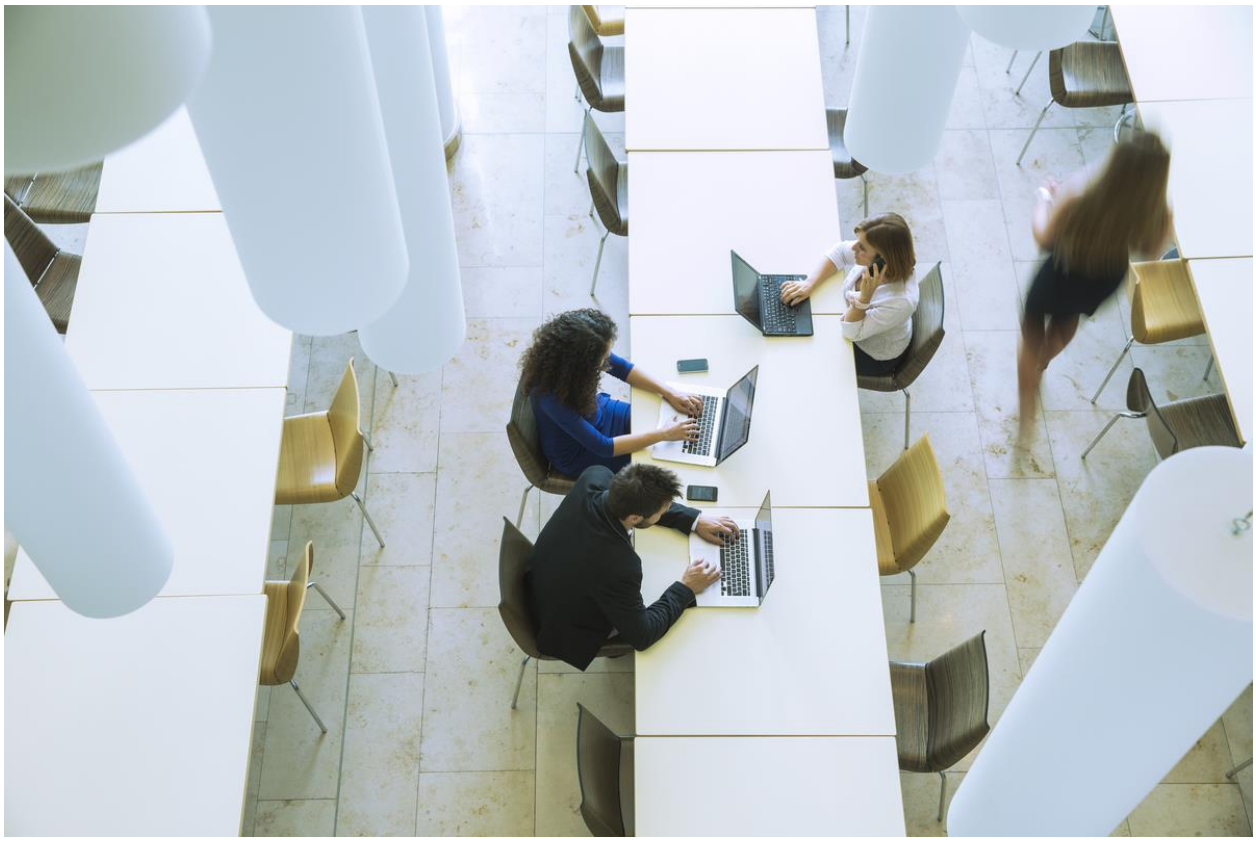

Cisco EasyQoS Solution Design Guide
APIC-EM Release 1.3

December, 2016

Table of Contents

Solution Overview

Customer Challenges

Today there is a virtual explosion of rich media applications on the IP network. This explosion of content and media types, both managed and unmanaged, requires network architects to take a new look at their Quality of Service (QoS) designs.

Virtually all businesses are looking to increase the productivity of their employees though the effective and innovative use of collaborative applications, regardless of the hardware-platforms these applications run on (PCs, laptops, tablets, smartphones, etc.), the physical or geographical location of the collaborating employees, or the types of media they wish to share on-demand. However, enabling and providing seamless Quality of Experience (QoE) around such services has traditionally placed challenging demands on network administrators, to the point where the foremost barrier to enabling QoS/QoE for collaborative applications is the not technical abilities of the infrastructure. Instead, it is the intrinsic complexity of enabling these features across disparate devices in a comprehensive, yet cohesive, manner.

In order for QoS to be effective, it needs to be deployed end-to-end, the same way a chain needs to be deployed end-to-end between a source and a target in order to have utility. Every link in the chain must have a cohesive, compatible QoS policy in order to achieve an end-to-end service level. However, it is the platform-by-platform variations in customizing, optimizing, and tuning that present the biggest barrier to QoS/QoE deployments.

Enter the network controller. The network controller helps by simplifying and abstracting platform-specific complexity from the network operator. Specifically, the network controller is programmed with all the link-specific information and Cisco best-practice knowledge so as to construct optimal end-to-end QoS "chains." An operator does not need to know the hardware or software queuing structures of the underlying infrastructure, nor do they need to know the QoS implications of interconnecting wired and wireless networks, nor do they need to know how the applications are to be recognized, despite the fact that an increasing number of these are encrypted. All the operator needs to know is which applications are important to his/her business. End-to-end provisioning is done in minutes (vs. months), leveraging industry standards and Cisco Validated Designs (CVDs).

Solution Description

Cisco Application Policy Infrastructure Controller Enterprise Module (APIC-EM) is Cisco's enterprise Software-Defined Networking (SDN) controller. APIC-EM provides the automation functionality within Cisco's new enterprise architecture, called the *Digital Network Architecture (DNA)*. A high level overview of the architecture is shown in the following figure.

Figure 1 **High-Level Overview of the Cisco Digital Network Architecture**

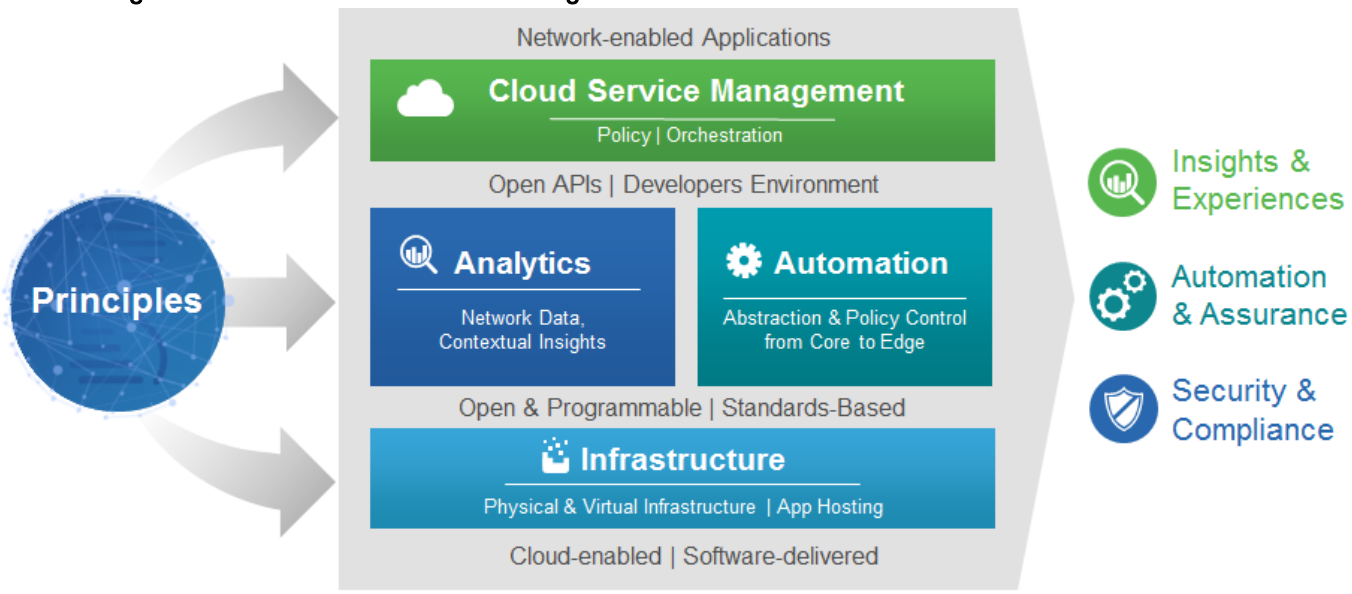

EasyQoS is an application that runs on top of APIC-EM. Hence, it is an integral part of the overall DNA architecture. The EasyQoS solution abstracts QoS policy by using a declarative model as opposed to an imperative model. A *declarative model* focuses on the intent or WHAT is to be accomplished, without describing HOW it is to be accomplished. For example, a network operator may express that an application such as Cisco Jabber is business-relevant—meaning that it is to be treated with the appropriate service—but he/she does not specify the details of how the QoS/QoE policies are to be configured in order to achieve this intent.

In contrast, an *imperative model* focuses on the execution of the intent (describing in detail HOW the objective is to be realized). For example, an imperative policy may include assigning Cisco Jabber to a hardware priority queue with a given bandwidth allocation percentage on a specific network switch interface.

Using a declarative model for policy-expression, rather than an imperative model, frees the network operator from having to spend extensive time-consuming cycles to deploy QoS policies. With a network controller, changes can be made in minutes, rather than months, resulting in agile networks that are tightly-aligned with evolving business requirements.

The following figure provides an overview of how the solution works.

Figure 2 High-Level Overview of the EasyQoS Solution

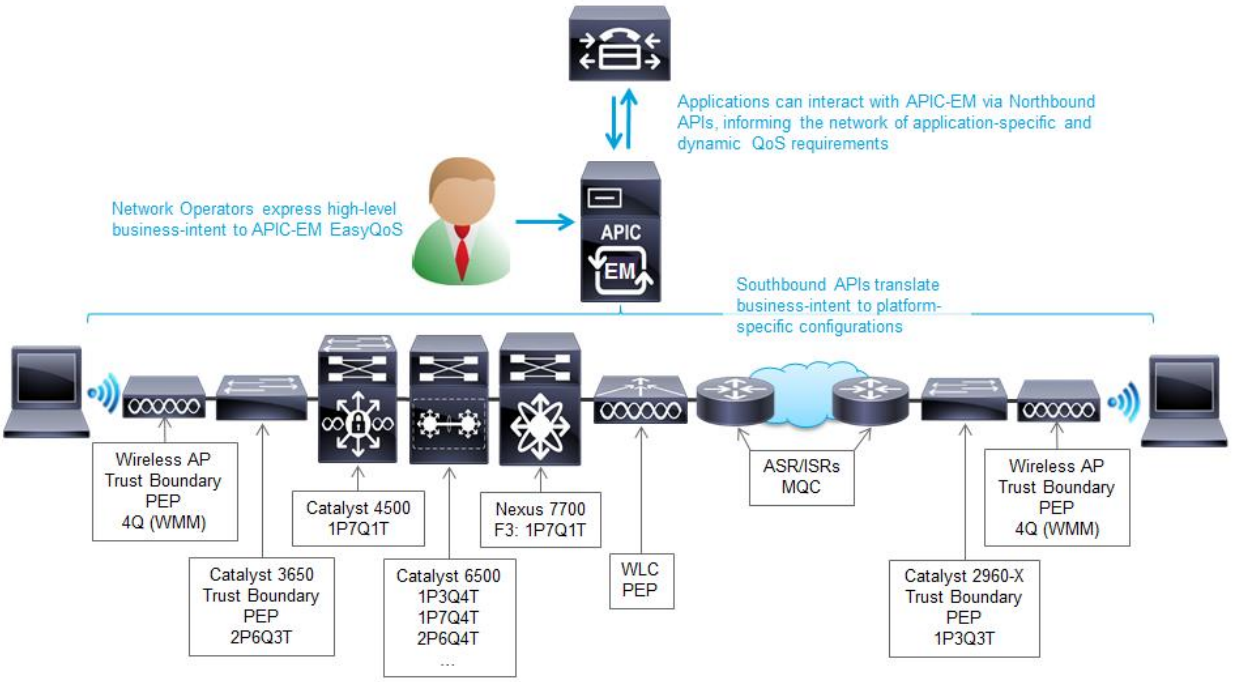

In the center of the figure is the APIC-EM controller with the EasyQoS application running on top of it. Network operators express their business intent directly through a web-based graphical user interface (GUI). EasyQoS then translates this business intent into platform specific configurations that are provisioned via southbound Application Programming Interfaces (APIs) onto groups of network infrastructure devices (referred to as policy scopes), based upon the application-level business intent. This functionality is referred to as *Static QoS* within this document.

Not only can a network controller simplify QoS/QoE deployments and accelerate these like never before, but it can also deliver completely new functionality in the form of application-integration. Traditionally applications have been separate and at arms-length from the network infrastructure, often with dedicated yet distinct teams of IT personnel to administer each. However, the role of the network isn't primarily to forward packets but rather to interconnect users via applications. As such, the network controller can play a crucial new role as the broker or intermediary between applications and the network. In order to do so, it has to understand the languages of each, which it does via two main types of APIs:

- Northbound API (NB API)/Northbound Interface (NBI): this interface allows for applications to communicate with the network controller, informing it of network policy requirements in real-time. Northbound APIs are commonly deployed with Representational State Transfer (REST) models.
- Southbound API (SB API)/Southbound Interface (SBI): this interface allows for the controller to communicate to individual network devices to configure the application policy-requirements. Southbound APIs include NETCONF/YANG models, as well as more traditional methods such as command line interface (CLI) and Simple Network Management Protocol (SNMP).

Specific to the context of QoE for collaboration, the network controller can receive information from the call-manager of the collaborative application–such as Cisco Unified Communications Manager (CUCM) for Cisco Jabber or Cisco WebEx or Cisco Spark–via the Northbound APIs, in order to inform it of any voice and/or

video calls that are proceeding on the network, providing it with the details of these flows. With this information, the controller can then quickly deploy QoS end-to-end across the enterprise for these voice and video calls, via the Southbound APIs. This functionality is referred to as Dynamic QoS within this document.

In summary the following is the business value of the EasyQoS solution:

- The EasyQoS solution provides end-to-end orchestration of QoS in the Enterprise network.
- The EasyQoS solution makes QoS policy simple and easy to deploy with an operator expressing business relevance for applications and the controller doing the rest under the hood.
- The EasyQoS solution works for both greenfield and brownfield deployments.
- The EasyQoS solution provides a declarative model that is business-intent driven, while abstracting away the platform/media/capability details.

Strategic QoS Policy

Over the past several years there has been an evolution in how Cisco approaches the deployment of QoS within organizations–revolving around the concept of policy abstraction. Traditionally when approaching QoS, the discussion quickly turns toward the tools that are used to implement QoS within the network infrastructure. The higher level conversation regarding the overall purpose for implementing QoS–that is, what you want to build with QoS–was often skipped. The challenge is to step back and see the bigger picture of how QoS connects to the business requirement first, before jumping in with the tools, as illustrated in the figure below.

Figure 3 **What Do You Consider First?**

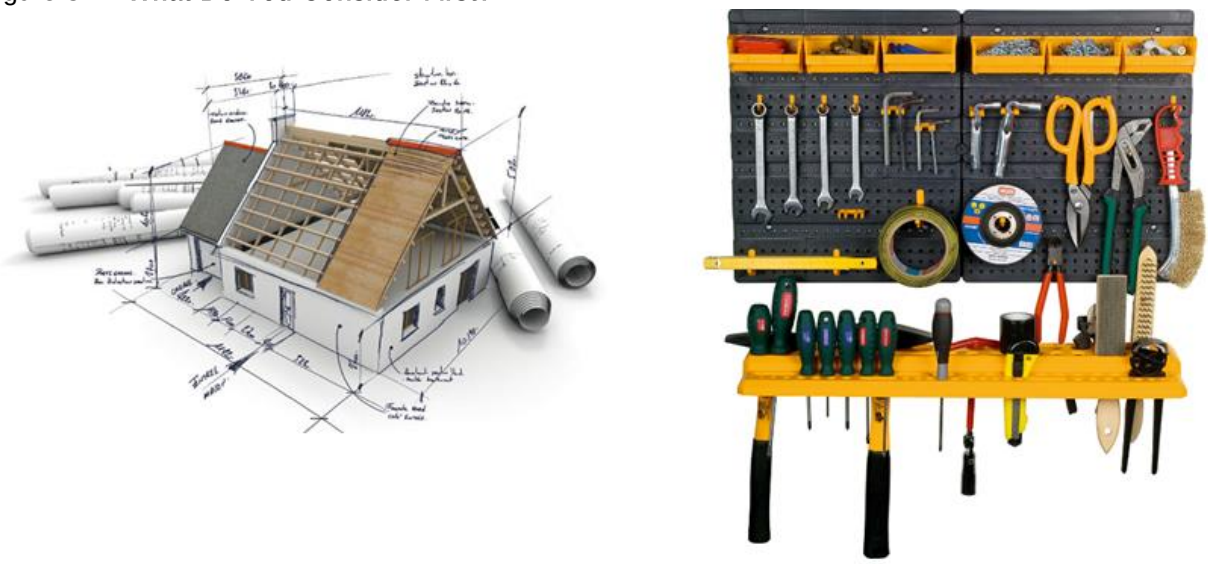

Where to Start?

The first step may seem obvious and superfluous, but in actuality it is crucial: clearly define the business objectives that your QoS policies are to enable. These may include any or all of the following:

- Guaranteeing voice quality meets enterprise standards
- Ensuring a high QoE for video
- Increasing user productivity by increasing network responsiveness for interactive applications
- Managing applications that are “bandwidth hogs”
- Identifying and de-prioritizing consumer applications
- Improving network availability
- Hardening the network infrastructure

Determining Application Business-Relevancy

With these goals in mind, network architects can clearly identify which applications are relevant to their business and which are not. There are three main states of business-relevance:

- Business-Relevant–these applications are known to contribute to business objectives of the organization and may include voice, multimedia applications, collaborative applications, database applications, email applications, file/content transfer applications, backup applications, etc., as well as control plane, signaling, and network management protocols.
- Default–these applications may or may not contribute to business objectives. For example, HTTP/HTTPS at times may be used for work or for personal reasons. As such, it may not always be possible to assign a static business-relevant designation to such applications (especially not without deeper packet inspection capabilities, which are not always available on all platforms).
- Business-Irrelevant–these applications are known to have no contribution to business-objectives and are often personal or entertainment-oriented in nature. Such applications may include video-on-demand (for example, Netflix, Hulu, YouTube, etc.), gaming traffic, peer-to-peer file-sharing applications, personal communication apps (for example, Skype, FaceTime, etc.) and other applications.

Figure 4 **Determining Application Business Relevance**

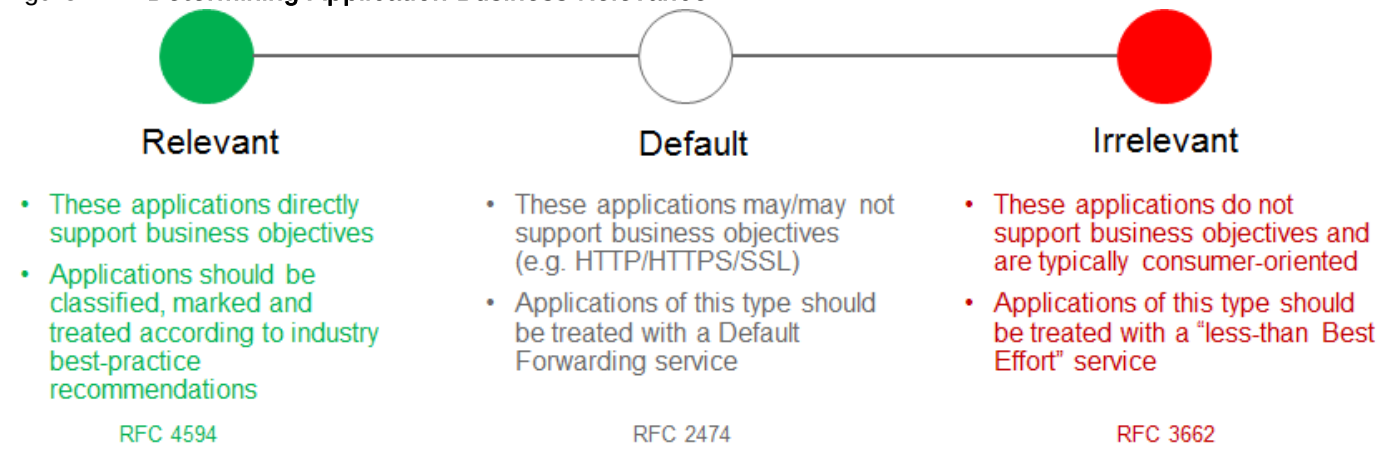

The 1300+ applications in Cisco's Network Based Application Recognition (NBAR) library have already been pre-programmed according to their most commonly-deployed level of business-relevance. This saves an operator from having to exhaustively go down a lengthy list and configure business-relevance one application at a time. However, the operator can override the default setting for business-relevance of any given application.

Mapping Business-Relevance to QoS Treatments

Cisco's RFC 4594-Based Strategic QoS Model

After applications have been defined as business-relevant (or otherwise), then the network architect must decide how to mark and treat these applications over the IP infrastructure. To this end, Cisco advocates following relevant industry guidelines, as this extends the effectiveness of your QoS policies beyond your

direct administrative control. That being said, it may be helpful to overview a relevant RFC for QoS marking and provisioning: RFC 4594, "Configuration Guidelines for DiffServ Service Classes."

These guidelines are to be viewed as industry best-practice recommendations. As such, enterprises and service providers are encouraged to adopt these marking and provisioning recommendations with the aim of improving QoS consistency, compatibility, and interoperability. However, it should be noted that these guidelines are not standards; as such, modifications can be made to these recommendations as specific needs or constraints require. Thus, to meet specific business requirements, Cisco has made a minor modification to its adoption of RFC 4594: specifically the swapping of Call-Signaling and Broadcast Video markings (to CS3 and CS5, respectively). A summary of Cisco's implementation of RFC 4594 is presented in the following figure.

Figure 5 **Cisco (RFC 4594-Based) QoS Recommendations**

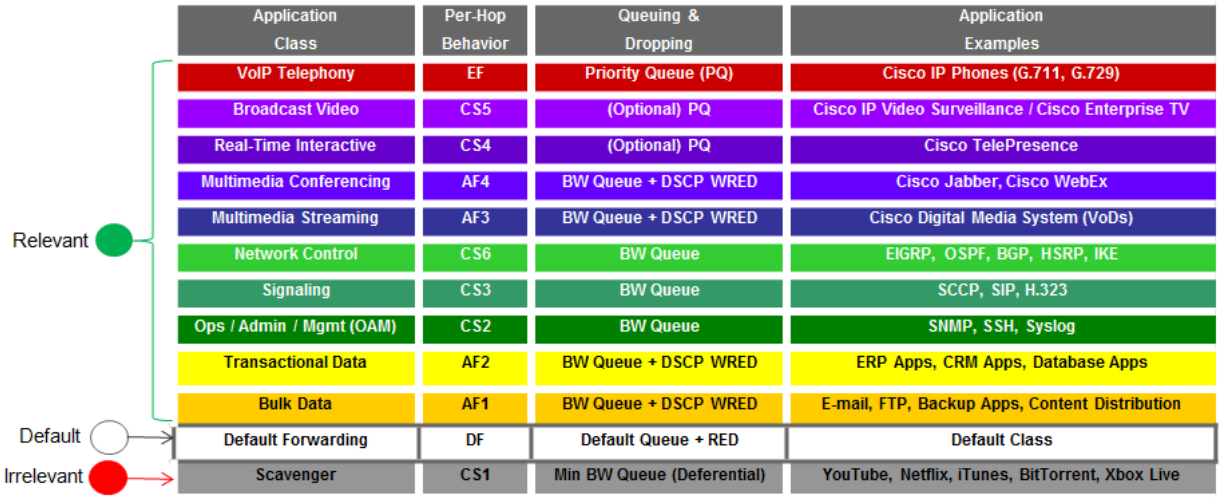

RFC 4594 also provides some application classification rules to help network architects to assign applications to the optimal traffic classes; these are summarized in the following sections.

QoS treatment for Business-Relevant Applications

Business relevant application can be grouped into one of four main categories:

- Control plane protocols
- Voice applications
- Video applications
- Data applications

Beginning with the control plane protocols, these may be subdivided further, as shown in the following figure.

Figure 6 **Control Plane Traffic Classes**

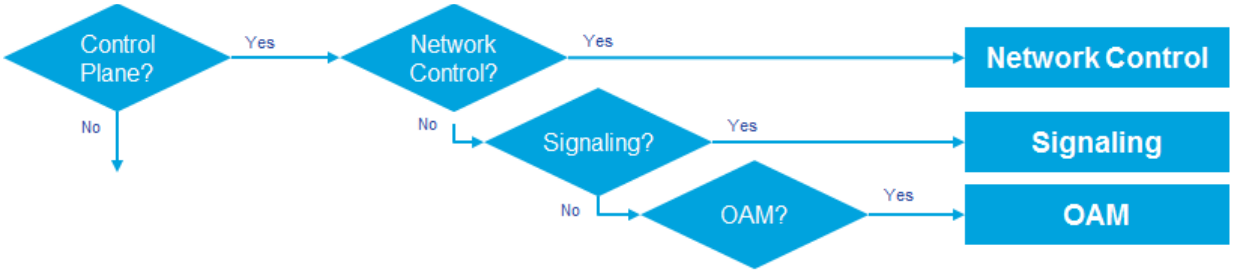

Network Control–This traffic class is intended for network control plane traffic, which is required for reliable operation of the enterprise network. Traffic in this class should be marked CS6 and provisioned with a (moderate but dedicated) guaranteed bandwidth queue. Weighted Random Early Detection (WRED) should not be enabled on this class, because network control traffic should not be dropped. Example traffic includes EIGRP, OSPF, BGP, HSRP, IKE, etc.

Signaling–This traffic class is intended for signaling traffic that supports IP voice and video telephony. Traffic in this class should be marked CS3 and provisioned with a (moderate but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, because signaling traffic should not be dropped. Example traffic includes SCCP, SIP, H.323, etc.

Operations/Administration/Management (OAM)–This traffic class is intended for network operations, administration, and management traffic. This class is critical to the ongoing maintenance and support of the network. Traffic in this class should be marked CS2 and provisioned with a (moderate but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, because OAM traffic should not be dropped. Example traffic includes SSH, SNMP, Syslog, etc.

Provisioning for voice is relatively straightforward, as shown in the following figure.

Figure 7 **Voice Traffic Class**

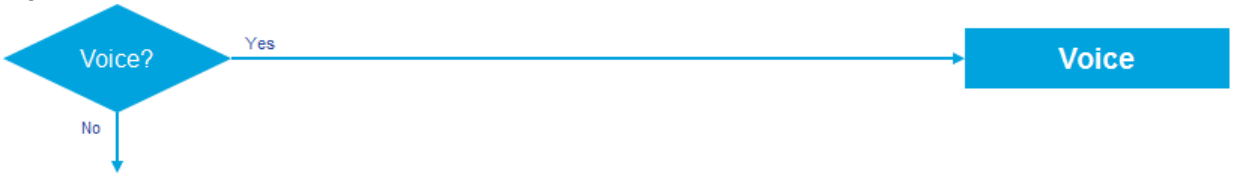

Voice–This traffic class is intended for voice/audio traffic (VoIP signaling traffic is assigned to the Call-Signaling class). Traffic assigned to this class should be marked EF. This class is provisioned with an Expedited Forwarding (EF) Per-Hop Behavior (PHB). The EF PHB defined in RFC 3246-is a strict-priority queuing service and, as such, admission to this class should be controlled. Example traffic includes G.711 and G.729a, as well as the audio components of multimedia conferencing applications, such as Cisco Jabber, WebEx, and Spark.

Video–This traffic class may have unique QoS requirements depending on the type of video, as illustrated in the following figure.

Figure 8 **Video Traffic Classes**

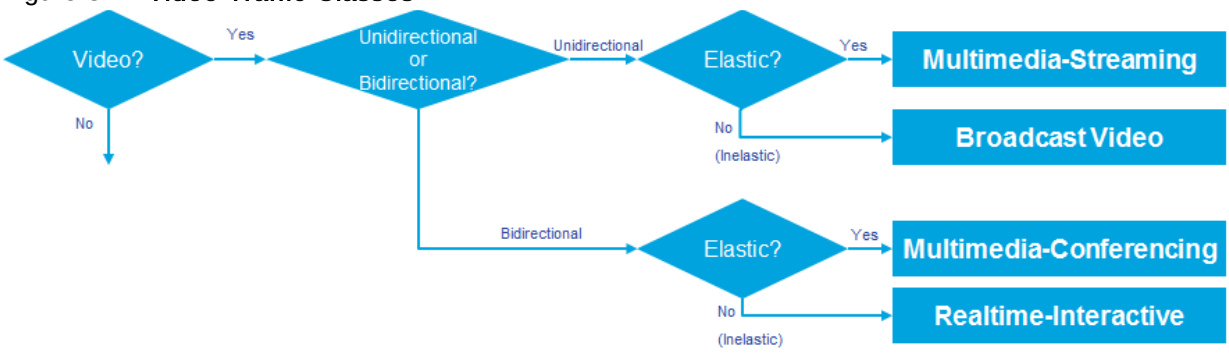

To determine the optimal traffic classification for a video application, two key questions need to be answered:

- Is the video unidirectional or bidirectional?
- Is the video elastic or inelastic?

Elastic flows are able to adapt to network congestion and/or drops (by reducing frame rates, bit rates, compression rates, etc.). *Inelastic* flows either do not have such capabilities or–in order to meet specific business requirements–are configured not to use these.

With these two questions answered, video applications may be assigned to their respective traffic classes, including the following.

Broadcast Video–This traffic class is intended for broadcast TV, live events, video surveillance flows, and similar inelastic streaming video flows. Traffic in this class should be marked CS5 and may be provisioned with an EF PHB; as such, admission to this class should be controlled. Example traffic includes live Cisco Enterprise TV streams, and Cisco IP Video Surveillance.

Real-Time Interactive–This traffic class is intended for inelastic interactive video applications. Whenever possible, signaling and data sub-components of this class should be separated out and assigned to their respective traffic classes. Traffic in this class should be marked CS4 and may be provisioned with an EF PHB; as such, admission to this class should be controlled. An example application is Cisco TelePresence.

Multimedia Conferencing–This traffic class is intended for elastic interactive multimedia collaboration applications. Whenever possible, signaling and data subcomponents of this class should be separated out and assigned to their respective traffic classes. Traffic in this class should be marked Assured Forwarding (AF) Class 4 (AF41) and should be provisioned with a guaranteed bandwidth queue with Differentiated Services Code Point-based Weighted-Random Early Detect (DSCP-WRED) enabled. Traffic in this class may be subject to policing and re-marking. Example applications include Cisco Jabber, WebEx, and Spark.

Multimedia Streaming–This traffic class is intended for elastic streaming video applications, such as Video-on-Demand (VoD). Traffic in this class should be marked AF Class 3 (AF31) and should be provisioned with a guaranteed bandwidth queue with DSCP-based WRED enabled. Example applications include Cisco Digital Media System VoD streams, ELearning videos, etc.

Figure 9 **Data Traffic Classes**

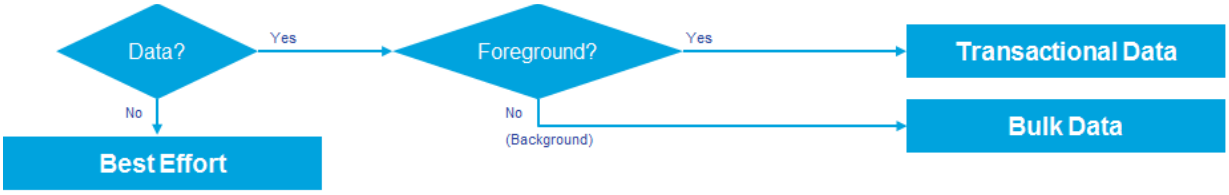

When it comes to data applications, there is really only one key question to answer (as illustrated in the figure above): is the data application foreground or background?

Foreground refers to applications from which users expect a response–via the network–in order to continue with their tasks. Excessive latency to such applications directly impact user productivity. Conversely, *background* applications–while business relevant–do not directly impact user productivity and typically consist of machine-to-machine flows.

Transactional Data–This traffic class is intended for interactive, foreground data applications. Traffic in this class should be marked AF Class 2 (AF21) and should be provisioned with a dedicated bandwidth queue with DSCP-based WRED enabled. This traffic class may be subject to policing and re-marking. Example applications include data components of multimedia collaboration applications, Enterprise Resource Planning applications, Customer Relationship Management applications, database applications, etc.

Bulk Data–This traffic class is intended for non-interactive background data applications. Traffic in this class should be marked AF Class 1 (AF11) and should be provisioned with a dedicated bandwidth queue with DSCP-based WRED enabled. This traffic class may be subject to policing and re-marking. Example applications include: email, backup operations, FTP/SFTP transfers, video and content distribution, etc.

With all business-relevant applications assigned to their respective traffic classes, only two types of traffic classes are left to be provisioned–Default and Scavenger traffic classes.

QoS Treatment for Default-Business Relevant Applications

Best Effort–This traffic class is the default class. The vast majority of applications will continue to default to this Best-Effort service class. As such, the default class should be adequately provisioned. Traffic in this class is marked Default Forwarding (DF or DSCP 0) and should be provisioned with a dedicated queue. It is recommended that you enable WRED on this class.

QoS Treatment for Business-Irrelevant Applications

Scavenger–This traffic class is intended for all applications that have been previously identified as business-irrelevant. These may include video applications that are consumer and/or entertainment-oriented. The approach of a "less-than Best-Effort" service class for non-business applications (as opposed to shutting these down entirely) has proven to be a popular political compromise.

Applications within the Scavenger traffic class are permitted on business networks when bandwidth is available. However, as soon as the network experiences congestion, this class is the most aggressively dropped. Traffic in this class should be marked CS1 and should be provisioned with a minimal bandwidth queue, which is the first to starve should network congestion occur. Example traffic includes Netflix, YouTube, Xbox Live/360 Movies, iTunes, BitTorrent, etc.

Tactical QoS Policy

Translating QoS Strategy into Tactical Designs

To meet the demands of today's media-rich networks, administrators should articulate a QoS strategy that reflects their business intent. This strategy details which applications are business relevant and which applications are not business relevant, as well as how these applications are to be marked and treated over the IP network. Furthermore, this QoS strategy is end-to-end and is not constrained by any technical or administrative limitation.

While defining such an unconstrained QoS strategy is an important part of the deployment process, when it comes to practical deployment, various technical constraints have to be taken into account, including the following:

- Hardware constraints
- Software constraints
- Media capability constraints
- Bandwidth constraints
- Service provider constraints

Thus the goal of tactical QoS design is to adapt the QoS strategy to the maximum of each platform's capabilities, subject to all relevant constraints.

The following are additional recommendations to keep in mind during the tactical design phase:

- Only enable QoS features if these directly contribute to expressing the QoS strategy on the given platform.
- Leverage QoS design best-practices to generate platform specific configurations that reflect the QoS strategy with maximum fidelity.

QoS Design Best Practices

The following sections discuss generic best practices for QoS design.

Hardware vs. Software QoS Design

Some Cisco routers, such as Cisco Integrated Services Routers (ISRs), perform QoS in software, which places incremental loads on the CPU. The actual incremental load depends on the numerous factors, including: the complexity and functionality of the policy, the volume and composition of the traffic, the speed of the interface, the speed of the CPU, the memory of the router, etc. On the other hand, other devices (such as Cisco Catalyst switches) often perform QoS in dedicated hardware–Application Specific Integrated Circuits (ASICs). As such, these switches can perform even the most complex QoS policy on maximum traffic loads at line rates on GE/10GE/40GE/100GE interfaces–all without any marginal CPU tax. Thus, whenever a choice exists, Cisco recommends implementing QoS policies in devices that perform QoS

operations in hardware—rather than software—as this will result in more efficient utilization of network infrastructure resources.

For example, suppose an administrator has the option of deploying classification and marking policies in a branch network in either a Catalyst switch (in hardware) or at the LAN-edge interface of an ISR router (in software). Since a choice exists as to where the policy should be deployed, it would be more efficient to classify and mark within the Catalyst switch.

However, there may be cases where such a choice doesn't exist. Continuing the example: there may be a business need to perform deep-packet inspection on branch-originated traffic (which isn't currently supported on Catalyst switches), and as such the administrator would then have to apply the required classification and marking policies on the ISR router.

Classification and Marking Best Practices

When classifying and marking traffic, a recommended design best practice is to classify and mark applications as close to their sources as technically and administratively feasible. This principle promotes end-to-end differentiated services and PHBs.

In general, it is not recommended that you trust markings that can be set by end users on their PCs or other similar devices because end users can easily abuse provisioned QoS policies if permitted to mark their own traffic. For example, if an EF PHB has been provisioned over the network, a PC user can easily configure all their traffic to be marked to EF, thus hijacking network priority queues to service their non-real-time traffic. Such abuse could easily ruin the service quality of real-time applications throughout the enterprise. On the other hand, if enterprise controls are in place to centrally administer PC QoS markings, then it may be an acceptable design option to trust them.

Following this rule, it is further recommended that you use DSCP markings whenever possible, because these Layer 3 IP-header markings are end-to-end, more granular, and more extensible than Layer 2 markings. For example, IEEE 802.1p, IEEE 802.11e (now part of the IEEE 802.11 standard) and MPLS EXP only support three bits (values 0-7) for marking. Therefore, only up to eight classes of traffic can be supported with these marking schemes and inter-class relative priority (such as RFC 2597 Assured Forwarding drop preference markdown) is not supported. On the other hand, Layer 3 DSCP markings allow for up to 64 distinct classes of traffic.

As the line between enterprises and service providers continues to blur and the need for interoperability and complementary QoS markings is critical, you should follow standards-based DSCP PHB markings to ensure interoperability and future expansion.

Policing and Remarking Best Practices

There is little reason to forward unwanted traffic only to police and drop it at a downstream node. Therefore, it is recommended that you police traffic flows as close to their sources as possible.

Whenever supported, markdown should be done according to standards-based rules, such as RFC 2597, the Assured Forwarding PHB. For example, excess traffic marked to AFx1 should be marked down to AFx2 (or AFx3 whenever dual-rate policing—such as defined in RFC 2698—is supported). Following such markdowns, congestion management policies, such as DSCP-based WRED, should be configured to drop AFx3 more aggressively than AFx2, which in turn should be dropped more aggressively than AFx1.

Congestion Management (Queuing) Best Practices

Business-critical applications require service guarantees regardless of network conditions. The only way to provide service guarantees is to enable queuing at any and every node that has the potential for congestion.

In addition, because each application class has unique service level requirements, optimally each should be assigned a dedicated queue. In such a manner, specific bandwidth allocations and dropping policies can be assigned to each discrete application class to meet its distinctive QoS requirements. Otherwise, if multiple application classes are assigned into a common queuing bucket, the administrator no longer can control if bandwidth resources are being shared among these application classes according to their individual requirements.

At a minimum, however, the following standards-based queuing behaviors should be supported:

- Real-time queue(s)—to support an RFC 3246 Expedite Forwarding service
- Guaranteed-bandwidth queue(s) —to support RFC 2597 Assured Forwarding services
- Default queue—to support an RFC 2474 Default Forwarding service
- Bandwidth—constrained queue-to support an RFC 3662 Scavenger service

Cisco offers design recommendations for each of these types of queues. These queuing best practices are illustrated in the following figure.

Figure 10 **Queuing Best Practices**

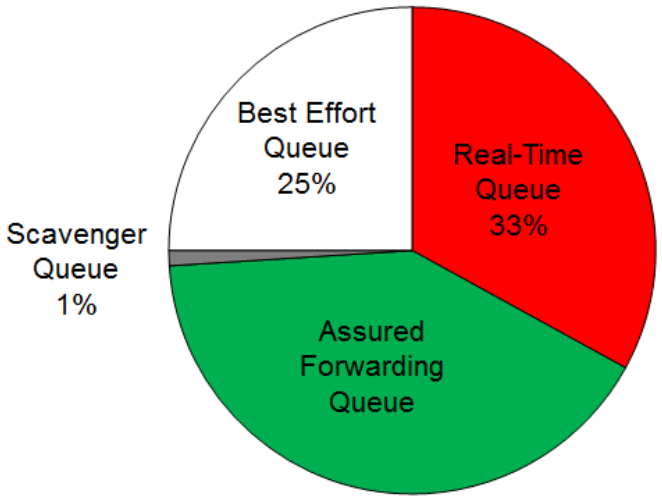

Real-Time Queue

The Real-Time queue corresponds to the RFC 3246 EF PHB. The amount of bandwidth assigned to the Real-Time queue is usually variable. However, if the majority of bandwidth is provisioned with strict-priority queuing (which is effectively a first-in, first-out queue), the overall effect is a dampening of QoS functionality. Remember the goal of convergence is to enable voice, video, and data applications to transparently coexist on a single network. When real-time applications dominate a link, non-real-time applications fluctuate significantly in their response times, destroying the transparency of the converged network.

Cisco has done extensive testing and has found that a significant decrease in non-real-time application response times occurs when real-time traffic exceeds one-third of link bandwidth capacity. In fact, both testing and customer deployments have shown that a general best queuing practice is to limit the amount of strict priority queuing to 33% of link bandwidth capacity. This strict priority queuing recommendation is a conservative and safe design ratio for merging real-time applications with data applications.

Finally, WRED—or any similar congestion avoidance mechanism—should never be enabled on the strict priority queue. Traffic assigned to this queue is often highly drop sensitive; therefore, early dropping should never be induced on these flows.

Assured Forwarding Queue

At least one queue should be provisioned as an Assured Forwarding Queue. Per RFC 2597, up to four queues can be provisioned with this service:

- AF Class 1—AF11, AF12, AF13
- AF Class 2—AF21, AF22, AF23
- AF Class 3—AF31, AF32, AF33
- AF Class 4—AF41, AF42, AF43

These queues should have bandwidth guarantees that correspond with the application class requirements of the traffic assigned to it.

In addition, DSCP-based WRED should be enabled on these queues, such that traffic marked AFx3 is (statistically) dropped sooner and more often than AFx2, which in turn is (statistically) dropped more aggressively than AFx1.

Best Effort Queue

The Best Effort Queue is the default treatment for all traffic that has not been explicitly assigned to another queue. Only if an application has been selected for preferential/deferential treatment is it removed from the default class. Because most enterprises have several thousand applications running over their networks, adequate bandwidth must be provisioned for this class as a whole to handle the sheer number and volume of applications that default to it. Therefore, Cisco recommends provisioning at least 25% of link bandwidth for the default Best Effort class.

In addition, it is recommended that you enable WRED on the default class to improve throughput and reduce TCP synchronization. Because all traffic destined to this class is to be marked to the same DSCP value (of 0), there is no "weight" component to the WRED dropping decision, and therefore the congestion algorithm is effectively random early detect.

Less-Than-Best-Effort (Scavenger) Queue

Whenever the Scavenger Queue is enabled, it should be assigned a minimal amount of bandwidth, such as 1% (or whatever the minimal bandwidth allocation that the platform supports).

WRED is not required on the Scavenger class queue because traffic assigned to this queue has no implied "good-faith" service guarantee or expectation. Therefore, there is little to gain by adding this feature and it may even be wasteful of router CPU resources.

APIC-EM and the EasyQoS Application

The Application Policy Infrastructure Controller–Enterprise Module (APIC-EM) is Cisco's enterprise SDN controller. EasyQoS is one of several applications which run on APIC-EM. The following sections discuss how to access APIC-EM, declaratively express QoS policies within the EasyQoS application, and then deploy those QoS policies to groups of network infrastructure devices.

Logging Into APIC-EM

APIC-EM provides a web-based GUI for configuring and monitoring the base APIC-EM functionality as well as the applications that reside upon it.

Note: APIC-EM also includes an extensive set of northbound REST-based APIs for configuring and monitoring APIC-EM functionality and the applications that reside upon it. This version of the APIC-EM EasyQoS Design Guide does not cover the northbound REST-based APIs. Future versions may include a discussion of the northbound APIs.

In order to access the APIC-EM login page, the network operator must launch a web browser and open an HTTPS connection to the IP address or fully qualified domain name of the APIC-EM server. An example of the login page is shown in the figure below.

Figure 11 **APIC-EM Login Screen**

Upon entering the proper login credentials (username and password) and clicking on the Log In button, the network administrator will be taken to the APIC-EM Home page, as shown below.

Figure 12 **APIC-EM Home Page**

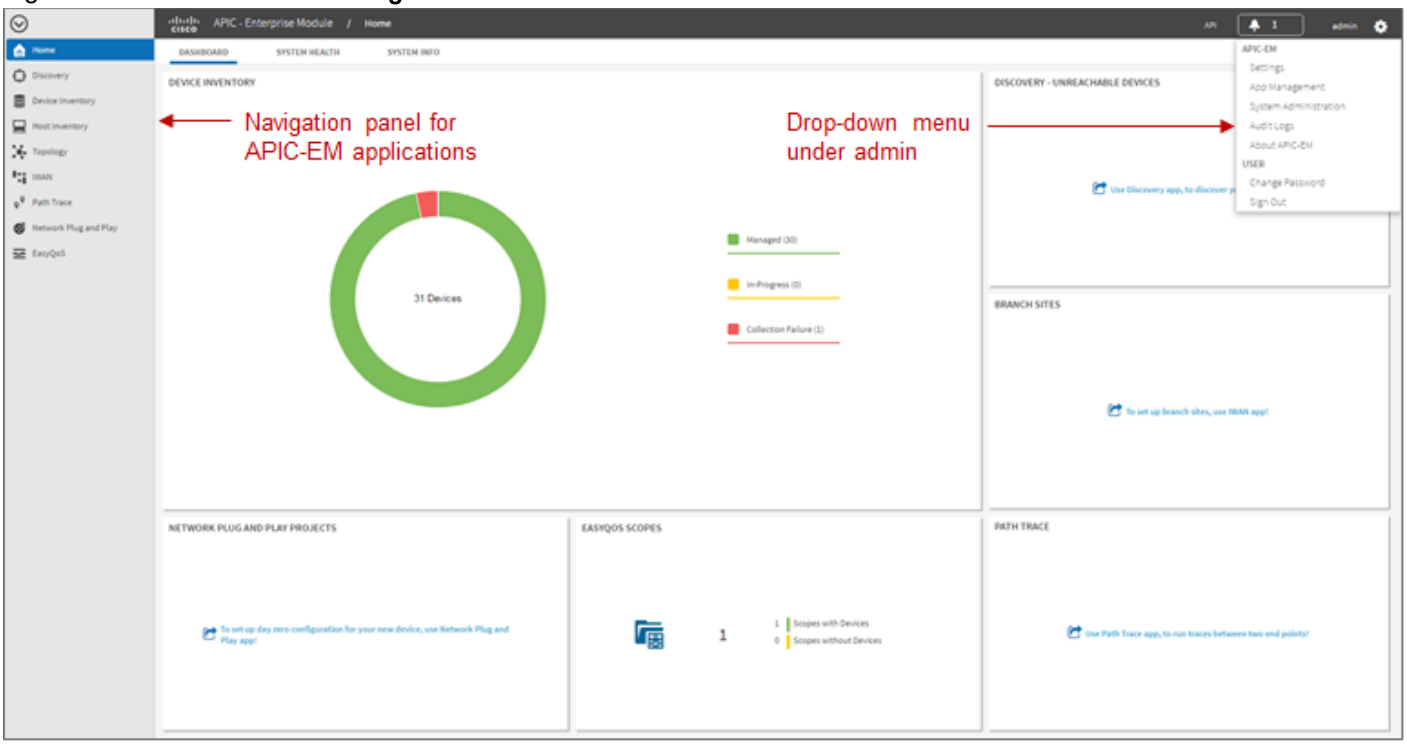

APIC-EM supports both integration with an external AAA server via the RADIUS protocol, as well as an Internal Users database locally administered on the APIC-EM server. Both are accessed from Settings within the drop-down menu that appears when clicking on "admin" in the upper right-hand corner of any APIC-EM page, as shown in the figure above.

The expandable navigation panel on the left-side of any APIC-EM page displays the various applications (or functions) available within APIC-EM. The following four applications (or functions) provided by the APIC-EM controller are not part of the EasyQoS application itself but are discussed within this document, because they provide necessary functionality for the EasyQoS application to operate.

- Network Device Discovery
- Device Inventory
- Host Inventory
- Topology

This document does not discuss the IWAN or Network Plug and Play applications, because they do not directly provide functionality that is required for the EasyQoS application. The Path Trace application is briefly mentioned at the end of the *EasyQoS Application* section of this document.

Network Device Discovery

In order to apply QoS Policies to network devices within the EasyQoS application, network devices must first be discovered, added the APIC-EM device inventory, and managed by APIC-EM. Hence, the network operator must first perform a Discovery in order to discover network devices and place them into the Device

Inventory database. Clicking on the Discovery icon with the expandable panel on left-side of any APIC-EM page takes the network operator to the Discovery page. An example is shown in the following figure.

Figure 13 **APIC-EM Discovery Page**

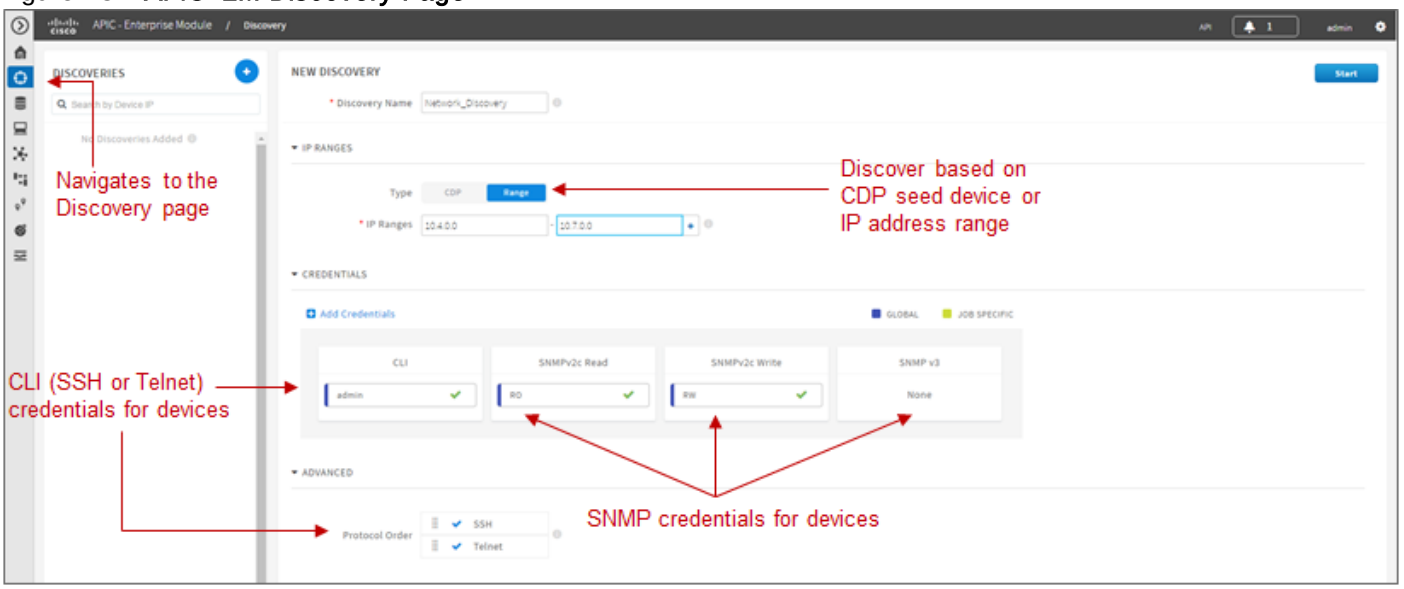

The Discovery page allows the network operator to create Discovery jobs based on either a Cisco Discovery Protocol (CDP) seed device or an IP address range. The Discovery process requires SNMP credentials (v2c, or v3) and CLI credentials (SSH or Telnet) to be entered. The network operator must ensure that CLI credentials–including an enable password–have been previously configured on the network device and that the device is network reachable from APIC-EM.

SNMP and CLI credentials for devices can be created and stored via Settings within the drop-down menu that appears when clicking on “admin” in the upper right-hand corner of any APIC-EM page. The following screen shot shows an example of where the CLI credentials are configured within APIC-EM.

Figure 14 **Adding CLI Credentials to APIC-EM**

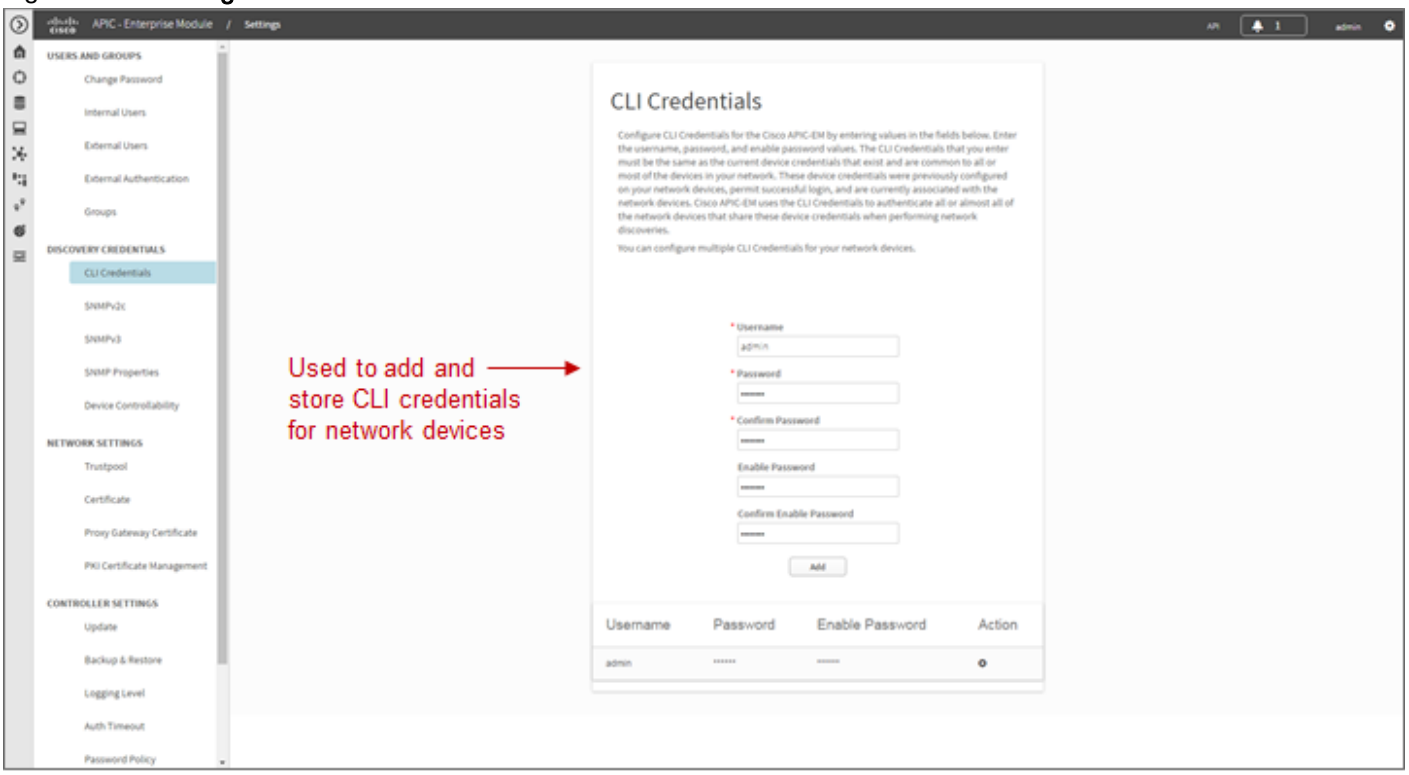

SNMP credentials can be added similarly. The saved credentials can then be referenced within the Discovery process.

APIC-EM release 1.3 has added a Device Controllability feature. This feature allows APIC-EM to configure the SNMP credentials onto network infrastructure devices–via the CLI interface. This feature can save time by not requiring the network operator to manually access each network infrastructure device and configure the SNMP credentials. The following figure shows the screen that appears when selecting Device Controllability.

Figure 15 **Device Controllability Feature**

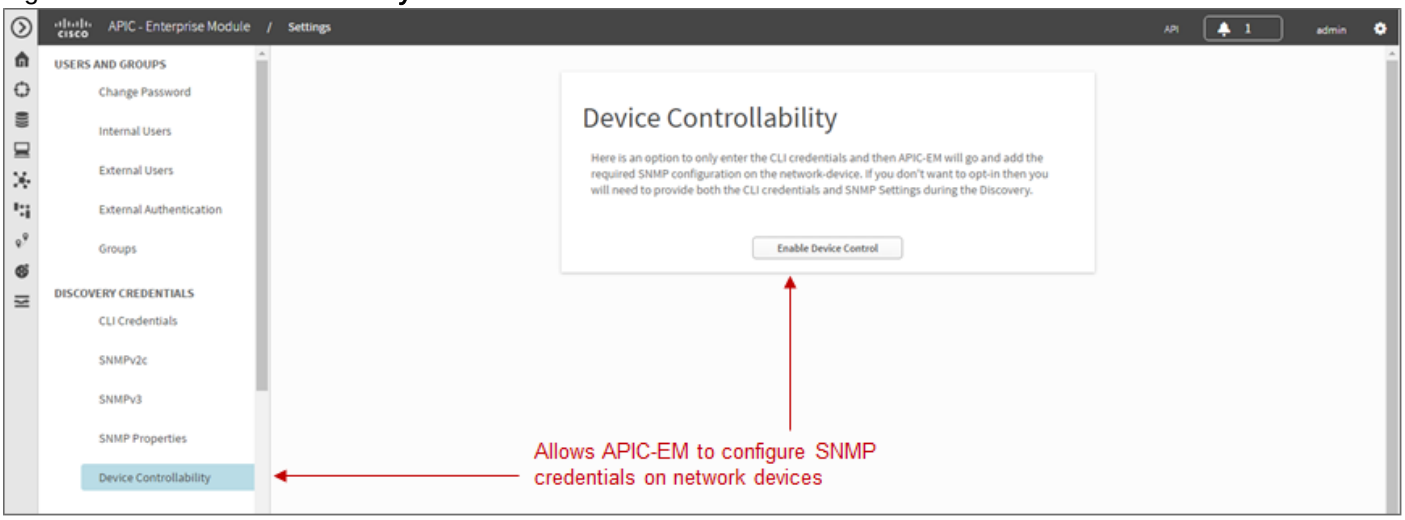

If the Device Controllability feature is disabled (which is the default setting), the network operator must ensure that SNMP access is configured on each network device and that each device is SNMP reachable from APIC-EM.

The SNMP Properties page can be used to modify the number of SNMP retries and the timeout between each attempt. The network operator may find it useful to modify the default settings if some network devices occasionally experience SNMP timeouts when APIC-EM attempts to synchronize its Device Inventory database with those devices. An example of the SNMP Properties page is shown below.

Figure 16 **SNMP Properties**

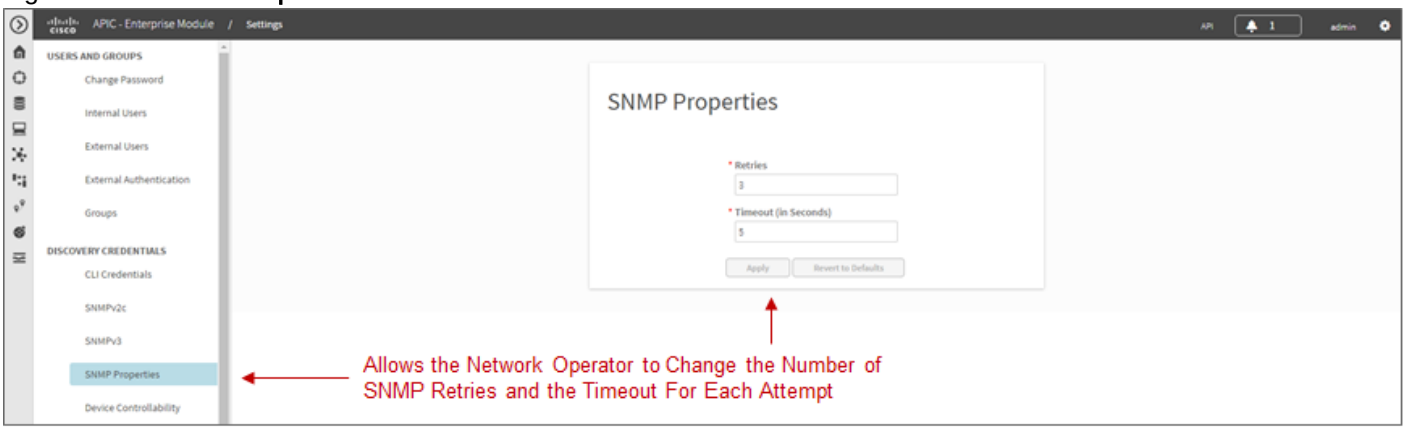

Device Inventory

Only after network devices have been discovered will those devices be added the APIC-EM Device Inventory database and managed by APIC-EM. Clicking on the Device Inventory icon within the expandable panel on left-side of any APIC-EM page takes the network operator to the Device Inventory page. An example is shown in the following figure.

Figure 17 **APIC-EM Device Inventory Page**

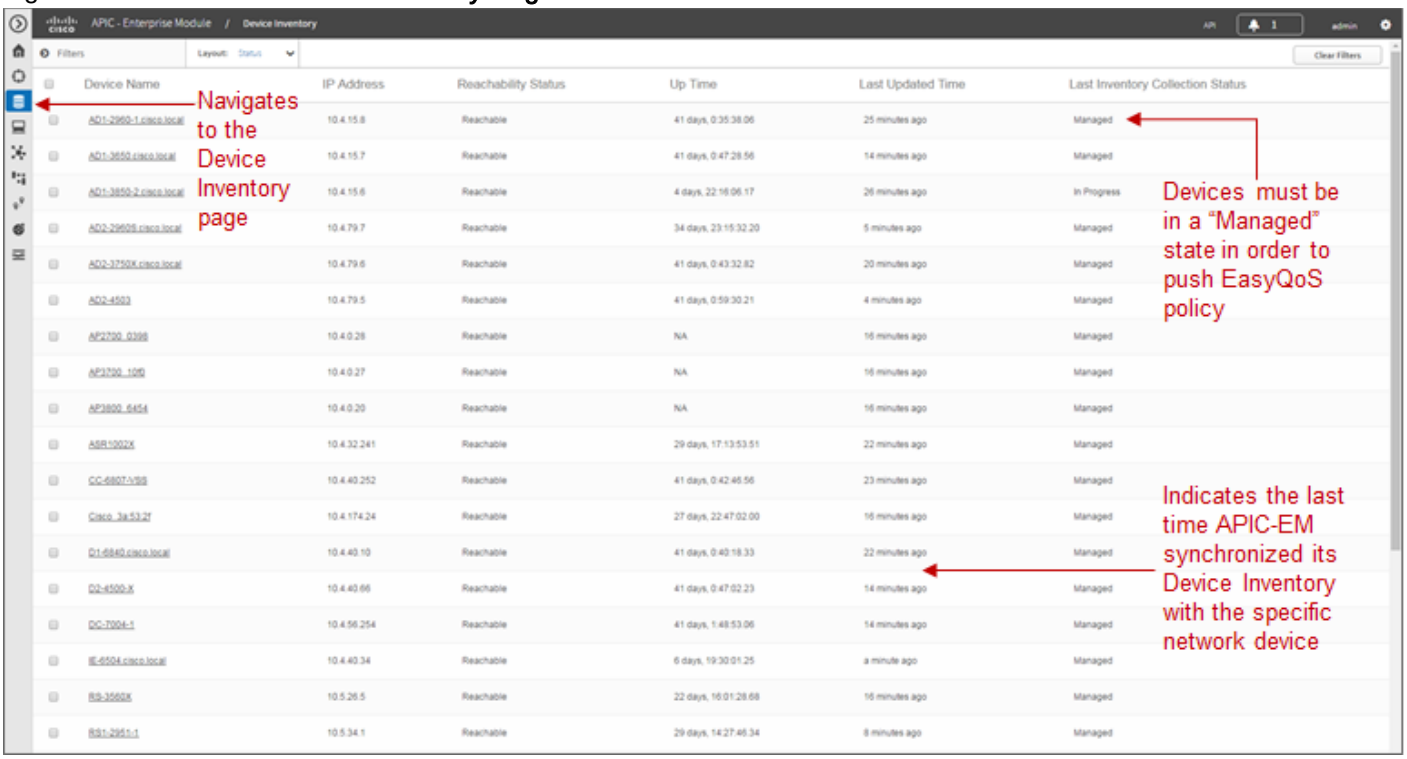

As mentioned previously, devices must be in a Managed state in order to provision EasyQoS policy. APIC-EM periodically (approximately every 30 minutes) synchronizes the Device Inventory database with each network device. There is currently no manual way of initiating this synchronization within APIC-EM. If changes to the configuration of a particular network device have been made either via CLI or via management platforms such as Prime Infrastructure, it is recommended that you wait until APIC-EM has re-synchronized with the device, in order to ensure the configuration changes have been identified by APIC-EM, before applying any changes to QoS policy through EasyQoS.

An example would be adding or changing a WAN service provider profile (SPP) tag manually via the CLI to an ISR or Cisco Aggregation Services Router (ASR) WAN interface. APIC-EM would have to re-synchronize with the ISR or ASR in order to be aware of the updated WAN SPP tag first, in order to apply the appropriate QoS policy to the WAN interface.

Host Inventory

Cisco Device Endpoints

APIC-EM also discovers certain Cisco hardware device endpoints which are then included within Static and Dynamic QoS policies provisioned to network infrastructure devices. These hardware endpoints include the following:

- Cisco IP phones
- Cisco TelePresence devices
- Cisco video conferencing endpoints

- Cisco video surveillance cameras

APIC-EM makes use of CDP running on Catalyst switches in order to discover these hardware endpoints. EasyQoS uses the IP addresses of the hardware endpoints collected through CDP information, along with the knowledge of which Catalyst switch and switch port the endpoint is connected to–in order to pre-populate access control entries (ACEs) within classification & marking access control lists (ACLs) for Static and Dynamic QoS on switching devices. In order for this functionality to operate, CDP must be enabled on the Catalyst switch ports that connect to hardware endpoint devices. By default CDP is enabled on Cisco Catalyst switch ports.

The Cisco hardware endpoint devices themselves must also support CDP. The CDP information provided by the hardware endpoint must also include its IP address. If the IP address is not included, APIC-EM will not know which switch port to populate with ACE entries.

Note: Older versions of Cisco TelePresence code may not support the sending of IP addresses within CDP when in a VLAN configuration–such as when a Voice VLAN is configured. Such systems may require an upgrade to TC7.3.6, CE8.0.2, or CE8.1.1 or higher in order for these devices to be statically populated within the correct ingress classification & marking ACLs when deploying EasyQoS. Please refer to Cisco defect CSCuy71139 for details.

Discovered Cisco endpoint devices are populated within the Host Inventory database within APIC-EM. These can be displayed by clicking on the Host Inventory icon with the expandable panel on left-side of any APIC-EM page. An example of the Host Inventory page is shown in the following figure.

Figure 18 **APIC-EM Host Inventory Page**

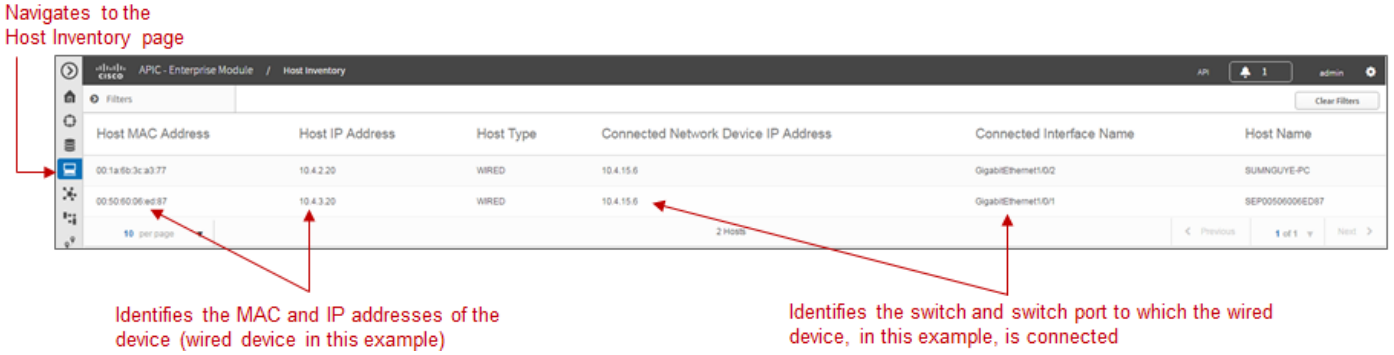

For wired Cisco device endpoints, after the endpoint information is collected, APIC-EM provisions ACE entries into the ACLs configured for Static QoS corresponding to the ingress classification & marking policy deployed across all access-edge ports on the switch to which the device is connected. This is discussed in detail in the *Cisco Device Endpoints* section of this document. If Dynamic QoS is enabled, EasyQoS will also push ACE entries into the Dynamic ACL policy shells corresponding to the dynamic ingress classification & marking policy for the specific switch port as well. Dynamic QoS is discussed in detail within the *Dynamic QoS Design* section of this document.

There are no equivalent ACE entries generated for wireless devices with the current EasyQoS solution. This is because the AireOS wireless LAN controller (WLC) EasyQoS ingress classification & marking policy uses Cisco Application Visibility and Control (AVC) profiles, rather than Layer 2-4 ACLs.

Topology

After network infrastructure devices have been discovered, the network operator can view the network via the Topology page. The Topology page is accessed by clicking on the Topology icon within the expandable panel on left-side of any APIC-EM page. An example is shown in the following figure.

Figure 19 **APIC-EM Topology Page**

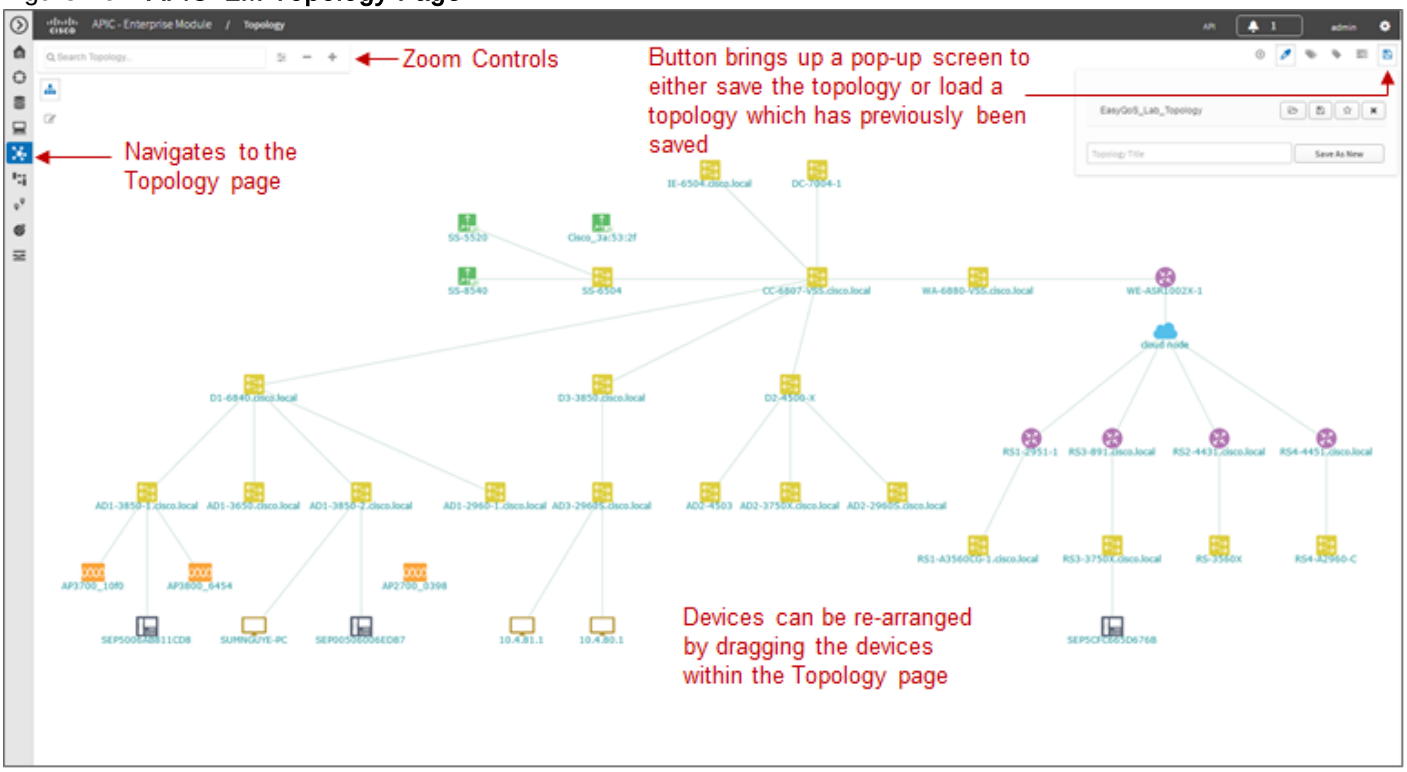

APIC-EM automatically discovers the relationship between devices and connects them together within the Topology page. Individual devices or groups of devices can be re-positioned by dragging them around within the page and by zooming-in and zooming-out as needed. When the network operator has arranged the devices as desired, he/she can save the layout via the Save or Load Topology icon in the upper right corner of the Topology page. This can be loaded in the future when visiting the Topology page, so that the network operator doesn't have to re-arrange the devices upon every visit to the Topology page.

