

Buffers, Queues, and Thresholds on the Catalyst 6500 Ethernet Modules

August 2007

Executive Summary

With the wide variety of Ethernet modules currently available on the Catalyst 6500 series switches, determining the size of the port buffers, the number and type of the receive and transmit queues, and the number of tail-drop or weighted random early detect (WRED) thresholds can be a daunting task.

This document provides a brief discussion of how the port buffers, queues, and thresholds are used in the Catalyst 6500 series switch Ethernet modules, and then describes the buffer sizes, queue structures, and thresholds supported on each of the modules.

This document does not discuss the Cisco 7600 Optical Services Module (OSM) interfaces or the new SIP and SPA modules that are also supported in the Catalyst 6500 series switches.

Overview of Buffers, Queues, and Thresholds

All of the Catalyst 6500 series Ethernet modules implement some form of receive and transmit buffering. These buffers are used to store frames as forwarding decisions are made within the switch, or as packets are enqueued for transmission on a port at a rate greater than the physical medium can support.

In the Catalyst 6500 architecture, access into the switch fabric itself is almost never the bottleneck. Rather, on the transmit side, one or several ports are the likely destination for a majority of the packets entering the switch. As such, the receive-side port buffers on the Ethernet modules are relatively small compared to the transmit-side port buffers.

Without Quality of Service (QoS) features enabled, all packets have equal access to the port buffers, regardless of the type of traffic. For example, Voice over IP (VoIP) packets get placed in the same buffer as HTTP (web) or FTP (file download) data packets. Furthermore, in the event of congestion (that is, a port buffer overflows), all traffic is equally subject to discard. To summarize, packets are serviced in the order in which they were received, and if the buffer is full, all subsequent packets are dropped. This is known as First In, First Out (FIFO) queuing with tail-drop.

When you enable QoS on the switch, the port buffers are carved into one or more individual queues. Each queue has one or more drop thresholds associated with it. The combination of multiple queues within a buffer, and the drop thresholds associated with each queue, allow the switch to make intelligent decisions when faced with congestion. Traffic sensitive to jitter and delay variance, such as VoIP packets, can be moved to the head of the queue for transmission, while other less important or less sensitive traffic can be buffered or dropped.

Port Queue and Drop Threshold Types

When QoS is enabled, the multiple queues and drop thresholds on the Ethernet module switch ports are enabled. There are several different configurations of queue types and thresholds, depending on the model of the Ethernet module.

Ingress and egress scheduling are always based on the COS value associated with the frame. By default, higher COS values are mapped to higher queue numbers. COS 5 traffic, typically associated with VoIP traffic, is mapped to the strict priority queue, if present.

In addition to the different queues, each standard queue has one or more drop thresholds. There are two types of drop thresholds:

- **Tail-drop thresholds:** On ports with tail-drop thresholds, frames of a given COS value are admitted to the queue until the drop threshold associated with that COS value is exceeded; subsequent frames of that COS value are discarded until the threshold is no longer exceeded.

For example, if COS 1 is assigned to queue 1, threshold 2, and the threshold 2 watermark is 60%, then frames with COS 1 will not be dropped until queue 1 is 60% full. All subsequent COS 1 frames are dropped until the queue is less than 60% full.

- **WRED drop thresholds:** On ports with WRED drop thresholds, frames of a given COS value are admitted to the queue based on a random probability designed to avoid buffer congestion. The probability of a frame with a given COS being admitted to the queue or discarded depends on the weight and threshold assigned to that COS value.

For example, if COS 2 is assigned to queue 1, threshold 2, and the threshold 2 watermarks are 40% (low) and 80% (high), then frames with COS 2 will not be dropped until queue 1 is at least 40% full. As the queue depth approaches 80%, frames with COS 2 have an increasingly higher probability of being discarded rather than being admitted to the queue. Once the queue is over 80% full, all COS 2 frames are dropped until the queue is less than 80% full. The frames the switch discards when the queue level is between the low and high thresholds are picked out at random, rather than on a per-flow basis or in a FIFO manner. This method works well with protocols, such as TCP, capable of adjusting to periodic packet drops by backing off and adjusting their transmission window size.

