MAC Authentication Bypass

Deployment Guide

May, 2011

© 2011Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Contents

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

1. Introduction

The need for secure network access has never been greater. In today's diverse workplaces, consultants, contractors, and even guests require access to network resources over the same LAN connections as regular employees, who may themselves bring unmanaged devices into the workplace. As data networks become increasingly indispensable in day-to-day business operations, the possibility that unauthorized people or devices will gain access to controlled or confidential information also increases.

The best and most secure solution to vulnerability at the access edge is to use the intelligence of the network. One access control technique that Cisco provides is called MAC Authentication Bypass (MAB). MAB uses the MAC address of a device to determine what kind of network access to provide.

This document focuses on deployment considerations specific to MAB.

To learn more about solution-level uses cases, design, and a phased deployment methodology, see http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6638/whitepaper_C11-530469.html.

For step-by-step configuration guidance, see http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6638/Whitepaper_c11-532065.html.

2. About MAB

2.1 Benefits and Limitations

MAB offers the following benefits on wired networks:

- **Visibility:** MAB provides network visibility since the authentication process provides a way to link a device's IP address, MAC address, switch, and port. This visibility is useful for security audits, network forensics, network use statistics, and troubleshooting.
- **Identity-based services:** MAB enables you to dynamically deliver customized services based on an endpoint's MAC address. For example, a device might be dynamically authorized for a specific VLAN or assigned a unique access list that grants appropriate access for that device. All the dynamic authorization techniques that work with IEEE 802.1X authentication will also work with MAB.
- **Access control at the edge:** MAB acts at Layer 2, allowing you to control network access at the access edge.
- **Fallback or standalone authentication:** In a network that includes both devices that support and devices that do not support IEEE 802.1X, MAB can be deployed as a fallback, or complementary, mechanism to IEEE 802.1X. If the network does not have any IEEE 802.1X–capable devices, MAB can be deployed as a standalone authentication mechanism.
- **Device authentication:** MAB can be used to authenticate devices that are not capable of IEEE 802.1X or that do not have a user.

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

MAB enables visibility and security, but it also has limitations that your design must take into account or address:

- **MAC database:** As a prerequisite for MAB, you must have a preexisting database of MAC addresses of the devices that are allowed on the network. Creating and maintaining an up-to-date MAC address database is one of the primary challenges of deploying MAB.
- **Delay:** When used as a fallback mechanism to IEEE 802.1X, MAB waits for IEEE 802.1X to time out before validating the MAC address. During the timeout period, no network access is provided by default. Delays in network access can negatively affect device functions and the user experience. A mitigation technique is required to reduce the impact of this delay.
- **No user authentication:** MAB can be used to authenticate only devices, not users. Different users logged into the same device will have the same network access.
- **Strength of authentication:** Unlike IEEE 802.1X, MAB is not a strong authentication method. MAB can be defeated by spoofing the MAC address of a valid device.

2.2 Functional Overview

2.2.1 What Is MAB?

MAB enables port-based access control using the MAC address of the endpoint. A MAB-enabled port can be dynamically enabled or disabled based on the MAC address of the device that connects to it. Figure 1 illustrates the default behavior of a MAB-enabled port.

Figure 1. Default Network Access Before and After IEEE 802.1X

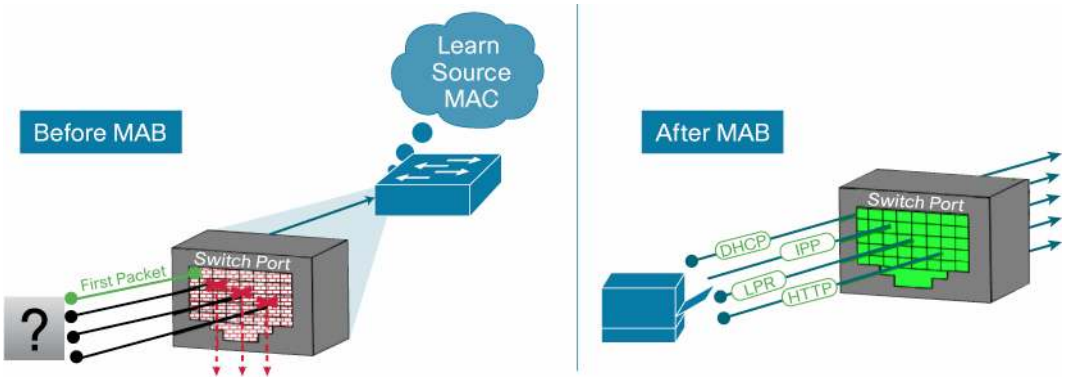

Prior to MAB, the endpoint's identity is unknown and all traffic is blocked. The switch examines a single packet to learn and authenticate the source MAC address. After MAB succeeds, the endpoint's identity is known and all traffic from that endpoint is allowed. The switch performs source MAC address filtering to help ensure that only the MAB-authenticated endpoint is allowed to send traffic.

MAC address authentication itself is not a new idea. An early precursor to MAB is the Cisco® VLAN Management Policy Server (VMPS) architecture. With VMPS, you create a text file of MAC addresses and the VLANs to which they belong. That file gets loaded into the VMPS server switch using the Trivial File Transfer Protocol (TFTP). All other switches then check with the VMPS server switch to determine to which VLAN those MAC addresses belong. MAB represents a natural evolution of VMPS. Instead of storing MAC addresses on a VMPS server switch, MAB validates MAB addresses that are stored on a centralized (and thus more easily managed) repository and that can be queried using the standard RADIUS protocol.

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

2.2.1.1 High-Level Functional Sequence

The high-level functional sequence in Figure 2 illustrates the way that MAB works when configured as a fallback mechanism to IEEE 802.1X. If IEEE 802.1X is not enabled, the sequence is the same except that MAB starts immediately after link up instead of waiting for IEEE 802.1X to time out.

