating Equivalent Basic Calls	155
	Equivalent Basic Calls	156
	Glossary	159
	Index	187

Chapter 1

Introduction

In this chapter

Overview	12
What's new in Release 5.0	13
Engineering a contact center	17
Skills you need	18
Related documents	19

Overview

Welcome

Nortel Networks presents Symposium Call Center Server. The server is designed to provide a contact center solution for varied and changing business requirements by offering a suite of applications that includes call processing and agent handling, management and reporting, networking, and third-party application interfaces.

Some advantages of Symposium Call Center Server are

- complete call control and reporting
- application flexibility
- state-of-the art user interface
- industry standard, client-server architecture
- open interfaces: Database, Real time, Host routing, and Meridian Link Services (MLS)
- comprehensive networking through public and private networks
- leveraged PBX switching reliability and client-server processing power

Who should read this guide

This guide is for Symposium Call Center Server system designers and technical support staff members. It is also intended to be used by administrators who are responsible for day-to-day management of the Symposium Call Center Server configuration.

Network information

This guide contains references to the Symposium Call Center Server Network Skill-Based Routing (NSBR) feature. However, this feature is not available for all switch types. For more information on networking, refer to the *Symposium Call Center Server Network Control Center Administrator's Guide*.

What's new in Release 5.0

Changes to product limits

Release 5.0 contains a number of capacity enhancements over Release 4.2. The following table shows the product limits that have been increased in Release 5.0, along with the original Release 4.2 limit.

Parameter	Release 5.0 limit	Release 4.2 limit
Number of logged-on agents		
Meridian 1/Succession 1000	2200	1500
DMS	3300	N/A
Number of configured agents	6000	3000
Maximum script size (characters)	50 000	30 000
Number of skillsets	1000 (996)	354 (350)
Number of network skillsets	1000	100
Number of inbound calls per hour	58 000	55 000
Number of IVR ports	1000	500
Number of MLS calls per hour	58 000	16 000
Number of servers per switch (Meridian 1/Succession 1000)	3	1
Number of CPUs	4	2

Notes:

- For a complete list of product limits, see Appendix A, “Product limits.”
- The product contains four predefined skillsets. Therefore, for Release 4.2, the customer can create 350 skillsets, and for Release 5.0, the customer can create 996 skillsets.

- The number of skillsets given here includes both local and network skillsets.
- Throughout this guide, the term DMS switch applies to the following switch types:
 - DMS switch
 - MSL-100 switch
 - Succession 2000 switch
 - Nortel Networks Communication Server 2100 (CS 2100) switch

Multiple servers on a Meridian 1/Succession 1000 switch

In Symposium Call Center Server Release 5.0, one Meridian 1/Succession 1000 switch can support up to three Symposium Call Center Server systems.

The three servers cannot use the same switch resources simultaneously.

The following table shows how Nortel Networks products and features interact with multiple servers configured on the same switch.

Product/ Feature	Interaction
Symposium Web Client	You can use a single Symposium Web Client to manage all of the servers installed on your switch. For each server, in the Configuration component, choose Server → Add Server. Then enter the server information.
Classic Client	You can use a single Classic Client to manage multiple servers. Therefore, you can use one Classic Client to manage all of the servers installed on your switch. In the SMI Workbench, add each server as a system (see the <i>Administrator's Guide</i>). You can only log on to one server at a time.
Network Skill-Based Routing (NSBR)	Networking of servers on the same switch is not supported.

Product/ Feature	Interaction
CallPilot	<p>If you are using CallPilot to provide front-end IVR, the same CallPilot server can support all three Symposium Call Center Server systems.</p> <p>If you are using Symposium Voice Services on CallPilot—that is, if CallPilot is providing Give IVR or ACCESS voice services (Open/Close Voice Session, Collect Digits, and Give Controlled Broadcast)—CallPilot can serve only one Symposium Call Center Server system.</p>
Meridian Mail	<p>If you are using Meridian Mail to provide front-end IVR, the same Meridian Mail system can support all three Symposium Call Center Server systems.</p> <p>If you are using Symposium Voice Services on Meridian Mail to provide IVR services (that is, with the Give IVR command), the same Meridian Mail can support all three Symposium Call Center Server systems. However, the following restrictions apply:</p> <ul style="list-style-type: none"> ■ You must allocate the Meridian Mail IVR ports between three IVR queues, and dedicate a queue to each server. ■ All of the servers must belong to the same customer group. (Therefore, you cannot network the servers together.) <p>If you are using Symposium Voice Services on Meridian Mail to provide ACCESS voice services (Open/Close Voice Session, Collect Digits, and Give Controlled Broadcast), Meridian Mail can serve only one Symposium Call Center Server system.</p>
HDX	<p>Multiple servers can access the same database if your HDX application allows all of the servers to register.</p>
IPML	<p>If you are using the Integration Package for Meridian Link (IPML), an IPML server can support only one Symposium Call Center Server. Each Symposium Call Center Server must have its own IPML server.</p>

Product/ Feature	Interaction
TAPI	If you are using the Telephony Application Program Interface (TAPI), a TAPI server can support only one Symposium Call Center Server. Each Symposium Call Center Server must have its own TAPI server.
Agent Greeting/ Remote Observe	Since these features are implemented on the switch, with no dependency on Symposium Call Center Server, multiple servers can share Agent Greeting/Remote Observe.

CapTool User's Guide

Instructions for using the Capacity Assessment Tool (CapTool) have been removed from the *Planning and Engineering Guide*. For information about CapTool, see the new *CapTool User's Guide*.

Engineering a contact center

Engineering tasks

When engineering a contact center, you must perform the tasks listed in the following checklist.

Description	✓
Determine requirements for Symposium Call Center Server. See Chapter 3, “Engineering Symposium Call Center Server.”	
Determine requirements for the client. See Chapter 4, “Engineering the client.”	
Determine switch requirements. See Chapter 5, “Engineering the switch.”	
Determine network requirements. See Chapter 6, “Engineering the network.”	
Determine the requirements of the voice processing system. See Chapter 7, “Engineering the voice processing system (Meridian 1/ Succession 1000 only).”	
Determine the requirements of the remote support system. See Chapter 8, “Setting up remote support with a VPN.”	

Skills you need

Requirements

To successfully engineer a call center, you must be familiar with

- call center operations and metrics
- Symposium applications/products
- computer, networking, and traffic engineering performance parameters and metrics

You need not be an expert in these areas, but you must have some familiarity with these concepts.

In addition, you must have experience running and using Windows applications.

Related documents

Introduction

This section lists the documents in which you can find additional information related to Symposium Call Center Server.

Note: Documentation is available online, at the following URL:

<http://www130.nortelnetworks.com/cgi-bin/eserv/cs/main.jsp?cscat=documentation>

Symposium Call Center Server installation

The following documents contain procedures for installing the Symposium Call Center Server hardware and software:

If you need information about	Refer to
■ installing your server software	<i>Nortel Networks Symposium Call Center Server Installation and Maintenance Guide</i>
■ planning the network configuration between the DMS switch and the WAN	<i>Nortel Networks Symposium Call Center Server DMS-100 ICM Router Guide</i> <i>Nortel Networks Symposium Call Center Server and DMS/MSL-100 Switch Guide</i>
■ installing the Network Control Center server	<i>Nortel Networks Symposium Call Center Server Installation and Maintenance Guide</i>

Symposium Call Center Server setup

The following documents provide instructions for the setup and configuration of Symposium Call Center Server and the Meridian 1/Succession 1000 or DMS family of switches:

If you need information about	Refer to
<ul style="list-style-type: none"> ■ configuring the server 	<p><i>Nortel Networks Symposium Call Center Server Setup Guide and the Nortel Networks Symposium Call Center Server Administrator's Guide</i></p>
<ul style="list-style-type: none"> ■ configuring the Network Control Center server 	<p><i>Nortel Networks Symposium Call Center Server Network Control Center Administrator's Guide</i></p>
<ul style="list-style-type: none"> ■ Meridian 1 or Succession 1000 switch configuration 	<p><i>Nortel Networks Symposium Call Center Server Symposium, M1/Succession 1000, and Voice Processing Guide</i></p>
<ul style="list-style-type: none"> ■ DMS switch configuration 	<p><i>Nortel Networks Symposium Call Center Server and DMS Switch Guide</i></p>

DMS documents

The following documents provide instructions for the administration of the DMS switch:

If you need information about	Refer to
<ul style="list-style-type: none"> ■ utilities used to manage and monitor the switch 	<p><i>DMS Utilities Guide</i></p>
<ul style="list-style-type: none"> ■ Ethernet Interface Unit (EIU) installation and configuration 	<p><i>EIU Installation and Configuration Guide</i></p>

Meridian 1/Succession 1000 switch documents

The following documents provide instructions for the administration of the Meridian 1/Succession 1000 switch:

If you need information about	Refer to
<ul style="list-style-type: none"> ■ determining the switch requirements for Meridian 1/Succession 1000 	<i>Large System Planning and Engineering or Small System Planning and Engineering</i>
<ul style="list-style-type: none"> ■ determining the ELAN requirements 	<i>Data Networking for VoIP</i>
<ul style="list-style-type: none"> ■ overlays used to manage and monitor the switch 	<i>Software Input/Output Guide: Administration</i> <i>Software Input/Output Guide: Maintenance</i> <i>Software Input/Output Guide: System Messages</i>

Symposium Call Center Server administration

The following documents provide instructions for the administration of Symposium Call Center Server with the Classic Client:

If you need information about	Refer to
<ul style="list-style-type: none"> ■ the support and administration of the call center application that runs on client PCs connected to the server 	<i>Nortel Networks Symposium Call Center Server Administrator's Guide</i>
<ul style="list-style-type: none"> ■ setting up real-time displays ■ managing reports 	<i>Nortel Networks Symposium Call Center Server Supervisor's Guide</i>
<ul style="list-style-type: none"> ■ accessing the database ■ entity relationship diagrams (ERDs) 	<i>Nortel Networks Symposium Call Center Server Historical Reporting and Data Dictionary</i>
<ul style="list-style-type: none"> ■ creating and administering call center scripts 	<i>Nortel Networks Symposium Call Center Server Scripting Guide</i>

If you need information about	Refer to
<ul style="list-style-type: none"> ■ support and administration of the Network Control Center server 	<p><i>Nortel Networks Symposium Call Center Server Network Control Center Administrator's Guide</i></p>

Symposium Web Client

The following documents provide instructions for installing and configuring Symposium Web Client and for using it to administer Symposium Call Center Server.

If you need information about	Refer to
<ul style="list-style-type: none"> ■ the installation and configuration of Symposium Web Client on the application server and client PCs 	<p><i>Nortel Networks Symposium Call Center Web Client Planning, Installation, and Administration Guide</i></p>
<ul style="list-style-type: none"> ■ setting up real-time displays ■ managing reports 	<p><i>Nortel Networks Symposium Call Center Web Client Supervisor's Reference Guide</i></p>
<ul style="list-style-type: none"> ■ administering Symposium Call Center Server with Symposium Web Client 	<p>Symposium Web Client online Help</p>

Contivity 1100

The following documents provide instructions for the installation and configuration of the Contivity 1100.

If you need information about	Refer to
<ul style="list-style-type: none"> ■ installing the Contivity 1100 	<p><i>Installing the Contivity 1010/1050/1100</i></p>
<ul style="list-style-type: none"> ■ installing the modem option on Contivity 1100 	<p><i>Installing the Contivity 1010/1050/1100</i></p>
<ul style="list-style-type: none"> ■ restricting user access to the network 	<p><i>Configuring Basic Features for the Contivity Secure IP Services Gateway</i></p>

If you need information about

Refer to

- configuring split tunneling on the Contivity 1100

Configuring Basic Features for the Contivity Secure IP Services Gateway

Chapter 2

Contact center architecture

In this chapter

Overview	26
Meridian 1/Succession 1000	27
DMS	34

Overview

The *Symposium Call Center Server Planning and Engineering Guide* provides information on how to determine the engineering requirements of your Symposium Call Center Server.

For information on using or administering other tools and features of Symposium Call Center Server, refer to the appropriate document. To find out which document you need, see “Related documents” on page 19.

This chapter describes the major components of the Symposium Call Center Server architecture for each switch type. The supported switch types are the following:

- Meridian 1 (includes Meridian 1 IE and Succession 1000)
- DMS (includes MSL-100, Succession 2000, and CS 2100)

For Meridian 1/Succession 1000 systems with the optional Network Skill-Based Routing (NSBR) feature, this guide also illustrates the major components in a Network Control Center (NCC) setup.

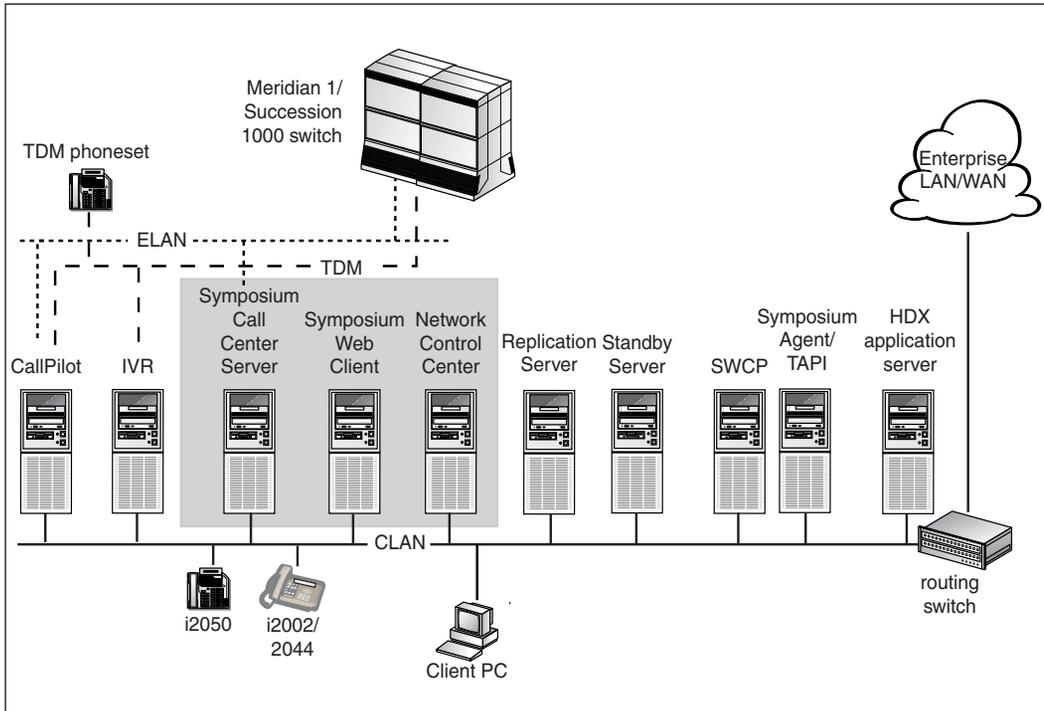
Note: The term *switch* may be used in this document as a generic term to refer to any of the previously specified telephony platforms that interoperate with Symposium Call Center Server.

Meridian 1/Succession 1000

Symposium Call Center Server components

Symposium Call Center Server consists of a number of core components, as shown in the following illustration. This guide focuses primarily on Symposium Call Center Server and Symposium Call Center Web Client but, where appropriate, it provides references to other documentation.

The following illustration shows a contact center based on a Meridian 1 or Succession 1000 switch. For an illustration of a contact center with a DMS switch, see “DMS” on page 34.



Telephony component

The telephony component is made up of the phonesets and the switch. On the Succession 1000 switch, the telephony component is purely IP-based; on the Meridian 1 switch, it is a more traditional TDM-based solution. Hybrid solutions can be deployed for businesses that want to adopt a more evolutionary approach to IP telephony rollout.

Contact center server components

The server components consist of the following elements:

- **Symposium Call Center Server:** The core contact center component, which provides intelligent call routing capability. This server operates under Windows 2000 Server or Windows 2000 Advanced Server and runs the Symposium Call Center Server application software. Symposium Call Center Server allows you to identify each agent's unique abilities, or skillsets. All calls arriving at the switch are routed to the agent with the appropriate skillset. Rules for call treatment and routing can be simple or complex.
- **Symposium Web Client application server:** A server that provides browser-based access to the contact center for administrators and supervisors. This server operates under Windows 2000, Windows 2000 Advanced Server, or Windows Server 2003 (Enterprise or Standard Edition), and runs the Symposium Call Center Web Client application server software.
- **Network Control Center (NCC) server (optional):** The server in a Symposium Call Center Server network that manages the Network Skill-Based Routing (NSBR) configuration and communication between servers. This server is required when multiple servers in Symposium Call Center Server are networked and operating as a single distributed contact center. It runs the NCC software application, which is a keycoded subset of the Symposium Call Center Server application software.
- **Replication Server and standby server:** An optional component in the Symposium Call Center Server system that provides additional redundancy. The Replication Server backs up the database on the active server to a standby server, in real time. If the active server fails, the standby server can be speedily deployed.

- **Symposium Web Center Portal (SWCP) (optional):** A client/server contact center application that expands contact center e-mail capabilities to allow agents to view, respond to, and track requests over the Internet. Unlike conventional e-mail requests to a single e-mail account, Symposium Web Center Portal lists all your customers' requests, and records all your agents' responses with the initial request. This allows you to measure and control the volume of traffic from the Internet. Supervisors and administrators can view real-time displays of contact center activities, as well as run historical reports.

The agent/client interface presents the agent with a browser-based graphical user interface. Agents can use it to respond to customers' requests over the telephone, by e-mail, or over the Internet.
- **Symposium Agent (optional):** An agent productivity tool that enables contact center agents to provide intelligent and personalized customer care. Agents use a personal computer to access the agent telephony functions. Call rules and dialing plans set up by the contact center administrator determine how calls are handled. Applications can be configured to pop up on the agent's screen to assist in call processing.
- **Telephony Application Program Interface Service Provider (TAPI) (optional):** A client/server application that integrates a telephone on a user's desktop with client and server-based applications. The telephone is physically connected to a switch but is not physically connected to a client PC. You do not need any special telephones, connectors, circuit packs, or additional wiring for the client PC.
- **Host Data Exchange (HDX) application server (optional):** A host computer running a third-party provider application that receives data (such as a credit card number) from Symposium Call Center Server, and returns data (such as account balance) to Symposium Call Center Server. Symposium Call Center Server Release 5.0 includes a provider application that coresides with the server.

Call center client components

- **Symposium Web Client PCs:** Client PCs used to administer the server and to monitor contact center performance using a browser-based interface. The number of these computers is usually proportional to the total number of agents in the contact center.

- **Classic Client PCs:** Client PCs used historically to administer the server and to monitor contact center performance. The Classic Client is now being superseded by the Symposium Web Client client/server architecture. However, it still exists as a product. Typically a contact center will have a very small number of these (one or two) to access functions that are not available in Symposium Web Client.

Voice services components

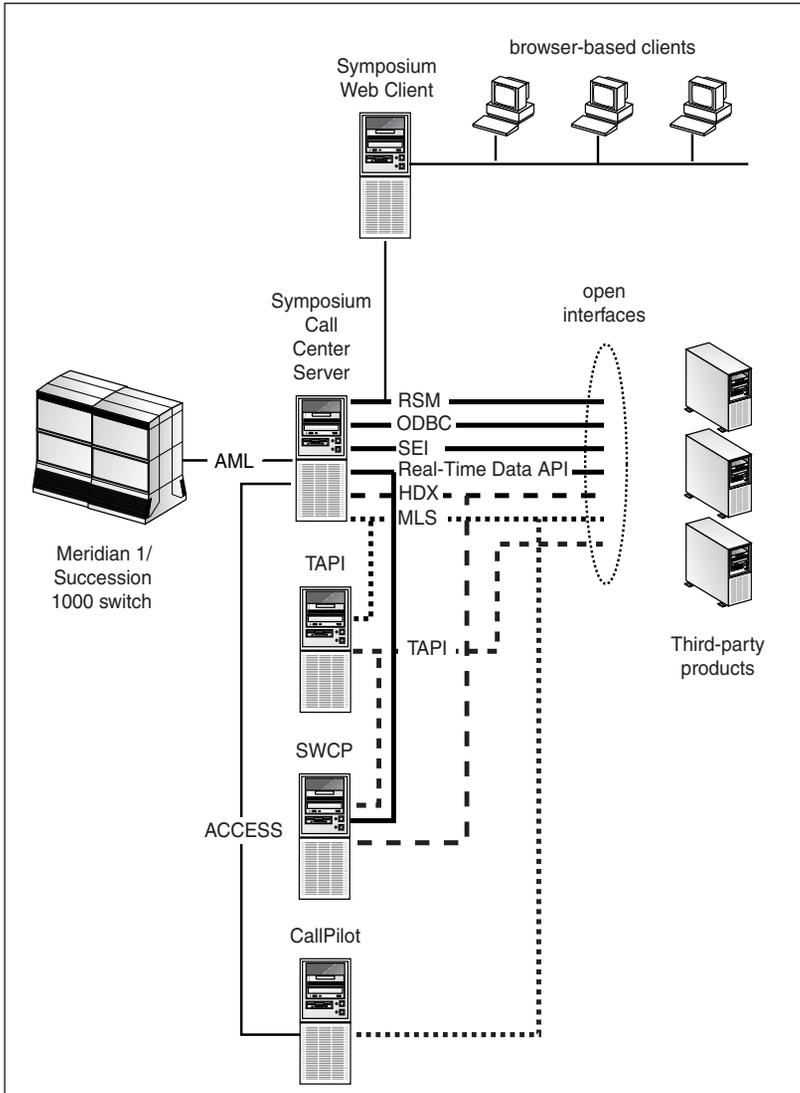
- **CallPilot/Meridian Mail:** Voice mail systems that can be used to provide front-end IVR or voice services to Symposium Call Center Server. If they are used for voice services—either Give IVR or ACCESS (Open/Close Voice Session, Give Controlled Broadcast, or Collect Digits)—the voice ports on these voice services platforms must be dedicated. (Symposium Call Center Server has direct access to them.)
Note: Meridian Mail is not supported with Succession 1000.
- **Interactive Voice Response (IVR):** An application that allows telephone callers to interact with a host computer using prerecorded messages and prompts. You can use Nortel Networks IVR or third-party IVR systems to provide front-end IVR to calls before they are handed over to Symposium Call Center Server.

Network infrastructure

- **Customer Local Area Network (CLAN):** The LAN to which your corporate services and resources connect. The Symposium Call Center Server client and server both connect to the CLAN. Third-party applications that interface with the server also connect to this LAN.
- **Embedded Local Area Network (ELAN):** A dedicated Ethernet TCP/IP LAN that connects the switch and the servers in Symposium Call Center Server.
- **Wide Area Network (WAN):** A computer network that spans a relatively large geographical area. Typically, a WAN consists of two or more local area networks (LANs). A WAN is required for Network Skill-Based Routing.

