CAMPUS NETWORKS REFERENCE ARCHITECTURE

Table of Contents

Table of Figures

Introduction

A typical campus environment consists of a plethora of legacy devices all requiring different operating systems and management tools. With layers of legacy devices that do not seamlessly integrate, network architects and administrators are burdened more than ever to maintain network performance to meet the increased demands of intensive bandwidth applications, savvy users, and increasing costs. Consequently, today's IT personnel are forced to spend more time and effort planning, configuring, deploying, and managing a myriad of network devices that run on different operating systems.

Typically, campus architecture functionally spans up to three layers, from desktop devices connected to wiring closet switches at the access layer, to the core layer at the center of a large campus LAN. The hierarchical topology segments the network into physical building blocks, which simplifies operation and increases availability, but also adds considerable management complexity, reduces network availability, and increases capital and operational expenses. This occurs when there are numerous incompatible devices that require additional management and operating systems. This three-layer functionality can be collapsed into a two-layered physical design by using high-performance and virtualization capabilities, resulting in a lower total cost of ownership and, by lowering the number of devices, a reduction in the environmental footprint. Juniper Networks® campus network infrastructure allows for collapsing the traditional three-layered infrastructure, which makes planning and deployment easier for architects and designers.

From Juniper's perspective, today's campus network must be a high-performance, highly scalable environment with unique qualifications, especially in the current economic climate. To design a modern campus environment that meets these requirements, Juniper focused primarily on three campus components: connectivity, security, and manageability. These components are discussed in detail later in this document.

A modern campus network design is essential since legacy solutions cannot offer the services mentioned above at a lower total cost of ownership with minimal operational overhead. The modern campus design must allow for a high-performance network that is scalable, easy to deploy and upgrade, and accommodates emerging computing trends such as extensible use of file sharing, downloading, additional network services, and emerging technologies like unified communications. Campus architects can achieve this if they adhere to the fundamental principles addressed in this reference architecture, namely LAN/WAN connectivity design considerations, security, and centralized management. If practiced, this reference architecture will aid as a design tool to help architects avoid complete campus network restructuring.

Trends and Challenges

The campus network of the present is a strategic location in the distributed enterprise acting as the headquarters for productivity by offering fast, secure, and reliable services at scale, all the time. The campus network has been forced to evolve from supporting traditional client/server data flows to supporting real-time application traffic such as video conferencing and multicast traffic flows while accommodating the ever-increasing number of devices, services, and users.

According to Gartner Research, global expenditure for campus LAN equipment reached $18 billion in 2008, making campus infrastructure one of the largest network equipment expenditures in the enterprise market. As enterprises leverage their network to increase productivity, there are opportunities to introduce significant innovation in the campus infrastructure by boosting performance of the network infrastructure, resulting in enhanced business productivity. Current technologies help boost performance by improving high availability (HA) and prioritizing the campus LAN to become more important to accelerate business growth. For example, bandwidth-hungry applications such as web-enabled video applications and emerging technologies such as Unified Communications increasingly exist in a modern enterprise campus.

In addition to rolling out unified communication and mobility applications, enterprises are also looking to consolidate infrastructure to simplify operations and lower the total cost of ownership (TCO). Existing campus infrastructure solutions cannot meet the requirements needed to provide secure and reliable high-performance access for campus users, nor do they provide the centralized management capabilities critical for reducing costs and streamlining operations.

A new campus LAN design that meets connectivity, security, and centralized management challenges while enabling key IT initiatives, is now required. It must also scale, offer operational simplicity, maintain high-performance, and be flexible and adaptable in order to accommodate new computing trends, without requiring an entire redesign.

Today's major trends impacting the campus network include:

- Unified communications
- Bandwidth-hungry applications
- User productivity
- Risk mitigation and compliance

Table 1 defines these components.

Table 1: Major Component Trends and Definition

Major Trends	Definition
Unified communications	A unified communication solution is typically deployed centrally, where call servers and unified messaging servers are located in a data center. Campus and/or headquarters are usually where major voice over IP (VoIP) devices reside with high bandwidth connections to the data center through a core connection or through a private WAN. The branch office is connected to the core using a private WAN or Internet.
Bandwidth-hungry applications	Many new unified communications applications require more bandwidth. Many popular business applications such as Oracle, SAP, PeopleSoft, and video conferencing have introduced Web-enabled versions that require, in some instances, more than 10 times the bandwidth of their LAN-based counterparts. This has seriously impacted performance, reliability, and availability. Other activities, such as data backup to local servers, can also be bandwidth intensive. However, these activities can be scheduled to take place during times of low usage to lessen their impact on the network.
User productivity	User productivity increases as network performance and accessibility improve. The campus network should be leveraged with services such as wireless coverage and remote access to maximize productivity.
Risk mitigation and compliance	Critical campus resources should not only be protected from external threats but from internal threats as well. This protection should cover large, multiple LANs and provide high-performance capabilities in unison with LAN/WAN accessibility.

Scope

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

- Network connectivity
- Network security
- Network management

In addition, this document provides design guidance for the campus LAN, WAN connectivity, security, and management, and also focuses on the following network devices:

- Routers
- Switches
- Firewalls
- Intrusion prevention systems
- VPN access devices
- WAN acceleration products

Target Audience

This guide is intended for the following audiences:

- Network architects evaluating the feasibility of new approaches in enterprise network design
- Network engineers and operators designing and implementing new campus networks
- Technologists researching and analyzing new approaches for implementing flexible robust networks

Enterprise Campus Network Design Considerations

This section summarizes some of the technical considerations for designing a modern campus network.

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

The critical attributes for designing a current campus network with extreme availability and superior performance are as follows:

- Connectivity—ubiquitous connectivity to disparate sets of resources
- Security—security and compliance
- Management—centralized policy and control
- Visibility—not only in network traffic and security events, but also in application traffic
- Quality of service (QoS)—for real-time applications such as unified communications and other critical applications
- High availability—to ensure business continuity

Connectivity—Ubiquitous Connectivity to Disparate Sets of Resources

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

- WAN connectivity to enable campus users to access the campus network applications and Internet connectivity
- Superior speed for campus network backbone connectivity, data replication, business continuity, and use of technologies such MPLS

The campus LAN hosts a large population of end users that require high-speed and highly available network connectivity to the resources residing at the data center and the Internet. In addition, there can be multiple LAN segments and networks deployed that differ in security and capacity levels and other services offered.

Security —Security and Compliance

The campus security architecture must employ layers of protection from the network edge through the core to the various endpoints for in-depth defense. A layered security solution protects critical network resources that reside on the network. If one layer fails, the next layer will stop the attack and/or limit the damages that can occur.

This level of security allows IT departments to apply the appropriate level of resource protection to the various network entry points based upon their different security, performance, and management requirements.

Layers of security that should be deployed at the campus network include:

- Denial of service (DoS) protection at the edge
- Firewalls to tightly control who and what gets in and out of the network
- VPN to protect internal communications
- Intrusion prevention system (IPS) solutions to prevent a more generic set of application layer attacks
- Insider threat protection through network access control

Management—Centralized Network Policy and Control

Policy-based networking is a powerful concept that enables devices in the network to be efficiently managed —especially within virtualized configurations—and used to provide granular network access control. The policy and control capabilities should allow organizations to centralize policy management and offer distributed enforcement. The centralized network policy and control management solution should ensure secure and reliable networks for all applications and users by providing appropriate levels of access control, policy creation and management, and network and service management.

Visibility

It is important to have visibility into network traffic and security events to effectively maintain and manage resources. It is also critical to collect IP traffic flow statistics to give enterprises insight into data flow, resource utilization, fault isolation, capacity planning, tuning, and offline security analysis. WAN utilization and user-level visibility can help IT better support application performance by leveraging network services and other resources. Security visibility is crucial to granularly view security events in order to help determine how these events are handled.

Extending this visibility to develop a deeper understanding of application-specific traffic is crucial for understanding a wide range of operational and performance information that can impact the users of these applications. For example, specific compression and acceleration technologies can be applied at the network layer to accelerate email applications such as Microsoft Exchange. Another example is preventing employee's access to services such as YouTube and social networking sites from impacting business applications. Understanding these applications and enforcing policies based on the application ensures that business-critical applications meet or exceed the performance expectations of end users.

