Quality of Service in Service Provider Networks

Technology Overview

Introduction

Quality of Service (QoS) has traditionally been used in connection-oriented ATM and Frame Relay networks to guarantee traffic delivery. In contrast, connectionless IP networks have addressed QoS by adding capacity as networks became congested. The expense of throwing capacity at the QoS issue and the blurring of switched and routed networks and services have now made QoS a critical requirement in all types of service provider networks.

QoS mechanisms are required to manage inevitable network congestion and guarantee applications sensitive to delay, jitter and packet loss. With QoS-enabled routers, providers can offer differentiated services, guarantee bandwidth, increase control over network resources, classify and prioritize traffic and manage network congestion in networks that were previously limited to best-effort Internet service.

QoS-enabled routers include a number of components

- ***Classifier/Marker*** to identify groups of packets to receive a specific service and carry information about packet class, drop precedence, or both. Classification can be simple (assign Gold class to packets received on a particular interface) or complex (assign Gold class to packets destined to particular address within 40 Mbps of traffic contract)
- ***Policer*** restricts some traffic to ensure subscribers don't send traffic beyond their subscribed rate
- ***Queuing*** protects and isolates traffic to make sure the most important traffic is handled appropriately
- ***Buffer Management*** defines discard policy if queue is full to ensure higher-priority traffic gets through
- ***Shaping*** controls jitter and ensures compliance with contracted bandwidth

QoS Requirements at the Service Provider Edge

The edge of service provider networks is where the customer meets the provider and a range of services and access speeds are aggregated for transport across the core. Sophisticated QoS is a must in service edge routers, to support ATM and Frame Relay traffic and enable IP-based differentiated services.

Sophisticated QoS capabilities are a function of hardware and software and cannot be retrofitted in edge routers designed for other purposes. Core router vendors have recommended the practice of redeploying core routers to the edge as higher-capacity core routers are introduced into the network. However, with a focus on capacity, core routers are not well suited to supporting the sophisticated QoS, bandwidth management and accounting functions required of a service edge router. Some core

routers attempt to implement these functions in software, but do so at the cost of performance.

In a multi-service edge router, the following types of QoS are required to deliver a complete range of data services including ATM, Frame Relay, Ethernet and IP-based services.

Wire-speed Multi-protocol QoS: Service edge routers must support ATM CBR/VBR/UBR, Frame Relay DE bit, Ethernet 802.1p and MPLS EXP QoS service.

QoS Granularity: Typical vendor equipment provides QoS on a port level. Port-based QoS is sufficient for the limited service requirements in the core, but not at the edge. To effectively provide differentiated services, carriers must provide QoS on a logical interface level. All VLANs/VCs/DLCIs/LSPs require independent treatment depending upon customer requirements.

Policing and Shaping: Ingress policing is required to support tiered services. For example, a customer may only use 200 Mbps of a Gigabit Ethernet connection. The service edge router must police down to 200 Mbps inbound and perform egress shaping outbound to the customer. This is especially critical for Internet traffic where the inbound traffic is small, but the outbound traffic (back to the customer) could use up the entire 1 Gbps of bandwidth. Therefore, the service edge router must perform policing and shaping on a per-logical interface (in this example an 802.1q VLAN) as the edge is an aggregation point for many Ethernet VLANs on a Gigabit Ethernet connection.

QoS-Enabled Service Provider Applications with the ST-series

ECI's Service Edge Router enables service providers to apply sophisticated QoS to any type of switched or routed data service. With the ST-series, a wide range of QoS-enabled applications are possible.

Scale and Bundle Existing ATM and Frame Relay Services: ATM and Frame Relay services are highly successful, yet ATM and Frame Relay equipment is running into capacity limitations. With ST-series QoS features, ATM and Frame Relay traffic can be mapped across a high-capacity IP/MPLS core to enable high-speed services as well as bundled services (such as IP-enabled Frame Relay).

Turn IP Networks into Multi-service Networks: For IP providers, the ST-series provides the QoS mechanisms to offer ATM and Frame Relay services over an existing IP/MPLS network without requiring the construction of separate Frame Relay and ATM networks. IP providers can then offer higher speed ATM and Frame Relay services as well as IP-enabled services, greatly expanding their service portfolio.

