JUNOS TRIO

Programmable Silicon Optimized for the Universal Edge

Table of Contents

Table of Figures

List of Tables

Executive Summary

Junos Trio is a new silicon chipset developed by Juniper Networks based on an advanced network instruction set designed specifically to support a universal edge. This chipset dramatically improves business economics with its 3D scaling for bandwidth, subscribers, and services, enabling a truly converged platform with exceptional power efficiency. Its fully programmable engine makes it highly extensible while enabling service innovation with superior end-user quality of experience.

This white paper covers the motivations that drove the need for a new approach to networking instruction set design, the key elements and benefits of the Junos Trio chipset, and its applicability to next-generation networks and emerging applications.

Introduction

The increased demand for sophisticated, media-rich services, exponential growth of mobile sessions, and the emerging trend of cloud computing require a networking infrastructure that supports massive numbers of subscribers, service types and instances, and bandwidth. With the growing need to support unbounded applications, especially with the new open software models on consumer devices, it has become imperative to build infrastructures that are service independent.

The challenges are threefold.

1. Create an infrastructure that is multipurpose to support the requirements of emerging applications.

2. Exponentially increase packet processing performance and scalability.

3. Deliver exceptional quality of experience for all subscribers.

Juniper Networks has developed a breakthrough technology to fundamentally address these requirements. We call this new technology the Network Instruction Set Processor (NISP), and it is based on a new silicon architecture unlike traditional ASICs and NPUs. This new architecture leverages customized network instructions to maximize performance and functionality (Figure 1).

This new Junos Trio chipset using NISP technology is a revolution in network instruction set design. In the early days of computing, the battle was being fought over CISC versus RISC instruction sets for CPU technology. For the networking domain, there is no standard or off-the-shelf instruction set that can be used to meet the demands of high-performance packet processing. Juniper has taken the lessons learned from 14 years of developing programmable ASICs to design a new networking instruction set that sets the bar in the industry for programmability at performance and scale.

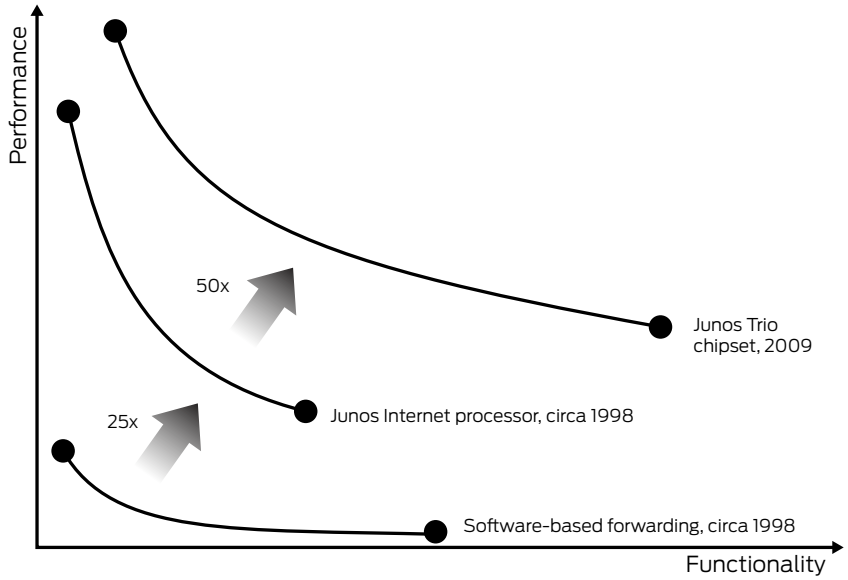

Figure 1: NISP architectural breakthroughs

Evolution of Next-Generation Silicon Chipsets

With the advancement of silicon technology, the parameters driving the embedded logic are being refined with each generation. The development of successful silicon depends on one key factor: getting the basic primitives right in order to ensure the longevity of the hardware architecture. In the area of silicon developed for packet processing, a few primitives are being optimized and new ones are emerging.