Contact center logical representation

The following illustration shows the contact center from a more logical perspective. The key to understanding how the contact center operates is a basic understanding of the interfaces used by the various components to interact with each other.



Meridian Link Services (MLS): A two-way communications facility that provides the interface between external computer applications and the Private Branch Exchange (PBX), to achieve computer-telephony integration (CTI). Meridian Link Services is a protocol exported as part of Symposium Call Center Server. An example of an MLS application is an inbound telemarketing contact center, where MLS provides the Calling Line ID (CLID) and Dialed Number Identification Service (DNIS) information from an incoming call to a third-party application. The application can use this information to retrieve data—both customer and product information—from a database, and present it to the agent’s PC before the call is even answered.

Host Data Exchange (HDX): A rich scripting language provided with Symposium Call Center Server to control treatment of calls. The scripting language can send information to, and request information from, a provider application (such as a database) over the Host Data Exchange (HDX) interface. You can use three commands (Send Info, Send Request, and Get Response) in a call script to interface with a third-party application to obtain information and influence the script operation. For example, Symposium Call Center Server can send the CLID to the provider application to determine whether the caller is a priority or regular customer; based on the response from the application, it queues the call appropriately. Similarly, Symposium Call Center Server can request caller-entered information from an IVR system for use in call handling or skillset selection, and then pass this information to a CTI application to ensure that an appropriate screen of information is presented when the call reaches an agent.

Open Database Connectivity (ODBC): A standard database interface that allows open access to the relational database in which Symposium Call Center Server stores its historical data. It is used with Structured Query Language (SQL), a standard language that allows data to be extracted from the database. With these two interfaces, customers can use any ODBC-enabled application (such as a report writer) to manipulate the historical data in the database. Therefore, contact center managers can use industry-standard report writers and open database connectivity to merge valuable contact center information with other corporate data for a complete view of their customer relationships.

Real-Time Statistics Multicast (RSM) and Real-Time Data Application Programming Interface (RTD API): Interfaces that provide real-time information to third-party applications. These applications can then display real-time statistics on wallboards and agent desktop displays, or write custom formulas.

Symposium Event Interface (SEI): An interface that provides third-party vendors with the information they need to create complementary applications by providing call progress and resource events. Communication is based on a client/server paradigm. The Event Server provided by Symposium Call Center Server acts as the source (or server) of events. Event consumers are the third-party applications (or clients) connected to the Event Server. Communication between client and server is expected to be in local area networks (LANs) using a connection-based (point-to-point) protocol. This interface ensures delivery and guarantees proper sequencing of events. The Event Server maintains one connection per client.

Telephony Application Program Interface (TAPI): An interface between the switch and an application that allows the application to control the telephone on a user's desktop.

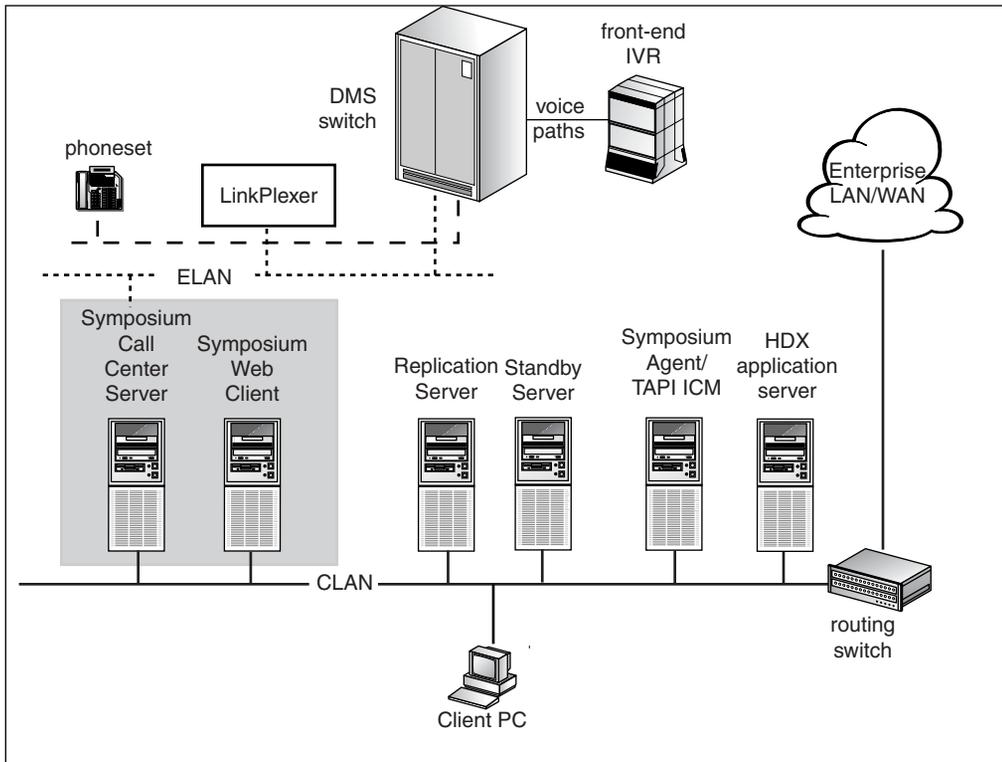
Application Module Link (AML): An internal protocol used by Symposium Call Center Server to communicate directly with the switch. Application Module Link runs over the Embedded LAN (ELAN). Symposium Call Center Server executes scripts and instructs the switch, using AML, to set up the speech paths necessary to connect calls to voice ports, agents, or RAN trunks, and to provide tone treatments (such as ringback and busy) to calls.

ACCESS: An internal protocol used by Symposium Call Center Server to directly control some of the voice services available on the CallPilot or Meridian Mail platform. When operating as a voice services system, CallPilot is used as a Voice Services Component (as opposed to a normal Messaging mode that is usually used in a non-contact center environment). Using ACCESS, a Symposium script may specify that certain announcements be played to the call on that channel, or even obtain any data entered by the caller over DTMF.

DMS

Symposium Call Center Server components

The following illustration shows a contact center based on a DMS switch. For an illustration of a contact center with a Meridian 1/Succession 1000 switch, see “Meridian 1/Succession 1000” on page 27.



Notes:

- Throughout this guide, the term DMS switch applies to the following switch types:
 - DMS Switch
 - MLS-100

- Succession 2000
- Nortel Networks Communication Server 2100 (CS 2100)
- This illustration shows the switch and Symposium Call Center Server connecting through the ELAN. For the DMS/CompuCALL, which uses the X.25 interface, the server connects to the switch through the LinkPlexer.

Differences from the Meridian 1/Succession 1000 architecture

The architecture in the DMS environment differs from the M1/Succession 1000 architecture in the following ways:

- Network Skill-Based Routing (NSBR) is not supported in the DMS environment. Therefore, the architecture does not include an NCC server.
- Front-end IVR is supported in the DMS environment, but integrated IVR (the Give IVR script command) is not supported. (Give RAN and Give Music treatments are supported.)
- Symposium Voice Services on CallPilot and Meridian Mail (Open/Close Voice Session, Give Controlled Broadcast, and Collect Digits) is not supported in the DMS environment.
- Only a subset of the MLS protocol is available.

LinkPlexer

LinkPlexer is a Windows 2000/NT application. For a DMS switch that uses the X.25 protocol, LinkPlexer is required to interface between Symposium Call Center Server and the switch.

Optionally, LinkPlexer can be used for an IP switch, to allow multiple applications to share the same switch resources. The CompuCALL/ICM interface limits the association of switch DNs to a single session. This means that the same DN cannot be associated with two host applications that have simultaneous ICM application sessions with a switch (whether they are on the same or different physical hosts). Linkplexer is used to overcome this limitation.

Chapter 3

Engineering Symposium Call Center Server

In this chapter

Overview	38
Hardware configurations	39
Call load	42
MLS and HDX performance impact	46
Guidelines to minimize capacity requirements	52
Capacity estimation	55

Overview

Introduction

You must ensure that the platform on which you plan to install Symposium Call Center Server satisfies the capacity requirements of your contact center. To help you do so, Nortel Networks provides the Capacity Assessment Tool (CapTool).

CapTool

CapTool is a stand-alone MS Windows software application used to determine the processor capacity requirements of the following components:

- Symposium Call Center Server
- Network Control Center (NCC) server
- Symposium Web Client application server

As well, it can estimate

- the number of voice ports required for a specified call complexity and call load
- the traffic impact on the CLAN due to real-time data, reporting, and other data-intensive activities
- (in a networked environment) the WAN bandwidth requirements between the local site and all remote sites due to network data traffic

After you enter specifications for contact center parameters, CapTool uses mathematical models to estimate the performance and capacity of the required components.

The quality of the results obtained from the tool is directly proportional to the quality of the input received from the user. To use CapTool effectively, you must ensure that the input is as accurate as possible.

Hardware configurations

Hardware platforms

Symposium Call Center Server Release 5.0 is a software-only contact center solution, which operates on any hardware platform that meets specified requirements. This solution is referred to as Platform Vendor Independence (PVI).

Symposium Call Center Server supports two of the High Availability platforms available from Stratus.

Platform Vendor Independence

Symposium Call Center Server does not require Nortel Networks-supplied hardware. It runs on any hardware platform with

- an Intel Pentium CPU
- Windows 2000 Server or Windows 2000 Advanced Server operating system and Microsoft-certified drivers
- CPU speed, RAM, hard drive capacity, and hard drive speed that satisfies the capacity requirements of the contact center

Note: You can use the CapTool application to determine the CPU speed and hard drive capacity required for your configuration.

For more detailed information about server requirements, see the *Symposium Portfolio Server And Operating System Requirements* document, available on the Partner Information Center web site, in the location

Products by Brand (Documentation) / Symposium Call Center Server 5.0/
Technical Guides and Reference

High Availability servers

Symposium Call Center Server runs on two High Availability servers available from Stratus:

- ftServer 3220
- ftServer 3300

Each of these servers is available as

- a strictly hardware high availability solution (H/W RAID)
- a hybrid software/hardware solution (Hardware Assisted Software Mirroring [HASM] or Hardware Assisted Software RAID). This option offers high availability at lower cost.

CPU requirements

Symposium Call Center Server requires a processor from the Intel Pentium suite. For optimal performance, average CPU utilization should not exceed 50 percent for any 15-minute period.

Note: It is expected and normal for CPU utilization to exceed 50 percent (with utilization as high as 100 percent) for short periods.

As contact center size and call loads increase, the speed of the processor required to maintain average CPU utilization below 50 percent will increase. The system capacity as a function of processor speed is given in the tables in “Capacity estimation” on page 55.

Memory requirements

RAM requirements

Symposium Call Center Server requires a minimum of 512 Mbytes of RAM. Large contact centers may require additional RAM.

To determine whether the amount of memory on your platform is adequate for your workload, use the Windows Performance Monitor. During steady state operation, the average value of the pages per second counter for a 20-minute period should not exceed 5. If it does, increase RAM and adjust the paging file size (see the following section).

You can use more than the recommended amount of RAM, but if you do, you must allow additional disk space to accommodate the increase in size of the paging file (see the following section).

Paging file

The following table shows the default paging file sizes set during the Windows 2000 installation:

Server RAM size	Minimum paging file size	Maximum paging file size
Less than 2 Gbytes	1.5*RAM	2*RAM
2 Gbytes or greater	2 Gbytes	2 Gbytes

For a system with 512 Mbytes of RAM, the default minimum paging file size is 768 Mbytes and the default maximum paging file size is 1 Gbyte.

To optimize performance, Microsoft recommends that the minimum paging file size equal the maximum paging file size. Therefore, Nortel Networks advises that both the minimum and maximum paging file sizes be set to 1.5 * RAM. For 512 Mbytes of RAM, therefore, the paging file requires 768 Mbytes of disk space.

Note: The maximum size allowed for one paging file is 4.095 Gbytes. To overcome this limit, you can use multiple paging files. For detailed instructions on how to set up this configuration, see the article “How to Overcome 4,095-Gbytes Paging File Size Limit in Windows” in the Microsoft Knowledge Base at www.microsoft.com.

Call load

Introduction

Together, call complexity and call rate determine the resource requirements (CPU, memory, and so on) due to the call load.

Call complexity

Call complexity is the number of each type of service used by a call.

Expected resource consumption

Over a period of time, the average number of each type of service per call can be used to estimate the expected resource consumption. For example, if a typical call is queued to an average of two skillsets, then the expected resource cost per call is two times the cost of queueing a call to one skillset (provided that the costs are a linear function of call rate).

Cost of a basic call

To be able to estimate the resource consumption on Symposium Call Center Server for different call rates, we must define the cost of a basic call, as well as the costs associated with the most typical call operations. These costs have been measured, and are incorporated in the CapTool calculations.

Note: The cost of a basic call is the resource consumption incurred due to basic call processing (assuming that the agent answers immediately).

The following table lists the most often used call services, and indicates the typical number used per call in the hybrid call model:

	Meridian 1/ Succession 1000	DMS
Parameter	Services per call	Services per call
Basic Call	1	1
Queue to Skillset	2	2.2

	Meridian 1/ Succession 1000	DMS
Parameter	Services per call	Services per call
Queue to Agent	0	0.1
Give Controlled Broadcast (S/S)	1	N/A
Voice Services Collect Digits	0	N/A
Give IVR	1	N/A
Give RAN	2	1
Give Music	1	1.5
HDX Send Info	1	1
HDX Request/Response	1	0
Intrinsics	5	5
If / Then's Executed	5	4
Proportion of Calls Transferred	5%	5%
Proportion of Calls Conferenced	5%	15%
Proportion of Calls Transferred to a DN	N/A	10%
MLS Screen Pops	1.2	1.2
MLS Messages	0	0
Queue to Network Skillset	2	N/A
Proportion of Calls Networked	10%	N/A

Note: For a description of call models, see Appendix B, “Standard call models.”

The number of services per call is an average value taken over all inbound (or outbound, if that is the context) calls. See the examples presented in the section, “MLS and HDX performance impact” on page 46.

Call rate

Call rate is the average rate of calls processed by the server. The call rate is measured in Calls Per Hour (CPH) and is a function of the average Call Arrival Rate and Mean Holding Time (MHT).

Note: Mean Holding Time is the time that the agent is involved in serving a call. It is the sum of

- average talk time
- time required for post-call processing, when the agent is not available to handle other calls
- inter-call interval (including union break time, if any)

Under heavy call loading, or during the busy time, when there is no agent idle time, Mean Holding Time is equal to Mean Time Between Calls (MTBC). (These definitions apply to both inbound and outbound calls.)

Call rate, number of active agents, and MHT are related: given the same call rate, the more agents there are, the longer the MHT can be. For example, if the call rate is 60 CPH and only one agent is available, then the MHT cannot be more than one minute. On the other hand, if there are 60 agents for the same call rate, then each agent can take up to an hour, on average, for a call. The following table shows the maximum possible MHT for each inbound and outbound combination:

Inbound and outbound call rate (CPH)	Maximum allowable mean holding time (seconds)				
	Entry (20 agents)	Small (100 agents)	Medium (200 agents)	Large (500 agents)	UprEnd (1500 agents)
1000	72	360	720	1800	5400
5000	14	72	144	360	1080
10 000	7	36	72	180	540

Inbound and outbound call rate (CPH)	Maximum allowable mean holding time (seconds)				
	Entry (20 agents)	Small (100 agents)	Medium (200 agents)	Large (500 agents)	UprEnd (1500 agents)
15 000	5	24	48	120	360
20 000	4	18	36	90	270
25 000	3	14	29	72	216
30 000	2	12	24	60	180
35 000	2	10	21	51	154

The values in the preceding table are based on successful call terminations (for example, treatment, available agent, call servicing, call termination), and do not take into account agent activity other than call handling. This table should be used to estimate combinations of call rates and workloads that would be “reasonable.” (For example, 20 Agents [Entry workload] handling 25 000 CPH will spend, on average, no more than 3 seconds per call. This is probably unreasonable for a human agent, but may be acceptable for an automated voting application).

MLS and HDX performance impact

Introduction

Symposium Call Center Server services also impact performance. This section describes the performance of two services, for which many contact centers require detailed information.

Meridian Link Services

Meridian Link Services (MLS) is an intelligent signaling link offering computer-telephony integration (CTI) applications access to Meridian 1/Succession 1000 call processing functions.

CTI applications

Many contact center customers have a requirement for third-party CTI applications that utilize MLS. Examples of these applications include software phones, Outbound Predictive Dialing, Host Enhanced Routing, and CTI applications such as Symposium Agent.

CPU impact

CapTool helps determine the impact of MLS on the Symposium Call Center Server performance. CapTool calculates the CPU impact of issuing passive screen-pops, as well as the general impact of MLS usage by applications.

Every CTI application that interfaces with MLS sends messages to and receives messages from the switch. The MLS software on the server takes messages from the application en route to the switch and translates them into the protocol understood by the switch, namely the Application Module Link (AML) protocol. Conversely, messages from the switch en route to the application are translated from the AML protocol to the Meridian Link protocol by the MLS software. The Symposium Call Center Server CPU impact, therefore, depends on the rate of messages exchanged between the switch and the application. This message rate is a function of the application and must be known to calculate the CPU usage. The CapTool user will have to determine the average number of MLS messages per call for the MLS applications being considered.

Example

Consider a predictive dialing application having the following message profile:

Message number	From application to switch	From switch to application
1	MakeCall	
2		Progress (Trunk seized)
3		Progress (Answered)
4	InitiateTransfer	
5		Progress (Ringing)
6		CallOffered
7	Answer	
8		Progress (Answered)
9		AnswerIndication
10		Answer Response
11	CompleteTransfer	
12		Progress (Transfer complete)
13	Release	
14		Release Response

If all outbound calls use this application, the number of MLS messages processed per outbound call is 14. To include the impact due to this application in the CapTool model, enter 14 in the Number of MLS messages per outbound call box on the MLS Services input page.

If only 75 percent of the outbound calls use this application, and the remaining 25 percent use another MLS application with an average of 20 messages per call, the overall average number of MLS messages per call is

$$(0.75*14) + (0.25*20) = 15.5$$

In this case, enter 15.5 in the Number of MLS messages per outbound call box.

CLAN impact

To calculate CLAN impact, CapTool requires the average message length. (If you do not know the average message length, use 50 bytes per message.) To calculate the average message length for the example, consider the following table:

Message type	Message length (bytes)	Number per call	Effective length (bytes)
MakeCall	46	1	46
Progress (Trunk seized)	52	1	52
Progress (Answered)	49	2	98
InitiateTransfer	50	1	55
Progress (Ringing)	54	1	54
CallOffered	36	1	42
Answer	28	1	28
AnswerIndication	63	1	63
Answer Response	28	1	41
Complete Transfer	52	1	52
Progress (Transfer complete)	52	1	52
Release	51	1	51
Release Response	50	1	52
Total		14	686
Average			49

Note: The message lengths in this example do not represent real data.

The average length per call serviced by MLS is 49 bytes. If all calls receive MLS service, enter this value into the MLS message size box on the MLS Services input page.

Host Data Exchange

The host data exchange (HDX) server allows the values of script variables to be sent to or received from a third-party provider application.

Notes:

- In Release 5.0, Nortel Networks provides a provider application that can coreside with Symposium Call Center Server. The Database Integration Wizard provides an easy-to-use tool for configuring and customizing the Nortel Networks provider application. For more information, see the *Symposium Database Integration User's Guide*.
- A provider application may reside on a third-party host computer, and, therefore, is often referred to as a host application.

For example, a script can

1. obtain a credit card number from a caller using IVR
2. query the provider application using the HDX API to determine the account balance of the caller
3. use the account balance as a variable in the script

An API known as the service provider API allows a Symposium Call Center Server user to write custom applications (provider applications) that register with the HDX server to handle back-end processing for the script elements.

Two service elements can be invoked in the script:

- Send Info
- Send Request/Get Response

The Send Info command sends data to the provider application or the HDX server. The Send Request/Get Response command sends information to and receives information back from the provider application. The Send Request/Get Response operation uses approximately twice as much CPU resources as the Send Info operation.

CapTool can estimate the CPU and CLAN load. On the Call Complexity input page, enter the average number of Send Info and Send Request/Get Response commands issued per call.

Note: This is the average value taken over all incoming calls.

Example

Suppose that the call rate is 20 000 CPH during the peak hour. Suppose further that 40 percent of incoming calls are treated with the HDX service, and of these calls

- 20 percent use one Send Info command
- 20 percent use two Send Info commands
- 30 percent use one Send Info and one Send Request/Get Response command
- 30 percent use one Send Request/Get Response command

The average number of Send Info commands issued per incoming call is

$$0.4 * (0.2 + 0.2 * 2 + 0.3) = .36$$

The average number of Send Request/Get Response commands issued per incoming call is

$$0.4 * 0.3 * 2 = 0.24$$

Enter these values into the appropriate boxes on the Call Complexity property sheet.

Cautions

If the provider application is running on a slow platform, or if it is running on the same platform as other CPU-intensive applications, it may not be able to handle the HDX messages (Send Request) fast enough. As a result, a high volume of messages may become queued in the HDX server. If the queue reaches its size limit, the HDX server terminates the provider session. When this situation occurs, the provider application receives a “DXM_SERVER_SHUTDOWN” message from the API.

As a result, the generation of a DXM_SERVER_SHUTDOWN message should be interpreted to mean either of the following:

- The session is terminated because the provider application is too slow to respond.
- The communication is down because the HDX server is terminated.

If the message results from the first reason above, either reduce the incoming Symposium Call Center Server call rate or run the provider application alone on a faster computer.

Guidelines to minimize capacity requirements

Introduction

The engineering models used to calculate the capacity requirements of your contact center assume that you follow certain guidelines to minimize the load on your server.

Steady state operation

Steady state refers to an operational state in which average values of the capacity parameters do not change with time. For example, CPU utilization may vary widely at different consecutive time instances; however, if we examine the average values of CPU utilization taken over consecutive 20-minute intervals, during a period of steady state operation, these average values are approximately the same.

Guidelines for steady state operation

To ensure trouble-free operation of the server, adhere to the following guidelines for steady state operation:

- **Processor CPU:** *Average CPU utilization* over any 20-minute period during the peak hour under steady state operation must not exceed 50 percent.
- **Server RAM memory:** *Average pages per second* (found in the Memory Object of the Performance Monitor) over any 20-minute period during the peak hour under steady state operation must not exceed 5.
- **Server virtual memory:** *Committed Bytes* (found in the Memory Object of the Performance Monitor) must not exceed 90 percent of the *Commit Limit* (also found in the Memory Object of the Performance Monitor).
- **Physical and virtual memory:** The Microsoft recommendations for physical RAM and virtual memory sizing must be adhered to for optimal performance. For more information, see “Memory requirements” on page 40.