Quality of Service

To truly assure application experience over large campus networks, QoS is a key requirement. It is critical to assign and manage QoS levels in order to ensure satisfactory performance of the various software applications including unified communication applications that are sensitive to jitter, packet loss, and latency. A minimum of three levels of QoS (each determines a priority for applications and resources) are as follows:

- Real-time
- Business critical
- Best effort

This is especially critical with voice and video deployments, since QoS can mitigate latency and jitter issues by sending traffic along preferred paths or by enabling a fast reroute to anticipate performance problems or failures. The campus network design should allow the flexibility to assign multiple QoS levels based on end-to-end assessment and allow rapid and efficient management to ensure end-to-end QoS for the enterprise.

High Availability

High availability disaster recovery is a key requirement for the campus network and must be considered not only by what is happening within the campus network, but also connections to critical off-campus resources like data center locations. Network high availability should be deployed by using a combination of link redundancy (both external and internal connectivity) and critical device redundancy to ensure uninterrupted network operations and business continuity. Moreover, devices and systems deployed within the confines of the campus network should support component-level high availability, such as redundant power supplies, fans, and routing engines. Another important consideration is the software/firmware running on these devices, which should be based on a modular architecture that provides features such as unified in-service software upgrades (ISSUs) to prevent software failures/upgrade events from impacting the entire device. Software failures/upgrades should only impact a particular module, thereby ensuring system availability.

Campus Architecture Overview

Juniper Networks delivers a proven IP infrastructure for the campus that meets these challenges, enabling the performance, scalability, flexibility, security, and intelligence needed to not just meet, but increase campus user productivity. Juniper Networks offers flexible configurations and price points that meet the needs of all campuses, while delivering high-performance throughput with services such as firewall, adaptive threat management, VPN, MPLS, IPV6, and Connectionless Network Service (CLNS).

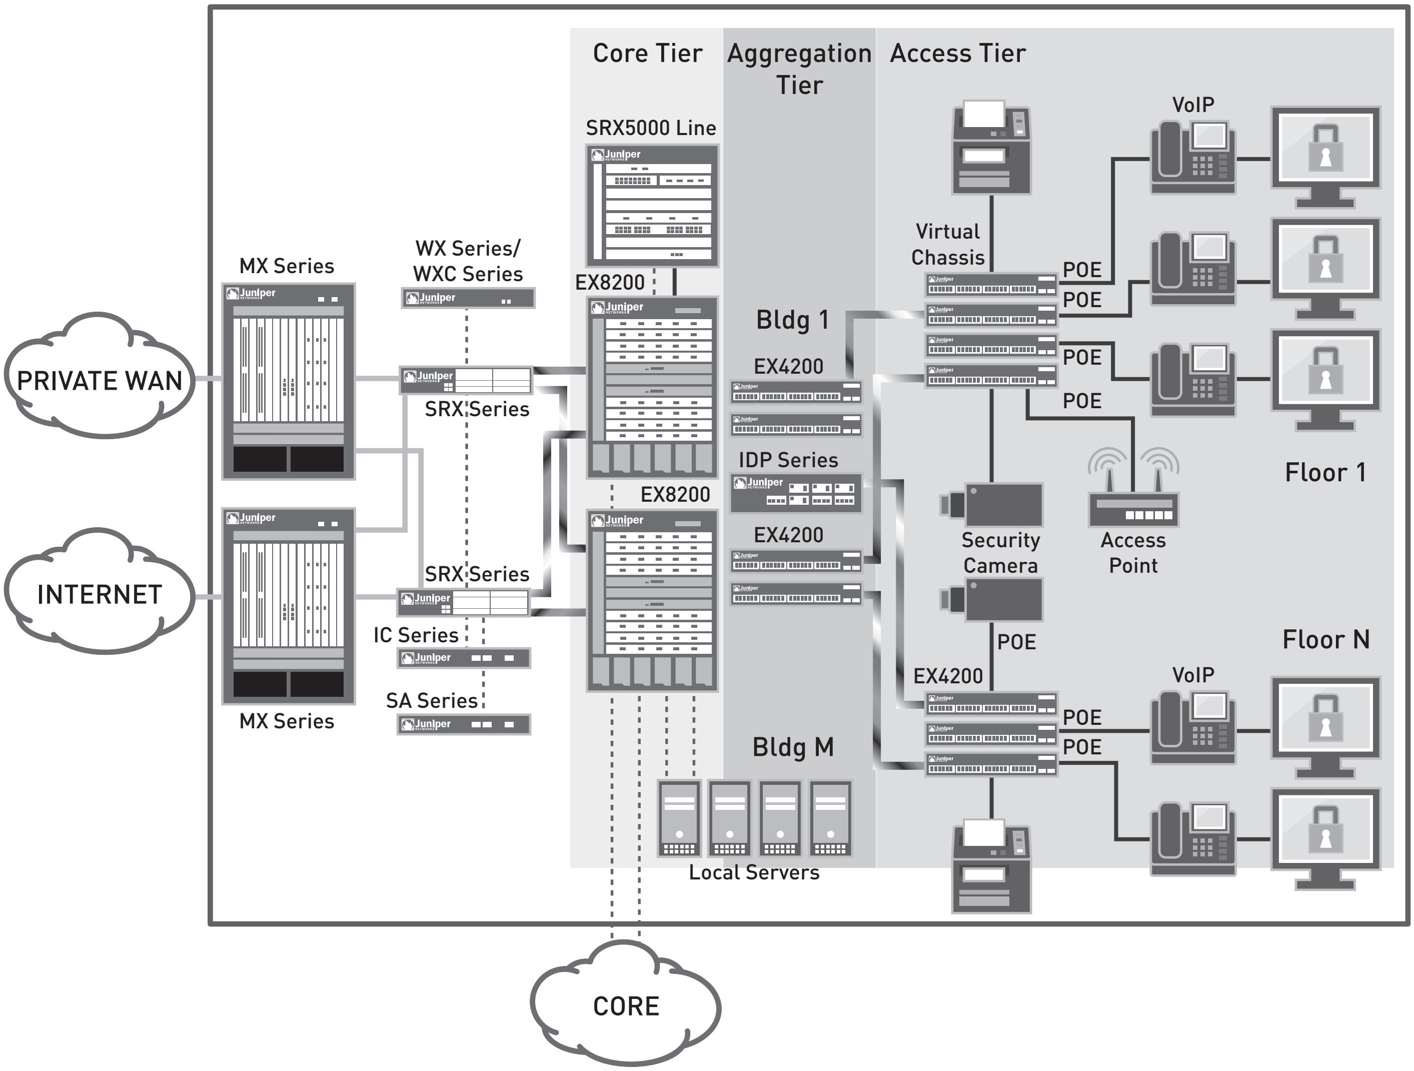

Figure 1: Juniper Networks campus reference architecture

Network Connectivity

Pertaining to connectivity, a Juniper Networks JUNOS® Software-based routing and switching infrastructure is an "always on" network infrastructure that provides security, reliability, and cost effectiveness by lowering the total cost of ownership.

LAN Connectivity

Layered Approach

An enterprise campus LAN architecture may span up to three functional layers, from desktop devices connected to wiring closet switches at the access layer, to the core layer at the center of a large campus LAN. The hierarchical topology segments the network into physical building blocks, which simplifies operation and increases availability. Each layer within the hierarchical infrastructure has a specific role. Figure 2 illustrates the layered approach.

