

53-1003786-01
07 April 2015

Designing an Agile, Automated, and Affordable Campus Network

Design Guide

BROCADE 

© 2015, Brocade Communications Systems, Inc. All Rights Reserved.

ADX, Brocade, Brocade Assurance, the B-wing symbol, DCX, Fabric OS, HyperEdge, ICX, MLX, MyBrocade, OpenScript, The Effortless Network, VCS, VDX, Vplane, and Vyatta are registered trademarks, and Fabric Vision and vADX are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned may be trademarks of others.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

The authors and Brocade Communications Systems, Inc. assume no liability or responsibility to any person or entity with respect to the accuracy of this document or any loss, cost, liability, or damages arising from the information contained herein or the computer programs that accompany it.

The product described by this document may contain open source software covered by the GNU General Public License or other open source license agreements. To find out which open source software is included in Brocade products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit <http://www.brocade.com/support/oscd>.

Contents

- Designing an Agile, Automated, and Affordable Campus Network.....4**
- Introduction..... 4
- The Campus Network..... 4
- Building a Tiered Campus Network.....6
- Brocade Campus Network Solutions..... 12
- Summary.....20
- References.....21

Designing an Agile, Automated, and Affordable Campus Network

- Introduction..... 4
- The Campus Network..... 4
- Building a Tiered Campus Network.....6
- Brocade Campus Network Solutions..... 12
- Summary.....20
- References.....21

Introduction

Enterprise networks are constantly changing. As networks respond to variable business demands, sudden growth can occur in some network areas, and constriction can occur in others. The proliferation of new data access technologies, such as smartphones and tablets, demands both increased wireless network access and more stringent security. Corporate reorganization can require relocation of users and workstations, while new applications can require higher performance network connectivity for some departments and lower performance connectivity for others. Accommodating these constantly changing requirements quickly-or at all-is no simple task, particularly when resources and budgets are limited.

At the same time, dynamic change should not undermine corporate network policies for availability, performance, and security. Is your campus network design adequate for the extraordinary challenges ahead? By building a business-optimized network infrastructure today, you ensure you can meet both current and future requirements and satisfy strategic corporate goals

The Campus Network

As the corporate data center network undergoes transformation with technologies such as server virtualization, network consolidation, and adoption of Ethernet fabrics, the focal point of change is currently centered in the user-facing campus side of the network. The physical scope of the campus network is typically one or more buildings that house key corporate departments and staff. Depending on the size of the company, a campus network can be quite extensive, servicing scores of buildings and thousands of users.

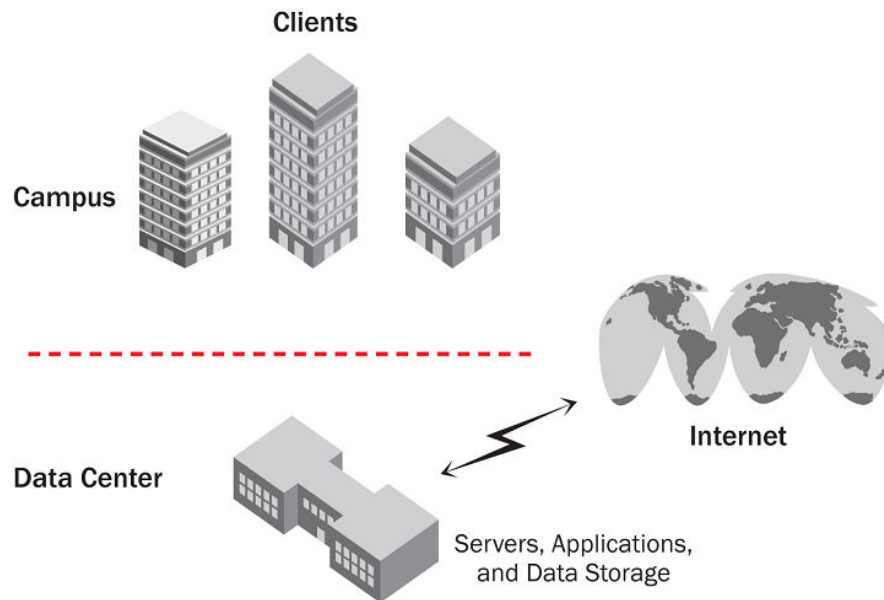
FIGURE 1 The campus network is the user-facing aspect of the corporate network

Fig1_campus_LAN

As illustrated in [Figure 1](#), the data center provides the "heavy lifting" of data processing by housing file and application servers, the data storage network infrastructure, and secure access to the external Internet and corporate intranet. The network design criteria for the data center are therefore significantly different from the design requirements of the campus network. By necessity, the data center network is highly centralized and tends toward uniformity of assets. The campus network, by contrast, is dispersed and must accommodate a wide variety of devices, including workstations, laptops, PDAs, smartphones, Wireless Access Points (WAPs), Voice over IP (VoIP), IP-based remote sensors, Radio Frequency Identification (RFID) readers, and security devices.

Although the data center and campus network are part of a single corporate network, the unique requirements of each domain must be understood, to ensure the harmonious integration of the entire network. The sudden introduction of a new technology (for example, wireless access) on the campus can have unintended consequences in the data center, just as data center restructuring can adversely affect campus access. Optimizing the campus infrastructure must therefore incorporate any downstream effects that require additional changes to the data center design or services.

Performance and Availability

In addition to the diversity of devices and greater mobility required for campus network access, a campus network must also accommodate a wide spectrum of performance and availability requirements for client application access. Many business applications are adequately supported by conventional Fast Ethernet (100 megabits per second [Mbps]) or Gigabit Ethernet (GbE) connectivity, although some very high-performance client applications might require 10 GbE links. With current improvements to wireless LAN technology, the 300+ Mbps enabled by IEEE 802.11n and 802.11ac is suitable for most laptop and mobile applications. Specific applications such as VoIP require additional performance guarantees in the form of Quality of Service (QoS) support and traffic policing.

Redundant network links for high availability are not required for all applications, but certainly mission-critical client applications must have failover capability in the event of an individual link or interface outage. A properly designed campus network infrastructure must be sufficiently flexible to provide the required bandwidth and availability per workgroup or application as business requirements change. The variability of client connectivity requirements also impacts other layers of the network infrastructure as client traffic is funneled to the network core and the data center.

Load balancers are deployed in a campus network to improve availability, security, and performance. Load balancers enhance the user experience by ensuring that users are not sent to failed servers. Load balancers guard against Denial of Service (DoS) attacks and provide offload services such as Secure Sockets Layer (SSL), allowing servers to process more transactions.

