Distributing the Converged Supercore: Scale Out with the World's First Fixed Configuration Core Router—Juniper's Perspective

A New Level of Scale and Cost Control Dispersed Across the Core

The Internet puts a wealth of information in the palms of our hands, transforming lives in incredible ways. For example, education is no longer available to just a privileged few; online services like Khan Academy provide world-class curriculum in math, science, economics, arts, and much more to anyone, anywhere around the globe. Medical breakthroughs can be reported and shared in near real time; video collaboration enables a botanist in the Brazilian rainforest to learn about new life-saving antibiotics from lab technicians in Montreal. Cosmological discoveries can occur at ever greater frequencies, as astronomers in Hawaii and France create sophisticated new 3D computer-generated maps defining the Laniakea supercluster's home in the Milky Way galaxy. The creativity and ingenuity of the human spirit is limited only by its imagination; what was once thought impossible is now routine, thanks to the constantly evolving Internet.

The volume of Internet traffic, however, combined with a demand for faster, more consistent application response times, is pushing current architectural limits. Larger and denser peering points are no longer the only solution. Providers are looking to solve congestion problems by increasing the number of peering points dramatically to deal with the data deluge. Evolving in unseen ways, however, peering points are no longer isolated to major Internet exchanges. In the early days, access to content was sparse; the entire scope of the Internet was extremely small, with very few network peering points. These peering points are "voluntary" interconnections of administratively separate Internet networks, established for the purpose of exchanging traffic between each network's respective users. Today, thanks to PeeringDB.com, the Internet is a vast, World Wide Web with more than 556 registered exchange sites supporting multi terabits worth of capacity. Larger exchange sites support more than 700 participants with a mix of public and private peering options. Every new idea leverages the vast network of Internet Exchange facilities, giving rise to new peering points and moving ever closer to achieving a simple mission: improving the human-to-human experience across the Internet while making way for new innovations.

Today, peering across the Internet is not only more ubiquitous, it is also becoming more distributed. The number of public and private peering participants has increased as content and cloud-based services are consumed at an exponential rate.

Both content providers and service providers have a profound impact on how users around the globe exchange ideas and information, whether for work, learning, or play. As new technologies emerge, network speeds increase; as the Internet penetrates unconnected population segments, this impact and influence will grow even greater.

The Converged Supercore: Power Dispersed at Scale

Trusted equipment vendors must break the mold and establish a long-term goal of reducing cost pressures and increasing profits while alleviating traffic deluge. Juniper's vision is rooted in three simple foundational principles:

- **Scale** the infrastructure elastically
- **Automate** to eliminate operational complexity
- **Create** innovative, value-added services

While easy to articulate, these foundational principles are hard to deliver. They require engineering innovation and commitment at different altitudes—a challenge Juniper welcomes every day, to earn a place as a trusted partner.

Scale

Content providers demand a high-performance network infrastructure that can scale both up and out to ensure a positive end-user experience while allowing them to manage huge amounts of machine-to-machine traffic. This requires pioneering core router innovation, evoking new scale-out architectures built with high-performance, purpose-built IP/MPLS silicon that includes the elegant craftsmanship needed to facilitate deployability anywhere and bring it all together.

- **Scaling Out the "Distributed Core:"** Core architectures—the backbone of a particular network—have historically only scaled in one direction: up. Interface speeds on core routers available in the early to mid 1990s started at a measly 155 Mbps, delivering a total system capacity of just a few gigabits per second. A mere 21 years later, core router interface speeds have reached 100 Gbps, with system capacity skyrocketing to a mindboggling 24 Tbps in a single chassis. This scale-up approach has resulted in a system architecture that is highly resilient to both hardware and software failures, enabling pay-as-you-grow scalability and broad interconnect flexibility to address all manner of core architectures. Content providers are looking for the freedom to scale their core backbones on the x-axis—another architectural option that was once unavailable.

  The Juniper Networks® PTX1000 Packet Transport Router breaks this mold, enabling content providers to cost-effectively scale out their core architecture across many sites, distributing the interconnections between many smaller routers without sacrificing overall network performance, manageability, or availability. The ability to maintain a consistent operational environment provides a congruent deployment experience that existing users are accustomed to with the PTX Series Packet Transport Routers family.

- **Scaling Out "Distributed Peering:"** Insatiable consumer and business demand for more and more content-rich services are creating unpredictable traffic patterns over service provider networks. Content providers are distributing their peering points further out in the network, helping to isolate traffic spikes in a particular region and simplifying intense traffic engineering modeling exercises that normally take months to implement. Service providers gain additional benefits by distributing their peering points closer to the end user, improving the overall content experience by shortening response times and reducing content contention over fixed and mobile access.

As more enterprises migrate toward the "cloudification" of corporate applications like Microsoft Exchange to Office 365, service providers are dispersing their peering points with Microsoft closer to the end user. These distributed peering points create a transit portal between the service providers and Microsoft, isolating Office 365 traffic and confining it on a direct path to the Microsoft public cloud. The number of hops between the enterprise and the cloud-virtualized applications is drastically reduced, ensuring a consistent user experience as they migrate from an on-premise environment to a cloud-based one.

