



Customer Guide to Passive VoIP Recording

www.incontact.com

Recording

Customer Guide to Passive VoIP Recording

Customer Guide to Passive VoIP Recording

- **Version** — This guide should be used with inContact WFO v5.6 or later
- **Revision** — July 2016
- **About inContact** — inContact (NASDAQ: SAAS) is leader in cloud contact center software, helping organizations around the globe create customer and contact center employee experiences that are more personalized, more empowering and more engaging today, tomorrow and in the future. inContact focuses on continuous innovation and is the only provider to offer core contact center infrastructure, workforce optimization plus an enterprise-class telecommunications network for the most complete customer journey management. inContact offers customers a choice of deployment options. To learn more, visit www.inContact.com.
- **Copyright** — ©2016 inContact, Inc.
- **Disclaimer** — inContact reserves the right to update or append this document, as needed.
- **Contact** — Send suggestions or corrections regarding this guide to documentationsrequest-discover@incontact.com.

Table of Contents

| | |
|---|-----------|
| Introduction | 5 |
| Audience..... | 5 |
| Goals..... | 5 |
| Assumptions..... | 5 |
| Need-to-Knows | 6 |
| Customer Responsibilities | 6 |
| Passive VoIP Recording | 7 |
| Passive VoIP Recording Overview..... | 7 |
| Control Protocols..... | 7 |
| Real-time Transport Protocol | 7 |
| Known Limitations and Considerations..... | 9 |
| Telephony Requirements..... | 9 |
| Supported Control Protocols | 9 |
| Supported RTP Codecs | 10 |
| inContact | 10 |
| Hardware | 10 |
| Software | 10 |
| Licensing | 10 |
| Customer Configuration Overview..... | 11 |
| VoIP Network Preparation for Passive Recording | 12 |
| Network Topology Considerations | 12 |
| Passive Recording with Mirror Ports | 14 |

Customer Guide to Passive VoIP Recording

- Mirror Port Limitations 15
- Mirror Port Guidelines and Best Practices 16
- Common Mirror Port Issues..... 17
- Effects of NIC Type on Mirror Port-Based Recording..... 18
- Passive Recording with Hardware-Based Network Taps..... 18
- Appendix: PBX-Specific Examples 19**
 - Avaya 19
 - AVAYA G650..... 19
 - Avaya G700 20
 - IP Station Configuration on an Avaya PBX 21
 - ShoreTel 22
- Customer Administration Tasks..... 23**
 - Voice Boards Overview 23
 - Voice Board Configuration 23
 - Channel Configuration Settings..... 24
- Appendix: Channel Assignment Settings Definitions 25**
- Document Revision History 26**

Introduction

Audience

This document is written for customers and prospective customers interested in using inContact Call Recording in a VoIP telephony environment, either without a CTI integration or with an integration that depends on a separate audio source. Readers who will perform procedures in this guide should have a basic level of familiarity with IP telephony and control protocols, general networking, the Windows operating system, and inContact WFO.

Goals

The goal of this document is to provide knowledge, reference, and procedural information necessary to understand a proposed inContact WFO implementation using passive VoIP recording, and to configure the telephony environment to support this method of audio acquisition.

This document is NOT intended as a specific system or network design document, although it does contain network design overviews for reference. If further clarification is needed, consult with your telephony vendor(s).

Assumptions

This document assumes the reader has access to an inContact WFO Sales Engineer, Project Manager, or other resource to assist in applying this information to the reader's environment.

Need-to-Knows



To facilitate ease of use, this document takes advantage of PDF bookmarks. By opening the bookmark pane, readers can easily refer to the portion(s) of the guide that are relevant to their needs. For example, the inContact WFO application administrator can click on the **Customer Administration Tasks** bookmark to jump directly to that section.

To expand and collapse the bookmark pane, click on the bookmark icon on the left side of the document window.

For information and procedures related to inContact WFO configuration, consult your inContact WFO installation team.

If you are using passive VoIP recording in conjunction with a CTI integration, refer to the customer guide for that integration for additional requirements, tasks, and procedures.

Customer Responsibilities

You are responsible for supplying the physical connection(s), IP connection(s), or both to your VoIP telephony network.

Passive VoIP Recording

Passive VoIP Recording Overview

inContact WFO Call Recording can passively monitor or “sniff” VoIP transactions and record those transactions for long-term storage and retrieval.

This method redirects the VoIP traffic on a network to the inContact WFO server in one of two ways:

- Using port mirroring technology built into your existing network switches or
- Inserting a passive Ethernet network-tap device into your voice network

On a typical installation, a single network interface and quad-core processor can support up to 200 simultaneous conversations. This is inclusive of all traffic arriving at the interface, regardless of whether or not the conversation is recorded. inContact WFO still incurs processing overhead while manipulating the packets of any conversation received on the interface. On implementations over this size, traffic should be split to multiple interfaces and multiple servers.

VoIP communication usually involves two different network protocols: the **control (or signaling) protocol** and the Real-time Transport Protocol (commonly known as RTP).

