EXHIBIT G

Release Notes

EdgeIron 24GS, 48GS, 24G-A, 48G and 8X10G

Releases 2.2.4.21



Release Date: November 2004

Publication Date: November 23, 2004

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

Feature Support

Table 5 Shows the features supported on EdgeIron devices.

Table 5: Features Supported in Release 2.2.4.21

Features		EdgeIron 24G-A EdgeIron 48G	EdgeIron 24GS EdgeIron 48GS	EdgeIron 8X10G
System	Software	Version 2.2.4.21	Version 2.2.4.21	Version 2.2.4.21
	Size	1U	1U	1U
	Uplink	mGBIC/TX Combo	XFP	XFP
	Redundant Power	External	External	External
	Jumbo Frames	Yes (up to 9 Kbytes)	Yes (up to 9 Kbytes)	Yes (up to 9 Kbytes)
	Flash Memory	8 MBytes	16 MBytes	32 MBytes
	DRAM	32 MBytes	64 MBytes	128 MBytes
	Buffer Architecture	24G-A - 2 Mbytes (per system)	24GS - 2 Mybtes (per system)	8X10G - 4 Mybtes (per system)
		48G - 2 Mybtes (per system)	48GS - 2 Mybtes (per system)	
	RADIUS and TACACS+ authentication	Yes	Yes	Yes
	BootP	Yes	Yes	Yes
	DHCP Clients	Yes	Yes	Yes
	IP Address Management	Yes	Yes	Yes
	Configuration Management	Yes	Yes	Yes
	TFTP	Yes	Yes	Yes
	Remote Logging	Yes	Yes	Yes
	HTTPS	Yes	Yes	Yes
	SSH (version 2.0)	Yes	Yes	Yes
	SSH (version 1.5)	Yes	Yes	Yes
	SNMP (v1, v2c and v3)	Yes	Yes	Yes
	SNTP	Yes	Yes	Yes

© 2004 Foundry Networks, Inc.

Table 5: Features Supported in Release 2.2.4.21

Features		EdgeIron 24G-A EdgeIron 48G	EdgeIron 24GS EdgeIron 48GS	EdgeIron 8X10G
MIB Support	MIB II	Yes	Yes	Yes
	Bridging MIB	Yes	Yes	Yes
	Interface MIB	Yes	Yes	Yes
	RMON MIB	Yes	Yes	Yes
	RADIUS MIB	Yes	Yes	Yes
	Ethernet-like MIB	Yes	Yes	Yes
	Entity MIB	Yes	Yes	Yes
	Extended MIB	Yes	Yes	Yes
	SNMP MIB	Yes	Yes	Yes
Layer 2	IEEE 802.1q Port-based VLANs	Yes	Yes	Yes
	VLANs per switch	255	255	255
	Private VLANs	Yes	Yes	Yes
	IEEE 802.1v Protocol VLANs	Yes	Yes	Yes
	IEEE 802.1D (STP)	Yes	Yes	Yes
	Fast Link Option	Yes	Yes	Yes
	IEEE 802.1s (MSTP)	Yes (up to 32 instances)	Yes (up to 32 instances)	Yes (up to 32 instances)
	IEEE 802.1w (RSTP)	Yes	Yes	Yes
	Foundry Per VLAN STP compatible	Yes (up to 32 instances)	Yes (up to 32 instances)	Yes (up to 32 instances)
	Dual-mode*	Yes	Yes	Yes
	GVRP	Yes	Yes	Yes
	Priority (IEEE 802.1p & IEEE 802.1q)	Yes	Yes	Yes
	Priority queues per port	8	8	8
	IP Precedence, IP DSCP, IP Port	Yes	Yes	Yes
	Strict scheduling for priority queue	Yes	Yes	Yes

 © 2004 Foundry Networks, Inc.

Table 13: Known Issues in Release 2.2.4.21

EIF24G-A, EIF48G, EIF8X10G	Flow control does not work across switch chips.	EIF4802CF-00354, EIF48G-00296
EIF24G-A, EIF48G	VLANs are not configurable for trunk groups or trunk members; however, VLANs are configurable for ports. To configure a VLAN for a trunk you must first configure the VLAN for the specified ports and then add the ports to the trunk.	EIF24GA-00177
EIF48G	There will be packet loss when running 10 pairs of gigabit speed port pair traffic. Using 9 pairs is OK. At non-wire speed running 10 pairs generates 1% packet loss.	EIF48G-00164
EIF48G	Failed to restart the switch after two days of continuous running.	EIF48G-00211
EIF48G	Spanning tree is only configurable at a global level. Per VLAN settings do not work.	EIF24GA-00177
EIF24G-A, EIF48G	Due to a hardware limitation the maximum number of VLANs created is limited to 32. Creating more than 32 VLANs at PVST enabled ports results in more BPDUs than the CPU can handle.	EIF48G-00216

© 2004 Foundry Networks, Inc.

