

CLOUD SERVICES AND CLOUD INFRASTRUCTURE: THE CRITICAL ROLE OF HIGH-PERFORMANCE NETWORKS

Over the past 40 years, the evolution of computing paradigms has typically been fueled by three driving forces: 1) The desire to deliver new services for businesses, consumers and/or governments; 2) The need to change the economics of how those services are delivered or used; and 3) The desire to improve the user experience of how those services and applications are received, consumed and shared.

As we approach the second decade of the 21st century, a new paradigm has emerged that has many in the computing industry all atwitter. And no, it's not just Twitter. Or Facebook. Or Salesforce.com. Or even Google. It's the underlying phenomenon of cloud computing.

Of course, much has been said and written about cloud computing already. As with all new megatrends, the hype grows with each passing day — and the “trough of disillusionment” gets deeper — making it difficult to discern the true definition and core drivers of the trend. The boundaries of what it encompasses are often a little, well, cloudy. So let's start by defining some terms...

Cloud computing and cloud services

For the purposes of this paper, we believe it's most useful to our customers and communities to build on the definitions put forth by Gartner Group, IBM, Amazon and other industry leaders.

In its simplest terms, cloud computing is a new mode of computing where technology-enabled services are delivered to a set of consumers from a set of providers over the network. The definition is inclusive of compute, storage, network, application and service elements, typically running in large data centers accessed by a variety of mobile and fixed devices. Unlike prior computing systems that were essentially dedicated to individual businesses or applications, cloud computing's infrastructure must be shared across multiple applications, users, departments or businesses in order to achieve economies of scale. Sharing may occur in private, public or hybrid cloud architectures, depending on the target audience, usage and security requirements.

While cloud computing clearly depends on a lot of powerful technology, the services being delivered are ultimately what matter most. Cloud services can take a variety of forms, ranging from the delivery of technology services themselves up to the delivery of applications, content or process services. Amazon presents an interesting example:

- **Application** – They provide the ability to buy books and other merchandise online via a transaction application that runs on Amazon's platform and infrastructure.
- **Platform** – They have enabled businesses and consumers to create their own applications, written via Amazon Web Services to run on Amazon's internal development platform and infrastructure.

- **Infrastructure** –They offer their infrastructure for use by businesses and consumers on demand, via Amazon’s EC2 (compute) and S3 (storage) services.
- **Process and Content** –They will also host a customer’s online store (application and content) while providing access to the Amazon customer database (platform) and marketing engine (content). The Amazon system will also provide payment handling and fulfillment capabilities (process).

Cloud Service Types

Process
Content
Application
Platform
Middleware, DBMS, shared application service
Infrastructure
Networking, Compute, Storage, Facilities

Cloud services are defined by vendors in a variety of ways, including “software as a service,” “platform as a service,” and “infrastructure as a service.” The revenue for providing services can be generated via direct subscription fees (typical of B2B services); pay-per-use fees (which can be paid directly by consumers or business, or transaction fees charged to sellers); advertising (typical of Internet services); internal chargeback mechanisms (typical of private providers); or subsidized delivery (typical of government, educational or charitable organizations).

Cloud Infrastructure: Service Production and Service Delivery

In the physical and digital worlds, infrastructure is often an unsung hero that makes it possible to produce and deliver services. As The New York Times Magazine recently wrote: “We have an almost inimical incuriosity when it comes to infrastructure. It tends to feature in our thoughts only when it’s not working. The Google search results that are returned in 0.15 seconds were once a stirring novelty but soon became just another assumption in our lives, like the air we breathe. Yet whose day would proceed smoothly without the computing infrastructure that increasingly makes it possible to navigate the world and our relationships within it?”

Cloud Infrastructure includes a broad array of computing, storage, networking, applications and other technology that reside in corporate data centers and service providers. It serves two primary functions:

1. **Service Production:** Where applications and services are developed, deployed, provisioned and managed, typically in data centers that serve as the “engine room” or “production facility” of the cloud; and
2. **Service Delivery:** Where applications and services are delivered, accessed and shared by users, typically via the global wide area network that serves as the “delivery channel” of the cloud.