Clicking on a device will bring up a side window with additional detail on the device–including the role of the device within the network–as shown in the figure below.

Figure 20 **Selecting the Role of a Device from Within the Topology Page**

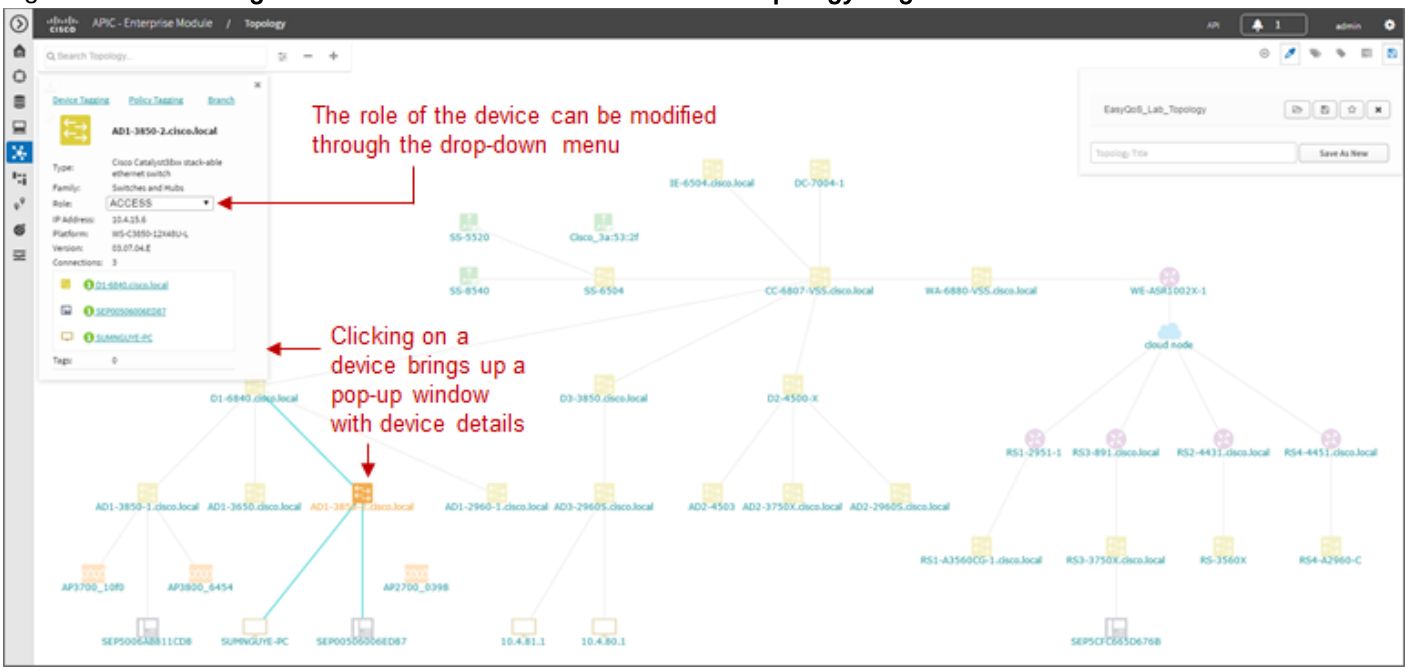

Each discovered network infrastructure device is automatically categorized into one of the following roles:

- Core
- Distribution
- Access
- Border Router
- Unknown

The Core, Distribution, and Access roles apply only to Catalyst switches. The Border Router role applies only to Cisco ISRs and ASRs. The network operator should verify that the particular device selected has been characterized with the correct role, in order to ensure the correct QoS policy is applied to the device by the EasyQoS application. This applies primarily to Catalyst switches. If necessary, the network operator can change the role within the side window. The policy applied to Catalyst switches based upon their role is discussed in the *Campus LAN Static QoS Design* section of this document.

EasyQoS Application

The Discovery, Device Inventory, Host Inventory, and Topology functions discussed in the previous sections are not part of the EasyQoS application. However, they were discussed because the functionality they provide is necessary for the EasyQoS application itself to operate. This section shifts the discussion to the specific functionality within the EasyQoS application itself.

The EasyQoS application is accessed by clicking on the EasyQoS icon within the expandable panel on left-side of any APIC-EM page. An example is shown in the following figure.

Figure 21 **APIC-EM EasyQoS**

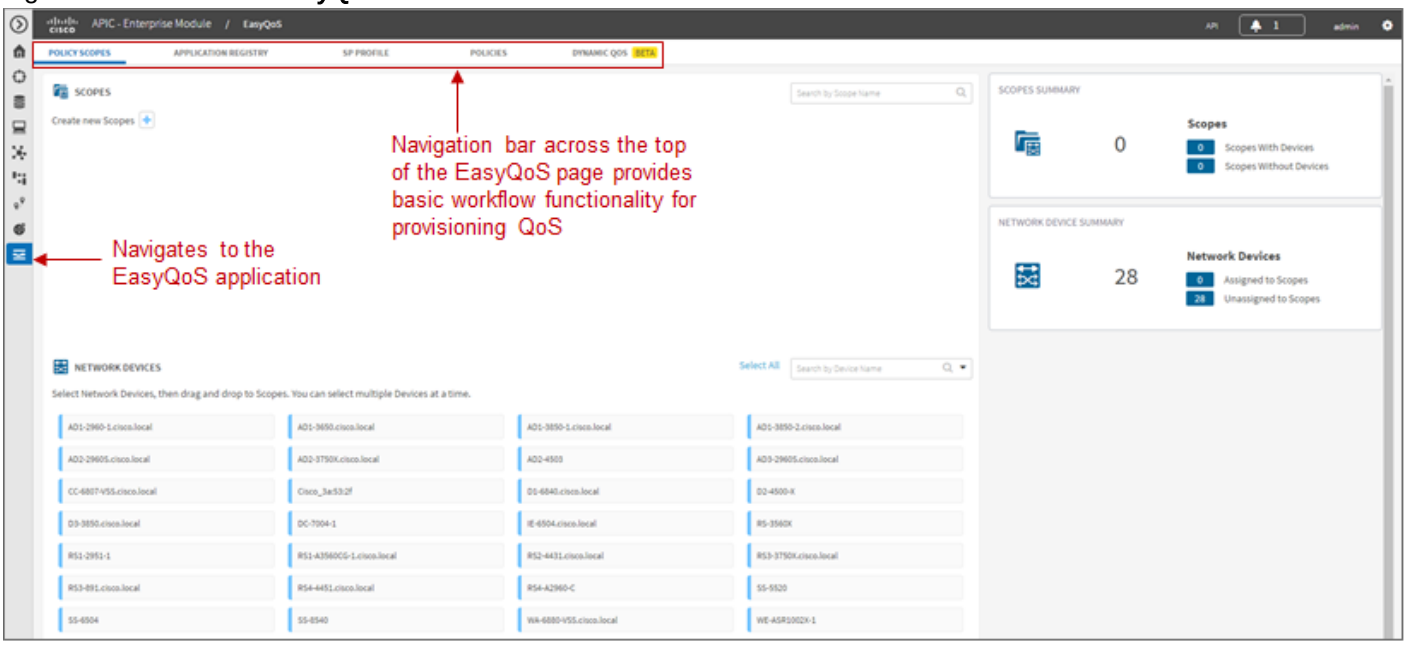

As shown in the figure above, the EasyQoS application has several tabs that appear as a bar across the top of the page.

- Policy Scopes
- Application Registry
- SP Profile
- Policies
- Dynamic QoS

The tabs are intended to roughly guide the network operator through something similar to a basic workflow for deploying QoS policy. Therefore, it is recommended that you access the tabs in order (from left to right) when deploying QoS policy, although the network operator is free to access the tabs in any order. Each of these tabs will be discussed in separate sections.

Policy Scopes

The network operator is by default automatically taken to the first tab—Policy Scopes—when clicking on the EasyQoS application icon within the expandable panel on left-side of any APIC-EM page.

The first step to deploying QoS policy through EasyQoS is to create one or more policy scopes. Policy scopes are simply a way of grouping one or more network devices together in order to apply QoS policy to the group all at once, rather than having to individually apply QoS policy to one network device at a time.

The network operator can define a single policy scope for all of the network devices under his/her administrative control. Alternatively, the network administrator is free to define multiple policy scopes—each of which contains one or more network devices. Either way, EasyQoS will deploy the appropriate QoS policy

to each device, based upon the network topology, and the role of the device within the network. Up to 2,000 devices can be configured in a single policy scope as of APIC-EM version 1.3.

In order to create a new policy scope, the network administrator can click on the + next to Create New Scopes on the upper left side of the Policy Scopes tab. An empty box representing the new policy scope name will appear. The network operator will be prompted to give the new policy scope a name and click on the green check mark in order to create it. An example is shown in the following figure.

Figure 22 **Creating a New Policy Scope**

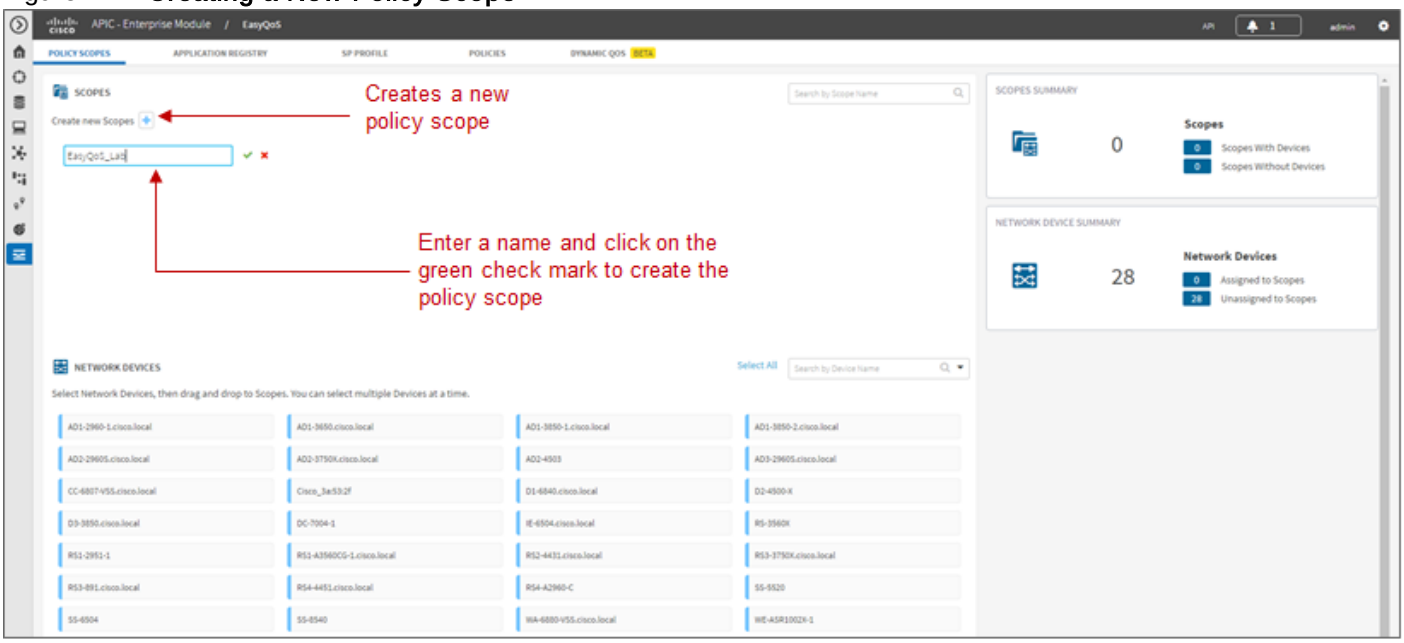

Note: The Policy Scope name cannot include any blank spaces. Use either an underscore, dash, or single word for the Policy Scope name.

When a new policy scope is created, it contains no network devices. In order to add network devices to a policy scope, the network operator must drag-and-drop one of the available devices within the Network Devices section into the policy scope. A network device can be a member of only one policy scope at a given time. Only network devices that have been Discovered and added to the Device Inventory of APIC-EM will appear within the Network Devices section, as shown in the figure above.

After a device has been dragged-and-dropped into a policy scope, it will no longer appear within the Network Devices section. Instead, it will appear within the right-hand panel within the display when the network operator clicks on the policy scope. An example is shown in the figure below.

Figure 23 **Displaying Network Devices within a Policy Scope**

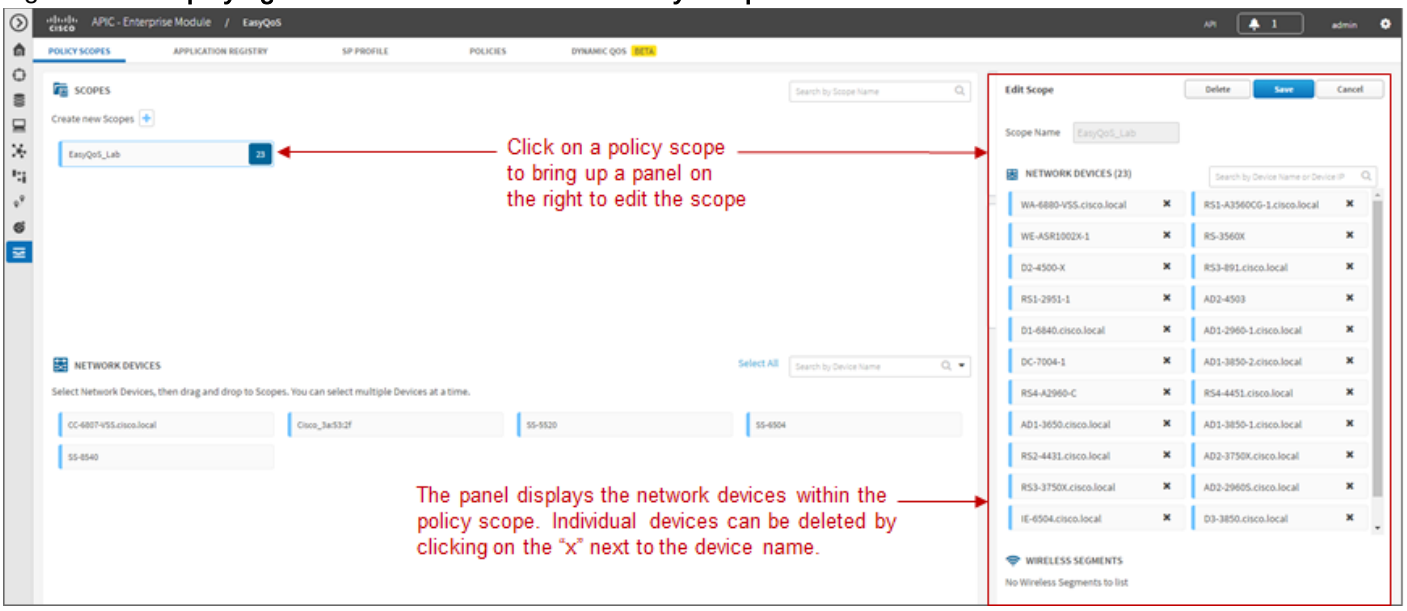

Individual network devices can be deleted from the policy scope by clicking on the "x" next to the network device name. The Save button must be clicked in order to save the changes. Devices removed from a policy scope automatically appear again within the Network Devices section underneath the Scopes section. The network operator can also delete the entire policy scope by clicking on the Delete button.

After the desired number of policy scope(s) are created and the desired network devices have been moved into the policy scope(s), the network operator can click on the Application Registry tab.

Application Registry

The second step in deploying QoS policy through EasyQoS is to access the Application Registry in order to select Favorite applications and to create Custom applications. The Application Registry serves as a common repository of applications known to APIC-EM via the NBAR taxonomy, Favorite applications, and Custom applications. It can be leveraged by various APIC-EM applications such as EasyQoS and IWAN. An example of the Application Registry is shown in the figure below:

Figure 24 **Application Registry**

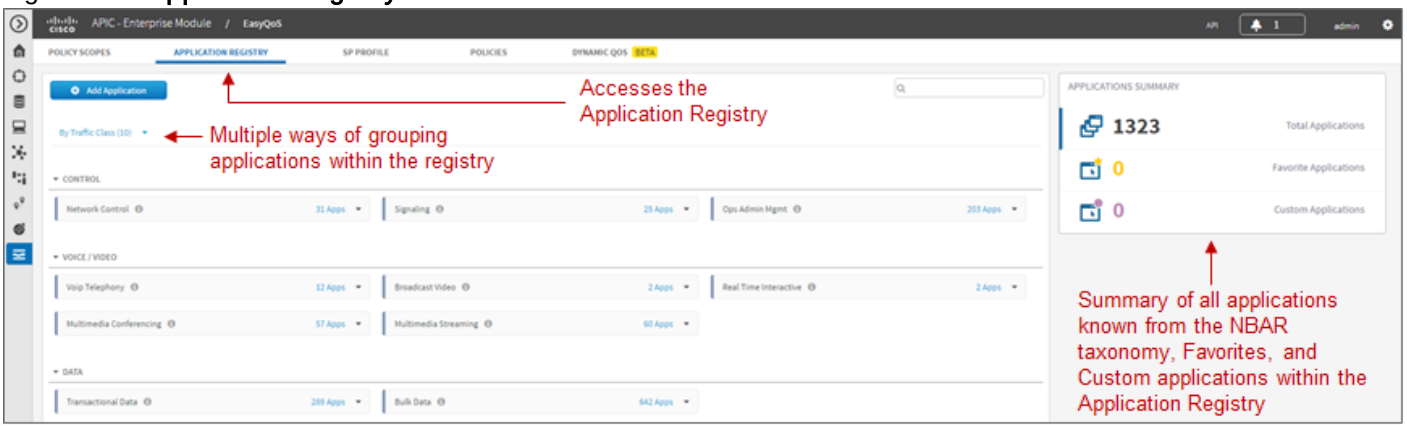

Applications can be grouped in multiple ways when viewing them within the left-panel of the Application Registry as follows:

- Applications–This lists all applications (both from the NBAR taxonomy and Custom applications) alphabetically.
- Application Groups–This lists all applications based on the NBAR application category attribute to which the applications belong.
- Traffic Class–This lists all applications based on the NBAR traffic-class attribute to which the applications belong.

The panel on the right provides a summary of all the 1300+ applications known via the NBAR taxonomy, all Favorite applications, and all Custom applications currently known and/or configured within APIC-EM.

Favorite Applications

The concept of Favorite Applications has been added to EasyQoS to address the issue that some platforms have limited ability to support applications. For instance, AireOS WLCs currently can support only 32 applications per AVC profile. Likewise, some older Catalyst switch platforms have limited TCAM space, and hence can only support a limited number of ACE entries within the ingress classification & marking ACLs deployed to these devices by EasyQoS.

By selecting an application as a Favorite, the network operator declares a preference for including that application within QoS policies provisioned by EasyQoS, over other applications. When EasyQoS creates QoS policies, it will select applications that have been marked as Favorites for inclusion within the policies before the remainder of the applications within the NBAR taxonomy. Note that by default Custom applications are automatically marked as Favorite applications when they are created.

Applications are selected as Favorites by clicking on the star next to the name of the application. Clicking on the star caused it to turn yellow, indicating application has been selected as a favorite. An example is shown in the figure below.

Figure 25 **Selecting Favorite Applications**

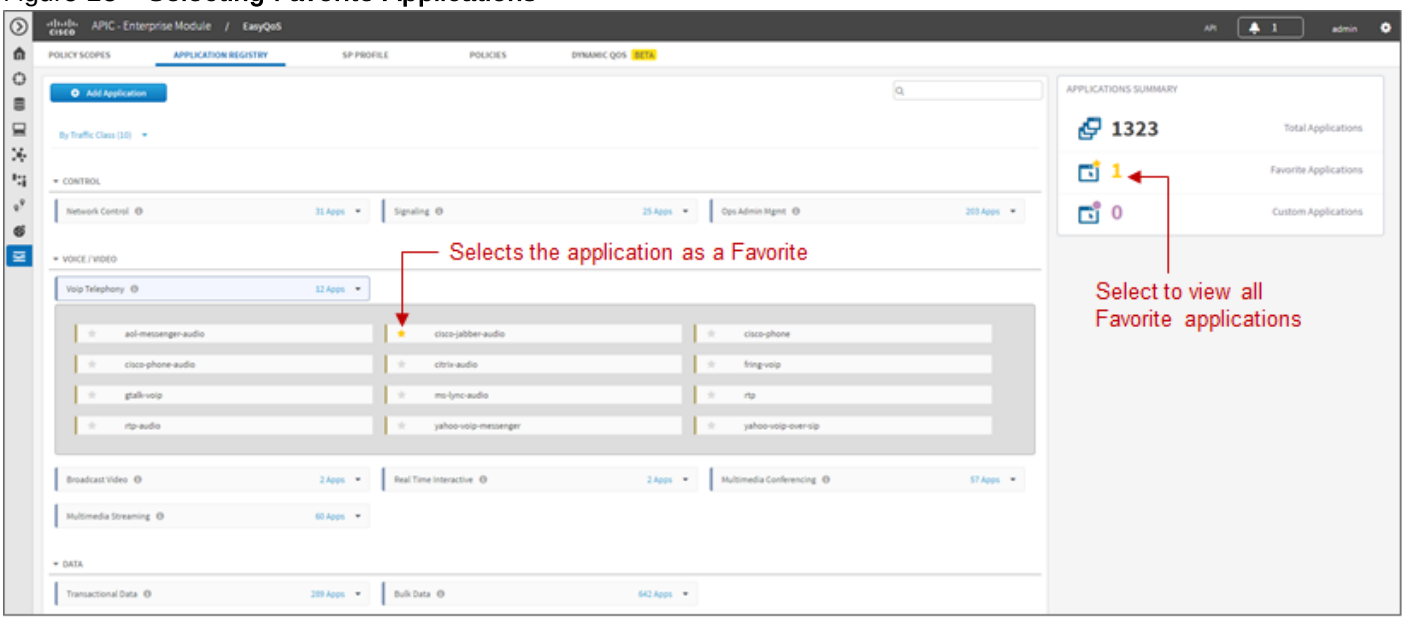

The list of Favorite applications is global to the APIC-EM EasyQoS deployment—meaning that Favorites are the same across all Policy Scopes. The list of Favorite applications can be displayed by clicking on Favorite Applications in the panel on the right side of the Application Registry.

Custom Applications

The Application Registry is also where the network operator can create Custom applications. Although AVC/NBAR currently identifies approximately 1300+ applications, organizations sometimes develop their own internal applications, which may not be recognized by AVC/NBAR. In order to identify and provide the proper QoS treatment for these applications across the network infrastructure, the network operator can create a Custom application for each of them.

Custom applications are added by clicking on the Add Application button within the Application Registry page. The right-hand panel of the page will change, allowing the network operator to add the application based upon a URL or a Server IP/Port range. An example is shown in the following figure.

Figure 26 **Custom Application Based on a URL**

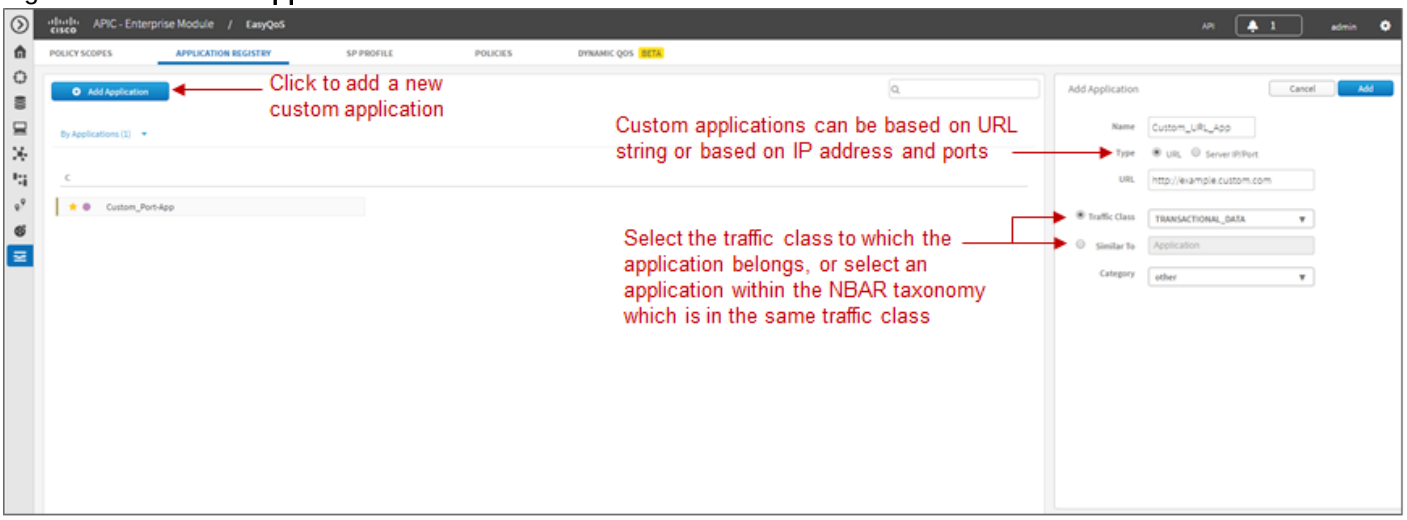

In the example above, the Custom application is based on a URL. For an EasyQoS Custom application based on a URL, the network operator must provide the following information:

- A name by which APIC-EM will know the application
- The URL string to which the application is matched against within AVC/NBAR policies
- The traffic class to which the application belongs

If the network operator does not know to which traffic class the Custom application should belong, he/she can simply select the Similar To check box and use the drop-down menu to select one of the applications known to EasyQoS via the NBAR taxonomy that has similar characteristics. By *similar characteristics*, we mean that the NBAR traffic-class and category attributes assigned to that similar application will also be assigned to the Custom application.

Custom applications that are based on URLs are not capable of being deployed on Catalyst switch platforms. They are only deployed onto ISR and ASR platforms that implement policy-maps that contain “match protocol attribute” statements. This is because the traffic-class attribute must be programmed into the Custom application, and the traffic-class attribute requires a “match protocol attribute traffic-class”

statement to be configured within the policy-map. An example of the policy configuration for a Custom application that is based on a URL is shown in the *Custom Applications on ASR and ISR Platforms* section of this document.

Alternatively, a Custom application can be based upon one or more IP addresses (or address ranges) and one or more IP, TCP, and/or UDP ports (or port ranges). An example using multiple IP addresses and port ranges shown in the following figure.

Figure 27 **Custom Application Based on Server IP Address and Ports**

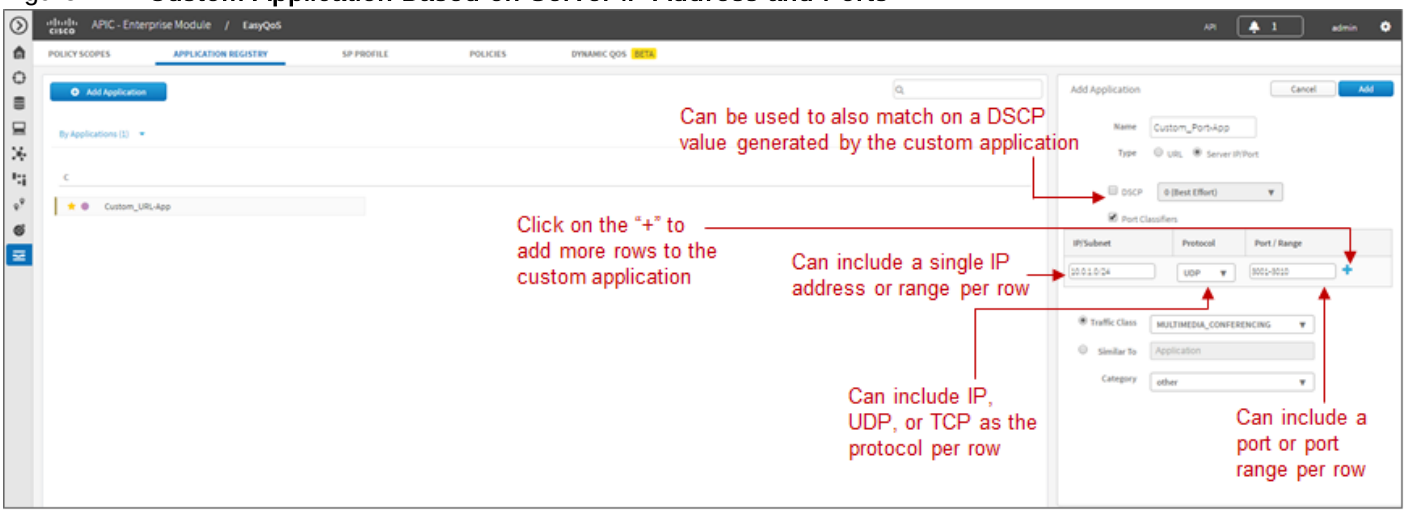

For an EasyQoS Custom application based on a server IP addresses and ports, the network operator needs to provide the following information:

- A name by which APIC-EM will know the application
- A DSCP value (optional). This field is used to match on a DSCP value generated by the Custom application within the QoS policy generated by EasyQoS.
- Port Classifiers (optional), which include one or more IP addresses or IP address ranges, along with one or more protocols (IP, TCP, or UDP), and one or more ports or port ranges.
- The traffic class to which the application belongs

The example in the figure above demonstrates the use of the subnet mask field—set to 24 bits (10.0.1.0/24) in the first row of the Port Classifiers—to include a full subnet as a destination IP address range. Likewise, the range (3000-3010) in the first row of the Port Classifiers shows how to include a range of ports (UDP ports in this example—based on the Protocol setting for the particular row). This IP address or address range as well as port or port range refers to a destination—also referred to as the producer. Additional rows can be added to include more individual IP addresses or IP address ranges, as well as more ports or port ranges to the Custom application.

As with URL-based Custom applications, if the network operator does not know to which traffic class the Custom application should belong, he/she can simply select the Similar To check box and use the drop-down menu to select one of the applications known to EasyQoS via the NBAR taxonomy that has similar characteristics. By *similar characteristics*, we mean that the NBAR traffic-class and category attributes assigned to that similar application will also be assigned to the Custom application.

Custom applications that are based on server (destination) IP addresses and ports are capable of being deployed on both Catalyst switch platforms and ASR or ISR platforms. For Catalyst switch platforms, the server (destination) IP address or address range, ports, and/or DSCP fields are translated into one or more ACEs that are populated within the ACL corresponding to the traffic-class to which the Custom application belongs. An example of this is shown in the *Custom Applications Provisioned within ACLs* section of this document.

For more complex applications, a source IP address or address range as well as a source port or port range can be added to the Custom Application. This is referred to as *adding a Consumer to the application*. Adding a Consumer is discussed in the *Policies* section below.

SP Profile

The configuration of custom SPPs is an optional step that is dependent upon the following two questions:

- Is a managed-service WAN implemented on any interface of any ISR or ASR router within the scope of the policy to be deployed?
- If there is a managed-service WAN, does the service match one of the four default SP profiles provided by EasyQoS?

The four default SP profiles provided by EasyQoS can be viewed by clicking on the SP Profile tab and then highlighting one of the four default profiles as shown in the following figure:

Figure 28 **Default SP Profiles Provided by EasyQoS**

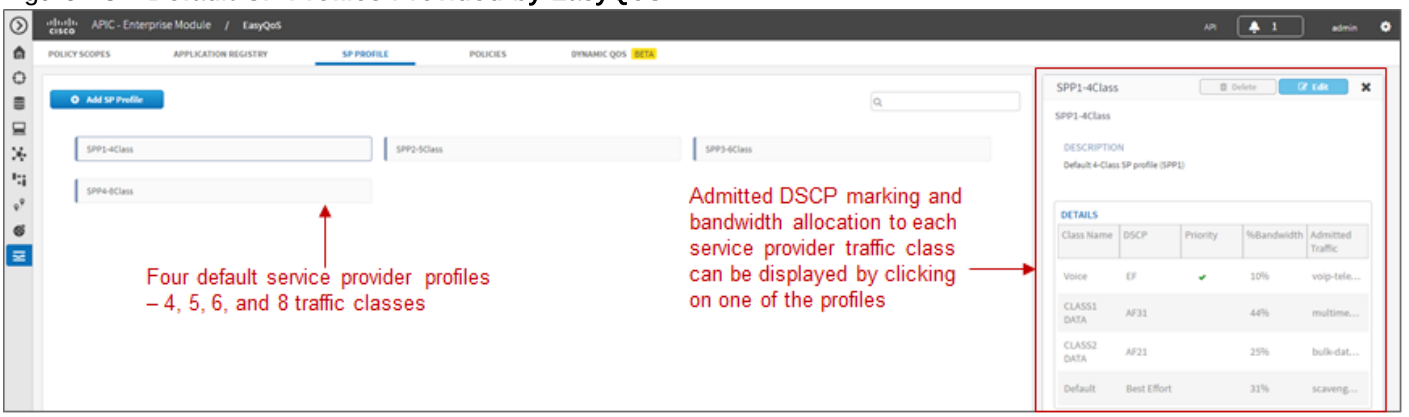

The network operator can view the bandwidth allocations and the admitted DSCP markings to each of the service provider traffic classes for each of the profiles by simply clicking on one of the four default SP profiles.

Note: If the QoS policy has previously been deployed, and if the selected default SP Profile has been deployed to devices within the policy scope, the devices and interfaces will be displayed within a panel at the bottom of the screen.

If the network operator determines that none of the four default SP profiles matches a managed-service WAN deployed on an interface on any of the devices within the policy scope, he/she can create a custom SP profile by clicking on the Add SP Profile button in the upper left corner of the page. This will bring up a page similar to the following:

Figure 29 **Creating a Custom SP Profile**

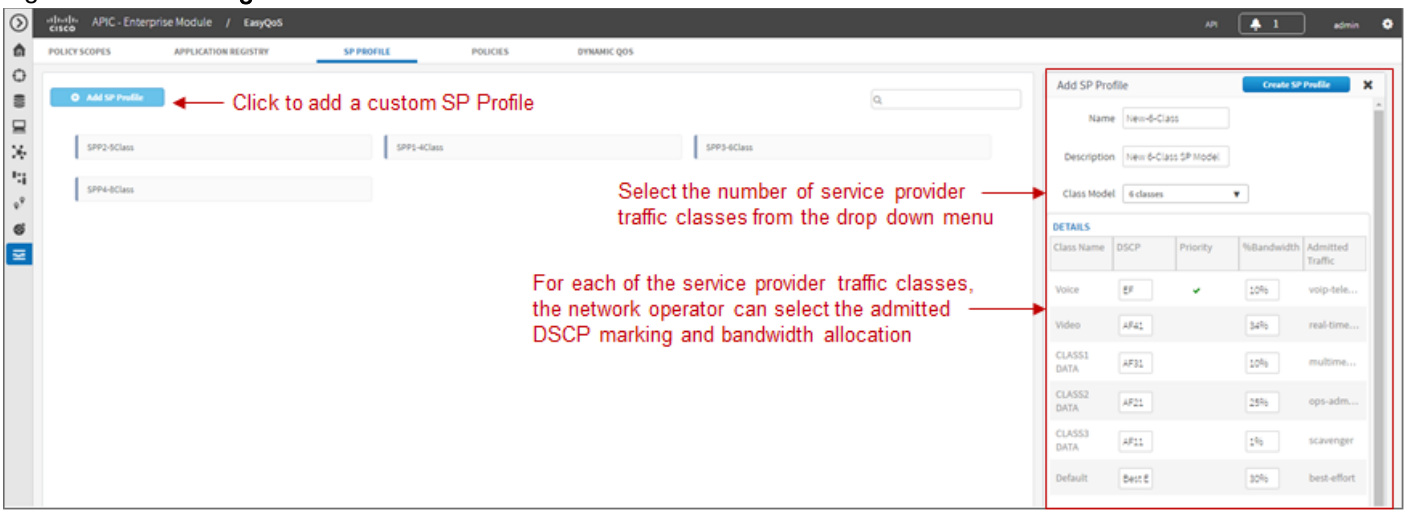

Custom service provider profiles are based on the four default SP profile templates–meaning that custom SP profiles can only have 4, 5, 6, or 8 traffic classes. The network operator must first select the number of classes in the custom SP profile through the drop down menu next to Class Model. This will change the Details panel at the bottom left of the page to reflect the number of traffic classes in the model. The network operator can change the admitted DSCP marking to the service provider traffic class or accept the default value. Likewise, the network operator can change the amount of bandwidth allocation to the particular service provider traffic class, or accept the default value.

The Voice traffic class is the only traffic class that is mapped to a low-latency queue (LLQ), otherwise known as a *priority queue*. Because the Voice traffic class is a priority queue, the remaining bandwidth allocations are technically bandwidth remaining allocations–which must total to 100%–regardless of the amount of bandwidth allocation provisioned to the Voice traffic class.

When the network operator is satisfied with the admitted DSCP markings and bandwidth allocations for the service provider traffic classes, he/she can click on the Create SP Profile button in the upper right corner of the page in order to create the custom SP Profile.

The application of service provider profiles–regardless of whether they are one of the four default SP profiles, or a custom SP profile–to WAN interfaces is done automatically by APIC-EM. It is based on the network operator having previously configured a specific tag within the description of the WAN interface connected to a managed service WAN. This is discussed in detail within the *Service Provider Managed-Service WAN QoS Design* section of this document.

Policies

The next step to deploying QoS policy through EasyQoS is to access the Policies tab in order to create a policy under a scope. An example is shown in the figure below.

Figure 30 **Creating a Policy Within a Policy Scope**

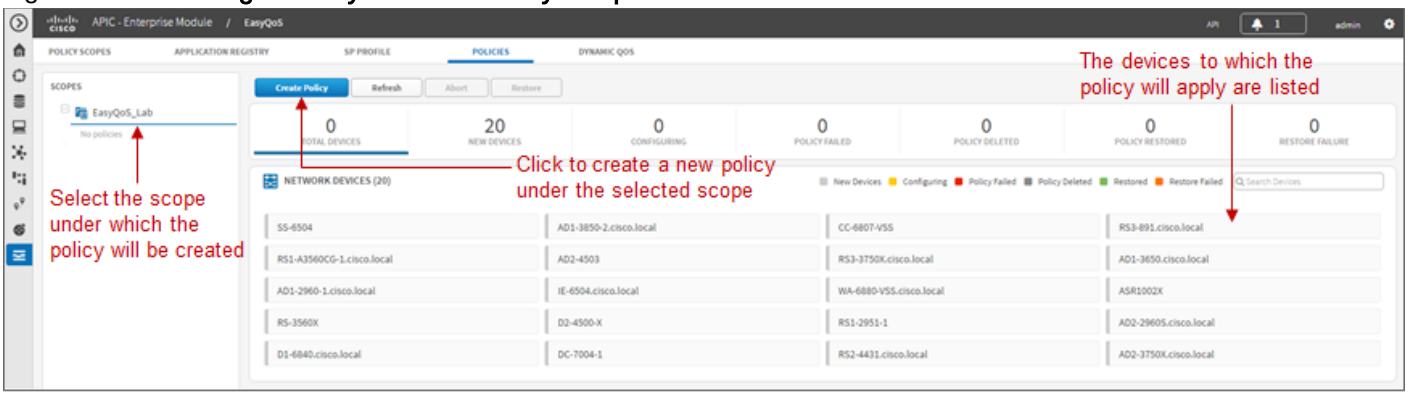

Policies are created and applied per policy scope. Therefore policies only affect those devices that are part of the particular scope to which the policy is applied.

Clicking on the Create Policy button brings up the screen used to create, configure, view, and apply the policy. An example is shown in the figure below.

Figure 31 **Configuring EasyQoS Policy**

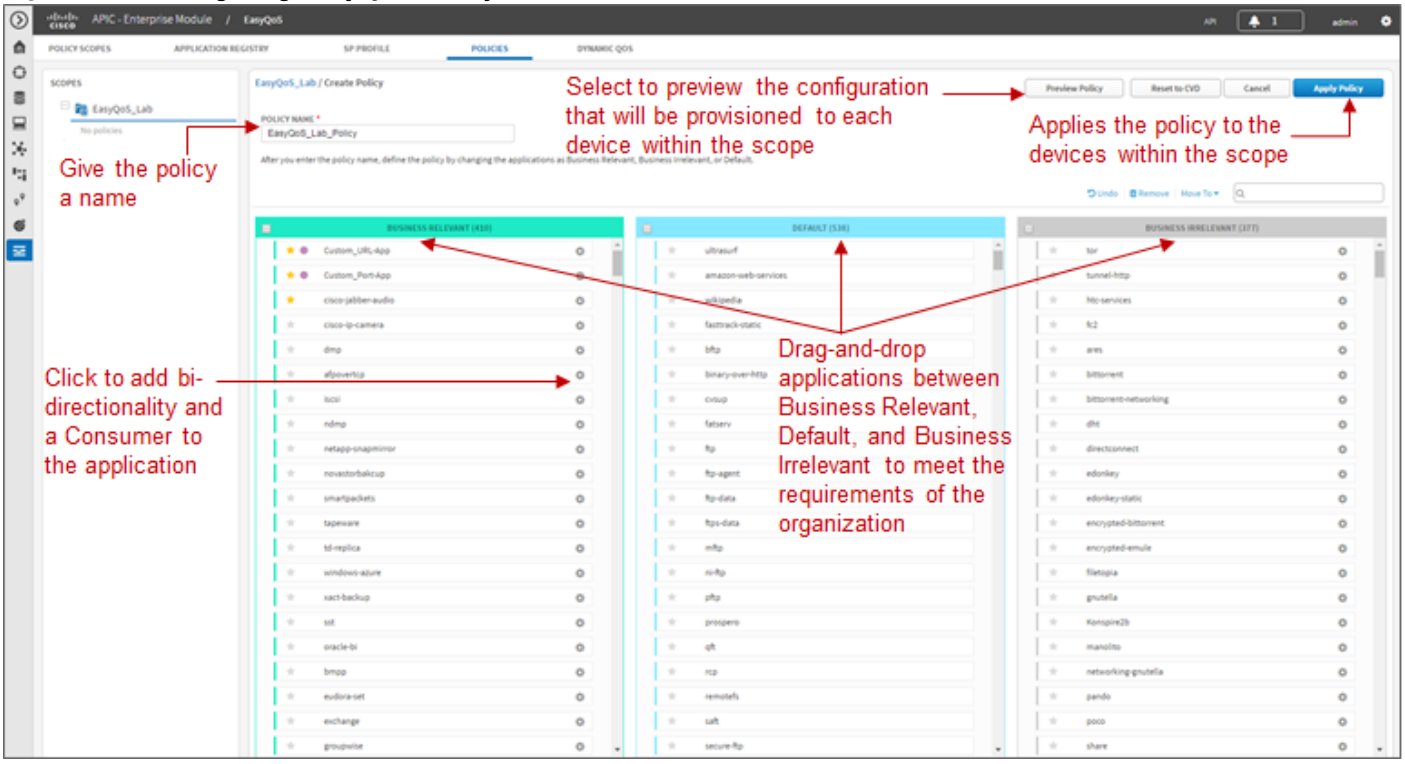

It is mandatory to name each QoS policy. The purpose of the policy is as follows:

- To capture the application-level business intent of the network operator
- To transform the business intent into network device QoS configuration for each device within the policy scope
- To apply the configuration to the devices within the policy scope

- Finally, to inform the network operator of the status of applying the policy to each device within the policy scope

The application-level business intent of the network operator is captured by dragging-and-dropping individual applications known via the Application Registry between the three business-relevance attribute values of Business Relevant, Default, and Business Irrelevant. All 1300+ applications known within the NBAR taxonomy have default settings for the business-relevance attribute. The network operator can simply choose to accept these default values, or customize as many applications as needed to meet the business requirements of the organization.

The network operator must also select the business-relevance attribute of Custom applications–either Business Relevant, Default, or Business Irrelevant, by dragging-and-dropping Custom applications into the appropriate grouping. By default, custom applications are Unassigned when they are created.

Changing the business-relevance of an application changes its QoS treatment across the network, as discussed in the *Mapping Business-Relevance to QoS Treatments* section of this document.

Note: If a policy is applied before changing the business-relevance of the Custom application, the Custom application will not be included within the policy. However, since applications with a business-relevance of Default do not have any actual configuration generated on network devices for those applications, the Custom applications will in effect be treated with a business-relevance of Default across the network infra-structure.

The application-level business intent is then transformed by APIC-EM EasyQoS into QoS configuration for each network device within the policy scope covered by the policy. The configurations are based upon best practice recommendations for QoS configuration, compiled through years of CVD guidance.

Consumers, Producers, and Bi-Directionality

The Policies page is where the network operator can choose to make the application bi-directional. As well as add a Consumer (a source IP address or IP address range, and/or source port or port range). This is done by clicking on the icon next to the name of the application, which brings up a pop-up screen to edit the details of the application. An example using the Custom application discussed earlier is shown below.

Figure 32 **Adding Bi-Directionality and a Consumer to an Application**

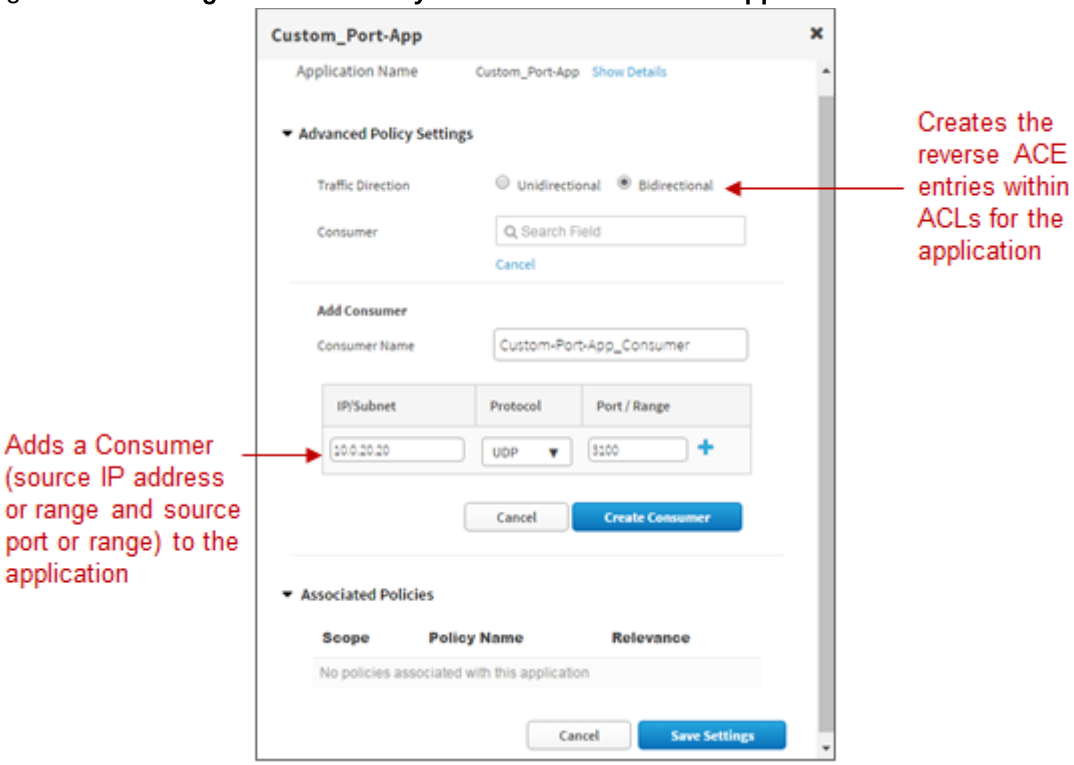

All Ingress classification & marking policies implemented on Catalyst switches use ACE entries within ACLs. Ingress classification & marking policies for Custom applications implemented on ISR and ASR platforms also use ACE entries within ACLs. These ACE entries are, by default, unidirectional. The bi-directional feature is intended to ensure that return traffic from an application is classified and marked correctly when the destination (the Producer) is not within a data center or on a server where the switch port to which the server is connected can be configured to trust the DSCP markings of traffic from the server.

Note: On ISR and ASR platforms, ingress classification and marking policies involving any of the 1300+ applications known to the NBAR taxonomy are handled by the AVC/NBAR engine and are bi-directional.

As mentioned in the previous paragraph, ingress classification and marking policies implemented on Catalyst switches, as well as custom applications implemented on ISR and ASR platforms use ACE entries within ACLs. These ACE entries by default only specify a destination IP address or range of IP addresses, as well as a destination port or port range. The source is by default “any” device. In order to specify a source IP address or IP address range as well as a source port or port range, a Consumer is created and added to the application. This can be for a Custom application or for any of the 1300+ applications know to the NBAR taxonomy. The network operator accomplishes this by simply giving the Consumer a name, specifying a source IP address or IP address range, specifying whether the transport layer protocol is UDP or TCP, specifying a port or port range, and clicking the Create Consumer button. This, along with the choice for bi-directionality must be saved before closing the Edit Application Details pop-up screen.

The effects on ASR or ISR router configuration and on Catalyst switch configuration of adding bi-directionality and a Consumer are discussed in the *Custom Applications on ASR and ISR Platforms* and *Custom Applications Provisioned within ACLs* sections of this document, respectively.

Reset to CVD

As applications are dragged-and-dropped between the Business Relevant, Default, and Business Irrelevant groupings within a given policy, the network operator may lose track of their original default settings. Likewise as bi-directionality and consumers are added to individual applications within a policy, the network operator may lose track of which applications have been set for bi-directionality and/or have consumers added. The network operator has the ability to reset the applications back to their original business-relevance attribute setting, and to remove bi-directionality and consumers within a given policy, by clicking on the Reset to CVD button. Note that the selection of Favorites is system-wide (that is, across policies and policy scopes) and therefore unaffected by the Reset to CVD button.

Policy Preview

Before applying the configuration, the network operator can optionally choose to preview the policy. This option is enabled by selecting the Preview Policy button within a policy. When the Policy Preview is selected, an additional panel will appear as shown in the following figure.

Figure 33 **Generating a Policy Preview for a Device**

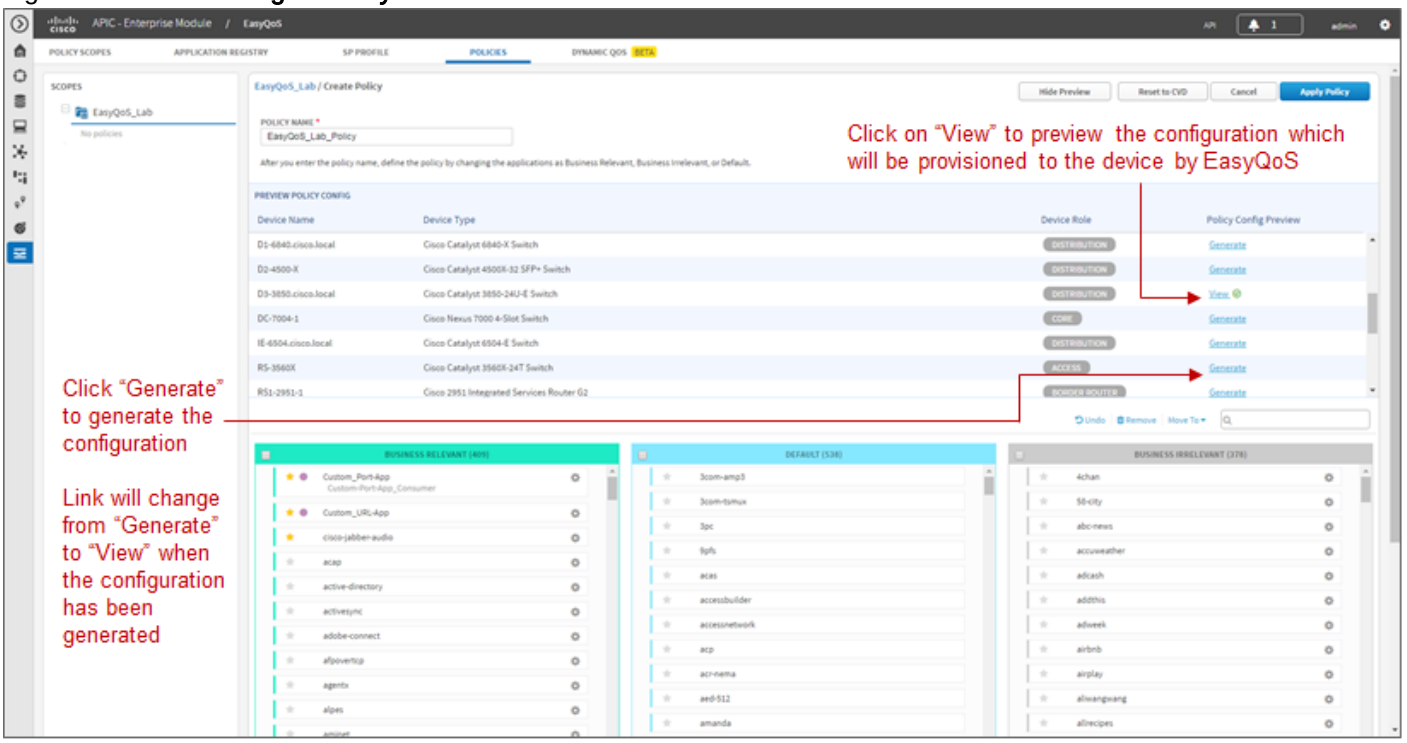

The Preview Policy Configuration panel allows the network operator to generate the actual commands that will be provisioned to each device by EasyQoS. This is done by clicking on the Generate link adjacent to the specific device. The Generate link will change to View when the configuration has been generated. Clicking on the View link will bring up a pop-up window in which the configuration commands will appear. The configuration commands can then be viewed by scrolling up and down within the panel. An example is shown in the figure below.

Figure 34 **Displaying the Preview Configuration**

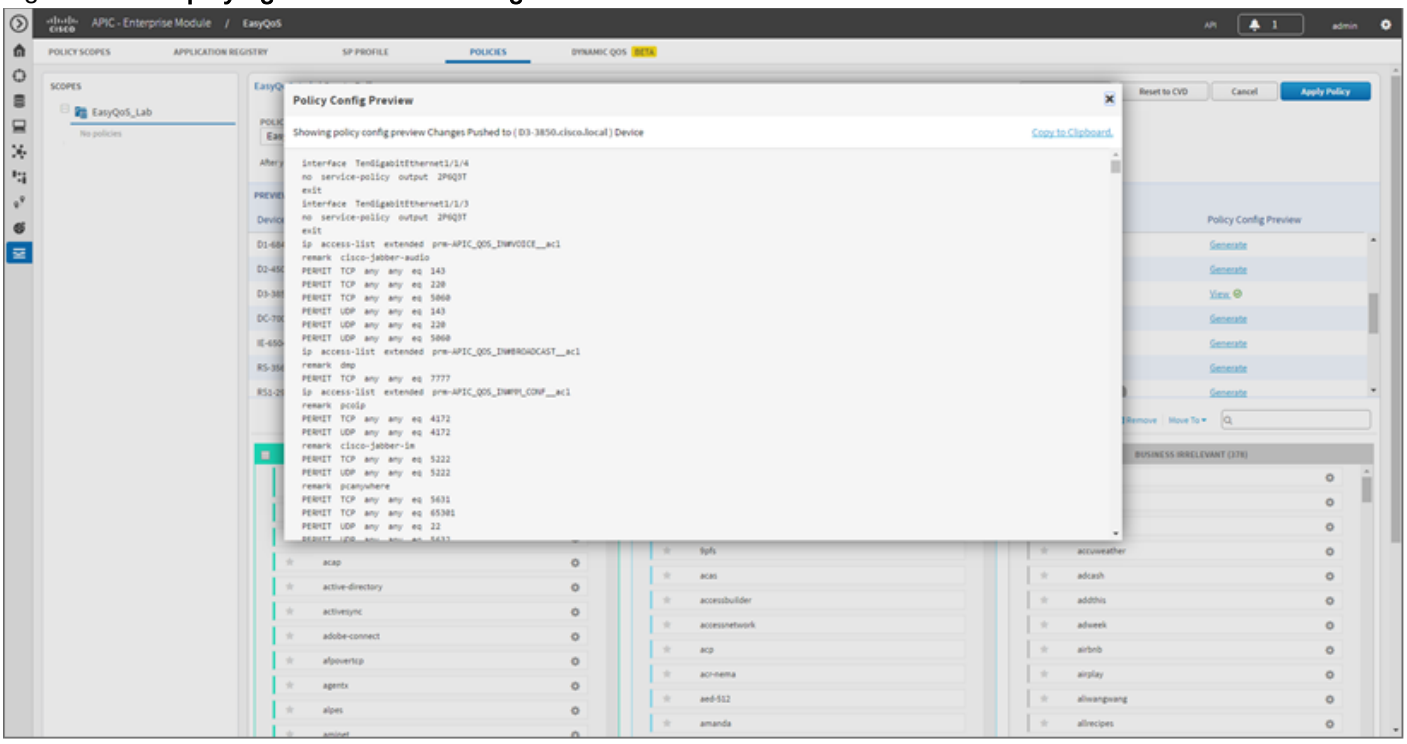

The preview policy option can be useful in uncovering potential errors in applying policy–such as an unsupported line card within a Catalyst 6500 Series switch for instance–before the policy is applied. The network operator can then take remedial actions, such as removing the device from the policy scope, or removing the line card from the switch–before applying the policy. Because the actual configurations that are provisioned by EasyQoS to each device are generated, this may also improve the time taken to deploy the policy to all of the devices within the scope, as well.

Policy Status, Abort, History, and Rollback

When the network operator is satisfied with the policy, he/she can click on the Apply Policy button in the upper right corner policy screen to apply the policy to the devices within the scope. EasyQoS provides the status of the policy on each device–as the policy is being applied. Initially each device will appear with a gray bar next to it–indicating that no policy is applied (if this device is new to EasyQoS with no policy). A yellow bar next to a device indicates that policy is currently being configured onto the device. Finally, a green bar next to the device indicates that the policy has successfully been provisioned onto the device. An example of the policy being applied to devices within the scope is shown in the following figure.

Figure 35 **Policy Status**

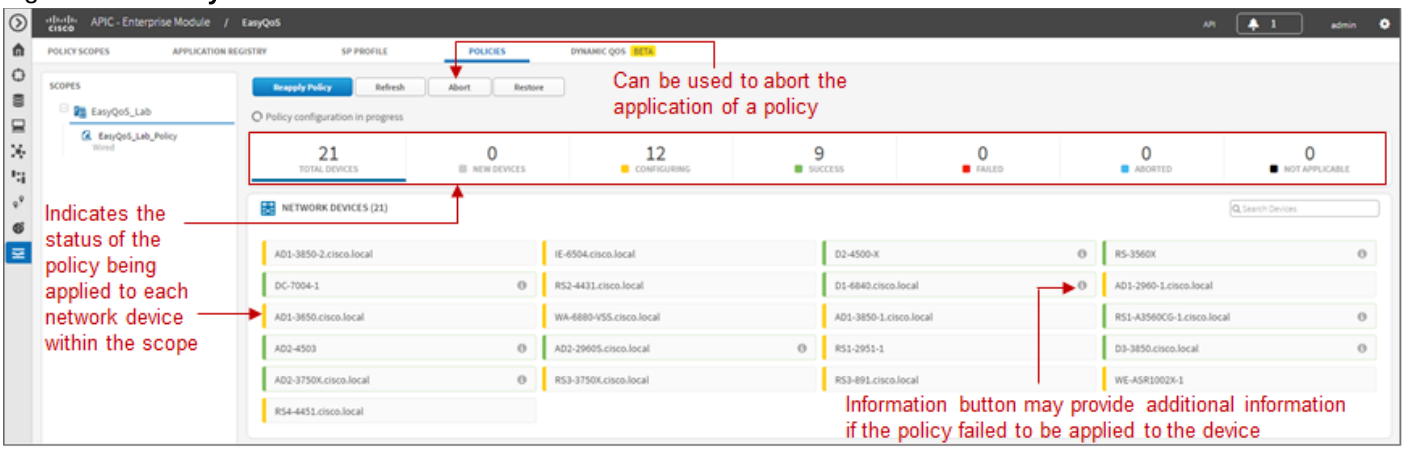

The network operator can abort the provisioning of the policy to network devices after the policy provisioning has begun, but before the policy provisioning process has completed, by clicking on the Abort button. EasyQoS provisions multiple (up to 40) devices at a time. Hence, the abort option is only useful when there are a large number of devices (more than 40) within a policy. Rather than waiting for the entire policy to be provisioned to each device, and then either rolling back the policy or restoring the configuration, the network operator can instead terminate the provisioning of the policy with the Abort button. For policies with a small number of devices, it may be more effective to allow the policy to complete and then either Rollback the policy or Restore the devices to their configuration before EasyQoS policy was applied.

When the Abort button is pressed, EasyQoS cancels the provisioning process only on network devices that have not yet been started to be configured. A light blue bar next to these devices will indicate a status of Policy Aborted for these devices. For devices that have started to be configured, EasyQoS will complete the provisioning of the policy. For devices for which the provisioning of the policy has been completed before the Abort button was pressed, EasyQoS will leave the policy on the device and will update the status of each of these devices–a green bar for Success or a red bar for Failed–based on the outcome of the provisioning of the policy to the device. The network operator can then either Rollback the policy or Restore these devices to their original configuration before EasyQoS policy was applied.

After a policy has been applied, clicking on the name of the policy within the left-hand panel displays the policy, as shown in the following figure.

Figure 36 **Displaying a Policy**

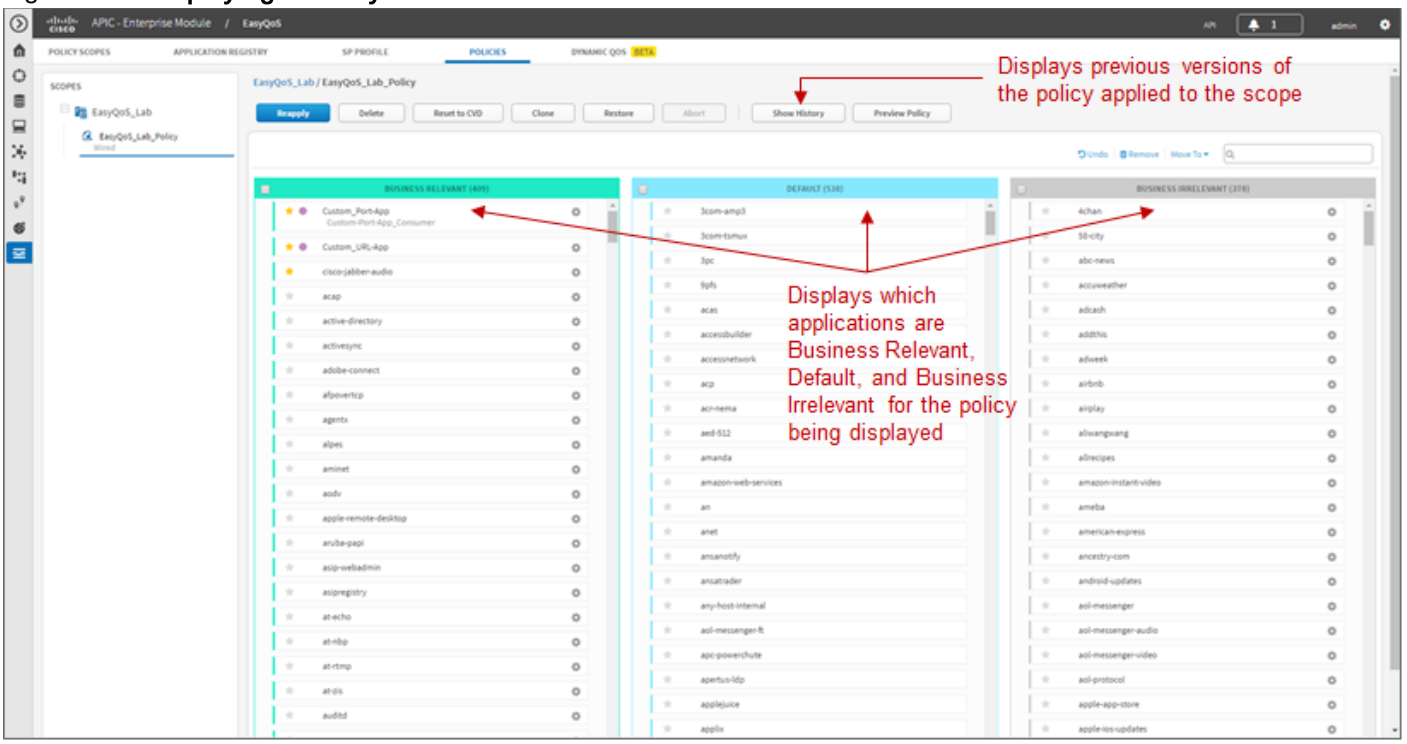

The network operator's business intent–in terms of the business-relevance of applications–is applied per policy. This means that applications can be assigned different business-relevance attributes in different policies. Here the network operator can view which applications are Business Relevant, Default, and Business Irrelevant for the policy being displayed.

Clicking on the Show History button opens a new Policy History panel in the center of the page, as shown in the following figure.

Figure 37 **Show History**

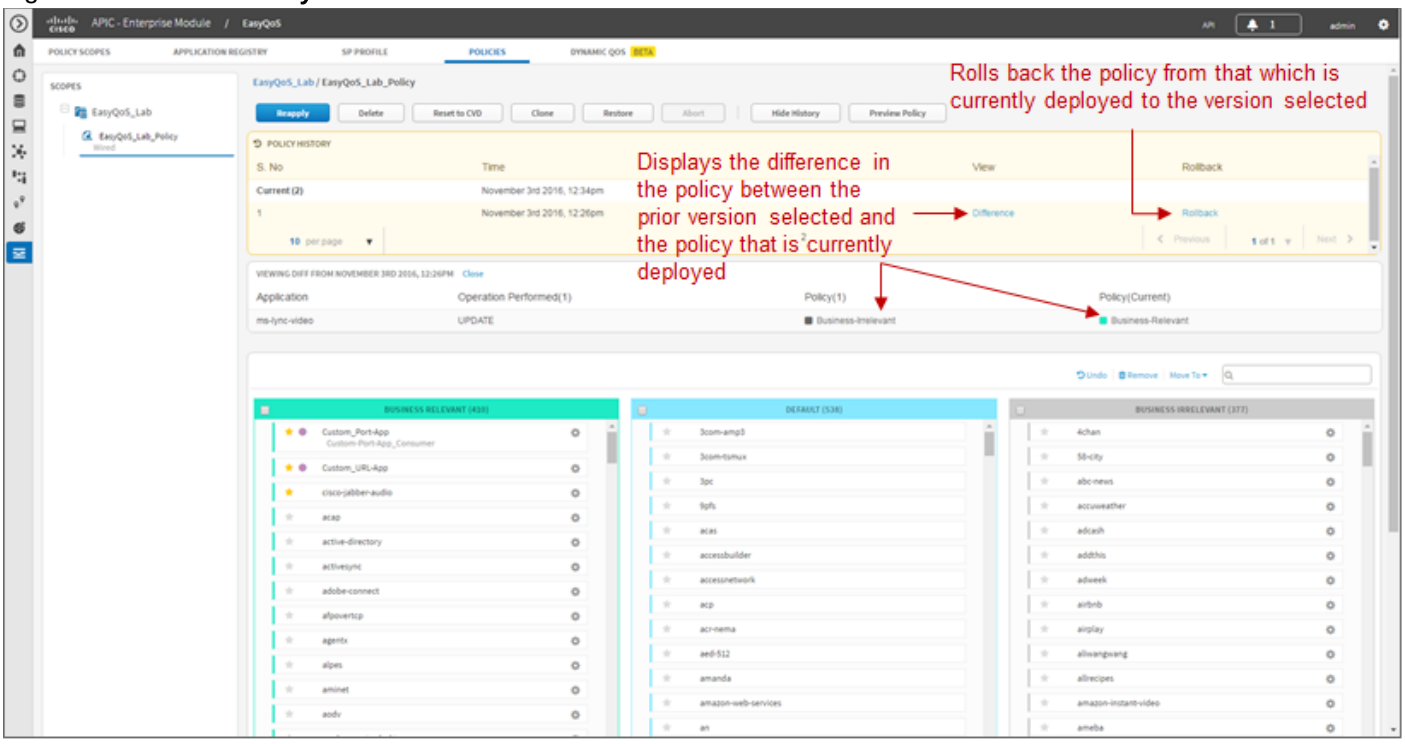

The Policy History panel displays previous versions of the policy selected. The network operator can view the changes in the policy that have been made in the various versions by selecting the Difference feature under a particular prior version. This will display the difference in the policy between the prior version selected and the policy that is currently deployed (it does not display the difference in policy between the prior version selected and the next lower prior version).

The difference in policy is represented in terms of applications–meaning certain applications may have been moved between business relevance, certain applications may have been added or deleted from the Favorites, or custom applications may have been added or deleted. The Difference feature does not display the difference in the actual configuration applied to each network device.

The Rollback feature under a particular prior version can be used to roll back the policy from that which is currently deployed to the particular prior version selected. This feature is useful in change management scenarios, where a particular change is found to be undesirable and the network infrastructure needs to be rolled-back to the state it was in prior to the change being implemented.

Finally, the Clone button can be used to copy the entire policy. Upon clicking the Clone button, the network operator will be asked to enter a new policy name for the cloned policy and to select a policy scope to which the new cloned policy will be applied. With complex policies the network operator can save administrative time by not having to duplicate the same policy across multiple policy scopes. After the policy is cloned, the network operator is free to modify it as needed for the particular policy scope.

Restore

The Restore button deletes an EasyQoS policy and attempts to restore the QoS configuration on all devices covered by the EasyQoS policy back to the original configurations before any EasyQoS policy was applied. Because the EasyQoS policy is deleted when the network operator selects the Restore button, there is no

ability to retry the Restore function if it does not succeed in restoring the configuration of all devices to their original (pre-EasyQoS) configuration in APIC-EM release 1.3. This behavior is similar to when the network operator selects the Delete button, in that the EasyQoS policy is deleted. There is no ability to retry the delete function, either.

Note: If the first attempt to provision an EasyQoS policy to a device (meaning the device initially has no EasyQoS policy) fails, EasyQoS will also automatically attempt to restore the QoS configuration on that device to its original (pre-EasyQoS) configuration.

The network operator should be aware that if a device is removed from an EasyQoS policy after the policy has been applied to the device, the EasyQoS policy will remain on that device. In other words, EasyQoS will not automatically attempt to delete the QoS policy provisioned to the device, nor will EasyQoS attempt to restore the QoS configuration on that device to the original (pre-EasyQoS) configurations.

The Restore button will also not restore the QoS configuration on that device to the original configuration if the original configuration was already an EasyQoS configuration. This situation may occur when upgrading from APIC-EM 1.2 to APIC-EM 1.3. APIC-EM did not collect the information required to restore the original configuration before provisioning policy in APIC-EM version 1.2.

Finally, there are some parts of the EasyQoS policy that may not be restored, depending upon particular network device platform. The *Pre-Existing QoS Configuration on ISR and ASR Router Platforms* section of this document details what is restored and not restored on router platforms when clicking on the Restore button. Likewise, the *Pre-Existing QoS Configuration on Switch Platforms* section of this document details what is restored and not restored on switch platforms when clicking on the Restore button.

Wireless Policies

Cisco AireOS WLCs can also be added to policy scopes by dragging-and-dropping the device into a particular policy scope. Optionally, a separate policy scope can be created for wireless devices. Wireless policies are deployed per WLAN/SSID. If there are multiple WLANs/SSIDs to which EasyQoS policies need to be applied, then the network operator must create a policy for each WLAN/SSID.

A wireless policy (separate from the policy applied for wired devices) must be created under the policy scope. This is done by clicking on the blue circular + icon adjacent to the name of the WLAN/SSID within the Wireless Segments section of the page. As with wired policies, the wireless policy must be given a name; individual applications can be moved between Business-Relevant, Default, and Business-Irrelevant groupings; bi-directionality can be selected for individual applications; and the policy can be previewed before being deployed.

An example of a wireless policy created within an existing policy scope is shown in the figure below.

Figure 38 **Adding a Wireless Policy to a Policy Scope**

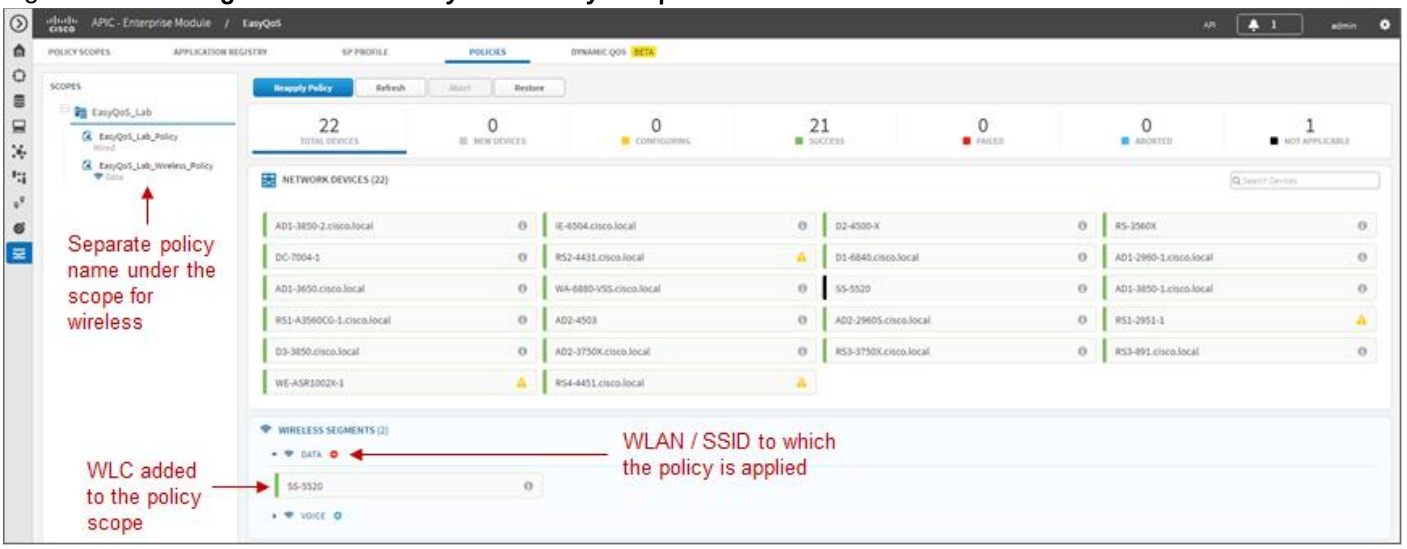

Clicking on the wireless policy displays the policy, as shown in the figure below.

