EXHIBIT C

Comparison of Claims 1, 2, 4, and 16 of Henrion U.S. Patent No. 5,301,192 entitled TEMPORARY INFORMATION STORAGE SYSTEM COMPRISING A BUFFER MEMORY STORING DATA STRUCTURED IN FIXED OR VARIABLE LENGTH DATA BLOCK
With Foundry EdgeIron, ServerIron, NetIron, BigIron, and FastIron products – see Attachment "Model Numbers" for additional specifics

1. Temporary information storage system comprising	Per White Paper: JetCore Based Chassis Systems - 2003 "The new JetCore chipset represents the third generation of ASICs created by Foundry Networks." "Foundry's chassis systems include three series: • FastIron: aimed at high performance layer-2/layer-3 switching for enterprise applications. • BigIron: aimed at high performance layer-2/layer-3 switching for data centers as well as large campus networks. • NetIron: aimed at high performance layer-2/layer-3 and MPLS switching for service providers, and wide area network backbones. All of the chassis series above are capable of leveraging the new JetCore architecture …." At FN A 00163049. "[T]he systems place the intelligent packet forwarding functions on the interface modules, i.e., all the interface modules incorporate their own switching engines, buffers, queues, and switching fabrics." At FN A 00163051 "The "BigIron Architecture Brief" white paper describes the high level architecture of Foundry's interface modules as per the 2nd generation chipset. The 3rd generation chipset, described here, consolidates all the required functions as well as some new functions into fewer ASICs to provide higher port density, reliability, cost savings, and added functionality." At FN A 00163051 Note: The "BigIron Architecture Technical Brief – 1999 FN A_00045748 describes the prior system which shares the same architecture and methods (switching engines, buffers, queues, and switching fabrics, etc). Similar quotes can be found there.
a buffer memory storing data structured as fixed or variable length blocks of data, each of said data blocks comprising one or more data elements, and	"The shared memory buffer pool provides storage for the received packet(s) till it is time to the send the packet out …." At FN A 00163059

Comparison of Claims 1, 2, 4, and 16 of Henrion U.S. Patent No. 5,301,192 entitled TEMPORARY INFORMATION STORAGE SYSTEM COMPRISING A BUFFER MEMORY STORING DATA STRUCTURED IN FIXED OR VARIABLE LENGTH DATA BLOCK
With Foundry EdgeIron, ServerIron, NetIron, BigIron, and FastIron products – see Attachment "Model Numbers" for additional specifics

a data block and outlet selection queue management logic comprising	"The buffer manager is responsible for managing the packet storage space inside the shared memory." Id. "Under the control of the buffer manager, the shared memory switch fabric, also, retrieves packets from the shared memory and sends them to their destination ports, as instructed by the buffer manager." At FN A 00163059
a plurality of queues for interlinking each of the data blocks to a respective destination of a plurality of destinations, such that when a data block is assigned to a destination it is sufficient to write its identity in the respective queue specific to this destination, and	"When the buffer manager is notified of the buffer number of a new packet in the shared memory, the buffer manager assigns the buffer number to the appropriate queue based on the packet's destination and priority." Id. "Foundry's architecture offers 4 queues per destination to allow up to 4 levels of traffic prioritization. The architecture honors those priority levels, and schedules service for the prioritized traffic flows through the use of: • 4 output queues per physical port within the same port group • 4 output queues per neighboring port group • 4 output queues per neighboring interface module" At FN A 00163062 and FN A 00163063 Note: In the Foundry device, a data block's identity is its buffer number
data block selection means for assigning each of said data blocks to a respective individual outlet of a respective plurality of outlets associated with said respective destination, said data block selection means further comprising	"When a packet reaches the head of its queue, the buffer manager instructs the shared memory switch fabric as to where to retrieve the packet from the shared memory (using the buffer number), and which physical port to send it to." At FN A 00163059 Note: in the Foundry device, a data block's destination is its destination port.

