



# Unified Access Design Guide

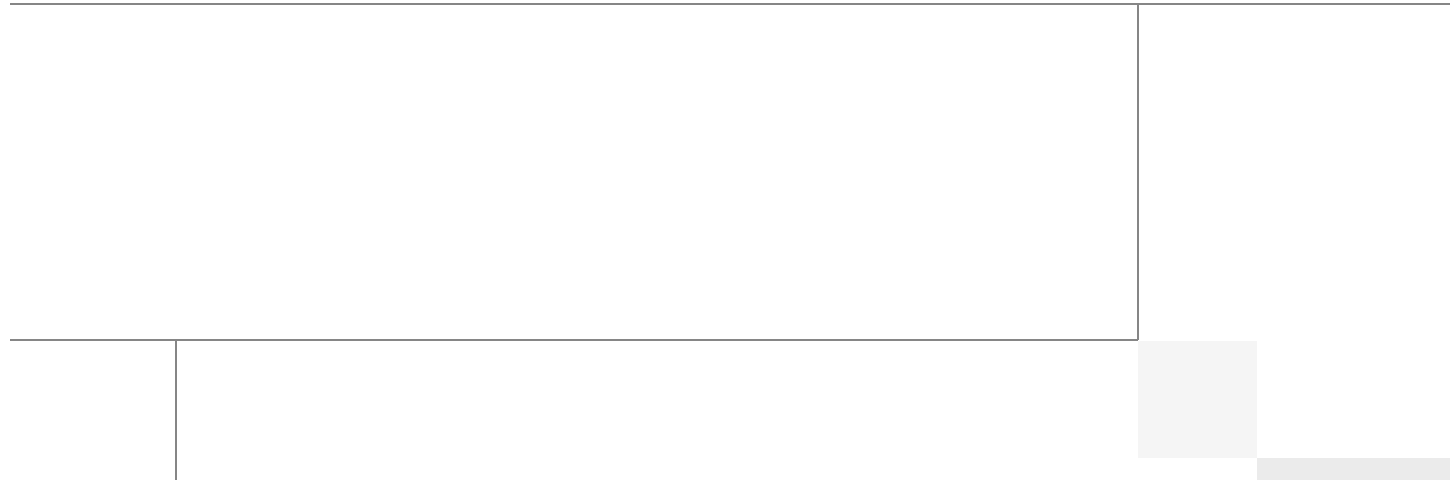
Last Updated: October 18, 2011



Cisco  
Validated  
Design



Building Architectures to Solve Business Problems



## About the Authors



Jay Cedrone

### **Jay Cedrone, Test Engineer, Systems Architecture & Strategy Unit (SASU), Cisco Systems**

Jay Cedrone is a Test Engineer focusing on Campus Architectures in Cisco's Systems Architecture & Strategy Unit. With 10 years experience at Cisco, he has worked on Layer 4 to 7 Content Switching, Cryptographic Protocols, and Global Server Load Balancing. Prior to his role as a Test Engineer within SASU, he worked in the Cisco Safe Harbor testing group. Previously he was a member of the Technical Marketing team for Layer 4 to 7 products.



Steve Gyurindak

### **Steve Gyurindak, Solutions Manager, Systems Architecture & Strategy Unit (SASU), Cisco Systems**

Steve Gyurindak is a Solutions Architect in the Campus group, creating architectures for Cisco's public sector market segment. Previously at Cisco he served as a systems engineer in several sales groups, mostly serving large enterprise accounts and public sector accounts in Georgia and neighboring states. Prior to working at Cisco, he worked as a senior network engineer for the Georgia Court System and Legislature for the State of Georgia. Gyurindak holds many industry certifications including CCNP, CCNA, CCDP, CCDA, MCSE, MCNE, CCIE Routing & Switching certification 9057, and CISSP certification 61046. He holds a bachelor of science degree in network engineering from the State University of New York at Buffalo.



Zeb Hallock

### **Zeb Hallock, Technical Marketing Engineer, Systems Architecture & Strategy Unit (SASU), Cisco Systems**

Zeb Hallock is in the Enterprise Systems Engineering group of Cisco, focusing on digital media systems. He is also pursuing creation and development of future based collaboration systems, holding two patents in the field. He has been with Cisco for 10 years working on enterprise system testing, system design and testing of H.323 based video conferencing, and network infrastructure. He has also been a specialist working on Cisco Unified IP Contact Center Cisco Unified MeetingPlace. Before Cisco he worked as a consultant designing and implementing local and wide area networks.

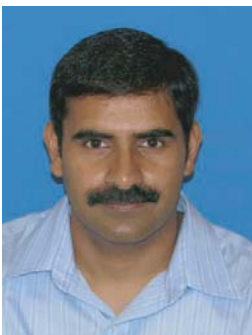


John Strika

### **John Strika, Technical Marketing Engineer, Systems Architecture & Strategy Unit (SASU), Cisco Systems**

John Strika is a Solutions Architect in Cisco's Systems Architecture and Strategy Unit and has authored white papers in the areas of mobility and location-based services. As a member of Cisco's Enterprise Architecture Board, he helps focus Cisco's vision and architectural direction and define Cisco's roadmap for context-aware mobility solutions. John was Cisco's first mobility consulting systems engineer, responsible for architecting creative wireless solutions for very large enterprise customers. His 30 years of experience spans network design and implementation, applications development, facilities planning and management, consulting, and general management. His past roles have included mission-critical telecommunications design and development at AT&T and systems programming and data communications management with Wall Street brokerages and commercial banks.

Prior to joining Cisco, John Strika was with the parent corporation of Cisco's Aironet wireless acquisition for nine years, where he was the Southern Division Vice President of Wireless Technologies and Services. John is a member of the IEEE and has held several Federal Communications Commission licenses in the use and modification of amateur and commercial radio. His educational background is in electrical engineering and computer applications programming from Columbia University and in computing quantitative methods and finance from Fordham University's College of Business Administration. He also holds a masters of communications technology certificate from the American Institute. Always seeking opportunities to use his mobility and advanced communications knowledge to improve public safety in our communities, John recently completed public safety and law enforcement training at Reinhardt University. He is active in his local community and has volunteered his time as a search and rescue team member, as well as a certified Public Safety Officer and EMS First Responder.



Srinivas Tenneti

### **Srinivas Tenneti, Technical Marketing Engineer, Systems Architecture & Strategy Unit (SASU), Cisco Systems**

Srinivas Tenneti is a technical marketing engineer for WAN and branch architectures in Cisco's Enterprise Solutions Engineering team. Previously he worked in the Commercial System Engineering team on producing design guides, and system engineering presentations for channel partners and system engineers. Tenneti holds CCIE certification 10483.

# About Cisco Validated Design (CVD) Program

---

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit <http://www.cisco.com/go/designzone>.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at <http://www.cisco.com/go/trademarks>. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Unified Access Design Guide

© 2011 Cisco Systems, Inc. All rights reserved.





# CHAPTER 1

## Unified Access Executive Summary

---

Cisco's Unified Access (UA) solution provides design guidance on solving key business problems related to network access—on-campus or remotely—with traditional devices such as laptops and desktops as well as non-traditional devices like iPhones, iPads, and other personal devices. Unified Access is an integrated solution that brings together security, mobility, management, and intelligent network infrastructure to target and solve key business problems. Unified Access clearly illustrates the importance of network infrastructure as the foundation for intelligently and dynamically solving network access challenges.

In Unified Access 1.0 we describe how businesses can deliver a consistent access experience for their employees regardless of network access location or device and deliver a consistent authentication method using wired or wireless devices. Context-awareness is introduced to the wired network to help businesses understand from where a user or device is physically accessing the network. Finally, UA helps businesses understand, in real-time, who and what are currently on the network.

## Business Problems

The Unified Access solution addresses the following business problems:

- [Network Access Experience](#)
- [Authentication for Wired Devices](#)
- [Securing Remote Access](#)
- [Securing Smart Phones](#)
- [Physical Location of Network Access](#)
- [Real Time Network Usage](#)

## Network Access Experience

One of the biggest frustrations employees face is the different methods required to access their employer's network based on their location. An employee sitting at a desk in their cubicle plugs their laptop into the network, logs in, and gets connected to the network. That same employee in a conference room logs in to the operating system, associates their computer to the wireless system, and enters their username and password to access to the network. At home with their laptop, they log into the operating system, associate their computer to their home wireless system, enter their wireless username and password to access their home network, launch a VPN client connect to their employer's network, enter a username and a one-time password, and then finally access the network. If access policies change,

software is modified, or their laptop is replaced, the employee must re-learn all these steps. If an employee received an iPad as a gift and wants to access the corporate network, then the employee must learn all of these different combinations for that device.

## Authentication for Wired Devices

When an employee accesses an employer's network on-campus using wireless technology, there is almost always a need to authenticate that device to make sure it should have access. Companies do not want unauthorized users on campus using their devices to access the internal company network. However this logic is rarely applied to wired ports in the company, which means anyone that gains access to a campus building can plug a device into a port and access the internal company network, without the need to actually authenticate in any way. Disgruntled employees, hackers, and people with bad intentions can bring personal wired devices, plug them in, and gain network access with the intent of gathering confidential information, launching attacks against the company, or generally disrupting network resources.

## Securing Remote Access

Companies generally have tight network controls about who can access what information when on campus. Network administrators use access control lists and operating system security settings to ensure compliance with security policies. However when the employee leaves the campus, enforcing a security policy becomes much more difficult so most companies resort to turning their remote access employees into second-class access entities by placing strict access rules in place, making them enter through a separate area called a DMZ. In general companies want to mitigate risk, however it makes some network resources unavailable to remote workers or, if the network as a whole is made available, then the remote worker becomes a higher risk due to the lack of policy controls.

## Securing Smart Phones

As more and more "smart" devices become available, it is natural for employees to want to use them to accomplish business goals. The invention of smart phones, such as the iPhone and Blackberry, as well as the introduction of new computing platforms, such as the iPad and Droid operating systems, can lead to severe security risks. The device being used is normally employee-owned and the "apps" running on it probably are not limited to productivity tools. In fact, many of the apps could have serious security implications. Should a company let an employee on their network with a personally-owned device? If the device is allowed on the network, from where should access be allowed? Only on-campus, remotely, or both? How will companies ensure that personal devices do not violate security policies? These devices can lead to greater employee response and benefit the company, but they also pose very serious security risks as well as the added IT support these devices require.

## Physical Location of Network Access

As companies become larger, their facilities also grow. It is easy to determine from where someone is physically accessing the company network when you only have 20 employees, but how about when you have 200, 2,000, 20,000 or 200,000? How do you know where an asset, such as a printer, laptop, or video conferencing device, actually plugs into the network? The lack of visibility can lead to business problems, such as how do you locate someone that is performing malicious activity on your network? How can you inventory the devices on your network and actually know where they reside? What happens

when a device gets misplaced and you have to find it? How can you enforce location-based security policies if you do not know the location of the person/device accessing the network? Understanding the context around where a device enters the network is the foundation of understanding and enforcing network access policies.

## Real Time Network Usage

Network administrators face a fundamental problem today in that they do not know who is actually connected to the network. Their servers can indicate who is accessing them by using programs that operate on the server operating system. However to determine what devices are physically connected to the network, you can enable authentication, which lets you know who entered the network, but rarely lets you know if they disconnected. You can enable troubleshooting tools to see what traffic is traversing the network, but it is very time consuming to piece all that traffic into a list of devices. You can enable SNMP management tools to send traps to a management server to keep a list of every time a port turns on, but this is very limited information. You can initiate a port sweep to try to discover every device, however most personal firewalls simply drop this traffic, treating it as unsecured traffic. Hence the ability to actually take a snapshot of what devices are on the network at any given time is a great tool to understand inventory, find outdated equipment that should be removed from the network, and generally understand the makeup of the network end points. Finally, what would be even more beneficial would be to actually know the user name associated with the device on the network.

## Pillars of Unified Access

The Unified Access solution brings together security, mobility, and management running over an intelligent routing and switching infrastructure to show how these items working together as a system can be used to solve specific business problems.



### Note

This document is version 1.0, which will be updated to include additional business problems that are specifically addressed by the Unified Access solution.

To fully understand the approach this solution takes, a brief description of the pillars of Unified Access is provided.

## Security

Security comprises many broad technologies that work together to ensure that the right user or device is allowed to access the company network and they are only allowed to access resources based on the security policy of the company. Sometimes security is mistakenly believed to be a specific device, such as a firewall or intrusion prevention system, but in reality those are just different tools in the security arsenal. A true security system is made up of network devices that have the ability to authenticate, authorize, and perform accounting on the network. It also includes having the right authorization servers to perform those tasks. The ability to create access controls on the network device is also critical for security policy enforcement. Devices such as firewalls also play a role in a true security system to ensure traffic traversing in or out of the company network complies with the security policy and that unwanted traffic is dealt with properly. All these items working together constitute what a network security system actually represents. So, as a part of the Unified Access solution, security means much more than a single device; it is everything described above working together to ensure that security policies are enforced.

## Mobility

Unified Access sees mobility as much more than a company's wireless system. Mobility truly represents the ability to work anywhere, whether on-campus, off-campus in a coffee shop, at the home office, or connected by wires or wirelessly. Mobility is the ability to offer the appropriate services to the end user based on where they are, what device they are using, and how they are connecting to the network. The services made available and the connection used should be dynamic and not require the end user to have to perform tedious tasks. In Unified Access mobility is made up of wireless access points, wireless controllers, mobility service engines, VPN devices, and mobility software working directly with the intelligent routing and switching infrastructure to deliver secured anytime, anywhere access to the network.

## Management

Unified Access defines management as:

- Having the ability to understand the network, how it is reacting, and what is being used on it
- Easing deployments
- Having a central location from which the network administrator can operate the company network

Management is a non-technology specific item. Management can mean:

- Deploying security features, such as identity-based networking using configuration templates, to routers and switches
- Having a list of every device on the network and where they are located
- Having a location where intelligent network devices actually send alerts to notify a network administrator of a situation on the network

The key to managing a network is to have the management station act as the coordinator of the network, allowing it to interface with routers, switches, security systems, mobility systems, and any other part of the network to build a holistic view of the network and to make operating the network easier.

## Unified Access 1.0 Solution

In Unified Access 1.0 we tie solving the above business problems into specific chapters. Some of the chapters address multiple business problems and include:

- [Chapter 3, “Bring Your Own Device—Unified Device Authentication and Consistent Access Experience”](#)
- [Chapter 4, “Context-Awareness for Wired Devices”](#)
- [Chapter 5, “User and Device Network Access Reporting”](#)

The following subsections provide an overview of each chapter.

## Bring Your Own Device—Unified Device Authentication and Consistent Access Experience

The chapter focuses on identifying and authenticating users connecting to the network from different places with different devices over different connection media. The devices can be categorized as laptops, desktops, and smart phones, the places can be on the campus or remote, and the connection media can be wired, wireless, or over an un-trusted network.

The specific business problems that are covered in this chapter include:

- Unified Access experience whether on-campus or remote, using either a traditional or non-traditional networking device
- Wired Device Authentication for networking devices that connect to the network, including describing how to tie the authentication into existing directory systems such as Microsoft's Active Directory
- Securing Remote Access for employees that work off the network. Detailed authentication and authorization information is provided to allow network administrators to implement highly-secure, policy-based remote access.
- Securing Smart Phones for company network access. Details describe how you can enforce security policies on smart phones as they connect to the company network.

## Context-Awareness for Wired Devices

This chapter focuses on applying best practices towards the enablement and use of Cisco Context-Aware Services for wired endpoint location tracking in Unified Access solutions. It is intended as a guide to producing functional designs that incorporate the Cisco Mobility Services Engine (MSE) and Cisco Wireless Control System (WCS) to help fulfill common business needs in enterprise environments. Very specific use cases of how a company can use context awareness are detailed, such as building asset inventories, recovering lost assets, and discovering user, server, and network device locations.

## User and Device Network Access Reporting

This chapter describes how to use Cisco's LAN Management Solution (LMS) to collect user and device information about who is accessing the company network, what device(s) they are using, and what specific networking device are they connecting into to provide information such as:

- Current port utilization of the network
- Movement of employees and devices on the network
- Suspicious or unauthorized access to the network

Additionally, this chapter shows how companies can use historical information to provide the network administrator with:

- Historical port utilization data
- Persistent records of when users and devices accessed specific locations in the network
- Searchable data of historical access for troubleshooting and asset tracking

# Overview of Topics Anticipated in UA 2.0

At the time of writing of the Unified Access 1.0 document, efforts were underway to expand the number of business problems that Unified Access will address. The Unified Access solution will be adding additional technologies and systems to address issues that companies and network administrators are experiencing. A preview of some of the high-level business problems that Unified Access 2.0 will address is provided below (subject to change).

## Dynamic Policy Enforcement

Building upon the work done in Unified Access 1.0, the next step is the evolution of security policy enforcement. Unified Authorization and Accounting, along with context awareness, an intelligent network infrastructure, and the Access Control solutions, will be brought together to allow a company to have the network dynamically enforce the company's security policies.

The network administrator will be given the tools to allow them to create security policies much as is done now for server operating systems. They will be able to create groups and users and define what each user or device can access on the network based on who they are, what they are, what they can do (from a device posture standpoint), when they are accessing the network, and from what location. This gives companies the ultimate amount of flexibility to implement detailed security policies into a centralized system, which in turn works with the intelligent network infrastructure to dynamically enforce network access security policies.

## Mobility Service Optimization

Unified Access 2.0 will explore business problems related to mobility services, such as providing high-quality uninterrupted video services and wideband audio to mobile devices over multiple network types:

- Build a system that will automatically mitigate RF interference in a campus wireless system and alert the network administrator to the source and physical location of that interference.
- Deliver context-aware information from the wireless environment so that network policies and management stations can understand where someone is accessing the network when they are connected via a wireless technology.

## Network Device Configuration Optimization

This high-level topic will explore business problems that network administrators face when deploying new features on existing or new network infrastructure equipment. Unified Access 2.0 will also highlight how network information is gathered and shared among mobility, security, and management systems to give the network administrator central control and information about their network. Finally, the business problem of easing the level of complexity when deploying intricate solutions, such as identity-based networks, will be explored.



## CHAPTER 2

# Unified Access Network Design and Considerations

---

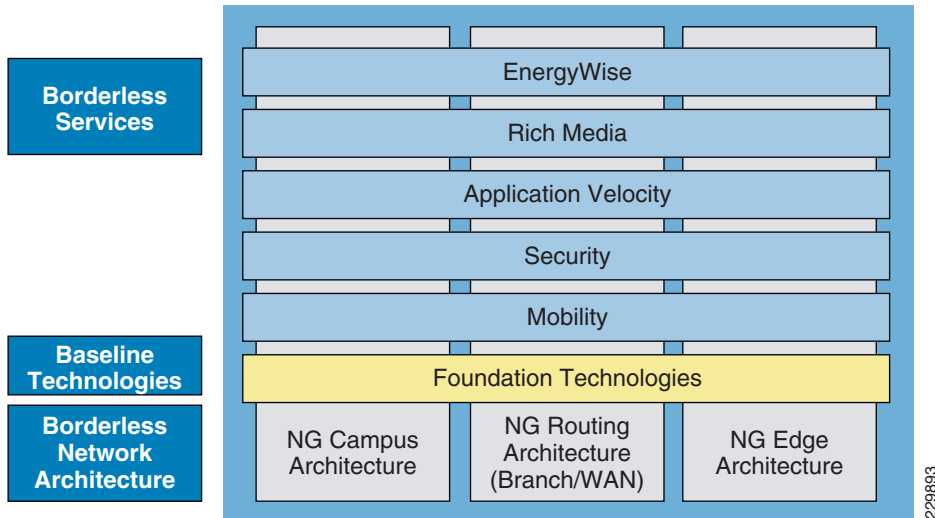
## Cisco Borderless Network Architecture

The Unified Access Solution uses at its foundation the Cisco Borderless Network architecture and the Cisco Borderless Campus Design principles. All Unified Access elements work on top of this architecture, which is specifically designed to provide the proper foundational features needed for the Unified Access services being deployed. This chapter describes the foundational principles of the Cisco Borderless Network Architecture and provides specific guidance to design an intelligent network infrastructure to handle borderless services. We discuss the design choices, switching platforms, and network features you need to understand about the intelligent network that Unified Access solutions use to solve business problems.

For an in-depth discussion of how to build and deploy this architecture, see the *Cisco Borderless Campus 1.0 Cisco Validated Design Guide* at:

[http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/Borderless\\_Campus\\_Network\\_1.0/Borderless\\_Campus\\_1.0\\_Design\\_Guide.html](http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/Borderless_Campus_Network_1.0/Borderless_Campus_1.0_Design_Guide.html).

The Cisco Borderless Network architecture is a next-generation architecture that allows different elements of the network, from access switches to wireless access points, to work together and allow users to access resources from anywhere at anytime. The Cisco Borderless Network integrates key services into the network fabric while increasing reliability and security and decreasing outages. For such an infrastructure, the enterprise network must be developed with an architectural approach that embeds intelligence, simplifies operations, and scales to meet future demands. The Cisco Borderless Network is composed of several modular components, as illustrated in [Figure 2-1](#).

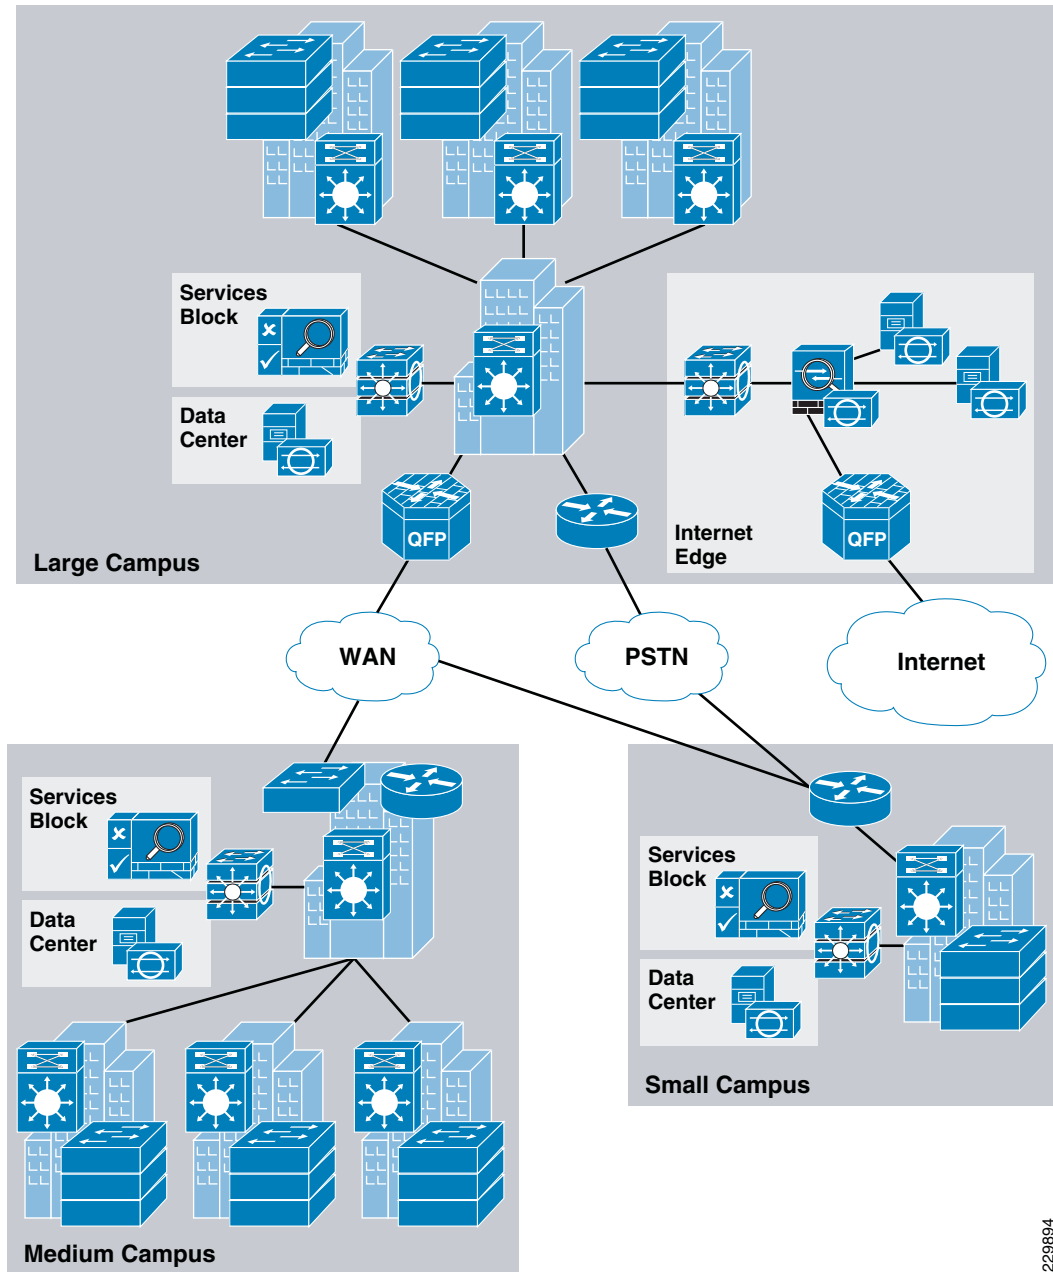
**Figure 2-1 Cisco Borderless Network Framework**

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

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

## Borderless Campus Network Design

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

**Figure 2-2** *Borderless Campus Network Design*

229894

The Cisco Borderless network architecture focuses on the campus framework and network foundation technologies that provide a baseline of routing, switching, and several key network services. The campus design connects infrastructure components, such as devices in the access layer, the services block, the WAN, and so on, to provide a foundation on which mobility, security, and management, as well as other key services, can be integrated into the overall design.

The Cisco Borderless Campus provides guidance on building next-generation enterprise networks, which with the addition of critical network technologies become the framework to deliver the foundation for Unified Access. This chapter details the approach of the Cisco Borderless Network Architecture and is divided into these sections:

- *Campus design principles*—Provides proven network design choices to build various types of campus infrastructure.
- *Campus design model for the enterprise*—Leverages the design principles of a tiered network design to facilitate a geographically-dispersed enterprise campus network made up of various elements.
- *Considerations of a multi-tier campus design model for enterprises*—Provides guidance for the enterprise campus LAN network as a platform with a wide range of next-generation products and technologies to seamlessly integrate applications and solutions.

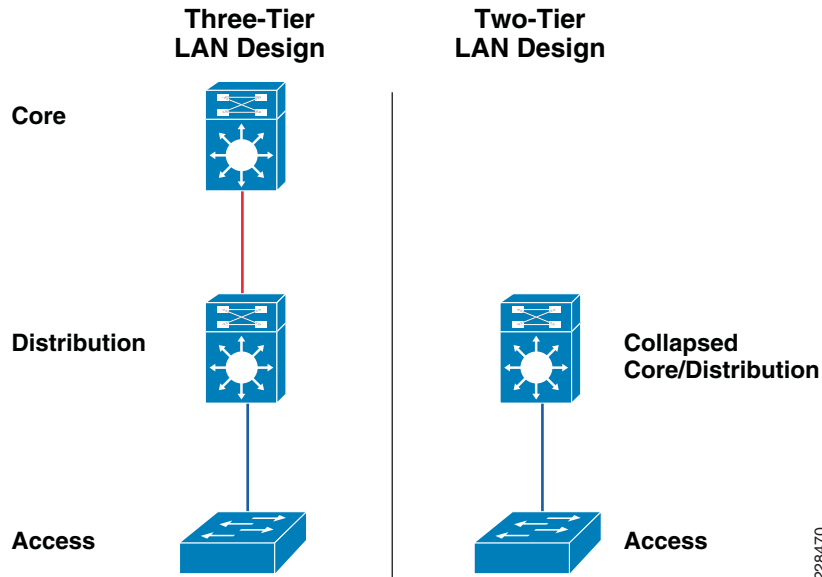
## Borderless Campus Network Design Principles

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

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

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

Designing the Borderless Campus network in a hierarchical fashion creates a flexible and resilient network foundation that allows network architects to overlay the security, mobility, and management features essential to the Unified Access Solution. The two proven, time-tested hierarchical design models for campus networks are the three-tier layer and the two-tier layer models, as shown in [Figure 2-3](#).

**Figure 2-3 Three-Tier and Two-Tier Campus Design Models**

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

- *Access layer*

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

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

- *Distribution layer*

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

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

- *Core layer*

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

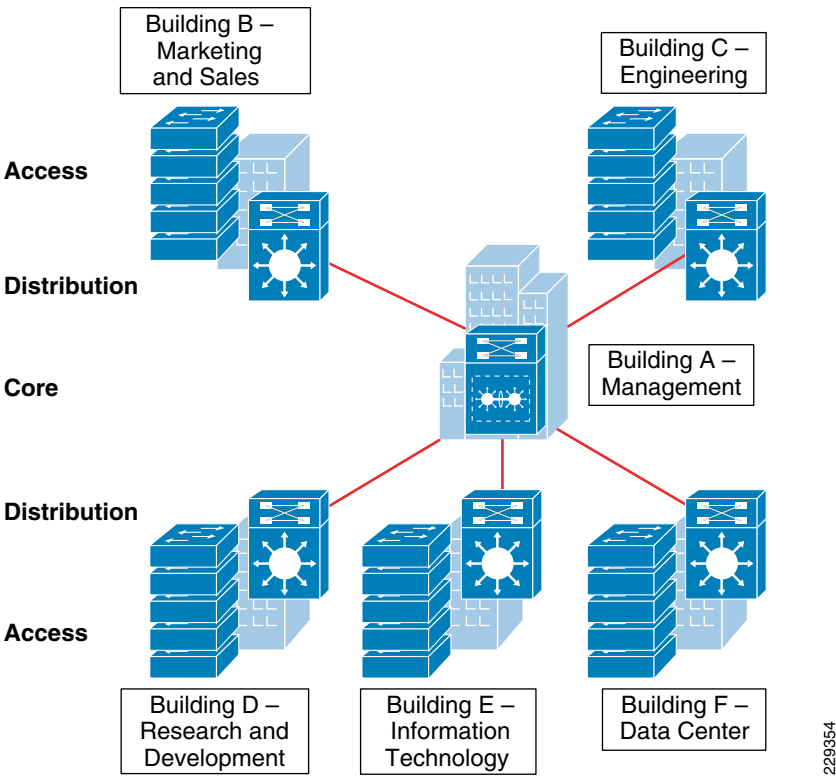


**Note**

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

Figure 2-4 shows a sample three-tier campus network design for enterprises where the access, distribution, and core are all separate layers. To build a simplified, scalable, cost-effective, and efficient physical cable layout design, Cisco recommends building an extended-star physical network topology from a centralized building to all other buildings on the same campus.

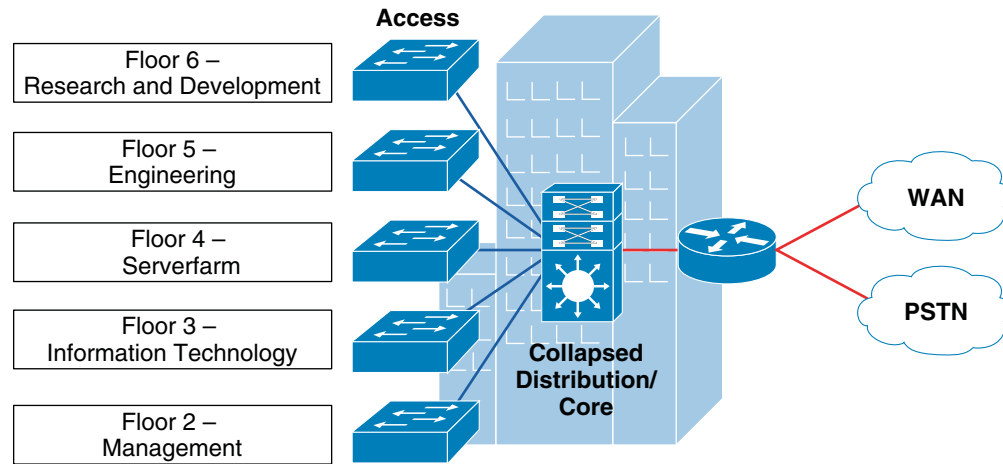
**Figure 2-4 Three-Tier Campus Network Design Example**



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

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

Figure 2-5 shows an example of a two-tier campus network design for small enterprise campus locations where the distribution and core layers are collapsed into a single layer.

**Figure 2-5 Two-Tier Network Design Example**

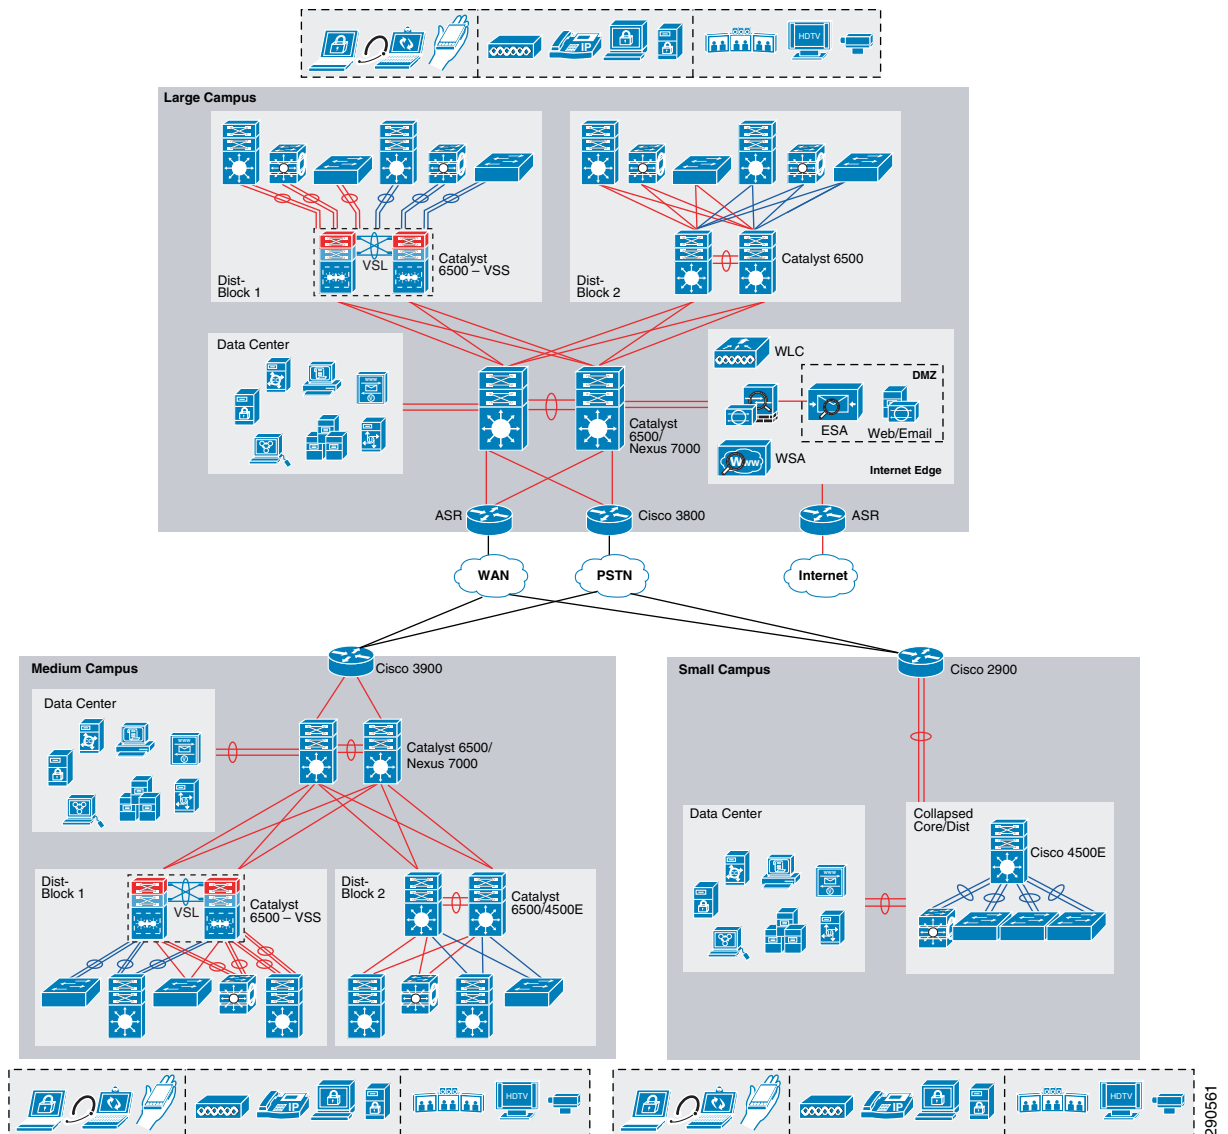
If using the small-scale collapsed campus core design, the enterprise network architect should understand the network and application demands so that this design ensures a hierarchical, modular, resilient, and flexible campus network.

## Borderless Campus Network Design Models

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

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

As shown in [Figure 2-6](#), multiple campuses can co-exist within a single enterprise system that offers borderless network services.

**Figure 2-6 Borderless Campus Network Design Model**

Depending on the size and number of users and devices on the medium and small campuses, their relative network size should be less than the large campus. Hence compared to the large campus network, the medium and small campus sites may have alternative network designs that can provide network services based on overall campus network capacity.

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

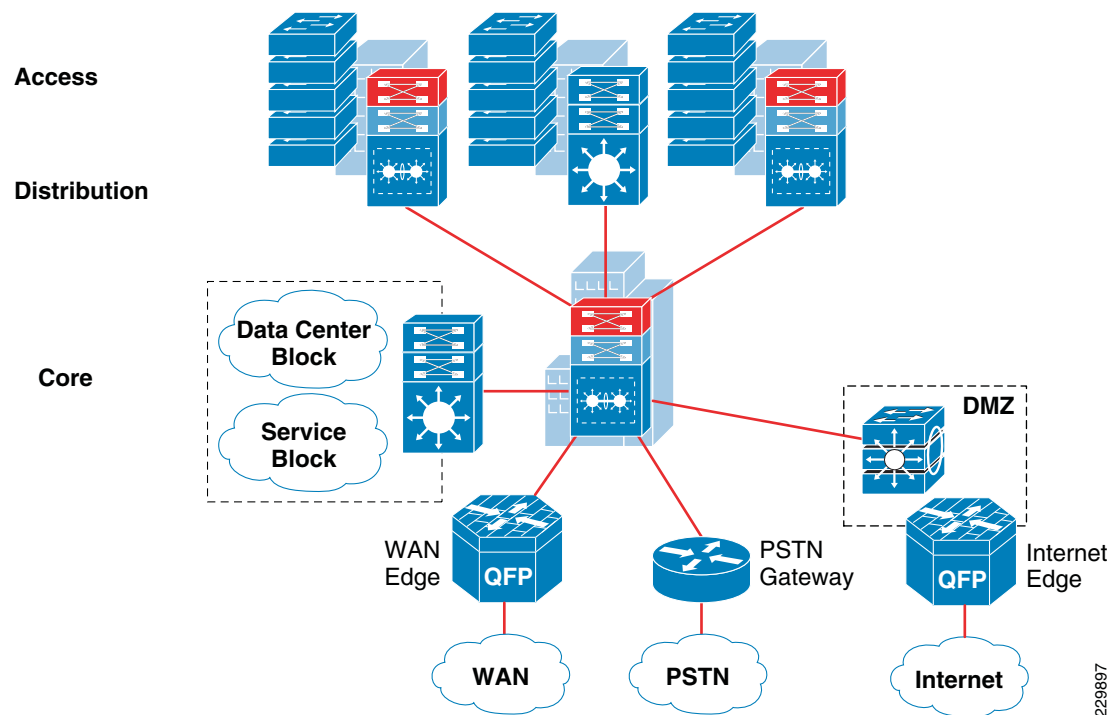
Table 2-1 shows a summary of the Borderless Campus Network design models as they are applied in the different enterprise network designs.

