Effect of Rivulet Based Traffic on Existing Data

By John Bartlett and Peter Sevcik

December 2004

In a previous paper (NetForecast Report 5075) we discussed the need for Quality of Service (QoS) capabilities, how this is addressed by priority based QoS systems and how it is managed with the synchronous approach used by Rivulet. The ability of the Rivulet technology to align packets in time to prevent contention or queuing delay means Rivulet can deliver packet streams with very low loss and jitter on converged networks.

What the previous paper left undone is to describe the effect of introducing either QoS-based or Rivulet-based high priority traffic on the existing best effort traffic in the network. Here we study the interaction between high priority and best effort traffic and then describe that effect. Results are derived using the NetForecast Network Simulator (see Sidebar.)

We find in this study that the introduction of either Rivulet or low latency QoS has an impact on the underlying best effort traffic. We see through the simulation results that the Rivulet impact is significantly less than the impact of QoS, because of the even bandwidth utilization Rivulet provides. While packet loss for QoS increases with utilization, Rivulet is able to carry traffic streams with almost no loss or jitter across converged networks, with only minimal impact to existing best effort traffic.

Understanding the Problem

We know from experience, and have verified both with mathematical models and simulation, that the bursty nature of data traffic causes packet loss at much lower utilization levels than are predicted by an M/M/1 queuing model. A typical data link running at 35% average utilization is nearing its peak usefulness. Beyond this level, packet loss increases to unacceptable amounts, TCP-based applications slow down, users become less productive and help-desk phones begin to ring. A real-time analysis of link utilization shows many high peaks or bursts of utilization where applications momentarily consume 10 times or even 100 times their average bandwidth. Supporting these bursts is critical to maintaining application performance. Thus we consider the unused portion of the link bandwidth as a necessary buffer to insure low loss and proper transaction performance.

When we use a priority QoS mechanism to support a preferred class of service, we change the dynamic of this headroom. As described in the previous paper, the high priority traffic behaves as if the entire link bandwidth is available to it. Whenever a high priority packet arrives in an output queue, it is forwarded onto the link, so the high priority traffic behaves as if the link is empty.

We would expect the high priority traffic to have the same loss curve as if it were the only traffic on the link. Except for the slight interference caused by best effort packets in the process of being clocked out, this appears to be true. Thus the primary impairment for high priority traffic is other high priority traffic.

Assume, for example, that this high priority traffic has a utilization equivalent to 10% of the total link bandwidth. Unless the traffic profile is perfectly regular, it will have bursts that consume portions of the 'unused' bandwidth. This high priority traffic is taking advantage of the headroom.

Now consider that the best effort traffic is also trying to operate in the bandwidth that remains. The average bandwidth available is 90% of the link capacity. However, this bandwidth is not constantly available, it comes and goes as bursts of high priority traffic

NetForecast Report 5076

swell and die down. So how does this varying bandwidth affect the best effort traffic? We turn to the simulator to find out.

Effect of QoS Traffic on Best Effort Traffic

The simulator was run with varying percentages of high priority and best effort traffic to produce the curves in Figure 1. Low latency QoS (priority) is used for the high priority traffic. The high priority traffic profile is voice over IP (VoIP) calls from a large base of users, who are starting and ending relatively short telephone calls. The curves represent the loss behavior of the best effort traffic; high priority loss curves are not shown.

