

# Nokia IPSO 4.1/4.2 QoS Feature Enhancements

April 2007

# Nokia IPSO 4.1/4.2 QoS Feature Enhancements

## Introduction

Traditional IP networks offer best-effort service for delivery of data packets. In this model, all packets compete equally for network resources. However, this best-effort service cannot provide reliability and predictability in end-to-end packet delivery, making it unsuitable for real-time and business critical applications. As the popularity of the “everything over IP” paradigm increases, a network requires Quality of Service (QoS) in addition to best-effort service in order to offer service differentiation based on requirements of users and applications such as high-quality video and delay-sensitive real-time voice.

## Differentiated Services

Differentiated Services (DiffServ) is a computer networking architecture that specifies a coarse-grained mechanism for classifying, managing network traffic and providing QoS guarantees on modern IP networks. DiffServ can provide low-latency service to critical network traffic such as voice or video while providing simple best effort traffic guarantees to non-critical services such as web traffic and file transfers. In the DiffServ model, each data packet is placed into a limited number of traffic classes. Each router on the network is configured to differentiate traffic based on its class. Each traffic class can be managed differently, ensuring preferential treatment for higher-priority traffic on the network.

Nokia IPSO software supports the DiffServ model of QoS. In addition to best-effort forwarding mechanism, Nokia IPSO uses two of the more popular forwarding mechanism proposed by the Internet Engineering Task Force (IETF): Expedited Forwarding (EF) Per Hop Behavior (PHB) and Assured Forwarding (AF) PHB group. Per Hop Behaviors define the packet forwarding properties associated with a class of traffic. Users can define different PHBs to offer low-loss, low-latency forwarding properties or best-effort forwarding properties. Forwarding classes are associated with Differentiated Services Code Points (DSCP); they are encoded in the Type of Service (TOS) field of each packet’s IP header.

The following table shows the values for the different queue levels.

Name of Queue Level	Priority	IETF DiffServ Code point	Queue Specifier Value
Internetwork Control	0	0xc0	7
Expedited Forwarding	1	0xb8	6

AF Class 4	2	Find DSCP codes in the below given table	5
AF Class 3	3		4
AF Class 2	4		3
AF Class 1	5		2
Best Effort	7		0

The AF PHB group includes four traffic classes. Packets with each AF class can be marked with one of the three possible drop-precedence values (high, medium or low) for further packet differentiation.

The table below shows twelve separate DSCP encodings derived from the combination of AF classes and drop precedence.

	AF Class 1	AF Class 2	AF Class 3	AF Class 4
Low Drop	001010xx	010010xx	011010xx	100010xx
Medium Drop	001100xx	010100xx	011100xx	100100xx
High Drop	001110xx	010110xx	011110xx	100110xx

Internetwork Control (IC) traffic, such as routing messages and keepalives, should be configured to use the IC queue class so that it receives precedence over regular IP traffic. Note that locally originated control traffic is automatically sent through this queue.

The following sections address the Expedited Forwarding PHB and Assured Forwarding PHB group in greater detail.

## IPSO 4.1 QoS Features Overview

Nokia IPSO 4.1 QoS functionality allows packet streams to be filtered, shaped, and prioritized. The prioritization mechanism conforms to RFC 2598, the Expedited Forwarding PHB specification of the IETF DiffServ Working Group. Three queue levels are pre-defined and used: the Internetwork Control, Expedited Forwarding (EF), and Best Effort.

Expedited Forwarding PHB has the characteristics of low delay, low loss and low jitter suitable for voice, video, and other real-time services. EF traffic receives strict priority queuing above all other traffic classes.

A differentiated services-compliant network node classifies traffic through Access Control Lists (ACLs). These ACLs select packets based on the value of the DS field, along with queue class and packet scheduling mechanisms capable of delivering the specific packet forwarding treatment indicated by the DS field value.

Any ACL rule that has the action prioritize must have an aggregation class applied to it. An Aggregation Class (AGC) determines whether the traffic stream meets certain throughput goals (aggregation classes are not the same as traffic classes). Traffic that meets these goals is conformant and is allowed to pass through. Traffic that does not meet these goals is non-conformant and is dropped. The combination of ACLs and AGCs form the basis of the filtering, shaping, and prioritization tools.

### How Packets are Filtered

Classified traffic is filtered immediately. The actions for filtering are:

Accept - The accept action forwards the traffic.

