New (And Different) Approaches To Bandwidth Management

John Bartlett

QOS requires bandwidth management, which the DiffServ standard does not provide. Four new companies are helping to solve this problem.

Enterprises are rushing to implement quality of service (QOS) in their networks, to support voice, video, interactive applications and mission-critical tasks. The cost savings of running a converged network, and the rapid transition of the telephony market from PBXs to IP-based voice are pushing enterprises down this converged network path. Bandwidth management of the high-priority traffic is a critical, complex and not well understood component of a successful strategy.

In this article we take a look at four young companies that are directly addressing this bandwidth management issue, and tackling it from four quite different perspectives. Look for these concepts to work their way into standard network deployments over the coming years.

John Bartlett is vice president at NetForecast, where he focuses on application performance, supporting real-time traffic, and quality of service (QOS). NetForecast provides consulting to enterprises and networking equipment vendors on application performance issues and convergence of voice and video conferencing on the IP network. John can be reached at john@netforecast.com.

The Problem

Implementing QOS means a lot more than turning on the QOS features in network routers. The four components required to ensure your network will reliably transport high-priority traffic are classification, priority mechanisms, bandwidth management and monitoring.

When we enable IEEE 802.1p/Q in our switches or DiffServ in our routers, we are creating a *priority mechanism. Classification* is the job of marking those packets that deserve high priority, and it is done by either the endpoint (phone or video codec), or by the edge router. *Monitoring* means constantly testing the network for real-time traffic support by either watching call quality, or by generating synthetic test traffic with agents and a server.

That leaves bandwidth management.

If you ask the network team about bandwidth management, they will tell you that high priority traffic is assigned a limited bandwidth on each network link to protect the network from over-utilization by high-priority streams. This is important in two ways.

It is important that lower-priority traffic, which is often the bread-and-butter data transactions of the enterprise, continue to have sufficient bandwidth to support adequate application performance. Secondly, it is important that high-priority traffic not become more than about 35 percent of a link's total capacity. Standard priority mechanisms don't work well beyond this value. So having these limits in place is good.

If you ask the application teams about bandwidth management, they *should* tell you that they are managing their demand to the available bandwidth provided by the network. However, as a consultant, this is often where I find a disconnect.

Managing demand from the application side means that application systems have to understand the network topology and the high-priority bandwidth limits assigned on each link. Some voice call managers or video gatekeepers have a static view of the network topology, and track bandwidth used. But topology can change based on link failures, costing changes or network upgrades. Voice and video systems usually don't share their information, so either inefficiencies occur due to static allocations, or conflicts occur when links are overutilized.

While voice and video streams have a predictable bandwidth, data application streams are more difficult. A data application often has a modest average bandwidth utilization, but causes bursty peaks of traffic when transactions between client and server devices occur. Supporting these bursty bandwidth peaks is critical to maintaining application performance.

The four companies discussed in this article are each tackling the issue of determining how much bandwidth is required, and vetting the network for the availability of that resource. But each company approaches this problem from a very different view. Let's take a look at how each com-

If Rivulet's solution sounds like TDM to you, you're not far off

pany approaches this bandwidth management issue, and the value that they bring to this tough network problem.

Corvil

Corvil is a new company based in Dublin, Ireland. The team at Corvil recognized that the understanding of bandwidth use needs to be much more dynamic than the usual link utilization statistics that most enterprises use. Momentary congestion events are barely noticed by the usual statistics, and yet they can have a profound impact on the performance of both transaction- and real-time-based applications.

Corvil has developed a monitoring tool that watches the network behavior at a millisecond level, and is able to capture that information through a mathematical model that does not require huge volumes of storage. The Corvil tools provide insight into bandwidth use behavior, help determine appropriate link sizing and aid in QOS design.

Watching bandwidth use behavior shows the IT professional how bursts of data traffic are affecting queues and link utilization in critical portions of the network. The key to this understanding is to view bandwidth demand both at macro and micro levels. Rare momentary congestion events have little effect on applications, and are expected in a packet-based network. But regular or frequent congestion indicates that applications are suffering performance degradation and link utilization is not being maintained efficiently.

Properly sizing a WAN link requires knowing how much bandwidth each application using that link requires. Under-sizing the link creates performance problems, while over-sizing the link wastes money. Through their capture and mathematical model approach, Corvil is able to determine an appropriate bandwidth for applications using a connection by watching application behavior.