Figure 2. High-Level MAB Sequence

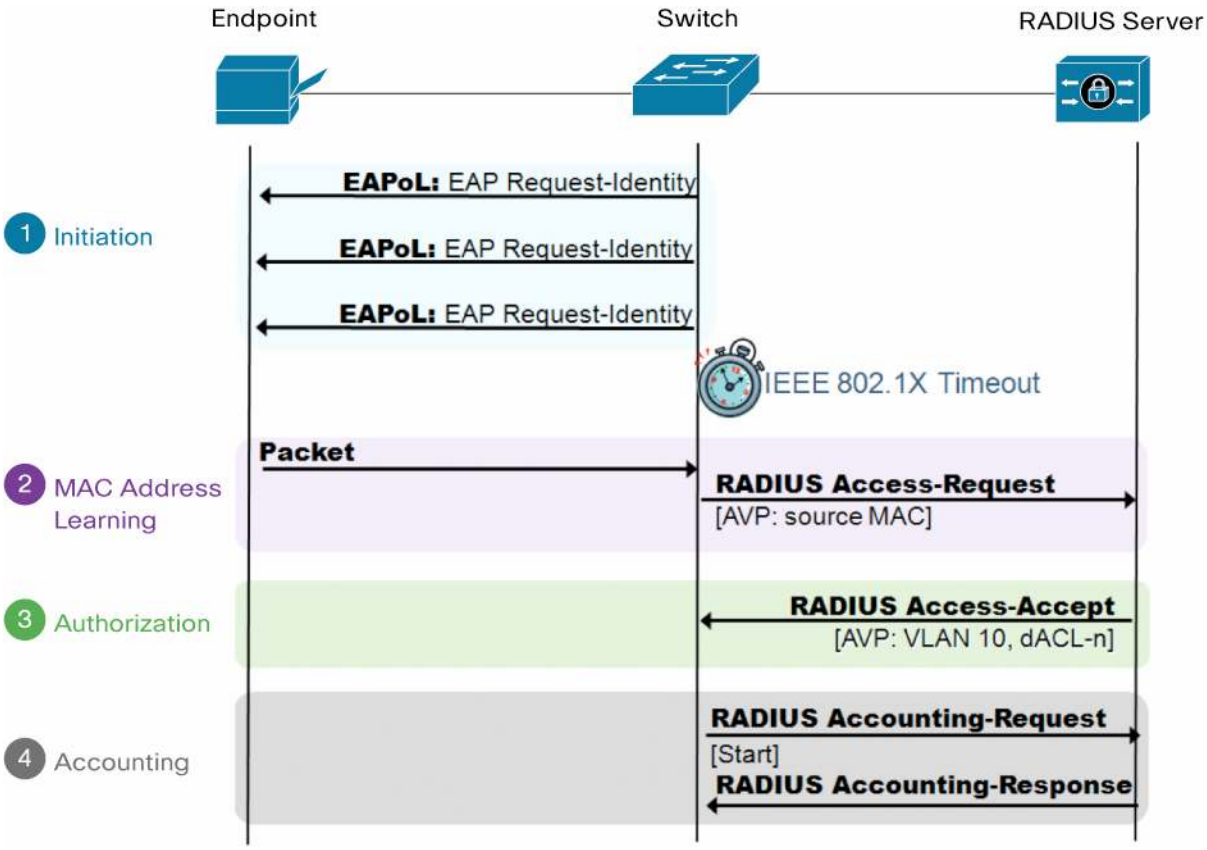

2.2.2 Session Initiation

From the switch's perspective, the authentication session begins when the switch detects link up on a port. The switch will initiate authentication by sending an Extensible Authentication Protocol (EAP) Request-Identity message to the endpoint. If the switch does not receive a response, the switch will retransmit the request at periodic intervals. If no response is received after the maximum number of retries, the switch will let IEEE 802.1X time out and proceed to MAB.

2.2.3 MAC Address Learning

During the MAC address learning stage, the switch begins MAB by opening the port to accept a single packet from which it will learn the source MAC address of the endpoint. Packets sent before the port has fallen back to MAB (that is, during the IEEE 802.1X timeout phase) are discarded immediately and cannot be used to learn the MAC address.

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

The switch can use almost any Layer 2 and 3 packets to learn MAC addresses, with the exception of bridging frames such as Cisco Discovery Protocol, Link Layer Discovery Protocol (LLDP), Spanning Tree Protocol, and Dynamic Trunking Protocol (DTP).[1]

After the switch learns the source MAC address, it discards the packet. Then the switch crafts a RADIUS Access-Request packet. A sample MAB RADIUS Access-Request packet is shown in the sniffer trace in Figure 3.

Figure 3. Sample RADIUS Access-Request Packet for MAB

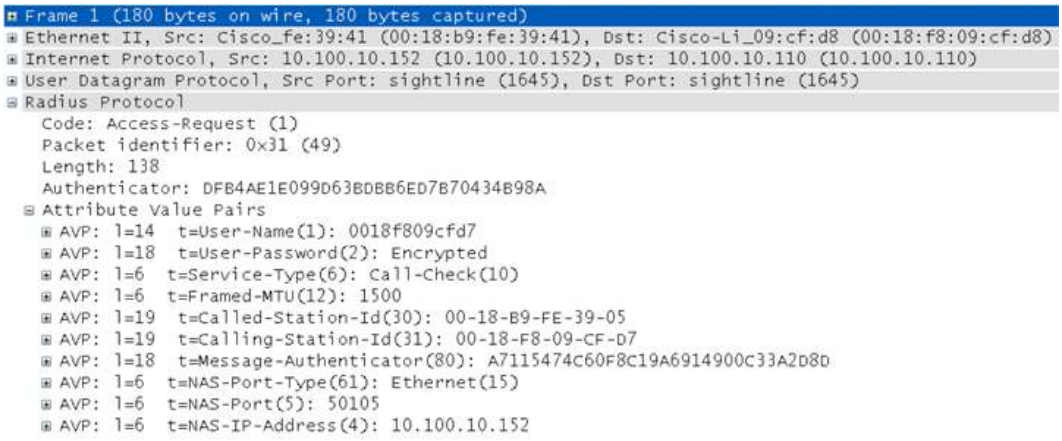

By default, the Access-Request message is a Password Authentication Protocol (PAP) authentication request, The request includes the source MAC address in three attributes: Attribute 1 (Username), Attribute 2 (Password), and Attribute 31 (Calling-Station-Id). Although the MAC address is the same in each attribute, the format of the address differs. This feature is important because different RADIUS servers may use different attributes to validate the MAC address. Some RADIUS servers may look at only Attribute 31 (Calling-Station-Id), while others will actually verify the username and password in Attributes 1 and 2.

Table 1 summarizes the MAC address format for each attribute.

Table 1. MAC Address Formats in RADIUS Attributes

RADIUS Attribute	Format	Example
1 (Username)	12 hexadecimal digits, all lowercase, and no punctuation	0018f809cfd7
2 (Password)	Same as the username but encrypted	\xf2\xb8\x9c\x9c\x13\xdd#,\xcaT\xa1\xcay=&\xee
31(Calling-Station-Id)	6 groups of 2 hexadecimal digits, all uppercase, and separated by hyphens	00-18-F8-09-CF-D7

Because MAB uses the MAC address as a username and password, you should make sure that the RADIUS server can differentiate MAB requests from other types of requests for network access. This precaution will prevent other clients from attempting to use a MAC address as a valid credential. Cisco switches uniquely identify MAB requests by setting Attribute 6 (Service-Type) to 10 (Call-Check) in a MAB Access-Request message. Therefore, you can use Attribute 6 to filter MAB requests at the RADIUS server.

[1] Consult platform documentation for any platform-specific exceptions.

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Optionally, Cisco switches can be configured to perform MAB as EAP-MD5 authentication, in which case the Service-Type attribute will be set to 1 (Framed). However, because the MAC address is sent in the clear in Attribute 31 (Calling-Station-Id), MAB EAP does not offer any additional security by encrypting the MAC address in the password. Also be aware that because the service type for MAB EAP is the same as an IEEE 802.1X request, the RADIUS server will not be able to easily differentiate MAB EAP requests from IEEE 802.1X requests.

2.2.4 Session Authorization

If the MAC address is valid, the RADIUS server will return a RADIUS Access-Accept message. This message indicates to the switch that the endpoint should be allowed access to the port. Optionally, the RADIUS server may include dynamic network access policy instructions (for example, a dynamic VLAN or access control list [ACL]) in the Access-Accept message. In the absence of dynamic policy instructions, the switch will simply open the port. No further authentication methods will be tried if MAB succeeds.

If the MAC address is not valid or is not allowed to access the network for policy reasons, the RADIUS server will return a RADIUS Access-Reject message. This message indicates to the switch that the endpoint should not be allowed access to the port based on the MAC address. Depending on how the switch is configured, several different outcomes are possible. If alternative authentication or authorization methods are configured, the switch may attempt IEEE 802.1X authentication or web authentication or deploy the guest VLAN. The interaction of MAB with these features is described in Section 2.4.

If no fallback authentication or authorization methods are configured, the switch will stop the authentication process and the port will remain unauthorized. You can configure the switch to restart authentication after a failed MAB attempt by configuring **authentication timer restart** on the interface. Enabling this timer means that unknown MAC addresses will periodically fail authentication until the endpoint disconnects from the switch or the address gets added to a MAC database. To prevent the unnecessary control-plane traffic associated with restarting failed MAB sessions, Cisco generally recommends leaving **authentication timer restart** disabled.

2.2.5 Session Accounting

If the switch can successfully apply the authorization policy, the switch can send a RADIUS Accounting-Request message to the RADIUS server with details about the authorized session.

2.2.6 Session Termination

Session termination is an important part of the authentication process. To help ensure the integrity of the authenticated session, sessions must be cleared when the authenticated endpoint disconnects from the network. Sessions that are not terminated immediately can lead to security violations and security holes. Ideally, session termination occurs as soon as the endpoint physically unplugs, but this is not always possible if the endpoint is connected indirectly (for example, through an IP phone or hub).

Multiple termination mechanisms may be needed to address all use cases. Table 2 summarizes the mechanisms and their applications.

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Table 2. Termination Mechanisms and Use Cases

Use Case	Typical Termination Mechanisms
All endpoints directly connected • Single endpoint per port • No IP phones	Link down
Endpoints connected through IP phone • At most two endpoints per port (one phone and one data)	Cisco Discovery Protocol enhancement for second-port disconnect (Cisco phones) Inactivity timer (phones other than Cisco phones)
Endpoints connected through hub • Physical hub • Bridged virtual hubs	Inactivity timer

The following sections discuss in more detail the ways that a MAB session can be terminated.

