Radyne's SkyWire Gateway Quality of Service

White Paper

WP026
Rev 1.0
April 2008

Radyne Corporation • 3138 E. Elwood St. • Phoenix, AZ 85034 • (602) 437-9620 • Fax: (602) 437-4811

Radyne's SkyWire Gateway Quality of Service

Introduction

SkyWire provides the same Quality of Service (QoS) functionality as many commercially available QoS switches costing hundreds of dollars. The four prioritized queues and programmable ingress and egress policies give network administrators the ability to control how the various types of real-time and non real-time services flow through the network. And because all of the QoS determinations and queuing are performed at line speed, there is no additional delay through the network.

The Relationship between IP and Other Protocols

Despite being a layer 2 Ethernet device, SkyWire provides Quality of Service functionality that works at layer 3, the IP level. Processing IEEE tags when they are available and interpreting the QoS information inherent in the IP header when they're not.

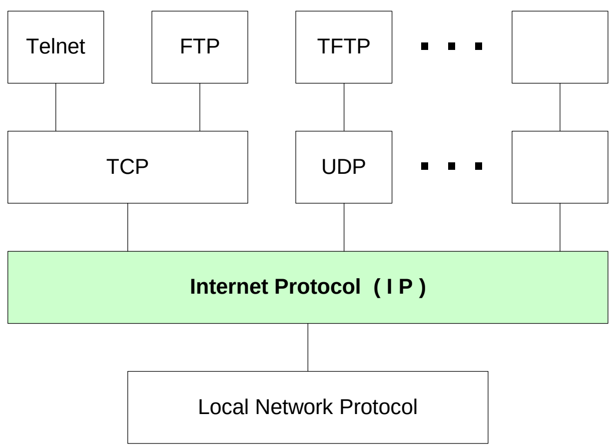

Protocol Relationships

Figure 1.

This IP level functionality is compatible with various methods of implementing QoS including:

- The use of an optional IEEE 802.3ac tag containing IEEE 802.1p priority information
- The Type of Service field or Differentiated Services field contained in an IPv4 header
- The Traffic Class field (RFC 2460) contained in an IPv6 header

By making use of the information inherent in the data stream, SkyWire offers Quality of Service when no other QoS source is available, while at the same time, providing transparent operation that won't interfere when external equipment is used for traffic shaping. The following paragraphs provide additional detail on the various types of QoS information available in the data stream and how SkyWire uses them to route the data to the appropriate queue.

IEEE Tagged Packets

When a packet contains an optional IEEE tag, SkyWire uses the priority information contained in the tag to determine the appropriate queue. The IEEE tag supports 8 priorities with 7 being the highest and 0 being the lowest. Table 1 and Figure 2 illustrate the format of the IEEE tag and the mapping of the various priorities to the four available queues.

IEEE Tag Format

First Octet 0x81		
Second Octet 0x00		
Priority 3 bits	CFI 1 bit	VID (start) 4 bits
VID (continued) 8 bits		

