



Network Virtualization—Access Control Design Guide

This document provides design guidance for enterprises that want to provide Internet and limited corporate access for their guests and partners. Several solutions for guest and partner access challenges are proposed and analyzed in this document, at both the architectural and functional levels. For related information, see the following documents:

- *Network Virtualization—Guest and Partner Access Deployment Guide—*
http://www.cisco.com/en/US/docs/solutions/Enterprise/Network_Virtualization/GuestAcc.html
- *Network Virtualization—Network Admission Control Deployment Guide—*
http://www.cisco.com/en/US/docs/solutions/Enterprise/Network_Virtualization/NACDepl.html
- *Network Virtualization—Path Isolation Design Guide—*
http://www.cisco.com/en/US/docs/solutions/Enterprise/Network_Virtualization/PathIsol.html
- *Network Virtualization—Services Edge Design Guide—*
http://www.cisco.com/en/US/docs/solutions/Enterprise/Network_Virtualization/ServEdge.html

Contents

Introduction	3
Technology Scope	5
Product Scope	6
Client-Based Authentication	7
802.1x Framework	7
Wireless Guest Access	7
Lightweight Access Point Deployment with the Cisco WLAN Controller	8
802.1x Authentication Failure VLAN (Wired)	11
Auth-Fail-VLAN Operational Overview	13
Auth-Fail-VLAN Configuration	13
Auth-Fail-VLAN Verification	14




Auth-Fail-VLAN Summary and Recommendations	27
Clientless-Based Authentication	28
Static VLAN Configuration	29
802.1x Guest-VLAN	29
802.1x Guest-VLAN Functionality	30
802.1x Guest-VLAN Configuration	30
Wake-on-LAN Primer	33
Guest-VLAN and WoL Interaction	35
Interaction with VoIP Deployments	36
Guest-VLAN Summary	43
MAC Authentication Primer	43
MAC Authentication Bypass Operational Overview	45
802.1x Rehearsal	45
Guest-VLAN Rehearsal	46
MAB Operation	47
Functional Details	49
MAC Authentication Bypass Configuration and Verification	50
Configuration	51
802.1x Timeout	51
Verification	55
MAC Authentication Bypass Feature Interaction	56
MAB and EAPOL Interaction	56
MAB and the Guest-VLAN	57
MAB and WoL Interaction	58
MAC Authentication Bypass Opportunities and Benefits	59
Location-Based Awareness	59
Network Access Profile Matching and Potential Value	60
Fallback Technique for New/Re-imaged Machines with WZCSVC	62
MAC Authentication Bypass Limitations and Challenges	64
Fallback Technique for Re-imaged Machines with CSSC	64
Network Admission Control	65
MAB EAP Option on ACS	70
ACS 4.1	70
Provisioning	72
Lack of Existing Identity Store	73
Lack of Voice Support	73
MAC Movement	74
MAC Authentication Bypass Policy Assignment	75
MAC Authentication Bypass Summary	76
Overall Summary	77

Introduction

The term *network virtualization* refers to the creation of logical isolated network partitions overlaid on top of a common physical infrastructure (see Figure 1). Each partition is logically isolated from the others, and must behave and appear as a fully dedicated network to provide privacy, security, and an independent set of policies, service levels, and even routing decisions.

Figure 1 Network Virtualization



Network virtualization provides multiple solutions to business problems and drivers that range from simple to complex. Simple scenarios include enterprises that want to provide Internet access to visitors (guest access). The stringent requirement in this case is to allow visitors external Internet access, while simultaneously preventing any possibility of unauthorized connection to the enterprise internal resources and services. This can be achieved by dedicating a logical “virtual network” to handle the entire guest communication path. Internet access can also be combined with connectivity to a subset of the enterprise internal resources, as is typical in partner access deployments.

Another simple driver for network virtualization is the creation of a logical partition dedicated to the machines that have been quarantined as a result of a Network Admission Control (NAC) posture validation. In this case, it is essential to guarantee isolation of these devices in a remediation segment of the network, where only access to remediation servers is possible until the process of cleaning and patching the machine is successfully completed.

Complex scenarios include enterprise IT departments acting as a service provider, offering access to the enterprise network to many different “customers” that need logical isolation between them. In the future, users belonging to the same logical partitions will be able to communicate with each other and to share dedicated network resources. However, some direct inter-communication between groups may be prohibited. Typical deployment scenarios in this category include retail stores (for example, Best Buy, Albertson’s, Wal-Mart, and so on) that provide on-location network access for kiosks or hotspot providers.

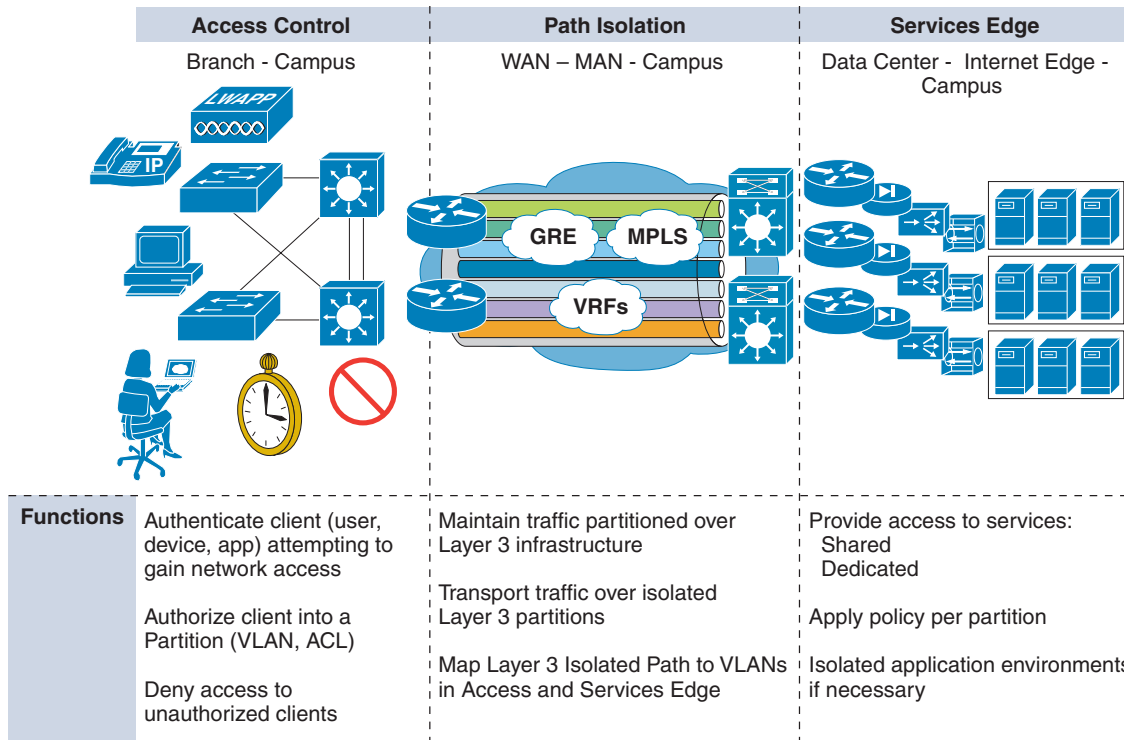
The architecture of an end-to-end network virtualization solution targeted to satisfy the requirements listed above can be separated in the following three logical functional areas:

- Access control
- Path isolation

- Services edge

Each area performs several functions and must interface with the other functional areas to provide the end-to-end solution (see Figure 2). This design guide focuses on the access control functional area to securely grant and control authorized access into any internal network system, while providing optional access to guests or partners.

Figure 2 Network Virtualization—Three Functional Areas



221036

The access control functional area identifies the users or devices logging into the network so they can be successfully assigned to the corresponding groups. An identity is an indicator of a client in a trusted domain. In this architecture, it is used as a pointer to a set of rights or permissions to allow for client differentiation. The model described in this document demonstrates how to use identities as not only a security mechanism, but also how to use identity to provide permissions to service within a domain. Although network services are arbitrary, this represents a linkage to path isolation techniques to provide a holistic form of differentiation between various types of clients. Access control also promotes authentication: the process of establishing and confirming the identity of the client requesting services. Authentication is crucial for network-based security benefits, and to establish corresponding authorization as well.

When identified, the endpoints must be authorized onto the network. To achieve this, the enterprise LAN edge port on which an endpoint connects is activated and configured with certain characteristics and policies. Examples of authorization include the configuration of the VLAN membership of a port based on the results of an authentication process, and the dynamic configuration of port ACLs based on the authentication.

**Note**

For wireless access, the concept of a port can be replaced by the association between client and access point (AP). When authorizing a wireless device, the association is customized to reflect the policy for the user or device. This customization can take the form of the selection of a different wireless LAN (WLAN), VLAN, or mobility group, depending on the wireless technology employed.

When an endpoint is authorized on the network, it can be associated to a specific group that typically corresponds to a separate partition or domain. Thus, the authorization method ultimately determines the mapping of the endpoint to an end-to-end virtual network. For example, when a VLAN is part of a virtual network, a user authorized onto that VLAN is therefore authorized onto the virtual network.

The main authentication scenarios for the enterprise are as follows:

- Client-based authentication for endpoints with client software
- Clientless authentication for endpoints with no client software

The current state of the technology provides broad support for VLAN assignment as an authorization alternative. In the cases where policy changes based on authentication are required and only VLAN assignment authorization is available, a static assignment of a policy to a VLAN provides the required linkage between the user authorization and the necessary policy. In effect, the policy is applied to the VLAN because users are subject to the policy when authorized onto the VLAN. The primary use of VLAN assignment promotes differentiation, and is critical to linkages to path isolation techniques. In essence, VLANs may be mapped into separate policy domains, which define the correct entrance criteria into the path isolation architecture alternatives.

Various access control technologies are discussed in this document: 802.1x, Guest-VLAN, Auth-Failed VLAN, MAC-Authentication Bypass (MAB), and so on. Note the following two important points:

- The various access control technologies are discussed in the context of network virtualization. This means, for example, that the reader should not expect to find here all the details regarding 802.1x deployments, but only the portions of that technology that have been validated and positioned as part of the network virtualization project to provide an answer to the business problems previously listed.
- Not all the technologies found in this design guide represent the right fit for each business problem. For example, the use of Guest and Auth-Failed VLAN features may be particularly relevant in guest and partner access scenarios, but not in deployments aiming to fulfill different business requirements (as for example, NAC quarantining). To properly map the technologies discussed here with each specific business problem, it is thus recommended to see the accompanying deployment guides:
 - Network Virtualization—Guest and Partner Access Deployment Guide (OL-13635-01)
 - Network Virtualization—Network Admission Control Deployment Guide (OL-13636-01)
 - Network Virtualization—Network Hosted Access Deployment Guide (OL-13634-01)

Technology Scope

The client-based framework focuses on 802.1x only as the access control method to provide holistic control over client access to the network. 802.1x always assumes a supplicant at the edge. 802.1x can give customers ubiquitous, port-based access control and provides them with the ability to manage access control on multiple levels for wired and wireless integration purposes. In support of network virtualization, 802.1x can also allow customers to leverage the notion of an authenticated identity with granular policy controls. Although out of this document scope, 802.1x can also provide auditing/accounting measures to network visibility and automate encryption techniques for end stations (wireless only today).

Upon evaluation of 802.1x, a customer must take Guest-VLAN interoperability into account. This design guide discusses recent changes in this arena. It also addresses the Auth-Fail-VLAN to provide wired topologies a method to provide clients network access that is illegitimate and be otherwise failed on any connection attempt into the networked system. The Auth-Fail-VLAN is positioned here as a means to provide access for the 802.1x-enabled partner or guest. It is not positioned as a de facto recommendation for any 802.1x deployment. This design guide also introduces other clientless methods of access control to provide access as well. This form of access control is device-specific in nature, and is discussed in the wired context only. This functionality is MAC-Auth-Bypass. In all cases, Windows Active Directory was used as the backend identity store as the verified directory infrastructure.

This document does not discuss the following technology areas:

- Web-Auth
- IPsec authentication/remote access
- In-depth concepts on identity management and single sign-on
- Privacy issues—Packet confidentiality and integrity
- Topology-independent access control
- In-depth policy administration
- In-depth authorization techniques
- Specific EAP methods
- X.509 certificates and PKI
- EAP over UDP (EAPoUDP)
- NAC posture assessment/remediation

Product Scope

The 802.1x supplicants discussed in both the wired and wireless contexts are the Cisco Secure Services Client (CSSC) version 4.0.5 and the embedded supplicant available in Windows XP SP2. The switches and corresponding Cisco IOS revisions examined as part of the demonstrated architecture are as follows:

- Cisco Catalyst 3750—12.2(25)SEE
- Cisco Catalyst 4500—12.2(31)SG
- Cisco Catalyst 6500—8.5(4)

Wireless was also tested in the scope of access control. To achieve this, a Cisco Wireless LAN Controller (WLC) 4402 version 3.2.150.6 was tested with Cisco 1240 access points. As the backend authentication server, the CiscoSecure ACS Appliance 1112 version 4.0 was tested.

No other supplicants, network infrastructure devices, or authentication servers or directory infrastructures were tested as part of this architecture and cannot be validated as working components of the solution at this time.

Client-Based Authentication

802.1x offers an efficient framework to a protected network for authenticating and administering user traffic. Together with technology extensions and supplemental authentication techniques, 802.1x builds on access control to establish a technology solution that can improve the security of physical and logical access to LANs.

802.1x Framework

The use of 802 LANs in public and semi-public places has dramatically increased. There is now a desire to provide a mechanism to associate identities with the port of access to the LAN to establish authorized access. 802.1x ties the Extensible Authentication Protocol (EAP) to both the wired and wireless LAN media and supports multiple authentication methods. 802.1x defines a generic framework that is able to use different authentication mechanisms without implementing these mechanisms outside the backend authentication infrastructure and client devices. 802.1x specifies a protocol framework between devices desiring access to a LAN (supplicants) and devices providing access to a LAN (authenticators). Various credentials, such as token cards, Kerberos, one-time password, certificates, and public key authentication can be used with 802.1x. Primarily, 802.1x is an encapsulation definition for EAP over an IEEE 802 media. This is known as EAP over LAN, or EAPOL. EAPOL transports authentication messages (EAP) between supplicant (user/PC) and authenticator (switch or access point). 802.1x always assumes a secure connection, and the actual enforcement is done via MAC-based filtering and port-state monitoring.

Although 802.1x is the recommended method to deploy access control in an enterprise environment, it is not the specific focus in this paper. The business problems that network virtualization is aimed to solve in this phase include the following:

- Guest access
- Partner access
- NAC remediation
- Hosted access

Hosted access and NAC remediation environments are not typically enabled for 802.1x at present. The need remains for some way to provide access to guest or partners when they are equipped with an unmanaged 802.1x supplicant. The 802.1x supplicant of the guest or partner may indeed be managed, but not by the IT staff that owns the network into which they plug. Thus, this design guide focuses only on what it takes to allow guest or partner online access in a virtualized environment when they are equipped with an 802.1x supplicant on the device.

Wireless Guest Access

Wireless users typically access the network differently than wired users. The paradigm of public access has extended to the enterprise. Mobility demands network connectivity. Enterprise guest access services are now a necessity in the corporate environment. The solution is made up of many components: access points, controllers, and management systems.

A detailed description and comparison of the various wireless deployment options is not within the scope of this document; a brief, high-level description of each scenario is provided in the following sections but only in the context of network access control. For more information on Cisco Integrated Wireless Networks, see the following URL:

http://www.cisco.com/en/US/netsol/ns340/ns394/ns348/ns337/networking_solutions_package.html.


A typical company security policy most likely requires the implementation of various types of authentication and encryption for various types of users. For example, open authentication and no encryption are the typical choices when providing guest access, whereas 802.1x authentication and strong encryption are usually adopted for internal employees. This is achieved by defining multiple service set identifiers (SSIDs) on each access point, with each SSID characterized by its own security policies.

End users associate with the closest access point by selecting a specific SSID to access the enterprise network. After this point, the WLAN Controller allows traffic to be logically separated from the traffic for users belonging to different groups. This is described in more detail in the following section.

Lightweight Access Point Deployment with the Cisco WLAN Controller

A WLAN controller system is used to create and enforce policies across many different lightweight access points in this architecture (see [Figure 3](#)). Security, mobility, quality of service (QoS), and other functions essential to WLAN operations can be efficiently managed across an entire wireless enterprise by centralizing intelligence within a controller system. Furthermore, by splitting functions between the access point and the controller, IT staff can simplify management, improve performance, and increase security of large wireless networks.

Figure 3 Cisco Wireless LAN Controller




LWAPP revolutionizes the way WLAN deployments are managed with the concept of split MAC, which means the ability to separate the real-time aspects of the 802.11 protocol from most of its management aspects. In particular, real-time frame exchange and certain real-time portions of MAC management are accomplished within the access point, while authentication, security management, and mobility are handled by WLAN controllers. The Cisco Centralized WLAN Solution, which uses LWAPP, is the first centralized WLAN system to use the split MAC.

From a traffic handling perspective, all data traffic originating from wireless clients associated to the distributed lightweight access points is encapsulated on the access points themselves and carried to a centralized wireless LAN controller, which aggregates the traffic and represents the single point of ingress and egress for IP traffic to and from the wired network. Traffic is tunneled from the access points to the centralized controller, leveraging LWAPP. The LWAPP tunnel is a Layer 2 tunnel (the original Ethernet frame is LWAPP-encapsulated), which carries both control and data traffic. Data traffic uses UDP port 12222, control traffic is encapsulated in UDP port 12223, and Radio Resource Manager uses ports 16666/16667. In addition, the control traffic is AES-encrypted, while the data is in the clear.

There is not a separate logical tunnel for each defined SSID; only a single logical tunnel is built between each access point and the centralized WLAN controller. This LWAPP tunnel is used to carry the data traffic for all the wireless clients associated to the access point, independently from the SSID they are using for this association.

[Figure 4](#) shows the deployment of the lightweight architecture in an enterprise campus network where two categories of users (employees and guests) are defined as an example.

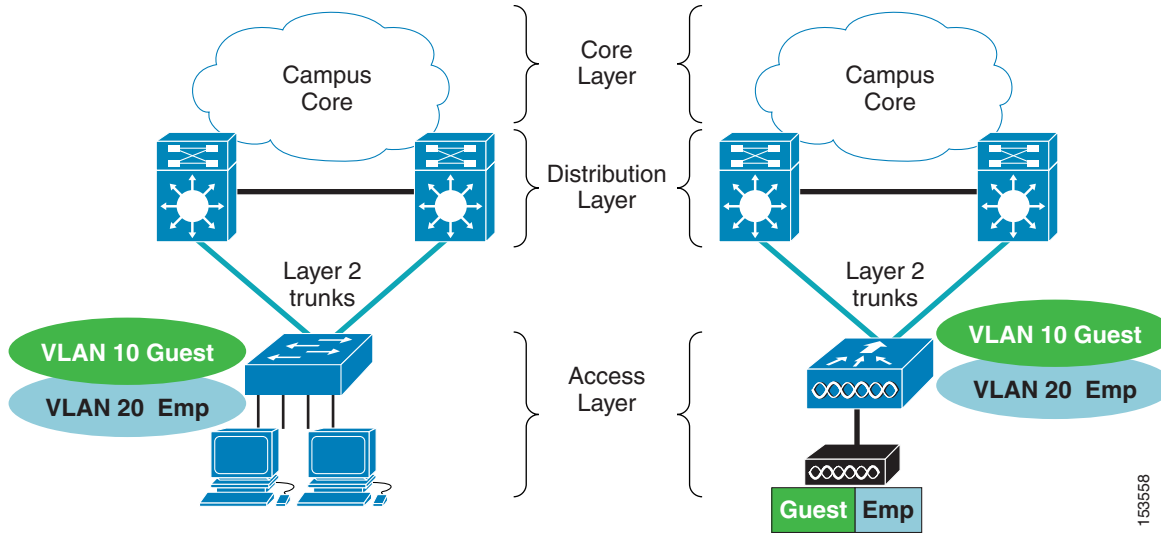
Figure 4 Lightweight Architecture Deployment



From the traffic isolation perspective, this scenario is very similar to the wired deployment in a traditional campus design. The reason is that traffic from various categories of users associating with their own SSID, after being aggregated to the main WLAN controller, is bridged to a corresponding VLAN and carried up to the first Layer 3 hop device.


Figure 5 shows how the use of VLANs allows maintaining separation between the guest traffic and the enterprise internal traffic in the Layer 2 domain, in a very similar way to the wired scenario for a traditional campus deployment (Layer 2 in the access).

Figure 5 *Similarities Between Wired and Wireless Deployments*



Several alternative designs can be deployed when positioning the WLAN controllers in the campus network. Cisco recommends placing the WLAN controllers in a centralized location (for example, a data center) to leverage the high availability and continuous monitoring characteristic of such an environment. (See [Figure 6](#).)

Figure 6 Cisco Wireless LAN Controller Deployment in a Campus Network

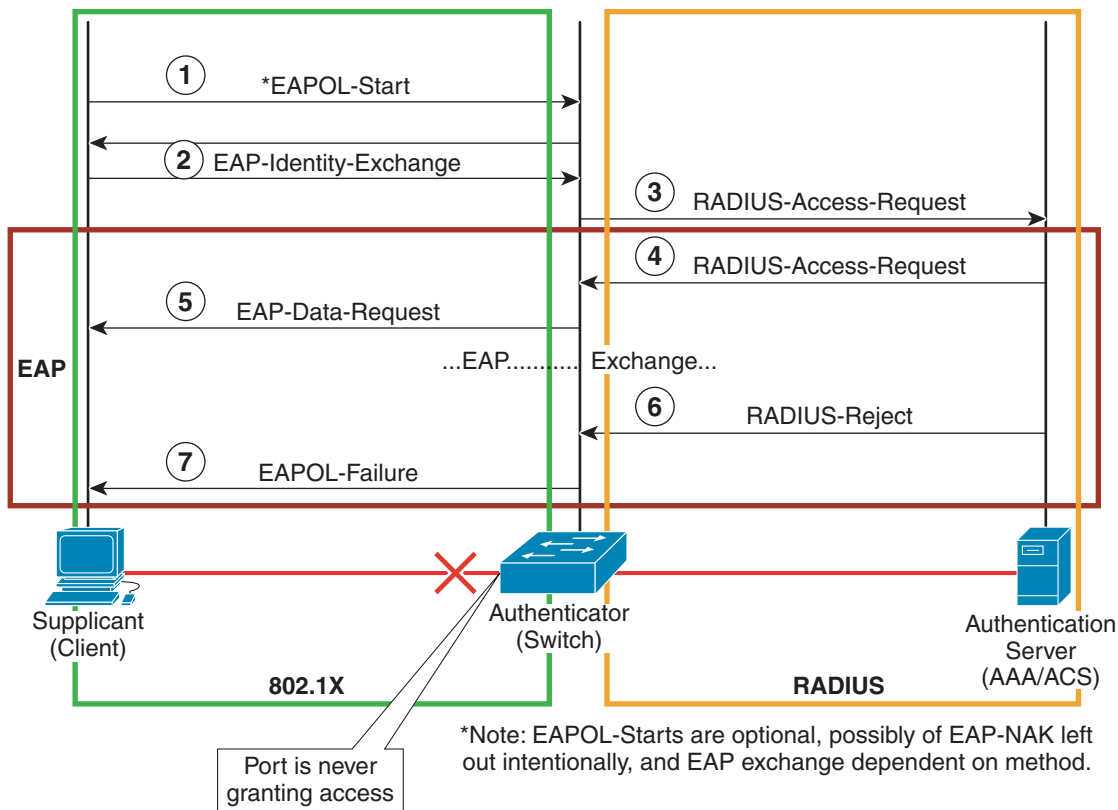


For more information on the design and implementation of the Cisco Unified Wireless Network based on the unified wireless architecture, which includes products operation with LWAPP, see the *Enterprise Mobility 3.0 Design Guide* at the following URL:
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob30dg/emob30dg-Book.html>.

802.1x Authentication Failure VLAN (Wired)

On a traditional 802.1x wired port, the switch does not provide access to the network until the supplicant connected to a port is authenticated, by verifying its identity information with an authentication server. There is no concept of an SSID for wired topologies today. For both media types, authentication failures work great in preventing rogue access to a network. This is a primary reason that some enterprises seek to enable 802.1x pervasively at the LAN edge. This default behavior is shown in [Figure 7](#).

Figure 7 Typical 802.1x Authentication Failures



However, for wired topologies, there must be a way to deal with the fact that an 802.1x-enabled guest or partner can plug into the enterprise LAN via wired ports.

The Auth-Fail-VLAN can be configured for an 802.1x port to provide limited services to clients. These clients are 802.1x-compliant and cannot access another VLAN because they fail the authentication process. A restricted VLAN allows users without valid credentials in an authentication server (typically, visitors to an enterprise) to access a limited set of services. The administrator can control the services available to the restricted VLAN.



Note

The same VLAN can be configured as both the Guest-VLAN and the Auth-Fail-VLAN when providing the same services to both types of users. The Guest-VLAN is discussed in detail in [Clientless-Based Authentication, page 28](#).

With the Auth-Fail-VLAN feature, you can configure the VLAN on a per-port basis and is enabled (by default) after three 802.1x authentication attempts. The port is then enabled and port forwarding is allowed in a VLAN where the supplicant can access the network. The Auth-Fail-VLAN can be configured for an 802.1x port to provide limited services to clients that are 802.1x-compliant and cannot access another VLAN because they fail the authentication process.

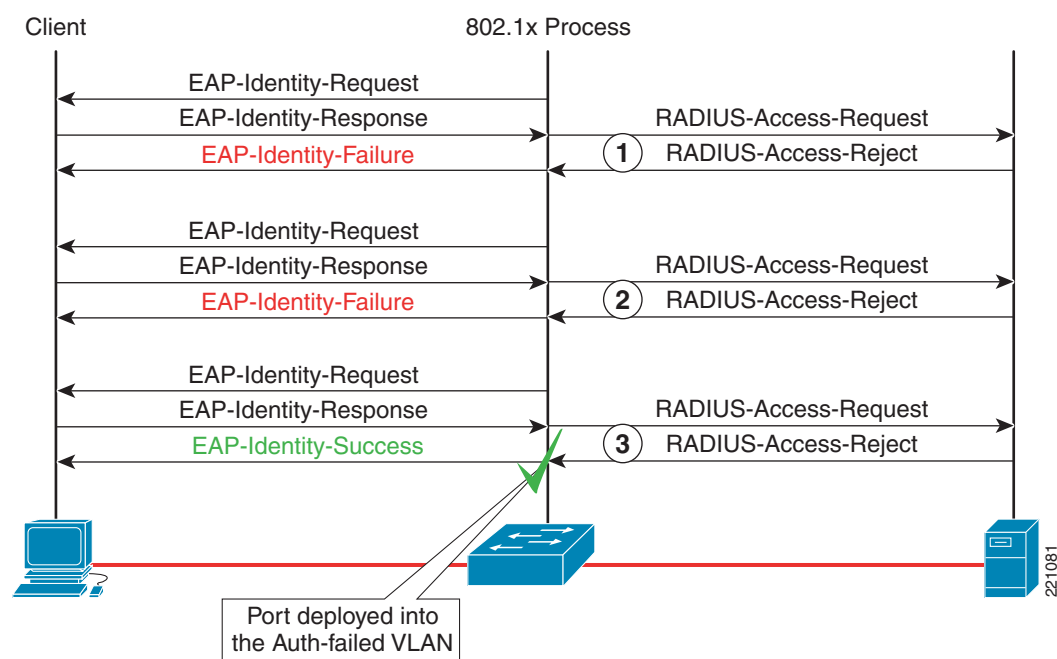
There may be several reasons why a user fails the 802.1x authentication. In addition, refer to an over-arching security policy to evaluate the deployment of the Auth-Fail-VLAN. The Auth-Fail-VLAN ultimately grants access to a device or end user that fails authentication. Although this authentication failure event can be differentiated from authorized devices, there is no chance to differentiate an 802.1x-enabled guest or partner who needs some form of network access from a hacker or illegitimate user. The same principle exists in wireless topologies. If 802.11 opens with no authentication provided

by a separate SSID, there is no way to keep an illegitimate user off the network. Wireless uses an SSID to differentiate the entire session. Wired can only use an actual authentication failure to attempt accomplish a similar task.