2.2.6.1 Link Down

The most direct way to terminate a MAB session is to unplug the endpoint. When the link state of the port goes down, the switch completely clears the session. If the original endpoint (or a new endpoint) plugs in, the switch will restart authentication from the beginning.

2.2.6.2 Cisco Discovery Protocol Enhancement for Second-Port Disconnect

For IP telephony deployments with Cisco IP Phones, the best way to help ensure that all MAB sessions are properly terminated is to use Cisco Discovery Protocol. Cisco IP Phones can send a Cisco Discovery Protocol message to the switch indicating that the link state for the data endpoint's port is down, allowing the switch to immediately clear the data endpoint's authenticated session.

> **Best Practice Recommendation: Use Cisco Discovery Protocol Enhancement for Second-Port Disconnect for IP Telephony Deployments**
>
> This feature works for all authentication methods, takes effect as soon as the endpoint disconnects, and requires no configuration. If you are using Cisco IP Phones and Cisco Catalyst® Family switches with the appropriate code release, this method offers the simplest and most effective solution. No other method works as well to terminate authenticated sessions behind Cisco IP Phones.

2.2.6.3 Inactivity Timer

When the inactivity timer is enabled, the switch monitors the activity from authenticated endpoints. When the inactivity timer expires, the switch removes the authenticated session.

The inactivity timer for MAB can be statically configured on the switch port, or it can be dynamically assigned using the RADIUS Idle-Timeout attribute (Attribute 28). Cisco recommends setting the timer using the RADIUS attribute because this approach lets gives you control over which endpoints are subject to this timer and the length of the timer for each class of endpoints. For example, endpoints that are known to be quiet for long periods of time can be assigned a longer inactivity timer value than chatty endpoints.

The inactivity timer is an indirect mechanism that the switch uses to infer that a endpoint has disconnected. An expired inactivity timer cannot guarantee that a endpoint has disconnected. Therefore, a quiet endpoint that does not send traffic for long periods of time (for example, a network printer that services occasional requests but is otherwise silent) may have its session cleared even though it is still connected. That endpoint will then have to send traffic before it can be authenticated again and have access to the network.

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

2.2.6.4 Reauthentication and Absolute Session Timeout

Reauthentication cannot be used to terminate MAB-authenticated endpoints. Absolute session timeout should be used only with caution.

The reauthentication timer for MAB is the same as for IEEE 802.1X. The timer can be statically configured on the switch port, or it can be dynamically assigned by sending the Session-Timeout attribute (Attribute 27) and the RADIUS Termination-Action attribute (Attribute 29) with a value of RADIUS-Request in the Access-Accept message from the RADIUS server. For IEEE 802.1X endpoints, the reauthentication timer is sometimes used as a keepalive mechanism. This feature does not work for MAB. Upon MAB reauthentication, the switch does not relearn the MAC address of the connected endpoint or verify that the endpoint is still active; it simply sends the previously learned MAC address to the RADIUS server. Essentially, a null operation is performed.

The absolute session timer can be used to terminate a MAB session, regardless of whether the authenticated endpoint remains connected. The session timer uses the same RADIUS Session-Timeout attribute (Attribute 27) as the server-based reauthentication timer described earlier with the RADIUS Termination-Action attribute (Attribute 29) set to Default. The switch will terminate the session after the number of seconds specified by the Session-Timeout Attribute and immediately restart authentication. If IEEE 802.1X is configured, the switch will start over with IEEE 802.1X, and network connectivity will be disrupted until IEEE 802.1X times out and MAB succeeds. This process can result in significant network outage for MAB endpoints. As an alternative to absolute session timeout, consider configuring an inactivity timeout as described in Section 2.2.6.3.

2.2.6.5 RADIUS Change of Authorization

RADIUS change of authorization (CoA) allows a RADIUS server to dynamically instruct the switch to alter an existing session. Cisco Catalyst switches support four actions for CoA: reauthenticate, terminate, port shutdown, and port bounce. The reauthenticate and terminate actions terminate the authenticated session in the same way as the reauthentication and session timeout actions discussed in Section 2.2.6.4. The port down and port bounce actions clear the session immediately, because these actions result in link-down events.

2.3 Design Considerations

This section discusses important design considerations that you should evaluate before you deploy MAB.

2.3.1 MAC Address Discovery

Before deploying MAB, you must determine which MAC addresses you want to allow on your network. There are several approaches to collecting the MAC addresses that will be used to populate your MAC address database.

The easiest and most economical method is to find preexisting inventories of MAC addresses. For example, in some companies the purchasing department keeps rigorous records of the MAC address of every device that has ever been approved for purchase. Another good source for MAC addresses is any existing application that uses a MAC address in some way. For example, Cisco Unified Communication Manager keeps a list of the MAC addresses of every registered IP phone on the network. VMPS users can reuse VMPS MAC address lists. After existing inventories of MAC addresses have been identified, they can be exported from the existing repository and then imported into a MAB database as discussed in Section 4.

In the absence of existing MAC address inventories, you may be able to use information from the network to discover the MAC addresses that exist in your network today. One option is to enable MAB in a monitor mode deployment scenario. In monitor mode, MAB is performed on every endpoint, but the endpoint's network access is not affected regardless of whether MAB passes or fails. In this way, you can collect MAC addresses in a

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

nonintrusive way by parsing RADIUS authentication records. See Section 2.4.15.1 for more information about monitor mode. Simple Network Management Protocol (SNMP) MAC address notification traps, syslogs, and network management tools such as CiscoWorks LAN Management Solution (LMS) may also contain MAC address information.

Another option is to use MAC address prefixes (or wildcards) instead of actual MAC addresses. When assigning MAC addresses to devices, vendors set the first three octets to a specific value called the organizationally unique identifier (OUI). OUIs are assigned by the IEEE and uniquely identify the manufacturer of a given device. If an endpoint vendor has an OUI (or set of OUIs) that is exclusively assigned to a particular class of device, then you can create a wildcard rule in your RADIUS server policy that allows any device that presents a MAC address beginning with that OUI to be authenticated and authorized.

With the exception of a preexisting inventory, the approaches described here tell you only what MAC addresses currently exist on your network. No automated method can tell you which endpoints are valid corporate-owned assets. If this is a necessary distinction for your security policy, some sort of manual process (such as an export from an existing asset inventory) will be required.

2.3.2 MAB Databases and RADIUS Servers

After you have discovered and classified the allowed MAC addresses for you network, you must store them in a database that can be accessed by the RADIUS server during the MAB attempt. Where you choose to store your MAC addresses will depend on many factors, including the capabilities of your RADIUS server. Deployment considerations for internal databases, external Lightweight Directory Access Protocol (LDAP) databases, and Microsoft Active Directory are discussed in this section.

2.3.2.1 Internal Databases

An obvious place to store MAC addresses is on the RADIUS server itself. With some RADIUS servers, you simply enter the MAC addresses in the local user database, setting both the username and password to the MAC address. Other RADIUS servers, such as Cisco Secure Access Control Server (ACS) 5.0, are more MAB aware. Cisco Secure Access Control System 5.0 stores MAC addresses in a special host database that contains only allowed MAC addresses. Instead of treating the MAB request as a PAP authentication, Cisco Secure ACS 5.0 recognizes a MAB request (by Attribute 6 [Service-Type] = 10) and compares the MAC address in the Calling-Station-Id attribute to the MAC addresses stored in the host database.

Before choosing to store MAC addresses on the RADIUS server, you should address several concerns. First, does your RADIUS server support an internal hosts database? For example, Microsoft Internet Authentication Service (IAS) and Network Policy Server (NPS) do not have the concept of an internal host database (they rely on Microsoft Active Directory as the identity store). Second, what is the capacity of your RADIUS server? For example, Cisco Secure ACS 5.0 supports up to 50,000 entries in its internal host database. If you plan to support more than 50,000 devices in your network, an external database will be required. Third, how will MAC addresses be managed? If MAC addresses are stored locally on the RADIUS server, then the people who need to add, modify, and delete MAC addresses will need to have administrative access to the RADIUS server. If that presents a problem to your security policy, an external database will be required.

