

# CLOUD-READY DATA CENTER REFERENCE ARCHITECTURE

## Table of Contents

Introduction .....	4
Scope .....	4
Framework .....	4
The Importance of Data Centers and Their Infrastructures .....	4
A Data Center by Any Other Name .....	4
Supporting Enterprise and Cloud Data Centers .....	5
Solution Profile Overview .....	5
Key Trends in Today's Data Center .....	6
Evolving Business Application Architectures .....	6
Server Virtualization .....	6
Reducing OpEx .....	6
Protecting Against Security Threats .....	6
Convergence of Fibre Channel and Ethernet Networks .....	7
Requirements .....	7
Functional Areas in the Cloud-Ready Data Center .....	8
Juniper's Approach to a Cloud-Ready Data Center .....	9
Network Infrastructure .....	9
Application Traffic Flows .....	10
Simplified Network Infrastructure .....	10
Distributed Data Centers .....	10
Juniper's Approach .....	11
Compute and Storage Infrastructure .....	11
Integrating Virtual Server Infrastructure .....	12
I/O Convergence .....	13
Fibre Channel and FCoE .....	14
Juniper's Approach .....	14
Services .....	16
Security Services .....	17
Application Services .....	21
Juniper's Approach .....	22
Management, Orchestration, and Automation .....	23
Profile of an Effective Orchestration Platform .....	24
Management Infrastructure Supporting Cloud-Level Orchestration .....	25
Juniper's Approach .....	26
Junos Space Juniper's Open Network Orchestration Platform .....	27
Automation Based on Junos OS .....	28

Data Center Types .....	28
Transactional Production Data Center Network .....	28
Content and Services Hosting Production Data Center Network .....	29
High-Performance Compute (HPC) Production Data Center Network .....	30
Enterprise IT Data Center .....	31
Small and Midsize Business IT Data Center .....	32
Conclusion .....	33
Appendix A Juniper Products for the Cloud-Ready Data Center Network .....	33
Switching .....	33
Routing .....	33
Security Services .....	34
Application Services .....	34
Operating System .....	34
Unified Network Client .....	34
Orchestration and Network Management .....	34
Junos Space Platform .....	34
Technical Services .....	34
About Juniper Networks .....	35

## Table of Figures

Figure 1. Data center reference framework .....	8
Figure 2. Reference architecture network infrastructure .....	10
Figure 3. Compute and storage .....	12
Figure 4. Consistent management of the physical and virtual network from Junos Space .....	15
Figure 5. Services functional area .....	16
Figure 6. Flow types in the new cloud infrastructure .....	18
Figure 7. Management, orchestration, and automation .....	23
Figure 8. Juniper Networks management infrastructure .....	26
Figure 9. Junos Space infrastructure .....	27
Figure 10. Content and services hosting production data center network .....	29
Figure 11. High performance compute production data center network .....	30
Figure 12. Enterprise IT data center network .....	32

## Introduction

The data center is an essential corporate asset that connects all servers, applications and storage services. Businesses rely on their data centers to support critical business operations and drive greater efficiency and value. As such, the data center is a key component that needs to be planned and managed carefully to meet the growing performance demands of users and applications. Juniper Networks offers a comprehensive data center network solution that combines best-in-class products with well-defined practices to build high-performance, robust, virtualized and cost-effective data center networks. This reference architecture proposes practices, technologies and products that help data center architects and engineers who are responsible for answering the requirements of designing modern data center networks that support business goals.

## Scope

This document introduces Juniper Networks architectural model and its offerings in support of data center and cloud computing networks. The purpose of this reference architecture is to communicate Juniper's conceptual framework and architectural philosophy in creating data center and cloud computing networks for our customers.

This reference architecture is intended for the following personnel:

- Customers in the enterprise and public sector
- Service providers
- Juniper partners
- IT and network industry analysts
- Individuals in sales, network design, system integration, technical support, product development, management, and marketing who have an interest in data center design

## Framework

### The Importance of Data Centers and Their Infrastructures

Data centers run the applications that deliver business processes and services. These applications provide critical information and rich, differentiated content for users. Users now demand an agile, responsive infrastructure that provides exactly the access that they need. This can be 24x7x365 for services that must be “always on and accessible from anywhere,” or a series of scheduled updates set to meet user needs for time-based information (hourly, daily, weekly, monthly, or quarterly).

For innovators and technology suppliers such as Juniper Networks, data center networks are central to the business mission, providing the focal point for solutions that unlock value in unique and compelling ways for businesses and their users.

### A Data Center by Any Other Name

Not all data centers are the same. Their use, size and design vary with the needs of the business and the results that must be achieved. Examples include:

- Online transaction processing centers meeting strict transaction time constraints and carrying financial obligations with transaction results (exchange trading platforms, online financial services, online retail sales)
- Multimedia content delivery with strict quality and consistency requirements (online entertainment and news, video conferencing, live meetings)
- Computationally intense workloads (homeland security, logistics and production control, flight control, scientific research and economic modeling)
- General enterprise-grade operations data processing (CRM, ERP, human resources, finance, and messaging/communication)
- Cost-effective, reliable and manageable data center infrastructures for basic business operations

Requirements vary widely across these data center types. Some demand the lowest possible latency and highest possible availability, while others require comprehensive attention to quality of service (QoS), scale and high availability (HA). Still others require cost minimization by opting for less sophisticated availability and robustness.

### Supporting Enterprise and Cloud Data Centers

As we have seen, data centers have grown to serve a wide range of business needs, and there are many factors to consider when designing a solution that meets different objectives. Within the past several years, a powerful new paradigm has emerged that has important implications for data center architectures and how they meet these varied objectives. This is the paradigm of cloud computing.

Cloud computing delivers services dynamically over networks from an abstracted set of resources. The resources are somewhere in the cloud and available on demand. The types of resources and their location are transparent to end users. End users primarily care that their applications, data and content are secure and available, with a desired level of quality.

From the infrastructure perspective, cloud computing heavily leverages resource pools in a variety of technologies—compute, storage and network—for dynamic allocation in an automated, orchestrated and logically diversified environment, accommodating a variety of applications. Using orchestration, resources can be pooled within and across multiple data centers to provide an environment that responds dynamically to user needs.

Many organizations have started adopting elements of cloud computing, although the form varies widely. Cloud models include:

- Private clouds, in which clouds deliver services to users and groups within an individual organization (business or government/public sector entity)
- Public clouds, in which services are delivered over publicly available networks to a range of users for a variety of applications, depending on the cloud provider's offerings
- Hybrid clouds, in which an organization blends the capabilities of its private cloud with those of a public cloud to deliver a composite services menu

The cloud-computing paradigm can be applied to data center network designs to meet a variety of business and application requirements, with differences that emphasize specific criteria in support of business goals.

### Solution Profile Overview

Reference Architecture for Data Centers and the Cloud

Juniper's Reference Architecture for Data Centers and the Cloud is a framework that allows architects to optimize data centers by:

- Scaling up and out without compromising performance or adding complexity
- Maintaining high availability
- Minimizing latency
- Maximizing capacity and throughput
- Dynamically providing protection against evolving threats
- Supporting diverse application flows (multiple services, directions, and controls)
- Minimizing energy consumption
- Incorporating standards
- Maintaining an open architecture for adding value

We have constructed our architecture to be robust enough to serve the range of customer environments that exist today. Our architecture can be applied to targeted implementations where architects focus on just a portion of the overall data center network; or it can provide a context for designing entire data center networks to meet needs across the range of categories we have described.

To accommodate the industry's embrace of cloud computing, our reference architecture is flexible enough to enable private, public or hybrid cloud computing services and to support the application environments that are critical to achieving the organization's business objectives.

This document is intended to help organizations that are considering cloud computing, whether or not they have started to implement any cloud elements. We will consider the implications of cloud computing for data center architectures and provide a reference for organizations who want to adopt cloud computing as they move forward.

It is our objective to provide a reference architecture that is unsurpassed in its ability to meet the needs of a diverse range of organizations, keeping in mind that the business objectives are primary, and that cloud computing is the enabling model, not the end in itself.

## **Key Trends in Today's Data Center**

In this section, we present some of today's key market and technology trends and examine how these trends inherently affect data center requirements.

### **Evolving Business Application Architectures**

Today's enterprises rely on their business applications. Business applications enable transactions for internal employees, collaboration with outside partners and customers and capabilities that improve the business' competitive advantage. In today's globally competitive world, applications must be available everywhere and at all times. When business applications perform on an "as needed" basis, the organization thrives; when they do not, business is lost. Concurrently, we also note the evolution of a rich mix of application architectures that must be supported in their own right. In many cases, these are blended into mixed or tiered designs with a range of resulting flows. Some are strictly constrained to a narrow, necessary content mix, while others are more fluid and involve a varying mix of content and transaction types depending on user choice. A key requirement of data center architectures is to support a wide range of applications successfully. Some of these application types include Service-Oriented Architecture (SOA), Software as a Service (SaaS), Web 2.0, Unified Communications (UC) and streaming services.

### **Server Virtualization**

Aligned with the trend toward more powerful servers, more open application designs, and the need to accomplish "more with less" in the data center infrastructure, the adoption of virtualization in the server infrastructure continues to increase. This produces a need to network the individual virtual machines with an additional layer of "virtual switching" within each server. Because multiple logical hosts now run on an individual server, it becomes necessary to differentiate their identities within the network and allow them to operate properly within their own logical domains. This trend creates the need to relate the virtual and physical network configurations, and it creates an interest in the ability to move application workloads in a flexible and seamless fashion.

#### **Increasing Demands on Bandwidth and Capacity**

Rich media applications, proliferation of users and device types, compute and storage utilization, and access methods continue to drive technology innovation. From a bandwidth perspective, we have seen a progression from GbE to 10 GbE to 40 GbE and 100 GbE links, and this evolution will continue to drive requirements in how data center networks are built.

### **Reducing OpEx**

Changes in the global economy and the desire to achieve greater business value associated with IT investment are creating more pressure to control costs. Despite more stringent requirements for high availability and resiliency, this is particularly relevant for the ongoing operational costs associated with maintaining IT and data center networks.

### **Protecting Against Security Threats**

New types of attacks are constantly surfacing, and attackers often employ new ways to exploit and hide in legitimate traffic. This places organizations in a continual mode of catch-up, trying to make sure that they have appropriate protection against the latest vulnerabilities and threats. With the emergence of new applications, the security landscape continues to change. Although existing intrusion prevention techniques are still applicable, simply identifying source and destination addresses and port combinations no longer offers sufficient protection. The concept of application fluency is required to address these evolving security threats.

## Convergence of Fibre Channel and Ethernet Networks

Design evolution has allowed storage to be “pooled” for access over networks by a diverse population of servers and computers. Distinct technologies have emerged to enable designs to handle server-to-storage communications, and this has led to a desire to work towards the design of a converged storage and Ethernet data center network that would allow storage and application traffic to share the same common network. This would ultimately save money and allow increased operational efficiency.

## Requirements

At the same time that we take note of the varying business objectives that drive organizations and their data centers, we also need to note the evolution of a rich mix of application architectures that must be supported in their own right, and in many cases are blended into mixed or tiered designs with a range of resulting flows. Some are strictly constrained to a narrow, necessary content mix; while others are more fluid and involve a varying mix of content and transaction types depending on user choice. For example, the most significant impact of SOA and Web 2.0 applications is the variability of traffic load and traffic patterns that both permit and often place demands on the network infrastructure. Without proper network planning, every new SOA or Web 2.0 mashup application is at risk of creating congestion, performance problems and even application failures. Latency, jitter and packet loss effects are important predictors of UC and streaming services.

To successfully support a range of application types is a central requirement of data center architectures. Following are some of the key requirements that are emerging for businesses as they plan for the evolution of their application infrastructure and anticipate the impact of these changes on their data centers.

**Performance**—To an enterprise’s customers, partners and employees, business applications are the means to an end, the ability to obtain information, complete transactions, or perform a job. High performance is essential to employee productivity, customer satisfaction and the enterprise’s bottom line. Application response time is the most fundamental component of understanding application and data center network performance.

**Scalability**—In existing computer and network environments, planning for growth and change is a costly and time-consuming effort. A successful organization must be able to readily and cost-effectively scale business applications, even when capacity limits are reached within existing data centers.

**Accessibility**—In today’s mobile and volatile world, users now require access anywhere in the world, on virtually any type of computer and network connection, 24 hours a day. Enterprises must support access from corporate headquarters, branch offices, other business establishments, home offices, wireless hotspots and cellular networks throughout the world.

**Agility**—As the pace of global economic activity continues to accelerate, organizations must be able to respond quickly to changes in demand and other market conditions. Agility improves with the user’s ability to reposition infrastructure resources rapidly and inexpensively. Business applications that support agility can help reduce time to market, strengthening the organization’s competitive position and increasing market share.

**Availability and Continuity**—No application is 100 percent failure proof. To protect an enterprise’s competitive edge, business applications must be at least as available as those of competitors, and productivity must not suffer when failures occur. Furthermore, when a disaster occurs, the organization should recover with minimal discontinuity, getting business applications online again quickly and ensuring that the associated user data is protected and available.

**Security**—Security is a multifaceted concern that touches upon almost every aspect of the business landscape. Organizations must respond effectively to evolving threats that can compromise business data or interfere with application availability. They must ensure secure operations in shared environments and meet industry compliance and regulatory requirements. Business applications must also support guaranteed service-level agreements (SLAs) and be consistent with stringent real-time requirements.

