

DATA CENTER LAN CONNECTIVITY DESIGN GUIDE

Design Considerations for the High-performance
Enterprise Data Center LAN

Table of Contents

Executive Summary	5
Introduction	5
Trends and Challenges	6
Centralization of Data Centers	6
Server Consolidation	7
Virtualization	7
Storage	7
Service Oriented Architecture (SOA)	7
Software as a Service (SaaS)	8
An Increasingly Decentralized Workforce	8
Green and Environmentally Friendly Data Centers	8
The Proliferation of Unified Communications	8
Increasing Focus on Security	8
Data Center Network Design Considerations	9
Services Required in the Data Center	9
High Availability (HA)	9
Visibility	10
Network Connectivity	10
Security	10
Policy and Control	11
Quality of Service (QoS)	11
High Performance	11
Juniper Network Design Approach	11
Data Center Architecture Overview	13
Layered Approach	13
Benefits	13
Challenges	13
A Network Revolution	14
Data Center Access Layer	14
Access Layer Design Considerations	15
Application and Server Architectures	15
Benefits and Challenges of the Three-Tier Model	15
Server Virtualization	16
Connectivity	16
Power over Ethernet (PoE)	16
High Availability (HA)	17
VLAN and Spanning Tree Protocol (STP)	17
Using Layer 2 versus Layer 3 at the Access Layer	18
Physical Deployment: Top-of-Rack vs. Middle of Row/End-of-Row	19
Storage Connectivity	19
Quality of Service (QoS)	20
Data Center Access Layer Design Recommendations	20
Scalable Configuration with Virtual Chassis Technology	20
EX4500 10GbE Switch	23
Modular Chassis Configurations	23
Data Center Aggregation Layer	24

Aggregation Layer Design Considerations	25
High Availability (HA)	26
Scalability	26
Network Virtualization	26
Application Visibility	26
Security and Threat Containment	26
Data Center Aggregation Layer Design Recommendations	26
Traditional Layered Approach	26
Collapsing the Aggregation Layer into the Core Layer	27
Data Center Core Layer	27
Data Center Core Design Considerations	28
High Availability (HA)	28
Data Center Core Layer Design Recommendations	29
Consolidating the Aggregation Layer and the Core Layer	29
WAN Edge Integration	31
WAN Edge Design Considerations	32
Connectivity	32
High Availability (HA)	32
Firewall/VPN	32
WAN Edge Layer Design Recommendations	32
M Series Routing Platform	32
Operational Simplicity and Unified Management	33
Achieving Operational Simplicity with Junos OS	33
The Power of Junos OS	34
Modular Processes	34
Rollback Capability	34
Advanced Features	34
Benefits	34
Impact	35
Unified Management with Juniper Networks Network and Security Manager (NSM)	35
Benefits	35
Remote Configuration and Management with J-Web	35
Benefits	35
Conclusion	36
About Juniper Networks	37

Table of Figures

Figure 1: The data center LAN in the enterprise network	6
Figure 2: Data center LAN functional design model.	9
Figure 3: Highly available data center LAN configuration.	12
Figure 4: The layered approach	13
Figure 5: Access layer of a highly available data center LAN.	14
Figure 6: The three-tier application model.	15
Figure 7: Virtualized server infrastructure	16
Figure 8: Layer 2 versus Layer 3 at access layer.	18
Figure 9: Top-of-rack vs. end-of-row switch deployments	19
Figure 10: Virtual Chassis technology	21
Figure 11: Top-of-rack deployment using Virtual Chassis technology.	22
Figure 12: End-of-row deployment using Virtual Chassis technology.	22
Figure 13: EX8200 line of modular chassis solutions.	23
Figure 14: End-of-row deployment using fixed chassis technology.	24
Figure 15: Aggregation layer in a highly available data center LAN	25
Figure 16: Core layer in a highly available data center LAN	28
Figure 17: Aggregation layer collapsed into the core layer in a highly available data center LAN.	30
Figure 18: WAN edge in a highly available data center LAN.	31
Figure 19: Junos OS—The three ones: one source code, one train, and one modular architecture.	34
Figure 20: Juniper switching solutions	36

Executive Summary

The data center LAN is a critical corporate asset, connecting servers, applications and storage services in the enterprise. This strategic tool supports vital day-to-day operations and is crucial for corporate success. The data center LAN faces a number of challenges as enterprises are centralizing applications and consolidating servers to simplify operations and reduce costs while business productivity increasingly depends on operations carried out at distributed branch offices. As businesses continue to expand across the globe, downtime is not an option—a data center LAN must efficiently operate 24x7.

These trends raise the density, scalability, throughput and high availability (HA) requirements of the data center LAN. Trying to support these needs with low-density, single-function legacy equipment is not only inefficient, it's not cost effective, adversely affecting performance, reliability, valuable rack and cabinet space as well as driving power and cooling costs higher. Enterprises are also moving towards applications that use a Service-Oriented Architecture (SOA) and also provide Software as a Service (SaaS), both of which present a new set of throughput, performance and HA requirements for the data center LAN. New technologies such as virtualization are needed to increase scalability, efficiency and lower total cost of ownership.

These changes, coupled with IT initiatives such as Unified Communications, require that data center LANs operate with the same carrier-class reliability and performance demanded by fee-based service providers. Existing data center infrastructure solutions cannot meet these requirements, nor do they provide the unified management capabilities critical for reducing costs and streamlining operations.

Simply designing a data center that only deploys more servers, more storage, and more devices significantly increases network complexity and cost. Legacy solutions are inefficient; for example, more than 50 percent of Ethernet switch ports within the data center are typically used for switch interconnectivity. A new data center LAN design that meets the growing performance demands of users and network-centric applications from a variety of locations is needed. It also must economically scale and flexibly accommodate new computing trends and IT initiatives without an entire redesign.

This document introduces the issues related to changing data center needs and also presents design considerations and recommendations for data center LANs. In addition, it shows how infrastructure solutions from Juniper Networks® advance the economics of networking, allowing businesses to “change the rules” with their IT investments and create a truly innovative and competitive environment that helps them increase revenue and raise productivity today and in the future.

Introduction

Data centers contain centralized computing resources vital to all employees in the enterprise, be they at headquarters, a large regional office, a remote branch office, a home office or at a customer site. As most critical business processes are carried out online, any data center LAN downtime or inefficiency has a negative impact on business processes and the corporate bottom line. The data center LAN must provide secure, high-performance, highly-available LAN services at scale to ensure that the network is always online and that the necessary resources are always available to maximize business productivity and customer satisfaction.

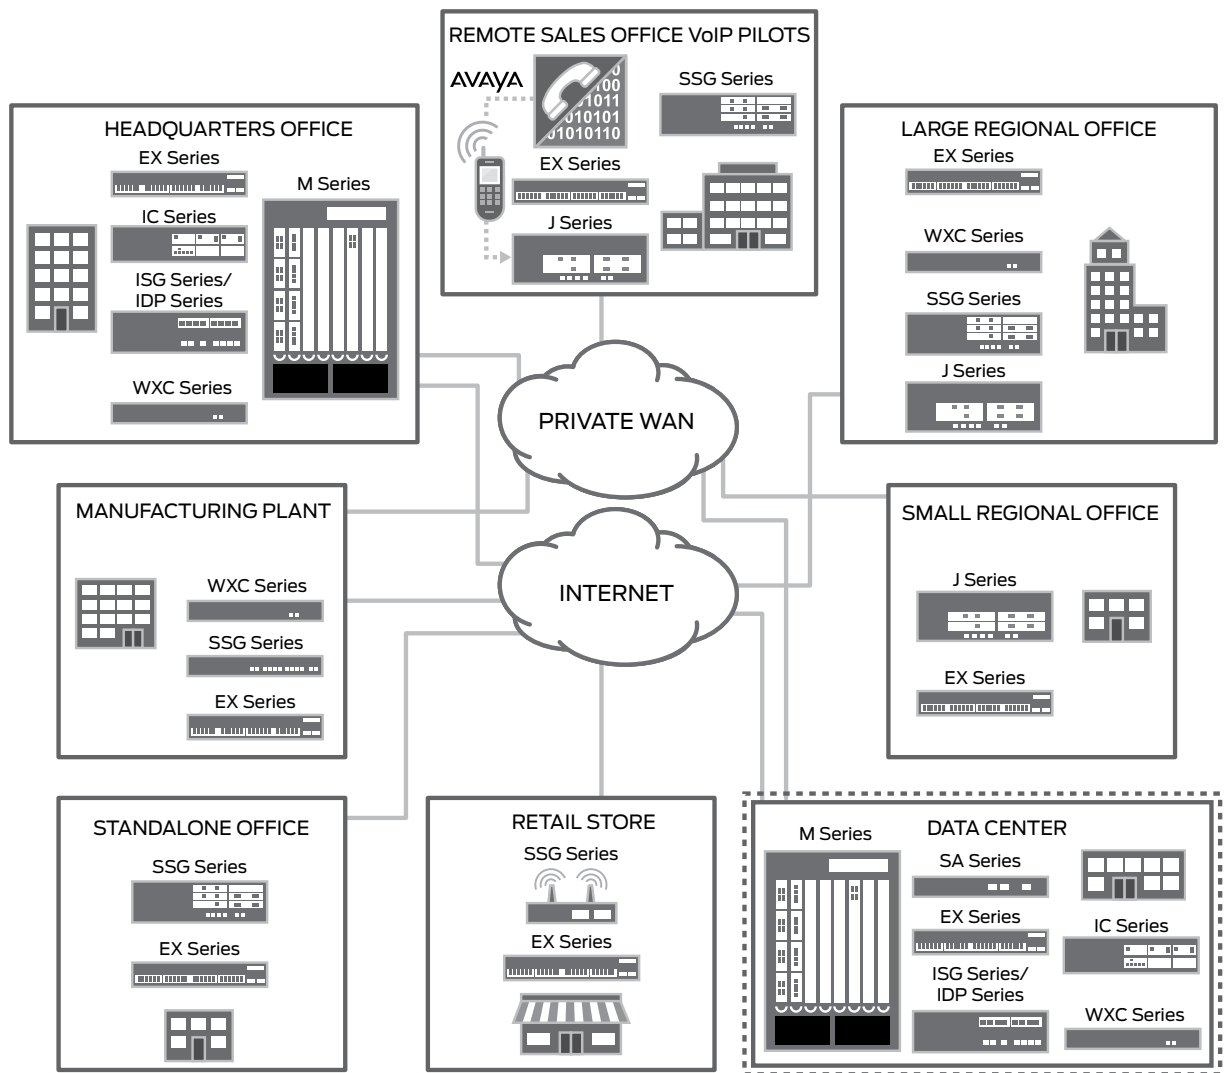


Figure 1: The data center LAN in the enterprise network

Trends and Challenges

In addition to the requirements previously mentioned, the following trends must be considered in a data center LAN design:

Centralization of Data Centers

To reduce costs, simplify operations and comply with regulatory guidelines, more and more enterprises are consolidating their data centers. According to a 2006 Nemertes Research report¹, 91 percent of companies interviewed were under compliance constraints and more than 50 percent of the companies consolidated their dispersed data centers into fewer larger data centers in the past 12 months, with even more planning to consolidate in the following 12 months. In addition to HA requirements ensuring nonstop operations, centralization raises new latency and security issues for the data center LAN.

Server Consolidation

Gartner (2007) asserted that servers are growing at an annual rate of 11 percent and that storage is increasing at 22 percent, both causing tremendous strain on the data center's power and cooling capacities. A 2007 Forrester report² states that 51 percent of all firms consider server centralization a key priority. Gartner also reports that most enterprise servers operate at 20 percent capacity; new technologies like virtualization are needed to better utilize these resources. Additionally, backup and security concerns must be addressed, and companies also demand consolidated, centralized management solutions that help reduce the time and resources devoted to keeping data centers online and operational.

Virtualization

Virtualization, a technology used to share resources, makes single physical resources appear as many individually separate resources. Conversely it also makes individually separate physical resources appear as one unified resource. Virtualization can also include making one physical resource to appear, with somewhat different characteristics, as one logical resource. The benefits of virtualization are in creating more complex systems with minimal effort. It takes advantage of commodity hardware to build modular systems that easily scale and accommodate consolidation, advanced automation, security and ease of management. It is used on four main resource categories: servers, storage, networks, and end-user desktops.

Server virtualization allows a single server using software such as VMware® or Microsoft Virtual Server to appear as many machines. Ideal for underused application servers such as Web servers, this technology is not as suitable for processor-intensive applications such as database servers. Server virtualization enables IT to flexibly manage workload and also provides basic HA and disaster recovery services.

Storage virtualization helps make many storage arrays and pools and systems appear as a single resource, providing for seamless scaling, easier migration, improved resource utilization and simplified management.

Virtualizing a network is enabled by various technologies that provide data-plane virtualization, control-plane virtualization and management-plane virtualization. An example of data-plane virtualization is using a using 802.1q VLAN tagging on single physical network interface to provide security to multiple network segments. Supporting multiple routing domains and protocol instances on a single router using Virtual Routers and/or VRF are examples of control-plane virtualization. Support for multiple logical firewall/VPN security systems using Virtual Systems (VSYS) in a single device is a management-plane virtualization example. Virtualization delivered via MPLS and VPLS also enable an ultra fast data center backbone network in order to meet the performance demands of the consolidated LAN architecture. Virtualization can enable multiple switches to act as one, simplifying device configuration and management while also increasing reliability and reducing potential choke points.

Client virtualization enables IT to provide instant and ubiquitous access to hosted desktops. Ideal for remote users or non-employees, such hosted corporate machines are fully secure and simple to manage and upgrade.

Storage

As businesses increasingly rely on vast stores of data to make business decisions and meet compliance regulations, scalable, high-performance storage solutions are becoming a necessity for today's enterprise. Fibre Channel still maintains a large portion of the SAN market, but the growing prevalence of gigabit Ethernet (GbE) and the simplicity of deploying and managing an Ethernet-based Network Attached Storage (NAS) are making iSCSI an attractive, low-cost alternative. Additionally, Ethernet-based NAS solutions more easily take advantage of virtualization to rapidly scale and provide HA. While 4 or 8 Gbps Fibre Channel offers a speed advantage over GbE, Network Interface Cards (NICs) offering TCP Offload capabilities greatly enhance iSCSI performance. In addition, the emergence and adoption of lower-cost 10 GbE allows iSCSI to outperform Fibre Channel and accommodate any high-speed storage needs.

Service Oriented Architecture (SOA)

Emerging enterprise applications are increasingly using a Service-Oriented Architecture (SOA) to unify business processes by structuring large applications as a collection of smaller independent modules called services. In this manner, IT can leverage key processes or technology assets across applications. In an SOA-based environment, services exchange messages to interoperate, in some instances generating millions of messages each, which can

impact LAN bandwidth needs. Web services are often used to implement SOA and provide ubiquitous access to the applications. Web services put extra processing demands on servers while also increasing network bandwidth requirements as Web-based applications use far more bandwidth than client-server applications. Virtualization is often used in SOA environments to increase the reliability of services and help scale capacity. SOA also broadens application access to internal and external users, raising security concerns. Additional security issues are raised as application services expose capabilities to other applications which require a different level of security.

Software as a Service (SaaS)

Many common enterprise applications, such as customer-relationship management (CRM), human-resource management (HRM) and supply-chain management (SCM), can now be delivered in the Software as a Service (SaaS) model. Many of these Web-based services require, in certain instances, more than 10 times the bandwidth of their LAN-based counterparts, seriously impacting performance, reliability, availability and bandwidth requirements.

An Increasingly Decentralized Workforce

The corporate data center LAN design needs to accommodate the delivery of HA, high-performance services to the estimated 89 percent of employees who work outside of headquarters in remote or branch offices (Nemertes Research 2006). As employees in remote or branch offices become increasingly dispersed across different time zones, HA time requirements also increase. In addition, virtualized operations have expanded enterprise user populations beyond employees to include contractors, consultants, business partners and customers who may be anywhere in the world. As a result, enterprises need to provide their end users with ubiquitous, secure connectivity while ensuring all corporate resources and applications are secure.

Green and Environmentally Friendly Data Centers

As old data center facilities are upgraded and new data centers are built, it is important to ensure that the data center network infrastructure is designed for maximum energy and space efficiency as well as a minimal environmental impact. Power, space and cooling requirements of all network components must be accounted for and compared with different architectures and systems so that the environmental and cost impacts across the entire data center as a whole can be ascertained—even down to the lighting. Many times, it might be more efficient to implement high-end, highly scalable systems that can replace a large number of smaller components, thereby delivering energy and space efficiency. Green initiatives that track resource usage, carbon emissions, efficient utilization of resources such as power and cooling are to be considered when designing a data center.

The Proliferation of Unified Communications

The adoption of Unified Communications systems that combine voice, video and data services is on the rise. According to Forrester Research (2006), 46 percent of all companies in North America have installed IP telephony systems and 39 percent use VoIP to communicate with their remote employees. Such deployments have a direct impact on the high-performance and HA requirements of a data center LAN. For example, not only must adequate LAN and WAN bandwidth be provisioned, but quality of service (QoS) rules must identify, classify and prioritize traffic to deliver effective VoIP communication services.

Increasing Focus on Security

FBI/CSI statistics show that 72 percent of all companies surveyed reported at least one security incident in 2006. Not surprisingly, a 2006 Forrester Research survey found that 57 percent of all firms consider “upgrading security environment” a top priority. As employees and non-employees are being granted an ever-widening range of network access, robust security is necessary at all levels in the corporate and data center LANs. IT must protect applications, data and infrastructure by applying appropriate access controls without inhibiting user efficiency or negatively impacting application performance. IT must also mitigate risks from untrusted sources such as non-employees, whose PCs and networks are not under IT control. The move to globalize and virtualize the enterprise puts new demands on IT to secure remote access communications and protect site-to-site communications, including connections between data centers and from data centers to backup sites. IT must also fortify the network perimeter as increasing volumes of Web and other traffic types flow across it.

Data Center Network Design Considerations

A new data center LAN design is needed as legacy solutions cannot meet these key requirements, nor reduce costs and streamline operations. The LAN design must also scale and accommodate emerging computing trends and additional network services without an entire redesign. The new design should be architected in order to maximize efficiency gains from technologies like virtualization.

Services Required in the Data Center

The following high-level services are required of data centers to provide carrier-class network service throughout the enterprise and thus optimize efficient business operations. Each of these areas is addressed in more detail in this document and, where appropriate, additional considerations or challenges for a specific service, feature or data center category are presented.




Figure 2: Data center LAN functional design model

High Availability (HA)

With the consolidation and centralization of servers and resources, HA is a key requirement from the data center LAN. Redundancy of critical subsystems and seamless failover are needed for routers, security appliances, and any other devices on the user-to-data center path. Designing HA into the data center network requires consideration of three key aspects — device availability, network availability and operational availability.

Table 1: The Three Aspects of Designing HA Into the Enterprise Network

DEVICE AVAILABILITY	NETWORK AVAILABILITY	OPERATIONAL AVAILABILITY
<ul style="list-style-type: none"> • Redundant components • Hot-swappable components • Modular operating system software • In service software upgrades 	<ul style="list-style-type: none"> • Network access control • Redundant devices and paths • Routed network designs • Quality of service 	<ul style="list-style-type: none"> • Open standards • Consistent software features • Automate operational tasks • Reduce complexity

Network devices deployed within the data center should support device-level HA with components such as redundant power supplies, fans and route engines. The operating system software running on data center network devices should have a modular architecture so that software failures will be isolated to a single process and not impact other critical operating system services, ensuring system and network availability. Features such as in-service software updates (ISSU) also maintain network availability while still providing network software updates.

Network availability should be enabled by using combinations of redundant devices and path (for both external and internal connectivity) and critical device redundancy to ensure network operations and business continuity.

Operational availability denotes a set of network operating system attributes that ensure simple and efficient operation of the data center network. Network devices must support open management standards and consistent software features for simple, error-free configuration that maintains network availability. Also, network devices should support scripting to enable automation of operational tasks that free resources for other, more critical tasks.

Visibility

Visibility into network traffic and security events is important in order to effectively maintain and manage network resources. Real-time and historical reporting enables IT to maximize performance and availability across the entire data center infrastructure, meet regulatory requirements, and plan for future capabilities and capacity. Collecting IP traffic flow statistics can give enterprises valuable insight into areas such as 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 operational and performance patterns that can impact bottom-line productivity. For example, compression and acceleration technologies can be applied at the network layer to accelerate email applications, or application-based policies can ensure that business critical applications meet or exceed performance requirements when other non-essential bandwidth hungry services like YouTube are accessed.

Network Connectivity

Customers, partners and employees all require fast access to applications and information. Connectivity has to be absolutely reliable, consistent and provide low latency. Modern applications, especially those provided as a Web service, demand significant network performance. At the same time, the challenge of working from any location in or out of the enterprise further increases complexity. The following critical aspects of external network connectivity need to be considered as part of the data center network design:

- High-speed (10 GbE) LAN connectivity for servers and storage devices
- WAN connectivity to enable branch office and campus users to access applications and shared resources
- Internet connectivity to enable partner access as well as secure remote access for remote and mobile users
- Super-fast data center backbone connectivity for purposes of data replication and business continuity and use of technologies like VPLS/MPLS

The data center LAN hosts a large number of servers that require high speed and highly available network connectivity. Multiple LAN segments and networks may be deployed with differing levels of security, capacity and other services. Local server connections of one gigabit per second or greater for local servers, with a forward view towards the proliferation of 10 GbE, and also utilizing 10 GbE for connecting to upstream or downstream devices should be a consideration.

Security

Security is critical to the entire corporate LAN and especially to the data center LAN. Access to centralized networks and applications must be ubiquitous and pervasive, yet remain secure and controlled. The security design must employ layers of protection from the network edge, through the core, and both in front of and between the application computing systems, providing in-depth defense. The protection must be integrated into the network operating system and not simply layered on top. A tiered, integrated security solution protects critical network resources that reside on the network. If one tier fails, the next tier will stop the attack and/or limit the damages that may occur. This allows an IT department to apply the appropriate level of resource protection to the various network entry points based upon their different security, performance, and management requirements.

Today's data center networks needs not only to effectively handle unmanaged devices and guest users attempting network access; they also need to support unmanageable devices, post admission control, and application access control, visibility and monitoring. In addition to Unified Threat Management (UTM) services, security policies supporting demilitarized zones (DMZs), ensuring quality of service, mitigating Denial of Service (DoS) and distributed DoS (DDoS) attacks and threats, and ensuring that the organization meets compliance criteria are needed. All security policies should be centrally managed and remotely deployed.

Policy and Control

Policy-based networking is a powerful concept that enables efficient management of devices in the network, especially within virtualized configurations, and 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 offer distributed and even layered 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. The data center network infrastructure also should easily integrate into customers' existing management frameworks and third-party tools such as IBM Tivoli and HP software and also provide best-in-class centralized management, monitoring and reporting services for network services and infrastructure.

Quality of Service (QoS)

For optimal network performance, QoS is a key requirement. QoS levels must be properly assigned and managed to ensure satisfactory performance for various applications through the data center and across the entire LAN. A minimum of six levels of QoS are recommended, each of the following determines a priority for application of resources:

- Gold Application Priority
- Silver Application Priority
- Bronze Application Priority
- Voice
- Video
- Control Plane

In MPLS networks, network traffic engineering capabilities are typically deployed to allow configuration of Label Switched Paths (LSP) with the Resource Reservation Protocol (RSVP), LDP, or BGP. 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 in anticipation of performance problems or failures. The LAN design should allow the flexibility to assign multiple QoS levels based upon end-to-end assessment and allow rapid and efficient management to ensure end-to-end QoS throughout the enterprise.

High Performance

To effectively address performance requirements related to virtualization, server centralization and data center consolidation, the data center network must offer high-capacity throughput and processing power with minimal latency. The data center LAN also must boost the performance of all application traffic, be it local or remote. The data center must offer a LAN-like user experience for all enterprise users regardless of their physical location. In order to accomplish this, the data center network must enable optimization for applications, servers, storage and network performance.

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

Beyond WAN optimization, critical infrastructure components such as routers, switches, firewalls, remote access platforms and other security devices must be built on non-blocking modular architecture. This ensures that they have the performance characteristics necessary to handle the higher volumes of mixed traffic types associated with centralization and consolidation, as well as the needs of users operating around the globe.

Juniper Network Design Approach

The network infrastructure in today's data center is no longer sufficient to satisfy these requirements. Instead of adding costly layers of legacy equipment and highly skilled IT resources to support the growing number of single-function, low-density devices and services in the enterprise, a new, more integrated and consolidated data center solution is needed. High-density, multifunction devices are needed in the new data center LAN. Such devices can help collapse costly latency-inducing layers, increase performance, decrease logical and physical cabling complexities, decrease choke points, decrease configuration and management tasks and increase reliability—all while decreasing TCO as well as ongoing rack and floor space, power, and cooling costs.

Juniper Networks delivers a proven IP infrastructure for the data center that meets these challenges, enabling the performance, scalability, flexibility, security and intelligence needed to not just meet but increase branch-office user productivity. Juniper offers flexible configurations and price points that meet the needs of all data centers while delivering high-performance throughput with services such as firewall, UTM, VPN, MPLS, IPV6 and CLNS-enabled.

Juniper provides an open systems approach that enables enterprises to design a high performance data center network that consolidates network elements into a single IP network and employs fewer network devices and fewer layers. This greatly simplifies the network architecture, and enables operational efficiencies and creates better data center networks.

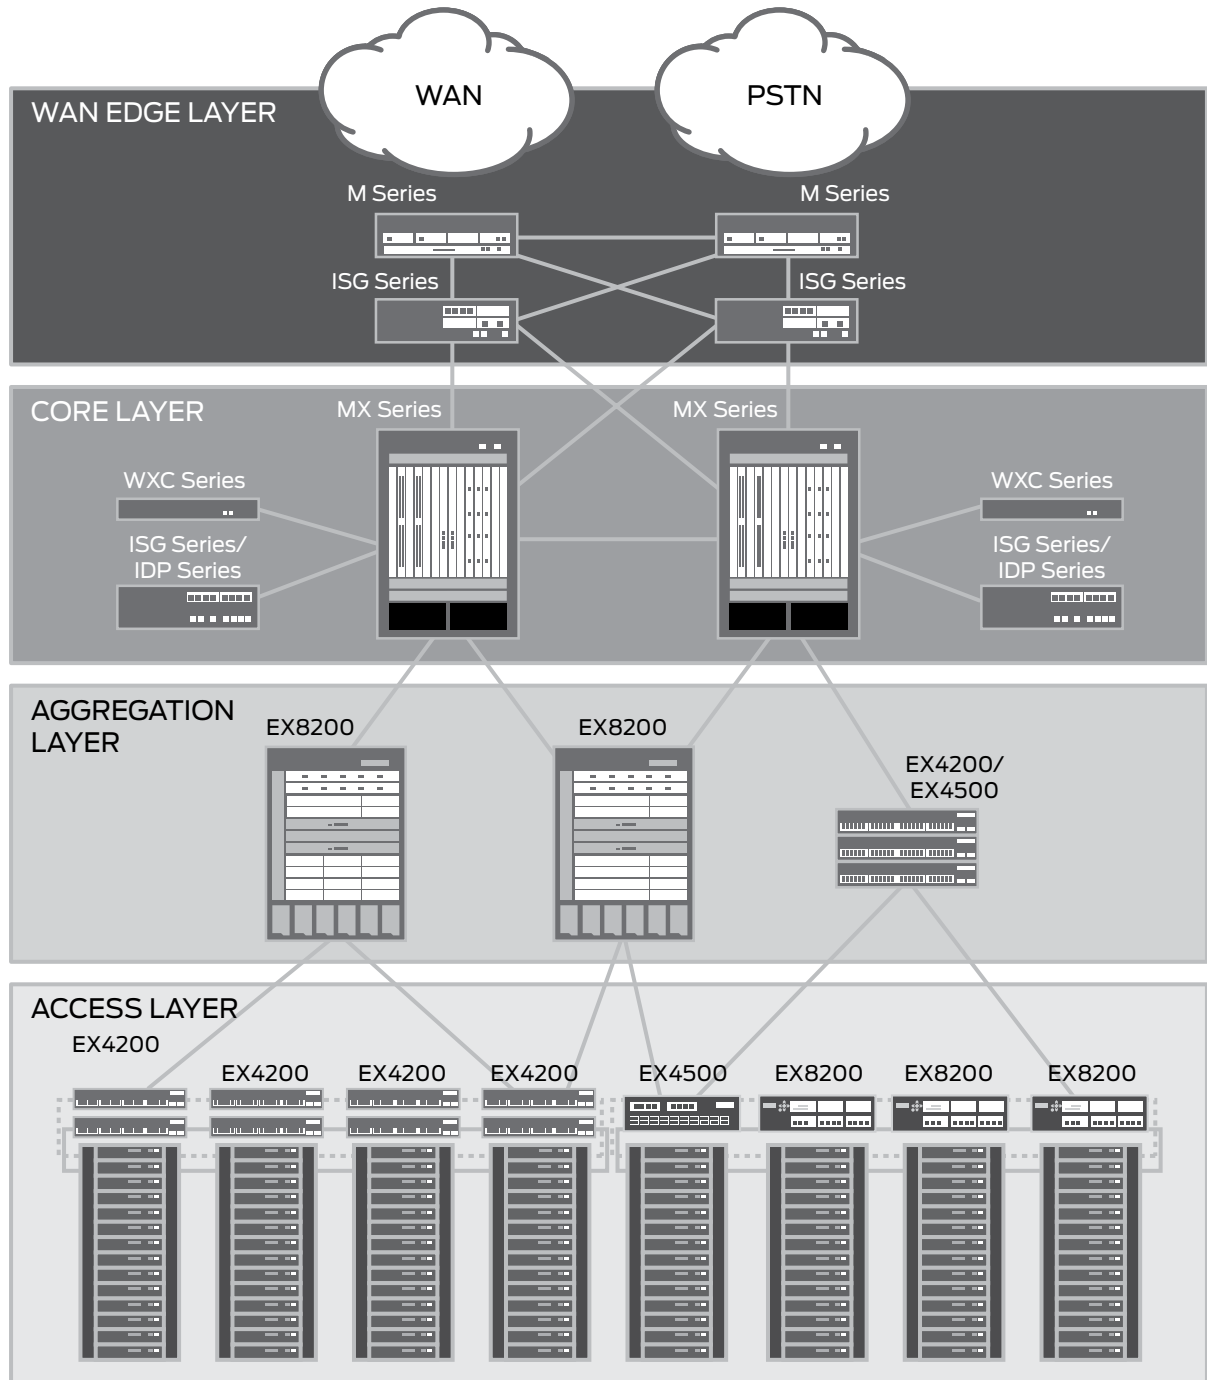


Figure 3: Highly available data center LAN configuration

Data Center Architecture Overview

Layered Approach

The typical enterprise network is built upon multiple levels of switches deployed in three general layers: access, aggregation and LAN core.

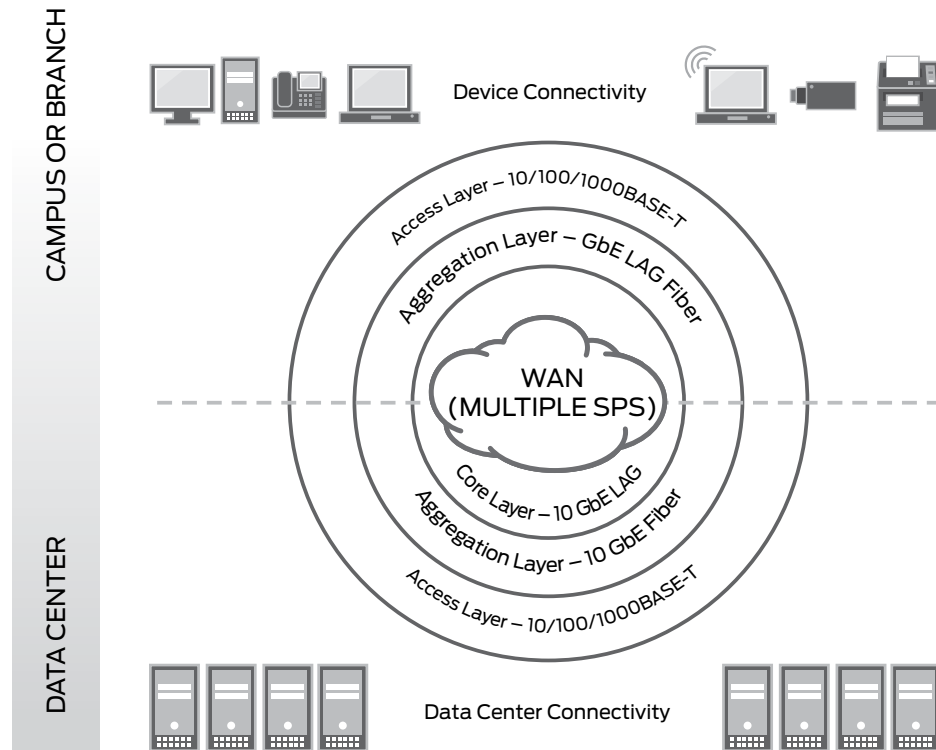


Figure 4: The layered approach

Providing vital LAN services, these layers exist at various locations throughout the network, including data centers, campus buildings and the data center. This document focuses primarily on the layers deployed in the data center. Areas outside of that scope are presented when relevant to the discussion. For example, some data centers may choose to collapse the aggregation layer into the core.

The access layer provides connectivity to the servers, applications, storage devices, and any IP or office automation devices required in the data center facility.

The aggregation layer aggregates connections and traffic flows from multiple access-layer switches to core-layer switches.

The core layer provides secure connectivity between aggregation-layer switches and the routers connecting to the WAN.

The WAN edge provides connectivity to the Internet and the WAN to enable remote connectivity.

Benefits

A multilayered architecture facilitates network configuration by providing a modular design that can rapidly and economically scale. It also creates a flexible network on which new services can be easily added without redesign. The layered approach also delivers separated traffic, balances load across devices and simplifies troubleshooting.

Challenges

Over the years, networks have grown bloated trying to address emerging bandwidth, throughput and port density requirements by deploying multiple layers of low density, single-function legacy hardware, many of which are redundant. These old solutions not only fail to meet the current data center requirements, but also add considerable management complexity, reduce network availability, and drive up capital and operational expenses.

A Network Revolution

Typically over 50 percent of Ethernet switch ports are used for switch-to-switch connectivity in the data center. High-density switches that eliminate layers reduce server-to-server latency by 50 percent, decrease bandwidth chokepoints and increase bandwidth capacity by 75 percent, require 50 percent less power with smaller thermal and physical footprints, increase growth capacity, and also simplify network cabling, topology and device management.

As a recent entrant into the switching market, Juniper Networks has factored lessons learned and other experiences into the development of a new portfolio of high-density Ethernet switch products and network solution designs that address contemporary issues and accommodate the future growth of high-performance networks. These new products are designed to eliminate unnecessary network layers while providing a platform for delivering higher availability, converged communications, integrated security and higher operational efficiency. With these solutions, Juniper Networks simultaneously advances the fundamentals and economics of networking by delivering greater value, increasing simplicity and lowering the total cost of network ownership.

Data Center Access Layer

The access layer provides connectivity to all of shared enterprise servers, applications, storage devices, and any IP or office automation devices required in the data center facility. Most data center access switches are deployed at the top of the rack or at the middle/end of row of server racks, with a minority deployed in the wiring closet of the data center facility, which supports local connectivity needs.

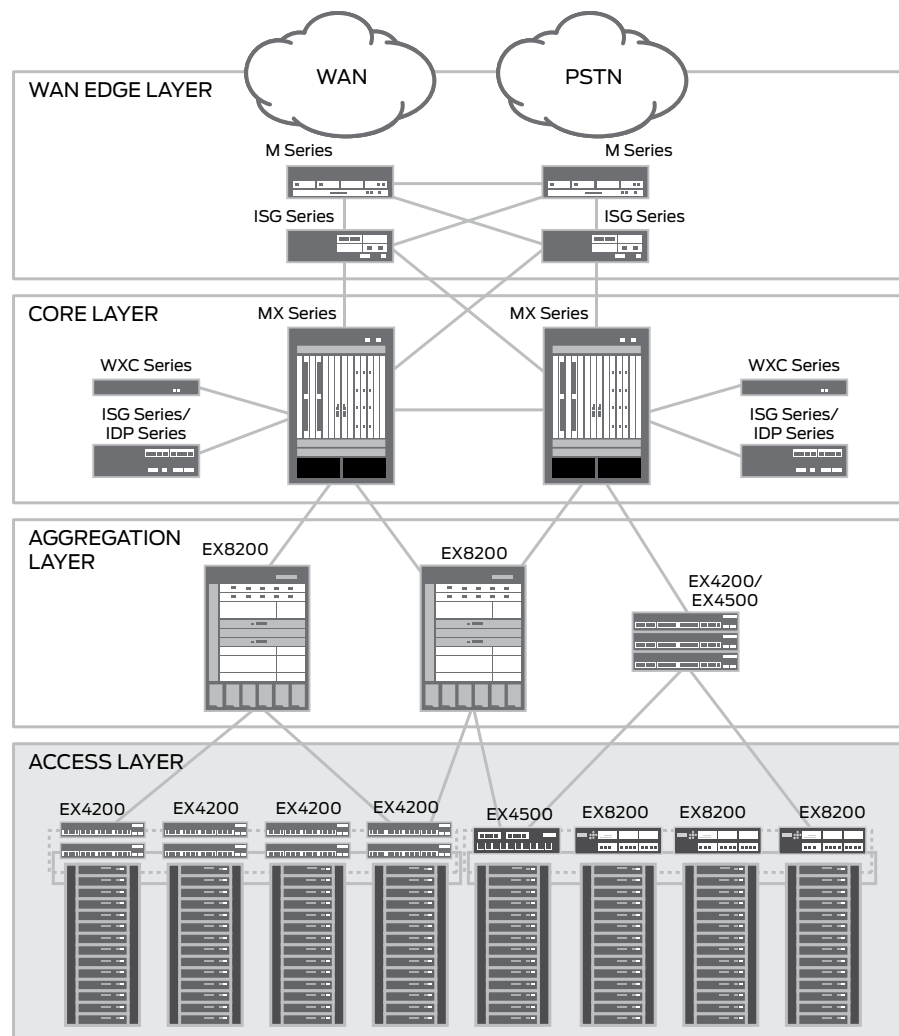


Figure 5: Access layer of a highly available data center LAN

Access Layer Design Considerations

Application and Server Architectures

Another way to look at the access requirements of the data center is via the common three-tier application model upon which a majority of Web-based applications are built. It defines application architectures in the following modular components:

1. Web
2. Application
3. Database




Figure 6: The three-tier application model

Today, most Web-based applications are built upon this model. This model runs separate processes on the same machine or across different networked servers. While Web servers and application servers may share the same machine or set of servers, it is common to separate the database on a separate machine or set of servers dedicated to that task.

Benefits and Challenges of the Three-Tier Model

When server farms are used, this model provides built-in HA because any individual server can be taken out of service without disrupting service since the same function runs on another server belonging to the same application tier. In that same manner, additional machines can be added to seamlessly scale capacity as needed. Load balancing the traffic between tiers improves performance and HA. Security is built in as attacks on one server are insulated from others. For example, a hacked Web server compromises only that server without gaining access to the application or database servers. Security can be further enhanced by placing firewalls between tiers of servers or virtualizing a high-end firewall to inspect traffic between the layers and enforce security policies. VLANs can also increase security by segmenting traffic and reduce the server farm complexity. For increased performance and security, physical segregation may be desired.

There are a few disadvantages to the three-tier application model. This model does not work as well as other topologies for computational-intensive applications such as financial modeling, animation, manufacturing and search engines. Another disadvantage is that often complex traffic engineering is required to optimize performance. Finally, the TCO of this architecture can be high due to inefficient use of physical server infrastructure with high power, cooling and space requirements.

Server Virtualization

Server virtualization capabilities such as those delivered by Microsoft Virtual Server or VMware Infrastructure are increasingly being deployed to increase the operational efficiency of server infrastructure and in turn lower power, cooling and space requirements. While delivering operational efficiency, the virtualized infrastructure places new demands on the access layer of the data center. A high-performance network infrastructure is critical in delivering the required levels of scalability, availability, performance and security required for virtualized operating systems and applications.

Understanding the density of the planned virtual server infrastructure is critical in defining key IP addressing. Typically most networks are designed to accommodate approximately 250 hosts per subnet, with around 2,000 hosts in a large data center LAN. If we are to consider those 2,000 servers with 5:1 virtual server density, this translates to 10,000 IP addresses and 10,000 MAC addresses. The scalability of the EX4200 Series Ethernet switches with Virtual Chassis technology easily meets the need of these dense environments with large route and MAC address tables and scalable wire-speed performance.

The granular QoS capabilities of the Juniper Networks EX Series switches, with eight queues per port, also enables differing QoS policies to be set per virtual operating system and application.




Figure 7: Virtualized server infrastructure

Connectivity

Properly accounting for the required number of high-speed wired access ports for servers and storage devices as well as all aggregation layer connections in the data center is vital. Not only must the port density be specified, but the appropriate number of GbE and 10 GbE ports must also be taken into consideration. It's also important to account for any WLAN access points, IP phones, CCTV cameras and other IP devices the data center must directly support when addressing port requirements. The logical segmentation required and the number of logically separate networks that should share the same LAN must also be determined. These considerations help establish what type of hardware configuration is needed.

Power over Ethernet (PoE)

Most highly available data center facilities will have WLAN access points, IP phones, security cameras and other IP-based office automation peripherals, many of which require PoE to function. Accounting for the correct number and location of PoE ports needed in the data center is important at the access layer.

High Availability (HA)

Since the data center servers connected at the access layer are utilized by all throughout the enterprise, it's critical that data center networks operate with maximum reliability and uptime. The following levels of HA may be implemented in the data center:

1. Device-level HA

Most device failures are due to power supply failures or mechanical cooling problems. It is important to always support business processes with high-performance, carrier-class network switching devices such as the Juniper Networks EX Series Ethernet Switches or MX Series 3D Universal Edge Routers. Purchasing equipment with internal dual load-sharing power supplies and redundant fans or blowers to minimize equipment failure is always recommended and raises the mean time to repair (MTTR). Additional device-level HA can be provided by doubling up on key devices to assure that there is a backup device to pick up in the event of a failed device. If budget doesn't support a full set of backup devices, purchasing extra key device components such as a backup set of field-serviceable or hot-swappable power supplies or fan trays, helps mitigate the impact of a component failure.

2. Link-level HA

Ensuring that the data center maintains the data flow vital to business processes through internal and external resources is achieved through link-level HA. At the data center, link-level HA requires that two links operate in an active/backup configuration, such that if one link fails, the other can take over or reinstate the forwarding of traffic that had been previously forwarded over the failed link. Other technologies such as Link Aggregation (LAG) can be utilized to bond multiple uplinks and load balance across them.

3. Network Software HA

Juniper Networks Junos® operating system is the consistent operating system software that powers all of Juniper Networks' switch, router and high-end firewall products. It provides carrier-class network software to highly available data centers of all sizes. Junos OS supports features like nonstop forwarding (NSF), graceful protocol restart, in-service software upgrade (ISSU), Bidirectional Forwarding Detection (BFD) and other features which together make IP networking as failure-safe and reliable as traditional PSTN telephony networks. The Junos OS modularity and uniform implementation of all features enables the smallest data center to benefit from the same hardened services in their Junos OS-based devices as the largest service providers.

VLAN and Spanning Tree Protocol (STP)

Data centers typically use VLANs to group any set of servers or storage devices into logical networks through software configuration instead of physically relocating devices on the LAN. VLANs help address issues such as scalability, security and network management, as was introduced in the three-tier application model.

VLANs are Layer 2 broadcast domains that exist only within a defined set of switches. Using the IEEE 802.1Q standard as an encapsulation protocol, packets are marked with a unique VLAN tag. Tagged packets are then forwarded and flooded only to stations in the same VLAN. Tagged packets must be forwarded through a routing device to reach any station not belonging to the same VLAN. Any switch or switch port can be dynamically or statically grouped into a VLAN. Alternately, traffic may be grouped into a VLAN and forwarded through specific ports based on the specific data protocol being sent over the LAN. For example, VoIP traffic from a soft phone can be segmented from other traffic and put into a VLAN that receives a higher QoS.

Spanning Tree Protocol (STP)

VLANs may create multiple active paths between network nodes, resulting in problematic Layer 2 bridge loops. The loops will cause the same MAC addresses to be seen on multiple ports causing the switch forwarding function to fail. Also, the loop may cause broadcast packets to be forwarded endlessly between switches, consuming all available network bandwidth and switch CPU resources.

The IEEE 802.1D STP standard, ensures a loop-free topology for any Layer 2 bridged LAN. STP is designed to leave a single active path between any two network nodes by first creating a tree within a mesh network of connected LAN switches and then blocking the links which are not part of that tree. STP thus allows a network design to include redundant links to provide automatic backup paths if an active link fails, without the danger of bridge loops, or the need for manual enabling/disabling of these backup links.

Issues with STP

Troubleshooting may be challenging with STP due to complicated routing, incorrect configuration, or mis-cabling. Since every packet must go through the root bridge of the spanning tree, routing performance with STP can also be non-optimal. STP often creates underutilized links and lacks a load-balancing mechanism as well. In addition, STP has a slow convergence of up to 30 to 40 seconds after a topology change. The Rapid Spanning Tree Protocol (RSTP) was created to combat this, providing sub-second convergence but only on point-to-point links. The IEEE 802.1s Multiple Spanning Tree Protocol (MSTP) standard supports multiple instances of STP, but it also increases configuration complexity.

Using Layer 2 versus Layer 3 at the Access Layer

Access switches can be configured to use Layer 2 STP bridging protocols or Layer 3 routing protocols.

Using Layer 2 at the access layer

Using Layer 2 at the access layer is a traditional configuration providing plug-and-play configuration and making the deployment in smaller networks easier to implement and manage.

Since this option typically requires Spanning Tree with legacy solutions, troubleshooting can be more difficult in more complex networks, and convergence in case of a switch or link failure often takes too long for larger highly-available data center LANs.




Figure 8: Layer 2 versus Layer 3 at access layer

Using Layer 3 at the access layer

Routing is enabled on the switch when using Layer 3 at the access layer, but it still provides the ability to put data flows into different VLANs. Routing at Layer 3 to the access layer eliminates the creation of layer 2 loops and the need for spanning tree. Furthermore, Layer 3 routing is more deterministic. In this scenario STP can be disabled, making it easier to troubleshoot, which is important in larger networks. Using OSPF or other open-standard protocols for rapid convergence, delivers sub-second convergence. For larger or more complex networks, using layer 3 routing to the access layer lowers maintenance and administrative costs in comparison to using Layer 2 at the access layer.

Deploying layer 3 routing to the access layer is often more costly to deploy with legacy network equipment as it usually requires the additional purchase of a layer 3 software license.

Unlike competitive products, the Juniper Networks data center solution provides the ability to implement either Layer 2 or Layer 3 at the access layer without any added expense—Layer 3 features are built into the base license, and no extra license required. Instead of STP, the Juniper data center solution also uses open-standard protocols such as OSPF for rapid convergence. LAN designs using the EX4200 or EX4500 Ethernet Switches with Virtual Chassis technology also benefit from Redundant Trunk Group (RTG) protocol as a built-in, optimized replacement to STP for sub-second convergence. And, according to an independent 2007 Lake Partners³ study, time spent operating Juniper Networks solutions running Junos OS can be up to 25 percent lower than competitive solutions. Since cost is not an issue, LAN size and complexity best determine when each solution is most appropriate.

1. Small Data Center LANs

For small data centers with few devices and a simple topology, Juniper recommends using Layer 2 at the access layer. Such a LAN design can implement EX4200 Virtual Chassis at the core/aggregation and access layers, eliminating the need for STP and increasing convergence response while reducing CapEx and OpEx.

2. Most Data Center LANs

Since the LAN design for most highly available data centers has a series of redundant devices and connections, Juniper Networks recommends using Layer 3 to the access layer, which is included in the EX Series at no extra cost. In this design, Juniper recommends switches with Virtual Chassis technology to deliver high performance load balancing and simplified device management. This equates to lower CapEx and OpEx compared to competing solutions.

Physical Deployment: Top-of-Rack vs. Middle of Row/End-of-Row

In traditional top-of-rack (TOR) or the middle-of-row (MOR) or end-of-row (EOR) deployments, modular-chassis access-layer switches are used to provide high-performance, HA services and high-density GbE and 10 GbE connections to servers in the data center.




Figure 9: Top-of-rack vs. end-of-row switch deployments

Top-of-Rack (TOR)

This configuration places high-performance switches at the top of the server rack in a row of servers in the data center. Cabling run lengths are minimized in this deployment and simpler than middle-of-row (MOR) or end-of-row (EOR) configurations. However, each legacy switch must be managed individually, complicating operations and adding expense as multiple discreet 24- or 48-port switches are required to meet connectivity needs in TOR configurations.

Middle-of-Row (MOR)/End-of-Row (EOR)

In this configuration, high-density switches are placed in the middle or at the end of a row of servers in the data center. Traditional modular chassis switches have commonly been used in this deployment, where cabling is quite complex. Switch port utilization is suboptimal with traditional chassis-based switches, and most consume a great deal of power and cooling, even when not fully configured or utilized. In addition, these large chassis-based switches are usually large and take up a great deal of valuable data center space.

Storage Connectivity

Increased productivity and intelligent decision making both rely on instant access to valuable business data. With a critical impact on the bottom line, enterprise data storage must be fast, reliable and always available. It also must be secured against unauthorized access, unwanted modification and loss. Additionally, it must easily scale to meet compliance regulations and maintain important business records. Storage networks such as Fibre Channel, InfiniBand, iSCSI and NDAS should be included as part of the data center design. Virtualization technologies may be used to provide seamless and unlimited storage. Critical application servers, such as those from NetApp, directly connect to the storage devices through a separate host bus adapter (HBA) to ensure fast access to data. Other servers connect via Ethernet or

another interface to get access to the storage facilities. For high-performance data access needs, iSCSI solutions with TCP Offload capabilities or using 10 GbE for NAS solutions should be considered. Not only do these alternatives cost less than Fibre Channel, but they provide higher performance and are easier to manage. Additionally, separate QoS queues can be used to ensure critical data flows are prioritized appropriately. For example, a database application should be prioritized over other less important data flows such as archived document data. Storage must also be backed up on a regular basis without impacting LAN performance and be accounted for in the disaster recovery plan.

Quality of Service (QoS)

Each application on the LAN has different QoS requirements. Unified Communications have real-time requirements that are not necessary for most data applications. VoIP packets, for example, must be efficiently transported throughout the LAN and WAN to ensure high-quality voice communications, even when the network is experiencing high utilization or congestion. Simply adding more bandwidth doesn't make the network voice-friendly. Latency, jitter and packet loss are common VoIP challenges that must be addressed with QoS queuing and scheduling to ensure toll-quality VoIP communications.

Traditional applications such as Web browsing and e-mail work fine with the best-effort delivery standard on IP networks. However, additional requirements must be met to ensure effective delivery of voice, video conferencing and other real-time applications. Unlike streaming video, for example, real-time voice data can't be cached nor can lost voice packets be retransmitted, since both would add an unacceptable delay and ruin the quality of the communication and result in a poor experience. Voice packets, therefore, must be given top priority when creating QoS policies.

To facilitate QoS, data can be classified by a combination of physical port, device and protocol. For example, a block of IP phones connected to a specific LAN segment could be placed in a VLAN designated for voice traffic based on their port numbers. Or Link Layer Detection Protocol-Media Endpoint (LLDP-MED) may be used to discover an IP phone and automatically place it on a VLAN using IEEE 802.1X access control. Or traffic from a soft phone can be analyzed at the protocol level, with voice data given top priority regardless of the source port. Once the data is classified with the appropriate Differentiated Services Code Point (DSCP), it needs to be queued and scheduled. Most importantly, the same QoS rules need to be enforced consistently throughout the LAN and WAN.

QoS or Class of Service (CoS) features are built into all Juniper infrastructure, security and application acceleration solutions. Junos OS comes standard with a full complement of QoS services; for example, all EX Series Ethernet Switches support eight QoS queues per port and offer a range of policing options from best-effort delivery to enhanced delivery to assured delivery. Since the same Junos OS runs all Juniper router and switch products, the same QoS policies can be implemented throughout the data center LAN and across the WAN for easy and consistent traffic management. In addition, ASICs in all Juniper routers and switches support QoS by processing prioritized data and minimizing CPU load.

Data Center Access Layer Design Recommendations

To meet the access requirements of any sized data center, Juniper provides a scalable chassis or a traditional chassis-based solution.

Scalable Configuration with Virtual Chassis Technology

A data center LAN must be able to accommodate growth and adapt to new technologies. This needs to be done economically with respect to capital expense, network overhead and network operational expense perspectives. Juniper Networks addresses these requirements with a true innovation: the EX4200 Ethernet Switches with Virtual Chassis technology. This innovation advances the economics of networking by delivering the HA and high port densities of a modular chassis in a compact, cost-effective, pay-as-you-grow platform.

1. Features and Benefits

Each compact EX4200 Series switch offers either 24 100BASE-FX/1000BASE-X ports, 24 10/100/1000BASE-T ports or 48 10/100/1000BASE-T ports. The 10/100/1000BASE-T platforms offer either full or partial PoE options. Switches with the partial PoE option provide PoE on the first eight ports of the switch while switches with the full PoE option provide PoE on all ports. Each PoE port delivers up to 15.4 watts of power and is compatible with class 0-3 IP phones. The EX4200 Series switches' built in Link Layer Discovery Protocol-Media Endpoint Discovery

(LLDEP-MED) services provide a standards based mechanism to automate and extend the power management of these PoE endpoints as well as assist with inventory management and directories. The switches with the partial PoE option are ideal for data center access deployments where PoE for a small number of IP phones, WLAN access points or other devices needing power and purchasing full PoE is overkill.

Each EX4200 line of switch supports optional front-panel uplink modules supporting either four GbE or two 10 GbE ports for high-speed connections to aggregation or core switches. These uplinks support online insertion and removal.

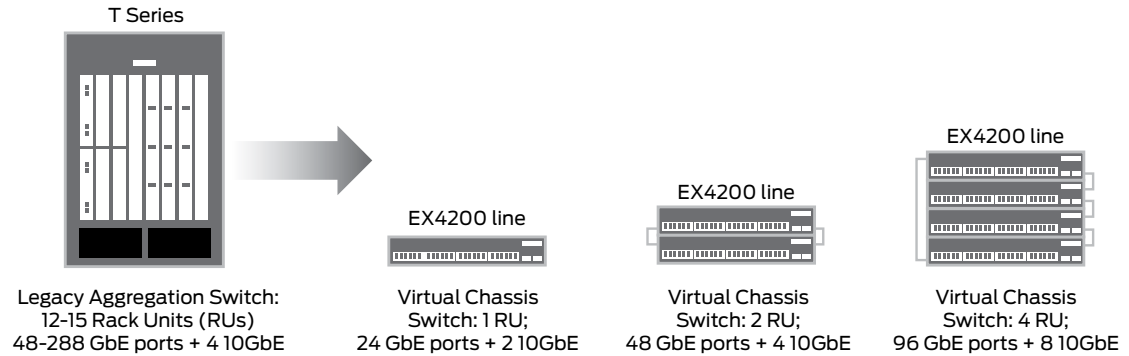


Figure 10: Virtual Chassis technology

2. Pay-As-You-Grow Scalability

The Juniper Networks Virtual Chassis technology enables a data center to add as many EX4200 line of switches as needed to meet its connectivity needs while delivering true chassis-like functionality. Juniper Networks' unique pay-as-you-grow model allows a single 1 RU top-of-rack EX4200 switch to be deployed and incrementally add up to nine more switches for a total of 10 switches. Resiliently interconnected via a 128 Gbps virtual backplane or 10 GbE uplink module, a fully-loaded Virtual Chassis configuration supports up to 240 100BASE-FX/1000BASE-X ports, 480 10/100/1000BASE-T ports, or any combination of the two, plus up to 20 10 GbE uplink ports. Not only does Virtual Chassis technology lower capital expenses when compared to legacy chassis systems by requiring less upfront investments, but it dramatically reduces operating expenses by enabling any group of interconnected switches to appear and be managed as a single switch. Coupled with the incremental, pay-as-you-grow model, the compact form factor of the EX4200 switches enables the data center to save not only on upfront and recurring rack space usage but also on costly power and cooling fees. Additionally, with the virtual-chassis configuration, cabling is greatly simplified.

3. Carrier-class Reliability

The EX4200 Ethernet Switches with Virtual Chassis technology provide the same HA features as modular chassis-based systems. Each switch supports internal redundant, load-sharing, hot-swappable AC or DC power supplies, as well as a field-replaceable hot-swappable fan tray with redundant blowers, any of which can fail without affecting operations.

Virtual Chassis technology provides unparalleled device and link HA utilizing the virtual backplane protocol and Junos OS. Each set of interconnected switches with Virtual Chassis technology automatically takes full advantage of the multiple route engines present to deliver graceful protocol restart. Graceful Route Engine Switchover (GRES) and non-stop forwarding ensure uninterrupted operation in the rare event of any individual switch failure. For added device and link HA, the EX4200 switches can be configured to address any requirements. For example, a single virtual-chassis configuration of 10 switches could instead be configured as two five-switch virtual-chassis configurations, or in any other desired combination.

4. Location Independence

Another key feature of the Virtual Chassis technology is that the virtual chassis protocol can also be extended across the 10 GbE uplink ports to interconnect switches that are more than a few meters apart, creating a single virtual switch that spans multiple wiring closets, floors or even data center server racks. Even when separated by long distances, interconnected switches with Virtual Chassis technology can be managed, monitored, upgraded and otherwise treated as a single resilient switch, dramatically reducing recurring management and maintenance costs. This enables either top-of-rack or end-of-row deployment with the EX4200 switches.

a. Top-of-Rack Deployments

Taking full advantage of Virtual Chassis technology, a scalable top-of-rack deployment takes the minimum amount of space with small form-factor switches that scale with high-density wire-speed ports as needed, lowering heating and cooling costs while conserving space. Virtual Chassis technology enables up to 10 units to interoperate and be managed as a single device, dramatically simplifying configuration and management while reducing operational costs and simplifying cabling. By configuring 10 top-of-rack switches as a single virtual chassis, fewer uplinks are necessary which further lowers cost, cable complexity and troubleshooting. Most importantly, the servers attached to the top-of-rack switches are all interconnected by a single, high-bandwidth low-latency switch and do not need to rely on traffic going to an aggregation switch for server-to-server communications—valuable for improving performance in an SOA environment.

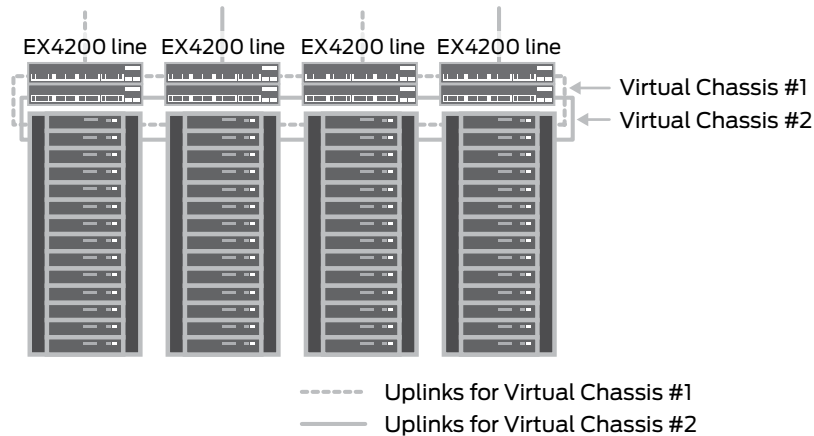


Figure 11: Top-of-rack deployment using Virtual Chassis technology

b. Middle-of-Row/End-of-Row Deployments

Configurations requiring middle-of-row or end-of-row deployments can also take advantage of Virtual Chassis technology with a small form factor that scales with high-density wire-speed ports as needed, lowering heating and cooling costs while conserving space. Dramatically simplifying operations and configuration, Virtual Chassis technology enables a set of up to 10 units to be managed as one device and lowers operations expense.




Figure 12: End-of-row deployment using Virtual Chassis technology

5. Reducing CapEx and OpEx

At one-eighth the footprint and less than one third the cost of the most commonly purchased chassis-based switch offering 480 1000BASE-T GbE ports and twenty 10 GbE wire-speed ports, the EX4200 switches with Virtual Chassis technology represents the new generation of GbE access switching.

The Juniper EX4200 Series switches include standard features that require costly add-ons in competitive solutions. For example, the EX4200 Series includes Layer 3 features in the base software license, offers built-in 10 GbE uplink capability, delivers partial or full PoE, provides built-in redundant power supplies and fans, and more

in a single cost-optimized platform. OpEx savings include the unified Junos OS feature set and remote mirroring capability for full troubleshooting, maintenance, upgrades and debugging from a central NOC.

Not only does Juniper Networks lower capital and operational expense by collapsing layers and therefore reducing the number of devices in the network that need to be purchased and managed, but Virtual Chassis technology saves on valuable rack space, as well as recurring power and cooling costs. Delivering greater value while reducing capital and operational expenses, Virtual Chassis technology frees up precious IT budget dollars that can be invested in new technologies that improve business productivity.

For a full set of features, benefits, and specifications, please view the Juniper Networks EX4200 Ethernet Switch with Virtual Chassis Technology Data Sheet.

EX4500 10GbE Switch

The EX4500 is another viable option for top-of-rack deployments, providing both 1GbE and 10GbE connections.

1. EX4500 Features and Benefits

The EX4500 supports up to 48 10GbE ports in a 2RU form factor, delivery a high-density solution for the data center. A non-blocking architecture ensures the highest throughput for web-based applications that require a lot of bandwidth. In addition, the EX4500's fiber ports are speed auto-sensing, meaning they automatically configure themselves as either 10GbE or 1GbE, based on the optical interfaces inserted. As a result, the EX4500 delivers a highly effective 1GbE-to-10GbE migration path for racks including a mix of 1GbE and 10GbE NICs.

The EX4500 also offers the same HA as the other EX Series platforms, with redundant fan blowers and power supplies.

Modular Chassis Configurations

The Juniper Networks EX8200 Ethernet Switch is recommended as an access layer solution for those requiring modular chassis configurations.




Figure 13: EX8200 line of modular chassis solutions

1. EX8200 Features and Benefits

To meet the access demands of even the largest data center, the top-of-the-line EX8200 Terabit-chassis switch delivers a powerful, high-density, high-performance solution. Capable of up to 3.2 Tbps throughput, the EX8200 line of switches offer up to 368 (eight-slot chassis) or 764 (16-slot chassis) wire-speed GbE ports, up to 64 (eight-slot chassis) or 128 (16-slot chassis) wire-speed 10 GbE ports, or up to 320 or 640 5:1 over-subscribed 10GbE ports. With a redundant control plane, the EX8200 line also runs Juniper Networks' top-rated Junos OS for carrier-class HA. Other features include inline application visibility, DDOS protection and anomaly-based threat detection.

2. Middle-of-Row/End-of-Row Deployments

The EX8200 line of switches enable traditional middle-of-row or end-of-row deployments with a scalable fixed form factor with either high-density wire-speed ports or 5:1 over-subscribed 10GbE ports.




Figure 14: End-of-row deployment using fixed chassis technology

Data Center Aggregation Layer

The aggregation layer, sometimes referred to as the distribution layer, aggregates connections and traffic flows from multiple access layer switches to provide connectivity to the LAN core or WAN edge layer switches.

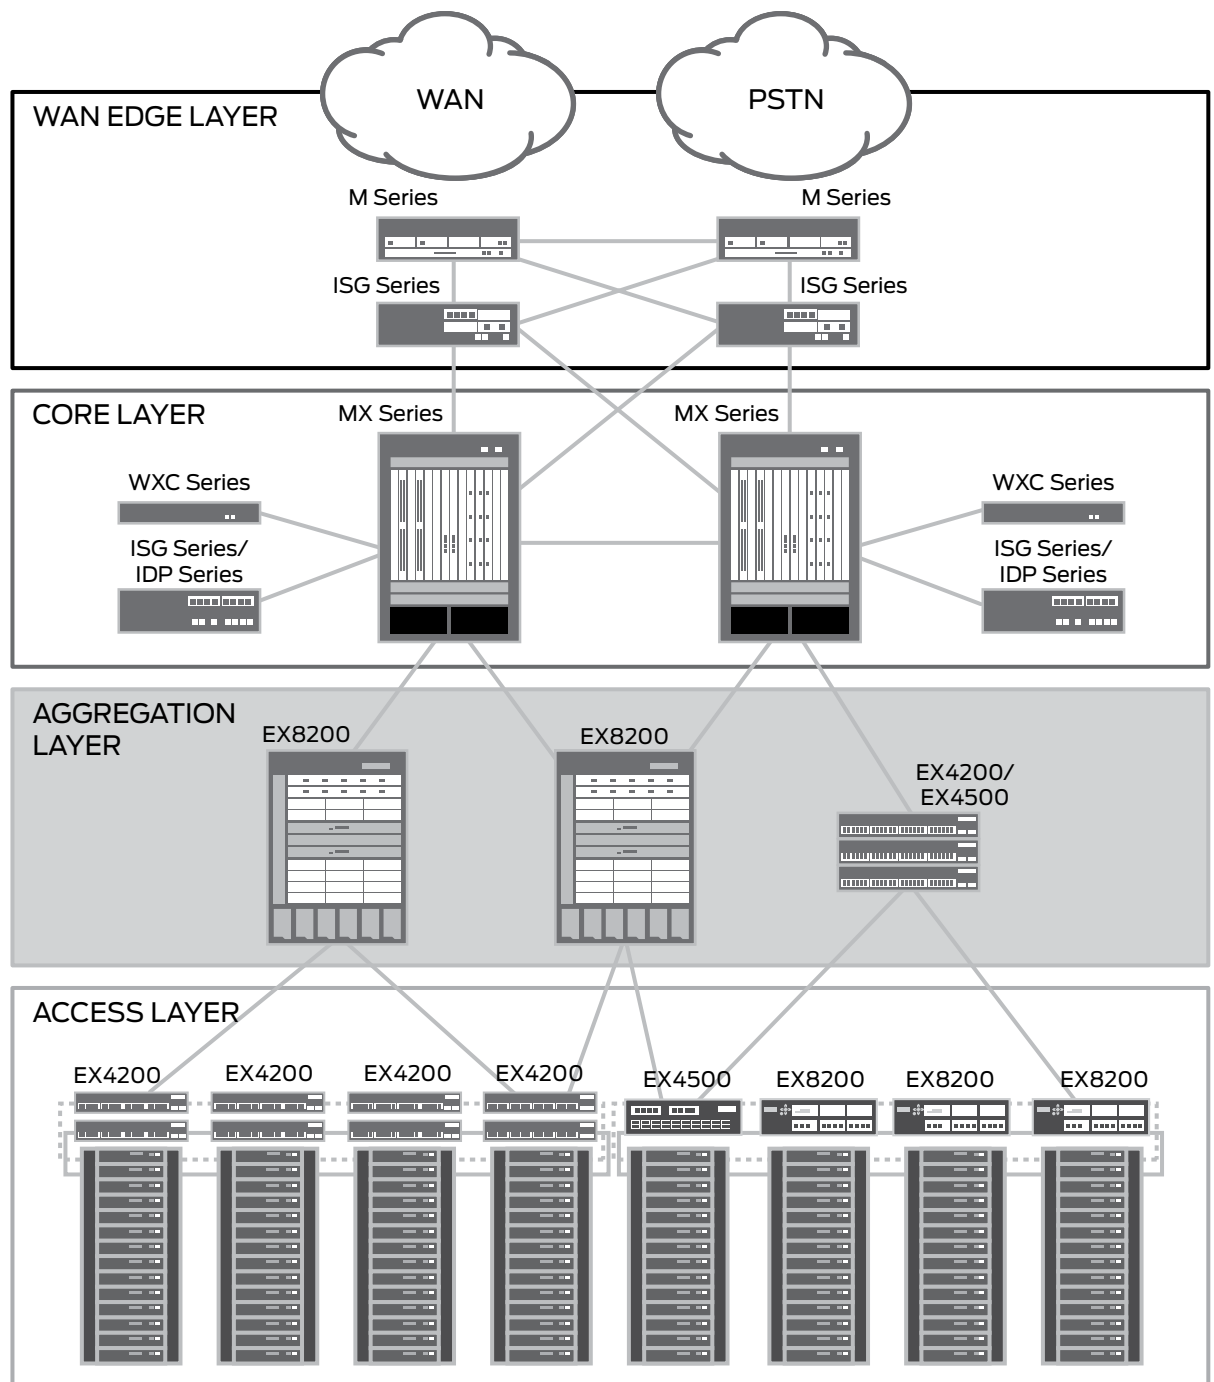


Figure 15: Aggregation layer in a highly available data center LAN

Aggregation Layer Design Considerations

Due to their location in the network, aggregation-layer switches must provide scalable, high-performance, high-density, wire-rate ports, and HA hardware and software features that deliver carrier-class reliability and robustness. The aggregation layer is also a location from which to deploy additional services, such as threat containment. Layer 3 should be provided at the aggregation layer for route summarization, fast convergence, and load-sharing.

In some instances, based on port density, aggregation throughput, and oversubscription requirements, the aggregation layer may be eliminated and collapsed into the core layer. For more detail on this configuration, please view the Data Center Core Layer Design Recommendations section.

High Availability (HA)

It's crucial that data center networks operate with maximum reliability and uptime. Device redundancy is required, and all devices must have robust HA features such as redundant, load-sharing power supplies and cooling fans, and in some cases, fully redundant hardware. Redundant GbE downlinks to the access layer and 10 GbE uplinks to the core layer are also required.

Scalability

The aggregation layer must provide high-density port connectivity to the core layer and be able to easily handle peak throughput while adding minimal latency.

Network Virtualization

Aggregation switches should also support generic routing encapsulation (GRE) tunneling for sending mirrored traffic to monitoring devices in the network operations center for centralized troubleshooting and analysis, or to build segregated overlay networks without the challenges associated with Spanning Tree.

Application Visibility

To successfully manage a network, it's important to know how it's being used so that application deliver may be optimized and efficiency maximized. Real-time information and detailed reporting are needed to provide rapid access into LAN wide application information that can help identify patterns or applications that are disrupting performance or in need of QoS support.

Security and Threat Containment

It's vital that the aggregation layer include integrated security features to guard against intruders or other external threats such as distributed denial of service (DDoS) attacks. It should deliver an extra layer of security by first authenticating users and performing virus checks, then enforcing precise, end-to-end security policies that determine who can access what network resources, as well as quality of service (QoS) policies to ensure delivery of business processes.

Data Center Aggregation Layer Design Recommendations**Traditional Layered Approach**

For a traditional three-layer network design, Juniper Networks recommends the EX8200 line of switches for aggregation layer deployment. All Juniper solutions at the aggregation layer offer the following features and benefits:

1. High Availability (HA)

The EX8200 line of switches offer fail-safe operations. Redundant links to each core layer device are provided in the event of a device or link failure. The EX8200 line also offers a redundant control plane as well as redundant power supplies and fans. All equipment runs Junos OS, providing HA features such as QoS and Graceful Routing Engine Switchover, preserving forwarding and routing operations during device events with non-stop forwarding and automatic load balancing.

2. Scalable Performance

- a. EX8200 line

To meet the aggregation demands of even the largest data center, the EX8200 line of Terabit-chassis switch delivers a powerful, high-density, high-performance solution. Capable of up to 3.2 Tbps throughput, the EX8200 modular Ethernet switches offer up to 64 (eight-slot chassis) or 128 (16-slot chassis) wire-speed 10 GbE ports. The EX8200 line delivers 200 Gbps of switching capacity per slot, enabling the future addition of 100 Gbps uplinks. By providing capacity now, the EX8200 line of switches allow users to easily migrate to higher-speed connections when they are ready—without requiring any changes to the switch fabric, Route Engines, backplane, power supplies or cooling system. The EX8200 line also offers a redundant control plane and runs Junos OS for maximum software HA.

The EX8200 line of switches include integrated security features to guard against intruders or other external threats. Integrated anomaly-based threat detection provides additional protection by identifying and blocking distributed denial of service (DDoS) attacks. Taking advantage of behavioral threat detection algorithms, the EX8200 line of switches are also capable of identifying and closing half-open sessions—important for defending against zero-day threats for which no signatures exist.

b. EX4500

Similar to the EX4200 Ethernet Switches with Virtual Chassis technology, the EX4500 can be utilized in smaller aggregation deployments.

c. Virtual Chassis

The EX4200 Ethernet Switches with Virtual Chassis technology can be utilized in smaller aggregation configurations requiring high-density 1000BASE-X fiber GbE ports. For typical aggregation environments requiring 48 GbE SFP fiber ports and four 10 GbE uplinks, two 24-port EX4200 Ethernet Switches deliver the same wire-speed port densities and functionality as the most popular chassis-based solution—at one-sixth the size, one-fifth the power, and one-third the cost.

3. CapEx and OpEx Savings

Typically more than two layers of legacy Layer 3 switches are required to achieve the wire-speed port densities demanded by today's high-performance data center. The Juniper Networks EX Series Ethernet Switches, however, meet these needs and also enable the collapse of the number of aggregation layers, creating a direct positive impact on the economics of networking.

Junos OS also simplifies network operations and lowers operating expense on all fronts, from upgrades and moves, adds and changes to troubleshooting and problem resolution.

The EX Series switches deliver greater value while reducing capital and operating expenses. This frees up valuable IT resources that may be invested in new technologies to improve business productivity and further streamline operations.

For a full set of features, benefits and specifications, please view the Juniper Networks EX Series Ethernet Switches data sheet.

Collapsing the Aggregation Layer into the Core Layer

An aggregation layer is not always necessary and may be eliminated in some data center LAN configurations. Based on port density, aggregation throughput, and oversubscription requirements, the aggregation layer may be collapsed into the core layer. For more detail, please view the Data Center Core Layer Design Recommendations section.

Data Center Core Layer

The core layer provides a fabric for high-speed packet switching between multiple aggregation devices or the access layer in a collapsed network. It serves as the gateway to where all other modules meet, such as the WAN Edge. The core typically requires 10 GbE interface for high level throughput, and maximum performance to meet oversubscription levels.

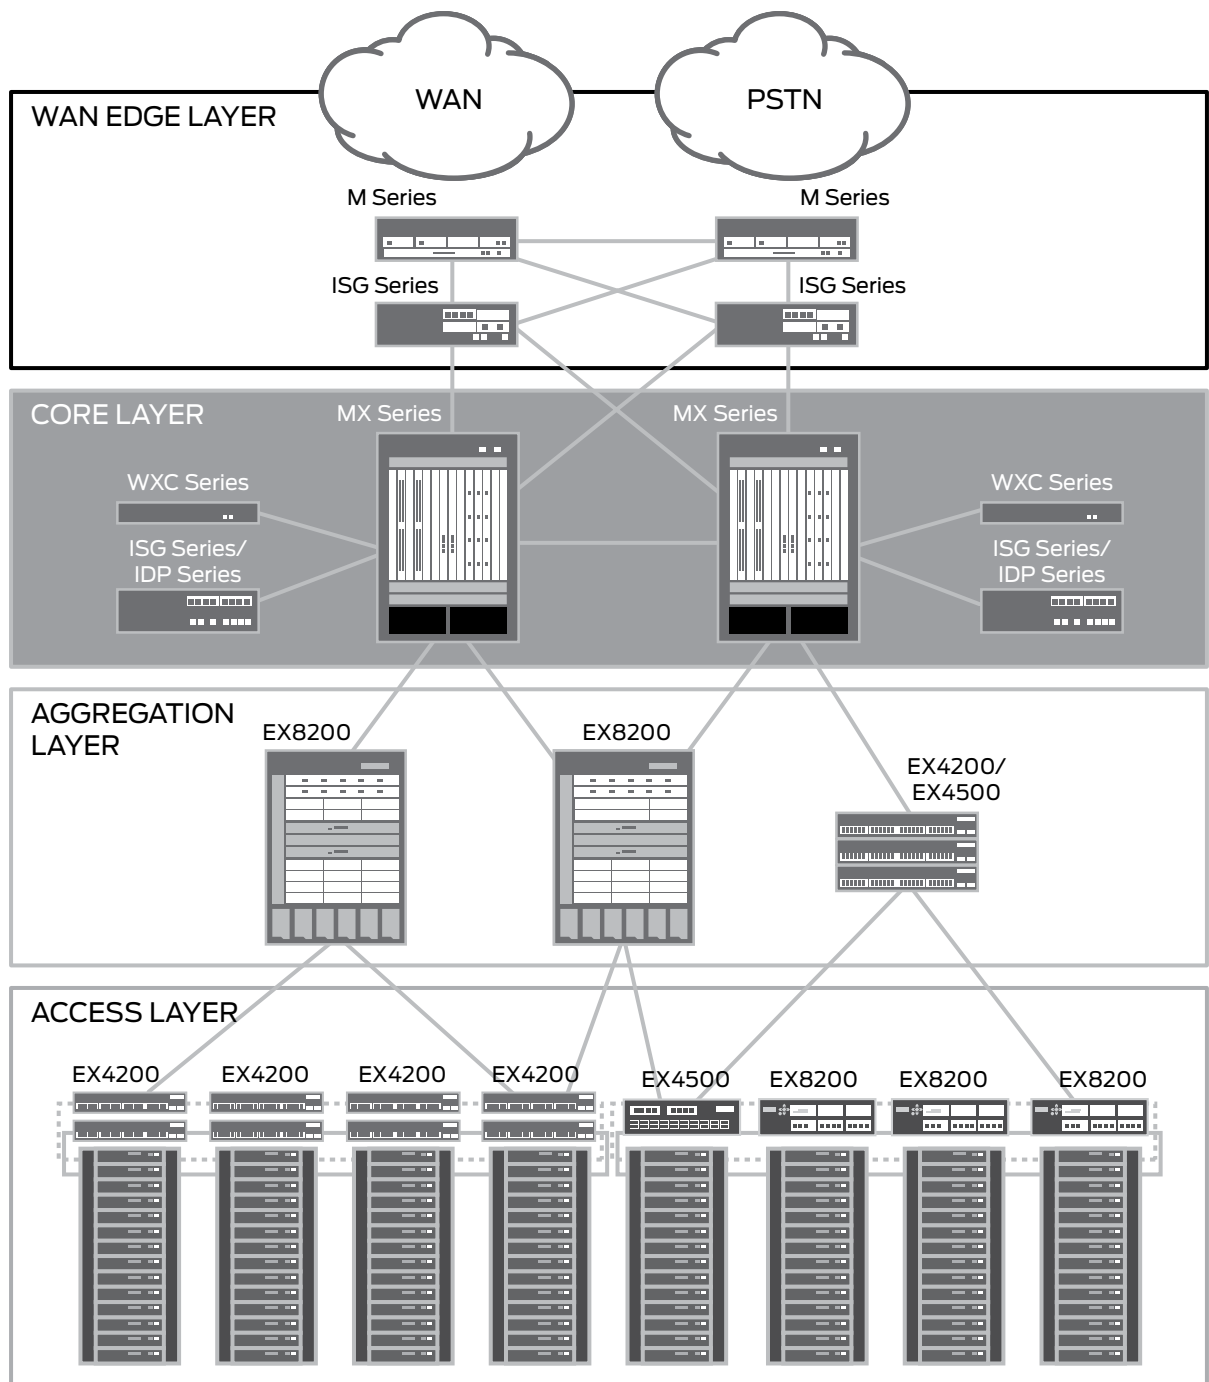


Figure 16: Core layer in a highly available data center LAN

Data Center Core Design Considerations

The core provides high-speed throughput for all data going in and out of the data center. The core layer must provide resilient, fail-safe Layer 3 connectivity to multiple aggregation devices.

High Availability (HA)

All core layer devices in the data center must provide a full complement of HA services to maintain critical uplink connectivity. The devices must be robust and offer fully redundant hardware. Core layer devices should be load balanced for optimal performance and also run OSPF or another open protocol for fail safe connectivity between layers.

Data Center Core Layer Design Recommendations

The EX8200 line of switches or MX Series routers are recommended as core layer solutions.

1. High Availability (HA)

Both Juniper core layer solutions offer fail-safe operations. Redundant links to each core layer device are provided in the event of a device or link failure. The MX Series offers fully redundant hardware. The EX8200 line offers a redundant control plane as well as redundant power supplies and fans. All equipment runs Junos OS, providing HA features such as graceful protocol restart and Graceful Routing Engine Switchover, preserving forwarding and routing operations during device events with non-stop forwarding and automatic load balancing.

2. Scalable Performance

The decision to select the EX8200 line of switches or the MX Series routers depends on throughput, high-density non-oversubscribed 10 GbE port, and high scaling requirements for MAC, IP, or IP multicast tables and/or ACL entries.

c. EX8200 Series

The EX8200 delivers a powerful, high-density, high-performance solution. Capable of up to 3.2 Tbps throughput, the EX8200 line of switches offer up to 64 (eight-slot chassis) or 128 (16-slot chassis) wire-speed 10 GbE ports. The EX8200 line delivers 200 Gbps of switching capacity per slot. The EX8200 line also offers a redundant control plane and runs Junos OS for maximum HA.

The EX8200 line provides wire-speed application visibility into more than 150 applications via integrated high-performance ASICs. The EX8200 line of switches also include integrated security features to guard against intruders or other external threats. Integrated anomaly-based threat detection provides additional protection by identifying and blocking DDoS attacks.

When high-density, non-oversubscribed 10 GbE ports are required, the MX Series chassis routers are recommended. Built on a flexible modular chassis with fully redundant hardware capable of up to 960Gbps throughput, the MX Series offers high scaling of IP, IPMC, MAC or ACL. Running Juniper's common Junos OS, the MX Series also provides fully featured Layer 2 and Layer 3 Dense Port Concentrators, high performance multicast support, MPLS, L2/L3 including MPLS, NSR, ISSU, GRES, and more.

3. CapEx and OpEx Savings

Typically more than two layers of legacy Layer 3 switches are required at the core to achieve the wire-speed port densities demanded by today's high-performance data center. Enabling the collapse of the number of core layers, the high-density, high-performance EX8200 or MX Series routers create a direct positive impact on the economics of networking. The switches also lower operating expense and simplify all network operations via Junos OS.

Delivering greater value while reducing capital and operating expenses, the EX8200 line and MX Series devices free up valuable IT resources that may be invested in new technologies to improve business productivity and further streamline operations.

For a full set of features, benefits and specifications, please view the Juniper Networks EX Series Ethernet Switches data sheet and the Juniper Networks MX Series 3D Universal Edge Routing Platforms data sheet.

Consolidating the Aggregation Layer and the Core Layer

Based on port density, aggregation throughput, and oversubscription requirements the aggregation layer may be collapsed into the core. When determining whether to collapse the aggregation layer, the throughput and port density of available 10 GbE connections should be considered. It's also important to consider future growth. In some instances, capacity may be exceeded in the near future and thus dictate that it might be simpler not to collapse layers as adding a layer later on can be time consuming and disruptive to LAN operations and uptime.

Aggregation at the core also allows for more flexibility and easier support of virtualization but requires very high-speed processing and HA levels. One of the biggest advantages of this 2-layer design is a dramatic reduction of the number of devices which offers significant power savings, reduces the facilities footprint of the system, offers simplified device management, and allows tighter security control. In addition, it also reduces the number of system failure points. The scalability limitation of this architecture is typically limited by the scalability of the core network devices.

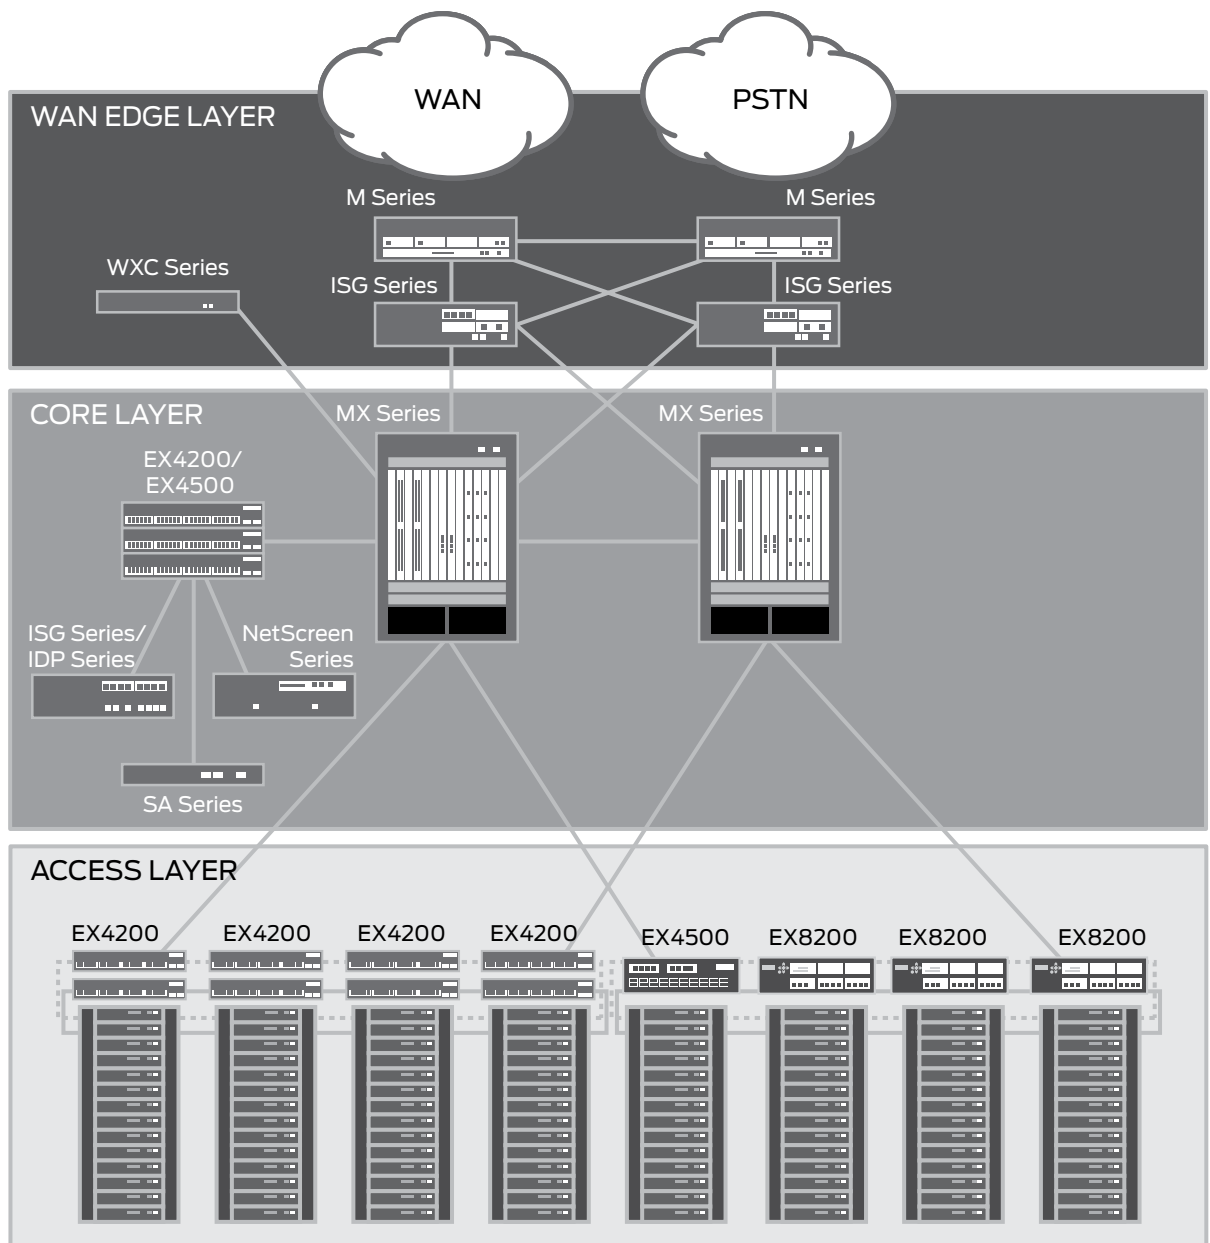


Figure 17: Aggregation layer collapsed into the core layer in a highly available data center LAN

4. Features and Benefits

When collapsing the aggregation layer using an MX Series router, the resulting configuration creates operational efficiencies and cost savings with fewer devices to manage and a reduction in power usage and cooling expenses. The MX Series high throughput ensures optimal performance and HA features while providing all the functionality provided at the core

WAN Edge Integration

WAN connectivity provides the vital link to centralized services and resources through which all campuses, remote branch offices, and end users connect. This document is not intended to cover all aspects of WAN Edge which can be found in other publications, but to introduce some of the challenges that all high performance organizations face when designing and scaling a data center LAN for assured network connectivity and performance.

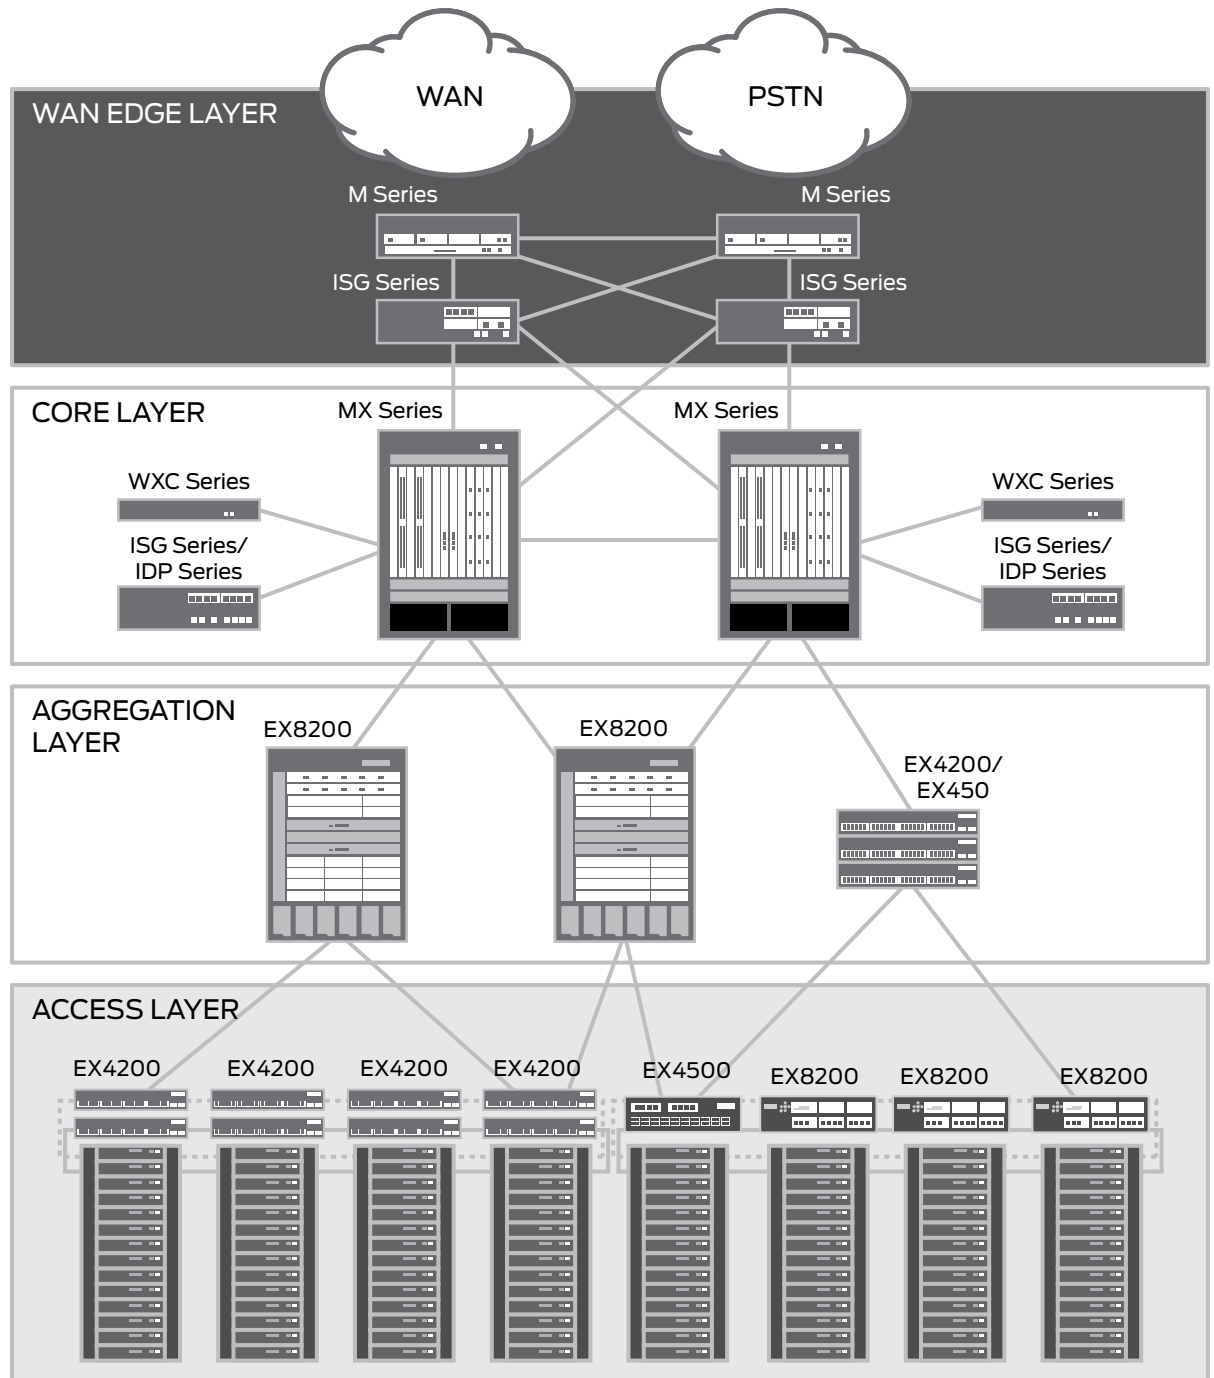


Figure 18: WAN edge in a highly available data center LAN

WAN Edge Design Considerations

The following WAN edge routing platform must offer sufficient high-speed Ethernet ports to provide connectivity between the WAN and the core or aggregation layer. It also must provide high-performance throughput to the Internet and WAN.

Connectivity

A WAN edge routing platform must offer sufficient high-speed Ethernet ports to provide connectivity between the WAN and the core or aggregation layer. It also must provide high-performance throughput to the Internet and WAN.

High Availability (HA)

All WAN edge devices must provide a full complement of HA services to maintain critical WAN connectivity. The hardware must be robust and offer redundant power supplies and cooling fans. Devices should be paired in active/active routing states for optimal HA. And an alternate connection to the Internet or WAN must be maintained.

