

a White Paper from  
Vanguard Communications

# Which VoIP Architecture Makes Sense For Your Contact Center?

by  
Areg Gharakhanian

August 2002

Vanguard Communications Corporation  
100 American Road  
Morris Plains, New Jersey 07950  
+1 (973) 605-8000  
[www.vanguard.net](http://www.vanguard.net)



Entire contents © 2002 Vanguard Communications Corporation.  
All rights reserved. Storage or reproduction of this publication in  
any form or media without prior written permission is prohibited.

# Which VoIP Architecture Makes Sense For Your Contact Center?

By Areg Gharakhanian, Vanguard Communications

Convergence of voice and data is becoming a reality, and Voice over Internet Protocol (VoIP) is taking center stage. The early momentum centered on carrier-based VoIP services, with providers setting their sights on the \$100+ billion long distance services market. Now, the VoIP-based premises equipment market is heating up, with Cisco and other data networking companies staking claim to this territory. Traditional switch vendors have responded by shifting their product development focus away from proprietary systems to software-based solutions that run on converged architectures.

It's time for enterprises to determine *when* and *how* to transition their voice communications infrastructures to the next generation technology. A starting point for migration planning is a solid understanding of the architectural options. This paper describes traditional circuit-switched, IP-enabled, and IP-centric architectures, and addresses the pros and cons for each alternative.

## Life In The TDM World

A traditional Private Branch Exchange (PBX) or Automatic Call Distribution (ACD) system is comprised of a central processing unit running call processing software, a Time Division Multiplexing (TDM) switching matrix, network interface cards (e.g., analog, T1, PRI trunks), end user telephones (stations), and associated cabling. The TDM switching matrix is responsible for establishing connections between endpoints. It uses *circuit-switched* technology to allocate pre-defined bandwidth (typically 64 Kbps before compression) for the duration of each call.





Figure 1: TDM Single Site Deployment

User telephones, voice mail servers, and Interactive Voice Response (IVR) systems are connected to the PBX or ACD. A separate infrastructure supports *packet-switched* data connections to the local area network (LAN) and wide area network (WAN) to provide access to enterprise computing resources and the Internet (see Figure 1). Both the voice and data infrastructures require technical support to install, configure and service all of the endpoints, and administrative support to maintain the user database.

Figure 2 provides an example of a “virtual call center” configuration in which the headquarters ACD receives the bulk of the incoming calls and directs the call routing activities for all locations in the network. When qualified agents are available at headquarters, the ACD switches incoming calls to local stations. When calls need to be routed to either Site A or Site B, the headquarters ACD extends them out to the sites via dedicated, point-to-point voice connections, usually T-1 circuits. In this configuration, Site A is configured with a remote carrier capable of direct network access. Calls can be received directly into Site A (e.g., local calls), and routed to the appropriate agents as directed by the central processor located at Headquarters. Since the principal non-voice business applications reside on servers at the headquarters data center, the sites must also be configured with data access over a WAN.



**Figure 2: TDM Multi-site Deployment**

In a simple network with few sites, the duplication in voice and data infrastructure can be managed. But what about operations that support a large number of sites? Or what if all the locations receive incoming calls and then selectively route them to other sites on a peer-to-peer basis? This is another common multi-site configuration, frequently used when formerly autonomous centers with their own switches are networked into a virtual operation. In these scenarios, multi-site connectivity requirements are often implemented via dedicated point-to-

point circuits among the sites. The expense and complexity of this arrangement becomes significant with as few as four locations, and grows dramatically with more.

The pros and cons of the traditional TDM world are summarized below.

Advantages of TDM	Disadvantages of TDM
<ul style="list-style-type: none"> <li>• Proven and reliable technology.</li> <li>• Addresses call queuing and music-on-hold within the main chassis. Many TDM systems can provide announcements and collect caller entered digits using telephony cards.</li> <li>• Scales well for large call centers, especially for centers requiring call treatment for a large number of callers in queue.</li> </ul>	<ul style="list-style-type: none"> <li>• Requires significant investment in a single vendor's proprietary technology that is no longer the focal point of that vendor's R&amp;D efforts.</li> <li>• Requires separate infrastructures for voice and data with separate hardware and software components, and increasing administration and maintenance costs.</li> <li>• Increases connectivity expenses and complexities in multi-site configurations.</li> </ul>