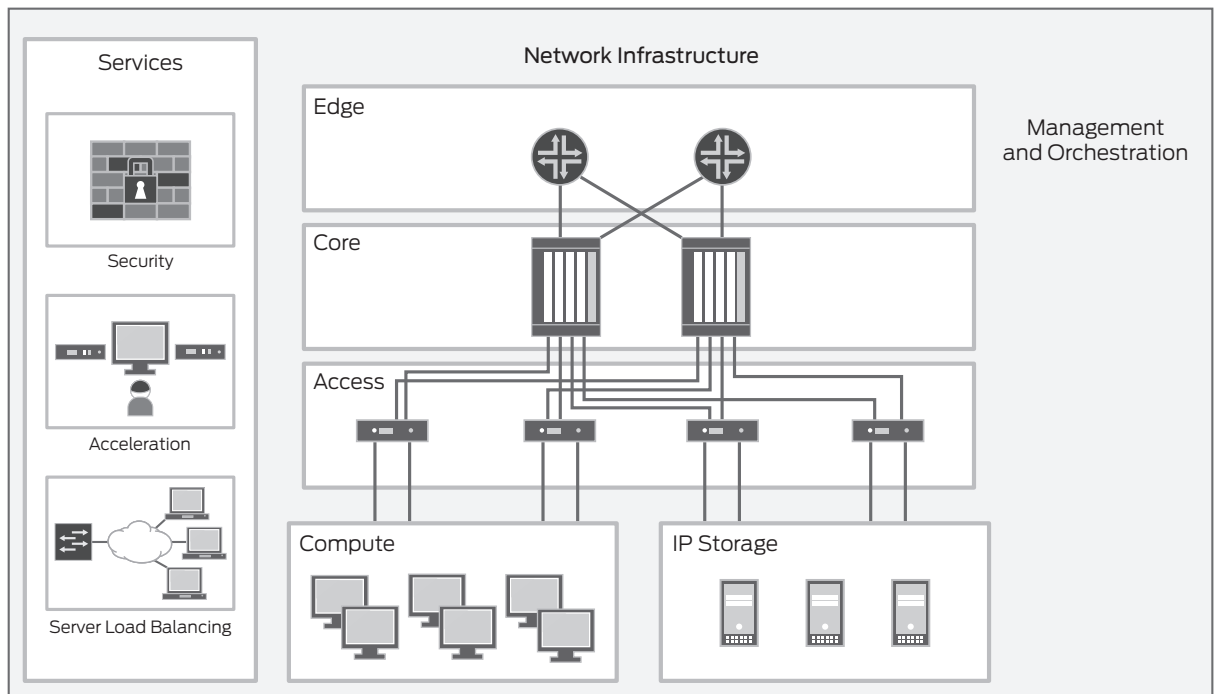
**Manageability**—To help reduce OpEx, the data center network should be orchestrated to simplify the management tasks associated with configuration, monitoring, maintenance and other administrative tasks.

## Functional Areas in the Cloud-Ready Data Center

To deliver applications from the cloud data center, organizations must divide the required tasks into optimized functional areas. Effective choices within each functional area can help designers meet application goals with respect to latency, availability, security and scale.

Figure 1 illustrates the framework we employ to envision the data center network at its highest level. It includes the following areas and their functional interrelationships:

- **Network Infrastructure**—provides connectivity and transport for applications and services between users and the data center, within the data center and across multiple data centers. The Network infrastructure has three main sub components, namely the access network, the core network and the edge network.
- **Compute and Storage**—represents the compute and storage infrastructure appropriate for applications (rack-mount and chassis-based, cost-effective and multi-core, with unstructured content and highly structured transaction databases). The compute and storage functional area hosts all business applications such as Enterprise Resource Planning (ERP), SaaS, SOA and Web 2.0 applications (among others).
- **Services**—supports applications with security, user verification, and entitlement, and application support, including application acceleration, deep packet inspection (DPI), and load balancing
- **Management and Orchestration**—ties together all of the elements of the cloud-computing infrastructure, enabling efficient and responsive monitoring, management, and planning



**Figure 1. Data center reference framework**

While each component has its own characteristics, specific requirements and enabling technologies, Juniper Networks packages them all together with a common cloud-computing architecture that meets the individual and combined requirements with powerful enabling technologies. Let us take a closer look at each of the functional components beginning with business applications.



## Juniper's Approach to a Cloud-Ready Data Center

To maximize effectiveness of the data center across the major functional areas, Juniper has embraced a strategy to optimize designs in multiple dimensions: to simplify, share, and secure the data center network to the maximum extent possible and to provide a powerful suite of automation tools. Each dimension brings concrete value to solution designers, enabling data centers to meet important application delivery objectives:

1. **Simplify.** By simplifying the data center network, we mean minimizing the number of network elements required to achieve a particular design, thus reducing both capital and operating costs. Simplifying also means streamlining data center network operations with consistently implemented software and controls.
2. **Share.** By sharing the data center network, we mean intelligently (and in many cases dynamically) partitioning the infrastructure to support diverse applications and user groups and to interconnect large pools of resources with maximum agility. In many cases, this involves powerful virtualization technologies that allow multiple logical operations to be performed on individual physical entities (such as switches, routers and appliances).
3. **Secure.** When we secure the data center network, we must extend protection to support the rich, distributed architectures that many applications currently use. This requires a robust, multidimensional model that enhances and extends the traditional perimeter defense. By increasing the granularity and agility of security policies, we can enable trusted sharing of incoming information and resident data within the data center, while complementing the functions embedded in operating systems and applications.
4. **Automate.** By automating, we mean capturing the key steps involved in performing management, operational, and application tasks, and embedding task execution in software that runs as an intelligent added value to the overall data center operation. Tasks can include synchronizing configurations among multiple disparate elements, starting and stopping critical operations under various conditions, and diagnosing or profiling operations on dimensions important for managers to observe.

With the high-level framework in mind, now we can discuss the individual functional components and their associated requirements and enabling technologies.

### Network Infrastructure

When designing the data center network, we must consider all communications occurring within the data center itself, between the data center and its users, and among data centers within the cloud. The infrastructure consists of a combination of elements in three domains, integrated in a variety of ways based on customer needs:

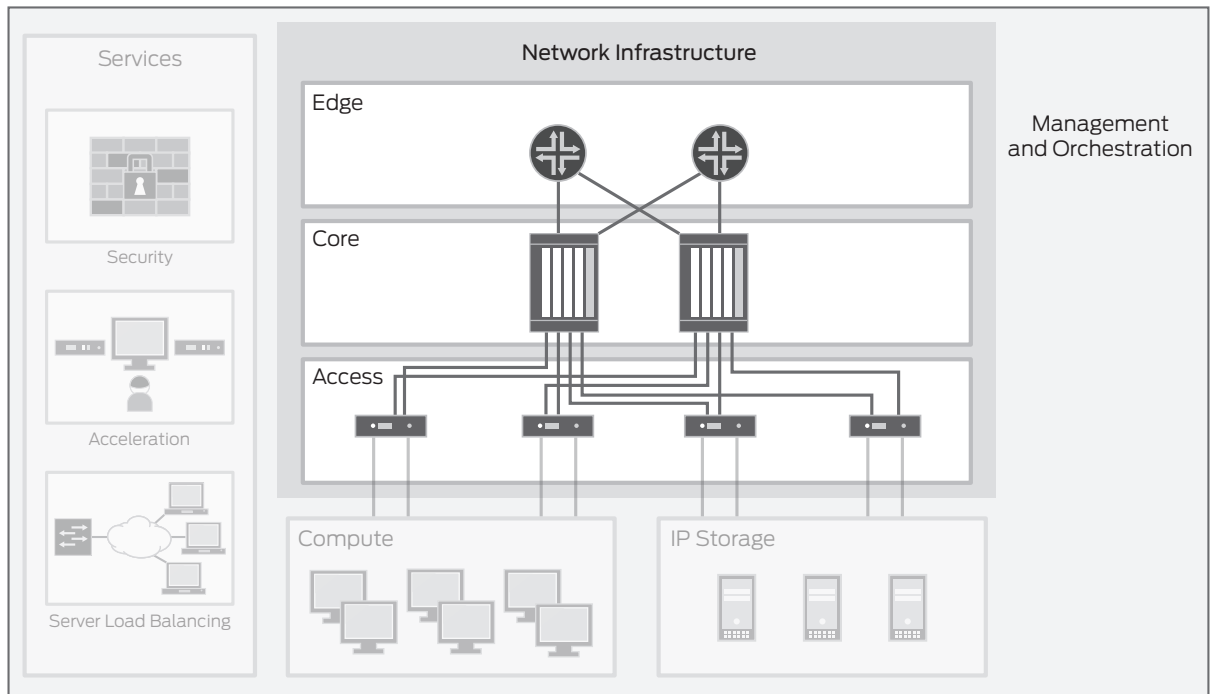
- Access network
- Core network
- Edge network

The access network provides connectivity to all shared enterprise servers, applications, storage devices, and any IP or office automation devices required in the data center facility. Most data center access switches are deployed at the top of the rack or at the end of the row of server racks.

The core network provides a fabric for high-speed packet switching between multiple access network devices. Due to their location in the network, core-layer switches must provide scalable, high-performance, high-density, wire-rate ports, and HA hardware and software features that deliver carrier-class reliability and robustness. The core serves as the gateway where all other modules such as the WAN edge meet. It typically requires a 10 GbE interface for high-level throughput, and maximum performance to meet oversubscription levels. The core provides high-speed throughput for all data going into and out of the data center, and it must provide resilient, fail-safe Layer 3 connectivity to multiple access layer devices.

The edge network provides the communication links to end user networks of various types. These can be private WAN or campus backbones, mobile access networks, VPNs, or other types of Internet access. The high performance and reliability of these connections improve user experience. Agility ensures that users will have access to applications and services where and when they are needed. In addition, multilayered security controls ensure that users, applications and data are protected at appropriate levels.

Figure 2 shows the network infrastructure functional area of the reference framework.



**Figure 2. Reference architecture network infrastructure**

### Application Traffic Flows

In the past, applications were designed with a very specific traffic flow. Typically, requests would originate from a client system and be routed to a single application server, which would then respond directly back to the client. This client/server model was, in effect, a single direction north-south scheme. Because of demands for greater application performance and response time, and the continued adoption of virtualization technologies, application architecture has changed. A more distributed model has also had an impact on application traffic flows. Today, a request originates from a client system and is routed to an application, but the processing of the request results in information sharing across multiple servers, prior to responding to the original request. Furthermore, these servers can exist across multiple physical machines and locations. Because of this shift, the network infrastructure should optimize the ability of the application infrastructure to handle the increasing levels of server-to-server communication streams.

### Simplified Network Infrastructure

Another significant trend in data center networks is the continual need to provide scale and agility for growth, while simultaneously controlling costs. As new applications and business models emerge, the network design that worked well for businesses may not be able to support new demands on the IT infrastructure and, most importantly, new business requirements. Networks built on fragmented and oversubscribed tree structures have scaling and consistent performance problems. As more devices are added, design and management complexity and costs increase exponentially. A simplified network infrastructure can help meet these requirements of scale, while mitigating the concerns of cost and complexity.

### Distributed Data Centers

Due to rapid growth, bandwidth, and latency considerations as well as space, power, or cooling capacity requirements, data center locations continue to multiply. While this has catalyzed a desire for improved simplification and consolidation, organizations also are considering ways that will enable the network infrastructure to connect these different locations together.

Similar to the application infrastructure, as the demands for agility, responsiveness and scale continue to grow, it is critical for the same agility, responsiveness and scale to be available to services running in two or more distinct data centers. The network should serve as a transparent platform for elastic and reliable delivery of services to distributed users and applications.

### Juniper's Approach

Juniper Networks is applying the expertise it has gained from high-performance global networking to address the requirements resulting from these trends. Juniper's approach leverages a combination of products, technologies, services, and design expertise to build the most efficient, scalable, and cost-effective data center network possible. Some highlights include:

At the access tier, Juniper switches provide a flexible Virtual Chassis configuration that connects all of the physical switches to a high capacity backplane. This reduces the number of uplinks—in some cases down to zero—that traffic must traverse in the case of server-to-server traffic patterns. This measurably improves performance for server-to-server communications in SOA, Web services, and other distributed application designs, and allows organizations to better address today's application infrastructure while simultaneously simplifying the data center network.

In the core, Juniper simplifies the data center network and eliminates layers of cost and complexity. Using fabric technologies such as Virtual Chassis technology, Juniper helps flatten data center networks, often reducing them from three layers to two. This technology enables a simpler, spanning-tree free topology, which enables a network that is optimized to scale as the business grows.

At the edge, Juniper Networks leads the way, making it possible for enterprises and service providers to implement network architectures and services based on MPLS. Juniper platforms provide a wide range of MPLS features and functionality powered by Juniper Networks® Junos® operating system. Combined with Juniper's superior silicon and platform architecture designs, Junos OS delivers advantages over alternative platforms that are either too immature to support the breadth of MPLS features or architected in a monolithic fashion that is too complicated or unwieldy for efficient management. Juniper's platforms offer the most complete, advanced routing features in the industry without compromising performance. These features include traffic segmentation and virtualization with MPLS, ultra low latency multicast and comprehensive security and QoS implementations to accelerate delivery of time-sensitive applications and services. The carrier-class reliability and HA features available on Juniper Networks MX Series 3D Universal Edge Routers include graceful restart, nonstop routing, MPLS fast reroute, unified in-service software upgrade (ISSU), and VPLS multihoming. Together, these capabilities provide a unique and superior foundation for creating a transparent virtual cloud to interconnect multiple data centers.