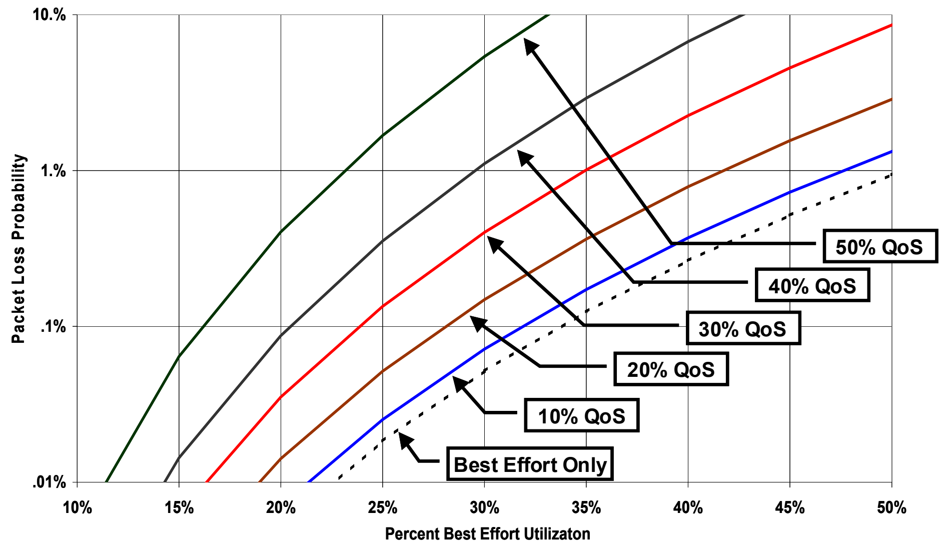

Figure 1 - Best Effort Traffic Loss Using QoS

The X-axis of Figure 1 shows the percentage of link bandwidth used by the best effort traffic stream. The Y-axis of Figure 1 indicates the probability of packet loss for the best effort traffic.

The curve in Figure 1 marked "Best Effort Only" is the packet loss probability when no high priority traffic is present. This is the same loss experienced when no QoS mechanism is implemented.

The curve marked "20% QoS" is the packet loss probability when 20% of the link bandwidth is consumed by high priority traffic, leaving 80% for the best effort traffic and its headroom. The 30% point on the X-axis now represents 30% best effort traffic. Where the 20% QoS line intersects the 30% point on the X-axis, we find the operating point for a link carrying 20% QoS and 30% best effort traffic, for a total link utilization of 50%. The Y-axis shows the packet loss probability for the best effort traffic. Thus the QoS traffic causes the best effort loss rate to increase from 0.05% to 0.14%.

Effect of Rivulet Traffic on Best Effort Traffic

A second simulation run implements high priority traffic using the Rivulet scheme. Because Rivulet synchronizes data flows to prevent contention, the bandwidth consumption of high priority traffic is not as variable, and the impact on best effort traffic is lower. Figure 2 shows the curves resulting from the Rivulet-based simulation run. Note that the area between the curves is smaller, indicating a lower loss rate for best effort traffic using the Rivulet approach.

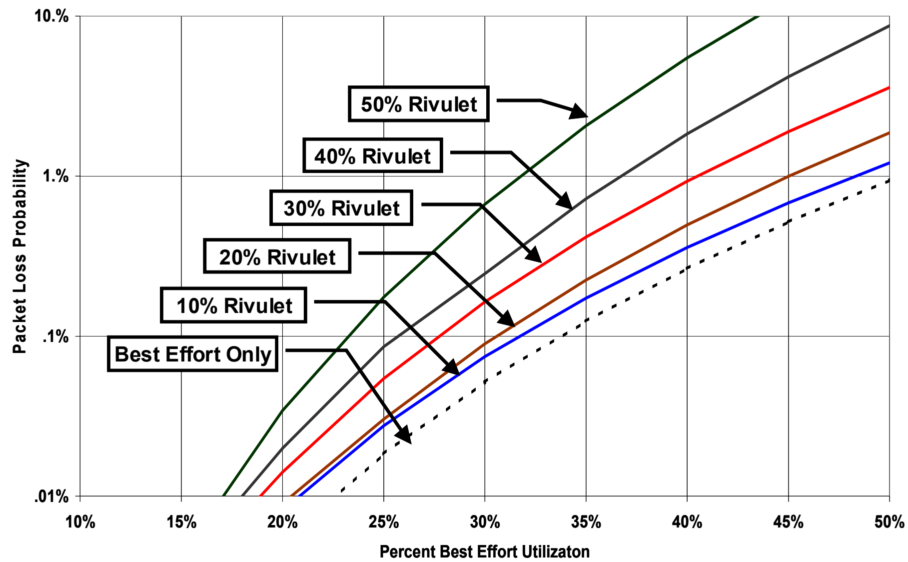

Figure 2 – Best Effort Traffic Loss using Rivulet

An Enterprise Example

Lets now consider an example to better illustrate this effect on a typical carrier connection. For our example (see Figure 3) we will model an Enterprise-to-Enterprise connection, which uses two T3-hops to gain access to the carrier core at each enterprise location. We will use the curves presented above to model packet loss in the access links.

