Optimizing the VoIP-Enabled Contact Center and Enterprise

October 2003

WITNESS SYSTEMS

Introduction

Voice over IP ('VoIP') is a hot topic in telecommunications circles, and increasingly within contact centers. From a world without VoIP just a few years ago, it is now becoming inevitable that VoIP will be deployed in everything from carriers' backbones to our offices and homes. Indeed, leading industry analysts project that within the next four years traditional call center telephony infrastructure suppliers like Avaya, Alcatel, Nortel Networks and Siemens will announce their intention to discontinue support within five years for system architectures based on Time Division Multiplexing (TDM).

The move from traditional circuit-switched to packet switched network technology will have a significant business impact on the applications used to record customer, supplier and other interactions, and the associated benefits such applications can deliver. IP telephony will enable smaller businesses to access application functionality previously afforded only by large organizations, yet also affording larger customers greater flexibility in their infrastructure with resulting cost savings.

IP Telephony is also enabling the development of enterprise-wide contact management systems, which embrace every department that touches the customer, and every department that deals with the outside world. These external interactions with customers, suppliers and other third parties can can include enormous amounts of valuable information and insight that, if made easily available to the right people in your business, may dramatically improve the service that you offer to your customers, and the overall business performance.

Until now, however, the valuable insight contained in these telephone calls has been, at best, only captured and shared within the contact center. With IP telephony, IP recording and applications such as the Witness eQuality suite, it is becoming viable to deploy recording both in the IP-enabled contact center and beyond. Analyzing all these external communications is more important than ever as every business seeks to gain a competitive advantage from better understanding their customers and acting on insights contained in interactions with them.

This document explains how IP Telephony impacts customer interaction recording and analysis, as well as how emerging solutions from companies like Witness Systems can support organizations as they migrate to or deploy IP telephony.

What is IP Telephony?

Traditional telephony uses circuit switched networks, which involves dedicating end-to-end bandwidth to each individual connection for the duration of the call, as in Figure 1. The guaranteed bandwidth and audio quality this offers is good, but it isrelatively inefficient. In contrast, VoIP is a means of transmitting speech over a packet-based infrastructure such as a data network or the Internet. The audio signals are chopped into small packets and are sent across the data network to their destination, as in Figure 2. By sending a stream of such packets in each direction, organizations can achieve two-way telephony with multiple conversations sharing common bandwidth, which is considerably more efficient in resource usage.

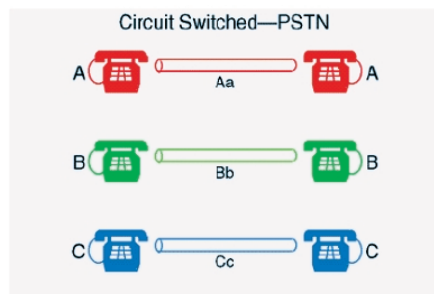

Figure 1: Topology of a circuit-switched telephony network

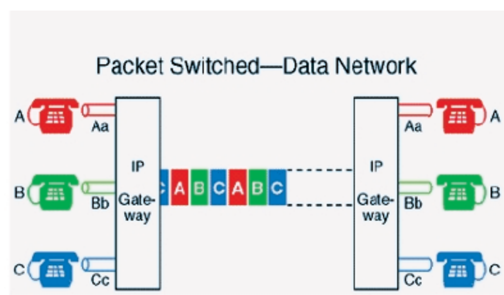

Figure 2: Topology of an IP Telephony network

As a basic technology, VoIP is therefore fairly simple to understand. The challenge is where and when to deploy it in the overall scheme of things. Companies are taking several approaches, the principal drivers behind which are these:

VoIP 'in the cloud'

The original hype around VoIP centered around using it to bypass the cost of international and long distance phone calls, exploiting the fact that data networks typically have a flat or per byte charging structure that ignores where and how far the packets travelled across the myriad public data networks – commonly referred to as the network "cloud." Unfortunately, because of the congestion and high latency on the Internet, it is a less than ideal medium for business quality telephone calls.

