

Scenario-Based TrustSec Deployments Application Note

Last Updated: September 6, 2011



Building Architectures to Solve Business Problems

About Cisco Validated Design (CVD) Program

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

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

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

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

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

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

Scenario-Based TrustSec Deployments Application Note

© 2011 Cisco Systems, Inc. All rights reserved.



Scenario-Based TrustSec Deployments

Cisco Trusted Security (TrustSec) is an integrated system including several Cisco products that offer authentication and identity-based access control to secure network connectivity and resources. To simplify the process of deploying TrustSec, this document provides a high-level description of a phased, scenario-based deployment strategy that can be used to roll out TrustSec with minimal impact to end users. This application note includes the following sections:

- Deployment Scenarios, page 3
- Planning a Deployment, page 4
- Implementation Details, page 5
- Monitor Mode, page 12
- Low Impact Mode, page 15
- High Security Mode, page 19
- References, page 23

This document first discusses important steps in the planning process and then explains the implementation details that form the basis for the scenarios that are described later in the document.

Deployment Scenarios

.......

CISCO

With Cisco TrustSec, you can facilitate greater security and enjoy cost-effective management of changes throughout your organization. Having a secure TrustSec framework in place helps enterprises better manage employee mobility, reduce network access expenses, and boost overall productivity while lowering operating costs.

TrustSec deployments are often most successful when they are implemented in phases, gradually adding network access restrictions to minimize the impact for end users. Three scenarios are described in this document, as summarized in Table 1:

- Monitor mode
- Low impact mode
- High security mode



Many customers have successfully deployed 802.1X using these deployment scenarios. By starting your deployment in monitor mode and adding access control in a phased transition to scenario, you can deploy 802.1X with minimal impact to end users.

Deployment Scenario	Best for	Auth Types	Host Mode	Pre-Auth	Successfu I Auth	Failed Auth
Monitor mode	All customers (initial deployment)	802.1X and MAB	Multi-auth	Open	Open	Open
Low impact mode	Customers seeking simple access control with minimal impact to end users and network infrastructure	802.1X and MAB	Single-auth (non-IPT) multi-domain (IPT)	Selectively Open	Dynamic ACL	Selectively open
High security mode	Customers seeking the security of traditional 802.1X with L2 traffic isolation and/or network virtualization	802.1X and MAB	Single-auth (non-IPT) Multi-domain (IPT)	Closed	Dynamic VLAN	Closed (or Auth-Fail VLAN)

Table 1 Deployment Scenario Summary

Each scenario leverages specific combinations of features and configurations to satisfy a particular set of use cases. By starting with these scenarios, you can follow a well-defined and well-understood blueprint for implementing TrustSec. Instead of starting from the very beginning, you can follow the guidelines for a particular deployment scenario and then, if necessary, customize it to suit your network requirements.

Planning a Deployment

In its simplest formulation, TrustSec enables you to grant customized access to the network based on the identity of a person or device. Therefore, when planning a TrustSec deployment, the first and most fundamental question that must be answered is deciding who gets what kind of access.

To answer this, start with your corporate security policy. A well-defined security policy should provide broad guidelines for who can access which corporate resources under what conditions. Next, determine the categories of users and devices that connect to your network.

<u>}</u> Tip

Best Practice Recommendation—Keep it Simple. For initial deployments, use broad categories for users and devices, such as Employee, Managed Asset, and Unknown Asset. The simpler your policy, the smoother and more successful your initial deployment will be.

After basic access control has been successfully deployed, you can easily add more granular groups and policies if you need more highly differentiated access control.

After you have determined the high-level categories of users and devices on your network, use the guidelines from your security policy to map each category to a network access level. The example shown in Table 2, although simple, has been used as the basis for many successful deployments.

Category	Network Access Level
Employee	Full access (intranet and Internet)
Managed asset	Full access (intranet and Internet)
Unknown device	Connectivity services (DHCP, DNS) only
Pre-authentication	Connectivity services only
Failed authentication	Connectivity services only

Table 2	Network Access	evels

The last two categories, *pre-authentication* and *failed authentication*, deserve a bit more explanation. Pre-authentication refers to the level of access that a user or device gets before its identity has been determined. Under a very strict security policy, this level could be *no access*. Alternatively, it could be limited to a small set of services, such as DHCP, DNS, and TFTP. The latter policy allows a device to download an operating system or perform web authentication before its identity has been established.

Failed authentication refers to the level of access that a user or device gets if it fails to provide valid credentials. Again, under a very strict security policy, this could be *no access*, which is the default. Be aware, however, that with this policy, you must provide a manual process by which legitimate users can update their credentials, such as when renewing an expired password. As an example of a manual process, employees would have to take their laptops to a physically secure location where an IT administrator can update the credential. Such a process can be resource intensive, both in terms of end user education and IT support. Therefore, if your security policy permits, you might want to consider modifying the default failed authentication policy to permit a small set of services that would allow a device to automatically update or request credentials.

After you have identified the categories of users and devices on the network and determined the appropriate level of network access for each category, it is time to move on to some lower-level questions and issues that determine how you implement the defined policy.

Implementation Details

The following sections the questions you need answer and describe the configuration elements that are required to implement a solution that best suits your network:

- Implementation Overview, page 6
- Who is Connecting?, page 6
- What Level of Authorization?, page 7
- How are Hosts Connected?, page 10

Implementation Overview

To implement your network access policy, you must define the following:

- 1. Who is connecting to the network?
 - **a.** What methods will be used to identify them as they come onto the network?
 - **b.** What kinds of credentials will be accepted as valid forms of identification?
- 2. How will network access be controlled?
 - a. Will VLANs be used to isolate traffic?
 - **b.** Will access control lists (ACLs) be used to isolate traffic?
- 3. How are users connected to the network?
 - a. Are they connected directly to the access switchport?
 - **b.** Are they connected via an IP phone?
 - **c.** Are they connecting via a host operating system from a virtual machine or using a secondary desktop switch?

