Avaya Contact Recorder

Release 15.2
Planning, Installation and Administration Guide

Issue 1.13
February 2021

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

"Documentation" means information published in varying mediums which may include product information, operating instructions and performance specifications that are generally made available to users of products. Documentation does not include marketing materials. Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of Documentation unless such modifications, additions, or deletions were performed by Avaya. You agree to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by You.

Link disclaimer

Avaya is not responsible for the contents or reliability of any linked websites referenced within this site or documentation provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Avaya provides a limited warranty on Avaya hardware and software. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya's standard warranty language, as well as information regarding support for this product while under warranty is available to Avaya customers and other parties through the Avaya Support website: http://support.avaya.com/helpcenter/getGenericDetails?detailId=C20091120112456651010 under the link "Warranty & Product Lifecycle" or such successor site as designated by Avaya. Please note that if You acquired the product(s) from an authorized Avaya Channel Partner outside of the United States and Canada, the warranty is provided to You by said Avaya Channel Partner and not by Avaya. **"Hosted Service"** means an Avaya hosted service subscription that You acquire from either Avaya or an authorized Avaya Channel Partner (as applicable) and which is described further in Hosted SAS or other service description documentation regarding the applicable hosted service. If You purchase a Hosted Service subscription, the foregoing limited warranty may not apply but You may be entitled to support services in connection with the Hosted Service as described further in your service description documents for the applicable Hosted Service. Contact Avaya or Avaya Channel Partner (as applicable) for more information.

Hosted Service

THE FOLLOWING APPLIES ONLY IF YOU PURCHASE AN AVAYA HOSTED SERVICE SUBSCRIPTION FROM AVAYA OR AN AVAYA CHANNEL PARTNER (AS APPLICABLE), THE TERMS OF USE FOR HOSTED SERVICES ARE AVAILABLE ON THE AVAYA WEBSITE, HTTP://SUPPORT.AVAYA.COM/LICENSEINFO UNDER THE LINK "Avaya Terms of Use for Hosted Services" OR SUCH SUCCESSOR SITE AS DESIGNATED BY AVAYA, AND ARE APPLICABLE TO ANYONE WHO ACCESSES OR USES THE HOSTED SERVICE. BY ACCESSING OR USING THE HOSTED SERVICE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE DOING SO (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THE TERMS OF USE. IF YOU ARE ACCEPTING THE TERMS OF USE ON BEHALF A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT THAT YOU HAVE THE AUTHORITY TO BIND SUCH ENTITY TO THESE TERMS OF USE. IF YOU DO NOT HAVE SUCH AUTHORITY, OR IF YOU DO NOT WISH TO ACCEPT THESE TERMS OF USE, YOU MUST NOT ACCESS OR USE THE HOSTED SERVICE OR AUTHORIZE ANYONE TO ACCESS OR USE THE HOSTED SERVICE.

Licenses

THE AVAYA GLOBAL SOFTWARE LICENSE TERMS FOR VERINT SOFTWARE PRODUCTS AVAILABLE ON THE AVAYA WEBSITE, HTTP://SUPPORT.AVAYA.COM/LICENSEINFO, OR SUCH SUCCESSOR SITE AS DESIGNATED BY AVAYA, ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS THE SOFTWARE (AS DEFINED IN THE AVAYA GLOBAL SOFTWARE LICENSE TERMS FOR VERINT SOFTWARE PRODUCTS), AND WHO PURCHASED THE LICENSE FROM AVAYA OR AN AVAYA CHANNEL PARTNER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AVAYA CHANNEL PARTNER. REFER TO THE AVAYA SOFTWARE LICENSE TERMS FOR VERINT SOFTWARE PRODUCTS FOR INFORMATION REGARDING THE APPLICABLE LICENSE TYPES PERTAINING TO THE SOFTWARE.

All Rights Reserved

Avaya and/or its licensors retain title to and ownership of the Software, Documentation, and any modifications or copies thereof. Except for the limited license rights expressly granted in the applicable Avaya Global Software License Terms for Verint Software Products, Avaya and/or its licensors reserve all rights, including without limitation copyright, patent, trade secret, and all other intellectual property rights, in and to the Software and Documentation and any modifications or copies thereof. The Software contains trade secrets of Avaya and/or its licensors, including but not limited to the specific design, structure and logic of individual Software programs, their interactions with other portions of the Software, both internal and external, and the programming techniques employed.

Virtualization

The following applies if the product is deployed on a virtual machine. Each product has its own ordering code and license types. Note that each Instance of a product must be separately licensed and ordered. For example, if the end user customer or Avaya Channel Partner would like to install two Instances of the same type of products, then two products of that type must be ordered.

Third Party Components

Certain software programs or portions thereof included in the Software may contain software (including open source software) distributed under third party agreements ("Third Party Components"), which may contain terms that expand or limit rights to use certain portions of the Software ("Third Party Terms"). Information regarding distributed Linux OS source code (for any Software that has distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the Software, Documentation or on Avaya's website at: http://support.avaya.com/Copyright (or a successor site as designated by Avaya). The following applies only if the H.264 (AVC) codec is distributed with the product. THIS PRODUCT IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO (i) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD ("AVC VIDEO") AND/OR (ii) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE HTTP://WWW.MPEGLA.COM

Service Provider

THE FOLLOWING APPLIES TO AVAYA CHANNEL PARTNER'S HOSTING OF AVAYA PRODUCTS OR SERVICES. THE PRODUCT OR HOSTED SERVICE MAY USE THIRD PARTY COMPONENTS SUBJECT TO THIRD PARTY TERMS AND REQUIRE A SERVICE PROVIDER TO BE INDEPENDENTLY LICENSED DIRECTLY FROM THE THIRD PARTY SUPPLIER. AN AVAYA CHANNEL PARTNER'S HOSTING OF AVAYA PRODUCTS MUST BE AUTHORIZED IN WRITING BY AVAYA AND IF THOSE HOSTED PRODUCTS USE OR EMBED CERTAIN THIRD PARTY SOFTWARE, INCLUDING BUT NOT LIMITED TO MICROSOFT SOFTWARE OR CODECS, THE AVAYA CHANNEL PARTNER IS REQUIRED TO INDEPENDENTLY OBTAIN ANY APPLICABLE LICENSE AGREEMENTS, AT THE AVAYA CHANNEL PARTNER'S EXPENSE, DIRECTLY FROM THE APPLICABLE THIRD PARTY SUPPLIER. WITH RESPECT TO CODECS, IF THE AVAYA CHANNEL PARTNER IS HOSTING ANY PRODUCTS THAT USE OR EMBED THE G.729 CODEC, H.264 CODEC, OR H.265 CODEC, THE AVAYA CHANNEL PARTNER ACKNOWLEDGES AND AGREES THE AVAYA CHANNEL PARTNER IS RESPONSIBLE FOR ANY AND ALL RELATED FEES AND/OR ROYALTIES. THE G.729 CODEC IS LICENSED BY SIPRO LAB TELECOM INC. SEE WWW.SIPRO.COM/CONTACT.HTML. THE H.264 (AVC) CODEC IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO: (I) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD ("AVC VIDEO") AND/OR (II) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE.ADDITIONAL INFORMATION FOR H.264 (AVC) AND H.265 (HEVC) CODECS MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE HTTP://WWW.MPEGLA.COM.

Compliance with Laws

You acknowledge and agree that it is Your responsibility for complying with any applicable laws and regulations, including, but not limited to laws and regulations related to call recording, data privacy, intellectual property, trade secret, fraud, and music performance rights, in the country or territory where the Software is used.

Preventing Toll Fraud

"Toll Fraud" is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud intervention

If You suspect that You are being victimized by Toll Fraud and You need technical assistance or support, call Technical Service Center Toll Fraud Intervention Hotline at +1-800-643-2353 for the United States and Canada. For additional support telephone numbers, see the Avaya Support website: http://support.avaya.com, or such successor site as designated by Avaya.

Security Vulnerabilities

Information about Avaya's security support policies can be found in the Security Policies and Support section of https://support.avaya.com/security Suspected Avaya product security vulnerabilities are handled per the Avaya Product Security Support Flow (https://support.avaya.com/css/P8/documents/100161515).

Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the Documentation, any Hosted Service(s), and product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the Documentation, Hosted Service(s) and product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party.

Avaya is a registered trademark of Avaya Inc.

All non-Avaya trademarks are the property of their respective owners.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support website: http://support.avaya.com, or such successor site as designated by Avaya.

Contact Avaya Support

See the Avaya Support website: http://support.avaya.com for product or Hosted Service notices and articles, or to report a problem with your Avaya product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: http://support.avaya.com (or such successor site as designated by Avaya), scroll to the bottom of the page, and select Contact Avaya Support.

Avaya Contact Recorder
Release 15.2

Planning, Installation and Administration Guide

About This Guide 15
- Intended audience 15
- Summary of information included in this guide 16
- Conventions used in this guide 18
- Additional references 20
- Support 21
- Document revision history 21

Chapter 1: System Overview 24
- Introduction 25
  - What's New 26
  - Recording Options 30
- Server components 42
  - Avaya Contact Recorder Server 42
  - Optional Server Applications 43
- End-User tools 45
  - Workforce Optimization ("WFO") 45
  - Search and Replay 45
  - Agent Initiated Monitoring 46
  - Administration Tools 47
- Recording Functionality 47
  - Sampled Recording for Quality Assessment 47
  - Bulk Recording 47
  - Ad-hoc or Occasional Recording Modes (Communication Manager only) 49
  - Replay Options 50
- Miscellaneous 50
  - Beep Tone 50
  - International support 51
  - Liability 51

Chapter 2: Planning and Prerequisites 52
- Introduction 53
- Recording Bandwidth 54

Voice Recording 54
Screen Recording 56
Storage Requirements 57
Storage at Each Recorder 57
Workforce Optimization ("WFO") 59
Central Database Storage 59
Archive Call Storage 59
Backup Storage 59
TDM Interfaces 60
Cards Supported 60
Chassis Requirements 60
Platform Restrictions 60
Server Platform 61
Sizing 61
Antivirus and other 3rd Party Applications 64
Common Problems 64
Virtualization 65
Network Issues 66
Load 66
Ports Used 66
Network Address Translation Routing 67
Licensing 67
Recording Limit 67
Backup Recording Channels Limit 67
Concurrent Screen Recording Limit 67
Quality Monitoring Seat Limit 68
Telephone Replay Channel Count 68
Dialer Integration 68
Secure Call Recording 68
Selective Recording 68
Live Monitoring 68
Search and Replay API 69
Mass Export 69
Timed Trials 69
Communication Manager System Prerequisites 69
Communication Manager 69
Gateway Resources 70
Session Border Controller (SBC) 71

AE Services 71
Expansion Interface Boards (TN570) 72
C-LAN 72
VoIP Resources 73
Multi-Connect Capacity 76
DMCC (IP_API_A) Licenses 76
TSAPI Licenses 76
VoIP Network Design 76
CS1000 System Prerequisites 77
Contact Center Requirements 77
CS 1000 Systems and IP Client Requirements 77
Avaya Oceana™ System Prerequisites 78
AACC System Prerequisites 78
Supported Topologies 78
Required Components 78
Topologies 79
Bulk Recording System 79
Bulk Recording + Quality Monitoring System 79
Large Bulk Recording Systems 80
External Recording Control 83
Business Rule Driven Recording 83
Desktop Process Analytics ("DPA") Control 83
Agent Initiated Monitoring ("AIM") 84
Auto-dialer Integrations 84
Custom Integrations 84
Chapter 3: Installation 85
Overview 86
Avaya System Configuration 87
Prerequisites 87
Communication Manager Configuration 87
CS1000 Configuration 95
AACC Configuration 98
Test Phonesets 99
Order in which to Install Applications 99
Platform Prerequisites 100
Linux 100
Windows 105
Antivirus and Backup Software 106

Time Synchronization 107
Java Timezone (TZ) Update 107
Network Connectivity 107
Installing Avaya Contact Recorder 108
Linux 108
Windows 110
Updating Components 112
PostgreSQL 112
Java 113
Tomcat 113
Installing Workforce Optimization ("WFO") 114
Installing Screen Capture and Agent Initiated Monitoring (AIM) Software 114
Chapter 4: Configuration 115
Accessing the System 117
URL 117
Initial User Account 117
Key Points 118
Licensing 119
Terminology 119
Obtaining a License Activation Key 120
Central Replay Servers 121
Standby and Slave Servers 121
Adding additional licenses 121
Reinstalling on the same PC 122
Reinstalling the Recorder on a new PC 122
Security 122
Securing the System 124
Windows Authentication 125
Windows Accounts for Screen Recording or AIM 125
General Setup 126
Server 126
Data Sources 135
TDM Tap Points 155
Email Configuration 158
System Monitoring 159
Via the Administration Pages 159
Via Email 160
Application Logs 160

Tomcat Logs 161
Remote logging via Syslog Server(s) 161
SNMP 162
Operations 163
Common Settings 163
Assigning Ports 165
On Demand Recording (Communication Manager only) 176
Meeting Recording (Communication Manager only) 178
(Telephone) Replay Ports (Communication Manager only) 180
I/O Jobs: Archive, Import, Mass Export 181
Managing I/O Jobs 182
File Storage 185
Exported Files 190
Advanced Settings 191
Search, Replay and Live Monitor 198
Access Rights 199
Layout Builder 202
Forwarding Replay Requests with NAT and/or SSL 208
Client Prerequisites 209
Restricting Access to Replay Layouts 211
Miscellaneous Security Features 211
Locking Recordings 211
Replay Authorization Process 213
Modify Default Behavior 215
Backup/Restore 216
Application 216
Backing up the Database 216
Restoring data to a new PostgreSQL database 218
Backing up Voice Recordings 219
Distributing User Instructions 221
Those Using Recording 221
Those entitled to replay calls 222
Configuring Avaya Support Remote Access 222
Chapter 5: Operations, Administration & Maintenance 223
Introduction 224
Status Monitoring 224
System 224
Server 227

CTI Monitors 229
Ports 231
Alarms 233
Audit Trail 234
Preventative Maintenance 235
Daily 235
Weekly 236
Monthly 237
Restarting the System 238
Chapter 6: System Security 239
Access to the Recorder 240
Windows Domain Authentication 240
Single Login 244
Dual Sign-in 244
Use of SSL 245
Allow search and replay from this server? 247
Session Inactivity Timeout 247
Minimum Password Length 247
Force strong password 247
Password expires after (days) 247
Password cannot be reused within (days) 248
Minimum changes between reuse of same password 248
Repeated Login Attempts 248
Legal Notice on Login Page 248
Disable Autocomplete 248
Replay Authorization Process 249
Server Hardening 249
Linux 249
Windows 250
ACR Firewall ports 250
Changing Passwords 251
User Accounts 251
PostgreSQL Database Owner 251
Encrypting Connections 252
Computer Telephony Integration 252
Session Initiation Protocol (SIP) 253
Audio Streams 253
Screen Recording 254

Encrypted File Storage 255
Enabling Encryption 255
Slow Start-up under Linux on Virtual Machines 259
I/O Jobs 259
Finger-print Validator 260
To Enable Finger-printing 260
To Validate a File 260
I/O Jobs 260
Manual, Ad Hoc Requests 261
Mass Export 261
Distributed Replay Server 261
PCI Compliance 262
Chapter 7: Advanced Configuration 264
Properties File 265
Slave Server 280
Installation 280
Licensing 280
Configuration 280
Major Switch Configuration Changes 281
User Accounts 281
Search and Replay 281
Standby Server 281
Design and Planning 281
Installation 281
Licensing 282
Configuration 282
Central Replay Server 284
Archives 284
“Turbo” Mode 285
Installation 287
Configuration 287
Configuring other Recorders 288
Distributed Replay Server 288
Features 288
Configuration 289
Deliverables 290
Usage Report 291
Enabling the Report 291

Content 291
Accessing through URL: 291
Accessing the Usage report in a log file 292
Party Statistics Usage Report 292
Enabling the Report 292
Content 292
Accessing through URL 293
Accessing the Party Statistics report in a log file 293
Automated Party Statistics Reporting 294
Configuration Report (Communication Manager only) 295
Background 295
Automatic Configuration Reporting 295
Configuration 296
File Management 296
File Contents 296
Column Definitions 297
Selective Record Barring 297
Configuration 297
Example 298
Limitations 298
After Call Work without Business Rules 298
Altering Translations 298
Appendix A: Technical Reference 299
Recording files 301
WAV files 301
XML files 301
SCN files 301
Internal Database 302
Recording details 302
Configuration details 302
Recorder Interfaces 302
HTTP/HTTPS Interfaces Offered 303
Communication Manager 304
CS1000 305
Avaya Aura® Contact Center 306
Avaya Oceana™ 306
Other Data Sources 306
Screen Recordings 306

Workforce Optimization ("WFO") 307
Phone Replay 307
Live Monitor (Audio) 307
Other Recorders 308
External Control Interface 308
AET/DPA Interface 308
Database Upload Interface 308
Summary 309
Recording Attributes 312
Overview 312
Definitions 313
Call Identifiers 317
User Defined Fields 318
Search and Replay Attributes 319
WFO Integration 328
Appendix B: Troubleshooting 338
Hints and Tips 339
Where to Look for Clues 339
Determining Current Version 339
Advanced Diagnostics 341
Installing PSTools (Windows only) 341
Finding the ACR process ID 341
Linux 341
Windows 341
Creating thread dumps 341
Dumping the Java heap 342
Specific Problems 343
System Administration page problems 343
Connectivity 344
Search and Replay problems 344
Recording Problems 349
Screen Recording Problems 350
Live Monitor Problems 350
Appendix C: Alarms 352
Alarms 353
Alarms Table 353
Appendix D: High Availability (“HA”) 375
Fully Parallel Systems 376

Communication Manager 376
CS1000 377
Central Replay Server for Parallel Systems 377
High Availability Recording 378
Multiple Switches 381
CS1000 381
Communication Manager 381
High Availability Storage 382
RAID Storage 382
NOT Standby Recorders 382
NOT Disk Shadowing or Backup 382
Storage Attached Network (SAN) 382
Archival to Network Attached Storage (NAS) 383
High Availability Retrieval 384
High Availability CM Topology 385
Server Failure 385
Media Processing Resource Failure 387
Network Failure 387
Survivable Recording 388
Known limitations 394
Standby recorders and Unify/External Control 394
Controller connected to Master & Standby 394
Controller connected to Master only 394
Mode of operation 395
Fall-back Detection Beacons 395
Primary v Secondary Data Sources 396
Power-On 396
Standby mode 396
Control Priority 397
Failure Detection 397
Disk Space Monitoring 398
Active mode 398
Return to Standby mode 398
Switchover Implications 399
Restoring the Master 399
Comparison with hardware switch-over units 400
Appendix E: Non-standard Hardware 401
Overview 402

Disks 402
NICS 402
Appendix F: Advanced Security Settings 403
SSL Certificates 404
Backing up the Keystore files 404
Creating a new Certificate 404
Generating a Certificate Signing Request 405
Importing the CA's certificates 405
Backing up the keystore file 407
Creating and Signing Certificates Internally for RSA KMS and SSL 407
Adding AES CA Root Certificates to ACR 414
Changing Tomcat Port Numbers 414
Encrypting Properties File entries 415
Changing the Windows Service Account 415
Appendix G: External APIs 416
HTTP/HTTPS access 416
Authentication and Authorisation 416
WebXAPI 416
Error/Success codes 416
Fault tolerance 417
Controllable devices 417
Commands 418
Appendix H: GDPR 420
Introduction 420
GDPR Delete 421
Authentication and Authorisation 423
Search and Delete 424
Restful API 425
Status 425
Purging Delete Status Job Information 425
Cancelling Jobs 425
Viewing Job Details 426
Audit Trail 426
Limitations 426
Glossary 427

About This Guide

The *Avaya Contact Recorder Planning, Installation and Administration Guide* provides details of the Avaya Contact Recorder system, as well as recommended and required components.

Intended audience

This guide is intended to be used by:

- Pre-sales Systems Engineers developing system topologies and designs
- Professional Services staff installing and deploying systems
- Systems Administrators
- Support personnel

The reader is expected to be familiar with:

- System administration of Microsoft Windows servers and/or Linux servers
- TCP/IP Networking and Voice over IP (VoIP)
- Avaya contact center systems administration

Summary of information included in this guide

The following table provides information about this guide.

Chapter Title	Description
Chapter 1: System Overview	This chapter provides an overview of the design options for an Avaya Contact Recorder system.
Chapter 2: Planning and Prerequisites	This chapter gives details of the prerequisites for an Avaya Contact Recorder system. You should also review Chapter 6 System Security as some of the optional elements described there may also require additional cost and/or effort.
Chapter 3: Installation	This chapter gives details of the steps to install an Avaya Contact Recorder system.
Chapter 4: Configuration	This chapter gives details of the steps to configure an Avaya Contact Recorder system.
Chapter 5: Operation, Administration and Maintenance	This chapter provides details of regular maintenance required for an Avaya Contact Recorder system.
Chapter 6: System Security	Security of customer recordings is very important. This chapter discusses the various features - some optional - that you can use to ensure the safety and integrity of recordings.
Chapter 7: Advanced Configuration	This chapter provides an overview of the more complex and rarely used options for an Avaya Contact Recorder system
Appendix A: Technical Reference	This appendix provides technical details about the Avaya Contact Recorder system.
Appendix B: Troubleshooting	This appendix covers general troubleshooting tips and specific common issues.

Chapter Title	Description
Appendix C: Alarms	This appendix provides details of the alarms that can be raised by the system.
Appendix D: High Availability	In addition to using fault tolerant components within servers as described in the High Availability Systems section, recording systems can be made tolerant of many server and network failure conditions. This appendix details how such systems are designed and configured, how they handle failures, and how to upgrade them.
Appendix E: Non-standard Hardware	This appendix discusses considerations for non-standard hardware such as blade servers
Appendix F: Advanced Security Settings	This appendix discusses some features and prerequisites for advanced security.
Appendix G: External APIs	This appendix discusses how ACR can be controlled by external commands over HTTP/HTTPS.
Appendix H: GDPR	This appendix discusses how the ACR deletion process works and its main features.
Glossary	The glossary defines the terms you need to understand this manual.

Conventions used in this guide

The following table shows how user input, output and instructions are highlighted in this guide, as well as special notations that you will see as you use this guide.

To show...	This style is used	For example...
Information shown on screen	`Fixed width`	You should see the prompt below: `login:`
Characters that you should type exactly as shown	**`Fixed width, bold`**	Enter the following command: **`mount /mnt/cdrom`**
Characters that you should replace with appropriate information	***`Fixed width, bold italic`***	Browse to the new server by entering **`http://`*`servername`*`:8080`**
Menu selections, buttons and tabs	**Sans Serif, Bold**	Click on the **Install** button.
Helpful hints that can improve the efficiency or effectiveness of your work	**Tip:**	**Tip:** If no part-time licenses are available, a full time license may be used instead.
Important details that we want to make sure that you do not overlook	**Note:**	**Note:** Media Encryption may or may not show up on this form.
Advice that can help you avoid undesirable results	**Important:**	**Important:** If the network does not meet the three conditions listed, there will be no media resources.
Situations that can result in: • Harm to software • Loss of data • An interruption in service	**CAUTION:**	**CAUTION:** Perform this procedure only after normal business hours. This procedure restarts all links on the interface, and can cause a temporary loss of service.

To show...	This style is used	For example...
Situations that can result in harm to hardware or equipment	**WARNING:**	**WARNING:** Make sure that the disks are the Update you require. Red Hat and other vendors still sometimes supply Update 0 disks.

Additional references

- Avaya Contact Recorder Version 15.2 User Guide
- Avaya Contact Recorder Version 15.2 Upgrade Guide
- Avaya Contact Recorder VMWare Guide, Release 15.2
- Avaya WFO 15.2 Distributor Technical Reference (DTR)
- Avaya Communication Manager Call Recording: A Design Approach for Device Media and Call Control (DMCC, previously called CMAPI) (Compas ID 128862)
- Avaya WFO Agent Initiated Monitoring (AIM) Quick Reference Guide
- Avaya Contact Recorder: Migrating from Central Archive Manager
- Avaya WFO Security Configuration Guide
- Avaya Contact Recorder Integration to Workforce Optimization Guide
- Avaya Aura® Contact Center, Planning and Engineering Guide (NN44400-210)
- Avaya Communication Manager Guide to ACD Contact Centers
- Administrator's Guide for Avaya Communication Manager
- Administration for Network Connectivity for Avaya Communication Manager
- Avaya Aura Communication Manager Survivability Options
- Avaya Aura® Communication Manager Feature Description and Implementation
- Avaya Desktop Applications Deployment Reference and Installation Guide

Note:

Avaya Communication Manager documentation is available through the Avaya online support Web site, http://www.avaya.com.

The latest information on Version Compatibility is at:

https://secureservices.avaya.com/compatibility-matrix/menus/product.xhtml?name=Avaya+Aura+Workforce+Optimization&version

Support

- Go to the Avaya Support site at http://support.avaya.com for the most up-to-date documentation, product notices, and knowledge articles. You can also search for release notes, downloads, and resolutions to issues.
- Use the online service request system to create a service request. Chat with live agents to get answers to questions, or request an agent to connect you to a support team if an issue requires additional expertise.

Document revision history

Revision	Description of change
1.00	Initial Publication
1.01	The following changes have been made in this revision: • Added section on Tracking Delete Request jobs on the new status tab • Update the TSAPI configuration section • Section on 16 digit support in CM 8.0 added • Major updates to Enabling Encryption section, including the details for using Thales Data Security Manager • Updated the Alarms table with new warnings • Updates to the certificates information in the Advanced Security Settings section • Replaced any references to KMS with RSA KMS • Updated the GDPR Delete section • Added details of the new Status section in the GDPR section.

Revision	Description of change
1.02	The following changes have been made in this revision: • Updated Install on Windows instructions. • Updated Install on Linux instructions. • In ‘Locking Recordings’ section, under ‘How it works’, add a note about calls recorded in a V12.0 system or earlier with screen retention period shorter than the audio period. • In ‘Locking Recordings’ section, under ‘Multiple Server Systems’, update that the locking feature is only allowed on the recorder that holds the consolidated database. • In ‘Exported Files’, under ‘“Stitching” of Recording Files’, update step #2. • Appendix C, add a new alarm: ‘alarms.privviolation’ • Removed references to Java Unlimited Strength Policy.
1.03	• Windows (SMB) Fileshare port for both SMB V1 and SMB V2. Also added to the What’s new in 15.2 list. • Purging Delete Status Job Information section moved under Status heading. • New limitation to the Limitations section regarding manually deleting contacts. • Under On Demand Recording and Meeting Recording, added note about calls not being segmented.
1.04	• Fix typos
1.05	• Added RSA KMS to Thales KMS key migration • Renamed Appendix G: WebXAPI Control Interface to Appendix G: External APIs • Added a new WebXAPI heading under External APIs • Removed <basicauthenticationtag> from Appendix G: External APIs
1.06	• Under ‘Dumping the Java heap’, fixed the command for dumping memory on Windows
1.07	• Added additional entries in tables Chapter 7 • Added additional response information in WebXApi • Appendix G External APIs Error/Success Codes Section

Revision	Description of change
1.08	• Under Controllable Devices section, replaced whole section with updated content. • Under Commands section, in “Start” and “Stop” sections, replaced command endings with “&user=user1@test.com”
1.09	• Update that soft phone set type can be configured with 96xx in addition to 4624.
1.10	Replace ‘VMWare ESXi 5.0 or 5.5’ with ‘VMWare ESXi’.
1.11	• Under Purging the Locked Folder section, replaced text on button to “Purge Locked Recordings Folder” • Under Domain Name Server (DNS) Entries, additional Important Note. • WebXAPI addition, under Pause/Resume section, additional sentence added to end of section. • Under Properties File section, replaced entry in rec.maskallowed and in recorder.pool. • Under Windows section, removed sentence from Note. • Under AMS Zone to Recorder/Pool Mappings replaced section.
1.12	Under Error/Success codes, added response code 501 NOT IMPLEMENTED.
1.13	Replace ACR Activation site https://oaccess.verint.com/acr with https://licensing.verint.com/acractivation.

Chapter 1: System Overview

This chapter provides an overview of the design options for an Avaya Contact Recorder system.

Important:

Ensuring that a call is recorded and stored reliably in accordance with the needs of the business depends on careful administration, on-going checks, and for mission critical deployments, redundant configurations. It is important to read and follow the advice given in this manual and in others such as those listed under Additional references on page 20, and to understand the importance of robust recorder and archive deployment, configuration and management.

As with any system, external system changes, equipment failure, errors within the system, errors in the system configuration, and human error can result in the system either not keeping the calls you want or recording a call that you are not supposed to record. To protect against critical system failure, redundant recording and CTI solutions should be deployed where possible. In all cases, careful monitoring of alarms and systems must be routinely undertaken. To protect against human error or unexpected behavior from any configured recording or archive rules, checks should be made regularly to ensure that the system is recording and archiving, and that only the required calls are recorded and archived.

Failure to do so may result in data loss, data corruption, or recording of unintended calls.

Introduction

Avaya Contact Recorder provides an extremely efficient and scalable, voice recording platform, running on standard PC Hardware.

This chapter describes how Avaya Contact Recorder can be used to record calls on several different Avaya switch types. It describes the Avaya Contact Recorder (which is a mandatory component) and the optional components.

Example Topology. ACR deployed against Communication Manager.

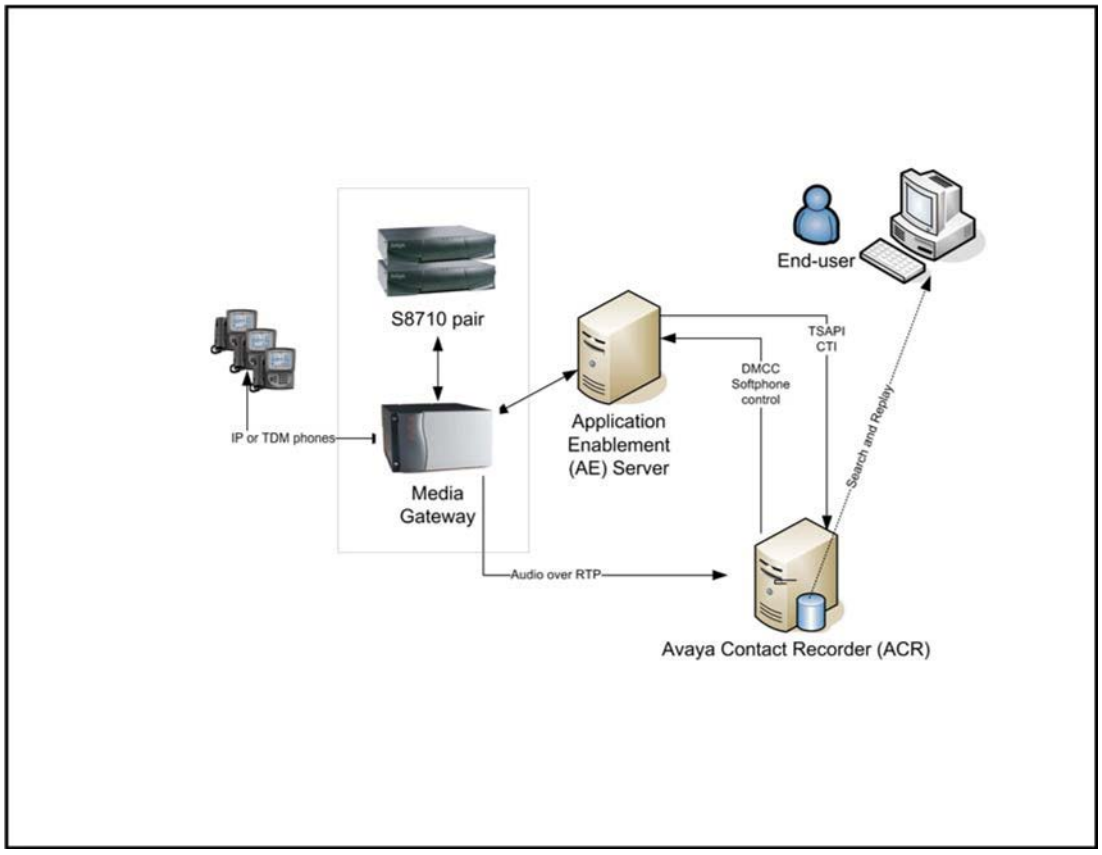

What's New

In Version 12.1

- Faster retrieval from NAS archive storage.
- Can retrieve recordings from EMC that were archived there using Central Archive Manager (CAM) 7.8 (See separate Technical Note for import instructions).
- Supports "After Call Work" setting on WFO business rules.
- Supports Agent Initiated Monitoring (AIM). This replaces the previous "Contact Recording Desktop" application.
- Supports screen only recording under AET/DPA control
- Supports percentage sampling of screen recordings and designating specific recorders/pools to record screen content from particular populations of workstations.
- Passive IP recording supported in multi-server topologies.
- Consistent left/right assignment of agent/other(s) audio in stereo recordings.
- Live monitor of voice and screen content for bulk recordings (optional, additional license required).
- Support for Workforce Optimization's live monitoring service allowing real-time evaluation of calls as they are recorded.
- **System Admin.** and **Restricted Admin.** User account details now passed automatically to slave recorders.
- Layouts and User account details now passed automatically from master Central Replay Server to secondary Central Replay Server.
- Screen recording content encryption between recorded desktop and Avaya Contact Recorder (optional).
- Option to block **Restricted Admin.** role from administering user accounts.
- Support for Avaya Enterprise Manager (Communication Manager Elite).
- Search and Replay by session as alternative to call segment.
- AACC SIP recording enhancements: AMS Zone support - allowing specific recorders/pools to be assigned to each AMS zone for better network traffic management; encryption of signaling ("sips") and audio streams ("sRTP").
- Support for SIPREC recording with Session Border Controllers on Communication Manager (only).

In Version 15.1

The following features, developed in 15.1 have also been back-ported into 12.1:

- Allows multiple “owners” of recordings, which may be additional to or replace the default ownership.
- Supports archive to SFTP.
- Support for Java 8 and Red Hat Enterprise Linux (RHEL) 7.
- Linux “kickstart” process now via http and supports additional disk and network options.
- Tolerates (does not attempt to record) video streams within SIPREC calls.
- Automatic busy-out of failing DMCC softphones.
- Screen recording tolerates Windows 8.x “Desktop Window Manager” dummy accounts.

These features are available from 15.1 onwards:

- Compatible with Communication Manager 7.
- Compatible with Avaya WorkForce Optimization (WFO) 15.1.
- Compatible with RSA Key Management Server (KMS) 3.2.1.
- Multi-zone support for Proactive Outreach Manager (POM) 3.0.1.
- Backward compatibility of 15.1 Slave recorder with 12.0 and 12.1 recorders (CRS must be upgraded to 15.1 first).
- Central Replay Server supports “turbo” search mode across any search and replay layouts (increase disk space required).
- Archiving, Import and Mass Export functions consolidated and enhanced into comprehensive new “I/O Jobs” framework. Note that Mass Export is an optional, licensable feature.
- “Distributed Replay Server” allows outsourcers to provide their customers with a dedicated search and replay server and to feed it with the appropriate subset of recordings.
- Maintenance mechanism to allow repopulation of WFO database with recordings.
- Alternative beep-tone generation option.

Note that 15.1 removes support for:

- Red Hat Enterprise Linux 5
- Windows Server 2008 R2
- Communication Manager 5.x

In Version 15.1FP1

- Support for Avaya Oceana™
- On Demand and Meeting recording modes now use TSAPI for call tagging.
- SNMP traps can be sent to multiple destinations and use the configured community name.
- Options to allow the **Restricted Admin.** role to use the **Status > CTI Monitors** and/or **Ports** pages.
- A patching tool is provided to simplify and reduce risks associated with patching the recorder.
- DMCC softphones can now use srtp-aescm128-hmac32-unauth encryption on systems where the AES supports this scheme.
- Archiving to removable optical media (DVD and Blu-ray) is deprecated.
- After Call Work automatically stops as soon as the next contact is connected so as to avoid overlap of screen recording (revert to previous behavior with property setting **`acw.earlyterminate=false`** if required).
- You can specify a default After Call Work duration using property setting **`acw.default=`*nn*** where ***nn*** is in seconds. This applies to all recordings for which no WFO Business Rule has fired.
- A remote standby will only apply TSAPI observers while it is active – reducing the load on the system during “sunny-day” operation.
- Updated RSA KMS client with improved features and better security.

In Version 15.1FP2

The following features, developed in 15.1FP2, have also been back-ported into 15.1FP1:

- Stronger encryption.

These features are available from 15.1FP2 onwards:

- Search and Replay (but *not* Live Monitor, Telephone Replay nor Export from the replay screen) are possible via additional browsers (previously only via Internet Explorer). ACR 15.1FP2 has been evaluated and tested with VPAT standards when using Internet Explorer 11. VPAT testing for additional browsers will be included in a future release.
- Within the Linux “Kickstart” tool, the option to partition the disk optimally for a Central Replay Server – giving more space to the calls database and less to the calls buffer.
- Previously, you could only search for recordings in Session-based layouts if your replay rights included the party which the session was tracking. This was problematic

for parties with many digits in their address e.g. international numbers. Now, you can search for sessions if you have replay rights over any of the underlying recording segments in the session.

- High Availability configuration and operation has been significantly enhanced. See Appendix D: High Availability on page 375 for full details. This includes:
  - Explicit **Full** versus **Partial Standby Coverage**.
  - Explicit **Warm** (TSAPI observers only when needed) versus **Hot Standby Readiness**.
  - Explicit relative **Standby Priority**.
  - Geo-redundant High Availability (GRHA) AES support.
  - Automatic handling of fragmentation and subsequent fallback via **Beacon Softphones**.
  - External control via HTTP/HTTPS as detailed in Appendix G: External APIs on page 416

The following features are deprecated as of 15.1FP2:

- AACC calls are no longer to be recorded via SIP and the AMS. Use DMCC ports or SIPREC instead. Existing users will find their systems continue to use SIP recording until they deliberately disable this via the AACC's administration page under **General Setup** which they should do only after they have provided sufficient alternative recording capacity.

  Note:

  The use of DMCC recording instead of SIP recording should only be used on AACC 7.x or later. Another consideration is the availability of DMCC recording capacity on the CM. Note that CM 7 offers the ability to use AMS for DMCC recording which may also be a relevant factor. The only reason for a new system to use SIP via AMS is if either (a) beep-tone is to be injected or (b) AACC pre-dates Version 7.0 or (c) the solution specifically requires the AMS zoning feature.
  In this case, set property **`aacc.sipenable=true`** and turn on SIP recording via the AACC's **General Setup** administration page.

- The **Designated Recorder/Pool(s)** setting now only supports named pools – not individual recorder serial numbers. These will continue to work in this release if already configured but should be migrated to pool names. Specifying individual recorders leaves the system subject to problems on upgrade should a recorder be retired or replaced with another.

15.1FP2 removes support for:

- Archiving to removable optical media (DVD and Blu-ray). Existing customers already using these media may continue to do so until the next release. Replay from media already written will still be supported but newly recorded calls will have to be archived to fixed storage instead (NAS, mounted drive or symlink, SFTP or EMC).

- Avaya Aura® Contact Centre running on switches other than Communication Manager or CS1000.
- Locally configured Standby recorders (**`standby.localconfig=true`**). Use the **Designated Pool(s)** setting instead to determine behavior during failover. See Appendix D: High Availability on page 375 for more details.

In Version 15.2

These features are available from 15.2 onwards:

- The screen replay window can now be toggled between Scale-to-Fit (where the screen is scaled to display in full) or Native (where the screen is displayed unadjusted).
- Reporting and Deletion of call records via the Administration screen, as part of the support of the EU 'General Data Protection Regulation'.
- Tracking of GDPR Delete Request jobs via a new Status tab, including view and cancel options.
- Support for the latest version of RSA Data Protection Manager (3.5.2).
- Support for Thales Data Security Manager (6.0.2).
- Deprecation of use of JKS keystores files in favour of PKCS12 files.
- Support for Windows SMB V2 fileshares (and deprecation of SMB V1 which can still be enabled with a property setting).

Recording Options

Avaya Contact Recorder can be used to record telephony calls made on one or more Communication Manager, CS1000 or Avaya Aura® Contact Center systems. The ways in which calls may be recorded are described below. Where more than one recording mechanism is possible, the highest priority mechanism is normally used, as shown in the table below.

Priority	Description
Highest Priority	TDM Station-side tap
	TDM Trunk Side Tap
	SIPREC
	Passive IP

Lowest Priority	DMCC Single-step conference (Communication Manager) or Duplicate Media Streaming (CS1000)

Communication Manager

Avaya Contact Recorder can record calls made on a Communication Manager system in several ways:

Trunk-side Tap

Where the bulk of telephone sets are not IP or the number of trunks carrying recordable calls is limited, tapping into the trunks may be appropriate. Using E1 or T1 passive tap cards, Avaya Contact Recorders can record any of the time slots as shown below. When using this method, internal calls cannot be recorded. This makes it unsuitable for quality monitoring.

Trunk-side Recording on Communication Manager

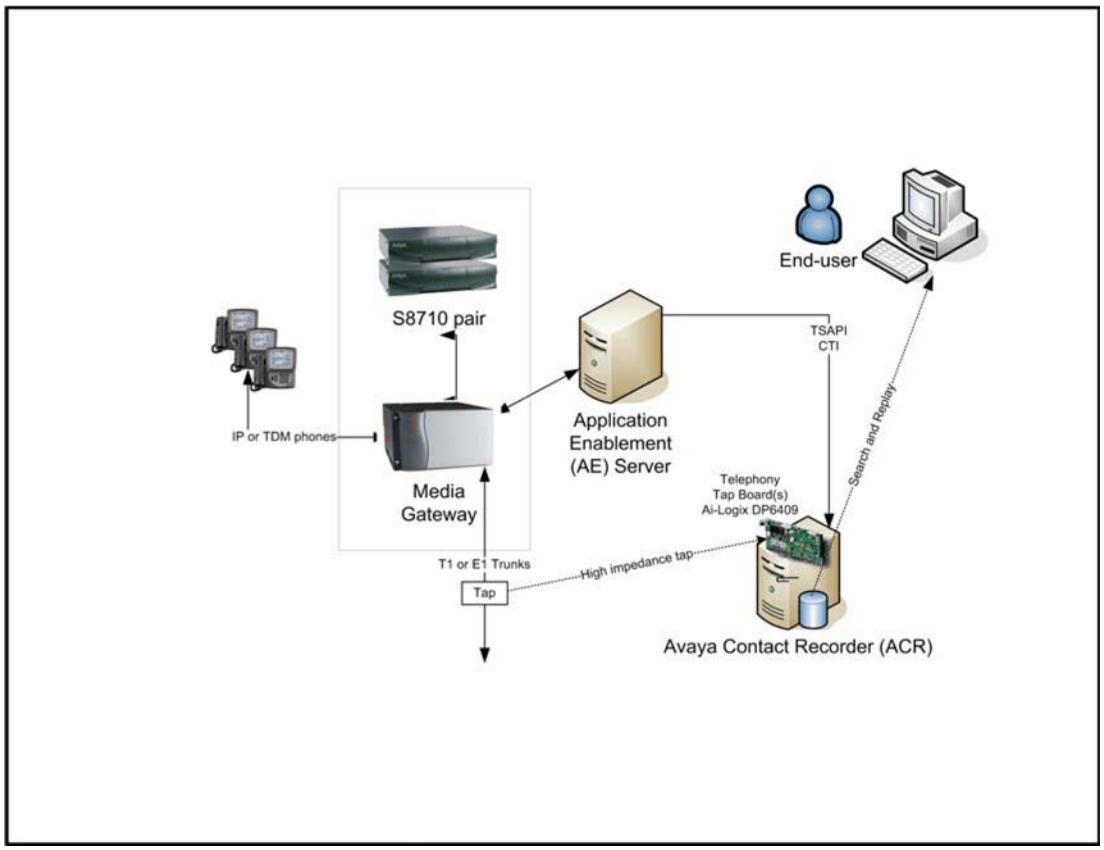

TDM Station-side Tap

Where the telephone sets are digital and you only need to record a limited number of telephones, tapping into the extensions may be appropriate. Using digital extension tap cards allows an Avaya Contact Recorder to access the audio passing to or from the tapped telephones as shown below.

TDM Station Recording on Communication Manager

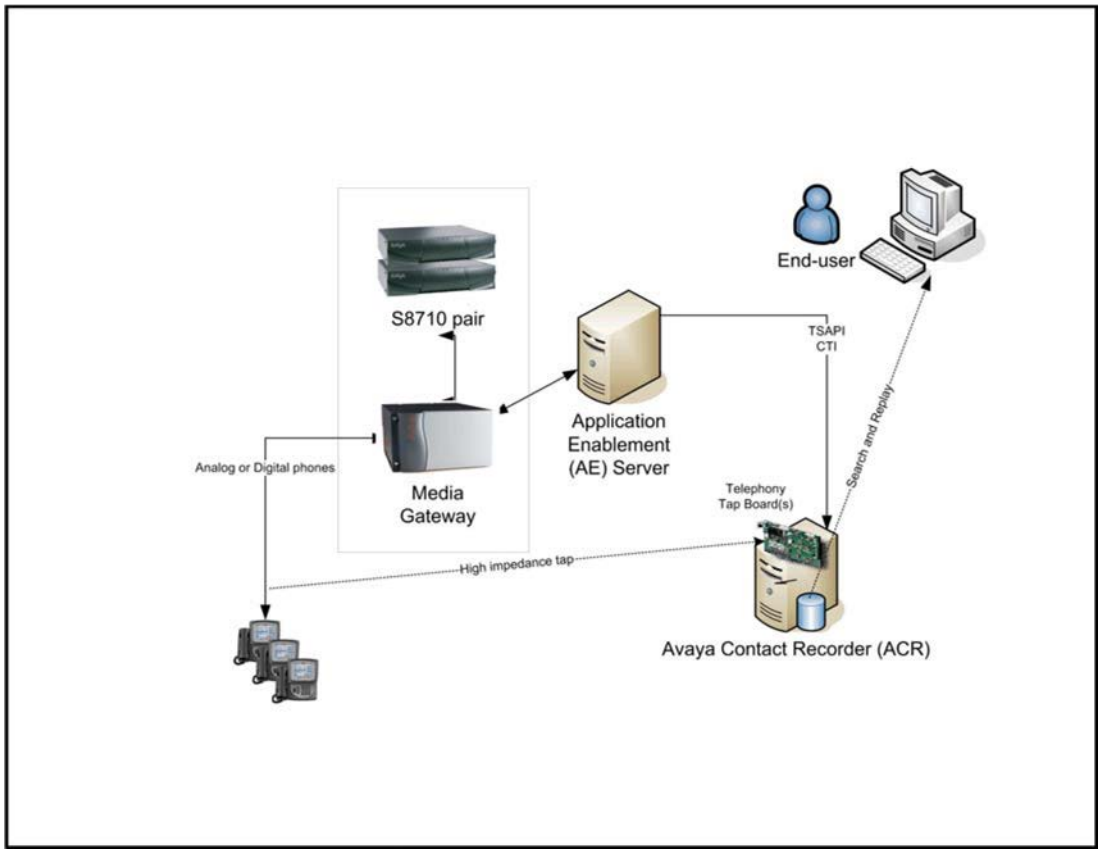

Passive IP Tap

Where IP phones are used and the audio transmitted is unencrypted, a passive IP tapping solution may be appropriate. This does not load the Communication Manager itself as a conferenced approach would. In this approach, one or more additional Network Interface cards (NICs) tap into the Ethernet segment(s) over which the voice traffic is flowing - as shown below.

Passive IP Tapping on Communication Manager

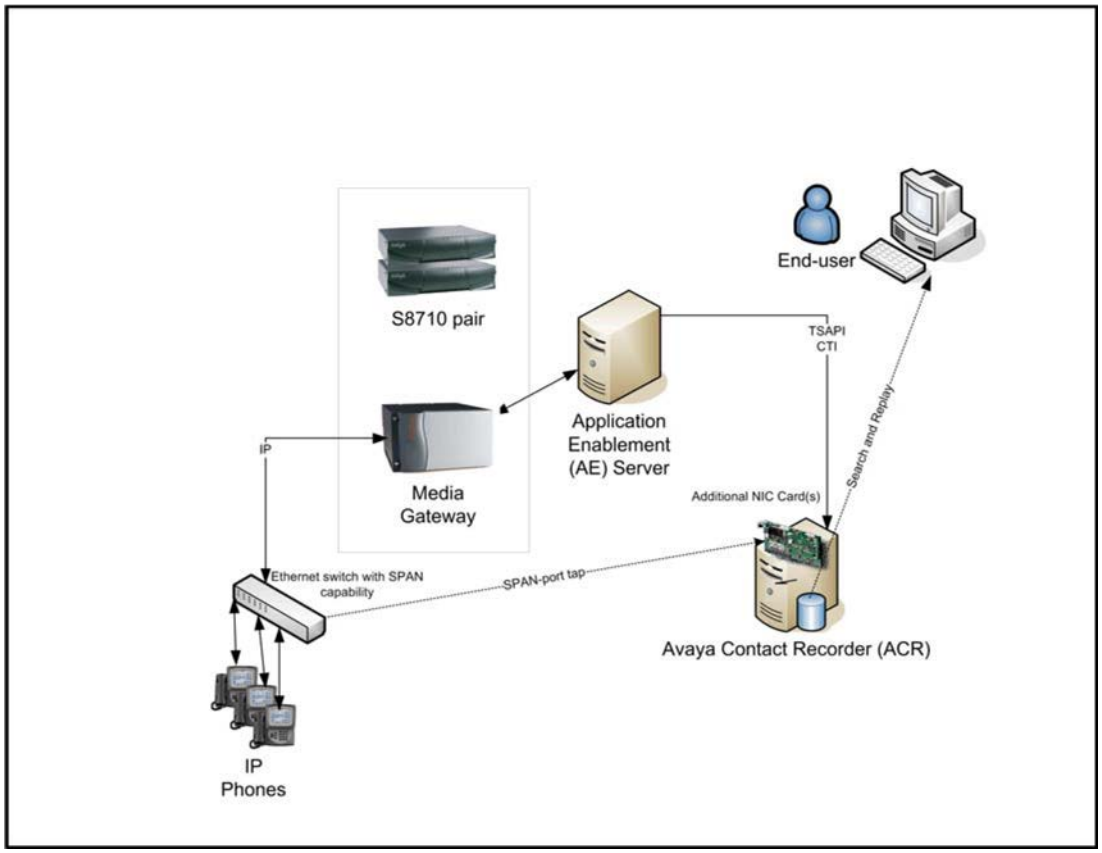

Limitations

1. As the passive IP recorder first has to learn the location of IP phones, it may not be able to record the first call made on each station after the recorder starts up. You should therefore also provide a small number of DMCC Softphone recording ports which will be used automatically in cases where passive IP recording cannot be performed.
2. OneXAgent softphones can be recorded but must be H.323 (not SIP) and Version 2.5.5 or later.
3. SIP phones cannot be recorded in this way, only H.323 phones.

SIPREC Recording via Session Border Controller (SBC)

Avaya Contact Recorder can record calls that are routed to or from the Communication Manager (only) via a Session Border Controller ("SBC") using that device's SIPREC recording capabilities as shown in the diagram below. These recordings are made in stereo - with the internal party(ies) audio in one channel and that of the external party(ies) in the other.

SIPREC Recording on Communication Manager

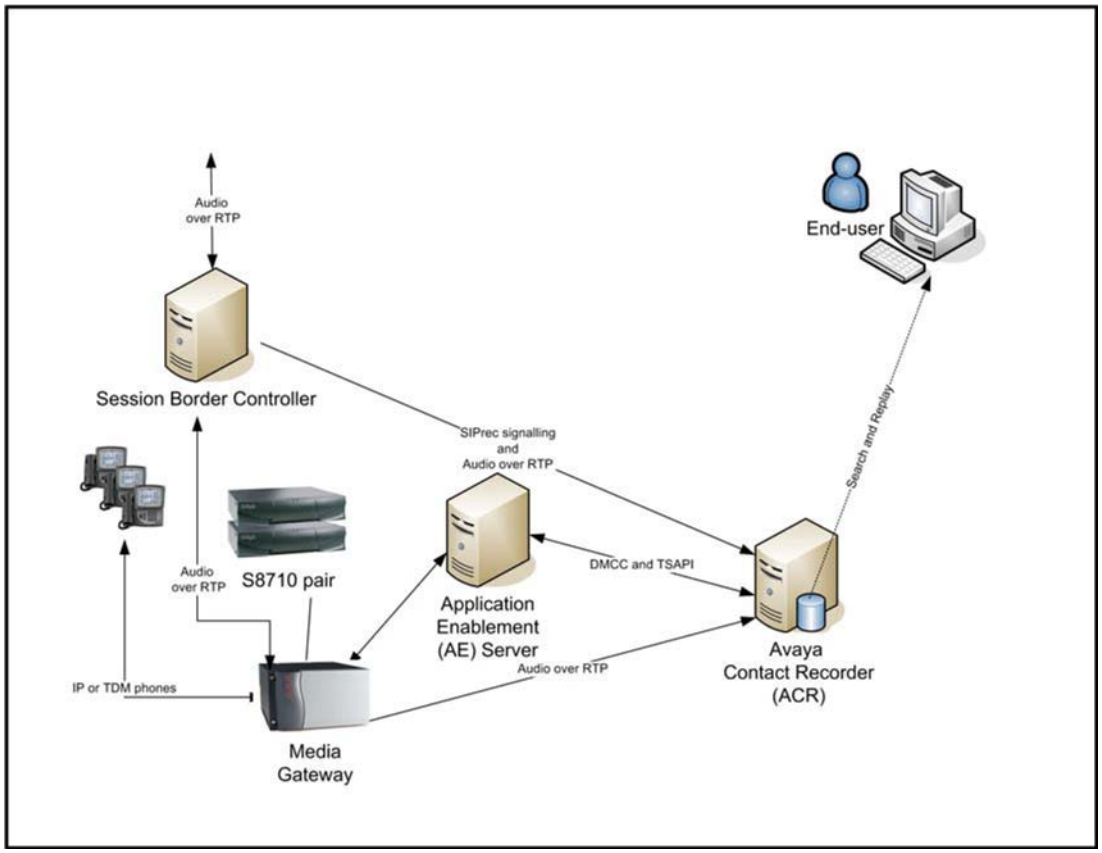

Internal calls, however, do not pass through the SBC and so cannot be recorded in this way. Hence the diagram above also shows some DMCC recording. Avaya Contact Recorder will automatically fall back to other means of recording calls that it has not been invited to record via SIPREC.

Conferencing into Calls ("DMCC recording")

If SIPREC recording is not possible, the preferred approach to recording calls on this switch is to use Avaya's Device, Media and Call Control (DMCC) features to provide a wide range of recording modes with all the benefits of VoIP-based recording but without the limitations of passive tap IP recording systems. This approach to recording offers the following benefits:

- The recorder can record potentially any call on the switch. Traditional trunk and extension modes cannot record internal and tandem calls respectively.
- There is no cabling to maintain as new trunks or extensions are added to the switch.
- Uses standard PC servers with no proprietary cards.

The recorder uses two different methods to record calls. The table below shows which mode is used by each mode.

Recording Mode	**Uses**
On Demand Recording	Conference
Meeting Recording	
Bulk Recording	Single-step Conference
On Demand with External Controller	
Quality Monitoring	

Note:

Regardless of which recording method is used, when a recorder port joins a Communication Manager call to record it counts as an additional party on that call. Hence your normal limit of 6 parties on a call includes one party through which all of the recordings are made. This reduces the number of real parties on the call to five. Also note that it means that a call that has reached 6 parties without being recorded cannot then be recorded as the recorder's port will be unable to join until one of the other parties leaves the call.

Single-step Conference

Single-step conferencing as used by DMCC recording has the following characteristics and limitations:

Exclusion

Prior to Communication Manager 6.2, this recording mode cannot be used to record calls in which any user invokes the Exclusion feature.

Timeslots

Single-step conferences require an additional timeslot on the switch if (and only if) the recorder is configured to inject beep tone.

Device Names

To avoid excessive load on the system, the recorder caches the "name" of each device rather than request it on every call. These names are refreshed overnight so any changes are picked up the next day unless you restart the recorder.

Call Segmentation

"Bulk" recordings are broken into separate segments whenever the parties on the call change. The recorder continues to record only so long as the real parties on the call are connected. If the call is on hold, recording stops.

If an external controller is used, it may segment calls at points it chooses.

Bridged Lines

Because of an inherent limitation in the underlying call/connection model, calls involving more than one instance of the same (bridged) line may not be recorded correctly.

If the station whose number is being bridged is not actually registered or (in the case of a digital set) physically present, the recorder will start recording the call before it is answered. If using beep-tone injection, this results in the ANI being replaced by “Conferenced.”

Beep Tone

If you choose to inject beep tone on a single-step conferenced recording, the recorder becomes a full member of the call rather than a listen-only member and therefore:

- Becomes visible to the agent, who can see that the call is conferenced
- Uses an extra timeslot on the switch

Conference

In On Demand and Meeting recording modes, a user or an external application will dial a port on the recorder. When this happens, the recorder answers the call and is therefore a normal party on the call.

Timeslots

As "just another party" on the call, the recorder port will use the single additional timeslot that any other phone would use when added to a call.

Call Segmentation

As the recorder port is a normal party on the call, it is still connected even if one or more other parties on the call places the call on hold. It will receive the same audio that the parties remaining on the call receive. This may include music on hold or silence.

If an external controller is used, it may segment calls at points it chooses.

Attendant Consoles and Polycom Conference Bridges

Because TSAPI cannot observe these types of devices, ACR cannot record calls made on these devices.

CS1000

Avaya Contact Recorder can record calls made on the CS1000 switch via IP streaming (preferred) or via traditional TDM tapping - or a combination of the two. If using TDM, Avaya Contact Recorder must be installed on Windows.

Duplicate Media Streaming over IP

This mechanism overcomes the difficulties inherent in passive-tap IP approaches and allows greater flexibility in the location of recorders as well as increased capacity and reliability of recorders. It requires only a standard full-duplex Ethernet card in the recorder. It can be used on any type of switched Ethernet system, as it does not require port mirroring. Ensure that the bandwidth is adequate. Above 100 channels, gigabit Ethernet will be required.

This is the preferred method of recording calls on CS1000 and should be used where possible as it provides the simplest, most flexible and usually the most cost effective approach to recording. In this approach, the recorder instructs the telephone system to stream a duplicate copy of the audio directly to it over IP. n CS1000 systems, the audio is streamed by the IP phone itself; this approach is only available for systems using Phase II (or later) IP phones. The recorder sends instructions via a Meridian Link Services (MLS) interface as shown below.

Duplicate Media Streaming in CS 1000

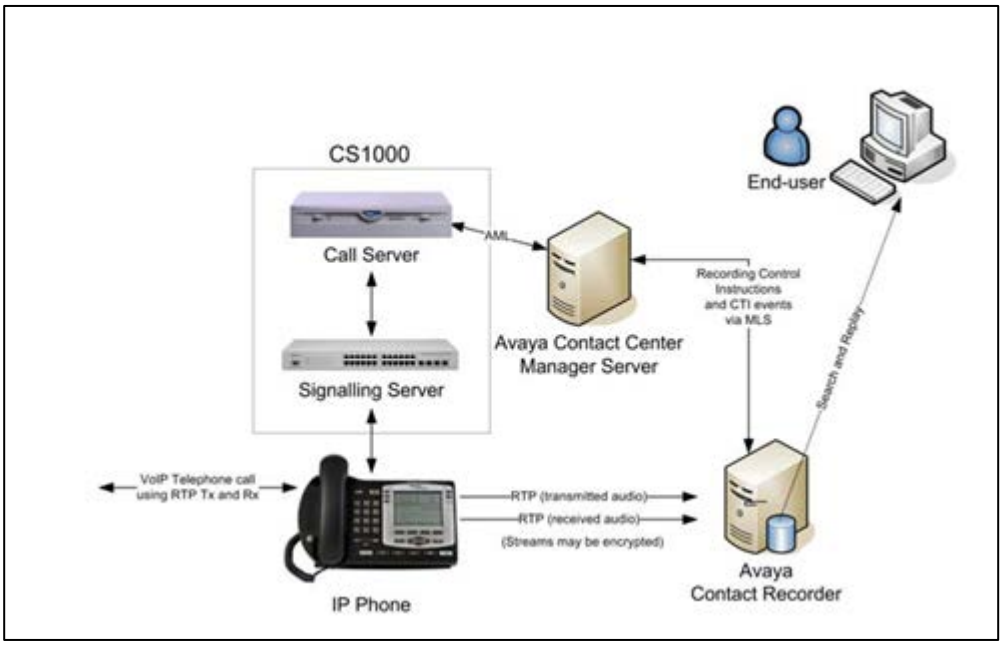

You should use this option for Bulk Recording or Quality Monitoring if both of these requirements are met:

- Your Avaya system meets the prerequisite requirements for Duplicate Media Streaming as shown in CS1000 System Prerequisites on page 81.
- The required bandwidth is available between the Avaya Contact Recorder(s) and the IP phonesets.

The recorder can be used with either G.711 (A- or µ-law), or compressed G.729A data streams - the latter with or without Voice Activity Detection (VAD).

Trunk-side tap

Where the bulk of telephone sets are not IP or the number of trunks carrying recordable calls is limited, tapping into the trunks may be appropriate. Using E1 or T1 passive tap cards, Avaya Contact Recorder can record any of the time slots as required. When using this method:

- Internal calls cannot be recorded. This makes it unsuitable for quality monitoring.
- An Avaya Contact Center Manager Server is required

Trunk-side Recording on CS1000

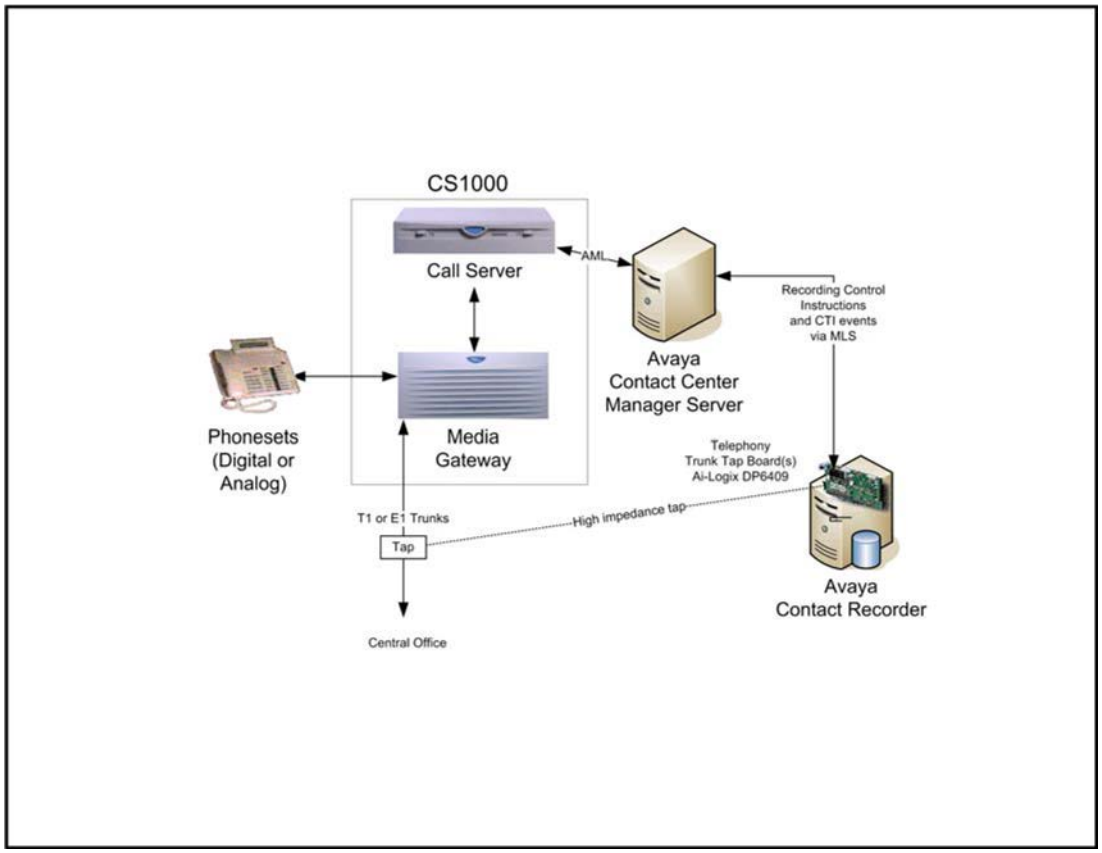

TDM Station-side tap

Where the telephone sets are digital and you only need to record a limited number of telephones, tapping into the extensions may be appropriate. Using digital extension tap cards allows an Avaya Contact Recorder to access the audio passing to or from the tapped telephones as shown below.

Digital Extension Recording on CS1000

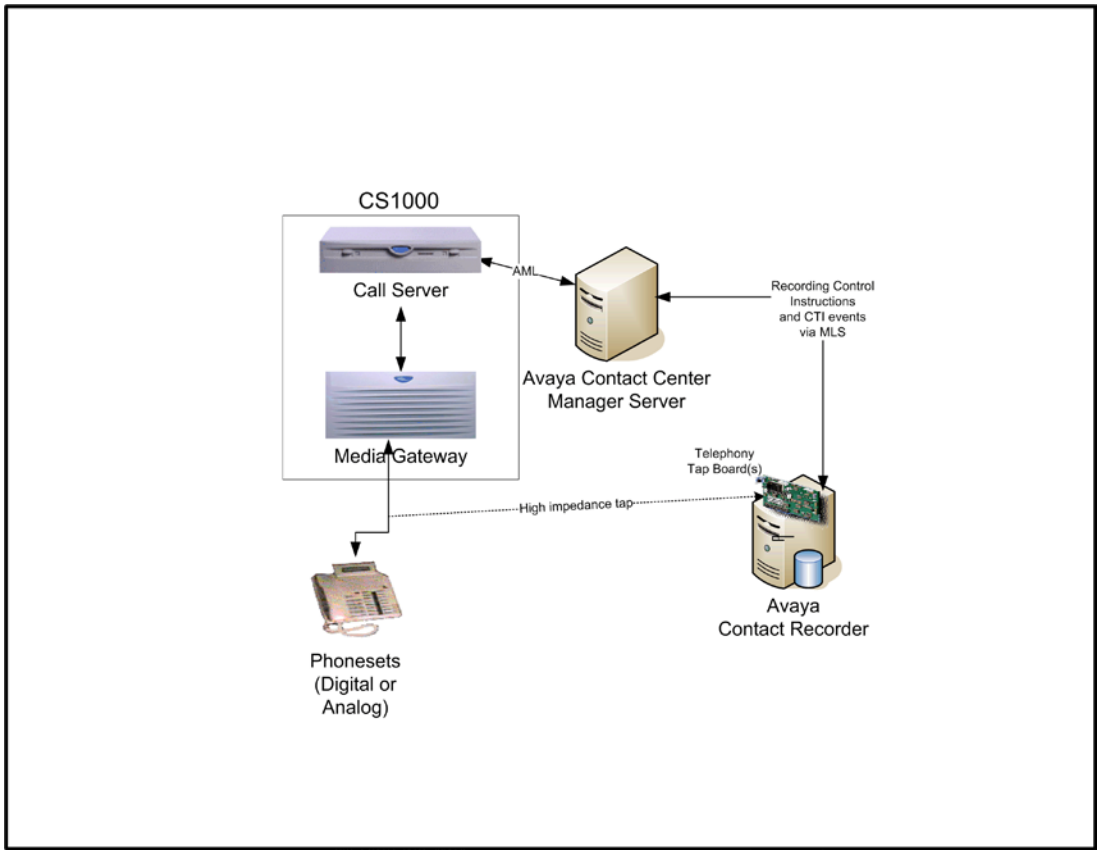

CAUTION:

CS1000 switches support many different types of extension. Check the Ai-Logix cards' manuals to confirm precise models supported.

Limitations

Actions by Unobserved Positions/DNs

CTI events are only received for the positions/DNs being recorded. When another, unrecorded and hence unobserved position/DN is involved in a call, not all actions taken by that other party are visible to the recorded DN. If such an internal "far-end" places the call on hold, transfers to another party or conferences in another party, the recorder will not be aware and hence cannot tag the call with the details of the third parties to whom the call is transferred or who are conferenced into the call. This can also result in the

recorder not being told when a call ends. In this case, the contact may remain open (and hence not copied to WFO) until timed out (3-4 hours later) – at which point it will be assumed to have ended and the details flushed through to WFO.

Emergency Calls

Calls received by a phone as a result of someone using the Emergency key may be recorded more than once.

Supervisor Calls

Calls received by a phone as a result of someone using the Supervisor call feature may be recorded more than once.

Beep tone

Beep tone is only supported in Duplicate Media Streaming recording modes. It is not supported in any of the TDM recording modes.

Agent Status

In CC V6, agent status is not visible to the recorder at startup. Agents must log out and in again after the recorder has started before calls can be recorded by and tagged with agent numbers.

MARP/MADN

Multiple Appearance DNs ("MARP" and "MADN") can only be recorded when in Bulk Recording mode in a knowledge worker environment. These are not supported by the Quality Monitoring application, or the Avaya Contact Center environment.

Recording of calls made from one instance of a number to another line key on the same number are NOT supported. The CSTA computer telephony model which underpins the call state tracking cannot handle the same address being involved in two separate connections on the same call.

Avaya Aura® Contact Center

Avaya Contact Recorder can tag calls controlled by an Avaya Aura® Contact Center. When this is hosted on a Communication Manager, ACR uses the CTI information provided via the CCT feed to tag recordings more fully than can be done via TSAPI alone. All calls are recorded using one of the methods described above for the appropriate underlying switch platform. In such cases, you must plan and configure your system to support both the underlying recording mechanism appropriate to your switch platform and with the CTI links shown below.

ACR recording Avaya Aura® Contact Center (based on Communication Manager)

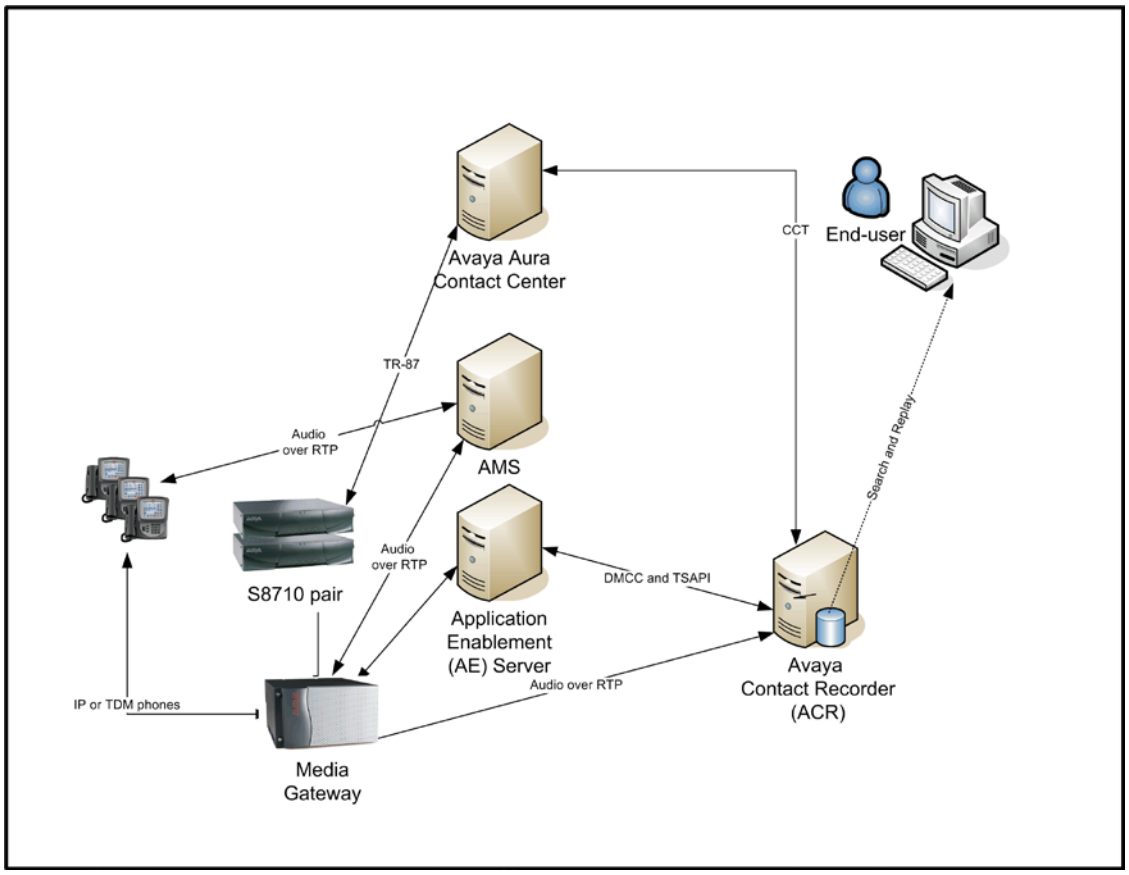

Limitations

ACR merges CTI information from both the AACC and the underlying Communication Manager, waiting a short while (typically 250ms) to allow both systems to settle before deciding on the next recording action. Where CTI events are separated by more than this, short recording segments may be created during these transient conditions.

Server components

The Avaya Contact Recorder system can be installed as a single server solution providing recording and replay of calls. Large systems can be built in a Master/Slave or Master/Standby/Slave topology.

You can extend the scope of the system by adding additional optional server applications to create a comprehensive Workforce Optimization system.

The Avaya Contact Recorder runs on Windows or Linux but certain recording modes are only supported on one or other, as shown in the table below.

Environment	Recording Mode	Windows	Linux
Communication Manager	DMCC	Yes	Yes
	SIPREC (of calls via SBC)	Yes	Yes
	TDM tap	Yes	No
	IP Passive tap	No	Yes
CS1000	Duplicate Media Streaming	Yes	Yes
	TDM	Yes	No

All of the other applications in the suite require Windows. The optional server applications normally require their own physical server.

Each of the other components has a corresponding guide (as detailed in Additional references on page 20) which you should refer to for more detail.

Avaya Contact Recorder Server

In most small to medium-sized (up to several hundred channels) Bulk Recording systems, one of these applications provides the entire recording and replay system (as shown in Introduction on page 25).

In larger systems (where a single physical server is not powerful enough), you should install multiple instances of this application on different physical servers, each providing a subset of the system's overall functionality. The Avaya Contact Recorder can:

- Connect to your Avaya switches CTI feed(s) and control all voice recordings.
- Record and store telephone calls.

- Record and store screen content of Windows desktops during phone calls.
- Archive the recordings it makes to one or more file stores.
- Provide search and replay services to users connecting via their browser or via their telephone.
- Provide voice recording services to Workforce Optimization.
- Control other Avaya Contact Recorders.
- Be controlled by another Avaya Contact Recorder.
- Act as a **Full** or **Partial Standby** to an Avaya Contact Recorder Master.
- Act as a centralized replay server (or, be a dedicated Central Replay Server), holding details of recordings made by other Avaya Contact Recorders.

Optional Server Applications

Workforce Optimization ("WFO")

This application suite provides comprehensive Workforce Optimization tools that can control and exploit the voice and screen recordings made by Avaya Contact Recorder. These include Quality Monitoring, coaching, and Workforce Management tools among others.

See Avaya Contact Recorder Integration to Workforce Optimization Guide for how to integrate it into the overall system.

Centralized Replay Server(s)

In any system with more than one Avaya Contact Recorder, details of recordings are (by default) uploaded to the Master (and **Full Standby** if present). This allows users to search for and replay calls recorded on any of the servers in a single query. On large (>5000 channels or more than 10 concurrent replays) systems, you should dedicate a server to this search and replay task.

This server should be another instance of the Avaya Contact Recorder, licensed and configured as a “Central Replay Server" (CRS).

Note that the recording system is designed to continue recording regardless of the state of these central servers. This means that the availability of the central applications cannot affect the reliability of the recorders themselves. However, should you wish to deploy independent and hence fault tolerant central search and replay servers, this is also supported.

