

# Cisco ASR 1000 Series Aggregation Services Routers: QoS Architecture and Solutions

## Introduction

Much more bandwidth is available now than during the times of 300-bps modems, but the same business principles of saving as much bandwidth as possible still applies. It is particularly important for enterprises to conserve bandwidth, but it also important for service providers. Both segments must have a network architecture that can offer higher levels of service for critical customers or Layer 7 applications. For this reason, routers capable of offering quality of service (QoS) that is flexible, scalable, and high-performing is more important than ever.

QoS is defined as the techniques that give applications on a network a predetermined performance level. When a network has QoS implemented, the performance of an application is determined by a policy rather than by the network load or its components. Most agree that the performance of an application is determined by the following factors: packet loss, delay, jitter (delay variation), and throughput. Each network administrator determines acceptable values, depending upon the specific application and user requirements. More details on these definitions and general information about Cisco IOS® QoS are available at: <http://www.cisco.com/go/qos>. The following are enterprise and service QoS design guides, which offer detailed information about how to deploy QoS.

[http://www.cisco.com/application/pdf/en/us/quest/netsol/ns432/c649/ccmigration\\_09186a008049b062.pdf](http://www.cisco.com/application/pdf/en/us/quest/netsol/ns432/c649/ccmigration_09186a008049b062.pdf) and

[http://www.cisco.com/en/US/partner/netsol/ns590/networking\\_solutions\\_white\\_paper09186a00801b1c5a.shtml](http://www.cisco.com/en/US/partner/netsol/ns590/networking_solutions_white_paper09186a00801b1c5a.shtml).

QoS on Cisco ASR 1000 Series Aggregation Services Routers offers powerful advanced and flexible hardware QoS capability. This paper offers an overview of the Cisco ASR 1000 Series Router hardware, which includes the Cisco QuantumFlow Processor (Cisco QFP) and its two chips: Cisco QFP Engine and Cisco QFP Traffic Manager. It also addresses Cisco ASR 1000 Series software QoS features and capabilities and how they apply in several different architectures and deployments.

## Cisco ASR 1000 Series QoS Overview

Cisco IOS® Software has always been an industry leader with comprehensive and flexible QoS. In addition, Cisco hardware-based platforms have performed some of these QoS features at a very high throughput. Cisco ASR 1000 Series Routers combine these two paradigms with very complex QoS configurations at a very high throughput.

Cisco ASR 1000 Series Router architecture addresses the basic QoS performance requirements, as illustrated in Table 1.

**Table 1.** QoS Performance of Cisco ASR 1000 Series Router

QoS Parameters	Cisco ASR 1000 Series Router Architecture Features
Packet loss	Low drop probability for high-priority applications, even during failovers
Delay	Carrier-class low-latency design of data path

<b>Jitter (delay variation)</b>	Ingress QoS and ability to monitor egress and ingress buffer levels
<b>Throughput</b>	Capability of scheduling the entire system bandwidth with no performance penalties

Cisco ASR 1000 Series QoS advanced features include:

- Fully dedicated and different silicon for both scheduling (Cisco QFP Traffic Manager) and packet processing (Cisco QFP Engine)
- Multiple and flexible levels of Hierarchical Scheduling and Queuing
- Three-parameter scheduling with high-priority traffic protection under very high levels of oversubscription
- Deep Packet Inspection features such as Network Based Application Recognition (NBAR) at a very high throughput, implicit in the Cisco QFP architecture
- QoS integration with the Cisco Session Border Controller (SBC), IP Security (IPsec), IP Multicast, NetFlow, Compressed Real Time Protocol (cRTP), etc. at a very high throughput and scalability
- Flexible software architecture that allows very high feature velocity and versatility

### Cisco ASR 1000 Series QoS Hardware Architecture

Cisco ASR 1000 Series Routers offer a centralized architecture when it comes to the forwarding plane. The 5- and 10-Gbps ASR 1000 Series Embedded Services Processors (ESP5 and ESP10, respectively) represent the forwarding plane in this case. All network traffic flows through ESP10, making it the main engine for all QoS features. The main component of both ESPs is the Cisco QFP and its two pieces of silicon, the Cisco QFP Engine and Cisco QFP Traffic Manager. Cisco QFP Engine is in charge Modular QoS CLI (MQC) classification, marking, Weighted Random Early Detection (WRED), policing, etc. Cisco QFP Traffic Manager handles shaping, priority, and Class-Based Weighted Fair Queuing (CBWFQ) scheduling, etc. The Cisco ASR 1000 Series SPA Interface Processor 10 (SIP10) is also capable of ingress classification, buffering, and scheduling with a non-MQC configuration. For more information on these components please visit <http://www.cisco.com/go/asr1000>.