**Table 2-1 Enterprise Recommended Campus Design Models**

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

## Large Campus Network Design

The large campus in the Borderless Campus consists of a centralized hub that interconnects several medium and small campuses to provide end-to-end access to resources and borderless services. The network in the large campus is larger than the medium and small campuses and includes end users, devices, servers, security, mobility, and management devices. Multiple buildings of various sizes exist in one location, as shown in [Figure 2-7](#).

**Figure 2-7 Large Campus Reference Design**

The large campus utilizes a three-tier campus design model to meet all key technical requirements and provide a strong, well-structured network foundation. The modularity and flexibility of the three-tier campus design allows for easier expansion of the large campus network and keeps all network elements protected and available.

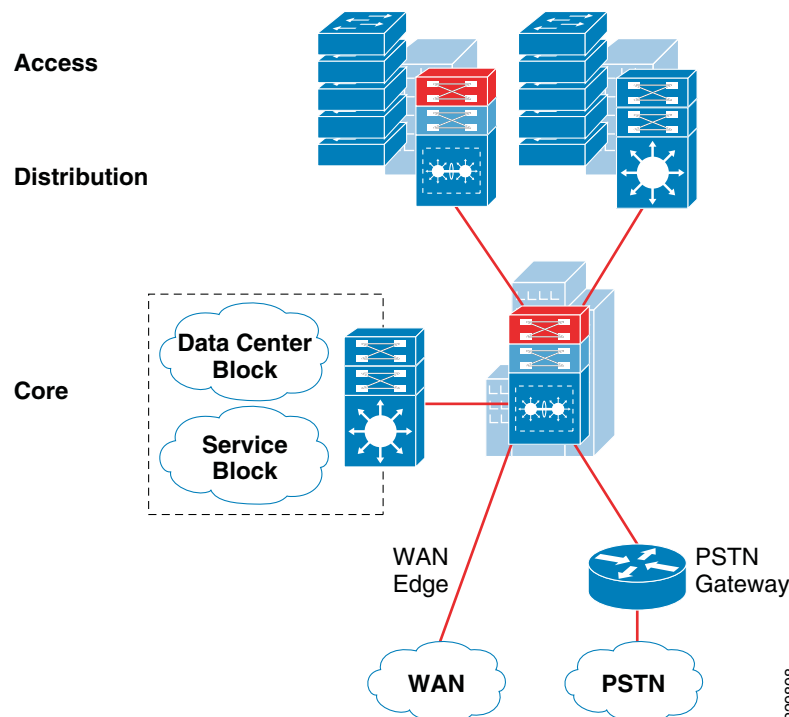
To enforce external network access policy, the large campus also provides external gateway services to employees to access the network to and from the Internet.

## Medium Campus Network Design

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

Similar to the large campus network design, Cisco recommends the three-tier campus design model for the medium campus, as shown in [Figure 2-8](#).

**Figure 2-8 Medium Campus Reference Design**



## Small Campus Network Design

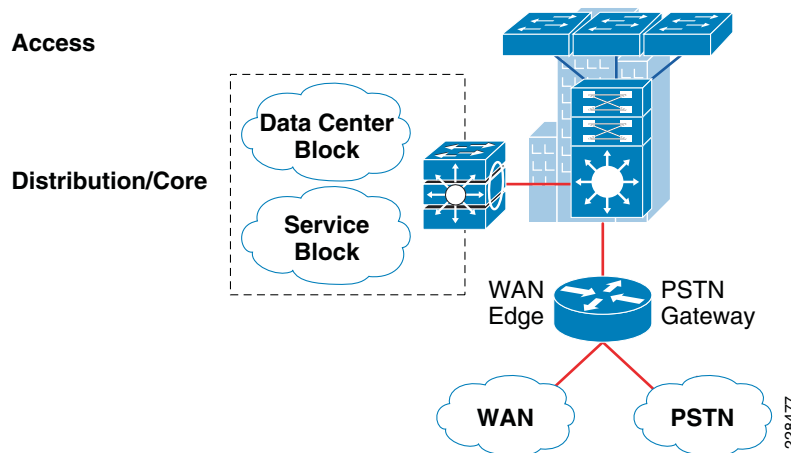
The small campus is typically confined to a single building that spans across multiple floors with different organizations. The network scale in this design is smaller compared to the large and medium campuses, however application and borderless services demands are the same as in the medium and large campuses.

In a smaller campus network deployment, the distribution and core layer functions can be collapsed into the two-tier design without compromising basic network requirements. Prior to deploying the collapsed core and distribution system, network architects must consider scale, expansion, and manageability of the network to ensure the network meets current and future enterprise requirements.

The necessary WAN bandwidth must be assessed appropriately for the small campus network design. Although the network size is reduced compared to other campuses, sufficient WAN capacity is needed to deliver an appropriate collapsed core and distribution design. This alternative and cost-effective

network design is recommended only in smaller locations and only when WAN traffic and application needs must be considered before choosing this model. Figure 2-9 shows the small campus network design in more detail.

**Figure 2-9 Small Campus Reference Design**



## Multi-Tier Borderless Campus Design Models

This section provides more detailed network infrastructure guidance for each tier in the campus design model. Each design recommendation is optimized to keep the network simplified and cost-effective without compromising network scalability, security, and resiliency.

### Campus Core Layer Network Design

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

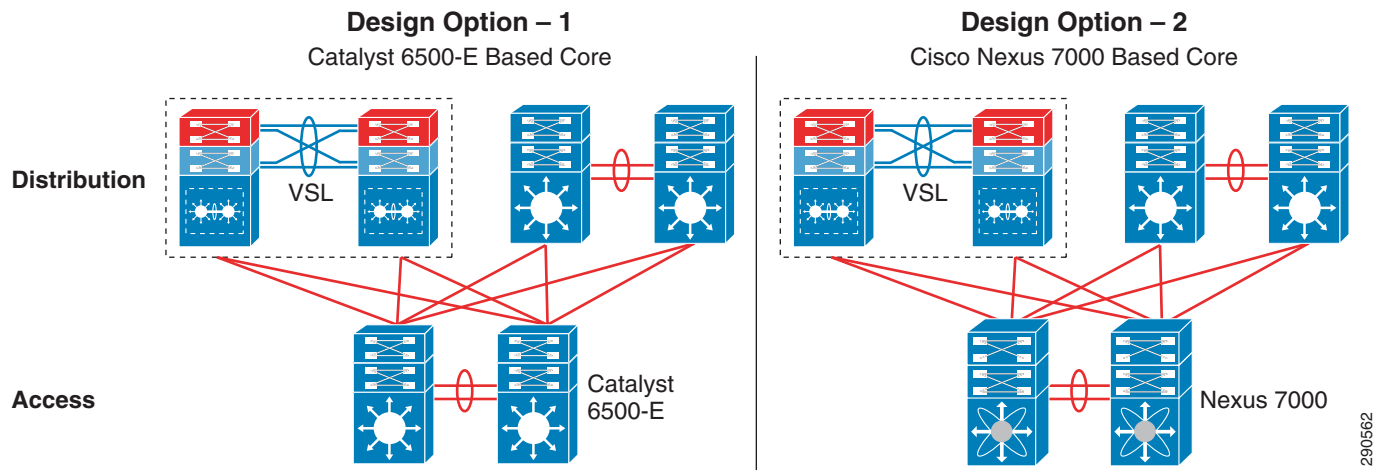
When network architects are designing a campus core, it becomes imperative take into consideration network scalability, capacity, and reliability to allow for high-performance end-to-end borderless services. Determining the core layer scalability and performance may be challenging as it varies depending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

Cisco recommends building the next-generation borderless campus core with following principles:

- The architecture should be designed to support modern technologies that enable advanced networking and integrated services to solve key business problems.
- Scalability to adapt to enterprise network needs, as well as the ability to provide for intelligent borderless network services.
- Flexible design options that maximizes return on investment (ROI) and reduces total cost of ownership (TCO).

These design principles are important when designing the core network so that the core is capable of addressing current and future borderless network demands. Cisco recommends the Cisco Catalyst 6500-E and Nexus 7000 switching platforms for the core of the next generation borderless campus. These multi-terabit switching platforms are designed with a robust hardware architecture that exceeds the foundational borderless campus requirements. [Figure 2-10](#) illustrates core designs for building the next-generation Borderless Campus Core.

**Figure 2-10** Core Layer Design Model Options



## Cisco Catalyst 6500-E

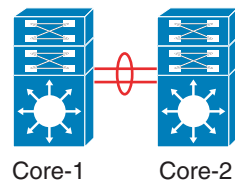
The industry-leading and widely-deployed Cisco Catalyst 6500-E series platform is the lead system to deploy in a borderless campus core role. Because of its advanced hardware and software innovations, the Catalyst 6500-E switching platform is the preferred way to build an enterprise-class borderless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

To provide mission-critical network services, it is recommended that the core layer be deployed with high resiliency, such as using dual Cisco Catalyst 6500-E systems. Deploying resilient standalone core layer switches with redundant hardware provides constant network availability for business operation during faults and also provides the ability to load share high-speed network traffic between different blocks (e.g., distribution and service block). A redundant core network design can be deployed in a traditional standalone model or in a Virtual Switching System (VSS) model. The campus core layer network design and operation broadly differ when the core layer is deployed as a standalone, which operates all three planes (forwarding, control and data planes) in isolation. However with Cisco VSS technology, two core systems are clustered into a single logical system and the control and management planes get combined on the systems to produce a single logical Catalyst 6500-E core system.

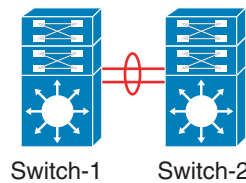
The standalone/VSS physical and operational view is shown in [Figure 2-11](#).

**Figure 2-11 Standalone/VSS Physical and Operational View****Catalyst 6500-E Core – Standalone Mode**

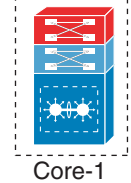
Physical View

**Catalyst 6500-E Core – VSS Mode**

Physical View



Logical View



290563

**Note**

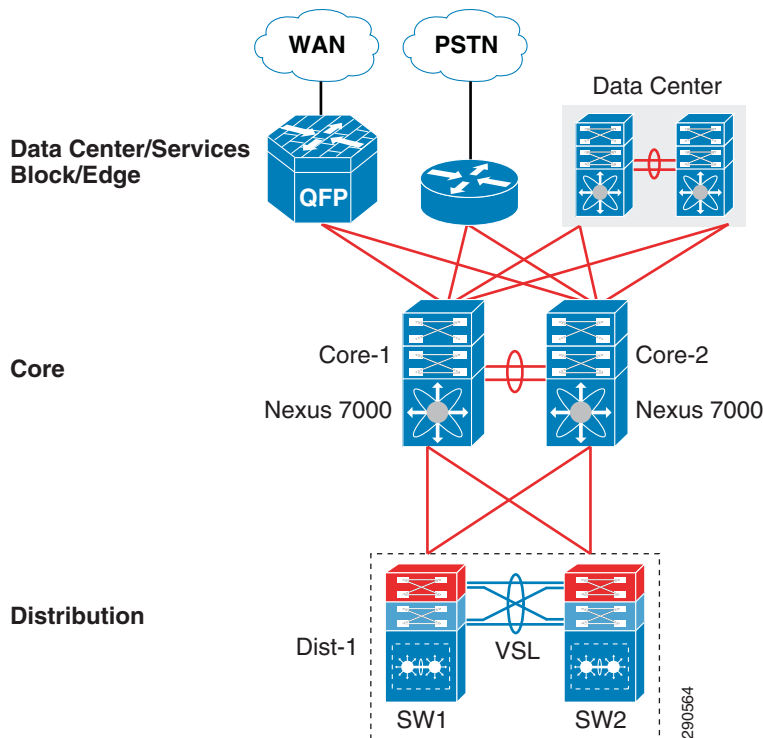
For more detailed VSS design guidance, see the *Campus 3.0 Virtual Switching System Design Guide*: [http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/VSS30dg/campusVSS\\_DG.html](http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/VSS30dg/campusVSS_DG.html).

## Cisco Nexus 7000

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

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

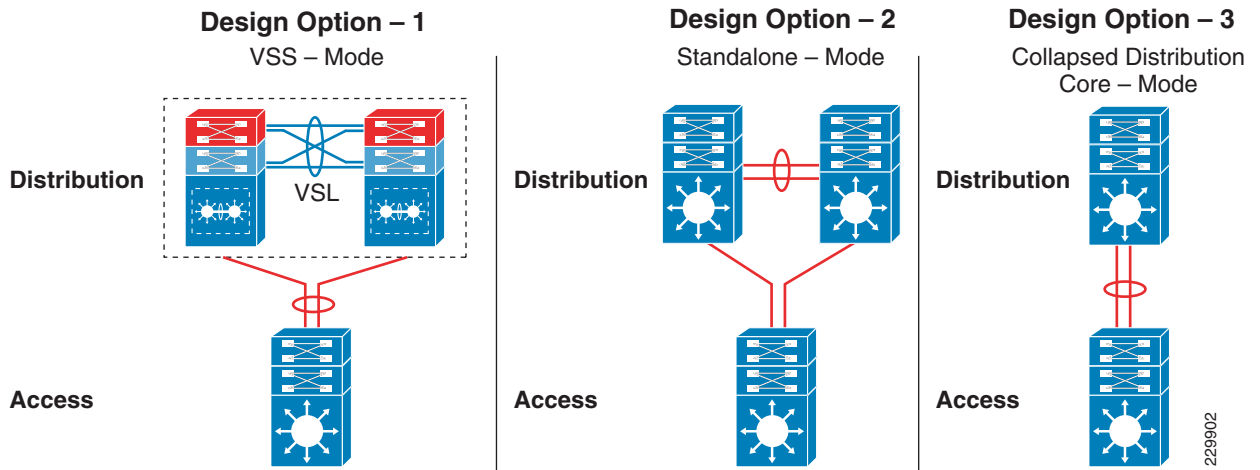
Figure 2-12 illustrates core network design options with the Cisco Nexus 7000 peering with other Cisco platforms to enable end-to-end business communication:

**Figure 2-12 Cisco Nexus 7000 Campus Core Design**

## Campus Distribution Layer Network Design

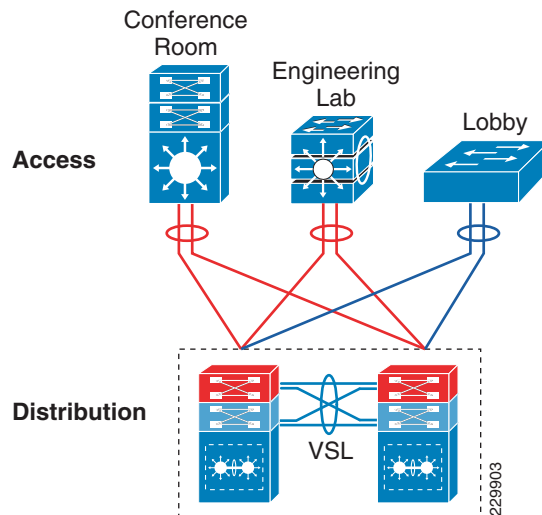
The distribution or aggregation layer is the network demarcation boundary between wiring closet switches and the campus core network. The framework of the distribution layer system in enterprise design is based on best practices that reduce network complexities, increase reliability, and accelerate network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

The distribution layer design options provide consistent network operation and configuration tools to enable various borderless network services. Three simplified distribution layer design options can be deployed in large, medium, and small campus locations, depending on network scale, applications used, borderless services demands, and cost, as shown in [Figure 2-13](#). All distribution design models offer consistent network foundation services, high availability, expansion flexibility, and network scalability. However each enterprise network is different, with its own unique business challenges that require the appropriate aggregation solution. Factors that should be taken into consideration when selecting the right distribution model include scalability, high-speed network services, virtualized systems, and cost. Depending on network designs and key technical requirements, the network architect must make appropriate aggregation layer design choices to enable end-to-end borderless network services.

**Figure 2-13** *Distribution Layer Design Model Options*

## Distribution Layer Design Option 1—VSS Mode

Distribution layer design option 1 is intended for the large and medium campus network design and it is based on deploying Cisco Catalyst 6500-E Series switches using Cisco VSS, which lowers the management burden and allow multiple switches to work as one single virtualized switch, as shown in Figure 2-14.

**Figure 2-14** *VSS-Enabled Distribution Layer Network Design*

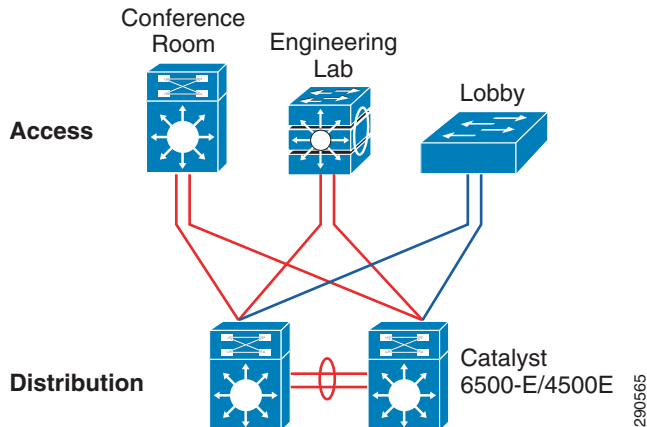
## Distribution Layer Design Option 2—Standalone Mode

Distribution layer option 2 is the traditional and proven network design used in many enterprise campus networks. It can be deployed with resilient Cisco Catalyst 6500 or 4500E switches to operate as standalone switches. This is an alternative distribution network deployment design if there is no desire

to virtualize the aggregation layer switches using Cisco VSS technology. The Cisco Catalyst 6500 without Virtual Switch Link (VSL) capable supervisors can be deployed as a standalone solution, as well as the Catalyst 4500E switches.

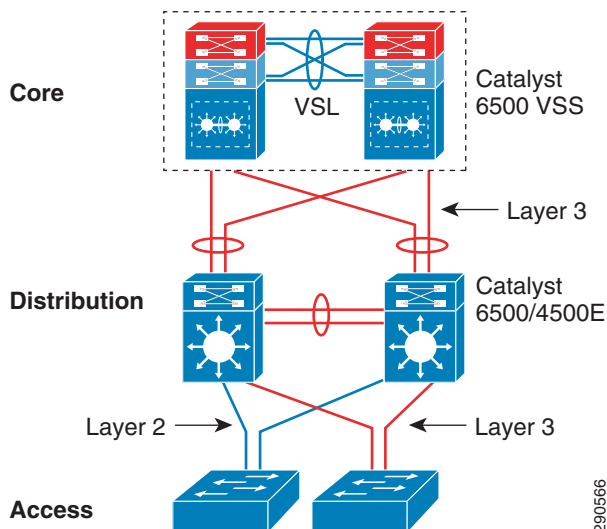
The two single-chassis standalone mode distribution layer design options are shown in [Figure 2-15](#).

**Figure 2-15 Standalone Mode Distribution Layer Network Design**



In standalone mode each Catalyst distribution switch operates independently and builds local network adjacencies and forwarding information with the access and core layers. The Layer 2 and Layer 3 protocols operate over each physical interface between the standalone distribution switches and the access layer switches. Since the core layer in the large and medium campus networks is simplified using Cisco VSS technology, the network administrator can simplify the core network topology by bundling Layer 3 interfaces into a logical EtherChannel, as shown in [Figure 2-16](#).

**Figure 2-16 Network Design with Distribution in Standalone Mode**



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

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

Cisco innovated VSS technology to mitigate such challenges, hence it is recommended to deploy a Cisco VSS-based distribution layer infrastructure that simplifies the multilayer network and increases network capacity and performance, resulting in a highly-reliable network that provides consistent and deterministic network recovery. The traditional standalone-mode distribution layer network is an alternative solution that does not introduce any fundamental design changes.

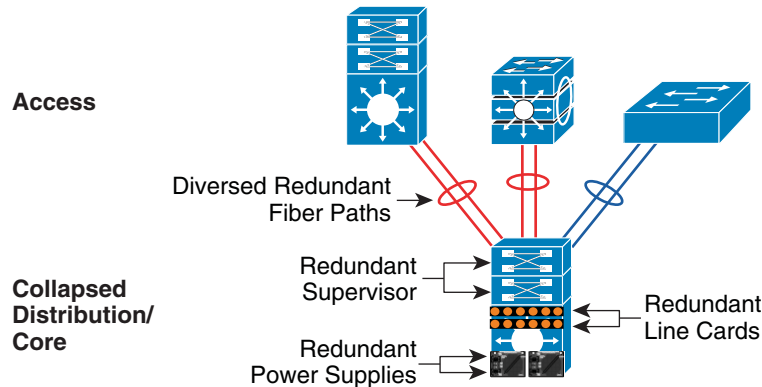
**Note**

For more information on configuring and deploying standalone mode distribution layer Catalyst switches, see the *Campus Network for High Availability Design Guide*:  
[http://www.cisco.com/en/US/partner/docs/solutions/Enterprise/Campus/HA\\_campus\\_DG/hacampusdg.html](http://www.cisco.com/en/US/partner/docs/solutions/Enterprise/Campus/HA_campus_DG/hacampusdg.html).

## Distribution Layer Design Option 3—Collapsed Distribution/Core Mode

The small remote campus may have several departments working on various floors within a building. Network administrators can consider collapsing the core function into the distribution layer switch for a small campus where there may only be a single distribution block. The collapsed distribution/core switch can provide network services to a small number of wiring closet switches and directly connect to the WAN edge to reach the large campus for centralized data and communication services. This solution is manageable and cost effective as long as it meets the needs of the network users and endpoints; capacity is the main factor that restricts using this model.

The collapsed distribution/core network can be deployed with two resilient systems as recommended in [Distribution Layer Design Option 1—VSS Mode](#) or alternatively in standalone mode as described in [Distribution Layer Design Option 2—Standalone Mode](#). In a space-constrained small campus environment, a single Cisco Catalyst 4500E series platform can be deployed with resilient hardware to build a highly-available, collapsed distribution/core system that has the appropriate network performance, availability, and reliability required to run borderless services. With resilient hardware this solution can provide 1+1 in-chassis protection against hardware and software failure. Deploying the network in a recommended design provides consistent sub-second network recovery in the event of an unplanned outage. A single Cisco Catalyst 4500E with multiple resilient system components can be deployed as shown in [Figure 2-17](#).

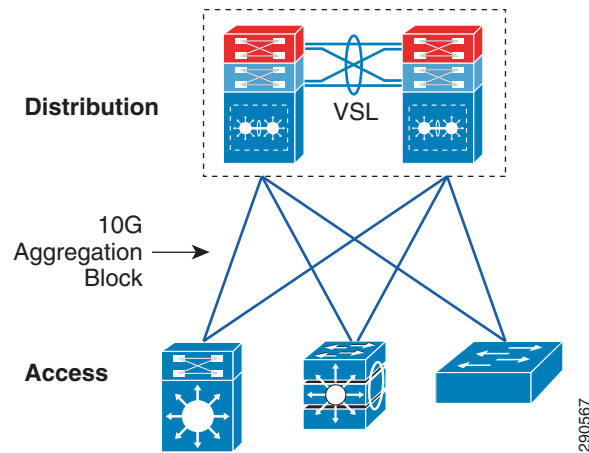
**Figure 2-17** *Highly Redundant Single Collapsed Distribution/Core Design*

229906

## Campus Access Layer Network Design

The access layer is the first tier or entry point into the campus network; it is where end devices such as PCs, printers, cameras, phones, and so on attach to the wired and wireless campus network. The wide variety of possible types of devices that can connect and the various services and dynamic configurations that are necessary make the access layer one of the most feature-rich parts of the Borderless Campus network. Not only does the access layer switch allow users to access the network, the access layer switch provides network protection so that unauthorized users or applications do not access the network. The challenge for the network architect is determining how to implement a design that meets this wide variety of requirements, the need for various levels of mobility, and the need for a cost-effective and flexible environment, while being able to provide the appropriate balance of security and availability expected in more traditional, fixed-configuration environments. The next-generation Cisco Catalyst switching portfolio includes a wide range of fixed and modular switching platforms, each designed with unique hardware and software capabilities to function in the access layer.

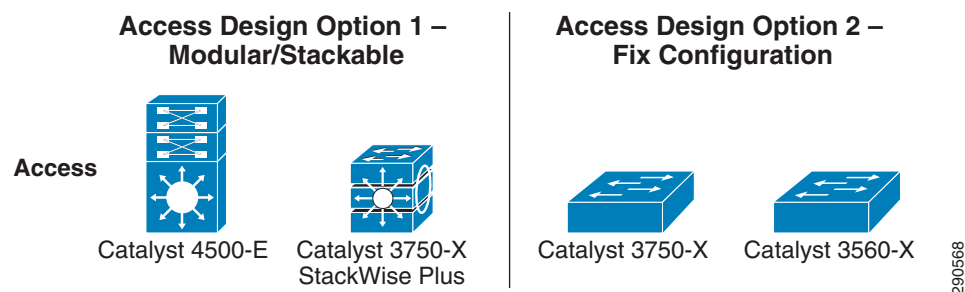
Enterprise campuses may deploy a wide range of network endpoints which all have different requirements on the network; low-latency, link speed, and low-jitter rates are just some of those requirements. The network architect must take into consideration network requirements, as well as the planned growth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

**Figure 2-18 High-Performance Distribution-Access Block**

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

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

Based on the broad-range of business communication devices and endpoints, network access demands, and capabilities, two access layer design options can be deployed, as shown in [Figure 2-19](#).

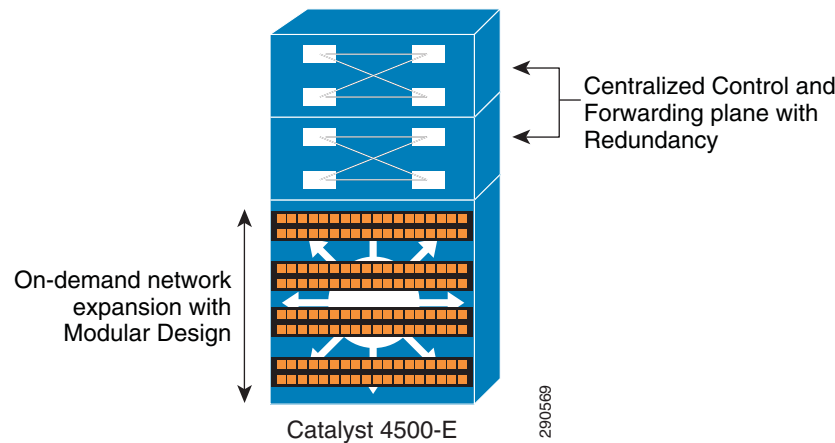
**Figure 2-19 Access Layer Design Models**

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

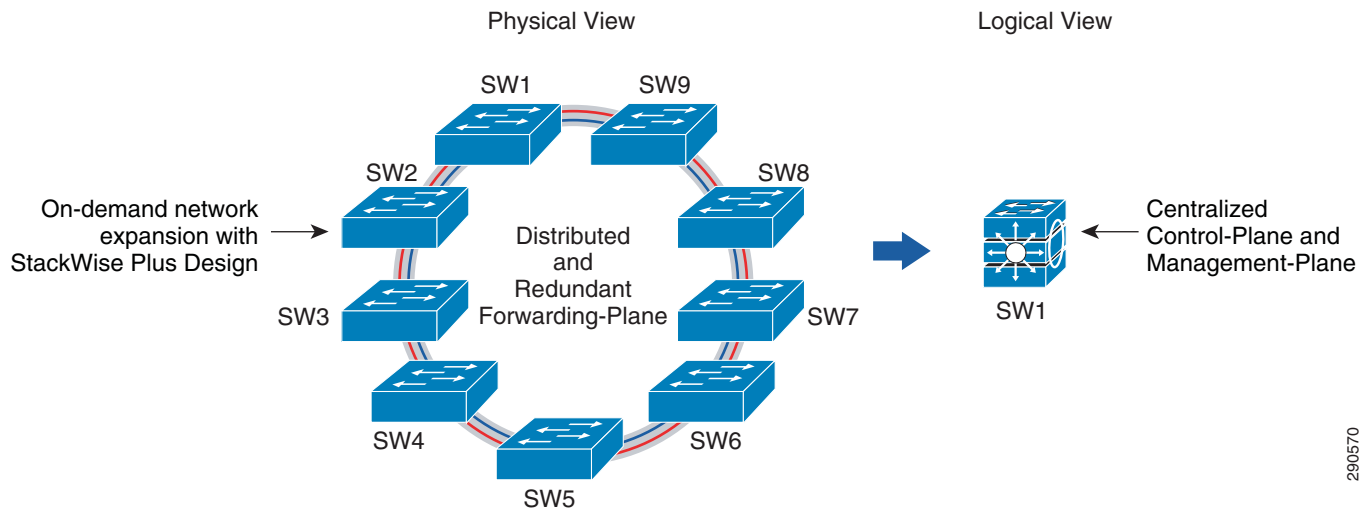
Access layer design option 1 is intended to address network modularity, performance, scalability, and availability as well as the support for advanced Borderless Network services needed by the multitude of endpoints accessing the network. Implementing a modular and stackable Cisco Catalyst switching platform provides the flexibility to expand or contract the number of ports needed at the access layer, as well as allowing for modular upgrades for future requirements.

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

**Figure 2-20 Network Edge Expansion with Modular Design**



The Cisco Catalyst 3750-X Series is an alternative Cisco access layer switching platform. Using Cisco StackWise Plus technology provides flexibility and availability by clustering multiple Cisco Catalyst 3750-X Series Switches into a single high-speed stack ring that simplifies operation and allows incremental access layer network expansion or contraction. Catalyst 3750-X switches deployed in Cisco StackWise Plus mode changes network operation compared to standalone mode. When deployed in StackWise Plus mode, the switches become a single logical access layer switch, the control plane processing becomes centralized, and because of the distributed forwarding architecture, all the hardware resources gets fully utilized across all stack member switches (see [Figure 2-21](#)). Cisco StackWise Plus provides high-speed multigigabit switching capacity for network traffic switching within the stack ring and the distribution-access block can be built with multiple parallel 10Gbps uplinks for load sharing and network resiliency. The network is optimized and simplified when the cross-switch uplink ports are bundled into single logical interface using EtherChannel technology. This network design provides non-stop network communication in case of an individual stack member switch failure.

**Figure 2-21 Network Edge Expansion with StackWise Plus Design**

## Access Layer Design Option 2—Fixed Configuration Access Layer Network

The fixed configuration switch access layer design is widely chosen for enterprise environments today as it enables quick redeployment of networking resources. The fixed configuration Cisco Catalyst switching portfolio supports a wide range of access layer technologies that allow seamless service integration and enable intelligent network management in the access layer. The fixed configuration Cisco Catalyst switches deployed as standalone switches are an ideal design choice for a small wiring closet to provide consistent borderless network services for up to 48 endpoints.

The Cisco Catalyst 3750-X and 3560-X Series switches are the recommended platforms for wired network access that can be deployed in a mixed configuration for critical and non-critical end point devices, such as Cisco IP phones, PCs, printers, and so on. For non-stop network operation during power outages, the Catalyst 3560-X requires an internal or external redundant power supply such as the Cisco RPS 2300. The increased power capacity allows flexibility when deploying enhanced Power-over-Ethernet (PoE+) on a per-port basis and with its wire-speed 10G uplink forwarding capacity, this design reduces network congestion and latency to significantly improve application performance.

To provide a consistent end-to-end user experience, the Cisco Catalyst 3750-X and 3560-X Series platforms support critical network control services to secure the access layer and intelligently provide differentiated services to various types of traffic, as well as simplified management. The Cisco Catalyst should be deployed with dual uplinks to the distribution layer for increased bandwidth capacity and network availability.

Both the modular and fixed design options offer consistent borderless network services in the access layer as well as differentiated, intelligent, and secured network access to trusted and un-trusted endpoints/devices.

## Summary

The intelligent network infrastructure is at the center of the Unified Access Solution. Its proper architecture and deployment, as well as the features and configuration, are critical to the ability of the network to solve business problems. Without proper design, bottlenecks can occur, services will not be differentiated, and priorities will not be given to different types of network traffic. For too long the

panacea to any network issue has been to just increase bandwidth; this is no longer the case. With network traffic increasing exponentially due to more and more services moving to the network, the network can no longer act as a dumb packet pump. It must be aware of the different traffic traversing it and it must be able to treat that traffic based on the requirements necessary to ensure that traffic reaches its destination within the requirements of the application utilizing the network. The network needs intelligence, and with the design guidance provided along with Cisco switching platforms, the network becomes the intelligence that can deliver solutions to key business problems.



## CHAPTER 3

# Bring Your Own Device—Unified Device Authentication and Consistent Access Experience

---

This chapter focuses on identifying and authenticating users who connect to the network with different devices and at different places. The devices can be categorized as laptops, desktops, and smart phones and the different places could be on a campus or at a remote location. This chapter first outlines the core components that are needed to identify and authenticate users and then gives significant details on how to configure the core components that authenticate and identify users. This chapter is intended to build on existing best practices of authentication. Where needed this chapter provides references to additional information.

## Executive Summary

### Drivers for Bring Your Own Device—Unified Device Authentication and Consistent Access Experience

The rapid propagation of mobile devices and their introduction into the enterprise workplace has caused a paradigm shift in regards to the term “end points” and to the phrase “being at work”. The distinction between a device used exclusively for “work” and a device used exclusively for “personal use” is fading, a trend often referred to as Bring Your Own Device (BYOD). Additionally some enterprise environments are shifting away from providing their employees with a “standard imaged computing device” to an environment where the employee owns and controls the computing device that is used for work. The challenge for IT organizations is to provide the end user with the freedom to use any device while enforcing stringent security policies to protect corporate intellectual property when the user device joins the secure campus network. In addition to providing a proper security policy, it is also imperative to provide a consistent access experience.

There is also a need to provide the end user and end point devices with transparent access to corporate resources while enforcing the proper security policy for wired, wireless, and remote user authentication. The location of the end point and user (from a logical perspective as well as a physical perspective) should have no bearing on the experience that the user observes. The user should have the capacity to seamlessly move from wired, wireless, and remote locations, where each environment is employing different technologies. The user should not have to distinguish or understand the technology being

adopted that is providing authentication and access to the secure corporate network. Any user should be able to connect from any device at any time. The security controls in place, however, must be able to understand how the employee is connecting and apply the proper security policy for that connection.

The following are the challenges faced by an IT organization today:

- How to provide [Secure Campus Wireless Access](#)
- How to provide [Secure Campus Wired Access](#)
- How to provide [Remote VPN Access](#)

## Secure Campus Wireless Access

### Challenges to achieving secure access to the corporate campus network via a wireless (IEEE 802.11) connection using 802.1x authentication

Utilizing native operating system supplicants requires IT organizations to maintain and support a broad matrix of operating systems and 802.1x supplicants. Adding to the complexity, when an end point device physically or logically moves to other wireless networks such as airports, home networks, coffee shops, as well as other secure campus networks, additional network profiles are added to the end point device's configuration. This sprawl of network profiles on each device locally can directly affect the user experience by forcing users to know the network to which they want to connect. IT organizations are challenged with maintaining the integrity of the secure campus network profile for their organization's configuration which resides locally on the end point device for each user.

In addition to the many usability issues in utilizing native 802.1x supplicants in a secure campus environment, as previously stated, there are also severe security policy violations that can manifest.

With the proliferation of Personal Area Networks (PANs) and Wireless Personal Area Networks (WPANs), the need to have control of the end point's security policy becomes even more critical to the integrity of the IT organization's resources and intellectual property. An end point device may be authenticated and have access to the secure corporate network and at the same time have other devices tethered to it using different wired and wireless technologies. It is imperative to maintain the integrity of the data as it terminates on the end point device. In a second example, imagine a scenario where a user is physically connected via a wired connection in the secure campus network and simultaneously associated to a different wireless network. It could be argued that this end point device is now logically acting as a gateway or router as part of the secure campus network.

Over time, as these end point devices physically move and become part of a wide variety of disparate networks, all with unique security policies, a major problem manifests itself again from both a security and usability perspective. As the "scan list" within these 802.1x supplicants begins to grow, it is imperative for the IT organization to enforce and maintain a security policy for these end point devices. As mentioned earlier, the explosion and onset of PANs and WPANs increases the need to maintain control of the policies for these devices, which one could argue are becoming the edge of the network and, in some situations, these devices are becoming a forwarding engine on the network.

In many customer environments users are instructed to access a Web page/wiki or other resource in order to obtain the 802.1x settings to use (EAP-PEAP, EAP-TTLS, etc.). Granted, there is a solution for such settings for IT-owned assets as the IT organization has the ability to roll out pre-configured end point images, but as mentioned previously the pre-deployment of such images may not pertain to personal devices that users continue to bring into enterprise environments.

## Secure Campus Wired Access

### **Challenges to achieving secure network access to the corporate campus network via a wired (IEEE 802.3) connection using 802.1x authentication**

Gaining access to a secure campus network via an 802.1x wired connection has some challenges that are similar to those discussed in [Secure Campus Wireless Access](#), but it also has its own set of unique challenges. Again, using the native supplicants provided by Microsoft® can be a great operational challenge. For example, the native 802.1x wired supplicant for Windows XP is not enabled by default, thus the end user must enable this service manually. As discussed earlier, this type of configuration change on the client can be solved by the IT organization rolling out pre-deployed images to IT end point assets; however, this may not solve the problem for non-IT owned assets. As suggested earlier, the paradigm shift that is taking shape in the industry is forcing IT organizations to support a wide variety of consumer-based products and provide enterprise-class security.

Some of the same usability and security issues discussed with accessing the secure corporate campus network via native 802.1x supplicants also exist for accessing the secure corporate campus network via 802.1x native wired supplicants.

A security policy within an organization encompasses a broad spectrum of entities, ranging from enforcing help desk social conduct, logical security, management policies on all enterprise devices, physical security, etc. The need to secure access ports is crucial. For some organizations, providing port-level security for each access port becomes a compliance issue.

## Remote VPN Access

### **Challenges to achieving secure network access to the corporate campus network via a virtual private network (VPN) connection using both x.509 digital certificates and two-factor authentication**

Just as the lines are being blurred between personal and work end point devices, so are the lines between how users access the secure corporate campus network via wireless, wired, and remote locations.

Traditionally, users that were on a public network would need to understand that they are not on the corporate network, manually launch a VPN software client, choose the correct location to terminate the VPN connection (head-end) based on proximity, and then finally authenticate.

This should no longer be the case. In order for users to access a resource on the secure corporate campus, the user should not have to take preemptive measures to do so. The user should not have to know the exact network where the resources they are attempting to access reside. Instead, the client end point should dynamically detect that the user is attempting to access a resource on a remote campus network and build the appropriate secure connection.