Table 1

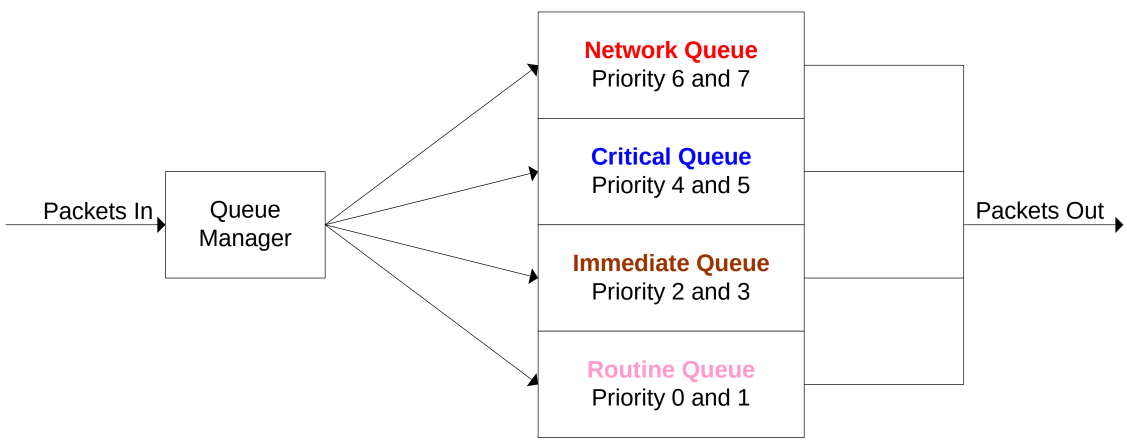

IEEE Tag Prioritized Traffic Flow

Figure 2

IPv4 Packets with a Type of Service field (RFC 791)

When a packet doesn't contain an IEEE tag, and it's an IPv4 packet with a Type of Service field (TOS), SkyWire uses the information contained in the TOS Precedence field to determine the appropriate queue. The Type of Service field is contained in the second byte of the IPv4 header and supports 8 levels of precedence with 7 (Network Control) being the highest and 0 (Routine) being the lowest. Table 2 and Figure 3 illustrate the format of the IPv4 Type of Service information and their mapping to the available queues. Please note, in IPv4, the Type of Service field and Differentiated Services field (discussed in the next section) occupy the same location in the IP header, therefore, only one of these fields will be present in any particular IPv4 header.

IPv4 Type of Service Field

Precedence 3 bits	D 1 bit	T 1 bit	R 1 bit	Reserved 2 bits

R = Reliability
T = Throughput
D = Delay

Precedence
- **111 = Network Control**
- **110 = Internetwork Control**
- **101 = Critical / ECP**
- **100 = Flash Override**
- **011 = Flash**
- **010 = Immediate**
- **001 = Priority**
- **000 = Routine**

Table 2

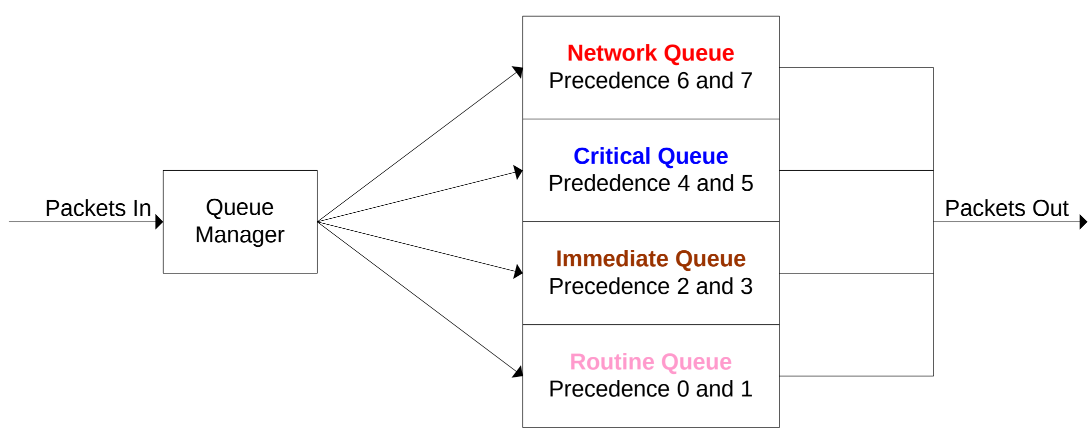

IPv4 Type of Service Traffic Flow

Figure 3

IPv4 Packets with a Differentiated Services field (RFC 2474)

When a packet doesn't contain an IEEE tag, and it's an IPv4 packet with a Differentiated Services field (DS), SkyWire uses the information contained in the Differentiated Services Codepoint field to determine the appropriate queue. The Differentiated Services field is contained in the second byte of the IPv4 header and supports 64 codepoints with 63 being the highest priority and 0 being the lowest. Table 3 and Figure 4 illustrate the format of the IPv4 Differentiated Services information and their mapping to the available queues. Please note, in IPv4 the Type of Service field (discussed in the previous section) and Differentiated Services field occupy the same location in the IP header, therefore, only one of these fields will be present in any particular IPv4 header.