Highlights of QoS features in the ESP10 and ESP5 follow:

- 128,000 hardware queues
- Multiple levels of hierarchies
- Two levels of high-priority traffic per policy

QoS performance with the ESP10 depends upon the QoS configuration and the combination with other features. Nevertheless, the performance with a very complex QoS configuration, several thousand subinterfaces, IPv6 Multicast, and access control lists (ACLs) is several million packets per second. The huge advantage with the Cisco ASR 1000 Series is the fact that Cisco QFP Traffic Manager is not affected by any of the features that Cisco QFP Engine needs to perform. So QoS shaping and Low Latency Queuing (LLQ) configurations has minimum impact performance in the Cisco ASR 1000 Series Routers.

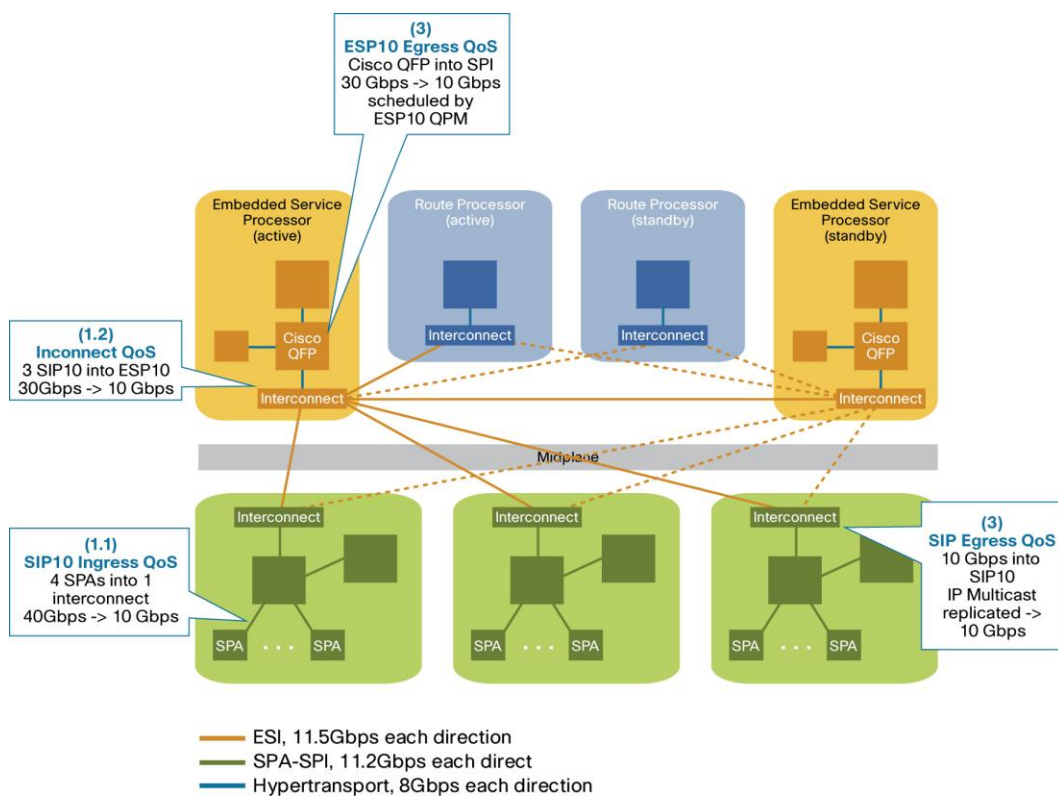
The best approach to understand Cisco ASR 1000 Series Router QoS is by following the flow of the packets from the ingress to the egress (Figure 1). Note that the Cisco ASR 1000 Series Router platform is designed to support high levels of oversubscription, and in the figure each of the steps through the box is a possible congestion point. The Cisco ASR 1000 Series Router QoS always tries to transmit high-priority traffic before any other type of traffic. This feature is especially

important when customers oversubscribe the system bandwidth with high- and low-priority traffic. Following is the flow of the packet in a Cisco ASR 1000 Series Router from a QoS congestion point of view:

1. SIP10 and ESP10 ingress classification and scheduling
  - 1.1. SIP10 ingress classification and scheduling
  - 1.2. ESP10 interconnect scheduling
2. ESP10 egress QoS
3. SIP10 buffering

Figure 1 shows a graphical representation of these steps.

**Figure 1.** Cisco ASR 1000 Series Router QoS Stages



It is important to note that this oversubscription is possible and desirable in most customer networks because the whole system is designed to preserve and propagate the high priority of, for example, voice and video packets.