## Challenges in Distributing Digital Certificates

Many IT organizations would like to provide identity and authentication services which include both digital certificates and a form of two-factor authentication for remote user access. However deploying client certificates to remote VPN end users can be extremely challenging. One of the challenges with deploying digital certificates to client end points accessing the campus network remotely is that the user and end point device may need to access the company's certification authority (CA) server directly (after being authenticated to the corporate campus network) in order to manually install the client certificate. This method requires the end user to manually install the client certificate while ensuring the certificate is installed in the proper certificate store on the local end point. As mentioned earlier, from a usability perspective, the user should not be required to know where on the end point device a digital certificate should be installed.

Deploying digital certificates on non-PC based devices can require a different process as many of these devices do not natively support all the features and functionality in order to create/download and install digital client certificates in the same manner as traditional PC-based devices.

Some end point devices do not support Simple Certificate Enrollment Protocol (SCEP) out of the box.

For example, in order for users to install digital client certificates using SCEP on Apple iOS devices, the IT administrator needs to manually create the configuration profile using the iPhone Configuration Utility and distribute the profile to user devices via E-mail, USB, or Web deployment. For additional information regarding the iPhone configuration utility, see the Apple Configuration Utility ([http://developer.apple.com/library/ios/#featuredarticles/FA\\_iPhone\\_Configuration\\_Utility/Introduction/Introduction.html](http://developer.apple.com/library/ios/#featuredarticles/FA_iPhone_Configuration_Utility/Introduction/Introduction.html)). The iPhone configuration utility can be downloaded from:

- iPhone Configuration Utility 3.3 for Mac OS X: <http://support.apple.com/kb/DL851>
- iPhone Configuration Utility 3.3 for Windows: <http://support.apple.com/kb/DL926>

Tradition full-featured PC-based devices are more apt to take advantage of the many services, such as Microsoft's NDES, to provide certificate enrollment. However, with the onset of so many non-PC-based mobile devices on the market today, it cannot be assumed that natively these devices can interoperate with many of the enterprise services currently deployed.

This solution provides guidance on enabling a unified access system in order to address the issue of achieving secure access to the corporate campus network via a wireless (IEEE 802.11) and wired (IEEE 802.3) connection using 802.1x authentication. This design also presents a migration from the existing Microsoft 802.1x native supplicants to a Cisco solution in order to solve many of the challenges facing a secure wireless and wired solution for campus and remote locations.

[Consistent Access Experience](#) provides design guidance on implementing a consistent network access experience on the campus or when remotely accessing the campus network.

This solution discusses the migration from the existing native 802.1x supplicants to the Cisco AnyConnect 3.0 Secure Mobility Client and describes the dynamic Web deployment of the Cisco AnyConnect 3.0 Secure Mobility including client profiles.

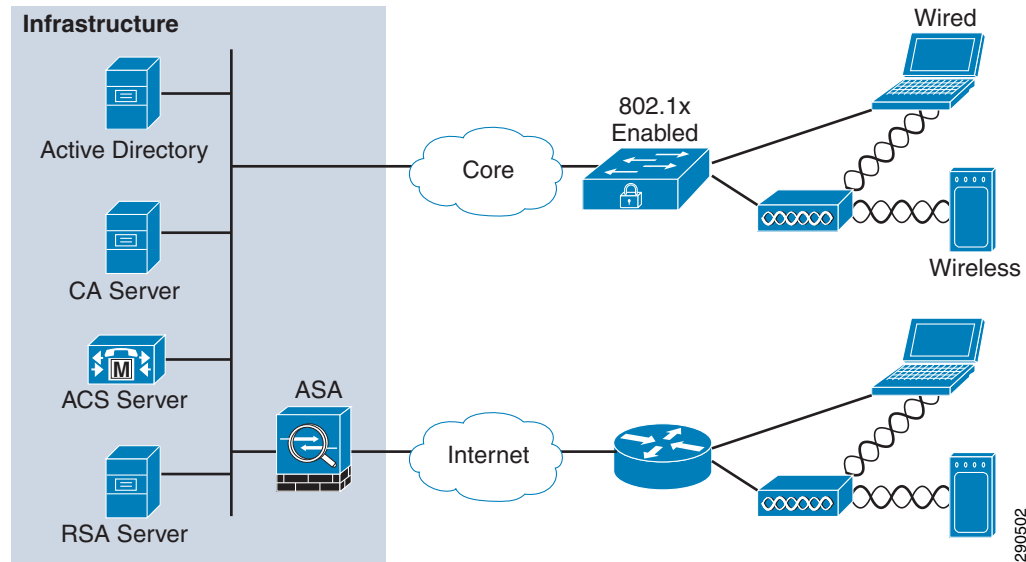
The solution discusses how the Cisco AnyConnect Secure Mobility Client, along with other integral products such as the Cisco Secure ACS, Cisco's next generation Layer 2 Access Switches, Cisco ASA 5500, and Cisco Access Points, enable end users and devices with secure and unified access to the corporate campus network using wired, wireless, and remote technologies.

## Unified Device Authentication

The main objectives of unified device authentication are:

- Enforce control on all the access layer switches that users use to connect to the network, including both wired and wireless users.
- Maintain centralized access control policy.

Traditionally wireless users are subjected to authentication when they access the network. Wired users connect and obtain their IP address without any kind of authentication. However, in certain organizations they are required by compliance to enforce access controls for wired users also. Unified Device Authentication helps organizations to authenticate both wired and wireless users using the same system and in the same way.

**Figure 3-1 Core Components of Unified Device Authentication**

This chapter focuses on properly designing and configuring the following items to enable an enterprise to utilize a Unified Device Authentication solution:

- Cisco AnyConnect Client
- Cisco Network Access Manager (NAM)
- Microsoft Active Directory (AD)
- RSA SecurID Authentication Server
- Microsoft Certificate Authority (CA)
- Cisco Adaptive Security Appliance (ASA)
- Cisco Secure Access Control System (ACS)

## Active Directory Environment

In the Unified Access Solution, Microsoft Active Directory (AD) has been implemented to authenticate user credentials. To obtain more information on implementing Microsoft AD services, see:

<http://technet.microsoft.com/en-us/library/cc754438%28WS.10%29.aspx>

The following services were configured in the Active Directory environment on the campus network:

- Enterprise CA Server with Network Device Enrollment Service—The internal Certificate Server running on the secure network segment of the campus network. The Network Device Enrollment (NDES) is enabled in order to provide clients with the ability to automatically enroll for client x.509 digital certificates.
- DHCP Server—The DHCP server on the secure campus network provides DHCP services for internal wired and wireless users.
- DNS Server—The DNS server provides internal name resolution for secure wired and wireless devices as well as providing name resolution for remote VPN users. As discussed later in this document, the AnyConnect client uses this internal name server in order to determine if a VPN connection should be established when using the Trusted Network Detection (TND), a feature within the AnyConnect client.

- IIS Server—The Web service is enabled so that users can use HTTP/HTTPS protocol to request certificates from a CA server via NDES as well as to provide internal users with a Web page link to download AnyConnect modules and profiles.

All the users in the campus network must be able to enroll their machines with the Active Directory Domain Controller. In our design, the only network device that is enrolled with the CA server is the Cisco Secure ACS Server.

## RSA Server

VPN security is only as strong as the methods used to authenticate users (and device end points) at the remote end of the VPN connection. Simple authentication methods based on static passwords are subject to password “cracking” attacks, eavesdropping, or even social engineering attacks. Two-factor authentication, which consists of something you know and something you have, is a minimum requirement for providing secure remote access to the corporate network ([http://www.cisco.com/web/about/security/intelligence/05\\_08\\_SSL-VPN-Security.html](http://www.cisco.com/web/about/security/intelligence/05_08_SSL-VPN-Security.html)). This design includes the RSA SecurID Authentication Server 7.1 along with RSA SecurID hardware tokens in order to provide two-factor authentication. The passcode that the user presents is a combination of their secret PIN and the one time password (OTP) code that is displayed on their token at that moment in time. The OTP on the RSA SecurID token changes every 60 seconds. This design utilizes both RSA SecurID (two-factor authentication) in conjunction with the deployment and use of x.509 client digital certificates.

For information on how to configure the RSA Secure Authentication Manager, see: <http://www.rsa.com/node.aspx?id=1166>

When a remote user establishes a VPN connection to the ASA via the AnyConnect Secure Mobility Client and attempts to authenticate using two-factor authentication with a RSA SecurID token, the ASA communicates with the ACS Server using the RADIUS protocol. The ACS Server then communicates with the RSA Authentication Manager using the Security Dynamics Incorporated (SDI) protocol in order to provide authentication. Since the ASA is communicating through the ACS server to the RSA Server, it is critical that the RSA “text message prompts” be properly translated by the ACS Server to the ASA and finally back to the end user. Also, because the SDI messages are configurable on the SDI server, the message text on the ASA must match (in whole or in part) the message text on the SDI server. Otherwise, the prompts displayed to the remote client user may not be appropriate for the action required during authentication.

## Certificate Authority Server

Identification is a critical part of an authentication strategy. Without a good authentication strategy, security does not provide any value to the network. Digital Certificates distributed by a CA provide the best form of identification for all devices on the network. Once users obtain their certificates, they can use them to identify themselves at different points in the network. Network devices such as the ASA, ACS, and wireless access points can also identify themselves using certificates. The best form of identification happens when both network users and network devices use certificates to identify themselves. Hence, we recommend that digital certificates be deployed at all places in the network. The Unified Access Solution is designed such that users and network devices authenticate each other using digital certificates.

Even though digital certificates provide many benefits to enterprise network deployments, the usual challenge has been in distributing these certificates to a large number of users. In the Unified Access Solution, this challenge has been addressed by providing two different methods of enrollment:

- Enrolling directly with the CA server using SCEP
- SCEP Proxy feature of ASA firewall

## Enrolling Using SCEP

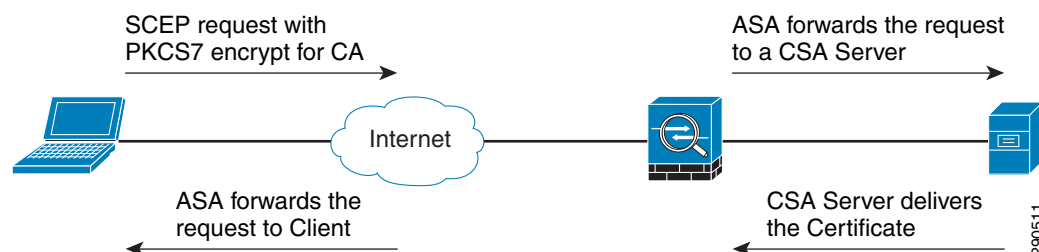
The user can use a Web browser to initiate certificate enrollment. The CA server can be set up either to issue certificates automatically or to be issued manually by an administrator. Granting certificates automatically is efficient for deployment, but the CA server must be set up with the right policies to verify user credentials before granting the certificate. If the granting mode is manual, then an administrator must issue certificates after verifying user credentials. CA server functionality is supported by several vendors. In the Unified Access Solution, the Microsoft CA server is deployed. To obtain more information about SCEP feature on the Microsoft CA server, see:

<http://www.microsoft.com/downloads/en/details.aspx?familyid=E11780DE-819F-40D7-8B8E-10845BC8D446&displaylang=en>

## Enrolling Using SCEP Proxy Feature of ASA Firewall

Even though SCEP is a great protocol for certificate enrollment, the major caveat is the requirement of a CA server be accessible to all clients. This requirement is fine with clients located at the campus site, but for clients that are outside the campus, such as remote clients, it is difficult to make the CA server accessible to them. In the Unified Access Solution, ASA's SCEP Proxy feature is implemented to address the connectivity problem for remote clients. This feature helps remote clients to make requests to the ASA, ASA forwards the requests to the CA server, and responds back to the clients with the certificates issued by the CA server. Figure 3-2 shows how the SCEP Proxy feature is deployed.

**Figure 3-2 SCEP Proxy Feature Deployment**



## NDES Server Configuration for SCEP

The Network Device Enrollment Service (NDES) is the Microsoft implementation of the SCEP, a communication protocol that makes it possible for network devices to enroll for X.509 certificates from a CA. In order to distribute and deploy digital x.509 client certificates to remote users, the Microsoft Network Device Enrollment Service (NDES) was utilized in conjunction with a Microsoft CA Server. To implement NDES service, see:

<http://technet.microsoft.com/en-us/library/cc753784%28WS.10%29.aspx>

By default, the NDES service is configured to present one-time enrollment passwords for certificate enrollment. The use of one-time passwords by the NDES service is typically used to allow network and IT administrators to enroll certificates for network devices within the IT organization. However, in the Unified Access Solution this feature is disabled because remote endpoints are authenticated by using their RSA SecureID tokens.

Disabling the “one-time password” on the NDES server is configured in the following registry key:  
Computer\HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\Cryptography\MSCEP\EnforcePassword

EnforcePassword value data is set to “0”. “0” ensures no password is requested by NDES.



#### Note

In Windows Server 2003, Microsoft SCEP (MSCEP) was a Windows Server 2003 Resource Kit add-on that had to be installed on the same computer as the CA. In Windows Server 2008, MSCEP support has been renamed NDES and is part of the operating system; NDES can be installed on a different computer than the CA (<http://technet.microsoft.com/en-us/library/cc753784%28WS.10%29.aspx>).

The NDES extension to IIS uses the registry to store configuration settings. All settings are stored under one registry key:

HKEY\_LOCAL\_MACHINE\Software\Microsoft\Cryptography\MSCEP

## Certificate Template

Digital certificates can be used for different purposes like server authentication, secure E-mail, encrypting the file system, and client authentication. Hence, it is important that a client is issued a certificate which serves its purpose. For example, a Web server may need a certificate that it uses primarily for server authentication. Similarly, a normal client needs a certificate mainly for client authentication. Therefore, certificate templates are needed to properly distribute certificates to users based on their specific needs. In the Unified Access Solution, a security template has been created on the Microsoft Windows 2008 CA server so that users can obtain the proper certificate. This section describes important steps to set up the certificate template on the Windows CA server and specific actions needed on the user.

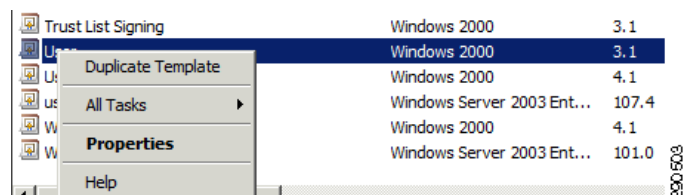
For more information on certificate templates, see:

<http://technet.microsoft.com/en-us/library/cc730826%28WS.10%29.aspx>

Since we are using SCEP for auto enrollment of client end points using the AnyConnect VPN Client, we are utilizing the properties of the “User” template that is a default template in the Microsoft Server 2008 R2 CA Server deployment. Default templates in Microsoft Server 2008 R2 cannot be edited. Therefore, a customized template can be built that gives an administrator more flexibility in defining the certificate options. This section describes how to create a customized template, which is named user2 in this example.

The first step is to create a duplicate template from the pre-defined list of templates. [Figure 3-3](#) shows how to create a duplicate template.

**Figure 3-3** Creating a Duplicate Template



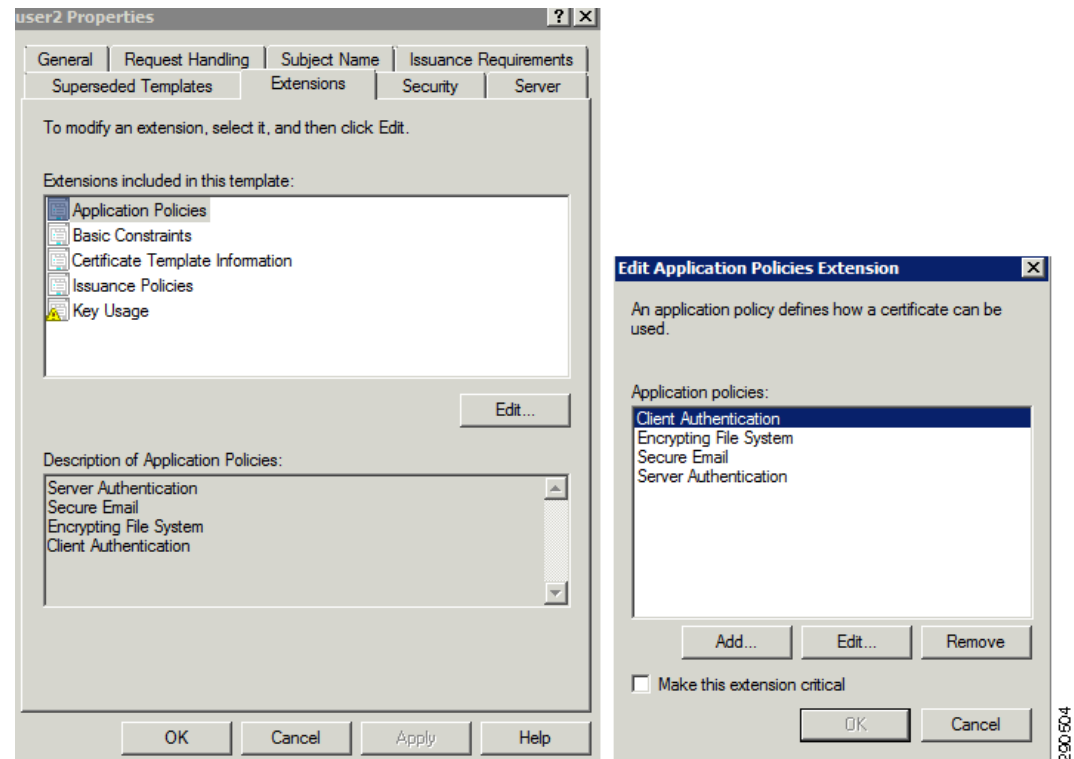
The default “User” template was copied and renamed as “user2”. Then the “user2” template was used to auto-enroll AnyConnect VPN clients with client certificates using this newly created template.

The next step is to configure the extensions of the certificates that are derived from the “user2” template. The EKU extension and extended property specify and limit the valid uses of a certificate. The extensions are part of the certificate itself. They are set by the issuer of the certificate and are read-only. Certificate-extended properties are values associated with a certificate that can be set in an application. To obtain more information about extended properties, see:

<http://msdn.microsoft.com/en-us/library/aa380252%28v=vs.85%29.aspx>

Figure 3-4 describes how to configure the extended properties for the certificates.

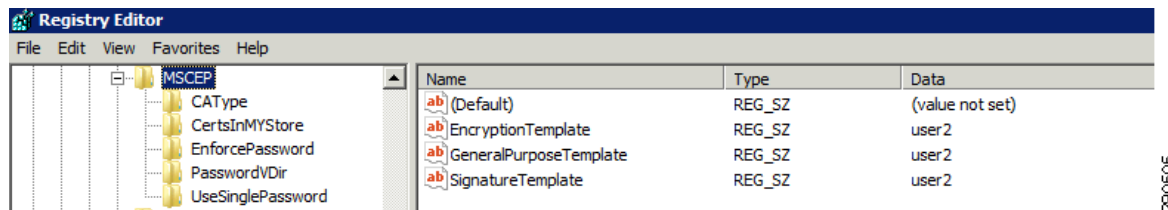
**Figure 3-4** *Configuring Extended Properties for Certificates*



Notice the template named “user2”. This value must be set in the registry as it correlates to the “user2” template, which was copied from the “User” template in the “Certificate Templates Console” on the CA Server.

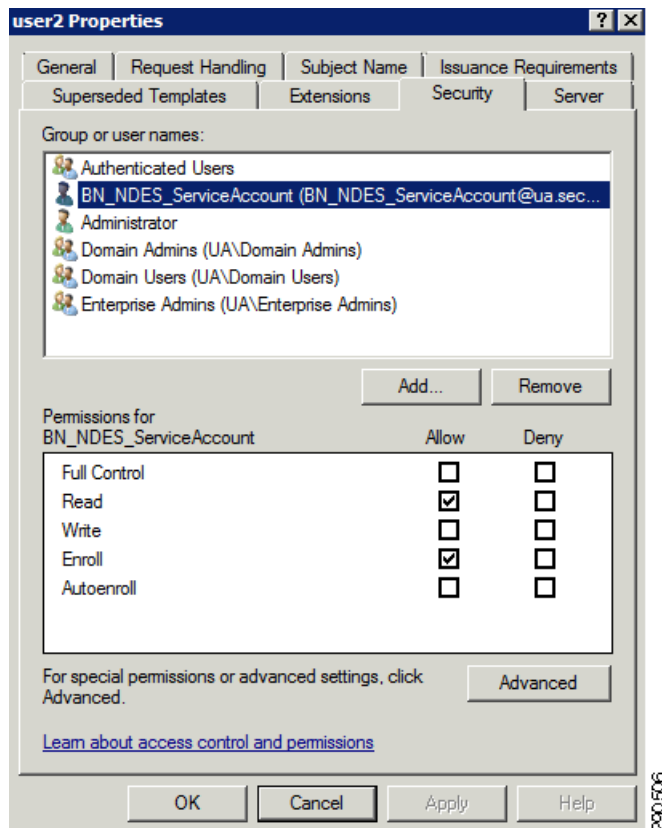
Figure 3-5 describes how the registry setting must be modified to reflect the newly-created template “user2”.

**Figure 3-5** *Modifying the Registry*

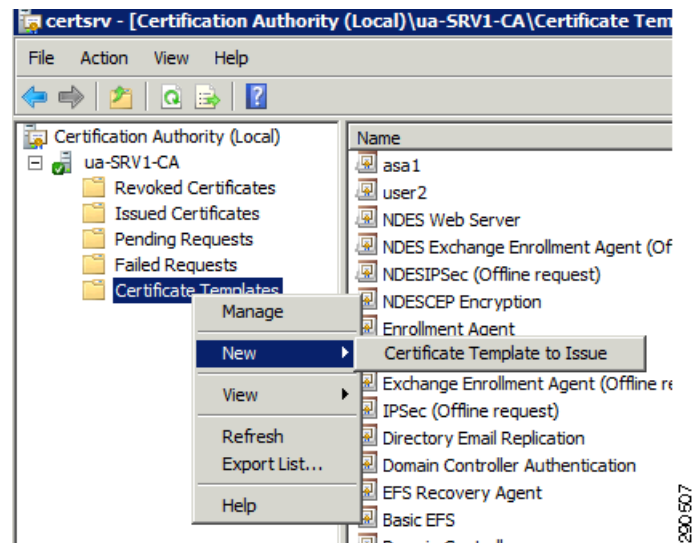


Once the template has been duplicated, the permissions are set for the NDES\_ServiceAccount on the “user2” template to Read and Enroll. Figure 3-6 displays the Read and Enroll permissions that have been set for the NDES\_ServiceAccount on the “user2” template.

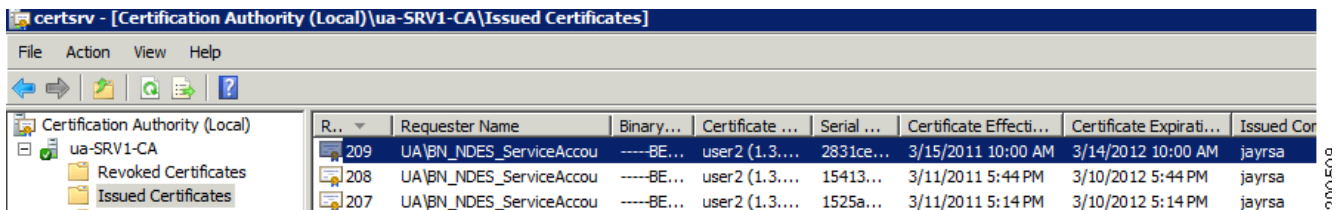
**Figure 3-6** Read and Enroll Permissions



Ensure that the newly created “user2” template is available to be issued via the CA. Right click <Certificate Templates> | <Certificate Template to Issue> and choose newly-created “User2 Certificate”.

**Figure 3-7 Ensuring Template is Available From CA**

Now the certificate template is fully configured and can be used by users to submit enrollment requests. Figure 3-8 shows a successful enrollment request to the “user2” template that was submitted by a user, “jayrsa”.

**Figure 3-8 Successful Enrollment Request**

A successful auto-enrollment request has occurred on the CA Server. Notice that the requester name is the NDES Service Account that is configured for Read and Enroll permissions and also notice that the “user2” certificate template was chosen.

Figure 3-9 shows how a user can submit a certificate enrollment request using a Web browser.

**Figure 3-9** Submitting a Certificate Enrollment Request with a Web Browser

Microsoft Active Directory Certificate Services - Windows Internet Explorer

http://192.168.1.102/certsrv/certqma.asp

Microsoft Active Directory Certificate Services

Microsoft Active Directory Certificate Services - ua-SRV1-CA

### Advanced Certificate Request

**Certificate Template:**

User

**Key Options:**

CSP: Web Server

Key Usage: IPSec (Offline request) CEP Encryption

Key Size: user2

Use existing key set

Geographic Provider v1.0

Automatic key container name

Mark keys as exportable

Enable strong private key protection

**Additional Options:**

Request Format: CMC

Hash Algorithm: sha1

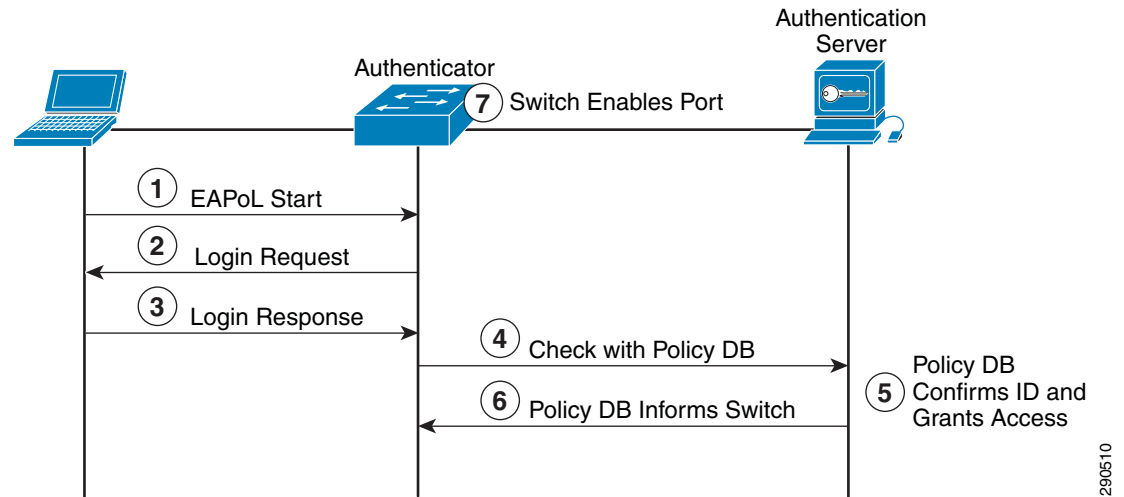
Only used to sign request.

Save request

## Authentication Protocols

The IEEE 802.1X protocol allows Cisco Catalyst switches to offer network access control at the port level. Every port on the switch is individually enabled or disabled based on the identity of the user or device connecting to it. When 802.1X is first enabled on a port, the switch automatically drops all traffic received on the port except the request to start 802.1X authentication. After the 802.1X authentication successfully completes, the switch starts accepting other kinds of traffic on the port.

The high-level message exchange shown in [Figure 3-10](#) illustrates how port-based access control works within an identity-based system.

**Figure 3-10 High-Level Message Exchange**

The following steps describe the port-based access control flow:

1. A client, such as a laptop with an 802.1X supplicant, connects to an IEEE 802.1X-enabled network and sends a start message to the LAN switch (the authenticator).
2. When the start message is received, the LAN switch sends a login request to the client.
3. The client replies with a login response.
4. The switch forwards the response to the policy database (authentication server).
5. The authentication server authenticates the user.
6. After the user's identity is confirmed, the policy database authorizes network access for the user and informs the LAN switch.
7. The LAN switch then enables the port connected to the client.

The user or device credentials are processed by a AAA server. The AAA server is able to reference user or device profile information either internally, using the integrated user database, or externally using database sources like Microsoft Active Directory, Lightweight Directory Access Protocol (LDAP), Novell Directory Services, or Oracle databases. This enables the IBNS solution to be integrated into existing user management structures and schemes, which simplifies overall deployment.

## 802.1X and EAP

When authenticating users for network access, the client must provide user and/or device identification using strong authentication technologies. IEEE 802.1X does not dictate how this is achieved. Instead, the 802.1X protocol defines an encapsulation for the transport of the Extensible Authentication Protocol (EAP) from the client to the switch. The 802.1X encapsulation is sometimes referred to as EAP over LAN (EAPoL). The switch in turn relays the EAP information to the authentication server using the RADIUS protocol (EAP over RADIUS).

RFC 3748 defines EAP, which is a framework and not a specific authentication method. It provides a way for the client and the authentication server to negotiate an authentication method that they both support. There are many EAP methods, but the ones used more frequently for 802.1X wired authentication include EAP-TLS, EAP-PEAP, and EAP-FAST.

For the Unified Access Solution, we have chosen EAP-TLS and PEAP as the authentication protocols. We highly recommend using the EAP-TLS protocol for customers for the following reasons:

- Authentication using certificate provides the highest level of trust among all the methods available today.
- All the other protocols except EAP-TLS do not provide mutual authentication—only the access layer switch presents its identity, but not the client. EAP-TLS allows for mutual authentication—both the access layer switch and the user authenticate each other.
- When EAP-TLS is deployed and the user obtains the certificate, then the same certificate can be presented as a source of identity for remote connections also, which is an added benefit of deploying EAP-TLS.

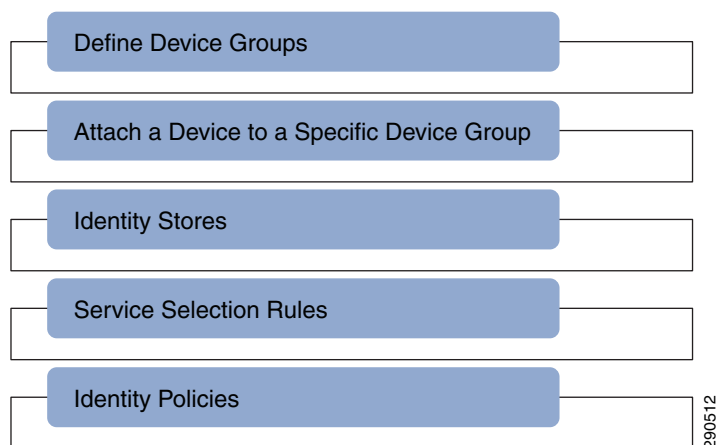
## ACS Policy Definitions

Configuring ACS is an important step in developing a centralized access control system. Cisco Secure ACS is used in the solution such that all the authentication requests come to the ACS server; the ACS server re-directs the requests to various additional authentication servers such as Active Directory and the RSA server. All the authentication requests coming from various sources, such as access layer switches, access points, and remote users, come to the ACS server first; based on the policy defined on the ACS server, they are further directed to either Active Directory or the RSA secure server. The following section describes how to deploy this centralized access control system.

## ACS Authentication Overview

In Cisco Secure ACS, when an authentication request comes from the client into the ACS server, the ACS server first tries to identify to which Network Device group the client belongs and, based on the Network Device Group, the request is directed to an appropriate Access service. In each Access service there are rules defined to further authenticate the request. [Figure 3-11](#) shows an overview of ACS authentication.

**Figure 3-11 ACS Authentication Overview**



## Device Groups

As shown in [Figure 3-11](#), the first step is to define the network device groups. In the Unified Access Solution, three different types of device groups are defined:

- WLC—All wireless LAN controllers

- Layer 2 access—All access layer switches
- AC VPN termination—VPN gateways

Figure 3-12 shows how these device groups are defined on ACS.

**Figure 3-12**      *Device Groups on ACS*



## Mapping Devices to Device Groups

The second step in ACS configuration is to map the individual devices to a specific group. For example, a WLAN Controller must belong to the device group WLC, which is shown in Figure 3-12. Figure 3-13 shows how different devices are mapped to different device groups.

**Figure 3-13**      *Mapping of Devices to Device Groups*

Name	IP / Mask	NDG:Location	NDG:Device Type
<input type="checkbox"/> cr22-2960s	10.125.130.2/32	All Locations	All Device Types:Layer 2 Access
<input type="checkbox"/> cr22-3560x	10.125.130.130/32	All Locations	All Device Types:Layer 2 Access
<input type="checkbox"/> cr22-3750x	10.125.131.2/32	All Locations	All Device Types:Layer 2 Access
<input type="checkbox"/> cr22-4500	10.125.131.130/32	All Locations	All Device Types:Layer 2 Access
<input type="checkbox"/> cr22-ap1	10.125.130.3/32	All Locations	All Device Types:WLC
<input type="checkbox"/> cr24-asa1	192.168.1.100/32	All Locations	All Device Types:AC VPN Termination

## Identity Stores

The ACS server passes the identification requests to different identity stores to identify the users. The identity stores are external authentication devices that can perform the authentication of users. In the Unified Access Solution, there are three main identity stores:

- Active Directory
- RSA secure server
- Certificate Profile

To use Active Directory as an identity store, the ACS must first join the active directory domain. Figure 3-14 shows how ACS should be configured to join the domain.

**Figure 3-14 ACS Configuration for Active Directory**

Users and Identity Stores > External Identity Stores > Active Directory

**General** | Directory Groups | Directory Attributes

**Connection Details**

✱ Active Directory Domain Name:

Please specify the credentials used to join this machine to the Active Directory Domain:

✱ Username:

✱ Password:

You may use the Test Connection Button to ensure credentials are correct and Active Directory Domain is reachable.

Click on 'Save Changes' to connect to the Active Directory Domain and save this configuration. Once you have success Groups and Directory Attributes to be available for use in policy rules.

**End User Authentication Settings**

☒ Enable password change

☒ Enable machine authentication

☐ Enable Machine Access

**Restrictions**

Aging time (hours):

**Connectivity Status**

Joined to Domain: ua.secbn.com

Connectivity Status: CONNECTED

✱ = Required fields

5/15/2015 2:00:15

The RSA secure server is used to authenticate remote users. When remote users connect to the ASA to establish an AnyConnect VPN session, they are required to enter a one-time password using RSA tokens. ACS must be configured to join the RSA server, as shown in Figure 3-15.

**Figure 3-15**      **ACS Configuration for RSA**


Users and Identity Stores > External Identity Stores > [RSA SecurID Token Servers](#) > Edit: "rsa1"

RSA Realm

ACS Instance Settings


Advanced

### General

 Name:

Description:

### Server connection

 Server Timeout:  Seconds

☐ Reauthenticate on Change PIN

### Realm Configuration File


The RSA Configuration file (sdconf.rec) should be provided by your RSA administrator after they have registered all the ACS Instances

Current File: Download...

Timestamp: 12:05 04.01.2011

File Size: 1024 bytes

Import new 'sdconf.rec' file:  Browse...

 = Required fields

When the RSA Manager communicates with Cisco Secure ACS, it uses UDP port 5500. [Figure 3-16](#) shows the communication between the RSA Manager and ACS during authentication.

The trace in [Figure 3-16](#) is taken on the RSA Authentication Manager (the IP address 192.168.1.101 is the ACS Server and 192.168.1.103 is the RSA server). Note that the ACS Server and the RSA Server are communicating on UDP:5500.

**Figure 3-16** *ACS/RSA Manager Communication During Authentication*

[illegible]

Configure the appropriate SDI messages in the AAA Server Group for the AnyConnect Clients on the ASA, as shown in [Figure 3-17](#).

**Figure 3-17** *Configure SDI Messages*

**Edit AAA Server**

Server Group: ACS

Interface Name: inside

Server Name or IP Address: 192.168.1.101

Timeout: 10 seconds

**RADIUS Parameters**

Server Authentication Port: 1645

Server Accounting Port: 1646

Retry Interval: 10 seconds

Server Secret Key: •••••

Common Password:

ACL Netmask Convert: Standard

Microsoft CHAPv2 Capable: ☒

**SDI Messages**

Message Name	Message Text
ready-for-sys-pin	ACCEPT A SYSTEM GENERATED PIN
new-pin-reenter	Reenter PIN:
new-pin-meth	Do you want to enter your own pin
next-ccode-and-rea...	new PIN with the next card code
next-code	Enter Next PASSCODE
new-pin-sys-ok	New PIN Accepted
new-pin-sup	Please remember your new PIN
new-pin-req	Enter your new Alpha-Numerical PIN

(Double-click in a text cell to make changes.)

Restore default message texts

Help Cancel OK

The Cert Authentication Profile feature allows ACS to match the certificates presented by the user to the Active Directory. [Figure 3-18](#) shows how to configure it on this ACS server.

**Figure 3-18** Configuring Cert Authentication Profile

Users and Identity Stores > Certificate Authentication Profile > Edit: "Cert-Auth-AD"

**General**

Name: Cert-Auth-AD

Description:

**Certificate Definition**

Principal Username X509 Attribute: Common Name

☒ Perform Binary Certificate Comparison with Certificate retrieved from LDAP or Active Directory

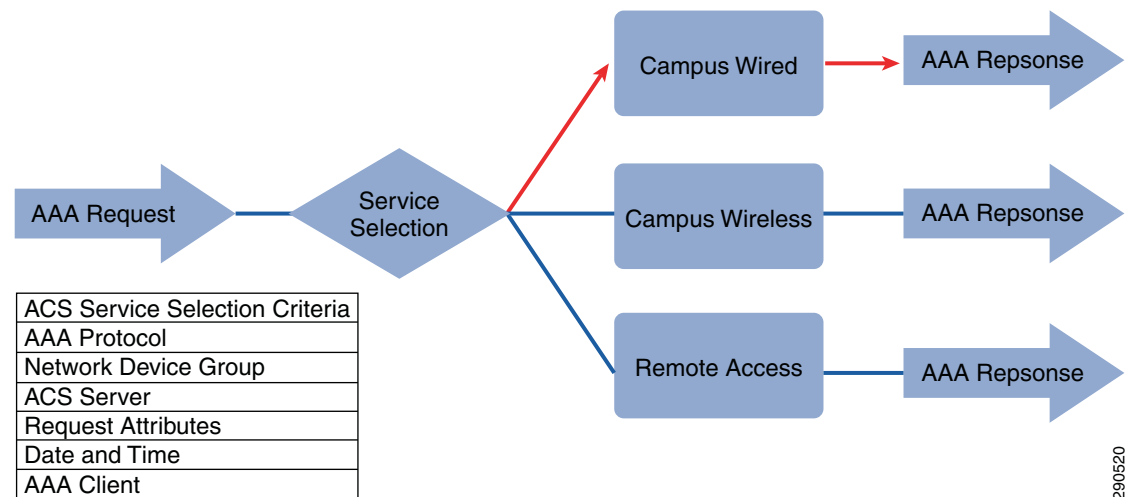
Directory Name: AD1 Select

\* = Required fields

290519

## Service Selection Policy

The Service Selection Policy acts like a gateway that forwards authentication requests to the various identification policies. In the Unified Access Solution, there are three different service selection policy elements—wired access policy, wireless access policy, and remote access policy. Figure 3-19 shows how the service selection policy works in the Unified Access Solution.

**Figure 3-19** Service Selection Policy in UA Solution

290520

Figure 3-20 shows how the Service Selection Policy looks in Cisco ACS server.