IPv4 Differentiated Services Field

Differentiated Services Codepoint 6 bits	Reserved 2 bits

Table 3

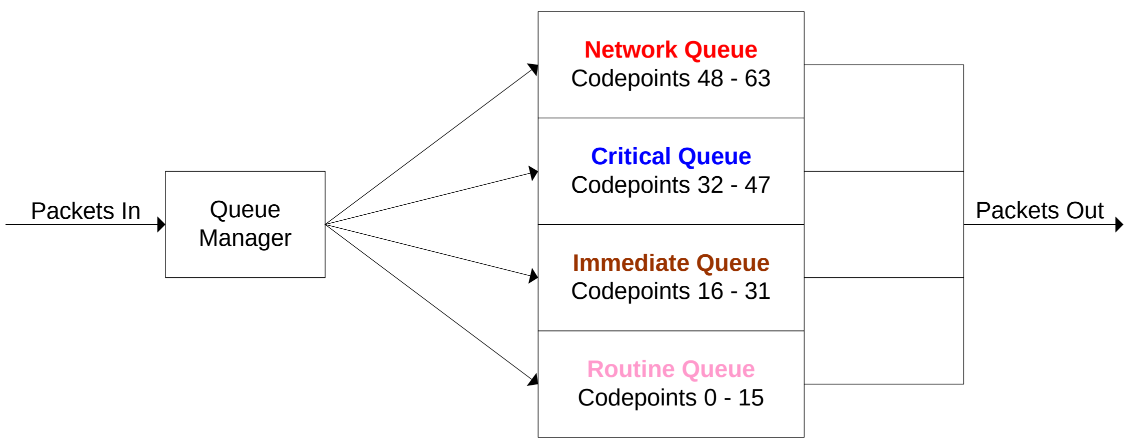

IPv4 Differentiated Services Traffic Flow

Figure 4

IPv6 Traffic Class (RFC 2460)

When a packet doesn't contain an IEEE tag, and it's an IPv6 frame, SkyWire uses the information contained in the Traffic Class field to determine the appropriate queue. The Traffic Class field is contained in the 1st and 2nd bytes of the IPv6 header and supports 64 priority levels with 63 being the highest and 0 being the lowest. Table 4 and Figure 5 illustrate the format of the IPv6 Traffic Class field and the mapping of the various priorities the available queues.

IPv6 Traffic Class Field

IP Version 4 bits		Traffic Class 4 bits
Traffic Class 2 bits	Reserved 2 bits	Flow Label 4 bits

