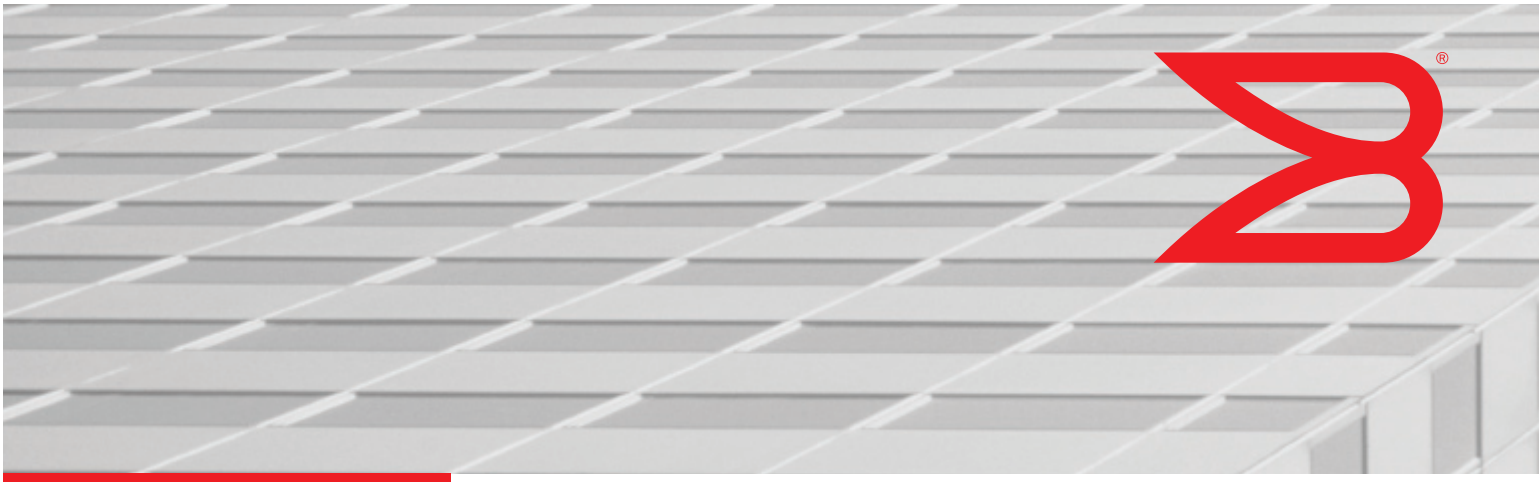

IP NETWORK

Brocade IronWare OS Powers Brocade Routers and Switches

Brocade IronWare OS provides the intelligence behind the high-performance switches, routers, and application load balancers from Brocade. Leveraging over 12 years of experience in powering global networks, Brocade IronWare OS has continually innovated to ensure non-stop network operation.

BROCADE

This white paper provides an overview of the high-level architecture of the multi-service Brocade® IronWare® OS, which powers the Brocade family of routers and switches. This operating system (OS) is designed to efficiently address the diverse needs of today's and tomorrow's infrastructures, while providing a flexible framework for powering multiple Brocade platforms.

Note that the same operating system is leveraged across different platforms and Brocade IronWare OS capabilities are fine tuned to specific products. This paper also discusses advanced capabilities such as high availability, modularity, fault tolerance, network monitoring, security, and the role of Brocade IronWare OS in the inevitable migration to IPv6.

OVERVIEW

The expanding role of IP networks in powering a variety of end-user applications in both enterprise and service provider networks, makes network uptime and, by extension, uptime of individual nodes extremely critical. Based on the experience of network administrators, it's clear that software plays a vital role in ensuring network uptime. Non-stop operation capabilities in an operating system, act as insurance against the unexpected and ensure that failures in a component of a node are quickly detected and addressed via a corrective action.

Brocade IronWare OS incorporates several capabilities that make it ideal for high performance and high availability:

- Distributed operating system with a very high level of distribution for maximum efficiency
- Modular operating system that prevents corruption or interference of different modules in Brocade IronWare OS
- Multi-threaded operating system with different software modules running as light-weight threads
- Pre-emptive operating system for predictable performance by individual threads
- High-availability capabilities to ensure non-stop operation of the system
- Advanced protective checks in the operating system to ensure reliable operation of the different software modules

MODULAR ARCHITECTURE

The modular architecture of the Brocade IronWare OS incorporates a clear separation between different modules and logical components in the operating system. For example, an abstraction layer hides the underlying hardware layer from the upper layers, such as the protocol layer. The underlying modular design of the OS greatly enhances the robustness of the system and makes it easily portable to new platforms. For network operators, this gives a common look and feel to all systems running Brocade IronWare OS, thereby decreasing the operational costs of running a network.

