Cisco Systems

White Paper

Quality of Service on Cisco Advanced ATM Multiservice Switches: Benefits of the Cisco Multiservice Architecture

Service providers want a single network that supports multiple services. Their greatest challenge in establishing this multiservice network is the ability to provide a wide variety of IP services to their customers without affecting the existing ATM and Frame Relay investments. Service providers who are looking beyond basic IP connectivity have a growing interest in technologies that can simultaneously provide high reliability, high-aggregate bandwidth, and flexibility of service offerings while ensuring appropriate quality of service (QoS) for all of the services provided.

Traditional service provider networks transport many applications, including delay-sensitive voice, bandwidth-intensive video, and best-effort data. Networks are challenged to provide predictable, secure, and reliable communications for the applications they transport. Various mechanisms have been used to provision and preserve the integrity of traffic flows. Providers are increasingly focusing on QoS while they add Layer 3 capabilities such as traffic engineering and business services such as virtual private networks (VPNs) to their networks. QoS has become key to the success of these networks, especially as separate networks converge to become a true multiservice transport network with applications such as voice, video, and data.

ATM has historically been recognized as the technology of choice for creating an infrastructure that not only provides very granular levels of QoS, on a per-connection basis, but also offers intelligent routing, signaling, and control mechanisms. In recent years, Multiprotocol Label Switching (MPLS) has brought the notion of connectivity to connectionless IP networks and provided such networks with intelligent traffic engineering and routing enhancements. MPLS has been embraced as a mechanism for forwarding packets through a network, and is used to perform label switching for transporting packets across packet- or cell-based networks. MPLS is designed to work with numerous protocols—a primary reason service providers have been successful in adopting MPLS on their existing ATM infrastructures. Advantages of MPLS include eliminating the hop-by-hop packet-processing overhead and facilitating explicit route computation. ATM and MPLS offer exactly what providers are looking for—differentiated levels of QoS, bandwidth optimization, and load balancing over a common network infrastructure.

Service providers need more than a multiservice network, however. They require a service-quality network that can provide the highest level of QoS and compliance with their service-level agreements (SLAs) with low risk, minimal operational complexity, and a high return on investment.

This paper describes the advantages of Cisco multiservice switches for the evolving QoS needs of service providers, and examines the necessary mechanisms for preserving end-to-end QoS in emerging carrier multiservice networks. It presents the unique architectural benefits of Cisco multiservice switches—benefits that make the provisioning of IP and ATM QoS on a single switch possible without compromising either kind of service.

Advanced QoS on the Cisco Multiservice Platform

As networks grow in size and complexity, the integrity of the underlying transport infrastructure becomes critical in delivering the services provisioned on these networks. Robust QoS solutions include numerous technologies that interoperate to deliver scalable end-to-end performance and compliance with SLAs. Network QoS can be characterized and guaranteed by:

- Reliability, availability, and serviceability
- Resource partitioning and management, especially under conditions of congestion
- Cell-switching architecture
- Availability of QoS features from access to edge to core

The Cisco multiservice QoS solution offers the most robust and mature technology for meeting each of the above characteristics.

Reliability, Availability, and Serviceability

Reliability and availability are fundamental to delivering QoS. Although reliability and availability are traditionally not considered QoS metrics, their combined role in providing QoS is critical. Switches with low uptime and availability cannot be expected to provide high QoS and compliance with SLAs. Packet voice and video require high uptime and have strict latency bounds. Retransmission and other failure-recovery mechanisms cannot guarantee that packets will cross the network in the time needed to deliver high QoS.

Reliability, availability, and serviceability are significantly more robust in switches with parallel and redundant operating devices, while providing similar functionality. However, reliability, availability, and serviceability go beyond basic redundancy. While redundancy of the switch and service modules is critical, high reliability, availability, and serviceability relate not only to minimizing failures and downtime, but also to switch recovery time in the event of a failure, failure reporting, diagnostics, and carefully planned upgrades. Cisco multiservice switches achieve the highest levels of uptime in carrier networks. In evaluating multiservice switch architectures, service providers should consider reliability, availability, and serviceability as key QoS metrics for their networks. A switch that provides large buffers and queues but implements poor reliability, availability, and serviceability will deliver poor QoS.

