



CHAPTER 1

Borderless Campus Design and Deployment Models

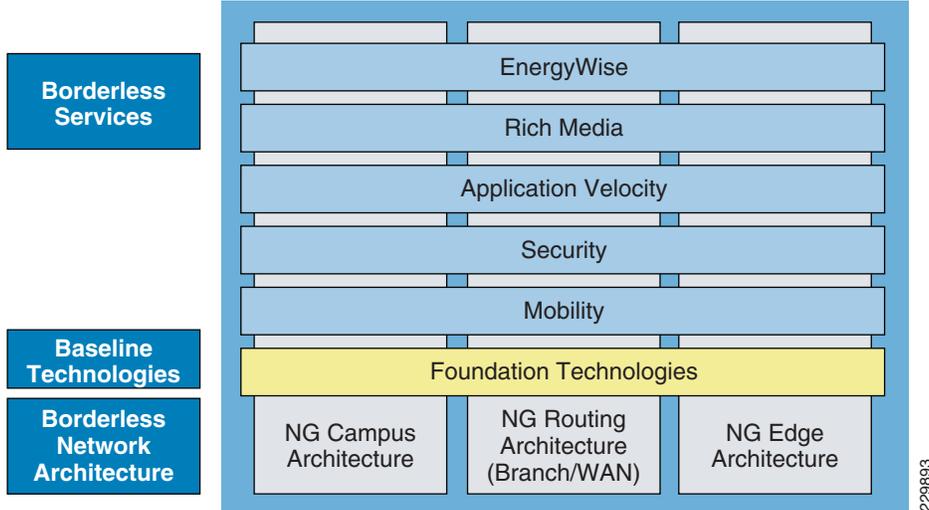
Executive Summary

Enterprises are making a fundamental shift, with employees no longer confined to physical offices, geographical locations, and time zones. In today's globalized workplace, work can occur anywhere in the world and enterprise information needs to be virtually accessible, on-demand, through every part of the network. These requirements drive the need to build next-generation networks that are secure, reliable, and highly available. Adding to the impetus for next generation networks is the transformation of the network from traditional data and voice transport to super-highways for video, building systems, physical security systems, and other non-traditional systems that now use IP networks to communicate. As more and more systems converge on the network, network complexity increases and the capability to handle traffic in a responsible manner is essential as the network becomes even more mission critical.

This network transformation presents many challenges to enterprise IT staff. They must be able to transform their networks to allow secured network access from anywhere, but enforce security policies based on how users are accessing the network. They must allow multiple systems or services to simultaneously traverse the network, while allocating sufficient network resources to ensure those systems do not negatively impact each other. The network must be agile to adapt to future requirements without requiring an overhaul of the existing network. And of course the scalable, resilient, highly available, service differentiating, adaptable network they create must be cost effective and protect investments in the network.

The Cisco Borderless Network architecture is designed to directly address these IT and business challenges. It offers a seamless user experience with a next-generation network that allows different elements of the network, from access switches to wireless access points, to work together and allow users to access resources from anyplace at anytime. The Cisco Borderless Network uses a tiered approach to virtually collapse the network as a single borderless network. It also integrates key services into the network fabric while increasing reliability and security and decreasing service time. For such an infrastructure, the enterprise network must be developed with an architectural approach that embeds intelligence, simplifies operations, and is scalable to meet future demands. The Cisco Borderless Network is a next-generation network architecture that combines several innovations and architectural design considerations to offer a new workspace experience. The Cisco Borderless Network is composed of several modular components, as illustrated in [Figure 1-1](#).

Figure 1-1 Cisco Borderless Network Framework



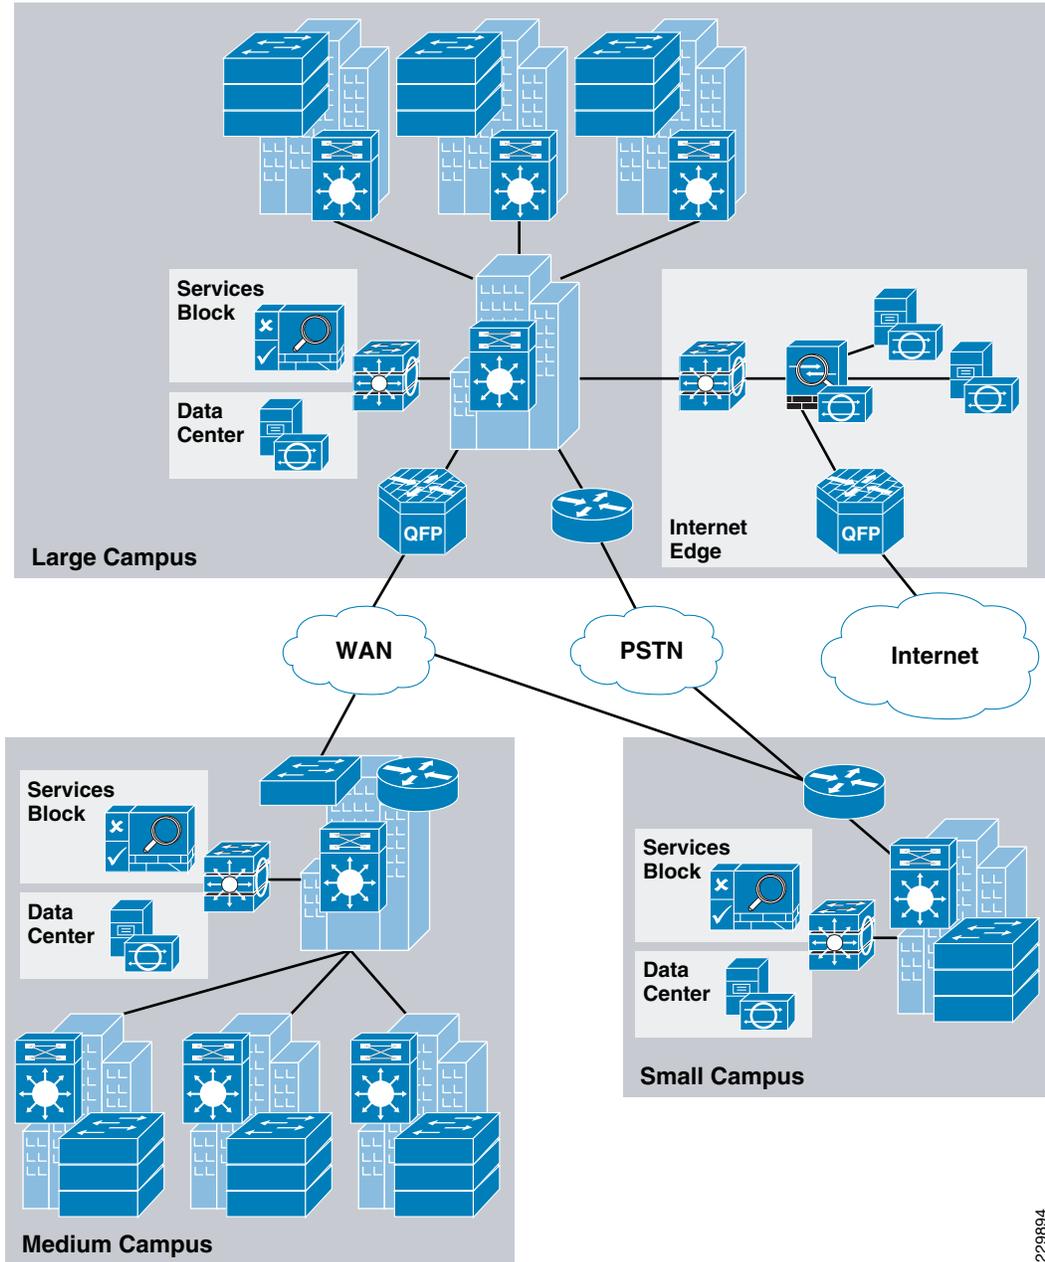
Each building block in the Cisco Borderless Network framework is designed to offer the following components:

- **Network Infrastructure**—Builds enterprise campus, WAN, and edge networks as an open platform that can provide secure and intelligent services at the network edge, aggregation scalability, and a high-performance backbone solution to enable end-to-end borderless services and applications.
- **Foundation Technologies**—Common baseline technologies that are integrated across various enterprise architectures to optimize service delivery, intelligently differentiate between various applications, and build the highly-available network infrastructure.
- **Borderless Services**—Enables the end-to-end borderless user experience to provide ubiquitous connectivity with security, reliability, and sustainability to the enterprise workspace users and the network edge elements. Empowers network architects to leverage the network as a platform to offer rich services to reduce business operational costs, increase efficiency through green practices, and much more.

