White Paper

The Anatomy of Today's Core Routing Operating System

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
408 745 2000 or 888 JUNIPER
www.juniper.net

Part Number: 200114-001

Contents

Introduction

The demands on Internet Protocol (IP) networks have evolved dramatically over the last decade. Nowhere is this more apparent than at the heart of service provider networks: the core. Service providers continue to make the transition from best-effort Internet to a singular IP / MPLS infrastructure capable of delivering assured experiences for not only data, but also voice and video. This new network paradigm – an Infranet – only increases the technical pressure on core routing hardware and software. Today's changing environment is particularly demanding on the routing operating system (OS). As many service providers and vendors are finding, perhaps the greatest attribute of a core routing OS is an intangible, rather than a technical one: foresight.

In 1998, Juniper Networks launched the first purpose-built carrier-class router, the M40 Internet backbone router. The JUNOS Software, which ran on the M40 router, was ground-breaking given that it was the first operating system to have a modular architecture and be purpose-built for service provider environments. Since its initial introduction with the M40 router, JUNOS software has attained over 7 years operational experience in the world's largest service provider networks and now executes on all M-series, T-series, and J-series routers.

Recently, other router vendors have started to implement the router architecture and modular software designs first introduced by Juniper Networks in 1998. Although these implementations emulate the fundamental design principles pioneered by Juniper Networks, they are not equal to the Juniper Networks implementations. The ASIC hardware used in the Packet Forwarding Engine (PFE) of Juniper Networks routers supports all features on all interfaces and has constantly demonstrated its merit in production networks over the years. JUNOS software is a full featured implementation, providing regular and reliable updates while simplifying the software upgrade management process. While the initial success of JUNOS software demonstrates Juniper Networks understanding of core requirements and design expertise, the subsequent longevity and extensibility of JUNOS highlights Juniper Networks vision and foresight.

Router Operating Systems

A router operating system is responsible for implementing and supporting the control plane of the network. The key functions performed by a router operating system include:

- User interface management
- Chassis management
- Network interface management
- Routing protocol management
- Local packet management
- Network management

Each router vendor releases periodic operating system updates that run on their hardware platforms to introduce new features that service providers require to support growth and provide revenue generating services. Since the software upgrade process is an essential part of network operations, it should be simple to manage and feature updates should not

negatively impact the delivery of customer services.

The operating system executes on the routing engine (i.e., computer) and is closely coupled to the vendor's custom processor hardware that implements the packet forwarding engine. Because the processor hardware implements the control plane that is defined by the router's configuration, it is very difficult to evaluate an operating system without also evaluating the processor hardware that supports the packet forwarding process. For example, a feature rich operating system that is not efficiently and accurately supported by hardware will provide inadequate network control. If the processor hardware does not provide adequate performance, then enabling certain features can severely compromise the forwarding performance of the router. The router hardware must scale with the operating system if it is to sustain the intent of the software and provide feature consistency on both low speed *and* high speed interfaces.

Implementation Requirements

The numerous requirements for a carrier-class operating system include:

- A modular software design
- Industrial strength routing protocols
- Stability
- Scalability
- Security
- Multiple platform support
- Dependable software updates
- Engineering discipline

Modular Software Design

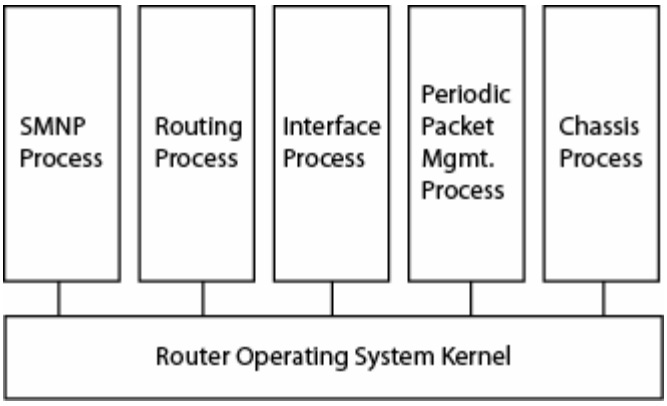

Figure 1: Figure 1: Modular Operating system

A modular operating system allows each individual process to be independently restarted or changed without affecting the operation of the router and protects the entire operating system from crashing due to the failure of a single module. Modularity also allows users to upgrade a specific software process without rebooting the entire system and permits router vendors to make modular upgrades to their software.