When referring to the individual queues and thresholds on a port, a rather terse nomenclature is used. This terminology describes the number of strict priority queues (if present), the number of standard queues, and the number of tail-drop or WRED thresholds within each of the standard queues. The different queue and threshold types on the Catalyst 6500 Ethernet modules are shown in Table 1.

Table 1. Receive and Transmit Port Queue and Drop Threshold Types on Catalyst 6500 Series Ethernet Modules

Port Queue and Drop Threshold Structure with QoS	Description
Receive Queues	
1q2t	One standard queue with two tail-drop thresholds
1q4t	One standard queue with four tail-drop thresholds
1q8t	One standard queue with eight tail-drop thresholds
2q4t	Two standard queues with four WRED drop thresholds per queue

Port Queue and Drop Threshold Structure with QoS	Description
2q8t	Two standard queues with eight tail-drop thresholds per queue
8q4t	Eight standard queues with four WRED drop thresholds per queue
8q8t	Eight standard queues with eight WRED drop thresholds per queue
1p1q4t	One strict-priority queue, one standard queue with four tail-drop thresholds
1p1q0t	One strict-priority queue, one standard queue with one non-configurable (100%) tail-drop threshold
1p1q8t	One strict-priority queue, one standard queue with eight configurable WRED drop thresholds and one non-configurable (100%) tail-drop threshold
Transmit Queues	
2q2t	Two standard queues with two tail-drop thresholds per queue
1p2q2t	One strict-priority queue, two standard queues with two WRED drop thresholds per queue
1p3q1t	One strict-priority queue, three standard queues with one WRED drop threshold and one non-configurable tail-drop threshold per queue
1p2q1t	One strict-priority queue, two standard queues with one WRED drop threshold and one non-configurable (100%) tail-drop threshold per queue
1p3q4t	One strict-priority queue, three standard queues with four WRED drop thresholds per queue. Supports both DWRR and SRR.
1p3q8t	One strict-priority queue, three standard queues with eight WRED drop thresholds per queue. Supports both DWRR and SRR.
1p7q8t	One strict-priority queue, seven standard queues with eight WRED drop thresholds per queue
1p7q4t	One strict-priority queue, seven standard queues with four WRED drop thresholds per queue. Supports both DWRR and SRR.

Buffer Size, Queues, and Thresholds by Ethernet Module

Table 2 provides the following information for each of the Catalyst 6500 series Ethernet modules:

- Total buffer size per port (Total Buffer Size)
- Overall receive buffer size per port (Rx Buffer Size)
- Overall transmit buffer size per port (Tx Buffer Size)
- Port receive queue and drop threshold structure (Rx Port Type)
- Port transmit queue and drop threshold structure (Tx Port Type)
- Default size of receive buffers per queue with QoS enabled (Rx Queue Sizes)
- Default size of transmit buffers per queue with QoS enabled (Tx Queue Sizes)

Table 2. Buffer Size, Queues, and Thresholds by Ethernet Module

Module Model Name	Module Description	Total Buffer Size	Rx Buffer Size	Tx Buffer Size	Rx Port Type	Tx Port Type	Rx Queue Sizes	Tx Queue Sizes
Supervisor Engine Modules (Gigabit Ethernet and Ten Gigabit Ethernet Uplink Ports)								

