

Introduction

Web-authentication, (also known as Captive Portal), is a simple way to provide secure guest-user access to a network. It is used in a wide range of environments including Wi Fi hot spots, hotels, universities, and business centres.

In basic terms, if the switch detects an unauthorised user web browsing, then irrespective of the IP configuration on their PC, they are re-directed to a Web-authentication login page. At this point, the user is required to enter a username and password before they can begin to web browse.

The main benefits of this solution come from not requiring additional customer knowledge, software or special configuration. Users are able to quickly and easily gain access to the network regardless of the type of device or operating system used.

This How To Note:

- Describes what Web-authentication is
- Explains how Web-authentication is used
- Examines the operation of Web-authentication
- Presents some examples of how to manage traffic of unauthenticated supplicants

List of terms:

ARP intercept

When ARP intercept is enabled, the switch will process ARP packets coming in from supplicants, and respond to those ARPs as though the switch possessed the IP address being requested in the ARP. The response will be sent even if the switch does not possess that IP address.

Guest VLAN

In a secure network, the default behaviour is to deny any access to supplicants that cannot be authenticated. However, it is often convenient to allow limited access to unauthenticated users. A popular solution is to define a limited-access VLAN, called the Guest VLAN, and assign unauthenticated users into that VLAN.

Auth-fail VLAN

In some networks, it is useful to be able to differentiate between users who have not even attempted to authenticate themselves, and those who have tried and failed. This is achieved by defining an auth-fail VLAN in addition to the Guest VLAN. When a user has had sufficient unsuccessful authentication attempts, they are assigned to the Auth-fail VLAN.

Related How To Notes

You also may find the following AlliedWare Plus How To Notes useful:

- http://www.alliedtelesis.com/media/datasheets/howto/howto_aw+_local_radius_with_certificates.pdf
- http://www.alliedtelesis.com/media/datasheets/howto/config_swi_max_secure_against_attack.pdf

Which products and software version does it apply to?

This How To Note applies to the following Allied Telesis managed Layer 3 switches:

- x600 series switches
- x900 series switches
- SwitchBlade x908 switches

It requires AlliedWare Plus software version 5.3. 4 or later.

| | | |
|--------------------------|---|----|
| Table of Contents | Introduction | 1 |
| | Related How To Notes | 2 |
| | Which products and software version does it apply to? | 2 |
| | What is Web-authentication? | 3 |
| | Web-authentication Basics | 3 |
| | Configuring Web-authentication | 5 |
| | Choosing the Web-authentication-server address | 5 |
| | Starting a Web-authentication Session | 7 |
| | Understanding the Web-authentication Features | 9 |
| | Support for protocols underlying Web-authentication | 9 |
| | Secure Authentication | 14 |
| | Copying a certificate onto the switch | 14 |
| | Ping-poll Monitoring of Supplicant Presence | 15 |
| | Checking the IP addresses of the supplicants | 17 |
| | Ping-poll and promiscuous mode | 17 |
| | Managing Traffic of Unauthenticated Supplicants | 18 |
| | No Guest VLAN or Auth-fail VLAN | 18 |
| | Guest VLAN but no Auth-fail VLAN | 19 |
| | Auth-fail VLAN, but no Guest VLAN | 20 |
| | Auth-fail VLAN, and Guest VLAN | 20 |
| | Monitoring the operation of Web-authentication | 21 |

What is Web-authentication?

Web-authentication is a convenient alternative to 802.1x authentication, it's commonly used to authenticate users in educational institutions, where regular users' workstations are not managed by the network administrator. Web-authentication enables the switch to detect an unauthenticated workstation web browsing into the network, then redirect the user's web browser to its own authentication web page.

Web-authentication works like this:



- The user enters their username and password into the web page, which the switch then sends to a RADIUS server for checking.
- If the RADIUS server accepts the user's credentials, the switch then allows their traffic into the network.

The Web-authenticating switch interacts with a RADIUS server in the same way as an 802.1x authenticator. So the two methods can easily be used together in the same network, using the same RADIUS server.


