



Ontario's Gateway to Research and Discovery

QoS Support on ORION

Version:	DRAFT 2.0
Date:	October 28, 2008
Written By:	Sam Mokbel Senior Director, Engineering
Classification:	General Distribution

The new ORION routing architecture is capable of supporting QoS, MPLS, and Multicast in addition to IPv6 and IPv4 applications and services. Current members' traffic can be classified into two types: general use best effort, or jitter and loss sensitive. The overwhelming majority of most members' external traffic would belong to the general use best effort type. Traffic that may be sensitive to jitter, delay and packet loss include video conferencing, remote classroom, and intra institutional IP telephony between campuses. Since none of the ORION circuits is experiencing any congestion both traffic types experience similar quality performance on the network.

ORION currently does not recognize, characterize, or differentiate between traffic types due to three reasons: The first reason is related to the service pricing model which is membership based as opposed to bandwidth usage based. Secondly, the network is based on a minimum hop design and operated on a congestion free basis so prioritizing packets has no material effect on traffic flows and application performance. Thirdly, none of the members have so far requested a specific performance related customization on their IP traffic.

Members do not currently classify their external traffic into any class and therefore the traffic defaults into the general best effort scheme implemented on ORION, the R&E global network, or the public Internet. Applications such as remote education and telehealth, voice and video over IP, and measurements and sensory data monitoring are becoming more prolific and require more deterministic transport characteristics than generally available on public networks. ORION is currently well suited for the delivery of such applications and will continue to support them by providing customized service when required. ORION Engineering will maintain a low hop count streamlined network design operated as congestion free with low jitter and small buffers. However, as the number of members and network applications increases some may require a highly controlled and guaranteed packet delivery environment.

During its April 5th, 2007 meeting, ORION's Technical Committee discussed ways to support members' QoS based applications. The Committee agreed on ORION supporting two service levels for members' IP traffic and one level for internal network management. The traffic types are Premium IP, Best Effort IP, and Internal Network Control. Support for the Premium IP traffic will consist of supporting two sub-types one is focused on Expedited Forwarding and the other is oriented towards Assured Forwarding of packets. The Committee also recommended that ORION does not classify, mark, and police its members traffic and only support a QoS application when requested.

For the foreseeable future, the requirement for a deterministic and guaranteed packet delivery is expected to be very limited in terms of origin/destination pairs and number of applications and users. Under these circumstances the best approach to supporting QoS sensitive applications between members is by transporting their selected traffic on ORION custom MPLS VPN services. This approach also integrates well with members internal networks since most of them use VLAN/VPN type transport mechanisms to

differentiate between their various traffic types. A member that requires the delivery of QoS enabled application across ORION to one or multiple destinations should associate the traffic with a VLAN/VPN and hand it over to an ORION VLAN/VPN at the existing ORION/Member interface at the PoP. The ORION VPN will have the required QoS characteristics and will be transported accordingly to the destination ORION/member interface. Marking will be done by the source network and will be trusted by ORION and transported as is to the destination end of the VLAN/VPN. The QoS parameters required by the application will be associated with the VLAN/VPN and all ORION network equipment will handle that VPN traffic accordingly. ORION will not establish trust boundaries at which it will mark all incoming packets. All traffic reaching ORION at the interface, with the exception of the agreed on VLANs/VPN, will continue to be treated on a best effort basis and routed accordingly.

Processing of QoS Sensitive Traffic

Network devices have buffers that temporarily hold packets as they arrive and get processed. Congestion occurs when packets arrive at a device faster than they can exit. To meet predefined service levels, queuing and congestion avoidance algorithms manage packets as they arrive or wait inside the buffers according to their priority levels and arrival rate. These algorithms, when processing the packets, act on special QoS or traffic flow header information.

ORION's core routers are capable of selectively handling and processing traffic by checking source/destination addresses, IPP bits, DSCP bits, 802.1Q/p CoS, or the MPLS EXP bits of each data packet received.

Traffic Scheduling Tools

Various scheduling tools are used by the core routers to enforce low latency or guaranteed throughput per service. Here is a description of some of these tools that have common industry names and functions:

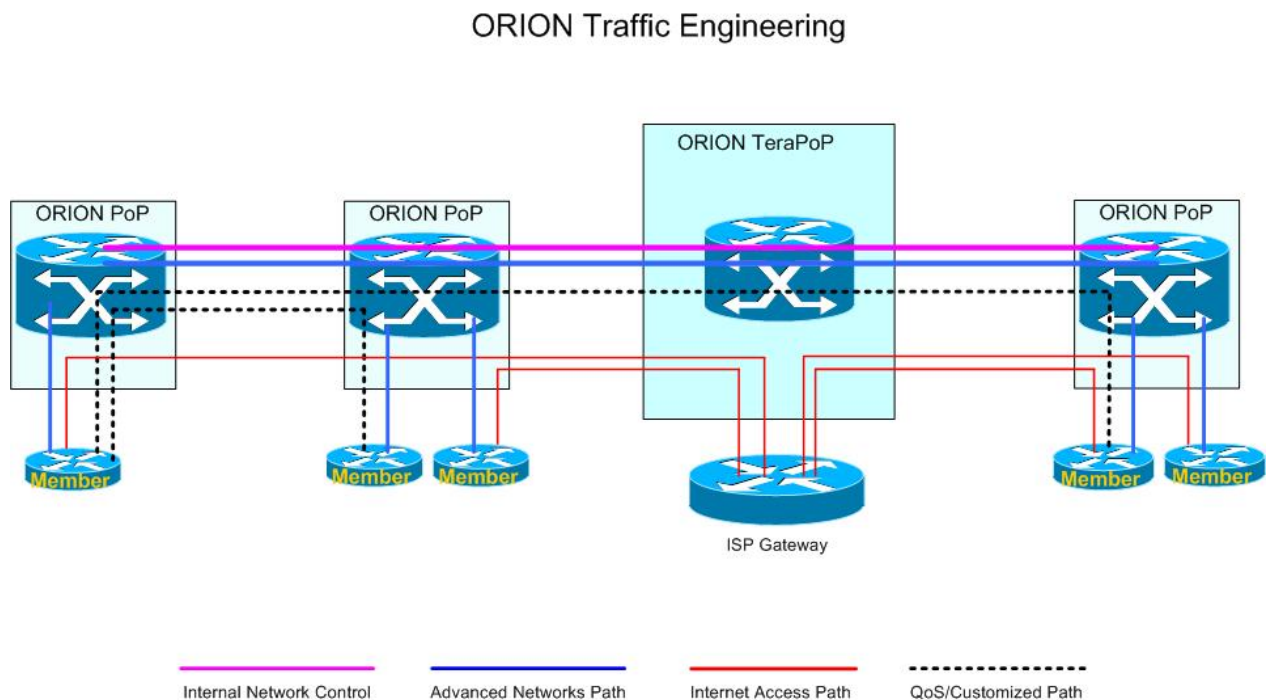
CAR: Committed Access Rate. CAR limits the input or output transmission rate on an interface or sub-interface based on predetermined criteria. CAR is usually configured on interfaces at the edge of a network to limit traffic in or out of the network. Traffic is rate limited based on certain matching criteria, such as incoming interface, IP Precedence, or IP access lists. The network manager configures the actions CAR should take when traffic conforms or exceeds the rate limit.

WRED: Weighted Random Early Discard. WRED monitors network traffic, trying to anticipate and prevent congestion at common network and inter-network bottlenecks. RED can selectively discard lower priority traffic when an interface begins to get congested. It can also provide differentiated performance characteristics for different classes of service.

CBWFQ: Class Based Weighted Fair Queuing. WFQ is an automated scheduling system that provides fair bandwidth allocation to all network traffic. WFQ uses weights (priorities) to determine how much bandwidth each class of traffic is allocated.

MDRR: Modified Deficit Round Robin Queuing is a low latency class-based queuing mechanism suitable for handling delay-sensitive traffic such as voice and video.

Service Types



Internal Network Control: Used for route distribution and network management and control. Most ORION network equipment is manageable via SNMP or TL1 through secure private IP or optical transport circuits that coexist with other circuits carrying member traffic. Network control traffic always receives the highest priority in terms of speed and guarantee of delivery to ensure timely and uninterrupted network management and operation. TL1 network control traffic uses dedicated wavelengths on the optical layer systems and thus is guaranteed its own dedicated best performance circuits and will not be affected by a QoS implementation. SNMP network control traffic is used for all IP routing, switching, and PoP management devices and uses dedicated VLANs and VPNs for delivery.

Premium IP: This traffic class would receive the second highest priority when processed by ORION's IP devices. Premium IP traffic would be tagged as such by the originating

member network and would be processed by ORION accordingly. ORION will not police or retag traffic. Real time applications such as Voice over IP could receive an Expedited Forwarding service. Interactive Video Conferencing for example could receive an Assured Forwarding service. Members typically configure their internal networks to transport an application's data flow as Premium service with certain QoS characteristics and if that data flow is to reach externally, would ask ORION to provide a similar set of QoS while the data flow traverses ORION. Assured Forwarding services will have a guaranteed throughput level as a minimum committed bit rate associated with the service.

Best Effort IP: All Internet and R&E data exchange would fall into this best effort service category. It is important to state that ORION is managed and operated as a congestion free and minimal hop count network which means that there is zero packet loss and minimum propagation delay introduced by intermediate routing and switching devices. The difference between the Best Effort and Premium services is that one provides a guarantee and the other does not.

Costing

ORANO will offer these Premium IP traffic services at the same price levels currently used for other Inter-member transport services. For example a 10Mbps CIR service will be priced at \$250/month, a 100Mbps CIR service at \$500/months, and 1Gbps CIR service at \$1,250. A \$1000 setup charge applies to all services. Pricing for low latency and other customized QoS services will be made available when required.