QoS-Enabled IP VPNs: MPLS IP VPNs provisioned using the ST-series include the same level of QoS as ATM and Frame Relay services, with wire-speed performance. Software-configurable per-customer queues meet SLAs while providing multiple

distinct service classes for each customer. Per-customer traffic shaping delivers precisely the purchased bandwidth. Packet classification and filtering provide differentiated services based on packet markings, IP source/destination or application. Traffic policing and marking allow carriers to create burstable services with guaranteed bandwidth. Weighted random early detection (WRED) congestion management maximizes link utilization and ensures delivery under congestion.

QoS-enabled Ethernet Service: Enterprises have long relied on Ethernet as a primary method for transporting IP-based traffic in the LAN. That familiarity has led to demand for wide area Ethernet services. The ST-series enables carriers to extend Ethernet, a widely used and understood technology, beyond the reach of local and metropolitan networks. Wide area Ethernet service takes advantage of ST-series QoS capabilities to offer ATM or Frame Relay-like QoS to customers.

Differentiated Internet Services: The ability to deliver differentiated service levels charged at different billing rates can expand service provider revenue opportunities beyond low-margin Internet access. Unfortunately, most network equipment can't provide a sensible model for delivering and accounting for differentiated services. The ST-series enables the delivery of premium Internet services by establishing connections through the network with guaranteed bandwidth, classifying and mapping customer traffic onto those connections, and incrementing corresponding billing counters.

ECI's Service Edge Router: Unmatched QoS

The ST-series provides robust QoS and traffic-management capabilities for both Layer 2 and IP traffic and follows established standards including the ATM Forum Traffic Management Specification Version 4.1 and the DiffServ model for Internet traffic.

- ***Wire-speed packet classification*** provides differentiated services based on Layer 2 classifiers such as the ATM CLP, Frame Relay DE and Ethernet 802.1p bits and Layer 3-4 IP multifield classifiers such as source/destination address and Type of Service (ToS)
- ***Standard and custom policing schemes*** enable the creation of burstable services with guaranteed bandwidth. Weighted random early detection (WRED) congestion management maximizes link utilization and ensures traffic delivery
- ***Flexible queuing strategies*** enable combinations of strict priority and weighted fair queuing (WFQ) queues, allowing further prioritization and differentiation
- ***Three levels of scheduling*** (CBR, VBR, UBR) ensure that customers receive only their purchased bandwidth – nothing more, nothing less

ST-series Queuing

Per-VC queuing is required to maintain QoS while delivering services to many customers. The ST-series supports up to 32,000 VC queues per interface card, with separate classes of service for CBR, VBR and UBR traffic so that service providers can offer premium Internet services. In addition, the ST-series supports per sub-interface output traffic shaping so rate adaptation can occur when a customer's source interface speed is larger than its destination interface speed.

ST-series queuing is performed on the egress data path using queue blocks

A queue block is assigned to a logical or physical interface. It can consist of one to eight output queues. One queue block can be configured per logical port. Service providers can implement both strict priority queuing and WFQ within a queue block.

- ***Strict Priority Queuing*** is determined by the order in which the queues are configured. Queues are processed in descending order. The first queue configured is serviced until it is empty. The remaining queues are then serviced in sequence
- ***Weighted Fair Queuing*** is a queuing algorithm that combines fair queuing and preferential weighting. The fairness aspect of WFQ functions similarly to round-robin queuing, with queues serviced in a continuously repeating sequence from top to bottom, and then starting at the top again. The weighting aspect of WFQ applies a "weight" to a queue that indicates the importance of the queue in relation to the available resources. The weight is used to ensure that more important queues get serviced more often than other less important queues. With WFQ, queues are first sorted in order of their increasing weighted value. Then, each queue is serviced in order of its weighted proportion to the available resources

For each queue within a queue block, the following parameters may be defined

- ***Weight (for WFQ only):*** A ratio (from 1-255) that indicates how the queue is to be serviced in relation to the available resources.
- ***Maximum buffer size:*** The maximum amount of buffer space (in bytes, from 0 to 16,776,960) to be reserved for traffic in the queue. This buffer prevents the traffic in a single queue from consuming an excessive amount of buffer resources.
- ***Drop precedence:*** The drop precedence that is associated with the queue. Up to two drop precedences can be configured for each queue; the priority is determined by the order in which they are configured. The drop profile associated with each drop precedence is based on a WRED algorithm.

