Enterprise Data Center Network Reference Architecture

Using a High Performance Network Backbone to Meet the Requirements of the Modern Enterprise Data Center

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408.745.2000
1.888 JUNIPER
www.juniper.net

Part Number: 803001-001 May 2008

Table of Contents

Executive Summary

The data center is an extremely critical corporate asset. As such, the data center network that connects all critical servers, applications and storage services is increasingly a key component that needs to be carefully planned and managed in order to meet the growing performance demands of users and many network-centric 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, cost-effective and business supported data center networks.

This reference architecture proposes practices, technologies and products that help data center architects and engineers responsible for answering the requirements of designing modern data center networks.

Introduction

Trends and Challenges

According to research conducted by Nemertes (2006), 91 percent of benchmarked companies were under compliance constraints, and more than 50 percent of companies consolidated their dispersed data centers into fewer but larger data centers in the last 12 months, with even more planning to consolidate in the upcoming 12 months. While enterprises are consolidating their data centers and centralizing their servers, the opposite is happening with the employees themselves. More than 90 percent of employees work remotely and more companies are opening a larger number of branch offices to get closer to their customers. These two divergent trends are causing tremendous strain on the enterprise WAN connectivity as more people are trying to access applications that are highly centralized. Further, performance of these applications is becoming a critical bottleneck in terms of employee productivity.

Another interesting trend is that servers are continuing to grow at a high annual rate of 11 percent, while storage is growing at an even higher rate of 22 percent; both of which are causing tremendous strain on the data center's power and cooling capacity. According to Gartner, OS and application instability is increasing the server sprawl with utilization rates of 20 percent. Also, CIOs are increasingly demanding utilization/efficiency reports on servers and storage, which is leading to an increased adoption of virtualization technologies such as VMware and XenSource.

The major challenges identified by customers regarding their data centers include the following:

- Power capacity
- Cooling
- Increasing growth of the data centers
- Availability
- Disaster recovery
- Operational issues concerning change management and controlling operational costs

Enterprises are demanding "zero-downtime" or "extreme availability" as businesses are becoming increasingly global and are functioning 24x7.

Gartner (2007) identifies the following trends from a list of the top ten disruptive technologies in the data center.

- Unified Communications
- The Web as a platform to deliver applications
- Virtualization that goes beyond consolidation
- Mashups and composite applications
- Green IT

Emerging applications that use Service Oriented Architecture (SOA) and Web services are increasingly computational and network intensive.

While businesses are attracted to the cost savings of consolidating data centers and therefore are reducing the number of facilities and operating locations, architects are faced with the challenge of designing a data center that centralizes servers and applications that are accessible from a variety of locations (see Figure 1). Throughout all of these challenges, today's data center must meet the performance requirements needed to ensure satisfactory user experience without compromising security and compliance.

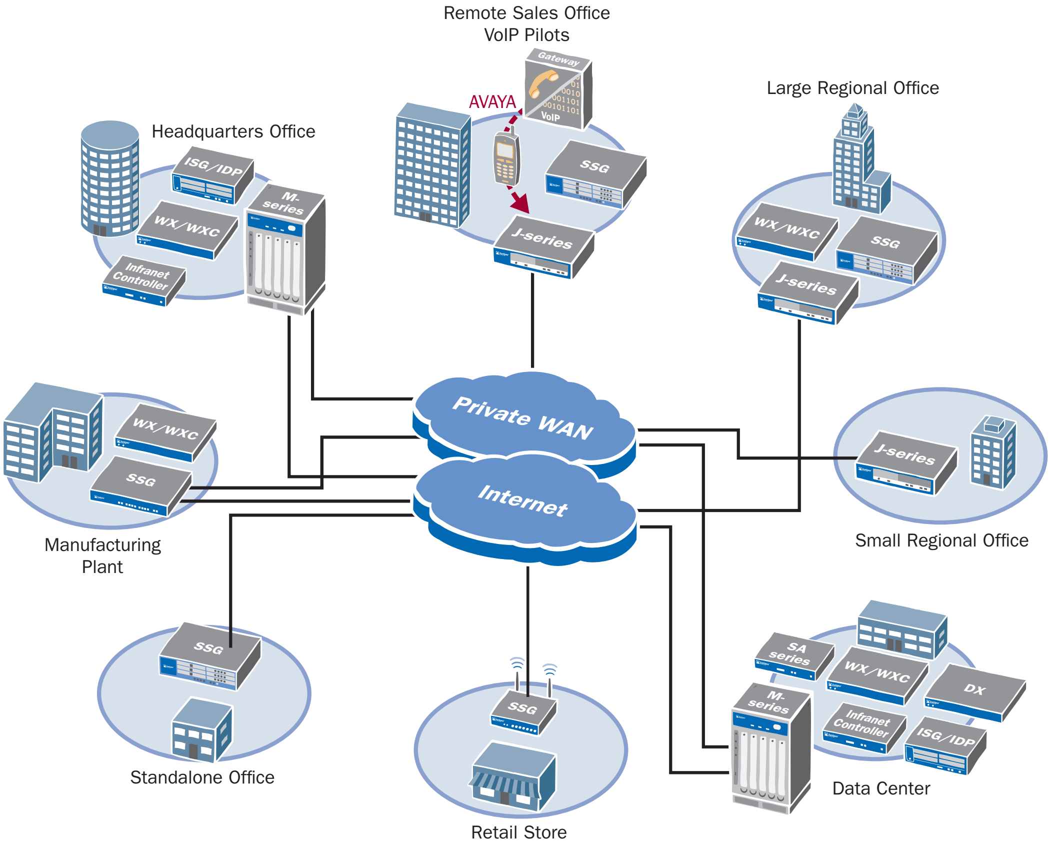

Figure 1: Location-Based Perspective of the Enterprise Network

In addition, Gartner (2007) asserts that 50 percent of the Ethernet switch ports within the data center are used for switch interconnectivity.

Simply designing a data center that only deploys more servers, storage and devices significantly increases network complexity and cost. Organizations must change the way they view their data center network architecture in order to maximize efficiency gains from technologies such as virtualization. The architecture must use virtualization capabilities such as MPLS and virtual private LAN service (VPLS) to enable an extremely fast, high-performance data center backbone network, in order to meet the performance demands of the consolidated architecture. The data center network also must offer required components such as security, performance acceleration, high density and a resilient network infrastructure. These critical components help ensure that users sustain the performance needed to succeed in their jobs, and that the network supports their business goals. This document shares Juniper Networks best practices in designing a highly efficient, secure, scalable and flexible data center network. This document also showcases advanced network technologies such as high-density next-generation Ethernet switches, application delivery controllers and WAN acceleration that can be employed to create a seamless user experience, irrespective of the location on the network.

Juniper Networks Approach and Solution

The Juniper Networks strategy for designing the data center network uses an open systems approach that enables enterprises to design a high-performance data center network that consolidates network elements into fewer networks and employs fewer network devices. This approach simplifies network architecture, enables operational efficiencies, and offers data center networks that are agnostic to multiple media types.

The architecture virtualizes critical network infrastructure components and functionalities such as security, load balancing and applications acceleration, and it deploys and manages based on a combination of business as well as technical heuristics. The architecture optimizes network performance and increases efficiencies within the network infrastructure. The architecture also automates management of the network infrastructure by connecting smoothly into the customer's existing management frameworks and third-party tools like IBM Tivoli, as an example.

Scope

The purpose of this document is to provide our partners, customers and potential customers with a data center network architecture that mitigates business risk and supports the modern, consolidated data center. This document addresses the following topics:

- Network infrastructure
- Security
- Connectivity
- Performance aspects of the data center infrastructure

In addition, this document provides design guidance for the data center network, the inter-data center and associated connectivity. Discussions focus on the following network devices:

- Routers
- Switches
- Firewalls
- Intrusion prevention systems
- VPN access devices
- Application front ends
- WAN acceleration products.

NOTE: Because application-specific components such as operating systems, processing machines, databases and storage arrays are out of scope of this solution, they are not addressed in this document.

Target Audience

- IT managers and security managers
- Systems engineers
- Network analysts and engineers
- Network administrators

Enterprise Data Center Network Design Considerations

The following section summarizes some of the technical considerations for designing a modern day data center network that must support consolidated and centralized server and storage infrastructure as well as enterprise applications.

NOTE: The design considerations discussed are not necessarily specific to Juniper Networks solutions and can be applied universally to any data center network design, regardless of the vendor.

The functional data center network design model (Figure 2) considers key design attributes. Each of these attributes is summarized in the sections that follow.

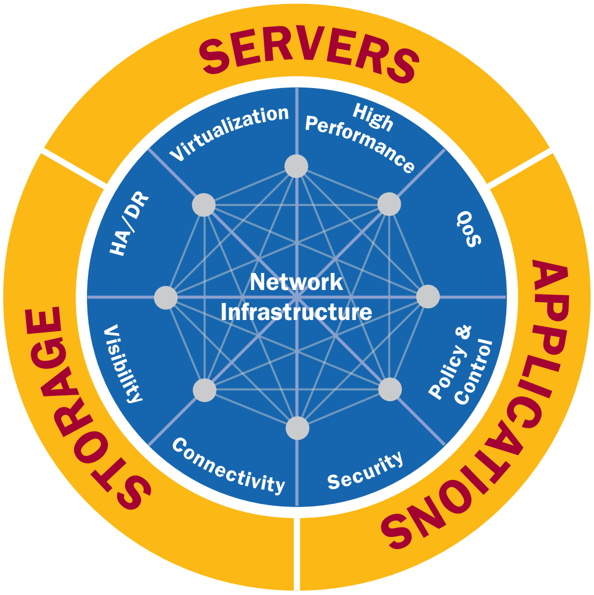

Figure 2: Center Network Functional Design Model

As data centers become consolidated, more servers become centralized. The data center can be viewed from the perspective of the different groups of people interacting to create a highly available and functional end user requirement for the enterprise. These groups typically comprise storage, server, application and network groups. Observing all of the installed devices in the data center, we obviously see large racks of servers (X86 servers, blade servers or mainframe systems), different types of storage switches that use Fibre Channel and InfiniBand, and a variety of applications (Oracle, SAP, Microsoft) that utilize these resources to deliver business requirements. These three silos are connected through a fast, secure and reliable data center network fabric which forms the fourth silo of systems and devices in the data center. The critical attributes for designing today's data center for extreme availability and superior performance are as follows:

- Virtualization—network, server and storage
- High Availability/Disaster Recovery (HADR)
- Visibility—not only in the network traffic and security events, but also into application traffic
- Connectivity—ubiquitous connectivity to disparate sets of resources
- Security —security and compliance
- Policy and Control—centralized policy and control
- Quality of Service (QoS)
- High Performance—applications, storage, servers and the network

Virtualization

As described in Wikipedia, virtualization is a technique for hiding the physical characteristics of computing resources from the way in which other systems, applications or end users interact with those resources. This means making a single physical resource such as a server, an operating system, an application or a storage device appear to function as multiple logical resources; or making multiple physical resources such as storage devices or servers appear as a single logical resource. Virtualization also means making one physical resource appear, with somewhat different characteristics, as one logical resource.