Comparison of Claims 1, 2, 4, and 16 of Henrion U.S. Patent No. 5,301,192 entitled
TEMPORARY INFORMATION STORAGE SYSTEM COMPRISING A BUFFER MEMORY STORING DATA STRUCTURED IN FIXED OR VARIABLE LENGTH DATA BLOCK
With Foundry EdgeIron, ServerIron, NetIron, BigIron, and FastIron products – see Attachment "Model Numbers" for additional specifics

first means for identifying a specific said queue associated with said respective destination with which said respective individual outlet is associated,	"The fabric … reports the following information to the buffer manager: Buffer number - where the packet is stored Forwarding Identifier (FID) …." "The buffer manager uses the above information to determine the location of the packet in the shared memory, … the destination(s) of the packet, and consequently, which queue(s) it … [is] placed on." At FN A 00163058 Note: in the Foundry device, the CAM & PRAM translate a destination port which implies a specific outlet and its associated queue. "The content addressable memory (CAM) is used by the packet classifier to provide reverse lookups based on the content of the packet. The packet classifier extracts a portion(s) of the packet headers - such as destination IP address, destination MAC address, source and destination IP addresses, etc. - and hands it over to the high speed CAM to perform a lookup. In case a match is found, the CAM returns an index of a location in the PRAM, where information on how to forward the packet is stored. The PRAM is a high speed memory that takes the index returned by the CAM lookup and returns to the packet classifier information on how to forward the packet: FID (destination port or ports), MAC address of next hop, VLAN ID that should be attached to the packet as it goes out a given interface, etc. The information retrieved by the PRAM is used to build the internal Foundry header that gets attached to the packet. This information is used by other modules of the hardware to determine how to forward the packet." At FN A 00163060
second means for obtaining from said specific queue an address of said each data block and	"When a packet reaches the head of its queue, the buffer manager instructs the shared memory switch fabric as to where to retrieve the packet from the shared memory (using the buffer number),…" At FN A 00163059

Comparison of Claims 1, 2, 4, and 16 of Henrion U.S. Patent No. 5,301,192 entitled
TEMPORARY INFORMATION STORAGE SYSTEM COMPRISING A BUFFER MEMORY STORING DATA STRUCTURED IN FIXED OR VARIABLE LENGTH DATA BLOCK
With Foundry EdgeIron, ServerIron, NetIron, BigIron, and FastIron products – see Attachment "Model Numbers" for additional specifics

third means for transferring said each data block to said respective individual outlet.	"Under the control of the buffer manager, the shared memory switch fabric, also, retrieves packets from the shared memory and sends them to their destination ports, as instructed by the buffer manager." At FN A 00163059 Note: The individual outlet is the hardware associated with the chosen queue.
2. Temporary information storage system in accordance with claim 1, wherein	Per White Paper: JetCore Based Chassis Systems - At FN A 00163047 and subsequent pages: "Foundry chassis switching routers deliver high performance and high port density suitable for several applications … , the systems place the intelligent packet forwarding functions on the interface modules, i.e., all the interface modules incorporate their own switching engines, buffers, queues, and switching …. This leaves the chassis' switch fabric with the simple job of non-blocking forwarding of packets between the individual interface modules. A packet classifier operating at the full line rate is responsible for parsing and examining the incoming packets at layers 2, 3, and 4, in order to make a decision on how the packet should be forwarded. … The packet classifier's processing results in a forwarding identifier (FID) that describes the destination port(s) for the packet processed. … it attaches an internal Foundry header to the packet. This is an optimized header that describes to the hardware how the packet should be handled from that point onwards. The FID is included as part of that header. The shared memory switch fabric is responsible for transferring the packets to and from the high speed shared memory. … When the shared memory switch fabric receives the packet from the packet classifier, it transfers the packet to a free location in the shared memory. The fabric then reports … to the buffer manager: … The buffer manager uses the above information to determine … , which queue(s) it should be placed on. …. When a packet reaches the head of its queue, the buffer manager instructs the shared memory switch fabric as to where to retrieve the packet from the shared memory (using the buffer number), and which physical port to send it to. "