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

2.3.2.2 LDAP Databases

A common choice for an external MAC database is a Lightweight Directory Access Protocol server. LDAP is a widely used protocol for storing and retrieving information on the network. After you have collected all the MAC addresses on your network, you can import them to the LDAP directory server and configure your RADIUS server to query that server.

Because external databases are dedicated servers, they can scale to greater numbers of MAC addresses than can internal databases. They can also be managed independently of the RADIUS server.

The first consideration you should address is whether your RADIUS server can query an external LDAP database. Although LDAP is a very common protocol, not all RADIUS servers can perform LDAP queries to external databases. For example, Microsoft IAS and NPS servers cannot query external LDAP databases.

Because the LDAP database is external to the RADIUS server, you will also need to give special consideration to availability. Since the LDAP database is essential to MAB, redundant systems should be deployed to help ensure that the RADIUS server can contact the LDAP server. To address the possibility that the LDAP server may become completely unavailable, the RADIUS server should be configured with an appropriate failback policy (for example, fail open or fail closed, based on your security policy).

2.3.2.3 Microsoft Active Directory

Microsoft Active Directory is a widely deployed directory service that many organizations use to store user and domain computer identities. If centralizing all identities in a single store is important to you, Active Directory can be used as a MAC database. In fact, in some cases, you may not have a choice. For Microsoft NPS and IAS, Active Directory is the only choice for MAC address storage. In any event, before deploying Active Directory as your MAC database, you should address several considerations.

Starting with Microsoft Windows Server 2003 Release 2 (R2) and Windows Server 2008, Microsoft Active Directory provides a special object class for MAC addresses called ieee802Device. By using this object class, you can streamline MAC address storage in Active Directory and avoid password complexity requirements. Unfortunately, in earlier versions of Active Directory, the ieee802Device object class is not available.

In the absence of that special object class, you can store MAC addresses as users in Microsoft Active Directory. Unfortunately, this method adds unnecessary attributes and objects to the Users group and will not work in an Active Directory forest in which a password complexity policy is enabled. Remember that for MAB, username = password = MAC address, a situation that is intentionally disallowed by password complexity requirements in Active Directory. Another option that avoids the password complexity requirements is to load your MAC addresses as text (TXT) records in a Domain Name System (DNS) zone that is stored inside Active Directory.

If you are going to store MAC addresses in Microsoft Active Directory, make sure that your RADIUS server can access account information in Active Directory. Microsoft IAS and NPS do this natively. Some RADIUS servers, such as the Cisco Secure ACS, accomplish this by joining the Active Directory domain. Alternatively, you can create a lightweight Active Directory instance that can be referred to using LDAP.

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

2.4 Feature Interaction

2.4.1 IEEE 802.1X

MAB is an important part of most IEEE 802.1X deployments. MAB is one of the features Cisco provides to accommodate non–IEEE 802.1X endpoints. After IEEE 802.1X times out or fails, the port can move to an authorized state if MAB succeeds.

Figure 4 shows the MAB process when IEEE 802.1X times out because the endpoint cannot perform IEEE 802.1X authentication.

Figure 4. MAB as Fallback Mechanism for non–IEEE 802.1X Endpoints

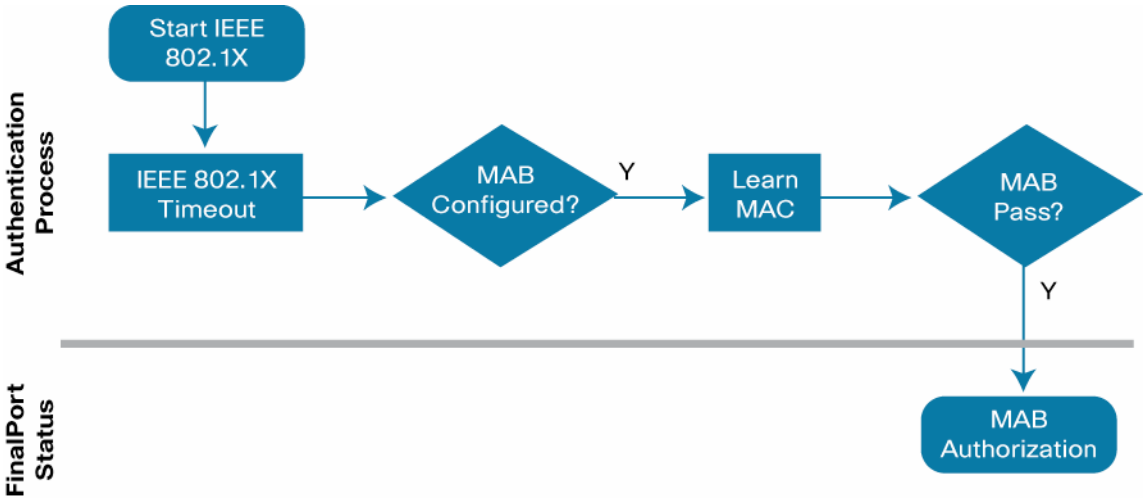

When configured as a fallback mechanisms, MAB is deployed after IEEE 802.1X times out. If the endpoint's Pre-eXecution Environment (PXE) process times out, or if Dynamic Host Configuration Protocol (DHCP) gets deep into the exponential backoff process before the timeout occurs, the endpoint may not be able to communicate even though the port has been opened. To the end user, it will appear as if network access has been denied. There are three potential solutions to this problem:

- Decrease the IEEE 802.1X timeout value. See Section 2.4.1.1.1 for more information about relevant timers.
- Use a low-impact deployment scenario that allows time-critical traffic such as DHCP prior to authentication. See Section 4 for additional reading about deployment scenarios.
- If your network has many non–IEEE 802.1X–capable endpoints that need instantaneous access to the network, you can use the Flexible Authentication feature set that allows you to configure the order and priority of authentication methods. Instead of waiting for IEEE 802.1X to time out before performing MAB, you can configure the switch to perform MAB first and fallback to IEEE 802.1X only if MAB fails. See Section 4 for additional reading about Flexible Authentication.

MAB can also be used as a failover mechanism if the endpoint supports IEEE 802.1X but presents an invalid credential. Figure 5 illustrates this use of MAB in an IEEE 802.1X environment.

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Figure 5. MAB as a Failover Mechanism for Failed IEEE Endpoints

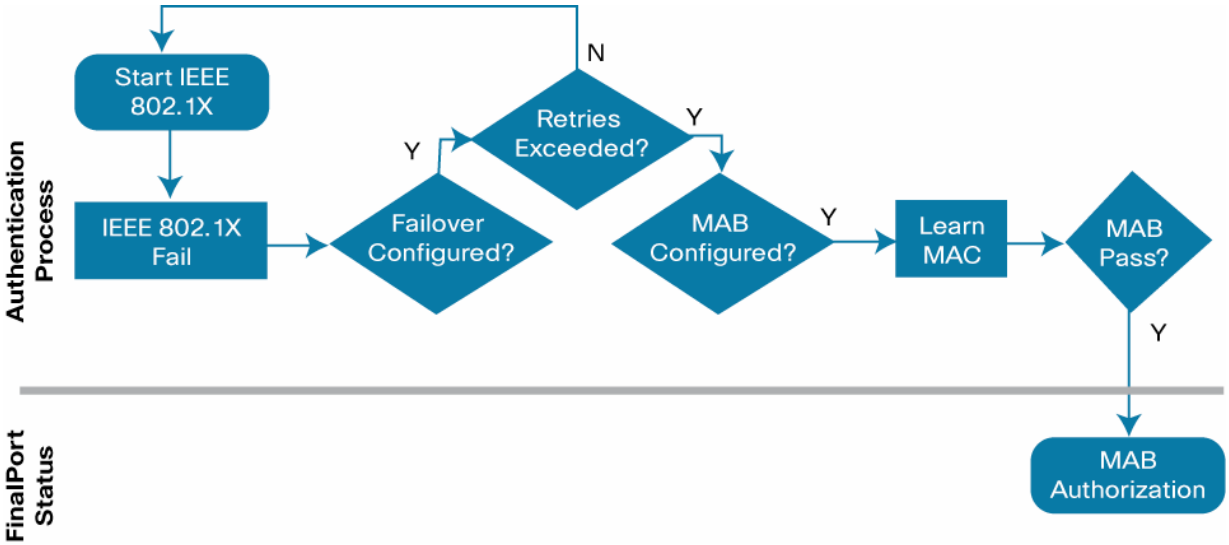

Because MAB begins immediately after an IEEE 802.1X failure, there are no timing issues. However, to trigger MAB, the endpoint must send a packet after the IEEE 802.1X failure. In other words, the IEEE 802.1X supplicant on the endpoint must fail open.