Centralized Search and Replay

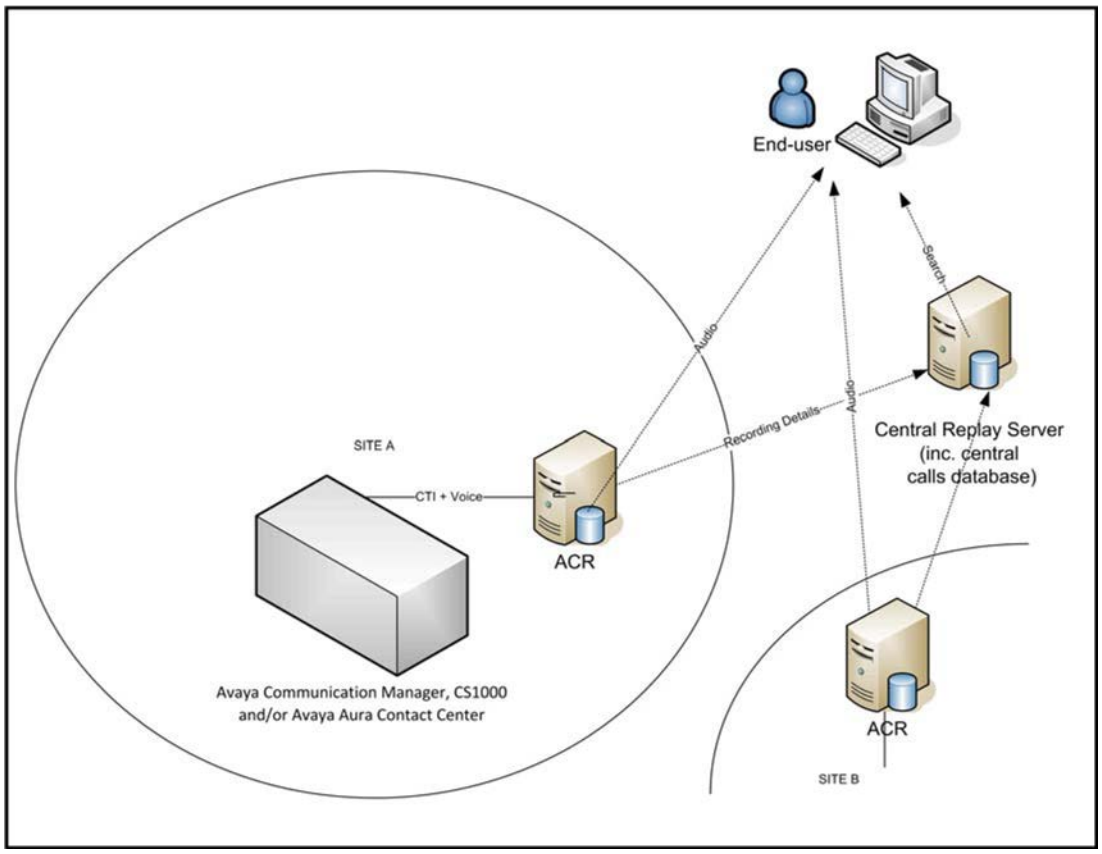

External Control

In addition to the data provided by the Avaya switches main CTI link(s), you may wish to control and/or tag your recordings with details from other CTI feeds or application interfaces. These may include third party systems and/or your own in-house applications. The recorder supports a wide range of systems and allows them to be connected to your recording system. These include:

- Avaya Proactive Contact and other auto-dialers as described in Proactive Contact (PCS/PDS Dialer) on page 149 onwards.
- WebXAPI commands as detailed in Appendix G: External APIs on page 416 onwards.

End-User tools

To access the recordings held in the system, users have a variety of options.

Workforce Optimization ("WFO")

If you have integrated Avaya Contact Recorder with a WFO system, the user interfaces of that system are available to search for, replay, monitor in real-time and evaluate the recordings.

Search and Replay

Integral Search, Replay and Live Monitoring

The Avaya Contact Recorder includes a search, replay and (optionally) live monitoring application within it. This replay mechanism is a very simple and intuitive browser-based interface, requiring the user to access it via Internet Explorer Version 11. For further details, see *Avaya Contact Recorder User Guide*.

Other browsers that support the Web Audio API (such as Edge[1] and Chrome) may be used for Search and Replay via the user's PC. However, Live Monitoring, Telephone Replay and Export of recordings from the Replay screen require Internet Explorer V11.

The Search and Replay application is hosted on a web server running on the recorder itself. It uses a local database of recordings to allow users to search for recordings at two different levels - either for individual recording files ("call segment" level) or at an overall "session" level. This latter level groups recordings together so that consultation calls, transfers and conferences are shown as integral parts of an overall "contact" within which each person involved in the call experienced their own "session". Users typically search for calls according to:

- Call start date/time
- The name(s) and number(s) where provided of any party on the call or through which the call was routed. This includes stations, ANI, DID, Skill or hunt group, VDN etc. where provided by the switch.
- Agent ID and name
- Call duration
- Call Identifier or Contact Identifier
- User defined fields supplied by external controllers

[1] Edge has such low performance with this API that it may be unusably slow.

A full description of the attributes that recordings are tagged with is given in Recording Attributes on page 312. Access restrictions determine which recordings individual users are able to replay. Each recording is assigned one or more "owners" at recording time (see Access Rights on page 199 for further details).

The user can search for and play any recording that matches their access rights and can search for and view the details of any contact containing at least one recording that they are allowed to play.

When a call is played, a graphical representation of the audio level of the call, the audio wave form, is displayed. The audio wave form shows silence and tones, so the user can click beyond irrelevant sections and pinpoint parts of the call that are of interest. Screen content, where recorded, can also be viewed. See the accompanying User Guide for further details of this application.

If licensed and permitted by the administrator, the replay screen also includes a **Live Monitor** link which, when using Internet Explorer V11, opens a separate window. The window allows a user to monitor a station or agent - hearing the call they are on and, if configured, seeing their workstation's screen.

Central Search and Replay

Where multiple Avaya Contact Recorders are deployed, an additional server can be nominated as a Central Replay Server. This server is not used for recording but can provide telephony replay ports (on Communication Manager only).

The other recorders upload details of the recordings they have made into this server's database allowing users to search for and replay recordings made on any recorder without having to know which one recorded a particular call.

Agent Initiated Monitoring

This is a very simple desktop application that lets users control recording of a line that they are using alongside the PC on which they are running this application. The application typically shows as an icon in the tool tray. Users can then use the application's popup menu to:

- start or stop recording
- request that a recording is deleted
- pause and resume recording around sensitive security questions
- tag a recording with additional details

Note:

Users with CS1000 IP phones can perform the first two directly from keys on their phone.

Avaya Desktop Applications Deployment Reference and Installation Guide describes how to configure Avaya Contact Recorder for use with this application.

Administration Tools

As the suite is designed specifically for Avaya systems, much of the complexity associated with generic recording systems has been removed resulting in a system that is easy to configure and maintain. The recorders are administered via a web interface. The detailed use of this interface is the subject of later Sections in this guide.

Alternatively, on Communication Manager systems, much of the day-to-day administration can be performed on Avaya Contact Center Control Manager and pushed to the Avaya Contact Recorder.

Recording Functionality

The first task in designing any recording system is to define what is to be recorded. This in turn is often driven by the reason for wanting the recordings. This section introduces the various ways in which recorders' ports can be used. Refer to Operations on page 163 for detailed functionality and limitations of each mode.

Sampled Recording for Quality Assessment

If the end goal is to record only a sample of calls in order to assess the quality of your interactions with customers, Workforce Optimization ("WFO") can be used alongside the Avaya Contact Recorder. See *Avaya Contact Recorder Integration to Workforce Optimization Guide* to determine how you will use the WFO system to control recording. Note that this document also highlights a number of limitations.

Licensing

The Avaya Contact Recorder must be provided with a license key that specifies how many different telephone stations can be recorded by WFO.

Bulk Recording

If you need to record all of the calls taken by specific stations, agents, skill groups, or Vector Directory Numbers (VDNs) then you require Bulk Recording.

CS1000

You can record specific Position IDs and DNs.

Communication Manager

CTI Information received over an AES TSAPI link allows the Avaya Contact Recorder to record calls on specific stations, agents, skill groups, VDNs or CoRs. Advanced configuration options let you filter the calls by VDN or Skill Group rather than have to record every call on the nominated addresses.

Avaya Aura® Contact Center (AACC)

Calls are recorded as described above for the relevant underlying platform. In the case of Communication Manager, an additional CTI feed from AACC provides further tagging details.

Externally Controlled Recording

More complex recording requirements may be met by customized or specialist applications that interface to other CTI feeds or customers' own applications. Such an application can control ports on the Avaya Contact Recorder, allowing it to record exactly what and when it requires. The application may also "tag" the recordings with additional details such as customer number or account number.

Screen Recording

In addition to recording the audio content of a call, you can also associate a telephone with a Windows PC and record the content of its screen in Bulk Recording mode. Alternatively, the recorder can track where an agent logs in and record the screen of that agent's Windows desktop. In this case, as well as a dedicated Windows PC the recorder can record the user's desktop on several thin client topologies as detailed in the *Avaya Desktop Applications Deployment Reference and Installation Guide*.

Note that the screens of more than one party on a call may be recorded.

Licensing

The recording capacity across the whole system is restricted according to the license key entered. (This encompasses all voice recordings being made except the number of screen recordings which are separately licensed.)

Ad-hoc or Occasional Recording Modes (Communication Manager only)

Although Bulk Recording can be configured to delete most calls and only retain those a user selects during a call, this still requires a recording port to be assigned throughout the duration of a call. Two other recording modes are provided for those requiring occasional recording:

On Demand Recording

This mode lets users of any phone on your system dial into or conference in a recording port as and when they want to start recording a call.

One or more "pools" of ports on the recorder can be assigned to this recording mode and accessed via Hunt Group numbers so that callers automatically reach an available port. The recorder automatically answers the incoming call on its port and starts recording.

Refer to On Demand Recording (Communication Manager only) on page 176 for a full description of this mode and how to configure it.

Meeting Recording

A novel use for recording is in taking a detailed log of a meeting, either as an audio record for those attending, or as a way to include non-attendees later. You can use any meeting room or office with a telephone that has a speakerphone or, ideally, conference phone capabilities to record the meeting.

Important:

The audio recorded with Meeting Recording is the same as someone dialing in would hear it on the phone used to record it. Place the phone so it picks up the speech of all participants. They should speak loudly and clearly. Experiment with this recording mode before relying on it to provide full and complete records of your meetings. Avaya cannot be held responsible for the failure to pick up all of the audio intelligibly. Use this recording mode as an aid to note taking, not a replacement for it.

One or more "pools" of ports on the recorder can be assigned to this recording mode and accessed via Hunt Group numbers so that callers automatically reach an available port. The user follows the spoken instructions to start the recording and specify which user(s) can access it.

Refer to Meeting Recording (Communication Manager only) on page 178 for a full description of this mode and how to configure it.

Port Requirements

Ports assigned to On Demand or Meeting recording

- Can be assigned to one or more hunt groups making them easily shared across one or more user populations.
- Can be used not only by stations on the switch but also from outside the switch if they are made accessible via a DID number.

Licensing

- Each recording is counted as part of the overall recording capacity and is restricted according to the license key entered.

Replay Options

Soundcard Replay

Many users choose to replay recordings via their browser and the soundcard on their PC. This does not use any ports on the recorder and does not require any additional licensing.

Telephone Replay (Communication Manager and IE11 only)

However, the recorder also supports replay via the user's telephone on Communication Manager when using Internet Explorer V11 and this does use a port on the recorder.

Refer to (Telephone) Replay Ports (Communication Manager only) on page 180 for a full description of this mode and how to configure it.

Licensing

To use telephone replay ports, you will need to purchase the required number of replay channel licenses.

Miscellaneous

Beep Tone

Some states and countries require that both parties on a call be made aware that the call is being recorded. One way to do this is to apply a tone to the line. Note, however, that most telephone users are unaware of the significance of this tone and might, in fact,

regard it as a fault. There are more effective means of informing the user that the call is being recorded. For example, you can inform users on advertising and in contract literature or play a recorded announcement before the call is connected. Check the legal position in your jurisdiction and apply the appropriate settings.

See the discussions of how each recording mode works earlier in this chapter to determine whether or not you need to turn on this warning tone within the recorder. In most cases, it is recommended that you set this option to No and use the other mechanisms described there.

As an alternative to having the recorder inject beep-tone using the above setting, as of Version 15.1 you can enable beep-tone injection for Communication Manager when a recorder's DMCC port joins a call. This does not require an additional timeslot. However, it is also always on – any time a recorder's DMCC port is conferenced into a call – whether or not the audio is actually being recorded. Using this method (and only this method), the beep-tone is also recorded. To enable this mode, set the property **`dmcc.addbusyverify=true`** and restart the recorder.

In the few cases where the recorder is able to inject or control beep tone, you can set this to be on, off, or on only when the recording will be retained.

International support

The recorder's browser-based Administration and Search/Replay interfaces are provided in several languages. Check with Avaya for availability of specific languages. International support includes:

- Time zone and DST support. All dates and times are stored in the database in Coordinated Universal Time (UTC). However, when you view records using the search and replay application, these are converted to local time of the recorder or Central Replay Server that you are accessing. If you view the records using a database query tool, the times will be shown in the time zone of the client machine, which may be different from the server time. Note that the XML files relating to the recordings include ISO standard timestamps, giving both UTC and offset from Greenwich Mean Time (GMT).
- Number ranges are stored in the system database in left-to-right format (e.g. 100-200) unless you configure the property **`system.forcertl=true`** in the properties file. This affects alarms and audit trail entries.

Liability

Liability of Avaya for failure to record any calls is limited under the terms of supply.

Chapter 2: Planning and Prerequisites

This chapter gives details of the prerequisites for an Avaya Contact Recorder system. You should also review Chapter 6 System Security as some of the optional elements described there may also require additional cost and/or effort.

The main sections in this chapter are:

- Introduction on page 53
- Recording Bandwidth on page 54
- Storage Requirements on page 57
- TDM Interfaces on page 60
- Server Platform on page 61
- Network Issues on page 66
- Licensing on page 67
- Communication Manager System Prerequisites on page 69
- CS1000 System Prerequisites on page 77
- Avaya Oceana™ System Prerequisites on page 78
- AACC System Prerequisites on page 78
- Topologies on page 79

Introduction

Unfortunately, there is no one "right" order in which to plan a system. A number of the requirements and pre-requisites are inter-dependent and it may be necessary to iterate through this section several times, refining your plans each time. This Chapter assumes that you have read the previous one and know the type and quantity of recording and replay that you require - and hence your total license requirements.

This Chapter will now guide you through:

- identifying what is to be recorded, the options for doing so and the relevant license requirements for the preferred option
- determining the storage needed to hold your recordings
- sizing the servers needed for each application
- quantifying the network bandwidth needed

You can then determine an appropriate system topology.

Recording Bandwidth

Each type of recording requires a certain amount of data to be passed between components of the systems, processed and ultimately stored on disk and/or archive drives. This section explains the implications of each type of recording and - where choices can be made in terms of configuration or topology - the implications of these.

Voice Recording

The Avaya Contact Recorder supports two types of audio codec: G.711 (64kbps) and G.729 (8kbps). The choice of format depends on your switch type and network topology.

G.711

If the recorder receives audio in G.711, recordings will be compressed by the recorder and stored as G.729 (8kbps) (unless compression is deliberately disabled by setting `acr.disablecompress=true` in the properties file.)

The implications of receiving G.711 (and compressing it) rather than receiving G.729 in the first place are:

- Several times more bandwidth will be needed between the recorder and the source of the audio.
- Bulk recording capacity of a given server will be reduced by about 30% if the server does not support SSSE3.

G.729A

If the recorder receives G.729, the recorder will not have to perform any compression tasks. The pros and cons are therefore roughly reversed from those shown above.

Communication Manager

With TDM and passive IP tapping, the recorder does not impose any load on the switching fabric of the Communication Manager. When using DMCC softphones, however, calls are recorded using resources on the Avaya Communication Manager to conference a recorder port into a live telephone call. One of the benefits of this approach is that the recorder can ask to receive audio in a format that suits it - without impacting the experience of the parties actually speaking on the call.

The recorder requests either G.711 (µ-law), or compressed G.729A data streams. The above approach also means that audio is received at the recorder already mixed - so a G.729 recording requires much less processing power than a G.711 call (which is compressed to G.729A by the recorder).

G.729 is definitely preferable from the recorder's point of view but uses about twice as much audio processing resource within the Communication Manager. You must choose which is appropriate for your recording needs, network and server sizing.

Note:

When using G.711, the recorder always requests µ-law, never A-law. This applies to all countries regardless of that country's national preference.

If using IP passive tap recording, the Avaya Contact Recorder receives packets in whatever format you have configured your Communication Manager to use. This mode results in "stereo" recording of the call[1].

SIPREC recordings are also received in stereo.

TDM recordings are compressed by the line interface cards which should be configured to provide G.729A (8kbps).

CS1000

On a CS1000 system with IP call recording, IP phones stream copies of the audio packets they are sending and receiving to the Avaya Contact Recorder. As packets are simply copied to the recorder, they are always in the same format (G.711 or G.729) that is being used on the call itself. If you wish to take advantage of the benefits of G.729 recording, you must configure your CS1000 system to use this format.

Duplicate Media Streaming results in "stereo" recording of the call. Two separate streams of audio are received. In the case of G.711 these are mixed and compressed to a G.729A (8kbps) file. In the case of G.729, they are kept as two separate streams and so occupy 16kbps. Hence the storage requirement is doubled but the recorder will support a much higher channel count if it receives G.729.

Note that a phoneset using bandwidth of X kbps to make a phone call will require an additional 2X kbps (twice X) for the streams being copied to the recorder. This results in a new total bandwidth of 3X kbps (three times the bandwidth of the unrecorded call) from the phone being recorded. The quality of service (QoS) of the network between the IP phones being recorded and the recorder must be the same QoS as required for a voice call.

If you have limited wide area bandwidth, consider placing slave recorders on the same sites close to populations of phones to be recorded. You can force recordings to be made on specific recorders.

TDM recordings are compressed by the line interface cards which should be configured to provide G.729A (8kbps).

---

[1] Unless you set **Keep Audio in stereo wherever possible** on the **General Setup > Server** page to **No** – in which case stereo G.711 streams will be mixed before being compressed as a single (mono) recording.

Screen Recording

Screen content is also transmitted via IP. In this case, when a screen is associated with a phone that is being bulk recorded, or a user is seen to log on with a domain account for which the recorder has an Agent ID set the screen content is sent from the workstation being recorded to the recorder. Note that, in the case of thin client workstations, the recorder interacts with the terminal server providing the services, not directly with the user's desktop display device.

Typical storage requirements for screen recordings are shown below but these can vary enormously according to how dynamic the users' screens are. These figures are with "Color Reduction" applied - which is mandatory for ACR recordings.

Screen Resolution	Typical Storage
1024 x 768	200 KB / min
1280 x 1024	239 KB / min
1600 x 1200	301 KB / min
1920 x 1280	335 KB / min

Screen recordings must be taken into account when sizing storage requirements as well as network bandwidth. Note that you can configure the recording system such that screens are recorded on designated recorders - helping you to control bandwidth demands across your WAN. See Designated screen recorders on page 132 for more details.

To avoid screen recording files becoming too large and unmanageable, screen recordings are automatically terminated after a maximum duration. You can adjust this threshold using the property value **`screen.maxdurationmins`**. The default value is 720 minutes (12 hours).

Storage Requirements

Having determined the type, size and location of your recording capacity, you must now determine how much storage is required within the system. Storage requirements of many terabytes are not uncommon.

Storage is required for the:

- operating system and applications installed on each server
- audio and screen content of recordings
- database holding the searchable details of these recordings

In each case you should consider requirements locally at each recorder and centrally.

To do this, you will need to know:

- How many recordings will be made on the recording channels already identified (typically expressed in channel hours per day). In other words, how many hours of audio and/or screen you expect to record on each recorder every day.
- How long you wish to retain recordings (typically expressed in days).
- The average duration of a call (typically expressed in seconds).

Storage at Each Recorder

In addition to the operating system, installed software and its configuration data, each recorder stores:

- details of the recordings it has made in a local database
- the recordings it has made as files on its hard disk

In all cases, the system automatically notes the location of archived copies it makes of recordings so that when a user wishes to replay a call that is no longer in local online (disk) storage, it knows where the recording can be found.

Voice recordings are stored in G.729A format - in either mono (8kbps) or stereo (16kbps) according to the recording method and format in which the audio was received as described above. This allows for high volumes of recordings to be stored on the available disk space. Where recordings are stored in stereo, the replay screen shows the appropriate (left or right) channel waveform for each party that is present in the call. On a simple external call with one internal and one external party, the internal party is the left channel.

Avaya strongly recommends RAID arrays or fault tolerant Storage Attached Network (SAN) devices for online storage of recordings on the recorder platform. The operating system and calls database should be RAID 1 and the (typically much larger) recording storage area should be RAID 5.

The recorder automatically manages the available recording storage space. This is used as a circular buffer providing instant access to the most recent recordings and deleting

the oldest, as space is required for new recordings or as they reach the retention period configured.

Rather than storing bulk recordings on RAID arrays in each recorder, many customers prefer to use Storage Attached Networks. These must be connected directly to recorders.

Hierarchical File Storage (HFS) Systems and Network Attached Storage (NAS), however, can only be supported via the recorder's Archive features. These act as secondary recording storage.

Important:

The storage used for recordings **must** be local to the server making the recordings. It must **not** be accessed via a network connection.

The table below summarizes the requirements for Bulk Recording on an Avaya Contact Recorder.

	Details of Recordings	**Content of Recordings**
Stored	In local postgreSQL database	Stored as files on local disk in a hierarchical folder structure
Purged	Nightly, after user defined period (default 60 months)	100 at a time as disk space is needed for new recordings
Volume	Approximately 2KB per recording. So total ~ 2KB x recordings/day x days retained. Double this if “turbo” mode is to be used on a Central Replay Server.	Approximately 7.2MB (G.729 stereo) or 3.6MB (G.729 mono) per channel hour of audio. So total ~ 7.2MB (or 3.6MB) x channel hrs/day x days retained
Location (Linux)	/var (in addition to the normal contents of /var)	/calls
Location (Windows)	Install path	User Configurable. Dedicated partition required.
Type	RAID 1 (mirrored) or 5 (striped) strongly recommended. Local hard drive or SAN, not NAS.	RAID 1 (mirrored) or 5 (striped) strongly recommended. Local hard drive or SAN, not NAS.

Workforce Optimization ("WFO")

If Avaya Contact Recorder is integrated with a WFO system, then it is likely to require more storage space for both screen and audio content.

By default, all recordings made on Avaya Contact Recorder are made available to the WFO database. You can override this by setting the property entry

`core.consolidateall=false`. Thereafter, only details of recordings involving parties configured appropriately within WFO will be sent to the WFO database. Further details are given in Sessions Visible to WFO on page 330. You must first determine what proportion of recordings will be made visible to WFO. For those calls, you will require additional storage.

WFO stores details of recordings by "session" - that is, according to the internal party involved in the call. Each internal party needs a separate session - and hence separate recordings. Normally, Avaya Contact Recorder makes a single recording of each call segment regardless of how many parties are on the call. However, where WFO must be advised of multiple sessions, each one requires a separate copy of the audio. Additional storage capacity for these recordings must be provided. Where a high proportion of internal calls are to be recorded this can be significant.

Central Database Storage

Follow the guidance above to size the storage needed for the call details using the total call volumes across all recorders that are feeding into the Central Replay Server.

It is advisable to provide at least double the calculated storage space in case “turbo” mode is required later. This mode trades disk space for search speed. See “Turbo” Mode on page 285 for further details.

Archive Call Storage

Where you wish to retain calls for longer than it takes to fill the hard disk storage on a recorder, provide one or more file-shares, SFTP servers, mount points and/or EMC Centera file stores onto which the recorders can archive their recordings.

Backup Storage

See Backup/Restore on page 216 for a discussion of Backup and Restore options. You should determine whether additional storage space is required in your corporate backup system to accommodate the new recording system.

TDM Interfaces

Cards Supported

The following Ai-Logix cards can be used:

- DP Series (trunk taps)
- NGX Series (digital extension taps)
- LDA Series (analog taps)

Chassis Requirements

As far as performance is concerned, up to 510 channels can be supported in a single chassis but the capacity is normally limited by the number of card slots available. An extension chassis may be used if the motherboard does not provide sufficient slots.

Platform Restrictions

As the software tools provided for these cards are 32 bit Windows applications, any Avaya Contact Recorder using TDM cards must run on Windows (not Linux) regardless of which telephony switch is being recorded.

A separate (32 bit) process runs as part of the Avaya Contact Recorder and passes audio to the main (64 bit) service.

Server Platform

Taking the above factors and the potential location(s) of your recorders into consideration, you must determine how many channels of each type of recording you wish to deploy and on which site.

Having decided the total recording capacity at each of your locations, you must translate this into one or more server platforms capable of handling the load identified.

Apart from the specific exceptions listed in Virtualization on page 65, dedicated server(s) must be provided with no other applications running on them.

Sizing

The benchmarks given below are for the following server specifications. Type A is a typical mid-range server while Type B is the minimum required for new installations. Existing servers can be assumed to continue supporting the load for which they were sized on the version originally installed.

Type A – Mid-range Server

- Dual 2.67GHz six-core CPU with hyper-threading and Intel Supplemental SSE3 (SSSE3) support
- 1Gbps Ethernet NIC port
- 16GB RAM
- RAID 1 or 5 strongly recommended
- DVD drive for installation of software
- Red Hat Enterprise Linux 7 (64 bit) or Windows 2012 R2 (Standard or Data Center Edition)

Type B – Entry-level Server

As above except for CPU and memory:

- Single Intel Xeon E5620, 2.4 GHz, quad-core CPU with Hyper threading
- 8 GB RAM

Important:

The table below shows the officially supported load on a server of at least the specification given above. While a more powerful server will probably have a higher sustainable capacity, you cannot simply extrapolate the figures below. For example, a server with twice as many cores and twice as much memory will still be transferring all data via a single Gigabit NIC port - and that may become the bottleneck. The table below therefore gives the officially supported load per physical server.

Recording Method (all supported codecs)	**Max concurrent channels** (Notes 1, 2)		
	Type A (Windows)	**Type A (Linux)**	**Type B (either O/S)**
Communication Manager using DMCC, SIPREC or Passive IP(Note 4)	1000	1200	650
CS1000 Duplicate Media Streaming			
Either switch via TDM	510 (hardware limited)		

BUT WHEN USING	Reduce max concurrent channels by
24-hour loading rather than 8 hour day	30%
Screen Capture	If your screen recordings occupy significantly more than the typical figures shown in Screen Recording on page 56, you should reduce the maximum load on each server.
More than 10 concurrent replays per recorder	N/A. Install dedicated Central Replay Server instead. (Note 3)
Average call duration less than 1 minute	25%
Virtualized environment (under Linux or Windows)	10% but see Virtualization on page 65
G.711 input Server that does not support Intel Supplemental SSE3 (SSSE3). Most AMD servers do not support this.	30%

Notes:

1. A single “channel” includes voice *and* any associated screen recording.
   Check the limits imposed by your switch infrastructure (later in this chapter) as these may limit the capacity of individual servers and/or the overall system to lower figures than the recorder hardware does.
2. Where more than one factor from the second table applies, the effects are cumulative. For example, the acceptable channel count for AACC is 1200 but if calls are under a minute long on average and a virtualized environment is to be used, the capacity is 1200 X 75% x 90% = 810 channels.
3. A dedicated Central Replay Server can support systems of up to 5000 channels and (Communication Manager only) 40 phone replay ports.
4. Passive IP figures assume no more than 480,000 packets per second on a Type A server or 240,000 on a Type B server AND no more than 120,000 packets per second (whether recorded or not) on any one NIC.

Antivirus and other 3rd Party Applications

The overarching requirement for an Avaya Contact Recorder - which is a high performance, real-time system that makes heavy demands on CPU, disk and network I/O - is that it is not materially impacted by other applications running on the same server. Sizing of servers and hence the correct operation of the systems assume that substantially all (i.e. 90% or more) of specified resources are available to the Avaya Contact Recorder at all times.

This applies to all other applications running on the same ***virtual*** and ***physical*** machine i.e. includes but is not limited to:

- Anti-virus software.

  The recorder's files must be excluded from real time and scheduled scans

  The recorder and underlying postgreSQL database processes must be completely excluded from real time and scheduled scans. Note that some tools still access directories even when these have been excluded and others do not have real-time exclusion features at all.
- Backup tools (including Virtual machine snapshotting and recovery mechanisms) - see Backup/Restore on page 216 for details.
- System management tools
- Virtualization frameworks

Common Problems

Common causes of material impact that MUST be avoided include:

1. Anti-virus software interacting ***at all*** with the recorder's database or recording folders (as is known to happen with e.g. Trend Micro even when these folders are excluded and real-time scanning is turned off). At present, only AVG version 3485 Business Edition has been confirmed to avoid all contact with the recorder's database files (when the recorder's install path and call recording buffer are correctly excluded from its scope for both real-time and scheduled scans).

   All virus scanners impact disk performance significantly when running scans. The performance of the recorder during such scanning is not guaranteed. If these are run out of hours, it is the customer's responsibility to ensure that the recorder has not been affected by them and any impact found to be caused by them will be chargeable. Current recommendations with respect to the underlying database used (PostgreSQL) can be found at:-

   https://wiki.postgresql.org/wiki/Running_%26_Installing_PostgreSQL_On_Native_Windows#Antivirus_software
2. Any heavy and, especially, sustained disk activity that reduces the recorder's ability to create, read, write or rename files - such as scheduled virus scans, backups, Windows disk indexing and VM snapshots. Backup/Restore on page 216 explains why these are not required. The recorder's performance is not guaranteed while

these are running - and note that problems arising from them even when carried out in idle periods may carry forward into business hours.

3. Other processes using up available physical memory and hence triggering page file swapping must be avoided at all costs. This degrades apparent memory performance by at least an order of magnitude and also impacts disk I/O dramatically.

The above not only reduce the sustainable capacity of a recorder, they also push many of the tight real time constraints (e.g. 20ms packet intervals in most VoIP environments) out of acceptable bounds and increase latency such that events are not processed within the sub-second intervals needed to avoid losing significant audio content at the start of a recording.

Virtualization

ACR *can* run on Virtual Servers and where your total recording load is significantly less than the throughput provided by a single server you can use a VM to share *that* server with other non-ACR applications - subject to the above constraints.

However, as soon as your load exceeds 70% of one server's capacity there is a significant *downside* in deploying on VM guests; namely:

1. The ACR topology inherently supports failover. Virtualization adds nothing to the overall resilience of the solution as a whole. It merely complicates it.
2. The resultant solution loses (typically) 10% of the available capacity due to virtualization.
3. There are many ways to configure VMs incorrectly and various combinations of hardware and software make it impossible to configure correctly. Specifically, the real time clock must be within one second of real time at all times and on multi-core or multi-CPU machines must be consistent to within one second across all cores at all times. This problem has been seen repeatedly on many Virtual Machines where the underlying hardware, host operating system and guest operating system are either incorrectly configured or where heavy CPU load causes missed ticks. Any blockages or time sync issues not only impact recording itself but make it virtually impossible for a *set of ACRs* to operate reliably together as none of them can trust that the others are immediately available or have the same view of any timeout or heartbeat.

As a general rule, large systems requiring more than a whole server do not benefit from virtualization and are generally much more reliable and simpler to maintain when running directly on physical hardware.

It is never possible to support more channel capacity on a single physical server by deploying multiple ACR slices on it. The actual resultant capacity will always be less than that achieved by installing a single ACR process directly on the physical server.

Even though a physical server may have ten times as many cores as the minimum specification machine for which sizing is given, it cannot be assumed that the capacity of such a server is ten times that quoted. The bottleneck determining its capacity will be something else e.g. NIC card throughput or memory bandwidth.

Important:

Do not configure more than one Avaya Contact Recorder slice on any physical server.

For detailed instructions on how to configure VMWare, refer to the *Avaya Contact Recorder VMWare Guide, Release 15.2*

Network Issues

In planning the network that will support your Avaya Contact Recorder system, you must consider:

- the additional load imposed on the network
- the IP ports used - so that firewalls can be configured appropriately
- (passive IP recording only) the paths over which audio is transmitted - so that the recorder is able to tap into all the calls that are to be recorded.

Load

You must design your network topology to accommodate the additional traffic created by the recording system. See Recording Bandwidth on page 54 for a discussion of the bandwidth required for each type of recording.

Ports Used

The components of the system use a number of IP ports to communicate:

- between each other
- with various other Avaya components
- with end-users and administrators

Recorder Interfaces on page 302 provides a diagram and table listing all of the interfaces to and from the Avaya Contact Recorder software. Your network and firewalls must be configured to permit traffic to pass over these links.

Network Address Translation Routing

The IP address of an Avaya Contact Recorder is sent to the telephone system as part of the recording process. It is therefore essential that when the telephone system components transmit to this address, the packets are routed correctly to the recorder. The recorder must therefore be visible to the media processing resources, IP phones, Media Access Server or Border Control Point as appropriate on the IP address that it uses itself.

Additionally, if the recorder has more than one NIC card and these are not "teamed" or "bonded" into a single address, it is imperative to ensure that all VoIP packets are transmitted over the same NIC card (i.e. the network route for all recorded audio streams must be the same).

Licensing

A license key is only needed for each recording system. This may be a "standalone" recorder or any number of Slaves and/or Standby recorders connected to a Master recorder. Slaves and Standbys are controlled by the Master and do not require their own license key. Central Replay Servers require a license key each.

Recording Limit

The license key determines the maximum number of Recordings that can occur in the system as a whole. Recording a call uses one recording channel license - regardless of whether this was triggered by Bulk recording configuration, WFO Business Rules or both. Duplicate copies of recordings with multiple sessions, made for WFO's benefit do not use up additional licenses.

Where a recording system includes multiple servers (Master plus Standby and/or Slaves) the limit is applied across all of these recorders.

Backup Recording Channels Limit

To use one or more Standby recorders to provide fault tolerant backup capability, you must license as many backup recording channels as are to be configured on Standby recorders.

Concurrent Screen Recording Limit

The license determines the maximum concurrent number of screens that can be recorded across the system as a whole.

Quality Monitoring Seat Limit

The license determines how many different stations can be recorded by the WFO application. This is not a concurrent limit. Each station that is recorded for Quality monitoring counts against this license.

Telephone Replay Channel Count

The license determines how many recorder channels can be assigned to telephone replay. This feature is only available on Communication Manager based systems.

Dialer Integration

The license key will enable or disable integration to predictive dialers using the recorder's integral support for these. It is not required where integration is via a separate (and hence separately licensed) controller, nor is it required for integration between a CS1000 system and the SER dialer.

Secure Call Recording

For CS1000 systems, the license determines whether or not the audio streams sent to the recorder can be encrypted.

Communication Manager based systems support this as standard when using DMCC recording but this cannot be recorded using passive IP tapping.

Selective Recording

The default for Bulk recording mode is that 100% of the calls that match the recording criteria are recorded (subject to capacity limitations). If screen recording is licensed and the recorder can determine which recordable screen is associated with a call, then 100% of these screens will also be recorded.

If licensed for this optional feature, you can specify what percentage of such calls have their audio recorded and, of those, what percentage also have their screen recorded (where possible). These two thresholds can be set overall and/or for specific bulk recording targets.

Live Monitoring

Optionally, users can monitor voice and (if installed) screen recordings in real time.

Search and Replay API

If licensed for this optional feature, external applications can search for and retrieve recordings matching various criteria. Full details of this API are available on request.

Mass Export

If licensed for this optional feature, the recorder can automatically export recordings as they are made. This is typically used to deliver copies of recordings to external parties and/or to feed third party applications such as speech recognition.

Timed Trials

Avaya can, at its discretion, issue an activation key which includes an expiration date. This allows for timed trials of any combination of features and capacity. As the expiration date is fixed within the license, the server will stop operating at that date regardless of when you enter the license key.

To extend a timed trial or to upgrade to a full license, contact us for an updated license activation key. Contact details are in the section titled Support on page 21.

A five-day trial license is available automatically from the license key entry page. This will allow you to try out Avaya Contact Recorder in a single server topology.

CAUTION:

The five-day trial license must not be used for production recording. When a full license is installed, any trial recordings become unplayable.

Communication Manager System Prerequisites

To use the Avaya Contact Recorder system with a Communication Manager, you will need to ensure that your Avaya system meets the following requirements. This section discusses the various hardware, software and configuration requirements.

Communication Manager

Avaya Contact Recorder requires AE Services and hence is only supported on the models and versions of Communication Manager that support this platform.

Model

The Avaya media server running Communication Manager must be an S8300 through S8800 system.

Station Count

Each DMCC recording port on a recorder is an additional IP phone on the switch. Do not exceed the total station count for the switch in question.

Loading

Each DMCC recording adds as much load to a switch as a normal call. Hence you can only record 100% of calls using this method if your switch is running at no more than 50% of its design load. Please refer to the current Avaya product documentation on http://support.avaya.com.

Software Version

Please refer to the current Avaya product documentation on http://support.avaya.com.

Gateway Resources

These house the media switching components of the Avaya system. You must ensure that the system has, or is expanded to have sufficient:

- Card Slots for the C-LANs and Media Processing Resources described below
- Time-slots for the original calls and, where needed, the recording channels.

Card Slots

Each C-LAN and Media Processing card must be located in the appropriate gateway and therefore in an available card slot alongside the existing cards.

Time-slots

When using DMCC softphones, the recorder conferences into calls in order to record them.

- ANY recording in which the recorder is injecting beep-tone will require one additional time-slot per concurrent recording.
- On Demand and Meeting modes use normal conferencing (as opposed to single-step) and therefore use a timeslot per recording.

Where additional timeslots are needed, the total timeslot count must not exceed the maximum available on that port network (484). Therefore, for a 100% recorded system, using beep tone injection do not equip any port network with more than:

- 6 x T1s (=144 calls, 432 timeslots) or
- 5 x E1s (=150 calls, 450 timeslots)

These guidelines allow for reasonable additional timeslot usage for conferences with other port networks, shared tones and so on.

Rebalance port networks or add new ones to reduce the timeslot requirement on each to this level.

Session Border Controller (SBC)

If your system includes one or more SBCs as detailed below, you can record calls that pass through these using SIPREC. The Avaya Contact Recorder will use SIPREC in preference to other IP recording modes when invited to do so by the SBC.

Compatibility

Avaya Contact Recorder can be used with the following devices:

1. Avaya Session Border Controller for Enterprise (ASBCE) 7.0

Please check the latest Release Note for updates to the list of supported devices/versions.

Codecs

All calls recorded via SIPREC are stored in stereo in G.729A format (16kbps total) but may be received by the recorder in G.729A (with or without Annex B silence suppression) or G.711 (μ- or A-law).

If any audio codecs apart from these three are used, the calls will not be recorded. Any video streams present will be ignored rather than recorded.

Further detail is provided in the Avaya WFO 15.2 Distributor Technical Reference (DTR)

AE Services

Each DMCC recording port on an Avaya Contact Recorder uses a DMCC softphone. The recorder also makes use of TSAPI services. If you wish to configure recording of stations according to the CoR of the station, the recorder will also use the SMS Web Service to determine which CoR each station uses.

These are all provided by Avaya's AE Services platform.

Loading

Note:

To avoid overloading an AE Server, do not attempt to record more than 20,000 calls per hour through each AE Server. (20,000 BHCA).

Note:

You must not use more than 1,000 softphones (recorder ports) through a single AE Server.

If several small recorders are used, you may connect them to a shared AE Server, but only if the total load on the AE Server does not exceed this figure. If the load imposed by a single recorder exceeds this figure, you must split the load across multiple smaller recorders, spreading the load across multiple AE Servers.

Multiple AE Servers

Most Communication Managers can support up to 15 AE Servers but this is a total count - not just those associated with recording. You may have other AE Servers associated with other CTI applications.

Location

In a multi-site system, you should always aim to install an AE server on the same site as the recorder(s) that is (are) using it. This minimizes the chance of system failure due to loss of connectivity between recorder and switching system.

Vintage

Avaya Contact Recorder 15.2 requires AE Services 6.3 or above. Ensure you are running the latest recommended load of AE Services.

Expansion Interface Boards (TN570)

All Expansion Interface Boards must be TN570C Vintage 3 or later.

C-LAN

C-LANs (TN799 DP) are used for two purposes:

- CTI information may be passed through them
- DMCC softphones register through them.

Number of C-LANs

To ensure that a C-LAN does not become a single point of failure in a recording system, you should always provide at least two C-LANs for each AE server. As the CTI load and channel count increases, you should provide more C-LANs as shown below.

C-LANs per AE Server	Maximum BHCA through the AE Server	Maximum Recording Channels through the AE Server
1	NOT SUPPORTED	NOT SUPPORTED
2	12,000	200
3	20,000	400
4	20,000 (AE Server limited)	600
5	20,000 (AE Server limited)	800

Location of C-LANs

For maximum resilience, spread C-LANs across multiple port networks.

To avoid bottlenecks between the port network and the switch, do not connect more than 400 DMCC recording ports to the C-LANs in any port network.

Vintage

Refer to the switch documentation for the release of Communication Manager you are running.

Firmware

The latest update is recommended but these cards must be at least at Firmware update 132.

VoIP Resources

Although TDM and passive IP recording do not use any VoIP resources, each DMCC port on the recorder will use media processing resources on the Avaya system when it is active (recording or replaying via the telephone). You must ensure that sufficient media processing resources are available for the recording and replay load - in addition to any existing use of these resources within your system.

Resource Requirements

G.711 recording or replay uses less resource than G.729A recording. Note that replay is always performed using G.711. Depending on the type and version of your Communication Manager, you may require one or more of the devices shown in the table below.

Resource Type	Comments	G.711 Recording or Replay Ports	G.729A Recording Ports
Media Processing Resource TN2302AP		64	32
MM760 VoIP Module	Included within S8300	64	32
Media Processing Resource 2602AP		320	280

These requirements are solely for the recorder's ports and are in addition to any used by other IP phones or other switch components.

Note:

It is not recommended to dedicate Media Processing resource to recording so it is important to over- rather than under-provide as other users of this resource could otherwise starve the recorder of this capability. On the plus side, you may use existing spare capacity in the switch for recording - but check the location of the resource as well as the amount.

Location of Resources

When adding recording to systems with multiple port networks, it is vital to check that the recordings do not overload the interconnects between port networks.

If a call cannot be recorded using VoIP resources within a port network that the call would have been routed through anyway, then the call must be routed to another port network to reach the VoIP resource. This varies according to whether the phones are digital (DCP) or IP and, with IP phones, whether the system is IP- or Multi-connect based.

Site the VoIP resources according to the table below.

<table>
<tr><th>Recording System</th><th>DCP Phones</th><th>IP Phones Multi-Connect</th><th>IP Phones IPConnect</th></tr>
<tr><td rowspan="2">Station-side - High % (>25%) of calls on trunks recorded. (Bulk or Unify applying station-based rules.)</td><td>Same Port Network as the Phones being recorded.</td><td colspan="2">Same Port Network as the trunks carrying the calls.</td></tr>
<tr><td>N x VoIP resources per phone being recorded.</td><td colspan="2">N x VoIP resources per trunk channel on that port network that could be recorded concurrently.</td></tr>
<tr><td rowspan="2">Station-side - Low % (<25%) of calls on trunks Recorded. (Bulk or Unify applying station-based rules.)</td><td colspan="2">Dedicated port network(s).</td><td>As above</td></tr>
<tr><td colspan="2">VoIP resources equal to N x total number of trunk channels that could be recorded concurrently</td><td></td></tr>
</table>

Where, N=1 for G.711 recording and N=2 for G.729A recording.

Vintage

Refer to the switch documentation for the release of Communication Manager you are running.

Firmware

The latest update is recommended but Media Processors must be at least at Firmware update 105.

Fault Tolerance

You should consider providing one additional board per port network. In the event of a board failing, a spare is then available to handle the full recording load, without having to look outside the Media Gateway - which could introduce sub-optimal use of back-plane timeslots and potentially impact recording in other gateways.

Further Information

For more information, refer to Chapter 2: Administering C-LAN and IP Media Processor circuit packs, in the Administering Converged Networks section of the *Administration for Network Connectivity for Avaya Communication Manager* manual.

Multi-Connect Capacity

Keeping in mind the number and location of recorders and VoIP resources as defined above, confirm that the capacity of the Multi-connect switch (if present) is not exceeded.

DMCC (IP_API_A) Licenses

As long as you are running CM5.1 or later, the recorder does not require or use any of your existing IP_API_A licenses.

TSAPI Licenses

As long as you are running CM5.1 or later, the recorder does not require or use any of your existing TSAPI licenses.

VoIP Network Design

The recorder hosts the equivalent of 1 x Avaya 96xx (recommended) or 4624 IP Phone per DMCC port - whether used for recording or replay. You must therefore design the connectivity between it and the rest of the Avaya switch infrastructure as if there were a bank of this many IP phones at the location of the recorder. Follow Avaya's network design guidelines for this number of IP softphones operating in either G.711 or G.729A mode, but with 60ms packet intervals.

If the bandwidth between recorder and the media processing resources that it uses is less than LAN speeds (100Mbps full duplex) then you should use G.729A recording only.

CS1000 System Prerequisites

To use the Avaya Contact Recorder with a CS1000 system you will need to ensure that your Avaya system meets the following requirements. This section discusses the various hardware, software and configuration requirements.

Contact Center Requirements

To use Duplicate Media Streaming you require SCCS 5.0 or CCMS 6.0 or higher.

If you want to use the MultiDN Recording capability, then a minimum of CCMS 7.0 or later (AACC) is required. MultiDN capability for IP sets allows all physical keys to be recorded on a single IP Phone (Previously there was a 2 key restriction). The CCMS/AACC license for MultiDN must contain the number of DNs or position ID's that are being recorded, including Multiple Appearance (MADN). AST licenses are no longer required on the CS1000 when the multiDN capability is configured.

If you want to use the Record on Demand / Save Conversation feature, CCMS 7.0 or later (AACC) is required. A on/off license key is required on the CCMS/AACC to enable this functionality.

CS 1000 Systems and IP Client Requirements

To use Avaya Contact Recorder you need:

- CS 1000 Release 4.5 or higher.
- For Record on Demand/Save conversation a minimum of CS1000 6.0 is required.
- If Secure Call Recording is being deployed, then Unistim 4 or higher is required. Also, a minimum of CS 1000 release 6.0 is required. Note that the Secure Call recording feature is only supported on specific IP Phones equipped with this feature, currently the 11xx phones.
- If not using the MultiDN feature, AST licenses are required on the Phoneset keys that you wish to record. Note that the MultiDN feature is only available with CCMS 7.0 or later (AACC) and CS1000 release 6.0 or later. When using MultiDN, a license is required on the CCMS server to enable this functionality. Also, MultiDN is only applicable to IP phones, so AST licenses will definitely be required for TDM recording.

  Note:

  Even if you are recording trunk-side, the system still monitors the phonesets you wish to record (not the trunks). You therefore require as many AST ISMs as you have phonesets to be recorded.

To use Duplicate Media Streaming you require Avaya IP Client Phase 2 sets loaded with firmware that supports Duplicate Media Streaming.

Avaya Oceana™ System Prerequisites

See the *Avaya WFO 15.2 Distributor Technical Reference (DTR)* for details.

AACC System Prerequisites

Avaya Aura® Contact Center (AACC) is the next generation Contact Center product from Avaya. Refer to the *Avaya Aura® Contact Center, Planning and Engineering Guide (NN44400-210)* which provides all Server Requirements and prerequisites for AACC.

Supported Topologies

Note that AACC can be configured in multiple ways which can influence the type of recording mode required, as follows:

Avaya MBT or Communication Manager (CM) environment

In this environment, the AACC itself has only one mode of operation and it is installed and configured as an “AACC -SIP” Contact Center. This basically means that the Contact Center interfaces to the Switch infrastructure using SIP trunking for voice sessions (via the SES), and using SIP TR87 (via AES) for CTI events from the switch. From a Call Recorder perspective, the connectivity to AACC is via CCT Web Services, and this is the mechanism by which the Recorder receives additional CTI information for all agent related call events and agent events (such as Login, Logout, Ready, Not Ready).

Avaya Communication Server 1000 (CS1000) environment

In a CS1k environment, the CTI link to the Switch is via a proprietary AML link and all calls are tracked via the MLS interface. This is used in an AACC-AML configuration, but can also be used in a knowledge worker environment. Recording is possible using either Duplicate Media Stream or TDM recording.

Required Components

Note that for a SIP Contact Center solution, the Avaya Media Server (AMS) and Communications Control Toolkit (CCT) components are always required for the base Contact Center functionality.

Topologies

This document has so far discussed functionality in terms of "applications" without being specific as to the physical location of these. As the individual components of the recording system interconnect using IP-based mechanisms, you may distribute the components across your Enterprise's network in a wide range of topologies.

In small systems, a single server can perform recording, storage and retrieval but in larger systems, you can separate these tasks onto different servers in a variety of ways as discussed below. A number of the basic topologies have already been shown in the diagrams of Recording Options on page 30. The following paragraphs define the rules under which each of these basic topologies can be used and introduce the more advanced topologies required for larger and more fault tolerant systems.

Bulk Recording System

At its simplest level, an Avaya Contact Recorder system consists of a single server running the Avaya Contact Recorder software and configured as a Master. This application provides:

- system administration functions via a browser
- voice recording of Avaya phonesets
- screen content recording of Windows desktops
- integration to the Avaya switch via a CTI feed for real-time control and tagging of recordings
- integral search and replay capabilities

See the diagram in Introduction on page 25 for this basic Bulk Recording Topology.

For most small to medium sized bulk recording requirements, this single server is all that is required.

Bulk Recording + Quality Monitoring System

Bulk recording and Quality recording can be combined on a single Avaya Contact Recorder but the Quality Monitoring application is provided by the WFO suite which resides on a separate server. It is possible to configure a virtualized environment using VMWare ESXi (on Windows or Linux) where the Quality Monitoring System and the ACR are installed in their own guest O/S on the same hardware. Care should be taken to follow VMWare recommendations for ensuring accurate timekeeping in the virtualized guest systems. On such a system, WFO may be configured with "Business Rules" that instruct the Avaya Contact Recorder to record specific calls. These recordings are made *in addition* to the bulk recordings and all recordings are made available to WFO as each contact ends.

Large Bulk Recording Systems

Where your recording load exceeds the capacity of a single server, or where a distributed topology is more appropriate, you may deploy multiple recorders - in one of two ways.

Partitioned Systems

If your recording requirements can be completely separated, you can deploy multiple independent recorders, each unaware of the others. This is only appropriate if you are Bulk recording isolated populations of phone-sets on CS1000 systems.

Master/Slave Recorders

Where you are recording Communication Manager calls on the basis of Station, CoR, Agent ID, Skill hunt group or VDN, there must be a single recorder in charge of all recordings. The same is true of all AACC systems and all CS1000 systems except for completely isolated populations of stations.

To support large systems, or to control traffic over your wide area network, you can add further recorders to increase the capacity of the system. Recordings will normally be load-balanced across the Master recorder and any additional Slave and/or Standby recorders.

In such cases, the "Master" recorder is connected to the main CTI feed and is aware of the recording rules - and of the type and locations of the other, "Slave" recorders as shown below for a system recording calls on a Communication Manager.

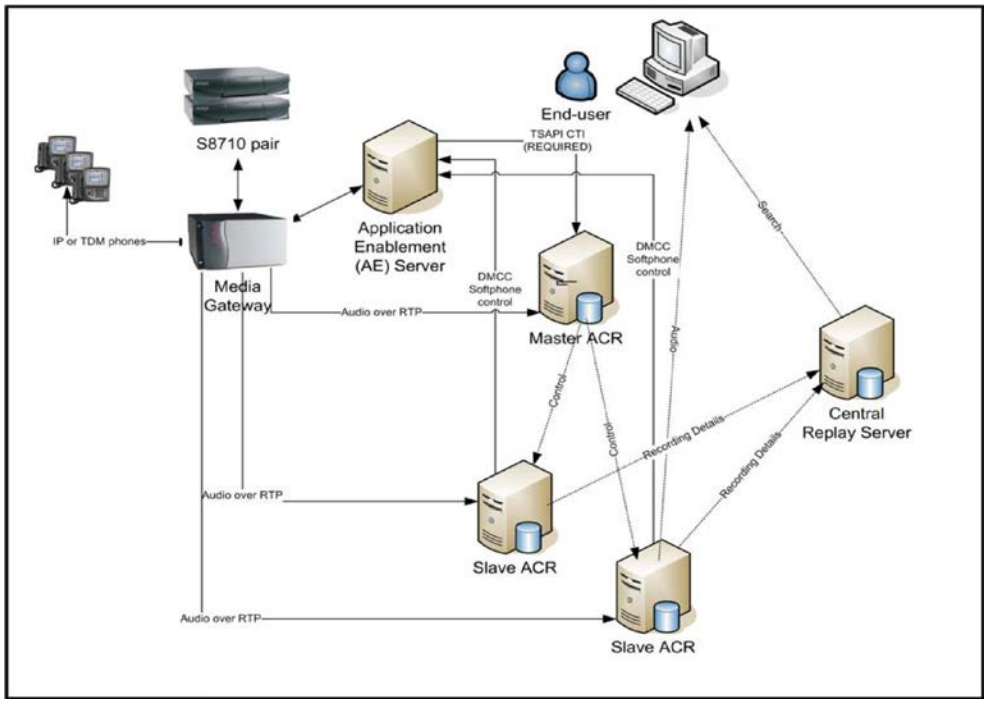

The Master recorder communicates with each Slave via a TCP/IP link. It instructs the slaves to tag the required data/voice streams with the details it learns from the CTI link. Preferably, one or more of the slave recorders is actually designated as a **Full Standby** recorder and can take over should the Master fail (though this requires additional backup channel licenses). Recordings will be made across the Master, Standby and Slaves according to load balancing settings, regardless of which server is active at the time.

Recorder Type and Location

You may distribute recorders across your network. This lets you trade off network load versus security of storage.

For example, if you wish to record calls on an overseas site to which you have limited bandwidth, you can locate a recorder and media processing resources on that site.

High Availability Systems

Because the recording system is based on industry standard PC hardware, you can spend as much or as little as you like on fault tolerant hardware to increase the reliability of each server. Avaya recommends that you use fault tolerant, hot-swappable RAID disk arrays, redundant power supplies and fans as a matter of course. Bonded and/or dual NIC cards are strongly recommended.

See Appendix D: High Availability on page 375 for a full discussion of how to provide fault tolerant recording, storage and retrieval if required.

Standby Recorders

Instead of configuring all additional recorders as "Slaves", you can configure one of these as a "Full Standby" recorder. The standby copies configuration details automatically from the master via a fault-tolerant TCP/IP connection between the two. Over these same links, the two recorders exchange "heartbeats" every few seconds. The standby will take over should the heartbeat fail or should the master request that it do so. This latter case will occur, for example, due to the master's disk filling or connection to the CTI link failing. The figure below shows a typical configuration.

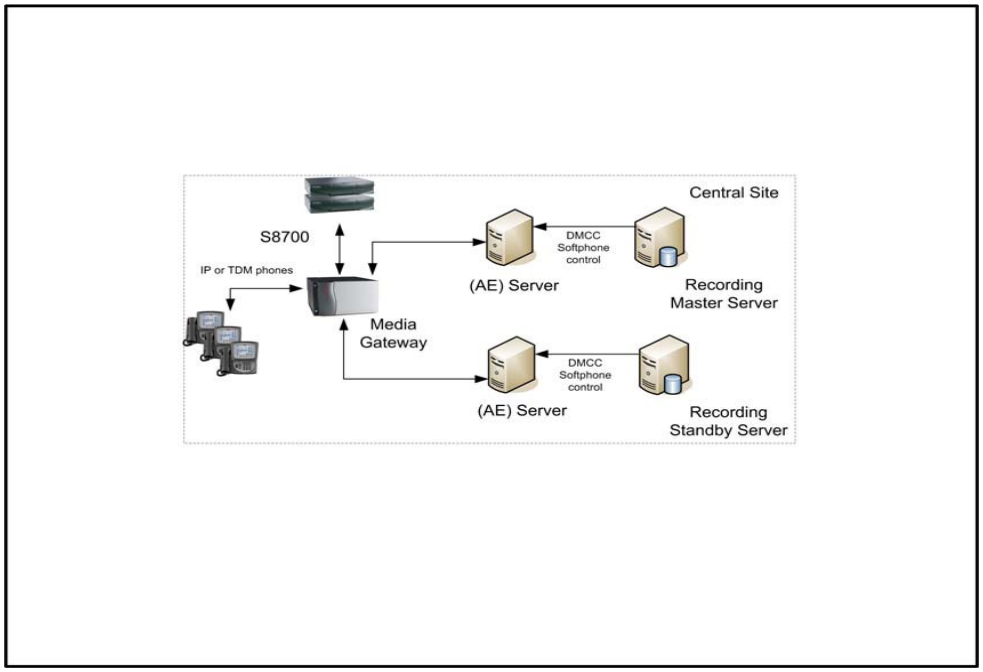

In small systems consisting of just a Master and one **Full Standby**, details of recordings made on each server are (by default) passed to the other so that users can search for calls from a single location regardless of which of the two recorders actually recorded the call they are looking for. In larger systems, where a Master recorder controls one or more "Slave" recorders, the use of a **Full Standby** recorder ensures that the Master does not become a single point of failure for the whole system. In this case, each of the recorders permanently connects to both Master and Standby. This allows the Standby to take over rapidly should the Master fail. In such systems one or more dedicated servers can be provided solely to support search and replay. Details of recordings from all of the recorders are uploaded to these servers allowing a single search to cover all recordings.

More complex, multi-site standby topologies are described in Appendix D: High Availability on page 375.

Fault Tolerant Pools

A "pool" of recorders can be equipped with one more recorder than is needed for the projected loading. In this configuration, if one fails, the Master recorder will reassign the calls being recorded by the failed unit across the remaining servers.

Note:

The exception to this is TDM recording which relies on physical taps onto the extensions or trunks being recorded. In this case, you can provide a fully parallel system instead, providing additional taps into the same points.

Centralized Applications

While recorders are often distributed around the network, most other applications are centralized and available to all from one location. See Optional Server Applications on page 43 for further details on these.

External Recording Control

The operation of your Avaya Contact Recorder can be controlled by a number of external influences.

Business Rule Driven Recording

"Business Rules" configured on an Avaya Workforce Optimization Suite are passed to the Avaya Contact Recorder to control which calls are recorded. These rules allow more sophisticated and flexible control of recording than is available through the recorder's own administration pages.

Desktop Process Analytics ("DPA") Control

This component of the Avaya Workforce Optimization Suite can be installed on one or more agent desktops and used to control recording of telephone calls (with or without screen recording) according to the actions taken on that workstation. A typical use for this is to automatically mask recording (of both audio and screen content) during sensitive security questions.

This same tool can also be configured to start and stop screen-only recordings - for example, of how your agents handle E-mails or Instant Messaging sessions - regardless of whether or not they are also involved in a telephone call.

For further details, see the *Avaya Desktop Applications Deployment Reference and Installation Guide*.

Agent Initiated Monitoring ("AIM")

Where simple, manual control of recording is required, the Agent Initiated Monitoring ("AIM") application can be configured on the appropriate agents' desktops. This is an integral part of the workstation recording client that is installed on a workstation to enable screen recording.

Agents can use AIM to "tag" details onto recordings; to start and stop recordings; mask and unmask recordings or to block recordings (deleting any portions of the contact that have already been recorded).

For further details, see the *Avaya Desktop Applications Deployment Reference and Installation Guide* and the Avaya WFO Agent Initiated Monitoring (AIM) Quick Reference Guide.

Auto-dialer Integrations

A number of "auto-" or "predictive-" dialers work by having an agent make several customer calls over a single, persistent call between the agent and the dialer. For Avaya Contact Recorder to break this recording into the appropriate number of separate "calls" and to label each with the appropriate counterparty details, the recorder listens to events from the dialer. Dialers are configured as described from Proactive Contact (PCS/PDS Dialer) onwards, starting on page 149 (was previously covered in separate Appendix). Note that these typically require an additional license option.

Custom Integrations

Avaya Contact Recorder implements a sub-set of the "Recorder Integration Service API" - through which recordings can be controlled. This is the mechanism through which DPA (detailed above) controls recordings.

The recorder also supports a sub-set of the "Unify API" and associated "Java API Toolkit". These have been used to integrate previous versions of Avaya Contact Recorder with a wide range of other CTI feeds and applications.

However, while existing integrations based on these two interfaces should continue to work, no new integrations should be implemented against them given they are deprecated as of 12.1.

Customers upgrading from earlier versions of Quality Monitoring (which may have included an additional CTI link directly into the Quality Monitoring application) should review their existing integration to determine if it provides the necessary contact and session level detail required for quality assessments to be performed via the Avaya Workforce Optimization suite.

Chapter 3: Installation

This chapter gives details of the steps to install an Avaya Contact Recorder system.

The main sections in this chapter are:

- Overview on page 86
- Avaya System Configuration on page 87
- Order in which to Install Applications on page 99
- Platform Prerequisites on page 100
- Installing Avaya Contact Recorder on page 108
- Updating Components on page 112
- Installing Workforce Optimization ("WFO") on page 114
- Installing Screen Capture and Agent Initiated Monitoring (AIM) Software on page 114

Note:

Always refer to the Release Notes for the specific patch version you are due to install. These may contain additional information that was not available at the time this manual was prepared.

Note:

Upgrading Existing Systems:
This chapter describes how to install the software on a new system. You cannot simply upgrade any Version earlier than 12.1 in place. You must backup the database, install 15.2 on a newly installed operating system platform and then restore the database. Please refer to the release notes for the new version and to the *Avaya Contact Recorder, Release 15.2, Upgrade Guide*. These highlight issues when migrating from earlier releases and may require you to upgrade via an intermediate release - using the release note for the latest patch of that version.

Overview

Installation of a complete recording system requires:

- Configuration of your Avaya system to support recording
- Planning the order in which to install the application servers
- Preparing each server
- Installing the Avaya software on each server
- Installing screen capture software on workstations that are to be recorded (if any)

Avaya System Configuration

Before installing any Avaya Contact Recorder components, you must ensure that your Avaya telephony system is correctly configured and, where necessary, upgraded to support the recording system. As you complete these steps you will be asked to note a number of details which you will need later when configuring the recorder.

Prerequisites

Refer to the appropriate Avaya documentation to apply any prerequisite upgrades and/or additional licenses as detailed in Chapter 2. Then continue with this Chapter, applying the configuration checks and/or changes for your particular switch type(s) as described in the following sections.

Communication Manager Configuration

Use the Avaya administration interface to configure the following items:

Note:

Where page numbers are mentioned on the Avaya administration interface, these are a rough guide only. As new settings are added from one version of Communication Manager to the next, these page numbers tend to increase.

Common UCID Configuration

UCID (Universal Call ID) is a common call id passed between elements in the Avaya network. This must be configured as below to enable ACR to combine CTI feeds from different sources in order to record calls using the most appropriate method.

Customer-options

Set the required system parameters as follows:

1. Run the following command:

   `display system-parameters customer-options`

2. On page 4, verify that Enhanced Conferencing is set to **`y`**.

Features

You must set the following system-wide CM parameters.

1. Enter the following command line:

   `change system-parameters features`

2. On Page 5, set **Create Universal Call ID (UCID)** to **y** and allocate a number to the switch if it does not already have a unique reference. If there is only one switch, set it to **1**.
3. On Page 13, set Send UCID to ASAI to y.

Trunk Group

If using SIPREC recording, identify the Trunk Group connecting Communication Manager to the Session Border Controller and set its parameters as follows:

1. Enter the following command line (replacing **NN** with the appropriate Trunk Group number):

   change trunk-group NN
2. On Page 3, set UUI Treatment to shared.
3. Also on Page 3, set **Send UCID?** to **y**.

Device Names

You are strongly advised to ensure that all devices to be observed, recorded or used in recorder (i.e. stations, softphones, agents, splits and VDNs) have a Name configured. Otherwise you may encounter errors when attempting to record these stations.

Adding IP softphones

You must add a station on the Communication Manager for each DMCC recording or replay port you require on the recorder. Create all stations identically. You will subsequently use the recorder's Administration pages to assign them to the various modes.

1. Use the **add station** command to add as many stations as there are ports on your recorders. Note the station numbers as you will need to enter these into the recorder later.
2. Run **add station *xxxx,*** where ***xxxx*** is the new station's extension that you want to administer.
3. For **Station Type**, enter **96xx** (recommended) or **4624**.
4. For **Security Code**, enter the numeric security code the recorder must use to register softphones with the Communication Manager. Note this security code as you will need to enter it into the recorder later. Use the same code for all stations you create for the recorder.
5. In the **IP Softphone** field, enter **y**.
6. On Page 2:

Set **IP-IP Audio Connections** and **IP Audio Hairpinning** to **n.**

```
change station 11001                                          Page   1 of   4
STATION

Extension: 11001                          Lock Messages? n          BCC: 0
     Type: 96xx                           Security Code: 12345       TN: 1
     Port: S00081                       Coverage Path 1:            COR: 1
     Name: CCE Line 01                  Coverage Path 2:            COS: 1
Hunt-to Station:

STATION OPTIONS
              Loss Group: 19            Personalized Ringing Pattern: 1
                                                   Message Lamp Ext: 11001
           Speakerphone: 2-way                  Mute Button Enabled? y
Display Language: english
 Survivable GK Node Name:
          Survivable COR: internal                Media Complex Ext:
              Survivable Trunk Dest? y
          IP SoftPhone? y
          IP Video Softphone? n

change station 11001                                          Page   3 of   4
STATION
 SITE DATA
       Room:                                          Headset? n
       Jack:                                          Speaker? n
      Cable:                                         Mounting: d
      Floor:                                       Cord Length: 0    Building:
Set Color:



ABBREVIATED DIALING
     List1:                    List2:                    List3:

BUTTON ASSIGNMENTS
 1: call-appr                             7:
 2: call-appr                             8:
 3: call-appr                             9:
 4:                                      10:
 5:                                      11:
 6:                                      12:
```

Administering hunt groups

If you want to use On Demand Recording or Meeting Recording modes, consider grouping these ports into one or more hunt groups for each mode. Users can then access the recording functionality through the number of the hunt group rather than through individual ports.

You could assign all ports for a recording mode to a single hunt group to provide a single shared pool of ports, available on a "first come, first served" strategy. Alternatively, you could split the pools into several independent hunt groups - or even leave some individual ports. For example, a dedicated port to be used only by the conference phone in the board room would ensure that meetings there could always be recorded.

For more information about hunt groups, refer to Managing Hunt Groups in Chapter 7: Handling incoming calls, in Volume 1 of the *Administrator Guide for Avaya Communication Manager.*

Note which ports you have assigned to which hunt groups.

Configuring tone detection

The recorder uses tone detection for **Meeting Recording** and for the delete and retain commands in **Bulk Recording**. To configure tone detection:

1. Type **`change system-parameters ip-options`**.
2. In the `Intra-System IP DTMF Transmission Mode` field on Page 2, enter
   rtp-payload

```
change system-parameters ip-options                              Page   2 of   2
IP-OPTIONS SYSTEM PARAMETERS

 Always use G.711 (30ms, no SS) for intra-switch Music-On-Hold? n

 IP DTMF TRANSMISSION MODE
   Intra-System IP DTMF Transmission Mode: rtp-payload
                     Inter-System IP DTMF: See Signaling Group Forms
```

Network Region setup

The recorder requires the following:

- The DMCC softphones used as recorder ports must be in an IP network region that supports G.729A and G.711MU with 60ms packet intervals and NO OTHER CODECS.
- There must be a media gateway or a media processor resource in the same network region or in an interconnected network region.

To set up a network region as above:

Create a Codec Set

Create a new codec set specifically for the recorder(s) as follows:

1. Choose an unused codec set number for the recorders.
2. Use the **`change ip-codec-set setnumber`** command to create a codec set that uses G.729A and G.711MU.
3. Verify that G.711MU and G.729A are the ONLY codecs in the codec set.
4. Set **Silence Suppression** to **n**.
5. Set **Frames Per Pkt** to **6** - which will show a **Packet Size** of 60 (ms).
6. Set the Media Encryption options to **none**, **aes** and **srtp-aescm128-hmac32-unauth** (the latter option is only supported in recent AES versions).
7. On page 2, ensure that **FAX, Modem**, **TDD/TTY** are all set to **`off`**.
8. Also on Page 2, ensure that **Allow Direct-IP Multimedia** is set to **n**.

```
                        IP Codec Set

    Codec Set: 4

   Audio           Silence       Frames          Packet
   Codec           Suppression   Per Pkt
Size(ms)
1: G.711MU              n           6               60
2: G.729A               n           6               60
3:
4:
5:
6:
7:

    Media Encryption
1: none
2: aes
3: srtp-aescm128-hmac32-unauth
```

Refer to the following for further information:

- For an explanation of administering IP codec sets, refer to the Administering IP Codec.
- Sets Section, in Chapter 4: Network Quality Administration, of the *Administration for Network Connectivity for Avaya Communication Manager* guide.
- For a screen reference, refer to the IP Codec Set Section, in Chapter 19: Screen Reference, of the *Administrator Guide for Avaya Communication Manager*.

Create a network region

Create a new network region and assign the previously created codec set to it as follows:

1. Choose an unused network region for the recorders' softphones.
2. Type **`change ip-network-region`** ***`region`*** where ***`region`*** is the number of the chosen network region.
3. Specify the Codec Set created in the previous step.
4. Set the two **IP-IP Direct Audio** options to **No**.
5. Set **IP Audio Hairpinning** to **n**.
6. Now update the region-region codec table so that all existing regions will use this newly created codec set when communicating with the network region that you have just created for the recorder's softphones.

Refer to the following for detailed information:

- Administering IP Network Regions Section, in Chapter 4: Network Quality Administration of the *Administration for Network Connectivity for Avaya Communication Manager* guide.
- For a screen reference, refer to the IP Network Region Section, in Chapter 19: Screen Reference of the *Administrator Guide for Avaya Communication Manager.*

Assign softphones to the network region

To ensure that the recorder's softphones register in the network region created above, use the **`change ip-network-map`** command.

Add the IP address of the AE Server to the new network region so that all the recorder's softphones - which register via that AE Server - are created in this network region.

The following example shows how to complete the form. In this example, the AE Server is using network region 10 with an IP address of 192.168.2.100.

```
change ip-network-map                                          Page   1 of  32
IP ADDRESS MAPPING

                                                              Emergency
                                        Subnet                Location
 From IP Address  (To IP Address   or Mask)  Region    VLAN   Extension
 192.168.2  .100  192.168.2  .100             10       n
    .   .   .        .   .   .                          n
    .   .   .        .   .   .                          n
    .   .   .        .   .   .                          n
```

Important:

If you are configuring a system for high availability, it is essential that:

a. Calls can be placed from the softphones used as beacons on the master and any full standby servers to and from the beacon softphones on all partial standby servers.

b. You must ensure that the network region settings Direct IP-IP and Hairpinning are both Set to No (else beacon phones will not work at all).

c. You must ensure that there is a compatible codec set between these network regions.

AE Server Configuration

The AE Server provides the recorder with DMCC client services, TSAPI services and (if specifying which CoRs are to be recorded) SMS Web Services. The following instructions assume the AE Server has been installed specifically for recording and that OAM and User Management administrative accounts have been created. If the AE Server to be used is already in use for other purposes, check these settings and confirm that its current configuration is appropriate.

Important:

If you have more than one AES and intend to use ACR's fault tolerant capabilities you must name the Switch Connection the same on each AES.

Administer C-LANs/H.323 Gatekeepers

Use the AES Administration Screens to add the C-LANs that will be used to register softphones.

1. From the **CTI OAM Admin OAM** main menu, select:

Communication Manager Interface > Switch Connections

2. Click **Edit H.323 Gatekeeper**. OAM displays the Edit H.323 Gatekeeper - Switch1 page.
3. In the **Name or IP Address** field, type the ***hostname*** or ***IP Address*** of the switch C-LAN (for non-CLAN systems, enter the PROCR's IP address), and then click **Add Name or IP**.

TSAPI Configuration

Even if there are no explicit CTI clients, you must configure TSAPI (previously known as Avaya CT). Do this after AE Services have been configured as above. If you are using a Security Database as part of your AES TSAPI setup, you must ensure that the recorder is granted access to all the addresses (stations, VDNs and skill hunt groups) that it will need to observe.

1. On AES, go to **Security > Security Database > CTI Users > List All Users**.
2. Select the user setup for WFO and edit it.
3. In User Profile, check the **Unrestricted Access** option.
4. Apply the changes.

SMS Web Services Configuration

If you wish to control which stations are recorded by assigning them to specific CoRs you must enable SMS Web Services to allow the recorder to determine which CoR each station is in. The recorder will use HTTPS to port 443 on the AE Server to determine which stations are in each CoR.

User Account

If you are using the Security Database for Authentication, create a CTI user account on the AE Server as follows:

1. Go to: **User Management > Add User**.
2. Complete all of the required fields (indicated by an asterisk).
3. Select **userservice.useradmin** from the **Avaya Role** drop-down menu.
4. Select **Yes** from the **CT User** drop-down menu.
5. Ensure that the new CTI user has access to all TSAPI-controlled devices by going to: **Administration > Security Database > CTI Users > List All Users**.
6. Click **Enable** next to the **Unrestricted Access** option.

16 Digit Support in CM 8.0 or later

CM 8.0 introduces the ability to support 16 digit extensions. For this to operate correctly with a call recorder, the ASAI link version should be set to a value that is not less than 7.

If the ASAI link version set for the Tlink is less than 7, the CM truncates certain numbers to 15 digits. The ASAI link version number must be set to the highest value supported by the installed AES.

The ASAI Link Version parameter can be configured at AE Services > TSAPI > TSAPI Links.

This value does not get updated with an upgrade of AES. The previously configured value is retained. If you want to add support for 16 digit extensions you must update this parameter manually.

By default, TSAPI is set to use the latest ASAI version supported on the AES server when new Tlinks are added.

ASAI link versions do not match CM or AES release versions, as they have independent numbering schemes.

We recommend that for TSAPI, the ASAI link version should always be set to the highest supported version on AES. AES negotiates with CM to get the highest mutually supported ASAI link version. TSAPI already has an API version to handle discrepancies between the ASAI link version and the TSAPI client version.

SIPREC Configuration

To record calls using SIPREC, you must configure the Avaya Contact Recorder(s), Communication Manager and Session Border Controllers as follows. Note that although SIPREC recording does not require any special licensing of the Avaya Contact Recorder, it may require an additional license on your Session Border Controller.

Avaya Contact Recorder Configuration

If using TLS, configure the certificate as described under Session Initiation Protocol (SIP) on page 253. Note that for SIPREC recording, ACR is the "server" and ASBCE is the "client".

Communication Manager Configuration

You must configure UCID sharing as described under Trunk Group on page 88.

Session Border Controller Configuration

Configure your SBC as described in the *Avaya WFO 15.2 Distributor Technical Reference (DTR)*.

CS1000 Configuration

Update Call Server

In addition to the prerequisites detailed in Chapter 2, you must ensure that your call server has the latest SU and PEPs.

Check/Set Parameters

Ensure the following parameters are set:

Setting	Description	Value
ISAP	Integrated Services Digital Network/Application Protocol (ACD messages sent across the ISDN/AP link). (Overlay 23)	Default=NO. Set to YES only for Meridian Mail applications. Hence for AML/ELAN messages it should remain NO.
SECU	Security Setting for Meridian Link applications. (Overlay 17, VAS configuration)	When set to NO, the host computer must specify both the TN and DN of the associate set in connect, answer and release messages. For AML/ELAN messaging this should be set to YES.
IAPG	Meridian Link Unsolicited Status Message (USM) group. IAPG assigns AST DNs to a status message group defined in LD 15. These groups determine which status messages are sent for an AST set.	The default Group 0 sends no messages, while Group 1 sends all messages. For AML/ELAN messaging it is better to set it to 1.
ICRA (Rel 6) or RECA (Rel 7)	IP Call recording allowed. (Overlay 11)	Set to YES for each IP phoneset that is to be recorded using duplicate media streaming.

Configure ROD and SAVE Keys

To configure either a ROD (Record on Demand) or SAVE key on an IP telephone, use the KEY entry in Overlay 11.

Example:

- **KEY 03 ROD** This configures Key 3 with the ROD button
- **KEY 03 SAVE** This configures Key 3 with the SAVE button

Determine Meridian Link Services (MLS) Connection Details

Note the following details that you will need when configuring the Avaya Contact Recording Master later in the installation process:

- the IP address of the Avaya Contact Center Manager Server
- the IP address of any fallback Avaya Contact Center Manager Server to be used in the event of failure of the default server
- your Meridian1 Customer Number (if using a multi-tenanted system)
- your Meridian1 Machine Name (if using a multi-tenanted system)

Record on Demand and Save Keys

Customers running CCMS release 7.0 or higher (AACC) and CS1000 release 6.0 or higher can give users with IP phones control over recording and deletion/retention of recordings.

(See the *Distributor Technical Reference Bulletin* for additional information on this feature).

To use these features:

1. Apply the necessary CCMS licenses.
2. Add ROD and/or SAVE buttons to the appropriate IP phones

As you set up Bulk Recording on page 168, ensure that you

1. Set the appropriate **Recording Control** options. For standard "on demand recording" you should set:
   - **Start recording automatically at start of call** off
   - **Allow user/external start/restart** on
   - **Allow user/external stop** on

   To have recordings deleted unless deliberately saved you should set:
   - **Allow user/external delete** on
   - **Retain ONLY if requested by user/external** on
2. If using a SAVE key, indicate that this is present on the appropriate DN/Position IDs.