### Compute and Storage Infrastructure

We begin by describing the characteristics and trends for determining how compute and storage elements participate in the overall data center infrastructure. In many ways, the data center exists to support optimal operation of these elements, particularly because they drive the way application logic is performed. As we will see, there is a symbiotic relationship among the compute, storage and network elements in the data center, and successful service delivery requires the optimization of designs within each area.

Server virtualization is enabling deployment of large application workloads with greater agility, efficiency and reliability than has been possible in the past. In the storage infrastructure, complementary virtualization and efficiency, combined with high-speed data transfers for data processing and replication, are helping to improve overall application performance, agility and availability. In addition, I/O convergence is evolving to simplify overall data center networking, with the promise of eliminating costly redundant infrastructures and supporting a variety of cloud-enabled applications on a single, unified network. Let us consider each of these advances in turn. Figure 3 shows the compute and storage functional area of the reference framework.

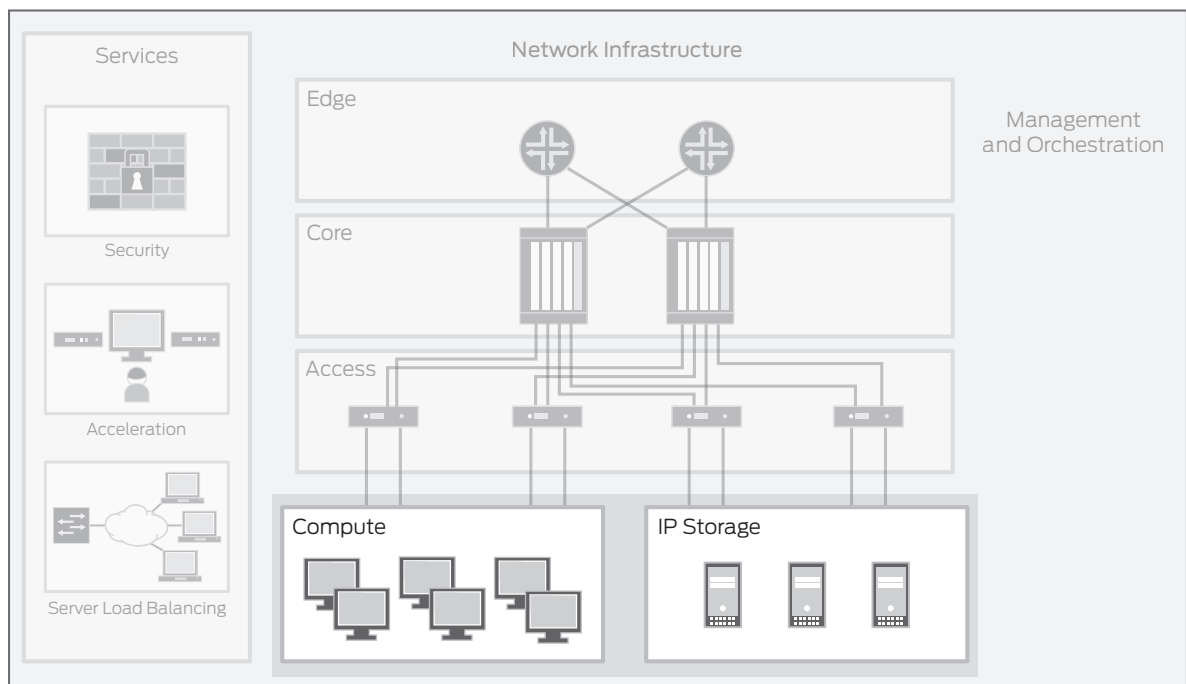


Figure 3. Compute and storage

### Integrating Virtual Server Infrastructure

Server virtualization reduces the number of physical servers in the data center and provides greater flexibility to meet rapidly changing business needs. However, server virtualization introduces challenges as well, some of which directly involve the data center network. Virtual machines increase the density of traffic loads to and from individual machines (because each virtual machine has its own operating system and applications). This increases network link utilization and places additional demands on the network fabric, especially when we consider dynamic creation and migration of virtual machines.

The use of virtual machines also creates an additional logical (or virtual) layer of networking within the server endpoints and between the virtual machines. A “virtual network extension” allows separation and connection of traffic to and from individual virtual machines, both within the physical servers and between the physical servers and the rest of the network. This creates a need for configuration, state and policy integration between the physical and virtual parts of the network.

As workloads change, the data center infrastructure must support rapid, on demand reassignment of resources in a way that is completely transparent to end users. Compute capacity must scale to meet the demands for applications and services without disruption. Scaling must encompass high-density deployment within the data center, and it must provide processing power flexibly across multiple data centers.

With virtualization technology now supported on multiple operating systems and computing platforms, the data center network architect must evaluate the impact of the virtualized server environment on network architecture.

- Increased capacity due to higher link utilization (multiple virtual machines now running on an individual physical server) and associated resources (increased media access control (MAC) and IP addresses and applications per physical server)
- Expanded availability requirements due to increased operational risk (loss of one server means the loss of numerous virtual machines)
- Increased relevance of standards and automation in the integration of physical and virtual networks, dynamically and at scale
- Increased importance of network-based services and their relation to virtual infrastructure such as firewalls, intrusion prevention systems (IPS), and load balancers, all of which affect network performance.

Provisioning sufficient bandwidth to meet application SLAs is a primary consideration. A conventionally oversubscribed network design becomes unacceptable in the face of increased link utilization and dynamic traffic flows. To meet SLAs, architects must consider increasing link bandwidth in the server and network infrastructures.

To provide load balancing in the cloud, some virtual machines may need to be moved across physical machines within the data center or to other data centers, and the network must have the agility to support this move.

In this environment, it is important that the network and virtual servers be synchronized automatically with respect to virtual machine configuration and policies. This is critical for managing SLAs and meeting audit and compliance requirements. For successful virtual server networking, the architecture must embrace the emerging extensions to the IEEE 802.X family of Ethernet protocols that enables synchronization of physical and virtual network configurations under the name of Virtual Ethernet Port Aggregator (VEPA). These standards help customers maximize choice in deployment of virtual servers and confidently support networking them with agility and high performance, regardless of the number of applications and hypervisor vendors used.

An additional subtlety in successfully supporting the virtual server endpoints is enabling a successful end-to-end security architecture for the applications. In the virtual server environment, conventional security practices such as monitoring network activity, inspecting and filtering traffic, and maintaining strictly separate security domains are often absent. Inter-virtual machine communication is a particular blind spot. Virtual machine traffic does not touch the physical network and is not protected by physical network monitoring or security.

Filtering traffic to and from a virtual server (or cluster) is only one part of the solution. To truly mitigate the risks within the virtual environment, especially those related to inter-virtual machine communication, an in-depth defense at the level of individual virtual machines is required. An effective, multilayered defense is only feasible if it maintains the productive capacity of the host servers and remains independent of the malware it defends against. An approach that integrates the capabilities of virtual appliances running within hypervisor environments with the security capabilities of the physical data center network is the type of integrated, multitiered, and multilayered design required for end-to-end success with virtual machines and the cloud.

New data centers also require managing virtualized network and security profiles and virtual machine configurations as they migrate across physical hosts. Managing profiles across physical hosts is difficult and may prevent organizations from taking server virtualization efforts beyond server consolidation and into dynamic resource allocation. Juniper addresses this requirement with Juniper Networks Junos Space applications such as Virtual Control, which allows for management of virtual machine configurations and switch port profiles on an integrated basis between the physical and virtual domains.

## I/O Convergence

The rising cost and complexity of building and operating modern data centers have led organizations to seek new ways to make the data center infrastructure simpler and more efficient. Although the cost of data center networking equipment is relatively small compared to the cost of server hardware and software, the underlying network fabric is the linchpin that connects all mission critical resources. A simpler, more streamlined data center fabric means greater efficiency and productivity and lower operating costs. In addition, shared (centralized or distributed) storage, be it file-based ((Network Access Storage (NAS), or block-based (storage area network (SAN) using Internet Small Computer System Interface (iSCSI), Fibre Channel (FC), and Fibre Channel over Ethernet (FCoE)) are essential elements of an effective compute and storage solution for data centers and the cloud. They can be used in concert to support advanced virtual systems and the overall virtual networking infrastructure.

Traditionally, servers are deployed with multiple I/O cards to connect to multiple separate physical network segments or even completely separate network infrastructures: dual SAN for disk access, another SAN or LAN for backup, dual LAN for client/server or campus LAN connection, out-of-band management, VMotion and cluster traffic. I/O convergence helps to reduce the number of such interfaces and networks. It has been promoted along with Ethernet or IP-based storage technologies such as iSCSI NAS and more recently FCoE.

With the increased affordability and rapid adoption of 10 GbE in the data center, Ethernet is poised to take on the connectivity tasks formerly relegated to InfiniBand and Fibre Channel to become the dominant data center networking technology. Reducing the number of I/O cards and network ports drives many potential savings.

**Data Center Bridging (DCB)** is a set of proposals developed by the IEEE to support I/O convergence, including the necessary capabilities to support the stringent requirements of storage network traffic. To drive the numerous enhancements required for lossless and delay constrained transport, a consortium of network, storage, CPU, ASIC, server, and network interface card (NIC) vendors have joined forces to create the Converged Enhanced Ethernet (CEE) proposals. The original version of this proposal, which was submitted to the IEEE in March 2008, includes recommendations on four vital components of the IEEE 802 standard that comprise DCB and governs Ethernet's physical and logical properties:

1. Priority Flow Control (PFC)—IEEE 802.1Qbb
2. Enhanced Transmission Selection (ETS)—IEEE 802.1Qaz
3. Ethernet Congestion Management (ECM)—IEEE 802.1Qau
4. Data Center Bridging Exchange Protocol (DCBX)—IEEE 802.1 DCBX

DCB's Enhanced Ethernet transport will benefit existing Ethernet or IP-based storage networking solutions such as NAS and iSCSI, and like Fibre Channel, will benefit from a lossless and delay constrained transport.

### **Fibre Channel and FCoE**

Fibre Channel was designed initially as a transport mechanism for the Small Computer System Interface (SCSI) protocol to connect and share a set of remote disks with a group of servers. The technology was progressively enhanced over the years, and today it includes sophisticated services and management tools that make it the preferred technology for moving mission critical storage traffic within data centers. Scale, performance and cost pose challenges to the future viability of this technology in the cloud data center. However, as application workloads grow and the pressures on efficiency remain strong, scale, performance and cost pose challenges to the future viability of this technology in the cloud data center.

Currently, standards development in the International Committee for Information Technology Standards (INCITS) organization is addressing the challenge of transporting Fibre Channel-based SAN traffic over Ethernet.

The T11 Technical Committee is a subgroup of INCITS developing the FCoE protocol. The first version of the standard was completed in the summer of 2009 as part of the Fibre Channel Backbone 5 (FC-BB-5) project. Since then, the committee has started work on a new version of the standard as part of Fibre Channel Backbone 6 (FC-BB-6). Although the development of FCoE as an industry standard will bring the deployment of unified data center infrastructures closer to reality, FCoE by itself is not enough to complete the necessary convergence. It is, however, highly dependent upon the success of the aforementioned standards designed to enhance the behavior of Ethernet within the data center.

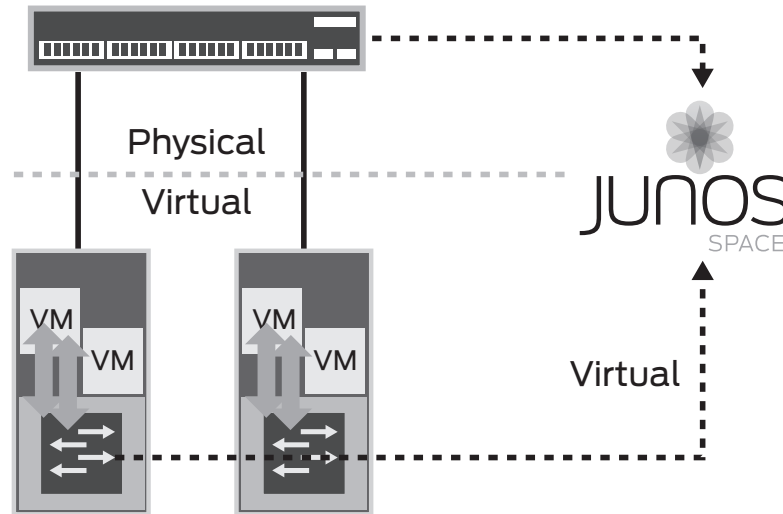
CEE, FCoE and NAS subsystems are essential elements of an effective compute and storage solution for data centers and the cloud. They will be used in concert to support advanced virtual systems and the overall virtual networking infrastructure.