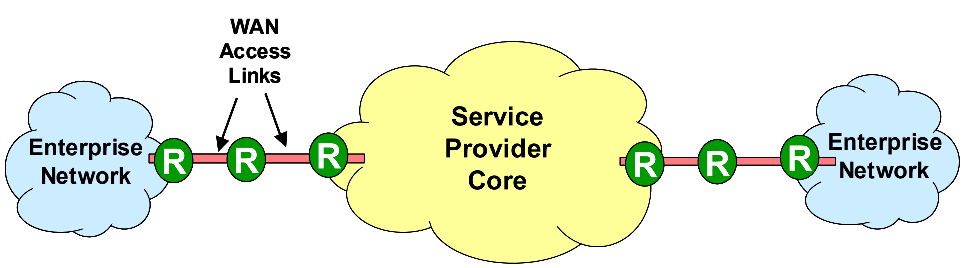

Figure 3 - Example Enterprise-to-Enterprise Network

For packet loss in the Service Provider core we assume that the carrier is running at a utilization level that provides no worse than a 0.01% or 10^{-4} packet loss rate. If we use QoS in the access links, we assume that QoS is provided on all access link routers for traffic both entering and leaving the service provider network. We do not assume QoS is implemented in the service provider core.

For the Rivulet-based approach we assume that a low latency 3-level priority QoS mechanism is implemented throughout the service provider network, including the core, and that the Rivulet technology is the sole user of the higher 2 priority levels.

The enterprise customer, in our example, is currently using only best effort service, but would like to add high priority voice traffic amounting to 20% of their existing data rate.

Table 1 lists three possible utilization rates for this customer that we will explore in this example, and shows the resulting loss rate the customer currently experiences.

Table 1 - Example Utilization and Best Effort Loss Rates with no Real Time Traffic

	Link Utilization Rates		
Total Link Utilization	30%	40%	50%
Best Effort Loss Rate	0.05%	0.23%	0.80%

There are a number of ways to think about the addition of the high priority traffic impacting the network, which depend on the priorities of the customer and of the network provider. We explore three cases here and show the results for each case.

Case A:

In this case we make the assumption that as the enterprise adds high priority traffic, they will reduce the best effort traffic volume so that their total link utilization does not change. This scenario is typical of an enterprise that is running voice or video over IP protocols, and determines that QoS is needed because current voice/video quality is suffering from packet loss or jitter. The existing voice or video traffic will be moved to the higher priority level; overall utilization will remain the same.

The first line of Table 2 shows the three total link utilization cases, and the second line shows the best effort loss rates before QoS or Rivulet is introduced (the same as in Table 1.) Line 3 of Table 2 shows the line utilization of high priority traffic (20% of the total traffic), and line 4 shows the utilization of best effort traffic, reduced to keep total link utilization constant.

Table 2 - Maintaining the Same Utilization

	Low Latency QoS			*Rivulet*		
Total Link Utilization	30%	40%	50%	30%	40%	50%
Initial Best Effort Loss	0.05%	0.23%	0.80%	0.05%	0.23%	0.80%
High Priority Utilization	6%	8%	10%	6%	8%	10%
Best Effort Utilization	24%	32%	40%	24%	32%	40%
High Priority Loss Rate	0.01%	0.01%	0.01%	10^{-8}	10^{-8}	10^{-8}
Best Effort Loss Rate	0.02%	0.15%	0.48%	0.02%	0.09%	0.35%

The 5th line of Table 2 shows the packet loss probability for the high priority traffic. High priority loss for the QoS case is set by the network core at 0.01%. The access link loss rates are below 10^{-6} (0.0001%), but because the core rate is higher (at 0.01%) it becomes the determining factor.

High priority packet loss for the Rivulet case is determined by the bit error rate of the links in the connection, since no queuing loss occurs. This value is estimated here at 10^{-8}.

The best effort loss rate is shown in line 6 of Table 2. Note that the best effort loss rates have improved with the introduction of QoS (compare to line 2 in Table 1), because a higher percentage of the traffic is well-behaved voice traffic, and total utilization has not increased. The Rivulet case performs better for best effort loss than the QoS case, because Rivulet does not create as much variation in its bandwidth usage as described earlier.

<u>Case B</u>:

A second approach to managing the introduction of high priority traffic is to continue to support the best effort traffic at its current loss level. This scenario is important where mission critical data applications are using the link, and their transaction performance is important. With this approach we have to reduce best effort traffic slightly as the high priority traffic is introduced, but not as much as in Case A. Overall link utilization should rise.