- One of the metrics that determines the performance of silicon is the number of **instructions per second**. For example, the multiprotocol lookup performance can be optimized based on the number of instructions it takes to perform the function. The ability to perform complex functions with the least number of instructions per second while also supporting nx100 Gbps throughput rates has been a challenge for most silicon implementations.
- The **number of transistors packed in a single engine** allows for more embedded logic that can be enabled for new features in forwarding software.
- Historically, silicon engines were specialized for specific target applications or types of functions. For example, there were separate engines designed for fabric switching, edge functions, core functions, service-based processing, and general-purpose computing. Finding the **right balance between programmability, scale, and performance** is vital for developing an engine for the service-enabled edge.
- The choice of silicon process being used in today's system plays an important role in achieving **power, thermal, and space efficiency**.
- **Embedded Layer 4 through Layer 7 packet processing** requires deeper inspection into the packet, not just the header. The flexibility and ability to do this inline and natively within the forwarding path, along with stateful flow-based functions and **service-aware QoS**, is being demanded by application-aware networking models.

Additionally, next-generation, silicon-based systems need to take into account recurring trends in networking requirements.

- Scale
- Exponential growth
- Automation
- More capability at lower cost
- Richer services with higher feature velocity
- Greater control
- Simpler management

Optimizing and refining silicon from one generation to the next is about a simple, but crucial process: *develop, deploy, optimize,* and *repeat,* something Juniper has been able to successfully adopt and execute on.

Junos Trio, built using 65 nm silicon process technology, plays a significant role in fulfilling each of the above requirements. This advanced chipset was built on the foundation of three generations (Figure 2) of chipset architectures with 78 chips in total that have been widely deployed across diverse applications.

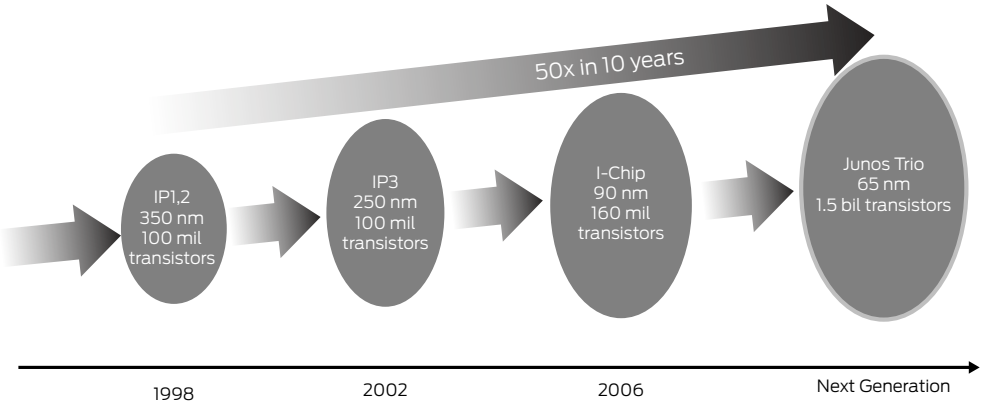

Figure 2: Juniper Networks silicon advancement

Junos Trio Architecture

Junos Trio chipset uses a highly modular and extensible architecture. Its building blocks can be flexibly put together to create highly optimized and cost-efficient systems for targeted applications. To achieve needed scale and flexibility, Junos Trio's architecture effectively uses the design principle of correctly partitioning the forwarding task into building blocks that are optimized by function. This partitioning is critical for balancing the size of each chip within the overall Packet Forwarding Engine (PFE).

Junos Trio Functional Blocks

There are four chips that constitute the Junos Trio chipset (Figure 3), each serving distinct functions that are tightly coupled with each other and with clean interfaces between them.

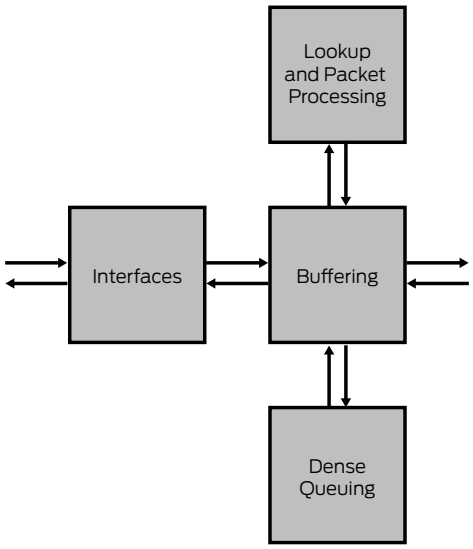

Figure 3: Junos Trio functional blocks

Lookup and Packet Processing

The Junos Trio lookup and packet processing engine, with its multicore and multithreaded design, uses a high level of parallelism. In addition to being able to do Layer 2 and Layer 3 lookups at scale, it is programmable to do Layer 4 through Layer 7 packet processing. The entire instruction set is new with its core primitives being much more general purpose. The lookup database is more compact than what was used in previous generations of chipsets.

