

Mobility Design

The Cisco Community College reference design is intended to assist community colleges in the design and deployment of advanced network-based solutions within twenty-first century learning environments.

The reference design addresses the business challenges currently facing community colleges. At the heart of the reference design is the network service fabric, which is a collection of products, features, and technologies that provide a robust routing and switching foundation upon which all solutions and services are built. Operating on top of the network service fabric are all the services used within the community college network to solve business problems, which include the following:

- Safety and security
- Virtual learning
- Secure connected classrooms
- Operational efficiencies

Community college students are dynamic, mobile, and technology-savvy. When on campus, they move about while equipped with an array of mobility-enabled devices including PDAs, phones, and laptops. In contrast to the typical enterprise business environment, community colleges consist of a large student population that typically experiences a complete turnover every few years. Typical community college students tend to use new applications and the network for many aspects of their lives, demanding connectivity wherever they are. This connected generation is untethered from wired access connectivity and assumes the presence of a high-performance, reliable wireless LAN (WLAN) in all major campus areas.

The mobility design implemented by a community college must meet the needs of this mobile generation while also addressing the requirements of faculty, staff, administrators, and visitors. The challenge for community colleges is to create a robust, end-to-end, mobility-enabled network that supports their requirements at a cost that is within their often constrained budgets. Community colleges should be equipped with a mobility solution that supports the following:


- Secure communications between local and remote campus sites to support students, faculty, staff, administrators, and visitors, using the new generation of mobility-enabled devices and applications in the current marketplace
- A scalable design model that can easily accommodate the addition of new campus buildings as well as existing building modifications
- Built-in support for bandwidth-intensive, high-speed multimedia applications
- Simplified management tools to facilitate maintenance of the system-wide mobility solution
- The use of new tools and applications for mobile learning, collaboration, and campus operations
- Effective communication with public safety first responders in the event of an emergency

In addition, each community college must remain competitive, differentiating itself from its peer institutions so as to attract and retain the best students and faculty. Students want to attend quality community colleges that provide technology services relevant to the way they live, work, and learn. They want to take full advantage of community college capabilities to facilitate their success while they are students, as well as when they are pursuing post-graduation placement. A community college with a pervasive, high-speed wireless network not only demonstrates technological leadership and innovation, but enables the deployment of innovative applications that improve learning, the streamlining of operations, collaboration enhancements, and productivity improvements.

This mobile campus lifestyle helps to drive the need for careful wireless capacity and coverage planning. Keep in mind that the traditional scenario of a mass of students filing into a large lecture hall within a monolithic campus building is no longer the only learning environment seen within higher educational institutions. High performance, secure wireless technologies can enable “virtual classrooms” even in non-traditional settings, such as leased space in shopping malls, retail plazas, and even from homes and offices. School administrators need secure access to tools, records, and resources, as well as access to mobile voice capabilities throughout the campus. In addition, the expectation for secure, reliable, high-performance guest access by contractors, vendors, and other guests of the community college establishment has become a standard and expected component of doing business.

To meet these and other student, faculty, and guest needs, community colleges must evolve into mobility-enabled campuses and twenty-first century learning centers. The primary objectives of this document are the design considerations surrounding the requirements and expectations that must be considered when integrating mobility into the Cisco Community College reference design, as well as the tradeoffs required to facilitate the four service requirements stated previously. These design considerations form a critical part of the overall service fabric design model, shown in [Figure 1](#).

Figure 1 Service Fabric Design Model




Given the mobility of students, staff, and visitors, wireless LANs have emerged as one of the most effective and high performance means for these mobile users to access the campus network. The Cisco Unified Wireless Network (Cisco UWN) is a unified solution that addresses the wireless network security, deployment, management, and control

aspects of deploying a wireless network. It combines the best elements of wireless and wired networking to deliver secure, scalable wireless networks with a low total cost of ownership.


Figure 2 shows a high-level topology of the Cisco Unified Network, which includes access points that use the Control and Provisioning of Lightweight Access Points (CAPWAP) protocol; the Cisco Wireless Control System (WCS); and the Cisco Wireless LAN Controller (WLC). In addition to the traditional standalone WLAN controller, alternate hardware platforms include the Cisco ISR router Wireless LAN Controller Module (WLCM) or the Cisco Catalyst 6500 Wireless Services Module (WiSM). The Cisco Access Control Server (ACS) and its Authentication, Authorization, and Accounting (AAA) features complete the solution by providing Remote Authentication Dial-In User Service (RADIUS) services in support of user authentication and authorization.

Figure 2 Cisco Unified Wireless Network Overview



The Cisco Community College reference design accommodates a main campus and one or more remote smaller campuses interconnected over a metro Ethernet or managed WAN service. Each of these campuses may contain one or more buildings of varying sizes, as shown in Figure 3.

Figure 3 Community College Reference Design Overview



Operating on top of this network are all the services used within the community college environment such as safety and security systems, voice communications, video surveillance equipment, and so on. The core of these services are deployed and managed at the main campus building, allowing each remote campus to reduce the need for separate services to be operated and maintained by community college IT personnel. These centralized systems and applications are served by a data center in the main campus.

As Figure 3 shows, the Cisco Community College reference design uses a centralized approach in which key resources are centrally deployed at either the campus or college level. The key feature of this integration is the use of one or more WLAN controllers at each campus, with the overall WLAN management function (the Cisco WCS) located at the main campus. This approach simplifies the deployment and operation of the network, helping to ensure smooth performance, enhance security, enhance network maintainability, maximize network availability, and reduce overall operating costs.

The Cisco Community College reference design takes into account that cost and limited network administrative resources are common limiting factors for most community colleges. The topologies and platforms are carefully selected to increase productivity while minimizing the overall cost and complexity of operation. In certain instances, tradeoffs are necessary to reach these goals, and this document points out such areas.

The Cisco mobility approach within the Cisco Community College reference design focuses on the following key areas:

- *Accessibility*

- Enabling students, staff, and guests to be accessible and productive on the network, regardless of whether they are in a traditional classroom setting, collaborating in a study hall, having lunch with colleagues within campus eating areas, or simply enjoying a breath of fresh air outside a campus building
- Enabling easy, secure guest access to college guests such as alumni, prospective students, contractors, vendors and other visitors.

- *Usability*

In addition to extremely high WLAN transmission speeds made possible by the current generation of IEEE 802.11n technology, latency-sensitive applications (such as IP telephony and video conferencing) are supported over the WLAN using appropriately applied quality-of-service (QoS) classification. This gives preferential treatment to real-time traffic, helping to ensure that video and audio information arrives on time.

- *Security*

- Segmenting authorized users and blocking unauthorized users
- Extending the services of the network safely to authorized parties
- Enforcing security policy compliance on all devices seeking to access network computing resources. Faculty and other staff enjoy rapid and reliable authentication through IEEE 802.1x and Extensible Authentication Protocol (EAP), with all information sent and received on the WLAN being encrypted.

Note For information on how security design is addressed within the Cisco Community College reference design, see the *Community College Security Design Considerations* document.

- *Manageability*

A relatively small team of college network administrators must be able to easily deploy, operate, and manage hundreds of access points that may reside within a multi-campus community college. A single, easy-to-understand WLAN management framework provides small, medium, and large community colleges with the same level of WLAN management scalability, reliability, and ease of deployment demanded by traditional enterprise business customers.

- *Reliability*

- Providing adequate capability to recover from a single-layer fault of a WLAN access component or controller wired link
- Ensuring that WLAN accessibility is maintained for students, faculty, staffs and visitors in the event of common failures

Accessibility


This section provides a brief introduction to the fundamental protocol used for communication between access points and WLAN controllers, followed by a discussion of mobility design considerations pertaining to those aspects of the Cisco Community College reference design relevant to accessibility, such as the following:

- WLAN controller location
- WLAN controller connectivity
- Access points

The basic mobility components involved with providing WLAN access in the Cisco Community College reference design consists of WLAN controllers and access points that communicate with each other using the IETF standard CAPWAP protocol. In this arrangement, access points provide the radio connection to wireless clients, and WLAN controllers manage the access points and provide connectivity to the wired network.

Figure 4 shows the use of CAPWAP by access points to communicate with and tunnel traffic to a WLAN controller.

Figure 4 CAPWAP Access Point to WLC Communication



CAPWAP enables the controller to manage a collection of wireless access points, and has the following three primary functions in the mobility design:

- Control and management of the access point
- Tunneling of WLAN client traffic to the WLAN controller
- Collection of 802.11 data for overall WLAN system management

CAPWAP is also intended to provide WLAN controllers with a standardized mechanism with which to manage radio-frequency ID (RFID) readers and similar devices, as well as enable controllers to interoperate with third-party access points in the future.

In controller software Release 5.2 or later, Cisco lightweight access points use CAPWAP to communicate between the controller and other lightweight access points on the network. Controller software releases before Release 5.2 use the Lightweight Access Point Protocol (LWAPP) for these communications. Note that most CAPWAP-enabled access points are also compatible with the preceding LWAPP protocol. An exception is that the Cisco Aironet 1140 Series Access Point supports only CAPWAP.

The mobility approach in the Cisco Community College reference design is based on the feature set available in Cisco Wireless LAN Controller software Release 6.0, which uses CAPWAP.

For detailed CAPWAP protocol information, see the following URL:

<http://www.ietf.org/rfc/rfc5415.txt>.

WLAN Controller Location

WLAN campus deployments are typically categorized into two main categories, *distributed* and *centralized*:

- *Distributed controller*—In this model, WLAN controllers are located throughout the campus network, typically on a per-building basis, and are responsible for managing the access points resident in a given building. This technique is commonly used to connect controllers to the campus network using distribution routers located within

each building. In the distributed deployment model, the CAPWAP tunnels formed between access points and WLAN controllers are typically fully contained within the confines of the building.

- *Centralized controller*—In this model, WLAN controllers are placed at a centralized location in the network. Because centralized WLAN controllers are typically not located in the same building as the access points they manage, the CAPWAP tunnels formed between them must traverse the campus backbone network.

The Cisco Community College reference design is based on the centralization of WLAN controllers, on a per-campus basis, and follows established best practices, such as those contained in Chapter 2 of the *Enterprise Mobility 4.1 Design Guide* at the following URL: http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns820/landing_ent_mob_design.html.

Figure 3 shows the planned deployment of WLAN controllers within distinct per-campus service blocks, each associated with the main, large remote, medium remote, and small remote campus sites respectively. Service blocks tend to be deployed at locations in the network where high availability routing, switching, and power is present. In addition, these areas tend to be locally or remotely managed by network staff possessing higher skill sets.

Some of the advantages underlying the decision to centralize the deployment of WLAN controllers on a per-campus basis include the following:

- *Reduced acquisition and maintenance costs*—By servicing the needs of all campus users from a central point, the number of WLAN controller hardware platforms deployed can be reduced compared to that required for a distributed, per-building design. Similarly, incremental software licensing costs associated with WLAN controllers are reduced as well. These economies of scale typically increase with the size of the campus WLAN.
- *Reduced administrative requirements*—By minimizing the total number of WLAN controllers deployed, the controller management burden imposed on community college campus network administrators is minimized.
- *Cost-effective capacity management*—The use of a centralized WLAN controller model allows the designer the ability to centrally service access points located in multiple building locations and efficiently manage controller capacity.
- *Simplified network management and high availability*—Centralized WLAN controller designs simplify overall network management of controllers, as well as facilitate cost-effective controller high availability approaches. This can protect the campus from a loss of WLAN access in the rare event of a controller failure, without the expense of 1:1 controller duplication.
- *Reduced component interaction points*—Centralizing WLAN controllers minimizes the number of integration points that must be managed when interfacing the controller with other devices. When integrating the WLAN controller with the Network Admission Control (NAC) appliance on any given campus, for example, only one integration point must be administered.
- *Increased performance and reliability*—Centralized WLAN controller deployments usually lead to highly efficient inter-controller mobility. For large campuses, there is also an incremental economy of scale that occurs as the network grows larger. By centralizing WLAN controllers on a per-campus basis, CAPWAP tunneling between access points and WLAN controllers is not normally required to traverse WAN links (except during controller failover), thereby conserving WAN bandwidth and improving performance overall.