From a network virtualization perspective, there are various technologies that provide data, control and management plane virtualization. A data plane virtualization example is a single physical interface to provide security to multiple network segments using 802.1q VLAN tagging. From a control plane virtualization perspective, multiple routing domains and protocol instances are other examples. A management plane virtualization example supports multiple logical firewall/VPN security systems that use Virtual Systems (VSYS) for true multi-department or multi-customer environments, such as large enterprises or service providers who offer managed security services all in a single device.

High Availability Disaster Recovery

High Availability Disaster Recovery (HADR) is a key requirement from the data center network perspective and must be considered not only from what is happening within the data center, but also across multiple data centers. Network HA should be deployed by using combinations of link redundancy (both external and internal connectivity) and critical device redundancy to ensure network operations and business continuity. In addition, using site redundancy (multiple data centers) is critical to meeting disaster recovery and regulatory compliance objectives. Moreover, devices and systems deployed within the confines of the data center should support component-level HA, such as redundant power supplies, fans and routing engines. Another important consideration is the software/firmware running on these devices, which should be based on a modular architecture that provides features such as in-service software upgrades (ISSUs) to prevent software failures/upgrade events from impacting the entire device. Software failures/upgrades should only impact a particular module, thereby ensuring system availability.

Visibility

It is important to have visibility into network traffic and security events in order to effectively maintain and manage the resources. It is critical to collect IP traffic flow statistics to give enterprises insight into data flow, resource utilization, fault isolation, capacity planning, tuning and offline security analysis. WAN utilization and user-level visibility can help IT better support application performance by leveraging network services and other resources. Security visibility is crucial to granularly view security events to help determine how these events get handled. Further, extending this visibility to develop a deeper understanding of application-specific traffic is crucial for understanding a wide range of operational and performance information that can impact the users of these applications. For example, specific compression and acceleration technologies can be applied at the network layer to accelerate email applications such as Microsoft Exchange. Another example is to prevent employee access to services such as YouTube and social networking sites from impacting business applications. Understanding the application (YouTube, Instant Messaging) and enforcing policies based on the application ensures that business critical applications meet or exceed the performance expectations of end users.

Network Connectivity

Customers, partners and employees all require immediate access to applications and information. Modern applications such as supply chain applications, IP telephony, Customer Relationship Management (CRM), Enterprise Resource Planning (ERP), or sales force automation applications demand significant network performance. Concurrently, the challenge of working from any location in the enterprise further increases the complexity.

As part of the data center network design, the following critical aspects of external network connectivity must be considered:

- WAN connectivity to enable branch office and campus users to access applications
- Internet connectivity to enable partner access as well as secure remote access for remote and mobile users
- Superior speed for data center backbone connectivity, data replication, business continuity, and use of technologies such as VPLS/MPLS

The internal data center comprises one or more server network(s) or data center LANs . The data center LAN hosts a large population of servers that require high-speed and highly available network connectivity. In addition, there may be multiple LAN segments and networks deployed that differ in security and capacity levels and other services offered. Typically, connections of 1Gbps and higher (while 10Gbps are becoming the standard) will be available in the data center network, providing at least 1Gbps to the server and preferably 10 Gbps at network choke points.

Security

The critical resources in any enterprise location are typically the applications themselves and the servers and supporting systems such as storage and databases. Financial, human resources (HR), and manufacturing applications with supporting data typically represent a company's most critical assets and, if compromised, can create a potential disaster for even the most stable enterprise. The core network security layers must protect these business critical resources from unauthorized user access and attacks, including application-level attacks.

The security design must employ layers of protection from the network edge through the core to the various endpoints, for example defense in depth. A layered security solution protects critical network resources that reside on the network. If one layer fails, the next layer will stop the attack and/or limit the damages that can occur. This level of security allows IT departments to apply the appropriate level of resource protection to the various network entry points based upon their different security, performance and management requirements.

Layers of security that should be deployed at the data center include the following:

- Denial of Service (DoS) protection at the edge
- Firewall(s) to tightly control who and what gets in and out of the network
- VPN to protect internal communications
- Intrusion Prevention System (IPS) solutions to prevent a more generic set of application layer attacks

Further, application-layer firewalls and gateways also play a key role in protecting specific application traffic such as XML.

Policy and Control

Policy-based networking is a powerful concept that enables devices in the network to be efficiently managed, especially within virtualized configurations, and it can be used to provide granular network access control. The policy and control capabilities should allow organizations to centralize policy management while at the same time offering distributed enforcement. The network policy and control solution should provide appropriate levels of access control, policy creation and management, and network and service management, ensuring secure and reliable networks for all applications.
In addition, the data center network infrastructure should integrate easily into customers' existing management frameworks and third-party tools such as Tivoli, and provide best-in-class centralized management, monitoring and reporting services for network services and infrastructure.

Quality of Service (QoS)

In order to truly assure application experience over large networks, QoS is a key requirement. It is critical to make sure that QoS levels are assigned and managed to ensure satisfactory performance of the various software applications. A minimum of three levels of QoS (each of which determines a priority for applications and resources) are as follows:

- Real-time
- Business critical
- Best effort

MPLS networks and network traffic engineering capabilities are typically deployed to configure Label Switch Paths (LSPs) with RSVP or LDP. This is especially critical with voice and video deployments, as QoS can mitigate latency and jitter issues by sending traffic along preferred paths or by enabling fast reroute to anticipate performance problems or failures. The data center network design should allow the flexibility to assign multiple QoS levels based on end-to-end assessment, and allow rapid and efficient management to ensure end-to-end QoS for the enterprise.

High Performance

To effectively address performance requirements related to virtualization, server centralization and data center consolidation, the data center network should boost the performance of all application traffic, whether local or remote. The data center should offer LAN-like user experience levels for all enterprise users irrespective of their physical location. To accomplish this, the data center network should optimize applications, servers, storage and network performance.

WAN optimization techniques that include data compression, TCP and application protocol acceleration, bandwidth allocation and traffic prioritization improve performance network traffic. These techniques can also be applied to data replication, backup and restoration between data centers and remote sites, including disaster recovery sites.

Within the data center, Application Front Ends (AFEs) and load balancing solutions boost the performance of both client-server and Web-based applications, as well as speeding Web page downloads. In addition, designers must consider offloading CPU-intensive functions, such as TCP connection processing and HTTP compression, from backend applications and Web servers.

Beyond application acceleration, critical infrastructure components such as routers, switches, firewalls, remote access platforms and other security devices should be built on non-blocking modular architecture, so that they have the performance characteristics necessary to handle the higher volumes of mixed traffic types associated with centralization and consolidation. Designers also should account for remote users.

A Green and Environmentally Friendly Data Center

A green data center is a repository for the storage, management and dissemination of data in which the mechanical, lighting, electrical and computer systems provide maximum energy efficiency with minimum environmental impact. As older data center facilities are upgraded and newer data centers are built, it is important to ensure that the data center network infrastructure is highly energy and space efficient. Network designers should consider power, space and cooling requirements of all network components, and they should compare different architectures and systems so that they can ascertain the environment and cost impacts across the entire data center. In some environments, it may be more efficient to implement high-end, highly scalable systems that can replace a large number of smaller components, thereby promoting energy and space efficiency. Green initiatives that track resource usage, carbon emissions and efficient utilization of resources such as power and cooling are important factors when designing a data center. Appendix B presents an analysis of the Juniper Networks MX960 Ethernet services router's effects on reductions in energy consumption and footprint within the data center. One can use this appendix as an example for comparative analysis against other core solutions.

Juniper Networks Data Center Network Architecture

The intent of Juniper Networks approach for building the enterprise data center network is to allow enterprises to take advantage of the most advanced technologies, offer a design model that supports the current as well as future applications and data processing requirements of the enterprise, while at the same time reducing risk and total cost of ownership.

Open Systems Approach—Juniper Networks Enterprise Framework

Juniper Networks uses a simplified version of the Open Systems Interconnection (OSI) model that includes three functional layers controlled by a Policy and Management domain (Figure 3). These functional layers are as follows:

- Applications
- Services
- Infrastructure

The Applications layer provides support to the various software applications that are required to run the business. It provides the environment that allows applications to run and interoperate. The Services layer combines the traditional presentation, session and transport layers and provides support to users and applications. It includes security services, applications interfaces, and acceleration and optimization services. The Infrastructure layer combines the network, data link and physical layers and consists of routing and switching features that manage the network, connection management, data flow and QoS.

The Policy and Management domain integrates with the customer's centralized policy and management functions to help reduce operations costs while simultaneously enabling compliance. All three layers are interconnected with open standards-based interfaces that allow enterprises to seamlessly deploy a multivendor solution that provides flexibility to use the best technologies to meet business requirements.

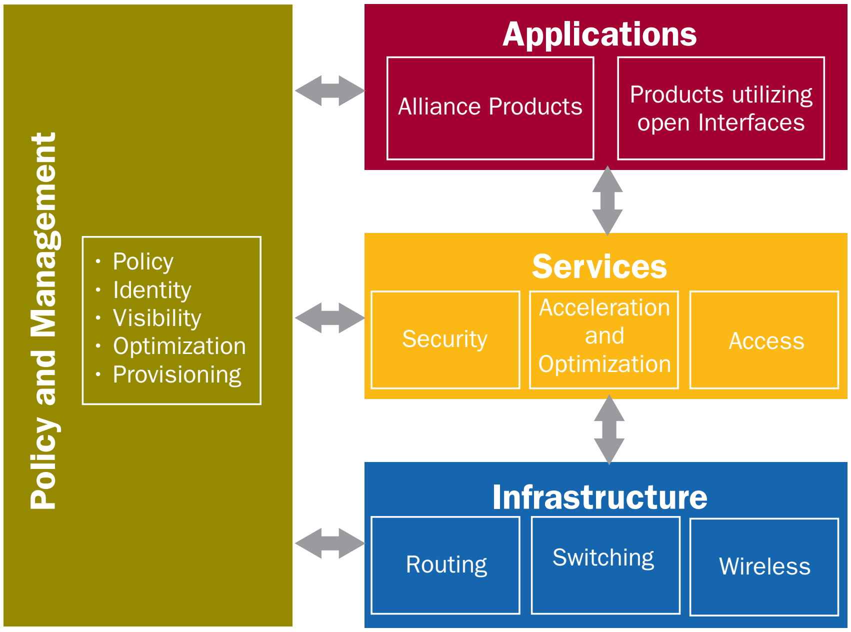

Figure 3: The Juniper Networks Enterprise Framework

The Juniper Networks Enterprise Framework supports the next-generation data center network by providing a best-in-class network environment that uses open, standards-based and industry-accepted interfaces. Enterprises can use this framework to logically view their network infrastructure and applications in order to make decisions that best serve the requirements of deploying enterprise applications.

Juniper Networks takes a holistic approach to next-generation networking and takes into account the user, network and applications perspectives. Our understanding of applications and how they are accessed from a variety of locations enable us to provide an architecture that meets the demands of a variety of users.

Location-Based Approach

In this section, we describe how the data center network interconnects and how data remains available to the enterprise regardless of physical locations and geographies from which data clients attempt to connect. Enterprises typically have campuses, regional offices, branch offices, a private WAN and data centers.

The key intention of the data center is to offload “always on” requirements from various enterprise locations to a central, stable location that always contains the enterprise’s most recent application data. By decoupling the information store from the physical location of the user, enterprises derive greater efficiencies by creating a centralized pool of resources. This trend of centralizing applications and consolidating multiple facilities makes the WAN or other external networks extremely critical, because users now need to traverse a larger network in order to gain access to data. As such, a great deal of emphasis has been given to the design of the enterprise Private WAN and Internet edge, which in many cases hosts branch office connectivity and remote user connections.

The data center does not typically host users and most certainly does not accommodate data center application users. However, this data center network design model can support different operational requirements that are unique to certain enterprises. Options such as administrative user access can be built into any data center design.

WAN services should extend to all of the remote location connections. Among these services are stateful firewalls, intrusion prevention and WAN acceleration. Figure 4 depicts a high level perspective, illustrating the overall enterprise connectivity into the data center and connectivity between data centers.

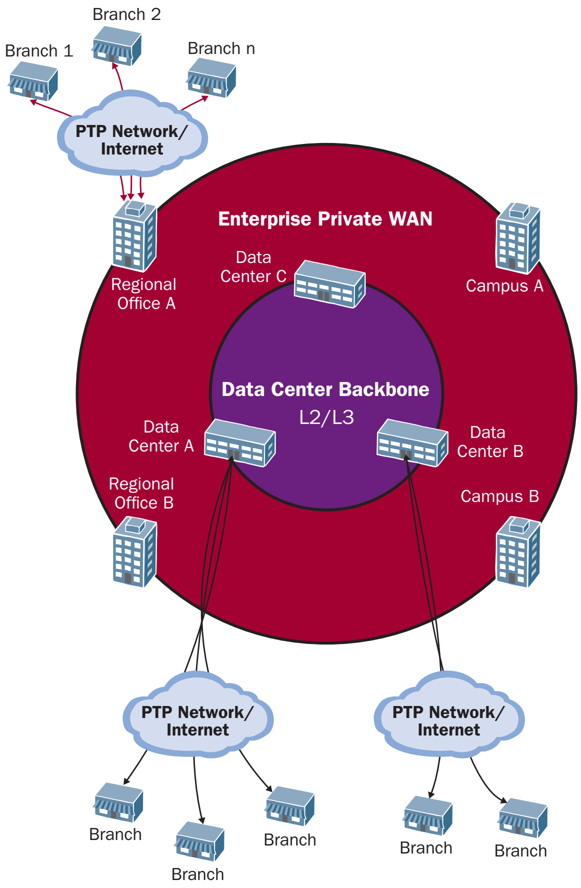

Figure 4: Enterprise Network Connectivity to the Data Centers

Design Principles

Key design principles originate from business and technical reasons. The business reasons are fairly clear—optimize capital expenditures and reduce operation expenses. The top level technical requirements include the following:

- Leverage shared infrastructures
- Employ virtualization technologies to increase utilization and efficiencies
- Ensure scalability, flexibility, security and application performance over the network

Juniper Networks key design principles are as follows:

Consolidation of Data Centers and Centralization of Services from Multiple Business Offices—This principle imposes a variety of technical requirements on the data center network. Centralizing services typically does not improve overall processing time nor data availability, but it often increases overall utilization and allows for more streamlined IT operations. Additionally, centralizing services requires maintenance of the unique aspects of legacy distributed processing configurations such that different processing instances may belong to different business entities, such as finance and HR. Uniqueness and operational freedom should remain “virtually” independent.

Virtualization—The virtualization of processing has introduced a new standard in resource pooling and resource utility optimization. Virtualization technologies at various levels are introduced in the data center from virtualization of large storage arrays and servers to network virtualization and network service. The network infrastructure manifests virtualization through VPNs, labels and tags of forwarding plane traffic, while the network services manifest virtualization through the definition of service instances and application of unique processing logic to the different instances. The overall data virtualization capabilities of the data center are key requirements that effectively drive network virtualization.

High Availability (HA)—Consolidating and centralizing resources, as well as virtualizing technologies, makes guaranteeing data access all the more critical. Data should be available regardless of the location from which it is being served. The four key vectors that address network HA include the following:

- Component
- Device
- Link
- Site

Streamlined Operation and Management of Data Center Services—In a consolidated and virtualized environment, one of the key elements is a single management platform based on open standards that knows how to control servers, applications, storage and network infrastructure as one. Hence, it is critical to use devices and systems that support open standards-based interfaces and protocols, so that these devices and systems can all be controlled from existing and from evolving customer management systems.

High-Level Architecture

Figure 5 illustrates the Juniper Networks data center network architecture; the following lists the major architectural tiers:

- Edge Services Tier—hosts all WAN services connecting to “non-data center” locations
- Core Network Tier—connects all data center networks within and across data centers
- Network Services Tier—supports WAN acceleration, Intrusion Prevention and other network services
- Applications and Data Services—provides network connectivity to the data center server and application infrastructure
- Data Center Backbone—provides connectivity between data center facilities for HA, replication and disaster recovery

In the paragraphs that follow, we explore the different network tiers in greater detail.

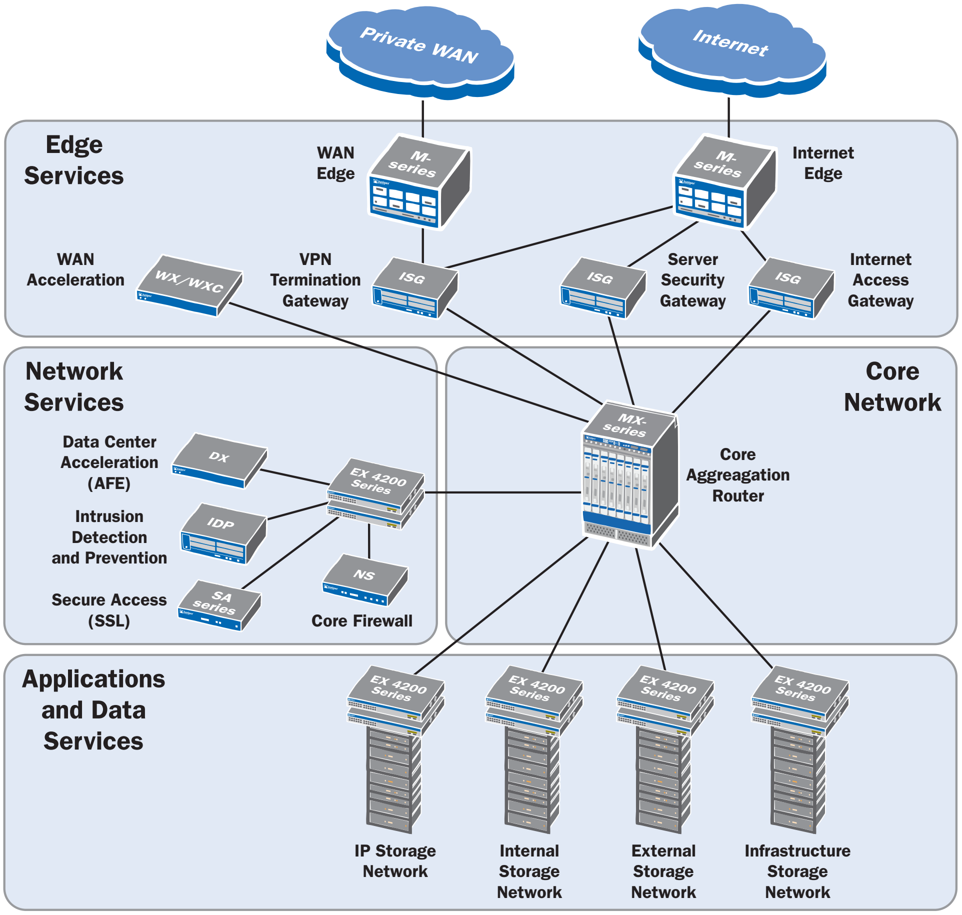

Figure 5: Juniper Networks Data Center Network Architecture

The data center network scalability requirements are rather significant because they must support centralized applications and data center consolidation. Hosting a large network in one location requires some modularization that allows certain services to be re-applied to applications and areas as needed.

As the data center edge network serves as the key boundary to the data center, it is responsible for maintaining reachability with all other external networks.

It is important to support network applications with an extremely fast core network that is capable of forwarding the total aggregate traffic at line rate. Effectively, the network core can span across multiple locations and multiple devices. Logically, the network core connects all data center networks directly to itself. This attribute is critical in order to offer rack/location-agnostic server-to-network binding, which is a key element in building a virtualized data center fabric that supports automatic repurposing of compute resources. Another benefit of this approach is that it maintains a more controllable HA design, so that a single device includes its own redundancy component to augment an additional device (or set of devices) as a backup system.

Extending all networks to the data center core allows flexibility to enable or disable services to each of the networks independently, in addition to allowing scalable services initiated from demand and available capacity. A virtualized approach for enabling network services optimizes performance and efficiency. A common example is a stateful firewall that provides virtual domain security by directly connecting to the core and by securing multiple physical networks. This approach proves highly useful in segmenting the network by firewall policy.

Edge Services Tier

The Edge Services tier is responsible for all connectivity and network level security aspects (up to Layer 4) to connect the data center to the outside world. Typically, routers and firewall/VPNs are located in this tier. It is likely that the data center connects to various leased lines connecting to partners, branch offices and to the Internet. For connecting all of these networks, it is important to plan for the following:

- Internet routing isolation, for example separating the exterior routing protocols from the interior routing protocols
- Network Address Translation (NAT) to convert your private IP addresses to public Internet routable IP addresses
- IPSec VPN tunnel termination for partner, branch and employee connections
- Border security to enforce stateful firewall policies and content inspection
- Quality of service (QoS)

Network architects have often used Layer 2 switches at the edge to form a hierarchical mesh, with the intention of allowing multitude links to provide fault protection during failure. The Juniper Networks solution employs Juniper Networks M-series multiservice edge routers and Integrated Security Gateway (ISG) firewalls. Juniper Networks leverages the routing functionality of the ISGs to provide a routed connectivity solution instead of a traditional switched mesh. This method places failure detection and correction into a domain that is solely routed, providing more effective and intelligent network resource use. The direct protocol interaction between the routers (without intervening switches) eliminates the typical layer of Ethernet switches that are commonly used at the edge.

Edge Services Connectivity

Figure 6 shows the Juniper Networks Edge Services design and illustrates how the Edge Services tier connects multiple, external networks to the data center. Edge Services provide all connectivity and network level (up to Layer 4) security aspects for connecting the data center to the outside world. The edge routers and firewall VPNs reside in this tier.

