Products Product Literature

Using QoSWorks™ to Manage SNA

Business requirement:

A large institution with 60 remote facilities has made a strategic commitment to IP. They have an IBM S/390 mainframe at headquarters; clients access the mainframe using TN3270 and TN5250 emulation software. The company is using Datalink Switching (DLSw) to encapsulate the SNA traffic within the TCP/IP protocol. In addition to SNA, the enterprise network carries SMTP e-mail and Web traffic between the remote locations and headquarters. The network manager is concerned that the mission-critical SNA traffic will time out when its flows are impeded by the bursty e-mail and Web traffic.

Network topology:
The 60 remote facilities are linked to headquarters via a Frame Relay network with port speeds of 64Kbps. The permanent virtual circuit committed information rate (CIR) for each location is 50% of the port speed, or 32 Kbps.

Application requirements:

SNA applications are designed to work on low-speed lines (below 10Kbps). These applications are delay sensitive and will time out if a response is not received within a couple of seconds. Keep-alive timers (RR polls) used in SNA applications help determine if end devices are still in session. Typically, these timers are set in the 1-3 second range.

When encapsulated SNA flows are mixed with SMTP or Web flows, the average delay experienced by the SNA traffic will go up significantly. This happens for two reasons. First, SNA applications often experience delays because TCP-based applications such as e-mail and Web browsing are bursty in nature. For example, e-mails with large attachments or graphics-intensive Web pages will cause a burst, in which long packets travel very close to each other and queue up in the router. This leaves little room for SNA packets, which in the worst case will be dropped, and in the best case will experience a delay.

The second reason that SNA applications experience delays is that the average packet size of the SMTP or Web transfer is large (usually 1500 bytes), and the larger the packet, the longer the router will take to send it. For example, the amount of time required to send a 1500-byte packet on a 64Kbps line is 187ms. If an IP packet carrying SNA data gets in the queue behind 10 SMTP or Web packets, the SNA packet will have to wait 1.8 seconds. The acknowledgement packet will have to wait the same amount of time on its way back to the branch. The total amount of time - over 3 seconds - is long enough to cause problems with the SNA application.

