

Business to Business Videoconferencing

Achieving Secure Universal Connectivity in
the Digital Age

January 2011

Study sponsored by:



Table of Contents

Introduction	1
IP Networking Challenges for B2B Video	1
B2B Networking / Connectivity Options	2
ISDN Solutions.....	3
The Internet.....	3
QoS-enabled B2B Networks.....	6
B2B Deployment Options	7
The Enterprise Option.....	7
The Service Provider Option.....	8
Solution Spotlight: Providea's Universal Connectivity	8
Summary	11
About Wainhouse Research	12
About the Author(s).....	12
About Providea Conferencing	12

Table of Figures

Figure 1: The Providea Universal Connectivity Platform.....	9
---	---

Introduction

The use of personal and conference room videoconferencing, as well as fully integrated immersive telepresence suites, is growing rapidly. Businesses of all sizes are finding these solutions to be invaluable tools for enhancing productivity; reducing the time, cost, and stress associated with business travel; deploying business continuity strategies; and increasing the speed of business decision making in general. Today, many enterprises are looking to use videoconferencing to strengthen partnerships and to communicate more effectively with customers and suppliers. These savvy companies are looking for ways to use videoconferencing for mission-critical, business-to-business (B2B) communications.

At its outset over twenty years ago, B2B videoconferencing was straightforward because video systems communicated using ISDN (Integrated Services Digital Network), a service that operates over the public switched telephone network (PSTN). Using the PSTN provides access to a well-established, common addressing scheme (known to the world as telephone numbers) and allows enterprises to interconnect and participate in B2B video calls. However, ISDN videoconferencing, whether used for internal-only or external (B2B) sessions, often delivers poor reliability and a lower quality experience limited by bandwidth practicalities. In addition, ISDN services are sold on a metered, “cost per minute of use” basis; a business model that discouraged the use of videoconferencing, and especially high quality calls / HD calls.

Videoconferencing over IP networks has gained considerable traction over the past decade to the point that IP-compatible hardware now dominates the market, and IP-based video calls now account for the majority of usage. Videoconferencing over IP addresses ISDN's primary shortcomings by providing a superior user experience and higher reliability. In addition, most IP services are sold for a fixed monthly fee. More importantly, since IP is the foundation of the enterprise network (as well as the Internet), deploying videoconferencing over IP allows organizations to leverage their existing network and integrate videoconferencing with other enterprise applications and the overall workflow process. However, IP videoconferencing, and specifically B2B IP videoconferencing, also introduces a new set of challenges.

Authors Note: Although this paper focuses on the challenges associated with B2B (business to business) video communication, in many cases the same problems impact internal communications between offices (or video systems) using different network carriers.

IP Networking Challenges for B2B Video

For IP-based videoconferencing, the key B2B challenges include the need to:

- Traverse corporate firewalls without compromising security - video calling across the firewall is problematic because of the way in which data connections are opened and closed. Firewalls are asymmetric; they have a trusted and an untrusted side. Videoconferencing is symmetric; it requires bidirectional video and audio stream connections that must cross the firewall in order for the video call to be successful.

Many firewall suppliers have upgraded their devices to account for voice and videoconferencing protocols, but not all customer firewalls in the field have been upgraded. In addition, NAT / firewall traversal solutions are available from a range of others including many videoconferencing vendors.

- Provide adequate end-to-end IP bandwidth and throughput – unlike non-real time data traffic (e.g. web browsing, email), the user experience during a videoconference is severely impacted by bandwidth capacity and performance issues (e.g. packet loss, delay, jitter). In general, there are three ways to overcome this problem:
 - Overprovision the networks to avoid contention and packet loss
 - Implement QoS / prioritization technologies to maximize the throughput for videoconferencing traffic
 - Utilize videoconferencing solutions / protocols that are less sensitive to network issues (e.g. SVC) or provide some form of error correction / concealment (e.g. FEC). This method is less attractive to enterprises with an existing deployment of standards-based video systems as it may require the purchase and deployment of all new videoconferencing endpoints and infrastructure.

This situation is compounded by the inability to peer private MPLS networks between carriers without the loss of QoS.