Control Protocols

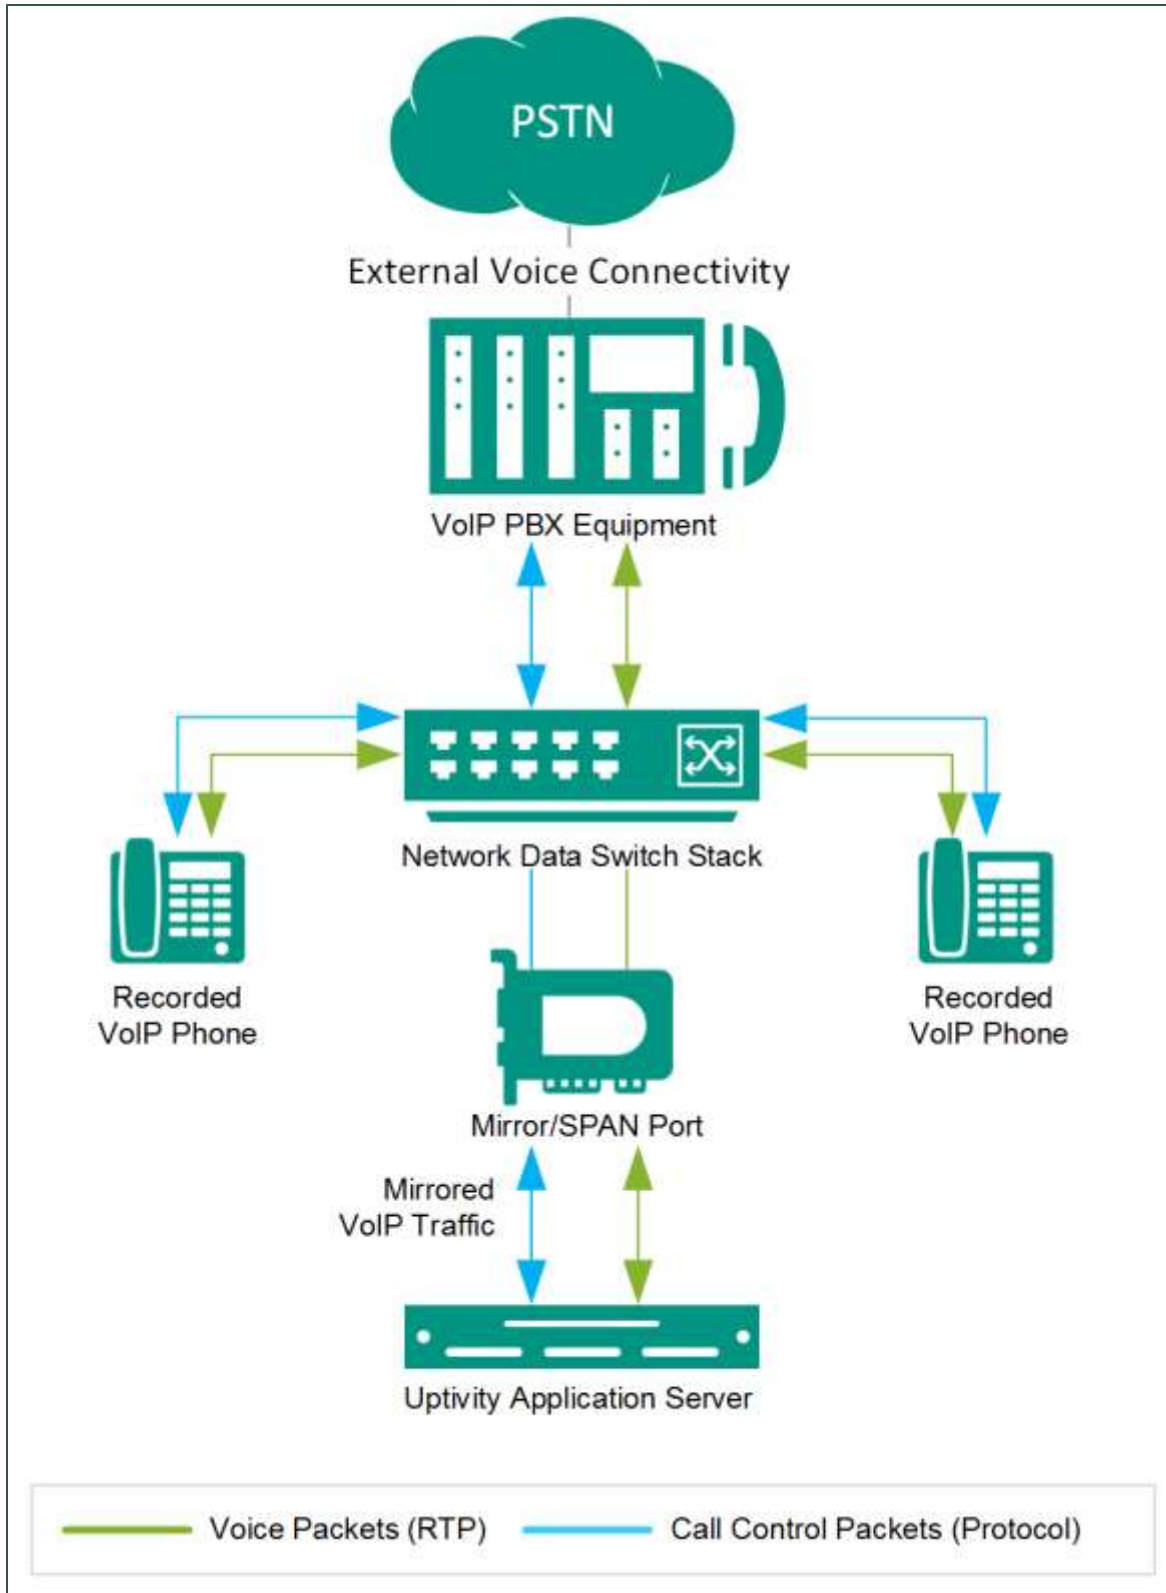
The control protocol handles call setup, phone features, number dialed and the parties involved in the call. It tells the VoIP phone how to set up the audio stream and where to send this audio stream.

The control protocol also communicates information between the phone and the PBX, such as the features the phone supports and the model and firmware version of the VoIP phone. The PBX also uses the control protocol to send information to the phone, such as LED states or text to display. For related information, see [Supported Control Protocols](#).

Real-time Transport Protocol

The Real-time Transport Protocol is a standard protocol used for sending audio (sometimes video as well). Occasionally RTP also sends DTMF (dual-tone multi-frequency) tones. This layer typically takes place within the framework of audio codecs. For related information, see [Supported RTP Codecs](#).