Note For additional information on inter-controller mobility and roaming, see the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/h2_Arch.html#wp1028197.

The choice of WLAN controller for the Cisco Community College reference design is the Cisco 5508 Wireless Controller, as shown in Figure 5.

Figure 5 Cisco 5508 Wireless Controller



The Cisco 5508 Wireless Controller is a highly scalable and flexible platform that enables system-wide services for mission-critical wireless in medium to large-sized enterprises and campus environments. Designed for 802.11n performance and maximum scalability, the Cisco 5508 Wireless Controller offers the ability to simultaneously manage from 12 to a maximum of 250 access points per controller. Base access point controller licensing provides the flexibility to purchase only the number of access point licenses required, with the ability to add additional access point licenses in the future when community college campus growth occurs. In campuses requiring more than 250 total access points, or load sharing/high availability is required, multiple controllers can be deployed as necessary.

More information on the Cisco 5508 Wireless Controller can be found at the following URL:

http://www.cisco.com/en/US/prod/collateral/wireless/ps6302/ps8322/ps10315/data_sheet_c78-521631.html.

WLAN Controller Connectivity

This section discusses WLAN controller connectivity, including the following:

- Controller connectivity to the wired network
- Controller connectivity to the wireless devices
- Defining WLANs and Service Set Identifiers (SSIDs)
- WLAN controller mobility groups
- WLAN controller access point groups
- WLAN controller RF groups

Controller Connectivity to the Wired Network

WLAN controllers possess physical entities known as *ports* that connect the controller to its neighboring switch (the Cisco 5508 Wireless Controller supports up to eight Gigabit Ethernet Small Form-Factor Pluggable [SFP] ports). Each physical port on the controller supports, by default, an 802.1Q VLAN trunk, with fixed trunking characteristics.

Note For more information concerning the various types of ports present on Cisco WLAN controllers, see the *Cisco Wireless LAN Controller Configuration Guide, Release 6.0* at the following URL:
<http://www.cisco.com/en/US/docs/wireless/controller/6.0/configuration/guide/Controller60CG.html>.

Interfaces are logical entities found on the controller. An interface may have multiple parameters associated with it, including an IP address, default gateway, primary physical port, optional secondary physical port, VLAN identifier, and Dynamic Host Configuration Protocol (DHCP) server. Each interface is mapped to at least one primary port, and multiple interfaces can be mapped to a single controller port.

Note For more information concerning the various types of interfaces present on Cisco WLAN controllers, see the *Cisco Wireless LAN Controller Configuration Guide, Release 6.0* at the following URL:
<http://www.cisco.com/en/US/docs/wireless/controller/6.0/configuration/guide/Controller60CG.html>.


A special type of controller interface is known as the *AP manager interface*. A controller has one or more AP manager interfaces, which are used for all Layer 3 communications between the controller and its joined access points. The IP address of the AP manager interface is used as the tunnel source for CAPWAP packets from the controller to the access point, and as the destination for CAPWAP packets from the access point to the controller. The AP manager interface communicates through a distribution system port by listening across the Layer 3 network for CAPWAP “join” messages generated by access points seeking to communicate with and “join” the controller.

Link aggregation (LAG) is a partial implementation of the 802.3ad port aggregation standard. It bundles all of the controller distribution system ports into a single 802.3ad port channel, thereby reducing the number of IP addresses needed to configure the ports on your controller. When LAG is enabled, the system dynamically manages port redundancy and load balances traffic transparently to the user. LAG bundles all the enabled distribution ports on the WLAN controller into a single EtherChannel interface.

Currently published best practices specify either multiple AP manager interfaces (with individual Ethernet links to one or more switches) or link aggregation (with all links destined for the same switch or switch stack) as the recommended methods of interconnecting WLAN controllers with wired network infrastructure. For more information, see the following URL:
<http://www.cisco.com/en/US/docs/wireless/controller/6.0/configuration/guide/c60mint.html#wpl277659>.

In the Cisco Community College reference design, the Cisco 5508 Wireless Controllers are interconnected with the modular switches or switch stacks found in the services block using link aggregation and EtherChannel exclusively, as shown in [Figure 6](#).

Figure 6 WLAN Controller Link Aggregation to Services Block



In this way, one or more centralized WLAN controllers are connected via the services block to the campus core. This design can make use of up to eight Gigabit Ethernet connections from the Cisco 5508 Wireless Controller to the services block. These Gigabit Ethernet connections should be distributed among different modular line cards or switch stack members as much as possible, so as to ensure that the failure of a single line card or switch stack failure does not result in total failure of the WLAN controller connection to the campus network. The switch features required to implement this connectivity between the WLAN controller and the services block are the same switch features that would otherwise be used for EtherChannel connectivity between switches in general.

Further discussion of the advantages of using controller link aggregation, as well as the considerations concerning its implementation in the Cisco Community College reference design can be found in [Controller Link Aggregation, page -19](#).

The key advantage of using link aggregation in this fashion instead of multiple AP manager interfaces is design performance, reliability, and simplicity:

- With the Ethernet bundle comprising up to eight Gigabit Ethernet links, link aggregation provides very high traffic bandwidth between the controller and the campus network.
- With link aggregation, if any of the controller ports fail, traffic is automatically migrated to one of the other controller ports. As long as at least one controller port is functioning, the system continues to operate, access points remain connected to the network, and wireless clients continue to send and receive data. Terminating on different modules within a single Catalyst modular switch, or different switch stack members (as shown in [Figure 6](#)), provides redundancy and ensures that connectivity between the services block switch and the controller is maintained in the rare event of a failure.
- Link aggregation also offers simplicity in controller configuration; for example, configuring primary and secondary ports for each interface is not required.

Controller Connectivity to Wireless Devices

This section deals with the design considerations that involve provisioning wireless access for the various user groups that reside within the community college campus system, such as the faculty, administrators, students, and guests. These considerations include the WLAN controllers deployed in the campus services blocks, as well as the access points that are located in the campus buildings.

Defining WLANs and SSIDs

In most community colleges, various campus user groups likely require access to the WLAN for a variety of purposes. Although peaks in usage may occur at different times, it is safe to assume that a large portion of these groups will likely want access to the WLAN at the same time. Thus, in designing for mobility within the Cisco Community College reference design, the physical campus wireless infrastructure needs to support logical segmentation in such a fashion that a reasonable proportion of all users can be serviced simultaneously and with an appropriate degree of security and performance.


One of the basic building blocks used in the WLAN controller to address this need is the ability to provision logical WLANs, each of which are mapped to different wired network interfaces by the WLAN controller. These WLANs are configured and assigned a unique SSID, which is a sequence of characters that uniquely names a WLAN. For this reason, an SSID is also sometimes referred to simply as a *network name*.

Note Each set of wireless devices communicating directly with each other is called a basic service set (BSS). Several BSSs can be joined together to form one logical WLAN segment, referred to as an extended service set (ESS). An SSID is simply the 1–32 byte alphanumeric name given to each ESS.

To promote ease of administration, the value chosen for the SSID should bear some direct relationship to the intended purpose of the WLAN.

Figure 7 provides a high-level illustration of the four logical WLANs that provide mobility within the Cisco Community College reference design, and how they are mapped to WLAN controller network interfaces or tunneled to another controller. For ease of administration and the support of students, faculty, and guests that frequent multiple campuses, the names chosen for the WLAN SSIDs should be consistent within each campus in the community college system. For example, student wireless access should be available anywhere there is WLAN RF coverage within this particular community college system using the SSID entitled *student*.

Figure 7 WLAN SSIDs



In the Community College reference design, the set of WLAN SSIDs provide access to the following WLANs:

- A *secured staff* WLAN network with dynamically generated per-user, per-session encryption keys.

This WLAN would be used by college faculty, staff, and administration using managed client devices, such as laptops, PDAs, and so on. The secured staff WLAN is designed to provide secure access and good performance for devices controlled by the community college network administration staff. Unlike the student and guest access WLANs, devices that are used on the secured staff WLAN are usually procured and deployed by (or with the knowledge and cooperation of) the community college network administration staff on behalf of faculty and other university staff users. Faculty and staff users are typically prohibited from bringing their own personal PDAs, laptops, or voice over WLAN (VoWLAN) phones to use on the secured staff WLAN. This allows, for example, a baseline level of authentication and encryption to be deployed for the secured staff WLAN without concern for whether or not the devices using the secured staff WLAN can support this level of authentication and encryption.

The characteristics of this WLAN include the following:

- Wi-Fi Protected Access 2 (WPA2) encryption with 802.1x/EAP authentication, and Cisco Centralized Key Management (Cisco CKM, also referred to as CCKM) for enhanced roaming.
- Most modern WLAN client devices being produced today support this level of authentication and encryption. The addition of Cisco CKM in this case provides for faster roaming by enabling Cisco CKM-equipped clients to securely roam from one access point to another without the need to re-authenticate after the roam completes.
- Broadcast SSID enabled. Enabling this helps to avoid potential connectivity difficulties with some clients. There is no real disadvantage to enabling broadcast SSID.
- QoS profile setting of *silver* (best effort delivery).

Note For more details on WLAN QoS, see the references contained at the end of [Quality-of-Service, page -14](#).

- Wi-Fi Multimedia (WMM) policy of allowed. This allows devices and applications that can support 802.1e enhanced QoS prioritization to do so. Enabling the use of WMM in this way is also in compliance with the 802.11n.
- Mandatory IP address assignment via DHCP. Eliminating the configuration of static IP addresses helps to mitigate the risk of IP address duplication.
- Radio policy set to allow clients to use either 2.4 GHz or 5 GHz to access this WLAN. This allows clients that can take advantage of benefits of 5 GHz operation (such as increased capacity and reduced interference) to do so.

Note The 802.11b and 802.11g physical layers (PHYs) are applied in the unlicensed 2.4 GHz industrial, scientific, and medical (ISM) frequency band, whereas the 802.11a PHY is applied in the unlicensed 5 GHz ISM band. “Dual-band” 802.11a/b/g clients are capable of operating in either 2.4 or 5 GHz frequency bands because they are capable of using any of the three PHYs. Selection between PHYs is typically achieved via software configuration.

Clients using the very high speed 802.11n PHY may be designed to operate in a single band, or they may be 802.11n “dual-band” clients. Unlike the 802.11b, 802.11g, and 802.11a PHYs, simply stating that a client is 802.11n does not precisely indicate what frequency bands the client is capable of operating within.

For more information about the 802.11n PHY and its application to the 2.4 and 5 GHz frequency bands, see the following URL:
http://www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns348/ns767/wHITE_PAPER_80211n_design_and_deployment_guidelines.html.

- A *secured VoWLAN* network that is optimized for VoWLAN usage by college faculty, staff, and administration using managed client devices.

As was the case with the secured staff WLAN, this WLAN is designed to provide secure access and good performance when used with VoWLAN devices (such as the Cisco Unified Wireless IP Phone 7925G) that are usually procured, deployed, and managed by (or with the knowledge and cooperation of) the community college network administration staff. Such procurement is usually conducted on behalf of faculty and other university staff users. To assure proper security and promote effective device management, faculty and staff users are typically prohibited from

bringing their own personal VoWLAN phones and using them on this WLAN. This allows, for example, a baseline level of authentication and encryption to be deployed for this WLAN with the knowledge that the devices using this WLAN can support that level of security. The key differences between this WLAN and the secured staff WLAN include the following:

- The security policy on this WLAN is WPA with Cisco CKM, which is recommended as a best practice for the Cisco 7921G and 7925G VoWLAN phones.
- WLAN controller QoS profile setting of *platinum*, which assigns the highest prioritization to voice traffic.
- WMM policy is *required* (this precludes the use of clients that do not support WMM).
- Load-based Call Admission Control (CAC) should be specified for this WLAN. This prevents VoWLAN calls from being added to an access point that is unable to accept them without compromising call quality.
- The radio policy should be set to allow clients to access only this WLAN using 5 GHz. This helps to ensure that all secured voice devices take full advantage of the robust call capacity and reduced co-channel interference characteristics associated with 5 GHz.