**Figure 3-20** Service Selection Policy Cisco ACS Server

3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<a href="#">Rule-3</a>	match Radius	in All Device Types:WLC	Campus_Wireless	217
4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<a href="#">Rule-4</a>	-ANY-	in All Device Types:Layer 2 Access	Campus_Wired	359
5	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<a href="#">Rule-5</a>	-ANY-	in All Device Types:AC VPN Termination	Remote Access	728

290521

As shown in Figure 3-20, all the devices that belong to device type group WLC match the individual policy element Campus\_Wireless. Similarly, all the devices that belong to Layer 2 access match Campus\_Wired and all the devices that belong to the group VPN Termination match the group Remote\_Access.

## Individual Policy Elements

The service selection policy directs the authentication request to each individual policy. In the Unified Access Solution, the individual policies are Campus Wired, Campus Wireless, and Remote Access. These individual policies must be constructed carefully because the default behavior of an access policy is to reject an authentication if there is no match in the policy. In the Unified Access Solution, there are three fields that are matched for an authentication request—EAP Authentication Method, Authentication Method, and EAP Tunneling Protocol. Figure 3-21 shows how the Campus Wired Policy is constructed.

**Figure 3-21** Campus Wired Policy

Access Policies > Access Services > Campus_Wired > Identity								
<input type="radio"/> Single result selection <input checked="" type="radio"/> Rule based result selection								
Identity Policy								
Filter: Status Match If: Equals Enabled Clear Filter Go								
		Status	Name	Conditions			Results	Hit Count
				Eap Authentication Method	Authentication Method	Eap Tunnel Building Method	Identity Source	
1	<input type="checkbox"/>	●	<a href="#">Rule-1</a>	-ANY-	-ANY-	match PEAP	AD1	79
2	<input type="checkbox"/>	●	<a href="#">Rule-2</a>	-ANY-	match x509_PKI	-ANY-	CN Username	20

As shown in Figure 3-21, the selection policy has two different rules. The first rule is used to match PEAP with Active Directory as an identity store and the second rule matches EAP\_TLS with digital certificates as an authentication method. The main point to note above is that PEAP is matched on the Tunneling method rather than on the EAP Authentication method because the PEAP protocol works differently than the EAP\_TLS protocol. When PEAP is used, the first step is to establish a TLS tunnel and actual credentials are sent in the second stage, which is protected by the TLS tunnel that has been established. Therefore PEAP has an outer method and an inner method. The outer method is PEAP and the inner-method authentication type that is negotiated during phase two can be either EAP-MSCHAPv2 or EAP-GTC. The combination of the outer PEAP method with a specific inner EAP method is denoted using a slash character; for example, PEAP/EAP-MSCHAPv2 or PEAP/EAP-GT. The following link gives more information on deploying PEAP in ACS 5.2:

[http://www.cisco.com/en/US/docs/net\\_mgmt/cisco\\_secure\\_access\\_control\\_system/5.2/user/guide/eap\\_pap\\_phase.html#wp1031414](http://www.cisco.com/en/US/docs/net_mgmt/cisco_secure_access_control_system/5.2/user/guide/eap_pap_phase.html#wp1031414)

The second policy is Campus Wireless. The Unified Access Solution assumes that wireless users can use either EAP-TLS or PEAP protocol. The EAP\_TLS protocol uses digital certificates as identity store and the PEAP protocol uses Active Directory as its identity store. Figure 3-22 shows how this policy is configured in ACS.

**Figure 3-22 Campus Wireless Policy**

Access Policies > Access Services > Campus\_Wireless > Identity

☐ Single result selection ☒ Rule based result selection

**Identity Policy**

Filter: Status Match if: Equals Enabled Clear Filter Go

	<input type="checkbox"/>	Status	Name	Eap Authentication Method	Authentication Method	Eap Tunnel Building Method	Results Identity Source	Hit Count
1	<input type="checkbox"/>	●	<a href="#">Rule-1</a>	match EAP-TLS	match x509_PKI	-ANY-	CN Username	65
2	<input type="checkbox"/>	●	<a href="#">Rule-2</a>	-ANY-	-ANY-	match PEAP	AD1	73

The last policy is Remote Access. In the Unified Access Solution, remote users use two-factor authentication along with client digital certificates. The initial authentication happens between the ASA and the remote user. After the initial authentication is successful, the ASA prompts the user for an RSA code. Once the user inputs the code, the ASA passes the code to the ACS server. The ACS server in-turn passes the user response to the RSA server. Once the RSA server successfully authenticates the code, the user is allowed to access the network. Figure 3-23 depicts the remote access policy configuration in ACS.

**Figure 3-23 Remote Access Policy**

Access Policies > Access Services > Remote Access > Identity

☐ Single result selection ☒ Rule based result selection

**Identity Policy**

Filter: Status Match if: Equals Enabled Clear Filter Go

	<input type="checkbox"/>	Status	Name	Compound Condition	NDG:Device Type	Authentication Method	Results Identity Source	Hit Count
1	<input type="checkbox"/>	●	<a href="#">Rule-2</a>	-ANY-	in All Device Types:AC VPN Termination	match PAP_ASCII	rsa1	365

## ASA Configuration

Some challenges faced by network architects when deploying a remote VPN solution for their network include:

- How do remote users trust the VPN gateway?
- How does the VPN gateway identify remote users?
- How to organize different types of users in groups so that different kinds of services can be provided?
- What kind of mobility client solution is needed for a particular client?
- Once the right kind of VPN solution is identified, how will the mobility client be installed on the remote device?
- How to centralize the policy settings for VPN users? It is always very annoying for remote users to configure a mobile device for VPN functionality.

The Cisco ASA coupled with the Cisco AnyConnect client addresses the above challenges. The Cisco AnyConnect client 3.0 is used to meet the need of wired, wireless, and remote users. The Cisco AnyConnect Secure Mobility client is the next-generation VPN client, providing remote users with

secure IPsec (IKEv2) or SSL VPN connections to the Cisco 5500 Series Adaptive Security Appliance (ASA). AnyConnect provides end users with a connectivity experience that is intelligent, seamless, and always-on, with secure mobility across today's proliferating managed and unmanaged mobile devices.

The Cisco AnyConnect Secure Mobility client, Version 3.0, integrates new modules into the AnyConnect client package:

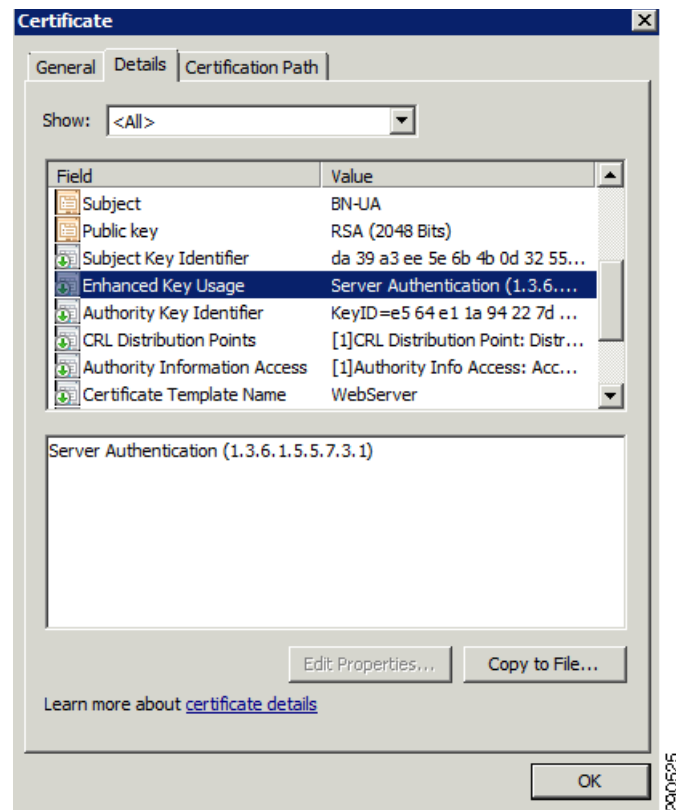
- **Network Access Manager (NAM)**—Formerly called the Cisco Secure Services Client, this module provides Layer 2 device management and authentication for access to both wired and wireless networks.
- **Posture Assessment**—The AnyConnect Posture Module provides the AnyConnect Secure Mobility Client with the ability to identify the operating system, antivirus, antispymware, and firewall software installed on the host prior to creating a remote access connection to the ASA. Based on this prelogin evaluation, you can control which hosts are allowed to create a remote access connection to the security appliance. The Host Scan application is delivered with the posture module and is the application that gathers this information.
- **Telemetry**—Sends information about the origin of malicious content detected by the antivirus software to the Web filtering infrastructure of the Cisco IronPort Web Security Appliance (WSA), which uses this data to provide better URL filtering rules.
- **Web Security**—Routes HTTP and HTTPS traffic to the ScanSafe Web Security scanning proxy server for content analysis, detection of malware, and administration of acceptable use policies.
- **Diagnostic and Reporting Tool (DART)**—Captures a snapshot of system logs and other diagnostic information and creates a .zip file on your desktop so you can conveniently send troubleshooting information to Cisco TAC.
- **Start Before Logon (SBL)**—Forces the user to connect to the enterprise infrastructure over a VPN connection before logging on to Windows by starting AnyConnect before the Windows login dialog box appears.

To obtain more information about Cisco AnyConnect 3.0, see:

[http://www.cisco.com/en/US/docs/security/vpn\\_client/anyconnect/anyconnect30/administration/guide/ac01intro.html](http://www.cisco.com/en/US/docs/security/vpn_client/anyconnect/anyconnect30/administration/guide/ac01intro.html)

## Remote Users Trusting VPN Gateway

The ASA can present a self-signed certificate, but presenting a certificate that is signed by a CA server increases the trust of the ASA. However, the certificate that the ASA presents must be a certificate that can be used for server authentication. In our testing scenario, we have used a Microsoft CA with the certificate template configured for Web Server. [Figure 3-24](#) shows the certificate that can be used for server authentication.

**Figure 3-24** Certificate for Server Authentication

ASA can obtain the certificate from the CA server by using SCEP or by a manual cut-and-paste method. SCEP was used to obtain the certificate in this testing. To obtain more information on deploying certificates on the ASA, see:

[http://www.cisco.com/en/US/docs/security/asa/asa80/configuration/guide/cert\\_cfg.html](http://www.cisco.com/en/US/docs/security/asa/asa80/configuration/guide/cert_cfg.html)

Example 3-1 is the ASA configuration for certificate enrollment.

#### **Example 3-1 ASA Configuration for Certificate Enrollment**

```
crypto ca trustpoint f5
enrollment url http://192.168.1.102:80/certsrv/mscep/mscep.dll
subject-name CN=asa1
no client-types
crl configure
```

## Identification of Remote Users

Identification is very important, particularly for remote users. Digital certificates play an important part in identifying remote users, however it is often difficult to deploy the certificates for remote users because they cannot access the CA server directly through the Internet. They can access the CA server through a VPN connection, but without having a digital certificate in first place, they cannot establish a VPN tunnel. To solve this problem the ASA supports the SCEP proxy feature.

SCEP enables network devices to enroll for x509 version 3 certificates from a CA. The ASA can proxy SCEP requests between AnyConnect and a third-party CA. The ASA supports SCEP-Proxy for AnyConnect clients.

The CA only needs to be accessible to the ASA, as the ASA is acting as the proxy for the client. For the ASA to provide this service, the user must authenticate using any of the methods supported by AAA before the ASA sends an enrollment request. Enrollment occurs only after a VPN tunnel has been established between the AnyConnect client and the ASA. The AnyConnect client enrolls to the ASA Client Services interface using SSL. After the SSL VPN tunnel is established, the ASA proxies the SCEP enrollment to the CA on behalf of the client.

The VPN tunnel is disconnected after an enrollment success or failure. If the enrollment was a success, the AnyConnect client re-launches, allowing the user to authenticate with their credentials while using the newly-enrolled client certificate. ASA supports this feature only with an AnyConnect SSL or IKEv2 VPN session. The ASA supports all SCEP-compliant CAs, including IOS CS, Windows Server 2003 CA, and Windows Server 2008 CA.

## Configuring SCEP for AnyConnect Clients on the ASA

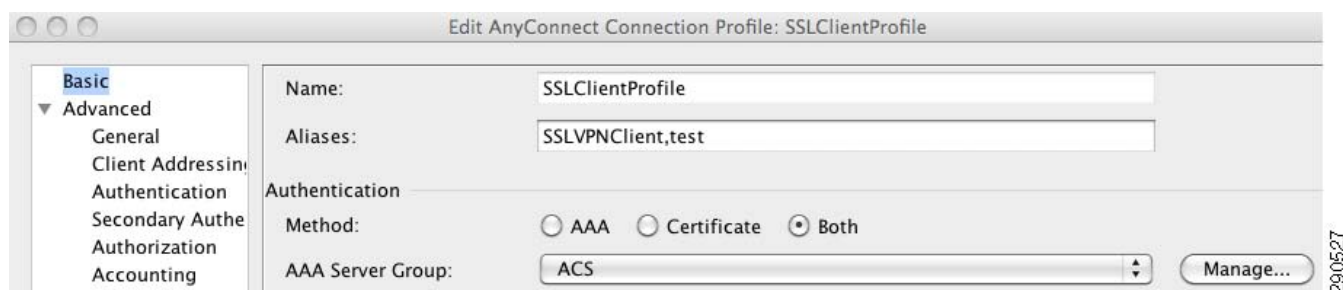
Enable SCEP for the Connection Profile:

**Figure 3-25** Enable SCEP



Configure the Connection Profile for authentication (configure the Authentication Method for “Both” Active Directory username/password and PKI x.509 Certificates):

**Figure 3-26** Configuring the Connection Profile



Configure the Group Policy for Network Client Access.

The SCEP forwarding URL points to the SCEP/NDES resource service.

**Figure 3-27**     **Configuring the Group Policy**

Edit Internal Group Policy: SSLClientPolicy		
Name:		SSLClientPolicy
Banner:	<input checked="" type="checkbox"/> Inherit	
SCEP forwarding URL:	<input type="checkbox"/> Inherit	http://192.168.1.102/certsrv/mscep/mscep.dll
Address Pools:	<input type="checkbox"/> Inherit	testpool1
IPv6 Address Pools:	<input checked="" type="checkbox"/> Inherit	

Configure the Client Profile with the Automatic SCEP Host which is the IP Address of Hostname of the ASA for which the ASA will proxy SCEP requests on behalf of the client. 172.16.2.1 is the outside IP address of the ASA.

The prompt for password is left blank as the purpose of our SCEP enrollment is for client end points performing auto-enrollment and not requiring the use of a password for each end point enrolling for a certificate.

Name(CN) is %USER%, which is used to create the enrollment with this CN.

**Figure 3-28** Configuring the Client Profile

The screenshot shows the 'AnyConnect Client Profile Editor - vpn1' window. The left sidebar lists configuration categories: VPN, Preferences (Part 1), Preferences (Part 2), Backup Servers, Certificate Matching, Certificate Enrollment (selected), Mobile Policy, and Server List. The main area is titled 'Certificate Enrollment' and contains the following fields and options:

- ☒ Certificate Enrollment
- Certificate Expiration Threshold (days): 30
- Automatic SCEP Host: 172.16.2.1
- CA URL: http://192.168.1.102/certsr
- ☐ Prompt For Challenge Password
- CA Thumbprint: (empty field)
- Certificate Contents:
 

Name (CN): %USER%	Qualifier (GEN):
Department (OU):	Qualifier (DN):
Company (O):	City (L):
State (ST):	Title (T):
State (SP):	CA Domain: ua.secbn.com
Country (C):	Key Size:
Email (EA):	<input type="checkbox"/> Display Get Certificate Button
Domain (DC): ua.secbn.com	
SurName (SN):	
GivenName (GN):	

At the bottom are buttons for Help, Cancel, and OK. A vertical text '290529' is visible on the right edge of the window.

## Creating Groups for Different Types of Users

Group Policy is an important building block for designing an effective access mechanism for users. The needs of specific users can differ. For example, one user might like to have a domain value of xyz.com and have 1.1.1.1 and 2.2.2.2 as their DNS servers. Another user might have similar requirements, but in addition might need a proxy server configured for his or her user name. If you have to attach all these attributes to each individual user, the configuration might become very large and complex. To solve this problem, multiple groups are created, each with its set of individual attributes. In this case you can simply associate a user with a group name, rather than the large number of attributes, thus minimizing the configuration complexity when you have multiple users.

By default, the Cisco ASA creates DftGrpPolicy and the other group policies that inherit most of the common attributes; only very specific attributes need to be configured explicitly for each group.

For more information about configuring tunnel groups, group policies, and users, see:  
<http://www.cisco.com/en/US/docs/security/asa/asa80/configuration/guide/vpnggrp.html>

In our Group policy definition the main attributes you need are vpn-tunnel protocol, split\_tunnel\_acl, and address pool location. [Example 3-2](#) shows how this group policy was defined.

**Example 3-2 Group Policy on the ASA**

```
group-policy SSLClientPolicy internal
group-policy SSLClientPolicy attributes
  banner value Welcome to the SASU Group!
  vpn-tunnel-protocol ssl-client ssl-clientless
  split-tunnel-policy tunnelspecified
  split-tunnel-network-list value split_Acl
  default-domain none
  user-authentication disable
  address-pools value testpool1
  scep-forwarding-url value http://192.168.1.9/certsrv/mscep/mscep.dll
```

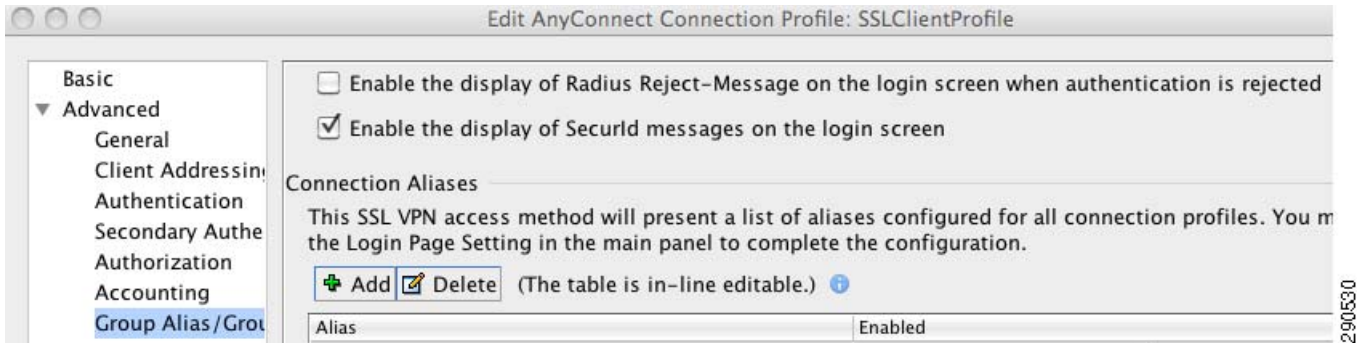
## Connection Profile Configuration

Group policies define the attributes for a group, but the connection profile specifies the attributes specific to a connection. For example, a connection profile for AnyConnect will specify if the users belonging to this connection are authenticated by a radius server or locally. The connection profile also points to the group profile to which it belongs. If no connection profile is defined on the system, the ASA points to a default connection profile, but to make administration simple it is better to define a specific group and connection profiles. [Example 3-3](#) is the configuration on ASA for Any Connect connection profile.

**Example 3-3 AnyConnect Connection Profile**

```
tunnel-group SSLClientProfile type remote-access
tunnel-group SSLClientProfile general-attributes
  address-pool testpool
  authentication-server-group ACS
  default-group-policy SSLClientPolicy
  scep-enrollment enable
tunnel-group SSLClientProfile webvpn-attributes
  authentication aaa certificate
  proxy-auth sdi
  group-alias RSA enable
  group-alias SASU enable
  group-alias SSLVPNClient enable
  group-alias test enable
!
```

Configure the AnyConnect Connection Profile to Enable the display of SecureID Messages on the login screen of the AnyConnect Client.

**Figure 3-29** *Configuring the AnyConnect Connection Profile*

Configure the appropriate SDI messages in the AAA Server Group for the AnyConnect Clients.

**Figure 3-30**      **Configuring SDI Messages**

The screenshot shows the 'Edit AAA Server' configuration window. The 'Server Group' is set to 'ACS'. The 'Interface Name' is 'inside'. The 'Server Name or IP Address' is '192.168.1.101'. The 'Timeout' is '10' seconds. Under 'RADIUS Parameters', the 'Server Authentication Port' is '1645', the 'Server Accounting Port' is '1646', the 'Retry Interval' is '10 seconds', the 'Server Secret Key' is masked with dots, the 'Common Password' is empty, the 'ACL Netmask Convert' is 'Standard', and 'Microsoft CHAPv2 Capable' is checked. The 'SDI Messages' section contains a 'Message Table' with the following entries:

Message Name	Message Text
ready-for-sys-pin	ACCEPT A SYSTEM GENERATED PIN
new-pin-reenter	Reenter PIN:
new-pin-meth	Do you want to enter your own pin
next-ccode-and-rea...	new PIN with the next card code
next-code	Enter Next PASSCODE
new-pin-sys-ok	New PIN Accepted
new-pin-sup	Please remember your new PIN
new-pin-req	Enter your new Alpha-Numerical PIN

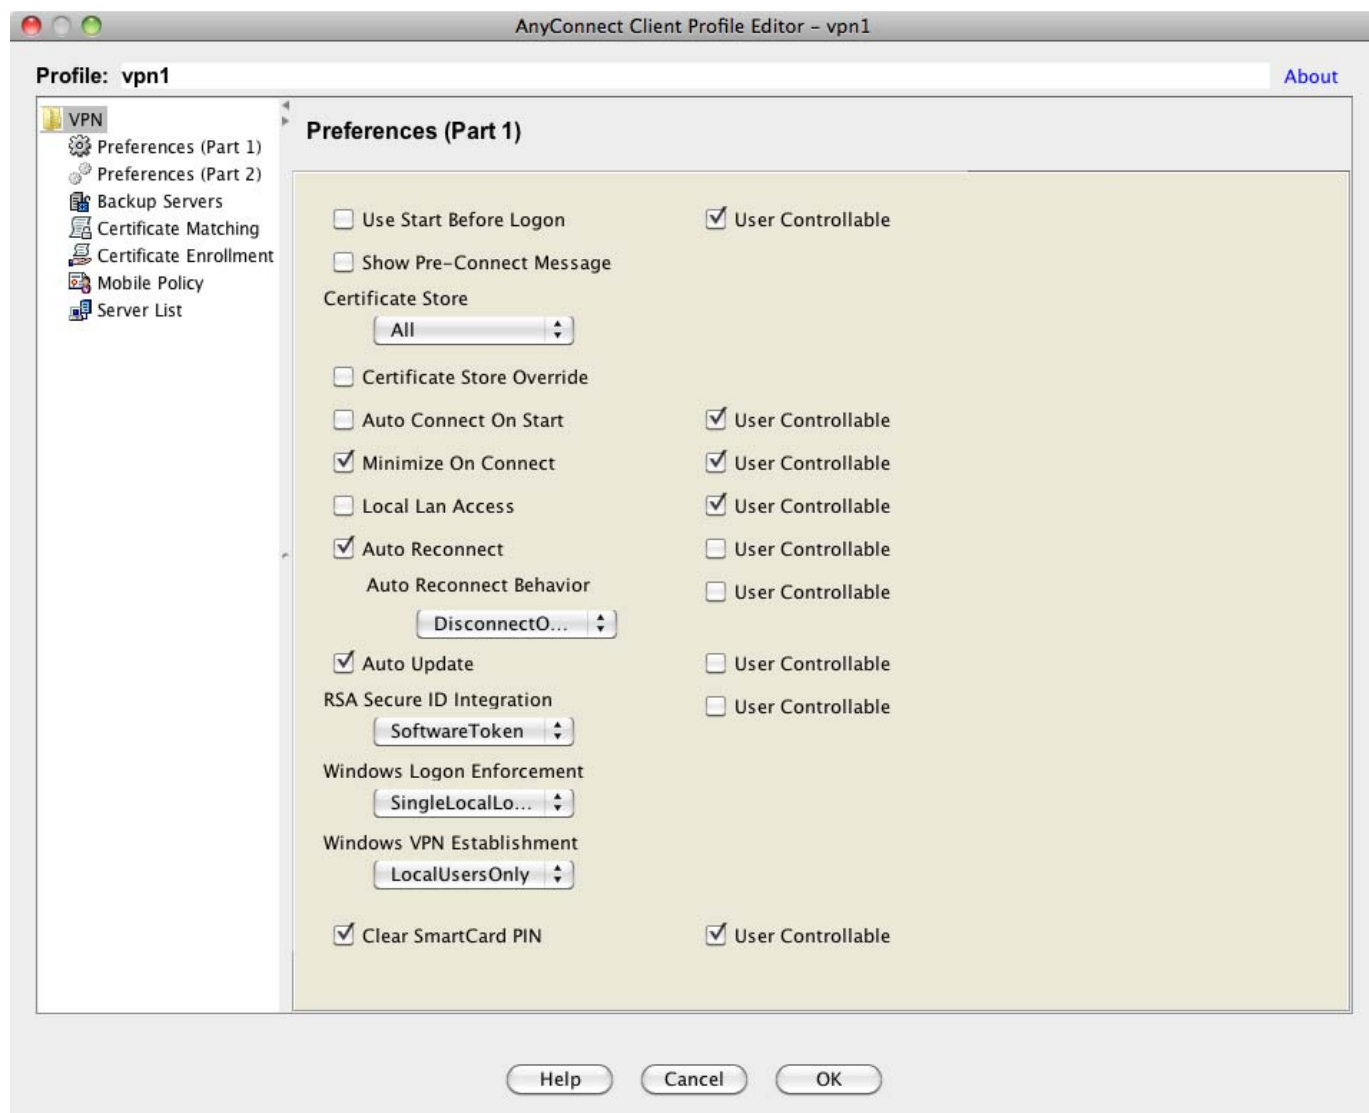
Below the table is a button labeled 'Restore default message texts'. At the bottom are 'Help', 'Cancel', and 'OK' buttons.

## Centralizing Policy Setting for VPN Users

Remote users need more in-depth security settings than corporate users because they normally connect from an un-trusted network, such as through the Internet. In the Unified Access Solution, we have enhanced the security for VPN users by mandating that remote users enter a RSA hardware token one-time password as a means to authenticate with the ASA. This section describes how to centralize policy settings by using Cisco's Adaptive Security Device Manager (ASDM). To obtain more information about ASDM, see:

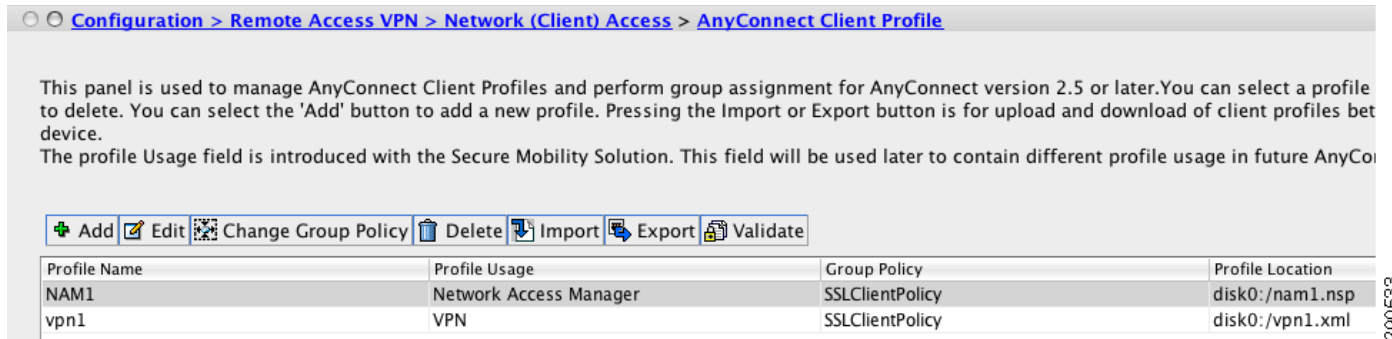
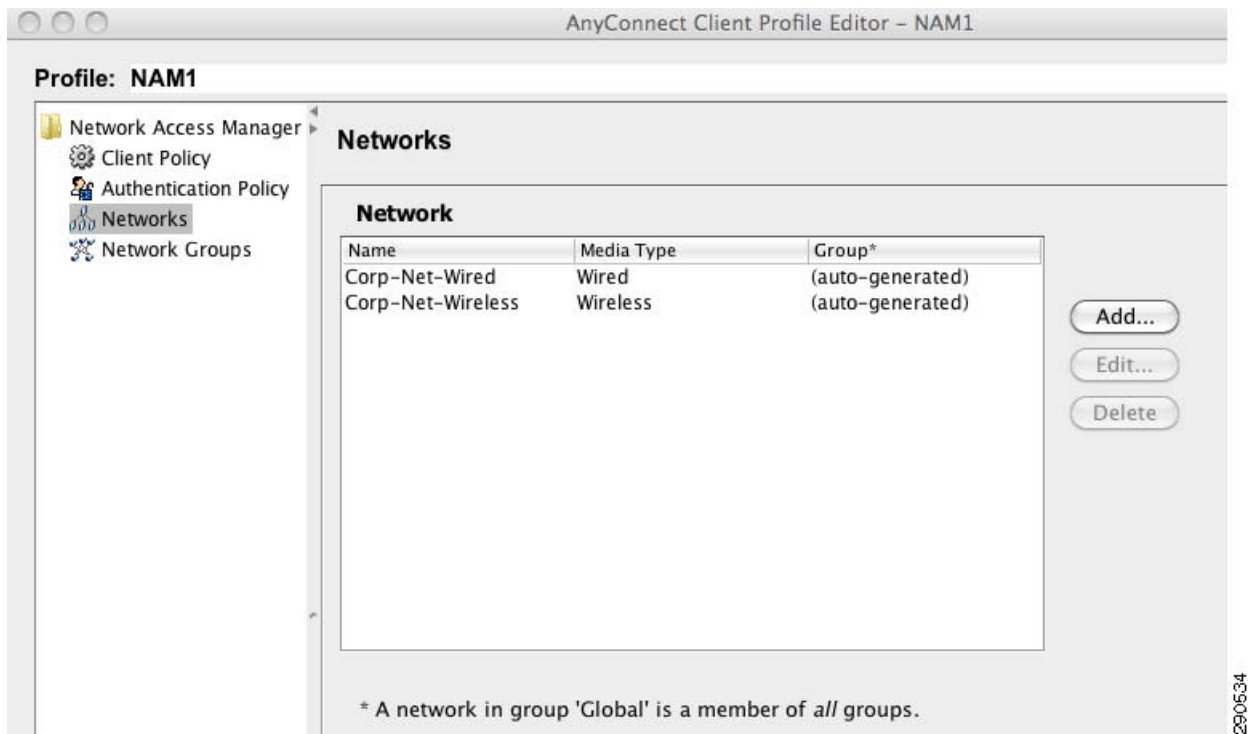
<http://www.cisco.com/en/US/docs/security/asa/asa80/asdm60/user/guide/usrguide.html>

Figure 3-31 illustrates how to configure the VPN profile on the ASA.

**Figure 3-31** Configuring the VPN Profile on the ASA

## NAM Profile

The NAM profile editor is designed to allow you to create NAM configuration profiles and create pre-configured client profiles. This configuration is deployed on the endpoints so that NAM can enforce administratively-defined end user and authentication policies and make the pre-configured network profiles available to end users. To use the profile editor, create settings for a profile, save it, and then place the configurations on the client. AnyConnect includes the profile editor inside ASDM, but a standalone version is also available. Refer to the *Cisco AnyConnect Secure Mobility Client Administrator Guide*, Chapter 2, “Deploying the AnyConnect Secure Mobility Client” ([http://www.cisco.com/en/US/partner/docs/security/vpn\\_client/anyconnect/anyconnect30/administratio n/guide/ac02asaconfig.html#wpxref89319](http://www.cisco.com/en/US/partner/docs/security/vpn_client/anyconnect/anyconnect30/administratio n/guide/ac02asaconfig.html#wpxref89319)) for profile editor requirements and deployment instructions.

**Figure 3-32 AnyConnect Client Profiles—VPN and NAM****Figure 3-33 AnyConnect Client Profile Editor for NAM**

The NAM profiles are loaded onto the AnyConnect client machine as XML files in the following directories:

- Windows XP:  
C:\Documents and Settings\All Users\Application Data\Cisco\Cisco AnyConnect Secure Mobility Client\Network Access Manager\newConfigFiles
- Vista/Windows 7:  
C:\ProgramData\Cisco\Cisco AnyConnect Secure Mobility Client\Network Access Manager\newConfigFiles

**Figure 3-34** Example of the NAM XML File

```

<group>
  <groupName>Default</groupName>
  <allowUserToSeeScanlist>true</allowUserToSeeScanlist>
  - <wiredNetwork>
    <displayName>Corp-Net-Wired</displayName>
    <connectionTimeout>40</connectionTimeout>
  - <authenticationNetwork>
    - <userAuthentication>
      - <authenticationMethod>
        - <eapPeap>
          <doNotValidateServerCertificate />
          - <unprotectedIdentityPattern encryptContent="true">
            - <EncryptedData xmlns="http://www.w3.org/2001/04/xmlenc#"
              Type="http://www.w3.org/2001/04/xmlenc#Content">
              - <CipherData>
                <CipherValue>SBIQc9l1ri5P0mOfMPnI/82D2lPWc+5gGEff8lqUsc=</CipherValue>
                </CipherData>
              </EncryptedData>
            </unprotectedIdentityPattern>
          - <enableFastReconnect>
            <alwaysAttempt />
            </enableFastReconnect>
          - <authMethods>
            - <builtinMethods>
              - <authenticateWithPassword>
                - <protectedIdentityPattern encryptContent="true">
                  - <EncryptedData xmlns="http://www.w3.org/2001/04/xmlenc#"
                    Type="http://www.w3.org/2001/04/xmlenc#Content">
                    - <CipherData>

```

200635

## Enabling AnyConnect VPN on the ASA

After defining the group-policy and connection profile on the ASA, the last piece is to enable the AnyConnect VPN feature on the ASA. After enabling AnyConnect, the administrator can also configure additional features, such as pointing to the AnyConnect image software, NAM profile, and VPN Profile. [Example 3-4](#) is the configuration for enabling AnyConnect modules.

### Example 3-4 AnyConnect Modules Definition

```

webvpn
  anyconnect keep-installer installed
  anyconnect modules value nam,telemetry,posture
  anyconnect profiles value NAM1 type nam
  anyconnect profiles value vpn1 type user
  anyconnect ask none default webvpn
  always-on-vpn profile-setting

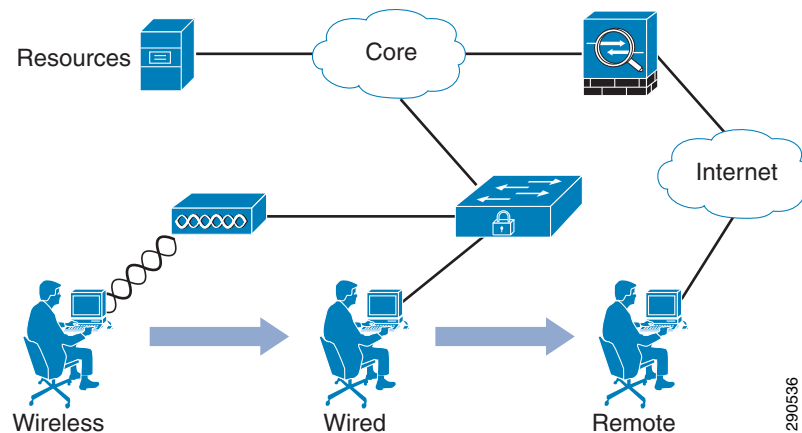
```

## Consistent Access Experience

As mentioned in the executive summary, the Unified Access Solution addresses the user experience pain points by providing:

- Migration from an existing 802.1x wireless network that uses native supplicants to Cisco's AnyConnect mobility client.
- Deploying 802.1x solution for wired users using Cisco's AnyConnect mobility client.
- Deploying a remote access solution using Cisco's AnyConnect mobility client.
- Distributing certificates for remote users using Cisco's AnyConnect mobility client.

[Figure 3-35](#) shows how a user can move around in an enterprise network.

**Figure 3-35 User Movement in an Enterprise Network**

Users on the secure campus network currently use the existing native Microsoft 802.x wireless supplicant in order to authenticate and gain access to the network. The user is on the trusted campus network. The user has authenticated via 802.1x wireless connection from a Windows 7 client. The client is using the native Microsoft 802.1x supplicant in order to authenticate and gain access to the trusted campus wireless network. The Windows 7 Client is authenticating using PEAP (MSCHAPv2) to a Cisco 1142 AP authenticator. The authentication policy is located on the Authentication Server (Cisco ACS 5.2).

Figure 3-36 shows a successful authentication by a Microsoft 802.1x supplicant.

**Figure 3-36 Successful Authentication**

### Campus Wireless

	<input type="checkbox"/>	Status	Name	Eap Authentication Method	Conditions	Eap Tunnel Building Method	Results	Hit Count
1	<input type="checkbox"/>		Rule-1	match EAP-TLS	match x509_PKI	-ANY-	CN Username	15
2	<input type="checkbox"/>		Rule-2	-ANY-	-ANY-	match PEAP	AD1	21

#### Most Recent Authentication

Time: February 16, 2011 8:43:59.970 PM  
 RADIUS Status: Authentication succeeded  
 NAS Failure:  
 Username: winxp  
 Network Device: cr22-ap1 : 10.125.130.3 : 1043  
 Access Service: Campus\_Wireless  
 Authorization Profiles: Permit Access  
 CTS Security Group:  
 Authentication Method: PEAP(EAP-MSCHAPv2)

## Migration from an Existing Native MS 802.1x Supplicant to a Cisco AnyConnect Solution

Once users have authenticated via the native 802.1x wireless supplicant, they are instructed to start the deployment of Cisco AnyConnect NAM/VPN Supplicant by clicking on a link via an internal campus Web page which redirects to the Cisco ASA trusted interface.

The users login to the ASA with either their RSA SecureID or Active Directory username/password. Once the user has authenticated to the ASA, the ASA determines which AnyConnect components are required to be downloaded to the client based on the client profile created by the administrator.

The NAM Supplicant must be deployed as part of the AnyConnect Secure Mobility Client VPN. The VPN is the core module for which all the other modules are installed. In order to install the AnyConnect NAM (802.1x Supplicant), the user may visit an internal Web page on the trusted campus network. This Web page points to the trusted (inside) interface of the ASA for which the ASA images are stored/maintained.