Figure 1 shows the key components in Brocade IronWare OS.

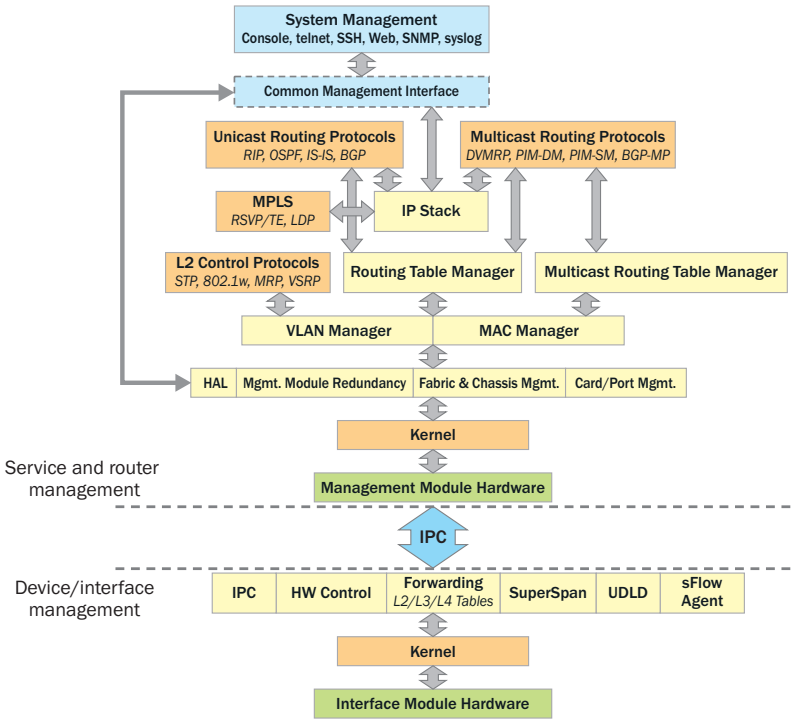

Figure 1.
Components in Brocade IronWare OS.

The modular and distributed architecture is also reflected in the industry-leading rapid cold-restart time for routers and switches powered by this operating system. For example, high-end, feature-rich systems such as Brocade BigIron® RX switches and Brocade MLX routers have a cold restart time of less than 60 seconds from cold-start to service activation, even after extensive diagnostic tests that are run during system bring-up.

Note that in Figure 1:

- Distribution of Layer 2/3/4 forwarding tables to the interface module, which allows packets to be forwarded by the hardware on the interface module without any processing by software.
- Distributed sFlow agent on every interface module which allows Brocade's switches and routers to provide very high-performance traffic monitoring even when sFlow is enabled on all ports concurrently.
- Uni-Directional Link Detection (UDLD) is performed completely on the interface module for speedy detection of link failures and resulting corrective action.

SuperSpan is also performed by the interface module for scalable and high performance tunneling of customer Spanning Tree BPDUs.

CLEAR SEPARATION BETWEEN CONTROL AND DATA PLANES

Brocade IronWare OS is architected with a clear separation between the control plane and the data plane, which ensures a high-availability system at all times. Control messages between the interface module and management module are sent on a separate out-of-band redundant link, which is distinct from the path used for data traffic. Consequently, even during times of very high utilization, control messages, which are crucial for correct system operation, can still be exchanged reliably between the management and interface modules without any loss.

Common Denial of Service (DoS) attacks in the network often target vulnerabilities in a system when system utilization becomes high. A clear separation between control and data planes greatly decreases the probability of a system being brought down by a DOS attack.

HIGH AVAILABILITY—ALWAYS

Brocade IronWare OS addresses high availability of a network from multiple angles:

- Protection against faults within the system
- Protection against faults within the network

Just as a chain is only as strong as its weakest link, a network should protect against faults in the individual nodes as well as connectivity-related faults at the physical and higher network layers.

The root cause for software faults in a system can often be traced to issues such as lost control messages, buffer overflows, or poor congestion handling mechanisms in the control plane. It is therefore extremely important for the underlying kernel to provide robust mechanisms and take preventive measures against these conditions.

- Protection against faults in the system is achieved with advanced protection capabilities of the operating system's robust kernel and fault-tolerant capabilities to ensure non-stop operation in higher layer applications.
- Protection against faults in the network is accomplished by incorporating redundant links/paths in the network architecture with the use of appropriate protocols to detect and handle network faults.

Each of these goals is accomplished using several mechanisms, described in the following section.

