INTEROP LABS

9 in a Series

Switch Features

Switch Functionality for 802.1X-based NAC

As an IEEE standard, 802.1X is a critical building block of the major NAC architectures. Before deploying a NAC architecture, the first step is to roll out 802.1X. This whitepaper will cover the switch and access point features that support an 802.1X environment.

802.1X has been a standard feature of switches and access points for many years. In its purest sense, the switch requires a device to authenticate successfully before the port is opened. The original 802.1X implementations in switches had numerous shortcomings that led to painful deployments. Many of these issues have been addressed by new features over the past several years.

These features can be divided into two primary groups: features that make the network easier to use (supplicant-bypass mechanisms) and features that make the network more manageable (access enforcement mechanisms).

The features discussed in this paper are applicable to switches, WLAN controllers, and to a limited degree, freestanding access points. For the purposes of this paper, the term "switch" will be used to generically refer to this entire group of devices.

Supplicant-Bypass Mechanisms

Access to an 802.1X network requires the end-user's device to have an 802.1X supplicant installed and correctly configured. Traditionally, a switch would deny access if the client was incorrectly configured. This created support issues in 802.1X deployments.

Today, a new set of switch features, coined supplicant-bypass mechanisms, have emerged and relaxed this restriction. These new features are designed to allow access to the network, typically provisioned in a limited fashion, in the following situations:

- The device does not have an 802.1X supplicant installed.
- The device has a supplicant but it is disabled.
- The device has an incorrectly configured 802.1X supplicant, such as using PEAP vs. TTLS.
- The end-user does not have proper credentials (i.e. wrong password) and fails authentication.
- The end-user does not have an account (i.e. guest user).

The following list of supplicant-bypass mechanisms is limited to features present in switches. It is common for these features to be used in combination with other network devices to provide more flexible authentication mechanisms. The switch-based supplicant bypass mechanisms include:

Default/Guest Policy

This feature was introduced several years ago using a default or guest VLAN and has become very popular. The original implementations of this feature placed the port into a default (or guest) VLAN if a supplicant was not detected within a defined period of time.

The original intent of this feature was to provide guests with an Internet-only VLAN, either by directly routing the VLAN to the Internet or by routing the VLAN through a gateway appliance. This feature can also be used as the basis for additional functionality. By combining the default policy with a captive portal (either on or off the switch), a support portal can be created that allows for functionality like guest registration, software & how-to distribution, and configuration assistance.

Newer implementations have extended this concept to include functionality beyond VLAN assignment, such as ACL or packet filter assignment.

A couple notes are worth mentioning regarding switch behavior:

- Some switches leave the port closed until 802.1X times out while others open the port immediately on the default VLAN and/or ACL.
- Some switches allow the configuration of a different default VLAN and/or ACL for a failed authentication situation versus a situation where a supplicant is not present.

MAC Authentication Bypass

When MAC authentication bypass is configured, the port waits a defined period of time for an 802.1X authentication to occur, and if it does not occur, the switch generates a RADIUS request on behalf of the end-user device, using the device's MAC address as the user name. This functionality can be used in conjunction with enforcement options such as VLAN and ACL assignment.

The most obvious benefit of MAC authentication bypass is that it allows MAC addresses to be managed centrally. When queried, the RADIUS server can recognize a MAC address as a phone and place it on the phone network regardless of what switch is used. A less obvious benefit is that it allows the RADIUS server the opportunity to apply a default policy to every new connection, potentially replacing the default/guest policy feature.

Similar to the default/guest policy feature, existing implementations differ in behavior:

- Some switches leave the port closed during the timeout period while others leave it open on the default VLAN.
- Switches differ in how the MAC address is formatted in the user name field of the RADIUS request. Some switches support multiple formatting options. Switches use different passwords in the RADIUS request. Some switches use the MAC address while others use a switch-wide configurable string.
- Switches differ in how they specify the timeout period. Some switches use a single value, such as 15 seconds from link up. Others use a combination of values, such as waiting 5 seconds per attempt and trying 3 times.

Web-Login Bypass

The final supplicant-bypass mechanism is a web-based login hosted on the switch. The switch performs an HTTP hijack (Universal Access Method) on web requests until a successful supplicant-based or web-based authentication occurs. If the user submits the web login page, the switch packages the information into a RADIUS request and forwards it to the RADIUS server for authentication.