For further information on best practices for voice applications, see the *Voice over Wireless LAN 4.1 Design Guide* at the following URL:
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book.html>.

- A *student WLAN* that uses web authentication for wireless access to the network using unmanaged and privately owned clients such as laptops, PDAs, iPod Touch, iPhones, and so on.

This method of access is normally simple enough for all WLAN users and all platforms, regardless of manufacturer or model. A key challenge in managing wireless access for any large population of users possessing the freedom to choose their wireless clients is how to provide ubiquitous access while still providing an acceptable level of security. Because the ratio of students to network administrative staff is so heavily skewed in favor of the number of students, any student access WLAN solution should require virtually “zero touch” from campus community college network staff, while allowing the vast majority of devices on the marketplace to successfully connect to the network. Characteristics of the student WLAN include the following:

- 802.1x /EAP authentication is not used. For simplicity of configuration across all devices, encryption is not configured on the student WLAN. Transport-level or application-layer encryption may be used if deemed applicable.
- To provide access control and an audit trail, the student access WLAN authenticates the user via a web portal (“web authentication”) where all network access, apart from DHCP and Domain Name Service (DNS), is blocked until the user enters a correct username and password into an authentication web page.
- The student WLAN client device user is re-directed to the web authentication web page whenever the client attempts to open any web page before successful web authentication. This authentication web page can be provided either by an internal WLAN controller web server or the NAC appliance in the Cisco Community College reference design. Usernames and passwords for authentication can reside on a RADIUS AAA server (such as Cisco ACS),

- Broadcast SSID is enabled.
- QoS profile setting of *silver* (best effort delivery).
- WMM policy is set to *allowed*.
- Radio policy should be set such that client access is allowed using either 2.4 GHz or 5 GHz.

- A *guest access WLAN* that uses web authentication for guest users of the campus network.

Traffic to and from this guest access WLAN is tunneled to the DMZ transparently, with no visibility by, or interaction with, other traffic in the enterprise. The Cisco Community College reference design uses the Cisco Unified Wireless Network to provide a flexible, easy-to-implement method for deploying wireless guest access by using Ethernet in IP (RFC3378) within the Cisco Community College reference design. Ethernet in IP is used to create a tunnel across a Layer 3 topology between two WLAN controller endpoints (known as the *foreign* and *anchor* controllers). The foreign controller is the controller resident in the respective campus services block described earlier, whereas the anchor controller is resident within the network DMZ. The benefit of this approach is that no additional protocols or segmentation techniques must be implemented to isolate guest traffic travelling within the tunnel from all other enterprise traffic.

See [Guest Access, page -14](#) for further information regarding considerations surrounding the products and techniques used to provide guest access when designing for mobility in the Cisco Community College reference design.

For technical information on Guest Access best practices in wireless networks, see the Guest Access section in the *Enterprise Mobility 4.1 Design Guide* at the following URL:
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch10/GuAc.html>.

Similar to the requirements stated earlier for the student access WLAN, the guest access WLAN must also be designed to accommodate campus guests (such as alumni, vendors, contractors, prospective students, parents, and so on) as well as the wide variety of WLAN guest clients they may bring onto the campus. Although their numbers will likely be much less compared to that of students, the WLAN clients brought onto campus by guest users are typically not managed or directly supported by community college campus network administrative staff. Because of the lack of control over the type of device used, mandating the use of 802.1x authentication and WPA or WPA2 encryption is usually not practical for guest access.

Characteristics of the guest access WLAN include the following:

- The guest access WLAN uses web authentication in a fashion similar to what was described in the student access WLAN, in order to provide access control and an audit trail.
- The guest access WLAN user is re-directed to a web authentication web page whenever the user attempts to open any web page before successful authentication via the web portal. This authentication web page is provided by an internal WLAN controller web server in the Cisco Community College reference design. However, there is an option of using a non-controller-based web authentication server, such as the Cisco NAC Appliance. Usernames and passwords for authentication can reside on a RADIUS AAA server (Cisco ACS).
- Broadcast SSID is enabled.

- The guest access WLAN uses a QoS profile setting of *bronze* (less than best effort).
- WMM policy is set to *allowed*.
- Radio policy should be set such that client access is allowed to use either 2.4 GHz or 5 GHz.

Additional information about the definition of controller WLANs and SSIDs can be found in the *Enterprise Mobility 4.1 Design Guide* at the following URL:

<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/emob41dg-wrapper.html>.


WLAN Controller Mobility Groups

A *mobility group* is a group of WLAN controllers that behave as a single virtual WLAN controller, sharing essential end client, access point, and RF information. A given WLAN controller is able to make decisions based on data received from other members of the mobility group, rather than relying solely on the information learned from its own directly connected access points and clients. The WLAN controllers in a mobility group form a mesh of authenticated tunnels between themselves, affording any member controller the ability to efficiently communicate with any other member controller within the group.

Mobility groups are used to help facilitate seamless client roaming between access points that are joined to different WLAN controllers. The primary purpose of a mobility group is to create a virtual WLAN domain (across multiple WLAN controllers) to provide a comprehensive view of a wireless coverage area. Typically, two WLAN controllers should be placed in the same mobility group when an inter-controller roam is possible between access points. If the possibility of a roaming event does not exist, it may not make sense to put the WLAN controllers in the same mobility group.

For example, consider the scenario illustrated in [Figure 8](#). Here we see a large and a medium building located on the same campus, in relatively close proximity to one another, with a small building located on a remote campus some distance away. Assume for the purposes of this example that the access points of each building are joined to a different WLAN controller, with the controllers servicing the large and medium building located within the main campus service block, and the WLAN controller servicing the smaller building located on the remote campus. The circular and oval patterns surrounding each building are intended to represent a very simplistic view of hypothetical outdoor RF coverage.


Figure 8 Campus Roaming



[Figure 8](#) shows that there is overlapping coverage between the large and medium buildings, but not between the small building and any other building. This is because users must leave the main campus and traverse through a part of the town to get to the smaller remote campus, and vice versa. Because roaming is clearly possible between the medium and large building, but not between the small building and any other building, only the WLAN controllers servicing the medium and large building are required to be in the same mobility group. The WLAN controller servicing the small building may be configured to be a member of the same mobility group, but it is not mandatory in this case. In applying the concept of mobility groups to the Cisco Community College reference design, consider the following:

- Within a community college or community college system comprised of one or more campuses, it is assumed that intra-campus roaming is possible between all buildings resident on the same campus. This may not actually be the case in all campuses, as some may have buildings co-located on the same campus where areas of non-coverage exist between them. However, assuming that intra-campus roaming is possible between all buildings allows us to make a design assumption that is generally applicable to both situations. Thus, in our Community College reference design, all WLAN controllers servicing access points deployed on the same campus are placed within the same mobility group.
- It is also assumed that in the vast majority of cases, remote campuses are sufficiently distant from the main campus (as well as from one another) to render inter-campus roaming impractical. Allowing for the rare exception that two campuses may be adjacent or otherwise overlap one another, for the most part it is assumed that roaming between buildings located on different campuses is very unlikely.

[Figure 9](#) provides a high-level illustration of how mobility group assignment can be handled in the Community College reference design. Note that *MG* refers to the mobility group name assigned for the campus.

Figure 9 Community College Mobility Groups

mobility groups as indicated in [Figure 9](#), design simplicity is maintained. Given the large capacity of the Cisco 5508 Wireless Controller, the limitation on the maximum number of controllers per mobility group is not a significant tradeoff.

Additional information about WLAN controller mobility groups, including best practice information, can be found in the *Enterprise Mobility 4.1 Design Guide* at the following URL: http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch2_Arch.html#wp1028143.

WLAN Controller Access Point Groups

Typically, each WLAN defined on the controller is mapped to a single dynamic interface (as shown earlier for the secure staff, VoWLAN, and student access WLANs). Consider the case however, where the Cisco 5508 Wireless Controller is deployed and licensed for 250 access points. Assume also that there are 10 users associated to each access point, using the same WLAN and SSID. This would result in 2500 users sharing the single VLAN to which the WLAN is mapped. A potential issue with this approach is that, depending on the particular overall network design, the use of subnets large enough to support 2500 users may not be possible.


To address this issue, the WLAN can be divided into multiple segments using the AP grouping capability of the WLAN controller. AP grouping allows a single WLAN to be supported across multiple dynamic VLAN interfaces on the controller. This is done by assigning a group of access points to an access point group at the WLAN controller, and then mapping the group to a specific dynamic interface. In this way, access points can be grouped logically, such as by building or set of buildings. [Figure 10](#) shows the use of AP grouping based on site-specific VLANs.

The following are some of the key design considerations concerning mobility groups:

- The controllers present at each campus are defined as members of a mobility group unique to that campus. Each controller in the same mobility group is defined as a peer in the mobility list of all controllers for that mobility group.
- If inter-campus roaming between two campuses is possible, the controllers at both campuses should be assigned into the same mobility group and defined as peers in the mobility list of all controllers for that mobility group.
- Because of high-speed WAN/MAN connectivity between campuses, access point failover to a remote backup controller resident at the main campus becomes feasible. To support this, access points can be configured to failover to a WLAN controller outside of their mobility group. This is discussed further in [Controller Redundancy, page -20](#) and [AP Controller Failover, page -21](#).
- A single mobility group can contain a maximum of 72 WLAN controllers. The number of access points supported in a mobility group is bound by the number of controllers and the access point capacity of each controller. Thus, for the Cisco 5508 Wireless Controller, a mobility group can have up to 72 times 250, or 18,000 access points.

The advantage of this approach to mobility group use is clarity and simplicity in deployment and administration. This is a key point when keeping in mind that the typical community college has a limited network administrative staff that is usually resource-constrained and very busy. By dividing the community college system into

Figure 10 Access Point (AP) Groups



As shown in Figure 10, three dynamic interfaces are configured, each mapping to a site-specific VLAN: VLANs 61, 62, and 63. Each site-specific VLAN is mapped to a group of access points that uses the same WLAN/SSID (AP groups one, two, and three). Thus, a faculty member associating to the WLAN using an access point that is part of AP group one is assigned an IP address from the VLAN 61 IP subnet. Likewise, a faculty member associating to the WLAN using an access point that is part of AP group two is assigned an IP address from the VLAN 62 IP subnet, and so on. Roaming between the site-specific VLANs is then handled internally by the WLAN controller as a Layer 3 roaming event. As such, the WLAN client maintains its original IP address.

Cisco 5508 Wireless Controllers can contain up to 192 access point group definitions, with up to 16 WLANs defined in each group. Each access point advertises only the enabled WLANs that belong to its access point group. Access points do not advertise disabled WLANs that are contained within its access point group, or WLANs belonging to another access point group.

In implementations of the Cisco Community College reference design where addressing limitations are present, the use of access point grouping to allow a single WLAN to be supported across multiple dynamic VLAN interfaces on the controller can be extremely beneficial.

WLAN Controller RF Groups

The strategy behind how *RF groups*, otherwise known as *RF domains*, are deployed within the Cisco Community College reference design represents another important deployment consideration that can affect overall accessibility. An RF group is a cluster of WLAN controllers that collectively coordinate and calculate their dynamic radio resource management (RRM) settings. Grouping WLAN controllers into RF groups in this way allows the dynamic RRM algorithms used by the Cisco Unified Wireless Network to scale beyond a single WLAN controller. In this way, the benefits of Cisco RRM for a given RF group can be extended between floors, buildings, and even across campuses.