Auth-Fail-VLAN Operational Overview

The authenticator (access switch) counts the failed authentication attempts for a client. When this count exceeds the configured maximum number of authentication attempts (the default is 3), the port is deployed into the Auth-Fail-VLAN. After a port is moved to the Auth-Fail-VLAN, an EAP success message is sent to the client, as shown in Figure 8.

Figure 8 Auth-Fail-VLAN Operation



The behavior of the client from this point is affected by three variables: the specific 802.1x supplicant that is used (Microsoft or CSSC), the EAP protocol, and the specific switch platform acting as authenticator. These interdependencies are clarified in the following sections of this paper.

Any active VLAN can be configured as Auth-Fail-VLAN with the exception of an RSPAN VLAN, or a voice VLAN (VVID). In addition, the Auth-Fail-VLAN feature is not supported on internal VLANs (routed ports) or trunk ports; it is supported only on access ports.



Note

Authentication has to actually fail for this process to complete. Any sort of timeout condition (supplicant/authenticator or authenticator/authentication server) is not addressed by the Auth-Fail-VLAN feature.

Auth-Fail-VLAN Configuration

Following are configuration samples that enable the Auth-Failed VLAN feature on IOS and CatOS authenticators:

- IOS:

```
interface FastEthernet0/1
switchport access vlan 2
switchport mode access
dot1x pae authenticator
dot1x port-control auto
dot1x auth-fail vlan 5
spanning-tree portfast
spanning-tree bpduguard enable
```

- CatOS:

```
set vlan 2 2/1
set port dot1x 2/1 port-control auto
set port dot1x 2/1 auth-fail-vlan 5
set spantree portfast 2/1 enable
set spantree bpdu-guard 2/1 enable
```

**Note**

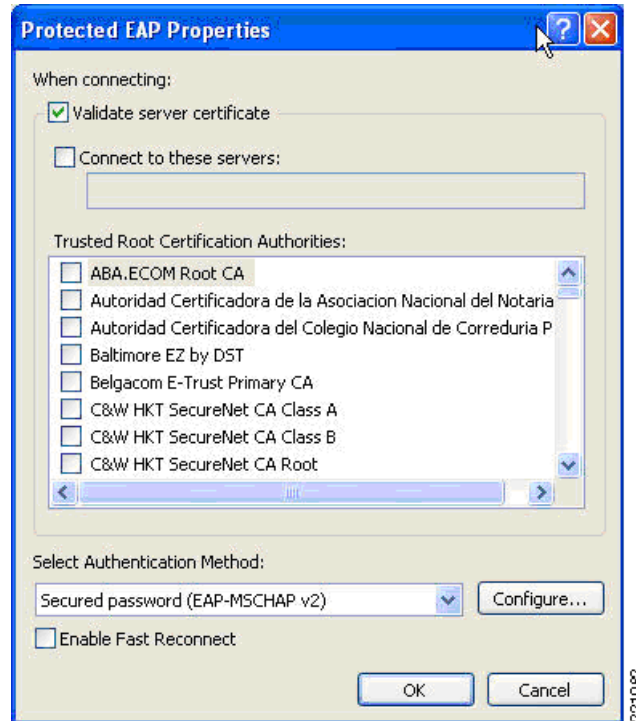
Although not verified as part of the solution, the number of failures to deploy a port into the Auth-Fail-VLAN can be configured in IOS via the **dot1x auth-fail max-attempts** command. The default value for this parameter is 3, and 3 is the hard-coded parameter in CatOS.

Auth-Fail-VLAN Verification

Reasons why a user may fail the 802.1x authentication, causing the switch port to be deployed in the Auth-Fail-VLAN, include the following:

- A tunneled EAP method (PEAP or EAP-FAST) is used for authentication, and the supplicant is configured for validating the server certificate. In this case, there are the following two scenarios:
 - Most supplicants are configured to use the Certificate Trust List (CTL) to validate a server certificate with a tunneled method. In this case, the authentication fails, unless the certificate sent by the RADIUS server is trusted by the supplicant. This means the supplicant trusts the intermediary that has signed or issued the server certificate. An example of a pre-populated CTL is shown in [Figure 9](#). This is the Trusted Root Certification Authorities list available with Microsoft supplicants.

Figure 9 Microsoft Supplicant CTL Example



In many situations, including guest access scenarios, the certificate authority (CA) that provided the certificate sent by the RADIUS server is most likely not part of the client CTL (especially in deployments where a private CA is used). As a consequence, the TLS handshake tried in the tunnel establishment phase fails. The client denies the authentication attempt by being unable to verify the backend server to establish an SSL tunnel between client and server. On ACS, the message appears as indicated in Figure 10.

Figure 10 Authentication Failure From Client

Date ↓	Time	Message-Type	User-Name	Group-Name	Caller-ID	Network Access Profile Name	Authen-Failure-Code	Au Fai C
02/15/2006	11:57:25	Authen failed	GUEST-MAX\max	Default Group	00-0D-9D-89-74-E8	(Default)	EAP-TLS or PEAP authentication failed during SSL handshake	..

Note that by default, the Microsoft Supplicant (WZCSVC) and the Cisco Secure Services Client (CSSC) validates the server certificate by default when tunneled methods are configured. Older versions of the Meetinghouse AEGIS client did not trust a server certificate by default.

- Alternatively, a non-default configuration for WZSVC with Windows XP SP2 enables the supplicant to conditionally validate the server certificate. This way, the end user is presented a popup window to inform the user to accept the certificate (similarly to what happens on HTTPS transactions). An example of this capability offered WZCSVC is shown in Figure 11.

Figure 11 *Conditional Trust for a Server Certificate*

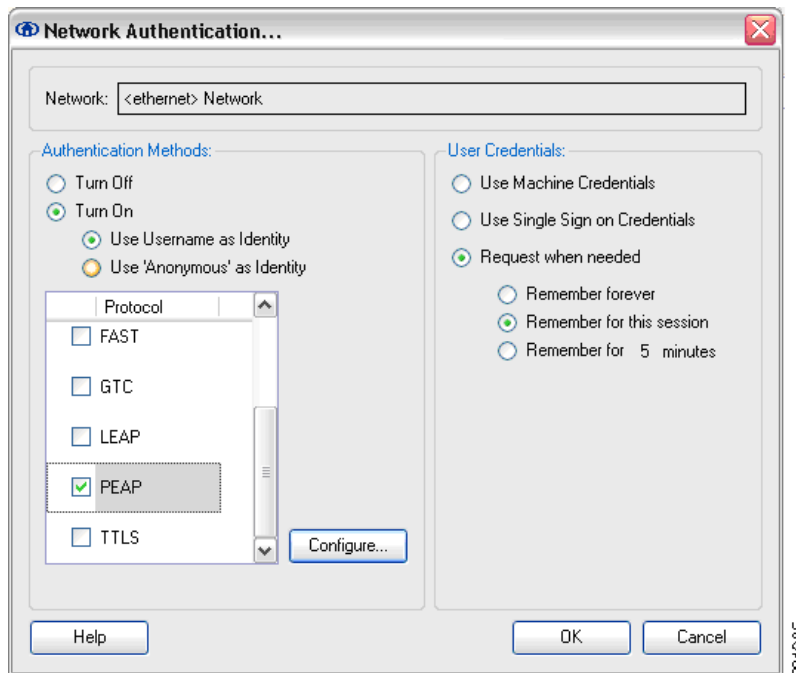


Note This functionality of conditional trust is not available with CSSC.

When the popup is displayed, an end user can manually accept the certificate sent by the authentication server, and avoiding failing the authentication because of SSL handshake. The authentication at this point may still fail for one of the two reasons discussed in bullets two and three below.

- The client is sending wrong credentials to the RADIUS server (or backend authentication server, as Active Directory). Note that most commonly the same credentials used for Windows are also used for 802.1x authentication. This is the default behavior for WCZSVC. For CSSC, the client can be configured to operate via Windows credentials or not, as shown in [Figure 12](#).

Figure 12 *CSSC Options for Authentication*

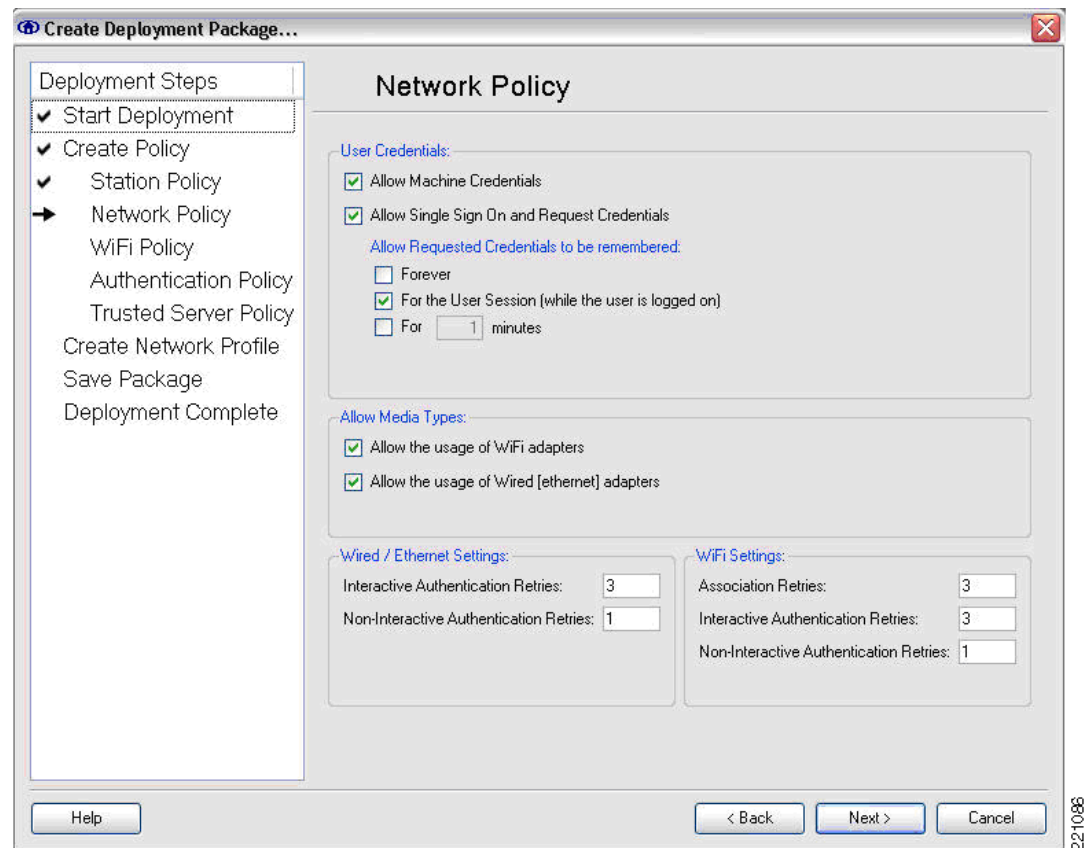


The supplicant is configured to try an EAP type that is not supported by the RADIUS server.

CSSC Behavior

WZCSVC implements PEAPv0, whereas the CSSC client implements PEAPv1 (“Cisco PEAP”). This subtle difference does not change the overall behavior, as shown below. The behavior verified with CSSC for an authentication failure was far different than WZCSVC. CSSC does not work with the Auth-Fail-VLAN as of version 4.0(5). CSSC sends an EAPOL-Logoff frame as a result of an authentication failure event. Two default values in a standard deployment of CSSC 4.0(5) pertain to this: “Interactive Authentication Retries” and “Non-Interactive Authentication Retries” failures, which are shown in Figure 12.

Figure 13 CSSC Deployment Package



The screen shown in Figure 13 is from a CSSC deployment package. Creating a deployment package and changing the default mode of operation is the only way to change these values. “Interactive” means when the supplicant is not using any cached credentials and prompts the end user. “Non-interactive” means when the supplicant is using cached credentials for Single Sign On (SSO), or for machine authentication where the same Windows credentials are also sent to the RADIUS server for 802.1x authentication.

The values demonstrated above are when EAPOL-Logoffs are sent out. The value of these timers represents the number of EAP Failure messages that the supplicant can receive before sending an EAPOL-Logoff message to the switch. When considering the Auth-Fail-VLAN, it is important to correlate the value of the Auth-Fail-Max-attempts parameter that is configured on each switch port (using the **dot1x auth-fail max-attempts** command). This parameter dictates how many RADIUS Reject messages the switch must receive from the RADIUS server before sending an EAP Success to the supplicant. Its default value is 3, which essentially causes the switch to behave as shown in Figure 8.

When receiving the third RADIUS Reject message from the AAA server, an EAP Success is sent to the client. As a consequence, the supplicant never sees more than two EAP Failure messages from the authenticator, which leads to the following two scenarios:

- Interactive retries—The switch port is deployed into the Auth-Fail-VLAN, because the client never sees more than two EAP Failure messages, and thus does not send the EAPOL Logoff.
- Non-interactive retries—After the client receives the first EAP Failure, it sends an EAPOL Logoff. As a consequence, the switch clears the session state for that client on that port. This causes the creation of an endless loop where the supplicant fails authentication and consequently sends an EAPOL Logoff, as shown in the trace in [Figure 14](#).

Figure 14 EAPOL-Logoff after Failure

No. -	Time	Source	Destination	Protocol	Info
1	0.000000	Ibm_ac:53:75	Spanning-tree-(for-	EAPOL	Start
2	0.215226	Cisco_be:43:03	Spanning-tree-(for-	EAP	Failure
3	0.215261	Cisco_be:43:03	Spanning-tree-(for-	EAP	Request, Identity [RFC3748]
5	0.449211	Ibm_ac:53:75	Spanning-tree-(for-	EAP	Response, Identity [RFC3748]
6	0.449619	Cisco_be:43:03	Spanning-tree-(for-	EAP	Request, Identity [RFC3748]
7	0.637029	Ibm_ac:53:75	Spanning-tree-(for-	EAP	Response, Identity [RFC3748]
8	0.647441	Cisco_be:43:03	Spanning-tree-(for-	EAP	Request, PEAP [Palekar]
10	0.778554	Ibm_ac:53:75	Spanning-tree-(for-	TLS	Client Hello
11	0.783844	Cisco_be:43:03	Spanning-tree-(for-	TLS	Server Hello, Certificate, Server Hello Done
12	0.798214	Ibm_ac:53:75	Spanning-tree-(for-	EAP	Response, PEAP [Palekar]
13	0.803450	Cisco_be:43:03	Spanning-tree-(for-	TLS	Server Hello, Certificate, Server Hello Done
14	1.001791	Ibm_ac:53:75	Spanning-tree-(for-	TLS	Client Key Exchange, Change Cipher Spec, Encrypted Handshake Message
15	1.016427	Cisco_be:43:03	Spanning-tree-(for-	TLS	Change Cipher Spec, Encrypted Handshake Message
16	1.075714	Ibm_ac:53:75	Spanning-tree-(for-	EAP	Response, PEAP [Palekar]
17	1.078861	Cisco_be:43:03	Spanning-tree-(for-	TLS	Application Data
19	1.220293	Ibm_ac:53:75	Spanning-tree-(for-	TLS	Application Data, Application Data
20	1.223765	Cisco_be:43:03	Spanning-tree-(for-	TLS	Application Data
21	1.368629	Ibm_ac:53:75	Spanning-tree-(for-	TLS	Application Data, Application Data
22	1.372462	Cisco_be:43:03	Spanning-tree-(for-	TLS	Application Data
23	1.495800	Ibm_ac:53:75	Spanning-tree-(for-	TLS	Application Data, Application Data
24	1.506419	Cisco_be:43:03	Spanning-tree-(for-	TLS	Application Data
25	1.541565	Ibm_ac:53:75	Spanning-tree-(for-	TLS	Application Data, Application Data
26	1.544968	Cisco_be:43:03	Spanning-tree-(for-	EAP	Failure
27	1.545100	Cisco_be:43:03	Spanning-tree-(for-	EAP	Request, Identity [RFC3748]
28	1.758133	Ibm_ac:53:75	Spanning-tree-(for-	EAPOL	Logoff
29	1.758565	Cisco_be:43:03	Spanning-tree-(for-	EAP	Failure
30	1.758618	Cisco_be:43:03	Spanning-tree-(for-	EAP	Request, Identity [RFC3748]
36	4.317971	Ibm_ac:53:75	Spanning-tree-(for-	EAPOL	Start
37	4.318279	Cisco_be:43:03	Spanning-tree-(for-	EAP	Request, Identity [RFC3748]
38	4.595061	Ibm_ac:53:75	Spanning-tree-(for-	EAP	Response, Identity [RFC3748]
39	4.602809	Cisco_be:43:03	Spanning-tree-(for-	EAP	Request, PEAP [Palekar]
41	4.737771	Ibm_ac:53:75	Spanning-tree-(for-	TLS	Client Hello
42	4.745026	Cisco_be:43:03	Spanning-tree-(for-	TLS	Server Hello, Certificate, Server Hello Done
43	4.775698	Ibm_ac:53:75	Spanning-tree-(for-	EAP	Response, PEAP [Palekar]
44	4.781179	Cisco_be:43:03	Spanning-tree-(for-	TLS	Server Hello, Certificate, Server Hello Done
45	5.028536	Ibm_ac:53:75	Spanning-tree-(for-	TLS	Client Key Exchange, Change Cipher Spec, Encrypted Handshake Message
46	5.044233	Cisco_be:43:03	Spanning-tree-(for-	TLS	Change Cipher Spec, Encrypted Handshake Message
47	5.069848	Ibm_ac:53:75	Spanning-tree-(for-	EAP	Response, PEAP [Palekar]
48	5.073029	Cisco_be:43:03	Spanning-tree-(for-	TLS	Application Data
49	5.200130	Ibm_ac:53:75	Spanning-tree-(for-	TLS	Application Data, Application Data

221087



Note

The behavior indicated above has been observed on Catalyst 6500 platforms running CatOS software. The tested Cisco IOS releases for Catalyst 3750 and 4500 are affected by a bug (CSCsg75620) that allows the deployment of the switch port into the Auth-Fail-VLAN, regardless of the Logoff. For IOS, the switch does not reset the Auth-Fail-Max-attempts parameter after receiving the EAPOL Logoff. Cisco recommends for any Auth-Fail-VLAN deployment that these values both be set to at least 4 to allow successful deployment by reconfiguring the default values through the deployment package demonstrated above. Recommended values for both the interactive and non-interactive values should be set to at least 4.

By default, when authentication fails without the Auth-Fail-VLAN, sending an EAPOL-Logoff does not matter, because the switch has a default held timer of 60 seconds anyway. However, the problem with the transmission of an EAPOL-Logoff frame in an Auth-Fail-VLAN scenario is that upon the first failure, there is practically no held timer, and the switch tries to authenticate the session again immediately. Furthermore, the EAPOL-Logoff frame sent by the supplicant is now processed normally by the switch. For IOS, this also resets the Auth-Fail count toward the deployment of the Auth-Fail-VLAN back to zero, and this process repeats again immediately. This is now a potential DoS attack on RADIUS. A single machine may be sitting idle and failing authentication repeatedly with no held state to protect or delay the event from occurring. However, if you modify the default deployment

package, as described above, and increase the number of default retries for both the parameters above to something greater than 3, it can allow the Auth-Fail-VLAN to work. In addition, the default values of these are being changed to a number greater than three in the 4.1 release of CSSC code.

CSSC with Catalyst 6500

The sequence of events when a client fails the 802.1x authentication with the RADIUS server is displayed in the sniffer trace shown in [Figure 15](#).

Figure 15 Auth-Fail-VLAN Deployment

No.	Time	Source	Destination	Protocol	Info
1	0.000000	Cisco1c:ec:ff	Spanning-tree-(for	EAP	Request, Identity [RFC3748]
2	0.787626	10.1.13.10	Spanning-tree-(for	EAPOL	Start
3	0.788315	Cisco1c:ec:ff	Spanning-tree-(for	EAP	Request, Identity [RFC3748]
4	1.116104	10.1.13.10	Spanning-tree-(for	EAP	Response, Identity [RFC3748]
5	1.120029	Cisco1c:ec:ff	Spanning-tree-(for	EAP	Request, PEAP [Palekar]
6	1.283298	10.1.13.10	Spanning-tree-(for	TLS	Client Hello
7	1.286994	Cisco1c:ec:ff	Spanning-tree-(for	TLS	Server Hello, Certificate, Server Hello Done
8	1.506049	10.1.13.10	Spanning-tree-(for	TLS	Alert (Level: Fatal, Description: Bad Certificate)
9	1.508371	Cisco1c:ec:ff	Spanning-tree-(for	EAP	Failure
10	1.511617	Cisco1c:ec:ff	Spanning-tree-(for	EAP	Request, Identity [RFC3748]
11	2.611668	0.0.0.0	255.255.255.255	DHCP	DHCP Discover - Transaction ID 0x4d0599b0
12	5.708307	10.1.13.10	Spanning-tree-(for	EAP	Response, Identity [RFC3748]
13	5.711795	Cisco1c:ec:ff	Spanning-tree-(for	EAP	Request, PEAP [Palekar]
14	5.900472	10.1.13.10	Spanning-tree-(for	TLS	Client Hello
15	5.904022	Cisco1c:ec:ff	Spanning-tree-(for	TLS	Server Hello, Certificate, Server Hello Done
16	6.099383	10.1.13.10	Spanning-tree-(for	TLS	Alert (Level: Fatal, Description: Bad Certificate)
17	6.120882	Cisco1c:ec:ff	Spanning-tree-(for	EAP	Failure
18	6.124147	Cisco1c:ec:ff	Spanning-tree-(for	EAP	Request, Identity [RFC3748]
19	6.608477	0.0.0.0	255.255.255.255	DHCP	DHCP Discover - Transaction ID 0x4d0599b0
20	11.171984	10.1.13.10	Spanning-tree-(for	EAP	Response, Identity [RFC3748]
21	11.175570	Cisco1c:ec:ff	Spanning-tree-(for	EAP	Request, PEAP [Palekar]
22	11.340172	10.1.13.10	Spanning-tree-(for	TLS	Client Hello
23	11.343581	Cisco1c:ec:ff	Spanning-tree-(for	TLS	Server Hello, Certificate, Server Hello Done
24	11.606658	10.1.13.10	Spanning-tree-(for	TLS	Alert (Level: Fatal, Description: Bad Certificate)
25	11.631460	Cisco1c:ec:ff	Spanning-tree-(for	EAP	Success
56	15.606880	0.0.0.0	255.255.255.255	DHCP	DHCP Discover - Transaction ID 0x4d0599b0
57	15.609204	10.1.13.3	10.1.13.10	DHCP	DHCP Offer - Transaction ID 0x4d0599b0
58	15.609478	0.0.0.0	255.255.255.255	DHCP	DHCP Request - Transaction ID 0x4d0599b0
59	15.613644	10.1.13.3	10.1.13.10	DHCP	DHCP ACK - Transaction ID 0x4d0599b0
316	41.698191	10.1.13.10	Spanning-tree-(for	EAPOL	Start

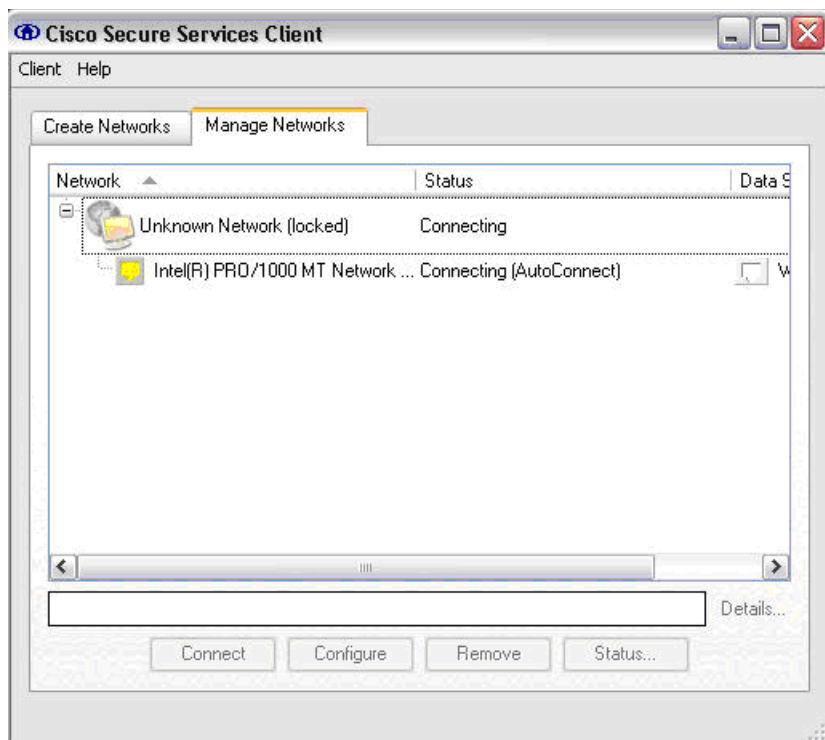
221088

The client receives two EAP Failures (frames 9 and 17), but after the third authentication failure, the switch is actually sending an EAP Success (frame 25).

At this point, the switch port is open and deployed to the Auth-Failed VLAN, so the client can pull an IP address from the DHCP server (frames 56–59) and achieve network connectivity.

From a user experience perspective, the supplicant appears to be still trying to connect, as displayed in [Figure 16](#).


Figure 16 Supplicant GUI Upon Auth-Fail-VLAN Deployment



The reason for this behavior is that the supplicant still knows that it failed the inner authentication (inside the SSL tunnel), despite the fact that the switch sent an EAP Success outside of the tunnel.

After 30 seconds from receiving the first EAP-Success, the supplicant sends a new EAPOL-Start (frame 316) to try again to authenticate. The authenticator does not honor this frame, but simply discards it (because the port is in the Auth-Failed VLAN), and at this point the supplicant “gives up” and enters the connected and authenticated state, as shown in [Figure 17](#).

Figure 17 Supplicant GUI After Timeout of Trying to Authenticate Again



Note

The last consideration above is just a cosmetic one. From a connectivity perspective, the client can communicate with the network as soon as the switch port is deployed into the Auth-Fail-VLAN, because the supplicant does not exhibit the notion of a controlled port.

CSSC with Catalyst 3750 and Catalyst 4500

The sequence of events when the supplicant fails the 802.1x authentication with the AAA server is displayed in the sniffer trace shown in [Figure 18](#).