Customer Guide to Passive VoIP Recording



General architectural diagram of a passive VoIP recording implementation

Known Limitations and Considerations

- VoIP recording is only supported for ShoreTel SIP phones with TAPI integration. ShoreTel uses a proprietary encryption for the SIP traffic, making inContact WFO dependent on the TAPI messaging for recording.
- Traffic on your VoIP network cannot be encrypted.
- inContact WFO must support the control/signaling protocol used by your PBX.
- inContact WFO must support the audio codec used to encode your VoIP traffic.
- You must configure your VoIP network to deliver all necessary traffic to the inContact WFO system for processing (for more information, see [VoIP Network Preparation for Passive Recording](#)).
- Passive VoIP recording does not support the real-time blackout functionality in inContact WFO.
- The API command RECORDSTART is not supported. This command is used to create duplicate recordings, and is available only in Premises deployments.
- Virtual recording servers are not supported for passive VoIP recording. See [Effects of NIC Type on Mirror Port-Based Recording](#).

Telephony Requirements

Passive VoIP recording is dependent on the PBX and network topologies employed in the phone system. Due to the varying configurations and complexities possible, an inContact WFO Sales Engineer must determine whether the VoIP integration is viable, and how to deploy it properly.

Supported Control Protocols

inContact WFO supports passive VoIP recording in Avaya, Cisco, or ShoreTel environments using DHCP and any of the following VoIP protocols:

- **Avaya H.323:** Device reboot while inContact WFO is monitoring traffic
- **CISCO SCCP (Skinny):** Registration upon delivery of initial call to device
- **Avaya (Nortel) UNISTIM:** Registration upon delivery of initial call to device
- **SIP:** Registration upon delivery of initial call to device
- **ShoreTel MGCP:** Device reboot while inContact WFO is monitoring traffic

Customer Guide to Passive VoIP Recording

Supported RTP Codecs

inContact currently supports the following VoIP codecs:

- G.711 (A-law)
- G.711 (μ -law)
- G.729a
- iLBC
- L16
- G.722

Traffic to be recorded in your environment must be configured to use one of these codecs.

inContact

Hardware

inContact WFO and VoIP hardware requirements vary depending on system configurations and requirements. Appropriate hardware is identified during the system implementation process.

Software

This guide covers the following release:

- inContact WFO, 16.1 or later

Additional third-party software is required for passive VoIP recording:

- CACE WinPcap version 4.1.x

Licensing

- One (1) Voice Seat license per trunk channel to be recorded **or**
- One (1) Voice concurrent session license for each simultaneous call to be recorded
- Additional licensing may be required if the system includes optional features (for example, inContact Screen Recording)

Customer Configuration Overview

The following table provides a high-level overview of the customer configuration steps in passive VoIP recording. Links are provided for procedures covered in this guide.

| Customer Configuration Steps for Passive VoIP Recording | |
|---|---|
| 1 | Prepare your VoIP network for recording via port mirroring or hardware-based network taps. For more information, see VoIP Network Preparation for Passive Recording . |
| 2 | Complete all necessary physical and IP connections between the recording server(s) and the VoIP network taps/mirror ports. |
| 3 | Complete any applicable PBX configuration. For example, see IP Station Configuration on an Avaya PBX . |

VoIP Network Preparation for Passive Recording

Network Topology Considerations

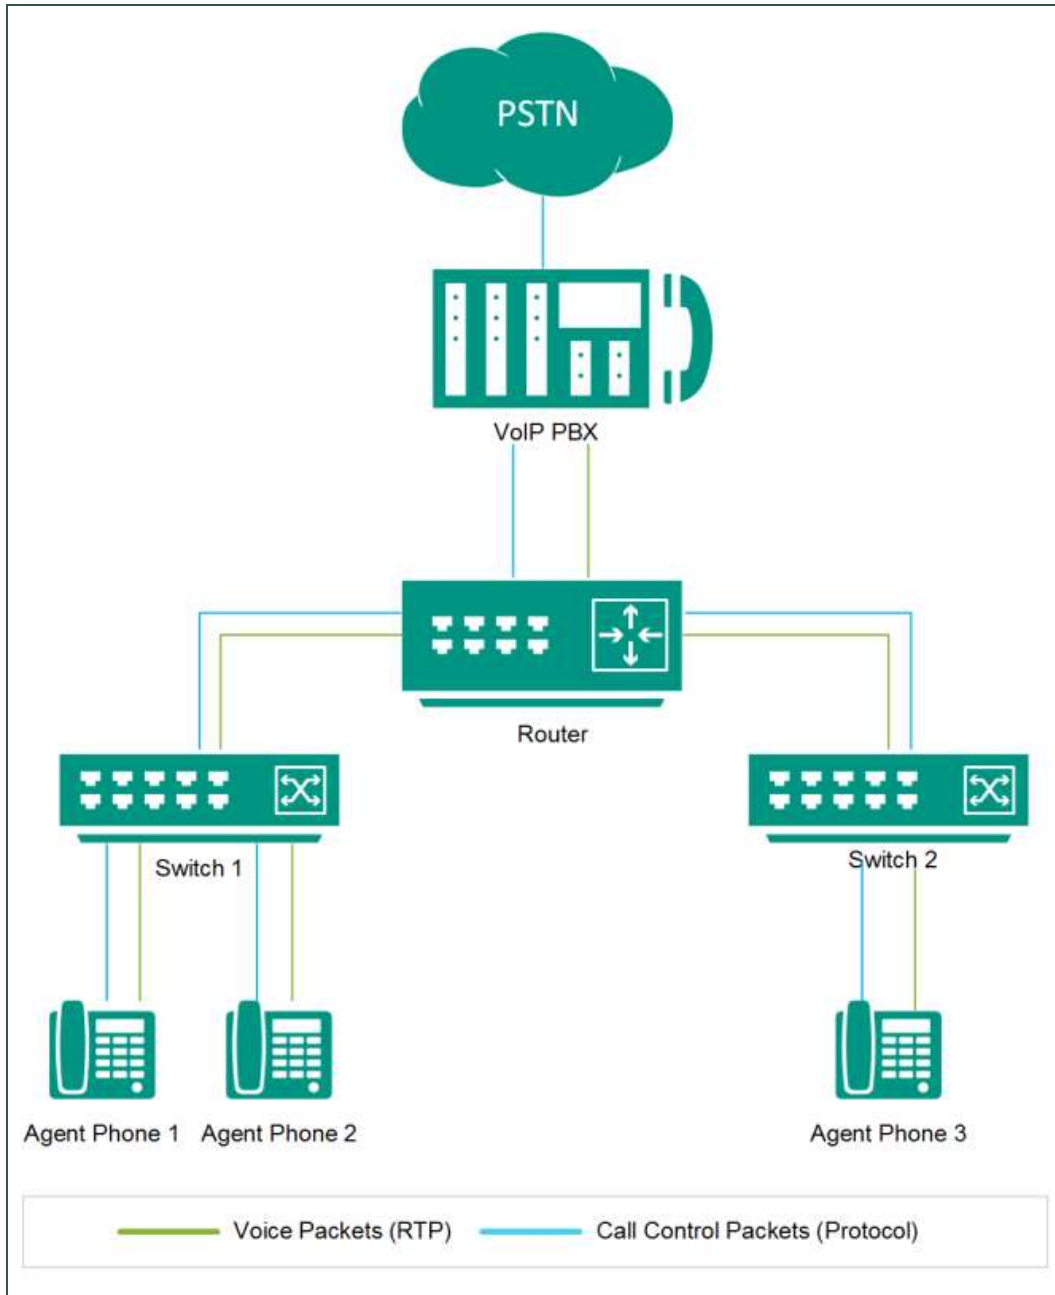
On traditional wired telephone networks, all voice and call control information passes through a central location—the PBX. Each channel on the network is tapped individually, and a single tap point obtains all voice data and call control information.