Drop - The drop action drops the traffic without any notification

Reject- The action drops the traffic and sends an ICMP error message to the source.

### How Traffic is Shaped

Classified traffic is shaped to a mean rate. The traffic-shaping tool uses a token bucket algorithm which means that you can configure a burst size from which bursts can “borrow.” Measured over longer time intervals, the traffic will coerce to the configured mean rate. Over shorter intervals, traffic is allowed to burst to higher rates. The coercion works by adding delay to packets that must wait for more tokens to arrive in the bucket. Traffic drops when more bursts arrive than can be accommodated by the shaping queue. Both outgoing and incoming traffic streams can be shaped.

### How Traffic is Prioritized

Classified traffic can be given preferential treatment. Higher-priority traffic must be policed to prevent starvation of lower-priority service traffic. Traffic that conforms to the configured policing rate is marked with the supported DSCP associated with the EF class. When such traffic is processed by the output queue scheduler, it receives favorable priority treatment. Prioritization is only relevant for outgoing traffic. Incoming traffic is never marked. In Nokia IPSO 4.1, the supported prioritization mechanism is Expedited Forwarding Per Hop Behavior with strict priority scheduling mechanism.

## QoS Feature Enhancements in IPSO 4.2

The QoS feature enhancements in IPSO 4.2 include:

1. Assured Forwarding PHB
2. Scheduling Mechanism: Weighted Round Robin (WRR) and Cascade
3. Dropping Mechanism: Tail drop and Weighted Random Early Detection (WRED)
4. Policing of ICMP and multicast packets
5. DSCP Marking Before Encryption
6. Integration with SecureXL

### Assured Forwarding PHB

There are five classes of packet marking in the DiffServ architecture: Expedited Forwarding (EF) and Assured Forwarding Classes 1-4. EF is already implemented in the Nokia IPSO 4.1. The Assured Forwarding Classes 1-4 (AF1 –AF4) is implemented in IPSO 4.2. Any traffic that does not meet the requirements of any other defined classes is placed in the Best Effort class.

The Assured Forwarding PHB group provides delivery of IP packets in four independent forwarded AF classes. Within each AF class, an IP packet can be assigned one of three different levels of drop precedence. In case of congestion, the drop precedence of a packet determines the relative importance of the packet within the AF class. A congested Differentiated Services node tries to protect packets with a lower drop precedence value from being lost by preferably discarding packets with a higher drop precedence value.

### Scheduling Mechanism

Rather than using strict priority queuing (supported in IPSO 4.1), more balanced queue scheduling algorithms such as weighted round robin (WRR) and cascade scheduling can be used in IPSO 4.2 QoS.

#### Weighted Round Robin

Weighted round robin (WRR) is an algorithm that uses priority based scheduling in which packets are sent from the highest priority level first. The WRR scheduler uses weights proportional to a traffic class' bandwidth allocation. The weight determines the number of bytes/packets that a traffic class is allowed to send in a scheduling round. Higher bandwidth queues can send more packets each time they go through a service round. When WRR is enabled, the default queue weight and queue length for each of the eight individual queues are equal.

## Cascade Scheduling

When Cascade scheduler is used, queues with priorities 0 and 1 should have Strict Priority weight (0), and WRR is used for priority queues ranging from 2 to 6. Queues with priorities from 2 – 6 must have weights configured in descending order. Cascade scheduling is a combination of Strict Priority and WRR scheduling, and offers a varying degree of granularity, jitter, and network administrator control.

## Dropping Mechanism

Tail drop and weighted random early detection are algorithms users can use to drop packets.

### Tail Drop

Because queues are not of infinite size, they can fill and overflow. When a queue is full, any additional packets cannot get into the queue and will be dropped until space is available in the queue. This is tail drop. Tail drop can reduce the number of packets that are dropped, particularly if the queue has been allocated a large amount of memory. However, long queues result in the end-to-end delay. Tail drop doesn't discard packets until the queue is 100% full and resources are completely exhausted.

### Weighted Random Early Detection

To prevent issues associated with tail drop, weighted random early detection (WRED) is another dropping mechanism option available. WRED is an extension of random early detection (RED). RED is an active queue management technique deployed in large IP networks. With RED, when a single packet is discarded, a router sends an implicit warning to a source TCP that the discarded packet experienced congestion at some point along the path to the destination. As a response to this warning, the source TCP reduces its transmission rate so the router's queue buffer does not overflow.