### Life In The VoIP World

VoIP-based PBX and ACD systems convert voice conversations into packets and stream<sup>1</sup> them between endpoints using the Internet Protocol (IP). Voice IP packets traverse the data infrastructure using one of a variety of standard VoIP protocols (e.g., H.323, SIP, MGCP, Megaco/H.248) and vendor-specific protocols. *See Sidebar on page 8.*

The main challenge with VoIP is implementing the appropriate Quality of Service (QoS) to ensure that voice packets arrive at their destinations in an acceptable time frame (within 150 msec) and with minimum jitter (variation in inter-packet delay), so that users hear a continuous flow of speech. To accommodate these requirements, most enterprises use high-speed (10/100 Mbps) Ethernet switches to connect endpoints on their LANs. They also configure Virtual LANs (VLANs) to segment voice endpoints from data endpoints, thereby minimizing the potential for data packets to adversely impact voice packet delivery. In addition, the IEEE standard 802.1p/Q allows packet prioritization, giving voice frames higher priority than data frames.

The basic building blocks of VoIP-based PBXs and ACDs are gateway devices that convert calls from TDM to IP, call control processors, IP phones, and the associated data communications infrastructure (routers, Ethernet switches, and cabling). This basic architecture comes in two flavors: IP-enabled and IP-centric. The two are primarily differentiated by:


1. Where the TDM to IP conversion takes place
2. Where supporting call center functions, such as queuing, queue slots, prompting, music-on-hold, and announcements, take place.

<sup>1</sup> Voice packet transmission is accomplished using the Real time Transport Protocol (RTP) over IP. This protocol is focused on delivering packets to the distant end without retransmissions. It conserves time that might otherwise be consumed re-sequencing packets that were delivered out of order, or retransmitting corrupted data — both common occurrences in data packet transmission. The continuous transmission of voice packets is often referred to as *streaming*.

## IP-enabled Architecture

An IP-enabled system has essentially the same core architecture as a TDM system (switching matrix, common control, call processing software), with the added ability to service both TDM and IP phones via different station-side cards, and terminate IP network connections (see Figure 3). The CPU/common control and call control software may reside in the TDM switching matrix, or on an external server.

In an all-IP station deployment, the IP-enabled architecture leverages the enterprise's data infrastructure for both IP phone and PC connectivity. Most IP phones are equipped with a mini-Ethernet switch that connects the user's IP phone and PC to the Ethernet switch using a single port. The IP phone's mini-switch transmits packets with specific QoS and VLAN settings such that voice and data are given appropriate transmission priority in the network. Call queuing, recorded announcements, and music-on-hold are provided on the TDM switch. Since an IP-enabled architecture can service both TDM and IP phones and network connections, it provides a low risk transition plan for businesses that want to migrate users to IP slowly.



**Figure 3: IP-enabled Single Site Deployment**


In a multi-site environment, an IP-enabled architecture simplifies inter-site connectivity. The separate communications paths used for voice and data in a TDM world are combined into a consolidated data infrastructure.

Figure 4 shows a VoIP-enabled multi-site environment. Here, the headquarters location equipment provides all of the routing intelligence, and directs inbound calls to the appropriate location over the WAN. Note that in this example, Site B's equipment is stripped down to a bare minimum — a router, an Ethernet switch, and IP phones. Calls destined for Site B are terminated at the headquarters switch, converted to IP via gateway cards, and transported to Site B via the data infrastructure and WAN.

Site A can receive inbound calls from headquarters over the WAN (similar to Site B), or directly from the PSTN with remote call control performed by the headquarters IP-enabled switch. For most vendors, this set-up requires the addition of a media gateway at Site B. Some media gateways support TDM-based recorded announcements and touch-tone detection via resource

cards to allow for localized caller treatment in queue. This arrangement requires that the media gateway and IP-enabled headquarters switch be provided by the same vendor. Other vendor solutions handle caller treatment at the headquarters site by streaming audio signals associated with announcements or music over the WAN.

A WAN with adequate QoS performance is critical for a successful multi-site VoIP deployment. While the public Internet does not support QoS, most private WAN service providers offer service level agreements and performance guarantees. Some providers also allow users to prioritize their packets as they traverse the WAN by using standard QoS techniques (examples include 802.1 p/Q, DiffServ, MPLS, RSVP, and WFQ).




**Figure 4: IP-enabled Multi-site Deployment**

The pros and cons of IP-enabled solutions are summarized below.