With VoIP networks, only the call control information is guaranteed to pass through the IP PBX. Once call setup is complete, voice packets are often routed along a different path from the signaling data. Thus, IP networks may not offer a central location to tap all voice and call control information.

For example, incoming calls typically enter the external-facing, VoIP-enabled PBX. The PBX negotiates the call with the IP phone and voice packets pass directly to the phone once the call is connected. Positioning a tap for inContact WFO between the router and each workgroup switch allows for capturing both control and voice packets.

However, in the case of peer-to-peer (internal calls) where both phones are on the same workgroup switch, the RTP voice packets are routed directly between the endpoints. Consequently, inContact WFO would not record these calls as the voice packets never leave the local subnet.

The following diagram illustrates this scenario.

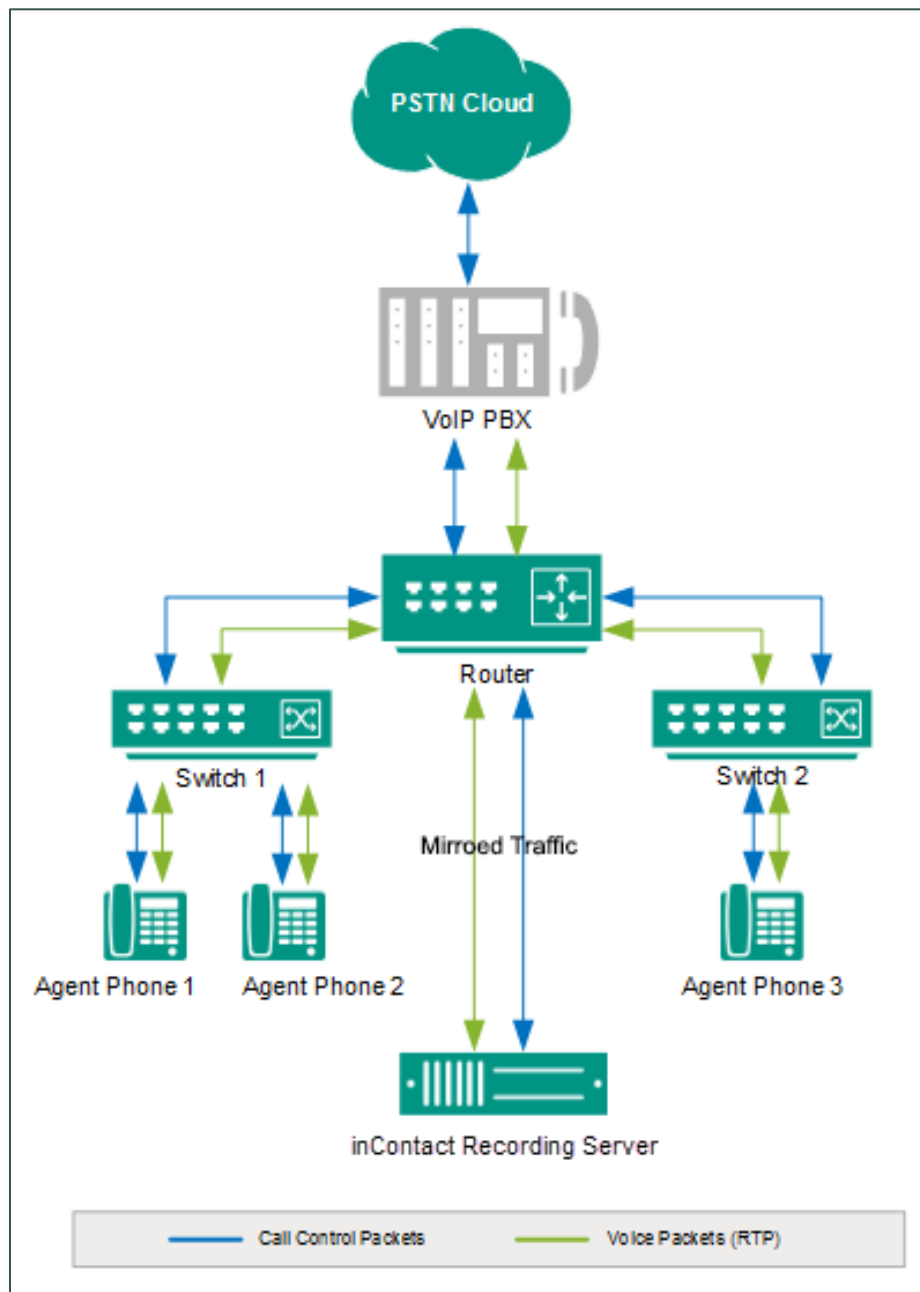


General example of VoIP network topology. A tap placed between the Router and Switch 1 would capture incoming calls, but not internal calls between Agent 1 and Agent 2.

