

Implementing and Deploying Quality of Service (QoS) for EtherNet/IP

Brian Batke
Rockwell Automation

Presented at the ODVA
2009 CIP Networks Conference & 13th Annual Meeting
February 25, 2009
Howey-in-the-Hills, Florida

Abstract

Greater connectivity to Industrial Ethernet brings more applications, more traffic, and therefore more potential for congestion and resulting latency and packet loss. To ensure that critical EtherNet/IP traffic is delivered according to application needs, standard mechanisms for Quality of Service have been defined for EtherNet/IP devices and can be used in conjunction with network infrastructure for optimal performance.

1 Introduction

Industrial Ethernet has grown from being just a faster way to upload and download programs to a legitimate control network to a converged IT-manufacturing network supporting many different services. What once carried a small amount of traffic for obscure industrial protocols now may carries larger amounts of data that include standard applications and protocols. The network is faster, devices are faster, and application uses have multiplied. What was once considered absurd might soon be common: high speed motion control sharing network bandwidth with voice and video.

For such a converged network to be feasible, the critical traffic – that which enables widgets to be made and controls moving machinery – must be delivered according to the needs of the application. In an application-rich, open networking environment, standard Quality of Service (QoS) mechanisms can be employed in order to meet the application's data delivery needs.

This paper presents Quality of Service mechanisms that are defined for EtherNet/IP implementations, and shows how the mechanisms may be used in an application network.

2 Quality of Service Overview

2.1 A QoS Analogy

You are at the airport, waiting for your flight's boarding process to begin. As the boarding time nears, people start making the inevitable surge toward the door. Before long there is a large, unorganized mass of passengers, each of whom wants to board early in order to wedge a suitcase into a space designed to fit a small knapsack.

But not everyone is forced to fight the masses elbowing for position. You stand and watch in envy as the first class passengers stroll through the gate. Next, the elite membership holders are called. They experience a slight delay but quickly disappear down the tunnel.

Finally the economy passengers are called. Boarding by row is intended to smooth out the congestion, but still often results in a long wait. And during the process the first class/elite line remains open so that one of the privileged can walk past the angry line of common-folk to the front of the queue.

When the traffic is light, say 10 people on a 737, there would be no issue with everyone boarding at once. But with a full plane, a "QoS" mechanism allows those important passengers to board with as little delay and trauma as possible.

2.2 QoS Principles

The above analogy demonstrates important QoS principles that are also employed in networks:

- Distinguishing one traffic stream from another (classification). The airline distinguishes between first class, elite, and regular economy passengers who are further classified by row.
- Assigning a label to each piece of traffic (marking) and making sure the traffic is within the traffic profile (policing). Each passenger's ticket identifies whether it is first-class or elite.
- Providing different treatment to the different traffic classes (queuing and scheduling). First class is called before elite, who in turn are called before economy. At any time, a first-class passenger can move to the front of the queue.

In a network, the above QoS principles can be used to provide levels of service appropriate to the different needs of traffic types. Tolerance to loss, delay, and jitter are the primary factors in determining the QoS requirements for different types of traffic.

The following table shows the tolerance to loss, delay, and jitter for EtherNet/IP-related traffic.