The edge routers are Juniper Networks M-series routers and are the edge devices for both Internet and private WANs. The M-series routers were selected for two primary reasons: interface capacity and throughput.

Each router has a single connection to the Internet (or private WAN). Connectivity between the edge routers to each ISG firewall creates a fully meshed network. You can link the edge routers to each other using a single gigabit Ethernet link that provides a transit path around a less preferred or failed path. In addition, Juniper Networks uses redundant hardware, Dynamic Routing Protocols (DRPs) and fully meshed links to minimize the amount of failure cases that could impede business continuity.

Edge Services High Availability

The Edge Services tier should provide HA at three levels where appropriate:

- Link
- Device
- Component.

Link-level HA should be applied at all Internet connections and in cases where additional data centers are available, it is best to keep a single leased line/private WAN connection in each data center. Device-level HA is relevant only when we enable the link-level HA setting, as multiple devices cannot utilize a single link themselves. Hence, Internet facing routers and devices located behind these should support device-level HA. Additionally, component level HAs (multiple power supplies, fans, route engines) should be mandatory for edge-deployed devices.

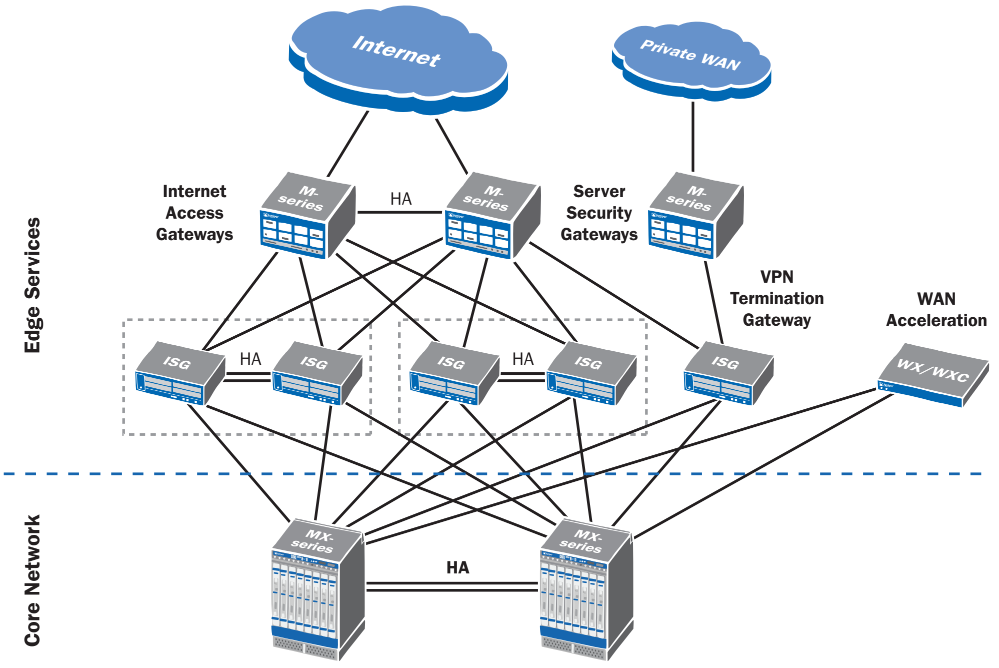

Figure 6: Data Center Network Edge Services

In this solution, dynamic routing determines the flow of traffic. Each tier is deployed as a fully meshed solution. As a result, redundant paths are provided on each redundant device. During a link failure, a single device is not lost and this increases environment uptime by avoiding bringing down a viable path, unless necessary.

During a failure, the network requires a minimum of one additional redundant path to route around the failure. While this design itself offers HA, the addition of a second data center further ensures HA, as an entire data center could be lost without losing network operability.

Edge Services Performance

As in any other major server concentration, the data center should terminate a large number of WAN acceleration tunnels. These tunnels correspond to as many remote sites as may be appropriate for optimal user experience and performance. Some of the WAN acceleration technologies include redundant WAN acceleration tunnels and load balanced WAN acceleration clusters. Both technologies integrate by using intelligent traffic rerouting techniques in the data center.

Edge Services Security

The Edge Services network serves three major security functions. First, it protects against Denial of Service (DoS) attacks that are most efficiently controlled at the data center edge without using other valuable processing resources. Second, the edge tier firewalls can perform stateful inspection. Third, we implement VPN secure connectivity services. This section covers the design guidelines for these three security functions.

For large data centers, Juniper Networks recommends using three sets of firewalls in the Edge Services tier. The first set, the Internet firewalls, must connect to the Internet and receive routing information from the edge routers to enable outbound traffic routing to the Internet. The second set, the Secure Services Gateways (SSGs), secure the server and data resources and software applications for inbound traffic originating from the Internet. The third set, the IPSec VPN firewalls, comprise the connectivity hub for all remote sites and terminate IPSec VPNs from the Internet as well as from the private WAN. The IPSec firewalls also terminate VPN tunnels for all of the remote branches over the private WAN. To provide services to the remote branches, the IPSec VPN firewalls must connect to the network core. Although these firewalls are shown as three sets, for smaller capacities and performance requirements, it is possible to consolidate the three firewalls into one or two sets.

General DoS protection to all data center services should be performed at the Edge Services tier. This moves the security intelligence closer to the provider edge, thereby decreasing the number of devices that can potentially be compromised, especially with DoS attacks. A large flood can present challenges to any network, as it can consume all available network bandwidth and may require extra processing by stateful firewalls. Large floods result in high CPU usage and slow response times.

While stateful firewalls provide much needed visibility and fine-grade protection against a variety of floods, all stateful firewalls have an upper limit in their capacity to deal with certain types of floods such as SYN or Internet Control Message Protocol (ICMP). If a firewall is overwhelmed with a flood, the firewall will experience high-CPU load and may drop legitimate traffic. The specific rate per attack varies per firewall depending upon its configuration and software version. To protect the firewall and network against massive floods, rate limits should be implemented on routers protecting all firewall interfaces. The goal is to limit certain types of traffic, such as TCP control traffic and ICMP types, to rates which will not impact available bandwidth and overwhelm the firewall.

As part of the VPN design and encryption protocols selection, there are trade-offs that must be made. Organizations should choose the strongest encryption that does not compromise the performance requirements for the network. Encryption algorithms should be based on a balancing act between security and performance. A longer key length provides more security against brute force attacks yet may require more computational power. Therefore, this approach provides less performance for encrypting large amounts of data. Note that performance considerations should be made for all devices participating in the VPN, not only devices that terminate at the headend. Satellite devices may not be as powerful as the application-specific integrated circuit (ASIC)-accelerated, crypto powered headend systems. When analyzing the elements, it is important to acknowledge the handshake protocol encryption requirements. These typically use asymmetric encryption algorithms for improved security and may affect the devices dramatically, especially with many VPN peers.

One also must consider bulk encryption algorithms. Typically, they must be symmetrical and least influenced by design due to hardware assistance and the lower cost of hand shakes. However, if the design presents few VPN peers and large amounts of data transfer, this element should be considered; the lowest common denominator will be the speed that determines the VPN capacity. Finally, one should consider hashing algorithms. This selection is primarily done based on security requirements, but if hardware assistance is involved, then design considerations diminish.

Core Network Tier

The Juniper Networks design employs a data center network architecture consisting of two logical forwarding tiers rather than a traditional

3-tier model. Traditional 3-tier networks add an aggregation network between access networks and core networks, and are the primary method to extend networks because of scalability limitations with most available core network devices. Aggregation at the core allows for more flexibility and easier support for virtualization, but it requires high-speed processing and High Availability levels. A 2-tier network is one core network with all of the access networks connecting directly to it.

One of the biggest advantages of a 2-tier design is a dramatic reduction in the number of devices. Reducing the number of devices provides the following advantages:

- Produces significant power savings
- Reduces the facilities footprint of the system
- Offers simplified device management
- Allows tighter security control
- Reduces the number of system failure points

The scalability of the 2-tier model is typically limited by the scalability of the core network devices. The more traditional 3-tier design, which allows for high scalability requirements, is not discussed in this paper.

Core Network Connectivity

The core network provides the key data center fabric connectivity by connecting routers, servers, appliances and storage devices. It does not directly allow connections between the different networks that connect to the core, as each network must be contained in a separate routing instance of Virtual Routing and Forwarding (VRF). In cases where traffic should traverse between the VRFs, the core firewall performs the forwarding according to the security policy. Effectively, the core firewalls should connect between the different networks that reside on the same data center (see Network Services Tier).

Core Network HA

The core network is a key component in enabling HA in the data center network. By connecting all networks to the core network with full redundancy at the core, HA is achieved without added complexity and dependency on the network protocols and convergence. Traditionally, adding HA requires redesign of the network, whereas by using standards-based redundancy protocols and a core network approach, HA is provided at easier operational overhead. In addition to adding redundant devices, it is extremely important to ensure that the core data center devices support in-service operations such as hot-swap interfaces and software upgrades.

Core Network Virtualization

To achieve network virtualization from the server through the network core, there are a variety of options to consider. For all options, a key assumption is the deployment of network systems that deliver line-rate throughput even when all features are turned on. In general, there are two possible approaches:

1. Extend VLANs from the access layer or server all the way to the network core.
2. Use VLANS between servers and access devices, and divide the network by using MPLS from that point on through the core.

There are advantages and disadvantages to each approach, and each approach may be more practical on a different scale and for data center operators with different skill sets. VLANs that extend all the way to the core are more appropriate for smaller networks. However, in the case of larger networks that require VLAN scaling limitations with more elaborate QoS requirements, MPLS is the preferred choice. Regardless of the approach, the Juniper Networks data center network architecture and solution components provide both approaches without sacrificing performance.

Multiple instances of a single VLAN, residing in different physical access networks, can be joined at the core network across line cards (or not) without impacting performance. Additionally, multiple distinct VLANs, all connecting to a single access switch port, can be seamlessly reclassified and associated with MPLS Label Switched Paths (LSP s) with unique QoS and connectivity characteristics. The Juniper Networks 2-tier architecture provides for a more flexible design option (see Figure 7).The MX-series Ethernet services routers (ESRs) reside in the Core Network and Juniper Networks EX-series Ethernet switch platforms reside in the access layer.

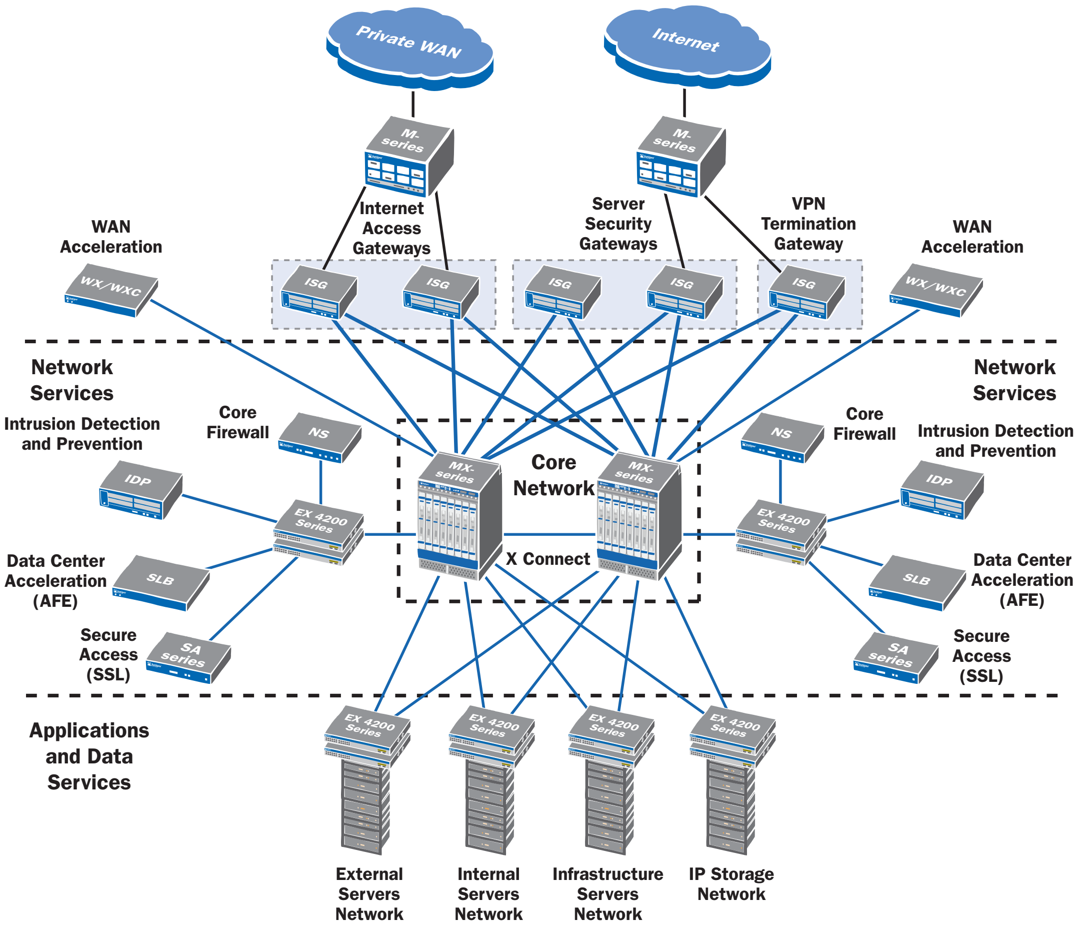

Figure 7: Data Center Core Network and Network Services

Network Services Tier

The network services tie closely to the network protocols that support data center applications. Network services are generally divided into two main categories:

- Security services
- Application services

Throughout this section, we will describe both families in greater detail and address the key elements that comprise the data center network architecture.

The Network Services tier should extend itself to any of the server networks hosted in the data center, and apply a network-specific policy and set of configurations to behave appropriately with the traffic in that particular network section. For example, using a security service such as traffic SYN checking/sequence number checking may only be required for servers available to the outside world. Therefore, the architecture should support the application of these features only to those systems or networks. Most importantly, key characteristics are enabled by direct logical attachment to the network core of the data center.

Leveraged throughout is the Network Services tier's ability to extend a shared pool of network services to any of the server and data networks, while allowing for granular and specific network service settings for each one of the services. The network services are virtually available for each of the backend data and service networks while sharing the network service resources across the entire data center. This approach allows the network designer to intelligently deploy network services to different applications and networks in the data center. Virtual instances are a key consideration in designing the Network Services tier.

Figure 8 illustrates the connectivity systems (MX-series ESR), application systems (Network N), and the network service systems (depicted on the far right of the diagram). This diagram provides a core network perspective and shows the interaction between the core and the pooled service devices.

The network services include the following systems:

- Security services—such as firewalls and Intrusion Prevention
- Application frontend services—such as server load balancing, SSL offload, HTTP cache, TCP multiplex, and global server load balancing (GSLB)

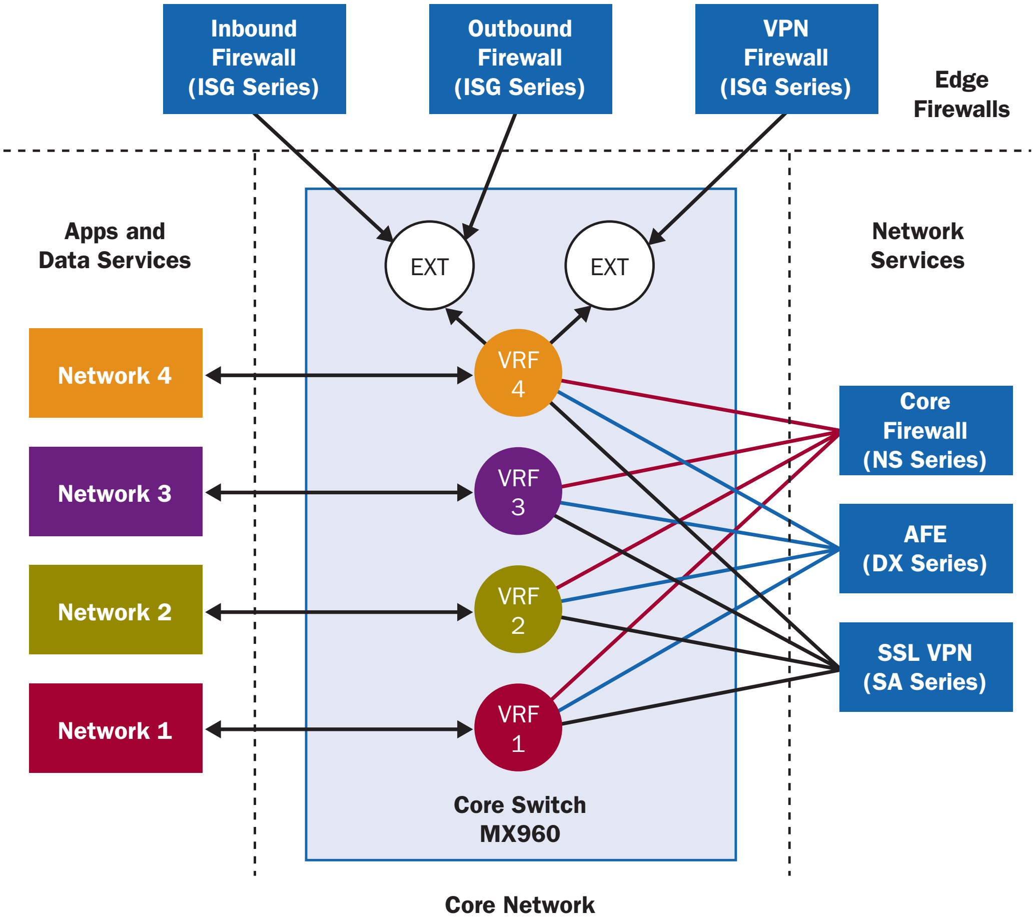

Figure 8: Connectivity Systems, Application Systems and Network Service Systems

Data Center Security Services

One of the most important services of the Network Services tier is the security service. The security service essentially controls segmentation of the data center into separate networks, and it enables secure connectivity between the different networks. Because security services are broadly used, multiple devices participate in the application of security services to the data center server network.

Stateful firewalls are the cornerstone of the data center network's security service. Stateful firewalls enforce a security policy that aligns with business and operational requirements through the identification and classification of networks. In addition to being the primary L4 access control system, the firewalls help with many security functions in the data center, such as service DoS or quota protections, Deep Inspection to specific applications where it makes sense, and also potential network address translation.

Generally, the first layer of defense inside the data center is the stateful firewall. However, it is important to recognize that the firewall must be capable of extending a logical subset of its functionality as you dedicate it to a specific data center network. The minimum amount of resources the firewall must dedicate will be a separate control and forwarding engine (virtual router) such that all traffic streams are totally isolated, and forwarding decisions will not mistakenly puncture the security protections. An additional attribute in designing a consolidated data center services instance is HA capabilities that must extend themselves particularly at the services layer, in order to truly design a network that depends on the services for its core functionality.

The Juniper Networks NetScreen firewall systems can split into separate virtual domains of control and forwarding instances or Virtual Systems (VSYS), creating separate virtual domains that allow autonomy to different departments so that they can control their security policies. In order to connect all of the core networks, the core firewall must participate in routing protocols within the data center network.

Application Security—In addition to assuring secure connectivity at Layer 4, the Network Services tier should employ application security services such as Intrusion Prevention to protect the data center infrastructure. Because these application services are available to all users coming from insecure locations, the risk of application misuse or application DoS increases. In addition, because multiple applications are co-located, this creates a chain effect in which each application is affected by the risk to which another is exposed.

The platforms should support the level of performance required by the data center and be able to inspect L7 information at line-rate speeds. It is necessary to understand that the protocols deconstruct the data streams and build the right context to look for application threats. Therefore, a powerful and rich application protocol decoder is necessary. Also, the integration of the application protocol decoding to firewalls is a key consideration to help reduce the number of devices and to increase overall effectiveness. Finally, virtualization or context-based security policy application, in which the security systems are able to uniquely treat different networks and applications, is another important consideration.

Application Front Ending Services

It is important to find ways to scale the data center services without a linear increase in the hardware footprint and to ensure that the design does not increase the operational complexity. A key component of the Network Services tier is a solution that enables offloading of non-specialized services from the data center servers.

One should consider deploying a system that supports acceleration for the different application tiers and also provides comprehensive capabilities around the more common and emerging application areas like Web 2.0. A data center acceleration solution should boost the performance of both client-server, Web-based, and server-to server-applications, as well as speeding Web page 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. For its part, the application acceleration platform should be seamlessly expandable through stacking or clustering of multiple devices. In addition to advanced traffic management and acceleration, the application front ending service should serve as a standard load balancer. This means forwarding traffic to its destination address from a pool of available addresses.

Organization/business requirements drive the need to allow different applications to be treated differently, and to allow different departments to control and define what acceleration and front ending characteristics they require from the network service. Our solution addresses these requirements.

Applications and Data Services Tier

The Core Network tier connects to the Applications and Data Services tier which hosts all of the servers, databases and storage. Generally, there are four types of networks and there can be multiple instances of each type. Primary reasons for the multiple instances are separation of duties within the company, and differentiated business objectives and IT requirements for the different networks. Figure 9 illustrates the four networks. Description of the four networks is as follows:

- External Applications Network—can be multiple external networks serving separate network segments. These typically include applications such as the public Web site, public mail transfer agent (MTA), Domain Name System (DNS) services, remote access and potential file services that are available through unfiltered access.
- Internal Applications Network—multiple internal networks serving different levels of internal access from within the campus of branch locations. These networks typically connect internal applications such as finance and human resource (HR) systems. Partner applications also reside in the internal network and any specific applications that are exposed to partners such as inventory systems and manufacturing information.
- Infrastructure Services Network—only servers that are accessible to users are allowed to access infrastructure networks. These are intended to operate only on an automatic basis and performance usually is quite predictable. Common examples of infrastructure services include Lightweight Directory Access Protocol (LDAP), databases, file shares, content management and middleware servers.
- Storage—storage networks, such as Fibre Channel, InfiniBand or Internet Small Computer System Interface (iSCSI) are part of the storage networks. Critical application servers directly connect to the storage devices through a separate Host Bus Adapter (HBA) to ensure fast access to data. Other servers connect using Ethernet to access storage facilities.

