

GOVERNMENT DATA CENTER NETWORK REFERENCE ARCHITECTURE

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

Table of Contents

Executive Summary	4
Target Audience	4
Introduction	4
Trends and Challenges	5
Juniper Networks Approach and Solution	7
Government Data Center Network Design Considerations	7
Virtualization	8
A Green and Environmentally Friendly Data Center	8
High Availability Disaster Recovery	8
Visibility	9
Network Connectivity	9
Security	9
Policy and Control	10
QoS	10
High Performance	10
Juniper Networks Data Center Network Architecture	11
Open Systems Approach – Juniper Networks Government Framework	11
Location-Based Approach	12
Design Principles	13
High-Level Architecture	13
Edge Services Tier	15
Edge Services Connectivity	15
Edge Services HA	15
Edge Services Performance	16
Edge Services Security	16
Core Network Tier	17
Core Network Connectivity	18
Core Network HA	18
Core Network Virtualization	18
Network Services Tier	20
Data Center Security Services	21
Application Front-Ending Services	21
Applications and Data Services Tier	22
Storage Area Networks (SANs)	24
Fibre Channel SANs	24
iSCSI SANs	24
Data Center Backbone	25
Data Center Network Management	26
Conclusion	27

Appendix A: Juniper Networks Data Center Network Solution Tables	28
Partner Products	29
Symantec	29
SurfControl and Websense	29
Avaya IG550	29
Appendix B: Juniper Networks Core Network Power Efficiency Analysis	30
About Juniper Networks	30

Table of Figures

Figure 1: Location-based perspective of the government agency network	6
Figure 2: Data center network functional design model	7
Figure 3: The Juniper Networks government framework	11
Figure 4: Network connectivity to the data centers	12
Figure 5: Juniper Networks data center network architecture	14
Figure 6: Data center network edge services	16
Figure 7: Data center core network and network services	19
Figure 8: Connectivity systems, application systems and network service systems	20
Figure 9: Data center application network types/purposes	22
Figure 10: Application and data services network view	23
Figure 11: Data center backbone connectivity	25
Figure 12: Network management framework built on Juniper Networks products	27

Executive Summary

The next-generation data centers currently being used by enterprises and data hosting facilities also allow governments to reduce operational cost while delivering transparency to citizen services through an open architecture. Leveraging the Federal Enterprise Architecture (FEA), the U.S. government has sought industry input to develop IT solutions that optimize investments in commercial off-the-shelf (COTS) technology to cost-effectively build high-performance data centers. The Juniper Networks® “Innovation in Government” concept helps agencies to provide a responsive and trusted environment that drives mission assurance. Government agencies trust Juniper Networks to provide a comprehensive approach to building next-generation data centers that leverage the FEA framework by utilizing best-in-class products with well-defined practices that can be replicated across the government enterprise.

Target Audience

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

The remainder of this paper references the Juniper Network architecture approach and outlines best practices, technologies and products that support data center architects and engineers responsible for answering the requirements of designing government agencies’ data center networks.

Introduction

The purpose of this document is to provide government IT managers and administrators with a data center network architecture that mitigates 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, it 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 the scope of this solution, they are not addressed in this document.

Trends and Challenges

Federal, state and local government data centers experienced rapid growth during the late 1990s and early part of this decade to improve service effectiveness and achieve mission objectives. To deliver results, IT departments have increasingly become dependent on high-performance networks to meet the challenges of escalating demand for secure networks that can cost-effectively scale, have lower operational cost, can react quickly to changing technology trends such as virtualization, support Software as a Service (SaaS), and facilitate unified communications and consolidation.

Effective access to massive amounts of data is critical to government operations—from highly confidential federal, state and local tax information to geospatial data for military operations. Unfortunately, the escalating complexity and risk of managing today's geographically distributed legacy network infrastructures are an increasing impediment to cost-effective data center management. With budget a constant concern, government data centers tend to consist of a heterogeneous assortment of multiple platforms, network tiers and management tools, making data center consolidation and administration more complex and time-consuming. In most cases, agencies choose to control costs by using open standards-based COTS technologies that must also be integrated with these legacy systems.

To complicate matters, government agencies face a multitude of connectivity mandates such as migrating to IPv6 and reducing the number of Internet connections. Given that, for example, the U.S. Department of Energy will soon be utilizing a full fiber of data, it is clear that networking infrastructures need to become more efficient and effective to provide 24/7 availability and six nines of performance. Mobile workers and third-party contractors are also adding new pressures to the data center complex, requiring access to data from phones, handhelds and laptops from anywhere, at any time.

At the same time, data needs to be kept completely secure. As of February 2008, a report from the Government Accountability Office (GAO) stated that, "Federal agencies are still not fulfilling all the mandates for protection of personal information." Notorious incidents such as multiple thefts of information from the Veterans Administration brought to light the fact that in many cases agencies do not know what information they have, who has access to it or the location of the devices on which it has been stored. With data centers being managed using a wide variety of security policies and tools, weaknesses can too easily lead to a security breach. As a result, government data center administrators are under intense public scrutiny not only to secure information, but also to provide transparency for audits to show that they are meeting regulatory requirements such as the Federal Information Security Management Act.

Another major trend in government is an increased emphasis on "green IT." After heating, ventilation and air conditioning (HVAC) systems, data centers consume the most power in any government building. Server use is growing at an annual rate of 11 percent, with storage at an even higher rate of 22 percent, causing tremendous strain on the data center's power and cooling capacity. The use of service-oriented architectures (SOAs) and Web services is also becoming increasingly network intensive.

When, at the end of 2007, President George W. Bush signed into law the Energy Independence and Security Act (EISA) to reduce energy usage, he required government agencies to seek out new methods of improving energy efficiency. As a result, agencies are instituting green initiatives that track resource usage, carbon emissions and efficient utilization of resources such as power and cooling for their data centers. Based on this information, they are considering new data center networking technologies to help improve power capacity, reduce cooling costs, maintain availability and support disaster recovery, while sustaining constant growth and operational management in the data center.

These technologies also support the trend toward consolidation and virtualization of government data centers, which is reducing the number of facilities and operating locations. However, this also means that architects are faced with the challenge of designing data centers that centralize servers and applications while keeping them accessible from a variety of locations (see Figure 1).

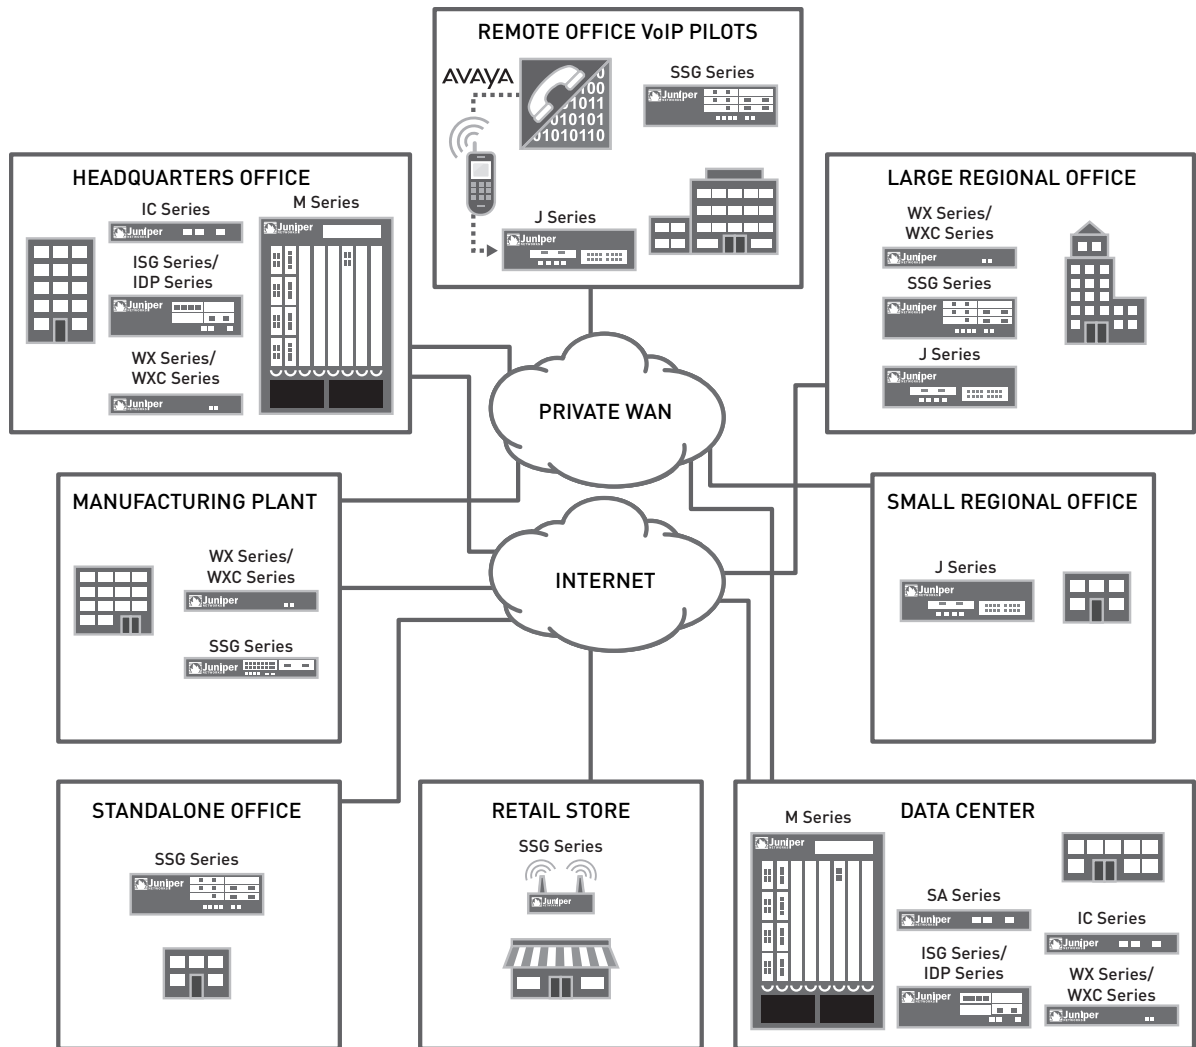


Figure 1: Location-based perspective of the government agency network

Simply designing a data center to deploy more servers, storage and devices significantly increases network complexity and cost. Government agencies must change the way they view their data center network architecture in order to maximize efficiency gains. This new architecture uses virtualization capabilities such as MPLS and virtual private LAN service (VPLS) to enable an extremely fast, high-performance data center backbone network.

A government data center network must also offer security, performance acceleration, high density and a resilient network infrastructure. These critical components help ensure that users have the support they need to succeed in their work. This document shares Juniper Networks best practices in designing a highly efficient, secure, scalable and flexible data center network. It also showcases advanced network technologies such as high-density, next-generation Ethernet switches, application delivery controllers and WAN acceleration, which can be employed to create a seamless user experience irrespective of their location on the network.

Juniper Networks Approach and Solution

Juniper Networks strategy for designing the data center network uses an open systems approach. This enables agencies to design a high-performance data center network that consolidates network elements into fewer networks that employ fewer network devices, consuming significantly less power, cooling and physical space. Simplifying the network architecture enables operational efficiencies and offers data center networks that are agnostic to multiple media types.

This architecture also virtualizes critical network infrastructure components and functionalities such as security, load balancing and application acceleration, which it deploys and manages based on a combination of operational and technical heuristics. It also optimizes network performance and increases efficiencies within the infrastructure. Management of the network infrastructure is automated by connecting smoothly into existing management frameworks and third-party tools such as IBM Tivoli.

Juniper Networks single network operating system—Juniper Networks JUNOS® Software—further reduces the complexity, cost and risk associated with deploying a new network infrastructure, features and functionality. Its consistent, standards-based management optimizes provisioning, new service delivery and network administration. The architecture also improves data center availability and security in terms of network protection, connectivity failures and disaster recovery, while integrating security services that can be virtualized across the data center, the agency and to other federal and state agencies as needed.

Government Data Center Network Design Considerations

The following section summarizes some of the technical considerations for designing a modern data center network that supports a consolidated and centralized server and storage infrastructure, as well as agency 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.

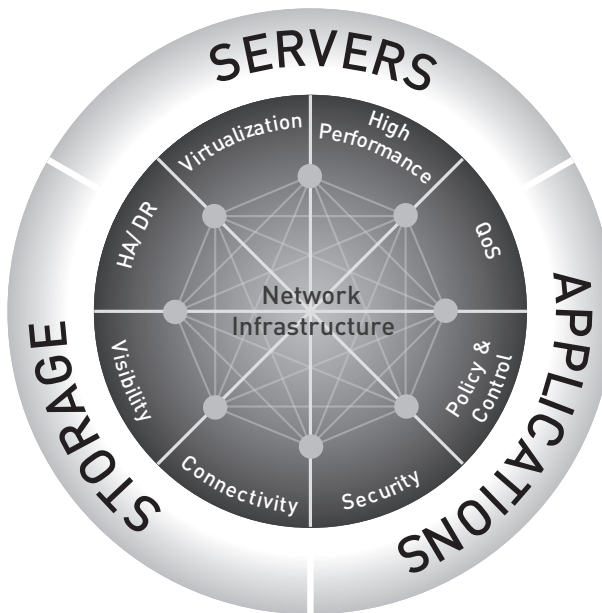


Figure 2: Data center network functional design model

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. These typically comprise the storage, server, application and network groups.

Observing all of the installed devices in the data center, one can 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 agency requirements. These three silos are connected through a fast, secure and reliable data center network fabric that 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 include:

- 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 – data security and regulatory compliance
- Policy and Control – centralized policy and control
- Quality of Service (QoS)
- High Performance – applications, storage, servers and the network

Virtualization

Virtualization is a technique for hiding the physical characteristics of computing resources from other systems, applications or end users interacting with those resources. A single physical resource—such as a server, operating system, application or storage device— thus appears to function as multiple logical resources; multiple physical resources, such as storage devices or servers, may appear as a single logical resource; or, one physical resource may 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. An example of data plane virtualization is a single physical interface that provides security to multiple network segments using 802.1q VLAN tagging. Control plane virtualization could include multiple routing domains and protocol instances. An example of management plane virtualization supports multiple logical firewall/VPN security systems that use Virtual Systems (VSYS) for true multi-agency environments, such as when state and federal agencies must work together to enact a Homeland Security directive.

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 ones are built, it is important to ensure that the 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 to 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. 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. This appendix can be used as an example for comparative analysis against other core solutions.

High Availability Disaster Recovery

HADR is a key requirement for the data center network and must be considered for all existing data center facilities. Network HA is deployed by using combinations of link redundancy (for both external and internal connectivity) and critical device redundancy to ensure network operations and business continuity. In addition, site redundancy (multiple data centers) is critical to meeting disaster recovery and regulatory compliance objectives. Moreover, devices and systems deployed within 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 ISSUs to prevent software failures/upgrade events from impacting the entire device. Assuring that software events impact only a particular module helps maintain system availability.

Visibility

It is important to have visibility into traffic and security events in order to effectively maintain and manage network resources. This includes the ability to collect IP traffic flow statistics to give organizations 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 granular viewing of security events to help determine how these are being handled. Extending this visibility to develop a deeper understanding of application-specific traffic provides a wide range of operational and performance information that can impact application users. For example, specific compression and acceleration technologies can be applied at the network layer to accelerate email applications such as Microsoft Exchange. Or, it may be necessary to bar employee access to services such as YouTube and social networking sites, as they may impact internal application performance or violate agency security procedures. Understanding the application (YouTube, instant messaging) and enforcing appropriate policies ensure that performance meets or exceeds the expectations of end users.

Network Connectivity

Agency employees and associated third-party contractors all require immediate access to applications and information. Citizen services applications also demand significant network performance. The challenge of working from multiple locations further increases the complexity of providing consistent data access. As part of the data center network design, the following critical aspects of external network connectivity must therefore be considered:

- WAN connectivity to enable distributed agency users to access applications
- Internet connectivity to enable secure remote access for remote and mobile users
- Superior speed for data center backbone connectivity and use of technologies such as VPLS and 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 requires high-speed, highly available network connectivity. In addition, LAN segments and networks may be deployed that require different security and capacity levels and services. Typically, connections of 1 Gbps and higher (while 10 Gbps is becoming the standard) should be available in the data center network, providing at least 1 Gbps to the server and preferably 10 Gbps at network choke points.

Security

The most critical resources in any agency location are typically the applications themselves and their servers and supporting systems, such as storage and databases. Financial, human resources (HR) and citizen-facing applications with supporting data can, if compromised, create a potential operations and public relations disaster.

The core network security layers must therefore protect these mission-critical resources from unauthorized user access and attacks, including at the application level. The security design needs to employ layers of protection from the network edge through the core to the various endpoints. Multiple layers of security protect critical network resources: If one layer fails, the next steps up to stop the attack and/or limit the damage. This security approach 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 within virtualized configurations, and it can be used to provide granular network access control (NAC). Policy and control capabilities allow agency IT organizations to centralize policy management while offering distributed enforcement. The network policy and control solution permits 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 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.

QoS

In order to assure a high-quality application experience over large networks, QoS levels are assigned and managed to ensure satisfactory performance. A minimum of three levels of QoS (each of which determines a priority for applications and resources) is as follows:

- Real-time
- Mission-critical
- Best effort

MPLS networks and network traffic engineering capabilities are typically deployed to configure label switched 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, as well as rapid and efficient management to ensure end-to-end QoS for the agency.

High Performance

To effectively address performance requirements related to virtualization, server centralization and data center consolidation, the data center network needs to boost performance of all application traffic, whether local or remote. Providing a LAN-like experience for all users irrespective of physical location, the data center network should optimize applications, servers, storage and network performance.

WAN optimization techniques include data compression, TCP and application protocol acceleration, bandwidth allocation, and traffic prioritization to improve performance network traffic. These techniques can also be applied to data replication, as well as 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, 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 can be built on a nonblocking 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 should also account for remote users.

Juniper Networks Data Center Network Architecture

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

Open Systems Approach – Juniper Networks Government Framework

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

- Applications
- Services
- Infrastructure

The Applications layer provides support to the various software applications required to operate the agency and carry out its mandates. The Services layer combines the traditional presentation, session and transport layers. It also supports users and applications. This layer 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 agency's centralized policy and management functions to help reduce operations costs while enabling compliance. All three layers are interconnected with open standards-based interfaces that allow the organization to seamlessly deploy a multi-vendor solution, providing the flexibility to use the best technologies to meet organizational requirements.

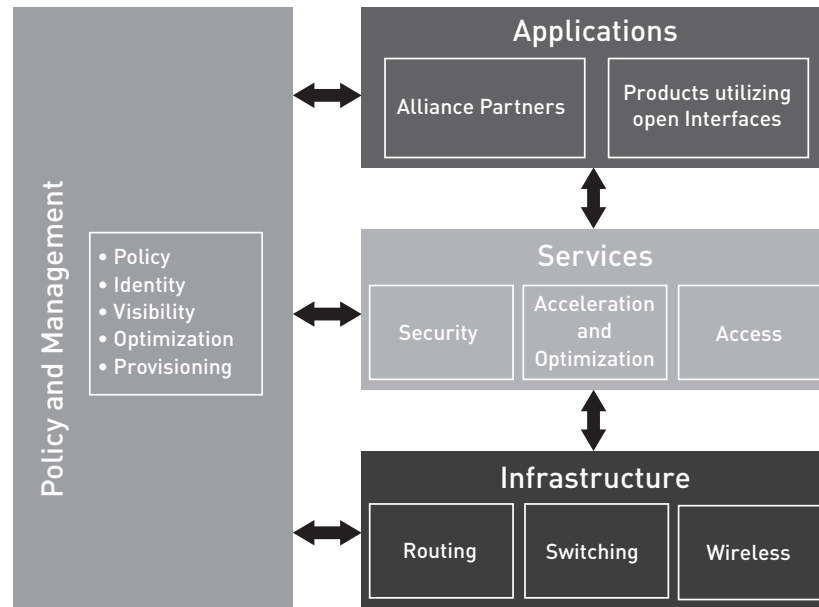


Figure 3: The Juniper Networks government framework

The Juniper Networks Government Framework supports the next-generation data center network by creating a best-in-class network environment that uses open standards-based and industry-accepted interfaces. Government agencies can use this framework to logically view their network infrastructure and applications in order to make decisions that best serve user requirements.

Juniper Networks takes a holistic approach to next-generation networking that allows for the user, network and applications perspectives. Its understanding of applications and how they are accessed from a variety of locations enables an architecture that meets the demands of a variety of users.

Location-Based Approach

The key function of the data center is to offload “always on” requirements from various locations to a central, stable location that contains the most recent application data. By decoupling the information store from the physical location of the user, agencies derive greater efficiencies by creating a centralized pool of resources. This trend of centralizing applications and consolidating multiple facilities increases the importance of the WAN and other external networks, as users 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 agency’s private WAN and the Internet edge that hosts remote user connections.

The data center does not typically host users and most certainly does not accommodate data center application users. However, this model can support different operational requirements unique to each agency. 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 connectivity into the data center and connectivity between data centers.

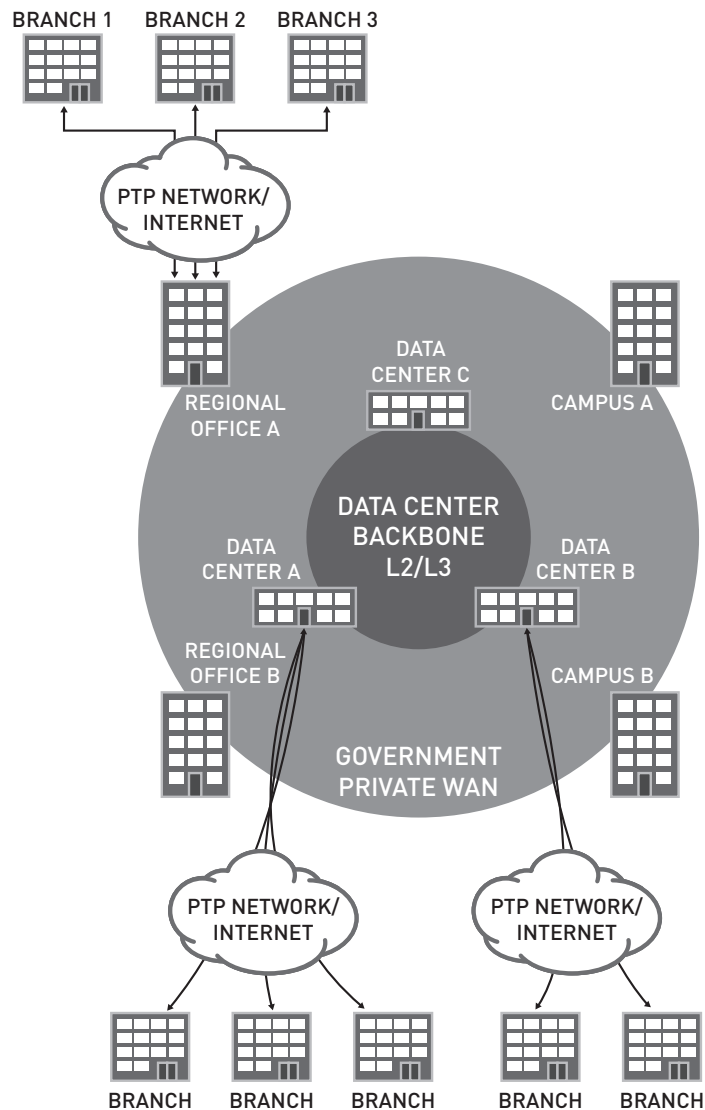


Figure 4: Network connectivity to the data centers

Design Principles

Key design principles are derived from operational and technical objectives. The operational objectives are fairly clear—reduce operation expenses, maintain security, adhere to “green IT” principles and so on. The top-level technical requirements include:

- 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 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 agency entities, such as contracts management or tactical operations. Uniqueness and operational freedom must remain “virtually” independent.

Virtualization – The virtualization of processing has introduced a new standard in resource pooling and resource utility optimization. Such technologies at various levels are introduced into the data center, from large storage arrays and servers to network virtualization and 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.

HA – Consolidating and centralizing resources, as well as virtualizing technologies, make 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:

- Component
- Device
- Link
- Site

Streamlined Operation and Management of Data Center Services – A consolidated and virtualized environment relies on a single management platform that can control servers, applications, storage and network infrastructure as one. Hence, devices and systems need to support open standards-based interfaces and protocols, so that they can all be controlled from existing and evolving management systems.

High-Level Architecture

Figure 5 illustrates the Juniper Networks data center network architecture. The major architectural tiers include:

- 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, the different network tiers are explored in greater detail.

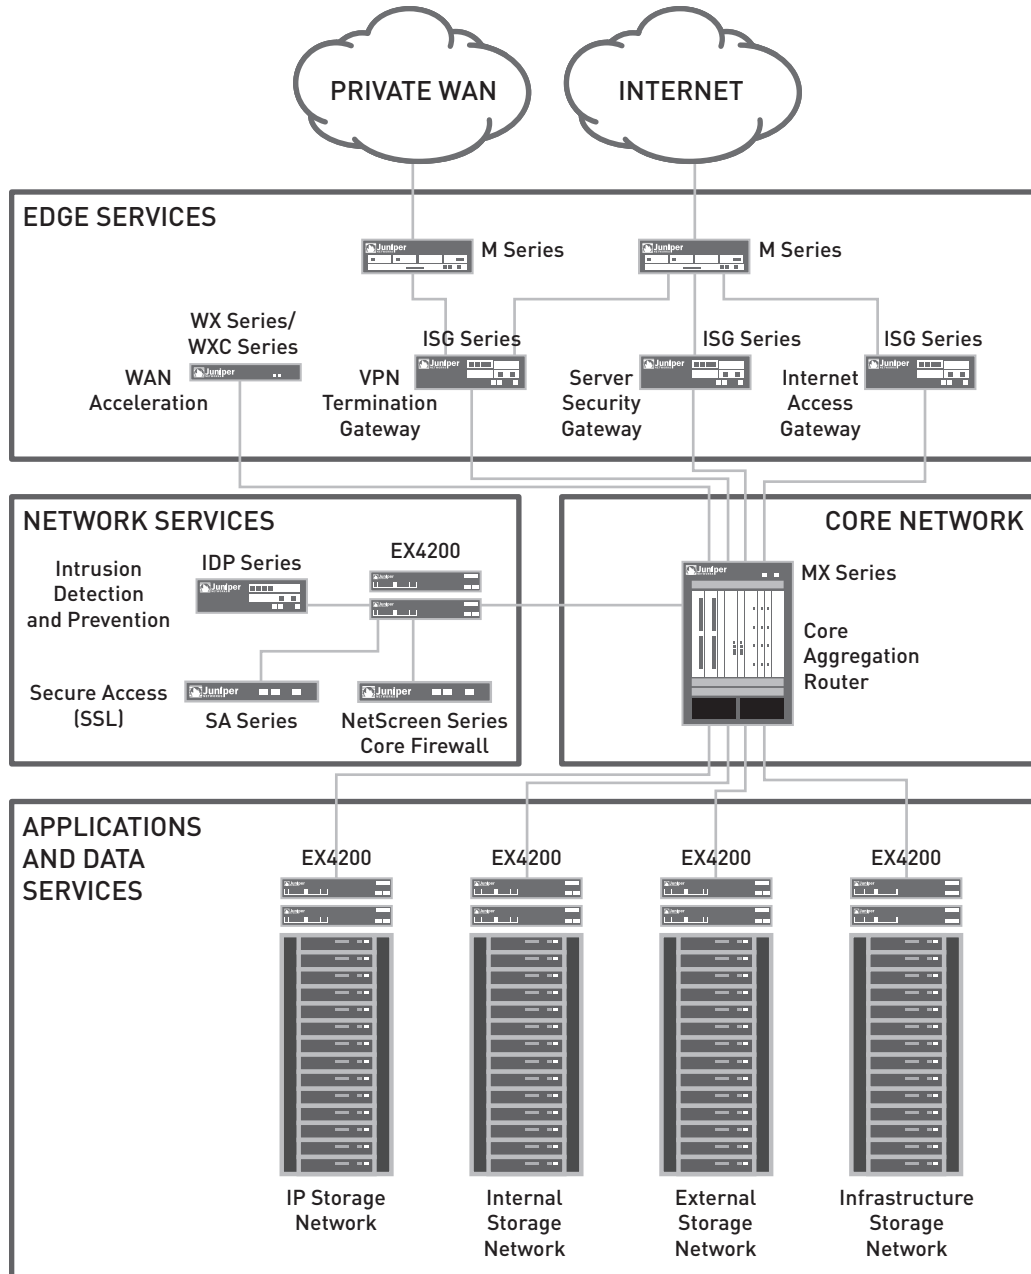


Figure 5: Juniper Networks data center network architecture

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

It is important to support network applications with an extremely fast core network capable of forwarding total aggregate traffic at line rate. Effectively, the network core spans across multiple locations and devices. Logically, the network core connects all data center networks directly to itself. This attribute enables rack/location-agnostic server-to-network binding, which is a key element in building a virtualized data center fabric that supports automatic repurposing of computing 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 supporting 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, which 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. As the data center edge network serves as the key boundary to the data center, it is responsible for maintaining reachability into all other external networks.

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 other agencies, third-party contractors and to the Internet. For connecting all of these networks, it is important to plan for:

- Internet routing isolation – for example, to separate the exterior routing protocols from the interior routing protocols
- Network address translation (NAT) to convert private IP addresses to public Internet routable IP addresses
- IPsec VPN tunnel termination for agency and employee connections
- Border security to enforce stateful firewall policies and content inspection
- QoS

Network architects have often used Layer 2 switches at the edge to form a hierarchical mesh, with the intention of allowing a multitude of links to provide fault protection during failure. The Juniper Networks solution employs Juniper Networks M Series Multiservice Edge Routers and Juniper Networks ISG Series Integrated Security Gateways. This leverages the routing functionality of the ISG Series to provide a routed connectivity solution instead of a traditional switched mesh. It also 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 commonly used at the edge.

Edge Services Connectivity

Figure 6 shows the Juniper Networks Edge Services design and illustrates how this 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 M Series Multiservice Edge Routers and are the edge devices for both Internet and private WANs. The M Series routers were selected for both their interface capacity and throughput.

Each router has a single connection to the Internet (or private WAN). Connectivity between the edge routers to each ISG Series firewall creates a fully meshed network. Administrators 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 number of failure cases that could impede operational continuity.

Edge Services HA

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. 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 enabling 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) are mandatory for edge-deployed devices.

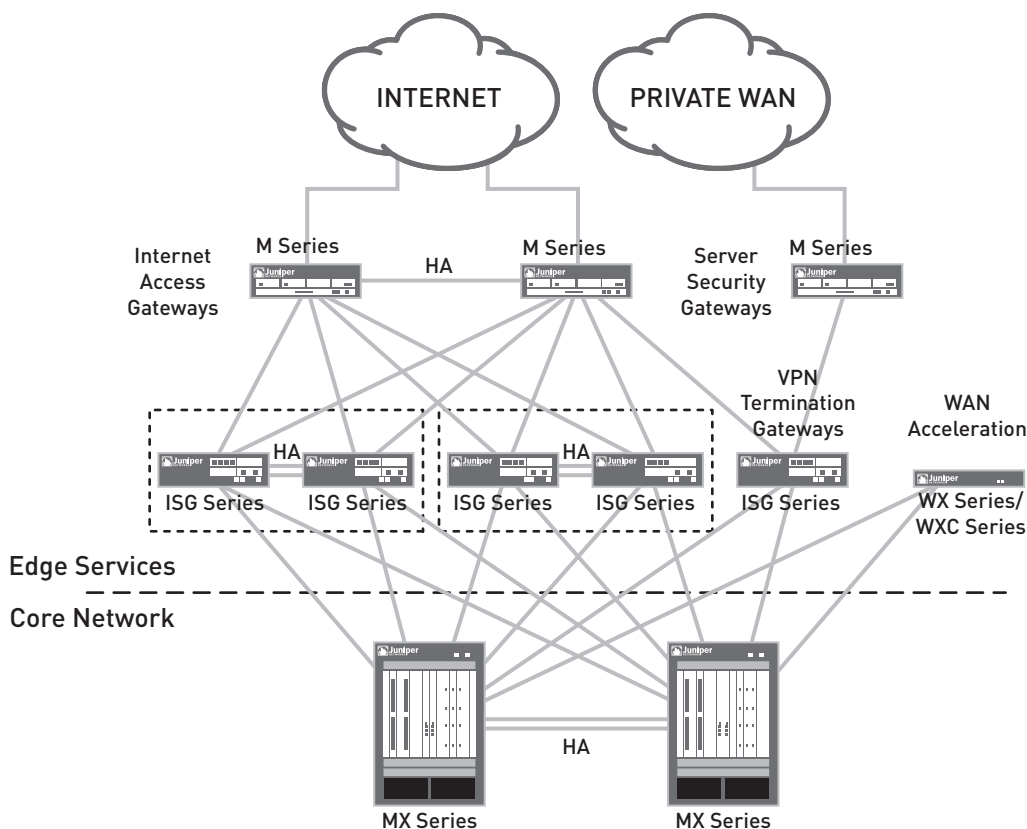


Figure 6: Data center network edge services

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. A single link failure, therefore, will not usually bring down the device, and with it, a viable path.

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 provides further insurance, 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 WAN acceleration technologies include redundant tunnels and load-balanced 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 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, VPN secure connectivity services are implemented at the edge. 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.

The second set, Juniper Networks SSG Series Secure Services Gateways secure the server and data resources and software applications for inbound traffic originating from the Internet.

The third set, the IPsec VPN firewalls, comprises the connectivity hub for all remote sites and terminates IPsec VPNs from the Internet, as well as from the private WAN. The IPsec firewalls also terminate VPN tunnels for all remote users over the private WAN. To provide remote services, 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.

By performing general DoS protection at the Edge Services tier, security intelligence is moved closer to the provider edge, decreasing the number of devices that can be potentially compromised during attacks. A large flood presents challenges to any network, as it can consume all available network bandwidth and 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 by a flood, it will experience high CPU load and may drop legitimate traffic. The specific rate 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 kinds of traffic, such as TCP control traffic and ICMP types, to rates that will not impact available bandwidth and overwhelm the firewall.

In selecting VPN design and encryption protocols, trade-offs must be made. Organizations should choose the strongest encryption algorithm that does not compromise performance requirements for the network while maintaining security. A longer key length provides more security against brute force attacks, yet may require more computational power. Therefore, this approach lowers performance when 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 devices dramatically, especially those with many VPN peers.

One also must consider bulk encryption algorithms. Typically, they are symmetrical and least influenced by design due to hardware assistance and the lower cost of handshakes. However, if the design presents few VPN peers and extensive data transfer, this element should be considered: The lowest common denominator will be the speed that determines VPN capacity. Finally, one should consider hashing algorithms. This selection is primarily done based on security requirements, but if hardware assistance is involved, 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 three-tier model. Three-tier networks add an aggregation network between access networks and core networks, and are the primary method for extending networks because of the scalability limitations of most available core network devices. Aggregation at the core allows more flexibility and easier support for virtualization, but it also requires high-speed processing and HA levels. A two-tier network is one core network with all of the access networks connecting directly to it.

One of the greatest advantages of a two-tier design is a dramatic reduction in the number of devices. This 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 two-tier model is typically limited by the scalability of the core network devices. The more traditional three-tier design, which allows for high scalability requirements, is not discussed in this paper.

Core Network Connectivity

The core network provides data center fabric connectivity by unifying 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 VPN 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

By connecting all networks to the core network with full redundancy at the core, HA is achieved without added complexity and dependency to data center network protocols and convergence. Traditionally, adding HA requires redesign of the network, but by using standards-based redundancy protocols and a core network approach, HA is enabled with a lower operational overhead. As well as adding redundant devices, it is extremely important to ensure that the core data center devices support in-service operations, such as hot-swappable interfaces and software upgrades.

Core Network Virtualization

To achieve network virtualization from the server through the network core, a key assumption is that even when all features are turned on, network systems are deployed that deliver line-rate throughput. 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, depending on the scale of the data center and on its administrator's skill set. VLANs that extend all the way to the core are more appropriate for smaller networks. In the case of larger networks that require VLAN scaling limitations with more elaborate QoS requirements, MPLS is the preferred choice. The Juniper Networks data center network architecture and solution components support 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 LSPs with unique QoS and connectivity characteristics. The Juniper Networks two-tier architecture provides for a more flexible design option (see Figure 7). Juniper Networks MX Series Ethernet Services Routers reside in the core network and the Juniper Networks EX Series Ethernet Switches reside in the access layer.

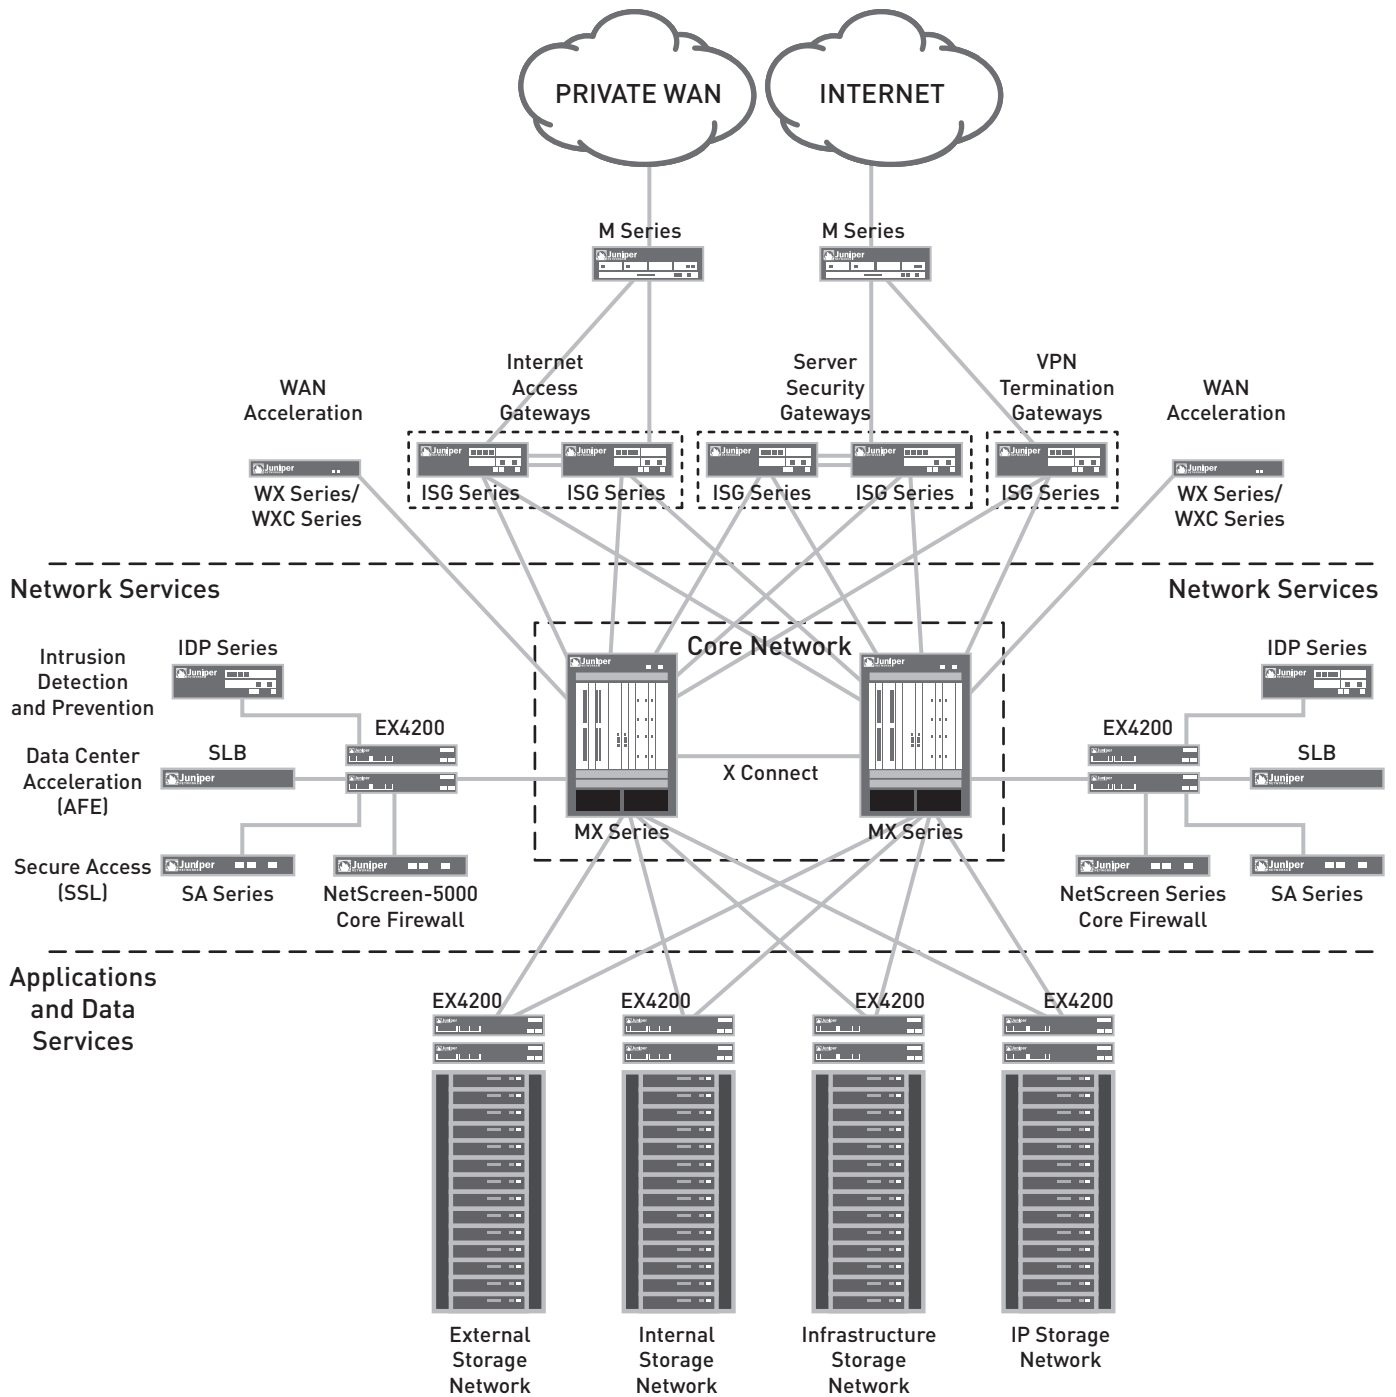


Figure 7: Data center core network and network services

Network Services Tier

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

- Security services, such as firewalls and intrusion prevention
- Application front-end services, such as server load balancing, SSL offload, HTTP cache, TCP multiplex, and global server load balancing (GSLB)

Throughout this section, both will be described in greater detail and the key elements that comprise the data center network architecture will be addressed.

The Network Services tier should extend to any of the server networks hosted in the data center, and apply a network-specific policy and set of configurations to appropriately interact 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 data center’s network core.

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 service. 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 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, application systems (Network N) and the network service systems (depicted on the far right of the diagram). This provides a core network perspective and shows the interaction between the core and the pooled service devices.

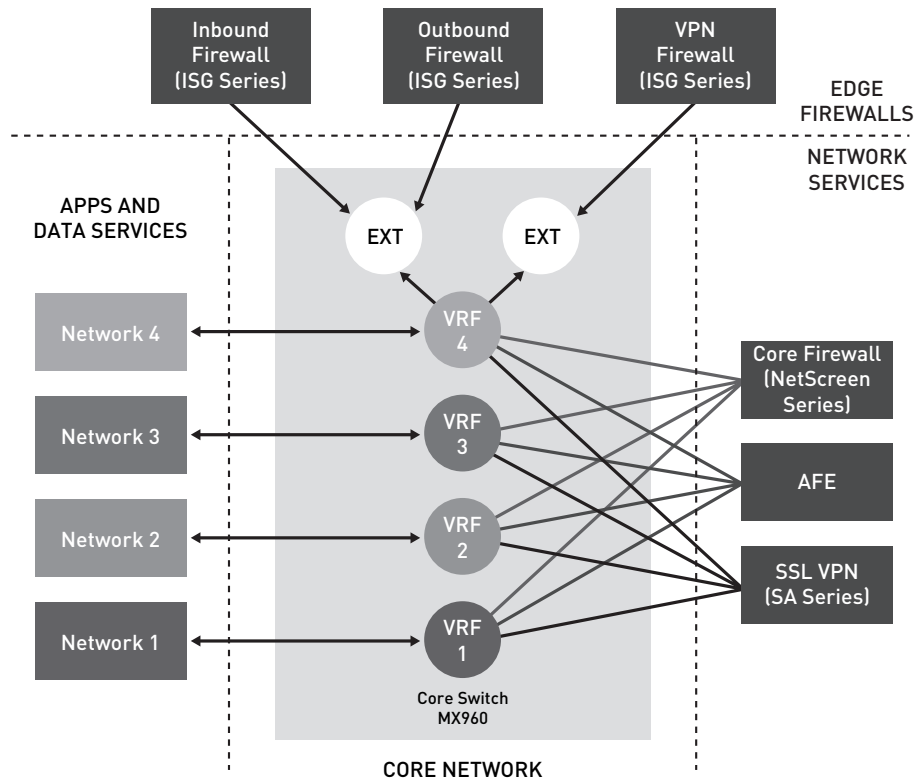


Figure 8: Connectivity systems, application systems and network service systems

Data Center Security Services

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

Stateful firewalls are the cornerstone of the data center network's security service. They enforce a security policy that aligns with agency mandates and operational requirements through the identification and classification of networks. In addition to being the primary L4 access control system, firewalls support many security functions in the data center—such as service DoS or quota protections, Deep Inspection to specific applications as needed and potential NAT.

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 an administrator dedicates 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, which must be extended to the services layer in order to design a network that depends on services for its core functionality.

Juniper firewall systems can split into separate virtual domains of control and forwarding instances or VSYS, creating separate virtual domains that allow security policy autonomy to different departments. To connect all of the core networks, the core firewall must participate in routing protocols within the data center network.

In addition to assuring secure connectivity at Layer 4, the Network Services tier employs application security services such as intrusion prevention to protect the data center infrastructure. Because these 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 colocated, 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 detect 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 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

Administrators need 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 data center servers.

The best answer is to deploy a system that supports acceleration for the different application tiers and also provides comprehensive capabilities around the more common emerging application areas such as Web 2.0. A data center acceleration solution should boost the performance of client/server, Web-based, and server-to server-applications, as well as speeding Web page downloads. In addition, the acceleration solution needs to 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 may serve as a standard load balancer. This means forwarding traffic to its destination address from a pool of available addresses.

Organizational 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. The Juniper solution addresses these requirements.

Applications and Data Services Tier

The Core Network tier connects to the Applications and Data Services tier that hosts all of the servers, databases and storage. Generally, there are four types of networks, with multiple instances of each type. The primary reasons for the multiple instances are separation of duties within the organization, and differentiated objectives and IT requirements for the different networks. Figure 9 illustrates the four networks:

- External Applications Network – There 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 serve different levels of internal access from within the organization’s various locations. These networks typically connect internal applications such as finance or healthcare services systems.
- 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 – The storage network is built on technologies including Fibre Channel, the InfiniBand serial link, and the Internet Small Computer System Interface (iSCSI) protocol. Critical application servers directly connect to storage devices through a separate Host Bus Adapter (HBA) to ensure fast access to data. Other servers connect using Ethernet to access storage facilities.

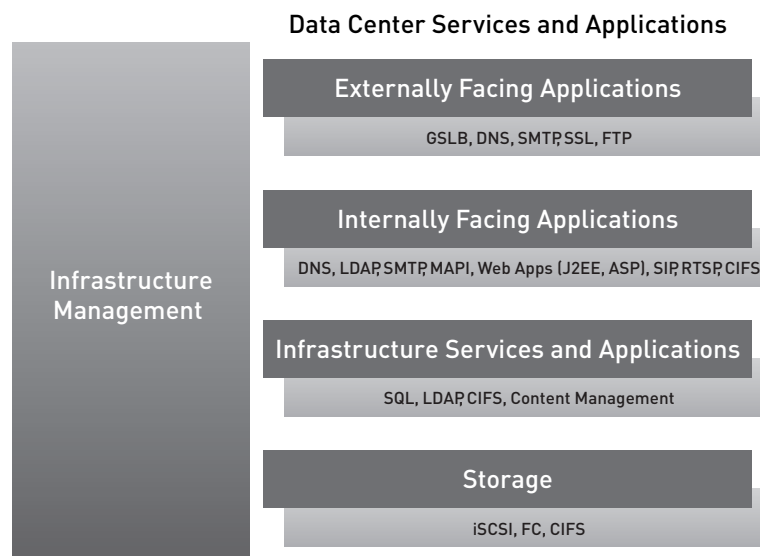


Figure 9: Data center application network types/purposes

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 aggregate multiple Layer 2 networks to use fewer connections.

Each internal and external applications network can be segmented into several subnetworks (see Figure 10). The servers that host these applications connect with at least a 1 Gbps (currently moving towards 10 Gbps) link to the EX Series switch with Virtual Chassis technology. The EX Series switch connects to the network core via a 10 Gbps connection. Depending on the number of servers, multiple EX Series 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.

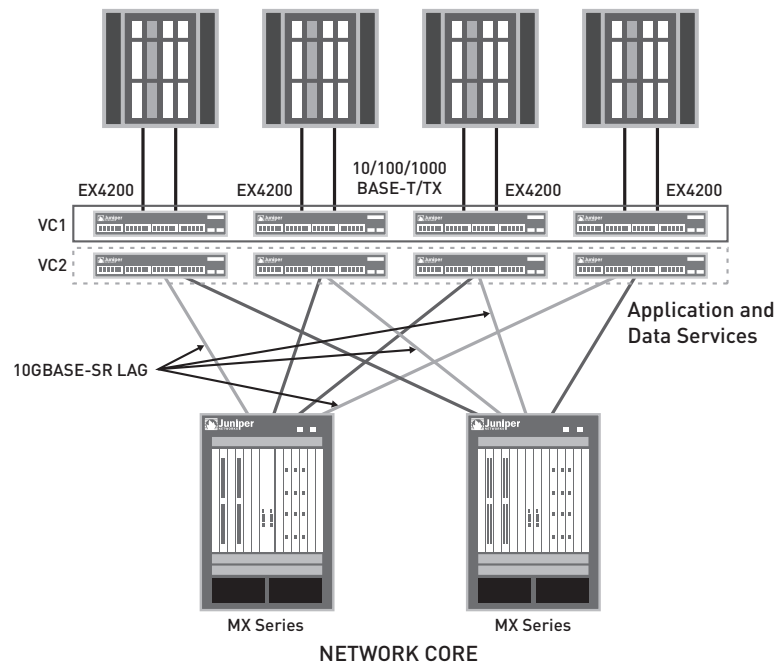


Figure 10: Application and data services network view

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. The Virtual Chassis technology in the Juniper Networks EX4200 Ethernet Switch offers several advantages when deployed as a top-of-rack access switch. The EX4200 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 EX4200 offers additional wire-rate uplink interfaces, with a maximum of four Gigabit Ethernet or two 10 Gigabit Ethernet uplink modules for interconnecting from the top rack back to the data center core.

The EX4200 also supports the virtual chassis concept, whereby a maximum of 10 EX4200 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 10, significantly simplifying operations and reducing costs associated with maintaining large numbers of legacy access switches. Additionally, uplinks can be distributed across multiple EX4200 switches in a single Virtual Chassis, providing uplink performance flexibility and added redundancy levels not otherwise found in legacy switches.

