VoIP - Building a Sound Foundation for Voice over IP

Jason Lackey, Riverstone Networks

Abstract

Conventional wisdom holds that currently there are two types of network operators - those who are currently running VoIP and those who are thinking about it. Even those not planning on immediate implementation of Voice over IP would be reluctant to roll out infrastructure incapable of supporting VoIP. This solution note explores the attributes and characteristics of networks capable of properly supporting Voice over IP or other similar delay sensitive applications.

Introduction

In the past, voice and data were carried on different networks. People also lived in caves.

Traditional network architecture in enterprises and other organizations called for a strict Church and State separation of voice and data traffic. At first, there was no other choice - data networking did not deliver the end-to-end QoS and low latency required to support delay sensitive applications such as voice and video conferencing. Later, although there was a choice, there was reluctance for various financial, technical and political reasons to converge the two networks.

However, as the reliability and capability of modern data networks continues to increase and the cost of bandwidth decreases, unified or converged voice and data networks become more feasible. Indeed, depending upon the method of analysis, some would say that there are compelling cost related arguments for a converged network.

VoIP - What is Voice over IP?

Voice over IP describes the process of converting analog sound into digital packetized data, which is then transported over an IP network. Typical implementations involve some sort of end user terminal, which can either be a physical appliance looking much like a telephone handset or the IP phone can be implemented in software running on a PC with a sound card, speakers and a microphone.

The VoIP Terminal can be a computer or a dedicated appliance

The terminal (computer or IP telephone handset) converts sound into a data stream for outgoing calls and data into sound for incoming calls, most often using the PCM codec, the same algorithm used to encode music on CDs. The handset or virtual handset interfaces with an IP telephony gateway which can either transmit data over your IP network to another station on your network or to a gateway at a different location or connect the call to the Public Switched Telephone Network (PSTN), depending on the location and equipment of the party you are calling.

Standards - H.323 and SIP

There are various standards applicable to VoIP, including H.323 (the ITU-T backed standard) and SIP (Session Initiation Protocol, the IETF standard), which are the most common. For the most part, these standards define the manner in which the handsets and gateways communicate with each other but since they are relatively high level protocols they are not directly concerned with the underlying network infrastructure so long as that infrastructure meets the requirements spelled out below.

Firewalls are an important consideration, as H.323 uses RTP (Real Time Protocol) for data transmission, which instead of using predefined port numbers for the data and the control stream, instead negotiates a random even number port for data and a random odd number port for control traffic. This can make firewall management a challenge. One way around this is to specify a firewall that is H.323 aware when building a VoIP ready network.

Key Properties of a VoIP Ready Network
The four basic characteristics that any VoIP-ready network will have are support for end-to-end QoS, high availability, low latency and low jitter.

End-to-End QoS: QoS (Quality of Service) is the ability to assign different priorities to different types of traffic. For example, it is useful to be able to assign high priority to latency sensitive applications such as video or voice. Other applications, such as email, are not sensitive to delay and can be assigned lower priority. For these applications a little latency is fine as long as the information gets through eventually.

Voice and video traffic, in contrast, are extremely sensitive to delay and misordered packets. The loss or delay of only a few packets can result in disruption and an unacceptable user experience. Because of this it is imperative to have the ability to prioritize voice traffic over other forms of traffic such that someone reading an email does not cut short your phone call.

Riverstone has extensive experience building networks for carriers and service providers who have stringent QoS requirements. Coupled with QoS features implemented in hardware rather than slower and less scalable software implementations, Riverstone delivers end-to-end QoS. 5th generation custom ASICs (application specific integrated circuits, often called chips) enable network operators to turn on QoS and other features without impacting overall network performance.

High Availability: With converged networking it is critical that the network be reliable and available. Any network outage will result in lose of all traffic types; voice, video and data.

HPS Benefits:
- Zero Lost Traffic
- Reduced CAPEX
- VRRP Complement

Features such as HPS (Riverstone's Hitless Protection System) allow control module failover without any dropped packets or lost traffic. This means that not only control module failures but also software upgrades and reboots are transparent to the end user community. This is a critical feature for the end user, but it is also good for the network administrator who can schedule software upgrades at a time convenient to him rather than a time chosen not to impact the user community.

Redundant power supplies, control modules and switch fabrics help eliminate single points of failure while support for Virtual Router Redundancy Protocol (VRRP), which provides a way to assign backup routers in a LAN, allows network operators to cluster routers for even greater network availability. These software and hardware features combined result in what is known as a carrier class platform, ideal for carriers, service providers and those running mission critical networks.

Low Latency: latency is the delay that a network device creates between the receipt of a packet and its retransmission out a different port. In interactive applications low latency is critical to the end user experience.

250 miliseconds is considered by many to be the maximum acceptable latency allowable in a VoIP network. In many situations, Riverstone delivers 8 microsecond latency, delivering a several thousand fold margin of safety.

Low Jitter: jitter is defined by the ITU as "Short-term variations of the significant instants of a digital signal from their ideal positions in time" . Digital communications can be thought of as 1's or 0's sent out on a regular beat. The variation of the timing of the sending of the 1 or 0 off the hypothetical perfect beat would be jitter.

The Riverstone hardware architecture helps insure low jitter with a non-blocking switch fabric and a design that requires all packets, even those entering and leaving on the same linecard, to pass over the backplane. This increases consistency and reduces jitter, optimizing the platform for use in jitter sensitive applications such as IP telephony or multimedia broadcast.

Conclusion
To realize the promise of a converged data/voice network, enterprises and organizations need an infrastructure that is optimized to support voice and other media over IP. Riverstone's RS router family provides the end-to-end QoS, low latency, low jitter, and high availability that allows network operators to build next-generation network infrastructures that will support VoIP and other advanced network applications in the future.

Privacy Statement | Terms of Use | Feedback