

Implementing Trusted Endpoint Quality of Service Marking

Introduction

This document describes how to integrate Cisco Security Agent (CSA) 5.0, Cisco Network Admission Control Phase 2 (NAC2), and Cisco Quality of Service (QoS) features into a Trusted Endpoint QoS marking solution. It discusses high level design goals and provides example configuration procedures for Cisco IOS and host software devices within an enterprise environment. In particular, it provides guidance on how the campus network trust boundary can be extended to hosts.

CSA 5.0 introduces the capability to apply QoS markings to host application flows as specified by CSA policy rules. These markings are Differentiated Services Code Point (DSCP) values inserted into the IP headers of transmitted packets. They can be used by Cisco IOS devices upstream in the enterprise network to classify the packets and apply QoS service policies such as policing and queueing.

In a best practice deployment, NAC2 is implemented to identify which hosts are running CSA 5.0 and to configure the campus access ports to ensure that only trusted endpoints have full network access. Non-conforming hosts can be quarantined until remediation is performed and they are brought into compliance.

The solution that results from the integration of these three technologies can improve delivery of mission-critical traffic as well as enhance the self defending nature of the enterprise network by providing mitigation against DoS and worm attacks.

Intended Audience

This document is intended for system engineers and administrators responsible for the implementation of QoS and/or network security policy. This document assumes you know how to configure Cisco IOS devices and are already familiar with QoS, CSA, and NAC2 deployment issues.

How This Document is Organized

This document consists of a design guide with example configurations, followed by three appendices that provide overviews of QoS, CSA, and NAC2 features, respectively. These appendices are included as a basic background for understanding the implementation guidance and example configurations presented in the design guide chapter. If you need additional background, consult the references cited below.

Obtaining Additional Documentation

If you want to review the most recent Cisco enterprise QoS policy design guidance in more detail, refer to *Enterprise QoS Solution Reference Network Design* at:
http://www.cisco.com/application/pdf/en/us/guest/netso/ns432/c649/ccmigration_09186a008049b062.pdf

QoS tools and design concepts are presented in even greater depth in the Cisco Press book *End-to-End Quality of Service Network Design* (ISBN: 1587051761).

Implementation of CSA assumes that you are familiar with Microsoft Windows operating systems and host machines and with the configuration and operation of Management Center for Cisco Security Agents. It also assumes you are familiar with certificate authorities and the trust models provided by digital certificates. The following documents are available on Cisco.com:



Installing Management Center for Cisco Security Agents 5.0
Using Management Center for Cisco Security Agents 5.0

Implementation of NAC2 assumes that you are familiar with Microsoft Windows operating systems and host machines and with the configuration and operation of Cisco Secure ACS. It also assumes you are familiar with certificate authorities and the trust models provided by digital certificates. The following documents are available on Cisco.com:

Network Admission Control Documentation Reference
Cisco Trust Agent Administrator Guide, Version 2.0
User Guide for Cisco Secure ACS for Windows Server, Version 4.0
Configuring Network Admission Control feature documentation for Cisco IOS Release 12.2(25) SEC

Implementation of Trusted Endpoint QoS Marking

This chapter describes the details of the Trusted Endpoint QoS marking solution. It discusses deployment issues and provides example configurations for each of the three technologies integrated in this solution - CSA, NAC2, and QoS.

Solution Overview

The object of the Trusted Endpoint QoS Marking solution is to apply packet classification and marking policy at the source of the traffic, namely in the host. The applied DSCP markings can then be used by upstream network devices to classify packets. This is a more scalable and granular approach than previous methods for trusted marking of host traffic such as installing ACLs on the access switch ports. Marking at the host also works over VPN tunnels, where the original IP header is opaque to network devices.

The problem up until now has been the susceptibility of general purpose hosts to virus and worm attacks, making the QoS markings originating from such hosts untrustworthy. CSA solves this problem by hardening the host against infection and preventing harmful behaviors. CSA is also an ideal platform for providing proxy capabilities such as QoS marking on behalf of legacy applications. CSA QoS policies are typically deployed by the same organization that administers the desktop applications.

CSA protects against infections and performs the role of a QoS marking policy enforcement point. In cases where the correct installation and operation of CSA can be guaranteed, such as on a server in a data center, installing CSA may provide sufficient confidence to extend trust to the endpoint. In best practice deployments, NAC2 will be implemented to ensure that the required version of CSA is installed and running on the hosts.

Where is Trusted Endpoint QoS Marking Applicable?

This solution is applicable to enterprise or campus networks that implement the DiffServ architecture as described in the Cisco references cited at the beginning of this document. The goal of such networks is to separate application traffic into different service classes, providing either less than, equal to, or better than Best Effort according to application performance needs. QoS has been succinctly described as “a system of managed unfairness.”

The network access switches must be capable of supporting a fairly functional and granular policing and queuing policy, such as the Cisco Catalyst 2970 switch or better. Applicable networks will also have Cisco network security products deployed including CSA and ideally NAC2. The remainder of this document assumes you have at least a minimal lab environment in which you can deploy and test QoS and security policies.

How Should Trusted Endpoint QoS Marking be Deployed?

It is important to deploy Trusted Endpoint QoS Marking first at the hosts and then move toward the access edge. That is, first implement the conditions necessary to extend trust to the hosts and then modify access switch port configurations to permit marked traffic. This is summarized by the following tasks:

- Deploy CSA 5.0 on all hosts and at a minimum define a default QoS rule to mark all traffic to Best Effort. Define additional rules to apply markings to identified mission critical applications.
- (Recommended best practice) Deploy NAC2 and define the “Healthy” posture to include at least the conditions that: a) CSA 5.0 or later is installed and b) CSA is running.
- Implement IOS access switch port policy to trust DSCP markings but police and markdown excess traffic. In best practice deployments, trust is conditionally extended based on host posture assessment.

Information on how to accomplish each of these tasks is presented in the following sections.

Example Network Architecture

The remaining sections in this chapter assume a limited lab network like the one shown below in Figure 1:




Figure 1: Example Network Architecture

CSA QoS Policy Configuration Example

This section provides instructions for configuring and distributing a simple Trusted Endpoint QoS policy to Cisco Security Agents. The following example assumes you have at least a limited CSA deployment for characterizing application behavior and testing QoS policies. It describes how to create rules to classify and mark data flows for the Cisco SoftPhone application and to distribute this policy to the agents that are installed on end user systems. For a full discussion of application classes, rule modules, and policies, you should refer to *Using Management Center for Cisco Security Agents 5.0*.

Configuring a Static Application Class

Access control rules are application-centric. The focus of this example is to create a policy consisting of network access control rules that monitor or mark particular traffic flows. The first step in generating a Trusted Endpoint QoS policy is to specify which application(s) the policy will pertain to by creating an application class. This provides fine-grained control over the scope of the policy. An application class for network access control policy might be quite broad, for example all applications that transmit to port 80, but for QoS marking rules the class will often be quite specific.

To create an application class, do the following:

- Step 1** Move the mouse over **Configuration** in the menu bar and select **Applications>Application Classes** (Windows or UNIX) from the pulldown list that appears. The list of existing Application classes is displayed. CSA MC ships with several pre-configured applications. Some Application classes appear within brackets. These are built-in CSA MC application classes and you cannot edit them.
- Step 2** Click the New button to create a new application class. This takes you to the application class configuration view (see Figure 2).
- Step 3** Enter a **Name** for the application class you are creating. It is important to use a descriptive name that you can easily recognize in the application selection list that appears in the rule views. For this example you will enter *SoftPhone*.
- Step 4** Enter a **Description** for your application class. This description becomes visible in the application class list view. You will enter *Windows application class for SoftPhone QoS*.
- Step 5** **Operating System**—When you create an application class, you must select to either create a UNIX or a Windows application class. Your application class is then designated for all UNIX or all Windows platforms. Optionally, you may target an operating system more narrowly by selecting a specific UNIX or Windows operating system from the **Target** pulldown menu. For this example you do not need to change this setting.
- Step 6** You do not need to change the **Display only in Show All mode** checkbox setting, which is unchecked by default.
- Step 7** Under **Add process to application class**, for a static application class, do the following:
Leave the default **when created from one of the following executables** radio button selected. Then enter the executable file names (one per line) for the applications you are grouping together in this application class. For this example, enter ***\SoftPhone.exe*. This syntax indicates a specific application with an arbitrary pathname.
- Step 8** When you are finished, click the **Save** button. This application class name, *SoftPhone*, now appears in the application list view and in the application selection fields for rule configurations. When you select it in a rule, you are indicating all the executables that comprise it (which, in this case, is a single application.)




Figure 2: Static Application Class

Configure a Rule Module

When you configure a policy, you are combining rule modules under a common name. Those rule modules are then attached to a policy. That policy is attached to a group of hosts and it uses the rules that comprise the policy to control the actions that are permitted and denied on those hosts.

For this example, you will configure a rule module containing network access control rules that monitor (log) all network activity for the SoftPhone application and also set DSCP markings for its voice and call-signaling flows.

Note: Cisco recommends that you do not edit the preconfigured policies shipped with the Management Center for Cisco Security Agents, but instead add new policies to groups for any changes you might want.

To configure the SoftPhone QoS rule module, do the following.

- Step 1** Move the mouse over **Configuration** in the menu bar and select **Rule Modules [Windows]** from the drop-down list that appears. The Windows Rule Module list view appears.
- Step 2** Click the **New** button to create a new module. This takes you to the Rule Module configuration page. See Figure 3.
- Step 3** In the configuration view, enter the **Name** *SoftPhone*. Note that names are case insensitive, must start with an alphabetic character, can be up to 64 characters long. Spaces are also allowed in names.
- Step 4** Enter a **Description** of your module. You will enter *Windows rule module for SoftPhone QoS*.

Step 5 Click the **Save** button. (You will not use State Sets in this example.)

Now you will add network access control rules to this module.




Figure 3: Rule Module Creation View

Create a Network Access Control (Monitor) Rule

- Step 1** From the Rule Module configuration page (Figure 3), click the **Modify rules** link at the top left corner of the page. You are now on the Rules page.
- Step 2** In the Rule page, click the **Add rule** link. A drop down list of available rule types appears.
- Step 3** Click the **Network access control** rule from the drop down list (see Figure 4). This takes you to the configuration page for this rule.
- Step 4** In the Network access control rule configuration view (see Figure 5), enter the following information:
 - **Description**— You will enter *SoftPhone network monitor access*.
 - **Enabled**—(This is selected by default. Don't change this setting for this example.)
- Step 5** Select **Monitor** from the **Take the following action** pulldown list.

Normally a QoS policy would be developed in stages. First, the application traffic would be characterized and then a classification and marking policy would be specified. A monitor rule can help in the first stage because it can be used to log all application traffic. In this example we include it to show how it is configured.




Figure 4: Add Rules to Module

- Step 6** Select a preconfigured **Application class** from the available list to indicate the applications whose network access you want exercise control over. For this example, use the *SoftPhone* class created in the previous section. Note that when you click **Save**, selected application classes move to the top of the list.




Figure 5: Network Access Control (Monitor) Rule

- Step 7** The **But not in the following class** field can be used to limit the scope of the rule. It is not used in this example.
- Step 8** Now you will enter information that classifies the traffic of interest (see Figure 6.) In the case of a monitor rule, you wish to log all traffic. In the **Attempt to act as a** pulldown, select **client or server**. This will capture traffic either to or from the service ports specified in the next step.
- Step 9** In the **for network services** field, enter *UDP/0-65535* and *TCP/0-65535*. This specifies you wish to log traffic to/from all service ports. If you were only interested in monitoring, say, connections to TCP port 80 from the host then you would specify **client** in the previous step and *TCP/80* here. This would have the effect of logging all connections to TCP port 80 on a remote server along with the corresponding replies (i.e. replies to the *originating* port on the host.)
- Step 10** In the **Communicating with host addresses** and **Using these local addresses** fields, you have the ability to constrain the destination and source IP addresses. For this example, you will leave these as they are.
- Step 11** Click the Save button.

You have now classified traffic of interest based on application name (by specifying the class) and by IP 5-tuple (protocol, source and destination port, and source and destination network address.) The ability to specify ranges for these fields provides a wild-card capability in the classification. Next, you will create marking actions based on **set** network access control rules.




Figure 6: Network Access Control (Monitor) Rule, continued

Create a Network Access Control (Set Attribute) Rule

- Step 1** Return to the Rule page; you can see the effects of the previous save. Once again, click the **Add rule** link.
- Step 2** Click the **Network access control** rule from the drop down list. This takes you to the configuration page for this rule.
- Step 3** In the Network access control rule configuration view (see Figure 7), enter the following information:
 - **Description**— You will enter *SoftPhone mark call-signaling with CS3*.
 - **Enabled**—(This is selected by default. Don't change this setting for this example.)
- Step 4** Select **Set** from the **Take the following action** pulldown list.
- Step 5** From the **Attribute** pulldown menu, select **Differentiated Service**.
- Step 6** From the **Value** pulldown menu, select **Call Signaling (24, CS3)**. The values in this pulldown list all of the standards-based DSCP values shown in Table A-1. (Recall that a DSCP can be described by either its numeric value or PHB label.) However, there are additional choices. **Application specified** is the default QoS marking action in the absence of any **set attribute** rules to the contrary. This means that CSA will pass any markings the application may generate. Cisco IP Communicator, the replacement for SoftPhone, is an example of an application that marks its own traffic. But recall from the last chapter that the access switch ports will trust these markings, subject only to aggregate edge policing. Therefore, in a production environment it is a

best practice to create a *default* marking rule that sets the Differentiated Service value of all application network traffic to **Best Effort**.

If multiple network access control rules are written for the same policy then a more specific classification will take precedence over a less specific one. In this way, exceptions to the default rule can be generated. If the classification parameters (IP 5-tuple) are equivalent then the precedence depends on attribute value. In this case, the values are ordered in the pulldown menu from highest (top) to lowest. The **priority Best Effort** and **priority Scavenger** values enable you to generate rules that override existing rules.

Step 7 Select the **Log** checkbox.

This means that the system action in question is logged and sent to the server. Because this example is intended to generate test QoS rules, you will want to turn logging on so you can monitor event activity.

Step 8 For **When an enforcement action of the following type [] occurs**, accept the default **Allow** choice.




Figure 7: Network Access Control (Set Attribute) Rule

Step 9 Select the **SoftPhone Application class** from the drop down list as before (see Figure 8.)

Step 10 The **But not in the following class** field can be used to limit the scope of the rule. It is not used in this example.

Step 11 Now you will enter information that classifies the traffic of interest. In the **Attempt to act as a** pulldown, select **client or server**. This will capture traffic either to or from the service ports specified in the next step.

