TELOIP

Quality of Service Manifesto: Next-Generation QoS Gives Customers the Confidence to Run their Business over IP.

It's time VoIP providers put their customers first and their products second. Only when technology is built with business in mind can Quality of Service be achieved. TELoIP rebuilt the way QoS should be defined—by mapping it against the requirements of the business. If a VoIP provider cannot offer the following list of technological advancements, we believe they haven't got the best interests of their customers in mind.

Quality of Service Manifesto: Change what you think about QoS

Quality of Service Manifesto:

1. Your QoS should avoid latency and jitter.
2. Your QoS shouldn't discriminate: It should support Bi-directional traffic.
3. High-Priority Traffic should never be buffered.
4. QoS should be managed on-site, and should support connections from multiple carriers.

Author
Cameron Couch

When VoIP (Voice over Internet Protocol) started to gain momentum in the early 2000s, everyone was excited about the prospect. The young technology that enabled people to speak over the Internet garnered widespread adoption as people were able to speak almost costlessly to friends and relatives around the world. Organizations, however, were slow to adopt the new technology.

Traditional phone lines were expensive, but were reliable. Business would not and could not tolerate dropped calls. And so, VoIP providers rallied to come up with the answer to breaking into the corporate landscape.

Using SLAs to sell reliability

The perceived risks of "going IP" kept most businesses away from adopting VoIP technology. The VoIP providers were determined to quell customer fears and introduced the concepts of SLAs and Quality of Service. They began by building pieces into their technology to ensure that voice packets reigned supreme across an organization's network. Initially, they focused their efforts on Packet Prioritization—in which they created a hierarchy of packet-types to ensure that some were treated differently as they passed up and down the network connections. These providers used Packet Prioritization to queue packets of information in order to guarantee Quality of Service (QoS). Packet prioritization opended the door for VoIP—without it, VoIP would have no legs in today's fast-paced business world.

A number of algorithms came out of the need to prioritize packets. Network stability meant that sometimes the algorithms were complex. Developers began to implement Schedulers—algorithms that queued packets according to an organization's prioritization hierarchy. The developers started with simple "First-in, First-out" queues and soon evolved to class-based queuing (CBQ) and Priority Queuing (PRIQ) to up-the-ante on the competition. These algorithms made certain that an organization's important packets (Voice) would have priority over the less-important packets (Web-browsing).

Quality, though, was lagging behind. Some packets would get priority, but many would get lost despite the best efforts of the schedulers. Some VoIP providers went as far as creating CODECs to conceal packet loss. The codecs, however, could not solve both quality and bandwidth usage issues together. It was a tug-of-war: customers had to decide whether to increase quality or decrease bandwidth usage. Ultimately, the CODECs were a quick-

fix and failed to deal with the very real issues that threatened the corporate-world's universal adoption of VoIP technology: congestion, layer errors, and network element errors. Changes were afoot, but Quality of Service remained tied to First-Generation packet prioritization techniques, that failed to offer substantial advancements in quality or reliability.

One major error VoIP technology providers were making, we believe, was that they had built technology and were determined on tweaking it to be more palatable to business customers. To respond to the fear of their Internet "going down", the VoIP technology industry began using SLA, 99.999% SLA, and were relying on QoS to sell the product.

QoS is about Business, not about Networking

At TELoIP, we believe that quality comes from a variety of improvements in IP networking—improvements on which companies should never compromise if Quality of Service is important to how they do business. Networking technology, we believe, should support all of the critical applications of any business—regardless of that business's connections, geography, or size. Instead of tweaking an existing technology, we developed our "All over IP" Ana™ technology so that it meets the demands of the business.

In order to satisfy the needs of all organizations—big or small—a Next-Generation QoS needs to incorporate technological enhancements that make Business over IP possible, not just a single application. If IP solutions are going to truly succeed at lowering network infrastructure-related costs, changes are going to have to be made. We've put together a QoS Manifesto (if you will) for IP network infrastructure that enables the delivery of Next-Generation QoS.

Principle 1:
Your QoS should avoid latency and jitter.

Networks are complicated, there's no question. And frequently they're being heavily used. Traffic is running up and down your connections all day—sometimes the traffic is so great that the connections fail (or deliver poor quality) as they're hitting a usage threshold. It happens frequently, and we believe it's one of the major reasons why companies are hesitant to adopt IP-based voice applications.

First-Generation:
If the connection fails, the call (or session) is dropped.