Advantages of IP-enabled	Disadvantages of IP-enabled
<ul style="list-style-type: none"> <li>• Includes TDM components and software that are considered proven and reliable technology.</li> <li>• IP-enabled switch cabinet handles caller treatment efficiently with queue slots, announcements, music-on-hold, and caller entered digits using telephony cards.</li> <li>• Scales very well for large call centers that must support large number of callers in queue.</li> <li>• Begins migration to IP while leveraging existing investment and minimizing risk.</li> </ul>	<ul style="list-style-type: none"> <li>• Requires significant investment in a single vendor's proprietary technology that is no longer the focal point of that vendor's R&amp;D efforts.</li> <li>• Multi-site configurations may require additional proprietary components (e.g., media gateways), further increasing investment in a single vendor solution.</li> <li>• Mixed architecture includes TDM and IP based components with greater management complexity than IP-centric architectures.</li> </ul>

## IP-centric Architecture

IP-centric systems completely decompose the TDM based architecture. The TDM switching matrix is replaced with the enterprise's data infrastructure, and call control, call queuing, announcements, touch tone prompting, and music-on-hold functionality is provided via separate server processes that may or may not reside on separate physical servers (see Figure 5).

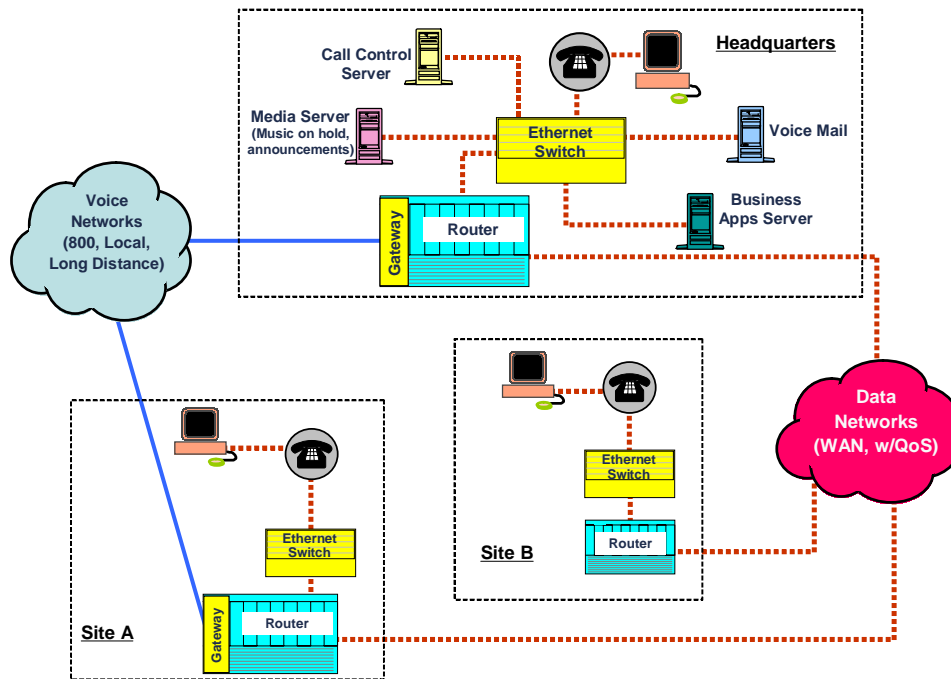


**Figure 5: IP-centric Single Site configuration**

In a single site deployment, a router blade acts as the media gateway. Thereafter, all call processing and switching is done via packets. Call control functionality is usually performed by a separate server, and may be configured in a cluster for redundancy. Call queuing, music-on-hold, announcements, and collection of caller entered digits are performed via a separate media server. Additional media servers may be required based on the number of simultaneous sessions/connections. Currently, each call in queue — whether listening to music or an announcement — is allocated a dedicated VoIP stream, which uses server CPU cycles. IP multicast techniques, which would allow multiple callers to listen to one resource simultaneously, have not yet been implemented.

An IP-centric system applies the same decomposition to a multi-site deployment. A centralized call control server controls VoIP streams, whether they arrive at the main facility or the remote sites (see Figure 6). In this deployment, the call control server at headquarters controls all inbound communications independent of location. Call queuing, music-on-hold, announcements, and prompting functionality are provided centrally via the headquarters media server.