- Step 12** In the **for network services** field, enter *TCP/2000-2002*. This range specifies the protocol and ports used to transmit SCCP call setup packets to the Cisco Call Manager.
- Step 13** In the **Communicating with host addresses** and **Using these local addresses** fields, you have the ability to constrain the destination and source IP addresses. For this example, you will leave these as they are.
- Step 14** Click the Save button.
- Step 15** Repeat these steps and generate a second **Set** rule that will mark SoftPhone's voice data packets. Use the following values in the Network access control rule configuration view:
- **Description**— enter *SoftPhone mark voice data with EF*
 - **Attribute**— enter *Differentiated Service*
 - **Value**— enter *Voice (46, EF)*
 - **Log**— select
 - **Class**— select *SoftPhone*
 - **Act**— select *client or server*
 - **Network Services**— enter *UDP/16384-32767*

Step 16 Click the Save button.

You have now created three network control rules - one monitor rule and two set attribute rules. Next, you will create a policy to attach the rule module to.




Figure 8: Network Access Control (Set Attribute) Rule, continued

Configure a Policy

Generally, when you configure a policy, you are combining multiple rule modules under a common name. That policy name is then attached to a group of hosts and it uses the rules that comprise the policy to control the actions that are allowed and denied on those hosts. You can have several different types of rules in a rule module and consequently within one policy.

The policy level is the common ground by which host groups acquire the rules that make up their security policy. You can attach rule modules of differing architectures to the same policy. This way, you can configure task-specific, self-contained, inclusive policies across all supported architectures (Windows, Solaris, and Linux) for software that is supported on all platforms.

To configure a policy, do the following

- Step 1** Move the mouse over **Configuration** in the menu bar of CSA MC and select **Policies** from the drop-down menu that appears. The policy list view appears.
- Step 2** Click the **New** button to create a new policy entry. This takes you to the policy configuration page.
- Step 3** In the available policy configuration fields, enter the following information:
 - **Name**—This is a unique name for this policy grouping of rule modules.
Names are case insensitive, must start with an alphabetic character, can be up to 64 characters long and can include alphanumeric characters, spaces, and underscores. For this exercise, enter the name *SoftPhone*.
 - **Description**—This is an optional line of text that is displayed in the list view and helps you to identify this particular policy. For this exercise, enter *Windows policy for SoftPhone QoS*.
- Step 4** Click the **Save** button.

Attach a Rule Module to a Policy

To apply the configured SoftPhone QoS rule module to the policy you've created, do the following.

- Step 1** From Policy edit view, click the **Modify rule module associations** link. This takes you to a view containing a swap box list of available modules.
- Step 2** Select the **SoftPhone** module from the list box on the left and click the **Add** button to move it to the right side box.
- Step 3** Select the **Cisco Trust Agent Module** from the list box on the left and click the **Add** button to move it to the right side box. This will be used in the next chapter.

The rule module is now attached to this policy.

Attach a Policy to a Group

To apply the configured QoS policy to a particular group of host systems, you must attach the policy to this group.

- Step 1** Move the mouse over **Systems** in the menu bar and select **Groups** from the drop-down menu that appears.
- Step 2** From the group list view, click the link for the group you want to attach the policy to. This brings you to that group's edit view.
- Step 3** From the edit view, click the **Modify policy associations** link. This takes you to a view containing a swap box list of available policies.
- Step 4** Select the **SoftPhone** policy from the list box on the left and click the **Add** button to move it to the right side box.
- Step 5** Select the **Cisco Trust Agent Module** from the list box on the left and click the **Add** button to move it to the right side box. This will be used in the next chapter.

Step 6 The policy is now attached to this group.

Generate Rule Programs

Now that you've configured a QoS policy and attached it to a group, you'll next distribute the policy to the agents that are part of the group. You do this by first generating the rule programs.

Click **Generate rules** in the bottom frame of CSA MC. All pending database changes ready for distribution appear.

If everything looks okay, you can click the **Generate** button that now appears in the bottom frame. This distributes your policy to the agents.

You can ensure that agents have received this policy by clicking **Hosts** (accessible from **Systems** in the menu bar) and viewing the individual host status views. Click the Refresh button on your browser and look at the host Configuration version data in the host view to make sure it's up-to-date.

Note: Hosts poll the CSA MC to retrieve new policies. You can shorten or lengthen this polling time in the Group configuration page. You can also send a hint message to tell hosts to poll in before their set polling interval. See the User Guide for details.

Now your policies are installed and will mark host traffic for the application you have configured.

NAC Phase 2 Configuration Example

This section provides instructions on how to configure an example NAC2 deployment to implement the Trusted Endpoint QoS marking solution. The recommended deployment option is the integrated NAC-L2-802.1x access method, which combines identity and posture validation in a single exchange and supports dynamic VLAN assignment based on host posture. The VLAN assignment of the host can then be used as the basis for trusting QoS markings applied by CSA running on that host.

Configuring NAC2 to support Trusted Endpoint QoS consists of three main objectives. The first objective is to configure the Access Control Server v4.0 (ACS) with parameters common to all deployment options. The second objective is to configure the parameters specific to the NAC-L2-802.1x deployment option, which includes configuring IOS on the Network Access Devices (access switches) to create VLANs and enable AAA association between the NAD client and the ACS. The VLANs that are created should correspond to the possible posture states of the endpoints and the VLAN names must match those referenced in ACS policies. The third objective is to install the Cisco Trust Agent (CTA) on the PCs and validate NAC2 system operation.

This example is taken from the NAC2 Lab Guide (available to Cisco employees and partners) and has been abbreviated to highlight only those elements that are required for Trusted Endpoint QoS marking. For a complete discussion of NAC2 deployment, see the references cited earlier. If you do not have access to the lab software, you must first install ACS v4.0 according to the instructions provided in *Installation Guide for Cisco Secure ACS for Windows*, which can be found at http://www.cisco.com/en/US/products/sw/secursw/ps2086/products_installation_guide_book09186a0080533d5e.html. This example assumes an ACS server named “w2ks” located at 10.0.200.20 in the test network.

Access Control Server (ACS) Common Configuration

The following section will walk you through the basic configuration of ACS 4.0 for all Network Admission Control (NAC) deployment scenarios. Those elements that are not required for NAC-L2-802.1x have been omitted for clarity.

Note: ACS v4.0 or later software is required for NAC phase 2 deployment. This deployment guide assumes the use of the ACS v4.0 Software for Windows and does not cover steps specific to the ACS Solution Engine.

Vendor Attribute-Value Pairs (AVPs)

NAC introduces the ability to authorize network hosts not only based upon user and/or machine identity but also upon a host’s posture compliance. The posture compliance is determined by the comparing the host’s credentials to a compliance policy which you create from attribute-value pairs (AVPs) defined by Cisco and other vendors who are NAC partners. Since the range of NAC attributes extends across many vendors and applications, ACS does not include any non-Cisco attributes by default. Therefore, you must import a NAC attribute definition file (ADF) from each vendor application that you would like to validate in your NAC compliance policies.

Task 1: Import Trend Micro AVPs (Omit)

This task is omitted, since we are only interested in Cisco attributes in this example.

Network Configuration

Task 2: Network Device Group (Optional)

This task is provided as a reference to ensure you are aware of this option. It is recommended that you leave the Network Device Group unassigned.

If you want to group your Network Access Devices (NADs) into Network Device Groups (NDGs) for location or service-based filtering, you must first enable the use of Network Device Groups. This can be

done by selecting **Interface Configuration** from the main ACS menu and picking **Advanced Options**, and then click the box at the bottom of the page to enable **Network Device Groups**. Otherwise you may leave them unassigned.

Select **Network Configuration** from the main ACS menu and select the **Add Entry** and provide the Network Device Group Name and Key.

Network Device Group Name	Key
Switches	cisco123

Task 3: AAA Clients

From the **Network Configuration** screen select the hyperlink under **Network Device Group**. If you did not previously assign a name you will see “**Not Assigned**”. This will take you to the **AAA Client** screen.

Step 1 Configure the AAA Clients by selecting the **Add Entry** button. You can define all NADs as single AAA client using IP address wildcards.

Step2 Click on submit and apply to save the changes.

(Not Assigned) AAA Clients				
AAA client Hostname	AAA Client IP Address	Key	Network Device Group	Authenticate Using
Any	*.*.*.*	cisco123	(Not Assigned)	RADIUS (Cisco IOS/PIX6.0)

Note: AAA client definitions with wildcards CANNOT overlap with other AAA client definitions, regardless of authentication types.

Task 4: AAA Servers

Note: Your AAA Server is automatically populated during the installation of ACS, using the hostname assigned to the host operating system.

The AAA Server information is populated with the hostname and IP address of the machine ACS is installed on. After ACS installation, you will notice the Server Name **w2ks** and IP address **10.0.200.20** is already configured as shown below.

Step 1 Configure the Key as shown below for the AAA server. Configure this by selecting the AAA Server Name hyperlink [w2ks](#).

(Not Assigned) AAA Servers				
AAA Server Name	AAA Server IP Address	AAA Server Type	Key	Network Device Group
w2ks	10.0.200.20	Cisco Secure ACS	cisco123	(Not Assigned)

Note: You can optionally assign the ACS server to a previously configured Network Device Group.

Interface Configuration

The items configured in the **Interface Configuration** section, such as RADIUS attributes, must be enabled here in order to be available in other portions of the ACS configuration.

Task 5: Configure RADIUS Attributes

Select the Interface configuration button from the main menu and select **RADIUS (IETF)** and make the noted selections and then select **RADIUS Cisco IOS/PIX6.0** and make the proper selections.

- Step 1** Select the following required RADIUS attributes. Only the attributes checked below are necessary for NAC. All other attributes should be unchecked in order to save time in later configuration steps.

RADIUS (IETF)	☒ [027] Session-Timeout ☒ [029] Termination-Action ☒ [064] Tunnel-Type ☒ [065] Tunnel-Medium-Type ☒ [081] Tunnel-Private-Group-ID
RADIUS (Cisco IOS/PIX6.0)	☒ [026/009/001] cisco-av-pair

Note: Attributes 64, 65, and 81 are necessary for VLAN assignments.

- Step 2** Enable the following under the **Interface Configuration menu > Advanced Options**

Advanced Options:	☒ Group-Level Shared Network Access Restrictions ☒ Group-Level Network Access Restrictions ☒ Group-Level Downloadable ACLs ☒ Network Access Filtering ☒ Distributed System Settings ☒ Cisco Secure ACS Database Replication ☒ Network Device Groups
--------------------------	---

Note: The Group Downloadable ACLs box must be checked here or downloadable ACLs will not work in later tasks. (N/A for NAC-L2-802.1x)

System Configuration

To access the ACS Certificate Setup menu select **System Configuration** from the main menu and select the **ACS Certificate Setup** link.

Task 6: ACS Certificate Setup

ACS should be configured with a digital certificate for establishing client trust when challenging the client for its credentials.

Note: Using a production PKI and certificates signed by the production CA or RA(s) is highly recommended for the most scalable NAC deployments. We have significantly compressed and abbreviated this part of a NAC implementation – customers will need to use an existing PKI (internal or outsourced) to securely identify the ACS infrastructure to endpoint devices (e.g. CTA). If you are a customer with an Internet-facing presence, you will already have (at minimum) an understanding as to how to obtain SSL certificates (e.g. for a web server).

Note: If your deployment is going into an Active Directory domain you must consider which authentication mechanism, if any, that your deployment requires - domain certificates or domain credentials?

The remainder of this section assumes you have access to the NAC lab materials and the supplied pre-generated digital certificates. These certificates are located in the ACS virtual machine in `c:\files\certs\`

- Step 1** Select the ACS Certificate Authority Setup link. Specify the location of the CA certificate, and click on the **Submit** button.

ACS Certificate Authority Setup
Add new CA certificate to local certificate storage Certificate file: C:\files\certs\ca.nac.cisco.com.cer

- Step 2** Restart ACS after adding the new CA certificate. Go to System Configuration, Service Control and select Restart.
- Step 3** After installing the CA certificate, you should add it to the Certificate Trust List (CTL) as a trusted authority. To do this, select the **Edit Certificate Trust List**, link from the ACS Certificate Setup screen and locate the name of your CA in the list and check the box next to it and click **Submit** to save the changes.

Edit the Certificate Trust List (CTL)
☒ ca

- Step 1** Changing the CTL will require an ACS restart so you go to **System Configuration > Service Control** and click on the **Restart** button.
- Step 2** Select the Install Certificate link. Specify the location of the ACS certificate, and click on the **Submit** button.

Install New Certificate
Read certificate from file Certificate file: C:\files\certs\ACS-1.nac.cisco.com.cer Private key file: C:\files\certs\ACS-1.PrivateKey.txt Private key password: cisco123

- Step 3** After a successful installation of the ACS certificate, you must restart ACS. To do this select **System Configuration from the main menu, select Service Control** and click on the **Restart** button. This completes the ACS certificate installation process.

Task 7: Global Authentication Setup

ACS supports many different protocols for securely transferring credentials from the host to the ACS for authentication and authorization. You must tell ACS which protocols are allowed and what the default settings will be for each protocol.

Note: Unless you have a limited deployment environment or specific security concerns, it is highly recommended that you enable *all* protocols globally. You will have an opportunity to limit the actual protocol options later when you create the Network Access Profiles for NAC. But if they are not enabled here, they will not be available in the Network Access Profiles.

- Step 1** Select System Configuration from the main menu and pick Global Authentication Setup.
- Step 2** Select the following global authentication parameters to make them available in the Network Access Profile authentication configuration.

EAP Configuration	
PEAP	
☒	Allow EAP-MSCHAPv2
☒	Allow EAP-GTC
☒	Allow Posture Validation
	Cisco client initial message: <empty>
	PEAP session timeout (minutes): 120
	Enable Fast Reconnect: Yes
EAP-FAST	
EAP-FAST Configuration (see below)	
EAP-TLS	
☐	Allow EAP-TLS
Select one or more of the following options:	
☒	Certificate SAN comparison
☒	Certificate CN comparison
☒	Certificate Binary comparison
	EAP-TLS Session Timeout (minutes): 120
LEAP	
☐	Allow LEAP (For Aironet only)
EAP-MD5	
☐	Allow EAP-MD5
	AP EAP request timeout (seconds): 20

MS-CHAP Configuration	
☒	Allow MS-CHAP Version 1 Authentication
☒	Allow MS-CHAP Version 2 Authentication

- Step 1** Click “Submit + Restart” to save these changes.
- Step 2** Click “EAP-FAST Configuration” to enter the EAP-FAST screen.