- **CLAN traffic:** Average CLAN utilization must not exceed the limit specified in “CLAN requirements” on page 92.
- **ELAN traffic:** Average ELAN utilization must not exceed the limit specified in “ELAN requirements” on page 94.

Guidelines for non-steady state operations

A number of non-steady state processes can have a significant impact on the steady state call processing activity of the server. To minimize their impact, Nortel Networks recommends a number of restrictions:

- All non-steady state processes
 - Run only one non-steady state process at any given time.
 - Do not run non-steady state processes between midnight and 12:30 a.m. During this time, the Historical Data Manager (HDM) service performs data consolidation for monthly, weekly, and daily data. CPU usage for this activity is high.
- Activation of the Master script
 - Do not activate the Master script during a busy period.
 - If you must activate the Master script during a busy period, activate all primary and secondary scripts first.

Note: If the server is not performing call processing, you can activate the Master script without first activating the primary and secondary scripts.

- Validation of large scripts
Do not validate the Master script or any large script during a busy period.
- Agent-to-supervisor assignments
Do not run multiple agent-to-supervisor assignments concurrently.
- Agent-to-skillset assignments
Do not run multiple agent-to-skillset assignments concurrently.
- Generation of large reports
Generate large reports one after the other rather than concurrently.
- Extraction of large amounts of data from the database
Generate large data extractions one after the other rather than concurrently.
- En masse logon and logoff of agents

Spread agent logon/logoff activity over a 5- to 15-minute period, and do not perform this activity during the peak busy hour.

- Database backup

Perform database backups during off-peak hours.

- Checking files for viruses

Perform this activity during off-peak hours. For more details, see the *Symposium Portfolio Server And Operating System Requirements* document, available on the PIC web site, in the location

Products by Brand (Documentation) / Symposium Call Center Server 5.0/Technical Guides and Reference

Capacity estimation

Introduction

The following tables show how the Symposium Call Center Server capacity varies with different call loads and standard workloads. The performance metrics are the outputs from the capacity models (the same ones used in the CapTool), which are based on controlled measurements (calibration measurements), as well as high-capacity testing validation results.

The tables have been constructed using capacity model extrapolations. The capacity models are based on Symposium Call Center Server Release measurements.

Rated capacity for call processing for different processors

Rated capacity is the maximum load that can be sustained at steady state such that the average CPU utilization does not exceed 50 percent. The capacity limits for different hardware platforms and different Mean Holding Times (MHT) are shown in the following table. For these calculations, the Classic Client is assumed to be the only client (Symposium Web Client is not in use), having RTD refresh rates of 3 seconds for the agent screen and 10 seconds for skillset, application, and contact center summary screens. It is assumed that networking is not enabled and that RSM is turned off for these capacity estimations. The call complexity model is assumed to be the standard one given in “Call complexity” on page 42.

Processor	MHT (minutes)	Agents	Peak call rate (CPH)
PIII733	2	536	16 080
	3	683	13 660
	4	791	11 865

Processor	MHT (minutes)	Agents	Peak call rate (CPH)
PIII1.0B	2	696	20 880
	3	886	17 720
	4	1026	15 390
Xeon 1.5 GHz/ PIV1.5 G	2	1000	30 000
	3	1275	25 500
	4	1473	22 095

Peak call rates for different processors and real-time display refresh rates

The following table demonstrates the impact of different real-time display refresh rates on the capacity limits of the different processors. The refresh rates used are for the agent screen and the other three common screens, namely the application screen, skillsets screen, and contact center summary screen. The Mean Holding Time (MHT) is assumed to be 3 minutes. The convention used is Agent Refresh Interval in minutes/Other Displays Refresh Interval in minutes.

Processor	Refresh rate (seconds)	Agents	Peak call rate (CPH)
PIII733	2/3	415	8298
	3/5	531	10 628
	4/10	691	13 818
PIII1.0G	2/3	527	10 532
	3/5	683	13 661
	4/10	897	17 939

Processor	Refresh rate (seconds)	Agents	Peak call rate (CPH)
Xeon 1.5 GHz/ PIV1.5 G	2/3	741	14 817
	3/5	973	19 461
	4/10	1289	25 786

Peak call rates for different processors and percentage of calls networked

The following table presents the capacity impact under networking of different percentages of networked calls. The performance limits for different hardware platforms are shown in the following table. The Mean Holding Time is assumed to be 3 minutes. The number of remote sites is six.

Processor	Percentage of calls networked	Agents	Peak call rate (CPH)
PIII733	10	567	11 342
	30	418	8350
	50	324	6485
PIII1.0G	10	737	14 741
	30	544	10 870
	50	422	8448
Xeon 1.5 GHz/ PIV1.5 G	10	1062	21 237
	30	785	15 700
	50	611	12 224

Peak sustainable capacity

The next table shows the upper limit on processing calls with the hybrid call model described in Appendix B, “Standard call models.”

	Peak capacity per workload summary (CPH)				
Processor	Entry	Small	Medium	Large	UprEnd
PIII733	23 750	23 100	21 900	17 700	N/A
PIII1.0B	31 400	30 700	29 350	25 100	7200
Xeon 1.5 GHz/ PIV1.5 G	46 200	45 400	43 900	39 500	21 400

Chapter 4

Engineering the client

In this chapter

Overview	60
Section A: Symposium Web Client	63
Section B: Classic Client	75

Overview

Introduction

Two client applications are available for administration of Symposium Call Center Server: Symposium Web Client and the Classic Client.

Symposium Web Client

Symposium Web Client is a browser-based thin client for administrators and supervisors using Symposium Call Center Server Release 4.0 or higher.

Since Symposium Web Client uses IP multicasting to distribute real-time display data, it generates less LAN traffic, minimizes server load, and reduces waits for real-time data. (For more information about IP multicasting, see Appendix C, “IP Multicast Networking.”)

Symposium Web Client provides access to the following administration functions:

- Contact Center Management
- Access and Partition Management
- Configuration
- Scripting
- Real-Time Reporting
- Historical Reporting
- Emergency Help
- Audit Trail

Note: Symposium Web Client does not provide access to the server maintenance utilities provided by the Classic Client (see the following section).

Symposium Web Client also includes a separate tool, Agent Desktop Displays, which provides agents with real-time skillset monitoring.

Classic Client

In addition to the administration functions that are available on Symposium Web Client, the Classic Client also provides access to the server maintenance utilities:

- Backup/Restore
- Alarm Monitor
- Event Browser
- Server Performance Monitor
- Voice Prompt Editor

To access these functions, you must use the Classic Client.

Note: Release 5.0 networking features cannot be configured with the Classic Client.

Section A: Symposium Web Client

In this section

Architecture	64
Requirements	65
Application server performance	66
Client performance	72
Minimizing CPU load	73

Architecture

Introduction

Symposium Web Client uses a three-tiered Internet-based architecture with functionality distributed among various components. The major components of Symposium Web Client include the following:

- **Application server:** The middle layer that communicates with Symposium Call Center Server and makes information available to the client PCs.
- **Client PCs:** Employ a web-based browser to interface with the application server. They are used to administer the server and to monitor contact center performance.
- **Symposium Call Center Server:** Responsible for functions such as the logic for call processing, call treatment, call handling, call presentation, and the accumulation of data into historical and real-time databases.

Requirements

Application server

The application server runs on any platform with

- an Intel Pentium CPU
- Windows Server 2003 or Windows 2000 Server operating system
- Internet Explorer
- CPU speed that satisfies the capacity requirements of the contact center

Note: You can use the CapTool application to determine the CPU speed required for your configuration.

For more detailed information about application server requirements, see the *Symposium Portfolio Server And Operating System Requirements* document, available on the PIC web site.

Client PC

The client PC runs on any platform with

- an Intel Pentium CPU
- Windows Server 2003, Windows 2000 (Professional, Server, or Advanced Server), or Windows XP operating system
- Internet Explorer
- CPU speed that satisfies the capacity requirements of the contact center

Note: You can use the CapTool application to determine the CPU speed and hard drive capacity required for your configuration.

Application server performance

Application server CPU impact

For optimal performance, average CPU utilization should not exceed 70 percent over a 15-minute period. It is expected and quite normal for the CPU utilization to exceed the 70-percent limit (up to 100 percent) for short periods of time.

Model

Precise CPU measurements were carried out in Nortel Networks labs using a Pentium III 733 MHz PVI platform. The IIS component was found to be the dominant loading factor. IIS loading is measured in hits per second. The relationship between hits per second and CPU utilization is

$$U_{CPU} = U_{BG} + HitRate * Hit_Cost / 3600$$

where

- U_{CPU} is the CPU utilization on the PIII 733 MHz processor
- U_{BG} is the background CPU utilization
- $HitRate$ is the hit rate in hits per second
- Hit_Cost is the CPU cost per hit in CPU-seconds

The following table summarizes the values obtained from the measurements:

Parameter	Value	Units	Notes
U_{BG}	0.0446	% (= 4.46%)	Overall background CPU (%) on PIII733
Hit_Cost	0.037	CPU-Sec	CPU cost of each Web hit on PIII733

Note: All CPU-Sec values are in PIII733 CPU-Seconds.

Sample values

CapTool uses this model to estimate the maximum number of client PCs that can be supported by different processors. The estimates assume a maximum CPU utilization of 70 percent on a dedicated processor (only application server), and a maximum IIS hit rate per client of 17 hits per minute. The results are provided in the following table:

Processor	Maximum number of client PCs
Single processors	
PIII733	47
PIII750	47
PIII800	50
PIII866	54
PIII933	59
PIII1.0G	62
PIV1.3G	78
PIV1.5G	92
Dual Processors	
2PIII733	72
2PIII750	72
2PIII800	77
2PIII866	83
2PIII933	89
2PIII1.0G	95
2PIV1.3G	119
2PIV1.5G	139

Notes:

- This model is not affected by the number of Classic Client PCs.
- This model makes the following assumptions:
 - The ratio of agents to supervisors does not exceed 10 (that is, 10 agents per supervisor).
 - Average CPU utilization does not exceed 70 percent over at least a 15-minute period during peak usage loads.
 - The number of requests from each user to the application server does not exceed 17 per minute.
- This model assumes that the entire CPU is dedicated to IIS. Other coresident applications that may impact the CPU, such as Agent Desktop Displays, are planned for future inclusion in the model.
- To determine the processor required for an application server in your environment, use the CapTool. Based on the number of client PCs, CapTool recommends a processor and predicts the CPU impact. Alternatively, you can use the Symposium Web Client CPU utilization analysis spreadsheet, which is also available on the Partner Information Center web site.
- To determine the processor required for a client PC, use the CapTool. Based on the amount of real-time display traffic, CapTool recommends a processor and predicts the CPU impact. You can also use the Symposium Web Client CPU utilization analysis spreadsheet to obtain this information.

Additional parameters

Refresh rates: The minimum refresh rate for real-time statistics on the application server is .5 seconds. You can adjust this rate to achieve optimal balance between latency and CPU consumption.

Historical reports: The combined number of ad hoc or scheduled reports that you can generate simultaneously is limited to five. You can schedule as many historical reports as required; however, only five scheduled reports are processed simultaneously while the others wait in queue. Likewise, for ad hoc reports, only five reports can be generated at the same time. For example, five supervisors can generate an ad hoc report, but the sixth supervisor to do so receives a message saying the system could not process the request. This supervisor must try to generate the ad hoc report again later, after the first five reports have been

generated (or schedule the report to run later). This limitation applies to the *total* of the ad hoc and scheduled reports that can be generated at a particular time. For example, if two reports are scheduled to be output at noon, then only three ad hoc reports can be generated at this time, bringing the total to five.

Parameters not included in the model

The CPU utilization on the application server may be impacted by the following parameters, which are not accounted for in the preceding model:

- scheduled historical reports
- antivirus scanning
- backup/restore procedures

Multiple application servers

It is possible to split Symposium Web Client users across multiple application servers. When using the CPU model, each application server must be handled individually to determine the CPU loading on each one.

Symposium Call Center Server CPU impact

In the worst case, each IIS hit on the application server has an associated Symposium Call Center Server direct access database cost because the OAM database access API is bypassed. This additional Symposium Call Center Server CPU incurred cost has been measured to be 0.01811 CPU-seconds on a Pentium II 300 MHz, and represents the cost of connecting to and disconnecting from the database.

The cost of data extraction is already accounted for in the Symposium Call Center Server CPU model.

Application server CLAN/WAN impact

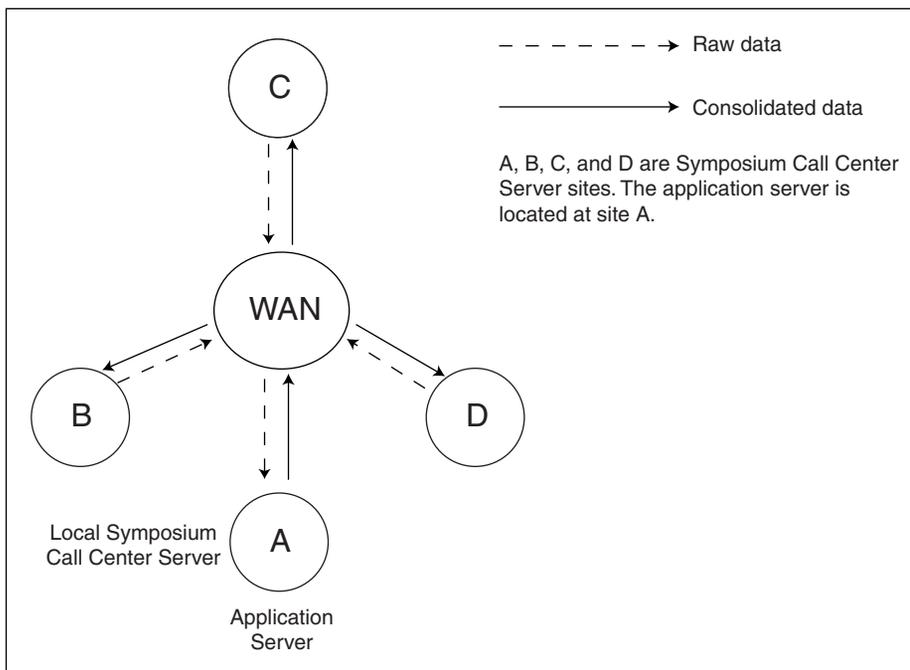
The LAN/WAN impact from the application server can be divided into two parts, as shown in the following illustration:

- RSM multicast data sent from Symposium Call Center Server to the application server
- Consolidated Real-Time Display (CRTD) data

The application server consolidates multicast traffic into a single stream, and sends it to the client PCs in either multicast or unicast format.

Note: Since the unicast option has a significant impact on network bandwidth requirements and CPU usage, Nortel Networks recommends that you use multicast connections if possible.

In a network environment, the application server consolidates traffic from multiple call center servers. The RSM multicast data streams may originate at local and remote sites, and may be directed to both local clients and remote clients. In this environment, the consolidated display data is known as Networked Consolidated Real-Time Display (NCRTD) data.



(N)CRTD multicast characterization

The inputs required to characterize the (N)CRTD multicast traffic are

- Send rates (time intervals in seconds) for each of the following statistics:
 - Agent
 - Application
 - Skillset

- Nodal
 - IVR
 - Route
 - The number configured for each of
 - Active agents
 - Applications
 - Skillsets
 - IVR queues
 - Routes
- Note:** Number of nodes is always equal to 1.
- The number of data streams sent for each of the above statistics. This value is 0, 1, or 2 for each type of statistic. The two types of data streams are Moving Window and Interval-to-date.

(N)CRTD unicast characterization

The inputs required to characterize unicast traffic are the same as those for multicast traffic, with the following additional input: number of unicast connections for each type of statistic (Agent, Application, Skillset, Nodal, IVR, and Route). A separate unicast data stream is required for each unique unicast display on each client. The number of possible unique displays per client is 12—six for Moving Window statistics and six for Interval-to-date statistics. If more than one identical display for a particular statistic type is required on a given client, then only one unicast stream is sent for both.

For example, if two Agent/Moving Window displays are opened by the same client, only one Agent/Moving Window data stream is sent. However, if another client PC opens an Agent/Moving Window data stream, a new unicast stream is sent from the server. Two identical streams are open at this point.

Client performance

Symposium Web Client CPU impact

The largest impact on CPU performance on the thin client are the real-time displays. The input parameters used in calculating Symposium Web Client CPU requirements are

- the refresh rate (assumed identical for each display)
- the number of lines being displayed (over all displays, including fixed header rows)

Minimizing CPU load

Application server

- Reduce real-time display refresh rates.
- Stagger scheduled historical reports so that they are not scheduled to run at the same time.
- Schedule large reports to run at off-peak hours.
- Schedule antivirus scanning to occur at off-peak hours.
- Perform backup/restore procedures at off-peak hours.

Client PC

- Reduce real-time display refresh rates.
- Configure the client to display less data by using data partitioning and filtering.

Note: If the parameters are exceeded, then you can use more than one application server, and you can split Symposium Web Client users across the multiple application servers.

Section B: Classic Client

In this section

Classic Client configuration

76

Classic Client configuration

Introduction

You use Release 4.0 of the Classic Client application to connect to Symposium Call Center Server Release 5.0. The Classic Client PC employs the System Management Interface (SMI) Workbench.

Client configuration

The client PC runs on any platform with

- an Intel Pentium CPU
- Windows 2000 (Professional or Server) or Windows XP operating system

For more detailed information, refer to the *Installation and Maintenance Guide*.

Hard disk requirements

The client software requires 130 Mbytes of hard disk space and an additional 10 Mbytes of disk space on the hard drive where the operating system is installed. An optional set of NCC report templates requires 150 Mbytes of disk space. (These templates are only applicable in a networking environment.)

Virtual memory

Nortel Networks strongly recommends that the operating system manage the virtual memory resources of the PC. This prevents memory problems caused by insufficient disk space for swapping. Nortel Networks recommends that at least 650 Mbytes of free disk be available at run time (after all applications are loaded), on the drive where the swapfile is located.

Temporary files

The report generation process can create large temporary files in the operating system's default temporary (TEMP) directory. Reports from the Call Detail Reporting feature can create temporary files of one Gbyte or more, depending on the circumstances. The Classic Client PC must have enough capacity to generate whatever report is being run.

Chapter 5

Engineering the switch

In this chapter

Meridian 1/Succession 1000	78
DMS	84

Meridian 1/Succession 1000

Number of servers supported

A single Meridian 1/Succession 1000 can support up to three Symposium Call Center Server systems. Engineer each Symposium Call Center Server system independently of each other, but engineer the switch as a shared resource.

Note: Servers cannot share switch resources.

Supported switches

Symposium Call Center Server requires one of the following types of switches:

- Meridian 1 Options 11, 11C Mini, 51C, 61C, 81, or 81C
- Meridian 1 Internet Enabled (IE)
- Succession 1000
- Succession 1000M

Switch software versions

Symposium Call Center Server requires one of the following versions of switch software:

- X11 Release 25 or greater
- Succession 1000, Release 1.1 or 2.0
- Succession 1000M, Release 3.0

Switch capacity

The capacity of Symposium Call Center Server is a factor not only of Symposium Call Center Server itself but also of the physical capacity of the switch. The call throughput of Symposium Call Center Server depends on the following factors:

- rated capacity of the switch
- call complexity

- expected call rate

You can use the M1 Switch Capacity Spreadsheet to calculate call throughput for the Meridian 1/Succession 1000 switch. (This spreadsheet is available from the Partner Information Center web site.)

Rated capacity of the switch

The capacity of the switch is specified as the number of Equivalent Basic Calls (EBCs) per hour. An EBC is a measure of the switch CPU real time required to process a Basic Call. A Basic Call is defined as a simple unfeatured call between two 2500 sets, on the same switch, using a four-digit dialing plan.

The EBC capacity of the switch depends on the processor type, as shown in the following table:

Processor type	EBC capacity
Option 11C/CSE 1000	42 000 (TDM mode) 35 000 (IP-enabled mode)
CP2	54 782
CP3	72 000
CP4	100 800
CPP	315 000

Notes:

- For more information on switch application engineering, see the *Large System Planning and Engineering Guide* (NTP 553-3021-120).
- CP2 is not supported by Meridian 1 in IP-enabled mode, or by Succession 1000M, Release 3.0.

Call complexity

The complexity of a Symposium Call Center Server call is defined as the number of each type of service used by the call. All calls have an EBC cost, with calls of higher complexity (that is, using a greater number of services) costing more EBCs.

For example, a basic call costs 2.40 EBC; Give Music costs 0.25 EBC; Give IVR (including transfer) costs 2.29 EBC. Therefore, a call that receives IVR and Music treatments costs:

$$2.40 + 0.25 + 2.29 = 4.94 \text{ EBC}$$

To quantify levels of call complexity, Nortel Networks has defined several call models, which represent simple, complex, and front-end IVR systems (see Appendix B, “Standard call models”). You can calculate the EBC cost using the M1 Switch Capacity Spreadsheet, which is available on the PIC web site.

Simple model: Front-end IVR system

The cost of processing IVR is removed from the switch with this call model. The following table shows the number of each type of treatment per call:

Service name	Number of treatments per call
Queue to Skillset	2
Give RAN	3
Give Music	1
Intrinsics Accessed	5
If-Then-Else Executed	5

The EBC cost of this call model is 5.42 EBC.

Average complexity model: hybrid

This call model uses features from both Symposium Call Center Server and the switch. The following table shows the number of each type of treatment per call:

Service name	Number of treatments per call
Queue to Skillset	2
Give Controlled Broadcast Start/Stop	1
Give IVR	1
Give RAN	2
Give Music	1
HDX Send	1
HDX request/Response	1
Intrinsics Accessed	5
If-Then-Else Executed	5

The EBC cost of this model is 8.78 EBC.

Complex model: Symposium Voice Processing (SVP)

The following table shows the number of each type of treatment per call under this model:

Service name	Number of treatments per call
Queue to Skillset	2
Give Controlled Broadcast Start/Stop	3
Collect Digits Voice Session	1
Give IVR	1
Give RAN	1
Give Music	1

Service name	Number of treatments per call
HDX Send	1
HDX request/Response	1
Intrinsics Accessed	5
If-Then-Else Executed	5

The EBC cost of this model is 13.84 EBC.

Maximum achievable call rates

To determine the maximum achievable call rates for different switch models, all contributions resulting from the following parameters must be considered:

- the call complexity
- the MLS commands issued by CTI applications
- any other applications that may be communicating over the ELAN with the switch

You can determine the call rate by calculating the total Equivalent Basic Call (EBC) value for all incoming traffic per switch type.

Sample calculations using the Meridian 1 switch Capacity Tool

This calculation considers the Meridian 1 Option 81 with a CPP processor, which has an EBC capacity of 315 000.