Table 4

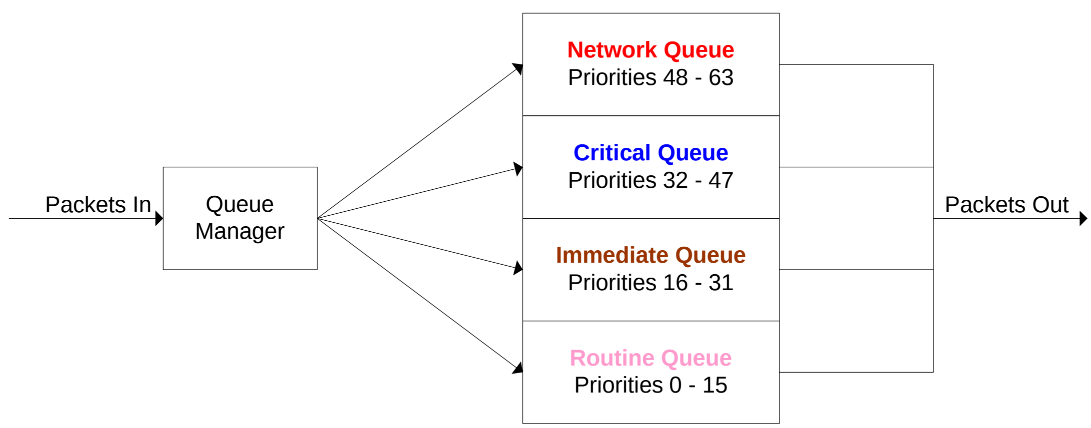

IPv6 Traffic Class Flow

Figure 5

Programmable Ingress Policies

SkyWire provides the network administrator with the following QoS settings that control how the incoming data is routed to the four prioritized queues.

Normal QoS

When Normal QoS is selected, the determination of traffic priority and queuing is based on information contained in the Ethernet packet as follows:

- When a packet contains an IEEE802.3ac Tag (this tag is optional), SkyWire uses the IEEE802.1p priority information embedded in the Tag.
- When a packet does not contain the optional Tag, priority is determined using the
    - IPv4 Type of Service / Differentiated Service field
      or
    - IPv6 Traffic Class field

Port Based QoS

When Port Based QoS is selected, SkyWire ignores and QoS information embedded in the headers and instead priority and queuing is based on the port on which the traffic arrives. Ethernet packets arriving on Data Port J3 have the highest priority and those arriving on Data Port J4 have the lowest. This type of QoS is typically used when the operator wants to control bandwidth allocation based upon the data ports.

Port Based QoS can also be used to create a FIFO like operation where the data is transmitted in the exact order it is received. This can be extremely useful when traffic shaping equipment or crypto devices are used with SkyWire. In these cases, only one of the data ports (J3 or J4) would be utilized.

Programmable Egress Policies

SkyWire provides the network administrator with the following QoS settings that control how the data is routed out of the four prioritized queues and into the modulator.

Fair Weighted Queuing

A fair weighted priority scheme allows higher priority traffic to move quickly through the system, while at the same time insuring that lower priority traffic isn't stalled indefinitely. When Fair Weighted queuing is selected, packets will be transmitted at a rate of 8, 4, 2, and 1 from the highest to lowest priority queues respectively.

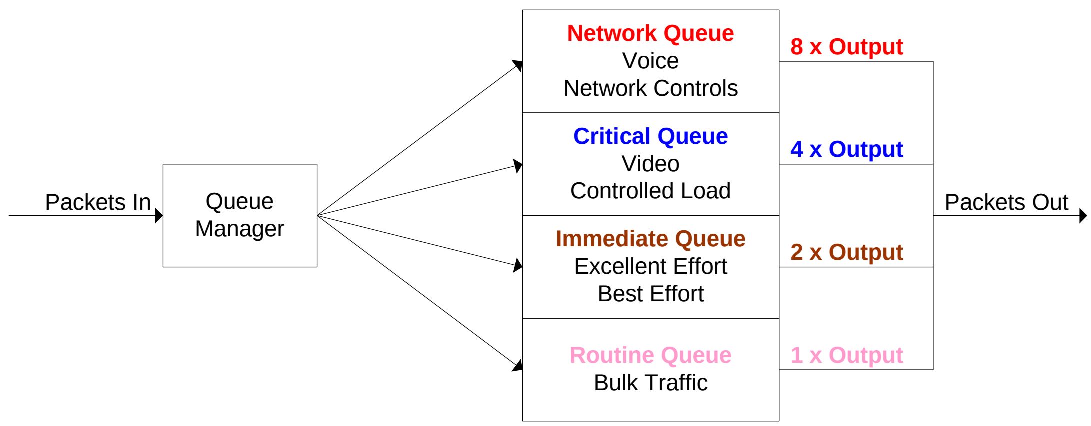

Fair Weighted Traffic Flow

Figure 6

Strict Priority Queuing

A strict priority scheme means that the highest priority traffic will always be transmitted first, even if it means that the lower priority traffic is stalled indefinitely. When Strict Priority queuing is selected, packets in the highest priority queue will be transmitted until that queue is empty. At that time, packets from the next highest priority queue will be transmitted until it is either empty or a higher priority packet arrives. If a higher priority packet arrives before a lower priority queue is empty, it will immediately go to the head of the line and be the next packet transmitted. Care needs to be taken when selecting this option as it is entirely possible that data in a lower priority queue might never be transmitted.