Comparison of Claims 1, 2, 4, and 16 of Henrion U.S. Patent No. 5,301,192 entitled TEMPORARY INFORMATION STORAGE SYSTEM COMPRISING A BUFFER MEMORY STORING DATA STRUCTURED IN FIXED OR VARIABLE LENGTH DATA BLOCK
With Foundry EdgeIron, ServerIron, NetIron, BigIron, and FastIron products – see Attachment "Model Numbers" for additional specifics

said respective individual outlet is associated with a plurality of said destinations, and	"The buffer manager, responsible for managing the priority queues, assigns the buffer number of the packet to the appropriate queue(s) of the destination port(s). When the buffer number gets to the front of its queue, the buffer manager instructs the shared memory switch fabric to retrieve the packet from the shared memory and send it to the destination port. The destination(s) of the packet could be a port on the same module (within the same port group or not), a port on another module, a mirror port, or the management module. The buffer manager then deletes the packet from the shared memory." At FN A 00163062
said data block selection means further comprises destination selection means for identifying a selected destination from said plurality of said destinations, said destination selection means further comprising fourth means for identifying all associated destinations associated with said respective individual outlet having at least one data block to transfer, fifth means responsive to said fourth means for identifying from said all associated destinations, a selected outgoing destination having at least one data block to transfer and sixth means responsive to said fifth means for establishing the identity of the queue associated with said selected destination.	"BY MR. GREENSPOON: Q What is the mechanism in JetCore for sending a packet to multiple ports on a different module? THE WITNESS: So what we do is that the destination will include multiple ports in the port mask Q By port mask you're talking about a certain portion of the FID? A So the FID basically is a index to an area which is representative of the ports and that's what I mean by port mask. … Q So the FID is an index into a table like a hardware table which then would output the port mask? A It's an index into a into the memory which will give out the port mask … Q So how many bits wide is the FID? A FID is 16 bit quantity. … Q So instead the FID itself is used as an index and then the output of the index does or contains the port map? A Yes, FID is an index to aid in the memory which will give more details of the ports that it's going to be sent out on. Q What is the name of that area in memory? A It's called FID RAM." – Deposition of Nitin Jain, 1/23/2006

Comparison of Claims 1, 2, 4, and 16 of Henrion U.S. Patent No. 5,301,192 entitled
TEMPORARY INFORMATION STORAGE SYSTEM COMPRISING A BUFFER MEMORY STORING DATA STRUCTURED IN FIXED OR VARIABLE LENGTH DATA BLOCK
With Foundry EdgeIron, ServerIron, NetIron, BigIron, and FastIron products – see Attachment "Model Numbers" for additional specifics

4. Temporary information storage system in accordance with claim 1, wherein said destination selection means further comprise	Per White Paper: JetCore Based Chassis Systems - At FN A 00163047 and subsequent pages: "Foundry chassis switching routers deliver high performance and high port density suitable for several applications … , the systems place the intelligent packet forwarding functions on the interface modules, i.e., all the interface modules incorporate their own switching engines, buffers, queues, and switching …. This leaves the chassis' switch fabric with the simple job of non-blocking forwarding of packets between the individual interface modules. A packet classifier operating at the full line rate is responsible for parsing and examining the incoming packets at layers 2, 3, and 4, in order to make a decision on how the packet should be forwarded. … The packet classifier's processing results in a forwarding identifier (FID) that describes the destination port(s) for the packet processed. … it attaches an internal Foundry header to the packet. This is an optimized header that describes to the hardware how the packet should be handled from that point onwards. The FID is included as part of that header. The shared memory switch fabric is responsible for transferring the packets to and from the high speed shared memory. … When the shared memory switch fabric receives the packet from the packet classifier, it transfers the packet to a free location in the shared memory. The fabric then reports … to the buffer manager: … The buffer manager uses the above information to determine … , which queue(s) it should be placed on. …. When a packet reaches the head of its queue, the buffer manager instructs the shared memory switch fabric as to where to retrieve the packet from the shared memory (using the buffer number), and which physical port to send it to. "

Comparison of Claims 1, 2, 4, and 16 of Henrion U.S. Patent No. 5,301,192 entitled
TEMPORARY INFORMATION STORAGE SYSTEM COMPRISING A BUFFER MEMORY STORING DATA STRUCTURED IN FIXED OR VARIABLE LENGTH DATA BLOCK
With Foundry EdgeIron, ServerIron, NetIron, BigIron, and FastIron products – see Attachment "Model Numbers" for additional specifics