Loss, Delay, Jitter Tolerance for EtherNet/IP

Traffic Type	Traffic Characteristics	Tolerance to		
		Loss	Delay	Jitter
IEEE 1588	Fixed size messages, 44 or 64 bytes payload. Produced on a cyclic basis, once per second	High performance apps not tolerant to loss.	For high performance apps, less than 50 microseconds.	Up to the maximum delay.
CIP Motion	Fixed size messages, typically 80– 220 bytes. Usually produced according to a cyclic rate. High performance apps: current target: 32 messages in 500 usecs (32 axes in 500usecs). Ultimate target is close to wire speed (100 axes in 1ms).	Can tolerate occasional loss of up to 2 consecutive packets. Target: 0 loss.	For high performance apps, less than 100 microseconds.	Up to the maximum delay.
CIP I/O	Fixed size messages, typically 100 – 500 bytes. Usually produced according to a cyclic rate. Can also be produced on application change of state. Typical cyclic rate per stream: 1ms to 500 ms.	Application dependent. Generally can tolerate occasional loss. CIP connection typically times out if 4 consecutive packets lost. Target: 0 packet loss.	Application dependent. Tolerance proportional to the packet rate. Target: < 25% of the packet interval.	Application dependent. Generally can tolerate jitter up to the maxim tolerable delay.
CIP Safety I/O	Fixed messages, typically on the order of 16 bytes payload. Produced according to a cyclic rate. Typical cyclic rate: 5-10ms or greater	Can tolerate occasional loss of 1 packet in a safety period (1 out of 4 transmissions).	Dependent on the packet rate; in general can tolerate delay on the order of 5 ms	Up to the maximum delay.

HMI / Messaging	<p>Variable size messages.</p> <p>Typical size: 100 – 500 bytes; likely to be larger in the future.</p> <p>Produced under application control. Can be at regular cyclic intervals, or based on application state or user action.</p>	Can tolerate packet loss so long as TCP connection remains.	Can tolerate delay so long as the TCP connection remains.	Can tolerate large degree of jitter.
-----------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------	-----------------------------------------------------------	--------------------------------------

2.3 QoS Protocol Mechanisms

2.3.1 Overview of QoS Protocol Mechanisms

Many different QoS mechanisms have been defined for different network protocols and media. Of the many QoS mechanisms, two are important for EtherNet/IP: Differentiated Services and IEEE 802.1D/Q.

Differentiated Services, or “Diffserv” operates at layer 3 (IP) and is defined in the following RFCs:

- RFC 2474 – Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers
- RFC 2475 – An Architecture for Differentiated Services
- RFC 2597 – Assured Forwarding PHB Group
- RFC 3140 – Per Hop Behavior Identification Codes
- RFC 3246 – An Expedited Forwarding PHB (Per-Hop Behavior)
- RFC 4594 – Configuration Guidelines for DiffServ Service Classes

IEEE 802.1D/Q operates at layer 2, and is defined in the following IEEE specifications:

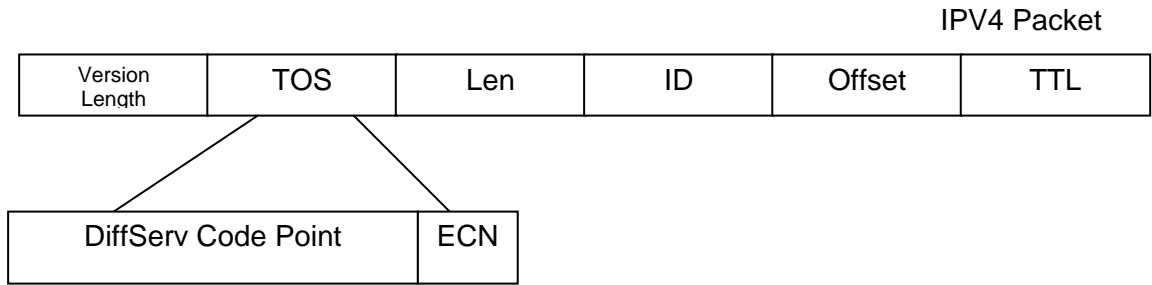
- IEEE Std 802.1D – 2004 (defines the use of priority in the 802.1Q frame format)
- IEEE Std 802.1Q – 2005 (defines VLAN operation including the tagged frame format)

The following sections present a brief summary of Diffserv and 802.1D/Q based QoS.