Cisco multiservice switches achieve very high hardware and software reliability through many mechanisms, such as inter- and intraprocess message passing, transmit/receive priority queuing for messages, and data traffic protection. Features such as memory protection and active monitoring of system components and databases ensure that the switch does not cause degradation of QoS for applications needing uninterrupted control and data protection. In addition, switch design ensures that adequate data protection is maintained even if the controller suffers a fault. Increasing reliability and availability makes it possible to ultimately support advanced QoS on the switches and ensures a higher uptime for the service provider network.

Resource Partitioning and Management

Multiservice implies not only the existence of multiple traffic types on the network, but also the ability of a single network to support these applications without compromising QoS. Resource partitioning makes it possible to support more refined diversification among traffic classes as a result of unambiguous partitioning of resources. Initial attempts to interwork IP and ATM focused on attempting to map IP onto established ATM networks. Mechanisms such as Multiprotocol over ATM (MPOA) were only able to partially address the need to make IP and ATM work together. Existing mechanisms were not sufficiently scalable or flexible, suffered from being too complex to operate, and had severe performance limitations. The ATM control plane still needed to be a part of the process, so an IP flow would have to be mapped to a switched virtual circuit (SVC) or to a pre-established Layer 2 multicast group. To meet performance requirements, this would have to be implemented in hardware, which often was not practical.

Cisco avoided all of these issues by implementing the Multiservice Switching Forum (MSF) architectural approach. The MSF (http://www.msforum.org) advocates an open, standards-based architecture that allows flexibility in terms of the adaptation, switching, and control planes used to build a multiservice network. This is shown in the diagram of the Cisco Virtual Switch Architecture (CVSA) in Figure 1, derived from the MSF documentation.

Figure 1
Cisco Virtual Switch Architecture

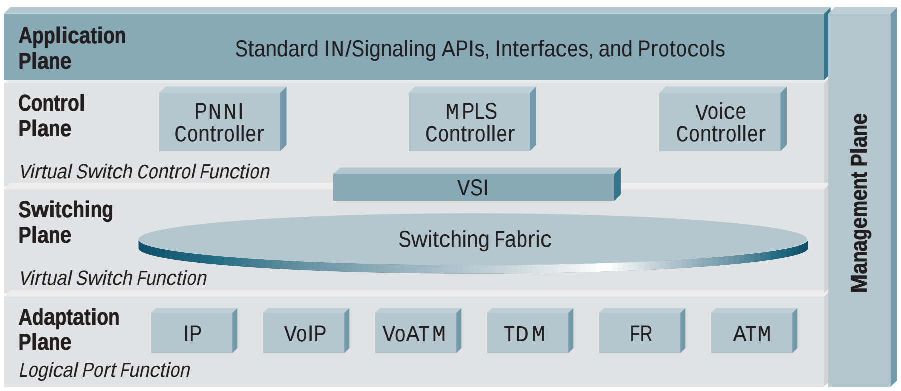

Resources are allocated to each controller, in this case Private Network-to-Network Interface (PNNI), MPLS, and Media Gateway Control Protocol (MGCP). Additional controllers can be added as support for applications grows. Each controller can act independently and concurrently with the other controllers. This gives Cisco a number of immediate advantages:

- IP services are managed by an MPLS controller, eliminating the requirement to map IP flows to ATM SVCs or permanent virtual circuits (PVCs). This removes the interaction between controllers as described above and preserves QoS signaling in its native IP format.
- Cisco IOS® Software implementation of QoS features (for IP and MPLS) is immediately applicable to the cell-based switching infrastructure.
- The multiple control plane capability on Cisco ATM multiservice switches allows service providers to build a complete MPLS network with a mix of cell- or packet-based switches as appropriate.
- Common management tools are applicable across product lines. A single network management system is needed to manage one physical switch operating as two separate logical switches.

A key feature of the CVSA is the separation of IP flows and ATM-based connections into separate queues. A switch that has implemented separate controllers for MPLS, ATM, and IP switching on a single switch, but does not provide separate queues, will have problems maintaining the QoS for each traffic type. As IP flows start getting mapped into the same queues that are being used for ATM connections, the resources that have been carefully allocated to meet per-VC service requirements will no longer be available to those connections, and SLAs are likely to be affected. The CVSA partitions resources in such a way that resources *between* partitions can be static (dedicated number of queues for ATM and IP), but resources *within* partitions can be dynamic (one IP queue can use unused bandwidth on another IP queue). This ensures that unused bandwidth is not wasted and that bandwidth allocation remains flexible.