This feature allows the switch to provide dynamic enforcement through VLANs and ACLs, but reduces the password security built into 802.1X.

Access Enforcement Mechanisms

The majority of a switch's role in NAC involves enforcing the policy determined by the RADIUS server during or after authentication. The least common denominator for access enforcement is dynamic virtual LAN (VLAN) assignment. A newer approach is to assign the device a packet filter or an access control list. This functionality is supported in a growing list of switches. For the purposes of this paper, the term "ACL" will be used generically to refer to both access control lists and packet filters.

The standard approach for communicating enforcement information from the RADIUS server to the switch is in the RADIUS response. IETF RFC 3580 specifies the standard attributes currently used for VLAN assignment (tunnel-type, tunnel-medium-type, private-tunnel-group-id) and for ACLs (filter-id). Today, most vendors support the standard attributes for VLAN assignment. Most vendors also use the standard filter-id attribute for specifying the ACL, but the format of the attribute's value differs by vendor.

Multi-Authentication

According to the 802.1X standard, a single port should have one and only one device attached to it. However, this restriction has caused deployment issues and switch manufacturers have responded with the multi-authentication feature. The multi-authentication feature allows multiple devices to attach to a single port, typically through a hub or switch. Authentication and enforcement occur on a per-device basis, utilizing the MAC address as a device's identifier. The result is that three PCs may be connected to a hub, which uplinks to the switch. Each PC will be forced to authenticate individually and each PC will receive its own VLAN assignment.

The number of devices supported on a single port varies from a handful to thousands based on the switch.

Due to the shared medium nature of wireless, this functionality is inherent in wireless devices.

VLAN Assignment by RADIUS

VLAN assignment allows an end-user device to be dynamically placed on a VLAN based on the response from the RADIUS server. If the end-user device unplugs from the switch, the port is immediately returned to its default state.

For the current generation of switches, specifying a VLAN in the RADIUS response requires three attributes. This is because VLAN assignment utilizes the tunnel concept standardized by IETF RFC 3580 as a generic mechanism for specifying network access capabilities. To specify the VLAN, the RADIUS server returns three attributes (numeric values in parenthesis):

1. tunnel-type (64) which is "VLAN" (13)
2. tunnel-medium-type (65) which is "IEEE-802" (6)
3. private-tunnel-group-id (81) which is the VLAN id or name ("corporate")

Together, these three attributes say that we want to place the user in a tunnel, which will utilize a VLAN with the id "corporate" within an 802 network. In this example, the switch is responsible for translating "corporate" into a numeric VLAN id.

Along with these three attributes, there is the concept of the tunnel-tag. The tunnel-tag is simply a byte that occurs at the beginning of the value for each of the three attributes. This tunnel-tag byte provides a mechanism to group the three attributes together in the event that multiple tunnels are specified in the RADIUS response.

Before the standardization of the VLAN attributes occurred, vendors implemented VLAN assignment using vendor-specific attributes (VSAs). Today, most vendors have retired these VSAs in favor of the standard VLAN attributes.

The VLAN attributes defined in RFC 3580 do have a limitation. The structure of the VLAN attributes in RFC 3580 limits the RADIUS response to specifying a single, untagged VLAN id (PVID). In many cases, this is sufficient. To allow support for RADIUS configuration of 802.1Q trunk ports, IETF 4675 was recently drafted. RFC 4675 defines standard RADIUS attributes for granting a device access to an array of VLANs and permits control over tagging. This new capability will allow full control of IEEE 802.1Q functionality and should begin appearing in switches later this year.

The following is a list of differences in current implementations and general notes.

- The VLAN needs to be defined on the switch before it can be dynamically assigned. Defining the VLAN on the uplink trunk is occasionally sufficient. This may also be accomplished through the use of the GARP VLAN Registration Protocol (GVRP), which allows VLAN information to be automatically propagated for 802.1Q trunk ports. GVRP allows edge switches to "learn" VLAN information from the core switch, reducing the amount of configuration necessary on the edge switch.
- Some switches require that the private-tunnel-group-id provide the VLAN name while others use the VLAN id. The former requires predefinition of the VLAN name on the switch but is convenient for allowing the RADIUS server to say "put the user on the QUARANTINE VLAN." As such, the switch can handle the translation between "QUARANTINE" and the local VLAN id.
- Implementations vary on the treatment of the tunnel-tag portion of the VLAN attributes. During initial configuration, if a switch fails to assign the port to the specified VLAN, investigate the tunnel-tag behavior of the switch.