Service Categories for Packet-based Traffic

The ST-series is the first edge router with the ability to precisely match the service **guarantees of existing ATM, Frame Relay or Ethernet services.**

Real-Time Variable Frame Rate (rt-VFR): The highest service category, rt-VFR specifies low delay, low delay variation and low frame loss service. The rt-VFR service category is best suited to real-time traffic.

Non Real-Time Variable Frame Rate (nrt-VFR): This medium service category provides a guaranteed loss ratio for Frame Relay traffic. The nrt-VFR category is typically used for delay-tolerant, bursty traffic.

Unspecified Frame Rate (UFR): This is the lowest service category, providing no rate guarantee. UFR is generally used for non-critical delay- and loss-tolerant traffic.

The traffic in these categories is strictly prioritized, meaning that rt-VFR is always serviced first, followed by nrt-VFR, and finally UFR.

Scheduling Modes

Scheduling modes determine how the ST-series services traffic from logical paths. There are two scheduling modes: fair-share and shaped. Fair-share scheduling allows the ST-series to service logical paths within a service category equally, regardless of bandwidth requirements. For example, if you have three logical paths in a service category, the ST-series takes a packet from the first interface's queue, then the second interface's queue, then the third, and so forth, until the queues of the logical paths in that service category are empty. Then the ST-series services the next priority service category.

Shaped scheduling allows the ST-series to service logical paths within a service category according to a peak bandwidth setting by applying a shaping parameter that sets the peak amount of bandwidth available to each logical path. The interface receives no more bandwidth than what is configured for its peak bandwidth setting. For example, if there are three logical paths in a service category (one with 10 MB of bandwidth, one with 5 MB of bandwidth, and one with 1 MB of bandwidth) the ST-series services 10 MB of packets from the first interface, 5MB of packets from the second, and 1 MB of packets from the third. Then the ST-series services the next priority service category. Both rt-VFR and nrt-VFR service categories can be individually configured for either shaped or fair-share mode. These modes are configured per physical port. UFR is always configured for fair-share mode.

Applying QoS Profiles on the ST-series

Service profiles define a way to tag incoming packets through the ST-series. Service profiles are mapped to the packet fabric VoQ (Virtual Output Queue) and the output queues are set per-interface to mark incoming traffic.

To provide QoS-enabled services, service providers should assign multiple service profiles that are mapped to different traffic characteristics. Once a service profile is created, traffic is assigned to it via Layer 2 classifiers, Layer 3 classifiers, policy-based classification, incoming logical interfaces and policers.

Using the ST, a service provider can configure up to eight service profiles. In the example below, six service profiles are set up. Five are for customer traffic, with the sixth used for network control traffic.

- ***Control*:** The highest priority, supported via a strict scheduling mechanism, is reserved for network control traffic, such as routing data and MPLS protocols. This service profile includes any message traffic where a failure to deliver that traffic could affect the operation and stability of the network
- ***Platinum:*** A second strict queue should be used for emulating circuit switching. This queue, which supports the Platinum service profile, is reserved for voice and video traffic, as well as any other traffic which requires guaranteed service with low levels of latency and jitter. Examples include ATM, Frame Relay or Voice over IP traffic
- ***Gold+:*** This could be used for applications including ATM VBRrt and realtime video
- ***Gold:*** This could be used for traffic requiring low jitter, low max delay and guaranteed delivery
- ***Silver:*** This could be applied to traffic such as wide area Ethernet and ATM UBR. Characteristics include high max delay, high jitter and guaranteed delivery
- ***Best Effort:*** This is used for high jitter and Internet traffic. Characteristics include high max delay, high jitter and lowest priority