This same sizing function also helps when creating a complex QOS policy to support a range of applications. Mixing voice, video, interactive applications, business critical applications and background tasks on the same link and getting good performance for all of them means creating the right prioritization and bandwidth allocation for each application class. The Corvil tool provides insight into appropriate bandwidth values for each class, allowing better optimization of the limited WAN resource.

Rivulet Communications

Imagine a network in which high-priority packets are spaced out in time such that when they enter a potential bottleneck link, they never collide. Their arrival is always timed so that only one packet at a time requires the use of the link. No congestion, no jitter, and no packet loss because the queue is always empty and the link is available when the high priority packet arrives. This approach is implemented by Rivulet Communications, based in Portsmouth, NH.

Rivulet accomplishes this by synchronizing the sending of packets from the edge of the network. Rivulet's solution is implemented either in an appliance that lives near the source of the real-time traffic, or with software that resides on the phone or videoconferencing endpoint itself.

All Rivulet edge devices are synchronized to a common time source to create a common timing window. Within that window, each edge device chooses a timing slot in which it sends its packets across the network. By ensuring that each edge device chooses its timing slot properly, Rivulet can guarantee that packet streams never enter the same queue at the same time.

If you think this system sounds a lot like time division multiplexing (TDM), you're essentially right—yet it complies with all the IP protocols.

Priority for the Rivulet-timed packets is provided by the normal priority mechanisms in existing routers and switches. To use the Rivulet approach, the highest priority available in the network must be assigned only to Rivulet based traffic. Rivulet endpoints mark the packets for high priority so routers will recognize them and give them the high priority they require.

So priority is handled in a familiar way. But how does this approach solve the bandwidth problem?

Rivulet does not help if the demand for carrying real-time traffic exceeds the capacity of the network to carry it. The Rivulet method will recognize that this has occurred, because it will not be able to find an open time slot for the traffic stream that pushes the network over the available limit. However, Rivulet does help in the case where bursty traffic can cause loss or jitter, and in the case where we would like to carry more than 35 percent of a link's capacity in real-time traffic. Let's take a look at these two cases.

We think of real-time traffic as being constant-bit-rate. A voice codec sends out one packet every 20 or 40 milliseconds for the entire duration of the call. Compressed voice and video are not quite as regular, because the compression algorithm can increase or decrease the size of the packets based on the amount of information it needs to send. But regardless, this is not the effect that can cause queues to fill up.

In a larger network, where we have many real-time sources, each source sends its packets out according to its own internal clock. These clocks are not synchronized, and some may be slightly faster or slower than others. So some of these endpoints can be sending out their packets such that they arrive at a particular router output queue at the same time. There is a statistical probability that many packets from different endpoints will arrive simultaneously, causing the router queue to be slow, which creates jitter, or to overflow, which causes packet loss. By synchronizing each endpoint in the network, Rivulet prevents this occurrence.

The second case is where we want a link to carry more than 35 percent of its capacity in real-time traffic. Thirty-five percent is the recommended limit for priority traffic on a link because a normal priority mechanism starts to degrade at about this level. Degradation actually occurs on a curve from 0 to 100 percent, but we draw the line at about 35 percent. We know priority works well for a small percentage of high-priority traffic. We also know that if 100 percent of the traffic is high-priority, then there is no priority, since all traffic is treated the same.

Using the Rivulet technology, we don't run into this 35 percent limit. Again because each packet entering a high-priority output queue on a router finds that queue empty, and is forwarded as the very next packet to go because it is high priority, the 35 percent limit does not affect these streams. Much higher ratios of real-time traffic to data-traffic are possible using the Rivulet technology, than are possible with a standard priority approach.

Prominence Networks

The team at Prominence Networks, based in Holmdel NJ, recognized the same dilemma described above, and is addressing it via the two main IETF QOS approaches—IntServ and DiffServ.

IntServ or the Resource Reservation Protocol (RSVP) provides all the priority and bandwidth management functions needed to provide good quality-of-service transport, but does not scale well in the core of the network. But the solution, DiffServ, gives up the concept of built-in bandwidth management and hands that function back to the application, which you may or may trust to make that decision properly. So Prominence Networks set out to provide solutions to both these problems.

Prominence provides a network appliance, the Clear Call Controller, that connects near the edge of the network, and aggregates media flows from voice and videoconferencing endpoints. The Clear Call Controller converses with the voice call server or the videoconferencing gatekeeper to become a media proxy between corporate locations. When endpoints request a connection, the call server or gatekeeper directs their media flows to the Clear Call Controller closest to their location. At the far end of the connection, another Clear Call Controller unit is dispersing the flow to the local endpoint, and collecting the packets of the reverse-direction media streams.