Borderless Campus Network Design

The Borderless Campus Network architecture is a multi-campus design, where a campus consists of multiple physical buildings with a wide range of network services that offer the capability for anyone to securely access network resources from anywhere at anytime, as shown in [Figure 1-2](#).

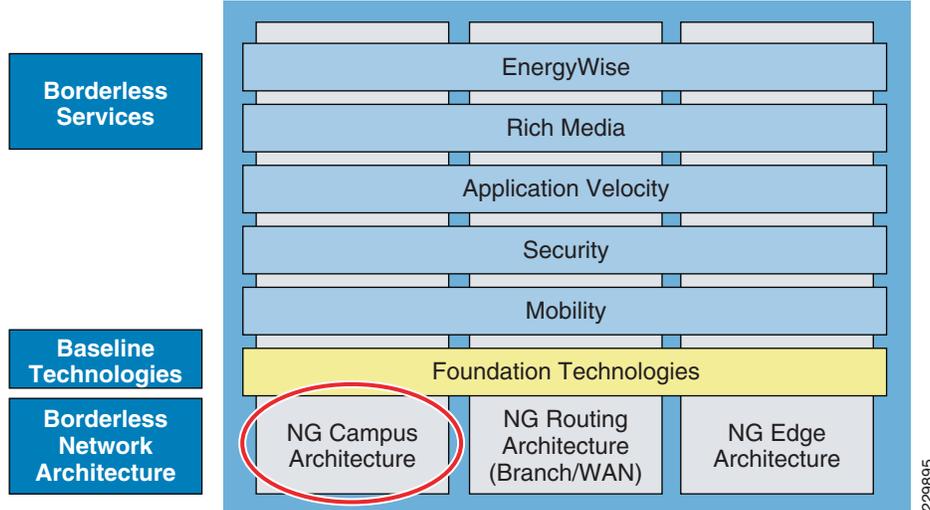
Figure 1-2 Borderless Campus Network Design



229894

Figure 1-3 shows the service fabric design model used in the Borderless Campus.

Figure 1-3 Borderless Campus Architecture



This document describes the campus framework and network foundation technologies that provide a baseline of routing, switching, and several key network services guidelines. The campus design interconnects several other infrastructure components, such as endpoints at the network edge, data center, WAN, and so on, to provide a foundation on which mobility, security, video, and unified communications (UC) can be integrated into the overall design.

This campus design provides guidance on building the next-generation enterprise network, which becomes a common framework along with critical network technologies to deliver the foundation for the service fabric design. This chapter is divided into the following sections:

- *Campus design principles*—Provides proven network design choices to build various types of campus infrastructure.
- *Campus design model for the enterprise*—Leverages the design principles of the tiered network design to facilitate a geographically-dispersed enterprise campus network made up of various elements, including networking role, size, capacity, and infrastructure demands.
- *Considerations of a multi-tier campus design model for enterprises*—Provides guidance for the enterprise campus LAN network as a platform with a wide range of next-generation products and technologies to seamlessly integrate applications and solutions.
- *Designing network foundation services for campus designs in the enterprise*—Provides guidance on deploying various types of Cisco IOS technologies to build a simplified and highly-available network design to provide continuous network operation. This section also provides guidance on designing network-differentiated services that can be used to customize the allocation of network resources to improve user experience and application performance and to protect the network against unmanaged devices and applications.

Borderless Campus Network Design Principles

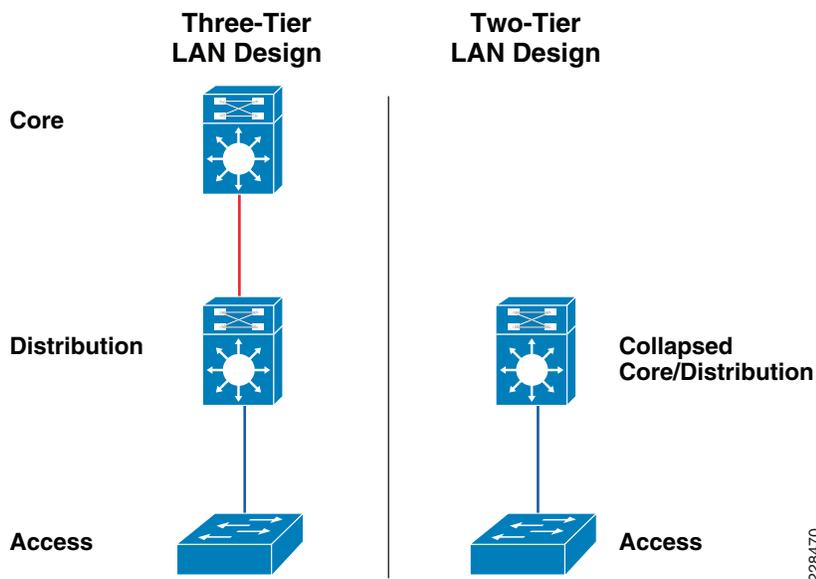
Designing the borderless campus requires that sound network design principles are used to ensure maximum availability, flexibility, security, and manageability. The use of sound network design ensures that the network will deliver on current requirements as well as be well prepared for future services and technologies. This document provides design guidelines that are built upon the following principles to allow the enterprise network architect to build a geographically-dispersed borderless network:

- *Hierarchical*
 - Facilitates understanding the role of each device at every tier
 - Simplifies deployment, operation, and management
 - Reduces fault domains at every tier
- *Modularity*—Allows seamless network expansion and integrated service enablement on an on-demand basis
- *Resiliency*—Satisfies user expectations for keeping the network always on
- *Flexibility*—Allows intelligent traffic load sharing by using all network resources

These are not independent principles. The successful design and implementation of a campus network requires an understanding of how each of these principles applies to the overall design. In addition, understanding how each principle fits in the context of the others is critical in delivering the hierarchical, modular, resilient, and flexible networks required by enterprises.

Designing the Borderless Campus network in a hierarchical fashion creates a flexible and resilient network foundation that allows network architects to overlay the security, mobility, and unified communication features essential to the service fabric design model. The two proven, time-tested hierarchical design frameworks for campus networks are the three-tier layer and the two-tier layer models, as shown in [Figure 1-4](#).

Figure 1-4 Three-Tier and Two-Tier Campus Design Models



The key layers are access, distribution, and core. Each layer can be seen as a well-defined structured module with specific roles and functions in the campus network. Introducing modularity into the campus hierarchical design further ensures that the campus network remains resilient and flexible to provide critical network services as well as to allow for growth and changes that may occur over time.

- *Access layer*

The access layer represents the network edge, where traffic enters or exits the campus network. Traditionally, the primary function of an access layer switch is to provide network access to the user. Access layer switches connect to distribution layer switches, which perform network foundation technologies such as routing, quality of service (QoS), and security.

To meet network application and end-user demand, the next-generation Cisco Catalyst switching platforms no longer simply switch packets, but now provide more converged, integrated, and intelligent services to various types of endpoints at the network edge. Building intelligence into access layer switches allows applications to operate on the network more efficiently, optimally, and securely.

- *Distribution layer*

The distribution layer interfaces between the access layer and the core layer to provide many key functions, including:

- Aggregating large-scale wiring closet networks
- Aggregating Layer 2 broadcast domains and Layer 3 routing boundaries
- Providing intelligent switching, routing, and network access policy functions to access the rest of the network
- Providing high availability through redundant distribution layer switches to the end-user and equal cost paths to the core, as well as providing differentiated services to various classes of service applications at the edge of network

- *Core layer*

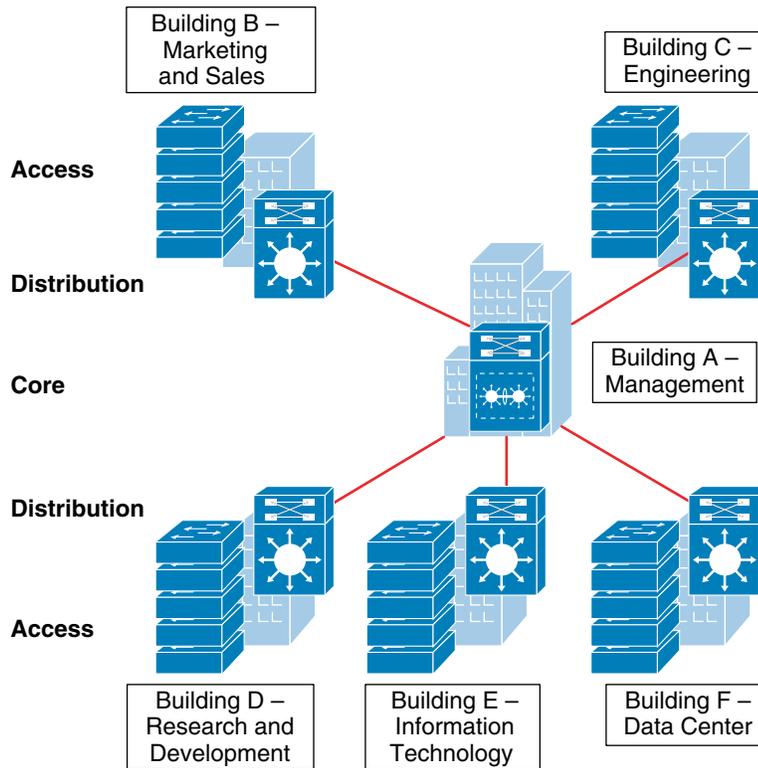
The core layer is the network backbone that hierarchically connects several layers of the campus design, providing for connectivity between end devices, computing and data storage services located within the data center and other areas, and services within the network. The core layer serves as the aggregator for all of the other campus blocks and ties the campus together with the rest of the network.