AACC Configuration

AACC Installation

Refer to the *Avaya Aura® Contact Center SIP Commissioning Guide* (NN44400-511) which provides full installation instructions.

Ensure that AACC is running the correct version, contains the required components, and is licensed in accordance with the pre-requisite requirements shown on page 83.

CCT Web Services

CCT Web Services are required for AACC contact center call and agent events, irrespective of whether the underlying recording mechanism is via SIP AMS or DMCC recording. The Avaya Contact Recorder communicates with the AACC (when in SIP mode) via these web services which must be enabled and configured as follows:

- Click on **Start > All Programs > Avaya > Contact Center > Communication Control Toolkit > CCT Console**.
- In the left hand pane of the CCT console, select **Communication Control Toolkit > Server Configuration > CCT Web Services**.
- Ensure the checkbox **Enable CCT Web Services** is ticked.
- Ensure the checkbox **TLS Security** is ticked, if TLS is required for CCT Web Services. Interactions will then be over HTTPS (with TLSv1.2). Both the CCT server and the Avaya Contact Recorder server must have server side certificates, which must both be signed by the same CA. This CA will typically be your in-house CA. To create ACR's certificate:
    1. Create a new certificate as detailed in Creating a new Certificate on page 404 – replacing **<*keystorename*>** with **`ccttls.p12`** and **<*alias*>** with **`cct`**.
    2. Have the certificate signed as described in Generating a Certificate Signing Request on page405.
    3. Import the signed certificate as detailed in Generating a Certificate Signing Request on page 405.
- Note that for AACC 6.1 or later:
    - The **Session Timeout** value on the CCT console should be set to a small value, greater than 1 minute. The default value of 120 minutes (2 hours) is acceptable. There are additional entries in the **CCT Web Services** page for the Call Recording UserID, Domain and Password that must also be configured. With AACC 6.0, a fixed User ID ("CallRecorderUser") must be used. From AACC 6.2, this issue is resolved and any normal windows userid can be used.
- Note also that the entry "**Domain Authentication Server**" is the actual server name of the Server that is running the Domain Controller Software.

- For Domain Authentication Method, use **Simple**. (If the alternative **Digest-MD5** is used, this then requires that the reversible encryption option is enabled on the Domain Controller for the CallRecordUser account).

Test Phonesets

You should provide three Avaya phones close to the recorder - in the same or a neighboring rack. Configure these phones with all of the features in use on the phones that you intend to record. You can then use these to place test calls while working on the recorder.

Order in which to Install Applications

In many cases, the recording system will consist solely of a single Avaya Contact Recorder Master server. However, in topologies that are more complex it is important that you install the basic recording infrastructure first and then layer the other applications on top.

Note:

If you have multiple recorders, install and configure all Avaya Contact Recorders BEFORE connecting these to a WFO system.

Install server components in the order shown below:

1. For each server, follow the procedures in Platform Prerequisites on page 100 then Installing Avaya Contact Recorder on page 108. Install the Avaya Contact Recorder(s) in this order:
    a. Master
    b. Standby(s) – see Standby Server on page 281
    c. Slave(s) – see Slave Server on page 280

Central Replay Server(s) – see Central Replay Server on page 284

2. If using a Key Management Server for encrypted recording storage, install this and configure the recorder(s) as described in Encrypted File Storage on page 255.

Platform Prerequisites

The table in the section Server components on page 42 shows which operating system can be used with each recording mode. Before installing the Avaya Contact Recorder software on the designated server(s), you must prepare each server as described below.

IMPORTANT: The disk partitions used by the recorder MUST be local to the recorder (i.e. directly or SAN connected), and NOT connected across a LAN/WAN (i.e. NOT NAS).

Linux

Version

The operating system for new installations must be Red Hat Enterprise Linux Version 7 update 4 (64-bit). See the Release Note for details of which Update is required. The operating system must be installed using the Red Hat kickstart process. Avaya supplies a tool to generate the kickstart script automatically. If you cannot use the kickstart process you must contact us (as described in the section Additional references on page 20) for guidance.

Disk Storage

You must plan the partitioning of your server's disk(s) in line with the storage needs outlined in Storage Requirements on page 57. The kickstart script partitions the disk(s) as follows:

Mount Point	Use	Size
`/boot`	Bootstrap	200 MB
`/`	Linux Software	at least 5GB
Swap	Virtual Memory	Depends on amount of RAM
`/opt/witness`	Avaya Software and Logs	At least 10GB. 100GB recommended as DEBUG logs can use several GB per day on a large system.
`/var/lib/pgsql`	PostgreSQL Database	See Storage Requirements on page 57. Note 1
`/calls`	The recordings	Remainder of the disk. Note 1

Notes:

1. Select **'CRS' layout** to assign 10GB to `/calls` and the rest to the PostgreSQL Database.

WARNING:

The size of `/var/lib/pgsql` must be calculated carefully based on the number of recordings per month and how many months the database records will be retained. Additional space provided now can save significant effort later if more calls are recorded than expected.

Creating the kickstart script

Kickstart is a Red Hat process that greatly simplifies installation of Linux. You use a Windows, Linux or Mac OS X desktop to run a program that generates an installation script unique to the machine you are installing. When you install, the Linux installer follows the script rather than prompting you with many pages of questions. By using the kickstart script, you not only simplify installation, but assure a standardized deployment.

If you are installing on new hardware for the first time, you will probably need to start a non-kickstart installation before creating the kickstart script. You need to discover the exact device names that Red Hat will give your NIC(s) and disk(s). Specifically:

- On the **Network and Hostname** screen, note the device name(s) of the **Ethernet Controller**(s) (e.g. `ens123`).
- On the **Installation Destination** screen, note the device name(s) of the **Local Standard Disks** (e.g. `sda`).

Once you have discovered these device names, cancel out of the interactive installation process and prepare a kickstart server as follows:

1. Copy the ks3.jar file from the installation media to a Windows, Linux or Mac OS X machine - which must have Java Version 7 or higher installed. This machine will become your "Kickstart server".
2. Check that port 8080 is available, and open port 8080 through the Kickstart server's firewall. The new Avaya Contact Recorder must be able to route through your network to the Kickstart server.
3. Double-click the ks3.jar file.
4. Select the appropriate version of Red Hat.
5. Fill in the form with how the new Avaya Contact Recorder is to be configured:
   - Select the keyboard layout.
   - Select or type in the time-zone.
   - If you have a corporate NTP server, specify its IP address or fully qualified domain name. If you leave the entry empty, the tool picks a suitable public server.

WARNING:

Good time synchronization is vital. You must take care with this setting and test that it is working after installation.

- Fill in the (mandatory) fixed IP address and netmask for the first NIC.
- Specify the address of the default router, which must be in the same subnet as the address that you specify for the first NIC.
- Specify the hostname, preferably as a fully qualified domain name (e.g. acrmaster.bigcorp.com).
- Specify the IP address of a DNS server.
- Optionally specify the root and witness user passwords, take a careful note of these if changing from the default. In the event these are forgotten it will not be possible for Avaya to support the product.
- Click the **'CRS' layout** check-box if you are preparing to install a replay server.

6. Click **Generate Script.**
7. Review the script carefully.
8. (OPTIONAL) If you need to edit the script for any of the reasons in Expert kickstart options on page 103, check the **Allow Edits** box and make the changes.
9. Click **Run HTTP Server**. You can test that the script is ready by using a browser on a different machine to connect to **`http://KickstartServer:8080/ks.cfg`** (replacing *`KickstartServer`* with the actual IP address of your Kickstart server. If you prefer to save the file, click the **Save As** button.

Performing the kickstart installation

When you have proved that the kickstart files is accessible, you are now ready to use it to install the operating system on the new server as follows:

1. Boot the new server using the Red Hat distribution disk. Make sure you have a supported Red Hat version (as detailed in the latest Avaya Contact Recorder Release Note).
2. At the initial screen, select the **Install Red Hat** option.
3. Press the TAB key to bring up the default command line.
4. Add a space and the following parameter to the end of the line, replacing ***`Kickstartserver`*** with the actual IP address of your Kickstart server: **`ks=http://KickstartServer:8080/ks.cfg`**
5. If you have more than one NIC card, add a further space and the following parameter to the end of the line, replacing ***`xxxx`*** with the device name of an

Ethernet card that can reach your Kickstart server (one of the ones you noted earlier and that you have connected to the network): **ksdevice=*xxxx***

6. If you do not have a DHCP server, also append the following three parameters with a space before each
   ip=*x.x.x.x*
   netmask=*y.y.y.y*
   gateway=*z.z.z.z*
   replacing the addresses with the same addresses you supplied when generating the kickstart script.
7. Press the **ENTER** key.
8. Wait for the automated install to complete.

   You may need to confirm that the disks will be completely erased.

Note:

You *must* make a careful note of the passwords you set for the root and witness user accounts during the kickstart process.

Expert kickstart options

NIC bonding (RHEL 7 only)

1. Comment out the un-bonded network line and uncomment the bonded network alternative.
2. Change the slave NIC names to match the hardware.
3. After installation, review the files in
   `/etc/sysconfig/network-scripts.`
4. Delete any referring to **ens0** or **ens1**. Leave the ifcfg-bond* files.

Example:

```
network --device=bond0 --bondslaves=ens1,ens2--
bootproto=static
--ip=10.10.11.44 --netmask=255.255.255.0 --gateway=10.10.11.1
--nameserver=10.10.11.11 --hostname=acrmaster.bigcorp.com
```

Second NIC in different subnet

Uncomment the prototype second NIC line, and manually change the IP address, net mask and device name.

Example:

```
network --device=ens3 --bootproto=static --ip=10.10.12.123 --
netmask=255.255.255.0
```

Two or more disks (designated disks)

You must designate a disk for each partition. The typical way to do this is to create `/calls` on the largest disk and the other partitions on the smaller one. Allow `/var/lib/pgsql` to grow to fill the smaller disk.

Add the **`ondisk`** suffix to each partition.

Example: (note how the default size of **`/opt/witness`** has been increased to allow for larger log files):

```
part /boot --fstype=ext4 --size=200 --ondisk=sda
part / --fstype=ext4 --size=10000 --ondisk=sda
part /opt/witness --fstype=ext4 --size=33000 --
ondisk=sda
part /var/lib/pgsql --fstype=ext4 --size=50000 --grow -
-ondisk=sda
part  swap --recommended --ondisk=sda
part /calls --fstype=ext4 --size=1000 --grow --
ondisk=sdb
```

One or more disks (volume group)

Red Hat Enterprise Linux can join multiple disks into one large volume. This allows greater flexibility going forward. To create a logical volume group, comment out one set of part lines and uncomment the others.

Allow one part **`pv.xx`** line per physical disk.

Set the **`ondisk`** to the device name for each one.

Make sure that each **`pv.xx`** appears on the volgroup line. This joins them together into one large volume group.

Example: (note how the default sizes of `/opt/witness` and `/var/lib/pgsql` have been increased to allow for larger logs and a larger database):

```
part /boot --fstype=ext4 --size=200 --ondisk=sda
part pv.01 --size=1 --grow --ondisk=sda
part pv.02 --size=1 --grow --ondisk=sdb volgroup acrvg
pv.01 pv.02
logvol swap --recommended --vgname=acrvg --name=swap
logvol / --fstype=ext4 --size=10000 --vgname=acrvg --
name=root
logvol /opt/witness --fstype=ext4 --size=30000 --
vgname=acrvg
--name=witness
logvol /var/lib/pgsql --fstype=ext4 --size=100000 --
vgname=acrvg
--name=postgres
```

```
logvol /calls --fstype=ext4 --size=1000 --grow --
vgname=acrvg --name=calls
```

Minimal package installation

By default, kickstart installs the Gnome desktop. If you do not want a desktop or the X windows system, delete the packages associated with the desktop, leaving just the minimal packages required. Note that lines starting "-" tell kickstart not to install a package. These should all be left intact.

Example:

```
%packages
@ base
@ core
perl
mkisofs
perl-libwww-perl
chrony/ntp depending on version
systemd-libs
tigervnc-server
-redhat-lsb
-ash
-aspell
-cups
-cups-libs ...
```

Windows

Version

For all new installation use either:

- Windows Server 2012 R2 (Standard or Data Center Edition), or
- Windows Server 2016 (Standard or Data Center Edition)

Disk Storage

You must plan the partitioning of your server's disk(s) in line with the storage needs outlined in Storage Requirements on page 57. Partition the disks of all servers so that the operating system and recordings storage are both separated and secure. For the Avaya Contact Recording application, prepare three partitions, shown below as C:, D: and F:.

(The E: partition needs no preparation.)

- C: will hold the operating system and other tools/applications - this need only be a few GB. 10 or 20GB is recommended.

- D: will hold the Avaya Contact Recording application. This will include the local call details database and should be sized at 10GB + 2GB per million recordings that you want to keep accessible in the local database. Do not forget that the Master and Full Standby[1] (or dedicated Replay server if you are deploying one) will hold details of recordings made on all recorders, not just their own.
- E: CD/DVD drive.
- F: will hold the actual recordings. See Storage Requirements on page 57 for sizing guidance.

Antivirus and Backup Software

Antivirus Software, Application Backup (including Virtual Machine snapshotting) and PostgreSQL reporting software should NOT be installed on the recorder.

Antivirus software can interfere with PostgreSQL's operation, because PostgreSQL requires file access commands to behave exactly as documented by the underlying operating system, and many antivirus programs contain errors or accidental behavior changes that cause these commands to misbehave subtly. Most programs do not care because they access files in fairly simple ways. Because PostgreSQL is continuously reading from and writing to the same set of files from multiple processes, it tends to trigger programming and design mistakes in antivirus software, particularly problems related to concurrency. Such problems can cause random and unpredictable errors, or even data corruption.

Antivirus software is also likely to dramatically slow down PostgreSQL's operation. For that reason, you should NOT install Antivirus software on the recorder.

Due to the huge volume of files created by the recorder traditional backup software which looks for deltas between backups is not appropriate on the recorder. The heavy disk I/O load caused by such tools severely impacts the recorder's ability to create recording files as well as the speed of the underlying database.

Please refer to Backup/Restore on page 216 for recommendations regarding backing up the recorder and to

[1] The one with highest priority if more than one.

Antivirus and other 3rd Party Applications on page 64 for further details.

Time Synchronization

You must synchronize all servers to the same source as your telephony switch. This will minimize any time differences if you need to compare events on the telephone system with timestamps of recordings or entries in log files.

Java Timezone (TZ) Update

Avaya ensures that the version of Java installed with Avaya Contact Recorder is up to date. However, governments sometimes change time zone rules. If your time zone rules change, you should update the Java time zone rules using the TZ Updater patch for Java. This patch provides updates to the time zone rules and, without it, the previous rules will be applied. This patch is located at http://www.oracle.com/technetwork/java/javase/tzupdater-readme-136440.html

Please check the instructions carefully. Sometimes the operating system must also be patched.

Network Connectivity

Domain Name Server (DNS) Entries

Ensure that the IP node names of all servers that make up the recording system are stored in the appropriate Domain Name Servers. Subsequent configuration can then be done by using the host name rather than having to use numeric addresses.

Important:

The Domain Name Server(s) used by the recorders must support reverse name resolution.

Important:

All IP addresses must be STATIC. Use of dynamic IP addresses provided by a Dynamic Host Configuration Protocol (DHCP) server is NOT recommended and can lead to loss of recording.

Network Routes

You should ensure that valid IP paths exist between each of the servers and from servers to the CTI interfaces, audio sources and recorded workstations. See Recorder Interfaces on page 302 for details of ports used.

Bonded or Teamed NICs

These are strongly recommended in all cases.

If you are deploying any Standby servers on the same site as the master, it is imperative that you provide bonded or teamed (same effect, different operating system) NICs in the Master and Standby servers and fully fault tolerant network connectivity between the Master and Standby. Failure to do so will leave the system vulnerable to failure should a common component in the network paths between the two fail.

In such a case, both recorders will think the other has failed and could attempt to take control of all recording, with unpredictable results.

Installing Avaya Contact Recorder

Note:

You should always refer to the latest Release Note that accompanies the most recent patch in case details have changed since this manual was produced. In particular, the supported versions of operating system change over time.

If you are installing a Central Replay Server, a Standby recorder or Slave recorder, you should first read the appropriate section in Chapter 7: Advanced Configuration on page 264.

Linux

Note:

1. In the following steps you must perform some as the witness user and some as the root user.
2. When connecting remotely to the ACR server you should always log in initially as the `witness` user. When you need to temporarily become the root user, type **`su -`**. Use `exit` to drop back to the `witness` user.
3. When using the console of the server, you should log in as `root`. Having logged in, you should create at least two Terminal windows. Leave one logged in as `root`. In the other type `su - witness` to become the `witness` user. Use the `root` window for root tasks and the `witness` window for witness tasks.
4. First read the Release Notes for the version you are installing. These may contain updates to the information in this manual and specify which Version(s) and Update(s) of Red Hat Enterprise Linux are currently supported.

To install Avaya Contact Recorder on a Linux server on which you have already installed the Red Hat Enterprise Linux Operating System using the kickstart script generator as described in the previous section:

1. As `witness`, copy the required files from the DVD to the `/home/witness` folder on the server:

   `acr-15.2-1.rhel7.x86_64.rpm` - the Avaya Contact Recorder.
   `rpm.postgresql966_rh7.run` - the PostgreSQL self-extracting file.
   `jdk1.8_161.run` - the Java self-extracting file.
   `tomcat8523.run` - the Tomcat self-extracting file.

2. As `root`, install the Avaya Contact Recorder rpm by running
   `rpm –Uvh acr-15.2-1.rhel7.x86.rpm`

3. As `root`, install the PostgreSQL database by running
   `sh ./postgresql966_rh7.run`

4. As `witness`, install the Java environment by running
   `sh ./jdk1.8_161.run`

5. As `witness`, install Tomcat by running
   `sh ./tomcat8523.run`

6. Using the **Patch Tool Utility**, apply Patch000, followed by the latest patch, using the instructions in the Release Notes.

7. Create a certificate for HTTPS access as detailed in Creating a new Certificate on page 404 - replacing **`<keystorename>`** with **`tomcat.p12`** and **`<alias>`** with **`tomcat`**.

8. If you require any of the advanced or non-standard features that are controlled by entries in the properties file, set them now. e.g. **`log.level=debug`**

9. Start the recorder using **`systemctl start acr`** and complete the system configuration.

10. Check the **Status > Server** page for recorder status and the **Alarms** page for any alarms. Refer to Appendix B: Troubleshooting on page 338 to resolve any issues.

11. Continue the system configuration as described in Chapter 4: Configuration on page 115.

Note:

Should you need to uninstall Avaya Contact Recorder, follow the Uninstalling section of the *Avaya Contact Recorder Upgrade Guide*.

Windows

CAUTION:

Before installing on Windows, please double-check these very important pre-requisites from Chapter 2.

For both the install folder (in which the database resides) and the partition you will use as your call recording buffer:

1. Disable Windows Indexing. Right click the drive in Windows Explorer and clear the checkbox on the **General** tab that says **Allow files on this drive to have contents indexed in addition to file properties**.
2. Ensure any Anti-Virus is supported (e.g. not Trend Micro) and excludes these locations.
3. Ensure there is NO backup mechanism attempting to snapshot these locations:
   - Type **backup** into the **Start** box at bottom left and see what programs appear.
   - Open **Control Panel** and type **backup** in the search box at top right to identify applications including Windows own.
   - Remove or confirm that all such applications - including **Windows Backup** - are OFF or exclude the two locations.
4. Check via Control Panel > Programs for any third party monitoring tools installed e.g. SNMP, Disk-space monitoring etc. Either remove or disable these applications or explicitly exclude scanning of the two locations.
5. DISABLE the Volume Shadow Copy Service from the Services window. Do not just set it to Manual as it will start on its own any time Windows decides it is needed. The impact of turning this off is that the built in Microsoft System Restore functions will not run but those should not be running on the ACR server anyway as it is known to cause system failures and is not supported.
6. If using a VM, check there is no disk imaging or mirroring configured for these partitions. Confirm that all partitions are configured as “Thick, Eager”.

ANY of the above can cause catastrophic database corruption.

To install Avaya Contact Recorder on a Windows server on which you have installed the operating system as described in the previous section:

1. Log on to the server using an Administrator's account.
2. Insert the Windows Avaya Contact Recording DVD. This contains the Avaya Contact Recording installation kits.
3. Copy the following install kits to the recorder:

   `acr15.2.exe` - the Avaya Contact Recorder install kit.

   `acrdb961.exe` - the PostgreSQL install kit.

   `acrjre8112.exe` - the Java install kit.

   `acrtomcat859.exe` - the Tomcat install kit.
4. Install the Avaya Contact Recorder by running **`acr15.2.exe`**.
   Follow the instructions on screen to set the path where you want to install the application, this should be the D partition.
5. Install the PostgreSQL database by running **`acrdb961.exe`**.
   Follow the instructions on the screen to set the path where you want to install the database. This this must have adequate space for your call details.

   Note:

   You MUST make a note of the database superuser password as it will be required to backup, restore and upgrade the database.
6. Install the Java environment by running **`acrjre8112.exe`**.
7. Install Tomcat by running **`acrtomcat859.exe`**.
8. Using the **Patch Tool Utility** apply Patch000, followed by the latest patch, using the instructions in the Release Notes.
9. Create a certificate for HTTPS access as detailed in Creating a new Certificate on page 404 - replacing **`<keystorename>`** with **`tomcat.p12`** and **`<alias>`** with **`tomcat`**.
10. If you require any of the advanced or non-standard features that are controlled by entries in the properties file, set them now. e.g. **`log.level=debug`**
11. Start the recorder and complete the system configuration.
12. Check the **Status > Server** page for recorder status and the **Alarms** page for any alarms. Refer to Appendix B: Troubleshooting on page 338 to resolve any issues.
13. Continue the system configuration as described in Chapter 4: Configuration on page 115.

Updating Components

The three underlying components on which the Avaya Contact Recorder depends can now be updated independently of recorder updates. Check the latest Release Notes for any compatibility issues before upgrading as instructed below.

PostgreSQL

Note:

Avaya Contact Recorder 15.2 ships with the current minor release of PostgreSQL Version 9.6. You can only update to later minor releases of 9.6 without requiring a backup and restore. See https://www.postgresql.org/support/versioning/ for further details.

To upgrade the PostgreSQL database:

Linux

1. Stop the recorder service by running **`systemctl stop acr`**
2. Stop the database service by running **`systemctl stop postgresql-9.6`**
3. As `root`, install the PostgreSQL database by running the self-extracting file (example below is for 9.6.6 – this part of the filename will differ for later versions).

   `sh ./postgresql`*`966`*`_rh7.run`
4. The database service starts automatically. Start the recorder service by running **`systemctl start acr`**

Windows

1. Stop the recorder service using Windows Services manager.
2. Stop the PostgreSQL database service using Windows Services manager.
3. Run the updated PostgreSQL install kit - `acrdb966.exe` (where 966 will be replaced by the later version number).
4. The database service starts automatically. Start the recorder service using Windows Service manager.

Java

To upgrade Java:

Linux

1. Stop the recorder service by running **`systemctl stop acr`**
2. As `witness`, install the Java environment by running (***161*** will change to reflect the update number)
   `sh ./jdk1.8_`*`161`*`.run`
3. Start the recorder service by running **`systemctl start acr`**

Windows

1. Stop the recorder service using Windows Services manager.
2. Run the updated Java install kit – `acrjre8161.exe` (where ***161*** will be replaced by the later update number).
3. Start the recorder service using Windows Service manager.

Tomcat

To upgrade Tomcat:

Linux

1. Stop the recorder service by running **`systemctl stop acr`**
2. As `witness`, install Tomcat by running (***8524*** will change to reflect the version being installed)
   `sh ./tomcat`*`8524`*`.run`
3. Start the recorder service by running **`systemctl start acr`**

Windows

1. Stop the recorder service using Windows Services manager.
2. Run the updated Tomcat install kit – `acrtomcat8524.exe` (where ***8524*** will be replaced by the later version number).
3. Start the recorder service using Windows Service manager.

Installing Workforce Optimization ("WFO")

To configure Avaya Contact Recorder and WFO, please refer to the Technical Note *Avaya Contact Recorder Integration to Workforce Optimization Guide*.

Installing Screen Capture and Agent Initiated Monitoring (AIM) Software

To configure Avaya Contact Recorder and Screen Capture (which includes Agent Initiated Monitoring, "AIM"), please refer to the *Avaya Desktop Applications Deployment Reference and Installation Guide.*

Chapter 4: Configuration

You must now configure the recording suite to suit your requirements. This section guides you through the various tasks in a logical order. You should follow its steps immediately after installation of the Avaya Contact Recorder application.

The main sections in this chapter are:

- Accessing the System on page 117
- Licensing on page 119
- Security on page 122
- General Setup on page 126
- System Monitoring on page 159
- Operations on page 163
- /O Jobs: Archive, Import, Mass Export on page 181
- Search, Replay and Live Monitor on page 198
- Backup/Restore on page 216

- Distributing User Instructions on page 221
- Configuring Avaya Support Remote Access on page 222

Accessing the System

Before you can configure the system, you must first:

- access the administration web-interface via its URL
- log in

URL

You administer the Avaya Contact Recorder system via a web interface.

1. Open Internet Explorer and navigate to **`http://servername:8080`** using the name of the server you wish to administer.
2. Enter a username of your choice and leave the password field blank.
3. Click **OK**.

Note:

The login page uses Javascript. If you see the login page but nothing happens when you click OK, your Internet Explorer settings may be blocking this. See Forwarding Replay Requests with NAT and/or SSL on page 208 for detailed instructions on this and other necessary settings.

Initial User Account

The application will accept any username during the first log on attempt after installation and will automatically create a local application account for you under that name and give it full system administrator rights.

As the password of this account has not yet been set, the web application immediately directs you to a page asking you to set the password for this account. In this instance, leave the **Old Password** field blank and enter a new password of your choice into the two other fields. This password must be at least 8 characters long, include upper and lowercase characters, at least one digit and at least one special character (#, @, %, ! or $). Click **OK**.

Important:

Make a note of this username and password otherwise you will not be able to access the web application in the future. Note that both username and password are case sensitive.

Key Points

Before using the System Administration pages, familiarize yourself with the following key points.

Invalid settings

Any of the system's settings that are known to be invalid are shown in red. Use the information in this guide to change the settings to valid values. If you change a setting, but submit an invalid entry, a message indicates the reason that the entry is rejected and you are prompted to re-enter it. To quit without changing a parameter, click on the **Close Window** link.

Show All

At the top of pages that show a list of entries that spans more than one page, the **Show All** link appears next to the page selection tags. When you click this link, all the search results are presented on a single scrollable page.

Page at a Time

If you have clicked the **Show All** link described above, you can return to seeing one page of entries at a time by clicking this link.

Impact of changes

When you change a setting, the window into which you enter the new setting explains the meaning of that setting and the consequences of changing it. Read these notes carefully. Some settings require you to restart the recorder while others may truncate current recordings.

Licensing

Until you enter a valid License Activation Key, connect to a licensed Master recorder or select the Five day, timed trial license, the application will only show you the license entry screen. Note that the timed trial license only allows you to run a single (master) ACR server.

Terminology

License Generation Key

This is a three-digit number that is specific to a particular server. This is shown on the license entry page. You will need it to obtain a valid License Activation key for a Master recorder and Central Replay Server.

License Activation Key

This is a long (30 or more characters) string containing the serial number, server type and other options that you have licensed. You must obtain this key and enter it into the administration pages before your master recorder or Central Replay Server will operate.

Recorder Serial Number

This is a unique identifier for every Avaya Contact Recorder. For Avaya Contact Recorders, this is a 6-digit number starting with 8. The serial number of each Master and Central Replay Server is allocated by Avaya as part of the licensing process.

The serial number of a Slave or Standby recorder is chosen by you. You should assign each Slave and Standby recorder in your enterprise a unique[1] number from 1 to 9999. The serial number for these servers will then be shown as 880000 plus this locally chosen recorder number. The serial number defines the first 6 digits of the unique reference number given to each call recorded by the recorder. For example, the recorder with serial number 800001 records its first call into the following files:

800001000000001.wav

800001000000001.xml

The serial number is encoded within every activation key issued. Once a Master recorder or Central Replay Server has been configured with its initial activation key, subsequent keys must have a matching serial number.

[1] Some older systems assigned the master a serial number also beginning with 88. In this case, ensure that the number you assign does not result in a clash with the master's serial number.

Important:

The 5-day license option uses serial number 800000. This temporary and non-unique serial number is the only serial number that you can subsequently override with the recorder's correct serial number, which is included in the full license key.

Obtaining a License Activation Key

Obtain a license activation key for a Master or standalone recorder or Central Replay Server as follows:

1. Open another browser window (on this or another PC) and ensure that pop-ups are not blocked.
2. Navigate to the Verint licensing website at https://licensing.verint.com/acractivation.
3. Enter your Username and Password.
4. Click **Log in**.
5. Click **License Activation**.
6. Choose the Serial Number from the drop-down list.
7. Enter the three-digit License Generation Key from the License page in the Administration application.
8. Enter the appropriate information for the end user.
9. Click **Generate Key**.
10. Your license activation key is:
    a. Displayed on the screen
    b. Sent to you through email

Activating the License

1. Return to the **System > License** page that you have open in another browser window.
2. Enter the License Key. The license activation key is not case-sensitive, and you can omit the dashes. If you use a browser on the same machine to obtain the activation key, you can copy and paste the number between the browser windows.
3. Click **Enter.** The page displays the licensed serial number, server type and channel capacity.
4. After entering or updating a license key you should always restart the recorder to allow any changes to take full effect.

Make a note of the license key and store it safely in case you need to reinstall the application on the same server - in which case you will be able to reuse the key. To reinstall on a different server, you will need a new key, because the MAC address, to which it is tied via the three-digit license generation key, will be different.

Once you have successfully entered a license key, you will be able to access the other pages of the administration interface.

Central Replay Servers

Enter the provided license key into the box shown. If you are configuring a secondary replay server, enter the IP address or hostname of the primary replay server that it is to shadow.

Standby and Slave Servers

These do not require a license key. Simply enter the Recorder Number you wish to assign to the server and provide the IP address(es) of the already licensed Master server and appropriate Standby server(s) as instructed on the lower half of the **System > License** administration page and explained in further detail in Licensing (Slave servers) on page 282 and Licensing (Standby servers) on page 280.

Adding additional licenses

Follow the procedure in Activating the License on page 120 for the installation of additional licenses. You must restart the recorder after changing any license settings so this is best done out of hours.

Note:

Additional licenses may require more switch components (for example, in a Communication Manager system, more C-LAN, VoIP resources, media processing boards).

Reinstalling on the same PC

If you reinstall the recorder software on a new hard disk in the same chassis, you can reuse your existing activation keys.

If you reinstall the recorder from the installation kit, you must re-enter the activation keys.

WARNING:

If you reinstall the software (having already backed up the database as described in Backing up the Database on page 216), you must restore the database as described in Restoring data to a new PostgreSQL database on page 218 before starting recording. Otherwise the recorder will reuse recording identifiers that have already been used.

Reinstalling the Recorder on a new PC

For license security, the installation is tied to the first Network Interface Card (NIC) in the server on which it is installed. To reinstall the recorder on another server, do one of the following:

- Move the first NIC to the new server and use your existing activation keys.
- Note the license generation key on the new server and request new activation keys as outlined in Obtaining a License Activation Key on page 120.

WARNING:

If you reinstall the software (having already backed up the database as described in Backing up the Database on page 216), you must restore the database as described in Restoring data to a new PostgreSQL database on page 217 before starting recording. Otherwise the recorder will reuse recording identifiers that have already been used.

Security

Security of recordings is very important and is discussed at length in

Chapter 6: System Security on page 239. At this stage in the configuration of your system you should immediately create appropriate user accounts as described below.

Securing the System

The system automatically creates an initial system administrator account as you log in for the first time. If you wish to create additional system administrator accounts, you should do so now. Until you do so, the only means of accessing the system is with the initial account.

To create a new user account:

1. Click on **System > Manage Users** at the top of the Administration screen.
2. Click on **Add User**.
3. Enter the user's name (with domain name and entirely in uppercase if using Windows domain accounts).
4. If using local (non-domain) accounts, you must also enter a temporary password for this account. The user will be forced to change this when they log in for the first time. You must tell the new user what this temporary password is.

Select the appropriate **Role** or **roles** for this user. **System Admin**istrators have full access to the system including all configuration. **Restricted Admin**istrators cannot change the system's recording configuration but may (by default) view the overall status, alarm and archive pages and configure non-administrator accounts.

Notes:

1. The Comment field is for your own notes.
2. See Access Rights on page 199 for an explanation of the remaining fields on this form.
3. You may extend the **Restricted Admin.** role – granting access to the **Status > CTI Monitors** and/or **Ports** pages. Enable each by adding the respective property file setting:
   `restrictedadmin.mayviewctistatus=true`
   `restrictedadmin.mayviewportstatus=true`
4. You may restrict the **Restricted Admin.** role - blocking access to the user accounts page. Do this by adding the property file setting:
   `restrictedadmin.mayeditaccounts=false`
5. You need only administer user accounts on Master and (Primary) Central Replay Servers as user accounts are automatically copied to other servers as appropriate.
   - All user accounts are copied from a Master to its Standby recorder(s).
   - **System Admin.** and **Restricted Admin.** accounts (only) are copied from a Master to its Slave recorder(s).
   - All user accounts are copied from a primary Central Replay Server to any secondary Central Replay Servers.

Note:

As user account details are only passed out from the Master and Primary Central Replay Server, it is important that users are logged onto these servers (and not a Standby, Slave or secondary Central Replay Server) when changing their password. This applies to the initial login where their temporary password is changed and to any subsequent changes.

Windows Authentication

If you enter a simple username without a domain qualifier (for example, ***admin*** rather than ***ADMIN@DOMAIN.COM***) then the account is administered by the recorder application. Administrators may change this account's password and the user may change their account password and log off from the application using the administration web pages.

However, if you specify a domain and username, the system will attempt to authenticate you via Kerberos. In this case, you will not see **Logoff** or **Change Password** links on the web pages presented by the recorder.

CAUTION:

When adding user accounts that correspond to Active Directory accounts, enter the username and domain all in uppercase - regardless of the case used in Active Directory.

This feature is also known as "Single Sign On" (SSO) as users only have to log onto their Windows workstation, but do not need to log on to the search and replay application separately. See Windows Domain Authentication on page 240 for details of how to enable this feature.

Windows Accounts for Screen Recording or AIM

In addition to entering the domain accounts of those who will administer the system and replay calls on it, you may also need to enter account details for the domain accounts of any agents (including Citrix agents) that you want to

- Record whenever they are logged on at a Windows desktop configured for screen recording, and/or
- Use Agent Initiated Monitoring (“AIM”) to control recordings.

Simply create the account (with the username in the traditional ***domain\username*** form rather than the Active Directory style used above as it has to match the login style reported by the Windows desktop itself) and set the Agent ID field. This is used to identify which workstation or thin client session they are logged on at and whenever audio is being recorded, then so long as the desktop they are using has screen capture installed, screen content will be recorded too.

Note:

Preferably, if you are using WFO, you can *instead* define the relationships between agents, employees and user accounts within WFO. Avaya Contact Recorder will use this information when determining which screen is associated with which telephone call for both bulk and business rule driven recordings. The mechanism described here is only required for systems that do not use WFO.

General Setup

You should now configure how the system makes recordings and how it interfaces with your Avaya telephony system. Follow the procedures below, clicking on the appropriate tab at the top of the Administration screen for each section.

Server

This is the section where you specify information about the recorder hardware and environment and the role this server is to take.

Note:

It is important that you review and change any of the settings that are not correctly defaulted. Some settings only apply to servers with particular roles and will not appear on other servers. Many settings are automatically copied from the Master or may be fixed due to the value of other settings. Such read-only settings do not show an **Edit** link next to them.

Recorder Pool

Each recorder must be assigned to a specific **Recorder Pool** for the purposes of **Designated Pool(s)** assignment and failover. Typically, all recorders on a site form a single **Recorder Pool** named after that site.

Normally, recorders are assigned to pools that match the Communication Manager “community” they are located within.

Recorder Pool names are case insensitive but you must take care not to accidentally create a separate pool by mistyping the name on one server.

Sunny-day Share of Undesignated Recordings

Unless you use the **Advanced** setting to designate specific recorder pools, Bulk and WFO Business Rule driven recordings are spread across available recorders according to this setting. You can select one of three ways in which recordings that don't have a specific Pool are allocated to the recorders during normal (or "sunny day" operation):

- **Fair**: This recorder will take its fair share of such recordings such that the percentage of its capacity in use across all recordings is similar to that of all recorders that are configured this way.
- **Nominal**: This recorder will be assigned up to but not more than one concurrent recording channel. This puts a minimal load on your WAN and storage and ensures that any problems with recording are spotted early – before the full capacity of this recorder is needed. Use this setting for example, if you have a recorder in a Disaster Recovery (DR) site but don't want it to take a significant load until it has to.
- **Zero**: This recorder will not be assigned any recordings that are not explicitly designated to its pool. This is appropriate for recorders that you are explicitly designating some recording targets to. Otherwise, a **Nominal** load is recommended.

Maximum Concurrent Recordings

The load on a recorder is the total of the active voice recording channels and screen recording channels. Although the overall system limits on these are each set as part of the system's license, this setting determines the maximum capacity that will be attempted on this particular server. It may be entirely voice, entirely screen or a mix of the two. There may be other factors that restrict its capacity before this level is reached. For example, the license may not include this many channels; there may not be enough DMCC softphones configured, and so on.

It is very unusual to even approach this figure on a server but it is also used in the load balancing algorithm to determine what percentage load each server is currently experiencing. So, for example, a server with this value set twice as high as on another will be given twice as many recordings as the other server.

Leave this on the default unless one or more of the following applies:

- The server specification is lower than that of the Mid-range server recommended for this channel count. See Sizing on page 61 for details.
- Master or Full Standby recorder on a very large system with high calling rates (more than 100,000 Busy Hour Call Completions.
- Recorders in a pool have substantially different capacities and you wish to assign more load to the more powerful servers.

Beacon Softphone (Communication Manager only)

Pick one of the softphones entered at the bottom of the Communication Manager's admin page to be the dedicated Beacon Softphone for this server. Each server registers its own, unique Beacon Softphone whenever it can and uses it to try and call other servers' Beacon Softphones. This lets each server determine whether it is in the same Communication Manager "fragment" as the other servers and whether or not they too can register softphones.

This, in turn, is used to ensure that one and only one server tries to control recording at a time – even when a fragmented system subsequently falls back to single system.

As softphones typically have a 5-minute keep-alive timer, it can take 5 to 10 minutes to determine changes to fragmentation and/or server status via this mechanism.

Note:

You must perform a complete system restart after changing the beacon softphone on any server – as that phone may already be in use for bulk recording on any of the other servers. To minimize disruption, during installation, set all servers' beacon softphones out of hours then restart the entire system.

See Appendix D: High Availability on page 375 for further details.

Important:

If you are configuring a system for high availability, it is essential that:

a. Calls can be placed from the softphones used as beacons on the master and any full standby servers to and from the beacon softphones on all partial standby servers.

b. You must ensure that the network region settings Direct IP-IP and Hairpinning are both set to No (else beacon phones will not work at all).

c. You must ensure that there is a compatible codec set between these network regions.

Standby Coverage

This setting and the following three are only visible on servers that have been configured as Standby servers via their License entry page. See Appendix D: High Availability on page 375 for a full discussion of Standby servers before altering these settings from their defaults.

If any of the following apply, any Standby server is implicitly a **Full** Standby and this setting becomes read-only:

- There is only one **Recorder Pool** configured.
- The system is recording more than one Communication Manager.
- The system is recording any CS1000 switches.

If, however, you have configured more than one **Recorder Pool** with a single Communication Manager, you have options. In the same way that Communication Manager supports “survivable core servers”[1] that can control the entire system when needed and “survivable remote servers”[2] that can only control their local resources, so you can select from:

- **Full**: This server will attempt to control the entire system when needed.
- **Partial**: This server will attempt to control recorders *within* its own **Recorder Pool** to perform those recordings that include its **Recorder Pool** within their **Designated Pool(s)** setting.

Standby Readiness

This setting is only visible on servers that have been configured as Standby servers via their License entry page. See Appendix D: High Availability on page 375 for a full discussion of Standby servers before altering this setting.

Standby servers are always **Hot** on CS1000 systems.

Warm is mandatory for all **Partial Standby** servers but optional for **Full Standby** servers on Communication Manager systems.

To know what to record on a Communication Manager system, a server (Master or Standby) must register TSAPI observers on stations, VDNs and Skill groups. There are system-wide limits on the number of different AE Servers that can observe each. For example, only four servers can observe a station – and other CTI applications may already be using some of those.

Partial Standby servers only control recording when their site is running in Remote Survivable mode – and that requires a restart of TSAPI connections and re-establishment of all observers anyhow – so there is no point having these observers

[1] Formerly Enterprise Survivable Servers (ESS).

[2] Formerly Local Survivable Processors (LSP).

present until then. **Partial** Standby servers are therefore always **Warm** and this setting is read-only.

Full Standby servers can be configured as:

- **Hot** – with all TSAPI observers registered from start-up. This can save up to an hour of partially interrupted recording on failover. It also reduces the load imposed on the Communication Manager during failover.
- **Warm** – with no TSAPI observers registered until the server has to take control of recording. This puts a significant load on the Communication Manager (at a time when it may already be stressed) and can take up to an hour or more to complete on a large, heavily used system with thousands of stations, VDNs and Splits to be registered. Only use this option as needed to avoid hitting limitations on the number of concurrent TSAPI observers.

  Note:

  If the Master is letting a **Full Standby** server remain in control, it will act in whatever mode that **Full Standby** is configured – whether **Hot** or **Warm** – so as not to exceed the number of TSAPI observers that would be present when the Master is in control.

Standby Priority

This setting is only visible on servers that have been configured as Standby servers via their License entry page. See Appendix D: High Availability on page 375 for a full discussion of Standby servers before altering this setting.

You can provide more than one standby server in a given pool. For example:

- You may have all recorders in a single pool across a main and disaster recovery (DR) sites, sharing load across both sites in normal operation. It makes sense to have a **Full** Standby co-located with the Master server to protect against failure of that server and to have a further **Full** Standby in the DR site in case the main site is out of action or isolated.
- If a remote site has two or more servers, two may be configured as **Partial** Standby servers so as to ensure recording control should either server fail.

In such cases, more than one Standby server *may* be able to take control. This setting determines which will do so. A lower number indicates a higher priority – so a Standby server with **Priority** 2 will defer to one with **Priority** 1. If two servers have the same priority, the one with the higher serial number will defer to its (presumably more aged) peer. All Standby servers defer to the Master.

Note that the **Beacon** mechanism described in

Beacon Softphone on page 128 automatically detects and resolves cases where more than one server is in control.

Call storage path

On Windows servers, you must specify the path into which the recorder will store its recordings. This should be the F drive.

On Linux servers this is automatically set to the /calls partition.

Warn when available channel license count falls BELOW

If the recorder is loaded to the point where there are very few licenses available, it will raise an alarm.

Warn when available screen recording license count falls BELOW

If the recorder gets close to the licensed limit for concurrent screen recordings it will raise an alarm.

IP Address to use on this server for Recordings

You must specify which IP address (and therefore which NIC) on the server should receive any audio streams that are directed to the recorder. Even if you do not intend to use RTP recording, set this to an appropriate NIC card anyway.

This setting does not affect passive IP tapping. If you configure the recorder to use passive tap IP recording, you must provide one or more Network Interface Cards that are dedicated to passive tapping. The recorder will automatically accept packets over any Network Interface Card that does not have an IP address. It is therefore important that you check how the cards have been configured as they may well have been given an IP address during installation of the operating system. Remove the IP address from the card(s) you intend to use as promiscuous mode taps.

Designated screen recorders

By default, each new screen recording is started on the lightest loaded recorder in the system. Use this setting to force specific groups of workstations to be recorded, instead, on specific recorders - for example, those in the same location as the workstations themselves.

Use Classless Inter-domain Routing ("CIDR") notation (e.g. "`10.10.10.0/24`") to specify IP address ranges for workstations that should be recorded at ("@") a specific recorder (e.g. "`890001`") or recorder pool (e.g. "`london`"). Separate entries with a semi-colon. For example:

`10.10.1.0/24@london;10.10.10.0/23@newyork;10.10.20.0/25@891234`

The above example specifies that:

- workstations 10.10.1.0 to 10.10.1.255 should be recorded on a recorder in the "london" pool
- workstations 10.10.10.0 to 10.10.11.255 should be recorded on a recorder in the "newyork" pool
- workstations 10.10.20.0-10.10.20.127 should be recorded on recorder with serial number 891234

You must enter an appropriate value for this setting on each Standby recorder as well as on the Master. Take into account the likely network partitioning and hence availability of recorders under the failure scenarios that would cause each Standby to become active.

Keep audio in stereo wherever possible

Some recording mechanisms provide stereo streams to the recorder. If these are not already compressed (in G.729A format), the recorder will compress them to G.729A[1]. By default, the recorder keeps these as two separate streams stored in a stereo file.

Setting this to **No** results in an incoming stereo G.711 media format being mixed into a mono stream before compressing them – giving a mono file.

Setting this option to **No** therefore loses the speaker separation on recordings that were received in uncompressed stereo but does:

1. Save storage space (8kbps instead of 16kbps).
2. Reduce CPU usage significantly.
3. Allow export of recordings. See Recording Format Restrictions on page 190.

Note:

For an incoming stereo G.729A media stream, due to the overhead involved, the recorder will not decompress, mix and re-compress the two audio streams together so the result will still be a stereo file – regardless of this setting.

Maximum recording segment duration (mins)

To optimize the playback experience, the recorder cuts long recordings into segments of the designated length. A typical value for call centers is 120 minutes. This is the default. However, if your switch regularly handles longer calls, you may increase this value.

SNMP Settings

See SNMP on page 162 for notes on these settings.

[1] Unless property acr.disablecompress=true has been set.

Replay Server(s)

By default, each Avaya Contact Recorder server will pass details of calls it has recorded to the Master and **Full Standby** (if one and only one is present). If you have other Standby server(s) or have installed one or more dedicated replay servers, you must override this default behavior by entering the address(es) of the replay server(s) here.

- Separate the addresses with a semi-colon.
- If you have two Central Replay Servers, you should configure each with the address of the other here - so that recording lock/unlock requests will be passed between them.

Once you have set this value, it will be acted on within one minute and subsequent recordings will be queued to be sent to the address(es) specified.

Note:

Do not specify the server itself - as there is no need to upload recording details that are already present locally.

When explicitly specifying a Master or Standby server here you must use the same notation (either hostname or IP address) as used when you configured the server under the License entry page.

If you intend to install a Central Replay Server and want all call details to be uploaded to it, enter its address in this field before you start recording. Do this even if the replay server is not yet ready. This way, all recordings' details will be queued ready for upload when the server is available. The recorder will raise an alarm daily until you install and correctly configure the replay server. Ignore this alarm.

Key Management Server

If you intend to encrypt recordings, refer to Encrypted File Storage on page 255 and set the name of your Key Management Server here.

Key Management Certificate Passphrase

If you intend to encrypt recordings, refer to Encrypted File Storage on page 255 and set the passphrase for your Key Management Certificate here.

URL(s) of external control port(s) to connect to

Specify the IP node name of each external controller. The port number will default to 1414 but to override this add a colon then the port number. If the recorder is supporting multiple servers, list their names separated by semi-colons.

For every URL entered here, the **Status > Server** page will monitor the status of the link and the Alarms page will show any problems with the link.

Data Sources

Avaya Contact Recorder can be connected to one or more Avaya systems (so long as the total recording load is within the guidelines given in Sizing on page 61). These can be any mix of:

- Communication Manager (including "Elite")
- CS1000
- Avaya Aura® Contact Center (AACC)

If licensed for dialer integration,

- Avaya Proactive Contact (PCS) and/or
- Proactive Outreach Manager (POM) can also be connected to the system as described in this section. Each of these provides one or more CTI feeds and are therefore collectively known as "data sources."

Configuring Data Sources

A single admin page, (initially entitled **General Setup > New Data Source**) is provided automatically. Use this to configure the first switch. You must restart the recorder if you add, delete or change the type of any data source but such major configuration changes are normally only required during installation.

Network Configuration

Data Source connections are normally via TCP/IP and are configured using hostnames. You must therefore ensure that the ACR server has been added to your name server, or to the server's local name service (i.e. the hosts file). All machines must be able to resolve their hostname to an IP address.

Adding a Data Source

If your Avaya Contact Recorder is to connect to additional data sources, click the **Add Data Source** button at the bottom left of the **General Setup > Server** page to create further pages as needed. When you configure the Data Source, the title of the page will change to identify the data source.

Deleting a Data Source

If you add a data source by mistake and need to delete it, first ensure that its page has a name other than **New Data Source** - (by adding dummy configuration details if necessary). Note that name and enter it into the **Delete Switch** field on the Maintenance page (`/servlet/acr?cmd=mtce`).

Conflicting Number Plans

If the same number (say, "1234") can occur on more than one switch, instruct the recorder to let you set a switch-specific string that will be added to internal numbers by adding the property file setting to all recorders in your system:

```
cti.switchidentifiers=true
```

The Administration page for each Data Source will then let you set a specific identifier. Choose a (very) short identifier - just long enough to uniquely identify the data source from any others. You only need to set this field on the Data Sources that have overlapping number plans. You can use a few digits (e.g. the local public prefix or the prefix you plan to give these switches when they are fully integrated into your numbering plan. Typically, a 1- to 3-digit number or two- or three-character abbreviation are used. Use an under-bar ("_") rather than dash ("-") if you want to visibly separate the prefix from the number. For example, "NY_" or "SF_". Commas are not allowed either.

You must restart the recorder after setting the identifier on a switch. It is not recommended to change identifiers after the system has gone live.

Data Source Type

Use the top setting on the **General Setup > New Data Source** page to define which type of data source the recorder is connected to. Note that dialer integrations are licensed separately and will not appear as options unless the server has been licensed for them.

You must restart the recorder once you have set this field so that the page is recreated with the details that are appropriate to the specific type of data source.

WFO Configuration (where present)

Each data source configured in Avaya Contact Recorder must have a corresponding data source (one with exactly the same name) configured in WFO.

Communication Manager

When recording a Communication Manager, configure these settings:

Switch Identifier Prefix

This setting is only present if you have configured the recorder (as described above) to support conflicting numbering schemes. It lets you specify a few characters that will be added to the front of each internal phone number that a recording is tagged with. You should only set this for switches whose numbers could be confused with those on other switches. Do not use commas or dashes within the string. Restart the recorder after setting this value. Changing this setting after recordings have been made is not advised as search and replay rights depend on it. Note that you will need to explicitly assign search and replay rights to these extended addresses if you use this option.

Minutes after which call information indexed by Call ID is discarded

Each call is given a reference number or "Call ID" but these are reused after a period and hence any information gleaned by the recorder must be discarded before that call ID is reused. The default of three hours is appropriate for many switches but if you have, for example, an auto-dialer or very high calling rates you may need to decrease this setting.

The disadvantage of a low setting is that a genuinely long call, which is active for longer than the time specified may not be tagged fully should it subsequently be transferred to another station.

Apply Beep Tone

This setting determines whether or not beep-tone is injected by the recorder during recording. It only applies to DMCC recordings as the recorder is only a passive observer in other recording modes (passive IP, TDM, and SIPREC). You should use the switch infrastructure's own beep tone injection features where these are available rather than this setting – as this will use an additional timeslot and will make the recorder visible as a conference party on the call.

Alternate Beep Tone Mechanism

As an alternative to having the recorder inject beep-tone using the above setting, you can enable beep-tone injection within DMCC. This does not require the additional timeslot. However, it is also always on – any time a recorder's DMCC port is conferenced into a call – whether or not the audio is actually being recorded. Using this method (and only this method), the beep-tone is recorded. To enable this mode, set the property **`dmcc.addbusyverify=true`** and restart the recorder.

Time between beeps (secs)

If Beep Tone is to be injected by the recorder, this setting determines the interval between beeps.

Optional Admin Pages Enabled

Use this setting to enable the additional administration pages that are needed if you want to use **On Demand** Recording, **Meeting** Recording, **TDM tap points** and/or **Phone Replay** on this switch. You must restart the recorder after changing this setting.

Audio Format

You can choose to stream audio from the switch in G.711 or G.729 format.

Avaya Communication Manager Name

Set this to the name of the AES switch connection you are using.

Maximum total call duration (hours)

To avoid having recording channels permanently active, the recorder will reset a channel that has been continuously recording for this many hours. Set this field to a value that you can be sure will not occur for a real call - typically just longer than any one person would ever be present on a shift. This setting only affects recordings which are under the direct control of the recorder. (Recordings may run over this duration by up to the maximum segment duration setting - on the **General Setup > Server** page -as they are only reassessed after the maximum segment duration expires.

AE Server Address(es)

This and the following two settings are required on each server that is providing DMCC based recording (as opposed to TDM, SIPREC and passive IP tapping). Enter the IP address of the Avaya Application Enablement Server on which the Device, Media and Call Control (DMCC) API is running.

See Appendix D: High Availability on page 375 for details of which AES to assign to which server.

To specify multiple addresses, which will be used in priority order (earliest in the list most preferred) add a semicolon before each subsequent address.

To specify an HA pair of which only one is active at a time, separate its two IP addresses with a forward slash. For example:

`192.168.130.1/192.170.120.24;192.100.100.1`

Note:

If you specify an HA pair, the recorder will raise a warning daily until both of the addresses specified have been connected at least once. This is to ensure that the addresses entered are both valid – before you need to use them in anger. During a maintenance window, force the AES onto its normally silent address so that the recorder connects to it. This alarm will then stop.

DMCC User Name

The user name that the recorder should use to log in to the Device, Media and Call Control API.

DMCC Password

The password that the recorder should use to log in to the Device, Media and Call Control API.

Media Stream Encryption

Use this setting to select the type of encryption (if any) for the audio between the recorder and VoIP resources. You can still record calls on which the other party's audio is encrypted with this set to **none**. If you set it to any of the other options, ensure that you have configured the codec set used by the recorder's softphones to support this type of encryption. **srtp-aescm128-hmac32-unauth** is only supported in recent AES Versions but is recommended over the older **aes** mode.

IP Station Security Code

All IP softphones that register with Communication Manager must provide a security code. The code you enter must match the code entered for all the stations that you created earlier on Communication Manager for the recorder to use. This field entry is masked for security purposes.

The following settings are only shown on Master and Standby servers as they are not required by Slave or Replay servers.

AES TSAPI Server(s)

Enter the IP address of the AE Server that is configured to provide TSAPI services to the recorder.

See Appendix D: High Availability on page 375 for details of which AES to assign to which server.

To specify multiple addresses, which will be used in priority order (earliest in the list most preferred) add a semicolon before each subsequent address.

To specify an HA pair of which only one is active at a time, separate its two IP addresses with a forward slash. For example:

192.168.130.1/192.170.120.24;192.100.100.1

Note:

If you specify an HA pair, the recorder will raise a warning daily until both of the addresses specified have been connected at least once. This is to ensure that the addresses entered are both valid – before you need to use them in anger. During a maintenance window, force the AES onto its normally silent address so that the recorder connects to it. This alarm will then stop.

AES TSAPI Switch Name(s)

Enter the name of the switch as configured in TSAPI. If you have specified multiple TSAPI servers and the name differs on each, enter the name that each server uses in order, separated by semi-colons.

AES TSAPI Service Login ID

Enter the login identifier that the recorder should use when accessing TSAPI.

AES TSAPI Service password

Enter the password that the recorder should use when accessing TSAPI.

Non-recorded Stations/IVR ports to Observe via TSAPI

If any of the calls that you wish to record are initially handled by stations that are not themselves recorded (such as a bank of IVR ports) the tagging of those calls may be incomplete. Some information is only available via TSAPI when the call is first answered. To avoid losing this information, you can specify one or more ranges of stations that the recorder will monitor via TSAPI and hence learn the additional details it needs to tag and record these accurately. You can enter contiguous ranges and individual stations. Separate these with semi-colons. For example: 1201-1209;1346;4000-4099

In the case of AACC, if there are any agent stations that are not configured in bulk recording for a particular reason, these stations should be entered here. This ensures that complex calls involving these agents can be accurately tracked by the recorder.

Similarly, stations of users who call the On Demand and/or Meeting ports should also be included here so as to ensure as much is known about their calls as possible.

Agent Skill Group(s) to Observe via TSAPI

TSAPI does not let the recorder observe AgentIDs directly. If you wish to record calls based on AgentID - or even to tag recordings with AgentIDs and names, you MUST configure this setting. Enter enough skill groups to ensure that each agent you wish to record is in at least one of these groups. Separate skill groups with a semi-colon. You can enter ranges of skills - but only if all values within the range are valid skill hunt groups.

Tip:

Consider creating one dummy skill hunt group and assign all agents to this skill. You then need only enter that one skill hunt group here. This also minimizes the number of TSAPI licenses you need.

VDN(s) to Observe via TSAPI

To ensure accurate tagging of recordings, you must enter all of the VDNs in use here. Separate VDNs with a semi-colon. You can enter ranges of VDNs - but only if all values within the range are valid VDNs.

Tip:

If you have ranges of VDNs within which some numbers are not used, consider creating these as VDNs to allow you to enter the whole range here.

Tag calls with which VDN?

Each recording can only be tagged with a single VDN. You can choose whether this is the first or last VDN that a call went through. As with the following setting, "first" and "last" are "as far as the recorder is aware" i.e. restricted to those VDNs you tell it to observe.

Add VDN number as additional "owner" of calls

Access to recordings is normally controlled according to the station or agent that is the subject of the recording. You may, however, choose to control access to recordings on the basis of which VDN the call was tagged with by setting this option to **Yes**.

Address of the Communication Manager

This setting and the following two are only required if you intend to control recording in Conferenced mode according to the Class of Restriction (CoR) in which stations are placed. The recorder needs these settings to use the SMS Web services.

Set this entry to the IP address or hostname of the Communication Manager. If using hostname, the AES must be able to resolve this to an IP address.

Username for Switch Administration

See immediately above for when this setting is required. If it is needed, enter the username the recorder should use when accessing the automated switch administration (SMS) services.

Password for Switch Administration

See above for when this setting is required. If it is needed, enter the password the recorder should use when accessing the automated switch administration (SMS) services.

Record with Passive IP Taps

Normally calls on Communication Manager will be recorded using SIPREC or DMCC conferencing if there is no physical TDM tap available. Set this option to force the recorder to prefer passive IP tapping instead. Ensure that you have sufficient NIC cards in the servers to tap all of the possible paths that calls may follow. Remove the IP stack from NIC cards that are to be used for passive tapping (unless using a single NIC card in a server for both passive tap and normal server interactions). If you change this setting you must restart the recorder.

Note:

Passive tapping is only supported on Communication Manager and ACR must be deployed on a Linux server.

Avaya Oceana™

If you use Avaya Oceana™ to route calls on this Communication Manager, specify the address of its reporting interface here. Enter its IP address or hostname followed by a colon and the port number[1] of this interface.

[1] Defaults to 443 for HTTPS but can be changed. Check configuration on Avaya Oceana™.

Interactions are over HTTPS (with TLSv1.2). Both the Avaya Oceana™ and Avaya Contact Recorder server must have server side certificates, which must both be signed by the same CA. This CA will typically be your in-house CA, or the CA built into Avaya System Manager. To create ACR's certificate:

1. Create a new certificate as detailed in Creating a new Certificate on page 404 – replacing **<*keystorename*>** with **aoctls.p12** and **<*alias*>** with **aoc**.
2. Have the certificate signed as described in Generating a Certificate Signing Request on page 405.
3. Import the signed certificate as detailed in Importing the CA's certificates on page 405.

The status of the link to Avaya Oceana™ will show on the **Status > Server** page of the master. If you have a **Hot, Full Standby**, it will also attempt to connect to the Avaya Oceana™ at this address. If you have other (**Warm** and/or **Partial**) **Standby** servers, they will only attempt to connect as and when they go active. You should confirm each such server is able to do so.

Tagging information is received from Avaya Oceana™ when a call is first answered. You must ensure that all stations taking such calls are being observed if you want this tagging to be noted. This happens automatically for stations that are being recorded. Any others – for example those that might transfer calls on to other, recorded stations - must be entered under the **General Setup > <*Communication Manager Name*> > Non-recorded Stations/IVR Ports to Observe** setting.

Each Avaya Oceana™ attribute used to assign a resource to a call will then be tagged as a user defined field on recordings.

If there are attributes that you do not wish to tag calls with or whose names clash with other fields, you can modify how these are stored as user defined fields by setting properties of the form.

aoc.attribute.*xxxx*=*yyyy*

where ***xxxx*** is the (lowercased) name[1] of the attribute as used in Avaya Oceana™ and ***yyyy*** is the name of the user defined field you want it to be stored as (instead of xxxx). Leave ***yyyy*** empty if you do not want the attribute to be stored as a user defined field.

For example, the properties:

aoc.attribute.language=

aoc.attribute.customertype=cust

result in the Avaya Oceana™ attribute “language” not being stored and the attribute “customertype” being stored as user defined field “cust”.

If using WFO, you will probably want to map at least some of the resulting user defined fields to WFO’s Custom Data Fields.

[1] The names of all attributes are forced to lowercase on receipt within the recorder.

If you are not using WFO but do need to restrict Bulk recording to or exclude recording on the basis of an Avaya Oceana™ attribute, you should set the property:

`aoc.serviceattribute=`*`xxxx`*

where ***xxxx*** is the (lowercased) name of the Avaya Oceana™ attribute that you want to add a Bulk Recording Rule on. See Recording Rules on page 172 for further details. If you require more complex recording rules (for example, based on multiple Avaya Oceana™ attributes), you should deploy WFO and configure appropriate Business Rules to control recording.

Note:

If you set this property, this AOC attribute then behaves exactly like an additional VDN would. It is therefore subject to the setting on the Communication Manager Data Source that determines whether the first or last one is stored. As these are interleaved with VDNs, neither of these two options is ideal. You are advised to remove the Service column from your replay layout and add a column showing the specific Oceana attribute instead – as this will be unaffected by the call's passage via genuine VDNs.

Extensions assigned to recorder(s)

To use DMCC recording, you must provide enough softphones to support the total recording load, plus one port per server (to be used as described under

Beacon Softphone on page 128).

To make subsequent administration as easy as possible, assign sufficient stations to handle the ultimate load on each server, even if you do not intend to allocate all of these ports immediately.

To add a port range

1. Click **Add Port(s)** at the bottom left.
2. Enter a range of station numbers.
3. If you have more than one ACR server, twenty softphones (by default) will be immediately assigned to each server for Bulk recording and additional ones added to maintain this headroom as the peak load on each server grows. Normally a single pool is provided and used across all recorders. In some topologies you may need to assign specific softphones to particular ACR servers. To do this, click the **Advanced** button and assign the range of softphones to a particular recorder by entering its serial number in the **Designated Pool(s)** field. In large, distributed systems using "follow the sun" loading you may not be able to provide sufficient softphones to allow all servers to be assigned enough for their peak load at the same time. In this case, you can have ACR recover spare softphones from servers that do not require them. Set property value **`dmcchandlers.maxfree=nn`** where ***nn*** is the maximum number of free softphones to be maintained on any site. This *must* be greater than the minimum headroom of 20 per server.

Note:

The above applies to **Bulk** recording ports only. **On Demand** and **Meeting** recording ports are always created on the active recorder (Master or Standby) and are unaffected by this setting.

Important

When configuring a fault tolerant system, you should provide enough softphones for the remaining servers to handle the full design load even if one server has failed in such a way that it retains the softphones that it has been allocated already. So in a 3 server, 1000 channel system, for example, provide enough softphones to allow two remaining servers to handle the full load even if a third of the ports are stuck on the dead server - i.e. 1000 + 1000/3 + 20 per server headroom = 1393 ports at least.

To edit a port range,

1. Click on the **Edit** link in the right-hand column.
2. Change lowest or highest port number or both.

To delete one or more port ranges:

1. Click the checkbox in the Select column for each station range you want to delete.
2. Click on the **Delete selected port(s)** link.

Once you have configured a range of softphones, use the tabs under **Operations** to assign ports to the various tasks that your license allows.

Note:

You can only delete a port range if none of its stations is assigned to any of the recording modes.

Important:

Set these port ranges up correctly before proceeding to the subsequent pages on which you allocate these ports to the various recording and replay tasks.

For a basic, small system, this is all that you need enter here. However, on larger systems you may want to click on the **Advanced** link to set the following:

- You can (although this is NOT recommended) force the ports to register through a specific C-LAN(s) rather than allowing AES to determine this automatically. Enter the IP address of the C-LAN they should use.

Important:

These settings will not take effect until you restart the recorder.

Once you have set any of these **Advanced** options, this setting will be shown on the list of ports.

CS1000

When recording a CS1000, configure these settings:

Switch Identifier Prefix

This setting is only present if you have configured the recorder (as described above) to support conflicting numbering schemes. It lets you specify a few characters that will be added to the front of each internal phone number that a recording is tagged with. You should only set this for switches whose numbers could be confused with those on other switches. Do not use commas or dashes within the string. Restart the recorder after setting this value.

Changing this setting after recordings have been made is not advised as search and replay rights depend on it. Note that you will need to explicitly assign search and replay rights to these extended addresses if you use this option.

Minutes after which call information indexed by Call ID is discarded

Each call is given a reference number or "Call ID" but these are reused after a period and hence any information gleaned by the recorder must be discarded before that call ID is reused. The default of three hours is appropriate for many switches but if you have, for example, an auto-dialer or very high calling rates you may need to decrease this setting. The disadvantage of a low setting is that a genuinely long call, which is active for longer than the time specified may not be tagged fully should it subsequently be transferred to another DN or Position ID.

Apply Beep Tone

This setting determines whether or not beep-tone is applied during recording. (Note that TDM recording does not support this feature).

Optional Admin Pages Enabled

Use this setting to enable the additional administration page that is needed if you want to use TDM tap points on this switch. You must restart the recorder after changing this setting.

Avaya Contact Center Manager Server Address(es)

Enter the IP address of the Contact Center Manager Server(s) that the recorder is to establish an MLS link to. There is no need to add a port number as MLS always uses a fixed port number. If you have a standby CCMS, this should be configured to use the same "managed" IP address as the master.

To find the correct IP address of the CCMS, login in to CCMA, choose **Config** and pick the CCMS server. You will find the IP address in **Properties**.

Meridian 1 Machine Name and Customer Number

If you share a multi-customer system, you must specify your Machine Name and Meridian 1 Customer Number. Otherwise, leave these fields at their default values (blank and 0 respectively).

Option 11

If your switch is an Option 11, you must indicate this here.

CC6 Mode

Avaya Contact Recorder assumes that Contact Center Manager Server 7.0 (or later), CS1000 release 6 (or later) and the appropriate number of MultiDN licenses on CCMS (appears as "Multiple DN Registration" on the CCMS License Manager Real Time Usage Screen) are present. If any of these three conditions are not present, then you must set **CC6 Mode** to **Yes**.

Avaya Aura® Contact Center Interface

If you are using an Avaya Aura® Contact Center on Communication Manager, you must enter the following information to allow the recorder to communicate with it.

Add domain name to identifiers

This setting is only present if you have configured the recorder (as described above) to support conflicting numbering schemes. If set to **Yes** the recorder will add the domain name (specified further down this page) to each address - so "1234" is stored as "1234*@mydomain.com*". You should only set this for switches whose numbers could be confused with those on other switches. Restart the recorder after setting this value.

Changing this setting after recordings have been made is not advised as search and replay rights depend on it. Note that you will need to explicitly assign search and replay rights to these extended addresses if you use this option.

Underlying Switch

Avaya Aura® Contact Center is only supported when hosted on a Communication Manager or CS1000 defined on another Contact Centre Interface tab. Use this setting to specify which switch this is. You must configure and name the underlying switch on its own administration page before selecting it here.

CCT Server

Enter the IP address of the server running CCT. The port number defaults to 9080 but can be changed by adding a colon followed by the required port number.

Note:

If using a Partial Standby on a remote site, configure it with the main CCT server first then the backup CCT server on that site (separated by a semi-colon). The Standby will therefore connect to the (normally live) main feed but fall back to the local connection when isolated – at which point this feed should be manually activated.

CCT Username

Enter the Windows username that the recorder should use when connecting to the server running CCT. This defaults to "CallRecordUser". Note that with AACC 6.2 or later, any normal username can be used.

CCT Password

Enter the Windows password the recorder should use when connecting to the server running CCT.

Windows Domain

Enter the Windows domain name your AACC uses, in which the above **CCT Username** and **Password** are configured.

AMS Zone to Recorder/Pool Mappings

Note:

This setting is only applicable when using SIP recording via the AACC AMS (as this release of ACR now offers – and, in fact, defaults to - the alternative approach of always recording AACC CDN calls via DMCC instead).

Avaya Aura® Contact Center 6.4 onwards supports AMS zoning. If you are using SIP recording and have configured AMS zones, you can minimize WAN traffic by instructing Avaya Contact Recorder to use a specific recorder or pool of recorders for each AMS zone.

To place a recorder in a pool (or set of pools), use the "Designated recorder/Pool(s)" advanced setting, in the **General Setup**>**Server** page.

To match an AMS zone to a recorder or pool of recorders, enter the AMS zone name followed by “@” and the serial number of a specific recorder, or the name of a pool of recorders. Separate successive entries with a semi-colon. For example:
dublinnorth@dublin;dublinsouth@dublin;luxembourg@amsterdam;
amsterdam@amsterdam;hague@890123

The above example instructs Avaya Contact Recorder to record calls on AMS servers in zones dublinnorth and dublinsouth on the dublin pool of recorders; AMS zones

luxembourg and amsterdam on the amsterdam pool of recorders and AMS zone hague on recorder 890123.

CDN(s)

List the CDNs used in your AACC - either individually or in ranges - so that the recorder can recognise these addresses as CDNs when they appear in CTI events. Separate entries with a semi-colon. For example: 1234;5500-5599

Proactive Contact (PCS/PDS Dialer)

To correctly separate and tag calls made via one or more PCS dialers, you must connect the recorder to each dialer so that it can receive details of agent logins and calls in progress.

Firewalls

If there are any firewalls between the PCS and the recorder, these must allow connection from the recorder to the dialer's name and event services on ports 23200 and 23120 (or if encrypted link is used, 23201 and 23121) on the PCS server.

When the PCS establishes a return connection, by default it uses a random port number on the recorder. Hence if a firewall is present you must specify a particular port number and open the firewall to that port on the recorder.

Limitations

1. As PCS only advises the recorder of an agent’s location when agents log in, you should restart the recorder out of hours or have agents log out and in again after the recorder starts.
2. Consult calls, Transfers and Conference calls made via the hard dialer's user interface are not reflected in the hold, transfer and conference count parameters for each call.

Dialer Name

You can choose what to name the dialer here - but this must match the name of the corresponding data source defined in WFO (if present).

Underlying Switch

Specify which switch this PCS is connected to.

Dedicated Headset Port Trunk Range(s)

If this is a "hard" dialer (as opposed to "soft") specify the trunk range(s) that dialer uses on the underlying switch for its "nailup" calls (i.e. not those used for transfers, consult calls etc.).

To determine this, look at the dialer’s configuration file, `/opt/avaya/pds/config/dgswitch.cfg`

Find the section listing the “Headset Ports” where lines start “H:” and ignore all other sections. (It is always worth checking that there are no out of place “H:” lines in the other sections though – as these will also need to be included – and should be moved within the file to their proper section!)

```
#Headset Ports
H:1:505:0::#H:15:1:1-1-10-4-10
H:2:506:0::#H:15:1:1-1-10-4-11
H:3:507:0::#H:15:1:1-1-10-4-12
H:4:508:0::#H:15:1:1-1-10-4-13
H:5:520:0::#H:15:1:1-1-11-1-1
H:6:521:0::#H:15:1:1-1-11-1-2
H:7:522:0::#H:15:1:1-1-11-1-3
H:8:523:0::#H:15:1:1-1-11-1-4
H:9:524:0::#H:15:1:1-1-11-1-5
H:10:525:0::#H:15:1:1-1-11-1-6
#Outbound Ports
...
```

Each line within this section defines a channel on a trunk. You must determine which Communication Manager Trunk Group and Member each of these corresponds to and enter ALL of them as instructed. In the above file, for example, if the channels corresponded to (all 4 members of) Trunk Group 17 (mapped to T1 1-10 channels 10-13) and members 1 thru 6 of Trunk Group 25 (which has other members too) then you would enter:

17;25:1-6

Leading Digits to be removed from dialer Agent ID

If your dialer is configured with superfluous leading digits, you can remove as many as required by setting this value. The resultant (trimmed) ID must match those configured in WFO (if present).

Hostname

The hostname of the dialer - and is case sensitive.

IP Port to Connect via

The port number the recorder will attempt to connect to on the PCS. If left at 0, the recorder will use the default port (which varies according to whether the connection is encrypted or not).

Username

Specify the username the recorder should connect as.

Password

Specify the password the recorder should use when connecting to this PCS.

Dialer Version

Specify whether the dialer is running PCS Version 4, 5, 5.1 or V5.1.1.

Note:

ALL PCS dialers connected to an ACR must be of the same version.

Use Encrypted Connection

Whether or not to use encryption on the connection to the PCS.

Name Service IOR (optional)

For a single dialer, leave this setting blank to have the recorder look up the Name Service via corbaloc.

However, if there is more than one dialer and they are configured as a POD, you must set this parameter. Use the name service on the master dialer to determine this value for each dialer in the POD.

IP Address on Recorder to use for Return Connection

If the recorder has multiple NIC cards or if the dialer cannot resolve the recorder's hostname, you must specify the IP address it should use to contact this recorder.

Otherwise, leave this field blank.

IP Port on Recorder for Return Connection

Set this if you need to have PCS use a specific port on the recorder. The recorder's firewall must permit connections from PCS on this port. If left at zero (the default) any port may be used. If using multiple PCS dialers, specify a different port for each one.

Use Dialer Agent's Name Instead of ID

If, and only if, dialer agent IDs overlap those on the underlying switch but do not map one to one, you can use this option to identify dialer agents by name rather than their ID. If used, the agent IDs in WFO (if present) must also be configured by name not ID.

Note:

If you select this option, the “owner” of each dialer call becomes that agent’s name rather than numeric ID. As there is no mechanism to allow a user to search/replay across a range of alpha-numeric strings (as there is with number ranges) you will have to explicitly add *each and every* agent name to user accounts that need to replay these calls from ACR (even administrator accounts).

Completion Code to be Stored as

Completion codes will be stored as a user defined field of this name. Set it blank to disable completion code storage.

User Defined Fields

In addition to the Completion Code, the recorder automatically tags recordings with the following information if provided by the dialer:

- Jobname
- Jobnumber
- Any "IDENT:" fields configured

Each of these is stored in a User Defined Field of the same name unless overridden by a property setting of the form:

dialer1name.field.originalfieldname=newfieldname

where

- *dialer1name* is the name of the dialer.
- *originalfieldname* is the name of the field provided by the dialer.
- *newfieldname* is the fieldname that you want the recorder to store this as. (To suppress the field altogether, leave *newfieldname* empty).

The fieldname will be used as an XML tag name so should only include a-z. You should not include spaces in the field name but if you do, these will be converted to underbars.

Maintenance Window

This dialer is normally configured to restart nightly - typically shortly after midnight. You can suppress the alarms that this would otherwise cause by setting the hours during which such a restart will be ignored - so long as the link recovers with 5 minutes (by default). You can override this time threshold if needed with property setting **`cti.restartmins=nn`** where ***nn*** is in minutes.

Proactive Outreach Manager (POM) Dialer

To correctly separate and tag calls made via one or more Proactive Outreach Managers, you must connect the recorder to each so that it can receive details of agent logins and calls in progress.

Version

POM 3.0.5 is currently supported.

Firewalls

If there are any firewalls between the POM and the recorder, these must allow connection from the recorder to the POM server (defaults to port 7999 and for TLS 1.2 connections, will require port 7998).

Property Settings

In addition to the settings on the administration pages, you should also set the following property values in the `acr.properties` file:

`sessions.ignored=`*`nnnn`*

where ***nnnn*** is the CLID shown on the nail-up calls to POM.

`mls.unknownmaybeinternal=true` `(for CS1000 systems only)`

If the ACR is using a secure link to POM via TLS 1.2, then add the following line to the acr.properties file:

`pom.secure=true`

Dialer Name

You can choose what to name the dialer here - but this must match the name of the corresponding data source defined in WFO (if present).

Underlying Switch

Specify which switch this POM is connected to.

Leading Digits to be removed from dialer Agent ID

If your dialer is configured with superfluous leading digits, you can remove as many as required by setting this value. The resultant (trimmed) ID must match those configured in WFO (if present).

Hostname

The hostname of the dialer.

