Real-time Applications on IP Networks: Overcoming Economic Constraints on Quality of Service

By John Bartlett and Peter Sevcik
November 2004

NetForecast Report 5075

The allure of cost savings drives the need for network convergence - yet the technical difficulties of merging real-time traffic streams (voice and video) with data streams has stymied the industry for over a decade. The simplicity of operating a single network, the lower cost structure of IP, and the promise of new applications that can leverage both the voice and data infrastructure, entice to pursue the converged network.

In this paper we look at the behavior of data and voice/video traffic on an IP network, and examine the tradeoff between connection quality and link utilization. We then assess a new technology, Rivulet Queuing, which overcomes the economic barrier by enabling high utilization and very high quality connections in a standard IP network environment. Lastly we present two examples illustrating the value this new technology delivers.

Understanding the Problem

IP networks drop packets. The Internet Protocol (IP) is classified as an unreliable transport because packets may get dropped. This is a well-understood and accepted fact of life. Most data applications that require fully reliable transport use the Transmission Control Protocol (TCP) on top of IP (TCP/IP) to ensure packet delivery. TCP recovers from packet loss by tracking packet sequence numbers, and requesting lost packets to be resent. Packets are then handed off to the receiving application in order and with 100% delivery.

Packet loss in an IP network is not just an occasional error condition, but a regular occurrence. The TCP congestion management algorithm, which slows down transmission when bandwidth resources are scarce, relies on packet loss to let it know when congestion is occurring. Each TCP flow constantly increases its transmission rate, only backing down when packet loss is detected. After backing down, TCP slowly increases its transmission rate again until loss occurs. Thus TCP causes packet loss as a way of managing its flow rate.

Real-time traffic streams do not have the luxury of resending lost packets, because the data they carry is being consumed on arrival. The packets of a voice conversation are played within milliseconds of their arrival at the receiving phone or computer. A missing packet causes an interruption in the audio content because there is insufficient time to detect the missing packet, request a new copy, and wait for its arrival.

When real-time streams are mixed with best effort data streams, the quality of the real-time streams is compromised by packet loss. The extent of the loss depends on the utilization of the link, but loss is inevitable. Quality of Service (QoS) mechanisms overcome this loss by granting priority to the real-time streams (more on this later).

But first let's look at the relationship between link utilization and packet loss. Figure 1 shows curves developed mathematically and verified through simulation (see "Economics of QoS on WAN Access Lines," *Business Communications Review*, October 2004, pp. 16-22). The X-axis in Figure 1Figure 1 represents average link utilization, and the Y-axis represents packet loss probability. Four curves are shown. The curve marked '1-Hop Bursty' shows loss probability for one network link carrying typical computer data traffic, which is by nature bursty. The curve marked '1-Hop Well Behaved' shows loss probability for voice, video or TDM emulation traffic. These streams have much more regular bandwidth use, and so do not cause as much packet loss. The two remaining curves show the loss for a 4-hop network. Each link adds additional packet loss so the numbers increase as network distances increase.