Module Model Name	Module Description	Total Buffer Size	Rx Buffer Size	Tx Buffer Size	Rx Port Type	Tx Port Type	Rx Queue Sizes	Tx Queue Sizes
VS-S720-10G-3C/XL	Supervisor 720-10G TenGigabitEthernet Uplink ports in 10Gig-only mode	191.8MB	104.2MB	87.6MB	8q4t	1p7q4t DWRR, SRR	Q8—20.8MB Q7—0MB Q6—0MB Q5—0MB Q4—0MB Q3—0MB Q2—0MB Q1—83.4MB	SP—13.9MB Q7—0MB Q6—0MB Q5—0MB Q4—0MB Q3—13.0MB Q2—17.3MB Q1—43.4MB
	Supervisor 720-10G Ten Gigabit Ethernet Uplink ports	191.8MB	104.2MB	87.6MB	2q4t	1p3q4t DWRR, SRR	Q2—20.8MB Q1—83.4MB	SP—13.9MB Q3—13.0MB Q2—17.3MB Q1—43.4MB
	Supervisor 720-10G Gigabit Ethernet Uplink ports	17.7MB	9.6MB	8.1MB	2q4t	1p3q4t DWRR, SRR	Q2—1.9MB Q1—7.7MB	SP—1.2MB Q3—1.2MB Q2—1.6MB Q1—4.1MB
WS-SUP720 WS-Sup720-3B WS-SUP720-3BXL	All Supervisor Engine 720 uplink ports	512KB	80KB	432KB	1p1q4t	1p2q2t WRR	SP—10KB Q1—70KB	SP—64KB Q2—64KB Q1—304KB
WS-SUP32-10G	Supervisor 32-10GE Ten Gigabit Ethernet uplink ports	193MB	105MB	88MB	2q8t	1p3q8t DWRR, SRR	Q2—1MB Q1—104MB	SP—14MB Q3—13MB Q2—17MB Q1—44MB
	Supervisor 32-10GE Gigabit Ethernet Uplink port	17.7MB	9.6MB	8.1MB	2q8t	1p3q8t DWRR, SRR	Q2—1.9MB Q1—7.7MB	SP—1.2MB Q3—1.2MB Q2—1.6MB Q1—4.1MB
WS-SUP32-GE	All Supervisor Engine 32 uplink ports	17.7MB	9.6MB	8.1MB	2q8t	1p3q8t DWRR, SRR	Q2—1.9MB Q1—7.7MB	SP—1.2MB Q3—1.2MB Q2—1.6MB Q1—4.1MB

Module Model Name	Module Description	Total Buffer Size	Rx Buffer Size	Tx Buffer Size	Rx Port Type	Tx Port Type	Rx Queue Sizes	Tx Queue Sizes
WS-X6K-S2U-MSFC2 WS-X6K-S2-MSFC2 WS-X6K-S2-PFC2	All Supervisor Engine 2 uplink ports	512KB	80KB	432KB	1p1q4t	1p2q2t WRR	SP— 10KB Q1— 70KB	SP— 64KB Q2— 64KB Q1— 304KB
WS-X6K-S1A-MSFC2 WS-X6K-SUP1A-MSFC WS-X6K-SUP1A-PFC WS-X6K-SUP1A-2GE	All Supervisor Engine 1A uplink ports	512KB	80KB	432KB	1p1q4t	1p2q2t WRR	SP— 10KB Q1— 70KB	SP— 64KB Q2— 64KB Q1— 304KB
WS-X6K-SUP1-2GE	Supervisor Engine 1 uplink ports	512KB	80KB	432KB	1q4t	2q2t WRR	Q1— 80KB	Q2— 80KB Q1— 352KB
Ethernet and Fast Ethernet Modules								
WS-X6524-100FX-MM	24-port 100BaseFX fabric-enabled with MT-RJ connectors	1116KB	28KB	1088KB	1p1q0t	1p3q1t DWRR	SP— 6KB Q1— 22KB	SP— 272KB Q3— 272KB Q2— 272KB Q1— 272KB
WS-X6548-RJ-21	48-port 10/100BaseTX fabric-enabled with RJ-21 connectors	1116KB	28KB	1088KB	1p1q0t	1p3q1t DWRR	SP— 6KB Q1— 22KB	SP— 272KB Q3— 272KB Q2— 272KB Q1— 272KB
WS-X6548-RJ-45	48-port 10/100BaseTX fabric-enabled with RJ-45 connectors	1116KB	28KB	1088KB	1p1q0t	1p3q1t DWRR	SP— 6KB Q1— 22KB	SP— 272KB Q3— 272KB Q2— 272KB Q1— 272KB
WS-X6324-100FX-MM	24-port 100BaseFX with MT-RJ connectors	128KB	16KB	112KB	1q4t	2q2t WRR	Q1— 16KB	Q2— 22KB Q1— 90KB
WS-X6324-100FX-SM	24-port 100BaseFX with MT-RJ connectors	128KB	16KB	112KB	1q4t	2q2t WRR	Q1— 16KB	Q2— 22KB Q1— 90KB
WS-X6348-RJ-45 WS-X6348-RJ-45V	48-port 10/100BaseTX with RJ-45 connectors	128KB	16KB	112KB	1q4t	2q2t WRR	Q1— 16KB	Q2— 22KB Q1— 90KB