EXHIBIT H

FOUNDRY NETWORKS

BigIron RX architecture brief

The BigIron® RX Series is the industry's most powerful family of Layer 2/Layer 3 Ethernet switches. IP networks now play a critical role in ensuring satisfactory end-user experience, irrespective of the application being run on the network. It is therefore imperative to ensure non-stop operation, high performance, low latency and QoS guarantees for the respective applications that are concurrently running on the network. Further, the ability to adapt to evolving technologies such as 40-Gig or 100-Gig interfaces in the not too distant future has imposed a need for administrators to install future-proof systems. The BigIron RX is a cost-efficient and future-proof solution that has been specifically architected to address the build-out of enterprise and service provider networks around today's technologies such as 10 Gigabit Ethernet, and IPv4/IPv6 while allowing growth into evolving technologies such as 40-Gig or 100-Gig interfaces. Its robust system architecture, versatile feature set and availability in 3 different sizes allows network designers to standardize on a single product family for needs ranging from aggregation to backbone switching.

Designed with state of the art packet processing technology, the BigIron RX has a non-blocking switching capacity of up to 3.84 Tbps and routing performance of up to 1.14 billion packets per second. Its advanced distributed hardware architecture with fine-grained QoS support allows uncompromised full-duplex, wire-speed performance to be achieved for any mix of IPv4, IPv6 and Layer 2 services. These capabilities are made possible by an innovative system architecture that has several distinguishing characteristics:

- Clos-based[1] self-routing, distributed, non-blocking architecture provides the foundation for a robust, scalable platform
- Distributed packet processing and advanced QoS capabilities across the system allow a rich set of features to be implemented at wire-speed rates
- High availability architecture with a clear separation between control and data planes
- Fully redundant architecture with redundant power supplies, management modules, fan trays and switch fabric modules to avoid any single point of failure

The BigIron RX is available in three different configurations:

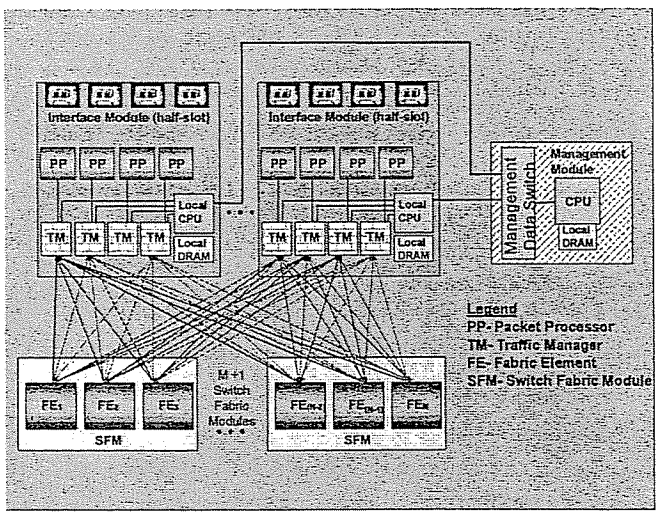
- BigIron RX-16, a 14 RU, 16 interface-slot system
- BigIron RX-8, a 7 RU, 8 interface-slot system
- BigIron RX-4, a 4 RU, 4 interface-slot system

All modules on the BigIron RX are hot pluggable. The management and interface modules can be interchangeably used across any of these systems, thereby decreasing inventory and maintenance costs for network administrators.

Industry-leading density

The BigIron RX is scalable to an industry leading density of 64 non-blocking 10 Gigabit Ethernet ports or 768 non-blocking Gigabit Ethernet ports in a single chassis. In a standard 7' rack, the BigIron RX can support up to 192 10 Gigabit Ethernet ports or 2,304 Gigabit Ethernet ports. These capabilities give the BigIron RX an extraordinary density of 55 Gigabit Ethernet ports per rack unit (RU) and 110 Gbps of full-duplex switching capacity per RU.

Scalable Clos fabric architecture



The BigIron RX uses a Clos fabric architecture that provides a high level of scalability, redundancy and performance. As shown in the Figure above, there are multiple switch fabric

[1] Named after the groundbreaking work by researcher Charles Clos, the Clos architecture has been the subject of much research over several years. The resiliency of this architecture makes it the ideal building block in the design of high availability, high performance systems.

©2005 Foundry Networks, Inc.

modules (SFMs) in the system. An SFM has multiple fabric elements, each of which has multiple connections to every interface slot.

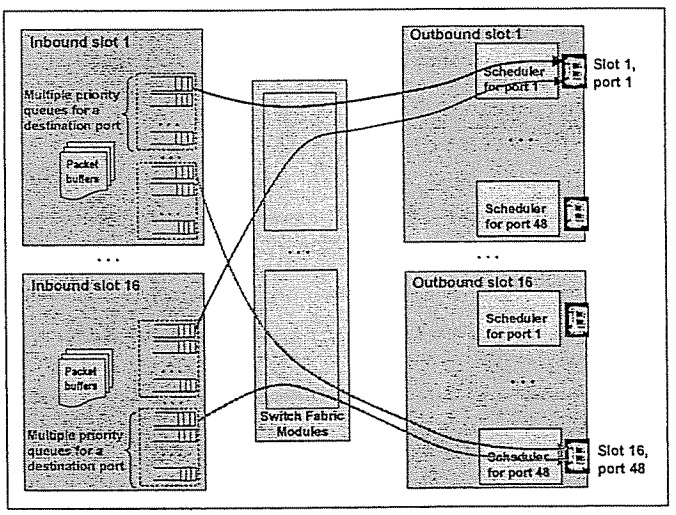
The Clos architecture uses data striping to ensure optimal utilization of fabric interconnects at all times. This mechanism always distributes the load equally across all available links between the input and output interface modules. By using fixed-size cells to transport packets across the switch fabric, the BigIron RX's switching architecture ensures predictable performance with very low and deterministic latency and jitter for any packet size. The presence of multiple switching paths between the input and output interface modules also provides an additional level of redundancy.