VoIP within a building
In contrast, within a building such as a contact center, transmission charges and bandwidth limitations are not an issue. Companies are reducing the cost, complexity and administrative overhead of the building wiring and desktop equipment through the use of VoIP. Using VoIP, either the PC on the desktop can take over all telephony functions or, if users prefer, they can connect a dedicated teleset directly to the Ethernet network. The control of telephony interactions is then performed by a software application running on PC servers, rather than a dedicated, proprietary and less flexible telephony switch. Based on these savings, we are therefore starting to see deployments of VoIP in internal environments, such as contact centers, particularly in "green field" sites.

VoIP in the enterprise
To large multi-site enterprises with existing, but separate, voice and data networks for inter-site traffic, VoIP is attractive as it provides a means of integrating these two disparate networks – reducing costs and getting better utilization from a shared infrastructure. The major vendors have developed technologies that allow organizations to deploy Virtual Private Networks (VPNs) with guaranteed Quality of Service (QoS) standards to support business quality telephone calls. However, replacing a network of legacy PBXs is far from trivial and full integration is a major exercise. Hence, companies using an IP carrier in this way may opt to use their existing circuit-switched PABXes within their individual premises, but route traffic between sites using an IP network.

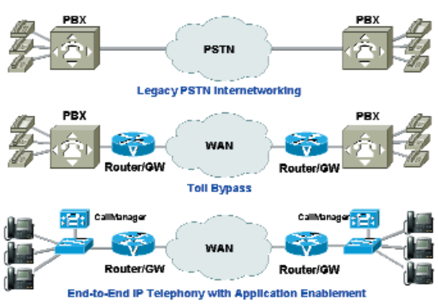

Figure 3: From traditional telephony to end-to-end IP telephony

VoIP Throughout
When VoIP is used in all the above ways, there are further savings to be gained through the elimination of costly gateways between IP and the Public Switched Telephone Network (PSTN), along with increased flexibility. However, with the huge investment in circuit-switched technologies around the world, gateways between the IP and circuit switched world will be with us for some time to come. In the meantime, companies deploying IP telephony within the contact center still communicate with their customers through a PSTN gateway and may route traffic between their sites using a circuit-switched network. In this case current recording systems can still be deployed on the circuit-switched side of the gateways, making IP recording an option but not a necessity. It is only when the complete speech path from the contact center desk to end customer is wholly over IP that IP recording becomes essential. Nevertheless, companies can gain many benefits from deploying IP recording within the IP telephony infrastructure, and the remainder of this document describes these features and benefits.

Recording of VoIP

Having set the scene for VoIP's gradual but inevitable introduction, how does this affect recording solutions such as that offered from Witness Systems?

Current recording systems are connected to the speech systems they record through two paths: the 'speech path' and the 'control path'. The former is the route by which the actual content of the conversation is recorded – typically a passive, high-impedance tap onto the extensions or trunks carrying the calls or, in call centers, optionally through a dedicated 'Silent Monitoring' port. The control path is typically a Computer Telephony Integration (CTI) output from the switch controlling the calls. The Witness Systems architecture actively controls the recording system by interpreting this CTI information, determining which speech paths to record at any given time and how to 'tag' these recordings with call attributes that will aid in their subsequent retrieval, and then stores both the call content and the call attribute details together.

Recording Control
As we move to VoIP, the control path is, conceptually, little affected. Even IP-based call controllers, such as Cisco's CallManager, have 'CTI' outputs – often supporting standards such as TAPI or JTAPI – through which the activity of the system can be observed. So long as these interfaces are rich enough to provide information on the IP socket addresses being used for transmissions, or provide means by which to establish conference calls onto the conversation to be recorded, organizations can use such interfaces to govern the recording and attributes in the same way that they support traditional switches.