Ultimately, inContact WFO can only process packets that it can see. If your organization wants to record all voice packets on the network (including agent-to-agent conversations), taps must be positioned deeper in the network. Depending on your needs, your inContact Sales Engineer may recommend positioning taps between phones and the workgroup switch, or may recommend relying on the switch's mirror port (also known as a SPAN port in Cisco environments).

Passive Recording with Mirror Ports

When mirror ports are used, data packets processed by the switch are replicated and passed off the network via this specialized port. The mirror port is connected to the inContact WFO recording server, which is capable of decoding signaling packets and processing RTP media.



General example of a VoIP network with an inContact WFO server connected to the router via a mirror (SPAN) port.

Mirror Port Limitations

There are two significant limitations to tapping a VoIP network with mirror ports. First, capacity limitations can be reached which result in packet loss. When mirror ports are used, data from multiple ports on the switch is aggregated and passed onto a single port. Access to full-duplex traffic is constrained by the mirror port's capacity. If a single 10/100 Mbps port attempts to monitor multiple 10/100 Mbps ports, data packets are dropped when the mirror port reaches its capacity.

Many switch manufacturers now design mirror ports with larger capacities to compensate for this. You and your inContact Sales Engineer should work together to determine the capacity of mirror port(s) on your system before assuming a single port can handle the traffic load.

The second consideration is the inherent design of the switch itself. In normal usage, the switch moves all network traffic. When packet mirroring is used, the switch has the added responsibility of duplicating packets passed to the mirror port. This added burden to the switch's processors impacts its performance. Some switch manufacturers, such as Cisco, build safeguards into their products which result in the switch giving a lower priority to mirrored data.

In other words, if any resource under load must choose between passing normal traffic and mirrored data, the normal traffic is passed and the mirrored frames are arbitrarily discarded. When the network is running at high capacity it can overload the switch. As a result, packet loss can occur even if the switch is designed with a gigabit mirror port.

Ultimately, the best design is a tapping solution that reduces processing needs placed on the switch and minimizes load on the mirror port. This type of system design reduces potential for packet loss. Depending on the capabilities of the switch, one or more of the following configuration options are recommended:

- Eliminate traffic from ports not connected to VoIP endpoints.
- Distribute packets across multiple mirror ports.

Eliminating Packets

This approach considers the types of devices connected to each port. For example, inContact WFO is not interested in data packets intended for a network device, such as a printer. To eliminate this traffic, only the ports connected to employee work stations equipped with a VoIP phone should be mirrored.

Most switches can be configured to enable mirroring control on a port-by-port basis. As a result, only the data packets of specified ports are replicated and pushed into the mirror port. This reduces overall load placed on the switch while at the same time reducing the potential of reaching the mirror port's capacity.

Customer Guide to Passive VoIP Recording

It is important to understand the control options supported by your switch model. Some switches allow users to select mirroring options based on a virtual LAN (VLAN). In this case, all ports assigned to a single VLAN are mirrored. Other models allow configuration based on individual ports.

Ultimately, the best design is one that eliminates unnecessary packets and limits the total number of packets being duplicated. While this is an effective strategy, it is important to remember that local networks can change (for example, a printer can be removed from a port and a phone can be added).

Distributing Packets

A tapping solution can be designed using a single switch with multiple mirror ports. Distributing traffic across multiple ports is an effective method when working on high-density networks. In this scenario, traffic from a few ports is duplicated to one mirror port, while the rest of network traffic is mirrored to another port.

This design option does not protect the processing needs of the switch, but it does reduce the load on each mirror port. This option is only available if the tapped switch supports multiple mirror sessions.

Additionally, your inContact hardware solution will require a NIC per mirror port. This option may incur additional hardware costs.

Mirror Port Guidelines and Best Practices

Before designing a tapping solution that relies on mirror ports, it is imperative to evaluate whether the port's capacity is large enough to handle all of the packets. Consider these guidelines in evaluating your mirror port(s):

- A single 10Base-T is considered to be running at full capacity when network rates reach about 6-7Mbps (60-70%). If this limit is breached, errors can be noted due to collisions. Most corporate networks are designed with this in mind.
- On low traffic networks, where each port remains at 10-20% capacity (1-2Mbps), a single 10Base-T mirror port is capable of monitoring a total of three ports (a total of 3-6Mbps of traffic). To monitor more ports, a 100Base-T mirror port is required.
- The same rule applies with a 100Base-T switch. If three ports are monitored, each running at 10-20Mbps, the total amount sent to the mirror port would be approximately 30-60Mbps. This is the maximum load that should be passed to the mirror port without risking packet loss. For high-capacity networks, multiple mirror ports should be used, with each connected to a single monitoring component.

- Where possible, mirror only the received (Rx) packets from each device/host on a port for individual phone ports. Add a second mirror for only the Rx packets from each switch gateway that passes RTP and control packets. This will result in mirroring both sides of the conversation without duplicating traffic.
- Mirroring both transmitted (Tx) and Rx packets on all ports is not advisable unless you do not intend to monitor the PBX endpoints. Otherwise, you will duplicate all the traffic since you will set it at both the phone and PBX taps.
- In most cases, VLAN mirroring is preferred to individual port mirroring. VoIP traffic is normally passed on a separate VLAN, so this is usually easiest to implement. On Cisco devices, this also allows you to use RSPAN to span the same VLAN on multiple switches. There is usually a limit to the number of individual ports you can mirror, so VLAN spanning may be the only way to capture traffic from all required ports. inContact WFO supports recording tagged VLAN traffic.