### SIP10 Ingress Classification and Scheduling

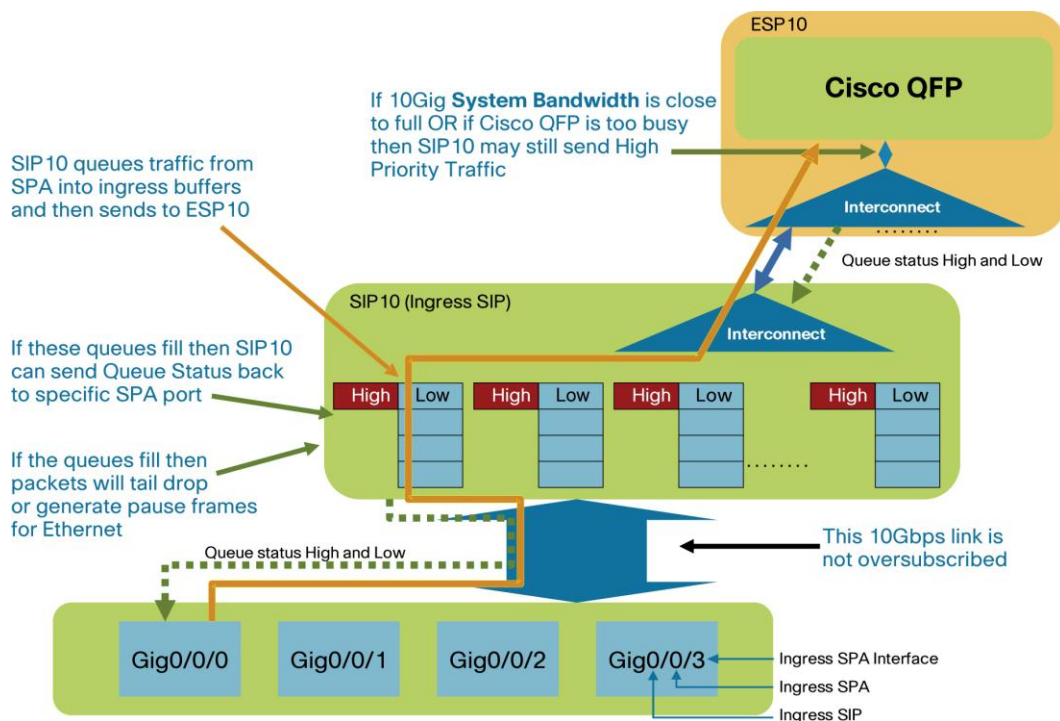
Ingress traffic can be classified and scheduled at the SIP10, so the high-priority traffic reaches the ESP10 before any other traffic. The classification could be done based on the priority of the ingress packets. No ingress remarking is done by default. The ingress packets classification is based on 802.1p (Ethernet priority), Multiprotocol Label Switching experimental bits for priority (MPLS EXP), and IP Precedence or IP differentiated services code point (DSCP) for both IPv4 and IPv6. Priority is then determined based on this classification. Cisco ASR 1000 Series Routers have a specific non-MQC command-line interface (CLI) to customize this classification; for example: `plim qos input map ip dscp <dscp-value | dscp-range> queue <low-latency | 0>`. In the case of Ethernet, you can configure this classification per VLAN.

Based on the previous classification and priority, packets are then moved into 2 queues (high and low priority). There are 2 queues per port and two levels of scheduling (minimum bandwidth and excess bandwidth), per port and per SIP10. The minimum bandwidth or committed information rate (CIR) is the rate at which traffic from each SIP10 is sent to the ESP10. The remaining bandwidth, or excess weight, is basically the bandwidth that is unused by any type of traffic. As with the SIP10 classification, there is a platform-specific CLI to modify the SIP10 scheduling; for example, `plim qos input [bandwidth <Kbps> [low-latency]] [weight <weight>]`. If the low-latency keyword is used, only high-priority traffic gets a minimum bandwidth guarantee per SIP or per port, making this basically a strict priority queuing scheme.