- Deploy compatible addressing schemes – in order to make a connection between two devices (telephones, video systems, etc.), the addresses of these devices must be unique. If the connections are internal only (e.g. phone calls from internal extension to internal extension, or data connections from one internal device to another), the addresses need to be unique within the internal environment only. However, in order to support connections between internal and external devices, the addresses between all involved environments must be unique. Unfortunately, no truly global addressing scheme for IP videoconferencing exists today. Instead, addressing schemes are typically managed at the enterprise level.

B2B Networking / Connectivity Options

To deploy B2B videoconferencing, all enterprises wishing to communicate must first make two decisions:

- 1) What network(s) they will use to interconnect – at least today, the choice of which networks an enterprise uses for its B2B video traffic may determine its ability to reach other enterprises. This is likely to be resolved in the future as more B2B exchange providers and network service providers interconnect their exchanges and networks; WR has already observed quite a bit of activity in this area.
- 2) What technology solutions they will use to connect their local networks to the interconnect network

The wide range of networking / connectivity options available to organizations seeking to conduct B2B videoconferencing can be organized into three general categories:

- ISDN networks
- The Internet
- QoS-enabled B2B networks

ISDN Solutions

ISDN (Integrated Services Digital Network) is a PSTN-enabled network technology that provides bandwidth greater than that of a single voice circuit. Because ISDN is a public service with publicly assigned E.164 addresses (E.164 is the international numbering plan for public telephone systems), ISDN avoids the addressing issues common in today's B2B videoconferencing environment. In addition, ISDN is totally separate from the corporate data network, thereby eliminating hacker-based security concerns.

Enterprise videoconferencing units may be connected to an ISDN circuit directly, through a PBX, or through an ISDN switch, but a more common method is to use an IP-to-ISDN gateway to connect IP-based video systems on the corporate LAN to the ISDN network. The gateway may be located on the enterprise IP network directly, or it may be provided by a service provider. This configuration is similar to a common VoIP situation in which gateways are used to connect IP-based PBXs and telephony systems to the PSTN, thereby enabling B2B voice calls.

The Internet

The Internet provides IP connectivity between almost all businesses today. In cases where sufficient capacity is available and security constraints do not block connectivity, the Internet can be used to conduct high quality videoconferencing. In fact, the Internet is used by companies, both large and small, every day to conduct B2B IP video calls. In some cases, however, using the Internet to host IP video calls is less than ideal due to the following significant limitations:

- Lack of guaranteed bandwidth - in general, the bandwidth available is highly dependent on each party's access link to the Internet. In addition, despite its exceptionally high capacity, the Internet backbone itself may be subject to uncontrollable congestion at any point in time. In addition, in many cases Internet bandwidth is shared by many subscribers, and as a result capacity and performance can fluctuate widely.
- Lack of Quality of Service (QoS) support – the Internet lacks the ability to prioritize specific streams of data traffic over other streams. This means that real-time voice and video traffic may suffer from long latency and packet loss, resulting in a compromised user experience.

- Lack of global video conferencing addressing scheme - a common method for connecting devices (like video endpoints) to the Internet is to use a NAT router. The advantage of using NAT is that numerous devices with private IP addresses can share a single public IP address. The disadvantage, however, is that without some type of address translation device (e.g. a gatekeeper, SBC), the systems will be unreachable from the public Internet.
- Firewall challenges - firewalls are designed to block unauthorized incoming data traffic. When Enterprise A places an IP video call to Enterprise B, the firewall at Enterprise B may block the connection since the firewall opens connections that are initiated on the trusted side (inside), but blocks connections initiated on the untrusted side (Internet side). For video calls, of course, calls coming from an external company would appear as calls from an untrusted source and would, by default, be blocked by the firewall - unless a video-friendly firewall or some other NAT / firewall traversal method is used.