CPP utilization per number of active agents per call model:

Active agents	2200	1500
Calls per hour	44 000	30 000
Symposium Voice Processing	193%	131%
Hybrid	123%	84%
Front-end IVR	76%	51%

This table implies that if there are 44 000 Symposium Voice Processing type calls then the CPP uses 193 percent of the processor capacity. In contrast, the Front-end IVR call model uses only 76 percent of the CPP for an equivalent call rate.

Note: CPU utilization greater than 100 percent is not recommended.

Meridian 1/Succession 1000M Networked ACD

The usage of Networked ACD (NACD) is transparent to Symposium Call Center Server. The call rates used in Symposium Call Center Server engineering are the total calls arriving to Symposium Call Center Server from the local switch, either directly or from Networked ACD.

Meridian 1/Succession 1000M ISDN

The ISDN circuits to the PSTN must be provisioned to handle the network call traffic to and from each switch. It is assumed that these circuits are provisioned in a similar manner to that of NACD.

DMS

Supported software loads

The following Call Center Module (CCM) software loads are supported with Symposium Call Center Server Release 5.0: CCM010, CCM011, CCM012, CCM013, CCM014, CCM015, CCM016, CCM017. Refer to the DMS engineering guidelines to properly configure the DMS switch.

Note: Throughout this guide, the term DMS switch applies to the following switch types:

- DMS Switch
- MLS-100
- Succession 2000
- Nortel Networks Communication Server 2100 (CS 2100)

Number of servers supported

A single DMS can support up to 16 Symposium Call Center Server systems. Engineer each Symposium Call Center Server system independently of each other, but engineer the DMS as a shared resource.

DMS switch requirements

To analyze the impact from one or more Symposium Call Center Server systems on the DMS switch, you must first calculate the workload on each of the servers in Symposium Call Center Server. You can then derive the workload generated against the DMS switch from each Symposium Call Center Server.

The following factors may limit the maximum call rate that can be achieved, depending on the details of a particular call processing scenario:

- For communication between the DMS switch and Symposium Call Center Server, 1024 invoke IDs are available. IDs in the range from 0 to 511 are reserved for communications sent from Symposium Call Center Server to

the switch, while the switch can use IDs in the range from 512 to 1023 to send communications to Symposium Call Center Server.

- For each ICM/SCAI link, 128 buffers are available to process incoming and outgoing messages.

Workload characterization

DMS workload from Symposium Call Center Server is described in terms of the number and types of ICM messages being sent to the DMS.

Normal operation

The types of ICM messages that Symposium Call Center Server uses during normal operations are Give_Treatment and Route_Call.

Call processing

The types of ICM treatments that Symposium Call Center Server uses during call processing operations are described in the following table:

Treatment type	ICM message
Ringback	Give_Treatment(Ringback)
RAN	Give_Treatment(RAN)
Music	Give_Treatment(Music)

The ICM messages sent to the DMS depend on the script commands that are executed by the scripts. It is assumed that each script includes the following:

- one Give_Ringback command
- one Queue_To_Skillset command
- one Quit command

This results in one GT_Cuc(Ringback) command for the Give_Ringback, and one Route_Cuc command for the completion of the Queue_To_Skillset being sent for the Basic Call. If a script starts with either the Give_RAN or Give_Music command, Symposium Call Center Server automatically sends a

Give_Ringback command. For engineering purposes, it is assumed that the script always starts with the Give_Ringback command. The mapping of Symposium Call Center Server script commands to ICM messages is summarized in the following table:

Script command	ICM messages
Basic Symposium Call Center Server call operations services	
Queue To Skillset	0
Queue To Agent	0
Give Ringback	1 - Give_Treatment(Ringback)
Give RAN	1 - Give_Treatment(RAN)
Give Music	1 - Give_Treatment(Music)
Route Call	1 - Route_Call
Data Exchange Send Info	0
Data Exchange Request / Response	0
Script Intrinsic Reference	0
“If Then Else” treatments	0

Notes:

- For the Queue To Skillset command, this model assumes that the Route_Call message sent after the Remove From Skillset command is included in the definition of the Basic Call.
- For the Queue To Agent command, this model assumes that the Route_Call message sent after the Remove From Agent command is included in the definition of the Basic Call.

The DMS workload of the predefined Symposium Call Center Server call model is based on the expected number of call services per call, and the cost of the individual script command. The following table shows the resulting number of ICM messages being sent to the DMS for each Symposium Call Center Server call:

ICM message	Number per call
Route_Call	1
Give_Treatment(Ringback)	1
Give_Treatment(Music)	User-defined
Give_Treatment(RAN)	User-defined

The DMS workload is also a function of the number of external events that happen per call. This information is summarized in the following table:

Switch event	Number per call
IVR Call Processed	1 if external IVR; 0 otherwise
Call Transfer	User-defined
Conference Call	User-defined

Number of audio routes required

Symposium Call Center Server can use up to 512 preconfigured audio routes. The audio routes are classified as either music routes or RAN routes, depending on whether the last give treatment command in the audio route is Give Music or Give RAN respectively. The DMS must have the capability to assign an audio route to every call waiting for an agent. Symposium Call Center Server supports up to 3000 waiting calls.

Number of CDNs required

The calls arriving at any Symposium Call Center Server are held in a series of CDNs. Each CDN holds up to 511 calls. Symposium Call Center Server Release 5.0 supports 1500 active agents with active calls. The server can support up to 3000 waiting calls. For 3000 waiting calls, you require at least six CDNs (3000 / 511).

Impact of MLS traffic

Symposium Call Center Server supports a third-party CTI interface. For every CTI command sent to Symposium Call Center Server per call, a corresponding command is sent to the DMS switch. To analyze this activity, CapTool analyzes the activity of each application that sends CTI commands to Symposium Call Center Server. The only MLS messages supported on the DMS switch are

- Initiate Transfer
- Complete Transfer
- Login
- Logout
- Ready
- Not Ready

For performance modeling purposes, only the transfer operations are considered, since the overall contribution due to number of agent interruptions per shift (that is, Login/Logout, Ready/not Ready) is expected to be insignificant when compared with the transfer events.

LinkPlexer and session sharing

LinkPlexer allows Symposium Call Center Server and other ICM applications (such as an external IVR system) to control the same DN on the DMS switch.

The DMS switch allows a DN to be associated with only one DMS/host session, and multiple applications must log on to the switch using separate sessions. Therefore, different applications cannot control the same DN.

LinkPlexer overcomes this limitation by opening a single session to the switch, and allowing multiple applications to use this session. In the case of an external IVR, LinkPlexer is not required if the external IVR systems uses the MLS feature of Symposium Call Center Server. However, the MLS feature only provides support for the following CTI commands: Login, Logout, Ready, Not Ready, Initiate Transfer, and Complete Transfer (the last two commands pertain to digital transfer).

Configuring the LinkPlexer system

Refer to the engineering guidelines of the LinkPlexer system to properly configure the LinkPlexer system and the DMS switch.

Chapter 6

Engineering the network

In this chapter

CLAN requirements	92
ELAN requirements	94
ELAN connection to CLAN	99
WAN requirements (Meridian 1/Succession 1000 only)	101
Recommended network equipment	103

CLAN requirements

Introduction

The Customer LAN is an Ethernet link between the server in Symposium Call Center Server and the client PCs.

CLAN traffic

CLAN traffic consists of the following elements:

- real-time display traffic
- real-time data API traffic
- Graphical Real-Time Display traffic
- Real-time Statistics Multicast traffic
- MLS traffic
- Host Data Exchange traffic
- reporting-related traffic
- Meridian 1/Succession 1000 elements
 - Event Interface traffic
 - Network Control Center traffic
 - network CBC traffic to the NCC
 - network call processing traffic
 - Network Consolidated Real-Time Display traffic
- DMS elements
 - external IVR traffic
- non-Symposium Call Center Server customer traffic

Notes:

- IVR caller-entered data (CED) can use either the CLAN or ELAN.
- Nortel Networks recommends that you use an IP router to separate the traffic on the ELAN and CLAN.

- CapTool calculates the CLAN bandwidth requirements for traffic generated by Symposium Call Center Server only. You must calculate the bandwidth required for other customer applications to determine the total bandwidth required.

Maximum acceptable utilization

Total utilization of the CLAN must not exceed 30 percent in a shared network environment. Symposium Call Center Server utilization of the CLAN can be as high as 9 percent for a system with 500 agents. Make sure that the CLAN has enough spare capacity to accommodate Symposium Call Center Server traffic in addition to customer traffic.

ELAN requirements

Introduction

The Embedded LAN is an Ethernet link between the PBX switch and the server in Symposium Call Center Server. When you implement Symposium Call Center Server onto your data network, you must keep the ELAN simple, protected, and local. (For example, you must not allow communication between Symposium Call Center Server and the switch to traverse a WAN.)

ELAN traffic

The ELAN carries the following traffic:

Meridian 1/Succession 1000 switch

- call processing AML traffic
- for Symposium Voice Services on CallPilot, ACCESS traffic

DMS switch

- call processing ICM traffic (ICM_Utilization)
- external IVR traffic (DMS only)

Note: IVR caller-entered data (CED) can use either the ELAN or CLAN.

Maximum acceptable utilization

The maximum acceptable utilization of the ELAN depends on the amount of traffic on the LAN, the length of the cable, and the size of the messages. The probability of collision of packets depends on these factors and affects the average delay within the network.

To minimize excessive network message transfer delays due to network congestion, adhere to the following guidelines:

- For a Meridian 1/Succession 1000 switch on a shared network, steady state ELAN utilization must not exceed 10 percent. (This assumes that the

switch and the server in Symposium Call Center Server are both connected locally to the ELAN.)

- For a DMS switch connected locally to a shared network, steady state ELAN utilization must not exceed 30 percent.

If the combined ELAN traffic is at or near this limit, replace the shared media hub with an Ethernet switch. For more information about configuring your ELAN, refer to *Data Networking for VoIP* (NTP 553-3001-160).

Multiple servers on a PBX switch

In an environment where multiple servers connect to the same PBX switch, each switch contributes a separate Application Module Link (AML) or Intelligent Call Manager (ICM) traffic stream. To find the total ELAN impact, add together the total ELAN bandwidth required for each server (as calculated by CapTool). Ensure that the total impact does not exceed the maximum specified in the preceding section.

Technical problems

The following sections describe the problems that can arise if the ELAN is not kept simple, protected, and local:

Propagation/queuing delays

AML or ICM traffic between the switch and Symposium Call Center Server is real-time sensitive. Network devices (such as routers and firewalls) and distance cause network propagation and queuing delays. These delays are dynamic, and at a certain threshold, cause the AML or ICM to time out and initialize.

Impact: Inability to treat calls

Lack of reliability or robustness

When an ELAN is extended across a WAN, at least three physical networks are interposed between the switch and the server.

- The ELAN was designed for mission-critical purposes. Additional network devices increase the number of potential points of failure and, therefore, increase the chances of failure of the ELAN.

- The ELAN was designed to be secure and protected. Allowing external physical connectivity exposes the ELAN to potential security threats.

Note: The Symposium Call Center Server documentation states that the ELAN must be physically and logically isolated from any other network (such as the CLAN or PSTN).

Impact: Switch initialization failure, contact center outages, unauthorized access to sensitive data

Increased maintenance and support effort

If more network devices are added to the ELANs (for example, for WAN connectivity), they will require additional

- configuration (for example, routing)
- maintenance (for example, firmware and software upgrades)
- support (it takes more time to troubleshoot a more complex network)

These additional maintenance and support activities may result in a greater number of interruptions to the communication between the switch and Symposium Call Center Server.

Impact: Contact center outages and recovery delays

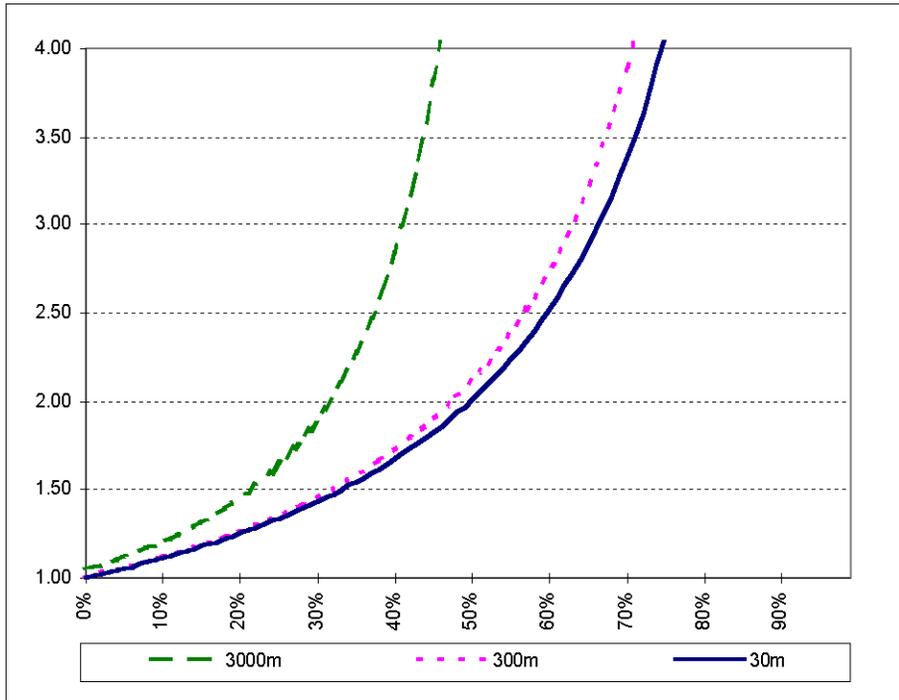
Bandwidth contention

Under the normal supported ELAN configuration, bandwidth contention on the ELAN is engineered by Nortel Networks and is not an issue. However, in an out-of-specification environment, one cannot take this for granted.

- The switch is sensitive to heavy ELAN traffic. Heavy ELAN traffic (such as broadcast storms and multicast traffic) caused by other devices on the ELAN (such as defective NICs or misconfigured devices) can cause the switch to initialize.
- Propagation delays lower the effective bandwidth availability.

The following illustration shows the relationship between the average delay factor and the LAN utilization for different wire lengths. For example, for a system located in a single room (wire length between components is under 30 meters), the delay factor is 2 whenever the LAN utilization is 50 percent. That is, it takes the data packet twice as long to travel between the components as it would on an idle system. The delay factor X effectively

reduces the LAN bandwidth by the factor of X. For example, for a delay factor of 2, the effective bandwidth of the Ethernet LAN is 5 Mbps instead of 10 Mbps. This illustration shows the performance characteristics for Ethernet.



Plan CLAN and ELAN traffic so that the delay factor is never greater than 2. Use this illustration to determine the maximum allowable utilization given the distance between Symposium Call Center Server components. For example, if the distance between the Symposium Call Center Server components is expected to be 3000 meters, then ELAN utilization should not exceed 30 percent. If all of the Symposium Call Center Server components are placed in the same building and the wire length does not exceed 300 meters, then the maximum ELAN utilization can be as high as 45 percent. CLAN utilization is estimated based on the maximum distance between Symposium Call Center Server components as well as your own components.

Impact: Switch's call handling operations and contact center outages

Other problems

Multiple groups from multiple companies manage the ELAN

In many companies, the IT group, Network group, and the Telecom group are not in the same reporting structure, or they are outsourced organizations. With a simple embedded LAN between the switch and Symposium Call Center Server, the servicing group does not have to involve all of these groups when troubleshooting ELAN problems. However, once the ELAN is connected to a WAN, multiple groups and companies are required to solve any ELAN-related problems.

Impact: Delays to ELAN network problem resolution

Process complications due to (security) policies when the ELAN requires external access (for example, a WAN)

When end customers expose their internal network to the external network, they normally apply and enforce security policies. The resulting additional security devices (firewall, VPN, and so on) add overhead and propagation delays between the switch and Symposium Call Center Server. Additional security policies also add time delays to accomplish tasks. For example, a simple IP address change (for troubleshooting purposes) requires the approval of many groups and the signatures of many managers.

Impact: Delays in maintenance activities and problem resolutions

Conclusion

The ELAN was designed as a mission-critical link between the switch and Symposium Call Center Server. Therefore, the focus is not on the average uptime, but on the single time that the ELAN could fail or cause a failure.

The goal is to keep the ELAN simple: to minimize potential points of failures and hindrances. If you connect Symposium Call Center Server and the switch with a WAN, simplicity is replaced by complexity, thus degrading the mission-critical level of the ELAN for the server and the switch.

ELAN connection to CLAN

Introduction

The ELAN is used for different purposes with different Nortel Networks products. Some products (such as OTM) use the ELAN in a standard burst-mode (transaction-based) communication, while others (such as Symposium Call Center Server) depend on the ELAN for a stream-mode (real-time based) communication.

OTM on the ELAN

In the case where there is no Symposium Call Center Server connected to the switch, the ELAN traffic may be used with Optivity Telephony Manager (OTM) for switch-management purposes. While this type of communication is considered standard data communication, the following are required when OTM is connected to the ELAN:

- Use an Ethernet switch rather than a shared-media hub.
- If the ELAN is connected to the CLAN, install a filtering router to protect the ELAN. This is to protect the ELAN from unintended traffic from the CLAN, which may, in turn, interrupt the operation of the switch.

Symposium Call Center Server on the ELAN

In the case where Symposium Call Center Server is connected to the switch, the ELAN traffic between the server and the switch is categorized as mission-critical. This is due to the link's real-time sensitivity and potential impacts to the contact center. In such a situation, the ELAN is to be protected through physical and logical isolation from any other network.

Symposium Call Center Server and OTM on the ELAN

In situations where both Symposium Call Center Server and OTM are present, extra care and precaution must be taken in the treatment of the ELAN. There are various OTM configurations to satisfy the needs of most of its features and, at the same time, to satisfy the requirement of isolating the ELAN from the CLAN.

However, depending on the features used, the networked switch configuration, the data network layout, and so on, there may be situations where the ELAN must be connected to the CLAN. In such situations, you follow the OTM's strict guidelines for filtering and routing.

The usage of the ELAN was designed and tested for inter-Nortel Networks product communications. Any communication with non-Nortel Networks equipment has not gone through Nortel Network's testing and proper engineering analysis. These external communications over the ELAN, therefore, present an unknown factor and, thereby, a potential negative impact to the overall operation of the switch and its auxiliary processors.

WAN requirements (Meridian 1/Succession 1000 only)

Introduction

A WAN can be used to provide communication between multiple nodes in the networked Symposium Call Center Server environment.

WAN traffic

The WAN carries the following types of data:

- network call processing-related traffic between the servers
- network call events recording traffic between the servers and NCC
- NCC routing table update traffic between the NCC and the individual servers
- call-by-call (CBC) and consolidated reporting traffic between the client PCs and the NCC or servers
- nodal Real-Time Display multicast data between the server in Symposium Call Center Server and the application server
- Network Consolidated Real-Time data between the Symposium Web Client application server and client PCs

Network call processing (NCP) traffic must take into account all activity during the peak busy hour of incoming calls. Reporting traffic must take into account all traffic during the period of highest reporting activity. These two times are usually mutually exclusive. (Nortel Networks recommends against running large reporting activities during the peak busy hour.)

Note: CapTool calculates the WAN bandwidth required for Symposium Call Center Server. If you use the WAN for other applications, you must determine the bandwidth required for those applications, and determine the total bandwidth required for Symposium Call Center Server and all other applications.

Dedicating the WAN network

Ideally, the WAN network is dedicated to Symposium Call Center Server call processing, although this is not always possible. In a shared WAN environment, network administrators may not have enough control over the network traffic to prevent a large file transfer from impacting other traffic, and to guarantee that latency time requirements are met. In an uncontrolled environment, it is difficult to engineer a system that meets specified performance constraints.

Timeouts

The primary factors that determine the maximum acceptable latency time of the NCP messages are the timeouts defined in the networking code. The timeout set for NCP traffic is 10 seconds. This includes the time to send a message from one node to another and receive a response. (Responses are not received for every message, but the exceptions can be ignored.) The largest NCP messages are approximately 400 bytes, including TCP/IP and Media Access Control (MAC) overhead. However, testing was done with a simulated latency time of less than 1 second. As such, it has been concluded that the maximum acceptable latency time to transmit a single message from node to node through the CLAN over a WAN connection is 1 second.

Recommended network equipment

Introduction

This section lists the layer two and layer three switches that are recommended for use on the ELAN and CLAN.

Recommendations are based on the number of ports required on the LAN.

Note: The list is up-to-date at the time of publication of this guide. However, it is constantly under review. For more current information, contact your distributor.

For small- to medium-sized installations

- BayStack 425-24T (layer 2)
- BayStack 470 (layer 2)
- Passport 1648T (layer 3)

For medium- to large-sized installations

- BayStack 470 (layer 2)
- BayStack 5510 (layer 2; gigabit capable)
- Passport 8600 (layer 3)

Chapter 7

Engineering the voice processing system (Meridian 1/Succession 1000 only)

In this chapter

Port usage	107
ACCESS requirements	109
Symposium Voice Services on CallPilot requirements	110
Symposium Voice Services on Meridian Mail requirements	111

Overview

This chapter provides information for determining the number of voice ports required to provide voice processing services on a Symposium Call Center Server system, as well as the requirements for CallPilot and Meridian Mail, if they are providing voice services to Symposium Call Center Server.

Note: Symposium Voice Services on Meridian Mail is not supported for the Succession 1000 switch.

Port usage

Introduction

The number of voice ports required depends on

- the rate of port requests
- the duration of voice sessions
- the Grade of Service (GOS)

Grade of Service refers to the probability that requests will be delayed by more than a certain number of seconds. For CallPilot and Meridian Mail, the standard GOS used is 5 percent probability that the calls will be delayed for more than 6 seconds, and 95 percent of the calls will incur a delay of less than 6 seconds.

Note: Voice ports must be dedicated to Symposium Call Center Server. They cannot be shared with other services.

ACCESS port usage

Symposium Call Center Server can support a single ACCESS connection to control voice processing. A single ACCESS connection supports up to 96 voice ports. This may limit Symposium Call Center Server performance by limiting the rate of calls that require Symposium Call Center Server control of voice processing.

Notes:

- CallPilot supports a maximum of 96 voice ports. However, one voice port must be reserved for messaging. Therefore, 95 voice ports are available to provide voice services for Symposium Call Center Server.
- None of the predefined applications (and, therefore, workloads) require controlled voice services; therefore, none of them result in ACCESS traffic.

Non-ACCESS port usage

Symposium Call Center Server voice services that do not require local voice port control (such as Give IVR) do not result in ACCESS usage and, therefore, are not subject to the 96-port limitation. Additional voice ports may be required, however, to support these services.

Note: For the predefined workloads with an SVP model (see Appendix B, “Standard workload models”), it is estimated that Symposium Call Center Server requires 64 to 96 voice ports or more, depending on the call rate.