EAP-FAST Settings	
EAP-FAST	
☒	Allow EAP-FAST
	Active master key TTL: 1 month
	Retired master key TTL: 3 month
	Tunnel PAC TTL: 1 week
	Client Initial Message: <empty>
	Authority ID Info: cisco

☐ Allow anonymous in-band PAC provisioning
☒ Allow authenticated in-band PAC provisioning

☒ Accept client on authenticated provisioning

☐ Require client certificate for provisioning
☒ Allow Machine Authentication

Machine PAC TTL **1 week**

☒ Allow Stateless Session Resume

Authorization PAC TTL **1 hour**

Allow inner methods
☒ EAP-GTC
☒ EAP-MSCHAPv2
☒ EAP-TLS
 Select one or more of the following EAP-TLS comparison methods:
☒ Certificate SAN comparison
☒ Certificate CN comparison
☒ Certificate binary comparison
 EAP-TLS session timeout (minutes):
☒ EAP-FAST master server
 Actual EAP-FAST server status: **Master**

Task 8: Configuring Attributes for Logging

Note: In order to log any non-Cisco NAC attribute values from the hosts, the attribute definitions must first be imported into ACS then selected for logging.

Step 1 To configure which log files are enabled and which event attributes are recorded within them, select the **System Configuration** option from the main menu then select **Logging**.

The recommended log files and their logged attributes for NAC are shown below. Your actual list of logged attributes will probably be longer depending upon which NAC vendor attributes are of interest in your deployment. The “Cisco:HIP:CSAxxx” attributes provide information on CSA posture, which will be used later in rules that determine overall host posture and therefore VLAN assignment.

CSV Failed Attempts

☒ Log to CSV Failed Attempts

Logged Attributes
Message-Type
User-Name
Caller-ID
Authen-Failure-Code
NAS-Port
NAS-IP-Address
AAA Server
Network Device Group
Access Device
PEAP/EAP-FAST-Clear-Name
Logged Remotely
EAP Type
EAP Type Name
Network Access Profile Name
Shared RAC
Downloadable ACL
System-Posture-Assessment
Application-Posture-Assessment
Reason
cisco-av-pair
Cisco:PA:PA-Name
Cisco:PA:PA-Version
Cisco:PA:OS-Type
Cisco:PA:OS-Version
Cisco:Host:ServicePacks
Cisco:Host:Hotfixes
Cisco:Host:Package
Cisco:HIP:CSAVersion
Cisco:HIP:CSAOperationalState
Cisco:HIP:CSAMCName
Cisco:HIP:CSAStates
Cisco:HIP:CSADaysSince
LastSuccessfulPoll

CSV Passed Authentications

☒ Log to CSV Passed Auths

Logged Attributes
Message-Type
User-Name
Caller-ID
NAS-Port
NAS-IP-Address
AAA Server
Filter Information
Network Device Group
Access Device
PEAP/EAP-FAST-Clear-Name
Logged Remotely
EAP Type
EAP Type Name
Network Access Profile Name
Outbound Class
Shared RAC
Downloadable ACL
System-Posture-Assessment
Application-Posture-Assessment
Reason
Cisco:PA:PA-Name
Cisco:PA:PA-Version
Cisco:PA:OS-Type
Cisco:PA:OS-Version
Cisco:Host:ServicePacks
Cisco:Host:Hotfixes
Cisco:Host:Package
Cisco:HIP:CSAVersion
Cisco:HIP:CSAOperationalState
Cisco:HIP:CSAMCName
Cisco:HIP:CSAStates
Cisco:HIP:CSADaysSince
LastSuccessfulPoll

CSV RADIUS

Accounting

☒ Log to RADIUS

Logged Attributes
User-Name
Group-Name
Calling-Station-Id
Acct-Status-Type
Acct-Session-Id
Acct-Session-Time
Acct-Input-Octets
Acct-Output-Octets
Acct-Input-Packets
Acct-Output-Packets
Framed-IP-Address
NAS-Port
NAS-IP-Address
Class
Termination-Action
Called-Station-Id
Acct-Delay-Time
Acct-Authentic
Acct-Terminate-Cause
Event-Timestamp
NAS-Port-Type
Port-Limit
NAS-Port-Id
AAA Server
ExtDB Info
Network Access
Profile Name
cisco-av-pair
Access Device
Logged Remotely

Administration Control

Task 9: Add Remote Administrator Access

To remotely administer your ACS via a web browser you must enable this feature by selecting the **Administration Control** button from the main menu. By adding one or more accounts, you can login to your ACS with HTTP.

- Step 1** Select the **Add Administrator** button and add the following information in the Administration Control section.

Administrator Name:	Administrator
Password:	cisco123
Administrator Privilege:	Grant All

Shared Profile Components

Shared Profile Components are configurations that can be reused across many different Network Access Profiles for filtering within ACS or for network authorizations within RADIUS. These will need to be defined before configuring Network Access Profiles.

Note: Network Access Profiles are new to ACS 4.0. They allow you to create and map individual authentication, posture validation, and authorization components depending on the access method being used.

The most useful include Downloadable IP Access Control Lists (ACLs) and RADIUS Authorization Components (RACs).

Task 10: Configure Downloadable IP ACLs (Omit)

This task is omitted since it applies to NAC-L2-IP or devices such as VPN concentrators and routers which are not used in this example.

Task 11: RADIUS Authorization Components (RACs)

RADIUS Authorization Components (RACs) are sets of RADIUS attributes that are applied to Network Access Devices during network authorizations.

Step 1 To configure RACs, go to **Shared Profile Components in the main menu and select RADIUS Authorization Components** and click on the **Add** button for each new RAC. Each RAC may contain one or more vendor RADIUS attributes including Cisco IOS/PIX 6.0, IETF, and Ascend.

Note: The Session-Timeout values used for NAC deployments can have a significant impact on ACS performance. It is strongly recommended that you adjust it for the scale of your network and ACS transaction capacity.

Step 2 Create the following RAC entries, attribute assignments, and values NAC-L2-802.1x:

RAC Name	Vendor	Assigned Attributes	Value
L2_1x_Healthy_RAC	IETF	Session-Timeout (27)	3600
	IETF	Termination-Action (29)	RADIUS-Request (1)
	IETF	Tunnel-Type (64)	[T1] VLAN (13)
	IETF	Tunnel-Medium-Type (65)	[T1] 802 (6)
	IETF	Tunnel-Private-Group-ID (81)	[T1] healthy
L2_1x_Transition_RAC	IETF	Session-Timeout (27)	30
	IETF	Termination-Action (29)	RADIUS-Request (1)
L2_1x_Quarantine_RAC	IETF	Session-Timeout (27)	3600
	IETF	Termination-Action (29)	RADIUS-Request (1)
	IETF	Tunnel-Type (64)	[T1] VLAN (13)
	IETF	Tunnel-Medium-Type (65)	[T1] 802 (6)
	IETF	Tunnel-Private-Group-ID (81)	[T1] quarantine

The following table is only for reference and lists all of the attributes that may be sent from ACS in a RADIUS-Accept response for NAC:

NAC-L2-802.1x	NAC-L2-IP	NAC-L3-IP	#	Attribute Name	Description
√			1	User-Name	Copied from EAP Identity Response in Access Request
	√	√	8	Framed-IP-Address	IP address of host
	√	√	26	Vendor-Specific Cisco (9,1) CiscoSecure-Defined-ACL	ACL name. Automatically sent by ACS.
√			26	Vendor-Specific Cisco (9,1) sec:pg	Policy-based ACL assignment. Only applies to Catalyst 6000. sec:pg = <group-name>
	√	√	26	Vendor-Specific Cisco (9,1) url-redirect	Redirection URL. url-redirect=<URL>
	√	√	26	Vendor-Specific Cisco (9,1) url-redirect-acl	Apply the named ACL for the redirect URL; ACL must be defined locally on the NAD. Only works on switches with IOS. url-redirect-acl=<ACL-Name>
√	√	√	26	Vendor-Specific Cisco (9,1) posture-token	Posture token/state name. Automatically sent by ACS.
	√	√	26	Vendor-Specific Cisco (9,1) status-query-timeout	Sets Status Query timer
	√	√	26	Vendor-Specific Cisco (9,1) host-session-id	Session identifier used for auditing. Automatically sent by ACS.
?	√	√	26	Vendor-Specific Microsoft = 311	Key for Status Query: MS-MPPE-Recv-Key. Automatically sent by ACS.
√	√	√	27	Session-Timeout	Sets Revalidation Timer (in seconds)
√	√	√	29	Termination-Action	Action on Session Timeout (0) Default: Terminate session (1) Radius-Request: Re-authenticate
√			64	Tunnel-Type	13 = VLAN
√			65	Tunnel-Medium-Type	6 = 802
√	√	√	79	EAP Message	EAP Request/Response Packet in Access Request and Access Challenge: - EAP Success in Access Accept - EAP Failure in Access Reject
?	?	?	80	Message Authenticator	HMAC-MD5 to ensure integrity of packet.
√			81	Tunnel-Private-Group-ID	VLAN name

Group Setup

For this example, local usernames and groups defined in ACS will be used for authentication. There is a supplemental lab in the NAC2 Lab Guide that provides information on integrating active directory with ACS for user and group authentication.

Group and User Setup			
Group Number	Group Name	Local ACS Users	Password
1: Group 1	Employees	Administrator	cisco
1: Group 1	Employees	employee1	cisco
2: Group 2	Contractors	contractor1	cisco
3: Group 3	Guest	guest1	cisco
4: Group 4	Utilities	Utilities1	cisco

Step 1 Click on **Group Setup** in the main ACS menu.

Step 2 Click on **Rename** and rename the first three default groups as shown in the table below under Group Setup from the main menu. You will also need to rename another group to Utilities for use with application-specific devices (ASDs).

User Setup

Step 1 Click on **User Setup** in the main ACS menu. In the **User** dialog box type in the first username as shown below: **employee1** and select the **Add/Edit** button.

Step 2 In the **User: employee1 (New User)** screen under **User Setup** enter **cisco** as the user's password. In the **Group to which the user is assigned** drop down box assign the user to the **Employees** group. Scroll to the bottom and click the **Submit** button.

Step 3 Repeat these steps for creating the remaining users: **contractor1**, **guest1**, **utilities**. (Optional)

Note: The individual RADIUS attributes will be configured and applied in the Network Access Profile section and do not need to be configured for each individual group.

Posture Validation

Posture Validation is a core component of the NAC configuration. In the Posture validation section rules are created to validate hosts and client posture compliance. Tokens are delivered to the NAD granting or denying network access as a result of this compliance. The resulting tokens could be healthy, checkup, transition, quarantine, infected, and unknown.

ACS can use the following to perform posture validation:

- locally within ACS
- externally using the HCAP protocol to one or more posture validation servers (PVS)
- externally using the GAME protocol to an audit server for NAC Agentless Host (NAH) support

Note: You can perform both local and external posture validation at the same time. However you can not perform local and external posture validation for the same NAC credential types (vendor/application combinations). Example: Verifying Trend Micro information locally in ACS and externally in the Trend Policy Server.

Posture validation policies are configured in ACS under **Posture Validation** in the main menu. These policies are later selected and applied to network access profiles. The policies are defined separately so that you may mix and match or reuse them to provide differentiated access for multiple network services across many locations.

Task 12: Internal Posture Validation Setup

Posture validation policies consist of rules, and these rules are built from a set of conditions. Each of these conditions can match a received credential from the client, and result in a potential policy assessment.

- Step 1** To create the policy requirements of the reference network locally on the ACS, the set of NAC posture validation policies should be defined, as per the table below. To create these policies, select **Posture Validation** from the main menu and then **Internal Posture Validation Setup**.
- Step 2** Select **Add Policy** to create a new policy.
- Step 3** Enter a name, and optionally a description, for your new Posture Validation Policy, and then click on **Submit**.
- Step 4** You will now enter the specific Posture Validation Rules for your new policy. You will build a set of Conditions for the given policy that match to a specific posture assessment result. To create this rule, click on **Add Rule**.
- Step 5** Click on **Add Condition Set** to move to the screen where you will define the conditions that compose the posture validation rule.
- Step 6** Using the table below, add the appropriate Attributes, Operators and Values for the needed condition and click on **Submit**. For example, to set a condition for validating the CTA version on the client, choose the “Cisco:PA:PA-Version” credential from the Attribute menu, change the Operator to “>=”, enter “2.0.0.25” into the Value field and click “enter”. If you need to evaluate multiple credentials together, continue to add those conditions into the rule. It is important to note that to evaluate these conditions together as a single rule, after selecting **Submit**, choose the modal option “Match ‘OR’ inside Condition and ‘AND’ between Condition Sets” and click **Submit** again.
- Step 7** Click **Done** to return the original Posture Validation Rules screen.
- Step 8** After any and all changes have been made in the Rules, click on **Apply and Restart** at the bottom of the page.

Policy Name	#	Condition	Posture Assessment	Notification String
CTA	1	Cisco:PA:PA-Version >= 2.0.0.25 AND Cisco:PA:Machine-Posture-State >= 1	Cisco:PA:Healthy	
	2	Default	Cisco:PA:Quarantine	
Windows	1	(Cisco:PA:OS-Type contains Windows XP AND Cisco:Host:ServicePacks contains 2) OR (Cisco:PA:OS-Type contains Windows 2000 AND Cisco:Host:ServicePacks contains 4)	Cisco:Host:Healthy	
	2	Default	Cisco:Host:Quarantine	
CSA	1	Cisco:HIP:CSAOperationalState = 1 AND Cisco:HIP:CSAVersion >= 5.0.0.0	Cisco:HIP:Healthy	
	2	Default	Cisco:HIP:Quarantine	

Note: You may choose to add additional AVPs to the CSA policy to increase your level of trust, such as Cisco:HIP:CSAMCName contains <your MC URL> or Cisco:HIP:CSADaysSinceLastSuccessfulPoll <= 5.

Network Access Profiles

Initial ACS configuration is now done. Network Access Profiles will be configured in the next section.

NAC-L2-802.1x Configuration

The following sections will cover NAC-L2-802.1x configuration for switches running IOS and creation of individual Network Access Profiles for the NAC-L2-802.1x access method.

NAC-L2-802.1x for IOS Switches

In this section you will configure the various components to enable the base functionality of NAC-L2-802.1x. The following overview shows the steps to be performed in this section:

- Step 1** Configure the NAD for NAC-L2-802.1x
- Step 2** Verify port goes from connecting to authenticating to verify 802.1x NAD port configuration
- Step 3** Configure the ACS for NAC-L2-802.1x
- Step 4** Install Cisco Trust Agent (if not already installed as part of CSA)
- Step 5** Test NAC-L2-802.1x

NAC-L2-802.1x Deployment Method Overview

Before beginning the configuration section for 802.1x NAC it is important to understand there are two ways to deploy 802.1x; NAC-L2-802.1x and traditional 802.1x using a legacy (e.g. Windows) supplicant.

The first method, NAC-L2-802.1x, uses a NAC-enabled 802.1x supplicant to perform identity *and* posture credential validation within an 802.1x access control conversation.

The second method for deployment uses a non-NAC enabled 802.1x supplicant to do the identity credential validation for port access and then use NAC-L2-IP to do a posture credential validation after the endpoint has an IP address and triggered a NAC-L2-IP interrogation. In this method, access is enforced using ACLs downloaded to the access port rather than through VLAN assignment.