IP Port to Connect via

The port number the recorder will attempt to connect to. On the POM server, port 7999 is the default port used for unsecure connections over TCP. Port 7998 is for secure connections over TLS 1.2.

Username

Specify the username the recorder should connect as. Typically this is a username which has administrator privileges on POM.

Password

Specify the password the recorder should use when connecting to this POM using the above account.

Zone

Which POM zone the recorder should monitor.

- If zones are not in use, leave this at "Default". Do not clear the field.
- For multi-zone operation set this to “all” (not case sensitive).

User Defined Fields

The recorder automatically tags recordings with any "params" fields present in the MediaInfoEvents sent by POM. Each of these is stored in a User Defined Field of the same name as the "param" unless overridden by a property setting of the form:

Each of these is stored in a User Defined Field of the same name unless overridden by a property setting of the form:

dialer1name*.field.*originalfieldname=newfieldname

where

- *dialer1name* is the name of the dialer.
- *originalfieldname* is the name of the field provided by the dialer.
- *newfieldname* is the fieldname that you want the recorder to store this as. (To suppress the field altogether, leave *newfieldname* empty).

The fieldname will be used as an XML tag name so should only include a-z. You should not include spaces in the field name but if you do, these will be converted to underbars.

TDM Tap Points

If you are recording any TDM trunks and/or phones, you must install the appropriate Ai-Logix cards into one or more Avaya Contact Recorders. These servers must be running Windows (not Linux). Use the supplied SmartView tools to confirm that your cards are connected and are receiving the appropriate audio from your trunks and/or phones. You must then use this administration page to specify which trunks or phones you have connected to the input channels on these cards. To enable the required administration page(s), check **TDM tap points** on the **Optional Admin Pages Enabled** setting for the appropriate switch(es) under **General Setup**. Enter all information into the Master Avaya Contact Recorder as follows.

Note:

As TDM tap points relate to physical connections, these details are normally only changed when wiring changes occur. You MUST make changes out of hours and restart the recorder after making any TDM configuration changes.

1. On the **Operations** tab, select the **TDM Tap Points** sub-tab for the switch being tapped.
2. Click the **Add Tap(s)** button to enter details of the phones and/or trunks to which the TDM recorders' ports are connected.
3. Enter the connections according to the table below:

	Extension-side Taps	Trunk-side Taps
Communication Manager	Enter the numbers of the stations connected to the TDM cards. If you are tapping a contiguous range of stations and you have connected these to a contiguous range of input channels, you can enter these as a single entry. If your connections are not sequential, enter each one individually.	Enter the trunk group number and how many members it has. You must connect the entire trunk group to a single contiguous block of channels within one recorder.
CS 1000	Enter the TNs of the phone sets connected to the TDM cards. If you are tapping a contiguous range of TNs and you have connected these to a contiguous range of input channels, you can enter these as a single entry. If your connections are not sequential, enter each one individually.	First determine whether the trunk is a T1, E1 or DPNSS trunk and if T1, whether 23 or 24 channels of audio are present. Then determine the loop number for each trunk. These are determined from the switch using LD21 on the switch console and entering LTM (List Trunk Map). Once these have been found, set the type of trunk and enter the loop number for each trunk being tapped.

Tip:

If extension side tapping, connect phone sets with consecutively numbered TNs/DNs/PositionIDs/Stations) to consecutively numbered media channels so that you can enter ranges of phones rather than have to specify each one separately.

1. Enter an optional comment to describe the device or trunk being tapped.
2. Select the serial number of the Avaya Contact Recorder that is tapping this device/trunk.
3. Each trunk or range of phones will be connected to the appropriate number of input channels on a TDM card. These number from 1 as shown on SmartView. Enter the lowest (first) channel number used here. The recorder will then associate the appropriate number of channels starting with this one to the trunk or phone(s) entered above.
4. Repeat steps 2 through 6 for each trunk and phone or group of consecutively numbered phones.

Note:

(For CS 1000 Systems) Even though you enter TNs in the TDM tap points page to tell Avaya Contact Recorder which physical phones or loops are being tapped, you still enter DNs when specifying what to record. The recorder uses CTI information to determine which TN is involved when a recordable DN is active, and enables the relevant port on the TDM recorder. The DNs are administered in **Operations > Bulk Recording** as described in Bulk Recording on page 168.

"Unable to Record" Alarms

When using trunk-side recording, the system may raise alarms warning you that it could not record specific calls. These occur when the recorder attempts to record a call that it believes includes an external party but none of the devices to which they are connected are recordable. The reasons for this, and the appropriate actions to take are listed below:

1. The call is actually internal. Set **Record internal calls** to **No** on the **Operations > Bulk Recording** page.
2. The call does go via an external trunk, but one that has not been tapped so cannot be recorded. If you have genuine external trunks that do not have a recorder connected to them, you should either connect a recorder or accept that calls over these trunks will not be recorded. You can suppress the alarms on a specific trunk as shown below.
3. (CS1000 only) The TN reported in the CTI message is actually a conference bridge. In four (or more) way conferences, where an external party is present the recorder needs to be told to ignore the loops which are actually conference bridge loops. (The call should be recorded on the external trunk loop and hence does not need to be recorded on the conference loop as well.)

You can suppress these alarms on a particular trunk by creating an appropriate entry in the properties file (see Properties File on page 265).

On a	Add this to the properties file	Where...
Communication Manager	**`trunkgroup.ignore.nnn=true`**	***nnn*** is the trunk group number to be ignored. Do not use leading zeroes.
CS 1000	**`loop.ignore.nnn=true`**	***nnn*** is the loop number to be ignored - from 1 to 255. Do not use leading zeroes.

(CS1000 only) Note that the recorder assumes that trunks will be digital when it converts from the "packed TN" provided by the switch into a "loop number". This algorithm differs for analog trunks and the recorder may report a trunk number that does not actually

exist. This is the number that the TN would represent if it were a digital trunk. You should examine the packed TN to determine what this really is. When setting the "loop.ignore" value in the properties file, use the loop number reported by the recorder.

Email Configuration

The Avaya Contact Recorder sends Emails as part of the Replay Authorization process and can also be configured to email support staff with details of alarms that occur. If you require either of these features, create an email account that the recorder can use to send emails. Click the **System › Email Server** tabs at the top of the Administration page and enter the following details:

SMTP Mail "From" Address

Set the name from which alarm email messages should originate, for example, *recorder1@alt.bigcorp.com*.

SMTP Mail Server

Enter the name of the SMTP mail server on which you have established the email account that the recorder will use to send email messages. If you leave this blank, the system will not send email messages when alarms occur - and you can then leave the remaining settings on this page blank. The system uses the standard SMTP port (25) unless overridden by the property setting **`smtp.port=`*`nnnn`***.

SMTP Username

Leave this blank if your SMTP server allows unauthenticated sending. If it requires authentication, set the username of the SMTP account here.

SMTP Password

Leave this blank if your SMTP server allows unauthenticated sending. If it requires authentication, set the password of the SMTP account here. The password is masked when entered in this field.

Send Alarm/event emails to

Specify the email address(es) to which alarm and event messages should sent. Separate multiple addresses with a semi-colon (;).

Note:

Confirm that you are receiving emails correctly after you make any changes to these settings.

System Monitoring

There are several ways to track the operation of the Avaya Contact Recording system:

- You can browse the status and alarms pages via the administration pages
- The system can proactively send emails to warn of problems
- You can analyze log files produced by the application and its Tomcat web servlet container
- You can have application logging information sent to a Syslog server
- You can interrogate the system and have it send notifications via SNMP

These options are described in detail below. You are strongly advised to set up one of the proactive mechanisms to ensure that any problems are brought to your attention in a timely fashion. This can be done via Email, SNMP or Syslog monitoring.

Via the Administration Pages

The status of your recording system is shown on several web pages beneath the **Status** link on the Administration interface. Here you can see:

- overall **System** Status (Master and Standby recorders only). This shows the status of all recording servers in your system and provides links to the administration interface of the other servers should you wish to investigate any issues that are highlighted there. Problems are highlighted in red or amber according to their severity.
- overall status summary for the **Server** and its interfaces to other components
- loading levels - both current and peak
- **CTI Monitors** and current call states
- the state of all recorder **Ports**

Having just configured the recorder's main settings, you should now review the status and address any problems highlighted on these pages or the **Alarms** page.

Alarms and events occurring within the system are stored in its local database. You can view them via the **Alarms** tab at the top of the Administration page. If this link is red, it indicates that there are one or more uncleared alarms. As you address the cause of an alarm, you should clear it by clicking the checkbox to the left of it and then clicking the **Clear Selected Event(s)** link.

Via Email

Use the **Send Alarm/event emails to** setting on the **System > Email Server** page to specify the email address(es) to which alarm and event messages should sent. Separate multiple addresses with a semi-colon (;). The email recipient can be a local system administrator, a manned help-desk and/or suppliers' support desks if you have a support agreement that includes this facility.

The system sends an email message each time an alarm occurs or is resolved. It also sends an email once per day as a "heartbeat" to let you know it is still operating. This email is sent overnight and also advises you of the available disk space after the log files have been purged. You should investigate any failure to receive the daily heartbeat message as it could indicate that the server has failed. On a Windows server, alarms are translated into the language that your server is configured for. To do the same on a Linux server, you must set a property in the properties file e.g. to have alarms sent in Spanish set **`server.locale=es`**

Note:

The system will batch emails for up to 10 minutes to avoid flooding users' inboxes. The default configuration is that all alarms and events are sent via E-mail. You can change this so that only alarms above a specified severity are sent by E-mail. Be aware that if you do so, you may be unaware of issues that are affecting your system. See **`"email.minalarmlevel"`** in Properties File on page 265.

Application Logs

The Avaya Contact Recorder writes log files to the /logs directory beneath its install path (typically `/opt/witness` on Linux, `D:\Avaya\ACR152` on Windows). The current day's log file is called **`acr.log`**. At midnight the current log file is closed, renamed to **`acr.log`**.***`<date>`***, and a new log file opened. These log files are automatically purged after 30 days by default but this can be overriden with a properties file entry.

Setting the server log level

You can set the level of messages logged by the Avaya Contact Recorder to **DEBUG**, **INFO** (the default), **WARN**, **ERROR** or **FATAL** in one of two ways. Note that the logging levels of some components within the Avaya Contact Recorder are controlled by other settings - which are not to be altered by end users.

Permanently from next restart

To change the level permanently, enter the following line in the properties file **`acr.properties`** and restart the Avaya Contact Recorder service to have it take effect.
`log.level=DEBUG`

If you change the log level in the properties file, it remains set on subsequent restarts. When logging at DEBUG level, note that the log files grow very quickly and can overflow

the disk if left at this level. You may therefore also wish to change the number of days for which log files are retained. Set this with acr.logkeepdays=nn in the properties file.

Temporarily, immediately

To change the level temporarily, without restarting the service, simply use a browser to request the URL:

`http://`*`myrecorder`*`:8080/log?level=DEBUG`

using the name of your server in place of *myrecorder*.

Alternatively, on Linux systems if you do not have access to the web administration screens, but do have access via secure shell, execute the following command:
`perl /opt/witness/bin/loglevel.pl DEBUG`

You do not have to stop recording in order to change the logging level. To set it back again, enter the same URL, replacing **`DEBUG`** with **`INFO.`** The command is case-sensitive. Using this method changes the log level temporarily. It will revert to normal the next time that the system is rebooted.

Tomcat Logs

Avaya Contact Recorder uses the Tomcat web servlet container, which writes log files to **`/logs`** beneath the install path (which is **`/opt/witness`** on Linux, **`D:\Avaya\ACR152`** on Windows).

Remote logging via Syslog Server(s)

A subset of the information sent to the application logs can also be forwarded to one or more Syslog servers. To configure this, use the maintenance page at

`/servlet/acr?cmd=mtce`

There you can

- enter the IP hostname or address of one or more syslog servers
- filter the events to only those at or more severe than INFO, WARN or ERROR level

The recorder announces these events as coming from facility "LOCAL1" but this can be overridden using the property file entry **`syslog.facility`**.

SNMP

You can use an SNMP monitoring system such as HP OpenView to monitor Avaya Contact Recorders. To do so, you must first set the name of the **SNMP Read Community** in the **General Setup › Server** page. The recorder will then respond to SNMPV1 Get messages using Version 1, 2c or 3 as configured with the **SNMP Version** setting on this page. Version 3 is recommended but may require additional settings to be made via the properties file as shown in the table below.

If you change any of these settings, you should restart the recorder.

For security reasons, recorders:

- do not allow "well known" community names like "private" or "public"
- do not respond to SNMP Gets until a community name has been set
- do not use the usual port of 161, but instead use 2161

Property	Default	Description
snmp.port	2161	The port number to use for SNMP.
snmp.authtype	SHA	Set to "MD5" to override the default SHA authorization type with MD5.
snmp.privtype	AES128	Set to "DES", "3DES" or "AES256" to change the privacy type from AES128.
snmp.mainusername	acrsnmpuser	The main username the Network Management System will use to connect to the recorder.
snmp.username.*nn*	no default value	Additional usernames can be entered. Replace *nn* with 1, 2, 3 etc.
snmp.password.*nn*	no default value	Encrypted password to be used in conjunction with the corresponding snmp.username.*nn* entry.

SNMP Traps

Traps will be sent to the address(es) you configure as **SNMP Notification Destination(s)** on the **General Setup > Server** page. To specify multiple destinations, separate these with a semi-colon. SNMP traps are sent using the **SNMP Read Community** name entered on this page.

Operations

After initial setup, configuration of the recorder on a day-to-day basis largely consists of changing what is being recorded. You should also regularly check the status of any archive locations that you have configured here.

Beneath this tab therefore are administration pages that let you manage Archiving and each of the different recording and replay functions that the recorder provides.

Common Settings

The following settings are used on more than one **Operations** page and have the same meaning on each. This does not imply that all of these settings are applicable to ALL modes.

Apply Beep Tone within recorder

In some recording modes, you can specify whether or not the recorder injects or controls beep tone. For a full discussion, see Beep Tone on page 50 and refer to the configuration details for your particular Data Source type.

Stop recording if the call drops to just one other party

In some modes, users call a port on the recorder in order to record a call. Because the recorder port is then itself a party to the call being recorded, the call will stay active if any one of the other parties fails to hang up.

To avoid the port getting stuck in this way, you can set the recorder to hang up if other parties on the call hang up - leaving only the recorder and one other party left on the call.

Ports Configured

This figure shows how many ports you have allocated to this recording mode - as detailed in the table at the bottom of the page.

Designated Pool(s)

Use this field to force recordings to be made in or more **Designated Pools.** This feature is used both to force recording to take place on specific site(s) if possible but also allows a **Partial Standby** recorder to provide backup capability for these recording targets if isolated from the rest of the system. See Recorder Pool on page 126 for how to determine which **Recorder Pool** a particular server is in and Appendix D: High Availability on page 375 for a full discussion of **Partial Standby** servers.

In summary, if you specify a **Designated Pool**, and there is an AES on the remote site and the system is licensed for backup channels, you can configure one or more recorders there as **Partial Standby** servers rather than slaves. When needed, one of these will then record the targets for which its **Recorder Pool** is one of the **Designated Pool(s)**.

If you list more than one **Recorder Pool** (separated by semi-colons) or multiple rules triggering on a call cause more than one **Recorder Pool** to be listed, the recording will happen on the least loaded recorder in the highest priority pool that is reachable and has available capacity.

To explicitly prioritize between **Recorder Pools** add `/1` to the most preferred pool, `/2` to the next most etc. e.g. designate `atlanta/1` on one station's rule and `newyork/2` on another. This ensures that on a call between these two stations, the `atlanta` rule will prevail if possible.

This mechanism also controls failover behavior within a single recording target's settings so that if the top (lowest) priority pool or recorder is not available, the next one specified will be used.

Warn when free port count falls BELOW

Where a pool of ports is used as a shared resource, this setting triggers an alarm warning that a pool is running short of available capacity. The warning message identifies the pool of ports using the **Comment** field (or the port number range if no comment was entered). The default for this setting is zero - and as a pool can never have less than zero ports available, the warning is effectively turned off by this setting. If you set it to 1 the recorder will alarm when the free count falls below 1 - that is to zero. And so forth.

Recording owner(s)

See Access Rights on page 199 for details of how a recording's "owner" determines which user(s) can replay it. You can use this setting to override the default ownership (phone number or agent number).

Note:

This sets the owner of a recording or “segment” and hence controls who can replay this segment and search for it in a segment-based search/replay layout. In a session-based layout, users can search for and see the details of sessions that contain at least one recording that they are entitled to play – but will still only be able to play the specific recordings that they are entitled to. The presence of other recordings will be shown but these will not be playable.

For example, if you are configuring a pool of ports for use by the human resources department, you might find it more convenient to have all recordings made on them owned by "hr" rather than several different phone numbers, as would be the case if you left the default ownership in place. You can then control access to all of these recordings by assigning replay rights over "hr" to those users entitled to play them.

You can enter any alphanumeric string (with the exception of the semi-colon and dash (“-”) characters) or a number.

Leading zeroes, if entered, are significant and are retained.

All characters are stored in lowercase, so replay rights are applied case insensitively.

To assign multiple owners, separate successive owners with a semi-colon.

Assigning Ports

Each of the sub-tabs beneath **Operations** (except the **Archive** one) lets you assign ports to any of the uses that are appropriate given the type of Avaya switches you are recording, and the server and channels licenses you have installed. You allocate these ports to specific uses in two ways:

- **Explicit Allocation** - where you specify exactly which ports should be used. This approach is used where users dial in to or conference in to the ports in order to use them. An example is a pool of On Demand recording ports. The station numbers assigned to this pool must match those you have placed in the hunt group that you are going to tell your users they should dial for a recorder port.
- **Automatic Allocation** - where you specify what you want the ports to record. This approach is used for modes where users do not need to dial into ports and hence it does not matter which ports the recorder selects for a given task. An example is Bulk recording of a particular station. You tell the recorder which station you want to record but don't care which of the recorder's ports is used.

Read the text in the pop-up dialogs carefully to see whether you are being asked to specify recorder ports or targets to be recorded.

Ranges

When assigning ports or specifying which stations or addresses are to be recorded, you can enter single numbers or ranges of numbers. For example, entering "4000-4099" is much easier to enter than 100 different numbers. Typically, you need to enter some contiguous ranges and some individual port numbers.

A set of recorder ports is referred to as a "pool" and is configured on the appropriate page of the administration application by entering one or more port ranges. As you enter port numbers, keep the following in mind:

- A range can be a single number or a set of contiguously numbered ports or phone numbers.
- The lowest and highest port numbers in a range must have the same number of digits. Therefore, if you have some 4 and some 5-digit ports, you must enter these as two separate ranges.
- It is possible to have port or phone numbers that have one or more leading zeros. It is therefore important that you enter any leading zeros. The start and end ports in any port range must have the same number of digits, even if some of these are leading zeros.

Entering a range

Enter a range of ports as follows (phone and address ranges are entered in the same way):

1. Click the **Add port(s)/address(es)** button. The **Station Range** dialog is displayed.
   - To enter a single port, type the number in the top text box.
   - To enter a range, type the number of the first port in the range into the top text box and the final port in the range into the second text box.
2. Add a **Comment** (optional).

   For pooled modes, you can use this field to name a range of ports and to note any hunt group number that you assigned to these ports. The text you enter appears in status reports and warning messages as labels for the specified range. For more information about pooled modes, see Using pooled port modes on page 168.
3. To quit without entering the port numbers, click **Close Window**.

   To enter the port numbers and keep the window open to specify additional port numbers, click **Enter and Stay Open**. Your previous settings are retained. Change those that differ for the next range and repeat as necessary.

To enter the port numbers and close the window, click **Enter and Close**.

Port ranges are checked for consistency with other ranges and with license conditions as they are entered. If an entry is invalid, a message indicating the error is displayed; you can change the information as necessary.

Editing ranges

You can either change or delete the ranges listed. You can change the following fields without interrupting any recordings that might be active on the port(s) affected:

- **Comment**
- **Prompt User in...**
- **Warning Level**
- **Recording Owner**
- **Recording Rules**

Some changes require that the recording channel be reset and hence active recordings will be truncated. These are:

- **C-LAN address**
- **Designated Recorder/Poolname(s)**
- Any change to the number of ports in the range
- Any change to the port numbers in the range
- A change to the Codec (will only take effect on next restart) To edit a range:
    1. Click the **Edit** link to the right of the range you want to alter.
    2. Edit the range in the port entry form.
    3. Click **Enter**.

Deleting ranges

Delete one or more ranges as follows:

1. Click the checkbox in the **Select** column for each range that you want to delete.
2. Click **Delete selected port(s)/address(es)**. In some recording modes, this will truncate any recording in progress on the port(s).

Advanced settings

To implement advanced settings on a range of ports, phone numbers or addresses:

1. Add or edit a range as above.
2. Click the **Advanced** button.
3. Enter data in the **Advanced** fields.
4. To quit without saving the settings, click **Close Window**.
5. To enter the port numbers and keep the window open to specify additional port numbers, click **Enter and Stay Open**.
6. To enter the port numbers and close the window, click **Enter and Close**.

Using pooled port modes

If you use On Demand Recording, Meeting Recording or Phone Replay, you can manage each range of ports dedicated to a recording mode separately. For example, if you assign two ranges of ports to Meeting Recording, you can use one for English speakers and one for French speakers. You can also track usage of individual pools. You could, for example, set up two different On Demand pools, one for a particular department on one hunt group and one for everyone else on another hunt group. You can track the status of these pools through the **Status** pages. For each port range assigned to a mode, you can view activity for that port range in the **Status > System** page.

Bulk Recording

This tab is only visible if your license includes bulk recording channels. You configure this screen on the Master recorder only. The top part of the screen lets you choose how bulk recording works for all ports but you can override most of these settings for specific recording targets - using the **Advanced** settings of the phone number ranges which are shown in the bottom section of the screen. The type of address that you can enter as a "recording target" in the bottom section of the screen depends on the switch type you are using as follows:

	Recording Targets Can Be	**Notes**
Communication Manager (To switch between CoR, Style and other target types you must first delete all recording targets then	Station	Ideally, choose one type of recording target rather than mix these as the results of the latter can be complex and confusing. The recorder assumes that an address never changes type. Address names are refreshed nightly but if you change the type of device associated with an address you must restart the recorder for this to be noticed.
	Agent	
	Split	
	VDN	

change the **Specify recording targets** setting.)	CoR	Requires SMS services as described under the Communication Manager section within Data Sources on page 135. If you change recording targets, click the **Refresh CoR membership** button (otherwise this only happens overnight as it may take several minutes). Check the **Alarms** page after 10 minutes to confirm this completed successfully. Ensure each **Full Standby** is also successful.
	Recording Style	Use the recorder's administration page to define the **Advanced** settings for each recording "style" but use Avaya Contact Center Control Manager to specify which phones are to be recorded in a given style.
CS1000	DN	Single occurrence on a single phone set.
	PositionID	NB. Not AgentID. (And cannot be used with AACC agents as they log on to DNs not positions).
	IDN	Single occurrence on a position.
	MARP	Only on Knowledge Worker sets. On CC6, enable MARP/MADN for specific recording targets using **Advanced** setting. Always on for CC7 and higher.
AACC (only)	Agent	Only applicable if neither Communication Manager nor CS1000 is connected.

The advanced settings, for the recording mode as a whole and/or specific recording targets are:

Designated Recorder/Pool(s) (specific recording targets only)

If set to a recorder's 6-digit identifier or pool name, this will determine which recorder or pool of recorders should record this/these addresses. Leave blank for automatic recorder allocation. To specify a list of recorders or pools to be used in descending priority order, separate recorder numbers or pool names with semi-colons. (Note that TDM channels are physically connected to specific recorders and must, therefore be recorded on those servers. This setting only applies to IP recordings.)

Screen to Record (specific recording targets only)

To record screen content at a particular Windows workstation whenever audio is recorded on a Communication Manager station, CS1000 DN or Position or (if running AACC without CM or CS1000) AACC Agent, enter the IP address or host name of the workstation. Where a range of phones is used, enter the screen name for each in order, separated by a semi-colon. This cannot be used when targeting Communication Manager Agents, Splits or VDNs as it indicates a physical relationship between a telephone and a screen.

If, on the other hand, you wish to record the screen content for particular agents, you can use **System > Manage Users** to associate an AgentID with a Windows domain account. If you do this, the screen of that agent will be recorded along with the audio whenever they are logged in to a screen that has been configured with the "Capture Service". This approach works not only for physical workstations but also for a range of thin client desktop topologies.

Note:

When using WFO Business Rules to drive screen recordings, the Agent/Employee/User Account mappings defined within WFO are used to determine which screen should be recorded - rather than this setting.

Recording owner(s)

For Bulk recordings, the “normal” ownership is determined by the address being targeted for recording. Where an Agent, Skill or VDN is targeted, that address becomes the default owner of the recording. Where a station is recorded[1] or a recording is the result of a WFO Business Rule then the default owner is the agent logged on at the time (if any) or the station if no agent was present.

As described in Recording owner on page 164, this setting will override the default owner of any recording made in this mode.

Normally any entry made in this field, will replace this default owner. To change this (for Bulk recording only) and have it add additional owner(s) instead, add property setting **`owners.additive=true`**.

Note:

See Access Rights on page 199 for full details of recording ownership.

Apply Beep Tone

You can override the default beep tone set for the switch or (for recording modes other than bulk/business rule recording) for the recording mode as a whole. Beep tone can be on, off or on only when the recording is being made and will be retained at the end.

[1] Whether explicitly or as the result of being in a recorded Class of Restriction (CoR).

Record Internal Calls

Whether or not to record internals calls. You must set this to **No** if you only have trunk-side recording taps.

Internal calls that meet all other recording requirements will be recorded if any station or agent on the call is configured to record internal calls (either explicitly on the target's Advanced settings or, in the absence of an override there, falling back to the Bulk record mode's overall setting).

Note:

Do not attempt to override the default setting on specific Skill or VDN recording targets – only stations and/or agents are supported.

Percentage of calls to record

If your system has been licensed for this optional feature, you can choose what percentage of calls that meet the criteria for bulk recording are actually recorded. The decision to record or not is taken when a call first becomes recordable and will persist for subsequent segments of that call and any related calls i.e. those to which people connected to the original call are also connected. You can set this percentage for the recording mode as a whole. You can also override this default by setting it as an **Advanced** property for one or more recording target ranges.

Percentage of screens to record

If your system has been licensed for this optional feature, you can choose what percentage of calls that are recorded also have their screen content recorded (where possible). You can set this percentage for the recording mode as a whole. You can also override this default by setting it as an **Advanced** property for one or more recording target ranges.

Note:

If users' Windows accounts are defined in WFO and a specific Business Rule has not triggered, ACR will attempt to record 100% of screens for these users as it infers that WFO will be interested in the screen recordings. To stop this, and allow the percentage set here to take effect, set the property **`screen.bulkdominates=true`**.

Recording Control

By default, recording occurs whenever a call is connected to the phone number(s) specified and cannot be influenced by external controls such as Unify, instructions from the phone or desktop applications. Using this setting you can change this behavior:

- To **Trigger on alerting** - i.e. to record even if the phone only rang and was subsequently answered by another party.

- To not **Start recording automatically at start of call** - i.e. to wait for a manual or external start signal.
- To **Follow the call** - i.e. continue recording even if the phone number that triggered recording is no longer on the call. Note that this only supported on Communication Manager (and then only when using DMCC recording) and requires the recorder to keep a recording port connected to the call at all times. This may restrict calls to 4, rather than 5 way conferences as there may, briefly, be two recording ports present on a call as a consult call merges into a conference call. Also note that the scope of this capability varies from switch to switch and is also limited where calls are controlled by external dialers.
- to **Allow user/external start/restart** - i.e. to act on START commands from manual or external control applications.
- to **Allow user/external stop** - i.e. to act on STOP commands from manual or external control applications.
- to **Allow user/external delete** - i.e. to act on DELETE commands from manual or external control applications. Not supported for TDM recording.
- to **Retain ONLY if requested by user/external** - i.e. to delete the recording unless a RETAIN command has been received. Not supported for TDM recording.

Recording Rules

These settings let you control whether or not to record on the basis of a particular piece of CTI information. The fields available vary according to the switch type as shown below.

Switch Type	Fields	Notes
Any	DNIS	
Communication Manager	Agent	Must be observing an appropriate set of skills (as configured on the **General Setup** page for this switch).
	Split	Must be observing an appropriate set of skills (as configured on the **General Setup** page for this switch).
	VDN	Must be observing an appropriate set of VDNs (as configured on the **General Setup** page for this switch).
	Avaya Oceana™ Attribute	Use property setting **`aoc.serviceattribute`** to specify which Avaya Oceana™ attribute is to be used for bulk recording decisions.

CS1000	CDN	
	Activity Code	Set during call so call has to be recorded and then deleted at the end if not wanted. Manual entry is prone to error. An activity code entered by one user does not always affect the call once it has been transferred to another user. **LIMITATION:** If screen recording is enabled, the screen recording will be retained even if the audio is not retained. Use with caution.
	Agent	In CC6 mode, agent location cannot be determined at recorder startup - only when the agent next logs in.
	Skillset	Alphanumeric name
AACC	Agent	
	CDN	
	Skillset	Alphanumeric name

For each filter setting, you can set one of the following options:

- Record only those where the field has specific values.
- Record only those where the field does NOT have specific values.
- (only when overriding a rule on the recording mode as a whole) to ignore this field.

You can also choose:

- whether to record or not if this attribute is blank on a call.
- whether the rule applies to the first, last, current or any values of this attribute on the call. Note that some - such as split/skillset and VDN - are never actually connected to the call while it is active, hence "current" is not an option.

You can add rules to bulk recording as a whole or to specific recording targets using the **Add Rule** buttons at the bottom of the **Bulk Recording** and **Advanced** settings pages respectively. A rule assigned to a specific set of recording targets will override an overall rule if and only if the two rules apply to the same type of address.

To delete a rule, click its **Edit** link and then click the **Delete Rule** button.

Each recording is only tagged with a single split/skillset or VDN - though as a call is handled and transferred it may actually have passed through several of these. You can choose whether the recorded segment is tagged with the "first" (earliest) or "last" (latest) VDN but in the case of splits/skillsets it is always the latest that is tagged. Because recording rules on these fields can be applied to the "first", "last" or "any" of the

splits/skillsets or VDNs on the call, the resulting recordings will not necessarily be tagged with the split/skillset or VDN that triggered the recording.

Note:

Where an overlay system such as AACC or a dialer is controlling calls the recording rule only checks the call connections known to that overlay system. For example, a rule specifying VDNs will not see any VDNs on dialer calls even if an underlying consultation call is routed via one. For more complex combinations of rules, use WFO's Business Rules.

Specify Recording targets (Communication Manager only)

On Communication Manager, you can choose to record specific Stations, Agents, Splits and VDNs or you can choose to specify one or more Classes of Restriction (CoR). In the latter case, all stations in that CoR will be recorded. You can also choose to administer recording targets via Avaya Contact Center Control Manager.

Delete Recording by entering (Communication Manager only)

Use this setting to specify a digit string that a user can dial during a recorded call to instruct the recorder to delete a recording. This will only work if you have also set **Allow user/external delete** under the **Recording Control** setting. You must restart the recorder after changing this setting to allow it to take effect.

Retain Recording by entering (Communication Manager only)

Use this setting to specify a digit string that a user can dial during a recorded call to instruct the recorder to retain a recording. This will only work if you have also set **Retain ONLY if requested by user/external** under the **Recording Control** setting.

Once this command has been entered, the current and subsequent segments of the same call will be retained (after hold and retrieve for example). Any recording segment already completed (e.g. before a previous hold) will not be retained. Other calls taken while that call is still active (such as consultation calls) will not be retained unless you also enter the retain code during those calls.

This mechanism should not be used with screen recording. Use AIM control instead when screen recording is installed. You must restart the recorder after changing this setting to allow it to take effect.

Block IP recording (force TDM)

If there is a TDM tap onto an extension being recorded, that will be used. If not, the recorder will try to use IP recording if possible and only use TDM trunkside as a last resort. Stop it using IP recording using this setting.

Save/Delete Key Present (specific recording targets on CS1000 CC7 and higher only)

If you have configured a Save/Delete key on a CS1000 phone, check this setting to ensure the recorder updates the lamp on the button appropriately.

Allow MARP/MADN (specific recording targets on CS1000 CC6 only)

By default, CC6 DNs are assumed to be single line appearances. If multiple appearances are used, check this setting.

Number of Addresses Targeted

This is not set directly, but summarizes how many different recording targets have been configured in the table at the foot of the page. Note that some targeted addresses (e.g. bridged lines, Splits, MADNs) may result in multiple calls being in progress at the same time so the maximum number of concurrent recordings that may be attempted could be higher.

Delayed Call Deletion (Communication Manager only)

Optionally, you can configure the recorder to wait before deleting a recording made in this mode. If a user decides to retain a call after it has completed, he can do so by ringing a specific number (the "Retain Port") on the recorder.

Note:

This feature requires one additional DMCC port on the recorder. Delayed Call Deletion is incompatible with the "Block" command which requires deletions to act immediately. If using screen recording, Agent Initiated Monitoring (AIM) is also available and should be used in preference to these basic delete/retain capabilities. Screen recordings will NOT be deleted even if the associated voice recording segment is deleted.

To implement this configuration, you must add the following line to the properties file on ALL recorders (Master, all Standby and Slaves):

`execmode.deletedelaymins=`*`NN`*

Where ***NN*** is a number of minutes from the end of the call within which the user can call the "retain" number to retain the call.

And this property on the Master and all Standby server(s):

`execmode.retainnumber=`*`NNNNN`*

Where ***NNNNN*** is the station number of an otherwise unallocated port on the recorder that will be dedicated to receiving Retain commands after hang up.

If you set these properties, call segments recorded in Bulk mode with the **Retain ONLY if requested by user/external** setting enabled are not deleted until the time specified has elapsed after the end of that recording segment. A retain command entered during

or after the recording ends (within the specified period) will preserve the calls. When determining an appropriate value for this delay, consider the following:

- You want to maximize the chance of retaining a call that the Station Executive user chooses to Retain, but
- Recordings cannot be uploaded into the WFO database until this long after the whole contact ends (even if you only retained an early segment of it).
- You should increase the delay on the **Operations > I/O Jobs** page to exceed this delay so that retained recordings are archived shortly after this period.
- If using After Call Work on WFO or via property setting `acw.default`, you must set this delay to be greater than the longest After Call Work period used.

The **Retain** command applies only to the most recent call on a station. So, if a call is placed on hold and a consultation call is made, this consultation call is now the most recent.

Calling the retain number when the previous call is still on hold results in the consultation call being retained. However, if the user resumes the held call, hangs up, then dials the retain number, the original (and final) call is retained. In this case, the segment before the call was placed on hold is retained, as long as it ended less than **NN** minutes before the retain command was given.

On Demand Recording (Communication Manager only)

This type of recording uses a pool of ports on the recorder that you can access via one or more hunt groups. To use this recording mode, first enable the required administration page(s) by checking **On Demand Recording** on the **Optional Admin Pages Enabled** setting for the appropriate switch(es) under **General Setup**. The configuration page will appear when you restart the recorder.

Note:

Calls recorded in this way are typically already in progress before the recorder's port joins them. If the station(s) and/or VDNs from which the call originates are not being observed by the recorder, it cannot always identify all parties on the call nor determine which call a consultation call is associated with.

On Demand calls are not segmented, and therefore an absent party can be observed within a replay result with no waveform.

To observe a station that uses this recording mode but is not itself being recorded, either add it to the Non-recorded Stations/IVR ports to Observe as described on page 140 or include it in your WFO configuration.

If this recording mode is also accessible to external calls that arrive via VDNs, add those VDNs to the VDN(s) to Observe setting as described on page 140.

Mode Setup

At the top of the page are the following settings:

- **Recording Owner(s)**
- **Apply Beep Tone within recorder**
- **Warn when available channel count falls below**
- **Stop recording if the call drops to just one other party**
- **Ports Configured**

All of these are explained in Common Settings starting on page 163.

Ports Assigned

The table at the bottom of the page lets you assign specific recorder ports to this recording mode. Refer to Assigning Ports on page 165.These are created on the active recorder (Master or Standby).

In most cases, you will want to place each of these ports in a hunt group and make your users aware of this hunt group number and/or set up a key on their phones to access it.

Note:

If you assign any ports that have already been created as part of a Phone Replay pool, you will have to restart the recorder to allow them to be reassigned. Use the **Status > Ports** page to check this.

Advanced Settings

Click on the **Advanced** link when entering or editing a port range to set:

- **Recording owner(s)**
- **Warn when free port count falls below**

All of these are described under Common Settings starting on page 163.

Audix

To allow a user easy access to On Demand Recording from a station, combine On Demand Recording with the "One-Step Recording via Audix" Communication Manager feature. To do this, enter the hunt group used for some On Demand ports as the parameter for the **Audix-rec** feature button you assign to the user's station.

Note:

Do not configure an Audix-rec button on a station that is to be Bulk recorded.

When using the **Audix-rec** feature, keep in mind:

- If beep tone set to yes on ACR, then only the user pressing the button will hear this tone and not the other party on the call. For this reason it is recommended to use the CM beep tone in environments where Audix-rec and beep tone are required.
- If you are using the specific tone options on CM for Audix-rec, both parties on the call will hear this tone.
- Calls recorded using this feature are only indexed with the party that pressed the button, not the other party on the call.

For more information, refer to the AUDIX One-Step Recording Chapter in the Avaya Aura® Communication Manager Feature Description and Implementation guide.

Meeting Recording (Communication Manager only)

This type of recording uses a pool of ports on the recorder that you can access via one or more hunt groups. To use this recording mode, first enable the required administration page(s) by checking **Meeting Recording** on the **Optional Admin Pages Enabled** setting for the appropriate switch(es) under **General Setup**. The configuration page will appear when you restart the recorder.

Note:

Calls recorded in this way are typically already in progress before the recorder's port joins them. If the station(s) and/or VDNs from which the call originates are not being observed by the recorder, it cannot always identify all parties on the call nor determine which call a consultation call is associated with.

Meeting calls are not segmented, and therefore an absent party can be observed within a replay result with no waveform.

To observe a station that uses this recording mode but is not itself being recorded, either add it to the Non-recorded Stations/IVR ports to Observe as described on page 140 or include it in your WFO configuration.

If this recording mode is also accessible to external calls that arrive via VDNs, add those VDNs to the VDN(s) to Observe setting as described on page 140.

Mode Setup

At the top of the page are the following settings:

- **Recording Owner(s)**
- **Apply Beep Tone within recorder**
- **Warn when available channel count falls BELOW**
- **Stop recording if the call drops to just one other party**

- **Prompt Users in** - sets the default language for spoken prompts
- **Ports Configured**

All of these except for **Prompt Users in** are explained in Common Settings on page 163.

Note that all Meeting recordings are performed in G.711 to make the production of custom voice prompts simpler.

Ports Assigned

The table at the bottom of the page lets you assign specific recorder ports to this recording mode. Refer to Assigning Ports on page 165. These are created on the active recorder (Master or Standby).

In most cases, you will want to place each of these ports in a hunt group and make your users aware of this hunt group number and/or set up a key on their phones to access it.

Note:

If you assign any ports that have already been created as part of a Phone Replay pool, you will have to restart the recorder to allow them to be reassigned. Use the **Status > Ports** page to check this.

Advanced Port Settings

Click on the **Advanced** link when entering or editing a port range to set:

- **Recording owner(s).**

Tip:

On Meeting Recording ports, a user can enter a list of owners manually using the dial pad.

- **Warn when free port count falls BELOW**

Both of the above are described under Common Settings starting on page 163.

- **Prompt users in** - Sets the language for spoken prompts. If using multiple languages, you should configure a hunt group to correspond with each range of ports that uses a different language.

Custom Prompts

If the language you require is not offered, select **Customer defined prompts** and provide your own set of prompts. These files are located in **`/wav`** beneath the folder into which you installed the recorder and are called:

- `welcome_custom.wav`
- `owners_custom.wav`
- `recording_custom.wav`

- `help_custom.wav`

You can replace these four files with your own recordings. Listen to the corresponding file in a language you understand (for example, **`welcome_english.wav`**) and record the equivalent messages in your own language.

WARNING:

You MUST use the same G.711 µ-law encoding that the supplied files use.

(Telephone) Replay Ports (Communication Manager only)

The **Operations > Replay** page is only available if you have purchased one or more telephone replay port licenses. To use this feature, enable the required administration page(s) by checking **Phone Replay** on the **Optional Admin Pages Enabled** setting for the appropriate switch(es) under **General Setup**. The configuration page will appear when you restart the recorder.

Mode Setup

The top of the page simply shows:

- **Number of Softphones Assigned** - set this to the number of ports you want to make available for replay via the phone.
- **Warn when available channel count falls below** - which is explained in Common Settings on page 163.

You must restart the recorder after setting the number of ports available for phone replay.

Important:

Replay ports impose a significant load on the recorder. Be sure that you have specified a powerful enough server.

Note:

Phone replay is not supported on Standby servers. If you require fault tolerant phone replay, you must install a pair of Central Replay Servers and configure a pool of phone replay ports on each.

I/O Jobs: Archive, Import, Mass Export

Avaya Contact Recorder lets you configure a range of fully automated input and output ("I/O") jobs that perform several "roles" - letting you archive, import, export and share your recordings with others.

The **Operations > I/O Jobs** page also lets you specify how long recorded files and the details about them are to be retained.

The table below explains what the various types of job are and which server(s) they may run on.

Input/Output Role…	**Notes**
ARCHIVE Copies all or selected recordings and their details to designated location(s) shortly after they are made. Notes details of these copies so that they can be retrieved later for replay. Advises replay server(s) of these copies. Optionally, purges archives older than a specified number of days.	A Central Replay Server does not record calls, so can only access archived recordings for replay, not archive new recordings. A Distributed Replay Server can be configured to archive recordings that it receives into longer term storage.
EXPORT Determines which recordings make up each telephone contact for all or selected calls. Copies "stitched" files closely matching CMS records to a designated location. Names these files using call details.	This optional feature is licensed separately and runs on Central Replay Servers and Distributed Replay Servers only. All recordings must be in G.729 Mono format and are decrypted on export.
DISTRIBUTED REPLAY SERVER FEED Copies all or selected recordings and their details to a designated staging location in a format suitable for a Distributed Replay Server.	For further details, refer to Distributed Replay Server on page 288.
DISTRIBUTED REPLAY SERVER IMPORT Reads each file that appears in the specified staging location (which is configured as a DRS Feed on another Avaya Contact Recorder). Imports the recordings and details to its own database and calls buffer then deletes the file.	For further details, refer to Distributed Replay Server on page 288.

Input/Output Role…	Notes
CENTRAL ARCHIVE MANAGER (CAM) IMPORT Examines all CAM archive files in a specified location. Notes the details and location of the recordings within them, making them accessible for search and replay directly within Avaya Contact Recorder. Optionally, purges archives older than a specified number of days.	CAM was a separate archive server used in Version 11 and earlier. This import mechanism allows users who upgrade to continue to access archives created by it. See *Avaya Contact Recorder: Migrating from Central Archive Manager* for further details.

Managing I/O Jobs

You can add, edit, delete and view the status of I/O jobs on the **Operations > I/O Jobs** page of the Master, CRS and DRS servers. Most configuration is performed on the Master recorder only. The information is disseminated to other recorders automatically though Central Replay Servers require some additional configuration.

Overall Settings

In the upper part of the page are settings that relate to all I/O Jobs.

Minutes to wait before considering new recordings

It is not uncommon for additional details to be added to call recordings shortly after the recording ends. This setting instructs the recorder to wait for the specified number of minutes after a recording has been made (or, in the case of a replay server, consolidated) before examining the call to see if it should be archived and if so, to where. This affects archive, Mass Export and DRS Feed jobs.

Service Archive Requests

If set, this server will service any archive requests that it can rather than forward them to the recorder that made the recording. That recorder may still have the original file and not need to access the (typically slower) archive copy. If the archive storage is co-located with this server then it is normally best to have it attempt to service replay requests from there rather than add load to an active recorder.

Retain Call Details for (days)

At some point, the size of the call details database will become either unmanageable or will fill the available disk space. Specify, in days, how long the system should retain call detail records before they are purged from the system. This ensures that the database stabilizes at a finite size. Purging is carried out at, or shortly after, 1:00 a.m. each night and does not affect recording or replay.

This figure must be at least as large as the one below and must be at least 31 days more than the retention period configured on any individual I/O Job below.

If you reduce this figure significantly, such that suddenly many months of recordings are due to be purged, it may take several days for the recorder to enforce the new limit fully as it will only attempt to purge a maximum of 500,000 recordings each night.

Maximum Days to Retain Recorded Files on Local Disk (if space permits. 0 = as long as possible)

Enter the number of days after which recorded files will be deleted from the disk buffer. Set to 0 to disable this feature and keep calls until the space is needed for newer calls.

You cannot set this higher than the **Retain Call Details** setting above.

CAUTION:

Calls will still be deleted sooner if the disk buffer fills up and the space is needed for new recordings. Use Archive jobs for fixed term storage.

Recordings Awaiting Consideration

This (read only) figure shows how many recordings have been made or received and will be processed by the relevant I/O jobs once the Wait before archiving timeout has expired on each.

Input/Output Jobs Table

The main body of the page shows the details of each I/O Job that you have configured. Note that the information is static. To update it click the **Refresh** button.

Select

Click this checkbox and then **Delete selected job(s)** to delete one or more jobs.

#

The unique reference number for each job.

Role

Which of the I/O job roles this job is performing.

Detail

The **Job Title** shows here, along with any **Advanced** settings that you have adjusted away from their defaults.

Path(s)

This shows the file storage location(s) you have specified for the job and, where appropriate, the state of each (for example, the available space on an archive drive).

Status

The number, age and volume of recordings waiting to be written or “pending” (in the case of jobs that write data) and the number and volume of files that have been read or written.

Edit

Click this link to change how the job is configured.

Adding I/O Jobs

At the foot of the page will be one or more **Add…** buttons – according to the type of server and your licensed options. Simply click to add a new job of the type indicated and fill in the details on the dialog that pops up. For all job types:

- You can enter a **Job Title** field to remind you of the purpose of each I/O Job that you create.
- You will also need to specify a **File Storage** location (or, in some cases, locations) that it will use. This is discussed in more detail below.
- If you need to refine how the job operates, click the **Advanced** button to see the additional options that are available. These are discussed in detail below for each job type.

If you have a Central Replay Server, you will find that it automatically creates Archive jobs that it learns about from the recorders that feed it. However, it only learns basic information in this way and you must edit the configuration manually. For example, you will need to provide the path and credentials if you want the CRS to be able to service replay requests directly rather than forwarding these on to the individual recorders.

Editing I/O Jobs

To change the configuration of an I/O Job, click the **Edit** link to the right of its entry in the table.

Deleting I/O Jobs

To delete one, or more, I/O jobs, click the **Select** checkbox (left hand column) for each and then the **Delete selected job(s)** button (bottom left). However, if a job has already succeeded in archiving calls or importing them from CAM it cannot be deleted as this would stop you replaying the files it has already processed. You should use the **Advanced** option **Disable** instead. This will stop it processing any further recordings.

File Storage

Recordings are read from and written to a wide range of storage devices - not all of which are suitable for or supported by all types of I/O job – as the table below shows.

File Storage Type	Archive	Mass Export	DRS (feed and import)	CAM Import
Windows (“smb”) file-share	Yes	Yes	Yes	Yes
Local (or mounted) disk	Yes	Yes	Yes	Yes
SFTP Server	Yes	No	Yes	No
EMC Centera	Yes	No	No	Yes

Note:

All Avaya Contact Recorders controlled by a Master recorder (whether Standby or Slave) will attempt to use the same path and credentials. If these vary, you must create a separate I/O job for each variant and specify which server(s) each is to run on.

In all cases, the account running the Avaya Contact Recorder service must have the appropriate access rights to the location(s) specified (read and write for all except CAM import which only requires read access).

Never configure more than one I/O Job pointing to the same destination.

Capacity

The recorder makes no attempt to manage the available storage space. You may use a Hierarchical File Storage (HFS) system to do so but you should test that the retrieval time is acceptable.

If instructed, the recorder will delete archived or imported tar files that have passed the set retention period but it will not delete them merely to make room for new archives as the recorder's main circular buffer does.

Increasing Capacity

For Archive jobs (only) you may "top up" the available storage space by making additional directories available as the initial one(s) fill up. Add these paths to the end of the list of paths and the recorder will start using them as the earlier one(s) fill.

Note:

The recorder will only start writing to a folder that is *completely* empty. If it has been used previously, even if all the tar files have been purged, there will still be a **".uid"** media identifier file left. This will stop the folder being reused. If the tar files have all been deleted, you must delete this **".uid"** file if you want the folder to be filled again.

Local (or Mounted) disk

A key role of most I/O jobs is normally to copy recordings to or from somewhere other than the recorder itself so that those copies are safe should the recorder or the site it is in be destroyed or stolen. However, a "mounted" disk of almost any type can be used – and these may actually be remote. Supporting a local path allows such mounted drives to be used – even if the file system on them is not one of those supported directly by Avaya Contact Recorder. For example, NFS and SMBV3 may be used.

When using a local (or mounted drive) you merely need to provide the path to be used. You must specify an absolute path as relative paths may change should subsequent upgrades use different working directories.

If you have more than one recorder, standard practice is to mount a single, shared folder with the same local path on each recorder – and replay server. That way, each recorder or replay server can not only write its own tar files to its sub-folder but can also read the tar files from all the other recorders – so providing replay without having to forward requests on to the recorder that made the recording. If you do not do this, and each local path points to a different folder, you should set the **Service Archive Requests** setting to **No** – forcing each server to forward replay requests to the recorder that made and archived the recording.

Windows (SMB) Fileshare

The most common storage location used with Avaya Contact Recorder. The recorder uses the JCIFS-NG library to access Windows file shares. This library supports both SMB V1 and SMBV2. Since support for SMB V1 is deprecated it must be enabled using the property setting **archive.forcesmb1=true.** To use later versions, mount the file-share as a local path as above.

When using an SMB fileshare, you must specify:

- The **Path** in UNC form. You may include subfolders within a file-share e.g. **\\mynas\avayarecordings\campaign_X**
- A valid **Username** (in the form ***domain**username***) that has read/write access to the path.
- The **Password** for the above account.

Distributed File System Disabled

By default, Avaya Contact Recorder configures JCIFS-NG with Distributed File System (DFS) disabled – to avoid common problems with DNS configuration that slow access dramatically in systems where DFS is *not* actually used (most systems).

If you need to enable this feature, edit the `.conf` file in the `wrapper` folder beneath the install path. Remove the line containing `-Djcifs.smb.client.dfs.disabled=true` and if this was not the highest numbered line, renumber those below it to remove the newly created gap in the numbering.

SFTP Server

This provides secure transmission and is well suited to external connections, for example, if delivering files to another company. However, the round-trip delays make it unsuitable for Mass Export jobs which typically write huge numbers of small files.

The SFTP server used must support:

- SFTP V3 or later
- The statvfs@openssh.com extension

You can check a server's capabilities using WinSCP. Establish an SFTP connection, then select the menu item Session > Server- protocol capabilities. On the **Capabilities** tab, the **Can check available space** entry must be **Yes** and the above extension must be listed as shown below.

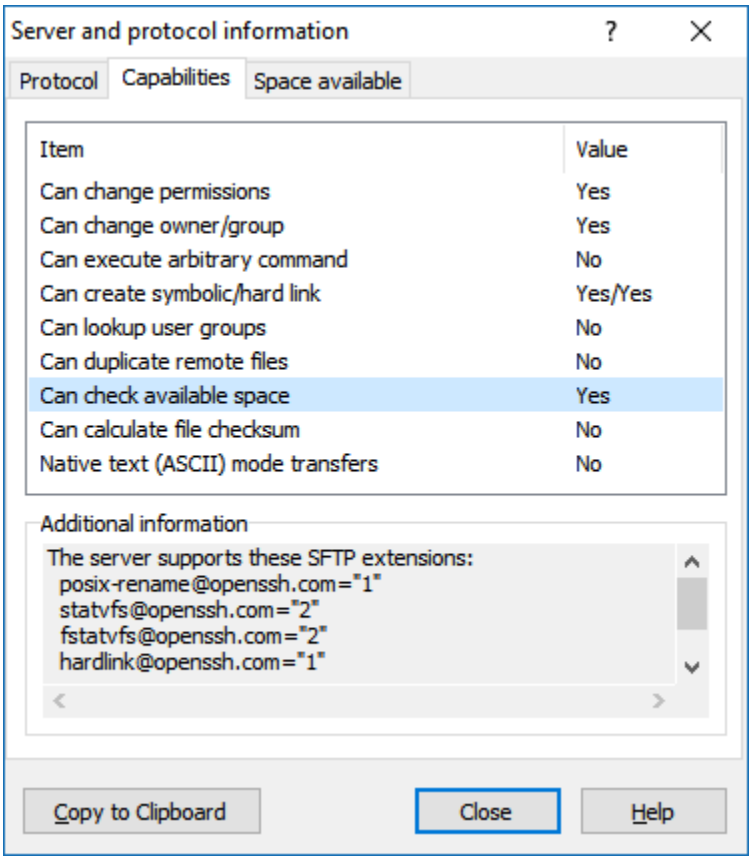

When using an SFTP Server, you must specify:

- The **URL** of the SFTP server in the form **`sftp://`*`hostname_or_ipaddress`*`[:`*`port`*`]/`*`absolute_path`***. You may include subfolders. For example **`sftp://sftp.bigco.com/avayarecordings/campaign_X`**
- A valid **Username** that has read/write access to the folder specified.
- The **Password** for the above account.

EMC Centera

If you wish to use an EMC Centera file store, you must first download the appropriate driver files onto each Avaya Contact Recorder and then restart the recorder.

Downloading the EMC drivers

EMC provides drivers for Centera for Windows and Linux at their support website. The package is called Centera SDK. ACR is tested with version 3.2p5 of the Centera SDK.

- For Linux choose the version for gcc 4.4.
- For Windows choose the 64-bit version.

Installing the drivers on Linux

Download the Centera SDK.

Perform the following steps as root (replacing the specific SDK version with the one you are installing if different):

```
cd /tmp
tar xvzf downloadedFile.tgz
cd Centera_SDK_Linux-gcc4/Centera_SDK-3.2.705/install/
./install
Accept the default installation location of
/usr/local/Centera_SDK
cd /usr/local/Centera_SDK/lib
chmod 755 FPLibrary.jar
cd /usr/local/Centera_SDK/lib/64
chmod 755 *so.3.2.*
Perform the following step as witness
cp /usr/local/Centera_SDK/lib/FPLibrary.jar
/opt/witness/lib/
```

Installing the drivers on Windows

Download the Centera SDK.

Unzip the downloaded file to a temporary location.

Under the temporary location, locate the lib and lib64 folders.

Copy `FPLibrary.jar` from the `lib` folder to the `lib` folder under the ACR installation.

Copy all the `dll` files (not the `lib` file) from the `lib64` folder to the \wrapper folder under the ACR installation.

Configuring

When specifying EMC Centera file storage, you must enter the pool access string for the EMC Centera server you want to use. You should include multiple IP addresses of the pool and the full path to the PEA file. A typical connection string takes the form

`10.10.10.10,11.11.11.11?/home/witness/atlantaemc.pea`

Note:

If you remove or change one or more paths from the destination list you must restart the recorder.

Exported Files

“Stitching” of Recording Files

A single exported file typically contains one recorded file, or “call segment”. However, in the following cases, two or more files from the recorders are “stitched” together – with the appropriate amount of silence in between:

1. When a call is placed on hold and then retrieved, two separate recordings are made – but these are joined back together for export. The silence inserted between them is as long as the call was held for.
2. When a consultation call is made, and the original call is then transferred to the new party – these two separate recording files are joined together for the export. The section between the consultation and the transfer is not included in the file.
3. When a party leaves a conference call, the recordings will be joined together unless the party leaving was the last party to join the call.

Recording Format Restrictions

Because multiple recordings may be “stitched” together as described above, it is important that *all* recordings stored on the recorder are in the same format. A single “.wav” file cannot contain a mixture of mono and stereo sections or multiple codec types.

Export only supports mono G.729 files at present – so the following stereo recording mechanisms can cause problems:

- CS1000 Duplicate Media Streaming
- Communication Manager Passive IP recording
- SIPREC recording

If your system includes any of the above and you wish to export recordings, you must ensure all segments of a call recording are in mono by doing the following (as export does not support stereo recordings):

1. Ensure that the RTP is encoded in G.711 rather than G.729a as it reaches the recorder.
2. Set the **Keep audio in stereo wherever possible** setting on the **General Setup > Server** page on each recorder to **No**.
3. You must not set the property `acr.disablecompress=true`

Replaying Exported Files

The exported wav files are in G.729a (mono) format and can be played via Media Player if you have previously installed the WFO replay tools.

Advanced Settings

For each of the different types of I/O Job, there is a number of **Advanced** settings that can be altered by clicking on the **Advanced** button while Editing their properties. Each of the settings is discussed below.

Filter

This setting applies to output jobs only (archive, Mass Export and DRS Feed). By default, such jobs will process all recordings that they make or learn about.

To restrict it to a specific set of calls, such as those from a particular campaign, you first select, from the drop-down list shown, one of the same Search and Replay Layouts that you are familiar with from the Replay screen. Choose the one that either already filters the calls as you require (layouts with fixed filter fields can do this) or which provides the filter field or fields that you need to specify your required criteria.

Note:

At present you can only use Segment-based Layouts for archive jobs as these act on each individual recording segment shortly after it completes. As the archiving decision is made for each recording segment you must ensure that the criteria you set do retrieve all segments of calls that you want to archive and not, for example, just the initial segment.

Note:

If you are exporting CSV entries for each recording that you archive, you may wish to specify a particular layout even if you are not filtering the calls at all. This is because the details written to the csv file are the columns of a search and replay layout – and if you want to include details not present in the default Avaya layout you should specify the appropriate one here.

As you select a particular layout, the dialog expands to show the filter fields for that layout (with the exception of Call Set), exactly as they would appear on the left of the **Replay** page.

Note:

Archives cannot be filtered by Call Set. Calls are placed into call sets by users of the replay application – but this is typically long *after* the call has been considered for archive. To keep long term copies of calls that you have assigned to particular call sets you should, instead, use the export feature of the Replay interface described in the *Avaya Contact Recorder User Guide*. This provides not only the recordings but also their details in a standalone form suitable for long term retention.

Enter the criteria exactly as you would do when searching from the **Replay** page. In fact, it is a good idea to first use the **Replay** page to determine which layout and which criteria you should be entering – as there you can see the results of the filtering immediately.

When it is selecting the calls you expect, simply transfer the settings that you have tested into these **Filter** settings.

Please refer to the User Guide for more details on the available fields and search operators.

Note:

Although you can specify a date and time range for calls, this is seldom used as an archive destination is typically left running indefinitely. It would only be of use in selecting calls relating, for example, to a specific campaign to be run over a known time period.

There is therefore no need to (and it harms performance if you do) set a date/time range for most I/O jobs. Archive and Mass Export jobs act on recordings shortly after they are made. ONLY set a date/time range if you genuinely need the job to be run against calls that had already been recorded when you created the job and/or the job must stop processing calls at some point in the future. When setting a historic range, take care not to have it look back beyond the oldest call on disk or in the server's database as calls cannot be archived once they or their details have been purged.

If you do specify a date/time range that extends into the past, the recorder will show on the **I/O Jobs** table that it will check for historical recordings overnight. This ensures that the database load incurred does not impact recording. It also gives you a chance to change or correct your filter parameters before they are used in anger. Overnight, the recorder will run a query to identify existing recordings that meet the criteria. These will be added to the archive queue alongside new recordings.

When setting the timespan, you may notice that the SQL statement shown as a result of this appears to be one second out on the end date and time. This is deliberate - so as to avoid missing calls that you thought would be included. The start date and time displayed on screen is actually rounded to one second granularity so the stored time could be a fraction of a second later and would therefore not be included in the result of a query run to exactly that time.

Bundle into tar Files

This setting applies to output jobs only (archive, Mass Export and DRS Feed) and specifies whether recordings are exported individually (in the case of Mass Export jobs) or combined into a much smaller number of larger ".tar" files (in the case of archive and DRS feed jobs).

Although you cannot change between these two options, you can, in the case of archive and DRS feed jobs, specify two parameters that affect when tar files are written. These determine how large (in MB) each tar file should be and the maximum interval (in minutes) between tar files being committed to the **File Storage** location.

The default settings (100MB and 24 hours) are appropriate for most live systems. On a lab or test system, you can reduce these to make testing quicker (you don't have to wait till you've recording 100MB or 24 hours have passed before you see the next tar file appear). If you are particularly concerned about making an off-site copy of recordings as soon as possible, you may want to reduce the maximum interval but less than 15 minutes is not recommended as this may generate a large number of very small and hence inefficient tar files, particularly overnight.

Content to be Written

This setting applies to output jobs only (archive, Mass Export and DRS Feed). By default, all recordings (whether audio or screen) and the full metadata about them (their XML file) are archived or exported.

Audio Recordings

Determines whether or not audio recordings are written.

Screen Recordings

Determines whether or not screen recordings are written.

XML Metadata

Including XML metadata associated with each audio or screen recording that is output is strongly recommended (and is actually required for all DRS feeds) as without it, you are entirely reliant on the recorder's database to know what is in each recording file. By including the XML file, you allow subsequent reimport of the full recording detail to a fresh database. This makes the archive self-consistent and not reliant on any other backup.

CSV Entry

You can also choose to export a row to a "csv" (actually tab not comma separated) file for each file that is archived. This includes basic details identifying the call, and which tar file it is in plus the column entries from the replay layout specified under the **Filter** setting above (or the default layout if none is specified there).

The files generated:

- Use UTF-8 format with, a byte-order mark; tab field delimiters and Carriage Return + Line Feed line delimiter.
- Do NOT contain a header row.
- Reside in the root folder specified under File Storage and are named after the recorder that generated them and the date they cover. For example, recorder 801234 would generate file `2015-05-01.801234.csv` for recordings starting on 1st May 2015.
- Contain one row per recording archived or exported.
- May contain blank rows.

- Should be manually purged if required. There is no automatic deletion as the size of these is generally insignificant compared to the size of the recordings being archived.
- Are appended to by the recorder. Do not assume these are complete until 8 hours into the next day (or longer if I/O jobs are backlogged). If you set a historical date range, this will result in that range's files being updated when the range is processed overnight.
- Details can only be written to this file once they are fixed. In the case of an archive, this is when a tar file is written to it. In the case of Mass Export, this is (by default) 2 hours 15 minutes after the end of the recording as a subsequent segment of the same call may complete within that period and result in the previous file being appended to and its name changed.

TIP:

To make lab testing with short calls quicker, you can force these to flush through in a shorter period e.g. 20 mins by setting property `csv.contactpurgemins=20` either in the property file or via the maintenance page (…`/servlet/acr?cmd=mtce`). Revert to the same value as the maximum call segment on the master recorder plus 15 minutes for production.

Each row in the file contains the fields shown in the table below:

#	Name	Example	Notes
1	inum or inums (in the case of a stitched export file built from multiple recordings)	`801234000123456`	15 digit number(s) identifying the recording(s). First 6 digits are serial number of recorder that made the recording. Last 9 digits are unique reference within that recorder. If multiple INums are listed, these are comma separated.
2	filename	`801234/1023.tar`	Relative filename of the file within which this recording has been archived or the actual ".wav" file if not using ".tar" files.
3	callid	`00001012341412433428`	The identifier given by the switch responsible for the call (UCID for CM calls, dialer call id for PCS calls)
4	startedat	The remaining fields in each row are the same as those of the Search and Replay "Layout" used to filter the archive or export job. This is specified in the setting **Operations > I/O Jobs >**	
5	duration		
6	agents		

7	parties	**Advanced > Filter**. Those shown in the second column are typical (from the default layout).
8	service	
9	skills	
10	callid	
11	udflist	

Sub-folders

This setting applies to output jobs only (archive, Mass Export and DRS Feed). Within the **File Storage** path specified, files can be separated into several levels of sub-folder according to the recorder that made the recording or (in the case of Mass Export) the details associated with the recording.

In the case of an archive job, you can select:

- A **Folder per recorder** – in which each recorder places its tar files in a sub-folder named after its serial number. This was the only option available prior to Version 15.1 and existing archives will continue to use this option on upgrade.
- A **Folder per recorder per hour** – in which case additional levels of sub-folder are created for the year, month, day and hour in which a tar file is written. This ensures that no one folder ever has more than a few files in it and is therefore the default for newly created archive jobs.

In the case of Mass Export jobs, you can select:

- A **Folder per Skill** – in which case sub-folders are created for each skill on each day of each month. For example, a call recorded on 29 June 2015 using skill 12345 would be written to `2015-07/29-06-15/12345` while a call with no skill would be in the folder `2015-07/29-06-15`
- A **Folder per Owner** – in which case files are assigned to sub-folders initially on the basis of the (one and only) owner of each recording[1]. This is appropriate when **Advanced** recording settings have been used to assign each recording to a (single) specific alphanumeric owner. For example, recordings made on 29 June 2015 and owned by "sales" would be written to sub-folder `sales/29-06-15`. A further level of sub-folder is used to separate calls on the basis of the skill associated with the call. Two special cases are also handled – those with no skill and those made by a dialer. The former are stored in sub-folder `NoSkill` and the latter in sub-folder `APC`.

[1] If you do have multiple owners, any one of these may be used for a given recording – so this option is not recommended.

File Names

This setting applies to Mass Export jobs only. In all other cases, files are named according to the 15-digit reference number assigned by the recorder. In the case of exports, however, each exported file may have to be identifiable without the benefit of the associated database and may include more than one recording file. Two options are available.

Note that there is a 255-character limit on filenames so filenames will be truncated if longer than this.

Similarly, if there are any characters in the filename that are not valid in a Windows NTFS filename ('/' – forward slash, '\' – backslash, ':' – colon, '*' asterisk,' '?' question-mark, ' "' – double-quote, '<' – less than, '>' – more than and '|' – pipe character) these are replaced by an exclamation mark ("!") and a WARN log entry added to the log file.

CTI with tags

In this case, the filename consists of the following items concatenated:

- The call identifier of the original call
- "_ANI_" + the ANI (defined here as customer number if available for an inbound call, station for an outbound call)
- "_DNIS_" + the DNIS (defined here as answering station for inbound call, customer number for an outbound call)
- "_DATE_" + the date of the call in format dd-MM-yy
- The start and end times of the recording file in format HH-mm-ss and separated by an underbar
- "_GE_" + the hunt group the call was routed through
- "_AGE_" + the (latest) agent on the call
- Any user defined fields the call is tagged with, as "_***udfname***_" + ***udfvalue***

So a recording of call 00001012341412345678 from 0137286900 to station 12345 on 29th June 2015 from 09:00:12 to 09:04:25 via hunt group 2001 taken by agent 3001 with userdefined fields custype=gold and compcode=73 would be named:

```
00001012341412345678_ANI_0137286900_DNIS_12345_29-06-
15_09-00-12_09-04-
25_GE_2001_AGE_3001_custype_gold_compcode_73.wav
```

CTI but no tags

This naming convention is the same as above but with just a single underbar between the fields rather than the labels listed above. So the same call as above would be named:

```
00001012341412345678_0137286900_12345_29-06-15_09-00-
12_09-04-25_2001_3001_gold_73.wav
```

Runs on Servers(s)

Not every I/O Job can or should run on every server in your system. This field allows you to choose from the appropriate servers for each type of job. Use it to specify which server(s) each runs on – for example, if not every server has the same drive letter for a mounted drive or if you want to use different credentials or paths on some.

By default:

- An Archive or DRS Feed job that you enter on a Master will run on it and be copied to and also run on all Standby and Slave servers that it controls.
- An Archive job that a CRS learns about from a recorder that feeds it will attempt to run on that server (but you will have to manually add the **File Storage** details to the archive job that it creates automatically when it sees recorders using an archive).
- Mass Export jobs only run on Replay Servers (CRS and DRS).

Note:

The default settings for each type of job are applied when you create an I/O Job. If you subsequently add another server – for example an additional slave – you will need to explicitly enable the appropriate job(s) to run on that server. You will notice that the Detail column of the I/O Jobs table shows the set of servers if this does not match what would now be the default.

Retention

An archive or CAM import job notes the date of the most recent recording stored within each tar file. If can therefore purge tar files in which all recordings are more than a specified number of days old. If you select this option, then overnight (normally shortly after 1AM) the recorder will determine which files are old enough to be deleted and delete them.

CAUTION:

If you enable this feature, deleted files cannot be retrieved.
This feature continues to run even on a **Disabled** I/O Job.
To stop it, uncheck the **Purge overnight** checkbox.

Hours of Operation

By default, I/O jobs run 24 hours a day but you can choose which hours they run. This can be useful to limit bandwidth to a remote **File Storage** location; to allow a maintenance window on the file server or to reduce load on the server.

Simply clear the checkbox(es) for the hour(s) you want the job to pause.

Disabled

If an I/O Job is no longer required, you can disable it. No more recordings will be read from or written to it but those already there will remain accessible and nightly purging will continue if a retention period is set.

If there are any files pending for the job, you should flush these through before disabling the job. Reduce the time interval before which tar files are written to a few minutes and wait for them to be written. Otherwise, an alarm will be raised when this tar file becomes ready for writing but finds that its job has been disabled.

CAUTION:

Do NOT disable a job in order to reboot, upgrade or expand the file storage it is writing to. Recordings made while a job is **Disabled** will NOT be archived when you re-enable it. Instead, either pause the job by reducing its **Hours of Operation** (see above) or simply be prepared for the Alarms that will occur when the **File Storage** location is taken off line. The recorders are designed to tolerate this and will keep jobs pending until the storage is accessible again, at which point they will resume archiving and catch up fully automatically.

A disabled job will continue to check for and purge files older than the retention period (if set). To stop this, you must clear the **Purge overnight** checkbox.

Search, Replay and Live Monitor

The instructions below relate to the integral search, replay and live monitor application that is part of the recorder itself. The accompanying User Guide provides end user instructions on this application. As the system administrator, however, you may wish to:

- control which users can search and replay recordings and/or live monitor
- customize the search fields and/or display using **Layout Builder**
- ensure that the Java Applet used with Internet Explorer 11 can be downloaded to your clients' PCs
- restrict access to certain replay layouts
- enable other security features
- allow users to "lock" and "unlock" recordings
- force some users to obtain authorization before they can replay recordings
- modify default behavior via the properties file

Access Rights

Access rights are used to control which recordings can be replayed, which sessions can be searched for and which devices can be monitored in real-time.

Replay Rights

When a recording is made, the recorder assigns its "owner" or, in some cases "owners" as follows:

Audio Recordings

- If the **Recording owner** field is set for a recording mode as a whole, or on the **Advanced** settings for the address being recorded, the owner will be the number(s) or name(s) specified there (with any address specific setting overriding the overall setting).

  Note:

  By default, this will override the lower priority options listed below. For Bulk recording mode only, you can change this by setting the property **`owners.additive=true`** – in which case the following options may also add one or more owners.

- If the recording is made because a VDN or Skill group is configured to be bulk recorded, then that VDN or Skill group address will be the owner. This applies to Communication Manager recordings only.
- Otherwise, the owner will be the station(s) or agent(s) that were on the call and led to it being recorded. If an agent is logged on at a station, that agent becomes the owner (regardless of whether it was the agent or station's address that triggered the recording). Failing that, the station is the fallback owner. Note that internal calls often therefore have multiple owners as more than one party on the call may be recorded.
- Where **Follow the Call** is enabled – either explicitly or implicitly due to a WFO Business Rule firing, the owner(s) of one segment are carried forward to successive recorded segments (which may have no other owner added and hence would otherwise not be playable by anyone). This ensures that someone authorized to replay a segment of such a call will be able to "follow the call" in replay as well.
- Communication Manager recordings have additional variants:
  - If **Add VDN as additional "owner" of calls** is set on the **General Setup** page for a particular Data Source, then calls routed via a VDN will have a VDN as another owner. A further setting on that page determines whether this is the first or last VDN a call went through.
  - **Meeting Recording** however, does not follow the above rules. In this mode, the voice prompts advise the caller to enter one or more owners.

Tip:

Use dummy station identifiers to allocate owners to calls made in Meeting Recording mode. All members of a particular team can be configured with replay rights for a particular number, even though this is not a valid station number. When prompted at the start of the call, mark meetings recorded for and by a team with this "owner" so that all members of the team can access the recording.

Screen Recordings

Where associated with a voice recording, these run for the entire duration of a session rather than a single “segment” of a call as the voice recordings do. A single screen recording therefore may encompass multiple voice recordings, each of which could have different owners. The “owner” of a screen recording associated with one or more voice recordings is therefore always the telephony Agent logged on at the associated telephone or, if that is not known, the associated telephone.

Screen recordings that are not associated with voice calls, such as those started by DPA control, are owned by the agent or device specified in the command that started them.

Controlling Access to Recordings

You control which recordings your users can search for and replay by adding a user account for each person that needs to use the Replay page and specifying which range(s) of owner they are entitled to see.

To add a user account, follow the same procedure that you used in

Securing the System on page 124 but do not select either of the Administrator roles. Other role options let you control whether the user can monitor recordings in real-time, lock and/or unlock calls, must be authorized and/or be allowed to authorize others and whether or not the user is allowed to export recordings as files.

Each user's rights are shown on the **System > Manage Users** page and are set when adding the user account. You can change these by clicking on the **Edit** link next to them. These rights determine:

- which recordings can be played (the user must have rights over at least one owner of the recording)
- which call segments can be searched for (the user must have rights over at least one owner of the recording)
- which sessions can be searched for (the user must have rights over at least one owner of at least one recording within the session)
- which addresses the user can live monitor (assuming live monitor is enabled at all for this user)

The initial administrator's account is automatically given access rights to all number ranges up to 14 digits. As you add other users you must specify which ranges of owners each user is entitled to replay. The number of digits is significant. A user with replay rights over 0000-9999 cannot replay calls made by and "owned by" agent 567 though they could play calls owned by agent 0567. In this example, you might grant the user replay rights over 0000-9999, 000-999.

Typical examples of how to use replay rights are:

- A user allowed to play calls made on his own station (1234) would be given replay rights 1234.
- An Agent who logs on as AgentID 5012 and is allowed to replay his own calls may be given replay rights **`5012`**.
- A supervisor who logs on as AgentID 5050 and manages AgentIDs 5010-5019 and 5025-5028 may be given replay rights **`5050,5010-5019,5025-5028`**.
- All recorded stations used by the HR staff have their **Recording Owner** set to HR on the **Advanced** settings. The Human Resources Manager, who uses station 5678 may be given replay rights **`HR, 5678`**.

Note:

If you have enabled switch specific identifiers in order to resolve conflicting numbering scheme ambiguities, you must explicitly grant replay rights to the resultant address ranges. For example, to grant access to all addresses on a switch using 4 digit numbers with prefix "ATL" you should enter "ATL0000-ATL9999". This ensures that this user cannot access 4 digit numbers on other switches (which would have different prefixes). In the case of an AACC switch, the domain is uses as its switch specific identifier - so you must assign replay rights such as

`0000@myaacc.com-9999@myaacc.com`

Search Rights

When searching with a call segment-based layout, each call segment that is found relates to a single recording. These results will be limited to those recordings that the user has replay rights over - as described above.

When using a session-based layout, however, it is possible that a user may have replay rights over some, but not all of the recordings within a contact. Search results are restricted to sessions containing at least one recording that the user is entitled to replay. Therefore, should the user select a session for replay, they may find that some of its recordings show up as grey boxes on the display - indicating that they do not have authorization to replay those recording.

Live Monitor Rights

A user who is assigned the "**May monitor"** role will be able to use the integral live monitor application to observe devices that fall within their assigned **Access Rights**. Note that this may allow them to hear calls that they are not subsequently allowed to replay (for example, if the recording rules result in the recording's owner being set to something other than the agent or station's address).

Layout Builder

The **Replay** page in Avaya Contact Recorder provides users with features to search through all the recordings made by the recorder (or recorders in a multi-recorder system). Replay users are able to find the recordings that they are interested in playing by filtering the selection of recordings based on criteria like date and time, length of recording, skill groups, phone numbers, agent names, etc.

The formatting (or *layout*) of the Replay page is normally fixed by ACR. This means that the selection criteria, the order of the display columns, the labels used to describe filters, etc. are pre-determined. The panel at the left of the replay screen is a list of *filters* that select a subset of the available recordings based on the selection criteria. The results are shown in tabular form in the right panel. The *columns* of the table show a subset of the full information available about each recording, some of which is grouped for ease of display. For example, the 'agents' column shows Communication Manager agents, Proactive Contact agents and Proactive Outreach Manager agents all merged into one column.

Two default layouts are provided automatically. The "Avaya" layout operates on individual recordings, each of which captures a "call segment". When a call is placed on hold and retrieved again, this breaks it into two segments - as does any other change to who is talking on the call. The "Session Layout", on the other hand, operates at a higher level, looking at the whole of an individual's interaction with a particular contact. This can include not only successive segments of the same call but also related calls - such as consultation calls.