See Section 4 for more information about IEEE 802.1X.

2.4.1.1 Timers and Variables

This section describes the timers on the switch that are relevant to the MAB authentication process in an IEEE 802.1X–enabled environment.

2.4.1.1.1 **dot1x timeout tx-period** and **dot1x max-reauth-req**

This section discusses the timers that control the timeout and retry behavior of a MAB-enabled port in an IEEE 802.1X–enabled environment.

If IEEE 802.1X is enabled in addition to MAB, the switch will send an EAP Request-Identity frame upon link up. The switch will wait for a period of time defined by **dot1x timeout tx-period** and then send another Request-Identity frame. The number of times it will resend the Request-Identity frame is defined by **dot1x max-reauth-req**.

Figure 6 shows the effect of the **tx-period** timer and the **max-reauth-req** variable on the total time to network access.

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Figure 6. Tx-period, max-reauth-req, and Time to Network Access

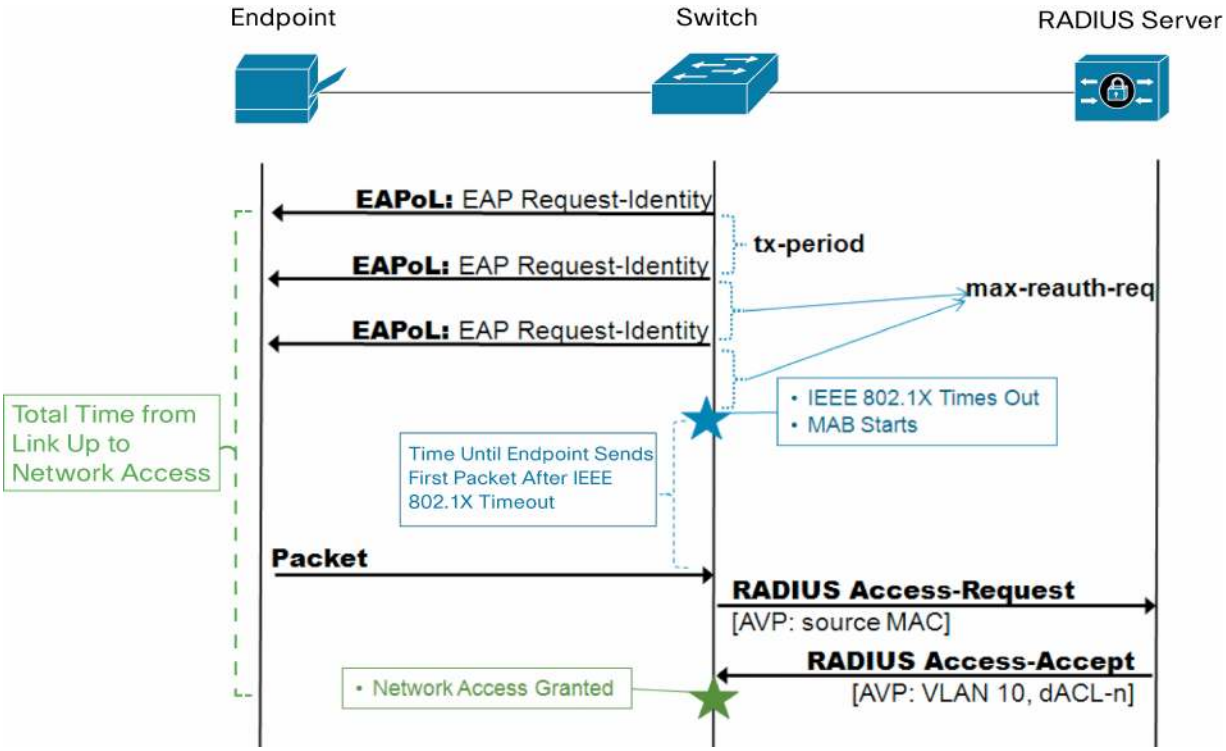

The combination of **tx-period** and **max-reauth-req** is especially important to MAB endpoints in an IEEE 802.1X–enabled environment. MAB endpoints must wait until IEEE 802.1X times out before attempting network access through a fallback mechanism. The total time it takes for IEEE 802.1X to time out is determined by the following formula:

```
Timeout = (max-reauth-req +1) * tx-period
```

Cisco Catalyst switches have default values of **tx-period** = 30 seconds and **max-reauth-req** = 2. Applying the formula, it takes 90 seconds by default for the port to start MAB. By modifying these two settings, you can decrease the total timeout to a minimum value of 2 seconds.

Unlike with IEEE 802.1X, there is no timeout associated with the MAC address learning phase. The switch will wait indefinitely for the endpoint to send a packet. So while the time needed for IEEE 802.1X to time out and fall back to MAB is determined precisely by the configured IEEE 802.1X timeout value and retry count, the time needed for the MAC address to be learned is indeterminate, since the time depends on the endpoint's sending of some kind of traffic. Therefore, the total amount of time from link up to network access is also indeterminate. For chatty devices that send a lot of traffic, MAB will be triggered shortly after IEEE 802.1X times out. But for quiet devices, or for devices that have gone quiet because, for example, the DHCP client timed out before IEEE 802.1X did, MAB may not occur for some time.

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Because of the impact on MAB endpoints, most customers change the default values of **tx-period** and **max-reauth-req** to allow more rapid access to the network. When modifying these values, consider the following:

- A timer that is too short may cause IEEE 802.1X–capable endpoints to be subject to a fallback authentication or authorization technique. Although IEEE 802.1X–capable endpoints can restart IEEE 802.1X after a fallback has occurred, you may still be generating unnecessary control-plane traffic. In addition, if the endpoint has been authorized by a fallback method, then that endpoint may temporarily be adjacent to guest devices that have been similarly authorized. If your goal is to help ensure that your IEEE 802.1X–capable assets are always and exclusively on a trusted network, then you should make sure that the timer is long to allow IEEE 802.1X–capable endpoints time to authenticate.
- A timer that is too long can subject MAB endpoints to unnecessarily long delays in getting network access.

Best Practice Recommendation: Test tx-period and max-reauth-req in Your Network

Since the optimal value for the timeout will depend on the specifics of your network, Cisco recommends that you use your deployment planning phase to test whatever value you select. Pay particular attention to DHCP clients, PXE clients, and the specifics of your managed desktop infrastructure.

2.4.2 Web Authentication

MAB is compatible with Web Authentication (WebAuth). Cisco Catalyst switches can be configured to attempt WebAuth after MAB fails. Access to the network is granted based on the success or failure of WebAuth. The sequence of events is shown in Figure 7.