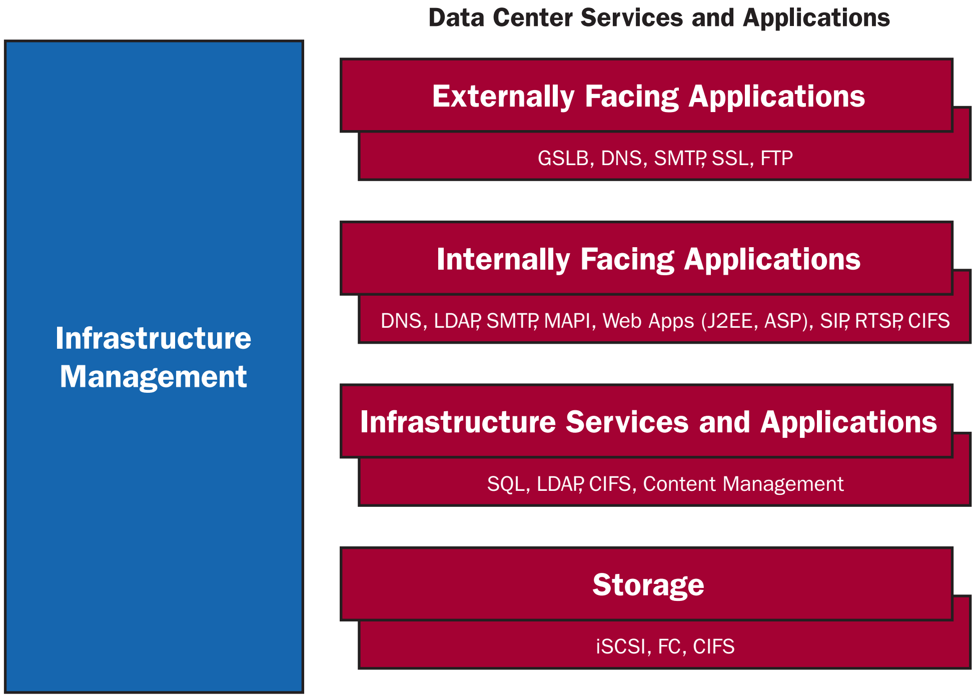
Data Center Services and Applications
Infrastructure Management
Externally Facing Applications
GSLB, DNS, SMTP, SSL, FTP
Internally Facing Applications
DNS, LDAP, SMTP, MAPI, Web Apps (J2EE, ASP), SIP, RTSP, CIFS
Infrastructure Services and Applications
SQL, LDAP, CIFS, Content Management
Storage
iSCSI, FC, CIFS

Figure 9: Data Center Application Network Types/Purposes

The Applications and Data Services tier is primarily responsible for connecting and wiring all servers. Essentially, this tier requires two high-speed, independent, top-of-rack switches that connect to the core network as a solution. In data center environments, servers are interconnected to access switches deployed within server racks. These access switches are often referred to as "top-of-rack" switches due to their location within the data center. Top-of-rack switching provides increased levels of availability because of multiple independent operating characteristics and physical power sources. Servers connect to two different physical switches, each part of a separate virtual chassis ring. Each ring in turn connects to the core network while using a loop detection and HA L2 protocol.

Data center application connection is as follows:

- Each server has two, 1 Gbps access network switches; each server connects to a separate access switch for redundancy purposes.
- The access switching layer connects to the core network using 10 Gbps uplink; each access switch has separate 10 Gbps links.
- The server connection links and access switch uplinks can use VLAN trunking technology to support both server virtual location and aggregation; all aggregating multiple Layer 2 networks then use fewer connections.

Each internal and external applications network can be segmented into several sub networks (see Figure 10). The servers that host these applications connect with at least a 1 Gbps (currently moving towards 10 Gbps) link to the Juniper Networks EX-series Virtual Chassis™ switch. The EX-series Virtual Chassis switch connects to the Network Core via a 10 Gbps connection. Depending on the number of servers, multiple EX-series Virtual Chassis may be required, as shown in Figure 10. Juniper Networks recommends dual homing the access layer switches using L3 with OSPF equal-cost multipath (ECMP) instead of the Spanning Tree Protocol for deterministic behavior for minimal packet loss.

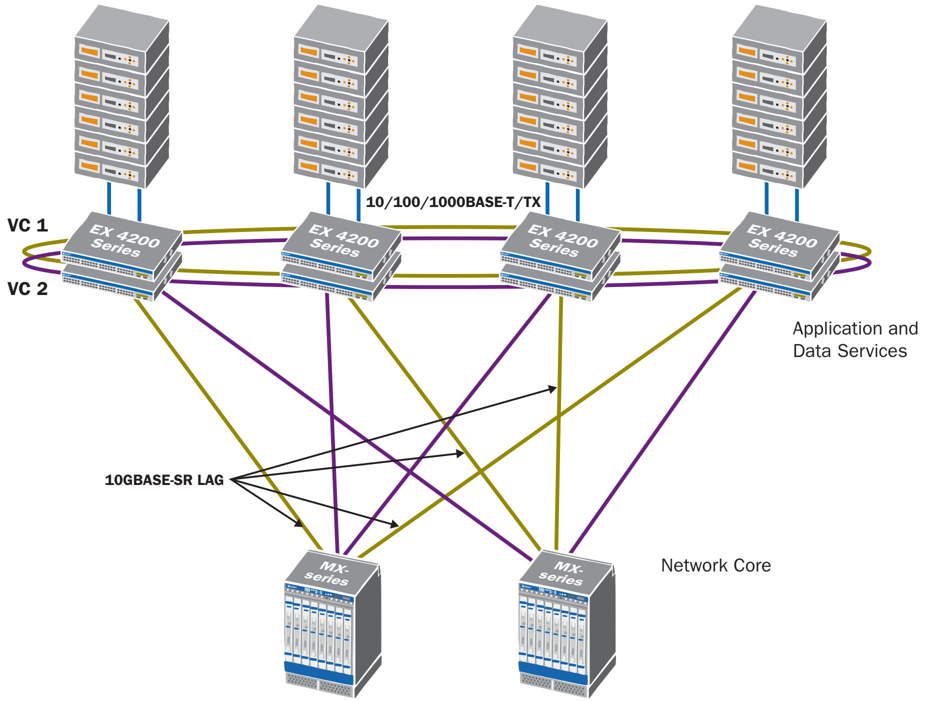

Figure 10: Application and Data Services Network View

In data center environments, servers are interconnected to access switches deployed within server racks. Typically, top-of-rack access switches are deployed in pairs to redundantly support servers within a single rack. Juniper Networks EX 4200 Virtual Chassis offers several advantages when deployed as a top-of-rack access switch. The EX 4200 supports a maximum of 48 10/100/1000BASE-T/TX interfaces for attached server devices at 1 Gbps wire-rate per interface. As a result, performance is not compromised. Also, each EX 4200 offers additional wire-rate uplink interfaces, with a maximum of four GbE or two 10 GbE uplink modules for interconnecting from the top rack back to the data center core.

The EX 4200 also supports the virtual chassis concept, whereby a maximum of ten EX 4200 switches can be interconnected through a redundant, high-speed 128 Gbps interconnect, yet still be managed and maintained as a single logical device. With the Virtual Chassis technology , the number of managed devices can be reduced by a factor of ten, significantly simplifying operations and reducing costs associated with maintaining large numbers of legacy access switches. Additionally, uplinks can be distributed across multiple EX 4200s in a single Virtual Chassis, providing uplink performance flexibility and added redundancy levels unfounded in legacy access switches.

Storage Area Networks (SANs)

A Storage Area Network (SAN) connects servers and storage devices across a packet-switched network. SANs allow arbitrary block level access from servers to storage devices and storage devices to other storage devices. Multiple servers can therefore share storage for clustering and HA applications. In addition, the storage devices themselves can implement data protection services, such as synchronous data replication, asynchronous data replication or data snapshots by directly moving data to another storage device. SANs also provide a set of configuration, directory, discovery and notification services to attached devices.

A data center typically contains multiple SANs, each serving a different application, set of applications, work group or department. Depending upon the specific requirements, these SANs can be either FC (Fibre Channel) or iSCSI-based deployments. Both Fibre Channel Protocol (FCP) and iSCSI allow block access to storage devices using SCSI commands. FCP uses the Fibre Channel communication structure of exchanges, sequences and frames. The iSCSI protocol uses TCP/IP with an overlay of iSCSI Protocol Data Units (PDUs) to implement SCSI commands and data framing.

Fibre Channel SANs

A Fibre Channel fabric has link-level credit-based flow control making it essentially lossless without equipment failure. Link speeds are 1/2/4 Gb with 8 Gb on the horizon. FC host bus adaptors (HBAs) are FC protocol offload engines that handle most of the exchange management and all of the frame transmission or other low level protocol work. Frame forwarding is based on an equal cost multipath link state protocol—Fabric Shortest Path First (FSPF). Switch implementation does not reorder frames unless a failure occurs. The set of FC fabric services are distributed throughout the switches in the fabric.

iSCSI SANs

An iSCSI SAN can be based upon any network supporting the IP protocols. In practice, this means iSCSI SANs are built from Ethernet switches. Because iSCSI is based upon TCP/IP, it can in principle run on any switching infrastructure. In practice, depending upon the features of the Ethernet switches, the performance characteristics of TCP/IP in the face of dropped frames can limit iSCSI deployments to low performance SANs. In addition, most iSCSI deployments presently only use 1 Gb Ethernet with software drivers, and the resulting performance does not compare favorably to FC at 2 or 4 Gb with an offload HBA. However, iSCSI SANs can be considerably less expensive than FC SANs. The Internet Storage Name Service (iSNS) server provides all fabric services in an iSCSI SAN.

Where iSCSI-based SANs are desirable, Juniper Networks switches and core routers are excellent platforms for creating the underlying network, because they support symmetric flow control using 802.3X pause frames, RED (random early detection), QoS and logical partitioning. Discards due to RED only occur in congested environments, and most SANs are designed to avoid all but transient congestion. QoS allows traffic priority to be set so that storage traffic can have improved throughput and delivery characteristics during congestion. Logical partitioning allows the networking equipment that implements the SANs to be tailored to fit the needs of the specific data center and its applications.

SANs are often linked to remote data centers so that data can be replicated as part of a BC/DR (Business Continuity/Disaster Recovery) design. The inter-data center connections can run across direct optical repeater circuits such as dense wavelength-division multiplexing (DWDM), private IP-based WAN connections or the Internet.