Note

For more information on each of these layers, see the enterprise class network framework at: <http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/campover.html>.

Figure 1-5 shows a sample three-tier campus network design for enterprises where the access, distribution, and core are all separate layers. To build a simplified, scalable, cost-effective, and efficient physical cable layout design, Cisco recommends building an extended-star physical network topology from a centralized building location to all other buildings on the same campus.

Figure 1-5 Three-Tier Campus Network Design Example

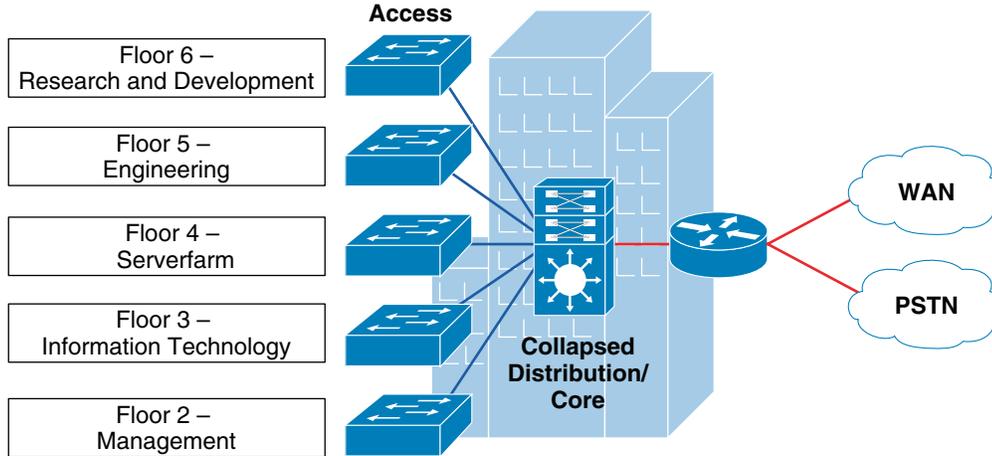
229354

The primary purpose of the core layer is to provide fault isolation and high-speed backbone connectivity with several key foundational services. Isolating the distribution and core into separate layers creates a clean delineation for change control activities affecting end devices (laptops, phones, and printers) and those that affect the data center, WAN, or other parts of the campus network. A core layer also provides for flexibility in adapting the campus design to meet physical cabling and geographical challenges. If necessary, a separate core layer can use a different transport technology, routing protocols, or switching hardware than the rest of the campus, providing for more flexible design options when needed.

In some cases, because of either physical or network scalability, having separate distribution and core layers is not required. In smaller campus locations where there are fewer users accessing the network or in campus sites consisting of a single building, separate core and distribution layers may not be needed. In this scenario, Cisco recommends the alternate two-tier campus network design, also known as the collapsed core network design.

Figure 1-6 shows a two-tier campus network design example for an enterprise campus where the distribution and core layers are collapsed into a single layer.

Figure 1-6 Two-Tier Network Design Example



If the small-scale collapsed campus core design is used, the enterprise network architect must understand network and application demands so that this design ensures a hierarchical, modular, resilient, and flexible campus network.

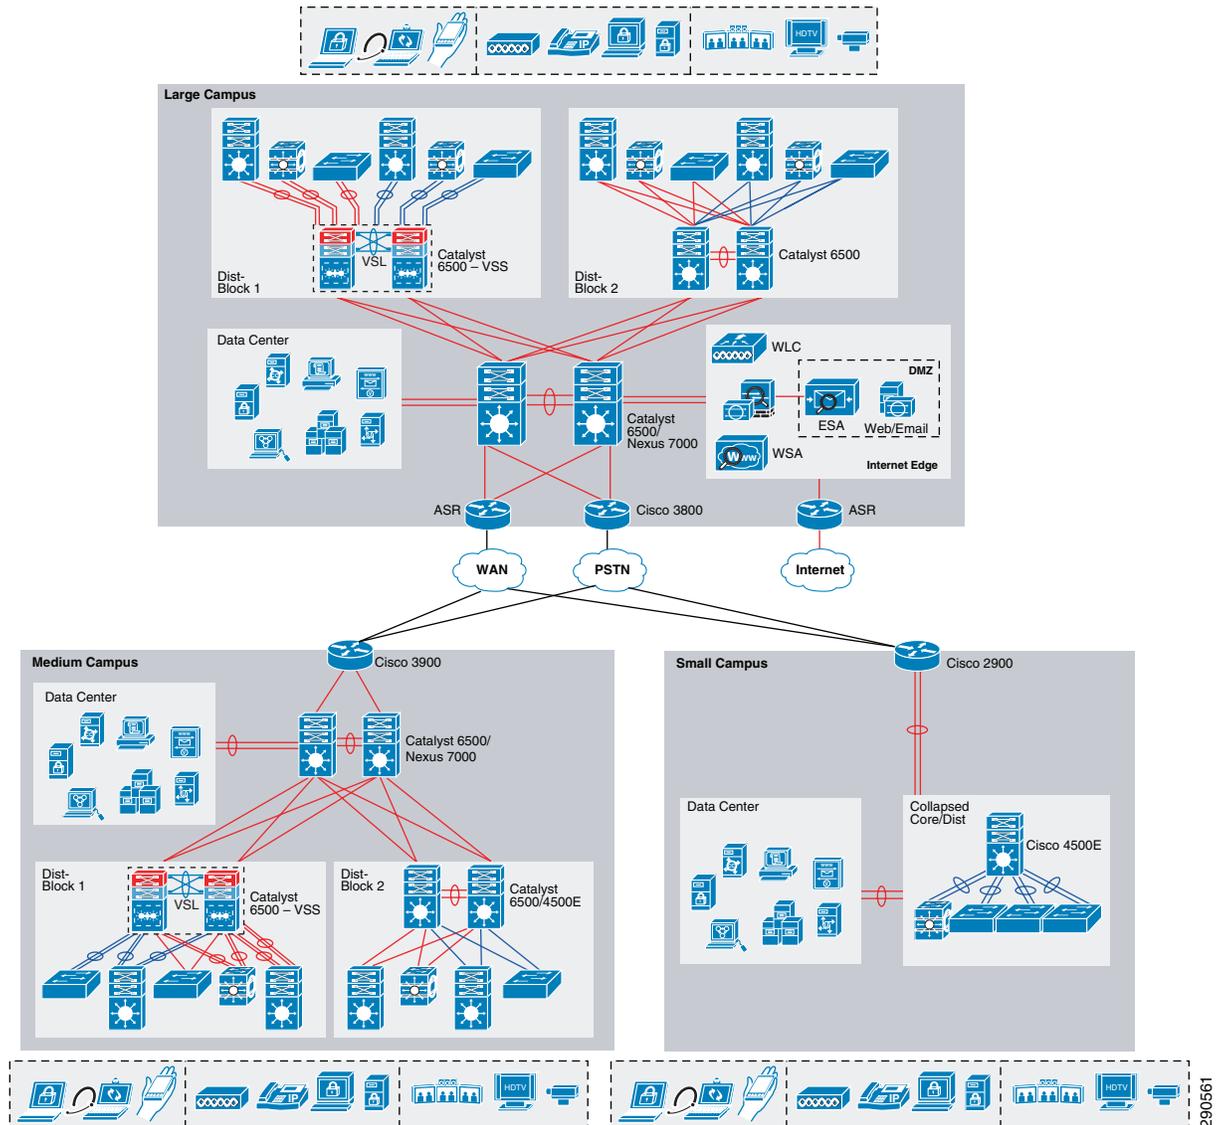
Borderless Campus Network Design Models

Both campus design models (three-tier and two-tier) have been developed with the following considerations:

- *Scalability*—Allowing for network speeds from 100mb to 10gb, the ability to scale a network based on required bandwidth is paramount. The network provides investment protection by allowing for upgradability as bandwidth demand increases.
- *Simplicity*—Reducing operational and troubleshooting cost by the use of network-wide configuration, operation, and management.
- *Resiliency*—Ability to provide non-stop business communication with rapid sub-second network recovery during abnormal network failures or even network upgrades.
- *Cost-effectiveness*—Integrated specific network components that fit budgets without compromising design principles and network performance.