There are several advantages of a Clos architecture over a traditional architecture:

- **Common architecture across the product family** as the same fabric elements are used on all three chassis of the BigIron RX Series. This demonstrates the superior scalability of the architecture from a small 4 slot system to a large 16 slot system.
- **No head-of-line blocking** at any point irrespective of traffic pattern, packet size or type of traffic.
- **Optimal utilization of switch fabric resources** at all times. The data striping capability ensures that there is fair utilization of the switch fabric elements at all times without overloading of any single switch fabric element.
- **"Intra-SFM" redundancy**: An SFM can withstand the failure of some of the fabric elements and yet continue to operate with the remaining fabric elements. This unique capability provides a very high level of redundancy even within an SFM.
- **Exceptional high availability**: The BigIron RX SFMs have (N+1) redundancy. This allows the BigIron RX to gracefully adapt to the failure of multiple switch fabric elements. Moreover, because there are multiple fabric elements within an SFM, the failure of a fabric element does not bring down the entire SFM.

Distributed queuing for fine-grained QoS

A distinguishing characteristic of the BigIron RX architecture is the use of a distributed queuing scheme that maximizes the utilization of buffers across the whole system during congestion. This scheme marries the benefits of input-side buffering (Virtual Output Queuing) with those of an output-port driven scheduling mechanism. Input queuing using virtual output queues ensures that bursty traffic from one port does not hog too many buffers on an output port. An output-port driven scheduling scheme ensures that packets are sent to the output port only when the port is ready to transmit a packet. Each interface module maintains multiple, distinct priority queues to every output port on the system. Packets are "pulled" by the outbound interface module when the output port is ready to send a packet. Switch fabric messaging is used to ensure that there is tight coupling between the two stages. This closed loop feedback between the input and output stages ensures that no information is lost between the two stages. The use of such "virtual output queues" maximizes the efficiency of the system by storing packets on the input module until the output port is ready to transmit the packet. In all, there are 512K virtual output queues on the BigIron RX chassis that are distributed across the system.



Complementing the large number of virtual output queues is a rich set of scheduling algorithms that can be applied to each Egress port. These mechanisms can either be used individually or in combination to deliver tiered QoS guarantees for several applications on the same port:

- Strict priority
- Enhanced strict priority scheduling
- Weighted Fair Queue (WFQ) destination-based scheduling

©2005 Foundry Networks, Inc.

- Weighted Fair Queue (WFQ) source-based scheduling
- Maximum rate-based scheduling
- Minimum rate-based scheduling

For example, multiple queues can be serviced in the order of decreasing priority in this architecture. These advanced QoS capabilities make the BigIron RX architecturally superior to existing legacy architectures.

Congestion avoidance is handled by applying Weighted Random Early Discard (WRED) or tail-drop policy.

The QoS subsystem on the BigIron RX has extensive classification and packet marking capabilities that can be configured:

- Prioritization based on Layer 2 (802.1p), TOS, or DSCP of an input packet
- Mapping of packet/frame priority from Ingress encapsulation to Egress encapsulation
- Remarking of a packet's priority based on the result of the 2-rate, 3-color policer.

IPv6 capabilities

The increasing migration to converged network architectures and the ubiquity of IP-enabled user devices makes the eventual migration to IPv6 inevitable. To ease this migration, the BigIron RX supports IPv6 over IPv4 tunnels. Every interface module of the BigIron RX has native hardware support for dual-stack IPv4/IPv6 routing with a route information base capacity of 200K IPv6 routes. The hardware FIB per interface module allows wire-speed, full-duplex IPv6 performance to be achieved even at full capacity. Each interface module also supports IPv6 ACLs in hardware. IPv6 routing protocols that are supported today include MP-BGP-4, OSPFv3, IS-ISv6 and RIPng. These capabilities make the BigIron RX an ideal system for IPv6 networks.

High Availability

Both the hardware and software architecture of the BigIron RX are designed to ensure very high Mean Time Between Failures (MTBF) and low Mean Time To Repair (MTTR). Cable management and module insertion on the same side of the chassis allows ease of serviceability when a failed module needs to be replaced or a new module needs to be inserted.

The ability to handle the failure of not only an SFM but also elements *within* an SFM ensures a robust, redundant system ideal for non-stop operation. The overall system redundancy is further bolstered by redundancy in other active system components such as power supplies, fans, and management modules. The passive backplane on the BigIron RX chassis increases the reliability of the system. Self-adjusting, variable-speed fans in the BigIron RX chassis help to maintain the optimal operating temperature. Cards that cross a pre-set temperature threshold will be powered off by the management module so as not to completely disrupt switch function.

The BigIron RX also supports the ability to gracefully shut down a switch fabric module with zero packet loss for a scheduled maintenance event. When this facility is invoked, the system will not use the links between the interface modules and the decommissioned SFM.

