

WHITE PAPER

PacketAssist™

An Architectural Approach to Delivering
Quality of Service (QoS) for IP Video



VCON

1 CLICK
AWAY
From Being There.

TABLE OF CONTENTS

Overview	3
Internet Protocol and the H.323 Standard	4
VCON's PacketAssist Architecture	4
Bandwidth Management Issues	5
Bandwidth Management Solutions	6
Adaptive Bandwidth Adjustment	7
Packet Loss Recovery	7
Quality of Service Issues	8
Quality of Service Solutions	8
End-to-End QoS Architecture	12
Summary	12
Contact Us	13

OVERVIEW

The market's deployment of IP-based voice and video applications is well past the early adopter stage and now into the deployment phase in most geographies around the world. One of the driving forces behind this adoption is the opportunity for convergence at both the network and application levels. Convergence brings about exciting new applications that are driving customer demand, while IP brings ubiquitous transport over customers' existing data network infrastructure.

Convergence continues to serve as a significant growth driver for the whole networking industry today. By consolidating voice, video and data applications onto a single "rich media" network, enterprise customers can create an environment that saves costs because it is easy to manage and administer. According to industry analyst Christine Perey, President of Perey Research and Consulting, "Rich media conferencing and interactive collaboration is the goal towards which all convergence is aiming. With widespread upgrades to the infrastructure we are crossing an important frontier so that we can support voice, video and data in a secure yet manageable environment. Network managers worldwide are joining forces with the telecommunications world and making significant progress towards a seamless integration. When we can leverage a common network, and yet distinguish traffic and prioritize time sensitive data types, everyone will experience substantial cost savings as well as deploy new multimedia-rich forms of communications."

This white paper describes both the key issues and the solutions associated with deploying video applications on an IP network. As such, it is targeted primarily at the network manager either in the consideration or implementation stages of an IP video deployment. As an early innovator and technology leader in IP-based videoconferencing, VCON has learned a great deal about deploying video applications on an IP network. We have gained much experience from successful deployments. This white paper breaks down the various deployment issues by describing how they apply to video applications and how the network manager can resolve them. But like many network design and deployment issues, there is no single solution that addresses all network topologies and configurations. Therefore, the issues and solutions described in this white paper should be treated as guidelines or design considerations. Ultimately, it is up to the network manager to make the final decision about how to best meet the service level commitments to his/her end-user constituency.

INTERNET PROTOCOL AND THE H.323 STANDARD

The most common protocol for network convergence is, without a doubt, the Internet Protocol (IP). It offers the best combination of interoperability at the LAN level and connectivity over the WAN.

The International Telecommunications Union (ITU) approved the H.323 standard in 1996. H.323 is the globally accepted standard for audio, video, and data communication over packet-based networks. It specifically describes how multimedia communications occur between user terminals, network equipment, and assorted services on Local and Wide Area Internet Protocol (IP) networks. Another protocol, SIP, is sponsored by the IETF and offers an increasingly viable solution for IP video deployments. While most of this white paper makes specific mention of H.323, the concepts and solutions apply equally to both protocols.

IP-based video applications vary in their tolerance of “best effort” network conditions. The video application that network managers should concern themselves with when designing new networks, is interactive video. This is not to say that streaming video applications (like those that utilize unicast or IP multicast) don’t have their own set of design considerations. Rather, the real-time nature of interactive video (namely videoconferencing) imposes an additional set of challenges that network managers have not had to face prior to the introduction of real-time services.

However, videoconferencing is not the only real-time application driving this demand. Many network managers today are either investigating or deploying voice services over their IP network. Regardless of which application will be deployed first in a given enterprise, the network design considerations are very similar. In other words, if the network is going to be readied for one of these real-time applications, it might as well be readied for both.

There are many different network topologies that can be used to carry interactive videoconferencing traffic. The scope of this white paper is specifically those topologies and transports for IP – both LAN and WAN transports.

VCON’S PACKETASSIST ARCHITECTURE

Committed to improving the quality and manageability of IP video, VCON developed PacketAssist – an architectural approach to delivering Quality of Service (QoS) for interactive video over IP networks. PacketAssist is a technology firmly embedded into VCON’s various application solutions. The PacketAssist Architecture has two primary design objectives:

- Bandwidth Management – ensuring that interactive video applications do not unnecessarily saturate the network
- Quality of Service (QoS) – providing the best possible audio and video quality to the end user at a given data rate.