For control and indexing of IP calls which will allow for sophisticated analysis and evaluation using the eQuality application suite, the above CTI integration is required in the same way it would be for a traditional, circuit-switched telephony system. However, for many basic recording needs, such as liability recording or emergency services call recording, very simple indexing is all that is needed. One of the major benefits of VoIP is that the call control information is normally transmitted over the same route as the speech packets themselves. Whether standard protocols such as H323, MGCP, SIP or lighter-weight proprietary protocols are used, there is often sufficient information contained in the packets used to establish the call to provide adequate attribute tagging of the call as well. Consequently, a simple analysis of these packets by the VoIP recorder provides a basic recording solution without the need to configure additional, often complex and costly, CTI solutions.

Speech Recording
In traditional telephony systems the circuit-switched voice path follows a predictable path through the telephony switch or 'ACD' that distributes the calls from incoming 'trunks' to 'customer service representatives' (CSRs') handsets or headsets. The recording system can then connect either to, or on either side of the ACD in order to record the calls. However, with a packet-switched network, the speech packets are transmitted over a path determined by the routers and switches in the network. In a complex network with multiple paths, it is difficult to predict where packets will be sent. Successive packets on a given connection can even follow different routes or arrive out of sequence.

As a result, companies can take two basic approaches to recording calls in an IP environment: “passive” – in which the recorder attempts to ‘sniff’ or ‘snoop’ the network without any of the parties involved in the call needing to be aware of this; “active” – in which at least one of the parties involved in the call is ensuring that the required data is deliberately sent to the recorder.

Given the gradual adoption of VoIP as described earlier, the passive approach is ideally suited to the majority scenario, where there are a limited number of points in the system through which calls enter and leave the contact center. By ‘tapping’ the gateways between the contact center’s CSRs and the outside world, companies can deploy a functional and effective recording solution. The passive approach is in fact the only possible approach in some cases – either because of the nature of the recording or due to limitations in the system being recorded. For example, some VoIP systems may not support the required volume of conferencing needed to perform bulk recording of all channels. The main limitation of this approach, apart from some constraints on topologies, is that internal calls cannot be recorded. However, this limitation is the same as for traditional, trunk-side recording schemes and is adequate for many call centers.

Overall, Witness Systems designed its ContactStore for IP recorder to use passive tap primarily because it has the benefit of having no impact on end-users; imposes no delay in setting up the call; no reliability impact on the system; scales well for bulk recording and makes it easy to record in stereo.

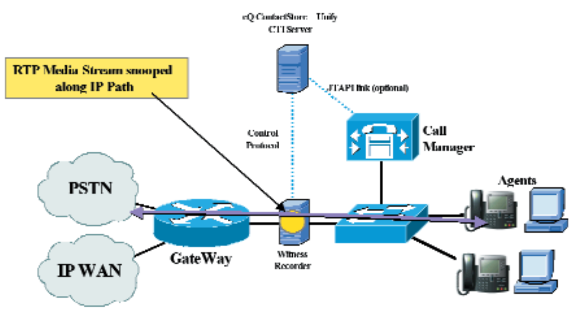

Figure 4: VoIP ‘Passive Tap’ Recording

‘SPAN’ and ‘VLAN’ Configurations

Where sophisticated routers and switches are employed, rather than physically connecting onto the same port as a Gateway (say, through an unswitched hub) it is often possible to configure the network to take advantage of Virtual LANs (VLANs) and the ‘SPAN’ feature. The leading VoIP products allow the logical partitioning of components on the network, then they can set up a switch output to ‘SPAN’ this network and copy all packets required to the port on which the recorder is connected.

This solution allows one or more recorders to monitor the packets associated with all telephony calls on the system – both internal and external - with the sole exception of those going from one PC to another (e.g., ‘softphone’ to ‘softphone’ calls).

Entry Level Solution

As discussed in the “Recording Control” section, where recordings need only be indexed with basic call information that can be gleaned from the call control packets, there is no need to provide the CTI link (shown as dotted lines in the above diagram) and a simpler, single connection recording node can provide the complete solution. The precise details that can be gleaned from these packets depend on the protocols used for call control.