Figure 39 **Displaying a Wireless Policy**

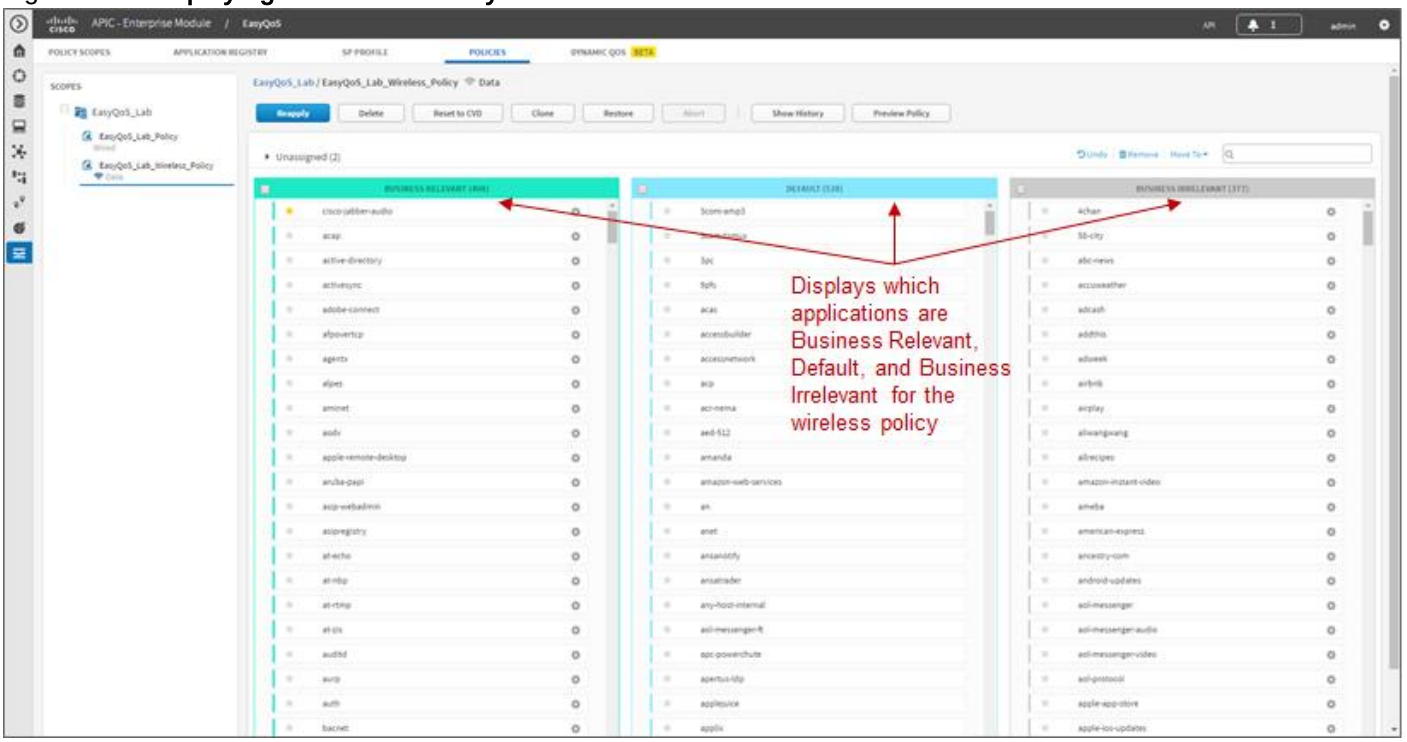

Applications known via the AVC/NBAR engine within the AireOS WLC are displayed in groupings of Business Relevant, Default, and Business Irrelevant. The network operator can drag-and-drop the applications between the three groupings in order to match the business intent of the organization.

The AVC/NBAR based classification & marking policy deployed to AireOS WLCs specifies an action of mark, rather than drop. For Business Relevant applications, the DSCP marking is based on the value of the traffic-class attribute assigned within the NBAR taxonomy. For Business Irrelevant applications, the DSCP marking is set to CS1 (DSCP 8). Applications with a business-relevance attribute of Default are not programmed into

the AVC/NBAR policy. The AVC/NBAR policy overrides the QoS Profile applied to the WLAN/SSID. Because EasyQoS sets the Maximum Priority field within the QoS Profile to a setting of Voice, applications with a business-relevance attribute of Default are not reset to a DSCP value of Best Effort (DSCP 0). Instead, the WLC allows such applications to pass through with their DSCP values unaltered.

AireOS WLCs support up to 32 applications per QoS Policy. This is a current limitation of the AVC/NBAR engine within AireOS WLCs. By default, EasyQoS will select the 32 applications that get programmed into the wireless policy based upon which applications are selected as Favorites and then based upon the popularity attribute pre-configured for all of the 1300+ applications within the NBAR taxonomy.

Because the network operator has no view of the popularity attribute for any given application within the NBAR taxonomy, there are two methods by which the network operator can guarantee which 32 applications are provisioned into the AireOS AVC/NBAR-based classification & marking policy. Note that the AVC/NBAR-based policy can have less than 32 applications as well, if desired by the network operator.

- In the first method, the network operator can select up to 32 applications as Favorites. APIC-EM will provision applications marked as Favorites before provisioning other applications within the NBAR taxonomy. However, because the choice of Favorites is a global setting—meaning the selection of Favorite applications is the same across all policies in all policy scopes—this may not be an ideal solution.

- In the second method, the network operator can highlight all applications within each of the three groupings—Business Relevant, Default, and Business Irrelevant—via the checkbox at the top of each group. The network operator can then remove all of the applications. This will place all applications for the wireless policy in the Unassigned group. The network operator can then use the search field in the upper right corner of the page to search for each application he/she wants to add back into the policy. Upon locating the application, the network operator must drag-and-drop that application from the Unassigned group into either the Business Relevant or Business Irrelevant grouping. In this manner, the network operator can add up to 32 applications into the wireless policy and ensure they will be provisioned to the AireOS WLC.

As of APIC-EM version 1.3, the default behavior of the AVC/NBAR-based classification & marking policy is to mark in the upstream direction only. In order to implement bi-directional policies, the network operator must configure bi-directionality for the application. This is done the same way as discussed for Custom applications in the *Policies* section above.

Finally, the network operator should note that Custom applications—regardless of whether they are URL-based or port-based—are not provisioned into the AVC/NBAR-based classification & marking policy of WLC platforms. This is a current limitation of AireOS WLC platforms, in that they do not support the ability to define custom applications within the AVC/NBAR policy.

Dynamic QoS

For the APIC-EM 1.3 release, Dynamic QoS is still a Beta application within EasyQoS. In order to enable Dynamic QoS, the network operator must access the Dynamic QoS tab to bring up the screen shown in the figure below.

Figure 40 **Enabling Dynamic QoS**

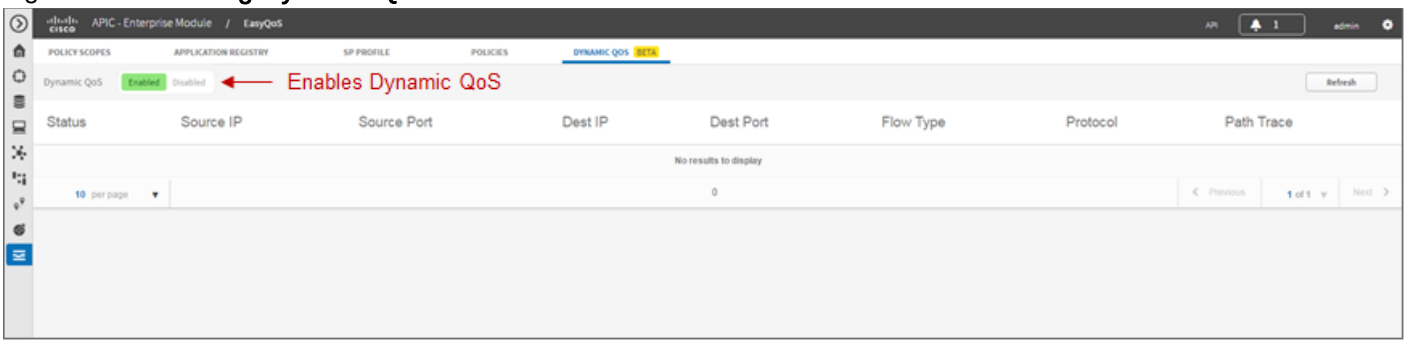

For APIC-EM 1.3 Dynamic QoS is a feature that is enabled globally–meaning across all policy scopes–through the sliding button shown in the figure above. Upon re-applying static QoS policy to a given policy scope, EasyQoS will then provision Dynamic policy-map shells to access-layer switches within that policy scope. The Dynamic policy-map shells are discussed within the *Dynamic QoS Design* section of this document.

The Dynamic QoS screen displays the status of Dynamic QoS flows when they are active. Dynamic QoS flows are initiated when call signaling systems, such as CUCM, use the northbound REST-based API to signal to APIC-EM that a call has been established. An example of an active Dynamic QoS session–representing a voice call between two endpoints–is shown in the figure below.

Figure 41 **Example of an Active Dynamic Voice Call**

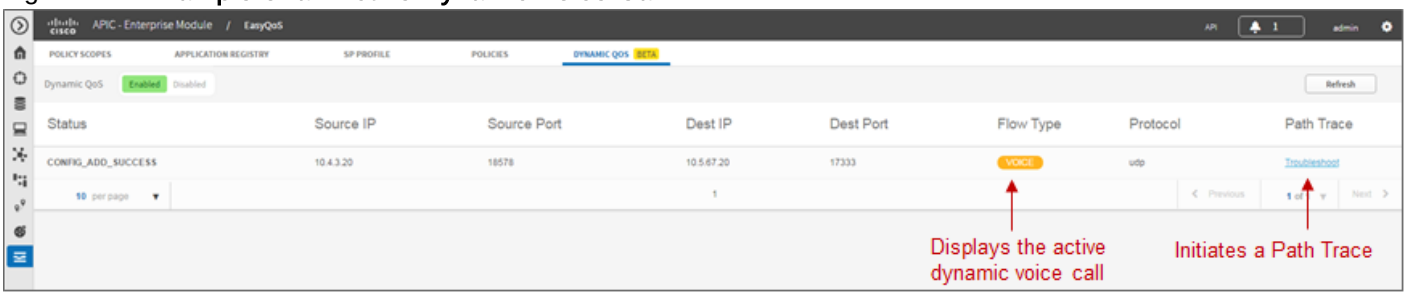

The Path Trace tool can be used to troubleshoot active Dynamic QoS flows by displaying the path of the traffic through the network infrastructure between the endpoints. An example of the Path Trace tool is shown in the figure below.

Figure 42 **Example of the Output of the Path Trace Tool for a Dynamic Flow**

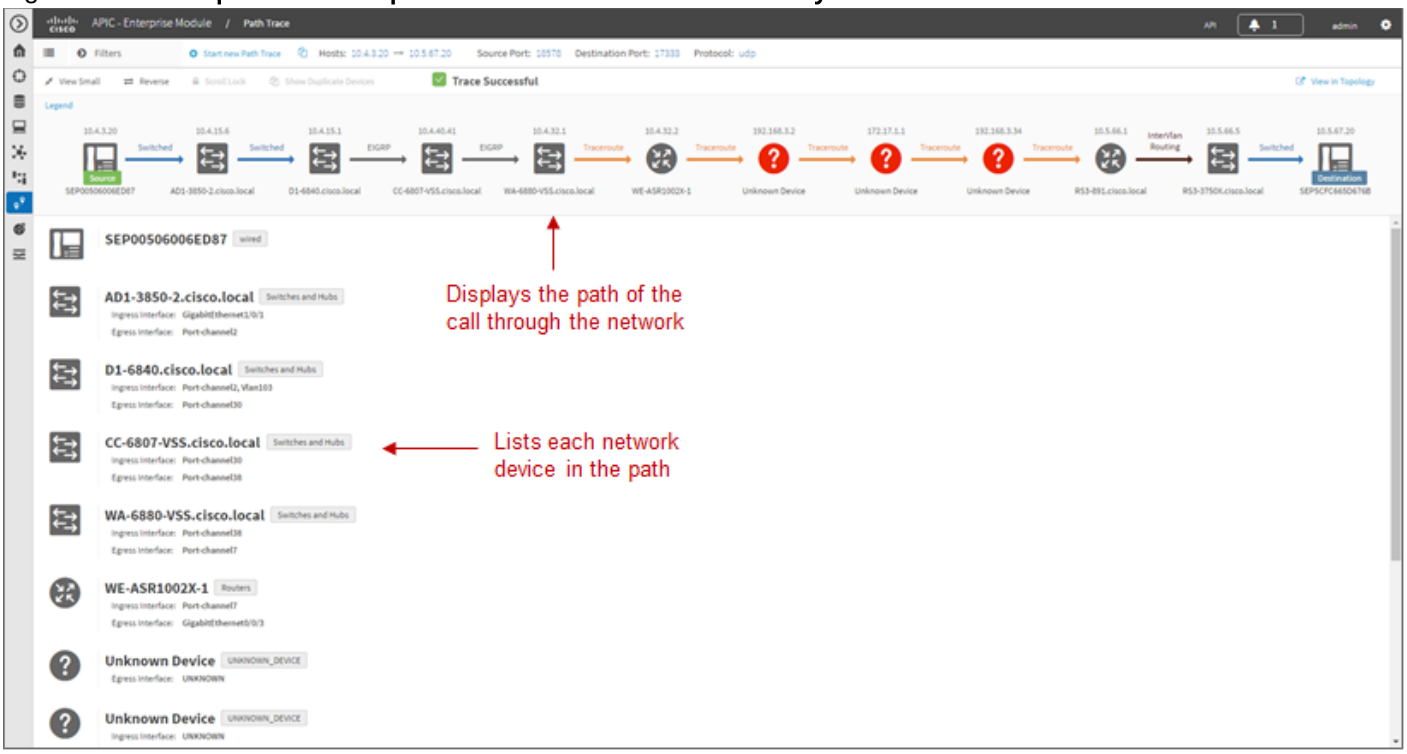

The Path Trace tool displays the path of the traffic associated with the Dynamic QoS flow through the network infrastructure within the top panel of the display. The lower panel provides more detail regarding each of the network devices through which the Dynamic QoS flow passes.

Note: In the example in the figure above, several network devices are listed as Unknown Devices. These devices represent Service Provider WAN network devices. Devices not under the administrative control of APIC-EM are not discovered and added to the Network Devices database within APIC EM. APIC-EM may be able to determine the existence of such devices within the output of the Path Trace tool. However, these devices may appear as Unknown Devices within the Path Trace tool.

When call signaling systems, such as CUCM, use the northbound REST-based API to signal to APIC-EM that a call has been terminated, APIC-EM will remove the entry for the Dynamic QoS flow. The *Dynamic QoS Design* section of this document has further details regarding the operation of Dynamic QoS.

WAN and Branch Static QoS Design

Within the EasyQoS solution, different network devices implement the ingress classification & marking QoS policies to the best of their abilities. Cisco ISR G2 Series, ISR 4400 Series, and ASR 1000 Series router platforms implement the following QoS policies:

- Ingress classification & marking policies based on AVC/NBAR2 policy-maps that contain either "match protocol attribute" or "match protocol" statements.
- Egress queuing policies

EasyQoS Policy Based on Platform, NBAR2 Protocol Pack, and Licensing

The following table summarizes the ingress classification & marking policy provisioned by EasyQoS to Cisco IOS and IOS XE platforms based upon software release, NBAR protocol pack version, and protocol pack license (Standard versus Advanced). Platforms that run IOS software releases include Cisco ISR G2 (3900 Series, 2900 Series, and 800 Series) platforms. Platforms that run IOS XE software releases include Cisco ISR 4400 Series and Cisco ASR 1000 Series platforms.

Table 1 Ingress Classification & Marking Policy for ISR and ASR Platforms

Platform Type	IOS Release	Protocol Pack Version	Ingress Classification & Marking Policy
IOS XE or IOS	Any IOS XE or IOS release	Standard Protocol Pack	No ingress classification & marking policy Ingress classification & marking policies are not supported on devices running Standard Protocol Pack on both IOS and IOS-XE platforms
IOS XE or IOS	IOS XE 3.12 or below Or IOS XE versions 3.13.6 to IOS XE 3.14 (excluding) Or IOS versons below 15.5(1)T	Any Advanced Protocol Pack	Ingress classification & marking policy using "match protocol" statements Custom applications that include a hyphen will not be programmed
IOS XE	IOS XE versions 3.13.1 to 3.13.5 and 3.14 to 3.16 Or IOS 15.5(1)T and 15.5(2)T	Any Protocol Pack	No ingress classification & marking policy (Cisco software defects–see note below)
IOS XE	IOS XE versions 3.16.1 to 3.16.3 Or IOS versions 15.5(3)M to 15.5(3)M3	Advanced Protocol Pack versions below 22.0.0	No ingress classification & marking policy (Cisco software defects–see note below)

IOS XE	IOS XE versions 3.16.1 to 3.16.3 Or IOS versions 15.5(3)M to 15.5(3)M3	Advanced Protocol Pack versions 22.0.0 or higher	Ingress classification & marking policy using “match attribute” statements
IOS XE	IOS XE versions 3.16.4 or later OR IOS 15.5(3)M4 or later	Advanced Protocol Pack versions 14.0.0 or higher	Ingress classification & marking policy using “match attribute” statements

ISR G2 Series platforms require a Data license for NBAR2 Advanced Protocol Pack. ISR 4000 Series platforms require an Application Experience license for NBAR2 Advanced Protocol Pack. ASR 1000 Series platforms require an Advanced Enterprise Services or Advanced IP Services license for NBAR2 Advanced Protocol Pack. EasyQoS will always push an egress queuing policy to a supported ISR or ASR router platform, regardless of the IOS XE or IOS software version, NBAR protocol pack version, and protocol pack license (Standard or Advanced).

Note: Although the business-relevance and traffic-class attributes are supported with IOS XE software versions that support Advanced Protocol Pack 14.0.0 and higher, due to CSCva30089, ingress classification & marking policies are not provisioned to Cisco ISR 4400 and ASR 1000 Series routers by EasyQoS unless the IOS XE software version is upgraded as shown in the table above.

NBAR2 QoS Attributes

Cisco NBAR Protocol Pack 14.0.0 introduced two new attributes–“traffic-class” and “business-relevance.” All 1300+ applications known to NBAR have been given a default value for each of these attributes.

Traffic-Class Attribute

Every application within the NBAR taxonomy for Protocol Pack14.0.0 and higher has also been assigned to one of the following 12 traffic classes:

- VoIP Telephony
- Broadcast Video
- Real Time Interactive
- Multimedia Conferencing
- Multimedia Streaming
- Network Control
- Ops Admin Mgmt
- Signaling

- Transactional Data
- Bulk Data
- Scavenger
- Default

These 12 traffic classes correspond to the 12-class QoS model recommended in IETF RFC 4594 with minor modifications (Signaling traffic marked CS3 and Broadcast Video traffic marked CS5 with the Cisco model). An example of the Cisco RFC 4594-Based 12-Class QoS model was shown in Figure 5 earlier in this document.

Business-Relevance Attribute

Every application within the NBAR taxonomy for NBAR Protocol Pack 14.0.0 and higher has one of the following three settings for the business-relevance attribute:

- Business relevant–these applications directly support business objectives.
- Business irrelevant–these applications do not support business objectives and are typically consumer-oriented.
- Default–these applications may/may not support business objectives (e.g. HTTP/HTTPS/SSL).

Business-relevant applications are intended to be serviced within their respective RFC 4594 traffic-class. Business-irrelevant applications are intended for a RFC 3662 lower than best effort or Scavenger traffic-class treatment. Applications with business-relevancy settings of default are intended for a RFC 2474 Default Forwarding treatment.

Ingress Classification & Marking Policies

As discussed in the *EasyQoS Policy Based on Platform, NBAR2 Protocol Pack, and Licensing* section above, the ingress classification & marking policy pushed by EasyQoS to ISR and ASR router platforms is dependent upon the IOS or IOS XE software version, the NBAR protocol pack version, and the NBAR protocol pack licensing of the platform. The following sections provide details regarding the policy.

Class-Map Definitions with "Match Protocol Attribute" Statements

The following is an example of the class-map definitions for the ingress classification & marking policy deployed by EasyQoS to ISR and ASR Series routers–based upon the use of "match protocol attribute" statements.

```
!
class-map match-all prm-MARKING_IN#VOICE
 match protocol attribute traffic-class voip-telephony
 match protocol attribute business-relevance business-relevant
class-map match-all prm-MARKING_IN#BROADCAST
```

```
 match protocol attribute traffic-class broadcast-video
 match protocol attribute business-relevance business-relevant
class-map match-all prm-MARKING_IN#REALTIME
 match protocol attribute traffic-class real-time-interactive
 match protocol attribute business-relevance business-relevant
class-map match-all prm-MARKING_IN#CONTROL
 match protocol attribute traffic-class network-control
 match protocol attribute business-relevance business-relevant
class-map match-all prm-MARKING_IN#SIGNALING
 match protocol attribute traffic-class signaling
 match protocol attribute business-relevance business-relevant
class-map match-all prm-MARKING_IN#OAM
 match protocol attribute traffic-class ops-admin-mgmt
 match protocol attribute business-relevance business-relevant
class-map match-all prm-MARKING_IN#MM_CONF
 match protocol attribute traffic-class multimedia-conferencing
 match protocol attribute business-relevance business-relevant
class-map match-all prm-MARKING_IN#MM_STREAM
 match protocol attribute traffic-class multimedia-streaming
 match protocol attribute business-relevance business-relevant
class-map match-all prm-MARKING_IN#TRANS_DATA
 match protocol attribute traffic-class transactional-data
 match protocol attribute business-relevance business-relevant
class-map match-all prm-MARKING_IN#BULK_DATA
 match protocol attribute traffic-class bulk-data
 match protocol attribute business-relevance business-relevant
class-map match-all prm-MARKING_IN#SCAVENGER
 match protocol attribute business-relevance business-irrelevant
class-map match-all prm-MARKING_IN#TUNNELED-NBAR
 match protocol capwap-data
!
```

The meaning of the "match-all" expression within class-map definitions that contain two "match" statements is that both lines must be true in order for traffic to be classified into the traffic class. For example, for the prm-MARKING_IN#SIGNALING class-map definition, matching traffic has to have both an NBAR traffic-class attribute of "signaling" and an NBAR business-relevance attribute of "business-relevant."

The prm-MARKING_IN#SCAVENGER class-map definition, is the only class-map definition that matches on an NBAR business-relevance attribute of "business-irrelevant." In other words, all applications marked as "business-irrelevant" within the APIC-EM EasyQoS GUI will match the prm-MARKING_IN#SCAVENGER class-map definition.

The addition of the prm-MARKING_IN#TUNNELED-NBAR class-map definition serves two purposes. First, it preserves the DSCP marking of Control And Provisioning of Wireless Access Points (CAPWAP) encapsulated data traffic. The DSCP marking of CAPWAP data traffic is based upon DSCP marking of the IP packet sent by the wireless client, in the upstream direction, and hence should be preserved. Second, the prm-MARKING_IN#TUNNELED-NBAR traffic-class is used for Custom applications as discussed in a following section.

Class-map Definitions with "Match Protocol" Statements

The following is an example of the class-map definitions for the ingress classification & marking policy deployed by EasyQoS to ISR and ASR Series routers–based upon the use of "match protocol" statements.

```
!
class-map match-any prm-MARKING_IN#TUNNELED-NBAR
 match protocol capwap-data
class-map match-any prm-MARKING_IN#VOICE
 match protocol cisco-jabber-audio
 match protocol cisco-phone
 match protocol cisco-phone-audio
 match protocol citrix-audio
 match protocol ms-lync-audio
…
class-map match-any prm-MARKING_IN#BROADCAST
 match protocol cisco-ip-camera
 match protocol dmp
class-map match-any prm-MARKING_IN#REALTIME
 match protocol cisco-phone-video
 match protocol rtp-video
 match protocol telepresence-media
class-map match-any prm-MARKING_IN#MM_CONF-NBAR
```

```
 match protocol adobe-connect
 match protocol cisco-jabber-im
 match protocol cisco-jabber-video
 match protocol ms-lync
 match protocol ms-lync-video
…
class-map match-any prm-MARKING_IN#MM_STREAM-NBAR
 match protocol apple-remote-desktop
 match protocol citrix
 match protocol citrix-static
 match protocol dameware-mrc
 match protocol gotomypc
…
class-map match-any prm-MARKING_IN#CONTROL
 match protocol aodv
 match protocol aurp
 match protocol bgmp
 match protocol bgp
 match protocol capwap-control
…
class-map match-any prm-MARKING_IN#SIGNALING
 match protocol cisco-jabber-control
 match protocol rtsp
 match protocol sip
 match protocol sip-tls
 match protocol skinny
…
class-map match-any prm-MARKING_IN#TRANS_DATA
 match protocol activesync
 match protocol banyan-rpc
 match protocol clearcase
 match protocol coauthor
```

```
 match protocol corba-iiop
…
class-map match-any prm-MARKING_IN#BULK_DATA
 match protocol afpovertcp
 match protocol bmpp
 match protocol cifs
 match protocol corba-iiop-ssl
 match protocol dicom
…
class-map match-any prm-MARKING_IN#SCAVENGER
 match protocol 4chan
 match protocol 58-city
 match protocol abc-news
 match protocol accuweather
 match protocol adcash
 match protocol addthis
…
```

The specific protocols that appear within the "match protocol" statements within the class-map definitions will vary, depending upon the deployment. This is based upon whether the network administrator has selected the particular protocol as having a business relevance of "business-relevant," "business-irrelevant," or "default" within the EasyQoS application for the particular scope to which the ISR G2 Series router belongs.

Modifying the Business Relevance of an Application

Network administrators have the ability to modify the business-relevance of applications within the EasyQoS graphical user interface and include these changes within policies pushed by APIC-EM to router and switch platforms. The *APIC-EM and the EasyQoS Application* section of this document shows how to modify the business relevance of applications within EasyQoS.

Modifying Business Relevance—Policy-maps with "Match Protocol Attribute" Statements

When the business-relevance of an application is modified and pushed to an ASR or ISR router platform that implements a policy-map containing class-map definitions that include "match protocol attribute business-relevance" or "match protocol attribute traffic-class" statements, EasyQoS will generate additional configuration within ISR and ASR router platforms.

First, EasyQoS creates one or all of the following attribute-map definitions shown below. The names of attribute-map definitions match the three values of the business-relevance attribute—APIC-A_M_RELEVANT, APIC-A_M-DEFAULT, and APIC-A_M-SCAVANGER.

```
!
ip nbar attribute-map APIC-A_M-RELEVANT
 attribute business-relevance business-relevant

ip nbar attribute-map APIC-A_M-DEFAULT
 attribute business-relevance default

ip nbar attribute-map APIC-A_M-SCAVENGER
 attribute business-relevance business-irrelevant
!
```

Under each of these attribute-map definitions, EasyQoS sets the business-relevance attribute.

- For the APIC-A_M-Relevant attribute-map definition, the business-relevance attribute is set to business-relevant.
- For the APIC-A_M-Default attribute-map definition, the business-relevance attribute is set to default.
- For the APIC-A_M-SCAVENGER attribute-map definition, the business-relevance attribute is set to business-irrelevant.

EasyQoS then maps each application that has been modified from whatever its default setting is within the NBAR2 taxonomy to one of the three attribute-map definitions above. This is accomplished via the "ip nbar attribute-set" command. In the example below, the application "ms-lync-video" has been mapped to a business-relevance of "business-irrelevant."

```
!
ip nbar attribute-set ms-lync-video APIC-A_M-SCAVENGER
!
```

Modifying Business Relevance—Policy-maps with "Match Protocol" Statements

When the business-relevance of an application is modified and pushed to an ISR or ASR router platform that implements a policy-map containing class-map definitions that include "match protocol" statements, the "match protocol" statement for the application will be modified as follows:

- If an application is moved from "business-relevant" or "default" to "business-irrelevant," the "match-protocol" statement for the application will appear under the prm-MARKING_IN#SCAVENGER traffic-class.
- By default, no applications within the NBAR taxonomy are classified with the traffic-class attribute of "scavenger". Therefore, if an application is moved from either "business-irrelevant" or "default" to "business-relevant", the "match-protocol" statement for the application will appear under one of the following nine class-map definitions—depending upon the traffic-class attribute of the particular application.

  - prm-MARKING_IN#VOICE
  - prm-MARKING_IN#BROADCAST
  - prm-MARKING_IN#REALTIME
  - prm-MARKING_IN#CONTROL
  - prm-MARKING_IN#OAM
  - prm-MARKING_IN#MM_CONF
  - prm-MARKING_IN#MM_STREAM
  - prm-MARKING_IN#TRANS_DATA
  - prm-MARKING_IN#BULK_DATA

- If the application is moved from either “business-relevant” or “business-irrelevant” to “default,” no “match-protocol” statement for the application will appear under any of the class-map definitions. This is because “match protocol” statements are not programmed for applications with a business-relevance of “default.”

Custom Applications on ASR and ISR Platforms

Network operators have the ability to add Custom applications within the EasyQoS graphical user interface and to include these Custom applications within policies pushed by APIC-EM to router and switch platforms. The *APIC-EM and the EasyQoS Application* section of this document shows how Custom applications are created and added to policy scopes within EasyQoS. Custom applications can be either be specified by a URL string or by one or more server IP addresses and UDP/TCP ports.

URL-Based Applications–Policy-maps with “Match Protocol Attribute” Statements

For Custom applications that are specified based on a URL string, EasyQoS will generate additional configuration within ISR and ASR router platforms similar to the following example:

```
!
ip nbar attribute-map Custom_URL-App
 attribute traffic-class transactional-data
 attribute business-relevance business-relevant
 attribute category other
 attribute sub-category other
!
~
!
ip nbar custom Custom_URL_App http url "http://example.custom.com" id 16299
!
```

```
~
!
ip nbar attribute-set Custom_URL_App Custom_URL_App
!
```

This first block of configuration creates an attribute profile (named Custom_URL_App in the example above). The name of the attribute profile corresponds to the name of the Custom application specified by the network administrator when creating the Custom application definition within the EasyQoS web-based GUI. The configuration then assigns the attribute profile several attributes, including a traffic-class attribute and a business-relevance attribute. In the example Custom application, a traffic-class attribute of "transactional-data" and a business-relevance attribute of "business-relevant" have been assigned to the attribute profile.

The second block (single line) of configuration above defines a web-based custom protocol match, specifying the URL string that is used to match on the name of the custom protocol (also named Custom_URL_App in the example above).

Finally, the third block (single line) of configuration maps the attribute profile to the web-based custom protocol match—both defined in the previous two blocks of configuration. In other words, the custom protocol is assigned the attributes specified within the attribute profile.

The effect of this configuration example is that the custom protocol defined by the URL string "http://example.custom.com" will match the prm-MARKING_IN#TRANS_DATA class-map definition and be treated as Transactional Data traffic. Additional URL-based Custom applications will generate additional configuration blocks similar to those shown in the example above.

URL-Based Applications–Policy-maps with "Match Protocol" Statements

URL-based applications are not programmed into ISR and ASR router platforms that implement a policy-map containing class-map definitions that include "match protocol" statements.

Server IP/Port Based Applications

Custom applications that are specified by one or more server IP address and UDP/TCP port numbers are handled through the creation of an ACL with ACEs on ISR and ASR router platforms. Specifically an ACL for the prm-MARKING_IN#TUNNELED-NBAR traffic-class, called prm-MARKING_IN#TUNNELED-NBAR__acl, is generated and populated with ACE entries. An example is shown below.

```
!
ip access-list extended prm-MARKING_IN#TUNNELED-NBAR__acl
 remark Custom_Port-App
 permit udp any 10.0.10.0 0.0.0.255 range 3001 3010
 permit udp 10.0.10.0 0.0.0.255 range 3001 3010 any
!
```

In the example above, the Custom application—based on a destination server IP address range and port range (also referred to as the producer)—has been specified to be bi-directional by the network operator

through the EasyQoS web-based GUI. Hence, the reverse of the ACE entry is also generated to allow traffic from the server IP address and port range to also be treated the same.

In the example above, a server IP address range (10.0.10.0-10.0.10.255) and port range (UDP 3001-3010) is configured. Custom applications also support single IP addresses and ports, or the use of "any" specified as the destination IP address. Although a single UDP port range is specified in the example above, multiple UDP and/or TCP ports can be configured as well–each of which would appear as a separate "permit" statement.

Additional IP Address/Port-based Custom applications will generate additional ACE entries within the prm-MARKING_IN#TUNNELED-NBAR__acl, similar to those shown in the example above.

A more sophisticated example shown below, adds a source IP address or range (referred to as the consumer) as well as the destination IP address or range (referred to as the producer) to the Custom application. Again, this is configured bi-directionally via the APIC-EM EasyQoS web-based GUI by the network operator. An example of the same application–but with a consumer–is shown below.

```
!
ip access-list extended prm-MARKING_IN#TUNNELED-NBAR__acl
 remark Custom_Port-App_Consumer__Custom_Port-App
 permit udp host 10.0.20.20 eq 3100 10.0.10.0 0.0.0.255 range 3001 3010
 permit udp 10.0.10.0 0.0.0.255 range 3001 3010 host 10.0.20.20 eq 3100
!
```

The combination of the producer and consumer, along with the ability to apply the policy bi-directionally, essentially gives the network operator the ability to use nearly the full CLI functionality in terms of being able to configure QoS ACE entries.

After the ACL and ACE entries have been generated, EasyQoS adds the ACL entry to the prm-MARKING_IN#TUNNELED-NBAR class-map definition via a "match access-group" statement. This is regardless of whether the class-map definitions within the ingress classification & marking policy-map uses "match protocol attribute" or "match protocol" statements. An example is shown below.

```
!
class-map match-all prm-MARKING_IN#TUNNELED-NBAR
 match protocol capwap-data
 match access-group name prm-MARKING_IN#TUNNELED-NBAR__acl
!
```

Notice that the "match protocol capwap-data" entry is maintained. The "match access-group" statement is simply added to the class-map definition. The prm-MARKING_IN#TUNNELED-NBAR class is used for Custom applications because it is the only class-map definition within the policy-map that specifies no action. Therefore, any DSCP markings for the Custom application specified at the ingress access-edge switch are maintained, when the traffic passes through an ISR or ASR router platform.

Policy-map Definition

The following is an example of the policy-map definition for the ingress classification & marking policy deployed by EasyQoS to ISR and ASR routers—regardless of whether the class-map definitions within the ingress classification & marking policy-map uses “match protocol attribute” or “match protocol” statements.

```
!
policy-map prm-MARKING_IN
 class prm-MARKING_IN#TUNNELED-NBAR
 class prm-MARKING_IN#VOICE
  set DSCP ef
 class prm-MARKING_IN#BROADCAST
  set DSCP cs5
 class prm-MARKING_IN#REALTIME
  set DSCP cs4
 class prm-MARKING_IN#MM_CONF
  set DSCP af41
 class prm-MARKING_IN#MM_STREAM
  set DSCP af31
 class prm-MARKING_IN#CONTROL
  set DSCP cs6
 class prm-MARKING_IN#SIGNALING
  set DSCP cs3
 class prm-MARKING_IN#OAM
  set DSCP cs2
 class prm-MARKING_IN#TRANS_DATA
  set DSCP af21
 class prm-MARKING_IN#BULK_DATA
  set DSCP af11
 class prm-MARKING_IN#SCAVENGER
  set DSCP cs1
 class class-default
  set DSCP default
!
```

The policy-map sets the DSCP marking, and hence the per-hop behavior, for traffic matching the particular traffic class to meet Cisco's RFC-4594 based recommendations for a 12-class QoS model, shown in Figure 5 earlier in this document.

Application of the Ingress Classification & Marking Policy to Interfaces

The ingress classification & marking policy is applied to all Ethernet interfaces on the ISR or ASR router platform. An example of the application of the ingress classification & marking policy is as follows:

```
!
interface GigabitEthernet0/1
 service-policy input prm-MARKING_IN
!
```

For brownfield deployments, EasyQoS will remove any existing ingress classification & marking service-policy statements that appear on the interface, before applying the prm-MARKING_IN service-policy. However, policy-map and class-map definitions for the existing policy will remain within the configuration of the ASR or ISR router platform.

WAN-Edge Egress Queuing Policy

The WAN-edge egress queuing policy is deployed to the following interfaces:

- WAN links that are not connected to service provider managed-service offerings requiring the support of sub-line rate speeds and the re-marking of traffic to meet the traffic classes provided by the service provider.
- LAN links between the ISR or ASR router and the Catalyst switch

The following figure shows the WAN bandwidth allocation model for the WAN-edge egress queuing policy.

Figure 43 **WAN Bandwidth Allocation Model for the WAN-Edge Egress Queuing Policy**

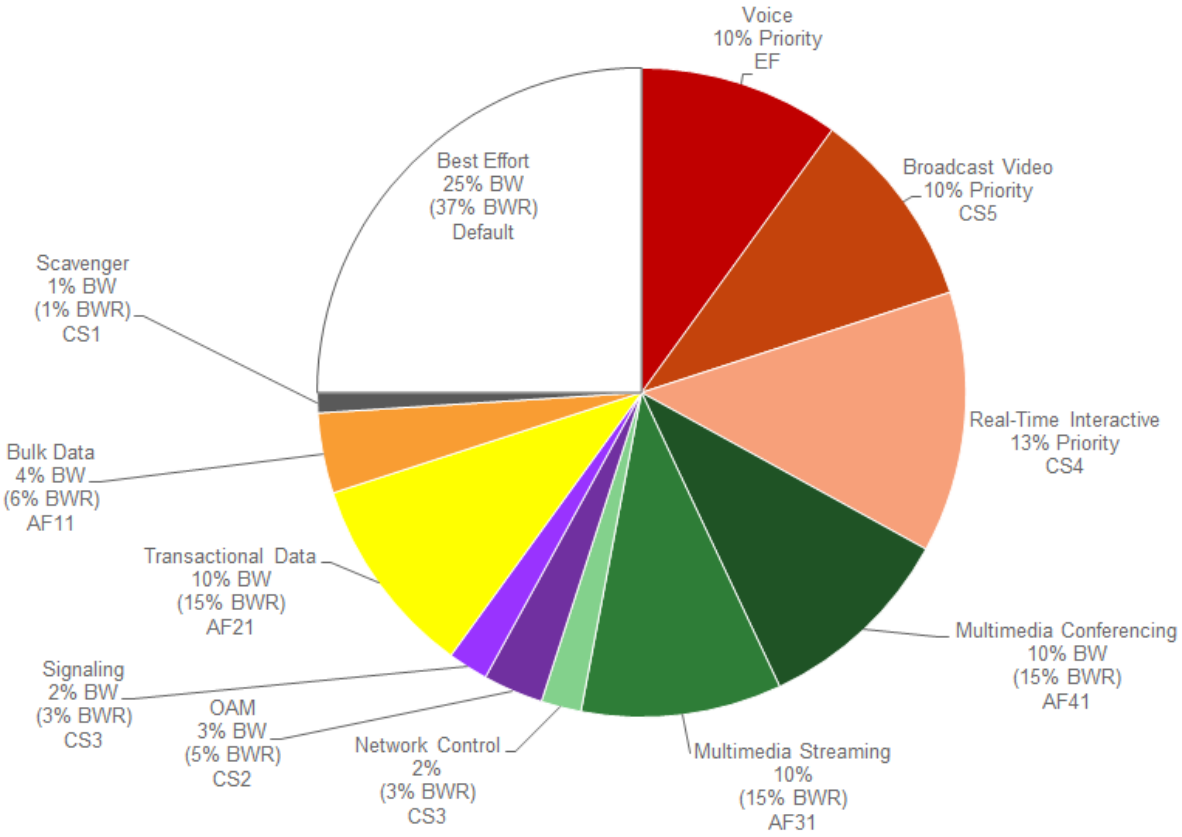

The bandwidth allocations for the WAN edge queuing policy are fixed and cannot be modified in APIC-EM version 1.3. Because queuing is done in software on ISR and ASR router platforms, the WAN-edge egress queuing policy implements a 12 queue model–meaning a queue for each of the classes within the RFC 4594-based 12-class QoS model shown in Figure 5 earlier in this document.

Class-map Definitions

The following are the class-map definitions for each of the 12 queues provisioned by EasyQoS.

```
!
class-map match-any prm-EZQOS_12C#VOICE
 match DSCP ef
class-map match-any prm-EZQOS_12C#BROADCAST
 match DSCP cs5
class-map match-any prm-EZQOS_12C#REALTIME
 match DSCP cs4
class-map match-any prm-EZQOS_12C#CONTROL
 match DSCP cs6
class-map match-any prm-EZQOS_12C#SIGNALING
```

```
 match DSCP cs3
class-map match-any prm-EZQOS_12C#OAM
 match DSCP cs2
class-map match-any prm-EZQOS_12C#MM_CONF
 match DSCP af41
 match DSCP af42
 match DSCP af43
class-map match-any prm-EZQOS_12C#MM_STREAM
 match DSCP af31
 match DSCP af32
 match DSCP af33
class-map match-any prm-EZQOS_12C#TRANS_DATA
 match DSCP af21
 match DSCP af22
 match DSCP af23
class-map match-any prm-EZQOS_12C#BULK_DATA
 match DSCP af11
 match DSCP af12
 match DSCP af13
class-map match-any prm-EZQOS_12C#SCAVENGER
 match DSCP cs1
!
```

Policy-map Definition

The following is an example of the policy-map definition for the WAN-edge egress queuing policy for an ISR or ASR router, provisioned by EasyQoS.

```
!
policy-map prm-DSCP#QUEUING_OUT
 class prm-EZQOS_12C#VOICE
  police rate percent 10
  priority
 class prm-EZQOS_12C#BROADCAST
  police rate percent 10
```

```
 priority
class prm-EZQOS_12C#REALTIME
 police rate percent 13
 priority
class prm-EZQOS_12C#MM_CONF
 bandwidth remaining percent 15
 fair-queue
 random-detect DSCP-based
class prm-EZQOS_12C#MM_STREAM
 bandwidth remaining percent 15
 fair-queue
 random-detect DSCP-based
class prm-EZQOS_12C#CONTROL
 bandwidth remaining percent 3
class prm-EZQOS_12C#SIGNALING
 bandwidth remaining percent 3
class prm-EZQOS_12C#OAM
 bandwidth remaining percent 5
class prm-EZQOS_12C#TRANS_DATA
 bandwidth remaining percent 15
 fair-queue
 random-detect DSCP-based
class prm-EZQOS_12C#BULK_DATA
 bandwidth remaining percent 6
 fair-queue
 random-detect DSCP-based
class prm-EZQOS_12C#SCAVENGER
 bandwidth remaining percent 1
class class-default
 bandwidth remaining percent 37
 fair-queue
 random-detect DSCP-based
```

```
    random-detect DSCP 0 50 64        ! ISR G2 and 800 Series platforms only.
  !
```

The WAN-Edge egress queuing policy implements a three LLQ policy, meaning a separate LLQ for each of the Voice, Broadcast-Video, and Realtime-Interactive traffic classes. The Voice queue supports traffic with an EF per hop behavior. The Broadcast-Video queue supports traffic with a Class Selector 5 (CS5) per hop behavior. The Realtime-Interactive traffic class supports traffic with a CS4 per hop behavior. Broadcast-Video and Realtime-Interactive traffic classes are meant to support traffic flows that are inelastic–meaning the endpoints generating the flows do not down-speed their transmission rate when packet loss occurs. Because of the inelastic nature of these flows, they are eligible for LLQ treatment on the ISR or ASR router platforms, along with Voice traffic. Explicit policers (10%, 10%, and 13% respectively) ensure that each of the LLQs can use no more than the percentage of the bandwidth of the WAN link allocated to the traffic class, regardless of whether there is available bandwidth.

The remaining eight queues share the remaining bandwidth based on a percentage allocation of bandwidth. This is accomplished via the "bandwidth remaining percent" command. Each of these queues can use more than its percentage allocation, if more bandwidth is available–meaning if the one or more of the other queues is not using its full allocation of remaining bandwidth percentage.

The Multimedia-Conferencing, Multimedia-Streaming, Transactional Data, and Bulk Data queues support traffic with Assured Forwarding (AF) per-hop behaviors (AF4x, AF3x, AF2x, and AF1x respectively). Fair-queuing, along with DSCP-based WRED is implemented for these traffic classes. The minimum and maximum WRED thresholds for these queues is left at their default values.

For ISR G2 and 800 Series platforms, the default queue limit size is 64 packets per queue. The minimum and maximum WRED thresholds are expressed in terms of the number of packets. The default WRED thresholds for the AF per-hop behaviors, and the drop probability is shown in the following table.

Table 2 ISR G2 WRED Minimum & Maximum Threshold and Drop Probability Default Values

Per-Hop Behavior and DSCP Value	Minimum Threshold (Packets)	Maximum Threshold (Packets)	Drop Probability
AF11 (DSCP 10)	32	40	1/10
AF12 (DSCP 12)	28	40	1/10
AF13 (DSCP 14)	24	40	1/10
AF21 (DSCP 18)	32	40	1/10
AF22 (DSCP 20)	28	40	1/10
AF23 (DSCP 22)	24	40	1/10
AF31 (DSCP 26)	32	40	1/10
AF32 (DSCP 28)	28	40	1/10
AF33 (DSCP 30)	24	40	1/10
AF41 (DSCP 34)	32	40	1/10

AF42 (DSCP 36)	28	40	1/10
AF43 (DSCP 38)	24	40	1/10
Default (DSCP 0)	20	40	1/10

The Default queue also implements fair-queuing, along with DSCP-based WRED. WRED is effective here at preventing TCP synchronization of flows, which can result in overall lower throughput and bandwidth utilization. For the ISR G2 and 800 Series platforms, the default WRED thresholds for the Default queue are considered to be too aggressive–meaning the minimum drop threshold is set lower than desired. Hence, the minimum drop threshold has been adjusted to 50 packets, and the maximum drop threshold adjusted to the depth of the queue–64 packets. For the ISR 4400 and ASR 1000 Series platforms, the default WRED thresholds for the Default queue are left at their default values.

The Control, Signaling, OAM, and Scavenger queues each support a single Class Selector (CS) per hop behavior (CS6, CS3, CS2, and CS2, respectively). For the Control, Signaling, and OAM traffic classes, WRED is not implemented. Randomly discarding network control, signaling, or operational traffic when a minimum queue depth threshold is exceeded, may simply result in degraded network performance. Hence these queues implement tail-drop at the back of the queue, because the objective is to not drop traffic in these queues by provisioning sufficient remaining bandwidth percentage allocation to these queues.

The Scavenger queue is considered to be a bandwidth-constrained queue for less-than-best-effort treatment. WRED is not implemented for this queue, because the consideration is not to optimize the use of this queue but simply to provision some minimal amount of bandwidth for support of traffic within this queue.

Application of the Egress Queuing Policy to Interfaces

The egress queuing policy is applied to all Ethernet interfaces on the ISR or ASR router platform. An example of the application of the egress queuing policy is as follows:

```
!
interface GigabitEthernet0/1
 service-policy output prm-DSCP#QUEUING_OUT
!
```

For brownfield deployments, EasyQoS will remove any existing egress queuing service-policy statements that appear on the interface, before applying the prm-DSCP#QUEUING_OUT service-policy. However, policy-map and class-map definitions for the existing policy will remain within the configuration of the ASR or ISR router platform. This provides the network operator the option to restore the configuration of the ISR or ASR router platform to its original non-EasyQoS policy, should that be necessary.

ASR-1000 Series Specific Interface-Level Commands

For the ASR-1000 Series platforms, additional interface-level configuration commands are provisioned by APIC-EM EasyQoS. Network input/output on the ASR-1000 Series platforms consists of shared port adapters (SPAs) controlled by one or more SPA interface processors (SIPs). Ethernet and ATM SPAs perform Layer 2 and Layer 3 packet classification, and they also decide on the internal priority of the packet–

high priority or low priority. High-priority packets are sent on separate channels to the embedded services processor (ESP) than low-priority packets. QoS is then performed within the ESP. The SPA queues packets on high channels to high-priority buffers, and packets on low channels to low-priority buffers. Internal classification of packets can be based on DSCP, IPv6 traffic class, MPLS EXP or 802.1Q/P class of service (CoS) values.

APIC-EM EasyQoS enables SPA-based internal scheduling and classification by provisioning the following commands on ASR 1000 Series platforms:

```
!
plim qos input map ip DSCP-based
plim qos input map ip DSCP 32 40 46 queue strict-priority
!
```

The first command enables DSCP-based classification within the SPA. By default, EF (voice) traffic is mapped to the strict-priority internal queue, and all other DSCP values are mapped to the low-priority internal queue. The second command modifies this by mapping CS4 (real-time interactive) and CS5 (broadcast video) traffic to the strict-priority internal queue.

Note: When restoring the configuration on an ASR 1000 Series platform to a pre-EasyQoS policy, APIC-EM will not remove any of the "plim qos" commands configured on the platform. This must be manually removed by the network operator if desired.

Pre-Existing QoS Configurations on ISR and ASR Router Platforms

This section discusses how EasyQoS handles prior QoS configurations on ISR and ASR router platforms when deploying a QoS policy. IOS (ISR G2 and 800 Series routers) and IOS XE (ISR 4400 and ASR 1000 Series routers) platforms implement both ingress classification & marking policies and queuing policies by applying a service-policy definition across interfaces. The service-policy definition references an existing policy-map definition. For both ingress classification & marking policies and queuing policies, EasyQoS will remove any existing service-policy definition from the interface and replace it with its service-policy definitions. The previous class-map and policy-map definitions will not be deleted by EasyQoS. This is necessary for restoring the original pre-EasyQoS (before any EasyQoS configuration was applied) configuration back to the switch platform. Clicking on the Restore button within an EasyQoS policy will cause the pre-EasyQoS classification & marking and queuing service-policy statements to be re-applied to the interfaces. The Restore feature is a new feature added to APIC-EM EasyQoS release 1.3.

Note: If the network operator as manually deleted the original QoS configuration (that is, the Pre-EasyQoS configuration)—meaning the policy-map and class-map definitions—the Restore feature will not be able to restore the QoS configuration on the device to its original configuration.

The Restore feature will not remove any "plim qos" commands on the ASR 1000 Series platforms discussed in the *ASR-1000 Series Specific Interface-Level Commands* section of this document.

Service Provider Managed-Service WAN QoS Design

Challenges

WAN connectivity to a service provider managed-service offering may involve sub-line rate bandwidth provisioning—meaning that the provisioned bandwidth is below the physical interface of the ISR or ASR router platform. For example, it is common to provision a managed-service offering in which the physical connectivity between the Customer Edge router and the Provider Edge router is a Gigabit Ethernet connection. However, the contracted rate between the service provider and the organization is only provisioned for perhaps 50 Mbps or 100 Mbps of total bandwidth.

The contracted rate may be further sub-divided into multiple traffic classes. It is common for service providers to offer between four and eight traffic classes. Some of these traffic classes provide Service Level Agreements for support of real-time (priority) traffic such as voice and video support, while others provide data or best effort service. The number of traffic classes supported by the service provider, the percentage bandwidth allocation between the traffic classes, and the supported DSCP markings of those traffic classes—are collectively referred to as the *service provider profile* (SPP).

In order to support deployments that have managed-service offerings, APIC-EM must determine the following when deploying QoS policy to ISR/ASR router platforms:

1. Is a WAN interface connected to a managed-service offering?
2. If so, what is the sub-line rate of the managed-service offering (if any)?
3. What is the service provider profile for this managed-service offering—meaning how many traffic-classes are implemented by the service provider, are any eligible for priority treatment, what is the expected mapping of the DSCP values from the traffic classes within the organization to the traffic classes within service-provider, and what are the percentage bandwidth allocations between the service provider traffic classes?

EasyQoS supports four default SPP models. Each of the default SPP models supports the following:

- A fixed number of traffic classes (4, 5, 6, and 8 classes)
- A fixed mapping of the DSCP values and priority treatment from the traffic classes within the organization to the traffic classes within the service-provider network
- Fixed bandwidth allocations between the service provider traffic classes.

Additionally, the APIC-EM 1.3 release of EasyQoS supports the ability to create custom SPPs based on the default 4, 5, 6, and 8 class SPP models. Custom service provider profiles allow the network administrator to specify the mapping of the DSCP values from the traffic classes within the organization to the traffic classes within the service-provider network, as well as to specify the percentage bandwidth allocations between the service provider traffic classes.

EasyQoS requires the network administrator to tag WAN interfaces with a specific string in the interface-description in order to identify the items listed in the three questions above. This must be configured before deploying a QoS policy to the platform via EasyQoS.

There are up to three important fields within the tag. Each of the fields within the tag is delineated via a "#". The meaning of the fields within the tag is discussed in the following sections.

Identifying WAN Interfaces

EasyQoS requires the network administrator to tag WAN interfaces that connect to a service-provider managed service with a specific string in the interface-description: #WAN#. This is the first field in the overall tag discussed in the previous section, and is a required field.

An example of the configuration is shown below with the first part of the tag shown in bold.

```
!
interface GigabitEthernet0/0
 description CIRCUIT TO WE-ASR2 GIG-0-0-1 #WAN#50M#SPP:test1#
!
```

Note: If the network administrator has configured no tag on a WAN interface, EasyQoS applies the WAN-edge egress queuing policy discussed in the ***WAN and Branch Static QoS Design*** section of this document. This is because the WAN-edge egress queuing policy is also applied to LAN connections between the ISR or ASR platform and the Catalyst switch, and LAN connections are not required to have any tag within their interface descriptions.

Currently the #WAN# part of the overall tag provides no additional functionality in the context of a service-provider managed-service other than to identify the interface as a WAN connection to a service-provider managed-service.

Identifying Sub-Line-Rate WAN Interfaces

Optionally, when connecting to a service provider managed-service using sub-line rate connectivity EasyQoS requires the network administrator to tag WAN interfaces with the sub-line rate–meaning the overall provisioned bandwidth of the service contracted from the service provider. The sub-line rate is tagged with a specific string in the interface-description: #rate#. This is the second field in the overall tag discussed previously. If a sub-line rate service is not provisioned, this field can be omitted within the overall tag.

An example of the configuration is shown below with the second part of the tag shown in bold.

```
!
interface GigabitEthernet0/0
 description CIRCUIT TO WE-ASR2 GIG-0-0-1 #WAN#50M#SPP:New-6-Class#
!
```

The rate is specified using abbreviations–"M" for Mbps. In the example above "50M" stands for a sub-line rate of 50 Mbps contracted from the service provider. This rate is read by APIC-EM during inventory process and is then used by EasyQoS to provision the shaper at the top-level of the hierarchical egress queuing policy. This shaper is necessary to provide the back-pressure in order for QoS to be engaged on the WAN link, when implementing a sub-line rate service.

Identifying WAN Interfaces Mapped to SP Class-of-Service Models

EasyQoS requires the network administrator to tag WAN interfaces with either the name of one of the four default SPPs or the name of a custom profile, when connected to a service provider managed-service. The format of the tag is dependent upon whether one of the four default service provider profiles is to be attached to the interface or whether a custom service provider profile is to be attached to the interface.

Default Service Provider Profiles

When implementing one of the four default service provider profiles, the format of the tag can take one of the two forms.

The first form provides backward compatibility with prior versions of APIC-EM EasyQoS.

```
#WAN#rate#SPPx#
```

The "x" in "SPPx" refers to one of the four default service provider profiles, which are discussed in detail in the following sections.

The second form is uses the same format as custom service provider profiles.

```
#WAN#rate#SPP:SPPx-yClass
```

The "x" in SPP:SPPx-yClass" refers to one of the four default service provider profiles, which are discussed in detail in the following sections. The "y" in "SPP:SPPx-yClass" refers to the number of traffic-classes supported by the service provider.

The two forms can be used to express the same service provider profile as shown below:

- SPP1 = SPP:SPP1-4Class
- SPP2 = SPP:SPP2-5Class
- SPP3 = SPP:SPP3-6Class
- SPP-4 = SPP:SPP4-8Class

An example of the configuration of an interface description using the default service provider profile SPP1 within the tag is shown below, with the third part of the tag shown in bold.

```
!
interface GigabitEthernet0/0
 description CIRCUIT TO WE-ASR2 GIG-0-0-1 #WAN#50M#SPP1#
!
```

Custom Service Provider Profiles

When implementing one of the custom service provider profiles, the format of the tag is as follows:

```
#WAN#rate#SPP:custom_profile_name#
```

The "custom_profile_name" refers to the name of the custom service provider profile created within EasyQoS. This is discussed in the *APIC-EM and the EasyQoS Application* section of this document. An example of the configuration of an interface description using a custom service provider profile named "New-6-Class" within the tag is shown below, with the third part of the tag highlighted.

```
!
interface GigabitEthernet0/0
 description CIRCUIT TO WE-ASR2 GIG-0-0-1 #WAN#50M#SPP1:New-6-Class#
!
```

Each of the four default service provider profiles, as well as custom service provider profiles is discussed in the *Service Provider Default Class of Service Models* section below.

Note: If the tag within the interface description is added, removed, or modified, the network administrator must wait until APIC-EM re-synchronizes the configuration of the ISR or ASR router by running its inventory process again before re-applying any QoS policy to the device. APIC-EM will synchronize the configuration of network devices approximately every 30 minutes. If a QoS policy is re-applied to the device by EasyQoS before APIC-EM has re-synchronized the configuration with its internal database, the EasyQoS policy may not reflect the desired changes to the policy, based on the changes to the interface description.

Service Provider Default Class of Service Models

EasyQoS supports connectivity to service provider managed-service offerings, using one of the following four default SPP class-of-service models for ISR and ASR router platforms:

- SPP1/SPP:SPP1-4Class
- SPP2/SPP:SPP2-5Class
- SPP3/SPP:SPP3-6Class
- SPP4/SPP:SPP4-8Class

Because queuing is done in software on ISR and ASR router platforms, all SPP models implement an egress queuing policy consisting of 12 egress queues—one for each of the traffic classes as shown in Figure 5 earlier in this document.

APIC-EM determines which of the four SPP models to deploy based on the #SPPx#" field (where x is from 1 to 4) or "#SPP:SPPx-yClass#" (where x is from 1 to 4 and y is 4, 5, 6, or 8) within the description configured on the ISR or ASR WAN interface connected to the service provider managed-service.

SPP1 or SPP:SPP1-4Class

The SPP1/SPP:SPP1-4Class is based on managed-service offerings with four traffic classes. These traffic classes are specified as follows within this document:

- SP-Voice
- SP-Class1Data
- SP-Class2Data
- SP-Default

The following figure shows the WAN bandwidth allocation for the SPP1/SPP:SPP1-4Class service provider profile.

Figure 44 **WAN Bandwidth Allocation for the SPP1/SPP:SPP1-4Class**

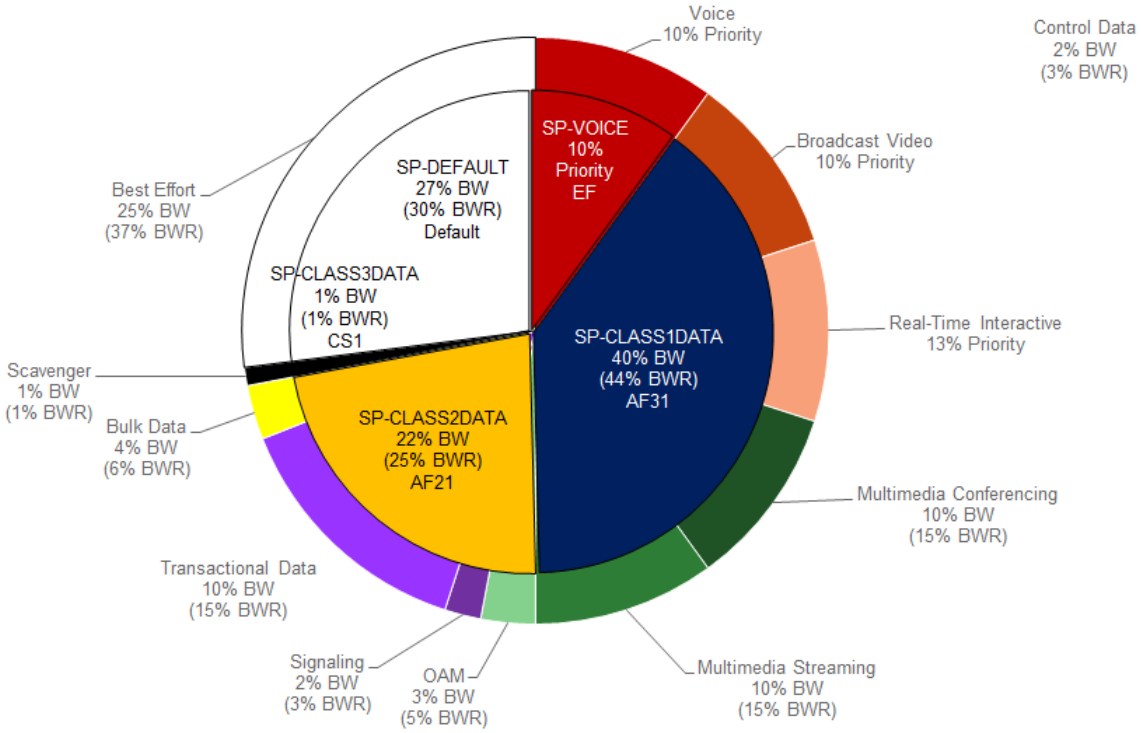

The class-map definitions provisioned by EasyQoS for all of the SPPs discussed here are the same as was as discussed in the *WAN-Edge Egress Queuing Policy* section of this document and will not be duplicated here.

For the SPP1/SPP:SPP1-4Class model, EasyQoS must map the RFC 4594-based 12-class QoS model implemented within the organization into the four traffic classes provide by the service provider. The following figure shows this mapping with the bandwidth allocations and traffic re-marking for the service provider traffic classes.

Figure 45 EasyQoS Marking Mappings Into SPP1/SPP:SPP1-4Class

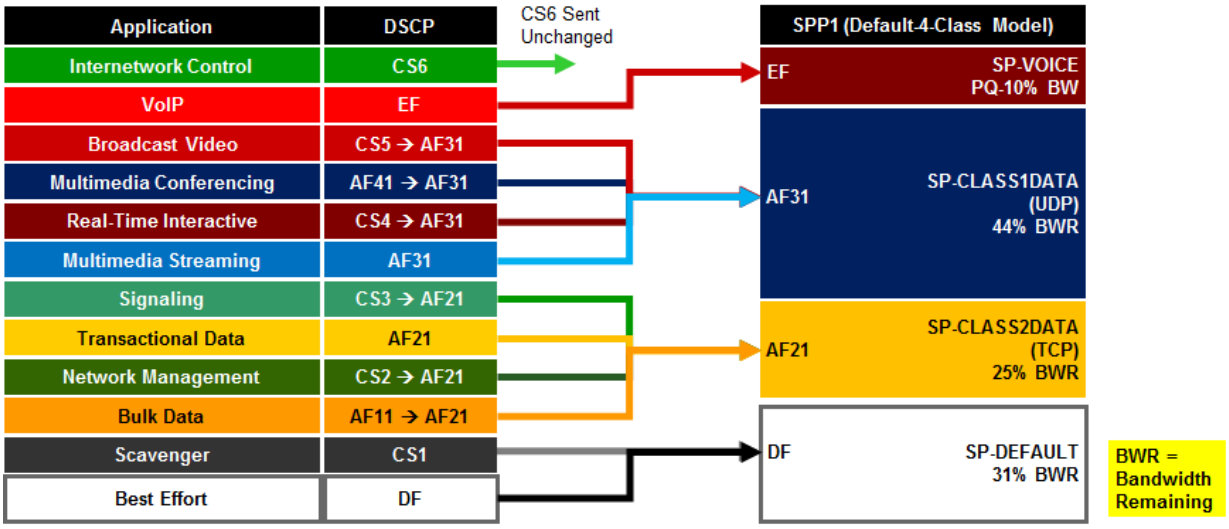

The following is an example of the hierarchical policy-map definition pushed by EasyQoS that implements the SPP1/SPP:SPP1-4Class queuing policy with the bandwidths and traffic re-marking for each of the service provider traffic classes. It assumes a sub-line rate of 50 Mbps to the service provider network.

```
!
policy-map prm-DSCP#QUEUING_O_1#shape#50.0
 class class-default
  shape average 50000000
   service-policy prm-DSCP#QUEUING_O_1


policy-map prm-DSCP#QUEUING_O_1
 class prm-EZQOS_12C#VOICE
  police rate percent 10
  priority
  set DSCP ef
 class prm-EZQOS_12C#BROADCAST
  bandwidth remaining percent 8
  set DSCP af31
 class prm-EZQOS_12C#REALTIME
  bandwidth remaining percent 11
  set DSCP af31
 class prm-EZQOS_12C#MM_CONF
```

```
 bandwidth remaining percent 12
 fair-queue
 set DSCP af31
 random-detect DSCP-based
class prm-EZQOS_12C#MM_STREAM
 bandwidth remaining percent 12
 fair-queue
 set DSCP af31
 random-detect DSCP-based
class prm-EZQOS_12C#CONTROL
 bandwidth remaining percent 3
class prm-EZQOS_12C#SIGNALING
 bandwidth remaining percent 3
 set DSCP af21
class prm-EZQOS_12C#OAM
 bandwidth remaining percent 4
 set DSCP af21
class prm-EZQOS_12C#TRANS_DATA
 bandwidth remaining percent 12
 fair-queue
 set DSCP af21
 random-detect DSCP-based
class prm-EZQOS_12C#BULK_DATA
 bandwidth remaining percent 5
 fair-queue
 set DSCP af21
 random-detect DSCP-based
class prm-EZQOS_12C#SCAVENGER
 bandwidth remaining percent 1
 set DSCP default
class class-default
 bandwidth remaining percent 29
```

```
  fair-queue
  set DSCP default
  random-detect DSCP-based
  random-detect DSCP 0 50 64        ! ISR G2 Series platforms only.
!
```

The names of the parent and child policy-maps reflect the SPP configured for the WAN interface. When using the older method where the service provider profile is indicated via the #SPP1#" tag, the format of the parent and child policy-maps will be as follows:

```
!
policy-map prm-DSCP#QUEUING_O_1#shape#50.0
policy-map prm-DSCP#QUEUING_O_1
!
```

This is the format shown in the configuration example above.

When using the newer method where the service provider profile is indicated via the #SPP:SPP1-4Class# tag, the format of the parent and child policy-maps will be as follows:

```
!
policy-map prm-DSCP#QUEUING_O_SPP1-4Class#shape#50.0
policy-map prm-DSCP#QUEUING_O_SPP1-4Class
!
```

If a sub-line rate service has been provisioned, the top-level of the SPP1/SPP:SPP1-4Class hierarchical egress queuing policy-map simply implements shaping to an average rate that matches the sub-line bandwidth rate of the managed-service offering provisioned by the service provider. This rate is learned via the #rate# field within the tag, which must be pre-configured within the description of the interface connected to the managed service WAN link.

Note: If a sub-line rate service has not been provisioned, EasyQoS will not configure a hierarchical policy-map with a shaper at the parent-level. Instead, the policy-map will only have a single level with the configuration similar to the child-policy discussed below.

The child-policy of the SPP1/SPP:SPP1-4Class egress queuing policy-map implements a single LLQ policy, meaning a separate LLQ for the Voice traffic class. An explicit policer (10% of bandwidth) for the Voice queue ensures that the LLQ can use no more than the percentage of the bandwidth of the WAN link allocated to the Voice traffic class, regardless of whether there is available bandwidth.

The remaining eleven queues share the remaining bandwidth based on a percentage allocation of bandwidth. This is accomplished via the "bandwidth remaining percent" command. Each of these queues can use more than its percentage allocation, if more bandwidth is available—meaning if one or more of the other queues is not using its full allocation of remaining bandwidth percentage.

The following traffic is admitted (mapped) to the service provider SP-Voice traffic class. Traffic mapped to this service provider traffic class is remarked to EF.

- Traffic exiting the Voice queue

The bandwidth allocated to the Voice queue (10% priority and policed) is meant to match the 10% bandwidth allocation of the service provider SP-Voice traffic class as shown in Figure 44.

The following traffic is admitted (mapped) to the service provider SP-Class1Data traffic class. Traffic mapped to this service provider traffic class is remarked to AF31.

- Traffic exiting the Broadcast-Video queue re-marked from CS5
- Traffic exiting the Realtime-Interactive queue re-marked from CS4
- Traffic exiting the Multimedia-Conferencing queue re-marked from AF4x
- Traffic exiting the Multimedia-Streaming queue re-marked from AF3x

The sum of the bandwidths allocated to four queues–Broadcast-Video (8% bandwidth remaining), Realtime-Interactive (11% bandwidth remaining), Multimedia-Conferencing (12% bandwidth remaining), and Multimedia-Streaming (12% bandwidth remaining)–is meant to roughly match the 44% bandwidth remaining allocation of the service provider SP-Class1Data traffic class as shown in Figure 44.

The following traffic is admitted (mapped) to the service provider SP-Class2Data traffic class. Traffic mapped to this service provider traffic class is remarked to AF21.

- Traffic exiting the Signaling queue re-marked from AF3x
- Traffic exiting the OAM queue remarked from CS2
- Traffic exiting the Transactional-Data queue marked from AF2x
- Traffic exiting the Bulk-Data queue re-marked from AF1x

The sum of the bandwidths allocated to the four queues–Signaling (3% bandwidth), OAM (4% bandwidth), Transactional-Data (12% bandwidth remaining), and Bulk-Data (5% bandwidth remaining)–is meant to roughly match the 25% bandwidth remaining allocation of the service provider SP-Class2Data traffic class, as shown in Figure 44.

The following traffic is admitted (mapped) to the service provider SP-Default traffic class. Traffic mapped to this service provider traffic class is remarked to Default (Best Effort).

- Traffic exiting the Scavenger queue is re-marked from CS1
- Traffic exiting the Default queue

The sum of the bandwidth allocated to two queues–Scavenger (1% bandwidth remaining) and Default (29% bandwidth remaining)–is meant to roughly match the default 31% bandwidth remaining allocation of the service provider SP-Default traffic class, as shown in Figure 44.

Fair-queuing, along with DSCP-based WRED is implemented for the following queues:

- Multimedia-Conferencing

- Multimedia-Streaming
- Transactional-Data
- Bulk-Data
- Default

For ISR 3900, 2900, and 800 Series (ISR G2) platforms only, with the exception of the Default queue, minimum and maximum WRED thresholds for the queues are left at their default values. Table 2 summarized these minimum and maximum thresholds. The default WRED thresholds for the Default queue are considered to be too aggressive—meaning the minimum drop threshold is set lower than desired. Hence, the minimum drop threshold has been adjusted to 50 packets, and the maximum drop threshold adjusted to the depth of the queue—64 packets. For ISR 4400 and ASR 1000 Series platforms the minimum and maximum WRED thresholds for the queues are left at their default values.

Traffic within the Control queue is sent unchanged to the service provider network and is not considered to be mapped into one of the four service provider traffic classes.

The SPP1/SPP:SPP1-4Class egress queuing policy is applied to WAN interfaces that include the #WAN#rate#SPP1# or #WAN#rate#SPP:4-Default-Class-Model# tag within the interface description.

An example of the application of the egress queuing policy is as follows:

```
!
interface GigabitEthernet0/0
 description CIRCUIT TO WE-ASR2 GIG-0-0-1 #WAN#50M#SPP1#
 service-policy output prm-DSCP#QUEUING_O_1#shape#50.0
!
```

SPP2/SPP:SPP2-5Class

The SPP2/SPP:SPP2-5Class is based on managed-service offerings with five traffic classes. These traffic classes are specified as follows within this document:

- SP-Voice
- SP-Class1Data
- SP-Class2Data
- SP-Class3Data
- SP-Default

The following figure shows the WAN bandwidth allocation model for the SPP2/SPP:SPP2-5Class service provider profile.