FC traffic uses DWDM for metro-to-regional distances and uses specialized FCIP tunnel gateways for regional to longer distances. Using DWDM requires FC switches with FC credits sufficient to span the distance at the desired throughput. Fibre Channel over IP (FCIP) gateways create complete WAN acceleration services such as compression, large buffering, security, encapsulation and tunneling for FC traffic.

The iSCSI traffic can directly traverse the WAN connection without requiring a gateway, but iSCSI implementations do not generally provide sufficient buffering to fully utilize high-speed connections. The iSCSI implementations do not contain compression or other WAN optimization features. Therefore, iSCSI WAN traffic can often benefit from a WAN acceleration device such as the WX application acceleration platforms. The iSCSI traffic also can benefit from a data security gateway providing IPSec and VPN tunnels.

Data Center Backbone

In considering the scale of processing performed at data centers and the requirements for regulatory compliance, the data center backbone is a key component in the architecture and design, primarily for disaster recovery reasons. As such, the data center backbone supports a variety of computational services like data mirroring to ensure that accurate data is represented at multiple data centers. All of the functions that rely on a high-performance data center backbone include the following:

- Data replication that supports application clustering and compliance
- Data backup and restore services
- The reach to a variety of location-specific services using fast and secure connectivity across data centers to support service-oriented architecture applications
- Legacy clustering technology support that requires L2 connectivity (Figure 11).

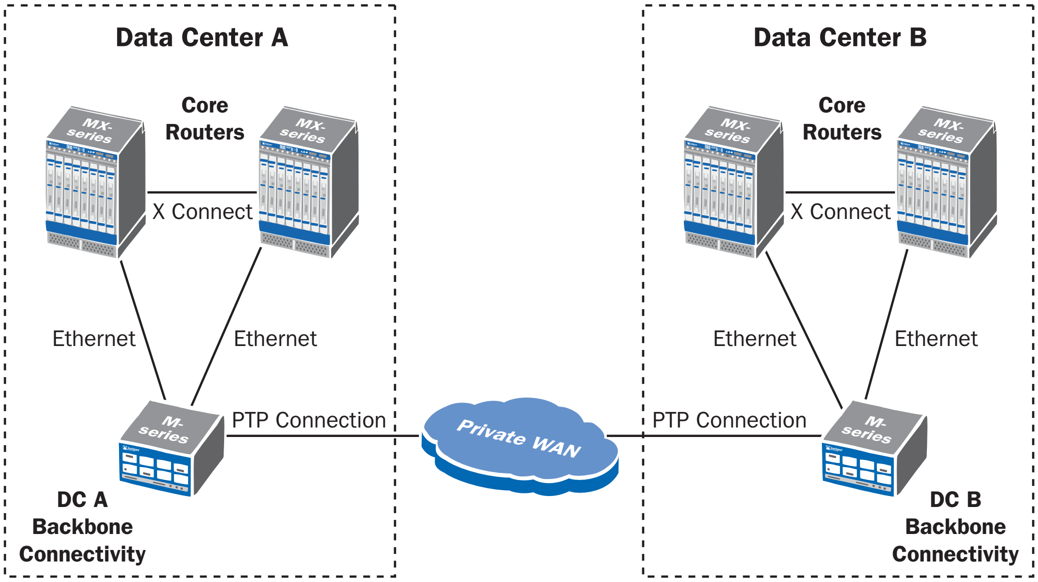

Figure 11: Data Center Backbone Connectivity

Different from any other WAN interconnect, the data center backbone must offer high-speed connectivity, especially because of the real-time and consistent representations of data that needs to be available across multiple data centers at any given time. In many cases, high volumes of data are exchanged between the data centers to accommodate mirroring, replication and backup operations.

The Juniper Networks data center network architecture defines a few key technical elements as enablers for the data center backbone which network architects can leverage to satisfy the requirements of their enterprise. These elements include the connectivity links encompassing high-speed data connectivity between data centers (most likely using fiber-based high-speed transport). The second element is the interconnect protocol, on top of the interconnecting link, that supports the separation of traffic types for QoS and security reasons. These services can be obtained through MPLS technologies most effectively. In addition to the interconnect elements, there are the availability protocols that extend the data presence beyond a single physical location. To achieve this, an

L3-based reachability or routing protocol peering exchange with Internet routers can be used to create resilient connectivity at the IP level. Additionally, employing an L7 DNS as a global load server balancing mechanism provides resilience at the service level.

Interconnectivity between data centers can be implemented using MPLS or VPLS as routing and forwarding technologies. This allows distinct IP routing information to be shared across data centers, and forwarding can be performed based on the unique, per-domain logic exchanged across the data center facilities. MPLS technologies allow for the exchange of the forwarding and routing information base to achieve consistent forwarding across all networks that interconnect using MPLS. In addition to MPLS, L2 extensions and technologies can be used so that non-IP or broadcast domain dependent/attached protocols can be connected as part of a single network. For such applications, pseudowires, data link switching (DLSw) and VPLS technologies should be used with the MPLS implementation.

Ensuring that the service is globally available and is enabled by the Network Services tier is a task that extends beyond the network forwarding layer. The key premise is that applications and users connect and associate themselves to name conventions that are other than IP (HTTP, SIP, CIFS, FTP and so on), typically through the Domain Name System (DNS). To present available services and data regardless of data center location and device availability, a GSLB technology should be applied so that queries regarding an IP-resident service will always have an answer and that service will always remain available.

Border Gateway Protocol (BGP) multi-homing is also important for the following two reasons.

The first reason is for the transitory phase in which end-service clients still maintain DNS information obtained from the GSLB service that does not represent changes to the network (potentially 24 hours, depending on time to live or TTL). The second reason is for cases where certain services are tied to a specific data center and Internet or where WAN connectivity is lost. The latter is the more common and is an important use case. Obviously, it is paramount that we assume that the data center backbone connectivity layer exists in order to support service availability when data center connectivity is lost.

To summarize, the four key elements that construct the data center backbone are as follows:

- Optical transport
- Network virtualization technology that interconnects the data centers
- IP-level availability/resilience scheme
- GLSB

All four elements support the services associated with backbone connectivity and utilization.

Data Center Network Management

IT managers want to streamline operations, deliver better service to end users and ensure compliance. Customers are increasingly adopting best practices as recommended by Information Technology Infrastructure Library (ITIL) and are increasingly investing in automation technologies that make it easier to rapidly deploy new services.

Critical requirements for data center network management systems are as follows:

- The network devices should smoothly integrate into the customer's management framework with minimal or no retraining of network/security operations center (NOC/SOC) staff. The network designers should be able to easily provision, configure, monitor and troubleshoot the network infrastructure.
- All network devices should support centralized policy management and distributed policy enforcement.
- Device management systems should leverage open standards, such as Trusted Network Computing (TNC) and the Internet Engineering Task Force (IETF), to ensure smooth interoperability with existing and future enterprise management systems.

Figure 12 illustrates a network management framework built on Juniper Networks products.

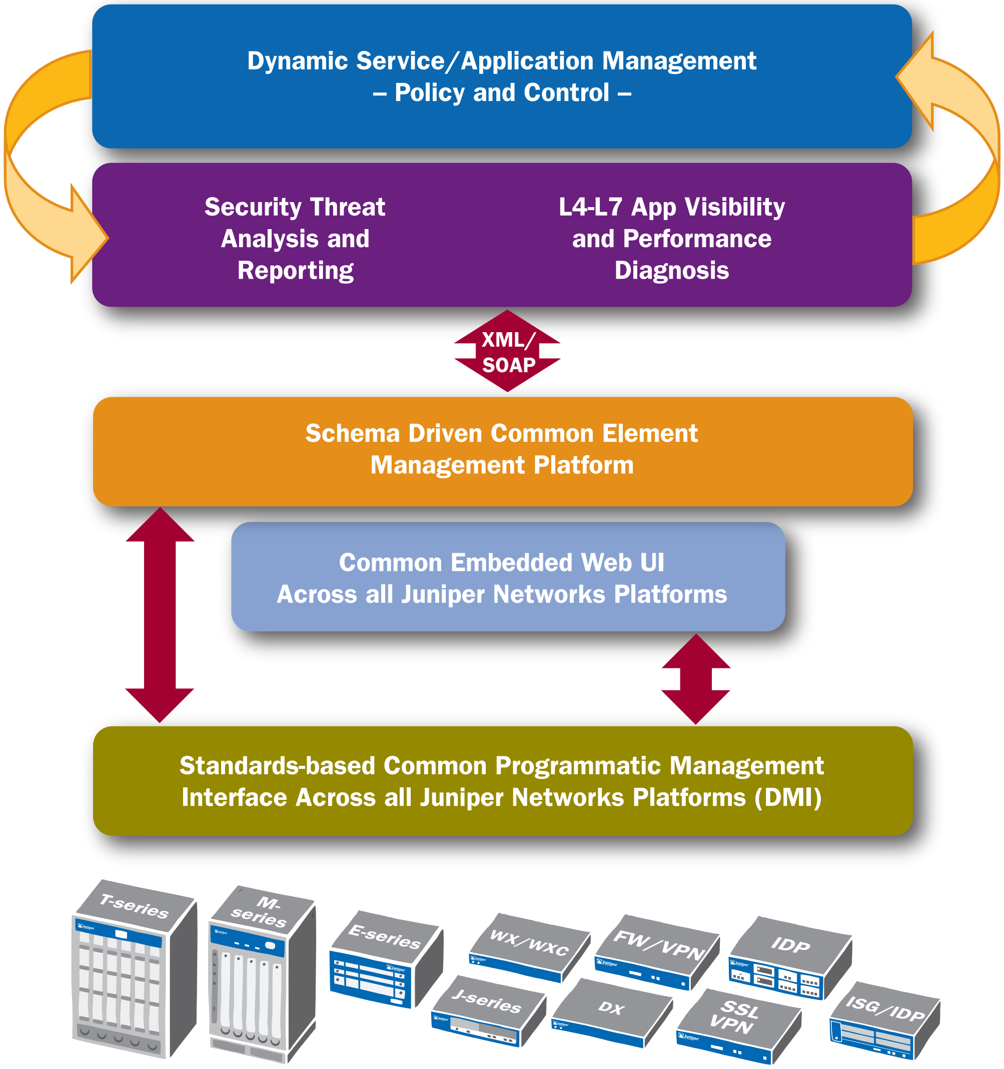
Dynamic Service/Application Management
– Policy and Control –
Security Threat Analysis and Reporting
L4-L7 App Visibility and Performance Diagnosis
XML/ SOAP
Schema Driven Common Element Management Platform
Common Embedded Web UI Across all Juniper Networks Platforms
Standards-based Common Programmatic Management Interface Across all Juniper Networks Platforms (DMI)
T-series
M-series
E-series
WX/WXC
FW/VPN
IDP
J-series
DX
SSL VPN
ISG/IDP

Figure 12: Network Management Framework Built on Juniper Networks Products

Juniper Networks provides a comprehensive set of manageability, network management tools and partnerships for end-to-end management of the next-generation data center.