Service Name	Service Profile	Queue Order	Application	Q Depth, Drop Strategy	Servicing WFQ or Strict	Characteristics
Control	0-7	0	Routing/MPLS protocols	Xms/hop tail drop or WRED	%, WFQ or strict	Control traffic queue
Platinum	0-7	1	ATM/FR CBR, VoIP	Xms/hop tail drop or WRED	%, WFQ or strict	Lowest jitter, no delay, guaranteed delivery
Gold +	0-7	2	ATM VBRrt CLP, Frame DE, realtime video excess tagged	Xms/hop tail drop or WRED	%, WFQ or strict	Out of contract gold traffic
Gold	0-7	2	ATM VBRrt, Frame Relay, realtime video	Xms/hop tail drop or WRED	%, WFQ or strict	Low jitter, low max delay, guaranteed delivery
Silver	0-7	3	EPL, Frame 0 CIR, ATM UBR	Xms/hop tail drop or WRED	%, WFQ or strict	High max delay, high jitter, guaranteed delivery
Best Effort	0-7	4	Internet traffic, high jitter	Xms/hop tail drop or WRED	%, WFQ or strict	High max delay, high jitter, lowest priority

Non-conforming traffic, or traffic that exceeds the customer's service level agreement, is not discarded but is given a lower priority than all conforming traffic.

Conclusion

Sophisticated QoS is a critical requirement of any multi-service edge router. The ST-series is the first service edge router to offer the robust QoS and traffic-management capabilities required to deliver all types of switched and routed traffic. For service providers, the result is an ability to offer a complete range of QoS-enabled data services from a single, flexible platform.

For more information on ST-series products and ShadeTree Management Suite, go to http://www.ecitele.com/dnd or contact one of ECI's local offices listed here:

Corporate Headquarters/Research & Development Center
ECI Telecom Ltd.
30 Hasivim Street
Petach Tikva, 49133 Israel
Tel: +972 3926 6555
Fax: +972 3928 7100

US Research & Development Center
ECI Data Networking Division
Omega Corporate Center
1300 Omega Drive
Pittsburgh, PA 15205, USA
Tel: +1 412 809 4200
Fax: +1 412 809 4201

Europe
ECI Telecom Gmbh (Germany)
Buropark Oberursel, In der Au 27,
61 440 Oberursel, Germany
Tel: +49 6171 6209 0
Fax: +49 6171 6209 88

ECI United Kingdom
ISIS House, Reading Road, Chineham
Basingstoke, Hampshire, RG24 8TW, UK
Tel: +44 1256 388 000
Fax: +44 1256 388 144

ECI Telecom France
Espace Velizy "Le Nungesser"
13 Avenue Morane Saulnier, 78140, Velizy,
France
Tel: +33 (1) 3463 0480
Fax: +33 (1) 3946 2118

North America
ECI Telecom Inc., USA
1201 West Cypress Creek Rd
Fort Lauderdale, FL 33309, USA
Tel: +1 954 772 3070
Fax: +1 954 351 4404

Latin America
ECI Telecom do Brasil Ltda.
Av. Dr. Cardoso de Melo, 1460 - cj. 101/2
Vila Olimpia, 04548-005 - Sao Paulo - SP - Brasil
Tel: +55 11 3512 1600
Fax: +55 11 3512 1601

Asia Pacific
ECI Telecom Singapore
150 Beach Road #28-07/08
Gateway East, Singapore 189720
Tel: +65 6297 7335
Fax: +65 6299 2716

ECI Telecom India
301, Boston House
Suren Road, Andheri - East
Mumbai - 400 093
Tel: +91 22 5675 8971
Fax: +91 22 5675 8973

About ECI Telecom
ECI Telecom offers future-ready telecommunications solutions that leading carriers and service providers rely on for delivering revenue-generating services to their business and residential customers. With its current products, ECI can deliver a full complement of access-to-edge IP transport solutions today. Known for its ability to translate customer needs into scalable, flexible, cost-effective solutions, ECI helps companies increase the value of their network infrastructure and reduce operating expenses. The company's single-shelf networking systems simplify network deployment and enable Build-as-You-Grow™ next generation telecommunication networks.

The Data Networking Division
The Data Networking Division (DND) adds next-generation IP/MPLS edge routing technology to ECI Telecom's product and services portfolio. DND's edge routers offer full-featured, multi-service support and complete Internet routing in a carrier-class, IP-based platform. ECI's ST-series™ routers provide the automated subscriber management, reliability, and performance that service providers need to implement advanced, revenue-generating broadband applications, like video on demand or voice over IP.