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.

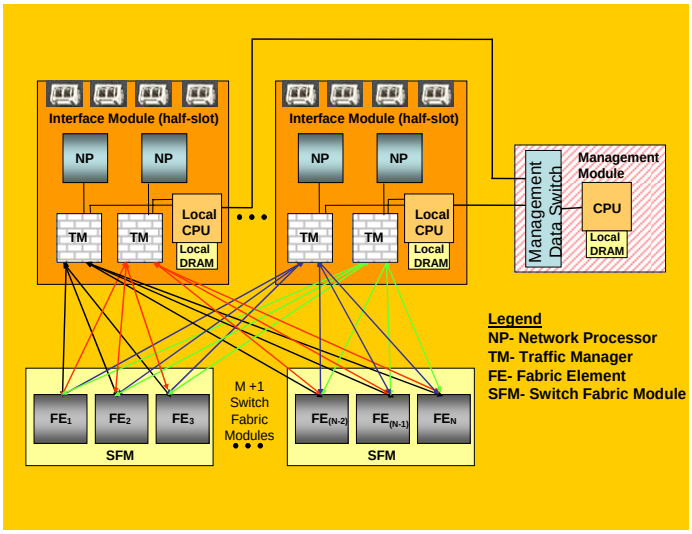

The Clos architecture uses data striping technology to ensure optimal utilization of fabric interconnects. This mechanism always distributes the load equally across all available links between the input and output interface modules. By using fixed-size cells to transport packets across the switch fabric, 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

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

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

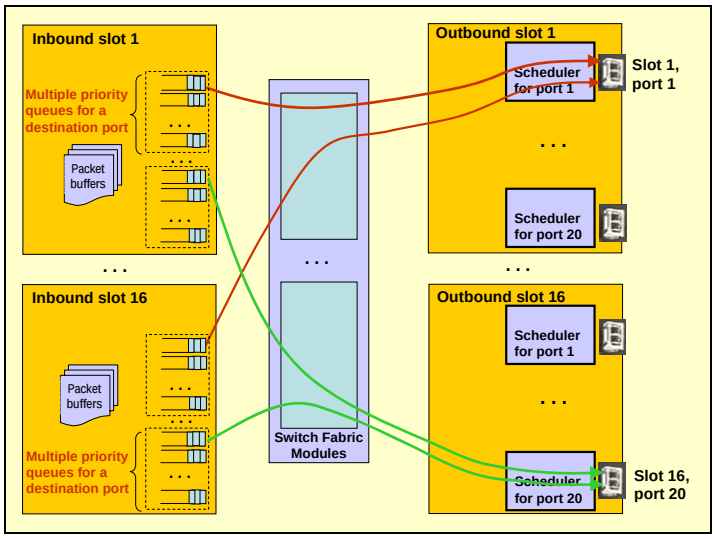

A unique characteristic of the NetIron XMR Series is the use of a distributed queuing scheme that maximizes the utilization of buffers across the whole system during congestion. This scheme marries the benefits of input-side buffering (Virtual Output Queuing) with those of an output-port driven scheduling 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.

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.

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