Note Complete information regarding Cisco Radio Resource Management can be found in the *Cisco Radio Resource Management under Unified Wireless Networks* at the following URL:
http://www.cisco.com/en/US/tech/tk722/tk809/technologies_tech_note09186a008072c759.shtml.

If there is any possibility that an access point joined to one WLAN controller may receive RF transmissions from an access point joined to a different WLAN controller, the implementation of system-wide RRM is recommended, to include both controllers and their access points. In this way, RRM can be used to optimize configuration settings to avoid 802.11 interference and contention as much as possible. In this case, both WLAN controllers should be configured with the same RF group name.

In general, Cisco prefers simplicity in the configuration of RF groups within the mobility design. Thus, all WLAN controllers in the Community College reference design are configured with the same RF group name. Although it is true that geographically disparate WLAN controllers have very little chance of experiencing RF interaction, and thus need not be contained in the same RF domain, for most community college deployments there is no disadvantage to doing so. An exception to this would be in extremely large deployments, as the maximum number of controllers that can be defined in a single mobility group is twenty. A clear advantage to this approach is simplicity of configuration and better support of N+1 controller redundancy (see [Controller Redundancy, page -20](#) for further details).

A more detailed discussion as well as best practice recommendations regarding the use of RF groups can be found in the *Enterprise Mobility 4.1 Design Guide* at the following URL:
http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch2_Arch.html#wp1028184.

Access Points

In the Cisco Community College reference design, it is anticipated that each campus building requiring WLAN access will be outfitted with dual-band 802.11n access points providing RF coverage in both the 2.4 and 5 GHz bands. It is generally assumed that campus users will require WLAN access in most building interior areas, plus a 50–75 yard outdoor perimeter area surrounding each building. Of course, it is important to consider that most buildings will almost certainly contain some areas not intended for human entry or occupancy at any time. Similarly, some buildings may possess areas within the aforementioned outdoor perimeter that simply may not be accessible to campus users at any time. During your initial mobility design, these vacant areas may not be identified, so the precise subset of interior and exterior areas requiring WLAN access will likely be better determined during the site survey planning process that is an integral part of any wireless network deployment.

Note For more information on site survey planning, see the *Cisco 802.11n Design and Deployment Guidelines* at the following URL:

http://www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns348/ns767/wite_paper_80211n_design_and_deployment_guidelines.html.

In most community colleges, the vast majority of interior building WLAN access can be provided by the Cisco Aironet 1140 Series 802.11n access point (see [Figure 11](#)), which delivers pervasive wireless connectivity while blending in seamlessly with the aesthetics of most modern campus learning environments.

Figure 11 Cisco Aironet 1140 Series 802.11n Access Point (AIR-LAP1142N)



To deliver the right mix of style and performance, the Cisco Aironet 1140 Series 802.11n access point contains six integrated omni-directional antenna elements that incorporate the use of three hidden discrete elements for each frequency band. Ideal for indoor environments such as classrooms, corridors, libraries, faculty offices, and so on, the Cisco Aironet 1140 Series 802.11n access point has a visually pleasing metal housing covered by a white plastic shell that blends with the most elegant learning environments. The Aironet 1140 series 802.11n access point provides the ability to be powered directly from 802.3af power-over-Ethernet (PoE) while sustaining full-performance 802.11n connections on both of its radios simultaneously. In the Cisco Community College reference design, the model of the Cisco 1140 Series 802.11n access point recommended for most interior campus building locations is the AIR-LAP1142N.

Note Complete information (including country-specific ordering information) regarding the Cisco Aironet 1140 series 802.11n Access Point can be found at the following URL: <http://www.cisco.com/en/US/products/ps10092/index.html>.

Although the Cisco Aironet 1140 Series 802.11n access point is capable of servicing the bulk of all community college interior wireless access needs, there are some tradeoffs to consider in specialized situations. For example, in situations where the results of pre-site survey planning indicate that the use of external antennas are required to best meet specific RF coverage requirements, an access point providing external antenna connectors will be necessary. This can be a situation where a focused directional antenna pattern is required, or simply one where aesthetic requirements demand that the access point be completely hidden, with only a small antenna footprint exposed to public view. In other cases, perhaps one or more access points will need to be deployed in laboratory environments where the anticipated operating temperature extremes are not within common norms. Here, extended operating temperature tolerances beyond that of the Cisco Aironet 1140 Series 802.11n access point may be required.

To assist in addressing these and other rare but still significant deployment challenges that may be encountered on the community college campus, the Cisco Aironet 1250 Series 802.11n access point is recommended (see [Figure 12](#)).

Figure 12 Cisco Aironet 1250 Series 802.11n Access Point (AIR-LAP1252AG)



Designed with a next-generation ruggedized modular form factor, the Cisco Aironet 1250 Series 802.11n access point is intended for no-compromise performance in combination with the inherent expandability and customizability required to address challenging deployment situations. With robust modularized construction and six RP-TNC antenna jacks that allow antennas to be positioned independently of the access point itself, the Cisco Aironet 1250 Series 802.11n access point can be used to address situations requiring focused directional coverage patterns, extended operating temperature capabilities or minimal-footprint installations where it is highly preferable that the access point chassis is totally hidden from view. In the Cisco Community College reference design, the AIR-LAP1252AG model of the Cisco 1250 Series of access points is recommended for those and other types of demanding deployments.

Note To help discourage theft and vandalism, both the Cisco 1140 as well as 1250 Series 802.11n access points are manufactured with a security slot machined into the access point casing. You can secure either model access point by installing a standard security cable (such as the Kensington Notebook MicroSaver, model number 64068) into the access point security cable slot.

Complete information regarding the Cisco Aironet 1250 series 802.11n access point can be found at the following URL:

<http://www.cisco.com/en/US/products/ps8382/index.html>. Additional information concerning the antenna options available for the Cisco Aironet 1250 Series 802.11n access point can be found at the following URL:

http://www.cisco.com/en/US/prod/collateral/wireless/ps7183/ps469/at_a_glance_c45-513837.pdf

Note that Cisco Aironet 1140 Series 802.11n access points can power both 802.11n radios, at full transmit power running two spatial streams with encryption, while drawing only 15.4 watts of power from an 802.3af PoE Catalyst switch. A tradeoff associated with the use of Cisco Aironet 1250 Series 802.11n access points is that the AP-1250 Series requires slightly more power to reach its peak levels of performance, approximately 18.5 to 20 watts of power from a switch capable of providing enhanced-PoE (ePoE). Keep in mind, however, that if the full performance capability of the Cisco Aironet 1250 series access point is not necessary in your particular deployment, or you wish to support only a single RF band (i.e., either 2.4 GHz or 5 GHz) the Cisco Aironet 1250 Series 802.11n access point can also operate with 15.4 watts from a 802.3af PoE Catalyst switch.

To provide the Cisco Aironet 1250 Series 802.11n access point with 20 watts of input power, Cisco recommends the following power options:

- An ePoE Cisco Catalyst switch or switch blade module (such as the 3560-E, 3750-E, 4500E and 6500E Series).

- The use of a mid-span ePoE injectors (Cisco part number AIR-PWRINJ4). This option allows the Cisco Aironet 1250 series 802.11n access point to deliver full 802.11n performance while connected to any Cisco Catalyst switch. Power is injected directly onto the wire by the AIR-PWRINJ4 mid-span injector without reliance on the power output level of the switch itself.

Although its deployment flexibility is unparalleled within the marketplace, in most community college installation cases, the Cisco Aironet 1250 series 802.11n access point is typically only deployed only in those locations where they are necessary to address challenging situations. Other tradeoffs include a higher total cost per access point because of the added cost of external antennas, a larger footprint, and a heavier mounting weight as compared to the Cisco Aironet 1140 series 802.11n access point.

Note For the Cisco Aironet 1250 Series 802.11n access point, Cisco recommends performing your site survey using the same levels of PoE input power as you expect to use in your final deployment. For example, if you plan to deploy Cisco Aironet 1250 Series 802.11n access points with 15.4 watts of PoE, it is recommended for consistency and accuracy that perform your site survey using the same PoE input power levels.

The following design considerations regarding dual-band access points should be kept in mind when designing networks for dense user environments (for example, interior classrooms and lecture halls within community college campus buildings):

- *Use the 5 GHz band whenever possible*

In general, this applies for both 802.11n as well as pre-802.11n wireless clients. The characteristics of 5 GHz operation make it advantageous for most users, and especially 802.11n users, for the following reasons:

- Despite the maturity of 802.11 wireless LAN technology, the installed base of 5 GHz 802.11a clients generally is not nearly as widespread as 2.4 GHz 802.11b and 802.11g clients. A smaller installed base of users translates into less contention with existing clients and better operation at higher throughput rates.
- The number of non-802.11 interferers (such as cordless phones and wireless personal networks) operating in the 5 GHz band is still just a fraction of the number found within the 2.4 GHz band.
- The amount of available bandwidth found in the 5 GHz band is much greater than that of the 2.4 GHz band. In the United States, there are twenty-one 5 GHz non-overlapping channels that can be deployed. This translates into the ability to deploy with density and capacity in mind, and allow background resources such as Cisco RRM to handle channel and power output requirements accordingly.

- *Design and survey for capacity, not just maximum coverage*

It is a natural tendency to try to squeeze the most coverage from each access point deployed, thereby servicing as much of the campus as possible with the lowest total access point investment. When designing networks for high-speed applications, attempting to design for maximum coverage at maximum transmitter output power can be counter-productive, as the maximum coverage footprint is typically attained using lower data rates and degraded signal-to-noise ratios. In addition, such false economies often sacrifice the ability to effectively make use of advanced tools such as Cisco RRM to address anomalies such as “coverage holes” and other deficiencies. Instead, the successful designer should design for capacity and generally aim to have access points installed closer together at lower power output settings. This approach allows for access point transmitter power to be dynamically managed via

Cisco RRM. It also allows the practical use of higher data rates, provides RRM with the necessary transmission power “headroom” to allow for the ability to compensate for environmental changes, and facilitates the use of advanced capabilities such as location-based context-aware services.

- *Mount access points or antennas on the ceiling when possible*

Cisco Aironet AP-1140 Series 802.11n access points should be mounted on ceilings only. Ceiling mounting is recommended in general for the types of indoor environments found within community colleges, especially for voice applications. In the majority of carpeted indoor environments, ceiling-mounted antennas typically have better signal paths to handheld phones, taking into consideration signal loss because of attenuation of the human head and other obstacles.

Ceiling mounting locations are usually readily available, and more importantly, they place the radiating portion of the antenna in open space, which usually allows for the most efficient signal propagation and reception. Cisco Aironet 1250 Series 802.11n access points can be mounted as deemed necessary during pre-site survey planning or during the actual site survey process. However, ceiling mounting of Cisco Aironet 1250 Series access point antennas is highly recommended, especially for omni-directional style antennas.

- *Avoid mounting on surfaces that are highly reflective to RF*

Cisco recommends that all antennas be placed one to two wavelengths from surfaces that are highly reflective to RF, such as metal. The separation of one or more wavelengths between the antenna and reflective surfaces allows the access point radio a better opportunity to receive a transmission, and reduces the creation of nulls when the radio transmits. Based on this recommendation, a good general rule of thumb then is to ensure that all access point antennas are mounted at least five to six inches away from any large metal reflective surfaces. Note that although recent technological advances have helped greatly in mitigating problems with reflections, nulls, and multipath, a sensible antenna placement strategy still is very important to ensure a superior deployment.

- *Disable legacy and low speed data rates*

Globally disable any unnecessary low speed 802.11a/b/g data rates. Clients operating at low data rates (for example, 1, 2, and 5.5 Mbps) consume more airtime when compared to clients transmitting the same data payloads at higher data rates such as 36 Mbps and 54 Mbps. Overall system performance in any given access point cell drops significantly when a large percentage of low data rate frames tend to consume available airtime. By designing for capacity and disabling lower data rates, aggregate system capacity can be increased.

Unless you are aware of specific reasons why one of the data rates described below are required in your deployment (such as the presence of clients that can transmit or receive *only* at these rates), the following actions are recommended:

- For 2.4 GHz, disable the 1, 2, 5.5, 6, and 9 Mbps rates.
- For 5 GHz, disable at a minimum the 6 and 9 Mbps rates.

A common question concerning 2.4 GHz is why not disable 802.11b entirely? In other words, why not disable the 1, 2, 5.5, and 11 Mbps 2.4 GHz rates altogether? Although this certainly may offer advantages relating to better performance for 802.11g users, this approach may not be entirely practical, especially on guest access WLANs where a visitor might attempt to gain access using a device with embedded legacy radio technology that may not support 802.11g. Because of this, depending on the mix of clients in the environment, it may be wiser to simply disable only the three

802.11b data rates below 11 Mbps. Only if you are completely confident that the situation just described is entirely not applicable in your environment should you consider completely disabling all 802.11b data rates.

Additional best practice guidelines for access point and antenna deployments can be found in the following reference documents:

- *Enterprise Mobility 4.1 Design Guide*—
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/emob41dg-wrapper.html>
- *Voice Over Wireless LAN 4.1 Design Guide*—
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book.html>

To provide outdoor WLAN access around the immediate perimeter area of each campus building, the Cisco Aironet 1520 Series Lightweight Outdoor Access Point is recommended (see [Figure 13](#)).

Figure 13 Cisco Aironet 1520 Series Lightweight Outdoor Access Point



As part of the Cisco Community College reference design, the Cisco Aironet 1520 Series Lightweight Outdoor Access Point provides an outdoor extension to the campus wireless network, with central management provided through WLAN controllers and the Cisco Wireless Control System. A very rugged enclosure allows for deployment outdoors without the need to purchase additional housings or third-party National Electrical Manufacturers Association (NEMA) enclosures to provide protection from extreme weather. The robust, weatherized housing of the Cisco Aironet 1520 Series Lightweight Outdoor Access Point can be painted to adapt to local codes and aesthetics.

Although the Cisco Aironet 1520 Series Lightweight Outdoor Access Point is part of the outdoor mesh series of Cisco access point products, a full outdoor mesh campus infrastructure is beyond the scope of the Cisco Community College reference design at this time. Rather, in this design Cisco Aironet 1520 Series Lightweight Outdoor Access Points are deployed only as root access points (RAPs), located outdoors on each building in such a manner that a satisfactory outdoor perimeter area is established. The precise location of these outdoor access points, as well as antenna choices, depends on the characteristics associated with the required coverage area and other particulars, and should be determined during pre-site survey planning.

For readers who wish to augment the recommendations made in this design guide and deploy a full campus outdoor mesh configuration, see the *Cisco Aironet 1520, 1130, 1240 Series Wireless Mesh Access Points, Design and Deployment Guide, Release 6.0* at the following URL:

http://www.cisco.com/en/US/docs/wireless/technology/mesh/design/guide/MeshAP_60.html.

In choosing among the various models of Cisco Aironet 1520 Lightweight Outdoor Access Points, readers may also wish to consider whether local campus, municipal, state or other public safety agencies are currently using or otherwise plan to deploy compatible 4.9 GHz public safety equipment (see note below) in emergency response vehicles. If this is the case, it may be wise to plan ahead in conjunction with campus and local public safety agencies to accommodate the use of this licensed band for connectivity from properly equipped first responders and emergency vehicles to your campus WLAN. In the event of a campus emergency, the ability to connect to and monitor in-building events, or access key safety and security applications, can significantly enhance the ability of law enforcement and other agencies to locate and combat threats.

Note In 2003, the U.S. Federal Communications Commission (FCC) allocated 50 MHz of spectrum in the 4.9 GHz band to public safety services. Public safety agencies can use this 4.9 GHz band to implement wireless networks with advanced services for the transmission of mission-critical information. Because of the limited number of transmitters and the requirement for licensing, interference on the 4.9 GHz band tends to be below that of other bands, such as 2.4 GHz and 5 GHz. Communications using the 4.9 GHz public safety band must be related to the protection of life, health, or property. Examples include WLANs for incident scene management, mobile data, video surveillance, VoWLAN, fixed point-to-point, and so on.

Even if 4.9 GHz access is not available on campus, public safety agencies may still be able to access the campus WLAN using standard 2.4 GHz or 5 GHz unlicensed bands. This depends on whether the emergency response vehicles of the agencies in question are equipped to do so, as well as the configuration of their equipment. Keep in mind that when public safety users access campus WLANs using unlicensed 2.4 GHz and 5 GHz frequencies during crisis events, they must also contend for access with other unlicensed users of these frequencies, as well as deal with any interference from other sources located within those bands.

With this in mind, the particular model of outdoor access point recommended for outdoor perimeter building coverage, depending on the inclusion of 4.9 GHz as follows:

- The Cisco Aironet 1524PS (Public Safety) Lightweight Outdoor Access Point includes 4.9 GHz capability and provides flexible and secure outdoor WLAN coverage for both public safety and mobility services. The Cisco Aironet 1524PS Public Safety Lightweight Outdoor Access Point is a multiple-radio access point that complies with the IEEE 802.11a and 802.11b/g standards, as well as 4.9 GHz public safety licensed operation parameters. This access point can support independent data exchanges across all three radios simultaneously. The main tradeoff with the Cisco Aironet 1524PS Public Safety Lightweight Outdoor Access Point is the added purchase and deployment cost. However, in environments where public safety agencies are already equipped with compatible 4.9 GHz clients, the added benefits and advantages afforded by the 1524PS are often considered worthwhile. The model of Cisco Aironet 1524PS Public Safety Lightweight Outdoor Access Point recommended in the Cisco Community College reference design is the AIR-LAP1524PS.
- The Cisco Aironet 1522 Outdoor Lightweight Access Point is a dual-radio, dual-band product that is compliant with IEEE 802.11a (5-GHz) and 802.11b/g standards (2.4-GHz). Designed for demanding environments, the Cisco Aironet 1522 provides high performance device access through improved radio sensitivity and range performance. The tradeoffs of deploying this model are the lack of 4.9 GHz licensed

public safety support in environments where 4.9 GHz is in use among public safety agencies. The model of Cisco Aironet 1522 Lightweight Outdoor Access Point recommended in the Cisco Community College reference design for deployments without 4.9GHz is the AIR-LAP1522AG.

Cisco offers a wide array of antenna options for the entire range of Cisco Aironet 1520 Series Lightweight Outdoor Access Points. Information on these antenna options can be found in the *Cisco Aironet 1520 Series Lightweight Outdoor Access Point Ordering Guide* at the following URL:

http://www.cisco.com/en/US/prod/collateral/wireless/ps5679/ps8368/product_data_sheet0900aecd8066a157.html.

All models of the Cisco Aironet 1520 Series Lightweight Outdoor Access Point can be powered from a multitude of sources, include PoE, direct DC, or direct AC. The entire range of power input options is described in the *Cisco Aironet 1520 Series Lightweight Outdoor Access Point Ordering Guide*.

Note Although the Cisco Aironet 1520 Series Lightweight Outdoor Access Point can be conveniently powered via PoE, a power injector (Cisco AIR-PWRINJ1500-2) specific to this product line must be used. Do not use any other power injector or Ethernet switch PoE capability (including enhanced PoE switches) in an attempt to directly provide PoE to Cisco Aironet 1520 Series Lightweight Outdoor Access Points. The Cisco Aironet 1520 Series Lightweight Outdoor Access Point is approved for use only with the Cisco AIR-PWRINJ1500-2 power injector. Keep in mind that although the Cisco Aironet 1520 Series Lightweight Outdoor Access Point is intended to be installed exposed to outdoor weather elements, the AIR-PWRINJ1500-2 power injector is approved for indoor installation only.

Usability

This section discusses the mobility design considerations pertaining to those aspects of the Cisco Community College reference design that are relevant to overall usability, such as the following:

- Quality-of-service (QoS)
- Guest access
- Traffic and performance

Quality-of-Service

The WLAN controller should be configured to set the 802.1p marking of frames received and forwarded onto the wired VLAN to reflect the QoS policy used on this WLAN. Therefore, if the WLAN controller is connected to a switch that is configured to trust the class-of-service (CoS) and maintain a translation table between CoS and Differentiated Services Code Point (DSCP), the translation between wireless QoS policy and wired network QoS policy occurs automatically.

In the Cisco Community College reference design, WLAN traffic is prioritized based on the QoS profiles (platinum, silver, bronze, and so on) applied to each WLAN. However, this does not change the IP QoS classification (DSCP) of the client traffic carried, which means that client traffic leaving WLAN controllers may need to be reclassified based on network policy.

This may be achieved via one of following approaches:

- Applying policy at each of the switch virtual interfaces (SVIs) connecting the WLAN controller to the wired network

- Learning the QoS policy that has already been applied by the wireless networking components, because this should already be in alignment with the overall network policy

In the Cisco Community College reference design, the plan is to use the latter approach, because it provides both the advantage of initial configuration simplicity as well as ongoing ease of maintenance. This technique requires only that the QoS profiles be maintained on the WLAN controllers themselves, without the need to configure explicit policies on adjacent switches. Switches need to be configured to trust only the QoS of frames forwarded to them by the WLAN controller.

To implement this approach, the WLAN controller should be configured to set the 802.1p marking of packets forwarded onto wired VLANs to reflect the QoS policy used on the specific WLAN from which they were received. Therefore, if the WLAN controller is connected to a switch that is configured to trust CoS and maintain a translation table between CoS and DSCP, the translation between wireless and wired network QoS policy occurs automatically.


For example, assume a packet received originates from a WLAN to which a platinum QoS profile has been assigned. This translates to a DSCP value of EF; therefore, the WLAN controller assigns a CoS value of 5 in the header of the frame that carries this data to the wired switch. Similarly, if the same packet originates from a WLAN assigned a QoS profile of silver, the translated CoS value is 0.

For more information on WLAN QoS, see the following URLs:

- *Voice over Wireless LAN 4.1 Design Guide 4.1*—
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book.html>.
- *Enterprise Mobility 4.1 Design Guide*—
http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/ch5_QoS.html

Guest Access

The Cisco Community College reference design uses the Cisco Unified Wireless LAN Guest Access option to offer a flexible, easy-to-implement method for deploying wireless guest access via Ethernet over IP (EoIP), as described in RFC3378. EoIP tunneling is used between two WLAN controller endpoints in the centralized network design. The benefit of this approach is that there are no additional protocols or segmentation techniques necessary to achieve guest traffic isolation in relation to other internal traffic. [Figure 14](#) shows a high-level view of guest access using this technique with a centralized WLAN controller design.

Figure 14 Guest Access Solution High-Level Overview

As shown in [Figure 14](#), a WLAN controller with a specific purpose is located in the main campus DMZ, where it is referred to as an *anchor controller*. The anchor controller is responsible for terminating EoIP tunnels originating from centralized campus WLAN controllers, and interfacing the traffic from these controllers to a firewall or border router. As described in earlier sections of this document, the centralized campus WLAN controllers are responsible for termination, management, and standard operation of the various WLANs provisioned throughout the enterprise, including one or more guest WLANs. Instead of being switched locally to a corresponding VLAN on the campus controller, guest WLANs are instead transported via the EoIP tunnel to the anchor controller in the DMZ.

When an access point receives information from a WLAN client via the guest access WLAN/SSID, these frames are encapsulated using CAPWAP from the access point to the campus WLAN controller. When received at the WLAN controller, they are encapsulated in EoIP from there to the anchor controller. After reaching the anchor controller, these frames are de-encapsulated and passed to a firewall or border router via the guest VLAN. The use of EoIP and an anchor WLAN controller in the DMZ allows guest user traffic to be transported and forwarded to the Internet transparently, with no visibility by, or interaction with, other traffic in the enterprise.