### **Juniper's Approach**

In today's virtualized server environments, there are multiple software implementations of the host operating system (hypervisor) which host a number of guest operating systems (virtual machines) in a physical server. There is also a virtual switch on that server that does the switching between virtual machines. This adds complexity to the network because security and configuration (of virtual switches) must be consistent with physical switches. There are often situations now where part of the network is run on the server and is being managed by a server team. This can result in inconsistencies in the way that servers and the network are configured, and can cause a rift between the server administration and network management teams in a data center.

To enable consistent configuration and visibility of the virtual and physical network, Juniper Networks' chosen solution is a web-based software application called Junos Space Virtual Control. Residing on Junos Space <http://www.juniper.net/us/en/products-services/software/junos-platform/junos-space/>, Juniper's Orchestration Software platform, the Virtual Control application enables end-to-end network topology, configuration and policy management from a single pane of glass. Junos Space Virtual Control dramatically simplifies data center management, reducing total cost of ownership (TCO) by providing operational consistency and visibility throughout the network.

Figure 4 shows Junos Space Virtual Control application managing both the physical and virtual network.



**Figure 4. Consistent management of the physical and virtual network from Junos Space**

**Junos Space Virtual Control**—Junos Space Virtual Control allows network operators to discover, configure, provision, and monitor a VMware vNetwork Distributed Switch (vDS) as well as a Juniper switch platform. This single pane of management facilitates synchronous configuration changes for both physical and virtual switching environments, and it simplifies network operations by dynamically mapping port profiles to support VM mobility. Junos Space Virtual Control leverages VMware open APIs to achieve this functionality, while similar integration with Junos Space can be achieved with other virtual switching environments (Xen, PowerVM, Hyper-V) with similar open interfaces. An emerging standard is being developed to define the interface for virtual and physical switching environments called Virtual Ethernet Port Aggregator (VEPA). VEPA is a nondisruptive and cost-effective solution to inter-VM communications. Implementation requires minimal changes to the software running on the physical switch, not wholesale replacement of the existing networking infrastructure. VEPA allows virtual switching to be extracted from the server, improving server performance and increasing the number of VMs that can run on each server. Finally, because VEPA is based on open standards and is server- and hypervisor-agnostic, customers have maximum flexibility in deploying server virtualization. VEPA will enable rapid innovation in services for users, as well as operational consistency, simplicity, and efficiency.

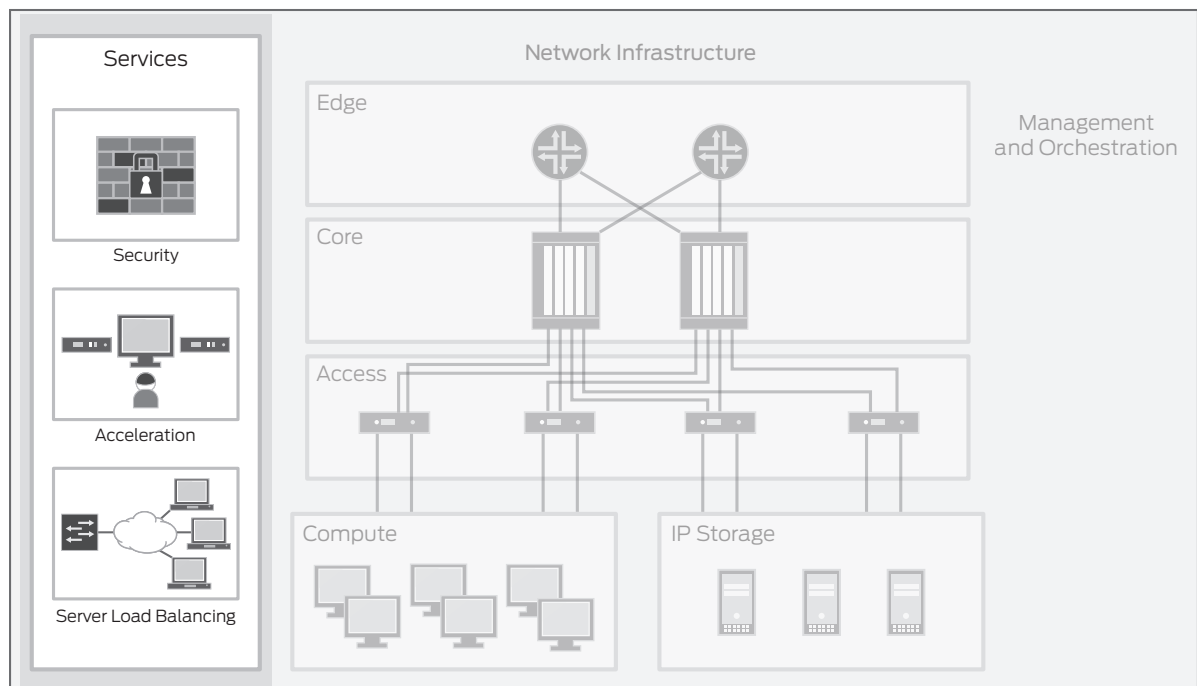
The pending VEPA standard also contains a critical feature known as multicasting. Because many virtual servers contain more than one virtual network switch, physical switches must be able to identify the virtual switch source of traffic coming to them. While this advanced feature will require some hardware upgrades, the basic VEPA technology can be supported with a simple software upgrade.

**I/O Convergence**—Converged data center networks will require a robust and complete implementation of FCoE and DCB standards to be viable in supporting the critical application and data integrity requirements of data center applications. Because of the timing of ratification of the respective standards (FCoE having preceded DCB by approximately a year) and because of the incremental progress in cost effectiveness in the related infrastructures (early implementations not truly passing the cost effectiveness test), implementation of converged data center networks will occur in two phases. In phase one, convergence within the rack will enable partial gains while supporting separate LAN and SAN infrastructures using FCoE gateways between the two. In phase two, networks will be converged fully by virtue of support of the full DCB standards suite and by allowing adequate support for all traffic types in an optimized data center network.

By maintaining active participation in the related standardization efforts and by rethinking the technology and economics of the data center network from the ground up, Juniper Networks provides customers with a winning set of platforms. Juniper also offers a pragmatic, innovative strategy to develop a single, converged data center fabric with the flexibility and performance required in a fully virtualized infrastructure, while continuing to drive down the cost and complexity of enabling it properly.

## Services

As we have seen, data centers are increasing agility and versatility for service delivery in highly virtualized environments. While this enables managers to function responsively, it also exacerbates risks that—if not addressed—can compromise the effectiveness of the newly tuned environment. These risks are principally in the areas of security and application performance. Forward-looking network architecture in the virtualized world includes functionality embedded in the network itself that controls and mitigates many of the risks and facilitates optimum performance. The idea is that protection and acceleration capabilities can run in the network on behalf of or in concert with functionality that executes in the application endpoints, for overall effective, secure and responsive system architecture. Figure 5 shows the services functional area of the reference framework.



**Figure 5. Services functional area**

Juniper's data center reference framework includes a functional area dedicated to delivery of virtualized services. We describe the capabilities of that area in this section.

The Services functional area allows data center managers to address the following critical challenges:

- **Evolving threat landscape**—Data center and cloud service operators must address ever-escalating threats to application delivery, integrity and privacy. Major threats include service disruption, application denial-of-service (DoS) attacks, data leakage to the outside world, attacks on data integrity, and identity fraud.
- **Sharing of resources**—Resource sharing allows organizations to realize economies of scale that are essential to success with virtualization and the cloud. However, to realize this potential, operators must be confident that shared resources such as virtual machines, applications and supporting platforms will not be compromised.
- **Managing virtualization risks**—In the traditional data center where resources and applications map directly to physical equipment, security is straightforward because physical boundaries still exist. However, as resources become virtualized, traditional security controls are insufficient. Intelligence is required to help operators understand and limit the risks that arise when physical boundaries are replaced by virtual boundaries.
- **Granular policy control**—Managers must secure the entire path between the end user's source and the destination application. This requires extensive, granular policy control for security and entitlement throughout the infrastructure. By using granular controls, managers can ensure service integrity and meet SLAs.



- **Traffic integrity and confidentiality**—In many cases, traffic flowing through the cloud must be secured to prevent unwarranted data disclosure and ensure the confidentiality of user information.
- **QoS**—To assure the quality of the end user experience, services are required to apply QoS metrics, such as preferred service for VoIP traffic.
- **Compliance and SLA**—Cloud data center operators must meet auditing and risk assessment requirements mandated by regional regulatory authorities. Services can be deployed to ensure that cloud providers meet these requirements and to demonstrate that security controls are effective in enforcing security policies.
- **Application acceleration**—Application acceleration services can boost the performance of major applications within an enterprise. Applications may be business critical (for example, ERP applications such as SAP and Oracle), or contribute to employee productivity (such as Microsoft Outlook).

In prior designs, services were required only at the data center edge, with gateways securing the connection between end users and the data center interior. This has been called a “perimeter defense.” With current trends toward consolidation and virtualization in the data center and management of flows between data centers, security- and application-related services are now required at a greater number of control points in the data center and virtual systems, not only at the entry gate. A comprehensive and agile architecture of services intelligence must be deployable from applications and hypervisors running in virtual machines, to critical protection points in the core of the network, to the data center edge, and ultimately to the end user.

A well-designed infrastructure applies services where needed and enforces policies dynamically on network traffic. Delivery of services can be optimized by using resource pools that are shared across the network.

As with the network infrastructure, services in the data center must meet stringent requirements for performance, scalability and availability. They must also support granular policy control and the intelligence to meet user- and application-specific SLAs.

The Services functional area comprises two major groups of capabilities: security and application services. Security services control access to resources and protect traffic within the cloud. Application services improve the performance, scalability and agility of applications and infrastructure and simplify operations.

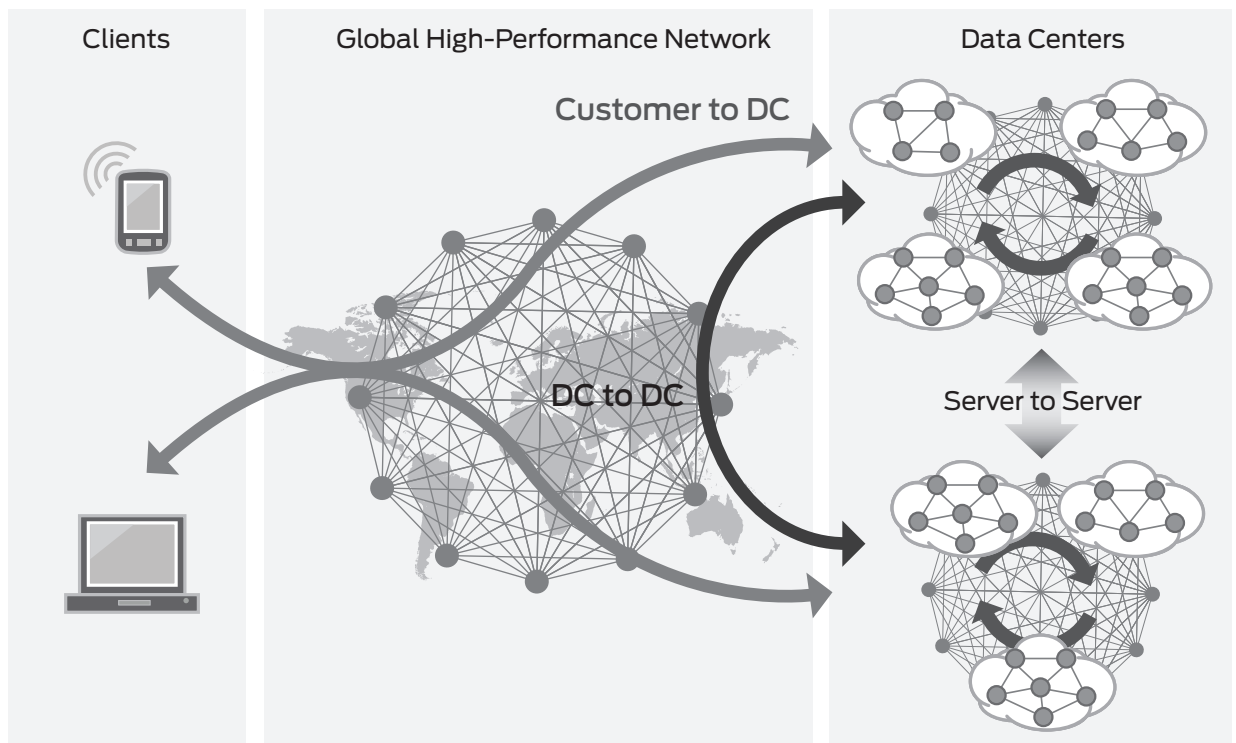
The following section focuses on these types of services.

## Security Services

A single data center can include many thousands of physical compute and storage arrays that enable hundreds of thousands of virtual endpoints used by tens or even hundreds of thousands of clients. The result is a complex set of flows between servers and clients (north-south) and among compute and storage systems (east-west). Comprehensive, effective security must be deployed to scrutinize all traffic and weed out any traffic that can pose a risk to traffic flows or data integrity.

Figure 6 shows the major types of traffic flows that must be secured in the virtualized data center world:

- East-west traffic between servers within the data center and between compute and storage systems (server to server).
- North-south traffic between servers and end user systems, where the end users can be anywhere in the world, use virtually any type of client device, and obtain access through almost any type of commercial access network (customer to data center).
- Traffic between data centers for fast response to changes in demand and load conditions (data center to data center).