ACCESS requirements

Symposium Call Center Server generates ACCESS traffic when it communicates with the integrated voice processing system (CallPilot or Meridian Mail) to obtain the following controlled voice services:

- Give Controlled Broadcast command
- Collect Digits command
- Open/Close Voice Session commands

For Symposium Voice Services on Meridian Mail, ACCESS traffic is transmitted over a dedicated high-speed serial connection. For Symposium Voice Services on CallPilot, ACCESS traffic travels on the ELAN.

Symposium Voice Services on CallPilot requirements

CallPilot version

Symposium Voice Services on CallPilot requires CallPilot Release 2.0 or later.

CallPilot and multiple servers on the same switch

If you are using CallPilot to provide front-end IVR, the same CallPilot server can support all three Symposium Call Center Server systems.

If you are using Symposium Voice Services on CallPilot—that is, if CallPilot is providing Give IVR or ACCESS voice services (Open/Close Voice Session, Collect Digits, and Give Controlled Broadcast), CallPilot can serve only one Symposium Call Center Server system. Therefore, each Symposium Call Center Server system must be connected to a separate CallPilot.

CPU impact

Symposium Voice Services on CallPilot uses MLS for communication between CallPilot and Symposium Call Center Server. To estimate the additional CPU load generated by Symposium Voice Services on CallPilot, use the CapTool application.

ELAN impact

For Symposium Voice Services on CallPilot, ACCESS traffic is carried on the ELAN. In an environment with Symposium Voice Services on CallPilot, the CapTool application automatically determines the additional load on the ELAN.

CLAN impact

Symposium Voice Services on CallPilot results in additional MLS traffic on the CLAN. When you use CapTool to perform a capacity assessment in an environment with Symposium Voice Services on CallPilot, the application automatically calculates the impact of the additional MLS traffic on bandwidth.

Symposium Voice Services on Meridian Mail requirements

Software release

Meridian Mail, Release 11 or later, must be used with Symposium Call Center Server.

Meridian Mail and multiple servers on a switch

If you are using Meridian Mail to provide front-end IVR, the same Meridian Mail can support all three Symposium Call Center Server systems.

If you are using Symposium Voice Services on Meridian Mail to provide IVR services (that is, with the Give IVR command), the same Meridian Mail can support all three Symposium Call Center Server systems. However, the following restrictions apply:

- You must allocate the Meridian Mail IVR ports between three IVR queues, and dedicate a queue to each server.
- All of the servers must belong to the same customer group. (Therefore, you cannot network the servers together.)

If you are using Symposium Voice Services on Meridian Mail to provide ACCESS voice services (Open/Close Voice Session, Collect Digits, and Give Controlled Broadcast), Meridian Mail can serve only one Symposium Call Center Server system.

Meridian Mail platforms

The following table shows the four Meridian Mail platforms, the number of ports available on each of these platforms, and the increments for port additions:

Meridian Mail platform	Ports	Port increments	Approx. maximum CCS	Maximum port requests at 1 min MHT	Maximum port requests at 30 sec MHT
Card Opt	2–12	2 ports	247	412	824
EC 11	4–48	4 or 8 ports	1342	2237	4474
Modular Opt/ Modular Opt GP	4–64	4 or 8 ports	1858	3097	6194
Modular EC	4–96	4 ports	2912	4853	9706

Notes:

- 36 hundred call-seconds (CCS) is the equivalent of 1 Erlang and is the amount of traffic one port can handle if it is busy all the time.
- MHT is not to be confused with call rate. A single call can create more than one port request.

ACCESS link

The bandwidth of the ACCESS link ranges from 4.8 Kbits/sec to 38.4 Kbits/sec. The maximum utilization of the link is assumed never to exceed 50 percent. The recommended ACCESS link speed is 19.2 Kbps.

Installation grounding

To avoid damage that can occur to the server in Symposium Call Center Server, the switch, or the voice processing system as a result of poor grounding, electro-optical isolators should be installed for use on the RS-232 ACCESS cable. Use this type of isolator to ensure that no surges occur during electrical disturbances.

ACCESS link utilization

The following table shows the utilization of the ACCESS link for different call rates. Each call is assumed to include one Collect Digits treatment and one Give Controlled Broadcast treatment:

Call rate	ACCESS utilization (%)
1000	2.0%
5000	9.8%
10 000	19.6%
15 000	29.3%
20 000	39.1%
25 000	48.9%
30 000	58.7%
35 000	68.5%

Note: Maximum utilization is 70 percent.

The following formula calculates utilization of the ACCESS link:

$$\text{ACC_Utilization} = 100\% * \text{ACC_BW_Required_KbitsSec} / \text{ACC_Bandwidth_KBitsSec}$$

ACCESS link capacity

The following is the computation of the maximum rate of ACCESS-related calls that the link can support for 100 percent Give Controlled Broadcast (GCB) calls and 100 percent Voice Session Collect Digits (VSCDG) calls:

$$\text{Max_AC_GCB_PerHour} = \frac{(\text{ACC_Bandwidth} * \text{ACC_Max_Utilization} * 1000 * 3600)}{(\text{GCB_Acc_Size} * 8)}$$

$$\text{Max_AC_VSCDG_PerHour} = \frac{(\text{ACC_Bandwidth} * \text{ACC_Max_Utilization} * 1000 * 3600)}{(\text{VSCDG_Acc_Size} * 8)}$$

CSL

Command and status link (CSL) traffic is used for communication between Meridian 1/Succession 1000 and Meridian Mail. CSL traffic is generated only when voice services are required for a call. CSL traffic is transmitted over a dedicated high-speed serial connection.

Note: For Symposium Voice Services on CallPilot, CSL traffic travels on the ELAN, and is included in computations of ELAN bandwidth.

CSL traffic cost

The following table shows the variables and their values used in the CSL traffic calculations.

Variable	Definition	Value
CSL_Bandwidth_KBitsSec	CSL Bandwidth (kbps)	9.6
CSL_Max_Utilization	CSL Maximum Utilization	0.7 (70%)
nGCB_Simultaneous	Average # simultaneous calls per port on GCB	User-defined

The following is the computation of the bandwidth required for CSL traffic:

$$\text{CSL_BW_Required_KbitsSec} = \frac{((\text{PeakCallRate}) / \text{nGCB_Simultaneous}) * \text{CSL_Bytes_PerSession} * \text{AvgGCBCall} * 8) / 1000}{3600}$$

The bandwidth of the CSL is 9.6 kbps. The maximum utilization of the CSL is 70 percent.

CSL utilization

The following table shows the utilization of the CSL based on workload and the call rate:

Call rate	CSL utilization per workload (%)
1000	1.2
5000	6.2

Call rate	CSL utilization per workload (%)
10 000	12.4
15 000	18.6
20 000	24.8
25 000	31.0

Note: Maximum utilization is 70 percent.

The following formula is used to calculate utilization of the CSL:

$$\text{CSL_Utilization} = 100\% * \text{CSL_BW_Required_KbitsSec} / \text{CSL_Bandwidth_KBitsSec}$$

CSL capacity

The following formula calculates the maximum rate of CSL-related calls (voice) supported by the link. Using the computations below, the maximum CSL call rate is estimated to be 56 523 calls per hour, if all calls require voice service.

$$\text{Max_CSL_Sessions_PerHour} = \frac{(\text{CSL_Bandwidth} * \text{CSL_Max_Utilization} * 1000 * 3600)}{(\text{CSL_Bytes_PerSession} * 8)}$$

NLI link

The network loop interface (NLI) link facilitates the voice path between Meridian 1/Succession 1000 and Meridian Mail. It is used only for calls requiring IVR service. The number of voice ports needed for this link is calculated based on the number of voice sessions required by Symposium Call Center Server.

Chapter 8

Setting up remote support with a VPN

In this chapter

Overview	118
Guidelines for the remote support VPN at the customer's premises	119
VPN configurations	120

Overview

Standard remote support over a direct-connect modem

If you need technical support, your distributor or Nortel Networks customer support staff must be able to connect to your server. To facilitate support, support staff require the following:

- a modem connected to each server
- Remote Access Services (RAS) configured on the server
- pcAnywhere Host software installed on the server

Remote support in a Virtual Private Network

In some enterprises, a modem attached to a server is viewed as a security risk. Rather than using a modem and RAS, these enterprises prefer to use a Virtual Private Network (VPN). Many VPN technologies, and configurations within the technologies, are available. To facilitate remote support, Nortel Networks recommends a technology based on the Contivity 1100 (as a minimum) in a host-to-gateway configuration.

This chapter provides guidelines for this recommended VPN configuration.

Guidelines for the remote support VPN at the customer's premises

When creating your VPN for remote support, follow these guidelines:

- Create a dedicated subnet for Nortel Networks voice application servers, and treat this subnet as mission-critical. (It is a good network engineering practice, even in a non-VPN environment, to optimize network traffic by grouping servers that need to communicate with each other on a subnet.)
- Install, at a minimum, Nortel Networks' Contivity 1100, version 4.8 or later, with the modem option. Configure the modem as a user tunnel to listen on the PSTN.
- Connect the Contivity VPN Switch to the Nortel Networks server subnet.
- Configure Contivity, as well as any network routers and firewalls, to give inbound remote support users unrestricted access to the Nortel Networks application servers.
- Optionally, restrict remote support users' access to other subnets of your LAN/WAN. If you do, make sure that the Nortel Networks application servers have unrestricted access to the enterprise LAN/WAN.
- If you must connect the ELAN to the CLAN (for example, if you are using Symposium Call Center Server in a networked OTM environment), take the additional precaution of configuring the routing switch to allow *only* OTM-related traffic, ftp traffic, rlogin traffic, and SNMP traffic through into the ELAN.
- Activate split tunneling on the Contivity VPN Switch. Concerns over access into the corporate network may be alleviated by restricting access of remote support staff from other subnets.

VPN configurations

Introduction

This section describes recommended configurations that meet the needs of most customers. However, since every network is different, the exact configurations may not be practical in all environments. Use them as a starting point when creating your VPN.

Benefits

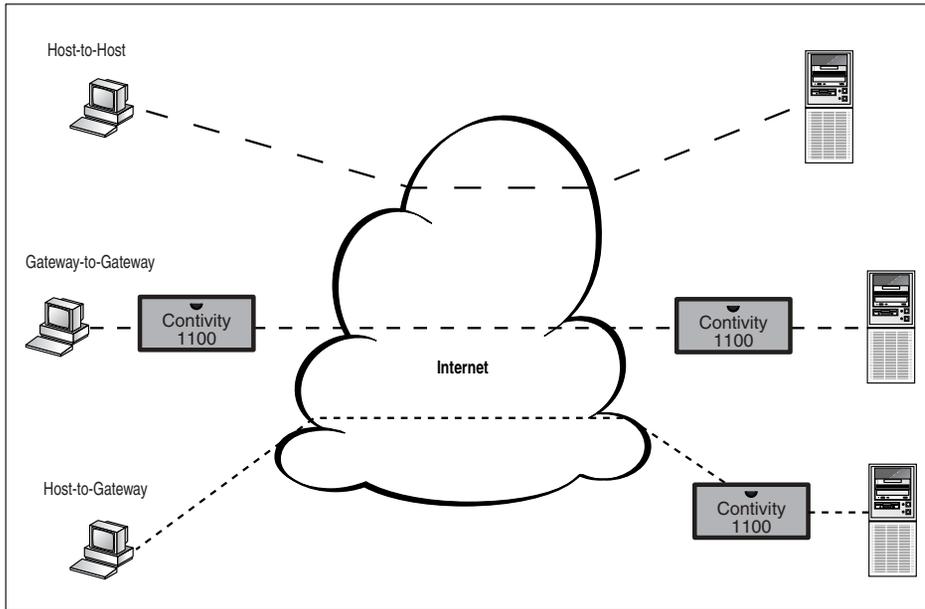
The recommended configurations provide the following benefits:

- They protect the customer's network.
- They are accessible from any location, even through an analog line.
- They provide a flexible design that can be extended to non-Nortel Networks products, and that can accommodate customer-specific network requirements.
- The VPN equipment is local to the equipment it serves, resulting in a simple network and management simplicity.
- The solution is cost-effective.

The recommended configuration is provided as a starting point in the process of VPN design. However, when you deviate from the recommended configurations, you sacrifice some of these benefits.

Configuration types

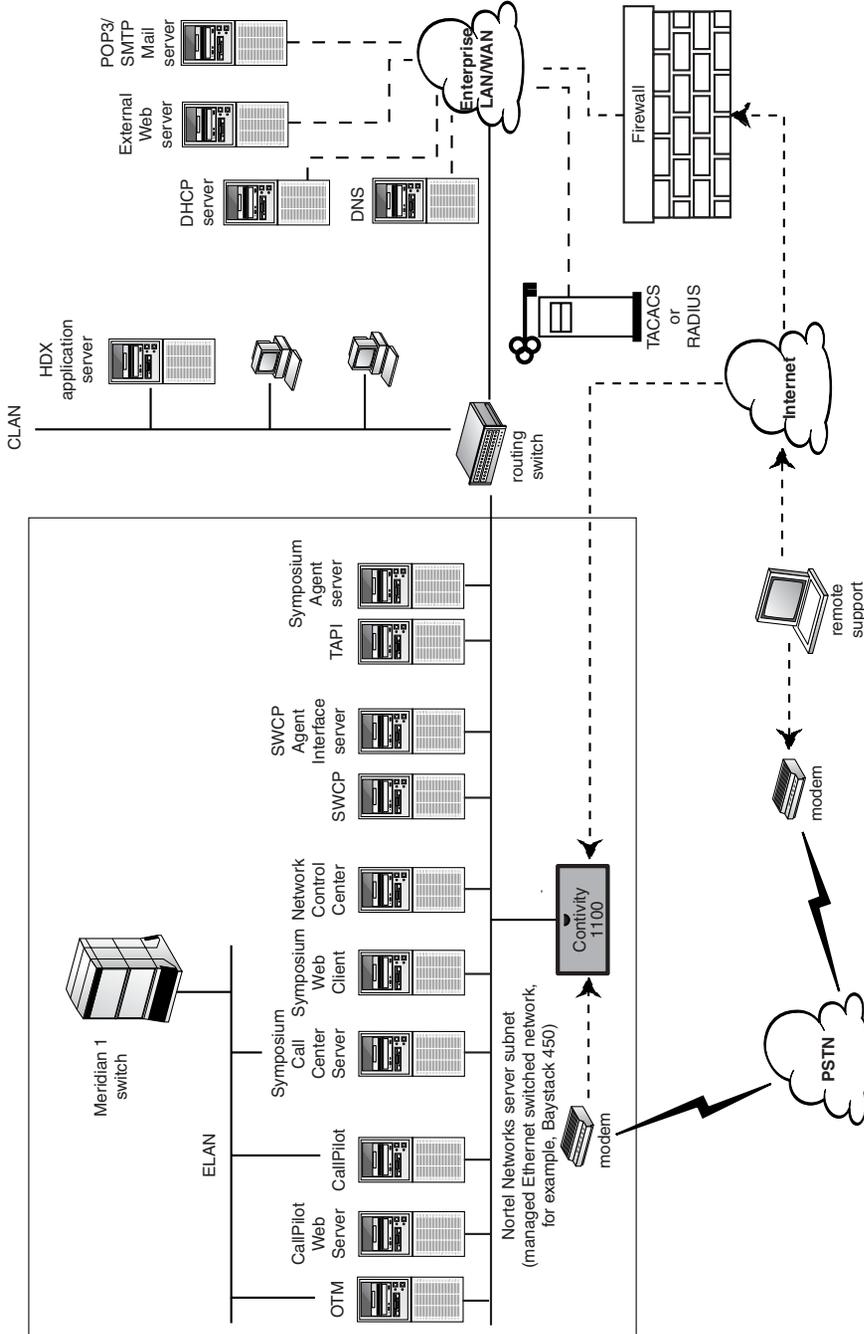
A VPN can be configured in three ways, as shown in the following illustration. Nortel Networks recommends a host-to-gateway configuration for the remote support VPN.



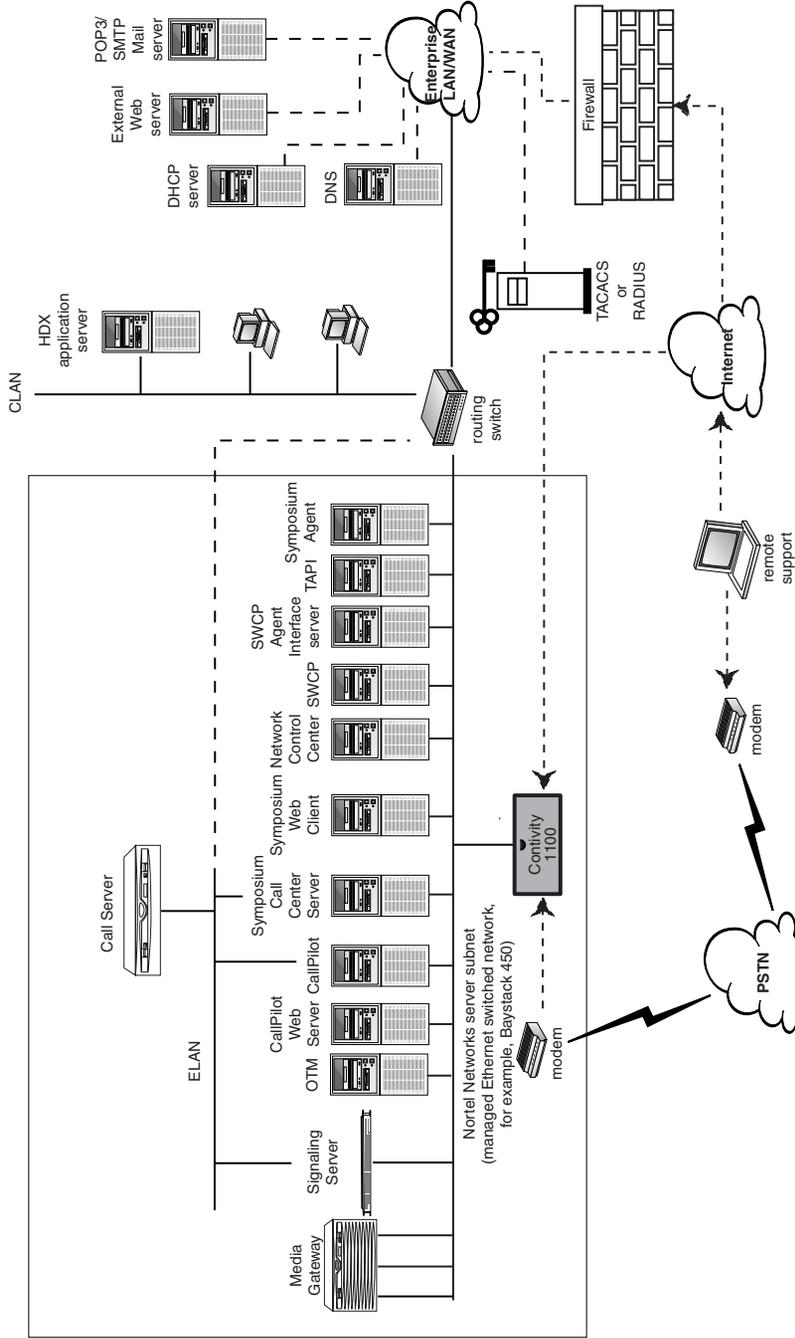
Illustrations

The following illustrations show recommended VPN configurations. The first illustration shows a VPN in a non-Voice Over IP (VoIP) environment. In this illustration, the ELAN is isolated. The second illustration shows a VPN in a VoIP environment with the ELAN connected to the customer's network.

Non-VoIP with isolated ELAN



VoIP



Appendix A

Product limits

In this appendix

Product limits

126

Product limits

Maximum capacity values

The following table specifies the maximum capacity values supported by Symposium Call Center Server:

Notes:

- The capacities supported on a given server are limited by the server platform. To determine the capacity of your server, use the CapTool application.
- The Release 4.2 values apply to the Meridian 1/Succession 1000 switch only. Release 4.2 does not support DMS.
- These are the values supported by Symposium Call Center Server. Capacity values are also limited by switch capacity. To find out the limits for your switch, check your switch documentation.