### ESP10 Interconnect Scheduler

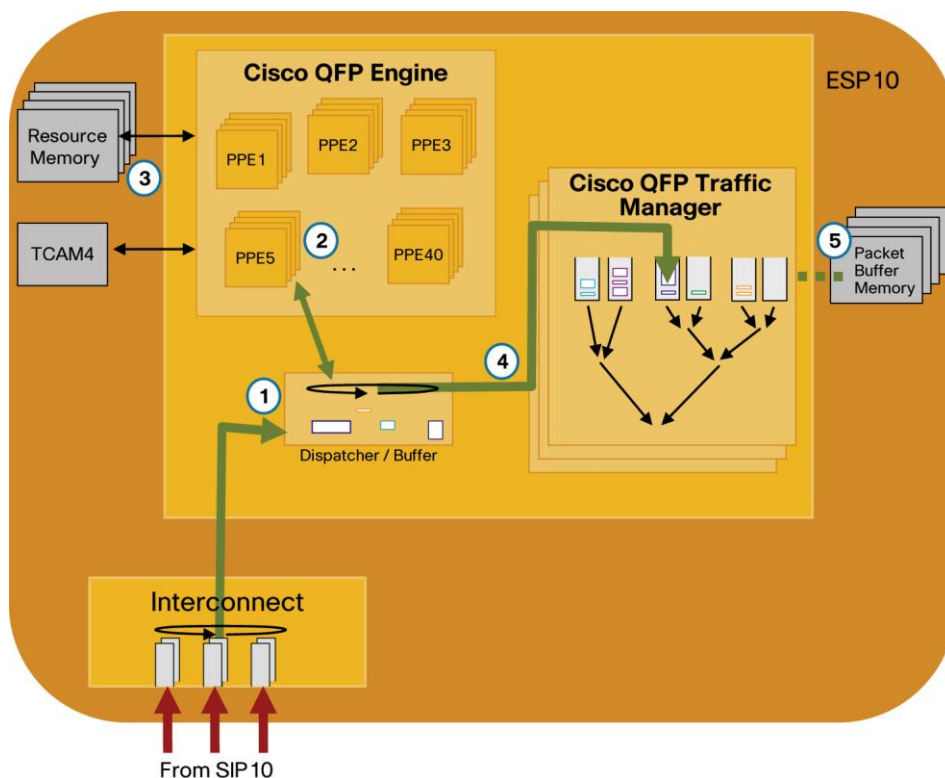
After the SIP10 classification and scheduling, ingress traffic is then scheduled by the ESP10 interconnect, which basically selects among the different SIP10s in the system, the high-priority traffic to be processed first, followed by the low-priority traffic. Obviously if there is no oversubscription all SIPs can send traffic toward the ESP10. There is also a platform-specific CLI to modify this scheduler; for example `hw-module platform slot <0-3> qos input [(bandwidth <> kbps [low-latency][weight <1-100>][strict-priority])]`.

Figure 2 shows a basic overview of how the system is constantly aware of the queue status, when the 10-Gbps system is overcommitted (based on Cisco QFP Engine usage). Even under these conditions, high priority can still be processed, since the interconnect uses backpressure from Cisco QFP and this backpressure is independent per priority queue and per ESI. Low priority can be backpressured while high priority packets continue to flow. The status of these ingress queues in the SIP allows the system to generate pause Ethernet frames traffic in the case of Ethernet, or tail-drop packets if SIP10 queues are almost full.

**Figure 2.** Cisco ASR 1000 Series Router Ingress Queues**ESP10 Egress QoS and SIP10 Backpressure**


After the ingress classification and scheduling at the SIP10 and the ESP10 interconnect scheduling among all the SIP10s in the system, ingress packets are then sent to Cisco QFP, first to the Cisco QFP Engine as shown in Figure 3 and described as follows:

1. Ingress packets are temporarily stored in Cisco QFP in a small internal packet buffer until they are processed
2. A free Cisco QFP Engine thread is allocated for each packets and processing of the packets begins, with features such as MAC classification, QoS classification, ACLs, forwarding lookup, WRED, cRTP, etc., including modifying packet contents. This process also includes NBAR at a very high throughput since the Cisco QFP can process the whole packets. It also includes policing; the policing algorithm is a single- or double-rate, three-color policer. This algorithm is defined in RFC 2697 and RFC 2698 and is fully supported in Cisco QFP
3. Tables are accessed in resource DRAM and ternary content addressable memory (TCAM) to perform lookups for features enabled for these packets, update statistics, update state for stateful features, etc.
4. When packet processing is complete and the packet has been modified, a request is issued to enqueue the packet to an output queue
5. The packet content is copied from the internal packet buffer to the deep output packet buffer where it is stored until scheduled for output

**Figure 3.** Cisco ASR 1000 Series Router QoS Forwarding Path: ESP10

Then packets are sent to the Cisco QFP Traffic Manager, as shown in Figure 4, where the following takes place:

1. From the previous step, packets were stored in the main output buffer
2. The Cisco QFP Traffic Manager Scheduler performs packet-scheduling decisions and:
  - a. It dynamically creates the different levels of hierarchy to match the MQC configuration
  - b. Because it is a three-parameter scheduler, it can enforce minimum (bandwidth in Cisco IOS MQC), excess (bandwidth remaining), and maximum (shape) bandwidth rates for each queue. A total of 128,000 hardware queues are possible
  - c. High-priority packets can pass lower-priority packets in the hierarchy
3. The bandwidth scheduler then allocates Cisco QFP's output bandwidth among the SIP10s
  - a. It selects the highest priority packets among the SIP10s not backpressured
  - b. It shares outbound bandwidth accordingly. High priority packets always take precedence regardless of SIP destination
4. Packet data is then transferred to the SIP10 simultaneously
5. Shallow buffers on the SIP10 and share port adapter (SPA) are used to allow simultaneous packet transfer out multiple ports
6. Backpressure from shallow buffers is used to control Cisco QFP Scheduler at the corresponding hierarchy node

**Figure 4.** Cisco ASR 1000 Series Router QoS: SIP10 Egress

It is important to note that shaping in Cisco QFP is implemented through by using of a token bucket algorithm with its additional burst parameters. As with the policer, the bandwidth for shaping is calculated including some of the Layer 2 fields of the packet, depending on the encapsulation (such as Packet over SONET Cyclic Redundancy Check (POS CRC), or Ethernet preamble).

It also important to highlight some of the Cisco QFP Traffic Manager Scheduler characteristics:

- There are multiple levels for egress hierarchical queuing, and there's backpressure and queue status at all these levels: SIP10 and some SPAs. These levels are additional to the ones in the Cisco IOS MQC configuration. So we have: MQC 3 level + SPA + SIP10, giving the Cisco ASR 1000 Series Routers an innovative QoS multilayer hierarchy in the industry. This hierarchy is used not only for networks interfaces, but also for internal communication with the Cisco ASR 1000 Route Processor 1 (RP1) and the built in ESP10 encryption engine
- High priority traffic is always processed first, two levels of high priority traffic are supported (priority level 1 and 2) this ensures low latency, as long HP traffic is less than the total bandwidth capacity of ESP10
- SPA+SIP10 +ESP10 queue status communication helps ensure a very high level of high priority traffic protection. The Cisco QFP is aware of and can receive queue status from all interfaces on all SPAs. It uses this status to control how the MQC queues are serviced. The Cisco QFP does not send more than the interface bandwidth. It can also listens to both high- and low-priority queue-status messages from egress queues on both the SIP10 and the SPA

## Cisco ASR 1000 Series QoS Software Architecture

Cisco IOS XE Software and its components constitute the Cisco ASR 1000 Series software architecture. One of the components of Cisco IOS XE Software is Cisco IOS Software, so most of the QoS features are derived from Cisco IOS Software, including traditional MQC QoS features such as Classification, NBAR, Marking, CBWFQ, Class-Based Shaping, Hierarchical Shaping and Policing, Percentage-Based Shaping and Policing, MQC Frame Relay Traffic Shaping, WRED, etc. The Cisco ASR 1000 Series Router QoS implementation adheres strictly to the Cisco IOS MQC framework to configure QoS policies. MQC provides a uniform template-based command syntax across Cisco IOS Software platforms. For more information about MQC, please visit: [http://www.cisco.com/univercd/cc/td/doc/product/software/ios122sr/cr/srqos\\_r/index.htm](http://www.cisco.com/univercd/cc/td/doc/product/software/ios122sr/cr/srqos_r/index.htm).

Two major features that are important to highlight follow:

- Multilevel priority queues

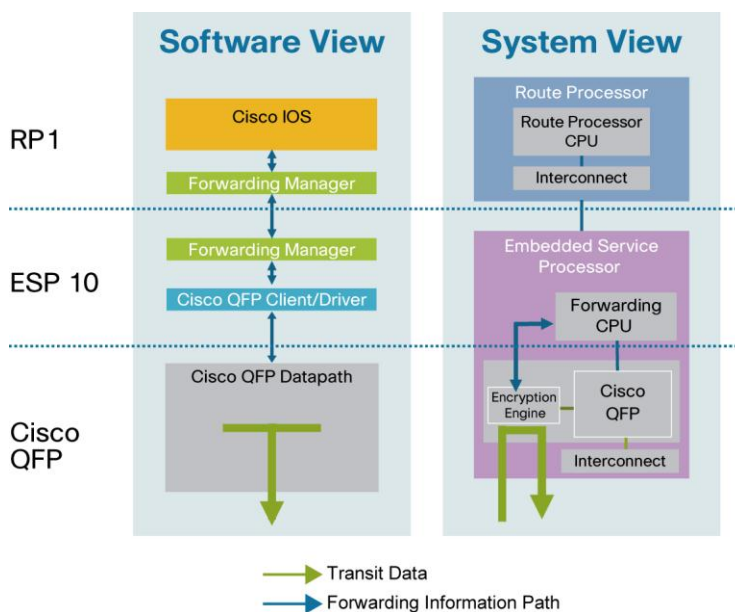
<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122sb/newft/122sb31/mpq.htm>

- Bandwidth remaining ratio

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122sb/newft/122sb31/bwratio.htm>

In basic terms, all these features configured in the familiar Cisco IOS MQC CLI are translated into Cisco QFP instructions to map exactly the functions and requirements of the most flexible QoS schemes. This process is illustrated in the Figure 5.