Between the Clear Call controllers, Prominence is able to leverage either the IntServ or the DiffServ protocol through the corporate network. Let us look first at the IntServ approach.

The key issue with IntServ is that the router must maintain state information about each flow passing through that router. This information includes its source and destination, port numbers, bandwidth request and current bandwidth utilization. The connection must be refreshed regularly to ensure it is still active, and quality information about the flow must be passed back and forth between the endpoints. This is not much work for a few media streams, but if the router is in the core of the network, it could be required to support thousands of streams. In service provider networks, these numbers are even larger.

Prominence uses IntServ to manage bandwidth

To avoid this scaling issue, the Prominence approach is to open a single IntServ stream between each of the Clear Call Controller appliances. The bandwidth requested for this connection is scaled up and down based on the number of phone calls or videoconferences that need to use that path. In this way a much smaller number of flows is handled by the core routers, and just updated periodically during the day to reflect current usage. If insufficient capacity exists in the network to support the voice or video requests, the Clear Call Controller system signals to the call server or gatekeeper, and the call is either denied, carried without QOS or rerouted via the PSTN, depending on predefined policies.

If the network between Clear Call Controller systems uses DiffServ for priority, Prominence takes a different approach. With DiffServ, the network pre-allocates a limited amount of bandwidth to each defined priority class. The bandwidth management problem is to ensure that the combination of voice and videoconferencing across any one link of the network does not exceed this predefined capacity.

To support this environment, the Clear Call Controller appliances rely on a centralized Clear Call Director appliance, which maintains a complete view of the network topology. The Clear Call Director product understands the network topology, the current routing paths, and the bandwidth available for real-time traffic on each link. This information is collected and updated automatically by querying routers.

When new media flows request a connection through the Clear Call Controller appliances, the Clear Call Director maps that flow against the network topology, counts up the bandwidth already allocated on those links, and determines if sufficient bandwidth exists to support this call. Like the IntServ case, if there is insufficient bandwidth, the Clear Call Director signals the call server or gatekeeper and the call is redirected.

Xelor Software

Xelor Software is an early-stage company in Salem, NH, with a technical team in Western Australia. Xelor provides a software solution that converses with the voice call server, and uses the capabilities of the existing network routers to manage the bandwidth of voice traffic.

The Xelor application initially scopes out the

Xelor's solution talks to the routers and the call server

network by reading SNMP MIBs from the routers. The collected information allows Xelor to build a complete understanding of the network topology, routes, and available link bandwidths. This information is updated over time, allowing the Xelor application to always know how voice traffic is flowing through the network.

Xelor then configures QOS capabilities in each router and switch. Using industry best-practices, the Xelor software enables QOS, creates policy maps and assigns bandwidth limits on each router that will support real-time traffic.

Next, Xelor connects to the call server, where it is able to learn when a call is being established. When a call is started, Xelor maps that call to the network topology, and determines if sufficient high-priority bandwidth exists along that path. If the answer is yes, Xelor programs the two edge routers involved in the call with an Access Control List (ACL) entry. This ACL allows the traffic from those two endpoints to be given high-priority treatment.

Xelor adds the bandwidth being used to its database of active calls and available bandwidth, so it can keep track of how much bandwidth remains available for the next caller. When a call ends, Xelor reverses the process by tearing down the ACL and removing that call's link utilization from its accounting of bandwidth in use.

So Xelor has addressed the bandwidth management issue by watching and accounting for each call. Calls are only given high-priority treatment if sufficient bandwidth exists to support it. Thus the voice traffic demand for high priority is kept within bounds, and those calls using high priority will get good network transport, resulting in high-quality voice connections.

Summary

The IETF simplified the task of supporting QOS in core routers when it defined DiffServ, by removing the bandwidth management requirement inherent in IntServ (RSVP). But the problem did not go away.

New solutions are being created to address this problem, and it will be interesting to see how these solutions evolve, which become standards, and which best address the issue of balancing bandwidth use and performance in our networks□

Companies Mentioned In This Article

Corvil (www.corvil.com)

Prominence Networks (www.prominencenetworks.com)

Rivulet Communications (www.rivulet.com)

Xelor Software (www.xelorsoftware.com)