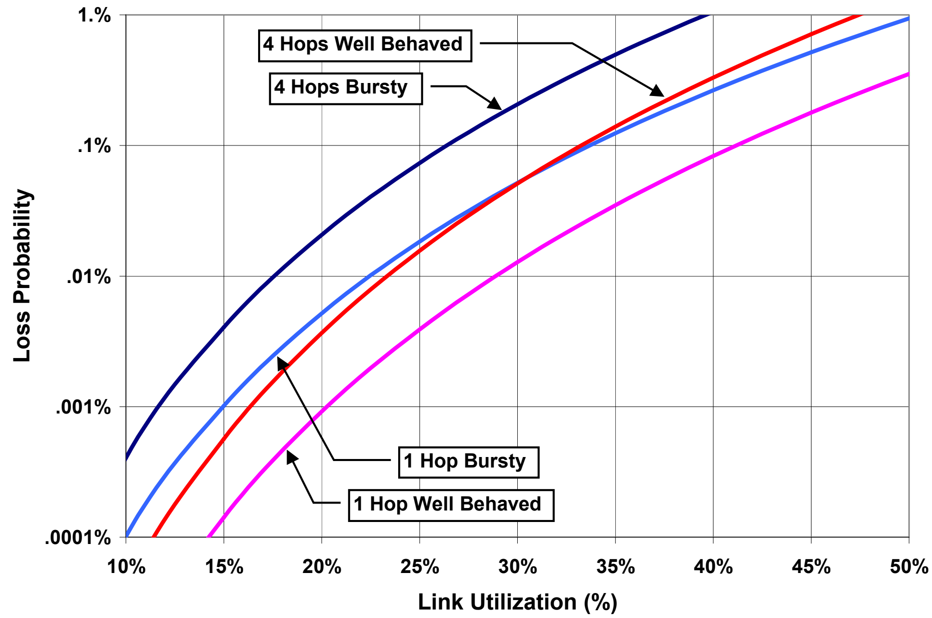

Figure 1 - Effect of Link Utilization on Packet Loss Probability

The loss curves for well-behaved traffic indicate packet loss for a network carrying *only* voice, video or TDM emulation flows. When data traffic is mixed with real-time traffic, the bursty nature of the data pushes the loss back towards the bursty curve, causing loss to both data and real-time flows.

Note that Figure 1 uses a log scale on the Y-axis. You may be more familiar with a loss curve that rises exponentially as utilization approaches 100%. Using the log scale slows more detail in the lower utilization ranges.

To use these curves to predict a level of packet loss, find the network's current utilization level on the X-axis, and note the position of the relevant curve on the Y-axis. Conversely to design a link to maintain a packet loss maximum, choose the loss value on the Y-axis and note the position of the relevant curve on the X-axis to find the maximum allowed link utilization.

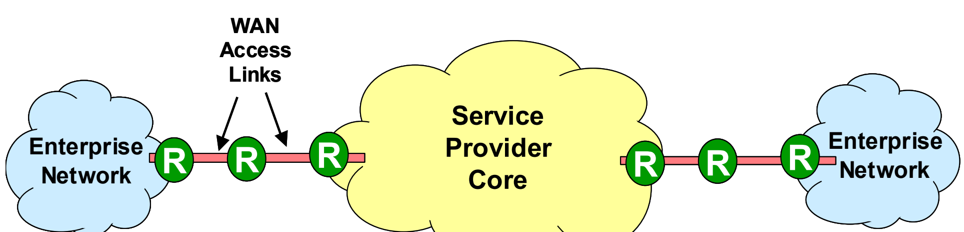

Figure 2 - Example Enterprise-to-Enterprise Network

Suppose we are designing a network for mixed audio/video and bursty traffic, and we wish to keep loss probability below 0.01%. The network has two access link hops on either side of a service provider cloud, as shown in Figure 2. The behavior of the mixed voice/video/data traffic will lie somewhere between the two curves in Figure 1, but let's assume the worst case (with minimal voice) and use the blue bursty-data curve.

A loss rate of 0.01% shows that utilization must be kept below 18%. If we relax the specification to 0.1% we can push utilization up to about 27%. In either case this constraint means the assets involved are poorly utilized, which raises the overall cost of supporting both data and voice/video on this network.

Using Quality of Service (QoS)

The traditional answer to this packet loss problem has been to give priority to the loss sensitive streams. There are many variations of QoS, but the one most commonly implemented for real-time streams is a simple priority mechanism called low-latency QoS. The router separates high-priority traffic into a separate high-priority queue at each output port. Whenever there is a packet in the high priority queue, it is sent ahead of any packets in the lower priority queue.