RED uses packet drop profile to control the aggressiveness of its packet discard process. The drop profile defines a range of drop probabilities across a range of queue lengths. If a queue length stays below the configured minimum threshold (minth), a packet is never dropped from the queue. If the length exceeds the configured maximum threshold (maxth), a packet is dropped using the tail drop technique. If the queue length remains between the minth and the maxth thresholds, a packet is dropped according to the configured drop probability. Packets are also dropped at random to avoid dropping packets from the same source.

WRED drops packets that are of lower priority before dropping higher-priority packets. Edge routers assign IP Precedence to packets as they enter the network. WRED is useful on any output interface where

you expect to have congestion. WRED uses these precedence values to determine how it treats different types of traffic. WRED tries to make sure that the queue does not fill up, so that there is room for high-priority packets.

Note: The default WRED parameter values are based on the best available data. It is recommended to not change the default parameter values unless you have determined that your applications will benefit from the changed values.

## Policing

Excessive multicast packets can potentially monopolize Route Processor resources, starving other important processes. Policing rates limit the flood of multicast packets to prevent this issue from occurring.

Rate limiting of Internet Control Message Protocol (ICMP) mitigates the effects of certain ICMP based attacks. ICMP packets with TYPE “destination unreachable” are rate limited to mitigate resource crunch and ICMP based attacks.

## DSCP Marking Before Encryption

If a queued packet is destined for an IPSec tunnel, IPSO 4.2 applies the appropriate DSCP before the packet is encrypted. When the firewall encrypts the packet, it copies the DSCP value to the header of the encrypted packet. The DSCP value is therefore visible to the routers, which allows the packet to receive prioritized service as it transits the tunnel.

Because DSCP values must be applied to packets before they are encrypted and packets are encrypted before egress, IPSO 4.2 can apply DSCP marking on ingress. Previous versions of IPSO do not allow packets to be marked on ingress.

## Integration with SecureXL

In IPSO 4.1, when SecureXL is enabled, traffic bypasses the QoS layer completely. Enabling SecureXL completely disables QoS. IPSO 4.2 corrects this problem by first allowing packets to go through the QoS layer for any packet treatment (mark, shape, prioritize) before handing them over to SecureXL for fast forwarding. QoS and SecureXL can co-exist in IPSO 4.2. This is a very important enhancement in IPSO, because Check Point’s Floodgate completely disables SecureXL which results in poor performance.

## Supported Configurations

QoS functions on a per-interface basis which means high and low priority traffic must share the same interface.

**Note:** IPSO 4.2 QoS is not supported on ADP platforms nor supported for Link Aggregation.

The supported interfaces are:

- T1/E1
- Ethernet
- HSSI
- X.21
- V.35

**Note:** Only Ethernet interfaces are supported on newer platforms. ATM, FDDI, Token Ring, and ISDN interfaces are not supported.

## Benefits for Customers

Different applications have varying needs for delay, delay variation (jitter), bandwidth, packet loss, and availability. These parameters form the basis of QoS. Nokia IPSO software is designed to provide the requisite QoS to these applications. Now Nokia IP platforms can participate in an end-to-end DiffServ QoS implementation. The choice of QoS with minimal latency is the most costly in terms of forwarding performance, but it allows the least amount of blocking for high priority traffic. Now with SecureXL integration, customers can prioritize delay-sensitive traffic like VoIP and IPTV without impacting performance too much.

## For More Information

Nokia Inc.  
102 Corporate Park Drive  
White Plains, NY 10604 USA  
[www.nokia.com](http://www.nokia.com)

**Americas**  
Tel: 1 877 997 9199  
Email: [mobile.business.na@nokia.com](mailto:mobile.business.na@nokia.com)

**Asia Pacific**  
Tel: +65 6588 3364  
Email: [mobile.business.apac@nokia.com](mailto:mobile.business.apac@nokia.com)

**Europe, Middle East, and Africa**  
France: +33 170 708 166  
UK: +44 161 601 8908  
Email: [mobile.business.emea@nokia.com](mailto:mobile.business.emea@nokia.com)

### About Nokia

Nokia is the world leader in mobile communications, driving the growth and sustainability of the broader mobility industry. Nokia is dedicated to enhancing people's lives and productivity by providing easy-to-use and secure products like mobile phones, and solutions for imaging, games, media, mobile network operators, and businesses. Nokia is a broadly held company with listings on five major exchanges.

For more information, please visit <http://www.nokia.com/forbusiness>.