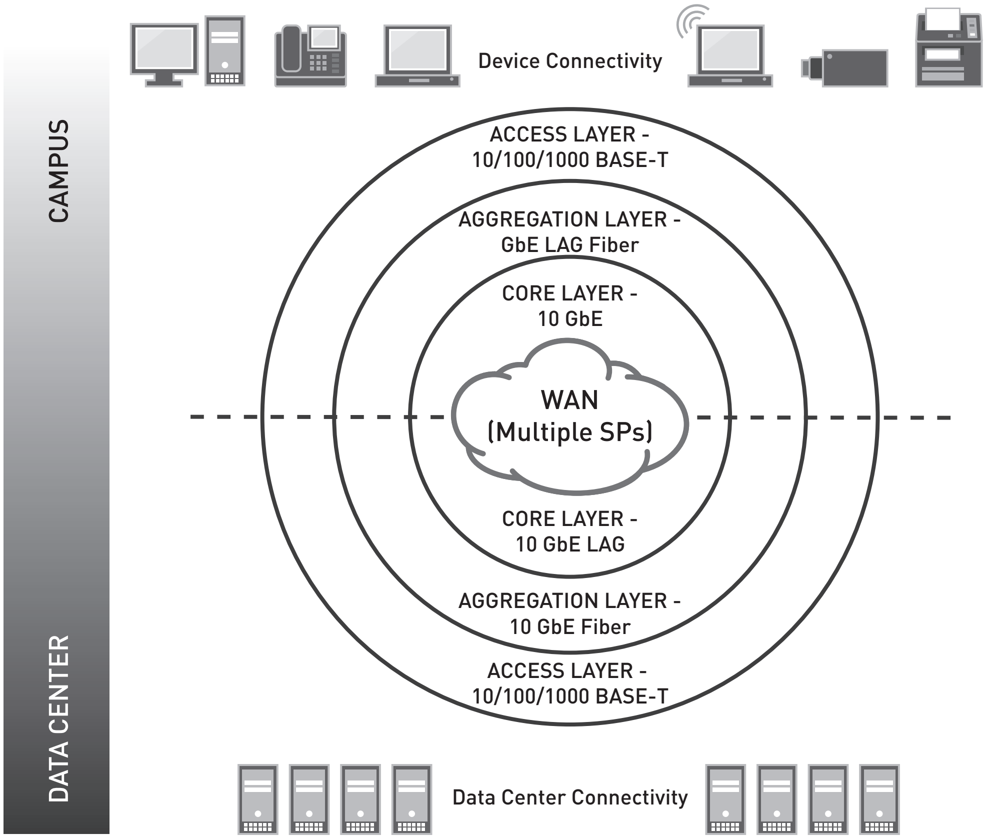
CAMPUS
DATA CENTER
Device Connectivity
ACCESS LAYER - 10/100/1000 BASE-T
AGGREGATION LAYER - GbE LAG Fiber
CORE LAYER - 10 GbE
WAN (Multiple SPs)
CORE LAYER - 10 GbE LAG
AGGREGATION LAYER - 10 GbE Fiber
ACCESS LAYER - 10/100/1000 BASE-T
Data Center Connectivity

Figure 2: The layered approach

The access layer provides an access control boundary and delivers network connectivity to end devices (client machines, printers, IP telephony, and cameras) in a campus.

The aggregation layer aggregates connections and traffic flows from multiple access-layer switches, providing a core enforcement perimeter as it delivers traffic to core-layer switches.

The core layer provides secure connectivity between aggregation-layer switches and the routers connecting to the WAN and the Internet, which enable business-to-business collaboration.

Benefits and Challenges to the Layered Approach

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

This three-layered approach traditionally requires additional hardware and can be costly to configure, deploy, and administer. Based on port density requirement and geographical distribution of the campus, the three functional layers can be collapsed into two layers. Networks have previously attempted to address emerging bandwidth, throughput, and port density requirements. As a result, these networks have grown bloated with extra layers of inefficient, ill-suited legacy hardware that not only fails to meet these needs, but added considerable management complexity, reduced network availability, and driven up capital and operational expenses.

Access Layer

In a campus network, the access layer provides network connectivity to end users by connecting devices such as PCs, printers, IP phones, and CCTV cameras to the corporate LAN through wired or wireless LAN (WLAN) access points. Access-layer switches typically reside in the wiring closets of each floor in each campus facility.

A campus network architect, who is considering integrating unified communications at the access layer, should concentrate on the following attributes:

- Port density
- Flexibility
- Scalability
- High availability
- Power over Ethernet (PoE)
- QoS
- Segmentation
- Security infrastructure integration

To meet primary connectivity requirements, one of the most important aspects for access layer devices are port density for client connection, as well as an uplink to the aggregation/core layers to reduce the client-to-uplink oversubscription ratio. Also, the flexible scalability on a need-to-grow basis is important to reduce capital and operating expenditures. Virtual chassis technology provides high port density and meets the flexible scalability requirement since the access layer device can be added to the virtual chassis with minimum operational and management costs.

Another important aspect of the access layer is high availability issues such as component level redundancy, power supply, control modules, etc. This can be achieved by a chassis-based or virtual chassis-based access layer solution. The PoE functionality at the access layer simplifies the deployment of unified communication services such as VOIP telephony and CCTV. In addition and pertaining to real-time applications is end-to-end QoS. At the access layer, QoS functionality of classification, marking and prioritization is important for end-to-end QoS support.

The access layer serves as a primary boundary of access control for security requirements. Virtualization capabilities like virtual LANs and virtual routers are important to support required segmentation of the access layer network. In addition, integrating the network security infrastructure with unified access control is another important aspect of the security features at the access layer. As a first line of defense, security controls such as broadcast storm control, (Dynamic Host Configuration Protocol (DHCP)-snooping, and Address Resolution Protocol (ARP)-spoofing protection are important features at the access layer. With increasing use of multicast applications and multicast feature support, such as Internet Group Management Protocol (IGMP)-snooping, IGMP and multicast routing protocol support are also important considerations for access layer devices.

Ideal for employees meeting in dispersed conference rooms or areas other than their offices, as well as a necessity for supporting contractors, partners, and guests, wireless access must be provided across the campus. With the plethora of IP devices currently available on the market and used in the workplace, especially by unknown guests, a comprehensive security policy must ensure that only trusted devices access the campus network. Furthermore, the appropriate LAN resources must be restricted and made available only to people with the proper credentials. This is especially true for contractors, partners, and other guests. Seamless coverage enabling a user to roam the campus with the same login credentials is also expected.

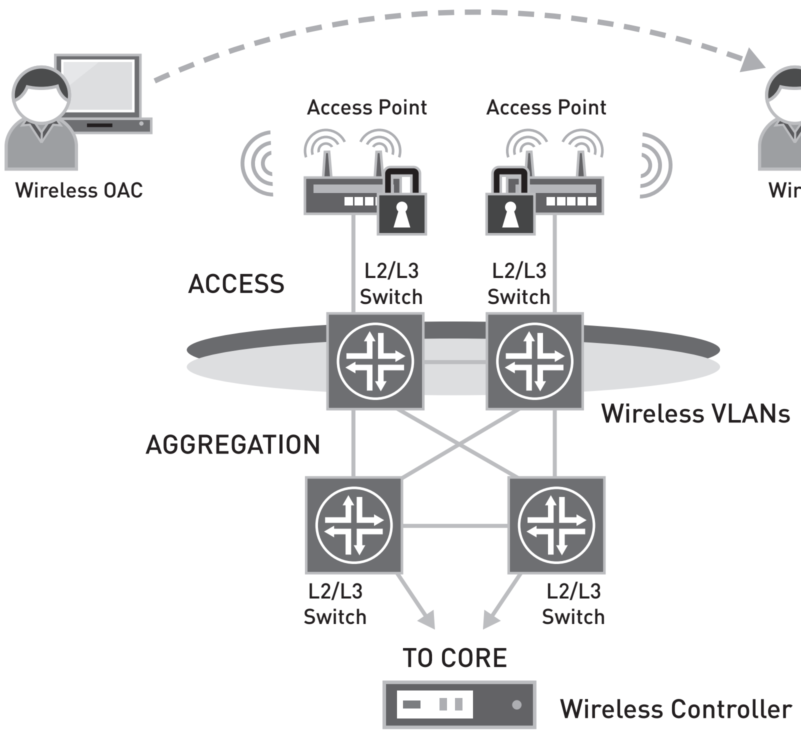

Figure 3: Flexible and roaming wireless access solutions

There are two main designs for flexible and roaming wireless solutions:

- Non-controller-based wireless access
- Controller-based wireless access

Non-Controller-based Wireless Access

In this design, an 802.1Q trunk for access point to switch is required. Roaming requires spanning at least two VLANs between access layer switches.

Controller-based Wireless Access

This design uses a virtualized and centralized wireless controller. Access point VLANs are placed local to the access switch. Roaming does not require spanning VLANs across the campus network. Campus facilities are likely to also have:

- IP phones
- PoE for security cameras and WLAN devices
- Support for multicast protocols such as IGMP snooping and QoS
- Access layer for connectivity–basic security functionality, QoS, link aggregation, and connection
- Component level high availability with redundant power source

Juniper Networks EX Series Ethernet Switches infrastructure is recommended at the access layer. Customers can choose Juniper Networks EX3200 Ethernet Switch, Juniper Networks EX4200 Ethernet Switch with Virtual Chassis capability, or the Juniper Networks EX8200 line of Ethernet switches for appropriate port density requirements.

Aggregation Layer

The aggregation layer aggregates connections and traffic flows from multiple access layer switches to provide high-density connectivity to the LAN core. The following attributes should be considered for aggregation layer design:

- Scalability
- High-performance and throughput
- HA
- Network services integration
- QoS support

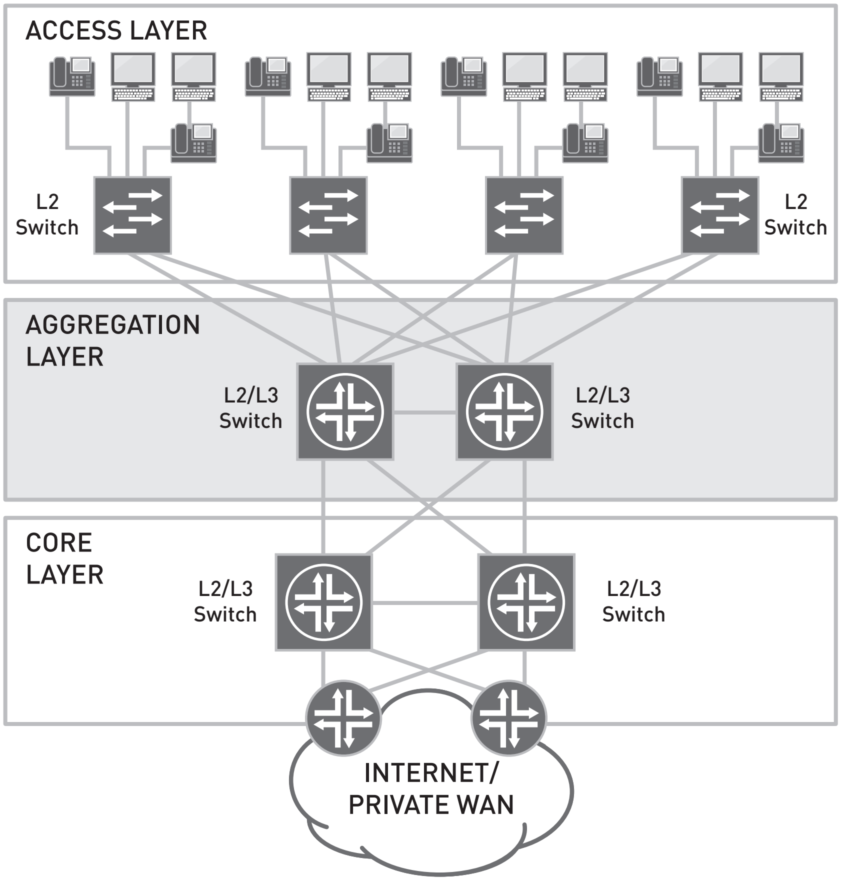

Figure 4: Aggregation layer in a highly available campus LAN

Due to their location in the network, aggregation-layer switches must offer high-density ports to provide maximum scalability at the aggregation layer, along with wire-rate forwarding for maximum throughput. Also, a nonblocking architecture at the aggregation layer is important to minimize the oversubscription ratio since a large number of client connections are supported through the aggregation layer. Therefore, it is critical to have high availability hardware and software features that deliver high reliability and robustness (further details on high availability features are covered in the next section). For device-level redundancy, the aggregation layer devices should be deployed in pairs to serve access-layer devices. The primary function of the aggregation layer infrastructure is to provide high throughput and nonblocking switching/routing fabric. The dynamic routing protocol support, high-performance control plane, and high capacity data plane are important features of aggregation layer devices.

Because the aggregation layer is not as distributed as the access layer, it is an ideal place to locate your security defenses and an ideal place to create segments using virtual routers or VLANs to contain threats. The network service integration capabilities, such as segmentation using a VLAN or a virtual router and traffic redirection capability with required policy control, are important aspects of aggregation layer devices. To support real-time application and prioritize critical application and control traffic QoS capabilities, such as multiple queues, queue capacity and integration with end- to-end QoS infrastructure should also be evaluated. In order to support multicast applications, multicast routing protocol and efficient multicast replication techniques are important factors for aggregation layer devices.

Core Layer

The core layer provides a fabric for high-speed packet switching between multiple sets of aggregation devices, or the access layer devices in a collapsed aggregation/core layer deployment. The core layer serves as the gateway where all other modules meet, such as the WAN edge. Functionally, the core layer is where high speed connections to all campus networks occur —such as different buildings, departments, and server areas—and connects these to a perimeter or WAN edge network as displayed in Figure 5.

The following attributes should be considered for core layer design:

- High-performance
- High throughput
- High availability

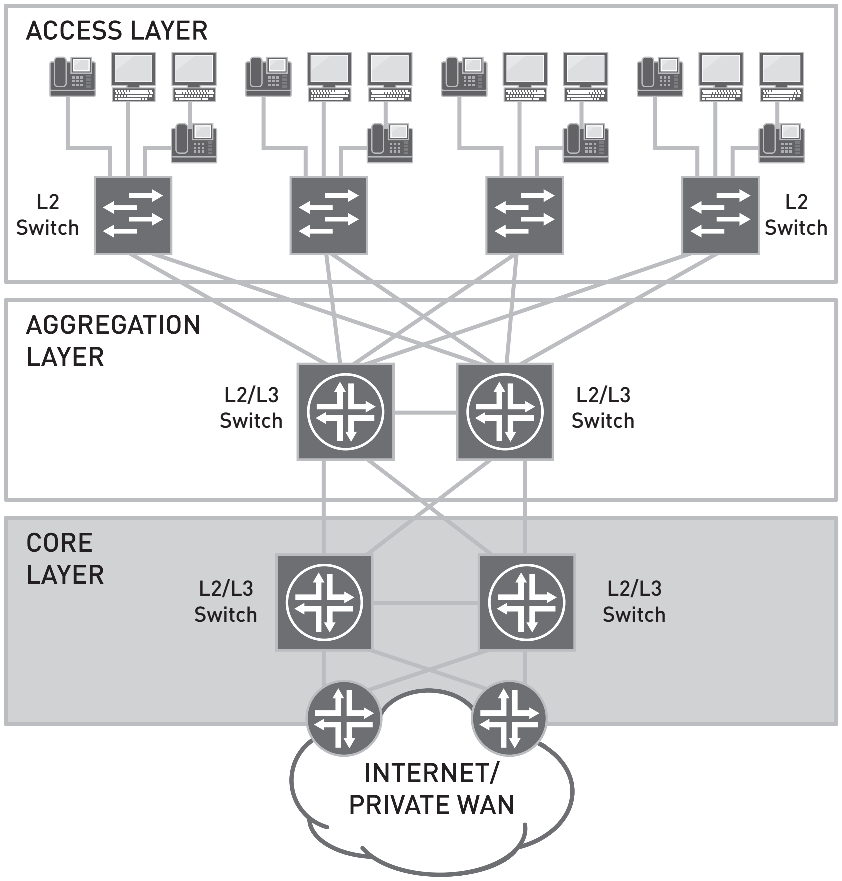

Figure 5: Core layer in a highly available campus LAN

As the name implies, the core layer serves all campus users and therefore any failure at the core layer should be minimized. High availability software features such as unified ISSU, nonstop forwarding/routing, graceful restart capabilities, and a modular operating system design should be enforced to limit the impact of any module failure. For link bandwidth and redundancy, core connections should deploy aggregated links in multiples of 10-Gigabit Ethernet connections from aggregation layer devices.

Core layer devices should be deployed in pairs corresponding to each aggregation layer device. Device level redundancy capabilities like redundant power supply, fan modules, control modules, and switching fabrics are required at core layer devices.

Since any performance degradation at the core layer affects the entire campus network, high-performance, non-blocking switching/routing architecture is extremely important.

Depending on the size of the network, the aggregation and core layer functionality can be collapsed within one set of devices, since it reduces capital and operating expenditures and reduces latency of the traversing traffic. The integration of network services at the collapsed core/aggregation layer should have a minimum impact on performance of core layer devices.

WAN Connectivity

WAN connectivity provides a vital link from the campus to centralized services and resources. Designing and scaling a campus LAN for assured network connectivity and performance is a challenge that every high-performance organization faces.