There are multiple methods for enterprises to connect to the Internet for B2B videoconferencing, including:

- Connecting directly to the Internet - this relatively unsecure approach is often used by mobile users communicating from hotels or public Wi-Fi hotspots. It is also used by some small offices with single systems, and in cases where the enterprise Internet connection is totally separate from the enterprise data network. To be able to receive calls, the videoconferencing system should have a fixed, public IP address or be registered to a SIP server or H.323 gatekeeper that will bind the IP address to a higher level, fixed address.
- Placing an endpoint in the network DMZ - in this case, the video endpoint is connected to either the DMZ port of the firewall directly, or the edge router assigns the internal IP address of the endpoint to be the DMZ device and thus passes all incoming external traffic directly to the videoconferencing endpoint.
- Connecting through the firewall - with this approach, the videoconferencing system(s) resides on the LAN and behind the enterprise firewall, and the firewall must be configured to:
 - allow the video traffic from the video system to traverse the firewall and go out to the Internet, and
 - permit incoming traffic directed to the video system to pass through the firewall and make its way to the correct internal endpoint.

In order to support incoming video calls from external systems, a public IP address must be assigned to each videoconferencing unit. The IP address is statically mapped through the firewall to connect the external IP address to the internal IP address of the videoconferencing endpoint.

- Deploying a multipoint control unit (a.k.a. MCU or video bridge) with connectivity both inside and outside the firewall - an MCU with two IP media ports can be used to successfully cross the firewall border. The MCU is deployed with one media port connected to the enterprise network (inside the firewall), and a second port connected to the Internet or the firewall DMZ. The result is a bridge that can be reached either from the enterprise network or from the Internet, with neither IP connection having to cross the NAT portion of the firewall. In operation, users see this solution as a “meet-me” bridge with users on the enterprise network dialing out and poking a hole in the firewall. No other firewall traversal technique is needed, and outside callers never gain access the private network.

- Deploying a session border controller (SBC) - an SBC is a voice and video-specific device designed to enable secure network-to-network traversal. The SBC has two ports - a WAN-side port with a public address, and a LAN-side port with an internal address. The SBC manages address translation and QoS marking as traffic enters and leaves the enterprise network. It also decodes the H.323 and SIP signaling protocols, provides network address translation, and acts as a secure gateway to allow video (and voice) traffic in and out of the enterprise network without compromising network security. The SBC only passes traffic that, a) meets the strict specifications of the protocol, and b) is targeted to devices of the right type within the organization. Business partners that connect frequently can connect their SBCs together through ‘neighboring.’ This capability allows two or more SBCs to share directory information.

- Deploying an H.460 traversal server - H.460 is an ITU industry standard that enables H.323 videoconferencing signaling and media streams to cross the firewall using several well-defined, outbound TCP and UDP ports. H.460 requires two elements:
 - 1) Client software operating behind the firewall. In many cases, this software is embedded within the videoconferencing endpoints. In other cases, a proxy server (e.g. a video gatekeeper) provides the client function for many endpoints on the network.

 - 2) An H.460 traversal server located outside the firewall in a publicly addressable location. Endpoints tunnel through the firewall, register with the H.460 traversal server, and maintain an open channel to the server to exchange signaling messages. The external traversal server (Internet side) acts as a relay point for video calls from inside the firewall, and creates video tunnels between the internal endpoints and the external server using the dedicated firewall ports defined by the protocol. These ports can be safely opened through the firewall by limiting access to just the IP address of the external H.460 server.

While all of the Internet-based options above are potentially viable, each method poses a range of drawbacks including:

- Lack of security due to need to open many network ports and/or put systems in the DMZ
- Need to provide public IP addresses for each system
- Need to purchase, deploy, and manage SIP server, gateway, SBC, or other registration system
- Need to deploy a video-friendly firewall or an H.460-type NAT / firewall traversal solution
- Need to deploy an expensive video bridge / MCU and utilize MCU ports for every B2B connection

The net is that the Internet can be a viable transport medium for both internal and external (B2B) videoconferencing. However, for a variety of reasons (e.g. lack of guaranteed bandwidth, lack of QoS, lack of global addressing scheme, etc.), the Internet may not be the best B2B video choice for some enterprises in certain situations.

QoS-enabled B2B Networks

Unlike the Internet, QoS-enabled networks recognize QoS class markings and treat high-priority streams with the right forwarding behavior to achieve low packet loss, low latency, and minimal jitter. The result is a more consistent visual collaboration experience. QoS-enabled B2B networks today are either dedicated IP private video networks or VPN-to-VPN connections.