We mentioned how QoS and SLAs were built to sell VoIP and other IP-based systems, but that the industry failed to truly address the needs of today's business. This comment is never more apparent than in relation to how the technology addresses failover. First-Generation technology frequently reacts to disturbances in the network connections and will failover to either another link (if Bonding is the chosen solution) or another area of a dedicated amount of bandwidth. Often, the system's reaction is late (and sometimes non-existent)—and the user experiences undue jitter or latency, or in some cases, the call or application session is dropped entirely. The result of a failed connection could be as simple as having to reload a web-page, or as critical as a dropped sales-call.

Next-Generation:
Getting in front of the problem.

Next-Generation providers truly appreciate and support all of the critical applications of their customers. They offer a mechanism that detects increased network disturbances and fails-over before the user experiences any jitter or latency during the session. A Next-Generation system is able to monitor link quality and failover to another connection (if it's a converged solution) in order to keep the application running smoothly with high quality. These same systems that failover pre-emptively should also be able to self-correct and fail-back to the primary connection once the disturbance goes away. A customer on a next-generation solution should never have to worry about an application.

Principle 2:
Your QoS shouldn't discriminate:
It should support Bi-directional traffic.

Not all bandwidth is symmetrical. Not all bandwidth is asymmetrical. Your network should be able to put Quality of Service on all of your traffic, regardless of whether it's inbound or outbound.

First-Generation:
Wasted Bandwidth with Symmetrical QoS.

Your business is your business. We're not going to tell you what applications can and can't be QoS'ed. We believe you should be able to QoS them all. Too often, we see network solutions that can only QoS symmetrically—it wouldn't matter what amount of bandwidth your connection offers for inbound data, Quality of Service would only be placed on whatever was available for outbound data. *Diagram 1* shows the two options that are traditionally available for QoS over cloud connections. In the first image, Quality of Service is offered on a standard symmetrical link.

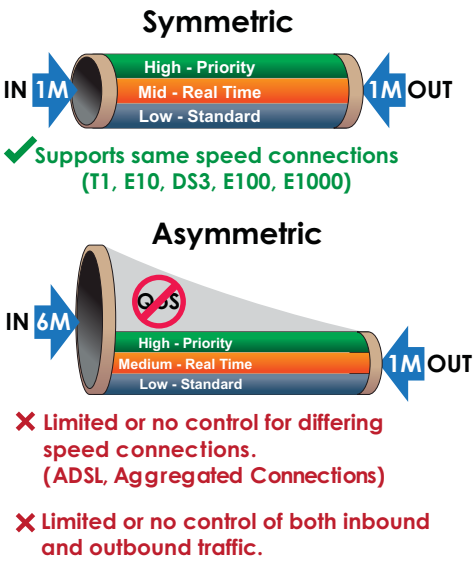

Diagram 1:
Symmetrical QoS limits bandwidth availability

The bottom graphic shows an example of MPLS networks that we see frequently where the service provider has used an ADSL for the network connection—but can't offer Quality of Service on the assymetrical download bandwidth. In this case, the provider would only be able to QoS the inbound bandwidth at whatever the outbound rate is—regardless if the ADSL has an inbound bandwidth capacity of 10Mbps, if the outbound rate is 1.5Mbps, the QoS is on 1.5Mbps for both. Frequently, the customer is not given the unQoS'ed bandwidth at all. The extra 8.5Mbps inbound can't be used by the customer—as it may disrupt the network.

The Next-Generation:
Supporting the connection's entire capacity.

We believe that Quality of Service should exist on all of the bandwidth, regardless of the direction of traffic. The Next-Generation of Quality of Service should be advanced enough to support symmetrical or asymmetrical connections and offer the QoS on all traffic in the network.

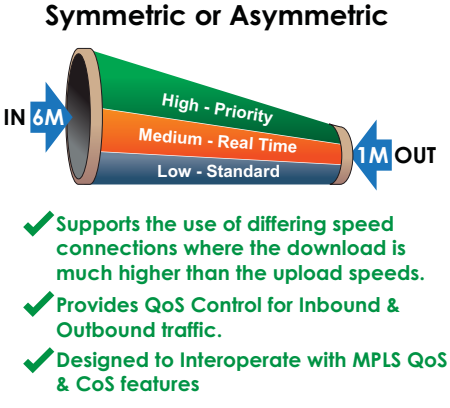

Diagram 2:
Asymmetrical QoS can control & use all available bandwidth

Principle 3:
High-Priority Traffic should never be buffered.