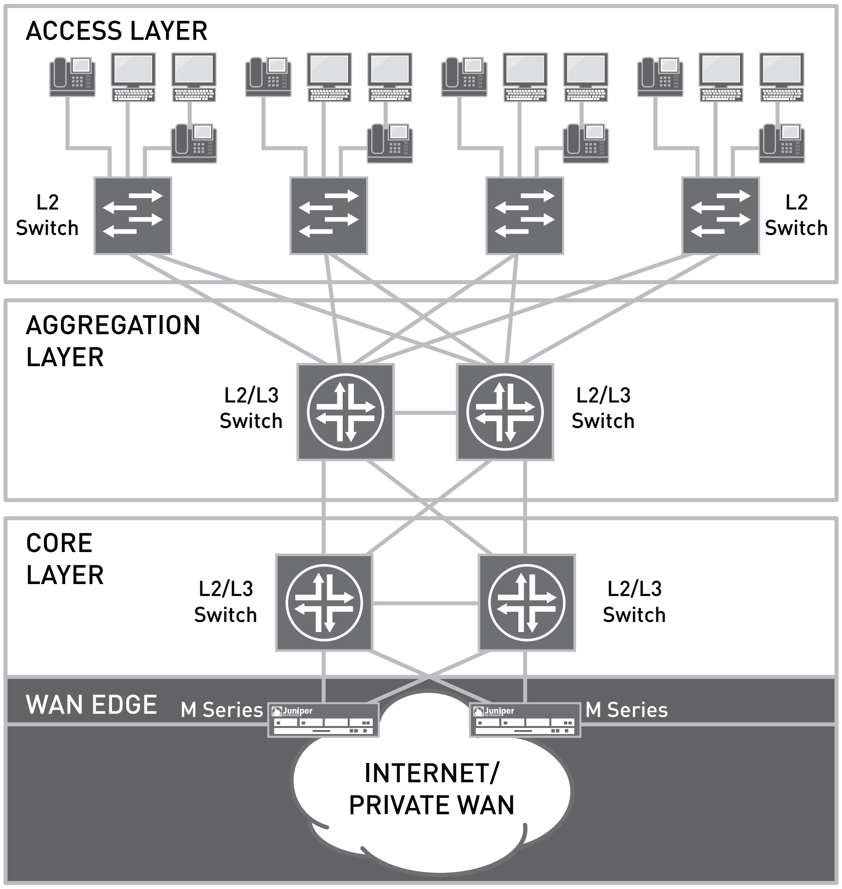

Figure 6: WAN edge in a highly available campus LAN

WAN Design Considerations

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

All WAN edge devices must provide a full complement of high availability 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 high availability using dynamic routing protocols with minimal convergence times. Also, an alternate connection to the Internet or WAN must be maintained.

Secure and optimized voice services should be provided at the WAN edge to enable effective communications across the LAN and WAN. Either an integrated or standalone VoIP gateway may be implemented.

WAN Acceleration

Adding more bandwidth does not automatically deliver LAN-like performance across the WAN. Acceleration services are needed to optimize performance of centralized applications across the WAN at all times, even when bandwidth is constrained.

WAN optimization products seek to accelerate a broad range of applications accessed by distributed enterprise users through eliminating redundant transmissions, staging data in local caches, compressing and prioritizing data, and streamlining chatty protocols like CIFS.

High Availability in the Campus Network

Downtime should not be considered an option in today's campus network with some customers requiring at least 99.999 percent reliability. Any failure in a campus network can impact the availability of business critical applications, resulting in reduced user productivity and corporate revenues.

It is critical to have a robust architecture to minimize the number of failures in a campus network. Network architects should consider high availability of the campus network to ensure that any failure that occurs has minimal impact on application accessibility. Network architects should also strive towards achieving high resiliency similar to critical carrier class networks. In the event of a rare failure, the device and network architecture should have built-in resiliency features to minimize or avoid any disruption to the network services. This ensures that even if a failure occurs, users' business application accessibility is not impacted.

As explained earlier, the campus network consists of switches, routers, security devices, etc. connected by network links. Architects and designers should consider failures in the following four major categories:

- Device component-level failures
- Network link-level failures
- Network software failures
- Network link-level failures

Device Component-Level Failure

Regarding critical components, Juniper recommends the device architecture should support component-level redundancy to minimize the impact of component failure on device functionality. Power supply modules, cooling modules, control modules, and switching fabrics are examples of such critical components. High availability is achieved by having an N+1 hot standby module. Hot-swappable and hot pluggable capabilities ensure that device maintenance does not impact network services. Such component-level redundancy ensures that critical component failure does not result in device failure and keeps the failure transparent to the users, while avoiding control plane and data plane convergence.

Network Link-Level Failure

Network links connect all devices to the network and to one another. Link-level high availability ensures that business processes maintain vital data flow through internal and external resources.

Since all traffic flows through network links, any failure in the physical link should be transparent to maintain high availability of network services. This can be achieved by connecting devices using multiple links and using link aggregation technologies, as shown in Figure 7. All links can be active and provide high bandwidth and all uplink connections between devices on different functional layers should use link aggregation group (LAG). Devices in aggregation and core layers should also use LAG.

Ideally, the physical routing path of the links in a LAG should be different to minimize impact of any physical failures. The links can also be terminated on different line cards located on each device to minimize the impact of line card failures. Link aggregation ensures that a physical link failure does not result in data plane convergence and hence keeps the failure transparent for application accessibility.

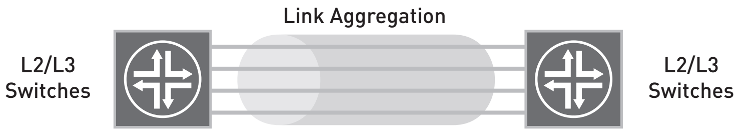

Figure 7: Link aggregation group (LAG)

Network Software Failure

Network software is a critical component for device functional operations. A modular network software system ensures that failure of one module does not impact the functionality of all other modules. Network software should support features such as nonstop forwarding, graceful restart, and unified ISSU to ensure that software module failure is transparent to the users.

JUNOS Software's modularity and uniform implementation of all features enables even the smallest campus to benefit from the same hardened services in devices running JUNOS Software as compared to the largest service providers.

Network Device Failure

A network device failure can be addressed by deploying 1+1 redundant network devices and having partial or full mesh connectivity among such devices. Redundant network devices can be deployed in an active/active configuration where both devices load balance the traffic. It can also be an active/passive deployment where one device serves as a hot standby for the active device. The traffic convergence is achieved by Layer 2 link redundancy techniques such as spanning tree, redundant trunk group, etc., and Layer 3 dynamic routing protocols. A device failure may result in control and data plane convergence. Therefore, the design should evaluate the convergence time for failure recovery. Below are a few options revealing the benefits of full and partial mesh connectivity.

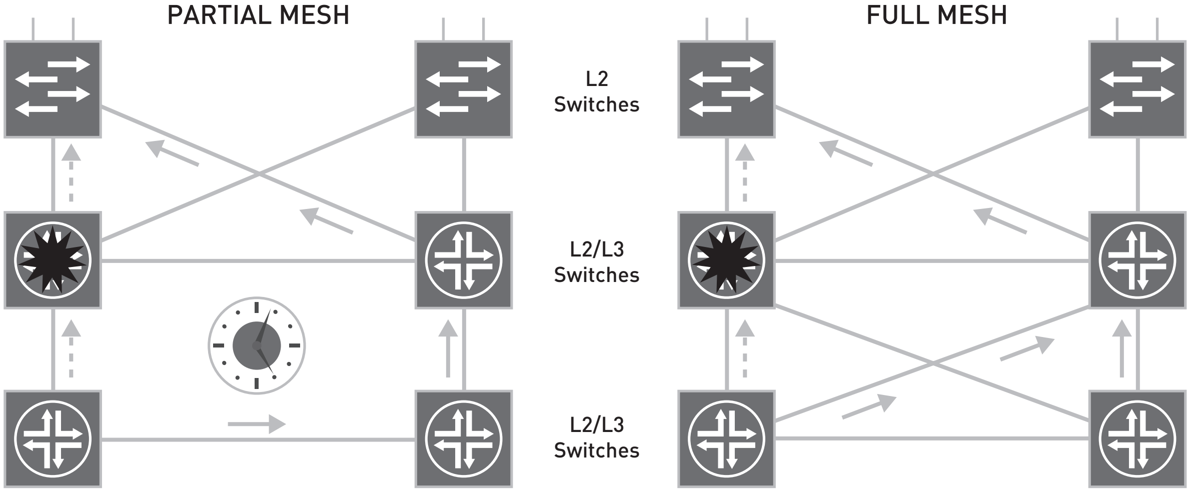

Figure 8: Full mesh versus partial mesh connectivity

Partial Mesh Configuration

In this design, a Layer 3 peering square is configured between the aggregation and core layers. Route peering provides a redundant path. Link failure requires Layer 3 protocol convergence, which may vary since the route is non-deterministic. The result of this deployment is dropped sessions and/or lost packets, which delivers sub-optimal performance.

Full Mesh Configuration

In this design, a Layer 3 dual-homed triangle is configured between the aggregation and core layers. Equal-cost multipath (ECMP) provides redundant, load-sharing path. Any link failure results in a fast failover time since the route is deterministic. The result is optimal performance with minimal packet loss.

Best Practices for Campus Network High Availability

Putting this all together, Juniper recommends the following high availability configuration, as shown in Figure 9.

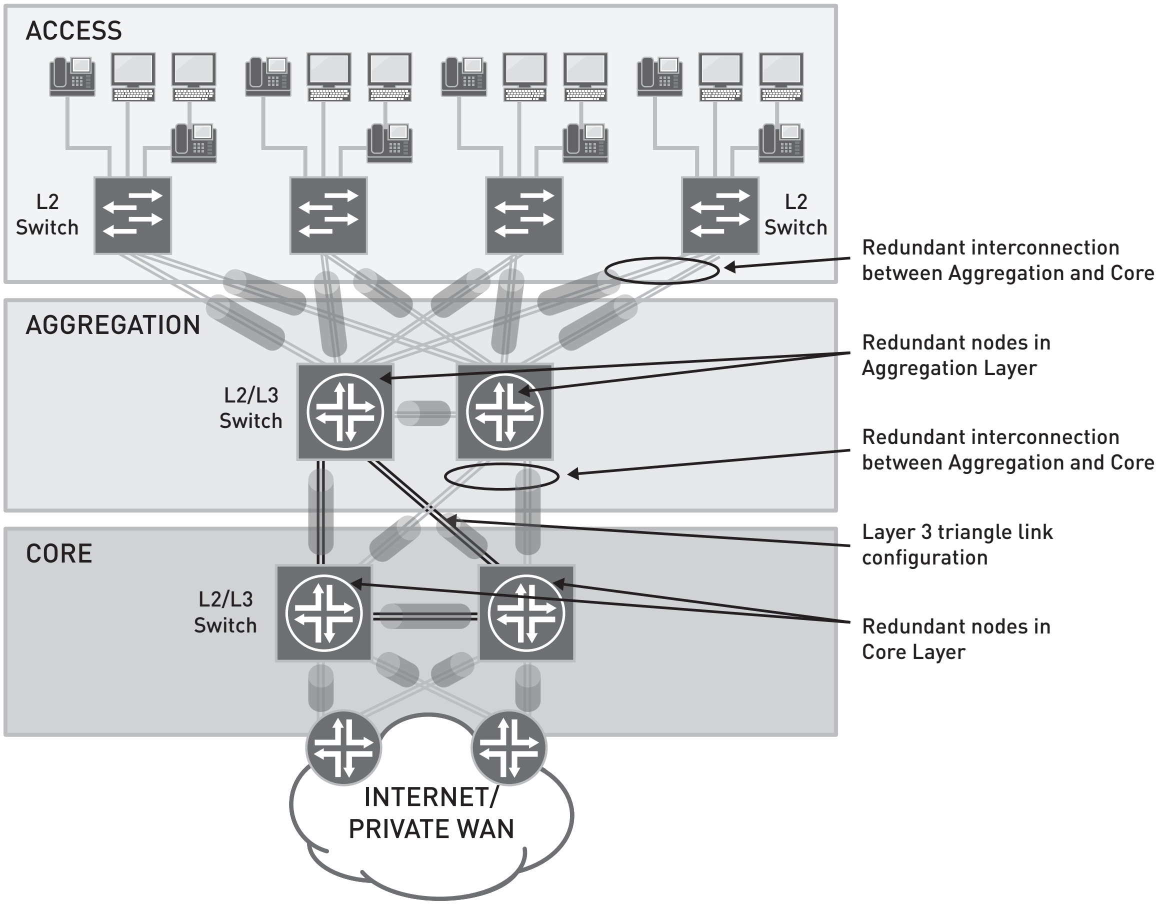
ACCESS
L2 Switch
L2 Switch
Redundant interconnection between Aggregation and Core
AGGREGATION
Redundant nodes in Aggregation Layer
L2/L3 Switch
Redundant interconnection between Aggregation and Core
CORE
Layer 3 triangle link configuration
L2/L3 Switch
Redundant nodes in Core Layer
INTERNET/ PRIVATE WAN

Figure 9: Campus network high availability

The access layer switches should be "dual-homed" to redundant nodes in the aggregation layer. The aggregation and core layers are both built with dual-homed interconnects. Each alternate path uses Layer 3 for optimal convergence. The core layer switches are also dual homed to WAN edge routers. At all layers, link bandwidth and node capacity are designed to withstand link or node failure.

Network Security

Pertaining to security, Juniper Networks offers a JUNOS-based security policy, enterprise-wide access control and Juniper Networks Adaptive Threat Management Solutions that extend to distributed enterprise locations like the campus. This integrated, multi-functional security approach provides the right-sized fit for all campus sizes.

Security

Campus LAN security issues are intensified by the increased mobility of campus network users, the growing utilization of contractors, the co-location of partners on site, visiting guests, the proliferation of unified communications, and the demand for wireless access.. IT must protect valuable campus resources from internal and external threats across large or multiple LANs as it delivers high-performance with secure and ubiquitous LAN and WLAN access.

Increasing security threats and risks force campus LANs to remain secure and controlled on all fronts while providing open and pervasive access to maintain and increase productivity. The most effective security architecture to ensure maximum protection from network and application layer threats is based on multi-layered protection that's appropriate for each location of the network. Holistic solutions that offer comprehensive security features, proven reliability, and exceptional performance are needed. 802.1X and network access control should be used to effectively handle unmanaged devices and guest users attempting network access, as well as to support unmanageable devices, post admission control, application access control, visibility, and monitoring. Firewalls and intrusion prevention systems are also needed to help ensure security across the LAN. In addition, QoS can be used as a security tool to identify, classify, and queue traffic. For example, QoS policies can protect access to departmental resources or ensure that high-priority data flows are unaffected by malicious traffic.

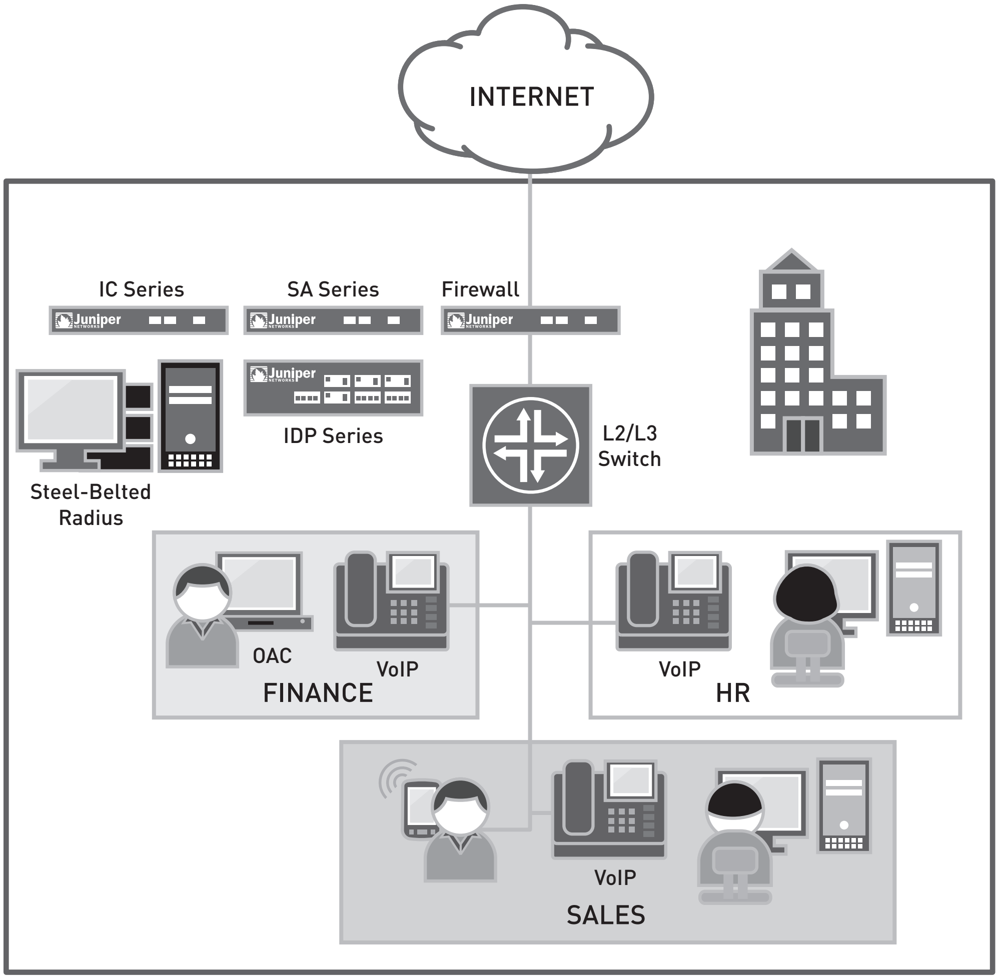

Figure 10: Security architecture in a campus environment

A multilayered security architecture facilitates network configuration by providing a modular design that can rapidly and economically scale based on the number of users in a campus environment. It also creates a flexible network where new security services can be easily added without total redesign.