Research indicates that these two issues are the most significant inhibitors to the deployment of video applications over an IP network. The remainder of this white paper focuses specifically on these two deployment issues.

BANDWIDTH MANAGEMENT ISSUES

Adding interactive or streaming video to the network can sometimes be perceived as a threat to a network manager's mission. Their concern is that video traffic may unnecessarily saturate network capacity and adversely affect service delivery for other mission-critical applications.

How much bandwidth does an interactive video application take? The answer lies primarily in the audio and video quality required by the end user. Many videoconferencing solutions offer support for a wide range of data rates. In fact, VCON leads the industry in video quality by offering data rates up to 4 Mbps. These data rates describe the traffic in each direction of a conversation between the videoconferencing participants. Therefore, for a point-to-point conference the actual demand on the network is twice the data rate of the conference, plus some overhead (10-15% as a rule of thumb). This means that a typical 384Kbps conference will demand approximately 850-900Kbps on the network ($384 \times 2 + \text{overhead}$) with full motion by the participants. During times when there is low motion by the participants, the bandwidth demand on the network is typically less. While the data rates mentioned might not pose a significant network design change at the LAN level, it almost certainly has implications at the WAN level where bandwidth is typically more costly, and therefore more fully utilized on average to begin with.

Of course, the instinct of the typical end user is to conduct their conferences at the highest possible data rate to get the best possible audio and video quality. If enough bandwidth is available, in fact, this will be the typical use in order to reproduce a face-to-face meeting. However, if the conference data rate causes the network to become saturated enough to produce packet loss and excessive network jitter, the user would get a more satisfying experience by dropping down to a lower conference data rate in which packet loss and jitter are eliminated. Said a different way, a 768Kbps videoconference without packet loss and jitter will almost certainly have better audio and video quality than a 1.5Mbps conference with packet loss and jitter.

BANDWIDTH MANAGEMENT SOLUTIONS

In response to the above concerns, the data networking industry had developed a plethora of bandwidth management solutions. While applying additional network capacity (bandwidth) at a problem network segment certainly helps alleviate most of the bandwidth-related issues, it can get very costly – especially in the WAN. Furthermore, it does not address most of the QoS issues described later in this paper. Four complementary approaches to solving the bandwidth constraints include:

- Upgrading to switched LAN segments
- Upgrading from 10 to 100Mbps LANs
- Implementing gatekeeper functions
- Deploying systems with VCON PacketAssist's adaptive bandwidth adjustment feature.

Taking the next step of also upgrading to 100Mbps or Gigabit subnet segments is truly a function of overall segment utilization, which can easily be tested and monitored. Many customers find that switched 100Mbps to the desk and Gigabit backbone links between the switches is very suitable. Once again, this exercise is a very standard practice for most network managers, and it should be driven by overall segment utilization.


The H.323 standard defines an important component called a Gatekeeper. The gatekeeper provides many management functions for an H.323 network, including bandwidth management, access control, address translation, and registration. From a bandwidth management perspective, most gatekeepers allow bandwidth to be controlled in the following policy level settings:

- Limit the maximum conference data rate by a user
- Limit the maximum conference data rate by a group of users or a network segment
- Limit the aggregate conference data rate to a stated percent of total LAN bandwidth
- Limit the number of H.323 terminals allowed to simultaneously use services on the network.

In conjunction with network analysis and monitoring, the gatekeeper can be a critical component for managing bandwidth for H.323 conferences. Careful evaluation of the bandwidth management features of the various vendors' gatekeepers will prove valuable in the long run. Additionally, since many gatekeeper functions are optional enhancements to those required in the H.323 specification, it is important that the videoconferencing endpoint can exploit these advanced bandwidth management features.

ADAPTIVE BANDWIDTH ADJUSTMENT

Wouldn't it be nice if the endpoints had enough intelligence to manage bandwidth for themselves? With VCON PacketAssist, they can. One important feature of VCON's PacketAssist Architecture is called Adaptive Bandwidth Adjustment (ABA). Just as the name implies, adaptive bandwidth adjustment enables the endpoint to automatically adapt to the network's capacity and performance after detecting unfavorable or highly favorable conditions. As such, it dynamically adjusts the video data rate during the conference in an effort to provide the best possible video quality to the end user without saturating the network. The audio channel stays fixed at a given data rate to provide consistent high quality audio throughout the conference. Also, ABA is intelligent enough to control both the inbound and outbound data streams – even if one of the endpoints is not a VCON endpoint.