**Figure 5.** Software and System Stack on Cisco ASR 1000 Series Routers



As mentioned previously a small set of platform-specific commands are used to control the ingress classification and scheduling in the SIP10. For more information please visit

<http://www.cisco.com/go/asr1000>.



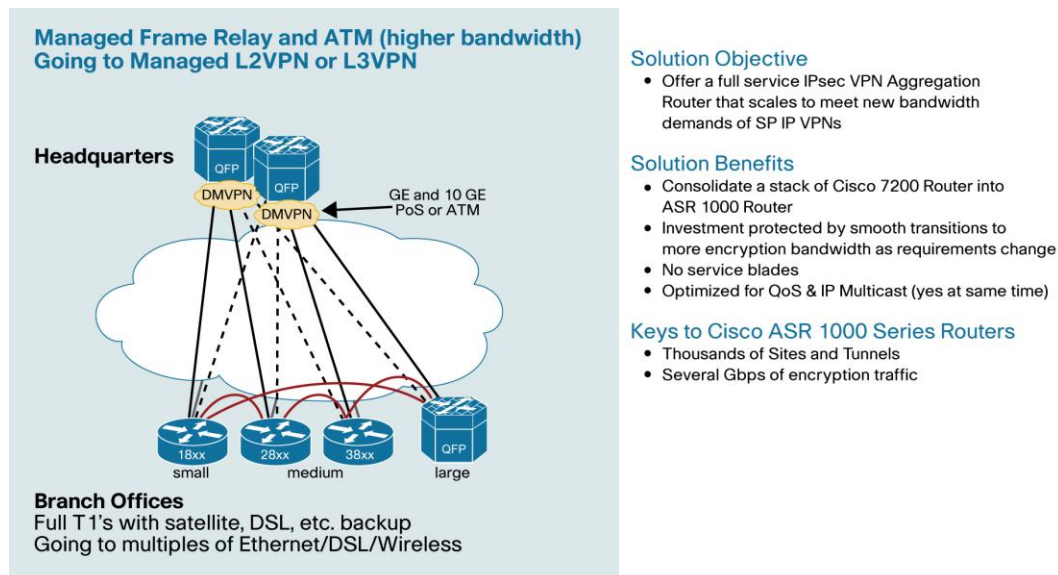
## Cisco ASR 1000 Series Powerful QoS and Its Benefits

Quality of service is just one of the many IP services that Cisco ASR 1000 Series Routers and Cisco QFP can deliver today. The great Cisco ASR 1000 Series QoS performance and versatility can and should be integrated with your other IP Services. Following is a sample of some of the benefits that QoS and other services offer.

### Enterprise and Service Provider Edge QoS: NBAR, cRTP, QoS and IPsec

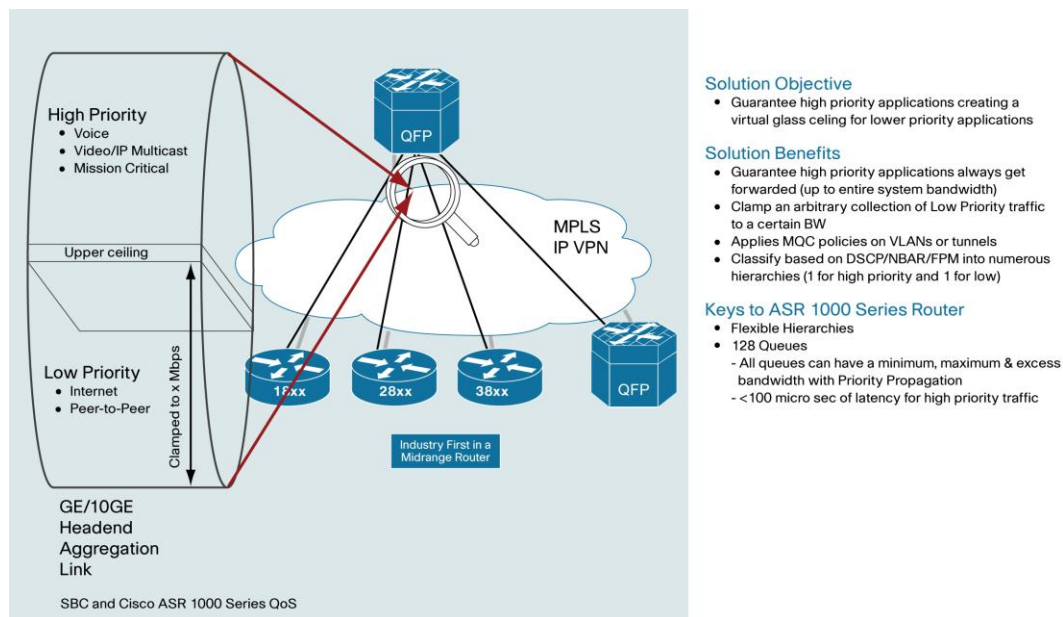
The next generation of enterprise networking architectures will definitely benefit from a midrange platform that can integrate several IP services at several Gbps of performance and throughput. Figure 6 illustrates how a headend platform at headquarters could have full IPsec VPN, QoS and IP multicast for several thousand branches in the same Cisco ASR 1000 Series Router. As mentioned previously, the multilayer hierarchy and queuing status applies not only to the network interfaces and Cisco QFP Traffic Manager, but also to the communication with the encryption engine, allowing QoS to integrate and buffer traffic in perfect coordination with all the pre- and postencryption services that are available in Cisco QFP. The additional integration with IP Multicast and its bursts is not only possible but also doable at a very high throughput. More information is available at <http://www.cisco.com/go/asr1000>.