Common Mirror Port Issues

- **Mirroring only one direction:** Some switches only support mirroring one direction (Tx or Rx) on an interface. This can be overcome by deploying an aggregator or by using the Windows operating system to bond together multiple NICs (in other words, transmit on one NIC and receive on another).
- **Packet duplication:** inContact WFO has the ability to turn on a packet filter to eliminate duplicate packets, but this can impact performance on the recording server. inContact therefore recommends eliminating packet duplication as much as possible.
- **Not mirroring all phone-to-phone traffic:** When not all phone-to-phone traffic is mirrored, audio will be missed in some scenarios. This commonly occurs on larger distributed networks. For related information, see [Network Topology Considerations](#).
- **Saturating the mirror port:** Exceeding the interface's bandwidth rate limit causes missed audio and call events. The result can be gaps in audio, recordings that run together, or both. To overcome this issue, split off network segments to multiple switches and use an aggregator or multiple NICs.

Customer Guide to Passive VoIP Recording

Effects of NIC Type on Mirror Port-Based Recording

inContact WFO relies on the Windows operating system to provide data from any network interface. Thus, the application should be able to monitor traffic from any interface that can run in promiscuous mode. If the Windows operating system treats the device as a standard network interface, inContact WFO can interface with the NIC and record the network traffic.


inContact WFO Call Recording has been used successfully in bonded NIC environments, where two interfaces are bonded together on the Windows server. Bonded NICs are commonly deployed to aggregate transmit and receive traffic. inContact WFO has also successfully recorded fiber-based network interfaces.

i Virtual machines and VMWare applications typically do not allow for an interface to be accessed in promiscuous mode. Due to limitations and inconsistencies in the performance of these platforms, inContact WFO does not support passive VoIP recording in a virtualized environment.

Passive Recording with Hardware-Based Network Taps

inContact has partnered with DataCom Systems to provide customers with a physical, hardware-based tapping solution. These hardware devices can be passively injected into a network environment (for example, in between a PBX uplink port and a network switch). DataCom also has data aggregation units available if multiple tapping points are required.

These solutions vary in complexity and deployment configuration but are an option for organizations that are unable to or prefer not to use port mirroring. For more information on hardware-based network taps, ask your inContact WFO team.




Example: 10/100 Ethernet tap unit

Appendix: PBX-Specific Examples

The sample diagrams and configurations in this appendix illustrate how to capture VoIP traffic for specific PBX systems.