One of the challenges when designing a modular operating system is to determine the specific functions that should be performed within each module. For example, should each individual routing protocol execute as its own dedicated process or should all of the routing protocols execute within a single shared process? This basic design decision can severely impact the complexity of the runtime system and memory utilization. If each individual routing protocol executes as an independent process, the control plane chatter between each routing protocol processes is increased and more memory is required to maintain duplicate state and support the additional inter process communication. Excessive modularization of an operating system can potentially negate any advantage of modularizing the routing protocols.

JUNOS Software was the first operating system to be specifically designed for deployment in rapidly-growing and highly-stressful service provider networks. It is based on BSD Unix and is highly modular. JUNOS Software consists of a series of system processes that handle the router's management, routing protocol, and control functions. The JUNOS kernel, which is responsible for scheduling and device control, underlies the support of these individual processes.

With years of experience in the world's largest production networks, JUNOS Software has set the standard against which all other operating systems are judged. The fact that other router vendors are just beginning to launch modular operating systems acknowledges the limitations of the monolithic approach and recognizes Juniper Networks foresight and leadership in core operating system design. issues.

Industrial Strength Routing Protocols

Routers implement routing protocols to exchange the information that is used to calculate forwarding paths through the provider's network and across the Internet. The forwarding paths must support the delivery of IPv4, IPv6, and Multi-protocol Label Switching (MPLS) traffic. Among the standards-based routing protocols that must be supported by a carrier class router are the BGP-4, RIPv2, OSPF, and IS-IS. Since MPLS is used within core networks to manage bandwidth and implement quality of service for IP flows, MPLS constraint based routing, traffic engineering, and fast reroute features must also be supported.

There is a huge difference between simply providing a workable implementation and correctly and efficiently implementing these routing protocols. Core operating systems are not just enterprise operating systems executing on larger routers. They must be specifically designed to support the unique hardware, software, and service delivery requirements of provider networks. JUNOS software benefits from an experienced engineering staff and from receiving industry-feedback on our implementations. While many router vendors claim to support these routing protocols, none of the competing implementations can match the feature set, scalability, performance, security, and robustness of the JUNOS software routing protocol implementations.

Stability

As providers evolve their infrastructures to support multiservice networks and deliver reliable services that satisfy network availability targets, the stability and availability of core routers has never been more critical. Stability is concerned with the capacity of a router and its operating system to withstand the pressures of operating in large heterogeneous networks and to continue running for long periods of time without failure.

The crash of a router's operating system is more detrimental to network operations than the failure of an individual link or line card. When a link or line card fails, there are generally alternative paths available and the stress of routing protocol updates on the control plane is minimal. However, when an operating system crashes a large part of the network may become disjoint from the rest of the network, and the amount of stress placed on the control plane resulting from the flood of routing table updates can be extreme. Depending on the stability of the other routers in the network, the crash of a single router can result in a cascade of router failures causing the entire network to fail.

The stability of an operating system does not occur by accident, it must be designed into the system architecture during the earliest stages of development. JUNOS software takes a holistic approach to providing stability that is based on the four cornerstones of stable router design:

- An extremely resilient system architecture that provides a clean separation between the control plane and the packet forwarding plane.
- A modular operating system comprised of loosely coupled functional modules controlled and monitored by a master executive module.
- Robust scalable routing protocol implementations that are designed to converge quickly and support nonstop packet forwarding.
- A comprehensive set of industry-standard protection mechanisms such as SONET Automatic Protection Switching (APS 1+1), MPLS fast reroute, Bidirectional Forwarding Detection (BFD), Multilink PPP, and the Virtual Router Redundancy Protocol (VRRP).