Figure 46 **WAN Bandwidth Allocation for the SPP2/SPP:SPP2-5Class**

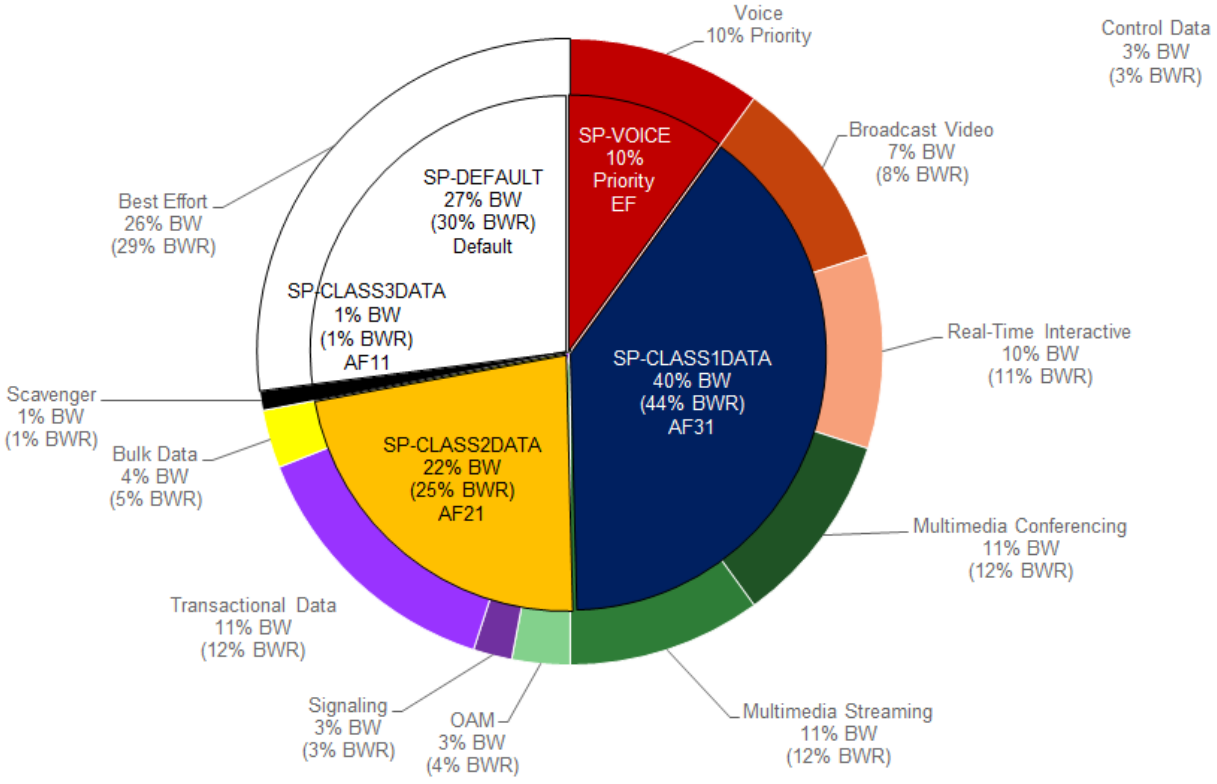

For the SPP2/SPP:SPP2-5Class, EasyQoS must map the RFC 4594-based 12-class QoS model implemented within the organization into the five traffic classes provide by the service provider. The following figure shows this mapping with the bandwidth allocations and traffic re-marking for the service provider traffic classes.

Figure 47 **EasyQoS Marking Mappings into SPP2/SPP:SPP2-5Class**

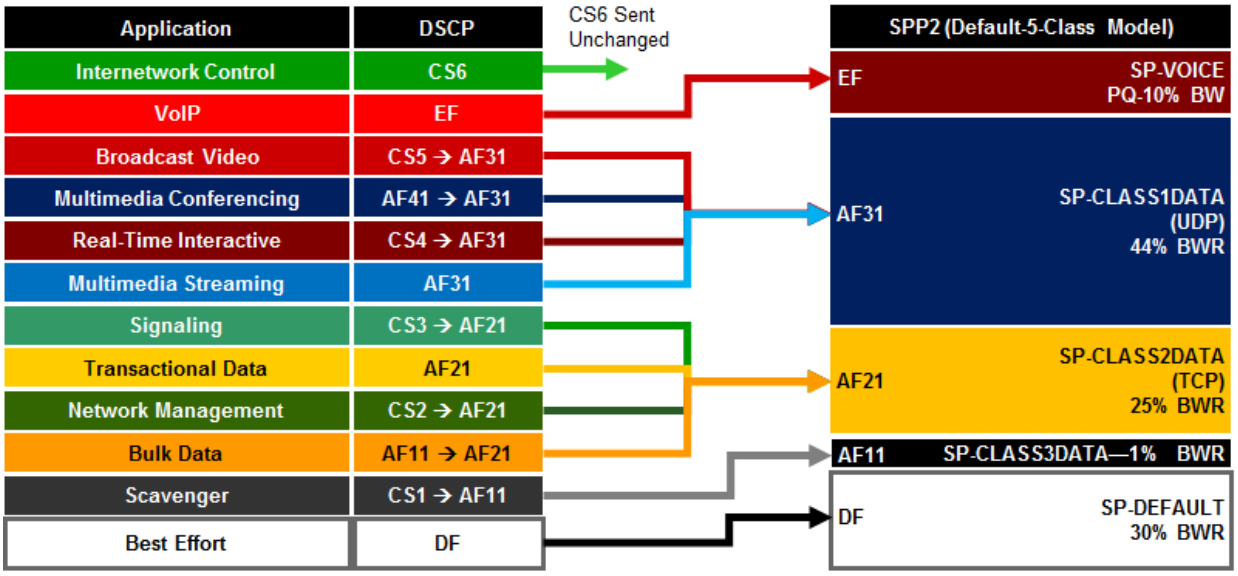

The following is an example of the hierarchical policy-map definition pushed by EasyQoS that implements the SPP2/SPP:SPP2-5Class queuing policy with the bandwidths and traffic re-marking for each of the service provider traffic classes. It assumes a sub-line rate of 50 Mbps to the service provider network.

```
!
policy-map prm-DSCP#QUEUING_O_2#shape#50.0
 class class-default
  shape average 50000000
   service-policy prm-DSCP#QUEUING_O_2


policy-map prm-DSCP#QUEUING_O_2
 class prm-EZQOS_12C#VOICE
  police rate percent 10
  priority
  set DSCP ef
 class prm-EZQOS_12C#BROADCAST
  bandwidth remaining percent 8
  set DSCP af31
 class prm-EZQOS_12C#REALTIME
  bandwidth remaining percent 11
  set DSCP af31
 class prm-EZQOS_12C#MM_CONF
  bandwidth remaining percent 12
  fair-queue
  set DSCP af31
  random-detect DSCP-based
 class prm-EZQOS_12C#MM_STREAM
  bandwidth remaining percent 12
  fair-queue
  set DSCP af31
  random-detect DSCP-based
 class prm-EZQOS_12C#CONTROL
  bandwidth remaining percent 3
 class prm-EZQOS_12C#SIGNALING
```

```
  bandwidth remaining percent 3
  set DSCP af21
 class prm-EZQOS_12C#OAM
  bandwidth remaining percent 4
  set DSCP af21
 class prm-EZQOS_12C#TRANS_DATA
  bandwidth remaining percent 12
  fair-queue
  set DSCP af21
  random-detect DSCP-based
 class prm-EZQOS_12C#BULK_DATA
  bandwidth remaining percent 5
  fair-queue
  set DSCP af21
  random-detect DSCP-based
 class prm-EZQOS_12C#SCAVENGER
  bandwidth remaining percent 1
  set DSCP af11
 class class-default
  bandwidth remaining percent 29
  fair-queue
  set DSCP default
  random-detect DSCP-based
  random-detect DSCP 0 50 64        ! ISR G2 Series platforms only.
!
```

Again, the names of the parent and child policy-maps reflect the SPP configured for the WAN interface. When using the older method where the service provider profile is indicated via the #SPP2#" tag, the format of the parent and child policy-maps will be as follows:

```
!
policy-map prm-DSCP#QUEUING_O_2#shape#50.0
policy-map prm-DSCP#QUEUING_O_2
!
```

This is the format shown in the configuration example above.

When using the newer method where the service provider profile is indicated via the #SPP:SPP2-5Class# tag, the format of the parent and child policy-maps will be as follows:

```
!
policy-map prm-DSCP#QUEUING_O_SPP2-5Class#shape#50.0
policy-map prm-DSCP#QUEUING_O_SPP2-5Class
!
```

The difference between the SPP1 and SPP2 models is that a 5th service provider traffic class–SP-Class3Data–is provisioned specifically for handling traffic with a lower than best-effort treatment (Scavenger traffic).

The policy-map admits and re-marks traffic exiting the Scavenger queue from CS1 to AF11, corresponding to the service provider SP-Class3Data traffic class. The bandwidth allocated to the Scavenger queue (1% bandwidth remaining) is meant to match the 1% bandwidth remaining allocated to the service provider SP-Class3Data traffic class, as shown in Figure 46.

The SPP2/SPP:SPP2-5Class egress queuing policy is applied to WAN interfaces that include the #WAN#rate#SPP2# or #WAN#rate#SPP:SPP2-5Class# tag within the interface description.

An example of the application of the egress queuing policy is as follows:

```
!
interface GigabitEthernet0/0
 description CIRCUIT TO WE-ASR2 GIG-0-0-1 #WAN#50M#SPP2#
 service-policy output prm-DSCP#QUEUING_O_2#shape#50.0
!
```

SPP3/SPP:SPP3-6Class

The SPP3/SPP:SPP3-6Class is based on managed-service offerings with six traffic classes. These traffic classes are specified as follows within this document:

- SP-Voice
- SP-Video
- SP-Class1Data
- SP-Class2Data
- SP-Class3Data
- SP-Default

The following figure shows the WAN bandwidth allocation model for the SPP3/SPP:SPP3-6Class service provider profile.

Figure 48 **WAN Bandwidth Allocation for the SPP3/SPP:SPP3-6Class**

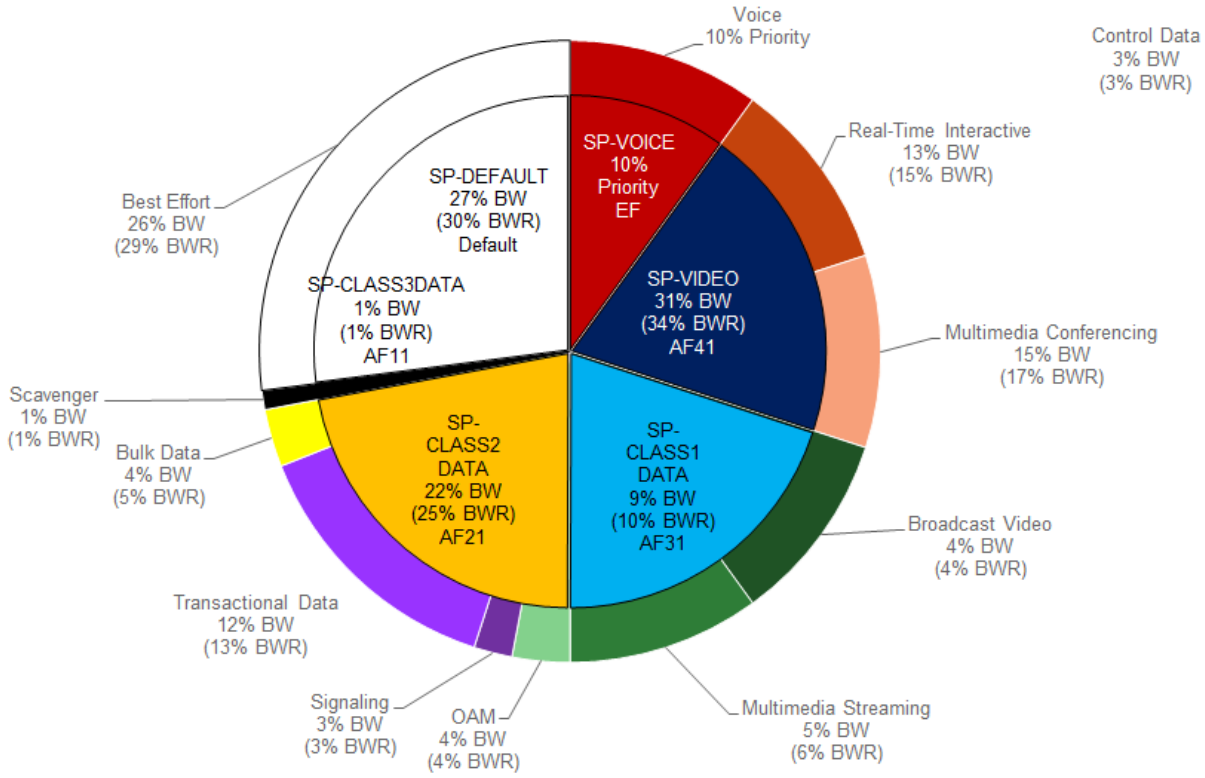

For the SPP3/SPP:SPP3-6Class, EasyQoS must map the RFC 4594-based 12-class QoS model implemented within the organization into the six traffic classes provide by the service provider. The following figure shows this mapping with the bandwidth allocations and traffic re-marking for the service provider traffic classes.

Figure 49 **EasyQoS marking Mappings Into SPP3/SPP:SPP3-6Class**

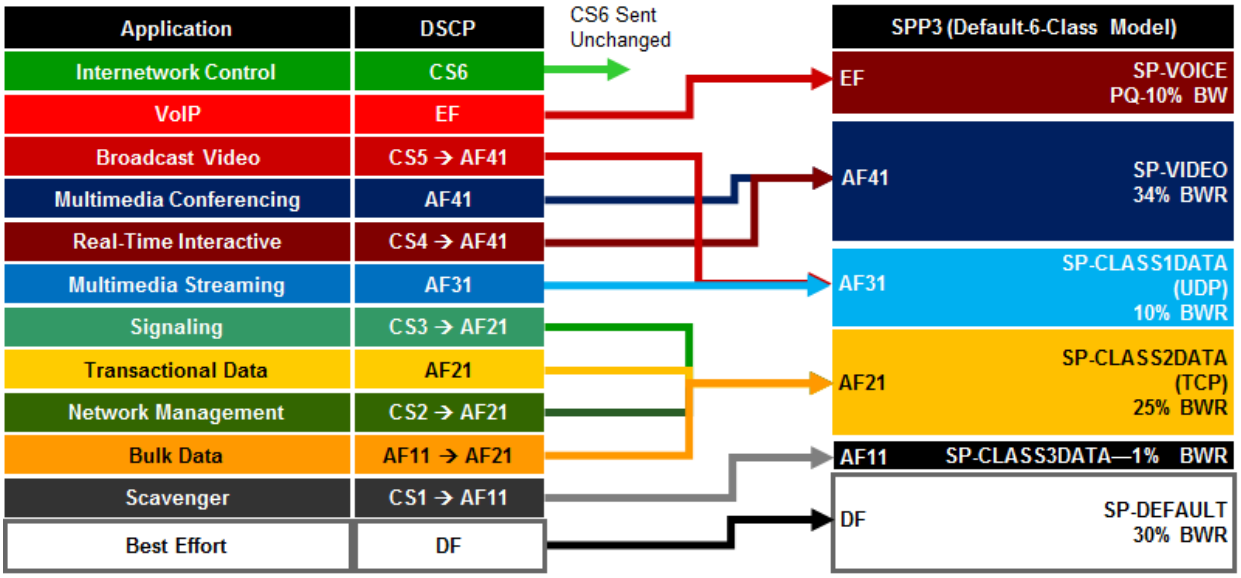

The following is an example of the hierarchical policy-map definition pushed by EasyQoS that implements the SPP3/SPP:SPP3-6Class queuing policy, with the bandwidths and traffic re-marking for each of the service provider traffic classes. It assumes a sub-line rate of 50 Mbps to the service provider network.

```
!
policy-map prm-DSCP#QUEUING_O_3#shape#50.0
 class class-default
  shape average 50000000
   service-policy prm-DSCP#QUEUING_O_3


policy-map prm-DSCP#QUEUING_O_3
 class prm-EZQOS_12C#VOICE
  police rate percent 10
  priority
  set DSCP ef
 class prm-EZQOS_12C#BROADCAST
  bandwidth remaining percent 4
  set DSCP af31
 class prm-EZQOS_12C#REALTIME
  bandwidth remaining percent 15
  set DSCP af41
 class prm-EZQOS_12C#MM_CONF
  bandwidth remaining percent 17
  fair-queue
  set DSCP af41
  random-detect DSCP-based
 class prm-EZQOS_12C#MM_STREAM
  bandwidth remaining percent 6
  fair-queue
  set DSCP af31
  random-detect DSCP-based
 class prm-EZQOS_12C#CONTROL
  bandwidth remaining percent 3
 class prm-EZQOS_12C#SIGNALING
```

```
  bandwidth remaining percent 3
  set DSCP af21
 class prm-EZQOS_12C#OAM
  bandwidth remaining percent 4
  set DSCP af21
 class prm-EZQOS_12C#TRANS_DATA
  bandwidth remaining percent 13
  fair-queue
  set DSCP af21
  random-detect DSCP-based
 class prm-EZQOS_12C#BULK_DATA
  bandwidth remaining percent 5
  fair-queue
  set DSCP af21
  random-detect DSCP-based
 class prm-EZQOS_12C#SCAVENGER
  bandwidth remaining percent 1
  set DSCP af11
 class class-default
  bandwidth remaining percent 29
  fair-queue
  set DSCP default
  random-detect DSCP-based
  random-detect DSCP 0 50 64        ! ISR G2 Series platforms only.
!
```

The names of the parent and child policy-maps reflect the SPP configured for the WAN interface. When using the older method where the service provider profile is indicated via the #SPP3#" tag, the format of the parent and child policy-maps will be as follows:

```
!
policy-map prm-DSCP#QUEUING_O_3#shape#50.0
policy-map prm-DSCP#QUEUING_O_3
!
```

This is the format shown in the configuration example above.

When using the newer method where the service provider profile is indicated via the #SPP:SPP3-6Class# tag, the format of the parent and child policy-maps will be as follows:

```
!
policy-map prm-DSCP#QUEUING_O_SPP3-6Class#shape#50.0
policy-map prm-DSCP#QUEUING_O_SPP3-6Class
!
```

The difference between the SPP2 and SPP3 models is that a 6th service provider traffic class–SP-Video–is provisioned specifically for handling video traffic.

The following traffic is admitted (mapped) to the service provider SP-Video traffic class. Traffic mapped to this service provider traffic class is remarked to AF41.

- Traffic exiting the Realtime-Interactive queue is re-marked from CS4
- Traffic exiting the Multimedia-Conferencing re-marked from AF4x

The sum of the bandwidth allocated to two queues–Realtime-Interactive (15% bandwidth remaining) and Multimedia-Conferencing (17% bandwidth remaining)–is meant to roughly match the 34% bandwidth remaining allocation of the service provider SP-Class1Data traffic class as shown in Figure 48.

The mappings are also adjusted so that the following is admitted (mapped) to the service provider SP-Class1Data traffic class. Traffic mapped to this service provider traffic class is remarked to AF31.

- Traffic exiting the Broadcast-Video queue is re-marked from CS5
- Traffic exiting the Multimedia-Streaming queue is re-marked from AF3x

The sum of the bandwidth allocated to two queues–Broadcast-Video (4% bandwidth remaining) and Multimedia-Streaming (6% bandwidth remaining)–is meant to roughly match the 10% bandwidth remaining allocation of the service provider SP-Streaming-Video traffic class as shown in Figure 48.

The SPP3/SPP:SPP3-6Class egress queuing policy is applied to WAN interfaces that include the #WAN#rate#SPP3# or #WAN#rate#SPP:SPP3-6Class# tag within the interface description.

An example of the application of the egress queuing policy is as follows:

```
!
interface GigabitEthernet0/0
 description CIRCUIT TO WE-ASR2 GIG-0-0-1 #WAN#50M#SPP3#
 service-policy output prm-DSCP#QUEUING_O_3#shape#50.0
!
```

SPP4/SPP:SPP4-8Class

The SPP4/SPP:SPP4-8Class is based on managed-service offerings with eight traffic classes. These traffic classes are specified as follows within this document:

- SP-Voice
- SP-Interactive-Video
- SP-Streaming-Video
- SP-Net-Ctrl-Mgmt
- SP-Call-Sig
- SP-Critical-Data
- SP-Scavenger
- SP-Default

The following figure shows the WAN bandwidth allocation model for the SPP4/SPP:SPP4-8Class service provider profile.

Figure 50 **WAN Bandwidth Allocation for the SPP4/SPP:SPP4-8Class**

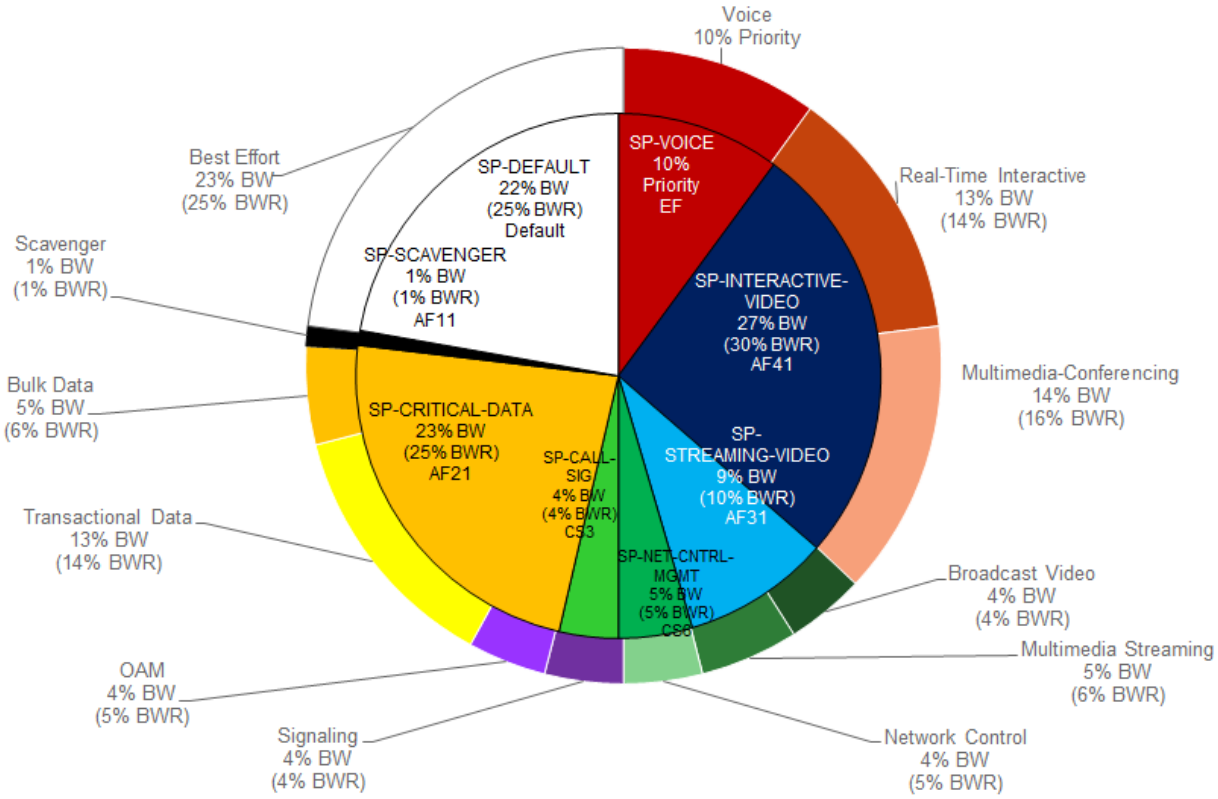

For the SPP4/SPP:SPP4-8Class, EasyQoS must map the RFC 4594-based 12-class QoS model implemented within the organization into the eight traffic classes provide by the service provider. The

following figure shows this mapping with the bandwidth allocations and traffic re-marking for the service provider traffic classes.

Figure 51 EasyQoS Marking Mappings into SPP4

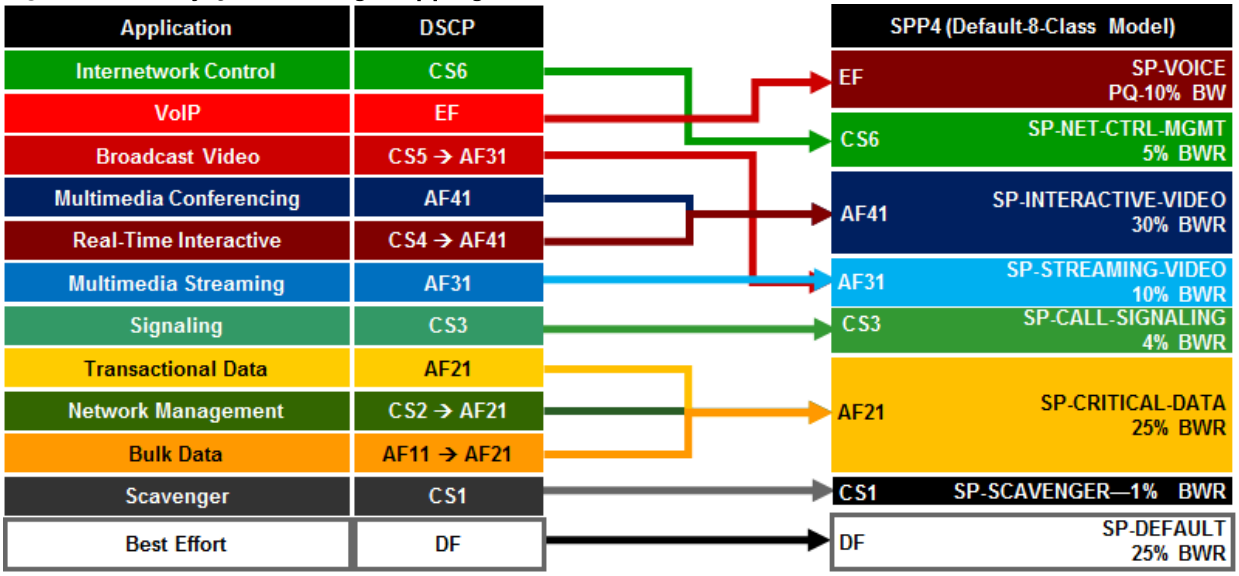

The following is an example of the hierarchical policy-map definition pushed by EasyQoS that implements the SPP4/SPP:SPP4-8Class queuing policy with the bandwidths and traffic re-marking for each of the service provider traffic classes. It assumes a sub-line rate of 50 Mbps to the service provider network.

```
!
policy-map prm-DSCP#QUEUING_O_4#shape#50.0
 class class-default
  shape average 50000000
   service-policy prm-DSCP#QUEUING_O_4


policy-map prm-DSCP#QUEUING_O_4
 class prm-EZQOS_12C#VOICE
  police rate percent 10
  priority
  set DSCP ef
 class prm-EZQOS_12C#BROADCAST
  bandwidth remaining percent 4
  set DSCP af31
 class prm-EZQOS_12C#REALTIME
  bandwidth remaining percent 14
```

```
 set DSCP af41
class prm-EZQOS_12C#MM_CONF
 bandwidth remaining percent 16
 fair-queue
 set DSCP af41
 random-detect DSCP-based
class prm-EZQOS_12C#MM_STREAM
 bandwidth remaining percent 6
 fair-queue
 set DSCP af31
 random-detect DSCP-based
class prm-EZQOS_12C#CONTROL
 bandwidth remaining percent 5
 set DSCP cs6
class prm-EZQOS_12C#SIGNALING
 bandwidth remaining percent 4
 set DSCP cs3
class prm-EZQOS_12C#OAM
 bandwidth remaining percent 5
 set DSCP af21
class prm-EZQOS_12C#TRANS_DATA
 bandwidth remaining percent 14
 fair-queue
 set DSCP af21
 random-detect DSCP-based
class prm-EZQOS_12C#BULK_DATA
 bandwidth remaining percent 6
 fair-queue
 set DSCP af21
 random-detect DSCP-based
class prm-EZQOS_12C#SCAVENGER
 bandwidth remaining percent 1
```

```
  set DSCP cs1
 class class-default
  bandwidth remaining percent 25
  fair-queue
  set DSCP default
  random-detect DSCP-based
  random-detect DSCP 0 50 64        ! ISR G2 Series platforms only.
!
```

As with the previous service provider profile models, the names of the parent and child policy-maps reflect the SPP configured for the WAN interface. When using the older method where the service provider profile is indicated via the #SPP4#" tag, the format of the parent and child policy-maps will be as follows:

```
!
policy-map prm-DSCP#QUEUING_O_4#shape#50.0
policy-map prm-DSCP#QUEUING_O_4
!
```

This is the format shown in the configuration example above.

When using the newer method where the service provider profile is indicated via the #SPP:SPP4-8Class# tag, the format of the parent and child policy-maps will be as follows:

```
!
policy-map prm-DSCP#QUEUING_O_SPP4-8Class#shape#50.0
policy-map prm-DSCP#QUEUING_O_SPP4-8Class
!
```

If a sub-line rate service has been provisioned, the top-level of the SPP4/SPP:SPP4-8Class hierarchical egress queuing policy-map simply implements shaping to an average rate that matches the sub-line bandwidth rate of the managed-service offering provisioned by the service provider. This rate is learned via the #rate# field within the tag, which must be pre-configured within the description of the interface connected to the managed service WAN link.

The child-policy of the SPP4/SPP:SPP4-8Class egress queuing policy-map implements a single LLQ policy, meaning a separate LLQ for the Voice traffic class. An explicit policer (10% of the bandwidth) for the Voice queue ensures that the LLQ can use no more than the percentage of the bandwidth of the WAN link allocated to the traffic class, regardless of whether there is available bandwidth.

The remaining eleven queues share the remaining bandwidth based on a percentage allocation of bandwidth. This is accomplished via the "bandwidth remaining percent" command. Each of these queues can use more than its percentage allocation, if more bandwidth is available—meaning if one or more of the other queues is not using its full allocation of remaining bandwidth percentage.

The following traffic is admitted (mapped) to the service provider SP-Voice traffic class. Traffic mapped to this service provider traffic class is remarked to EF.

- Traffic exiting the Voice queue

The bandwidth allocated to the Voice queue (10% priority and policed) is meant to match the 10% bandwidth remaining allocation of the service provider SP-Voice traffic class as shown in Figure 50.

The following traffic is admitted (mapped) to the service provider SP-Net-Ctrl-Mgmt traffic class. Traffic mapped to this service provider traffic class is remarked to CS6.

- Traffic exiting the Control queue

The bandwidth allocated to the Control queue (5% bandwidth remaining) is meant to roughly match the 5% bandwidth remaining allocation of the service provider SP-Net-Ctrl-Mgmt traffic class as shown in Figure 50.

The following traffic is admitted (mapped) to the service provider SP-Interactive-Video traffic class. Traffic mapped to this service provider traffic class is remarked to AF41

- Traffic exiting the Realtime-Interactive queue re-marked from CS4
- Traffic exiting the Multimedia-Conferencing queue re-marked from AF4x

The sum of the bandwidth allocated to two queues—Realtime-Interactive (14% bandwidth remaining) and Multimedia-Conferencing (16% bandwidth remaining)—is meant to roughly match the 30% bandwidth remaining allocation of the service provider SP-Class1Data traffic class as shown in Figure 50.

The following traffic is admitted (mapped) to the service provider SP-Streaming-Video traffic class. Traffic mapped to this service provider traffic class is remarked to AF31

- Traffic exiting the Broadcast-Video queue re-marked from CS5
- Traffic exiting the Multimedia-Streaming queue re-marked from AF3x

The sum of the bandwidth allocated to two queues—Broadcast-Video (4% bandwidth remaining) and Multimedia-Streaming (6% bandwidth remaining)—is meant to roughly match the default 10% bandwidth remaining allocation of the service provider SP-Streaming-Video traffic class as shown in Figure 50.

The following traffic is admitted (mapped) to the service provider SP-Call-Signaling traffic class. By default traffic mapped to this service provider traffic class is remarked to CS3

- Traffic exiting the Signaling queue

The bandwidth allocated to the Signaling queue (4% bandwidth remaining) is meant to roughly match the 4% bandwidth remaining allocation of the service provider SP-Call-Signaling traffic class as shown in Figure 50.

The following traffic is admitted (mapped) to the service provider SP-Critical-Data traffic class. Traffic mapped to this service provider traffic class is remarked to AF21

- Traffic exiting the OAM queue re-marked from CS2
- Traffic exiting the Transactional-Data queue re-marked from AF2x

- Traffic exiting the Bulk-Data queue re-marked from AF1x

The sum of the bandwidth allocated to the three queues–OAM (5% bandwidth remaining), Transactional-Data (14% bandwidth remaining), and Bulk-Data (6% bandwidth remaining)–is meant to roughly match the 25% bandwidth remaining allocation of the service provider SP-Critical-Data traffic class, as shown in Figure 50.

The following traffic is admitted (mapped) to the service provider SP-Scavenger traffic class. Traffic mapped to this service provider traffic class is remarked to CS1.

- Traffic exiting the Scavenger queue

The bandwidth allocated to the Scavenger queue (1% bandwidth remaining) is meant to roughly match the 1% bandwidth remaining allocation of the service provider SP-Scavenger traffic class as shown in Figure 50.

The following traffic is admitted (mapped) to the service provider Default traffic class:

- Traffic exiting the Default queue

The bandwidth allocated to the Default queue (25% bandwidth remaining) is meant to roughly match the 25% bandwidth remaining allocation of the service provider SP-Default traffic class as shown in Figure 50.

Fair-queuing, along with DSCP-based WRED is implemented for the following queues:

- Multimedia-Conferencing
- Multimedia-Streaming
- Transactional-Data
- Bulk-Data
- Default

For ISR 3900, 2900, and 800 Series (ISR G2) platforms only, with the exception of the Default queue, minimum and maximum WRED thresholds for the queues are left at their default values. Table 2 summarized these minimum and maximum thresholds. The default WRED thresholds for the Default queue are considered to be too aggressive–meaning the minimum drop threshold is set lower than desired. Hence, the minimum drop threshold has been adjusted to 50 packets, and the maximum drop threshold adjusted to the depth of the queue–64 packets. For ISR 4400 Series platforms the minimum and maximum WRED thresholds for the queues are left at their default values.

The SPP4/SPP:SPP4-8Class egress queuing policy is applied to WAN interfaces that include the #WAN#rate#SPP4# or #WAN#rate#SPP:SPP4-8Class# tag within the interface description.

An example of the application of the egress queuing policy is as follows:

```
!
interface GigabitEthernet0/0
 description CIRCUIT TO WE-ASR2 GIG-0-0-1 #WAN#50M#SPP4#
 service-policy output prm-DSCP#QUEUING_O_4#shape#50.0
```

!

Custom Service Provider Profiles

Configuration of custom service provider profiles is discussed in the *APIC-EM and the EasyQoS Application* section of this document. Custom service provider profiles use one of the four default service provider profiles discussed in the sections above as a template for the custom profile. Hence the basic structure of the egress queuing policy is the same as discussed in the sections above.

The mapping of the internal queues to the service provider traffic classes is fixed, depending upon which of the four default service provider profiles has been selected as the template upon which to base the custom service provider profile. In other words, the internal traffic classes that are admitted to each service provider traffic class is fixed based upon the 4, 5, 6, or 8-class template chosen for the custom service provider profile. However, the DSCP value to which the internal traffic classes are re-marked as they enter the service provider network can be specified by the network administrator when configuring the custom service provider profile.

The percentage of bandwidth allocated to the service provider traffic classes can also be specified by the network administrator when configuring the custom service provider profile. If a sub-line rate service has been provisioned and the interface includes a #rate# tag within the description, then a hierarchical policy-map will be configured by EasyQoS, with a shaper matching the sub-line rate configured at the parent level.

The child-policy of the egress queuing policy-map will still implement a single LLQ policy, meaning a separate LLQ for only the Voice traffic class. An explicit policer for the Voice queue ensures that the LLQ can use no more than the percentage of the bandwidth of the WAN link allocated to the Voice traffic class, regardless of whether there is available bandwidth. For a custom service provider profile, the percentage of bandwidth allocated to the policer is directly dependent upon the amount of bandwidth allocated to the service provider Voice traffic class.

The remaining queues still share the remaining bandwidth based on a percentage allocation of remaining bandwidth. This is accomplished via the "bandwidth remaining percent" command. For the custom service provider profile, the percentage of remaining bandwidth allocated to each queue is dependent upon the amount of remaining bandwidth allocated to the service provider traffic class to which the queue is mapped. EasyQoS will automatically divide the bandwidth specified for the service provider traffic class among the among the various traffic classes of the organization that map to the particular service provider traffic class. Again, the amount of remaining bandwidth allocated to the service provider traffic class is specified by the network administrator when configuring the custom service provider profile. Each of these queues can use more than its percentage allocation, if more bandwidth is available–meaning if one or more of the other queues is not using its full allocation of remaining bandwidth percentage.

Server IP/Port-Based Custom Applications and Managed Service WANs

As discussed in the *Server IP/Port Based Applications* section of this document, Custom applications that are defined based on a server IP address and/or ports are added to the prm-MARKING_IN#TUNNELED-NBAR traffic-class in ASR and ISR router platforms by EasyQoS. No action is specified for the prm-MARKING_IN#TUNNELED-NBAR traffic class within the ingress classification & marking policy on ASR and ISR router platforms. Effectively, Custom applications based on server IP addresses and/or ports do not pass through the NBAR engine within the ASR or ISR router platform. When traffic from such a Custom application exits the organization's network, entering a service provider managed-service network, the

Custom application traffic may be remarked to match the allowed DSCP markings for the service provider traffic class to which the Custom application is mapped. However, when the Custom application traffic re-enters the organization's network at the other end of service provider manage-service network, the traffic will not be re-marked back to its original DSCP marking. Instead, Custom applications based on server IP addresses and/or ports will retain the DSCP markings to which they were mapped when entering the service provider managed-service network. Future version of APIC-EM EasyQoS may modify this behavior.

Campus LAN Static QoS Design

Within the EasyQoS solution, different network devices implement the QoS policy to the best of their abilities. Within the APIC-EM 1.3 EasyQoS solution, Catalyst switches implement one or more of the following QoS policies, depending upon their role within the network infrastructure:

- Ingress classification & marking policies on wired ports based on policy-maps that contain ACLs with Layer 2-4 ACEs.
- Ingress and/or egress queuing policies

The role of the Catalyst switch within the network infrastructure is discussed in detail in the *Catalyst Switch Roles* section of this document.

Note: Catalyst 3850 and 3650 Series switches support wired AVC/NBAR-based ingress classification & marking policies with IOS XE 16.3.1 and higher software versions. The APIC-EM 1.3 release of EasyQoS does not use AVC/NBAR-based ingress classification & marking policies. Future versions of EasyQoS may add this support.

TCAM Utilization Challenges

Cisco Catalyst switch platforms implement ingress classification & marking policies that include ACL entries in hardware for performance reasons. This hardware is commonly referred to as Ternary Content Addressable Memory (TCAM). Most Catalyst switch platforms have sufficient QoS TCAM space, such that the ingress classification & marking QoS policies provisioned by EasyQoS do not exceed the available QoS TCAM space. However, a few older Catalyst switch platforms supported by EasyQoS do have limited QoS TCAM space. Because QoS TCAM space is limited on these older Catalyst switch platforms, the ACEs that are provisioned by the EasyQoS ingress classification & marking policies across various platforms can vary slightly. The EasyQoS application is aware of the QoS TCAM space available in supported switch platforms; and will implement ACEs up to the limits of the TCAM-constrained platforms–leaving sufficient TCAM space for additional functionality such as Dynamic QoS and statically learned Cisco devices.

In order to minimize the impact of limited TCAM space on these older platforms, APIC-EM deploys Custom and Favorite applications first. Custom applications are by default marked as Favorite applications, as well. This ensures that the applications most relevant to the network operator are included within the ACEs deployed by EasyQoS.

The following table summarizes the QoS TCAM space for various Catalyst access-layer switch platforms supported by EasyQoS.

Table 3 QoS TCAM Space per ASIC for Various Catalyst Access-Layer Switches

Switch Platform	QoS TCAM Entries
Catalyst 2960-X Series	504 with the lanbase-default SDM template 384 with the default and lanbase-routing SDM Templates

Catalyst 2960-XR Series	512 with all SDM templates
Catalyst 2960-S Series	384 with all SDM templates
Catalyst 3560-C Series	384 with the default SDM template
Catalyst 3560-X Series	512 with all SDM templates other than the routing template 384 with the routing SDM template
Catalyst 3750-X Series	512 with all SDM templates other than the routing template 384 with the routing template
Catalyst 3850 Series	2816 with the Advanced (high scale) SDM template 3072 with the VLAN SDM template
Catalyst 3650 Series	3072 with both the Advanced (low scale) and the VLAN (high scale) SDM templates

Due to the algorithm for converting ACEs into QoS TCAM masks and entries, there is not necessarily a one-to-one correlation between a single ACE within an ACL and a single TCAM entry. However, a general guideline that can be used is that EasyQoS will require a single TCAM entry for each ACE. Put more simply, if a particular platform has room for 384 QoS TCAM entries (ignoring masks), then in general it can support approximately 384 ACEs within the ACLs of the class-maps within the ingress classification & marking policy, minus any TCAM entries reserved for the platform itself. However, it should be noted that individual applications may be identified via multiple TCP and UDP ports or port ranges. Bi-directionality will also double the ACE entries. Hence each application may result in multiple ACEs within the ingress classification & marking policy. Therefore, those platforms with limited QoS TCAM size have limited ability to support applications specified within the policy created by EasyQoS.

Methodology and Workflow

The general method by which the applications selected within the EasyQoS web-based GUI are translated into ACLs with ACE entries is discussed in the following sections.

CAPWAP Control and Data Traffic

EasyQoS will check to see if there is QoS TCAM space available for additional ACEs within any of the ACLs that are part of the ingress classification & marking QoS policy-map. If so, EasyQoS will generate and deploy ACEs for the prm-APIC_QOS_IN#TUNNELED__acl ACL first. This currently consists of only two ACEs–one for CAPWAP control traffic that uses destination UDP port 5246 and one for CAPWAP data traffic that uses destination UDP port 5247.

Custom Applications

Following the deployment of ACEs for the prm-APIC_QOS_IN#TUNNELED__acl, EasyQoS will again determine if there is sufficient QoS TCAM space available for additional ACEs within any of the ACLs that are

part of the ingress classification & marking QoS policy-map. If sufficient QoS TCAM space is available for additional ACEs, EasyQoS will check to see if all ACEs for all Custom applications have been deployed. EasyQoS assigns a Rank to all Custom and Favorite applications. Custom applications have a Rank of 1. Applications within the NBAR taxonomy do not have a Rank by default. However, they are given a Rank of 10,000 when assigned as a Favorite application. EasyQoS processes applications by Rank first–from lowest number to highest number.

If the ACEs for all of the Custom applications have not been deployed, EasyQoS will select the next Custom application. Only Custom applications consisting IP addresses, IP ports, and TCP/UDP ports can be provisioned onto Catalyst switch platforms. This is discussed in the *Custom Applications Provisioned within ACLs* section of this document. How EasyQoS determines which ACL to deploy ACE entries for applications, based on the business-relevance and traffic-class attribute values of the application, is discussed in the *Effects of Changing Business Relevance* section of this document.

EasyQoS will continue to check for available QoS TCAM space and continue to deploy ACEs for Custom applications until either all Custom applications have been deployed, or the available QoS TCAM space is exhausted.

Favorite Applications

After ACEs for all Custom applications have been deployed, EasyQoS will begin parsing the applications within the NBAR taxonomy. If sufficient QoS TCAM space is available for additional ACEs, EasyQoS will check to see if all ACEs for Favorite applications have been deployed. This is because EasyQoS processes applications by Rank first, and Favorite applications are assigned a rank of 10,000. If the ACEs for all of the Favorite applications have not been deployed, EasyQoS will select the next Favorite application.

EasyQoS will check to see if the Favorite application has the traffic-class attribute set to one of the following:

- VoIP Telephony
- Broadcast Video
- Real-Time Interactive
- Multimedia Conferencing.

If the traffic-class attribute for the Favorite application matches one of these, EasyQoS will check to see if any of the indicative ports for the Favorite application are duplicates. Many voice and video apps known to NBAR include indicative ports for signaling protocols such as SIP, Cisco SCCP, STUN, etc. Signaling protocols should not be configured into voice and video ACLs. Instead they should appear within the Signaling ACL. Hence they should not be duplicated within the voice and video ACLs. Additionally many collaboration applications include indicative ports for additional functionality such as IMAP, etc. Email protocols should appear within the Bulk Data ACL. Hence they should not be duplicated within the voice and video ACLs either.

For the purposes of this document voice and video ACLs refer to the following ACLs:

- prm-APIC_QOS_IN#VOICE__acl
- prm-APIC_QOS_IN#BROADCAST__acl
- prm-APIC_QOS_IN#REALTIME__acl

- prm-APIC_QOS_IN#MM_CONF__acl.

EasyQoS will also check to see if the Favorite application is identified by any other indicative TCP or UDP ports. If the ports by which the application is identified correspond to TCP destination ports 80, 443, or 8080, EasyQoS will again not implement ACEs for these ports. This is because many applications use the ports corresponding to HTTP (port 80 or 8080) and HTTPS (port 443). Hence, Layer 2-4 ACEs are not effective at identifying applications that use these ports.

If the Favorite application is identified by any other indicative UDP or TCP ports, EasyQoS will generate ACEs for that Favorite application. If the business-relevance attribute is set for Business Relevant, the ACE entry(s) will be generated under the class-map definition based on the traffic-class attribute to which the Favorite application belongs to within the NBAR taxonomy. If the business-relevance attribute is set for Business Irrelevant, the ACE entry(s) will be generated under the class-map definition for Scavenger traffic. If the business-relevance attribute is set for Default, no ACE entry(s) will be generated under any class-map definition. EasyQoS will continue to do this until either all Favorite applications have been deployed or the available QoS TCAM space is exhausted.

Other Applications within the NBAR Taxonomy

After all Favorite applications are deployed, EasyQoS will determine if there is any available QoS TCAM space left for additional ACE entries. If available space exists, EasyQoS will distribute the available space across the various traffic classes. By distributing the available TCAM space across the various traffic classes, EasyQoS ensures that at least some applications from the NBAR taxonomy for each traffic class are represented in the ACLs that are generated for each traffic class. EasyQoS selects applications from each of the traffic classes based upon popularity–otherwise known as the *NBAR commonly-used attribute*. Every application within the NBAR taxonomy is assigned a value for the commonly-used attribute. Values range from 10 (most popular) to 1 (least popular).

EasyQoS will select an application from one of the traffic classes based upon popularity. If multiple applications have the same popularity, EasyQoS will select the next application alphabetically from those that have the same popularity. EasyQoS will check to see if the application has the traffic-class attribute set to one of the following:

- VoIP Telephony
- Broadcast Video
- Real-Time Interactive
- Multimedia Conferencing

If the traffic-class attribute for the application matches one of these, EasyQoS will check to see if any of the indicative ports for the application are duplicates. Many voice and video apps known to NBAR include indicative ports for signaling protocols such as SIP, Cisco SCCP, STUN, etc. Signaling protocols should not be configured into voice and video ACLs. Instead they should appear within the Signaling ACL. Hence they should not be duplicated within the voice and video ACLs. Additionally many collaboration apps include indicative ports for additional functionality such as IMAP, etc. Email protocols should appear within the Bulk Data ACL. Hence they should not be duplicated within the voice and video ACLs either. For the purposes of this document voice and video ACLs refer to the following ACLs:

- prm-APIC_QOS_IN#VOICE__acl

- prm-APIC_QOS_IN#BROADCAST__acl
- prm-APIC_QOS_IN#REALTIME__acl
- prm-APIC_QOS_IN#MM_CONF__acl.

EasyQoS will also check to see if the application is identified by any other indicative TCP or UDP ports. If the ports by which the application is identified correspond to TCP destination ports 80, 443, or 8080, EasyQoS will again not implement ACEs for these ports. This is because many applications use the ports corresponding to HTTP (port 80 or 8080) and HTTPS (port 443). Hence, Layer 2-4 ACEs are not effective at identifying applications that use these ports.

If the application is identified by any other indicative UDP or TCP ports, EasyQoS will generate ACEs for that application. If the business-relevance attribute is set for Business Relevant, the ACE entry(s) will be generated under the class-map definition based on the traffic-class attribute to which the application belongs to within the NBAR taxonomy. If the business-relevance attribute is set for Business Irrelevant, the ACE entry(s) will be generated under the class-map definition for Scavenger traffic. If the business-relevance attribute is set for Default, no ACE entry(s) will be generated under any class-map definition. EasyQoS will continue to do this until either all applications within the traffic class have been deployed or the available QoS TCAM space for the traffic class is exhausted.

EasyQoS will continue to do this for all traffic classes until either all applications within all traffic classes have been deployed or the available QoS TCAM space for all traffic classes is exhausted.

Catalyst Switch Roles

Catalyst and Nexus switch platforms can function in one of the following three possible roles within APIC-EM—reflecting a traditional 3-layered campus architecture:

- Core-layer switch
- Distribution-layer switch
- Access-layer switch

When APIC-EM discovers and places network devices into the device inventory database, it will classify each network device in one of five roles, discussed in the *APIC-EM and the EasyQoS Application* section of this document. For Catalyst and Nexus switches, EasyQoS uses the device role in order to determine what, if any, ingress classification & marking QoS policy to apply to each switch port, based on the role of the switch within the network infrastructure. Hence, it is highly important that the network operator review (and if necessary modify) the role of each network device within APIC-EM before implementing EasyQoS policies.

The following figure shows the roles the various supported Catalyst and Nexus switches can participate within the EasyQoS Solution; as well as the QoS policies applied to each switch based upon its role.

Figure 52 **Catalyst and Nexus Switch Roles within the EasyQoS Solution**

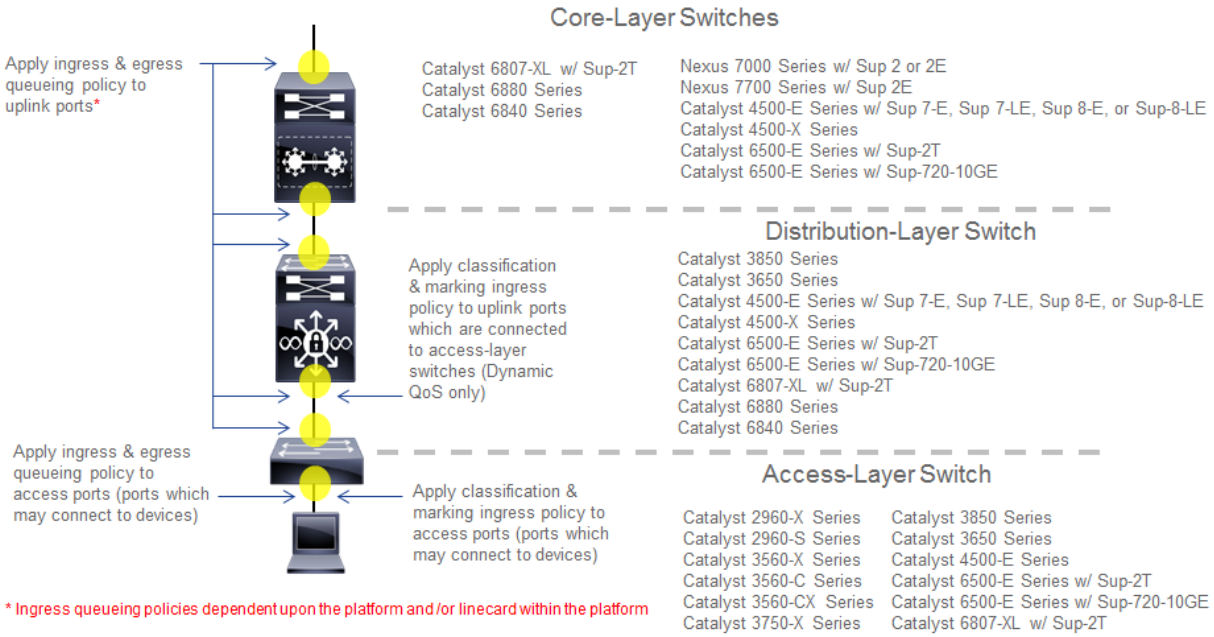

The following are the restrictions regarding the roles that the various supported Catalyst and Nexus switch platforms can have within the EasyQoS solution.

- Catalyst 6500-E Series switches with Sup-2T supervisors, Catalyst 6500-E Series switches with Sup-720-10GE supervisors, Catalyst 6807-XL switches with Sup-2T, and Catalyst 4500-E Series switches with Sup7-E, Sup7-LE, and Sup-8E supervisors are supported in the roles of a core-layer, distribution-layer, or access-layer switch.
- Nexus 7000 Series with Sup2 or 2E supervisors, and Nexus 7700 with Sup2E supervisors are supported only in the role of a core-layer switch.
- Catalyst 6880 Series switches, Catalyst 6840 Series switches, and Catalyst 4500-X Series switches are supported only in the roles of a core-layer or distribution-layer switch.
- Catalyst 3850 Series switches and Catalyst 3650 Series switches are supported in the roles of a distribution-layer or an access-layer switch.
- Catalyst 2960-X, 2960-XR, 2960-S, 3560-X, 3560-C, and 3560-CX switches are only supported in the role of an access-layer switch.

A single switch functioning as both a distribution-layer switch and an access-layer switch simultaneously is not supported. Multiple switch platforms of the same model can individually function in the role of a distribution-layer switch or access-layer switch within a single deployment.

Core-Layer Switch QoS Design

For devices operating as core-layer switches, EasyQoS will only apply ingress and/or egress queuing policies to the uplinks ports. *Uplink ports* refer to ports that connect to other core-layer switches or to distribution-layer switches. The specifics as to whether both ingress and egress queuing policies, or only egress queuing policies are applied, are dependent upon whether the particular Catalyst switch platform

and/or line card within the platform supports both ingress and egress queuing, or only egress queuing. This is discussed in detail for each platform and/or line card in the queuing design sections of this document.

Because only queuing policies are pushed to core-layer switches by EasyQoS, the QoS policy is the same for core-layer switches, regardless of whether the customer chooses to implement Static or Dynamic QoS. Dynamic QoS is discussed within the *Dynamic QoS Design* section of this document.

Access-Layer Switch QoS Design

Access-layer switch QoS design consists of the following policies:

- Ingress/egress queuing policies applied to access-edge switch ports and uplink switch ports
- Ingress classification & marking policies applied to access-edge switch ports

Ingress/Egress Queuing Policies

Regardless of whether the network operator has implemented Static or Dynamic QoS, queuing policies will always be pushed to access-layer switches. The EasyQoS application will apply ingress and/or egress queuing policies to both access-edge ports and uplink ports. Access-edge ports refer to ports that directly connect to end devices, such as laptops, PCs, IP Phones, wireless Access Points, etc. Uplink ports refer to ports that connect to distribution-layer switches. The specifics as to whether both ingress and egress queuing policies, or only egress queuing policies are applied, are dependent upon whether the particular Catalyst switch platform and/or line card within that platform supports both ingress and egress queueing, or only egress queueing. This is discussed in detail for each platform and/or line card in the queueing design sections of this document.

Ingress Classification & Marking Policies

The ingress classification & marking policy provisioned onto access-layer switches by EasyQoS is dependent upon whether the network operator chooses to implement Static QoS or Dynamic QoS. For Static QoS, EasyQoS will apply an access-layer ingress classification & marking policy to all access-edge ports. The ingress classification & marking policy consists of policy-maps, which contain class-maps, which in turn contain ACLs with Layer 3 & 4 ACEs. Layer 3 & 4 refers to IP addresses, protocols (that is, TCP, UDP, etc.) and higher-layer ports (HTTP, Telnet, FTP, etc.). The access-layer classification & marking policy establishes the QoS trust boundary and policy enforcement point at the ingress edge of the network.

The following are the class-map definitions for the ingress classification & marking policy pushed by EasyQoS to Catalyst switch platforms, when configured in the role of an access-layer switch within APIC-EM.

```
!
class-map match-any prm-APIC_QOS_IN#VOICE
 match access-group name prm-APIC_QOS_IN#VOICE__acl
class-map match-any prm-APIC_QOS_IN#BROADCAST
 match access-group name prm-APIC_QOS_IN#BROADCAST__acl
class-map match-any prm-APIC_QOS_IN#REALTIME
```

```
 match access-group name prm-APIC_QOS_IN#REALTIME__acl
class-map match-any prm-APIC_QOS_IN#MM_CONF
 match access-group name prm-APIC_QOS_IN#MM_CONF__acl
class-map match-any prm-APIC_QOS_IN#MM_STREAM
 match access-group name prm-APIC_QOS_IN#MM_STREAM__acl
class-map match-any prm-APIC_QOS_IN#SIGNALING
 match access-group name prm-APIC_QOS_IN#SIGNALING__acl
class-map match-any prm-APIC_QOS_IN#OAM
 match access-group name prm-APIC_QOS_IN#OAM__acl
class-map match-any prm-APIC_QOS_IN#TRANS_DATA
 match access-group name prm-APIC_QOS_IN#TRANS_DATA__acl
class-map match-any prm-APIC_QOS_IN#BULK_DATA
 match access-group name prm-APIC_QOS_IN#BULK_DATA__acl
class-map match-any prm-APIC_QOS_IN#SCAVENGER
 match access-group name prm-APIC_QOS_IN#SCAVENGER__acl
class-map match-any prm-APIC_QOS_IN#TUNNELED
 match access-group name prm-APIC_QOS_IN#TUNNELED__acl
!
```

The following is the policy-map definition for the ingress classification & marking policy pushed by EasyQoS to the switch platforms.

```
!
policy-map prm-APIC_QOS_IN
 class prm-APIC_QOS_IN#VOICE
  set DSCP ef
 class prm-APIC_QOS_IN#BROADCAST
  set DSCP cs5
 class prm-APIC_QOS_IN#REALTIME
  set DSCP cs4
 class prm-APIC_QOS_IN#MM_CONF
  set DSCP af41
 class prm-APIC_QOS_IN#MM_STREAM
  set DSCP af31
```

```
    class prm-APIC_QOS_IN#SIGNALING
     set DSCP cs3
    class prm-APIC_QOS_IN#OAM
     set DSCP cs2
    class prm-APIC_QOS_IN#TRANS_DATA
     set DSCP af21
    class prm-APIC_QOS_IN#BULK_DATA
     set DSCP af11
    class prm-APIC_QOS_IN#SCAVENGER
     set DSCP cs1
    class prm-APIC_QOS_IN#TUNNELED
    class class-default
     set DSCP default
   !
```

Eleven of the twelve classes defined within the RFC 4594-based Cisco 12-Class QoS model are defined within the class-maps and policy-map above. The 12th traffic class corresponds to Network Control traffic. The access-layer ingress classification & marking policy is intended to be applied to switch ports that connect directly to end-user devices–not network equipment, such as routers and other switches. Network Control traffic should never be seen by access-layer switch ports connected to end-user devices. Hence the ingress classification & marking policy does not define a class-map or traffic-class definition to account for Network Control traffic.

A Cisco wireless Access Point may be connected to an access-layer switch port. Becauuse conditional trust is not used by EasyQoS, the ingress classification & marking policy-map must take into account CAPWAP tunneled traffic from the Access Point. CAPWAP tunneled traffic can be either CAPWAP control traffic or CAPWAP data traffic. For CAPWAP data traffic, the DSCP marking of the outer CAPWAP header is set by the Access Point, and is based on the DSCP marking of the IP packet sent by the wireless client. This is discussed further in the *WLC QoS Design* section of this document. For CAPWAP control traffic, the DSCP marking of the outer CAPWAP header is set with a DSCP marking of Class Selector 6 (CS6). The prm-APIC_QOS_IN#TUNNELED traffic class is used to match on tunneled traffic, such as CAPWAP control and data traffic. Within the policy-map definition, no action is taken for prm-APIC_QOS_IN#TUNNELED traffic. This ensures that the DSCP marking of both CAPWAP control and CAPWAP data traffic is preserved as it traverses the access-layer ingress switch port connected to the Access Point.

The default-class within the policy-map definition is configured to set all traffic that does not match any of the previous traffic classes to a DSCP marking of default (Best Effort). This ensures that all traffic that does not match one of the traffic classes is bleached–in other words provided a best effort service.

The ingress classification & marking policy is applied to all access-edge switch ports on each switch in a stackable switch platform or line card in a modular switch platform; and all switches within a switch stack or a VSS pair. Access-edge switch ports are switch ports that are used to connect end-user devices.

Note: QoS policy cannot be applied to switch ports that are part of a virtual switch link (VSL) connecting the two switches of a Virtual Switching System (VSS) pair.

The following show an example of the application of the service-policy to a Gigabit Ethernet access-edge switch port.

```
interface GigabitEthernetx/x/x
 service-policy input prm-APIC_QOS_IN
```

For uplink ports on switches configured with the role of an access-layer switch within APIC-EM, no ingress classification & marking policy is applied to the switch port. Instead, the switch port is configured to trust DSCP markings from devices attached to the switch port. For MLS QoS-based switches (Catalyst 2960 Series, Catalyst 3560 Series, Catalyst 3750 Series and Catalyst 6500 Series switches with Sup-720 Supervisors), this must be explicitly configured via the following command, because the default port trust-state is untrusted on MLS QoS-based switches.

```
mls qos trust DSCP
```

For MQC-based platforms (Catalyst 3850 Series, Catalyst 3650 Series, and Catalyst 4500 Series), C3PL-based platforms (Catalyst 6500 Series with Sup-2T supervisors, 6807-XL, 6880 Series, and 6840 Series), and NX OS platforms (Nexus 7000 and 7700 Series), the default port trust-state is trusted. No explicit configuration command needs to be pushed from APIC-EM to these platforms.

Access-Control Lists

The following are the ACL definitions for the ingress classification & marking policy pushed by EasyQoS to all of the Catalyst switching platforms supported by EasyQoS, when the platform functions as an access-layer switch. The specific ACE entries within each ACL are not shown.

```
!
ip access-list extended prm-APIC_QOS_IN#VOICE__acl
ip access-list extended prm-APIC_QOS_IN#BROADCAST__acl
ip access-list extended prm-APIC_QOS_IN#REALTIME__acl
ip access-list extended prm-APIC_QOS_IN#MM_CONF__acl
ip access-list extended prm-APIC_QOS_IN#MM_STREAM__acl
ip access-list extended prm-APIC_QOS_IN#SIGNALING__acl
ip access-list extended prm-APIC_QOS_IN#OAM__acl
ip access-list extended prm-APIC_QOS_IN#TRANS_DATA__acl
ip access-list extended prm-APIC_QOS_IN#BULK_DATA__acl
ip access-list extended prm-APIC_QOS_IN#SCAVENGER__acl
ip access-list extended prm-APIC_QOS_IN#TUNNELED__acl
!
```

The following provides an example of what the ACL entries will look like after the ACE entries have been populated–with just a few of the ACE entries shown for compactness.

```
!
ip access-list extended prm-APIC_QOS_IN#VOICE__acl
ip access-list extended prm-APIC_QOS_IN#BROADCAST__acl
ip access-list extended prm-APIC_QOS_IN#REALTIME__acl
ip access-list extended prm-APIC_QOS_IN#MM_CONF__acl
ip access-list extended prm-APIC_QOS_IN#MM_STREAM__acl
…
 remark citrix—Citrix
 permit tcp any any eq 1494
 permit udp any any eq 1494
 permit tcp any any eq 2598
 permit udp any any eq 2598
…
ip access-list extended prm-APIC_QOS_IN#SIGNALING__acl
…
 remark skinny
 permit tcp any any eq 2000
 permit tcp any any eq 2001
 permit tcp any any eq 2002
 remark sip
 permit tcp any any eq 3478
 permit udp any any eq 3478
…
ip access-list extended prm-APIC_QOS_IN#OAM__acl
…
 remark dhcp—Dynamic Host Configuration Protocol
 permit udp any any range 67 68
 remark dns—Domain Name System
 permit tcp any any eq 53
 permit udp any any eq 53
```

```
 permit tcp any any eq 5353
 permit udp any any eq 5353
…
ip access-list extended prm-APIC_QOS_IN#TRANS_DATA__acl
…
 remark ibm-db2—IBM-DB2
 permit tcp any any eq 523
 permit udp any any eq 523
 remark sap—SAP
 permit tcp any any eq 3200
 permit tcp any any eq 3300
 permit tcp any any eq 3600
…
ip access-list extended prm-APIC_QOS_IN#BULK_DATA__acl
…
 remark ftp—File Transfer Protocol
 permit tcp any any eq 21
 permit tcp any any eq 21000
 remark imap—Internet Message Access Protocol version 4
 permit tcp any any eq 143
 permit udp any any eq 143
 permit tcp any any eq 220
 permit udp any any eq 220
…
ip access-list extended prm-APIC_QOS_IN#SCAVENGER__acl
…
 remark blizwow—World of Warcraft
 permit tcp any any eq 3724
 permit udp any any eq 3724
 remark call-of-duty—Call of Duty
 permit tcp any any eq 20500
 permit tcp any any eq 20510
```

```
 permit tcp any any eq 28960
 permit udp any any eq 20500
…
ip access-list extended prm-APIC_QOS_IN#TUNNELED__acl
 remark CAPWAP Control Traffic
 permit udp any any eq 5246
 remark CAPWAP Data Traffic
 permit udp any any eq 5247
!
```

Remarks are used in order to make it visually easy for the network operator to determine which applications have been deployed.

The specific applications that appear within each ACL are dependent upon the applications declaratively selected by the network operator as being business-relevant, default, or business-irrelevant, as well as any DSCP, IP address, or TCP/UDP port based Custom applications defined by the network operator within the EasyQoS web-based GUI.

Effects of Changing Business Relevance on ACLs

For Catalyst switch platforms, ACLs with ACE entries corresponding to the IP addresses and ports used to identify an application, are provisioned by EasyQoS in order to classify and mark the application as it enters the access-edge switch port. Therefore, changing the business relevance of an application within the EasyQoS web-based GUI simply changes the placement of the ACE entry within the ACLs that are referenced from the class-map definitions for each traffic class–based on the following rules:

- If an application is moved from having a business relevance attribute value of business-relevant to business-irrelevant (that is, moved from the business-relevant grouping to the business-irrelevant grouping within the application registry for the policy applied to the device), the ACE entry for the application will be provisioned within the prm-APIC_QOS_IN#SCAVENGER__acl ACL. Hence, all applications that have been identified as being business-irrelevant are classified into the Scavenger traffic class are re-marked to a Class Selector 1 (CS1) per-hop behavior. This assumes the application can be uniquely identified by IP addresses, IP ports, or UDP/TCP ports and that there is sufficient TCAM space available to provision the ACE entries.

- If an application is moved from having a business relevance attribute of either business-relevant or business-irrelevant, to default (that is, moved from either the business-relevant or business-irrelevant grouping to the default grouping within the application registry for the policy applied to the device), no ACE entry for this application will be provisioned in any ACL. All applications with a default business relevance are classified in the Default traffic class and re-marked to a best effort per-hop behavior.

- If an application is moved from having a business relevance attribute value of either business-irrelevant or default, to business-relevant (that is, moved from the business-irrelevant or the default grouping to the business-relevant grouping within the application registry for the policy applied to the device), the ACE entry for the application will be provisioned into the ACL corresponding to the traffic-class attribute for that particular application. All 1300+ applications identified within the NBAR taxonomy

have a default setting for the traffic-class attribute–meaning the traffic-class to which the application belongs. This attribute cannot be modified currently within the EasyQoS web-based GUI. All Custom applications created within the EasyQoS web-based GUI must have a traffic-class value assigned to them when they are created. Note that the traffic-class attribute value assigned to Custom applications and all 1300+ applications known by the NBAR taxonomy does not include values for Scavenger or Default traffic classes. Hence applications identified as being business-relevant have ACE entries generated within the traffic class to which the application belongs. This assumes the application can be uniquely identified by IP addresses, IP ports, or UDP/TCP ports and that there is sufficient TCAM space available to provision the ACE entries.

Custom Applications Provisioned within ACLs

EasyQoS is not able to deploy Layer 2-4 ACE entries for Custom applications that consist of URL strings. APIC-EM will therefore skip over the deployment of Custom applications consisting of URL strings when configuring Catalyst switch platforms. Hence Catalyst switches are unable to implement a Custom application that is based on the use of a URL to identify the application. Catalyst switches can only implement Custom applications that are based upon DSCP values, IP addresses, IP ports, and TCP/UDP ports. Custom applications based on IP addresses, ports, and/or DSCP values are simply added as additional ACE entries under the ACL corresponding to the particular traffic-class to which the Custom application has been defined by the network operator.

Custom applications are by default marked as a Favorite applications by EasyQoS. In order to include a Custom application within a QoS policy, the network operator must drag-and-drop Custom applications into one of the three business relevance groupings within the EasyQoS web-based GUI interfaces. This is discussed in the *APIC-EM and the EasyQoS Application* section of this document.

An example of a Custom application configured for the multimedia-conferencing traffic class is shown below.

```
!
ip access-list extended prm-APIC_QOS_IN#MM_CONF__acl
 remark Custom_Port-App
 permit udp any 10.0.10.0 0.0.0.255 range 3001 3010
 permit udp 10.0.10.0 0.0.0.255 range 3001 3010 any
!
```

In the example above, the Custom application, based on a destination server IP address range and port range–also referred to as the producer–has been specified to be bi-directional by the network operator, through the EasyQoS web-based GUI. Hence, the reverse of the ACE entry is also generated to allow traffic from the server IP address and port range to also be treated the same.

In the example above, a server IP address range (10.0.10.0-10.0.10.255) and port range (UDP 3001-3010) is configured. Custom applications also support single IP addresses and ports, and the use of “any” specified as the destination IP address. Although a single UDP port range is specified in the example above, multiple UDP and/or TCP ports can be configured as well–each of which would appear as a separate “permit” statement.

Additional IP Address/Port-based Custom applications will generate additional ACE entries within ACLs, similar to those shown in the example above, based on the rules discussed within the *Effects of Changing Business Relevance on ACLs* section above.

A more sophisticated example, shown below, adds a source IP address or range–also referred to as the *consumer*–as well as the destination IP address or range–referred to as the producer to the Custom application. Again, this is configured bi-directionally via the APIC-EM EasyQoS web-based GUI by the network operator. An example of the same application–but with a consumer this time–is shown below.

```
!
ip access-list extended prm-APIC_QOS_IN#MM_CONF__acl
 remark Custom_Port-App_Consumer__Custom_Port-App
 permit udp host 10.0.20.20 eq 3100 10.0.10.0 0.0.0.255 range 3001 3010
 permit udp 10.0.10.0 0.0.0.255 range 3001 3010 host 10.0.20.20 eq 3100
!
```

The combination of the producer and consumer, along with the ability to apply the policy bi-directionally, essentially gives the network operator the ability to use nearly the full CLI functionality in terms of being able to configure QoS ACE entries.

Cisco Device Endpoints

APIC-EM also discovers Cisco endpoints, such as Cisco IP phones, Cisco video surveillance cameras, Cisco TelePresence devices, and Cisco video conferencing endpoints. CDP information provided by the Cisco device endpoint also identifies the device type. This information is necessary because different device types are populated via ACE entries within different ACLs with different DSCP markings.

As part of Static QoS, the IP addresses of these endpoints, along with the appropriate DSCP markings for traffic generated by these devices are also added to the ingress classification & marking policy for each switch to which the endpoints are connected. The DSCP values populated into the ACLs for Static QoS is shown in the table below.

Table 4 Wired Cisco Device Endpoints and DSCP Markings

Wired Endpoint Device Type	Allowed DSCP Values	Static QoS ACL in Which the ACE Entry Will be Added	Description
Cisco IP Phone	EF AF41	prm-APIC_QOS_IN#VOICE__acl prm-APIC_QOS_IN#MM_CONF__acl	Cisco IP Phones typically send VoIP media (and associated RTCP flows) marked as EF when a call is audio only, and VoIP and video media both marked as AF41 wen a call is both audio and video.
Cisco Video Conferencing Endpoints	EF AF41	prm-APIC_QOS_IN#VOICE__acl prm-APIC_QOS_IN#MM_CONF__acl	Cisco Video Conferencing Endpoints typically send VoIP media (and associated RTCP flows) marked as EF

			when a call is audio only, and VoIP and video media both marked as AF41 when a call is both audio and video.
Cisco TelePresence Device	CS4 EF	prm-APIC_QOS_IN#REALTIME__acl prm-APIC_QOS_IN#VOICE__acl	Cisco TelePresence devices typically send VoIP and video media (and associated RTCP flows) both marked as CS4 when a call is audio and video; and VoIP media (and associated RTCP flows) marked as EF when a call is audio only.
Cisco Video Surveillance Cameras	CS5	prm-APIC_QOS_IN#BROADCAST__acl	H.264 or H.265 encoded streaming video surveillance typically uses the RTP protocol for transport. The network administrator may need to ensure that streaming video is sent with a CS5 marking.

The DSCP marking of voice and video media for devices under the control of CUCM can be modified via the CUCM GUI. Hence, the CUCM administrator should ensure that the markings of audio and video media are the same for the endpoints as listed in the table above.

APIC-EM populates both the prm-APIC_QOS_IN#VOICE__acl and the prm-APIC_QOS_IN#MM_CONF__acl within the static ingress classification & marking policy on the switch to which a wired Cisco IP Phone endpoint is discovered with permit statements for the source IP address of a Cisco IP Phone along with the expected media markings (DSCP values). An example is as follows:

```
!
ip access-list extended prm-APIC_QOS_IN#VOICE__acl
 permit ip host 10.4.81.21 any DSCP ef
!
ip access-list extended prm-APIC_QOS_IN#MM-CONF__acl
 permit ip host 10.4.81.21 any DSCP af41
!
```

Cisco IP Phones are expected to generate voice traffic with a DSCP marking of EF in an audio-only call, and voice & video traffic with a DSCP marking of AF41 in a video call.