Additional capabilities

In addition to the aforementioned QoS capabilities, the SkyWire Ethernet interfaces provides the following additional functionality.

Automatic Learning and Aging

The Automatic Learning and Aging function allows SkyWire to be connected to a network without requiring extensive setup and programming by the operator. The learning process insures that local LAN traffic stays on the local LAN and that only traffic that is not destined for other equipment on the local LAN gets forwarded over the satellite link.

The aging process allows equipment to be shut off, moved, or even completely removed from the network without tying up resources or requiring operator intervention. For instance, a particular user might carry a laptop from remote office to remote office and use it to check their e-mail back at the corporate office. The apparent movement of the laptop's MAC address from one position in the network to another is automatically handled by SkyWire.

The learned address database maintained by SkyWire can hold up to 1024 MAC addresses. When a new LAN address is learned, it is placed in the database and an aging timer is started. Then, every time this learned address is heard from, its aging timer is reset and restarted. Finally, when the learned address has not been heard from for more than 5 minutes, it is removed from the database. All of the processing required for learning, searching, and aging is handled at line speed and without any unnecessary delay at startup.

SkyWire also performs this automatic learning and aging function on its satellite WAN side. Learning the MAC addresses it receives via the satellite, and only outputting those destined for its LAN and those whose destination is still unlearned. This feature is particularly useful in mesh applications as it prevent unnecessary traffic from flooding every LAN in the mesh.

Satellite Packet Error Checking

Packet error checking is a standard part of any terrestrial Ethernet system and is performed using the CRC contained in the Ethernet wrapper. When a CRC error is encountered, the packet is discarded by the router, switch, hub, or other device in which the error was detected.

SkyWire uses a similar methodology over the satellite link where bit errors result in a corrupted packet with a bad CRC. And for the majority of the situations, the standard process of discarding these packets when they are encountered is the preferred methodology as the packet can simply be retransmitted.

However, there are some cases where the end device would rather receive the erred packet than no packet at all. Scenarios involving cryptography or where the end device has additional error correction capability are examples of two such situations. For these customers, SkyWire provides the ability to turn off the satellite packet error checking and packets with bit errors will be output with a valid CRC so that they will pass through any routers or switches in between the modem and the receiving device.

Daisy Chain capability

For large networks that require more than 4 carriers, the ability to daisy chain multiple SkyWire Gateways together provides a simple upgrade path. As you network grows, an addition SkyWire Gateway with up to four demodulators can be added to the chain without having to buy additional blades for your router or update the routing table to allow for multiple return paths. From your routers point of view, there will still only be a single one wire path to and from all of the destinations. Despite the fact that the outbound satellite path and some of the satellite return paths go through two boxes.