Module Model Name	Module Description	Total Buffer Size	Rx Buffer Size	Tx Buffer Size	Rx Port Type	Tx Port Type	Rx Queue Sizes	Tx Queue Sizes
WS-X6348-RJ21V	48-port 10/100BaseTX with RJ-21 connectors	128KB	16KB	112KB	1q4t	2q2t WRR	Q1—16KB	Q2—22KB Q1—90KB
WS-X6224-100FX-MT	24-port 100BaseFX with MT-RJ connectors	64KB	8KB	56KB	1q4t	2q2t WRR	Q1—8KB	Q2—16KB Q1—40KB
WS-X6248-RJ-45	48-port 10/100BaseTX with RJ-45 connectors	64KB	8KB	56KB	1q4t	2q2t WRR	Q1—8KB	Q2—16KB Q1—40KB
WS-X6248-TEL	48-port 10/100BaseTX with RJ-21 connectors	64KB	8KB	56KB	1q4t	2q2t WRR	Q1—8KB	Q2—16KB Q1—40KB
WS-X6248A-TEL	48-port 10/100BaseTX with RJ-21 connectors	128KB	16KB	112KB	1q4t	2q2t WRR	Q1—16KB	Q2—22KB Q1—90KB
WS-X6196-RJ-21 WS-X6196-21AF	96-port 10/100BaseTX with RJ-21 connectors	1116KB	28KB	1088KB	1p1q0t	1p3q1t DWRR	SP—6KB Q1—22KB	SP—272KB Q3—272KB Q2—272KB Q1—272KB
WS-X6148A-RJ-45 WS-X6148A-45AF	48-port 10/100BaseTX with RJ-45 connectors	5.3MB	60KB	5.3MB	1p1q4t	1p3q8t DWRR	SP—8KB Q1—52KB	SP—810KB Q3—810KB Q2—1.1MB Q1—2.6MB
WS-X6148-RJ-45 WS-X6148-RJ-45V WS-X6148-45AF	48-port 10/100BaseTX with RJ-45 connectors	128KB	16KB	112KB	1q4t	2q2t WRR	Q1—16KB	Q2—22KB Q1—90KB
WS-X6148-RJ-21 WS-X6148-RJ-21V WS-X6148-21AF	48-port 10/100BaseTX with RJ-21 connectors	128KB	16KB	112KB	1q4t	2q2t WRR	Q1—16KB	Q2—22KB Q1—90KB
WS-X6148X2-RJ-45 WS-X6148X2-45AF	96-port 10/100BaseTX with RJ-45 connectors	1116KB	28KB	1088KB	1p1q0t	1p3q1t DWRR	SP—6KB Q1—22KB	SP—272KB Q3—272KB Q2—272KB Q1—272KB