Because the anchor controller is responsible for termination of guest WLAN traffic and is positioned within the Internet DMZ, firewall rules must be established to limit communication between the anchor controller and only those controllers authorized to establish EoIP tunnels to them. Such rules might include filtering on source or destination controller addresses, UDP port 16666 for inter-WLAN controller communication, and IP protocol ID 97 (Ethernet over IP) for client traffic. Other rules that might be needed include the following:

- TCP 161 and 162 for SNMP
- UDP 69 for TFTP
- TCP 80 or 443 for HTTP, or HTTPS for GUI access
- TCP 23 or 22 for Telnet, or SSH for command-line interface (CLI) access

The following are other important considerations to keep in mind regarding the use of this guest access solution:

- For the best possible performance, Cisco strongly recommends that the anchor controller be dedicated to supporting EoIP guest access tunneling only. In other words, do not use the anchor controller for any other purpose but EoIP guest access tunneling. In particular, in addition to its guest access role, the anchor controller should not be used to control and manage other access points in the enterprise.
- When deploying a Cisco 5508 Wireless Controller as an anchor controller, keep in mind that because the anchor controller is not going to be used to manage access points, it can be licensed to support only a minimal number of access points. For example, a Cisco CT5508-12 (12 access point-licensed capacity) can function quite well as an anchor controller in the Cisco Community College reference design, even in networks where hundreds or thousands of access points may be joined to other campus Cisco 5508 Wireless Controllers.
- Multicast traffic is not supported over guest tunnels, even if multicast is enabled on wireless controllers.
- The mobility group name of the anchor controller should differ from that configured for campus controllers. This is done to keep the anchor controllers logically separate from the mobility groups associated with the general campus wireless deployment.
- The mobility group name for every campus WLAN controller that establishes EoIP tunnels with the anchor controller must be configured as a mobility group member in the anchor controller configuration.

Finally, although the focus for the Cisco Community College reference design is on the pure controller-based guest access solution, note that other, equally functional solutions are available that combine what is discussed in this section with the use of an access control platform external to the WLAN controller. For example, the guest access solution topology described in this section can be integrated with the Cisco NAC Appliance. This might be the case, for example, if the community college has already deployed the Cisco NAC Appliance within their Internet DMZ to support wired guest access services. As shown in [Figure 15](#), the wireless guest access topology remains the same, except that in this scenario, the guest VLAN interface on the anchor controller connects to an inside interface on the NAC Appliance, instead of to a firewall or border router.

Figure 15 Cisco UWN Guest Access with Anchor WLC and NAC Appliance





Figure 15 shows that the NAC Appliance is responsible for redirection, web authentication, and subsequent access to the Internet. The campus and anchor controllers are used only to tunnel guest WLAN traffic across the enterprise into the DMZ, where the NAC appliance is used to actually control guest access. The tradeoff here is the added cost of the external access control solution, versus the benefits it affords in relation to your particular deployment.

Note Additional information concerning the design and deployment of the Cisco Unified Wireless Network guest access solution can be found in the *Enterprise Mobility 4.1 Design Guide* at the following URL: <http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/c10GuAc.html#wp999659>.

The Cisco NAC Guest Access Server is another member of the Cisco Network Admission Control solution family that can further enhance the utility of your design by assisting network administrators in the provisioning of guest access user accounts. The NAC Guest Access Server facilitates the creation of guest accounts for temporary network access by permitting provisioning by authorized personnel in a simple and secure manner. In addition, the whole process is recorded in a single place and stored for later reporting,


including details of the network access activity. Cisco NAC Guest Server integrates with Cisco NAC Appliance through an application programming interface (API), allowing for guest accounts to be controlled via the Guest Server user interface, including creation, editing, suspension, and deletion of accounts. The Cisco NAC Guest Server then controls these accounts on the Cisco NAC Appliance through the API (shown in Figure 16). In addition, the Guest Server receives accounting information from the NAC Appliance to enable full reporting.

Figure 16 NAC Guest Server with NAC Appliance and WLAN Controller



Cisco NAC Guest Server can also integrate directly with Cisco WLAN controllers through the RADIUS protocol, allowing for guest accounts to be controlled via the Guest Server user interface, including the creation, editing, and deletion of guest accounts. In this case, the WLAN controller makes use of the NAC Guest Server to authenticate guest users (shown in Figure 17). In addition, the Guest Server receives accounting information from the WLAN controller to enable full reporting.

Figure 17 NAC Guest Server with WLAN Controller Alone



Note For more information on the Cisco NAC Guest Server, see the following URL: http://www.cisco.com/en/US/prod/collateral/vpndev/ps5707/ps8418/ps6128/product_data_sheet0900aecd806e98c9.html.

Traffic and Performance

When designing mobility solutions incorporating the tunneling of CAPWAP traffic across campus infrastructure, questions often arise concerning the impact of such tunneling on network performance. In examining the impact of CAPWAP traffic in relation to overall network traffic volume, the following three points should be considered:

- *CAPWAP control traffic volume*—CAPWAP control traffic volume can vary considerably depending on the current activity state of the network. For example, this type of traffic volume usually reaches a zenith during a software upgrade or WLAN controller reboot. In most campuses, however, this degree of sporadic loading is considered negligible, and is of no consequence when considering the merits of a centralized deployment model over other options.
- *Tunneling overhead*—A Layer 3 CAPWAP tunnel adds a relatively negligible amount of overhead to a typical IP packet traversing to and from a WLAN client.

Note A previous examination of the Light Weight Access Point Protocol (LWAPP), the predecessor to CAPWAP and similar in many ways, indicates that this overhead is approximately 44 bytes. With that said, traffic studies have concluded that the average load LWAPP control traffic places on the network is approximately 0.35 Kb/sec. Given that average packets sizes found on large scale network deployments are approximately 300 bytes, this represents an overhead of approximately 15 percent.

Once again, this is generally viewed as resulting in little to no consequence, especially in light of the considerable merits associated with a centralized deployment versus other options.

- *Traffic engineering*—WLAN traffic tunneled to a centralized controller is typically routed from the location of the WLAN controller to its final destination in the network. In the case of the Cisco Community College reference design, established best practices are followed concerning the placement of WLAN controllers within each per-campus centralized services block. With that said, the longer tunnels and traffic flows associated with a centralized deployment model can be mitigated by positioning the WLAN controllers in that part of the network where a large portion of the client traffic is already destined. In the Cisco Community College reference design, client-to-host/server traffic is typically destined for a local campus or main campus data center. This being the case, the overhead associated with any inefficiencies introduced because of centralized placement is not seen as adding significant delay or overhead.

Manageability

As mentioned earlier, each WLAN controller in the Cisco Community College reference design provides both a CLI as well as a graphical web user interface, which are primarily used for controller configuration and management. These user interfaces provide ready access to the network administrator. However, for a full-featured, centralized complete lifecycle mobility management solution that enables community college network administrators to successfully plan, configure, deploy, monitor, troubleshoot, and report on indoor and outdoor wireless networks, the use of the Cisco Wireless Control System (WCS) is highly recommended (see [Figure 18](#)).

Figure 18 Cisco Wireless Control System




The Cisco Wireless Control System allows very effective management of wireless networks supporting high-performance applications and mission-critical solutions. Effective management of these networks helps to simplify college network operation and improve the productivity of administrators, staff, and faculty. The comprehensive Cisco WCS platform scales to meet the needs of small, midsize, and large-scale WLANs across local and remote campuses. Cisco WCS gives college network administrators immediate access to the tools they need when they need them, wherever they may be located within the community college.

Operational costs are significantly reduced through a simplified and intuitive GUI, with built-in tools delivering improved efficiency and helping to reduce training costs, even as the campus network grows incrementally larger. Cisco WCS lowers operational costs by addressing the whole range of mobility management requirements (radio frequency, access points, controllers, mobility services, and so on) using a single unified management platform deployed in a centralized location, and with minimal impact on staffing requirements.

Cisco WCS can scale to manage hundreds of Cisco WLAN controllers, which in turn can manage thousands of Cisco Aironet access points. For installations where network management capabilities are considered mission-critical, WCS also supports a software-based high availability option that provides failover from a primary (active) WCS server to a secondary (standby). Adding mobility services such as context-aware software and adaptive wireless intrusion prevention systems (wIPS) is simplified through Cisco WCS integration with the Cisco Mobility Services Engine (MSE).

Note A detailed description of each management feature and benefit available in the Cisco Wireless Control System is beyond the scope of this chapter, but the information can be found at the following URL:
http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/product_data_sheet0900aecd802570d0.html.

In the Cisco Community College reference design, a centralized WCS management server located in the data center block within the main campus is used. The data center block was initially shown in [Figure 3](#). [Figure 19](#) provides greater detail and magnification.

Figure 19 WCS Within the Data Center Block

The current upper limit for scaling WCS on a high-end server is up to 3000 Cisco Aironet CAPWAP-based access points, and up to 750 Cisco WLAN controllers. As such, most implementations of the Cisco Community College reference design are well served by a mobility design using a WCS management server located on the main campus.

Note For further information on WCS hardware platforms and requirements, see the following URL:

http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0_wst.html#wp1061082.

The planning, configuration, deployment, monitoring, reporting, auditing, and troubleshooting made available by WCS are accessible to any authorized community college network administrator via standard secured web browser access.

Generally speaking, it is anticipated that access to WCS will be restricted to network administrators and their staff located at the main and remote campuses, as well as faculty administrators and staff. However, these groups will not all have equivalent resource and functionality access. It is anticipated that resource access will be limited further, based on administrative level and assigned campus or campuses. With few exceptions, it is not anticipated that most students will be required nor authorized to use the majority of services offered by WCS.

In this design, the ability to query and manage campus mobility resources is regulated using the virtual domain feature of WCS, in conjunction with the appropriate assignment of WCS user rights. Thus, although key members of the main campus central network administration staff may possess the authority to manage any and all mobility resources located on any campus throughout the college system, remote campus administrators may be limited by the following:

- *Campus resource management visibility policy*—This is performed by assigning the network mobility infrastructure components associated with each campus to a WCS virtual domain, and assigning the virtual domains to appropriate network administrators. Key members of the central administrative staff are assigned to the WCS root domain, granting them overall authority to view and configure all mobility infrastructure resources, on any campus, via their WCS management consoles. However, personnel responsible for local campus network administration are restricted to the discrete mobility infrastructure components associated with the virtual domain representing their local campus. These infrastructure components include WLAN controllers, access points, configuration templates, WCS events, reports, alarms, WLAN clients, and so on.
- *Campus resource management access policy*—Although the visibility of a resource is determined by WCS virtual domain assignment, the subset of acceptable actions that are allowed against any visible resources are further regulated by the

assignment of appropriate WCS user and group rights, which allow policies to be applied that further limit what actions each may be allowed against any visible resources.

Via the WCS GUI interface, virtual domains (as well as WCS user rights) can be assigned at the WCS server or using an external security manager such as Cisco Secure ACS.

Note Further information regarding how WCS virtual domains may be used to limit individual campus network administrator access to segments of the mobility network outside of their scope of responsibility, while still providing for overall “root” administrator control of the entire wireless network, may be found at the following URL:

http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/brochure_c02-474335.html.

Guest access credentials can be created and managed centrally using the Cisco WCS. A network administrator can create a limited privilege account within WCS that permits “lobby ambassador” access for the purpose of creating guest credentials. With such an account, the only function a lobby ambassador is permitted is to create and assign guest user credentials to controllers that have web-policy configured WLANs. In the rare event that a centralized WCS management system is not available because of a server failure, a network administrator can establish a local administrator account on the anchor WLAN controller, with lobby ambassador privileges, as a backup means of managing the guest access solution.