Parameter	Release 4.2 maximum	Release 5.0 maximum
General parameters		
Number of logged-on agents		
Meridian 1/Succession 1000	1500	2200
DMS	N/A	3300
Note: Configurations above 1500 agents require special consideration for CLAN bandwidth and disk requirements.		
Number of agents defined in the system	3000	6000
Number of phonesets	3000	3000
Number of supervisors logged on (Classic Client)	100	100
Number of supervisors logged on (Web Client)	150	150

Parameter	Release 4.2 maximum	Release 5.0 maximum
Number of supervisors defined in the system	300	300
Note: The number of configured supervisors defined in the system is not limited, but Nortel Networks only tests up to 300 configured supervisors.		
Number of scripts	1500	1500
Note: The number of scripts defined in the system is not limited, but Nortel Networks only tests up to 1500 scripts.		
Number of active scripts	1000	1000
Note: The product contains three predefined scripts. Therefore, the customer can create 997 scripts.		
Maximum script size (characters)	30 000	50 000
Number of applications (that is, exit points from the Master_Script)	505	505
Note: The product contains five predefined applications. Therefore, the customer can create 500 applications.		
Number of call variables	20	20
Number of skillsets	354 (350)	1000 (996)
Notes:		
<ul style="list-style-type: none"> ■ The product contains four predefined skillsets. Therefore, for Release 4.2, the customer can create 350 skillsets, and for Release 5.0, the customer can create 996 skillsets. ■ The maximum given here includes both local skillsets and network skillsets. 		
Number of skillset priority levels	48	48
Number of skillsets per call	20	20

Parameter	Release 4.2 maximum	Release 5.0 maximum
Number of activity codes	5000 (4997)	5000 (4998)
Note: The product contains three predefined activity codes. Therefore, the customer can create 4997 activity codes.		
Inbound calls per hour	55 000	58 000
Number of waiting calls	3000	3000
Call resources parameters		
Number of IVR queues (Meridian 1/Succession 1000)	150	150
Number of IVR ports	500	1000
Number of ACCESS ports (Meridian 1/Succession 1000)	96	96
Number of routes	513	513
Number of trunks (Meridian 1/Succession 1000)	3000	3000
Note: Nortel Networks has only tested 1000 trunk members. There are no plans to test the 3000-trunk limit.		
Number of CDNs		
DMS	N/A	1000
Meridian 1/Succession 1000	750	750
Number of RAN and music routes	512	512
Number of DNISs	10 000	10 000
Assignment parameters		
Number of agents in an agent-to-supervisor assignment	1000	1000

Parameter	Release 4.2 maximum	Release 5.0 maximum
Matrix size for agent-to-skillset assignments	5000	5000
<p>Note: An agent-to-skillset assignment contains a matrix with a row for each agent in the assignment, and a column for each skillset to which the agents belong. The matrix size is the number of agents multiplied by the number of skillsets.</p>		
Number of agent-skillset reassignments in an agent-to-skillset assignment	1000	1000
<p>Note: In an agent-to-skillset assignment, you can change an agent's status for multiple skillsets. For example, you can put the agent James Jones on Standby for the skillset Bookings, and give him priority 1 for the skillset European Vacations. Thus, you have two reassignments for the agent James Jones in the agent-to-skillset assignment.</p>		
<p>Networking parameters (Meridian 1/Succession 1000 only)</p>		
Number of call processing nodes in the network (including local node)	30	30
<p>Note: The number of configured nodes is 30; however, only 20 nodes can be configured in the routing table.</p>		
Number of network skillsets	100	1000
<p>Note: The maximum given for Release 5.0 includes the four predefined skillsets, local skillsets, and network skillsets.</p>		
Number of sites in the routing table for a network skillset	20	20
Number of network skillsets to which a call is queued	10	10
Number of agent reservation requests per call	30	30
Number of remote applications (applications accessible over the network)	6000	6000

Parameter	Release 4.2 maximum	Release 5.0 maximum
Network calls per hour for which CBC data is collected	10 000	10 000
Number of target nodes	3	20
<p>Note: For Release 4.2, the default number of target nodes is three; however, up to 20 nodes are available through Symposium Professional Service (SPS).</p>		
Real-time displays		
Number of RTD screens	4	4
Database parameters		
Number of client PCs and RTI applications connected to the database	100	100
Number of other applications connected to the database	100	100
Number of Fault Management messages in database	7500	7500
Maximum number of report clauses	255	255
<p>Note: The Sybase database server supports a maximum of 255 clauses on a single SQL statement.</p>		
Third-party interface parameters		
Number of MLS applications	16	16
Number of MLS calls per hour	16 000	58 000
<p>Note: The Release 5.0 maximum is subject to confirmation by testing.</p>		
Number of Symposium Event Interface (SEI) applications	3	3
<p>Note: SEI is available in SU09 of Symposium Call Center Server 4.2.</p>		

Parameter	Release 4.2 maximum	Release 5.0 maximum
Number of HDX connections	3	10
Note: When configured, Database Integration Wizard (DIW) uses a single HDX connection.		
Number of RTI client systems/applications	100	100
Other parameters		
Number of scripts activated under load	1	1
Note: Script activation supports activation cascading, where the activation of a parent script forces activation of all lower-level scripts. Do not use this feature on a system under load. Under load, activate scripts from the lowest level up, with the Master script being activated last.		
Number of CPUs	2	4
Steady state CPU	50%	50%
Recommended disk space (Gbytes)	64	64
Number of servers per switch		
DMS	N/A	16
Meridian 1/Succession 1000	1	3
Note: The number of active agents supported by the switch includes active agents at all connected servers.		

Appendix B

Standard call models

In this appendix

Inbound call models

134

Inbound call models

Introduction

For the purposes of the Symposium Call Center Server performance evaluation, five typical local inbound call models are defined. These models apply to calls that originate on the local node.

Meridian 1/Succession 1000 models

Symposium Voice Processing (SVP)

This call model assumes that the average call uses the following services:

- basic call
- queuing to two skillsets
- voice services controlled by Symposium Call Center Server (Give Controlled Broadcast, Collect Digits, and Open/Close Voice Session)

Meridian Voice Processing (MVP)

This call model assumes that the average call uses the following services:

- basic call
- queuing to two skillsets
- voice services controlled by the switch (Give RAN instead of Give Controlled Broadcast and Give IVR instead of Collect Digits and Open/Close Voice Session)

Hybrid

The hybrid call model is a combination of the SVP and MVP call models.

DMS models

Symposium Customer (Simple)

The customer uses a DMS switch with an external IVR system. Each call is given IVR treatment, and then it is routed to an agent with a particular skillset.

Busy Symposium Customer (Complex)

The customer uses a DMS switch with an external IVR system. Each call is given IVR treatment followed by multiple RAN or music treatments while waiting for an agent.

Number and types of services per call

The following table shows the average number and types of services assumed for calls in each model:

Parameter	M1/Succession 1000			DMS	
	SVP	MVP	Hybrid	Simple	Complex
Basic Call	1	1	1	1	1
Average number of skillset queues entered per inbound call	2	2	2	1	2.2
Average number of agent queues entered per inbound call	0	0	0	0	0.1
Average number of controlled broadcasts in Start/Stop mode per inbound call. Never with Give RAN.	3	0	1	N/A	N/A
Average number of controlled broadcasts in Continuous mode per inbound call	0	0	0	N/A	N/A
Average number of collect digit services per inbound call. Two digits each time (including voice session and play prompt).	1	0	0	N/A	N/A
Average number of Give IVR treatments per inbound call	1	1	1	N/A	N/A
Average number of Give RAN treatments per inbound call. Never with GCB.	1	3	2	0.2	1

Parameter	M1/Succession 1000			DMS	
	SVP	MVP	Hybrid	Simple	Complex
Average number of Give Music treatments per inbound call	1	1	1	0	1.5
Average number of Host Data Exchange Send Info treatments per inbound call. Only if Host Data Exchange is present.	1	1	1	1	1
Average number of Host Data Exchange Request/Get Response treatments per inbound call. Only if Host Data Exchange is present.	1	1	1	0	0
Average number of Intrinsic References per inbound call (Expected Wait Time, Longest Idle Agent, Oldest Call, Position in Queue)	5	5	5	2	5
Average number of If Then Else treatments per inbound call	5	5	5	2	4
Proportion of inbound calls that are transferred to another agent or DN	5%	5%	5%	0%	5%
Proportion of inbound calls that are conferenced with another agent or supervisor	5%	5%	5%	0%	15%
Proportion of conferenced calls completed by an MLS application (such as Symposium Agent)	0%	0%	0%	5%	10%
External IVR system connected to the DMS system	N/A	N/A	N/A	Yes	Yes

Parameter	M1/Succession 1000			DMS	
	SVP	MVP	Hybrid	Simple	Complex
Average number of screen pops per inbound call	1.2	1.2	1.2	1.2	1.2
Average number of MLink messages per inbound call (excluding screen pops)	0	0	0	0	0
Collected call-by-call statistics	Yes	Yes	Yes	Yes	Yes
Average number of network skillset queues entered per call	2	2	2	N/A	N/A
Proportion of calls arriving at the local node that are queued to a network skillset	10%	10%	10%	N/A	N/A

Appendix C

IP Multicast Networking

In this appendix

Overview	140
Multicast sending and receiving	141
Implementing IP multicasting for Symposium Call Center Server	151

Overview

What is IP multicasting?

IP multicasting provides multipoint communication by simultaneously delivering information from one sender to multiple receivers who want to receive the information. The greatest advantage to IP multicasting is its ability to transmit information to many recipients in a way that minimizes the bandwidth required to communicate across networks, and the resources required by the sender to carry out a transmission.

Traditional multipoint communications

Traditional methods of multipoint communication require that a source send a copy of information to each recipient: ten recipients require ten copies of the data. This method, called point-to-point unicast, creates two constraints:

- The source's system resources are used up in the duplication and distribution of multiple copies of a single piece of information.
- The combined size of the copies of data sent to recipients cannot be greater than the share of bandwidth available to the source.

IP multicasting multipoint communications

Both point-to-point unicast and broadcast communications are server-based concepts that negatively impact the source system and its network.

With IP multicasting, communication is receiver-based. Users who want to receive data join a multicast host group and become members of that group. Since duplication and distribution of the information is handled by a router, the source computer's resources and its designated bandwidth are utilized efficiently, allowing the source to distribute information quickly and with minimal impact on the network.

Multicast sending and receiving

Introduction

To send to multiple users, IP multicasting communicates with multicast host groups that are comprised of multicast group members. Recipients must be members of multicast groups to receive multicast data. A sender, however, does not need to be a member in a multicast host group to transmit multicast data. Anyone who can send information to a multicast IP address can send multicast information to a multicast host group. The following sections look at the building blocks of multicast communication in greater detail.

How sending and receiving works

Multicast IP sending is the same as unicast sending: the sender indicates the IP address that it wants to send to, and the information travels through the network and arrives at its destination.

It is more complex to receive multicast IP datagrams. When an application on a PC indicates that it wants to receive multicast data, several things must occur in the background for the data to travel through the network(s) and be received by the application. The section below looks at sending and receiving within the framework of Symposium Web Client's Real-Time Reporting component.

Sending

Sending begins when a user opens a browser, connects to the application server, and opens Real-Time Reporting. Real-Time Reporting on the client issues a request to join a host member group associated with Real-Time Reporting multicast data. The request is sent to the host member group's multicast group host.

Note: When a multicast host group is part of a permanent group, the host filters continuously for data coming from the multicast IP address. If the host is dynamic, it only begins filtering for data when it receives a request for membership. See "Multicast host groups," on page 144 for more information on the types of multicast groups.

In IP multicasting, there is an All-Hosts Group with the reserved address 224.0.0.1, whose function is to represent all hosts on the network. The All Routers Group with the reserved multicast IP address 224.0.0.2 represents the communication point for all routers on the network. The All-Hosts Group continuously sends out requests to its hosts and asks for a report: “Are there any groups that contain members who want to receive multicast data?”

Since the concept of IP multicasting rests upon the idea of virtual networks, an All-Hosts Group should be viewed only as representing all of the host groups, not a physical piece of hardware. The address 224.0.0.1 can designate

- a router
- or
- a system with an IP multicast-capable network interface card

If you are using IP multicasting in a very simple network, one router on a LAN can represent

- the All-Hosts Group
- the All-Routers Group
- and
- the host that the host group members join to receive their multicast data

In this example, the network consists of two servers in Symposium Call Center Server on one LAN. The application server and its client PCs reside on a separate LAN. Each server in Symposium Call Center Server and the application server are connected to multicast routers.

In this scenario, one of the routers is designated as the All-Routers Group (224.0.0.2). The application server acts as the host to the host group members, while one of the Symposium Call Center Server routers acts as the All-Hosts Group (224.0.0.1). At this stage, the All-Hosts Group waits to find out if there are hosts with members who want to receive multicast data.

The All-Hosts Group sends a query requesting that its hosts report on its membership, and the query travels from the All-Hosts Group to the host(s).

The host(s) report on their membership lists. These are all of the clients who requested membership in a host group by opening a browser, launching Symposium Web Client, and then opening Real-Time Reporting.

The report travels from each host back to the All-Hosts Group.

Receiving

At this stage, the scene has been set for multicast data to be received by the browsers that have Real-Time Reporting running. The hosts know who their members are. The All-Hosts Group knows who its hosts are. The routers that service the hosts are aware that their hosts are waiting for multicast data. Symposium Call Center Server now needs to provide that data.

Symposium Call Center Server delivers its real-time statistics data to its IP multicast-capable router on its LAN. The router puts together the data to be sent to the host groups, and maps the address of the multicast All-Hosts Group to the IP address that it uses to send data.

The data is sent from the LAN router to the All-Hosts Group. The All-Hosts Group then sends the data to the routers, whose job it is to receive data for hosts on their network or subnetwork. The routers for each host forward the data to their hosts, and each host forwards the data to its members.

Note: In traveling from the receiver to the sender, the request may travel through several routers. Only the routers nearest to the sender and receiver must be multicast-capable.

Multicast groups and members

Multicast hosts

Any system or router can be a host and can send multicast data to a multicast group if it meets the following conditions:

- The network interface card in the system is multicast-capable.
- The system or router is on a network with a local multicast router.

Note: The sender does not have to be a member of a multicast host group if it is only sending multicast data. Inclusion in a multicast host group is required only if receipt of multicast data is required.

Multicast host groups

Recipients of IP multicasting datagrams are called host groups. Host groups fall into the following two categories:

- permanent host groups
- transient host groups

Permanent host groups are groups with an assigned IP multicast group address. The number of members in the host group is irrelevant in that a permanent host group with no members still exists as long as its IP multicast address is defined.

A transient host group, by contrast, exists only if it has at least one member that requires its services. The multicast IP address for the transient host group is not permanently assigned to the host group; however, the addresses that can be dynamically assigned to a host group have two restrictions. The IP multicast address for a transient host group

- must be in the address range designated for IP multicasting
- cannot be the same as an address for a permanent host group

Multicast groups are virtual groups: they exist only from the point of view of multicast-capable routers or an All-Hosts Group. A host is simply a PC in a network that is designated to accept requests for multicast data from other PCs in the same network. This host conveys its membership status to its designated multicast-capable router. A group is formed when other PCs communicate their desire to join the host's group. The PCs that want to join the group can be from different networks or subnetworks. Their communication with the host makes them part of a single group.

The following groups are some of the permanent host groups that exist in an IP multicast-capable network:

- **The All-Hosts Group:** This group is used to identify all IP multicast hosts at your organization. When a host reports that it has members who want to receive multicast data, it sends this report to the All-Hosts Group. The multicast IP address for this group is 224.0.0.1.
- **The All-Routers Group:** This group is used to identify all IP multicast routers at your organization. The multicast IP address for this group is 224.0.0.2.

Multicast host group members

Host group members have few restrictions. They can

- reside anywhere on any network
- join or leave a host group at any time
- join more than one host group

To receive a multicast message

- the member must join the group to which the message is being sent and
- the group that the member has joined must belong to a network that is registered with a local multicast router

If the member joins a group that does not belong to a network registered with a local multicast router, the router receives the multicast message but has no way of distributing the message through the network to the member.

Multicast addresses

IP multicasting specifies multicast host groups using Class D Internet Protocol addresses. These host group addresses range from 224.0.0.0 through 239.255.255.255. While IP addresses identify a specific physical location, a multicast IP address identifies a transmission session—a request conveyed from a client to a host to join a multicast group.

However, when choosing IP multicast sending and receiving addresses, you must be aware of the following restrictions:

- The IP multicast addresses between 224.0.0.0 and 224.0.0.255 inclusive are reserved for routing protocols and topology discovery or maintenance protocols.
- Additional IP multicast addresses between 224.0.0.0 and 239.255.255.255 are also reserved for specific applications like Net News.

The IP multicast addresses that you select for IP multicasting groups at your organization *cannot* be within the 224.0.0.0 and 224.0.0.255 range. In addition, you must check to make sure that you do not select an IP multicast address that has already been reserved for a specific multicast application.

The following organizations maintain current information on IP multicasting addressing and can provide access to an extensive list of reserved IP multicast addresses. It is highly recommended that you review the information at one or both of these sites prior to assigning an IP address to a multicast group:

- Internet Engineering Task Force (<http://www.ietf.org>)
- Internet Assigned Numbers Authority (<http://www.iana.org>)

Multicast routing methods

The method that multicast routers use to interact with one another depends upon the routing protocol that has been set up for communications. All of these routing protocols use a routing method that moves a multicast packet from its source to its destination(s). There are several different routing methods:

- flooding
- spanning trees
- core-based trees
- reverse path broadcasting
- truncated reverse path broadcasting
- reverse path multicasting

A detailed description of each of these routing methods is beyond the scope of this document. The section below briefly discusses the spanning tree method, one of the more simple and efficient routing methods. To find out more about routing methods, visit the Internet Engineering Task Force (<http://www.ietf.org>), and Internet Assigned Numbers Authority (<http://www.iana.org>) web sites. Both sites provide additional information and articles that address IP multicast routing methods in greater detail.

Spanning trees

Multicast routing depends upon its multicast-capable routers to exchange information about neighboring routers and efficiently route multicast traffic. The Internet Group Management Protocol (IGMP) selects one router as the primary router for each physical network in a LAN. This primary router creates a routing method called a spanning tree that connects all other routers that belong to an IP multicast group.

A spanning tree is a loop-free network of paths between routers. Only one path is established between each router. When each router is aware of the branches in the spanning tree, it copies multicast datagrams only to those branches of the tree. With this method, datagrams are duplicated only when the spanning tree branches, keeping the amount of duplication required on a network to a minimum.

Multicast protocols

There are a variety of protocols available for multicast routing. The protocol that your network operations department chooses for your routers depends upon the type of delivery service that you must provide.

If your network configuration does not require the delivery of multicast packets between routers or across networks, you only need the Internet Group Management Protocol. If your multicast data recipients extend beyond a single network, your network operations department must define multicast routing protocols for your routers. These protocols create the spanning trees and forward the multicast packets that are required to get the data to the group members.

The following list includes some of the most common multicast protocols and a brief description of the routing features that each provides.

Internet Group Management Protocol

When clients indicate that they want to join a group, and hosts indicate to routers that they have group members, Internet Group Management Protocol (IGMP) is the protocol used to convey this information between host group members, hosts, and routers. See “How sending and receiving works” on page 141 for more information on how group membership occurs. IGMP must be available on any interface running a multicast protocol, as well as on any static interface over which you want to transfer multicast traffic.

Distance Vector Multicast Routing Protocol

Routers that use Distance Vector Multicast Routing Protocol (DVMRP) advertise the shortest-path routes to the networks on which a multicasting source resides. DVMRP is the opposite of RIP, which advertises routes to destination networks.

Multicasting Extensions to Open Shortest Path First

Routers using Multicasting Extensions to Open Shortest Path First (MOSPF) utilize an enhanced version of Open Shortest Path First (OSPF). This protocol allows a router to forward multicast IP traffic within an autonomous OSPF (v.2) system.

Protocol Independent Multicast

Protocol Independent Multicast (PIM) provides efficient routes for multicast traffic that must cross the Internet to reach members of sparsely distributed multicast groups. The Nortel Networks implementation of PIM supports sparse mode. PIM communicates with far-flung members by

- inviting downstream members to join a shared tree by sending explicit join messages
- using rendezvous points (RPs) for receivers to meet new sources. Sources announce their existence to RPs; receivers query RPs to learn about multicast sessions.
- establishing a shortest-path tree to create a data path between sources and receivers

Resource Reservation Protocol

Resource Reservation Protocol (RSVP)-capable routers allow their host systems in an IP network to reserve resources for unicast or multicast dataflows.

Packet migration between multicast and non-multicast networks

With the variety of networks that exist and the data that must travel between them, it is too expensive and difficult (if not impossible) to set up network infrastructures that carry only multicast packets, while unicast networks carry only unicast data. The implementation of multicasting in your network does not preclude the transmission of unicast packets.

You can configure your routers to allow tunneling — unicast packets that travel as multicast packets, and multicast packets that travel as unicast packets between multicast and non-multicast networks. The table below provides an overview of how different packet types can travel between multicast and non-multicast networks:

Router receives	On interface type	Forwarding Action and How to Enable
Unicast or broadcast packet	Multicast	<p>The multicast protocol running on the interface forwards the packet to a multicast destination address (or list of multicast destination addresses) dictated by an IP traffic filter.</p> <p>The IP traffic filter must be configured to convert the unicast or broadcast packets to multicast.</p>
Multicast	Multicast	<p>The router's multicast protocol forwards the packet to</p> <ul style="list-style-type: none"> ■ a multicast configured outbound interface (based on multicast protocol decisions) or ■ a non-multicast, IGMP static configured outbound circuit <p>In Site Manager, you must set the IGMP static forwarding entries policy for Dynamic to Static forwarding mode.</p>

Router receives	On interface type	Forwarding Action and How to Enable
Multicast	Non-multicast, IGMP static configured	<p>The router forwards multicast packet traffic to a multicast-enabled network if</p> <ul style="list-style-type: none"> ■ multicast protocols are running on the routers ■ the IGMP static forwarding policy is set to Static to Dynamic ■ the IGMP interface parameter Static Forward Cache Lifetime is set to a value in accordance with the multicast protocol (DVMRP or MOSPF) running on the router
Multicast	Non-multicast, IGMP static configured	<p>The router forwards the multicast traffic to a non-multicast, static configured interface if</p> <ul style="list-style-type: none"> ■ the IGMP static forwarding policy is set to Static mode

Implementing IP multicasting for Symposium Call Center Server

IP multicast requirements

The preceding sections discussed how multicasting works, the communication between software and hardware that multicasting generates, and the routing and related protocols that make the transmission of multicast data between sources and destinations possible. With this information, you can begin considering how to implement IP multicasting for your specific LAN or WAN, or both, to facilitate Symposium Call Center Server's real-time data multicasting requirements.

The following list is a checklist of the requirements that must apply to your network, network components, and multicast-capable applications for Symposium Call Center Server's multicasting capabilities to work in a simple LAN configuration:

Requirements for multicast communication on one LAN	✓
The sending and receiving nodes in your network must be multicast-enabled.	
The TCP/IP protocol stack must support IP multicast sending and receiving.	
The software used to communicate a request to join a multicast group must support the IGMP protocol.	
IGMP must be configured on all routers that receive or forward multicast or non-multicast datagrams.	

Requirements for multicast communication on one LAN	✓
<p>The network interface cards and their drivers at the sending and receiving nodes must be able to filter for LAN data link layer addresses that have been mapped from network layer IP multicast addresses.</p> <p>Note: If there are two network interface cards installed on the application server (one for the ELAN and the other for the CLAN), then you must manually configure the cards so the application server always sends multicast data through the CLAN card. The client PCs are located on the CLAN and, therefore, expect to receive multicast data on this network.</p>	

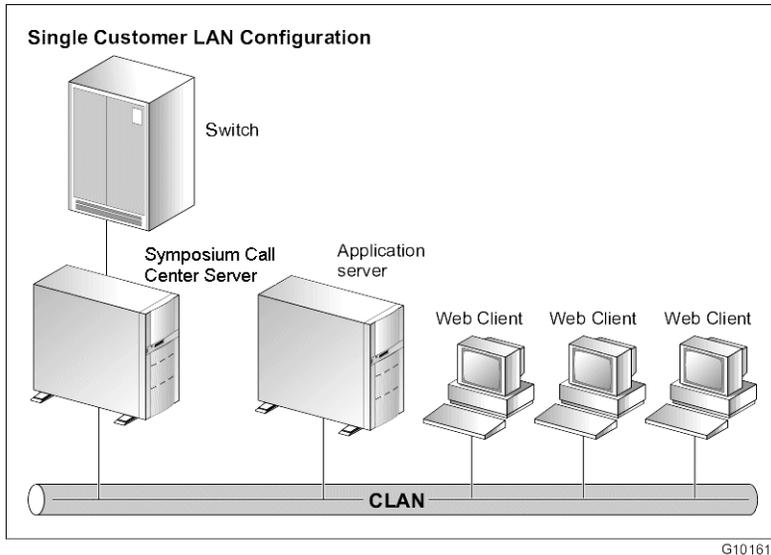
Routers are not required for a host to join a multicast group and share multicast data with other hosts on the same subnetwork. When multicast sending and receiving must travel between WANs and LANs, the list of requirements includes the above checklist, in addition to the items below:

Requirements for multiple LANs or LAN-to-WAN multicast communication	✓
Intermediate routers between sending and receiving nodes must be IP multicast-capable.	
Firewalls between LANs and WANs must be configured to permit IP multicast traffic.	
An IP traffic filter must be able to convert packets from unicast to broadcast or broadcast to unicast.	
An IP traffic filter must be able to convert packets from unicast to multicast or multicast to unicast.	
Configure an IGMP static forwarding policy for interfaces that multicast and for interfaces that do not multicast.	
Set policy filters to identify multicast protocol-compliant gateways, interfaces, tunnels, and networks for IGMP, DVMRP, and MOSPF.	
Configure the network interface cards on the application server so they always send multicast data through the CLAN card.	