Buffering

Juniper chose the option to use standard packet memory interfaces such as SRAM and DRAM that are long-lived and simple. Junos Trio chipset offers additional flexibility to use TCAM for optimizing firewall filters, as well as the capability to use optional DDR memory for flow table storage-based applications. The delay-bandwidth buffers can also be dynamically assigned for handling subscriber-level queuing with multiple levels of hierarchy and priority.

Dense Queuing

Dense queuing needed for hierarchical QoS has been optimized on a separate chip that can be optionally added according to the requirements of the application. The main attribute here is the programmability of schedulers, shapers, and buffer allocations that can be controlled independently of each other at multiple levels of the hierarchy. The entire QoS engine was designed with careful consideration for service-aware and hierarchical queuing models.

Interfaces

For external interfaces, Junos Trio is highly optimized for Ethernet, especially for the functions that make Ethernet carrier-grade and ready for emerging applications. Synchronous Ethernet support is built-in. Junos Trio enables the implementation of OAM at scale and has fully integrated Layer 2 Ethernet support. Junos Trio can also enable new technologies such as G.709, as well as legacy interfaces such as SONET/SDH, as required.

Junos Trio supports two operational modes: fabric and standalone, which are determined based on whether the internal interfaces of Junos Trio are connected to a switch fabric or not. These options provide the ability to create adaptable, highly integrated designs that are space- and power-efficient.

Delivering a Junos Trio-Based System

A Junos Trio-based system offers multidimensional performance, scale, and advanced features.

Performance

Junos Trio chipset allows the flexibility to provide throughput ranging from 30 Gbps to 120 Gbps full duplex. The initial line-card implementations with 120 Gbps of wire-rate performance support 16x10GbE interfaces, enabling industry-leading 10GbE densities with extremely low power consumption.

Scale

Junos Trio architecture allows 3D scalability for bandwidth, subscribers, and services (Figure 4). This 3D scaling allows you to build high-performance systems that address the broadband services edge supporting large numbers of fixed and mobile subscriber terminations, the business services edge with thousands of Layer 2 and Layer 3 VPN services, and the aggregation market with high-performance 10/100GbE dense configurations.

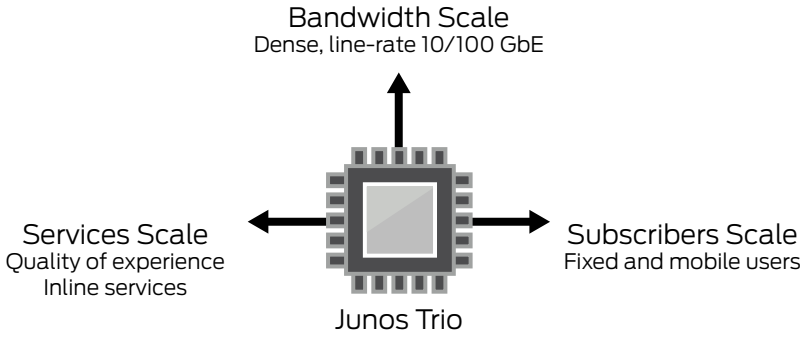

Figure 4: Junos Trio chipset with dynamic 3D scale

Junos Trio supports scalable subscriber management with up to 64,000 subscribers per slot, allowing over 2.3 million subscribers per rack using the currently shipping Juniper Networks® MX960 3D Universal Edge Router. Unlike other vendors' equipment, MX960 3D routers can support over 18.4 million fully usable egress queues per rack based on 256,000 egress queues per Junos Trio-based PFE. Junos Trio has evolved the per-subscriber QoS model to offer market leadership for high-count subscriber services with richer hardware instructions to more easily and quickly support new features.

Junos Trio provides an 8x memory increase over previous implementations to accommodate massive forwarding table entries, filters, policers, and counters. A major component of this innovation comes from the memory data structures that utilize storage in the most efficient manner possible.

Power Efficiency

Junos Trio is designed to offer a very efficient power and thermal system. With 16x10GbE line cards, power draw is as low as 37W/10GbE, creating up to 1.7x more power efficiency per 10GbE compared to similar competitive line cards (Figure 5).