Scalability

With IP networks becoming the fastest growing portion of a service provider's infrastructure, the scalability of core routers has become an essential carrier issue. To service providers, scalability means that both the router hardware and operating system are capable of satisfying network requirements during periods of extreme growth without requiring the installation of a new router. The hardware of a core router must be expandable via in-service upgrades and the operating system must scale in a linear fashion as the size and complexity of the network grows.

Among the many factors that determine the scalability of network operating systems are:

- The maximum number of physical and logical interfaces supported by each router.
- The speed of the routing table lookup algorithm.
- The maximum number of routes that can be stored in the routing table.
- The maximum number of adjacencies or peers that can be supported on each router.
- The maximum number of link state advertisements that can be stored in the router's link state database.

- The capacity of the routing policy definition language to easily and efficiently manage the import, export, and modification of enormous amounts of routing information.
- The maximum number of VPNs that can be supported by each router.
- The maximum number of virtual routers that can be configured on each router.
- The capacity of the underlying router hardware to consistently support all features (QoS, flow monitoring, rate limiting, packet filtering, and so forth) on both low speed *and* high speed interfaces.

Security

Due to the growing frequency and sophistication of cyber-threats, security issues are consistently at the top of any service provider's list of concerns. Security features implemented in an operating system are designed to protect users, their applications, and network resources from attack.

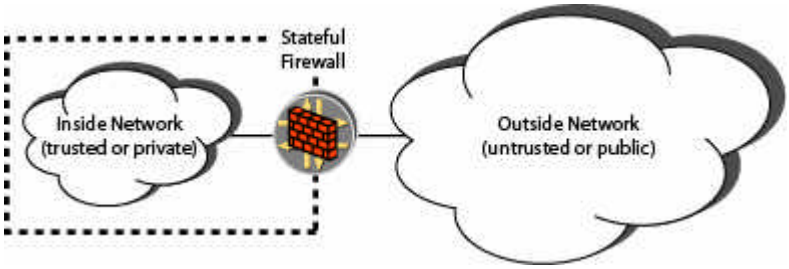

Figure 2: Figure 2: Stateful Firewall Service

Among the various features implemented by an operating system to enhance network security are:

- A router architecture that separates the control and data plane functions within each router and establishes a "virtual DMZ" between the disjoint control and data planes.
- Stateful firewall and Network Address Translation (NAT) capabilities to provide secure access links to the service provider's facilities.
- Virtual private LAN services (VPLS), Layer 2 virtual private networks (VPNs), Layer 3 (RFC 2547bis) VPNs, carrier-of-carrier VPNs, and inter-provider VPNs to support the logical separation of public and private traffic flows.
- Flow monitoring to support the sampling of traffic flows to examine them for suspicious activity.
- Discarding or rate-limiting the amount of traffic forwarded to specific destination prefixes to protect against Denial of Service (DoS) attacks.
- Unicast Reverse Path Forwarding (uRPF) to check and verify source addresses enabling providers to thwart address spoofing.
- The encryption of routing protocol traffic to eliminate bogus routing information

- The rate limiting and the stateful monitoring of all control traffic to protect against DoS attacks.
- Encrypted system administration, multiple user access levels, and centralized user authentication through RADIUS and TACACS+ to secure system administration when using the Command Line Interface (CLI).
- Configuration commit and rollback features to manage configuration changes and minimize errors.
- Event logging and audit trails to maintain records that identify who has accessed a router and what operations they performed during a given period of time.
- The ability to quickly add new hardware-based security features via specialized processors through the use of dedicated *service cards*.
- Flexible software licenses that allows a carrier to pay-as-it-goes and turn on new security features so it doesn't have to pay for features that it does not currently require.