Traditional ATM-based switches do not have dedicated IP queues, and therefore have issues when trying to map short-duration IP flows into ATM queues. Likewise, traditional IP routers do not have queues dedicated to ATM-based services. When supporting Layer 2 emulation services across an IP infrastructure, a close approximation of the required per-VC QoS for ATM needs to be supported. At higher speeds, and for services that require less specific QoS, such as a UBR (unspecified bit rate) service, this is less of an issue, and an IP's differentiated service profiles will be sufficient. Connection scale for the core may be met by using virtual trunking. However, to extend this technology out toward the edges of the network, and to support all of the various service types within ATM, traditional ATM switches are unable to provide the separate queues required to give QoS appropriate to each service.

Cell-Switching Architecture

A multiservice network is designed to handle packets from various applications and sources. It is expected that source traffic will vary not only in its requirements for the expected QoS, but also that source traffic flows will consist of packets of varying lengths and sizes. Packets to be transported may include Frame Relay traffic, ATM cells, and IP packets—all of which will need to be switched by the switch fabric and transported across the network. The key components of any switching architecture include the queuing model and the switching fabric itself. Queuing is used to alleviate the congestion that arises due to traffic bursts and unpredictable traffic flows in the switch. The switching fabric is the path the data takes to move from one port to another.

Cisco multiservice switches implement cell switching, or the ability to switch fixed-size data cells as opposed to packets or cells of varying size. This has several advantages. First, fixed-size cells (either standard ATM-size cells or other cells) have deterministic latency profiles across the fabric, and have a significant advantage in controlling jitter (interpacket or intercell delay variation), a key QoS metric when handling voice and video traffic multiplexed with data. Secondly, combined with intelligent scheduling, prioritization, and input/output (I/O) mechanisms, the risk of large packets (an FTP transfer packet) holding up service for a small packet (a voice packet) is minimized. Finally, the combination of virtual output queuing, load balancing, and intelligent arbitration mechanisms, used during resource contention, are key to minimizing delay across the switch fabric.

As mentioned earlier, an inherent advantage of the Cisco Virtual Switch Architecture is the separation of the control and forwarding planes. It is important to note that even though the control plane can be implemented to support services that are not necessarily cell-based (MPLS or voice, for example), the underlying forwarding and switching plane takes advantage of the benefits of cell switching. Since the fabric works with fixed-size cells, the separation of the control and forwarding plane allows non-cell-based services (Frame Relay or MPLS, for example) to take advantage of the inherent benefits of cell switching.

Cisco multiservice switches have dedicated hardware queues and a crossbar architecture. Crossbar switches are inherently non-blocking, and combined with virtual output queuing (VOQ) ensure that issues such as head-of-line blocking are eliminated. Schedulers can be individually tuned to ensure weighted scheduling for one class over another. This is crucial for maintaining service fidelity for high-priority traffic in the presence of oversubscribed low-priority traffic. The sophisticated queuing and buffering mechanisms implemented in hardware are especially relevant for emerging applications. These applications require very tight jitter control, can be "bursty," and are sensitive to delay. Cisco multiservice switches provide effective management and control of large buffers in hardware, a feature critical to establishing and sustaining the QoS of a flow.

Cell-based crossbar switches ensure guaranteed priority packet delivery while providing priority-based congestion control, dedicated low-latency queuing, and packet sequence integrity, all of which are essential for supporting real-time premium services such as voice over IP (VoIP).

Availability of QoS Features

MPLS (previously called Tag switching) was pioneered at Cisco. Cisco IOS Software, which underlies MPLS, provides advanced QoS mechanisms that can be used from the edge to the core of the network to ensure QoS of traffic flows. Various QoS mechanisms are implemented on Cisco multiservice switches to work in a "divide and conquer" style. Cisco multiservice switches can perform intensive packet lookups, identify flows, and classify packets according to user requirements. They can perform bandwidth management and use Cisco IOS Software QoS in the following ways:

- *IP Precedence*—Uses three bits in the IP header to indicate the service class of a packet. Classes are set at the edge and enforced at the core.
- *Committed Access Rate (CAR)*—Performs two functions: packet classification and bandwidth management. CAR analyzes the packet and assigns a service class based on the packet header information. This is done at Layer 3, facilitating use of a variety of attributes such as source, destination, protocol, or application to classify packets. CAR can also be used for bandwidth management by setting multiple Layer 3 thresholds.
- *Weighted Random Early Detect (WRED)*—Prevents network congestion before it occurs. WRED is a proactive mechanism that works by assigning weights to service classes (voice can be prioritized over data, for example).
- *Class-Based Weighted Fair Queuing (CBWFQ)*—Assigns different weights to different service classes allowing a switch to manage buffering and bandwidth for each service class. It can be used to constrain delay bounds.
- *Setting Experimental (EXP) Bits*—Satisfies the requirement of service providers who do not want the value of the IP Precedence field modified within IP packets transported through their networks. By choosing different values for the MPLS EXP field, packets can be marked based on characteristics such as rate or type, so certain packets have priority over others during congestion.
- *Application recognition and higher layer QoS*—Uses network-based application recognition (NBAR) and policy-based routing (PBR).

Traffic engineering plays a critical role in providing QoS. ATM implementation of intelligent routing and traffic engineering (route optimization, load balancing, and so on) is carried out using the PNNI Routing and Signaling Protocol. MPLS traffic engineering is based upon Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) protocols, and supports measurement-based engineering of connectionless traffic (such as IP traffic). MPLS traffic engineering not only improves bandwidth utilization, but also increases the robustness and reliability of service provider networks by providing alternate routes and load balancing. Cisco is unique in offering all of these advanced IP features on its ATM multiservice switches.

Example of Voice and Data in a Multiservice Network

This section of the paper highlights how two separate classes of traffic may be handled in a multiservice network. Voice needs to be transported across the network with very strict delay and jitter bounds. Low-priority data also has low tolerance for packet loss but can withstand larger bounds on delay and jitter.

It is imperative that voice traffic, when competing for network resources against lower priority traffic, receives output link scheduling priority at all times to maintain strict loss and jitter parameters.

Table 1 lists the tolerances for the various parameters that may be considered.

Table 1 Voice and Data needs for QoS

QoS Parameter	Packet Voice	Packet Data
Bandwidth requirement	Low to moderate	Moderate to high
Sensitivity to random drops	Low	High
Sensitivity to jitter (delay variation)	High	Low
Sensitivity to end-to-end delay	High	Low

Features such as Compressed Real-Time Transport Protocol (CRTP, RFC 2508), Multilink Point-to-Point Protocol (MLPPP), low-latency queuing (LLQ), and link fragmentation and interleaving (LFI) are critical to making sure that voice QoS is guaranteed in access networks over low-speed links.

Four specific QoS mechanisms may be used to ensure voice QoS in this scenario:

- *Classification/queuing/scheduling*—Features such as WFQ, CBWFQ, LLQ, and WRED may be used to ensure that the packet gets priority, in addition to marking the priority via IP Precedence, MPLS EXP bits.
- *Link bandwidth efficiency*—To limit delay on slow access, links may be used via fragmentation and interleaving (LFI and via compression such as CRTP and voice-activity detection (VAD). Voice uses Real-Time Protocol (RTP). IP, User Datagram Protocol (UDP) and RTP packet headers can be compressed from approximately 40 bytes down to 5 to 8 bytes.
- *Traffic shaping*—Mechanisms such as generic traffic shaping (GTS) and Frame Relay traffic shaping (FRTS) may be used to ensure that speed mismatches are handled appropriately and that jitter is reduced. Depending on the application, different shaping mechanisms are used. For example, in a pure VoIP application over serial links, no shaping is needed and bandwidth is guaranteed at line speed. For VoIP over Frame Relay (VoIPoFR) or voice over Frame Relay (VoFR), FRTS may be used. For voice over ATM (VoATM), ATM traffic shaping mechanisms may be used.
- *Bandwidth reservation and management*—Aimed at reserving bandwidth for certain flows. Guaranteed bandwidth mechanisms and Resource Reservation Protocol (RSVP) can be used to signal and reserve bandwidth in advance.

As the voice packets transit the switch, they need to be scheduled ahead of other low-priority traffic, and be given assured QoS. A cell-based crossbar-switching architecture with dedicated queues and tight jitter control mechanisms ensures that the voice packet is given assured QoS and minimizes crosspoint latency. The egress scheduler schedules the voice packet to an outbound interface for transit across the network. Data traffic queuing techniques such as CBWFQ and modified deficit round robin (MDRR) may be used. Further packets that violate service contracts may be tagged and dropped during times of severe congestion.