As shown in [Figure 1-7](#), multiple campuses can co-exist within a single enterprise system that offers borderless network services.

Figure 1-7 Borderless Campus Network Design Model



Depending on the medium and small campus office facility, the number of employees and the networked devices in remote campuses may be equal to or less than the large campus. Hence compared to the large campus network, the medium and small campus sites may have alternate network designs that can provide network services based on overall campus network capacity.

Using high-speed WAN technology, several medium and small enterprise campuses can interconnect to a centralized large campus that provides protected shared data and network services to all employees independent of their physical location.

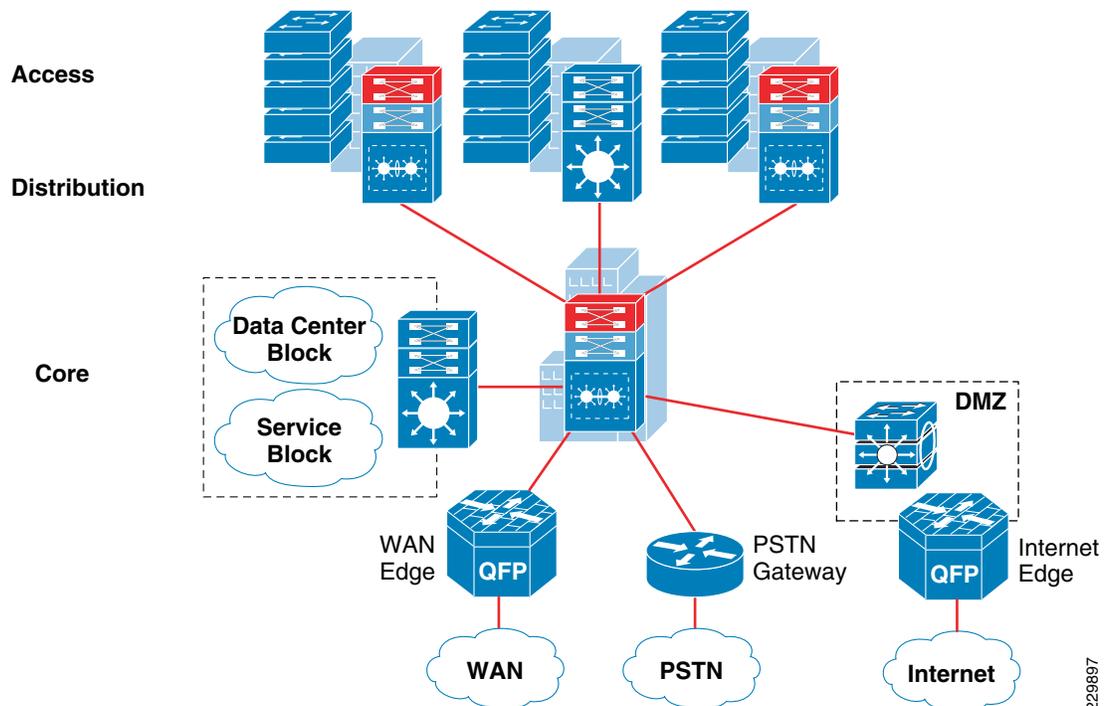
Table 1-1 shows a summary of the Borderless Campus Network design models as they are applied in different overall enterprise network designs.

Table 1-1 Enterprise Recommended Campus Design Models

Enterprise Location	Recommended Campus Design Model
Large campus	Three-tier
Medium campus	Three-tier
Small campus	Two-tier

Large Campus Network Design

The large campus in the enterprise design consists of a centralized hub campus location that interconnects medium and small campuses of several sizes to provide end-to-end shared network access of resources and borderless services. The large campus typically consists of various sizes of building facilities and various organizational and departmental groups. The network scale in the large campus is higher than the medium and small campus networks and includes end users, IP-enabled endpoints, servers, security, and network edge devices. Multiple buildings of various sizes exist in one location, as shown in [Figure 1-8](#).

Figure 1-8 Large Campus Reference Design

The three-tier campus design model for the large campus meets all key technical aspects to provide a well-structured and strong network foundation. The modularity and flexibility in a three-tier campus design model allows easier expansion and integration in the large campus network and keeps all network elements protected and available.

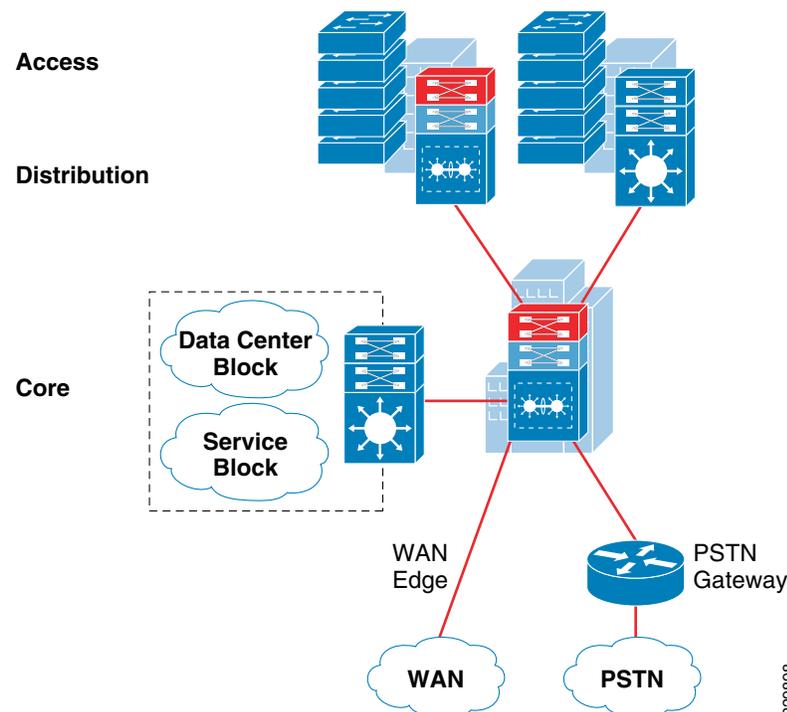
To enforce external network access policies for end users, the three-tier model in the large campus also provides external gateway services to employees for accessing the Internet.

Medium Campus Network Design

From a location, size, and network scale perspective, the medium campus is not much different than the large campus. Geographically, it can be distant from the large campus and require a high-speed WAN circuit to interconnect both campuses. The medium campus can also be considered as an alternate campus to the large campus, with the same common types of applications, endpoints, users, and network services. Similar to the large campus, separate WAN devices are recommended to provide application delivery and access to the large campus given the size and number of employees at this location.

Similar to the large campus network design, Cisco recommends the three-tier campus design model for the medium campus, as shown in [Figure 1-9](#).

Figure 1-9 Medium Campus Reference Design



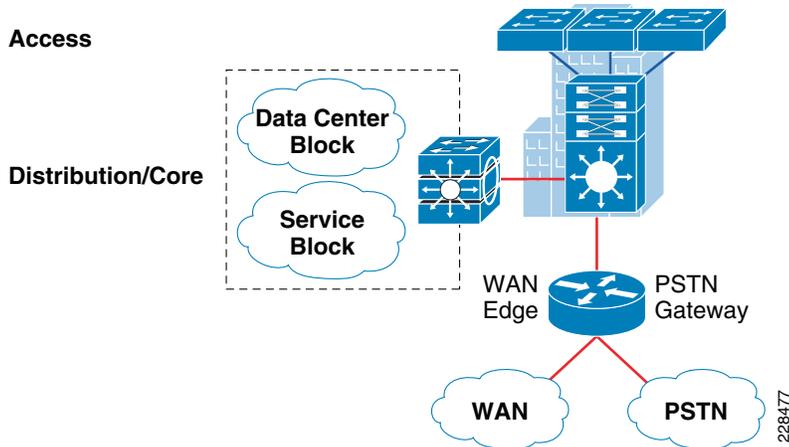
Small Campus Network Design

The small campus is typically confined to a single building that spans multiple floors with different organizations. The network scale factor in this design is reduced compared to other large and medium campuses. However, application and borderless services demands are still consistent across the enterprise geographical locations.

In such smaller scale campus network deployments, the distribution and core layer functions can be alternatively collapsed into the two-tier campus model without compromising basic network design principles. Prior to deploying the collapsed core and distribution system, network architects must consider the scale, expansion, and manageability factors which may reduce overall operational efficiency.

WAN bandwidth requirements must be assessed appropriately for the remote small campus network design. Although the network scale factor is reduced compared to other larger campus locations, sufficient WAN link capacity is needed to deliver consistent network services to users. A single Cisco platform in a highly-redundant configuration mode can provide collapsed core and distribution LAN layers. This alternate and cost-effective network design model is recommended only in smaller locations; WAN traffic and application needs must be considered. Figure 1-10 shows the small campus network design in more detail.

Figure 1-10 *Small Campus Reference Design*



Multi-Tier Borderless Campus Design Models

The previous section discussed various recommended campus design models for each enterprise location. This section provides more detailed network infrastructure guidance for each tier in the campus design model. Each design recommendation is optimized to keep the network simplified and cost-effective without compromising network scalability, security, and resiliency. Each campus design model for an enterprise location is based on the three parts of the campus network architecture—core, distribution, and access layers.

Campus Core Layer Network Design

As described in the previous section, the core layer is the center point of the network and becomes a high-speed transit point between multiple distribution blocks and other systems that interconnect to the services block, the WAN, and the campus edge. The common design in large networks is to build a high-performance, scalable, reliable, and simplified core.

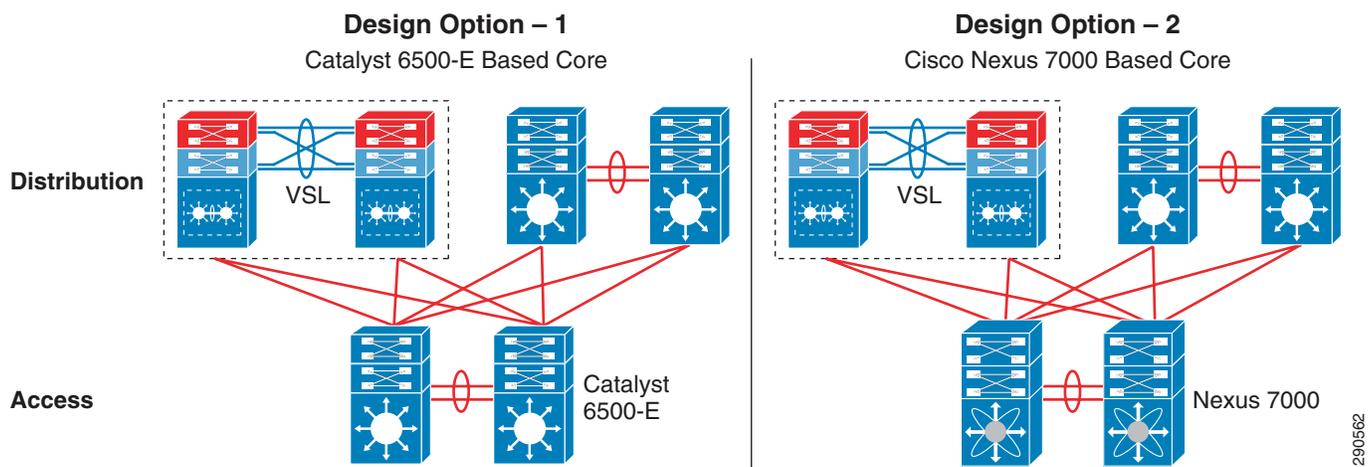
When network architects are designing a campus core, it becomes imperative to take into consideration network scalability, capacity, and reliability to allow for high-performance, end-to-end borderless services. Quantifying the core layer scalability and performance may be challenging as it varies depending on the needs of the enterprise. In campus core design, large enterprise networks are largely built with highly-resilient systems and high-speed 10Gbps links. Network architects must proactively foresee the expansion, evolution, and advancement of devices and applications on the network that may impact the core.

Cisco recommends building the next-generation borderless campus core with the following principles. The architecture should be:

- Designed to support modern technologies that enable advanced networking and integrated services to solve key business problems.
- Scalable to adapt to enterprise network needs and able to provide intelligent borderless network services.
- Flexible, with design options that maximize return on investment (ROI) and reduce total cost of ownership (TCO).

These design principles are important when designing the core network so that the core is capable of addressing current and future borderless network demands. Cisco recommends the Cisco Catalyst 6500-E and Nexus 7000 switching platforms for the core of the next generation borderless campus. These multi-terabit switching platforms are designed with a robust hardware architecture that exceeds the foundational borderless campus requirements. [Figure 1-11](#) illustrates core designs for building the next-generation Borderless Campus core.

Figure 1-11 Core Layer Design Model Options



Cisco Catalyst 6500-E

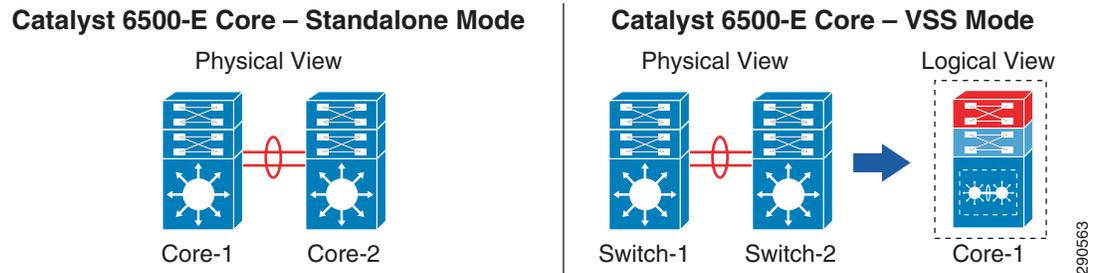
The industry-leading, widely-deployed Cisco Catalyst 6500-E series platform has advanced hardware and software innovations that make it the preferred system to build an enterprise-class borderless campus core network. Cisco Catalyst 6500-E switches have a flexible architecture that enables a rich set of features and advanced technologies, along with the high-speed interfaces needed for the borderless campus. In the large and medium campuses, bandwidth intensive and latency sensitive applications—such as real-time IP-based voice and video—are ubiquitous, so network architects must take this into consideration when selecting the appropriate core platform. As networks expand, the management and troubleshooting of the infrastructure increases, however administrators can leverage Cisco’s system virtualization technology to ease those burdens.

To provide mission-critical network services, it is recommended that the core layer be highly resilient. Deploying resilient, dual Cisco Catalyst 6500-E systems provides constant network availability for business operations during faults and also provides the ability to load share high-speed network traffic between different blocks (e.g., the distribution and service blocks). A redundant core network design can be deployed in a traditional standalone model or in a Virtual Switching System (VSS) model. The campus core layer network design and operation broadly differ when the core layer is deployed as a

standalone, which operates all three planes (forwarding, control and data planes) in isolation. However with Cisco VSS technology, two core systems are clustered into a single logical system and the control and management planes get combined on the systems to produce a single logical Catalyst 6500-E core system.

The Standalone/VSS Physical and Operational View is shown in Figure 1-12.

Figure 1-12 Standalone/VSS Physical and Operational View



Note

For more detailed VSS design guidance, see the *Campus 3.0 Virtual Switching System Design Guide*: http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/VSS30dg/campusVSS_DG.html.

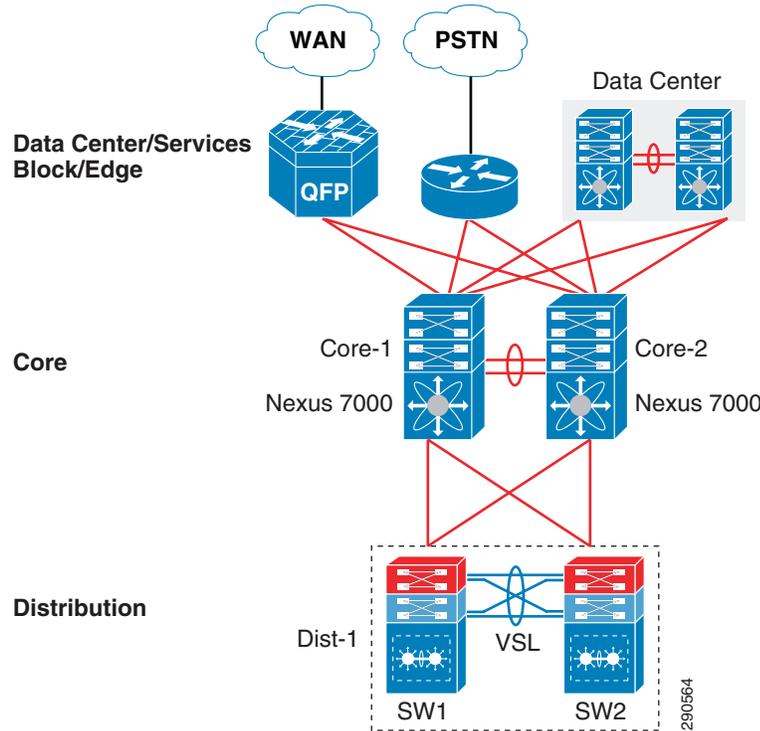
Cisco Nexus 7000

In high-speed and dense networking environments, enterprises require a simplified network architecture that expands the infrastructure's scalability, performance, and reliability. With this in mind, Cisco developed a powerful, multi-terabit switching platform, the Cisco Nexus 7000, to deliver these fundamental requirements. Next-generation data center architectures are built on the Cisco Nexus product family and the Cisco Nexus 7000 series platform leads in data center aggregation and in the data center core networking role.

Because of its unique architecture, technical advantages, and ability to deliver a baseline of campus core requirements, the Cisco Nexus 7000 series can be an alternative platform for deployment in the campus core. In the campus core environment, the Cisco Nexus 7000 offers un-paralleled 10G density to aggregate distribution blocks. It enables low-latency and wire-speed backbone connectivity between the service block and campus edge. The Nexus 7000 utilizes Cisco NX-OS as its operating system, which is a highly-evolved, multithreaded, and modular operating system to deliver core class networking services. NX-OS offers resilient network communication, system virtualization, and several other technical innovations that enable enterprises to have the capabilities needed for the next-generation Borderless Campus network. The Nexus 7000 platform operates in a standalone configuration that locally maintains the control, distributed forwarding, and management planes. For a resilient and mission critical campus core design, the Cisco Nexus 7000 system should be deployed with redundant hardware components that maintain backbone switching capacity and service availability during planned upgrades or un-planned network outages.

Figure 1-13 illustrates core network design options with the Cisco Nexus 7000 peering with other Cisco platforms to enable end-to-end business communication:

Figure 1-13 Cisco Nexus 7000 Campus Core Design

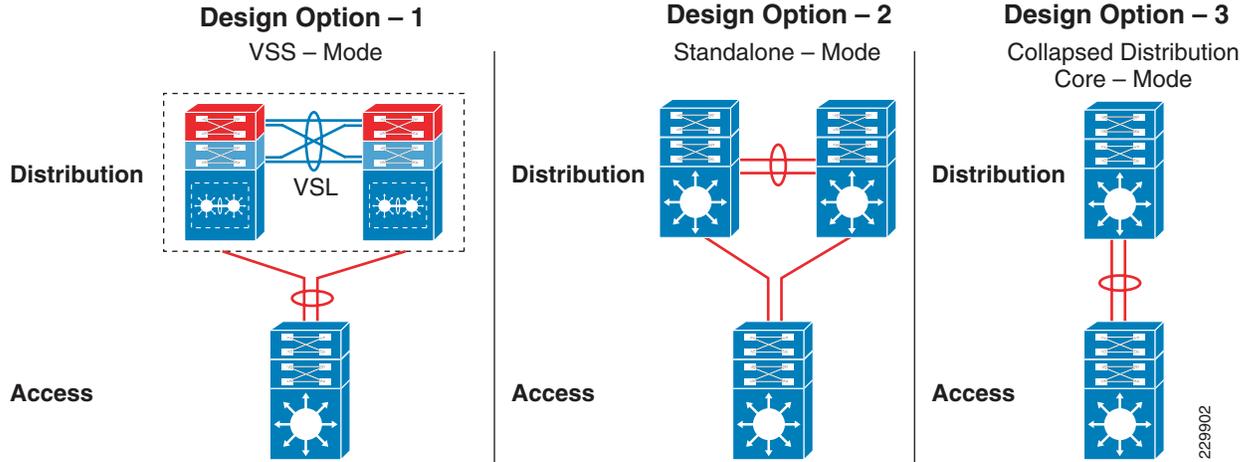


Campus Distribution Layer Network Design

The distribution or aggregation layer is the network demarcation boundary between wiring closet switches and the campus core network. The framework of the distribution layer system in the enterprise design is based on best practices that reduce network complexities, increase reliability, and accelerate network performance. To build a strong campus network foundation with the three-tier model, the distribution layer has a vital role in consolidating networks and enforcing network edge policies.

The distribution layer design options provide consistent network operation and configuration tools to enable various borderless network services. Three simplified distribution layer design options can be deployed in large, medium, and small campus locations, depending on network scale, application and borderless services demands, and cost, as shown in Figure 1-14. All distribution design models offer consistent network foundation services, high availability, expansion flexibility, and network scalability. However each enterprise network is different, with unique business challenges that require a cost-effective aggregation solution, scalability, high-speed network services, virtualized systems, etc., that can be enabled with advanced technologies. Depending on network designs and key technical requirements, network architects must make appropriate aggregation layer design choices to enable end-to-end borderless network services.

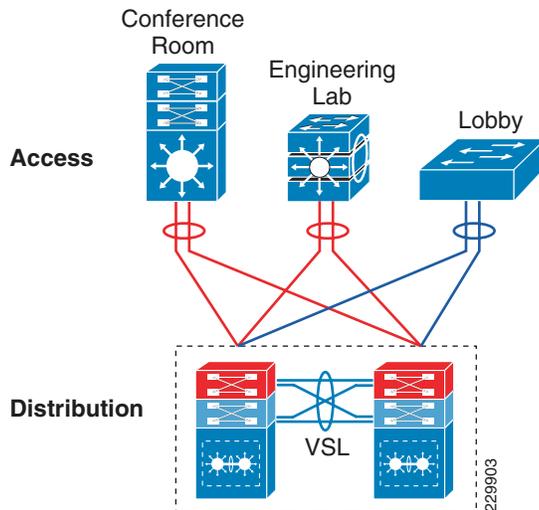
Figure 1-14 Distribution Layer Design Model Options



Distribution Layer Design Option 1—VSS Mode

Distribution layer design option 1 is intended for the large and medium campus network design and it is based on deploying the Cisco Catalyst 6500-E Series switches using Cisco VSS, as shown in [Figure 1-15](#).

Figure 1-15 VSS-Enabled Distribution Layer Network Design



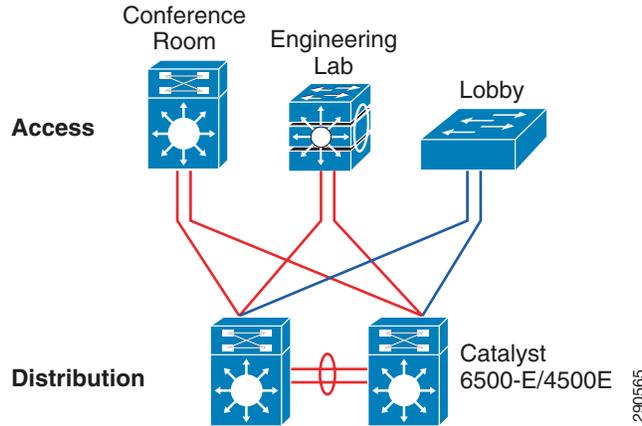
Distribution Layer Design Option 2—Standalone Mode

The distribution layer option 2 is a traditional and proven network design in the enterprise campus network. It can be deployed with redundant Cisco Catalyst 6500 or 4500E systems to operate in standalone mode. This is an alternative distribution network deployment model if there is no desire or capability to virtualize the aggregation layer switches using Cisco VSS technology. In the large campus,

the Cisco Catalyst 6500 with non-VSL capable supervisor modules can be deployed in standalone mode, whereas in the medium campus, network administrators can deploy the Catalyst 6500 or the alternative Catalyst 4500E system in standalone mode.