Web-authentication Basics

Conceptually, the operation of Web-authentication is quite simple:

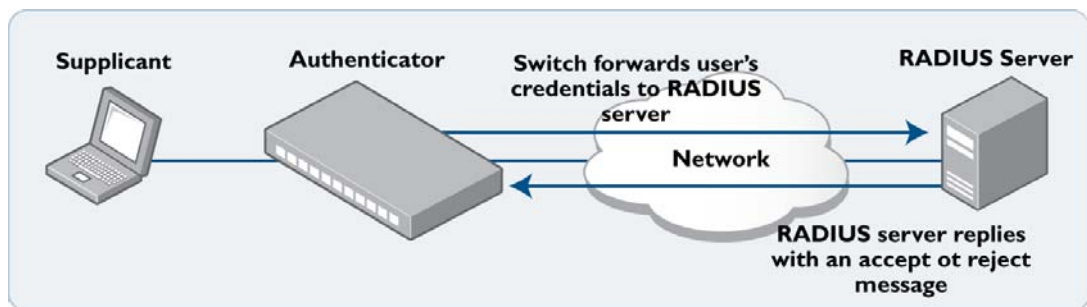
1. The authenticating switch receives HTTP or HTTPS traffic from an unauthenticated supplicant. It intercepts the supplicant's web session, and redirects it to its own internal web server.



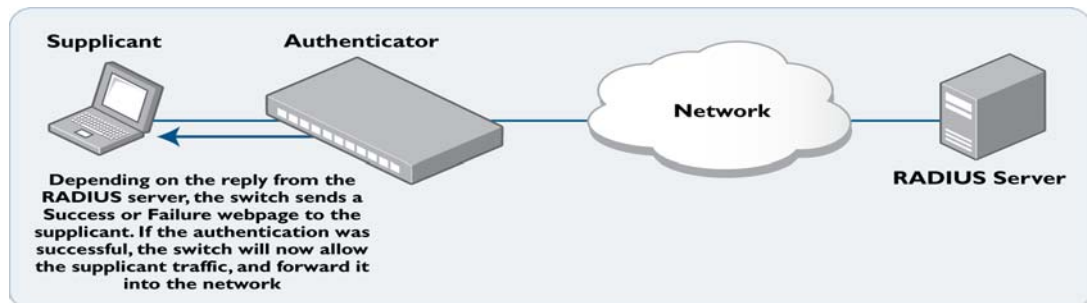
2. The web server serves up an authentication page into which the user may enter their username and password.



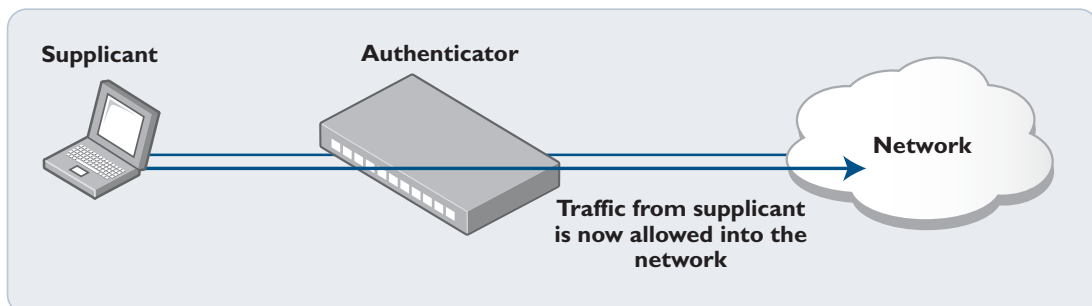
3. The username and password are sent to a RADIUS server, which informs the authenticating switch whether or not the supplicant is authenticated.



4. The user is then informed of the RADIUS server's verdict.



5. If the supplicant has been successfully authenticated, the authenticating switch will give the supplicant workstation access to the network.



Configuring Web-authentication

Web-authentication can be configured on a switch in four simple steps:

1. Configure a RADIUS server.

```
radius-server host <server-ip-address> key <shared secret>
```

2. Instruct Web-authentication to use the configured RADIUS server.

```
aaa authentication auth-web default group radius
```

3. Define the IP address that the Web-authentication service will be accessed on.

```
auth-web-server ipaddress <ip-address>
```

4. Configure ports for Web-authentication.

```
interface port1.0.1-1.0.20 auth-web enable
```

Choosing the Web-authentication server address

When you come to configure Web-authentication, you will need to answer the question:

- What IP address should I specify as the Web-authentication server address?
 - Is it OK to use just any IP address that is configured on one of the switch's VLANs, or is the choice more constrained than that?

The answer is that you must **use the IP address** that is configured on the VLAN that the **supplicant's packets** will **arrive** on.

