MPLS Quality of Service (QoS) Offering

A technical explanation from SDN Communications

SDN's MPLS network offers end-to-end quality of service (QoS) capabilities. The charge for this service is in one megabit increments. QoS is activated on an individual circuit level and is only active during times of congestion on the circuit.

SDN offers three queues: Real Time, Mission Critical and Best Effort. SDN will need to know from the customer how the traffic is being marked. The customer markings can be DSCP (Differentiated Services Code Print) and/or IP Precedence. SDN will honor these markings and place the traffic into the appropriate queue. SDN collaborates with the customer to determine these markings and what traffic type gets placed into the corresponding queue. SDN also works with the customer on determining the size of each queue.

The amount of bandwidth purchased for QoS is distributed amongst the Real Time and Mission Critical queues and the remaining circuit's bandwidth is allocated to the Best Effort queue. For example, if a ten megabit Ethernet circuit is purchased with one megabit of QoS traffic, the one megabit of QoS can be evenly dispersed with 500 kilobits for Real Time and 500 kilobits to Mission Critical, or any combination up to the purchased value. In this previous example the entire one megabit could be allocated to the Real Time queue. The remaining nine megabit (non-QoS purchased bandwidth) is automatically allocated to the Best Effort queue.

The Real Time queue is a strict priority queue. When this queue is active it will be serviced until the queue is completely empty. Once the Real Time queue is empty the remaining two queues will be serviced. Due to the Real Time queue having a higher precedence and being serviced until it is completely empty, it is generally used for time sensitive applications like voice and video. This queue has a "hard ceiling," meaning when this queue is active it cannot take available bandwidth from the remaining two queues. While this queue is active and if the customer's router attempts to send more bandwidth than what is currently allocated, the excessive bandwidth is discarded. For example, if the size of this queue (hard ceiling) is set at 500 kilobits and the network needs to send 600 kilobits of real time data, the excess of 100 kilobits is completely discarded.

Continued on page 2

Continued from page 1

The Mission Critical queue is a class-based queue. A class-based queue allocates a specific amount of bandwidth per the agreed upon markings. During times of congestion, traffic that meets the agreed upon markings is guaranteed this minimum amount of bandwidth. The Mission Critical queue can use available bandwidth from the Real Time and Best Effort queues, thus it is possible this queue can transmit above the minimum defined value. For example, if 500 kilobits are allocated to the Real Time queue and 500 kilobits to the Mission Critical queue and there currently isn't any traffic in the Real Time queue, the Mission Critical queue can use the entire one megabit of QoS purchased bandwidth. In this example the Mission Critical queue can also take from the Best Effort queue if excessive bandwidth exists. However, traffic taken from the Best Effort queue isn't guaranteed a specific amount of bandwidth. Only traffic within the Real Time and Mission Critical queues are guaranteed. Excessive bandwidth taken from the Best Effort queue will not be discarded if the bandwidth is available, it just won't be guaranteed.

The Best Effort queue treats all data allocated to this queue the same. The queue is a "catch all" for traffic that doesn't adhere to the Real Time and Mission Critical markings. The customer should ensure the Best Effort queue is large enough to provide service to applications that aren't assigned to the Real Time and Mission Critical queues.

SDN doesn't mark-down traffic that exceeds the specified queue sizes. In the event of excessive traffic in the Real Time queue it is completely discarded. In the situation of exceeding the traffic size of the Mission Critical queue, the Mission Critical queue can "steal" bandwidth from the other queues if it exists. If no free bandwidth is available from the other queues, the Mission Critical queue will drop any bandwidth needs that exceed its specified value. Due to these reasons it is important that the sizes are chosen correctly and periodically monitored by the customer for dropped packets within the queues.

Depending upon the MPLS circuit type that the customer purchases, QoS is handled differently. If the customer purchases a layer 3 VPN, QoS is very granular. It is granular down to the individual IP flow. With a layer 3 VPN and the corresponding DSCP and/or IP Precedence markings, any packet that is tagged with the agreed upon markings will be treated independently. This allows the customer to give preference down to an application level, or for that matter down to an individual IP address. If a layer 2 VPN is purchased, SDN can only give preferential treatment to the Ethernet VLAN or Frame Relay/ATM private virtual circuit (PVC). With a layer 2 VPN, SDN is unable to give preferential treatment down to an individual IP flow. However, multiple layer 2 circuits (VLAN/PVCs) can be provisioned on the same physical circuit and preference can be given to one VLAN or PVC versus another.

Questions about QoS? Contact Gary Fischer at 800-247-1442 or your account executive.