Impact of VoIP on Recording Systems

The increasing use of VoIP within contact centers and the networks to which they connect should bring a number of opportunities over the coming years with significant benefits as outlined below:

Hardware commoditization

Becauase the voice packets are typically carried over standard fast ethernet connections, there will no longer be a need for specialist telephony interface cards and other connection hardware. Hence we will see a gradual move to the use of standard PC chassis as recorders rather than the proprietary chassis of today. This gives greater flexibility and a lower hardware cost than traditional telephony solutions. The Witness Systems eQuality ContactStore for IP solution already uses such an open systems approach.

Stereo Recording

With VoIP communication, individual packet streams travel in each direction between the two parties on the call. This makes it straightforward to record the two streams as the two halves of a ‘stereo’ signal. The value of this is seen when the recordings need to be replayed. Graphically showing who was speaking when and the balance of speaking, listening, silence and interruption between each party can be very informative. For instance, conversations with a high incidence of interruption can indicate an angry customer call which is worth reviewing.

Furthermore, as speech recognition algorithms are deployed as part of a recording solution, keeping the speech from the two parties separate is critical to high recognition accuracy. A mixed signal is much more difficult to analyze than a separated pair of signals.

Recording ‘in the cloud’

One of the attractions of VoIP is that it makes it easier to provide complex voice and data applications ‘in the cloud’ without the need for dedicated hardware at the customers’ premises. In the same way that the traditional ACD and IVR components are being ‘virtualized’ so it will be possible to offer recording services as part of the functionality of the network – with no hardware on the customer’s site. Recorders will be connected either to the ‘up-link’ to the ISP or deeper in the carrier’s network.

The earlier discussion on the routing of packets within a customer’s network also applies and the carriers will rely on VLAN type techniques to ensure that packets from customers who have purchased recording services are routed to the recorder pool in the network. The system will perform search and replay through browser-based applications and the recordings will be streamed over the network to the customer whenever replaying a call.

A step along the way to this model is already becoming prevalent in the outsourced call center model, with calls routed through a gateway between the VPN and PSTN in the home country, then out to the outsourced call center using VoIP over a VPN. Again, search and replay is performed through browser-based applications to allow users in the home country to search for and replay calls.

Distributed Call Centers
One of the uses of the above Application Service Provider (ASP) model is in building less formal and more distributed call centers. Because there is no longer a requirement for expensive ACD, IVR and other systems, such as recorders, it becomes much more cost effective to deploy small teams of people at multiple sites whenever they are needed. For example, back office staff in a branch may be used as an 'overflow' resource to assist the main call center when it is overloaded. In such environments, recording nodes can be placed in the remote offices and managed over the WAN, or at the central office if all VoIP packets are routed past them there.

Home Working
The ultimate distributed call center is where individual users are working from home with nothing more than a data link into their company's network. Rather than requiring separate data and speech lines, organizations can use VoIP to allow operation over a single high-speed line, such as DSL or cable modem. In this environment, a shared recording node may be located at the central facility from which calls are routed out to CSRs' homes.

'PABX' Recording
Outside of the formal call center, many traditional PABX users have modest recording requirements, but can benefit from the ability to invoke recording on an ad-hoc or "on-demand" basis. For instance, the recording of "promises to pay" in a credit control department; the recording of switchboard calls for security threats; or recording of conference calls which other interested parties can subsequently access.

The move to VoIP makes it much easier to provide sophisticated but cost-effective recording services without the need for expensive hardware. We can therefore expect to see more widespread adoption of recording across the enterprise.

Witness Systems IP Leadership

Worldwide Customer Base
Witness Systems has deployed nearly 200 IP-specific solutions to customers worldwide over the past two years. On the back of this experience Witness Systems have built up the most extensive experience of recording in the IP environment of any vendor. Proven, field tested solutions in the most demanding and mission-critical environments, particularly in financial services, makes Witness Systems the right choice for IP recording.

Integrated part of the feature-rich eQuality suite of applications
With a flexible, resilient, reliable recording platform as a foundation, customers for ContactStore for IP can leverage more value from their investment by deploying applications that can exploit the customer interaction intelligence captured by the recording platform. Witness Systems offers the most complete suite of applications available, from quality monitoring and e-learning, call analytics and visualization, to call center performance management.