Module Model Name	Module Description	Total Buffer Size	Rx Buffer Size	Tx Buffer Size	Rx Port Type	Tx Port Type	Rx Queue Sizes	Tx Queue Sizes
WS-X6148-FE-SFP	48-port 100BaseFX with SFP transceivers	5.4MB	60KB	5.4MB	1p1q4t	1p3q8t DWRR	SP— 8KB Q1— 52KB	SP— 810KB Q3— 810KB Q2— 1.1MB Q1— 2.7MB
WS-X6024-10FL-MT	24-port 10Base-FL with MT-RJ connectors	64KB	8KB	56KB	1q4t	2q2t WRR	Q1— 8KB	Q2— 16KB Q1— 40KB
Gigabit Ethernet and 10/100/1000 Modules								
WS-X6816-GBIC	16-port 1000BaseX dual-fabric with GBIC transceivers	512KB	80KB	432KB	1p1q4t	1p2q2t WRR	SP— 10KB Q1— 70KB	SP— 64KB Q2— 64KB Q1— 304KB
WS-X6748-GE-TX	48-port 10/100/1000T dual-fabric with RJ-45 connectors	1.3MB	166KB	1.2MB	2q8t (with DFC3)	1p3q8t DWRR	Q2— 33KB Q1— 133KB (with DFC3)	SP— 175KB Q3— 175KB Q2— 233KB
					1q8t (with CFC)		Q1— 166KB (with CFC)	Q1— 583KB
WS-X6748-SFP	48-port 1000BaseX dual-fabric with SFP transceivers	1.3MB	166KB	1.2MB	2q8t (with DFC3)	1p3q8t DWRR	Q2— 33KB Q1— 133KB (with DFC3)	SP— 175KB Q3— 175KB Q2— 233KB
					1q8t (with CFC)		Q1— 166KB (with CFC)	Q1— 583KB
WS-X6724-SFP	24-port 1000BaseX single-fabric with SFP transceivers	1.3MB	166KB	1.2MB	2q8t (with DFC3)	1p3q8t DWRR	Q2— 33KB Q1— 133KB (with DFC3)	SP— 175KB Q3— 175KB Q2— 233KB
					1q8t (with CFC)		Q1— 166KB (with CFC)	Q1— 583KB
WS-X6548-GE-TX WS-X6548V-GE-TX WS-X6548-GE-45AF	48-port 10/100/1000T fabric-enabled with RJ-45 connectors	1.4MB (shared among 8 ports)	185KB (shared among 8 ports)	1.2MB (shared among 8 ports)	1q2t (shared among 8 ports)	1p2q2t (shared among 8 ports) WRR	Q1— 185KB	SP— 182KB Q2— 182KB Q1— 849KB

Module Model Name	Module Description	Total Buffer Size	Rx Buffer Size	Tx Buffer Size	Rx Port Type	Tx Port Type	Rx Queue Sizes	Tx Queue Sizes
WS-X6516-GBIC	16-port 1000BaseX fabric-enabled with GBIC transceivers	512KB	73KB	439KB	1p1q4t	1p2q2t WRR	SP—9KB Q1—64KB	SP—64KB Q2—64KB Q1—311KB
WS-X6516A-GBIC	16-port 1000BaseX fabric-enabled with GBIC transceivers	1MB	144KB	880KB	1p1q4t	1p2q2t WRR	SP—26KB Q1—118KB	SP—128KB Q2—128KB Q1—624KB
WS-X6516-GE-TX	16-port 10/100/1000T fabric-enabled with RJ-45 connectors	512KB	73KB	439KB	1p1q4t	1p2q2t WRR	SP—9KB Q1—64KB	SP—64KB Q2—64KB Q1—311KB
WS-X6408-GBIC	8-port 1000BaseX with GBIC transceivers	512KB	80KB	432KB	1q4t	2q2t WRR	Q1—80KB	Q2—80KB Q1—352KB
WS-X6408A-GBIC	8-port 1000BaseX with GBIC transceivers	512KB	73KB	439KB	1p1q4t	1p2q2t WRR	SP—9KB Q1—64KB	SP—64KB Q2—64KB Q1—311KB
WS-X6416-GBIC	16-port 1000BaseX with GBIC transceivers	512KB	73KB	439KB	1p1q4t	1p2q2t WRR	SP—9KB Q1—64KB	SP—64KB Q2—64KB Q1—311KB
WS-X6416-GE-MT	16-port 1000BaseSX with MT-RJ connectors	512KB	73KB	439KB	1p1q4t	1p2q2t WRR	SP—9KB Q1—64KB	SP—64KB Q2—64KB Q1—311KB
WS-X6316-GE-TX	16-port 1000BaseT with RJ-45 connectors	512KB	73KB	439KB	1p1q4t	1p2q2t WRR	SP—9KB Q1—64KB	SP—64KB Q2—64KB Q1—311KB
WS-X6148A-GE-TX WS-X6148A-GE-45F	48-port 10/100/1000T with RJ-45 connectors	5.5MB	120KB	5.4MB	1q2t	1p3q8t DWRR	Q1—120KB	SP—810KB Q3—810KB Q2—1.1MB Q1—2.7MB
WS-X6148-GE-TX WS-X6148V-GE-TX WS-X6148-GE-45AF	48-port 10/100/1000T with RJ-45 connectors	1.4MB (shared among 8 ports)	185KB (shared among 8 ports)	1.2MB (shared among 8 ports)	1q2t (shared among 8 ports)	1p2q2t (shared among 8 ports) WRR	Q1—185KB	SP—182KB Q2—182KB Q1—849KB