2.3.2 Layer 3: Differentiated Services

Differentiated Services (DiffServ) is a model for providing different treatment of traffic classes based on their characteristic behavior and tolerance to delay, loss, and jitter. The overall model is defined in RFC 2475. DiffServ allows nodes (typically switches and routers) to service packets based on class of traffic as defined by the DiffServ Codepoint

(DSCP) and associated Per-Hop Behavior (PHB) characteristics. The DSCP values appear in the ToS field in the IP header, as shown below:



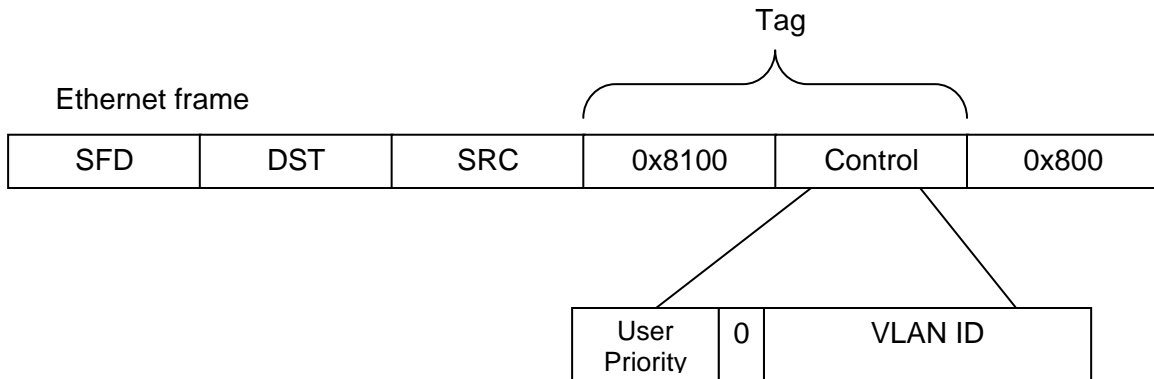
The RFCs define a number of different Per-Hop Behaviors for different traffic classes. The following shows some of the different network applications and their corresponding DSCP values and PHBs:

Application	Diffserv Classification		RFC
	PHB	DSCP	
Network Control	CS5	48	RFC 2474
VoIP Telephony	EF	46	RFC 3246
Call Signaling	CS5	40	RFC 2474
Multimedia Conferencing	AF41	40	RFC 2597
Real-Time Interactive	CS4	32	RFC 2474
Multimedia Streaming	AF31	26	RFC 2597
Broadcast Video	CS3	24	RFC 2474
Low-Latency Data	AF21	18	RFC 2597
High-Throughput Data	AF11	10	RFC 2597
Best Effort	DF	0	RFC 2474
Low-Priority Data	CS1	8	RFC 3662

EtherNet/IP uses DSCP values defined as “local use”, as described in section 3.

2.3.3 Layer 2: IEEE 802.1D/Q

IEEE 802.1Q defines an Ethernet frame format that allows inclusion of VLAN ID and priority. The 802.1Q frame has EtherType of 0x8100 and a 4-byte prefix between the Source and Type fields of the frame. The tagged frame defines a 3-bit field to specify 8 priority levels, further specified in 802.1D. Priority 7 is the highest. Priority 0 is the lowest.