As Priority Queuing has evolved, the most common networking technologies use a buffer to manage traffic. However, traditional priority queuing algorithms were used to apply a buffer as an answer to congestion issues over the network—the priority queue would essentially be "turned on" and would operate at certain times of the day, when traffic in the network was considerably high. These Priority Queuing (PRIQ) algorithms had no way to deal with shoulder traffic (traffic that occurs around those times, but not exactly during those times). The algorithm would not get applied to any traffic in the network occurring around that window of time.

Next-Generation:
High Priority traffic gets Right-of-Way always.

Next-generation PRIQ algorithms ensure that a certain amount of bandwidth is reserved for critical applications. Critical packets are not held-up by any type of congestion in the network.

At TELoIP, we've taken prioritization one step further and dedicated a "diamond" or "HOV" lane for high-priority traffic where the traffic passes freely across the allotted bandwidth with no congestion, ever. We like to consider it the HOV lane for Emergency Response Vehicles.

If your network is not configured in such a way as to optimize the traffic for critical applications, we believe you should be looking for a system that does—otherwise, you limit the quality of any of those applications. And if you're limiting the quality, you're limiting your business.

Principle 4:
QoS should be managed on-site, and should support connections from multiple carriers.

Service providers have been offering Quality of Service on their own links for years now. And frequently, they will offer up what they call 5-nine Service Level Agreements (SLAs) as a way to give their customers peace-of-mind that their network is reliable and won't go down. We believe that reliability should come from a combination of technology, diversity in connections and suppliers, and proper application management.

First-Generation:
Single-Supplier False Redundancy.

If Carrier A provides a customer with a connection, Carrier A's reliability is a major external risk to the customer's network. If the customer has only one cloud connection, if Carrier A goes down, so too does the customer's network. If that single connection is disrupted by any outside influence (ex. Construction), that's another risk borne by the customer. Some carriers have offered network convergence solutions by providing up to four connections, but the redundancy is still false: if Carrier A goes down, all of the connections are unusable, and the customer's network still fails. The only traditional way around this issue is to charge the customer for expensive dedicated links that makes QoS a privilege of the wealthy.

Next-Generation:
True Redundancy from Autonomous QoS.

The Next-generation Quality of Service technology allows for the customer's on-site equipment to act as a control mechanism that can eliminate the reliability risk. next-generation QoS should support connections from multiple carriers. This advancement in technology means that redundancy can be created by offering QoS across a network that is comprised of multiple connections from multiple carriers. If Carrier A goes down, Carrier B is still up to support the business. The Next-Generation technologies can offer QoS on all of the connections regardless of the supplier—whether it's Cable, Fiber, Fixed-Wireless, DSL, or even Dial-Up for the more remote customers. A next-generation network provides affordable reliability for customers.

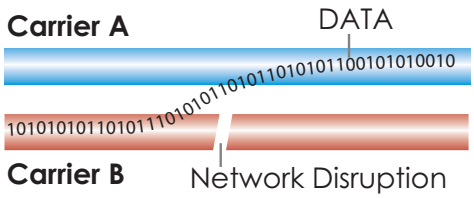

Diagram 3:
Using multiple carriers, DATA can pass from carrier to carrier uninterrupted.

How important is next-generation technology?

As more and more application providers move their technology to the cloud, the need for reliability on a company's network is going to grow substantially. Companies that currently have a Virtual Private Network (VPN) and one or two other IP-based applications are going to see the number of hosted applications they use every day rise dramatically. Customer Relationship Management tools, Risk Management solutions, Video and Voice-Conferencing Applications, and a myriad of other applications are all going to be accessed online by companies who are looking to streamline their growing IT departments. Having a reliable network is a must for these organizations.

First-generation Quality of Service technologies don't have the forward-thinking developments that can support such a massive transition. A company that is looking to truly take advantage of the benefits of IP-based applications needs to have a network that offers true reliability, a smart disturbance-monitoring mechanism, a super-highway for their most critical applications, and usage of all of their asymmetrical bandwidth. Times are changing and next-generation network technologies are leading the way in developing truly beneficial components that facilitate operating a Business over IP.

TELoIP Inc.
487 Adelaide Street West
Suite 200
Toronto, Ontario, Canada
M5V 1T4

TELoIP is a Toronto-based technology firm serving the most forward-thinking organizations in the North American communications market. TELoIP creates and sells network convergence solutions based on a technology that enables voice, data, and video to work simultaneously over the same infrastructure. Their patent-pending Ana™ system aggregates and accelerates any connection, from any carrier - combining the speeds of each connection and providing customers with reliability for their critical applications. Bandwidth requirements are growing at exponential rates, TELoIP's technology helps companies scale up quickly and reliably to meet tomorrow's growing bandwidth requirements. For more information, please visit www.teloip.com