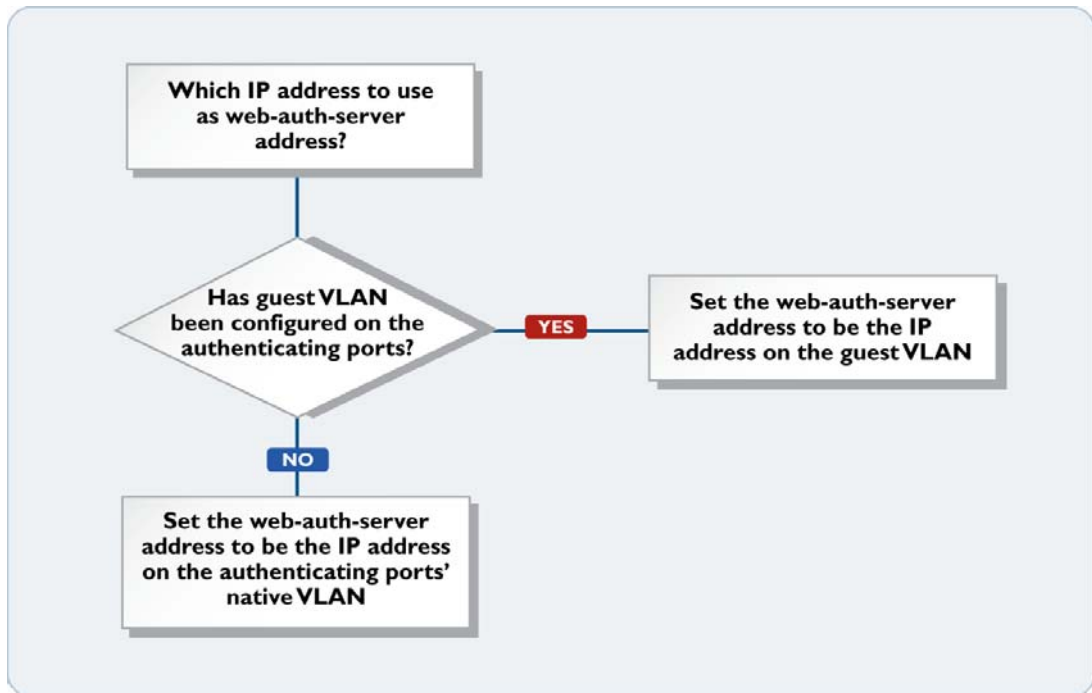
At first glance, that seems a simple answer. If the supplicant-connected ports are access ports in VLAN10, then you would expect that you configure the IP address on VLAN10 as the Web-authentication server address. In fact, that is the correct choice, UNLESS a guest VLAN has been configured on the supplicant-connected ports.

The logic that the switch uses in deciding which VLAN to associate non-authenticated supplicants' packets with is:

- If guest VLAN has been configured on the port where the packet arrives, then associate the packet with the guest VLAN.
- Otherwise associate the packet with the port's native VLAN.

To reiterate, if you configure the supplicant-connected ports with guest VLAN, then use the IP address on the guest VLAN as the IP address of the Web-authentication server. Otherwise use the IP address on the supplicant-connected ports' native VLAN.

Deciding which IP address to use as the Web-auth-server address:



Configuration Example 1: Using guest VLAN

```

VLAN database
  VLAN 20 name guest
  VLAN 10 name edge
  VLAN 30 name core

radius-server host 192.168.30.129 key verysecret
aaa authentication auth-Web default group RADIUS
auth-Web-server ipaddress 192.168.20.1

int vlan10
  ip address 192.168.10.1/24
int vlan20
  ip address 192.168.20.1/24
int vlan30
  ip address 192.168.30.1/24

int port1.0.1-1.0.20
  switchport access vlan 10
  auth-Web enable
  auth guest-vlan 20

int port1.0.21-1.0.22
  switchport access vlan 30
  
```

Configuration Example 2: When no guest VLAN is in use

```
VLAN database
  VLAN 10 name edge
  VLAN 30 name core

radius-server host 192.168.30.129 key verysecret
aaa authentication auth-web default group radius
auth-web-server ipaddress 192.168.10.1

int vlan10
  ip address 192.168.10.1/24
int vlan30
  ip address 192.168.30.1/24

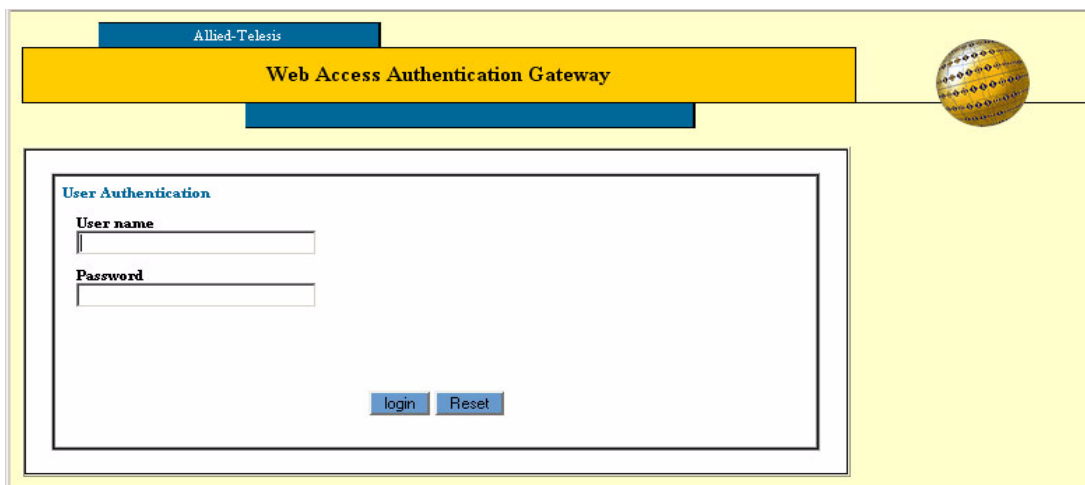
int port1.0.1-1.0.20
  switchport access vlan 10
  auth-Web enable

int port1.0.21-1.0.22
  switchport access vlan 30
```