The basic idea behind multilayered security architecture is to protect the "crown jewel" (data center resources) with multiple layers of defense, where if one should fail, another will provide crucial protection. Another important thing to remember is that everything cannot be defended, so our layered defense approach should be asset-centric rather than perimeter or technology-centric. While focusing on an asset-centric layered defense approach is clearly important, we must not forget to protect users who access those assets as well. Therefore, we must protect the end user from not only external threats but internal ones as well. This means that the endpoint must be secure at all times.

Access Control and Segmentation

The most vulnerable and most desired targets for attack in a campus environment are the endpoints themselves. Therefore, an initial line of defense is required to monitor who (and what) is coming in and out of the wired/wireless network. Authentication and access control should be in place to discourage opportunistic attacks from outsiders.

Authentication and authorization answers two very important questions: "who" is entering the network and "what" service is being delivered, respectively. Once the user and service is verified, the experience delivered for the application/service can be varied per user based on user subscription and profile. Device health and location data is then determined in order to deliver granular access control.

Holistic network access control should be deployed with support for all access technologies (wired, wireless, or remote access) so that only authorized users and applications from devices that adhere to your network security policies are permitted through the first layer. Endpoints (hosts) should be authenticated when they initially connect to a LAN. Authenticating endpoints before they receive an IP address from a DHCP server prevents unauthorized endpoints from gaining access to the LAN. Network access control should provide both standards-based 802.1X port-level access and Layer 2 through Layer 4 policy enforcement based on user identity.

To achieve differentiated role-based access from internal networks, enterprises should segment the network and the logical control points should be defined to control access to critical data, as well as contain any threat within as small segment of the network as possible. Access port security features such as dynamic ARP inspection, DHCP snooping, and Media Access Control (MAC) limiting should be leveraged to harden the access layer.

The network access control solution should combine user identity, device security state, and network location information for session-specific access policy by user and leverage the existing network infrastructure. The network access control should deliver comprehensive control, visibility, and monitoring, as well as be standards-based, reduce threat exposure, and decrease access control deployment costs and complexity. It should also be adaptable and scalable to meet the network access control requirements for campuses of any size.

Campuses typically have a number for visiting guests and contractors accessing the network from outside on a daily basis. Because of this, the network access control solution should address the common problem of how to provide appropriate access to temporary guests by using a web interface. Guests can be granted customizable, limited time access privileges on the network during the duration of their stay.

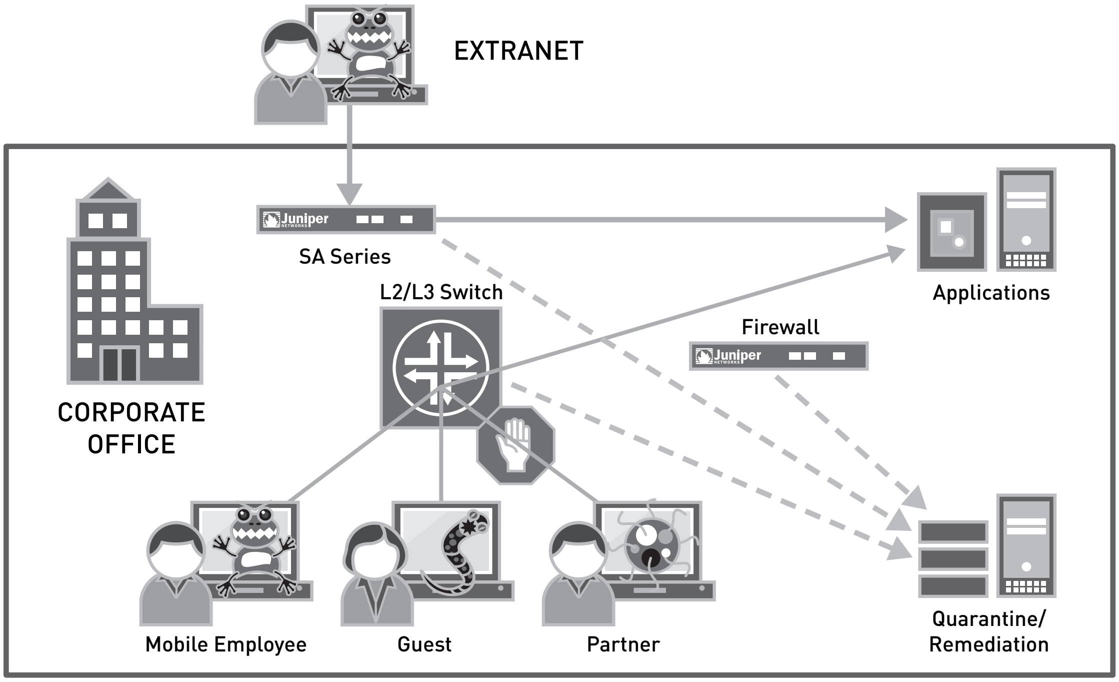

Figure 11: Enforcing endpoint health policy for all user types

The network and security infrastructure (switches, routers, wireless access points, firewalls) should integrate with inventory management and existing AAA systems, as well as network management and monitoring frameworks to gain unprecedented real-time visibility into the campus security environment.

For more information, refer to **www.juniper.net/us/en/products-services/security/uac/.**

Stateful Firewalls and Router-based Security

In this ever-changing threat landscape, smarter, more sophisticated attacks have the ability to penetrate the above mentioned lines of defense. Thus, a sturdy firewall with stateful inspection is necessary as an added layer.

These firewalls provide stateful inspection of traffic traversing different network segments. They also should be able to create VPNs using IPsec for authenticating and encrypting IP packets, Transport Layer Security (TLS), and SSL VPN capabilities to provide critical protection against DoS, distributed denial of service (DDoS), and other types of attacks deployed at the perimeter.

Firewalls must be able to scale to handle drastically increasing volumes of traffic flow when deployed at the network perimeter or at the core, so the network's performance is not negatively impacted during spikes.

There are several distinct features of firewall security:

- Scalable performance: The ability to leverage new services with appropriate processing capabilities without sacrificing overall system performance
- System and network resiliency: Carrier-class reliability
- Interface flexibility: Highly flexible I/O configuration and independent I/O scalability
- Network segmentation: Security zone, VLANs, and virtual routers allow administrators to tailor security and networking policies for various internal, external, and demilitarized zone (DMZ) subgroups
- Robust routing engine: Carrier-class routing engine provides physical and logical separation of data and control planes to allow deployment of consolidated routing and security devices and ensure the security of routing infrastructures
- Comprehensive threat protection: Integrated security features and services include a multi-gigabit firewall, intrusion prevention system, DoS/DDoS Detection and Mitigation, NAT, and QoS.

In both wireless and wired campus networks, intelligent routers should be deployed to prevent IP spoofing. On the data plane, routers should perform anti-spoofing by implementing access control lists (ACLs) and IP fragment filtering to drop all inbound traffic with suspicious source IP addresses or IP address ranges.

The Juniper Networks SRX Series Services Gateways are designed to meet the network and security requirements for campus LAN consolidation, rapid services deployment, and aggregation of security services. For more information, refer to http://www.juniper.net/us/en/products-services/security/srx-series/.

Application Layer Security

The most sophisticated network attacks require another logical layer of defense: an intrusion prevention system.

An intrusion prevention system provides important content inspection and antivirus/anti-spam capabilities. Content inspection is designed to stop L7 application attacks and is the only way to detect what is really running on the L7 application or the signaling application layers.

The intrusion prevention system detects unusual or suspicious behavior on the application layer by using customizable signatures based on stateful protocol inspection, attack patterns, and behavioral learning. This capability is vital for enterprises seeking to protect their networks against penetration and proliferation of worms and other malware including trojans, spyware, Keyloggers, and adware.

These systems should be designed to detect the presence of attacks within permitted traffic flow to the network by using stateful signatures that scan for attacks based on known patterns. Stateful signatures need to be easily customizable in order to fit into different provider requirements and specific concerns.