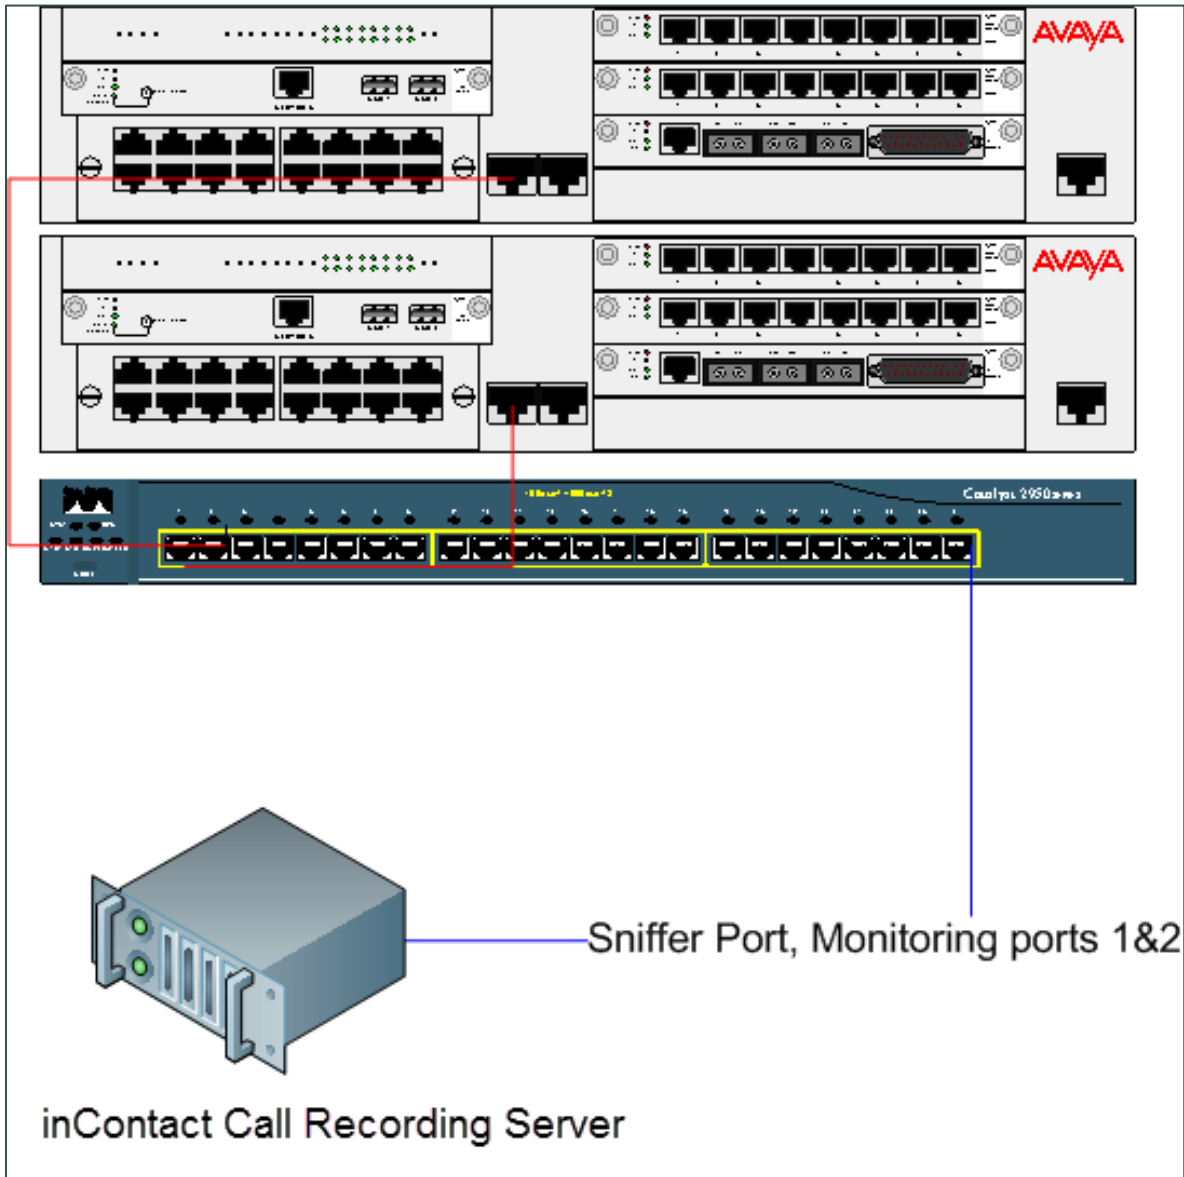
Avaya

AVAYA G650



Customer Guide to Passive VoIP Recording

Avaya G700



IP Station Configuration on an Avaya PBX

By default, IP phone to IP phone calls are permitted to bypass the phone switch. If the mirror port or network tap observes only traffic coming from the PBX and not the phones as well, internal calls will not record. For related information, see [Network Topology Considerations](#).

Using the Graphically-Enhanced DEFINITY Interface (commonly known as GEDI), you can typically force traffic through the PBX.

```

change station 7004                                     Page 2 of 3
STATION
FEATURE OPTIONS
  LWC Reception: spe                                     Auto Select Any Idle Appearance? n
  LWC Activation? y                                     Coverage Msg Retrieval? y
  LWC Log External Calls? n                             Auto Answer: none
  CDR Privacy? n                                       Data Restriction? n
  Redirect Notification? y                               Idle Appearance Preference? n
  Per Button Ring Control? n                             Bridged Idle Line Preference? n
  Bridged Call Alerting? n                             Restrict Last Appearance? n
  Active Station Ringing: single                         Conf/Trans on Primary Appearance? n

  H.320 Conversion? n                                 Per Station CPN - Send Calling Number? y
  Service Link Mode: as-needed
  Multimedia Mode: enhanced
  AUDIX Name: _____
  Audible Message Waiting? n
  Display Client Redirection? n
  Select Last Used Appearance? n
  Coverage After Forwarding? s
  Multimedia Early Answer? n
  Direct IP-IP Audio Connections? n
  Emergency Location Ext: 7004                         Always Use? n
  IP Audio Hairpinning? n

```

To configure this setting:

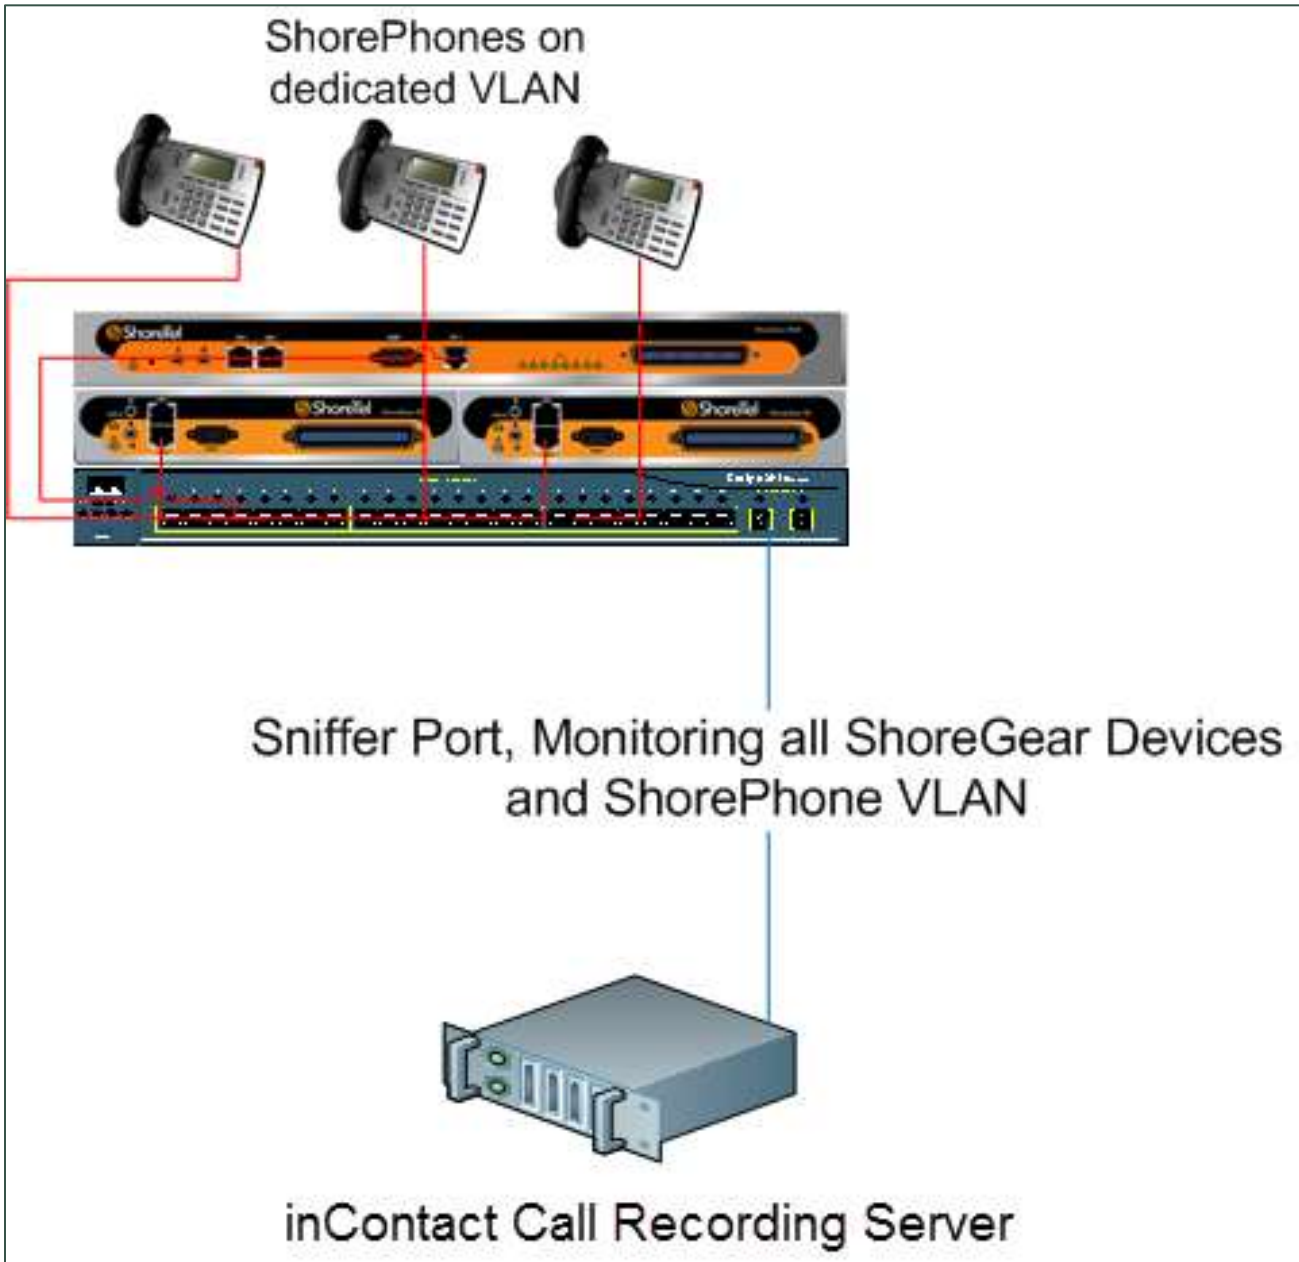
1. Log in to GEDI with an appropriately-permissioned account.
2. Type the command: **change station n**, where n is the station to be recorded.
3. On page 2 of the display, for **Direct IP-IP Audio Connections?**, type the value: **n** as shown in this image. This prevents the station from bypassing the PBX.