Ten-Gigabit Ethernet Modules

Module Model Name	Module Description	Total Buffer Size	Rx Buffer Size	Tx Buffer Size	Rx Port Type	Tx Port Type	Rx Queue Sizes	Tx Queue Sizes
WS-X6708-10G-3C WS-X6708-10G-3CXL	8-port 10GbE dual-fabric with X2 transceivers	256MB	109MB	92MB	8q4t	1p7q4t DWRR, SRR	Q8— 21.9MB Q7— 0.5KB Q6— 0.5KB Q5— 0.5KB Q4— 0.5KB Q3— 0.5KB Q2— 0.5KB Q1— 87.4MB	SP— 14.6MB Q7— 0.5KB Q6— 0.5KB Q5— 0.5KB Q4— 0.5KB Q3— 13.6MB Q2— 18.2MB Q1— 45.5MB
WS-X6704-10GE	4-port 10GbE dual-fabric with XENPAK transceivers	16MB	2MB	14MB	8q8t (with DFC3)	1p7q8t DWRR	Q8— 400KB Q7— 22KB Q6— 22KB Q5— 22KB Q4— 22KB Q3— 22KB Q2— 22KB Q1— 1.5MB	SP— 2.2MB Q7— 10KB Q6— 10KB Q5— 10KB Q4— 10KB Q3— 2.2MB Q2— 2.9MB Q1— 7.2MB
					1q8t (with CFC)		Q1— 2MB	
WS-X6502-10GE	1-port 10GbE fabric-enabled with OIM connectors	64.2MB	256KB	64MB	1p1q8t	1p2q1t DWRR	SP— 51KB Q1— 205KB	SP— 15.3MB Q2— 17.9MB Q1— 30.7MB
WS-X6501-10GEX4	1-port 10GbE fabric-enabled with SC connectors	64.2MB	256KB	64MB	1p1q8t	1p2q1t DWRR	SP— 51KB Q1— 205KB	SP— 15.3MB Q2— 17.9MB Q1— 30.7MB
Optical Services Modules**								
OSM-*	All Optical Services Modules (L2 GbE fabric-enabled ports only)	512KB	73KB	439KB	1p1q4t	1p2q2t WRR	SP— 9KB Q1— 64KB	SP— 64KB Q2— 64KB Q1— 311KB

** This document only discusses the four "standard" Gigabit Ethernet ports present on most of the OSMs. It does not discuss the 4-port GE-WAN OSM or any other OSM-related interfaces.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

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

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

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0807R)

Printed in USA

C11-386109-01 08/08