The use of a centralized WCS management server in the Cisco Community College reference design provides key advantages such as reduced initial deployment cost and ease of maintaining server resources in a centralized location, coupled with good performance across modern high-speed LANs and WANs. Of course, as with any design choice, certain tradeoffs exist, such as the following:

- *WCS server failure*

In the Cisco Community College reference design, the centralized mobility network management services provided by WCS are not regarded as being mission-critical for the majority of community college deployments. Thus, in the rare event of a WCS server failure, and given the cost constraints of most community college environments, it is assumed that direct WLAN controller management workarounds (such as that described earlier for guest access management) are an acceptable cost compromise. Any downtime realized because of a WCS server failure, although undoubtedly very inconvenient, would in most cases not be viewed as entirely catastrophic. This being the case, the Cisco Community College reference design does not at this time provide for the added cost of a secondary WCS management server in an N+1 software-based high-availability arrangement. However, deployments where WCS management services are critical to the mission of the community college should instead consider modifying the design to include the services of a secondary WCS management platform configured for N+1 software-based high-availability.

Note For more information on WCS high availability configurations, see the following URL:

http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0ad_min.html#wp1132580.

- *Unrecoverable WAN failure*

A catastrophic, unrecoverable WAN failure can interrupt management traffic between WCS and the WLAN controllers that are located on remote campuses. One way to protect against this is to distribute the WCS management server function out further into the network, and centralize WCS management on a per-campus basis. However, this increases the cost of WCS deployment significantly, requiring one WCS management server per campus, and preferably a Cisco WCS Navigator management aggregation platform located at the main campus site. Because it is believed that the centralized mobility network management services provided by WCS are not regarded as mission-critical to the majority of community colleges, these decentralized management options are not included in the Cisco Community College reference design at this time. Instead, it is assumed that in this type of a rare occurrence, the aforementioned ability to minimally manage WLAN controllers directly will suffice, should any network management intervention be required in such circumstances.

Note For more information on WCS Navigator, see the following URL:
<http://www.cisco.com/en/US/products/ps7305/index.html>.

Reliability


This section discusses the mobility design considerations pertaining to those aspects of the Cisco Community College reference design relevant to overall reliability, and includes the following:

- Controller link aggregation
- Controller redundancy
- AP controller failover

Controller Link Aggregation

An important capability used to enhance the reliability of WLAN controller interconnection to the wired network is *link aggregation (LAG)*. As mentioned earlier, LAG is a partial implementation of the 802.3ad port aggregation standard. It bundles all the controller distribution system ports into a single 802.3ad port channel, thereby reducing the number of IP addresses needed to make use of all controller wired ports. When LAG is enabled, the system dynamically manages port redundancy and load balances access points across each port, without interaction from the network administrator. With the Cisco 5508 Wireless Controller and the release 6.0 software used in the Cisco Community College reference design, all eight ports can be bundled together into a single Gigabit EtherChannel interface. LAG is effective in distributing access point traffic across all controller ports, as shown in [Figure 20](#). This can be especially important with high capacity controllers licensed for many access points, such as the Cisco CT5508-250.

Figure 20 LAG in the Cisco 5508 WLC



LAG simplifies controller configuration and improves the overall solution reliability. If any of the controller ports fail, traffic is automatically migrated to one of the remaining ports. As long as at least one controller port is functioning, the system continues to operate, access points remain connected to the network, and wireless clients continue to send and receive data.

The Gigabit Ethernet connections comprising the LAG (up to eight on the Cisco 5508 Wireless Controller) should be distributed among different modular line cards or switch stack members in the services block to the greatest degree possible. This is done to ensure that the failure of a single line card or switch stack member does not result in total failure of the WLAN controller interconnection to the campus network.

For example, if there are four switch stack members in the services block and LAG is configured using all eight WLAN controller interfaces, the Gigabit Ethernet links from the services switch block to the WLAN controller should be distributed two per services block switch stack member. In this way, if any switch stack member fails, six other Gigabit Ethernet links to the WLAN controller remain ready, active, and available to pass data.

The switch features required to implement this connectivity between the WLAN controller and the services block are the same switch features that are otherwise generally used for EtherChannel connectivity between switches.

When using a Cisco 5508 Wireless Controller with link aggregation enabled, it is important to keep the following considerations in mind:

- When the port channel is configured as "on" at both ends of the link, it does not matter if the Cisco Catalyst switch is configured for either Link Aggregation Control Protocol (LACP) or Cisco proprietary Port Aggregation Protocol (PAgP), because no channel negotiation occurs between the controller and the switch.

The recommended load balancing method for Cisco Catalyst switches is by use of the CLI command **src-dest-ip**.

- You cannot configure the controller ports into separate link aggregation groups. Only one link aggregation group is supported per controller. Therefore, you can connect a controller in link aggregation mode to only one neighbor switch device (note that this can be a switch stack with multiple member switches).
- When you enable link aggregation or make any changes to the link aggregation configuration, you must immediately reboot the controller.
- When you enable link aggregation, only one AP manager interface is needed because only one logical port is needed. The in-band management interface of the Cisco 5508 Wireless Controller can also serve as the AP manager interface.
- When you enable link aggregation, all Cisco 5508 Wireless Controller distribution ports participate in link aggregation by default. Therefore, you must configure link aggregation for all the connected ports in the neighbor switch that have been outfitted with small form-factor plug-in (SFP) modules.
- When you enable link aggregation, only one functional physical distribution port is needed for the controller to pass client traffic. Although Cisco 5508 Wireless Controllers have no restrictions on the number of access points per port, Cisco recommends that if more than 100 access points are connected to the controller, make sure that at least two or more Gigabit Ethernet interfaces are used to connect the controller to the services block.
- As mentioned previously, there are eight SFP interfaces on the Cisco 5508 Wireless Controller. These may be fully deployed to take full advantage of multilayer campus design guidelines regarding the oversubscription of access layer uplinks. By doing so, it is relatively straightforward to design a solution that delivers access layer uplinks from the WLAN controller with an oversubscription rate of between 8:1 and 20:1 (Note that these oversubscription rates are not unique to wireless products and are equivalent with what is typically seen in wired networks as well.)

Table 1 provides information for the Cisco 5508 Wireless Controller deployed with its maximum complement of 250 access points.

Table 1 Cisco 5508 Wireless Controller Oversubscription Rates

Throughput per AP (Mbps)	Cisco 5508 Wireless Controller Oversubscription Rate (8 Gbps)
25	1:1
50	2:1
100	4:1
150	5:1
200	7:1
250	8:1

Table 1 shows that even if designing for peak 802.11n throughput of 250 Mbps per access point, oversubscription is not expected to exceed campus design guidelines of 8:1 when using all the available controller interfaces with LAG.


Note For more information concerning WLAN controller link aggregation, see *Deploying Cisco 440X Series Wireless LAN Controllers* at the following URL: <http://www.cisco.com/en/US/docs/wireless/technology/controller/deployment/guide/dep.html#wp1062211>.

Controller Redundancy

The ability of the solution to recover from a reasonable degree of component failure is important in ensuring the reliability of any WLAN networking solution. This is especially important when there are many users that may rely on a centralized component, such as a WLAN controller, for access into the network. An easy solution is to have a “hot” standby secondary controller always at the ready for each primary controller in active service (otherwise known as 1:1 controller redundancy). Although this offers the highest degree of protection from any number of failed primary controllers, it is also the most costly approach.

In the Cisco Community College reference design, unforeseen controller failures are avoided using an “N+1” controller redundancy model, in which the redundant WLAN controller is placed in a central location and acts as a backup for multiple active WLAN controllers. Each access point is configured with the name or IP address of its primary WLAN controller, but is also configured with the name or IP address of the redundant controller as its secondary WLAN controller. The N+1 controller redundancy approach is based on the assumption that the probability of more than one primary WLAN controller failure occurring simultaneously is very low. Thus, by allowing one centralized redundant controller to serve as the backup for many primary controllers, high availability controller redundancy can be provided at a much lower cost than in a traditional 1:1 redundancy arrangement. Figure 21 provides a general illustration of the principle of N+1 controller redundancy.

Figure 21 General N+1 WLAN Controller Redundancy



The main tradeoff associated with the N+1 redundancy approach is that the redundant controller may become oversubscribed if multiple primary controllers fail simultaneously. In reality, experience indicates that the probability of multiple controller failures is low, especially at geographically separate site locations. However, when designing an N+1 redundant controller solution, you should assess the risk of multiple controller failures in your environment as well as the potential consequences of an oversubscribed backup.

controller. In situations where there is reluctance to assume even this generally small degree of risk, other controller redundancy approaches are available that can provide increasingly greater degrees of protection, albeit with associated increases in complexity and equipment investment.

Note For more details on controller redundancy, see *Deploying Cisco 440X Series Wireless LAN Controllers* at the following URL:
<http://www.cisco.com/en/US/docs/wireless/technology/controller/deployment/guide/dep.html#wp1060810>.

The configuration of N+1 redundancy in any mobility design depends greatly on the licensed capacity of the controllers used and the number of access points involved. In some cases, configuration is rather straightforward, emulating what is shown in [Figure 21](#) by having the access points of the main campus as well as all remote campuses address a common redundant controller located in the main campus services block. In other cases, there may be sufficient capacity on the primary controllers located on the main campus themselves to accommodate the access point and user load of a single failed controller on any of the remote campuses. This approach requires that main campus controllers be licensed for a greater number of access points than necessary for the support of the main campus alone. Additional licensing of existing controllers is performed in place of providing a dedicated additional controller platform at the main campus for system-wide redundancy. In this case, the available capacity of the primary main campus WLAN controllers allow them to act as the secondary destination for the access associated with the largest remote campus. Thus, in this particular case, the need to deploy hardware at the main campus site explicitly for the purposes of controller redundancy may be avoided.

For example, assume that the main campus shown in [Figure 3](#) contains a total of 250 combined access points across all main campus buildings, and the largest of the remote campuses also contains 250 combined access points across all remote campus buildings. In this case, if the main campus services block is equipped with two Cisco CT5508-250 WLAN controllers (the “-250” signifies that this particular Cisco 5508 Wireless Controller is licensed for 250 access points), the access point load of the main campus alone can be split equally between the two controllers (125 access points on each controller). This leaves ample capacity in the main campus for one of the following scenarios to occur:

- Either of the main campus controllers may fail and allow up to 125 joined access points to migrate (failover) to the other controller in the pair. This results in the remaining functional controller bearing the full load of 250 access points.
- Any remote campus controller may fail and allow its joined access points to migrate (failover) to the main campus controllers. In the case of a failure of the largest remote campus, this results in each of the main campus controllers operating at their full licensed capacity.

Further information regarding WLAN controller redundancy may be found in the following documents:

- *Deploying Cisco 440X Series Wireless LAN Controllers*—
<http://www.cisco.com/en/US/docs/wireless/technology/controller/deployment/guide/dep.html#wp1060810>
- *Enterprise Mobility 4.1 Design Guide*—
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/emob41dg-wrapper.html>

AP Controller Failover

The Cisco Unified Wireless Network provides for multiple failover options that can allow access points to determine which WLAN controller to migrate in the event of a controller failure, based on pre-configured priorities. When an access point goes through its discovery process, it learns about all the WLAN controllers in its mobility group. The access point can prioritize which controller it attempts to join based on its high availability configuration, or choose a WLAN controller based on loading.

In the Cisco Community College reference design, a high-speed WAN/MAN is present between campuses, thus making access point failover to a remote WLAN controller feasible, as described in the previous section. To accomplish this in the Cisco Community College reference design, access points can be configured to failover to a WLAN controller that is outside their mobility group. In this scenario, the remote WLAN controller is not in the mobility group that is learned during the AP discovery process, and the IP address of the remote WLAN controller must be provided in the HA configuration.

For this to be effective, however, a common WLAN SSID naming policy for key WLANs must be implemented to ensure that WLAN clients do not have to be reconfigured in the event of an access point failover to the main campus backup controller.