Security

The sheer number of client devices at the campus layer poses an ongoing security challenge. A network is only as secure as its weakest link, so a large, dispersed campus networks must be purposely provisioned with distributed security mechanisms to eliminate vulnerabilities to attack or inadvertent access. Access Control Lists (ACLs), authentication, Virtual Private Networks (VPNs), in-flight data encryption, and other safeguards restrict network access to only authorized users and devices and forestall attempts to penetrate the network within the campus itself. Protection of corporate data is both a business imperative and, for many sectors, a legal requirement. Financial and health-related industries, in particular, are now obliged to protect customer and patient information in order to comply with government regulations. Compared to the security mechanisms typically in place in the data center, the campus network is far more vulnerable to malicious attack. Security for the campus network must therefore be constantly reinforced and monitored to avoid exposure.

Management

Given the inherently dispersed nature of the campus and the diversity of client devices, centralized and uniform management is essential for maintaining performance and availability and for enforcing corporate security policies. As new technologies, such as wireless LAN, are introduced to facilitate user access, the campus network management framework must integrate new device and security features to ensure stable operation and provide the necessary safeguards against unauthorized intrusion. Comprehensive network management should also be able to monitor traffic patterns throughout the campus network to proactively identify potential bottlenecks for network tuning.

Reducing Operational Expense

With potentially thousands of workstations, laptops, PDAs, smartphones, and other end devices and hundreds of access points, network switches, and routers, the campus network represents a substantial hardware investment. One component is the initial cost of the equipment itself, but footprint, cooling, and power consumption also contribute to ongoing Operational Expense (OpEx). Due to the dispersed nature of the campus network infrastructure, these costs are less readily identified than comparable operational overhead in the data center. However, such costs should still be factored into the overall campus network design and product selection. Integrating more energy efficient network infrastructure elements and leveraging technologies such as Power over Ethernet (PoE) dramatically reduces ongoing OpEx and minimizes the impact of the network on the corporate budget. In addition, consolidation of network assets by using more efficient high-port-count switches both streamlines management and reduces energy consumption.

Building a Tiered Campus Network

Early networks were essentially large flat networks that enabled peer-to-peer communication at Layer 2 using Media Access Control (MAC) addressing and protocols. A flat network space, however, is vulnerable to broadcast storms that can disrupt all attached devices. This vulnerability increases as the population of devices on a network segment grows. Consequently, Layer 3 routing and the IP address scheme were introduced, to subdivide the network into manageable groups and to provide isolation against Layer 2 calamities. Multiple Layer 2 groups interconnected by Layer 3 routers

facilitate optimal communication within and between workgroups and streamline network management and traffic flows.

Over the years, this basic layered architecture has been further codified into a commonly deployed three-tier network design, as shown in [Figure 2](#). At the periphery, the *access layer* provides initial connectivity for devices to the network. For a large campus network, the access layer might be composed of multiple Layer 2 groupings, as dictated by application or departmental requirements. At the next tier, the *aggregation layer* (sometimes referred to as the distribution layer) concentrates the connectivity of multiple access layer switches to higher-port-count (and typically higher-performance) Layer 3 switches. The aggregation layer switches are in turn connected to the network *core layer* switches, which centralize all connectivity in the network. The trio of access, aggregation, and core layers enables the network to scale over time to accommodate an ever-growing number of end-user devices.

FIGURE 2 A tiered campus network architecture provides access, aggregation, and core layers

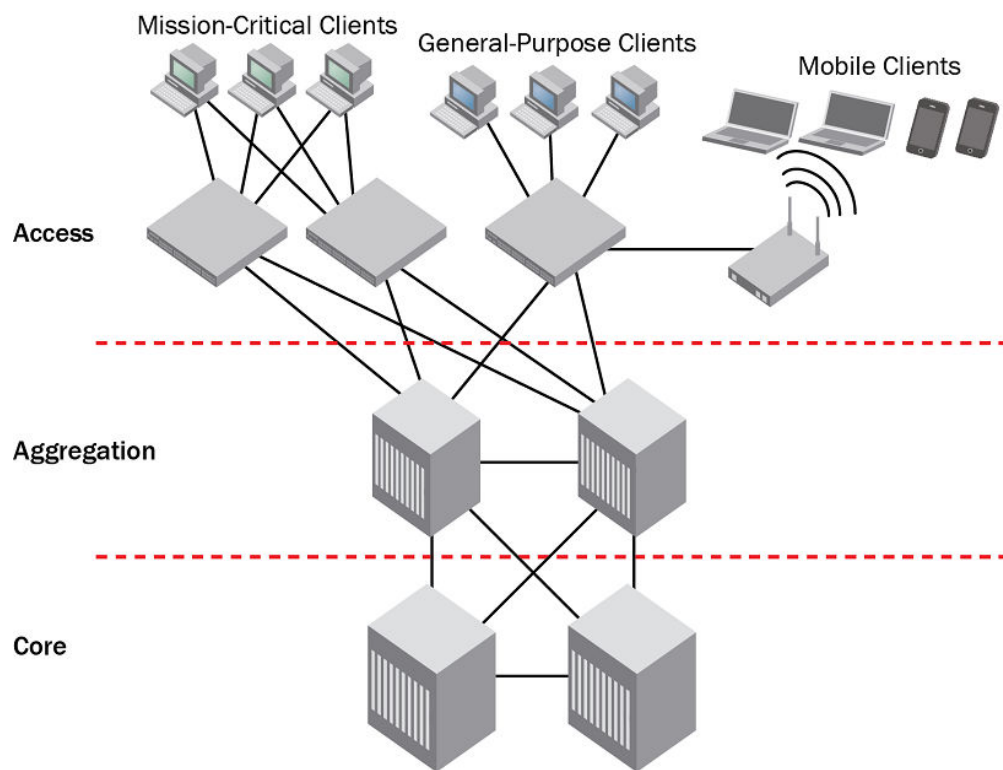


Fig2_tiered_campus_LAN_v2

A tiered campus network design provides the flexibility to support multiple capabilities at the access layer or network edge. Depending on application requirements, high-performance clients can be provisioned with multiple 1 or 10 GbE interfaces for maximum throughput to the aggregation and core layers. General-purpose clients might not require redundant connectivity or high-speed connectivity, so they are adequately serviced by single Fast or GbE links. Likewise, mobile or roaming clients can use a variety of 802.11 speeds via variable-speed wireless access points.