Because the high priority traffic is always forwarded first, it behaves as if it is the only traffic on the link*. The curves in Figure 1 can again be used to predict packet loss, but now the X-axis (utilization) represents the link utilization of only the high priority traffic. By moving traffic to high priority we have moved from the combined utilization of all traffic to the utilization of the high priority traffic only.

Figure 3 shows an example of traffic crossing a 4-hop access network, consisting of half voice and half data traffic, with a combined utilization of 40% of the link. Without QoS we expect about 1% packet loss. Implementing QoS for the voice portion of the data stream reduces the utilization for the high priority traffic to 20%, and moves us to the curve for well-behaved traffic. This results in an expected packet loss of only 0.003%. This is the advantage that QoS has touted, and in many applications the advantage this shift provides is adequate for supported applications.

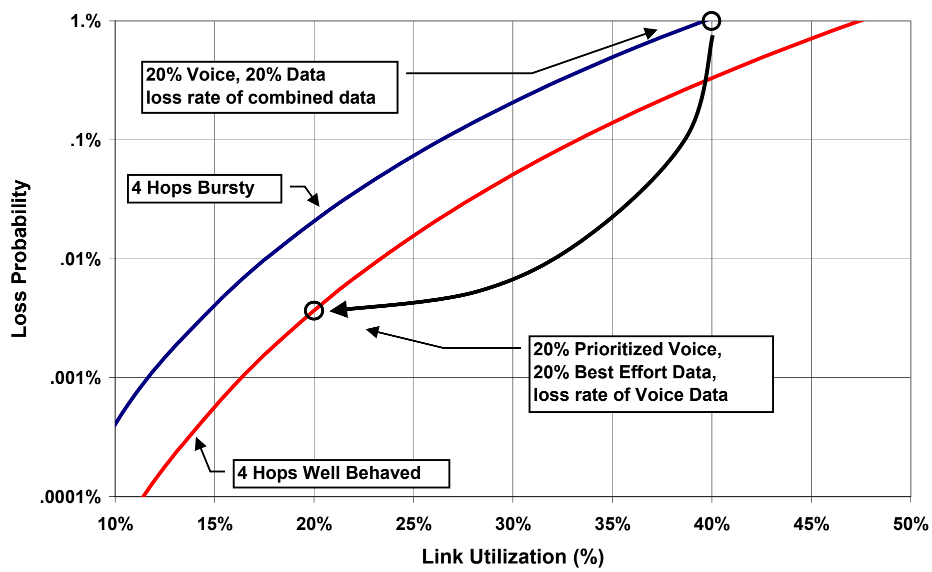

Figure 3 - Effect of Adding QoS on Voice/Video Packet Loss

* High priority traffic can be delayed if a lower priority traffic packet is in the process of being sent when the high priority packet arrives. This delay can be as long as the longest packet divided by the link rate. A maximum size Ethernet frame (1500 bytes) on a T1 link (1.54Mbps) can result in a 7.8 ms delay. At higher data rates this number quickly drops and becomes insignificant.

Over time, new applications have appeared that require a higher quality transport. Higher speed video, MPEG4, uncompressed broadcast video and TDM emulation require lower packet loss rates to function properly (see Table 1).

Table 1 – Application Packet Loss Requirements to Maintain Performance

Application	*Packet Loss Req.*
TCP (Data Applications)	0.1% (10^{-3})
Business Quality Videoconferencing and Voice	0.01% (10^{-4})
High Quality Video, SANs, and Cryptographic flows	0.001% (10^{-5})
Minimum Circuit Emulation	0.0001% (10^{-6})
High Quality Circuit Emulation	0.00001% (10^{-7})

QoS reduces but does not eliminate packet loss, and utilization still plays a role in the loss that does occur. The lower the desired packet loss, the lower the utilization required. Using the network shown in Figure 2, we calculate the percentage utilization possible for the increasing quality levels described in Table 1.