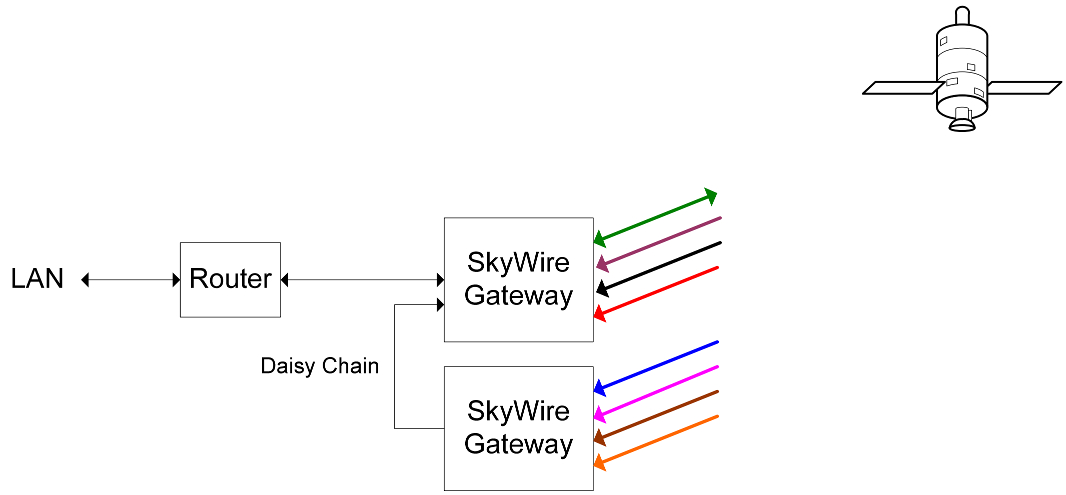

Eight share groups daisy chained on a single router connection

Figure 7

In-band Control

For users who wish to monitor and control the remote end of a link, SkyWire provides an internal connection between the normal data path and the M&C. When this capability is enabled, the Ethernet SNMP commands and responses to the far side ride along with the rest of the data traffic. At the far end they are internally passed to and from the M&C without requiring an external connection between the data port and the M&C port. Security for the control information that is riding along with the user traffic is provided by using SNMP version 3.

For users who never want to mix their control information with user traffic, disabling this feature breaks the internal connection. In this case, Ethernet control of the M&C can only be accomplished by using one of the M&C RJ-45 control ports on the modem.

Internal Buffer and Flow Control Throttle

SkyWire provides 120K bytes of internal data storage for buffering of Ethernet traffic. When this buffer is almost full, if flow control is enabled, SkyWire will attempt to slow down the incoming terrestrial data to prevent the buffer from being overrun. It does so using industry standard backpressure on half duplex links and an IEEE 802.3x Pause Frame on full duplex links. Together the combination of the buffer and flow control can provide an effective rate exchange and throttle between the Ethernet ports and the satellite link.

Auto Everything Ethernet ports

The auto crossover, auto polarity, auto negotiation functionality on the SkyWire Ethernet ports provide true "plug and play" connectivity and eliminates the need for special cabling.

Auto Crossover

SkyWire automatically determines whether or not it needs to cross over between pairs so that external crossover cables are not required. When interoperating with a device that can't automatically adjust for crossover, SkyWire makes the necessary adjustments prior to commencing auto-negotiation. When interoperating with a device that implements MDI/MDIX crossover, a random algorithm as described in IEEE 802.3 section 40.4.4 determines which device performs the crossover.

Auto Polarity

SkyWire automatically detects and corrects for polarity errors on the receive pairs in 1000Base-T and 10Base-T modes. In 100Base-T mode, the polarity does not matter.

Auto Negotiation

SkyWire automatically negotiates with its link partner to determine the speed and duplex at which to operate. If the link partner is unable to auto-negotiate, SkyWire utilizes a parallel detect mode to determine the speed of the link partner and then operates in half-duplex at the appropriate rate.

Transparent Operation

Because SkyWire is completely transparent to the applications that are being run, you don't have to worry about whether or not they will work over SkyWire. The answer is yes. Whether it's a VoIP call, or a DHCP server application, or a VPN to the office to check e-mail, or simply web-surfing, if it's transmitted via Ethernet, it will pass transparently though SkyWire from one point to another.

Adding Acceleration, Compression, Network Security, and Traffic Shaping

Today, there are various acceleration, compression, IP security, and traffic shaping solutions available from numerous third party vendors. Some are proprietary implementations, while others are based on open standards such as MIL-STD-2045-44000 and CCSDS 714.0-B-1. Some may even be implemented in your current network and routers. And while different solutions (including none at all) make sense for different types of networks, SkyWire, with its transparent mode of operation is compatible with all of them.