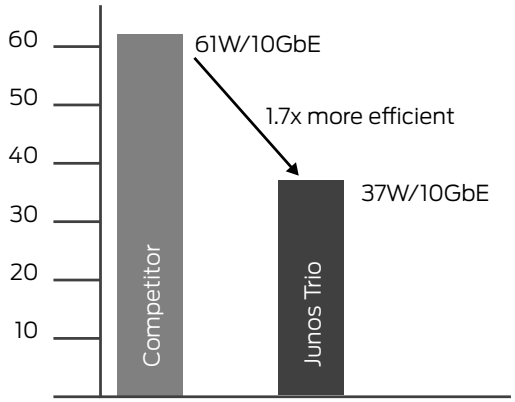

Figure 5: Power efficiency comparison

Advanced Features

A key innovation that Junos Trio offers is in the area of programmability, which accelerates the development of advanced features (Figure 6).

Programmable QoS

Junos Trio has a programmable QoS engine that creates a completely flexible processing pipeline to handle the sequence of QoS events. As an example, the queuing hierarchy can be built flexibly with the three parameters of minimum, maximum, and excess at each level.

This universal model can be leveraged for coarse, as well as fine-grained queuing, and it supports systems at the network edge for dense subscriber services, aggregation, business services, and other applications. Its *Dynamic Policy Engine* makes it extensible to a wide variety of intelligent applications. The related enhancements are significant and multidimensional.

- Large number of queues (256,000 per single PFE)
- Multiple levels of scheduling hierarchy with independent bandwidth guarantees at each level
- Up to five levels of queue priorities allow a single customer with multiple applications to be grouped to one node
- Fair portioning of bandwidth such that queues can use bandwidth if unused by other queues
- Intelligent oversubscription with lightweight input oversubscription mechanisms that conserve resources since they do not require full-blown input queuing
- Flexible hierarchical QoS mechanisms also available for host-bound traffic
- Optimal delay-bandwidth buffer
- Sub-30 microsecond latency
- Tunnel-aware, LAG-aware, and multicast-aware
- Granular statistics: per queue, per color, per drop decision

Junos Trio takes the hierarchical QoS model to the next level with advanced features (Table 1).

Table 1: QoS Features in Junos Trio

QoS FEATURES	DESCRIPTION
Stateless application detection	The ability to look deeply into the packet to detect applications and class of packets in a stateless manner.
Intelligent class-aware hierarchical rate limiters	The ability to both honor user-configured, rate-limit policies for multiple classes of traffic at the same time, and protect high-priority traffic from low-priority traffic bursts. These rate limiters can be applied to a variety of attachment points: ports, logical interfaces, arbitrary collection of interfaces, and a variety of user-configured policies.
Dynamic bandwidth profiles	The ability to arbitrarily group a set of interfaces and provide an aggregate shaping bandwidth control for them.
Class-aggregate bandwidth profiles	The ability to apply a policy to individual subscribers, as well as shape individual classes of traffic as an aggregate of all subscribers.
Dynamic priority protection	The ability to protect bandwidth of high -priority traffic even in the presence of bursty, low-priority traffic that has depleted a subscriber's bandwidth.

Inline Services

Traditional packet forwarding models implement a route lookup-based forwarding paradigm. However, with service processing, the matching criteria are far more diverse than just the IP destination address. The matching criteria might include other parts of the packet header such as fields from the Layer 4 header. Given the lookup criteria are complex, a flow-based architecture is required. Junos Trio chipset, with its integrated flow lookup capabilities, allows specialized service functions to be processed inline within the packet forwarding path without the need for external service modules. The following is only a subset of functions listed that can be done inline.

- Tunnel encapsulation and de-encapsulation
- Cflowd
- Carrier-grade NAT
- DPI

The programmable lookup engine opens up an array of options to perform other specialized packet processing functions demanded by emerging applications. The challenge is for the lookup engine to maintain high performance during service-specific processing. The next-generation Junos Trio services architecture addresses this issue by offering the option to use a hybrid approach where a combination of a separate service module and inline services can be used to minimize the performance vs. functionality trade-off. With this hybrid option, the service module can execute complex operations and offload standard forwarding functions to the PFE.

Based on the flexible pipeline in Junos Trio, there is no predefined sequence of operations between packet cracking, processing, and encapsulation. This creates an array of options to implement complex service processing tasks within the forwarding path.

Fully Integrated Layer 2 Functions

The Layer 2 functions are natively supported in Junos Trio and can be enabled at scale. For example, up to one million MAC entries can be enabled on a single PFE. Sampling of Layer 2 headers and native support for mirroring to multiple ports can also be enabled.