Cloud computing is distinct from traditional “hosting” services because it must offer “multi-tenancy” – the ability to deliver personalized or partitioned services running a common, shared

infrastructure. Cloud infrastructure also has to run multiple applications and services simultaneously; deliver fast access across multiple network protocols to a broad array of fixed and mobile devices; and adapt to real-time demands from businesses and consumers. Given that environment, we believe the essential requirements for cloud Infrastructure are:

- **Scale:** The ability to scale applications and computing infrastructure to support large customer populations – often in the millions of users, not just the thousands of users that most corporate data centers support – without adding complexity to the infrastructure.
- **Efficiency:** The ability to consolidate, standardize, virtualize and/or share infrastructure to reduce or minimize the costs of providing services at the required scale. This has always been a focus for data center managers, but becomes increasingly critical given the enormous scale and flexibility involved with cloud services.
- **Elasticity:** The ability to automatically provision and dynamically reassign resources (compute, storage and network) to meet peak or changing application loads/demands. In the past this was achieved by static overprovisioning of resources by individual applications. In the cloud world, dynamic allocation enables new services to be delivered faster by quickly adding or decreasing infrastructure resources to meet acceptable service levels on demand.

The final distinction for cloud Infrastructure is how users “experience” the cloud services they consume. In general, users select their provider and may choose the quality of service they want to experience, but they generally don’t see or know what’s required to generate and deliver the service. This is the same way most people experience the electrical power grid. On the other hand, providers are very conscious of the infrastructure as it is their responsibility to build, fund, and manage it. Thus while “cloud” describes the consumer’s experience, “computing” describes the provider’s experience.

High-Performance Networking = The Backbone of Cloud Infrastructure

The continued explosion of services, richness of digital media and appetite for real-time information will continue to drive demands for better and faster networks. Within cloud computing, networking plays an even more vital role as the backbone for both the data center (enabling service production) and the global network (enabling service delivery). Put simply: **without the network, there is no cloud.**

Within the data center, networks are quickly becoming the change agent as virtualization, consolidation and standardization dominate the agenda. Servers are increasingly being pooled, abstracted and virtualized using network load balancing, virtual machines and high-performance computing grids. The same is true for storage. Data is migrating back to the data center where it can be professionally and efficiently managed, while additional copies of data are maintained temporarily in the client or the network to optimize the user experience.

For service delivery, the role may be more obvious since the global high-performance network is the delivery vehicle. In reality, the global network plays multiple roles as a delivery and access vehicle from users to the cloud, between data centers and even between different clouds. Concentration of resources in larger data centers enables more efficient resource pooling and management, but that only works if the global network can deliver the goods. Consumers and clients are becoming more mobile and expect the same user experience regardless of their location or the device they use to access services.

Given the need to drive economies of scale while delivering rich user experiences, cloud infrastructure places a premium on the networks involved. If the cloud services are valuable – to the provider and/or the consumer – any old network won't do. What's needed is a high-performance network – one that is reliable, secure, fast and scalable yet simple:

1. **Scalable:** The ability to grow and connect large numbers of endpoints, any to any.
2. **Fast:** The ability to support any communications application or service, addressing both bandwidth and latency.
3. **Reliable:** The network's ability to function and deliver services despite failures, whether they are natural, accidental or malicious.
4. **Secure:** The network's ability to protect itself and the information it carries.
5. **Simple:** Taking complexity out of the customer environment, and automating it into the network itself. This is critical to making the technology and the economics work.

Juniper's Strategy and Capabilities for Cloud Computing

Juniper Networks has a common vision for helping customers embrace cloud computing via both the data center and the global network. Since cloud computing is not an overnight move, we believe customers should explore several steps to pursue their own cloud-like infrastructure.

Service Production: Networking the Data Center

Juniper is applying its expertise from the high-performance global network to help customers rethink the basic data center architectures that were built over the last 20 years. We are guided by a vision of a single "logical switch" in the data center. We want to take the simplicity and low latency of a single switch and scale it reliably and securely across a data center.

Today we help our customers move from three or more tiers to a simplified two-tier model. This reduces the number of elements in the network resulting in reduced equipment costs; reduced management complexity and operational costs; dramatic improvements in network latency and consistency of performance; and reduced power, space, and cooling costs.

We then integrate the simplified switching environment with virtualized security services, powerful routing capability to connect across and between clouds, and a single "pane of glass" to

manage the network. The end result is a better user experience delivered at a low total cost of ownership.

In the future, Juniper's "Stratus" project will deliver a flat single-tier, non-blocking, converged data center fabric. It is being designed to provide the inherent simplicity of a single switch and yet scale across even the largest cloud data centers. It aims to provide the perfect foundation to dynamically connect the virtualized infrastructure of the modern data center and deliver the economies of scale promised by cloud infrastructure.