**Figure 6.** Cisco ASR 1000 Series Routers: Multiservice, Scalable, and Secure at the Edge



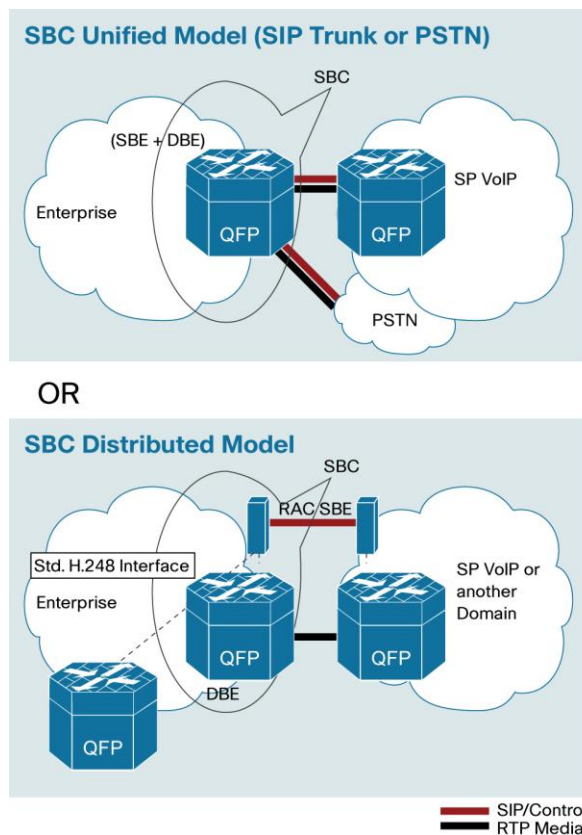
As mentioned previously, Cisco ASR 1000 Series Router QoS follows Cisco IOS MQC strictly, but it has also enhanced it with a new QoS feature (for more information on this new command, please visit <http://www.cisco.com/go/asr1000>) to allow network administrators and service provider operators to offer real bandwidth reservation to the most important applications in today's networks. This new QoS feature can be delivered under very high levels of congestion but at the same time, at a very high throughput and it is available not only to the traditional service provider operators and their voice-over-IP (VoIP) offerings, but also in general to network administrators who can identify other more specific Layer 4-Layer 7 applications. Cisco ASR 1000 Series Routers are capable of several Gbps throughput with features such as Cisco NBAR because of its full packet processing, which can identify features such as Skype, Citrix and peer-to-peer applications such as Bittorrent. Figure 7 summarizes these benefits.

**Figure 7.** Cisco ASR 1000 Series Routers: Superior Application Availability at the Edge



### SBC and Cisco ASR 1000 Series QoS

The Cisco Session Border Controller is a critical component for interconnecting VoIP and multimedia networks of enterprises and service providers. For more information please visit: <http://www.cisco.com/go/asr1000>. It is important to note that the Cisco ASR 1000 Series Router combination of SBC and QoS gives network operators a highly scalable VoIP solution, with thousands of VoIP calls in combination with IP Unicast data, IP Multicast video, etc. Figure 8 summarizes this solution.

**Figure 8.** Cisco ASR 1000 Series Routers: Next-Generation Voice and Multimedia at the Edge**Solution Objective**

- Interface with Service Provider-provided voice and multimedia (E.g. Telepresence) services directly at the edge.