The primary difference between the two options at the protocol level is the EAP method used to combine identity and posture in the client to server communication. A NAC enabled 802.1x must use EAP-FAST for the EAP method, since it has been modified to carry identity and posture credentials in a TLS tunnel. The CTA supplicant supports EAP-GTC, EAP-MSCHAPv2 and EAP-TLS for client side authentication.

The following sections will guide you through the configuration of the first method: NAC-L2-802.1x utilizing the NAC-enabled supplicant in CTA 2.0 *This is the recommended option for Trusted Endpoint QoS marking.*

In order to gain a better understanding of NAC-L2-802.1x lets begin by discussing two aspects of the credentials which can be sent from the client to the network. In a Microsoft Windows environment there are two set of identity credentials that can be presented to the network.

NAC-L2-802.1x Credential Overview

The first credential involves the concept of machine authentication where the machine is authenticated in advance of the user of the computer. Microsoft introduced the machine authentication facility to allow the client system to authenticate using the identity and credentials of the computer at boot time so that the client can establish the required secure channel to update and participate in the domain GPO (Group Policy Objects) model.

Machine authentication allows the computer to authenticate itself to the network using 802.1x, just after a PC loads device drivers at boot time. This allows the computer to subsequently communicate with Windows domain controllers in order to pull down machine group policies. This was designed to alleviate the problem of domain GPOs being broken by the introduction of 802.1x.

The second type of credential type of credential used for 802.1x is referred to as user authentication. After the GINA (login screen) is presented, a user can login to the computer or the Windows domain and the username/password used for login can be used as the identity credentials for 802.1x authentication..

In a NAC-L2-802.1x environment, the CTA supplicant uses EAP-FAST to perform machine and user authentication. EAP-FAST uses a Protected Access Credential (PAC) to mutually authenticate the client and RADIUS server. The PAC is a unique shared credential used to mutually authenticate client and server. It is associated with a specific client username and a server authority ID. A PAC removes the need for Public Key Infrastructure (PKI) and digital certificates. EAP-FAST is detailed here:

<http://www.cisco.com/univercd/cc/td/doc/product/wireless/airo1200/accsspts/techref/eapfast/eapfast.htm>

EAP-FAST comprises three basic phases:

- Phase 0 (optional): the PAC is initially distributed to client.
- Phase 1: using the PAC, a secure tunnel is established.
- Phase 2: the client is authenticated via the secure tunnel.

In the EAP-FAST specification there are two ways to provision the PAC, out-of-band-provisioning or in-band-provisioning. With the NAC-L2-802.1x CTA supplicant you can only provision a PAC with in-band-provisioning. The CTA supplicant will only provision a PAC on the host if the ACS server has been configured to allow in-band-provisioning and if the client side authentication is a successful machine authentication using a certificate assigned to the machine (machine certificate) or a successful user authentication. Out-of-band provisioning is not supported with the NAC-L2-802.x CTA supplicant.

For simplicity, only user authentication will be configured in this example. Authentication will be performed against local username and password database in ACS.

Configure the Catalyst 3750 for NAC-L2-802.1x

Task 1: VLANs for NAC-L2-802.1x

Configure the following minimum VLANs and VLAN interfaces on the NAD for NAC-L2-802.1x:

VLAN Name	VLAN	3750 Subnets
healthy	50	10.7.50.*
quarantine	80	10.7.80.*
voice	110	10.7.110.*
servers	200	10.0.200.*

Note: The VLAN names defined here must exactly match any **IETF Tunnel-Private-Group-ID (81)** parameters defined in the section ACS Common Configuration: Shared Profile Components: Task 11: RADIUS Authorization Components (RACs)

```

vlan 50
  name healthy
!
vlan 80
  name quarantine
!
vlan 110
  name voice
!
vlan 200
  name servers
!

```

```

interface Vlan50
  description Healthy VLAN
  ip address 10.7.50.1 255.255.255.0
  ip helper-address 10.0.200.10      ! Address of DHCP server
  !
interface Vlan80
  description Quarantine VLAN
  ip address 10.7.80.1 255.255.255.0
  ip helper-address 10.0.200.10
  !
interface Vlan110
  description Voice VLAN
  ip address 10.7.110.1 255.255.255.0
  ip helper-address 10.0.200.10
  !
interface Vlan200
  description Servers VLAN
  ip address 10.0.200.1 255.255.255.0
  ip helper-address 10.0.200.10
  !

```

Task 2: Configure AAA on NAD for NAC-L2-802.1x

This section describes the minimum steps required to enable AAA for NAC-L2-802.1x on an IOS switch.

Step 1 Enable the switch AAA service using the **aaa new-model** global configuration command.

```
CAT3750(config)#aaa new-model
```

Step 2 Configure the switch to use RADIUS for 802.1x authentication using the **aaa authentication dot1x default group radius** global configuration command.

```
CAT3750(config)#aaa authentication dot1x default group radius
```

Step 3 Configure the switch to run authorization for all network-related service requests using the **aaa authorization network default group radius** global configuration command.

```
CAT3750(config)#aaa authorization network default group radius
```

Step 4 Enable AAA accounting for 802.1x accounting using the **aaa accounting dot1x default start-stop group radius** global configuration command.

```
CAT3750(config)#aaa accounting dot1x default start-stop group radius
```

Step 5 Specify the NAD interface for all outgoing RADIUS packets using the **ip radius source-interface** global configuration command.

```
CAT3750(config)#ip radius source-interface Vlan200
```

Task 3: Configure the RADIUS Server

These are the minimum steps required to configure a RADIUS server on Cisco IOS:

Step 1 Specify the hostname or IP address of the RADIUS server (and optionally the authentication and accounting ports) using the **radius-server host** global command. The default RADIUS port number for authentication is 1645. The default RADIUS port number for accounting is 1646.

```
CAT3750(config)#radius-server host 10.0.200.20
```

- Step 2** Specify the RADIUS server encryption key using the **radius-server key** global command. Note that this key must match the key configured in the Cisco Secure ACS server for this NAD. If they do not match, the NAD and the Cisco Secure ACS will not be able to communicate posture validation information. See the section ACS Common Configuration: Network Configuration: Task 2: AAA Clients.

```
CAT3750(config)#radius-server key cisco123
```

- Step 3** Configure the switch to send Framed-IP-Address RADIUS attribute (Attribute[8]) in access-request or accounting-request packets using the **radius-server attribute 8 include-in-access-req** global command.

```
CAT3750(config)#radius-server attribute 8 include-in-access-req
```

- Step 4** Configure the NAD to recognize and use vendor-specific attributes using the **radius-server vsa send authentication** global command.

```
CAT3750(config)#radius-server vsa send authentication
```

- Step 5** Specify the NAD interface for all outgoing RADIUS packets using the **ip radius source-interface** global command.

```
CAT3750(config)#ip radius source-interface Vlan200
```

Note: Step 5 is optional; however it is recommended if there are multiple paths between the NAD and the Cisco Secure ACS. Assigning a source interface allows the Cisco Secure ACS to know from which NAD the RADIUS messages originated. This is the same IP address that must be configured in the Cisco Secure ACS AAA client record that represents this NAD.

Task 4: Enable 802.1x on the NAD

- Step1** Enable 802.1x using the **dot1x system-auth-control** global configuration command.

```
CAT3750(config)#dot1x system-auth-control
```

Task 5: Configure 802.1x on the NAD Interfaces

- Step 1** Enable the 802.1x port control to auto on Fast Ethernet ports using **dot1x port-control auto** interface command.

```
CAT3750(config)# interface range fa1/0/1-24
```

```
CAT3750(config-if-range)#dot1x port-control auto
```

- Step 2** Set the 802.1x reauthentication timer to use the timer set in ACS using the **dot1x timeout reauth-period server** interface command.

```
CAT3750(config-if-range)#dot1x timeout reauth-period server
```

- Step 3** Enable 802.1x reauthentication for the interface using the **dot1x reauthentication** interface command.

```
CAT3750(config-if)#dot1x reauthentication
```

- Step 4** If you have multiple hosts connecting per interface you can enable multi-host support using the **dot1x host-mode multi-host** interface command. (Optional.)

```
CAT3750(config-if-range)#dot1x host-mode multi-host
```

Test basic 802.1x flow between the client and the NAD

Without ACS there is not a lot of information that can be gathered from the NAD. However, you can verify the switchport port goes from connecting to authenticating to verify 802.1x communication is occurring between the client and the NAD.

```
CAT3750#sh dot1x interface fa1/0/1
```

```
Supplicant MAC 000d.80cd.cda6
```

```
AuthSM State = CONNECTING
```



```
BendSM State      = IDLE
Posture            = N/A
ReAuthPeriod      = None (From Authentication Server)
ReAuthAction      = N/A
TimeToNextReauth  = N/A
PortStatus        = UNAUTHORIZED
MaxReq            = 2
MaxAuthReq        = 2
HostMode          = Single
PortControl       = Auto
ControlDirection  = Both
QuietPeriod       = 60 Seconds
Re-authentication = Enabled
ReAuthPeriod      = From Authentication Server
ServerTimeout     = 30 Seconds
SuppTimeout       = 30 Seconds
TxPeriod          = 30 Seconds
Guest-Vlan        = 0
```

```
CAT3750#sh dot1x interface GigE 1/1
Supplicant MAC 000d.80cd.cda6
AuthSM State      = AUTHENTICATING
BendSM State      = RESPONSE
Posture           = N/A
ReAuthPeriod      = None (From Authentication Server)
ReAuthAction      = N/A
TimeToNextReauth  = N/A
PortStatus        = UNAUTHORIZED
MaxReq            = 2
MaxAuthReq        = 2
HostMode          = Single
PortControl       = Auto
ControlDirection  = Both
QuietPeriod       = 60 Seconds
Re-authentication = Enabled
ReAuthPeriod      = From Authentication Server
ServerTimeout     = 30 Seconds
SuppTimeout       = 30 Seconds
TxPeriod          = 30 Seconds
Guest-Vlan        = 0
```

Network Access Profile Configuration for NAC-L2-802.1x

In the following section you will configure a Network Access Profile (authentication, posture validation, and authorization) to support NAC-L2-802.1x. In ACS 4.0 there are two methods of configuring Network Access Profiles:

- Add an empty profile and configure all the necessary information
- Using the Template Profiles, customize the Network Access Profile desired with the base information included in the template.


There are seven Network Access Profile templates pre-defined in ACS 4.0:

- NAC L3 IP
- NAC-L2-IP
- NAC-L2-802.1x
- Microsoft IEEE 802.1x
- Wireless (NAC-L2-802.1x)
- Authentication Bypass (802.1x fallback)
- Agentless Host

In this section we will use the NAC-L2-802.1x Network Access Profile template to create a base profile and will then make the necessary changes to customize this template.

Task 6: Create the NAC-L2-802.1x profile from the template.

- Step 1** Click on Network Access Profiles from the main menu and select **Add Template Profile**.
- Step 2** Create a template for NAC-L2-802.1x by selecting it from the **Template** drop down box. Name the template with something similar to the one shown below. Be sure to select **Active** to enable the profile.



- Step 3** Click **Submit**.

Note: You will see sample RADIUS Authorization Components created as part of the ACS 4.0 templates. You can ignore these for this example.

Task 7: Authentication

- Step 1** On the **Network Access Profiles** screen select the **Authentication** link for the new profile.

☐	Lab NAC L2 802.1X	Authentication Posture Validation Authorization	Template for NAC Lab	YES
-----------------------	-----------------------------------	---	----------------------	-----

Step 2 Notice that a portion of the EAP-FAST configuration is already selected as part of the base template.

Network Access Profiles

EAP-FAST ☒ Allow EAP-FAST ☐ Allow anonymous in-band PAC provisioning ☒ Allow authenticated in-band PAC provisioning ☒ Accept client on authenticated provisioning ☐ Require client certificate for provisioning ☐ Allow Stateless session resume Authorization PAC TTL 1 hours Allowed inner methods ☒ EAP-GTC ☒ EAP-MSCHAPv2 ☐ EAP-TLS Posture Validation: ☐ None ☒ Required ☐ Optional - Client may not supply posture data. Use token Unknown ☐ Posture only EAP-TLS ☐ Allow EAP-TLS EAP-MD5 ☐ Allow EAP-MD5	Submit Cancel
---	-----------------------------

Step 3 Click **Submit**.

Note: If you enable EAP-GTC and use the CTA supplicant you will always be prompted for user credentials. This will occur even if you have configured the supplicant to use single sign-on.

Task 8: Posture Validation

The posture validation configuration below is for the reference network.

Step 1 Select the **Posture Validation** link for the profile from the **Network Access Profile** screen.

Step 2 Add the following posture policies to the template:

Name: L2-Posture				
Required Condition Types	Cisco:PA Cisco:Host			
Posture Validation Policies	CTA Windows			
Assessment ResultConfiguration	Result	Message	URL Redirect	
	Healthy	NAC-L2-802.1x Healthy		
	Checkup	Please update your software to prevent being quarantined by the network.		
	Transition	Computer under audit...		
	Quarantine	NAC-L2-802.1x Quarantined		
	Infected	Infected		
	Unknown			
Audit Selection				
Audit Server	None			

Step 3 Click **Submit**.

Task 9: Authorization

Step 1 Select the **Authorization** link from the template.

Step 2 Enable authorization.

User Group	Assessment Result	Deny Access	Shared RAC	Downloadable ACL
Employees	Healthy	No	L2_1x_Healthy_RAC	
Contractors	Healthy	No	L2_1x_Healthy_RAC	
Any	Healthy	No	L2_1x_Healthy_RAC	
Guests	Any	No	L2_1x_Quarantine_RAC	
Utilities	Any	No	L2_1x_Quarantine_RAC	
If a condition is not defined or there is no matched condition:			L2_1x_Quarantine_RAC	
Include RADIUS attributes from user's group:				No
Include RADIUS attributes from user record:				No

Step 3 Click **Submit**.

Cisco Trust Agent (CTA) Installation and Configuration

CTA is required to perform posture validation of the client. The Cisco Trust Agent is normally installed as part of a CSA agent kit but may also be installed using one of the ctasetup.exe files.