Figure 18 Switch Honors EAPOL-Starts after Supplicant is Confused

No.	Time	Source	Destination	Protocol	Info
2	0.386204	Cisco_6b:52:04	Spanning-tree-(for	EAP	Request, Identity [RFC3748]
7	1.461527	10.1.13.10	Spanning-tree-(for	EAPOL	Start
8	2.478135	0.0.0.0	255.255.255.255	DHCP	DHCP Discover - Transaction ID 0x36da32a1
9	2.494528	Cisco_6b:52:04	Spanning-tree-(for	EAP	Request, Identity [RFC3748]
10	2.711313	10.1.13.10	Spanning-tree-(for	EAP	Response, Identity [RFC3748]
11	2.723494	Cisco_6b:52:04	Spanning-tree-(for	EAP	Request, PEAP [Palekar]
12	2.878860	10.1.13.10	Spanning-tree-(for	EAP	Client Hello
13	2.888482	Cisco_6b:52:04	Spanning-tree-(for	EAP	Request, PEAP [Palekar]
14	2.893245	10.1.13.10	Spanning-tree-(for	EAP	Response, PEAP [Palekar]
15	2.900853	Cisco_6b:52:04	Spanning-tree-(for	TLS	Server Hello, Certificate, Server Hello Done
16	3.115976	10.1.13.10	Spanning-tree-(for	TLS	Alert (Level: Fatal, Description: Bad Certificate)
17	3.128037	Cisco_6b:52:04	Spanning-tree-(for	EAP	Failure
19	4.145625	Cisco_6b:52:04	Spanning-tree-(for	EAP	Request, Identity [RFC3748]
21	5.476083	0.0.0.0	255.255.255.255	DHCP	DHCP Discover - Transaction ID 0x36da32a1
23	7.942596	10.1.13.10	Spanning-tree-(for	EAP	Response, Identity [RFC3748]
24	7.956184	Cisco_6b:52:04	Spanning-tree-(for	EAP	Request, PEAP [Palekar]
25	8.125818	10.1.13.10	Spanning-tree-(for	TLS	Client Hello
26	8.135476	Cisco_6b:52:04	Spanning-tree-(for	EAP	Request, PEAP [Palekar]
27	8.143841	10.1.13.10	Spanning-tree-(for	EAP	Response, PEAP [Palekar]
28	8.151483	Cisco_6b:52:04	Spanning-tree-(for	TLS	Server Hello, Certificate, Server Hello Done
29	8.372899	10.1.13.10	Spanning-tree-(for	TLS	Alert (Level: Fatal, Description: Bad Certificate)
30	8.385654	Cisco_6b:52:04	Spanning-tree-(for	EAP	Failure
32	9.415358	Cisco_6b:52:04	Spanning-tree-(for	EAP	Request, Identity [RFC3748]
35	11.865307	10.1.13.10	Spanning-tree-(for	EAP	Response, Identity [RFC3748]
36	11.875611	Cisco_6b:52:04	Spanning-tree-(for	EAP	Request, PEAP [Palekar]
37	12.051290	10.1.13.10	Spanning-tree-(for	EAP	Client Hello
38	12.068701	Cisco_6b:52:04	Spanning-tree-(for	EAP	Request, PEAP [Palekar]
39	12.083294	10.1.13.10	Spanning-tree-(for	EAP	Response, PEAP [Palekar]
40	12.090996	Cisco_6b:52:04	Spanning-tree-(for	TLS	Server Hello, Certificate, Server Hello Done
41	12.257901	10.1.13.10	Spanning-tree-(for	TLS	Alert (Level: Fatal, Description: Bad Certificate)
42	12.272526	Cisco_6b:52:04	Spanning-tree-(for	EAP	Success
44	13.479237	0.0.0.0	255.255.255.255	DHCP	DHCP Discover - Transaction ID 0x36da32a1
45	13.484564	10.1.86.1	10.1.86.100	DHCP	DHCP Offer - Transaction ID 0x36da32a1
46	13.484845	0.0.0.0	255.255.255.255	DHCP	DHCP Request - Transaction ID 0x36da32a1
47	13.497825	10.1.86.1	10.1.86.100	DHCP	DHCP ACK - Transaction ID 0x36da32a1
162	42.383709	10.1.13.10	Spanning-tree-(for	EAPOL	Start
165	43.417711	Cisco_6b:52:04	Spanning-tree-(for	EAP	Request, Identity [RFC3748]
167	43.581046	10.1.13.10	Spanning-tree-(for	EAP	Response, Identity [RFC3748]
168	43.590543	Cisco_6b:52:04	Spanning-tree-(for	EAP	Request, PEAP [Palekar]
171	43.799851	10.1.13.10	Spanning-tree-(for	TLS	Client Hello
172	43.809412	Cisco_6b:52:04	Spanning-tree-(for	EAP	Request, PEAP [Palekar]
173	43.832679	10.1.13.10	Spanning-tree-(for	EAP	Response, PEAP [Palekar]
174	43.840362	Cisco_6b:52:04	Spanning-tree-(for	TLS	Server Hello, Certificate, Server Hello Done
175	44.024243	10.1.13.10	Spanning-tree-(for	TLS	Alert (Level: Fatal, Description: Bad Certificate)
176	44.040063	Cisco_6b:52:04	Spanning-tree-(for	EAP	Success
291	74.185003	10.1.13.10	Spanning-tree-(for	EAPOL	Start
295	75.213289	Cisco_6b:52:04	Spanning-tree-(for	EAP	Request, Identity [RFC3748]
296	75.416013	10.1.13.10	Spanning-tree-(for	EAP	Response, Identity [RFC3748]
297	75.425750	Cisco_6b:52:04	Spanning-tree-(for	EAP	Request, PEAP [Palekar]
299	75.601757	10.1.13.10	Spanning-tree-(for	TLS	Client Hello
300	75.617479	Cisco_6b:52:04	Spanning-tree-(for	EAP	Request, PEAP [Palekar]
301	75.636136	10.1.13.10	Spanning-tree-(for	EAP	Response, PEAP [Palekar]
302	75.644061	Cisco_6b:52:04	Spanning-tree-(for	TLS	Server Hello, Certificate, Server Hello Done
307	75.860103	10.1.13.10	Spanning-tree-(for	TLS	Alert (Level: Fatal, Description: Bad Certificate)
308	75.874199	Cisco_6b:52:04	Spanning-tree-(for	EAP	Success


22/1091

The client receives two EAP Failures (frames 17 and 30), but after the third authentication failure, the switch is actually sending an EAP Success (frame 42).

At this point, the switch port is open and deployed to the Auth-Fail-VLAN, so the client can pull an IP address from the DHCP server (frames 44–47) and achieve network connectivity.

From a user experience perspective, the supplicant appears to be still trying to connect, as displayed in [Figure 19](#).

Figure 19 Supplicant GUI Upon Auth-Fail-VLAN Deployment



The reason for this behavior is that the supplicant still knows that it failed the inner authentication (inside the SSL tunnel), despite the fact that the switch sent an EAP Success outside of the tunnel.

After 30 seconds from receiving the first EAP-Success, the supplicant sends a new EAPOL-Start (frame 162) to try again to authenticate. Differently from what is observed with the Catalyst 6500 platform, the Catalyst 3750 and 4500 actually honor the EAPOL Start message and initiate RADIUS communication with the RADIUS server (frames 165–307). The authentication keeps failing, so the authenticator keeps manufacturing indefinitely EAP-Success frames to be sent to the client (frame 308). There are two consequences for this behavior:

- The client performs a sort of DoS attack to the RADIUS server, forcing a re-authentication exchange every 30 seconds, as shown in [Figure 20](#).
- The GUI of the supplicant indefinitely remains in the “Connecting” state, despite the fact that the client has actually obtained network connectivity.

Figure 20 DoS Attack on RADIUS via Auth-Fail-VLAN

Date ↓	Time	Message-Type	User-Name	Group-Name	Caller-ID	Network Access Profile Name	Authen-Failure-Code	Auth-Failu-Cor
02/20/2006	11:22:52	Authen failed	GUEST-MAX\max	Default Group	00-0D-9D-89-74-E8	(Default)	CS user unknown	..
02/20/2006	11:22:21	Authen failed	GUEST-MAX\max	Default Group	00-0D-9D-89-74-E8	(Default)	CS user unknown	..
02/20/2006	11:21:50	Authen failed	GUEST-MAX\max	Default Group	00-0D-9D-89-74-E8	(Default)	CS user unknown	..
02/20/2006	11:21:16	Authen failed	GUEST-MAX\max	Default Group	00-0D-9D-89-74-E8	(Default)	CS user unknown	..

221093

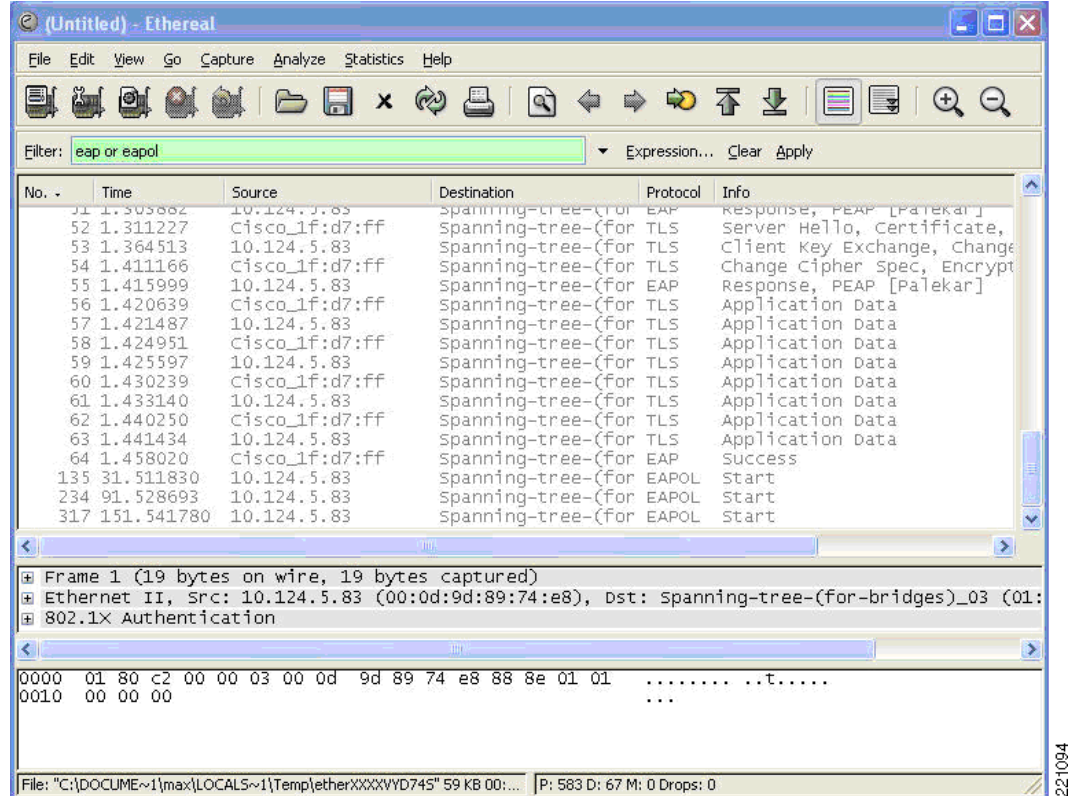
This is broken behavior (CSCsd42505). This will be developed in release 12.2(35)SE and 12.2(31)SGA, and will allow the 3750 and 4500 to behave like the 6500, silently discarding EAPOL Start messages received from the client when the switch port is deployed in the Auth-Fail-VLAN.

The impact of this behavior when deploying a large numbers of stations into the Auth-failed VLAN should be considered.

WZCSVC Behavior

The behavior of the WZCSVC supplicant that ships on Windows XP platforms is almost identical to the CSSC supplicant. (See Figure 21.) The behavior is also identical to that observed on CatOS with CSSC clients, with the only difference that the supplicant sends EAPOL-Start every 60 seconds (and not 30) after the switch port is deployed into the Auth-Failed VLAN (frames 135, 234, and 317 below). In addition, before giving up, it sends three EAPOL-Start messages (that are discarded by the switch), whereas only one EAPOL-Start was sent by CSSC.

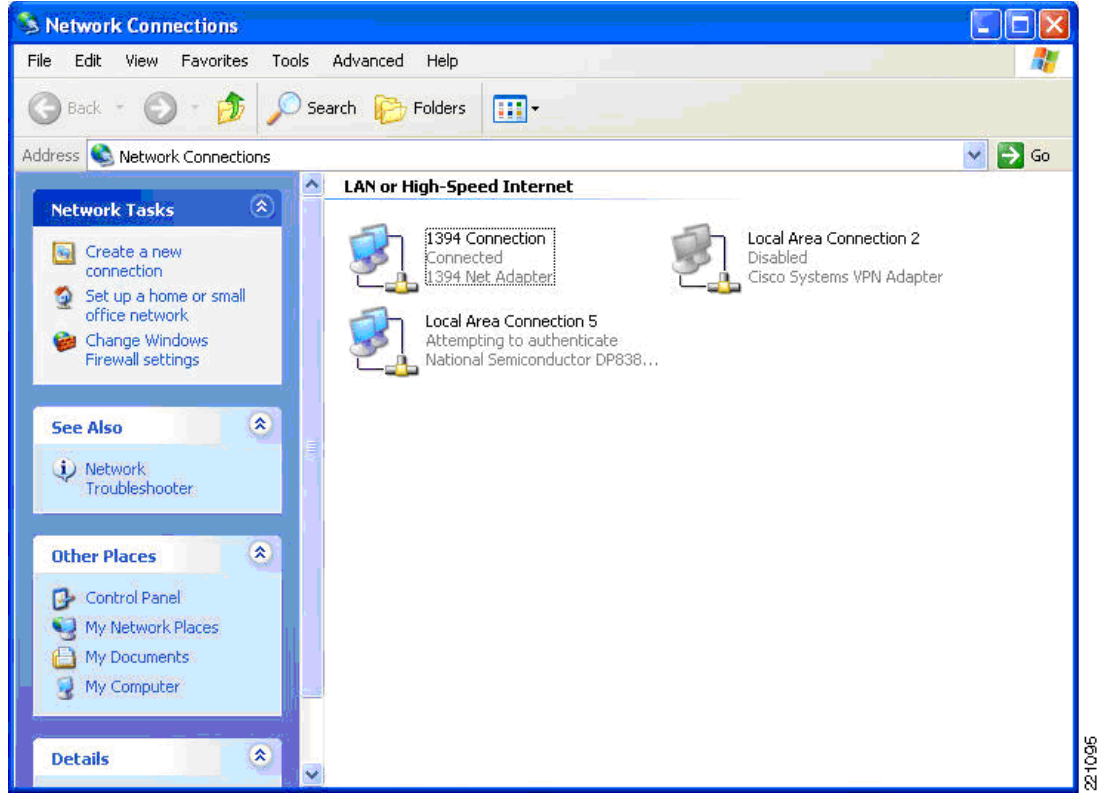
Figure 21 Auth-Fail-VLAN Deployment with Confused WZCSVC Supplicant



After sending the third EAPOL-start without receiving response from the authenticator, the supplicants give up trying to authenticate and enter the “open” state. (See [Figure 22](#).) WZCSVC has no control over the IP stack.

There is no direct indication to the end user that after sending the third EAPOL-Start, the supplicant is actually giving up the authentication attempts, because the network interface still says “Attempting to authenticate” or ultimately “connected”.

Figure 22 Confused WZSCVC Supplicant



However, the client would be able to pull a network address from the pool corresponding to the auth-failed VLAN and to gain network connectivity.

In addition, if Active Directory (AD) is used as a backend database with the WZCSVC supplicant and the user is not found in AD, or the password sent is wrong, the supplicant may get stuck during the communication with the RADIUS server. The consequence is that the attempt does not even fail, preventing the deployment of the switch port into the Auth-Fail-VLAN. This implies that no network access at all can be provided in this scenario.

It was noticed that this supplicant does not return a failure response to the failure message of the server. For ACS, this prevents the EAP state machine from getting to the next stage of sending the final EAP-MSCHAP failure code and a RADIUS-Reject. As such, ACS does not send a RADIUS-Reject to the switch immediately. Nonetheless, it should result in a failure. Operationally, the first two authentication failures of each conversation result in a RADIUS challenge as opposed to a reject. Then, a reject is sent on the third unsuccessful attempt. With respect to the Auth-Fail-VLAN, this means an end user may actually have to fail nine times before the Auth-Fail-VLAN activates, because it is enabled upon the receipt of RADIUS-Rejects. In addition, note that the WZCSVC supplicant does not display meaningful messages such as "Account Expired", "Bad Logon Hours", and so on. The only failure scenarios that work with this supplicant are a bad password (where the user is otherwise known) or an expired password. This behavior occurs only with PEAP for machines and users who either blindly trust a server certificate from ACS, or who conditionally trust the server certificate and the credentials have actually been removed from a domain.

Auth-Fail-VLAN Summary and Recommendations

In the wired media for 802.1x, there exists a need to provide access to devices that fail authentication. This is the Auth-Fail-VLAN. This can also serve as a method to provide 802.1x-enabled guests and partners network access in a network virtualization architecture.

This is not a problem for wireless, because of culture, the secondary nature of the wireless media, and 802.11 station authentication. With 802.11 open, no authentication, and a broadcast SSID, the guest or partner problem can be solved easily without the need to attempt to provide access to devices that actually fail 802.1x authentication. It is typically not a problem for IPsec, PPP, or dial-up environments either. It is a problem for the wired media at the enterprise LAN edge, however. The Auth-Fail-VLAN can serve as a way to deal with the wired 802.1x-enabled entity that is unknown to the hosting enterprise.

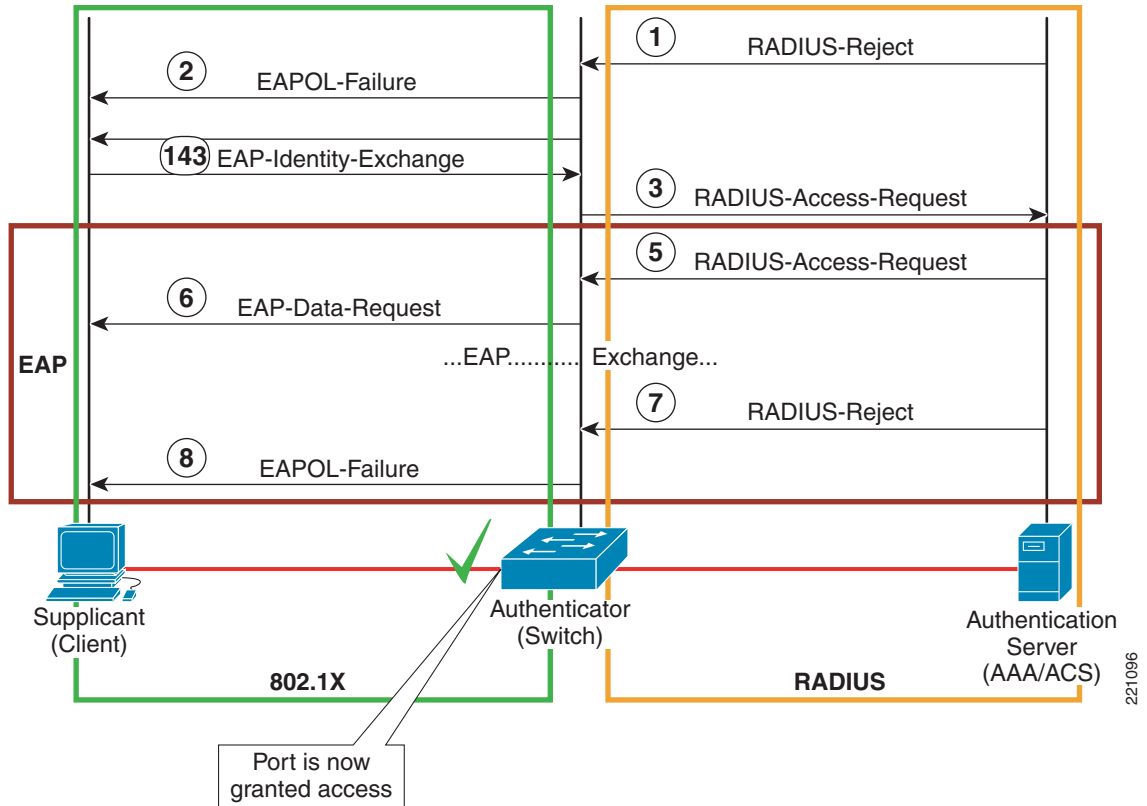
However, there are some architectural problems with the Auth-Fail-VLAN as verified above. EAP is between a supplicant and an EAP-Server in the 802.1x framework. Although not precluded by the EAP architecture, Cisco switches today are not EAP-Servers, but authenticators only. Primarily, this means that they serve as an EAP transport via 802.1x and RADIUS, and rely on an authentication server to be an EAP-Server. Because switches operate in pass-through mode for EAP, attempting to modify the result of the authentication conversation from the authenticator alone can be challenging. Specifically, this is the operation of the authenticator sending an EAP-Success when it would otherwise send an EAPOL-Failure. The first issue is that the switch does not have visibility into what is actually going on, nor should it. Fundamentally, the switch is launching a man-in-the-middle attack on EAP.

For EAP methods that support mutual authentication, the supplicant authenticates the network as well as the network authenticating the supplicant. For these methods, if the authentication fails, it is reasonable to assume that the supplicant will not authorize its port to allow traffic anyway, nor should it. Testing has proven that the CSSC and WZCSVC supplicants do not exhibit the notion of a controlled port, however, although this could change in the future. For tunneled EAP methods, if the supplicant receives a failure inside a tunnel (that a switch can do nothing about) and a clear channel success as demonstrated above, this becomes a conflicting message. In the future, supplicants could be configured to treat this behavior as an invalid authentication attempt that should not be trusted or allowed.

Verification testing has proven the current operation to work with most EAP methods, given these constraints. However, as demonstrated above, this does not work completely for tunneled methods such as PEAP or EAP-FAST. The verification demonstrates that in this scenario, supplicants get confused, and try to begin authentication over again, which results in a timeout condition.

Based on the series of issues discovered as part of the system verification, the code revisions tested on switches should not be recommended as part of a network virtualization solution. However, switches have already addressed most of the architectural issues with the behavior tested through CSCse71105 and CSCse58950. This changes the anatomy of the Auth-Fail-VLAN. There is no longer a transmission of an EAP-Success. Switches process failures rather normally. This new behavior is shown in [Figure 23](#) (which begins at the end of a second consecutive failure).

Figure 23 Auth-Fail-VLAN—New Behavior



After a third consecutive failure, the port is enabled rather silently. There is no new CLI added for this change in functionality. At this point of enabling the port, any supplicant can choose to access the network or not (which is ultimately out of the control of a switch). This provides a predictable supplicant behavior, works with any EAP method, and provides a supplicant-agnostic solution. There remain systemic-level gaps (such as the agreement of end-to-end session state) but it should be deployable based on fixes to the DDTs referenced above. In addition, not all customers run ACS or CSSC, so this should only impact internal deployments that attempt to get the Auth-Fail-VLAN to work. With CSSC, an end user also has the ability to temporarily stop the supplicant from the system tray anyway when they know they are traveling to a foreign network. This disables 802.1x on the supplicant, and it can be treated as a clientless session.

Clientless-Based Authentication

Currently, 802.1x is the recommended port-based authentication method at the access layer in enterprise networks. It has the following three primary components:

- Supplicant
- Authenticator
- Authentication server

Typically, the authenticator tries to authenticate the host device running the supplicant software to the authentication server. With some operating systems, the 802.1x supplicant capability is enabled by default (for example, Windows XP), but not all devices have this supplicant capability embedded into their operating system. For example, most printers, IP phones, fax machines, and so on, do not have this

capability but still need to be allowed into the network even without 802.1x authentication. A supplemental authentication technique should be employed as the basis of the non-responsive host issue with 802.1x. This solution-based feature set is MAC Authentication Bypass (MAB). In addition, exception lists on routers or switches are not scalable for large enterprises. Thus, a method is needed for supporting these hosts.

For network virtualization, access control must also focus on clients who do not possess 802.1x capability, or whose 802.1x capability may be temporarily suspended to support mobility into environments where the end user/client may not be otherwise known to the authentication infrastructure in advance. When 802.1x is implemented in such an environment, a customer typically needs the ability to dynamically provision individual MAC addresses (without impacting service availability) for network authentication of non-responsive devices such as printers, video conferencing units, satellite receivers, faxes, and so on. MAB is intended to control network access based on a MAC address. The goals of MAB are to provide network access control on a port basis, based on a MAC address, and to dynamically apply policy to a client session based on a MAC address.

The Guest-VLAN may also be used to provide access for clients incapable of 802.1x and where the client MAC address may be unknown in advance. Although originally designed as a deployment enabled for 802.1x supplicant functionality on end stations, the Guest-VLAN provides an option for mobile guest users as well.

In addition, this document reflects updates to changes in recent functionality across the Cisco Catalyst switching product line that may impact the related architecture to support network virtualization.

Static VLAN Configuration

In this approach, each switch port on the access layer switches in the campus needs to be manually assigned to a VLAN. There are multiple drawbacks to this approach, including the lack of mobility capabilities across the enterprise network and the lack of any mechanism to identify the user before allowing connectivity to the network. In a design supporting multiple user groups that need to remain isolated from each other, there is also the drawback of increased costs because each switch port is reserved for a specific user group, even when not used to capacity.

802.1x Guest-VLAN

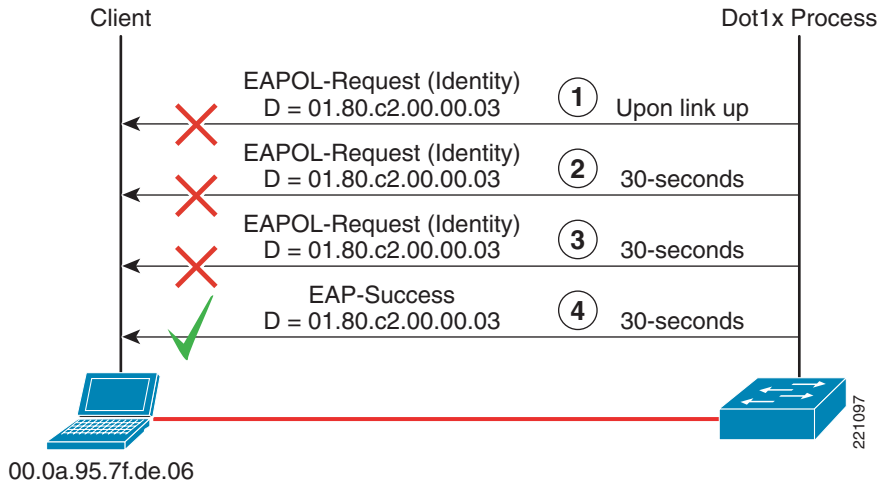
For enterprises that are starting to deploy 802.1x in their networks, leveraging Guest-VLAN functionality is a key element in providing network access to clients that are not equipped with an 802.1x supplicant. The 802.1x Guest-VLAN functionality was initially deployed as a migration tool to allow enterprises to easily migrate client devices to support 802.1x, while still providing network connectivity.

Any VLAN can be configured as the Guest-VLAN, voice VLANs (VVID), and the VLAN used for Remote SPAN (RSPAN). The Guest-VLAN feature is currently supported across all Cisco Catalyst platforms (4500, 3750, 3560, 2950 running Cisco IOS, and 6500 running CatOS); it will be integrated into Cisco IOS software releases for Catalyst 6500 platforms in the near future.

802.1x Guest-VLAN Functionality

Figure 24 shows the functionality of the 802.1x Guest-VLAN feature.

Figure 24 Guest-VLAN Feature



Currently, when a switch port initially receives a link, an EAP-Identity-Request message is transmitted to actively look for an 802.1x supplicant. This happens regardless of whether the device that connected to the port is actually equipped with the supplicant.

Assuming that the user does not have the 802.1x capability on their machine, the request from the switch goes unanswered. After the expiration of a timer (**tx-period**), the switch sends a new EAP-Identity-Request frame. This behavior is dictated by the 802.1x specification. This process continues until the third request from the switch goes unanswered. The number of retries is determined by the value of the **max-reauth-req** parameter. After the maximum number of retries is exceeded, and if the switch port has been configured with the 802.1x Guest-VLAN functionality, the port is moved to the Guest-VLAN and the switch sends an EAP-Success message to the client. This message is ignored and discarded by the client.

From the perspective of the 802.1x process, the port has become authorized and the 802.1x state machine has entered the authenticated state; no further security or authentication mechanisms are applied (the 802.1x state machine stops running). It is basically as if the administrator has disabled 802.1x and hard-set the port into that specific VLAN.

802.1x Guest-VLAN Configuration

The behavior illustrated in [Figure 24](#) is valid when using default values for the 802.1x parameters that affect Guest-VLAN functionality: **max-reauth-req** and **tx-period**.

The **max-reauth-req** parameter sets the maximum number of times that the switch retransmits an EAP-Identity-Request frame on the wire before receiving a response from the connected client. This value is set to two by default. This is why [Figure 24](#) shows two retries (at Steps 3 and 4) after the initial EAP-Identity-Request frame sent at link-up. The commands used to change this parameter (in CatOS and IOS) are as follows:

- CatOS

```
cat6500> (enable) set dot1x max-reauth-req ?
<max-reauth-req>      maximum number of retries to supplicant (1..10)
```

- Cisco IOS

```
cat3750(config-if)#dot1x max-reauth-req ?
<1-10> Enter a value between 1 and 10
```

The **tx-period** parameter sets the number of seconds that the switch waits for a response to an EAP-Identity-Request frame from the client before retransmitting the request. The default value is 30 seconds and is configurable as follows:

- CatOS

```
cat6503> (enable) set dot1x tx-period ?
<tx-period> tx period (1..65535 seconds)
```

- Cisco IOS

```
cat3750(config-if)#dot1x timeout tx-period ?
<1-65535> Enter value between 1 and 65535
```

The **max-req** parameter is part of the configurable 802.1x parameter in Cisco IOS. The **max-req** parameter is different from the **max-reauth-req** parameter and represents the maximum number of retries a switch performs for EAP-Request frames of types other than EAP-Identity-Request. Basically, this parameter refers to EAP-Data frames, which are the EAP frames exchanged after the supplicant has replied to the initial EAP-Identity-Request frame. For this reason, the **max-req** parameter is effective only when there is a valid 802.1x supplicant connected, and it does not apply to Guest-VLAN services.

The configurable values for the parameters shown in the preceding configuration example are consistent between the various Catalyst switch platforms, when running the following minimum Cisco IOS software releases (previous releases are characterized by platform-specific configurable values):

- Catalyst 3560—12.2(25)SEE
- Catalyst 3750—12.2(25)SEE
- Catalyst 4500—12.2.(31)SG

For a Catalyst 6500 running CatOS software, the situation is different; the main distinction is the fact that in CatOS releases earlier than 8.5, there is no **max-reauth-req** parameter. This implies that the same parameter described above (**max-req**) is used to tune both the number of retries for the EAP-Identity-Request and EAP-Data frames. Note also that the configurable values are consistent with the one detailed for Cisco IOS: **max-reauth-req** (and **max-req**) can vary from 1 to 10 and **tx-period** from 1 to 65535.

The overall configuration of the 802.1x Guest-VLAN is relatively simple but differs on switches running IOS and CatOS software releases, as follows:

- Cisco IOS

```
interface FastEthernet0/1
switchport access vlan 2
switchport mode access
dot1x pae authenticator
dot1x port-control auto
dot1x guest-vlan 10
dot1x max-reauth-req 2
dot1x timeout tx-period 30
spanning-tree portfast
spanning-tree bpduguard enable
```

- CatOS

```
set vlan 2 2/1
set port dot1x 2/1 port-control auto
set port dot1x 2/1 guest-vlan 10
set spantree portfast 2/1 enable
set dot1x max-reauth-req 1
set dot1x tx-period 30
set spantree portfast 2/1 enable
set spantree bpdu-guard 2/1 enable
```

**Note**

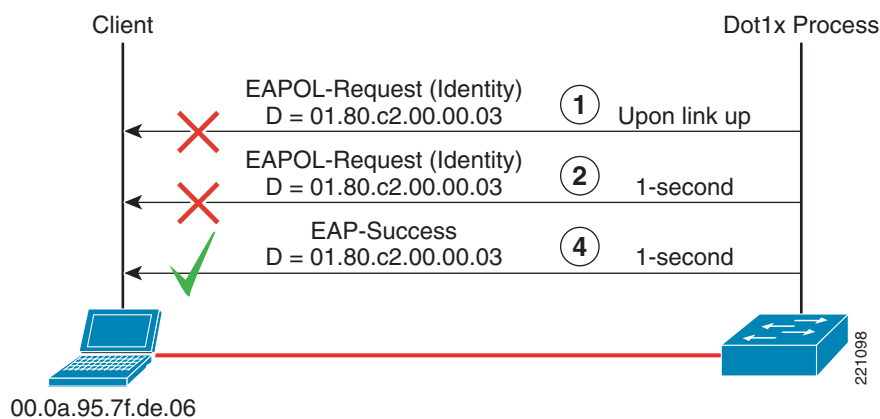
In CatOS systems, the values for *max-reauth-req* and *tx-period* are set at a global level, and not per port, as they are in Cisco IOS software.

The following formula calculates the time interval before the Guest-VLAN is enabled:

$$[(\text{max-reauth-req} + 1) * \text{tx-period}]$$

The end station must then attempt to send traffic into the network, so the specific time to ultimately authenticate the end device varies. The operation of tweaked timers to timeout 802.1x quickly as indicated above is shown in Figure 25.

Figure 25 Guest-VLAN with Tweaked Timers



This configuration should be attempted only after considering the consequences that this can have on the regular functionality of 802.1x. Analyzing the integration issues between 802.1x and DHCP at startup time helps in understanding this.

If a user starts up a machine equipped with an 802.1x supplicant, two possible scenarios can occur in relation to the use of 802.1x machine authentication after connecting to a switch port configured for Guest-VLAN. A complete description of machine authentication is not within the scope of this document.

The following two scenarios are possible:

- The 802.1x supplicant is enabled for machine authentication and the switch port is configured with a *max-reauth-req* setting of 0 and *tx-period* setting of 1. At system startup, and subsequent port link-up, the switch immediately sends an EAP-Identity-Request frame in an attempt to find a supplicant online. As a consequence of the 802.1x parameter settings defined here, the switch port is deployed into the Guest-VLAN after two seconds and the 802.1x state machine stops before the supplicant can authenticate. At a certain point during the startup process, the supplicant on the clients initializes and, because machine authentication is enabled, it can send an EAPOL-Start frame to restart the authentication process. This message is sent by default with CSSC, but not with the native Windows XP 802.1x supplicant, which requires a specific setting of the Windows registry. This is described at the following URL:

<http://www.microsoft.com/WindowsServer2003/techinfo/overview/wififaq.msp#EAAAA>

However, even assuming that the 802.1x supplicant is enabled to send EAPOL-Start frames, the DHCP and the 802.1x processes are completely asynchronous. Therefore, the machine can acquire an IP address from the DHCP pool associated to the Guest-VLAN even before sending the EAPOL-Start frame. In this case, the IP address must be released and renewed after the machine

authentication process completes successfully, because the port can now be deployed in a different VLAN from the Guest-VLAN. In this case, things can break if the supplicant running on the machine is not able to trigger this DHCP renewal. The machine would not be able to get an IP address in the correct subnet.

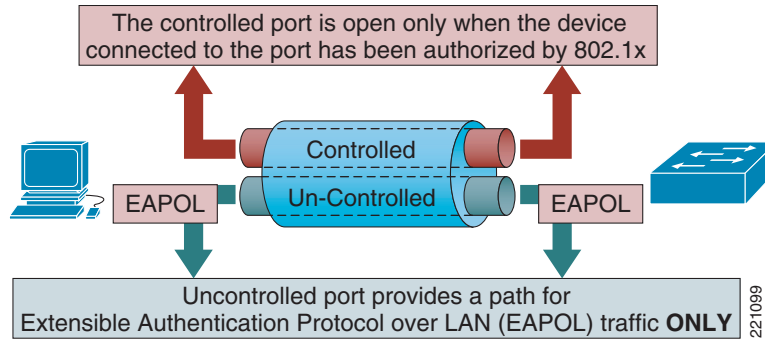
Tests run with various supplicants showed that all of them are able to renew the IP address after the machine authentication process completes. However, this happens by default with CSSC but not with the Microsoft client. Windows XP requires the registry settings described previously or the machine does not send the EAPOL-Start and therefore is stuck in the Guest-VLAN. The same situation occurs when the user logs in through the Graphical Identification and Authentication (GINA) interface (which serves as the gateway for interactive logons).

- The 802.1x supplicant is not enabled for machine authentication. In this case, during the startup process the switch port could be deployed into the Guest-VLAN and the machine could get an IP address from the Guest-VLAN pool. This happens not only when setting the *max-auth-req* and *tx-period* parameters at the minimum possible values, but also every time the startup process is longer than $(\text{max-reauth-req} + 1) * \text{tx-period}$ seconds. A typical Guest-VLAN security policy limits communications to internal resources; this can break the startup process in all the scenarios where Windows clients need to participate in a Windows Active Directory (AD) networking environment. Even assuming the connectivity with an AD domain controller is not required, after the user logs in and successfully authenticates, there is the same need to renew the IP address as described previously. Validation tests run with various supplicants show that, by default, the CSSC clients are able to renew their IP address, whereas the Microsoft supplicant requires the registry setting to be modified to initiate the EAP authentication process after the user logs in.

In conclusion, it is possible to set the *tx-period* and *max-reauth-req* parameters to the minimum configurable values to reduce the time interval required for the deployment of a switch port in the Guest-VLAN. To avoid breaking the 802.1x functionality when using the Microsoft XP supplicant, Cisco recommends that you modify the default Windows registry values to allow the Microsoft supplicant to send EAPOL-Start frames. This is the default behavior for CSSC clients.

Wake-on-LAN Primer

Wake-On-LAN (WoL) is an industry standard, which is the result of the Intel-IBM Advanced Manageability Alliance. WoL creates a power management wake-up event. This is an advanced power management capability on many network interface cards (NICs) in the industry today. NICs that support WoL have an extra connector and cable to connect to the motherboard. After a machine goes into suspend mode, it can be automatically reactivated when data from the network is received by the NIC. This capability can be used to wake up a mail server machine to deliver mail, for software management pushes, to deploy patches overnight, and so on. By default, 802.1x and WoL are mutually exclusive, because of the architecture of 802.1x, as shown in [Figure 26](#).

Figure 26 Standard 802.1x Operation

As indicated above, a switch exerts control over a virtual port in both directions. This is known as a bi-directional controlled port. This means only EAPOL should come into or go out of the switch port until authenticated. However, the operational direction of the controlled port can be changed per section 6.4(b) of the IEEE spec for 802.1x. Thus, in an effort to interoperate with WoL environments, most Cisco Catalyst switches provide unidirectional controlled port functionality as an optional configuration. WoL is a per-port feature. Operationally, the controlled port should then only operate in one direction. A WoL “magic packet” can now exit the network to wake a machine up if necessary. The machine still must 802.1x authenticate before successfully send traffic into the network. This corresponding configuration is as follows:

- Cisco IOS

```
interface FastEthernet0/1
  switchport access vlan 2
  switchport mode access
  dot1x pae authenticator
  dot1x port-control auto
  dot1x control-direction in
  spanning-tree portfast
  spanning-tree bpduguard enable
```

- CatOS

```
set vlan 2 2/1
set port dot1x 2/1 port-control auto
set port dot1x 2/2 port-control-direction in
set spantree portfast 2/1 enable
set spantree bpdu-guard 2/1 enable
```

The configuration above represents a weaker deployment of the technology because it allows outgoing traffic on a port before being secured by 802.1x, while still dropping all the incoming traffic on a port that has not yet authenticated. However, a subtle change is that spanning tree is now placed in a forwarding state for any ports that are not yet authorized.

**Note**

A best practice is to enable WoL functionality along with 802.1x only on the ports where it is needed.

Minimum releases for the support of this per-port functionality on Catalyst switches are as follows:

- Catalyst 6500—CatOS 8.3(1)
- Catalyst 4500—12.2(31)SG
- Catalyst 3750-2970—12.2(25)SEC
- Catalyst 2960—12.2(25)FX

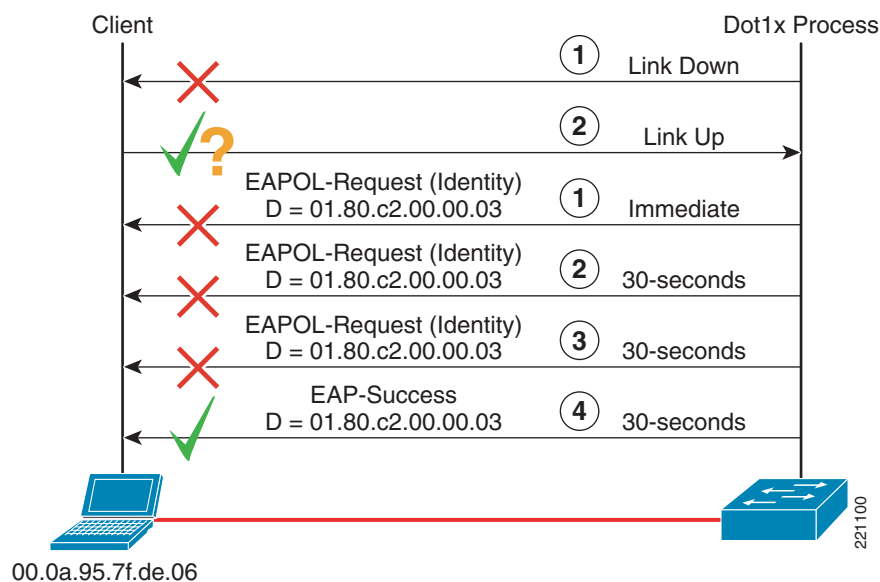
- Catalyst 2940-2950—12.1(22)EA5

A recommended best practice for any deployment of 802.1x, MAB, the Guest-VLAN, and WoL are to plan ahead of time. Test how specific Network Driver Interface Specification (NDIS) functionalities or configurations residing on end devices should impact link change.

Guest-VLAN and WoL Interaction

A switch port is down conditionally after a link-down event is processed by an authenticator as a machine goes to sleep. The link should then come back up on the port immediately. The link-up event is then processed on the port as well. If the Guest-VLAN is configured, a port is enabled into the Guest-VLAN soon after the original “go to sleep” event. This process is shown in [Figure 27](#).

Figure 27 Machine Going Into Power Save Mode with the Guest-VLAN



As shown above, a machine that goes into power save mode with the Guest-VLAN also enabled bounces link state, and then is deployed into the Guest-VLAN. There may be differences between “hibernate” and “standby” settings on end stations, so specific functionality must be examined in detail to evaluate the impact 802.1x may have on the environment. In addition, it is critical to understand whether an EAPOL-Logoff is, or needs to be, sent by an 802.1x supplicant on the specific implementation when going to sleep.

The operational behavior above exists on ports with the following configurations:

- Cisco IOS

```
interface FastEthernet0/1
  switchport access vlan 2
  switchport mode access
  dot1x pae authenticator
  dot1x port-control auto
  dot1x guest-vlan 22
  dot1x control-direction in
  spanning-tree portfast
  spanning-tree bpduguard enable
```

- CatOS

```

id1-6503-1> (enable) set port dot1x 2/2 guest-vlan 605
Port Control Direction set to IN on 2/2. Guest-VLAN can not be enabled.

id1-6503-1> (enable) set port dot1x 2/2 port-control-direction in
Port 2/2 is guest-vlan enabled, Port Control Direction can not be set to IN.

```

**Note**

The CatOS configuration example above represents an attempt to enable the Guest-VLAN on a port already enabled with WoL functionality (or vice versa). For CatOS switches, this configuration combination cannot be achieved.

For IOS-based switches, any combined deployment of WoL and the Guest-VLAN renders WoL specific functionality needless. Operationally, if the port is enabled into the Guest-VLAN already, a specific port configuration for the allowance of a unidirectional controlled port itself is not needed.


A best practice for a combined environment is to support WoL functionality from the Guest-VLAN itself, and not to bother configuring the unidirectional controlled port functionality. However, if the Guest-VLAN is being used as a means to provide third-party unauthenticated access, any WoL servers elsewhere in the network must be able to reach the Guest-VLAN. If the Guest-VLAN is configured as a separate VLAN from the access VLAN already configured on the port, this means that the WoL implementations may need to be re-configured to deploy 802.1x to reach the subnets associated to the Guest-VLAN deployed at the edge.

Network virtualization may also impact the operation of WoL functionality for PCs. For example, if the Guest-VLAN on a switch serves as entrance criteria to a separate VRF or VPN for guest access, this guest partition is not able to reach the rest of the enterprise network by design. However, this may conflict the requirement of a WoL server to be able to reach the Guest-VLAN, because it is safe to assume that the WoL server is privately owned and operated by enterprise IT staff. Thus, this needs to be planned for as part of the services edge design of a network virtualization architecture. See the *Services Edge Design Guide* for more details.

Interaction with VoIP Deployments

The integration of 802.1x and IP phones is based on the switch configuration of multi-VLAN access ports. Multi-VLAN ports belong to two VLANs: native VLAN (PVID) and auxiliary VLAN (VVID). This allows the separation of voice and data traffic and enables 802.1x authentication only on the PVID. [Figure 28](#) shows the type of communication that occurs on these two VLANs.

Figure 28 Multi-VLAN Port



When 802.1x is enabled on a multi-VLAN access port, a supplicant must complete the authentication process before getting access to the data (native/PVID) VLAN. The IP phone can get access to the voice (auxiliary/VVID) VLAN after sending the appropriate Cisco Discovery Protocol (CDP) packets, regardless of the dot1x state of the port. The use of CDP with Cisco IP phones is typically required, given the lack of pervasive support for an embedded 802.1x supplicant. However, newer IP phones models have attained supplicant capability. In addition, multi-domain authentication (MDA) on switches provides a means to authenticate a phone and data client on the wire via 802.1x and/or MAB.

**Note**

MDA and 802.1x supplicant capability on IP phones are not within the scope of this document.

The configuration commands for Cisco IOS and CatOS that are required to enable multi-VLAN functionality, in conjunction with 802.1x and Guest-VLAN functions, are as follows:

- Cisco IOS

```
interface FastEthernet0/1
  switchport access vlan 2
  switchport mode access
  switchport voice vlan 2
  dot1x pae authenticator
  dot1x port-control auto
  dot1x guest-vlan 10
  spanning-tree portfast
  spanning-tree bpduguard enable
```

- CatOS

```
set vlan 2 2/1
set port dot1x 2/1 port-control auto
set port dot1x 2/1 guest-vlan 10
set port auxiliaryvlan 2/1 2
set spantree portfast 2/1 enable
set spantree bpdu-guard 2/1 enable
```

The multi-VLAN and the Guest-VLAN features are not currently supported on Catalyst 6500 platforms running native Cisco IOS images. The Catalyst 6500 requires CatOS or Hybrid (CatOS and Cisco IOS). These features are supported on all other Catalyst platforms, beginning with the following minimum software releases:

- Catalyst 2950—12.1(12c)EA1
- Catalyst 3560—12.1(19)EA1
- Catalyst 3750—12.1(14)EA1
- Catalyst 4500—12.2(20)EWA
- Catalyst 6500—7.6(1)

This section describes the circumstances in which the Ethernet port on the IP phone can be shared between users equipped with 802.1x supplicants and users that do not have supplicants.

When accessing the network by connecting to the port on an IP phone, no link-down event occurs on the switch port when the PC is later removed. Therefore, the switch is unaware of this event, which poses security vulnerabilities.

**Note**

The same considerations are valid in cases where hubs are deployed between the end users and the access layer switch. Cisco does not officially support the deployment of hubs in conjunction with 802.1x, so the topic here is limited to the use of IP phones.

To determine the conditions under which the same IP phone PC port can be sequentially used by various categories of users (users equipped with an 802.1x supplicant and those who are not), two main aspects must be considered:

- The capability of the IP phone to send an EAPOL-Logoff message on behalf of the client after it detects a link-down event on the PC port (this functionality is called proxy EAPOL-Logoff).

**Note**

Proxy EAPOL-Logoff was introduced in the Cisco 7940 and 7960 phones in firmware 7.2(2) and the Cisco 7911, 7941, 7961, 7970, and 7971 phones in firmware 7.0(1). Both images were released in June 2005. More information can be found at the following URL:
http://www.cisco.com/en/US/products/hw/phones/ps379/prod_release_note09186a0080461f84.html.

- The mode of operation of the 802.1x process on Catalyst switch ports after receiving the EAPOL-Logoff message. This differs depending on the switch platform and on the specific software image. Based on this, it is possible to distinguish the two scenarios discussed in the next subsections.

Scenario 1

The considerations in this scenario are valid for Cisco Catalyst switch platforms running Cisco IOS software versions previous to the following:

- Catalyst 2950 (IOS)—12.1(22)EA2
- Catalyst 3560 (IOS)—12.2(25)SE
- Catalyst 3750 (IOS)—12.2(25)SE
- Catalyst 4500 (IOS)—12.2(25)EWA

The same considerations are valid for the Catalyst 6500 running the following CatOS release (and newer):

- Catalyst 6500 (CatOS)—8.4(4)

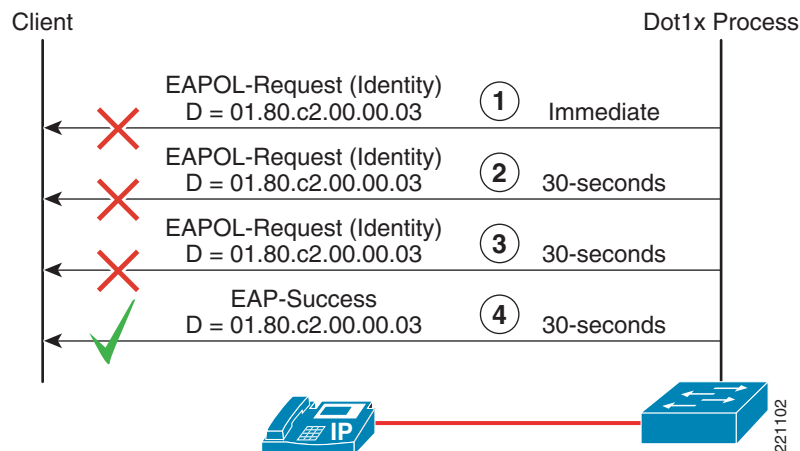
As previously mentioned, when a user connects to the Ethernet port on a Cisco IP Phone, the switch port is unable to detect when the user later disconnects. The proxy EAPOL-Logoff capability is provided on Cisco IP phones to communicate to the 802.1x process that the user (who might have been previously authenticated) has left the port, so that the switch port can be moved from the authenticated to the unauthorized state.

In describing the interaction of the Guest-VLAN with an IP phone, it is essential to distinguish between the two cases where the Cisco IP Phone supports EAPOL-Logoff.

IP Phone That Supports EAPOL-Logoff

When an IP Phone is plugged into a switch port configured as shown in the preceding configuration example, the port is assigned to the Guest-VLAN (see [Figure 29](#)).

Note that the Cisco IP phone is not able to reply to the EAP messages originated from the switch given the lack of 802.1x supplicant.

Figure 29 IP Phone and the Guest-VLAN

The 802.1x state machine running on the switch port has now entered into the authenticated state. At this point, what happens depends on who is connecting to the Ethernet port on the IP phone. The port remains configured in the Guest-VLAN if one of the following two situations occurs:

- A guest connects using a machine that does not have an 802.1x supplicant.
- An internal employee (or even a guest) connects using a machine that is equipped with an 802.1x supplicant that is not able to send EAPOL-Start (remember that by default Microsoft Windows XP supplicant does not send an EAPOL-Start). In this specific case, the user might wonder why the 802.1x authentication completes successfully when connecting to a switch port, but does not when plugging into the Ethernet port on the IP phone.

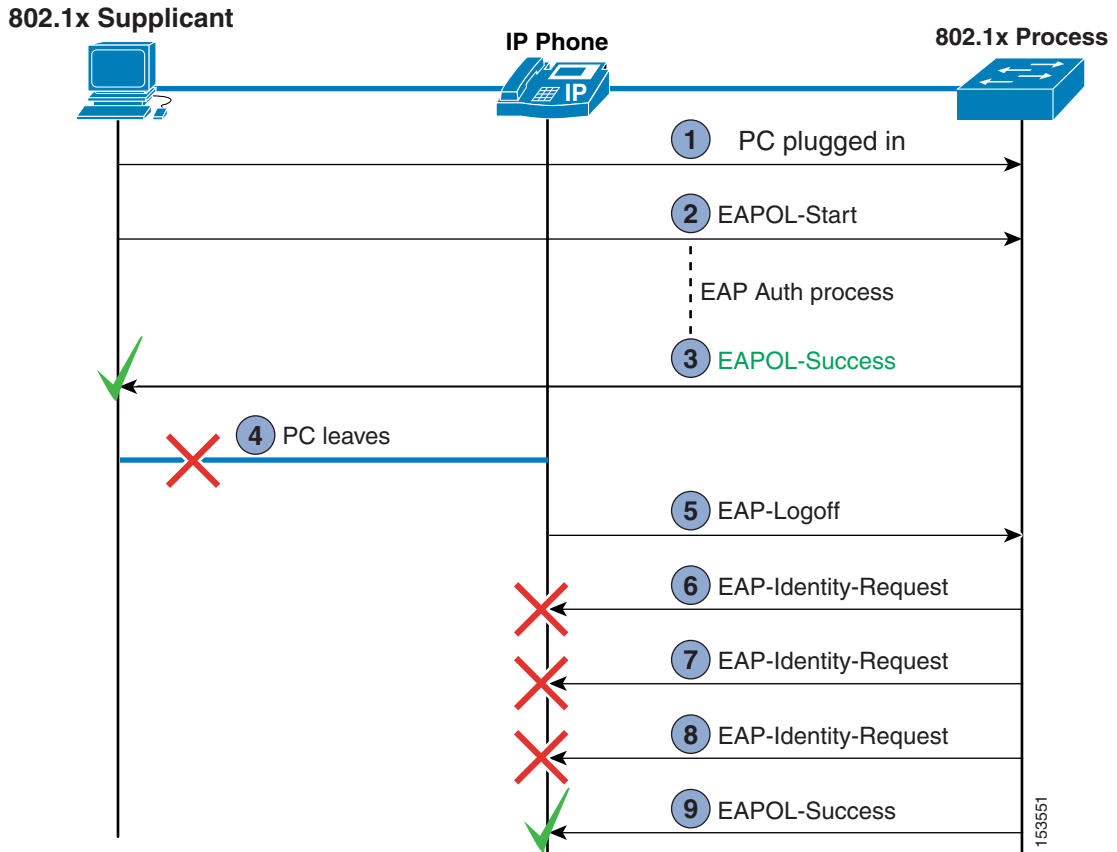
In either of these two situations, the 802.1x state machine remains in the authenticated state (it basically stops running); this is true even when the user disconnects from the IP phone.

A different situation occurs if an employee connects to a machine that is equipped with a supplicant that is able to send EAPOL-Start. The reception of this packet restarts the 802.1x state machine on the switch port, and assuming that the authentication is successful, the port is deployed in the proper data VLAN.

However, with a Catalyst 4500, running 12.2(31)SG, an 802.1x supplicant that plugs into a phone in the scenario above may not receive an EAP-Success. In error, an EAP-Success is 802.1Q-tagged and dropped by an IP phone. CSCse49193 accounts for this. In addition, on the 4500, 802.1x may successfully authenticate the user, but the user may remain authorized on the Guest-VLAN. CSCsh47641 accounts for this. Both of these bugs will be fixed in 12.2(31)SGA, so it is highly recommended to upgrade to this release upon general availability.

When the employee disconnects, the IP phone sends an EAPOL-Logoff message to the switch to inform the 802.1x process of this event. As a consequence, the 802.1x state machine on the switch port enters the unauthorized state and 802.1x stays active. After sending a few EAPOL-Identity-Request frames (this number is dictated by the value of the *max-reauth-req* parameter) looking for a supplicant, the switch port is finally deployed into the Guest-VLAN. This sequence of events is shown in [Figure 30](#).

Figure 30 EAPOL-Logoff Capability with Older Software Releases



In summary, this case represents the best scenario because it allows deploying a user into the Guest-VLAN or into a data VLAN, depending on the capability of the user client to initiate an 802.1x conversation with the switch after connecting to the IP phone. In addition, because of the EAPOL-Logoff capability of the IP phone, the switch port is always deployed into the Guest-VLAN when a previously connected employee disconnects from the IP phone, making the port reusable by a guest.