JUNOS software provides a fully integrated, scalable, and uniform set of security tools for network operators. These tools have been purpose-built to run at both the edges and in the core of service provider networks, and on all high-speed interfaces. Our comprehensive approach offers network operators a foundation for delivering new revenue generating security services and service packages.

Multiple Platform Support

To successfully manage their software upgrade schedule, service providers require an operating system that is based on a single release train and that executes on multiple platforms. Support for a single release train eliminates the problem of trying to determine if a particular service-enabling feature is supported by a specific software release, hardware chassis, and line card combination. Providing a consistent set of features that execute on multiple platforms means that the services supported by the operating system run reliably in the core, at the edge, and at individual customer sites.

Unlike operating systems from other vendors, JUNOS software simplifies the mix of knowledge that is necessary to remain current with router code updates. The operating systems from other core router vendors operate on only a single platform, force providers to add yet another operating system to their matrix of operating systems, and do not simplify the management of the router software upgrade process. In contrast, JUNOS software allows service providers to leverage a single operating system that runs on core routers (T-series platforms), edge routers (M-series platforms), and CPE devices (J-series platforms). This approach simplifies network operations and provides service continuity. For example, the configuration of VPN services is consistent across all platforms.

Dependable Software Update Schedule

Service providers demand a reliable and consistent software update schedule to successfully manage growth and support the rollout of new services in their networks. Because software features and router hardware support are tightly coupled, router vendors must have the foresight to envision future carrier requirements and support those software needs in hardware. For example, if a router vendor had the vision to implement IPv6 forwarding in its processor hardware *before* carrier demand, the router vendor can simply implementing the IPv6 protocol suite in its operating system and deliver support for IPv6

traffic without requiring special builds, software branches, or hardware updates.

Unlike vendors that issue software updates on a haphazard schedule with missing or late features, JUNOS software updates are released on a regular schedule with the promised feature set. Service providers demand a reliable and consistent software update schedule to successfully manage growth and support the rollout of new services in their networks. Additionally, Juniper Networks has always had the foresight to anticipate future carrier requirements and provide early support for those needs in its router processor hardware. JUNOS software provides not only an extremely reliable and powerful operating system, but also delivers a feature-rich IP services toolkit. And the JUNOS software updates run on all T-series, M-series, and J-series routers to facilitate the seamless and rapid deployment of new services by supporting P, PE, and CPE functionality.

Engineering Discipline

Last, but certainly not least, a carrier-class operating system requires engineering discipline. Good engineering practices involve all aspects of router development including hardware engineering, software engineering, system test, release management, and product management. Once a solid architecture is developed and deployed, a router vendor must be willing to spend a considerable amount of time maintaining the architecture, making incremental improvements to the architecture, and cleaning up the implementation.

Unlike the operating systems developed by other vendors, JUNOS software has demonstrated its engineering discipline by steadfastly maintaining a single release train and having a long history of periodic enhancements. Many of these improvements are directly visible to service providers by the support for new features but many others are less obvious because they are designed to improve the operating system infrastructure or streamline the implementation. For example, some daemons have been rewritten, new algorithms have been implemented, and the kernel has been updated. Good engineering practices along with periodic enhancements are among the most important reasons for the success of JUNOS software.

The Benefits of JUNOS Software

JUNOS software was first released in 1998 with the launch of the M40 Internet backbone router. In fact, JUNOS software underwent extensive pre-release testing in a number of service provider networks to enable software certification before the M40 router hardware was available for shipment. The release of JUNOS software set the standard against which all future core operating systems are judged.