The modular architecture of IronWare operating system has several distinguishing characteristics that differentiate it from legacy operating systems:

- Industry-leading cold restart time of less than a minute
- Support for hitless software upgrade
- Hitless Layer 2 and Layer 3 failovers
- Sub-second switchover to the standby management module if a communication failure occurs between active and standby management modules.
- Support for graceful OSPF restart and graceful BGP restart

Distributed forwarding for wire-speed performance at any packet size

The BigIron RX has a distributed forwarding architecture that combines state of the art packet processing technology with a very fast switch fabric to ensure uncompromised, full-duplex, wire-speed performance at any packet size. These packet processors implement a variety of features such as access control lists, policing, classification, and multicast replication. The use of fast packet processors on each interface module allows wire-speed performance to be maintained, independent of the features that have been enabled.

©2005 Foundry Networks, Inc.

There are several capabilities that have been implemented in the IronWare® operating system software to facilitate distributed packet forwarding and security:

- **Distributed maintenance of the Layer 2 MAC address table on each interface module**: The management module maintains all the learned MAC addresses and distributes the information to be locally maintained on the interface modules. Each interface module locally handles aging of its local MAC addresses and updates the management module in order to keep the MAC table consistent across the entire system.
- **Foundry Direct Routing (FDR)** technology stores the entire forwarding table in each interface module to allow for hardware forwarding of all traffic.
- **Distributed Access Control List (ACL) maintenance**: Each interface module has hardware support for input ACLs on all ports.

Hardware assists for a robust, scalable infrastructure

The BigIron RX has several innovative hardware assists that enables network designers to create a scalable and stable network. These include:

- **CPU protection:** Layer 2 services require support for efficient replication of packets to the entire broadcast domain. For example, traditional architectures handle Ethernet frames with unknown MAC address by sending them to a processor to replicate the packet to the broadcast domain. The involvement of the CPU makes the system vulnerable to a potential denial of service attack. In contrast, the BigIron RX handles this scenario very efficiently by performing such flooding in hardware.
- **Complete separation of control and data plane**: The isolation of traffic to control plane ensures that control packets to be processed by the system's CPU are efficiently processed even during periods of high data traffic.
- **sFlow**: By performing sFlow collection in hardware, L2-L7 flows can be precisely monitored. This facilitates rapid trouble-shooting and accurate isolation of faults in a network. The BigIron RX supports sFlow v5.

Traffic policers and ACLs

All interface modules support a large number of inbound traffic policers in hardware. Both single-rate 3-color and 2-rate 3-color policers are supported. The single-rate, 3-color marker meters flows that exceed a configured compliant (CIR) rate while the 2-rate, 3-color policers meter subscriber flows by classifying them into compliant (CIR) rates or Peak (PIR) rates. These capabilities are especially useful when mixing traffic flows with different characteristics on the same port.

Input ACLs (Access Control Lists) are supported by the system on all interface modules.

Separation of control and data plane

The BigIron RX has a dedicated out-of-band management link between each interface module and the management module to isolate control traffic from data traffic. Multiple queues to the management module allow different types of control traffic to be prioritized. These capabilities, together with secure management (via SSH, SCP and SNMPv3) and ACLs, are immensely useful in securing the system from potential DoS attacks in the network.

Spatial multicast support

The BigIron RX architecture has native support for spatial multicast, a critical requirement for offering video services in a network. The input interface module sends one copy of an incoming multicast packet to the switch fabric. The switch fabric then replicates the packet within itself to multiple output interface modules in the system, which in turn replicate the multicast packet to the destination ports and VLANs.

100-Gig ready slots

The BigIron RX has a future-proof architecture. Every interface slot has over 48 Gbps of full-duplex switching bandwidth available. The divider between two adjacent interface slots can be removed to convert the half slots into a full slot. The full slot is equipped to handle 100 Gbps of full-duplex bandwidth from the backplane. This makes the BigIron RX chassis capable of migration to 40 Gigabit Ethernet or 100 Gigabit Ethernet interfaces in future.

©2005 Foundry Networks, Inc.

Industry-leading feature set

Foundry Networks has built on the cumulative experience gained in powering enterprise and service provider networks for over 7 years to create the IronWare software that runs on BigIron RX. The software complements the BigIron RX architecture to offer the following capabilities:

- Support for BGPv4, OSPF, IS-IS and RIP routing protocols in IPv4 networks
- Support for IPv6 including MP-BGP-4, OSPFv3, IS-ISv6 and RIPng routing protocols
- IGMP, MLD, PIM-SM/-DM, and DVMRP support to power multicast applications
- Layer 3 redundancy protocols such as Virtual Router Redundancy Protocol (VRRP), and Virtual Router Redundancy Protocol- Extended (VRRP-E)
- Layer 2 redundancy protocols such as Virtual Switch Redundancy Protocol (VSRP)
- Support for MAC layer service protection protocols such as Metro Ring Protocol (MRP), Rapid Spanning Tree Protocol (RSTP)
- Support for secure management via SSH (v1 and v2), SCP (v1 and v2) or SNMPv3
- Support for 802.1x authentication
- sFlow-based L2-L7 traffic monitoring of activity on the node with underlying hardware support for reliable packet sampling

The BigIron RX architecture allows both Layer 2 and Layer 3 services to be offered on the same device and the same port **concurrently**. This ability gives unprecedented flexibility to the network administrator in tailoring the system to meet end user needs.