Ironclad Protection in the Kernel

In order to ensure non-stop operation, the OS incorporates advanced protection and fairness mechanisms. These underlying constructs ensure high availability of the system, predictable and deterministic performance, and processing prioritization of high-impact events at peak load, some of which are described below:

- Enforcement of read/write rights by the kernel on shared data structures that are accessed by multiple processes. With this capability, a process can register to access a certain shared data structure with read-only or read-write privileges. Violations of these access privileges can then be spotted immediately by the kernel.
- Detection of memory violations by using separate virtual memory spaces. Key components of the OS such as routing protocols have their databases in a secure virtual memory space to prevent accidental corruption/access of those tables by errant threads.
- Detection of errant threads. The kernel watches for symptoms of errant threads. Examples of monitored symptoms include spotting of stack usage violations, excessive memory usage by a thread, or excessive CPU "hogging" by a single thread.

- Multiple communication paths during Inter-Process Communication (IPC) with tiered priority levels between the interacting threads.
- Lightweight, reliable transport of IPC messages exchanged between processes on different slots.
- In distributed operating systems, it is imperative that the main table on the central management module and the cached copy in the interface module remain in sync. Brocade IronWare OS incorporates an integrity check mechanism wherein the data table integrity is checked periodically between the module hosting the main copy (for example, the management module) and the module hosting the cached copy (for example, the interface module).

Fault-Tolerant Capabilities

Brocade IronWare OS includes several capabilities to enhance the tolerance to faults in both the system as well as the network:

- Hitless Layer 2 and 3 failover ensures that Layer 2 and 3 protocols running on the primary management module smoothly fail over to the redundant management module with virtually no negative impact on traffic forwarding or the routing/switching domain throughout the network.
- Graceful restart for BGP and OSPF allows a router to cooperate with its adjacent routers in updating its own routing table (for example, after a management module failover) without causing network-wide disruptions due to routing protocol re-convergence.
- In-service OS upgrade allows new software patches or versions to be downloaded to the system without any loss in traffic.
- Switch fabric/fabric element/single lane failure detection in just a few milliseconds.
- Removal of switch fabric module with no loss of traffic, a capability that is extremely useful in planned network upgrades.
- Rapid detection of ECMP failures, whether the constituent links are on the same interface module or on different modules.
- Rapid detection of trunk failures, whether the constituent links in the trunk group are on the same interface module or on different modules.
- Support for protocols such as VRRP/VRRP-E, RSTP, LACP (via IEEE 802.3ad), ECMP, and BFD.
- In Multiprotocol Label Switching (MPLS) networks, fast re-route and standby LSP Paths are also invaluable in protecting against faults in the network. Brocade IronWare provides both capabilities on MPLS-enabled platforms.

The ability to perform in-service upgrades is particularly critical in ensuring high availability. Several studies have shown that planned upgrades of the network constitute a significant portion of network downtime. On the other hand, the ability of downloadable modules, wherein specific processes can be restarted, in-service upgrades with multiple management modules allows for almost uninterrupted operation of the system during upgrades.

SECURITY IN BROCADE IRONWARE OS

Brocade IronWare OS security features assist in guarding against malicious attacks, which left unattended, could eventually lead to a compromised infrastructure. Support for Secure Shell (SSH-v2) and Secure Copy (SCP) ensures that management traffic is encrypted. Similarly, use of SNMPv3 and MD5 authentication in routing protocol exchanges ensures that session or protocol exchanges are not exchanged in the clear. User authentication can be performed using mechanisms such as RADIUS or TACACS+. 802.1x authentication is also supported on switch ports with advanced features such as multiple 802.1x client support (including limiting the number of clients that can be authenticated per interface) and MAC port security.

The operating system allows configuration of the underlying hardware to detect and prevent DoS attacks such as SYN attacks or Smurf attacks by monitoring interfaces for unusual activity and setting thresholds for various traffic types. When a DoS attack is detected, the port is automatically shut off by the OS and an operator alert is logged.

Brocade's high-end platforms provide a large capacity for Access Control List (ACL) entries that are distributed on each individual interface module, which is extremely valuable in ensuring security of the network. In addition to the large numbers of hardware-based ACLs on the platforms, the OS also has extensive support for ACL accounting and ACL logging. With these capabilities, network administrators can track the number of hits to an ACL clause and also log the information in an IP packet header when a hit is encountered. Such functions are completely distributed to interface modules to increase efficiency.

