WHITE PAPER

Understanding the Protocols of Class-Based QoS (Quality of Service)

Nick Sandoval
Product Marketing Manager

MOXA®

Introduction

"When everything is a priority, nothing is a priority."

-Simon Fulleringer

Quality of service (QoS) refers to a network's ability to achieve maximum bandwidth and deal with other network performance elements such as latency, error rate and uptime. As important of a topic QoS is, surprisingly it's one of the least understood in networking.

QoS is comprised of three different methodologies: Predictive, Flow-based, and Class-based. Of the three, Class-based is the best understood and one of the most utilized. With it, you will be able to build more deterministic and resilient networks that can ensure your important data, like control signals, get through even when the network is in a period of high utilization. Class-based is also the most useful as it specifies traffic prioritization. This means that some traffic is deemed more important than others, and thus is treated differently in a network.

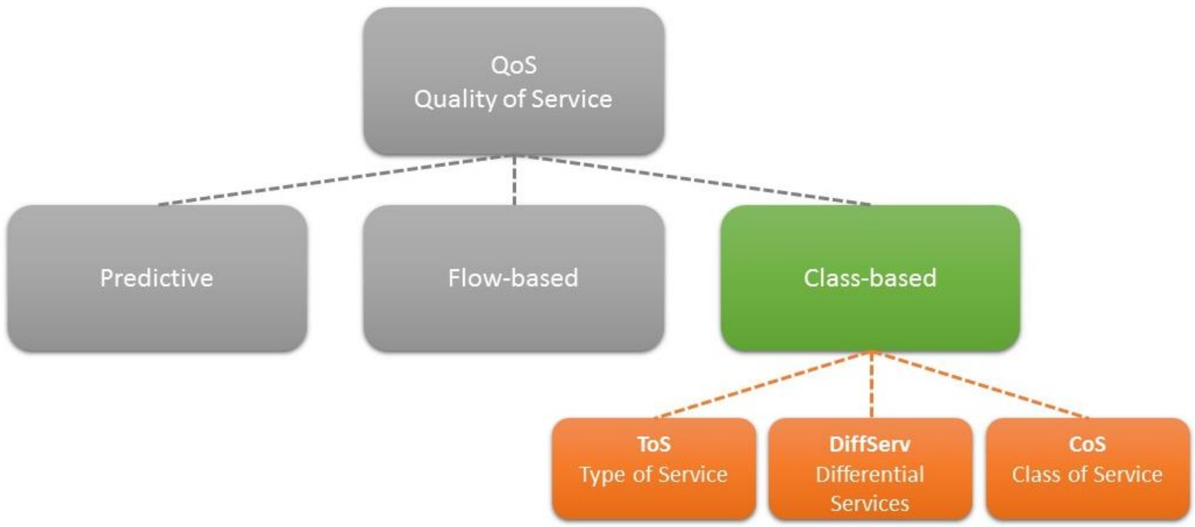

Released on January 2, 2018

How to contact Moxa
Tel: 1-714-528-6777
Fax: 1-714-528-6778

With Class-based QoS, packets are grouped into queues based on similar quality requirements. Each queue is prioritized differently and will receive different access levels to network resources. Within each queue, the packets are transmitted in the order they are received. For Class-based QoS to be effective, packets are marked before being transmitted. They are marked so the receiving device knows which queue the packet belonged to. Based on these markings, the receiving network device chooses how to prioritize traffic. Higher priority traffic will be placed into a queue that has first access to available bandwidth. Lower priority traffic will be placed into a queue where it will yield to or wait for the higher priority traffic to transmit first.

To better explain Class-based QoS, we will use a busy airport scenario. In this example, let's say there are three lines for going through security, each with a different priority level. Line 1 never waits, line 2 will move when line 1 is empty, but line 3 will move only when line 1 and 2 are empty. Line 1 is for pilots and stewards, line 2 is for first class flyers, and line 3 is for economy flyers. Boarding tickets assign each passenger a priority and a designated line.