- Device Manageability Tools—Juniper Networks provides an open standards (XML and NETCONF-based) Device Management Interface (DMI) to manage all security and network devices. The DMI-based interface allows the device to express its manageability in an XML schema, thereby allowing existing enterprise network management systems to quickly discover and adapt to the new Juniper Networks infrastructure.
- Centralized Management—In addition to standards-based device manageability tools and interfaces, Juniper Networks provides centralized network management applications for comprehensive device management, policy and configuration, event and application visibility management. Juniper Networks has also partnered with IBM Tivoli to provide an end-to-end systems management solution including applications and servers. This is truly a complementary partnership, providing the necessary tools for the network infrastructure management of devices that are integrated using open XML/Web Services Description Language (WSDL) interfaces that communicate with the IBM Tivoli end-to-end Systems Management interfaces.

Summary

Juniper Networks uses an open systems approach for designing the data center network infrastructure that enables enterprises to design a high-performance data center network, as well as devices and elements to achieve a more efficient, manageable, and flexible network infrastructure. This greatly simplifies the design and enables operational efficiencies by deploying networks that are agnostic to multiple media types.

The architecture virtualizes critical network infrastructure components and functionalities, for example security, load balancing, and applications acceleration deployed and managed using a combination of business as well as technical heuristics. The architecture optimizes network performance and increases efficiencies of the network infrastructure. The architecture also automates network infrastructure management by plugging smoothly into the customer's existing management frameworks and third-party tools such as IBM Tivoli.

Glossary

3DES	Triple Data Encryption Standard
AES	Advanced Encryption Standard
AFE	Applications Front End
ATM	Asynchronous Transfer Mode
CIFS	Common Internet File System
CRM	Customer Relationship Management
CPE	Customer Premises Equipment
DHCP	Dynamic Host Configuration Protocol
DMZ	Demilitarized Zone
DNS	Domain Name System
DoS	Denial of Service
ERP	Enterprise Resource Planning
FR	Frame Relay
GRE	Generic Routing Encapsulation
GSLB	Global Server Load Balancing
HA	High Availability
HTTP	Hypertext Transport Protocol
HTTPS	Hypertext Transport Protocol over Secure Sockets Layer
IC	Infranet Controller
ICA	International Communications Association
IDC	Internet Data Center
IDP	Intrusion Detection and Prevention
IKE	Internet Key Exchange
IM	Instant Messaging
IMAP	Internet Message Access Protocol
IPS	Intrusion Prevention System
IPSec	Internet Protocol Security
ISDN	Integrated Services Digital Network
ISP	Internet Service Provider
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
L2	Layer 2
L2TP	Layer 2 Transport Protocol
L2VPN	Layer 2 VPN
MAPI	Message Application Programming Interface
MPLS	Multi protocol Label Switching
NAC	Network Access Control
NAT	Network Address Translation
NBT	Net BIOS over TCP/IP
NFS	Network File System
NLS	Network Layer Services
NSM	NetScreen Security Manager
NTP	Network Time Protocol
PE	Provider Edge
POE	Power over Ethernet

POP3	Post Office Protocol Version 3
PPVPN	Provider Provisioned VPN
PTP	Point-to-point
QoS	Quality of Service
RDP	Remote Desktop Protocol
RTSP	Real-Time Streaming Protocol
SAML	Security Access Markup Language
SFA	Sales Force Automation
SIP	Session Initiation Protocol
SLB	Server Load Balancing
SMB	Server Message Block
SMTP	Simple Mail Transfer Protocol
SOAP	Simple Object Access Protocol
ISG	Integrated Security Gateway
TE	Traffic Engineering
TELNET	Teletype Network
TNC	Trusted Network Computing
UAC	Unified Access Control
URL	Uniform Resource Locator
URPF	Unicast Reverse Path Filtering
UTM	Unified Threat Management
VNC	Virtual Network Computing
VoIP	Voice over Internet Protocol
VPN	Virtual Private Network
VRF	Virtual Routing and Forwarding
WF	WiFi Wireless
WINS	Windows Internet Name Service
WLAN	Wireless Local Area Network
WX	WAN Application Acceleration
WX CMS	WX Central Management System

Appendix A Juniper Networks Data Center Network Solution Tables

Data Center Product Tables

Infrastructure			Services	Policy and Management
Routing	**Switching**	**Security/VPN**	**Secure Access**	**Policy & Management**
MX960	EX 3200	ISG 2000*	SA 6000	IC 6000
M320	EX 4200	ISG 1000*		NSM
M120		NetScreen-5400		WX CMS
M10i		NetScreen-5200		Odyssey® Access Client (OAC)
		IC 6000		Steel-Belted Radius® (SBR)
		IC 4000		

* - Optional Integrated IDP

Data Center Product Tables (by Tier)

	Edge	Core	Network Services
Routing	M320 M120 M10i	MX 960	–
Switching		EX 3200 EX 4200	EX 3200 EX 4200
Firewall	ISG 2000* ISG 1000*	–	NetScreen-5400 NetScreen-5200 IC 6000 IC 4000
Secure Access (SSL VPN)	–	–	SA 6000
WAN Optimization	–	–	WXC 500 Stack
Policy and Management	–	–	IC 6000 NSM WX CMS OAC SBR

Partner Products

Symantec

Juniper Networks has teamed with Symantec Corporation to leverage its market-leading anti-spam solution for Juniper Networks' small-to-medium office platforms, helping to slow the flood of unwanted email and the potential attacks they carry. Part of a complete set of UTM features available on Juniper Networks firewall/VPN gateway, the anti-spam engine filters incoming email for known spam and phishing users to act as a first line of defense. When a known malicious email arrives, it is blocked and/or flagged so that the email server can take an appropriate action.

Kaspersky

By integrating a best-in-class gateway antivirus offering from Kaspersky Lab, Juniper Networks integrated security appliances can protect Web traffic, email and Web mail from file-based viruses, worms, backdoors, trojans and malware. Using policy-based management, inbound and outbound traffic can be scanned, thereby protecting the network from attacks originating from both outside and inside the network. Unlike other integrated antivirus solutions that are packet or network signature-based, the Juniper-Kaspersky solution deconstructs the payload and files of all types, evaluating them for potential viruses and then reconstructs them, sending them on their way. The Juniper-Kaspersky solution detects and protects against over 100,000 viruses, worms, malicious backdoors, dialers, keyboard loggers, password stealers, trojans and other malicious code. Included in the joint solution is a best-in-class detection of spyware, adware and other malware-related programs. Unlike some solutions that use multiple non-file based scanners to detect different types of malware, the Juniper-Kaspersky solution is based upon one unified comprehensive best-of-breed scanner, database and update routine to protect against all malicious and malware-related programs.

SurfControl and Websense

All Internet content that is read, sent or received carries inherent risks. Employee access to the Internet continues to introduce new dangers and content that can negatively impact an enterprise in four fundamental ways:

- Security Threats: Viruses, spyware and other malware can all enter an enterprise's network through Web-based email, file downloads, instant messaging, P2P applications and other non work-related sites.
- Legal Threats: Content that is inappropriate can lead to gender, minority or religious harassment and discrimination. Illegal downloading and distribution of copyrighted or illegal material over an enterprise's network has legal liability issues as well.
- Productivity Threats: Temptations of non work-related Web destinations are endless. Just 20 minutes of recreational surfing a day can cost a company with 500 employees over $8,000 per week (at $50/hour/employee).
- Network Threats: Employees can crash an enterprise's network just by logging in to the wrong Web site. Other activity like recreational surfing and downloading MP3 files can divert valuable bandwidth from critical business needs.

To regulate inappropriate Web usage, Juniper Networks has teamed with both SurfControl and Websense to provide either an integrated (on-box) or redirect (two boxes) Web filtering solution.

- Integrated Web Filtering: Integrated Web Filtering leverages an "in the cloud" architecture hosted by SurfControl's certified hosting partner that allows enterprises to build Web access policies from the largest URL database (over 6 million pages) spread across more than 40 categories. From the WebUI or NetScreen Security-Manager (NSM), an administrator can assemble firewall policies that incorporate and enforce Web access rights.
- Redirect solution with SurfControl or Websense: Traffic is redirected from any of the firewall/VPN appliances to a customer-hosted server running the Web filtering software where Web access grant/deny decisions are made and executed. The customer is responsible for the server, the software and the associated management of the solution. Redirect Web filtering is supported across the entire product line.

Avaya IG550

The Avaya IG550 Integrated Gateway provides an additional choice in the Avaya line of Media Gateways. Enterprises can now consolidate the number of devices that they deploy and manage in their branch offices. This solution provides high-sustained network performance when under load, integrated voice and data security, and multilevel business continuity options. This best-in-class solution is available through Avaya direct channel and certified Avaya and Juniper Networks resellers.

The Avaya IG550 Integrated Gateway consists of two primary components: a Telephony Gateway Module (TGM) and Telephony Interface Modules (TIMs).

The TGM550 module inserts into any slot in the J4350 or J6350 router and delivers a rich telephony feature set to the branch office. This feature set includes:

- Access to central Avaya Communication Manager and other communications applications
- Support for call center agents
- 6-party meet-me conferencing
- Local survivability in the event of a WAN failure
- Local music-on-hold and voice announcements
- Full encryption of voice traffic

The TGM operates as any other Avaya H.248-based gateway and includes a two-analog trunk/two-analog station module, modular Digital Signal Processors (DSPs) and a memory expansion slot.

There is a choice of several TIMs with analog, T1/E1/PRI and BRI options. The TIM514 analog module contains four trunks (FXO) and four stations (FXS); the TIM510 DS1 module supports T1/E1 and ISDN-PRI; and the TIM521 module supports four ISDN-BRI interfaces.

Appendix B Juniper Networks Core Network Power Efficiency Analysis

Characteristics	Juniper Networks Core MX960 2x Chassis
Line-rate 10GigE (Ports)	96
Throughput per Chassis (Mpps)	720
Output Current (Amps)	187.84
Output Power (Watts)	9020.00
Heat Dissipation (BTU/Hr)	36074.33
Chassis Required (Rack Space)	2 Chassis
Rack Space (Racks)	2/3rd Racks

About Juniper Networks

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

CORPORATE HEADQUARTERS
AND SALES HEADQUARTERS FOR
NORTH AND SOUTH AMERICA
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

EUROPE, MIDDLE EAST, AFRICA
REGIONAL SALES HEADQUARTERS
Juniper Networks (UK) Limited
Building 1
Aviator Park
Station Road
Addlestone
Surrey, KT15 2PG, U.K.
Phone: 44.(0).1372.385500
Fax: 44.(0).1372.385501

EAST COAST OFFICE
Juniper Networks, Inc.
10 Technology Park Drive
Westford, MA 01886-3146 USA
Phone: 978.589.5800
Fax: 978.589.0800

ASIA PACIFIC REGIONAL SALES HEADQUARTERS
Juniper Networks (Hong Kong) Ltd.
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803