**Figure 6. Flow types in the new cloud infrastructure**

Security services must also take into account the fact that Web 2.0, UC, and rich media applications lead to unpredictable, challenging traffic patterns. A client no longer communicates with a server using a single stream of data (or TCP session) to complete a request. New applications present more challenging traffic patterns involving communications between multiple servers over multiple sessions to fulfill a single user request. A single application interaction now requires much more server-to-server communication and higher levels of module performance to meet requirements. Traffic flows involve multi-node applications, server virtualization and storage over IP. Furthermore, throughout the process, user and data identity and integrity are at stake.

Ultimately, at the user and operator levels, security of information handling is about trust. The fast evolving world of virtualization and the cloud has been inhibited partially in its uptake by the more slowly emerging architectures designed to ensure trust. By putting appropriately broad and effective security services in place, organizations can increase trust levels among end users and potential subscribers, and drive increased satisfaction and demand.

Traditional security platforms, including routers, firewalls, IPS, VPN, and network access control (NAC) continue to be central to the security of the data center. However, existing use of these elements is not sufficient to meet new requirements of the cloud. To meet security challenges of the cloud, an extended portfolio of interconnected security services is required, including stateful firewalls, IPS, application and identity awareness, secure remote access, NAC, Domain Name System/Dynamic Host Configuration Protocol (DNS/DHCP) services, and authentication, authorization, and accounting (AAA) services.

We highlight the role of each of these critical technologies within the security functional area of the cloud in the following sections.

**Stateful firewalls**—Data center operators have traditionally deployed numerous firewalls to separate servers by function or tier in their system designs, for example, the database, application, and Web tiers of a system. Multiple firewalls were often deployed at the same level of the network and inhibited overall performance. In many cases, firewalls were bypassed when concerns about their ability to pass real-time traffic were considered paramount and the desire for performance outweighed security concerns. Concerns about firewall impact on raw bandwidth, connections per second, and sustained connections caused some data center operators to limit firewall use or even dispense with firewalls in some areas. All of these limitations of prior architectures have resulted in high-risk compromises that are not sustainable in the current privacy, compliance and high-performance end user environments.

The challenge is to move from this existing situation, in which appliances are solution inhibitors, to the cloud data center network, in which firewalls are solution enablers. This necessitates the introduction of high-performance, stateful firewalls at the data center core.

High-performance, stateful services are the cornerstone of security in the virtualized data center. Stateful services enforce policies that align with business and operational requirements through the identification and classification of networks. In addition to being the primary Layer 4 access control system, stateful firewalls can support many additional security functions such as DoS or quota protections, DPI on specific applications, and Network Address Translation (NAT).

With stateful firewalls, it is possible to introduce fine-grained control over all traffic flow types (intra-data center, inter-data center, and data center WAN) and to support key security functions such as NAT, Application Layer Gateway (ALG) services, IPsec VPN services, distributed DoS, as well as unified threat management (UTM), which includes antivirus, anti-spam, and Web filtering.

Additional capabilities can be inserted in a modular manner on top of this foundation. When done in a modular way per policy zone, this approach provides maximum agility, efficiency and performance.

**Intrusion prevention systems**—Network and application level attacks are an ongoing concern, and the data center network must be able to detect and prevent attacks in traffic flows by supporting versatile, high-performance IPS functionality as part of the security service. Because applications must be available to users at locations that are not inherently secure, the risk of misuse or application DoS will always be high. Moreover, because applications are colocated in virtualized data center infrastructures, a chain effect (in which an application is affected by the risk to which another application is exposed) can be created too easily.

IPS must be highly accurate in its detection and prevention capabilities, with low numbers of false positives and false negatives. Effective intrusion detection and prevention requires a multidimensional approach involving protocol analysis, anomaly detection, and signature analysis. IPS should support multiple detection modes and accommodate placement of sensors in different parts of the network.

As an example, sniffer modes involve network taps that passively observe the flow of traffic and identify potential threats, whereas inline systems are deployed with traffic flows and can potentially prevent attacks in real time. Mixed mode solutions can deliver the benefits of both sniffer and inline methods. Actions that are triggered when an attack is detected should include the traditional allow/deny along with finer grained actions such as rate limiting, setting Differentiated Services code point (DSCP) marking, closing client connections/server connections, and performing TCP resets.

The IPS platforms should support the performance and capacities required in data centers of varying sizes and inspect Layer 4 through Layer 7 information at line rates. They should coordinate threat responses with other access control gateways (SSL VPN and NAC) by sharing attacker information, so attacks can be mitigated closest to their source. Because protocol decoders in the IPS deconstruct streams and build the right context to look for threats, a powerful and rich protocol decoder must be in place. Finally, network-based security services, including intrusion detection, attack prevention, encryption, and monitoring, should be consolidated into highly scalable, virtualized security platforms to reduce security device sprawl.

**Fluency with application and user identities**—Historically, attack prevention has focused on identifying and thwarting malicious activity within allowed traffic, as evidenced by content security technologies such as antivirus and anti-spyware. These mechanisms have been a vital part of the network fabric and offer protection by identifying known attack patterns or behaviors that deviate from the norm.

Unfortunately, new types of attacks are constantly occurring, and attackers often employ new ways to exploit and hide in allowed traffic. This places organizations in a continual mode of catch-up, trying to make sure that they have appropriate attack coverage against the latest vulnerabilities and threats. Organizations need tighter control over what can and cannot be done within a given application. In other words, the solution must evolve from a reactive approach to a more proactive security stance.

Juniper has introduced stateful application filters such as stateful signatures and protocol anomalies. These filters control the commands that are used within an application, so that organizations can reduce the opportunities for exploitation and increase the availability of information and networking services.

However, with the emergence of new applications, the application networking and security landscape continues to change. Although existing intrusion prevention techniques are still applicable, simply identifying source and destination addresses and port combinations no longer offers sufficient protection.

Traditional stateful security devices assume that an application uses a service that runs over a fixed, predetermined, and publically acknowledged TCP/UDP port number, and that the traffic being processed can be identified by looking at the first packet in a session. This approach no longer works because the relationship between port numbers and applications is simply a convention that may not apply, and because it is necessary to examine subsequent packets to establish reliably the actual application and specific functions or commands that are being used.

The concept of fluency with application and user identities is intended to address these evolving security threats. The idea is to go beyond traditional security approaches to identify exactly what actions are allowed by specific users in specific application instances. Application and identity fluency are essential for applications such as BitTorrent, Skype and YouTube that are enabled on top of HTTP and use nonstandard ports (or even randomly assigned ports).

Fluency is also important to maintain agility in the data center. If an IT organization wants to shut down one application and bring up a new one, it must be able to do so quickly. If firewalls support only protocol and port mappings, doing so becomes a time-consuming and tedious task. To enable agility, firewall configuration must be supported at the application level with a fluency that is independent of ports and protocols.

To support application and identity fluency, network security platforms such as enforcement gateways, firewalls and monitoring systems must identify application context and user conversations with thorough and intelligent signature-based classification. They must provide visibility into the application infrastructure, making it possible to determine application usage profiles and other valuable application-level information. It must be possible to control application and resource access based on user identity, not just source IP address. With a mobile, dynamic workforce that connects to application elements that reside on multiple servers within the cloud, organizations can no longer assign access privileges based on a well-controlled and fixed user location represented by an IP address. Services must be application and identity aware.

Juniper's approach to enabling security services in the data center and cloud-computing environment enables all of these capabilities comprehensively.

**Secure remote access**—Given the trend towards consolidating applications into fewer numbers of data center sites, as well as the trend towards enabling modular applications to connect with each other in distributed application designs within and between those sites, users must securely access diverse resources from a variety of remote access points. At the entry points to the cloud, all endpoint access must be checked for compliance before access is granted, and the security status of the endpoint must be monitored throughout the time that a session is in progress. Notification of security issues must be done in a timely manner. IPsec VPNs are effective for site-to-site connectivity; however, they are not ideal for all remote access situations. For example, many employees must access corporate resources from unmanaged devices such as home PCs, public kiosks or PDAs. By contrast with IPsec, SSL VPNs allow granular access from any type of endpoint device (unmanaged or managed) if it complies with the minimum security policy that is in place in the organization. The SSL VPN maintains productivity for employees by enabling them to work from anywhere using any type of device.

Blending secure remote access into a comprehensive, modular security design is an important goal of Juniper's security services architecture.

**NAC**—NAC controls access to a network by way of policies, including pre-admission endpoint security checks and post-admission controls over where users and devices can go on a network and what they can do. NAC services control a user's initial access to the network and verify the integrity of the user's system. For example, NAC services can verify that the user's system has up-to-date antivirus software installed. NAC services should include support for remote software upgrades, including pushing upgrades to the user system (for example, to download a Windows service pack). NAC should support policies that determine the types of endpoints or user roles allowed to access designated areas of the network, and should enforce them in switches, routers and firewalls. NAC services should also coordinate with IPS for real-time detection and prevention of attacks that can originate from sharing within the internal network.

**DNS/DHCP**—The data center network infrastructure must support fast and reliable DNS and DHCP services. Issues with DNS cache refreshes and persistent DHCP bindings can be a potential security issue when customers are using cloud services. DNS/DHCP services must be configured correctly and run all the time so that policies that are tied to IP addresses can be applied quickly and accurately. DNS/DHCP services also are necessary to support VLAN operation and address pool reservations.

**AAA**—AAA services control whether users can log into data center systems and they determine which resources each user is permitted to access. The network security infrastructure should be able to leverage existing identity data stores, including Active Directory (AD) and Lightweight Directory Access Protocol (LDAP) servers.

Standard technologies exist to help different types of networks exchange identity and privilege information and share common notions of user identities. Standards include Security Assertion Markup Language (SAML), eXtensible Access Control Markup Language (XACML), and Interface for Metadata Access Point (IF-MAP). These technologies make it easier for network security devices to coordinate and enforce policies based on identity attributes. Products and solutions that provide security services should support these standards to ensure that identity and access information is shared among different networks.

To summarize, the integrated and virtualized security services resident in the network can provide benefits to users and applications that share the infrastructure. The comprehensive protection provided by these services can secure data flows into, within and between data centers. All of these services should be managed centrally and the infrastructure should enable distributed enforcement through the application itself and the supporting identity-aware security policies. As a group, security services increase the confidence, trust and agility with which virtualized services can be delivered.

## Application Services

In some cases, applications running on multiple hosts can benefit from network-resident services that can be spread across them efficiently to improve their performance and distribute loads. By including such services as part of an intelligent network infrastructure, the pooled resources of the cloud can operate much more efficiently. An important way to do this is to provide specialized services from systems logically and physically embedded in the network that offload work from other data center servers. These application services include application acceleration, DPI and global server load balancing.

### Application Acceleration

Application acceleration speeds performance for repetitive actions. For example, if a user accesses a document from a website, the initial download might take several minutes. With application acceleration, the document is cached following the initial download and subsequent requests can be done in seconds. Application acceleration can be tied to specific applications. For example, an application acceleration service can be configured to recognize and accelerate requests from an organization's SAP system.

Data center architects should consider deploying a system that supports acceleration for the different application tiers, and provides comprehensive capabilities in support of current and emerging application areas such as Web 2.0, SOA and SaaS. The acceleration solution should boost the performance of client/server, Web-based, and server-to-server applications, and it should speed webpage downloads. In addition, the acceleration solution should offload CPU-intensive functions such as TCP connection processing and HTTP compression from backend applications and Web servers. The application acceleration platform should be seamlessly expandable through stacking or clustering of multiple devices.

### Deep Packet Inspection

QoS is important to ensure application experience over large networks. QoS levels should be assigned and managed to ensure satisfactory application performance. DPI technology helps deliver advanced services by identifying applications based on key characteristics and by applying policies appropriate to them. For example, a DPI-enabled network element can apply QoS policies to an application to ensure preferred quality for video streams. Instead of the application adapting to network constraints, the network can adapt to application needs, providing a better user experience.

### **Global Server Load Balancing**

It is important to find ways to scale data center services without a linear increase in the hardware footprint, and to ensure that the design does not increase operational complexity. Global load balancing adds flexibility and adaptability to the data center network, so users always have access to applications and data, even if service to the primary data center is interrupted. This type of technology helps organizations support the technical and business goals of application and data availability without sacrificing performance. Server overload also can be reduced by using SSL offload and acceleration services.

### **Integrated Virtual Services**

Numerous and diverse services are needed to support the rich, complex network structure at the core of the virtualized data center. Delivering these services on existing single or limited purpose platforms can easily lead to appliance proliferation in the data center, as more and more platforms are introduced to deliver a richer set of security and application services. The resulting duplication of costs, physical space constraints, management overhead and organizational complexity can seriously inhibit growth of a successful data center or cloud. Many of these concerns can be resolved by introducing high-performance service processing platforms that support multiple services and stitch together with a common policy architecture and management structure.

### **Juniper's Approach**

Traditionally, organizations have faced a difficult trade-off between providing network security and delivering performance for applications. Juniper Networks eliminates this trade-off, making it possible for data centers to have the robust network security they require with performance that meets the most demanding application and user environments. Going further, Juniper Networks can consolidate network security for the data center into fewer devices—with centralized policy and visibility to improve significantly the operational efficiency of the data center environment.