As the client population grows, additional access layer switches and WAPs can be deployed—as well as, when necessary, additional aggregation layer switches to accommodate fan-in to the network core. The network core itself, in turn, can be expanded by the addition of core switches with link aggregation to enhance switch-to-switch bandwidth.

The data center network mirrors the campus network, as shown in [Figure 3](#). In this case, however, the access layer provides connectivity to servers, not clients, and so it typically has much higher bandwidth and availability requirements per device.

FIGURE 3 The data center network uses the access layer for centralized server connectivity




Fig3_access_for_central

In terms of Layer 2 (MAC) and Layer 3 (IP) network protocols, the campus network access layer can be designed for either one of these. Using IP between aggregation and access layer switches enables the use of Open Shortest Path First (OSPF) and other Layer 3 routing protocols instead of the conventional Spanning Tree network protocol (STP) used for Layer 2 bridged networks. This used to be advantageous, given the faster network reconvergence time of OSPF versus STP in the event of a link or switch failure. Rapid Spanning Tree Protocol (RSTP), however, provides subsecond recovery and so is still viable for Layer 2 designs. Routing between the aggregation layer and the core, as well as from the core to any external network or the Internet, is based on Layer 3 IP protocols.

The logical division of the campus network into access, aggregation, and core layers does not necessarily require three physical tiers. By using high-port-count switches, for example, the access and aggregation functions can be collapsed onto a single physical infrastructure, which is a commonly considered option. Smaller campus environments, in particular, benefit from this consolidation, since there are fewer physical assets to manage, and connectivity can be centralized. A consolidation strategy, though, should accommodate growth requirements over time so that the flexibility of a layered architecture is preserved.

An increasingly popular approach campus network architects are considering is the Brocade® HyperEdge® Architecture. HyperEdge Architecture combines innovative new technologies such as Mixed Stacking with Distributed Services and Distributed Access Point Forwarding with existing technologies to collapse the Aggregation and Access tiers and eliminate legacy protocols such as STP. This enables better support of mobility, and reduces application deployment times and operating costs. For more information on this evolutionary and cost-efficient networking approach, download *The Effortless Network: HyperEdge Architecture for the Campus Network* solution brief at: www.brocade.com/forms/getFile?p=documents/position-papers/the-effortless-network-ga-pp.pdf.

Access Layer Requirements

Because client application requirements can vary dramatically between different corporate departments, the access layer represents the most complex tier of the campus network. Unfortunately, network managers are often in a reactive mode to user requests and do not have the luxury of designing and deploying comprehensive architectures that selectively address unique departmental requirements. Consequently, over time, portions of the access layer might be spontaneously expanded with an eclectic mix of network equipment that becomes difficult to manage and unable to accommodate changing user needs. A technology refresh (for example, from Fast to Gigabit Ethernet speeds) is often the only opportunity to streamline access layer design and simplify management. Developing a flexible and scalable wired and wireless access layer design requires first and foremost an understanding of the business applications that need to be supported, or are expected in the foreseeable future, the work habits of users in various departments, and the rate of growth (or contraction) of departmental transactions.

Because high-availability, high-performance, and security requirements can vary from one department to the next, the access layer switch infrastructure should provide multiple speeds, high availability, and VPN and other security protocols, as required. Unified communications such as concurrent VoIP, streaming media, and conventional data transactions can require additional functionality for QoS delivery and Power over Ethernet (PoE). In addition, applications requiring wireless connectivity need both intelligent Wireless LAN (WLAN) access points, as well as centralized management, to ensure stable and secure connectivity.

At the access layer, the goal should be to create an *intelligent edge infrastructure* that automates QoS configuration on a per-user/per-port basis (in other words, different users might need different QoS priorities for mission-critical applications) and that accommodates unified communications with no loss of performance or degraded user experience.

WAPs are deployed on ceilings, walls, and sometimes outdoors throughout the campus environment to provide wireless network access for the organization's mobile users. Wireless traffic connects to the wired access layer switches using Fast-E, Single-GigE, or Dual-GigE uplinks. Uplink capacity is determined based on the type of applications used by mobile users and the expected number of concurrent users during peak loads.

Meanwhile, access layer switches are typically housed in wiring closets distributed on multiple floors of each building on the corporate campus. These, in turn, are connected to aggregation layer switches that feed traffic to other segments or to the network core. To accommodate the fan-in of multiple access layer switches to the aggregation layer, high-performance uplinks are required. Currently, these are typically 10 GbE links or multiple 10 GbE links combined into a single logical link via link aggregation, as shown in [Figure 4](#).

FIGURE 4 Link aggregation provides high-performance uplinks between access and aggregation layers

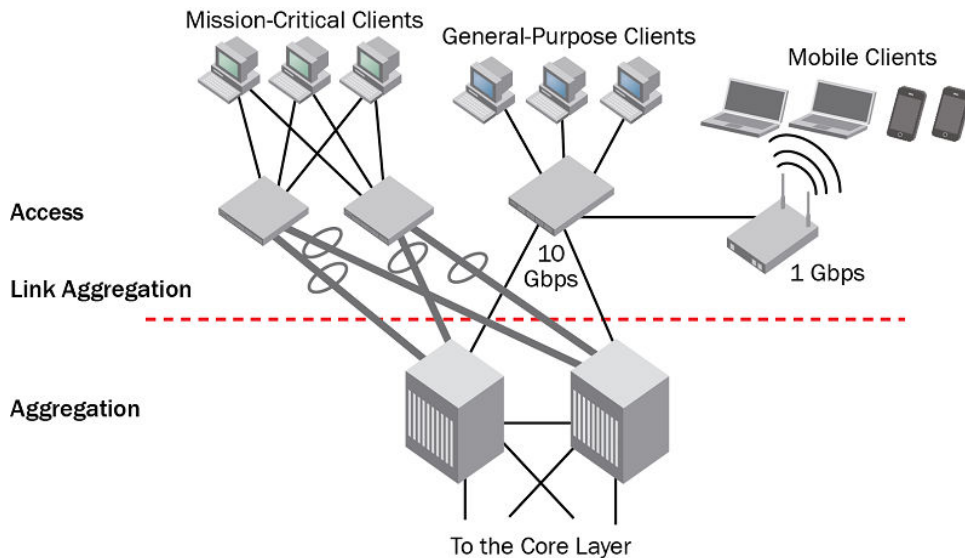


Fig4_link_aggregation_v2

This is another critical design consideration for both access layer and aggregation layer switch selection and sizing. Over-provisioning uplinks can result in wasted bandwidth and additional cost, but under-provisioning can adversely impact client application performance and provoke packet loss due to congestion. The actual uplink bandwidth required does not depend on the fan-in ratio of client devices to uplink ports, but rather on the traffic loads that those clients generate. The uplink connectivity between access and aggregation layers should therefore be sized to support peak traffic volumes, to ensure optimum operation. Best practice is to utilize traffic policing and rate limiting in these links, allowing the connection to limit the amount of a certain type of traffic and ensuring that the link does not become congested.