The airline ticketing system is the same principle we explained earlier about marking packets. The Passenger Service Agent for the airline will check each ticket and direct passengers to their assigned lines. From each line, flyers will board using the FIFO (first in first out) method. One important caveat is that this prioritization method is only relevant when there are queues or lines. Similar to the airline boarding analogy, there is no benefit of having a higher boarding priority if there is no one in line to board. This means that an individual with an economy ticket will board the plane as if they had a first class ticket. Thus, in networks, Class-based QoS is only effective in periods of congestion.

When to use Class-based QoS

Class-Based QoS greatly benefits networks that are deterministic, or mission-critical. It can be used on its own or in conjunction with other QoS methodologies; however, the more deterministic or mission-critical a network, the more important it is to use more than one QoS schema. Class-based QoS will only be relevant in periods of congestion, thus if a network has little to no traffic, there will be little benefit. Nonetheless, even in such networks, it provides an extra level of assurance that the SLA[1] (Service Level Agreement) is maintained.

[1] The SLA is the agreement that the user has with the service provider in regards to peak-bit rate, traffic prioritization and more. A caveat with SLA is that it is not well-defined or standardized from one service provider to the next.

The Class-based QoS Protocols

Type of Service (ToS), Differential Services (DiffServ) and Class of Service (CoS) are three of the most widely used protocols for Class-based QoS. All of these protocols use markings/tags in order to determine priority. While CoS and ToS both accomplish the same objective of prioritizing traffic, CoS achieves this in the Ethernet frame (Layer 2), while ToS does so in the IP packet (Layer 3).

In short, Class-based QoS works in the following manner:

1. When data is received (ingress), the IP header is read and the DSCP header is interpreted. Because this is in the IP header, hardware optimization can still take place so QoS will have minimal impact. Using the IP header is unique to ToS and DiffServ protocols, while CoS uses the Ethernet header.
2. If the QoS value is trusted, the packet will be placed in the appropriate outgoing queue (egress). This queue contains traffic of similar prioritization. If it is not trusted, then SLA requirements will dictate how the packet is to be handled. For example, the QoS value may be changed to best-effort. Best-effort describes how there is no guarantee that the data is delivered or prioritized.
3. Each differentiated queue is typically handled using best-effort. There are several different methodologies in which queues are managed. A strict methodology with high priority traffic will always transmit while pausing the transmittal of lower priority traffic. If lower priority traffic is paused for too long, the data might lose pertinence to the receiver due to time lag. To prevent this, a weighted-fair queue scheme can be implemented. Which scheme to implement is determined by the networking device. The process will repeat itself at the next router or network device.

It is best practice for larger network deployments to use trust boundaries and classifiers. A classifier is a data traffic manager that regulates data coming into the inner layers of a network. Trust boundaries are set at the lowest layer under then network architect's full control. For example, if a network is fully controlled by the network architect, it can be set to the edge layer[2]. This means that any prioritization value inside the boundary is trusted, and conversely prioritizations outside the boundary are not trusted. The classifier is what regulates entry through this boundary and ensures that the SLA is maintained.

The classifier may choose to trust or not trust Class-based QoS tagging that requests to enter the trust boundary. If the classifier trusts the tagged values, then the data will pass through the boundary and into the network unmodified. If it does not trust the values, it can change the tagged values to something that adheres to the SLA. It may even reset all QoS values to best-effort should that be the SLA.

[2] Networks are built with several levels of hierarchy. The end devices like computers, PLC, and such are known as edge devices and are collectively known as the edge layer. Interconnecting all the edge devices is the access layer. Inside the network there may be a distribution layer and in larger networks, a core layer. A trust boundary would usually reside at the access layer.

Type of Service (ToS):

As a Layer 3 protocol, ToS functions independently of Layer 1 or Layer 2. The ToS field has 8 bits reserved in an IPv4[3] header. This header is comprised of two major components: IP Precedence (3 bits) and DTR - Delay, Throughput, and Reliability (3 bits).