PSTN calls inbound to the headquarters are converted from TDM to IP via the router blade. The call control server establishes a VoIP stream with the media server for queue treatment and prompting functionality, and then connects the call with the appropriate IP phone. Communication between the headquarters and the remote locations is similar to the IP-enabled solution.



**Figure 6: IP-centric Multi-Site configuration**

The pros and cons of IP-centric solutions are summarized below.

Advantages of IP-centric	Disadvantages of IP-centric
<ul style="list-style-type: none"> <li>• Provides a flexible and consistent architecture for call routing.</li> <li>• Completely leverages the data infrastructure, and minimizes investment in TDM components.</li> <li>• Provides greater flexibility in distributed environments, especially in configurations with local direct calling into multiple smaller sites.</li> <li>• Provides choice in technology element purchases, enabling a mixed vendor environment for routers, servers, software applications, etc.</li> </ul>	<ul style="list-style-type: none"> <li>• Decomposed architecture requires multiple server platforms.</li> <li>• Leading edge technology and newer code set are not “battle tested”.</li> <li>• Requires a highly reliable and robust data infrastructure.</li> <li>• Does not scale well today for providing queue treatment for large number of callers (e.g., music on hold, announcements, etc.).</li> <li>• Can be considered less reliable; may not be well suited for 24X7 call center operations.</li> </ul>

## Summary

VoIP will be the technology of the future for contact centers. The string of VoIP-based product announcements by traditional equipment vendors fully supports this prediction. Convergence of voice and data simplifies the infrastructure, and decouples the call control function by transitioning it to an application resident on an enterprise network. The benefits are magnified in a multi-site environment.<sup>2</sup>

IP-enabled and IP-centric approaches provide similar capabilities with some subtle differences. Although the evolution of these offerings may further blur the differences, buyers should be aware of the basic characteristics of each architecture to choose wisely the VoIP approach that best supports their business strategy.

A key deciding factor in determining when and how to move to VoIP is the migration strategy. IP-enabled approaches will appeal to the more conservative and heavily invested call center environment, while IP-centric methods will be the choice of smaller, newer, and more aggressive centers who can bear some reliability risk to move faster to a more advanced platform.

### **Sidebar: VoIP Standards**

VoIP standards are still evolving without a clear winner. Many vendors currently support or plan to support several of the standards. There are also a few proprietary approaches that claim better performance and functionality. Often vendor products are based on a mixture of proprietary and standards-based call control features.

The main VoIP standards today are:

**H.323** is the most mature and most widely implemented standard to date. H.323 was developed by the International Telecommunications Union (ITU) in 1996, and is now in its fourth revision (i.e., H.323 version 4). Microsoft's NetMeeting is based on H.323.

**SIP** (Session Initiation Protocol) is a more recent standard, developed by the Internet Engineering Task Force (IETF). SIP focuses on call initiation and termination between endpoints. It is a simpler standard than H.323, but not as feature rich. SIP has been adopted by some carriers; a SIP-based softphone is included in Microsoft's Windows XP.

**MGCP** (Media Gateway Control Protocol) is another recently developed standard by the IETF. MGCP was primarily defined to address how a softswitch (a software-based IP-switch generally sitting in the public network, such as an IP Central Office switch), can control gateway type equipment, including IP phones.

**Megaco/H.248** is a standard jointly developed by the ITU and IETF. Megaco/H.248 improves upon MGCP, and is intended to be its successor.

---

<sup>2</sup> See Vanguard's companion paper "Should Contact Centers Jump on the VoIP Bandwagon?" to understand possible applications and business benefits of migrating your call center to VoIP.

Vanguard Communications is the leading independent consulting company in the contact center industry. Founded in 1980, Vanguard helps clients meet business goals by planning, designing, and implementing solutions for customer contact and for managing customer and partner relationships.

Core services include strategy development, assessment and design, requirements planning, and implementation management, for both traditional call centers and emerging ecommerce technology. We work extensively with ACDs, IVR, CTI, Web access, network architecture, CRM and other software systems, and business process design.

Please contact us if you have any questions about this White Paper, or about how some of these ideas might apply to your organization. The author, Areg Gharakhanian, can be reached at [areg@vanguard.net](mailto:areg@vanguard.net).

And visit our website to read other articles and White Papers. There, you can also learn about our recently published book, *"Call Center Technology Demystified."*

Vanguard Communications Corp.  
100 American Road  
Morris Plains, New Jersey 07950  
+1 (973) 605-8000  
[www.vanguard.net](http://www.vanguard.net)