Feature Availability: IP and ATM QoS

The Cisco Differentiated Services (DiffServ) traffic engineering solution extends MPLS traffic engineering, enabling service providers to implement constraint-based routing (CBR) of guaranteed bandwidth traffic, satisfying a restrictive bandwidth constraint (similar to CBR traffic). This allows service providers to perform admission control and separate route computation for discrete subsets of traffic, such as voice, video, or data.

MPLS auto bandwidth for tunnels allows you to automatically adjust the bandwidth allocation for traffic based on measured load.

Table 2 QoS Mechanisms for IP and ATM

QoS Scheme	IP QoS Features	ATM QoS Features
Connection Admission Control (CAC)	RSVP or CBR using Label Distribution Protocol	PNNI CAC
Traffic packet classification	Default copy of IP Precedence to EXP bits; matching on MPLS EXP bits	Mark cell loss priority (CLP) bit
Congestion management	LLQ/CBWFQ support for MPLS packets; MDRR support for MPLS	Available bit rate (ABR) flow control, ABR virtual source/virtual destination (VS/VD), closed loop congestion management
Congestion avoidance	WRED based on EXP bits	Early packet discard, partial packet discard
Traffic/packet marking	Set MPLS EXP bits using CAR	
Traffic conditioning	Policing with EXP bit matching and setting using CAR; GTS	Standards-based traffic shaping and policing
Compression and fragmentation for low-speed links	CRTP and LFI	N/A

Summary and Conclusion

Customers of service provider networks are end users, and QoS is ultimately driven by the applications that these users employ. Network applications, such as mission-critical data, and real-time applications will continue to require very strict QoS standards, while applications such as peer-to-peer networking and best-effort services will both need support on the same network. Service providers are therefore interested in multiservice networks that can provide a wide spectrum of QoS, from finely granular on a per-VC basis to aggregate flow-based QoS. The success of multiservice networks is based largely on the ability to flexibly add such services on the network and to provision QoS for these applications.

Key components of such networks are reliability, availability, and serviceability; resource partitioning and management; a robust and scalable cell-switching architecture; and rich QoS feature availability. Underlying reliability is critical. In fact, reliability is itself a metric that must be considered in evaluating the overall QoS. Key features of resource partitioning include the ability to have multiple non-interfering resource partitions (MPLS or PNNI, for example) on a single switch, and the ability to dynamically adjust resource allocations to the applications that use these partitions. The importance of a switch architecture designed for handling tight QoS—including jitter and delay control, intelligent scheduling, and buffer management, as well as patented hardware traffic management—is critical to making sure the switch provides the highest level of QoS support. Finally, the availability of a rich set of features and a QoS toolkit that can be customized for specific applications is key to ensuring that different applications have access to the most pertinent and applicable QoS techniques. The availability of a large number of QoS features and switching capacity, while integral to providing QoS, cannot provide high QoS in isolation. These must be orchestrated with intelligent traffic engineering, routing, management, and high reliability to ensure that service providers can continue to guarantee their SLAs.

Cisco multiservice switches deliver a strong suite of IP and ATM QoS capabilities, ensuring true multiservice QoS on a single infrastructure without the need to manage two separate networks for both IP and ATM.

Cisco Advanced ATM Multiservice Portfolio

Cisco BPX® 8600 Series Multiservice Switches

The Cisco BPX 8600 Series Multiservice Switches are the most widely deployed multiservice carrier switch. Designed specifically for the service provider environment, the BPX 8600 delivers "five-nines" or 99.999-percent service availability in the world's largest Frame Relay and ATM networks. The Cisco BPX 8600 Series also leads the industry in traffic management capabilities and multishelf scalability.

Cisco MGX® 8000 Series Multiservice Switches

The Cisco MGX 8000 Series Advanced ATM multiservice switches offer the widest range of high-availability Frame Relay, ATM, IP, circuit emulation, and voice services. They are also the most cost-effective multiservice solution for a wide range of sites with the ability to scale from DS0 to OC-192C/STM-64.

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe
11 Rue Camille Desmoulins
92782 Issy-les-Moulineaux
Cedex 9
France
www-europe.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 317 7777
Fax: +65 317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe