

IP NETWORK

Brocade NetIron MLX Architecture

The Brocade NetIron MLX Series of routers is the industry's most powerful suite of IPv4/IPv6/MPLS/Multi-VRF switching routers. Its robust system architecture, versatile feature set, and choice of form factors makes it capable of scaling from the edge to the core.

BROCADE

The Brocade® NetIron® MLX Series of routers is a cost-effective solution that is purpose-built to handle the most demanding of service provider applications with non-blocking, wire-speed performance.

OVERVIEW

Designed with state-of-the-art network processing technology, the Brocade NetIron MLX has a non-blocking switching capacity of 7.68 Tbps, data forwarding capacity of 5.12 Tbps, and a total routing performance of 7.6 billion packets per second (Bpps). 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 Multi-Protocol Label Switching (MPLS) services. These capabilities are made possible by an innovative system architecture that has several distinguishing characteristics:

- Clos-based, 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 a single point of failure.

The NetIron MLX is available in four different configurations:

- NetIron MLX-32, 32 interface-slot system
- NetIron MLX-16, 16 interface-slot system
- NetIron MLX-8, 8 interface-slot system
- NetIron MLX-4, 4 interface-slot system

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 Netron MLX is scalable to an industry-leading density of 256 x 10 Gigabit Ethernet (GbE) ports, 1536 x GbE (mRJ-21) ports, or 640 x GbE (RJ-45 or SFP) ports in a single chassis. In a standard 7' telco rack, the Netron MLX can support up to 384 x 10 GbE ports, 2304 GbE (mRJ-21) ports, or 960 x GbE (RJ-45 or SFP) ports. The Netron MLX also supports up to 256 x OC-12/48 ports or 64 x OC-192 ports in a single chassis.

SCALABLE CLOS FABRIC ARCHITECTURE

The Netron MLX uses a Clos fabric architecture that provides a high level of scalability, redundancy, and performance. As shown in Figure 1, there are multiple Switch Fabric Modules (hSFMs) in the system. An hSFM has multiple fabric elements, each of which has multiple connections to every interface slot.

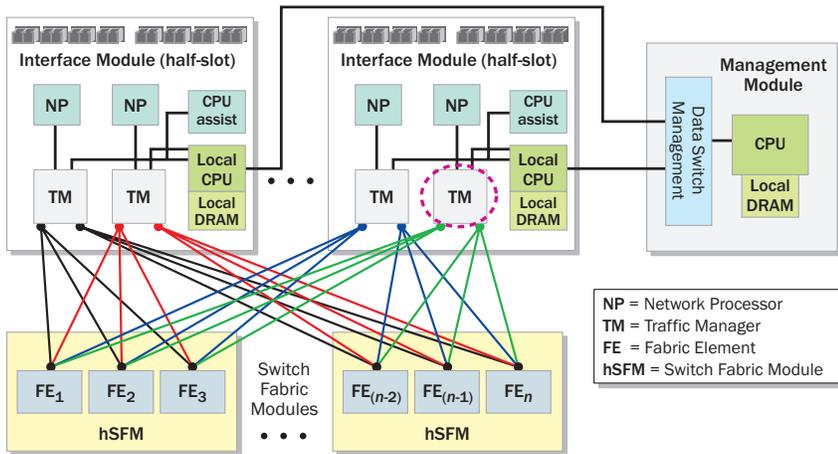


Figure 1.
Scalable Clos
fabric architecture.

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 Netron MLX switching architecture ensures predictable performance with very low and deterministic latency and jitter for any packet size. In addition, the Netron MLX offers a Turbo Mode, which increases switch fabric efficiency by using variable-size cells to transport packets. 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 traditional architecture:

- Common architecture across the product family. The same fabric elements are used on all four chassis of the Netron MLX Series. This demonstrates the superior scalability of the architecture from a small 4-slot system to a large 32-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 hSFM 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 hSFM.
- Exceptional high availability. The Netron MLX supports redundant hSFMs, which allows the Netron MLX to gracefully adapt to the failure of multiple switch fabric elements. Moreover, because there are multiple fabric elements within an hSFM, the failure of a fabric element does not bring down the entire hSFM.

CLOS ARCHITECTURE

Named after the ground-breaking 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.

DISTRIBUTED FORWARDING: WIRE-SPEED PERFORMANCE AT ANY PACKET SIZE

The NetIron MLX 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 and requirements emerge in the future.

There are several capabilities implemented in the multi-service Brocade 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.
- **Brocade Direct Routing (BDR)** technology stores the entire forwarding table in each interface module to allow for hardware forwarding of traffic.
- **Distributed Access Control List (ACL).** Each interface module can support up to 114,688 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 MLX 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 hSFM but also elements within an SFM ensures a robust, redundant system ideal for high-availability 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 MLX 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 NetIron MLX 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 does not use the links between the interface modules and the decommissioned hSFM.

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

- Industry-leading cold restart time of less than a minute
- Support for hitless Layer 2 software upgrade within a release
- 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

DISTRIBUTION QUEUING FOR FINE-GRAINED QOS.

A distinguishing characteristic of the NetIron MLX Series is the use of a distributed queuing scheme, which 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.

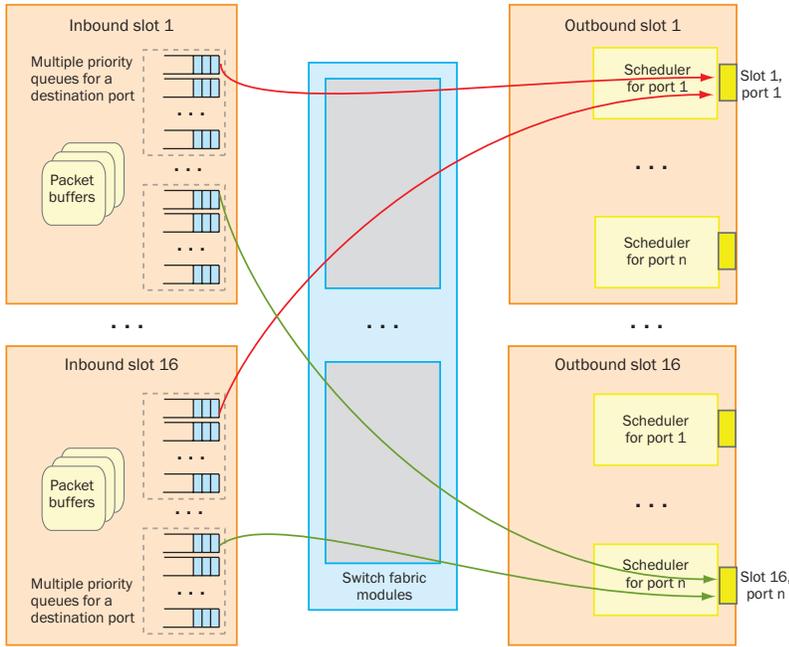


Figure 2.
Distributed queuing.

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 ensure that no information is lost between the two stages. The use of “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 NetIron MLX 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 MLX has extensive classification and packet marking capabilities that can be configured in several ways:

- 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 512K traffic policers can be concurrently configured in the system. The 2-rate, 3-color policers meter subscriber flows by classifying them into compliant (CIR) 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 Access Control Lists (ACLs) and output ACLs are supported by the system on every interface module. Up to 114,688 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 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 (DoS) attack. In contrast, the Netron MLX handles this scenario very efficiently by performing the flooding in hardware.

The Netron MLX 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 Netron MLX 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 associated with the multicast group, which in turn replicate the multicast packet to the destination ports and VLANs.

INDUSTRY-LEADING, MULTI-SERVICE FEATURE SET

Brocade has built on the cumulative experience gained in powering enterprise and service provider networks for over 10 years to create the IronWare software that runs on the Netron MLX. The software complements the Netron MLX 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-IS, and RIPng routing protocols
- Support for MPLS, including signaling protocols such as RSVP-TE, and LDP
- Extensive traffic engineering support for MPLS
- MPLS fast re-route support
- Layer 2 VPN using VPLS or VLL
- Layer 3 VPN using RFC 2547bis or multi-VRF
- IGMP, MLD, PIM-SM/-DM, PIM-SSM, and DVMRP support to power IPv4 and IPv6 multicast applications
- Hitless Layer 2 and Layer 3 failover with support for hitless software upgrades
- 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) and 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

SCALABILITY

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

- Up to 4K VPLS instances and up to 256K VPLS MAC addresses
- 8K VLLs per system
- Support for 4094 VLANs and up to 1 million MAC addresses
- 512K IPv4 routes in hardware
- 112K IPv6 routes in hardware
- 2 million BGP routes
- 400 BGP/MPLS VPNs and up to 256K VPN routes

INVESTMENT PROTECTION

The NetIron MLX 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 160 Gbps of full-duplex bandwidth per full slot. In addition, with the ability to offer multiple services including dual-stack IPv4/IPv6 and MPLS services in hardware, the NetIron MLX offers excellent investment protection for a service provider.

SUMMARY

The NetIron MLX is the industry's most advanced service delivery platform to offer IPv4/IPv6/MPLS/Multi-VRF services. Its robust, scalable architecture coupled with a rich feature set in multi-service IronWare software makes it the leading router in its class. An industry-leading density of Ethernet ports in a single rack makes it an excellent investment for service providers planning to build a converged multi-service network for the future.

To learn more, visit www.brocade.com.

Corporate Headquarters

San Jose, CA USA
T: +1-408-333-8000
info@brocade.com

European Headquarters

Geneva, Switzerland
T: +41-22-799-56-40
emea-info@brocade.com

Asia Pacific Headquarters

Singapore
T: +65-6538-4700
apac-info@brocade.com

© 2010 Brocade Communications Systems, Inc. All Rights Reserved. 04/10 GA-WP-1370-01

Brocade, the B-wing symbol, BigIron, DCX, Fabric OS, FastIron, IronView, NetIron, SAN Health, ServerIron, and Turbolron are registered trademarks, and Brocade Assurance, DCFM, Extraordinary Networks, and Brocade NET Health are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned are or may be trademarks or service marks of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

**BROCADE**