Type of Service (ToS): 8 bites = 1byte

0	1	2	3	4	5	6	7
Precedence			D	T	R	ECN Field	

ToS is a legacy protocol that never gained much traction. In its onset, ToS had only 3 bits reserved for traffic prioritization, known as IP Precedence. The other bits remained largely unused. Because of this, ToS was redefined to incorporate most the bits for traffic prioritization through DiffServ.

Differential Services (DiffServ)

DiffServ is also a Layer 3 protocol. It renames the 8 bits used for ToS to DSCP (Differentiated Services Code Point). Despite the renaming to DSCP, traffic prioritization through ToS is preserved. This is because the 3 bits used in IP Precedence are remapped into DSCP, and therefore devices that utilize IP Precedence will be backwards compatible with newer DSCP devices. DSCP has a total of 64 priority levels. This means in the airport scenario, you now have 64 passenger lines/queues compared to the original 8. However, most networking hardware have limited queues, which means that although DSCP has 64 priority levels, you are limited to what is available through your hardware. To make up for this discrepancy, the DSCP priority levels are mapped to match the number of queues provided by the hardware.

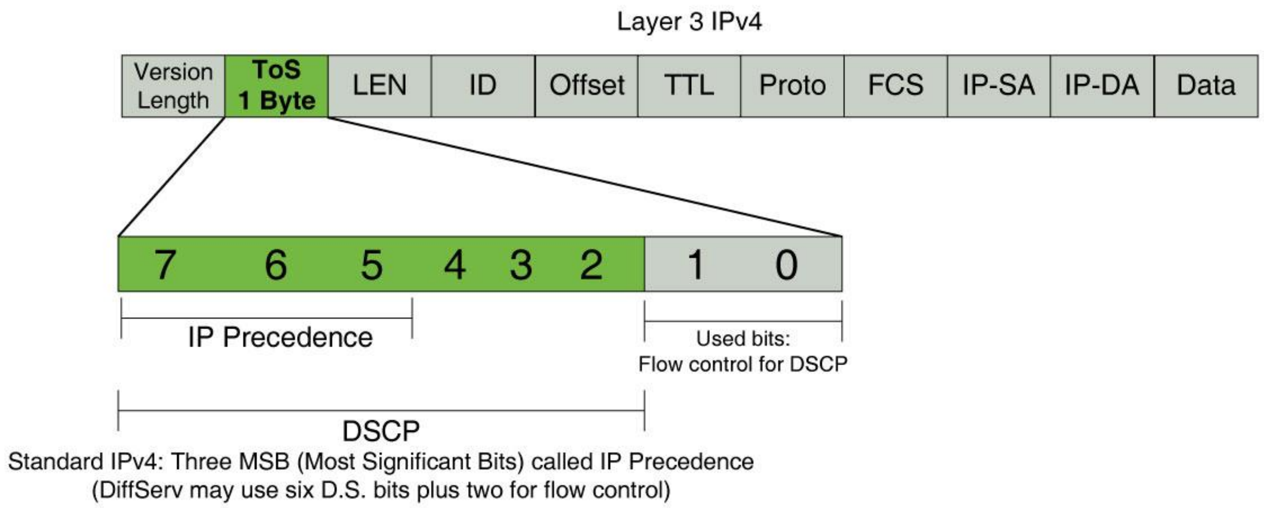

[3] IP version 4 or Internet Protocol version 4 (IPv4) is a networking protocol used in data communication. Based on the best-effort model, it's a connectionless protocol for use on packet-switched layer networks. It provides a logic based connection between network devices by providing identification for each device.

Class of Service (CoS)

Class of Service can be one of two things: CoS can be used to identify all Class-based QoS protocols, or it can be a protocol itself. This is synonymous with naming a car company "Car". Nonetheless, when most refer to CoS, the reference is usually to the protocol. In this white paper, we will refer to the Class of Service protocol as CoS.