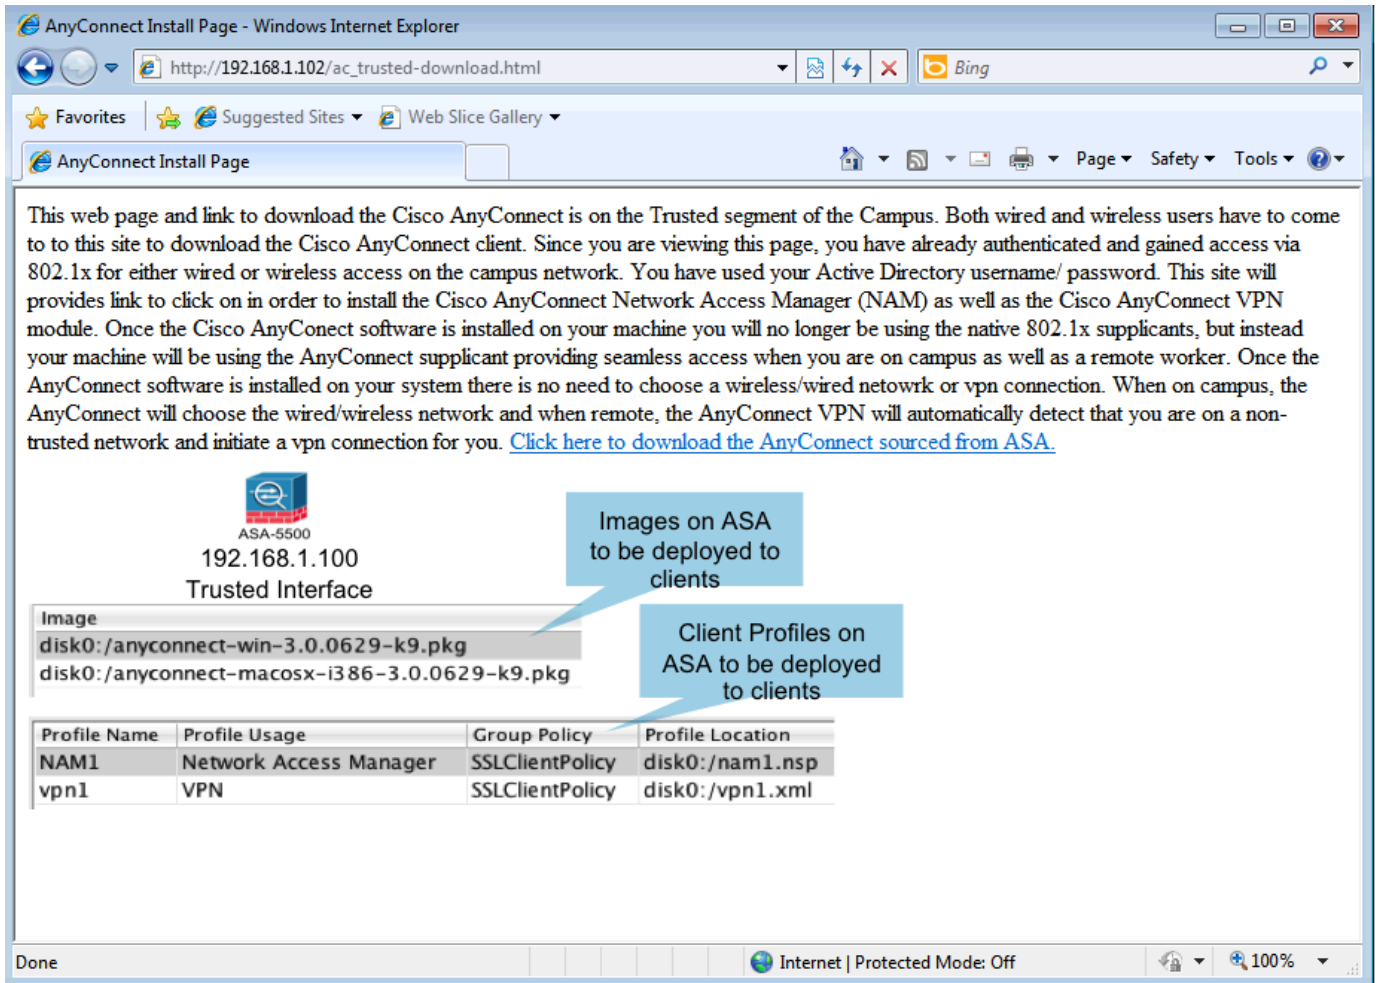
**Note**

---

A Cisco ASA is not required in order to use the AnyConnect Network Access Manager. The AnyConnect Network Access Manager can be deployed without the use of an ASA. This solution uses the ASA in order to deploy both the NAM as well as the NAM profiles to the client end point. When the Cisco AnyConnect NAM module is installed, it takes over and assumes control of all 802.1x supplicant activity. The native 802.1x supplicant is no longer able to provide 802.1x services.

---

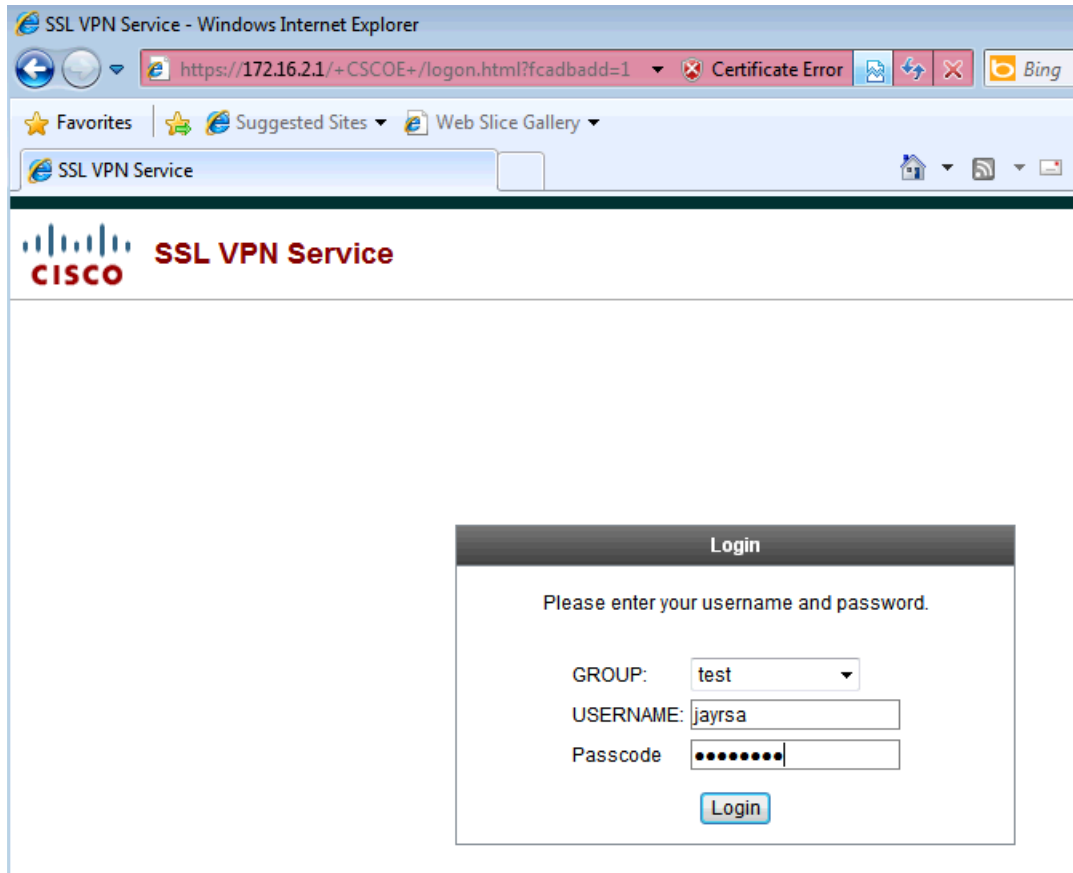
This allows administrators to maintain a repository of the latest AnyConnect modules, which are staged and ready to be dynamically deployed to users. Administrators could also choose to deploy the AnyConnect NAM via a software distribution system. However, having all the AnyConnect modules installed on the ASA allows for automatic unattended installs for all the AnyConnect modules to each user. [Figure 3-37](#) shows an internal Web page presented to a user on the trusted campus network.

**Figure 3-37** Web Page on Trusted Campus Network

The Web page redirects clients to download the AnyConnect VPN and AnyConnect NAM module.

On the screen shown in [Figure 3-37](#), the user clicks the link **Click here to download the AnyConnect sourced from ASA**, which redirects the user to the inside interface of the ASA for which all the AnyConnect modules are installed. Clients can also install the AnyConnect VPN and NAM modules from a VPN connection. If a client would like to install the AnyConnect modules via a VPN connection, the client simply accesses the outside interface (untrusted interface) of the ASA via HTTPS in order to begin the install process.

Once the user is redirected to the ASA via the trusted interface (users that are inside the ASA) or if users access the ASA on the untrusted interface (users residing on the public facing side of the ASA), they are both presented with the SSL VPN login screen shown in [Figure 3-38](#).

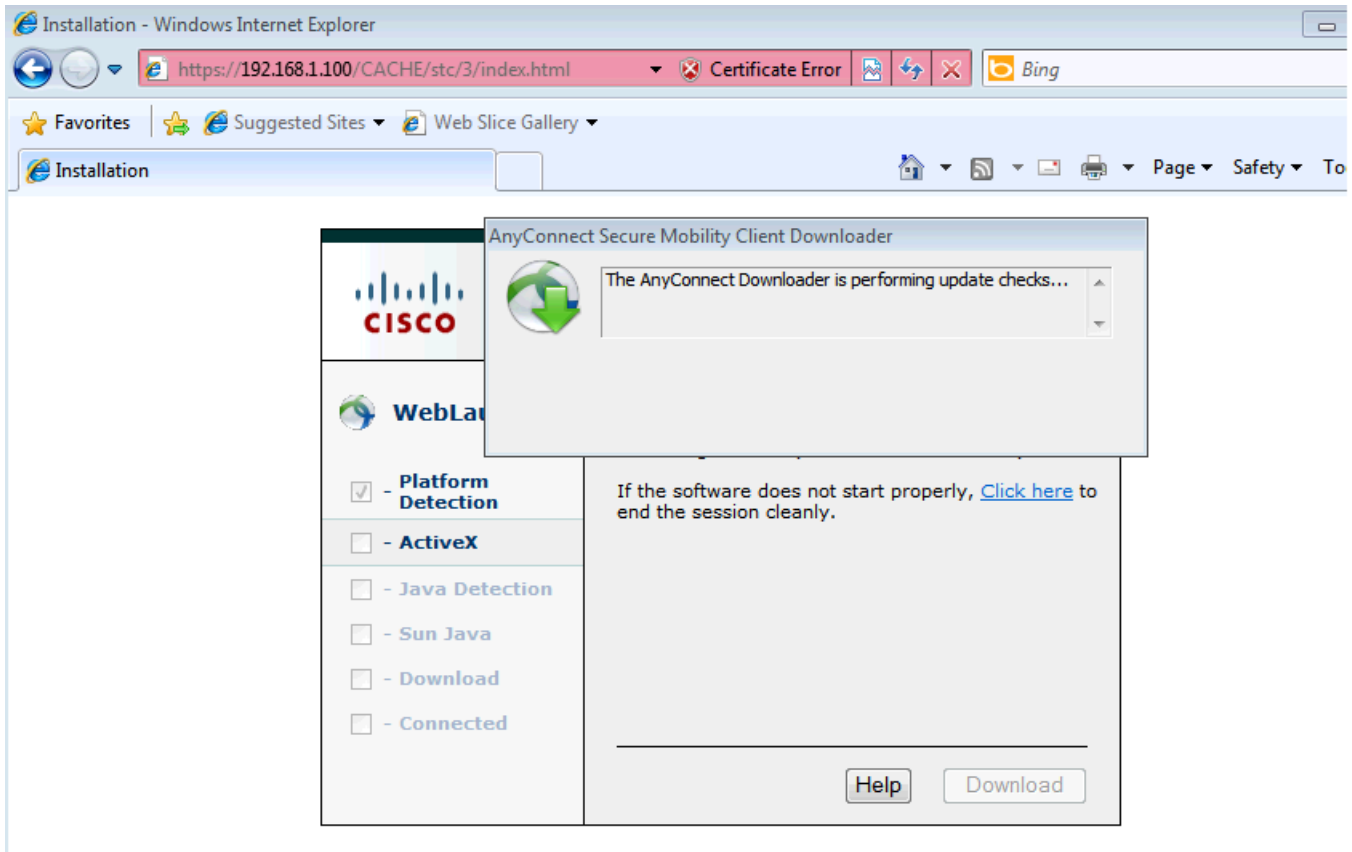
**Figure 3-38** SSL VPN Login Screen

In [Figure 3-38](#), notice the Group “test” is the user’s assigned group, which will determine which type of authentication method and which AnyConnect software modules will be loaded onto the client machine.

Based on the Web page, the user is instructed to login using either their SecureID token or their Active Directory username/password. In this scenario, the user is logging in with their RSA SecureID token.

Once the user has authenticated, there are a number of checks that are performed on the client and the ASA in order to determine which AnyConnect software must be installed or updated and if any profile updates are needed for the client.

[Figure 3-39](#) shows the ASA performing such checks.

**Figure 3-39** Update Determination

290540

The AnyConnect VPN module is the core AnyConnect module; thus, it must be installed prior to any other modules being installed. The AnyConnect VPN module is a requirement for all other AnyConnect modules. Based on the user's Connection Profile, the ASA determined that the user required the AnyConnect NAM supplicant to be installed and that two profiles needed to be sent to the client in the form of an XML file loaded onto the client's machine.

The location of the XML file, which is loaded onto the client machines, specifies the features and attribute values for each user type.

The location of the XML files are:

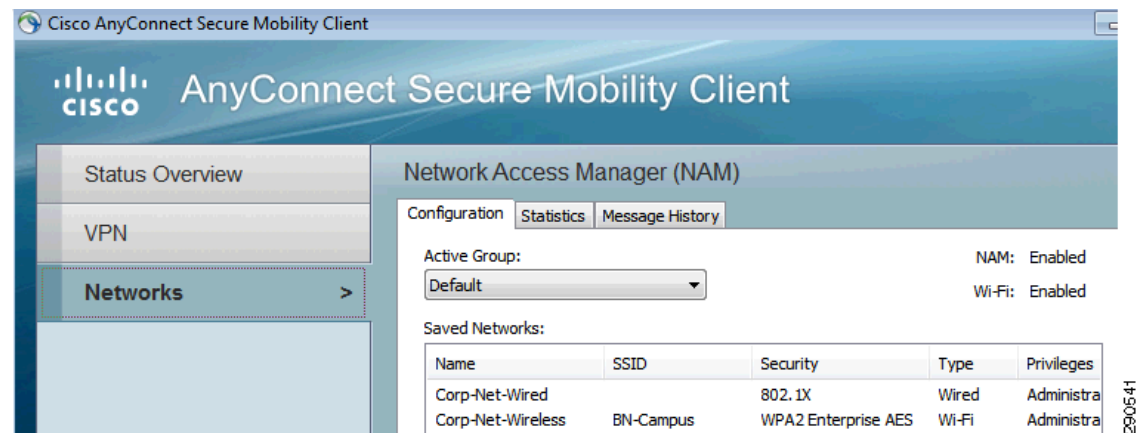
- Windows XP
  - Core client with VPN  
%ALLUSERSPROFILE%\Application Data\Cisco\ Cisco AnyConnect Secure Mobility Client\Profile
  - NAM (Network Access Manager)  
%ALLUSERSPROFILE%\Application Data\Cisco\ Cisco AnyConnect Secure Mobility Client\NetworkAccessManager\newConfigFiles
- Windows Vista
  - Core client with VPN  
%ProgramData%\Cisco\Cisco AnyConnect Secure Mobility Client\Profile

- NAM NAM (Network Access Manager)  
%ProgramData%\Cisco\Cisco AnyConnect Secure Mobility Client\NetworkAccessManager\newConfigFiles
- Windows 7
  - Core client with VPN  
%ProgramData%\Cisco\Cisco AnyConnect Secure Mobility Client\Profile
  - NAM NAM (Network Access Manager)  
%ProgramData%\Cisco\Cisco AnyConnect Secure Mobility Client\NetworkAccessManager\newConfigFiles
- Mac OS X  
All modules  
/opt/cisco/vpn/profile
- Linux  
All modules  
/opt/cisco/vpn/profile
- iOS Devices—Apple iPhone/iPodTouch/iPad  
The Apple iOS device supports only one AnyConnect XML profile. The contents of the generated configuration always match the most recent profile.

This AnyConnect NAM replaces the need for the user to use the native 802.1x supplicant and instead use the AnyConnect NAM, which has greater security and administrative control and allows the user to seamlessly connect to the campus network from both trusted and untrusted network segments using both wired and wireless networks. The client is connected to the trusted campus network dynamically when on the public Internet. The AnyConnect client understands the user is on an untrusted network and automatically establishes a VPN connection to the campus trusted network without any user intervention.

There are two profiles that are loaded onto the client machine during this process, one for secure wired access on the trusted campus and one for secure wireless access on the trusted campus network. Note that the privileges for both of these profiles (as seen below) are set to “Administrator”. This means that the security administrator on the ASA created these profiles and that the end user does not have access to alter or remove these campus trusted wired and wireless connection profiles.

Figure 3-40 shows the two client connection profiles that were loaded into the NAM on the client machine.

**Figure 3-40** Two Client Connection Profiles

Once the AnyConnect NAM and AnyConnect VPN modules are installed, there is no need for the user to manually reconnect to the wireless campus. The NAM dynamically launches and connects to the Corp-Net-Wireless network (the same network for which the client was authenticated when using the Microsoft Native 802.1x Supplicant). The user is prompted to enter their username/password credentials into the AnyConnect supplicant software in a similar manner as the user was presented a username/password for the Native Microsoft Supplicant.

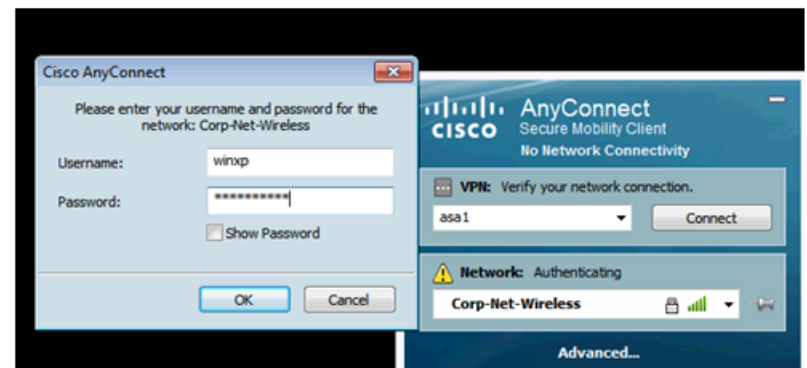
**Figure 3-41** Username/Password Screen**Figure 3-42** Connection Established Screen

Figure 3-42 shows that the user has authenticated to the campus wireless network and that the AnyConnect VPN Module has determined that a VPN connection is not necessary. The AnyConnect profile, which was created by the administrator on the ASA and pushed out to the client, is utilizing the feature Trusted Network Detection (TND). When TND is enabled for a connection profile, the AnyConnect supplicant determines if it is on a trusted network by sending TCP/UDP DNS queries in order to reach pre-configured secure campus name servers. If the AnyConnect supplicant is unable to reach the name server, the supplicant assumes it is on an un-trusted network and automatically launches a VPN connection.

## User Moves to a Wired Connection

When an end point that is using the AnyConnect NAM supplicant moves to a wired network connection, the AnyConnect NAM supplicant prefers all wired connections over wireless connections when AnyConnect is configured for “automatic mode”; this functionality can be overridden when the AnyConnect supplicant is in “manual mode”. AnyConnect automatically chooses the “Corp-Net-Wired” network profile over the existing “Corp-Net-Wireless” network profile. Only one connection is allowed at a time; all other connections on the end point machine are blocked. The AnyConnect UI automatically prompts the user if the EAP method requires (MSCHAPv2) and if the credentials are not cached. If the connection profile is configured to use SSO credentials and the credentials are incorrect, the user is required to change their password. Since the administrator on the ASA configures the network profiles, the user does not have access to edit/remove the network profiles. This greatly enhances the overall security and robustness of having AnyConnect as an 802.1x supplicant.

Figure 3-43 shows AnyConnect supplicant automatically selecting a wired connection.

**Figure 3-43 Automatic Selection of Wired Connection**



Example 3-5 shows the required configuration of the wired ports on the 802.1x-enabled switch.

**Example 3-5 Configuration of Wired Ports on 802.1x-Enabled Switch**

```
interface GigabitEthernet1/0/2
description 1X-ONLY
switchport access vlan 140
switchport mode access
switchport block unicast
switchport voice vlan 141
ip arp inspection limit rate 100
load-interval 30
srr-queue bandwidth share 10 10 60 20
```

```

queue-set 2
priority-queue out
authentication port-control auto
authentication fallback WEB-AUTH
mab
mls qos trust cos
dot1x pae authenticator
auto qos voip trust
storm-control broadcast level pps 1k
storm-control multicast level pps 2k
storm-control action trap
spanning-tree portfast
spanning-tree bpduguard enable
ip verify source
ip dhcp snooping limit rate 15
end

```

## User Moves to a Remote Location

The user has now moved to an untrusted segment (the public Internet) and is no longer on the campus trusted wired or wireless network. Once the wireless AnyConnect client is associated and obtains a valid IP address for the new network they just joined, the AnyConnect client again tries to establish contact with the internal name servers configured in the AnyConnect client profile. This time, the AnyConnect client is not able to establish connectivity to the trusted name servers and thus the AnyConnect client assumes it is on an untrusted segment and launches the AnyConnect VPN module. Figure 3-44 shows how AnyConnect client initiates a VPN connection when the user is on a un-trusted network, provided Trusted Network Detection is enabled.

**Figure 3-44** AnyConnect Initiates VPN Connection



Figure 3-45 shows a successful RSA SecureID login via AnyConnect SSL VPN on the RSA Server.

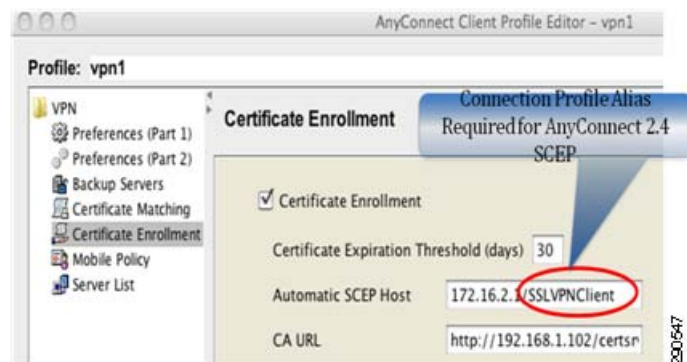
**Figure 3-45 Successful RSA SecureID Login**

Time	Activity Key	Description	Reason	User ID	Agent	Server Node IP	Client IP
2011-03-02 11:12:49.682	Principal authentication	User "jayrsa" attempted to authenticate using authenticator "SecurID_Native". The user belongs to security domain "SystemDomain"	Authentication method success	jayrsa	192.168.1.101	192.168.1.103	192.168.1.101

## AnyConnect Secure Mobility Client 2.4 for iOS Devices

We discussed earlier that the AnyConnect Security mobility client 3.0 supports Web deployment. The Security Mobility client version 2.4 which runs on iOS mobile devices does not support Web deployment; instead, users download and install the AnyConnect Security mobility client via the Apple iTunes Store.

AnyConnect 2.4, which is the supported version on the iOS devices (iPad, iPhone, and iTouch), supports SCEP by having the client mobile device send a SCEP request directly to the CA server once the VPN has been established. [Figure 3-46](#) shows the required configuration on AnyConnect Client Profile Editor on ASA for SCEP.

**Figure 3-46 Configuration on AnyConnect Client Profile Editor on ASA for SCEP**

[Figure 3-47](#) shows a trace on the CA server of the AnyConnect 2.4 client making a SCEP request directly to the CA server using the client source IP address. As mentioned the AnyConnect 2.4 client does not support the new 3.0 SCEP-Proxy feature; instead, the AnyConnect 2.4 clients sends a SCEP request directly to the CA server, rather than having the ASA proxy the SCEP, as shown in [Figure 3-47](#).

[Figure 3-47](#) shows an AnyConnect 2.4 client running on an iPad. Notice the source IP address is that of the iPad and not the ASA. This trace was taken on the CA server itself.

**Figure 3-47** Trace on CA Server

The user logs in using RSA Secured ID credentials in order to set up a SSL VPN so that the certificate enrollment process is conducted over a secure channel. Figure 3-48 shows how the user selects a connection profile after entering RSA credentials.

**Figure 3-48** Connection Profile Selection

In Figure 3-48, notice that the Group “SSLVPNClient” must match the group configuration in the above VPN Profile configuration. This is a requirement in order to support SCEP on AnyConnect 2.4. AnyConnect 3.0 does not require the Group “SSLVPNClient” to be specified in the VPN Profile configuration, however it is recommended in order to support both AnyConnect 3.0 SCEP-Proxy as well as AnyConnect 2.4 SCEP forwarding.

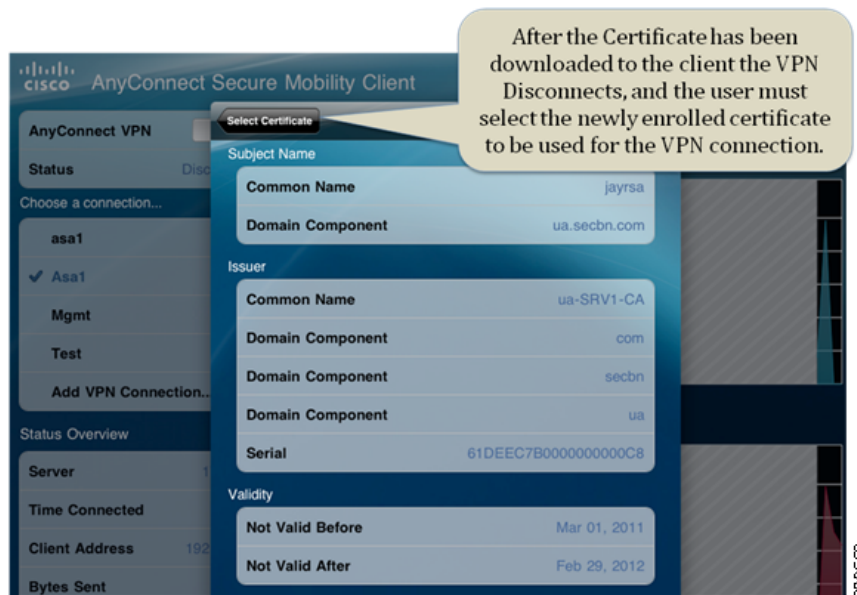
In Figure 3-49, the AnyConnect is displaying a message which indicates the enrollment process using SCEP on the AnyConnect client has succeeded.

**Figure 3-49** Enrollment Process Success

A valid client certificate is now available to be used for the next SSL VPN connection. Figure 3-50 shows the IP address obtained by the iPad.

**Figure 3-50 iPad Client IP Address**

In [Figure 3-51](#), the user chooses the newly enrolled client certificate, which the AnyConnect client sends to the ASA during authentication of the SSL VPN connection.

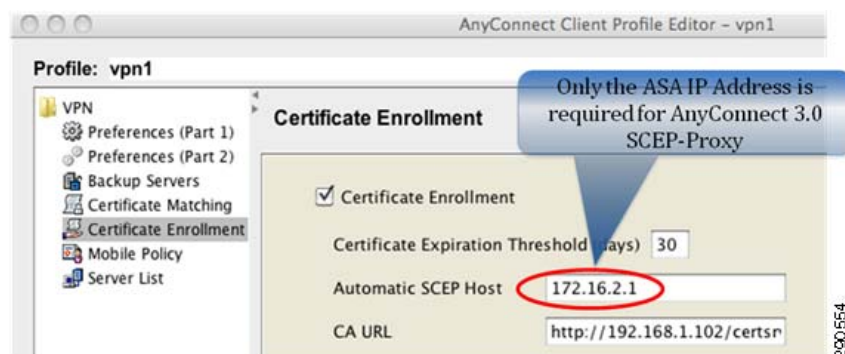
**Figure 3-51 Selecting the Newly Enrolled Client Certificate**

In [Figure 3-52](#), the AnyConnect client is displaying the client IP address that is the source IP address of the end point for the SSL VPN Connection. Notice that this IP address (192.168.1.11) is the source IP address which was used as the source of the SCEP request in order to obtain the client certificate during the certificate enrollment process.

**Figure 3-52** *Client IP Address*

## AnyConnect Secure Mobility Client 3.0

This section discusses a scenario where a remote user is in possession of an RSA SecureID token, but still does not have possession of a digital certificate. When the user initiates an SSL VPN session using the AnyConnect client, the ASA which is configured for two-factor authentication, as well as client certificate authentication, authenticates the user initially using the RSA two-factor authentication token. After validating the user/token, the ASA then sends the AnyConnect client a request for a certificate. However, since the client does not yet have a certificate, the ASA begins the SCEP process on behalf of the client. The SCEP process is completed over the SSL VPN connection. Once the certificate enrollment is successful (the client has received the certificate from the CA), the VPN connection is terminated. Immediately following the termination of the VPN connection, the AnyConnect client again initiates a connection to the ASA, but this time the user logs in with both two-factor authentication as well as presenting the newly-obtained client certificate to the ASA. [Figure 3-53](#) shows the configuration of the ASA acting as a SCEP-Proxy. Also note that the IP address 172.16.2.1 shown in [Figure 3-53](#) is the outside interface of the ASA.

**Figure 3-53** *AnyConnect Client Issuing a SCEP Request*

[Figure 3-54](#) displays a packet capture, which was taken on the CA Server. This packet capture shows the client IP address which is the trusted interface of the ASA. The ASA is proxying the SCEP request on behalf of the client. The source IP address of the SCEP request is 192.168.1.100, which is the inside/trusted interface of the ASA. The destination IP address (192.168.1.102) is the CA server.

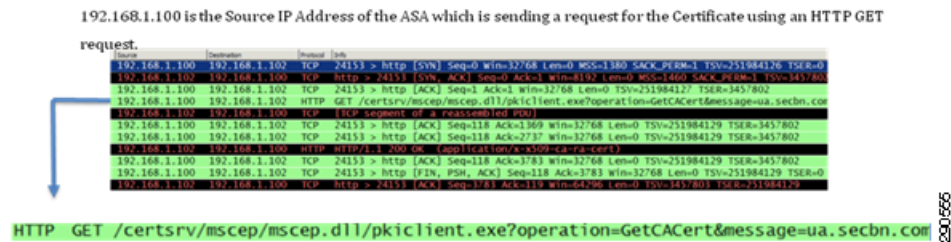
**Figure 3-54** Packet Capture on CA Server

Figure 3-55 displays a Windows 7 client running AnyConnect Secure Mobility Client 3.0 using the same VPN Group (SSLVPNClient) as the AnyConnect Secure Mobility Client 2.4 used.

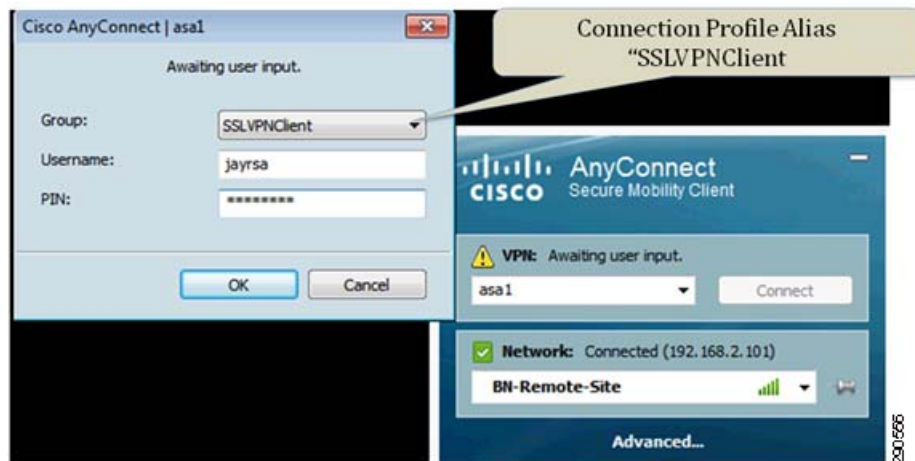
**Figure 3-55** Windows 7 Client Running AnyConnect Secure Mobility Client 3.0 Using VPN Group

Figure 3-56 displays a Windows 7 client running AnyConnect Secure Mobility Client 3.0 has successfully received a client certificate via the SCEP-Proxy feature on the ASA.

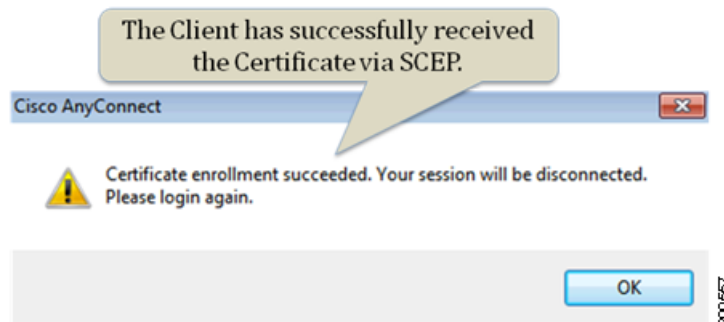
**Figure 3-56** Windows 7 Client Running AnyConnect Secure Mobility Client 3.0—Successful Client Certificate

Figure 3-57 displays a Windows 7 client running AnyConnect Secure Mobility Client 3.0 where the user is prompted to choose the certificate to send to the ASA during authentication for the SSL VPN connection.

**Figure 3-57** *Windows 7 Client Running AnyConnect Secure Mobility Client 3.0—Certificate Selection*

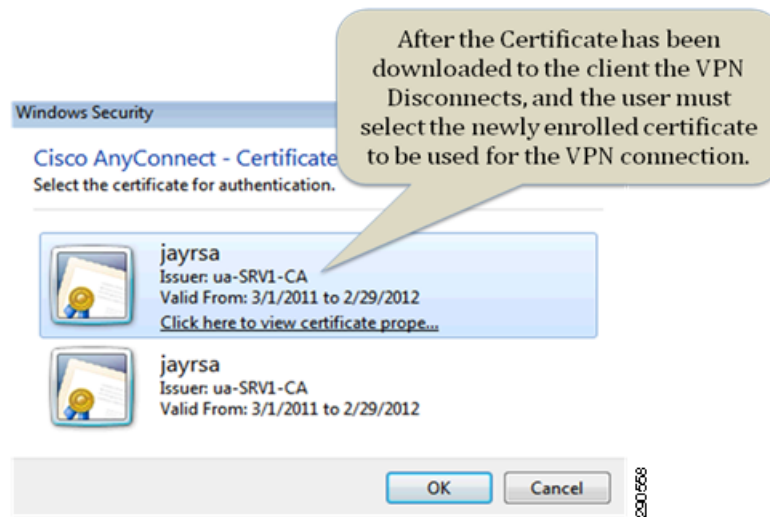
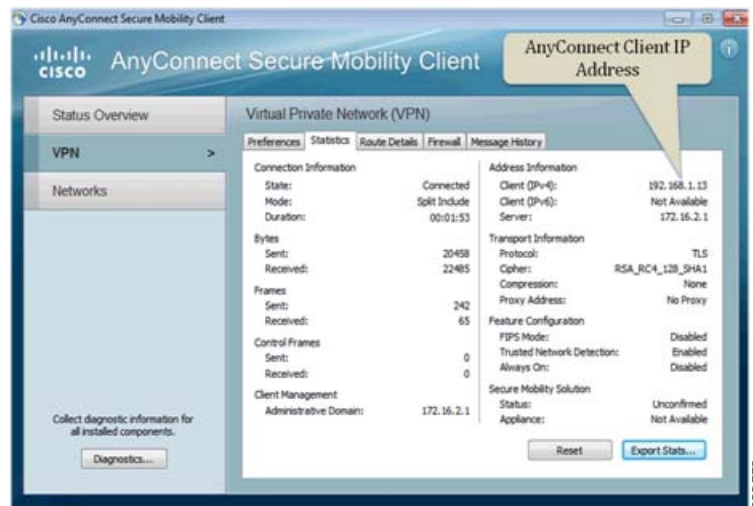


Figure 3-58 displays a Windows 7 client running AnyConnect Secure Mobility Client 3.0 and displaying that the source IP address of the AnyConnect client is 192.168.1.13. Note that during the SCEP process, this source IP address (192.168.1.13) was not used as the SCEP source; rather, the trusted/inside interface of the ASA was used.

**Figure 3-58** *Windows 7 Client Running AnyConnect Secure Mobility Client 3.0—Source IP Address*



## Summary

There has been major adoption of 802.1x technology for wireless devices because of the many benefits the technology provides in authentication and authorization. For wired users, the deployment of 802.1x technology is challenging because of the inconsistent behavior of supplicants on the native operating systems. Moreover, with the evolution of mobile devices in the market today, many employees require

access to business resources using both their traditional devices, such as laptops and desktops, and their personal devices, such as smart phones and tablets. While the productivity of employees is greatly improved with the ability to access the network with different end points, it also increases security risks. To balance the productivity gains versus the security risks, network architects must design the network such that users are securely authenticated and identified and their transactions are logged with a single solution for wired, wireless, and remote users.

The Unified Access Solution for security enables the network architect to:

- Maintain a centralized access control on all the access points of the network by authenticating both wired and wireless users.
- Identify users with digital certificates along with strong two-factor authentication.
- Provide digital certificates to users who can access the CA server directly with the help of the SCEP protocol and also to remote users by the SCEP-PROXY feature supported by ASA.
- Support different flavors of EAP authentication methods.
- Provide a method to easily migrate from an existing 802.1x environment with native supplicants to the Cisco AnyConnect client.
- Enforce two-factor authentication for remote users by using RSA SecureID tokens.
- Centralize the network profiles for 802.1x clients for both wired and wireless users.



## CHAPTER 4

# Context-Awareness for Wired Devices

---

This chapter focuses on best practices to enable and use Cisco Context-Aware Services for wired endpoint location tracking in Unified Access solutions. It is intended as a guide to producing functional designs that incorporate the Cisco Mobility Services Engine (MSE) and Cisco Wireless Control System (WCS) in order to help fulfill common business needs in enterprise environments.

Note that while this chapter is comprehensive, it is first and foremost intended as a design considerations document and is not intended to be a configuration guide to Cisco Catalyst switches, the Mobility Services Engine, or the Wireless Control System. For comprehensive configuration and general deployment guidance, the reader should refer to the various in-depth configuration guides referenced throughout this chapter.

## Introduction

### What Are Context-Aware Services?

Context-Aware Services provide value-added functionality to business applications by capturing, integrating, and consolidating intelligence about users and their endpoint devices from various points in the network. Together with Unified Access wired and wireless networks, Cisco Context-Aware Services help enhance application functionality by making this information readily available via an established application programming interface.

Cisco Context-Aware software is a mobility service with these characteristics:

- Acts across multiple edge technologies, such as 802.11 wireless and 802.3 wired networks.
- Provides a value-add function across multiple network elements.
- Provides an interface to a value-add function for external applications and servers using a well-defined API.
- Adds intelligence to the network and enhances usability.
- Provides visibility into the network that applications and servers would not otherwise easily obtain.
- Integrates with other network services to deliver higher-order functionality and value.
- Can be managed using other Cisco tools or the mobility service API.
- Can be deployed across multiple engines to scale the function it provides.

Location Services associated with the Cisco Context-Aware Mobility Solution provide a single unified view of contextual information through the Mobility Services Engine API. Cisco Context-Aware Location Services enable both queries for contextual information as well as server registration for asynchronous events occurring among both wired and wireless access controllers.