"**Layout Builder**" provides a mechanism to create additional layouts, based on either of the two defaults and each customized for a particular purpose. Some typical uses for customized layouts are:

- Changing the names of columns or filters (e.g. 'associates' instead of 'agents')
- Changing the order of columns or filters
- Pre-setting (and, optionally, locking) one or more filters to specific values (e.g. duration longer than 30 seconds)
- Hiding certain columns, which may contain confidential or irrelevant information
- Adding columns to display information that is not normally of interest (e.g. a customer-specific user-defined field)

Each replay user can be restricted to seeing just a subset of customized layouts. This means, for example, that managers may be given access to a layout showing certain confidential information, while other users may only be able to see a layout that has been customized to hide that confidential information.

Only users with full **System Administrator** role may use the **Layout Builder** and assign layouts to users.

Note:

At the time of release, Layout Builder did not support editing the layouts using Arabic (right-to-left). If you experience problems opening Layout Builder with your browser configured for Arabic, please change your browser to another language while editing layouts.

Layouts

Each layout consists of these elements:

- An internal name, which allows administrators to differentiate between layouts. (The default layout that is provided with the system is called 'Avaya'.)
- A set of filters that select from the database just those recordings that match the filter. (By default, each filter selects all recordings, but replay users may use them to select just a subset.) Some filters may be hidden from users.
- A set of columns that define which information from each recording is displayed in the results table and how that information is grouped and ordered. Some columns may be hidden from users.
- Internationalized names (in whichever languages you actively use) for the layout itself, each filter and each column.
- A setting that controls whether the layout is hidden from all users (typically used during construction or if a layout is no longer in use).
- A setting that controls whether, by default, all replay users are automatically given access to the layout.

- A setting that controls whether the layout uses “turbo” mode. This increases the storage space required by a factor of two but increases search speeds by (typically) a hundred times. This feature only applies to Central Replay Servers and must be explicitly enabled. See “Turbo” Mode on page 285 for further details.

Administrators can use **Layout Builder** to change each of the elements of a layout, as well as creating new layouts and deleting layouts.

Note:

Only the two control settings may be changed in the two system-provided layouts. The internal name, filters, columns and internationalized text are fixed. If you want to change a system-provided layout you should copy it, change the copy, make the copy visible and available and then hide the system-provided layout.

Using Layout Builder

Note:

Layout Builder is designed for use in 'standards' mode of Internet Explorer 9 to 11 only. Make sure IE is not set in 'compatibility' mode.

On the **System > Manage Users** page, click the **Edit Layouts** button. **Layout Builder** launches in a new window.

Title Bar

The title bar shows a list of the layouts defined in the system. On a new install there are just the default call segment-based "**AvayaSegment**" and default session-based "**AvayaSession**" layouts. Use the drop-down list to select a layout to work on.

Save

When you have made changes to a layout, you must click save to commit the changes. If you try to move away from a layout that you have changed without saving you will be asked if you wish to discard those changes.

Copy

Click copy to make a new layout. The new layout is a copy of the currently selected layout, except that the internal name is changed to 'New Layout'. You should change the internal name to something meaningful to you before saving.

Delete

Click delete to delete the currently selected layout. You may not delete the system-provided layout; if you do not want users to see it, you may simply hide it.

Close

Click close or close the popup window when you have finished using **Layout Builder**.

Left and Top Panels

The top and left panels show the columns and filters respectively. They are shown in the order that they are displayed in the Replay page.

Note:

Filters and columns are always paired, because, in the majority of cases, if you want to see a column with a certain piece of data you may want to filter that data, and vice versa. This pair is called a field. There will be times when you do not want to see one or other of the filter and the column (see examples below). In this case you can disable the filter or the column so that it will not be shown.

To change the order of the filters, drag the filters in the left panel into the correct order. To change the order of the columns, drag the columns in the top panel into the correct order.

Disabled filters and columns are shown faded. Disabled filters and columns are always shown at the end. To disable a filter or column clear the check box. To enable it, set the checkbox.

To change the settings of a field (filter-column pair) click the pencil in either the filter or column. This shows the field editor.

Tip:

If your layout contains many fields, you may need to zoom out (press Ctrl and "-" keys) or use a higher resolution screen to avoid these panels overflowing.

Field Editor

Clicking the pencil of either the filter or column of a field brings up the field editor for that field.

First Row

Use the drop list in the first row to select the field type. Refer to Search and Replay Attributes on page 319 for a full description of each field type. If the field type you select requires a parameter (e.g. UDF), enter it in the second box on the first row.

Second Row

Use the second row to set defaults for the filter, if required.

If you wish to default the search criterion of the filter, set the first drop list to:

- **Normal** - sets the default search criterion, but allows the Replay user to change it.
- **Fixed** - sets the default search criterion, but does not allow the Replay user to change it.
- **Hidden** - same as fixed, but does not even show the filter in the Replay page Select the search operator in the second drop list.

Enter the search criterion in the third box.

Examples

- **Normal < 5** - searches for recordings where this field is less than 5, but lets the user change both "less than" and "5" to other values.
- **Fixed Contains 12345** - searches for recordings where this field contains 12345, and does not let the user change this selection.

Internationalization Rows

Enter the display name of the filter in the left box and the display name of the column in the right box. Fill these in for your own language (highlighted in yellow) together with any other languages used by your replay users. Often the display names of the filter and the column are the same. However, since space is typically limited in the column display, it often makes sense to make this display name shorter.

Buttons

- **Save** - saves your changes
- **Cancel** - abandons your changes
- **Delete** - removes the field (both filter and column) altogether

Middle Panel

Use the middle panel to:

- **Change the internal name of the layout**. You should make sure that all internal names are unique so that you can differentiate between layouts as you edit them.
- **Hide a layout**. Replay users will not see this layout, even if they are given rights to access it. This is useful while you are building and testing a layout.
- **Make a layout available by default**. If the layout is not hidden, all users will be able to see and use the layout without you having to assign them rights to it individually.
- **Create a new field**. As noted above, a field is a filter-column pair. The new field is added to both the filter panel and the column panel. After you create a new field, you should set its type and give it a name by clicking the pencil in either the filter or column panel.

Right Panel

The right panel shows the internationalized name for the layout that users of each language will see. You should fill in the name for your own language (highlighted in yellow) together with any other languages used by your replay users. Unlike the internal name, there is no particular requirement for these names to be unique.

Examples

Removing Irrelevant Information

If your call center does not use agents, you will notice that the system-provided layout always contains a blank column. You could save that space by deleting the **Agents** field.

Copy the **Avaya** layout, rename it, click the pencil on the **Agent** filter or column, click delete. Uncheck the **hidden** box so your Replay users will see your new layout. Finally **Save** the layout. Select the **Avaya** layout, check the **hidden** box and **Save** that too.

You have created a new layout without Agents, made it visible and hidden the system-provided layout. Your users will now just see your new layout.

Changing Display Names

You may prefer to refer to your Agents as "Associates". To change the display name follow the instructions for removing irrelevant information, but do not delete the **Agent** field.

Instead, in the **Field Editor**, change the two yellow highlighted names from **Agent** to **Associate**. Continue with the instructions above.

Fixing a search criterion

You may want to create a layout that only ever displays recordings that were tagged with a particular Vector Directory Number (VDN). This could be as a convenience to Replay users who make this search frequently (saving them from repeatedly searching for **Service Starts *xxxxx***), or it could be a security measure to ensure that a particular set of users may only see recordings tagged with VDN *xxxxx.*

Create a new layout and name it. Choose a name that reflects that this layout only ever selects recordings tagged with VDN *xxxxx*. Edit the **Service** field. On the second row select: **Fixed, Starts** and enter the VDN number. (If you don't want your users to see that the VDN in question is *xxxxx*, changed **fixed** to **hidden**.)

Since the column called **Service** will now always contain the value xxxxx it becomes largely irrelevant. Disable the column by clicking off the checkbox on the column called **Service**. It will be greyed out, and move to the end of the list of columns.

Displaying a User Defined Field

If you have an external integration, you may be tagging your recordings with User Defined information. For example, you may be tagging recordings with a customer id (as UDF 'custnum').

Create a new layout by copying the system-provided default.

Click the **Add Field** button to create a new field. Click the pencil on the newly added field to edit that field. In the first row, change the **Field Type** to **User Defined Field** and enter **`custnum`** to the right.

Fill in the relevant language translations (particularly the yellow highlighted ones). You might choose "Cust ID" for the column (where space is tight) and "Customer ID" for the filter.

If you want to display the UDF, but not allow Replay users to search by it, click the check box off on the Customer ID filter.

Forwarding Replay Requests with NAT and/or SSL

If your recording system consists of more than one server, some replay requests will be forwarded to the server holding the actual recording. If any of your Replay users are separated from these servers by a Network Address Translation (NAT) point and/or use SSL, additional configuration is needed.

For each recorder that your users access directly (Master, Standby or dedicated Central Replay Server) add a property identifying each other server as follows:

`fqdn.nnnnnn=<fqdn>`

where nnnnnn is the serial number of the recorder that the server may have to forward requests on to and ***<fqdn>*** is the fully qualified domain name that this server uses.

Client Prerequisites

Browser

Replaying calls through your PC uses the Web Audio API on browsers that support this standard (most except Internet Explorer) and a Java Applet on Internet Explorer.

Note:

The latest generally available versions of Chrome and Edge are tested extensively. Other browsers typically work and may be used but in the event of problems, you should revert to Internet Explorer V11, Edge or Chrome.

Live Monitoring, Exporting from the Replay screen and Telephone Replay are implemented as Java Applets only and so require Internet Explorer V11.

Java Runtime

The latest version of Java 8 should be installed on the client PC.

When running under Windows 8.1 you should install *both* the 32-bit and 64-bit versions of the Java runtime environment (JRE).

Applet Certificate Revocation Checks (IE 11 only)

If the client's browser cannot access the Internet, you must either disable revocation checks for all applets or, preferably, just for those downloaded from the Avaya Contact Recorder.

If the client cannot access the internet, either:

1. Click **Start > Control Panel > Java > Advanced**
2. Uncheck **Perform certificate revocation checks on**

or

1. Click Start > Control Panel > Java > Security
2. Click Exception Site List > Edit Site List
3. Add the URL used to access Avaya Contact Recorder e.g. `myacrserver:8080`

Applet Permissions (IE 11 only)

If prompted, the user must allow the applet to be installed, to run and to access the recorder and (if export is required) the local hard disk.

Internet Explorer (V11) Configuration

Internet Explorer (IE) determines rights by putting web servers into zones and then granting those zones specific rights. To access and use the application, including the

Replay and **Monitor** pages with their Java applet, the recorder must reside in a zone with the following rights:

- Scripting of Java applets (replay, live monitor and export use these)
- Active scripting (all web pages use this)

Your Intranet zone and/or the Trusted Site zone may already have these rights. If so, you need to verify that the Avaya Contact Recorder is in one of these zones. If the recorder is not in the local Intranet Zone or is not a Trusted Site, you can add it as follows:

1. From your browser's tools menu (Alt+X in IE 11), open **Internet Options…**
2. Click the **Security** tab.
3. Click the **Trusted Sites** icon.
4. Click **Sites**.
5. Uncheck the **Require server verification (https:) for all sites in this zone** box (unless you are forcing users to use https)
6. Enter the URL of the Avaya Contact Recorder server and click **Add**.

Note:

Internet Explorer does not recognize that a certain Fully Qualified Domain Name (FQDN) and IP address are the same; you must add the URL to the list exactly as your end users are expected to type it in the address bar.

Before you advise end users of the URL of the recorder, you should make sure that your users can access the Replay page through your network and that the applet downloads and runs successfully. To test this:

1. Create a user account without either **System Admin.** or **Restricted Admin.** role checked and assign it some replay rights.
2. From a typical client machine, enter the URL for the recorder in the form:

   `http://myservername:8080/` (using the recorder's IP address or hostname - assuming you have entered it in your DNS server). If using https, replace **`8080`** with **`8443`**.
3. When prompted, enter the account's Username and the temporary password set in step 1 above.
4. Set a new password as directed.
5. Confirm that the Replay page displays correctly and that the applet is downloaded.

Tip:

You may wish to copy steps 2 through 4 above, fill in your URL, and send them out as instructions to your end users.

Restricting Access to Replay Layouts

Using **Layout Builder,** you can choose whether a layout is made available to all replay users or none initially. You can then see which users are able to access each layout - and refine this using the **System > Manage Users page**.

As you click **Add User** to create a new user account or **Edit** an existing one, you can set the **Search/Replay layout(s) available** to each user. Those they are entitled to use are listed and can be changed by clicking the **Edit** link and altering the checkbox next to a layout's name. Click **Enter** to save these changes.

Miscellaneous Security Features

Additional security features can be enabled be setting **`viewerx.secure=true`** in the properties file.

Locking Recordings

The recording system normally keeps recordings for a planned time-span or until disk space is required for new recordings - at which point old recordings are purged from the recorder's disk buffer and/or archive path. However, specific recordings may be of interest for longer periods and it is important that these are retained for as long as needed. For example, those required by a "legal hold" order must not be deleted. This is achieved by "locking" the required recordings.

Enabling Lock/Unlock

To use this feature, assign the **May lock recordings** and **May unlock recordings** roles to at least one user account. Add the roles from the **System > Manage Users** page as you add a new user account or **Edit** an existing one. Typically, only one or two users (often the System Administrator) are allowed to unlock recordings while more people (supervisors, investigators, compliance staff) are allowed to lock recordings.

Unless completely blocked from administering user accounts, a **Restricted Admin.** is allowed to assign the **May lock recordings** role to non-administrator user accounts but not the **May unlock recordings** role.

How it Works

Three system-wide call sets are created automatically and made visible to users with either of these roles assigned. Authorized users have lock and/or unlock icons on their search/replay screens with which they can select one or more recordings and give a reason why they are locking or unlocking the set of recordings. Recordings to be locked are placed in the **Lock Pending** call set.

A background process retrieves these recordings from disk buffer, archive or the recorder on which they are held and places a copy of the recording, (which may consist

of multiple audio and screen files) and its XML file in a special sub-folder within the **Call Storage Path.** This folder is not purged like the normal recording folders so anything placed there remains there until the recording is unlocked.

Note:

Do NOT attempt to delete or otherwise manage the files in this path by hand. Use the Lock and Unlock features provided for you.

All lock and unlock requests and associated actions are audited as **Call Storage** actions.

The *Avaya Contact Recorder User Manual* includes instructions for end users on how to use the Lock and Unlock features. Note that the latter can only be done from call segment-based layouts.

Note:

If you recorded a voice and screen recordings in a V12.0 system or earlier, but the screen recordings retention period is shorter than the audio component of the same recording, if you want to lock the audio recordings for which the screen content may not be available, set the property **lock.acceptaudioonly=true**.

Multiple Server Systems

The locking features is only permitted on the recorder(s) that hold the consolidated database of recordings from all your recorders. Where you have installed one, this will be on the Central Replay Server. Otherwise, this will be the Master recorder - which will be providing these centralized replay services. The feature will not be available (UI disabled) on other server types.

Each recorder will forward a lock or unlock request to any other recorder to which it is consolidating its own recordings. A Master/**Full Standby** pair (in the absence of dedicated Central Replay Server(s)) will update each other automatically. However, if you have a backup Central Replay Server, it is important that each Central Replay Server is configured with the address of the other Central Replay Server (on the **General Setup > Server** page) so that it knows to forward lock related requests to the other server.

Purging the Locked Folder

It is deliberately easier to lock a recording than it is to delete one - as the latter is irreversible. Calls can be locked and unlocked by appropriately authorized users but this does not actually delete them from the locked folder. This gives the System Administrator one last chance to reverse the unlock decision if that turns out to have been made in error. You should be absolutely certain that the unlock decision was correct before purging the locked folder as described below. If a mistake has been made, simply select the relevant calls and lock them again.

Overnight, as part of the daily purging process, the recorder will raise a Warning level alarm if it finds one or more recordings in the locked folder that no longer need to be

locked. To purge these unlocked recording, use the maintenance page (at url **`/servlet/acr?cmd=mtce`**) and click the **Purge Locked Recordings Folder** button. Where locked recordings are retained on multiple servers, you must purge each one in turn.

Replay Authorization Process

You can, optionally, require specific users to obtain explicit authorization before they are allowed to replay any recording. How this process impacts users is described in the *Avaya Contact Recorder User Guide*.

Important:

The Replay Authorization process currently only supports Call Segment-based layouts. If you require a user to request authorization, you must not grant that user access to any session-based layouts.

Email Server

The process relies on email to advise users of requests and responses. You must therefore configure and test valid settings on the **System > Email Server** page as described under Email Configuration on page 158.

Configuring the Process

To enable and configure this process, on the **System > Manage Users** page, set **Replay authorization process** enabled to **Yes**. A number of additional settings will then become visible. These let you specify:

Replay authorization to be approved by

By default, a single user must authorize each replay request but, by setting this to a higher number you can require two or more users to accept a request before the recording can be played by the user requesting it. This setting applies to all users who require authorization.

Replay authorization falls back after (hours)

If those asked to authorize a request are absent or do not respond after this number of hours (default 72) then the request will be copied to those you define as "fallback" authorizers so that they can make a decision on it.

Replay authorization falls back to

You can select any of the users you have given the **May authorize replay** role to act as "fallback" authorizers. (See immediately above).

Replay authorization expires in (hours)

Once authorization is granted, the requesting user has a certain time in which to play the recording (default 48 hours). After this period, the authorization expires and he would have to request it again if he needs to play the recording again.

Configuring Authorizers

To use this process, a full administrator must give at least one user account the **May authorize replay** role and enter an email address through which that user can be contacted and identified to other users. **Add** or **Edit** an existing user account to assign this role and set their email address.

You should also consider adding at least one such authorizer to the list of fallback authorizers.

Requiring Authorization

For each user account that you want to force to use this authorization process, first ensure that the user(s) you wish to handle their requests have already been configured as authorizers (see above) and that you know the user's email address.

Log in as either a **System Admin.** or **Restricted Admin.** and use the **System > Manage Users** page to add or edit the user's account. Set their **Email address** and click the **Edit** link opposite **Replay must be authorized by**. Select one or more authorizers from the list of email addresses shown and click **Enter**.

The email addresses of the selected authorizers are shown in the table at the foot of the **System > Manage Users** page. All of the authorizers that you select here will be notified immediately of requests by this user. Any "fallback" authorizers configured will also be notified if requests are still pending after the configured fallback period.

Make sure that the user only has access to call segment-based layouts. Replay Authorization is not supported from session-based layouts.

Audit Trail

All steps in the authorization process are logged to the Audit Trail as Replay actions. Where a step involves both a requesting and an authorizing user, two entries are made, one for each. This allows you to search for all replay actions by requestor or by authorizing user.

Standby Server

Although user account settings are copied to each Standby server, the actual replay authorization requests and responses are not maintained on any Standby. Should the system fail over to a Standby, new requests would have to be issued. Similarly, when the Master is forced to take over again, any requests made on a Standby will be lost and must be re-requested.

Backup Central Replay Server

Neither user configuration nor details of authorization requests are copied to a backup Central Replay Server. You should be backing up the contents of the PostgreSQL database on the main Central Replay Server so that these details (which include all authorization requests and their results) can be restored in the event of server failure or corruption.

Modify Default Behavior

You can add or modify lines in the properties file as described in Properties File on page 265 to change the behavior of the Replay page to:

- limit the number of results returned
- save filter settings from session to session
- limit maximum duration of searches

These settings affect all users and are only updated when the Avaya Contact Recorder Service starts.

Limit results returned

Add the following line to the properties file to set the maximum number of calls returned from a query:

`viewerx.limit=`*`nnnn`*

where ***nnnn*** is the maximum number of calls to return. Increasing this maximum over the default of 100 results in more CPU use and higher network traffic if users choose to view or accidentally request a large number of calls.

Limit Query Duration

A query spanning a very long time period may take many minutes to run on a large system. When it does eventually return, it is likely that the user has given up on it and refined the timespan or his browser has timed out. There is therefore little reason to let queries run for more than a few minutes. By default, searches time out after 5 minutes. You can adjust this using the property setting:

`viewerx.timelimit=`*`nnn`*

where ***nnn*** is the timeout period in seconds. Increasing this period allows longer queries to complete but if many of these are left running concurrently, other users will have to wait longer for a database connection to free up. If you increase this setting you may also need to increase your browser's time out so that it does not time out before the results are returned.

If searches across realistic timespans are consistently very slow, please contact Avaya for tuning and performance improvement options.

Save filter settings from session to session

Normally, the entries in the search filter pane are blank when you first access it. If you wish, the application can remember the last used settings and apply these instead. This feature is controlled by a setting in the properties file: **viewerx.savesettings=true**

Backup/Restore

Due to the huge volume of new files created every day, a voice recorder is not backed up in the same way as most application servers. This section guides you through the issues around backing up the application, the call details database and the recordings.

Application

The recorder's configuration is stored in its database (using PostgreSQL), alongside the details of the call recordings. To preserve the configuration of the server, back up the database frequently as described below.

If you have not installed other applications on the server, there is no need to back up the operating system or the recorder software. It is faster to reinstall these server components in the event of disk failure. You should therefore retain the installation media and license key that you used.

Backing up the Database

You can back up your recorder's database using a command line procedure. The procedure uses the PostgreSQL pg_dump command to extract data from the database. It must be executed while the database is running. Do not stop the Avaya Contact Recorder service or the Postgresql service before proceeding.

Linux

To back up your PostgreSQL database:

1. Log on as root.
2. Become the database owner by typing **su - postgres**
3. Create a backup file by entering the command:

pg_dump --format=c --compress=5 --exclude-table=*_i
eware > backupfile

You should specify a full path for the backupfile, and consider moving the resulting backup file to external media or another machine.

The exclude_table parameter prevents the turbo-mode instantiated views ((tables `segment_i` and/or `session_i` if present) from being backed up. These are derived

from other tables and will be recreated automatically in the background if ACR starts on a restored database.

Please observe the following guidelines concerning the compression factor:

- 5 is a modest compression factor.
- Using a higher number (maximum is 9) makes the backup slower and uses more resources. However, it results in a smaller backup file.
- Using a smaller number makes the backup faster and uses fewer resources. However, it results in a larger backup file.

Windows

To back up your PostgreSQL database:

1. Log on as the Administrator.
2. Open a command window.
3. Navigate to the pgsql\bin path below where you choose to install the database.

   cd D:\Avaya\ACRDB96\pgsql\bin

4. Create a backup file by entering the following command:

   pg_dump --format=c --compress=5 --username=postgres --exclude-table=*_i --host=127.0.0.1 --file=*backupfile*eware

5. Enter the postgres master password when prompted.

You should specify a full path for the backupfile, and consider moving the resulting backup file to external media or another machine.

The `exclude_table` parameter prevents the turbo-mode instantiated views ((tables `segment_i` and/or `session_i` if present) from being backed up. These are derived from other tables and will be recreated automatically in the background if ACR starts on a restored database.

Please observe the following guidelines concerning the compression factor:

- 5 is a modest compression factor.
- Using a higher number (maximum is 9) makes the backup slower and uses more resources. However, it results in a smaller backup file.
- Using a smaller number makes the backup faster and uses fewer resources. However, it results in a larger backup file.

Restoring data to a new PostgreSQL database

Important:

You can only restore data to the server from which you dumped it because the dump file stores the software serial number and license key information. These are tied to a MAC address on the recorder. Unless you can move the original NIC into the new server, you will need to obtain a new license key if you wish to restore to different hardware.

The following process erases the default database that exists after a complete re-installation and replaces it with the database that you have backed up. The previous database is renamed rather than deleted - so if you need to restore the database again, you must first rename, backup or delete the previously retained database.

Linux

To restore the database:

1. Re-install the operating system.
2. Log on as root and install the recorder as described in Installing Avaya Contact Recorder on page 108.
3. Stop the Avaya Contact Recorder service.
4. Become the database owner by typing **`su - postgres`**
5. Drop the existing database by entering the following command: **`dropdb eware`**
6. Create an empty copy of the PostgreSQL database by entering the following command: **`createdb eware`**
7. Restore the data by entering the following command:

`pg_restore --dbname=eware --use-set-session-authorization`

`backupfile`

Note:

The restore process may raise numerous "`Error from TOC entry`" errors regarding internal PostgreSQL functions. It may also report that functions or languages already exist. These can be safely ignored.

8. Start the Avaya Contact Recorder service by running **`systemctl start acr`**
9. Test to verify the restore is successful.

Windows

To restore the database:

1. Log into the server as the Administrator.
2. Stop the Avaya Contact Recorder service.
3. Open a command window.
4. Change directory (cd) to the postgreSQL binaries directory – e.g. D:\Avaya\ACRDB96\pgsql\bin.
5. Drop the existing database by entering the following command:
   `dropdb -U postgres eware`
6. Create an empty copy of the PostgreSQL database by entering the following command:
   `createdb -U postgres eware`
7. Restore the data by entering the following command:

`pg_restore -U postgres --dbname=eware --use-set-session-authorization` ***`backupfile`***

Note:

The restore process may raise numerous "`Error from TOC entry`" errors regarding internal PostgreSQL functions. It may also report that functions or languages already exist. These can be safely ignored.

8. Enter the postgres master password when prompted.
9. Start the Avaya Contact Recorder service ("ACR152").
10. Test to verify the restore is successful.

Backing up Voice Recordings

The Avaya Contact Recorder stores voice recordings in a single partition (/calls on Linux, selectable on Windows). This partition quickly fills up with thousands of directories and millions of files. When the partition is nearly full, the recorder maintains only a tiny amount of free space on the partition by deleting batches of 100 recordings (and the directory that catalogued them) at a time, as it requires space for new recordings. This causes a huge churn of files every day.

Limitations of full and incremental backup procedures

On an Avaya Contact Recorder server, two issues make it difficult to back up voice files:

- the file size
- the rate of change of the voice recording files

Together these issues make most traditional backup strategies and Virtual Machine "snapshotting" for the voice recordings ineffective. Traditional full backups are required more frequently than normal, which wastes backup media, and incremental backups are larger than expected because of the large churn of creations and deletions. For a backup strategy to be successful, it must be easy to restore the data if necessary.

Traditional "full plus incremental" backup solutions are ineffective because these backup solutions cannot complete fully. In the event of a complete disk failure, the process restores the full backup, then the increments in chronological order. This procedure immediately overflows the disk when the restore program tries to create the increments because the partition holding the call is almost at capacity to begin with. The full plus incremental backup will fail because it runs out of disk space before it has processed the "removals" part of the procedure.

Traditional restore procedures are also ineffective. If you use this solution to review a recording that has been deleted because of age, the recorder immediately deletes any restored file as part of its disk maintenance.

Finally, traditional backup solutions often require locks on the disk while they work. This can seriously disrupt the working of the recorder.

Backing Up Recordings

This simplest and cheapest strategy is to use the built in archive mechanism. This is not only fully integrated with the workings of the recorder and its search and replay mechanism, but also is well suited to the incremental recording required for a recorder. As recordings are added to the calls path they are copied to one or more archive destinations in an efficient manner. Even when they have been deleted from the hard disk, the recorder is still able to play them because it knows which folder they are in and can replay directly from there, without an intervening 'restoration' step. If you use NAS storage to archive your recordings, the data on these centralized disks is:

- organized in a more permanent way
- subject to less "churn"

It is possible to schedule gaps in the archiving process (e.g. 1am-4am), so, if a backup process requires a disk lock, the downtime does not cause a problem with the server's operation. This scheduling feature, together with the way the archive process organizes the audio on disk, makes this data much more appropriate for traditional full/incremental backup solutions.

Distributing User Instructions

Once you have configured the recorder, you should ensure that the end users know how to use it. Some users may need to know how to use and control the recording modes. Some will need to know how to search for and replay recordings.

Those Using Recording

You will need to advise users of some or all of the following:

Mode	If you…	You should tell…	This information
On Demand Recording	Use this mode at all	All potential users of On Demand recording	The station and/or hunt group number(s) to dial to reach an appropriate On Demand recording port. How to use this mode.
	Configured any Audix-rec buttons	Users of these stations	How to use the Audix-rec button.
Meeting Recording	Configured any Meeting Recording ports	All potential users of Meeting recording	The station and/or hunt group number(s) to dial to reach a Meeting recording port (with prompts in the appropriate language).
Bulk Recording	Enabled the delete command	Users who may need to delete recordings during a call	The digits to dial during a call to have a recording deleted.
	Enabled the retain command	Users whose stations are configured with this setting	The digits to dial during a call to have a recording retained.
	Have configured delayed retention	Users who may need to retain a call after the call has ended	The number of the retain port and explain that only the previous call will be retained. They must retain the call before making another call from the same station.

Mode	If you…	You should tell…	This information
	Use Record on Demand (ROD) or Save (SAV) buttons on CS1000 phones	Users of those phones	How to use these buttons and how to interpret the lamp conditions.
	Use Agent Initiated Monitoring (AIM)	Users of workstations where it is installed	How you expect them to control and/or tag recordings.

Those entitled to replay calls

The integral search and replay application is very straightforward and the online help within it is rarely needed. If you do not distribute the manual to all users you should still advise them of the following:

- The url (http or https) they should use to access search and replay
- Their username and how to log in for the first time (unless you are using Windows authentication)
- Any tips on which data fields would be particularly useful for them to search on (especially if you have populated any user defined fields)

Configuring Avaya Support Remote Access

Refer to the instructions in the document RemoteAccess.pdf provided in the docs folder of the installation CD.

Chapter 5: Operations, Administration & Maintenance

This chapter provides details of regular maintenance required for an Avaya Contact Recorder system.

The main sections in this chapter are:

- Introduction on page 224
- Status Monitoring on page 224
- Preventative Maintenance on page 235
- Restarting the system on page 235

Introduction

In addition to initial configuration, there are a number of tasks that need to be performed on an ongoing basis. This section discusses

- the use of the Status monitoring pages
- the Audit Trail
- preventative maintenance tasks that should be carried out on a regular basis

Status Monitoring

The status of the Avaya Contact Recorder server(s) is shown on pages that are accessed under the **Status** tab at the top left of the Administration web interface. These show the current:

- **System** status - the overall status of each recorder in the system and a table showing the overall loading across the whole system.
- **Server** status - shows the status of network links between this server and the others in the system plus an overview of what has been recorded.
- **CTI Monitors** - the devices that are being observed via the CTI link(s).
- **Ports**' status - the state of each recording port in the system.

Note:

A Slave server only shows its own **Server** status, not the other three pages.

System

This page shows summary information about the current state of the system and is therefore the default page shown when you first log in as an administrator to a Master or Standby recorder. Each server that this one is directly connected to is shown at the top of the page - with Replay server(s) first, then Master and Standby servers and finally any slave servers. Within each type of server, any that have alarms are brought to the top of the (potentially long) list.

The top left of each row shows the server's static configuration, including:

- **Serial Number**
- **Role** in the system
- (Recording servers only) **Maximum load** and what proportion of the **undesignated recording load** they will normally take.
- (Standby servers only) will show whether they are "**Hot**" or "**Warm**" and their **Priority**.

The state of each server is shown as regards

- **Connection** – whether or not this server is communicating directly with the other server and which IP address(es) it is using.

When the server is connected, it also shows:

- **Recording** state – which may be **Out of Service, Faulty** or, normally, show the number of recordings currently being made.
- **Control** state (on Master and Standby servers) – may be **In Control** or **Standing By**. The entry referring to the server you are connected to may also show that it is Starting or Shutting Down.
- If a critical event has occurred recently, it may show that it will reassess its state in a specified time. You should refresh the screen after that time to confirm what state it has adopted.
- On the right-hand side, alarms are shown. Do NOT assume that a server that is **In Control** is completely healthy. It may just be less unwell than the others and hence taking control rather than having them do so.

Note:

When a Master or Standby recorder makes contact with another recorder, only "persistent" alarms (ones that are still an issue) are refreshed. Any transient alarms that occurred while the recorder was disconnected will only be visible on that recorder's Alarms page.

Buttons shown against each server let you:

- Jump to the **Admin Pages** of a server to drill down into any problems that are showing.
- **Clear Alarms** from this server's summary display.
- **Clear Persistent Alarms** from this server's summary display.

And, if logged in as a full Administrator, you may also see:

- **Take Out of Service**:
  The server that is currently **In Control** will show this button against every other recording server that it is in contact with. If you click it, no new recordings will be started on that server. Existing recordings will continue until the end of the call.

  As long as there are other servers available to handle all the recordings normally assigned to the one you take **Out of Service**, this feature allows you to stop, patch and restart that server during business hours - without losing any recordings. The procedure is as follows:

  1. Take a server **Out of Service** by clicking this button.
  2. Use the **Status > Ports** page to see when the current recordings on that server's ports have all ended.
  3. Shut down the recording service and patch the server as normal.
  4. Restart the recording service. This automatically returns it to service and allows new recordings to be assigned to it.

Notes:

1. The server that is **In Control** cannot, itself, be taken **Out of Service** as this would inevitably interrupt current recordings.
2. Only use this feature when the system is completely healthy. A system split into isolated partitions will have inconsistent views of which servers are in service.
3. You must either restart an **Out of Service** recorder to bring it back into service or click the **Return to Service** button (which replaces the **Take Out of Service** button once that has been clicked).
4. If you have taken the Master recorder **Out of Service**, be sure to restart it or explicitly return it to service before attempting to have it **Take Control** (see below).

- **Take Control**: This button is only shown on a Master that is currently **Standing By**. This *will* interrupt recordings and **should only be done out of hours.**

Notes:

1. If you have taken the Master **Out of Service**, do NOT click this **Take Control** button to bring it back into control until you have either restarted it or explicitly brought it back into service and confirmed all servers are connected and healthy.
2. If you *do* attempt to have the Master **Take Control**, or if any server that is in control makes contact with another server that is also **In Control**, this will automatically cancel all **Out of Service** flags – as these may be inconsistent.

Load

On the Master (or Standby if it is **Active**), the table at the bottom of this page shows the current and peak loading levels of the overall system that the server is controlling. For each recording mode configured, it shows:

- The number of ports allocated.
- How many are currently active.
- The maximum number that have been active concurrently since midnight.
- The maximum number that have been active concurrently since the date and time shown above the right-hand column.

To reset the monitoring period, select the **Restart Peak Activity Count** button. For example, if your business has a weekly cycle, you may want to reset the monitoring period at the start of each week. Use this page to predict when to expand your recording capacity.For more detail of contention and "busy" events on the hunt groups associated with the pooled recording modes, use the Communication Manager's call detail recording tools. These tools might be useful if, for example, you are required to provide 98% availability of Meeting Recording ports.

Server

This page shows the status of the server displaying it. It includes a row for each network link that the server requires to function as well as basic counts of calls recorded.

See Recorder Interfaces on page 302 for a comprehensive list of interfaces.

The following items are not relevant to and hence are not shown on Central Replay Servers.

Total calls observed via CTI since startup

Shows how many calls have been tracked via CTI information (whether recorded or not) since the recorder was last started. Where overlay CTI feeds are used, the same underlying phone call may be advised to the recorder more than once (e.g. from TSAPI and via CCT). In such cases each appears as a separate "call" in this count.

Total Media Files recorded to date

This value shows the total number[1] of files recorded by this server and any others that treat it as a Replay Server. Note that each recording in the database may result in more than one file. As this takes some time to calculate on startup, it may not be available for a few minutes.

Total calls observed via CTI today

Shows how many calls have been tracked via CTI information (whether recorded or not) since midnight (or since the recorder was last started if that was today). (As with the "to date" figure, multiple CTI feeds can make one underlying call appear as multiple calls).

Total Media Files recorded today

Use this value to confirm that recordings are being made today. If you have restarted the server today, this will show the number of media files recorded by this server and any others that treat it as a Replay Server since that restart. As above, there may be more than one media file per recording.

Date of oldest call held on disk

Until your disk has filled for the first time, you should monitor the available space on the drive. Check that the rate at which space is being consumed is in line with your predictions. You should be able to estimate when the disk will reach capacity and when

[1] Recordings already purged from the database when 15.1FP2 was installed will not be included. Audio recordings made prior to 12.1 count as 1 regardless of the number of associated screen capture files.

the first calls recorded will be deleted to make way for new calls. This occurs when the free space drops below 1GB.

Once the disk has started to "wrap" and calls are being deleted daily, use this figure to monitor the online retrieval capacity. If the figure starts to fall, recordings are using up your disk space more rapidly than before. The recording volumes are increasing, so you may need to expand the disk capacity before the duration of calls it can hold falls below your minimum requirement.

Server Status

On any recorder that is acting as a central replay server (i.e. a dedicated Central Replay Server or a Master or Standby on a multi-server system that does not have a dedicated Central Replay Server), this page also shows the status of each other server that is contributing recordings to its database.

Where contact has not been made for more than 2.5 minutes (at least 2 consecutive updates failed) this is considered an inactivity problem and will show in red and raise an alarm at a configurable level.

Where the latest recording is more than a configurable threshold old, then this is considered a recording interval problem and similarly raises an alarm. These thresholds are set according to the day of the week and hour of the day. As with all alarms, these can be sent via Email and/or SNMP as well as viewed on the CRS's Alarms page.

A server will not show as a row on the status page at all until it has made contact with the replay server (which it should be doing at least once a minute). It is therefore important to check that all (Master, Standby and Slave) servers *are* showing and to wait to see ALL appear following a restart of the replay server.

Following a restart, each server that makes contact may show in red until a recording has been received. Although this may be out of hours, a test call should be made to each server following a restart anyway and it is better to flag that no recordings have been received than to give a false impression that all is well when it may not be.

The recording count is only approximate as the number of times each recording is sent to the replay server depends on how the call is handled and tagged.

The inactivity threshold is checked by the recorder once a minute - or when the status page is viewed. Since there is then a 5-minute window in which other alarms are collected before being emailed, this results in alarms being sent within 10 minutes of a blockage (5 minute recommended threshold + 5-minute email consolidation buffer). Any lower than 5 mins is likely to trigger a lot of false alarms for little benefit (6 mins rather than 10mins to alert).

Configuration

Basic status reporting requires no special configuration. Setting thresholds to trigger a red background based on time since last recording end is done as follows for each day of the week. Set a property of the form:

inactivity.day_n=n0,n1,n2,n3,n4,n5,n6,n7,n8,n9,n10,n11,
n12,n13,n14 ,n15,n16,n17,n18,n19,n20,n21,n22,n23

where:

- ***n*** is the day of week (1-7 as per locale of server)
- ***n0*** through ***n23*** are how many minutes to allow in each hour of the day since the previous recording before showing a server's status in red. (Or set to 0 to suppress recording interval alarms in those hours)
- Trailing zeros can be omitted (turning off alarming at the end of the day
- An entire day may be omitted (no recording interval alarms will be raised that day)

For example, a site in the UK (where Sunday = day 1) taking a trickle of calls from 6am but not really ramping till 9am and tailing off between 5 and 6pm on weekdays; late opening on Friday; shutting at lunchtime on Saturdays and no traffic on Sundays might configure:

i**nactivity.day_2=0,0,0,0,0,0,0,30,20,10,10,10,10,10,10,10,20,30**
inactivity.day_3=0,0,0,0,0,0,0,30,20,10,10,10,10,10,10,10,20,30
inactivity.day_4=0,0,0,0,0,0,0,30,20,10,10,10,10,10,10,10,20,30
inactivity.day_5=0,0,0,0,0,0,0,30,20,10,10,10,10,10,10,10,10,2
0,30
inactivity.day_6=0,0,0,0,0,0,0,30,20,10,10,10,10,10,10,10,20,30
inactivity.day_7=0,0,0,0,0,0,0,30,20,10,10,20,30

Other configurable parameters are:

- **inactivity.alarmlevel=*n*** where ***n*** is 0 for INFO, 1 for WARN, 2 for MINOR and 3 for MAJOR alarm levels (default i.e. not present = no alarms)
- **inactivity.alarmlevel.nnnnnn=n** where ***nnnnnn*** is a recorder serial number overrides the alarm level for that recorder with the level specified here (same values as the overall level above)

CTI Monitors

This page is only populated on Master and Standby recorders (not Slaves or Central Replay Servers). It shows each of the phone numbers that the recorder needs to monitor via the main CTI link. Note that not all recording modes use this CTI link so it may not be populated.

The detail shown and terminology used varies slightly from one telephone system to another.

Device

Shows each device that needs to be observed via the main CTI link.

Observing

Shows whether or not the device is being observed successfully. In the case of a CS1000, there may be mutiple rows showing against a single DN or position ID. Each row shows the DN or positionID and the physical TN on which the number occurs.

Part of Range

This field shows which number or number range that has been configured for recording has resulted in this particular phone number being monitored. Any advanced settings that apply to this range will determine how calls on this phone number are recorded.

Agent

If an agent has logged on to this device, this field shows that agent's ID.

Active Calls

For all calls that this station is connected to, this column shows:

- the call type and ID
- whether the call is considered to be incoming or outgoing
- whether the call is considered to be internal or external
- the Automatic Number Identfication (ANI - also known as Calling Line Identifier, CLI) if present

Each party on the call is then listed - showing

- the state of the connection (held, ringing etc.)
- whether this party is an internal or external party

If a call is shown in bold typeface, it implies that the recorder believes that this particular call is active and therefore may need to be recorded.

Note:

The details about each call are not translated. The detail within this information is of most value to second and third line support staff who will be familiar with the terse English abbreviations used to convey a lot of information in a small space on this screen.

Port

When a call is active and needs to be recorded, a port on a recorder is assigned to it. The name includes both recorder and port identifiers.

Observer(s)

Shows any parties that are observing this device. For example, this may include a live monitor client.

Rec

This column shows whether or not the channel assigned is recording the call at the moment. In the case of external controller or dialer integrations, a channel may still be assigned but not actually recording all the time. If the recording includes screen content, this is shown by a screen icon for each recording that is being made. Hover your mouse cursor over the icon to see the IP address being recorded.

Events

This column shows how many CTI events have been observed on this CTI monitor. To update the page, click the Refresh button.

Ports

This page shows the current state and configuration of each port in the system that is visible to this server. It therefore includes the server's own ports plus the ports on any slave to which it is connected. Note that channels are dynamically assigned and released as calls come and go. The columns show:

Recorder/Port

Identifies each port in the system.

Assigned To

Shows the recording target(s) that the port is being used to record.

Recording

Shows whether or not the port has been told to record active calls. If the recording is being deliberately masked this will also show here. If the recording includes screen content, this is shown by one or more screen icons. Hover your mouse cursor over the icon to see the IP address(es) being recorded.

State

Shows the current state of the port. Possible states are:

Faulty (DMCC ports only)

These ports are incorrectly configured or have experienced an error. They will be reregistered two seconds after the initial problem, in an attempt to recover them. If the problem persists, they will back-off, doubling the time between retries until this reaches 1 minute. They attempt to reregister every minute thereafter.

Starting (DMCC ports only)

These ports are registering or queuing to register with the Device, Media and Call Control API.

Idle

These ports have registered successfully, but are not in use.

Setup

These ports are placing or answering a call or receiving instructions from the caller, for example, with Meeting Recording.

Connected

These ports are in one of the following states:

- Have established a connection but are not actively recording
- Are replay ports that have placed a call but are not currently playing a file

Active

These ports are:

- Recording ports that are recording a call
- Replay ports that are actively playing a file

Stopped

A recording port may be in this state if it has stopped recording but has not hung up yet.

Resetting Ports (DMCC ports only)

To reset an individual port, select the **Reset** link to the right of the port. This action aborts any current recording attempt on that port allowing the recorder to attempt to restart recording on another port if possible.

To force a reset on all ports in quick succession, select the **Reset All** button at the bottom left.

Alarms

The Alarms page shows system warnings, alarms and events. The recorder stores alarms in its database. It deletes them when they are more than a month old.

Check the box at the top of the page if you want to see all Alarms, including those that have been "cleared". The default is to show only those alarms that have not yet been cleared.

The radio buttons at the top of the page let you select the minimum severity of alarms that are shown. The Alarms page will only show the most recent 1000 alarms (in 100 pages of 10 alarms each). If there are Informational or Warning level messages, you may need to restrict your view to Minor or Major alarms if you want to see further back in time.

Use the buttons above or below the table to refresh the page and to clear some or all alarms.

Note:

As long as you have set up mail account information on the **System > Email Server** page, an email message detailing alarms and events will be sent to the address(es) listed there.

Refer to Alarms on page 353, for a list of alarms and events that may be generated and what to do about them.

Viewing alarms and events

The default on the page is to show all alarms and events that are uncleared. You can see any new or outstanding issues on first viewing the page. To change the set of events shown:

1. Click on the check box and/or radio button to specify your preferences.
2. Click on the **Refresh** button above or below the table.

Clearing specific events

New alarms and events are initially "active". To "clear" an individual alarm or event so that it no longer shows:

1. Click the check box to the left of the event.
2. Click the **Clear selected events** button.

Clearing all events

Be careful using the **Clear All Events** button. Clearing an alarm without fixing the problem may lead to system problems being "hidden" without your knowledge.

Audit Trail

The **System > Audit Trail** page shows administrator and user actions over a specific period. The default reporting period is the current day. You can also filter this report according to Event Type and Username. To generate a report for a different period, enter the date range in the calendar controls, and click the **Refresh** button.

The Audit Trail functions track the following user actions:

- Successful user logins
- Failed user logins
- Password changes (although, for security reasons, the actual password is not stored)
- End user searches on the database

Replay requests (including steps in the replay authorization process) Call storage actions (lock, unlock, purge and delete).

It also tracks all administrator actions that affect recording, such as configuration changes, manual port resets and creation or deletion of user accounts.

Note:

Editing a station range is logged as a deletion followed by an addition.

The Detail column includes the SQL statement used in searching for calls. It also uses the internal name of a setting rather than the user-friendly, localized name. This avoids any change of meaning that could occur in internationalization.

Each report is restricted to a maximum of 1000 audit records. To report on more, break your reporting period into a number of smaller date ranges.

You can either click the **Export** button to obtain a file in comma separated variable (".csv") format or use your browser's print, save, or email features to provide a permanent record of the details. To create a summary that presents all results on a single page, click the **Show All** link at the top. The **Show All** and **Page at a Time** links are not shown if the list of audit entries is less than one page long.

Configuration records (which include the audit trail) are retained for three months by default. Each night after that period has elapsed, a background job deletes any records older than the retention period. If you want to retain the records longer, back up the database as described in Backing up the Database on page 216. You can change this default value of three months using the **`audit.purgemonths`** property as described in Properties File on page 265.

Preventative Maintenance

This section highlights a number of administrative tasks that should be performed on a regular basis to ensure the system continues to operate smoothly.

Daily

Unless you have fully automated alerting of these conditions, you should carry out the following procedures at the start of each day:

Alarms

Check the Alarms page for new problems.

Disk Capacity

Check the available disk space. The disk where recordings are stored will appear to be at or near capacity. However, the system consistently maintains a level of 1 GB of free space by deleting older files. This maximizes the number of recordings that are available online to you.

The Avaya Contact Recorder's disk manager thread deletes files on a FIFO (First In First Out) basis.

System Status

It is difficult to detect some problems automatically. Check the system status regularly via the **Status > System** and **Server** pages and verify that all figures are in line with expectations as described in Status Monitoring on page 224.

Check the contents of the log files as described in Where to Look for Clues on page 339 and examine any errors logged since the previous check. Look at all error and warning messages, not just those generated by the Avaya Contact Recorder services.

Confirm Port Status

Use the **Status > Port** page of the Administration application to confirm that the recording ports are in the appropriate states.

Confirm Recording and Replay

To confirm recording and replay:

- Verify that calls are being uploaded into the database.

Use the **Replay** page to select the most recent calls to verify that calls are accessible.

Confirm that the start time of these calls matches expectations. Verify that the start time corresponds to the most recent calls made on the extensions being recorded.

- Confirm that these calls are playable and that audio quality is good.

Archive

Check the available space and increase it or add additional folders as needed.

Weekly

As you become comfortable with the normal operation of your recorder, you can reduce the frequency of the daily tasks. For example, if you know that the rate at which your disk is filling is not going to fill the available space for several months, you can check it weekly.

Perform the following tasks each week:

Disk capacity: main recording store

When your recorder is first installed, the disk is almost empty. As it gradually fills, you should note the rate at which it is being used (at least weekly) and extrapolate to estimate when the disk will be full. At this point, the Avaya Contact Recorder will begin deleting the oldest calls to make room for new ones. If this happens to calls that are younger than planned, check the configuration of the recorder to ensure that only the anticipated calls only are being recorded. Add additional disk capacity to the partition before it fills.

Disk capacity: other partitions

Check the available space on any other disk partitions. Verify that these other drives have sufficient space. The recorder will warn you if they fall below 500MB of free space. Accumulated temporary files or log files can account for this drop in available space. You may need to purge them manually.

CAUTION:

When you are purging files on Windows, remember that files you delete go to the Recycle Bin and that the space they occupy is not freed until you empty it.

Call detail database purging

If you have enabled automatic purging of aged call detail records, you should still monitor the size of the calls database during the first few months of use. You can then predict how large the database will get by the time old records begin to be purged. Many customers plan never to purge call detail records, but choose instead to add disk capacity every year or two as the database grows. If you do this, you should upgrade your server every few years to compensate for the increasing size of the database and the reduction in search and update speed.

Configuration Backup

Changes to system configuration that affect user access rights are stored in the PostgreSQL database. This means that the system configuration is backed up whenever the call detail records are. See Backing up the Database on page 216.

Monthly

Check the following aspects of the system on a monthly basis:

Loading trends

Note the total call volumes recorded every month to be aware of gradually increasing traffic trends. To do this:

- Note the number of calls recorded at the end of each month and compare with previous month's accumulated total.
- Note the age of the oldest call on the disk (only applicable once the disk has filled for the first time).
- Note the CPU load during busy hour.

If it appears that the load is increasing, consider purchasing extra licenses if required and/or increasing server specification or disk space.

Restarting the System

Occasionally, you will need to restart your Avaya Contact Recorders.

Linux

To start, stop, restart and monitor the Avaya Contact Recorder service use the linux **`service`** command. To start, stop, restart and monitor the ACR service use the following commands:

```
systemctl start acr
systemctl stop acr
systemctl restart acr
systemctl status acr
```

You will be prompted for the root password.

Windows

To start, stop, restart, and monitor the Avaya Contact Recorder service use the Services Microsoft Management Console. The service name is `ACRService`.

Be patient

You should always make some test calls following a restart. However, it can take many minutes for a large system to register all its CTI observers. Use the Status pages to confirm that ports are available and that the appropriate CTI monitors have been created before attempting to place test calls.

Do not just keep restarting the system. Check the status pages to see if registrations are progressing and wait for them to complete.

CS1000 Agents

Remember that, in CC 6.0, Call Center agents must log out and log back in again before their status can be determined. If an agent logs out while the system is restarting, the system will not recognize that they logged out; it will still show them as logged in and as a result, may continue to try and record against that agent.

Chapter 6: System Security

Security of customer recordings is very important. This Chapter discusses the various features - some optional - that you can use to ensure the safety and integrity of recordings. This chapter assumes that you have suitable firewall, and physical access procedures in place (but refer to Antivirus and other 3rd Party Applications on page 64 for guidance on these topics).

The main sections in this chapter are:

- Access to the Recorder on page 240
- Server Hardening on page 249
- Changing Passwords on page 251
- Encrypting Connections on page 252
- Encrypted File Storage on page 259
- Finger-print Validator on page 259
- I/O Jobs on page 260
- PCI Compliance on page 262

Access to the Recorder

When the system is installed, it enforces a strong set of password criteria. This and a number of other settings can be changed at the top of the **System > Manage Users** page.

These are discussed below.

Note:

If you change any password strength policy settings, be aware that, because passwords are not stored in plain text, the recorder cannot retrospectively determine whether or not an existing password meets the new requirements. Consider forcing all existing users to update their password by resetting each to a new (unique) temporary password.

Windows Domain Authentication

You can create local user accounts within the recorder application itself. However, it is more secure to use Windows domain accounts and you may wish to enable this feature - or even restrict access so that *only* windows domain accounts have access to the system.

In an environment with a Central Replay Server you may choose to enable Windows Domain Authentication only on the CRS (which typically has lots of users) and stick with local user accounts on the recorders (which typically have few users).

To ENABLE Windows Domain Authentication, you must:

- Create an Active Directory account for the recorder(s).
- Configure each Avaya Contact Recorder server to use Active Directory.
- Add the appropriate domain users' accounts to ACR and assign each one appropriate role(s) within the recording system.
- Configure these users' browsers appropriately and test.

You may then, optionally, enforce access via Active Directory ONLY

Create User Account for Avaya Contact Recorder System

The Avaya Contact Recorder system itself needs to be recognized by the Active Directory in order to use it. If you have multiple ACRs they can share this single username. Create this as follows:

1. Create a user in the Active Directory for ACR (e.g. ACR@BIGCO.COM).
2. Set the password on the account, setting it to never expire.
3. Make sure that the account is enabled.
4. Confirm that the Active Directory is time synchronized.

Configure Each Avaya Contact Recorder

To let an Avaya Contact Recorder server use the Active Directory:

1. In your Active Directory, first check that the SPN you intend to create does not already exist by using this command to see the existing SPNs:

   setspn -L

2. Create a unique SPN for this server with the command:

   setspn -A HTTP/*RecorderFQDN username*

   where ***RecorderFQDN*** is the recorder's fully qualified domain name and ***username*** is that of the user created in step 1 of the previous section. For example:
   setspn -A HTTP/acr1.bigco.com BIGCO\ACR

3. Confirm that the server is time synchronized.

4. Set the Kerberos settings of the server using the command shown below, where:
   server is the IP address or FQDN of the Active Directory
   realm is the Active Directory realm in all **uppercase**
   username is the username added in the previous section

Linux Server

From a terminal session running as user "witness", change directory ("**cd**") to **/opt/witness** and enter

jdk8/jre/bin/java -cp lib/*
com.swhh.os.KerberosSettings *server realm username*

For example:

jdk8/jre/bin/java -cp lib/*
com.swhh.os.KerberosSettings ad1.bigco.com BIGCO.COM
ACR@BIGCO.COM

Windows Server

From a command window, change directory ("**cd**") to the installation directory and enter:

jdk8\bin\java -cp lib*
com.swhh.os.KerberosSettings *server realm username*

For example:

jdk8\bin\java -cp lib*
com.swhh.os.KerberosSettings ad1.bigco.com
BIGCO.COM ACR@BIGCO.COM

5. Enter the password for the ACR user when prompted.

6. The tool produces several pages of diagnostics, but if it succeeds the last line output will be:

```
Success - Kerberos configuration tested and stored
in database
```

7. Restart the recorder to have the new settings take effect.

Configure Users’ Roles

For each Active Directory user that you want to give access to Avaya Contact Recorder, you must create a user account and assign roles and access rights as described in

Securing the System on page 124).

Use all UPPERCASE in the user names that you enter into ACR. For example, **JSMITH@BIGCORP.COM**.

By default, the recorder forces usernames it receives to uppercase before comparing them with those you have entered. To change this, you can set **krb5.exact=true** in the properties file - but you *must* then enter usernames and domains with *exactly* the same cases as stored in the Active Directory.

Browser Configuration

The instructions below refer to Internet Explorer V11. If using an alternative browser, you must configure its equivalent settings.

On your users' browsers, assuming they are already logged in as a domain user that you have configured to have access to Avaya Contact Recorder as above:

1. Ensure that the URL of each Avaya Contact Recorder they will need to access is added to the **Security > Local Intranet Zone** via the **Sites > Advanced** button.
2. Confirm that the **Custom Level... User Authentication > Logon > Automatic logon only in Intranet zone** option is checked.
3. Access the recorder using its Fully Qualified Domain Name. For example:

http://*acr1.bigco.com*:8080/

You should get straight to the recorder main page. If you see the recorder login page, something has gone wrong. Check the recorder log for errors; check your Active Directory's reporting and configuration and review all previous steps.

Tip:

If users are prompted for their domain passwords when they access the web interface, make sure that the recorder is either part of the intranet zone, or make it a trusted site and configure Internet Explorer to automatically log on to trusted sites.

Enforce Active Directory Authentication

To enforce Windows Domain Authentication only (blocking local recorder accounts):

1. Enable Windows Domain Authentication as above.
2. Log in to the Avaya Contact Recorder using a domain account that you have given ACR Administrator rights.
3. On the **System > Manage Users** page, set **Allow local user accounts?** to **No**.

Single Login

When using local accounts, you can restrict users to a single session if required. To do this, add the following line to the Properties file as described in Properties File on page 265.

```
acr.singlelogin=true
```

Dual Sign-in

Security can be enhanced by forcing some (or even all) users to sign in with the aid of an additional user. This is sometimes known as a "four eyes" approach.

How it Works

This works by presenting users with a second log-in screen asking them to have an "authorizing user" enter their username and password to validate their attempt to log in. Only after the second set of credentials has been checked can the user access the application.

CAUTION:

The secondary or "authorizing" user account must be a local (not Windows Domain) account and must have replaced their original (temporary) password with one of their own choosing before they can authorize others.

Applying this Mode

To require a user to have a second, authorizing user help them log in, you must first assign at least one user the **May authorize logins** role. You will then be able to assign the **Login must be authorized** role to other user accounts.

Making this the Default

This "four eyes" approach can be made the default for new user accounts that are created. To do this, set the following entry in the properties file:
`foureyes.default=true`

Audit Trail

The normal audit entry that records a user login now also includes the name of the authorizing user. Audit records relating to subsequent actions of that user do not explicitly mention the other user.

Using this Mode

You can use this mode in two ways:

1. If an authorizing user signs someone in and then leaves them alone, you have improved the security of login (access) but not restricted or deterred the user from accessing areas of the application and/or recordings that they do not need to see.
2. If the authorizing user signs someone in and then observes their actions up to and including a deliberate "Log Out" you can deter users from doing or replaying more than their job requires them to do.

The second approach can be particularly effective with administrator accounts - as it is difficult to restrict their actions (someone has to be able to change things). As long as the observer is vigilant - ensuring, in particular, that the administrator has not removed the requirement for a second login on their own account - security can be significantly enhanced as a rogue individual can no longer do as they wish. This obviously makes option (a) inappropriate for administrator accounts.

Use of SSL

You should consider whether you wish to enforce the use of Secure Sockets Layer (SSL).

By default, users can access the recorder via HTTP (on port 8080) or by encrypted HTTPS (on port 8443). You can force users to use the secure https port, by setting **Allow unencrypted (http) access?** to **No**. When you do this, any user who attempts to access the recorder through the unsecured (HTTP) route is automatically redirected to the secure (HTTPS) address.

Note:

If you have more than one server and you enforce https, you must do so on ALL servers. The way in which recorders forward requests to each other must also be changed to use the fully qualified domain name of the other server(s). To do this, add a property to the properties file as follows: **fqdn.*nnnnnn*=<*fqdn*>** - where ***nnnnnn*** is the serial number of a server and **<*fqdn*>** is the fully qualified domain name that it uses.

If you created a self-signed certificate when installing the recorder, the browser will warn your users that the name on the certificate does not match the name of the server using it. You can either advise your users that this is acceptable and should be ignored or, for greater security, you may acquire and install your own SSL certificate as follows:

1. Backup the existing certificate as described in Backing up the Keystore files on page 404.
2. Create a new certificate as detailed in Creating a new Certificate on page 404 - replacing **<*keystorename*>** with **tomcat.p12** and **<*alias*>** with **tomcat**.
3. Restart the Avaya Contact Recorder service.
4. Access the Administration pages via **HTTPS**.
5. Check that the certificate matches the information entered.
6. Double click the padlock icon. Internet Explorer should warn you that the certificate is unsigned. However, it should no longer display a message that indicates the certificate does not match the web server name.

Tip:

If you do get a warning that the certificate does not match, check that the Common Name matches the URL. Double click the padlock, select the details tab, and click the Subject line. This displays the Common Name.

7. Have the certificate signed as described in Generating a Certificate Signing Request on page 405.
8. Import the signed certificate as detailed in Importing the CA's certificates on page 405.
9. Restart the Avaya Contact Recorder service.
10. Access the Administration pages via **HTTPS**.
11. Double click the padlock icon and confirm that the warning is no longer present.

Allow search and replay from this server?

If the details of the recordings made on a particular server are consolidated into the database on another server (Master, Standby or Central Replay Server) you may want to block users from replaying calls directly from this server - forcing them to access a server that has access to recordings from all servers. To do this, set **Allow search and replay from this server** to **No**.

Session Inactivity Timeout

You can specify how long a browser session will remain active before the user is asked to log in again. The effect of this is only noticed with local account access as when using Windows Domain authentication, the recorder simply requests authentication rom the domain again.

Minimum Password Length

This applies to local accounts only - not windows domain accounts. This defaults to 8 characters but can be increased or reduced to match your corporate standards.

Force strong password

This applies to local accounts only - not windows domain accounts. By default, the recorder forces all passwords to be "strong" - which means that they must contain at least one of each of:

- Uppercase characters • Lowercase characters
- Digits
- Special characters (#,@,%,! or $)

Changing this setting is not advised.

Password expires after (days)

This applies to local accounts only - not windows domain accounts. By default, passwords expire and must be changed on the next login that is 90 days after they were set. You can change this period to match your corporate standards.

Password cannot be reused within (days)

This applies to local accounts only - not windows domain accounts. By default, you cannot change your password to one that was originally set (rather than expired) less than 180 days ago. You can change this period to match your corporate standards.

Minimum changes between reuse of same password

This applies to local accounts only - not windows domain accounts. By default, you cannot change your password to one that is the same as any of your previous four passwords.

You can change this count to match your corporate standards.

Repeated Login Attempts

After a number (default 10) of failed attempts to log in within the space of a predetermined period (default 60 mins), an account is locked out for a pre-determined period (default 60 minutes or until recorder is restarted, whichever is sooner).

These parameters can be altered via the properties:

`acr.login.lockout=nn` `(default 10)`
`acr.login.lockmins=nn` `(default 60)`

A delay of 2 seconds is imposed before the login page is re-presented so as to throttle the rate at which a brute force attack can proceed.

Legal Notice on Login Page

You can display additional text (such as “Only authorized users...”) on the login page if you wish. Place the text in a file called `acr_loginwarning.html` in the properties folder (`\properties` under the installation path on Windows, or `/opt/witness/properties` on Linux). This can include basic formatting such as italics, bold etc. using HTML tags.

Disable Autocomplete

If not already blocked by your corporate policy, you should consider disabling Auto Complete to make it harder for someone to learn from previous users’ logins.

For Internet Explorer 11, click the **Tools** cog (top right) then **Internet options > Content > AutoComplete > Settings.** Uncheck **User names and passwords on forms**, and, if desired, **Forms**.

Replay Authorization Process

If you enable this setting, you can force some or all users to request authorization before they are allowed to replay calls. See Replay Authorization Process on page 213 for a full description of this feature and its parameters.

Server Hardening

The approach to server hardening for ACR is similar for Linux and Windows and consists of:

- running with minimal account privileges
- removing and/or disabling unused software to reduce the attack surface
- (optionally) blocking access to unused IP networking ports

The following sections provide recommendations for Linux and Windows.

Linux

User Accounts

ACR runs as an unprivileged user. The start-up script runs as root, but launches the recorder application under the 'witness' user. Almost all maintenance (e.g. gathering log files, editing properties files, patching, etc.) of the recorder should be carried out logged on at the witness user. Root access should only be needed for:

- initial installation of the RPMs
- starting and stopping the service

Never log on as root unless needed. Do not install patches as root.

Tip:

To start the recorder when logged on as witness type:

```
systemctl start acr
```

and enter the root password when prompted.

This correctly makes use of root privileges only on a per use basis.

Unused Services

You should disable or remove any services that are not used. The recommended starting point for this is a kickstart installation with the minimum package selection. Even this will install some RPM packages that are not needed and may increase the attack surface.

Peer to Peer Protocol (PPP) is required by Avaya's Remote Access Tools and hence is installed. You may wish to disable this if not using remote access.

Firewall

You may optionally enable the Red Hat firewall. ACR requires several ports to be open to operate correctly. See ACR Firewall ports on page 250 for a table of ports that must be opened depending on your configuration.

Windows

You should consider using Windows Specialized Security - Limited Functionality (SSLF). This disables most optional services on the server. SSLF is typically used by organizations where security is more important than functionality. When using SSLF you should make an exception for Remote Desktop connectivity.

Use the group policy editor to design a template for your ACR servers based on the SSLF template. You will need to open ports in the Windows Firewall (See ACR Firewall ports on page 250 for a table of ports that must be opened depending on your configuration). Unfortunately the group policy editor does not allow for port ranges. You will probably find it necessary to "allow local port exceptions" in the group policy editor and then use the Windows Firewall configuration screen on each ACR to open the ports required for that server. (The local configuration screens allow for port ranges to be entered.)

ACR Firewall ports

You should ensure that a firewall is enabled and block all ports, bar those below.

Port	Protocol	Use
8443	TCP	HTTPS admin access
123	UDP	NTP
1209	TCP	Inter-server TLS. Open on Master and Standby
10000-19999	UDP	RTP

Other ports are needed if your Avaya Contact Recorder is connected to other system components. These are described in Summary on page 309.

Changing Passwords

The recorder and related applications use a number of user account settings that are installed with a hard-coded default. You may change these as described below but be sure to note the new passwords securely as remote support staff will need these to maintain your system. Should you lose these passwords, your system will become completely unmaintainable.

User Accounts

Linux

The root and witness user accounts are installed with default passwords. Change these by logging on to the server and typing **passwd**.

Windows

The installation process does not create any user accounts. The PostgreSQL and recorder services run as Network Service.

PostgreSQL Database Owner

You may wish to change the password used by the recorder to access its local PostgreSQL database. Before doing so, you must obtain the Avaya encryption tool as described in Encrypting Properties File entries on page 415. You can then configure the recorder to use an alternative password as follows:

1. Choose a new password.
2. Encrypt the new password using the WitsBSUserCredentials tool.
3. Add the new password (in encrypted form) to the acr.properties file as **db.password=*encryptednewpassword***

Then follow the appropriate procedure below to set that password within PostgreSQL before restarting the recorder to have the above setting take effect.

Linux

Follow the steps below:

1. Log in as root.
2. Switch to the PostgreSQL user account by typing**su - postgres**

3. Access the database by entering:

   psql

4. Change the password by entering:

   alter user eware with encrypted password '*newpassword*';

5. Quit the database by entering:

   \q

Windows

Follow the steps below:

1. Log in as an Administrator.
2. Open a command window and "cd" to the \postgresql\bin directory beneath the recorder's install path.
3. Access the database as the recorder would by entering:

   psql -U eware

4. Enter the default password (must be obtained on request from Avaya at the time and not written down)
5. Change the password on the account that the recorder uses by entering:

   alter user eware with encrypted password '*newpassword*';

6. Quit the database by entering

   \q

Encrypting Connections

The Avaya Contact Recorder establishes connections with numerous other systems (as described in Recorder Interfaces on page 302. Some connections pass control instructions and metadata while others contain the actual content of telephone calls or screen recordings. Many links are encrypted by default but some require deliberate action. Where links are not encrypted, you should take their routing into account when designing your network topology.

Computer Telephony Integration

Communication Manager

TSAPI, DMCC and SMS links are encrypted by default.

CS1000

The MLS link does not support encryption.

AACC

The CCT link may be encrypted. See CCT Web Services on page 98.

Avaya OnmiCenter

This will use encryption by default (HTTPS).

PCS Dialer

This will use an encrypted connection by default.

POM Dialer

The POM link supports encryption with POM 3.0.5.

Session Initiation Protocol (SIP)

This can be used to record calls that pass through a Session Border Controller (SBC).

To encrypt these SIP interactions - using SIP/TLS - you must create (or re-use) a server certificate for your Avaya Contact Recorder as follows:

1. Create a new certificate as detailed in Creating a new Certificate on page 404 - replacing **<*keystorename*>** with **siptls.p12** and **<*alias*>** with **sip**.
2. Have the certificate signed using the Certificate Authority already used by the SBC or AACC as described in Generating a Certificate Signing Request on page 405.
3. Import the signed certificate, together with the Certificate Authority's root certificate as detailed in Importing the CA's certificates on page 405. Do this on each Avaya Contact Recorder.

After the next restart, ACR will use SIP over TLS for SIP messages.

Audio Streams

This varies according to the switch and type of recording being performed.

Communication Manager

DMCC Recording

Use the **Media Stream Encryption** setting on the **General Setup** page for your Communication Manager to select an encryption mode. The older **aes** mode is deprecated and should not be used if your system supports XXX mode (from AES Version XXX onwards).

Passive IP Recording

Unfortunately, for passive IP recording to work, the RTP streams from the phones being recorded (and mirrored to the recorders' Network ports) must **NOT** be encrypted.

SIPREC Recording

Avaya Contact Recorder will automatically use encrypted media streams if your SBC is configured to do so. As the keys for these streams are exchanged via the SIP signaling path, the recorder will alarm periodically if this has not also been secured using SIP/TLS.

CS1000

The RTP streams used in Duplicate Media Stream (DMS) recording will be encrypted if you have enabled this feature on your CS1000 and have installed an Avaya Contact Recorder license key that includes the **Encrypted Media Streaming** option.

Live Monitoring of Audio

Avaya Contact Recorder's integral Live Monitor feature streams monitored audio from the recorder to the client's workstation. You can ensure this is encrypted - regardless of which recording method and switch is involved as described under Use of SSL on page 245.

Screen Recording

Screen recording uses up to three TCP/IP sockets between the Avaya Contact Recorder and the workstation (or terminal server) whose screen is being recorded.

Screen Content

This is in a proprietary, compressed form and hence not easily read from a network trace. However, it is good practice to encrypt this link as instructed in the *Avaya Desktop Applications Deployment Reference and Installation Guide.*

Note that the presence of a server.wss file in the recorder's /keystore folder automatically enables encryption of all screen recording. Avaya Contact Recorder only supports a single such file hence all clients must be configured with the same credentials and all screen recording will be encrypted.

Control Socket

This socket informs the recorder when Windows users log on and off. It also passes any Agent Initiated Monitoring commands from workstation to recorder. This link is not encrypted.

Live Monitoring of Screen Content

Encrypting the screen content link as described above automatically enforces the same level of encryption between the recorder and any client using its integral Live Monitor capability to view a screen in real-time.

Encrypted File Storage

Recordings (WAV and SCN files) and their associated XML data files are not normally encrypted but can be - using the AES256 algorithm and RSA Security's Enterprise Key Manager or Thales e-Security’s Data Security Manager. Encrypted audio files are also "fingerprinted" to avoid tampering.

See one of the following manuals:

- *Avaya WFO Security Configuration Guide* for details of how to deploy an RSA Key Management Server (KMS).
- *Avaya WFO Thales KMS 6.0.2 Installation and Configuration Guide* for details of how to deploy a Thales Key Management Server.

Enabling Encryption

There are 2 methods of enabling encryption on the recorder, depending on whether a Thales KMS or an RSA KMS is in place.

Note:

Keys are never cached to disk by the ACR.

If the KMS is not accessible when the ACR starts, then any recordings ACR makes are stored unencrypted until the ACR makes contact with the KMS. Replay of encrypted calls is not possible.

If the KMS becomes inaccessible once the ACR has retrieved the current key, ACR continues to use that key to securely encrypt calls. Replay of calls encrypted with that key is possible, any other calls are unplayable.

Note:

If you use WFO, you must configure it to use the same Key Management Server.

Thales KMS

The installation of Thales requires a separate client install on each ACR, there are separate instructions for Windows and Linux installations.

Note:

The install path must be set to the KMS subdirectory within the ACR install path on Windows. For example: *D:\Avaya\ACR152\KMS* where *D:\Avaya\ACR152* is the install directory of the ACR. On linux the install path is forces to */opt/Vormetric/DataSecurityExpert/*

Note:

High Availability is supported on the ACR with Thales KMS. See *Avaya WFO Thales KMS 6.0.2 Installation and Configuration Guide* for details of how to deploy a Thales Key Management Server in a High Availability environment.

To install on Windows

1. Install the VAEClient on each ACR in the topology as follows:
    a. Double click the **vee-key-n.n.n-nn-win64.exe.**
    b. Click **Next** twice.
    c. Accept the **End User License Agreement**.
2. During the installation, change the install path to the KMS subdirectory within the ACR install path on Windows.
3. Once the installation is complete, select **Register Vormetric Agent (64 bit) now**.

Note:

If this check box is not selected, the ACR is not able to connect to the Thales KMS until the register_host.exe is executed. This executable can be found at *D:\Avaya\ACR152\KMS\Agent\shared\bin\register_host.exe*.

4. On the first screen of the registration, click **Next**.
5. When prompted for the type of installation, select **Key Agent** and **Use Shared Secret Registration** and click **Next**.
6. Complete the following fields, then click **Next**.
    a. **Shared Secret**, this must be set to the shared secret already configured as part of the configuration detailed in the *Avaya WFO Thales KMS 6.0.2 Installation and Configuration Guide.*
    b. **Domain name**, must be set to the domain name created on the Thales KMS rather than the network domain.
    c. **Host name** and **Host description** (optional)
7. Enter a value for **Primary Data Security Manager** and click **Next**.

The value entered must be the fully qualified domain name configured for the Thales KMS as detailed in the *Avaya WFO Thales KMS 6.0.2 Installation and Configuration Guide.*

8. Enter a value for **Password** and click **Register**. This password is required by the user interface and as such must be secure and stored securely.

To install on Linux

1. Install the VAEClient on each ACR in the topology as follows:
    a. Copy the **vee-key-n.n.n_nn-rh7-x86_64.bin** file into the ACR.
    b. Run the command **chmod +x vee-key-n.n.n_nn-rh7-x86_64.bin** in the directory holding the copied file.
    c. Run the command **./vee-key-n.n.n_nn-rh7-x86_64.bin** in the directory holding the copied file.
    d. Click **Enter** to open the **End User License Agreement** and begin the installation.
2. Upon completion of the installation, when prompted to register the agent, press *y*. Continue registering the Agent with Thales KMS.

    You can perform this step at any time by running, though the ACR cannot connect to the Thales KMS until it has been done. `/opt/Vormetric/DataSecurityExpert/agent/pkcs11/bin/register_host`
3. Enter a value for a **primary security server host name**, when prompted. This must be the hostname of the Thales KMS already configured as detailed in *Avaya WFO Thales KMS 6.0.2 Installation and Configuration Guide*. You must ensure that the name resolves correctly.
4. Enter a hostname for the ACR. This must be the fully qualified domain name of the ACR.
5. Select *shared secret* as the **Registration Method**. This should already be configured on the KMS. Enter the shared secret when prompted.
6. Enter the domain name for the host. This value must be the domain created on the Thales KMS not the network domain.
7. Skip the next prompt for the host group name.
8. Enable the hardware association functionality by default.
9. Enter a password which is required by the user interface and as such must be secure and stored securely.
10. Go to **General Setup > Server** on the recorder and do the following:
    a. Enter the fully qualified domain name of the Key Management Server.
    b. Set the **KMS Type** to *Thales*.

c. Set **Password** with the password entered during registration in step 9.

d. Restart the recorder. A restart of the recorder is required after any changes to the KMS settings or passphrase.

RSA Key Management Server

Note:

High Availability is NOT supported on the ACR with RSA KMS.

1. We strongly recommended to create a new RSA KMS client certificate; the default KMS client certificate must NOT be used.

Note:

The same client certificate must be deployed to all ACRs in the system. You must ***not*** reuse the ACR's HTTPS server certificate as the RSA KMS client certificate.

2. Use the KMS Client Certificate keystore (**`I360KMClientCertKey.p12`**) uploaded to the KMS Server. **This certificate must be present on every Avaya Contact Recorder**. The certificate is provided by the CA (or generated as part of the instructions in Appendix F: Advanced Security Settings).
3. Be sure to use the server-specific certificate created using the Certificate Signing Request (or from Appendix F: Advanced Security Settings) called **`tomcat.p12`**. This certificate is specific to each server and a new certificate must be created for each specific server.
4. Install these keystores on the recorder by logging in as witness and copying the file to the /keystore directory beneath the recorder's install path (/opt/witness on Linux systems). The keystore's filenames must be **`tomcat.p12`** (for the server-specific certificate) and **`I360KMClientCertKey.p12`** (for the client certificate used on each server). The filenames are case sensitive on Linux.
5. On the **General Setup > Server** page of the recorder, enter the following:

    a. IP address of the Key Management Server

    b. Key Management Server Type as *RSA*

    c. strong passphrase for the client certificate (**`I360KMClientCertKey.p12`**).

    You must restart the recorder after making any changes to the KMS settings, keystore, or passphrase.

Slow Start-up under Linux on Virtual Machines

In some environments where a Linux based recorder is hosted in a Virtual Machine using RSA KMS it has been noticed that recorder startup is significantly slower. This can be improved by adding an additional JVM parameter to the file `/opt/witness/wrapper/acrl.conf`.

For example:

```
wrapper.java.additional.9=-
Djava.security.egd=file:/dev/./urandom
```

Replace the suffix (***9*** in the example), with the number that is one greater than the highest suffix on an existing entry that starts with `wrapper.java.additional.` i.e. the example is correct for a file that already contains entries ending **`.1`** to **`.8`** inclusive.