a counter for each destination, giving the number of data blocks to be transferred for said destination, and detection means for specifying whether said number is zero or different from zero.	"Fixed Rate Limiting counts the number of Ethernet packet bytes that a port either sends or receives, in one second intervals. The direction that the software monitors depends on the direction you specify when you configure the rate limit on the port. If the number of bytes exceeds the maximum number you specify when you configure the rate, the port drops all further packets for the rate-limited direction, for the duration of the one- second interval." At FN A 00055505 Note: many Foundry products support rate limiting – the following is only a sample from a Release Notes published 9/2000: "Rate Limiting Support on Foundry Products NetIron400 and 800 BigIron 4000 and 8000 TurboIron/8 Layer 2 Switch TurboIron/8 Layer 3 Switch" At FN A 00055504
16. Temporary information storage system in accordance with claim 1, wherein said queues are FIFO queues.	Note – see introduction portion of quotes in Claims 1,2,and 4. Note: Foundry has used FIFO queues from the very beginning, the following excerpt is from a 1997 Functional Specification from which Foundry switches were born: " Fifo Descriptions Xmt Fifo: The Xmt fifo provides the software with a single managed interface for … Rcv Fifo: The rcv fifo is a single interface for fetching the pointer and header data for each frame which has been completely received." At FN A 00040733 Note: many other documents refer to the use of FIFOs as well. All of Foundry's queues are FIFO.

EXHIBIT D

Comparison of Claim 8 & 9 of Pauwels, U.S. Patent No. 5,506,840 entitled ASYNCHRONOUS SWITCHING NODE AND ROUTING LOGIC MEANS FOR A SWITCHING ELEMENT USED THEREIN
with all Foundry L2 Devices including AccessIron, EdgeIron, FastIron Edge, FastIron Edge X Series, BigIron, BigIron MG8, BigIron RX Series – See attachment "Model Numbers" for additional detail

8. Routing logic means for a switching element (S11, ..., S3N) having at least one inlet and a plurality of outgoing routes each comprising at least one outlet, said routing logic means (MC1; MC2; RL) selecting for cells received at said at least one inlet at least one of said plurality of outgoing routes as a selected route based on an internal routing label associated with said cells and routing said cells to said at least one outlet of said selected route,	Foundry document Foundry Switch and Router Installation and Basic Configuration Guide (2004) ("Configuration Guide"), "Foundry devices support ... 802.1p/q VLAN Tagging." C-1. Per IEEE 802.1Q (1998), "In port-based VLAN classification within a Bridge, the VID associated with an untagged or priority-tagged frame (i.e. a frame with no tag header, or a frame with a tag header that contains the null VLAN ID), is determined, based on the Port of arrival of the frame into the Bridge ...." p.31. Per Foundry document Switch and Router Command Line Interface Reference (2004) ("CLI Reference"), "When you enable hardware flooding for a VLAN, unknown unicast packets are flooded by hardware to all ports in the VLAN." At FN A 00080263. "The PRAM is a high speed memory that takes the index returned by the CAM lookup and returns to the packet classifier information on how to forward the packet: FID (destination port or ports), MAC address of next hop, VLAN ID that should be attached to the packet as it goes out a given interface, etc. The information retrieved by the PRAM is used to build the internal Foundry header that gets attached to the packet. This information is used by other modules of the hardware to determine how to forward the packet. From White Paper "JetCore Based Chassis Systems – An Architecture Brief "... At FN A 00163060 Note: the internal routing label is the internal Foundry header as mentioned above.
wherein said routing logic means (MC1; MC2; RL) includes preselection means (MB; M; RM) which preselect under control of said internal routing label, at least one preselected set having at least one of said plurality of outgoing routes	"A Layer 2 entry is made in the Foundry device's CAM that matches all packets for that VLAN. Any traffic for the VLAN that doesn't match a CAM entry for an individual destination matches this CAM entry." At FN A 00080263 Per Foundry document Foundry Security Guide (2004), "CAM entries contain Layer 2, Layer 3, or Layer 4 next-hop information .... Traffic that matches the match-all CAM entry is hardware-flooded to all ports in the VLAN ...." At FN A 00133095 Note: The internal Foundry header therefore allows pre-selection for such a frame of all outgoing ports in the inlet port's VLAN.