After determining how users and devices will be authenticated, what network access they will be granted before and after authentication, and how devices will be allowed to connect to the network, you put all the answers together in one of the three deployment scenarios described later in this document.

Who is Connecting?

Authentication is the process by which the network establishes the identity of devices and users that are attempting to connect. To be able to authenticate users and devices, you must first decide what kind of credentials you will accept as valid identification. Ideally, you will use strong forms of identification such as digital certificates or properly encrypted usernames and password. A much weaker form of identification would be the MAC address of the connecting device.

The type of credentials that you accept determines in large part what method you use to validate those credentials. The authentication method, in essence, determines how a device submits its credentials to the network. 802.1X is an IEEE standard that defines a process by which a device can submit strong credentials such as digital certificates and passwords using a client, called a *supplicant*. 802.1X is a strong authentication method and is preferred in all cases where the device can support the required client. The specific details of which credentials are accepted and how they are submitted are determined by what is known as the Extensible Authentication Protocol (EAP) method. Common EAP methods include PEAP-MSCHAPv2 for username and password credentials, and EAP-TLS for certificate-based credentials.

For devices that cannot support the required 802.1X client, a supplementary form of authentication must be used. MAC Authentication Bypass (MAB) is a secondary authentication method in which the access switch detects the MAC address of the device and submits it as a form of identification.

For users that cannot support the required 802.1X client or do not have valid credentials, web authentication (WebAuth) may be used a secondary authentication method. WebAuth authenticates the user at the access edge by providing a web-based login page on which users can enter their credentials. The switch can host the login page locally, or it can redirect the browser to a centrally managed page on the Cisco NAC Guest Server. After the user is identified, the user identity can be mapped to a policy that grants or denies granular network access.

Whatever kind of credential you choose to accept, you must be able to validate it. This means that you must have a database of allowed devices, users, their credentials and a certificate chain of trust for certificate-based authentication. If you authenticating based on MAC address, you need a database of valid MAC addresses.

<u>P</u> Tin

Best Practice Recommendation—Leverage existing identity databases In most cases, there is no need to build a credential database from the beginning. For example, many organizations already possess a database of valid users and computers in the form of Microsoft Active Directory. Some organizations also maintain databases with the MAC addresses of corporate assets. Very often, these databases can be re-used for 802.1X, MAB, and WebAuth. Using these databases greatly simplifies your 802.1X deployment.

For guest access, you may want to consider establishing a database that allows the creation of temporary credentials that may be used for authenticated access (for example, using WebAuth) for a limited amount of time. The Cisco NGS provides a sponsored guest portal that allows authorized sponsors to create credentials for guest users. When guests enter those credentials at the web login page, the credentials are forwarded to the Cisco NGS for validation.

After you have determined the credential types, authentication methods, and credential databases that you will use, you will be able to fill out a simple table such as Table 3.

Authentication Method	Credential Type	Credential Database
802.1X	Username/password ¹	Active Directory
MAB	MAC address	LDAP database
WebAuth	Username/password	Active Directory (employees)
		NAC Guest Server (guests)

Table 3 Credentials

 The credential type depends on what EAP method you choose to implement. Username/password can be used for PEAP-MSCHAPv2 and others.

What Level of Authorization?

Authorization is the process by which an endpoint is granted a certain level of access to the network. In an identity-enabled network, network access should correspond to the authenticated identity of the endpoint. However, the network access of an endpoint can also depend on where the endpoint is in the authentication process. When planning a deployment, consider what access the endpoint should have at each of the following stages:

- Pre-Authentication, page 8
- Successful Authentication, page 8
- Failed Authentication, page 9

The authorization options available in each stage are also discussed in each of these topics.

Pre-Authentication

By default, endpoints are not authorized for network access prior to authentication. Before a device successfully authenticates via 802.1X or MAB, the port allows no traffic other than what is required for authentication. Access to the port is effectively *closed*. Although very secure, this method of access control can cause problems for devices that need network access before they authenticate, such as PXE devices that boot an OS from the network. It can also create problems for devices that are sensitive to delays in network access, which may result from the authentication process.

As an alternative, it is possible to configure Cisco switches for two other levels of pre-authentication authorization:

- Open access—Open access is the opposite of the default pre-authentication authorization. With open access, all traffic is allowed through the port before authentication. Although open access is obviously not an effective way to enforce network access control, it does have an important role in the initial stages of deploying 802.1X. This is discussed in the "Monitor Mode" section on page 12.
- Selectively open access (also known as low-impact)—Selectively open access represents a middle-ground between the default closed access and open access. With selectively open access, you can use a default port ACL to permit or deny specific traffic. For example, this can permit TFTP and DHCP traffic to allow PXE devices to boot before they authenticate.

Table 4 lists the three pre-authentication access levels.

Pre-Authentication Access Level	Implementation
No access	Closed (default)
Open access	Open
Selectively open access	Open with port ACL

Table 4 Pre-Authentication Access Levels

Be aware that the choices that you make regarding pre-authentication access influence your choices for post-authentication and failed-authentication access. This is discussed in more detail in the following sections.

Successful Authentication

After a successful authentication, the port is, by default, opened up completely and all traffic is allowed into the configured native VLAN of the port. This is a simple binary decision: anyone who successfully authenticates by any method is granted full access. To achieve differentiated access based on the identity of the authenticated user or device, it is necessary to use dynamic access control with VLANs and/or ACLs.