Figure 7. MAB and Web Authentication After IEEE 802.1X Timeout

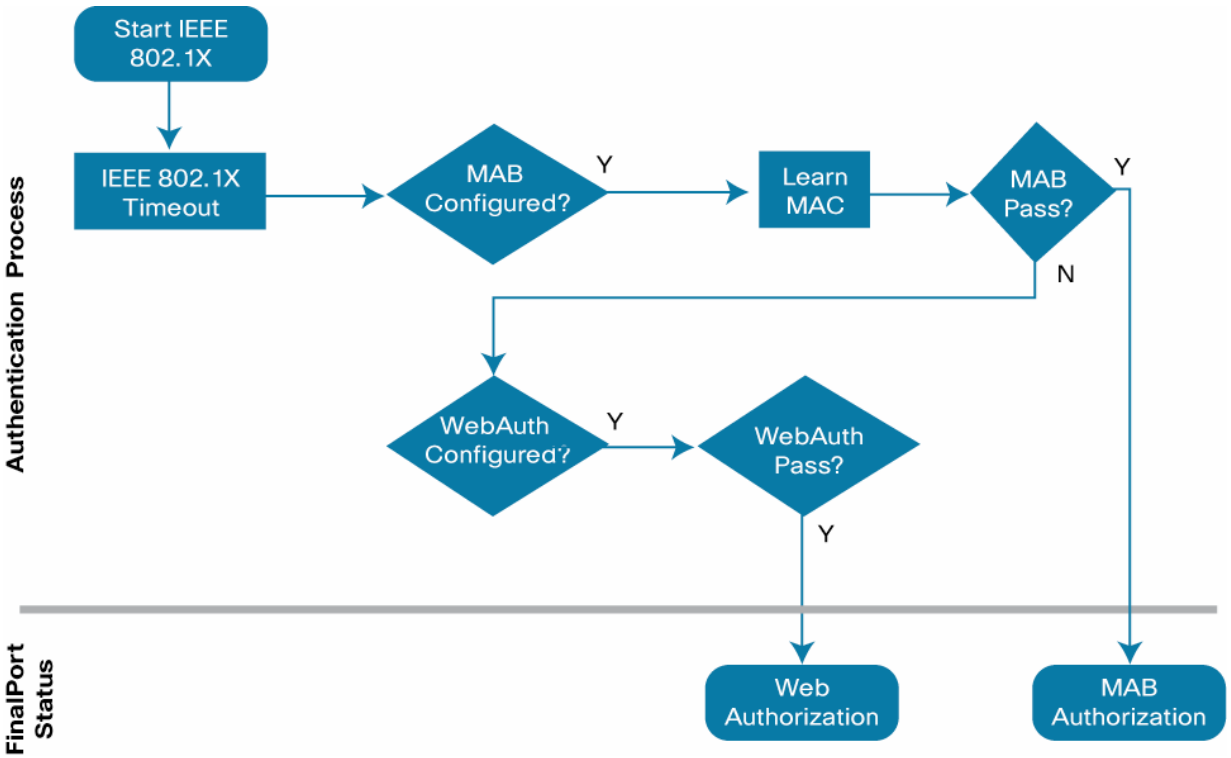

See Section 4 for more information about Web Authentication.

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

2.4.3 Guest VLAN

MAB is compatible with the Guest VLAN feature (Figure 8). If IEEE 802.1X times out (or is not configured) and MAB fails, the port can be moved to the Guest VLAN, a configurable VLAN for which restricted access can be enforced. Using the Guest VLAN, you can tailor network access for endpoints without valid credentials. For example, the Guest VLAN can be configured to permit access only to the Internet.

Figure 8. MAB and Guest VLAN After IEEE 802.1X Timeout

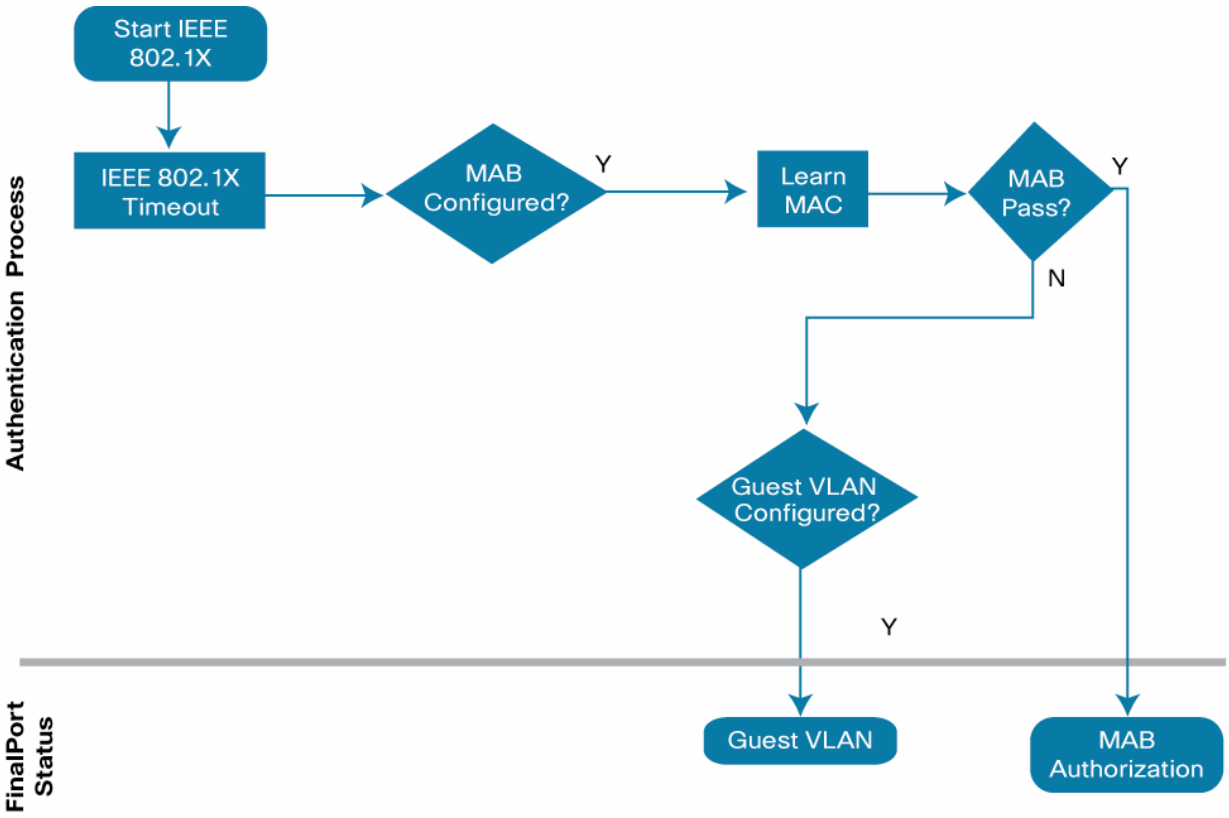

2.4.4 Authentication Failure VLAN

After an IEEE 802.1X authentication failure, the switch can be configured to either deploy the Authentication Failure (AuthFail) VLAN or proceed to the next authentication method (MAB or WebAuth). Figure 9 illustrates this process.

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Figure 9. AuthFail VLAN or MAB after IEEE 802.1X Failure