Note: The CTA Administrator Guide 2.0 provides detailed information regarding CTA files and installation.
http://www.cisco.com/en/US/partner/products/ps5923/products_maintenance_guide_book09186a008059a40e.html

Cisco Trust Agent Windows .exe versions

CTA for Windows provide several options for deploying and packaging CTA. Administrators can deploy the scripting interface and/or the supplicant for Windows as noisy or silent installs. The available packages for Windows are as follows:

CTA .exe Files	Description
ctasetup-win-[version].exe	If you use this package, the install is “noisy”. This means that the end user will be prompted to accept a license agreement, choose the install destination folder, and other general installation options. Additionally, this package only installs the CTA scripting interface. The supplicant is not installed using this package.
ctasetup-suppliant-win-[version].exe	If you use this package, the install is “interactive”. This means that the end user will be prompted to accept a license agreement, choose the install destination folder, and other general installation options. Additionally, this package can install both the CTA scripting interface and the supplicant. The end user is prompted to select which CTA features they want to install.
CtaAdminEx-win-[version].exe	If you use this package, you are choosing to create a silent installation package for the end user. You extract a file named ctasilent-win-[version].exe from this package. As the administrator, you accept the license agreement for endusers and then deploy the ctasilent-win-[version].exe file as a completely silent install that does not prompt the end user for any options. The supplicant is not installed using this package.
CtaAdminEx-suppliant-win-[version].exe	If you use this package, you are choosing to create a silent installation package for the end user. You extract a file named ctasilent-suppliant-win-[version].exe from this package. As the administrator, you accept the license agreement for endusers and then deploy the ctasilent-suppliant-win-[version].exe file as a completely silent install that does not prompt the end user for any options. The supplicant is installed using this package.

Note: The supplicant is required for clients to connect to and access a network that is protected by the IEEE 802.1x security protocol. Only after successful client-server authentication will the port access control on the 802.1x-enabled access device (the Ethernet switch) allow the end-user to connect to the network.

Task 10: Client Certificate for CTA Install

CTA must install the certificate you have installed on ACS to properly authenticate. There are two methods available to add the certificate to CTA on the client. The first method which will be shown here should be done prior to CTA install. The other can be used to add the certificate to the root store after CTA is installed. The second method will be shown at the end of this section. This section assumes you have certificates from the lab, or have generated them.

- Step 1** Move the \certs folder into the folder containing the CTA .exe installation file. The \certs folder contains the CA certificate that must be used by CTA to authenticate the client to ACS.

Note: CTA will import any public certificate located in the \certs subdirectory. This folder must be located in the same directory as the cta.exe file.


Task 11: Install CTA 2.0

- Step 1** Open the CTA build folder located on the client desktop and double click the appropriate ctasetup file. (For this example this will be ctasetup_suppllicant-win-[version].exe.).

The **Cisco Trust Agent Installation Wizard** appears:




- Step 2** Click **Next**.
- Step 3** Accept the license agreement by clicking **Next**; the **Destination Folder** window appears:




Step 4 Accept the default **Destination Folder** location and click **Next**.

Step 5 The **Select Installation Type** dialog box appears.



- Step 6** Click the **Complete** radio button.
- Step 7** Select your features and then click **Next**.
- Step 8** Click **Next**.
- Step 9** The application installs to the selected directory.
- Step 10** The following message will be displayed when the certificate is successfully imported during the install. Click **OK**.



- Step 11** When the installation is completed, the installer displays the **Installation Completed** window.
- Step 12** Click **Finish** to close the installation application. You may need to restart your system. You will be prompted to do so if it's necessary.

Task 12: (Optional) Manual install of root certificate for CTA

If you did not copy the "certs" folder into the CTA setup folder in [Task 1: step 1](#), you need to install the root certificate before using Cisco Trust Agent.

You can manually install the certificate using the following steps

- Step 1** Copy the certificate to the network client.
- Step 2** Open a command prompt on the network client.
- Step 3** Change directory to where Cisco Trust Agent is installed. By default, the location is C:\Program Files\Cisco Systems\CiscoTrustAgent\.
- Step 4** Enter `ctaCert.exe /add "cert_pathname" /store "Root"`, where cert_pathname is the full path and file name to the certificate.

The certificate is added to the trusted certificate store on the network client.

CTA Configuration for NAC-L2-802.1x

There are no configuration changes to be done on the supplicant. The supplicant will ALWAYS try to do machine authentication. ACS will issue a RADIUS Access-Reject in response to the first RADIUS Request from the network access device in order to disallow machine authentication. The network administrator should be aware that this will cause the 802.1x state machine on the switch to move the port into the “held” which will disallow the switch from accepting any EAPOL-Starts from the client. If the supplicant sends an EAPOL-Start while the switch is in the held state then the end user may experience a slower than normal login experience since the supplicant will wait to complete the user login until the switch moves the 802.1x state machine to a “connecting” state and accepts the EAPOL-Starts from the supplicant to initiate a successful 802.1x exchange.


Verify NAC-L2-802.1x Functionality

This section will help you validate NAC-L2-802.1x is configured properly, that you are being passed the correct posture token from ACS, and the correct VLAN assignment from ACS are being applied on the NAD.

To be considered healthy and to be placed in the healthy role the client must correctly pass back the required credential information to the NAD and then to ACS. The posture validation requirements for each credential created in the previous Network Access Profile section must be met for the client to be passed an application posture token of “healthy”. The credentials include: **Cisco Trust Agent**, the agent version is **>=2.0.0.25**, and the OS-Type contains **Windows XP**.

Note: It is important to remember, the client will only pass to ACS the credentials ACS is specifically requesting.

- Step 1** First, re-enable the switchport to which the client is connected by issuing the **no shut** command.
- Step 2** On the client you should see a credential request from the supplicant similar to the following:





Step 3 Issue the **show dot1x all** command to verify the current status of the client.

```
CAT3750#sh dot1x all
Dot1x Info for interface FastEthernet 1/0/1
-----
Supplicant MAC 000d.80cd.cda6
  AuthSM State      = AUTHENTICATED
  BendSM State      = IDLE
  Posture            = Healthy
  ReAuthPeriod       = 3600 Seconds (From Authentication Server)
  ReAuthAction       = Terminate
  TimeToNextReauth   = 3570 Seconds
PortStatus          = AUTHORIZED
MaxReq               = 2
MaxAuthReq           = 2
HostMode             = Single
PortControl          = Auto
ControlDirection     = Both
QuietPeriod          = 60 Seconds
Re-authentication    = Enabled
ReAuthPeriod         = From Authentication Server
ServerTimeout        = 30 Seconds
SuppTimeout          = 30 Seconds
TxPeriod             = 30 Seconds
Guest-Vlan           = 0
```

Step 4 Verify the client switchport has been placed in the correct VLAN.

```
CAT3750#sh vlan
```


VLAN	Name	Status	Ports
1	default	active	Fa1/0/5, Fa1/0/6, Fa1/0/7, Fa1/0/8 Fa1/0/9, Fa1/0/10, Fa1/0/13, Fa1/0/15 Fa1/0/16, Fa1/0/17, Fa1/0/18, Fa1/0/19 Fa1/0/20, Fa1/0/21, Fa1/0/22, Fa1/0/23 Fa1/0/24, Gi1/2
10	employees	active	Fa1/0/3
20	contractors	active	
30	utilities	active	
40	guests	active	
50	healthy	active	Fa1/0/1
60	checkup	active	
70	transition	active	
80	quarantine	active	

90	infected	active	
100	unknown	active	Fa1/0/4, Fa1/0/11, Fa1/0/14
110	voice	active	
200	servers	active	Fa1/0/12
255	nads	active	

Step 5 You can see the client switchport has been placed in VLAN 50 (healthy). Alternatively, you can also use the following interface command to view the switchport information for the interface:

```
CAT3750#sh int fa1/0/1 switchport
Name: Fa1/0/1
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: native
Negotiation of Trunking: Off
Access Mode VLAN: 50 (healthy)
...
```

Step 6 On the client, CTA should provide a pop-up message similar to the following:



Step 7 Click on the ACS **Reports and Activity** button to verify the client information in the appropriate report. Because it appears that we have correctly established communications between the client and ACS, the most appropriate report to check is **Passed Authentications**.

IOS QoS Configuration Examples

The IOS QoS configuration examples presented here are for the Catalyst 3750 switch. The QoS features and configuration syntax are identical to the Catalyst 3650 and 2970 and pertain to those switches as well. For complete details, see the “Configuring QoS” chapter of the *Catalyst 3750 Switch Software Configuration Guide, 12.2(25)SE* at:

http://www.cisco.com/en/US/products/hw/switches/ps5023/products_configuration_guide_chapter09186a00802c1100.html

The main idea behind the Trusted Endpoint QoS marking solution is that classification and marking are now performed on the host by CSA 5.0. The access edge switch port is configured to trust the DSCP markings set by CSA but it still applies access edge policing and queueing policies. This access model falls between the trusted and untrusted host models discussed earlier and might be described as **trust but verify**.

QoS is globally disabled by default on the Catalyst 3750. While QoS is disabled, all frames/packets are passed through the switch unaltered (which is equivalent to a trust DSCP state on all ports). When QoS is globally enabled, however, all DSCP values are set to 0 by default (which is equivalent to an untrusted state on all ports).

QoS must be enabled globally for configured policies to become effective. The example below shows how to verify if QoS has been enabled or not and also how it can be globally enabled.

Enabling QoS Globally on the Catalyst 3750

```
CAT3750#show mls qos
QoS is disabled

CAT3750#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
CAT3750(config)#mls qos
CAT3750(config)#end
CAT3750#

CAT3750#show mls qos
QoS is enabled

CAT3750#
```

The first step in defining the QoS policy is to set up the classification configuration on the ingress ports. Since the result of the NAC-L2-802.1x exchange is a dynamic VLAN assignment based on the security posture of the host, this fact can be used to extend trust to hosts or not. Some switches (e.g. Catalyst 3550) support a **match vlan** command in the **class-map** configuration, which simplifies this task. On the Catalyst 3750, we instead use the one-to-one correspondence between VLAN and subnet to create ACLs that match incoming traffic. Note that in this example we only show policies for typical IP telephony voice traffic.

Define ACLs that Match Marked Traffic From Trusted VLANs on the Catalyst 3750

```
ip access-list extended VVLAN-VOICE
 permit udp 10.1.120.0 0.0.0.255 any range 16384 32767 dscp ef
 ! Voice is matched by VVLAN subnet and DSCP EF
ip access-list extended VVLAN-CALL-SIGNALING
 permit tcp 10.1.120.0 0.0.0.255 any range 2000 2002 dscp af31
 permit tcp 10.1.120.0 0.0.0.255 any range 2000 2002 dscp cs3
 ! Call-signaling is matched by VVLAN subnet and DSCP AF31 or CS3
ip access-list extended VVLAN-OTHER
 permit ip 10.1.120.0 0.0.0.255 any
 ! Matches all other traffic sourced from the VVLAN subnet
ip access-list extended HEALTHY-VOICE
 permit udp 10.1.50.0 0.0.0.255 any dscp ef
 ! Voice is matched by healthy subnet and DSCP EF
ip access-list extended HEALTHY-VIDEO
 permit udp 10.1.50.0 0.0.0.255 any dscp af41
 ! Voice is matched by healthy subnet and DSCP AF41
ip access-list extended HEALTHY-CALL-SIGNALING
 permit ip 10.1.50.0 0.0.0.255 any dscp cs3
 ! Call-signaling is matched by healthy subnet and DSCP AF31 or CS3
```

```
ip access-list extended HEALTHY-OTHER
permit ip 10.1.50.0 0.0.0.255 any
! Matches all other traffic sourced from the healthy VLAN subnet
```

A class map is a mechanism that is used to name a specific traffic flow (or class) and to uniquely identify it from all other traffic. The criteria can include matching the access group defined by an ACL or matching a specific list of DSCP or IP precedence values. If you have more than one type of traffic that you want to classify, you can create another class map and use a different name. After a packet is matched against the class-map criteria, you assign actions through the use of a policy map.

Map the Traffic Match ACLs onto Classes on the Catalyst 3750

```
class-map match-all VVLAN-VOICE
match access-group name VVLAN-VOICE
class-map match-all VVLAN-CALL-SIGNALING
match access-group name VVLAN-CALL-SIGNALING
class-map match-all VVLAN-OTHER
match access-group name VVLAN-OTHER
class-map match-all HEALTHY-VOICE
match access-group name HEALTHY-VOICE
class-map match-all HEALTHY-VIDEO
match access-group name HEALTHY-VIDEO
class-map match-all HEALTHY-CALL-SIGNALING
match access-group name HEALTHY-CALL-SIGNALING
```

A policy map specifies which traffic class to act on. Actions can include trusting the CoS, DSCP, or IP precedence values in the traffic class; setting a specific DSCP or IP precedence value in the traffic class; or specifying the traffic bandwidth limitations and the action to take when the traffic is out of profile. In the example below, incoming DSCP markings may either be trusted or reset, and excess traffic may either be dropped or marked down, depending on the traffic class. For a policy map to become effective, you must attach it to a port.

Define Ingress Policing Policy on the Catalyst 3750

```
policy-map IPPHONE+CSA-PC
class VVLAN-VOICE
trust dscp
police 128000 8000 exceed-action drop
! Only one voice call is permitted per switchport VVLAN
class VVLAN-CALL-SIGNALING
set ip dscp 24
police 32000 8000 exceed-action policed-dscp-transmit
! Out of profile Call-Signaling is marked down to Scavenger (CS1)
class VVLAN-OTHER
set ip dscp 0
police 32000 8000 exceed-action policed-dscp-transmit
! Unauthorized VVLAN traffic is marked down to Scavenger (CS1)
class HEALTHY-VOICE
trust dscp
police 128000 8000 exceed-action drop
! Only one voice call is permitted per switchport DVLAN
class HEALTHY-VIDEO
trust dscp
police 500000 8000 exceed-action drop
! Only one video call is permitted per switchport DVLAN
class HEALTHY-CALL-SIGNALING
set ip dscp 24
police 32000 8000 exceed-action policed-dscp-transmit
! Out of profile Call-Signaling is marked down to Scavenger (CS1)
class class-default
set ip dscp 0
police 5000000 8000 exceed-action policed-dscp-transmit
! Unauthorized VVLAN traffic is marked down to Scavenger (CS1)
```

During policing, QoS can assign another DSCP value to an IP packet if the packet is out of profile and the policer specifies a markdown (rather than drop) exceed-action. This configurable map is called the policed-

DSCP map. You configure this map by using the **mls qos map policed-dscp** global configuration command. In this example, the specified out of profile DSCPs will be remapped to Scavenger class.

Define Markdown Behavior on the Catalyst 3750

```
mls qos map policed-dscp 0 10 18 24 25 34 to 8
! Excess DVLAN traffic marked 0, AF11, AF21, CS3,
! DSCP 25 or AF41 will be remarked to Scavenger (CS1)
```

A full discussion of queueing models is beyond the scope of this document; please refer to the references cited above. In brief, the Catalyst 3750 supports four egress queues each having three drop thresholds (where the third threshold is fixed at 100%.) In addition, the first queue can be used as a strict priority queue. In the example below, all eleven service classes defined in the Cisco QoS Baseline are mapped onto the four output queues of the Catalyst 3750 according to the queueing rules discussed in chapter one. A strict priority queue is used for Realtime (voice) traffic and the other three queues are assigned to Critical Data, Best Effort, and Scavenger/Bulk traffic, respectively. The drop thresholds are configured in order to limit the bandwidth of certain classes within the queues.