Figure 4 shows the maximum utilization values for high priority traffic with decreasing packet loss probabilities. The bottom portion of each bar represents the percentage of the link that can be used by high priority traffic for each specified packet loss. The upper portion of the bar shows the overhead portion, which can be used by low priority traffic, or must remain unused. We see that as the required quality level increases (as packet loss maximums decrease) a smaller and smaller percent of the link is available for high priority traffic. This limitation makes bandwidth allocation difficult, and results in low link utilization for high priority traffic, which raises the cost of transport.

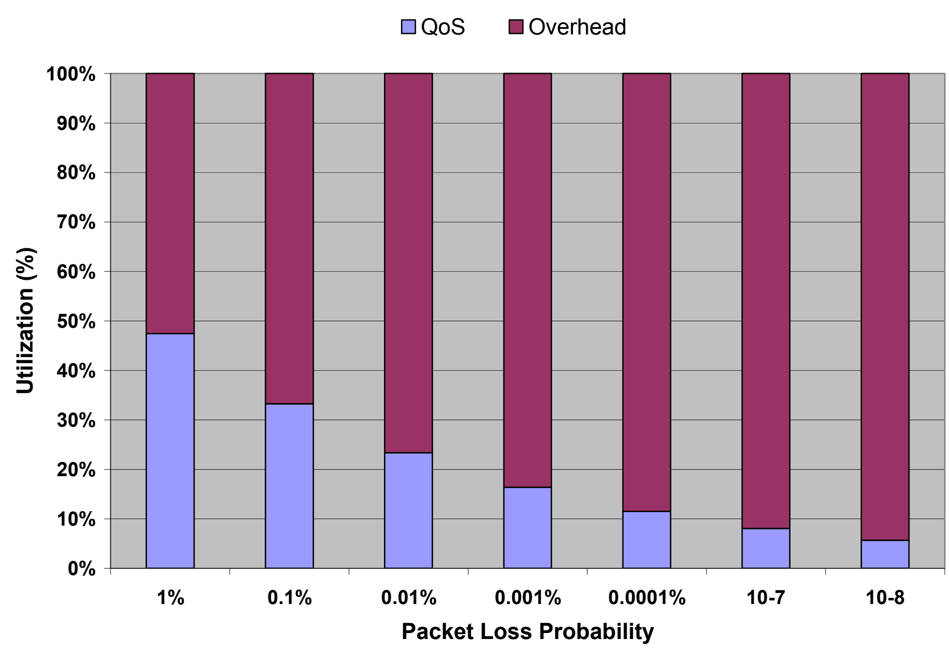

Figure 4 - Overhead Required for 4-Hop Network Connection Using QoS

Rivulet Queuing Technology

Rivulet has introduced a new technology, which overcomes the constraints of QoS. Rivulet queuing is able to carry much higher percentages of real-time traffic through a shared network infrastructure, while maintaining extremely low loss and jitter. An added benefit of this technology is that overall network utilization can be higher without impacting the quality of the high priority traffic streams.

Packet loss and jitter occur when the packets of multiple streams flow together in a router and fill or overwhelm the queue at the router output port. Traffic streams in a network are not synchronized to each other, so traffic bursts or even individual packets from voice or video streams may arrive at an output queue simultaneously. The probability of this happening is described by the packet loss curves in Figure 1.

Rivulet Queuing overcomes this problem by synchronizing IP packet flows so that high priority traffic streams never contend for an output port queue. By arranging the arrival time of packets such that there is no contention, queuing is nearly zero, jitter is nearly zero, and packet loss due to queuing is eliminated.

Consider the queuing delay that occurs on a highway during rush hour traffic. If cars were perfectly synchronized to enter an open slot at precisely 70 mph, cars could drive bumper-to-bumper in concert with other traffic at the same speed, like a train. This would maximize utilization and eliminate contention. This is essentially what Rivulet does with appliances at the edge of the network or software in endpoints, and it does so using standard routers and the standard IP protocol.

More information about how Rivulet Queuing accomplishes this synchronization can be found in white papers on the Rivulet web site (www.rivulet.com).