The more advanced features, like integrated routing and bridging (IRB) that enable support for both Layer 2 bridging and Layer 3 routing on the same interface, are now demanded of the next-generation edge routers. Features like IRB can be easily enabled at scale on a Junos Trio system.

Enhanced and Flexible Multicast

Junos Trio supports a distributed multicast architecture that enables enhanced scaling and fault tolerance. The dynamic replication structure supports very fast convergence. The user configurable QoS policies for multicast traffic provide greater control over the quality of experience. Similar to previous implementations, Junos Trio continues to

enable 10GbE line-rate multicast.

Faster Convergence in the Forwarding Plane

Junos Trio forwarding structures efficiently support indirection such that minimal structures are updated in the event of network topology changes. This allows much faster convergence during LSP reroutes, firewall filter changes, and LAG bundle failures.

Enhanced Load-Balancing Options

An enhanced hashing algorithm for load-balancing IPv4/IPv6, Layer 2, or even encapsulated traffic is possible based on the deeper lookup capability of Junos Trio. It also provides a symmetric hash for the forward and reverse flows, which is crucial for stateful services that have a requirement to see the forward and reverse flows for the same conversation. The symmetry in a hash calculation ensures that the direction of the flow does not influence the choice of traffic path.

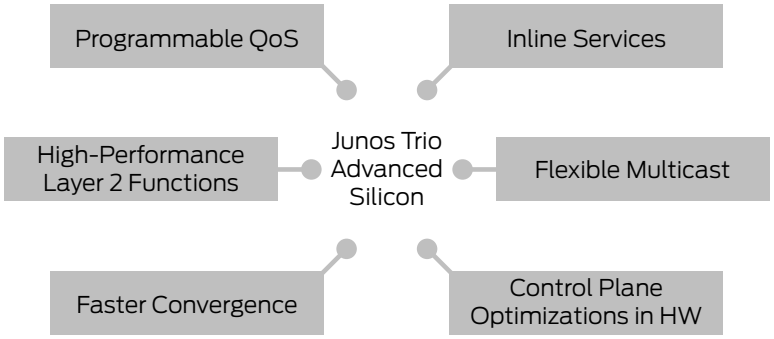

Figure 6: Junos Trio advanced features

Future-Proof Characteristics

An important element of a strong silicon-based architecture is its ability to support multiple generations of system designs and continuously adapt to new requirements created by emerging applications.

Programmable

Junos Trio's forwarding data structures are programmable, allowing for fast microcode changes in the hardware. The task of updating data structures at scale and implementing control plane functions for specific protocols in hardware becomes significantly easier in this new model. Keepalive processing can be enabled in hardware for all kinds of Layer 2, Layer 3, and services applications such as OAM, BFD, RPM, and others that need this functionality for enhanced scale and performance.

Extensible

There are several provisions carefully architected in Junos Trio that makes it extensible.

- Ability to look much deeper into the packet opens up tremendous potential for optimizing forwarding lookups, granular classification, sampling Layer 2 packets, and Layer 4 through Layer 7 processing as needed.
- Provisioned with large memory capacity to allow large volumes of scaling forwarding table capacity, policers, firewall filters, MAC entries, and counters as applications demand.
- No limitations on the number of labels that can be pushed onto the stack.
- Scalable beyond 120 Gbps throughput.

Modular

- Fabric and standalone operation modes enable multiple, customizable hardware configurations that are targeted for the network role and capacity requirements.
- Modular functional components of the architecture lend themselves to configurations that are optimized for different applications.
  - Forwarding and storage
  - Lookup and processing
  - Advanced per-subscriber queuing
- Different ratios of the above components to meet the requirements of the applications. For example, some applications require no per-subscriber queuing, and therefore a configuration needs to be optimized accordingly. For transport, the depth of packet processing is reduced, and so the forwarding engine can be light on processing performance in order to be optimized for cost efficiency. At the service edge where advanced services are required, more processing logic is a must.

Silicon and Software

Silicon design requires careful planning and foresight to ensure that the problems it is intended to solve are addressed at the root and that the solutions are applicable at a broad level. The trick is to ensure the basic primitives needed for the solution are embedded at the heart of the silicon and made programmable enough for general-purpose use. Junos Trio has fundamentally adopted this approach.