The IP core was very different in 1998 then than it is today:

- With the exception of JUNOS software (which had a modular architecture and was purpose-built for the IP core), all other operating systems had an enterprise heritage and a monolithic architecture.
- The number of hosts supported by large IP networks was a small percentage of the number of hosts supported today.
- Network speeds and feeds were significantly lower in 1998 than they are today. In 1998 service providers were considering migrating their core trunks to OC-12 / STM-4 (622 Mbps) or OC-48 / STM-16 (2.4 Gbps) while today's providers are considering the migration to OC-192 / STM-64 (10 Gbps) and greater speeds.

- IP-based voice and video applications were in their infancy so latency and jitter requirements were less critical than they are today.
- The Internet had not yet penetrated the public consciousness so security issues and the need to enforce strict service level agreements (SLAs) for mission-critical applications were of less importance than today.

Since 1998 the Internet has experienced tremendous growth causing the number of hosts, the speed of core trunks, application complexity, and the economic value and importance of IP services to increase. Throughout this expansion, service providers have required frequent routing software updates to provide new features and support the rollout of revenue generating services. Service providers have always been under enormous competitive and economic pressures to deploy the latest feature set to enhance the reliability and security of their network, satisfy subscriber demands for new services, retain existing customers, and attract new subscribers. JUNOS software has always demonstrated its engineering discipline by providing the on-time delivery of new features and production-ready implementations of new technologies such as:

- MPLS to provide the "circuit-like" delivery of IP services across provider-defined explicit routes.
- Traffic engineering and constraint-based routing to effectively use available bandwidth resources and avoid situations where some parts of a network are congested while other parts remain underutilized.
- QoS to ensure that mission-critical applications are given access to network resources during periods of congestion and other types of network stress.
- VPN services (Layer 2, Layer 3, carrier of carrier, and inter-carrier, VPLS) to enhance security by separating private data from public data that is transmitted over a common infrastructure.
- ATM and TDM interworking to allow providers to maintain their existing revenue streams and support the transport of these services over a converged IP infrastructure.
- Flow monitoring and DoS attack prevention to provide a secure networking environment.
- Generalized MPLS (GMLPS) to provide a common control plane between the optical transport and IP data layers resulting in faster and simpler provisioning, lower overall cost of operations, and improved capacity utilization.
- Logical routers so that customers can collapse many layers of the network and many service types into a single router.

Since 1998, Juniper Networks has taken the leadership role in developing the IETF standards that defined how these essential technologies should be implemented and deployed in service provider networks. JUNOS software and the features it supports have benefited from its years of operational experience in rapidly expanding service provider networks. The IP core was less complex in 1998 and the challenges of introducing a full-featured, stable, and scalable operating system were less demanding. Today, the IP core is significantly more complex and the technical challenges and economic risks are even greater. New operating system have not benefited from the experience of growing with the Internet and learning from that growth. Upon its first release, a new operating system has to remain stable, support all required features, scale, and perform appropriately in an extremely hostile environment. Within the service provider community "failure is not an option." Any false step can have a huge economic impact and severely damage a service

provider's hard earned reputation for delivering quality services.

JUNOS software has set the benchmark for high-performance core operating systems since 1998. It was the first operating system designed specifically for the carrier environments, runs on all Juniper Networks T-series, M-series, and J-series routers, and is currently deployed in the largest and fastest growing networks worldwide. Its full suite of industrial strength routing protocols, flexible routing policy language, and leading MPLS implementation efficiently scale to large numbers of network interfaces and routes.

Conclusion

While other router vendors are redesigning their core routing platforms and software to keep pace with service provider requirements, Juniper Networks is extending its technology leadership with JUNOS software. With today's IP core making a crucial transition from best-effort Internet to secure and assured Infranets, the operating system has never been under greater scrutiny. Juniper Networks design foresight and engineering discipline created an operating system capable of evolving and scaling with the growing demands of the IP core. And, in doing so, JUNOS software has and continues to benefit from the in-production experience and tuning it has received in the world's most complex networks. In the anatomy of a core operating system, there is simply no substitute for experience.