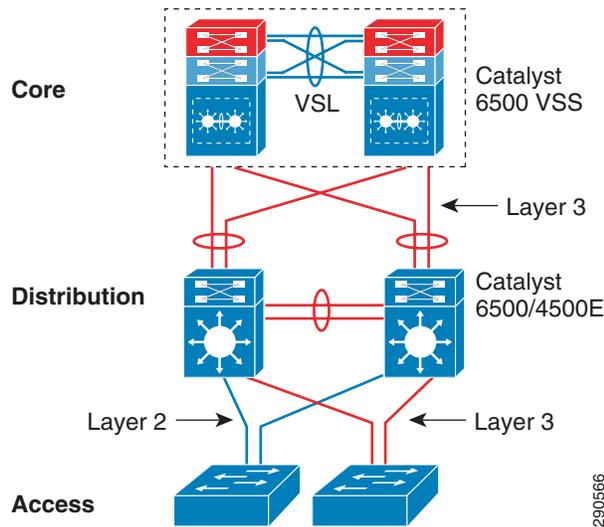
The two single-chassis standalone mode distribution layer design options are shown in Figure 1-16.

Figure 1-16 Standalone Mode Distribution Layer Network Design



In the standalone mode, each Catalyst distribution system operates in independent mode and builds local network adjacencies and forwarding information with access and core layer peering devices. Layer 2 and Layer 3 protocol operation is done over each physical interface between the standalone distribution switches and the access layer switches. Since the core layer network in large and medium campus networks is simplified using Cisco VSS technology, the network administrator can simplify the core network topology by bundling Layer 3 interfaces into a logical EtherChannel, as shown in Figure 1-17.

Figure 1-17 Network Design with Distribution in Standalone Mode



This network design does not raise any significant concerns in Layer 3 network designs. Each standalone distribution system will establish Layer 3 adjacencies with core and access layer (routed access) devices to develop routing topologies and forwarding tables. The traditional multilayer network design faces the following challenges when the access layer switches communicate with two distinct distribution layer switches:

- The multilayer network uses simple Spanning-Tree Protocol (STP) to build Layer 2 loop-free network paths, which results in a sub-optimal and asymmetric forwarding topology.
- It requires per-VLAN virtual gateway protocol operation between aggregation systems to provide high availability. For large networks, First Hop Redundancy Protocol (FHRP) protocols may limit network scalability and consume more system and network resources.
- For a stable, secure, and optimized multilayer network, each distribution and access layer system will require advanced network parameters tuning.
- Layer 2 network recovery becomes protocol type- and timer-dependent. The default protocol parameters could result in network outages for several seconds during faults. Protocol timers can be tuned aggressively for network recovery within a second range, however it cannot meet the high-availability baseline for business-class video applications like Cisco TelePresence.

Cisco innovated VSS technology to mitigate such challenges. Hence it is recommended to deploy a Cisco VSS-based distribution layer infrastructure that simplifies the multilayer network and increases network capacity and performance, resulting in a highly-reliable network that provides consistent and deterministic network recovery. The traditional standalone-mode distribution layer network is an alternative solution that does not introduce any fundamental design changes. Deployment guidance documented in various design guide remains consistent and can be leveraged, hence the standalone distribution network design is not fully re-validated in this document. For more information on configuring and deploying standalone-mode distribution layer Catalyst switches, see the *Campus Network for High Availability Design Guide*:

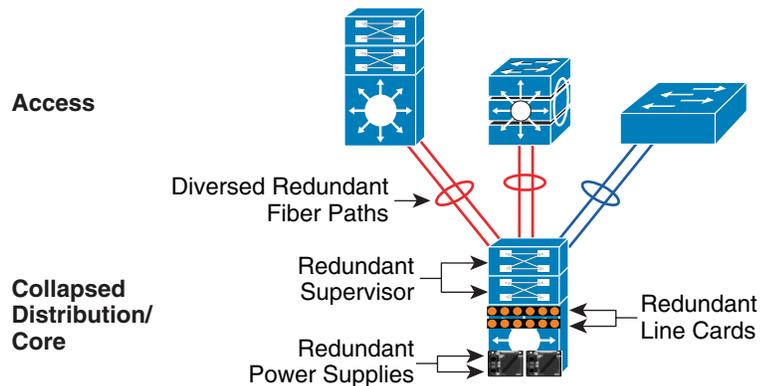
http://www.cisco.com/en/US/partner/docs/solutions/Enterprise/Campus/HA_campus_DG/hacampusdg.html.

Distribution Layer Design Option 3—Collapsed Distribution/Core Mode

The small remote campus location may have several departments working on various floors within a building. Network administrators can consider collapsing the core function into the distribution layer switch for such a small campus where there may only be a single distribution block. The collapsed distribution/core system can provide network services to a small number of wiring closet switches and directly connect to the WAN edge system to reach a large campus for centralized data and communication services. Deploying a two-tier network model for a single distribution block in a small campus does not break the three-tier campus design principles required in large and medium campus networks. This solution is manageable and cost effective as it meets the needs of network users and the number of endpoints is not as large as large or medium enterprise campuses.

The collapsed distribution/core network can be deployed with two redundant systems as recommended in [Distribution Layer Design Option 1—VSS Mode](#) or alternatively in standalone mode as described in [Distribution Layer Design Option 2—Standalone Mode](#). In a space-constrained small campus environment, a single Cisco Catalyst 4500E series platform can be deployed with multiple redundant hardware components. Building a single, highly-available, collapsed distribution/core system will ensure the network performance, availability, and reliability required to run borderless services. With various redundant hardware components, this solution can provide 1+1 in-chassis protection against various types of hardware and software failure. Deploying the network in a recommended design will provide consistent sub-second network recovery. A single Cisco Catalyst 4500E with multiple redundant system components can be deployed as shown in [Figure 1-18](#).

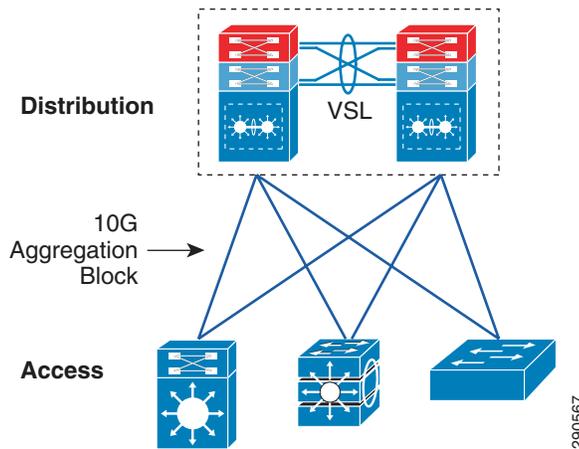
Figure 1-18 *Highly Redundant Single Collapsed Distribution/Core Design*



Campus Access Layer Network Design

The access layer is the first tier or edge of the campus, where end devices such as PCs, printers, cameras, Cisco TelePresence, etc. attach to the wired portion of the campus network. It is also the place where devices that extend the network out one more level, such as IP phones and wireless access points (APs), are attached. The wide variety of possible types of devices that can connect and the various services and dynamic configuration mechanisms that are necessary make the access layer one of the most feature-rich parts of the campus network. Not only does the access layer switch allow users to access the network, the access layer switch provides network protection so that unauthorized users or applications do not enter the network. The challenge for the network architect is determining how to implement a design that meets this wide variety of requirements—the need for various levels of mobility, the need for a cost-effective and flexible operations environment, etc.—while being able to provide the appropriate balance of security and availability expected in more traditional, fixed-configuration environments. The next-generation Cisco Catalyst switching portfolio includes a wide range of fixed and modular switching platforms, each designed with unique hardware and software capabilities to function in a specific role.

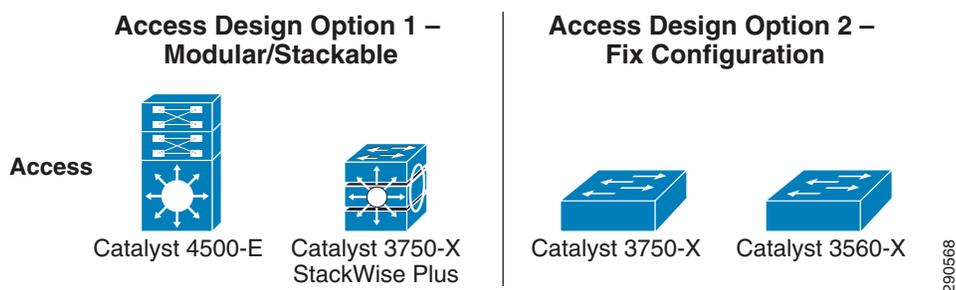
Enterprise campuses may deploy a wide range of network endpoints which all have different requirements on the network; low-latency, link speed, and low-jitter rates are just some of those requirements. The network architect must consider network requirements, as well as the planned growth of network resources, when determining bandwidth requirements for the access layer to distribution uplinks. To build a high-performance distribution-access block, Cisco access layer switching platforms are designed with 10Gbps uplinks to provide borderless network services at wire-rate.