Dedicated IP private video networks interconnect multiple enterprises via a direct connection between each enterprise and the interconnecting network. Private video networks are similar to enterprise IP overlay networks in that they are deployed to support only certain kinds of data, in this case videoconferencing traffic. Many IP overlay networks have been deployed over the last few years to support high definition videoconferencing and multi-codec / telepresence deployments because the enterprise data network lacked the capacity or performance necessary to host the visual collaboration traffic. Unlike overlay networks which tend to carry traffic within one enterprise, private video networks allow more than one company to connect. Once connected to the private video network, enterprises can participate in guaranteed bandwidth, high QoS videoconference sessions with other endpoints / enterprises connected to the same private video network. This approach is best suited for groups of organizations that, a) are already working together, b) are already leveraging the same network service provider, and c) have already resolved any potential NAT / firewall traversal issues. In other situations, this approach can be expensive and difficult to scale.

VPN-to-VPN connections allow the MPLS VPN of one enterprise to connect to the MPLS VPN of another enterprise without violating network security. Depending upon the situation and specific requirements, the connection may be direct between the two networks or may be realized via an intervening network or single connection point. VPN-to-VPN capabilities are generally delivered by external service providers and include the ability to interconnect

enterprise MPLS VPNs using an external QoS-supported network that provides a security boundary, address translation, and QoS mapping between enterprises. The key benefit of this approach is that it allows multiple enterprises to communicate without sacrificing security. However, the nuances of these types of interconnections can be quite complicated, which is why such an arrangement is best handled by a service provider.

B2B Deployment Options

Any enterprise wishing to support B2B connectivity for its IP-based videoconferencing systems or telepresence suites has two basic options:

- 1) The Enterprise Option – under this approach, the enterprise owns, controls, and manages the necessary infrastructure itself, or
- 2) The Service Provider Option – using this method, the enterprise relies on a B2B service provider to handle the details of B2B interconnections.

The Enterprise Option

Going the “do-it-yourself” route, no matter whether the infrastructure is located on the customer premise or in a co-location facility, enables the enterprise to have complete control over the access method involved and all call permissioning policies. In addition, this approach gives the enterprise the option of using upfront investment capital to reduce the on-going usage fees.

The enterprise option, however, has three significant disadvantages.

- 1) The enterprise must purchase and manage complex equipment that requires an understanding of IP protocols, videoconferencing standards and dialing practices, video call set-up procedures, and other technical challenges. This burden is somewhat ongoing as many of the components will require software upgrades and enhancements to remain current because B2B technology is changing rapidly.
- 2) The enterprise will need to deploy and manage multiple network connections since B2B partners are likely to be on different networks. In addition to the IP networks, the enterprise will likely need to deploy and maintain ISDN bandwidth for B2B calls where the IP networks cannot connect securely and for calls with enterprises supporting only ISDN video. Furthermore, as B2B conferencing becomes more common, the enterprise will need to deploy more bandwidth to support the greater call volume, higher call speeds, multi-codec / telepresence sessions, and multipoint video calls.
- 3) The enterprise will need to hire and maintain a highly skilled support staff. These specialists will need to understand videoconferencing endpoints and infrastructure products, as well as IP security and networking technology. The support load on these resources will also increase over time as B2B complexity increases and as B2B usage and the number of B2B partners grow.

The Service Provider Option

Communications has always been a service-provider-oriented business. Whether enterprises are using mobile or wired voice systems, room videoconferencing systems, or a network of telepresence suites, savvy customers have looked to service providers to provide the security, reliability, and quality of experience that users expect. Now, with the growing interest in B2B conferencing, enterprises rightly see service providers as key B2B enablers. Leveraging a service provider for B2B connectivity provides several significant benefits.