Filter/ACL Assignment by RADIUS

The dynamic application of a packet filter or an access control list to an end-user device is widely supported but less standardized than VLAN assignment. While packet filters and access control lists differ in structure, they will be referred to generically as ACLs.

ACLs have two major benefits. First, they provide finer-grain control than VLANs. Second, they can reduce the number of VLANs required. The number of rules supported by a switch varies widely. ACLs can be used by themselves or in combination with VLANs.

ACL support comes in two flavors. Most switches require that the ACL be predefined on the switch, such that the RADIUS server simply returns the name of the ACL. The benefit of this model is that it simplifies the RADIUS configuration and allows the ACLs to be customized on a per-switch basis. IETF RFC 3580 reserves the "filter-id" attribute for use in specifying filters. The syntax of the filter-id attribute varies by vendor.

Other switches support the specification of the ACL in the RADIUS response, allowing ACLs to be created by the RADIUS server. The benefit of this model is that it allows centralized configuration of the ACLs on the RADIUS server. Current implementations of this functionality utilize VSAs, but IETF RFC 4849 has standardized the "NAS-Filter-Rule" attribute for this purpose.

The following is a list of differences in current implementations and general notes.

- The format of the "filter-id" attribute's value is not standardized and varies from vendor to vendor. In a heterogeneous environment, you may find that two vendors use the "filter-id" attribute differently. If this occurs, there are workarounds. First, you can ask the vendors if they have a VSA available that provides the same functionality as the "filter-id" attribute. Second, you can specify separate policies on the RADIUS server for the vendors, so that each receives the "filter-id" attribute in the appropriate format. This can be accomplished by matching on attributes in the RADIUS request packet, such as the switch's friendly name.
- Blocking the server portion of the DHCP and DNS protocols is a good place to begin with ACLs.
- Some vendors have added the ability to group together VLAN, ACL and QoS information, such that an attribute (either private-tunnel-group-id or filter-id) in the RADIUS request is translated by the switch into a VLAN, an ACL, and/or a QoS priority. This allows the switch to implement the policy locally as needed. For example, the filter-id "QUARANTINE" may translate to VLAN 10 with no ACL on one switch while it is VLAN 20 plus an ACL on a second switch.
- Within RFC 4849, the format for the NAS-Filter-Rule attribute is as follows: "Permit in 6 10.0.0.18 to any 80, 443" where 10.0.0.18 is permitted to send TCP (6) traffic to any ip address on ports 80 and 443.

QoS/Rate Limiting Assignment

Switches are now offering the ability to assign a device a Quality of Service (QoS) value or a rate limiting value based on the RADIUS response. QoS is useful, for example, for prioritizing VOIP traffic. Rate limiting is useful, for example, to limit a user's total bandwidth or to mitigate an ICMP denial of service (DoS) attack.

QoS and rate limiting information is currently specified in the RADIUS response using vendor-specific attributes (VSAs). IETF RFC 4675 has standardized the "User-Priority-Table" attribute for use in specifying QoS information in alignment with IEEE 802.1D.

Utilization in Interop Labs

The network built in the lab utilized most of the features listed above. By doing so, we were able to use a unified configuration where every port on every switch had the same configuration. With the exception of uplinks, every port was configured with the following:

- 802.1X as the primary authentication mechanism.
- Dynamic VLAN &/or ACL assignment as the primary enforcement mechanism.
- Multi-Authentication to allow multiple devices to use a single port.
- MAC Authentication Bypass to allow the RADIUS server to specify the default VLAN/ACL for every device that doesn't authenticate by 802.1X.
- Default/Guest VLAN to apply a VLAN in the event that the RADIUS server is offline.

The end result is a network with the following characteristics:

1. Security-related configuration is centralized into the RADIUS server.
2. The switch configurations are very minimal and static. Other than the uplinks, every port has an identical configuration. There are no ports specified "phone port" or "printer port".
3. Any device can plug into any port or attach to any access point and be applied a consistent policy (as specified by the RADIUS server).

Summary

Switch and wireless manufacturers have rapidly introduced new features related to access control over the past several years. While some capabilities require new hardware, much of the functionality is available in your network today, or available as a firmware upgrade. As you plan a NAC deployment, carefully consider how current and future switch features help create a network design that ensures your network is both secure and supportable.