In [Figure 4](#) and other diagrams, connectivity between access layer switches and the aggregation layer is depicted in a dual-linked configuration. At Layer 2, one of those links is blocked by RSTP and activated only in the event of a primary link failure. This preferred high-availability design, however, is not mandatory for all access layer switches, depending on the business criticality of the applications supported by a particular access layer switch. Likewise, although WLAN access points are often configured to provide overlapping wireless coverage to guard against single access point failures, less mission-critical clients can be adequately served by a less resilient design. Most campus clients, however, require at least continuous access to e-mail and other corporate communication tools, and so they are best supported by a redundant connectivity scheme capable of failover in the event of a switch or link failure.

At the access layer, logical segregation of traffic in different workgroups or departments is achieved by Virtual LAN (VLAN) segmentation. Because IEEE 802.1Q VLAN tagging can span multiple switches, members of a specific VLAN do not have to be physically collocated. By providing a limited Layer 2 broadcast domain for specified groups of users, VLANs help enforce security and performance policies and simplify network management.

Aggregation Layer Requirements

The campus aggregation or distribution layer funnels transactions from multiple access layer switches to the network core. Because each aggregation layer switch is responsible for multiple upstream access layer switch traffic flows from hundreds of users, aggregation layer switches should have high-availability architectures-including redundant power supplies, hot-swappable fans, high-performance backplanes, redundant management modules, and high-density port modules. Aggregation layer

switches are typically Layer 2/3 switches with support for robust routing protocols, to service both the upstream access and downstream core layers. Currently, aggregation layer switches should support full IPv4 and IPv6 protocols, Routing Information Protocols RIPv1/v2 and RIPv6, OSPF v2/v3, Intermediate System-to-Intermediate System (IS-IS) protocol, Border Gateway Protocol (BGP)-4, Distance Vector Multicast Routing Protocol (DVMRP), and Protocol Independent Multicast Sparse and Dense Modes (PIM-SM/DM). To provide adequate bandwidth to the core, aggregation layer switches typically provide multiple 10 GbE ports and link aggregation for higher throughput. Optionally, aggregation layer switches might provide advanced security or other services to support upper layer applications.

Today, the primary driver for robust aggregation layer design is the dramatic increase in access layer clients, the diversity of applications run by those clients, and the increased use of traffic-intensive protocols for multimedia delivery and rich content. Higher volumes of traffic at the access layer require much higher performance at the aggregation layer, as well as mechanisms for maintaining traffic separation and security, as required.

Core Layer Requirements

The core layer represents the heart of the data network infrastructure. Transactions from campus clients to data center servers or to external networks must pass through the core with no loss in data integrity, performance, or availability. Core switch architectures are therefore designed to support 99.999 percent ("five nines") or greater availability and high-density modules of high-performance ports. Because the core layer is also the gateway to the extended corporate network and the Internet, core switches can also be provisioned with 10 GbE ports for access to Carrier Ethernet networks and OC12 (622 Mbps) or OC192 (9.6 gigabit-per-second [Gbps]) high-speed WAN interfaces.

In addition to the robust routing protocols supported by the aggregation layer, the core layer might also provide advanced Multiprotocol Label Switching (MPLS), Virtual Private LAN Service (VPLS), and Multi-Virtual Routing and Forwarding (Multi-VRF) protocols to enforce traffic separation and enforce QoS policies. Network virtualization at the core not only simplifies management; it also ensures that the unique requirements of diverse client populations are met in a common infrastructure. Although it is possible to collapse the aggregation and core layers into a single layer, using large chassis-based switches, the high availability of the core layer is so critical for most business operations that the majority of corporate networks maintain a dedicated core infrastructure.

Building a Fully Meshed Campus Infrastructure

As shown in [Figure 5](#), you should configure the aggregation and core layers in a fully meshed design, with multiple links provisioned between all appropriate switches within and between layers. A fully meshed design provides additional insurance against the loss of an individual link, port module, or entire switch, while minimizing network reconvergence time and ensuring non-disruptive traffic flow.

FIGURE 5 A fully meshed configuration minimizes recovery time and packet loss

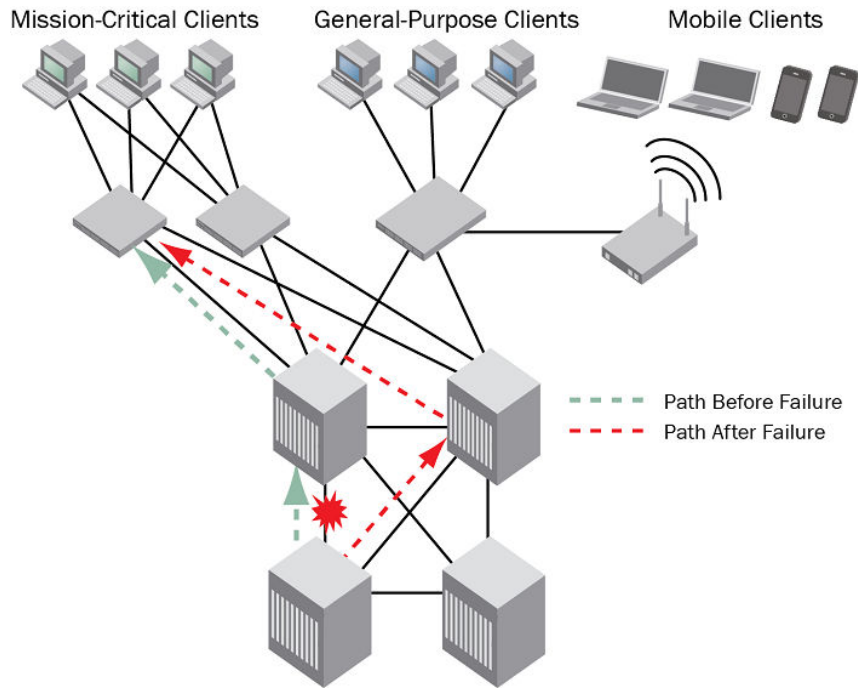


Fig5_fully_meshed_v2

With IP-based routing protocols such as OSPF, link availability and cost are automatically calculated by the switches themselves. With more available links between source and destination switches, traffic can be more quickly rerouted in the event of an individual link failure. While a meshed configuration dedicates more ports to inter-switch connectivity and reduces the number of ports available for device support, the cost is minimal compared to the potential application disruption if a network reconfiguration is triggered.

NOTE

For more detailed information about campus networking, see the References section at the end of the document.

Brocade Campus Network Solutions

Brocade offers a full suite of Layer 2 and Layer 3 network solutions engineered for enterprise-class applications, combining high performance with resiliency and industry-leading energy efficiency. Whether deploying a campus network using a traditional multi-tiered architecture, Brocade HyperEdge Architecture, or both, Brocade offers a broad spectrum of access, aggregation, and high-performance core switches to build a complete end-to-end solution for both large and medium enterprises.

Brocade HyperEdge Architecture Overview

To achieve the required flexibility, automation, and reductions in cost of ownership requires a new vision. Brocade calls this vision The Effortless Network, and it is built from the innovations found in the HyperEdge Architecture.

The Brocade HyperEdge Architecture collapses network layers to radically simplify networks and eliminate legacy protocols such as spanning tree. HyperEdge Architecture integrates innovative new features with existing network technologies to streamline application deployment, simplify management, and reduce operational costs.

HyperEdge Architecture Key Design Principles

Three key design principles drive the HyperEdge Architecture for modernizing and simplifying the network to achieve better business agility and productivity.

1. **Consolidated management:** Reduces the number of network layers to create large HyperEdge management domains that eliminate individual switch touchpoints to ease maintenance time and costs.
2. **Shared network services:** Allows premium and entry-level switches that share a common HyperEdge management domain to share advanced L2/L3 services to achieve lower price-per-port functionality.
3. **Scale-out networking :** Integrates high-performance fixed form factor switches to create a single logical device independent of physical location by scaling ports when and where needed across the campus.

HyperEdge Architecture Implementation Options

Brocade has developed unique enabling technologies, called Mixed Stack and Switch Port Extender, to achieve the benefits of the HyperEdge Architecture design principles (see [Figure 6](#)). The Brocade ICX® series of high-performance fixed switches embodies these enabling technologies with flexible distributed chassis configuration deployment options. These powerful deployments deliver equivalent or better functionality than large rigid modular chassis systems, but with significantly lower costs and carbon footprints.

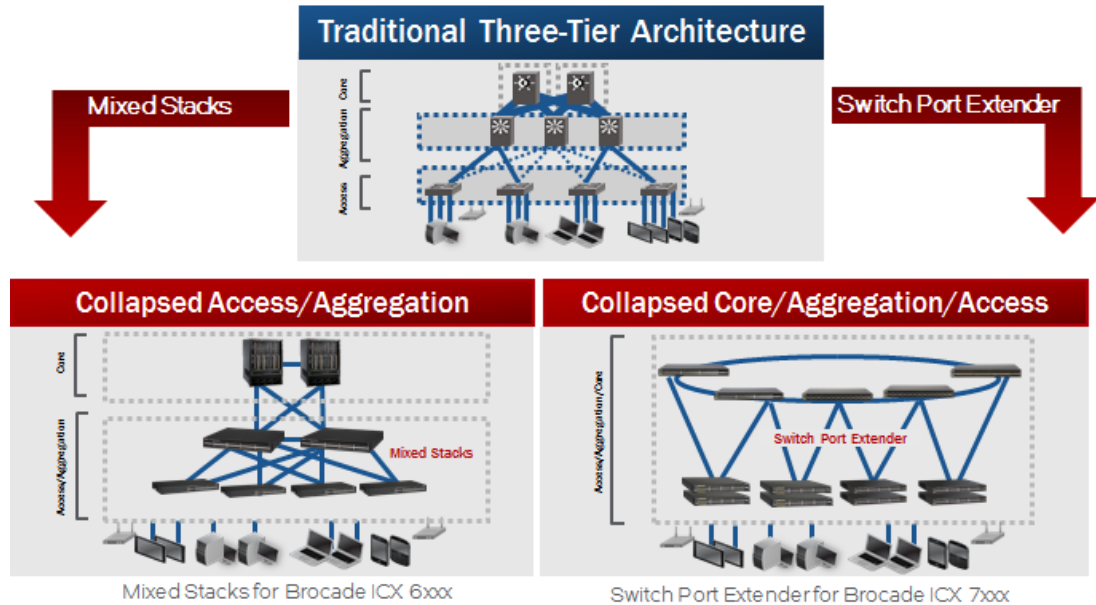
- **The Mixed Stack** enabling technology, offered for the Brocade ICX 6xxx Switches, integrates premium Brocade ICX 6610 and entry-level Brocade ICX 6450 Switches, which collapse the network access and aggregation layers into a single HyperEdge domain that shares services while reducing management touch points and network hops across fewer network layers as compared to legacy three-tier designs.
- **The Switch Port Extender** enabling technology, offered for the Brocade ICX 7xxx Switches, integrates premium Brocade ICX 7750, midrange Brocade ICX 7450, and entry-level Brocade ICX 7250 Switches, which collapse the network access, aggregation, and core layers into a single HyperEdge domain. Similar to Mixed Stacks, Switch Port Extender shares services across multiple levels of switches while reducing management touch points and network hops in a single-layer design across the entire campus network.

NOTE

Switch Port Extender support will be available in a future release.

These implementation options can also be integrated into existing traditional networks for incremental adoption. As Brocade continues to further develop the HyperEdge architecture, new implementation options will be available, delivering the benefits of the key HyperEdge design principles of single-point management, shared services, and scale-out networking across an ever increasing number of network ports.

FIGURE 6 The Brocade HyperEdge Architecture offers multiple implementation options



Benefits of HyperEdge Architecture

The HyperEdge Architecture enables organizations to build networks that are:

- **Agile:** By eliminating STP between HyperEdge Domain switches through a flatter Layer 2 design, the HyperEdge Architecture increases link utilization and reduces application deployment complexity.
- **Automated:** By grouping premium and entry-level switches into a consolidated management domain, HyperEdge Domains eliminate the need to provision and manage devices individually—offering as single point of management and simplifying network deployment.
- **Affordable:** The HyperEdge Architecture enables the propagation of advanced features and services from premium switches to entry-level switches, allowing IT organizations to purchase only what they need today and add intelligent services as the business evolves.

For information on Brocade's HyperEdge Architecture, download *The Effortless Network: HyperEdge Architecture for the Campus Network* solution brief at: <http://www.brocade.com/forms/getFile?p=documents/position-papers/the-effortless-network-ga-pp.pdf>.

Brocade Campus Network Access Layer Solutions

Brocade campus stackable access switches scale from 24 and 48 ports—with switching capacity from 12 Gbps for the Brocade ICX[®] 6430-C Switch, to 576 Gbps for the Brocade ICX 6610 Switch. Brocade campus access switches can be deployed to support modest port requirements and then to scale over time to accommodate higher populations of users and devices as business requirements evolve, demand grows, and new technologies emerge.

FIGURE 7 Brocade offers a range of options for campus access connectivity

As shown in [Figure 7](#), the Brocade campus access product family offers a range of products to fit any business and any budget—from the entry level Brocade ICX 6430 to the high-performance Brocade ICX 6610 Series.

The Brocade ICX 6430/6450 Series switches are enterprise-class stackable switches at an entry-level price. The Brocade ICX 6430 targets the entry level with 12, 24, or 48 × 10/100/1000 Mbps RJ-45 ports and 4 × 1 GbE SFP uplink ports. The Brocade ICX 6450 Switch represents the next level, adding 4 × 10 GbE SFP uplinks. Both models support hitless stacking. The Brocade ICX 6430 supports stacks up to 4 units, and the Brocade ICX 6450 supports up to 8 units with 40 Gbps of stacking bandwidth, offering up to 384 ports per stack for the Brocade ICX 6450. Both switches offer PoE+ and non-PoE+ models and provide an advanced external power supply option for redundancy and to extend PoE+ power. This enables organizations to tailor their network to their exact needs and budget.

The Brocade ICX 6430-C and 6450-C Compact Switches offer enterprise-class network switching capabilities, performance, reliability, security, and manageability in a small form factor with fanless operation for deployment outside the wiring closet. The Brocade ICX 6430-C is ideal for deployment in classrooms, retail locations, factories, small offices, workgroups, and space-constraint environments. The Brocade ICX 6430-C is available in a 12-port 10/100/1000 Mbps model with IEEE 802.3af PoE and 802.3at PoE+ support on 4 ports plus 4 additional Gigabit Ethernet uplink ports.

The Brocade ICX 7250 Series Switches delivers the performance, flexibility, and scalability required for enterprise Gigabit Ethernet (GbE) access deployment. It raises the bar with up to 8×10 GbE ports for uplinks or stacking and market-leading stacking density with up to 12 switches (576×1 GbE) per stack. In addition, the Brocade ICX 7250 combines enterprise-class features, manageability, performance, and reliability with the flexibility, cost-effectiveness, and "pay as you grow" scalability of a stackable solution.

The Brocade ICX 7250 Switch provides enterprise-class stackable LAN switching solutions to meet the growing demands of campus networks. Designed for small to medium-size enterprises, branch offices, and distributed campuses, these intelligent, scalable edge switches deliver enterprise-class functionality at an affordable price—without compromising performance and reliability. The Brocade ICX 7250 is available in 24- and 48-port 10/100/1000 Mbps models with 1 GbE or 10 GbE dual-purpose uplink/stacking ports—with or without IEEE 802.3af PoE and 802.3at PoE+—to support enterprise edge networking, wireless mobility, and IP communications without the need for additional power outlets or power injectors.

The Brocade ICX 7450 Series switches are the first in their class to offer 40 GbE uplinks, enabling enterprises to dramatically increase their network capacity while using their existing optical wire infrastructure.

The unique design of the Brocade ICX 7450 provides three modular slots, offering up to 12 x 1/10 GbE SFP/SFP+ ports, 12 x 10GBASE-T ports, or up to three 40 GbE QSFP+ ports for uplink or stacking. As a result, the Brocade ICX 7450 can easily deliver sufficient bandwidth between the edge and aggregation layers to support expanding video traffic, VDI adoption, and highspeed wireless 802.11ac deployment. The Brocade ICX 7450 is an ideal network solution for campus network 1 GbE access or small aggregation deployment with 10 GbE or 40 GbE uplinks to the core. The Brocade ICX 7450 also makes a very suitable data center Top-of-Rack (ToR) solution, delivering a mix of 1 GbE and 10 GbE server connectivity ports with 10 GbE or 40 GbE uplinks to the data center aggregation or core.

The Brocade ICX 6610 Series stackable switches deliver chassis-like capabilities in a stackable form factor, offering class-leading performance, scalability, and availability. The Brocade ICX 6610 provides

320 Gbps of stacking bandwidth through 4 × 40 Gbps redundant stacking links. Up to 8 × 10 GbE SFP+ ports can be configured as uplinks or used as high-performance access ports. With dual 1000 W power supplies, the switch can deliver up to 30 W of PoE+ power to all ports. The Brocade ICX 6610 also offers advanced IPv4 and IPv6 L3 capabilities.

By combining high port density with a compact rackable form factor, the Brocade campus access switches occupy minimal space in wiring closets and simplify deployment of both conventional and PoE/PoE+ campus devices-such as VoIP phones, web conferencing devices, surveillance cameras, and IEEE 802.1n WAPs.

The support offered by Brocade for both IEEE 802.3af PoE and 802.3at PoE+ standards at the access layer provides both data and power connectivity for multiple classes of PoE-capable phone systems, WAPs, environmental sensors, security cameras, and other equipment that uses standard CAT-5 cabling. The integration of both data and PoE connectivity adds another step in the network design process to ensure that sufficient power is allocated to the appropriate end devices.

Brocade access layer switches support advanced Layer 2 and Layer 3 protocols, including edge security, IEEE 802.1x port-based authentication, QoS delivery, and VLAN-based ACL policy enforcement. These intelligent edge services provide secure connectivity for client workstations and edge devices and enable network designers to deploy QoS as needed for high-priority or latency-sensitive applications.

Brocade Distributed Chassis Architecture for Scale-Out Networking at the Aggregation and Core

Traditional three-tier network design with "big-box" chassis at the aggregation and core layers require a significant up-front investment and offer limited deployment flexibility and future-proofing.

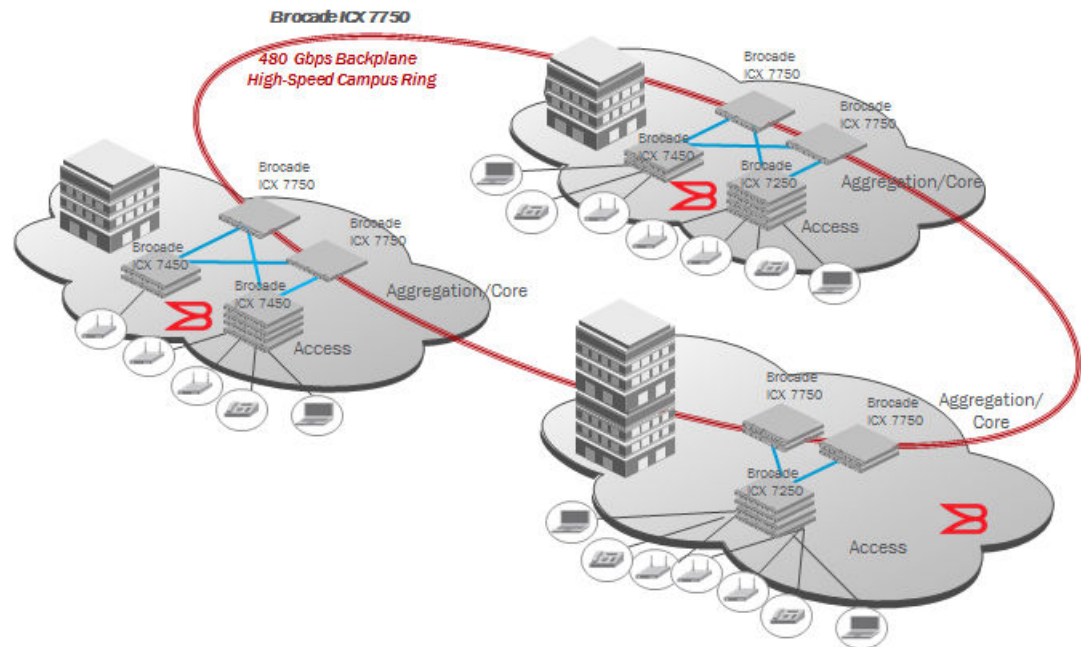
In contrast, a distributed "multi-box" architecture at the aggregation and core can deliver much greater scalability and future-proofing with an easier "upgrade as you go" model. This type of architecture enables network architects to add capacity exactly where it is needed in the network, unlike a "big-box" chassis approach with all ports located in the same closet.

Thanks to rapid technology evolution and innovative thinking, Brocade is able to offer the first stackable solution for the campus aggregation and small core that delivers higher performance and port density than a traditional midsize chassis, while offering the same level of reliability and availability.

The Brocade ICX 7750 stackable aggregation switch offers a level of flexibility, ease of deployment, and total cost of ownership unmatched by traditional aggregation and small-core chassis solutions.

Brocade offers a truly distributed chassis architecture, where all the components can be distributed and spread over the entire campus-due to the use of long-distance optical links-but the system as a whole can be managed as a single entity.

For maximum flexibility, multiple topologies are supported. [Figure 8](#) shows a campus ring where all aggregation switches are connected across the campus in a ring configuration. With this topology, it is actually possible to eliminate the need for a separate core layer on a medium-size campus. This greatly simplifies the management and troubleshooting of the entire campus network, with a single point of management used for the aggregation layer.

FIGURE 8 Brocade ICX 7750 deployed in a "campus ring" topology combining aggregation and core

Brocade offers a unique, distributed scale-out architecture that enables a true chassis replacement at the aggregation and core of the campus, with an increase in availability, scalability, and performance over a traditional chassis-based solution.

Brocade Campus Aggregation Layer Solutions

Brocade campus network aggregation layer switches include the Brocade ICX 6610 and ICX 7750 Series switches and the Brocade FastIron SX. Both series incorporate modular designs that enable customers to align aggregation deployments to the appropriate port density and performance requirements of different campus sizes.

The Brocade ICX 6610 Series: The unprecedented stacking performance and availability of the Brocade ICX 6610 make it an ideal solution for the mid-sized campus aggregation layer. The four Brocade ICX 6610 fully redundant 40 Gbps stacking links, combined with Brocade hitless stacking technology, deliver the level of availability required at the aggregation layer. The Brocade ICX 6610 family offers a fiber port option switch with 24×1 GbE SFP ports that can be configured to connect to access layer switches. It also offers 8×10 GbE SFP+ ports for connections to the core, delivering a "pay as you grow" stackable solution for campus aggregation.

The Brocade ICX 7750 Series: The Brocade ICX 7750 Switch is a 1U fixed form factor 10/40 GbE Ethernet switch that delivers a chassis experience for campus LAN aggregation and core. It offers unprecedented port density and chassis-level performance, availability, and scalability. The Brocade ICX 7750 distributed chassis stacking technology enables scale-out networking. The Brocade ICX 7750 redefines the economics of enterprise networking by delivering a unique 10/40 GbE campus aggregation solution in a stackable form factor, enabling new levels of performance, availability, and flexibility. It provides the capabilities of a chassis with the flexibility and cost-effectiveness of a stackable switch.

The Brocade FastIron SX Series delivers industry-leading price and performance value for campus aggregation and core switching and is available in 8- and 16-slot chassis. The high-performance architecture offers up to 132×10 GbE and 384×1 GbE ports in a single chassis, supporting IPv4/IPv6-capable Layer 2/3 switching and routing. The chassis is designed to deliver maximum availability with

Multi-Chassis Trunking (MCT), which enables two chassis to appear as a single logical switch at Layer 2 in active/active mode and delivers uninterrupted traffic flow in the event of node failover. Additional High Availability (HA) capabilities include redundant and hot-pluggable management modules (with hitless failover), redundant and hot-pluggable switch fabrics, power supplies, fans, stateful OSPF redundancy, graceful BGP and OSPF restarts, and hitless (in-service) software upgrades. The Brocade FastIron SX Series also supports up to 384 PoE-enabled ports for aggregating WAPs or other PoE and PoE+ devices.

FIGURE 9 Brocade offers a range of stackable and chassis products for campus aggregation, which scale from midsize to large deployments



Because the aggregation layer typically requires more robust routing support, Brocade aggregation solutions provide full Layer 3 IPv4 and IPv6 routing, including RIPv1/v2, RIPng, OSPFv2/v3, IS-IS, BGP-4, DVMRP, and PIM-SM/DM, as described in a previous section. With wire-speed performance and a rich suite of routing protocols, Brocade FastIron SX Series switches bring thousands of devices into an integrated and more easily managed campus network architecture.