- Easing the burden of managing video - B2B service providers solve complex IP connectivity issues as their core business proposition and their core competency. They have invested in highly skilled, specialized central resources to provide 24 x 7 support to multiple clients and to isolate customers from issues around staff turnover, holidays, sick leave, training costs, etc.
- Maximizing ROI – leveraging a service provider experienced in installing, debugging, and delivering inter-enterprise videoconferencing services can drive additional usage, minimize the need for capital investments, and shorten the “time-to-benefits.” All of these factors lead to higher ROI for the customer.
- Improved cost control – leveraging a service provider can improve budgeting and cost control by providing access to fixed price agreements independent of usage.

Solution Spotlight: Providea’s Universal Connectivity

Providea, the sponsor of this white paper, has introduced a B2B videoconferencing service intended to make it fast and easy for customers to conduct inter-enterprise video and telepresence calls - regardless of the networks or equipment deployed by the participants.

The center of the Universal Connectivity Platform (referred to as “the Platform” for the remainder of this document) is an array of videoconferencing bridges and other video infrastructure devices (see diagram below). These items are located in a secure network space accessible only to Universal Connectivity customers.

To connect to the B2B service, customers simply add the Platform to their existing wide area networks. Since the Platform already has connections in place with more than 150 network service providers, this is a relatively fast, inexpensive, and pain free process. As shown below, in some cases (see Carrier A) a dedicated cross-connection¹ from the carrier’s cloud to the Platform will be required. In other cases (see Carrier B), the connection can be made by partitioning off a portion of an existing cross-connection between the carrier and the Platform.

¹ Note that a cross-connection is NOT the same as a local loop. Within this context, a cross-connection is a connection made from one communication rack to another rack within the same co-location (co-lo) facility. A local loop connection involves a connection from one physical location / co-lo to another. Local loop connections are almost always more expensive and require more time to deploy than cross-connections.

This method of partitioning an existing cross-connection is especially convenient as it requires only a carrier software update to provide the customer with connectivity.²