Restart the recorder to have this setting take effect.

I/O Jobs

If you have installed and configured a KMS as above, the I/O Jobs pages will show explicitly whether each job results in encrypted or decrypted files being written.

Archives

By default, all files written to archive destinations will also be encrypted using the KMS. If you need to create an archive that will be accessible without access to the KMS you can specify that the files should be decrypted when configuring an Archive I/O Job.

However, you should consider using “Mass Export” and/or a Distributed Replay Server instead. The latter supports re-encryption using a separate KMS.

Mass Export

As exported files are intended for use away from the recording system, they are always decrypted.

Distributed Replay Server Feeds

As the Distributed Replay Server runs outside your system, it will not have access to the KMS. Files are therefore always decrypted as they are written. You are therefore advised to use encrypted transmission (e.g. SFTP) for the staging area.

If the Distributed Replay Server is, itself, configured with a KMS, then it will encrypt the incoming files itself.

Finger-print Validator

This is a standalone application used to validate the fingerprint created in the WAV audio recording files. It can only be used if fingerprinting was enabled at the time of the recording.

To Enable Finger-printing

Fingerprinting is enabled automatically if KMS is configured but can also be turned on without KMS by setting the property `acr.fingerprintwithoutkms=true`

To Validate a File

1. Unzip the ValidateFP utility from utils directory in the installation disk onto a Windows system.
2. From a command prompt navigate to the directory where you have placed the file.
3. Copy the recording file in question (***nnnnnnnnnnnnnn*.wav**) from the recorder into this folder.
4. Validate the fingerprint using the command below, replacing ***nnnnnnnnnnnnnn*** with the name of the recording file.

   ValidateFP *nnnnnnnnnnnnnn*.wav

   Note:

   ValidateFP cannot be used to decrypt KMS encrypted files.

I/O Jobs

As part of designing any export/retention scheme you will need to consider the security of the recordings and how you will control access to them. You will probably need to comply with at least:

- National legal requirements (and often local or state regulations)
- Industry standards (such as PCI compliance or financial authority regulations) and
- Your internal corporate security policies.

This section summarizes how the security considerations discussed in this Chapter apply to each of the three commonly used export mechanisms.

Whichever approach you take, you should consider *at least* how your approach satisfies the above requirements with regard to:

- Encryption in transit to your staging point(s) – SFTP is preferred
- Encryption in transit within your network

- Encryption at rest on your storage devices
- Platform security and hardening
- User Access controls including password policies
- Retention and deletion rules
- Audit Trail of searches, replays, deletions…

Manual, Ad Hoc Requests

The exported files are unencrypted and are in industry standard formats so as to be easy to view and play. How these are managed, stored and secured is up to the recipient.

Mass Export

The exported files are unencrypted and are in industry standard formats so as to be easy to view and play. How these are managed, stored and secured is up to the recipient.

Distributed Replay Server

This supports the same range of security options that are available on the Avaya Contact Recording system that recorded the calls in the first place. This rest of this Chapter addresses these. Most are simply a matter of configuring your system but encrypted file storage requires an additional component.

In the same way that the original recording system may support Encrypted File Storage (by including a Key Management Server), so can the remote Distributed Replay Server. This optional (additional cost) feature allows your customer to keep all his files encrypted and hence only accessible via the Distributed Replay Server. Individual users may be granted export rights to allow them to export unencrypted files from the system.

Note:

Even if your recording system includes a KMS and hence encrypts files within the system, the files written to the DRS Feed have to be decrypted (as you cannot share the overarching keys from your KMS with your customer). It is therefore important that the files are transferred and stored securely (SFTP recommended) until they have been processed and, if a KMS is present, re-encrypted on the Distributed Replay Server.

PCI Compliance

To make your system compliant with PCI recommendations, you should adopt *all* of the above features. In addition:

1. Do not store sensitive data in the User Defined Fields of recordings.
2. Use the PAUSE and RESUME recorder control features to avoid recording sensitive information. This will require integration with your other systems.
3. Increase the default 3 months for which audit records are kept to 13 months.
4. Use encrypted audio to and from your phone system
5. Force users to use the Master, Central Replay Server or WFO rather than replaying calls from each recorder server.
6. Ensure all other components of your system (e.g. WFO) are configured in accordance with the *Avaya WFO Security Administration Guide*.
7. PCI dictates a session timeout of 15 minutes. When using Windows Domain
8. Authentication (as required to meet other PCI requirements) the recorder will accept a valid Windows logon (on an appropriate account) as sufficient to gain access to the recorder's administration pages. You should therefore ensure that all users' PCs are configured to launch a screen saver after 15 minutes of idle time and have Password Protection on Resume enabled. To ensure this, domain administrators should lock down these settings in the group policies of the domain controller.
9. Review and adjust your Windows Domain policies for user accounts if required. For instance, the PCI specification mandates the following rules:
   - Immediately revoke access for terminated users.
   - Remove inactive user accounts every 90 days.
   - Do not use group, shared or generic accounts and passwords.
   - Change user passwords at least every 90 days.
   - Passwords should be at least 7 characters long and should include both numeric and alphabetic characters.
   - Do not allow individuals to submit a new password that is the same as any of the last 6 they have used.
   - Lock out the user account after not more than six unsuccessful access attempts (have a lockout duration of 30 minutes or until admin enables the account).
10. To ensure no recorded data is stored anywhere in an unencrypted format, including on the supervisor's PCs, Internet Explorer's Advanced Security **Do not save encrypted pages to disk** and **Empty Temporary Internet Files folder when browser is closed** settings on the supervisors' PC must be enabled and locked down. (**Internet Options > Advanced > Security)**. If using other browsers, configure the corresponding settings on them.

11. To ensure proper authentication when SSL is enabled in a recording system, the following advanced Internet Explorer security settings on Supervisor PCs must be enabled
    - Check for the publisher's certificate revocation.
    - Check for server certificate revocation.
    - Warn about invalid site certificates.
    - Use TLS 1.2

Chapter 7: Advanced Configuration

This chapter provides an overview of the more complex and rarely used options for an Avaya Contact Recorder system.

The main sections in this chapter are:

- Properties File on page 265
- Slave Server on page 280
- Standby Server on page 281
- Central Replay Server on page 284
- Distributed Replay Server on page 288
- Usage Report on page 291
- Party Statistics Usage Report on page 292
- Configuration Report (Communication Manager only) on page 295

  Selective Record Barring on page 297

  After Call Work without Business Rules on page 298
- Altering Translations on page 298

Properties File

A number of system settings can be changed from their default values by placing entries in the properties file as described below:

- This is a plain text file, located in the installation path (which is /opt/witness on Linux) and with filename: **`/properties/acr.properties`**
- You should edit this file using notepad on Windows servers or, on Linux, using the "vi" text editor. On Linux you MUST be logged on as witness - NOT root.
- The Avaya Contact Recorder service reads this file as it starts, so any changes made to the file will not take effect until you next restart the service. Some properties are cached at startup but others are consulted as required. These can be changed temporarily at run time using the Maintenance page (**`/servlet/acr?cmd=mtce`**). You should only use this option under the direct guidance of Avaya support staff.
- The file is normally empty as all settings default as shown in the table below.

Most of these entries are discussed elsewhere in this manual, in the appropriate context.

- The table below provides a summary of the available settings.

Entry (case sensitive)	Default	Meaning
12_1Upgrader.rowlimit	5000	Use this property to vary the number of entries in a single 12.1 upgrade batch for performance purposes.
12_1Upgrader.vacuumafterbatch	200	Use this property to vary the number of batches of upgrade entries before a vacuum database is run (for performance purposes).
15_2Upgrader.rowlimit	5000	Use this property to vary the number of entries in a single 15.2 upgrade batch for performance purposes.
acr.disablecompress	false	Normally, recordings made in G.711 will be compressed to G.729 8kbps. Set this to true to disable compression.
acr.disablereporting	false	Set true to block usage and party statistics reports.
acr.diskmanager	true	Whether to delete oldest calls when needed to create space for new recordings.
acr.localport	8080	The http port used by the recorder. Must match the entry in server.xml.
acr.log.datePattern	'.'yyyy-MM-dd	Sets the format and frequency of log file name rotation. Default provides daily file. For heavy load testing, set to **`'.'yyyy-MM-dd-HH`**

Entry (case sensitive)	Default	Meaning
acr.logkeepdays	30	Number of days log files to retain before purging.
acr.login.lockmins	60	Minutes for which to lock out a user account following maximum allowed number of login attempts.
acr.login.lockout	10	Number of login attempts after which an account will locked out.
acr.singlelogin	False	Set true to allow each user to be logged in at only one device at a time.
activitycode.fieldname	activitycode	(CS1000 only) The name of the user defined field to store (last) activity code in. Set blank to disable storage of activity code.
acw.default	0	Specifies the number of seconds of After Call Work (i.e. how long screen recording continues after call ends) for cases where no WFO Business Rule fired.
acw.earlyterminate	true	Terminate After Call Work screen recording as soon as another contact is connected.
aoc.attribute.xxxx	xxxx	Store Avaya Oceana™ attribute xxxx as a user defined field of the name specified. If set to empty string, do not store this attribute. All attribute names are forced to lower case so this property must also be all in lower case.
aoc.serviceattribute	null	If set, use the Avaya Oceana™ attribute specified as a pseudo “VDN” allowing a bulk recording rule to be added to control recording on the basis of this one attribute. All attribute names are forced to lower case so this property must also be all in lower case.
apachehttpclient.level	ERROR	Use this property to display different levels of the HTTP client messages in the acr log. Any messages lower than the selected level are discarded. Possible entries are: ERROR/WARN/INFO/DEBUG
arc.limit	100	This property was added to limit the number of inums in a single historical archive batch job and prevent massive archive batches blocking the archive process. Use this property to vary that number if required.

Entry (case sensitive)	Default	Meaning
archive.negspaceworkaround	false	Some NAS servers erroneously report negative free space available. This would normally prevent ACR from reactivating an existing folder. Set to "true" to ignore the freespace value and attempt to write to such paths.
archive.oldmechanism	false	Several improvements to classify and report archiving errors could have unexpected results on site and generate multiple alarms. Set this property to true to revert to the previous behaviour.
archive.writebuffersize	64000	Use this property to vary the amount of data being written at each attempt. This can improve performance writing to sFTP targets on high latency networks. This setting has a minimum of value 64000, and a maximum value of 256000.
audioserver.inactivitytimeout	30	The number of seconds of inactivity before an audioserver port is released back to the pool of available ports.
audit.purgemonths	3	The number of months to keep audit trail entries in the database.
beaconsoftphone.timeout	3	On HA systems only, use this property to set the delay (in seconds) before the attempt to check the status of the local beacon softphone times out and an alarm is raised.
bulk.dominates	false	Set this property to true to ensure that a recording decision affected by **both** a bulk record rule **and** a RetainOnlyOnCodeEntry rule will allow the bulk record rule to decide the record decision.
call.persistudfs	true	Set false to stop User Defined Field tags rolling forward onto the surviving call when calls merge on transfer and conferencing.
call.tracing	false	Set true to provide additional tracing of call and connection states.
cct.conntimeout	10000	Set this property to a higher value (in milliseconds) if connection to CCT is timing out.

Entry (case sensitive)	Default	Meaning
cct.socktimeout	15000	Set this property to a higher value (in milliseconds) if CCT requests (such as getAgents) are timing out.
conferenced.ignore	null	Set to one or more ranges of phones that should be ignored if present on calls. Use this to stop the recorder interpreting other recorders' softphones as additional parties on the call. Separate ranges with a comma. For example: **`10000-10099, 12000-12049`**. Can also be used to ignore IVR ports used to provide pre-recorded announcements on calls.
config.headers	true	Set false to disable header row from nnnnnn_config.log report (if enabled).
config.reporting	false	Set true to have recording targets automatically added to file ***nnnnnn*`_config.log`** (where ***nnnnnn*** is the recorder's serial number) in logs folder at midnight each day and on any changes being made.
controllerlistener.localport	8232	TCP Port number the master listens on for connections from Contact Recording Desktop applications.
cors.override	*	Set to comma separated list of Central Replay Servers' fully qualified domain names if required to tighten Cross Origin Resource Sharing (CORS) security.
cpu.blockedms	2000	How many milliseconds to tolerate between executions of the (supposedly) once per second thread before raising an alarm.
crs.checklocal	false	Set true to have a CRS check its own calls folder for calls. Can be used to hold recordings from one or more mothballed recorders - avoiding having to go to Archive for recordings that are still present.
crs.flush	10	The time to wait (in seconds) before CRS updates its local cache to the database (and thereby makes calls available for replay).
crs.maxinumsperjob	100	Use this property to change the maximum size that the cache of inums will grow to before being pushed to CRS.

Entry (case sensitive)	Default	Meaning
crs.v12forward	false	Set true to have a Central Replay Server forward requests on to Avaya Contact Recorders in V12.0 format – allowing CRS to be upgraded ahead of the recorders feeding it.
cs1k.immediateclear	true	Set false to delay acting on a disconnection, true to act immediately.
cs1k.showacddn	true	Set to false to stop the ACDDN to which an agent logs on showing alongside the agent's identifier.
cti.logfile	NULL	The filename (within the /logs folder) to read CS1000 CTI input from or write it to
cti.mssettling	250	How many milliseconds to wait after a CTI event before considering recording actions. Allows transient intermediate states to be ignored.
cti.offline	false	If true, read MLS or ICM input from the file specified in cti.logfile. Otherwise connect to live CTI feed(s).
cti.restartmins	5	How many minutes of outage to allow during a maintenance window before raising an alarm.
db.password	Not disclosed	Encrypted password for recorder's normal PostgreSQL user account
db.tickle	30	Number of seconds between database watchdog queries. See db.timeout.
db.timeout	180	Number of seconds within which the PostgreSQL database must respond to a basic select query. If it does not, the recorder will exit, forcing a restart. Protects against database lockups.
defaultstatusperiod	1	Set to the number of months of delete requests to display when reviewing GDPR Delete status.
device.tracing	false	Set true to provide additional tracing of device management functions.
dddddd.field.***fffff***	NULL	For dialer named ***dddddd***, store field ***fffff*** as User Defined field specified.

Entry (case sensitive)	Default	Meaning
dialedtodnis.udfname	NULL	The name of the user defined field to store a copy of the “Dialed to (DNIS)” field sent to WFO if required.
dialer.tracing	false	Set true to provide additional tracing of dialer CTI.
disk.stopatMB	10	Stop writing to disk if free space is less than this many MB
disk.warnAtMB	500	Raise warning if free disk space falls below this many MB
diskmanager.mingb	1	The number of GB that diskmanager will attempt to keep free on the calls partition.
dmcc.addbusyverify	false	Set true to have beep-tone injected into call whenever a recorder’s DMCC port is conferenced into the call. Beep-tone will also appear in recording. Does not require additional time-slot as normal beep-tone injection does.
dmcc.badcallthreshold	3	How many consecutive failed recordings to tolerate before busying out a DMCC softphone – which must then be **Reset** and hence reregistered before being used again.
dmcc.port	4722	The port number on the AES for DMCC.
dmcc.secure	true	Whether to connect to DMCC using SSL.
dmcc.tracing	false	Set true to provide additional tracing of DMCC interactions.
dmcc.trustall	false	Allows different SSL certificates to be installed on the AES.
dmcchandlers.maxfree	100	Maximum number of free ports to retain in the DMCC recording pool before clawing back some.
dn.regpersec	10	How many DNs/Position IDs to attempt to Associate per second on MLS or ICM links. Increasing this number speeds up initial startup but will impact the switch more during this period.
dpa.tagoneonly	false	On CS1K environments, set to "true" to change the DataEvent to only tag most recent INum across controllable devices.

Entry (case sensitive)	Default	Meaning
dtls.nonstandard	3	Change to 2 to use recorder with CS1K phones using Unistim 5.3 firmware or later.
email.minalarmlevel	0	Alarm level at or above which, emails will be sent. Default, 0 is INFO level. Set to 1, 2 or 3 to restrict emails to Warnings, Minor and Major Alarms respectively.
execmode.deletedelaymins	0	Minutes to wait after a recording ends for a "retain" command before deleting a call on a DN or Position ID requiring manual retention. Must be set identically on ALL recorders.
execmode.deletedelaysecs	0	Number of seconds to wait before deleting recordings. Must be set identically on ALL recorders. DEPRECATED. Use execmode.deletedelaymins instead.
execmode.retainnumber	no default value	The station number of the recorder port to be used as the "Retain" port.
fallbacktoany.enable	True	Set this property to false to prevent recording falling back across recorder pools, for example, where recording of calls must happen within country of origin.
foureyes.default	false	Set true to force dual sign-in.
fqdn.***nnnnnn***	Null	Use this setting to specify the Fully Qualified Domain Name of the server with serial number ***nnnnnn***. Required if https is used in multiple server systems (as requests must be forwarded to other servers using the FQDN that matches their SSL certificate).
gdpr.statuspagesize	50	Set to the number of rows to be displayed on the GDPR Delete status page.
https.socket	8443	socket to use for https access. Must match that defined in server.xml
inactivity.alarmlevel	-1	The level of alarm to raise if recorder inactivity is detected. Disabled by default. Set 0 for INFO, 1 for WARNING, 2 for MINOR and 3 for MAJOR alarm levels.
inactivity.alarmlevel.***nnnnnn***		Override inactivity.alarmlevel for recorder with serial number ***nnnnnn***.

Entry (case sensitive)	Default	Meaning
inactivity.day_n	null	Maximum time between recordings on day "n" of the week (according to server's locale). One integer per hour of the day, separated by commas E.g. inactivity.day_2=0,0,0,0,0,0,0,30,20,10,10,10,10,10,10,10,20,30 Trailing zeroes can be omitted. Zero implies no alarms during that hour of day (no recordings expected).
krb5.exact	false	If false, change incoming Windows domain/usernames to uppercase and compare with the configured user accounts. If true, require a case sensitive match between the received domain/username and a configured user account.
krb5.tracing	false	Enables additional output to the log file
lock.acceptaudioonly	false	Set to "true" to allow locking of calls where only the audio (and not the screen) can be retrieved.
log.annotate	false	Set true to have each administration web page display an annotation field and buttons allowing the user to submit entries for insertion into the log file(s) from their browser.
loop.ignore.nn	false	Ignore calls on CS1000 loop number nn (1-255). Add one entry for each loop that is not recorded on a partially equipped trunk-side recording system. Otherwise alarms will be generated as the recorder attempts to record calls on these loops.
meeting.defaultrecord	true	Whether or not Meeting recording ports will start recording automatically at the start of a call or wait for a START command from an external controller.
meeting.ocpneeded	true	Whether or not Meeting recordings should detect other call parties using the conference display feature. Set false if an external controller is to provide this information instead.
mls.acceptabledelay	200	Report warning if MLS response takes longer than this (in milliseconds)
mls.pollinginterval	20	Time (in seconds) between poll messages with MLS.

Entry (case sensitive)	Default	Meaning
numberdialed.udfname	NULL	The name of the user defined field to store a copy of the “NumberDialed” field sent to WFO if required.
ocp.alwaysrefresh	false	In recording modes where the conference display button is used to determine who is on the call, this process will be restarted if the display indicates that a party has joined the call. In some cases, however, an existing party may be updated e.g. a trunk name/number may be replaced by a (late arriving) ANI. Set this property true to have ANY change to the display trigger a refresh of all parties on the call.
ocp.strictparsing	false	When forwarding via coverage answer groups, the conference display may not contain a single numeric field in the end of the display. Previously this would lead to the recorder continually pressed the Conf Disp button - resulting in the party being removed from the call. Now default is not to check for spaces or valid number in this string. To restore previous behaviour if this causes problems, set this property = true.
oemcomms.connecttimeout	120	Seconds a recorder will wait on startup before reporting that a previously connected recorder has failed to connect.
oemcomms.inactivitytimeout	60	Seconds the Master will wait between polling messages before reporting another server as having failed.
oemcomms.tracing	false	Enable tracing of messages between recorders
ondemand.defaultrecord	true	Whether or not On Demand recording ports will start recording automatically at the start of a call or wait for a START command from an external controller.
ondemand.ocpblocked	false	If enabled, prevents OCP with On Demand recording.
ondemand.ocpneeded	true	Whether or not On Demand recordings should detect other call parties using the conference display feature. Set false if an external controller is to provide this information instead.

Entry (case sensitive)	Default	Meaning
owners.additive	false	Set to true to have **Advanced Owner(s)** setting in Bulk mode add to default owner(s) rather than replace them.
perday.tracing	false	Set true for additional tracing within the once per day background task that occurs just after midnight.
perhour.tracing	false	Set true for additional tracing within the once per hour background task that occurs on the hour.
perminute.tracing	false	Set true for additional tracing within the once per minute background task.
permorning.tracing	false	Set true for additional tracing within the once per morning background task. See reminders.hourofday for when this fires.
pernight.tracing	false	Set true for additional tracing within the once per night background task (defaults to 1AM).
persecond.tracing	false	Set true for additional tracing within the once per second background task.
purge.hourofday	1	Hour at which database and log file purges will take place. Default is 1am. Valid range 0 - 23. These actions will NOT take place if the recorder is restarted during this hour.
queue.fullname	false	If true, present "queue" to WFO with description if available (e.g. from Communication Manager) as "nnnn (description)".
queue.xxx.threads	1	Number of threads to use for job queue xxx.
rec.maskallowed	false	You must set this to true to enable the "pause/resume" features of any external APIs.
rec.mincallduration	250	Calls shorter than this many milliseconds are deleted. In hybrid AACC systems, set this to 1000 to avoid short recordings that can result from the difference in arrival times of CTI events over the two CTI feeds being used. Deleting such sub-second recordings avoids confusion when searching and playing recordings.
record.persistmode	device	Use this setting to choose between "call", "device" and "segment" persistence of recording commands.

Entry (case sensitive)	Default	Meaning
recorder.pool	NULL	This property setting is now deprecated and should be removed from the properties file. Designate a pool (or set of pools) for a recorder from the General Setup > Server in the "Designated Recorder/Pool(s)" Advanced setting.
reminders.hourofday	9	Hour of day when daily reminders are sent for Replay Authorization process.
remoteslave.***n***	NULL	Add to Master's property file to have a remote slave included in WFO and I/O Job server lists. N must start at 1 and be contiguous. The value of the property is the remote slaves 6-digit serial number. For example: remoteslave.1=8801200 remoteslave.2=8801201
restrictedadmin.mayeditaccounts	true	Set to false to stop **Restricted Admin.** accounts from having access to any user account administration.
restrictedadmin.mayviewctistatus	false	Set to true to allow **Restricted Admin.** accounts to use the **Status > CTI Monitors** page.
restrictedadmin.mayviewportstatuss	false	Set to true to allow **Restricted Admin.** accounts to use the **Status > Ports** page.
rtp.highport	20000	Upper limit on ports user for SIP and Duplicate Media Stream IP recording. Ports up to but not including this port will be used. Note: each recording uses up to FOUR ports.
rtp.lowport	12000	Lowest port number used for SIP and Duplicate Media Stream IP recording. Note each recording uses up to FOUR ports.
rtp.packetlog	false	If true, log details of every RTP packet received.
rtphandlers.maxfree	10000	Maximum number of free ports to retain in any recording pool (except DMCC) before clawing back some. The default value effectively disables this feature.
rtplegacy.highport	12000	Highest port number used for On Demand, Meeting and Phone Replay ports.
rtplegacy.lowport	10000	Lowest port number used for On Demand, Meeting and Phone Replay ports.

Entry (case sensitive)	Default	Meaning
screen.bulkdominates	false	Set to true to force the Bulk recording mode's screen percentage to override WFO's implicit 100% requirement for calls that did not trigger a business rule.
screen.ieport	29522	The port number on which the recorder listens for connections from agent desktop screen capture client connections.
screen.maxdurationmins	720	Maximum duration of a screen recording. Any screen recording reaching this duration will be terminated.
server.locale	en	The locale the server will use for automated output - i.e. where the user's locale cannot be determined by the user's browser settings. This is used for automated reporting and emailing of alarms.
siprec.mode	2	Which type of SBC is sending SIPREC requests. 2=ASBCE. 1=Other
sip.tracing	false	Set true to provide additional tracing of SIP interactions.
skill.fullname	false	If true, present "skill" to WFO with description if available (e.g. from Communication Manager) as "nnnn (description)".
slave.replayenabled	false	The administration setting "Allow search and replay from this server?" on a slave can no longer be edited via the admin pages. Replay is normally performed via Master/Standby or Central Replay Server(s). To enable replay directly from a slave, (for diagnostic purposes only) set this property to "true" (can be done at run-time via the maintenance page). Note that the admin page will continue to show that replay is not allowed.
smtp.port	25	The port used for emailing.
snmp.authtype	SHA	Set to "MD5" to override the default SHA authorization type with MD5.
snmp.mainusername	acrsnmpus er	The main username the Network Management System will use to connect to the recorder.
snmp.password.nn	no default value	Encrypted password to be used in conjunction with the corresponding snmp.username.nn entry.

Entry (case sensitive)	Default	Meaning
snmp.port	2161	The port number to use for SNMP.
snmp.privtype	AES128	Set to "DES", "3DES" or "AES256" to change the privacy type from AES128.
snmp.username.***nn***	no default value	Additional usernames can be entered. Replace ***nn*** with 1, 2, 3 etc.
sp.tracing	False	Set true for enhanced logging of On Demand and Meeting recording modes.
system.forcertl	false	Set true to have number ranges in audit entries and other database entries use right-to-left order (e.g. 200-100 instead of 100-200).
trunkgroup.ignore.nnn (Communication Manager only)	false	Set to "true" to suppress warnings about not being able to record calls on trunk group nnn. Add one entry for each trunkgroup that is not recorded on a partially equipped trunk-side recording system.
tsapi.cachecallingparty	true	The default behaviour is now to cache calling party and use to identify unknown dynamic in subsequent snapshot. Set to "false" if original behaviour is required.
tsapi.numplanlength	16	Force any address of more than this number of digits to be treated as an external number. Only required if problems with external numbers being misidentified as internal. Note that previous default (12.0 and earlier) was 8.
tsapi.retries	300	Number of times to keep trying to establish single-step conference.
tsapi.retrydelayms	1000	Interval in milliseconds between single-step conference attempts.
tsapi.timeout	65	How long (in seconds) the recorder will tolerate TSAPI failure to respond to heartbeats.
tsapi.tracing	false	Set true to provide additional tracing of TSAPI interactions.
unify.required	false	Whether a link to Unify is required for the recorder to be viable.
unify.xmlonstart	true	"Set to false to restore behavior of STARTED message to that of 10.0 and earlier (where XML is not sent in the STARTED message but, rather, in an UPDATE message immediately after the STARTED).

Entry (case sensitive)	Default	Meaning
usage.dailystats	false	Set true to enable automatic nightly reporting of Party Statistics.
usage.header	true	Set false to disable the header row in the Party Statistics report.
usage.partystats	false	Set true to enable a usage report detailing what has been recorded.
usage.partystats	false	Set true to enable Party Statistics reporting.
usage.reporting	false	Whether to track actual usage to the ConfigHistory table and show the Usage report.
uui.fieldname	no default value	User defined field name to store Avaya User to User data. If NULL, this data is not stored. Applies to Conferenced mode recording only.
viewerx.limit	100	Maximum number of recordings returned for each search.
viewerx.savesettings	false	Whether to save and reuse users' previous entries in search filter pane.
viewerx.secure	false	Set to true for enhanced security features at the expense of usability.
viewerx.timelimit	300	How long a user's search query will run before timing out. If you increase this setting you may also need to increase your browser's timeout so that it does not time out before the results are returned.
vx.servers.main	null	On a CRS that consolidates recordings from two parallel recording systems, set this to the (comma separate) list of serial numbers in your "main" system i.e. excluding those in the "secondary" system. Session layouts will then default to playing the calls from the main servers and fall back automatically to the secondary if no recordings found for the specified contact.
windowsuser.ignore	null	To ignore specific user account logon/logoff events from the screen capture client, set this property to the regular expression that will match the account names to be ignored. Note that you must double-escape back-slashes as this is read from the properties file. For example HQ\\\\svc.* will ignore all accounts in the HQ domain that begin with svc.

Entry (case sensitive)	Default	Meaning
xxx.tracing	false	Set to true to enable more detailed tracing of specific functionality. ***xxx*** can be any of: aim, aoc, call, iojobs, iojob.*nn*, cti, device, dialer, dmcc, dtls, eqc, eventlock, ha, krb5, mas, mdl, monitor, oemcomms, pip, recend, replay, screen, session, sip, sp, tag, tsapi.

Slave Server

Installation

(IP) Slave servers are installed in the same way as a standalone recorder (see Installing Avaya Contact Recorder on page 108).

Licensing

Instead of entering a license key, follow the instructions on the lower half of the license entry page. Allocate the recorder a unique number (2 to 9999) and specify the IP addresses of the servers that could control it.

These are:

- The Master
- All Full Standby recorders
- Any Partial Standby recorders in the same pool as this Slave

Restart the Avaya Contact Recorder service so the Slave can connect to and be controlled by whichever controller is **Active**.

Configuration

Work through Chapter 4: Configuration on page 115 configuring just those fields with **Edit** links next to them. Note in particular:

General Setup > Server lets you specify a Beacon softphone as described in

- Beacon Softphone (Communication Manager only) on page 128 and subsequent sections.

Important:

Note that a Slave on a remote survivable site may be better configured as a Partial Standby so that it can attempt to record calls locally should connection to the Master and all Full Standby servers be lost.

Major Switch Configuration Changes

If you add or delete a Data Source on the Master recorder, or change the type of a Data Source, these details are automatically passed to the Slave servers within one minute but you must then restart all servers for such a major change to take effect.

User Accounts

System Admin. and **Restricted Admin.** user accounts are automatically copied to all Slave servers - but not accounts that are purely for replay. Slave servers should not be used for replay. Use the server(s) they upload their recording details to instead.

Search and Replay

The administration setting **Setup > Manage Users > Allow search and replay from this server?** can no longer be edited via the administration pages on a Slave server. To enable replay directly from a Slave (for diagnostic purposes only) set the property file entry: **`slave.replayenabled=true`**

This can be done at run-time (via the maintenance page) but, regardless of how it is set, the administration page will continue to show that replay is not enabled.

Standby Server

Design and Planning

First, determine the location, number and type of Standby servers required using Appendix D: High Availability on page 375.

Installation

Standby servers are installed in the same way as a standalone recorder (see Installing Avaya Contact Recorder on page 108).

Licensing

Instead of entering a license key, follow the instructions on the lower half of the license entry page. Allocate the recorder a unique number (2 to 9999) and specify the IP address(es) of the servers it will connect to, namely:

- The Master
- All (other) **Full Standby** recorders
- Any (other) **Partial Standby** recorders in the same pool as this server.

Configuration

Standby Servers copy most of the configuration of the Master recorder and you need only enter a small number of details. Work through the Configuration chapter of this manual configuring just those fields with **Edit** links next to them. Note in particular:

General Setup > Server has options for Full versus Partial Coverage, Hot versus Warm Readiness and Priority as described in

- Beacon Softphone (Communication Manager only) on page 128 and subsequent sections.
- **General Setup > Data Source(s)** - Some of the settings on these pages are derived from the Master and will not show an **Edit** link next to them. The others must be configured on the Standby as these can (and in some cases should) differ from the Master. However, the settings that must be made here are ones that are unlikely to change once your system is running.
- **System > Email Server**

  You are strongly advised to create a separate email account, using a separate SMTP server from that used by the Master. This way, failure of an SMTP server will be noticed as you continue to receive e-mails from the Standby but not the nightly heartbeats from the Master - or vice versa. It is critical that you ensure that the Standby is ready to take over at any time.
- **System > Manage Users** - All user account details are copied from the Master. You are only able to edit these on the Standby when it is active - and even then, you should only do so when really necessary as any changes you make on the Standby will be overwritten when the Master next makes contact with the Standby.
- **Operations** - These pages are provided but their contents are not editable as these are copied from the Master recorder (once contact has been established with it). If using a **Partial Standby**, ensure you have configured (on the Master) the **Operations > Bulk Recording > Advanced > Designated Pool(s)** for the recording targets that this Standby is to control when needed.
- **Alarms > View Alarms** - This page is identical to that on the Master though some different alarms will be generated - when the Standby establishes contact with the Master and should it ever try to take over from the failed Master. You should use the email settings on the **System > Email Server** page to ensure that you are advised of these alarms within 10 minutes of one occurring.

Central Replay Server

You can deploy the recorder application onto an additional server which collects details of recordings made on one or more other servers and therefore can act as a dedicated search and replay server. Users can then search for recordings from any of the recorders feeding it with a single url and a single search.

Archives

As “tar” files are written to archive destinations, the Central Replay Server automatically tracks which archive destination this set of calls has been copied to. Because a single replay server can be uploading calls from more than one recording system, it cannot simply use the same **I/O Job** id (normally starting at 1) that is used on the corresponding Master recorder. These could clash across systems. Instead it creates a new id consisting of the master’s six digit serial number followed by the archive’s job identifier shown on that master. Hence archive 1 on 890001 becomes archive 890001001 on the Central Replay Server.

A Central Replay Server may therefore be aware of many different archive destinations, not all of which are necessarily accessible to it. This has several implications:

1. Unlike a normal recorder – which needs to actually archive its recordings – the Central Replay Server makes no recordings so will ONLY run an archive thread if that archive is explicitly set to run on it.
2. For NAS archives, the Central Replay Server may need to use a different path and/or credentials to access the share. For example, the share seen by the recorders may actually be a disk on the replay server and hence accessible to it via a local drive letter; the archive may be obsolete and has been moved to slower/cheaper storage on a different path.

The default mode of operation is that the Central Replay Server will forward all replay requests on to the recorder that made the recording. That recorder will then access an archive destination if necessary but, for recent recordings, is likely to be able to satisfy the request from its own calls buffer.

This can be overridden to reduce the load on the recorders themselves but at the expense of slower replay as the replay server will then have to fetch the call from a (large) tar file within the archive. To do this, use the administration pages on the Central Replay Server as follows:

1. Click on the **Operations > I/O Jobs** tab.
2. Click **Edit** on the **Service Archive Requests Locally** row and check the box if you want this server to service any requests that it can access directly from the archive locations rather than forward them to the recorder that made the call.
3. Then, for each archive that you want the replay server to access directly,

- Click the **Edit** link to the right of its path.
- Click the **Advanced** button.
- Check this Central Replay Server's six-digit serial number in the **Run on Server(s)** field to make it run there.
- If the path is to a NAS, set the **Path** and add **Username** and **Password** credentials to allow it to access the share. Similarly, for EMC archives, configure the appropriate **Connection string** and credentials so that the Central Replay Server can access the archived files itself.
- Click **Enter** to save the new settings.

"Turbo" Mode

On large systems, performing searches across more than a few days' recordings can take some time. A Central Replay Server can be configured to provide much more responsive searches – in return for extra disk space being used.

On a Central Replay Server, any or all search Layouts can be marked with **Turbo** mode enabled using **Layout Builder**. You should do this on large systems for layouts where users often need to perform searches across timespans of more than a few days and hence are waiting more than a few seconds for results. This mode works by pre-populating one or two additional database tables with the fields needed for all of the layouts that are set to use **Turbo** mode. This technique is known as an "instantiated view".

How it Works

As recordings are made, either or both tables (`segment_i` for segment- and `session_i` for session-based **Turbo** layouts) are updated with the appropriate rows for these new recordings. The end user's search then only has to scan this table rather than the several joined tables that are needed to populate it in the first place.

When more than one layout is marked for **Turbo** mode, these pre-populated tables are optimized such that any fields common to multiple layouts are mapped to a single field that is used for all layouts that need it.

Storage Implications

As a rule of thumb, when you enable turbo mode on any *one* layout, you increase the storage requirement of your Central Replay Server by approximately 25%. Having turbo mode on at least one segment and one session-based layout doubles this to approximately 50%. If you have many layouts, with lots of fields, each one increases storage requirements further as these additional fields are supported.

Hence it is a good idea to provide several times the initial estimated database size to avoid having to extend this disk partition in future.

Enabling Turbo Mode

Because of the additional disk space required, Turbo Mode is not enabled by default. To enable it, add the property **`turbo.enabled=true`** to the server's properties file. This setting will only take effect on the next restart.

Pausing Turbo Mode

If you find that creating the turbo tables is impacting the performance of the server, or disk space is running low, you can pause it by setting the property **`turbo.enabled`** back to **`false`** from the maintenance page (at `/servlet/acr?cmd=mtce`).

You can restart it again out of hours by setting the property back to **`true`**. Note that this only works if the property was set in the property file when the server started. If not, there will be no background thread running to look for the pause/resume command.

When Changes Are Applied

You may make several changes to your search layouts in rapid succession. It is not appropriate for the server to start modifying the (huge) pre-populated tables as soon as a single change is detected.

Therefore, if you make a change to a layout that means the existing pre-populated tables are inadequate (for example, you add a brand new field – that has not been used in any turbo mode layout to date) then queries on that layout will automatically revert to normal mode until the tables can be modified.

To force the server to repopulate the turbo mode tables, you can either restart the service or wait at least two minutes after your last layout change (to allow configuration to ripple to other servers) then click the **Refresh Turbo View(s)** button on the maintenance page (at `/servlet/acr?cmd=mtce`). Do this on each Central Replay Server. Also, note that turbo tables will not be populated if your server is still migrating recordings from an earlier version of Avaya Contact Recorder. This can also take considerable time (up to several days).

Be Patient Following Changes

When you first mark a layout as requiring turbo mode, or if you subsequently change the fields in a layout, the pre-populated table(s) may have to be created or modified – for all existing recordings. This can take some time. Days in some cases. During this period, the server notes what timespan the new table structure covers and will automatically use it for queries within that range. The tables are created or updated with the most recent recordings first so searches for recent calls will be back up and running in turbo mode as soon as possible while searches for older calls will take longer to return to turbo mode.

Limitations

Turbo mode is only supported on dedicated Central Replay Servers and cannot be used in conjunction with:

- Locking Recordings (see page 211)Replay Authorization Process (see page 213)

Installation

Install a Central Replay server in the same way as the Master recorder (see Installing Avaya Contact Recorder on page 108- steps 1-7).

Configuration

Work through Chapter 4: Configuration on page 115, configuring just those fields with **Edit** links next to them. Then configure the user accounts and replay rights for those entitled to use the application.

Telephone Replay (Communication Manager only)

Add a Data Source for the Communication Manager and complete its configuration page.

Add as many softphones as you need for replay ports. Then use the corresponding **Operations > Phone Replay** tab to specify the number of ports to use and the level at which to raise an alarm should the number of available ports drop to an unacceptable level.

Archive Destinations

See Archive on page 284 for details of how to configure archive destinations on a CRS.

Live Monitor

The Monitor link on the Replay page will automatically redirect requests to the active Master or Standby server (which has access to the live audio). The address that the replay servers uses to forward such requests is learned automatically as recordings from that Master/Standby are uploaded into the CRS. If these servers do not actually record, you will need to explicitly set their addresses using the property-file setting:

`fqdn.nnnnnn=xxxxxx`

where ***nnnnnn*** is the serial number of the recording server and ***xxxxxx*** is its (preferably fully qualified) hostname.

Note:

Live Monitor is not supported via the Central Replay Server if there is more than one Master connected to the server.

Configuring other Recorders

On each of the other IP recorders (Master, Standby and Slaves) enter this replay server's IP address as the **Replay Server(s)** on the **General Setup > Server** page. This will override their default configuration (which is to use Master and first Full Standby as Central Replay Servers).

Distributed Replay Server

Many owners of Avaya Contact Recording systems are outsourcers, recording calls on behalf of multiple end customers for whom they handle calls. Often each such customer requires a copy of and/or ongoing access to the recordings of their calls. In response to the limitations and costs of traditional file export mechanisms, Avaya has introduced the option of providing them with their Avaya Contact Recorder server which is licensed as a "Distributed Replay Server". This is then fed via a "**DRS Feed**" configured on the **I/O Jobs** page of the main recording system.

Note:

In this section "customer" refers to the customer of an outsourcer i.e. the company for whom recordings have been made.

Features

This solution is designed for any customer who has ANY of the following requirements:

1. Wants to take control of at least a significant portion of their recordings.
2. Wants to manage retention and deletion of recordings automatically – without having to design their own system.
3. Wants to provide quick, easy, powerful access to any subset of their recordings – regardless of selection criteria.
4. Frequent access to recordings – especially by multiple people and/or processes.
5. Needs to control access to recordings so that users can only search for and replay the recordings they are entitled to.
6. Wants to satisfy multiple overlapping demands for access to the recordings – i.e. anywhere a single recording might be of interest to more than one person/role/inquiry.
7. Is at all uncertain what their current and/or future requirements will be – and hence wants flexibility now and in the future.

Configuration

With this approach:

1. Customer installs a copy of Avaya Contact Recorder on a dedicated server or VM slice as shown below:

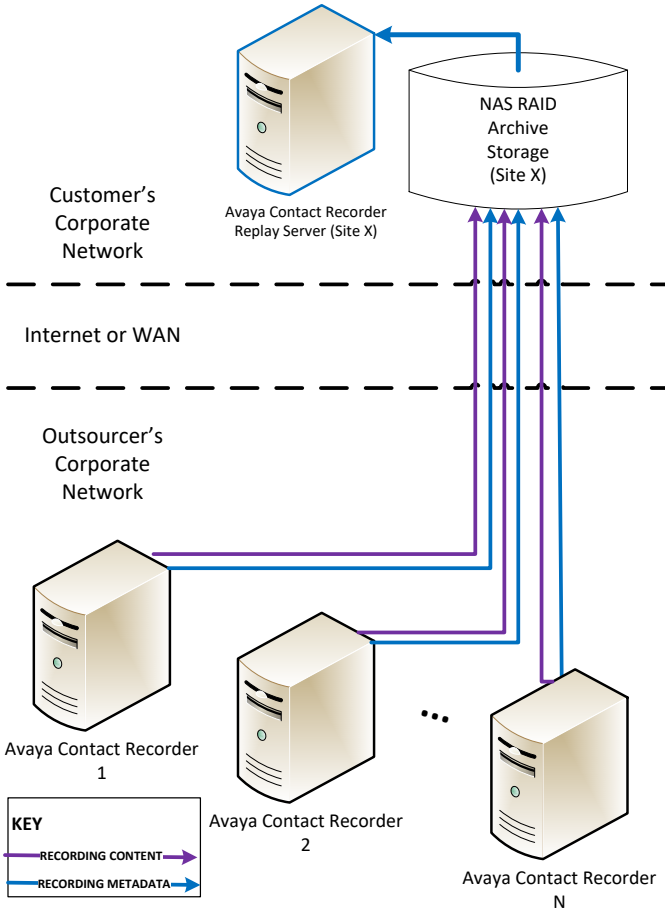

2. Customer enters a special license key that enables only the replay and archiving aspects of the system (i.e. it hides all the complicated recording control).
3. (If running on Windows) Customer specifies the call recording buffer location into which recordings will be copied. (This is fixed at /calls on Linux).
4. You configure a **DRS Feed** on the **I/O Jobs** page on the recording system – that selects all (or, if you set a **Filter**, a specified subset) of your recordings and copies them to a network accessible location. The Advanced settings available to you are a controlled subset of those for a normal archive – ensuring that the appropriate information is provided to the DRS.

5. Customer configures a DRS Import on their **I/O Jobs** page to look for incoming files in that same location. This has no Advanced settings. It takes what your system feeds it.
6. Your recordings are transferred automatically to the customer Distributed Replay Server

Deliverables

Having performed the above, the customer then has the same capabilities as you do to:

1. Grant appropriate search and replay rights to their staff
2. Be alerted should any problems occur with the system
3. Have a full audit trail of searches and replays performed by their staff
4. Archive or export subsets of recordings to other destinations by configuring other **I/O Jobs**
5. Automatically purge these recordings when they expire
6. Have a complete database of ALL tagging details – not just those that could fit into a filename

They *also* have a better level of insurance than in simple file export solutions. They have copies of their recordings and all associated data in industry standard forms – i.e. NOT reliant on your Avaya Contact Recorder to continue to access them.

As an inherent part of this solution they also have:

7. Their recordings in industry standard wav files – as with the file export mechanism
8. All details of all recordings (far more than can fit in a filename) in an industry standard, easily readable XML file alongside every WAV file
9. All details of the recordings stored in an industry standard PostgreSQL database – for which the simple schema used to store recordings can be provided should they need it and for which backup/restore instructions are provided

Usage Report

The recorder is occasionally sold to service providers with a "per use" license. Charges are levied based on the actual usage made of a recorder each month. This report can be enabled in such cases to provide the necessary billing information.

Enabling the Report

To enable this feature, you must add the following line to the properties file:
`usage.reporting=true`

Content

The Usage Report page shows a summary of the licenses that have been used over a specific period. The default reporting period is the previous calendar month.

To generate a report for a different period, enter the date range in the calendar controls, and click **Refresh**.

The report shows

- The maximum number of concurrent recordings
- The maximum number of telephone replay ports (if any) configured at any time during the reporting period

Configuration records are retained for 3 months by default. Each night after that period has elapsed, a background job deletes any records older than this.

Accessing through URL:

This report can be accessed directly, without having to use the normal system administration pages.

To access the report by way of its URL, enter the following line in the navigation bar of your browser:

`http://`*`myservername`*`:8080/servlet/report?from=`*`t1`*`&to=`*`t2`

Where ***`t1`*** is the start time in UTC seconds and ***`t2`*** is the end time in UTC seconds.

Note:

These times are in seconds not milliseconds and do not attempt to correct for leap-seconds. Unless an administrator made configuration changes within a few seconds of midnight, this will not affect reporting on monthly boundaries.

The names of recording modes and other information that can be localized are returned according to the language preferences established for the interface. If localized terms are not available, the returned values are in English.

Accessing the Usage report in a log file

If you request the usage report by way of its URL and the request is successful, the recorder writes the usage data to a log file called **usage.log** in the **/logs** directory beneath the install path (which is typically **/opt/witness** on Linux and **D:\Avaya\ACR152** on Windows).

You can access and view this log file in any text viewer. Its content is the same as the report returned for the URL request. Each time you request a usage report, a new log is created that overwrites the previous one.

If a request for the Usage report is not successful, no log file is written for it. You should examine the return value of the URL request for an indication of the error conditions (bad time parameters, configuration has been tampered with).

Party Statistics Usage Report

The recorder is occasionally sold to customers who need to know how many recordings have been created per station, agent or skill over a period of time. In such cases this report can be enabled. Note this report is limited to Avaya Communication Manager users and is not available for other party types on other switches.

Enabling the Report

To enable this feature, you must add the following line to the properties file:
usage.partystats=true

The default value is `false`

Content

The Party Statistics Report shows a summary of the recordings that have been made over a specific period against different party types. The party types supported are Station, Agent ID, VDN and Split for the Avaya Communication Manager Switch.

The report shows

- The total number of recordings across each party type.
- The usage in time across each party type.

The report shows the number of calls and the number of recording segments as well as total duration of these. These may differ slightly from the figures produced by CMS reports due to very short calls (<0.25s) not being stored (so CMS counts 1, ACR 0) and how a call being taken again on the same skill by the same station is counted (CMS counts 2, ACR 1) and if a call is transferred to stations or queues that ACR cannot see, it will be counted again should it return to ACR's view without having changed skill (CMS counts 1, ACR counts 2). There are many different reasons as to why these reports will differ: calls and recording segments are not the same; call IDs may change on transfer; really short recordings are deleted.

Accessing through URL

This report can be accessed directly, without having to use the normal system administration pages. To access the report by way of its URL, enter the following line in the navigation bar of your browser:

http://*myservername*:8080/servlet/report?from=*t1*&to=*t2*&cmd=partystats

Where ***t1*** is the start time in UTC seconds or in yyyy-MM-dd'T'hh:mm:ss format (e.g. 2014-11-01T00:00:00) and ***t2*** is the end time in UTC seconds.

Note:

These times are in seconds not milliseconds and do not attempt to correct for leap-seconds.

The names of the party types and other information that can be localized are returned according to the language preferences of the user's browser (when requested interactively) or those set for the server (by property setting **server.locale**) for automated reports. If localized terms are not available, the returned values are in English.

Accessing the Party Statistics report in a log file

If you request the usage report by way of its URL and the request is successful, the recorder writes the usage data to a log file called `partystats.log` in the `/logs` directory beneath the install path (which is typically `/opt/witness on Linux` and `D:\Avaya\ACR152` on Windows).

You can access and view this log file in any text viewer. Its content is the same as the report returned for the URL request. Each time you request a report, a new log is created that overwrites the previous one. If a request for the report is not successful, no log file is written for it. You should examine the return value of the URL request for an indication of the error conditions (bad time parameters, file access errors).

Automated Party Statistics Reporting

If you set the following property:

`usage.dailystats=true`

Then, each night, between 1AM and 4AM, a text file is created that lists the number of contacts, number of individual recording files ("segments") and total duration recorded for each station, skill, agent and VDN. This file is located in the recorder's "logs" folder and is named partystats.*yyyy-mm-dd*.csv where *yyyy-mm-dd* is the date of the day just gone.

A second file is updated every night with the same information but covering the month to date - so a running total during the month which, after the end of the month, gives the total for that month. This file is also located in the logs folder but is named **`partystats.yyyy-mm.csv`** where ***`yyyy-mm`*** is the month the report covers.

Tip:

The recorder purges any files in its logs folder that have not been updated for more than 30 days. So if you want to keep a permanent record of these reports, you must take a copy before they are 30 days old.

These automated reports are generated in the language determined by property setting **`server.locale`**. Use Java standard locales e.g. **`server.locale=en_US`**.

Report Format

The interactive report includes a preamble row identifying the recorder and the time range covered. This is not present in the automated reports as they are located on the server they come from and cover the time-span indicated by their filename. (This makes it easier to concatenate multiple such files into a single spreadsheet).

Optionally, there is a header row (which you can disable with **`usage.header=false`**) describing the columns which are:

1. Type of device
2. Device Address (phone number)
3. Date column (in automated reports this is either the day or the month to which the report refers)
4. Calls - the number of distinct phone calls (or "contacts") recorded - i.e. a consultation call does not count as an extra call)
5. Rec - recordings - the number of files (hold/retrieve makes a fresh file, consultation call is a fresh file etc.)
6. Usage - the number of seconds of audio recorded from the device.

An example is shown below. The original is a tab separated ".csv" file but for ease of reading here, has been formatted as a table:

`8309732014-10-17T00:00:002014-10-18T00:00:00`

Type	Address	YYYY-MM	Calls	Rec	Usage
Agent ID	10173	2014-10	87	104	21258.475
Agent ID	10463	2014-10	24	29	12544.921
Split	20029	2014-10	31	49	2587.234
Station	24148	2014-10	50	56	5676.456
VDN	27142	2014-10	54	61	6456.325

Configuration Report (Communication Manager only)

Some customers need to report daily on their recording configuration - i.e. which stations, agents, skills or VDNs are configured to be recorded in Bulk recording mode and which stations are used for other recording modes - on a daily basis and to be able to see easily the changes made to this list during the day.

Background

Recording configuration is done on the Master and rippled out to the Standby and Slave servers. This information is already optionally written to the Audit Trail - with a snapshot nightly at midnight (if **`usage.reporting=true`**) and individual entries for each change. However, the audit trail contains a lot of extraneous information and while a user can filter by date range and configuration type there is no facility for automatically doing so daily.

Automatic Configuration Reporting

Each ACR can be optionally configured to log the bulk recording targets and assigned on demand or meeting recording ports that it is configured with for each data source:

- into to a tab separated ".csv" file (which is thus easy to import into Excel (use **Data > From Text**)
- named after the recorder's serial number e.g. **`890001_config.csv`**
- written to the recorder's logs folder alongside **`acr.log`**

Configuration

To enable this automated report, set the following property settings: **`usage.reporting=true config.reporting=true`**

Optionally, the header row can be turned off by setting: **`config.headers=false`**

Although only required on the Master, (as all configuration is done via the Master), applying these properties to at least the Standby server(s) as well ensures that a nightly snapshot is written there even if the Master is down overnight. Manual merging of the appropriate rows from the Standby's files into the overall report would be required in this case to give an uninterrupted report. This will be a very rare event so manual intervention should be acceptable.

File Management

In comparison with the other logs in this folder, the size of this configuration log is insignificant but it should probably be purged (manually) at least annually.

Users *must* take a *copy* of the file as required rather than open it in situ - as some applications, (e.g. Word) will apply a write lock to it - preventing fresh entries from being appended. This will result in error messages in the acr.log file (containing "**`IOException writing config log:`**").

File Contents

The file is appended to (or created if not already present) and contains:

1. Optional header row (default on - disable with property setting **`config.headers=false`**)
2. A snapshot of the current configuration every time the recorder is restarted and at midnight each day. This results in one row in the file for each type of device being targeted (stations, splits, VDNs, agents or CoRs). The row lists each device of this type that is configured (as opposed to showing ranges which are harder to search within). If a recorder has just started, an address has only just been entered or if an invalid address is included, this will show on a separate row listing devices of type "Unknown". 3. Additions to and deletions from the recording mode are noted as they occur.
3. Where recording is targeted by Class of Restriction (CoR) the list of stations in the CoR is shown on a row labeled "Refresh" each time the list is refreshed via SMS.
4. Changes to the extent of existing ranges are treated as a deletion of the previous range and addition of the new range - at the same timestamp.
5. The text within the file is in the language set by the **`server.locale`** property. Use Java standard locales e.g. **`server.locale=en_US`**

Note that this is not intended to replace the Audit trail - which includes more detail. As such, it does not cover all changes that can impact what is recorded. For example, **Advanced** settings being changed can affect what is recorded for a given target. Note that the Audit trail can also be exported as a csv file. This report merely gives a simple summary of recording targets as these are changed over time.

Column Definitions

The columns show:

1. Date - (in short date format determined by **`server.locale`** setting)
2. Time - (in short time format determined by **`server.locale`** setting)
3. Mode - Recording Mode
4. Action - Type of entry (Snapshot, Added or Removed)
5. Delta Targets - Number of recording targets impacted (0 on Snapshot, positive on Added, negative on Removed)
6. Net Targets - how many targets are configured after this event
7. Device Type - which type of device this row describes
8. Address(es) - name(s) of target(s) impacted by a change or of all targets present when a startup or midnight snapshot is taken. Each individual address is listed even if these were entered as contiguous ranges

Selective Record Barring

Recent legislation changes mean that some customers may need to block recording of calls to or from specific area codes (e.g. California).

Configuration

It is possible to bar recording of calls to or from certain numbers. To configure such a “recording bar”, add the property file entry:

```
recording.barred=<regular expression>
```

Where <regular expression> is a regular expression that will match the digit strings to be barred from recording.

Tip:

See http://java.sun.com/docs/books/tutorial/essential/regex/index.html for instructions on how to form a regular expression.

Example

The following example shows how to bar calls to or from area codes 234, 567 and 890 – where the recorder is situated in area code 234 (which therefore does not have a 1 in front of it, unlike the others which may or may not). The trailing periods (there are seven of them) are important – as this forces the pattern to match only numbers with 7 digits following the area code. **`recording.barred=((234)|(1?567)|(1?890)).......`**

Any recording that is barred due to matching the digit pattern specified will cause an INFO level message to appear in the log file.

Limitations

1. This feature does not apply to On Demand or Meeting Modes.
2. To bar incoming only or outgoing only, first determine the digit patterns that are used. You may be able to change the outbound dial plan to always prefix with a ’1’. The "1?" matches calls with or without a preceding ’1’. Remove the ’?’ to require a ’1’ before the pattern is matched.

After Call Work without Business Rules

If the only reason you are adding a WFO Business Rule is to enable a fixed duration of After Call Work recording, this can be achieved instead by setting a property value.

Set property **`acw.default=`*nn*** where ***nn*** is the number of seconds to continue screen recording after the voice call ends. This will then be applied to all calls that have NOT triggered a WFO Business Rule.

Altering Translations

The Avaya Contact Recorder supports a number of different languages. The string displayed to a user is determined by looking up a "resource" in a particular language. The translations that the recorder ships with can be corrected without needing to wait for a subsequent release. To do this, log in as a System Administrator and access the Maintenance page (at `/servlet/acr?cmd=mtce`)

On this page you can enter the name of the resource to be changed, the string to be displayed instead of the one shipped with the program and the language to which this translation relates. **Custom Translations** entered are stored in the database and are permanent. You should only use these fields under direct instruction from Avaya support staff who will advise you of the resource name to enter.

Appendix A: Technical Reference

This appendix provides technical details about the Avaya Contact Recorder system.

The main sections in this appendix are:

Recording files on page 302

- Internal Database on page 302
- Recorder Interfaces on page 302
- Recording Attributes on page 312

Recording files

Voice recordings are stored in an industry standard **wav** file. Screen recordings are stored in a proprietary format in files with extension **scn**. When each call is completed and as each recorded call segment becomes available, the recorder updates its local database with a record of the call segment. These files are stored in a hierarchy of folders beneath **/calls** on a Linux system or the path specified by the administrator on a Windows system.

Every recording results in an XML file and zero or more media files (.wav, .scn etc.).

WAV files

The **wav** files contain the actual audio of the recording. You can double-click some **wav** files to play them directly. Others are in audio formats that are not directly supported by Microsoft's Media Player. This applies to most recordings made by this recorder. These must be converted into a supported format before they can be played. Since the recorder's Search and Replay application does this conversion automatically, you do not need to access these files directly.

XML files

The **xml** files contain details about the recorded call segments. Although most users typically search against the recorder's database of calls, you can view these files directly in a browser if required. Within each **xml** file there is:

- All the details known about this recording. Most of the information, but not all, is inserted into the calls database. Some of the information is only of interest for diagnostic and maintenance purposes.
- Start and end time in ISO format giving local time and offset from GMT.

SCN files

The **scn** files store the recorded screen content in a proprietary format. These can only be played using the search and replay tools provided with Avaya Contact Recorder.

Internal Database

If you have retained all of the **xml, wav** and **scn** files as described above, then you have kept all of the details about the recordings you have made. However, the system uses an industry standard database (PostgreSQL) to hold this information in more readily accessible forms. This database is located on the recorder itself. The database stores details of the recordings as well as details of the recorder's configuration.

Recording details

The call details database uses approximately 2KB per recording (in the absence of user defined fields). Each voice segment counts as one recording regardless of whether it is in mono or stereo. Each screen recording also counts as one recording.

To allow you to search for calls easily, the details of recordings are inserted into this database. It contains one record for each call segment recorded and additional records for each party or owner of the recording. The information stored for each call is described in detail under Recording Attributes on page 312.

Configuration details

Several tables hold details of system configuration, such as port assignments, file paths, timeouts and user authorization rights.

Recorder Interfaces

The interfaces supported by the recorder are described below (working clockwise round the diagram, starting at top left).

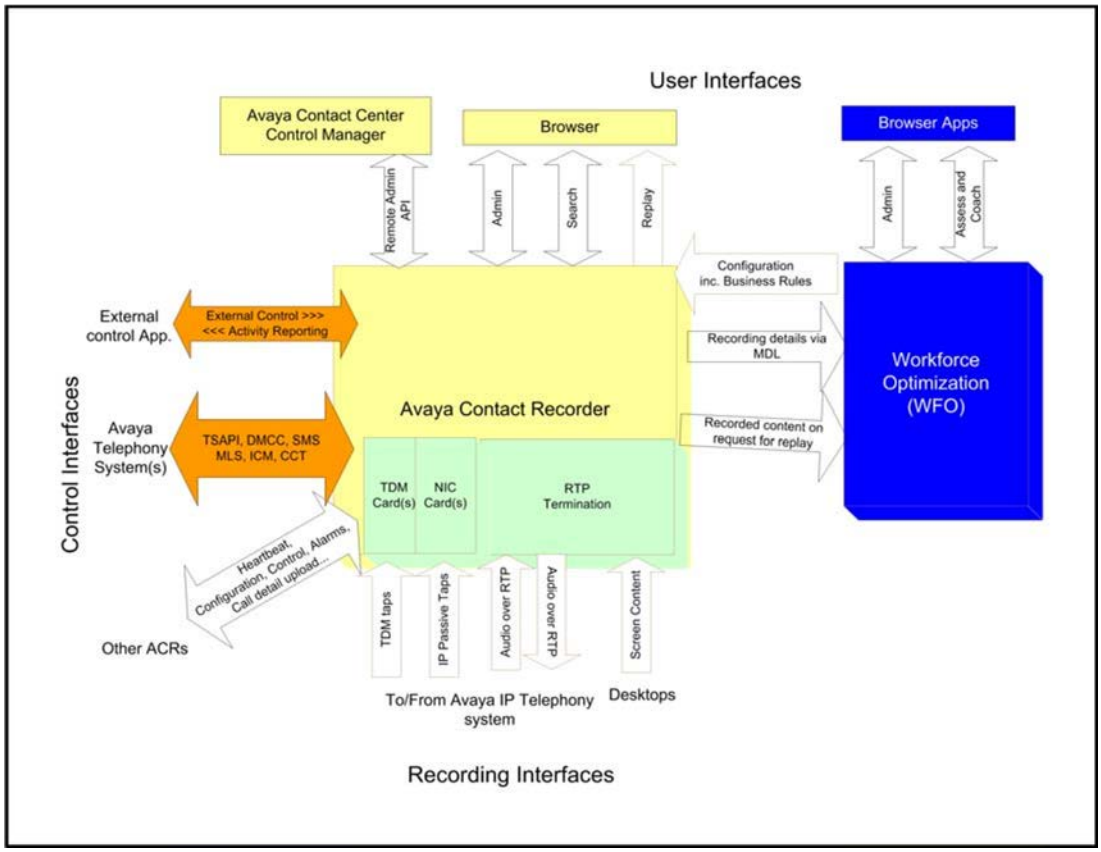

HTTP/HTTPS Interfaces Offered

The recorder uses the Tomcat (see www.apache.org) web servlet container to offer a number of services via HTTP and/or HTTPS (on ports 8080 and 8443 respectively). These provide:

Avaya Control Manager Interface

This SOAP interface allows administrators to configure many recorder features through Avaya Control Manager. These include:

- settings on the **General Setup** pages
- user accounts
- stations to be recorded in Bulk recording mode

This interface is located at **`/remoteadmin`**

The service definition is available from the service itself at **`/remoteadmin?wsdl`**

Administration Interface

This provides administrators with access to configuration and status monitoring pages.

Search Interface

End users access this to search for call recordings that match specific criteria.

Replay (Retriever) Interface

End users and other applications (such as WFO) use this interface to retrieve the content of a specific recording.

Call Details Interface

The details about a recording can be provided on request. This is used by the bulk export feature of the integral search and replay application and the Central Replay Server to populate its database.

Search and Replay API

External applications may search for and retrieve the content of recordings using this interface. Full details are available on request.

Communication Manager

The recorder interfaces to the Avaya components via several mechanisms:

DMCC

DMCC runs on an AE Server and provides softphone registration and signalling services.

TSAPI

The Avaya Contact Recorder software exchanges TSAPI messages with an AE Server if Bulk recording is used or if an external controller asks the recorder to establish single-step conferences.

SMS Web Services

If using CoR based recording, the recorder uses this interface to determine which stations are in which CoR.

Audio over RTP

The softphones on the recorder terminate RTP streams over which the audio to be recorded flows. These sockets are connected to whichever VoIP resource (for example, Media Processor) is providing the conferencing bridge through which the call is recorded.

The recorder uses ports in the range 10000-20000 (by default) using two ports per channel.

IP Passive Tap

Ethernet packets can be picked up from a "SPAN" or "Mirror" port on an Ethernet switch and connected to one or more NICs for recording.

TDM Taps

One or more trunks and/or phones can be tapped and connected to the appropriate Ai-Logix card(s) to allow recording directly from these devices.

CS1000

The recorder interfaces to these Avaya components via several mechanisms:

Meridian Link Services (MLS)

Master and **Full Standby** recorders connect using TCP/IP to the MLS feed from the Avaya Contact Center Manager Server. These recorders interpret the telephony events occurring in real-time and whichever is active instructs the server to forward call data from the required IP phones to the required ports on the recorder(s) when calls are to be recorded.

Audio over RTP or DTLS

When the recorder instructs IP phonesets (via MLS) to perform Duplicate Media Streaming, the voice packets are sent using RTP to the recorder. If licensed for encryption, these streams will use DTLS.

The recorder uses ports 12000 upwards, allocating two ports per recording channel.

TDM Taps

One or more trunks and/or phones can be tapped and connected to the appropriate Ai-Logix card(s) to allow recording directly from these devices.

Avaya Aura® Contact Center

The recorder interfaces to the Avaya components via several mechanisms:

CCT

The recorder uses the CCT web services to track activity on the switch and to ask for recorder ports to be invited into calls that need to be recorded. The recorder listens on port 9010 and sends to port 9080 by default.

Avaya Oceana™

If this is configured, the Master and Standby servers may use the Subscription Interface to establish a link to the Reporting Interface on the specified address and port. The former is (by default) an HTTPS interface on the port configured on Avaya Oceana™ (default 443) and the latter, an inbound HTTPS connection to port 9110 (or next available port above this if already used) on the Master and Full Standby Avaya Contact Recorders.

Other Data Sources

Where configured, these can include:

- Avaya Proactive Outreach Manager (POM) - default link over TCP using port 7999 on the POM server or secure link over TLS 1.2 via port 7998 on the POM server.
- Avaya Proactive Contact (PCS/PDS) uses multiple sockets. See PCS documentation for details.

Screen Recordings

If licensed for screen recording, the system will also use this interface.

Screen Content

Each recorder communicates with the workstations whose screen it is recording. This is a TCP connection to port 4001 (by default) on the client workstation (or, in the case of thin client recording, with the server hosting the Windows session).

Windows Account Logon/off

Workstations that have been configured to track logons will send details of account logon/off to port 29522 on the recorder.

Workforce Optimization ("WFO")

In addition to the Replay (Retriever) Interface already described on page 280, Avaya Contact Recorder interacts with a WFO system over the following interfaces:

EMA Interface

System configuration information is transferred via TCP/IP ports 7001 and 7002 on the WFO server hosting the "EM" component. It uses ports 8080 or (with SSL) 8443 on the Avaya Contact Recorder to do this.

MDL Interface

Details of recordings made are passed to WFO via HTTPS on port 443 or HTTP on port 80 of the WFO server configured with the "APP" role.

MAS Interface

Real-time device and call status is provided via HTTP on port 29520 (default).

Phone Replay

The replay applet establishes contact with RTP ports on the recorder that are in the same range (10000-20000) as recording ports.

Live Monitor (Audio)

Screen monitoring traffic flows between the client viewing the screen and the workstation being monitored – but audio monitoring is done via the recorder.

ACR Live Monitor

The client establishes contact with RTP ports on the recorder that are in the same range (10000-20000) as recording ports. The RTP then flows back to the client – without the need to open the firewall on the client.

WFO (“MAS”) Live Monitor

MAS tells the ACR which ports to send audio to on the client. These are fixed ports (8500-8503) and must be opened in the firewall on the Windows PC that is trying to monitor audio.

Other Recorders

Avaya Contact Recorder servers establish links between each other. Configuration, alarms and recording instructions are passed over these links. These are TCP/IP socket interfaces.

Master and Standby recorders listen for connections from other recorders on TCP/IP port 1209 by default and communication over these links is encrypted.

External Control Interface

The recorder supports a simple TCP/IP command interface (now deprecated). This allows other applications to control recording directly and/or add further user defined data fields to recordings.

AET/DPA Interface

The recorder listens for incoming HTTP commands on port 3020.

Database Upload Interface

When configured with either a Central Replay Server or (defaulting to) a Master and/or Full Standby server performing this role, the recorder uploads details of calls into a database on the appropriate server.

To Central Replay Server

To upload call details to the Central Replay Server (CRS) a recorder uses a web service on the CRS, to inform it that a new recording is available. The CRS in turn uses a web service on the recorder to retrieve the details. Both of these interactions default to HTTP on port 8080 or HTTPS on port 8443.

Summary

System	Interface	Protocol	Local	Remote	Direction
All	HTTP - Admin and Replay etc.	TCP	8080		Inbound
	HTTPS - Admin and Replay etc.	TCP	8443		Inbound
	RTP Audio	UDP	10000 - 20000		Bidirectional (CM) Inbound (CS1000, AACC) Outbound (Live Monitor)
	Between Avaya Contact Recorders	TCP		1209	Towards master and standby
If installed and configured	Screen Capture	TCP		4001, 29522 (default)	Inbound
	Central Replay Server	TCP		8080 or 8443	Bidirectional
	Dialers, External controllers	See the external system's manuals for interface details.			
If WFO present	EMA	TCP		7001, 7002	Outbound
	MDL	TCP		80 or 443	Outbound
	MAS	TCP	29520		Inbound

System	Interface	Protocol	Local	Remote	Direction
	Live Monitor	UDP	10000-20000	8500-8503	Firewall on Windows client trying to monitor must be open to ACRs on 8500-8503
When recording Communication Manager	DMCC	TCP		4722	Outbound
	TSAPI	TCP		450, 1066[1]	Outbound
	SMS web services	TCP		80	Outbound
When recording Avaya Aura® Contact Center	CCT web services (client)	TCP		9080	Outbound
	CCT web services (server)	TCP	9010		Inbound
When recording Avaya Aura® Contact Center via SIP and/or SBC(s) via SIPrec	SIP proxy	TCP		5060	Inbound
		TLS		5061	Inbound
When controlled by AET/DPA	EQConnect	TCP	3020		Inbound
When Avaya Oceana™ configured	Reporting Interface (server)	TCP	9110 (or next higher if already used)		Inbound

[1] 450 (service port), 1050–1065 (unencrypted TLINK), 1066–1081 (encrypted TLINK). Typically, this means 450 and 1066, but is configurable in the AES.

System	Interface	Protocol	Local	Remote	Direction
	Subscription Interface (client)	TCP		As explicitly configured in Oceana™ but default to 80 (HTTP) or 443 (HTTPS)	Outbound

Recording Attributes

The Avaya Contact Recorder stores some "tagging" information against each recording that it makes. Other information pertains to the overall telephone call or to one parties' involvement with that call. This section:

- Explains where and how data is stored
- Defines the terms "Call", "Contact", "Session", "Segment" and "Party"
- Explains how calls are identified by the recorder and how this relates to the underlying switches own call identifier
- Describes the available attributes; how they can be added to a search and replay layout and, where appropriate, consolidated into an associated WFO database

Overview

Avaya Contact Recorder tracks the phone calls that are made on a switch in real-time. This results in one or more recordings. These in turn result in entries being made in the recorder's database and (optionally) in the database of one or more other Avaya Contact Recorders acting as "Central Replay Servers".

Real-time Tracking

Avaya Contact Recorder:

- Tracks CTI events from one or more data feeds in real-time.
- Builds a CSTA-like in-memory model of the state of each **CALL**.
- Models each call as a current (and historical) set of connections ("**CONN**"s).

Each connection (**CONN)** identifies a specific **DEVICE** which is linked to the **CALL** for a certain period of time in a specific connection **STATE**.

Some connections are **TRANSIENT** while others **PERSIST.**

A call is **SEGMENT**ed into time periods during which the set of (implicitly persistent) **CONN**ections is constant i.e. any change to the connections results in a new **SEGMENT** starting. Some segments are recordable and others are not.

XML Files

For each recording, a snapshot of some of the call information is taken at the start of the recording. Some fields may be added during and at the end of the recording segments.

This information is stored in an XML file alongside the content (audio and/or screen capture) of that recording - and is copied along with the content should that be archived. This gives a fallback record of the data associated with the recording allowing the retrieved files to be of use even in the absence of any associated database records.

Recorder's Database

The recorder's (PostgreSQL) database is updated at the end of a recording with most - but not all - of the information present in the XML file.

Implicit in the above are the facts that

- **Nothing** is stored in the database during a call segment.
- **Nothing** is stored for call segments that did not result in a recording.

Central Replay Server(s) Database(s)

In a system with more than one Avaya Contact Recorder, it is normal practice for all other servers to upload details of their recordings to one or more servers that fulfil the role of "Central Replay Server". Such servers allow a user to search against their database to find recordings made on any of the recorders that feed this server with copies of their recordings' details.

In many systems, the Master (and Standby if present) act as Central Replay Servers for each other and for their population of connected Slave servers. In very large systems, dedicated Central Replay Server(s) may be provided to reduce the load on the Master and Standby.

The other reason for having a dedicated Central Replay Server is if custom integration is required into the PostgreSQL database. (e.g. additional tables or data feeds populating it). This is NOT supported on a live recorder as the load is indeterminate and out of Avaya's control.

The information uploaded to the Central Replay Server is always the same as is stored on the Avaya Contact Recorder that is feeding it - but is typically delayed by at least several seconds (10 second batch job interval by default) and may be much further behind if the CRS or path to it is unavailable.

Definitions

Before listing the attributes stored against each recording, it is important to understand the concepts of "Calls", "Contacts", "Sessions", "Segments" and "Parties".

Call

Each telephone call that is active on an Avaya switch has a call identifier and this is stored as an attribute of a recording - its Call ID. The format of this identifier varies from switch to switch. In the case of Communication Manager, this Identifier is extended to form a UCID (Unique Call ID). Other switches may reuse call identifiers over and over again. In this case, such an identifier is not sufficient to uniquely identify a call unless it is combined with some form of timestamp information to distinguish it from previous and subsequent uses of the same identifier. This combined, globally unique identifier is also stored but is kept internal to the recorder.

Contact

Note:

Contact related information is only stored for calls recorded using Bulk recording mode. Recordings made in other modes will not display this information.

When the first party on a call (the "calling party") starts to make a call, this is treated as a new "contact". A contact extends into other calls made while connected to this original call. Only when all such connected calls are over does the contact end. The total duration of this contact is tracked. Each recording is marked as being part of one and only one contact - with its Contact ID.

Session

Note:

Session related information is only stored for calls recorded using Bulk recording mode. Recordings made in other modes will not display this information.

Within a contact, several other parties may be involved with the call(s) made. Each of them cares about how they were treated during the contact. Did the other parties leave them listening to ringing tone; did they get put on hold many times or for an excessive period; were they transferred many times etc. Within the search and replay application, the user can search for fields that reflect the first session within the contact i.e. how the calling party was treated. (All sessions with appropriately configured internal parties are available to the optional Quality Monitoring application).

Segment

Each "call" is recorded as one or more recordings - each of which normally extends for only one "segment" of the overall call. A segment ends whenever the parties on the call change. This is because the recording rules typically operate on the basis of which parties are on the call and the security mechanisms that determine who will be allowed to replay a recording also operate on the basis of who was on the call - hence any change to these forces a new segment and, if appropriate, a new recording to start.

Party

Many of the fields available refer to the parties or devices involved with the phone call - or, more specifically, with a particular recording - hence call segment. These range from agents to VDNs to public phone numbers. The table below shows the types of party that may be associated with a recording in the database; which type of switch each relates to and the numerical "party type" identifier that is used in the recorder's database to indicate this sort of party.

Each party has:

- an "address" (previously referred to as its "number") which is normally, but not always numeric
- and *may* also have an alphanumeric "description" (previously referred to as its "name"). Not all switch CTI feeds provide this.

Switch	PartyTypeID	Party Type	Notes
All	50	DNIS	Some switches (e.g. AACC) provide DNIS for all calls, others only for incoming calls. ACR follows the switches definition of "DNIS" (whereas WFO treats "DNIS" as synonymous with "called party".
	51	External Number	
	52	Unknown	
CM	100	Agent	Alphanumeric description or "name" as configured on the switch is usually provided via TSAPI and stored along with the (normally) numeric "address".
	102	Split	
	103	Station	
	104	VDN	
	108	Dynamic	
	109	Other	
	111	Announcement	

Switch	PartyTypeID	Party Type	Notes
CS1K	201	CDN	Name information not available via MLS. Numeric address only stored.
	230	DN	
	232	Position ID	
	234	Line appearance	
	235	Agent	
	236	Skillset	Skillset name provided and should be used in preference to the numerical Skillset identifier.
	240	Activity Code	May be overwritten during a segment. Only the last code is stored.
AACC	301	AACC Agent	
	302	AACC Skillset	
	303	AACC CDN	When a recorder is restarted, it may first encounter these numbers in CTI from the underlying switch - identifying them as external numbers. Only once an AACC event has shown these to be CDNs can they be immediately identified as such. This can lead to the CDN number appearing as an additional redundant "external party" on the first call to each AACC CDN after a restart. You can avoid this by specifying the CDNs used on the AACC's administration page so that they are recognized as such from startup.
	304	Unknown party	
Dialers	401	Dialer Agent	
	402	Dialer Skill	
	403	Campaign	

Call Identifiers

Each switch identifies its calls differently. The recording system must have a way of uniquely identifying each call that it has recorded - and not be confused by the same id being used again by the switch - which really only needs to distinguish each live call from all other live calls.