Conclusion

The BigIron RX Series is the industry's most powerful IPv4/IPv6 Ethernet switch family. Its robust, scalable architecture coupled with the feature-rich IronWare networking software make it an ideal infrastructure solution for a wide range of customer environments including enterprise backbone, data centers, service provider infrastructure and high performance and cluster computing centers. Featuring the highest density wire-speed Gigabit and 10 Gigabit Ethernet solution in a single system and designed with a redundant switch fabric supporting 100 Gbps capacity per full slot, the BigIron RX delivers the scalability, reliability and performance necessary to meet the needs of the most demanding networking environments.

Document version 2.6

Foundry Networks, Inc.
Headquarters
2100 Gold Street
P.O. Box 649100
San Jose, CA 95164-9100
U.S. and Canada Toll-free: (888) TURBOLAN
Direct telephone: +1 408.586.1700
Fax: +1 408.586.1900
Email: info@foundrynet.com
Web: http://www.foundrynet.com

Foundry Networks, AccessIron, BigIron, EdgeIron, FastIron, IronPoint, IronView IronWare, JetCore, NetIron, ServerIron, Terathon, TurboIron, and the "Iron" family of marks are trademarks or registered trademarks of Foundry Networks, Inc. in United States and other countries. All other trademarks are the properties of their respective owners.

Although Foundry has attempted to provide accurate information in these materials, Foundry assumes no legal responsibility for the accuracy or completeness of the information. More specific information is available on request from Foundry. Please note that Foundry's product information does not constitute or contain any guarantee, warranty or legally binding representation, unless expressly identified as such in a duly signed writing.

©2005 Foundry Networks, Inc. All Rights Reserved.

EXHIBIT I

FOUNDRY NETWORKS

NetIron XMR architecture brief

The NetIron® XMR Series of routers is the industry's most powerful suite of IPv4/IPv6/MPLS multi-service routers. It is a cost-efficient solution that is purpose-built to handle the most demanding of service provider applications with non-blocking, wire-speed performance. The NetIron XMR's robust system architecture, versatile feature set and availability in 3 different sizes makes it capable of scaling from the edge to the core.

Designed with state of the art network processing technology, the NetIron XMR has a non-blocking switching capacity of 1.536 Tbps and routing performance of 960 Mpps. Its advanced distributed hardware architecture with fine-grained QoS support allows uncompromised full-duplex, wire-speed performance to be achieved for any mix of IPv4, IPv6 and MPLS services. These capabilities are made possible by an innovative system architecture that has several distinguishing characteristics:

- Clos-based[1] self-routing, distributed, non-blocking architecture provides the foundation for a robust, scalable platform
- Distributed network processing and advanced QoS capabilities across the system allow a rich set of features to be implemented at wire-speed rates
- High availability architecture with a clear separation between control plane and data plane
- Fully redundant architecture with redundant power supplies, management modules, fan trays and switch fabric modules to avoid any single point of failure

The NetIron XMR is available in three different configurations:

- NetIron XMR 16000, a 14 RU, 16 interface-slot system
- NetIron XMR 8000, a 7 RU, 8 interface-slot system
- NetIron XMR 4000, a 4 RU, 4 interface-slot system

[1] Named after the groundbreaking work by researcher Charles Clos, the Clos architecture has been the subject of much research over several years. A multi-stage Clos architecture has been mathematically proven to be non-blocking. The resiliency of this architecture makes it the ideal building block in the design of high availability, high performance systems.

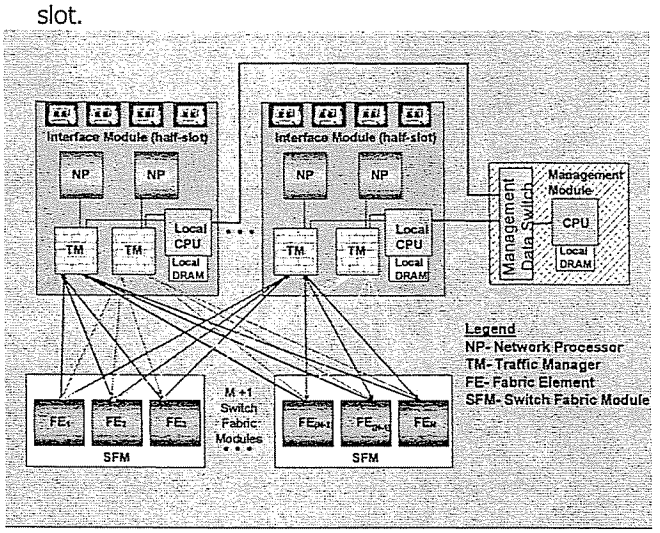
The management and interface modules can be interchangeably used across any of these systems, thereby decreasing inventory and maintenance costs for service providers across the network. All modules are hot pluggable.