In some cases, it may be necessary to track an asset with a high degree of accuracy throughout an enterprise, such as when a valuable missing asset must be located. On the other hand, some applications using Context-Aware Services may only require general indication of whether an asset is in or out of a permissible zone (such as the confines of a shipping and receiving dock, for example).

Context-Aware Services can also provide location and other contextual information for wired devices attached to Cisco Catalyst LAN switches, such as the 3560, 3750, 4500, and 4900 series. Catalyst switches can provide civic location details for wired devices to the Cisco Mobility Services Engine based on pre-configured information specified for each switch port. This information can then be presented to users in a tabular format combined with other contextual information, such as user name, device serial number, and emergency location identifier numbers (ELINs).

**Note**

A civic location specifies the civic address and postal information for a physical location using fields such as the number, street or road name, community, and county assigned to residential, commercial, institutional, and industrial buildings (e.g., 31 Main Street, Alpharetta, Georgia 30004). An emergency location identifier number (ELIN) is a number that can be used by the local public safety answering point (PSAP) to look up the geographic location of the caller in a master database known as the automatic location information (ALI) database. The ELIN also allows the PSAP to contact the emergency caller directly in the event the phone call is disconnected.

Additional information regarding civic address location is available from the IETF in the following RFCs:

- DHCP Option for Civic Addresses Configuration Information:  
<http://www.rfc-editor.org/rfc/rfc4776.txt>
- Revised Civic Location Format for Presence Information Data Format Location Object:  
<http://www.rfc-editor.org/rfc/rfc5139.txt>

For a detailed explanation of Context-Aware Services, see:  
<http://www.cisco.com/en/US/netsol/ns788/index.html>

## Why Should I Use Context-Aware Services?

Cisco Context-Aware Services provide the capability to determine physical location in the network, as well as additional contextual information such as the MAC address, IP address, serial number, or 802.1x username associated with the tracked entity. For example, at the time this document was published, tracked entities include wired switches, wired or wireless endpoints such as phones, physical computers, virtual machines (VM), Digital Media Players (DMP), and IP video cameras. This is by no means intended to be an exhaustive list, but rather is intended to illustrate the flexibility of Cisco Context-Aware Services in tracking a variety of different types of wired physical and virtual devices that are capable of connecting to Ethernet networks.

In order to provide delivery of a best-in-class network experience to end users, a Unified Access network solution should be aware of the dynamics surrounding users and endpoint devices, such as their current location. A wide variety of devices can appear on the network, both wired (switches, routers, IP phones, PCs, access points, controllers, video digital media players, and so on) and wireless (mobile devices, wireless tags, rogues, and so on). Often, locating missing assets in modern-day enterprises becomes a time consuming, error-prone, and costly manual process. In these cases, it has all too often become

simply more cost effective to replace the asset rather than expend employee productivity attempting to locate it. However, shrinkage costs and productivity losses will mount over time and can take their toll on the balance sheet. The inability to quickly and efficiently locate missing assets and ensure their availability when and where they are needed in the business cycle can severely throttle the productivity of even the best-managed organizations.

Cisco Context-Aware Services have been a part of Cisco wireless networks for some time and with recent enhancements regarding wired location capabilities, they are poised to play an increasingly important role in Unified Access networks. For example, future Context-Aware Services will enable location to be considered as a security policy attribute. Thus, a phone in the lobby of an office building can be assigned different policies than a phone in a conference room or in an employee's office. To date, authorization policies have been statically administered for endpoint devices based on characteristics such as the endpoint's MAC address or the user's personal access credentials. The location of the endpoint has traditionally not been considered in the determination of which policies might be most appropriate depending on location. Context-Aware Services enable security policies to evolve such that policy assignment will be based not only on the user's credentials and MAC address, but will include location and other contextual information. This culminates in the ability to offer security policies offering greater differentiation than what has been traditionally possible.

While integration of Context-Aware Services with security authorization and authentication network services has the potential to revolutionize how we grant access to network users, the context-aware story does not stop there. The information about wired and wireless endpoints contained within the context-aware databases can be used to bring contextual awareness to a host of other applications, including Unified Communications and Unified Presence. Thus, integration between the pillars of the Cisco Borderless Networks Architecture will make use of shared contextual information to change how voice and instant messaging users, for example, might detect user presence on the network. Instead of simply knowing if a user is present and available, the user's presence status could change based on contextual information such as their current location, whether they are stationary or in motion, or other characteristics.

In some cases, endpoints themselves may already be aware of their location. Context-Aware Services must be capable of making use of such endpoint intelligence. Examples of endpoints that can dynamically determine their own location would include those equipped with onboard Global Positioning System (GPS) capabilities. In other cases, endpoints derive their location from the network (either wirelessly or through a wired medium). In these cases, Context-Aware Services can play a key role in learning or supplying these endpoints with their current location via well known protocols such as Cisco Discovery Protocol (CDP), Link Layer Discovery Protocol (LLDP), and Link Layer Discovery Protocol for Media Endpoint Devices (LLDP-MED). And in other cases, endpoints may already be statically configured with their location information beforehand or receive such static location configuration from application servers. Dynamic location update provides better support of mobile endpoints, whether they are wired or wireless.

For example, if a wired phone is provisioned for Seat 1A/02, Room 100, Floor 1, Sorenson Building, and if the person whose phone this is moves his cubicle to Seat 3F/16, Room 306, Floor 3, Paine Building, the location information associated with this phone would need to be manually updated if there is no means by which it can be updated dynamically. However, if the endpoint is capable of "learning" its location from its point of attachment in the enterprise network, this location can be used by the endpoint as well as other network applications with which the endpoint interacts. The contextual information associated with the device can be recorded by a location server and be made available via an API to administrative and other applications wishing to track the whereabouts of the device, even if the device itself is not able to interact directly at the current time.

While the information that resides within the context-aware databases of the Mobility Services Engine is, in and of itself, highly valuable, the true power of Context-Aware Services blossoms when it is put to work in a true cross-pillar integration scenario and combined with the power of Cisco Unified Access solutions and Cisco Network Services. This results in enhanced capabilities that have the potential to shatter the boundaries of what we have come to accept as traditional enterprise computing networks.

## Use Cases For Context-Aware Services and Wired Endpoints

Context-Aware Services can help businesses answer important questions concerning enterprise assets as well as the users of those assets. In this way, Context-Aware Services can help contribute towards improving an organization's efficiency, increasing productivity, and helping to lower costs.

There are several important use cases where Context-Aware Services can yield real-world business benefits in Unified Access networks today, including:

- **Asset Recovery**—Context-Aware Services provide a central portal to the location of wired endpoint devices on the network. In the wired environment, context-aware location is capable of tracking not only the MAC address, assigned IP address, and 802.1x user name of attached wired devices, but the serial number or Unique Device Identifier (UDI)<sup>1</sup> information of endpoints compatible with CDP or LLDP-MED. This information can be used individually or in combination to determine not only the switch port to which the wired endpoint is attached, but also the civic location or ELIN information that is associated with this switch port. Having all of this information at the administrator's disposal can significantly accelerate recovery of missing assets. For example:
  - Recover misplaced assets that have been temporarily relocated within the enterprise legitimately, but never returned to their original location after such temporary use was completed.
  - Recover stolen assets or assets that have been relocated without proper authorization (perhaps between departments or divisions). This might include:
    - Assets purchased by an enterprise and “borrowed” by another business unit. While these assets are likely still on company property and in legitimate use, the original business unit in all likelihood wishes to have the asset returned and not have to incur the cost of asset replacement.
    - Assets that have been the target of theft or otherwise converted to unauthorized use outside the scope of the enterprise business mission. A good example might be a laptop that was removed by a student resident employed part-time by a university. The student removes the laptop from a laboratory and converts it to personal use in his or her dormitory, attaching it to the high-speed university wired network. Unknown to the student, the entire university campus (including all dormitories) makes use of Context-Aware Services as part of its deployment of a Cisco Validated Design (CVD), Unified Access Network Design. Thus, whenever the student uses the laptop on the network, network administrators and university loss prevention staff have the ability to identify the laptop as well note its location using the Cisco WCS and the Mobility Services Engine. Additionally, the Mobility Services Engine via its API allows for third-party applications to access historical asset location and identify a pattern of movement. This makes it possible to establish a record of illegal asset movement and determine the time and date assets were last seen on the network and where they were located at that time. Using the API, third-party applications can use the current

1. The Unique Device Identifier (UDI) is the Cisco Systems product identification standard for hardware products. More information about UDI can be found at:  
[http://www.cisco.com/en/US/products/products\\_identification\\_standard.html](http://www.cisco.com/en/US/products/products_identification_standard.html).

and historical location databases to issue alerts and notifications and apprise loss prevention and other authorities as to the appearance of a “wanted” asset on the ports of any monitored Catalyst switch.

Figure 4-1 is a visual depiction of how Context-Aware Services for wired endpoints could be used to support asset recovery efforts within an enterprise. This example assumes we are attempting to locate a missing Cisco 9971 IP phone with the MAC address E8:04:62:EA:75:7E. This device was originally installed on port 1/0/4 in a Catalyst switch with the IP address 192.168.84.11. However, it is no longer there. If this 9971 phone is present on the network, we would like to know its physical location so that it can be retrieved and returned to its original location and owner. Using Context-Aware Services, identifying and locating this wired IP phone is a straightforward procedure. First, we access the list of all wired clients of which the Mobility Services Engine is currently aware. We then search on the MAC address of the IP phone. We can see in image #2 that the device is currently connected on switch port 1/0/3 of the switch with IP address 192.168.84.30. Clicking on the device MAC address in image #2, we are presented with an array of information about the device and its current location. We can verify the serial number of the phone in image #3, allowing us to increase our confidence that this is indeed our missing asset and reassure ourselves that we are not the victim of MAC address spoofing. Image #4 and #5 in Figure 4-1 provide us with the information we seek, namely, that port 1/0/3 on this switch terminates at cubicle/seat identifier 1C/1-3 in room 103 on the first floor of building 300 of ABC Corporation. This building appears to be used by ABC Corporation for shipping and receiving. In case we are unfamiliar with where building 300 is, we can see that it is located at 9300 Kimball Bridge Road, in the City of Alpharetta, Fulton County, Georgia, postal zip code 30022.

In summary, with the information provided by Context-Aware Services in Figure 4-1, we can:

1. Search for the missing asset by MAC address.
2. Determine if the asset is currently connected to the network and, if so, identify its connected switch and switch port.
3. Optionally verify this is indeed the asset we seek by comparing its reporting serial number against any serial number for the asset that we may have in our records from its time of purchase or original deployment.
4. Determine the physical location where the switch port in #2 above is terminated. If the information in the switch is provisioned accurately and correctly, the location of the device can be determined right down to the floor, room and cubicle, office, or seat location. This is often sufficient to expedite efficient retrieval of the asset.

Figure 4-1 Example of Wired Asset Identification and Location

**1**

Wired Clients: sh-mse3300  
Services: Mobility Services > sh-mse3300 > Context Aware Service > Wired > Wired Clients

MAC Address*	IP Address	Username (802.1x)	Serial Number	Status	Switch IP Address	Port Type	Slot	Module	Port	VLAN Id	Civic Address
00:04:23:b3:34:b0	10.125.103.213			Connected	192.168.84.15	10Gbit	1	0	2	102	ABC Corporation 18/2 1 Building 210 1100 Westside Parkway Alpharetta Georgia 30009 US
00:07:85:13:0f:df	10.125.115.149			Connected	192.168.84.30	10Gbit	9	0	4	105	ABC Corporation 30/9-4 33 Building 300 9300 Kimball Bridge Road Alpharetta Georgia 30022 US

**2**

Wired Clients: sh-mse3300  
Services: Mobility Services > sh-mse3300 > Context Aware Service > Wired > Wired Clients

Search results for Wired Clients with <IP,User Name,MAC,VlanId> matching 'e8:04:62:ea:85:7e'

MAC Address*	IP Address	Username (802.1x)	Serial Number	Status	Switch IP Address	Port Type	Slot	Module	Port	VLAN Id	Civic Address
e8:04:62:ea:85:7e	10.125.115.22		FCM43585ST	Connected	192.168.84.30	10Gbit	1	0	3	103	ABC Corporation 1C/A-3 1 Building 300 9300 Kimball Bridge Road Alpharetta Georgia 30022 US

**3**

Wired Clients: "e8:04:62:ea:85:7e": sh-mse3300  
Services: Mobility Services > sh-mse3300 > Context Aware Service > Wired > Wired Clients

Device Information

MAC Address	e8:04:62:ea:85:7e
IP Address	10.125.115.22
Username (802.1x)	
Serial Number	FCM43585ST
UDI	
Model No.	CP-9971
Software Version	sg9971.9-1-1SR1
VLAN Id	103
VLAN Name	VLAN0103

**4**

Wired Clients: "e8:04:62:ea:85:7e": sh-mse3300  
Services: Mobility Services > sh-mse3300 > Context Aware Service > Wired > Wired Clients

Civic Address

Name	ABC Corporation
Street	Kimball Bridge Road
House Number	9300
House Number Suffix	
Address Line 2	A10-2 Jack-1
City	Alpharetta
State	Georgia
Postal Code	30022
Country	US

**5**

Wired Clients: "e8:04:62:ea:85:7e": sh-mse3300  
Services: Mobility Services > sh-mse3300 > Context Aware Service > Wired > Wired Clients

Advanced

ELIN	19789363023	Road Branch	-
Floor	1	Road Sub-branch	-
Building	Building 300	Road Pre-modifier	-
Apartment		Road Post-modifier	-
Room	103	Leading Street Direction	-
Place Type	Shipping & Receiving	Street Trailing Suffix	-
Neighborhood		Street Suffix	-
Landmark		Postal Community Name	-
Seat	1C/A-3	Post Office Box	-
Additional Code		City Division	-
Road		County	Fulton
Road Section			

- Asset Inventory—Context-Aware Services make use of CDP and LLDP-MED protocols to acquire serial number information from compatible wired endpoint devices. Any available wired endpoint serial number information acquired is displayed alongside location information for all wired endpoints that the system is currently tracking. Consequently, this capability can be used to conduct impromptu inventories of wired devices across the enterprise by examining wired device listings by serial number and comparing this information to actual deployment records. This can help enterprise IT administrators and others responsible for asset deployment to better determine whether assets have been moved between offices, buildings, or campuses improperly or without authorization. Figure 4-2 illustrates an example of how this capability could be used to provide a quick inventory of all Cisco IP phones that are attached to one or more Catalyst switches participating in Context-Aware Services. Clicking on the MAC address hyperlink for any of these phones would provide the administrator not only with the civic location of the phone, but also the assigned ELIN information that has been defined in the switch.

**Figure 4-2 Example of Quick Inventory of IP Phones**

System: mse1  
Services: Mobility Services > Context-Aware Service > Wired > Wired Clients

Wired Clients: mse1

MAC Address *	IP Address	Username (802.1x)	Serial Number	State	Switch IP Address	Port Type	Slot	Module	Port	VLAN Id
00:12:27:02:00:02	10.1.87.249		DNF10321CA2	Connected	10.1.96.41	1Gbit	1	0	20	96
00:1a:27:0b:53:07	10.1.87.233		DNF104511K7	Connected	10.1.96.25	1Gbit	1	0	1	56
00:1a:27:03:06:06	10.1.87.249		DNF10451DF6	Connected	10.1.96.25	1Gbit	1	0	1	56
00:1b:2a:06:4b:4b	10.1.87.244		FCH11098C87	Connected	10.1.96.25	1Gbit	1	0	1	56
00:1b:2a:06:46:7f	10.1.87.236		FCH11098CAD	Connected	10.1.96.25	1Gbit	1	0	21	56
00:1b:2a:06:05:43	10.1.87.253		FCH11098C0B	Connected	10.1.96.25	1Gbit	1	0	3	00

227624

- **Location by User Name**—In addition to locating the endpoint device that a user may be utilizing to access the network by its MAC or IP address, Context-Aware Services for wired endpoints also provide us with the ability to locate a wired user by the username entered during 802.1x authentication. This information is not only shown during a general listing of wired clients, but it is also a search parameter. While there are other sources that can locate the switch and switch port assigned to an 802.1x user in the Unified Access network, Context-Aware Services provide the administrator with the ability to quickly and efficiently determine not only what switch and port the user has been assigned, but the location in which the switch port terminates. Thus, the administrator is provided with a “one-stop shopping” source that can translate between a person’s 802.1x username and their location without any intermediate translation.

Figure 4-3 illustrates the process of searching on 802.1x username and how easily the search results are translated into an actual location. In image #1 we search for the username “1302280\_user1” across all of the wired clients of which the Context-Aware Services for this Unified Access network is currently aware. User 1302280\_user1 is located in image #2 and we can see the civic address information as well as any information pertaining specifically to how the Ethernet jack may be labeled. Image #3 allows us to retrieve additional information about where this user is located, including the building number, the ELIN, and the room and seat number assigned.

**Figure 4-3 Using for a Wired Client by 802.1x Username**

Wired Clients: mse1  
Services: Mobility Services > Context-Aware Service > Wired > Wired Clients

1 Search results for Wired Clients with <IP,User Name,MAC,VlanId> matching '1302280\_user1'

MAC Address *	IP Address	Username (802.1x)	Serial Number	State	Switch IP Address	Port Type	Slot	Module	Port	VLAN Id
00:15:58:32:c2:85	10.1.91.238	1302280_user1		Connected	10.1.96.41	1Gbit	1	0	3	88

Wired Clients: "00:15:58:32:c2:85": mse1  
Services: Mobility Services > Context-Aware Service > Wired > Wired Clients

2

Device Information	Port Association	Civic Address	Advanced
Name		ABC Corporation	
Street		Kimball Bridge Road	
House Number		9300	
House Number Suffix		-	
Address Line 2		A10-2 Jack 1	
City		Alpharetta	
State		Georgia	
Postal Code		30022	
Country		US	

Wired Clients: "00:15:58:32:c2:85": mse1  
Services: Mobility Services > Context-Aware Service > Wired > Wired Clients

3

Device Information	Port Association	Civic Address	Advanced
ELIN		19789363023	Road Branch
Floor		1	Road Sub-branch
Building		Building 300	Road Pre-modifier
Apartment		-	Road Post-modifier
Room		103	Leading Street Direction
Place Type		Shipping & Receiving	Street Trailing Suffix
Neighborhood		-	Street Suffix
Landmark		-	Postal Community Name
Seat		1C/1-3	Post Office Box
Additional Code		-	City Division
Road		-	County
Road Section		-	Fulton

- **Location of Host Servers:**

- **Location of Physical Servers**—In addition to wired end user devices such as laptops and desktop computers, Context-Aware Services for wired endpoints can also be used to efficiently track the location of physical servers and even the virtual machines that operate on these physical servers. Locating a specific physical server among a variety of often identical servers in massive and often monolithic computing environments can be a daunting task. Using Context-Aware Services and at least one known characteristic of the server (MAC address, IP address, 802.1x username, or VLAN ID), the location of the server can be quickly and efficiently retrieved using WCS and the MSE. The MSE can provide a very precise location for the server, right down to a civic address consisting of the building, floor, room, and even rack and occupied slot number where the server is housed. This is a tremendous benefit in large computing environments, where hours of employee productivity could be wasted attempting to locate server hardware whose placement records have not been kept up to date. A good example of how this can be applied is in ensuring that decommissioned physical hardware is accurately located, cleansed of all sensitive enterprise information, and transferred to its appropriate final destination. Context-Aware Services can help ensure that decommissioned hardware is not somehow returned to service without proper authorization. Such “cradle-to-grave” tracking in the enterprise is especially useful in environments where physical servers may be leased and must be decommissioned and returned to leasing companies at the conclusion of their lease and not misappropriated for other uses within the enterprise. This is an all too common occurrence that can result in unnecessary leasing penalties and added expense.
- **Location of Virtual Machines**—Locating a mislabeled physical server in an environment when there might be hundreds of other identical servers may seem like the quintessential “needle in the haystack” experience. Imagine then the challenge of searching for a specific virtual machine within this huge inventory of physical hardware. This might be the case if a virtual machine is relocated to a new physical host, with the same virtual IP address and MAC address that it has used elsewhere, but with no record of its movement. Context-Aware Services can prove useful to the harried administrator faced with resolving this situation quickly and efficiently. For example, a misplaced virtual machine operating on a VMware ESX physical host can be located using WCS and the MSE by searching on characteristics uniquely associated with the virtual machine, such as its virtual IP address or MAC address. If the virtual machine is moved from one physical host to another (which implies movement from one Catalyst switch port to another), Context-Aware Services will track its movement and its new location as if it were just another wired device. The virtual machine will appear in its new location using the civic location information associated with its current physical host. Because of its ability to search across all of the Catalyst switches attached to an entire server farm, Context-Aware Services can be used to locate a virtual machine, right down to a civic address consisting of the building, floor, room, and even rack and slot number of where the current parent physical server to the virtual machine is housed. No matter which connected physical host the virtual machine has been relocated to, Context-Aware Services for wired devices can contribute significantly to saving employee time and helping increase productivity.
- **Catalyst Switch Chassis Location**—In addition to quickly locating end user devices, physical servers, and even virtual machines, Context-Aware Services can also be used to quickly query the location of a deployed Catalyst switch chassis itself. By assigning civic location characteristics regarding its deployed location to the out-of-band management port of a Catalyst switch, a network administrator can use WCS and the MSE to query the location of the switch itself. If defined accurately and maintained judiciously, this capability may reduce the need to maintain separate lists of deployed switch locations in many cases.

Figure 4-4 provides a visual example of how context-aware wired endpoint location can be used to locate a Catalyst switch chassis in the Unified Access network. Imagine that a rookie network technician or administrator, unfamiliar with a large enterprise network, needs to quickly determine the location of a context-aware Catalyst 3750X switch with IP address 192.168.84.30. They can use the wired endpoint location capabilities of Cisco Context-Aware Services by simply finding this

switch in the list of wired switches in WCS and clicking on the civic location tab associated with it. As long as the civic location information assigned to the management port interface is kept accurate and up to date, all of the information to identify the rack slot, rack number, room, floor, and building associated with the 3750X switch can be obtained from the MSE via WCS. There is no need for the administrator or technician to query another database for the deployed location of the 3750X switch. Using the sequence of events shown in Figure 4-4, we see that it is simply a matter of a few mouse clicks in WCS to determine that the switch in question is located in rack 12, slot 2, in room 102 on the first floor of building 300. This building is operated by ABC Corporation for the purposes of Shipping and Receiving. It is located at 9300 Kimball Bridge Road, in the city of Alpharetta, Fulton County, Georgia, postal zip code 30022.

Context-Aware Services help to integrate under one roof not only the location information for deployed wired endpoints, but for the switch chassis itself, thereby placing all of it at the fingertips of the WCS user in an efficient manner. This reduces response times, increases productivity, and can help reduce the overall costs associated with maintaining the enterprise network.

**Figure 4-4 Locating a Catalyst Switch Chassis**

**Step 1: Wired Switches**

Wired Switches: sh-mse3300  
Services > Mobility Services > sh-mse3300 > Context Aware Service > Wired > Wired Switches

IP Address	Serial Number/UDI	ELIN	Civic Address
192.168.84.30	FDO1437P0XG / -		ABC Corporation 1 Building 300 9300 Kimball Bridge Road Alpharetta Georgia 30022 US
192.168.84.23	FDO1439K0K0G / -		ABC Corporation 1 BXB-300 300 Beaver Brook Boxborough MA 01719 US

**Step 2: Switch Information**

Wired Switches: 192.168.84.30: sh-mse3300  
Services > Mobility Services > sh-mse3300 > Context Aware Service > Wired > Wired Switches

Switch Information	Switch Ports	Civic	Advanced
IP Address	192.168.84.30		
MAC Address	-		
Serial Number/UDI	FDO1437P0XG / -		
Model Number	WS-C3750X-24P		
Software Version	Cisco IOS Software, C3750E Software (C3750E-UNIVERSALK9-M), Version 12.2(55)SE1, RELEASE SOFTWARE (fc1) Technical Support: http://www.cisco.com/techsupport Copyright (c) 1986-2010 by Cisco Systems, Inc. Compiled Thu 02-Dec-10 06:08 by prod_rel_team		
ELIN			
Client Count	Total Clients: 2 Connected: 2 Disconnected: 0 Unknown: 0		

**Step 3: Civic**

Wired Switches: 192.168.84.30: sh-mse3300  
Services > Mobility Services > sh-mse3300 > Context Aware Service > Wired > Wired Switches

Switch Information	Switch Ports	Civic	Advanced
Name		ABC Corporation	
Street		Kimball Bridge Road	
House Number		9300	
House Number Suffix		-	
Address Line 2		cr22-3750s-LB	
City		Alpharetta	
State		Georgia	
Postal Code		30022	
Country		US	

**Step 4: Advanced**

Wired Switches: 192.168.84.30: sh-mse3300  
Services > Mobility Services > sh-mse3300 > Context Aware Service > Wired > Wired Switches

Switch Information	Switch Ports	Civic	Advanced
ELIN	-		Road Branch -
Floor	1		Road Sub-branch -
Building	Building 300		Road Pre-modifier -
Apartment	-		Road Post-modifier -
Room	102		Leading Street Direction -
Place Type	Shipping & Receiving		Street Trailing Suffix -
Neighborhood	-		Street Suffix -
Landmark	-		Postal Community Name -
Seat	-		Post Office Box -
Additional Code	rack 012/002		City Division -
Road	-		County - Fulton
Road Section	-		

## Cisco Context-Aware Components

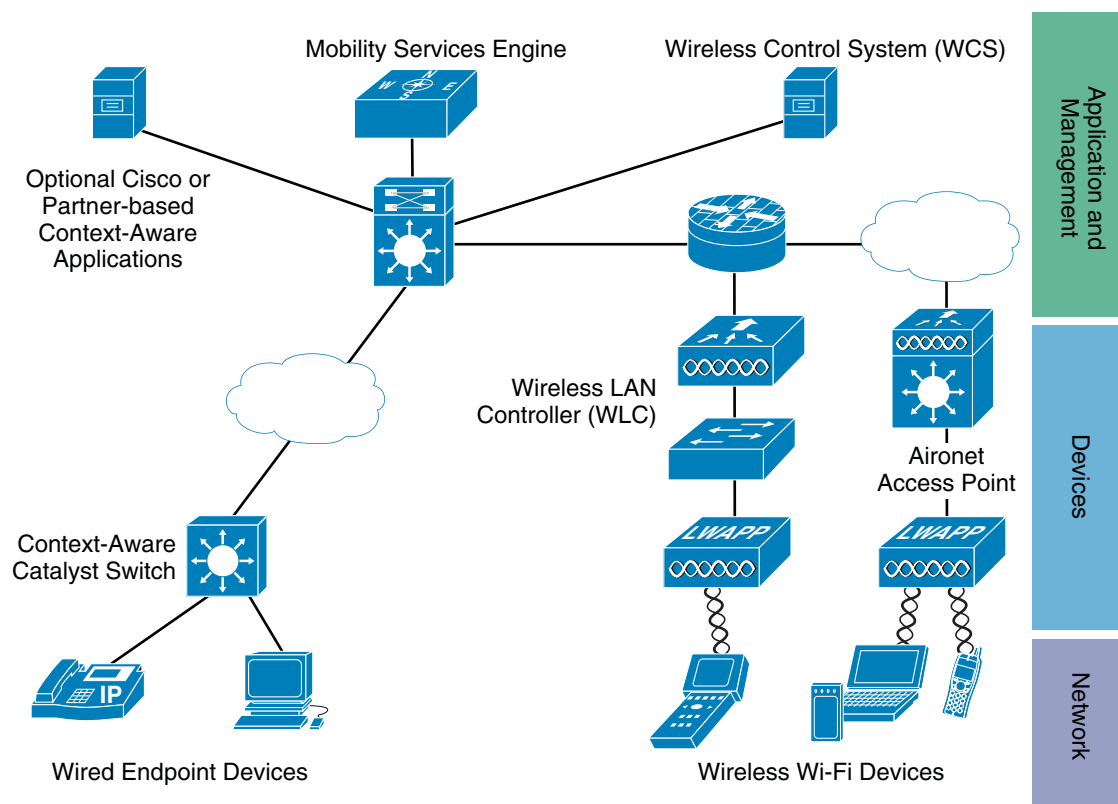
The components used to implement Cisco Context-Aware Services are shown in [Figure 4-5](#). Context-Aware Services are designed to record contextual information (including location) as it pertains to wired and wireless devices in an enterprise and maintain a repository of those locations in a central server. The overall solution has four fundamental components:

- Cisco 3300 Series Mobility Services Engine (MSE)
- Cisco Wireless Control System (WCS)
- Cisco Catalyst context-aware Ethernet switches for wired client access
- Cisco Wireless LAN Controllers for wireless client access

Included in [Figure 4-5](#) are optional context-aware applications. These might be value-added applications supplied by Cisco Technology Development Partners, such as RedSky and their E911 WiFi solution suite ([http://www.redsky911.com/e911\\_products/e911\\_manager/wifi\\_e911/cisco/](http://www.redsky911.com/e911_products/e911_manager/wifi_e911/cisco/)). Or these might be interfaces to other Cisco network services that represent other pillars of the Borderless Network Architecture, such as integrated security services or Unified Communications and Unified Presence applications.

It should be noted at this point that while [Figure 4-5](#) includes wireless LAN components, the majority of this chapter is focused on the use of Context-Aware Services in the Unified Access design as it pertains to wired endpoints.

**Figure 4-5 High Level View of Cisco Context-Aware Services**



290442

## Wired Endpoint Devices

These are wired Ethernet devices that interact with the network and whose location and other contextual parameters can be monitored by Context-Aware Services. Wired devices are equipped with an Ethernet interface and are attached to a Cisco Catalyst Ethernet switch (such as the 3560, 3750, 4500, and 4900 series) that is participating in Context-Aware Services. Via the Mobility Services Engine's API, civic and ELIN location information received from context-aware Catalyst switches can be provided to WCS, other Cisco Borderless Network services components, or Cisco partner-developed context-aware applications.

## Cisco Unified Access Network

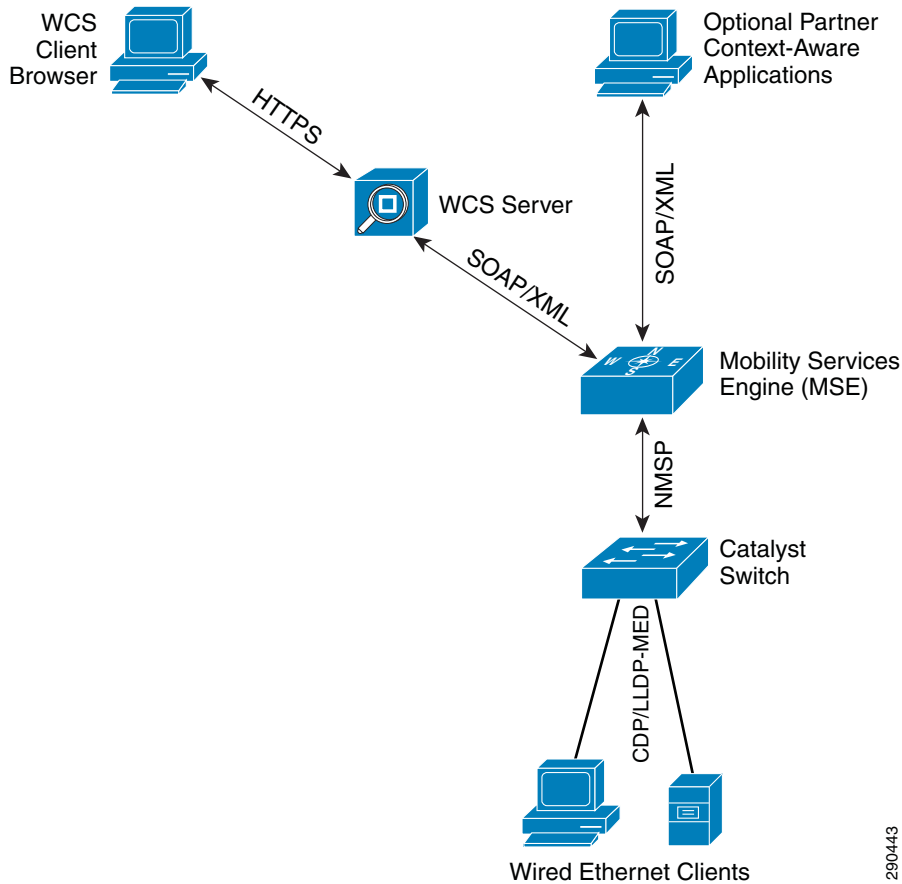
This multipurpose network contains the infrastructure required to address converged data, voice, and video requirements, as well as providing the foundation for Context-Aware Services. While the network may be composed of various routers, switches, firewalls, and so on, the main wired infrastructure network component actively participating in the passage of information back to the MSE is the context-aware Catalyst Ethernet switch. These switches support the specification of civic address and emergency location identification number (ELIN) information for each switch port.

Context-aware switches pass information for all attached devices to the Mobility Services Engine via communications sessions established between the two partners. This information may include the physical mailing or street address location associated with the attached device (the civic address) as well as other information such as the IP address, MAC address, port, VLAN, and 802.1x user name. Typically, this information is obtained using switch features such as IEEE 802.1x, Dynamic Host Configuration Protocol (DHCP) snooping, Dynamic Address Resolution Protocol (ARP) Inspection (DAI), and IP Source Guard. Additionally, if the end device supports CDP or LLDP-MED protocols, additional information, such as the version and serial number, can also be sent to the MSE.

Refer to the following data sheets for more information about the Cisco Catalyst switches and their ability to participate in Cisco Context-Aware Service as part of the Cisco Borderless Experience:  
[http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps6406/data\\_sheet\\_c78-584733\\_ps10744\\_Products\\_Data\\_Sheet.html](http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps6406/data_sheet_c78-584733_ps10744_Products_Data_Sheet.html)

## Management and Application Interaction

Figure 4-6 provides an illustration of the protocol interaction between the aforementioned Context-Aware Service components used with wired endpoints.

**Figure 4-6 Context-Aware Component Interaction**

- **Cisco Mobility Services Engine (MSE)**—The Cisco 3300 Mobility Services Engine can execute multiple independent services in support of Unified Access network infrastructures. Those services typically provide high-level capabilities, such as Cisco Context-Aware Services or the Wireless Intrusion Protection Service (wIPS). The Mobility Services Engine records contextual information for both wired and wireless assets on a continuous basis. It can do this by having the wired and wireless network infrastructure devices (WLAN controllers and context-aware Catalyst switches) send updated contextual information for attached assets to the MSE via NMSP as changes occur.

The Unified Access network communicates with the MSE using the Cisco Network Management Services Protocol (NMSP), which is a Cisco-defined protocol used for secure communication between the MSE and other context-aware network infrastructure components. The switch sends location and attachment tracking information for its connected devices to the MSE. The switch notifies the MSE of device link up and link down events via NMSP location and attachment notifications.

The MSE starts the NMSP connection to the switch, which opens a server port. When the MSE connects to the switch, there are a set of message exchanges to establish version compatibility and service exchange information followed by location information synchronization. After connection, the switch periodically sends location and attachment notifications to the MSE upon the expiration of an internal timer. Any link-up or link-down events detected during an interval are aggregated and sent at the end of the interval.

290443

When the switch determines the presence or absence of a device on a link-up or link-down event, it obtains the client-specific information, such as the MAC address, IP address, and 802.1x username. If the client is LLDP-MED or CDP capable, the switch obtains the serial number or UDI via these protocols.

Depending on the device capabilities, the switch obtains the following client information at link up:

- Slot and port on which connection was detected
- Detected client MAC address
- Detected IP address
- 802.1X username (if applicable)
- Device category (wired station)
- Device state (add new)
- Serial number or UDI
- Model number
- Time in seconds since the switch first detected the association

Depending on the device capabilities, the switch obtains this client information at link down:

- Slot and port on which connection was disconnected
- Detected client MAC address
- Detected client IP address
- 802.1X username (if applicable)
- Device category (wired station)
- Device State (delete)
- Serial number or UDI
- Time in seconds since the switch first detected the disassociation

When the switch shuts down, it sends an attachment notification with the state delete and the IP address before closing the NMSP connection to the MSE. The MSE interprets this notification as disassociation for all the wired clients associated with the switch.

When location information is changed on the switch, the switch sends an NMSP location notification message that identifies the affected ports and the changed address information.

After the MSE records the current location of a wired endpoint device, it provides this information to WCS, external Cisco-partner supplied context-aware applications, or other Cisco network components. The two primary approaches used for information retrieval from the MSE via the API are:

- Queries—The external system typically sends a query asking for a device's location, including optional query criteria such as to return data for only a specific MAC address. The MSE would respond with the answer immediately upon receipt of the query.
- Subscribe and publish—The external system typically registers an event subscription for device location based on a set of criteria such as changes in location or changes in location beyond a prescribed area.

Context-Aware Services software is capable of servicing up to a maximum of 18,000 simultaneously tracked devices per MSE-3350 or MSE-3355 appliance and 2,000 simultaneously tracked devices per MSE-3310 appliance. In Release 7.0, all civic and ELIN location information is configured

locally at the switch, and these changes are propagated to the MSE via NMSP. NMSP is used between the switches and the MSE to maintain synchronization and alert the switch as to the connection or disconnection of devices.

Refer to the following data sheet for more information regarding the Cisco 3300 Series Mobility Services Engines:

[http://www.cisco.com/en/US/prod/collateral/wireless/ps9733/ps9742/data\\_sheet\\_c78-475378.html](http://www.cisco.com/en/US/prod/collateral/wireless/ps9733/ps9742/data_sheet_c78-475378.html)

- **Wireless Control System (WCS)**—The Cisco Wireless Control System is a management platform that contains a context-aware client application interacting directly with the Mobility Services Engine. In this role the WCS provides access to the contextual information contained on the MSE using the MSE's context-aware application programming interface (API). WCS presents this information to the user in either a graphical or tabular format. The Cisco WCS also serves as a control client to be used to configure operational parameters on the MSE.

For more detailed information on the Cisco Wireless Control System management server and its capabilities, including its ability to serve as a context-aware application client to manage the MSE, refer to the documentation at:

[http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/product\\_data\\_sheet0900aecd802570d0.html](http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/product_data_sheet0900aecd802570d0.html).

- **Other Context-Aware Applications**—Although WCS provides an extremely powerful interface to the contextual information stored on the MSE, access to this information is not limited to only WCS. In fact, once contextual information has been captured, calculated, and stored in its resident databases by the Context-Aware software module on the MSE, it can be made available not only to WCS but to any business application that requiring access via the API, which is based on the Simple Object Access Protocol (SOAP) and XML protocol.

Cisco-developed applications have access to this API of course, but so do context-aware applications developed by Cisco Technology Development partners. Access to this API is available to any Cisco technology partner and allows full integration into enterprise business processes. When planning for what can be accomplished in your organization with the information contained within the context-aware service on the MSE, it is important to remember that context-aware applications developed by Cisco Technology Development partners often target specific industry needs, delivering value, functionality, and enhanced capabilities that are often not available from any other source.

For more information on the Cisco Context-Aware Services API, refer to:

<http://developer.cisco.com/web/contextaware/home>.