APIC-EM populates both the prm-APIC_QOS_IN#VOICE__acl and the prm-APIC_QOS_IN#MM-CONF__acl within the static ingress classification & marking policy on the switch to which a wired Cisco video conferencing endpoint is discovered with permit statements for the source IP address of the cisco video conferencing endpoint along with the expected media markings (DSCP values). An example is as follows:

```
!
```

```
ip access-list extended prm-APIC_QOS_IN#VOICE__acl
 permit ip host 10.4.81.22 any DSCP ef
!
ip access-list extended prm-APIC_QOS_IN#MM-CONF__acl
 permit ip host 10.4.81.22 any DSCP af41
!
```

Cisco video conferencing endpoints are also expected to generate voice traffic with a DSCP marking of EF in an audio-only call, and voice & video traffic with a DSCP marking of AF41 in a video call.

APIC-EM populates both the prm-APIC_QOS_IN#VOICE__acl and the prm-APIC_QOS_IN#REALTIME__acl within the static ingress classification & marking policy on the switch to which a wired Cisco TelePresence endpoint is discovered with permit statements for the source IP address of the Cisco TelePresence endpoint along with the expected media markings (DSCP values).

```
!
ip access-list extended prm-APIC_QOS_IN#VOICE__acl
 permit ip host 10.4.81.23 any DSCP ef
!
ip access-list extended prm-APIC_QOS_IN#REALTIME__acl
 permit ip host 10.4.81.23 any DSCP cs4
!
```

Cisco TelePresence endpoints are expected to generate voice traffic with a DSCP marking of EF in an audio-only call and voice & video traffic with a DSCP marking of CS4 in a video call.

Note: The default marking for Cisco TelePresence devices may change from CS4 to AF41 within future CUCM software versions. This reflects the fact that TelePresence video media has evolved over time from exhibiting a behavior more similar to an inelastic flow to exhibiting a behavior more similar to an elastic flow. There is currently no means for the network operator to change the value of DSCP markings populated in the static ACLs by APIC-EM for discovered endpoint devices. Therefore, the network operator must ensure that Cisco TelePresence devices mark video media as CS4 within the CUCM GUI, in order to correctly operate with APIC-EM EasyQoS.

Voice and video media from Cisco IP Phones, Cisco video conferencing endpoints, and Cisco TelePresence endpoints use RTP/UDP transport, typically in the port range from UDP ports 16384-32767, using even numbered ports. However, these endpoints may also generate other traffic, such as RTP Control Protocol (RTCP) traffic. RTCP traffic typically uses the next higher odd numbered UDP port. For example, if the audio media port is UDP 16384, the associated RTCP control stream is typically UDP 16385. RTCP provides feedback information regarding the quality of the media stream, including information regarding lost packets. Cisco IP Phones, Cisco video conferencing endpoints, and Cisco TelePresence endpoints send RTCP streams with the same DSCP marking as their corresponding media flow. Hence the ACE entries listed in Table 1 above apply to RTCP flows as well.

APIC-EM will populate the prm-APIC_QOS_IN-#BROADCAST__acl within the static ingress classification & marking policy on the switch to which a wired Cisco video surveillance camera is discovered with permit statements for the source IP address of the Cisco video surveillance camera along with the expected media marking (DSCP value). An example is as follows:

```
!
ip access-list extended prm-APIC_QOS_IN#VOICE__acl
 permit ip host 10.4.81.24 any DSCP cs5
!
```

Cisco video surveillance cameras are expected to generate video traffic with a DSCP marking of CS5. This video traffic is typically H.264 or H.265 encoded streaming video sent via the RTP protocol that uses UDP transport. Cisco video surveillance cameras may also sent RTSP control traffic using TCP ports 554 or 8554. This traffic should not be sent with a DSCP marking of CS5, because it is not streaming media (i.e. video). RTSP traffic should automatically be categorized as signaling traffic and an ACE entry for RTSP traffic placed into the prm-APIC_QOS_IN#SIGNALING-ACL.

APIC-EM will periodically re-discover devices on the network and automatically update the entries in the ACLs for devices that have been added/moved/changed. As a prerequisite for adds/moves/changes, the network operator will need to enable SNMP traps on the access switches to be sent to APIC-EM. After the interface connected to a Cisco IP Phone, Cisco video conferencing endpoint, Cisco Telepresence device, or Cisco video surveillance camera goes up or down APIC-EM will receive an SNMP trap and starts collecting information from the access switch that generated the SNMP trap, about the new Cisco endpoints. This takes approximately 80 seconds plus the time needed for the collection of the device information. After the Cisco endpoint information is collected, APIC-EM automatically pushes ACE entries containing the source IP address of the endpoint device to any destination, with the prm-APIC_QOS_IN#VOICE__acl, prm-APIC_QOS_IN#BROADCAST__acl, prm-APIC_QOS_IN#REALTIME__acl, and prm-APIC_QOS_IN#MM_CONF__acl entries with IP + DSCP in both static and dynamic policies.

Dynamic QoS

When the network operator has implemented Dynamic QoS, EasyQoS will configure a dynamic policy-map shell for ingress classification and marking of voice and video traffic only, for each switch port. These policy-map shells are dynamically populated with ACEs and dynamically placed/removed across the required switch port, based upon notification of calls beginning/ending from CUCM. This is discussed further in the *Dynamic QoS Design* section of this document.

Distribution-Layer Switch QoS Design

The QoS policy configuration pushed to distribution-layer switches by APIC-EM EasyQoS is not dependent upon whether the network operator chooses to implement Static or Dynamic QoS.

Ingress/Egress Queuing Policies

Regardless of whether the network operator has chosen to implement Static or Dynamic QoS, the EasyQoS application will always apply ingress and/or egress queuing policies to the uplinks ports. For a switch in the role of a distribution-layer switch, *uplink ports* refer to ports that connect to core-layer switches, to other

distribution-layer switches, or to access-layer switches. The specifics as to whether both ingress and egress queuing policies, or only egress queueing policies are applied, are dependent upon whether the particular Catalyst switch platform and/or line card within that platform supports both ingress and egress queueing, or only egress queueing. This is discussed in detail for each platform and/or line card in the *Catalyst and Nexus Switch Platform Queuing Design* section of this document.

Ingress Classification & Marking Policies

Regardless of whether the network operator has chosen to implement Dynamic QoS or Static QoS, the EasyQoS application will also additionally create and apply an ingress classification & marking policy to all uplink ports that connect to access-layer switches. This classification & marking policy is similar, but not identical, to the ingress classification & marking policy applied to switch ports connected to devices (access-edge switch ports) on switches functioning in the role of an access-layer switch within APIC-EM.

APIC-EM will create class-maps and policy-maps on every switch configured in the role of a distribution-layer switch within the policy scope. An example of the class-maps configured by EasyQoS are shown below.

```
!
class-map match-any prm-APIC_QOS_IN#CONTROL
 match access-group name prm-APIC_QOS_IN#CONTROL__acl
 match DSCP cs6
class-map match-any prm-APIC_QOS_IN#VOICE
 match access-group name prm-APIC_QOS_IN#VOICE__acl
 match DSCP ef
class-map match-any prm-APIC_QOS_IN#BROADCAST
 match access-group name prm-APIC_QOS_IN#BROADCAST__acl
match DSCP cs5
class-map match-any prm-APIC_QOS_IN#REALTIME
 match access-group name prm-APIC_QOS_IN#REALTIME__acl
match DSCP cs4
class-map match-any prm-APIC_QOS_IN#MM_CONF
 match access-group name prm-APIC_QOS_IN#MM_CONF__acl
 match DSCP af41
class-map match-any prm-APIC_QOS_IN#MM_STREAM
 match access-group name prm-APIC_QOS_IN#MM_STREAM__acl
class-map match-any prm-APIC_QOS_IN#SIGNALING
 match access-group name prm-APIC_QOS_IN#SIGNALING__acl
class-map match-any prm-APIC_QOS_IN#OAM
```

```
 match access-group name prm-APIC_QOS_IN#OAM__acl
class-map match-any prm-APIC_QOS_IN#TRANS_DATA
 match access-group name prm-APIC_QOS_IN#TRANS_DATA__acl
class-map match-any prm-APIC_QOS_IN#BULK_DATA
 match access-group name prm-APIC_QOS_IN#BULK_DATA__acl
class-map match-any prm-APIC_QOS_IN#SCAVENGER
 match access-group name prm-APIC_QOS_IN#SCAVENGER__acl
class-map match-any prm-APIC_QOS_IN#TUNNELED
 match access-group name prm-APIC_QOS_IN#TUNNELED__acl
!
```

With Dynamic QoS, policy-maps dynamically applied to access-edge switch ports classify & mark voice and video traffic based on signaling from CUCM. The marking of this traffic must be preserved as it enters the distribution-layer switch. Therefore, the distribution-layer policy contain the following additions:

- The class-map definition for prm-APIC_QOS_IN#VOICE also contains a “match DSCP ef” entry
- The class-map definition for prm-APIC_QOS_IN#MM_CONF also contains a “match DSCP af41” entry
- The class-map definition for prm-APIC_QOS_IN#REALTIME also contains a “match DSCP cs4” entry
- The class-map definition for prm-APIC_QOS_IN#BROADCAST also contains a “match DSCP cs4” entry.
- A class-map definition for control traffic (prm-APIC_QOS_IN#CONTROL) that matches on both an ACL and on DSCP markings (cs6) to preserve markings from the access-layer switches is also included.

All other class-maps match on just ACLs. An example of the policy-map configured by EasyQoS for distribution-layer switches is shown below.

```
!
policy-map prm-APIC_QOS_IN
 class prm-APIC_QOS_IN#VOICE
  set DSCP ef
 class prm-APIC_QOS_IN#BROADCAST
  set DSCP cs5
 class prm-APIC_QOS_IN#REALTIME
  set DSCP cs4
 class prm-APIC_QOS_IN#MM_CONF
  set DSCP af41
 class prm-APIC_QOS_IN#MM_STREAM
```

```
    set DSCP af31
   class prm-APIC_QOS_IN#SIGNALING
    set DSCP cs3
  class APIC_EM-CONTROL
    set DSCP cs6
   class prm-APIC_QOS_IN#OAM
    set DSCP cs2
   class prm-APIC_QOS_IN#TRANS_DATA
    set DSCP af21
   class prm-APIC_QOS_IN#BULK_DATA
    set DSCP af11
   class prm-APIC_QOS_IN#SCAVENGER
    set DSCP cs1
   class prm-APIC_QOS_IN#TUNNELED
   class class-default
    set DSCP default
  !
```

The policy-map at the distribution-layer switch also includes a class-map entry for control traffic. Control traffic may be generated by the access-layer switch and sent to the distribution-layer switch or generated by any Access-Points connected to an access-layer switch and sent to the distribution-layer switch.

The ACLs for the distribution-layer ingress classification & marking policy provisioned by EasyQoS to Catalyst switching platforms are basically the same as the ACLs provisioned by EasyQoS to Catalyst switching platforms–when configured in the role of an access-layer switch. These were discussed in the *Access-Control Lists* section above and will not be repeated here.

The service-policy is applied to all distribution-layer switch ports that connect to access-layer switches. An example of the configuration provisioned by EasyQoS is shown below.

```
!
interface TenGigabitEthernet1/0/1
 service-policy input prm-APIC_QOS_IN
!
```

APIC-EM EasyQoS does not apply the service-policy to distribution-layer switch ports that connect to other distribution-layer switches or to core-layer switches.

Catalyst and Nexus Switch Platform Queuing Design

This following sections discuss the ingress and/or egress queuing structures provisioned by EasyQoS to the various supported Catalyst and Nexus Series switch platforms.

Catalyst 2960-S/2960-X/2960-XR/3560-C/3560-CX Queuing Design

This section discusses the egress queuing structure provisioned by EasyQoS to the access-edge and uplink ports of the following switch platforms:

- Catalyst 2960-S Series
- Catalyst 2960-X Series
- Catalyst 2960-XR Series
- Catalyst 3560-C Series
- Catalyst 3560-CX Series

Catalyst 2960-S and 2960-X Series platforms must be running a LAN Base image in order to support the following QoS features:

- Policy maps
- Policing & marking
- Mapping Tables
- Weighted Tail Drop (WTD)

Because these features are used by EasyQoS, only Catalyst 2960-S and 2960-X Series platforms that run a LAN Base image are supported by EasyQoS.

Catalyst 2960-S and 2960-X Series platforms only support Layer 2 switching. Catalyst 2960-XR Series platforms run an IP Lite image, which includes Layer 3 switching support.

Catalyst 2960-S, 2960-X, and 2960-XR Series platforms are MLS QoS based, which require QoS to be enabled globally first before configuring any other QoS commands. The following global-level commands configured by EasyQoS enable QoS and set the internal COS-to-DSCP mapping table within the platform.

```
!
mls qos
! Globally Enables QoS
mls qos map cos-DSCP 0 8 16 24 32 46 48 56
! Maps CoS 5 to 46 (rest are default)
!
```

Note: The configurations shown for the Catalyst 2960 Series, Catalyst 3560 Series, and the Catalyst 3750 Series include comments, either before or after the actual commands. Comments are not part of the con-

figuration pushed by EasyQoS. They are included only to provide additional detail to the reader regarding the meaning of the command. Comment lines begin with a "!" within the configuration examples.

1P3Q3T Egress Queuing Design

Catalyst 2960-S, 2960-X, and Catalyst 2960-XR Series switches support only egress queuing. The 1P3Q3T egress queuing structure for these platforms implements DSCP-to-queue mapping with Weighted Tail Drop (WTD) for congestion avoidance, instead of WRED. Although there are three Weighted Tail Drop (WTD) thresholds, only two are configurable. The third threshold is by default set to the depth of the queue (that is, 100% queue depth).

The following figure shows the 1P3Q3T egress queuing model deployed by EasyQoS.

Figure 53 **1P3Q3T Egress Queuing for the Catalyst 2960-S, 2960-X, and 2960-XR Series Switches**

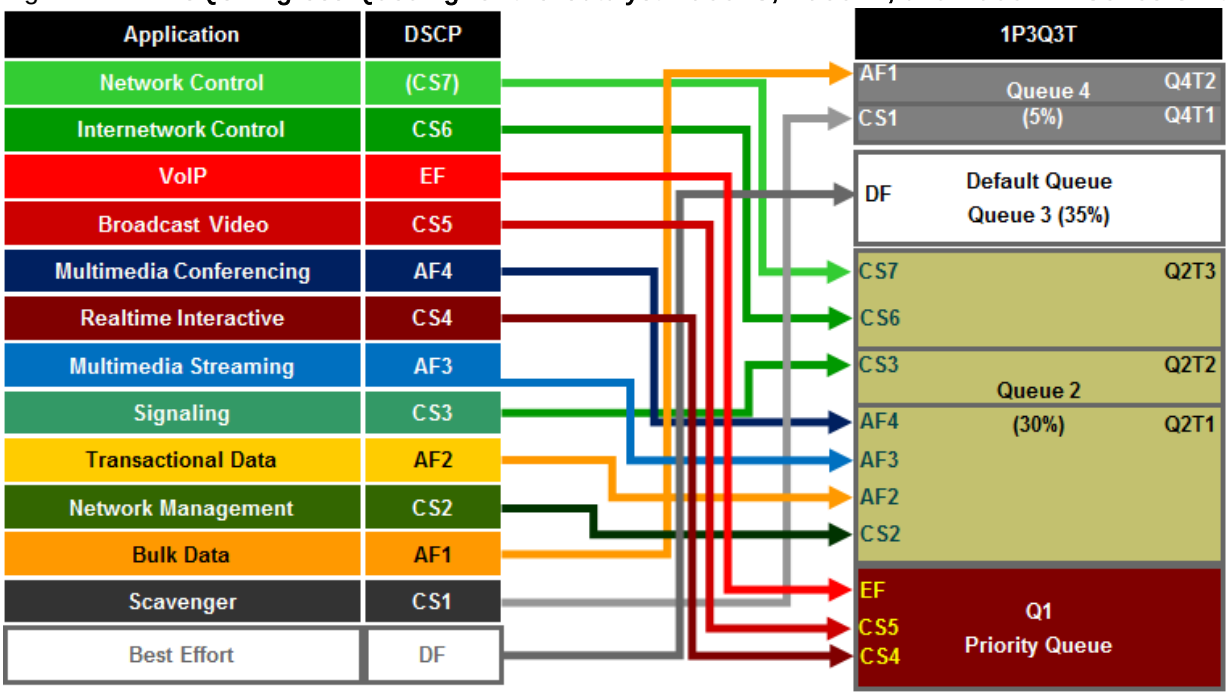

The following configuration provisioned by EasyQoS implements the 1P3Q3T egress queuing structure. Comments have been added to the configuration in order to explain each of the commands.

```
!
! Tunes Egress Queuing Buffers and Thresholds
mls qos queue-set output 1 buffers 15 30 35 20
! Allocates 15% for Q1-PQ; 30% for Q2; 35% for Q3-Default Queue; 5% for Q4
mls qos queue-set output 1 threshold 1 100 100 100 100
! No WTD thresholds for PQ; reserve full 100% of buffers; no need to borrow more
mls qos queue-set output 1 threshold 2 80 90 100 400
! Tunes Q2T1 for Call Signaling (CS); Tunes Q2T2 for Network Control (NC);
! reserve full 100% of buffers; borrow max as needed
```

```
mls qos queue-set output 1 threshold 3 100 100 100 400
! No WTD thresholds in BE queue (all packets have same CoS/DSCP weight of 0);
! reserve full 100%; borrow max as needed
mls qos queue-set output 1 threshold 4 60 80 100 400
! Tunes Q4T1 for Scavenger (SCV); Tunes Q4T2 for Bulk Data (BD)

! Maps CoS to Egress Queues (although we're not trusting CoS with APIC-EM;
! but for the sake of a comprehensive policy, this is included)
mls qos srr-queue output cos-map queue 1 threshold 3 4 5
! Maps Real-Time Interactive (RTI) and MM_CONF (MMC) and
! Broadcast Video (BVI) and Voice (VO)to PQ
mls qos srr-queue output cos-map queue 2 threshold 1 2
! Maps Operations-Administration-Management (OAM) and Transactional Data (TD) to
Q2T1
mls qos srr-queue output cos-map queue 2 threshold 2 3
! Maps Call Signaling (CS) + Multimedia Streaming (MMS) to Q2T2
mls qos srr-queue output cos-map queue 2 threshold 3 6 7
! Maps Network Control (NC) to Q2T3
mls qos srr-queue output cos-map queue 4 threshold 3 1
! Maps BD + SCV to Q4 (tail)

! Maps DSCP to Egress Queues
mls qos srr-queue output DSCP-map queue 1 threshold 3 32 40 46
! Maps RTI + BV + VO (DSCP EF)
mls qos srr-queue output DSCP-map queue 2 threshold 1 16 18 20 22
mls qos srr-queue output DSCP-map queue 2 threshold 1 26 28 30 34 36 38
! Maps MMS + TD + MMC to Q2T1
mls qos srr-queue output DSCP-map queue 2 threshold 2 24
! Maps CS to Q2T2
mls qos srr-queue output DSCP-map queue 2 threshold 3 48 56
! Maps NC to Q2T3 (Per RFC 4594 NC = DSCP CS6/48; but this class also includes
! CS7/56, which Cisco uses for internal DSCP of spanning tree & other protocols)
```

```
mls qos srr-queue output DSCP-map queue 3 threshold 3 0 1 2 3 4 5 6 7
! Maps BE + non-standard DSCPs to Q3 (tail)
mls qos srr-queue output DSCP-map queue 3 threshold 3 9 11 13 15
mls qos srr-queue output DSCP-map queue 3 threshold 3 17 19 21 23
mls qos srr-queue output DSCP-map queue 3 threshold 3 25 27 29 31
mls qos srr-queue output DSCP-map queue 3 threshold 3 33 35 37 39
mls qos srr-queue output DSCP-map queue 3 threshold 3 41 42 43 44 45 47
mls qos srr-queue output DSCP-map queue 3 threshold 3 49 50 51 52 53 54 55
mls qos srr-queue output DSCP-map queue 3 threshold 3 57 58 59 60 61 62 63
! Maps non-standard DSCPs to Q3 (tail)
mls qos srr-queue output DSCP-map queue 4 threshold 1 8 14
! Maps SCV + BD (AF13) to Q4T1
mls qos srr-queue output DSCP-map queue 4 threshold 2 12
! Maps Bulk (AF12) to Q4T2
mls qos srr-queue output DSCP-map queue 4 threshold 3 10
! Maps Bulk (AF11) to Q4 (tail)
!
```

The above configuration maps CoS 4 and 5, as well as DSCP values 46 (EF), 40 (CS5), and 32 (CS4) to queue 1, threshold 3 (Q1T3). By default, drop threshold 3 is set for 100% of the queue depth. Queue 1 is configured as a strict priority queue within the interface configuration show below. Queue 1 is allocated approximately 15% of the buffers.

CoS 7 and 6, as well as DSCP values 48 (CS6) and 56 (CS7) are mapped to queue 2, threshold 3 (Q2T3), because these are considered to be control traffic. Again, by default, drop threshold 3 is set for 100% of the queue depth. CoS 3, as well as DSCP value 24 (CS3) is mapped to queue 2, threshold 2 (Q2T2) with a drop threshold of 90% of the queue depth. CoS 2, as well as DSCP values 34 (AF41), 36 (AF42), 38 (AF43), 26 (AF31), 28 (AF32), 30 (AF30), 18 (AF21), 20 (AF22), 22 (AF23), and 16 (CS2) are mapped to queue 2 threshold 1 (Q2T1) with a drop threshold of 80% of the queue depth. Queue 2 is allocated approximately 30% of the buffers.

DSCP values 8 (CS1) and 14 (AF13) are mapped to queue 4 threshold 1 (Q4T1) with a drop threshold of 60%. DSCP value 12 (AF12) is mapped to queue 4 threshold 2 (Q4T2) with a drop threshold of 80%. CoS 1, as well as DSCP value 10 (AF11) is mapped to queue 4 threshold 3 (Q4T3) with a default drop threshold of 100%. Queue 4 is allocated approximately 20% of the buffers.

Finally all other CoS and DSCP values are mapped to queue 3, threshold 3 (Q3T3) with a default drop threshold of 100% of the queue depth. Queue 1 is configured for approximately 35% of the buffers.

In the configuration above, WTD is only applied to queues 2 and 4. Queue 1 is a priority queue, for real-time multimedia traffic, hence WTD is not necessary. Queue 4 is the queue with default traffic, hence WTD is not necessary there either.

The 1P3Q3T egress queuing structure is applied by EasyQoS to all Gigabit Ethernet and TenGigabitEthernet interfaces. An example of the provisioning to a single Gigabit Ethernet and single TenGigabitEthernet interface is shown below.

```
!
interface GigabitEthernetx/x/x
 no mls qos trust
! Default setting for access port interfaces (not explicitely provisioned by
EasyQoS)
 queue-set 1
! Default queue set for access port interfaces (not explicitely provisioned by
EasyQoS)
 srr-queue bandwidth share 1 30 35 5
! Sets the remaining queues to share the remaining bandwidth in a ratio of
! 30, 35, and 5
! Queue 1 bandwidth ratio is ignored when priority queueing is enabled
 priority-queue out
! Implements egress priority queuing for queue 1


interface TenGigabitEthernetx/x/x
 mls qos trust DSCP
! Explicitly sets access interface to trust DSCP markings
 queue-set 1
! Default queue set for access port interfaces (not explicitely provisioned by
EasyQoS)
 srr-queue bandwidth share 1 30 35 5
! Sets the remaining queues to share the remaining bandwidth in a ratio of
! 30, 35, and 5
! Queue 1 bandwidth ratio is ignored when priority queueing is enabled
 priority-queue out
! Implements egress priority queuing for queue 1
!
```

Within the interface configuration, the priority queue is specified. By default queue 1 is the priority queue when configured. The bandwidth allocation is shared–meaning that any given queue can use more than its allocated share if one or more of the other queues is not currently using their share and bandwidth is available. Queue 1 is a strict priority queue, so the bandwidth share is ignored for that queue. The remaining queues are allocated bandwidth in a ratio of 30 to 35 to 5. Gigabit Ethernet interfaces that connect to end-devices are set not to trust DSCP or CoS markings. TenGigabitEthernet interfaces that represent uplink interfaces are explicitly set to trust DSCP.

When implementing an EtherChannel connection on Catalyst 2960-S, 2960-X, and 2960-XR Series switches, queuing policies are applied to the physical interfaces. However, classification & marking policies are applied to the logical port-channel associated with the physical interfaces that make up the EtherChannel group. However, because these switches are only supported in the role of an access-layer switch for the EasyQoS solution, no distribution-layer ingress classification & marking policies are ever applied to the logical port-channel associated with the physical interfaces that make up the EtherChannel group.

An example of the configuration pushed by EasyQoS to a Catalyst 2960-S, 2960-X, and 2960-XR, Series switch when operating as an access-layer switch with EtherChannel connectivity to the distribution-layer switch is shown below.

```
!
interface Port-channelx
!
interface TenGigabitEthernety/y/y
 srr-queue bandwidth share 1 30 35 5
 priority-queue out
 mls qos trust DSCP
 channel-group x mode auto
!
```

Note that this configuration is no different than the configuration shown earlier, when EtherChannel is not implemented. EtherChannel interfaces could be Gigabit Ethernet interfaces or FastEthernet interfaces as well, depending upon the type of interface supported by the switch platform. A TenGigabitEthernet interface was shown only as an example in the configuration above.

Catalyst 3560-X/3750-X Queuing Design

This section discusses the ingress and egress queuing structures pushed by APIC-EM to the ports of the Catalyst 3750-X and 3560-X Series switches supported by EasyQoS.

Catalyst 3750-X and 3560-X Series platforms run one of the following three feature sets, all of which support QoS:

- IP Services Feature Set
- IP Base Feature Set (IP Base or Universal on the Catalyst 3560-X Series)
- LAN Base Feature Set

Catalyst 3750-X and 3560-X Series platforms are similar to Catalyst 2960 Series platforms but support Layer 2 and Layer 3 switching, as well as ingress queuing.

Catalyst 3750-X and 3560-X Series platforms are MLS QoS based platforms, which require QoS to be enabled globally first before configuring any other QoS commands. The following global-level commands configured by EasyQoS enable QoS and set the internal COS-to-DSCP mapping table within the platform.

```
!
mls qos
! Globally Enables QoS
mls qos map cos-DSCP 0 8 16 24 32 46 48 56
! Maps CoS 5 to 46 (rest are default)
!
```

Catalyst 3750-X and 3560-X Series switches support both ingress and egress queuing. Ingress queuing is a 1P1Q3T queuing structure. Egress queuing is a 1P3Q3T queuing structure. Both the ingress and egress queuing structures implement DSCP-to-queue mapping and WTD for congestion avoidance, instead of DSCP-based WRED. Although there are three WTD thresholds, only two are configurable. The third threshold is by default set to the depth of the queue (that is, 100% queue depth).

Note: Ingress queuing is not implemented by EasyQoS for the Catalyst 3560-X Series platform in APIC-EM version 1.3.

Ingress Queuing

The following figure shows the 1P1Q3T ingress queuing model deployed by EasyQoS.

Figure 54 1P1Q3T Ingress Queuing for the Catalyst 3750-X and 3560-X Series Platforms

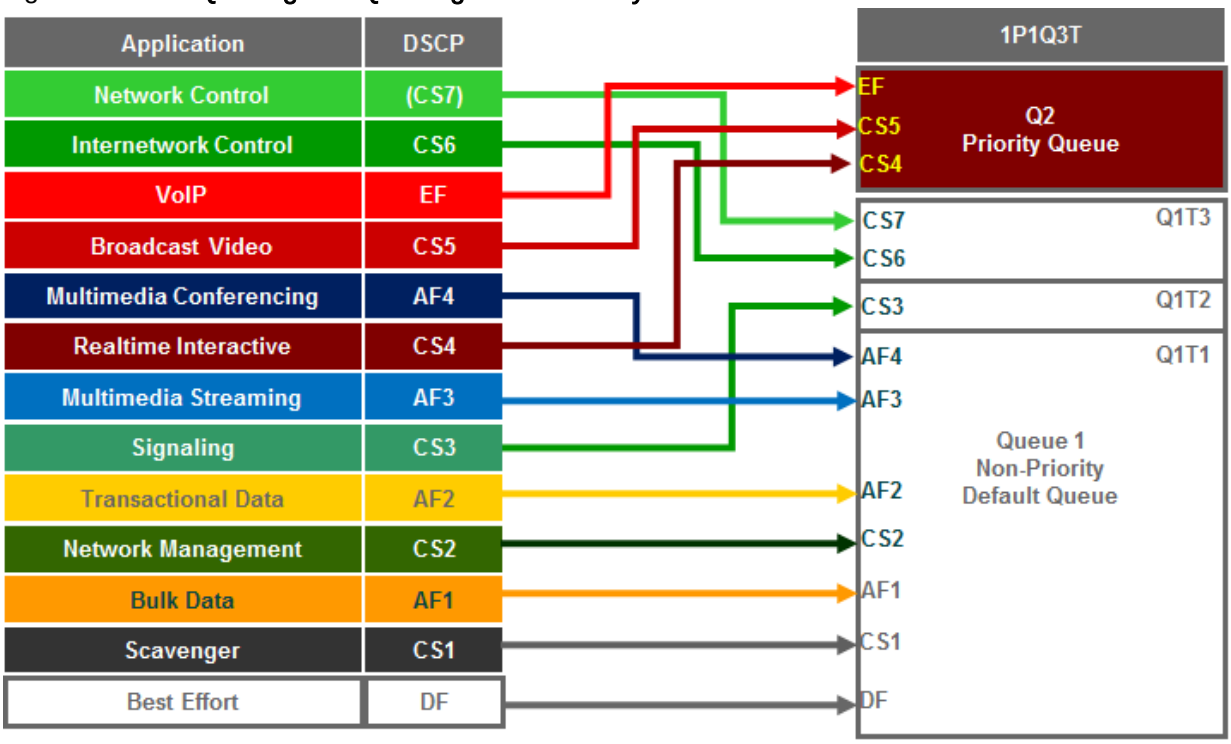

The following configuration implements the 1P1Q3T ingress queuing structure deployed by EasyQoS.

```
!
! This section configures the ingress queues and thresholds
mls qos srr-queue input priority-queue 2 bandwidth 30
! Q2 is enabled as a strict-priority ingress queue with 30% BW
mls qos srr-queue input bandwidth 70 30
! Q1 is assigned 70% BW via SRR shared weights—Q2 SRR shared weight
! is ignored (as it has been configured as a PQ)
mls qos srr-queue input buffers 90 10
! Q1 is assigned 90% of queuing buffers and Q2 (PQ) is assigned 10%
mls qos srr-queue input threshold 1 80 90
! Q1 thresholds are configured at 80% (Q1T1) and 90% (Q1T2)

! This section configures the ingress CoS-to-Queue mappings
mls qos srr-queue input cos-map queue 1 threshold 1 0 1 2
! CoS values 0, 1 and 2 are mapped to Q1T1
mls qos srr-queue input cos-map queue 1 threshold 2 3
```

```
! CoS value 3 is mapped to ingress Q1T2
mls qos srr-queue input cos-map queue 1 threshold 3 6 7
! CoS values 6 and 7 are mapped to ingress Q1T3
mls qos srr-queue input cos-map queue 2 threshold 1 4 5
! CoS values 4 and 5 are mapped to ingress Q2 (the PQ)

! This section configures ingress DSCP-to-Queue Mappings
mls qos srr-queue input DSCP-map queue 1 threshold 1 0 8 10 12 14
mls qos srr-queue input DSCP-map queue 1 threshold 1 16 18 20 22
mls qos srr-queue input DSCP-map queue 1 threshold 1 26 28 30 34 36 38
! DSCP CS2 and AF2 are mapped to ingress Q1T1
mls qos srr-queue input DSCP-map queue 1 threshold 2 24
! DSCP CS3 is mapped to ingress Q1T2
mls qos srr-queue input DSCP-map queue 1 threshold 3 48 56
! DSCP CS6 and CS7 are mapped to ingress Q1T3 (the tail of Q1)
mls qos srr-queue input DSCP-map queue 2 threshold 3 32 40 46
! DSCP CS4, CS5 and EF are mapped to ingress Q2T3 (the tail of the PQ)
!
```

The above configuration maps CoS 4 and 5, as well as DSCP values 46 (EF), 40 (CS5), and 32 (CS4) to queue 2, threshold 3 (Q2T3). By default, drop threshold 3 is set for 100% of the queue depth. Queue 2 is configured as a strict priority queue but is limited to 30% of the bandwidth and 10% of the buffers.

CoS 7 and 6, as well as DSCP values 48 (CS6) and 56 (CS7) are mapped to queue 1, threshold 3 (Q1T3), because these are considered to be control traffic. Again, by default, drop threshold 3 is set for 100% of the queue depth. CoS 3, as well as DSCP value 24 (CS3) is mapped to queue 1, threshold 2 (Q1T2) with a drop threshold of 90% of the queue depth. Finally all other CoS and DSCP values are mapped to queue 1, threshold 1 (Q1T1) with a drop threshold of 80% of the queue depth. Queue 1 is configured for approximately 70% of the bandwidth (because queue 2 is limited to 30% of the bandwidth) and 90% of the buffers, because Q2 is a strict priority queue and will be serviced first.

Egress Queuing

The 1P3Q3T egress queueing model deployed by EasyQoS for the Catalyst 3750-X and 3560-X Series switches is the same as was shown in Figure 53 and discussed in the *1P3Q3T Egress Queuing Design* section of this document.

Catalyst 3650/3850 Queuing Design

This section discusses the egress queuing structure pushed by EasyQoS to the ports Catalyst 3850 and 3650 Series switches.

Egress Queuing

Catalyst 3850 and 3650 Series switches support only egress queuing. The 2P6Q3T egress queueing structure for the Catalyst 3850 and 3650 Series implements DSCP-to-queue mapping with WTD for congestion avoidance. The following figure shows the 2P6Q3T egress queuing model.

Figure 55 **2P6Q3T Egress Queuing for the Catalyst 3650 and 3850 Series Switches**

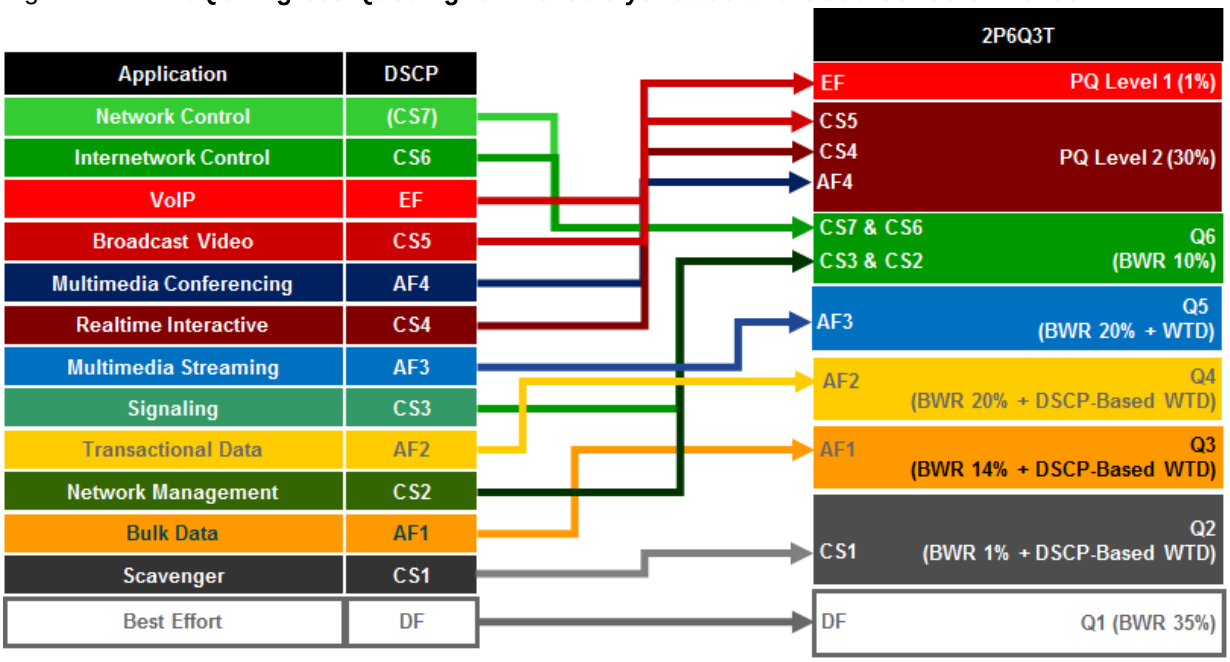

The following configuration, provisioned by EasyQoS, implements the class-maps for the 2P6Q3T egress queuing structure.

```
!
class-map match-any prm-EZQOS_2P6Q3T#VOICE-PQ1
 match DSCP ef
class-map match-any prm-EZQOS_2P6Q3T#VIDEO-PQ2
 match DSCP cs4
 match DSCP af41
 match DSCP af42
 match DSCP af43
 match DSCP cs5
class-map match-any prm-EZQOS_2P6Q3T#CONTROL-PLANE
 match DSCP cs2
 match DSCP cs3
 match DSCP cs6
 match DSCP cs7
```

```
class-map match-any prm-EZQOS_2P6Q3T#MULTIMEDIA-STREAMING
 match DSCP af31
 match DSCP af32
 match DSCP af33
class-map match-any prm-EZQOS_2P6Q3T#TRANSACTIONAL-DATA
 match DSCP af21
 match DSCP af22
 match DSCP af23
class-map match-any prm-EZQOS_2P6Q3T#BULK-DATA
 match DSCP af11
 match DSCP af12
 match DSCP af13
class-map match-any prm-EZQOS_2P6Q3T#SCAVENGER
 match DSCP cs1
!
```

The following configuration, provisioned by EasyQoS, implements the policy-map for the 2P6Q3T egress queuing structure.

```
!
policy-map prm-DSCP#APIC_QOS_Q_OUT
 class prm-EZQOS_2P6Q3T#VOICE-PQ1
  priority level 1 percent 1
  queue-buffers ratio 5
 class prm-EZQOS_2P6Q3T#VIDEO-PQ2
  priority level 2 percent 30
  queue-buffers ratio 5
 class prm-EZQOS_2P6Q3T#CONTROL-PLANE
  bandwidth remaining percent 10
  queue-buffers ratio 5
 class prm-EZQOS_2P6Q3T#MULTIMEDIA-STREAMING
  bandwidth remaining percent 20
  queue-buffers ratio 10
  queue-limit DSCP af31 percent 100
```

```
  queue-limit DSCP af32 percent 90
  queue-limit DSCP af33 percent 80
 class prm-EZQOS_2P6Q3T#TRANSACTIONAL-DATA
  bandwidth remaining percent 20
  queue-buffers ratio 10
  queue-limit DSCP af21 percent 100
  queue-limit DSCP af22 percent 90
  queue-limit DSCP af23 percent 80
 class prm-EZQOS_2P6Q3T#BULK-DATA
  bandwidth remaining percent 14
  queue-buffers ratio 20
  queue-limit DSCP af11 percent 100
  queue-limit DSCP af12 percent 90
  queue-limit DSCP af13 percent 80
 class prm-EZQOS_2P6Q3T#SCAVENGER
  bandwidth remaining percent 1
  queue-buffers ratio 5
 class class-default
  bandwidth remaining percent 35
  queue-buffers ratio 40
!
```

The 2P6Q3T egress queuing structure is applied by EasyQoS to all Gigabit Ethernet and TenGigabitEthernet interfaces. An example of the provisioning to a single Gigabit Ethernet and single TenGigabitEthernet interface is shown below.

```
!
interface GigabitEthernetx/x/x
 service-policy output prm-DSCP#APIC_QOS_Q_OUT
interface TenGigabitEthernetx/x/x
 service-policy output prm-DSCP#APIC_QOS_Q_OUT
!
```

The Catalyst 3850 and 3650 Series platforms support two strict priority queues. DSCP value EF (voice traffic) is assigned to the first priority queue (prm-EZQOS_2P6Q3T#VOICE-PQ1) which is allocated 1% of the bandwidth. 1% of the bandwidth may at first seem somewhat low. However, 1% of a 1 Gbps interface is

10 Mbps of bandwidth. Likewise, 1% of a 10 Gbps interface is 100 Mbps of bandwidth. Individual voice calls average under 100 Kbps, even with the overhead of Layer 2-4 headers. Hence, 1% of the bandwidth is sufficient for approximately 100 voice calls on a 1 Gbps interface and 1,000 voice calls on a 10 Gbps interface.

DSCP values CS5, CS4, AF41, AF42, and AF43 (broadcast video, real-time interactive video, and multimedia conferencing traffic) are assigned to the second priority queue (prm-EZQOS_2P6Q3T#VIDEO-PQ2), which is allocated 30% of the bandwidth. Both of these queues are allocated only approximately 5% of the buffers because they are priority queues. PQ1 is serviced first. If there are no packets in PQ1, then PQ2 is serviced. All other queues are serviced after the priority queues are serviced.

DSCP values CS6, CS7, CS3, and CS2 are mapped to the prm-EZQOS_2P6Q3T#CONTROL-PLANE queue, which is allocated 10% of the remaining bandwidth, after the priority queues are serviced. Because this queue holds control traffic (CS6 and CS7), signaling traffic (CS3), and OAM traffic (CS2), which is not expected to be a lot of traffic, approximately 5% of the buffers are allocated to the queue. WTD thresholds are not implemented for this queue because the objective is not to drop any of this traffic.

DSCP values AF31, AF32, and AF33 are mapped to the prm-EZQOS_2P6Q3T#MULTIMEDIA-STREAMING queue, which is allocated approximately 20% of the remaining bandwidth, after the priority queues are serviced. Because AF classes are specifically intended for marking down traffic, WTD is implemented within this traffic class. AF33 traffic is set with a drop threshold of 80% of the buffer depth, AF32 traffic is set with a drop threshold of 90% of the buffer depth, and AF31 traffic is set with a drop threshold of 100% of the buffer depth. Approximately 10% of the buffers are allocated to the queue.

DSCP values AF21, AF22, and AF23 are mapped to the prm-EZQOS_2P6Q3T#TRANSACTIONAL-DATA queue, which is allocated approximately 20% of the remaining bandwidth, after the priority queues are serviced. Because AF classes are specifically intended for marking down traffic, WTD is implemented within this traffic class. AF23 traffic is set with a drop threshold of 80% of the buffer depth, AF22 traffic is set with a drop threshold of 90% of the buffer depth, and AF21 traffic is set with a drop threshold of 100% of the buffer depth. Approximately 10% of the buffers are allocated to the queue.

DSCP values AF11, AF12, and AF13 are mapped to the prm-EZQOS_2P6Q3T#BULK-DATA queue, which is allocated approximately 14% of the remaining bandwidth, after the priority queues are serviced. Because AF classes are specifically intended for marking down traffic, WTD is implemented within this traffic class. AF13 traffic is set with a drop threshold of 80% of the buffer depth, AF12 traffic is set with a drop threshold of 90% of the buffer depth, and AF11 traffic is set with a drop threshold of 100% of the buffer depth. Approximately 20% of the buffers are allocated to the queue.

DSCP value CS1 is mapped to the prm-EZQOS_2P6Q3T#SCAVENGER queue, which is allocated approximately 1% of the remaining bandwidth, after the priority queues are serviced. This is a scavenger class, specifically intended for traffic that has a business relevancy of business-irrelevant. WTD is not implemented within this traffic class and approximately 5% of the buffers are allocated to the queue.

All other DSCP values are by default mapped to the class-default queue, which is allocated approximately 35% of the remaining bandwidth, after the priority queues are serviced. The default class is specifically intended for traffic that has a business relevancy of default. WTD is not implemented within this traffic class, and approximately 40% of the buffers are allocated to the queue.

There are no mappings of CoS values to queues with the Catalyst 3850 and 3650 Series platforms. These platforms only fall back to the use of CoS values for non-IP packets.

When implementing an EtherChannel connection on Catalyst 3850 and 3650 Series platforms, both queuing policies and classification & marking policies are applied to the physical interfaces that make up the EtherChannel group. An example of the configuration pushed by EasyQoS to a Catalyst 3850 or 3650 Series switch when operating as a distribution-layer switch with EtherChannel connectivity the access-layer switch is shown below.

```
!
interface Port-channelx
!
interface TenGigabitEthernety/y/y
 channel-group x mode auto
 service-policy input APIC_EM-MARKING-DIST-IN
 service-policy output prm-DSCP#APIC_QOS_Q_OUT
!
```

Catalyst 4500 Queuing Design

This section discusses the egress queuing structure pushed by APIC-EM to the ports of each of the line cards and supervisors supported by EasyQoS for the Catalyst 4500-E Series and Catalyst 4500-X Series switches.

Catalyst 4500-E Series switches with Supervisor 7-E, 7-LE, 8-E, and 8-LE; and Catalyst 4500-X Series switches support only egress queuing. The 1P7Q1T egress queueing structure for the Catalyst 4500-E and Catalyst 4500-X Series implements DSCP-to-queue mapping and Dynamic Buffer Limiting (DBL) for congestion avoidance instead of DSCP-based WRED or Weighted Tail Drop (WTD). The following figure shows the 1P7Q1T egress queueing model.

Figure 56 **1P7Q1T+DBL Egress Queuing for the Catalyst 4500 Series**

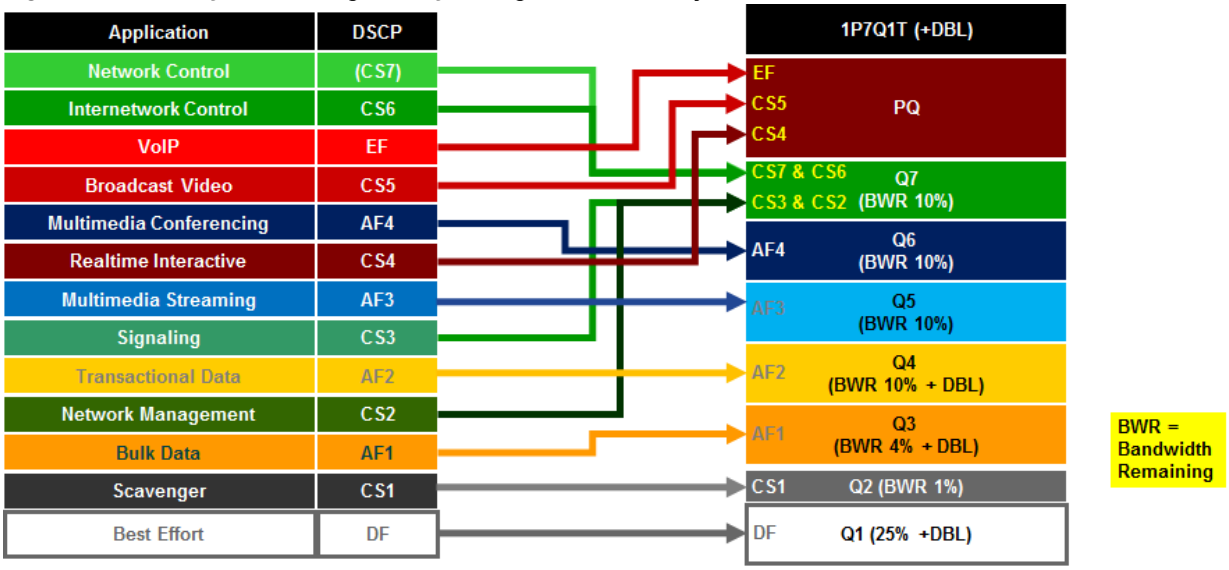

The following configuration, provisioned by APIC-EM EasyQoS, implements the class-maps for the 1P7Q1T egress queuing structure.

```
!
class-map match-any prm-EZQOS_1P7Q1T#REALTIME
  match  DSCP cs4
  match  DSCP cs5
  match  DSCP ef
class-map match-any prm-EZQOS_1P7Q1T#CONTROL
  match  DSCP cs2
  match  DSCP cs3
  match  DSCP cs6
  match  DSCP cs7
class-map match-any prm-EZQOS_1P7Q1T#MM_CONF
  match  DSCP af41
  match  DSCP af42
  match  DSCP af43
class-map match-any prm-EZQOS_1P7Q1T#MM_STREAM
  match  DSCP af31
  match  DSCP af32
  match  DSCP af33
class-map match-any prm-EZQOS_1P7Q1T#TRANS_DATA
  match  DSCP af21
  match  DSCP af22
  match  DSCP af23
class-map match-any prm-EZQOS_1P7Q1T#BULK_DATA
  match  DSCP af11
  match  DSCP af12
  match  DSCP af13
class-map match-any prm-EZQOS_1P7Q1T#SCAVENGER
  match  DSCP cs1
!
```

The following configuration, provisioned by APIC-EM EasyQoS, implements the policy-map for the 1P7Q1T egress queuing structure.

```
!
policy-map prm-DSCP#APIC_QOS_Q_OUT
 class prm-EZQOS_1P7Q1T#REALTIME
    priority
 class prm-EZQOS_1P7Q1T#CONTROL
    bandwidth remaining percent 10
 class prm-EZQOS_1P7Q1T#MM_CONF
    bandwidth remaining percent 15
 class prm-EZQOS_1P7Q1T#MM_STREAM
    bandwidth remaining percent 15
 class prm-EZQOS_1P7Q1T#TRANS_DATA
    bandwidth remaining percent 15
    dbl
 class prm-EZQOS_1P7Q1T#BULK_DATA
    bandwidth remaining percent 14
    dbl
 class prm-EZQOS_1P7Q1T#SCAVENGER
    bandwidth remaining percent 1
 class class-default
    bandwidth remaining percent 25
    dbl
!
```

The 1P7Q1T egress queuing structure is applied by EasyQoS to all Gigabit Ethernet and TenGigabitEthernet interfaces. An example of the provisioning to a single Gigabit Ethernet and single TenGigabitEthernet interface is shown below.

```
!
interface GigabitEthernetx/x-x
 service-policy output prm-DSCP#APIC_QOS_Q_OUT
!
interface TenGigabitEthernetx/x/x
 service-policy output prm-DSCP#APIC_QOS_Q_OUT
```

```
!
```

In the configuration above, DBL is only applied to the transactional data, bulk data, and default queues. DBL has a congestion-avoidance function, similar to WRED in that it can help prevent synchronization of TCP flows that result in under-utilization of the available bandwidth. DBL is not recommended to be deployed on real-time multimedia and control queues. DBL is not deployed on the scavenger queue, simply because the traffic is already considered to be scavenger traffic using whatever bandwidth is available.

When implementing an EtherChannel connection on Catalyst 4500-E and Catalyst 4500-X Series platforms, queuing policies are applied to the physical interfaces. However, classification & marking policies are applied to the logical port-channel associated with the physical interfaces that make up the EtherChannel group. An example of the configuration pushed by EasyQoS to a Catalyst 4500-E or Catalyst 4500-X series switch when operating as a distribution-layer switch with EtherChannel connectivity the access-layer switch is shown below.

```
!
interface Port-channelx
 service-policy input prm-APIC_QOS_IN
!
interface range TenGigabitEthernetx/x/x–xx
 channel-group x mode auto
 service-policy output prm-DSCP#APIC_QOS_Q_OUT
!
```

Catalyst 6500 Sup2T Queuing Design

The following sections discussing the ingress and egress queuing structures provisioned by APIC-EM EasyQoS to the ports of each of the line cards and supervisors supported by EasyQoS for the Catalyst 6500 Series platform with a Sup2T supervisor.

Ingress & egress queueing structures are dependent upon the following:

- The model of the line card
- Whether the line card supports a Centralized Forwarding Cart (CFC) or Distributed Forwarding Cart (DFC). This applies to WS-X6704, WS-X6724, and WS-X6748 series line cards.
- Whether the Sup2T supervisor (VS-2T-10G or VS-2T-10G-XL) Gigabit Ethernet ports are enabled or disabled

For the Catalyst 6500 Series and Catalyst 6807-XL switch with Supervisor 2T (Sup2T), if an unsupported line card is detected within the chassis, the behavior of APIC-EM EasyQoS is to skip over the line card–meaning not provision an QoS configuration for the unsupported line card–and attempt to continue provisioning the rest of the platform. The network operator will also be notified via the EasyQoS web-based GUI that an unsupported line card was detected within the chassis.

Note: Catalyst 6K and 4K switches are supported in VSS and non-VSS configurations. When operating in a VSS configuration, switch ports that belong to a port-group, which in turn are part of the VSL between the individual switches in the VSS configuration, cannot be configured with any QoS policy. APIC-EM EasyQoS has the ability to identify ports that are part of a VSL and not apply QoS policy.

1Q8T Ingress Queuing

1Q8T ingress queuing is supported by the following line cards:

- WS-X6704-10GE with CFC
- WS-X6724-SFP with CFC
- WS-X6748-SFP and WS-X6748-GE-TX with CFC

The WS-X6724-SFP, WS-X6748-SFP, WS-X6748-GE-TX, and WS-X6704-10GE line cards are supported in the Catalyst 6500 Series or 6807-XL with Sup-2T with either a CFC or DFC version 4 or 4-XL upgrade. The DFC is daughter card that sits on the line card itself.

Individual line card models can be identified within Catalyst 6500 Series or Catalyst 6807-XL switches via the “show module” exec-level command. An example of the output of the “show module” command is shown below, with a WS-X6748-GE-TX line card with a CFC in slot 1 highlighted.

```
o23-6500-1#show module
Mod Ports Card Type                                Model              Serial No.
--- ----- ---------------------------------------- ------------------ -----------
  1   48  CEF720 48 port 10/100/1000mb Ethernet  WS-X6748-GE-TX     SAL10478SWP
  2    8  DCEF2T 8 port 10GE                      WS-X6908-10G       SAL172682AK
  3    5  Supervisor Engine 2T 10GE w/ CTS (Acti VS-SUP2T-10G       SAL1702WNR0
  5   16  CEF720 16 port 10GE                     WS-X6716-10GE      SAL1228WYB7
  6    4  CEF720 4 port 10-Gigabit Ethernet       WS-X6704-10GE      SAL15013XBH
…
Mod  Sub-Module                  Model              Serial       Hw     Status
---- --------------------------- ------------------ ----------- ------- -------
  1  Centralized Forwarding Card WS-F6700-CFC       SAD074308C9  1.1    Ok
  2  Distributed Forwarding Card WS-F6K-DFC4-E      SAL17152T2R  1.2    Ok
  3  Policy Feature Card 4       VS-F6K-PFC4        SAL1638N3R3  1.2    Ok
  3  CPU Daughterboard           VS-F6K-MSFC5       SAL1702WNG1  1.5    Ok
  5  Distributed Forwarding Card WS-F6K-DFC4-E      SAL1541SQHX  1.1    Ok
  6  Centralized Forwarding Card WS-F6700-CFC       SAL1518CRZ3  4.1    PwrDown
```

Note: For Catalyst 6500 or 6800 Series switches in a VSS configuration, the exec-level command “show module switch all” can be used to display modules in both switches.

1Q8T ingress queueing for these line cards implements CoS-to-queue mapping, with CoS-based tail-drop for congestion avoidance. The following figure shows the 1Q8T ingress queueing model.

Figure 57 **1Q8T Ingress Queuing Models—CoS-to-Queue Mapping with CoS-Based Tail-Drop**

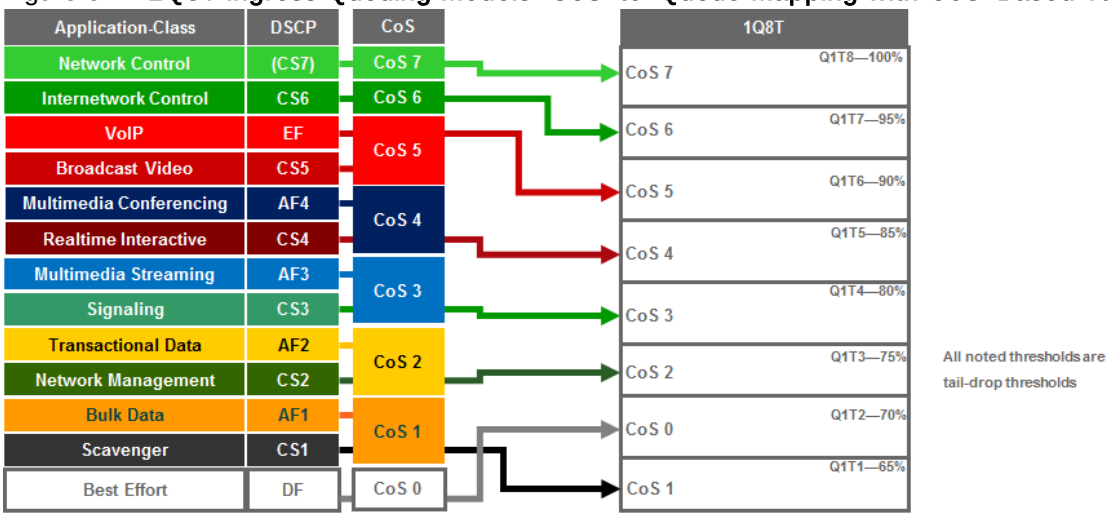

Because the 1Q8T ingress queuing structure only supports one queue, there are no class-map definitions. All traffic is mapped to class-default, corresponding to the default queue. The following configuration, provisioned by APIC-EM EasyQoS, implements the policy-map for the 1Q8T ingress queuing structure.

```
!
policy-map type lan-queuing prm-DSCP#EZQOS_1Q8T-IN
  class class-default
    queue-limit cos 0 percent 70
    queue-limit cos 1 percent 65
    queue-limit cos 2 percent 75
    queue-limit cos 3 percent 80
    queue-limit cos 4 percent 85
    queue-limit cos 5 percent 90
    queue-limit cos 6 percent 95
    queue-limit cos 7 percent 100
!
```

The 1Q8T ingress queuing structure is applied by EasyQoS to all Gigabit Ethernet and TenGigabitEthernet interfaces on the line card that support this queuing structure. An example of the provisioning to a single Gigabit Ethernet interface is shown below.

```
!
interface GigabitEthernetx/x
 service-policy type lan-queuing input prm-DSCP#EZQOS_1Q8T-IN
!
```

For the WS-X6724-SFP, WS-X6748-SFP, and WS-X6748-GE-TX line cards with CFCs, the ports are meant to be connected to end-user devices. Because line cards that support the 1Q8T queueing structure support only CoS-to-queue mapping, if the end-device does not send traffic with an 802.1p header, all traffic is treated with the default CoS mapping for the port (CoS 0) and mapped to Q1T2 with a tail-drop threshold of 70%.

Note, however, that interfaces may be configured with separate VLANs for voice, video, etc. by the network operator. If the end-device does send traffic with an 802.1p header–such as a Cisco IP phone that sends traffic marked as CoS 5–then ingress traffic from the IP phone will be mapped to Q1T6 with a tail-drop threshold of 90%, while a device chained-off the IP phone may have its traffic remarked to CoS 0 by the IP phone and mapped to Q1T2 with a tail-drop threshold of 70%.

Ports on the WS-X6704-10GE with CFC, are assumed to be uplink ports. If the Catalyst 6500 Series or Catalyst 6807-XL with Sup-2T is deployed as a distribution or core switch, and the link connecting the WS-X6704-10GE switch port is not a trunk port, then all ingress traffic will not have an 802.1p header. Hence all ingress traffic will be treated with the default CoS mapping for the port (CoS 0) and mapped to Q1T2 with a tail-drop threshold of 70%.

Due to the internal ASIC structure of the ports on the WS-X6748-SFP and WS-X6748-GE-TX line cards, the ingress and egress queueing structures of the ports cannot be configured independently. Instead, the queueing policy is applied to groups of ports on the line card by APIC-EM EasyQoS.

2Q4T Ingress Queuing

Note: Due to Cisco defect CSCvb72316, EasyQoS does not currently deploy ingress or egress queuing to switch ports on the VS-S2T-10G and VS-S2T-10G-XL (Supervisor 2T). Hence, this Catalyst 6K Sup2T queuing structure is currently not used within EasyQoS.

2Q4T ingress queuing is supported by the following line cards:

- All ports on the VS-S2T-10G and VS-S2T-10G-XL (Supervisor 2T) when the Gigabit Ethernet ports are enabled.

The Gigabit Ethernet ports on the Sup-2T are enabled with the following global configuration command.

```
!
no platform qos 10g-only
!
```

APIC-EM EasyQoS does not set this command but will look to see if this command has been set by the network operator, in order to determine the correct queuing structure to apply to ports on the Sup2T supervisor. The default setting is for the Gigabit Ethernet ports on the Sup-2T to be enabled, so this command will not appear in the configuration.

The status of whether the Gigabit Ethernet ports are enabled or disabled can be displayed by the network operator via the exec-level “show platform qos module x” command, where “x” refers to the slot with the Sup-2T. An example of the output from the command is shown below:

```
o23-6500-1#show platform qos module 3
  QoS is enabled globally
  Port QoS is enabled globally
  QoS serial policing mode enabled globally
   Distributed Policing is Disabled
   Secondary PUPs are enabled
 QoS Trust state is DSCP on the following interface:
EO0/2 Gi1/1 Gi1/2 Gi1/3 Gi1/4 Gi1/5 Gi1/6 Gi1/7 Gi1/8 Gi1/9
Gi1/10 Gi1/11 Gi1/12 Gi1/13 Gi1/14 Gi1/15 Gi1/16 Gi1/17 Gi1/18 Gi1/19
Gi1/20 Gi1/21 Gi1/22 Gi1/23 Gi1/24 Gi1/25 Gi1/26 Gi1/27 Gi1/28 Gi1/29
Gi1/30 Gi1/31 Gi1/32 Gi1/33 Gi1/34 Gi1/35 Gi1/36 Gi1/37 Gi1/38 Gi1/39
Gi1/40 Gi1/41 Gi1/42 Gi1/43 Gi1/44 Gi1/45 Gi1/46 Gi1/47 Gi1/48 Te2/1
Te2/2 Te2/3 Te2/4 Te2/5 Te2/6 Te2/7 Te2/8 Gi3/1 Gi3/2 Gi3/3
Te3/4 Te3/5 Te5/1 Te5/2 Te5/3 Te5/4 Te5/5 Te5/6 Te5/7 Te5/8
Te5/9 Te5/10 Te5/11 Te5/12 Te5/13 Te5/14 Te5/15 Te5/16 Te6/1 Te6/2
Te6/3 Te6/4 CPP CPP.1 Vl1
  QoS 10g-only mode supported: Yes [Current mode: Off]
  Global Policy-map: ingress[]
```

A setting of “Current mode: off” means that the Gigabit Ethernet ports are enabled.

2Q4T ingress queueing for the Sup-2T implements CoS-to-queue mapping, with CoS-based tail-drop for congestion avoidance. The following figure shows the 2Q4T ingress queueing model.

Figure 58 2Q4T Ingress Queuing Models—CoS-to-Queue Mapping with CoS-Based Tail-Drop

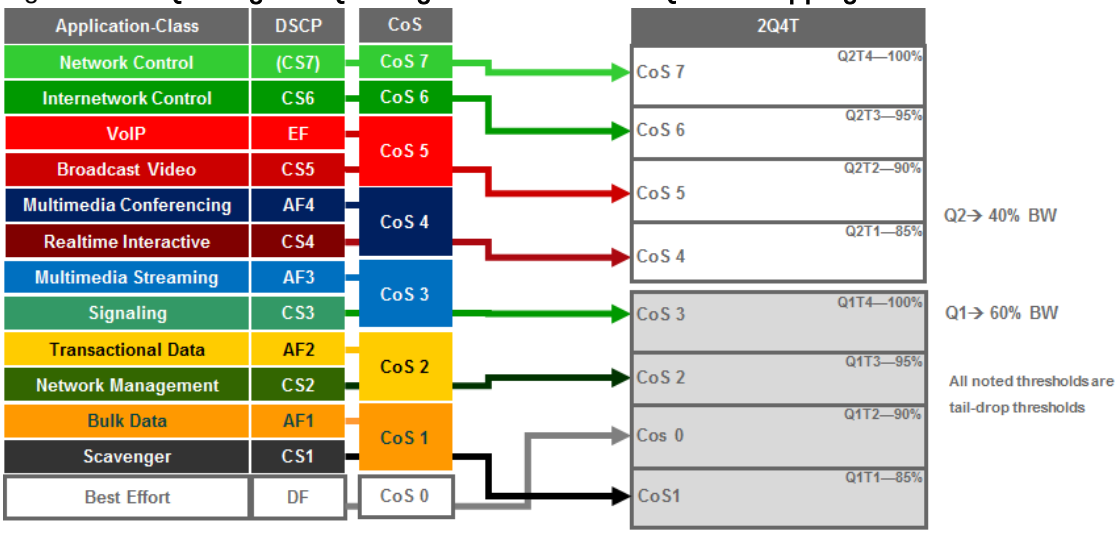

The following configuration, provisioned by APIC-EM EasyQoS, implements the class-map for the 2Q4T ingress queuing structure.

```
!
class-map type lan-queuing match-any prm-EZQOS_2Q4T#Q2
  match cos  7
  match cos  6
  match cos  5
  match cos  4
!
```

The following configuration, provisioned by APIC-EM EasyQoS, implements the policy-map for the 2Q4T ingress queuing structure.

```
!
policy-map type lan-queuing prm-DSCP#EZQOS_2Q4T-IN
  class prm-EZQOS_2Q4T#Q2
    bandwidth percent 40
    queue-limit cos 7 percent 100
    queue-limit cos 6 percent 95
    queue-limit cos 5 percent 90
    queue-limit cos 4 percent 85
  class class-default
    queue-limit cos 0 percent 90
    queue-limit cos 1 percent 85
```

```
    queue-limit cos 2 percent 95
    queue-limit cos 3 percent 100
!
```

The 2Q4T ingress queuing policy-map is provisioned by EasyQoS to all Gigabit Ethernet and TenGigabitEthernet interfaces on the Sup2T, when the Sup2T is configured for this queuing structure, and when the ports are not part of a VSL. An example of the provisioning to a single Gigabit Ethernet interface is shown below.

```
!
interface GigabitEthernetx/x/x
 service-policy type lan-queuing input prm-DSCP#EZQOS_2Q4T-IN
!
interface TenGigabitEthernetx/x/x
 service-policy type lan-queuing input prm-DSCP#EZQOS_2Q4T-IN
!
```

Cisco does not recommend connecting end-user devices to any ports on the Sup-2T. Hence, all ports on the Sup-2T are assumed to be uplink ports by EasyQoS. The Sup-2T ports with 2Q4T queueing structure only support CoS-to-queue mapping. If the Catalyst 6500 Series or Catalyst 6807-XL with Sup-2T is deployed as a distribution or core switch, and the link connecting the Sup-2T port is not a trunk port, then all ingress traffic will not have an 802.1p header. Hence all ingress traffic will be treated with the default CoS mapping for the port (CoS 0) and mapped to Q1T2 with a tail-drop threshold of 90%.

It should be noted that QoS policies cannot be applied to TenGigabitEthernet ports on the Sup-2T when the TenGigabitEthernet ports are part of a port-channel group that is part of a VSL. Therefore EasyQoS will not apply the 2Q4T ingress queuing policy when the switch port is part of a VSL. Switch ports can be identified as part of a VSL based upon the switch configuration. An example of this is shown below.

```
!
interface Port-channel63
 no switchport
 no ip address
 no platform qos channel-consistency
 switch virtual link 1
~
interface TenGigabitEthernet1/3/4
 no switchport
 no ip address
 no cdp enable
```

```
 channel-group 63 mode on
!
interface TenGigabitEthernet1/3/5
 no switchport
 no ip address
 no cdp enable
 channel-group 63 mode on
!
```

Physical interfaces are assigned to a port-group via the “channel-group” interface-level command. Port-channel interfaces are assigned to a VSL via the “switch virtual link” interface-level command, as shown in the figure above.

Note also that QoS policies cannot be applied to a Gigabit Ethernet port on the Sup2T when the TenGigabitEthernet ports on the Sup-2T are part of a port-channel group that is part of a VSL. Hence, EasyQoS will not apply the 2Q4T ingress queuing policy to Gigabit Ethernet ports on the Sup2T when the TenGigabitEthernet ports are part of a VSL.

Finally, due to the internal ASIC structure of the ports on the Sup2T, the ingress and egress queueing structure of the ports cannot be configured independently. Instead, when configuring an ingress or egress queueing policy to one of the interfaces, the same policy will be propagated to the rest of the interfaces on the Sup-2T. Because of the hardware design, EasyQoS applies the same ingress or egress queuing policy on all interfaces of the Sup2T.

2Q8T Ingress Queuing

2Q8T ingress queuing is supported by the following line cards:

- WS-X6724-SFP with DFC4/DFC4XL upgrade (WS-F6k-DFC4-A, WS-F6k-DFC4-AXL)
- WS-X6748-SFP and WS-X6748-GE-TX with DFC4/DFC4XL upgrade (WS-F6k-DFC4-A, WS-F6k-DFC4-AXL)
- WS-X6824-SFP-2T and WS-X6824-SFP-2TXL
- WS-X6848-SFP-2T, WS-X6848-SFP-2TXL, WS-X6848-TX-2T and WS-X6848-TX-2TXL

2Q8T ingress queueing for these line cards implements CoS-to-queue mapping, with CoS-based tail-drop for congestion avoidance.

The following figure shows the 2Q8T ingress queueing model.

Figure 59 2Q8T Ingress Queuing Models—CoS-to-Queue Mapping with CoS-based Tail-Drop

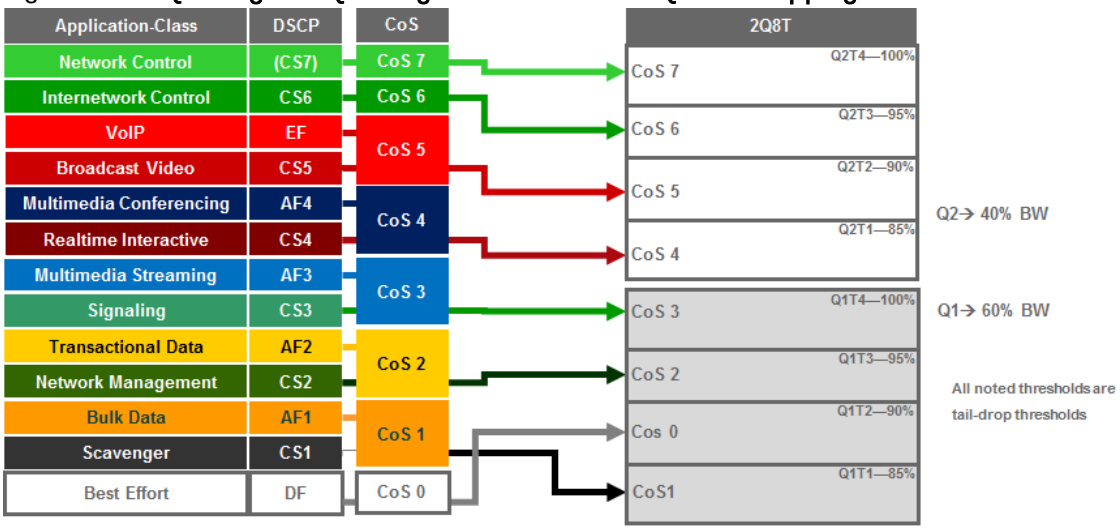

The following configuration, provisioned by APIC-EM EasyQoS, implements the class-maps for the 2Q4T ingress queuing structure.

```
!
class-map type lan-queuing match-any prm-EZQOS_2Q8T#Q2
  match cos  7
  match cos  6
  match cos  5
  match cos  4
!
```

The following configuration, provisioned by APIC-EM EasyQoS, implements the policy-map for the 2Q8T ingress queuing structure.

```
!
policy-map type lan-queuing prm-DSCP#EZQOS_2Q8T-IN
  class prm-EZQOS_2Q8T#Q2
    bandwidth percent 40
    queue-limit cos 4 percent 85
    queue-limit cos 4 percent 90
    queue-limit cos 6 percent 95
    queue-limit cos 7 percent 100
  class class-default
    queue-limit cos 0 percent 90
    queue-limit cos 1 percent 85
```

```
        queue-limit cos 2 percent 95
        queue-limit cos 3 percent 100
    !
```

The 2Q4T ingress queuing policy-map is provisioned by EasyQoS to all Gigabit Ethernet interfaces on the line cards that support this queuing structure. An example of the provisioning to a single Gigabit Ethernet interface is shown below.

```
!
interface GigabitEthernetx/x/x
 service-policy type lan-queuing input prm-DSCP#EZQOS_2Q8T-IN
!
```

For all of the line cards that support the 2Q8T ingress queueing structure listed above, the ports are meant to be connected to end-user devices. Because line cards that support the 2Q8T queuing structure support only CoS-to-queue mapping, if the end-device does not send traffic with an 802.1p header, all traffic is treated with the default CoS mapping for the port (CoS 0) and mapped to Q1T2 with a tail-drop threshold of 90%.

Note, however, that interfaces may be configured with separate VLANs for voice, video, etc. by the network operator. If the end-device does send traffic with an 802.1p header–such as a Cisco IP phone that sends traffic marked as CoS 5–then ingress traffic from the IP phone will be mapped to Q2T2 with a tail-drop threshold of 90%, while a device chained-off the IP phone may have its traffic remarked to CoS 0 by the IP phone and mapped to Q1T2 with a tail-drop threshold of 90%.

8Q4T Ingress Queuing

Note: Due to Cisco defect CSCvb72316, EasyQoS does not currently deploy ingress or egress queuing to switch ports on the VS-S2T-10G and VS-S2T-10G-XL (Supervisor 2T).