Comparison of Claim 8 & 9 of Pauwels, U.S. Patent No. 5,506,840 entitled
ASYNCHRONOUS SWITCHING NODE AND ROUTING LOGIC MEANS FOR A SWITCHING ELEMENT USED THEREIN
with all Foundry L2 Devices including AccessIron, EdgeIron, FastIron Edge, FastIron Edge X Series, BigIron, BigIron MG8, BigIron RX Series – See attachment "Model Numbers" for additional detail

and includes selection means (MX; MX'; SEL) to select said selected route from said at least one preselected set.	Per Configuration Guide, "A trunk group is a set of physical ports that are configured to act as a single physical interface. Each trunk group's port configuration is based on the configuration of the lead port, which is the lowest numbered port in the group. If you add a trunk group's lead port to a VLAN, all of the ports in the trunk group become members of that VLAN." At FN A 00135747 . "For switch trunk groups, load sharing occurs as follows: …. To select a port, the device determines the port in the trunk group that has the fewest number of flows. Traffic for the destination address is sent over the selected port." At FN A 00135639. Note: When a trunk group is part of a flooded VLAN, the device selects the outgoing port from among its pre-selected member ports.

Comparison of Claim 8 & 9 of Pauwels, U.S. Patent No. 5,506,840 entitled ASYNCHRONOUS SWITCHING NODE AND ROUTING LOGIC MEANS FOR A SWITCHING ELEMENT USED THEREIN
with all Foundry L2 Devices including AccessIron, EdgeIron, FastIron Edge, FastIron Edge X Series, BigIron, BigIron MG8, BigIron RX Series – See attachment "Model Numbers" for additional detail

9. Routing logic means (RL) according to claim 8, wherein said preselection means is a routing memory (RM) the memory locations (MSK) of which are addressed by a route reference number included in said internal routing label, each of said memory locations identifying routes of a distinct set.	The PRAM is a high speed memory that takes the index returned by the CAM lookup and returns to the packet classifier information on how to forward the packet: FID (destination port or ports), MAC address of next hop, VLAN ID that should be attached to the packet as it goes out a given interface, etc. The information retrieved by the PRAM is used to build the internal Foundry header that gets attached to the packet. This information is used by other modules of the hardware to determine how to forward the packet. From White Paper "JetCore Based Chassis Systems – An Architecture Brief "...At FN A 00163060 "BY MR. GREENSPOON: Q What is the mechanism in JetCore for sending a packet to multiple ports on a different module? THE WITNESS: So what we do is that the destination will include multiple ports in the port mask Q By port mask you're talking about a certain portion of the FID? A So the FID basically is a index to an area which is representative of the ports and that's what I mean by port mask. ... Q So the FID is an index into a table like a hardware table which then would output the port mask? A It's an index into a into the memory which will give out the port mask ... Q So how many bits wide is the FID? A FID is 16 bit quantity. ... Q So instead the FID itself is used as an index and then the output of the index does or contains the port map? A Yes, FID is an index to aid in the memory which will give more details of the ports that it's going to be sent out on. Q What is the name of that area in memory? A It's called FID RAM." – Deposition of Nitin Jain, 1/23/2006 Note: The FID RAM is the routing memory. The FID is the route reference number.

EXHIBIT E

REDACTED IN ITS ENTIRETY

EXHIBIT F

Foundry NetIron XMR/MLX Configuration Guide



2100 Gold Street
P.O. Box 649100
San Jose, CA 95164-9100
Tel 408.586.1700
Fax 408.586.1900

February 2006

Configuring an Interface as the Source for All TACACS/TACACS+ Packets 4-8
Configuring an Interface as the Source for All RADIUS Packets 4-8
Specifying a Simple Network Time Protocol (SNTP) Server 4-8
Setting the System Clock 4-10
Limiting Broadcast, Multicast, or Unknown-Unicast Rates 4-11
Limiting Broadcasts 4-11
Limiting Multicasts 4-11
Limiting Unknown Unicasts 4-12
Configuring CLI Banners 4-12
Setting a Message of the Day Banner 4-12
Setting a Privileged EXEC CLI Level Banner 4-13
Displaying a Message on the Console When an Incoming Telnet Session Is Detected 4-13
Configuring Terminal Display 4-13
Checking the Length of Terminal Displays 4-14
Enabling or Disabling Routing Protocols 4-14
Displaying and Modifying System Parameter Default Settings 4-15
Enabling or Disabling Layer 2 Switching 4-17
Changing the MAC Age Time 4-18
Configuring Static MAC Addresses 4-18
Configuring Static ARP Entries 4-19