When post-authentication access is implemented with VLANs, the switch dynamically assigns a VLAN to a port based on the identity of the user or device that authenticated. For example, engineers could be assigned the *ENG* VLAN while accountants could be assigned the *FINANCE* VLAN. Although this form of dynamic authorization is a powerful tool for differentiating access for different user groups, it comes at a cost. Supporting multiple VLANs on every switch may require changes to the network architecture and addressing scheme. In addition, VLANs isolate traffic at Layer 2 in the OSI stack; therefore, dynamic VLAN assignment by itself cannot restrict access to specific subnets (Layer 3) or applications

(Layer 4 and above). However, dynamic VLAN assignment does provide the foundation for virtualizing IT resources using network virtualization solutions.

When an authorization is implemented with ACLs, the switch dynamically assigns an ACL to a port based on the identity of the device that authenticated. Engineers could be assigned an ACL that permits access to engineering subnets and applications while accountants get a different ACL. Although ACLs do not achieve the same level of logical isolation that VLANs provide, dynamic ACLs (dACLs) can be deployed without changing the existing network architecture and addressing schemes. On the other hand, care must be taken to ensure that the dACLs do not overwhelm the Ternary Content Addressable Memory (TCAM) capacity of the access switch. Well-summarized networks and good network design are essential to the creation of short but effective ACLs.

When deciding between dynamic VLANs and dACLs, another factor to consider is the form authorization to apply when a port is in the pre-authentication stage. dACLs work well with any kind of pre-authentication authorization. Dynamic VLAN assignment, on the other hand, does not typically work well with open or selectively open pre-authentication authorization.



Note

Why not use Dynamic VLAN assignment with open pre-authentication? When pre-authentication authorization is open, devices can receive IP addresses on the switchport VLAN subnet at link up. If a different VLAN is assigned as the result of an authentication, the old address is not valid on the new VLAN. 802.1X-capable devices with modern supplicants can typically detect the VLAN change and request a new address on the new VLAN but clientless devices, such as printers, are not able to do this.

Table 5 lists the various kinds of authorizations available after successful authentication and the deployment considerations for each method.

Post-Authenticatio n Authorization Method	Impact to Network Architecture	TCAM Impact	Compatible Pre-Authenticatio n Methods	Notes
Default Open	Minimal	None	Closed	May be sufficient for simple deployments or as a first step for more complex deployments
Dynamic VLAN	Significant	None	Closed	Required for network virtualization Provides logical isolation of traffic at L2
Dynamic ACL	Minimal	Significant	All	Does not support network virtualization Provides access control at L3 and L4

Table 5Authorization Types

Failed Authentication

After a failed authentication, the port is, by default, left in the same state as it was before authentication was attempted. If the port was in the default closed state before authentication, it remains closed after a failed authentication. If the port was in a selectively open state before authentication, it remains that way: that is, open in the statically configured VLAN and subject to the default port ACL.

Because failed authentications revert to the pre-authentication authorization, it is necessary to decide whether your chosen pre-authentication network access is adequate for endpoints that fail authentication. If not, it may be necessary to modify your pre-authentication network authorization policy or to use some of the mechanisms available for modifying the default failed authentication network access levels. Specific mechanisms are discussed in the deployment scenarios below.

How are Hosts Connected?

This section includes the following topics:

- Host Connection Overview, page 10
- Desktop Switches, page 11
- Host Movement, page 11

Host Connection Overview

Hosts connect to the access layer in all kinds of ways. The simplest connection is a direct point-to-point connection of one host to one switch port. In TrustSec deployments, this is sometimes referred to as *single host mode*. Single host mode is the most secure form of connection and the default mode on Cisco switches enabled for 802.1X and MAB. A switch running 802.1X in single host mode allows only one device at a time on that port. If another device appears on the port, the switch shuts down the port as a security precaution. Because only one device is allowed to connect to the port, there is no possibility of another device snooping the traffic from the authenticated device. A port in single-host mode effectively prevents casual port piggybacking.

Although it is the most secure mode, single host mode is not always sufficient. One common exception to the point-to-point connection assumption is IP telephony. In IP telephony deployments, two devices are often connected to the same switch port: the phone itself and a PC behind the phone. In this case, a new host mode, *multi-domain* (MDA), is required.

When properly enabled for MDA, the switch divides the switchport into two virtual domains. A domain is equivalent to a VLAN on a wired network. The switch independently and asynchronously authenticates the phone and the device behind the phone. When the phone authenticates successfully, it is given access to the voice domain. When the device behind the phone is authorized, it is given access to the data domain.

Some deployments include devices that contain multiple virtual machines even though there is physically only one connected device. Cisco switches support a third host mode, *multi-auth*, that allows each virtual machine to access the port after being authenticated. The multi-auth host mode is a superset of multi-domain host mode, meaning that the multi-auth host mode allows one voice device in the voice VLAN and any number of data devices in the data VLAN.

When planning your access control strategy, examine your network and determine how your hosts are connecting to the network. This will help you determine the most appropriate host mode to configure on the switch.

<u>)</u> Tip

Best Practice Recommendation—Use the most restrictive host mode that suits your network.

802.1X is most effective when it is most restrictive. If you do not have IP telephony, do not configure multi-domain host mode. If you do not have virtual machines with unique MAC addresses on the same physical host, do not configure multi-auth host mode. If possible, use the same host mode throughout your network. Using a standardized configuration minimizes operational costs.

Table 6 lists the host modes.

Table 6 Host Modes

With this Endpoint	Use this Host Mode
Point-to-point only (PC, printer, and so on)	Single-host

IP telephony	Multi-domain
Virtual machines (with or without IP telephony)	Multi-auth

Desktop Switches

Although the various host mode settings can handle most deployments, a special case exists when a secondary switch connects to the access switch in the wiring closet. For example, users that need more ports in a conference room might connect a small desktop switch to the access switch. Connecting a secondary switch in an IEEE 802.1X-enabled network introduces several conditions that must be addressed:

- Like any other device in an IEEE 802.1X-enabled network, the desktop switch must authenticate itself to gain any access at all.
- The access switch sees the MAC addresses of any devices that connect to the desktop switch, potentially triggering security violations or causing the access switch to try to authenticate those devices, even if the desktop switch has already authenticated them.
- If the devices connected to the desktop switch are in different VLANs, the port between the desktop switch and the access switch might need to be a trunk instead of an access port.

To address these conditions, Cisco switches support Network Edge Access Topology (NEAT). NEAT is a set of features that extends identity outside the wiring closet. NEAT allows you to configure a switch to act as a supplicant to another switch. Thus, with NEAT enabled, the desktop switch can become a supplicant switch and authenticate itself to the access switch. After authentication, NEAT leverages the Client Information Signaling Protocol (CISP) to send the MAC addresses connecting to the supplicant switch to the authenticator switch. CISP ensures that only traffic from authorized hosts, which connect to the switch with the supplicant, is allowed on the network. Lastly, NEAT can optionally be configured to change the authenticator switch port type from access to trunk and enable 802.1X trunk encapsulation.

Host Movement

In modern networks, hosts do not simply connect to the network once and stay connected forever. Wired hosts are mobile, moving from cube to conference room to cafeteria and back again, all in the course of a day. For hosts that are directly connected to 802.1X-enabled access switch ports, this kind of movement is not an issue: the switch detects a link down when the host unplugs and clears the authenticated session from the port. However, hosts that are indirectly connected, such as those behind an IP phone or a hub, can be problematic. Not knowing that an indirectly connected host has disconnected, the switch may trigger a security violation when a new host plugs into the same port or when the original host plugs into a different port on the same switch.

Cisco IP phones can work in conjunction with Cisco switches to facilitate host movement using the Cisco Discovery Protocol (CDP) Enhancement for Second Port Disconnect feature. This feature allows the phone to send a CDP message to the switch when a host unplugs from behind the phone. The switch is then able to clear the authenticated session for the indirectly connected host, exactly the same as if the host had been directly connected and the switch had detected a link down event.

For non-Cisco phones or other intermediary devices that cannot signal a link down event, Cisco switches support two additional features that facilitate host movement: MAC Move and MAC Replace.

MAC Move allows a host to move to another port on the switch, even if an authenticated session already exists on a different port. For example, suppose you unplug your laptop from behind a non-Cisco phone in your cube and plug directly into a port in a nearby conference room that is connected to the same switch as the phone. Not knowing that the laptop had unplugged, the switch would detect the same MAC

address on two ports and, by default, trigger a security violation. Although this default behavior helps prevent MAC spoofing, it also impedes host movement. If MAC Move is enabled, however, the switch deletes the session on the first port and re-authenticates the laptop on the second port.

MAC Replace allows one host to replace a previously authenticated host on the same port. For example, suppose you had temporarily plugged your laptop into a phone in a quiet room. After you leave the quiet room, another user tries to plug in behind the phone. Not knowing that the original laptop had disconnected, the switch detects a second MAC on the port and, by default, triggers a security violation. Although this mitigates against port piggybacking, it also impedes host movement. If MAC Replace is enabled, however, the switch deletes the first session on the port and authenticates the second device, effectively replacing the first authenticated session with the second authenticated session.

- Monitor Mode, page 12
- Low Impact Mode, page 15
- High Security Mode, page 19

Monitor Mode

This section includes the following topics:

- Overview, page 12
- Implementing Monitor Mode, page 13
- Next Steps, page 14

Overview

Authenticating the identity of a user without enforcing some form of authorization is like using a webcam with an unlocked door. You cannot physically prevent anyone from gaining access, but you can see who goes in and out. This kind of visibility is a non-trivial asset. Having visibility onto the network gives you insight into who is getting access, who has an operational 802.1X client, who is already known to existing identity stores, who has credentials, and so on. As a side benefit, some intruders may be deterred by the simple knowledge that someone is watching. These, in essence, are the goals of monitor mode.

When you deploy TrustSec in monitor mode, you enable authentication (802.1X and MAB) without enforcing any kind of authorization. There is no impact to users or endpoints of any kind: they continue to get exactly the same kind of network access that they did before you deployed TrustSec. The authorization level pre-authentication is the same as after successful authentications and failed authentications: completely open. In the background, however, the network is querying each endpoint as it connects and validating its credentials. By examining the authentication and accounting records at the AAA server, it is possible to determine the information listed in Table 7.

Endpoints on Your Network	How Determined
All endpoints/users with 802.1X clients and valid credentials	Passed 802.1X authentication records
All endpoints/users with 802.1X clients and invalid credentials	Failed 802.1X authentication records

Table 7Endpoint Authentication Information

All endpoints without 802.1X clients and known MAC addresses	Passed MAB authentication records
All endpoints without 802.1X clients and known MAC addresses	Failed MAB authentication records
Ports with multiple connected devices	Multiple authentication records for the same port on same switch.

Table 7Endpoint Authentication Information

Combining the information in authentication and accounting records results in very detailed knowledge of each endpoint that connects to your network: username, IP address, MAC address, port and switch where connected, time of connection, and so on. Even more information about the endpoint type, such as distinguishing an IP phone from a Windows PC, can be obtained in monitor mode by using active or passive profiling tools, such as Cisco NAC Profiler.

Implementing Monitor Mode

Monitor mode requires configuration on all components in the network: the AAA server, 802.1X-capable endpoints, and the Cisco switch.

The AAA server, such as Cisco Secure Access Control Server (ACS), should be fully configured for 802.1X and MAB authentication. The following checklist offers a summary of the tasks that are required to configure the AAA server:

- 1. Configure the ACS server to accept RADIUS authentication requests from all switches that are part of the TrustSec deployment. Make a note of the RADIUS shared secret, or *secret key*, which is configured for each switch.
- 2. Configure the ACS server to communicate with all available identity credential stores that will be used to validate identities, such as Active Directory, asset databases with known MAC addresses, and so on.
- **3.** Configure the ACS server to process 802.1X and MAB authentications using the identity credential stores. For 802.1X, configure all the EAP types you want to support. For each EAP type, ensure that you have the proper pre-requisites for that EAP type on the server. For example, PEAP and TLS both require a root CA certificate and a server certificate on the ACS.
- **4.** Ensure that all dynamic authorization, such as dynamic VLAN and dACL assignment, is *disabled* on the ACS server. Any form of dynamic authorization *will* impact end users and thus undermine the goal of monitor mode, which is end user transparency.

All endpoints that are capable of 802.1X should be supplied with a pre-configured supplicant as well as any additional credentials or certificates that may be required by the EAP types you support. Pre-configured supplicants include the Cisco Secure Services Client or the native operating system supplicant. EAP types may include PEAP and TLS, which both require a root CA certificate on the client, and TLS which also requires a client certificate.

The switches should be fully configured for 802.1X and MAB authentication. The following checklist summarizes configuration tasks on the switch:

- 1. First validate that your switch is in normal operating mode. All endpoints, including phones, should have full connectivity.
- **2.** Globally enable AAA and configure the switch to send 802.1X authentication requests to your AAA server. Be sure the RADIUS secret key that you configure on the switch matches what you configured in the AAA server.

- **3**. Globally enable 802.1X.
- **4.** On each access port, enable open access authentication and multi-auth host-mode. Without multi-auth, the switch runs in single host mode and disable any ports with multiple devices, including phones. Monitor mode requires multi-auth host-mode to be transparent to end users.
- 5. If needed, enable MAC Move and MAC Replace to facilitate host movement.
- 6. If needed, configure NEAT on supplicant (desktop) and authenticator (access) switches.
- 7. On each access port, enable 802.1X and MAB.
- **8.** Verify that all endpoints have exactly the same access as before authentication was enabled on the switch.

Next Steps

After completing the configuration steps, your network immediately begins authenticating users and devices and you have visibility into who and what is connecting to your network. You know the endpoints, such as PCs, printers, cameras, and so on that are connecting to your network, where they are connected, whether they are 802.1X-capable or not, and whether they have valid credentials. Additionally, you know whether endpoints have known valid MAC addresses as a result of any failed MAB attempts.

One primary benefit of monitor mode is that it enables you to proactively address any issues that would affect end users after access control is enabled. Table 8 lists the things you should do before moving to the next phase of deployment.

To Do	Key Issue	Remediation
Analyze 802.1X failures.	Are these valid devices or users that should be allowed access but are failing 802.1X?	Update credentials for valid devices and users so they will pass 802.1X.
Analyze MAB success.	Are there any devices doing MAB that should be capable of 802.1X?	Update those devices with supplicants and credentials so they can authenticate using 802.1X
Analyze MAB failures.	Are there managed assets that should be allowed access to the network but are failing MAB?	Update your asset database with these MAC addresses.
Analyze ports that have multiple devices on them.	Are these rogue hubs, valid virtual machines, or NEAT switches?	Remove rogue devices. Note ports that may legitimately require support for multiple hosts per port.

Table 8 Next Steps

After addressing all the issues uncovered by deploying TrustSec in monitor mode, deploy identity-based access control. The next two scenarios describe common ways to deploy access control. Many customers will choose to implement low impact mode only. Others may start with low impact mode to help assess the impact of access control to end users and then later move on to high security mode. Other customers may move straight from monitor mode to high security mode. The rest of this paper helps provide an understanding of the access-control deployment scenario and sequence that is right for your network.

I

Low Impact Mode

This section includes the following topics:

- Overview, page 15
- Deployment Considerations, page 16
- Implementing Low Impact Mode, page 17
- Next Steps, page 18

Overview

Low impact mode allows you to incrementally increase the level of port-based access control without impacting the existing network infrastructure. Low impact mode does not require the addition of any new VLANs nor does it impact your existing network addressing scheme. With low impact mode, you can add as little or as much access control as you want.

Low impact mode builds on top of monitor mode. In monitor mode, the pre-authentication authorization level was completely open. In low impact mode, the pre-authentication level is *selectively* open. The difference is that low impact mode adds an ingress port ACL that specifies exactly what traffic will be allowed before authentication. This ACL can be as restrictive or permissive as your network requires.

<u>}</u> Tip

Best Practice Recommendation—Use the ingress ACL to permit traffic that might be sensitive to authentication-related delays.

If your goal is to enable PXE machines to boot before authentication completes, you could permit DHCP, DNS, and TFTP in your ingress ACL.

In low impact mode, a successful authentication causes the switch to download a dACL that is applied on top of the ingress port ACL. The contents of the dACL are determined by the identity of the user or device that authenticated. In a simple deployment, an employee or managed asset that authenticates successfully could receive a **permit ip any any** dACL that fully opens up the port. More complex deployments could assign different ACLs based on different classes of employee. For example, engineers might be assigned a dACL that permitted all engineering related subnets, whereas accountants could be assigned a different dACL that denied access to engineering subnets but permitted all other access.