Juniper Networks Junos® software, which sits on top of the Junos Trio chipset, takes advantage of the massive flexibility in programming forwarding path functions, such as lookup, route updates for convergence, multicast and firewall filter processing, granular subscriber-based interface control, QoS functions, and many more.

The notion of programmability and general purpose provides more control and vastly improves the rate at which features can be enabled in the embedded hardware forwarding path. The ability to exploit this advantage in hardware with a mature and proven operating system that is consistent across all routers provides a powerful advantage for operators looking to add revenue-generating features reliably and at scale.

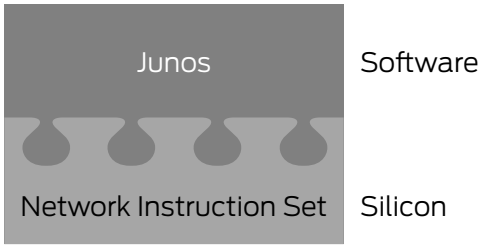

Figure 7: Junos software over Junos Trio

Junos Trio and Universal Edge

The universal edge is about increased interface density, optimizations for Ethernet, integrated services, scalable subscriber management, advanced QoS, high-performance multicast, and an unlimited service delivery platform.

Junos Trio is at the foundation of building the universal edge based on its capabilities to support massive scale, high-performance packet processing, inline services, and most of all, the ability to adapt to a constantly evolving feature set required at the network's edge.

With its leading capabilities to support large-scale subscriber termination, advanced queuing engine, and forwarding and services scalability across a common operating system, Junos Trio can truly enable business, residential, and mobile service convergence.

The Junos Trio Advantage

Breakthrough innovations clearly put Junos Trio far ahead of the competition in the area of service-enabled universal edge routing (Table 2). From scalability, to flexibility, to feature richness, Junos Trio gives you a solid option to reliably build large next-generation networks to sustain exponential growth. Moreover, the programmable element of Junos Trio creates a unique investment protection vehicle to accelerate and significantly broaden the possibilities of new services at scale.

Table 2: The Junos Trio Advantage

FEATURE HIGHLIGHT	JUNOS TRIO	VENDOR A	VENDOR B
State-of-the-art silicon	65 nm silicon process	90 nm silicon process	90 nm silicon process
Rich multi-tasking	16 cores with 320 simultaneous processes	Limited	Limited
Power efficiency	37W/10GbE	90W/10GbE	61W/10GbE
Forwarding performance	Line rate unicast/multicast, IPv4/IPv6, MPLS: 120 Gbps/slot	Limited throughput	Very limited multicast performance: 30 Gbps/slot
Massive scale			
Forwarding table routes	8 mil	512k	3 mil
Queues	512k/slot; 6+ mil/system	256k/slot; 3 mil/system	320k/slot; 3.2 mil/system
Subscribers	64k/slot, 768k/system	16k/slot	40k/slot; 64k/system
MAC addresses	1 mil/chipset	512k/system	512k/system
Programmable QoS with scale			
Flexibility	Advanced H-QoS with Dynamic Policy Engine	Limited	Limited
Subscriber queuing	64k subscribers per slot with up to 8 queues per subscriber	16k interfaces/slot	Separate interface module required for dense subscriber queuing; 20k subscribers/module
Granularity	Independent shaping at every hierarchy	Limited	Limited

Conclusion

Junos Trio is the industry's first Network Instruction Set Processor. It combines the proven performance of ASICs with the flexibility of general-purpose processors. This unique combination takes edge routing to the next level of network intelligence and sophistication while delivering exceptionally high performance and efficiency. Junos Trio is perfectly poised to take on the complex requirements of emerging applications for many years to come.

Acronyms

ASIC	application-specific integrated circuit
BFD	Bidirectional Forwarding Detection
CISC	complex instruction set computer
DDR	double data rate
DPI	deep packet inspection
DRAM	dynamic random access memory
IRB	integrated routing and bridging
LAG	link aggregation group
LSP	label-switched path
MAC	media access control
MPLS	Multiprotocol Label Switching
NAT	Network Address Translation
NISP	Network Instruction Set Processor
NPU	Network Processing Unit
OAM	Operation, Administration, and Maintenance
QoS	quality of service
PFE	Packet Forwarding Engine
RISC	reduced instruction set computer
RPM	real-time performance monitoring
SRAM	static RAM
TCAM	ternary content addressable memory
WRED	weighted random early detection

About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at **www.juniper.net**.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

2000331-001-EN Oct 2009