In high-end platforms, a set of mechanisms called CPU protection is employed. These mechanisms protect the Local Processor (LP) residing on the interface module, against CPU hogging tasks. Under this framework, more functions that traditionally required CPU processing are delegated to the hardware, such as broadcast/multicast/unknown unicast flooding over VPLS. In addition, non-critical traffic requiring some LP processing for completing a specific task, such as ACL logging or uRPF logging, is automatically throttled.

The amount of traffic going to the control plane is a common Trojan horse that is employed by hackers to gain control of the system. Brocade IronWare OS allows ACL-based hardware traffic policers to be applied to traffic going to CPUs within the system. Further, multiple priorities are used by the OS to handle control traffic so that higher-priority control traffic is given preferential treatment in the event of a congestion.

Brocade IronWare OS also has advanced capabilities that can be used to mitigate or prevent malicious Man-in-Middle (MiM) attacks. These capabilities are implemented with support in the underlying hardware to ensure high performance even when the security features are enabled on the platform. Some of these advanced facilities are:

- BPDU guard to prevent hijacking of networks running Spanning Tree Protocol (STP) by an errant switch
- Dynamic Address Resolution Protocol (ARP) inspection to detect malicious ARP packets and false replies to ARP requests
- IP source guard to prevent source IP address spoofing by malicious users; the OS automatically installs an anti-spoof filter in the underlying platform hardware after learning the IP address on a port.
- With Dynamic Host Configuration Protocol (DHCP) snooping to automatically inspect DHCP packets and learn and maintain IP address-to-MAC address bindings.
- DHCP option 82 (relay agent) functionality to ensure that DHCP requests from clients across untrusted ports are forwarded to the DHCP server after attaching information on the circuit-id corresponding to the port.

EASING THE MIGRATION FROM IPV4TO IPV6

As IP-based delivery of services continues its onward march, the demand for IP addresses will continue to increase at a rate that will make migration to IPv6 imperative. Nevertheless, it is essential for the transition to be seamless. Brocade IronWare OS offers two mechanisms to facilitate this transition:

- Dual-stack routing, in which routers run dual IPv4 and IPv6 stacks. In this scenario, the backbone can be a dual-stack backbone with both the IPv4 and IPv6 routing protocols and forwarding processes running in parallel. End systems can also run dual IPv4/v6 stacks in such a network.
- IPv6 over IPv4 tunneling allows different IPv6 domains to communicate via an intermediate IPv4 network. The OS supports the ability to create manually configured tunnels, automatic IPv4-compatible IPv6 tunnels, and automatic IPV6-to-IPv4 tunnels.

FEATURE-RICH BROCADE IRONWARE OS

Brocade IronWare OS supports a broad range of capabilities that makes it suitable for mission-critical applications in both enterprise and service provider networks. The ability to support multiple services on the same port simultaneously is one of its key attributes. The OS offers extensive support for unicast IPv4 and IPv6 routing protocols, multicast protocols, and great flexibility in configuring ACLs and traffic policers. On core router platforms such as Brocade MLX routers, MPLS routing protocols such as RSVP, LDP, and advanced traffic engineering capabilities permit the implementation of large-scale Layer 2 and 3 Virtual Private Networks (VPNs). Advanced capabilities such as GRE tunnels and virtual routing without MPLS provide additional choices to the network designer in designing a VPN.

For details on the feature set that is enabled on a platform, consult the data sheets on www.brocade.com.

THE POWER OF SFLOW

sFlow, specified in RFC 3176, is a powerful network monitoring technology that can be used for network anomaly detection, fault management, performance management, service accounting, billing, and more. sFlow uses a statistical packet sampling approach to accomplish its objectives and can be used to troubleshoot Layer 2 to Layer 4 flows in the network. Brocade IronWare OSs distributed implementation of the sFlow agent on the target platform allows real-time traffic monitoring on all ports without performance compromise. On platforms that support L2 VPNs or L3 VPNs, such as Brocade NetIron® CER 2000 or Brocade MLX Series, sFlow can also be used on VPN endpoints to provide valuable information on VPN traffic.

SUMMARY

Brocade IronWare OS is a feature-rich, multi-threaded, distributed operating system with advanced capabilities to ensure secure, non-stop operation and high availability of the network. The versatility of its design makes it an ideal operating system for powering many of the routers and switches available today from Brocade.

ABOUT BROCADE

Brocade provides innovative, end-to-end network solutions that help the world's leading organizations transition smoothly to a virtualized world where applications and information can reside anywhere. These solutions deliver the unique capabilities for a more flexible IT infrastructure with unmatched simplicity, non-stop networking, optimized applications, and investment protection. As a result, organizations in a wide range of industries can achieve their most critical business objectives with greater simplicity and a faster return on investment.

For more information about Brocade products and solutions, 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.co

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

BROCADE