Industry-leading density

The NetIron XMR is scalable to an industry leading density of 64 10 Gigabit Ethernet ports or 320 Gigabit Ethernet ports in a single chassis. In a standard 7' telco rack, the NetIron XMR can support up to 192 10 Gigabit Ethernet ports or 960 Gigabit Ethernet ports.

Scalable Clos fabric architecture

The NetIron XMR uses a Clos fabric architecture that provides a high level of scalability, redundancy and performance. As shown in the Figure below, there are multiple switch fabric modules (SFMs) in the system. A switch fabric module has multiple fabric elements, each of which has multiple connections to every interface slot.



The Clos architecture uses data striping technology to ensure optimal utilization of fabric interconnects. This mechanism always distributes the load equally across all available links between the input and output interface modules. By using fixed-size cells to transport packets across the switch fabric, the NetIron XMR's switching architecture ensures predictable performance with very low and deterministic latency and jitter for any packet size. The presence of multiple

©2005 Foundry Networks, Inc.

switching paths between the input and output interface modules also provides an additional level of redundancy.

There are several advantages of a Clos architecture over a traditional architecture:

- **Common architecture across the product family** as the same fabric elements are used on all three chassis of the NetIron XMR Series router. This demonstrates the superior scalability of the architecture from a small 4 slot system to a large 16 slot system.
- **No head-of-line blocking** at any point irrespective of traffic pattern, packet size or type of traffic.
- **Optimal utilization of switch fabric resources** at all times. The data striping capability ensures that there is fair utilization of the switch fabric elements at all times without overloading of any single switch fabric element.
- **"Intra-SFM" redundancy**: An SFM can withstand the failure of some of the fabric elements and yet continue to operate with the remaining fabric elements. This unique capability provides a very high level of redundancy even within an SFM.
- **Exceptional high availability**: The NetIron XMR SFMs have (N+1) redundancy. This allows the NetIron XMR to gracefully adapt to the failure of multiple switch fabric elements. Moreover, because there are multiple fabric elements within an SFM, the failure of a fabric element does not bring down the entire SFM.

Distributed forwarding for wire-speed performance at any packet size

The NetIron XMR has a distributed forwarding architecture that combines state of the art network processing technology with a very fast switch fabric to ensure uncompromised, full-duplex, wire-speed performance at any packet size. The use of fast network processors on each interface module allows wire-speed performance to be maintained, independent of the features that have been enabled. In contrast to custom ASICs, fast network processors provide flexibility in adding capabilities as new standards / requirements emerge in future.

There are several capabilities that have been implemented in Multi-Service IronWare operating system software to facilitate distributed packet forwarding and security:

- **Distributed Layer 2 MAC address table on each interface module**: The management module maintains all the learned MAC addresses and distributes the information to be locally maintained on the interface modules. Each interface module locally handles aging of its local MAC addresses and updates the management module in order to keep the MAC table consistent across the entire system.
- **Foundry Direct Routing (FDR)** technology stores the entire forwarding table in each interface module to allow for hardware forwarding of all traffic.
- **Distributed Access Control List (ACL)**: Each interface module can support up to 229,376 input ACL entries and 131,072 output ACL entries for ACL rules applied to local interfaces.

High Availability

Both the hardware and software architecture of the NetIron XMR are designed to ensure very high Mean Time Between Failures (MTBF) and low Mean Time To Repair (MTTR). Cable management and module insertion on the same side of the chassis allows ease of serviceability when a failed module needs to be replaced or a new module needs to be inserted.

The ability to handle the failure of not only an SFM but also elements *within* an SFM ensures a robust, redundant system ideal for non-stop operation. The overall system redundancy is further bolstered by redundancy in other active system components such as power supplies, fans, and management modules. The passive backplane on the NetIron XMR chassis increases the reliability of the system.

Temperature sensors on the system are used to automatically adjust the speed of the fans to maintain an optimal operating temperature. There is also a facility to automatically power-off a module if the configured temperature threshold is crossed.

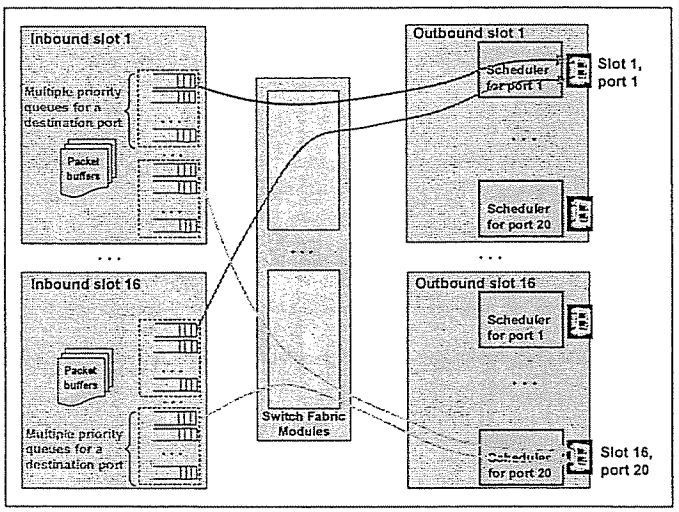
The modular architecture of Multi-Service IronWare® operating system has several distinguishing characteristics that differentiate it

©2005 Foundry Networks, Inc.

from legacy operating systems that run on routers:

- Industry-leading cold restart time of less than a minute
- Support for hitless software upgrade
- Hitless Layer 2 and Layer 3 failovers
- Sub-second switchover to the standby management module if a communication failure occurs between active and standby management modules.