IP Phone That Does Not Support EAPOL-Logoff

This situation is identical to what was previously described until Step 4 of [Figure 30](#). After this point, because the IP phone is not capable of sending an EAPOL-Logoff frame to the switch, the port remains in the 802.1x authenticated state. This clearly opens a security hole, because an illegitimate user can now gain access to the port by spoofing the authenticated MAC address, and completely bypass the 802.1x authentication phase. The situation grows worse because the process breaks even if a legitimate user connects to the IP phone when the switch port is in the authenticated state. In fact, assuming that the MAC address of the new user is different from the one of the user previously authenticated, the switch could treat the event as a security violation. As a consequence, the switch port can be placed in an error disable state, in which case not only is the new user prevented from connecting, but the IP phone is disconnected. The only scenario where the switch port might still be deployed into the Guest-VLAN is when 802.1x re-authentication is enabled on that port, and it initializes before a user connects to the IP phone. This is recommended as a workaround condition only, and is fundamentally not a reason to enable re-authentication alone.

Scenario 2

The considerations in scenario 2 are valid for Catalyst switch platforms running the following Cisco IOS software versions (and newer):

- Catalyst 2950 (IOS)—12.1(22)EA2
- Catalyst 3560 (IOS)—12.2(25)SE
- Catalyst 3750 (IOS)—12.2(25)SE
- Catalyst 4500 (IOS)—12.2(25)EWA

Once again, the situation is different for Catalyst 6500 platforms that require the CatOS release listed here (and previous):

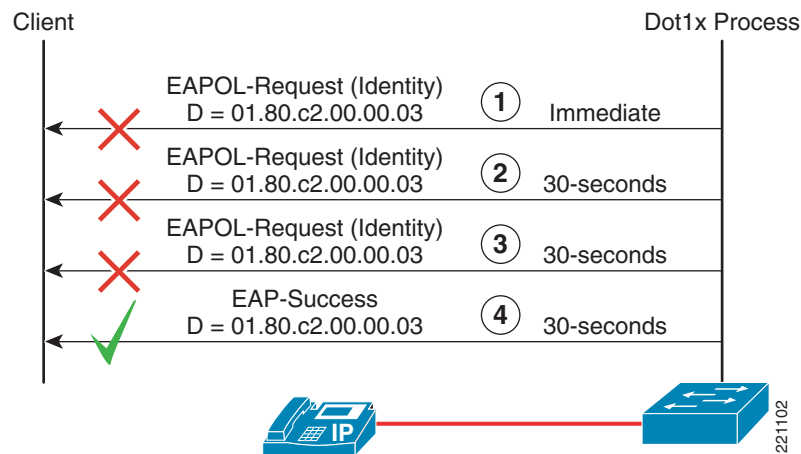
- Catalyst 6500 (CatOS)—8.4(3)

As described in Scenario 1, you need to make a distinction between the cases where IP phones have the capability to send EAPOL-Logoff frames.

IP Phone That Supports EAPOL-Logoff

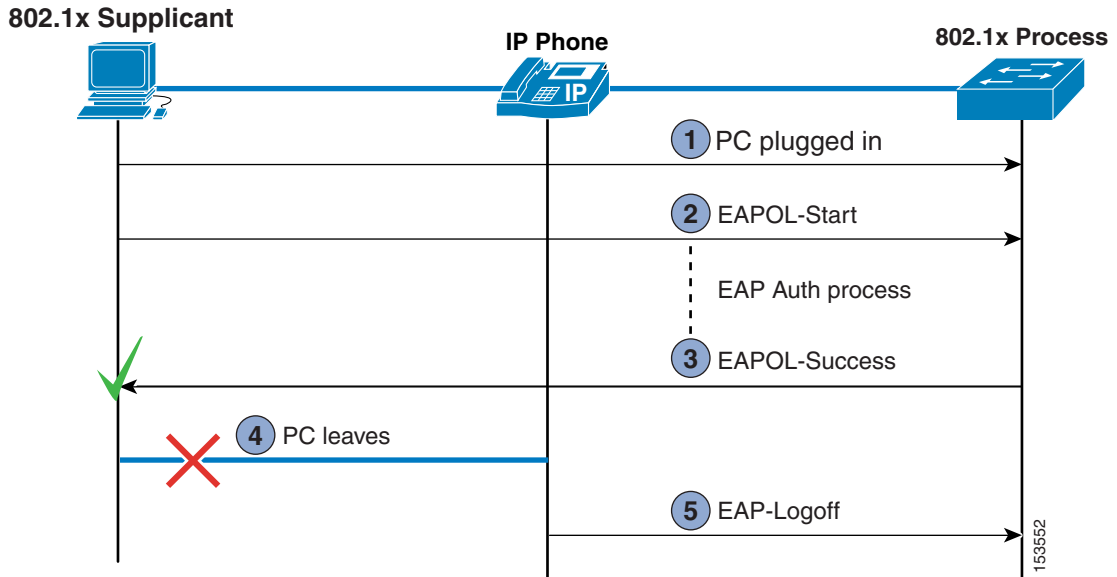
When an IP phone is connected to a switch port and configured as shown in the previous configuration example, the port is deployed into the Guest-VLAN. (See [Figure 31](#).)

Figure 31 IP Phone and the Guest-VLAN



This situation is identical to that previously described at the first step of Scenario 1. However, a different situation occurs if, starting from the initial situation shown in that scenario, an employee connects with a machine that is equipped with a supplicant capable to send EAPOL-Starts. The reception of this packet restarts the 802.1x state machine on the switch port, and assuming that the authentication is successful, the port is deployed into the appropriate data VLAN. If the employee disconnects, the IP phone sends an EAPOL-Logoff message to the switch to inform the 802.1x process of this event. As a consequence, the 802.1x state machine on the switch port enters the unauthorized state and stops running. This sequence of events is shown in [Figure 32](#).

Figure 32 EAPOL-Logoff Capability



Guests that attempt to connect a cable to the IP phone who are not equipped with a supplicant are not deployed into the Guest-VLAN because the state machine is stopped in the unauthorized state. The only event that unblocks this situation is when a machine with a supplicant that is able to send an EAPOL-Start message connects to the IP phone.

In summary, the IP phone capability of sending EAPOL-Logoff closes the security hole described in the following section, but at the same time jeopardizes the deployment of the switch port into the Guest-VLAN.

As previously mentioned, the description of events in Scenario 2 is valid for Catalyst switches running newer versions of Cisco IOS code. To maintain backward compatibility, a new global command, **dot1x guest-vlan supplicant**, was added in Cisco IOS software releases that allows you to manually configure the switches to behave as described in Scenario 1. This global command is available on Cisco IOS for all the Catalyst platforms, but it is not available on CatOS for Catalyst 6500, so these devices can only perform as described in this scenario when running software releases earlier than 8.4(4). In addition, the **dot1x guest-vlan supplicant** command has become hidden starting from the releases 12.2(31)SG for Catalyst 4500 and 12.2(25)SEE for Catalyst 3750. As of 12.2(35)SE, this command is still functional, but still hidden as well.

IP Phone That Does Not Support EAPOL-Logoff

This situation is identical to that shown in Steps 1 through 4 of the previous scenario. After this point, because the IP phone described here is not capable of sending an EAPOL-Logoff frame to the switch, the port remains in the 802.1x authenticated state. This clearly opens a security hole, because an illegitimate user can now gain access to the port by spoofing the authenticated MAC address, and completely bypassing the 802.1x authentication phase. Even worse, the process breaks even if a legitimate user connects to the IP phone when the switch port is in the authenticated state. In fact, assuming that the MAC address of the new user is different from the one of the user previously authenticated, the switch can treat the event as a security violation. As a consequence, the switch port can be placed in an error disable state, which not only prevents the new user from connecting, but also causes the IP phone to be disconnected. This differs from Scenario 1 in that, in this case, the switch port cannot be deployed into the Guest-VLAN even if re-authentication is enabled on that port and it initializes before a user connects to the IP phone.

Guest-VLAN Summary

In conclusion, when describing the integration of the 802.1x Guest-VLAN with IP telephony, the following are the recommended design choices:

- Use IP phones running a firmware version that enables them to send the EAPOL-Logoff frame when an authenticated user disconnects from the Ethernet port on the IP phone.
- Implement the Guest-VLAN functionality on Catalyst switches running either older IOS software versions or ones configured with the **dot1x guest-vlan supplicant** command. An improvement is to provide this command on a switch port level rather than globally. This allows individual ports to be configured for the operation and enhances operational flexibility by enabling this feature only on switch ports that are accessible from enterprise public areas (where guest access must be provided). It is not used in the other enterprise locations (a guest should not be able to achieve network connectivity when connecting, for example, to an IP phone that is located in an employee office). When using Catalyst 6500 platforms, Cisco recommends using the minimum software release 8.4(4).

MAC Authentication Primer

MAC address authentication itself is not a new idea. One classic type of this is port security. Another type is the Cisco VLAN Management Policy Server (VMPS) architecture. With VMPS, a customer can have a text file of MAC addresses and the VLANs to which they belong. That file gets loaded into the VMPS server switch via TFTP. All other switches then check with the VMPS server switch to see to which VLAN those MAC addresses belong after being learned by an access switch. Customers can also define actions for the switch to take if the MAC address is not in the MAC address text file. No other security is enforced.

Similar to VMPS, another type is the User-Registration Tool (URT), which uses the VLAN Query Protocol (VQP) and acts like a VMPS. Wireless also has a version of this support available on most access points (APs) and/or controllers. This base functionality for MAC address checking is already in place. For example, wireless access points can initiate a PAP authentication with a RADIUS server using a client MAC address as a username/password. APs can accomplish this based on the fact that initial associations have already been made (and based on that association traffic to/from a wireless NIC is blocked by the AP). No such association exists currently in the wired space. MAB as described in this document represents an attempt to make a wired equivalent of this functionality that is integrated with 802.1x. Similar to the operation examined here, MAB in the wireless space has its own similar security concerns, most notably the granting of network access on a MAC address. This is potentially a security risk for more enterprises, especially for wireless, because of the nature of the authentication method used. MAC addresses can be easily mirrored or spoofed.

With wireless, a MAC address check can even be done before 802.1x, so if a MAC address authentication fails, the user can still get on the network if they then pass 802.1x authentication. Cisco Clean Access (CCA) also provides a way to authorize users based on a MAC address. MAB makes an effort to leverage similar efforts that are already applied to other authentication schemes or mechanisms (802.1x/EAP). This should make deployments easier for customers to deploy and understand. MAB also represents a consolidation of current efforts toward identity, authentication, and security. These are some of the reasons why MAB is suited for network virtualization.

Other reasons to support MAB for access control are as follows:

- To provide a supplemental authentication technique using the EAP standard.
- To provide a supplemental authentication technique to be unified with 802.1x.
- Address the “all or nothing” specter of 802.1x.

- 802.1x + Guest-VLAN alone was not designed for what customers need here.
- There will always be wired devices that do not support 802.1x.
- To provide a migration path from port security.
- To provide a migration from URT and/or VMPS.

MAB, as described in this document, is intended to provide this controlled access to such devices based on their MAC address. MAB should allow non-802.1x compliant end devices to be governed by controlled access to the network in a transparent manner using a pre-populated database technique. The requirement for enabling access for clients that do not support 802.1x supplicant functionality is also applicable to the Network Admission Control (NAC) program, where a need exists to enable network access for all clients who may subsequently carry out a posture assessment. It is critical to network virtualization for MAB to leverage dynamic policy assignment as well. An overview of MAB is shown in Figure 33.

Figure 33 MAB Overview

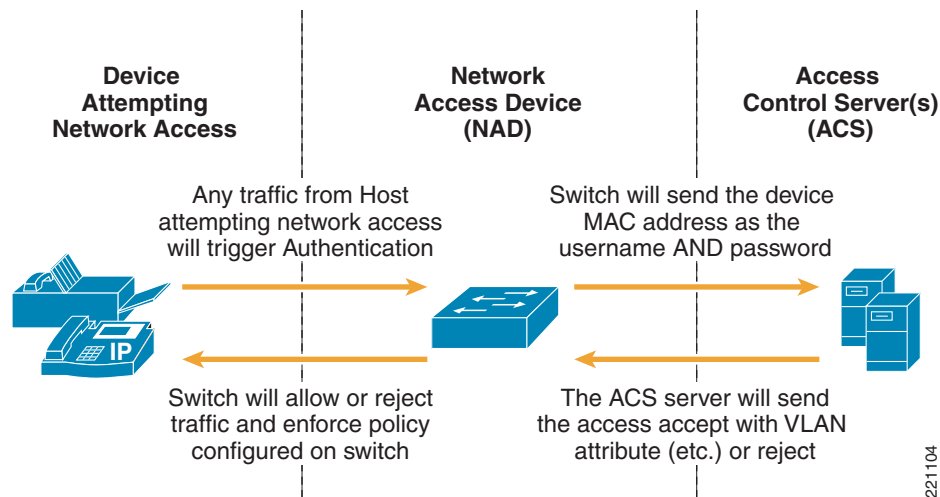


Figure 33 shows a device generating traffic (any traffic, such as DHCP, ARP, and so on), the switch captures the MAC address and forwards this as the username and password to ACS. MAB allows end users to authenticate (without any supplied credentials). As is discussed subsequently in this document, MAB is not intended to directly provide a MAC address learning mechanism. It is to be provided solely as a means of authentication and enforcement. Although MAB requires some form of a provisioning process, the described functionality is independent of any existing processes. This process alone assumes MAC addresses are already known. MAB should then allow clients that cannot/do not support 802.1x the functionality necessary to integrate into the current access control strategy for network virtualization. Like 802.1x, MAB is designed for the access layer and is supported on the following Cisco Catalyst switches referenced with minimum Cisco CatOS or IOS revisions:

- Catalyst 6500—CatOS 8.5(1)
- Cisco Catalyst 4500/4948—12.2(31)SG
- Cisco Catalyst 3750—2960—12.2(25)SEE
- Cisco Catalyst 2940—12.1(22)EA9

**Note**

Wireless LAN functionality is not examined further in the clientless context, primarily because of the nature of pervasive client capability in the overall wireless space. Because of the nature of the security threat model with the wireless media, MAC authentication is no longer recommended. There may, however, remain some cases to deal with this for wireless, such as Symbol handhelds, which may only support Wired Equivalent Privacy (WEP). For more details about wireless and MAC authentication capabilities, see the *LEAP/MAC Authentication Configuration Guide* at the following URL: http://www.cisco.com/en/US/products/sw/secursw/ps2086/products_configuration_example09186a00805e7a13.shtml.

**Note**

Branch router functionality is not examined further in the clientless context, primarily because of verification and testing resources. Cisco has traditionally provided 802.1x and its set of L3 authorization features on L3 ports popularly referred to as the spouse & kids (S&K) solution. S&K consists of 802.1x authentication, host-mode support (that is, single-host, multi-host, and multi-auth), Cisco IP phone support, guest or authentication failed handling using split-tunneling, and an implicit default behavior of MAB. This behavior is different from the behavior on Catalyst switches examined in this document. On branch routers, locally configured black and white lists based on MAC addresses can be configured as well. For more information, see *Configuring Cisco IOS Easy VPN Remote with 802.1x Authentication* at the following URL: http://www.cisco.com/en/US/tech/tk583/tk372/technologies_white_paper09186a00801fdef9.shtml#wp1002262

MAB is designed to address the market need for network edge authentication similar in nature and benefits to the functionality provided by the IEEE 802.1x framework, without the requirement for client-side code. It is intended to address a replacement technology for URT/VMPS environments. The target solution space is campus and enterprise switching. The goal of this feature set is to enhance the position of Cisco as a leader in that space by providing increased security and semi-automatic provisioning via the authentication of connected network clients.


MAC Authentication Bypass Operational Overview

The aspects of MAB operation need to be carefully considered. Before examining MAB, a rehearsal of the operation of 802.1x-enabled ports is provided for context.

802.1x Rehearsal

When 802.1x is enabled on a port, the MAC address of a machine is typically unknown until the port is authorized (or at the very least, until a supplicant sends EAPOL frames). This is because of the default operation of 802.1x, as shown in [Figure 34](#).

Figure 34 Standard 802.1x Operation

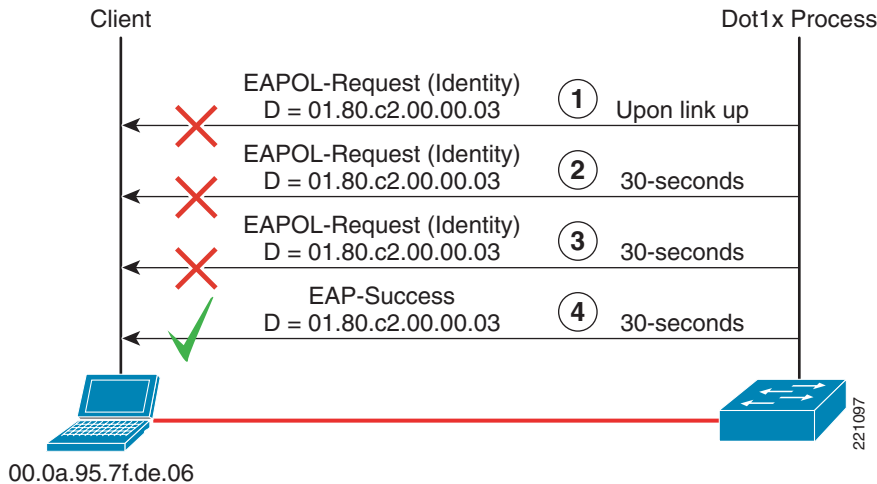


In [Figure 34](#), only EAPOL traffic is typically processed by a switch port while 802.1x is maintained in an operating and active state. Thus, any MAC address from any edge device may not be known until EAPOL frames are processed from it. These are security benefits of 802.1x, and do not change in any way with respect to any MAB implementation. Because it is noteworthy to this discussion, spanning tree is not even in a forwarding state on the port until it is authorized via 802.1x.

Guest-VLAN Rehearsal

Before MAB, the Guest-VLAN was the only alternative to provide network access to clients that do not speak EAPOL. This process is shown in [Figure 35](#).

Figure 35 Guest-VLAN Feature



There is no differentiation capability for the Guest-VLAN. If the client on the wire cannot speak 802.1x, the Guest-VLAN is enabled. Any device deployed into a Guest-VLAN may be a machine on the network that an administrator does not need or want to be placed in a Guest-VLAN. Thus, the ability to employ differentiated services based on the MAC address alone is advantageous for identification purposes. Upstream, the Guest-VLAN may also have access only to limited resources, as defined by the network administrator. Before MAB, a MAC address can be known to a switch port only after the port is enabled and placed into a Guest-VLAN. In addition, after a port is enabled and placed into a Guest-VLAN, no authentication (other than EAPOL initiation by a supplicant) takes place on the port directly, and the system can learn any number of MAC addresses on the port by default (which inherently does not

provide security). Thus, there are limitations to using the Guest-VLAN concept as a solution to provide access for any non-802.1x-enabled devices in the context of network virtualization that can be addressed through MAB functionality.

Therefore, what is needed is a way to update a switch CAM table with a (single) MAC address, while not circumventing the value added from a port-based 802.1x solution to begin with.

MAB Operation

Much like the Guest-VLAN, MAB operates based on an 802.1x timeout condition. After a switch port can ascertain that an 802.1x supplicant is not present on the port, it falls back to checking the MAC address (which is an authentication technique of lesser security). After timing out 802.1x on the port, a MAC address can be learned by the switch through classic MAC learning techniques. After a MAC address is learned by the switch, it can then be authenticated by RADIUS initiation. The RADIUS call transmits the attributes shown in [Table 1](#) as part of a RADIUS request to AAA.

Table 1 RADIUS Attributes

No.	Attribute Name	Request	Accept	Reject	Description
1	User-Name	1	0	0	MAC address sent in “hhhhhhhhhhhh” format, all lowercase with no meta characters or white spaces.
2	Password	1	0	0	Same as User-Name, but encrypted per PAP or MD5.
4	NAS-IP-Address	1	0	0	IP address of switch.
5	NAS-Port	1	0	0	Physical port of device acting as the authenticator.
6	Service-Type	1	0	0	Indicates framing to be used for framed access. This attribute indicates the type of service a user has requested, or the type of service to be provided. It <i>may</i> be used in both RADIUS-Access-Request and RADIUS-Access-Accept packets. It has been used on switches in the past to enable RADIUS exec authorization and to launch a user into enable mode. Currently set as Call-Check “10” in Access-Requests, and tracked by ACS in RADIUS Accounting logs.
12	Framed-MTU	1	0	0	Indicates MTU to be used by the user. Set to “1500”.
30	Called-Station-ID	1	0	0	MAC address of device acting as authenticator, as seen by the peer.
31	Calling-Station-ID	1	0	0	MAC address of client.
61	NAS-Port-Type	1	0	0	Indicates type of physical port on the authenticator. Set to “15” for Ethernet.
80	Message Authenticator	1	1	0	HMAC-MD5 to ensure integrity of packet.

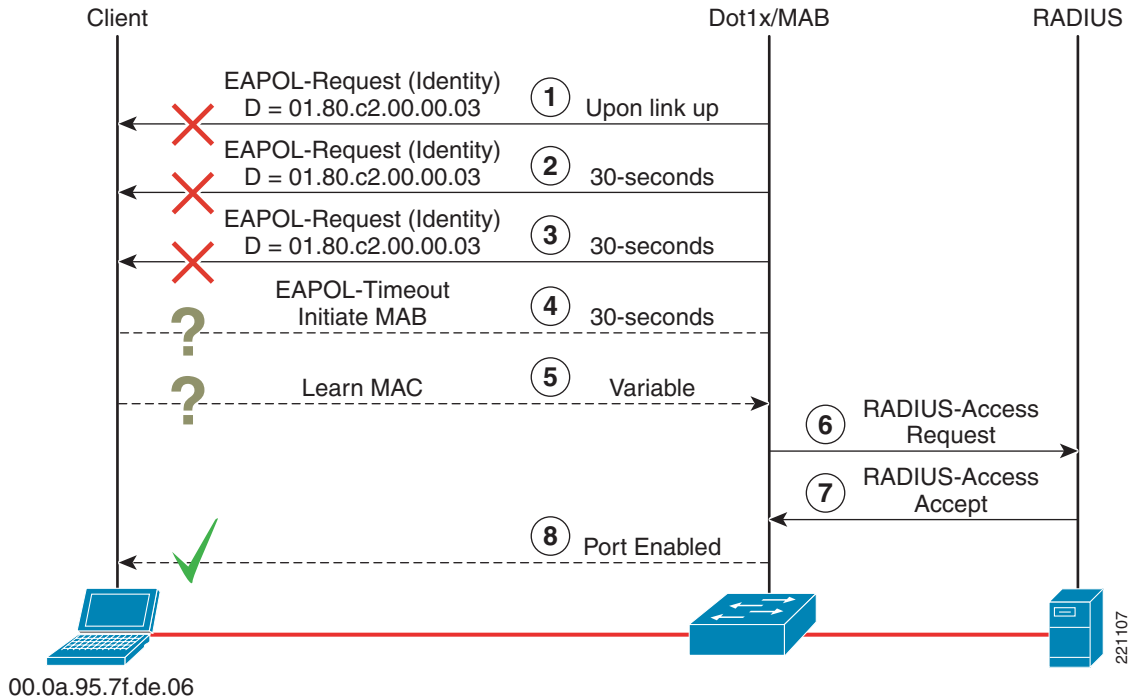


Note

CSCsh74068 has been filed for Attribute[61] above, because all IOS-based switches incorrectly set this to “5” for virtual at present. In addition, CatOS does *not* send the Framed-MTU attribute as part of a MAB request.

A complete operational flow of MAB is shown in [Figure 36](#).

Figure 36 MAB Operation



MAB initiates only after an 802.1x timeout. MAB then requires a variable amount of time for the end station to attempt to send traffic into the network for the MAC to be learned by the switch. After this occurs, RADIUS is initiated to the backend asking whether the MAC should be allowed network access.

As shown in Figure 36, after a host/device fails to supply 802.1x authentication credentials, the network access device takes the learned MAC address and hands it off to the authentication server as both the username and password. If the host/device fails to authenticate at this level, a user can optionally be placed into a pre-determined Guest-VLAN. Alternatively, the Guest-VLAN can be used as a means to support a provisioning process of the MAC address through scanning techniques, or captive portal techniques if end users are applicable to the devices seeking to be authenticated. One example of this is discussed subsequently in this document.

Ultimately, if the host/device passes with MAB credentials, the user can then be placed into the configured VLAN and can acquire an IP address to begin its desired functions. Optionally, dynamic policy can be downloaded from RADIUS the same way this is achieved with 802.1x in the form of VLAN assignment. This allows for consistent processing of authentication features to be applied in a consistent manner. Similarly, if MAB fails, the process continues indefinitely as it does with 802.1x. However, if the Guest-VLAN is also deployed, this serves as the direct failure criteria for MAB. This supports backward compatibility for existing techniques in place to provide network access to the Guest-VLAN solely in the absence of 802.1x.

Dynamic policy downloaded from an authentication server includes any capability currently available with 802.1x on the access switch in question, such as per-user ACLs, VLAN assignment, and so on. In addition, the validity of the authorized session is enforced on the switch much the same way it is enforced with 802.1x. This enforcement is achieved by restricting the traffic originating on the authenticated port to come from only the MAC address that was authorized. With MAB, only one host can be authenticated and locked down per port by default. Any new MAC address that is seen to attempt to pass traffic on a port is treated as a security violation.

Functional Details

As indicated previously, it is important to understand the format of the MAC address sent in any MAB request when MAB is used by the authentication infrastructure. Any RADIUS requests transmitted by MAB with Cisco Catalyst switches contains the following two RADIUS attributes:

- Attribute [30] (Called-Station-ID)—MAC address of the ingress interfaces of the switch or authenticator
- Attribute [31] (Calling-Station-ID)—MAC address of the 802.1x supplicant or the end-station

Both these attributes are sent in the format of “XX-XX-XX-XX-XX-XX” for all switches. This has recently been updated in switch code base to ensure both compatibility with legacy switch code and also compliance with RFC 3580. 802.1x requests operate the same way. Neither of these attributes, however, is necessarily expected to actually provide the authentication service provided by MAB, as discussed previously. Authentication and authorization are provided from RADIUS Attribute [1] (username) and RADIUS Attribute [2] (password). For IOS-based switches, and recent versions of CatOS, the format for the user-name and password attributes is simply “hhhhhhhhhhhh”; that is, an all lower-case version of “hhhh.hhhh.hhhh” with the punctuation stripped out. Therefore, if an identity infrastructure is to be built to support MAB, it should follow this format.

Figure 37 shows passed authentications on ACS from an IOS-based switch and a CatOS switch running MAB.

Figure 37 MAB from IOS 12.2(31)SG and CatOS 8.5(5)

Date ↓	Time	Message-Type	User-Name	Group-Name	NAS-IP-Address	Access Device	NAS-Port	Network Access Profile Name
02/12/2007	20:53:19	Authen OK	00145e426509	Default Group	10.137.71.4	id1-4503-2	50202	IBM-NICs
02/12/2007	20:44:13	Authen OK	00145e426509	Default Group	10.137.61.4	id1-6503-1	130	IBM-NICs



Note

Although not examined here, WLCs use the username attribute in the same manner reflected above for IOS-based switches.

However, before 8.5(5), CatOS did *not* follow this practice. CatOS transmitted the MAC address information to ACS using an “hh-hh-hh-h-hh-hh” format. ACS could not handle this if the user account is defined like the above for IOS. Figure 38 represents a passed authentication on ACS from CatOS with 8.5(4) demonstrating this condition.

Figure 38 MAB from CatOS 8.5(4) and Before

Date ↓	Time	Message-Type	User-Name	Group-Name	Caller-ID	NAS-Port	NAS-IP-Address
08/22/2005	17:21:18	Authen OK	00-d0- b7-1a- 76-0b	..	00-d0- b7-1a- 76-0b	101	172.26.198.135

As a result, if the same ACS server is used to authenticate MAB from these various types of switches, the MAC addresses need to be entered into a database twice. This is completely unmanageable, and a recommendation for 8.5(5) is a must in heterogeneous environments. This benefit is now fully realized because a single MAC need only be defined in a single location, while multiple authenticators can use it in the same format.