PACKET LOSS RECOVERY

Although VCON's ABA automatically adapts the data rate to the network's capacity, its down speeding mechanism alone may not sufficiently reduce packet loss on WANs with bursty or very erratic traffic, such as the Internet, or on networks with temporary equipment problems. To counter this condition, VCON devised a Packet Loss Recovery feature.

Packet Loss Recovery uses redundant information in data packets to detect and correct errors on the receiving side with no need to retransmit the lost packets. It dynamically adjusts the transmission of data according to various aspects of the network state, including network behavior and the percentage of packet loss.

On networks in which ABA is insufficient, the VCON end point automatically triggers the Packet Loss Recovery feature, which succeeds in recovering lost information from video and audio packets. This contrasts to other similar methods in the marketplace, which work on recovering only video streams, but not audio.

QUALITY OF SERVICE ISSUES

Quality of Service is described in the CCITT's Recommendation E.800 as "the collective effect of service performance, which determines the degree of satisfaction of a user of the service". This is a very subjective and user-centric definition that leaves the determination of "quality" to the user's judgment. Once interactive video applications are deployed to the end users, they will expect the best possible audio and video quality. To the end-user, this means good resolution video with smooth motion along with rich audio that is easily audible and without echo. No problem, right? Well, as it turns out, IP network characteristics that are not issues for applications like email can be big issues for applications like videoconferencing. Bandwidth utilization and the effect of packet collisions were discussed in the previous sections. In this section, we will describe typical packet behaviors in a routed IP network, and their resulting affect on QoS for interactive video.

It is no secret that routers manipulate packets as they flow through the network. In fact, this is the desired result of various queuing algorithms implemented in the router. However, a few of the natural side effects present in a routed network are not desirable for real-time applications such as videoconferencing. A few examples, described further in this white paper are packets out of order, packet duplication, and jitter.

Remember that with the UDP protocol, there is no way to inherently detect lost packets. There is also nothing to stop a UDP application (like interactive video) from sending data, even when the network congestion is causing data loss. As a result, the application itself must be explicitly enabled with functionality to compensate for the side-effects described above.

QUALITY OF SERVICE SOLUTIONS

As was described in the bandwidth management section, merely adding bandwidth is not enough to ensure a high degree of QoS for IP video applications. It might eliminate, or at least reduce, the need for packet prioritization and policy management, but it will not eliminate the issues described above. The reason is that it does not alleviate the effect of router packet manipulation.

Once again, VCON's new PacketAssist™ Architecture was developed with QoS in mind – specifically for IP video applications like videoconferencing. In this regard, PacketAssist ensures the user gets the best possible audio and video quality at a given data rate. It does this with a suite of features designed for QoS and implemented in the endpoint itself. Currently, PacketAssist QoS includes the following features:


- | | |
|---|--|
| <input type="checkbox"/> IP Precedence and DiffServ | <input type="checkbox"/> Overhead Prediction |
| <input type="checkbox"/> Packet Ordering | <input type="checkbox"/> Firewall Port Synchronization |
| <input type="checkbox"/> Packet Duplication Control | <input type="checkbox"/> Maximum Packet Size Adjustment |
| <input type="checkbox"/> Jitter Correction | <input type="checkbox"/> NAT IP Address Mask |
| <input type="checkbox"/> Lip Sync Correction | <input type="checkbox"/> Asymmetric Network Transmission |
| <input type="checkbox"/> Lip Sync Delay Adjustment | |

There are many more widely deployed QoS technologies in the networking industry. Before we describe how the VCON PacketAssist Architecture complements these, we will review the strengths and weaknesses of different QoS technologies. For example, the Resource Reservation Protocol (RSVP) is a true reservation setup and control protocol that is designed to achieve the characteristics similar to circuit switched networks on an IP network. It is primarily suitable for "quantitative" QoS requirements in which the application can discretely identify its quality requirements and parameters. The downside to RSVP is that it is not considered to be highly scalable since it handles QoS requests on a per-flow basis.

In contrast with RSVP, Differentiated Services (DiffServ) is a standard that assigns QoS classifications to different traffic based on service-level agreements between users and service providers. It defines two service levels – assured and premium – that can be applied to network traffic according to pre-determined policy criteria. As such, it is considered more highly scalable than RSVP, but does not offer fine-grain QoS classifications.