Service Delivery: Powering the Global Network

The seamless delivery of applications and computing resources 24x7 for users around the globe demands high scale, performance, reliability, and manageability of the underlying network. An always-on infrastructure is fundamental to cloud computing. Adjusting network resources dynamically to traffic volume, priorities, users and applications creates elasticity, while providing control and manageability. To enable this, Juniper's routing portfolio has reliably connected the world's largest networks with sophisticated network virtualization, intelligent services, and application and user level awareness.

- **Network virtualization:** Juniper provides the most sophisticated network virtualization capabilities in the industry, offering flexibility, scale and manageability to users. Virtualization capabilities are offered in hierarchies, which customers may deploy at any level of granularity.
- **Carrier-class reliability:** Through technologies such as non-stop forwarding, non-stop routing, In-Service Software Upgrade, rapid virtual path failovers, fast re-route and graceful restarts, Juniper ensures high availability is at the foundation of cloud infrastructures.
- **Intelligent services:** Juniper offers services to identify and uniquely treat traffic based on a wide variety of contextual criteria that include (alone or in combination) user and session identity; privilege, application and service requirements; and network and security state.
- **Application and user level awareness:** Juniper's edge routers and software help to define and automate policies that govern user authentication, security, bandwidth usage, priority settings, transactional needs, and service level agreements to ensure quality of experience.

Service Security: Securing the Cloud

Juniper offers a robust portfolio of security solutions that can help secure data flows within and between clouds, as well as directly to users. The solutions provide a distributed suite of protections operating at key points of vulnerability and at layers where network-based security adds to the quality and integrity of the services being delivered.

- Within each data center, we complement the traffic partitioning supplied by VPNs and VLANs in the network with corresponding security policies and enforcement mechanisms in the world's fastest firewall and intrusion detection platform, Juniper Networks SRX Series Services Gateways.

- Between data centers and throughout the network infrastructure, high-performance protection is deployed at the gateways to multiple internal data centers of the cloud infrastructure – anywhere from two to dozens of sites linked together in a distributed resource pool.
- At user access points to cloud services, central definition and distributed enforcement of cloud security policies is applied. Users and cloud providers alike benefit from a unified configuration management system enabling a diverse range of authentication, access permission, and integration into the forwarding infrastructures designed for carrying the cloud's multiple services (e.g., infrastructure as a service, software as a service, etc.).

Why Juniper for Cloud Computing?

Juniper is uniquely qualified to provide the network infrastructure for cloud computing. Based on our breadth of experience across service provider and data center networks, and our depth of innovation across silicon, systems and software, we have proven our ability to tackle the tough challenges at the economies of scale involved with cloud computing. From the network foundation to application and end-point enablement, we strive every day to ensure that technical innovation and lower cost of operations are not mutually exclusive experiences.

Juniper delivers its cloud computing technologies through an integrated portfolio of routing, switching and security products, fueled by a single powerful OS—Junos® operating system. Juniper's high-performance networks scale to meet the most demanding challenges of the Internet while presenting the highest levels of reliability and security to enable the constant delivery of service across the cloud. In the data center, we help customers build and manage a simple fast network that dynamically enables the "production facility" for cloud services. In the global network, we provide a simplified, virtualized network that powers shared infrastructure for scalable, efficient and elastic cloud services delivery.

From an economic standpoint, our elegantly engineered solutions deliver tangible benefits at an inherently lower total cost of ownership. Virtualization at the Juniper core of the cloud can provide 39% TCO savings over existing infrastructures.¹ Additionally, virtualization technologies and reduction in footprint and consumption are common Juniper features that make energy efficiency a measurable benefit of Juniper's cloud computing platform. In fact, our current data center solutions significantly reduce network complexity and TCO by up to 52% in capital expenditures, up to 44% in power, up to 44% in cooling, and up to 55% in rack space in the data center.²

For more information and updates on Juniper's cloud computing capabilities, please see Juniper's website at www.juniper.net/cloud or email us at cloud@juniper.net.

NOTE: Juniper's innovations have defined high-performance networking since 1997, and the company has become one of technology's fastest-growing stars. This position paper is intended to provide customers, partners, industry analysts, investors and media with a guide to Juniper's point of view on cloud computing, and a snapshot of Juniper's capabilities to help customers embrace the trend.

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.

Reference Citations:

¹ Business Case for Core Network IP Virtualization. Network Strategy Partners 2009

² Juniper competitive analysis, published November 2008, <http://www.juniper.net/datacenter>.

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.

Copyright 2009 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Junos, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. Junos is a trademark of Juniper Networks, Inc. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.