The Avaya Contact Recorder uses a single pool of Call objects so must ensure that the ids used to look calls up are unique across all call types that may be present (underlying switch, overlay switch such as AACC plus any number of heterogeneous dialers).

For each type of switch, the Avaya Contact Recorder derives its own variant of a CallID and how it maps the ID provided by its CTI feed(s) to this object - but each call identifier has:

- a 64-bit (signed) **long** value which must be unique over the lifetime of the recording system - and is used internally but never normally shown to the end user.
- a "**native**" string representation which is shown to the user on the replay screen and may be searched for as a string. This is in the form used by the switch itself e.g. a Communication Manager UCID such as 00001012341212345457.

In many cases the native call identifier is itself a representation of an underlying long (64-bit) value - and in such cases the Avaya Contact Recorder will use this value where it can. The various types of callID are summarized in the table below and discussed in more detail thereafter.

<table>
<tr><th rowspan="2">Call on</th><th rowspan="2">Type ID</th><th colspan="3">64-bit callid</th><th rowspan="2">"Native" Callid</th></tr>
<tr><th>Bits 48-63</th><th>Bits 32-47</th><th>Bits 0-31</th></tr>
<tr><td rowspan="2">CM or POM</td><td rowspan="2">1</td><td>Switch(S)</td><td>Callid (C)</td><td>UTC (T)</td><td rowspan="2">SSSSSCCCCCTTTTTTTTT
(standard Avaya UCID style)</td></tr>
<tr><td colspan="3">=Original CM provided UCID</td></tr>
<tr><td rowspan="2">CS1K</td><td rowspan="2">2</td><td>HLOC</td><td>Callid (C)</td><td>UTC (T)</td><td rowspan="2">HHHHCCCCC
See MLS spec for HLOC encoding. Not direct map of hex digits and may be fewer than 4 chars.</td></tr>
<tr><td colspan="2">= Original NetworkCallId provided by MLS</td><td>ACR’s Time</td></tr>
<tr><td>AACC</td><td>5</td><td>ACR assigned switch ID</td><td>Cyclic identifier assigned by ACR</td><td>ACR’s Time</td><td>May be a GUID, a UCID or other identifier passed in by originating switch.</td></tr>
</table>

Dialer (except POM)	4	ACR assigned switch ID	Low 48 bits of dialer's callid	Varies

User Defined Fields

In its internal PostgreSQL database, Avaya Contact Recorder can store:

- any number of user-defined fields (UDFs)
- each with a name of up to 50 characters
- each of arbitrary length (all are stored as TEXT fields)

Some of these fields can be populated by the Avaya Contact Recorder itself while others can be provided by external applications as described in Appendix G: External APIs on page 416.

Note:

As of 15.1FP1, any User Defined Fields associated with a call will be carried forward onto the resultant call when calls merge. This means that any tags on a call will persist through a transfer or conference. To revert to previous behavior, which requires that you explicitly tag the call again after transfer if you wish the fields to be marked against that segment, set the property **`call.persistudfs=false`**.

Built-in User Defined Fields

In addition to the basic CTI information stored with every call, other CTI fields may be accessible to the recorder as it processes events from the switch. These are not of interest to all users but can be stored in user-defined fields. Only the more commonly used ones default to being stored, are immediately available through the default search and replay layouts and are passed to WFO. Other, less useful ones default to not being stored but can be assigned to a user defined field using property file settings. Similarly, those that are on by default can be forced off (by placing nothing after the equals sign in the property file entry) or into differently named user defined fields with the property file settings shown below.

Attribute	Property Setting	Default	Notes
Activity Code	activitycode.fieldname	activitycode	CS1000 only. Last activity code set on each call segment.

Dialer call attributes	***dddddd.*** field.***fffff*** (where ***dddddd*** is the dialer name and ***fffff*** the field name).	Null	Available fields vary according to dialer type and how it is configured.
User to User Information (UUI)	uui.fieldname	Null	Communication Manager, text (not binary) UUI only.

Externally Provided User Defined Fields

The Avaya Contact Recorder automatically stores new user defined fields under whatever name is provided in the simple XML tags passed via the webXAPI or other control interfaces (AIM, DPA etc.).

Note:

As these field names are widely used within the product and other products to which it integrates, they should be short (50 char limit) and use only standard ASCII a..z characters. Names are forces to lowercase but comparisons are made case insensitively to reduce the potential for error.

Use of other characters may cause unpredictable results in other applications. The values stored within them and the internationalized strings used to display their names can, of course, use full 16 bit characters and spaces etc.

Search and Replay Attributes

The table below describes the various CTI fields that are stored within the recorder's database and hence can be easily added to a search and replay layout. The description also explains how to use the field selection box and its neighboring parameter field at to the top of the **Field Editor**. The fields are grouped logically below rather than alphabetically and follow the order shown in the drop-down list presented by the **Field Editor**. Where attributes vary according to the switch type or CTI available, this is also described.

Field	Param	Layout Type		Notes
		Segment-based	Session-based	
Start Time	N/A	The start date and time of a recorded segment.	The start date and time of the contact of which this session forms a part. Note that this session may start some way into the contact.	This field is mandatory (unless a specific call set is selected). As the search criteria for this field is normally a time range relative to the date the search is being done, it does not make sense to enter a fixed or pre-selected date/time range. Instead, enter the (integer) number of days back from today in the default search criteria box. So "0" sets the timespan to "today"; "1" to "Yesterday through to today"; "28" to "today and the previous four weeks" etc.
Session Type	N/A	N/A	Use to filter by the type(s) of session to be shown: External/Customer Internal parties Calling party Called party Other party All When shown as a column, this field shows the session number (1=first party etc.) and the type of entity that "owns" the session.	This field is mandatory in session-based layouts. Note that selecting "All" will give at least two rows in the results table for each contact - one for each party. If presetting the selection, use the appropriate ENGLISH word from this list "External, Internal, Calling, Called, Other, All".

Field	Param	Layout Type		Notes
		Segment-based	**Session-based**	
Owner(s)	N/A	The party or parties that "own" the recording. Your user rights must include at least one owner of a segment for you to be allowed to replay it.	The party or parties that "own" the recording(s) within the session. Your user rights must include at least one owner of a segment for you to be allowed to replay it.	
Participant	N/A	N/A	The party whose session this represents.	This field is mandatory in session-based layouts. Each party in a contact experiences the contact from their own perspective or "session".
Duration	N/A	The duration of the call segment to which the recording relates.	The duration for which this session owner was involved in the contact.	Duration shown in seconds or minutes and seconds.

Field	Param	Layout Type		Notes
		Segment-based	**Session-based**	
Direction	N/A	Direction of overall contact of which this segment is a part.	Direction of overall contact of which this session is a part.	If there are NO external parties (types 51 or 108) on the call, it is "Internal". Otherwise, if the first session on the contact was from an external party (type 51 or 108) then the call was "incoming" otherwise it was "outgoing". If you wish to "fix" this field's value, use one of the values 1 to 6 as follows: 1=Incoming only 2=Outgoing only 3=External only 4=Internal only 5=Not Outgoing 6=Not Incoming
Call ID	N/A	The switches own identifier for the call of which this segment is a part.	The switches own identifiers for the call or calls involved in this session.	The native identifier assigned by the switch. There may be several segments and sessions sharing the same Call ID. Internally, this id is supplemented with timestamp information to derive a unique identifier for the call. You can click on a Call ID to search for all segments of that call.
INum	N/A	Identifier of the recording associated with this segment.	Identifier(s) of any recording(s) of this session. Do not assume that the whole session was recorded.	An INum is a unique 15 digit reference number of a recording. The actual audio, screen and data files stored use this as their filename.

Field	Param	Layout Type		Notes
		Segment-based	**Session-based**	
Contact ID	N/A	Which contact this recording forms a part of	Which contact this session forms a part of.	You can click on a Contact ID to search for all segments or sessions of that contact.
All Parties	N/A	Displays all types of parties that have been stored in the database as being associated with the segment.	Displays all types of parties that have been stored in the database as being associated with any part of the contact of which this session is a part.	Can be used to make a very simple set of search results with a single column showing all the phone numbers, agents etc that were on the segment or contact. However, this can be confusing as not everyone recognizes which numbers are skills, agents, VDNs, stations etc. Normally the parties are spread across several columns using the fields immediately below.
Agents	N/A	Agent(s) that were connected to this segment.	Agent(s) that were connected to this contact.	Includes parties of types 100, 235, 263, 301 and 401.
Skills	N/A	The most recent skill or split that the call segment was routed via.	The skill/split(s) used during the contact.	Includes parties of types 102, 236, 302 and 402.
Services	N/A	The VDN or CDN used to route this segment.	The VDN/CDN(s) used during the contact.	For VDNs, a setting on the **General Setup** > Data Source admin page determines whether the first or last VDN is shown for each segment.

Field	Param	Layout Type		Notes
		Segment-based	**Session-based**	
DNIS	N/A	The DNIS that the call segment was directed towards.	The DNIS(es) that calls within the contact were directed towards.	Shows parties of type 50. (Technically, DNIS should only be present for an incoming call from the public network. Some switches treat the "called number" as the DNIS for internal calls). The recorder follows the lead of the switch(es) to which it is connected.
Other Parties	N/A	Other parties associated with the call segment.	Other parties associated with the contact.	Automatically excludes any party types that are being shown in other columns. It is a useful catch-all to ensure that any party tagging not shown by your other columns is presented to the user.
Specific Parties	xxx or xxx,yyy ,zzz	Parties of the type(s) specified in the parameter that were associated with the call segment.	Parties of the type(s) specified in the parameter that were associated with the contact.	Gives fine control over exactly which parties are to be displayed - e.g. allows separation of Dialer Agent from CM Agent.
Hold Count	N/A	How many times the calling party (first session) on the overall contact was placed on hold and hence unable to continue talking.	How many times the session's participant was placed on hold and hence unable to continue talking.	Note that a) This is NOT the same as the number of times the session participant put the call on hold. b) On a conference call, all other parties would have to put the party on hold to add to this count.

Field	Param	Layout Type		Notes
		Segment-based	**Session-based**	
Transfer Count	N/A	How many times the calling party (first session) on the overall contact was transferred to another party.	How many times the session's participant was transferred to another party.	Note that this is NOT the same as "how many times an agent transferred a call" (or attempted to transfer a call). It is how often the other party has left the call and been replaced by someone else.
Conference Count	N/A	How many times the calling party (first session) on the overall contact was involved in a conference with more than one party.	How many times the session's participant was involved in a conference with more than one party.	Dropping from three to two parties and returning to three counts as a new conference regardless of who the parties were.
Agent Count	N/A	How many agent sessions were involved with the contact as a whole.		Transferring back to a previous agent still counts as an additional 1 (so A transfers to B transfers to A gives a count of 3).
Ring Duration	N/A	How long the calling party spent unable to talk because the only other party on the current call had not answered.	How long the session's participant spent unable to talk because the only other party on the current call had not answered.	No one party on the contact need be present for the whole of this period. The originator may have dropped off the call before consultations finish for example.

Field	Param	Layout Type		Notes
		Segment-based	**Session-based**	
Total Duration	N/A	The total duration of the contact of which this is a part.	The total duration of the contact of which this is a part.	No one party on the contact need be present for the whole of this period. The originator may have dropped off the call before consultations finish for example.
Call History	N/A	The history of contact as seen by the calling party.	The history of the contact as seen by the session's participant.	
Bulk Trigger	N/A	True if the recording of this segment was made because of configuratio n under **Operations > Bulk** recording mode.	True if any part of the session was recorded because of configuration under **Operations > Bulk** recording mode.	
Rule Trigger	N/A	True if the recording of this segment was made because of a WFO Business Rule.	True if any part of the session was recorded because of a WFO Business Rule.	

Field	Param	Layout Type		Notes
		Segment-based	**Session-based**	
Call Sets	N/A	When used as a filter, this allows the user to select calls that are in a Call set - from a drop-down list of existing call sets. When shown as a column, this shows the names of any call sets into which the recording segment has been placed.	When used as a filter, this allows the user to select sessions for which at least one segment is in a Call set - from a drop-down list of existing call sets. When shown as a column, this shows the names of any call sets into which any recording segment with the contact has been placed.	If the user is permitted to lock or unlock recordings, the call sets shown will include these special call sets. Take care if using both segment and session-level layouts. A session will list the call sets that *any* of its segments have been placed in. This does not imply that all of its segments are in the call set as some may have been removed individually using a segment-based layout.
All UDFs	N/A	All user defined field tags on the call segment.	All user defined field tags on any call segment within the session. Note that a single session may have been tagged multiple times with the same udf field but different values - all of which will be listed here.	Places all populated user defined fields into a single string of the form *udfname1:udfvalue1 udfname2:udfvalue2* - i.e. a colon between udf name and value and a space between successive udfs. This allows a single filter to be used to search for any udf e.g searching for "Incl. custid:1" would find all recordings with a udf "custid" starting with digit 1.

Field	Param	Layout Type		Notes
		Segment-based	**Session-based**	
User Defined Field	*udfname*	The value (if any) of the specified user defined field for this call segment.	The value(s) (if any) of the specified user defined field for any call segment(s) associated with this session. Note that a single session may have been tagged multiple times different values - all of which will be shown here (space separated).	Unlike "All UDFs" above, these fields just show the value of the specified UDF in the results column. Use the filter and column titles to show the name of the UDF.
Custom Field (String)	*SQL*	Allows additional text fields to be derived from the data held within the database.		Please contact Avaya if you need to use these features.
Custom Field (Number)	*SQL*	As for Custom Field (String) above - but in this case the result is treated as a number for filtering purposes and in the column sort order.		

WFO Integration

Where the details of recordings are to be consolidated into a WFO system, a number of attributes are automatically provided, with fixed definitions, while other, user configurable fields can be added as required. See WFO documentation for procedures, limitations and definitions of these. The tables below describes how Avaya Contact Recorder supports the standard attributes defined in WFO.

Contacts

The following fields are always provided as part of the Contact data. These fields are used in a standard way within WFO but the **Notes** column highlights any known issues or differences in how Avaya Contact Recorder determines their content.

Field	Meaning	Notes
ID	Unique identifier for this contact	19 decimal digits, starting with "92" and ending with "1".
Start	Contact start date/time	
End	Contact end date/time	
Holds	Number of times all other parties on the call placed the call on hold	Calculated with respect to the original (calling) party on the call.
HoldDuration	Time spent with all other parties holding	Calculated with respect to the original (calling) party on the call.
Transfers	Number of times transferred	Calculated with respect to the original (calling) party on the call.
Conferences	Number of times call went from 2 parties talking to more than 2 parties	Calculated with respect to the original (calling) party on the call. If one party on a conference holds the call then retrieves, this counts as dropping to 2-way then back to 3-way again, hence adds to the number of conferences.
IsException	Whether or not marked as an exception	Only Business Rules configured to do so will mark a contact as being an exception.
ExceptionReason	0 unless IsException is true, in which case 1	
Dialed From (ANI)	Calling party address	WFO treats ANI as synonymous with "calling party" so this is provided, even on internal calls which, strictly speaking do not have an ANI.
Dialed To (DNIS)	Called party address	Originally called party (answering party otherwise). This is provided, even on internal calls which, strictly speaking do not have a DNIS.

Field	Meaning	Notes
		This field can also be stored as a User Defined Field in ACR if required. Set the property **`dialedtodnis.udfname=`*`xxx`* where *`xxx`* is the name of the field it will be stored as ("**`dialedtodnis`**" is recommended for consistency with WFO).
PauseDuration	Not supported	
WrapupTime	Not supported	

Sessions

Within Avaya Contact Recorder, every party that joins a call is assigned a Session. That party will have a device and may also be associated with an Agent.

Sessions Visible to WFO

By default, all sessions involving an internal party are passed to WFO but this can be overridden by setting set the property value **`core.consolidateall=false`**

Once this has been set, a session will only be sent to WFO if the logical device (station or DN) associated with the session is appropriately configured in WFO. This is summarized in the table below, where the right-hand columns show whether or not a session will be sent to WFO.

Session's Device	RecordingType configured in WFO for this device	Type of recording made	Sent to WFO if core.consolidateall=	
			true (default)	false
External	N/A	Any	No	No
Internal	Device Not configured	Any	Yes	No
	Record Always, Application Controlled, Override or TAG	Any		Yes

Session's Device	RecordingType configured in WFO for this device	Type of recording made	Sent to WFO if core.consolidateall=	
			true (default)	false
	Start at Business Rule	Bulk Recording		No
		Business Rule		Yes
	Any other recording type	Any		No

Standard Session Attributes

The following fields are always provided as part of the Session data. These fields are used in a standard way within WFO but the notes column highlights any known issues or differences in how Avaya Contact Recorder determines their content.

Field	Meaning	Notes
Telephony Contact Fields		
ID	The unique identifier of the contact of which this session is a part	Links this session to the appropriate contact.
User Fields		
StringExtension	The (normally numerical) address of the telephone device to which this session relates.	Either a Station (CM) , DN or Position ID (CS1K).
PbxLoginName	The agent identifier used to log in to the switch (if any).	Normally numeric but may be alphanumeric on e.g. dialers.
ID	The unique reference within EMA configuration for this employee (if known).	Found by looking up the PbxLoginName within the appropriate Data Source file.
NetworkLoginName	The computer (normally Windows) account name associated with this employee.	Found by looking up the Employee's details in EMA files.

Field	Meaning	Notes
DisplayName	Same as StringExtension.	
Telephony Session Fields		
SwitchCallID	The identifier used by the telephone switch and passed in CTI messages.	As the string was received from the switch. May be Avaya UCID, CS1K NetworkCallID, GUID or other. May not be unique.
SwitchID	The Data Source ID for the switch on which the call was received.	As defined in XML files received from EMA.
ANI	The number of the calling party (with respect to this session).	Supplied even on internal and outgoing calls where there is no "real" ANI.
DNIS	The number of the called party (with respect to this session).	Supplied even on internal and outgoing calls where there is no real "DNIS".
Direction	1 (Inbound), 2 (Outbound) or 3 (Internal).	Based on the direction of the call that initiated the session.
Holds	The number of times this party has placed the call on hold.	Includes hold implicit in transfer and conference setups.
HoldDuration	The total duration for which this party has left a call on hold.	Not the same as how long the other party has been kept on hold (for which, see the contact attribute). This timer stops when the party speaks on a consult call even though the original call is still on hold.
Audio Acquisition Fields		
Start	The timestamp of the start of the session.	Same as INumStart below i.e. does not include ringing or call setup period.
End	The timestamp of the end of the last recorded segment in the session.	Does not include e.g. wrap-up time or time on hold at end of session if session abandoned while on hold.

Field	Meaning	Notes
ParentInum (in Module and Channel fields)	The unique reference of the first recorded segment within this session.	Passed via AudioAcquisition Channel (9 least significant decimal digits) and Module (6 most significant decimal digits) parameters. The Module is also the serial number of the recorder making the recording.
Compression	Always "G.729A" unless default compress/mix behavior has been overridden in properties file.	Subsequent INums need not be the same compression format. Can be overridden by property setting "mdl.compression" if necessary.
WrapUpTime	0	Not currently supported.
Type	0	Not currently supported.
Screen Acquisition Fields		
ScreenInum	The unique identifier of the session-level screen recording (if any).	This screen recording persists for the duration of the session, including hold periods. Passed via ScreenAcquisition Exists (9 least significant decimal digits) and Module (6 most significant decimal digits) parameters. The Module is also the serial number of the recorder making the recording.
Additional Recorded Segments		
Inum	The unique reference of a further recorded segment within this session.	
InumStart	The timestamp of the start of this further recorded segment within this session.	
Private Data		
Up to 25 fields	See below for available fields.	

Private Data Fields

Note that User Defined Fields within Avaya Contact Recorder can also be mapped to WFO fields (subject to the overall limit of 25 such fields). These fields provide additional tagging of a session, over and above the standard data fields described above.

Note:

Conditional Custom Data tagging by business rules is not currently supported.

Agent Related Fields

The ExternalID (or "EmployeeID") is always provided as part of the standard session information for a logged on agent that is configured as an Employee in WFO. The fields below may also be configured into one or more Private Data fields. Most are read from the "organization-xxx.xml" files provided by EMA - so rely on

- ACR being able to match the (agent logon id + switch/Data Source ID) against the information given in EMA XML files.
- The information being configured in WFO and received in updated EMA XML cache files.

ID	Name	Type	Description
-919022	AgentID	String 32	The PBX logon id of the agent to whom this session relates. Empty string if no agent logged on to this device.
-919038	NetworkID	String 32	The agent's network logon id derived from EMA data in the same way that other Verint recorders do.
-919003	Skill	String 32	The skill used by the call that originated the session. Optionally, can include the description if available (e.g. from Communication Manager). Set property **`skill.fullname=true`** to enable this. Will then see "***`number (description)`***".
-919036	EmployeeGroup	String 64	Comma separated list of the EmployeeGroups of the agent to whom this session relates. Determined from EMA data in the same way that other Verint recorders do.
-919037	Organization	String 128	Organization of the agent to whom this session relates. Determined from EMA data in the same way that other Verint recorders do.

-919004	SupervisorName	String 128	The name of the Supervisor of the agent to whom this session relates. Determined from EMA data in the same way that other Verint recorders do.
-919001	AgentName	String 128	The name of the agent to whom this session relates. Determined from EMA data in the same way that other Verint recorders do.
-919002	LoggedOnDuration	Int	Not supported.
-919046	ScreenUnit	Int	Not supported.

User Defined Fields

Any User Defined Field stored in Avaya Contact Recorder can be configured as a private data field and sent to WFO. Define the field in EMA using the name of the User Defined Field.

The following fields are available where provided by the appropriate switch and relate to the specific session that has been recorded.

The data **Type** shown in the table below is the default that WFO uses. These fields will, however, be truncated according to the size of the private data field into which they are placed. The types shown are therefore a guide as to which of the 25 private data fields would be suitable for each rather than a strict limit on each. Where the data available in a field exceeds the length of the Private Data Field to which it has been assigned, the contents will be truncated.

The table below shows only those fields that are of value and are not already present in, or directly equivalent to the standard session or contact data fields shown above. Any field not listed is not supported.

ID	Name	Type	Description
-919032	CalledParty	String 32	The address (normally numeric) of the second party on the first call of this session.
-919031	CalledParty Name	String 32	The name of the second party on the first call of this session. Currently only provided on CM, for internal calls.
-919026	CallingParty	String 32	The address (normally numeric) of the second party on the first call of this session.
-919030	CallingParty Name	String 32	The name of the first party on the first call of this session. Currently only on CM, for internal calls.

ID	Name	Type	Description
-919008	Data Source Name	String 64	The name of the Data Source configured in EMA that provided the CTI for this session.
-919010	CallType	String 32	Not supported.
-919016	EventType	String 16	Limited support. One of: "Other","Alerting","Connected", "Held", "Retrieved", "Transferred", "Disconnected". Set at start of session only. Note that if you use this attribute in a business rule to control recording, the value stored when the recording is made may be different from what it was when the rule was triggered. For example, a rule set to trigger when EventType equals "Alerting" will recording incoming calls but by the time the recording starts, the EventType will typically have become "Connected".
-919035	Extended CallHistory	String 32	Comma separated string showing the progress of the session. Consists of one or more of "Ring", "Talk", "Held", "Transfer", "Disconncted", "Conf".
-919053	Fired BusinessRules	String 128	Comma separated list of the names of the Business Rules (if any) that fired on this session. May end in ",…" if the list is longer than 100 characters.
-919051	GlobalCallID	String 64	Globally unique call identifier assigned by the Avaya Contact Recorder.
-919013	NumberDialed	String 32	Number of the most recently explicitly called party. May not be the same as the party that answered the call e.g. if answered via coverage groups; bridged lines etc. This field can also be stored as a User Defined Field in ACR if required. Set the property **`numberdialed.udfname=`*`xxx`* where ***`xxx`*** is the name of the field it will be stored as ("**`numberdialed`**" is recommended for consistency with WFO).
-919040	Parties	String 128	Comma separated list of all parties involved in the contact.

ID	Name	Type	Description
-919007	Queue	String 32	The CDN (CS1K, AACC) or VDN (CM) if any through which the first call of the session was routed. Optionally, can include the description if available (e.g. from Communication Manager). Set property **`queue.fullname=true`** to enable this. Will then see "***`number (description)`***".
-919024	Thirdparty	String 32	The other party involved in a transfer or conference associated with this session.
-919014	Trunk	String 16	The trunk member carrying the call to/from the switch. Only populated for external calls on CM and CS1K. Not available on AACC.
-919015	TrunkGroup	String 16	The trunk group carrying the call to/from the switch. Only populated for external calls on CM and CS1K. Not available on AACC.
-919039	Workstation	String 32	The workstation name associated with the device or agent, as defined in EMA files or in screen capture client agent login data where available.
-919029	LastMessage	String 32	NOT SUPPORTED.
-919028	FirstMessage	String 32	NOT SUPPORTED.

Appendix B: Troubleshooting

This appendix covers two areas: general troubleshooting tips and some specific common issues:

The main sections in this appendix are:

- Hints and Tips on page 339
- Advanced Diagnostics on page 341
- Specific Problems on page 343

Hints and Tips

Where to Look for Clues

When problems occur, check the following:

- **Emailed Alarms and Events**. If you have been using the email settings to have alarms and events forwarded to one or more email addresses, you should check these carefully. As well as checking the contents of messages that you have received, also check for days when the nightly log file purge message has not been received.
- **Alarms Page**. This page within the administration application provides a wealth of information on problems that the system has detected. Review the alarms carefully. If the problem is not immediately apparent, consider viewing all alarms, including those that have previously been cleared. It may be that someone has cleared an alarm without addressing it or realizing its significance.
- **Log Files**. Check for errors being reported in log files within the following directories beneath your installation directory: (`/opt/witness` on Linux and typically `D:\Avaya\ACR` on Windows)
  - `logs`

Determining Current Version

When reporting problems, you should state precisely which version of software you are running.

Avaya Contact Recorder

To determine this, click on the **System > License** tab.

Note the precise version number shown.

Java and Tomcat

For Java and Tomcat please review the folder names at the following locations:

Windows

`ACR Installation Path``\packages (e.g. D:\Avaya\ACR152\packages)`

Linux

`/opt/witness/packages`

PostgreSQL

Windows

Use Control Panel > Programs to see the installed version.

Linux

1. Log in as root
2. Switch to the PostgreSQL user account by typing
   `su - postgres`
3. Access the database by entering
   `psql`
4. Query the version by entering
   `select version();`
5. Quit the database by entering
   `\q`

Advanced Diagnostics

Complex problems, such as memory leaks or thread deadlocks may require stack and/or memory heap dumps.

Installing PSTools (Windows only)

You will need to download and install the PsExec.exe tool to execute jstack.exe.

https://technet.microsoft.com/en-us/sysinternals/psexec

Finding the ACR process ID

You will need to know the process id that the recorder is running as.

Linux

Use the ps ("Process Status") command to find the id of the recorder's process. Filter the results using grep so that only the required process is shown:

```
ps -Af | grep witness.home
```

A typical output is shown below. The first number listed (bold below) is the process id.

```
witness  16190 16188  4 04:48 ?        00:09:25
/opt/witness/java/jdk1.8.0_40/bin/java -
Dwitness.home=../.. -Djava.io.tmpdir=../temp -
Djava.endorsed.dirs=../common/endorsed -
Dcatalina.base=.. -Dcatalina.home=.. -
Dorg.apache.coyote.USE_CUSTOM_STATUS_MSG_IN_HEADER=true
….
```

Windows

Use Task Manager to identify the recorder's PID.

The process will be named `java.exe` and have User name `ACRService`.

Creating thread dumps

First, identify the recorder's process ID as described above.

Linux

Issue the command:

kill -QUIT *nnnnnn*

where ***nnnnnn*** is the process ID.

The stack trace for each thread will be placed in wrapper.log.

Windows

1. Install PSTools as described above.
2. Run a command prompt as administrator.
3. cd to the folder containing pstools (unless included in your path).
4. Enter the command:

   psexec -s *"D:\Avaya\ACR152*\jdk8\bin\jstack.exe" *nnnnn* >dump1.txt

   replacing ***nnnnn*** wth the PID of the recorder's process and
   replacing ***D:\Avaya\ACR152*** with the full install path if different.

   The stack trace will appear in the file specified – dump1.txt in your PSTools folder in this example.

   "error code 0" means "success".

Dumping the Java heap

First, identify the recorder's process ID as described above.

Linux

Issue the command:

/opt/witness/jdk8/bin/jmap -dump:format=b,file=*xxxx* *nnnnn*

where ***nnnnnn*** is the process ID and ***xxxx*** is the name of the file you want the heap dump to be written to.

Windows

1. Install PSTools as described above.
2. Run a command prompt as administrator.
3. cd to the folder containing pstools (unless included in your path)
4. Enter the command

```
psexec -s "D:\Avaya\ACR152\jdk8\bin\jmap.exe" -
dump:format=b,file=xxxxx nnnnn
```

replacing ***nnnnn*** wth the PID of the recorder's process,

replacing ***xxxxx*** with the filename you want the heap written to and replacing ***D:\Avaya\ACR152*** with the full install path if different.

The heap dump will appear in the file specified – heap1,bin in your PSTools folder in this example.

"error code 0" means "success".

Specific Problems

System Administration page problems

You may encounter problems as you access and use the System Administration application. This section lists those problems and suggests steps to take to correct them.

Cannot access the System Administration pages

If you cannot access the System Administration pages, try the following:

- Ping the server to confirm that connectivity is possible. If not, confirm that the server is on and trace the network connections between client and server and double-check the server's IP address, default gateway etc.
- Use the numeric dot notation IP address instead of the hostname. If this works, then the hostname is wrong or cannot be translated by your DNS services. You may need to use a fully qualified node name, such as **recorder.bigco.com**.
- Confirm that the Avaya Contact Recorder service (named "acr" on Linux and "ACR152" on Windows) and the underlying postgreSQL database services are running. Check the recorder's log file for problems.
- Use the browser installed on the server itself to access the application at **http://localhost:8080**.

If this works, then the problem is in the network between server and client. If it does not work, then the problem may be with the Tomcat web server.

Cannot log in

If you have trouble logging in, double-check the state of Caps Lock and ensure the password is being entered with the correct case.

If you can log in under another account, set a new (temporary) password for the account having problems.

If nothing happens when you click the OK button, check that your Internet Explorer settings allow javascript to run.

These symptoms have also been seen when trying to access a server with an underbar in its node name. Note that this is not a valid IP name and should be changed.

Connectivity

Email alarm problems

Invalid entries in any one of the parameters used to define the email settings will result in errors. To check this:

- Try the settings you are using in a standard mail client, such as Outlook. Send a message using the account specified to prove that the settings are valid.
- If email messages have been working and then stop without any of the settings changing, verify that nothing has changed on the mail server. This problem occurs, for example, if your password has been reset or changed on the mail server.
- If the recorder is not sending email messages, it may be because it is not able to access the SMTP server. Check the network connections to the recorder.

CTI Link Connection Problems

The recorder communicates constantly with your switches CTI feed. If this link fails, alarms are raised. If you suddenly get multiple alarms, including those with other components, then the problem is more likely to be with the recorder's network connection than with a specific link.

Search and Replay problems

For most problems with Search and Replay, consider the following diagnostic approaches to narrow down the cause of the problem:

1. Search for a different call, for example, one that is more recent or older, shorter or longer.
2. Log in as a different user with different replay restrictions

Cannot access the replay application

If you cannot get to the login page, try accessing the page from a different machine:

1. From the same side of any firewalls.
2. On the same LAN if you are having problems with WAN access
3. From the same sub-net, if having problems from a different sub-net
4. From the recorder itself, if having problems from the same sub-net.

User replay restrictions do not work

If you have given a user account replay rights over a number of addresses but the calls from these stations are not listed when you enter a valid search that should include them, check the recording ownership. When an agent has logged on to a station that is being recorded, the calls recorded are "owned" by the agent number not the underlying station. Give the user replay rights to the range of agent IDs who have been using the stations in question. In fact, in most cases, you should restrict access to a set of Agent IDs rather than station numbers.

Cannot log in

If you see the login page but cannot get past it:

1. Verify that Caps Lock is off and that you are entering the password with the correct case.
2. Log in as a different user.
3. Confirm the spelling of your log in name with the system administrator and check that your account is still configured in the administration pages.
4. Ask the system administrator to reset your password. Log in with a blank password and change your password when redirected to the Change Password page.

Search returns no calls

If you get to the search page but no calls are returned when you perform a search:

1. Broaden your search criteria to confirm that you can at least find some calls. Start by requesting calls from any parties for today. If that shows no calls, extend the time period.
2. Try setting the date range back to at least the time you know you have seen call records for in the past.
3. Check that the system administrator has given you access to the correct calls. Your search and replay restriction may be wrong or too narrow for the search you are attempting.
4. Confirm that calls are being recorded. Follow the troubleshooting guidelines for recording problems if you suspect that the system is not actually recording or processing any calls.

Calls listed but cannot play them

If you can see the list of calls that matched your search criteria, but cannot actually play them, review

Client Prerequisites on page 209.

No Audio "graph"

This means that the call has not been retrieved from the recorder or archive or has not reached the client PC.

1. Check the server logs for errors.
2. Note the call's 15-digit reference number (shown if you hold the mouse pointer just to the right of the radio button that you click to retrieve the recording. Search for that wav file in the calls path to confirm that the recorded file exists.
3. Check connectivity and available bandwidth to the client PC.

Audio graph stops in mid call

This implies that the transfer of data from the Recorder to your client PC has been stopped or interrupted.

1. Request the same call again. There may have been a temporary network problem.
2. Request a different call. If the problem is only with one call, you may have a corrupt file on your hard disk.
3. Request the problem call from another PC on the same network. If the other PC can retrieve it successfully, assess the differences between the two client PCs; the problem is most likely at the client end.
4. Request the problem call from different sub-nets, ideally working closer to the recorder.
5. Request the call from the recorder server's own browser. If this works and the others don't, then the problem is likely to be in the network between server and clients.

Audio graph appears but no sound

The audio file has reached the client PC successfully; the problem is most likely to be with the PC's multimedia setup or current settings.

1. Verify that the PC has a sound card.
2. Play a wav file through Media Player or similar application to verify that that the sound card is set up correctly.
3. Adjust any hardware volume and/or mute controls on the speakers/headphones.
4. Double-click on the icon in the system tray at the bottom right-hand corner of the screen to verify that the PC's software volume controls are not set to mute or very low.

5. Ensure you are not running any other programs that may be locking the sound card exclusively. If in doubt, shut down all other programs.
6. Try another similar PC. If that works, look for differences in the multimedia setup of the two PCs.

No New Recordings Playable

If you can replay old recordings but not newly made calls, there may be a problem with the recording and/or storage components of the system. Follow these steps:

1. On the **Status > Server** page of the Recorder's System Administration application, look at the counts for total calls recorded and calls recorded today.
2. Use Bulk recording mode to make a test recording.
3. Complete the recording and hang up.
4. Return to the **Status > Server** page and note the **Total media files recorded today (or since restart if today)** and **Total media files recorded to date**. These counts should have increased by at least one, the recording that you just made. If the counts have increased, the recorder is processing recordings. This is probably a search/replay problem. See earlier sections for help.

Note:

If these counts have not increased, the recording has not been successfully stored on the Recorder or inserted into the call details database. Do the following:

5. Look for alarm messages that indicate problems with the recording channel or with file read/write or rename functions. The error message should indicate whether disk space or a directory access problem is the cause. Check that .wav files are appearing in the latest folder beneath the calls path as recordings are made.
6. Check disk space in all partitions. If any of these is 0 or less than 50MB, this may be the problem. Check for build-up of log files. Check that the call details database hasn't exceeded the available space. Consider reducing the number of months of calls kept-use the purge settings on the **General Setup > Server** configuration page to adjust this.
7. Look for alarm messages that indicate licensing problems. The recorder will not process any new calls if you have changed the MAC address, tampered with license settings or are running on a time-expired license. In all cases, you should obtain a new license key.
8. Check that the Avaya Contact Recorder service is running.
9. Check for messages in the log files.
10. Reboot the server and watch for error messages on startup.

Poor Audio Quality on Telephone Replay

If the network configuration is correct and there are no problems with its function, the most likely cause of this problem is that the recorder or the network is overloaded.

To look for a problem on a managed network switch, you should look at the diagnostics and configuration details. You should see ZERO errors on all ports. Any port showing more than 1 packet error in 10,000 is suspect and must be looked at.

If the C-LAN, MedPro and/or recorder ports are failing to auto-sense full/half-duplex properly, you can force each port to either full or half duplex so as to reduce the error count to zero.

Note:

Even though a port may show an error rate of less than 1 in 100 packets, the error counts are deceptive. A single packet error can trigger a full/half duplex negotiation during which all packets are lost in the servers, but none of these show as errors on the switch.

If your error counts are zero on all ports, then we must also consider overload of the recorder as a possible cause. You should monitor the CPU load of the recorder during busy hours. Replay and live monitor are very sensitive to overload. Recording may be unaffected but if the CPU load is too high, audio quality on replay can suffer.

Similar problems have also been seen on multi-CPU AMD Opteron servers. This is caused by an unstable system clock, which is addressed in Red Hat Version 5 update 4.

Recording Problems

Partial recording problems

Since no hardware component in the system is dedicated to specific ports, any hardware problem is likely to affect all recordings equally. Therefore, if some calls are being recorded and are playable but others are not, the problem is probably in the configuration.

1. Recording Mode versus Recording Channel? Determine whether your problems relate to all channels of one or more recording modes or just to certain ports.
2. Check the configuration pages for the affected recording mode(s).
3. Calculate the range of stations carefully. For example, 11000 to 11010 is a range of 11 addresses, not 10.
4. Use the **Status > Ports** page to observe the ports on the recorder during your test calls. The ports should go from idle to active and back again.

Meeting Recording (Communication Manager only)

Cannot Enter Owners

The recorder requires **rtp-payload** signalling as described in Configuring tone detection on page 90 in order to interpret dialled digits. Check that both IP phones and Digital phones are configured for this signalling mode as IP phones will default to it whereas digital phones may not.

Bulk Recording (Communication Manager only)

Cannot Enter Delete or Retain Command

The recorder requires **rtp-payload** signalling as described in Configuring tone detection on page 90 in order to interpret dialled digits. Check that both IP phones and Digital phones are configured for this signalling mode as IP phones will default to it whereas digital phones may not.

Screen Recording Problems

If you use One-X softphones, you may find that the Screen Capture client advises the recorder of logons from the Windows account these run under as well as the normal interactive user. You can block messages from such accounts so long as their names can be distinguished from regular users' accounts using a regular expression.

Set this property `windowsuser.ignore` to the regular expression that will match the account names to be ignored. Note that to match a backslash, it must be double-escaped as these are special characters in both the properties file and in the regular expression. For example **`HQ\\\\svc.*`** will ignore all accounts in the `HQ` domain that begin with `svc`.

Live Monitor Problems

Live monitoring is complex and can be affected by network topologies, NAT and firewalls. Start simple and then get more complex.

1. Concentrate on audio monitoring first and only add screen recording when that is working.
2. Start close to the recorder to eliminate NAT or routing problems.
3. Start with ACR's own Live Monitor client first then WFO's MAS if required.

Audio Monitoring

If ACR's live monitoring is working for audio but via WFO is not, the most likely cause is the firewall on the Windows PC that you are using to monitor audio. WFO live monitoring requires that you open UDP ports 8500-8503 in the firewall so that ACR can send to

these. ACR's own live monitor has the client establish the audio connection via an initial outbound request – avoid the need to open the firewall.

Screen Recording and Monitoring

You can only record and live monitor screens from which ACR is receiving Windows account logon/logoff events. Check that the screen being monitored has been configured with the IP address of the Master (and most appropriate Standby if present).

Appendix C: Alarms

This appendix provides details of the alarms that can be raised by the system.

The main sections in this appendix are:

- Alarms on page 353
- Alarms Table on page 353

Alarms

The recorder may generate the following Alarm or Event notification messages. These events are:

- shown on the **Alarms > View Alarms** page
- sent in email messages as specified on the **System > Email Server** page
- reported via SNMP
- logged to the recorder's log file `acr.log`

Alarms Table

Within messages, the strings XXX and YYY represent a specific parameter such as a station number, an IP address etc. The table shows both the English text of the message (that appears on the Alarms page) and the underlying resource string that appears in the log file.

Entries within the table are sorted according to the log file entry.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Major	alarms.aoc	Avaya Oceana™ Interface at XXX failed: YYY	XXX is IP address of Avaya Oceana™ feed. YYY is detail of error.	Depends on YYY
Major	alarms.applyingconfig	Error applying configuration from Master recorder: XXXX.		Note recent admin changes and submit details with log files from this server and the Master.
Major	alarms.archive.badlayout	Archive destination XXX cannot open layout YYY.	Layout name is set on Archive configuration.	Check this layout exists.
Major	alarms.archive.drive	Error accessing archive drive XXX. YYY.		Check the drive exists, is working and is not being used by any other application.
Minor	alarms.archive.filetoolarge	Error writing XXX to I/O job ID YYY. File too large. Take copy manually if required.	Tar file cannot hold individual files larger than 2GB.	Determine root cause of large files.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Major	alarms.archive.filter	Archive XXX filter incorrect. Can occur on upgrade. Check release notes and update layout details. Error: YYY.	Filter mechanism for archive changed in 12.0.	Bring this archive's Advanced settings in line with new Layout design.
Minor	alarms.archive.filterbad	One or mare archive filters are invalid.		Check this Archive's advanced filter settings.
Minor	alarms.archive.filtererr	Error determining if recording(s) XXX should be archived/exported: YYY.	Probably SQL problems. Check log file for full detail.	Use search and replay to look at these calls and apply same filter and layout as this IO jobarchive/export job has been told to use.
Warning	alarms.archive.notreinserted	Newly initialized disk XXX has not been reinserted on YYY.	De-iced disk should automatically reinsert.	Push disk in manually. Replace drive with motorized one.
Warning	alarms.archive.old	Call(s) XXX hours old have not been archived (oldest is YYY).	Recorder checks hourly for calls that should have been archived but haven't.	Check the archive jobs are running and not backlogged.
Major	alarms.beaconclash	Beacon phone XXX configured on multiple servers: YYY	YYY gives the serial numbers of the two servers that have been configured with the same beacon phone XXX	Change one or the other server's beacon phone configuration.
Major	alarms.beaconisdmcc	Beacon phone XXX has already been allocated to recording. Restart required.	Set up beacons before system goes live.	Choose another DMCC port as beacon or restart the system.
Minor or Major	alarms.bgnd.perXXXXX where XXXXX is second, minute, hour, day, morning or night	Error running XXX tasks: YYY ZZZ	Scheduled task either hit errors (Major) or was deferred as previous one had not completed (Minor)	Depends on frequency and nature of errors reported in YYY, ZZZ.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Warning or Major	alarms.blocked	CPU Blocked for XXXms (threshold YYY)	Warning if blockage less than twice the threshold. Otherwise Major.	Investigate cause of blockages.
Major	alarms.closeblocked	Disk I/O blocked for XXXms closing file YYY. Causes include VM "thin" disk provisioning, mirroring, snapshotting and anti-virus.	Recorder requires real-time file system. Cannot tolerate long blockages.	Review disk configuration and other processes accessing recorder's files. See technical note ACR on Windows (linux less susceptible).
Major	alarms.cmapi.up	Device, Media AND Call Control API running on XXX.	DMCC services are restored.	No action required.
Major	alarms.core.nodatasource	No Data Source configured for XXX.	WFO must be configured with Data Sources that match the name of each switch, dialer and (if present) AACC.	Bring WFO Data Source names in line with ACR.
Minor	alarms.crs.upload	Failed to upload details of recording XXX to central database. Response code YYY.	YYY is standard HTTP response code.	Look up HTTP response code.
Major	alarms.cti.badinitialsend	IO Exception on initial transmission.	Could not transmit over the socket connection.	Check far end is listening. Check network connection is up.
Major	alarms.cti.cannotconnect	Cannot connect.	Could not establish a TCP/IP link.	Check address.
Major	alarms.cti.clientrelease	Link dropped by far end.		Check far end for problems or manual shutdown.
Warning	alarms.cti.configchanged	CTI Configuration changed. Dropping previous connection.	System is responding to change in configuration.	Nothing.
Major	alarms.cti.connfailed	Connection failed.	An established link has been dropped.	Check for neighboring events at each end of the connection.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Major	alarms.cti.error	Link to CTI on XXX reporting error: YYY.	A problem occurred between the recorder and the Avaya Contact Center server.	Depends on error YYY reported.
Major	alarms.cti.heartbeatfailed	CTI link failed. Heartbeat lost.	No activity on the link for an unacceptable period.	Check network connectivity. Check health of component at far end.
Minor	alarms.cti.invalidcallid	Invalid call id: XXX.		Check CTI link configuration. Check switch patch levels.
Major	alarms.cti.invalidparams	Invalid CTI Address: XXX.	You have not specified an address for the Avaya Contact Center server. (CS 1000).	Enter a valid address on the **General Setup** > Data Source page.
Major	alarms.cti.ioexception	IO Exception.	A socket error occurred on the link.	Check network integrity.
Info	alarms.cti.providerok	CTI Services on XXX UP.	Connection to CTI link server has been established.	Check earlier error message to determine why link failed earlier.
Major	alarms.cti.shutdown	Service shut down.	The link has been shut down.	If not planned, check for reasons preceding this in the event log.
Warning	alarms.cti.slowresponse	Slow response from CTI link. CN XXX reg/unreg ltook YYYms.	Registration or unregistration of an address took longer than expected.	Check the recorder is not overloaded. Check the MLS link is not overloaded.
Major	alarms.cti.timeout	Connection timed out. Far end not responding.	The component specified has not responded; has not connected or reconnected within the expected time.	Check the network connectivity and the state of the component specified.
Major	alarms.dapi.aesvcsconfig	Cannot use CallInformationServices. Please check AES.	See AE Server manuals.	

Severity	Log File Entry	English Alarm text	Comment	What to Do
Major	alarms.database.cannotinsert	Cannot insert details of recording XXX into database.		Check disk space on the partition holding the call details database.
Major or Minor	alarms.database.migration	Error migrating database: XXX YYY.	Severity varies according to nature of problem.	Retain backup of database made before upgrade. Take backup of database now. Contact Avaya.
Major	alarms.dialer.key	Dialer(s) configured but not licensed.		Obtain and enter valid license key.
Major	alarms.dialer.linkdown	Dialer XXX link DOWN. Reason: YYY.	XXX is name of dialer.	Check dialer end and network.
Major	alarms.dialer.linkup	Dialer XXX link UP.	XXX is name of dialer.	Problem resolved.
Major	alarms.dialer.noclassname	Class name not defined.		Correct dialer class name in properties file.
Major	alarms.disk.full	Disk full on partition 'XXX.	'Check log files. If calls partition, check for files that cannot be purged.	Delete some files to make space.
Major	alarms.disk.notarchived	Recordings from XXX to YYY were purged from disk buffer but had not all been archived.		Provide more disk space. Check what is using up space.
Major	alarms.disk.purging	Disk buffer nearly full. Purging recordings regardless of required retention period.		Provide more disk space. Check what is using up space.
Warning	alarms.disknearlyfull	Disk nearly full on partition 'XXX'. Only YYY MB free.	Check log files. If calls partition, check for files that cannot be purged.	Delete some files to make space.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Minor	alarms.dn.notobservable	Cannot observe DN/Position ID XXX. CTI Cause Code YYY.	The recorder was not able to register for CTI events on the address shown.	CS 1000: Check the AST flag is set on the DN/Position ID. Check that the Meridian1 Machine name and number are correctly set. Check the MLS specification for specific cause codes.
Minor	alarms.dn.regdropped	Observation of DN/Position ID XXX dropped.	The recorder is no longer receiving CTI events for the position ID shown.	CS 1000: Check that the AST flag has not been removed. Check the MLS specification for meaning of each cause code.
Minor	alarms.dn.startrec	Recording Start failed on DN/TN XXX.Cause Code (hex):YYY.	Could not start Duplicate Media Streaming.	Check the specific cause code against MLS specifications.
Warning	alarms.dn.starttimeout	Timeout on recording start on DN XXX TN YYY.	Did not receive prompt confirmation of Duplicate Media Streaming stopping.	Check recorder is not overloaded. Check CTI is not overloaded. Check network connectivity.
Minor	alarms.dn.stopnotif	Recording stopped on DN/TN XXX. Reason Code (hex): YYY.	Duplicate Media Streaming stopped unexpectedly.	Check the specific cause code against the MLS specification.
Warning	alarms.dn.stoptimeout	Timeout on recording stop on DN XXX TN YYY.		Check recorder is not overloaded. Check MLS is not overloaded. Check network connectivity.
Major	alarms.drs.misconfigured	Recorder XXX at YYY is misconfigured. Do NOT include Distributed Replay Servers in list of Replay Servers.	XXX is serial number and YYY the IP address of a Distributed Replay Server.	Remove that server from the list of Replay Servers on the **General Setup > Server** page of the recorder raising the alarm.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Major	alarms.emc.notavailable	EMC Archive XXX not usable. Error: YYY.		Check configuration if it has never worked. Otherwise, check EMC Centera file store is available.
Minor	alarms.exclusion	Cannot record calls due to exclusion. Call ID XXX	Exclusion features are stopping the recorder from single-step conferencing into this call to record.	Determine parties on affected calls and review all switch exclusion settings.
Minor	alarms.file.compressfailed	Failed to compress audio file for recording XXX. Reason YYY.		Depends on reason.
Minor	alarms.file.deletefailed	Failed to delete file XXX. Reason: YYY.		Depends on reason.
Warning	alarms.file.notindb	XXX recordings up to and including YYY were not stored to database on last shutdown.	Implies last shutdown was not "clean".	Shut down service rather than kill.
Minor	alarms.file.wavwritefail	Failed to write file XXX.wav. I/O Error: YYY.		Depends on reason.
Major	alarms.file.xmlwritefail	Failed to write XML file XXX. Reason: YYY.		Depends on reason.
Major	alarms.filedetails	System is configured to retain recorded files for *DDD* days but only retaining call details for *XXX* days. Correct on I/O Jobs page.	Can occur on upgrading from 15.1FP1 or earlier. Relationship between these settings is now enforced automatically,	Increase days to retain call details or reduce days to retain recorded files on the **Operations > I/O Jobs** page.
Minor	Alarms.gdpr.cantsend	Unable to send Delete job #XXX to YYY. Ensure server is up and resubmit the delete request.		Ensure the server is up and resubmit the delete request.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Warn	alarms.hapair	AES at *XXX* is configured as part of a high availability pair but connection to YYY has never connected. Please confirm connection out of hours under failover conditions.	Will repeat until a connection has been established at least once on this address.	Perform testing on this AES failover IP address.
Warn	alarms.iojob.badfilter	Review the archive XXX as the filter can no longer locate new calls to archive. This causes needless processing.		Check listed archive job and disable or update as appropriate.
Warn	alarms.iojob.fromdate	Leaving a *from* date on an archive is bad practice. It forces the recorder to review ALL recording between then and now every night. As the historic check has been done already, you should remove the *from* date. XXX archive(s) with *from* dates active.		Check archive jobs and remove any redundant *From Date* criteria.
Warn	alarms.iojob.invalidpermissions	The user XXX does not have the required read, write and execute permissions for archive destination YYY. This causes archiving to this location to fail.		Check archive permissions on the network share.
Major	alarms.jobthread.stalled	Job thread XXX stalled. YYY.	XXX is the name of the job thread, YYY the cause of the problem.	Often a neighboring alarm giving more detail.
Major	alarms.license.channelcnt	Licensed bulk recording capacity reached. Cannot record call on XX.	Cannot record a call on the address shown. The recorder is already operating at its maximum licensed capacity.	Purchase and/or install additional capacity.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Minor	alarms.license.quality	Quality Monitoring agent station license exceeded. Cannot record XXX.	You have attempted to record more agent stations than your license permits.	Increase licensed capacity.
Minor	alarms.lockadvisor	Failed to advise server of XXX. Response code YYY.	Lock or unlock not passed to other server.	Check other server is available and network path is good.
Major	alarms.lockfailed	Failed to lock recording XXX. Reason: YYY.	A call has not been locked as requested.	Depends on reason.
Warning	alarms.lockinconsistency	Inconsistency in locked call information affecting XXX recordings.	May indicate attempts to tamper with recordings.	Report security concerns.
Warning	alarms.lockpurgeneeded	Lock directory contains XXX unlocked recording files. Use maintenance page to purge.	Nightly check of lock folder detects that calls are now unlocked and can be purged.	Purge as instructed.
Info	alarms.logpurge	Purged XXXKB of old log files. Now YYYMB free.	Nightly message.	No action required if received. If not received, check recorder is running.
Major	alarms.nobeacon	No Beacon phone configured for Communication Manager XXX	Your system topology requires a beacon softphone for each server. This server does not have one configured.	Configure a beacon softphone.
Major	alarms.node.overflow	Call failed to record XXX due to overload on recorder YYY.		Rebalance load or add capacity.
Info	alarms.node.portsavailable	XXX or more ports available on recorder YYY.		Problem resolved.
Major	alarms.node.portslow	Less than XXX ports available on recorder YYY.		Rebalance load or add capacity.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Minor	alarms.privviolation	Cannot record call XXX due to CoR violation. Devices on call: YYY	A Class of Restriction (CoR) has been breached and the call cannot be recorded. May really be a CoR issue or can be due to incorrect login.	Check called/calling party restrictions in that station/agent's COR. Check login details and try again.
Minor	alarms.props.obsolete	Properties file contains a setting for XXX. This is no longer set via the properties file. The value found when upgrading has been migrated to the administration settings. This item should now be removed from the properties file to avoid later confusion.		Remove this property.
Major	alarms.queue.acceptable	XXX Job Queue backlog reduced to acceptable level. Currently YYYms.		No action required.
Major	alarms.queue.copy	Failed to copy to XXX. Reason:YYY. NOTE: Further errors with the same root cause on the same path will only show here once every 24 hours. Check the log file if in doubt.	Will cache files until problem is resolved then copy all outstanding.	Determine why recorder cannot write to the share.
Warning	alarms.queue.slowjob	XXX Job Queue individual slow job - took YYYms.	This activity took longer than expected.	May occur under startup conditions. If occurring later, check recorder is not overloaded.
Warning	alarms.queue.tooslow	XXX Job Queue backlogged. Delay currently YYYms but may go higher.	The recorder cannot process these jobs as quickly as it should.	May occur under startup conditions. If occurring later, check recorder is not overloaded.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Minor	alarms.recfailed	Recording XXX failed. Reason: YYY.		Depends on reason.
Major	alarms.remoteids.serversock et	Error on server socket port XXX. YYY.	YYY is another Avaya Contact Recorder.	Check alarms on server YYY.
Major	alarms.replayapi.redirect	HTTP not allowed. XXX should be configured to use HTTPS.		Use the https url instead.
Major	alarms.restartneeded	RESTART required for recent configuration changes to take effect. Check other servers as they may also require a restart.		Restart server(s) as instructed.
Major	alarms.retainportclash	The Retain port, XXX is assigned to another recording mode.	Property setting execmode.retainnu mber specifies a port that you have assigned to On Demand or Meeting recording.	Change the property setting or remove this number from On Demand or Meeting recording.
Major	alarms.retention	I/O Job # *NN* has a retention period of *DDD* days. This is above or dangerously close to the overall call details retention period set on the I/O Jobs page.	Can occur if upgraded from 15.1FP1 or earlier. These settings now enforced automatically.	Reduce the retention period for the archive or increase the retention period for call details on the **Operations > I/O Jobs** page.
Major	alarms.rtp.misc	Error on RTP: XXX Parameters: YYY.		Check network.
Minor	alarms.rtp.noaudio	No audio in recording XXX of call YYY	An empty audio file has been recorded despite the call being more than 5s long.	Investigate using the INum shown in the alarm. Also shows (in YYY) call ID, duration of call and parties on the call.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Minor	alarms.rtp.nopackets	No audio packets received on call from XXX. Call failed or no network path to recorder.	See Avaya support site re PSN #345U - TN570C and TN570D Expansion Interface boards log chronic fiber alarms.	Check TN570 boards are correct vintage. (Communication Manager). Check network paths from phones to recorder (CS1000).
Warning	alarms.rtp.packetloss	Unacceptable packet loss from XXXX.	Only raised once per day per IP address.	Check network integrity between recorder and this IP address.
Major	alarms.sfi.certmissing	The Key Manager client side certificate file is missing.		Supply the file.
Major	alarms.sfi.initfailed	Key Manager initialization failed.Reason: XXX.		Depends on reason.
Major	alarms.sfi.kmsdown	Key Manager unreachable. Reason: XXX.		Check KMS server and network.
Major	alarms.sfi.needboth	Both Key Manager parameters must be specified.		Enter the missing parameter.
Major	alarms.sfi.needunlimited	Key Manager parameters are specified, but the Unlimited Strength Policy files are not installed.		Install unlimited strength policy files.
Major	alarms.sms.corserror	Error refreshing CoR List via AE Server XXX: YYY	Bulk recording by CoR is enabled but ACR cannot determine CoR members via the SMS Web Service on the AES XXX.	Check AES configuration and prove SMS Web Services is configured and functioning.
Major	alarms.softphone.beepstart	Port XXX. Failed to start beeptone. Reason: YYY.	Internal error.	Report problem.
Major	alarms.softphone.beepstop	Port XXX. Failed to stop beeptone. Reason: YYY.	Internal error.	Report problem.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Major	alarms.softphone.buttonlookup	Port XXX. Error looking up button Id. Reason: YYY.	Softphone misconfigured.	Check softphone configuration is as per installation section.
Major	alarms.softphone.buttonmissing	Port XXX. Unusable as button YYY is missing from softphone.	Softphone misconfigured.	Check softphone configuration is as per installation section.
Major	alarms.softphone.callstopped	Port XXX. Call dropped as it exceeded maximum permitted duration.	Recording port has been active for several hours.	Consider setting recording mode to release call on dropping to one other party.
Major	alarms.softphone.endrecfailed	Port XXX. Error ending recording. Reason: YYY.	Internal error.	Report Problem.
Major	alarms.softphone.hookswitch	Port XXX. Error setting hook switch. Reason: YYY.	Internal error.	Report Problem.
Major	alarms.softphone.inservice	Port XXX restored.		No action required.
Major	alarms.softphone.nullpointer	Port XXX. DMCC event 'YYY' fired with null pointer.	Internal DMCC error related to a particular recorder port.	Report occurrences with a copy of your log files.
Minor	alarms.softphone.outofservice	Port XXX out of service.	May recover within a few seconds.	Click **Reset** link on **Port Status** page.
Major	alarms.softphone.ownerreg	Port XXX. Error unregistering owner. Reason: YYY.	Internal error.	Report Problem.
Major	alarms.softphone.play	Port XXX. Error playing file. Reason: YYY.		Depends on reason shown.
Minor	alarms.softphone.processfailed	Error processing file XXX. Reason YYY.		Depends on reason shown.
Major	alarms.softphone.recordstartfailed	Port XXX. Failed to enable recording. Reason: YYY.		Depends on reason shown.
Major	alarms.softphone.recordtimeout	Port XXX. Recording failed to start. Resetting port.		Depends on reason shown.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Major	alarms.softphone.registrationfailed	Port XXX Registration failed. Reason: YYY.		Depends on reason shown.
Major	alarms.softphone.regtimeout	Port XXX. Registration timed out.		Check connectivity to AE Server.
Major	alarms.softphone.rtpthread	Port XXX. RTP handler failed. Reason: YYY.		Depends on reason shown.
Major	alarms.softphone.setuprecording	Port XXX. Error setting up recording. Reason YYY.		Depends on reason shown.
Warning	alarms.softphone.shortpacket	Port XXX. VoIP packet interval of YYYms less than recommended 60ms.	Default settings of 20 or 30ms are not efficient for recording.	Configure codec set and network region as per installation instructions.
Major	alarms.softphone.sscdropped	Port XXX. Single-step conference dropped unexpectedly while in state YYY.		Report problem if more than very occasional occurrences.
Major	alarms.softphone.sscfailed	Port XXX. Single-step conference failed despite multiple retries.		Consider increasing number of retries.
Major	alarms.softphone.ssctimeout	Port XXX. Single-step conference setup timed out. Trying to conference in to YYY.		Check system loading.
Major	alarms.softphone.startrecfailed	Port XXX. Recording failed to start. Reason: YYY.		Depends on reason shown.
Major	alarms.softphone.stopplaying	Port XXX. Error stopping file playing. Reason YYY.		Depends on reason shown.
Major	alarms.softphone.stoprecordfailed	Port XXX. Error disabling recording. Reason YYY.		Depends on reason shown.
Major	alarms.softphone.timedout	Port XXX. Timed out when in state YYY.		Report if more than very occasional occurrences.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Major	alarms.softphone.ttd	Port XXX. Failed to set up touch tone detect. Check you have OUT_BAND signalling configured. Failure Reason: YYY.		Configure signaling in accordance with installation instructions.
Warning	alarms.softphone.userreset	Port XXX. User 'YYY' reset the port.	Also logged to audit trail.	No action required.
Major	alarms.standby.invalidstns	Invalid Station list "YYY" received from master for station pool "XXX".	Internal error.	Report problem.
Major	alarms.standby.noswitch	No link to Communication Manager. All ports have failed.		Check AE server is up.
Major	alarms.standby.notlicensed	A Standby recorder is attempting to connect but this server is not licensed to support backup channels. Apply new license and restart server.		Enter a license allowing backup channels.
Major	alarms.standby.notviable	Recorder fatal error: XXXX.	The recorder has detected a major problem that means it cannot record.	Check the alarms preceding this one for the root cause of the problem
Major	alarms.standby.primarynotok	Master recorder requests that Standby(s) take over.	The master Avaya Contact Recorder has decided that it cannot record and requests that the Standby recorder should take control.	Check the alarms preceding this one for the root cause of the problem. This may be that a disk partition is full or the recorder cannot communicate with the CTI link.
Major	alarms.switchchanged	Switch type changed to XXX. Restart required.	Master configuration has changed and been copied to Standby.	Restart server

Severity	Log File Entry	English Alarm text	Comment	What to Do
Major	alarms.syslog	Cannot create syslog appender to host XXX. Error: YYY.		Check configuration and/or route to host according to error shown.
Major	alarms.timesynch	Clock synchronization error. XXX is YYYms adrift from this server.	Another ACR server's clock does not agree with this one's - by more than 5 seconds or messages have been delayed.	You MUST time synch all ACR servers (and any others they interact with) to real time via NTP.
Major	alarms.tsapi.backlogcleared	Single-step Conference Queue Backlogged. Flushing.	Queue of requests is too long.	Check system loading and increase capacity if needed.
Major	alarms.tsapi.backlogged	Single-step Conference Backlog Cleared.	Queue of requests now acceptable length.	No action required.
Warning	alarms.tsapi.conffailed	Single-step Conference Failed: XXX.	Very occasional failures are nothing to be concerned about. Regular failures should be reported.	Check the error code reported against those shown in TSAPI reference manuals.Check for latest versions of AES TSAPI.
Major	alarms.tsapi.configchanged	AES TSAPI Configuration changed. Dropping previous connection.	User changed configuration details.	No action required.
Minor	alarms.tsapi.csta44	Single-step Conference failed on address XXX. CSTA Error 44. Check that no other AE Server is controlling this call.	Always means that another application is trying to control the call.	Stop using Follow the call options or stop the other application.
Major	alarms.tsapi.error	AES TSAPI Services on XXX reporting error: YYY.		Depends on error shown.
Major	alarms.tsapi.heartbeatfailed	AES TSAPI Services failed. Heartbeat lost.	System should attempt to restart TSAPI services automatically.	Report if recurring.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Major	alarms.tsapi.invalidparams	Invalid Service, username or password parameters for AES TSAPI: XXX.		Check all parameters.
Major	alarms.tsapi.invalidskill	Invalid Skill Group: XXX.		Check that this is a valid skill hunt group.
Major	alarms.tsapi.invalidvdn	Invalid VDN: XXX.		Check that this is a valid VDN.
Major	alarms.tsapi.notaskill	Not a valid skill group (or AES TSAPI cannot reach switch).		Check that this is a valid skill hunt group.
Major	alarms.tsapi.observationended	Observation ended.		Report problem.
Major	alarms.tsapi.observer	AES TSAPI Observation of XXX reports error: YYY.		Report problem.
Major	alarms.tsapi.outofservice	Out of Service.		Check AES TSAPI services.
Major	alarms.tsapi.overflow	Conferenced Recording Mode failed to record a call because all ports were busy.		Increase number of concurrent recordings allowed on Conferenced mode.
Major	alarms.tsapi.portsavailable	Conferenced Recording Mode now has XXX ports available.		No action required.
Warning	alarms.tsapi.portslow	Conferenced Recording Mode has less than XXX ports available.		Increase number of concurrent recordings allowed on Conferenced mode.
Major	alarms.tsapi.providerok	AES TSAPI Services on XXX UP.		No action required.
Major	alarms.tsapi.shutdown	Service shut down.		Check AES TSAPI services.
Minor	alarms.unify.ioexception	File I/O error updating call details of XXX. Reason: YYY.		Depends on reason.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Minor	alarms.unify.parseexception	Parsing error updating call details of XXX. Reason: YYY.		Depends on reason.
Minor	alarms.unify.parserror	Error parsing Unify/External control information. XXX is not a number.		Correct external controller command.
Major	alarms.url.badport	Invalid IP port number for recording mode: XXX, Set to: YYY.		No action required.
Major	alarms.url.general	Error on link to 'XXX': YYY.		Check connectivity.
Major	alarms.url.linkup	Link established with XXX server on 'YYY'.		No action required.
Major	alarms.url.nourl	Ports are allocated to XXX recording but no url is set to communicate with the server.		Configure appropriate link's URL.
Major	alarms.url.socket	Error connecting to 'XXX'. YYY.	Problem connecting to other server.	Depends on reason shown. Typically network routing or security issues.
Major	alarms.url.unknownhost	Unknown host 'XXX' in url for YYY server.		"Check network, address is correct; add to DNS or use numeric IP address."
Major	err.alarmtag.exception	Error displaying alarm tag.	Internal error.	Report problem.
Major	err.alphas.invalidchar	Invalid character ('-') in a non-numeric string.		Correct the entry.
Major	err.alphas.toolong	Non-numeric list is too long. 10000 character maximum including commas.		Reduce the length of the entry.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Major	err.confighistory.recordsdeleted	One or more configuration history records have been deleted.		Do not trust config history or audit records.
Major	err.confighistory.tampered	One or more configuration history records have been altered.		Do not trust config history or audit records.
Major	err.database.purge	Error purging database. Reason: XXX.		Depends on reason shown.
Major	err.license.clocksetback	The system clock has been set back. Timed license is invalid.		Obtain and enter a non-timed license.
Major	err.license.create	Error creating license token object.	Internal error.	Report problem.
Major	err.mail.authentication	Authentication failed attempting to send e-mail.		Check email user and password entries.
Minor	err.mail.invalidaddress	Invalid email address: XXX.		Correct the address.
Minor	err.mail.send	Error sending e-mail.		Check all email account entries. Send a test email manually to verify settings.
Major	err.maxusagetag.exception	Error displaying peak usage details.	Internal error.	Report problem.
Major	err.mls.notconfigured	CTI Link not yet configured.	You have not yet provided the IP address of the Avaya Contact Center server.	Specify the address of the Avaya Contact Center server.
Major	err.portpool.nocurrent	No port pool specified.	Internal error.	Report problem.
Major	err.portpooltotaltag.exception	Error displaying port pool totals.	Internal error.	Report problem.
Major	err.settingstag.exception	Error in IterateSettingsTag.	Internal error.	Report problem.
Major	err.settingstag.invalid	Invalid settings group requested in IterateSettingsTag.	Internal error.	Report problem.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Major	err.settingtag.exception	Error in SettingTag.	Internal error.	Report problem.
Major	err.settingtag.invalidfield	Unrecognized field request in SettingTag.	Internal error.	Report problem.
Major	err.softphonestatetag.exception	Error showing port state details.	Internal error.	Report problem.
Major	err.stnpooltag.exception	Error in StnPoolTag.	Internal error.	Report problem.
Major	err.stnpooltag.invalid	Invalid station pool requested in StnPoolTag.	Internal error.	Report problem.
Major	err.stnrangetag.exception	Error displaying station range details.	Internal error.	Report problem.
Major	err.stnrangetag.invalidfield	Unrecognized field name in StnRangeTag.	Internal error	Report problem.
Warning	err.system.restart	System restarting. Running version XXX		No action required.
Info	err.system.shutdown	System shut down.		No action required.
Info	info.aoc	Avaya Oceana™ Interface at XXX connected.		No action required.
Info	info.archive.errorcleared	Wrote to archive disk correctly.		No action required.
Info	info.archive.rightdisk	Correct archive disk now inserted.	Now able to write to disk in drive.	No action required.
Info	info.core.datasource	Data Source XXX configured.	Configuration error has been corrected.	Nothing.
Info	info.dn.observed	CTI Monitor now established on XXX.		Problem resolved.
Info	info.emc.available	EMC Archive XXX available.	Previous error cleared.	Nothing.
Info	Info.goingstandby	This server entering standby mode.	Implies it has been in control.	Check this is appropriate given the state of other servers.
Info	info.jobthread.unstalled	Job thread XXX running OK.		Problem resolved.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Info	info.mail.sent	Mail sent successfully.		No action required.
Info	info.queue.copy	File copied successfully to XXX.		No action required.
Info	info.sfi.kmsup	Key Manager connection restored.		Problem resolved.
Info	info.standby.goingstandby	Standby recorder returning to idle.		No action required.
Info	info.standby.primaryok	Master recorder requests Standby(s) go idle.		No action required.
Info	info.standby.viable	Recorder fatal error resolved XXXX.		No action required unless another problem is highlighted.
Info	Info.takingcontrol	This server is taking control or recordings.		Check this is appropriate given the state of other servers.
Info	info.timesynch	Clock synchronization corrected. XXX is now YYYms adrift from this server.	Drift corrected.	Nothing.
Major	link.master.linkerr	Error on link to Slave recorder at XXX: YYY.		Check connectivity, check other node is up.
Major	link.primary.linkerr	Error on link to Standby recorder at XXX: YYY.		Check connectivity, check other node is up.
Major	link.reason.remotetimeout	Connection with XXX recorder YYY timed out.	XXX is server type. YYY is serial number.	Check server and network.
Major	standby.reason.connecttimeout	Initial connection timed out with XXX.		Check connectivity, check other node is up.
Major	standby.reason.inactivity	Heartbeat failure via XXX.		Check connectivity, check other node is up.

Severity	Log File Entry	English Alarm text	Comment	What to Do
Major	standby.reason.noactivity	Heartbeat failure.		Check connectivity, check other node is up.
Major	standby.reason.primaryreque st	Instruction from Master Recorder.		No action required.
Major	standby.reason.reconnecttimeout	Connection dropped and reconnection timed out with XXX.		Check connectivity, check other node is up.
Major	standby.reason.timeout	Initial connection timed out.		Check connectivity, check other node is up.
Info	stnrange.usage.nowarn	The pool of ports XXX has YYY or more port(s) available.		No action required.
Minor	stnrange.usage.warn	The pool of ports XXX has LESS THAN YYY port(s) available.		Consider allocating more ports to this mode.

Appendix D: High Availability (“HA”)

This Appendix assumes that the reader is already familiar with:

- High availability server options (redundant fans, PSUs, RAID storage).
- High availability network options (bonded/teamed NICs, salt-and-peppered network switches, redundant routing etc.).
- Design, installation and operation of Master and Slave recording systems as described in the body of this manual.
- The content of *Avaya Aura Communication Manager Survivability Options.*

The simplest and most robust – but also most expensive and potentially limiting – approach is to have two fully *independent* systems, neither of which need therefore be fault tolerant in itself.

If a *single* recording system is being deployed, the next most important thing is to carefully consider the general fault tolerance of each server in the system and the network between them.

However, there are specific options available to reduce the impact of a range of failure scenarios, leading to higher availability recording, storage and/or retrieval of recordings. Once all these options are understood, an appropriate topology can be designed. The procedure to do so is quite complex on a large Communication Manager system. The main sections in this appendix are therefore:

- Fully Parallel Systems on page 376
- High Availability Recording on page 378
- High Availability Storage on page 382
- High Availability Retrieval on page 384
- High Availability CM Topology on page 385
- Standby recorders and Unify/External Control on page 394
- Mode of operation on page 395

Fully Parallel Systems

There are a number of ways this can be achieved using two independent recording systems.

Communication Manager

Two systems, both using DMCC recording on a Communication Manager can be deployed so long as:

a) The limits on TSAPI observers are not exceeded.
b) The drop in maximum number of genuine parties on the call (now 4 plus 2 recorder ports) is acceptable.
c) Load on the AES and Communication Manager are acceptable – softphones and VoIP resources are doubled.
d) Each system is configured to ignore the recording ports of the other (see property setting **`conferenced.ignore`**.).
e) As with other HA options, the two systems are recording a single Communication Manager.
f) If records from both are uploaded to one or more shared CRSs then:
    - Ensure the database partition is sized for double the actual number of calls as each will be stored twice (in “sunny-day” running).
    - Ensure that no two recorders in your system share the same last 2 digits of their serial numbers (to allow ContactIDs to be matched unambiguously to servers).
    - Layouts must be added for “main” and “secondary” systems:
        i. Segment-based: selecting only segments whose INums are from the appropriate servers (INum/1000000000 == one of the serial numbers in required system).
        ii. Session-based: selecting only contacts whose ContactIDs are from the appropriate servers (ContactID/100 % 100) == (one of the serial numbers in required system % 100).
    - If the CRS contains recordings from 15.1FP1 or earlier, the CRS must remain configured with property setting **`vx.servers.main=`*`nnnnnn,mmmmm`* where ***`nnnnnn,mmmmmm`*** is a comma separated list of the recorder serial numbers in the “main” system – which provides WFO with ContactIDs (the other being “secondary”).

If (b) or (c) are problematic, one or both of the systems can be configured to use TDM taps or Passive IP taps.

WARNING:

The overall maximum of 6 parties on a conference might become a constraint if two of these are recording ports. Call scenarios in which two calls merge (transfer and conferencing) - and where each of these is already a 4 party call due to duplicated recording can exceed these limits and are not

supported. To avoid this, recording ports are removed from a call as soon as it goes on hold. In this way there are less parties present on the call when it later merges with a consultation call. (Note that this cannot be done with the "Follow the call" recording option as otherwise the call may be lost from the recorder's sight.)

CS1000

Each phone can only stream audio to a single recorder. The only way to achieve parallel recording systems is therefore if at least one of the systems uses TDM taps.

Central Replay Server for Parallel Systems

If both systems populate a shared CRS, you will get duplicated results in searches on Segment layouts and echo when playing calls in Session Layouts. To avoid this:

Segment-based Layouts

Create two similar segment layouts – each of which filters INums to those from one set of recorders only. Use the one that selects INums from the “main” server(s) normally and the other if that returns no results.

The first 6 digits of the 15 digit INum match the serial number of the recorder that made the call.

Session-based Layouts

Set property **`vx.servers.main=`*`aaaaaa,bbbbbb`*** where ***`aaaaaa,bbbbbb`*** is the comma separated list of server(s) in the “main” system. Do not include the servers in the secondary system.

Session layouts will then default to playing the calls from the main servers and fall back automatically to the secondary if no recordings are found for the specified contact on the main server(s).

High Availability Recording

If your license allows, you can configure one or more additional Avaya Contact Recorders to act as Standby to a Master recorder. The recording capacity provided by the Standby is available and is used by the Master in the same way as that of a Slave recorder but it can also take over the control of recording should the master fail. These would normally be **STANDING BY** and only become **IN CONTROL** when required.

Using additional recording servers as Standby and/or Slaves you can design in fault tolerance against server failure. These topologies are the same as those in previous versions but configuration is more explicitly and visibly controlled via the **General Setup > Server** page on each server rather than through the properties file and/or by inference from configuration history.

The fault tolerant options available depend on the type and number of telephone switches being recorded. The table below provides an overview of the options available to all systems. Systems recording a single Communication Manager[1] only can be configured to provide Network Failure Tolerance as well – as described in High Availability CM Topology on page 385.

[1]May include “overlay” systems such as Dialers, Avaya Oceana™ and AACC.

Ref	Topology	Recording	Recording Fault Tolerance	Notes
A	Single server (**Master**)	Centralized	Only that inherent in the server platform e.g. redundant fans, PSUs, RAID disks.	Server platform fault tolerance obviously applies to all servers in all topologies but is only highlighted here.
B	Higher Capacity (**Master** + **Slave**(s) co-located	Centralized	If “n” extra slaves are provided, tolerant to failure of up to “n” slave(s).	NOT tolerant to failure of the Master server. Strongly recommend purchase of Backup Channel Licenses and convert at least one slave to a **Full Standby** making this into Type D.
C	Distributed Recording (Master + slave(s) at least one of which is on separate site	Distributed	If “n” extra slaves are provided, tolerant to failure of up to “n” slave(s) - with option to specify alternate site that will provide cover if not available in same pool as failed slave(s). Ultimate fall-back to Master if possible.	Use **Designated Pool(s)** settings to specify preferred and fall-back recording pools for sets of recording targets. NOT tolerant to failure of **Master** server. Strongly recommend purchase of Backup Channel Licenses. Convert at least one **Slave** at main site to a **Full Standby** making this into Type D.
D	HA Pair (**Master** + co-located **Full Standby**) + zero or more slave(s)	Centralized or Distributed according to location of slaves and use (or not) of designated recorder setting	As type C **plus** Failure of **Master** server tolerated as **Full Standby** takes over control.	Requires Backup Channel Licenses. Master and Full Standby also provide recording capacity, so it is rare that the system needs many additional **Slaves**.