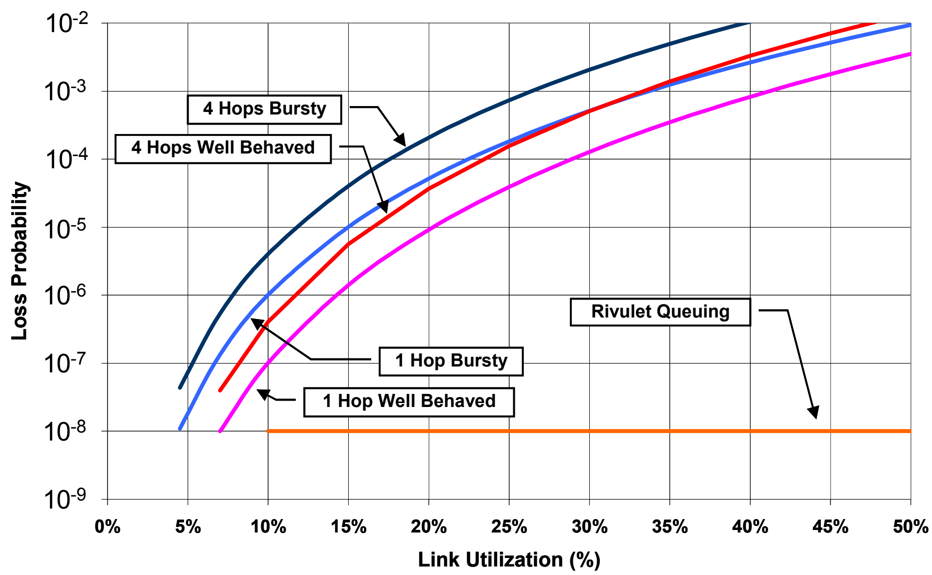

Figure 5 - Loss Curves Including Rivulet Queuing

The value of the Rivulet technology can be seen by looking at loss versus utilization curves. Figure 5 shows the four curves we saw in Figure 1, extended to lower loss levels, and includes a line for traffic carried using Rivulet Queuing. The loss probability for Rivulet based traffic no longer depends on utilization, so the line is flat. Loss is now determined by the link bit error rate which is at most 10^{-6} for a T1 line, and much less for fiber optic connections. Figure 5 shows that it is now possible to carry data streams with very low loss, and not be limited to very low utilization.

Rivulet Application Examples

Let's examine two examples to better understand the value provided by the Rivulet approach.

An Enterprise Voice Example

Example 1 is an enterprise that is deploying voice on their IP network, and is pleased with the flexibility and features it provides. However, they have experienced quality issues and see the need for QoS to ensure voice streams do not experience packet loss, especially during the busy periods of the IP network.

The network in this example is the same as in Figure 2. For this analysis we will ignore the loss in the enterprise network, assuming that utilization is sufficiently low to keep loss at least an order of magnitude below that in the WAN links. The WAN consists of two network hops into the service provider cloud on either end, plus the service provider network itself. The service provider does not implement QoS within the core, but does prioritize marked traffic as it leaves the core and enters the access links. The enterprise can prioritize packets going towards the WAN using the enterprise edge router. For the moment we will assume that the core packet loss is also negligible, and revisit this assumption at the end of the example.

Voice traffic using the G.711 codec, packetizing every 20 milliseconds, will consume 87Kbps of link bandwidth. If header compression is implemented, utilization can be reduced to 68Kbps.

Our preferred design goal is to keep loss below 0.1% to maintain high quality voice connections. The four-hop loss curves (Figure 1) show that well-behaved priority traffic must remain under 27% utilization to maintain our 0.1% loss specification. Without header compression we are limited to 4 simultaneous calls on this link. With header compression the number can be increased to 6.