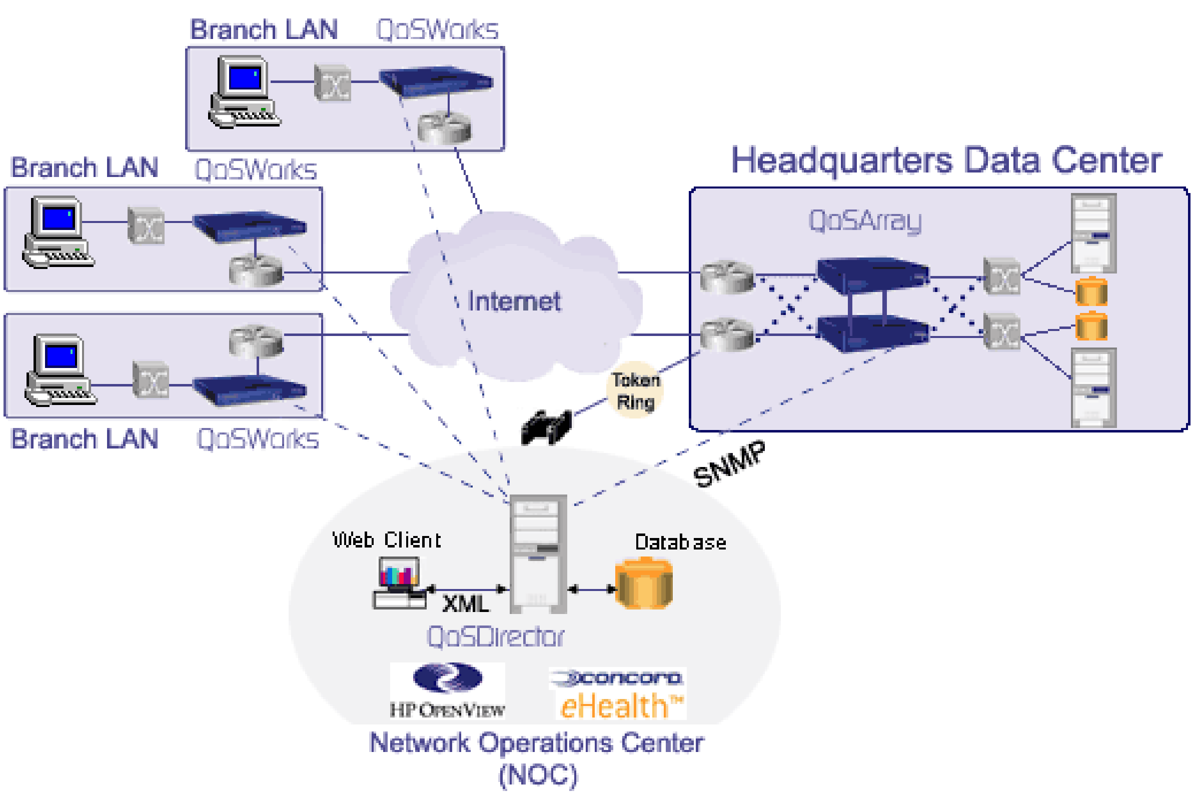

Fig. 1 Deploy QoSWorks behind branch routers to ensure improved, consistent response time for SNA-based applications.

The QoSWorks Solution

In order for the SNA application to function optimally, the network manager must control the latency experienced by the application and ensure that SNA packets do not get crowded out by large e-mail and Web packets. Only QoSWorks integrates all of the different mechanisms required to accomplish these objectives while at the same time providing sufficient resources for mail and Web traffic and increasing overall network efficiency.

Wire-speed classification:

DLSw commonly uses TCP ports 2065 and 2067. QoSWorks uses these ports to identify SNA applications. Network managers can assign high priority to all flows whose destination port matches 2065 or 2067. The e-mail and Web traffic is classified by the destination port number (e.g., 25 for SMTP and 80 for the Web).

Class-based queuing:

Class-based queuing (CBQ) provides the means of implementing traffic priority and bandwidth sharing. The network manager uses CBQ to place the DLSw, SMTP and Web traffic in three different classes according to destination port. The DLSw traffic is assigned priority over SMTP and Web traffic with the SMTP traffic prioritized over the Web traffic. This allows the DLSw flows to borrow excess bandwidth when necessary from the SMTP and Web. This minimizes the latency of DLSw traffic relative to e-mail and Web traffic.

CBQ is an excellent method of controlling application latency. However, CBQ alone cannot guarantee that SNA packets will never be queued behind large e-mail or Web packets, leading to delays. In addition, CBQ cannot guarantee that each connection within an application receives a fair share of bandwidth. These require two unique QoSWorks features developed by Sitara: packet-size optimization and fair allocation of bandwidth by connection.

Packet-size optimization:

Minimizes delays for SNA traffic by reducing the size of e-mail and Web packets. For example, it takes an average of 187ms to transmit a 1500-byte packet over a 64Kbps channel. This means that a high-priority packet, like DLSw, would have to wait 187ms before it is sent. Packet-size optimization solves this problem by forcing TCP connections with large packet sizes, like e-mail and Web, to reduce the maximum packet size to about 512 bytes. As a result, a DLSw packet only has to wait 64ms before being transmitted, for a round-trip time of 128ms, an acceptable level of delay.

Fair allocation of bandwidth by connection:

If there are multiple DLSw flows, each one can be assigned a minimum bandwidth (e.g., 10Kbps) using an algorithm created by Sitara that ensures fair allocation of bandwidth by connection. This will guarantee that each connection within the application does not time out.

TCP rate shaping:

In order to ensure a high quality of service for the SNA application, the network manager must not only prioritize DLSw flows but also control bursty TCP/IP-based e-mail and Web traffic. This requires a mechanism called TCP rate shaping. By manipulating the window sizes of SMTP and Web connections, TCP rate shaping helps manage the burstiness of the SMTP and Web traffic to ensure that it does not interfere with time-sensitive DLSw SNA traffic.

Clearly, TCP shaping should be used in concert with other QoS mechanisms such as queuing, as it controls only TCP/IP-based traffic. Also, it cannot ensure a consistent rate for applications (such as SNA keep-alive polls) where the amount of data sent by the server is small and the rate is low. This is where QoSWorks CBQ feature is particularly beneficial.

Caching:

In order for critical applications like SNA to function effectively, network managers must use every byte of available bandwidth as efficiently as possible, especially over low-speed links. This means eliminating redundant requests for Web objects. QoSWorks caching feature stores frequently accessed Web pages to reduce traffic load on the WAN and free up more bandwidth for the DLSw and e-mail traffic. Typically, 30-40% of Web content can be cached locally, leading to significant savings in bandwidth utilization.

991202-011-01