You will also need to verify that the system allows inContact WFO to observe and make two recordings for a single call (for example, agent-to-agent calls, conference calls with more than one participating agent, and so forth).

To enable this functionality:

1. Log in to GEDI with an appropriately-permissioned account.
2. Type the command: **change system-parameters features**.
3. On page 11 of the display, verify that **Allow Two Observers in Same Call?** is set to **y**.

ShoreTel



Customer Administration Tasks

During ongoing use of the system, your inContact WFO administrator may need to configure new channels or reconfigure existing channels. This integration requires changes to the Voice Boards page when channels are added or must be reconfigured. If you are using a CTI integration with passive VoIP recording, consult the customer guide for that integration to see if additional tasks are required.

Voice Boards Overview

Voice Boards control how inContact WFO acquires audio. This component provides **what** inContact WFO is to record. At least one Voice Board is required for most integrations. While Voice Boards can correspond to physical audio capture boards in some integrations, they are not those boards.

inContact WFO uses per-channel licensing, and each Voice Board software component maintains the count of licensed, used and available channels associated with it. The system will not use any Voice Boards or channels for which it is not licensed.

Voice Board Configuration

The basic procedure for configuring Voice Board channels is the same for all integrations and can be found in the online help for inContact WFO (<http://help.incontact.com/WFO/prem/help.htm>). For channel settings specific to this integration, see [Channel Configuration Settings](#). You must restart the Recorder service (cc_cticore.exe) after any Voice Board changes or channel changes.

Any Voice Board changes other than channel configuration should only be done under direct supervision from inContact WFO Support. Done incorrectly, Voice Board modifications can have serious negative impact to your system. In addition, altering the hardware configuration of your system may void your warranty.

Channel Configuration Settings

The following settings apply when configuring channels for passive VoIP recording:

| Setting | Definition | Value |
|---------------------|---|--|
| Assign | Used in deployments where physical devices and channels have a one-to-one correspondence, or to allocate specific channels to specific types of recording. For details, Appendix: Channel Assignment Settings Definitions . | Concurrent Licensing: Anything Per-Seat Licensing: Dedicated Record |
| Assign Value | If Assign is set to Anything , leave this field blank. If Assign is set to a Dedicated Record option, type the value for the corresponding device. | |
| Name | Type an optional name for the channel that can be used in channel scripting. | |

Appendix: Channel Assignment Settings Definitions

The following table lists and defines the values that can be selected for the **Assign** setting in channel configuration. Individual settings that appear in the inContact WFO Web Portal vary depending on the type of Voice Board and channel. Setting labels are affected by Terminology settings.

| Setting | Definition |
|---|---|
| Not in Use | Identifies a licensed channel that is not currently used. |
| Anything | Allows channel to be used for all recording and playback events, as determined by schedule priorities. |
| Playback Anything | Limits channel to playback of recordings via telephone. |
| Record Anything | Allows channel to be used for any scheduled or API-triggered recording. |
| Instant Record | Dedicates channel to instant recording requests from the API. |
| Dedicated Record ACD Group | Limits channel to recording only the specified ACD/PBX group (not inContact WFO Group), independently of any schedules. |
| Dedicated Record Device ID | Limits channel to recording a specific hardware resource (such as voice port or DN) on the ACD/PBX. |
| Dedicated Record Agent ID | Limits channel to recording a specific agent number or extension. |
| Dedicated Record Dialed Number | Limits channel to recording a specific inbound number, such as an 800-number carrying traffic to your facility. |
| Dedicated Record Caller ID | Limits channel to recording a specific ANI. Full or partial ANI matches may be used, for example, to limit to a matching area code. |
| Dedicated Record User1(2)(3)(4)(5) | Limits channel to recording a specific user-defined value as set by the API. Examples include Account and Case Number. |
| Playback and Instant Record | Limits channel to playback and instant recording requests from the API. |
| Playback and Record | Limit channel to scheduled recordings and playback. |
| Record and Instant Record | Limit calls to recording only, but of any recording type. |
| Unlicensed | Identifies a channel which is present (for example, on a physical audio capture card) but which is not licensed. |

Document Revision History

| Revision | Change Description | Effective Date |
|----------|----------------------------------|----------------|
| 0 | Initial release for this version | 2016-07-19 |