CoS as a protocol is considered part of the 4 byte field added to the Ethernet frame that contains both CoS traffic prioritization and VLAN (802.1Q) information. CoS as a protocol is in the Ethernet frame and thus Layer 2. Being a Layer 2 implementation means that it is applied at the switch level and not dependent on Layer 3 protocols. There are eight levels of priority classifications with CoS.

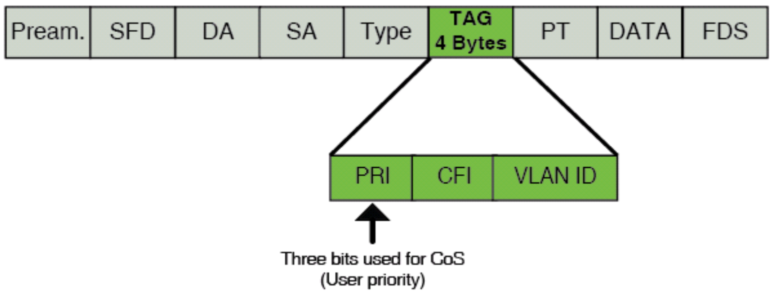

Of all the Class-based protocols, CoS and DiffServ are both widely used depending on the scale of the network. CoS is often found in networks that utilize VLANs but not routing. This is because CoS uses the same tag that VLANs use. DiffServ is used when VLANs are not used or cannot be forwarded from end to end.

Recommendations:

Trust boundaries

Setting up trust boundaries is imperative. Doing so ensures that data prioritization is not hijacked in a way where non-important data is prioritized over data that needs determinism or is mission-critical. In such cases, it is important to setup the classifier to retag any traffic outside of the trust boundary so that it is not prioritized over trusted data.

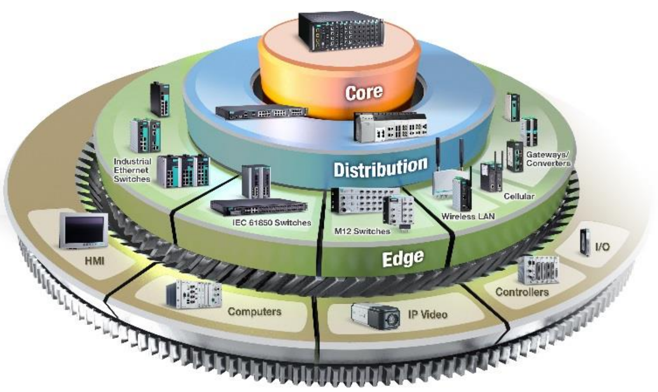

For networks that are controlled from edge to core, the trust boundary can be applied all the way to the edge-device. The advantage of an edge device being the classifier is that it allows for greater assurance that correct traffic is prioritized. Should a network device be the classifier, there may be limits to what is classified. For example, some network devices only operate on the port level, in which all traffic from a particular port will take on the classification level specified. However, since not all edge devices support QoS natively, it may still be necessary to setup your network devices as a classifier. This is especially the case in an industrial setting where few edge devices have the ability to support QoS protocols.

QoS Protocols

CoS is best used when VLANs are being utilized. This is because it uses the same VLAN tagging scheme. However, in larger networks that use Layer 3 routing without VLANs, DiffServ is the preferred method. DiffServ also provides many more QoS levels that allow for more granular prioritization control.

The QoS protocol applied depends on the protocol supported and the topology of the network. For example, if you were to utilize a camera that has the ability to use DiffServ, and no other protocol, then DiffServ should be the protocol utilized. This is unless it is remapped into another protocol like CoS. QoS mapping is a method for a classifier to apply the marking of one protocol to another and is used to ensure compatibility across a network.

Conclusion

Class-based QoS allows you to build networks that are deterministic and resilient when the network is in high utilization. Through prioritization, some traffic is deemed more important than others, and thus is treated as such in a network. Utilizing either Layer 2 or Layer 3, if not both, throughout the network will ensure greater network determinism.