Another popular standard, MultiProtocol Label Switching (MPLS), can deliver the required QoS while speeding up the transfer of packets over networks. MPLS places forwarding information on a label attached to outgoing packets, therefore reducing the time required by routers to decide their forwarding paths. MPLS also transfers the packets at the Open Systems Interconnection's (OSI) Layer 2 (switching) level instead of the Layer 3 (routing) level. To integrate MPLS with DiffServ, service providers have to add a few classes for deploying QoS. Examples of such classes are high-priority low-latency "Premium" service, guaranteed-delivery "Mission-Critical" service, and low-priority "Best-Effort" service.

Since this field is evolving so rapidly, VCON has developed a set of features underneath the PacketAssist umbrella, which address some of the most fundamental QoS issues for IP videoconferencing and easily coexist with the above technologies, regardless of their level of support in a given network. Since the PacketAssist features have been implemented in the endpoint, they work hand-in-hand with QoS features present in the network components (primarily switches and routers). Additionally, in many cases, the endpoint is the only place to correct for certain network anomalies like jitter, packet ordering and packet duplication. These features are described in further detail below.



Packet Ordering

With a real-time application like videoconferencing, it is a problem if packets arrive out of order. A missing or late packet can cause the video to freeze (if it is a video packet) or cause gaps in the audio (if it is an audio packet). And since it is very difficult to prevent packets from getting out of order through a routed IP network, VCON PacketAssist™ makes the correction at the endpoint. As packets arrive, their ordering is validated. Out of order packets, up to a limit, are put back in order to maintain continuity of both the audio and video streams for the end user. This compensation is made in the endpoint – not the network.

Packet Duplication Control

With real-time applications, it is a problem if multiple copies of the same packet arrive at the end destination. This too can cause the video to freeze (if it is a video packet) or cause gaps in the audio (if it is an audio packet). And since it is difficult to prevent packets from getting duplicated through the routed IP network, VCON PacketAssist makes the correction at the endpoint. Duplicate packets are eliminated in order to maintain continuity of both the audio and video streams for the end user. Once again, this compensation is made in the endpoint – not the network.

Jitter Correction

When audio and video packets leave the source endpoint, they are evenly spaced at regular intervals – just as they were encoded. However, through the routed IP network, this even spacing can easily be disrupted. Some packets may not have enough delay between them while others may have too much delay between them. This is referred to as “jitter”. Jitter can cause inconsistencies in both the audio and video streams for the end user at the destination. VCON PacketAssist automatically makes the delay uniform in order to maintain continuity of the audio and video streams for the end user. This adjustment is made in the endpoint – not the network.

Lip Sync Correction

When audio and video packets leave the source endpoint, they do so such that a given audio packet is paired with its corresponding video packet. However, through the routed IP network, the various queuing methodologies can treat audio packets differently than video packets. This causes the audio packets to get misaligned with their corresponding video packets. The end result is a loss of lip synchronization. And the more congested the network becomes, the worse this result can become. PacketAssist corrects for this using the RTP timestamp information in each packet. Using this RTP timestamp, the VCON endpoint can determine which audio and video packets should be paired together. Further realigning the packets accordingly ensures that lip synchronization is maintained – even through congested routers or network links. This function is performed in the endpoint – not the network.

Lip Sync Delay Adjustment

At the origin, the amount of time it takes to process audio is different than the time it takes to process video. Factors that affect this include the speed of sound versus the speed of light, the size and shape of the room, and the complexity of the encoding algorithms for audio and video. The end result is that a delay needs to be imposed on the audio stream at the origin in order for lip synchronization to occur. Each vendor typically implements a set of default delay values based on the audio algorithm selected. However, based on the various factors previously mentioned, this default delay value may not be suitable to achieve lip synchronization in all cases. Therefore, PacketAssist allows the VCON destination endpoint to add or subtract delay (from +500ms to –500ms) from the audio latency. The user is presented a slider from which to make corrections for poor delay settings in the remote endpoint.

Overhead Prediction

The reality is that with IP comes some amount of overhead. If the IP video application does not take this into consideration, it could result in packet loss - especially over network links with fixed bandwidth (like CIR frame relay). VCON’s PacketAssist Architecture takes IP overhead into account when establishing a call. In other words, when a user starts a 384K conference, the VCON application should consume as close to 384K as possible, including audio, video and IP overhead.