Map Eleven Traffic Classes onto Four Output Queues on the Catalyst 3750

```
mls qos srr-queue output dscp-map queue 1 threshold 3 46
! DSCP EF (Voice) gets all of Queue 1 (PQ)
mls qos srr-queue output dscp-map queue 2 threshold 1 16
! Maps DSCP CS2 (Network Management) to Q2 T1
mls qos srr-queue output dscp-map queue 2 threshold 1 18 20 22
! Maps DSCP AF21, AF22, AF23 (Transactional Data) to Q2 T1
mls qos srr-queue output dscp-map queue 2 threshold 1 26 28 30
! Maps DSCP AF31, AF32, AF33 (Mission-Critical Data) to Q2 T1
mls qos srr-queue output dscp-map queue 2 threshold 1 32
! Maps DSCP CS4 (Streaming Video) to Q2 T1
mls qos srr-queue output dscp-map queue 2 threshold 1 34 36 38
! Maps DSCP AF41, AF42, AF43 (Interactive Video) to Q2 T1
mls qos srr-queue output dscp-map queue 2 threshold 2 24
! Maps DSCP CS3 (Call-Signaling) to Q2 T2
mls qos srr-queue output dscp-map queue 2 threshold 3 48 56
! Maps DSCP CS6 and CS7 (Network/Internetwork) to Q2 T3
mls qos srr-queue output dscp-map queue 3 threshold 3 0
! DSCP 0 (Best Effort) gets all of Queue 3
mls qos srr-queue output dscp-map queue 4 threshold 1 8
! Maps DSCP CS1 (Scavenger) to Q4 T1
mls qos srr-queue output dscp-map queue 4 threshold 3 10 12 14
! Maps DSCP AF11, AF12, AF13 (Bulk Data) to Q4 T3
!
mls qos queue-set output 1 threshold 2 70 80 100 100
! Sets Q2 Threshold 1 to 70% and Q2 Threshold 2 to 80%
mls qos queue-set output 1 threshold 4 40 100 100 100
! Sets Q4 Threshold 1 to 40% and Q4 Threshold 2 to 100%
```

Finally, the QoS policy is applied to the access ports.

Map the Policies onto Ports on the Catalyst 3750

```
int range fa1/0/1-24
description ACCESS-EDGE IP PHONE + TRUSTED PC MODEL
switchport mode access
switchport voice vlan 110          ! VVLAN
service-policy input IPPHONE+CSA-PC
srr-queue bandwidth share 1 57 36 7
! Q1 is PQ; Q2 gets 57% of remaining BW; Q3 gets 36% and Q4 gets 7%
srr-queue bandwidth shape 30 0 0 0 ! PQ is BW limited to 30%
priority-queue out                  ! Q1 is enabled as PQ
mls qos trust device cisco-phone    ! Conditionally trust phone
dot1x port-control auto
dot1x host-mode multi-host
dot1x timeout reauth-period server
dot1x reauthentication
no mdix auto
spanning-tree portfast
```

Appendix A. Quality of Service Overview

The following chapter gives a high level overview of QoS and the tools required to implement it. This chapter is intended to provide background on how the Trusted Endpoint QoS Marking feature can be incorporated into an existing or planned campus QoS deployment. Example IOS QoS configurations appear in the first chapter, entitled “Implementation of Trusted Endpoint QoS Marking.”

What is QoS?

QoS is a measure of the service availability and transmission quality of a network.

Service availability is a crucial foundation element of QoS. Uptime and bandwidth are both aspects of availability. The network infrastructure must be designed to be highly available before you can reliably implement QoS. The target for High Availability is 99.999% uptime, with only five minutes of downtime permitted per year. Bandwidth is a measure of data rate but can also be viewed as the capacity of an interface or queue. Congestion occurs whenever the offered traffic load exceeds capacity.

The transmission quality of the network is determined by the following factors:

- **Loss**—A relative measure of the number of packets that were not received compared to the total number of packets transmitted. Loss is a function of availability. If the network is Highly Available, then loss during periods of non-congestion would essentially be zero. During periods of congestion, however, QoS policies and mechanisms can be used to determine which packets should be selectively dropped to alleviate the congestion.
- **Delay**—The finite amount of time it takes a packet to reach the receiving host after being transmitted from the sending host. In the case of voice, this is the amount of time it takes for a sound to travel from the speaker’s mouth to a listener’s ear. Delay is the sum of transit delay, the time it takes for a packet to travel end-to-end in the best case, and queuing delay, the time a packet spends in network device queues during periods of congestion.
- **Delay variation (Jitter)**—The difference in the end-to-end delay between packets. For example, if one packet requires 100 ms to traverse the network from the source host to the destination host and the following packet requires 125 ms to make the same trip, then the delay variation is 25 ms. Significant jitter may be indistinguishable from loss for some applications. For example, it can result in an audible degradation of call quality.

Where is QoS Needed in a Campus?

The case for QoS in WANs and VPNs is largely self-evident because of their low-bandwidth links compared to the high-bandwidth requirements of most applications. However, the need for QoS is sometimes overlooked or even challenged in high-bandwidth Gigabit/TenGigabit campus LAN environments.

Although network administrators sometimes equate QoS only with queuing, the QoS toolset extends considerably beyond just queuing tools. Classification, marking and policing are all important QoS functions that are optimally performed within the campus network, particularly at the access layer ingress edge (access edge).

Most campus links are underutilized. Some studies have shown that 95 percent of campus access layer links are utilized at less than 5 percent of their capacity. This means that you can design campus networks to accommodate oversubscription between access, distribution and core layers. Oversubscription allows for uplinks to be utilized more efficiently and more importantly, reduces the overall cost of building the campus network.

Typical campus oversubscription values are 20:1 for the access-to-distribution layers and 4:1 for the distribution-to-core layers, as shown in Figure A-1.




Figure A-1: Typical Campus Oversubscription Rates

It is quite rare under normal operating conditions for campus networks to suffer congestion. And if congestion does occur, it is usually momentary and not sustained as it might be at a WAN edge. However, critical applications like VoIP still require service guarantees regardless of network conditions.

The only way to provide service guarantees is to enable queuing at any node that has the potential for congestion—regardless of how rarely, in fact, this may occur. The potential for congestion exists in campus uplinks because of oversubscription ratios and speed mismatches in campus downlinks (for example, GigabitEthernet to FastEthernet links). The only way to provision service guarantees in these cases is to enable queuing at these points. Queueing enables prioritized network access for different service classes.

Queueing helps to meet network requirements under normal operating conditions, but enabling QoS within the campus is even more critical under abnormal network conditions such as DoS/worm attacks. During such conditions, network traffic may increase exponentially until links are fully utilized. Without QoS, the worm-generated traffic drowns out applications and causes denial of service through unavailability. Enabling QoS policies within the campus maintains network availability by protecting and servicing critical applications such as VoIP and even Best Effort traffic.

The intrinsic interdependencies of network QoS, High Availability and security are clearly evident in such worse-case scenarios.

So where is QoS required in campus?

Distribution and core switches require the following QoS policies:

- DSCP trust policies
- Optional per-user microflow policing policies (only on supported platforms)
- Queueing policies

Access switches require the following QoS policies:

- Appropriate (host-dependant) trust policies
- Classification and marking policies
- Policing and markdown policies
- Queueing and dropping policies

The Trusted Endpoint QoS marking solution described in this document integrates QoS and security policies at the access edge.

Access Edge Trust Models

The primary function of access edge policies is to establish and enforce trust boundaries. A trust boundary is the point within the network where QoS markings such as DSCP are first accepted. Previously set markings may be re-marked as required at the trust boundary.

You should enforce trust boundaries as close to the hosts as technically and administratively possible as shown in Figure A-2. The Trusted Endpoint QoS marking solution defines the trust boundary *at* the host.




Figure A-2: Establishing Trust Boundaries

The definition of the trust boundary depends on the capabilities of the hosts that are being connected to the access edge of the LAN. The following are the three main categories of hosts as they relate to trust boundaries:

- Trusted Hosts
- Untrusted hosts
- Conditionally-trusted hosts

Trusted Hosts

Trusted hosts have the capabilities and intelligence to mark application traffic to the appropriate CoS and/or DSCP values. Trusted hosts also have the ability to remark traffic that may have been previously marked by an untrusted device. Trusted hosts are typically static devices, meaning that the switch port into which they are plugged does not usually change. They are often physically secure and under administrative control; an example might be a server in a data center.

When trusted hosts are connected to a switch port, all that is typically required is enabling the following interface command: **mls qos trust dscp**. Optionally, if the traffic rate of the trusted application class is known, the network administrator could apply an access layer policer to protect against out-of-profile rates in case the trusted endpoint is somehow compromised.

Untrusted Hosts

Generally, Cisco recommends *against* trusting end users and their PCs because newer operating systems like Windows XP and Linux make it relatively easy to set CoS or DSCP markings on PC NICs. Such markings may be set deliberately or even inadvertently. In either case, improperly set QoS markings can affect the service levels of multiple users within the enterprise and make troubleshooting a nightmare.

If the host is running CSA 5.0 or later and correctly enforcing the enterprise QoS marking policy then the trust level of the host is increased dramatically. Such hosts can be considered trusted if the network administrator is confident that CSA is properly configured and running. Optionally, NAC2 can be utilized to ensure that trust of the host by the access switch is made conditional on the fact that CSA 5.0 or later is properly configured and running.

Conditionally Trusted Hosts

Thin clients like IP Phones are, in theory, less susceptible to infection or tampering and are therefore considered trusted endpoints. However, since IP Phones can change ports when their users move, a dynamic approach to port configuration is desirable (as it would be for other mobile devices as well.) The solution is to have an intelligent exchange of information between the switch and the devices plugged into their ports. If the switch discovers a device that is trustworthy, then it can extend trust to it dynamically.

Cisco IP Phones use this approach, relying on an exchange of messages between the phone and the switch to place the phone into the voice VLAN. NAC2 performs this function between PCs and the network; it can convey user or host identity as well as host configuration attributes and values referred to as host posture.

What is the Cisco QoS Toolset?

Cisco provides a complete toolset of QoS features and solutions for addressing the diverse needs of voice, video and multiple classes of data applications. Cisco QoS technology lets complex networks control and predictably service a variety of networked applications and traffic types. You can effectively control bandwidth, delay, jitter, and packet loss with these mechanisms. By ensuring the desired results, the QoS features lead to efficient, predictable services for business-critical applications. Using the rich Cisco QoS toolset, as shown in Figure A-3, businesses can build networks that conform to the Differentiated Services (DiffServ) architecture, as defined in RFC 2475.




Figure A-3: The Cisco QoS Toolset

Classification and Marking Tools

Aside from admission control (which is provided by NAC and discussed in chapter three) the next element of a QoS policy is to classify/identify the traffic that is to be treated differently. Following classification, marking tools can set an attribute in a frame or packet to a specific value. Classification and marking (or remarking) occurs at an administrative trust boundary that scheduling tools later depend on.

Classification tools typically operate at the ingress ports of Cisco network devices and may examine any of the following:

- Layer 2 parameters—802.1Q Class of Service (CoS) bits, Multiprotocol Label Switching Experimental Values (MPLS EXP)
- Layer 3 parameters—IP Precedence (IPP), Differentiated Services Code Points (DSCP), IP Explicit Congestion Notification (ECN), source/destination IP address
- Layer 4 parameters— L4 protocol (TCP/UDP), source/destination ports
- Layer 7 parameters— application signatures via Network Based Application Recognition (NBAR)

You can only apply QoS policies to traffic after it has been positively classified. To avoid the need for repetitive and detailed classification at every node, packets can be marked according to their desired service levels. Marking tools can be used to indicate respective priority levels by setting attributes in the frame or packet headers so that detailed classification does not have to be recursively performed at each hop. Within an enterprise, marking is done at either Layer 2 or Layer 3.

Layer 2 media may change as packets traverse from source to destination, so a more universal (end-to-end) and granular classification occurs at Layer 3. **For the Trusted Endpoint QoS solution described in this document, we are concerned only with DSCP values transmitted in the IP Type of Service (ToS) byte as shown in Figure A-4.** The IP Precedence (IPP) bits formerly defined in the ToS byte have been replaced by the DiffServ architecture, but backward compatible DSCP values are provided as described below.

DSCP values can be expressed either in numeric form or by standards-based keywords corresponding to Per-Hop Behaviors (PHBs.) For example, DSCP 46 is synonymous with DSCP EF and DSCP AF31 is synonymous with DSCP 26 (see Table A-1.) A Per-Hop Behavior refers to the differentiated level of QoS service that a marked packet *should* receive at each network device it passes through, as defined by the applicable RFC.




Figure A-4: The IP ToS Byte

There are four broad classes of DSCP PHB markings: Best Effort (BE or DSCP 0), RFC 2474 Class Selectors (CS1–CS7, which are identical/backward-compatible to IPP values 1–7), RFC 3268 Expedited Forwarding (EF), and RFC 2597 Assured Forwarding PHBs (AF_{xy}).

There are four Assured Forwarding classes, each of which begins with the letters “AF” followed by two digits. The first digit corresponds to the DiffServ Class of the AF group and can range from 1 through 4. The second digit refers to the level of Drop Preference within each AF class and can range from 1 (lowest Drop Preference) through 3 (highest Drop Preference).

Policing and Markdown Tools

Policing tools (policers) determine whether packets are conforming to administratively-defined traffic rates and take action accordingly. Such action could include passing, remarking or dropping a packet.

A basic policer monitors a single rate: traffic equal to or below the defined rate is considered to *conform* to the rate, while traffic above the defined rate is considered to *exceed* the rate. On the other hand, the algorithm of a dual-rate policer (such as described in RFC 2698) is analogous to a traffic light. Traffic equal to or below the principal defined rate (green light) is considered to *conform* to the rate. An allowance for moderate amounts of traffic above this principal rate is permitted (yellow light) and such traffic is considered to *exceed* the rate. However, a clearly-defined upper-limit of tolerance is set (red light), beyond which traffic is considered to *violate* the rate.

Policers complement classification and marking policies. For example, as previously discussed, RFC 2597 defines the AF classes of PHBs. Traffic conforming to the defined rate of a given AF class is marked to the first Drop Preference level of a given AF class (for example, AF21). Traffic exceeding this rate is marked down to the second Drop Preference level (for example, AF22) and violating traffic is either marked down further to the third Drop Preference level (for example, AF23) or simply dropped.

Scheduling Tools

Scheduling tools determine how a frame/packet exits a device. Whenever packets enter a device faster than they can exit it, such as with data link rate mismatches, then a point of congestion, or bottleneck, can occur. Devices have buffers that allow for scheduling higher-priority packets to exit sooner than lower priority ones, which is commonly called queueing.