Chapter 5
Configuring Interface Parameters 5-1
Assigning a Port Name 5-1
Assigning an IP Address to a Port 5-2
Modifying Port Speed 5-2
Modifying Port Mode 5-3
Disabling or Re-Enabling a Port 5-3
Changing the 802.3x Gigabit Negotiation Mode 5-3
Changing the Default Gigabit Negotiation Mode 5-4
Changing the Negotiation Mode 5-4
Disabling or Re-Enabling Flow Control 5-5
Specifying Threshold Values for Flow Control 5-5
Locking a Port to Restrict Addresses 5-5
Modifying Port Priority (QoS) 5-5
Assigning a Mirror Port and Monitor Ports 5-6
Configuration Guidelines for Monitoring Traffic 5-6
Syntax for Port Mirroring and Monitoring 5-6
Monitoring an Individual Trunk Port 5-6
Displaying Mirror and Monitor Port Configuration 5-7
Enabling WAN PHY Mode Support 5-7

Chapter 6
Configuring Trunk Groups 6-1
Trunk Group Connectivity to a Server 6-2
Trunk Group Rules 6-2

 © 2005 Foundry Networks, Inc.

Disabling or Re-Enabling Flow Control

You can configure full-duplex ports on a system to operate with or without flow control (802.3x). Flow control is enabled by default.

To disable flow control on full-duplex ports on a system, enter the following:

```
NetIron(config)# no flow-control
```

To turn the feature back on:

```
NetIron(config)# flow-control
```

Syntax: flow-control

Specifying Threshold Values for Flow Control

The 802.3x flow control specification provides a method for slowing traffic from a sender when a port is receiving more traffic than it can handle. Specifically, the receiving device can send out 802.3x PAUSE frames that request that the sender stop sending traffic for a period of time.

The NetIron XMR/MLX generates 802.3x PAUSE frames when the number of buffers available to a module's Buffer Manager (BM) drops below a threshold value. A module's BM can start running out of buffers when a port receives more traffic than it can handle. In addition, the device drops the lowest priority traffic when the number of available buffers drops below a second threshold. When the number of available buffers returns to a higher level, the device sends out another PAUSE frame that tells the sender to resume sending traffic normally. You can specify values for both thresholds, as well as the module where the thresholds are to take effect.

NOTE: To use this feature, 802.3x flow control must be enabled globally on the device. By default, 802.3x flow control is enabled on the NetIron XMR/MLX, but can be disabled with the **no flow-control** command.

To specify threshold values for flow control, enter the following command:

```
NetIron(config)# qd-flow sink 75 sunk 50 slot 1
```

Syntax: qd-flow sink <sinking-threshold> sunk <sunk-threshold> slot <slot>

The threshold values are percentages of the total number of buffers available to a module's Buffer Manager.

When the <sinking-threshold> is reached, the NetIron XMR/MLX sends out 802.3x PAUSE frames telling the sender to stop sending traffic for a period of time.

When the <sunk-threshold> is reached, the NetIron XMR/MLX drops traffic at the specified priority level.

The <slot> parameter specifies the location of the module where the thresholds are to take effect.

Locking a Port to Restrict Addresses

Address-lock filters allow you to limit the number of devices that have access to a specific port. Access violations are reported as SNMP traps. By default this feature is disabled. A maximum of 2048 entries can be specified for access. The default address count is eight.

EXAMPLE:

To enable address locking for port 2/1 and place a limit of 15 entries:

```
NetIron(config)# lock e 2/1 addr 15
```

Syntax: lock-address ethernet <portnum> [addr-count <num>]

Modifying Port Priority (QoS)

You can give preference to the inbound traffic on specific ports by changing the Quality of Service (QoS) level on those ports. For information and procedures, see "Configuring Quality of Service" on page 16-1.

 © 2006 Foundry Networks, Inc.

- You cannot enable this feature on the designated management VLAN for the device.
- If you enable this feature on a VLAN that includes a trunk group, hardware flooding for Layer 2 multicast and broadcast packets occurs only on the trunk group's primary port. Multicast and broadcast traffic for the other ports in the trunk group is handled by software.

Unicast Flooding on VLAN Ports

Unknown unicast packets do not have a specific (or unicast) recipient. In order for these "special" packets to reach their intended recipient, they needed to be sent on all ports of the VLAN (or "flooded" across the VLAN).

You must enable *hardware* flooding for unknown unicast packets on the NetIron XMR/MLX. It is disabled by default.