Multiple Switches

If your Avaya Contact Recorder is recording more than one Communication Manager and/or CS1000, there will be failure scenarios that cannot be handled well by a single server from a Master and Full Standby pair taking control. For example, the Master may be able to connect to one switch while the Standby is only able to connect to the other.

The first step to making this fault tolerant is to split the recording system into separate systems – one for each switch. The options for each are then as described below.

CS1000

A **Master** and a co-located **Full Standby** recorder can be connected to a single MLS feed to protect against failure of the Master recorder.

Alternatively, a **Full Standby** recorder may be co-located with and configured against a secondary call server to protect against failure of the main call server as well as failure of the **Master** recorder.

Communication Manager

If the recording system is connected to a single Communication Manager, recording servers can be deployed as **Full** and/or **Partial Standby** servers to mirror some or all of the Survivable Core and Remote Servers in the switch – giving the recording system very similar fault tolerance characteristics to the phone system itself.

High Availability Storage

There is absolutely *no* point designing a high availability recording system if the resultant recording is only present on a single disk drive and will be lost if that drive fails or is destroyed or stolen. There are three approaches that should be considered.

RAID Storage

All recorders must be configured with RAID1 or (preferably) RAID 5 arrays or be using a fault tolerant SAN on which recordings are kept.

Use of RAID arrays ensures the failure of a single drive does not compromise any recordings that have been made and stored on it. However, this does not protect against the destruction or loss of the entire RAID array or server, rack, room or building in which it is housed.

NOT Standby Recorders

Using Standby servers provides fault tolerant recording but does not, of itself, provide for fault tolerant *storage* of recordings. Until the recording exists somewhere other than inside the rack it was made in, it is at risk.

NOT Disk Shadowing or Backup

Automated snapshotting, mirroring or backing up of any of the recorders’ partitions is NOT an appropriate or supported approach. This will break the recorder.

See Backing up Voice Recordings on page 219 for an explanation.

Storage Attached Network (SAN)

These often provide mirrored storage in separate places – ensuring the loss of a server, rack, equipment room or even entire building or site does not result in the loss of both copies of the data.

If you have such a fault tolerant SAN, you can assign the calls storage path on each recorder to an area of this storage system. This is supported for locally connected SAN drives - NOT for Network Attached Storage (NAS).

You are then reliant on the mirroring or other redundancy and standby mechanisms associated with the storage network. These mechanisms typically include a tape library backup and might include hierarchical file storage (HFS), in which older files are replaced by small tokens that allow the system to retrieve the original content from the tape library.

The recorder runs successfully with Tivoli Storage Manager and might also support other similar systems. However, Avaya does not proactively test against these systems.

Connect to them at your own risk. If you want to use such a system, turn the recorder's own disk management function off. To do this, add the line **`disk.manager=false`** to the properties file. This stops the recorder from deleting the oldest files as the available disk capacity falls to 1GB. You must then ensure that the available space on the disk holding the calls directory does not fall below 1GB.

Archival to Network Attached Storage (NAS)

As noted above, NAS is not supported as the real-time storage location needed for a recorder’s calls buffer. The archival process, however, *is* designed to tolerate and recover from any network outages or bottlenecks between the recorder and the NAS.

The Archiving mechanism built into the recorder is designed to make copies of the recordings efficiently over even Wide Area Networks to Network Attached Storage (NAS). Archive destinations should be remote from the recorders, ideally on a separate site for maximum protection. Where fault-tolerant stores are used, ensure that the copies are in separate buildings and ideally separate sites.

There is a finite delay between the end of the recording and it being copied to the archive where it will be accessible even if the recorder is destroyed.

You can configure this delay – which defaults to 24 hours and 10 minutes from the end of the recording or when 100MB is waiting to be archived, whichever is sooner. This delay is the sum of the time the recorder waits before archiving (Minutes to wait before considering new recordings on page 182) and the maximum time before a file is written (Bundle into tar Files on page 192).

High Availability Retrieval

So long as recordings are stored securely, many users can tolerate not being able to retrieve them *immediately* – so long as failures are rare and the time taken to repair and reinstate search and retrieval functions is reasonable.

Any fault tolerant system is going to involve more than one recording server – so will therefore also require a centralised database of recordings from all recorders.

Each recorder may upload recordings to more than one Central Replay Server (in fact, in a Master + Full Standby configuration, this is the default). However, additional servers can be provided to create a regional and global hierarchy and/or to provide fault tolerant search and replay.

Independent Central Replay Servers

You can configure multiple servers as Central Replay Servers that act completely independently - allowing you to maintain separate sets of users and replay layouts. This is suitable for a partitioned, distributed approach but you will need to administer each separately and features such Locking recordings and Replay Authorization will be restricted to each server rather than shared.

Secondary Central Replay Server(s)

Alternatively, you can configure a Central Replay Server to be "secondary" to one already configured - and have the user accounts and replay layouts automatically copied between the two. This provides easier fallback and can also be used for load sharing.

To force a server to act as a **Secondary Central Replay Server**, enter the IP address of the CRS it is to shadow along with the license key.

You must also then configure *each* of the pair with the address of the other in their **General Setup > Server > Replay Server** settings.

When linked to a **Primary Central Replay Server** in this way:

- Each shows the link status between the two replay servers on the **Status > Server** page.
- Alarms should be raised if this link fails or cannot be established.
- The Secondary Central Replay Server should copy almost all of its configuration from the Primary Central Replay Server in the same way that a Standby copies it from a Master. However, Phone Replay port allocations (which are supported on Communication Manager systems only) must be entered separately into each server as each must be able to use its own ports regardless of the state of the other.

As with a Standby server, you cannot **Edit** or add configuration details that should be done on the primary.

Configuring Other Recorders

To configure either approach, set up each Central Replay Server as above and then add both addresses - separated by a semi-colon (';') - to the **Replay server(s)** setting on the **General Setup > Server** page.

High Availability CM Topology

This section describes how to design a recording system that has similar fault tolerance to the Communication Manager system it is recording.

It begins with a discussion of the various types of fault that can be designed for and the Beacon port mechanism that allows the recorder to infer when and how the Communication Manager system is fragmented.

Supported topologies that provide various levels of fault tolerance are then described.

Server Failure

The simplest – and not uncommon form of failure is a dead server. Note, however, that the more servers you have in a system, the more frequently you will experience server failure!

Redundant Controllers

As with Communication Manager, servers can be deployed as co-located, identical pairs. So long as these have fault tolerant[1] network connections between them, this provides simple and effective tolerance against failure of either server. Thus

- a **Full Standby** can be sited next to a **Master** on the main site.
- a **Full Standby** on a Disaster Recovery (DR) site where there is a Survivable Core Server can have an additional **Full Standby** sited next to it.
- a **Partial Standby** on a site with a Survivable Remote Server can have a second **Partial Standby** co-located with it.

Redundant Recording Capacity

Should a server fail, it must be possible to record the calls on the remaining server(s). This varies according to the type of recording but in general, the system will use whatever recording mechanism is available so a mixture of modes is possible. In the event of a recorder failing, live recordings are re-established on another server if possible.

[1] If the servers cannot communicate, each will have to assume the other is dead and try to take over so imperative this only happens if one really is dead, not just a single network fault.

DMCC

Provide one additional Slave server anywhere the design load is more than a single server can handle. This gives an “N+1 Pool” of recorders. For tolerance to “n” server failures, add “n” servers giving an “N+n Pool”.

Configure enough softphones for the remaining “N” servers to handle the full design load even if “n” servers have failed in such a way that they retain the softphones that they have been allocated already.

So in a 3 server (2+1), 1000 channel system, for example, provide enough softphones to allow the two remaining servers to handle the full load even if a third of the ports are stuck on the dead server - i.e. 1000 + 1000/3 + 20 per server headroom = 1393 ports at least.

Provide one additional DMCC port per recording server to act as its **Beacon** port.

SIPREC

Provide one additional Slave server anywhere the design load is more than a single server can handle. This gives an “N+1 Pool” of recorders. For tolerance to “n” server failures, add “n” servers giving an “N+n Pool”.

Passive IP

To record a call, a recorder must have at least one of its NIC cards connected to a SPAN port through which the RTP is being tapped. This can require that every SPAN port is duplicated and connected to two different servers.

Where multiple SPAN ports are required, these should be “salt and peppered” across the servers so that failure of any one server results in its load being shared across the remaining servers rather than all going to one.

The active controller automatically load-balances across the recorders that are able to record a given stream of RTP.

Alternatively, a sufficiently large pool of DMCC ports (and associated CM VoIP resources) will allow the system to continue recording (via DMCC) should any call not be recordable because the one and only recorder tapping its RTP is dead.

TDM

A call can only be recorded via TDM if the wires are physically connected to a port on a working recorder. TDM ports, being of no use for anything other than the call that is on them, are used in preference to other recording modes. This means that if the one and only recorder with a specific TDM tap has failed, recording will automatically fall back to DMCC or SIPrec recording if possible. A sufficiently large pool of DMCC ports (and associated CM VoIP resources) will allow the system to continue recording (via DMCC) should any call not be recordable because the one and only recorder tapping its TDM wires is dead.

Alternatively, you can use fully parallel, independent TDM systems as described in Fully Parallel Systems on page 376.

Screen Recording

Provide one additional Slave server anywhere the design load is more than a single server can handle.

Media Processing Resource Failure

In Communication Manager based systems using DMCC recording, you must provide one additional media processor per Media Gateway to ensure that failure of a single card does not impact the ability of the recording system to record at its design capacity.

Network Failure

In multi-site systems, there are often many more network outages than there are server failures. These range from a single Network Interface Card (NIC) failing, to the global system splitting into two or more “fragments” that cannot communicate with each other. Designing for this is more complex but can be done if the recording system is closely matched to the (one and only) Communication Manager system it is recording.

Complete Server Isolation

Where a server becomes completely isolated from the network (e.g. if its one and only NIC fails), then it may attempt to continue operating but will be unable to do so. Fortunately, it will also be unable to impact the operation of the other servers or the phone system – so has effectively failed in the eyes of the other servers and in this case fault tolerance is the same as if the server had failed completely.

Partial Loss of Connectivity

This can be dangerous and should be avoided wherever possible by using fault tolerant routing, bonded NIC cards, multiple NICs, salt-and-pepper Ethernet switch connectivity etc.

For example, if a **Standby** server is unable to connect to the **Master** server, it has to assume that the **Master** is “dead” and will try to take over control. If, however, the **Master** and **Standby** are both able to communicate with the phone system it is probably still trying to control recording. The two servers can then clash as they each try to use the same resources (e.g. softphones) and record the same calls. This is why the Master should be co-located with and hence easily provided with a fault tolerant network path to a **Full Standby** – ensuring that this scenario is extremely unlikely to occur.

Site Isolation

The most common failure in a WAN environment is for one or more sites to lose connectivity with one or more other sites – leading to multiple “islands” of connectivity – within which servers can communicate but cannot see into other islands until the breach is repaired. In Communication Manager, this can lead to servers controlling multiple, isolated “fragments”.

Co-located backup servers and N+1 pools do not provide tolerance to WAN failure modes. This can be provided for Communication Manager systems – which themselves are configured to handle such outages. However, due to the complexity of such systems, the first requirement is that the recording system is matched to the underlying Communication Manager’s failure modes – so can only be done for a single Communication Manager in a given recording system.

Best practice of how to use a combination of Full and Partial Standby servers matching the deployment of Communication Manager processors is described below.

Survivable Recording

For more sophisticated HA solutions the recording system can mirror the way in which the Communication Manager it is recording survives network outages. This section follows the terminology and approaches described in *Avaya Aura Communication Manager Survivability Options.*

Switching from “normal” operation to any survivable mode is a significant and disruptive event. It only occurs when it is extremely likely that a genuine network outage has actually occurred. Networks often self-heal around routing problems – typically within about a minute – and it is important that this is allowed to happen before control is switched to another site. Switching to survivable processor(s) therefore typically takes of the order of 3 minutes. There is no point the recording system trying to determine its best course of action until that has occurred. There is therefore a delay of 5 minutes deliberately introduced from when a recorder first learns of a problem affecting a different pool (hence, normally a different site) to the point at which it will decide whether or not to take control of recordings. This ensures that these servers will wait until the decision to switch to survivable mode (or not) has been made and that the Communication Manager has had time to come online and be ready to start accepting the recorder’s requests for TSAPI observers and DMCC softphones.

Note:

Full, distributed HA is provided for Bulk Recording only as remote sites are not fed WFO Business Rules.

Master and Standby Servers

Start by determining the number, location and configuration of the Master and each Standby server that you need. For optimum fault tolerance:

Locate Master ACR with Communication Manager’s Main server.

This server will normally control all recording in the system. It will normally also perform a fair share of recordings that are not explicitly assigned to any particular pool.

On the General Setup > Server tab:

1. Set its **Recorder Pool** to reflect the site’s role or location. For example: `main`, `hq`, or `london`

On the **General Setup** tab for the Communication Manager:

2. Configure a co-located AES (preferred) or a High Availability AES pair for both DMCC and TSAPI.
3. Specify one DMCC softphone as its **Beacon** port.

Protect against failure of the Master ACR (or its AES)

Add a Standby ACR server co-located with the Communication Manager’s Main Server’s duplex server. This will rapidly take over control of all recording should the Master fail or become non-viable (due to its AES failing, for example). It will also provide recording capacity for recordings that are not explicitly assigned to any particular pool.

On the General Setup > Server tab:

1. Set its **Recorder Pool** to be the same as that of the Master if co-located - or a different setting if remote from it and recording is not simply spread across both sites.
2. Set its **Standby Coverage** to **Full**.
3. Set its **Standby Readiness** to **Hot**.
4. Set its **Standby Priority** to **1**.

On the **General Setup** tab for the Communication Manager:

5. Connect it to a co-located AES for both DMCC and TSAPI. Unless part of an HA pair, this must be a different AES from the one the Master uses.
6. Specify one DMCC softphone as its **Beacon** port.

Recording Control with each Survivable Core Server[1]

Wherever there is a Survivable Core Server, we can provide a Full Standby server. On the **General Setup > Server** tab for each of these:

1. Set its **Recorder Pool** to the name of the site.
2. Set its **Standby Coverage** to **Full**, so that it will attempt to record everything the Master would have – though this will be inherently restricted to the set of port networks that have come under the control of the Survivable Core Server it is co-located with.
3. Set it to **Warm** Standby – so as to avoid unnecessary TSAPI observers being applied[2].
4. Set its **Standby Priority** to match that of the core server it is next to.
5. Connect it to a co-located AES.
6. Specify one DMCC softphone as its **Beacon** port.

With this configuration:

- Each **Full Standby** will provide recording services for whatever subset[3] of the overall system their Survivable Core Server is able to support.
- Each will ONLY go active if no viable, higher priority servers are visible to it.

Optional Backup Control with each Survivable Core Server.

If any of your Survivable Core Servers are a fault tolerant cluster, you can match this with a pair of recorders. For the second recorder on the site:

1. Configure it exactly as per the first **Full Standby** it is next to.
2. Ensure the co-located AES they connect to is HA or provide a second co-located one for this server.

The one with the higher serial number will claim priority, the other only going live if that one is dead.

Partial Standby with each Survivable Remote Server[4]

For each Survivable Remote Server, co-locate a Partial Standby server configured as follows:

1. Set its **Recorder Pool** to the name of the site.

---

[1] Formerly called Enterprise Survivable Server (ESS).

[2] As this server should only need to go active when ESS mode starts, it would have to apply observers again at that time anyway. There is a limit to how many observers can be present for each device so best not to have these present until needed.

[3] Which may be the entire system e.g. in many Disaster Recovery (DR) sites.

[4] Formerly called Local Survivable Processor (LSP)

2. Set its Standby Coverage to **Partial**, so that it will only attempt to record targets that have **Designated Recorder/Pool(s)** settings that include its pool name.
3. It will only offer **Warm** Standby – so as to avoid unnecessary TSAPI observers being applied.
4. Connect it to a co-located AES or to a Geo-redundant HA AES pair, the standby of which is co-located with it.

Each will ONLY go active if no viable **Master** or **Full Standby** servers are visible to it.

Recording Locations

Next, decide where you want recordings to be made – in both Sunny Day and Failover scenarios. Given the typical case of a “main” site, a disaster recovery (“DR”) site and a number of remote satellites sites, this could be:

1. **Centralized**:
   a. At main site in Sunny Day;
   b. DR if main is unreachable;
   c. remote sites only when these are isolated from both.
2. **Balanced**:
   a. Spread evenly across main and DR in Sunny Day;
   b. whichever of these is reachable if only one is reachable;
   c. remote sites only when isolated from both.
3. **Distributed**:
   a. stations recorded on same site as the station in Sunny Day;
   b. wherever possible in failover.

Often there are many more stations at the main and DR sites than there are in each remote site and a blanket rule can be used for these with only the exceptions on the remote sites being explicitly configured.

Default Server Selection

Start by specifying on each server’s General Setup > Server page whether it should take a Fair, Nominal or Zero **Sunny-day** Share of Undesignated Recordings.

Undesignated Recordings are all the recording targets for which you are not explicitly configuring any **Designated Pool(s)** using the Advanced settings.

The three options mean the following:

- **Fair**: will load balance recordings of targets with no **Designated Pool(s)** so that the load on each available recorder is similar.
- **Nominal**: will only make a single recording at a time on this server for targets with no **Designated Pool(s)**. This is just enough to catch problems when they occur – rather

than find out that a server cannot record only when it’s needed in failover mode. These recorders are only used for more undesignated recordings when they need to be – such as when a remote site has no other option as it has become isolated from all the recorders marked **Fair** or if all such recorders are at maximum capacity.

- **Zero**: will not normally record anything for targets that have no Designated Pool(s). Select this option if you plan to *explicitly* target some recordings onto this server’s pool.

So, for each of the recording strategies above you would configure as follows:

1. **Centralized**:
   a. Servers at main site take **Fair** share,
   b. Those at DR site **Nominal**
   c. Those at remote sites **Nominal**
2. **Balanced**:
   a. Servers at main and DR sites take **Fair** share
   b. Those at remote sites **Nominal**.
3. **Distributed**:
   a. All servers set to take **Zero**
   (though actually doesn’t matter as you will be explicitly setting the **Designated Pool(s)** for all targets so there will be no undesignated recordings).

Explicit Pool Designation

The **Designated Pool(s)** setting specifies which recording targets should be recorded where – and lets you explicitly set a priority order in case more than one option is available (as happens on internal calls where each end has a different setting, or where you deliberately list alternate pools).

This setting applies to sunny-day recording even if you have no Backup Channel License and hence no Standby servers. If you leave this setting blank (default) for any recording target, it/they will be load-balanced across the available servers as described above. However, if you enter a specific recorder pool, this pool will be used[1]. This lets you specify where recordings should be made in Sunny Day.

If you have deployed any **Partial Standby** servers, these will only attempt to record targets that include their own pool in the list of **Designated Pool(s)** setting - so you *must* specify this for those targets you want to be recorded by a **Partial Standby** should failover to it occur.

If you are happy for those recordings to be made on the remote site *always* – not just during failover to the remote standby – you just need to specify its pool. If, however, you

[1] There is an ultimate failover to a port on the controlling server itself if none of the specified resources are available

want recording to normally occur elsewhere – such as the “main” site – then you need to configure this other pool first in a list of pools.

For example, in a system with “main” (including Master), “dr” (including a Full Standby) and “backofbeyond” (where there is a **Partial Standby**) you could configure **Designated Pool(s)** for different sets of recording targets in several ways, including:

1. ***Left blank*** - to have targets load-balanced across all servers with Sunny-Day loading set to **Fair** (and a trickle to those on **Nominal**), without regard to which is recorded where.
2. **`Main/1;dr/2`** – for targets you want to record at the **`main`** site if possible, the **`dr`** site if not (and be ignored by the **Partial Standby**)
3. **`dr/1;main/2`** – for targets you want to record at the **`dr`** site if possible, the **`main`** site if not (and be ignored by the **Partial Standby**)
4. **`main/1;dr/2;backofbeyond/3`** – for those recording targets that are normally to be recorded at the main site, the dr site as a second-choice and the **`backofbeyond`** as third choice. The **Partial Standby** at **`backofbeyond`** would attempt to record these in its own pool of recorders if it went active[1].
5. **`Backofbeyond/1;dr/2;main/3`** – normally record these targets on the **`backofbeyond`** server pool, falling back to those in the **`dr`** pool if that’s not possible and to those in **`main`** if dr is not usable. Because backofbeyond is listed, the **Partial Standby** in that pool will try to record them if it ever takes control.

Add Servers to Support Required Load

You now know how many recording channels will be used where in both Sunny-day and Failover scenarios. Consider your peak load under Sunny-day conditions and then within each potential fragment that will have recording capability within it as follows:

1. Determine available capacity of the **Master** and **Standby** servers in each potential fragment. All can contribute to recording capacity as well as control recordings when needed.
2. Add additional server(s) on each site to support the design load that could be recorded there.
3. For “N+n” server fault tolerance within that fragment, add “n” (typically 1) additional server(s) to allow for the design load to be handled even during failure of “n” of the available servers.
4. If you now have more than one server on a site, but do not already have two that can control recording (Master or Standby) you may as well configure one

[1] Which it would only ever do if neither main nor dr pools were operative hence jumping straight to the third entry in the list as the only pool it can use.

of the remaining Slave servers identically[1] to the Standby on that site so that recording can be controlled locally even if *either* of those servers fails.

5. Configure further servers as Slave servers if you already have two Standby servers on a site.
6. Set each server’s **Sunny-day Share of Undesignated Recordings** as determined for the site.

Known limitations

Deploying Standby recorders:

- Does not guarantee access to recordings made by the failed server. You must also use centralized replay server(s) and archiving to achieve this.
- Does not guarantee "no loss" of recording on failure. See Mode of operation on page 395.
- Requires that the Standby recorder can contact the master recorder when it is designated as a Standby so that its licensing and/or configuration can be copied.
- Does not support the use of Telephone Replay ports on the Master and Standby recorders. Place these on a Central Replay Server (or pair of servers if fault tolerant replay is required).

Standby recorders and Unify/External Control

Controller connected to Master & Standby

When using Unify or an equivalent external controller with a Master and Standby recorder pair, it should send any START/STOP commands to the recorder that sent it port ONLINE messages.

Only the active recorder will send STARTED/STOPPED messages. Unify should only send TAG commands to this active recorder.

Controller connected to Master only

If the external controller is essential to the operation of the Master (for example, if it is setting up single step conferences) then the Master needs to recognize that it is not viable if it has no links to external controllers.

Set the property **`unify.required`** in the properties file to indicate this.

[1] If you configure Priority identically on both, they will choose between themselves on the basis of serial number.

You can then provide a second Unify/External controller connected to the standby recorder.

With this property set, the Unify/External controller can force a switchover to the standby by sending a **`1 BYE SHUTDOWN`** message and then not sending the **`HELLO`** response when the recorder re-establishes connection. When it is ready to take control again, it can send the **`HELLO`** message and the recorder will take over from the standby again.

Mode of operation

This section describes how recorders monitor each other; take over when needed and return to Standby mode when appropriate.

Fall-back Detection Beacons

Previous versions used the presence or absence of network links between recorders to infer that the Communication Manager system was similarly connected – or fragmented. This is not always the case – particularly after fragmented systems are restored into a single system. Failure to configure a Standby recorder correctly or to manually force it back to idle when normal operation is restored on fall-back to the main server could lead to ongoing conflicts between remote Standby servers and the central Master or Standby.

New in ACR 15.1FP2 is a fully automated mechanism that ensures that one *and only one* recorder is ever active within each “island” or “fragment” of connectivity that a large Communication Manager system can potentially be split into under failover conditions. Changes to the partitioning of the Communication Manager system are detected automatically within a few minutes and servers adjust their behavior accordingly.

This is achieved by using one dedicated DMCC softphone per recorder as a “beacon”:

1. Each recorder tries to register its own dedicated beacon port – and refuses to run if this fails.
2. If it succeeds in registering its beacon port, it knows it has DMCC services and could record calls[1]. It announces the beacon port in its heartbeat message to other servers along with other health indicators.
3. Each recorder notes the beacon port used by each other. This is persisted to the database and assumed to be valid on subsequent restarts.
4. Each recorder uses its own beacon port to telephone each of the other recorders’ beacon ports approximately once a minute.
5. If it succeeds, it knows that this recorder is in the same “island” or “fragment” of Communication Manager control as itself.

[1] Technically, it may be able to record via SIPREC or passive tap without DMCC but as DMCC ports are used as fall-back for these other methods, this connection is required for a server to be considered viable.

6. On receipt of a call from another recorder, the recorders exchange one or more DTMF digits that indicate their health with respect to the CTI feeds they require.
7. Each recorder therefore has a backup mechanism for the direct recorder to recorder links – from which it can tell not only the health of the other server but also that it is in the same Communication Manager.

Primary v Secondary Data Sources

The underlying Communication Manager is the “Primary” Data Source in this context. Without TSAPI and DMCC links to that, the server is essentially useless.

Secondary sources such as AACC, Avaya Oceana™ or PCS/PDS dialer feeds are highly desirable but not essential for all recording. A server with problems on one or more of these will only take control if there are no other servers with a full set of such data feeds functioning properly.

Power-On

A common case is that all recorders are powered on within a second or two of each other. To avoid destabilizing the system, Standby server(s) will always wait for a short pre-defined period longer than the Master would before deciding whether they need to go active or not. This ensures that on a “normal” power-up it is always the Master that takes control rather than a Standby.

Partial Standby recorders wait longer than Full Standby recorders – as they should only ever go live in a partitioned Communication Manager system and this only occurs after a few minutes of problems.

The time-period is determined by the **`oemcomms.connecttimeout=`*`xxx`*** property in the properties file. (***xxx*** is in seconds).

Standby mode

Once configured, a Standby server attempts to establish TCP/IP socket connections to the Master and any other Standby servers it is subservient to over the one or more IP addresses it has been configured with.

If it contacts the Master, it downloads the configuration details that are necessary for it to take over in the event of the Master failing. It continues to refresh these details every minute, thus keeping up to date with configuration changes.

Heartbeat messages are exchanged every few seconds.

Major changes to configuration (such as switching from Full to Partial standy) are advised in the “HelloMessage” only and hence require a restart of the server to take full effect.

Control Priority

Where more than one controller (Master or Standby) exists and is functional within a given Communication Manager fragment, the highest priority one takes control according to these rules:

1. Any server that is not viable will not take control (e.g. disk full).
2. A Master or Full Standby server with good primary CTI feed but problems on secondary CTI feeds will defer to one with good primary and secondary CTI feeds.
3. Any Standby will defer to a Master unless they are a Full Standby and are already **IN CONTROL** – in which case they will stay **IN CONTROL** until manual fallback is performed.
4. Any Partial Standby will defer to any Full Standby that has at least its Primary CTI viable.
5. Any time two Full Standby servers could both take control, the one with greatest Priority (lowest number) takes control.
6. Any time two Full Standby servers with the same Priority could both take control, the one with lowest serial number takes control.

Failure Detection

A Standby recorder will attempt to take control of recording in the following circumstances:

- The Fall-back Detection Beacon mechanism shows there are NO other higher priority active servers in this Communication Manager fragment.

AND

- it is the highest priority viable server in its Communication Manager fragment.

AND any one of

- The Master explicitly requests that the Standby take over because it has detected a fatal error such as hard disk full or switch connectivity lost completely and the Standby is reporting better health than it.
- On startup, if connection cannot be established with Master or other Standby servers within the allowed startup period.
- All TCP/IP sockets to the Master(s) and other Standby servers failing.
- Either or both TCP/IP sockets to the active controller are still active but the master does not respond to repeated heartbeat polls.

A master recorder is not aware of Standby recorder(s) until it has established contact with a Standby for the first time. Thereafter it will expect to establish contact with that Standby recorder always. (This can only be reversed by editing the underlying configuration database).

Note:

The loss of connectivity to any of the Data Sources configured is considered to be a fatal error and will cause a recorder to declare itself non-viable for recording control purposes (though it may still be able to record calls under the control of another server).

For example, a Master recorder that loses any one or more underlying Data Source connections will instruct a viable standby to take over from it. However, if that standby has problems with any one or more Data Sources, it will not have been declaring itself viable so the Master will retain control regardless of how many Data Sources it is connected to.

Disk Space Monitoring

Once a minute, each recorder will monitor the available disk space on the partitions used for the operating system, call details storage and recording storage. If any of these drops below 500MB a warning will be raised.

To override this default; set **`disk.warnAtMB=`*`nnn`*** in the properties file.

Should disk space fall to a critical low level (10MB) the server will declare itself not viable for recording or control allowing any other viable controller to take charge.

Active mode

On inferring failure of or being instructed to by the Master as above, a Standby server will try to start recording as per the automatically copied configuration or its local configuration.

A **Full Standby** will attempt to record in the same way that the Master was prior to failure.

A **Partial Standby** will only attempt to record targets that include its **Pool** in their **Designated Pool(s)** list.

Where softphones are marked with specific standby recorder numbers, only those designated for a given Standby will be registered. Note, however that this is no longer recommended.

Recordings will be made and their details uploaded to the Central Replay Server (if present).

Return to Standby mode

A Standby recorder will return to Standby mode:

- If the standby unit is shut-down before the failed Master is restored. When the Standby unit is rebooted, the master responds within the timeout mentioned above and the Standby recorder remains in Standby mode.

- If the Master recovers it may instruct the Standby recorder to return to idle immediately - or not until the administrator forces this from the **Status > System** page.
- If a server detects an active higher priority server in the same Communication Manager fragment as itself via the Beacon mechanism.

On relinquishing control, a **Partial Standby** will truncate all current recordings and un-register all softphones allowing the higher priority server to take over again. Note that this switch-back has the same impact on recordings as the initial switch-over. Hence you should only bring a failed unit back on-line in a controlled fallback out of hours.

Switchover Implications

It takes a few seconds to detect most of the failure modes. Although configurable using the properties file, this interval is a compromise between rapid detection of true failure versus risk of false alarms and/or "yo-yoing," a condition where the system goes unstable. The system defaults aim to detect failures within 10 seconds and should not normally be altered.

When a failure occurs and a Standby recorder becomes active:

- Recordings in progress may be interrupted. The partially completed .wav files for the recordings in progress might be manually recoverable (professional services chargeable). In G.729A mode, up to 1 minute of audio is buffered in memory and hence will not have been appended to the file. In G.711 mode, files are appended to every 30 seconds.

For **On Demand** and **Meeting** recording modes:

- Recordings in progress are dropped. The Standby recorder switches in within seconds, registers its softphones and awaits new calls arriving on these ports.

  The worst case loss of recording on a port is therefore determined by the speed with which the switch can re-register the failed softphones. The more channels on the failed recorder, the longer this process takes.

Restoring the Master

The fault condition that caused a Standby to become active can be removed in a number of ways, not all of which are ideal. These methods are:

- Power restored to Master
- Network connectivity restored to Master
- Connection to the switches CTI feed is restored
- Master repaired and reinstated

If the active recorder is a **Full Standby**, and is at least as healthy as the Master, then the Master recorder will allow it to continue so as not to interrupt recordings unnecessarily. The administrator may force the master to take control again by logging

into the administration web pages on the master and clicking the **Take Control** button on the **Status > System** page.

This will cause the Master to take over immediately. It should be performed out of hours as there will be a brief interruption in recording as the system switches over.

In cases where one or more **Partial Standby** recorders are, between them, providing distributed backups, this is less desirable than having the (single) Master regain control. It is also more difficult for the Master recorder to be confident that all recordings are being handled by these Standbys.

However, the Communication Manager system may still be fragmented so a Partial Standby will automatically revert to Standby mode only when the Beacon mechanism shows the Master or a Full Standby server is active in the same Communication Manager fragment as itself. Since fallback to a single system results in CTI links dropping and being reconnected, the longer delay that a Partial Standby implements before making its decision ensures that it will find the active server’s Beacon phone before it starts to go active and apply observers and register further softphones. This avoids any short term clash between the Partial Standby and the Master or Full Standby when the Communication Manager is healed back to a single system.

This is now fully automatic and works whether the Communication Manager heals automatically or manually.

Comparison with hardware switch-over units

Traditionally, high availability digital recording has been provided by inserting an "N+1 Switchover Unit" between the recorders and their audio sources. Although this provides a slightly faster switchover, it has several negatives in comparison to the software only approach used here:

- The cost of the switchover unit makes N+1 systems uneconomical for N=1 or 2.
- The additional cabling required to and from the switchover unit introduces further cost, failure mechanisms and opportunities for misconfiguration (swapped cables) that might not be noticed until after a failure has occurred.
- The switchover unit itself introduces a significant single-point of failure into the system.

Appendix E: Non-standard Hardware

This appendix discusses considerations for non-standard hardware such as blade servers. It covers the following topics:

- Overview on page 402

Overview

This manual and the automated installation processes assume a "typical" rack-mounted server as the server hardware. However, Avaya Contact Recorder has been installed on other hardware successfully, including blade servers. This appendix covers the considerations for such non-standard hardware.

Disks

SAN

Blade servers are normally equipped with locally attached SAN which is ideal for Avaya Contact Recorder. See Storage at Each Recorder on page 57 for further details.

Partitioning

The standard installation assumes either one or two physical disks, or one RAID array. In a more complex environment there could be separate arrays for different partitions. For example, a RAID 0 mirror for everything except call storage and a RAID 5 array for call storage. For Linux systems, you must edit the kickstart script to specify the appropriate disk device for each partition by changing the "on" parameter. See Expert kickstart options on page 103.

NICS

Blade servers often have redundant network configurations. It is therefore more likely that Avaya Contact Recorder should be installed with just one IP address (eth0 on Linux) instead of two (eth0 and eth1).

Appendix F: Advanced Security Settings

This appendix discusses some features and prerequisites for advanced security. It includes:

- SSL Certificates on page 404
- Creating and Signing Certificates Internally for RSA KMS and SSL on page 407
- Adding AES CA Root Certificates to ACR on page 414
- Changing Tomcat Port Numbers on page 414
- Encrypting Properties File entries on page 415
- Changing the Windows Service Account on page 415

SSL Certificates

Your SSL certificates must be specific to each installed server and should be signed by a trusted Certificate Authority.

In the instructions which follow:

1. Replace **<*installdir*>** with the location into which you installed Avaya Contact Recorder (always `/opt/witness` on Linux, typically `D:\Avaya\ACR152` on Windows).
2. Replace **<*keystorename*>** with the name of the keystore file being created.
3. Replace **<*alias*>** with the alias for the keystore file being created.

Backing up the Keystore files

The certificates and keys are stored beneath your installation folder in the folder: **<*installdir*>/keystore**

Because this folder contains any original, distributed certificates, it is important to make a backup of these. You may replace one of these files during the remaining steps. Should it be necessary to restore the original certificate, you can copy the backup to the original filename.

Creating a new Certificate

If you would like to test this implementation, you can practice this procedure with a certificate authority's 30-day trial certificate. Then, to implement real certificates, you can start over from this point.

To create a certificate:

1. Choose the public key algorithm. This can either be RSA or EC (elliptic curve). RSA is more compatible. EC is smaller and faster. If you choose EC, make sure that your CA can sign ECDSA certificates.
2. Log on to the server and change directory as follows:
   cd <*installdir*>/keystore
3. In all the following keytool commands, the exact syntax is operating system dependent. Use:

   Linux: **../jdk8/bin/keytool**

   Windows: **..\jdk8\bin\keytool**
4. If a file of the required name already exists, remove it as follows (replacing **<*keystorename*>** with the appropriate filename for this certificate).

   (Linux) **rm <*keystorename*>**

   (Windows) **del <*keystorename*>**

5. Run the java keytool utility with

   keytool **-genkeypair -keystore <*keystorename*> -alias <*alias*> -keyalg RSA –storetype PKCS12**

Replace RSA with EC as appropriate and specify the appropriate keystore filename.

6. Fill in the keytool prompts with the following:

   Password: Contact5tor3

Note:

You must type this password, exactly as shown. It is case sensitive.

i **First & Last Name**: enter the FQDN, IP address or intranet name

ii **Organizational Unit**: enter your division

iii **Organization**: enter your company name

iv **City/Location**: enter your location

v **State/Province**: enter your state

vi **Country Code**: enter the ISO 2 letter code for your country (for example, GB is the code for United Kingdom)

7. Enter **yes** if the information is correct.
8. Hit **enter** when prompted for the second password.

Generating a Certificate Signing Request

You need a Certificate Signing Request (CSR) as the first step of the signing process. When you have it, paste it into the Certificate Authority's web page. To generate a CSR:

1. Re-run the keytool command (still in the keystore directory as above).

   keytool **-certreq -keystore <*keystorename*> -alias <*alias*>**

2. Enter the password, which is **Contact5tor3**.
3. Copy and paste the output into the CA's web page. (Include the BEGIN and END lines.)
4. Complete the verification process.
5. Reply to the verification emails and other verification steps until you obtain a signed certificate back from the CA.

Importing the CA's certificates

Before you can import your certificate reply, you need to import the certificate authority's root certificate and any intermediate certificates between their root and your certificate.

To acquire these certificates:

1. Download these certificates from the certificate authority's website and save them in the keystore folder.
2. Save the root certificate as **rcert.crt** and any intermediate as **icert.crt**.

If you have more than one intermediate certificate, give them separate filenames. Save the certificate file the CA sent as cert.crt.

If it is not possible to save these certificates as .crt files, they may need to be converted. The certificates need to be in a .crt file type before eing imported into the keystore. To convert any certificates to .crt type, use openssl commands as shown in the following example:

- To convert a CA root certificate pem file (for example, rcert.pem) to a .crt file type use the command:

  openssl x509 –outform DER –in rcert.pem –out rcert.crt

- To convert a CA signed certificate of PKCS#7 type (for example, cert.p7b) to a .crt file type use the command:

  openssl pkcs7 –print_certs –in cert.p7b –out cert.crt

To import all your certificates:

1. Import the root certificate by running keytool:

   ***keytool* -importcert -keystore <*keystorename*>**

   -alias root -file rcert.crt

2. Enter the password - which is **Contact5tor3**.
3. Import the intermediate (if required).

   ***keytool* -importcert -keystore <*keystorename*> -alias inter -file icert.crt**

If you have more than one intermediate certificate, import them as **inter1**, etc.

4. Import your signed certificate.

   keytool –importcert –keystore <*keystorename*> -alias <*alias*> -file cert.crt

Note:

On Linux ACR, openssl is already available but on Windows ACR it is not. For Windows, you may need to install openssl separately.

Use the Microsoft Management Control (MMC) tool to import the root cert into the PC where the browser is running.

When these changes are complete, restart the ACR server for the changes to take effect.

Backing up the keystore file

The keystore file now contains:

- the random private key that is unique to this web server
- the signed certificate you just paid for

These two are linked and cannot be regenerated, so it is important to back up the keystore file. If either one of these components is lost, you must regenerate the certificate and pay again to get it signed.

Creating and Signing Certificates Internally for RSA KMS and SSL

Alternatively to having an external certificate authority, it is possible to host a certificate authority using open source software on a secure desktop environment. This allows to have a local certificate authority to authorise both the RSA KMS and SSL Certificates.

In the instructions which follow:

1. On a Secure Desktop Environment, create an "OpenSSL" directory which will be the working directory for the entire process.
2. Create the following file structure within this directory:
   a) OpenSSL\bin
   b) OpenSSL\conf
   c) OpenSSL\certs
   d) OpenSSL\keys
   e) OpenSSL\import
   f) OpenSSL\import\KMSClient
   g) OpenSSL\import\<acrFQDN>
   h) OpenSSL\import\<kmsFQDN>
   i) OpenSSL\db
   j) OpenSSL\db\database.txt (empty text file)
   k) OpenSSL\db\serial.txt (input "01" only into the text file)
   l) OpenSSL\signed_certs
   m) OpenSSL\signed_certs\KMSClient
   n) OpenSSL\signed_certs\<acrFQDN>
   o) OpenSSL\signed_certs\<kmsFQDN>

   Where <acrFQDN> is the Fully Qualified Domain name of each Avaya Contact Recorder which will connect to the RSA KMS.

And where <kmsFQDN> is the Fully Qualified Domain Name of the RSA Key Management Server

3. Download openssl-1.0.2xn-x64_86-win64.zip from https://indy.fulgan.com/SSL/ and save the unzipped contents to OpenSSL\bin.
4. Copy the following into their respective text files within OpenSSL\conf
   a. ca_conf.txt:

```
dir                             = .
certs_dir                       = $dir/certs
crl_dir                         = $dir/crl
db_dir                          = $dir/db
keys_dir                        = $dir/keys
req_dir                         = $dir/cert_req
RANDFILE                        = $keys_dir/private.rnd
#################################################################
[ ca ]
default_ca                      = ca_root

[ ca_root]
new_certs_dir                   = $certs_dir
database                        = $db_dir/database.txt
certificate                     = $certs_dir/cacert.pem
serial                          = $db_dir/serial.txt
crl                             = $crl_dir/crl.pem
private_key                     = $keys_dir/cakey.pem
x509_extensions                 = certificate_extensions
default_days                    = 730
default_crl_days                = 30
default_md                      = SHA256
preserve                        = no
policy                          = policy_match

[ certificate_extensions ]
basicConstraints                = CA:false

[ policy_match ]
countryName                     = optional
stateOrProvinceName             = optional
organizationName                = optional
organizationalUnitName          = optional
commonName                      = supplied
emailAddress                    = optional

#################################################################

[ req ]
default_bits                    = 2048
default_keyfile                 = $keys_dir/cakey.pem
```

```
default_md                      = SHA256
prompt                          = yes
distinguished_name              = root_ca_distinguished_name
attributes                      = req_attributes
x509_extensions                 = root_ca_extensions


[ root_ca_distinguished_name ]
commonName                      = Common Name (CN). No white
space
commonName_default              = RootCA
commonName_min                  = 2
commonName_max                  = 40
#organizationalUnitName         = Organization Unit Name (OU)
#organizationalUnitName_default = Recording System
#organizationalUnitName_min     = 2
#organizationalUnitName_max     = 20
organizationName                = Organization Name (O). No white
space
organizationName_default        = CompanyName
organizationName_min            = 2
organizationName_max            = 40

[ root_ca_extensions ]
basicConstraints                = CA:true

[ req_attributes ]
challengePassword               = A challenge password
challengePassword_min           = 4
challengePassword_max           = 20
```

b. req_conf.txt:

```
################################################################

[ req ]
default_bits                    = 2048
default_md                      = SHA1
prompt                          = yes
distinguished_name              = req_distinguished_name
#attributes                     = req_attributes
x509_extensions                 = req_extensions


[ req_distinguished_name ]
commonName                      = Common Name (CN)
commonName_default              = Localhost
commonName_min                  = 2
commonName_max                  = 40
```

```
organizationName                = Organization Name (O). Must
match CA's Organization Name
organizationName_default        = Company Name
organizationName_min            = 2
organizationName_max            = 40

[ req_attributes ]
challengePassword               = A challenge password
challengePassword_min           = 4
challengePassword_max           = 20

[ req_extensions ]
basicConstraints                = CA:false
```

5. Generate the Certificate Authority by opening the Windows Command Prompt in the OpenSSL Directory then writing the following command:

   bin\openssl req -x509 -sha256 -newkey rsa:2048 -days 3650 -config conf\ca_conf.txt -out certs\cacert.pem -outform PEM

   This command then prompts for the following details:

   a. "Enter PEM pass phrase:" - Choose a passphrase that will be used to sign each Certificate, so store this passphrase securely.

   b. "Common Name (CN). No white space [RootCA]:" - Choose the name for the Certificate Authority, which will appear on each Certificate issued.

   c. "Organization Name (O). No white space [CompanyName]:" - The organisation of the company, store this securely also as it will be required to sign each Certificates with this Certificate Authority.

6. Generate the RSA Key Management Server Certificate by opening the Windows Command Prompt in the OpenSSL Directory then writing the following commands.

   a. bin\openssl req -newkey rsa:2048 -config conf\req_conf.txt -keyout <kmsFQDN>key.pem -keyform PEM -passout pass:temppass -out <kmsFQDN>req.pem -outform PEM -subj "/CN=<kmsFQDN> /O=<OrganisationName>

   b. bin\openssl ca -notext -config conf\ca_conf.txt -in <kmsFQDN>req.pem -out <kmsFQDN>cert.pem

      Which will prompt for the following details:

      i. "Enter pass phrase for ./keys/cakey.pem:" - enter the passphrase given in Step 5.

      ii. "Sign the certificate? [y/n]:" - enter "y".

      iii. "1 out of 1 certificate requests certified, commit? [y/n]" - enter "y".

c. bin\openssl pkcs12 -export -in <kmsFQDN>cert.pem -inkey <kmsFQDN>key.pem -certfile certs\cacert.pem -passin pass:temppass -out signed_certs\<kmsFQDN>\I360KMClientCertKey.p12

Which will prompt for the following details:

i. "Enter Export password:" - enter a secure, possibly simple password that will be used to access the keystore.

Where: <kmsFQDN> is the Fully Qualified Domain Name of the RSA Key Management Server.
And where <OrganisationName> is the Organisation Name set in 5.a.

7. Generate the Key Management Client Certificate by opening the Windows Command Prompt in the OpenSSL Directory then writing the following commands:

a. bin\openssl req -newkey rsa:2048 -config conf\req_conf.txt -keyout kmsclientkey.pem -keyform PEM -passout pass:temppass -out kmsclientreq.pem -outform PEM -subj "/CN=kmsclient /O=<OrganisationName>

b. bin\openssl ca -notext -config conf\ca_conf.txt -in kmsclientreq.pem -out kmsclientcert.pem

Which will prompt for the following details:

i. "Enter pass phrase for ./keys/cakey.pem:" - enter the passphrase given in Step 5.

ii. "Sign the certificate? [y/n]:" - enter "y".

iii. "1 out of 1 certificate requests certified, commit? [y/n]" - enter "y" to commit changes for this new certificate.

c. bin\openssl pkcs12 -export -in kmsclientcert.pem -inkey kmsclientkey.pem -certfile certs\cacert.pem -passin pass:temppass -out signed_certs\kmsclient\I360KMClientCertKey.p12

Which will prompt for the following details:

i. "Enter Export password:" - enter a secure, complex password that will be used to access the keystore.

Where <OrganisationName> is the Organisation Name set in Step 5

8. Generate the Avaya Contact Recorder Certificate(s), by opening the Windows Command Prompt in the OpenSSL Directory then writing the following commands:

This process will need to be reproduced for each server which requires the connection of RSA KMS and SSL.

a. bin\openssl req -newkey rsa:2048 -config conf\req_conf.txt -keyout <acrFQDN>key.pem -keyform PEM -passout pass:temppass -out <acrFQDN>req.pem -outform PEM -subj "/CN=<acrFQDN> /O=<OrganisationName>"

b. bin\openssl ca -notext -config conf\ca_conf.txt -in <acrFQDN>req.pem -out <acrFQDN>cert.pem

   Which will prompt for the following details:

   i. "Enter pass phrase for ./keys/cakey.pem:" - enter the passphrase given in Step 5.

   ii. "Sign the certificate? [y/n]:" - enter "y"

   iii. "1 out of 1 certificate requests certified, commit? [y/n]" - enter "y" to commit changes for this new certificate

c. bin\openssl pkcs12 -export -in <acrFQDN>cert.pem -inkey <acrFQDN>key.pem -certfile certs\cacert.pem -passin pass:temppass -out signed_certs\<acrFQDN>\tomcat.p12 -passout pass:Contact5tor3 -name "tomcat”

Where: <acrFQDN> is the Fully Qualified Domain Name of each Avaya Contact Recorder which will connect to the RSA KMS.
And where <OrganisationName> is the Organisation set in 5.

9. Organise the Certificates and Keystores that will need to be imported onto the various servers in the instructions that follow:

   a. Copy "signed_certs\KMSClient\I360KMClientCertKey.p12" into "import\KMSClient\"

   b. Copy “signed_certs\<*kmsFQDN*>\I360KMClientCertKey.p12” into “import\<*kmsFQDN*>”

   c. Copy “signed_certs\<*acrFQDN*>\tomcat.p12” into “import\<*acrFQDN*>”

   d. Copy “certs\cacert.pem” to “import”

   e. Copy “certs\02.pem” to “import

      Where <acrFQDN> is the Fully Qualified Domain Name of each Avaya Contact Recorder which will connect to the RSA KMS
      And where <kmsFQDN> is the Fully Qualified Domain Name of the RSA Key Management Server

10. Import the relevant certificates into the RSA Key Management Server in the instructions that follow:

    a. Copy "import\02.pem" onto the RSA Key Management Server.

    b. Copy "import\cacert.pem" onto the RSA Key Management Server.

    c. Copy "import\<kmsFQDN>\I360KMClientCertKey.p12" onto the RSA Key Management Server.

    d. On the Key Management Server, open Start and search for "mmc"

    e. Select "File", "Add/Remove Snap-in…"

    f. Add "Certificates to "Selected snap-ins:"

    g. Select "Computer account" in the popup generated and click "Next"

h. Ensure "Local computer" is selected and click "Finish"

i. Now close the window by selecting "Ok".

j. Under "Console Root", expand "Certificates", and right-click "Personal" and select "All Tasks" > "Import…"

k. Click "Next"

l. Browse to the location where I360KMClientCertKey.p12 was copied to in step 10.c.

m. Input the secure export password, which was defined in Step 6.c and select "Next"

n. Close the window by selecting "Finish"

o. to "Personal" and find the newly imported certificates, and find the certificate which has both columns "Issued To" and "Issued By" set to the name of the Certificate Authority which was set in step 5a) as the Common Name.

p. Click and drag this certificate into the "Trusted Root Certification Authorities" section in the left pane, within "Certificates" and "Console Root".

q. Access Windows Run (Windows Key + R) and run "regedit", from here navigate to "HKEY_LOCAL_MACHINE" > "SYSTEM" > "CurrentControlSet" > "Control" > "Security Providers" > "SCHANNEL"

r. Add a new "DWORD (32-bit) Value" and label it "ClientAuthTrustMode" with a value of real value 2.

11. Import Certificates into the Avaya Contact Recorder(s) in the instructions that follow:

    a. Ensure the Avaya Contact Recorder Service is not active

    b. Copy "import\KMSClient\I360KMClientCertKey.p12" to the keystore directory on the Avaya Contact Recorder

    c. Copy "import\<acrFQDN>\tomcat.p12" to the keystore directory on the Avaya Contact Recorder

       The keystore directory always being /opt/witness/keystore on Linux and typically D:\Avaya\ACR152\keystore on Windows.

       And, where <acrFQDN> is the Fully Qualified Domain Name of each Avaya Contact Recorder which will connect to the RSA KMS.

    d. The Avaya Contact Recorder Service can now be active.

       Step 11 needs to be repeated for each Avaya Contact Recorder which is to connect with the RSA Key Management Server or have SSL enabled.

12. Configure the Avaya Contact Recorder(s) to enable RSA KMS by adding the RSA Key Management Server details to the Avaya Contact Recorder Web Service under "General Setup" > "Recorder" for the RSA Key Management

Server IP Address, and the Export Password for the RSA Key Management Client Certificate, set in step 7.c.

13. Considerations to be made during the Installation and Configuration for the RSA Key Management Software on the RSA Key Management Server:

    For the full installation and configuration of the RSA Key Management Server, please follow the "RSA Key Manager Server 3.5.2 Installation and Configuration Guide", available as part of the RSA Key Management Server Installation

    a. The "Install SSL Certificates" section under "Prerequisites for RSA Key Manager Server Installation has already been completed in Step 10.

    b. In the "Import Certificates into the KMS" section of the RSA Key Manager Guide, please upload 02.pem (originally copied onto the RSA Key Management Server in Step 10.a.) instead of the base certificate provided, I360KMClientCert.pem.

Adding AES CA Root Certificates to ACR

Go to the AES Server and navigate to **Security > Certificate Management > CA Trusted Certificates**.

Select each new CA Root Certificate in turn via the associated tick box.

Now click on the Export button and you will be presented with a window that contains the CA Root cert in a textual format.

Copy the entire contents of this text (including the -----BEGIN CERTIFICATE and END CERTIFICATE lines) and paste it into a filename of your choosing - a sensible choice would be the *CAname*`.pem`, for example, `inhouseca.pem`

Save this file in *installdir*/keystore/cacerts

Changing Tomcat Port Numbers

You can change the default http and https ports (8080 and 8443 respectively) by editing ***`installdir`*`/tomcat8/server.xml`**.

Locate the two **Connector** elements on roughly lines 20 and 30 and change 8080 to the chosen plain port number and 8443 to the chosen secure port number.

You must also set the following property in the properties file: **`acr.localport=`*`nnnn`*** specifying the replacement for 8080.

Note:

If you are using Central Replay Server, you must change all the recorders and the Central Replay Server. The port numbers must be consistent across all of these servers for upload to the Central Replay Server to work.

Note:

If you change the **HTTPS** port from the default of 8443 *and* you do not allow HTTP access, you must also set the following line in the properties file so that attempts to use HTTP are redirected to the correct HTTPS socket.
`https.socket=`*`nnnn`*

The server.xml may be overwritten on subsequent upgrades. You should keep a copy of the file after editing it so that this can be compared with any changed version and the appropriate set of merged changes determined after the upgrade.

Encrypting Properties File entries

Avaya provides a tool to encrypt your passwords so that they can safely be placed in properties files. You will need this tool if you want to, for example, change the PostgreSQL password.

Install the WitsBSUserCredentials tool provided in the utils folder of your installation disk following the instructions provided with it.

To encrypt a password for use in the properties file:

1. Launch the tool.
2. Select Other for the application type.
3. Enter a dummy username (it is not used) together with the password to be encrypted.
4. Use the encrypted password in the properties file.

Changing the Windows Service Account

By default, the Avaya Contact Recorder service is run as the Network Service account. To run the service under a domain account instead:

1. From the windows control panel launch the **Services** console.
2. Navigate to the ACR152 service.
3. Right-click select **Properties**.
4. Select the **Log On** tab.
5. Change the **Log on As: This Account** to the desired domain account.
6. Enter the **Password** and **Confirm password**.
7. Click **OK**.
8. Click **OK** on the **Services account granted Log On As A Service** right.
9. Ensure the account has read and write access to the install location and the folders below that location.
10. Start the service.

Appendix G: External APIs

Avaya Contact Recorder supports external APIs to control certain functions of the recorder.

Each API (excluding Search & Replay) must be enabled by setting a property value that is specific to the recorder serial number. Please contact your support representative for the exact details of each interface property.

HTTP/HTTPS access

The APIs allow both http and https connections. Both GET and POST requests are supported. For GET requests, the actual parameters can be added to the request URL. For POST requests, the parameters are encoded in the body of the request.

Authentication and Authorisation

The external APIs all use the same form of authentication and authorisation.

An external API role user must be created in ACR and given appropriate replay rights. Note that the ‘replay’ rights are actually used to determine which stations, agents etc. may be controlled.

Create the user with NO password set and then immediately attempt to use it – specifying the desired password in the Authorization header of the request. The password used on this first attempt will be noted and will be required for all subsequent commands from that user.

Alphanumeric owner replay rights are supported. This can help to simplify control of bulk recording targets. Such targets should be given an advanced owner field, with that owner also assigned as a replay right to the controlling username.

WebXAPI

Avaya Contact Recorder can be controlled by external command over HTTP/HTTPS as described below.

Error/Success codes

If the request is successful, WebXAPI returns a simply formatted response in XML, with HTTP response code 200.

If the request is unsuccessful, WebXAPI returns a non-200 HTTP response code.

503 SERVICE UNAVAILABLE

An attempt was made to use WebXAPI when it is not licensed.

501 NOT IMPLEMENTED

Returned if server is not active.

401 UNAUTHORIZED

WebXApi authentication failed.

- Invalid credentials
- Failed to decode credentials
- Failed to validate password

400 BAD REQUEST

- WebXApi command missing
- Unknown command
- Other unspecified error

Fault tolerance

Clients of the WebXAPI must be configured with the addresses of both the master and standby recorders. Clients can send commands to either server.

If the client cannot connect to one of the servers it should try the other.

If a command is received by an inactive server, it will send the client an HTTP redirect to the client with the IP address or FQDN of the active server. If the client detects that it was redirected it can send future commands to the other server for improved efficiency.

Controllable devices

XAPI uses the concept of controllable devices, which are typically stations or agents.

All commands must contain one of the following as a parameter:

- a **device** (the number of the station e.g. 2001)

- an **agent** (the number of the agent e.g. 2001)

- a **user** (the FQDN user configured with the required agent id e.g. user1@test.com)

These must be configured in ACR.

Commands

In all examples below, replace ***server*** with the name of the Avaya Contact Recorder server.

In addition, in all of the examples below, ensure that the Authorization header of the request has the appropriate basic authorization tag for the username you have granted API access and its associated password.

Pause/Resume

This pair of commands controls the masking of recordings that are happening / going to happen under the recorder's own control. Pause masks the audio, replacing it with a periodic double beep tone. Resume switches back to recording the actual audio.

A pause command can only be sent when a contact is in progress. If no resume is received before the start of the next contact, the next contact is not paused.

For example:

https://*server*:8443/webxapi?cmd=pause&agent=2001

https://*server*:8443/webxapi?cmd=resume&agent=2001

Note:

The property rec.maskallowed must be set to *true* for pause\resume to operate correctly.

Tag

This command tags the controlled device's contact with the additional metadata supplied.

The **scope** parameter controls whether the metadata is to be applied to the **current**, **previous** or **next** contact.

If **current** is specified but there is no current recording, it will be applied to the previous recording – on the assumption that the tag may have been received just too late.

Pairs of parameters specify the keys and values of user defined fields that the recording will be tagged with. The first key/value pair in the command should use parameters **udfname1** and **udfvalue1**. Further tags can be added using successively numbered **udfname** and **udfvalue** parameter pairs.

For further details on field naming, see User Defined Fields on page 335.

For example:

https://*server*:8443/webxapi?cmd=tag&device=1234&scope=current&udfname1=SSN&udfvalue1=111223333&udfname2=account&udfvalue2=1234567890

Start

Starts a new recording.

For example:

https://*server*:8443/webxapi?cmd=start&user=user1@test.com

Stop

Stops the current recording (if any).

For example:

https://*server*:8443/webxapi?cmd=stop&user=user1@test.com

Appendix H: GDPR

Avaya Contact Recorder supports the new General Data Protection Regulation (GDPR) and “right to be forgotten” legislation.

"Right to be forgotten" will be supported through deletion of calls (both current and archived) and user data (database) based on individual customers and employees.

This appendix provides details on how the ACR deletion process works and its main features.

The main sections of this appendix are:

- Introduction on page 420
- GDPR Delete on page 421
- Authentication and Authorisation on page 423
- Search and Delete on page 424
- Restful API on page 425
- Status and Alarms on page 425

Introduction

What is the General Data Protection Regulation?

The General Data Protection Regulation (GDPR) has been introduced by the European Commission to strengthen and unify data protection for individuals within the European Union. The Commission's primary objectives via the GDPR is to give citizens back the control of their personal data and to simplify the regulatory environment for international business.

The GDPR will come into effect on May 25th 2018, all companies and government organisations that offer goods and services to individuals within the European Union are subject to the new legislation. The GDPR applies to any organisation based inside, or outside of the European Union that collect and analyse personal data related to EU residents.

How will the GDPR affect Avaya ACR Customers?

Our customers, as a Data Controllers and/or Data Processors of customers’ and employees’ personal data, have an obligation to implement technical and organisational measures to demonstrate how data protection is integrated into its business activities. This includes the collection and processing of personal data according to the GDPR principles.

Purpose of this appendix?

This appendix focuses only on the implementation of the "Right to be forgotten" in the ACR. This is where an individual can ask that any data concerning them is removed. This covers all files/records and any copies.

GDPR Delete

The ACR GDPR delete feature enables the removal of all stored customer/employee details. These details include:

- Call and screen recordings stored on the Recorders
- Call and screen recordings archived within the ACR system
- Associated metadata stored within the ACR database and file system

The GDPR delete feature will not remove data outside the control of the system, i.e. remote logs, database backups, copied files.

Delete Selection

The interface to allow the selection of data to be deleted is based on a date range and at least one of the following:

- Participant
- Contact ID
- UDFs (searches all UDFs for matching patterns)
- Parties
- Specific UDFs (searches a specific UDF name for a matching value)

The output from the selection process will produce a list of Contact IDs and their corresponding details. The delete action will perform the deletion on all the Contact IDs identified in the search criteria.

Note: All information within the contacts identified will be removed.

The delete search will not select calls that have been locked (or are in the process of being locked) using the CallLock feature. Any locked calls must first be unlocked (again using the CallLock feature) in order to be included in the search results.

The delete search is limited to 100 results by default. The search limit can be increased using the property *delete.limit=nnn (where nnn is the number of results to display in the results pane).*

Deleted Information

The interface to delete calls, archives, and metadata is based on Contact IDs. All information pertaining to the contact IDs identified during the search will be deleted.

The delete process will provide the ability to purge data in the following areas:

- Archive
- Recorder Call Buffer
- Database

Delete Process

The GDPR delete process is initiated by the user issuing a delete request. The user will only be allowed to issue delete requests on the Central Replay Server (or the Master server if no CRS is configured).

Note:

If the system has multiple CRS servers configured, it may be necessary to delete contacts from each server individually, depending on your configuration. For example, where there are multiple recorder servers uploading to only one CRS in the set.

The ability to issue delete requests will be limited to users with the correct permissions (See Authentication and Authorisation).

To issue delete requests, the ACR provides the following interfaces:

1. **Restful API**

   A set of Restful APIs have been provided that allow delete requests and provide feedback on the status of the delete requests.

   Currently these are only exposed internally with the ACR system.

2. **Search and Delete**

   The Avaya Contact Recorder includes a search and delete application within it. The search and delete mechanism is a very simple browser based interface (See Search and Delete).

The GDPR delete request is processed by the server (Master or CRS) receiving the request in the following manner:

1. The delete request is validated, logged and stored persistently as a job to be actioned.
2. This job is now available for viewing on the Delete job status page, where the details can be viewed, or the job can be cancelled, if required.
3. The job is only actioned and the contact information deleted during overnight batch processing.

When the overnight batch process runs, any pending Delete jobs are processed; their status changes from Pending to Submitted, then to In-Progress. Once they have moved out of Pending status, they can no longer be cancelled.

Each individual server is responsible for performing the delete based on the criteria received. The following highlights the steps taken by each server to perform the deletion:

1. Identify the INums to delete based on delete criteria.
2. Remove associated calls from the servers' call buffer.
3. Delete from the archive(s).
4. Report status to issuer.

Note:

Deleting from archive overwrites the audio/screen binary information stored in the tar files rather than deleting them. The header information of the binary audio file will be retained.

Throughout the deletion process the ACR will provide status via the Status: Delete page, audit information, logs and validation that the operation has been executed.

The processing continues until all of the individual servers return their status.

Note:

Retired (mothballed) recorders must be removed from the topology before commencing the Delete operation. See "Retiring Recorders" section in the ACR Upgrade Guide for the correct procedure. This allows the CRS to send jobs to all servers (as the "mothballed" server is no longer present) without receiving any connection failures, which in turn allows it to delete any calls that it holds locally from the retired recorders and complete the delete job.

If they all report Success, then information about the contacts on the controlling host are added to the delete job and processed. If this also completes successfully, then the metadata details are deleted from the database and the delete job reports Success.

If any of the servers report Failure, then the overall job status is marked as Failure and the process completes.

Note:

i. In this case, the issuing server does not delete any of the metadata for that delete job. You must resubmit the entire delete request in order to get the issuing server to delete these items.

ii. You may not be able to replay any of the recordings that were successfully deleted within the request. The media for this item is likely to have been deleted already.

Authentication and Authorisation

To use this feature, a user must have the "May Delete" Role configured via the "Manage Users" page. This feature is only accessible from the CRS and/or Master server.

The user must also be granted one of the ACR administrative roles - system admin, or restricted admin.

Once a user has the "May Delete" Role, a new tab named “Delete” will appear to the right of the Replay tab.

Selecting this will open a new user interface labelled "Delete" which will allow selection and deletion of calls.

Additionally, a Delete sub-tab appears under the Status tab. This page displays the current status of Delete jobs that have been submitted.

The “May Delete” role works in conjunction with the ACR Access Rights. The Access rights control which recordings are searchable and will be displayed in the delete search results. Therefore, because the search results determine the data that will be deleted, the Access Rights control what the user is able to delete.

Search and Delete

The Avaya Contact Recorder includes a search and delete application within it (similar in operation to the Search and Replay application). This search and delete mechanism is a very simple and intuitive browser-based interface. For further details, see *Avaya Contact Recorder User Guide*.

The Search and Delete application is hosted on a web server running on the recorder itself. It uses a local database of recordings to allow users to search for recordings to be deleted. Users are required to search for calls according to:

- A date/time range

And at least one of the following:

- Participant
- Contact ID
- UDF
- Parties
- Specific UDF

Note: The number of ‘Specific UDF’s in a system will vary. By default the search will automatically present 5 UDFs (or less if fewer are configured). This number can be increased using the delete.udflimit=n property (where n is the number of user defined fields to display in the search filter pane).

A full description of the attributes that recordings are tagged with is given in Recording Attributes on page 312. Access restrictions determine which recordings ‘May Delete’ users are able to search and delete (see Access Rights on page 199 for further details).

The user can search for and delete any recording that matches their access rights.

To perform a delete, a user must first perform a search. Only once the search is performed can the user request the deletion. It is recommended to review the list of contacts displayed in the search because all data associated with the contact will be deleted.

Important: For each of the results displayed, all segments of the calls within each contact will be deleted. Deleting the call is irreversible and cannot be undone.

Restful API

The details of the restful API are internal to the ACR system. External applications may use this API to delete recordings from a single ACR on a per recorder basis. Details of this API are available upon request.

Status

Delete Job status is provided on a page at **Status > Delete** tab. This is only visible to Admin and Restricted Admin user that have the May Delete role enabled. The page gives a high level view of the delete jobs that have been issued.

Search criteria, based on the date range and status, can be used to restrict the number of jobs displayed.

Note:

The default display is the last month's worth of delete requests. This is controlled by the property setting `defaultstatusperiod=nn` (in months).

Purging Delete Status Job Information

Over time the status information collected on Delete Jobs accumulates and consumes database space. A purge batch process runs nightly to delete unwanted data and so limiting the amount of data that is stored. Any jobs that fall outside the retention period (if set) and don't have Pending status are permanently deleted.

A configuration option is available to set the limit. The **Retain Delete Job status** option is available from **Operations > I/O Jobs**. Specify how many months the data is held before it is permanently deleted. Specifying a value of 0, means that the data should be kept for as long as possible.

Any Delete Jobs that would be eligible for deletion but still have Pending status cause an alarm to be raised.

Cancelling Jobs

If a GDPR Job is still showing a status of Pending then it is possible to cancel that job.

1. Select the check box for the job to the right of the page.
2. Click **Cancel**.

Any GDPR jobs that have a status other that Pending cannot be cancelled. The check box for these jobs is disabled.

Viewing Job Details

There is a view link for each line of status presented. Clicking this link opens a popup that displays each of the items that were submitted for deletion on each server. The **Server #** column has a **View Alarms** link that takes you to the **Alarms** page of the corresponding server for that row. It is perfectly normal for a server to report 0/0 under the **Inums(s) Deleted** column when the job is successful, as this means there were no recordings to be deleted on that server.

In its initial state, 0/0 is displayed as the servers have not advised this server on how many deletions it requires or how many deletions it has done.

Audit Trail

A new audit trail event type ‘Delete’ can be selected from the ‘Audit Trail’ application.

The details of search and delete requests are added to the Avaya Contact Recorder audit trail.

An audit event is added on the completion (success, failure or cancellation) of a GDPR delete.

The details of the GDPR delete request and delete jobs (per server) are added to the ACR log file. The details of a delete jobs cancelled are also added to the ACR log file.

Limitations

1. This feature will process ACR data ONLY. Any data stored in the WFO domain will be handled via separate Verint-provided tools.

2. The feature will not handle the deletion of a single party from stereo recordings - both parties in the audio will be affected.

3. Deletion of contacts added to the system via CAM Import is not supported. These contacts must be deleted manually.

Glossary

ACD Automatic Call Distribution. This is a feature offered by the Avaya Communication Manager that queues and distributes incoming calls to waiting agents. Calls are queued until an agent is available. If multiple agents are available, calls are distributed on an equitable basis.

ACD DN The 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.

AE Services Avaya's Application Enablement Services are APIs to services such as telecoms, database, and so on. They include DMCC and TSAPI.

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.

AMS Avaya Media Server - a component of the Avaya Aura® Contact Center through which SIP calls are routed.

ANI Automatic Number Identification (Service). The provision of calling party information (typically, telephone number or billing/account number) to the called party.

Avaya CT Now known as AES TSAPI services. See TSAPI below.

BHCA Busy Hour Call Attempts. The number of calls that are attempted during the switches busy hour. Typically, slightly higher than, but often used interchangeably with BHCC.

BHCC Busy Hour Call Completions. The number of calls that are completed during the switches busy hour. Typically, slightly lower than, but often used interchangeably with BHCA.

CA Certificate Authority: An entity (typically a company) that issues digital certificates to other entities (organizations or individuals) to allow them to prove their identity to others.

CDN Controlled DN (CS1000 and AACC systems only). A CDN is similar to an ACD queue with no agents, a "holding place" for calls. Calls are queued in a CDN, and while in the queue, calls can receive treatment commands from a controlling application (for example, host application giving treatments such as music, ring-back, or silence, or routing the call). CDNs can operate in controlled or default

	(uncontrolled) mode. An active application controls calls in a CDN when the CDN is in controlled mode. In default mode, calls entering a CDN are immediately given the default treatment (for example, routed to an ACD DN and receive RAN, music, and so on), as specified in the default configuration in Overlay 23).
Codec	An abbreviation of COder/DECoder. A device or program that converts signals from one form to another. In this context, between different digital audio compression standards.
Communication Manager API	Now known as Device, Media and Call Control (DMCC).
CTI	Computer Telephony Integration - typically an interface through which a computer system can be advised of events occurring within a telephony system and/or control the telephony system. TSAPI is an example of such a link.
CUSTOMER	The CS1000 supports multitenant operations. The PBX can be partitioned into multiple, independent systems known as customers. Within each customer, DNs are unique and routes are private. Meridian Link is a system feature: that is, it covers all customers in the Meridian 1 PBX. However, when an application registers over Meridian Link (Application Registration message), it selects the customer it wants to work with.
Device, Media and Call Control (DMCC)	A software platform (part of AE Services, previously known as Communication Manager API or CMAPI) that applications such as the recorder use to create softphones that can participate in calls made on Avaya Communication Manager-based systems.
DID	Direct Inward Dialing. An attribute of a trunk. The CO passes the extension number of the called party over a DID trunk to the PBX when offering a call to the PBX. The PBX is then able to automatically route the call to that extension without requiring operator/attendant assistance. In this way, a single trunk can terminate calls for many different extensions (but not simultaneously).
DN	Directory number (CS1000 only). The number that identifies a phoneset on a PBX or in the public network. It is the number that a caller dials to establish a connection to the addressed party. The DN can be a local PBX extension (local DN), a public network telephone number, or an ACD-DN-the pilot or group number for an ACD queue.
Duplicate Media Streaming	A means of recording telephone calls by having the phoneset duplicate the packet streams that make up the real phone call - sending the copies to a specified IP address (the recorder).
Full Standby	A server that will attempt to take over control of the entire recording system should the Master (and any higher priority Full Standby servers) request it or be deemed to have failed.
HA	High Availability: more reliable and tolerant of failures than normal.

HFS	Hierarchical File Storage system. Typically, a combination of hard disk and tape drives plus controlling software that automatically migrates little used files to cheaper storage media.
IDN	Individual DN (CS1000). ACD agents using BCS sets can have additional DNs (in addition to their ACD key) configured on their phonesets. Key 0 is used to receive incoming ACD calls; other keys can support other DNs. IDNs are treated just like any other DN. 500/2500 set ACD agents can also have a (single) IDN configured.
IP	Internet Protocol. IP specifies the format of packets and the addressing scheme for internet data. The IP, like the postal system, allows you to address a package and drop it in the system. The packet may traverse multiple networks on the way to its ultimate destination.
IVR	Interactive Voice Response. A system/facility that plays voice menus to callers, and acts upon user input (typically, DTMF digits from a touch tone phone). It is sometimes called VRU (Voice Response Unit).
MADN	Multi-Appearance Directory Number (CS1000).
Master	An Avaya Contact Recorder that may be shadowed by one or more Standby Controllers. The Master will default to "Active" on startup unless a Full Standby recorder is already "Active". It will ask a Full Standby to take over its role if it determines that it cannot record e.g. a disk fills.
Meridian Link (ML)	Meridian 1's host interface. Meridian Link supports X.25 and LAPB communication protocols for host connectivity to CS1000.
NAS	Network Attached Storage. A term used for RAID, tape and other mass storage systems which have an integral network connection such as Ethernet or fiber-channel.
NIC	Network Interface Card. An expansion board that you insert into a computer so the computer can connect to a network. Most NICs are designed for a particular type of network, protocol, and media, although some can serve multiple networks.
Partial Standby	A server that will attempt to take over control of the servers in its own pool and perform those recordings designated to that pool should it determine that neither the Master nor any **Full Standby** server is in control.
Position ID	A unique identifier for a CS1000 phone set, used by the switch to route calls to the phones.
QoS	Quality of Service.
RAID	Redundant Array of Inexpensive Disks. RAID controllers use two or more hard disk drives together for fault tolerance and enhanced performance. RAID disk drives are used frequently on servers.

SAN	Storage Attached Network. A high-speed network that is typically part of an overall network of computing resources for an enterprise, in which the software knows the characteristics of storage devices and the quantity and value of the data stored in those devices.
Skill Hunt Group	A telephone number that is used to route calls to agents on the basis of the skills needed to handle the call. Agents are assigned to one or more such skill groups according to their expertise and appropriate calls are routed to them when they are available.
Slave	An Avaya Contact Recorder being controlled by the active Controlling Recorder (**Master**, **Full** or **Partial Standby**) recording what it is told to record.
Softphone	(SOFTware PHONE) In this context, a software-based emulation of an Avaya VoIP phone. Multiple such emulations run on the Avaya Contact Recorder and each can participate in a telephone call in order to record it.
Standby Controller	An Avaya Contact Recorder that shadows a Master Controller and takes over its role ("goes active") should the master fail or appear to fail (e.g. the standby loses contact with it).
Tagging	Adding details to the database of call recordings so that recordings can later be retrieved by searching through the available data fields.
TDM	Time Division Multiplexing. The traditional means of transmitting large numbers of voice calls over circuit switched networks - as distinct from VoIP.
TN	Terminal Number (CS1000 only). The physical address of a device (for example, a phone set, a trunk, an attendant) on the Meridian 1 PBX. The TN is composed of the loop, shelf, card, and unit IDs.
Trunk	A communications link between a PBX and the public central office (CO), or between PBXs, or between COs.
TSAPI	Telephony Services Application Program Interface. A CTI interface standard to which AES TSAPI conforms.
UTC	Universal Time Coordinated. A time scale that couples Greenwich Mean Time, which is based solely on the Earth's inconsistent rotation rate, with highly accurate atomic time. When atomic time and Earth time approach a one second difference, a leap second is calculated into UTC. UTC, like Greenwich Mean Time, is set at 0 degrees' longitude on the prime meridian.
VDN	Vector Directory Number. An extension number used in Avaya's ACD software to connect calls to a vector for processing. The VDN by itself can be dialed to access the vector from any extension connected to the switch. (See also Vector.)

Vector A list of steps that processes calls in a user-defined manner. The steps in a vector can send calls to splits, play announcements and music, disconnect calls, give calls a busy signal, or route calls to other destinations. (See also VDN.)

VoIP Voice over Internet Protocol. A means of transmitting telephone calls over the packet-switched IP network - as distinct from TDM.

VPN Virtual Private Network. Private, or restricted, communications networks which use encryption and other security measures to transmit information through a public network such as the Internet and avoid unauthorized use.

VRU Voice Response Unit. A device that plays voice menus to a caller and responds to caller instructions entered on a touch tone phone. Also known as Interactive Voice Response (IVR).

WFO Avaya Workforce Optimization Suite - an optional suite of applications that build on the underlying recording capabilities of the Avaya Contact Recorder.