Network deployment scenarios

Single LAN

In a single LAN environment, the clients, application server, and Symposium Call Center Server share a LAN. With no firewalls to potentially block access, this is the simplest environment to configure for IP multicasting.



When Symposium Call Center Server is installed, its IP multicast send and receive addresses are identified on the application server. Symposium Web Client uses the receive address to collect multicast data from Symposium Call Center Server. The IP multicast receive address on Symposium Web Client must be the same as the IP multicast send address of the server in Symposium Call Center Server. However, the IP multicast receive address on Symposium Call Center Server must be different from the IP multicast send address on Symposium Web Client.

The send address on the application server is the point from which multicast data is sent to the clients. The multicast-enabled router acts as both the host and the All-Hosts Group to the clients who become host group members when they open a browser and launch Real-Time Reporting.

Appendix D

Calculating Equivalent Basic Calls

In this appendix

Equivalent Basic Calls

156

Equivalent Basic Calls

The complexity of Symposium Call Center Server calls in terms of EBC is computed using the values from the following table for the appropriate switch software release:

Service	R25 EBC
Inbound Calls	
Basic call	2.40
QTS (Queue to Skillset)	0
QTNS (Queue To Network Skillset)	0
GCB (Give Controlled Broadcast Start/Stop)	1.70
GCB (Give Controlled Broadcast Continuous)	1.70
VSCDG (Collect Digits Voice Session)	2.29
GIVR (Give IVR, including transfer)	2.29
GRAN (Give RAN)	0.63
GMUS (Give Music)	0.25
Meridian Link messages/call (including screen pops)	0.60
Meridian Link calls transferred/conferenced	1.72
Conference/transfer	1.59
HDXSI (Data Exchange Send Info)	0
HDXRG (Data Exchange Send Request/Get Response)	0
INTR (script Intrinsic reference)	0
If-Then-Else	0

Service	R25 EBC
Incoming Accept Call	
M1 Basic Call + Trunks Incoming	1.18
Symposium Call Center Server scriptless overhead	1.33
Outgoing Accept Call	
M1 Basic Call + Trunks Outgoing	1.16
Symposium Call Center Server scriptless overhead	1.33
Outbound Calls	
Meridian Link calls transferred/conferenced	1.59
Calls conferenced/transferred out	1.72
Successful outbound call overhead	3.60
Unsuccessful outbound call overhead	1.88
Meridian Link messages per call (including screen pops)	0.60
Meridian Link messages/connection (unsuccessful call)	0.60
Meridian Link messages/unsuccessful connection	0.60

Note: The Meridian 1 Switch Capacity Tool automatically calculates the call rate based on an inputted call complexity model.

Glossary

A

accelerator key

A key on a phoneset that an agent can use to place a call quickly. When an agent presses an accelerator key, the system places the call to the configured number associated with the key. For example, if an agent presses the Emergency key, the system places a call to the agent's supervisor.

ACCESS

An internal protocol used by Symposium Call Center Server to directly control some of the voice services available on the CallPilot or Meridian Mail platform.

access class

A collection of access levels that defines the actions a member of the access class can perform within the system. For example, a member of the Administrator access class might be given a collection of Read/Write access levels.

access level

A level of access or permission given to a particular user for a particular application or function. For example, a user might be given View Only access to historical reports.

ACCESS link

A communication channel between Symposium Call Center Server and CallPilot or Meridian Mail.

ACCESS voice port

A voice port that is controlled by the ACCESS link.

ACD call

See Automatic call distribution call.

ACD-DN

See Automatic call distribution directory number.

ACD group

See Automatic call distribution group.

ACD routing table

See Automatic call distribution routing table.

ACD subgroup

See Automatic call distribution subgroup.

acquired resource

A resource configured on the switch that is under the control of Symposium Call Center Server. Resources must be configured with matching values on both the switch and Symposium Call Center Server.

activated script

A script that is processing calls or is ready to process calls. Before you can activate a script, you must first validate it.

active server

In a system with a Replication Server, the server that is providing call processing and administration services.

activity code

A number that an agent enters on his or her phoneset during a call. Activity codes provide a way of tracking the time agents spend on various types of incoming calls. They are also known as Line of Business (LOB) codes. For example, the activity code 720 might be used to track sales calls. Agents can then enter 720 on their phonesets during sales calls, and this information can be generated in an Activity Code report.

administrator

A user who is responsible for setting up and maintaining Symposium Call Center Server.

agent

A user who is responsible for handling customer calls.

agent logon ID

A unique identification number assigned to a particular agent. The agent uses this number when logging on. The agent ID is not associated with any particular phoneset.

agent to skillset assignment

A matrix that, when you run it, sets the priority of one or more agents for a skillset. Agent to skillset assignments can be scheduled.

agent to supervisor assignment

A definition that, when you run it, assigns one or more agents to specific supervisors. Agent to supervisor assignments can be scheduled.

AML

See Application Module Link.

API

See application program interface.

application

1. A logical entity that represents a Symposium Call Center Server script for reporting purposes. The Master script and each primary script have an associated application. The application has the same name as the script it represents. 2. A program that runs on a computer.

Application Module Link

An internal protocol used by Symposium Call Center Server to communicate directly with the switch.

application program interface

A set of routines, protocols, and tools that programmers use to develop software applications. APIs simplify the development process by providing commonly used programming procedures.

application server

The server on which the Symposium Web Client software is installed. This server acts as the middle layer that communicates with Symposium Call Center Server and makes information available to the client PCs.

associated supervisor

A supervisor who is available for an agent if the agent's reporting supervisor is unavailable. *See also* reporting supervisor.

Automatic call distribution

A means of automatically distributing an organization's incoming calls among a number of answering positions (ACD agents). Automatic call distribution is useful in operations where callers want a service rather than a specific person. Calls are serviced in the order they arrive and are distributed so that the workload at each answering position is approximately equal.

Automatic call distribution call

A call to an ACD-DN. ACD calls are distributed to agents in an ACD group based on the ACD routing table on the switch. *See also* Automatic call distribution directory number.

Automatic call distribution directory number

A primary or supplementary DN associated with an ACD group. Calls made to an automatic call distribution directory number are distributed to agents belonging to the group, based on the ACD routing table on the switch.

Automatic call distribution group

An entity defined on the switch for the purpose of call distribution. When a customer dials an ACD group, the call is routed to any agent who is a member of that group.

Automatic call distribution routing table

A table configured on the switch that contains a list of ACD-DNs used to define routes for incoming calls. This ensures that incoming calls not processed by Symposium Call Center Server will be queued to ACD groups and handled by available agents.

Automatic call distribution subgroup

An entity defined on the switch to assign supervisory responsibilities. Each subgroup has one supervisor phoneset and a number of agent phonesets associated with it. Agents can log on to any phoneset within their ACD subgroup. The supervisor must log on to the supervisor phoneset to monitor his or her assigned agents.

C

call age

The amount of time a call was waiting in the system before being answered by an agent.

call destination

The site to which an outgoing network call is sent. *See also* call source.

call intrinsic

A script element that stores call-related information assigned when a call enters Symposium Call Center Server. *See also* intrinsic, skillset intrinsic, time intrinsic, traffic intrinsic.

call presentation class

A collection of preferences that determines how calls are presented to an agent. A call presentation class specifies whether a break time between calls is allowed, whether an agent can put DN calls on hold for incoming ACD calls, and whether an agent phoneset displays that the agent is reserved for a network call.

call priority

A numerical value assigned in a script that defines the relative importance of a call. If two calls are in the queue when an agent becomes available, and one call is queued with a higher priority than the other, the agent receives the higher priority call first. *See also* skillset priority.

call source

The site from which an incoming network call originates. *See also* call destination.

call treatment

A script element that enables you to provide handling to a call while it is waiting to be answered by a call center agent. For example, a caller can hear a recorded announcement or music while waiting for an agent.

call variable

A script variable that applies to a specific call. A call variable follows the call through the system and is passed from one script to another with the call. *See also* global variable, script variable.

Calling Line Identification

An optional service that identifies the telephone number of the caller. This information can then be used to route the call to the appropriate agent or skillset. The CLID can also be displayed on an agent's phoneset.

CallPilot

A multimedia messaging system you can use to manage many types of information, including voice messages, fax messages, e-mail messages, telephone calls (including conferencing), calendars, and directories.

CDN

See controlled directory number.

CLAN

See Customer local area network.

Classic Client

The Windows-based client component for Symposium Call Center Server.

CLID

See Calling Line Identification.

client

The part of Symposium Call Center Server that runs on a personal computer or workstation and relies on the server to perform some operations. Two types of client are available, Classic Client and Symposium Web Client. *See also* server.

command

A building block used with expressions, variables, and intrinsics to create scripts. Commands perform distinct functions, such as routing a call to a specific destination, playing music to a caller, or disconnecting a caller.

Contivity VPN Switch

A Nortel Networks product that provides routing, firewall, bandwidth management, encryption, authentication, and data integrity for secure tunneling across managed IP networks and the Internet.

controlled directory number

A special directory number that allows calls arriving at the switch to be queued when the CDN is controlled by an application such as Symposium Call Center Server. When a call arrives at this number, the switch notifies the application and waits for routing instructions, which are performed by scripts in Symposium Call Center Server.

CTI

Computer Telephony Integration

customer administrator

A user who is responsible for maintaining Symposium Call Center Server.

Customer local area network

The LAN to which your corporate services and resources connect. The server in Symposium Call Center Server and the client both connect to the CLAN. Third-party applications that interface with the server also connect to this LAN.

D**DBMS**

Database Management System

deactivated script

A script that does not process any new calls. If a script is in use when it is deactivated, calls continue to be processed by the script until they are completed.

default activity code

The activity code that is assigned to a call if an agent does not enter an activity code manually, or when an agent presses the activity code button twice on his or her phoneset.

Each skillset has a defined default activity code.

default skillset

The skillset to which calls are queued if they have not been queued to a skillset or a specific agent by the end of a script.

desktop user

A configured user who can log on to Symposium Call Center Server from a client PC.

destination site

The site to which an outgoing network call is sent. *See also* source site.

DHCP

See dynamic host configuration protocol.

Dial-Up Networking

See Remote Access Services.

Dialed Number Identification Service

An optional service that allows Symposium Call Center Server to identify the phone number dialed by the incoming caller. An agent can receive calls from customers calling in on different DNISs and, if the DNIS is displayed on the phoneset, can prepare a response according to the DNIS.

directory number

The number that identifies a phoneset on a switch. The directory number (DN) can be a local extension (local DN), a public network telephone number, or an automatic call distribution directory number (ACD-DN).

directory number call

A call that is presented to the DN key on an agent's phoneset.

display threshold

A threshold used in real-time displays to highlight a value below or above the normal range.

DMS

Digital Multiplex Switch

DN

See directory number.

DN call

See directory number call.

DNIS

See Dialed Number Identification Service.

dongle

The attachment plugged into the parallel port of a server connected to a DMS/MSL-100 switch that authenticates the serial number required at the time of server installation.

dynamic host configuration protocol

A protocol for dynamically assigning IP addresses to devices on a network.

dynamic link library

A library of executable functions or data that can be used by a Windows application. Typically, a DLL provides one or more particular functions, and a program accesses the functions by creating either a static or dynamic link to the DLL. Several applications can use a DLL at the same time.

E**ELAN**

See embedded local area network.

embedded local area network

A dedicated Ethernet TCP/IP LAN that connects the server in Symposium Call Center Server and the switch.

Emergency key

A key on an agent's phoneset that, when pressed by an agent, automatically calls his or her supervisor to notify the supervisor of a problem with a caller.

event

1. An occurrence or action on Symposium Call Center Server, such as the sending or receiving of a message, the opening or closing of an application, or the reporting of an error. Some events are for information only, while others can indicate a problem. Events are categorized by severity: information, minor, major, and critical. 2. An action generated by a script command, such as queuing a call to a skillset or playing music.

expression

A building block used in scripts to test for conditions, perform calculations, or compare values within scripts. *See also* logical expression, mathematical expression, relational expression.

F**filter timer**

The length of time after the system unsuccessfully attempts to route calls to a destination site, before that site is filtered out of a routing table.

first-level threshold

The value that represents the lowest value of the normal range for a statistic in a threshold class. The system tracks how often the value for the statistic falls below this value.

G**global settings**

Settings that apply to all skillsets or IVR ACD-DNs that are configured on your system.

global variable

A variable that contains values that can be used by any script on the system. You can only change the value of a global variable in the Script Variable Properties sheet. You cannot change it in a script. *See also* call variable, variable.

H**HDX**

See Host Data Exchange

Host Data Exchange

A rich scripting language provided with Symposium Call Center Server to control treatment of calls.

ICM

See Intelligent Call Manager.

Incalls key

The key on an agent phoneset to which incoming ACD and Symposium Call Center Server calls are presented.

Intelligent Call Manager

A high capacity call center TCP/IP interface to the switch that enables the exchange of messages between the switch and a remote host computer.

Interactive voice response

An application that allows telephone callers to interact with a host computer using prerecorded messages and prompts.

Interactive voice response ACD-DN

A directory number that routes a caller to a specific IVR application. An IVR ACD-DN must be acquired for non-integrated IVR systems.

Interactive voice response event

A voice port logon or logoff. An IVR event is pegged in the database when a call acquires or de-acquires a voice port.

Internet Protocol address

An identifier for a computer or device on a TCP/IP network. Networks use the TCP/IP protocol to route messages based on the IP address of the destination. For customers using NSBR, site IP addresses must be unique and correct. The format of an IP address is a 32-bit numeric address written as four values separated by periods. Each value can be 0 to 255. For example, 1.160.10.240 could be an IP address.

intrinsic

A word or phrase used in a script to gain access to system information about skillsets, agents, time, and call traffic that can then be used in formulas and decision-making statements. *See also* call intrinsic, skillset intrinsic, time intrinsic, traffic intrinsic.

IP address

See Internet Protocol address.

IVR

See Interactive voice response.

IVR ACD-DN

See Interactive voice response ACD-DN.

IVR event

See Interactive voice response event.

IVR port

See voice port.

L**LAN**

See Local area network.

Line of Business code

See activity code.

LOB code

See activity code.

Local area network

A computer network that spans a relatively small area. Most LANs connect workstations and personal computers, and are confined to a single building or group of buildings.

local call

A call that originates at the local site. *See also* network call.

local skillset

A skillset that can be used at the local site only. *See also* network skillset, skillset.

logical expression

A symbol used in scripts to test for different conditions. Logical expressions are AND, OR, and NOT. *See also* expression, mathematical expression, relational expression.

M**M1**

Meridian 1 switch

M1 IE

Meridian 1 Internet Enabled switch

Management Information Base

A data structure that describes the collection of all possible objects in a network. Each managed node maintains one or more variables (objects) that describe its state. Symposium Call Center Server Management Information Bases (MIBs) contribute to the overall network MIB by

- identifying Nortel Networks/Meridian/Symposium Call Center Server nodes within the network
- identifying significant events (SNMP traps), such as alarms reporting
- specifying formats of alarms

Master script

The first script executed when a call arrives at Symposium Call Center Server. A default Master script is provided with Symposium Call Center Server, but it can be customized by an authorized user. It can be deactivated but not deleted. *See also* network script, primary script, script, secondary script.

mathematical expression

An expression used in scripts to add, subtract, multiply, and divide values. Mathematical expressions are addition (+), subtraction (-), division (/), and multiplication (*). *See also* expression, logical expression, and relational expression.

Meridian Link Services

A communications facility that provides an interface between the switch and a third-party host application.

Meridian Mail

A Nortel Networks product that provides voice messaging and other voice and fax services.

Meridian MAX

A Nortel Networks product that provides call processing based on ACD routing.

MIB

See Management Information Base.

MLS

See Meridian Link Services.

MM

See Meridian Mail.

MSL-100

Meridian Stored Logic 100 switch

music route

A resource installed on the switch that provides music to callers while they wait for an agent.

N**NACD call**

A call that arrives at the server from a network ACD-DN.

NCC

See Network Control Center.

network call

A call that originates at another site in the network. *See also* local call.

Network Control Center

The server on a Symposium Call Center Server system where NSBR is configured and where communication between servers is managed.

network interface card

An expansion board that enables a PC to be connected to a local area network (LAN).

network script

The script that is executed to handle error conditions for Symposium Call Center Server calls forwarded from one site to another, for customers using NSBR. The network script is a system-defined script provided with Symposium Call Center Server, but it can be customized by an authorized user. It can be deactivated but not deleted. *See also* Master script, primary script, script, secondary script.

Network Skill-Based Routing

An optional feature with Symposium Call Center Server that provides skill-based routing to multiple networked sites.

network skillset

A skillset that is common to every site on the network. Network skillsets must be created at the Network Control Center (NCC).

night mode

A skillset state in which the server does not queue incoming calls to the skillset, and in which all queued calls are given night treatment. A skillset goes into night mode automatically when the last agent logs off, or the administrator can put it into night mode manually. *See also* out-of-service mode, transition mode.

NPA

See Number Plan Area.

NSBR

See Network Skill-Based Routing.

Number Plan Area

Area code

O**object linking and embedding**

A compound document standard that enables you to create objects with one application, and then link or embed them in a second application.

ODBC

See Open Database Connectivity.

OEM

Original equipment manufacturer

OLE

See object linking and embedding.

Open Database Connectivity

A Microsoft-defined database application program interface (API) standard.

Optivity Telephony Manager

A Nortel Networks application used for switch management. It provides management simplicity and flexible control.

OTM

See Optivity Telephony Manager.

out-of-service mode

A skillset state in which the skillset does not take calls. A skillset is out of service if there are no agents logged on or if the supervisor puts the skillset into out-of-service mode manually. *See also* night mode, transition mode.

out-of-service skillset

A skillset that is not taking any new calls. While a skillset is out of service, incoming calls cannot be queued to the skillset. *See also* local skillset, network skillset, skillset.

P**PBX**

See private branch exchange.

pegging

The action of incrementing statistical counters to track and report on system events.

pegging threshold

A threshold used to define a cut-off value for statistics, such as short call and service level. Pegging thresholds are used in reports.

PEP

See Performance Enhancement Package.

Performance Enhancement Package

A Symposium Call Center Server supplementary software application that enhances the functionality of previously released software by improving performance, adding functionality, or correcting a problem discovered since the original release.

personal directory number

A DN on which an agent can be reached directly, usually for private calls.

phoneset

The physical device, connected to the switch, to which calls are presented. Each agent and supervisor must have a phoneset.

phoneset display

The display area on an agent's phoneset where information about incoming calls can be communicated.

Position ID

A unique identifier for a phoneset, used by the switch to route calls to the phoneset. Referred to as Telephony/Port Address in Symposium Call Center Server.

primary ACD-DN

A directory number that callers can dial to reach an ACD group.

primary script

A script that is executed or referenced by the Master script. A primary script can route calls to skillsets, or it can transfer routing control to a secondary script. *See also* Master script, network script, script, secondary script.

private branch exchange

A telephone switch, typically used by a business to service its internal telephone needs. A PBX usually offers more advanced features than are generally available on the public network.

R**RAN**

recorded announcement

RAN route

See recorded announcement route.

RAS

See Remote Access Services.

Real-time Statistics Multicast

An interface that provides real-time information to third-party applications in either multicast or unicast format.

recorded announcement route

A resource installed on the switch that offers a recorded announcement to callers.

relational expression

An expression used in scripts to test for different conditions. Relational expressions are less than (<), greater than (>), less than or equal to (<=), greater than or equal to (>=), and not equal to (<>). *See also* expression, logical expression, mathematical expression.

Remote Access Services

A feature built into Windows NT and Windows 95 that enables users to log on to an NT-based LAN using a modem, X.25 connection, or WAN link. This feature is also known as Dial-Up Networking.

Replication Server

A server that backs up the active server to the standby server in real time.

reporting supervisor

The supervisor who has primary responsibility for an agent. When an agent presses the Emergency key on the phoneset, the emergency call is presented to the agent's reporting supervisor. *See also* associated supervisor.

round robin routing table

A routing table that queues the first call to the first three sites in the routing table, then the second three sites, then the third three sites, and so on, until an agent is reserved at one of the sites. *See also* sequential routing table.

route

A group of trunks. Each trunk carries either incoming or outgoing calls to the switch. *See also* music route, RAN route.

routing table

A table that defines how calls are routed to the sites on the network. *See also* round robin routing table, sequential routing table.

RSM

See Real-time Statistics Multicast.

S**sample script**

A script that is installed with the Symposium Call Center Server client. Sample scripts are stored as text files in a special folder on the client. The contents of these scripts can be imported or copied into user scripts to create scripts for typical call center scenarios.

SCM

See Service Control Manager.

script

A set of instructions that relates to a particular type of call, caller, or set of conditions, such as time of day or day of week. *See also* Master script, network script, primary script, secondary script.

script variable

See variable.

second-level threshold

The value used in display thresholds that represents the highest value of the normal range for a given statistic. The system tracks how often the value for the statistic falls outside this value.

secondary directory number

A DN defined on the agent's phoneset as a Centrex line for incoming and outgoing non-ACD calls.

secondary script

Any script (other than a Master, network, or primary script) that is referenced from a primary script or any other secondary script. There is no pegging of statistics for actions occurring during a secondary script. *See also* Master script, network script, primary script, script.

SEI

See Symposium Event Interface.

sequential routing table

A routing table method that always queues a call to the first three active sites in the routing table. *See also* round robin routing table.

server

A computer or device on a network that manages network resources. Examples of servers include file servers, print servers, network servers, and database servers. Symposium Call Center Server is used to configure the operations of the call center. *See also* client.

service

A process that adheres to a Windows NT structure and requirements. A service provides system functionality.

Service Control Manager

A Windows NT process that manages the different services on the PC.

service level

The percentage of incoming calls answered within a configured number of seconds.

service level threshold

A parameter that defines the number of seconds within which incoming calls should be answered.