Firewall Port Synchronization

Many firewalls today have the ability to be configured to support H.323 and SIP traffic. However, the H.323 standard (for example) allows for a wide range of ports to be used for the audio and video streams (1024 - 65535). Because of this, many administrators desire to only configure a smaller range of ports (4800-5000 shown in this example).

This is all fine, however the video endpoint must also be configured to use this same range of ports. Otherwise, it will dynamically request ports in the 1024-65535 range. VCON's Firewall Port Synchronization feature allows the administrator to specify a specific port range for H.323 traffic and have that range synchronized between the firewall and the endpoint. VCON also has a fully comprehensive solution for firewall and NAT traversal, sold under the brand name "SecureConnect".


Maximum Packet Size Adjustment

This feature allows the administrator to change the default packet size used for their H.323 video streams. This can be very useful for mixed voice/video-over-IP deployments. The reason for this is that in some such mixed network environments, the large default packet size (1500 bytes) for video can choke out voice traffic in the routers. This feature is also useful for networks with IP at the edge and ATM in the backbone. The IP packet size can be adjusted so that it is better aligned with the ATM cell size – thereby reducing ATM cell overhead.

NAT IP Address Mask

Network Address Translators (NAT) can cause some real connectivity problems. The reason has more to do with the H.323 standard than the NAT itself. The H.323 standard places source IP address information in the payload of the packet. When the user is behind a NAT, this IP address is a Private (non-routable) IP address. When the NAT translates the IP address to a Public (routable) one, it only replaces the information in the header of the packet - not the payload.

This feature allows the user to manually enter their Public IP address, so that the VCON application can reflect this address in the payload portion of the packet. This feature does not solve the connectivity issues associated with all the various NAT/PAT configurations. It is most beneficial for pure NAT environments in which IP addresses are statically mapped.




Asymmetric Network Transmission

More and more, it is common to have a different outbound versus inbound data rate for network transmissions. Good examples include ADSL, cable modems, satellite transmissions, or even unbalanced IP networks. In this case, most systems would be forced to conference at the lowest common denominator. In other words, an ADSL link with 400Kbps outbound and 1.5Mbps inbound would require the user to conference at 400Kbps in each direction. However, VCON PacketAssist combines bandwidth management and QoS features such as the ones described in this paper to allow asymmetric conference data rates – a different outbound versus inbound data rate. This automatically optimizes audio and video quality based on available bandwidth in a given direction. In the example mentioned previously, the user would be able to send 400Kbps outbound while receiving 1.5Mbps inbound.

END-TO-END QoS ARCHITECTURE

True QoS is best achieved when all network components are working together. This paper has described QoS solutions from many different perspectives – the endpoint, gatekeeper, switch and router. The following diagram depicts how these components can be configured together for end-to-end quality of service. It also demonstrates the complementary nature of VCON PacketAssist running in the endpoint. This diagram is not intended to define an exact QoS solution for any particular network configuration. Rather, the features and components depicted are very common in many enterprise IP networks.



SUMMARY

Although real-time conferencing and collaboration applications offer a different set of deployment challenges than traditional data applications, their increasing acceptance and importance to the enterprise environment has heightened the desire for deployment solutions. And the network of choice for deployment is, increasingly, the enterprise IP network. The benefits of converging voice, video, and data applications onto a single enterprise network outweigh the added planning and design considerations associated with this convergence. In short, the enterprise IP network is a more suitable transport than ever before for real-time applications such as videoconferencing. Bandwidth management and QoS solutions are now very mature and designed just for this purpose – many of these leverage evolving standards such as H.323, SIP, and DiffServ, to name a few.

As a leader in IP-based video solutions, VCON is committed to delivering more than just audio and video algorithms. As this paper reveals, VCON is leading the industry in delivering mature video-based applications that can be effectively deployed on the enterprise IP network. By combining advanced management and administration technologies for the network manager with useful application technologies for the end-user, VCON offers solutions that allow people to communicate how they want and when they want.

CONTACT US

VCON Headquarters

Ph: +972-9-959-0059

Fx: +972-9-956-7244

VCON Americas

Ph: +1-201-883-1220

Fx: +1-201-883-1229

VCON Europe

Ph: +49-89 614 57 0

Fx: +49-89 614 57 399

VCON China

Ph: +86-10-6526-9791

Fx: +86-10-6526-9790

www.vcon.com

VCON