Figure 1-19 High-Performance Distribution-Access Block

Building a 10Gbps distribution-access block provides the following benefits:

- Increased throughput—Increases network bandwidth capacity ten-fold on a per-physical-port basis. The oversubscription bandwidth ratio in a high-density wiring closet falls within the recommended range.
- High performance—Accelerates application performance by multiplexing a large number of flows onto a single high-speed connection instead of load-sharing across multiple slow aggregate links.
- Reduced TCO—The cost of access switches becomes less per port; it reduces additional cost by deploying fewer cables and connectors when building parallel paths between two systems.
- Simplified design—Single high-speed link to manage, operate, and troubleshoot instead of multiple individual or aggregated bundled connections.

Based on the broad range of business communication devices and endpoints, network access demands, and capabilities, two access layer design options can be deployed, as shown in [Figure 1-20](#).

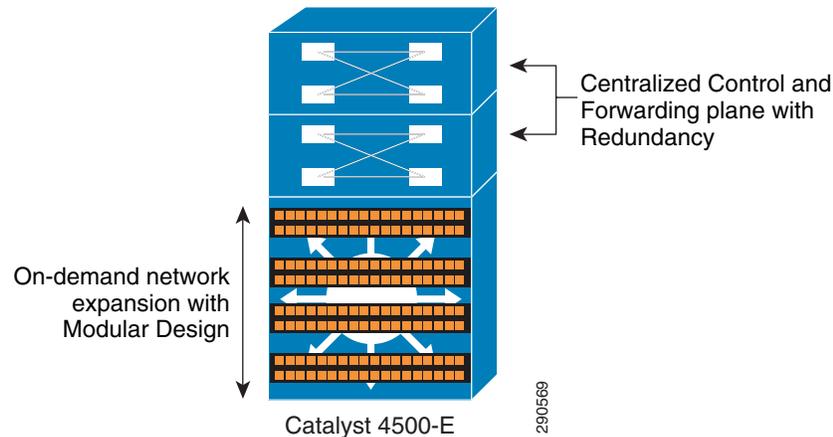
Figure 1-20 Access Layer Design Models

Access Layer Design Option 1—Modular/StackWise Plus Access Layer Network

Access layer design option 1 is intended to address network modularity, performance, scalability, and availability for IT-managed, critical voice and video communication edge devices. To accelerate the user experience and campus physical security protection, these devices require low latency, high performance, and a constantly-available network switching infrastructure. Implementing a modular and stackable Cisco Catalyst switching platform provides the flexibility to increase network scalability in the densely-populated campus network edge.

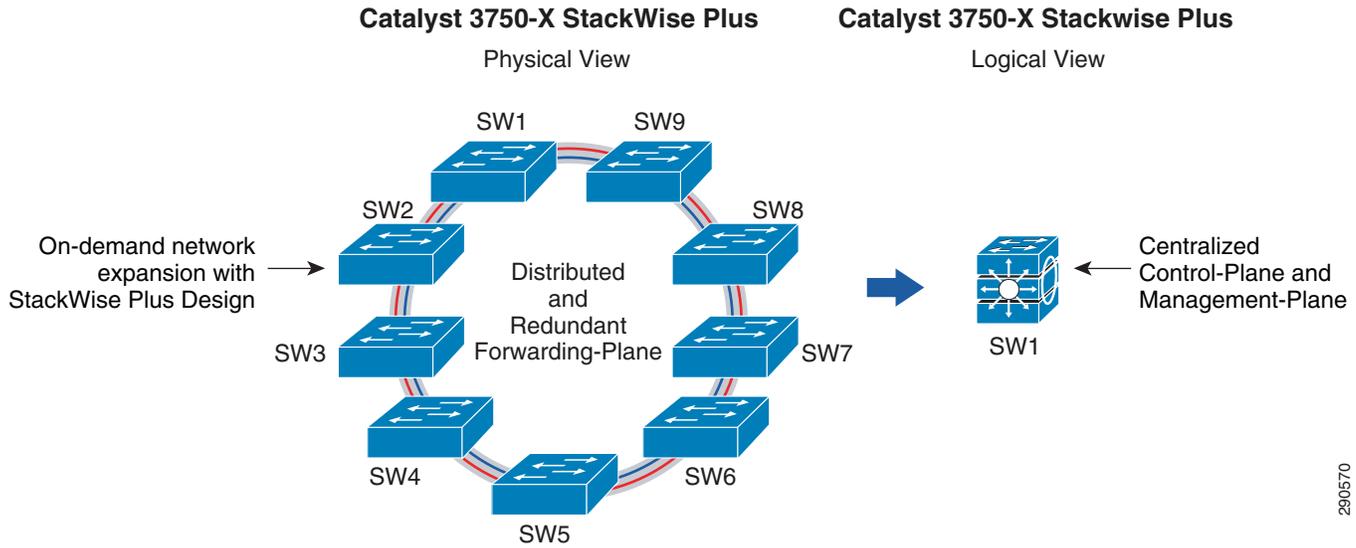
In large and medium campus deployments, the ubiquitous Cisco Catalyst 4500E Series platform provides a scalable, high-speed, and robust network solution. In a high-density access environment, it is imperative to simplify the management of hundred of end points through a single chassis. It is also essential during hardware or software failures to provide wire-speed network performance without compromising network reliability by using a non-stop forwarding architecture. The next-generation hardware architecture of the Cisco Catalyst 4500E in the campus access layer leverages new Cisco IOS software capabilities to enable several borderless network services at the campus network boundary.

Figure 1-21 Network Edge Expansion with Modular Design



The Cisco Catalyst 3750-X Series is the alternative Cisco access layer switching platform. Using Cisco StackWise Plus technology provides flexibility and availability by clustering multiple Cisco Catalyst 3750-X Series Switches into a single high-speed stack ring that simplifies operation and allows incremental access layer network expansion or contraction. Catalyst 3750-X switches deployed in Cisco StackWise Plus mode alters network operation compared to standalone mode. When deployed in StackWise plus mode, the switches become a single logical access layer switch, the control plane processing becomes centralized, and because of the distributed forwarding architecture, all the hardware resources gets fully utilized across all stack member switches (see [Figure 1-22](#)). Cisco StackWise Plus provides high-speed multigigabit switching capacity for network traffic switching within the stack-ring and the distribution-access block can be built with multiple parallel 10Gbps uplink paths for load sharing and network resiliency. The network is optimized and simplified when the cross-switch uplink ports are bundled into a single logical interface using EtherChannel technology. This network design provides non-stop network communication in case of the failure of an individual stack member switch.

Figure 1-22 Network Edge Expansion with StackWise Plus Design



290570

Access Layer Design Option 2—Fixed Configuration Access Layer Network

This entry-level access layer design option is widely chosen for enterprise environments. The fixed configuration Cisco Catalyst switching portfolio supports a wide range of access layer technologies that allow seamless service integration and enable intelligent network management at the edge. Fixed configuration Cisco Catalyst switches in standalone mode are an ideal design choice for a small size wiring closet to provide consistent borderless network services for up to 48 endpoints.

The next-generation fixed configuration Cisco Catalyst 3750-X and 3560-X Series are commonly deployed platforms for wired network access that can be in a mixed configuration with critical devices, such as Cisco IP phones, and non-mission critical endpoints, such as library PCs, printers, and so on. For non-stop network operation during power outages, the Catalyst 3560-X must be deployed with an internal or external redundant power supply solution using the Cisco RPS 2300. Increasing aggregated power capacity provides the flexibility to scale with enhanced Power-over-Ethernet (PoE+) on a per-port basis. With its wire-speed 10G uplink forwarding capacity, this design reduces network congestion and latency to significantly improve application performance.

To provide a consistent end-to-end enhanced user experience, the Cisco Catalyst 3750-X and 3560-X Series platforms support critical network control services to secure the network edge and intelligently provide differentiated services to various class-of-service traffic, as well as simplified management. The Cisco Catalyst must leverage the dual uplink ports to interconnect the distribution system for increased bandwidth capacity and network availability.

Both design options offer consistent network services at the campus edge to provide differentiated, intelligent, and secured network access to trusted and untrusted endpoints. The distribution options recommended in the previous section can accommodate both access layer design options.

Summary

As enterprises make a fundamental shift in their networks to meet the new demands of employees, customers, and partners, network infrastructure design decisions are critical. The Borderless Campus 1.0 CVD describes the design decisions required for an enterprise network campus. The Borderless Campus

1.0 architecture provides an architecture that showcases Cisco's best practices. This chapter discusses the design options for each layer in the Borderless Campus 1.0 Architecture. The remaining chapters describe implementation and deployment options.