Simple Network Management Protocol

A systematic way of monitoring and managing a computer network. The SNMP model consists of four components:

- managed nodes, which are any device, such as hosts, routers, and printers, capable of communicating status to the outside world via an SNMP management process called an SNMP Agent
- management stations, which are computers running special network management software that interact with the Agents for status
- management information, which is conveyed through exact specifications and format of status specified by the MIB
- Management Protocol or SNMP, which sends messages called protocol data units (PDUs)

site

1. A system using Symposium Call Center Server that can be accessed using SMI. 2. A system using Symposium Call Center Server and participating in Network Skill-Based Routing.

skillset

A group of capabilities or knowledge required to answer a specific type of call. *See also* local skillset, network skillset.

skillset intrinsic

A script element that inserts information about a skillset in a script. Skillset intrinsics return values such as skillsets, integers, and agent IDs. These values are then used in queuing commands. *See also* call intrinsic, intrinsic, time intrinsic, and traffic intrinsic.

skillset priority

An attribute of a skillset assignment that determines the order in which calls from different skillsets are presented to an agent. When an agent becomes available, calls might be waiting for several of the skillsets to which the agent belongs. The server presents the call queued for the skillset for which the agent has the highest priority.

SNMP

See Simple Network Management Protocol.

source site

The site from which an incoming network call originates. *See also* destination site.

standby

In skillset assignments, a property that grants an agent membership in a skillset, but makes the agent inactive for that skillset.

standby server

A server that contains an up-to-date version of the database, for use when the active server becomes unavailable.

supervisor

A user who manages a group of agents. *See also* associated supervisor and reporting supervisor.

supplementary ACD-DN

A DN associated with a primary DN. Any calls to the supplementary DN are automatically routed to the primary DN. A supplementary DN can be a toll-free (1-800) number.

SWCP

See Symposium Web Center Portal.

switch

The hardware that receives incoming calls and routes them to their destination.

switch resource

A device that is configured on the switch. For example, a CDN is configured on the switch, and then is used as a resource with Symposium Call Center Server. *See also* acquired resource.

Symposium Agent

An agent productivity tool that enables contact center agents to provide intelligent and personalized customer care. Agents use a personal computer to access the agent telephony functions.

Symposium Call Center Server

A client/server contact center solution for varied and changing business requirements. It offers a suite of applications that includes call processing and agent handling, management and reporting, networking, and third-party application interfaces.

Symposium Call Center Server call

A call to a CDN that is controlled by Symposium Call Center Server. The call is presented to the Incalls key on an agent's phoneset.

Symposium Event Interface

An interface that provides third-party vendors with the information they need to create complementary applications by providing call progress and resource events.

Symposium Standby Server

The server that contains an up-to-date back-up version of the Symposium Call Center Server database, for use if the active server fails. The database is kept up-to-date by the Replication Server.

Symposium Web Center Portal

A client/server contact center application that expands contact center e-mail capabilities to allow agents to view, respond to, and track requests over the Internet.

Symposium Web Client

A browser-based tool for call center administrators and supervisors used for managing and configuring a contact center and its users, defining access to data, and viewing real-time and historical reports. The Symposium Web Client software is installed on an application server. *See also* application server.

system-defined scripts

The Master_Script and the Network_Script (if NSBR is enabled). These scripts can be customized or deactivated by a user, but cannot be deleted. These scripts are the first scripts executed for every local or network call arriving at the call center.

T**TAPI**

See Telephony Application Program Interface.

target site

See destination site.

TCP/IP

See Transmission Control Protocol/Internet Protocol.

TDM

See Time-Division Multiplex.

telephony

The science of translating sound into electrical signals, transmitting them, and then converting them back to sound. The term is used frequently to refer to computer hardware and software that perform functions traditionally performed by telephone equipment.

Telephony Application Program Interface

An interface between the switch and an application that allows the application to control the telephone on a user's desktop.

threshold

A value for a statistic at which system handling of the statistic changes.

threshold class

A set of options that specifies how statistics are treated in reports and real-time displays. *See also* display threshold, pegging threshold.

Time-Division Multiplex

A method of transmission in which a signal is separated into multiple segments at the transmission source, and then reassembled at the receiving end.

time intrinsic

A script element that stores information about system time, including time of day, day of week, and week of year. *See also* call intrinsic, intrinsic, skillset intrinsic, traffic intrinsic.

Token Ring

A PC network protocol developed by IBM. A Token Ring network is a type of computer network in which all the computers are arranged schematically in a circle.

traffic intrinsic

An intrinsic that inserts information about system-level traffic in a script. *See also* call intrinsic, intrinsic, skillset intrinsic, time intrinsic.

transition mode

A skillset state in which the server presents already queued calls to a skillset. New calls queued to the skillset are given out-of-service treatment. *See also* night mode, out-of-service mode.

Transmission Control Protocol/Internet Protocol

The communication protocol used to connect devices on the Internet. TCP/IP is the standard protocol for transmitting data over networks.

treatment

See call treatment.

trunk

A communications link between a PBX and the public central office, or between PBXs. Various trunk types provide services such as Direct Inward Dialing (DID trunks), ISDN, and Central Office connectivity.

U

user-created script

A script that is created by an authorized user on the Symposium Call Center Server system. Primary and secondary scripts are user-created scripts.

user-defined script

A script that is modified by an authorized user on the Symposium Call Center Server system.

utility

A program that performs a specific task, usually related to managing system resources. Operating systems contain a number of utilities for managing disk drives, printers, and other devices.

V

validation

The process of checking a script to ensure that all the syntax and semantics are correct. A script must be validated before it can be activated.

variable

A placeholder for values calculated within a script, such as CLID. Variables are defined in the Script Variable Properties sheet and can be used in multiple scripts to determine treatment and routing of calls entering Symposium Call Center Server. *See also* call variable, global variable.

Virtual Private Network

A private network that is configured within a public network to take advantage of the economies of scale and management facilities of large networks.

voice port

A connection from a telephony port on the switch to a port on the IVR system.

VPN

See Virtual Private Network.

W

WAN

See also Wide area network.

Wide area network

A computer network that spans a relatively large geographical area. Typically, a WAN consists of two or more local area networks (LANs). The largest WAN in existence is the Internet.

workload scenarios

Sets of configuration values defined for typical patterns of system operations. Five typical workload scenarios (entry, small, medium, large, and upper end) are used in the Capacity Assessment Tool for capacity analysis for Symposium Call Center Server.

Index

A

ACCESS

- connection 107
- link 112
- protocol 33
- requirements 109
- voice ports 107

activating Master script 53

active server 28

addresses

- Class D Internet Protocol 145
- multicast 145
- restriction for IP multicast 145

agents

- number configured 13
- number logged on 13

agent-to-skillset assignments, running 53

agent-to-supervisor assignments, running 53

All-Hosts Group 143, 144

- and multicast data 142

All-Routers Group 144

- and multicast data 142

AML 33, 46

Application Module Link. *See* AML

application server 64, 69

- CLAN/WAN impact 69
- CPU utilization 66
- minimum refresh rate on 68
- multiple 69
- requirements 65

architecture

- DMS switch 34
- Meridian 1/Succession 1000 27

assignments, running 53

audio routes 87

average

- complexity model 81
- CPU utilization 52
- delay factor 96
- pages per second 52
- talk time 44

B

bandwidth

- contention 96
- required, CSL 114
- required, ELAN, for multiple servers on a switch 95

basic call 86, 135, 156

- cost 42, 80

buffers 85

busy Symposium customer model 135

C

Call Arrival Rate 44

Call Center Module loads supported 84

Call Center Seconds. *See* CCS

call complexity 42, 80

call load 42

call models 133

- busy Symposium customer 135
- hybrid 81, 134
- MVP 134
- simple 80
- SVP 81, 108, 134
- Symposium customer 134

call rate 44

- maximum 82

CallPilot 30

- and multiple servers on the switch 15, 110
- software release 110
- See also* Symposium Voice Services on CallPilot

calls per hour 44

capacity 55

- ACCESS link 113
- CSL 115
- enhancements in Release 5.0 13
- maximum 126

Capacity Assessment Tool. *See* CapTool

CapTool 38

- CCM loads supported 84
 - CCS 112
 - CDNs, calculating requirements for 88
 - CLAN 30
 - average utilization 53
 - connection to ELAN 99
 - impact of Symposium Web Client application server on 69
 - maximum utilization 93
 - requirements 92
 - traffic 92
 - Class D Internet Protocol addresses 145
 - Classic Client 30
 - and multiple servers on the switch 14
 - configuration 76
 - client PCs 64
 - Classic Client configuration 76
 - maximum supported by different processors 67
 - Symposium Web Client, requirements 65
 - clients
 - and multiple servers on a switch 14
 - Symposium Web Client 65
 - Close Voice Session. *See* voice sessions
 - Collect Digits command 109
 - Command and status link. *See* CSL
 - commit limit 52
 - committed bytes 52
 - Communication Server 2100 switch 14
 - See also* DMS switch
 - Complete Transfer message 88
 - complex call model 81, 135
 - CompuCALL 35
 - computer-telephony integration 32
 - configurations, VPN 120
 - configured agents 13
 - configuring
 - Contivity 119
 - LinkPlexer 89
 - consolidated real-time display data 69
 - contact center architecture
 - DMS switch 34
 - Meridian 1/Succession 1000 switch 27
 - Contivity 118, 119
 - configuring 119
 - Contivity 1100 documents 22
 - Controlled Directory Numbers. *See* CDNs
 - cost of basic call 42, 80
 - CPH 44
 - CPU
 - cost of Symposium Web Client on Symposium Call Center Server 69
 - impact, Symposium Voice Services on CallPilot 110
 - load, minimizing on Symposium Web Client 73
 - performance, Symposium Web Client 72
 - requirement, Classic Client 76
 - requirement, Symposium Call Center Server 39, 40
 - requirement, Symposium Web Client 65
 - requirement, Symposium Web Client application server 65
 - utilization, Symposium Call Center Server 52
 - utilization, Symposium Web Client application server 66
 - CRTD data 69
 - CS 2100 switch 14
 - See also* DMS switch
 - CSL 114
 - traffic cost 114
 - CTI 32, 46
 - Customer Local Area Network. *See* CLAN
- ## D
- data extractions 53
 - database backups 54
 - Database Integration Wizard 49
 - dedication of voice ports 107
 - delay factor 96
 - Distance Vector Multicast Routing Protocol 147
 - DIW 49
 - DMS switch 14
 - contact center architecture 34
 - documents 20
 - engineering 84
 - requirements 84
 - DMS/CompuCALL switch 35
 - DNs, sharing 88
 - drivers 39

duration, voice sessions 107
 DVMRP 147
 DXM_SERVER_SHUTDOWN message 50

E

EBCs 79, 155
 ELAN 30

- average utilization 53
- bandwidth required, for multiple servers on a switch 95
- connection to CLAN 99
- external communications on 100
- impact, Symposium Voice Services on CallPilot 110
- maximum utilization 94
- OTM on 99
- problems 95
- requirements 94
- traffic 94

 Embedded Local Area Network. *See* ELAN
 engineering

- network 91
- voice processing 105

 Equivalent Basic Calls 79, 155
 erlang 112
 Ethernet switch 99
 external

- communications on the ELAN 100
- events 87

F

fat client. *See* Classic Client
 features in Release 5.0 13
 filtering router 99
 front-end IVR 35, 80
 See also IVR
 ftp traffic 119
 ftServer 3220 39
 ftServer 3300 39

G

generating reports 53
 Get Response command 49
 Give Controlled Broadcast command 109
 Give IVR 108
 Give_Treatment message 85
 Give_Treatment(Music) message 85
 Give_Treatment(RAN) message 85
 Give_Treatment(Ringback) message 85
 GOS 107
 Grade of Service 107
 grounding, ACCESS link 112

H

H/W RAID 40
 hard disk capacity for Classic Client 76
 Hardware Assisted Software Mirroring 40
 Hardware Assisted Software RAID 40
 hardware high availability 40
 hardware platforms 39

- capacity limits 55
- maximum client PCs supported by 67
- Meridian Mail 112

 HASM 40
 HDX 32, 49

- and multiple servers on the switch 15
- application server 29

 High Availability servers 39
 historical reports, limit of simultaneously generated 68
 host application 49
 Host Data Exchange. *See* HDX
 Host Enhanced Routing 46
 host-to-gateway configuration 120
 hybrid model 81, 134

- and capacity 58

I

ICM messages 85
 IGMP 147

- impact
 - on CLAN/WAN of application server 69
 - Symposium Voice Services on CallPilot 110
 - inbound call models 134
 - increased product limits in Release 5.0 13
 - Initiate Transfer message 88
 - integrated IVR
 - and the DMS environment 35
 - See also* IVR
 - Interactive Voice Response. *See* IVR
 - inter-call interval 44
 - Internet Assigned Numbers Authority 146
 - Internet Engineering Task Force 146
 - Internet Group Management Protocol 147
 - invoke IDs 84
 - IP multicast addresses, restrictions 145
 - IP multicasting
 - implementing for Symposium Web Client 151
 - overview 140
 - requirements 151
 - IP telephony 28
 - IPML and multiple servers on the switch 15
 - ISDN 83
 - IVR 30
 - voice ports 108
- L**
- latency 102
 - limits, product 125
 - LinkPlexer 35, 88
 - configuring 89
 - logged on agents 13
 - Login message 88
 - logoffs 53
 - logons 53
 - Logout message 88
- M**
- M1 Switch Capacity Spreadsheet 79
 - Master script, activating 53
 - maximum
 - call rate, DMS switch 84
 - call rate, Meridian 1/Succession 1000 switch 82
 - capacity 126
 - client PCs supported by processors 67
 - utilization, ACCESS link 112
 - utilization, CLAN 93
 - utilization, ELAN 94
 - Mean
 - Holding Time 44
 - Time Between Calls 44
 - memory
 - management, Classic Client 76
 - Object 52
 - requirements, Symposium Call Center Server 40
 - See also* virtual memory
 - Memory Object 52
 - Meridian 1 switch
 - contact center architecture 27
 - documents 21
 - engineering 78
 - Meridian Link protocol 46
 - Meridian Link Services. *See* MLS
 - Meridian Mail 30
 - and multiple servers on the switch 15, 111
 - platforms 112
 - software release 111
 - See also* Symposium Voice Services on Meridian Mail
 - Meridian Voice Processing. *See* MVP
 - MHT 44, 112
 - MLS 32, 46, 88
 - in the DMS environment 35
 - messages supported on DMS switch 88
 - model
 - busy Symposium customer 135
 - call 133
 - hybrid 81, 134
 - inbound call 134
 - MVP 134
 - simple 80
 - SVP 81, 108, 134
 - Symposium customer 134
 - MOSPF protocol 148

MSL-100 switch 14
See also DMS switch

MTBC 44

multicast
 addresses 145
 data, sending and receiving 141
 group 141
 host group 141, 144
 host group, members of 145
 host group, permanent 144
 host group, transient 144
 hosts 143
 protocols 147
 routers 142
 routing methods 146
 traffic 70

multiple application servers 69

multiple servers on a switch
 and CallPilot 110
 and client 14
 and Meridian Mail 111
 ELAN bandwidth required for 95

multipoint communications
 and IP multicasting 140
 traditional 140

music routes 87

MVP 134

N

NACD 83

NCC 28

NCP messages 102

NCRTD data 70

Network ACD. *See* NACD

network call processing and peak busy hour 101

network consolidated real-time display data 70

Network Control Center. *See* NCC

network engineering 91

Network Loop Interface. *See* NLI link

Network Skill-Based Routing 28
 in the DMS environment 35

networked calls and capacity 57

networking and multiple servers on switch 14

new features in Release 5.0 13

NLI link 115

non-ACCESS voice ports 108

non-steady state 53

Nortel Networks Communication Server 2100
 switch 14

Nortel Networks voice application servers 119

Not Ready message 88

NSBR 28
 and multiple servers on switch 14
 in the DMS environment 35

number
 of active agents 44
 of lines displayed, CPU impact on
 Symposium Web Client 72
 of servers supported by a DMS switch 78, 84

O

ODBC 32

Open Database Connectivity. *See* ODBC

Open Voice Session. *See* voice sessions

operating system
 Classic Client 76
 Symposium Call Center Server 39
 Symposium Web Client 65
 Symposium Web Client application server 65

Optivity Telephony Manager. *See* OTM

OSPF protocol 148

OTM
 on the ELAN 99
 traffic 119

Outbound Predictive Dialing 46

P

pages per second 52

paging file 41

Performance Monitor 52

PIM protocol 148

Platform Vendor Independence 39

platforms 39
 capacity limits 55
 maximum client PCs supported by 67
 Meridian Mail 112

port request rate 107

post-call processing time 44
problems, ELAN 95
processors, rated capacity 55
product limit increases in Release 5.0 13
product limits 125
propagation delays 95, 96
Protocol Independent Multicast 148
provider application 49
PVI 39

Q

queuing delays 95

R

RAM requirements, Symposium Call Center Server 40, 52
RAN routes 87
rate, port request 107
rated capacity for processors 55
Ready message 88
Real-Time Data API. *See* RTD API
real-time displays
 and IP multicast 141
 minimum refresh rate on application server 68
 refresh rates and capacity 56
real-time reporting and IP multicast 141
Real-Time Statistics Multicast. *See* RSM
receiving multicast data 143
refresh rates
 and CPU impact on Symposium Web Client 72
 for real-time displays, and capacity 56
 minimum on application server 68
related documents 19
Release 5.0 features 13
reliability, lack of 95
remote support 117
 VPN guidelines 119
replication server 28
reporting
 and period of highest reporting activity 101
 real-time, and IP multicast 141

reports
 and temporary files 76
 generating 53
requirements
 ACCESS 109
 CLAN 92
 DMS switch 84
 ELAN 94
 Symposium Voice Services on CallPilot 110
 Symposium Voice Services on Meridian Mail 111
 Symposium Web Client 65
 voice ports 107
 WAN 101

Resource Reservation Protocol 148

rlogin traffic 119
robustness, lack of 95
Route_Call message 85
routers
 filtering 99
 multicast 142
routing methods
 multicast 146
 spanning tree 146
RSM 33
RSVP protocol 148
RTD API 33

S

scripts
 activating 53
 validating 53
SEI 33
Send Info command 49
Send Request command 49
sending IP multicast data 141
servers, support for multiple, on switch 14
service provider API 49
services per call 42, 44, 80, 135
session sharing 88
simple call model 80, 134
SMI Workbench 76
SNMP traffic 119
software phones 46
software/hardware high availability 40

- spanning tree routing method 146
 - split tunneling 119
 - standby server 28
 - steady state 52
 - subnet 119
 - Succession 1000 switch
 - contact center architecture 27
 - documents 21
 - engineering 78
 - See also* Meridian 1 switch
 - Succession 2000 switch 14
 - See also* DMS switch
 - support
 - for multiple servers 14
 - remote 117
 - SVP model 81, 108, 134
 - SWCP. *See* Symposium Web Center Portal switch 26
 - capacity, Meridian 1/Succession 1000 switch 78
 - engineering 77
 - multiple servers and CallPilot 110
 - multiple servers and Meridian Mail 111
 - supporting multiple servers 14
 - Symposium Agent 29, 46
 - Symposium Call Center Server 28
 - CPU impact of Symposium Web Client 69
 - hardware platforms 39
 - multiple on the same switch and CallPilot 110
 - multiple on the same switch and Meridian Mail 111
 - support for multiple servers on switch 14
 - Symposium customer model 134
 - Symposium Event Interface. *See* SEI
 - Symposium Voice Processing. *See* SVP
 - Symposium Voice Services on CallPilot
 - ACCESS traffic 109
 - and multiple servers on switch 15, 110
 - CLAN impact 110
 - CPU impact 110
 - ELAN impact 110
 - in the DMS environment 35
 - requirements 110
 - Symposium Voice Services on Meridian Mail
 - ACCESS traffic 109
 - and multiple servers on switch 15, 111
 - in the DMS environment 35
 - requirements 111
 - Symposium Web Center Portal 29
 - Symposium Web Client
 - and IP multicasting 151
 - and multiple servers on switch 14
 - architecture 64
 - client PCs 29
 - CPU impact on Symposium Call Center Server 69
 - documents 22
 - multiple application servers 69
 - requirements for client PCs 65
 - Symposium Web Client application server 28
 - CLAN/WAN impact 69
 - CPU utilization 66
 - impact on CLAN/WAN 69
 - requirements 65
- ## T
- TAPI 29, 33
 - and multiple servers on the switch 16
 - TDM telephony 28
 - Telephony Application Program Interface. *See* TAPI
 - temporary files 76
 - thick client. *See* Classic Client
 - thin client. *See* Symposium Web Client
 - timeouts 102
 - traffic
 - CLAN 92
 - ELAN 94
 - impact on DMS switch 88
 - WAN 101
 - typical call 42
- ## U
- unicast
 - sending and IP multicast sending 141
 - traffic 71
 - union break time 44

utilization

- ACCESS link 113
- maximum CLAN 93
- maximum ELAN 94

V

validating scripts 53

virtual memory

- Classic Client 76
- Symposium Call Center Server 52

virtual networks and IP multicasting 142

Virtual Private Network. *See* VPN

virus scanning 54

voice ports

- dedicating 107
- non-ACCESS 108
- number required 107

voice processing, engineering 105

voice sessions 109

- duration 107

VPN 117

- configurations 120
- remote support guidelines 119

W

waiting calls 87

WAN 30

dedicating 102

impact of Symposium Web Client application
server on 69

requirements 101

traffic 101

Wide Area Network. *See* WAN**X**

X11 releases 78



Reader Response Form

Nortel Networks Symposium Call Center Server
Product release 5.0
Planning and Engineering Guide

Tell us about yourself:

Name: _____

Company: _____

Address: _____

Occupation: _____ **Phone:** _____

1. What is your level of experience with this product?

- New user Intermediate Experienced Programmer

2. How do you use this book?

- Learning Procedural Reference Problem solving

3. Did this book meet your needs?

- Yes No

If you answered No to this question, please answer the following questions.

4. What chapters, sections, or procedures did you find hard to understand?

5. What information (if any) was missing from this book?

6. How could we improve this book?

Please return your comments by fax to 353-91-756050, or mail your comments to
Contact Center Documentation Research and Development Prime, Nortel Networks, Mervue Business
Park, Galway, Ireland.



Reader Response Form

Nortel Networks Symposium Call Center Server

Planning and Engineering Guide

Nortel Networks
Mervue Business Park
Galway, Ireland

Copyright © 2004 Nortel Networks, All Rights Reserved

Information is subject to change without notice. Nortel Networks reserves the right to make changes in design or components as progress in engineering and manufacturing may warrant.

The process of transmitting data and call messaging between the Meridian 1 and Symposium Call Center Server is proprietary to Nortel Networks. Any other use of the data and the transmission process is a violation of the user license unless specifically authorized in writing by Nortel Networks prior to such use. Violations of the license by alternative usage of any portion of this process or the related hardware constitutes grounds for an immediate termination of the license and Nortel Networks reserves the right to seek all allowable remedies for such breach.

Publication number:	297-2183-106
Product release:	5.0
Document release:	Standard 2.0
Date:	June 2004