Queueing algorithms are activated only when a device is experiencing congestion and are deactivated when the congestion clears. The main Cisco IOS software queueing tools are Low Latency Queueing (LLQ), which provides strict priority servicing and is intended for realtime applications such as VoIP; and Class-Based Weighted Fair Queueing (CBWFQ), which provides bandwidth guarantees to given classes of traffic and fairness to discrete traffic flows within these traffic classes.

How is QoS Optimally Deployed Within the Enterprise?

A successful QoS deployment is comprised of multiple phases, including:

- 1) Strategically defining the business objectives to be achieved via QoS.
- 2) Analyzing the service-level requirements of the various traffic classes to be provisioned for.
- 3) Designing and testing QoS policies prior to production-network rollout.
- 4) Rolling out the tested QoS designs to the production network.
- 5) Monitoring service levels to ensure that the QoS objectives are being met.

These phases may need to be repeated as business conditions change and evolve.

An in-depth treatment of each of these phases may be found in the Cisco QoS references cited earlier. The following sections provide a high level overview, in particular as they relate to deploying Trusted Endpoint QoS Marking.

1) Strategically Defining QoS Objectives

QoS technologies are the enablers for business/organizational objectives. Therefore, the way to plan a QoS deployment is by clearly defining the objectives of the organization. For example, among the first questions that arise during a QoS deployment are: How many traffic classes should be provisioned for? And what should they be?

To help answer these fundamental questions, organizational objectives need to be defined, such as:

- Is the objective to enable VoIP only or is video also required?
- If so, is video-conferencing required or streaming video? Or both?

- Are there applications that are considered mission-critical, and if so, what are they?
- Does the organization wish to demote certain types of traffic, and if so, what are they?

To help address these critical questions and to simplify QoS, Cisco has adopted a new initiative called the *QoS Baseline*. The QoS Baseline is a strategic document designed to unify QoS within Cisco. It provides uniform, standards-based recommendations to help ensure that QoS products, designs, and deployments are unified and consistent. The QoS Baseline defines up to 11 classes of traffic that may be viewed as critical to a given enterprise. A summary of these classes and their respective standards-based markings and recommended QoS configurations are shown in Table A-1.

- The **IP Routing** class is intended for IP Routing protocols such as Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), etc.
- **Voice** refers to VoIP bearer traffic only (and does not include Call-Signaling traffic.)
- **Interactive-Video** refers to IP Video-Conferencing.
- **Streaming Video** is either unicast or multicast unidirectional video.
- The (Locally-Defined) **Mission-Critical** class is intended for a subset of Transactional Data applications that contribute most significantly to the business objectives (this is a non-technical assessment).
- The **Call-Signaling** class is intended for voice and/or video signaling traffic, such as Skinny, SIP, H.323, etc.
- The **Transactional Data** class is intended for foreground, user-interactive applications such as database access, transaction services, interactive messaging, and preferred data services.
- The **Network Management** class is intended for network management protocols, such as SNMP, Syslog, DNS, etc.
- The **Bulk Data** class is intended for background, non-interactive traffic flows, such as large file transfers, content distribution, database synchronization, backup operations, and email.
- The **Scavenger** class is based on an Internet 2 draft that defines a “less-than-Best Effort” level of service. In the event of link congestion, this class will be dropped most aggressively (which is vitally important in mitigating DoS and work attacks.)
- The **Best Effort** is the default service level provided to traffic by the network in the absence of policies to mark packets to a higher or lower service level.

Application	Layer 3 Classification		Referencing Standard	Recommended Configuration
	PHB	DSCP		
IP Routing	CS6	48	RFC 2474	Rate-Based Queueing + RED
Voice	EF	46	RFC 3246	RSVP Admission Control + Priority Queueing
Interactive Video	AF41	34	RFC 2597	RSVP + Rate-Based Queueing + DSCP-WRED
Streaming-Video	CS4	32	RFC 2474	RSVP + Rate-Based Queueing + RED
Mission-Critical Data	AF31	26	RFC 2597	Rate-Based Queueing + DSCP-WRED
Call-Signaling	CS3	24	RFC 2474	Rate-Based Queueing + RED
Transactional Data	AF21	18	RFC 2597	Rate-Based Queueing + DSCP-WRED
Network Management	CS2	16	RFC 2474	Rate-Based Queueing + RED
Bulk Data	AF11	10	RFC 2597	Rate-Based Queueing + DSCP-WRED
Scavenger	CS1	8	Internet 2	No BW Guarantee + RED
Best Effort	BE	0	RFC 2474	BW Guarantee Rate-Based Queueing + RED

Table A-1: Cisco QoS Baseline Summary

Adopting standards-based marking recommendations helps position the enterprise for better integration with service-provider offerings as well as other internetworking scenarios. The QoS Baseline recommendations are intended as a standards-based guideline for customers-not as a mandate. Customers do not have to deploy all 11 traffic classes, but may start with simple QoS models and expand over time as business needs arise, as shown in Figure A-5.




Figure A-5: Example Strategy for Expanding the Number of Classes of Service over Time

2) Analyzing Application Service-Level Requirements

This topic is covered in depth in the Cisco QoS references cited previously. A high level summary is presented here. This task essentially consists of identifying the mission-critical applications and characterizing their traffic to establish a “normal” baseline for offered traffic.

Voice requires 150 ms one-way, end-to-end (mouth-to-ear) delay, 30 ms of one-way jitter and no more than 1% packet loss. Voice should receive strict priority servicing, and the amount of priority bandwidth assigned for it should take into account the VoIP codec, the packetization rate, IP/UDP/RTP headers (compressed or not) and Layer 2 overhead. Additionally, provisioning QoS for IP telephony requires that a minimal amount of guaranteed bandwidth be allocated to Call-Signaling traffic.

Video comes in two flavors: Interactive Video and Streaming Video. Interactive Video has the same service level requirements as VoIP because a voice call is embedded within the video stream. Streaming Video has much laxer requirements, because of the high amount of buffering that has been built into the applications. Control plane requirements, such as provisioning moderate bandwidth guarantees for IP Routing and Network Management protocols, should not be overlooked.

Data comes in a variety of forms, but can generally be classified into four main classes: Best Effort (the default class), Bulk (non-interactive, background flows), Transactional/Interactive (interactive, foreground flows) and Mission-Critical. Mission-Critical Data applications are locally-defined, meaning that each organization must determine the select few Transactional Data applications that contribute the most significantly to their overall business objectives.

3) Designing the QoS Policies

Once a QoS strategy has been defined and the application requirements are understood, end-to-end QoS policies can be designed for each device and interface, as determined by its role in the network infrastructure. A separate QoS design document delves into the specific details of LAN, WAN, and VPN (both MPLS and IPsec VPN) QoS designs. Because the Cisco QoS toolset provides many QoS design and deployment options, a few succinct design principles can help simplify strategic QoS deployments.

For example, one such design principle is to ***always enable QoS policies in hardware—rather than software—whenever a choice exists***. Cisco IOS routers perform QoS in software, which places incremental loads on the CPU, depending on the complexity and functionality of the policy. Cisco Catalyst switches, on the other hand, perform QoS in dedicated hardware ASICs and as such do not tax their main CPUs to administer QoS policies. This allows complex policies to be applied at line rates at even Gigabit or Ten-Gigabit speeds.

Other simplifying best-practice QoS design principles include:

- Classification and Marking Principles
- Policing and Markdown Principles
- Queueing and Dropping Principles

Also, the following topics should be considered as central to the design:

Classification and Marking Design Principles

When classifying and marking traffic, an unofficial Differentiated Services design principle is to ***classify and mark applications as close to their sources as technically and administratively feasible***. This principle promotes end-to-end Differentiated Services and PHBs. Do not trust markings that can be set by users on their PCs or other similar devices, because users can easily abuse provisioned QoS policies if permitted to mark their own traffic. For example, if DSCP EF received priority services throughout the enterprise, a PC can be easily configured to mark all the traffic of the user to DSCP EF, thus hijacking network priority queues to service non-realtime traffic. Such abuse could easily ruin the service quality of realtime applications like VoIP throughout the enterprise.

Following this rule, it is further recommended to ***use DSCP markings whenever possible***, because these are end-to-end, more granular and more extensible than Layer 2 markings. Layer 2 markings are lost when media changes (such as a LAN-to-WAN/VPN edge). There is also less marking granularity at Layer 2. For example, 802.1Q/p CoS supports only 3 bits (values 0–7), as does MPLS EXP. Therefore, only up to 8 classes of traffic can be supported at Layer 2, and inter-class relative priority (such as RFC 2597 Assured Forwarding Drop Preference markdown) is not supported. On the other hand, Layer 3 DSCP markings allow for up to 64 classes of traffic, which is more than enough for most enterprise requirements for the foreseeable future.

As the line between enterprises and service providers continues to blur and the need for interoperability and complementary QoS markings is critical, you should ***follow standards-based DSCP PHB markings to ensure interoperability and future expansion***. Because the QoS Baseline marking recommendations are standards-based, enterprises can easily adopt these markings to interface with service provider classes of service. Network mergers—whether the result of acquisitions, mergers or strategic-alliances—are also easier to manage when you use standards-based DSCP markings.

Policing and Markdown Design Principles

There is little reason to forward unwanted traffic only to police and drop it at a subsequent node, especially when the unwanted traffic is the result of DoS or worm attacks. The overwhelming volume of traffic that such attacks can create can cause network outages by driving network device processors to their maximum levels. Therefore, you should ***police traffic flows as close to their sources as possible***. This principle applies also to legitimate flows. DoS/worm-generated traffic can masquerade under legitimate, well-known TCP/UDP ports and cause extreme amounts of traffic to be poured onto the network infrastructure. Such excesses should be monitored at the source and marked down appropriately.

Whenever supported, markdown should be done according to standards-based rules, such as RFC 2597 (“Assured Forwarding PHB Group”). For example, excess traffic marked to AFx1 should be marked down to AFx2 (or AFx3, whenever dual-rate policing—such as defined in RFC 2698—is supported). Following such markdowns, congestion management policies, such as DSCP-based WRED, should be configured to drop AFx3 more aggressively than AFx2, which in turn should be dropped more aggressively than AFx1.

However, Cisco Catalyst switches do not currently perform DSCP-Based WRED, and so this standards-based strategy cannot be implemented fully at this time. As an alternative workaround, single-rate policers can be configured to markdown excess traffic to DSCP CS1 (Scavenger); dual-rate policers can be configured to markdown excess traffic to AFx2, while marking down violating traffic to DSCP CS1. Traffic marked as Scavenger would then be assigned to a “less-than-Best-Effort” queue. Such workarounds yield an overall effect similar to the standards-based policing model. However, when DSCP-based WRED is supported on all routing and switching platforms, then you should markdown Assured Forwarding classes by RFC 2597 rules to comply more closely with this standard.

Queueing and Dropping Design Principles

Critical applications such as VoIP require service guarantees regardless of network conditions. **The only way to provide service guarantees is to enable queuing at any node that has the potential for congestion**, regardless of how rarely this may occur. This principle applies not only to Campus-to-WAN/VPN edges, where speed mismatches are most pronounced, but also to Campus Access-to-Distribution or Distribution-to-Core links, where oversubscription ratios create the potential for congestion. There is simply no other way to guarantee service levels than by enabling queuing wherever a speed mismatch exists.

When provisioning queuing, some best practice rules of thumb also apply. For example, as discussed previously, the Best Effort class is the default class for all data traffic. Only if an application has been selected for preferential/deferential treatment is it removed from the default class. Because many enterprises have several hundred, if not thousands, of data applications running over their networks, you must provision adequate bandwidth for this class as a whole to handle the sheer volume of applications that default to it. Therefore, it is recommended that you **reserve at least 25 percent of link bandwidth for the default Best Effort class**.

Not only does the Best Effort class of traffic require special bandwidth provisioning consideration, so does the highest class of traffic, sometimes referred to as the “Realtime” or “Strict Priority” class (which corresponds to RFC 3246 “An Expedited Forwarding Per-Hop Behavior”). The amount of bandwidth assigned to the Realtime queuing class is variable. However, if you assign too much traffic for strict priority queuing, then the overall effect is a dampening of QoS functionality for non-realtime applications. Remember: the goal of convergence is to enable voice, video, and data to transparently co-exist on a single network. When Realtime applications such as Voice or Interactive-Video dominate a link (especially a WAN/VPN link), then data applications will fluctuate significantly in their response times, destroying the transparency of the converged network.

Cisco Technical Marketing testing has shown a significant decrease in data application response times when realtime traffic exceeds one-third of link bandwidth capacity. Extensive testing and customer deployments have shown that a general best queuing practice is to **limit the amount of strict priority queuing to 33 percent of link capacity**. This strict priority queuing rule is a conservative and safe design ratio for merging realtime applications with data applications.

Cisco IOS software allows the abstraction (and thus configuration) of multiple strict priority LLQs. In such a multiple LLQ context, this design principle would *apply to the sum of all LLQs to be within one-third of link capacity*.

Whenever a Scavenger queuing class is enabled, it should be assigned a minimal amount of bandwidth. On some platforms, queuing distinctions between Bulk Data and Scavenger traffic flows cannot be made because queuing assignments are determined by CoS values and these applications share the same CoS value of 1. In such cases you can assign the Scavenger/Bulk queuing class a bandwidth percentage of 5. If you can uniquely assign Scavenger and Bulk Data to different queues, then you should assign the Scavenger queue a bandwidth percentage of 1.

The Realtime, Best Effort and Scavenger queuing best practice principles are shown in Figure A-6.




Figure A-6: Queueing Rules for 4-Class Service Model

Because platforms support a variety of queuing structures, configure consistent queuing policies according to platform capabilities to ensure consistent PHBs.

For example, on a platform that only supports four queues with CoS-based admission (such as a Catalyst switch) a basic queuing policy could be as follows:

- Realtime (≤ 33%)
- Critical Data
- Best Effort Data (≥ 25%)
- Scavenger/Bulk (≤ 5%)

4) Rolling Out the QoS Policies

Once the QoS designs have been finalized and PoC tested, it is vital to ensure that the networking team **thoroughly understands the QoS features and syntax before enabling features on production networks**. Such knowledge is critical for both rollout and subsequent troubleshooting of QoS-related issues.

Furthermore, it is recommended to **schedule network downtime in order to rollout QoS** features. While QoS is required end-to-end, it does not have to be deployed end-to-end at a single instance. A pilot network-segment can be selected for an initial deployment, and pending observation, the **rollout can be expanded in stages** to encompass the entire enterprise.

A rollback strategy is always recommended, to address unexpected issues arising from the QoS deployment.

5) Monitoring the Service Levels

Implementing a QoS solution is not a one-time task that is complete upon policy deployment. A successful QoS policy **rollout is followed by ongoing monitoring of service levels and periodic adjustments and tuning** of QoS policies.

Short-term monitoring is useful for verifying that the deployed QoS policies are having the desired end-to-end effect. *Long-term monitoring* (trending) is needed to determine whether the provisioned bandwidth is still adequate for the changing needs of the enterprise. As business conditions change, the enterprise may need to adapt to these changes and **may be required to begin the QoS deployment cycle anew**, by redefining their objectives, tuning and testing corresponding designs, rolling these new designs out and monitoring them to see if they match the redefined objectives.