**Solution Benefits**

- No overlay network, no appliances, no service blades
- Easier operations/lower costs
- Flexible deployment models
- Aligns with ITU NGN/IMS/TISPAN
- Also supports protected signaling for both voice and video services

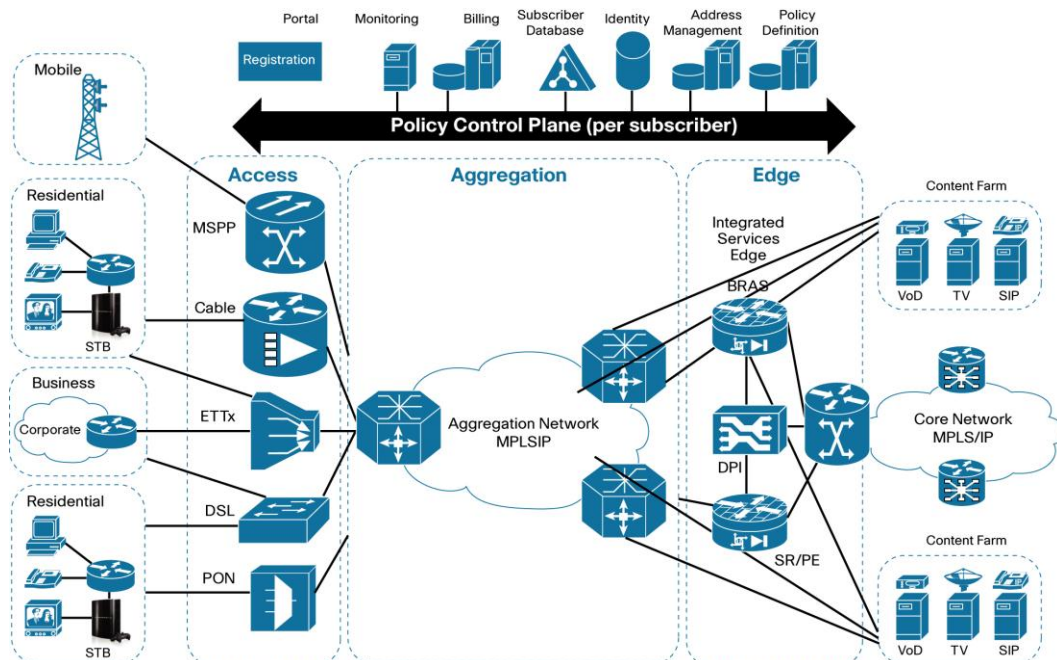
**Keys to Cisco ASR 1000 Series Router**

- IP-to-IP Gateway (SBC) with Security, QoS, IPv4/6 (IP unicast and multicast) - all enabled simultaneously
  - Thousands of VoIP calls + data up to 10 Gbps (w/accounting, firewall, call quality enabled)
  - Integrated with inbox HA infrastructure & DHCP Relay

**Service Provider (MPLS Provider Edge) and Broadband QoS**

Finally, Cisco ASR 1000 Series Routers offer a high throughput and highly scalable Broadband Remote Access Server (BRAS) solution, with complex hierarchical QoS schemes and IPv4 or IPv6 Multicast video streams. You could also combine this solution with an MPLS Provider Edge (MPLS PE) environment, where the different tunnels modes to offer Differentiated Services and QoS transparency. Figure 9 gives a high level overview of this possible solution.

**Figure 9.** Next-Generation Network Architecture: Service Integration and Architectural Distribution with Cisco ASR 1000 Series Routers



## Summary

- Cisco ASR 1000 Series Routers are designed to perform under highly oversubscribed conditions, with SIP10 ingress classification and scheduling and Cisco QFP flexible scheduler
- As long as high priority traffic does not oversubscribe ESP10 bandwidth, it will reach Cisco QFP for processing and will be transmitted before any other traffic
- The Cisco ASR 1000 Series Router software architecture is flexible for the rapid implementation of the latest Cisco IOS MQC QoS features
- The QoS architecture of the Cisco ASR 1000 Series and Cisco QFP satisfies all the voice, video and data requirements in today's distributed networks and their future generations



Americas Headquarters  
Cisco Systems, Inc.  
San Jose, CA

Asia Pacific Headquarters  
Cisco Systems (USA) Pte. Ltd.  
Singapore

Europe Headquarters  
Cisco Systems International BV  
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).

CCDE, CCVP, Cisco Eos, Cisco StadiumVision, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn is a service mark; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, IQ Expertise, the IQ logo, IQ Net Readiness Scorecard, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0801R)