While the intrusion prevention system should sometimes be physically integrated with routing and/or firewall systems, in this document we are focusing on the logical elements of multi-layered security.

For more information, refer to **www.juniper.net/us/en/products-services/security/idp-series/.**

Dynamic Security

A vast majority of current security solutions throughout the networking industry are static in nature. This static nature of the current security solutions is ineffective due to the constantly changing and unpredictable nature of modern attacks. Therefore, it is important that all security layers outlined above should communicate with the policy decision entity for real-time dynamic feedback and new security policies enforcements.

This real-time policy enforcement and correlation of information from all four logical security layers within policy management can make it possible to build dynamic user and service-aware security.

One Goal: Comprehensive Protection

In today's environment of constantly evolving threats, providers require solutions that can protect against unknown and known patterns. Many of the most significant threats involve "zero-day" attacks, or unknown pattern attacks that leverage vulnerabilities where there is no signature or software patch.

Furthermore, while external threats such as trojans, viruses, worms, buffer overflows, and SQL injections are the most publicized, internal threats are often overlooked and may be more common than external threats. Implementing multi-layered security helps to protect against both external and internal threats.

If one of these pieces of a comprehensive, multi-layer security approach is missing, enterprise campus networks are easily vulnerable to a loss of network integrity, revenue, and even corporate reputation.

Juniper Networks Adaptive Threat Management Solutions consist of best-in-class security products that work together and provide this high level of comprehensive protection to enterprise campus networks. Juniper Networks Adaptive Threat Management Solutions cooperate with each other, resulting in the networks' ability to dynamically adapt to changes in the environment without the need for manual intervention and always with an audit trail. This cooperative system also provides the real-time network-wide visibility needed to make the infrastructure more secure and efficient, yielding a significant competitive advantage.

For more information, refer to **www.juniper.net/us/en/solutions/enterprise/security-compliance/adaptive-threat-management/.**

Network Management

Campus Network Management

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

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

Critical requirements for campus Subscriber Identity Module (SIM)/ security information and event management (SIEM) systems are as follows:

- Efficiently manage logs and flows, identify prioritized offenses, and reduce operational complexity with a single appliance
- Compliance support for monitoring, reporting, and auditing processes of regulatory and security standards
- Detailed view into the systems that are available remotely and recommendations for appropriate changes
- Multi-vendor support for major network and security devices for correlation, collection, analysis, and reporting of logs
- Network behavior anomaly detection (NBAD) discovers aberrant activities using network flow data and enhances the ability to identify zero day threats by base lining network traffic patterns

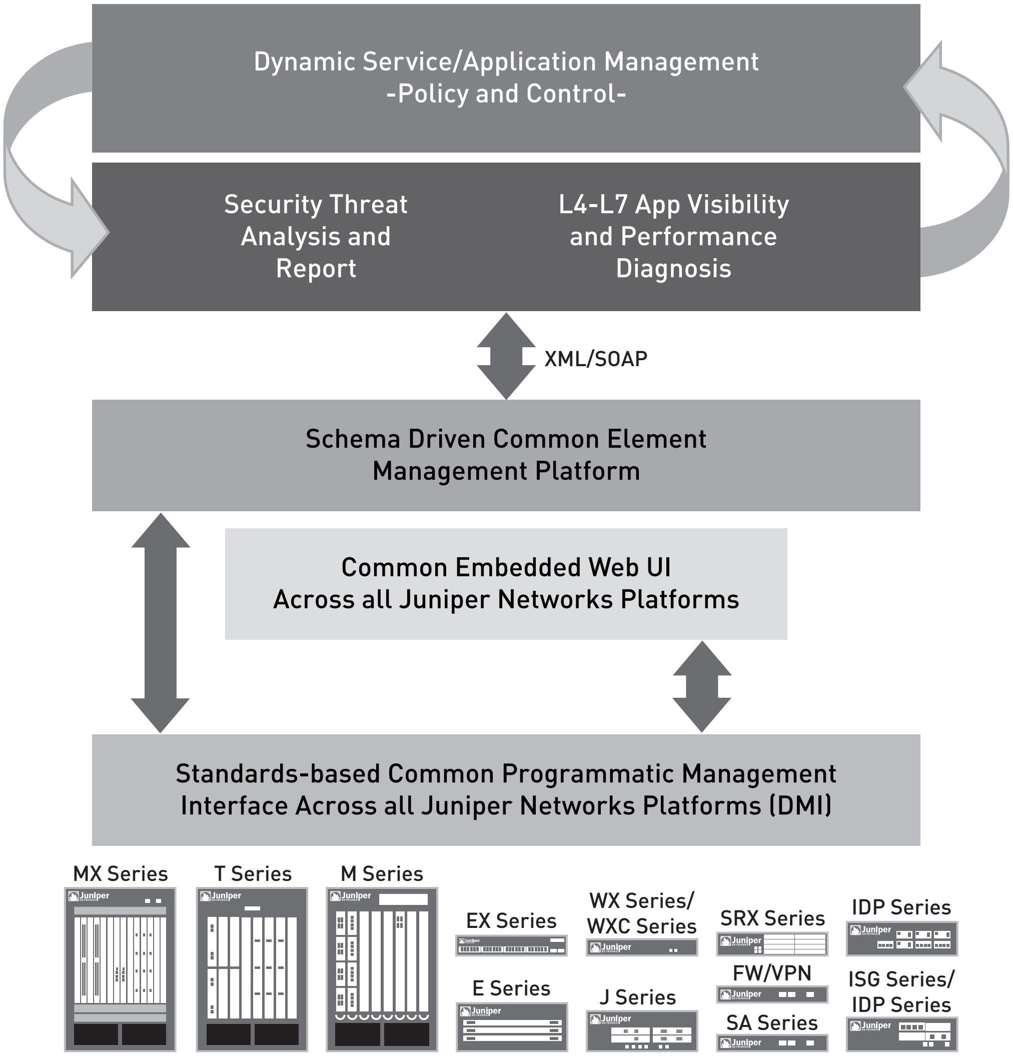
Dynamic Service/Application Management
-Policy and Control-
Security Threat Analysis and Report
L4-L7 App Visibility and Performance Diagnosis
XML/SOAP
Schema Driven Common Element Management Platform
Common Embedded Web UI Across all Juniper Networks Platforms
Standards-based Common Programmatic Management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

Figure 12: Juniper Networks network management framework

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

For more information about Juniper Networks STRM Series Security Threat Response Managers visit **www.juniper.net/us/en/local/pdf/datasheets/1000217-en.pdf**; for Juniper Networks Network and Security Manager visit **www-int.juniper.net/www-int-docs/ft/nm/NSM/110018-014.pdf.**

Conclusion

Today's modern campus network design must be a high-performance, adaptable, and scalable environment in order to meet stringent connectivity, security, and management challenges. Meeting these requirements will accommodate the ever-changing and increasing campus network trends without requiring a re-design of the entire campus network.

In order to help campus architects design and maintain a non-complex, cost-effective campus infrastructure, Juniper Networks offers a campus network solution based on its routers, firewalls, IDP Series/IPS, NAC/Juniper Networks Unified Access Control, and NSM/STRM Series management tools. All these campus network solutions are packaged through JUNOS, a single, consistent operating system that can be used across all campus network switches, routers, and firewall devices.

With this reference architect, designers can understand the principles for designing a campus infrastructure that is easy to deploy, configure, and upgrade. Simplifying the design enables operational efficiencies and allows architects to deploy a campus network that is agnostic to multiple media types. By focusing on the major campus attributes addressed in this reference architecture (security, connectivity, and management), architects can build a high-performance and scalable campus network.

Appendix A: Campus Product Reference List

Table 2: Campus Reference List

	Infrastructure		Services			Policy & Management	Integrated Routing & Switching
	Routing	Switching	Security/ VPN	Access Control	WAN Optimization	Policy & Management	Integrated Routing & Switching
CAMPUS	M Series	EX8200	ISG2000	SA6500	WX Series	IC6500	SRX5000
	MX Series	EX4200	ISG1000	SA4500		IC4500	SRX3000 line
		EX3200	NetScreen-5400			NSM, NSMXpress	
			NetScreen-5200			OAC	
			IDP8200			SBR Series	
			IDP800			STRM Series	

About Juniper Networks

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

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

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

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin,
Ireland
Phone: 35.31.8903.600
Fax: 35.31.8903.601

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

8030007-001-EN Apr 2009

Printed on recycled paper.