**Note**

ACS can not currently handle devices defined in a non-case-sensitive fashion. Any user account entered must take this into account also, or ACS processing fails completely. For more details, see CSCsd45762.

Timers are important to remember for MAB as well. For IOS-based switches, the standard timers for 802.1x are the same timers for MAB. For example, the timers to decrease the amount of time it takes to enable a port into the Guest-VLAN are *tx-period* and *max-reauth-req*. By default, it should take 90 seconds to enable a port in the Guest-VLAN. In the MAB case, this same timing is used as a signal to the switch platform that it should now open the port to learn the MAC address of an end station to begin the MAB authentication process. More timing details are discussed later in this document.

Re-authentication for MAB is supported the same way 802.1x supports it for IOS-based switches. Any re-authentication configuration that may currently exist on a switch impacts MAB clients. By default, however, if MAB re-authentication is enabled with a specific session-timeout through a port configuration, 802.1x needs to timeout again. After 802.1x times out, MAB then simply uses the MAC address the switch currently thinks is on the wire in its cache as a means to check whether the backend policy may have been disabled or changed for that MAC address. During this period, network access is persistent by default, however. Like 802.1x, MAB also incorporates the support of RADIUS attributes [27] and [29] as well. They can be set to have the switch deny access during the re-authentication event, or for the switch to re-learn a MAC instead of just using the one in the cache.

From a state machine perspective, the 802.1x state machine for an IOS-based switch also goes to an authenticated state after MAB successfully authorizes as MAC address and is updated accordingly.

By default, MAB also operates in single-auth mode like 802.1x. This means that only one MAC is allowed on the port to authenticate, and that any other MAC that appears on a port may be treated as a security violation. In addition, the host mode configured for 802.1x itself also impacts MAB. In other words, if 802.1x is configured to operate in multi-host mode, this allows any number of machines on the port subsequent to the port being authenticated. This is true for MAB as well via the same configuration.

MAC Authentication Bypass Configuration and Verification

The following network infrastructure devices were validated as part of an overall access control solution and an end-to-end network virtualization solution:

- Catalyst 6503—CatOS 8.5(4)
- Cisco Catalyst 4503—12.2(31)SG
- Cisco Catalyst 3750—12.2(25)SEE2
- Cisco Catalyst 2940—12.1(22)EA9

The following network infrastructure gear was *not* validated through testing as part of the overall solution:

- Any other 12.1E SKU (2950, 3550)
- Any other 12.2S SKU (3560, 2970, 2960, 4948)
- Any IP phone SKU
- 6500 IOS (does not support required feature set)
- Any other IP communications application
- Any branch/teleworker routers, ISRs

Configuration

MAB is a port-based feature and is required to be enabled on ports discretely. The following represents specific port configurations with MAB added.

- Cisco IOS

```
interface FastEthernet0/1
  switchport access vlan 2
  switchport mode access
  dot1x mac-auth-bypass
  dot1x pae authenticator
  dot1x port-control auto
  spanning-tree portfast
  spanning-tree bpduguard enable
```

- CatOS

```
set vlan 2 2/1
set port dot1x 2/1 port-control auto
set port mac-auth-bypass 2/1 enable
set spantree portfast 2/1 enable
set spantree bpdu-guard 2/1 enable
```



Note

The **dot1x pae authenticator** command above is not prevalent or applicable to the 2940. Where it can be applied, this command is used to enable 802.1x and for the specific type of operation under which the port should operate. the **dot1x port-control auto** command is now used as a means to configure the operating mode of 802.1x itself, assuming it has been enabled to begin with.

This is the only additional configuration required on a switch beyond an existing 802.1x configuration that may have already been deployed.

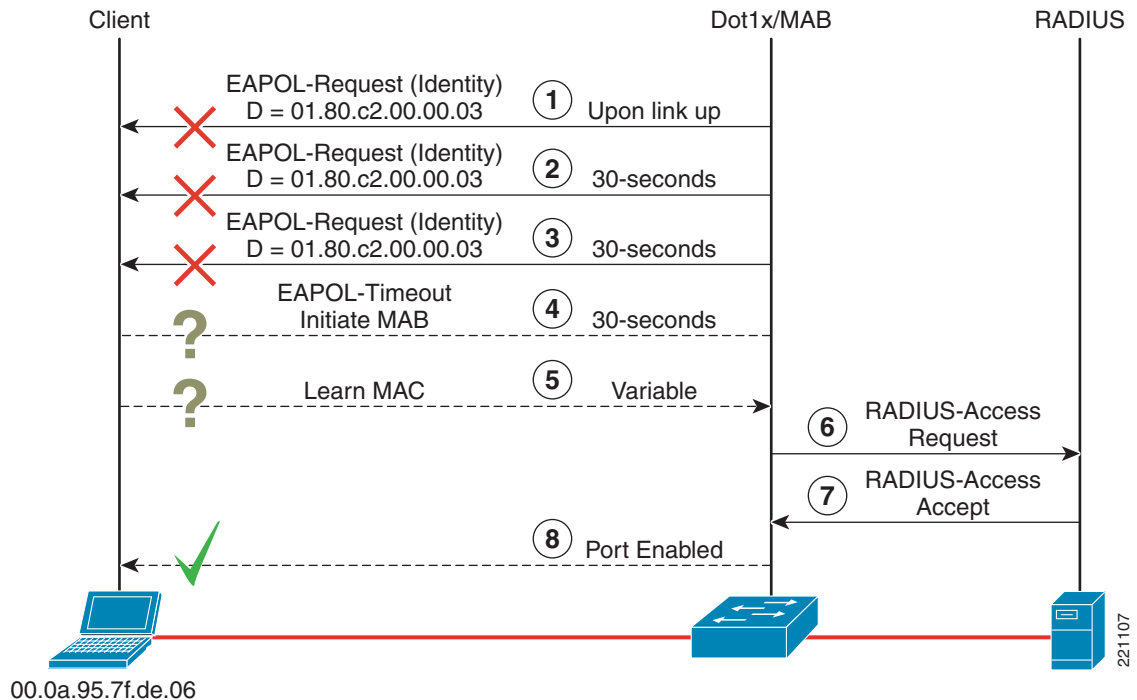
802.1x Timeout

Any port enabled for MAB at present must also timeout on 802.1x before execution of MAB can begin. The default operation of MAB clearly disrupts the boot-up process of standard PCs in an Active Directory environment because of the network delays that are imposed on the port when a machine first boots. By default, this time is over 90 seconds before a machine can begin forwarding traffic on the network successfully. MAB also cannot stand alone today as the only type of authentication configured on a port, so an 802.1x timeout must be employed for any IOS-based switched. For CatOS, however, MAB can be the only type of access control configured on a port, and is an optional configuration.

A best practice recommendation in this regard for an enterprise is to attempt to use MAB for corner cases only, and to allow 802.1x to handle the majority of controlled LAN access.

However, MAB is an ideal option for clients insensitive to delays upon boot-up or login, such as printers. An alternative to the timeout imposed by 802.1x is to reduce the timeout period. As discussed previously, the same timers and values to enable a port into the Guest-VLAN can be used for MAB to reduce the artificial delay imposed by 802.1x, and have MAB execute in a quicker manner if needed. The overall timeout process and MAB is rehearsed in [Figure 39](#).

Figure 39 MAB Operation



The *max-reauth-req* parameter sets the maximum number of times that the switch retransmits an EAP-Identity-Request frame on the wire before receiving a response from the connected client. This value is set to two by default. This is why MAB shows two retries (at Steps 2 and 3) after the initial EAP-Identity-Request frame sent at link-up. The commands used to change this parameter (in CatOS and IOS) are as follows:

- CatOS

```
cat6500> (enable) set dot1x max-reauth-req ?
<max-reauth-req> maximum number of retries to supplicant (1..10)
```

- Cisco IOS

```
cat3750(config-if)#dot1x max-reauth-req ?
<1-10> Enter a value between 1 and 10
```

The *tx-period* parameter sets the number of seconds that the switch waits for a response to an EAP-Identity-Request frame from the client before retransmitting the request. The responsibility of retransmitting the request unmodified when a response is accepted lies solely with an authentication within the confines of 802.1x. The default value for the *tx-period* is 30 seconds and is configurable as follows:

- CatOS

```
cat6503> (enable) set dot1x tx-period ?
<tx-period> tx period (1..65535 seconds)
```

- Cisco IOS

```
cat3750(config-if)#dot1x timeout tx-period ?
<1-65535> Enter value between 1 and 65535
```

The *max-req* parameter is also part of the configurable 802.1x parameter in Cisco IOS. The *max-req* parameter is different from the *max-reauth-req* parameter. The *max-req* parameter represents the maximum number of retries a switch performs for EAP-Request frames of types other than EAP-Identity-Request. Basically, this parameter refers to EAP-Data frames, which are the EAP frames exchanged after the supplicant has replied to the initial EAP-Identity-Request frame. For this reason, the *max-req* parameter is effective only when there is a valid 802.1x supplicant connected, and it does not apply to any method to deal with the timeout of 802.1x itself on the port.

The configurable values for the parameters shown in the preceding configuration example are consistent between the various Catalyst switch platforms, when running the following minimum Cisco IOS software releases (previous releases are characterized by platform-specific configurable values):

- Catalyst 3750—12.2(25)SEE
- Catalyst 4500—12.2(31)SG

For a Catalyst 6500 running CatOS software, the situation is different. In CatOS releases earlier than 8.5, there is no *max-reauth-req* parameter. This implies that the same parameter described above (*max-req*) is used to tune both the number of retries for the EAP-Identity-Request and EAP-Data frames. Note also that the configurable values are consistent with the one detailed for Cisco IOS: *max-reauth-req* (and *max-req*) can vary from 1 to 10, and *tx-period* from 1 to 65535.

The overall configuration of MAB is relatively simple but differs on switches running IOS and CatOS software releases. Complete configurations with tweaked timeouts are as follows:

- Cisco IOS

```
interface FastEthernet0/1
  switchport access vlan 2
  switchport mode access
  dot1x mac-auth-bypass
  dot1x pae authenticator
  dot1x port-control auto
  dot1x timeout tx-period 1
  dot1x max-reauth-req 1
  spanning-tree portfast
  spanning-tree bpduguard enable
```

- CatOS

```
set vlan 2 2/1
set port dot1x 2/1 port-control auto
set port mac-auth-bypass 2/1 enable
set dot1x max-reauth-req 1
set dot1x tx-period 1
set spantree portfast 2/1 enable
set spantree bpdu-guard 2/1 enable
```



Note

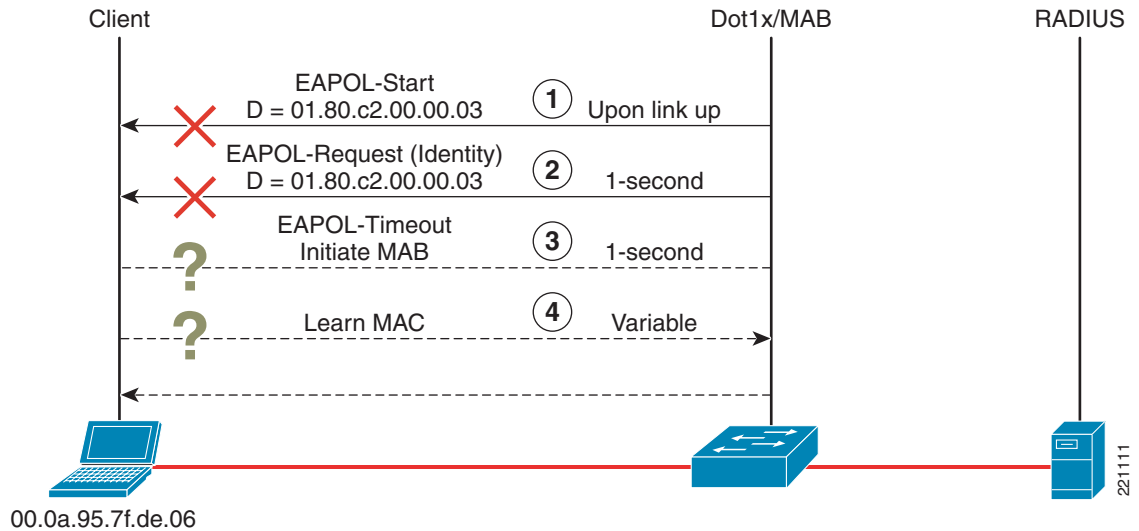
In CatOS systems, the values for *max-req* and *tx-period* are set at a global level, and not per port, as they are in Cisco IOS software.

As shown above, the timeout for 802.1x and MAB initiation can be configured as low as 2 seconds. The following formula calculates the time interval before MAB initiates:

$$[(\text{max-reauth-req} + 1) * \text{tx-period}]$$

As stated previously, MAB initiates only at this time. The end station must then attempt to send traffic into the network, so the specific time to ultimately authenticate the end device typically varies. The operation of tweaked timers to timeout 802.1x quickly as indicated above is shown in Figure 40.

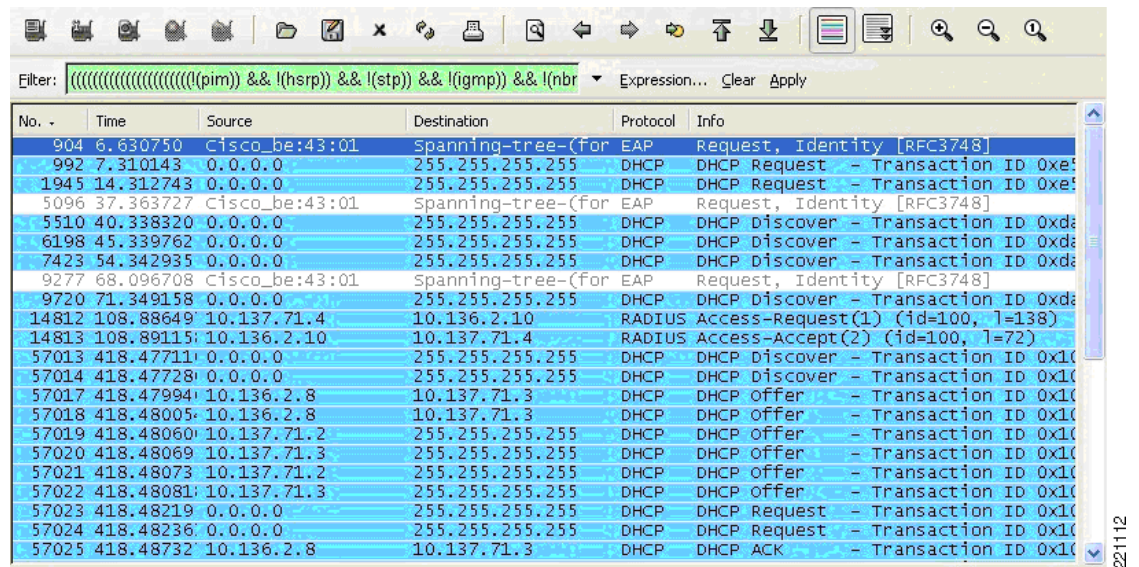
Figure 40 MAB Initiation with Tweaked Timers



This configuration should be attempted only after considering the consequences that this can have on the regular functionality of 802.1x.

Analyzing the integration issues between 802.1x and DHCP at startup time helps in understanding this. MAB was tested with default timers and Windows machines. As indicated, this causes DHCP to timeout entirely upon boot-up and any link-up condition. This process is shown in Figure 41.

Figure 41 MAB Impact on DHCP



As shown in the example above, 802.1x times and MAB is successful ~90 seconds after a link-up event. However, DHCP times out completely after approximately one minute. When Windows reverts to the internal address of 169.254.x.x, however, an IGMP report from this address causes L2 traffic to be learned by the switch, so that MAB can initiate. Therefore, although variable, MAB completes less than 40 seconds after DHCP has timed out in this example. Windows reverts to standard timeout procedures in this case, and does not attempt to renew its address for another five minutes. This is probably unacceptable to any end user experience, so timer tweaking may be needed here to enable this process to operate better for machines sensitive to this timeout condition.

**Note**

Although security feature interaction is not within the scope of this document, it was noticed that MAB did not learn an address if the packet was a DHCP frame. This may be because of DHCP Snooping also being enabled, and may cause even longer delays. CSCsg03626 was filed to account for this.

Proceed with caution for tweaking timers, however. If timers are tweaked too low, MAB (or the Guest-VLAN if configured) may execute on the device before 802.1x when the end station may be legitimately configured for 802.1x. An example of this is when a Windows machine boots. 802.1x may not execute on the machine two seconds after the machine starts trying to send traffic. Most 802.1x supplicants are applications themselves, so they also need time to load. This may be an undesired side effect, although nothing may be technically wrong about this operating condition. Security policies may need to dictate this as well. As a result, there may be no optimum for timer recommendations to make in this regard, because mileage varies based on requirements.

Verification

Following is an example of MAB working on a port of a CatOS switch:

```
id1-6503-1> (enable) sho port mac-auth-bypass 2/2
Port  Mac-Auth-Bypass State  MAC Address      Auth-State      Vlan
-----
  2/2  Enabled              00-14-5e-42-65-09  authenticated    601

Port  Termination action  Session Timeout  Shutdown/Time-Left
-----
  2/2  initialize           3600              NO                -

Port  PolicyGroups
-----
  2/2  -
```

Following is an example of MAB working on a port of an IOS-based switch:

```
id1-3750-2#sho dot1x interface g1/0/2 details
Dot1x Info for GigabitEthernet1/0/2
-----
PAE                               = AUTHENTICATOR
PortControl                        = AUTO
ControlDirection                  = Both
HostMode                           = SINGLE_HOST
ReAuthentication                   = Disabled
QuietPeriod                        = 60
ServerTimeout                      = 30
SuppTimeout                        = 30
ReAuthPeriod                       = 3600 (Locally configured)
ReAuthMax                          = 2
MaxReq                             = 2
TxPeriod                           = 30
RateLimitPeriod                    = 0
```

```

Mac-Auth-Bypass           = Enabled

Dot1x Authenticator Client List
-----
Supplicant                 = 0014.5e42.671b
  Auth SM State            = AUTHENTICATED
  Auth BEND SM Stat       = IDLE
Port Status                = AUTHORIZED
Authentication Method      = MAB
Authorized By              = Authentication Server
Vlan Policy                = N/A

```

The verified result from the 2940 implementation differs from the output of the IOS example above:

```

id1-2940-1#sho dot1x interface f0/2
Supplicant MAC 0014.5e42.6523
AuthSM State              = AUTHENTICATED
BendSM State              = IDLE
Posture                    = N/A
PortStatus                = AUTHORIZED
MaxReq                    = 2
MaxAuthReq                = 2
HostMode                  = Single
Port Control              = Auto
ControlDirection          = Both
QuietPeriod               = 60 Seconds
Re-authentication         = Disabled
ReAuthPeriod              = 3600 Seconds
ServerTimeout             = 30 Seconds
SuppTimeout               = 30 Seconds
TxPeriod                  = 30 Seconds
Guest-Vlan                = 0
AuthFail-Vlan             = 0
AuthFail-Max-Attempts    = 3
Mac-Auth-Bypass          = Enabled

```

MAC Authentication Bypass Feature Interaction

MAB and EAPOL Interaction

As shown above, MAB activates when 802.1x times out waiting for an EAPOL packet on the wire. The 802.1x state machine enters a waiting state and relinquishes control over to MAB to begin device authorization upon this timeout occurring. MAB runs passively and does not transmit any packets to detect devices. Again, the responsibility lies with the attached device to send traffic. If a device sends no traffic, then technically, a port could be listening for packets forever after MAB activates. Packets arriving on a port where MAB is active results in the switch forwarding the packets to the CPU. The source MAC address is gleaned off the packet and forwarded to the MAB process for authorization. The “trigger” packet itself is typically dropped. Before MAB activates, if an EAPOL packet is detected on the wire (such as an EAPOL-Start from an 802.1x supplicant), 802.1x never relinquishes control to MAB. The history of EAPOL packets seen on the wire is maintained as long as the port is physically connected. This “history” is lost upon physical link change.

When MAB activates, a port is typically in an unauthorized state (because 802.1x times out). Therefore, while waiting for a packet to glean a MAC address, if an EAPOL packet is detected, MAB de-activates and relinquishes complete control back to 802.1x entirely. 802.1x then attempts to authenticate the port. From then on, MAB never activates as long as the link is never lost on the port.

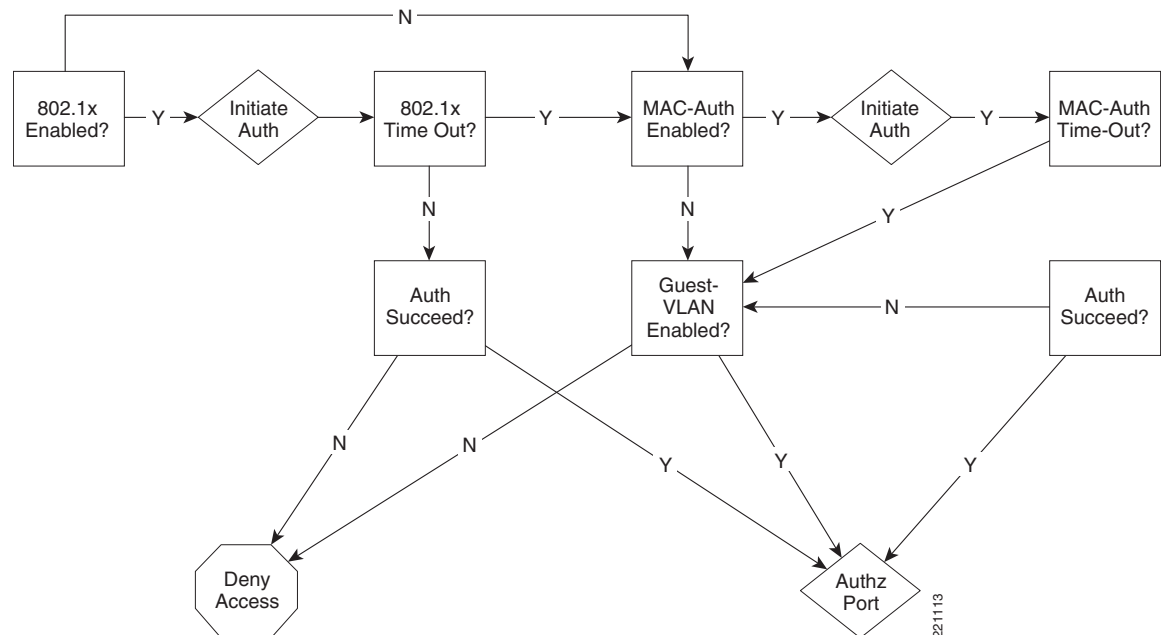
In some cases, MAB may have authorized a port already, and 802.1x is then seen on the wire. An example of this can be a successful MAB attempt before 802.1x has started on the client, or MAB being executed in an effort to assist the end station in downloading 802.1x supplicant software. Typically in this condition, the MAC addresses from both events match. However, if a port is authorized with MAC address A, and an EAPOL packet arrives with a source MAC address of B, this triggers a security violation by the switch.

On IOS-based switches, MAB cannot currently be enabled without 802.1x. This also impacts any potential re-authentication scenario. Technically, re-authentication for MAB does not exist on IOS. Therefore, if re-authentication is enabled (for 802.1x) when a switch port is authenticated via MAB, a switch sends out EAP requests upon re-authentication timer expiry, and 802.1x ultimately relinquishes control back to MAB for authorization only if no response is received. Because 802.1x has to timeout again for MAB re-authentication to occur, MAB re-authentication is not recommended. In addition, any port configuration involving 802.1x re-authentication is also not recommended. For any 802.1x re-authentication use case that may involve MAB, it is recommended that a RADIUS-supplied session-timeout be used to control the behavior for 802.1x devices only, and not for devices that have been authenticated via MAB.

MAB and the Guest-VLAN

The Guest-VLAN serves as a failure condition for MAB if configured on the same port as MAB. Otherwise, the failure process for MAB is to continually try to 802.1x authenticate the port again. For IOS-based switches today, this is primarily because of a MAB failure actually causing the port to go into the HELD state, much the same it would as if an 802.1x supplicant had failed authentication. Therefore, after the HELD state completes, 802.1x is attempted again, times out again, and MAB is attempted again. However, because the Guest-VLAN can serve as the failure criteria for MAB if configured along with MAB, this might provide a systemic value; for example, MAB and the Guest-VLAN could indirectly provide a means to provision credentials in an identity store for MAC addresses that may not be known in advance to the enterprise, as shown in [Figure 42](#).

Figure 42 802.1x, MAB and the Guest-VLAN



The operational nature of the feature interaction in [Figure 42](#) was designed primarily as part of MAB to support backward compatibility for devices that cannot speak 802.1x and have already deployed the Guest-VLAN.

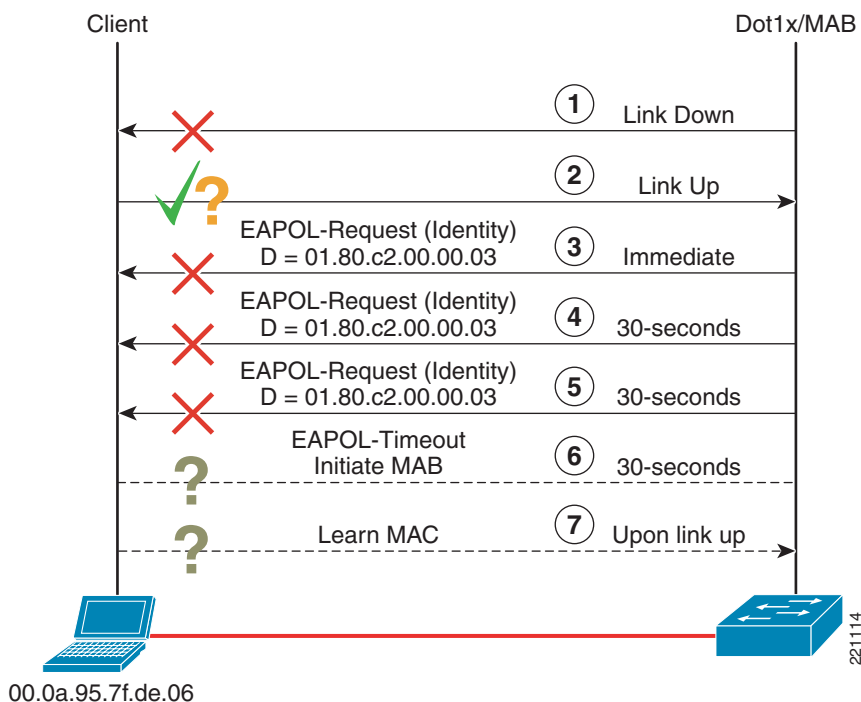
**Note**

If a port is initially configured for 802.1x with Guest-VLAN, and the port activates in Guest-VLAN, it remains there even though a network administrator enables MAB. The port link status must be flapped to initialize the 802.1x state machine.

MAB and WoL Interaction

If MAB is configured, a port is nailed up into a MAB state of “initiated” soon after the original “go to sleep” event. This process is shown in [Figure 43](#).

Figure 43 Machine Going Into Power Save Mode with MAB



As shown in [Figure 43](#), a machine that goes into power save mode with MAB also enabled bounces link state, and then is nailed up into a state of MAB needing to learn a MAC address to be able to authenticate it. There may be differences between “hibernate” and “standby” settings on end stations, so specific functionality must be examined in detail to evaluate the impact that 802.1x may have on the environment. Also critical to understand is whether an EAPOL-Logoff is, or needs to be, sent by an 802.1x supplicant on the specific implementation when going to sleep.