Best practice considerations regarding to AP controller failover include the following:

- After access points initially discover a WLAN controller, access points should be manually assigned to primary and secondary controllers. By doing this, AP assignment and WLAN redundancy behavior is deterministic.
- A common WLAN SSID naming policy is necessary to ensure that WLAN clients do not have to be reconfigured in the event of an access point failover to a central backup controller. The SSID used to access a particular WLAN throughout the multi-campus community college should be the same, regardless of the controller.
- WLAN controllers have a configurable parameter known as *AP Fallback* that causes access points to return to their primary controllers after a failover event, after the primary controller comes back online. This feature is enabled by default. However, leaving this parameter at the default value can have some unintended consequences. When an access point “falls back” to its primary controller, there is a brief window of time, usually approximately 30 seconds or so, during which service to wireless clients is interrupted because the access points are busy re-joining the primary controller. In addition, if connectivity to the primary WLAN controller becomes unstable for some reason, the access point might “flap” between the primary controller and the backup. For this reason, it is preferable to disable AP Fallback and, in the rare event of a controller failure, move the access points back to the primary controller in a controlled fashion during a scheduled service window.

Note For more information and best practices regarding AP controller failover, see the *Enterprise Mobility 4.1 Design Guide* at the following URL:
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob41dg/emob41dg-wrapper.html>.

Community College Mission Relevancy

This document attempts to present the mobility design considerations that comprise an important part of a successful implementation of the Cisco Community College reference design. The goal is to provide stakeholders with a reference design that assists in solving the complex business challenges that community colleges must face in the 21st century.

This closing section steps back from the technical intricacies of system design to examine how these design considerations relate to the foundation services described in the opening paragraphs of this document.

Safety and Security

The mission of the Cisco Community College reference design in this area is to enhance safety and security on campus by using a design model that proactively protects students, faculty, and staff. Maintaining safe buildings and grounds while keeping the network secure for today's community colleges. The Cisco Community College reference design helps to facilitate and enhance the effectiveness of physical campus security, track assets, protect the network, and prevent unauthorized network access.

The mobility aspects of Cisco Safety and Security Solutions includes the following three solution sets:

- *Campus physical safety and security*—Is the physical campus protected and safe? The Cisco Community College reference design helps enable community colleges to maintain safe buildings and grounds by the following:
 - Supporting the monitoring of unauthorized behavior and delivering alerts about detected events. Real-time monitoring helps campus security staff to prevent, deter, detect, and respond more quickly to incidents.
 - Providing reliable, secure, and high-performance WLAN communications throughout building interiors and outside buildings to students, faculty, administrators, and community college guests. This level of reliable wireless connectivity can be the key to ensuring rapid notification of campus personnel in the event of a safety incident.
 - Real-time tie-in to wired and wireless video surveillance systems as well as portable security devices and third-party campus safety systems, to ensure that unfolding events are detected quickly and monitored by the right personnel in the right location.
 - Offering WLAN connectivity from strategic campus locations to public safety emergency professionals using licensed 4.9 GHz frequencies as well as traditional 2.4 GHz and 5 GHz unlicensed frequencies. During periods of crisis, indoor and outdoor WLANs can provide first responders with vital tactical information about what is happening within the campus. The 4.9 GHz band that is available on the Cisco Aironet 1524PS Public Safety Lightweight Outdoor Access Point provides access via radio frequencies that are reserved by the FCC for public safety usage only.

Other Cisco products and solutions that work in collaboration with the Cisco Community College reference design to enable these and other capabilities include the Cisco Mobility Services Engine (MSE), Cisco Context-Aware Mobility Solution, Cisco Unified Communications, the Cisco Unified Wireless IP Phone 7925G, and Cisco Video Surveillance products.

- *Network and data security*—Is the wireless network secure? The Cisco Community College reference design addresses this issue by the following:
 - Protecting confidential data and transmissions by using the highest level of authentication and encryption applicable to the tasks at hand, helping to ensure that wireless transmissions remain secure and protected.

- Helping to prevent misguided students or malicious intruders from hacking into restricted servers or issuing attacks against the wireless network via the inclusion of the optional Mobility Services Engine with Wireless Intrusion Protection System (wIPS). It also helps quickly locate rogue access points anywhere on campus.
- Providing an economical guest access solution that furnishes safe and secure guest access for campus guests.
- *Context-aware mobility*—Where is an asset located on campus and what is its status? The Cisco Unified Wireless Network, in conjunction with Cisco Context-Aware Mobility solutions, supports the ability to do the following:
 - Capture and integrate into community college application and administrative processes, detailed contextual information about an asset such as its location, movement, status, and state. This solution helps community colleges automatically collect information about mobile assets, analyze it, and use it to reduce errors, improve asset security, prevent delays, improve scalability beyond manual processes, and enhance learning functions.
 - Any asset that is emitting a Wi-Fi signal can be monitored, tracked, and found with this solution. A Wi-Fi signal can be generated from a built-in wireless card or an attached Wi-Fi tag from third-party vendors including AeroScout, WhereNet, and others.
 - Expensive items such as projectors, televisions, portable plants, lab equipment, tools, laptops, or any asset that moves can be easily tracked.
 - Alerts can be issued about the movement of a device in or out of an area. Costs for misplaced items, loss, and theft can be reduced.
 - Faculty and staff can use context-aware mobility in conjunction with third-party applications to automatically send announcements, assignments, room change notifications, campus event updates, and emergency alerts to students as they roam on campus.
 - Security personnel can use this solution to receive silent alerts and notifications about asset movement and rogue devices, track the areas of the campus they have inspected or secured, and quickly learn the location and of emergency-triggered events.
 - Administrators can use this solution to quickly locate students, faculty, or staff anywhere on campus.

Virtual Learning

The traditional scenario of a mass of students filing into a large lecture hall within a large, monolithic campus building is by no means the only such model available to today's modern-day community college student. High performance, secure wireless technologies can enable "virtual classrooms" even in non-traditional settings, such as leased space in shopping malls, retail plazas and even from homes and offices.

School administrators need secure access to tools, records, and resources, as well as ubiquitous access to mobile voice capabilities throughout the campus.

Using the solutions and technologies presented within the Cisco Community College reference design, state-of-the-art instructional sites can be deployed in such non-traditional settings within urban, suburban, and rural venues. These types of facilities can help bring much-needed skills to areas that may not be within convenient reach of conventional community college campuses. For example, a community college location

at a shopping mall may operate as a science, technology, engineering, and math learning center. Such centers may range in size from one or two classroom sites to larger-scale deployments with ten or twelve classrooms, a hundred or more student computers, a science lab, two auditoriums, and even testing, conference, and office space.

Secure Connected Classrooms

Providing connectivity to students while attending class is the foundation of twenty-first century learning. However, it also presents several challenges for community colleges. For example, the density of wireless users in one location can be problematic. Wireless designs must take into consideration the number of users, radio interference, and network utilization.

The Cisco Community College reference design addresses these challenges in a variety of ways, including the following:

- High-performance dual-band access points that provide options to migrate users to 802.11n and better performing bands (5 GHz) that offer increased data rates with less interference.
- Advanced radio resource management algorithms and techniques that can automate the fine-tuning of transmit power and other parameters to best accommodate high-density user populations.
- Comprehensive wireless network management systems (WCS) that can assist in the identification of interference sources and rogue access points, including their location.
- Detailed reporting mechanisms that can enable administrators to better understand the points of congestion in the network and how best to address them.
- A high-performance controller platform, optimized for use with high-performance 802.11n access points, that offers aggregate wired interface bandwidth of up to 8 Gbps.

Operational Efficiencies

Delivering quick and cost-effective broadband access anywhere on campus extends learning beyond the classroom and improves campus operations, collaboration, and productivity. The Cisco Community College reference design supports secure, easy wireless network access to voice, video, and data applications for students, administrators, faculty, staff, and visitors as they roam about the campus.

The operational efficiencies enabled by the Cisco Community College reference design encompass the following solution sets:

- *Pervasive wireless on campus*—Is the WLAN available ubiquitously in all required indoor and outdoor areas? As a key component of the Cisco Community College reference design, the Cisco Unified Wireless Network delivers broadband access quickly and cost-effectively to all the required indoor and outdoor areas in the typical community college. The benefits of this are as follows:
 - When wireless access is available pervasively on campus, users do not need to hunt for wired ports because they can gain access to network resources using their wireless connection.
 - Users can stay connected to their applications as they roam, without having to re-log onto the network while they are in motion.
 - As long as an area is covered by the wireless infrastructure, faculty, students, and guests can work, share resources, collaborate, and communicate.

- With a pervasive wireless network, instruction is no longer limited to the classroom.
- Faculty can teach inside or outside the classroom, accessing the Internet and applications while on the move.
- Access to resources is improved because faculty and administrators do not have to return to their desk to perform online administration tasks, access research information, or check E-mail.
- Student satisfaction is increased and trouble calls are decreased because wireless access is predictable and consistent.
- With a pervasive Cisco WLAN, community colleges can deliver network access to locations where hardwiring is too expensive, too difficult, or implausible. Examples are refurbished buildings, older buildings with environmental concerns such as asbestos remediation, or sites with protected-building restrictions such as historical landmarks.
- Costs for cabling temporary spaces or for providing network access to new faculty or staff can be reduced or eliminated.

In fact, you may find that it is more cost-effective to provide wireless network access pervasively on campus than it is to install individual wired ports over the same geographic area.

- *High-speed wireless access*—Are bandwidth-intensive applications supported on the WLAN? The Cisco Unified Wireless Network facilitates the creation of solutions that accelerate the delivery of bandwidth-intensive applications and provides a better end-user experience.
 - The Cisco high-speed wireless network, based on the 802.11n standard, delivers unprecedented reliability, greater performance, and extended reach for pervasive wireless connectivity. It excels at supporting bandwidth-intensive applications that are used for research, learning, virtual environments, and social networking. This solution also delivers predictable and continuous WLAN coverage for areas with dense wireless usage such as lecture halls, auditoriums, open spaces, and social areas.
 - Community colleges that deploy 802.11n are demonstrating a commitment to technology innovation and leadership. They are building a solid technology foundation to attract new students and remain competitive in the ever-evolving global community college education marketplace.
- *Secure guest access*—Can visitors easily access the network? The Cisco Community College reference design supports secure wireless guest access that cost-effectively simplifies the process of providing temporary Internet access to visitors such as prospective students, alumni, parents, visiting lecturers, and temporary personnel. Wireless guest access eliminates the frustration that visitors experience when they are limited to wired-only ports in small areas on campus. It also eliminates the costs that community colleges might incur from wiring and maintaining wired ports to accommodate visitors. With the Cisco secure guest access solution, community colleges can do the following:
 - Enhance the community college experience for prospective students
 - Provide Internet access to guests attending campus events
 - Easily support network access for conference attendees and guest lecturers

- *Campus automation*—Are managing and tracking campus resources automated? The solutions enabled by the Cisco Community College reference design can help community colleges reduce costs by supporting Wi-Fi-enabled services that automatically manage, track, and maintain campus resources and assets.
 - The wireless network can assist with better management of real estate components to support green initiatives, improve energy efficiency, and create smart buildings.
 - Alarms, bells, and clocks can be wirelessly enabled to reduce the labor costs associated with managing them. Wi-Fi tags can be placed on assets to automatically track their movement and help reduce costs for misplaced items, loss, and theft.
- *Facilities management*—The Cisco Community College reference design includes adaptive power management capabilities that are built into the Cisco Unified Wireless Network through its Cisco Wireless Control System (WCS) management platform and software release 6.0. Cisco WCS adaptive power management allows community colleges to shrink their carbon footprint immediately through measurable reductions in energy usage and operational expenses.

By using Cisco WCS adaptive power management to turn access point radios on or off at scheduled intervals (hour, day, and week), power requirements and operating expenses can be reduced almost immediately. The power savings gained vary based on the Cisco Aironet access point model deployed. Using this feature can help organizations create a sustainable culture and gain momentum for “Green IT” initiatives