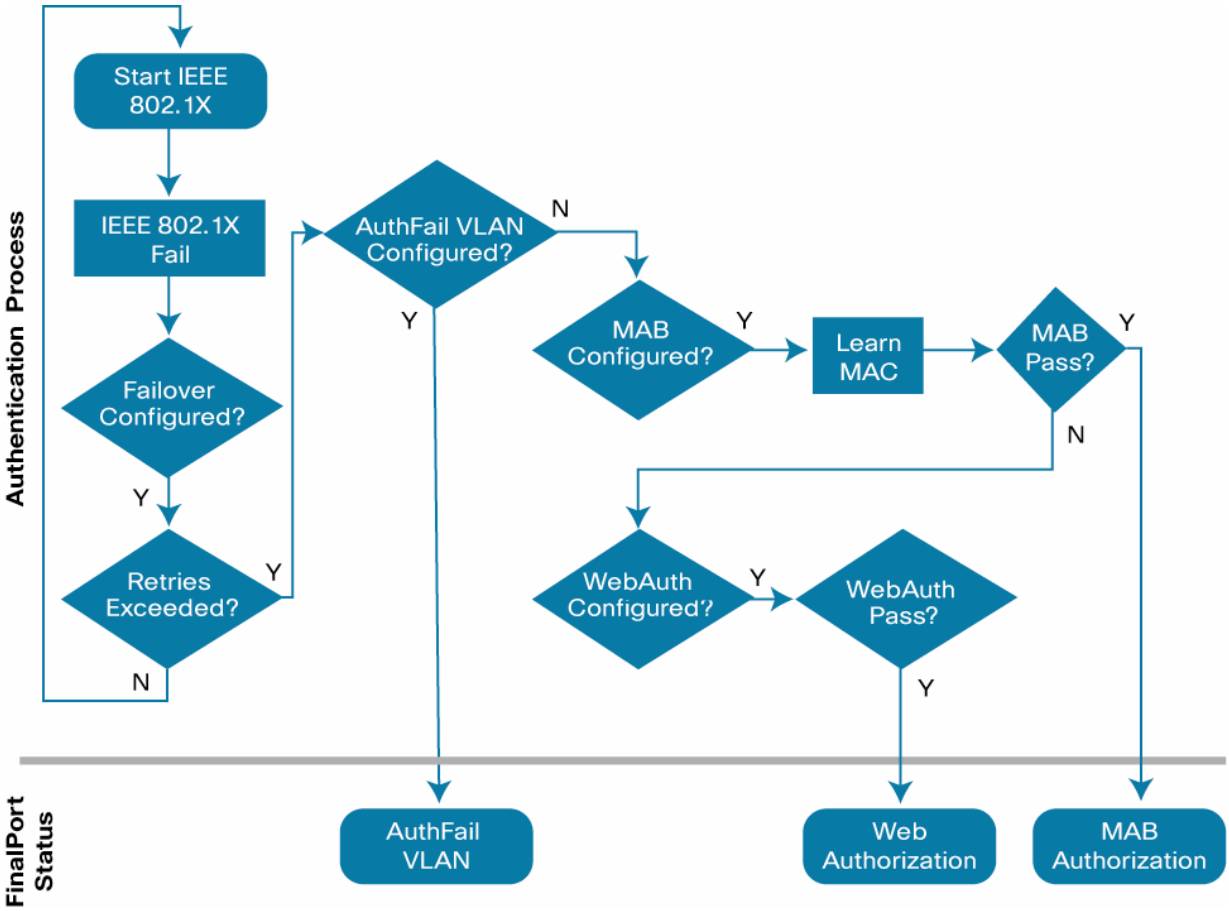

In this sense, AuthFail VLAN and MAB are mutually exclusive when IEEE 802.1X fails. However, you can configure the AuthFail VLAN for IEEE 802.1X failures (the client has a supplicant but presents an invalid credential as shown in Figure 9) and still retain MAB for IEEE 802.1X timeouts (the client has no supplicant as shown in Figures 7 and 8).

2.4.5 Dynamic Guest and Authentication Failure VLAN

Instead of using the locally configured Guest VLAN or AuthFail VLAN, another option is dynamic Guest and AuthFail VLANs, which rely on the RADIUS server to assign a VLAN when an unknown MAC address attempts to access the port after IEEE 802.1X times out or fails. In this scenario, the RADIUS server is configured to send an Access-Accept message with a dynamic VLAN assignment for unknown MAC addresses. The dynamically assigned VLAN would be one for which restricted access can be enforced. From the switch's perspective, MAB will pass even though the MAC address is unknown. The advantage of this approach over the local Guest VLAN and AuthFail VLAN is that the RADIUS server is aware of and in control of unknown endpoints. Centralized visibility and control make this approach preferable if your RADIUS server supports it.

2.4.6 Inaccessible RADIUS Server

When the RADIUS server is unavailable, MAB will fail and, by default, all endpoints will be denied access. In a highly available enterprise campus environment, it is reasonable to expect that a switch will always be able to communicate with the RADIUS server, so the default behavior may be acceptable. However, there may be some

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

use cases (for example, a branch office with occasional WAN outages) in which the switch cannot reach the RADIUS server, but endpoints should be allowed access to the network.

If the switch already knows that the RADIUS server has failed (either through periodic probes or as the result of a previous authentication attempt), a port can be deployed in a configurable VLAN (sometimes called the critical VLAN) as soon as the link comes up. Because the switch has multiple mechanisms for learning that the RADIUS server has failed, this outcome is the most likely. If the switch determines that the RADIUS server has failed during a MAB authentication attempt (for example, if this is the first endpoint to connect to the switch after connectivity to the RADIUS server has been lost), then the port will be moved to the critical VLAN after the authentication times out. Previously authenticated endpoints will not be affected in any way; if a reauthentication timer expires when the RADIUS server is down, the reauthentication will be deferred until the switch determines that the RADIUS server has returned.

When the RADIUS server returns, the switch can be configured to reinitialize any endpoints in the critical VLAN. This behavior poses a potential problem for a MAB endpoint. When the MAB endpoint originally plugged in and the RADIUS server was unavailable, the endpoint received an IP address in the critical VLAN. Because the MAB endpoint is agentless, it has no knowledge of when the RADIUS server has returned or when it has been reinitialized. If the device is assigned a different VLAN as a result of the reinitialization, it will continue to use the old IP address—an IP address that is now invalid on the new VLAN.

There are several ways to work around the reinitialization problem. You can disable reinitialization, in which case, critical authorized endpoints will stay in the critical VLAN until they unplug and plug back in. You also can set the critical VLAN to the data VLAN (essentially a fail-open operation) so that the MAB endpoints maintain a valid IP address across reinitialization. If neither of those options is feasible, consider setting the DHCP lease time in the critical VLAN scope to a short time (for example, 5 minutes) so that a MAB endpoint will have an invalid address for a relatively short amount of time.

2.4.7 Dynamic ACL Assignment

MAB is compatible with ACLs that are dynamically assigned by the RADIUS server as the result of successful authentication.

2.4.8 Dynamic VLAN Assignment

MAB is compatible with VLANs that are dynamically assigned by the RADIUS server as the result of successful authentication. Be aware that MAB endpoints cannot recognize when a VLAN changes. Therefore, if a MAB endpoint initially has an IP address in VLAN A and is later assigned to VLAN B without an intervening link-down or link-up event (for example, as the result of reauthentication), the unsuspecting MAB endpoint will continue to use the IP address from the old VLAN and hence be unable to get access on the new VLAN.

2.4.9 Wake on LAN

Wake on LAN (WoL) is an industry-standard power management feature that allows you to remotely wake up a hibernating endpoint by sending a “magic packet” over the network. Most WoL endpoints flap the link when going into hibernation or standby mode, thus clearing any existing MAB-authenticated session. By default, traffic through the unauthorized port will be blocked in both directions, and the magic packet will never get to the sleeping endpoint.

To support WoL in a MAB environment, you can configure a Cisco Catalyst switch to modify the control direction of the port, allowing traffic to the endpoint while still controlling traffic from the endpoint. This approach allows the hibernating endpoint to receive the WoL packet while still preventing the unauthorized endpoint from sending any

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

traffic to the network. After it is awakened, the endpoint can authenticate and gain full access to the network. Control direction works the same with MAB as it does with IEEE 802.1X.

Additionally, when a port is configured for open-access mode, magic packets are not blocked, even on unauthorized ports, so no special configuration for WoL endpoints is necessary.

2.4.10 Open Access

By default, the port drops all traffic prior to successful MAB (or IEEE 802.1X) authentication. This approach is sometimes referred to as closed mode. Cisco switches can also be configured for open access, which allows all traffic while still enabling MAB.

> **Best Practice Recommendation: Do Not Assign Dynamic VLANs to MAB Endpoints in Open Access**
>
> Endpoints should not be dynamically assigned a VLAN as the result of MAB in open-access mode. If an endpoint initially gets an IP address in the statically configured data VLAN in open-access mode and then is assigned to a new VLAN as the result of MAB, the endpoint will continue to use the IP address from the data VLAN and hence be unable to get access on the dynamically assigned VLAN.

Open access has many applications, including increasing network visibility as part of a monitor mode deployment scenario. It can be combined with other features to provide incremental access control as part of a low-impact mode deployment scenario. For more information about these deployment scenarios, see Section 4.

2.4.11 Multiple Endpoints per Port

By default, a MAB-enabled port allows only a single endpoint per port. Any additional MAC addresses seen on the port will cause a security violation.

Frequently, the limitation of a single endpoint per port will not meet all the requirements of real-world networks. Cisco Catalyst switches allow you to address multiple use cases by modifying the default behavior. The host mode on a port determines the number and type of endpoints allowed on a port. The various host modes and their applications are discussed in this section.

> **Best Practice Recommendation: Use the Most Restrictive Host Mode That Addresses Your Use Cases**
>
> Limiting the number of MAC addresses allowed on the port helps ensure the validity of the authenticated session and discourages casual port piggybacking.