Both ingress ACLs and dACLs can be Layer 3 or Layer 4.

Using ACLs allows you to permit or deny access to particular hosts, entire subnets, and/or specific applications.

Whatever the contents of the dACL, the switch substitutes the source address of each access list element with the source address of the authenticated host, ensuring that only that host is allowed by the dACL.

Devices that fail authentication through 802.1X or MAB continue to have their access limited by the ingress port ACL.

Deployment Considerations

Because low impact mode can be deployed with little or no change to the existing campus network design, it is attractive to customers looking to deploy access control without altering the existing VLAN infrastructure or IP addressing scheme. Some customers may not want the additional IT overhead needed to add and support additional VLANs, while others may not even control the VLAN infrastructure. The latter situation would occur at a branch office where the service provider owns the routers and has implemented MPLS. In such a case, ACL-based enforcement is the only choice for port-based access control. However, the following deployment considerations should be understood prior to deploying this model of access control:

• The current implementation of dACLs requires a pre-configured static port ACL on every access port that may download an ACL.

If the switch attempts to apply a dACL to a port without a pre-existing port ACL, the authorization fails and users are not able to gain access, even if they presented valid credentials and passed authentication.

 \mathcal{P}

Best Practice Recommendation—Leverage switch syslog messages for improved troubleshooting of ACL problems.

The switch generates syslog messages if authorization fails because the port ACL was not configured. Starting with version 5.1, the Cisco ACS can consume these syslogs and incorporate them into its authentication reports. Having a single, centralized view of the end-to-end process greatly enhances your ability to monitor and troubleshoot the network. Therefore, for the best centralized troubleshooting, enable the switch to send syslogs to ACS.

• Because the static port ACL is a port-based ACL, it applies to both the data VLAN and, if present, the voice VLAN.

Because the switch performs source address substitution on the dACL, traffic from the phone is not permitted by a dACL downloaded by a data device authenticating behind the phone. This means that both the phone and any devices behind the phone must authenticate individually and download their own dACL. This means that the host mode on the port must be multi-domain.

Cisco switches use TCAM to support wire-rate ACL enforcement.

TCAM is a limited resource that varies by platform. If the TCAM is exhausted, switching performance can be degraded. Therefore, it is important to verify that your access switches have sufficient TCAM to support the number and length of ACLs (static and dynamic) that your deployment will require.

- Dynamic (downloadable) ACLs (dACLs) extend the standard RADIUS protocol to optimize the downloading process and support ACLs of arbitrary size. Use the Cisco ACS server as your AAA server to support downloadable ACLs.
- Because the switch performs source substitution on the source address of the dACL, the switch does not enforce the dACL until it learns the source address of the authenticated host. IP address learning is enabled by the IP Device Tracking feature on the switch.
- When designing your ingress port ACL, be aware that this ACL restricts access before authentication *and* after failed authentications.

You should design the ingress port ACL with this in mind. For example, if you want employees that fail 802.1X because of an expired certificate to be able to download a new certificate, you should consider allowing access to the CA server in the ingress ACL. Or, if you want a contractor that fails 802.1X or MAB to be able to access the Internet or VPN to a home network, you should also allow that traffic in the ingress port ACL.

Implementing Low Impact Mode

Before transitioning to low impact mode from monitor mode, ensure that all endpoints that should authenticate, can authenticate. All identity store databases should be up-to-date and online. To enable a smooth transition from monitor mode, configure the switch first. To ensure that devices continue to get access during the transition, complete the following steps.

Procedure

Step 1 Configure an ACL on the switch that is completely open (**permit ip any any**).

This ACL is also called DEFAULT-ACCESS in the examples provided in this document.



The DEFAULT-ACCESS ACL is used for troubleshooting only and is modified in the final step.

- **Step 2** Configure the switch to accept authorization instructions from the AAA server.
- **Step 3** Apply the **PERMIT-ANY** ACL to all access ports on the switch.
- **Step 4** Enable IP Device Tracking.
- **Step 5** Configure the switch to send syslogs to ACS (version 5.1 only).

At this point, network access levels should not have changed from monitor mode because the *PERMIT-ANY* port ACL permits all traffic.

On the AAA server, complete the following steps.

Procedure

Step 1 Configure downloadable ACLs to match your policy.

A downloadable ACL can be as simple as a single dACL, such as **permit ip any any** for all users and devices that authenticate. Downloadable ACLs can include multiple dACLs if you want to differentiate between different classes of users or devices.

- **Step 2** Configure your ACS to apply the appropriate dACL to each class of users and devices, including phones.
- **Step 3** For phones, also ensure that the ACS is configured to send down the appropriate permission to allow the phone to access the voice VLAN.



This is a Cisco VPN Services Adapter: *device-traffic-class=voice*.

At this point, endpoints that connect to the network should start downloading the appropriate dACL based on their authenticated identity.

Step 4 Verify that the switch is successfully downloading and applying the dACLs.

The final step is to replace the PERMIT-ANY port ACL on the switch with a port ACL that restricts traffic before authentication.

- Step 5 Modify the DEFAULT-ACCESS ACL on the access ports of the switch so that it limits access before authentication in accordance with your policy.
- **Step 6** Unless you have a specific need to support multiple data devices on a single port, configure all access ports for single host-mode for non-IP-telephony deployments or multi-domain host-mode for IP telephony deployments.

When the endpoints re-authenticate, the dACL is applied on top of the DEFAULT-ACCESS ACL. Ensure that the combined ACL gives the endpoints the level of access you intended.