Brocade Campus Core Layer Solutions

While each aggregation layer switch provides reliable connectivity for discrete groups of access layer switches and access points, the core layer is responsible for the seamless operation of the entire campus network infrastructure, including connectivity to the data center and to external networks. So, in addition to the high availability and high performance functionality that is characteristic of aggregation layer switches, core layer platforms must provide further enhancements to performance. They must also provide WAN connectivity options and higher-level protocols for advanced services such as MPLS, VPLS, and virtual routing for network virtualization. These advanced features simplify management and enable the campus network to more easily meet diverse upper-layer application requirements and business goals.

The Brocade FastIron SX Series high-performance architecture delivers up to 128 × 10 GbE in a highly available chassis that supports IPv4/IPv6-capable Layer 2/3 switching and routing, providing leading price/performance value as a solution for the midsize campus core layer.

The Brocade MLX[®] Series switches are designed for the more rigorous requirements of the campus network core and are available in modular chassis configurations to meet specific performance and connectivity needs. As illustrated in the figure below, the Brocade MLX core switches provide from 4 to 32 slots for any combination of Gigabit Ethernet and 10 GbE and WAN port distribution. With an up to 7.68 terabits per second (Tbps) switching capacity and support for up to 128 × 10 GbE ports at wire speed, the Brocade MLX Series is designed to centralize campus network traffic management for continuous non-disruptive operation.

FIGURE 10 The Brocade MLX core layer switch series provides both aggregation layer and WAN connectivity



Conventional WAN connectivity for external networks is provided through 4- and 8-port OC48 (2.5 Gbps) and 2-port OC192 (9.6 Gbps) modules. In addition, you can use 10 GbE ports to connect to Carrier Grade Ethernet networks for metropolitan environments. The Brocade MLX architecture is future-proofed to support emerging 40 and 100 GbE standards, which provide greater flexibility for both mesh interswitch links and high-performance campus core backbones.

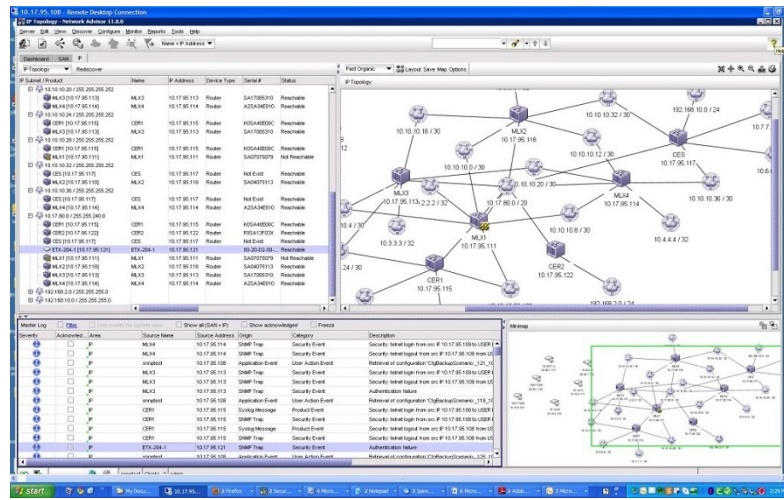
Advanced routing services at the core layer include MPLS and Multi-VRF advanced QoS for multiservice networks and sFlow management for more granular network traffic accounting. These Brocade MLX capabilities enable the construction of highly reliable core infrastructures that facilitate transport of multiple traffic protocols over an integrated network.

Building and Managing a Brocade Campus Network

The Brocade standards-based IP technology provides a wide variety of solutions for access, aggregation, and core layer connectivity, all centrally managed via a secure and robust management framework-Brocade Network Advisor.

Brocade ICX access layer switches, Brocade ICX and FastIron SX aggregation layer switches, and high-performance Brocade MLX core switches enable customers to selectively deploy the most efficient solutions that meet campus network requirements. you can configure these technologies for a wide range of client access requirements at the network edge, and you can scale through the aggregation and core layers to meet bandwidth and availability needs. Link aggregation between aggregation layer switches and the network core can be dynamically sized to accommodate traffic changes over time. Security policy enforcement is facilitated by ACL authentication, data encryption, IPsec, and VLAN enforcement throughout the network.

FIGURE 11 Brocade Network Advisor provides comprehensive campus network management



Brocade Network Advisor is the industry's first unified management solution across storage and IP networks that simplifies management of enterprise campus networks and metro and carrier Ethernet networks. Brocade Network Advisor supports Fibre Channel (FC), Fibre Channel over Ethernet (FCoE), IP switching and routing, wireless, and MPLS networks—with end-to-end network visibility across these different network types in a single application. Brocade Network Advisor supports comprehensive lifecycle management capabilities across these different networks via a seamless and unified user experience.

Summary

Campus networks have unique design criteria that demand greater flexibility in supporting a wide range of client devices and applications. Wired access, wireless access, and PoE devices must coexist at the network edge. Sophisticated client applications require high bandwidth, high availability, and security, while mid-tier applications might be adequately served with more modest connectivity. The traditional "tiered" campus network architecture with fan-out to client devices at the access layer, consolidation of access layer traffic through the aggregation layer, and centralized routing through the network core provides an adequate model for growing the campus network over time and accommodating higher traffic volumes and multiple protocols as required. For a more scalable, easier to deploy and manage and cost effective model, the Brocade HyperEdge Architecture collapses network layers to radically simplify networks and eliminate legacy protocols such as spanning tree. Brocade HyperEdge Architecture integrates innovative new features with existing network technologies to streamline application deployment, simplify management, and reduce operational costs.

The full suite of Brocade intelligent campus network IP infrastructure solutions and comprehensive network management tools enables customers to build and expand robust, cost-effective, and business-optimized campus networks using a traditional multitiered architecture, the Brocade HyperEdge Architecture or both to meet current and future corporate requirements.

References

For additional information about campus networking, including detailed information on the following:

- Campus Network Infrastructure, Base Reference Architecture
- Campus Network Infrastructure, Best Practices for High Availability Design
- Best Practices: Brocade HyperEdge Architecture with Mixed Stacking
- HyperEdge Architecture Deployment Guide for K12 Schools
- BYOD Solution Design Guide with Bradford Networks

For more information, go to <http://community.brocade.com/t5/Design-Build/tkb-p/DesignBuild>