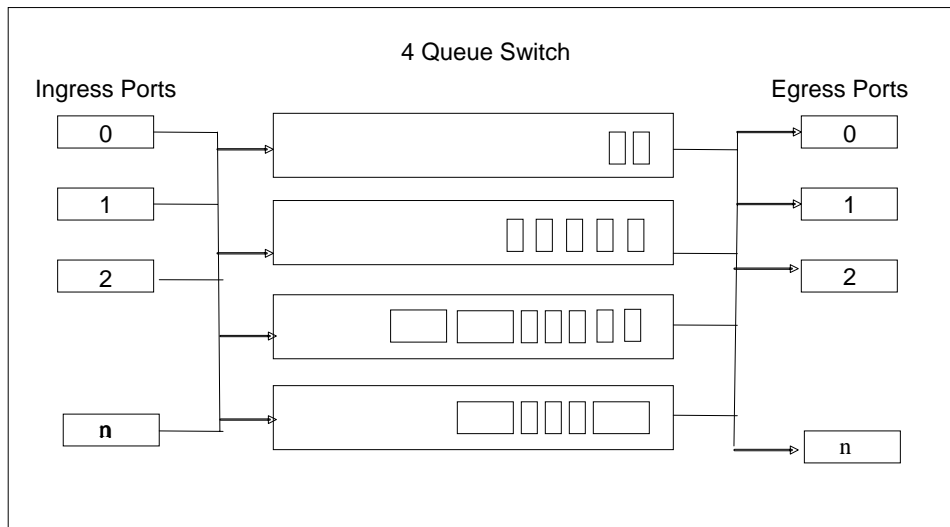

802.1D, in Annex G, suggests the following mapping of priority values to traffic types. Note that the mapping shows some consistency with Diffserv, but not an exact match:

User Priority	Acronym	Traffic Type
0 (default)	BE	Best Effort
1	BK	Background
2	Spare	
3	EE	Excellent Effort
4	CL	Controlled Load
5	VI	Video
6	VO	Voice
7	NC	Network Control

2.3.4 QoS Support in the Infrastructure

Simply marking packets with DSCP or 802.1D priority results in little value unless the network infrastructure is able to provide service based on those markings. Fortunately most, if not all, managed switches support multiple queues and differentiation based on 802.1D, and many support DSCP. Many, if not most, routers support Diffserv.

The following diagram illustrates a single-queue switch versus a multiple-queue switch.



In the single-queue switch, packets of all priorities are intermingled in a single queue, first-come, first-served. Higher priority packets may have to wait as lower priority packets are serviced.

In a multiple-queue switch, packets can be directed to different queues based on their priority markings. The different queues are then serviced according to a scheduling algorithm, such that higher priority packets are given precedence over lower priority packets.

It is beyond the scope of this paper to fully examine all the different QoS capabilities of switches and routers. As a general rule, QoS operation involves the following actions:

- **Classification:** examining packets and determining the QoS markings (e.g., DSCP and/or 802.1D priority). Many switches can also be configured to classify packets based on TCP or UDP port number, VLAN, or the physical ingress port. In such switches, EtherNet/IP traffic can be classified based on the TCP and UDP port numbers.
- **Policing:** per configuration, determine if the incoming packet is within the profile or outside the profile.
- **Marking:** the incoming packet may be further marked (e.g., upgraded or downgraded).
- **Queuing and Scheduling:** determine into which queue to place the packet; service queues according to the configured scheduling algorithm and parameters. Switches and routers may support scheduling on a strict priority basis, round-robin basis, or other.

Many switches and routers provide extensive configuration options, allowing different mappings of priorities to queues, selection of buffer space, and different scheduling algorithms.

3 QoS for EtherNet/IP Devices

3.1 Summary

The EtherNet/IP specification, starting with Edition 1.6, defines standard behavior for EtherNet/IP devices with respect to QoS. For devices that implement the Device Level Ring (DLR) protocol, QoS is required. For all other EtherNet/IP devices, QoS is highly recommended.

The overall approach for QoS for EtherNet/IP calls for devices to mark their packets with a priority value, using Diffserv Code Points and/or 802.1D priority values. By explicitly marking packets with a priority value, switches and routers are able to differentiate EtherNet/IP traffic from non-critical traffic, as well as differentiate specific EtherNet/IP traffic streams (e.g., IEEE 1588 vs. I/O vs. Explicit Messaging).

The following list summarizes the QoS behavior for EtherNet/IP:

- For CIP transport class 0 and 1 connections (i.e., UDP-based), there is a defined mapping of CIP priorities to 802.1D priorities and DiffServ Code Points.
- For UCMM and CIP transport class 3 connections (i.e., TCP-based), there is a single defined DiffServ Code Point and 802.1D priority value.
- For PTP (IEEE 1588) messages, there are DiffServ Code Points and 802.1D priority values corresponding to the two different types of PTP messages.