The operational behavior above exists on ports with the following configurations:

- Cisco IOS

```
interface FastEthernet0/1
```

```

switchport access vlan 2
switchport mode access
dot1x mac-auth-bypass
dot1x pae authenticator
dot1x port-control auto
dot1x control-direction in
spanning-tree portfast
spanning-tree bpduguard enable

```

- CatOS

```

set vlan 2 2/1
set port dot1x 2/1 port-control auto
set port mac-auth-bypass 2/1 enable
set port dot1x 2/2 port-control-direction in
set spantree portfast 2/1 enable
set spantree bpdu-guard 2/1 enable

```

For Catalyst switches, any combined deployment of WoL functionality and MAB does not impact the fundamental need to wake up machines from remote management locations. Operationally, when a machine wakes up, it must MAC authenticate because 802.1x has already timed out (while the machine was asleep). However, as discussed above for EAPOL and MAB interaction, a machine may also 802.1x authenticate when it wakes, which tears down all session state for the MAB context and 802.1x access is granted.

A best practice for a combined environment is to support WoL functionality from the statically configured access VLAN, the same way a customer would before 802.1x has been deployed.

Network virtualization may impact the operation of WoL functionality for PCs, as for example when the access VLAN on a switch serves as entrance criteria to a VRF or VPN separate from the global table, where a WoL server is typically deployed. In that case, you need to be sure that this partition can be reached from the segment on which the WoL server is located. Therefore, this does necessarily need to be planned for as part of the services edge design of a network virtualization architecture. This topology should allow WoL to work, while retaining most forms of separate network partitioning, after devices have been authorized into the networked system.

MAC Authentication Bypass Opportunities and Benefits

Location-Based Awareness

MAB can do a good job of providing MAC-based security, where only known MAC addresses are allowed access to the network, using a central RADIUS server (or identity store) to store the list of MAC addresses. This takes the burden of managing the MAC addresses off of any local switch, and is technically superior to port security in this respect. In support of network virtualization, VLANs can be assigned for granular policy as well. These benefits represent motivations behind the need for MAB. However, there is currently no easy way to have switches authenticate the device and at the same time limit the MAC to a specific location/switch. Although this functionality is not currently provided by any turnkey solution, similar capabilities exist in dial-up or WLAN models. A complete location-based system is not yet integrated into 802.1x or MAB itself for authorization purposes. However, some customer problems based on location can be solved. For example, if a customer has a device that should only be on the machine floor of a production plant (for example, robotic arm device), the authentication system may need to know that this device should only be connected to a single switch. This way, if the device shows up on another switch or location, the authentication system can realize this event and deny the authentication attempt on this basis. One way to technically achieve this is to configure ACS for

Network Device Groups (NDG). Then, as part of a Network Access Profile (NAP), Network Access Filter (NAF) can be set up based on the NDG. This can cause a MAB request to not match the NAP, because the request may originate from the wrong switch, as shown in Figure 44.

Figure 44 Deny MAB Request Based on Pre-determined Location

Date ↓	Time	Message-Type	User-Name	Group-Name	Caller-ID	Network Access Profile Name	Authen-Failure-Code	Author-Failure-Code
02/12/2007	21:14:08	Authen failed	00145e426509	Default Group	00-14-5e-42-65-09	(Default)	Access denied because there was no profile that matched	..

For more information on Network Access Profiles, see the following URL:
http://www.cisco.com/en/US/docs/net_mgmt/cisco_secure_access_control_server_for_windows/4.0/user/guide/sp.html

The above example may not suffice for general use cases. It is not intended to function as a true location-aware based service. However, because location-aware services are not yet prevalent in wired network authentication topologies, this can supplement the service for precedent. In summary, this can be technically achieved as shown above because of a specialized need.

Network Access Profile Matching and Potential Value

With ACS 4.0, a customer has the ability to match a NAP and also authenticate a MAB request.



Note

After defining a NAP in ACS, ACS tries to process NAPs first. Thus, Cisco recommends deferring to the global authentication settings for any requests that may not match any NAPs for any user request, as shown in Figure 45.

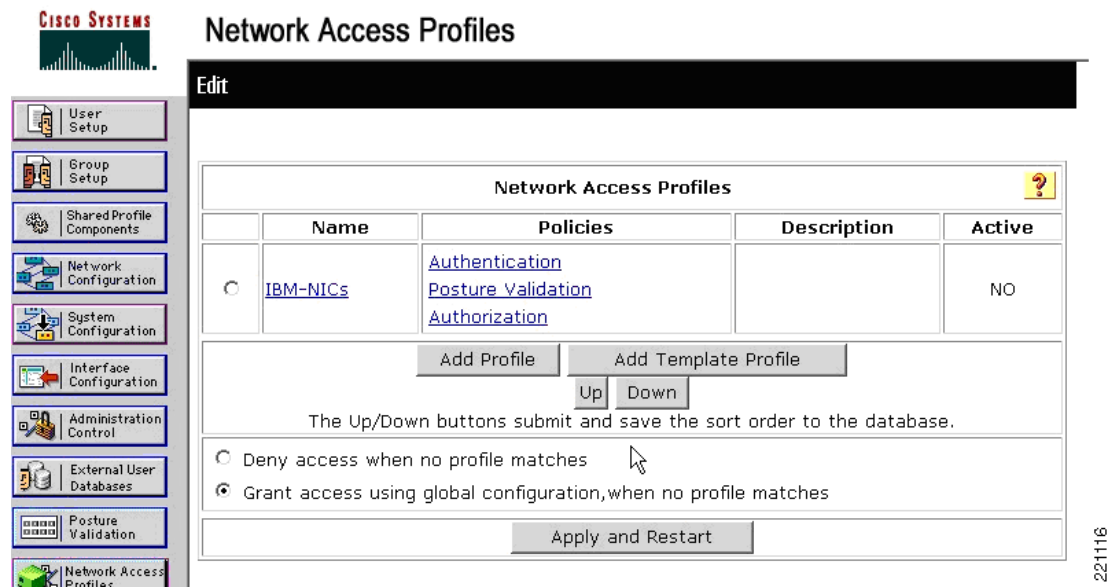
This technically provides support for the ability to incorporate MAB with a restriction that uses pattern matching to match requests even though a user request contains a MAC address as the username and password. Some customers may not have an existing database of all their MAC addresses. Therefore, as an initial rollout plan, they may be interested in wildcarding MAC addresses based on a vendor code. By leveraging a NAP, MAB can be made to work on ACS without actually defining a user, by doing the following:

1. Matching RADIUS-Attributes in such a way to distinguish MAB from something else.

There are multiple ways to do this. The only key is, by wildcarding, the request can be forced not to go to AD, not enable PEAP, or any other specific authentication settings. An example of this is in the following NAP:

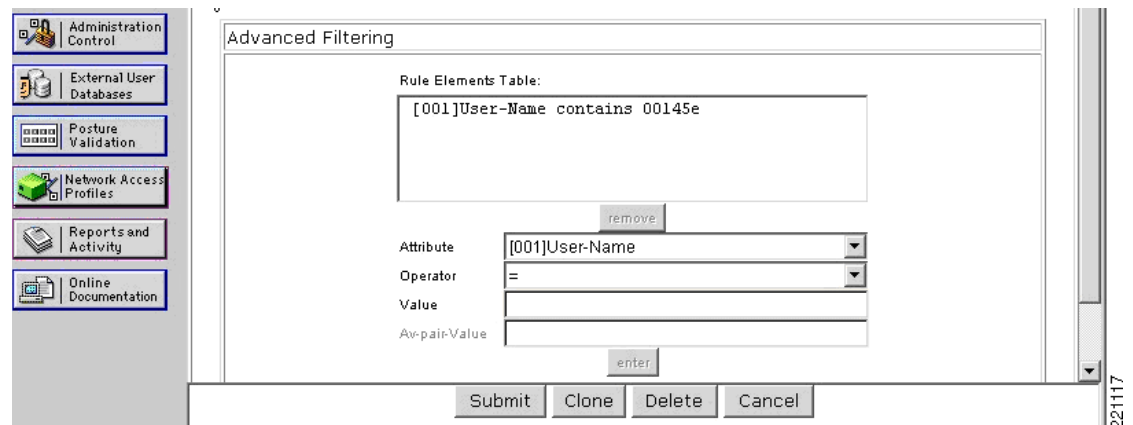
- a. First, you can create a NAP to match IBM NICs as an example, as shown in [Figure 45](#).

Figure 45 Example NAP to Match a NIC Vendor Code



- b. Then, filtering for the NAP itself should be set up. [Figure 46](#) shows an example for wildcarding a vendor code in the user-name attribute.

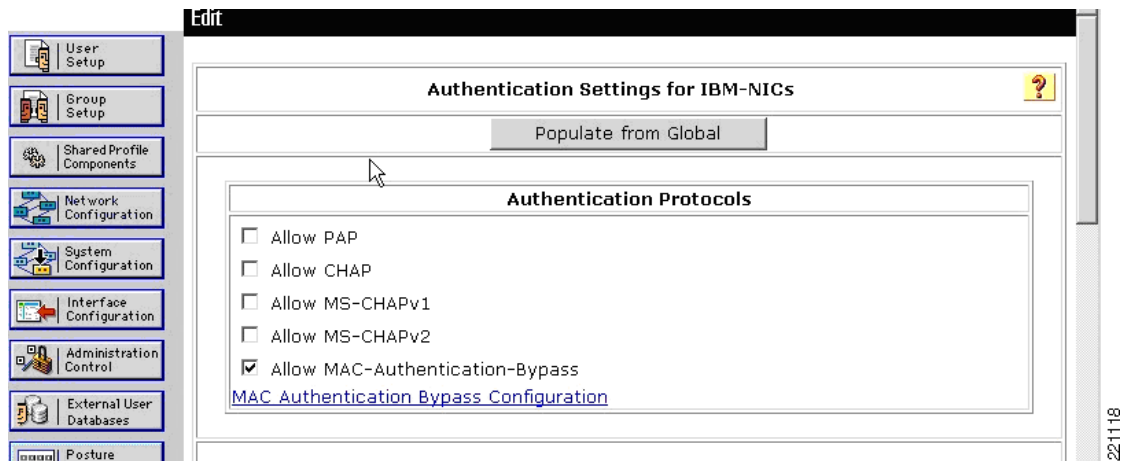
Figure 46 Example NAP to Match a NIC Vendor Code



This has now effectively wildcarded the MAB request itself. In other words, only MAC addresses matching the criteria are processed further. Otherwise, the MAB request does not match the NAP, and can be implicitly denied or fallback to global authentication settings.

2. Authentication rules should then be set up.

No special authentication configuration is needed, although the check box for MAC-Auth-Bypass must be enabled. Optionally, ACS can also be configured to place all of the MAC addresses that have matched the NAP per #1 above into a specific CiscoSecure Group. This is shown in [Figure 47](#).

Figure 47 Authentication Settings for the NAP

3. Posture rules for this NAP can be disabled, because this demonstration has nothing to do with NAC.
4. Authorization rules for the NAP can be configured. The group referenced from #2 above could then also deploy some sort of temporary access, like the dynamic assignment of the Guest-VLAN that may already be on this switch. MAC addresses that do not match the NAP to begin with can be implicitly denied.

The NAP can then be successfully verified much like the example in [Figure 45](#).

In summary, this is a demonstration of being able to wildcard a MAC address based on a vendor code or specific range of MAC addresses. At the same time, you have granted access (optionally, this access can be differentiated from full access) via MAB, although the MAC address may not have been known previously. This technique can be leveraged to help build a true identity database of MAC addresses. In addition, this technique can then be incorporated into a captive portal, for example, that an end user could interact with for registration purposes. Examples of this include home-grown applications a university may use to allow students to register their device for identification purposes onto the network in the absence of endpoint control. This identification technique can then allow for immediate network segmentation through network virtualization.


This is a demonstration of how MAB, along with network virtualization can be used to address TCO concerns of enterprises in endpoint management. This way, a network or server administrator need not play a part in the provisioning mechanism that may already be in place, even if a MAC address may not be known previously. In addition, the network access ultimately granted to the device follows the device wherever in the network the device plugs in through pervasive path isolation techniques.

Fallback Technique for New/Re-imaged Machines with WZCSVC

There are systemic challenges associated with MAB and the fallback nature of this supplemental technique in the absence of 802.1x. The first challenge is from Windows XP. A new or re-imaged PC is typically enabled for 802.1x by default. In addition, if the machine is running a default image for Windows XP, the 802.1x supplicant does not send EAPOL-Start frames even though 802.1x is enabled. This means that when the link comes up, the switch begins an 802.1x authentication event by transmitting an EAPOL-Identity-Request packets on the wire. However, although the PC is 802.1x-enabled, the supplicant is also enabled for EAP-TLS and the machine “knows” it does not have a certificate for either the machine or any user that happens to be logged into it. Operationally, a balloon message appears in the system tray at this point with “Windows was not able to find a certificate to log you onto the network”. Because a certificate is not on the device at all, Windows does not speak EAPOL to the switch. In addition, because the supplicant never sent the switch an EAPOL-Start, the switch has

no way of knowing the device is actually 802.1x capable. Therefore, this means that a brand-new machine can be initially deployed into the Guest-VLAN, or if the MAC address is known before the connection event, MAB can be used as a means to help deploy 802.1x, or at least provide network access to the device in a fallback method even though 802.1x is technically enabled on the client. An example of this is on the end station is shown in [Figure 48](#).

Figure 48 **802.1x Enabled by Default, Although Treated as Clientless**



The client, having no certificate provisioned before this event, does not reply to this request at all, and demonstrates the message above. The above scenario may hold true for machines that have been re-imaged as well, depending on the operational configuration or characteristics of the image itself.

Recommended best practices for these types of machines are to enable the 802.1x supplicant on the device to send EAPOL-Start frames (through registry setting) only *after* appropriate credentials have been loaded (such as any needed certificates).

MAC Authentication Bypass Limitations and Challenges

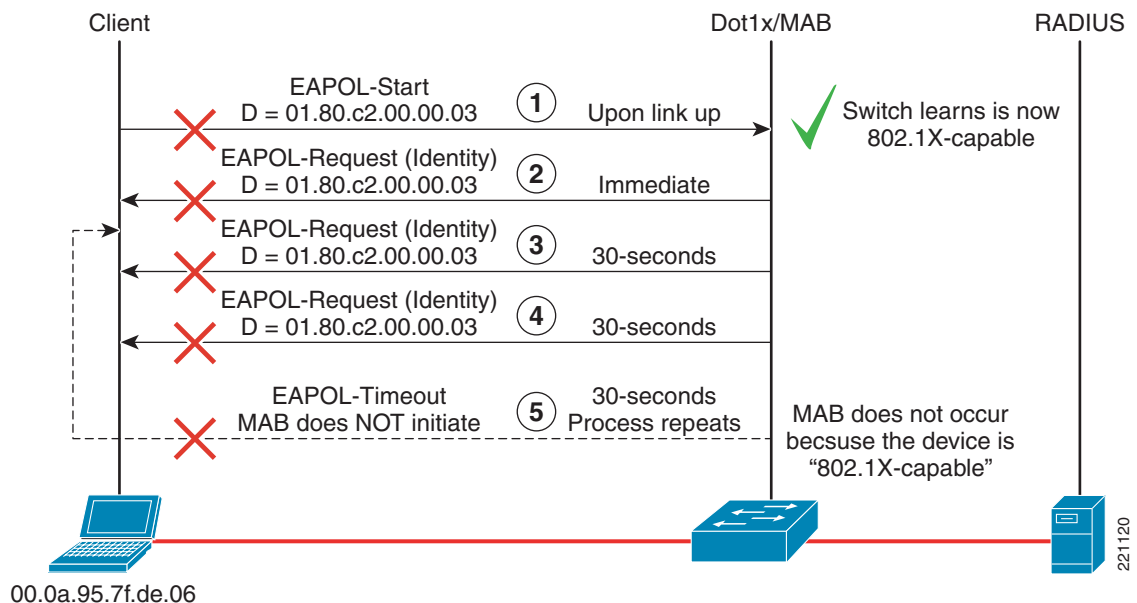
Fallback Technique for Re-imaged Machines with CSSC

A supplicant such as CSSC may be provisioned as part of a standard machine build. If the supplicant then sends an EAPOL-Start frame, with no existing certificate or temporary credential, 802.1x initiates and the 802.1x process times out. In addition, this process continues persistently and MAB never executes. This is because the client sent an EAPOL-Start frame to the switch initially, and a switch uses EAPOL to determine whether the device is supplicant capable (so 802.1x is tried always).

A recommended practice is to integrate with the provisioning process of the re-image of machines to disable 802.1x upon first boot unless 802.1x credentials can be built into the imaging process itself, such as one-time or temporary credentials to 802.1x authenticate, just to be able to attain appropriate network access for the purposes of downloading true user or device credentials.

For some cases, this may not always be the case in how the provisioning process occurs, especially because of re-imaged machines. Figure 49 shows an example of a Windows or Cisco Secure Services Client (CSSC) supplicant enabled for 802.1x that sends EAPOL-Starts and does not have prior certificate credentials.

Figure 49 802.1x and EAPOL-Starts Enabled without Credentials



This would be the same behavior if any other supplicant sent an EAPOL-Start and a screen was displayed to the user to input credentials for a challenge response-based EAP type such as PEAP. If the user does not respond to the credential notification, 802.1x times out and repeats transitively, and MAB is not initiated for these types of cases as well. Again, when a switch knows an 802.1x supplicant is on the wire through the device speaking EAPOL, MAB or the Guest-VLAN can not typically be leveraged. The only exception to the process indicated above is a global configuration available in IOS-based switches. Starting in 12.2(20)SE for Catalyst 3000 Series switches, the command is `dot1x guest-vlan supplicant`. In addition, this command has become hidden starting from the releases 12.2(31)SG for Catalyst 4500 and 12.2(25)SEE for Catalyst 3750. As of 12.2(35)SE, this command is still functional, but remains hidden as well. This command causes the EAPOL history not to be retained by the switch, so that after

the above process goes through at least once, the state machine continues to run and eventually the Guest-VLAN or MAB can be enabled on the port if configured. This serves as a workaround if a situation such as the above is encountered. There is no such workaround available in CatOS.

Network Admission Control

There is an inconsistent understanding of what MAB is in relation to NAC deployments. In addition, unless Network Admission Control (NAC) is needed for EAP over UDP (EoU), all the configurations provided in this section may not be needed. Another systemic challenge with MAB has been caused by the NAC program. The base functionality for MAC address authentication was already in place on ACS. For example, wireless APs and controllers can initiate a PAP authentication with a RADIUS server using a client MAC address as a username/password.




Note

For more information about MAC-Auth for WLAN devices, see the following URL:
http://www.cisco.com/warp/public/779/smbiz/prodconfig/help/eag/123-04.JA/1100/h_ap_howto_5.htm

The second part of this document is for local authentication (AP as a AAA server). A customer would just need to change from authentication local to ACS server.

Figure 50 shows what is achieved by the wireless implementation of MAC address checking.

Figure 50 **Origin of MAC Authentication**



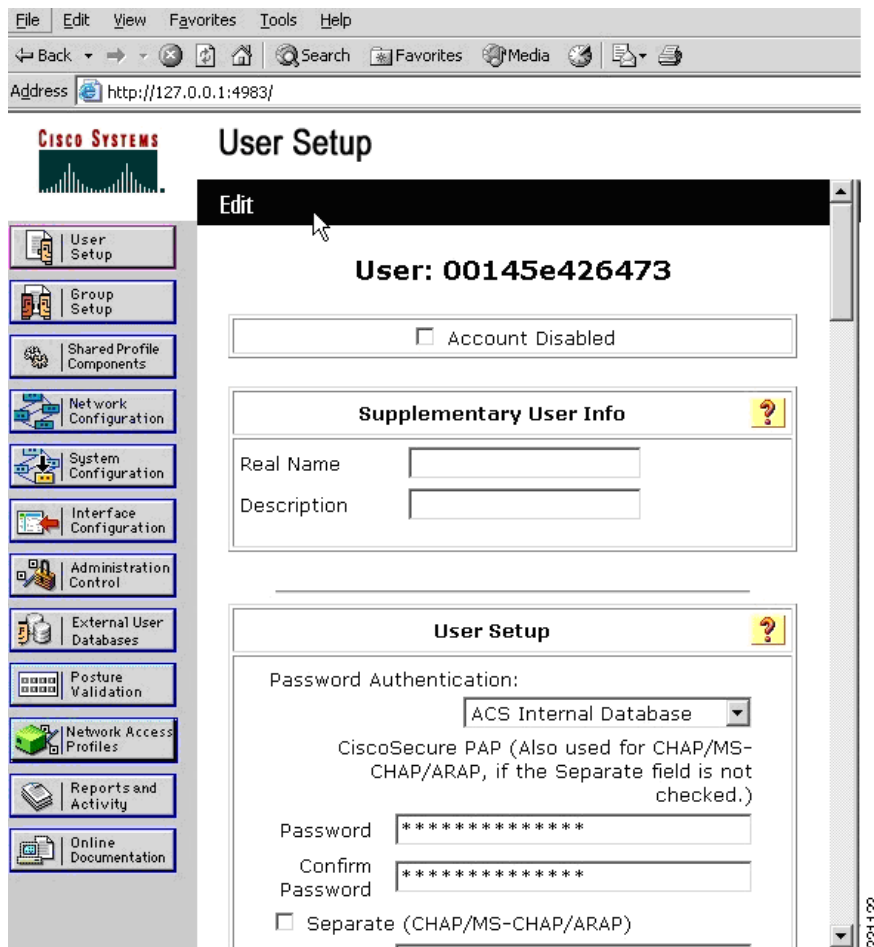
In an open authentication environment (or separate SSID), wireless can at least check the MAC address of an end client as an admission criteria. Because of client ubiquity, these topologies are no longer typically recommended. APs can accomplish this because initial associations have already been made (and based on that association traffic to/from a wireless NIC is blocked by the AP).

No such associations exist currently in the wired space, but MAB on switches provides a wired equivalent to the wireless functionality of checking a MAC address after an 802.1x timeout on a port. MAB on a wired switch is geared to satisfy the same requirement as the previous wireless example, the only difference being that on IOS-based switches, 802.1x must timeout before MAB is attempted as a supplemental authentication method. This means MAB can be used to deal with devices from an exception perspective only. In other words, MAB is used as an authentication technique in the absence of a preferred credential. This is *not* to be used as an exemption criteria, or as a way to deal with failures of the preferred credentials.

MAB is primarily geared toward providing a supplemental authentication technique in the absence of 802.1x. It should be simple by default. A minimum requirement is a database full of MAC addresses. Any database (including the local CiscoSecure database on ACS) supporting authentication requests to RADIUS via PAP or EAP-MD5 can suffice. MAC addresses can be imported/configured, with no punctuation, white spaces, or capitalization in these databases. This really just means the ability to

authenticate a MAC as the username/password combo from a RADIUS (PAP or EAP-MD5) exchange. It was designed to be no different than what APs have been doing on this for years. Figure 51 shows an example of a single user account being added to a local CiscoSecure database as a MAC address.

Figure 51 User Account Added to the ACS Local Database



Assuming default use cases, this configuration should allow ACS to successfully authenticate devices as switches proxy the authentication requests on behalf of the devices via MAB. No other special configuration is required to enable MAB for ACS other than what may have already been enabled for 802.1x. This functionality should work for ACS revisions back to ACS v2.6. For more information, see the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/access/acs_sofi/csacs4nt/csnt26/usergd26/userdb.htm#xtocid171961

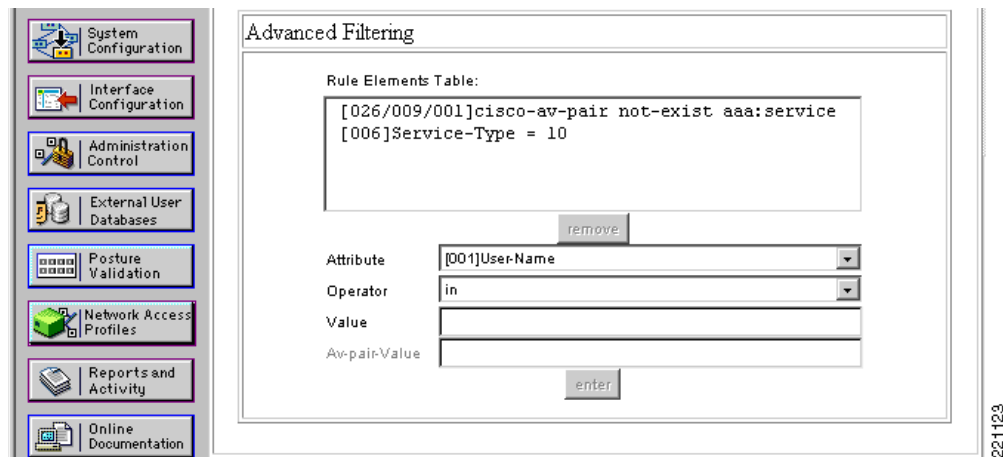
The service-type attribute from RADIUS is set in a MAB request from all switches and wireless APs/controllers as well. The calling-station-ID (CLID) attribute is also set to the MAC address in a MAB request from a switch and WLAN controller. However, NAC technologies such as EOU attempt to use the service-type attribute as a means to differentiate a request to be able to do agentless processing, and in some cases, spawn audit functionality.

From the documentation above, the service-type attribute is critical to enable NAPs on ACS for NAC, and is critical for service-type processing. This works fine for EOU requests. In addition, EOU requests (agentless or otherwise) do not contain user or identity information based on the current operational behavior of EOU on all switches and routers. However, by default, both service-type and CLID attributes

can effectively be irrelevant for MAB itself as the authentication event described in this document. The authorization based on the authentication event is rather typical as well, because it is expected to be user-name/password-based. The original ability to deal with a NAC agentless host for EOU was also termed “MAC Authentication Bypass” (the reasoning behind this is unclear).

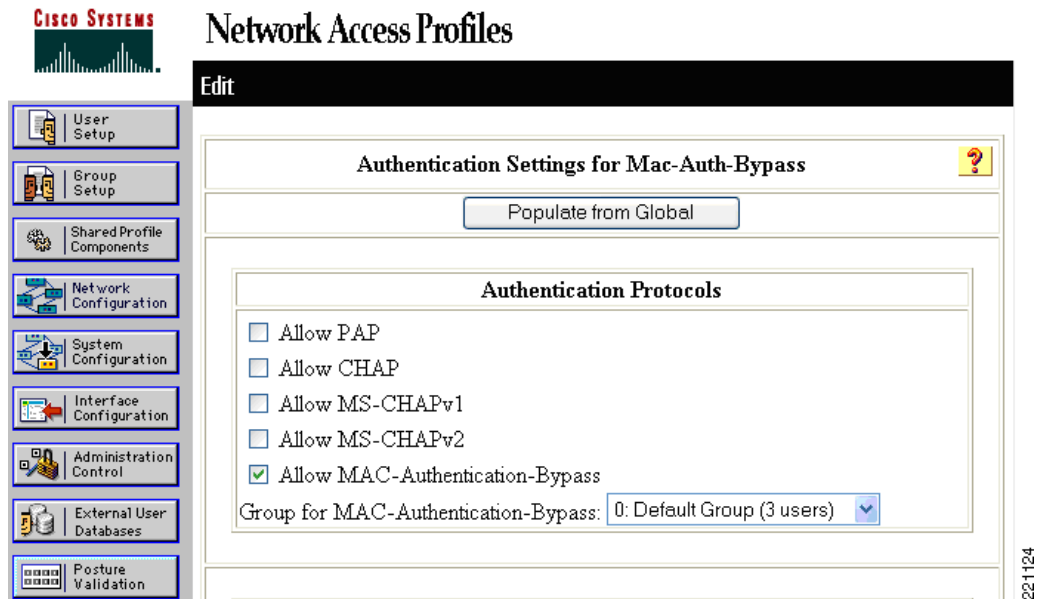
Per the documentation references above, a NAP that matches a RADIUS request should be set up and configured. It is expected that customers will create a new NAP to provide the accurate selection of the service profile for “MAC Authentication Bypass”. In the following example, a NAP with the name “MAC-Auth-Bypass” is created as demonstrated in the NAC configuration guide. However, any MAB request from the switch via RADIUS also has a service-type value of 10 just like NAC-L2-IP, but does not have a Cisco Attribute Value Pair (AVP) containing the keywords service (which is a VSA sent by EOU). This means any MAB request matches this NAP, whether intentional or not. Shown in [Figure 52](#) within the “Advanced Filtering” box are two entries; no match for “aaa:service” in the Cisco AVP and the “Service-Type” equal to 10.