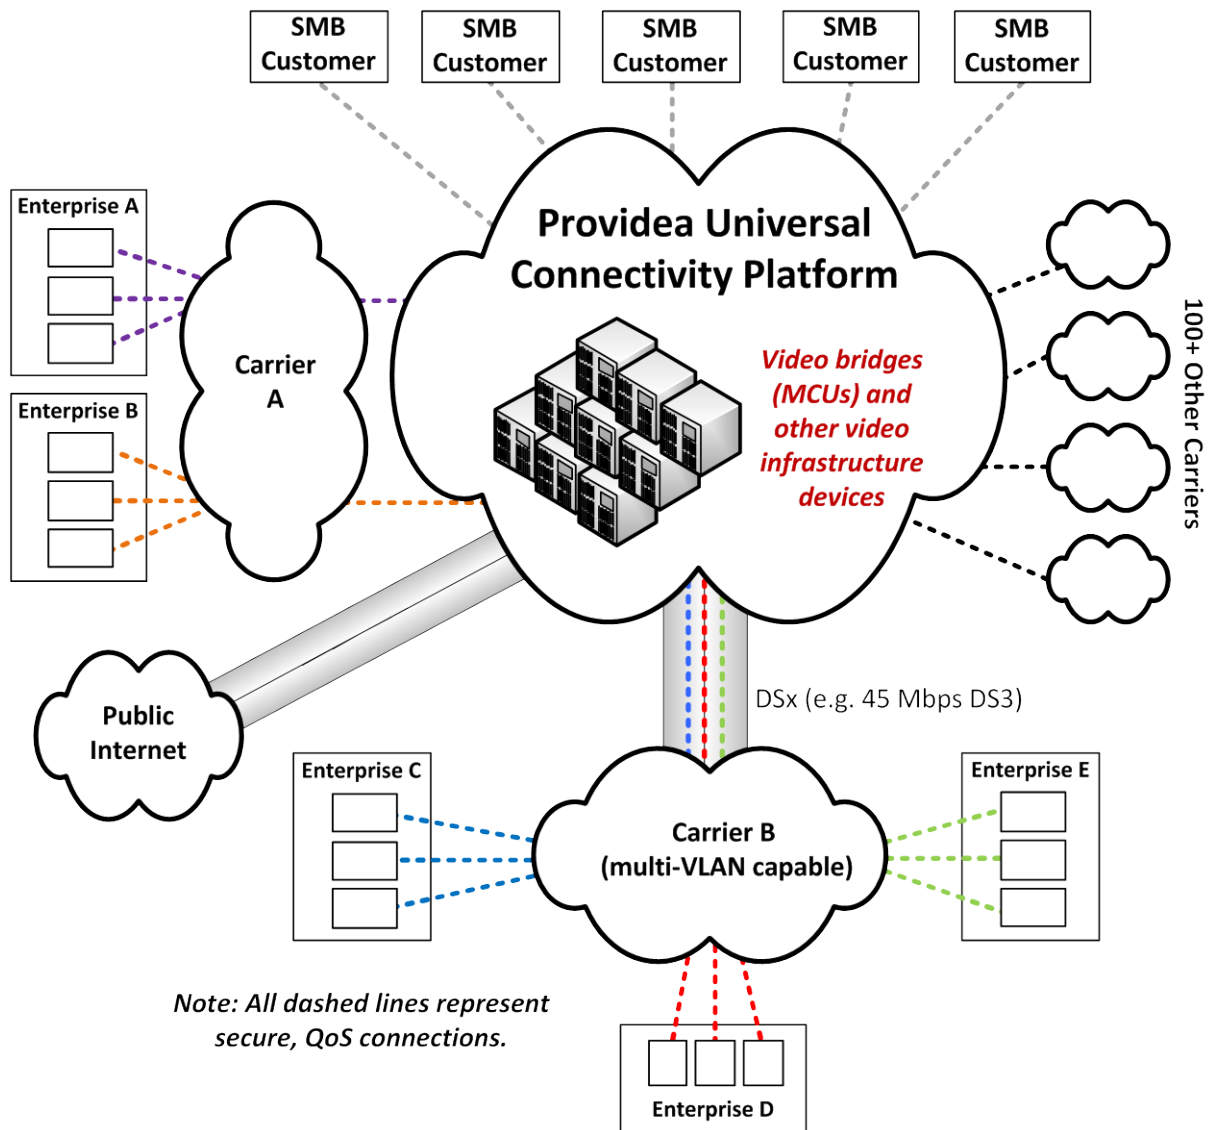


Figure 1: The Providea Universal Connectivity Platform

Once connected, each customer has access to the centrally deployed video bridges and infrastructure devices, and thus can conduct multipoint and B2B video calls with other customers connected to the Universal Connectivity Platform. This includes customers using, a) other private networks connected to the service, b) other accessible B2B exchange networks, c) the public internet, and d) ISDN for videoconferencing.

² In order to support this connection method, the carrier must support the dynamic allocation of bandwidth through a single network port into multiple customer-specific VLANs. As of this writing, only a handful of NSPs support this capability.

To maintain the privacy and security of each customer's data network, the customer connections to the Platform are configured to only allow traffic to flow between the customer's network and the centrally deployed video bridges and infrastructure devices within the Platform. In other words, traffic is NOT permitted to pass from one customer's network to another. If requested, a customer-specific firewall can be installed between the customer's network and the Platform network.

The Universal Connectivity Platform provides a number of key benefits:

- Access to reachable³ video systems on any of more than 600 global carriers
- Fast, low cost activation (typically requires only a cross-connection to gain access)
- Low monthly bandwidth costs (by leveraging existing bandwidth in place without the need for additional local loops)
- Access to Providea's video bridging service
- A reliable, high quality videoconferencing experience ...
 - via the use of redundant connections
 - by maintaining QoS tags when video traffic is routed back to customer's network
 - by over-provisioning the Platform's network to avoid congestion
- Maintains network security for each customer by ...
 - the use of VLANs and security policies to restrict access to other networks
 - the routing of all traffic through the MCUs (via static trusts)
 - the optional deployment of customer-specific firewalls
- Easy and cost-effective scalability by leveraging existing bandwidth instead of requiring customer-specific local loops
- Ability to leverage customer-owned infrastructure deployed within the customer premise of within Providea's video hosting center

Universal Connectivity customers also receive 24/7 support from Providea's video network operations center (VNOC); a video-centric help desk staffed by people with years of experience deploying and managing videoconferencing and telepresence solutions.