Storage Area Networks (SANs)

A 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 each other. 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 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 is 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. However, 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 Gigabit 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, random early detection (RED), 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 Business Continuity/Disaster Recovery (BC/DR) 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 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 Juniper Networks WX Series 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, such as data mirroring, to ensure that accurate data is represented at multiple data centers. The functions that rely on a high-performance data center backbone include:

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

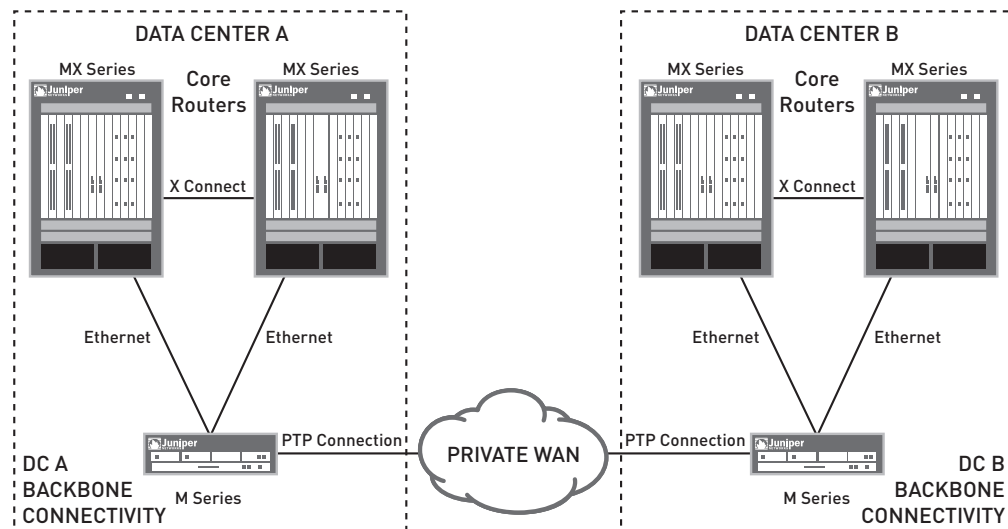


Figure 11: Data center backbone connectivity

Different from any other WAN interconnect, the data center backbone must offer high-speed connectivity because of the real-time and consistent representations of data that need 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 several key technical elements as enablers that network architects can leverage to satisfy organizational requirements for the data center backbone. First is the connectivity links encompassing high-speed 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 most effectively through MPLS technologies. In addition to the interconnect elements, availability protocols 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 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, L2 extensions and technologies can be used so that non-IP or broadcast domain dependent/attached protocols are 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 other than IP (HTTP, SIP, CIFS, FTP and so on), typically through the 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.

BGP multihoming is also important. The first 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 to 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
- GSLB

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. Government agencies are increasingly adopting best practices as recommended by the Information Technology Infrastructure Library (ITIL) and are 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 agency management framework with minimal or no retraining of network operations center/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 Connect (TNC) and the Internet Engineering Task Force (IETF), to ensure smooth interoperability with existing and future management systems.

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

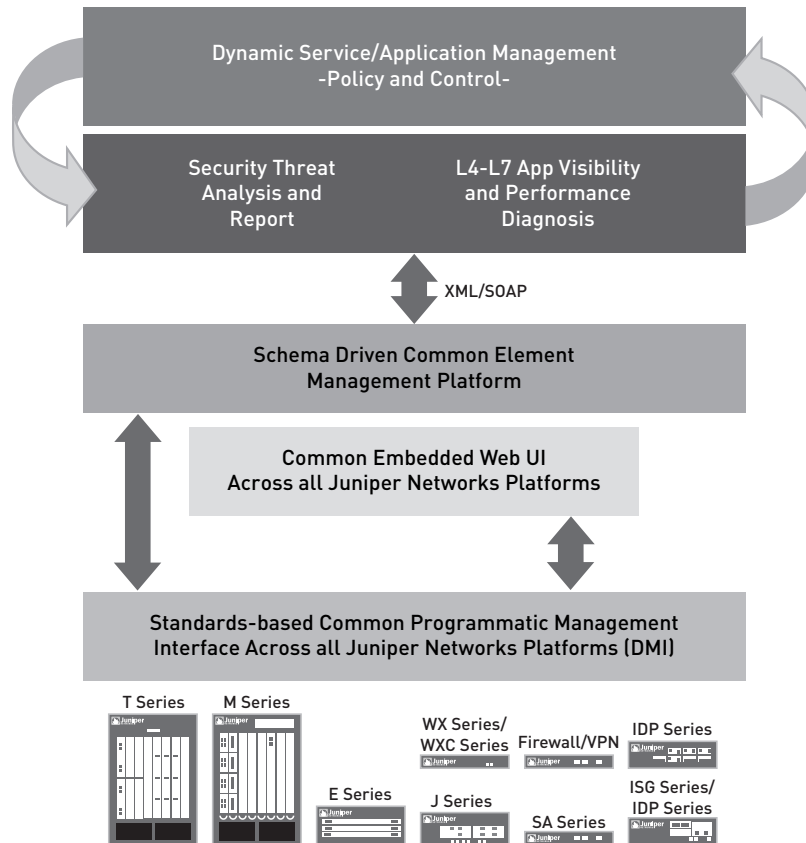


Figure 12: Network management framework built on Juniper Networks products

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, allowing existing network management systems to quickly discover and adapt to the new Juniper Networks infrastructure.
- Centralized Management – 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.

Conclusion

Juniper Networks offers an advanced data center network architecture that consolidates and simplifies the management and administration of government data center network infrastructures to deliver services throughout the distributed agency network. It is enabled by an open systems approach that supports devices and elements to achieve a more efficient, secure and cost-effective network infrastructure. This powerful solution greatly simplifies the design and enables operational efficiencies by deploying networks that are agnostic to multiple media types.

The Juniper Networks architecture virtualizes critical network infrastructure components and functionalities—for example, security, load balancing and applications acceleration, deployed and managed using a combination of organizational and technical heuristics. It also optimizes network performance and increases efficiencies of the network infrastructure. Finally, it automates network infrastructure management by plugging smoothly into existing agency management frameworks and third-party tools such as IBM Tivoli.

Appendix A: Juniper Networks Data Center Network Solution Tables

Table 1: Data Center Products Tables

Infrastructure			Services	Policy and Management
Routing	Switching	Security/VPN	Secure Access	Policy and Management
MX960	EX3200	ISG2000*	SA6000	IC6000
M320	EX4200	ISG1000*		NSM
M120		NetScreen-5400		WX CMS
M10i		NetScreen-5200		OAC
		IC6000		SBR Series
		IC4000		

Table 2: Data Center Product Tables (by Tier)

	Edge	Core	Network Services
Routing	M320 M120 M10i	MX960	—
Switching	—	EX3200 EX4200	EX3200 EX4200
Firewall	—	—	NetScreen-5400 NetScreen-5200 IC6000 IC4000
SSL VPN	—	—	SA6500
WAN optimization	—	—	WXC500 Stack
Policy and management	—	—	IC6000 NSM WX CMS OAC SBR Series

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.

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 a government agency in four fundamental ways:

- **Security Threats:** Viruses, spyware and other malware can all enter the 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 the network has legal liability issues as well.
- **Productivity Threats:** The temptations of non-work-related Web destinations are endless. Just 20 minutes of recreational surfing a day can cost an agency with 500 employees more than \$8,000 per week (at \$50/hour/employee).
- **Network Threats:** Employees can crash the network just by logging in to the wrong Web site. Other activities, such as recreational surfing and downloading MP3 files, can divert valuable bandwidth from critical agency 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 agencies to build Web access policies from the largest URL database (more than six million pages) spread across more than 40 categories. From the Web UI or Juniper Networks Network and Security Manager, 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 hosted server running the Web filtering software where Web access grant/deny decisions are made and executed. The agency 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. Agencies can consolidate the number of devices that they deploy and manage in remote sites. 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 the 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 Juniper Networks J4350 Services Router or Juniper Networks J6350 Services Router and delivers a rich telephony feature set to the branch office. This feature set includes:

- Central Avaya Communication Manager and other communications applications
- Call center agent support
- 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

Table 3: Appendix B

Characteristics	Juniper Networks Core MX960 2x Chassis
Line-rate 10 Gigabit Ethernet (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 Chasses
Rack space (racks)	2/3rds of a Single Rack

About Juniper Networks

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

Corporate and Sales Headquarters

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER
(888.586.4737)
or 408.745.2000
Fax: 408.745.2100

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
Fax: 35.31.8903.601

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

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