To enable unicast hardware flooding on a VLAN ports and enable software flooding, enter commands such as the following:

```
NetIron(config)# vlan 2
NetIron(config-vlan-2)# unknown-unicast-flooding
NetIron(config-vlan-2)# exit
NetIron(config)# reload
```

Syntax: [no] unknown-unicast-flooding

VLAN CPU Protection

This feature protects the line-card's CPU from being overwhelmed by excessive L2 packets that would require the CPU's attention, including known unicast, multicast packets and packets requiring source-MAC learning. Once this feature is enabled, all L2 multicast traffic will be hardware flooded. Furthermore, when the CPU is too busy, this feature will hardware flood unknown unicast traffic, as well as reduce the rate of source-MAC learning traffic to the line-card's CPU, so that the line-card's CPU will have enough resources to handle other types of packets.

Configuring VLAN CPU Protection

VLAN CPU protection is enabled per VLAN. To enable VLAN CPU protection on a VLAN, enter the following command:

```
NetIron(config)# vlan 24
NetIron(config-vlan-24)# vlan-cpu-protection
```

Syntax: [no] vlan-cpu-protection

By default, a maximum of 8 VLANs can have this feature enabled. If you want to increase the number, you must use the following command:

```
NetIron(config)# system-max hw-flooding 40
```

Syntax: [no] system-max hw-flooding <number>

The <number> variable is the number of VLANs that you want to support with Layer-2 CPU protection.

NOTE: You must reload your system for the **system-max** command to take effect.

Displaying VLAN Information

After you configure the VLANs, you can view and verify the configuration.

© 2006 Foundry Networks, Inc.

- Layer 2 CoS – The 802.1p priority in the Ethernet frame. The priority is a value from 0 – 7. The 802.1p priority is also called the ***Class of Service (CoS)***.
- Layer 3 IP Precedence – The value in the three most significant bits of the IP packet header's 8-bit ToS field. The IP Precedence is a value from 0 – 7. The IP Precedence values are described in RFC 791.
- Layer 3 DSCP – The value in the six most significant bits of the IP packet header's 8-bit ToS field. The DSCP is the six most significant bits in the IP packet header's ToS field. The DSCP value is sometimes called the ***DiffServ*** value, and can be from 0 – 63. The DSCP values are described in RFCs 2474 and 2475.

Marking

Marking is the process of changing the packet's QoS information for the next hop. For example, for traffic coming from a device that does not support DiffServ, you can change the packet's IP Precedence value into a DSCP value before forwarding the packet.

You can mark a packet's Layer 2 CoS value, its Layer 3 DSCP value, or both values. The Layer 2 CoS or DSCP value the device marks in the packet is the same value that results from mapping the packet's QoS value into a Layer 2 CoS or DSCP value.

Marking is optional and is disabled by default. When marking is disabled, the device still performs the mappings listed in "Classification" for scheduling the packet, but leaves the packet's QoS values unchanged when the device forwards the packet.

Scheduling

Scheduling is the process of mapping a packet to an internal forwarding queue based on its QoS information, and placing the packet into one of the eight hardware forwarding queues for forwarding. The Foundry device maps the packet's QoS value into a CoS or DSCP value, then maps that value to an internal forwarding queue. The device then places the packet in the hardware forwarding queue corresponding to the internal forwarding queue.

You can modify the scheduling by changing the following mappings:

- CoS –> DSCP
- IP Precedence –> DSCP
- DSCP –> DSCP
- DSCP –> internal forwarding priority

The first three mappings are the same ones described in "Classification" and are applicable for DSCP marking. The trusted QoS source (CoS, IP Precedence, or DSCP) is mapped to a DSCP value and the packet is marked with that DSCP value.

The DSCP –> internal forwarding priority mapping is used to translate the results of the first three mappings into a value that the Foundry device can use to select a hardware forwarding queue. In addition, if the outgoing interface is an 802.1q tagged interface, the DSCP value is mapped to an 802.1p value and the packet is marked with the 802.1p value.

Default QoS Mappings

The Foundry device can map an incoming packet's CoS (802.1p), IP Precedence, or DSCP value into a DSCP value, and can map the DSCP value to an internal forwarding priority. The mappings are used for forwarding the packet to an output queue within the Foundry device and for marking the packet.

The following tables list the default mappings. You can change the mappings if needed. See "Changing the QoS Mappings" on page 16-7.

 © 2006 Foundry Networks, Inc.