In addition to the above architecture, Universal Connectivity customers can also leverage their own videoconferencing infrastructure (MCUs, gateways, etc.).

³ Within this context, "reachable" video systems are those able to participate in B2B calls.

Summary

Enterprises that have invested in videoconferencing endpoints and infrastructure for internal communication will naturally want to use those tools for business-to-business communication as well. B2B videoconferencing greatly enhances the value of any enterprise's investment in visual collaboration tools.

At least today, B2B videoconferencing is a capability best handled by service providers rather than by the individual enterprise. By relying on video services specialists, the enterprise user can offload very technical video and network challenges to a partner for whom this is his core competence.

Independent B2B video exchanges can provide the network-to-network connectivity and the video interoperability that end user customers need, while maintaining the reliability and security that network managers demand.

The Universal Connectivity offering from Providea Conferencing, LLC, the sponsor of this white paper, is an innovative solution to the B2B challenge. Universal Connectivity eliminates hardware and network dependencies by providing an "MCU in the cloud" (using either Providea-owned or customer-owned equipment) for B2B, meet-me videoconferencing. Customers of this solution can enjoy reliable, high quality videoconferencing with companies on a wide range of carrier networks. In addition, the Universal Connectivity is well suited for customers needing to conduct frequent B2B calls with a handful of longstanding partners, as well as customers needing to conduct B2B calls on short notice with a wider range of organizations. Especially when coupled with Providea's extensive video managed service offering, the Universal Connectivity solution is worthy of consideration by any organization, large or small, interested in conducting high quality, business-to-business videoconferencing.

About Wainhouse Research

Wainhouse Research, LLC (WR) provides analysis and consulting on the market trends, technologies/ products, vendors, applications, and services in the collaboration and conferencing fields. Areas of coverage include hardware, software, and services related to audio, video, and web conferencing, unified communications, and enterprise social networking. The Company publishes market intelligence reports, provides customized strategic and tactical consulting and studies, and produces industry events (conferences). Additionally, the Company operates industry-focused and end user-focused Web sites and publishes a weekly sponsored bulletin for news and analysis. For more information on Wainhouse Research, visit www.wainhouse.com.

About the Author(s)

Ira M. Weinstein is a Senior Analyst and Partner at Wainhouse Research, and a 20-year veteran of the conferencing, collaboration and audio-visual industries. Prior to joining Wainhouse Research, Ira was the VP of Marketing and Business Development at IVCi, managed a technology consulting company, and ran the global conferencing department for a Fortune 50 investment bank. Ira's current focus includes IP video conferencing, network service providers, global management systems, scheduling and automation platforms, ROI and technology justification programs, and audio-visual integration. Mr. Weinstein holds a B.S. in Engineering from Lehigh University and can be reached at iweinstein@wainhouse.com.

Andrew W. Davis is a researcher, analyst, and opinion leader in the field of collaboration and conferencing. He is a co-founder of Wainhouse Research. Prior to Wainhouse Research, he held senior marketing positions with several large and small high-technology companies. Andrew has published over 250 trade journal articles and opinion columns on multimedia communications, videoconferencing, and corporate strategies as well as numerous market research reports and is the principal editor of the conferencing industry's leading newsletter, The Wainhouse Research Bulletin. A well-known industry guest speaker, Mr. Davis holds B.S. and M.S. degrees in engineering from Cornell University and a Masters of Business Administration from Harvard University and can be reached at andrewwd@wainhouse.com.

About Providea Conferencing

(Copy provided by Providea)

Headquartered in Camarillo, CA with offices nationwide, Providea Conferencing offers a full suite of Telepresence and high-definition video conferencing solutions, along with comprehensive system and network design consultation and world-class support. Founded in 1999, the Providea team boasts years of experience in video conferencing, professional services, multi-media design/build, multi-point bridging, network configuration and the best practices for successful deployment and utilization of these technologies. Leading organizations within financial, legal, technology, healthcare, entertainment, education and government sectors rely on Providea Conferencing. For more information, please visit <http://www.provideasolutions.com>.