## Network Mobility Services Protocol (NMSP)

The Network Management Service Protocol (NMSP) was designed to define intercommunication between Mobility Service Engines and network access controllers over a switched or routed IP network. An access controller can provide network access for either wired or wireless endpoints. Within the scope of this Unified Access chapter, access controllers are represented by context-aware Cisco Catalyst Ethernet switches.

NMSP is a two-way protocol that can be run over a connection-oriented or a connectionless transport. Context-aware switches can use NMSP to communicate with one or more MSEs. NMSP is based upon a bidirectional system of requests and responses between the MSE and access controllers.

**Note**

It is important to understand that the failure of an NMSP session has no direct impact on the ability of a Catalyst switch to pass normal client voice, data, and video session traffic to applications on the network. In other words, a failed NMSP session to a Catalyst switch may affect the ability of the MSE to provide updated contextual information for that switch and its resources, but it does not affect the ability of the attached devices to log on to applications residing on the network.

NMSP uses Transport Layer Security (TLS) and TCP port 16113 on the Catalyst switch. The MSE will initiate the connection to the Catalyst switch, after which messages may be transmitted in either direction. The TCP port (16113) that the Catalyst switch and MSE use for communication must be open on any firewall that exists between the switch and the MSE.

NMSP provides for a keep-alive protocol mechanism that allows either partner in a NMSP session to determine if the adjacent partner is still active and responsive. Should an MSE fail, the Catalyst switch will try to contact another MSE with which to communicate. If the Catalyst switch fails, all Context-Aware Services being provided to that Catalyst switch are disabled until the switch once again becomes active and re-establishes its NMSP session.

The MSE and the Catalyst switch use Echo Request and Echo Response control messages to maintain an active channel of communication so that the data messages can be sent. The Echo Request message is a keep-alive mechanism that allows either NMSP session partner to determine if the other partner remains active and responsive. Echo Requests are sent periodically (upon expiration of a heartbeat timer) by the MSE or its session partner to determine the state of the NMSP session. When the Echo Request is sent, a NeighborDeadInterval timer is started. The NeighborDeadInterval timer specifies the minimum time a session partner must wait without having received Echo Responses to its Echo Requests before the other session partner can be considered non-responsive and the NMSP session is placed in an idle state.

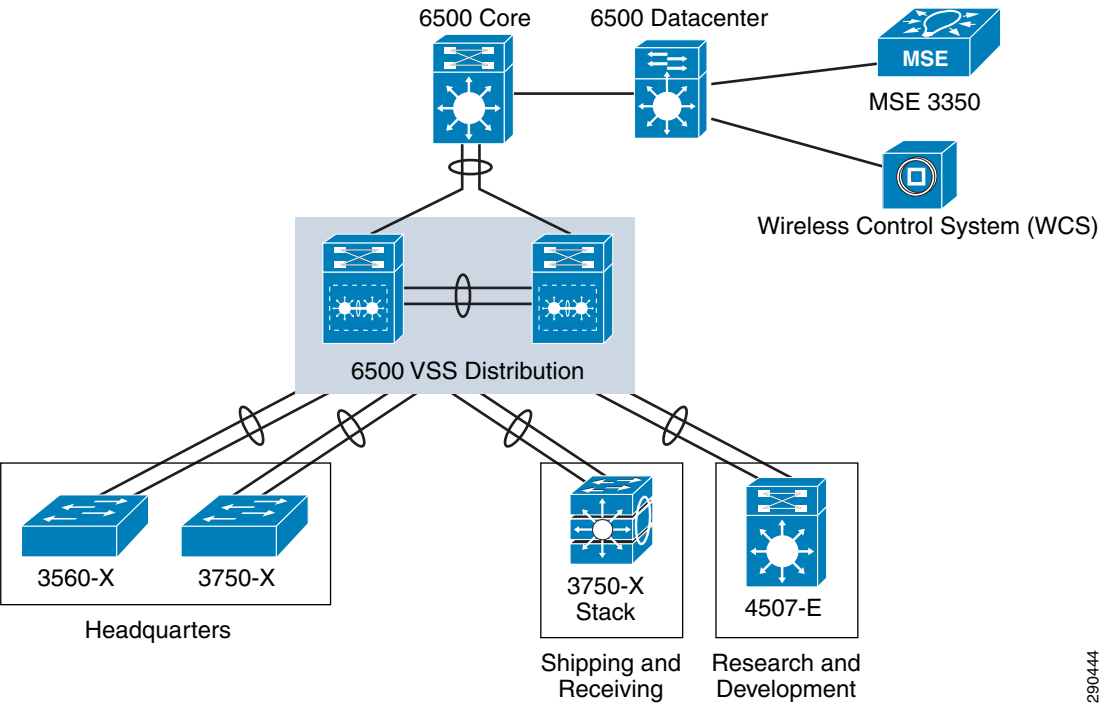
**Note**

Proper validation of certificates between context-aware service components requires the participants to possess sane clocks (clocks whose configured time does not differ from one another by large amounts). In order to facilitate this, it is highly recommended that the clocks in the MSE, WCS, and any Catalyst switches participating in Context-Aware Services be synchronized to a common time base using the Network Time Protocol (NTP). The lack of clock sanity among context-aware components in the network can cause NMSP sessions to fail if the configured date and time fall outside of the certificate's expiration.

## Context-Aware Services in Unified Access

Figure 4-7 illustrates the test bed used to validate the use of Cisco Context-Aware Services with wired endpoints in the Unified Access design.

Figure 4-7 Test Bed Configuration



290444

Table 4-1 Sample Civic Location Information for Headquarters 3750X

Identifier	bldg_210
County	Fulton
Street Group	Westside Parkway
Street number	1100
Name	ABC Corporation
Building	Building 210
Type of place	Headquarters
City	Alpharetta
State	Georgia
Postal code	30009
Country	US

Table 4-2 Sample Civic Location Information for Shipping and Receiving 3750X Stack

Identifier	bldg_300
County	Fulton
Street Group	Kimball Bridge Road
Street number	9300
Name	ABC Corporation

**Table 4-2 Sample Civic Location Information for Shipping and Receiving 3750X Stack**

Building	Building 300
Type of place	Shipping & Receiving
City	Alpharetta
State	Georgia
Postal code	30022
Country	US

**Table 4-3 Sample Civic Location Information for Research & Development 4507R**

Identifier	1/1
County	Fulton
Street Group	Webb Bridge Road
Street number	3680
Name	ABC Corporation
Building	Building 400
Floor	1
Room	101
Type of place	Research & Development
Seat	1A/1
City	Alpharetta
State	Georgia
Postal code	30023
Country	US

The access layer was configured so as to simulate a enterprise named “ABC Corporation” with local access switches spread among departments situated at various campus locations. Thus, we see in [Figure 4-7](#) that ABC Corporation possesses shipping and receiving, research and development, and headquarters locations, all of which have deployed context-aware Catalyst switches. These switches establish NMSP sessions with the MSE located at the central data center. A WCS has also been deployed at the data center and access to the WCS (and hence to the user GUI for Context-Aware Services) is available to authorized users throughout the enterprise.

The test network is based on a three-tier routed access campus design, with Layer 2 and Layer 3 switches at the access layer and Layer 3 switches in the distribution and core layers. As seen in [Figure 4-7](#), the Layer 2 access was provided in this design by a Catalyst 3750X switch. The Layer 3 switches used at the access layer were:

- Catalyst 3560X
- Catalyst 3750X, as a nine member switch stack
- Catalyst 4507
  - Chassis Type: WS-C4507R-E
  - Sup 6-E 10GE WS-X45-SUP6-E

- 10/100/1000BaseT (RJ45)V WS-X4548-GB-RJ45V used for client connections
- SFP, 10/100/1000BaseT (RJ45)V WS-X4506-GB-T used for uplink to distribution

All access switches were loaded with k9 IOS cryptographic images at a code train level supporting NMSP and civic address port location interface subcommands.

The distribution and core layers were composed of Catalyst WS-C6509-E switches with Sup 720 10GE. The distribution layer was configured to operate in VSS mode.

In the test bed, OSPF was used as the routing protocol. The campus core was configured to represent OSPF Area 0 and the access layer was configured to use stub areas.

In accordance with general industry network design best practices, out-of-band (OOB) network management was used where feasible. For example, the FastEthernet interfaces specifically designated for OOB network management on the Catalyst 3560 and 3750 switches were utilized for all NMSP and SNMP traffic to and from the switch. On the other hand, the management interface on the Catalyst 4507 was not used in this design. This interface is by default a VRF interface and does not currently support NMSP at this time. Therefore, all inbound and outbound NMSP to this switch was conducted in-band.

The test bed configuration was based foundationally upon the concepts established in the *Borderless Campus Design Guide 1.0*, which can be found at:  
[http://www.cisco.com/en/US/partner/docs/solutions/Enterprise/Campus/Borderless\\_Campus\\_Network\\_1.0/Borderless\\_Campus\\_1.0\\_Design\\_Guide.html](http://www.cisco.com/en/US/partner/docs/solutions/Enterprise/Campus/Borderless_Campus_Network_1.0/Borderless_Campus_1.0_Design_Guide.html)

A variety of wired endpoints were included in our testing:

- Windows XP PCs
- Windows 7 PCs
- Red Hat Enterprise Linux AS release 4
- FreeBSD 6.1
- Cisco 9971 IP phones
- Cisco Digital Media Player DMP-4310
- CIVS-IPC-2500 Video Surveillance Camera
- CIVS-IPC-4500 Video Surveillance Camera
- Cisco AP-3502 Access Point (AP3G1-K9W8-M)
- Cisco AP-1142N Access Point (C1140-K9W8-M)

These endpoints were moved between switch ports repeatedly, both on a intra-switch chassis and inter-switch chassis basis. Acceptable behavior of endpoints moving between switch ports was based on the following:

- For access switch ports with spanning-tree-portfast configured, device attachment status and initial location information should be reflected on WCS within a time interval not exceeding the default nmosp-attachment-interval of 30 seconds plus a maximum 10 second processing delay window, for a total time equal to or less than 40 seconds. Optimal performance under these circumstances would assume a processing delay window of zero seconds, for a total time equal to or less than 30 seconds.
- For access switch ports without spanning-tree-portfast configured, device attachment and initial location information should be reflected on WCS within a time interval not exceeding the switch's maximum spanning tree port transition delay plus the default nmosp-attachment-interval of 30 seconds plus a 10 second processing delay window.
- When evaluating the ability of the system to reflect a disconnected client, the spanning tree transition time is not applicable. Therefore, client disconnect status should be reflected on WCS within 40 seconds or less.

The Cisco DMP-4310 was used to verify that our test bed configuration could pass civic location information contained in CDP and LLDP-MED TLVs (type, length, and value attributes) to a compatible Cisco medianet endpoint device. This is an area of great development and excitement and is leading the way to allow for location-assisted auto-configuration of medianet devices in the future. More information on the DMP-4310 and its ability to receive civic location information from Cisco Catalyst switches may be found in the *Medianet Reference Guide*:

[http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/Medianet\\_Ref\\_Gd/chap7.html#wp1216490](http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/Medianet_Ref_Gd/chap7.html#wp1216490).

Two VMWare ESXi 4.0 Servers were used to allow for testing of virtual machines being moved from one physical host to another. The virtual machines tested included:

- Windows XP
- Ubuntu 10.4 Server
- CentOS 5.4

## Component Capacities

### Mobility Services Engine

Each Cisco Mobility Services Engine has a maximum combined device tracking capacity, which is a “hard” limit that is dictated by the licensing purchased for the Context-Aware software as well as the presence of any other applications on the MSE. Once a Mobility Services Engine has reached its maximum tracking capacity, any new devices the MSE becomes aware of that exceed that limit are simply not tracked. It is important to note that while this section discusses the maximum device tracking limits for the MSE, licenses can be purchased supporting device limits significantly lower than these maximums. Refer to the MSE Licensing and Ordering Guide

([http://www.cisco.com/en/US/prod/collateral/wireless/ps9733/ps9742/data\\_sheet\\_c07-473865.html](http://www.cisco.com/en/US/prod/collateral/wireless/ps9733/ps9742/data_sheet_c07-473865.html)) for more information regarding the various client tracking capacities available for the MSE.

For Release 7.0, the maximum device tracking capacities for the MSE when using only the Context-Aware Service are shown in Table 4-4.

**Table 4-4** Maximum Device Tracking Capacities

Mobility Service Engine	Maximum Tracked Device Capacity
MSE-3350/3355	18,000
MSE-3310	2,000

If you intend to use the MSE to deliver other services in addition to Context-Aware, the maximum capacities shown above will be reduced. See the *MSE Licensing and Ordering Guide* ([http://www.cisco.com/en/US/prod/collateral/wireless/ps9733/ps9742/data\\_sheet\\_c07-473865.html](http://www.cisco.com/en/US/prod/collateral/wireless/ps9733/ps9742/data_sheet_c07-473865.html)) for information on Context-Aware maximum tracked device capacities when used with other co-resident services.

When working within these maximum capacities, it is important to note that further category-specific limits can be instituted via the MSE configuration. This allows, for example, a maximum capacity of 2,000 tracked devices on a MSE-3310 to be further limited as 400 wired endpoints, 600 wireless endpoints, 500 RFID tags, 250 rogue access points and rogue clients, and 250 interferers. Partitioning the maximum tracking capacity of the context-aware software in this manner prevents any single device category from consuming more than its allotted share of the maximum tracking capacity. This is a very

important consideration when designing context-aware solutions that incorporate awareness of rogue access points, rogue devices, or interferers, since the anticipated size of each of these device categories is not always predictable. Failure by the context-aware solution designer to institute reasonable limits on the total number of rogues and interferers tracked could result in premature exhaustion of MSE capacity, thereby preventing the tracking of other categories of devices (such as wired or wireless client endpoints).

The Context-Aware System Performance chapter of the *Mobility Services Engine Context Aware Deployment Guide*

([http://www.cisco.com/en/US/products/ps9742/products\\_tech\\_note09186a00809d1529.shtml#casysperf](http://www.cisco.com/en/US/products/ps9742/products_tech_note09186a00809d1529.shtml#casysperf)) also points out that a single MSE can support up to 500 total NMSP connections. This has been validated to 100 total NMSP connections and includes not only NMSP sessions to Catalyst switches, but also to any WLAN controllers participating in Context-Aware Services with this MSE.

## Wireless Control System (WCS)

As mentioned earlier, WCS interacts with the MSE as both a Context-Aware Services application client as well as a control client. When used to locate endpoints in the network via the MSE, WCS does not execute any location algorithms itself. All location calculations are handled by the MSE. There are no NMSP sessions that are established between WCS and Catalyst switches or WLAN controllers. Because of this client-server relationship between WCS and the MSE, there are no direct constraints on the maximum number of tracked devices that are imposed by WCS.

There are, however, a few indirect constraints that currently exist of which the network designer should be aware:

- A single WCS can manage and interact with multiple MSEs. Cisco officially supports a single WCS managing and interacting with up to five (5) MSEs. While defining more than five MSEs to a single WCS is possible, Cisco has not validated this configuration.
- A single MSE should be managed by only one WCS. In other words, there is a 1:1 mapping that should be maintained between an MSE and the number of WCS systems attempting to manage and interact with that MSE. Care should be taken not to confuse this consideration with the bullet above.
- It is important to note that in release 7.0, non-root WCS virtual domain users cannot access WCS functions listed under Services > Mobility Services. This includes wired switch and wired endpoint device location. Therefore, since wired devices attached to context-aware Ethernet switches are displayed using Services > Mobility Services > Context Aware Service > Wired > Wired Clients, only users that are assigned to the WCS root virtual domain are able to display context-aware information for these devices.

Refer to Understanding Virtual Domains as a User, *WCS Configuration Guide 7.0*

([http://www.cisco.com/en/US/docs/wireless/wcs/7.0/configuration/guide/7\\_0virtual.html](http://www.cisco.com/en/US/docs/wireless/wcs/7.0/configuration/guide/7_0virtual.html)) for a complete list of network resources that are not available in non-root virtual domains.

## Catalyst LAN Switch

Each Catalyst LAN switches participating in Context-Aware Services within the Unified Access network design will participate in an NMSP session with the Mobility Services Engine. An important consideration to keep in mind is that while a single MSE can support concurrent NMSP sessions to many Catalyst LAN switches in the network, there is a 1:1 relationship between each participating Catalyst switch and the number of MSEs to which it communicates via NMSP.

## Integration with the Unified Access Network Architecture

### Clock Synchronization

All components participating in wired endpoint location using Context-Aware Services should have their internal clocks synchronized to a common time source. This includes the Mobility Services Engine, WCS, WLAN controllers, and any Catalyst switches. It is recommended that Coordinated Universal Time (UTC<sup>2</sup>) be utilized for this purpose. Because reliable certificate authentication relies on time-based consistency between participating components, it is important to ensure that context-aware components are time-synchronized throughout the network. In addition, having components synchronized to a common time source can help facilitate troubleshooting, especially when having to review events occurring in the logs of different network components. Any such log output, when coupled with accurate and consistent time stamps, will appear much more logical and will facilitate problem resolution.

The recommended approach to maintain time synchronization across all components is through the use of the Network Time Protocol (NTP).

### NTP Configuration of the Mobility Services Engine

Configuration of the NTP server addresses used by the MSE is handled during installation and the execution of the MSE automatic configuration script. An excerpt of that script is shown below. Detailed information regarding the automatic configuration script can be found in the Automatic Installation Script section of the *Mobility Services Engine Getting Started Guide*

([http://www.cisco.com/en/US/docs/wireless/mse/3350/quick/guide/mse\\_qsgmain.html#wp1057105](http://www.cisco.com/en/US/docs/wireless/mse/3350/quick/guide/mse_qsgmain.html#wp1057105)) and in Appendix A of the *Mobility Services Engine Context Aware Deployment Guide* ([http://www.cisco.com/en/US/products/ps9742/products\\_tech\\_note09186a00809d1529.shtml#appena](http://www.cisco.com/en/US/products/ps9742/products_tech_note09186a00809d1529.shtml#appena)).

If you choose to enable NTP, the system time will be configured from NTP servers that you select. Otherwise, you will be prompted to enter the current date and time.

NTP is currently disabled.

Configure NTP related parameters? (Y)es/(S)kip/(U)se default (S)kip: Y

Enter whether or not you would like to set up the Network Time Protocol (NTP) for this machine.

If you choose to enable NTP, the system time will be configured from NTP servers that you select. Otherwise, you will be prompted to enter the current date and time.

Enable NTP (yes/no) no : yes

Enter NTP server name or address: <IP address or DNS name of NTP server>

Enter another NTP server IP address (or none) none: none

### NTP Configuration of the Wireless Control System (WCS) Server

Configuration of the internal clock and the specification of which NTP servers to use for periodic time synchronization must be performed on the WCS server using the time and date capabilities of the WCS host operating system in use (either Windows or Linux).

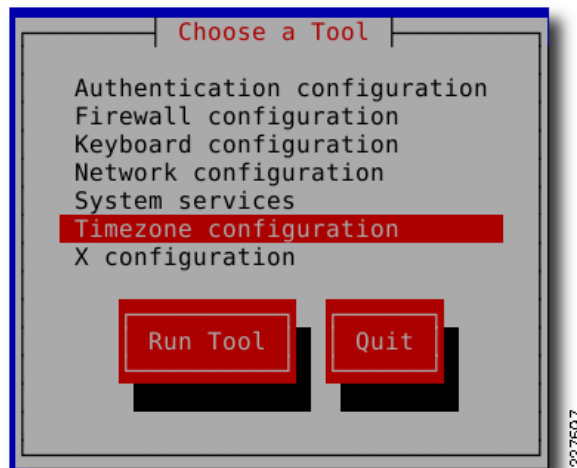
2. UTC may be considered as equivalent to Greenwich Mean Time (GMT) in this case.

## RHEL-Based WCS Server

For a Redhat Linux-based WCS server, login to the host OS as root and use the following procedure to synchronize the internal software clock to the NTP server, followed by synchronizing the software clock to the server's hardware clock, and then finally ensuring that synchronization is maintained by starting the ntpd client daemon:

- `clock`—Displays the current setting of the software clock.
- `/etc/init.d/ntpd stop`—Stops the ntpd client if it is already running.
- `ntpdate <ntp server name or address>`—Synchronizes the system software clock with the NTP server.
- `setup`—Brings up a setup utility that allows you to choose to set the time zone (shown in Figure 4-8).
- `hwclock --systohc`—Writes the software clock settings to the hardware clock.
- `/etc/init.d/ntpd start`—Starts the ntpd daemon to maintain clock synchronization going forward.<sup>3</sup>

**Figure 4-8** RHEL Setup Utility



**Note**

There are various other approaches that can be used to select the time zone on a Linux host system. The reader is encouraged to consult the Redhat documentation for methods involving the use of the TZ variable or symbolic links to the localtime file or a particular time zone file in the system's time zone directory.

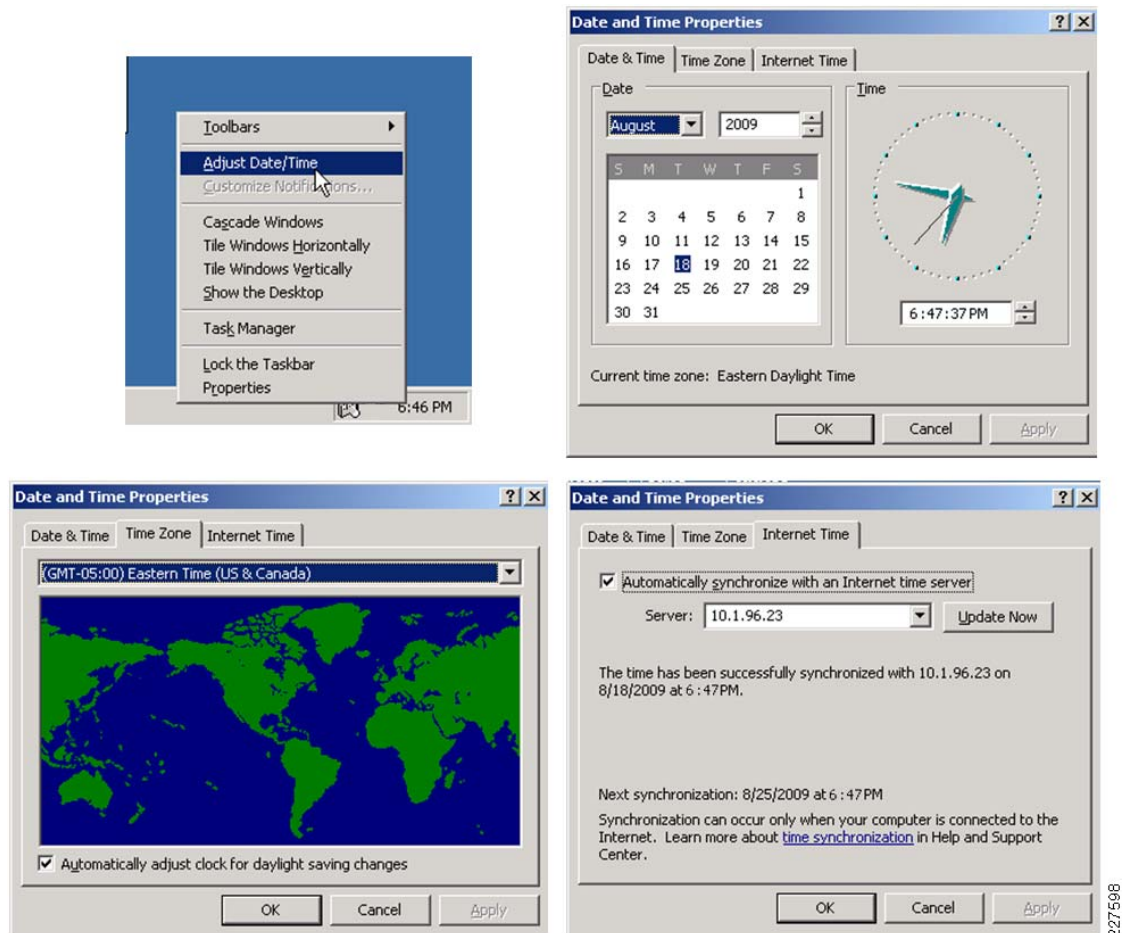
## Windows 2003-Based WCS Server

For a WCS server based on the Microsoft Windows 2003 Server OS, use the following procedure to synchronize and maintain the correct system time via the Windows Time service (see Figure 4-9):

1. Check **Settings>Control Panel >Administrative Tools>Services** for the Windows Time service and ensure that it has been started.
2. Right click on the **Task Bar** clock and select **Adjust Date/Time**.
3. If ntpd does not start as part of your system boot script, consider adding it using the command `chkconfig --add ntpd`.

3. Under the **Date & Time** tab, set the current date and clock time to the approximate time of your NTP server.
4. Set the Time Zone and Daylight Savings time selections appropriately.
5. Select the Internet Time tab, check the box to Automatically Synchronize With An Internet Time Server, type in the DNS name or address of your NTP server, and then click **Apply**.

**Figure 4-9** Setting Time and NTP Server on Windows 2003



227598

## NTP Configuration of Context-Aware Catalyst Ethernet Switches

In order to prevent any issues with certificate authentication and NMSP session initiation, Catalyst Ethernet switches participating in Context-Aware Services should be configured to utilize NTP in order to keep their clocks in synchronization with other context-aware components. NTP is configured similarly among the various switch models discussed in this chapter and the most comprehensive information on how to configure the NTP client within a Catalyst switch can usually be found in the configuration guide for the particular switch model. For example, the Catalyst 3560 NTP configuration is documented in the Configuring NTP section of the *Catalyst 3560 Switch Software Configuration Guide*

([http://www.cisco.com/en/US/docs/switches/lan/catalyst3560/software/release/12.2\\_55\\_se/configuration/guide/swadmin.html#wp1053923](http://www.cisco.com/en/US/docs/switches/lan/catalyst3560/software/release/12.2_55_se/configuration/guide/swadmin.html#wp1053923)).

Similar guidance regarding the use of NTP can be found in the configuration guides for the Catalyst 3750 and 4500 series.

It is best practice to ensure time synchronization of all network components whenever possible. However, from the perspective of Context-Aware Services in the Unified Access design, only those switches that are actually participating in an NMSP session with the MSE require mandatory clock synchronization.

## Context-Aware Service Parameters-Tracking

As mentioned earlier, Context-Aware Services can track up to a maximum of 18,000 licensed devices when using the MSE-3350/3355 hardware platform and up to a maximum of 2,000 licensed devices when using the MSE-3310 platform. The absolute limit on the number of devices that can be tracked is determined by the hardware platform used, the presence of any other applications co-residing on the MSE, and the level of licensing purchased. The WCS tracking parameters configuration panel (located at Services > Mobility Services > Context Aware Service > Administration > Tracking Parameters) allows the administrator to pre-determine just how much of the MSE's maximum licensed tracking capacity will be allocated towards the tracking of specific device categories. This is useful in order to allow the tracking of device categories such as rogue access points, rogue clients, or wireless interferers, but also limit these unpredictable categories such that an uncontrolled introduction of rogues or interferers is not allowed to consume all of the device tracking capacity on the MSE.

We can use the Context-Aware Service Tracking configuration to:

- Entirely enable or disable the tracking of wired or wireless client stations, asset tags, rogue access points, rogue clients, and interferers.
- Set limits on how much MSE tracked device capacity will be allocated to certain device categories. [Figure 4-10](#) provides us with an example of how this can be achieved, where the maximum number of rogue access points/clients and the maximum number of wireless interferers are capped at 2,000 devices each. No limit value has been placed on any other categories, which in effect means that the maximum number of wired and wireless devices tags tracked will be allowed to rise until the licensing limits are reached. But due to the limit values imposed on them, the MSE would never track more than 2,000 wireless rogues or interferers.



### Note

Any devices that are detected but excluded from tracking due to the enforcement of a tracking limit will be reflected in the “Not Tracked” device count column shown on the right side of the display. This is very useful to both the designer and the network administrator in that it allows for straightforward verification of license sufficiency post-deployment.

**Figure 4-10**      **Mobility Services Engine Tracking Parameters**

**Tracking Parameters**  
Services > Mobility Services > Context Aware Service > Administration > Tracking Parameters

The SNMP parameters and Polling Interval are applicable for Controller version 4.1 or below

**Tracking Parameters**

**Network Location Service Elements:** Licensed Limit = 12000

Enable	Tracking Parameters	Enable Limiting	Limit Value	Active Value	Not Tracked
<input checked="" type="checkbox"/>	Wired Clients	<input type="checkbox"/>	5	15	0
<input checked="" type="checkbox"/>	Wireless Clients	<input type="checkbox"/>	5	5	0
<input checked="" type="checkbox"/>	Rogue Clients and AccessPoints	<input checked="" type="checkbox"/>	2000	0	0
	<input type="checkbox"/> Exclude Adhoc Rogue APs				
<input checked="" type="checkbox"/>	Interferers	<input checked="" type="checkbox"/>	2000	0	0

**Asset Tracking Elements:** Licensed Limit = 3000

Enable	Tracking Parameters	Active Value	Not Tracked
<input checked="" type="checkbox"/>	Active RFID Tags	5	0

## Catalyst Switch Definition and Synchronization

To allow for proper tracking of the devices that may be registered or attached to them, context-aware Catalyst switches must be defined to WCS and then synchronized with the Mobility Services Engine.

Detailed information about adding Catalyst switch definitions to WCS using the WCS Configure > Add Ethernet Switches menu panel can be found in the *WCS Configuration Guide 7.0* ([http://www.cisco.com/en/US/docs/wireless/wcs/7.0/configuration/guide/7\\_0ctrlcfg.html#wp1089752](http://www.cisco.com/en/US/docs/wireless/wcs/7.0/configuration/guide/7_0ctrlcfg.html#wp1089752)).

Detailed information about synchronizing the MSE with context-aware Catalyst switches using the WCS Services > Mobility Services > Synchronize WCS and MSE(s) menu panel can be found in the Synchronizing Mobility Services Engines chapter of the *Context-Aware Service Configuration Guide 7.0* ([http://www.cisco.com/en/US/docs/wireless/mse/3350/7.0/CAS/configuration/guide/msecg\\_ch3\\_CAS.html#wp998995](http://www.cisco.com/en/US/docs/wireless/mse/3350/7.0/CAS/configuration/guide/msecg_ch3_CAS.html#wp998995)).

## Usage Considerations for Wired Endpoint Device Tracking

Beginning with release 6.0 of Cisco Context-Aware Services, civic and ELIN location attributes can be assigned to devices connected to Cisco Catalyst switch ports, such as the 3560, 3750, 4500, and 4900 series. As participants in Context-Aware Services, switches provide relevant contextual information for all the wired device endpoints that attach to them. These endpoints may include IP phones, PCs, host servers, access points, and so on. The NMSP protocol is used between the switches and MSE to deliver this contextual information to the MSE. Location information may include the physical street location address (also known as the civic address) as well as other information about endpoints such as their IP address, MAC address, port, VLAN, and 802.1x username. If the end device makes use of CDP or LLDP-MED, additional endpoint device information, such as the version number of its operating system and its hardware serial number, can also be sent to the MSE.

In the Unified Access design, redundancy is provided such that should a Catalyst switch stack member or modular component responsible for NMSP communications fail, the NMSP session will automatically recover once communication is re-established using a redundant component. Thus, in our test bed configuration, using the default NMSP parameter settings on the MSE, we observed that the total time

required for our nine member 3750X switch stack to recover from stack master failure and re-establish NMSP communications was under 60 seconds. Typically, about 50-55 seconds of this elapsed time was spent performing switch stack recovery and election of a new stack master, with only a three to seven second additional delay observed before the NMSP session itself was seen to be fully recovered and operational.

## Hardware and Software Requirements for Wired Device Tracking

As mentioned previously, at the current time wired device tracking is only performed on Catalyst switch hardware such as the 3560, 3750, 4500, and 4900 series.



### Note

Readers should note that cryptography-enabled (k9) switch images are mandatory in order to enable NMSP functionality in Catalyst switches.

Enabling the tracking of wired endpoint devices in a context-aware Unified Access network requires that the Mobility Services Engine be licensed for the expected number of wired devices you anticipate tracking. This is especially important in networks where wired device tracking is being utilized alongside the tracking of wireless devices, rogues, interferers, and RFID tags, as the license represents the upper limit on the total number of devices tracked.

Each context-aware switch that is enabled for wired device tracking in your network establishes one NMSP session to the MSE. Therefore, when enabling numerous switches for context-aware wired device tracking, it is recommended that you plan for the total number of MSEs that may be required to support the total number of NMSP sessions in your network. This is especially relevant in Unified Access network designs, where Context-Aware Services will likely be enabled for both wired Catalyst switches as well as wireless LAN controllers. In very large networks consisting of many switches and WLAN controllers, you may find that more than one MSE is required to accommodate all anticipated NMSP sessions. Although a single MSE is intended to support up to 500 NMSP sessions, Cisco scalability testing has validated this capability to an maximum of 100 simulated NMSP connections.

## Enabling Context-Aware Wired Device Tracking

In order to track wired devices on Catalyst switch ports, each switch whose devices we wish to track must be configured to enable NMSP and other important parameters and to contain the appropriate location information for each switch port. In addition, WCS must be configured to be aware of the context-aware switches in the network and to be able to communicate with them. WCS is also used to transmit information about the switches to the Mobility Services Engine and initiate the synchronization process.

A complete, step-by-step guide to configuring Catalyst switches and WCS for wired device tracking can be found in the *Context Aware Service Configuration Guide 7.0*

([http://www.cisco.com/en/US/docs/wireless/mse/3350/7.0/CAS/configuration/guide/msecg\\_ch7\\_CAS.html#wp1224011](http://www.cisco.com/en/US/docs/wireless/mse/3350/7.0/CAS/configuration/guide/msecg_ch7_CAS.html#wp1224011)).

In addition, the following chapters and documents provide valuable and detailed background information concerning the wired device tracking capability of Catalyst switches:

- Configuring LLDP, LLDP-MED, and Wired Location Service in the *Catalyst 3750 Switch Software Configuration Guide*  
[http://www.cisco.com/en/US/docs/switches/lan/catalyst3750/software/release/12.2\\_55\\_se/configuration/guide/swlldp.html](http://www.cisco.com/en/US/docs/switches/lan/catalyst3750/software/release/12.2_55_se/configuration/guide/swlldp.html)

- Configuring LLDP, LLDP-MED, and Wired Location Service in the *Catalyst 3560 Switch Software Configuration Guide*  
[http://www.cisco.com/en/US/docs/switches/lan/catalyst3560/software/release/12.2\\_55\\_se/configuration/guide/swlldp.html](http://www.cisco.com/en/US/docs/switches/lan/catalyst3560/software/release/12.2_55_se/configuration/guide/swlldp.html)
- Configuring LLDP and LLDP-MED in the *Catalyst 4500 Series Switch Software Configuration Guide*  
<http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/54sg/configuration/guide/swlldp.html#wp1097119>

Readers are reminded that:

- NMSP is disabled on Catalyst switches by default and must be explicitly enabled via the **nmosp enable** global configuration command.
- IP device tracking is disabled by default on Catalyst switches and must be enabled in order for Context-Aware wired device tracking to function properly. It can be enabled by issuing the **ip device tracking** in global configuration mode on the context-aware switch.
- The civic location identifier in the LLDP-MED TLV is limited to 250 bytes or less. To avoid receiving error messages regarding available buffer space during switch configuration, the total length of all civic location information specified for each individual civic-location identifier must not exceed 250 bytes.
- In Release 7.0, all wired device client and switch tracking is available only to the root WCS virtual domain user. Because of this, you may wish to limit the use of context-aware wired device tracking in this release to only those users with whom you are comfortable assigning WCS root virtual domain privileges.

## NMSP Attachment Notification Interval

After an NMSP session is established between the MSE and a context-aware Catalyst switch, the MSE transmits an echo response packet to the switch every echo interval time period. The echo interval is specified on the MSE using the WCS menu entitled Services > Mobility Services > System > NMSP Parameters and applies to all NMSP session partners. The echo interval can be set from 1 to 120 seconds, with the default being 15 seconds.

More information about NMSP session parameters can be found in the *Cisco Context-Aware Service Configuration Guide 7.0* at:

[http://www.cisco.com/en/US/docs/wireless/mse/3350/7.0/CAS/configuration/guide/msecg\\_ch4\\_CAS.html#wp1014368](http://www.cisco.com/en/US/docs/wireless/mse/3350/7.0/CAS/configuration/guide/msecg_ch4_CAS.html#wp1014368).

In addition to echo responses, the switch will periodically send attachment notifications to the MSE via the NMSP session. Any link-up or link-down events that are detected by the switch are aggregated during a configurable time interval and sent to the MSE via an attachment notification at the conclusion of that time interval. This interval is known as the **nmosp notification interval attachment interval-seconds** global command. The range of values for interval-seconds is from 1 to 30 seconds, with 30 seconds being the default.

The setting chosen for **nmosp notification interval attachment** will impact how quickly changes in device attachment are propagated from switches to the MSE. Shorter settings result in changes in device attachment being reflected faster, but at the cost of increased switch activity and NMSP network traffic. Longer settings result in more efficient aggregation of attachment information in NMSP packets, but changes in device attachment will not be reflected at the WCS as quickly.

In large networks where there are many NMSP sessions active to the MSE and the number of users connecting and disconnecting from each switch is high, configuring **nmosp notification interval attachment** to a very low number can increase the workload on each switch and increase the amount of NMSP traffic generated between switches and the MSE, creating an unnecessary burden on your switches, your network, and the MSE.

## Civic Address Configuration

The information contained in the *Context Aware Service Configuration Guide 7.0* ([http://www.cisco.com/en/US/docs/wireless/mse/3350/7.0/CAS/configuration/guide/msecg\\_ch7\\_CAS.html#wp1224011](http://www.cisco.com/en/US/docs/wireless/mse/3350/7.0/CAS/configuration/guide/msecg_ch7_CAS.html#wp1224011)) provides the guidance necessary to configure context-aware switches and the WCS for wired device tracking.

As can be seen in the *Context Aware Service Configuration Guide 7.0*, all configuration of civic and ELIN location information is performed on the Catalyst switch. This is normally done either directly using the switch CLI or by uploading a switch configuration text file that has been created offline. Once a switch is configured with the desired civic and ELIN location information, the switch will share all of the configured port information with the MSE when the NMSP session is initially established and will periodically update the MSE if any location updates are performed.

Readers should find the following IETF RFC documents helpful in better understanding the types of values that should be specified for the various civic location fields:

- RFC 4776 (<http://www.ietf.org/rfc/rfc4776.txt>)
- RFC 4589 (<http://www.ietf.org/rfc/rfc4589.txt>)
- RFC 5139 (<http://www.ietf.org/rfc/rfc5139.txt>).