8Q4T ingress queuing is supported by the following line cards:

- VS-S2T-10G, VS-S2T-10G-XL with Gigabit Ethernet ports disabled
- WS-X6908-10G-2T, WS-X6908-10G-2TXL

The Gigabit Ethernet ports on the Sup-2T are disabled with the following global configuration command.

```
!
platform qos 10g-only
!
```

APIC-EM EasyQoS does not set this command but will look to see if this command has been set by the network operator, in order to determine the correct queuing structure to apply to ports on the Sup2T supervisor. The default setting is for the Gigabit Ethernet ports on the Sup-2T to be enabled, so this command will appear in the configuration when the Gigabit Ethernet ports are disabled.

The status of whether the Gigabit Ethernet ports are enabled or disabled can be displayed by the network operator via the exec-level “show platform qos module x” command, where “x” refers to the slot with the Sup-2T. An example of the output from the command is shown below:

```
o23-6500-1#show platform qos module 3
  QoS is enabled globally
  Port QoS is enabled globally
  QoS serial policing mode enabled globally
   Distributed Policing is Disabled
   Secondary PUPs are enabled
 QoS Trust state is DSCP on the following interface:
EO0/2 Gi1/1 Gi1/2 Gi1/3 Gi1/4 Gi1/5 Gi1/6 Gi1/7 Gi1/8 Gi1/9
Gi1/10 Gi1/11 Gi1/12 Gi1/13 Gi1/14 Gi1/15 Gi1/16 Gi1/17 Gi1/18 Gi1/19
Gi1/20 Gi1/21 Gi1/22 Gi1/23 Gi1/24 Gi1/25 Gi1/26 Gi1/27 Gi1/28 Gi1/29
Gi1/30 Gi1/31 Gi1/32 Gi1/33 Gi1/34 Gi1/35 Gi1/36 Gi1/37 Gi1/38 Gi1/39
Gi1/40 Gi1/41 Gi1/42 Gi1/43 Gi1/44 Gi1/45 Gi1/46 Gi1/47 Gi1/48 Te2/1
Te2/2 Te2/3 Te2/4 Te2/5 Te2/6 Te2/7 Te2/8 Gi3/1 Gi3/2 Gi3/3
Te3/4 Te3/5 Te5/1 Te5/2 Te5/3 Te5/4 Te5/5 Te5/6 Te5/7 Te5/8
Te5/9 Te5/10 Te5/11 Te5/12 Te5/13 Te5/14 Te5/15 Te5/16 Te6/1 Te6/2
Te6/3 Te6/4 CPP CPP.1 Vl1
  QoS 10g-only mode supported: Yes [Current mode: On]
  Global Policy-map: ingress[]
```

A setting of “Current mode: On” means that the Gigabit Ethernet ports are disabled.

8Q4T ingress queueing for these line cards implements DSCP-to-queue mapping, with DSCP-based WRED for congestion avoidance. The following figure shows the 8Q4T ingress queueing model.

Figure 60 8Q4T Ingress Queuing Model—DSCP-to-Queue Mapping with DSCP-based WRED

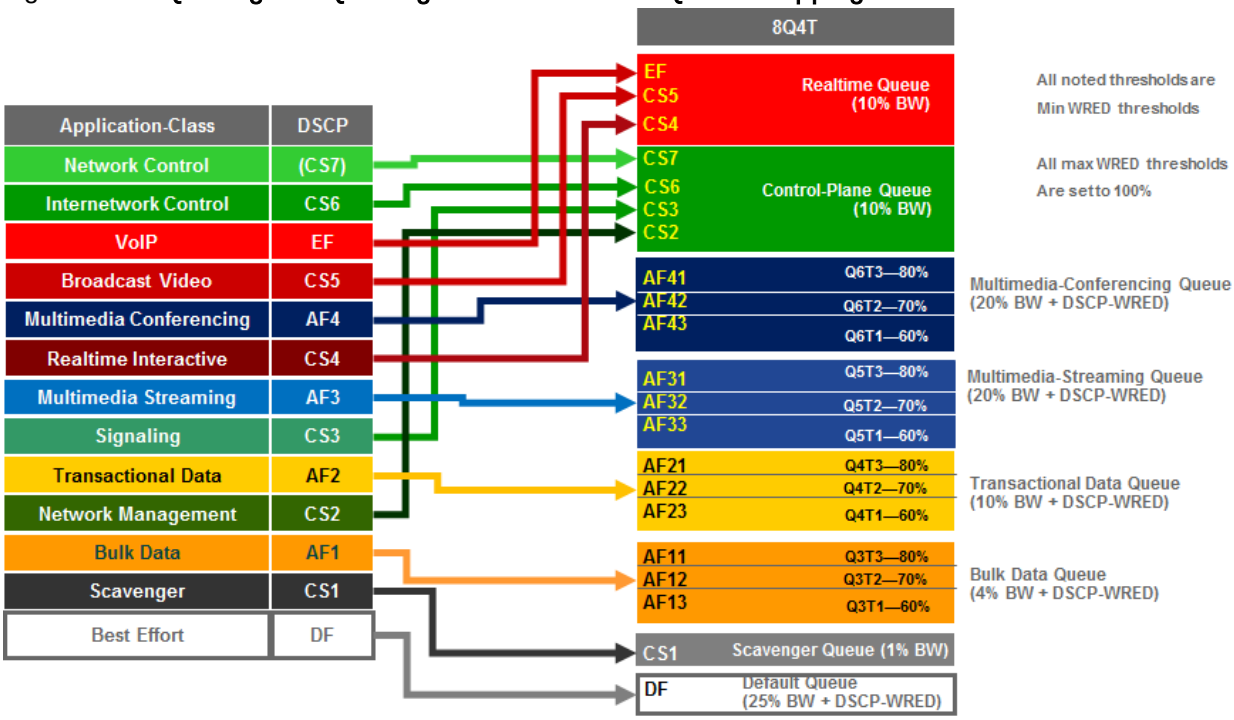

The following configuration, provisioned by APIC-EM EasyQoS, implements the class-maps for the 8Q4T ingress queuing structure.

```
!
class-map type lan-queuing match-any prm-EZQOS_8Q4T#REALTIME
 match DSCP cs4
 match DSCP cs5
 match DSCP ef
class-map type lan-queuing match-any prm-EZQOS_8Q4T#CONTROL
 match DSCP cs2
 match DSCP cs3
 match DSCP cs6
 match DSCP cs7
class-map type lan-queuing match-any prm-EZQOS_8Q4T#MM_CONF
 match DSCP af41
 match DSCP af42
 match DSCP af43
class-map type lan-queuing match-any prm-EZQOS_8Q4T#MM_STREAM
 match DSCP af31
```

```
 match DSCP af32
 match DSCP af33
class-map type lan-queuing match-any prm-EZQOS_8Q4T#TRANS_DATA
 match DSCP af21
 match DSCP af22
 match DSCP af23
class-map type lan-queuing match-any prm-EZQOS_8Q4T#BULK_DATA
 match DSCP af11
 match DSCP af12
 match DSCP af13
class-map type lan-queuing match-any prm-EZQOS_8Q4T#SCAVENGER
 match DSCP cs1
!
```

The following configuration, provisioned by APIC-EM EasyQoS, implements the policy-map for the 8Q4T ingress queuing structure.

```
!
policy-map type lan-queuing prm-DSCP#EZQOS_8Q4T-IN
 class prm-EZQOS_8Q4T#REALTIME
  bandwidth percent 10
 class prm-EZQOS_8Q4T#CONTROL
  bandwidth percent 10
 class prm-EZQOS_8Q4T#MM_CONF
  bandwidth percent 20
  random-detect DSCP-based
  random-detect DSCP 34 percent 80 100
  random-detect DSCP 36 percent 70 100
  random-detect DSCP 38 percent 60 100
 class prm-EZQOS_8Q4T#MM_STREAM
  bandwidth percent 20
  random-detect DSCP-based
  random-detect DSCP 26 percent 80 100
  random-detect DSCP 28 percent 70 100
```

```
  random-detect DSCP 30 percent 60 100
class prm-EZQOS_8Q4T#TRANS_DATA
  bandwidth percent 10
  random-detect DSCP-based
  random-detect DSCP 18 percent 80 100
  random-detect DSCP 20 percent 70 100
  random-detect DSCP 22 percent 60 100
 class prm-EZQOS_8Q4T#BULK_DATA
  bandwidth percent 4
  random-detect DSCP-based
  random-detect DSCP 10 percent 80 100
  random-detect DSCP 12 percent 70 100
  random-detect DSCP 14 percent 60 100
 class prm-EZQOS_8Q4T#SCAVENGER
  bandwidth percent 1
 class class-default
  random-detect DSCP-based
  random-detect DSCP 0 percent 80 100
!
```

The 8Q4T ingress queuing policy-map is provisioned by EasyQoS to all TenGigabitEthernet interfaces on the line cards that support this queuing structure. An example of the provisioning to a single TenGigabitEthernet interface is shown below.

```
!
interface TenGigabitEthernetx/x
 service-policy type lan-queuing input prm-DSCP#EZQOS_8Q4T-IN
!
```

For all of the line cards that support the 8Q4T ingress queueing structure listed above, the ports are assumed to be uplink ports, because they are all TenGigabitEthernet ports. If the Catalyst 6500 Series or Catalyst 6807-XL with Sup-2T is deployed as a distribution or core switch, and the link connecting the Sup-2T port is not a trunk port, then all ingress traffic will not have an 802.1p header. However, because these line cards support DSCP-to-queue mapping, ingress traffic will still be mapped into the correct queue based on the DSCP value of the IP packets.

8Q8T Ingress Queuing

8Q8T ingress queuing is supported by the following line cards:

- WS-X6704-10GE supported with a DFC4/DFC4XL upgrade (WS-F6k-DFC4-A, WS-F6k-DFC4-AXL).

The WS-X6704-10GE line card can support either a CFC or a DFC version 4 or 4XL upgrade, in order to operate with the Sup-2T. With a CFC, the WS-X6704-10GE line card supports 1Q8T ingress queuing. With a DFC4/4XL upgrade, the WS-X6704-10GE line card supports 8Q8T ingress queuing. Whether or not the WS-X6704-10GE line card has a CFC or DFC can be displayed via the exec-level "show module" command, as discussed in the *1Q8T Ingress Queuing* section of this document.

8Q8T ingress queueing for the WS-X6704-10GE implements CoS-to-queue mapping, with CoS-based tail-drop for congestion avoidance. The following figure shows the 8Q8T ingress queueing model.

Figure 61 **8Q8T Ingress Queuing Models—CoS-to-Queue Mapping with CoS-based Tail-Drop**

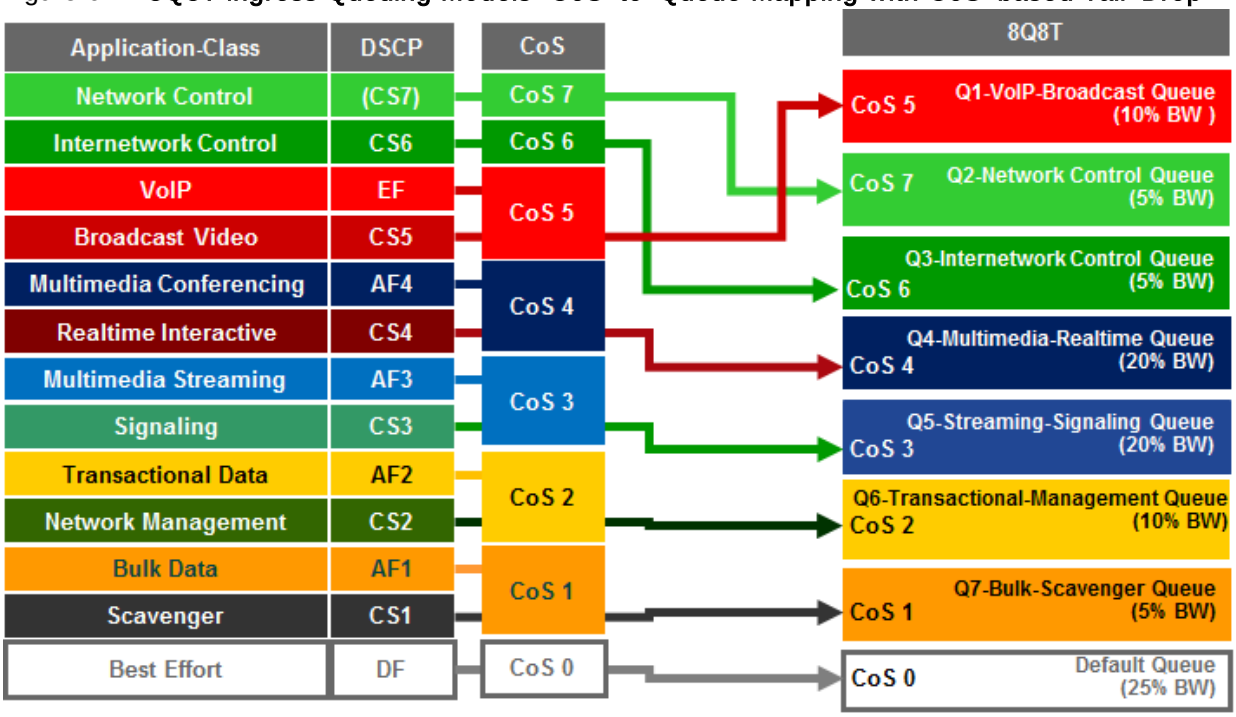

The following configuration, provisioned by APIC-EM EasyQoS, implements the class-maps for the 8Q8T ingress queuing structure.

```
!
class-map type lan-queuing match-any prm-EZQOS_8Q8T#Q1
 match cos 5
class-map type lan-queuing match-any prm-EZQOS_8Q8T#Q2
 match cos 7
class-map type lan-queuing match-any prm-EZQOS_8Q8T#Q3
 match cos 6
class-map type lan-queuing match-any prm-EZQOS_8Q8T#Q4
 match cos 4
```

```
class-map type lan-queuing match-any prm-EZQOS_8Q8T#Q5
 match cos 3
class-map type lan-queuing match-any prm-EZQOS_8Q8T#Q6
 match  cos 2
class-map type lan-queuing match-any prm-EZQOS_8Q8T#Q7
 match  cos 1
!
```

The following configuration, provisioned by APIC-EM EasyQoS, implements the policy-map for the 8Q8T ingress queuing structure.

```
!
policy-map type lan-queuing prm-DSCP#EZQOS_8Q8T-IN
 class prm-EZQOS_8Q8T#Q1
  bandwidth percent 10
 class prm-EZQOS_8Q8T#Q2
  bandwidth percent 5
 class prm-EZQOS_8Q8T#Q3
  bandwidth percent 5
 class prm-EZQOS_8Q8T#Q4
  bandwidth percent 20
 class prm-EZQOS_8Q8T#Q5
  bandwidth percent 20
 class prm-EZQOS_8Q8T#Q6
  bandwidth percent 10
 class prm-EZQOS_8Q8T#Q7
  bandwidth percent 5
 class class-default
!
```

The 8Q8T ingress queuing policy-map is provisioned by EasyQoS to all TenGigabitEthernet interfaces on the line cards that support this queuing structure. An example of the provisioning to a single TenGigabitEthernet interface is shown below.

```
!
interface TenGigabitEthernetx/x/x
service-policy type lan-queuing input prm-DSCP#EZQOS_8Q8T-IN
```

```
!
```

Ports on the WS-X6704-10GE with DFC4/4XL are assumed to be uplink ports. If the Catalyst 6500 Series or Catalyst 6807-XL with Sup-2T is deployed as a distribution or core switch, and the link connecting the WS-X6704-10GE switch port is not a trunk port, then all ingress traffic will not have an 802.1p header. Hence all ingress traffic will be treated with the default CoS mapping for the port (CoS 0) and mapped to the default queue. Because no other traffic will be in other queues, the default queue will essentially have 100% of the bandwidth available.

2P6Q4T Ingress & Egress Queuing

2P6Q4T ingress & egress queuing is supported by the following line cards:

- WS-X6904-40G-2T and WS-X6904-40G-2TXL
- C6800-8P10G, C6800-8P10G-XL
- C6800-16P10G, C6800-16P10G-XL
- C6800-32P10G, C6800-32P10G-XL

2P6Q4T ingress & egress queueing for these line cards implements DSCP-to-queue mapping, with DSCP-based WRED for congestion avoidance. The following figure shows the 2P6Q4T ingress & egress queueing model.

Figure 62 **2P6Q4T Ingress & Egress Queuing Model—DSCP-to-Queue Mapping with DSCP-based WRED**

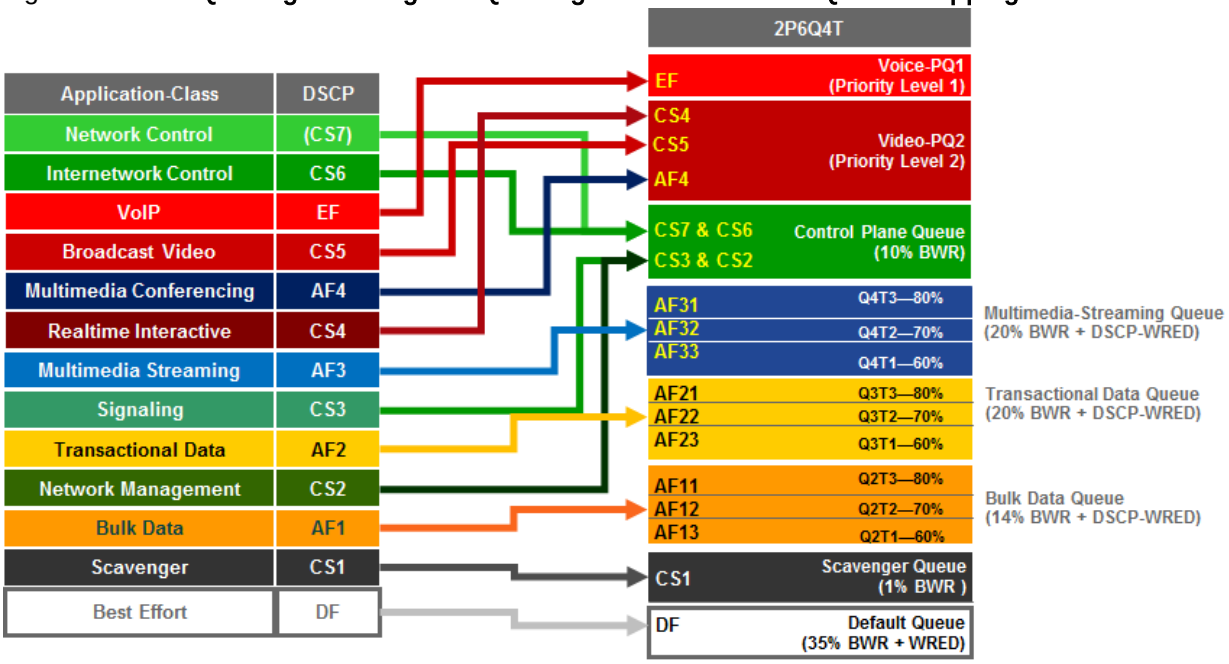

The following configuration, provisioned by APIC-EM EasyQoS, implements the class-maps for the 2P6Q4T queuing structure.

```
!
class-map type lan-queuing match-any prm-EZQOS_2P6Q4T#VOICE-PQ1
```

```
 match DSCP ef
class-map type lan-queuing match-any prm-EZQOS_2P6Q4T#VIDEO-PQ2
 match DSCP cs4
 match DSCP cs5
 match DSCP af41
 match DSCP af42
 match DSCP af43
class-map type lan-queuing match-any prm-EZQOS_2P6Q4T#CONTROL
 match DSCP cs2
 match DSCP cs3
 match DSCP cs6
 match DSCP cs7
class-map type lan-queuing match-any prm-EZQOS_2P6Q4T#MM_STREAM
 match DSCP af31
 match DSCP af32
 match DSCP af33
class-map type lan-queuing match-any prm-EZQOS_2P6Q4T#TRANS_DATA
 match DSCP af21
 match DSCP af22
 match DSCP af23
class-map type lan-queuing match-any prm-EZQOS_2P6Q4T#BULK_DATA
 match DSCP af11
 match DSCP af12
 match DSCP af13
class-map type lan-queuing match-any prm-EZQOS_2P6Q4T#SCAVENGER
 match DSCP cs1
!
```

The following configuration, provisioned by APIC-EM EasyQoS, implements the ingress policy-map for the 2P6Q4T queuing structure.

```
!
policy-map type lan-queuing prm-DSCP#EZQOS_2P6Q4T-IN
 class type lan-queuing prm-EZQOS_2P6Q4T#VOICE-PQ1
```

```
    priority level 1
 class type lan-queuing prm-EZQOS_2P6Q4T#VIDEO-PQ2
    priority level 2
 class type lan-queuing prm-EZQOS_2P6Q4T#CONTROL
    bandwidth remaining percent 10
 class type lan-queuing prm-EZQOS_2P6Q4T#MM_STREAM
    bandwidth remaining percent 20
     random-detect DSCP-based
     random-detect DSCP 26 percent 80 100
     random-detect DSCP 28 percent 70 100
     random-detect DSCP 30 percent 60 100
 class type lan-queuing prm-EZQOS_2P6Q4T#TRANS_DATA
    bandwidth remaining percent 20
     random-detect DSCP-based
     random-detect DSCP 18 percent 80 100
     random-detect DSCP 20 percent 70 100
     random-detect DSCP 22 percent 60 100
 class type lan-queuing prm-EZQOS_2P6Q4T#BULK_DATA
    bandwidth remaining percent 14
     random-detect DSCP-based
     random-detect DSCP 10 percent 80 100
     random-detect DSCP 12 percent 70 100
     random-detect DSCP 14 percent 60 100
 class type lan-queuing prm-EZQOS_2P6Q4T#SCAVENGER
    bandwidth remaining percent 1
 class class-default
     random-detect DSCP-based
     random-detect DSCP 0 percent 80 100
!
```

The following configuration, provisioned by APIC-EM EasyQoS, implements the egress policy-map for the 2P6Q4T queuing structure.

```
!
```

```
policy-map type lan-queuing prm-DSCP#EZQOS_2P6Q4T-OUT
 class type lan-queuing prm-EZQOS_2P6Q4T#VOICE-PQ1
    priority level 1
 class type lan-queuing prm-EZQOS_2P6Q4T#VIDEO-PQ2
    priority level 2
 class type lan-queuing prm-EZQOS_2P6Q4T#CONTROL
    bandwidth remaining percent 10
 class type lan-queuing prm-EZQOS_2P6Q4T#MM_STREAM
    bandwidth remaining percent 20
     random-detect DSCP-based
     random-detect DSCP 26 percent 80 100
     random-detect DSCP 28 percent 70 100
     random-detect DSCP 30 percent 60 100
 class type lan-queuing prm-EZQOS_2P6Q4T#TRANS_DATA
    bandwidth remaining percent 20
     random-detect DSCP-based
     random-detect DSCP 18 percent 80 100
     random-detect DSCP 20 percent 70 100
     random-detect DSCP 22 percent 60 100
 class type lan-queuing prm-EZQOS_2P6Q4T#BULK_DATA
    bandwidth remaining percent 14
     random-detect DSCP-based
     random-detect DSCP 10 percent 80 100
     random-detect DSCP 12 percent 70 100
     random-detect DSCP 14 percent 60 100
 class type lan-queuing prm-EZQOS_2P6Q4T#SCAVENGER
    bandwidth remaining percent 1
 class class-default
     random-detect DSCP-based
     random-detect DSCP 0 percent 80 100
!
```

The 2P6Q4T queuing policy-map is provisioned by EasyQoS to all TenGigabitEthernet and FortyGigabitEthernet interfaces, in both the ingress and egress directions, on the line cards that support this queuing structure. An example of the provisioning to a TenGigabitEthernet and FortyGigabitEthernet interface in the ingress direction is shown below.

```
!
interface TenGigabitEthernetx/x
 service-policy type lan-queuing input prm-DSCP#EZQOS_2P6Q4T-IN
!
interface FortyGigabitEthernetx/x
 service-policy type lan-queuing input prm-DSCP#EZQOS_2P6Q4T-IN
!
```

An example of the provisioning to a TenGigabitEthernet and FortyGigabitEthernet interface in the egress direction is shown below.

```
!
interface TenGigabitEthernetx/x
 service-policy type lan-queuing output prm-DSCP#EZQOS_2P6Q4T-OUT
!
interface FortyGigabitEthernetx/x
 service-policy type lan-queuing output prm-DSCP#EZQOS_2P6Q4T-OUT
!
```

For the line cards listed above that support the 2P6Q4T ingress queuing structure, the ports are assumed to be uplink ports, because they are either TenGigabitEthernet or FortyGigabitEthernet ports. If the Catalyst 6500 Series or Catalyst 6807-XL with Sup-2T is deployed as a distribution or core switch, and the link connecting the Sup-2T port is not a trunk port, then all ingress traffic will not have an 802.1p header. However, because these line cards support DSCP-to-queue mapping, ingress traffic will still be mapped into the correct queue based on the DSCP value of the IP packets.

1P3Q8T Egress Queuing

1P3Q8T egress queuing is supported by the following line cards:

- WS-X6724-SFP, WS-X6748-SFP and WS-X6748-GE-TX with CFC
- WS-X6724-SFP, WS-X6748-SFP, and WS-X6748-GE-TX with a DFC4 or DFC4XL upgrade (WS-F6k-DFC4-A, WS-F6k-DFC4-AXL)
- WS-X6824-SFP-2T and WS-X6824-SFP-2TXL
- WS-X6848-SFP-2T, WS-X6848-SFP-2TXL, WS-X6848-TX-2T and WS-X6848-TX-2TXL
- C6800-48P-SFP, C6800-48P-SFP-XL, C6800-48P-TX, and C6800-48P-TX-XL

1P3Q8T egress queueing for these line cards implements CoS-to-queue mapping, with CoS-based tail-drop for congestion avoidance.

The following figure shows the 1P3Q8T egress queueing model.

Figure 63 **1P3Q8T Egress Queuing Model—Cos-to-Queue Mapping with CoS-based Tail-Drop**

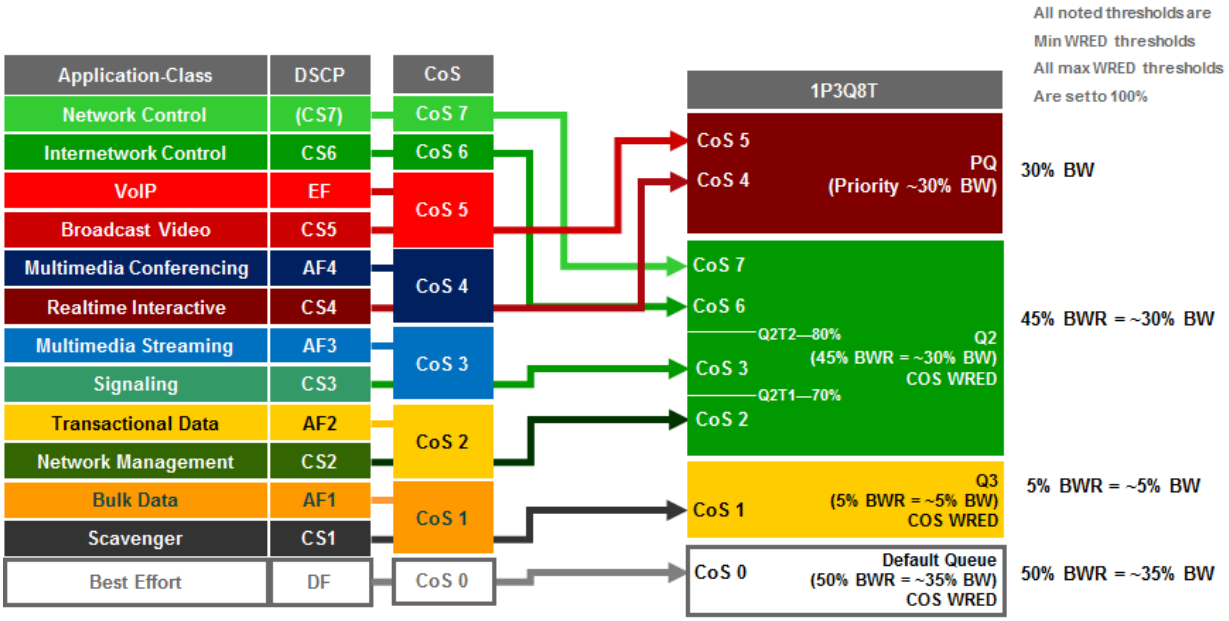

The following configuration, provisioned by APIC-EM EasyQoS, implements the class-maps for the 1P3Q8T egress queuing structure.

```
!
class-map type lan-queuing match-any prm-EZQOS_1P3Q8T#PQ
  match cos  4
  match cos  5
class-map type lan-queuing match-any prm-EZQOS_1P3Q8T#Q2
  match cos  2
  match cos  3
  match cos  6
  match cos  7
class-map type lan-queuing match-any prm-EZQOS_1P3Q8T#Q3
  match cos  1
!
```

The following configuration, provisioned by APIC-EM EasyQoS, implements the policy-map for the 1P3Q8T egress queuing structure.

```
!
```

```
policy-map type lan-queuing prm-DSCP#EZQOS_1P3Q8T-OUT
  class prm-EZQOS_1P3Q8T#PQ
    priority
  class prm-EZQOS_1P3Q8T#Q2
    bandwidth remaining percent 45
    random-detect cos-based
    random-detect cos 2 percent 70 100
    random-detect cos 3 percent 80 100
    random-detect cos 6 percent 100 100
    random-detect cos 7 percent 100 100
  class prm-EZQOS_1P3Q8T#Q3
    bandwidth remaining percent 5
    random-detect cos-based
    random-detect cos 1 percent 80 100
  class class-default
    random-detect cos-based
    random-detect cos 0 percent 80 100
!
```

The network operator should note that depending upon the particular line card, the policy-map configuration may look subtly different. Specifically the class configurations under the policy-map definition may include the words “type lan-queuing” within the definitions as shown below.

```
!
policy-map type lan-queuing prm-DSCP#EZQOS_1P3Q8T-OUT
  class type lan-queuing prm-EZQOS_1P3Q8T#PQ
    priority
  class type lan-queuing prm-EZQOS_1P3Q8T#Q2
    bandwidth remaining percent 45
    random-detect cos-based
    random-detect cos 2 percent 70 100
    random-detect cos 3 percent 80 100
    random-detect cos 6 percent 100 100
    random-detect cos 7 percent 100 100
  class type lan-queuing prm-EZQOS_1P3Q8T#Q3
```

```
    bandwidth remaining percent 5
    random-detect cos-based
    random-detect cos 1 percent 80 100
  class class-default
    random-detect cos-based
    random-detect cos 0 percent 80 100
!
```

Functionally, both variations of the policy-map definition are identical and simply represent minor differences in how the policy-map is displayed for various models of line cards. Other ingress and egress queuing structures may also exhibit these minor differences as well.

The following configuration implements the 1P3Q8T egress queuing structure. Additionally, it shows the application of the 1P3Q8T egress queueing structure to a Gigabit Ethernet interface.

The 1P3Q8T queuing policy-map is provisioned by EasyQoS to all Gigabit Ethernet interfaces, in the egress direction, on the line cards that support this queuing structure. An example of the provisioning to a Gigabit Ethernet interface is shown below.

```
!
interface GigabitEthernet x/x/x
 service-policy type lan-queuing output prm-DSCP#EZQOS_1P3Q8T-OUT
!
```

Due to the internal ASIC structure of the ports on the WS-X6748-SFP and WS-X6748-GE-TX line cards, the egress queueing structure of the ports cannot be configured independently. Instead, the queueing policy is applied to groups of ports on the line card by APIC-EM EasyQoS.

1P3Q4T Egress Queuing

Note: Due to Cisco defect CSCvb72316, EasyQoS does not currently deploy ingress or egress queuing to switch ports on the VS-S2T-10G and VS-S2T-10G-XL (Supervisor 2T). Hence, this Catalyst 6K Sup2T queuing structure is currently not used within EasyQoS.

1P3Q4T egress queuing is supported by the following line cards:

- VS-S2T-10G and VS-S2T-10G-XL with Gigabit Ethernet ports enabled

Enabling of the Gigabit Ethernet ports on the Sup-2T was discussed in the *2Q4T Ingress Queuing* section of this document.

1P3Q4T egress queueing for the Sup-2T implements CoS-to-queue mapping, with CoS-based WRED for congestion avoidance. The following figure shows the 1P3Q4T egress queueing model.

Figure 64 **1P3Q4T Egress Queuing Model—CoS-to-Queue Mapping with CoS-based WRED**

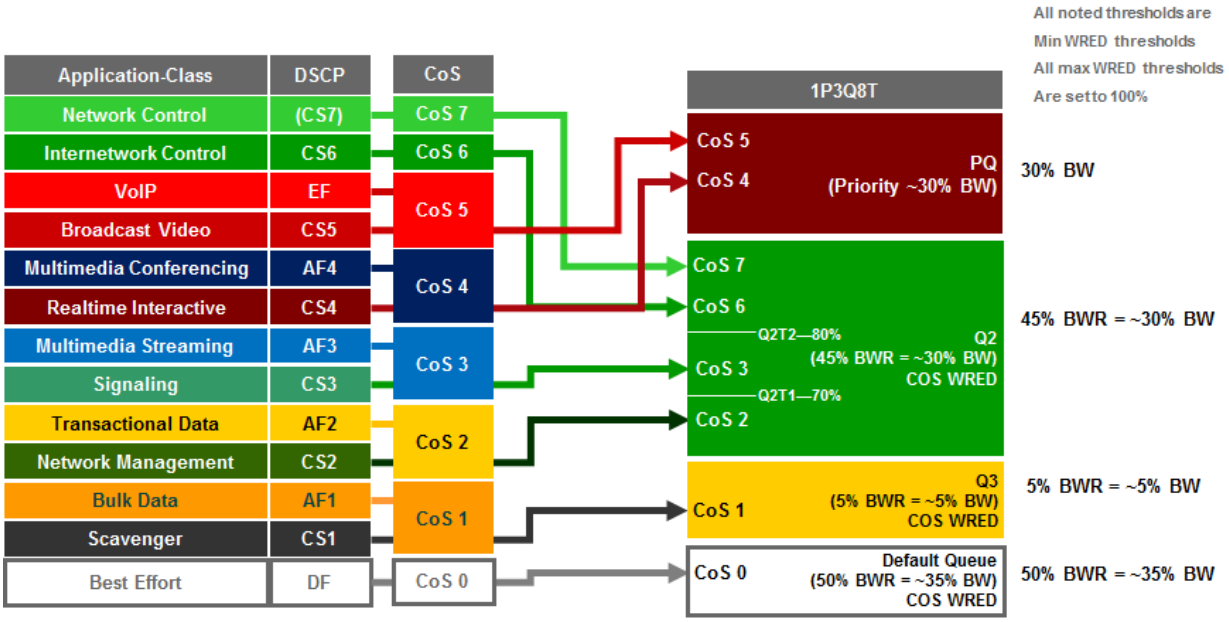

The following configuration, provisioned by APIC-EM EasyQoS, implements the class-maps for the 1P3Q4T egress queuing structure.

```
!
class-map type lan-queuing match-any prm-EZQOS_1P3Q4T#PQ
  match cos  4
  match cos  5
class-map type lan-queuing match-any prm-EZQOS_1P3Q4T#Q2
  match cos  2
  match cos  3
  match cos  6
  match cos  7
class-map type lan-queuing match-any prm-EZQOS_1P3Q4T#Q3
  match cos  1
!
```

The following configuration, provisioned by APIC-EM EasyQoS, implements the policy-map for the 1P3Q4T egress queuing structure.

```
!
policy-map type lan-queuing prm-DSCP#EZQOS_1P3Q4T-OUT
  class prm-EZQOS_1P3Q4T#PQ
    priority
```

```
    class prm-EZQOS_1P3Q4T#Q2
      bandwidth remaining percent 45
      random-detect cos-based
      random-detect cos 2 percent 70 100
      random-detect cos 3 percent 80 100
      random-detect cos 6 percent 100 100
      random-detect cos 7 percent 100 100
    class prm-EZQOS_1P3Q4T#Q3
      bandwidth remaining percent 5
      random-detect cos-based
      random-detect cos 1 percent 80 100
    class class-default
      random-detect cos-based
      random-detect cos 0 percent 80 100
!
```

The 1P3Q4T queuing policy-map is provisioned by EasyQoS to all Gigabit Ethernet interfaces, in the egress direction, on the supervisor that supports this queuing structure. An example of the provisioning to a Gigabit Ethernet interface is shown below.

```
!
interface GigabitEthernet x/x/x
 service-policy type lan-queuing output prm-DSCP#EZQOS_1P3Q4T-OUT
!
```

As discussed in the *2Q4T Ingress Queuing* section of this document, QoS service policies cannot be applied to TenGigabitEthernet ports on the Sup-2T when the TenGigabitEthernet ports are part of a port-channel group that is part of a virtual switch link (VSL). Also, QoS service policies cannot be applied to Gigabit Ethernet ports on the Sup-2T when the TenGigabitEthernet ports are part of a VSL.

Finally, due to the internal ASIC structure of the ports on the Sup-2T, the egress queueing structure of the ports cannot be configured independently. Instead, the queueing policy is applied to groups of ports on the line card by APIC-EM EasyQoS.

1P7Q4T Egress Queuing

Note: Due to Cisco defect CSCvb72316, EasyQoS does not currently deploy ingress or egress queuing to switch ports on the VS-S2T-10G and VS-S2T-10G-XL (Supervisor 2T).

1P7Q4T egress queuing is supported by the following line cards:

- WS-X6908-10G-2T and WS-X6908-10G-2TXL
- VS-S2T-10G and VS-S2T-10G-XL with Gigabit Ethernet ports disabled

Disabling of the Gigabit Ethernet ports on the Sup-2T was discussed in the *8Q4T Ingress Queuing* section of this document.

1P7Q4T egress queueing for these line cards implements DSCP-to-queue mapping, with DSCP-based WRED for congestion avoidance. The following figure shows the 1P7Q4T egress queueing model.

Figure 65 **1P7Q4T Egress Queuing Model—DSCP-to-Queue Mapping with DSCP-based WRED**

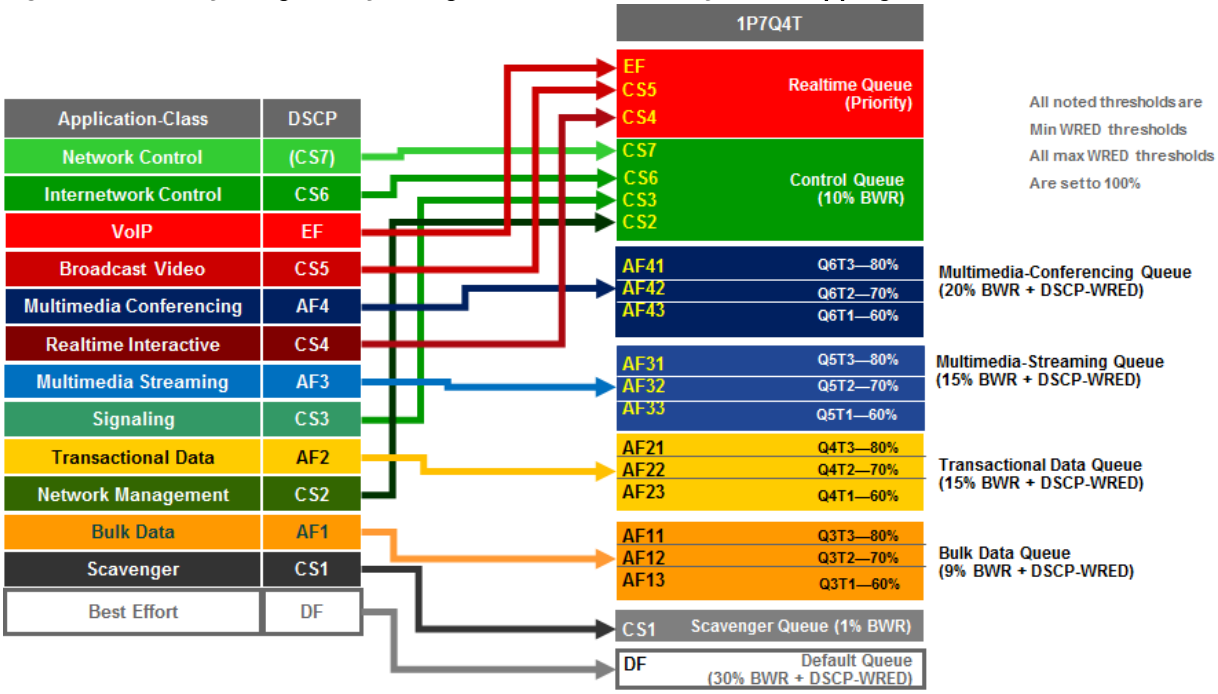

The following configuration, provisioned by APIC-EM EasyQoS, implements the class-maps for the 1P7Q4T egress queuing structure.

```
!
class-map type lan-queuing match-any prm-EZQOS_1P7Q4T#REALTIME
  match DSCP cs4
  match DSCP cs5
  match DSCP ef
class-map type lan-queuing match-any prm-EZQOS_1P7Q4T#MM_CONF
  match DSCP af41
  match DSCP af42
  match DSCP af43
class-map type lan-queuing match-any prm-EZQOS_1P7Q4T#MM_STREAM
  match DSCP af31
```

```
  match DSCP af32
  match DSCP af33
class-map type lan-queuing match-any prm-EZQOS_1P7Q4T#CONTROL
  match DSCP cs2
  match DSCP cs3
  match DSCP cs6
  match DSCP cs7
class-map type lan-queuing match-any prm-EZQOS_1P7Q4T#TRANS_DATA
  match DSCP af21
  match DSCP af22
  match DSCP af23
class-map type lan-queuing match-any prm-EZQOS_1P7Q4T#BULK_DATA
  match DSCP af11
  match DSCP af12
  match DSCP af13
class-map type lan-queuing match-any prm-EZQOS_1P7Q4T#SCAVENGER
  match  DSCP cs1
!
```

The following configuration, provisioned by APIC-EM EasyQoS, implements the policy-map for the 1P7Q4T egress queuing structure.

```
!
policy-map type lan-queuing prm-DSCP#EZQOS_1P7Q4T-OUT
  class prm-EZQOS_1P7Q4T#REALTIME
    priority
  class prm-EZQOS_1P7Q4T#CONTROL
    bandwidth remaining percent 10
  class prm-EZQOS_1P7Q4T#MM_CONF
    bandwidth remaining percent 20
    random-detect DSCP-based
    random-detect DSCP 34 percent 80 100
    random-detect DSCP 36 percent 70 100
    random-detect DSCP 38 percent 60 100
```

```
  class prm-EZQOS_1P7Q4T#MM_STREAM
    bandwidth remaining percent 15
    random-detect DSCP-based
    random-detect DSCP 26 percent 80 100
    random-detect DSCP 28 percent 70 100
    random-detect DSCP 30 percent 60 100
  class prm-EZQOS_1P7Q4T#TRANS_DATA
    bandwidth remaining percent 15
    random-detect DSCP-based
    random-detect DSCP 18 percent 80 100
    random-detect DSCP 20 percent 70 100
    random-detect DSCP 22 percent 60 100
  class prm-EZQOS_1P7Q4T#BULK_DATA
    bandwidth remaining percent 9
    random-detect DSCP-based
    random-detect DSCP 10 percent 80 100
    random-detect DSCP 12 percent 70 100
    random-detect DSCP 14 percent 60 100
  class prm-EZQOS_1P7Q4T#SCAVENGER
    bandwidth remaining percent 1
  class class-default
    random-detect DSCP-based
    random-detect DSCP 0 percent 80 100
!
```

The 1P7Q4T queuing policy-map is provisioned by EasyQoS to all TenGigabitEthernet interfaces, in the egress direction, on the line cards that support this queuing structure. An example of the provisioning to a TenGigabitEthernet interface is shown below.

```
!
interface TenGigabitEthernetx/x
 service-policy type lan-queuing output prm-DSCP#EZQOS_1P7Q4T-OUT
!
```

1P7Q8T Egress Queuing

1P7Q8T egress queuing is supported by the following line cards:

- WS-X6704-10GE with CFC
- WS-X6704-10GE with a DFC4 or DFC4XL upgrade (WS-F6k-DFC4-A, WS-F6k-DFC4-AXL)

1P7Q8T egress queueing for these line cards implements CoS-to-queue mapping, with CoS-based tail-drop for congestion avoidance. Note that due to the combination of 8 queues and only 8 CoS values, tail-drop thresholds are not used in this design.

The following figure shows the 1P7Q8T egress queueing model.

Figure 66 **1P7Q8T Egress Queuing Model—CoS-to-Queue Mapping with CoS-based Tail-Drop**

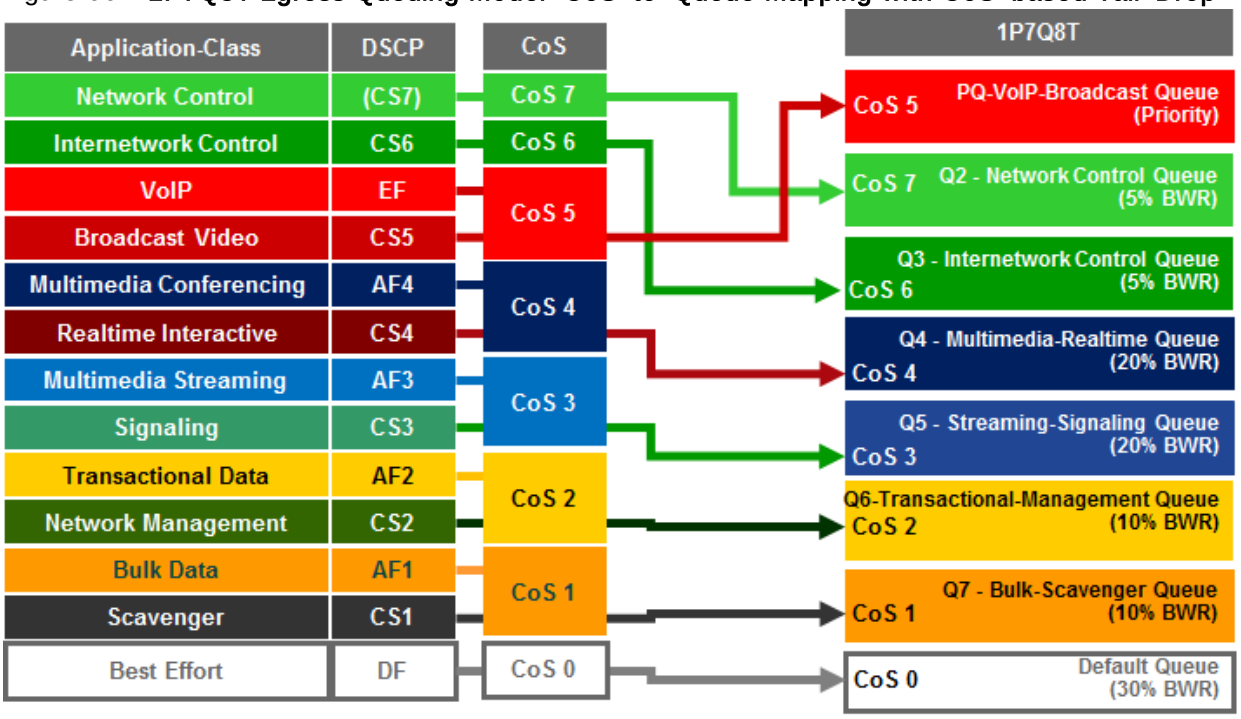

The following configuration, provisioned by APIC-EM EasyQoS, implements the class-maps for the 1P7Q8T egress queuing structure.

```
!
class-map type lan-queuing match-any prm-EZQOS_1P7Q8T#PQ
 match cos 5
class-map type lan-queuing match-any prm-EZQOS_1P7Q8T#Q2
 match cos 7
class-map type lan-queuing match-any prm-EZQOS_1P7Q8T#Q3
 match cos 6
class-map type lan-queuing match-any prm-EZQOS_1P7Q8T#Q4
```

```
 match cos 4
class-map type lan-queuing match-any prm-EZQOS_1P7Q8T#Q5
 match cos 3
class-map type lan-queuing match-any prm-EZQOS_1P7Q8T#Q6
 match  cos 2
class-map type lan-queuing match-any prm-EZQOS_1P7Q8T#Q7
 match  cos 1
!
```

The following configuration, provisioned by APIC-EM EasyQoS, implements the policy-map for the 1P7Q8T egress queuing structure.

```
!
policy-map type lan-queuing prm-DSCP#EZQOS_1P7Q8T-OUT
 class prm-EZQOS_1P7Q8T#PQ
  priority
 class prm-EZQOS_1P7Q8T#Q2
  bandwidth remaining percent 5
 class prm-EZQOS_1P7Q8T#Q3
  bandwidth remaining percent 5
 class prm-EZQOS_1P7Q8T#Q4
  bandwidth remaining percent 20
 class prm-EZQOS_1P7Q8T#Q5
  bandwidth remaining percent 20
 class prm-EZQOS_1P7Q8T#Q6
  bandwidth remaining percent 10
 class prm-EZQOS_1P7Q8T#Q7
  bandwidth remaining percent 10
 class class-default
!
```

The 1P7Q8T queuing policy-map is provisioned by EasyQoS to all TenGigabitEthernet interfaces, in the egress direction, on the line cards that support this queuing structure. An example of the provisioning to a TenGigabitEthernet interface is shown below.

```
!
interface TenGigabitEthernet x/x/x
```

```
 service-policy type lan-queuing output prm-DSCP#EZQOS_1P7Q8T-OUT
!
```

Cisco Catalyst 6880 and 6840 Series Queuing Design

Catalyst 6880 and 6840 Series switch platforms are only supported in the roles of a core-layer or distribution-layer switch within the APIC-EM EasyQoS solution. These platforms feature an embedded supervisor–similar to the Sup2T supervisor within Catalyst 6500 Series switches and the Catalyst 6807-XL switch. Switch ports on the Catalyst 6880 and 6840 Series platforms support a 2P6Q4T ingress and egress queueing structure. Hence the ingress and egress queuing design is the same as discussed in the *2P6Q4T Ingress & Egress Queuing* section of this document and will not be duplicated here for brevity.

Catalyst 6500 Sup720-10GE Queuing Design

The following sections discuss the ingress and egress queuing structures pushed by APIC-EM to the ports of each of the line cards and supervisors supported by EasyQoS for the Catalyst 6500 Series switch with Sup720-10GE supervisor.

Ingress & egress queueing structures are dependent upon the following:

- The model of the line card
- Whether the line card supports a CFC or DFC (applies to WS-X6704, WS-X6724, and WS-X6748 series line cards)
- Whether the Sup-720 (VS-S720-10G-3C and VS-S720-10G-3CXL) Gigabit Ethernet ports are enabled or disabled

For the Catalyst 6500 Series switch with Supervisor 720-10GE (Sup720), if an unsupported line card is detected within the chassis, the behavior of APIC-EM EasyQoS is to not provision any QoS configuration for the entire platform. For the Catalyst 6500 Series switch with Sup720-10GE, only line cards that have a DFC3C or DFC3CXL daughter card, or a CFC, are supported. Older line cards that support a DFC3A or DFC3B daughter card are not supported. Inserting the line card with one of these older DFC3A or DFC3B daughter cards into the Catalyst 6500 with Sup720-10GE may change the behavior of the platform. Specifically, the platform may “downgrade” its capabilities to the least common denominator for backward compatibility–meaning the lowest line card with a DFC3A or DFC3B installed. These capabilities may not be compatible with the QoS configurations provisioned by APIC-EM EasyQoS. Hence this is not supported, and APIC-EM EasyQoS will not provision any QoS configuration to Catalyst 6500 Series switch with Supervisor 720-10GE (Sup720), if an unsupported line card is detected within the chassis. The network operator will also be notified via the EasyQoS web-based GUI that an unsupported line card was detected within the chassis.

The following figure provides a flowchart that can be used to determine when ingress queuing is applied for the Catalyst 6500 Series with a Sup-720.

Figure 67 **When is Ingress Queuing Applied?**

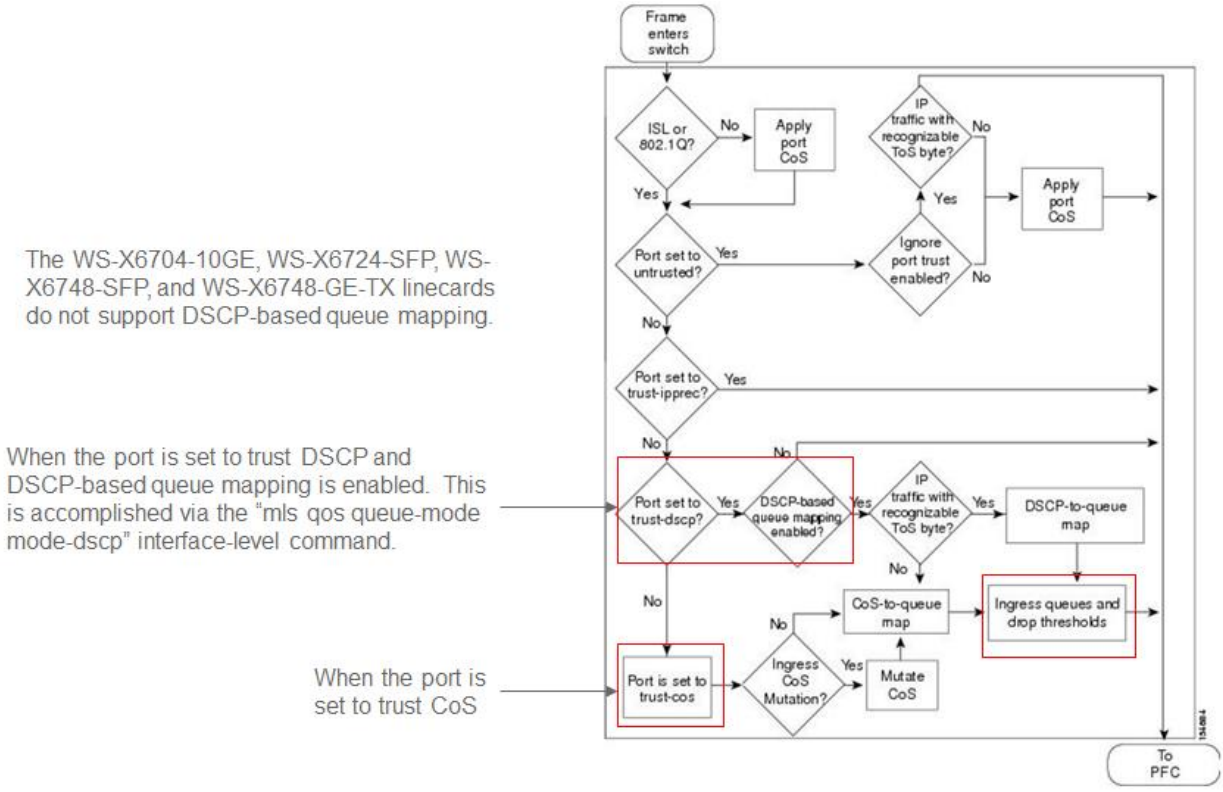

As can be seen in the figure above, ingress queuing applies under the following conditions:

- When a switch port is set to trust DSCP and DSCP-based queue mapping is enabled via the "mls qos queue-mode mode-DSCP" interface level command.
- When the switch port is set to trust CoS.

Several WS-X67xx Series line cards do not support DSCP-based queue mapping—they only support ingress CoS-to-queue mapping. With the EasyQoS solution, trust of ingress CoS markings is never enabled on switch ports. Trust of ingress DSCP markings is enabled on uplink ports. No trust is enabled for switch ports directly connected to end-user devices. Instead, an ingress classification & marking policy is used for switch ports directly connected to end-user devices.

As a result of this, there really is no need of for APIC-EM EasyQoS to push an ingress queueing policy from APIC-EM to the following line cards. Only DSCP trust will be configured by EasyQoS for these line cards, on uplink ports only.

- WS-X6724-SFP with CFC, WS-X6748-SFP with CFC, WS-X67480GE-TX with CFC, WS-6704-10GE with CFC. These line cards support a 1Q8T ingress queuing structure with CoS-to-queue mapping and CoS-based tail-drop for congestion avoidance.
- The Supervisor 720-10GE itself (VSS-720-10G-3C and VSS-S720-10G-3CXL) when the Gigabit Ethernet ports are active. The Sup720-10GE supports a 2Q4T ingress queuing structure with CoS-to-queue mapping and CoS-based tail-drop for congestion avoidance when the Gigabit Ethernet ports on the supervisor are active.

- WS-X6724-SFP with DFC3C or DFC3CXL, WS-X6748-SFP with DFC3C or DFC3CXL, or WS-X6748-GE-TX with DFC3C or DFC3CXL (WS-F6700-FDC3C or WS-F6700-DFC3CXL). These line cards support an 8Q4T ingress queuing structure with CoS-to-queue mapping and CoS-based tail-drop for congestion avoidance.
- WS-X6704-10GE with DFC3C or DFC3CXL (WS-F6700-FDC3C or WS-F6700-DFC3CXL). These line cards support an 8Q8T ingress queuing structure with CoS-to-queue mapping and CoS-based tail-drop for congestion avoidance.

The Catalyst 6500-E Series platforms with Sup-720 are MLS QoS based, which require QoS to be enabled globally first before configuring any other QoS commands. The following commands, provisioned by APIC-EM EasyQoS, enable QoS globally and set the internal COS-to-DSCP mapping table within the platform.

```
!
mls qos
! Globally Enables QoS
mls qos map cos-DSCP 0 8 16 24 32 46 48 56
! Maps CoS 5 to 46 (rest are default)
!
```

Note: Catalyst 6K and 4K switches are supported in VSS and non-VSS configurations. When operating in a VSS configuration, switch ports that belong to a port-group, which in turn are part of the VSL between the individual switches in the VSS configuration, cannot be configured with any QoS policy. APIC-EM EasyQoS has the ability to identify ports that are part of a VSL and not apply QoS policy.

1Q8T Ingress Queuing

1Q8T ingress queuing is supported by the following line cards:

- WS-X6704-10GE with CFC
- WS-X6724-SFP with CFC
- WS-X6748-SFP and WS-X6748-GE-TX with CFC

The WS-X6724-SFP, WS-X6748-SFP, WS-X6748-GE-TX, and WS-X6704-10GE line cards are supported in the Catalyst 6500 Series Sup-720 with either a CFC or a with a DFC version 3C or 3C-XL upgrade. The DFC is daughter card that sits on the line card itself. Which slots within the Catalyst 6500 Series hold these line cards can be displayed via the “show module” exec-level command. An example of the output of the “show module” command is shown below, with a WS-X6748-GE-TX line card with a CFC in slot 3 and a WS-X6748-GE-TX with a DFC3CXL in slot 4, shown in bold.

```
6504E-Sup720-10G#show module
Mod Ports Card Type                              Model              Serial No.
--- ----- -------------------------------------- ------------------ -----------
  1    5  Supervisor Engine 720 10GE (Active)    VS-S720-10G        SAL14480SAC
```

```
  2    5  Supervisor Engine 720 10GE (Cold)      VS-S720-10G        SAL12373BV5
  3   48  CEF720 48 port 10/100/1000mb Ethernet  WS-X6748-GE-TX     SAL1418GV6P
  4   48  CEF720 48 port 10/100/1000mb Ethernet  WS-X6748-GE-TX     SAL1440VP6T

Mod MAC addresses                       Hw    Fw           Sw           Status
--- ----------------------------------- ------ ------------ ------------ -------
  1  c47d.4ffe.68d8 to c47d.4ffe.68df   3.2   8.5(4)       12.2(33)SXI4 Ok
  2  001e.4af8.4498 to 001e.4af8.449f   2.0   8.5(2)       12.2(33)SXI9 Ok
  3  5475.d063.0340 to 5475.d063.036f   3.4   12.2(18r)S1  12.2(33)SXI4 Ok
  4  1cdf.0f90.4c30 to 1cdf.0f90.4c5f   3.4   12.2(18r)S1  12.2(33)SXI4 Ok

Mod  Sub-Module                  Model              Serial       Hw     Status
---- --------------------------- ------------------ ----------- ------- -------
  1  Policy Feature Card 3       VS-F6K-PFC3CXL     SAL14480QT0  1.2    Ok
  1  MSFC3 Daughterboard         VS-F6K-MSFC3       SAL14470GJ9  5.1    Ok
  2  Policy Feature Card 3       VS-F6K-PFC3CXL     SAL12383WAN  1.0    Ok
  2  MSFC3 Daughterboard         VS-F6K-MSFC3       SAL12373ADR  1.0    Ok
  3  Centralized Forwarding Card WS-F6700-CFC       SAL1412DA30  4.1    Ok
  4  Distributed Forwarding Card WS-F6700-DFC3CXL   SAL1436SZX7  1.6    Ok
```

1Q8T ingress queueing for these line cards implements CoS-to-queue mapping, with CoS-based tail-drop for congestion avoidance. The following figure shows the ingress queueing model implemented by the EasyQoS solution for these line cards.

Figure 68 **Queuing for WS-X6724-SFP with CFC, WS-X6748-SFP with CFC, WS-X6748-GE-TX with CFC, and WS-X6704-10GE with CFC** when "trust cos" is Not Configured

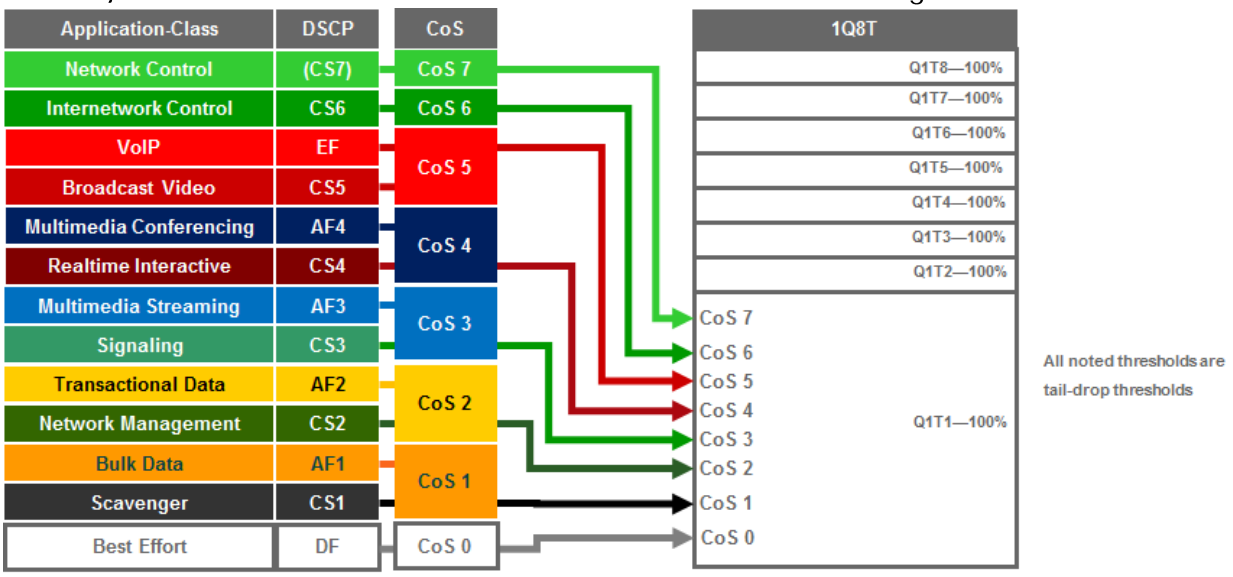

An example of the configuration pushed by APIC-EM for each of the ports on these line cards is shown below.

```
!
interface GigabitEthernet x/x/x
! Switch ports connected to end-user devices
!
interface TenGigabitEthernet x/x/x
 mls qos trust DSCP
! Switch ports functioning as uplink ports
!
```

As can be seen, DSCP trust is extended to switch ports functioning as uplink ports. These would typically correspond to TenGigabitEthernet ports on the line card. Because no trust is extended on ingress switch ports, all ingress traffic is treated as if it had a value of CoS 0 and mapped to Q1T1 with a tail-drop threshold of 100%.

2Q4T Ingress Queuing

2Q4T ingress queuing is supported by the following line cards:

- All ports on the VS-S720-10G-3C and VS-S720-10G-3CXL (Supervisor 720) when the Gigabit Ethernet ports are enabled.

The Gigabit Ethernet ports on the Sup-720 are enabled with the following global configuration command.

```
!
no mls qos 10g-only
```

```
!
```

APIC-EM EasyQoS does not set this command but will look to see if this command has been set by the network operator, in order to determine the correct queuing structure to apply to ports on the Sup720-10GE supervisor. The default setting is for the Gigabit Ethernet ports on the Sup-720 to be enabled, so this command will not appear in the configuration.

The status of whether the Gigabit Ethernet ports are enabled or disabled can be displayed via the exec-level "show mls qos module x" command, where "x" refers to the slot with the Sup-720. An example of the output from the command is shown below:

```
6504E-Sup720-10G#show mls qos module 1
  QoS is enabled globally
  Policy marking depends on port_trust
  QoS ip packet DSCP rewrite enabled globally
  QoS serial policing mode disabled globally
  Input mode for GRE Tunnel is Pipe mode
  Input mode for MPLS is Pipe mode
 QoS Trust state is DSCP on the following interface:
Gi1/1 Gi1/2 Gi1/3 Gi2/1 Gi2/2 Gi2/3 Gi3/31 Po10
  Vlan or Portchannel(Multi-Earl) policies supported: Yes
  Egress policies supported: Yes
  QoS 10g-only mode supported: Yes [Current mode: Off]


 ----- Module [1] -----
  QoS global counters:
    Total packets: 3920723
    IP shortcut packets: 2196428
    Packets dropped by policing: 0
    IP packets with TOS changed by policing: 2
    IP packets with COS changed by policing: 2
    Non-IP packets with COS changed by policing: 0
    MPLS packets with EXP changed by policing: 0
```

A setting of "Current mode: off" means that the Gigabit Ethernet ports are enabled.

2Q4T ingress queueing for the Sup-720 implements CoS-to-queue mapping, with CoS-based tail-drop for congestion avoidance. The Sup-720 interfaces do not support DSCP-based queue mapping when the 1 Gbps interfaces are enabled–they only support ingress CoS-to-queue mapping. With the EasyQoS solution,

trust of ingress CoS markings is never enabled. Trust of ingress DSCP markings is enabled on uplink ports. As a result of this, there really is no need of provisioning the 2Q4T ingress queueing policy from APIC-EM EasyQoS to the Sup-720 when the 1 Gbps ports are enabled.

The following figure shows the ingress queueing model implemented by APIC-EM EasyQoS for the Sup-720-10GE when the 1 Gbps ports are enabled.

Figure 69 **Queuing for VS-S720-10G-3C and VS-S720-10G-3CXL with Gigabit Ethernet ports enabled** and when "trust cos" is Not Configured

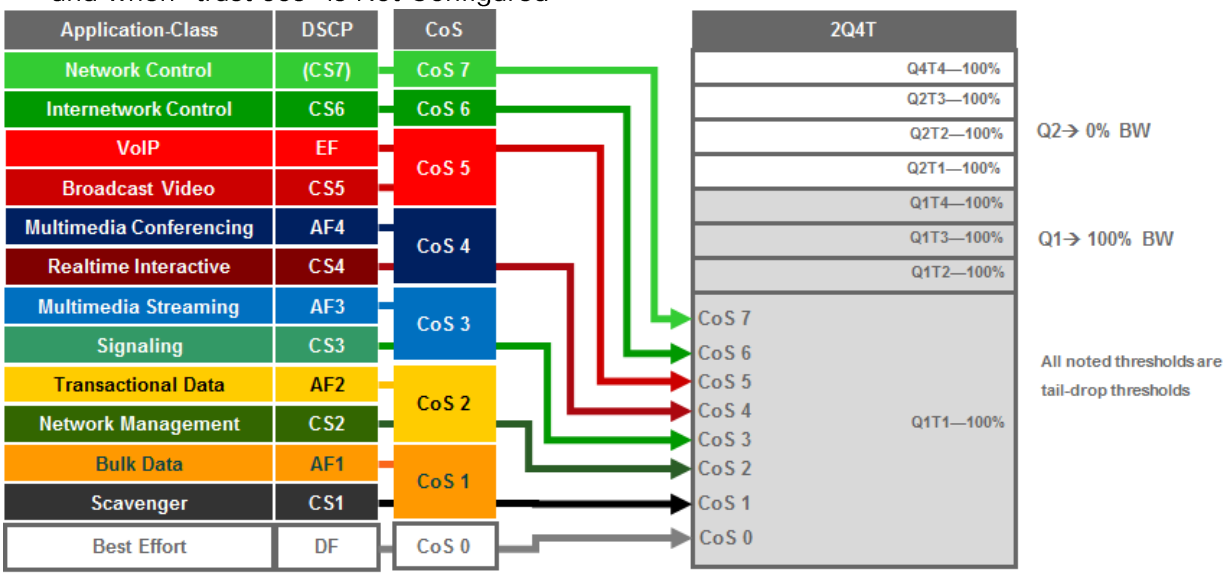

An example of the configuration provisioned by APIC-EM EasyQoS for the Sup720-10GE ports is shown below.

```
!
interface GigabitEthernet x/x/x
 mls qos trust DSCP
! Switch ports functioning as uplink ports
!
interface TenGigabitEthernet x/x/x
 mls qos trust DSCP
! Switch ports functioning as uplink ports
!
```

Cisco does not recommend connecting end-user devices to any ports on the Sup720-10GE. Hence, all ports on the Sup720-10GE are assumed to be uplink ports by EasyQoS.

Because no trust is extended on ingress switch ports, all ingress traffic is treated as if it had a value of CoS 0 and mapped to Q1T1 with a tail-drop threshold of 100%.

It should be noted that QoS service policies cannot be applied to TenGigabitEthernet ports on the Sup-720 when the TenGigabitEthernet ports are part of a port-channel group that is part of a VSL. Switch ports can be identified as part of a VSL based upon the configuration. An example of this is shown below.

```
!
interface Port-channel20
 no ip address
 switch virtual link 1
 no platform qos channel-consistency
~
interface TenGigabitEthernet1/1/4
 no ip address
 channel-group 20 mode on
!
interface TenGigabitEthernet1/1/5
 no ip address
 channel-group 20 mode on
!
```

Physical interfaces are assigned to a port-group via the "channel-group" interface-level command. Port-channel interfaces are assigned to a VSL via the "switch virtual link" interface-level command, as shown in the configuration above.

Note also that, as with the Sup-2T, EasyQoS cannot apply any QoS policy to either the TenGigabitEthernet or Gigabit Ethernet ports on the Sup-720 when the TenGigabitEthernet ports are part of a VSL.

2Q8T Ingress Queuing

2Q8T ingress queuing is supported by the following line cards:

- WS-X6724-SFP with a DFC3C or DFC3CXL (WS-F6700-DFC3C or WS-F6700-DFC3CXL)
- WS-X6748-SFP with a DFC3C or DFC3CXL
- WS-X6748-GE-TX with a DFC3C or DFC3CXL

How to determine if a line card has a CFC versus a DFC3C or DFC3CXL was discussed in the *1Q8T Ingress Queuing* section of this document.

2Q8T ingress queuing for the Sup-720 implements CoS-to-queue mapping, with CoS-based tail-drop for congestion avoidance. WS-X6724-SFP, WS-X6748-SFP, and WS-X6748-GE-TX with DFC3C/DFC3CXL line cards do not support DSCP-based queue mapping–they only support ingress CoS-to-queue mapping. With the EasyQoS solution, trust of ingress CoS markings is never enabled. Trust of ingress DSCP markings is enabled on uplink ports. As a result of this, there really is no need of provisioning the 2Q8T ingress queueing policy from APIC-EM EasyQoS to the line cards with this queueing structure.

The following figure shows the ingress queueing model implemented by APIC-EM EasyQoS solution for the WS-X6724-SFP, WS-X6748-SFP, and WS-X6748-GE-TX with DFC3C/DFC3CXL line cards.

Figure 70 **Queuing for WS-X6724-SFP, WS-X6748-SFP, WS-X6748-GE-TX with a DFC3C or DFC3XL** when “trust cos” is not configured

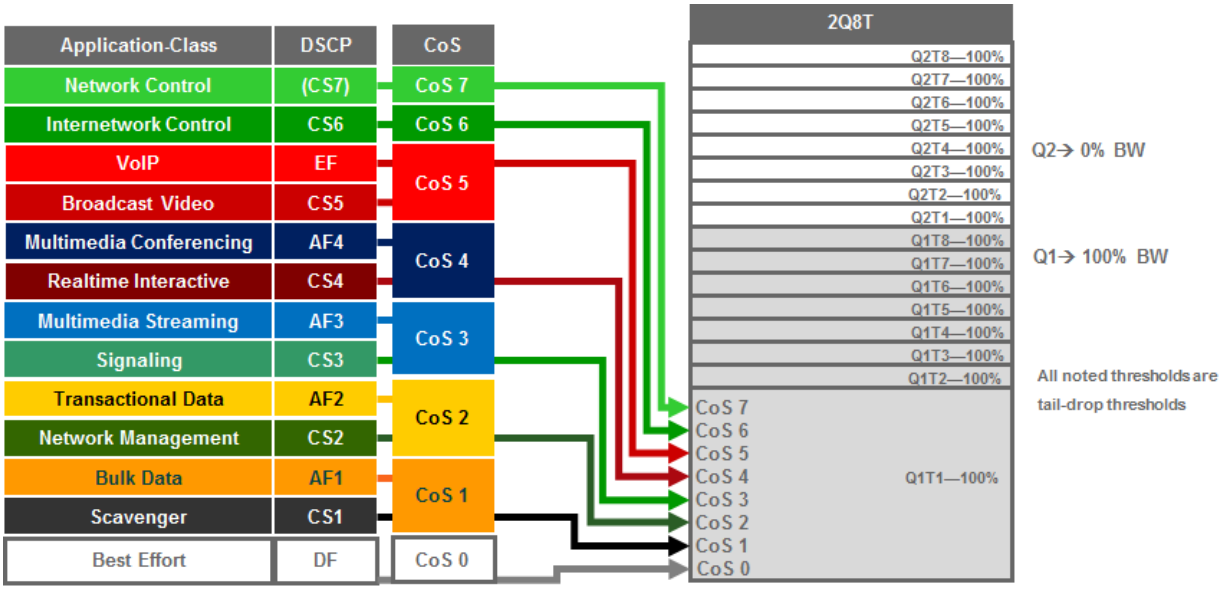

An example of the configuration pushed by APIC-EM EasyQoS for the WS-X6724-SFP, WS-X6748-SFP, and WS-X6748-GE-TX with DFC3C/DFC3CXL line card ports is shown below.

```
!
interface GigabitEthernet x/x/x
 mls qos trust DSCP
! Switch ports functioning as uplink ports
!
interface GigabitEthernet x/x/x
 mls qos trust DSCP
! Switch ports functioning as uplink ports
!
```

Because the WS-X6724-SFP, WS-X6748-SFP, and WS-X6748-GE-TX with DFC3C/DFC3CXL line card ports are all Gigabit Ethernet ports, generally these ports are assumed to be connected to end-user devices. When these ports are not used for uplinks, DSCP trust is not configured. If ports are used for uplinks, then DSCP trust is configured.