Table 3 - Maintaining the Same Best Effort Loss Rate

	Low Latency QoS			*Rivulet*		
Initial Link Utilization	30%	40%	50%	30%	40%	50%
High Priority Utilization	6%	8%	10%	6%	8%	10%
Best Effort Utilization	28.5%	36.5%	44%	29%	38.5%	46%
Total Utilization	34.5%	44.5%	54%	35%	46.5%	56%

Table 3 shows the resulting utilization values for this scenario. Best effort traffic loss values remain at the levels shown in Table 1. Line 1 of Table 3 shows the link utilization before introduction of QoS traffic. Line 2 of Table 3 shows the link utilization of high priority traffic, the same values as in Case A (Table 2). Line 3 of Table 3 shows the link utilization levels of best effort traffic, adjusted to maintain loss levels equivalent to those experienced with no high priority traffic. And line 4 of Table 3 shows the new total utilization levels. Again note that the Rivulet solution performs slightly better than the priority QoS solution.

<u>Case C</u>:

Lastly we consider the case where the enterprise customer wishes to maintain existing levels of best effort traffic, introduce new high priority traffic, and let the best effort packet loss rate degrade. This scenario is typical of an enterprise introducing a new application that requires high quality service, such as voice, video or TDM emulation. We use the curves here to determine the amount of loss degradation on the best effort traffic.

Table 4 shows the initial link utilization in Line 1, and the initial best effort loss rate, before the introduction of QoS traffic, in Line 2. Lines 3 and 4 show the high and low priority link utilization amounts and the total link utilization appears in line 5.

Table 4 - Maintain Existing Best Effort Traffic Levels

	Low Latency QoS			*Rivulet*		
Initial Link Utilization	30%	40%	50%	30%	40%	50%
Initial Best Effort Loss	0.05%	0.23%	0.80%	0.05%	0.23%	0.80%
High Priority Utilization	6%	8%	10%	6%	8%	10%
Best Effort Utilization	30%	40%	50%	30%	40%	50%
Total Utilization	36%	48%	60%	36%	48%	60%
Best Effort Loss Rate	0.28%	1.6%	6.4%	0.24%	1.2%	4.8%

Line 6 of this table shows the resulting best effort loss rates using this approach. The Rivulet solution again shows an advantage – it has a smaller effect on the loss rate of the best effort traffic as high priority traffic is introduced.

Because all three cases above use the same amount of high priority traffic, the high priority loss rate will remain the same for each case.

Summary

Our study has shown that the addition of high priority traffic to a network will cause the loss profile of best effort traffic to increase. Low latency QoS causes a significant impact on best effort traffic, whereas use of the Rivulet technology causes a lower impact as shown in the charts above. The Rivulet impact is 35% less than the impact of low latency QoS.

To improve the loss profile of best effort traffic, utilization levels must be reduced. Higher overall utilization amounts, or better loss rates for the best effort traffic can be achieved by using Rivulet technology. Rivulet is able to carry very high quality, low loss, low jitter traffic streams and has only a moderate impact on the remaining best effort data.

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

The NetForecast Simulator

NetForecast has built a sophisticated software-based simulator to model the behavior of data and voice traffic in varying network environments. The simulator is implemented in visual basic and uses Microsoft Excel to input parameters and display results.

The simulation engine implements three levels of traffic in a priority sequence. The higher priority queues are always emptied before allowing lower priority traffic to proceed. Queue simulation is done for a simulation interval, during which the simulator calculates how many packets it is possible to forward based on link speed and average packet length. If the momentary demand (queue input) exceeds the ability of the queue to forward packets, queue depth increases until its limit is reached. Further demand results in lost packets.

The traffic sources for the simulator come from packet captures of real applications in the field. Each packet capture represents the network usage of a single user. Specifying the number of users creates traffic demand. The simulator then determines how many users start a transaction during each simulation interval, using an algorithm designed to emulate the random and bursty nature of user demand.

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

Additional information on managing and improving application performance is available at:

www.netforecast.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 a leading authority on Internet traffic and performance. He has contributed to the design of more than 100 networks. Peter led the project that divided the Arpanet into multiple networks in 1984, which was the beginning of today's Internet. He can be reached at peter@netforecast.com.