How Can QoS Tools Be Used to Mitigate DoS/Worm Attacks?

A *reactive approach* to mitigating DoS/worm flooding attacks within enterprise networks is to reverse-engineer the worm and set up intrusion detection mechanisms and/or ACLs and/or NBAR policies to limit its propagation. However, the increased sophistication and complexity of worms make them harder and harder to separate from legitimate traffic flows. This exacerbates the finite time lag between when a worm begins to propagate and when analysis of the worm has been completed and a patch or ACL is distributed.

A *proactive approach* to mitigating such attacks is to immediately respond to out-of-profile network behavior indicative of a DoS or worm attack through access layer policers. Such policers meter traffic rates received from host devices and mark down excess traffic when these exceed specified thresholds (at which point they are no longer considered normal flows).

These policers are relatively “dumb” because they do not match specific network characteristics of specific types of attacks. Instead, they simply meter traffic volumes and respond to abnormally high volumes as close to the source as possible. The simplicity of this approach negates the need for the policers to be programmed with knowledge of the specific details of *how* the attack is being generated or propagated.

It is precisely this “dumbness” of such access layer policers that allow them to stay effective as worms mutate and become more complex. The policers do not care *how* the traffic was generated or *what* it looks like; they only care *how much* traffic is being put onto the wire. Therefore, they continue to police even advanced worms that continually change the tactics of how traffic is being generated.

The recommended method is to profile applications and establish normal thresholds for all implemented service classes. For example, in most enterprises it is quite abnormal (within a 95 percent statistical confidence interval) for PCs to generate sustained Best Effort traffic in excess of 5 percent of their link capacity. In the case of a FastEthernet switch port, this means that it is unusual in most organizations for an end-user PC to generate more than 5 Mbps of uplink traffic on a sustained basis.

These thresholds are coupled with access layer policers with hardware/software (campus/WAN/VPN) queuing policies. With this method, access layer policers mark down excess traffic to DSCP CS1 (Scavenger) and all congestion management policies (whether in Catalyst hardware or in IOS software) are provisioned with an end-to-end a less-than-Best-Effort Scavenger service class.

During the course of normal operation, individual hosts may burst above a threshold. Excess traffic will be remarked to Scavenger but, because the network is not experiencing congestion, the traffic will continue on normally to its destination.

In the case of illegitimate excess traffic, the effect of access layer policers on traffic caused by DoS or worm attacks is quite different. As many hosts become infected and traffic volumes multiply, congestion may be experienced in the campus up-links due to the aggregate traffic volume. For example, if just 11 end-user PCs on a single access-layer switch begin spawning worm flows to their maximum FastEthernet link capacities, the GigabitEthernet up-link to the distribution layer switch will congest, and queuing/reordering will engage. At such a point, VoIP and critical data applications, and even Best Effort applications, gain priority over worm-generated traffic. Scavenger traffic is dropped the most aggressively, while network devices remain accessible for the administration of patches/plugs/ACLs required to fully neutralize the specific attack.

WAN links are also protected. VoIP, critical data and even Best Effort flows continue to receive priority over any traffic marked down to Scavenger/CS1. This is a huge advantage, because WAN links are generally the first to be overwhelmed by DoS/worm attacks. Access layer policers thus *significantly mitigate* network traffic generated by DoS or worm attacks.

It is important to recognize the distinction between mitigating an attack and preventing it entirely. The strategy presented here *does not guarantee that no Denial of Service or worm attacks will ever happen, but serves only to reduce the risk and impact* that such attacks have on the campus network infrastructure and then, by extension, the WAN/VPN network infrastructure. Furthermore, while this strategy reduces the collateral damage to the network infrastructure caused by DoS/worm attacks, it may not mitigate other specific objectives of such worms, such as reconnaissance and vulnerability exploitation. Hence, a comprehensive approach must be used to address DoS/worm attacks, involving a holistic integration of security technologies with Quality of Service technologies.

Appendix B. Cisco Security Agent Overview

The following chapter gives a high level overview of Cisco Security Agent (CSA) and how it is deployed. Example CSA QoS configurations appear in the first chapter, entitled “Implementation of Trusted Endpoint QoS Marking.”

What Cisco Security Agent Does

Cisco Security Agents provides intrinsic, distributed security to your enterprise by deploying agents that defend against the proliferation of attacks across networks and systems. These Cisco Security Agents enforce a set of policies provided by Management Center for Cisco Security Agents (CSA MC) and selectively applied to system nodes by the network administrator. Operating under the direction of assigned policies, Cisco Security Agents provide strong system resource protection, tying together the auditing and control of multiple system and network resources.

How Cisco Security Agents Protect Against Attacks

The Cisco Security Agent differs from anti-virus and network firewall software in that it doesn't prevent users from accessing technologies they require. It assumes that users are going to put their systems at risk by making use of a wide range of Internet resources. Keeping this in mind, Cisco Security Agents install and work at the kernel level, controlling network actions, local file systems, and other system components, maintaining an inventory of what actions may be performed on the system itself. This way, malicious system actions are immediately detected and disabled while other actions are permitted. Both actions take place transparently, without any interruption to the user.

If an encrypted piece of malicious code finds its way onto a system via email, for example, as it attempts to unexpectedly execute or alter Cisco Security Agent-protected system resources, it is immediately neutralized and a notification is sent to the network administrator.

Cisco Security Agents use policies which network administrators configure and deploy to protect systems. These policies can allow or deny specific system actions. Cisco Security Agents must determine whether an action is allowed or denied before any system resources are accessed and acted upon.

Specifically, rule policies enable administrators to control access to system resources based on the following parameters:

- Which resource is being accessed.
- Which operation is being invoked.
- Which application is invoking the action.

The resources in question may be either system resources or network resources such as email servers.

When any system actions that are controlled by specific rules are attempted and allowed or denied accordingly, a system event is logged and sent to the administrator in the form of a configurable notification such as email, pager or custom script.

How is CSA Deployed?

Management Center for Cisco Security Agents contains two components:

- CSA MC—installs on designated Windows 2000 systems and includes a configuration database server and a web-based user interface.
- Cisco Security Agent (the agent)—installs on server and desktop systems across your enterprise network.

Using CSA MC, you assemble your network machines into specified groups and then attach security policies to those groups. All configuration is done through the web-based user interface and then deployed to the agents.




Figure B-1: CSA Policy Deployment

The network example shown in Figure B-1 illustrates a basic deployment scenario. CSA MC software is installed on a system which maintains all policy and host groups. The administration user interface is accessed securely using Secure Sockets Layer (SSL) from any machine on the network that can connect to the server and run a web browser. Use the web-based interface to deploy your policies from CSA MC to agents across your network.

Appendix C. Network Admission Control Overview

The following chapter gives an overview of the Cisco Network Admission Control architecture. An example IOS NAC2 configuration appears in the first chapter, entitled “Implementation of Trusted Endpoint QoS Marking.”

NAC Architecture

Network Admission Control assesses the state, or posture, of a host in order to prevent unauthorized or vulnerable endpoints from accessing the network. Typical endpoints are desktop computers, laptops, and servers but may include IP phones, network printers, and other specialized network-attached devices.

The posture validation process encompasses these major architectural components:

- **Host:** Machine accessing network for which NAC is enforced
- **Posture Plugin (PP):** A Cisco or third-party DLL that resides on a host and provides posture credentials to a posture agent residing on the same device.
- **Posture Agent (PA):** Host agent software that serves as a broker on the host for aggregating credentials from potentially multiple posture plugins and communicating with the network. The Cisco Trust Agent (CTA) is Cisco’s implementation of the posture agent. Posture Agent may use an L3 transport (EAPoUDP) or an L2 transport (802.1x).
- **Network Access Device (NAD):** Network devices acting as a NAC enforcement point. These may include Cisco access routers (1700-7200), VPN Gateways (VPN3000 series), PIX firewalls, Catalyst L2 and L3 switches, and wireless access points.
- **Authentication, Authorization and Accounting (AAA) Server:** The central server that aggregates one or more authentication and/or authorization decisions into a single system authorization decision, and maps this decision to a network access profile for enforcement on the NAD. Cisco Secure Access Control Server (ACS) is Cisco’s AAA server product that supports NAC.
- **Posture Validation Server (PVS):** A posture validation server from one or more third parties acts as an application-specific policy decision point in NAC for authorizing a set of posture credentials from one or more posture plugins against a set of policy rules.

Posture validation occurs when a NAC-enabled network access device detects a host attempting to connect or use its network resources. On detection of a new endpoint, the NAD sets up a communication path between the AAA server and the Posture Agent. Once the communication path has been established, the AAA server requests the endpoint for posture credentials from one or more posture plugins. The host responds to the request with sets of posture credentials from the specified posture plugins, which contain the state of the various hardware and software components on the host. The AAA server either validates the posture information locally or it may in turn delegate parts of the decision to external posture validation servers. Ultimately, the AAA server aggregates the posture validation results from one or more sources, which results in an authorization decision, or posture token, representing the host’s relative compliance to the network policy. The AAA server determines the appropriate network access profile for the host, and sends it to the network access device for enforcement of the host authorization.

All posture decision points, whether an AAA server or PVS, evaluate one or more sets of host credentials in rule-based policy engines which results in one or more application posture token (APTs). An APT represents a compliance check for a given vendor’s application on the host. The AAA server then merges all APTs from the delegated PVSes and its own policy engine into a single system posture token (SPT) representing the overall compliance of the host. Both APTs and SPTs are represented using the following pre-defined tokens:

Healthy - Host is compliant; no restrictions on network access.

Checkpoint – Host is within policy but an update is available. Checkup is used to proactively remediate a host to the Healthy state.

Transition – Host posturing is in process; give interim access pending full posture validation. This state is applicable either during host boot when all NAC-enabled applications may not be running or during an audit when posture information has not yet been obtained from the host.

Quarantine – Host is out of compliance; restrict network access to a quarantine network for remediation. The host is not an active threat but is vulnerable to a known attack or infection

Infected – Host is an active threat to other hosts; network access should be severely restricted or totally denied all network access.

Unknown - Host posture cannot be determined. Quarantine the host and audit or remediate until a definitive posture can be determined.

Identity authentication and posture validation occurs when a host requests access to a network. Through a Layer 2 or Layer 3 transport, a Cisco NAD retrieves Posture Credentials from the client device. How the host is admitted into the network is then based on the level of compliance with existing network policy rules. These posture credentials are typically based on the state of the device Operating System as well as applications such as Anti-Virus, IDS and Firewall. As an example, this helps customers implement security policies such as “Restricted access unless the CSA version is at least 5.0, CSA is operational, it is associated with the correct CSA MC, and the agent has polled the CSA MC within the last 24 hours.”

Figure C-1 shows the components of the NAC architecture and their relationship. The message flows used to gather host information and apply policy decisions are also shown.




Figure C-1: The Cisco NAC Architecture

1. A host attempts to access a NAC-compliant network. When the Network Access Devices (NAD) detects a host attempting to access the network it requests the user's identity. The NAD passes the user's identity information to the Cisco Secure Access Control Server (ACS) using a RADIUS session.

2. ACS requests posture credentials from the host. Posture plugins installed on the host gather posture credentials, pass them to CTA, and CTA forwards them to the NAD.
 - If the NAD is a router, the posture information is sent directly from CTA to the router using Extensible Authentication Protocol over Universal Datagram Protocol (EAPoUDP).
 - If the NAD is a switch using the NAC-L2-IP protocol, the posture information is sent directly from CTA to the switch using EAPoUDP.
 - If the NAD is a switch, using the IEEE 802.1x protocol, the posture information is sent to the switch by the Cisco Trust Agent 802.1x Wired Client, also known as the “supplicant.” (See the Cisco Trust Agent Administrator’s Guide for more information about the supplicant.)
3. The host’s credentials are forwarded from the NAD to the ACS.
4. (Optional NAC Configuration) Third party authentication servers may determine the posture of some of the applications which reported their credentials. There may be several different authentication servers which validate the posture information. Once the authentication servers determine their application postures, they return them to ACS.
5. ACS determines application postures for all those applications which do not require a third party authentication server to do so. ACS also aggregates all application posture tokens to define an overall system posture token for the host. The system posture token equals the least desirable posture of all the application posture tokens defined for the host.
6. Cisco Secure ACS maps the system posture token to a network access policy and, optionally, a user notification.
7. Cisco Secure ACS sends the security policy for the host to the NAD. The NAD enforces the policy for the host. Depending on how the policy is written, the host may be allowed to access the network it requested, it may be sent to a quarantine network where all it can access is a software update server, or it may be denied access to any network.
8. The host is notified of all of its application posture tokens and system posture token. The host may also pop-up a message informing the user of its overall posture depending on how ACS and CTA are configured.
9. The host’s posture is now defined.

How is NAC Deployed?

The Cisco Trust Agent is CSA’s Posture Agent and is installed automatically when CSA is installed. If admission control is to be based on user or machine identity, then the Cisco Trust Agent 802.1x Wired Client or other 802.1x supplicant must also be installed.

The fundamental difference between NAC deployment models has to do with the method used to exchange information between the host and the NAD and whether this occurs in one exchange or two. The process for deciding which model to deploy can be complex and is discussed in the Cisco NAC references cited earlier. For the purposes of this document it is sufficient to understand that NAC-L2-802.1x is supported on all Cisco switch platforms and the network access profile is enforced via dynamic VLAN assignment. For the NAC-L2-IP the network access profile is enforced using ACLs. The NAC-L2-802.1x model enables a more granular QoS policy at this time. It also places a reduced transaction load on the ACS by combining authentication and authorization in one exchange and is therefore more scalable.

The NAC deployment models are compared in Table C-1. The term “IBNS” stands for Cisco Identity-Based Network Services (802.1x). The NAC-L3-IP method refers to NAC as implemented on router NADs and is therefore not relevant to this solution, which is focused on the access edge. The term “NAH” refers



to Non-Authenticated (non-responsive) Hosts. These hosts can be audited by NAC but cannot participate in identity or posture exchange.

Deployment Model	Pros	Cons
Identity and Posture	Unified identity & posture with NAC-L2-802.1x L2 enforcement IBNS-compatible	Not supported with NAC-L2-IP and NAC-L3-IP Retail supplicant license for wireless support No audit support (Future)
IEEE 802.1x	IBNS-compatible	No posture No audit support
Posture Only	NAC-L2-IP and NAC-L3-IP NAH Audit support (L3-IP in Future) Supplicant optional	No identity
IEEE 802.1x and Posture	IBNS-compatible Posture Audit support (L3-IP in Future)	Disjointed Authorization (posture after VLAN assignment) Twice the load on the ACS server Multiple clients / management complexity

Table C-1: NAC Deployment Comparison