Because no trust CoS is extended on ingress switch ports and DSCP-to-queue mapping is not supported on these line cards, all ingress traffic is treated as if it had a value of CoS 0 and mapped to Q1T1 with a tail-drop threshold of 100%.

8Q4T Ingress Queuing

8Q4T ingress queuing is supported by the following line cards:

- WS-X6708-10G-3C, WS-X6708-10G-3CXL
- VS-S720-10G-3C with TenGigabitEthernet ports 4 & 5 when Gigabit Ethernet ports are inactive

The Gigabit Ethernet ports on the Sup-720 are enabled with the following global configuration command.

```
!
mls qos 10g-only
!
```

APIC-EM EasyQoS does not set this command but will look to see if this command has been set by the network operator, in order to determine the correct queuing structure to apply to ports on the Sup720-10GE supervisor. The default setting is for the Gigabit Ethernet ports on the Sup-720 to be enabled, so this command will appear in the configuration when the Gigabit Ethernet ports are disabled.

The status of whether the Gigabit Ethernet ports are enabled or disabled can be displayed via the exec-level "show mls qos module x" command, where "x" refers to the slot with the Sup-720. An example of the output from the command is shown below:

```
6504E-Sup720-10G#show mls qos module 1
  QoS is enabled globally
  Policy marking depends on port_trust
  QoS ip packet DSCP rewrite enabled globally
  QoS serial policing mode disabled globally
  Input mode for GRE Tunnel is Pipe mode
  Input mode for MPLS is Pipe mode
 QoS Trust state is DSCP on the following interface:
Gi1/1 Gi1/2 Gi1/3 Gi2/1 Gi2/2 Gi2/3 Gi3/31 Po10
  Vlan or Portchannel(Multi-Earl) policies supported: Yes
  Egress policies supported: Yes
  QoS 10g-only mode supported: Yes [Current mode: On]


 ----- Module [1] -----
  QoS global counters:
    Total packets: 3920723
    IP shortcut packets: 2196428
    Packets dropped by policing: 0
```

```
    IP packets with TOS changed by policing: 2
    IP packets with COS changed by policing: 2
    Non-IP packets with COS changed by policing: 0
    MPLS packets with EXP changed by policing: 0
```

A setting of “Current mode: On” means that the Gigabit Ethernet ports are disabled.

8Q4T ingress queueing for these line cards implements DSCP-to-queue mapping, with DSCP-based WRED for congestion avoidance. The following figure shows the 8Q4T ingress queueing model.

Figure 71 **8Q4T Ingress Queuing Models—DSCP-to-Queue with DSCP-based WRED**

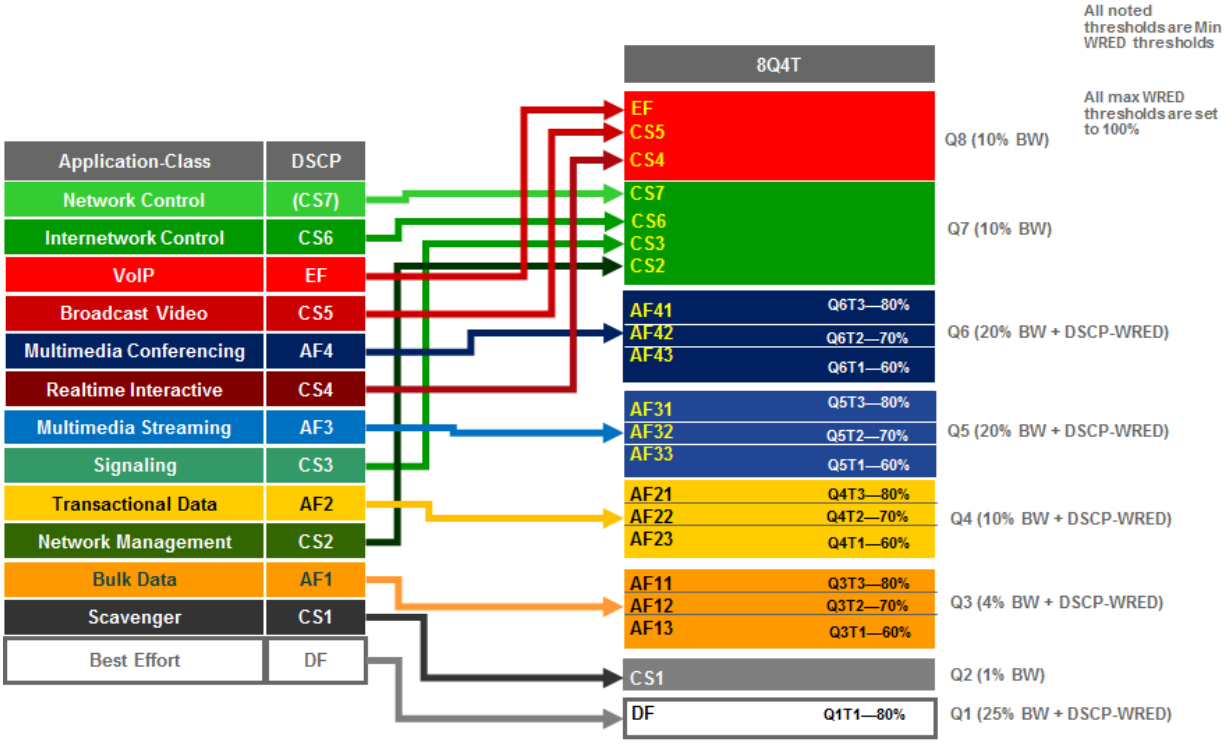

The following example configuration implements the 8Q4T ingress queuing structure to a TenGigabitEthernet interface.

```
!
interface TenGigabitEthernet x/x/x
 mls qos queue-mode mode-DSCP
 mls qos trust DSCP
 rcv-queue queue-limit 25 10 10 10 10 10 10 15
 rcv-queue bandwidth 25 1 4 10 20 20 10 10
 rcv-queue random-detect 1
 no rcv-queue random-detect 2
```

```
rcv-queue random-detect 3
rcv-queue random-detect 4
rcv-queue random-detect 5
rcv-queue random-detect 6
no rcv-queue random-detect 7
no rcv-queue random-detect 8
rcv-queue random-detect max-threshold 1 100 100 100 100
rcv-queue random-detect min-threshold 1 80 100 100 100
rcv-queue random-detect max-threshold 3 100 100 100 100
rcv-queue random-detect min-threshold 3 60 70 80 100
rcv-queue random-detect max-threshold 4 100 100 100 100
rcv-queue random-detect min-threshold 4 60 70 80 100
rcv-queue random-detect max-threshold 5 100 100 100 100
rcv-queue random-detect min-threshold 5 60 70 80 100
rcv-queue random-detect max-threshold 6 100 100 100 100
rcv-queue random-detect min-threshold 6 60 70 80 100
rcv-queue DSCP-map 1 1 0
rcv-queue DSCP-map 2 1 8
rcv-queue DSCP-map 3 1 14
rcv-queue DSCP-map 3 2 12
rcv-queue DSCP-map 3 3 10
rcv-queue DSCP-map 4 1 22
rcv-queue DSCP-map 4 2 20
rcv-queue DSCP-map 4 3 18
rcv-queue DSCP-map 5 1 30
rcv-queue DSCP-map 5 2 28
rcv-queue DSCP-map 5 3 26
rcv-queue DSCP-map 6 1 38
rcv-queue DSCP-map 6 2 36
rcv-queue DSCP-map 6 3 34
rcv-queue DSCP-map 7 1 16 24 48 56
rcv-queue DSCP-map 8 1 32 40 46
```

!

For all of the line cards that support the 8Q4T ingress queueing structure listed above, the ports are assumed to be uplink ports, because they are all TenGigabitEthernet ports. If the Catalyst 6500 Series with Sup-720 is deployed as a distribution or core switch, and the link connecting the switch port is not a trunk port, then all ingress traffic will not have an 802.1p header. However, because these line cards support DSCP-to-queue mapping, ingress traffic will still be mapped into the correct queue based on the DSCP value of the IP packets.

8Q8T Ingress Queuing

8Q8T ingress queuing is supported by the following line cards:

- WS-X6704-10GE with a DFC3C or DFC3XL (WS-F6700-DFC3C or WS-F6700-DFC3CXL)

How to determine if a line card has a CFC or DFC3C/3CXL was discussed in the *1Q8T Ingress Queuing* section of this document.

8Q8T ingress queueing for these line cards implements CoS-to-queue mapping, with CoS-based tail-drop for congestion avoidance.

WS-X67xx Series line cards do not support DSCP-based queue mapping–they only support ingress CoS-to-queue mapping. With the EasyQoS solution, trust of ingress CoS markings is never enabled on switch ports. Trust of ingress DSCP markings is enabled on uplink ports. As a result of this, there really is no need of provisioning the 8Q8T ingress queueing policy from APIC-EM EasyQoS to the WS-X6704-10GE with DFC3C/DFC3CXL line cards.

The following figure shows the ingress queueing model implemented by APIC-EM EasyQoS for these line cards.

Figure 72 **Queuing for WS-X6704-**10GE with a DFC3 or DFC3XL when "trust cos" is Not Configured

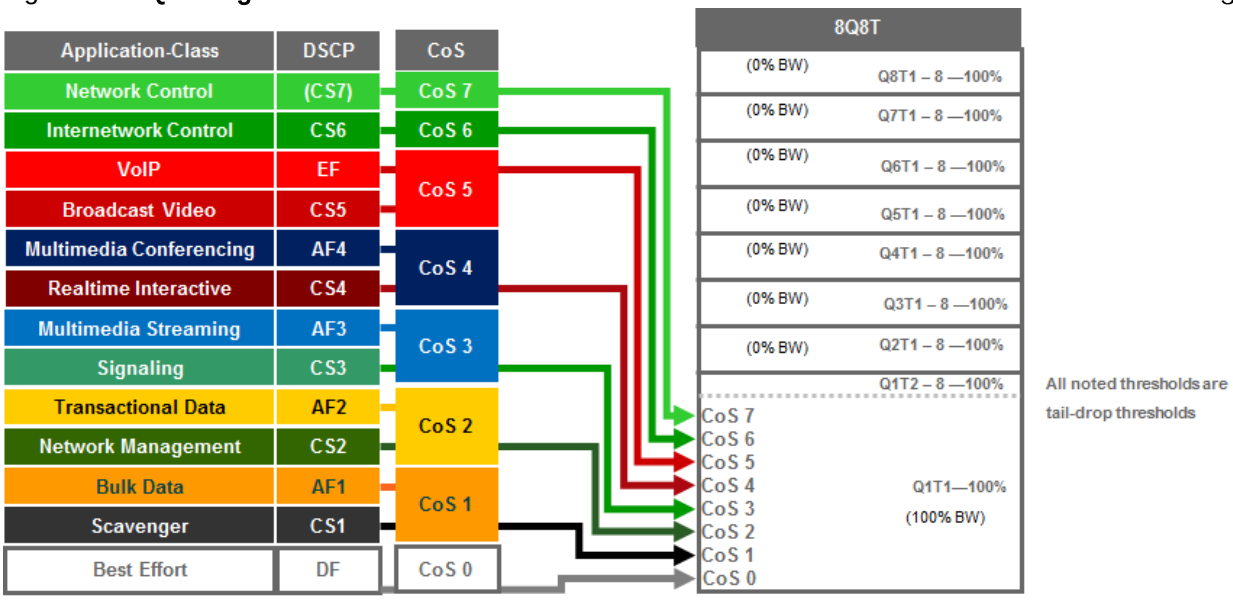

An example of the configuration provisioned by APIC-EM EasyQoS for these line cards is shown below.

!

```
interface TenGigabitEthernet x/x/x
 mls qos trust DSCP
! Switch ports functioning as uplink ports
!
```

As can be seen, DSCP trust is extended to switch ports functioning as uplink ports. Typically all ports on the WS-X6704-10GE line card would be assumed to be uplink ports, because they are all TenGigabitEthernet ports.

Because no trust CoS is extended on ingress switch ports and the line card does not support DSCP-to-queue mapping, all ingress traffic is treated as if it had a value of CoS 0, and mapped to Q1T1 with a tail-drop threshold of 100%.

1P3Q8T Egress Queuing

1P3Q8T egress queuing is supported by the following line cards:

- WS-X6724-SFP, WS-X6748-SFP, and WS-X6748-GE-TX with either CFC or DFC3/DFC3XL

1P3Q8T egress queueing for these line cards implements CoS-to-queue mapping, with CoS-based tail-drop for congestion avoidance.

The following figure shows the 1P3Q8T egress queueing model.

Figure 73 **1P3Q8T Egress Queuing Models—CoS-to-Queue Mapping with CoS-WRED**

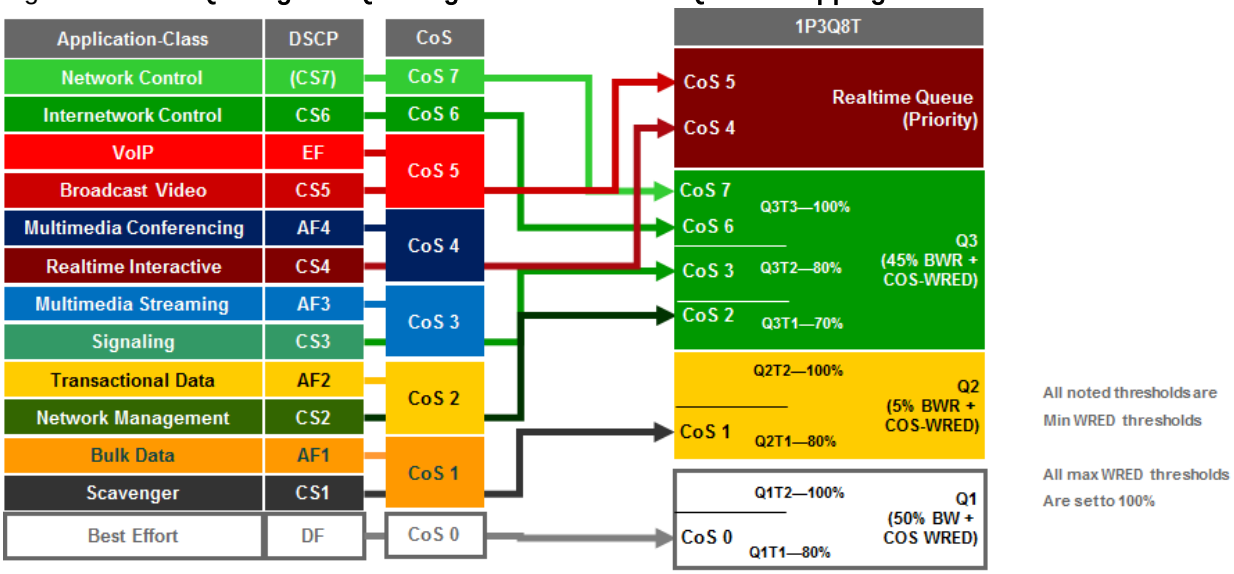

An example of the configuration provisioned by APIC-EM EasyQoS to a port on these line cards is shown below.

```
!
interface GigabitEthernet x/x/x
 wrr-queue queue-limit 40 15 40
 priority-queue queue-limit 15
```

```
 wrr-queue bandwidth 50 5 45
 wrr-queue random-detect 1
 wrr-queue random-detect 2
 wrr-queue random-detect 3
 wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
 wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
 wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
 wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
 wrr-queue random-detect max-threshold 3 100 100 100 100 100 100 100 100
 wrr-queue random-detect min-threshold 3 70 80 100 100 100 100 100 100
 wrr-queue cos-map 1 1 0
 wrr-queue cos-map 2 1 1
 wrr-queue cos-map 3 1 2
 wrr-queue cos-map 3 2 3
 wrr-queue cos-map 3 3 6 7
 priority-queue cos-map 1 4 5
!
```

Due to the internal ASIC structure of the ports on the WS-X6748-SFP and WS-X6748-GE-TX line cards, the egress queueing structure of the ports cannot be configured independently. Instead, APIC-EM EasyQoS will apply the queueing policy to groups of ports on the line card.

1P3Q4T Egress Queuing

1P3Q4T egress queuing is supported by the following line cards:

- All ports on the VS-S720-10G-3C and VS-S720-10G-3CXL when the Gigabit Ethernet ports are enabled

Enabling of the Gigabit Ethernet ports on the Sup-2T was discussed in the *2Q4T Ingress Queuing* section of this document.

1P3Q4T egress queueing for the Sup-720 implements CoS-to-queue mapping, with CoS-based WRED for congestion avoidance. The following figure shows the 1P3Q4T egress queueing model.

Figure 74 **1P3Q4T Egress Queuing Models—CoS-to-Queue Mapping with CoS-WRED**

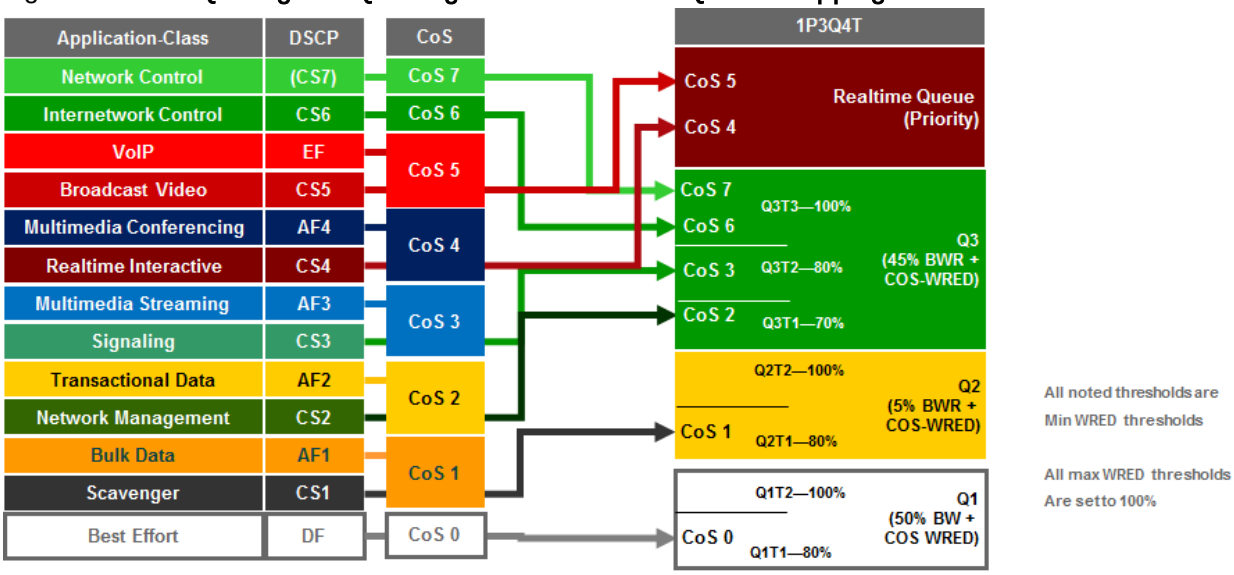

An example of the configuration provisioned by APIC-EM EasyQoS to both a Gigabit Ethernet and a TenGigabitEthernet port on these line cards is shown below.

```
!
interface GigabitEthernet x/x/x
 wrr-queue queue-limit 40 15 40
 priority-queue queue-limit 15
 wrr-queue bandwidth 50 5 45
 wrr-queue random-detect 1
 wrr-queue random-detect 2
 wrr-queue random-detect 3
 wrr-queue random-detect max-threshold 1 100 100 100 100
 wrr-queue random-detect min-threshold 1 80 100 100 100
 wrr-queue random-detect max-threshold 2 100 100 100 100
 wrr-queue random-detect min-threshold 2 80 100 100 100
 wrr-queue random-detect max-threshold 3 100 100 100 100
 wrr-queue random-detect min-threshold 3 70 80 100 100
 wrr-queue cos-map 1 1 0
 wrr-queue cos-map 2 1 1
 wrr-queue cos-map 3 1 2
 wrr-queue cos-map 3 2 3
 wrr-queue cos-map 3 3 6 7
```

```
 priority-queue cos-map 1 4 5
!
interface TenGigabitEthernet x/x/x
 wrr-queue queue-limit 40 15 40
 priority-queue queue-limit 15
 wrr-queue bandwidth 50 5 45
 wrr-queue random-detect 1
 wrr-queue random-detect 2
 wrr-queue random-detect 3
 wrr-queue random-detect max-threshold 1 100 100 100 100
 wrr-queue random-detect min-threshold 1 80 100 100 100
 wrr-queue random-detect max-threshold 2 100 100 100 100
 wrr-queue random-detect min-threshold 2 80 100 100 100
 wrr-queue random-detect max-threshold 3 100 100 100 100
 wrr-queue random-detect min-threshold 3 70 80 100 100
 wrr-queue cos-map 1 1 0
 wrr-queue cos-map 2 1 1
 wrr-queue cos-map 3 1 2
 wrr-queue cos-map 3 2 3
 wrr-queue cos-map 3 3 6 7
 priority-queue cos-map 1 4 5
!
```

As discussed in the *2Q4T Ingress Queuing* section of this document, QoS service policies cannot be applied to TenGigabitEthernet ports on the Sup-720, when the TenGigabitEthernet ports are part of a port-channel group that is part of a VSL. Also, QoS service policies cannot be applied to Gigabit Ethernet ports on the Sup-720 when the TenGigabitEthernet ports are part of a VSL.

Finally, due to the internal ASIC structure of the ports on the Sup-720, the egress queueing structure of the ports cannot be configured independently. Instead, APIC-EM EasyQoS will apply the queueing policy to groups of ports on the supervisor. This was discussed in the *2Q4T Ingress Queuing* section of this document, as well.

1P7Q4T Egress Queuing

1P7Q4T egress queuing is supported by the following line cards:

- WS-X6708-10G-3C and WS-X6708-10G-3CXL

- VS-S720-10G-3C and VS-S720-10G-3CXL TenGigabitEthernet ports when the Gigabit Ethernet ports are inactive

1P7Q4T egress queueing for these line cards implements DSCP-to-queue mapping, with DSCP-based WRED for congestion avoidance. The following figure shows the 1P7Q4T egress queueing model.

Figure 75 **1P7Q4T Egress Queuing Models—DSCP-to-Queue Mapping with DSCP-based WRED**

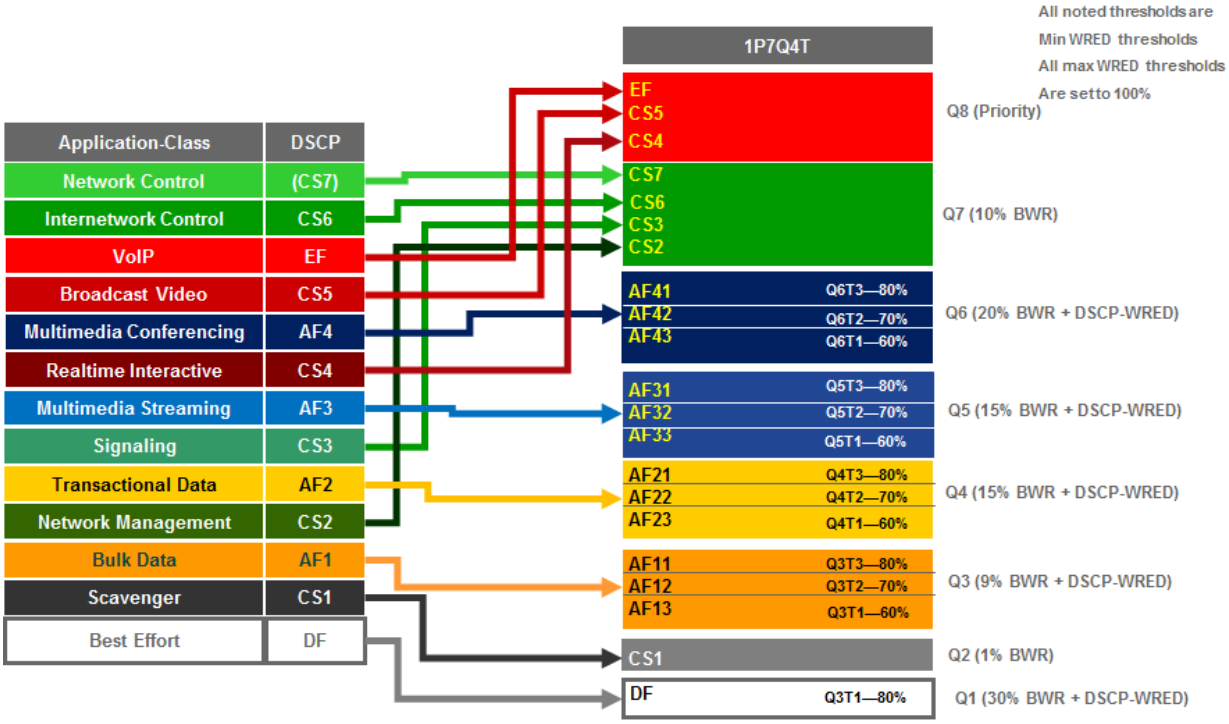

An example of the configuration provisioned by APIC-EM EasyQoS to a TenGigabitEthernet port on these line cards is shown below.

```
!
interface range TenGigabitEthernet x/x-x
 wrr-queue queue-limit 25 10 10 10 10 10 10
 wrr-queue bandwidth 30 1 9 15 15 20 10
 priority-queue queue-limit 15
 wrr-queue random-detect 1
 no wrr-queue random-detect 2
 wrr-queue random-detect 3
 wrr-queue random-detect 4
 wrr-queue random-detect 5
 wrr-queue random-detect 6
 no wrr-queue random-detect 7
```

```
 wrr-queue random-detect max-threshold 1 100 100 100 100
 wrr-queue random-detect min-threshold 1 80 100 100 100
 wrr-queue random-detect max-threshold 3 100 100 100 100
 wrr-queue random-detect min-threshold 3 60 70 80 100
 wrr-queue random-detect max-threshold 4 100 100 100 100
 wrr-queue random-detect min-threshold 4 60 70 80 100
 wrr-queue random-detect max-threshold 5 100 100 100 100
 wrr-queue random-detect min-threshold 5 60 70 80 100
 wrr-queue random-detect max-threshold 6 100 100 100 100
 wrr-queue random-detect min-threshold 6 60 70 80 100
 mls qos queue-mode mode-DSCP
wrr-queue DSCP-map 1 1 0
 wrr-queue DSCP-map 2 1 8
 wrr-queue DSCP-map 3 1 14
 wrr-queue DSCP-map 3 2 12
 wrr-queue DSCP-map 3 3 10
 wrr-queue DSCP-map 4 1 22
 wrr-queue DSCP-map 4 2 20
 wrr-queue DSCP-map 4 3 18
 wrr-queue DSCP-map 5 1 30
 wrr-queue DSCP-map 5 2 28
 wrr-queue DSCP-map 5 3 26
 wrr-queue DSCP-map 6 1 38
 wrr-queue DSCP-map 6 2 36
 wrr-queue DSCP-map 6 3 34
 wrr-queue DSCP-map 7 1 16 24 48 56
 priority-queue DSCP-map 1 32 40 46
!
```

1P7Q8T Egress Queuing

1P7Q8T egress queuing is supported by the following line cards:

- WS-X6704-10GE with either CFC or DFC3/DFC3XL

1P7Q8T egress queueing for these line cards implements CoS-to-queue mapping, with CoS-based tail-drop for congestion avoidance. Note that due to the combination of 8 queues and only 8 CoS values, tail-drop thresholds are not used in this design.

The following figure shows the 1P7Q8T egress queuing model.

Figure 76 **1P7Q8T Egress Queuing Models—CoS-to-Queue Mapping with COS-based WRED**

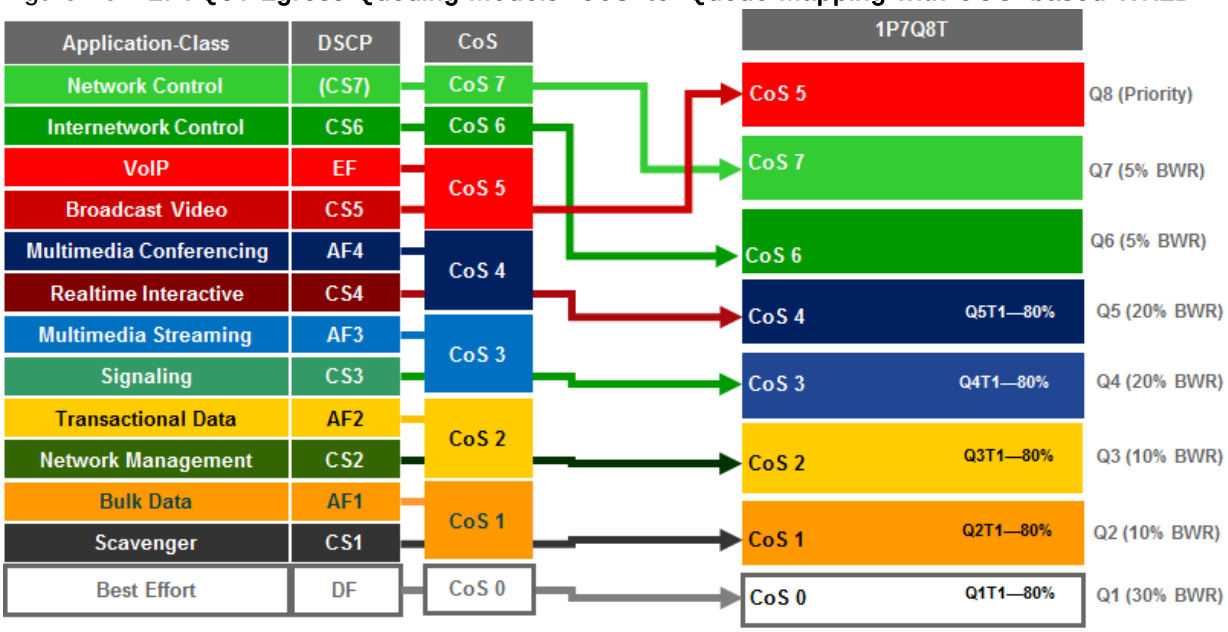

An example of the configuration provisioned by APIC-EM EasyQoS to a TenGigabitEthernet port on these line cards is shown below.

```
!
interface TenGigabitEthernet x/x/x
 wrr-queue queue-limit 25 10 10 15 15 5 5
 wrr-queue bandwidth 30 10 10 20 20 5 5
 priority-queue queue-limit 15
 wrr-queue random-detect 1
 wrr-queue random-detect 2
 wrr-queue random-detect 3
 wrr-queue random-detect 4
 wrr-queue random-detect 5
 no wrr-queue random-detect 6
 no wrr-queue random-detect 7
 wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
 wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
```

```
 wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
 wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
 wrr-queue random-detect max-threshold 3 100 100 100 100 100 100 100 100
 wrr-queue random-detect min-threshold 3 80 100 100 100 100 100 100 100
 wrr-queue random-detect max-threshold 4 100 100 100 100 100 100 100 100
 wrr-queue random-detect min-threshold 4 80 100 100 100 100 100 100 100
 wrr-queue random-detect max-threshold 5 100 100 100 100 100 100 100 100
 wrr-queue random-detect min-threshold 5 80 100 100 100 100 100 100 100
 wrr-queue cos-map 1 1 0
 wrr-queue cos-map 2 1 1
 wrr-queue cos-map 3 1 2
 wrr-queue cos-map 4 1 3
 wrr-queue cos-map 5 1 4
 wrr-queue cos-map 6 1 6
 wrr-queue cos-map 7 1 7
 priority-queue cos-map 1 5
!
```

Cisco Nexus 7000/7700 Queuing Design

The only role Nexus 7000 and 7700 Series switches have within the EasyQoS solution is as a campus-core switch. Only ingress and egress queuing policies are pushed from APIC-EM to these switches. No ingress classification & marking policies are pushed from APIC-EM to Nexus 7000 and 7700 Series switches

The following sections detail the ingress and egress queuing structures pushed by APIC-EM to the ports of each of the Nexus 7000 and 7700 Series modules supported by EasyQoS. Ingress & egress queueing structures are dependent upon the type of module and the chassis in which they are installed (Nexus 7000 versus Nexus 7700), as shown below:

- M2 Series modules are only supported on Nexus 7000
  - 8Q2T-IN/1P7Q4T-OUT queuing
- F2 Series modules are only supported on Nexus 7000
  - 4Q1T-IN/1P3Q1T-OUT queuing
- F2E Series modules are supported on both the Nexus 7000 and the Nexus 7700
  - F2E Series modules on the Nexus 7000 have identical queuing to F2 Series modules (4Q1T-IN/1P3Q1T-OUT)

  - F2E Series modules on the Nexus 7700 have identical queuing to F3 Series modules (4Q1T-IN/1P7Q1T-OUT)

- F3 Series modules are supported on both the Nexus 7000 and the Nexus 7700
  - F3 Series modules on the Nexus 7000 have identical queuing to F2 Series modules (4Q1T-IN/1P3Q1T-OUT)
  - F3 Series modules on the Nexus 7700 implement 4Q1T-IN/1P7Q1T-OUT queuing
- M3 Series modules are only supported on the Nexus 7700
  - 4Q1T-IN/1P7Q1T-OUT queuing

Note: In APIC-EM release 1.3, EasyQoS only supports the Nexus 7000 or 7700 Series platforms configured with a single default VDC. Changing network QoS requires being logged into the default VDC. Changes to the system class-maps are made only on the default VDC but take effect immediately across all VDCs. Queuing policy maps are defined per VDC.

M2 Series Modules on the Nexus 7000

M2 Series modules consist of the following:

- N7K-M224XP-23L
- N7K-M206FQ-23L
- N7K-M202CF-22L

These modules are only supported on the Nexus 7000 series chassis. M2 Series modules implement 8Q2T ingress queueing and 1P7Q4T egress queuing.

Nexus 7000 8Q2T Ingress Queuing

As of NX OS software version 6.2.2 and higher, ingress DSCP-to-Queue mapping is supported on M2 Series modules and must be configured via the following global configuration command.

```
!
hardware qos DSCP-to-queue ingress module-type all
!
```

EasyQoS will provision this command to enable DSCP-to-Queue mapping support on the M2 Series modules.

Trust of both CoS and DSCP values is implicit on the Nexus 7000 Series. No explicit configuration is required to enable trust. Because both CoS and DSCP values are trusted, queuing models for both CoS-to-queue mapping and DSCP-to-queue mapping need to be configured by EasyQoS.

The following figure shows the CoS-to-queue mapping for the 8Q2T ingress queuing model.

Figure 77 Ingress Queuing Model (8Q2T)—CoS-to-Queue Mapping

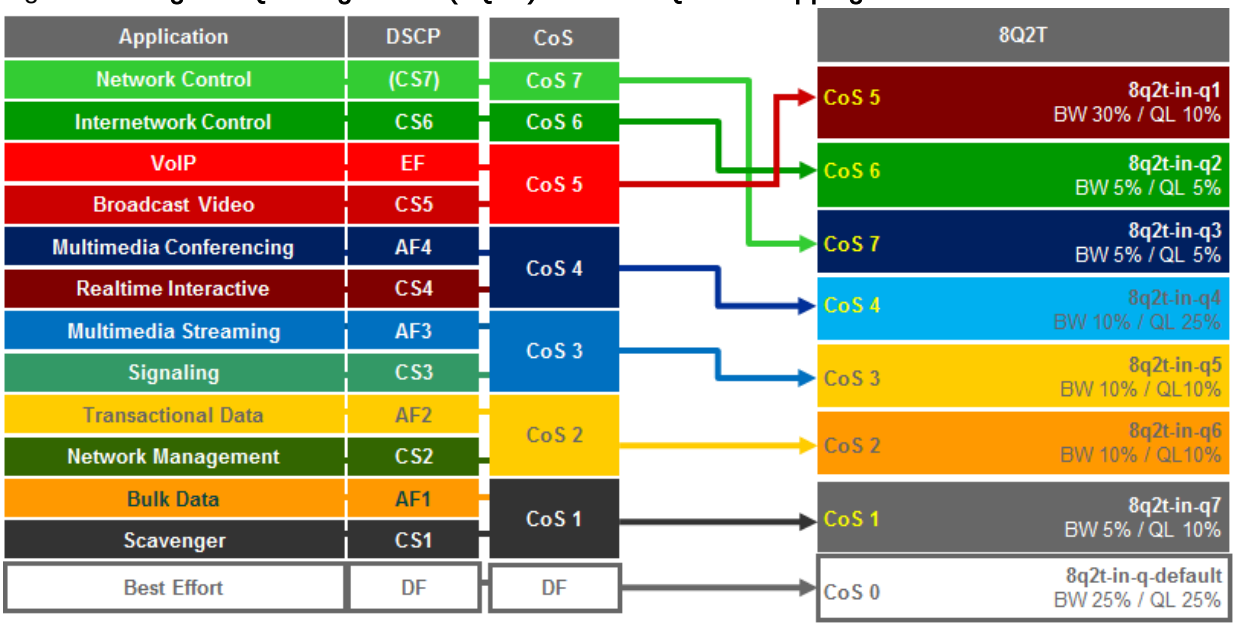

The following figure shows the DSCP-to-queue mapping for the 8Q2T ingress queuing model.

Figure 78 Ingress Queuing Model (8Q2T)—DSCP-to-Queue Mapping

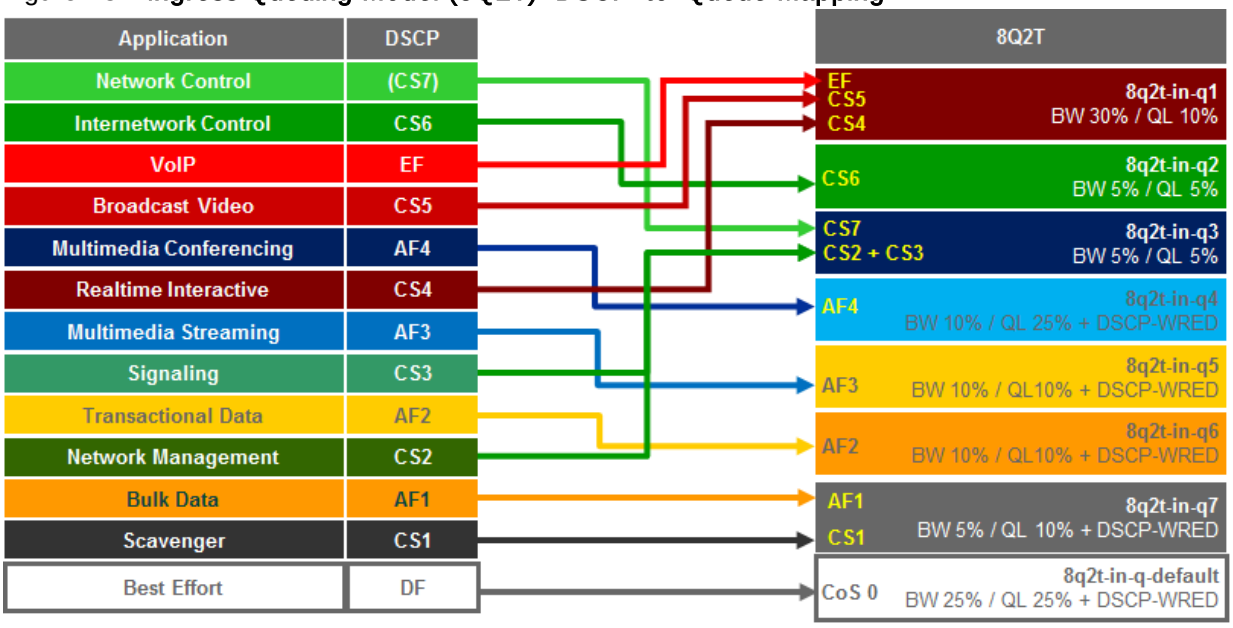

NX OS provides system-defined class-map names that cannot be renamed. The system-defined class-map names that match the eight ingress queues of the M2 Series modules are as follows:

- 8q2t-in-q1
- 8q2t-in-q2
- 8q2t-in-q3

- 8q2t-in-q4
- 8q2t-in-q5
- 8q2t-in-q6
- 8q2t-in-q7
- 8q2t-in-q-default

These class-maps have default and/or non-default (customer implemented) CoS and/or DSCP values attached to them. These values can be reset with “no match” commands.

A “no match DSCP” command causes one or more DSCP values (depending upon whether a single DSCP value “x” was entered or a range of DSCP values “x-y” was entered) attached to a particular class-map to be removed and automatically attached to the 8q2t-in-q-default class-map. However, a CoS value must first be mapped to the 8q2t-in-q-default class-map, or an error will be generated.

A “no match cos” command causes one or more CoS values (depending upon whether a single CoS value “x” was entered or a range of CoS values “x-y” was entered) attached to a particular class-map to be removed and automatically attached to the 8q2t-in-q-default class-map. However, all CoS values cannot be removed from a particular class-map if there are still DSCP values associated with the particular class-map. Attempting this will generate an error. In other words, all CoS values cannot be removed from a given class-map until all DSCP values have been removed from the class-map. Finally, “no match” commands are not accepted on the 8q2t-in-q-default class-map and will generate an error.

Because APIC-EM EasyQoS has no knowledge of the existing mapping of the DSCP and CoS values to the class-maps before configuring the ingress queuing policy, it will first set all CoS and DSCP values to the 8q2t-in-q-default class-map before moving CoS and DSCP values to their desired class-maps. EasyQoS then moves the DSCP and CoS values into the desired class-maps. This is accomplished via the following commands.

```
!
class-map type queuing match-any 8q2t-in-q1
  match DSCP 32, 40, 46
  match cos 5
class-map type queuing match-any 8q2t-in-q2
  match DSCP 48
  match cos 6
class-map type queuing match-any 8q2t-in-q3
  match DSCP 16, 24, 56
  match cos 7
class-map type queuing match-any 8q2t-in-q4
  match DSCP 34, 36, 38
  match cos 4
```

```
   class-map type queuing match-any 8q2t-in-q5
     match DSCP 26, 28, 30
     match cos 3
   class-map type queuing match-any 8q2t-in-q6
     match DSCP 18, 20, 22
     match cos 2
   class-map type queuing match-any 8q2t-in-q7
     match DSCP 8, 10, 12, 14
     match cos 1
   !
```

After the DSCP and CoS values have been moved into the appropriate class-maps, the following configuration provisioned by APIC-EM EasyQoS creates and configures the queueing policy-map with the 8Q2T ingress queuing structure on the Nexus 7000 Series M2 modules.

```
!
policy-map type queuing prm-DSCP#8q2t-in
  class type queuing 8q2t-in-q1
    bandwidth percent 30
    queue-limit percent 10
  class type queuing 8q2t-in-q2
    bandwidth percent 5
    queue-limit percent 5
  class type queuing 8q2t-in-q3
    bandwidth percent 5
    queue-limit percent 5
  class type queuing 8q2t-in-q4
    bandwidth percent 10
    queue-limit percent 25
    random-detect DSCP-based
    random-detect DSCP 34 minimum-threshold percent 80 maximum-threshold percent
100
    random-detect DSCP 36 minimum-threshold percent 80 maximum-threshold percent
100
    random-detect DSCP 38 minimum-threshold percent 80 maximum-threshold percent
100
```

```
  class type queuing 8q2t-in-q5
    bandwidth percent 10
    queue-limit percent 10
    random-detect DSCP-based
    random-detect DSCP 26 minimum-threshold percent 80 maximum-threshold percent
100
    random-detect DSCP 28 minimum-threshold percent 80 maximum-threshold percent
100
    random-detect DSCP 30 minimum-threshold percent 80 maximum-threshold percent
100
  class type queuing 8q2t-in-q6
    bandwidth percent 10
    queue-limit percent 10
    random-detect DSCP-based
    random-detect DSCP 18 minimum-threshold percent 80 maximum-threshold percent
100
    random-detect DSCP 20 minimum-threshold percent 80 maximum-threshold percent
100
    random-detect DSCP 22 minimum-threshold percent 80 maximum-threshold percent
100
  class type queuing 8q2t-in-q7
    bandwidth percent 5
    queue-limit percent 10
    random-detect DSCP-based
    random-detect DSCP 8 minimum-threshold percent 80 maximum-threshold percent
100
    random-detect DSCP 10 minimum-threshold percent 80 maximum-threshold percent
100
    random-detect DSCP 12 minimum-threshold percent 80 maximum-threshold percent
100
    random-detect DSCP 14 minimum-threshold percent 80 maximum-threshold percent
100
  class type queuing 8q2t-in-q-default
    bandwidth percent 25
    queue-limit percent 25
    random-detect DSCP-based
```

```
    random-detect DSCP 0 minimum-threshold percent 80 maximum-threshold percent
100
!
```

The policy-map with the 8Q2T ingress queuing structure is then applied by APIC-EM EasyQoS to Ethernet interfaces that connect to either other core-layer switches or to distribution-layer switches. Additionally, the policy-map can be applied to the logical port-channel interface when an EtherChannel connection is used for port-level resilience, instead of a single physical interface. An example of the configuration of the policy to an Ethernet and a Port-channel interface is shown below.

```
!
interface Ethernet x/x
 service-policy type queuing input prm-DSCP#8q2t-in
!
interface port-channel xxx
 service-policy type queuing input prm-DSCP#8q2t-in
!
```

Nexus 7000 1P7Q4T Egress Queuing

For egress queuing, only CoS-to-Queue mapping is supported on M2 Series modules. The following figure shows the Cos-to-queue mapping for the 1P7Q4T egress queueing model.

Figure 79 **Egress Queuing Model (1P7Q4T)—CoS-to-Queue Mapping**

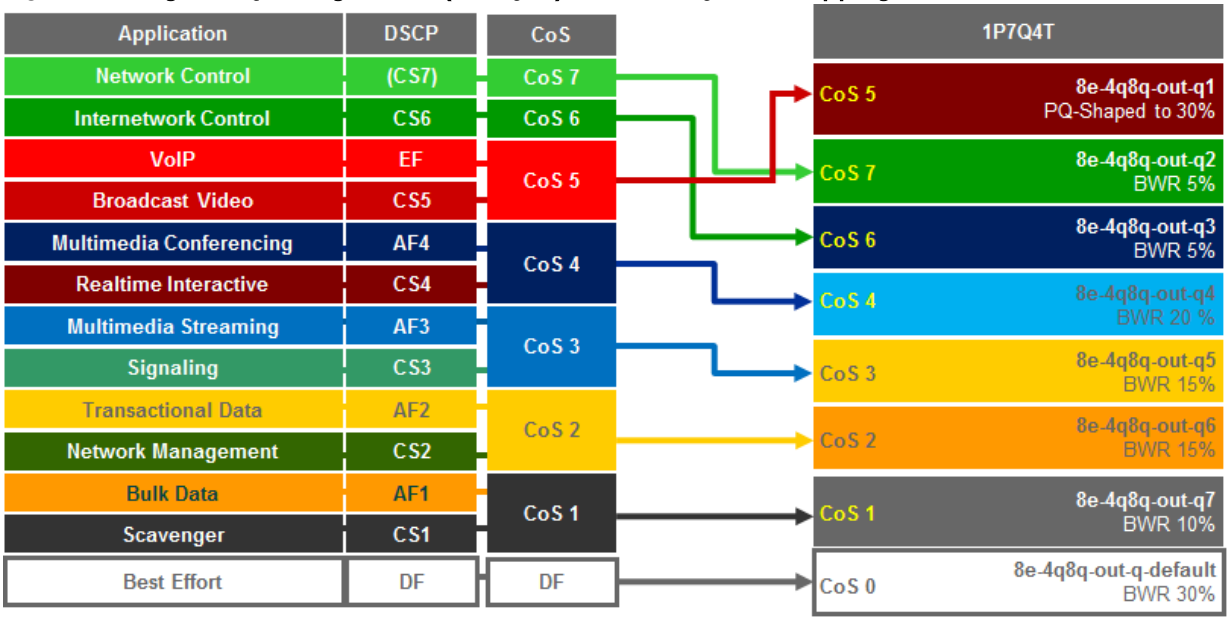

NX OS provides system-defined class-map names that cannot be renamed. The system-defined class-map names that match the eight output queues of the M2 Series modules are as follows:

- 1p7q4t-out-pq1

- 1p7q4t-out-q2
- 1p7q4t-out-q3
- 1p7q4t-out-q4
- 1p7q4t-out-q5
- 1p7q4t-out-q6
- 1p7q4t-out-q7
- 1p7q4t-out-q-default

These class-maps have default and/or non-default (customer implemented) CoS values attached to them. These values can be reset with “no match cos” commands. A “no match cos” command causes one or more CoS values (depending upon whether a single CoS value “x” was entered or a range of CoS values “x-y” was entered) attached to a particular class-map to be removed and automatically attached to the 1p7q4t-out-q-default class-map.

Because APIC-EM EasyQoS has no knowledge of the existing mapping of the CoS values to the class-maps before configuring the ingress queuing policy, it will first set all CoS values to the 1p7q4t-out-q-default class-map before moving CoS values to their desired class-maps. This is accomplished via the following commands.

```
!
class-map type queuing match-any 1p7q4t-out-pq1
  no match cos 0-7
class-map type queuing match-any 1p7q4t-out-q2
  no match cos 0-7
class-map type queuing match-any 1p7q4t-out-q3
  no match cos 0-7
class-map type queuing match-any 1p7q4t-out-q4
  no match cos 0-7
class-map type queuing match-any 1p7q4t-out-q5
  no match cos 0-7
class-map type queuing match-any 1p7q4t-out-q6
  no match cos 0-7
class-map type queuing match-any 1p7q4t-out-q7
  no match cos 0-7
!
```

Note that because there are no DSCP values within egress class-maps, there are no issues with moving CoS values to the default class-map, as was discussed in the *Nexus 7000 8Q2T Ingress Queuing* section.

APIC-EM EasyQoS can then move CoS values into the desired class-maps. This is accomplished via the following commands.

```
!
class-map type queuing match-any 1p7q4t-out-pq1
  match cos 5
class-map type queuing match-any 1p7q4t-out-q2
  match cos 7
class-map type queuing match-any 1p7q4t-out-q3
  match cos 6
class-map type queuing match-any 1p7q4t-out-q4
  match cos 4
class-map type queuing match-any 1p7q4t-out-q5
  match cos 3
class-map type queuing match-any 1p7q4t-out-q6
  match cos 2
class-map type queuing match-any 1p7q4t-out-q7
  match cos 1
!
```

After the CoS values have been moved into the appropriate class-maps, the following configuration provisioned by APIC-EM EasyQoS creates and configures the queueing policy-map with the 1P7Q4T egress queuing structure on the Nexus 7000 Series M2 modules.

```
!
policy-map type queuing prm-DSCP#1p7q4t-out
  class type queuing 1p7q4t-out-pq1
    priority
    shape average percent 30
    queue-limit percent 10
  class type queuing 1p7q4t-out-q2
    bandwidth remaining percent 5
    queue-limit percent 5
  class type queuing 1p7q4t-out-q3
```

```
    bandwidth remaining percent 5
    queue-limit percent 5
  class type queuing 1p7q4t-out-q4
    bandwidth remaining percent 20
    queue-limit percent 25
    random-detect cos-based
    random-detect cos 4 minimum-threshold percent 80 maximum-threshold percent
100
  class type queuing 1p7q4t-out-q5
    bandwidth remaining percent 15
    queue-limit percent 10
    random-detect cos-based
    random-detect cos 3 minimum-threshold percent 80 maximum-threshold percent
100
  class type queuing 1p7q4t-out-q6
    bandwidth remaining percent 15
    queue-limit percent 10
    random-detect cos-based
    random-detect cos 2 minimum-threshold percent 80 maximum-threshold percent
100
  class type queuing 1p7q4t-out-q7
    bandwidth remaining percent 10
    queue-limit percent 10
    random-detect cos-based
    random-detect cos 1 minimum-threshold percent 80 maximum-threshold percent
100
  class type queuing 1p7q4t-out-q-default
    bandwidth remaining percent 30
    queue-limit percent 25
    random-detect cos-based
    random-detect cos 0 minimum-threshold percent 80 maximum-threshold percent
100
!
```

The policy-map with the 1P7Q4T egress queuing structure is then applied to Ethernet interfaces that connect to either other core-layer switches or to distribution-layer switches. Additionally, the policy-map can be applied to the logical port-channel interface when an EtherChannel connection is used for port-level resilience, instead of a single physical interface. An example of the configuration of the policy to an Ethernet and a Port-channel interface is shown below.

```
!
interface Ethernet x/x
 service-policy type queuing output prm-DSCP#1p7q4t-out
!
interface port-channel xxx
 service-policy type queuing output prm-DSCP#1p7q4t-out
!
```

F2/F2e/F3 Modules on the Nexus 7000

F2, F2e, and F3 Series modules for the Nexus 7000 consist of the following:

- N7K-F248XP-25
- N7K-F248XP-25E
- N7K-F248XT-25E
- N7K-F348XP-25
- N7K-F312FQ-25
- N7K-F306CK-25

These modules are only supported on the Nexus 7000 series chassis.

The ingress and egress queuing structure for F2, F2e, and F3 Series modules on the Nexus 7000 is determined by the system network QoS policy. When an F2, F2e, or F3 Series module is inserted and becomes operational within the Nexus 7000 chassis, five network-qos policies (templates) are automatically added to the system. The following figure shows the five network-QoS templates available in the Nexus 7000 chassis.

Figure 80 **NX-OS Network-QoS Templates for the F2, F2e, and F3 Series Modules on the Nexus 7000**

Network QoS Template	*CoS Values That May Be Dropped*	*No Drop CoS Values*
default-4q-8e-policy (default-nq-8e-policy)	0,1,2,3,4,5,6,7	-
default-4q-7e-policy (default-nq-7e-policy)	0,1,2,4,5,6,7	3
default-4q-6e-policy (default-nq-6e-policy)	0,1,2,5,6,7	3,4
default-4q-4e-policy (default-nq-4e-policy)	0,5,6,7	1,2,3,4
default-8e-4q4q-policy (default-nq-8e-4q4q-policy)	0,1,2,3,4,5,6,7	-
default-8e-4q8q-policy (default-nq-8e-4q8q-policy)	0,1,2,3,4,5,6,7	-
default-7e-4q8q-policy (default-nq-7e-4q8q-policy)	0,1,2,4,5,6,7	3
default-6e-4q8q-policy (default-nq-6e-4q8q-policy)	0,1,2,5,6,7	3,4
default-4e-4q8q-policy (default-nq-4e-4q8q-policy)	0,5,6,7	1,2,3,4

These Network QoS Templates are only available on the Nexus 7700 Series

Network-qos policy template names contain the abbreviation “nq”, referring to network-qos, and “e,” referring to Ethernet. The numbers 4, 6, 7, and 8 denote the number of drop CoS values (in other words the number of CoS values that may be dropped during congestion) that are defined within the policy template.

Note: The bottom four templates are only available in the Nexus 7700 Series chassis and not the Nexus 7000 Series chassis.

For data center deployments, there may be requirements for a particular CoS value not to be dropped during congestion. However, campus core deployments typically do not have such requirements. Hence, the default-nq-8e-policy or default-nq-8e-4q4q-policy templates are more appropriate for campus core deployments featuring Nexus 7000 chassis with F2, F2e, and/or F3 Series modules. By default, the default-nq-8e-policy is applied by the system.

The “4q4q” in the default-nq-8e-4q4q-policy template refers to the fact that both the ingress and egress queuing structures use four queues. The default-nq-8e-policy template uses an ingress queuing structure that uses only two queues and an egress queuing structure that uses four queues. Because the default-nq-8e-4q4q-policy template better uses the available ingress queuing structure, this is the network-qos policy template that APIC-EM EasyQoS configures via the following global configuration commands.

```
!
system qos
 service-policy type network-qos default-nq-8e-4q4q-policy
!
```

This configuration change can be validated via exec-level “sow policy-map system” command. An example of the output is shown below.

```
show policy-map system
```

```
 Type network-qos policy-maps
 ============================
 policy-map type network-qos default-nq-8e-4q4q-policy template 8e-4q4q
   class type network-qos c-nq-8e-4q4q
     match cos 0-7
     congestion-control tail-drop
     mtu 1500
…
 Service-policy input:   default-8e-4q4q-in-policy
…
 Service-policy output:   default-8e-4q4q-out-policy
…
```

The following figure shows the default ingress queuing models for each of the network-qos policy templates.

Figure 81 **Default Ingress Queuing Models for the Network-QoS Templates**

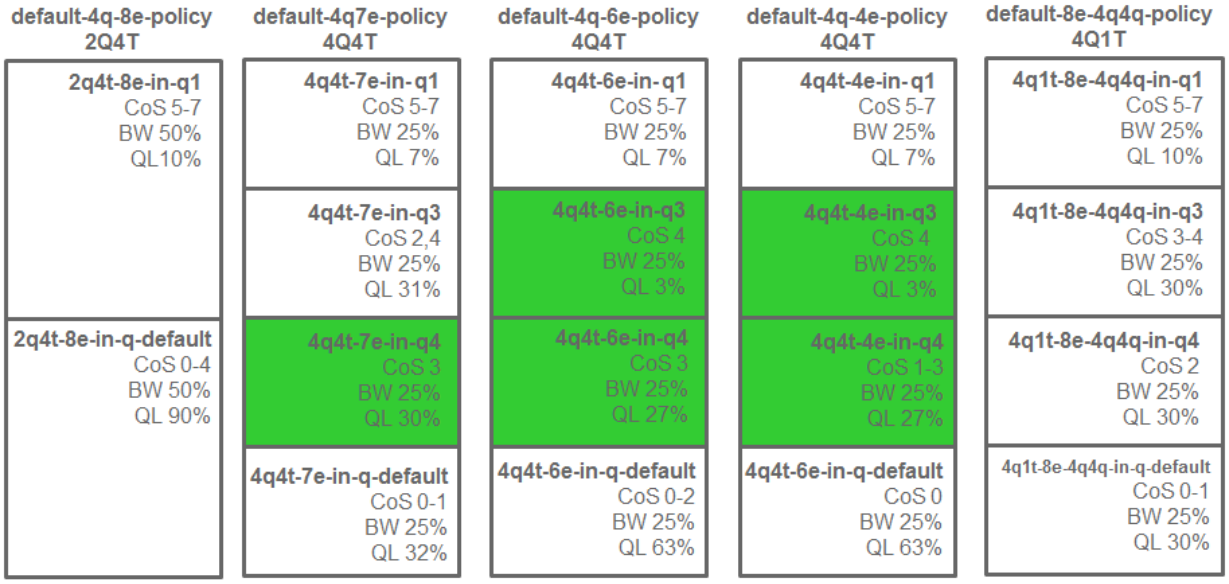

The following figure shows the default egress queuing models for each of the network-qos policy templates.

Figure 82 **Default Egress Queuing Models for the Network-QoS Templates**

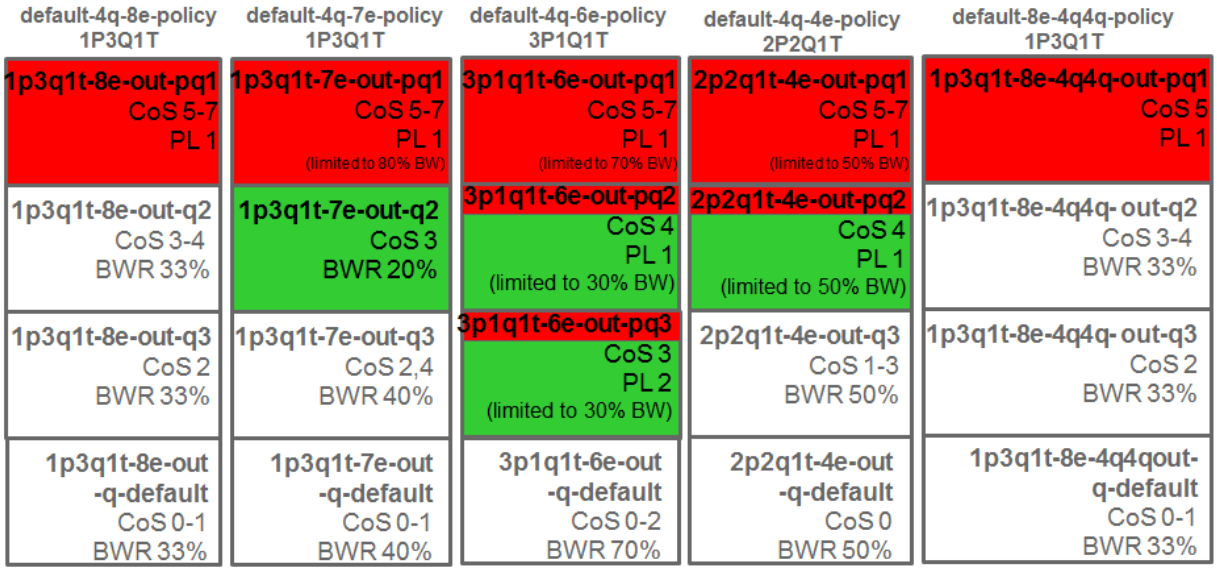

As can be seen from the figures above, the default-nq-8e-4q4q-policy implements a 4Q1T ingress queueing structure and a1P3Q1T egress queuing structure for F2, F2e, and F3 Series modules with the Nexus 7000.

Nexus 7000 4Q1T Ingress Queuing

As of NX OS software version 6.2.2 and higher ingress DSCP-to-Queue mapping is supported on F2, F2e, and F3 Series modules and must be configured via the following global configuration command.

```
!
hardware qos DSCP-to-queue ingress module-type all
!
```

APIC-EM EasyQoS will provision this command to enable DSCP-to-Queue mapping support on the F2/F2E/F3 Series modules for the Nexus 7000 Series.

Trust of both CoS and DSCP values is implicit on the Nexus 7000 Series. No explicit configuration is required to enable trust.

The network-qos default-nq-8e-4q4q-policy template has a default ingress queuing policy map that can be displayed via the exec-level “show policy-map type queuing default-8e-4q4q-in-policy” command. An example of the output is shown below.

```
DC-7010-2# show policy-map type queuing default-8e-4q4q-in-policy


  Type queuing policy-maps
  ========================


  policy-map type queuing default-8e-4q4q-in-policy
```

```
    class type queuing 4q1t-8e-4q4q-in-q1
      queue-limit percent 10
      bandwidth percent 25
    class type queuing 4q1t-8e-4q4q-in-q-default
      queue-limit percent 30
      bandwidth percent 25
    class type queuing 4q1t-8e-4q4q-in-q3
      queue-limit percent 30
      bandwidth percent 25
    class type queuing 4q1t-8e-4q4q-in-q4
      queue-limit percent 30
      bandwidth percent 25
```

For F2, F2e, and F3 series modules, the ingress queuing policy define queue-limit percentages and bandwidth percentages for each of the class-maps (which correspond to the queues). These percentages will be modified by APIC-EM for the EasyQoS solution.

The network-qos default-nq-8e-4q4q-policy template also provides default settings for the system-defined class-maps that can be displayed via the exec-level “show class-map type queuing” command for each of the system-defined class-map names. An example of the output for each of the system-defined class-maps is shown below.

```
DC-7010-2# show class-map type queuing 4q1t-8e-4q4q-in-q1

  Type queuing class-maps
  ========================

    class-map type queuing match-any 4q1t-8e-4q4q-in-q1
      Description: Classifier for Ingress queue 1 of type 4q1t-8e-4q4q
      match cos 5-7
      match DSCP 40-63

DC-7010-2# show class-map type queuing 4q1t-8e-4q4q-in-q-default

  Type queuing class-maps
  ========================
```

```
    class-map type queuing match-any 4q1t-8e-4q4q-in-q-default
      Description: Classifier for Ingress queue 2 of type 4q1t-8e-4q4q
      match cos 0-1
      match DSCP 0-15

DC-7010-2# show class-map type queuing  4q1t-8e-4q4q-in-q3

  Type queuing class-maps
  =======================

    class-map type queuing match-any 4q1t-8e-4q4q-in-q3
      Description: Classifier for Ingress queue 3 of type 4q1t-8e-4q4q
      match cos 3-4
      match DSCP 24-39

DC-7010-2# show class-map type queuing  4q1t-8e-4q4q-in-q4

  Type queuing class-maps
  =======================

    class-map type queuing match-any 4q1t-8e-4q4q-in-q4
      Description: Classifier for Ingress queue 4 of type 4q1t-8e-4q4q
      match cos 2
      match DSCP 16-23
```

The system-defined class-maps have default settings for mapping CoS and DSCP values to each of the class-maps. However, APIC-EM EasyQoS cannot guarantee that the customer has not already previously configured the default-nq-8e-4q4q-policy network-qos template and modified the default mappings of CoS and DSCP values to class-maps. Because both CoS and DSCP values are trusted, queuing models for both CoS-to-queue mapping and DSCP-to-queue mapping need to be configured.

The following figure shows the Cos-to-queue mapping for the 4Q1T ingress queueing model deployed by APIC-EM for the EasyQoS solution.

Figure 83 Ingress Queuing Model (4Q1T)—CoS-to-Queue Mapping

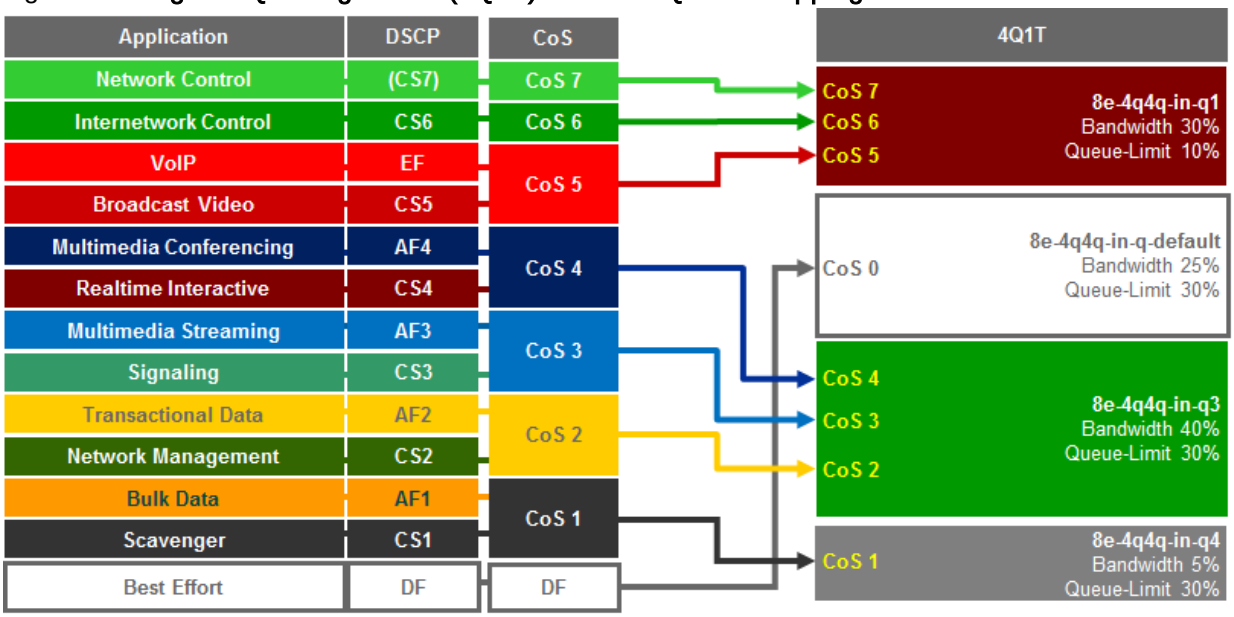

The following figure shows the DSCP-to-queue mapping for the 4Q1T ingress queueing model deployed by APIC-EM for the EasyQoS solution.

Figure 84 Ingress Queuing Model (4Q1T)—DSCP-to-Queue Mapping

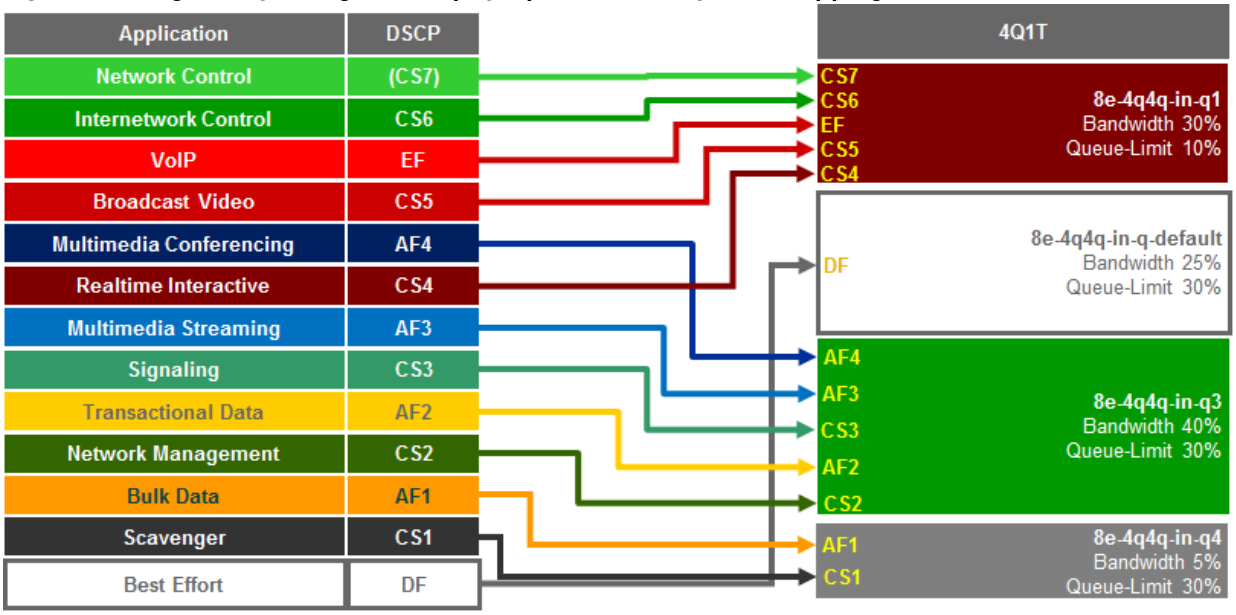

NX OS provides system-defined class-map names that cannot be renamed. The system-defined class-map names that match the four ingress queues of the F2, F2e and/or F3 Series modules for the Nexus 7000 are as follows:

- 4q1t-8e-4q4q-in-q1
- 4q1t-8e-4q4q-in-q3
- 4q1t-8e-4q4q-in-q4

- 4q1t-8e-4q4q-in-q-default

These class-maps have default and/or non-default (customer implemented) CoS and/or DSCP values attached to them. These values can be reset with “no match” commands.

A “no match DSCP” command causes one or more DSCP values (depending upon whether a single DSCP value “x” was entered or a range of DSCP values “x-y” was entered) attached to a particular class-map to be removed and automatically attached to the 4q1t-8e-4q4q-in-q-default class-map. However, a CoS value must first be mapped to the 4q1t-8e-4q4q-in-q-default class-map, or an error will be generated.

A “no match cos” command causes one or more CoS values (depending upon whether a single CoS value “x” was entered or a range of CoS values “x-y” was entered) attached to a particular class-map to be removed and automatically attached to the 4q1t-8e-4q4q-in-q-default. However, all CoS values cannot be removed from a particular class-map if there are still DSCP values associated with the particular class-map. Attempting this will generate an error. In other words, all CoS values cannot be removed from a given class-map until all DSCP values have been removed from the class-map. Finally, “no match” commands are not accepted on the 4q1t-8e-4q4q-in-q-default class-map and will generate an error.