Firewall/VPN

Security must be provided at the WAN edge, including VPN connections to remote locations and users as well as integrated firewall services to protect against worms, trojans, viruses and other malware. Such services should be centrally managed to facilitate rapid deployment and minimize ongoing operational costs.

Backup/Disaster Recovery

A data center backbone is a key component in the architecture and WAN Edge design primarily for disaster recovery reasons considering the scale of processing performed at data centers, and the requirements for regulatory compliance. As such, the data center backbone supports a variety of computational services such as data mirroring to assure highly accurate data is represented at multiple data centers. 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 services oriented architecture applications, and lastly the support for legacy clustering technologies that may require Layer 2 connectivity are all functions that are reliant upon the a high performance data center backbone.

WAN Edge Layer Design Recommendations

A WAN edge routing platform must offer sufficient high-speed Ethernet ports to provide connectivity between the WAN and the core or aggregation layer. The Juniper Networks M Series Services Router meets these requirements and more.

M Series Routing Platform

The M Series platform provides predictably high performance and a modular, carrier-class interface that delivers secure, reliable and scalable network connectivity.

1. Features and Benefits

Capable of throughput up to 320 Gbps, the M Series multi-service edge router offers a full breadth of connectivity options from DS0 to OC-192/STM-64 as well as 100 Mbps Fast Ethernet to 10 GbE. The platform also runs Junos OS, providing advanced carrier-class and field-proven routing features including advanced services such as MPLS, IPv6, hierarchical QoS and multicast in the base system at no additional license fee or upgrade.

2. HA

The M Series delivers carrier-class HA with fully redundant hardware, including redundant Routing Engines and Switching/Forwarding Engine Boards. Junos OS provides additional software HA features.

3. Integrated Services

The M Series solutions provide the essential security functions required for securely connecting sites over the Internet, including integrated firewall and IPsec VPN. The platform also supports centralized user security policy and enables a unique HA option in the form of dynamic route-based VPNs. Virtualization technologies allow segmentation of the network into many separate zones within a single platform for enforcing compliance to corporate security policies.

Built in QoS improves bandwidth utilization and Unified Communications performance, it also minimizes latency, jitter, and packet loss to ensure voice and data performance.

In addition to a command line interface (CLI), J-Web—built-in Junos OS—offers remote Web-based management of all M Series models. Built-in troubleshooting also minimizes network downtime and decreases operating expenses and revenue losses due to outages.

The M Series consolidates multiple services into a single platform, providing the lowest possible CapEx. The rich feature set allows customers to trial many different services without capital expenditure and scale successful services to larger populations.

For a full set of features, benefits and specifications please see the Juniper Networks M Series Services Routers data sheet.

For a more complete discussion on WAN Edge integration, please see the Juniper Networks Migrating to Next Generation WANs data sheet.

Operational Simplicity and Unified Management

Network operations form a large portion of any IT budget, and any methods of simplifying data center LAN operations help reduce operations expense. The four main challenges that complicate the streamlining of network operations are:

Inconsistent Feature Set

Most hardware solutions have different operating systems or feature implementations for each platform. One leading switch provider has hundreds of different operating systems in its product line, requiring IT to invest considerably in training to master a variety of interfaces. It also adds a layer of inefficiency and complexity while increasing the potential for misconfiguration when trying to apply consistent enterprise-wide services across the data center LAN, WAN, campus LAN, and remote branch LANs.

Upgrades and Deployments

Testing and deploying operating system upgrades or patches can be a time-consuming and ongoing process due to the number of different operating systems found in most legacy data center LAN solutions and the varying release schedules to which each adheres.

Unreliable Monolithic Operating Systems

Legacy hardware solutions have operating systems built on a monolithic architecture with each code function intertwined with the others. If any part of the monolithic operating system fails—for example, a bug in SNMP—the operating system crashes and reboots the system. Such a fault can cause the line cards to crash or restart, resulting in hundreds of seconds of downtime—which ripples across the enterprise, adversely affecting the bottom line.

Lack of Unified Management

The lack of unified features also impacts all aspects of setting and managing device configurations, network settings, and security policies. Not only do different interfaces increase the time of each task, but operations costs are further increased as IT needs to visit remote branch locations to configure devices, apply network settings and set security policies. What's needed instead is a set of unified and centralized management tools to address these types of operations remotely.

Juniper Networks addresses all of these issues and reduces costs by providing Junos OS, Juniper Networks Network and Security Manager (NSM), and J-Web.

Achieving Operational Simplicity with Junos OS

Junos OS is the common operating system on all Juniper Networks switches, routers, firewalls and acceleration solutions. Not only does Junos OS deliver advanced carrier-class network services, it provides a consistent feature set, and a centralized management capability which simplifies planning, speeds implementation, and enables intuitive day-to-day operations and management of any network.

The Power of Junos OS

Fundamental to the value of the Junos OS are the “three ones”—one source code, one release train and one modular architecture. By running a common operating system on all products, Juniper dramatically reduces maintenance and management overhead while ensuring interoperability and a consistent feature set across all products.

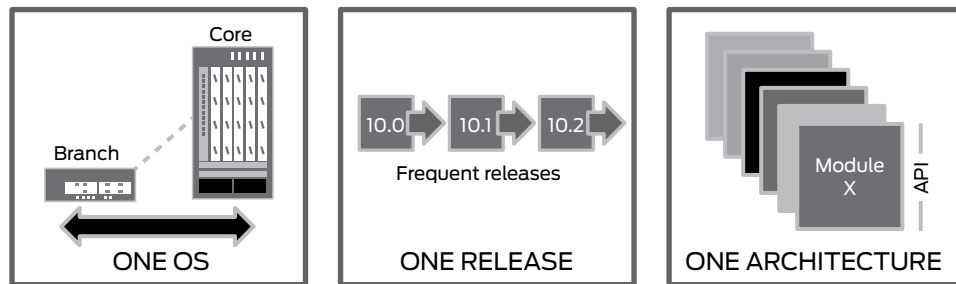


Figure 19: Junos OS—The three ones: one source code, one train, and one modular architecture

Modular Processes

The Junos OS is a completely modular operating system, enabling a functional division of labor for seamless development and operation of many advanced features and capabilities. By partitioning the software system, tasks are broken into manageable subsets that interact infrequently and provide new levels of fault-tolerance. Unlike monolithic operating systems, each key Junos OS function executes as an independent process and runs in its own protected memory space. Loading or executing one doesn't affect the others. One daemon can restart independently without disrupting another or forcing a full system crash or restart. A benefit of this approach is the ability to maintain full control of the switch or router at all times. Because of the separation of control, forwarding and services, filters can be added in real-time to thwart a Distributed Denial of Service (DDoS) attack.

Rollback Capability

Junos OS also offers error-resilient configuration that prevents operators from inadvertently bringing down the data center network. IT must explicitly commit changes after entering and reviewing all modifications. If a configuration change causes loss of connectivity to the device and no follow-up confirmation is provided, the device automatically reverts back to the previous configuration, restoring connectivity—saving time and ensuring Link-level HA. In addition to automatically checking for errors or incorrectly constructed configurations that could cause potential problems, Junos OS provides a rollback command to quickly restore any of the 50 prior configurations.

Advanced Features

The Junos OS also provides a broad spectrum of advanced routing and security software features such as stateful firewall, IPsec, MPLS and IPv6 without requiring an additional software license. In addition, the Junos OS provides comprehensive QoS functions to classify, prioritize and schedule traffic for applications such as VoIP. When Virtual Chassis technology is used, the Junos OS enables bidirectional forwarding detection for early detection of node or link failures.

Benefits

By running a common operating system, these Juniper solutions dramatically reduce maintenance and management overhead while ensuring a consistent feature set across all products, as well as a consistent implementation and management of those features. This equates to time savings in all categories of operations. In addition to a reduction in training time, the inherent interoperability across all platforms greatly simplifies new feature deployment, software upgrades and other network modifications. A single consistent code set also enables customers to qualify and deploy just one release. For many customers, the testing time of a new release is cut from what was months down to just a few weeks. Junos OS also provides features to facilitate fast restoration of previous configurations.

Impact

In an independent study conducted in 2007, Lake Partners quantified the time savings Juniper Networks customers experienced using the Junos OS across a number of common network operational tasks. The results are presented in Table 2:

Table 2: Junos OS Operating Efficiencies (Lake Partners 2007)

NETWORK OPERATIONS TASK	AVERAGE JUNOS EFFICIENCY
Adding Infrastructure	29%
Upgrading and Planned Events	23%
Troubleshooting and Unplanned Events	54%
Monitoring and Optimizing	24%
Average Time Saved With Junos software	25%

This time savings translates to a substantial, tangible cost savings. According to Lake Partners, an infrastructure of any size running Junos OS can save up to 29 percent on operational costs. Seeing that the IT department of a typical enterprise spends 40 to 60 percent of its budget to maintain and enhance basic IT services (McKinsey & Company 2006), this savings could be considerable.

Unified Management with Juniper Networks Network and Security Manager (NSM)

The Juniper Networks Network and Security Manager (NSM) product is a powerful, centralized management solution that controls the entire device life cycle of firewall/IPsec VPN and intrusion prevention system (IPS) devices, including basic setup and network configuration with local and global security policy deployment. Unmatched role-based administration allows IT departments to delegate appropriate levels of administrative access to specific users, thereby minimizing the possibility of a configuration error that may result in a security hole. NSM can easily scale to meet the needs of any enterprise with data centers. A wide range of reporting tools are available, enabling IT to view and analyze network traffic, device and VPN statistics, system resources, and other administrative information. IT can also customize templates for commonly used reports and generate these reports on a regularly scheduled basis.

Benefits

NSM lowers operational costs by presenting a GUI to simplify complex tasks such as device configuration, supplying device templates to minimize configuration errors, providing investigative tools for complete visibility into the network, and more.

Remote Configuration and Management with J-Web

In addition to a full-featured command-line interface (CLI), J-Web, a Web-based tool, is available to configure and manage any Junos OS-powered device.

Benefits

Built on Junos OS, J-Web offers highly-available data centers a graphical user interface for device management complementing the exiting suite of element and service management products from Juniper. J-Web provides IT administrators and network operators with simple-to-use tools to quickly and seamlessly monitor, configure, troubleshoot and manage any switch, router, or firewall.

J-Web allows non-technical users in data center/small office environments to commission and bring a router online quickly and easily. It offers seamless GUI access to all of the features and functions of Junos OS, reducing timelines for new service deployments. J-Web can be quickly integrated into existing network management or OSS (Operational Support System) applications such as Micromuse Netcool Omnibus, Dorado RedCell Manager, IBM Tivoli and HP OpenView, thereby minimizing complexity for the service provider or enterprise customer. Fast error-free service changes and upgrades can be made with J-Web's quick configuration wizards, and new services can be rapidly created and deployed with the use of configuration and QoS wizards that allow for real-time changes to service parameters.

Conclusion

The data center network is arguably the most valuable corporate asset. It plays an integral role in supporting key business processes and joining today's increasingly decentralized workforce. With a trend towards the centralization and consolidation of data centers and servers, a high-performance, highly available network is vital to overall business success. Legacy solutions cannot meet the growing data center LAN needs for security, connectivity, performance and high availability. A new data center LAN design that meets these needs while enabling key IT initiatives is required. It must also economically scale and flexibly accommodate new computing trends and leverage new technologies such as virtualization without an entire redesign.

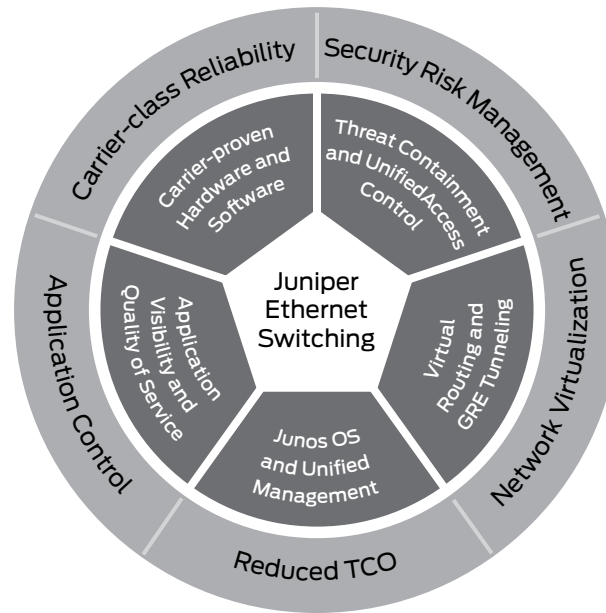


Figure 20: Juniper switching solutions

Juniper solutions, including a new family of high-performance Ethernet switches, redefine the way businesses build data center networks. Offering high port densities, wire-speed connectivity and high availability in compact, pay-as-you-grow platforms, Juniper switches represent a powerful yet cost-effective alternative to the aging and expensive solutions pushed by today's dominant switch vendors. By offering a smaller footprint in the data center, combined with lower power and cooling requirements, the Juniper switches represent the efficient and "green" solutions users are looking for to power their networks of the future. In addition to a full suite of secure services, Juniper products provide the end-to-end QoS required for latency sensitive and bandwidth-hungry applications such as Voice and Video.

Junos OS, a single, consistent operating system, is used across all Juniper switch, router and firewall products, making the network infrastructure exceedingly easy to deploy, configure and upgrade, saving considerable time and operating resources that can be reallocated to further improve business operations and maximize customer satisfaction.

Data center infrastructure solutions from Juniper Networks enable business today to deliver 24x7 carrier-class services at an enterprise price point. Juniper solutions advance the economics of networking, allowing businesses to "change the rules" with their IT investments and create a truly innovative and competitive environment that helps them increase revenue and raise productivity today and into the future.

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
www.juniper.net

APAC Headquarters

Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803


EMEA Headquarters

Juniper Networks Ireland
Airsides Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

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

Copyright 2012 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Junos, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

8020010-003-EN Nov 2012

 Printed on recycled paper