Distributed queuing for fine-grained QoS



A unique characteristic of the NetIron XMR Series is the use of a distributed queuing scheme that maximizes the utilization of buffers across the whole system during congestion. This scheme marries the benefits of input-side buffering (Virtual Output Queuing) with those of an output-port driven scheduling mechanism. Input queuing using virtual output queues ensures that bursty traffic from one port does not hog too many buffers on an output port. An output-port driven scheduling scheme ensures that packets are sent to the output port only when the port is ready to transmit a packet. Each interface module maintains multiple, distinct priority queues to every output port on the system. Packets are "pulled" by the outbound interface module when the output port is ready to send a packet. Switch fabric messaging is used to ensure that there is tight coupling between the two stages. This closed loop feedback between the input and output stages ensures that no information is lost between the two stages. The use of such "virtual output queues" maximizes the efficiency of the system by storing packets on the input module until the output port is ready to transmit the packet. In all, there are 256K virtual output queues on the NetIron XMR chassis.

Congestion avoidance is handled by applying Weighted Random Early Discard (WRED) or tail-drop policy. On the output ports, a variety of scheduling mechanisms such as strict priority, weighted fair queuing or a combination of these approaches can be applied to deliver tiered QoS guarantees for several applications.

The QoS subsystem on the NetIron XMR has extensive classification and packet marking capabilities that can be configured:

- Prioritization based on Layer 2 (802.1p), TOS, DSCP or MPLS EXP bit of an input packet
- Mapping of packet/frame priority from Ingress encapsulation to Egress encapsulation
- Remarking of a packet's priority based on the result of the 2-rate, 3-color policer.

Traffic policers and ACLs

All interface modules support a large number of both inbound and outbound traffic policers in hardware. Up to 256K traffic policers can be concurrently configured in the system. The 2-rate, 3-color policers meter subscriber flows by classifying them into compliant (CIR) rates or excess (EIR) rates. This capability is especially useful when mixing traffic flows with different characteristics on the same port.

For security purposes, both input ACLs (Access Control Lists) and output ACLs are supported by the system on every interface module. Up to 229,376 input ACL entries and 131,072 output ACL entries for ACL rules can be applied to local interfaces on every interface module.

Denial of Service (DoS) guards

Layer 2 services such as VPLS require support for efficient replication of packets to the entire broadcast domain. For example, traditional architectures handle Ethernet frames with unknown MAC address by sending them to a processor to replicate the packet to the broadcast domain. The involvement of the CPU makes the system vulnerable to a potential denial of service attack. In contrast, the NetIron XMR handles this scenario very efficiently by performing the flooding in hardware.

©2005 Foundry Networks, Inc.

The NetIron XMR has a dedicated out-of-band management link between each interface module and the management module to isolate control traffic from data traffic. Multiple queues to the management module allow different types of control traffic to be prioritized. These capabilities, together with secure management and ACLs, are immensely useful in protecting the system from potential DoS attacks in the network.

Spatial multicast support

The NetIron XMR architecture has native support for spatial multicast, a critical requirement for offering video services in a network. The input interface module sends one copy of an incoming multicast packet to the switch fabric. The switch fabric then replicates the packet within itself to multiple output interface modules in the system, which in turn replicate the multicast packet to the destination ports.

Industry-leading multi-service feature set

Multi-Service IronWare software leverages the cumulative experience that Foundry Networks has gained in powering service provider networks for over 7 years. The software complements the NetIron XMR architecture to offer the following capabilities:

- Support for BGPv4, OSPF, IS-IS and RIP routing protocols in IPv4 networks
- Support for IPv6 including MP-BGP-4, OSPFv3 and RIPng routing protocols
- Support for MPLS including signaling protocols such as RSVP-TE, LDP
- Extensive traffic engineering support for MPLS
- Layer 2 VPN using VPLS or VLL
- Layer 3 VPN using RFC 2547bis or multi-VRF
- IGMP, MLD, PIM-SM/-DM, and DVMRP support to power multicast applications
- Layer 3 redundancy protocols such as Virtual Router Redundancy Protocol (VRRP), and Virtual Router Redundancy Protocol- Extended (VRRP-E)
- Layer 2 redundancy protocols such as Virtual Switch Redundancy Protocol (VSRP)
- Support for MAC layer service protection protocols like Metro Ring Protocol (MRP), Rapid Spanning Tree Protocol (RSTP)
- Support for secure management via SSH (v1 and v2), SCP (v1 and v2) or SNMPv3
- sFlow-based L2-L7 traffic monitoring of activity on the node with underlying hardware support for reliable packet sampling

In contrast to some systems that limit the capabilities that can be concurrently enabled, the NetIron XMR's architecture allows both Layer 2 and Layer 3 services to be offered on the same device and the same port **concurrently**. This ability gives unprecedented flexibility to the service provider in tailoring the system to meet end user needs.