- **Silicon Scale:** Ultra-high network performance begins with advanced silicon—it's the atomic element from which the network is created. Unlike commodity or even merchant silicon, custom silicon is highly tuned and optimized to handle IP and MPLS traffic. "A designer knows perfection is achieved not when there is nothing left to add, but when there is nothing left to take away," said Antoine de Saint-Exupery. Diligent design optimizes the types of transistors, memory, and functions required to deliver the highest quality, lowest latency, fastest throughput, and lowest power draw in the tightest space possible.

  The all-new Juniper ExpressPlus™ ASIC is the latest example of Juniper's breakthrough silicon innovations. The ExpressPlus silicon features 6 billion transistors in a 28 nm configuration, 3D memory technology that reduces space by 20 times, and 1.6 billion advanced filtering operations per second for a whopping 500 Gbps of low latency, low power, IP/MPLS packet-processing throughput.

- **System Scale:** Silicon innovations that push the boundaries of performance and efficiency can't solve the problems of content providers and service providers alone. The networking platforms that use Juniper ExpressPlus silicon—the PTX5000, PTX3000, and the new PTX1000—also require intricate system design in order to deliver meaningful performance, economics, and ease of deployment.

  The PTX1000 provides all the same benefits as the PTX5000 and PTX3000 Packet Transport Routers in terms of performance and ease of deployment. Providing 3 Tbps of capacity, the PTX1000 achieves the same industry-leading efficiency as the PTX5000 and PTX3000: 0.55 watts per Gbps. With the addition of the PTX1000, providers get a consistent operating experience no matter how small or how large the application. Combining the industry's gold standard of efficiency with superior IP/MPLS performance, the PTX1000, at just 2 U in height, is a core router that uniquely unlocks additional applications such as distributed core architecture and distributed peering in Juniper's Converged Supercore® architecture. As new content continues to push ever growing traffic demands, content providers need to maximize their capital investments for today's and tomorrow's architectural requirements. The PTX1000 provides elegant engineering to facilitate seamless transitions from 10GbE to 40GbE interconnects, and even allows operators to grow into 100GbE interconnection to maximize investment protection and ensure smooth upgrade cycles.

  Juniper doesn't stop there, however. While these physical innovations are unparalleled, they are only part of the entire solution designed to control costs. Virtual innovations in SDN and automation extend cost control beyond individual platforms, preparing the entire network for optimization.

Automate

Networks must be efficient and simple to operate. Operational complexity plagues operators and slows any material cost reductions across the network. Virtual innovations bring SDN intelligence to complement physical strength and deliver network-wide optimization. The manner in which the service provider interacts with the network is transformed, reducing long planning cycles into real-time control. When what was once manually provisioned becomes automated, operators can eliminate unforeseen errors and emergencies, drastically increase network utilization, and optimize the network across layers to bring down costs.

For example, with SDN control, service providers can increase traffic utilization by as much as 35 percent, making the most out of current investments. They can also reduce unnecessary capital spending caused by redundant overprovisioning by as much as 40 percent.

With Juniper Networks NorthStar Controller, network changes can—for the first time—be made with the push of a button. NorthStar Controller uses open, standards-based protocols such as BGP-LS, NETCONF/YANG, Path Computation Element Protocol (PCEP), and REST APIs for third-party integration. Network operators can simulate, model, and implement changes with precision. Service providers can automate seamless network upgrades without fear of human error, SLA impact, and negatively influencing the end-customer experience. Rather than over-engineer the router's hardware and software to implement hitless upgrades, SDN can be used to remove any router from the network resource pool and automatically redistribute traffic across the remaining routers, eliminating node maintenance windows.

Furthermore, shifting transport topologies—even in multivendor environments—can be managed with certainty, making network upgrades and traffic optimization easy. The ability to predict the future means costs can be optimized, but many other amazing possibilities are also on the horizon. The network becomes flexible so that new services can be introduced much more rapidly and with lower investments than ever before.

Create

Juniper's Converged Supercore architecture enables both content providers and service providers to create virtualized services anywhere in the network by leveraging physical innovations to match traffic demands. Additionally, virtual innovations can be employed to precisely control traffic network-wide without sacrificing the services experience.

Conclusion: Expanding Core Scale Options While Driving Cost Down

Physical and virtual innovation requires dedicated engineering to holistically increase network scale up and out while driving costs down. The Converged Supercore from Juniper—networks that know how to scale up and out while driving costs down—liberates both content providers and service providers from the status quo of "one size fits all" core architecture. It does this by scaling capacity in multiple directions to efficiently handle increasing traffic levels with faster application response times, giving users a consistent experience across all access media. Users now have the freedom to scale out without sacrificing performance or efficiency, enabling the creation of wonderful new opportunities for service creation with newfound elasticity.

It's about time your network scaled as fast as your business. It's time you deployed the most efficient and easy to use core router available. It's time you unchained yourself from the burden of manual network provisioning. It's time for a Converged Supercore architecture from Juniper Networks.

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at www.juniper.net

Corporate and Sales Headquarters
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or +1.408.745.2000
Fax: +1.408.745.2100
www.juniper.net

APAC and EMEA Headquarters
Juniper Networks International B.V.
Boeing Avenue 240
1119 PZ Schiphol-Rijk
Amsterdam, The Netherlands
Phone: +31.0.207.125.700
Fax: +31.0.207.125.701

3200049-001-EN June 2015