The Junos operating system is the foundation of Juniper Networks security services. Junos OS provides a common language across Juniper's routing, switching, and security devices, reducing complexity in high-performance networks, speeding deployment, and simplifying provisioning and management. Because all Juniper networking products are built on Junos OS, data center architects can be confident that services will be compatible, and IT staff can draw on a common set of tools and experience.

Building on the Junos OS foundation, Juniper offers integrated solutions to meet the major security challenges in the cloud data center.

### **SRX Series Services Gateways for Comprehensive Security**

Juniper Networks SRX Series Services Gateways serve as the cornerstone of consolidated security within the data center, providing effective network segmentation, securing flows, delivering IPsec VPN encryption services, and offering IPS protection, NAT, and ALGs. By consolidating switching, routing and security in a single device, managers can economically deliver new applications, secure connectivity and deliver quality end user experiences. With its Dynamic Services Architecture, the SRX Series supports new services without sacrificing performance.

### **Unified Access Control to Secure LAN Access and Mitigate Insider Threats**

Juniper Networks Unified Access Control is a standards-based, scalable solution for adaptive access control that reduces threat exposure and mitigates risks. It guards mission critical applications and sensitive data, and it provides comprehensive control, visibility and monitoring.

The UAC approach to adaptive access control reduces the cost and complexity of delivering and deploying granular NAC. It also addresses challenges such as insider threats, guest access, outsourcing and off-shoring and regulatory compliance.

UAC is the industry's first NAC solution to offer full Layer 2 through Layer 7 enforcement capabilities. It is based on industry standards (802.1X, RADIUS, and IPsec) and open standards (Trusted Network Connect), including IF-MAP, which empowers UAC to integrate with third-party network and security devices.

### SSL VPN for Secure Remote Access

Juniper Networks SA Series SSL VPN Appliances provide enterprises and service providers with remote access and sophisticated partner and customer extranet features. SA Series appliances enable organizations to enforce differentiated access to resources based on user roles and groups. These appliances are available with a baseline software feature set or an advanced feature set that includes options for more complex deployments.

### WXC Series Application Acceleration Platforms and WXC Client

Juniper Networks WXC Series Application Acceleration Platforms accelerate mission critical applications over wide area links, providing compressed output that ranges from 2 Mbps to 155 Mbps rates. Each platform can support multiple remote sites, and multiple communities of WXC Series devices can be configured to support an unlimited number of locations.

The WXC Series uses compression and caching to reduce the amount of data actually flowing across wide area links. It does this by eliminating redundant data patterns and boosting connection capacity to accommodate a greater volume of traffic. It speeds the performance of specific applications and protocols over the WAN, cutting response times and optimizing traffic flows to deliver a more LAN-like experience for remote office users. Applications can make the most efficient use of available links and bandwidth to optimize performance and prioritize data traffic.

Juniper Networks WX Client is Windows-based software for mobile end users that provide LAN-like performance for applications. Installed on the end user's laptop, the WX Client improves application response times by applying disk-based caching, compression, and protocol acceleration techniques to WAN data traffic. Enterprises can now enable cost-effective, dynamically provisioned, pervasive application acceleration regardless of user location.

## Management, Orchestration, and Automation

With an understanding of the attributes of the three major traffic processing areas of the data center infrastructure—compute/storage, network, and security/application services—we can now turn our attention to the challenges of managing the data center cloud in the most efficient, flexible, and scalable manner. It is a formidable challenge to interconnect and supervise the growing number of physical and virtual devices in the cloud in a coherent, efficient way. Management complexity grows as more devices and users are added. To make data centers and the cloud truly responsive, all components must come together in a well-orchestrated ensemble under the IT organization's control. Figure 7 shows the management, orchestration and automation functional area of the reference framework.

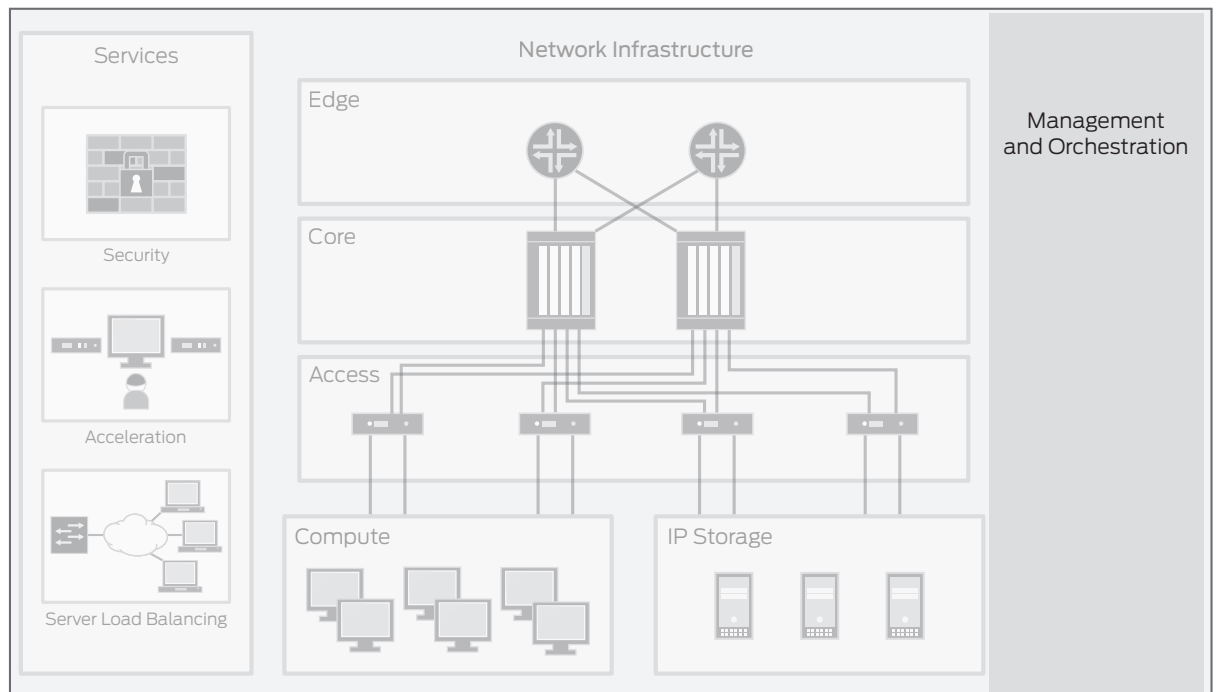


Figure 7. Management, orchestration, and automation



The term “orchestration” refers to the automated arrangement, coordination and management of components (compute, storage, network and service) to meet IT and business requirements.

In addition to the automation that is already an integral part of each component, orchestration requires that components interoperate with each other, business processes and rules are implemented properly, and end-to-end services are delivered completely and reliably. Orchestration takes the data center a major step beyond localized automation to encompass fully coordinated visibility and control over the data center’s disparate elements.

Because orchestration is complex and depends heavily on an organization’s specific systems, its requirements are best met by a network orchestration platform that is open and extensible for integration with diverse application and management systems. The network orchestration platform should support comprehensive network management functions and use industry standard APIs to enable integration with management and application systems. It should also provide development tools, including a software development kit (SDK), so that organizations can extend and adapt the platform to create their own orchestration environments.

### Profile of an Effective Orchestration Platform

To be effective in the virtualized data center, an orchestration platform should include the following capabilities.

**Auto-discovery**—To deliver on efficiency and productivity, a network orchestration platform should have robust, extensible and standards-based discovery functions that automatically recognize the network elements in the data center network.

**Modular, extensible platform and application architecture**—To be effective in a diverse, multi-vendor environment, an orchestration platform should allow simple and robust insertion of application modules based on well understood SOAs, Web services, and application design principles (for efficiency, agility, and scale). This should include live application hot insertion, deletion and change. It should also include well-understood methods for integration of functions between applications and between applications and the underlying platform, to enable extensibility and broader application relevance to users.

**Open application development environment for ecosystem innovation by partners and customers in addition to flexible enhancement of applications by the platform supplier**—A network orchestration platform should allow safe, efficient and flexible development of new application logic that can leverage the intelligent networked environment it is orchestrating. The new application logic can facilitate understanding of diverse network attributes, distribution of application functions across a network base, and export/import of network and application intelligence, thereby helping to optimize and tune application performance.

**Extensive and explicit use of standards**—To be effective, the orchestration platform must leverage open standards such as Worldwide Web Consortium (W3C), Internet Engineering Task Force (IETF), and Trusted Network Computing Group (TNC) standards for open communication between network and application elements, and it must ensure smooth interoperability with existing and future management systems.

**Resilient, scalable, and distributed architecture**—As virtual data centers and cloud computing infrastructures evolve, they often require distribution of management controls to multiple distributed sites, to share responsibilities among distributed teams, extend controls to new and distant infrastructures, and support HA and disaster recovery. An orchestration platform for this environment should be built on an architecture of replicated and distributed platforms and components, so that connectivity is maintained regardless of location, and access to control information is available despite changes in infrastructure availability and performance.



**Flexible, virtualized, and role-based user access**—Because network and application managers can be located anywhere and may be called upon to perform important functions at virtually any time, the orchestration platform should allow secure access from multiple network and device locations by leveraging flexible user interface and access technologies.

**Logically centralized, physically distributed**—The orchestration platform must support simple, effective and consistent management, monitoring and control over the extended physical infrastructure.

**Network-wide correlation**—The orchestration platform must be able to supervise the interplay of all network elements, balancing resource optimization and QoS.

**Application visibility**—The orchestration platform must promote application visibility for assured application delivery. Ultimately, orchestration is important because it maps technology components to service components and serves as a reference point during provisioning and operations for the cloud. Orchestration is an emerging area that will continue to develop as cloud solutions mature.

### Management Infrastructure Supporting Cloud-Level Orchestration

To realize the full benefits of orchestration, many elements of conventional network management are needed to support the higher level orchestration functions. A sampling of these includes:

- Consistent, standards-based programmatic interface for management functionality across the network infrastructure (examples include NETCONF and SNMP).
- Consistent, commonly available Web interface for access to management functions.
- Schema-driven element management.
- Standards-based APIs, using technologies such as XML/SOAP, to make management accessible from the orchestration gateway.
- Support for security services, including threat analysis, protection, and reporting, to identify risks, ensure reliable operations, meet regulatory compliance and satisfy SLAs. Security Information and Event Management (SIEM) and Security Information Management (SIM) should be employed to gain broad and deep visibility into the network.
- Layer 4 to Layer 7 application services to support visibility and performance diagnosis. Real-time, accurate and useful monitoring data must be obtained from the network to support network management and control. Data collection can be flow-based (using methods such as sFlow, jFlow or NetFlow) or event-based (using methods such as SNMP or system logging).
- Dynamic service and application management, including robust policy control. The central management of network and security events, network and application flow data, vulnerability data, and identity information greatly improves the ability to meet IT security objectives.

Figure 8 shows the entire Juniper Networks management framework.

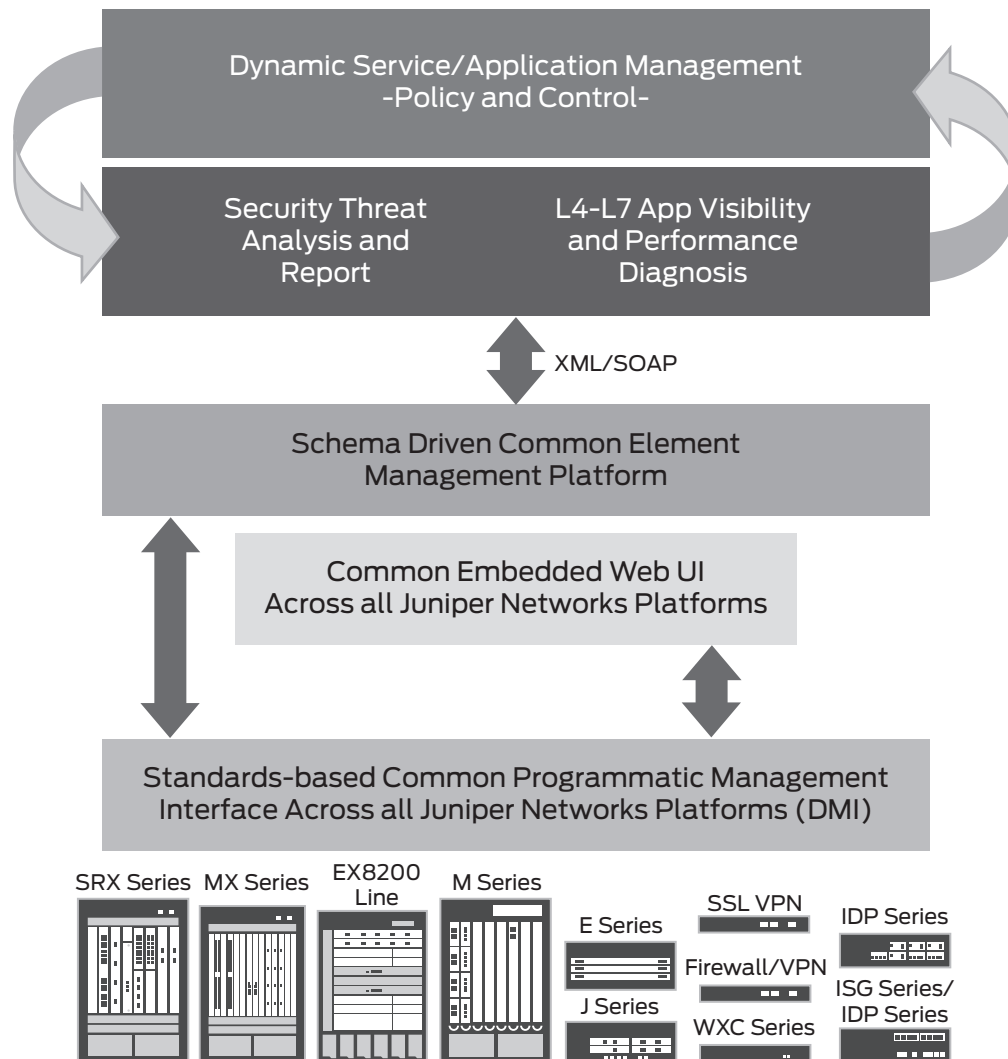


Figure 8. Juniper Networks management infrastructure

**Juniper’s Approach**

To enable flexible, effective and scalable orchestration of the network and its services, Juniper provides a comprehensive suite of orchestration, management and development/integration platforms to support management at multiple levels of infrastructure operation and control. These range from powerful device-level capabilities to rich, user-friendly management applications, to a broad, innovative and open platform for orchestration and integration of multiple applications and services.

The following section presents the highlights of the capabilities at each of these levels.

**Device-level manageability**—Juniper’s network elements use an open, standards-based XML/NETCONF Device Management Interface (DMI) to support all of the network and security platforms. DMI allows devices to maintain management information and structure in an XML schema, allowing broad discovery of attributes by Juniper and other vendors’ management applications. Over time, extensions to DMI will allow Juniper’s applications to discover and work with other suppliers’ network infrastructures.

**Robust, centralized management**—Juniper offers a comprehensive suite of applications for managing devices, network-level service functions (performance, security), user and application-related policies, and integration of multiple network-level applications. Juniper has worked with multiple partners to integrate important additional functionality into its management and orchestration environment. For example, Juniper has incorporated advanced

routing infrastructure visibility through its partnership with Packet Design on Route Insight Manager. Moreover, Juniper and IBM have collaborated extensively to enable applications in IBM's Tivoli suite of IT and data center management tools, such as Tivoli Netcool, to interwork efficiently with Juniper's management infrastructure. Further integration is possible as requirements evolve using the open XML/Web Services Description Language (WSDL) environment of Juniper's orchestration suite to communicate Tivoli's end-to-end aware systems management interfaces.

Centralized monitoring—Juniper provides powerful centralized monitoring and log analysis for visibility into traffic and security events. Juniper's architecture provides a central repository for traffic flow and security event collection for analysis and notification of important security events. This capability also addresses the logging requirements specified in the National Institute of Standards and Technology (NIST) recommended best practices, as described in the *NIST Guide to General Server Security*.

### Junos Space Juniper's Open Network Orchestration Platform

Junos Space is Juniper's groundbreaking, open and extensible network orchestration platform for developing, hosting and integrating applications that enable rapid delivery of services and reduce cost and complexity of operations. Junos Space allows partners, customers and independent software vendors (ISVs) to develop and deploy simple, smart network applications. The platform includes a scalable runtime environment with multitenant, hot-pluggable network application support, a complete rapid network application development framework and a Web 2.0 user interface. With Junos Space, orchestration applications can be supported along with other applications on the same platform. Junos Space opens up an entirely new realm of integration and orchestration possibilities never available before in network suppliers' portfolios. It creates the opportunity to flexibly and programmatically leverage an intelligent network cloud for the full range of network-enabled application deliveries required in emerging IT infrastructures.

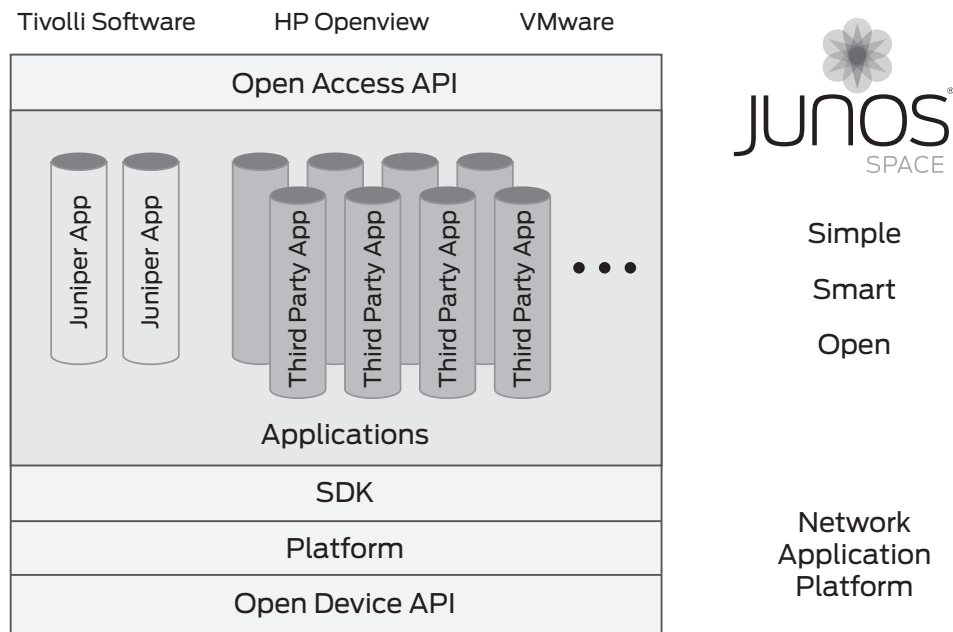


Figure 9. Junos Space infrastructure

Junos Space uses the same design principles and open systems approach of the Junos operating system to enable programmability at the network application layer. It is a multipurpose application platform designed to be simple, smart and open.

## Automation Based on Junos OS

To complement the management and orchestration capabilities available at the network level in the Junos Space and Juniper Networks Network and Security Manager platforms, Juniper offers a suite of scripting tools in Junos OS. The tools automate performance of specific operational tasks in ways that can augment functions available in the other application platforms.

Built directly into the Junos operating system, Junos OS script automation is a powerful and flexible onboard toolset available on all Junos OS platforms, including routers, switches, and security elements. Junos OS scripts support automated:

- **Configuration**—Simplify and enforce business rules to avert human errors and optimize network availability
- **Events**—Automate reactive and proactive actions in response to network events for self monitoring and self diagnostics
- **Operations tasks**—Customize and streamline manual tasks to increase operational efficiency and maximize staff effectiveness

To support Junos OS network administrators, an online script library is available as a repository for the most commonly used scripts. To read more about Junos OS script automation, visit: [www.juniper.net/us/en/community/junos/script-automation/](http://www.juniper.net/us/en/community/junos/script-automation/)

## Data Center Types

Not all data centers are the same. Depending on their use, size, and desired results, the requirements and design can vary widely. Some demand the lowest possible latency and highest possible availability, while others require comprehensive attention to QoS, scale, and HA. Still others require costs minimization by opting for fewer features. In this section, Juniper provides examples of the five most common data center types, along with recommendations on how to build those networks.

### Transactional Production Data Center Network

High-speed computing in particular is critical to the success and profitability of today's financial institutions. At high scale, every nanosecond of latency accounts for either a profit or loss. Therefore, businesses such as financial services cannot afford any performance or security risks.

Typically, a financial services network is extremely complex, using a myriad of devices and services to support a high-performance computing infrastructure. Unfortunately, the traditional data center design does not address these specific requirements for guaranteeing predictability and low latency associated with trading platforms, algorithmic trading applications, and market data distribution systems.

Juniper Networks realizes how essential it is to give a simplified solution serious consideration. Considering only high performance is insufficient and incomplete. Additional integral aspects/requirements such as scalability (future growth), simplification (no high learning curves), and security (intrusive threats) are needed. In short, a high-speed, secure, easy to operate and highly secure network infrastructure defines a transactional data center solution.

At the heart of this solution are Juniper's state-of-the-art Ethernet switches: Juniper Networks EX2500 Ethernet Switch, EX4500 Ethernet Switch, EX8216 Ethernet Switch, as well as MX Series 3D Universal Edge Routers and the SRX Series Services Gateways—offering the most scalable 10 GbE wire-rate data center solutions on the market today.

By deploying a transactional production data center solution, businesses can:

- Create Point of Delivery (PoD) in multiples of 500 10 GbE servers for modular scalability and optimal control, enabling support for 25 server racks per a single logical switching branch (EX8216 and MX Series)
- Achieve up to 120 Gbps firewall performance (per device) for maximum throughput capacity at the PoD
- Achieve up to 40 Gbps IPS performance, providing application-specific context recognition and security processing
- Implement Junos OS automation and instrumentation at the core
- Create 1:2.5 uplink oversubscriptions (nonblocking access and core) (Juniper Networks EX8200 line of Ethernet switches)
- Deploy rich multicast functionality
- Create deep buffers for reliable recovery of downlink saturation

## Content and Services Hosting Production Data Center Network


With the emergence of cloud computing and with more services being migrated to an IP-based delivery model, data centers are playing a more critical role for hosting content and services. For example, hosting providers are now offering cloud-based infrastructure as a service. A broad array of online retail options have emerged that position the Internet as a key enabler of their business. In addition, “pure play” cloud providers offer extensive development and delivery platforms as well.

These business models pose a very strict set of requirements to the network. From an end user perspective, high availability translates into whether the business is open or not, and the speed of the traffic effectively translates into the quality of the customer experience. Additionally, from the provider’s perspective, key functions in the network are needed in order to support these new business models. For example, workloads and the allocation of resources to specific workloads often change, and the network needs to support the process of changing workload distribution as well as delivering the new workloads to the end user. Security functionality needs to support virtualization technology and be granular enough to handle specific applications and users. The network also needs to seamlessly support applications running across multiple sites while retaining a consistent security policy across the whole environment.

Juniper’s content and services hosting production data center solution consists of the following products:

- MX Series devices for data center edge and collapsed core connectivity
- EX8200 line for data center core LAN connectivity building high capacity LAN PoDs
- Juniper Networks EX4200 Ethernet Switch for flexible GbE or 10 GbE server access connectivity
- Juniper Networks SRX5000 line for large scale security and services processing
- Juniper Networks SA6500 SSL VPN Appliance clusters for effectively and securely managing customer out of band management
- JUNOS Space

Figure 10 shows the content and services hosting production data center network.



**Figure 10. Content and services hosting production data center network**

The solution components described above render a data center that scales in increments of 5,000 10 GbE connected servers (per PoD). It supports full L3 and L2 functionality from the server access port throughout the network, over the WAN, and onto server access ports in additional data center locations using a combination of basic L3 functionality through L2 VPLS services. Additionally, this solution enables streamlined automation of network operations to align with overall data center operations at application and platform levels.

## High-Performance Compute (HPC) Production Data Center Network

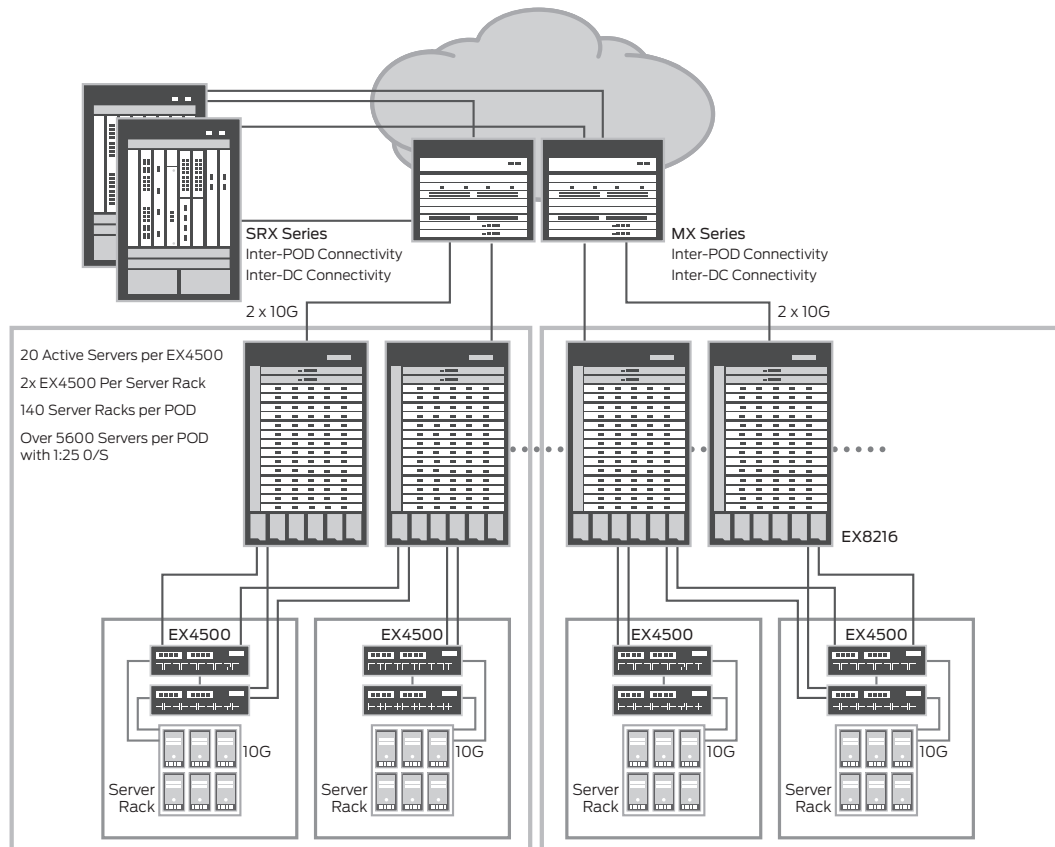
Scientific innovation in emerging fields such as nanotechnology, biotechnology, genetics and seismic modeling are driving specific requirements in data center networks. With the utilization of general-purpose hardware, many organizations are finding it more cost effective to leverage grid computing (High Performance Compute Clusters or HPC/C) for their intensive compute tasks. This technology is based primarily on grouping multiple CPUs together, distributing tasks among them and collectively completing large amounts of calculations. In the network, utilizing 10 GbE has distinct benefits not just for performance, but also in terms of the low cost-to-bandwidth ratio.

Typically, organizations trying to build these capabilities share common requirements for the network. The speed at which messages can transfer from one server to another has a significant incremental effect on the capacity and compute speed of the overall grid solution. Hence, high throughput and low latency are the most important characteristics of the network infrastructure. Additionally, achieving a form of local HA to avoid any interruption is highly desired. However, in customer types where a more cost-effective solution is required, device-level HA is an alternative.

Juniper Networks offers a high-speed, sub microsecond latency solution at a very attractive price point to accommodate different customer business models in this space. The components of this solution are the following:

- EX2500 switches connecting the server access ports
- EX4500 switches aggregating multiple EX2500 access switches, serving as the HPC PoD core
- EX8200 line switches serving inter-PoD connectivity in cases where larger scale is required.

Figure 11 shows the high-performance compute production data center network.



**Figure 11. High performance compute production data center network**

The above components create a data center network that scales in increments of 200 servers and supports sub microsecond latency across servers on the same rack with full L3 functionality and high availability.

## Enterprise IT Data Center

Traditionally in the enterprise, the data center has been designated as an IT “cost center” in that it does not directly contribute to generating revenue. In this scenario, the primary objective of the data center network is to be a business enabler, providing access to business applications, resources and services (such as Oracle, CRM, ERP, and others) for employees and other network users. With this in mind, the key requirements that surface include HA and low latency performance that are designed to improve the user experience and keep the workforce productive. From an IT perspective, the security of the network and centralized management capabilities are important for keeping operations consistent with other parts of the enterprise network. Finally, from a business perspective, this must be implemented while controlling costs and showing a direct impact to business value.

Typically, customers in this space will look to leverage network innovation and adopt technologies to address both types of requirements. For example, server virtualization, I/O convergence and MPLS/VPLS (virtual private LAN service) for seamless multisite connectivity are common in these environments. Consequently, the data center network must deliver features that support these technologies. Additionally, this type of data center is applicable for customers across a wide variety of industries, including healthcare, retail, manufacturing, education, and energy and utilities.

Juniper Networks offers a high-speed, sub microsecond latency solution at a very attractive price point to accommodate different customer requirements in this space. The components of this solution include the following:

- EX4200 and EX4500 switches connecting the server access ports
- EX8200 line switches aggregating multiple access switches, serving as the PoD core
- MX Series devices for data center edge and collapsed core connectivity
- SRX5000 line devices for large scale security and services processing
- SA6500 clusters for effectively and securely managing customer out of band management
- Junos Space

Figure 12 shows an example of an enterprise IT data center network.

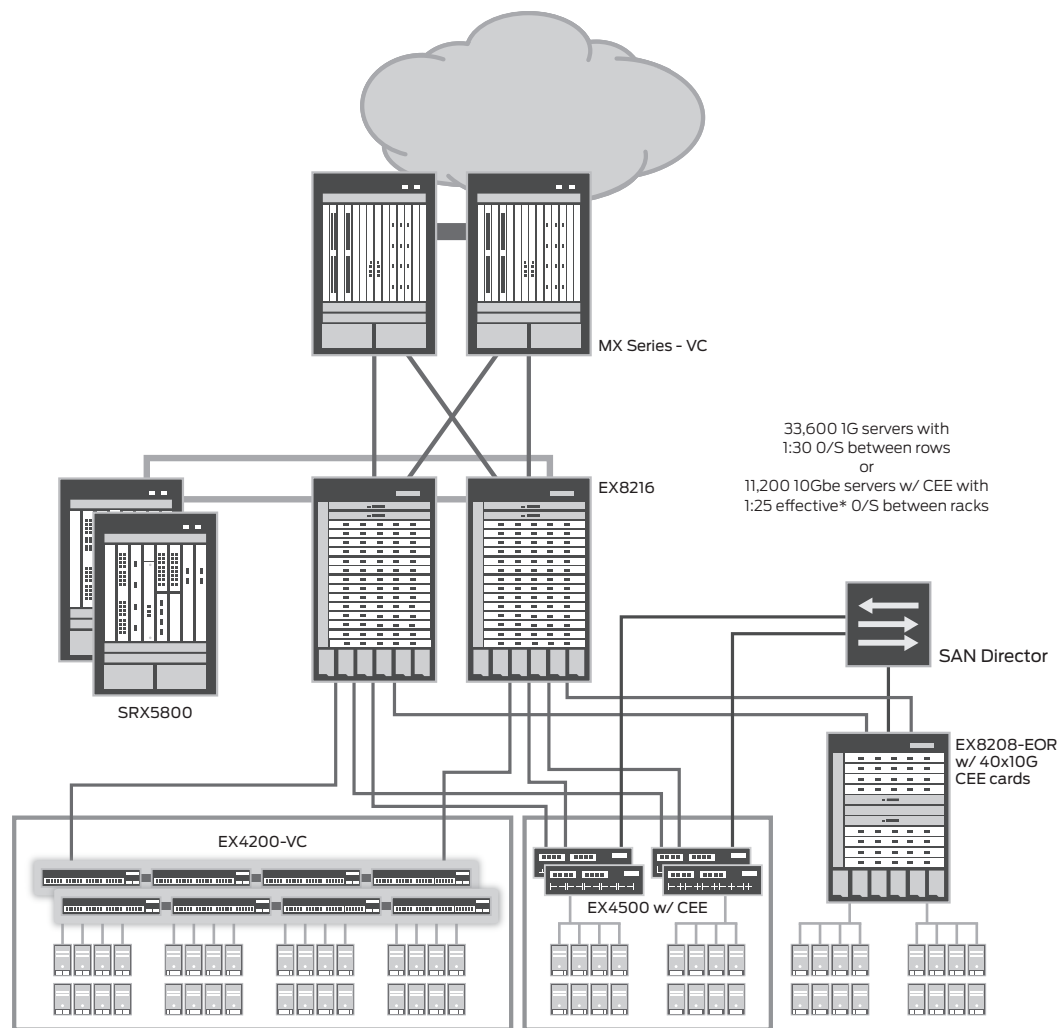


Figure 12. Enterprise IT data center network

The above components create a data center network that scales to over 30,000 GbE servers and over 10,000 10 GbE servers with HA, the ability to provide security services at scale, and an infrastructure for simplified and centralized management.

### Small and Midsize Business IT Data Center

With today's advanced, high priced networking technologies, most SMBs face serious challenges in being able to afford the latest data center infrastructure technologies while still trying to remain competitive and profitable. The typical SMB data center consists of a 3-tier architecture, which inherently adds complexity. In addition, huge operational overhead also remains a challenge for SMB IT teams, not only during the implementation phase but also during the adoption of new technologies, creating lag times and delays in ROI. Finally, scalability and security are additional major concerns for the cost conscious small-to-midsize business owner. Businesses in this situation require a reliable, easy to deploy and manage, and cost-effective option for their data center network buildout.

Juniper recognizes the importance of these SMB challenges by offering a low cost, less invasive approach for deploying a cloud-ready common switching and routing infrastructure based on fewer required devices, i.e., routing, switching, and security appliances that run on a single operating system across all platforms.

At the crux of this solution is Juniper's latest in switching functionality and network security—the Juniper Networks EX2200 Ethernet Switch, EX4200 Ethernet Switch, Juniper Networks SRX650 Services Gateway and Juniper Networks SA2500 SSL VPN Appliance. These devices are deployed in a 2-tier architecture that consists of access and core switches, as well as consolidated security functionality.



By deploying this hybrid cloud-ready solution, architects and designers can:

- Connect more than 480 single homed GbE servers (two EX4200s)
- Provide scalability with a maximum of 10 EX4200 Virtual Chassis switches at the core layer (EX4200)
- Support 7 Gbps stateful firewalls and 900 Mbps IPS (SRX650)
- Provide a maximum of 100 SSL VPNs (accelerated and secured) (SA2500 SSL VPN Appliance, Juniper Networks Junos Pulse, and the Juniper Networks WXC250 Application Acceleration Platform)
- Automate all EX Series Ethernet Switches and SRX Series Services Gateways (Junos OS).

## Conclusion

To succeed in the face of new economic, management, and service delivery imperatives, organizations need to embrace new architectures and designs to deliver flexible, secure and high-performance applications for their users. Juniper's Reference Architecture for Data Center Networks and the Cloud leverages fresh and powerful thinking throughout the data center network, bringing customer advantages that have not been available until now. Juniper's innovation and excellence in silicon and platform architecture, modular and extensible network operating software, security and application enhancement services and open, extensible, standards-based orchestration of applications enable a level of responsiveness and return on investment not achievable with legacy or alternative approaches. It is a new network world with Juniper's architecture for data centers and the cloud.

Leveraging our innovations to simplify, share, secure and automate the network, organizations can implement newly virtualized and cloud-ready solutions in their data centers. Moreover, this can be accomplished whether building a private cloud to serve users in their own organizations, a public cloud to deliver services over publicly available networks to a broad user population or a hybrid cloud that blends characteristics of public and private clouds.

By rethinking legacy designs and adopting the Juniper approach, organizations can build data center networks that support a continuously expanding user base and deliver a high quality experience for end users. They can deliver performance at scale and carrier-grade reliability to keep pace with escalating information processing demands, and they can advance the economics of high-performance networking by increasing speed, lowering the cost of application and service deployment, and reducing total cost of ownership.

Our architecture allows organizations to implement data center and cloud computing networks at a pace and with an approach compatible with overall business objectives. The architecture can be applied to targeted portions of the data center network or it can serve as a roadmap for designing entirely new data center infrastructures. Regardless of the existing environment, planning horizon or business objectives, Juniper provides the industry's most innovative and compelling reference architecture to serve business applications that run in the full range of data centers and the cloud.

## Appendix A Juniper Products for the Cloud-Ready Data Center Network

### Switching

EX Series Ethernet Switches: Deliver unmatched scale and performance for Layer 2 and Layer 3 data center networks. The high-density, high-performance EX8200 line of Ethernet switches supports data center and cloud computing environments. Virtual Chassis technology enables up to 10 interconnected EX4200 line switches to operate as a single device, reducing management overhead.

### Routing

Juniper Networks M Series Multiservice Edge Routers: Combine best-in-class integrated virtualization and traffic management with unmatched reliability, stability, security and service richness to connect the cloud data center to the WAN. The M Series routers provide a wide range of connectivity options without compromising performance.

MX Series 3D Universal Edge Routers: Provide Ethernet switching capabilities coupled with the carrier-class routing features that customers expect from Juniper to support advanced virtualization of network infrastructures and traffic management between data centers and WANs.

## Security Services

**SRX Series Services Gateways:** Combine routing, switching, application services, and user- and application-aware security within a single device.

**Unified Access Control:** Provides powerful identity- and role-based access control that increases agility in service deployment and overall quality of experience. UAC can be deployed within a data center or across an extended enterprise to protect networks and applications.

**SA Series SSL VPN Appliances:** Provide scalable, simplified and secure remote access from multiple remote networks and platforms (SSL VPN) to data center resources.

## Application Services

**WXC Series Application Acceleration Platforms:** Support mission critical applications by accelerating applications over WAN links.

## Operating System

**Junos OS:** Integrates routing, switching, and security services, and offers the power of one operating system to reduce complexity, achieve operational excellence, and deliver dynamic services with lower total cost of ownership.

## Unified Network Client

**Junos Pulse:** Delivers integrated connectivity, access, acceleration and security anytime/anywhere, while drastically simplifying the user experience with a dynamic, standards-based multiservice network client.

## Orchestration and Network Management

**Network and Security Manager:** Provides a single pane of management for the entire network infrastructure, including routing, switching and security devices.

**Juniper Networks STRM Series Security Threat Response Managers:** Collect events and alerts from different Juniper and third-party products, aggregating and delivering them to an enterprise-wide threat management view.

## Junos Space Platform

Simplifies application development and hosting, allowing you to develop and deploy simple, smart network applications. Junos Space includes a scalable runtime environment with multitenant, hot-pluggable network application support, a complete rapid network application development framework and a Web 2.0 user interface.

## Technical Services

**J-Care Technical Services:** Offer unlimited access to Juniper's support centers online or by telephone. Organizations get a family of support services that include immediate software updates and hardware replacement options.

**J-Care Efficiency, J-Care Continuity and J-Care Agility services:** Add the automation elements of Advanced Insight Services (AIS) to reduce operating expenses and simplify operations.

## 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](http://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](http://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 2010 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. 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.

8030001-003-EN May 2010

 Printed on recycled paper