During the course of lab testing, we noted the following behaviors that should prove useful to the network designer or network administrator planning to deploy Context-Aware Services and wired endpoint location:

- Civic and ELIN location configuration scope—In both code trains, civic location and ELIN information may be defined at a global level and then assigned to each switch interface using the appropriate civic or ELIN location identifier. The following example illustrates the process to define both a civic and ELIN location identifier and then assign both to an interface:

```
Switch(config)# location civic-location identifier bldg_17
Switch(config-civic)# number 3550
Switch(config-civic)# primary-road-name "Cisco Way"
Switch(config-civic)# city "San Jose"
Switch(config-civic)# state CA
Switch(config-civic)# building 19
Switch(config-civic)# room C6
Switch(config-civic)# county "Santa Clara"
Switch(config-civic)# country US
Switch(config-civic)# end
Switch(config-civic)# exit
Switch(config)# location elin-location 2122315596 identifier 1
Switch(config)# interface gigabitEthernet 1/0/1
Switch(config-if)# location civic-location-id bldg_17
Switch(config-if)# location elin-location-id 1
Switch(config-if)# exit
Switch(config)#exit
Switch#
```

If more than one switch port shares the same civic location or ELIN, then the same globally specified civic and ELIN location identifiers can be assigned. One such case where this would make sense might be where a user requires three separate switch ports to be terminated at the same physical location in a

building. In this case, a common civic-location identifier can be used, since all three ports will be present at the same building address, floor, room, and seat. Since each termination will bear a unique wiring jack identifier, this information can be uniquely applied to each interface using the interface location parameter **additional-location-information**. Any information specified using **additional-location-information** will appear in the “Address Line 2” area under the Civic Address panel for the device. [Example 4-1](#) is an excerpt of a switch configuration illustrating how a common civic-location identifier and **additional-location-identifier** can be used:

**Example 4-1 Common civic-location Identifier Shared among Multiple Interfaces with additional-location-information**

```
location civic-location identifier bldg_300
  building "Building 300"
  city Alpharetta
  country US
  county Fulton
  name "ABC Corporation"
  postal-code 30022
  state Georgia
  street-group "Kimball Bridge Road"
  number 9300
  floor 2
  room 212
  type-of-place "Shipping & Receiving"
!
location elin-location 19789363021 identifier 1009
location elin-location 19789363022 identifier 1010
location elin-location 19789363023 identifier 1011
.
.
interface GigabitEthernet1/0/9
  location civic-location-id bldg_300
  location elin-location-id 1009
  location additional-location-information "A10-Seat 5 Jack 1"
  switchport access vlan 106
  switchport mode access
.
.
interface GigabitEthernet1/0/10
  location civic-location-id bldg_300
  location elin-location-id 1010
  location additional-location-information "A10-Seat 5 Jack 2"
  switchport access vlan 106
  switchport mode access
.
.
interface GigabitEthernet1/0/11
  location civic-location-id bldg_300
  location elin-location-id 1011
  location additional-location-information "A10-Seat 5 Jack 3"
  switchport access vlan 106
  switchport mode access
```

In some cases, the switch might contain ports for which the assigned civic location information is largely unique. One approach to handling this would be to define unique civic-location-identifier definitions globally and apply those definitions to ports on a 1:1 basis. This technique was introduced in version 6.0, remains largely unchanged, and functions consistently across all supported Catalyst switching products. It is recommended when civic location identifiers differ greatly, such as when some ports are located in a building with a different name, address, organization, type, and so on. This is illustrated in

[Example 4-2](#), where a switch port terminates in an neighboring outbuilding (Building 301), as well as a building that is just across the adjacent state highway but technically in a different city and county (Building 400).

**Example 4-2 Unique civic-location-identifier Assigned to Each Port Interface**

```
location civic-location identifier bldg_300
  building "Building 300"
  city Alpharetta
  country US
  county Fulton
  name "ABC Corporation"
  postal-code 30022
  state Georgia
  street-group "Kimball Bridge Road"
  number 9300
  floor 2
  room 212
  type-of-place "Shipping & Receiving"
!
location civic-location identifier bldg_301
  building "Building 301"
  city Alpharetta
  country US
  county Fulton
  name "ABC Corporation"
  postal-code 30023
  state Georgia
  street-group "Waters Road"
  number 310
  floor 1
  room 000
  type-of-place "Receiving Overflow"
  additional-code "B13-Seat 22 Jack 9"
!
location civic-location identifier bldg_400
  building "Building 400"
  city Milton
  country US
  county Milton
  name "ABC Express"
  postal-code 30004
  state Georgia
  street-group "Freeman Road"
  number 55
  floor 4
  room 411
  type-of-place "Cust Pickup"
  additional-location-information "K64-Seat 7 Jack 13"
!
!
location elin-location 19789363024 identifier 1012
location elin-location 19789363025 identifier 1013
location elin-location 19789363026 identifier 1014
.
.
interface GigabitEthernet1/0/12
  location civic-location-id bldg_300
  location elin-location-id 1012
  location additional-location-information "A12-Seat15 Jack 1"
  switchport access vlan 106
  switchport mode access
```

```

switchport mode access
.
.
interface GigabitEthernet1/0/13
location civic-location-id bldg_301
location elin-location-id 1013
switchport access vlan 106
switchport mode access
.
.
interface GigabitEthernet1/0/14
location civic-location-id bldg_400
location elin-location-id 1014
switchport access vlan 107
switchport mode access
.
.

```

Note the use of **additional-location-information** embedded within the **civic-location-identifier** for Building 400. It is specified here instead of as a standalone statement at the GigabitEthernet1/0/14 interface. Also, note the use of **additional-code** within the **civic-location-identifier** for Building 301. We use this field in this example instead of **additional-location-information** simply to demonstrate its use. Whereas information specified for **additional-location-information** will appear in the device's "Address 2" field on the WCS Civic-Address panel, information specified for **additional-code** will appear in the "additional code" field under the WCS Advanced information panel.

**Note**

**Additional-code** cannot be specified as a standalone interface location command in the same manner as **additional-location-information**.

But what about the case where all ports in a high-density Catalyst switch terminate at a location where the civic location information remains largely the same, except for a few parameters such as floor, room, or seat<sup>4</sup>? Or the case where an entire switch is used to service a single floor of a building, where the only civic location parameters that vary are room and seat? Indeed, both of these are very common occurrences in modern enterprise office complexes. What is the most efficient manner to address this without creating inordinately long switch configurations or requiring a massive amount of redundant information to be entered?

## Civic Address Port-Location

Until the availability of IOS code trains supporting the **port-location** interface subcommand set, there were just two approaches to this situation. The choice of which approach to use depended on whether we required information to appear in their labeled areas on the Civic Address and Advanced information screens for each device. If we wanted to see the floor number display under the "Floor" heading, the room number display under the "Room" heading, and so on, then our only option was to specify a unique, fully-qualified civic-location-identifier for each and every port (as used in [Example 4-2](#)). The downside to this approach is the amount of labor required to create large switch configurations for high density, context-aware Catalyst switches, with a unique civic-location-identifier global definition for each interface.

4. It is assumed that "seat" can be used to specify a seat in a large open area, a cubicle, an office, or some type of combined location code such as pole number/cubicle (e.g., 3C/002). For equipment such as servers, seat could be used to indicate rack and slot number, although **additional-location-information** is probably better suited for that purpose.

If, however, we could suffice with the port-specific information being placed into a format that would fit within the **additional-location-information** field, then we could use as little as one civic-location identifier common to all ports in the switch and specify the floor/room/seat/jack in a concatenated fashion (e.g., 02-212-A45-01) using the **additional-location-information** field at the interface level. This approach was shown in [Example 4-1](#).

While both of these solutions are capable of fulfilling the task at hand, clearly neither is ideal. In our validation, however, we observed that the code trains tested contained a new **port-location** subcommand set within the **location civic-location** interface command:

```
location {additional-location-information word | civic-location-id id [port-location] |  
elin-location-id id}
```

**Port-location** is optional and is used to specify one or more port-specific location attributes. After entering the **location civic-location-id id port-location** command, the user is placed into the civic location port subcommand configuration mode. In this mode, the user can enter additional location attributes for every port using the same location parameters that were available under the global civic-location identifier.

The CLI command help feature (?) lists details regarding the location attributes that can be configured in this mode.

```
cr22-3750s-LB#conf t
cr22-3750s-LB(config)#interface gigabitEthernet 1/0/9
cr22-3750s-LB(config-if)#location civic-location-id bldg_300 port-location
cr22-3750s-LB(config-if-port)#?
Civic location configuration mode:
  additional-codeSet additional code, CA Type 32
  additional-location-information Set additional location info, CA Type 22
  branch-road-name Set branch road name, CA Type 36
  building Set building information, CA Type 25
  city Set city name, CA Type 3
  country Set the country id
  county Set county name, CA Type 2
  division Set city division name, CA Type 4
  floor Set the floor number, CA Type 27
  landmark Set landmark information, CA Type 21
  leading-street-dir Set leading street direction, CA Type 16
  name Set resident name, CA Type 23
  neighborhood Set neighborhood information, CA Type 5
  number Set the street number, CA Type 19
  post-office-box Set post office box, CA Type 31
  postal-code Set postal code, CA Type 24
  postal-community-name Set the postal community name, CA Type 30
  primary-road-name Set primary road name, CA Type 34
  road-section Set road section, CA Type 35
  room Set room information, CA Type 28
  seat Set seat information, CA Type 33
  state Set state name, CA Type 1
  street-group Set street group, CA Type 6
  street-name-postmodifier Set street name postmodifier, CA Type 39
  street-name-premodifier Set street name premodifier, CA Type 38
  street-number-suffix Set street number suffix, CA Type 20
  street-suffix Set the street suffix, CA Type 18
  sub-branch-road-name Set sub branch road name, CA Type 37
  trailing-street-suffix Set trailing street suffix, CA Type 17
  type-of-place Set type of place, CA Type 29
  unit Set unit, CA Type 26
```

If a civic-location attribute is configured globally as well as on the interface using **port-location**, the **port-location** configuration takes precedence. In this way, **port-location** could be thought of as a method with which to “override” the attributes specified by the global civic-location identifier.

**Example 4-3** illustrates the usefulness of **port-location** when configuring a high-density switch stack where all ports in the stack are used to service the same building. In this case, only the floor, room, and seat civic location parameters vary from port to port, with all other civic location attributes remaining constant. **Additional-location-information** is used at the port level to specify information relating to Ethernet jack numbering.

**Example 4-3 Use of Common civic-location Identifier with port-location**

```
location civic-location identifier bldg_405
  building "Building 405"
  city Milton
  country US
  county Milton
  name "ABC Corporation"
  postal-code 30004
  state Georgia
  street-group "Bethany Rd"
  number 693
  type-of-place "Test Facility"
!
!
location elin-location 19789363027 identifier 1015
location elin-location 19789363028 identifier 1016
location elin-location 19789363029 identifier 1017
location elin-location 19789363030 identifier 1018
location elin-location 19789363031 identifier 1019
.
.
interface GigabitEthernet1/0/15
  location civic-location-id bldg_405 port-location
    floor 1
    room 117
    seat 1A/027
  location elin-location-id 1015
  location additional-location-information "1A/117/J1"
  switchport access vlan 106
  switchport mode access
.
.
interface GigabitEthernet1/0/16
  location civic-location-id bldg_405 port-location
    floor 2
    room 222
    seat 1G/028
  location elin-location-id 1016
  location additional-location-information "1G/222/J2"
  switchport access vlan 106
  switchport mode access
.
.
interface GigabitEthernet1/0/17
  location civic-location-id bldg_405 port-location
    floor 3
    room 314
    seat 3F/007
  location elin-location-id 1017
  location additional-location-information "3F/314/J1"
  switchport access vlan 106
  switchport mode access
.
.
interface GigabitEthernet1/0/18
```

```

location civic-location-id bldg_405 port-location
  floor 4
  room 431
  seat 2Y/107
location elin-location-id 1018
location additional-location-information "2Y/107/J3"
switchport access vlan 106
switchport mode access
.
.
interface GigabitEthernet1/0/19
location civic-location-id bldg_405 port-location
  floor 5
  room 516
  seat 5E/044
location elin-location-id 1019
location additional-location-information "5E/516/J1"
switchport access vlan 106
switchport mode access
.
.

```

The utility of the **port-location** subcommand mode should be obvious from the preceding example. Using **port-location**, minor changes to global civic-location information can be easily made at the interface level. Common civic location parameters can now be used across the switch where necessary and they can be expanded on or overridden in a very straightforward fashion.

## Civic Location Caveats

During validation, we made note of a few minor caveats concerning the use of **port-location** with the access switches in the test bed:

- Civic-location information added at the interface level using the **port-location** subcommand set is not propagated from the switch to the Mobility Services Engine unless:
  1. A change is made to any of the global civic-location identifiers.  
Or:
  2. The Catalyst switch is unassigned and re-assigned to the MSE in WCS and then re-synchronized.  
Or:
  3. The Catalyst switch is rebooted.

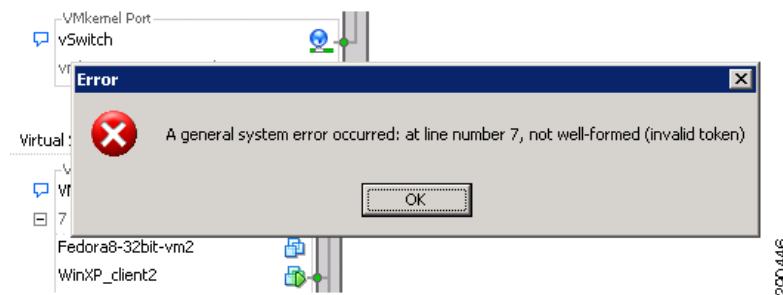
Since we consider the **port-location** capability to be a significant usability enhancement to the configuration of civic location information in Catalyst switches, we recommend #2 above as a temporary workaround. This resynchronization process needs to occur only once after all **port-location** changes have been made to the switch configuration. This workaround has the advantage of being non-disruptive to any data plane traffic to or from the switch.

- Although it is possible to specify “additional-location-information” as a **port-location** subcommand, we noticed during validation that doing so does not result in propagation of the information being added to the MSE. We identified a workaround to this caveat as simply being the use of the standalone **additional-location-information** interface level command instead.

## Other Considerations and Caveats

- If you are using Location MAC Filtering (Services > Mobility Services > Context Aware Service > Administration > Filtering Parameters) to specifically limit or block tracked wireless clients and tags, be advised that these address filters apply to wired device clients as well. Make sure that any filtering specifications that you set using Location MAC Filtering are flexible enough to allow tracking of not only your wireless clients and tags, but wired devices as well. Any devices that have been blocked from location tracking as a result of a defined filter will be viewable under the “Blocked MACs” listing on the Filtering Parameters page. Detailed information regarding how to configure Location MAC filtering can be found in the Modifying Filtering Parameters section of the *Cisco Context-Aware Service Configuration Guide 7.0* ([http://www.cisco.com/en/US/docs/wireless/mse/3350/7.0/CAS/configuration/guide/msecg\\_ch7\\_CAS.html#wp1100062](http://www.cisco.com/en/US/docs/wireless/mse/3350/7.0/CAS/configuration/guide/msecg_ch7_CAS.html#wp1100062)).
- In order to ensure that serial number information collected by Catalyst switches from attached Cisco IP phones is forwarded to the MSE, CDP should be disabled on the switch interface as a temporary workaround. This behavior was observed in testing with both 9971 and 7975 IP phones.
- ESX server 3.5 and above is installed with Cisco CDP enabled by default on the virtual switches. This can cause unreliable detection of virtual machines and errors at the ESX server (shown in Figure 4-11)

**Figure 4-11 VMWare ESX Error with CDP Enabled**



- In order to avoid such errors and allow reliable detection of virtual machine attachment and disconnection, CDP should be disabled. There are two approaches to accomplishing this:
  - Disabling CDP at the switch interface using the **no cdp enable** IOS interface configuration command.
  - Disabling CDP on the ESX server via the following procedure, replacing *yourVSwitch* by the name of the virtual switch installed on ESX server. Connect to the service console of the ESX server and enter these commands for each virtual switch:
    - **esxcfg-vswitch -B listen yourVSwitch**
    - **esxcfg-vswitch -B down yourVSwitch**
- If a change is made to the value specified for an ELIN, one of the following actions should be performed in order to propagate this change to the MSE:
  1. A change is made to any of the global civic-location identifiers.  
Or:
  2. (Recommended) The Catalyst switch is unassigned and re-assigned to the MSE in WCS, and re-synchronized.

Or:

3. The Catalyst switch is rebooted.

Since it has zero impact on other users that may be passing data through the switch, #2 workaround is recommended.

## Hardware and Software Releases

**Table 4-5** *Hardware and Software Releases Tested*

Component	Comments
Wireless Control System, Release 7.0	For licensing and part number information, see <a href="http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/product_data_sheet0900aecd804b4646.html">http://www.cisco.com/en/US/prod/collateral/wireless/ps5755/ps6301/ps6305/product_data_sheet0900aecd804b4646.html</a>
Mobility Services Engine 3300, Release 7.0	For licensing information, see <a href="http://www.cisco.com/en/US/prod/collateral/wireless/ps9733/ps9742/data_sheet_c07-473865.html">http://www.cisco.com/en/US/prod/collateral/wireless/ps9733/ps9742/data_sheet_c07-473865.html</a>
Catalyst 4500	Access switch; must use crypto (K9) image if Context-Aware Services for wired devices is desired
Catalyst 3750E Switch Stack	Access switch; must use crypto (K9) image if Context-Aware Services for wired devices is desired
Catalyst 3750E, X	Access switch; must use crypto (K9) image if Context-Aware Services for wired devices is desired
Catalyst 3560E, X	Access switch; must use crypto (K9) image if Context-Aware Services for wired devices is desired

## Context-Aware Services—General References

The following are recommended references with regard to general best practice deployment recommendations for Cisco Unified Networks and the use of Context-Aware Services release 7.0:

- Context-Aware Solution Deployment Guide  
[http://www.cisco.com/en/US/products/ps9742/products\\_tech\\_note09186a00809d1529.shtml](http://www.cisco.com/en/US/products/ps9742/products_tech_note09186a00809d1529.shtml)
- Context-Aware Services Configuration Guide, Release 7.0  
[http://www.cisco.com/en/US/docs/wireless/mse/3350/7.0/CAS/configuration/guide/CAS\\_70.html](http://www.cisco.com/en/US/docs/wireless/mse/3350/7.0/CAS/configuration/guide/CAS_70.html)
- Wireless Control System Configuration Guide, Release 7.0  
<http://www.cisco.com/en/US/docs/wireless/wcs/7.0/configuration/guide/WCS70cg.html>
- Catalyst 3560 Switch Software Configuration Guide  
[http://www.cisco.com/en/US/docs/switches/lan/catalyst3560/software/release/12.2\\_55\\_se/configuration/guide/3560\\_scg.html](http://www.cisco.com/en/US/docs/switches/lan/catalyst3560/software/release/12.2_55_se/configuration/guide/3560_scg.html)
- Catalyst 3750 Switch Software Configuration Guide  
[http://www.cisco.com/en/US/docs/switches/lan/catalyst3750/software/release/12.2\\_55\\_se/configuration/guide/scg3750.html](http://www.cisco.com/en/US/docs/switches/lan/catalyst3750/software/release/12.2_55_se/configuration/guide/scg3750.html)

- Catalyst 4500 Series Switch Software Configuration Guide  
<http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/54sg/configuration/guide/config.html>





## CHAPTER 5

# User and Device Network Access Reporting

---

Understanding who is accessing the corporate network, what they are using, and where they are connected allows customers to better understand:

- Current port utilization of the network
- Movement of employees and devices on the network
- Suspicious or unauthorized access of the network
- Location of missing or even stolen assets, such as on a college campus
- Location of unknown devices on the network

Adding historical logging of when users and devices are accessing the network allows:

- Historical port utilization data
- Persistent records of when users and devices accessed specific locations in the network
- Searchable data of historical access for troubleshooting and asset tracking

CiscoWorks LAN Management Solution 4.0 provides a single point for user and device monitoring and tracking data to be collected, organized, and displayed.

## CiscoWorks Lan Management Solution 4.0

CiscoWorks LAN Management Solution 4.0 (LMS) is an integrated suite of management functions that simplifies the configuration, administration, monitoring, and troubleshooting of an end-to-end borderless network. The solution aligns management functionality with the way network operators do their jobs.

This document focuses on the user tracking features of LMS to monitor and log who is accessing the network, from where, with what, and when.

## LMS User Tracking

User Tracking tracks wired end hosts, which might include PC workstations, IP phones, medianet endpoints, and other devices as well as Windows-based users on the network.

[Table 5-1](#) illustrates what can be expected from the user tracking features of LMS:

**Table 5-1** *User Tracking Features*

Who	Username of Windows-based clients
What	Device MAC, IP, and hostname
Where	Switch and port to which device is connected
When	Historical access logs showing when and where devices and users are connecting and disconnecting

## User Tracking—End Device Tracking and End User Tracking

While the overall feature is called User Tracking, it can really be described as a combination of End Device Tracking and End User Tracking, which are the terms used in this document for clarity. Within the LMS product and documentation, the term User Tracking is used to refer to all of the features and capabilities discussed in this document.

- End Device Tracking is of the actual end device hardware connecting to the wired network.
- End User Tracking adds the user name of users logging in and out through Windows-based end devices.

End Device Tracking may be implemented without End User Tracking, but not vice versa. End User Tracking relies on and compliments End Device Tracking information in the LMS database.

[Table 5-2](#) shows the primary information provided by End Device Tracking and End User Tracking.

**Table 5-2** *Information Provided by Tracking Types*

User Tracking	
End Device Tracking	End User Tracking
MAC address/IP address/hostname	MAC address/IP address/hostname
Switch/port/VLAN	Username

The first rows are identical, providing the MAC address, IP address, and hostname. The primary reason for the duplication is to match the user to the device in LMS. The secondary reason is to fill in any missing information.

For example, End Device Tracking relies on reverse DNS to populate the device hostname. The End User Tracking information is matched to an end device in the LMS database via the MAC address. If the device is not listed in DNS, such as a DHCP-enabled client workstation, End User Tracking will populate the hostname in LMS. If the IP address has not been acquired from ARP cache by End Device Tracking, End User Tracking will populate that as well.

# User Tracking Functionality Detail

## End Device Tracking

End Device Tracking represents the majority of the User Tracking functionality in LMS. Three main processes are used to gather and maintain end device information:

- Major acquisitions—Scanning all ports on all devices
- Minor acquisitions—Scanning for changes
- Dynamic updates—SNMP traps sent to LMS from access devices as end devices connect and disconnect

## End Device Scheduled Updates

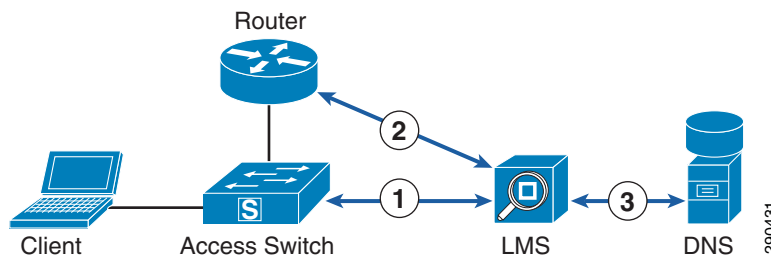
Scheduled updates include major acquisitions and minor acquisitions. For both of these operations, LMS accesses all devices it manages by default. Devices may be explicitly defined in LMS or discovered through a robust discovery mechanism. For more information about device management, see [Product Documentation](#).

Major acquisitions are executed based on a user-defined schedule, which could be anywhere from once per week to every few hours. Major acquisitions look at every port on every access device known to LMS. If all access devices are managed by LMS and set to send dynamic updates to LMS, the interval for major acquisitions may be set to once per day or longer.

Minor acquisitions are executed based on a smaller time interval, defaulting to once per hour. Minor acquisitions only look at changes on the access devices.

[Figure 5-1](#) illustrates the flow of events during scheduled updates.

**Figure 5-1** Event Flow During Scheduled Updates



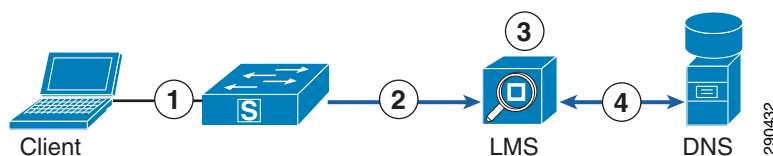
1. LMS executes a query of access devices to pull back MAC addresses associated with ports.
2. LMS executes a query of Layer 3 devices to pull back ARP tables.
3. LMS uses this information to correlate MAC addresses to IP addresses and performs reverse DNS lookups to obtain hostnames.
4. LMS populates the database with obtained data for User Tracking reporting.

## End Device Dynamic Updates

Dynamic updates are the most important, providing current and accurate information as to what devices are connected and where. These updates are initiated from the access devices via SNMP traps sent to a special listener on LMS. When an end device connects or disconnects from an access device, that access device immediately sends a trap to LMS. LMS then immediately updates the device tracking information in the database, recording the event for both historical tracking as well as reports reflecting current devices connected. The update is visible in current device tracking reports within 30 seconds of the event.

Figure 5-2 illustrates the flow of events during dynamic updates.

**Figure 5-2** Event Flow During Dynamic Updates



1. Access device generates a SNMP trap when a client device is connected to a port, sending the device MAC address with port information to LMS.
2. LMS receives this trap and enters it into the database.
3. LMS matches the MAC address to an IP address if an entry exists in the previously-obtained ARP tables.
4. If an IP address exists, LMS performs a reverse DNS lookup to obtain the hostname.
5. LMS populates the database with the obtained data for User Tracking reporting.

## End User Tracking

End User Tracking uses a Windows process on the client to track actual users as they log in and out from Windows-based end devices, matching their information to the end device they are using. This is accomplished by a process initiated on the end host when the user logs in and a login script executes.

The process, called UTLite, sends an initial update to LMS upon user login, followed by continuous updates every 10 minutes to maintain the active login state for that user in LMS. Upon logout, the UTLite process sends an update indicating the logout process occurred, removing the active state in LMS, and moving that login event to the archives for historical tracking purposes.

If the host is disconnected from the network without the user logging out, LMS will stop receiving updates from the UTLite process and will determine that the user is no longer active on the network.

End User Tracking is always dynamic since the end host sends the end user information to LMS unsolicited.

## End User Tracking and End Device Dynamic Updates

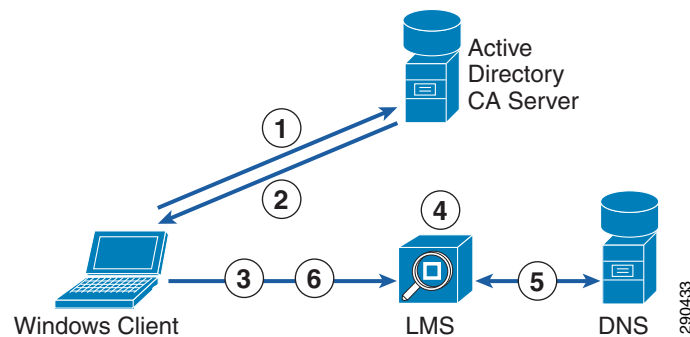
End Device Dynamic Updates, discussed earlier, must be implemented for End User Tracking to function correctly. The information delivered to LMS by End User Tracking is matched to an existing end device in the database. If the end device does not exist, the user tracking data is not recorded.

Without End Device Dynamic Updates implemented, devices connecting to the network will not be recorded until a Scheduled Minor Update runs. If the end device is disconnected before the scheduled update runs, the record of that user connecting to the network will be lost. Implementing End Device Dynamic Updates prevents this by inserting the device into the LMS database before the UTLite process sends the End User Tracking information to LMS.

## End User Tracking UTLite Process

Figure 5-3 illustrates the flow of events during logging of individual user access with the UTLite process.

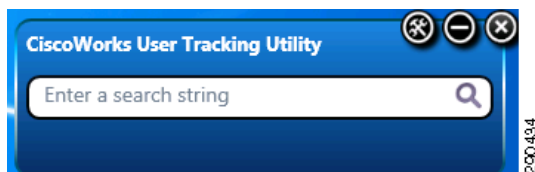
**Figure 5-3** Flow of Events During Logging of Individual User Access



1. Client logs into a domain.
2. Domain script runs, executing the UTLite process on the client.
3. The UTLite process running on the client sends UDP packets to the LMS server with the following client information:
  - User ID
  - Domain
  - MAC address
  - IP address
4. LMS then uses this information to correlate the user information to a device in its database by matching the MAC address and then consolidates the information, filling in User ID, Domain, and IP address.
5. LMS will then do a reverse DNS lookup to fill in the hostname, if it exists in DNS.
6. Upon logout, the UTLite process running on the client notifies LMS, terminates LMS functionality to record when the client logged out of the domain, and removes the username from the active device entry in the database.

## User Tracking Utility (End Device and End User)

The User Tracking Utility is a Windows-based utility used to search the User Tracking data in LMS. The User Tracking Utility lets an administrator using a Windows client workstation enter search criteria to locate end devices. The main interface of the User Tracking Utility is a simple search box as shown in Figure 5-4.

**Figure 5-4** User Tracking Utility Main Interface

After searching for a client, the results are displayed as illustrated in [Figure 5-5](#).

**Figure 5-5** Results Screen

The User Tracking Utility is a simple search utility, not a real-time tracking utility. Device and user information is shown only as a result of a search using some criteria already known to the administrator, such as client hostname.

## User Tracking Reporting (End Device and End User)

User Tracking Reporting allows you to see all hosts, some hosts, or a single host based on how you generate the report. Reports may be generated showing active end devices and users currently connected to the network as well as past connections, which is discussed in the next section.

Reports for active users and devices may be customized to include or exclude specific information about the clients as well as which clients to display. The core information shown by user tracking reports includes:

- End user name
- End device MAC address
- End device hostname
- End device IP address
- If the end device is currently active on the network

- End device subnet
- Access device hostname/IP address
- Access device port
- Access device port VLAN
- Timestamp of when the end device connected to the network

Figure 5-6 shows a single user and end device in a generic user tracking report.

**Figure 5-6 User and End Device in a Generic Tracking Report**

User Name	MAC Address	Host Name	IP Address	Status	Subnet	Device Name	Port	VLAN	Last Seen
User1@mycompany.com	00-15-5d-98-4e-01	z-win7client-1	172.26.152.84	Active	172.26.152.0/24	z-3560r1-4	Gi0/14	VLAN0999	26 Feb 2011, 15:31:08 EST

While Figure 5-6 shows only one user, reports may contain any number of users and devices.

## User Tracking History

The User Tracking End Host History Report may be used to reference the history of end devices connecting and disconnecting from the network and users logging into and out of the network.

Figure 5-7 shows an example of historical tracking of an individual user, showing when and where they logged in and out.

**Figure 5-7 Historical Tracking of an Individual User**

UserName	MAC Address	IP Address	Device	Port	VLAN	Port Connect	Port Disconnect
1. User1@mycompany.com	00-15-5d-98-4e-01	172.26.152.84	172.26.152.65	Gi0/14	VLAN0999	18 Feb 2011, 15:27:02 EST	19 Feb 2011, 14:56:55 EST
2. User1@mycompany.com	00-15-5d-98-4e-01	172.26.152.84	172.26.152.65	Gi0/14	VLAN0999	19 Feb 2011, 14:56:55 EST	20 Feb 2011, 14:56:48 EST
3. User1@mycompany.com	00-15-5d-98-4e-01	172.26.152.84	172.26.152.65	Gi0/14	VLAN0999	20 Feb 2011, 14:56:48 EST	21 Feb 2011, 13:34:36 EST

## Deployment Considerations

The following section is not intended to be a comprehensive installation guide, but rather a collection of significant implementation notes for User Tracking (End Device Tracking and End User Tracking) to supplement the core product documentation.

## Product Documentation

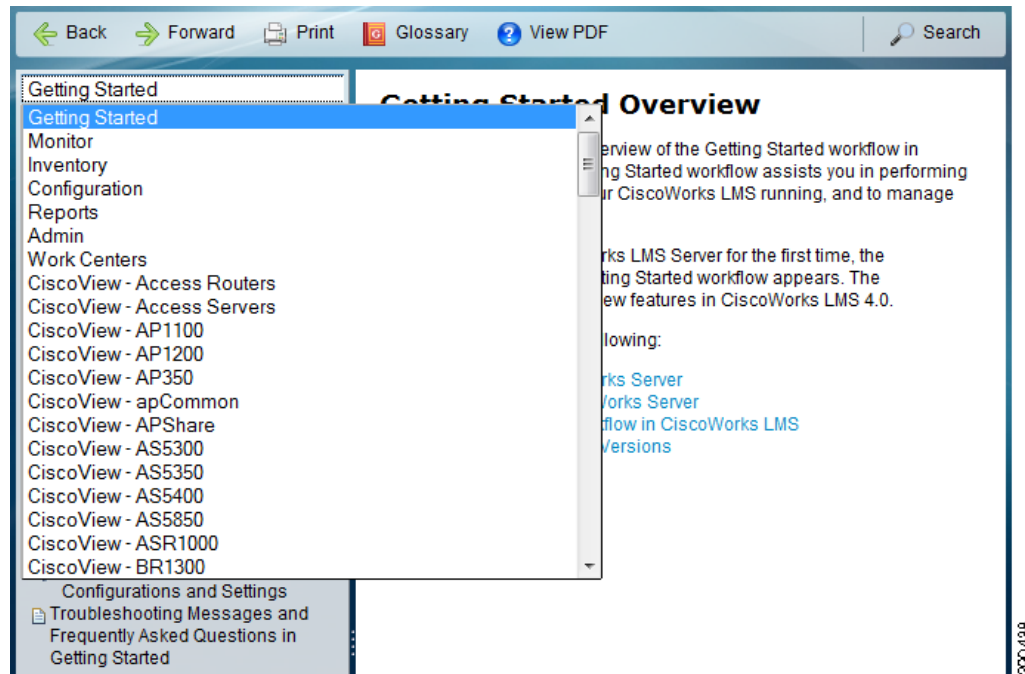
The most current implementation and configuration documentation can be found at: <http://cisco.com/go/lms>.

You should focus on the User Tracking sections of the *LMS 4.0 Administration Guide*.

In addition to documentation and information at <http://cisco.com/go/lms>, downloadable PDFs for all manuals and guides are available in the help section of LMS. Documentation is accessible after initial installation by clicking on the **Help** link in the top right corner of the Web interface.

The LMS help interface is shown in [Figure 5-8](#). Click the white box **Getting Started** and a drop down box shows all the different documents. Click on a document, such as Admin, and then click **View PDF** at the top to download the entire *LMS 4.0 Administration Guide*.

**Figure 5-8** LMS Help Interface



## Firewall Port Configuration

Ensure that all incoming ports have exceptions in the operating system's firewall. [Table 5-3](#) illustrates the critical ports that need to be open for LMS.

**Table 5-3** LMS Critical Ports

Incoming Traffic	Default Port
Default HTTP port	1741
Default HTTPS port	443
Dynamic update SNMP traps	1431
All other SNMP traps	162
UTLite updates from clients	16236

## Dynamic Updates for End Device Tracking

Dynamic updates require the access switch to send SNMP traps to LMS. Both global and interface level configuration must be completed on all access switches and their respective ports to enable dynamic updates.

At the interface level, SNMP traps must be enabled for any device connecting or disconnecting from the port.

### Interface Configuration

```
snmp trap mac-notification change added
snmp trap mac-notification change removed
```

### Global Configuration

```
snmp-server enable traps mac-notification change move threshold
snmp-server host [LMS server IP] [snmp string] udp-port 1431 mac-notification
mac address-table notification change
mac address-table notification mac-move
```

The default UDP port for SNMP traps related to User Tracking is 1431. This may seem to be a deviation from the standard port for SNMP traps, port 162, but there is a reason for using this port. LMS has two separate listeners, one for User Tracking-specific SNMP traps on port 1431 and one for all other SNMP traps on port 162. On the access devices, multiple “snmp-server” statements would be needed if LMS is to receive both User Tracking traps as well as other SNMP traps. LMS only needs to receive “mac-notification” traps for User Tracking.

## UTLite Process for End User Tracking

End user tracking relies on the UTLite33.exe process to be run as a service on Windows clients. When a user logs into a domain, the executable is copied to the client and started from a domain login script. The script is:

```
REM UTLite33.exe options are:
REM -domain <name>          NT/NDS domain name
REM -host <addr>             (Host IP address of ANI Server)
REM -port <num>              (listener port, default is 16236)
REM -sleep <num>             (default = 600, 10 minutes sync interval)
REM -nds                     Will try to send NDS names
REM -sleep and -nds are optional parameters
REM
REM Copy UTLite33.exe file from domain controller to local client.
REM
REM if NOT EXIST %WINDIR%\UTLite33.exe
copy %0\..\UTLite33.exe %WINDIR%
REM
REM Specify the parameters below
REM
start %WINDIR%\UTLite33 -domain <domain> -host <ipaddress> -port 16236
REM
REM WINDIR is where NT system is installed, usually it is: C:\WINNT
REM
REM %0\.. is where UTLite33 is installed on the domain controller, it is
REM the same directory where the login scripts are usually kept. it is:
```

```

REM \WINNT\system32\Repl\Import\Scripts          (for NT) and
REM \\Novell-Server-name\SYS\public\Your-Folder  (for NDS)
REM
REM The client machine copies UTLite33.exe to its local WINDIR directory.

```

The default script copies the executable into the Windows directory. Permission issues may be encountered due to additional protections on this directory. Altering the script to copy and start the executable from an alternative location, such as a temp directory with %temp%, will not impact functionality.

In addition to setting the destination location, the domain and host must be set as well:

- **Domain**—The domain can be set to any text value and not necessarily the domain the user is using. This field may be used to differentiate users in the User Tracking Reporting to meet your needs.
- **Host**—The host is simply the LMS server IP address.

## User Tracking Reporting

As noted earlier, there are many variations in generating User Tracking Reports, from individual users or devices to all users and devices on the network.

Custom reports may be created using the “Report Designer” in the reporting section of LMS. After creation, the location of the created reports may not be obvious. Custom reports reside in the following location.

Reports > Inventory > User Tracking

Click **User Tracking** and custom reports will be shown under all the standard User Tracking reports in the left menu. If the mouse is hovered over User Tracking instead of clicking it from the main navigation menu, all the standard User Tracking reports are shown, but not the custom reports.

## Locations of User Tracking Features in LMS Interface

**Table 5-4**      *Locations of User Tracking Features*

User Tracking Feature	Location
Reporting	Reports > Inventory > User Tracking
Custom Report Creation	Reports > Report Designer > User Tracking
Acquisition Schedules	Admin > Collection Settings > User Tracking > Acquisition Schedule
Acquisition Settings and UTLite port	Admin > Collection Settings > User Tracking > Acquisition Settings
Dynamic Updates Listener enable and port	Admin > Collection Settings > User Tracking > Trap Listener Configuration
Manual Major Acquisition Execution	Admin > Collection Settings > User Tracking > Acquisition Action
Acquisition Info Summary	Admin > Collection Settings > User Tracking > Acquisition Info

## User Tracking Utility Location

The User Tracking Utility must be downloaded from cisco.com. The location of the file is not obvious, however the following process should find it:

Go to <http://www.cisco.com> and then Support > Download Software.

Enter the following exactly as shown into the search box:

CiscoWorks Campus Manager 5.2

The manual path is:

Products > Network Management and Automation > Routing and Switching Management >  
CiscoWorks LAN Management Solution Products > CiscoWorks Campus Manager > CiscoWorks  
Campus Manager 5.2 > CiscoWorks Campus Manager User Tracking Utility

## LMS User Tracking Summary

LMS User Tracking allows current and historical reporting of who is accessing the network, with what, from where, and when. The current and historical data available from LMS provides a single reference point for user and device tracking on the wired network.