The default port ACL can always be modified.

The default port ACL, identified as DEFAULT-ACCESS in the example above, is not set in stone. You can start with a **permit any** statement, which is used in the example above to transition from monitor mode. You can then selectively add **deny** statements to prevent access to sensitive assets before authentication. Over time, you may transition from a default port ACL that denies access to a few resources and permits everything else to one that permits access to specific resources, such as DHCP, DNS, TFTP for PXE, and so on, and denies everything else. Evolving the default ACL in this way allows you to incrementally add access control without inadvertently blocking important traffic (see Figure 1).

Figure 1 Two Models for Default Port ACL



Next Steps

For many deployments, low impact mode is the final step in your TrustSec deployment. If this mode provides you with the access control that you need, the only "next step" is to monitor your network and to fine-tune ACLs as required by your policy.

However, if your network security requirements evolve and you no longer want to offer any form of pre-authentication access, you can move to the next phase: high security mode.

If low impact mode does not meet your network and design security requirements in the first place, you may skip low impact mode altogether and go directly from monitor mode to high security mode.

High Security Mode

This section includes the following topics:

- Overview, page 19
- Deployment Considerations, page 20
- Implementing High Security Mode, page 21

Overview

High security mode returns to a more traditional deployment model of 802.1X. In a properly prepared network, high security mode gives you total control over network access at Layer 2.

In high security mode, the port is kept completely closed until a successful authentication takes place. There is no pre-authentication access. For users and devices that successfully complete 802.1X, this is not typically an issue because 802.1X authentication is usually very quick.

Note

This assumes a single sign-on deployment, where credentials are automatically gleaned from the device or user. There may be some delay in an environment that requires manual sign on, which typically uses a pop-up window to enter the username and password.

For devices that cannot perform 802.1X, however, there may be a significant delay in network access. Because the switch always attempts the strongest secure authentication method first, non-802.1X-capable devices must wait until the switch times out the 802.1X authentication and falls back to MAB as a secondary authentication method.

One solution to the delays associated with MAB in high security mode is to configure the switch to perform MAB before 802.1X. This way, non-802.1X devices get immediate access after successful MAB.

After a successful authentication, network access, by default, changes from completely closed to completely open. To add more granular access control, high security mode may use dACLs to provide restrictions of available network services, as well as using dynamic VLAN assignment to isolate different classes of users into different broadcast domains.



It is possible to use dACLs and dynamic VLAN assignment at the same time. However, for the sake of simplicity, this deployment mode focuses only on dynamic VLAN assignment.

By isolating traffic from different classes of users into separate VLANs, high security mode provides the foundation for virtualized network services.

Note

For more information on Network Virtualization solutions, see http://www.cisco.com/en/US/netsol/ns658/networking_solutions_package.html

Devices that cannot authenticate or fail to authenticate retain the same level of access that they had before authentication; in other words, no access at all.

Deployment Considerations

Deploying high security mode with VLAN assignment can have a significant impact on network architecture. Understanding these potential impacts is essential to a successful deployment of this mode. This section describes some important deployment considerations and includes the following topics:

- Dynamic VLAN Assignment, page 20
- Supporting Multiple VLANs, page 20
- Changing Authentication Order, page 21
- Supporting Devices Without 802.1X or MAB, page 21

Dynamic VLAN Assignment

Dynamic VLAN assignment requires that every dynamic VLAN be supported on every access switch to which a user might connect and authenticate.

This requirement has several repercussions. For example, if you have three user groups to which you wish to assign unique VLANs, such as Engineering, Finance, and HR, every access switch must have those three VLANs defined by name. Note that the number assigned to the VLAN does not have to be the same. If the switch attempts to apply a non-existent VLAN to a port, the authorization fails and users are not able to gain access, even if they presented valid credentials and completed authentication.

Deploy the User Distribution feature to avoid renaming existing VLANs. User Distribution allows you to map multiple VLANs to a VLAN group name. This can be useful in large campus LANs because it allows the switch to load balance users in the same group across different VLANs, thus reducing the size of the broadcast domain for any single VLAN. The User Distribution feature was originally developed for this use case.

The User Distribution feature can be used to map the dynamic VLAN name sent by the RADIUS server to a different VLAN name on the switch. Suppose your existing environment has two switches. On Switch 1, the engineering VLAN is named ENG-SW1. On Switch 2, the engineering VLAN is named ENG-SW2. With User Distribution, you can map ENG-SW1 to ENG-SW on Switch 1, and ENG-SW2 to ENG-SW on Switch 2. That way, when the RADIUS server sends down ENG-SW to either switch, Switch 1 applies ENG-SW1 VLAN to the port and Switch 2 applies ENG-SW2 VLAN.

<u>}</u> Tip

Best Practice Recommendation—Leverage switch syslog messages for improved troubleshooting of VLAN problems.

The switch generates syslogs if authorization fails because the RADIUS-assigned VLAN name does not exist on the switch, or is not mapped to an existing VLAN name using the User Distribution feature. Starting with version 5.1, the Cisco ACS can consume these syslogs and incorporate them into its authentication reports. Having a single, centralized view of the end-to-end process greatly enhances your ability to monitor and troubleshoot the network. Therefore, for the best centralized troubleshooting, enable the switch to send syslogs to ACS.

Supporting Multiple VLANs

Supporting multiple VLANs per access switch is non-trivial from an IP addressing perspective.

Good campus design principles dictate a single subnet per VLAN with no VLAN spanning more than one switch. Your IP addressing scheme should support multiple subnets per switch in such a way that it does not over-burden the control and data planes of the campus distribution block.