- When QoS is implemented, the default behavior is to mark packets with DSCP values. Devices may optionally support sending and receiving 801.1Q frames with the corresponding priority values. If supported, sending tagged frames is disabled by default in order to prevent device interoperability problems. The end user is responsible to enabling the tagged frame behavior and ensuring interoperability between devices.
- The QoS Object provides a means to configure DSCP values, and a means to enable/disable sending of 802.1Q tagged frames.
- There are no requirements for devices to mark traffic other than CIP or IEEE 1588, but devices are free to do so.

3.2 Mapping CIP Traffic to DSCP and 802.1D

The following table shows the default priority values for CIP and IEEE 1588 traffic. The priority values can be changed via the QoS Object.

Traffic Type	CIP Priority	DSCP	802.1D Priority ¹	CIP Traffic Usage (recommended)
PTP Event (IEEE 1588)	n/a	59 ('111011')	7	n/a
PTP General (IEEE 1588)	n/a	47 ('101111')	5	n/a
CIP class 0 / 1	Urgent (3)	55 ('110111')	6	CIP Motion
	Scheduled (2)	47 ('101111')	5	Safety I/O I/O
	High (1)	43 ('101011')	5	I/O
	Low (0)	31 ('011111')	3	No recommendation at present
CIP UCMM CIP class 3 All other EtherNet/IP encapsulation messages	All	27 ('011011')	3	CIP messaging

¹ Sending 802.1Q tagged frames is disabled by default

As shown in the table, for CIP transport class 0 and 1 messages the DSCP and 802.1D values are based on the CIP priority value as indicated in the ForwardOpen service that established the connection. In contrast, the PTP and CIP explicit messages have DSCP and 802.1D values independent of CIP priority.

The default DSCP values are “local use” values, as defined in RFC 2474. Local use values were chosen in order to differentiate EtherNet/IP traffic from other, non-industrial traffic such as VoIP and video.

3.3 EtherNet/IP DSCP Usage

When an EtherNet/IP device supports QoS, EtherNet/IP messages are marked with the appropriate DSCP priorities as the default behavior. The above table shows the default DSCP values, which may be changed via the QoS object.

All EtherNet/IP devices – even those that don’t support sending packets with priority markings – must support receiving packets with non-zero DSCP values. This is a requirement of RFC 791 and RFC 1122. Some implementations have in the past erroneously assumed that the ToS field would always be 0 and as a result improperly computed UDP checksums.

3.4 EtherNet/IP 802.1D/Q Usage

In addition to using DSCP values, EtherNet/IP devices may optionally send traffic with 802.1Q tagged frames, using the default values in Table 1. When 802.1Q is supported, sending 802.1Q tagged frames must be off by default. The QoS Object allows sending tagged frames to be enabled as well as changing the default priority values.

The reason that sending 802.1Q frames is not the default behavior is that some devices do not support reception of tagged frames. When such a frame is received, the frame is dropped. In addition, some managed switches drop tagged frames unless the port is explicitly configured such that tagged frames are accepted.

Enabling sending of tagged frames is an action that must be taken by the end user, should such behavior be desired. The user is responsible for ensuring that both sending and receiving devices support 802.1Q frames, and for ensuring proper configuration of switches.

A scenario in which a user might wish to enable 802.1Q frames is when the user’s switches support QoS based on 802.1Q priority but not DSCP.

4 QoS in Application Networks

4.1 *Benefits of QoS*

QoS is not required for every EtherNet/IP application. For small, isolated networks with relatively low bandwidth requirements, QoS would offer little added benefit. In what situations would QoS offer benefits?

- Integrated networks where EtherNet/IP traffic may cross network boundaries and share bandwidth with other traffic. For example, a cell-level controller sends an interlock message through a router to another controller in a different cell (different subnet). By configuring QoS in the router, the EtherNet/IP interlock messages can be assured of being serviced before other less critical traffic.
- Applications with sensitivity to latency and jitter, for example CIP Motion and CIP Sync applications. Using QoS, in periods of network congestion, time-sensitive data can receive preferential treatment.
- Protection against network events. In the situations such as broadcast storms, QoS allows EtherNet/IP traffic to be given precedence, allowing the application to continue to function in spite of the network event.