Open Systems
eQuality ContactStore for IP is a software-only recorder that can be deployed on an industry-standard PC server without the need for any proprietary hardware. As the servers become more powerful, the scalability of the recorder can continue to increase, while minimizing the investment in specialized administration skills.

Desktop control
By exploiting the XML browser interface on the IP telephone, eQuality ContactStore for IP provides unprecedented control for recording calls, the "tagging" of calls with additional contextual information, and distribution of calls to the correct parties via e-mail if required. This greatly increases the ease with which organizations can act on and exploit their external communications with customers, suppliers, partners and staff.

Selective and / or bulk recording
eQuality ContactStore for IP can be configured to support the recording of calls on demand from the desktop, at the press of a button, or the recording of a selection of calls meeting pre-set parameters, or all calls.

Support for mixed traditional telephony and IP telephony environments
The eQuality ContactStore architecture can support mixed telephony environments within the same recording system, allowing the user to use the same application suite to search for, replay and evaluate calls irrespective of the platform or location at which the recording was made. The administrator can also manage and maintain both system types as a single entity.

Enterprise scalability
The ContactStore for IP recorder has a modular architecture such that all the components of a solution can be housed on a single server, but broken out into modular components to increase the scalability of the solution when required. This can include deploying multiple recorders at physically different locations and managing the system as a single entity to provide a single view of all customer interactions, wherever required, with minimal administrative overhead.

Support for screen data capture
In call center environments in particular, the ability of eQuality ContactStore for IP to capture and replay the screen data activity synchronously with the audio interaction can provide a more complete insight into the areas where agent skills and business processes can be improved.

Stereo recording and replay
Because ContactStore for IP is able to record and replay the two sides of a call separately or together, this improves the clarity of replay in mission critical emergency services and compliance environments, as well as enabling much greater accuracy of analysis by emerging speech recognition technologies.

Live monitor support
This feature, commonly required in contact centers, allows a supervisor or other user to listen in stereo to either side of a conversation, either locally or from a remote site.

Archiving
The ability to provide a secure and cost-effective archive for recorded interactions is often an imperative. ContactStore for IP allows customers to get maximum economies of scale from their corporate storage investments by archiving to any Windows file storage device.

Audit trail
eQuality ContactStore for IP provides the ability to audit which user has attempted to search for and retrieve which interactions, from where and when – an imperative in many sensitive applications and industries.

Cisco partnership

Witness Systems is a fully accredited AVVID and CCBU partner to Cisco Systems, and is the only recording vendor to have achieved this status. AVVID accreditation is rigorous and requires proof that the system itself is reliable and resilient, and will not impact on the mission critical telephony system. Furthermore accreditation requires that the vendor has the necessary skills, documentation and support to ensure successful deployments which can be made "TAC aware" within Cisco who can then provide end to end support.

Particular features which support the Cisco environment include:
Zero impact: Using 100% passive recording ensures that the ContactStore for IP recorder has no loading impact on the Cisco infrastructure and therefore no effect on any existing design or implementation.

No CTI required: Because ContactStore for IP decodes the Cisco SCCP Skinny protocol, a level of automatic call attribute "tagging" can be achieved "out of the box", which eliminates the cost of CTI and the associated integration required.

Resilience: Cisco have the ability to allow a remote site's telephony to continue to function in the event of a WAN failure, a feature known as Survivable Remote Site Telephony (SRST). Because ContactStore for IP decodes the Cisco skinny protocol and is not relying on a feed across the WAN, the recording can continue to function and maintain the same level of resilience as the telephony system itself.

Support for extension mobility: By decoding the Cisco Skinny protocol it is possible for the recording capability to follow an individual as they move around and log in to different extensions within the organization.

Award Winning: The XML applications on ContactStore for IP were voted to have the "Most Compelling RoI" of all the applications submitted to the global Cisco IP Telephony User Group.