Description Best Practice Recommendation—Use the minimum number of VLANs necessary.

The fewer VLANs you assign, the more manageable and scalable your solution will be. Indeed, some customers have found that, upon analysis, their security policy requirements could be met with very few VLANs, such as Employee, Guest/Fail, and Voice.

Changing Authentication Order

If you choose to change the order of authentication to perform MAB before 802.1X, be aware that this means that every device—even those capable of 802.1X—are subject to MAB. This could significantly increase the amount of control plane traffic in your network.

If there are devices in your network that might pass both 802.1X and MAB, either make sure that no 802.1X-capable devices are in the MAB database, or configure the switch to prioritize 802.1X over MAB so that the port processes 802.1X authentication after successful MAB authentication.



Best Practice Recommendation—Always prioritize 802.1X over MAB.

Some level of access may be needed for devices that fail 802.1X. For example, this might be required to allow employees with expired certificates to download a new certificate. It is possible to configure the solution to grant limited access based on the type of authentication method that failed.

If 802.1X fails, the switch can be configured to open the port into a special VLAN (the Auth-Fail VLAN) for this purpose.

The switch can also be configured to *fail back* to a MAB authentication if 802.1X fails. However, 802.1X with failover to MAB should typically not be deployed if the authentication order has been changed to do MAB first.

Supporting Devices Without 802.1X or MAB

There may be devices on your network that cannot perform 802.1X and cannot pass MAB. For example, a contractor with no supplicant may need to establish a VPN to a home network.

For unknown MAC addresses that fail MAB, it is possible to configure the Cisco ACS server with an unknown MAC address policy. Such a policy allows ACS to instruct the switch to allow devices with unknown MAC addresses into a dynamically assigned VLAN. In essence, an unknown MAC policy enables a dynamic version of the Auth-Fail VLAN for failed MAB attempts.

Implementing High Security Mode

Before transitioning to high security mode, you should ensure that all endpoints that should authenticate, can authenticate. All identity store databases should be up-to-date and online.

On the switch, complete the following steps, which assume you have previously configured the solution for monitor mode:

Step 1 Verify that all VLANs that may be assigned are defined by name on the access switch and that each VLAN has the expected connectivity. Use the User Distribution feature to map existing VLAN names if necessary. Step 2 Verify that the switch has been configured to accept authorization instructions from the AAA server. Step 3 Remove any ingress port ACLs from the switch. This deployment scenario does not require ACLs. Note Although ACLs are not required in this scenario, they are supported should you choose to customize this scenario with dACLs. Step 4 Disable open access authentication on all ports. If desired, configure the authentication order to perform MAB before 802.1X and modify the authentication priority so that 802.1X can preempt a successful MAB authentication. Step 5 Unless you have a specific need to support multiple data devices on a single port, configure all access ports in single host-mode for non-IP-Telephony deployments, or in multi-domain host-mode for IP telephony deployments.

On the ACS server, complete the following steps.

Procedure

Procedure

Step 1 Remove any dACL assignments.

This deployment scenario does not require ACLs.

Step 2 Configure your ACS to apply the appropriate VLANs to each class of users and devices except phones.Phones continue to use the statically configured voice VLAN on the port.



Many Cisco switches also support dynamic voice VLAN assignment. For simplicity, this scenario uses the static voice VLAN.

Step 3 For phones, ensure that the ACS is configured to send down the appropriate permission to allow that device to access the voice VLAN.

This is the Cisco *device-traffic-class=voice* VPN Services Adapter.

References

TrustSec 1.99 Documents

- Wired 802.1X Deployment Guide http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/TrustSec_1.99/Dot1X_Deployme nt/Dot1x_Dep_Guide.html
- IP Telephony for 802.1X Design Guide http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/TrustSec_1.99/IP_Tele/IP_Teleph ony_DIG.html
- MAC Authentication Bypass Deployment Guide http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/TrustSec_1.99/MAB/MAB_Dep_ Guide.html
- TrustSec Phased Deployment Configuration Guide http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/TrustSec_1.99/Phased_Deploy/Ph ased_Dep_Guide.html
- Local WebAuth Deployment Guide http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/TrustSec_1.99/WebAuth/WebAut h_Dep_Guide.html
- Scenario-Based TrustSec Deployments Application Note http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/TrustSec_1.99/Scenario_based_A ppNote/Scenario_based_AN.html
- TrustSec 1.99 Deployment Note: FlexAuth Order, Priority, and Failed Authentication http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/TrustSec_1.99/FlexAuthNote/flex auth-note.html
- TrustSec Planning and Deployment Checklist http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/TrustSec_1.99/TrustSec_Checklis t/trustsec-199_checklist.html

Related Documents

- Configuring WebAuth on the Cisco Catalyst 3750 Series Switches http://www.cisco.com/en/US/partner/docs/switches/lan/catalyst3750/software/release/12.2_55_se/ configuration/guide/sw8021x.html
- Configuring WebAuth on the Cisco Catalyst 4500 Series Switches http://www.cisco.com/en/US/partner/docs/switches/lan/catalyst4500/12.2/53SG/configuration/web auth.html
- Configuring WebAuth on the Cisco Catalyst 6500 Series Switches http://www.cisco.com/en/US/partner/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/gui de/webauth.html
- Cisco IOS Firewall authentication proxy http://www.cisco.com/en/US/products/sw/secursw/ps1018/products_tech_note09186a0080094eb0. shtml

• WebAuth with Cisco Wireless LAN Controllers http://www.cisco.com/en/US/partner/tech/tk722/tk809/technologies_configuration_example09186 a008076f974.shtml#external-process

1