Scalability

The NetIron XMR series of routers is a highly scalable family of routers. Some examples of its industry-leading scalability include:

- Up to 16K VPLS/VLL instances and up to 1 million VPLS MAC addresses
- Support for 4094 VLANs and up to 2 million MAC addresses
- 1 million IPv4 routes in hardware
- 256K IPv6 routes in hardware
- 10 million BGP routes
- 2000 BGP/MPLS VPNs and up to 1 million VPN routes

Investment protection

The NetIron XMR chassis uses a half slot design for interface modules. The divider between two adjacent half slots can be removed in future to combine them into a full slot. All chassis have 40 Gbps ready half slots and 100 Gbps ready full slots, making the system ready for handling interface speeds beyond 10 Gbps in future. In addition, with the ability to offer multiple services including dual-stack IPv4/IPv6 and MPLS services in hardware, the NetIron XMR offers excellent investment protection for a service provider.

Conclusion

The NetIron XMR is the industry's fastest IPv4/IPv6/MPLS router in a single shelf. Its robust, scalable architecture coupled with a rich feature set in Multi-Service IronWare software makes it the leading IPv4/IPv6/MPLS router in its class. An industry leading density of Ethernet ports in a single rack along with 40 Gbps / 100 Gbps ready slots makes it an excellent investment for service providers planning to build a converged multi-service network for the future.

©2005 Foundry Networks, Inc.



Document version 2.3

Foundry Networks, Inc.
Headquarters
2100 Gold Street
P.O. Box 649100
San Jose, CA 95164-9100
U.S. and Canada Toll-free: (888) TURBOLAN
Direct telephone: +1 408.586.1700
Fax: +1 408.586.1900
Email: info@foundrynet.com
Web: http://www.foundrynet.com

Foundry Networks, AccessIron, BigIron, EdgeIron, FastIron, IronPoint, IronView IronWare, JetCore, NetIron, ServerIron, Terathon, TurboIron, and the "Iron" family of marks are trademarks or registered trademarks of Foundry Networks, Inc. in United States and other countries. All other trademarks are the properties of their respective owners.

Although Foundry has attempted to provide accurate information in these materials, Foundry assumes no legal responsibility for the accuracy or completeness of the information. More specific information is available on request from Foundry. Please note that Foundry's product information does not constitute or contain any guarantee, warranty or legally binding representation, unless expressly identified as such in a duly signed writing.

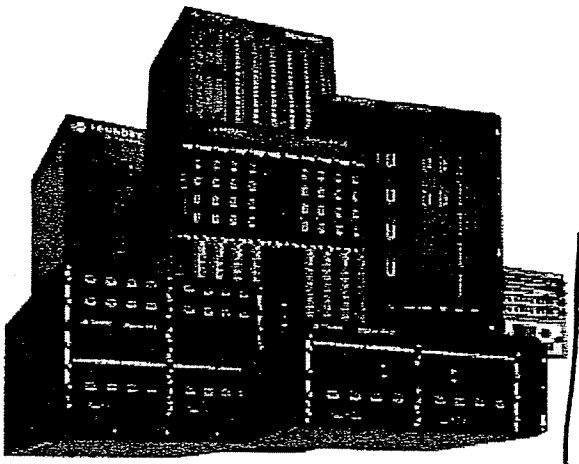
©2005 Foundry Networks, Inc. All Rights Reserved.

EXHIBIT J

Search

Quick find:

- BigIron
  - BigIron RX Series
  - BigIron MG8 Datasheet
  - JetCore BigIron Datasheet
  - White Papers
  - Application Notes
  - Customer Case Studies
  - Customer Testimonials
  - Test Reports
  - Documentation



BigIron RX, MG8, 4000, 8000, and 15000

Enterprise-class BigIron Layer 2/Layer 3 Switches

The BigIron product family now includes the powerful Terathon™ ASIC-based architecture, to deliver a new level of high-performance, distributed switching and routing for enterprise and service provider networks. The newest family member, BigIron RX-Series redefines terabit-capacity for routing and switching in the Enterprise and Service Provider environments. Flexibility, resiliency, security and performance are hallmarks of the design of the BigIron RX-Series family of 10 GbE routers.

The BigIron MG8, offers high density 10-Gigabit Ethernet switch for performance oriented Enterprise. All BigIron switches are purpose-built for bandwidth intensive enterprise networks to support business and research applications including high bandwidth campus backbone, data center aggregation, cluster and grid computing, converged (voice, video, image and data) communications, storage networking, financial transactions as well as government, healthcare and university super-computing research projects. BigIron's non-blocking architecture enables network managers to build scalable and highly available network designs. BigIron is the only chassis that has been 10-Gigabit ready since the day it started shipping in 1998.

"Enterprises and research institutions need both evolutionary as well as revolutionary migration paths to high-performance Ethernet networks," said Mike Warren, a computational astrophysicist and architect of the Space Simulator Beowulf cluster at Los Alamos National Laboratory. "Performance-oriented organizations will benefit from having a next-generation 10-Gigabit networking infrastructure which is highly scalable and resilient, with high Gigabit and 10-Gigabit Ethernet density. Foundry's BigIron MG8 will be one of the revolutionary products in this category which offers the performance required to support high-bandwidth networks such as our Beowulf clusters."

-Los Alamos

Copyright ©2006 Foundry Networks, Inc. All rights reserved. | Privacy | Terms of Use