Figure 52 NAP → Advanced Filtering



When this advanced filtering is defined, the authentication settings of the NAP must be changed (to enable “MAC Authentication Bypass” and a group or a pre-configured Network Access Restriction (NAR) for authorization purposes should then be selected. In the Authentication settings for a NAP, the check box next to “Allow MAC-Authentication-Bypass” should be selected, in addition to the selection of a group that has the correct NAR; in the example shown in [Figure 53](#), the group default Group.

Figure 53 Assigning “MAC-Authentication-Bypass” to a NAP



No other settings need to be changed in the NAP for “MAC Authentication Bypass” to function properly. However, this functionality is *not* MAB, as previously shown. It is effectively another feature entirely, because it looks only at the Calling-Station-ID field from RADIUS for authorization purposes.

The difference in this processing model is that the Calling-Station-ID field may be interrogated at any time as exemption criteria from any control plane (802.1x, PPP, and so on) to which the method is applied. In other words, MAB could be used as an authorization technique regardless of the existence of a preferred credential.

This functionality is potentially valuable in NAC use cases to support exemptions for EOU or agentless requests from EOU. For example, an EOU request for NAC posture can be completely bypassed even if Cisco Trust Agent (CTA) is loaded on a device, simply because of the MAC address of the end station. Alternatively, after EOU times out on a port in an effort to search for the CTA for NAC, an agentless request is then made with no user information, and the expectation is that the CLID field can be interrogated to determine the device access, or to start a NAC audit for the device in question. However, this is not the same feature set or functional goal of MAB as described in this document.

For more information, see the configuration guide for NAC at the following URL:

http://www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/413/cam/413_cam_book.html.



Note

Although ACS 4.1 is not within the scope of this document, the user interface shown above changed in the 4.1 release. It now appears in the GUI as “Agentless Request Processing” instead of “Allow MAC-Authentication-Bypass” under the NAP protocols page, as shown in [Figure 54](#).

Figure 54 Changes in ACS 4.1 GUI

Authentication Protocols	
<input type="checkbox"/>	Allow PAP
<input type="checkbox"/>	Allow CHAP
<input type="checkbox"/>	Allow MS-CHAPv1
<input type="checkbox"/>	Allow MS-CHAPv2
<input checked="" type="checkbox"/>	Allow Agentless Request Processing

The GUI changes should at least help alleviate confusion in this respect. NAPs are not needed to configure MAB. Therefore, when doing MAC address based exceptions such as this for NAC, ACS is looking at the Calling-Station-ID attribute from RADIUS, and not the username in the RADIUS request.

This issue is further complicated by the authentication processing of authentication requests of ACS v4.0. First, with ACS v4.0, if a NAP is *not* matched, a MAB session is still able to authenticate successfully as long as there is an actual user account in the database as demonstrated before the NAC complications here was introduced. The reason for this is that ACS reverts to legacy behavior for any session that does not match a NAP or the specific authentication settings in a NAP. Legacy behavior is also attempted if a NAP is not defined at all, of course, although this is available as a configurable condition. However, today with ACS 4.0, a NAP can actually be matched, and both behaviors can actually be made to work. The reason for this is actually because of a bug in ACS processing, because only the CLID field should be interrogated for authorization purposes. This can also be verified in a test such as one shown in [Figure 55](#).

Figure 55 Successful “MAC-Authentication-Bypass” that Matches a NAP

Date	Time	Message-Type	User-Name	Group-Name	NAS-IP-Address	Access Device	NAS-Port	Network Access Profile
02/12/2007	20:44:13	Authen OK	00145e426509	Default Group	10.137.61.4	id1-6503-1	130	IBM-NICs

As seen in [Figure 55](#), this is working “the legacy way” because the switch sends “00-14-5E-42-65-09” as the format of the Calling-Station-Id while the format of the user-name from the switch is “00145e426509” and is what is in the passed authentications log. In addition, the only other configuration in ACS for this case was a user account equal to “00145e426509”. This section has shown that a PAP request for MAB can be successful even though a NAP is matched. This is technically a bug in ACS v4.0 because it should not work.

As a best practice, if customers need only MAB to work (such as a URT replacement, or supplemental authentication technique for 802.1x), and they are not interested in NAC yet, it is not recommended to attempt to use a NAP for this by default in ACS.

MAB EAP Option on ACS

A wired IOS-based switch also has an EAP option available for MAB. It is configured via the following configuration example:

- Cisco IOS

```
interface FastEthernet0/1
  switchport access vlan 2
  switchport mode access
  dot1x mac-auth-bypass eap
  dot1x pae authenticator
  dot1x port-control auto
  spanning-tree portfast
  spanning-tree bpduguard enable
```

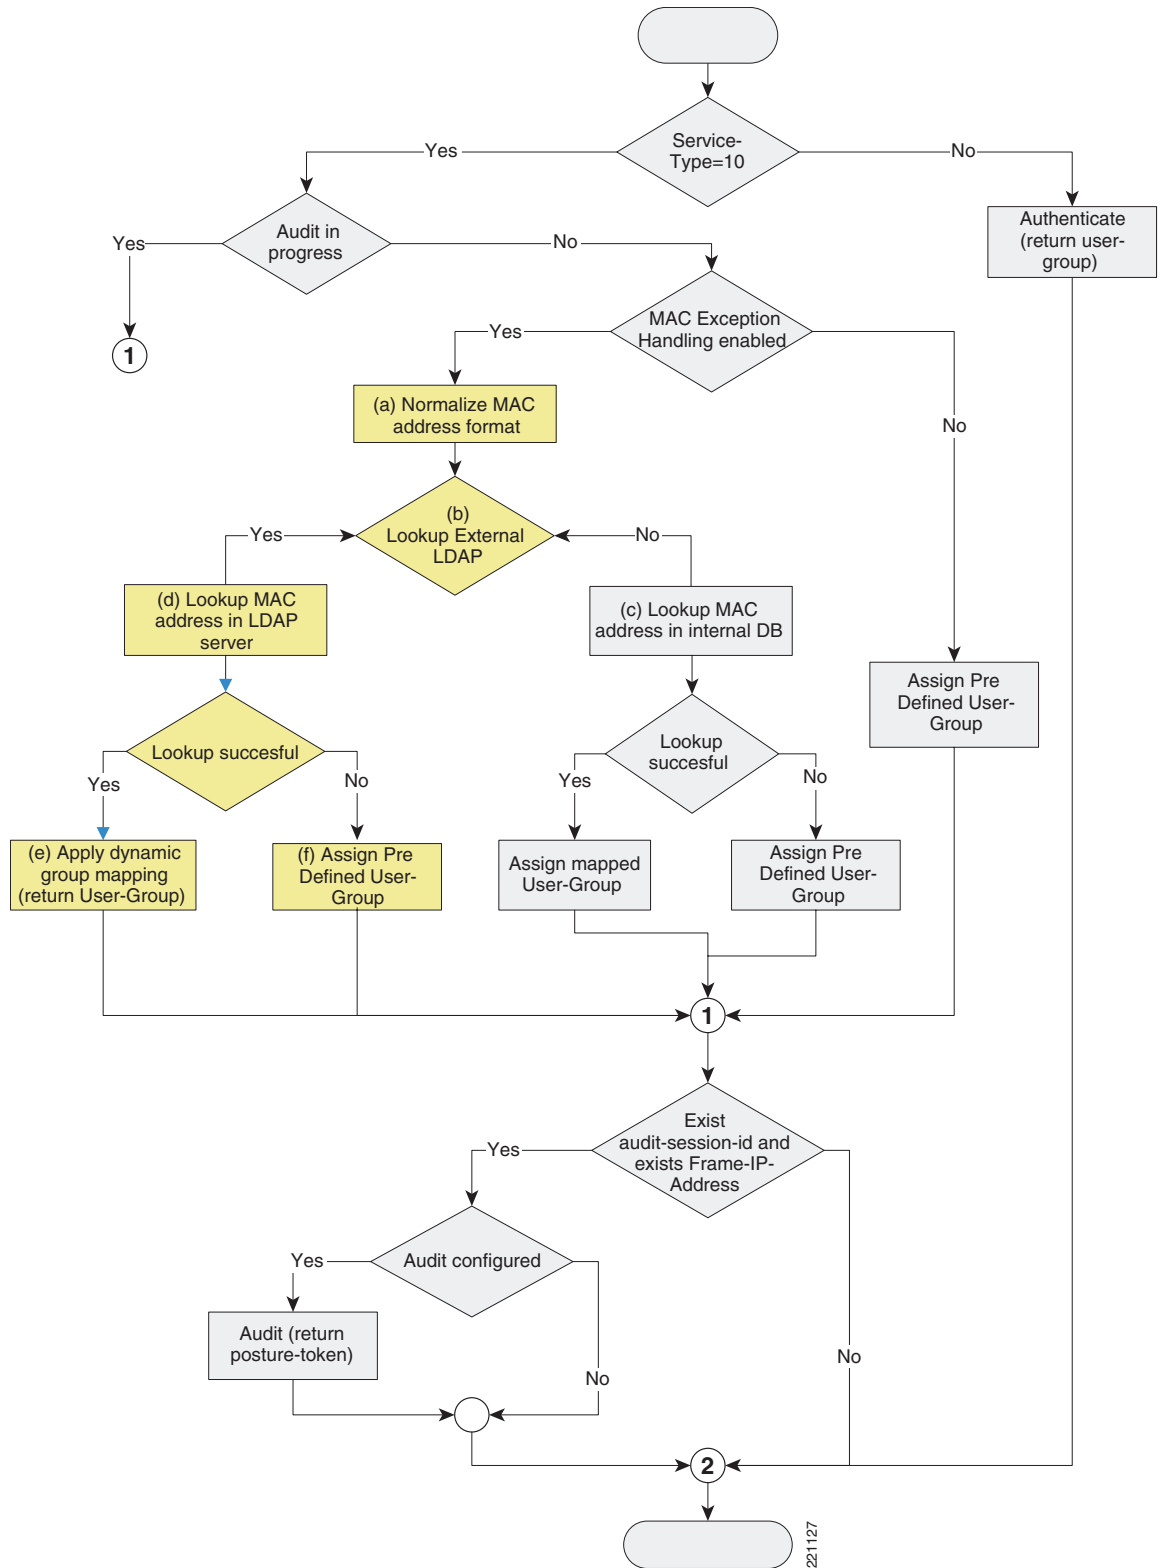
With the configuration above, the RADIUS exchange for MAB is not carried out in a single data exchange as it is with PAP. It is carried out in a challenge response mechanism with EAP-MD5 and makes a round trip for RADIUS before succeeding. The username and password for the request remain the MAC address. The MD5 option may be used in conditions where PAP does not meet a security policy, or PAP may have been disabled on the backend database server.

However, if you configure the same process as indicated in the previous section for ACS and a NAP, this does not work. In this condition, MAB neither passes or fails (see CSCse12673). Thus, if EAP-MD5 is needed on switches for MAB for any reason, this ACS configuration and the one for NAC are not compatible.

ACS 4.1

No form of MAB as explained in this document can currently be used with ACS 4.1. This includes all wired switches that run MAB as a supplemental authentication technique described here, or for any MAC authentication done by wireless APs/controllers. This is because of a change in processing for ACS 4.1 to better handle NAC and auditing capability. (See [Figure 56](#).)

Figure 56 ACS Processing as of ACS 4.1



As shown in [Figure 56](#), authentication cannot be performed if the service-type attribute is set. It is unclear what the purpose of an “Audit in Progress” even means, especially if the customer is not running NAC. There are more considerations not included in this document related to these changes for ACS 4.1. Although ACS 4.1 is not within the scope of this document, MAB cannot be used for this ACS revision at all, and is not recommended. CSCsh62641 and CSCsh63236 have been filed to account for this, however. The current plan is to fix these issues in ACS v4.1.3 (FCS by end of April 2007).

Provisioning

Provisioning is also a service of high concern to customers. A customer may not know what their MAC addresses are in advance. In addition, no turnkey solution is provided by Cisco to fill this void. Third-party products that provide asset management capabilities may help in this regard, such as products from Great Bay Software, Altiris, and so on.

Some customers have attempted to integrate learning techniques with their directory infrastructure. For example, a Cold Fusion front end can be used to force users to authenticate with Active Directory credentials. The front end then pulls the MAC address, host name, and user/machine details and puts them in an ODBC database. This is not only a potential MAC provisioning technique, but also a nice compromise of identity- and machine-based authentication without the complexities of 802.1x if the security model does not call for 802.1x.

However, the deployment of MAB itself can help elicit a provisioning mechanism. In addition, devices can be granted network access as well. An example of this is to use MAB along with the Guest-VLAN. Fundamentally in this scenario, a machine incapable of 802.1x always ends up in the Guest-VLAN. MAB does not necessarily change this, by the Guest-VLAN serving as a failure condition for MAB itself. Therefore, ultimately, a device can get into the Guest-VLAN much the same as it does without MAB, because it is incapable of 802.1x. However, if MAB fails “in the middle”, a failure of this event should be recorded on the AAA server. An example from ACS of this failure is shown in [Figure 57](#).

Figure 57 MAB Failure

The screenshot shows a web interface for viewing failed MAB attempts. At the top, there are filters for 'Regular Expression', 'Start Date & Time', and 'End Date & T'. Below the filters, there are 'Apply Filter' and 'Clear Filter' buttons. The main content is a table with the following data:

Date ↓	Time	Message-Type	User-Name	Group-Name	Caller-ID	Network Access Profile Name	Authen-Failure-Code	Author-Failure-Code
02/12/2007	21:00:27	Authen failed	00145e426509	Default Group	00-14-5E-42-65-09	(Default)	CS user unknown	..

As shown in [Figure 57](#), now the MAC address is effectively known to the authentication infrastructure. This MAC can now be potentially inserted into an asset management system or a primary directory infrastructure through various techniques.



Note

In-depth guidance on identity management is beyond the scope of this design guide.


However, remember that the gathering of MAC addresses does not extend trust explicitly. LMS from CiscoWorks can also help as a MAC address gathering tool. It also performs device name, IP address, and host name correction. However, none of these techniques necessarily ensure that the entity should be on the corporate network to begin with; they may only prove that it is already there. More work should be done for the verification of network MAC addresses to validate existing identified trusted machines.

Lack of Existing Identity Store

Specific MAC addresses are likely unknown to large enterprises. If they are known, they may not be incorporated directly into an existing directory infrastructure; they may be located only in an asset or inventory management system. For this management system to be used for authentication, it must be able to be interrogated by AAA, or the MAC addresses must be exported to a system that can be interrogated by AAA. For MAB, this means virtually any backend database into which ACS already has access. The identity store can be added onto, however. MAC addresses can be stored as user accounts on Windows Active Directory. The CiscoSecure ACS database can store MAC addresses as well. The IBM Tivoli agent can add/remove MAC addresses in an ACS NAP. If MAC addresses are being defined as users in ACS, in ACS 4.0, the limit is 300,000 entries.

Lack of Voice Support

The integration of 802.1x, MAB and IP phones is based on the switch configuration of multi-VLAN access ports. Multi-VLAN ports belong to two VLANs: native VLAN (PVID) and auxiliary VLAN (VVID). This allows the separation of voice and data traffic and enables 802.1x and MAB only on the PVID. The type of communication that occurs on these two VLANs is shown in [Figure 58](#).



When 802.1x or MAB is enabled on a multi-VLAN access port, a client must complete the authentication process before getting access to the data (native/PVID) VLAN. The IP phone can get access to the voice (auxiliary/VVID) VLAN after sending the appropriate Cisco Discovery Protocol (CDP) packets, regardless of the 802.1x state of the port. The use of CDP with Cisco IP phones may be required, given the lack of pervasive support for an embedded 802.1x supplicant.

The configuration commands for Cisco IOS and CatOS that are required to enable multi-VLAN functionality, in conjunction with 802.1x and MAB, are as follows:

- Cisco IOS

```
interface FastEthernet0/1
  switchport access vlan 2
  switchport mode access
  switchport voice vlan 2
  dot1x mac-auth-bypass
  dot1x pae authenticator
  dot1x port-control auto
  spanning-tree portfast
  spanning-tree bpduguard enable
```

- CatOS

```
set vlan 2 2/1
set port dot1x 2/1 port-control auto
set port mac-auth-bypass 2/1 enable
set port auxiliaryvlan 2/1 2
set spantree portfast 2/1 enable
set spantree bpdu-guard 2/1 enable
```

Therefore, although MAB can be configured on a port and be used to authenticate data device, customers should not use MAB in an attempt to authorize voice devices on a Voice-VLAN. MAB is designed at the moment to authorize devices on data VLANs only and support VLAN assignment. If the MAC address of a phone is provisioned in ACS and it sends out packets, the switch is able to glean the MAC address and begin authorization to grant the phone access into the network on the data VLAN (or VLAN assigned from RADIUS). A switch does not know or pre-suppose the type of device and does not know to put it on the voice VLAN as part of the authentication event, however. Thus, if the customer provisioned the phone to tag its packets on the voice VLAN, it fails as of today, because traffic on voice VLANs for MAB is explicitly ignored. Therefore, a customer cannot use MAB to attempt to authenticate a third-party phone. A potential workaround is to dynamically assign a data VLAN via RADIUS and MAB equal to a voice VLAN without the voice VLAN configured on the switch port. However, this is not recommended because single-auth mode would not allow any other MAC on the wire such as a client plugging into a phone. In essence, MAB shares the same rules in this space that 802.1x does.

As discussed in the documentation reference above, MAB can be enabled for data devices, and Cisco telephony devices can be ignored with CDP. However, similarly to 802.1x, MAB-authenticated session may disappear from the network without the network knowing about it explicitly. A client disconnecting from the back of an IP phone is not recognized as an event by the switch. The first problem with this behavior is that when a host disconnects from the phone, the host remains authorized on the switch port. In addition, for any new machines that plug into the phone, a security violation may be tripped, because the phone thinks another MAC has appeared on the wire other than the one it has authenticated. Catalyst 3000 switches recently delivered a MAB aging feature to address this in 12.2(35)SE, but could not be verified for this phase of network virtualization.

Further integration with IP Communications is planned for a later phase of network virtualization, which will examine Multi-Domain-Auth (MDA) and MAB aging. MDA is a new solution-based feature set that allows any phone to authenticate via 802.1x or MAB, and is also able to authenticate a client plugging in behind an IP phone via 802.1x or MAB starting with Catalyst 3000 switches in 12.2(35)SE, and 4500 switches in 12.2(37)SG.

MAC Movement

Like 802.1x, a MAC address authenticated by MAB is not allowed to move on a switch unless the port from which the device moved is unauthorized. This issue is exacerbated by the MAB aging issue introduced in the previous section with respect to IP telephony. Therefore, if a device is authenticated via MAB behind a phone and then moves to another port on the same switch, the port to which the user moved is err-disabled. This renders the phone on that port inoperable as well. With CatOS, there is a configurable nature for security violation behavior handling to restrict traffic from an offending MAC instead of shutting the port down. However, even this does not help in this case. This violation behavior handling would help only for the appearance of a second MAC address on the original port, not for the movement of the MAC address to begin with. This is typically not an issue for a MAB port with no IP telephony because the move drops link on the port and clear the binding of the address to MAB. This issue may persist in a hub-based topology, though this is not a supported design. In addition, in CatOS today, a port must be manually reset when this event occurs. There is no auto reset after a configurable interval.

MAC Authentication Bypass Policy Assignment

Based on the consistent architecture MAB promotes along with 802.1x, MAB can automatically leverage any specialized policy enforcement techniques that may already be available to 802.1x. Especially important to network virtualization is dynamic VLAN assignment via RADIUS. No special configuration on a switch is needed to achieve dynamic VLAN assignment.

Standard recommendations for 802.1x with VLAN assignment remain with MAB. It is highly recommended to plan and build out any supporting VLAN architecture in advance. VLAN assignment is done by name with MAB like it is with 802.1x. This can support flexible VLAN management techniques for various L2 or L3 VTP architectures, allowing for independence between separate L2 domains. The architecture also allows for policies to be applied to groups or down to a per-device level. Depending on the specialized need, MAB may be managed on a per-host basis like this in some cases.

Remember on IOS-based switches to make sure you enable AAA and specify the authentication and authorization methods:

```
aaa new-model
aaa authentication dot1x default group radius
aaa authorization network default group radius
```



Note

RADIUS attributes received in CatOS are automatically implemented if 802.1x is enabled. However, this is *not* the case for IOS. This is why you need the last configuration statement above, for the switch to accept configuration commands via RADIUS.

As mentioned above, none of the above applies to CatOS platforms, and these configuration steps are not needed by default. However, VLAN assignment, can be optionally disabled via the following configuration:

```
id1-6503-1> (enable) set dot1x radius-vlan-assignment ?
  disable          Disable dot1x Radius Vlan Assignment on the system
  enable           Enable dot1x Radius Vlan Assignment on the system
```

Nothing is needed on the ACS server, outside of what may already be in place for 802.1x as well. The following three standard RADIUS attributes defined by RFC 2868 are required:

```
[64] Tunnel-Type - "VLAN" (13)
[65] Tunnel-Medium-Type - "802" (6)
[81] Tunnel-Private-Group-ID - <VLAN name>
```

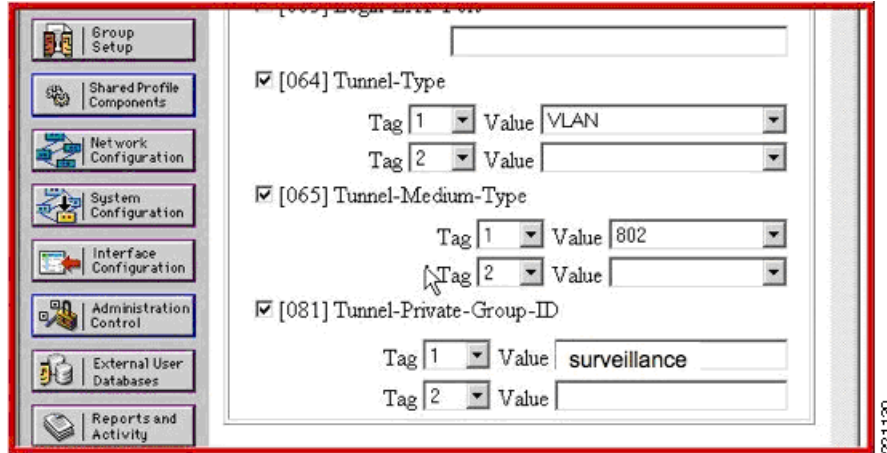


Note

Before ACS 4.0, these features were viewable by default. To enable group-level viewing, they needed to be viewed under the “RADIUS (IETF)” link under the Interface Configuration configuration button. There are check boxes for each attribute. With ACS 4.0, however, this configuration step is not needed, and the attributes are enabled by default via per-user, or a per-group deployment scenario.

[Figure 59](#) shows an example of configuring a certain group of devices for MAB to be deployed into the “surveillance” VLAN.

Figure 59 VLAN Assignment Configuration on ACS



This enables any user members of the group configured for VLAN assignment to be assigned into the named VLAN. The VLAN name must be present on the switch, and be the identical name of the configuration in ACS. This includes white spaces and capitalization. The VLAN must exist on the switch as well. If any of these are not valid, a switch denies authorization. The user may provide a credential authorizing the user to access the network on a VLAN. However, if the switch cannot verify the information about the VLAN itself (though any sort of VLAN name mismatch, type-o, and so on), a switch treats this as a user not in fact providing valid credentials.

The VLAN name is mapped to a VLAN number. Upstream, path isolation uses separate VLANs as entrance criteria into each separate network partition. With wireless, you may also optionally ensure the original request originated on the correct SSID to ultimately map a session into the correct VLAN.

By leveraging dynamic policy enforcement, this completes the ability of an enterprise to differentiate between clientless sessions on the network. Previously, network virtualization was incapable of leveraging this differentiation capability. Network virtualization could differentiate between client contexts with 802.1x, but could only default to providing a de facto level of access if 802.1x was not resident on an end device. By having MAB and policy enforcement available, network virtualization can now be expanded to included differentiated services among robotic arms on a factory floor, x-ray machines in a hospital, IP-enabled surveillance devices, or standard corporate PCs. This increases the end-to-end impact network virtualization provides with this additional, fine-grained access control.

MAC Authentication Bypass Summary

In summary, MAB functions as a port-based feature. It is primarily used as a fallback mechanism to 802.1x, although it is optionally available as a standalone authentication method with CatOS. There is no de facto ability to support more than one MAC per port. MAB is single host in nature just like 802.1x, and there is no multi-auth for MAB. A MAB port can be optionally enabled for multi-host mode just as is done with 802.1x. MAB cannot be used as a means to deal with failed 802.1x authentication attempts. MAB provides customers who do not or cannot do 802.1x, but who have also bought into port security with configured MAC addresses more options, and provides a migration path to customers running URT or VMPS technologies. MAB also works with any standard RADIUS server, with a default timeout of 30 seconds with three retries. This means that the total timeout period is at least 90 seconds by default, which is the same minimum default timeout of the Guest-VLAN. A device must also send traffic into a switch for the MAC to be learned after the 802.1x timeout. If MAB fails, network access is implicitly denied. If MAB fails and the Guest-VLAN is also configured, the Guest-VLAN is enabled (for backward compatibility). Additional network policy can be downloaded as well. This supports dynamic

virtualization, and the least common denominator is what 802.1x can currently do for the switch in question. A provisioning mechanism is not called for by MAB, although the Guest-VLAN can be used to assist in this process.

Overall Summary

With the increasing demands upon modern networks and the need to share information not only within an organization but as well with vendors and customers, security along with network access have become the top priority. The IEEE 802.1x specification for port-based network control has become the standard method for Layer 2 authentication access, not only with wireless but with the wired ports as well. 802.1x is a core technology component in support of access-control to promote end-to-end network virtualization.

One challenge in wired topologies is how to support failed authentications, especially for the 802.1x-enabled guest or partner. The Auth-Fail-VLAN is a way to grant access to 802.1x-enabled guest or partners, while providing differentiated access from internal machines to promote network virtualization. Additional challenges in implementing IEEE 802.1x are the requirements to support the cutting edge of yesterday, which is now the legacy of today. Most legacy devices, such as printers, VoIP phones, and new emerging devices such as IP security cameras, do not have the ability to support an 802.1x supplicant but must be included in any pervasive virtualized network architecture. MAC Address Authentication Bypass (MAB) is not meant to replace 802.1x; rather it is meant to allow an alternate means of authentication when a host or device does not respond to requests for credentials by network access devices.

The IEEE 802.1x standard and MAB allows the dynamic configuration of access ports as well as implementing the corporate security policy on the port level. MAB addresses the difficulty of deploying an 802.1x infrastructure throughout an enterprise network. An 802.1x supplicant is required to authenticate to an authentication server via a network access device. The MAB feature allows devices without this 802.1x capability to access the network and perform their desired function while allowing L2 authentication to occur and participate in the dynamic deployment of network policy to promote a virtualized network architecture.

The Guest-VLAN is also an option for devices incapable of 802.1x. By combining MAB and the Guest-VLAN, an enterprise can now differentiate between clientless stations in support of device-specific access control as an application of network virtualization. In addition, the access control methods described in this document provide multiple levels of user access, making it the first element of network security. In addition, by supporting end-to-end network virtualization, these levels of access can take on more of a matrix model, with departmental and divisional roles dictating where access can be applied in the future.

Network virtualization can help reduce overall risk, add value, and remove operational cost from the business because of its logical network overlays while promoting security. Enterprises should develop corporate strategies now to control network access. Future challenges such as closed user groups, dynamic project creation, or even more on-demand device mobility can then be evaluated with the network as a direct enabler of business services.