4.2 *Recommended Switch Capabilities*

In order to fully support QoS for EtherNet/IP applications, the following capabilities are recommended for switches:

- At least two queues, and preferably four.
- Support classification of packets according to DSCP markings, since this is the default behavior for QoS-enabled EtherNet/IP devices.
- Support classification based on TCP or UDP port number, to allow prioritization of EtherNet/IP packets from devices that are not QoS-enabled.
- Strict priority scheduling, for the high priority queue at least, to ensure lowest latency and jitter for high-performance applications.
- Follow the default queue mappings for EtherNet/IP, or at least allow configuration to support optimal mapping of priorities to queues for EtherNet/IP (see below).

4.3 Recommended Priority-Queue Mappings

4.3.1 Recommended Mapping

The following queue mappings are recommended, for applications other than CIP Motion applications.

Queue	DSCP	802.1D	Notes
1	59	7	Network control traffic, PTP
2	43, 46, 47, 55	4,5,6	CIP class 0/1 (UDP based) connections, VoIP
3	24, 27, 31	2,3	CIP explicit messages, VoIP
4	Everything else	0,1	Other traffic

4.3.2 Recommended Mapping for CIP Motion

The following queue mappings are recommended for CIP Motion applications.

Queue	DSCP	802.1D	Notes
1	59	7	Network control traffic, PTP
2	55	6	CIP Motion (CIP Urgent prio)
3	43, 46, 47	4,5	CIP class 0/1 (UDP based) connections, VoIP
4	Everything else	0,1	Other traffic

5 QoS Measurements

The following tables show results of QoS measurements, with the recommended queue mappings as shown above.

Traffic Profile:

Stream	% Traffic
01 - IEEE1588/PTP	1
02 - CIP I/O Unicast	40
03 - CIP I/O Mcast	40
04 - HMI/Messaging	10
05 - VOIP	3
06 - Standard TCP/IP	5
07 - Background Mcast	0.5
08 - Background Bcast	0.5

The above traffic profile was sent on two different ingress ports at 100% port utilization to a single destination port. This resulted in a 2:1 oversubscription for the egress port.

The following table shows the percentage of traffic loss seen at the egress port:

Stream	Percentage loss
IEEE1588/PTP	0
CIP I/O Unicast	49
CIP I/O Mcast	49
VOIP	49
HMI/Messaging	60
Standard TCP/IP	70
Background Mcast	80
Background Bcast	80

The highest priority traffic experienced no traffic loss. The CIP I/O traffic experienced roughly 50% loss, but it should be noted that the port was 2:1 oversubscribed. As the table shows, the lower priority traffic experienced significantly more loss.

The ideas, opinions, and recommendations expressed herein are intended to describe concepts of the author(s) for the possible use of CIP Networks and do not reflect the ideas, opinions, and recommendation of ODVA per se. Because CIP Networks may be applied in many diverse situations and in conjunction with products and systems from multiple vendors, the reader and those responsible for specifying CIP Networks must determine for themselves the suitability and the suitability of ideas, opinions, and recommendations expressed herein for intended use. Copyright ©2009 ODVA, Inc. All rights reserved. For permission to reproduce excerpts of this material, with appropriate attribution to the author(s), please contact ODVA on: TEL +1 734-975-8840 FAX +1 734-922-0027 EMAIL odva@odva.org WEB www.odva.org

CIP, Common Industrial Protocol, CIP Motion, CIP Safety, CIP Sync, CompoNet, CompoNet CONFORMANCE TESTED, ControlNet, ControlNet CONFORMANCE TESTED, DeviceNet, EtherNet/IP, EtherNet/IP CONFORMANCE TESTED are trademarks of ODVA, Inc. DeviceNet CONFORMANCE TESTED is a registered trademark of ODVA, Inc. All other trademarks are property of their respective owners.