Implementing Rivulet in this environment eliminates this upper bound on the number of voice calls possible, because Rivulet can carry much higher traffic levels without increasing packet loss. In this example, loss in the access links would be held to below 10^{-6} (0.0001%) by the Rivulet technology.

We assumed earlier that the effect of the core network was negligible. Most service providers today are offering a 0.1% or perhaps a 0.01% packet loss service level agreement (SLA), but no better. If the service provider were to implement Rivulet, loss for high priority traffic would drop to below 10^{-6} as well. Implementing Rivulet can also address losses in the Enterprise network which we chose to disregard in this analysis.

A Service Provider TDM Emulation Example

In our second example we look at a network service provider wishing to support voice, video and circuit emulation in concert with normal data traffic. The analysis focuses on circuit emulation, since it has the strictest loss requirements.

Circuit emulation requires a loss rate of no more than 10^{-6} and preferably 10^{-7} packets. This low loss rate must be maintained across access links and through the provider core. We assume an average of 4 access hops and 8 core hops to complete a national level connection.

The behavior of the service provider's core network may be much better than the curves in Figure 1 predict. Much longer queue lengths are possible in the network core because data rates are so high. A full 64-packet queue that takes 500ms to empty at T1 speed, empties in 1.2 ms at an OC12 rate. Jitter can be kept within reasonable bounds even with large queue sizes, which also work to keep loss low. Service provider SLAs, however, often don't specify loss levels below 0.1% or at best 0.01% packet loss, and this number is the average over a week or a month's time. But let's assume that the service provider can reach a loss level of no more than half our requirement, or 0.00005% loss.

To obtain the necessarily low loss rates in the access links, we use the same analysis as before. Using Figure 1 and referring to the 4-hop well-behaved curve, we find our specification of $5x10^{-7}$ is less than the chart describes, but at 10^{-6} we see that utilization must be below about 12%. Assuming a T3 access link, this allows only (12% x 45Mbps =) 5.4Mbps for all TDM emulation, voice and video traffic. (Remember the assumption we made earlier about the ability of the core network to maintain the low loss levels required.)

Two steps are required to use the Rivulet technology for this example. First, the service provider implements two high priority levels throughout his network, dedicated to Rivulet based traffic. These traffic priorities are higher than any other user traffic carried in the network. Second, Rivulet queuing appliances are deployed at each edge location (CPE) where voice, video or TDM emulation traffic will be carried. The Rivulet appliances use the two high priority traffic levels to synchronize traffic passing through the network, thereby eliminating packet loss for these streams.

With the Rivulet approach, much more voice, video and TDM emulation traffic can be carried, while maintaining the low loss levels required for proper TDM emulation.

Summary

Until now, low utilization levels were essential for high quality voice, video, or TDM emulation streams, and high link utilization caused high packet loss. Quality of Service can distinguish between high and low priority traffic streams, and improve loss - but utilization must still be bounded to maintain high quality streams.

Rivulet Queuing changes this equation by decoupling loss rates from link utilization, thus enabling high quality streams in a converged, yet highly utilized network. Rivulet supports real-time traffic streams with no queue packet loss at as high a utilization level as required. The tradeoff between high priority traffic volume and best effort traffic volume can now be determined by the application requirements, rather than by the economic constraint of expensive WAN bandwidth.

For more information on Rivulet, visit www.rivulet.com.

John Bartlett is Vice President of NetForecast, and has 26 years of experience designing at the chip, board, system and network levels, with a focus on performance. John led the team that built the first VLAN implementation, one of the first ATM switches, and he is a leading authority on the behavior of real-time traffic on the Internet. He can be reached at john@netforecast.com.

Peter Sevcik is President of NetForecast and is a leading authority on Internet traffic, performance and technology. Peter has contributed to the design of more than 100 networks, including the Internet, and holds the patent on application response-time prediction. He can be reached at peter@netforecast.com.

NetForecast develops customized analytic models to determine the business value of new technologies.

Additional information on network economics is available at:

www.netforecast.com