Because APIC-EM has no knowledge of the existing mapping of the DSCP and CoS values to the class-maps before configuring the ingress queuing policy, it will first set all CoS and DSCP values to the 4q1t-8e-4q4q-in-q-default class-map before moving CoS and DSCP values to their desired class-maps. APIC-EM EasyQoS then moves the DSCP and CoS values into the desired class-maps. This is accomplished via the following commands.

```
!
class-map type queuing match-any 4q1t-8e-4q4q-in-q1
 match cos 5-7
 match DSCP 32, 40, 46, 48, 56
class-map type queuing match-any 4q1t-8e-4q4q-in-q3
 match cos 2-4
 match DSCP 16, 18, 20, 22, 24, 26, 30, 34, 36, 38



class-map type queuing match-any 4q1t-8e-4q4q-in-q4
 match cos 1
 match DSCP 8, 10, 12, 14
!
```

After the DSCP and CoS values have been moved into the appropriate class-maps, system-defined policy-map is cloned by appending “prm-DSCP#” to the beginning of the system-defined policy map. This is done via the following config-level command.

```
!
qos copy policy-map type queuing default-8e-4q4q-in-policy prefix prm-DSCP#
```

```
!
```

After the cloned policy-map is created, the following configuration modifies the cloned queueing policy map with the 4Q1T ingress queuing structure on the Nexus 7000 Series F2, F2e, and F3 modules.

```
!
policy-map type queuing prm-DSCP#8e-4q4q-in
  class type queuing 4q1t-8e-4q4q-in-q4
    queue-limit percent 30
    bandwidth percent 5
  class type queuing 4q1t-8e-4q4q-in-q1
    queue-limit percent 10
    bandwidth percent 30
  class type queuing 4q1t-8e-4q4q-in-q3
    queue-limit percent 30
    bandwidth percent 40
  class type queuing 4q1t-8e-4q4q-in-q-default
    queue-limit percent 30
    bandwidth percent 25
!
```

The policy-map with the 4Q1T ingress queuing structure is then applied to Ethernet interfaces that connect to either other core-layer switches or to distribution-layer switches. Additionally, the policy-map can be applied to the logical port-channel interface when an EtherChannel connection is used for port-level resilience, instead of a single physical interface. An example of the configuration of the policy to an Ethernet and a Port-channel interface is shown below.

```
!
interface Ethernet x/x
 service-policy type queuing input APIC_EM-8e-4q4q-in
!
interface port-channel xxx
 service-policy type queuing input APIC_EM-8e-4q4q-in
!
```

Nexus 7000 1P3Q1T Egress Queuing

The network-qos default-nq-8e-4q4q-policy template has a default egress queuing policy map that can be displayed via the exec-level “show policy-map type queuing default-8e-4q4q-out-policy” command. An example of the output is shown below.

```
DC-7010-2# show policy-map type queuing default-8e-4q4q-out-policy

  Type queuing policy-maps
  ========================

  policy-map type queuing default-8e-4q4q-out-policy
    class type queuing 1p3q1t-8e-4q4q-out-pq1
      priority level 1
    class type queuing 1p3q1t-8e-4q4q-out-q2
      bandwidth remaining percent 33
    class type queuing 1p3q1t-8e-4q4q-out-q3
      bandwidth remaining percent 33
    class type queuing 1p3q1t-8e-4q4q-out-q-default
      bandwidth remaining percent 33
```

For F2, F2e, and F3 series modules, the egress queuing policy defines the priority queue and the default bandwidth remaining bandwidth percentages for each of the rest of the class-maps (which correspond to the queues). These percentages will be modified by APIC-EM for the EasyQoS solution.

The network-qos default-nq-8e-4q4q-policy template also provides default settings for the system-defined class-maps that can be displayed via the exec-level “show class-map type queuing” command for each of the system-defined class-map names. An example of the output for each of the system-defined class-maps is shown below.

```
DC-7010-2# show class-map type queuing 1p3q1t-8e-4q4q-out-pq1

  Type queuing class-maps
  =======================

    class-map type queuing match-any 1p3q1t-8e-4q4q-out-pq1
      Description: Classifier for Egress Priority queue 1 of type 1p3q1t-8e-4q4q
      match cos 5-7

DC-7010-2# show class-map type queuing 1p3q1t-8e-4q4q-out-q2

  Type queuing class-maps
```

```
  ========================

    class-map type queuing match-any 1p3q1t-8e-4q4q-out-q2
      Description: Classifier for Egress queue 2 of type 1p3q1t-8e-4q4q
      match cos 3-4

DC-7010-2# show class-map type queuing 1p3q1t-8e-4q4q-out-q3

  Type queuing class-maps
  ========================

    class-map type queuing match-any 1p3q1t-8e-4q4q-out-q3
      Description: Classifier for Egress queue 3 of type 1p3q1t-8e-4q4q
      match cos 2

DC-7010-2# show class-map type queuing 1p3q1t-8e-4q4q-out-q-default

  Type queuing class-maps
  ========================

    class-map type queuing match-any 1p3q1t-8e-4q4q-out-q-default
      Description: Classifier for Egress queue 4 of type 1p3q1t-8e-4q4q
      match cos 0-1
```

The system-defined class-maps have default settings for mapping CoS values to each of the class-maps. However, APIC-EM EasyQoS cannot guarantee that the customer has not already previously configured the default-nq-8e-4q4q-policy network-qos template and modified the default mappings of CoS and DSCP values to class-maps. Hence, the queuing models for CoS-to-queue mapping need to be configured.

For egress queuing, only CoS-to-Queue mapping is supported on F2, F2e, and F3 Series modules. The following figure shows the Cos-to-queue mapping for the 1P3Q1T egress queueing model.

Figure 85 Egress Queuing Model (1P3Q1T)—CoS-to-Queue Mapping

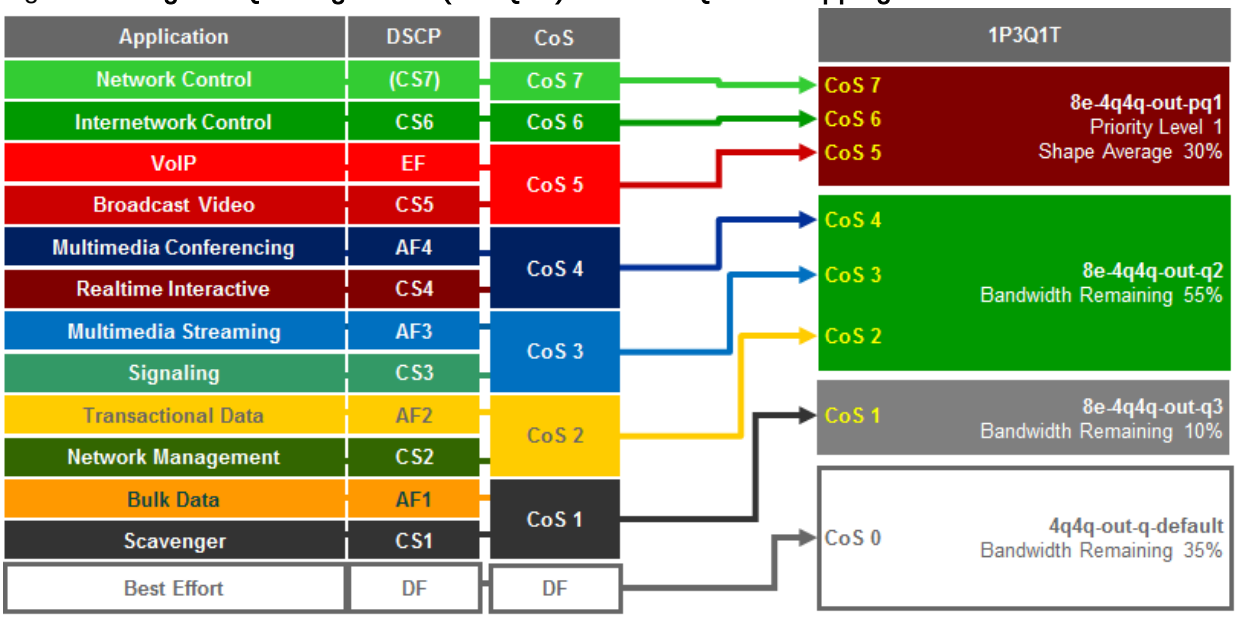

NX OS provides system-defined class-map names that cannot be renamed. The system-defined class-map names that match the four output queues of the F2, F2e, and F3 Series modules on the Nexus 7000 are as follows:

- 1p3q1t-out-pq1
- 1p3q1t-out-q3
- 1p3q1t-out-q4
- 1p3q1t-out-q-default

These class-maps have default and/or non-default (customer implemented) CoS values attached to them. These values can be reset with “no match cos” commands. A “no match cos” command causes one or more CoS values (depending upon whether a single CoS value “x” was entered or a range of CoS values “x-y” was entered) attached to a particular class-map to be removed and automatically attached to the 1p3q1t-out-q-default class-map.

Because APIC-EM has no knowledge of the existing mapping of the CoS values to the class-maps before configuring the ingress queuing policy, it will first set all CoS values to the 1p3q1t-out-q-default class-map before moving CoS values to their desired class-maps. This is accomplished via the following commands.

```
!
class-map type queuing match-any 1p3q1t-8e-4q4q-out-pq1
 no match cos 0-7
class-map type queuing match-any 1p3q1t-8e-4q4q-out-q2
 no match cos 0-7
class-map type queuing match-any 1p3q1t-8e-4q4q-out-q3
```

```
 no match cos 0-7
!
```

Note: Because there are no DSCP values within egress class-maps, there are no issues with moving CoS values to the default class-map, as was discussed in the ***Nexus 7000 8Q2T Ingress Queuing*** section.

APIC-EM can then move CoS values into the desired class-maps. This is accomplished via the following commands.

```
!
class-map type queuing match-any 1p3q1t-8e-4q4q-out-pq1
 match cos 5-7
class-map type queuing match-any 1p3q1t-8e-4q4q-out-q2
 match cos 2-4
class-map type queuing match-any 1p3q1t-8e-4q4q-out-q3
 match cos 1
!
```

After the CoS values have been moved into the appropriate class-maps, the system-defined policy-map is cloned by appending "prm-DSCP#" to the beginning of the system-defined policy map. This is done via the following config-level command.

```
!
qos copy policy-map type queuing default-8e-4q4q-out-policy prefix prm-DSCP#
!
```

After the cloned policy-map is created, the following configuration modifies the cloned queueing policy-map with the 1P3Q1T egress queuing structure on the Nexus 7000 Series F2, F2e, and F3 modules.

```
!
policy-map type queuing APIC_EM-8e-4q4q-out
  class type queuing 1p3q1t-8e-4q4q-out-pq1
    priority level 1
    shape average percent 30
  class type queuing 1p3q1t-8e-4q4q-out-q3
    bandwidth remaining percent 10
  class type queuing 1p3q1t-8e-4q4q-out-q2
    bandwidth remaining percent 55
  class type queuing 1p3q1t-8e-4q4q-out-q-default
```

```
    bandwidth remaining percent 35
!
```

The policy-map with the 1P3Q1T egress queuing structure is then applied to Ethernet interfaces that connect to either other core-layer switches or to distribution-layer switches. Additionally, the policy-map can be applied to the logical port-channel interface when an EtherChannel connection is used for port-level resilience, instead of a single physical interface. An example of the configuration of the policy to an Ethernet and a Port-channel interface is shown below.

```
!
interface Ethernet x/x
 service-policy type queuing output prm-DSCP#8e-4q4q-out
!
interface port-channel xxx
 service-policy type queuing output prm-DSCP#8e-4q4q-out
!
```

F2e and F3 Series Modules on the Nexus 7700

F2e and F3 Series modules for the Nexus 7700 consist of the following:

- N77-F248XP-23E
- N77-F348XP-23
- N77-F324FQ-25
- N77-F312CK-26

These modules are only supported on the Nexus 7700 Series chassis.

The ingress and egress queuing structure for F2e and F3 Series modules on the Nexus 7700 is determined by the system network QoS policy. When an F2e or F3 Series module is inserted and becomes operational within the Nexus 7700 chassis, all nine network-qos policies (templates) shown in Figure 78 are automatically added to the system.

Because the default-nq-8e-4q8q-policy template better uses the available ingress and egress queuing structure (4 ingress queues and 8 egress queues), this is the network-qos policy template that APIC-EM EasyQoS configures via the following global configuration commands.

```
!
system qos
 service-policy type network-qos default-nq-8e-4q8q-policy
!
```

The default-nq-8e-4q4q-policy implements a 4Q1T ingress queueing structure and a1P7Q1T egress queuing structure for F2, F2e, and F3 Series modules with the Nexus 7700.

Nexus 7000 4Q1T Ingress Queuing

DSCP-to-Queue mapping is by default enabled on the Nexus 7700 Series. However, in the event the customer has disabled this for some reason, APIC-EM EasyQoS will enable it via the following global configuration command.

```
!
hardware qos DSCP-to-queue ingress
!
```

Trust of both CoS and DSCP values is implicit on the Nexus 7700 Series. No explicit configuration is required to enable trust.

The network-qos default-nq-8e-4q8q-policy template has a default ingress queuing policy map that can be displayed via the exec-level “show policy-map type queuing default-8e-4q8q-in-policy” command. An example of the output is shown below.

```
N7700# show policy-map type queuing default-8e-4q8q-in-policy


  Type queuing policy-maps
  ========================


  policy-map type queuing default-8e-4q8q-in-policy
    class type queuing 8e-4q8q-in-q1
      queue-limit percent 10
      bandwidth percent 49
    class type queuing 8e-4q8q-in-q-default
      queue-limit percent 88
      bandwidth percent 49
    class type queuing 8e-4q8q-in-q3
      queue-limit percent 1
      bandwidth percent 1
    class type queuing 8e-4q8q-in-q4
      queue-limit percent 1
      bandwidth percent 1
```

For F2, F2e, and F3 series modules, the ingress queuing policy define queue-limit percentages and bandwidth percentages for each of the class-maps (which correspond to the queues). These percentages will be modified by APIC-EM for the EasyQoS solution.

The network-qos default-nq-8e-4q8q-policy template also provides default settings for the system-defined class-maps that can be displayed via the exec-level “show class-map type queuing” command for each of the system-defined class-map names. An example of the output for each of the system-defined class-maps is shown below.

```
N7700# show class-map type queuing 8e-4q8q-in-q1

  Type queuing class-maps
  ========================

    class-map type queuing match-any 8e-4q8q-in-q1
      Description: Classifier for Ingress queue 1 of type 4q1t8e
      match cos 5-7
      match DSCP 40-63

N7700# show class-map type queuing 8e-4q8q-in-q-default

  Type queuing class-maps
  ========================

    class-map type queuing match-any 8e-4q8q-in-q-default
      Description: Classifier for Ingress default queue of type 4q1t8e
      match cos 0-4
      match DSCP 0-39
```

The system-defined class-maps have default settings for mapping CoS and DSCP values to each of the class-maps. Note that no CoS or DSCP values are by default mapped to class-maps 8e-4q8q-in-q-2 and 8e-4q8q-in-q3. However, APIC-EM cannot guarantee that the customer has not already previously configured the default-nq-8e-4q8q-policy network-qos template and modified the default mappings of CoS and DSCP values to class-maps. Because both CoS and DSCP values are trusted, queuing models for both CoS-to-queue mapping and DSCP-to-queue mapping need to be configured.

The CoS-to-queue mapping for the 4Q1T ingress queueing model deployed by APIC-EM for the Nexus 7700 F2e and F3 modules is the same as was shown in Figure 81 above. Likewise the DSCP-to-queue mapping for the 4Q1T ingress queuing model deployed by APIC-EM for the Nexus 7700 F2e and F3 modules is the same as was shown in Figure 82 above.

NX OS provides system-defined class-map names that cannot be renamed. The system-defined class-map names that match the four ingress queues of the F2e, and F3 Series modules for the Nexus 7700 are as follows:

- 4q1t-8e-4q8q-in-q1
- 4q1t-8e-4q8q-in-q3
- 4q1t-8e-4q8q-in-q4
- 4q1t-8e-4q8q-in-q-default

These class-maps have default and/or non-default (customer implemented) CoS and/or DSCP values attached to them. These values can be reset with “no match” commands.

A “no match DSCP” command causes one or more DSCP values (depending upon whether a single DSCP value “x” was entered or a range of DSCP values “x-y” was entered) attached to a particular class-map to be removed and automatically attached to the 4q1t-8e-4q8q-in-q-default class-map. However, a CoS value must first be mapped to the 4q1t-8e-4q8q-in-q-default class-map, or an error will be generated.

A “no match cos” command causes one or more CoS values (depending upon whether a single CoS value “x” was entered or a range of CoS values “x-y” was entered) attached to a particular class-map to be removed and automatically attached to the 4q1t-8e-4q8q-in-q-default. However, all CoS values cannot be removed from a particular class-map if there are still DSCP values associated with the particular class-map. Attempting this will generate an error. In other words, all CoS values cannot be removed from a given class-map until all DSCP values have been removed from the class-map. Finally, “no match” commands are not accepted on the 4q1t-8e-4q8q-in-q-default class-map, and will generate an error.

Because APIC-EM EasyQoS has no knowledge of the existing mapping of the DSCP and CoS values to the class-maps before configuring the ingress queuing policy, it will first set all CoS and DSCP values to the 4q1t-8e-4q8q-in-q-default class-map before moving CoS and DSCP values to their desired class-maps. EasyQoS then moves the DSCP and CoS values into the desired class-maps. This is accomplished via the following commands.

```
!
class-map type queuing match-any 8e-4q8q-in-q1
  match cos 5-7
  match DSCP 32, 40, 46
  match DSCP 48, 56
class-map type queuing match-any 8e-4q8q-in-q3
  match cos 2-4
  match DSCP 16, 18, 20, 22
  match DSCP 24, 26, 28, 30
  match DSCP 34, 36, 38
class-map type queuing match-any 8e-4q8q-in-q4
  match cos 1
  match DSCP 8, 10, 12, 14
!
```

After the DSCP and CoS values have been moved into the appropriate class-maps, the following configuration creates and configures the queueing policy-map with the 4Q1T ingress queuing structure on the Nexus 7700 Series F2e and F3 modules.

```
!
policy-map type queuing prm-DSCP#4q1t-in
  class type queuing 8e-4q8q-in-q1
    bandwidth percent 30
    queue-limit percent 10
 class type queuing 8e-4q8q-in-q-default
    bandwidth percent 25
    queue-limit percent 30
  class type queuing 8e-4q8q-in-q3
    bandwidth percent 40
    queue-limit percent 30
  class type queuing 8e-4q8q-in-q4
    bandwidth percent 5
    queue-limit percent 30
!
```

The policy-map with the 4Q1T ingress queuing structure is then applied to Ethernet interfaces that connect to either other core-layer switches or to distribution-layer switches. Additionally, the policy-map can be applied to the logical port-channel interface when an EtherChannel connection is used for port-level resilience, instead of a single physical interface. An example of the configuration of the policy to an Ethernet and a Port-channel interface is shown below.

```
!
interface Ethernet x/x
 service-policy type queuing input prm-DSCP#4q1t-in
!
interface port-channel xxx
 service-policy type queuing input prm-DSCP#4q1t-in
!
```

Nexus 7700 1P7Q1T Egress Queuing

The network-qos default-nq-8e-4q8q-policy template has a default egress queuing policy map that can be displayed via the exec-level “show policy-map type queuing default-8e-4q8q-out-policy” command. An example of the output is shown below.

```
N7700(config)# show policy-map type queuing default-8e-4q8q-out-policy


  Type queuing policy-maps
  ========================
  policy-map type queuing default-8e-4q8q-out-policy
    class type queuing 8e-4q8q-out-q1
      priority level 1
    class type queuing 8e-4q8q-out-q2
      bandwidth remaining percent 14
    class type queuing 8e-4q8q-out-q3
      bandwidth remaining percent 14
    class type queuing 8e-4q8q-out-q4
      bandwidth remaining percent 14
    class type queuing 8e-4q8q-out-q5
      bandwidth remaining percent 14
    class type queuing 8e-4q8q-out-q6
      bandwidth remaining percent 14
    class type queuing 8e-4q8q-out-q7
      bandwidth remaining percent 14
    class type queuing 8e-4q8q-out-q-default
      bandwidth remaining percent 14
```

For F2e and F3 series modules, the egress queuing policy defines the priority queue, and the default bandwidth remaining bandwidth percentages for each of the rest of the class-maps (which correspond to the queues). These percentages will be modified by APIC-EM for the EasyQoS solution.

The network-qos default-nq-8e-4q8q-policy template also provides default settings for the system-defined class-maps that can be displayed via the exec-level “show class-map type queuing” command for each of the system-defined class-map names. An example of the output for each of the system-defined class-maps is shown below.

```
Xbow1# show class-map type queuing 8e-4q8q-out-q1


  Type queuing class-maps
  =======================
```

```
    class-map type queuing match-any 8e-4q8q-out-q1
      Description: Classifier for Egress priority queue of type 1p7q1t8e
      match cos 5

Xbow1# show class-map type queuing 8e-4q8q-out-q2

  Type queuing class-maps
  ========================

    class-map type queuing match-any 8e-4q8q-out-q2
      Description: Classifier for Egress queue 2 of type 1p7q1t8e
      match cos 7

Xbow1# show class-map type queuing 8e-4q8q-out-q3

  Type queuing class-maps
  ========================

    class-map type queuing match-any 8e-4q8q-out-q3
      Description: Classifier for Egress queue 3 of type 1p7q1t8e
      match cos 6

 Xbow1# show class-map type queuing 8e-4q8q-out-q4

  Type queuing class-maps
  ========================

    class-map type queuing match-any 8e-4q8q-out-q4
      Description: Classifier for Egress queue 4 of type 1p7q1t8e
      match cos 4

Xbow1# show class-map type queuing 8e-4q8q-out-q5
```

```
  Type queuing class-maps
  =======================

    class-map type queuing match-any 8e-4q8q-out-q5
      Description: Classifier for Egress queue 5 of type 1p7q1t8e
      match cos 3

Xbow1# show class-map type queuing 8e-4q8q-out-q6

  Type queuing class-maps
  =======================

    class-map type queuing match-any 8e-4q8q-out-q6
      Description: Classifier for Egress queue 6 of type 1p7q1t8e
      match cos 2

Xbow1# show class-map type queuing 8e-4q8q-out-q7

  Type queuing class-maps
  =======================

    class-map type queuing match-any 8e-4q8q-out-q7
      Description: Classifier for Egress queue 7 of type 1p7q1t8e
      match cos 1

Xbow1# show class-map type queuing 8e-4q8q-out-q-default

  Type queuing class-maps
  =======================

    class-map type queuing match-any 8e-4q8q-out-q-default
```

```
Description: Classifier for Egress default queue of type 1p7q1t8e

match cos 0
```

The system-defined class-maps have default settings for mapping CoS values to each of the class-maps. However, APIC-EM EasyQoS cannot guarantee that the customer has not already previously configured the default-nq-8e-4q8q-policy network-qos template and modified the default mappings of CoS and DSCP values to class-maps. Hence, the queuing models for CoS-to-queue mapping need to be configured.

For egress queuing, only CoS-to-Queue mapping is supported on F2e and F3 Series modules. The following figure shows the Cos-to-queue mapping for the 1P7Q1T egress queueing model.

Figure 86 **Egress Queuing Model (1P7Q1T)—CoS-to-Queue Mapping**

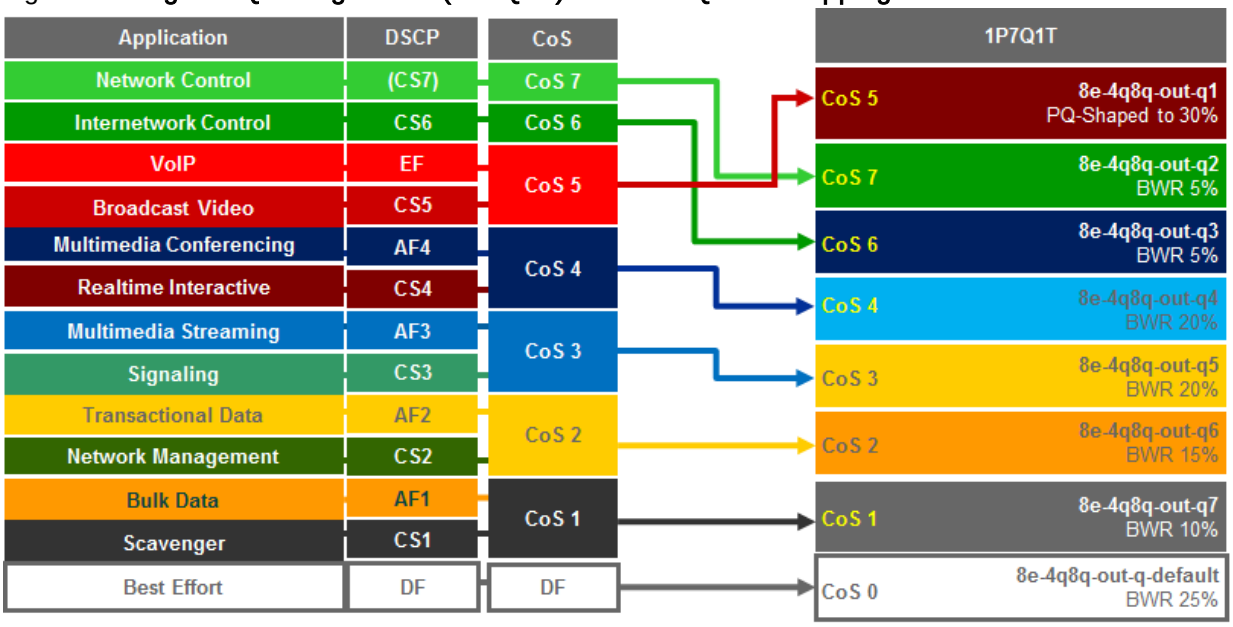

NX OS provides system-defined class-map names that cannot be renamed. The system-defined class-map names that match the eight output queues of the F2e and F3 Series modules on the Nexus 7700 are as follows:

- 1p7q1t-out-q1
- 1p7q1t-out-q2
- 1p7q1t-out-q3
- 1p7q1t-out-q4
- 1p7q1t-out-q5
- 1p7q1t-out-q6
- 1p7q1t-out-q7
- 1p7q1t-out-q-default

These class-maps have default and/or non-default (customer implemented) CoS values attached to them. These values can be reset with “no match cos” commands. A “no match cos” command causes one or more CoS values (depending upon whether a single CoS value “x” was entered or a range of CoS values “x-y” was entered) attached to a particular class-map to be removed and automatically attached to the 1p7q1t-out-q-default class-map.

Because APIC-EM EasyQoS has no knowledge of the existing mapping of the CoS values to the class-maps before configuring the ingress queuing policy, it will first set all CoS values to the 1p7q1t-out-q-default class-map before moving CoS and DSCP values to their desired class-maps. This is accomplished via the following commands.

```
!
class-map type queuing match-any 8e-4q8q-out-q1
 no match cos 0-7
class-map type queuing match-any 8e-4q8q-out-q2
 no match cos 0-7
class-map type queuing match-any 8e-4q8q-out-q3
 no match cos 0-7
class-map type queuing match-any 8e-4q8q-out-q4
 no match cos 0-7
class-map type queuing match-any 8e-4q8q-out-q5
 no match cos 0-7
class-map type queuing match-any 8e-4q8q-out-q6
 no match cos 0-7
class-map type queuing match-any 8e-4q8q-out-q7
 no match cos 0-7
!
```

Note: Because there are no DSCP values within egress class-maps, there are no issues with moving CoS values to the default class-map, as was discussed in the ***Nexus 7000 8Q2T Ingress Queuing*** section.

APIC-EM EasyQoS can then move CoS values into the desired class-maps. This is accomplished via the following commands.

```
!
class-map type queuing match-any 8e-4q8q-out-q1
 match cos 5
class-map type queuing match-any 8e-4q8q-out-q2
 match cos 7
```

```
class-map type queuing match-any 8e-4q8q-out-q3
 match cos 6
class-map type queuing match-any 8e-4q8q-out-q4
 match cos 4
class-map type queuing match-any 8e-4q8q-out-q5
 match cos 3
class-map type queuing match-any 8e-4q8q-out-q6
 match cos 2
class-map type queuing match-any 8e-4q8q-out-q7
 match cos 1
!
```

After the CoS values have been moved into the appropriate class-maps, the following configuration creates and configures the queueing policy-map with the 1P7Q1T egress queuing structure on the Nexus 7700 Series F2e and F3 modules.

```
!
policy-map type queuing prm-DSCP#1p7q1t-out
  class type queuing 8e-4q8q-out-q1
    priority level 1
    shape average percent 30
  class type queuing 8e-4q8q-out-q2
    bandwidth remaining percent 5
  class type queuing 8e-4q8q-out-q3
    bandwidth remaining percent 5
  class type queuing 8e-4q8q-out-q4
    bandwidth remaining percent 20
  class type queuing 8e-4q8q-out-q5
    bandwidth remaining percent 20
  class type queuing 8e-4q8q-out-q6
    bandwidth remaining percent 15
  class type queuing 8e-4q8q-out-q7
    bandwidth remaining percent 10
  class type queuing 8e-4q8q-out-q-default
    bandwidth remaining percent 25
```

```
!
```

The policy-map with the 1P7Q1T egress queuing structure is then applied to Ethernet interfaces that connect to either other core-layer switches or to distribution-layer switches. Additionally, the policy-map can be applied to the logical port-channel interface when an EtherChannel connection is used for port-level resilience, instead of a single physical interface. An example of the output for each of the system-defined class-maps is shown below.

```
!
interface Ethernet x/x
 service-policy type queuing output prm-DSCP#1p7q1t-out
!
interface port-channel xxx
 service-policy type queuing output prm-DSCP#1p7q1t-out
!
```

M3 Series Modules on the Nexus 7700

M3 Series modules for the Nexus 7700 consist of the following:

- N77-M348XP-23L
- N77-M324FQ-25L

These modules are only supported on the Nexus 7700 Series chassis.

The M3 Series modules implement a 4Q1T ingress queuing structure and a 1P7Q1T egress queuing structure. This queueing structure is the same as implemented by F2e and F3 modules on the Nexus 7700 Series, when using the default-nq-8e-4q4q-policy network QoS policy, as discussed in the *F2e and F3 Series Modules on the Nexus 7700* section of this document. The only difference between the QoS policy pushed by APIC-EM for the M3 modules and the F2e and F3 modules is that network-qos policy template does not need to be configured for the M3 modules.

Pre-Existing QoS Configurations on Switch Platforms

This section discusses how EasyQoS handles prior QoS configurations on switch platforms, when deploying an EasyQoS policy. For ingress classification & marking policies, EasyQoS will remove any existing service-policy definition from the interface and replace it with its service-policy definitions. The previous class-map and policy-map definitions will not be deleted by EasyQoS. This is necessary for restoring the original pre-EasyQoS (before any EasyQoS configuration was applied) configuration back to the switch platform. The Restore feature is a new feature added to APIC-EM EasyQoS release 1.3.

For queuing policies, the behavior depends on whether the platform is an MQC platform (Catalyst 3850, 3650, or 4500 Series), a C3PL platform (Catalyst 6K Series with Sup2T), an NX OS platform (Nexus 7000 or 7700 Series), or an older MLS QoS platform (Catalyst 3750, 3560, or 2960 Series, and older Catalyst 6K Series with Sup720).

- For MQC and C3PL platforms, queuing policies are applied via service-policy statements similar to ingress classification & marking policies. The behavior is the same as with ingress classification & marking policies. The previous class-map and policy-map definitions will not be deleted by EasyQoS. Clicking on the Restore button within an EasyQoS policy will cause the pre-EasyQoS queuing service-policy statements to be re-applied to the interfaces.
- For MLS QoS platforms, the queuing policy is configured directly on the interface. EasyQoS may change the policy, so there is no previous configuration saved on the switch platform. Therefore, clicking on the Restore button will not restore the pre-EasyQoS queuing policy for these platforms, although the ingress classification & marking policy will be restored, because it uses service-policy definitions applied to the interfaces.
- For NX OS platforms, the class-map definitions are system-defined, and not user-defined. EasyQoS may modify the mapping of DSCP and/or CoS values to the ingress and/or egress queues. This will not be restored to their pre-EasyQoS configuration. However, policy-map definitions are user-defined (or extended from the default template). Existing policy-map definitions are not deleted by EasyQoS. Therefore, clicking on the Restore button within an EasyQoS policy will cause the pre-Existing queuing service-policy statements to be re-applied to the interfaces.

EasyQoS does not currently remove Auto QoS statements. Depending on the platform and what form of Auto QoS is implemented, this can cause EasyQoS policy to not function properly. Therefore, the network operator should either completely remove Auto QoS configurations before applying any EasyQoS policy or not implement EasyQoS policy when Auto QoS is configured on the platform. Future versions of APIC-EM EasyQoS may remove Auto QoS configuration as well.

WLAN QoS Design

AireOS WLC QoS Design

This section discusses AireOS wireless LAN controller platforms within the EasyQoS solution. For the APIC-EM 1.3 release, only Cisco WLCs running the AireOS operating system are supported within the EasyQoS application. Dedicated IOS XE WLCs such as the Cisco 5760 are not currently supported. WLC functionality within Catalyst 3850, 3650, and Catalyst 4500-E Series switches are also not currently supported. Further, only centralized (local) mode configurations, in which traffic is tunneled from the Access Point to the WLC before being placed onto the LAN, are supported.

Disabling WLANs and Radios

In order to provision the global QoS configuration (meaning the QoS configuration which affects the entire WLC) required for EasyQoS Static QoS discussed in the sections below, all SSIDs and WLANs must be first disabled. In addition, the 802.11 b/g/n and the 802.11a/n/ac radios in all APs controlled by the WLC must also be disabled. APIC-EM uses an SSH session established to AireOS WLCs in order to provision QoS policy, instead of the web-based GUI.

The following global commands provisioned by APIC-EM EasyQoS disable all WLANS and radios on an AireOS WLC.

```
config 802.11a disable network

config 802.11b disable network

config wlan disable all
```

This means that the initial provisioning of EasyQoS Static QoS policy is a disruptive process to the WLCs that are part of a policy. Wireless connectivity itself will be disrupted for wireless clients connected to any WLAN/SSID on any Access Point serviced by a WLC that contains a WLAN/SSID to which EasyQoS policy is to be applied. To emphasize this again, the disruption will be to network connectivity and not just QoS marking. Hence, is it recommended that the initial provisioning of EasyQoS Static QoS policies deployed to WLANs should be scheduled during normal network change-control windows. The disruption will be to all WLANs/SSIDs serviced by the AireOS WLC.

Specifically the changes provisioned by EasyQoS that require the 802.11 b/g/n and 802.11a/n/ac and the WLAN/SSIDs to be disabled are as follows:

- Changes to the EDCA parameters
- Changes to Call Admission Control (CAC) settings
- Changes to the Global QoS Profile settings (specifically the Platinum QoS Profile)
- Changes to the QoS Map settings

Changes to the specific WLAN/SSID to which the EasyQoS policy is being applied require only the WLAN/SSID to be disabled. These are as follows:

- Applying the QoS Profile to the WLAN/SSID
- Enabling AVC on the WLAN/SSID
- Applying the AVC Profile to the WLAN/SSID

Therefore, after the initial provisioning of EasyQoS policy to any WLAN/SSID on the WLC, further modifications to an EasyQoS policy that simply involve changes to the applications within the policy may only be disruptive to the specific WLAN/SSID to which the policy is applied.

QoS Trust Boundaries and Policy Enforcement Points

QoS trust boundaries and policy enforcement points are more complicated with IEEE 802.11 wireless infrastructure due to the fact that over-the-air QoS is based on Layer 2 headers, not Layer 3 headers. IEEE 802.11 QoS consists of eight User Priorities (UPs) that are mapped to four Access Categories (ACs)–Voice, Video, Best Effort, and Background. Layer 3 DSCP values must be mapped to and from the eight Layer 2 IEEE 802.11 UPs, which are then mapped to the four ACs. Hence, multiple QoS trust boundaries can exist–depending upon whether the trust boundary is based on Layer 2 UP or Layer 3 DSCP marking. Additionally, there can be multiple policy enforcement points:

- Policy enforcement points for mapping Layer 3 DSCP value to Layer 2 UP values to-and-from wireless clients
- Policy enforcement points for re-mapping Layer 3 DSCP values to other Layer 3 DSCP values based on AVC profiles

An example of the various wireless trust boundaries and policy enforcement points is shown in the following figure.

Figure 87 **Wireless Trust Boundaries and Policy Enforcement Points**

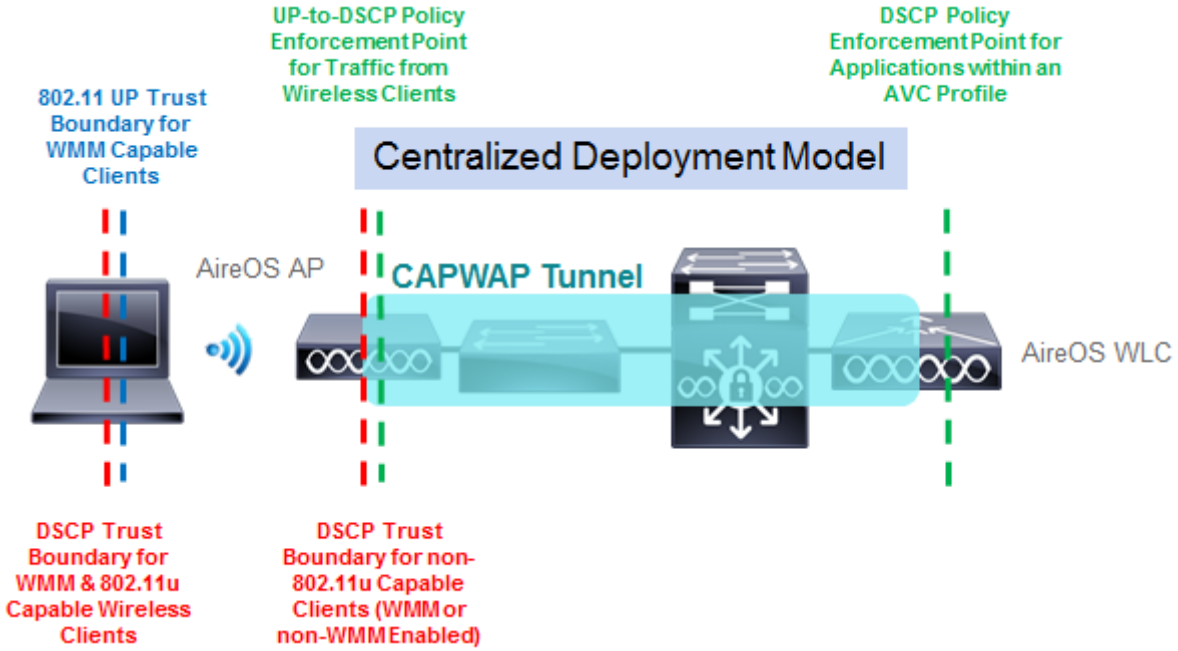

Explanation of the various trust boundaries and policy enforcement points can best be explained based on type of wireless client as follows:

- Non- WiFi Multimedia (WMM) capable wireless clients
- WMM capable wireless clients (non-802.11u capable)
- 802.11u capable wireless clients

Non-WMM Capable Wireless Clients

Non-WMM capable wireless clients are less common in deployments today. Such devices do not send wireless frames with the IEEE 802.11 QoS field. Hence there is no Layer 2 IEEE 802.11 QoS UP information included within frames sent by these clients, and therefore no Layer 2 QoS trust boundary. All traffic sent by these clients is by default placed into the Best Effort AC when it is sent over-the-air from the wireless client to the Access Point (AP).

From the AP to the WLC, IEEE 802.11 frames are encapsulated within an IP packet with a CAPWAP header and sent upstream. CAPWAP effectively tunnels IEEE 802.11 frames through an IP network between the AP and the WLC. Within the IEEE 802.11 frame itself, there is also an encapsulated IP packet from the wireless client. The ToS field of that inner IP packet can have any DSCP marking–assuming the application, the operating system, and the potentially the wireless drivers running on the wireless device allow such DSCP markings.

For the EasyQoS solution, the Platinum QoS Profile within the WLC is applied to the WLAN/SSID by EasyQoS. The global Platinum QoS Profile is modified by EasyQos to allow upstream and downstream traffic sent over the wireless medium to use up to the Voice AC. However, the global Platinum QoS Profile is also modified to set the unicast default priority and multicast default priority to the best effort Access Category. This means that any unicast and multicast traffic without an 802.11 QoS field are set to best effort. This is accomplished via the two "besteffort" parameters in the following global command provisioned by EasyQoS to AireOS WLCs:

```
config qos priority platinum voice besteffort besteffort
```

The net effect is that if the operating system and wireless drivers of the non-WMM wireless device allow an application running on the wireless device to mark IP packets with a DSCP value, the DSCP marking of the CAPWAP header will still be set to the default DSCP value (DSCP =0) as the traffic is sent from the AP to the WLC. If the operating system and wireless drivers of the wireless device do not allow an application to mark IP packets with a DSCP value, or if the application itself simply sends all traffic with a default DSCP value, the DSCP marking of the CAPWAP header also be set to the default DSCP value. In other words, all traffic from non-WMM clients will receive best effort (Default) treatment from the AP to the WLC.

With this configuration, the AP serves as a trust boundary and policy enforcement point for non-WMM devices. Regardless of the DSCP values of the IP packets sent by the wireless client, the DSCP value of the outer CAPWAP header will be set to the default DSCP value (DSCP = 0), when packets are sent from the AP to the WLC. Note, however, that the original DSCP markings of the IP packets within the CAPWAP tunnel are still preserved up to the WLC. The AVC policy applied at the WLC may, however, remark this traffic if the application is part of the AVC policy.

For downstream traffic, the DSCP marking of the outer CAPWAP header will match the DSCP marking of inner IP packet, also encapsulated within an IEEE 802.11 frame. This is because traffic up to the voice

Access Category is allowed with the Platinum QoS Profile as shown in the configuration example above. This preserves the Layer 3 DSCP QoS marking from the WLC to the AP. For non-WMM wireless clients, 802.11 frames are sent over-the-air with no UP value, because the 802.11 QoS field is not supported by non-WMM clients. However, the traffic is still scheduled into the appropriate AC queues, based on the mapping of the DSCP value of the outer CAPWAP header to the UP. So, effectively, QoS is preserved downstream.

WMM Capable Wireless Clients (non-802.11u Capable)

WMM capable wireless clients send wireless frames with the IEEE 802.11 QoS field. Hence there is Layer 2 UP information included within frames sent by these clients, and therefore the Layer 2 trust boundary is at the wireless client, as shown in Figure 87 above. However, the operating system and wireless drivers of the wireless device must allow application traffic to be marked with a Layer 2 UP. Further, the application must also be able to send traffic with DSCP markings. The mapping of those DSCP markings to the appropriate IEEE 802.11 UPs is then largely determined by the vendor of the wireless device if the device does not support IEEE 802.11u.

Traffic sent by these clients is placed into an IEEE 802.11 AC as it is sent over-the-air from the wireless client to the AP, based on the UP as shown in the figure below.

Figure 88 802.11 UP Values and ACs

802.11 UP Value	802.11 Access Category	WMM Designation	Cisco WLC Designation
7	AC_VO	Voice	Platinum
6			
5	AC_VI	Video	Gold
4			
3	AC_BE	Best Effort	Silver
0			
2	AC_BK	Background	Bronze
1			

The following global command configured on AireOS WLCs by APIC-EM EasyQoS will cause the AP to mark the ToS field of the outer CAPWAP header to match the DSCP marking of the encapsulated IP packet within the IEEE 802.11 frame sent by the wireless client.

```
config qos qosmap trust-DSCP-upstream enable
```

For the EasyQoS solution, the Platinum QoS profile within the WLC is configured to allow upstream and downstream traffic sent over the wireless medium to use up to the Voice AC. This is accomplished via the following global command on AireOS WLCs.

```
config qos priority platinum voice besteffort besteffort
```

The “voice” parameter within the command allows the AP to send traffic up to and including the IEEE 802.11 AC of Voice and UP values up to and including 7, to wireless clients that support WMM. It also honors traffic sent from wireless clients that support WMM to the AP that includes the IEEE 802.11 AC of Voice and UP value up to and including 7.

The net effect is that if the operating system and wireless drivers of the wireless device allow an application running on the wireless device to mark IP packets with a DSCP value, the DSCP marking of the CAPWAP header will match this DSCP value as the traffic is sent from the AP to the WLC. If the operating system and wireless drivers of the wireless device do not allow an application to mark IP packets with a DSCP value, or if the application simply sends all traffic with a default DSCP value (DSCP = 0), the DSCP marking of the CAPWAP header also be set to the default DSCP value.

Hence, with this configuration, the wireless client is trusted to send traffic marked with the correct DSCP value. In other words, the DSCP trust boundary is extended to the wireless client. The AP serves as a policy enforcement point, mapping the trusted DSCP values of the IP packets sent by the wireless client to the DSCP values of the CAPWAP headers when packets are sent from the AP to the WLC. Note that the original DSCP markings of the IP packets within the CAPWAP tunnel are also preserved up to the WLC. The AVC policy applied at the WLC may, however, remark this traffic if the application is part of the AVC policy.

For downstream traffic, the DSCP marking of the outer CAPWAP header will match the DSCP marking of the inner IP packet, also encapsulated within an IEEE 802.11 frame. For WMM-enabled wireless clients, IEEE 802.11 frames are sent over-the-air with UP values, because the IEEE 802.11 QoS field is supported by WMM clients.

Additional modifications to the global Platinum QoS profile provisioned by EasyQoS are as follows:

- 802.11p marking is disabled (all wired marking is DSCP-based) via the following command provisioned by EasyQoS.

```
config qos protocol-type platinum none
```

- Downstream UDP traffic is set to be unrestricted by EasyQoS through the following commands.

```
config qos burst-realtime-rate platinum per-ssid downstream 0
config qos average-realtime-rate platinum per-ssid downstream 0
```

The Platinum QoS Profile is applied to every SSID/WLAN controlled by EasyQoS. This is accomplished via the following SSID/WLAN-level command on AireOS WLCs:

```
config wlan qos x platinum
```

Note that “x” refers to the ID of the particular SSID/WLAN of the AireOS WLC.

For non-802.11u WMM-enabled clients, the AP will set the IEEE 802.11 UP value based on the QoS Map configuration within the WLC. As of AireOS software version 8.1.111.0, the QoS Map is now configurable, and applies globally to the entire WLC.

The following is the QoS Map Configuration (along with exceptions) that is provisioned by EasyQoS to AireOS WLCs for EasyQoS policies.

```
config qos qosmap disable
config qos qosmap default
config qos qosmap up-to-DSCP-map 0 0 0 7
config qos qosmap up-to-DSCP-map 1 8 8 15
config qos qosmap up-to-DSCP-map 2 16 16 23
```

```
config qos qosmap up-to-DSCP-map 3 24 24 31
config qos qosmap up-to-DSCP-map 4 32 32 39
config qos qosmap up-to-DSCP-map 5 34 40 47
config qos qosmap up-to-DSCP-map 6 46 48 62
config qos qosmap up-to-DSCP-map 7 56 56 63

config qos qosmap clear-all
config qos qosmap DSCP-to-up-exception 16 0
config qos qosmap DSCP-to-up-exception 8 1
config qos qosmap DSCP-to-up-exception 10 2
config qos qosmap DSCP-to-up-exception 12 2
config qos qosmap DSCP-to-up-exception 14 2
config qos qosmap DSCP-to-up-exception 18 2
config qos qosmap DSCP-to-up-exception 20 3
config qos qosmap DSCP-to-up-exception 22 3
config qos qosmap DSCP-to-up-exception 38 4
config qos qosmap DSCP-to-up-exception 36 4
config qos qosmap DSCP-to-up-exception 34 4
config qos qosmap DSCP-to-up-exception 30 4
config qos qosmap DSCP-to-up-exception 28 4
config qos qosmap DSCP-to-up-exception 26 4
config qos qosmap DSCP-to-up-exception 24 4
config qos qosmap DSCP-to-up-exception 40 5
config qos qosmap DSCP-to-up-exception 32 5
config qos qosmap DSCP-to-up-exception 46 6
config qos qosmap DSCP-to-up-exception 44 6

config qos qosmap trust-DSCP-upstream enable
config qos qosmap enable
```

The result of the QoS map is to map values as shown in the table below.

Table 5 EasyQoS QoS Map Values

DSCP Values	UP Values	802.11 Access Category

CS0 (DSCP 0) DSCP 1–7 CS2 (DSCP 16)	UP 0	Best Effort
CS1 (DSCP 8) DSCP 9, 11, 13, 15	UP 1	Background
AF11 (DSCP 10) AF12 (DSCP 12) AF13 (DSCP 14) DSCP 17, 19, 21, 23	UP 2	Background
AF21 (DSCP 18) AF22 (DSCP 20) AF23 (DSCP 22) DSCP 25, 27, 29, 31	UP 3	Best Effort
CS3 (DSCP 24) AF31 (DSCP 26) AF32 (DSCP 28) AF33 (DSCP 30) AF41 (DSCP 34) AF42 (DSCP 36) AF43 (DSCP 38) DSCP 33, 35, 37, 39	UP 4	Video
CS4 (DSCP 32) CS5 (DSCP 40) DSCP 41, 42, 43, 45, 47, 48	UP 5	Video
EF (DSCP 46) Admitted Voice (DSCP 44) CS6 (DSCP 48) DSCP 49–55	UP 6	Voice

CS7 (DSCP 56) DSCP 57–63	UP 7	Voice

AireOS WLCs must be running a minimum of software version 8.1.111.0 in order to support configurable DSCP-to-UP mappings (via the QoS Map) required for the EasyQoS solution. Downstream traffic is scheduled into the appropriate IEEE 802.11 AC queues based on the mapping of the DSCP value to UP, preserving QoS downstream across the Layer 2 wireless medium.

IEEE 802.11u Capable Wireless Clients

IEEE 802.11u capable wireless clients also send wireless frames with the IEEE 802.11 QoS field. Hence there is Layer 2 UP information included within frames sent by these clients, and therefore the Layer 2 trust boundary is at the wireless client, as shown in Figure 87 above. Again, the operating system and wireless drivers of the wireless device must allow application traffic to be marked with a Layer 2 UP. The application must also be able to send traffic with DSCP markings. However, the mapping of those DSCP markings to the appropriate IEEE 802.11 UPs is set via the QoS Map pushed from the AP to the IEEE 802.11u capable wireless client. Hence the Layer 3 DSCP QoS trust boundary is again at the wireless client, although the mapping of the DSCP values to User Priority is controlled now via the QoS Map configuration on the AireOS WLC. The upstream and downstream marking of traffic is identical to that discussed in the *WMM capable wireless clients (non-802.11u capable)* section.

AVC-Based Classification & Marking Policy

With the EasyQoS solution design, an AVC profile is also applied to the inner IP packet–meaning after the removal of the CAPWAP header and IEEE 802.11 frame for upstream traffic and before encapsulation with the CAPWAP header for downstream traffic. This applies to all types of wireless clients discussed in the sections above. In other words, classification & marking policies are applied in the upstream direction or in both the upstream and downstream directions, via AVC profiles applied to individual SSIDs/WLANs controlled by EasyQoS. This is accomplished via the following AireOS WLC SSID/WLAN-level commands.

```
config wlan avc x visibility enable

config wlan avc x profile EZQoS-WlanId-1 enable
```

Note that “x” refers to the ID of the particular SSID/WLAN. The profile name configured by the EasyQoS application will always be EZQoS-WlanId-1.

Note: AVC policies can specify either marking or dropping of traffic, based on an AVC profile. Only marking of traffic is used by EasyQoS.

Because AVC contains the NBAR2 engine, WLAN QoS policies consist of classification & marking policies that are based on the Cisco NBAR protocol pack supported by the WLC. These policies are currently applied at the WLC. Hence, the AVC Profile serves as another Layer 3 policy enforcement point for the original IP packet sent by the wireless client, as shown in Figure 87 above.

The specific applications within the profile are based upon the Favorites chosen by the network operator in the EasyQoS GUI, when applying a QoS Policy to a WLAN/SSID within Policy Scope that contains an AireOS WLC. This is discussed further in the *APIC-EM and the EasyQoS Application* section of this document.

AireOS WLCs are currently limited to only 32 applications per AVC policy. If less than 32 Favorites are chosen, APIC-EM EasyQoS will select the remaining applications for the AVC profile based upon applications that are most commonly used within the network. All 1300+ applications within the NBAR2 taxonomy have an attribute called "commonly-used". This attribute can have a value from 1 (least commonly used) to 10 (most commonly used). For applications that have identical values of the "commonly-used" attribute, EasyQoS will select the applications to be provisioned into the AVC-based policy based on the alphabetical name of the application.

An example of the commands provisioned by APIC-EM EasyQoS in order to create an AVC profile is shown in the configuration below.

```
config avc profile EZQoS-WlanId-1 create


config avc profile EZQoS-WlanId-1 rule add application cifs mark 10 UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application cisco-jabber-control mark
24 UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application datex-asn mark 24 UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application dnp mark 24 UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application exchange mark 10 UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application h323 mark 24 UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application imap mark 10 UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application lotus-notes mark 10
UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application mgcp mark 24 UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application ms-sms mark 10 UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application netvmg-traceroute mark 24
UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application nfs mark 10 UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application pop3 mark 10 UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application prm-nm mark 24 UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application prm-sm mark 24 UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application rpc2portmap mark 24
UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application rsvp_tunnel mark 24
UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application rtsp mark 24 UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application rtsps mark 24 UPSTREAM

config avc profile EZQoS-WlanId-1 rule add application secure-imap mark 10
UPSTREAM
```

```
config avc profile EZQoS-WlanId-1 rule add application sflow mark 24 UPSTREAM
config avc profile EZQoS-WlanId-1 rule add application sgcp mark 24 UPSTREAM
config avc profile EZQoS-WlanId-1 rule add application sip mark 24 UPSTREAM
config avc profile EZQoS-WlanId-1 rule add application sip-tls mark 24 UPSTREAM
config avc profile EZQoS-WlanId-1 rule add application skinny mark 24 UPSTREAM
config avc profile EZQoS-WlanId-1 rule add application smtp mark 10 UPSTREAM
config avc profile EZQoS-WlanId-1 rule add application snpp mark 24 UPSTREAM
config avc profile EZQoS-WlanId-1 rule add application spsc mark 24 UPSTREAM
config avc profile EZQoS-WlanId-1 rule add application ss7ns mark 24 UPSTREAM
config avc profile EZQoS-WlanId-1 rule add application svrloc mark 24 UPSTREAM
config avc profile EZQoS-WlanId-1 rule add application tpip mark 24 UPSTREAM
config avc profile EZQoS-WlanId-1 rule add application ups mark 24 UPSTREAM
```

By default, the AVC policy is unidirectional. This means that the AVC policy is applied to individual applications within the profile only in the upstream direction. In order to make the policy for individual applications within the AVC profile bi-directional, the network operator must select QoS policy to be applied bidirectionally for the given application within the EasyQoS policy screen. This is discussed *Policies* section of this document.

Note: FlexConnect designs are not supported with the APIC-EM 1.3 release of EasyQoS. Current support for AVC in FlexConnect deployments is below the requirements for EasyQoS. Hence only centralized (local mode) deployments are supported within EasyQoS.

EDCA Profile and Call Admission Control

All IEEE 802.11a/n/ac and 802.11b/g/n radios within APs controlled by an AireOS WLC in which at least one SSID/WLAN is controlled by EasyQoS, are configured with the following parameters.

- WMM EDCA profile. This is accomplished via the following global configuration commands pushed to the AireOS WLC by EasyQoS.

```
config advanced 802.11a edca-parameter wmm-default
config advanced 802.11b edca-parameter wmm-default
```

- Load-based call admission control (CAC) with up to 50% of the bandwidth reserved for voice calls and 6% of the allocated bandwidth reserved for roaming voice clients. Additionally, Expedited Bandwidth is enabled. This feature pertains to only CCXv5 compliant wireless clients. It allows such clients to indicate the urgency of a WMM traffic specifications request to the WLAN. This allows for some additional bandwidth to be used for emergency voice calls when usage exceeds 50%. This is accomplished via the following global configuration commands pushed to the AireOS WLC by EasyQoS.

```
config 802.11a cac voice acm enable
config 802.11b cac voice acm enable
```

```
config 802.11a cac voice max-bandwidth 50
config 802.11b cac voice max-bandwidth 50

config 802.11a cac voice roam-bandwidth 6
config 802.11b cac voice roam-bandwidth 6

config 802.11a exp-bwreq  enable
config 802.11b exp-bwreq  enable
```

Enabling WLANs and Radios

After the EasyQoS Static QoS configuration has been applied, the following commands provisioned by EasyQoS re-enable all WLANs and radios on an AireOS WLC.

```
config wlan enable all
config 802.11b enable network
config 802.11a enable network
```

Dynamic Qos Design

For the APIC-EM 1.3 release, Dynamic QoS is still a Beta application.

Need for Dynamic QoS

Current methods of discovering voice and video endpoints rely upon the use of protocols such as Cisco CDP, in order for the device to identify itself to the access-edge (that is, switch port) of the network. However two issues exist with regard to the use of CDP:

- CDP is not a secure protocol. CDP does not rely upon any mechanism for the endpoint device to authenticate with the network or for the network to authenticate to the endpoint device. Hence, anyone with some knowledge of programming could write an application allowing any device to "spoof" the access-edge network device (that is, switch) into thinking that the endpoint device is a Cisco IP phone or Cisco video conferencing endpoint.
- Mobile devices (such as smart phones and tablets), laptops, and PCs–running voice and video applications such as Cisco Jabber–typically do not use CDP to identify the device as being voice and/or video capable.

The APIC-EM EasyQoS solution discovers wired Cisco IP phones, Cisco IP video conferencing endpoints, Cisco TelePresence endpoints, and Cisco IP video surveillance cameras through the use of CDP. For wired devices, EasyQoS uses the IP addresses of the hardware endpoints collected through CDP information, along with the knowledge of which Catalyst switch and switch port the endpoint is connected to–in order to pre-populate ACE entries within classification & marking ACLs for Static QoS on switching devices. Because the wireless classification & marking policy deployed by EasyQoS relies on an AVC/NBAR profile, no equivalent ACE or ACL entries are generated for wireless devices. Wired devices with voice & video applications that do not use CDP also have no ACE entries generated within classification & marking ACLs for Static QoS on switching devices.

Dynamic QoS is designed to address these devices, as well as to provide a more authoritative source of information regarding whether flows should be allowed onto the network with voice and video markings, rather than simply trusting CDP information that can be easily spoofed.

Dynamic QoS Operation

For Dynamic QoS, a REST-based API has been implemented within APIC-EM. This API allows a call signaling device, such as CUCM, to inform APIC-EM EasyQoS when a voice and/or video call is established, and also when the voice and/or video call is terminated. Dynamic QoS applies to Catalyst switches (wired) currently.

Dynamic QoS for Wired Devices

Dynamic QoS for wired devices affects the ingress classification & marking policy of both access-layer switches and the distribution-layer switches to which those access-layer switches are connected.

APIC-EM EasyQoS Pre-Configuration for Access-Layer Switches

When Dynamic QoS is first enabled within EasyQoS–but before any voice and/or video calls are placed–APIC-EM will create class-map and policy-map shells for each switch port on every switch configured with the role of an access-layer switch within the policy scope. An example of the class-map shells and policy-map shells for one switch port is shown below.

```
!
class-map match-any prm-DYN-Gig1/0/13#DYN_VOICE
 match access-group name prm-DYN-Gig1/0/13#DYN_VOICE__acl
class-map match-any prm-DYN-Gig1/0/13#DYN_VIDEO
 match access-group name prm-DYN-Gig1/0/13#DYN_VIDEO__acl
class-map match-any prm-DYN-Gig1/0/13#DYN_REALTIME
!
policy-map prm-DYN-Gig1/0/13
 class prm-DYN-Gig1/0/13#DYN_VOICE
  set DSCP ef
 class prm-DYN-Gig1/0/13#DYN_REALTIME
  set DSCP cs4
 class prm-DYN-Gig1/0/13#DYN_VIDEO
  set DSCP af41
!
ip access-list extended prm-DYN-Gig1/0/13#DYN_VIDEO__acl
ip access-list extended prm-DYN-Gig1/0/13#DYN_VOICE__acl
!
```

The policy-maps have entries for voice, video, and real-time (that is, TelePresence) devices. Only voice and video class-map entries are currently used for Dynamic QoS by EasyQoS. The policy-map actions for each of the class-maps is to set the marking of the media to DSCP values that are consistent with those set for voice and video media by Static QoS.

Empty ACLs are created for dynamic voice and video calls. These are dynamically populated by ACE entries based on information passed to APIC-EM by call signaling platforms such as CUCM, via the northbound REST-based API.

APIC-EM EasyQoS Pre-Configuration for Distribution-Layer Switches

The ingress classification & marking policy applied to distribution-layer switch ports that are connected to access-layer switches was discussed in the *Distribution-Layer Switch QoS Design* section of this document.

Wired Dynamic QoS Workflow

After all of the pre-configuration has been completed by APIC-EM EasyQoS, Dynamic QoS is ready to accept signaling for voice/video calls via the REST-based API from call signaling agents, such as CUCM. The following figure provides a high-level overview of how Dynamic QoS operates for wired devices, when setting-up a call.

Figure 89 **Dynamic QoS—Wired Voice/Video Call Proceeding**

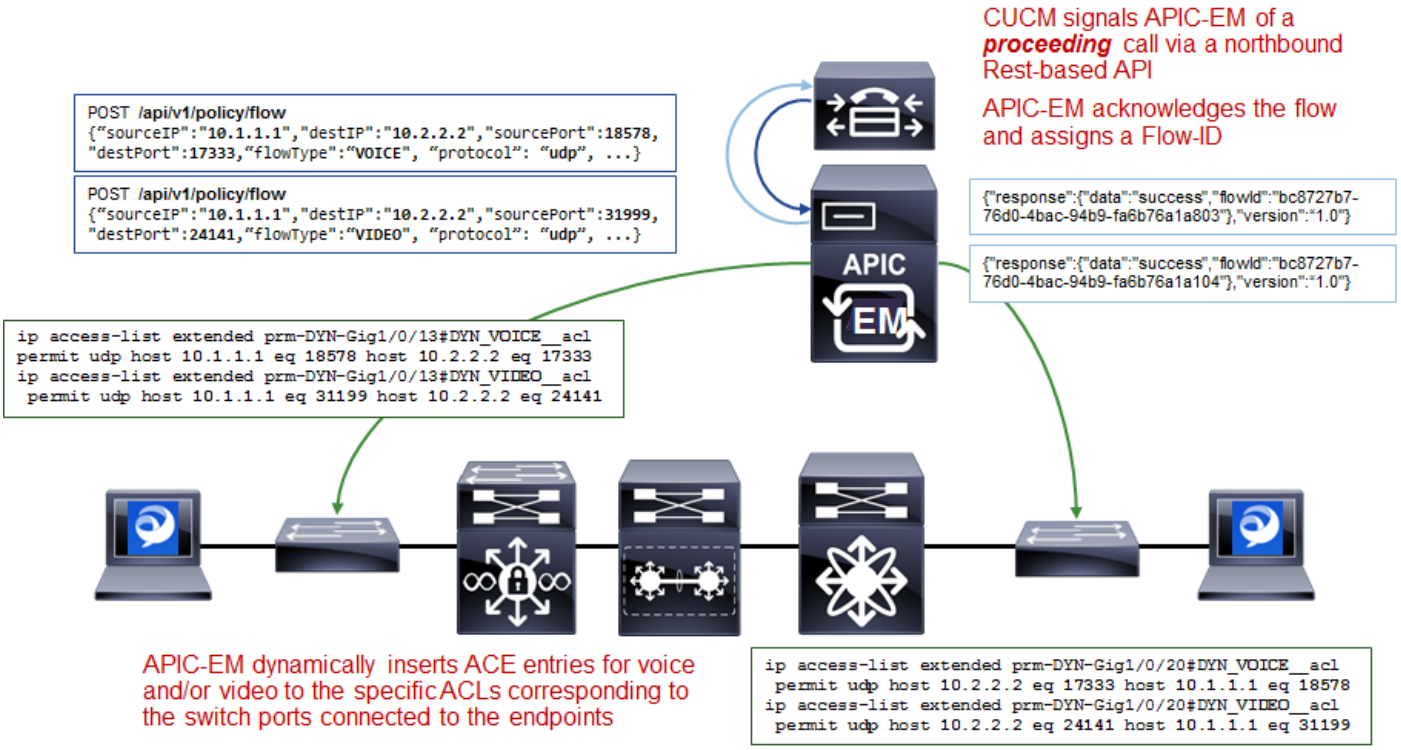

As shown in the figure above, call signaling agents such as CUCM signal to APIC-EM of a proceeding call via a northbound REST-based API. The information within the REST-based API call includes the source and destination IP addresses, the protocol (UDP), the media ports, and the type of media—VOICE or VIDEO. CUCM functions as a SIP back-to-back user-agent, meaning all SIP call signaling is between the endpoint and CUCM. Hence, CUCM has visibility into all call setup and call teardown signaling. Further, CUCM actually assigns the media ports that are used for the voice and video media sent between the endpoints. For a voice only call, only a single API call for audio media is needed. For a voice and video call, two API calls—one for audio media and one for video media is needed. APIC-EM will acknowledge each API call and will assign a flow identifier for each flow. The flow identifier is used to identify the flow during the teardown of the call.

APIC-EM then uses its southbound APIs (CLI and/or SNMP) to dynamically insert ACE entries for the voice and/or video IP addresses and media into the specific ACLs corresponding to the switch ports connected to the endpoints. The following provides an example of the configuration provisioned by EasyQoS that populates the ACE entries within the dynamic voice and video ACLs.

```
!
ip access-list extended prm-DYN-Gig1/0/13#DYN_VOICE__acl
permit udp host 10.1.1.1 eq 18578 host 10.2.2.2 eq 17333
```

```
ip access-list extended prm-DYN-Gig1/0/10#DYN_VIDEO__acl
 permit udp host 10.1.1.1 eq 31199 host 10.2.2.2 eq 24141
!
```

The ACLs are port-specific. The *Endpoints Running Cisco Jabber Soft Clients* section discusses how APIC-EM learns to which switch and switch port endpoint devices are connected.

Finally, APIC-EM EasyQoS APIC-EM EasyQoS dynamically swaps the Static QoS ingress classification & marking service-policy on the switch ports connected to the endpoints, to the service-policy used for Dynamic QoS. An example of the configuration provisioned by EasyQoS is shown below.

```
!
interface GigabitEthernet1/0/13
 no service-policy input prm-APIC_QOS_IN
 service-policy input prm-DYN-Gig1/0/13
!
```

This is the reason for the existence of the ingress classification & marking policy applied to distribution-layer switch ports connected to access-layer switches, when Dynamic QoS is enabled. When the Static QoS ingress classification & marking service-policy is replaced with the dynamic ingress classification & marking service-policy, only voice and video media are matched for that particular switch port, at the access-layer switch. The distribution-layer ingress classification & marking policy, applied by APIC-EM when EasyQoS is first enabled, is used to classify and mark all other applications that may still be sent by the endpoint device connected to the switch port. The distribution-layer ingress classification & marking policy is technically applied to all traffic from the access-layer switch. However, the policy is essentially a superset of the Static QoS policy, and hence applications from the other switch ports on the access-layer switch will not be affected, because they are already marked correctly.

The following figure provides a high-level overview of how Dynamic QoS operates for wired devices, when tearing-down a call.

Figure 90 **Dynamic QoS—Wired Voice/Video Call Termination**

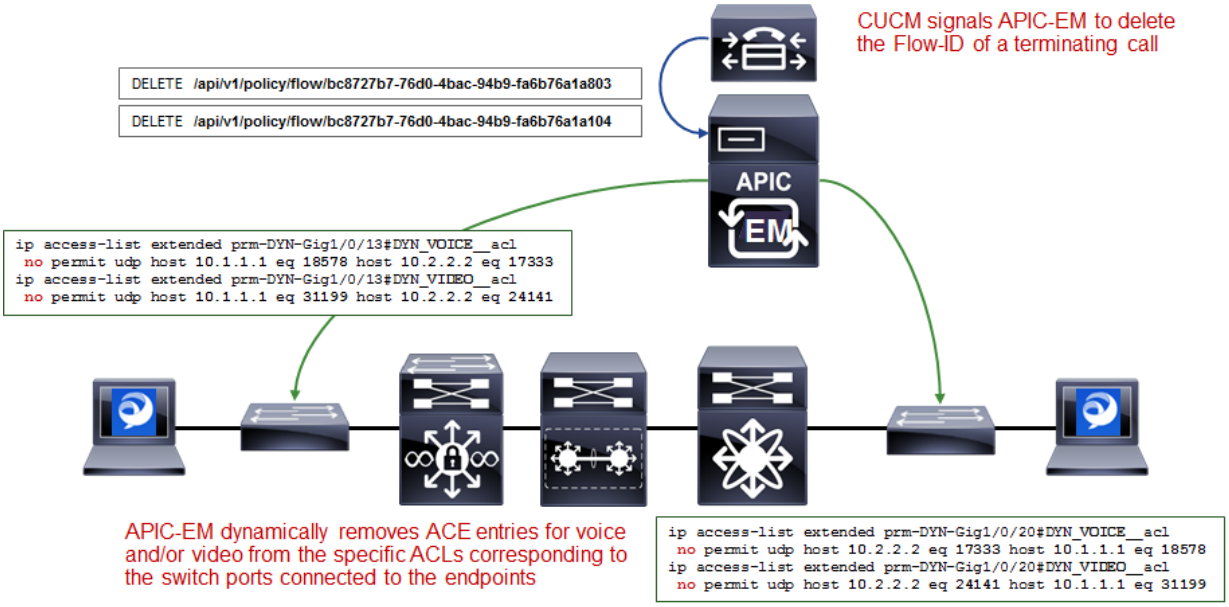

As can be seen in the figure above, CUCM signals to APIC-EM via the REST-based API that the call has terminated. For calls with voice and video media flows, two API calls may be sent by CUCM.

APIC-EM EasyQoS will then dynamically swap the ingress classification & marking service-policy on the switch ports connected to the endpoints, back to the service-policy used for Static QoS. An example of this is shown below.

```
!
interface GigabitEthernet1/0/13
 no service-policy input prm-APIC_QOS_IN
 service-policy input prm-DYN-Gig1/0/13
!
```

Finally, APIC-EM EasyQoS dynamically removes ACE entries from the dynamic ACLs, based on CUCM signaling.

```
!
ip access-list extended prm-DYN-Gig1/0/13#DYN_VOICE__acl
 no permit udp host 10.1.1.1 eq 18578 host 10.2.2.2 eq 17333
ip access-list extended prm-DYN-Gig1/0/10#DYN_VIDEO__acl
 no permit udp host 10.1.1.1 eq 31199 host 10.2.2.2 eq 24141
!
```

Endpoints Running Cisco Jabber Soft Clients

Cisco Jabber soft clients are also discovered by APIC-EM and populated within Dynamic QoS ACLs. However, Cisco Jabber soft clients do not typically send CDP information. Instead, Cisco IP Device Tracking (IPDT) must be enabled on Catalyst switches in order to support the ability for APIC-EM to discover and populate devices with Cisco Jabber soft clients.

Note: Devices running the Cisco Medianet Services Interface (MSI) may generate CDP information. However, the MSI interface is only supported on Windows 7 & 8, and Mac OS X 10.7 and 10.8 platforms, and no further development is expected for the MSI.

For a Catalyst 3850/3650 Series switch, IPDT is enabled via the following interface-level configuration command.

```
!
ip device tracking maximum x
!
```

"X" is the maximum number of devices to be tracked on the interfaces. The values range from 0 to 63,535 devices. A value of 0 disables IP device tracking on the switch port.

IPDT causes the switch port to send Address Resolution Protocol (ARP) probes (ARP request packets) periodically. The interval between ARP probes can be controlled via the following interface-level configuration command.

```
!
ip device tracking probe interval x
!
```

"X" is the interval between which ARP probes are sent. The default value is 30 seconds.

The number of ARP probes sent per interval can also be controlled via the following interface-level configuration command.

```
!
ip device tracking probe count x
!
```

"X" is the number of ARP probes sent per interval. The default value is 3 probes.

At a global configuration level, the following command can be used to delay the switch port from sending the ARP probes after a link up event or link flap (link down/link up).

```
!
ip device tracking probe delay 10
!
```

This command can be used to prevent the switch from sending an ARP probe while the device connected to the switch port checks for duplicate IP addresses.

As a prerequisite for supporting the ability for APIC-EM to dynamically discover adds/moves/changes of hardware endpoint devices and automatically update ACL entries for these devices, the network operator will need to enable SNMP traps (particularly the link up/link down trap) on access-layer switches to be sent to APIC-EM.

After the switch port connected to a hardware endpoint goes up/down, APIC-EM will receive an SNMP trap. APIC-EM will start collecting information from the access-layer switch that generated the SNMP trap, about the new endpoints. This will take approximately 80 seconds, plus the time needed for the collection of the device information.

Additional Resources

- Release Notes for Cisco Application Policy Infrastructure Controller Enterprise Module, Release 1.3.x

http://www.cisco.com/c/en/us/support/cloud-systems-management/one-enterprise-network-controller/products-release-notes-list.html

- Compatibility Information

http://www.cisco.com/c/en/us/support/cloud-systems-management/one-enterprise-network-controller/products-device-support-tables-list.html

- Install and Upgrade Guides

http://www.cisco.com/c/en/us/support/cloud-systems-management/one-enterprise-network-controller/products-installation-guides-list.html

- Configuration Guides

http://www.cisco.com/c/en/us/support/cloud-systems-management/one-enterprise-network-controller/products-installation-and-configuration-guides-list.html

- Programming Guides

http://www.cisco.com/c/en/us/support/cloud-systems-management/one-enterprise-network-controller/products-programming-reference-guides-list.html

- Configuration Examples and TechNotes

http://www.cisco.com/c/en/us/support/cloud-systems-management/one-enterprise-network-controller/products-configuration-examples-list.html

- Troubleshooting Guides

http://www.cisco.com/c/en/us/support/cloud-systems-management/one-enterprise-network-controller/products-troubleshooting-guides-list.html

About Cisco Validated Designs

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information, visit

http://www.cisco.com/go/designzone.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.