Single-Host Mode

In single-host mode, only a single MAC or IP address can be authenticated (by any method) on a port. If a different MAC address is detected on the port after a endpoint has authenticated with MAB, then a security violation will be triggered on the port. This is the default behavior.

Multidomain Authentication Host Mode

Multidomain authentication was specifically designed to address the requirements of IP telephony. When multidomain authentication is configured, two endpoints are allowed on the port: one in the voice VLAN and one in the data VLAN. Either, both, or none of the endpoints can be authenticated with MAB. Additional MAC addresses will trigger a security violation.

Multi-Authentication Host Mode

If the port is configured for multi-authentication (multi-auth) host mode, then multiple endpoints can be authenticated in the data VLAN. Each new MAC address that appears on the port will be separately authenticated. Any, all, or none of the endpoints can be authenticated with MAB. Multi-auth host mode can be used for bridged virtual environments or to support hubs.

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Multihost Mode

Unlike multi-auth host mode, which authenticates every MAC address, multihost mode authenticates the first MAC address and then allows an unlimited number of other MAC addresses. Because of the security implications of multihost mode, multi-auth host mode typically is a better choice than multihost mode.

2.4.12 IP Telephony

Cisco Catalyst switches are fully compatible with IP telephony and MAB. For a full description of features and a detailed configuration guide, see http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6638/config_guide_c17-605524.html.

2.4.13 Cisco Catalyst Integrated Security Features

This section describes Cisco Catalyst integrated security features.

2.4.13.1 Port Security

In general, Cisco does not recommend enabling port security when MAB is also enabled. Since MAB enforces a single MAC address per port (or per VLAN when multidomain authentication is configured for IP telephony), port security is largely redundant and may in some cases interfere with the expected operation of MAB.

2.4.13.2 DHCP Snooping

DHCP snooping is fully compatible with MAB and should be enabled as a best practice.

2.4.13.3 Dynamic Address Resolution Protocol Inspection

Dynamic Address Resolution Protocol (ARP) Inspection (DAI) is fully compatible with MAB and should be enabled as a best practice.

2.4.13.4 IP Source Guard

IP source guard is compatible with MAB and should be enabled as a best practice.

2.4.14 RADIUS Accounting

RADIUS accounting is fully compatible with MAB and should be enabled as a best practice. RADIUS accounting provides detailed information about the authenticated session and enables you to correlate MAC address, IP address, switch, port, and use statistics.

2.4.15 Deployment Scenarios

When deploying MAB as part of a larger access-control solution, Cisco recommends a phased deployment model that gradually deploys identity-based access control to the network. The three scenarios for phased deployment are monitor mode, low-impact mode, and high-security mode. Each scenario identifies combinations of authentication and authorization techniques that work well together to address a particular set of use cases. The interaction of MAB with each scenario is described in the following sections.

For more information about scenario-based deployments, see http://www.cisco.com/go/ibns.

2.4.15.1 Monitor Mode

MAB is fully supported and recommended in monitor mode.

The primary goal of monitor mode is to enable authentication without imposing any form of access control. This approach allows network administrators to see who is on the network and prepare for access control in a later phase without affecting endpoints in any way.

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

By enabling MAB in monitor mode, you get the highest level of visibility into devices that do not support IEEE 802.1X. In addition, by parsing authentication and accounting records for MAB in monitor mode, you can rapidly compile a list of existing MAC addresses on your network and use this list as a starting point for developing your MAC address database as described in Section 2.3.1.

2.4.15.2 Low-Impact Mode

MAB is fully supported in low-impact mode.

Low-impact mode builds on the ideas of monitor mode, gradually introducing access control in a completely configurable way. Instead of denying all access before authentication (as a traditional IEEE 802.1X or MAB deployment would require), low-impact mode allows you to use ACLs to selectively allow traffic before authentication. This approach is particularly useful for devices that rely on MAB to get access to the network. Waiting until IEEE 802.1X times out and falls back to MAB can have a negative effect on the boot process of these devices. Low-impact mode enables you to permit time-sensitive traffic prior to MAB, enabling these devices to function effectively in an IEEE 802.1X–enabled environment.

2.4.15.3 High-Security Mode

MAB is fully supported in high-security mode.

High-security mode is a more traditional deployment model for port-based access control, which denies all access prior to authentication. It also facilitates VLAN assignment for the data and voice domains. The primary design consideration for MAB endpoints in high-security mode is the lack of immediate network access if IEEE 802.1X is also configured. MAB endpoints that are not capable of IEEE 802.1X authentication will have to wait for IEEE 802.1X to time out and fall back to MAB before they get access to the network. To help ensure that MAB endpoints get network access in a timely way, you will need to adjust the default timeout value as described in Section 2.4.1.1. Alternatively, you can use Flexible Authentication to perform MAB before IEEE 802.1X authentication as described in Section 2.4.1

2.5 Deployment Summary for MAB

Table 3 summarizes the major design decisions that need to be addressed prior to deploying MAB.

Table 3. MAB Deployment Reference

Design Consideration	Relevant Section
Evaluate your MAB design as part of a larger deployment scenario.	2.4.15
Collect MAC addresses of allowed endpoints.	2.3.1
Store MAC addresses in a database that can be queried by your RADIUS server.	2.3.2
Modify timers, use low-impact mode, or perform MAB before IEEE 802.1X authentication to enable MAB endpoints to get time-critical network access when MAB is used as a fallback to IEEE 802.1X.	2.4.1
Use an unknown MAC address policy for the dynamic Guest or AuthFail VLAN.	2.4.5
Do not enable reauthentication.	2.2.6.4
Disable reinitialization on RADIUS server recovery if the static data VLAN is not the same as the critical VLAN.	2.4.6
Leave the restart timer disabled.	2.2.4
Decide how many endpoints per port you must support and configure the most restrictive host mode.	2.4.11
Eliminate the potential for VLAN changes for MAB endpoints.	2.4.8 and 2.4.10
Identify the session termination method for indirectly connected endpoints: • Cisco Discovery Protocol enhancement for second-port disconnect (Cisco IP Phones) • Inactivity timer with IP device tracking (physical or virtual hub and third-party phones)	2.2.6

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Design Consideration	Relevant Section
Enable RADIUS accounting.	2.4.14
Disable port security.	2.4.13.1

3. Conclusion

MAB offers visibility and identity-based access control at the network edge for endpoints that do not support IEEE 802.1X. With the appropriate design and well-chosen components, you can meet the needs of your security policy while reducing the impact on your infrastructure and end users.

4. For More Information

IEEE 802.1X Quick Reference Guide:
http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6638/whitepaper_c27-574041.pdf

IEEE 802.1X Design Guide:
http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6638/guide_c07-627531.html

IEEE 802.1X Deployment Scenarios Design Guide:
http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6638/whitepaper_C11-530469.html

IEEE 802.1X Deployment Scenarios Configuration Guide:
http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6638/Whitepaper_c11-532065.html

Basic Web Authentication Design and Configuration Guide:
http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6638/app_note_c27-577494.html

Advanced Web Authentication Design and Configuration Guide:
http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6638/app_note_c27-577490.html

Deploying IP Telephony in IEEE 802.1X Networks Design and Configuration Guide:
http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6638/config_guide_c17-605524.html

Flexible Authentication, Order, and Priority App Note:
http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6638/application_note_c27-573287_ps6638_Products_White_Paper.html

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

5. Sample Configuration for Standalone MAB

This section includes a sample configuration for standalone MAB. For configuration examples of MAB as a fallback to IEEE 802.1X, see the IEEE 802.1X Deployment Scenarios Configuration Guide in Section 4.

MAB requires both global and interface configuration commands. Note that even though IEEE 802.1X is not enabled on the port, the global authentication, authorization, and accounting (AAA) configuration still uses the dot1x keyword.

```
aaa new-model
aaa authentication dot1x default group radius
aaa authorization network default group radius
aaa accounting dot1x default start-stop group radius
!
interface FastEthernet2/48
switchport access vlan 40
switchport mode access
authentication port-control auto
mab
spanning-tree portfast
spanning-tree bpduguard enable
!
radius-server host 10.200.1.52 key cisco123
```


Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

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

Printed in USA C17-663759-00 05/11

© 2011 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.