Starting a Web-authentication Session

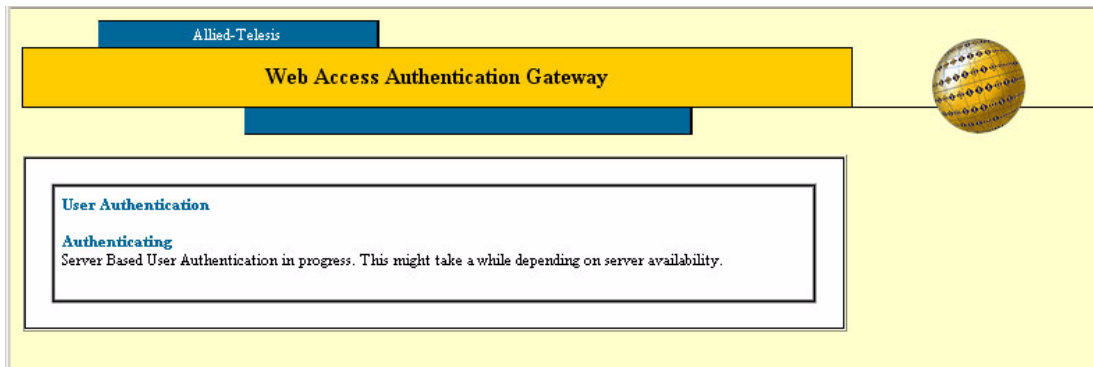
Let us look at what the user actually sees in a Web-authentication session:

1. The user starts their web browser, and browses to a page they wish to view. Shortly thereafter, the address in the browser's address bar automatically changes to the address of the authenticating switch's authentication page.
2. In the switch's authentication page, the user enters their **User name** and **Password**, and clicks **login**.

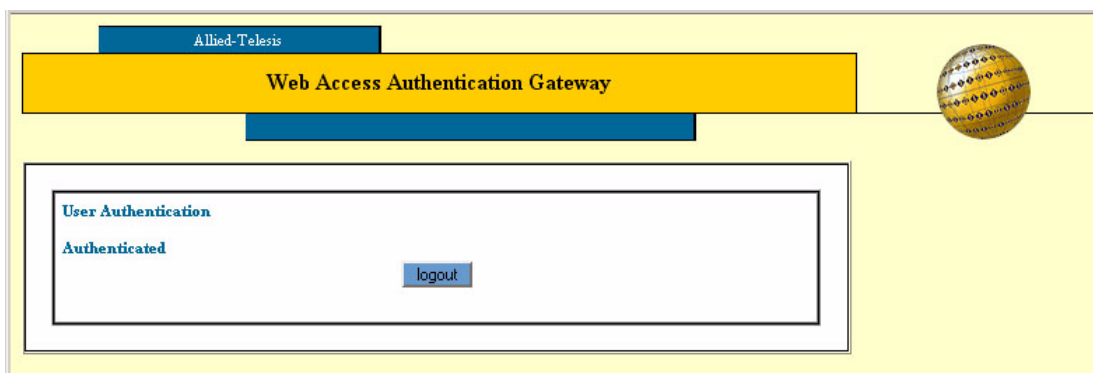


The screenshot shows a web browser window displaying the 'Web Access Authentication Gateway' page. The page has a yellow background and a blue header bar with the text 'Allied-Telesis' and 'Web Access Authentication Gateway'. A globe icon is visible in the top right corner. The main content area is titled 'User Authentication' and contains two input fields: 'User name' and 'Password'. Below the input fields are two buttons: 'login' and 'Reset'.

3. The switch displays a page that informs them that authentication is in progress.

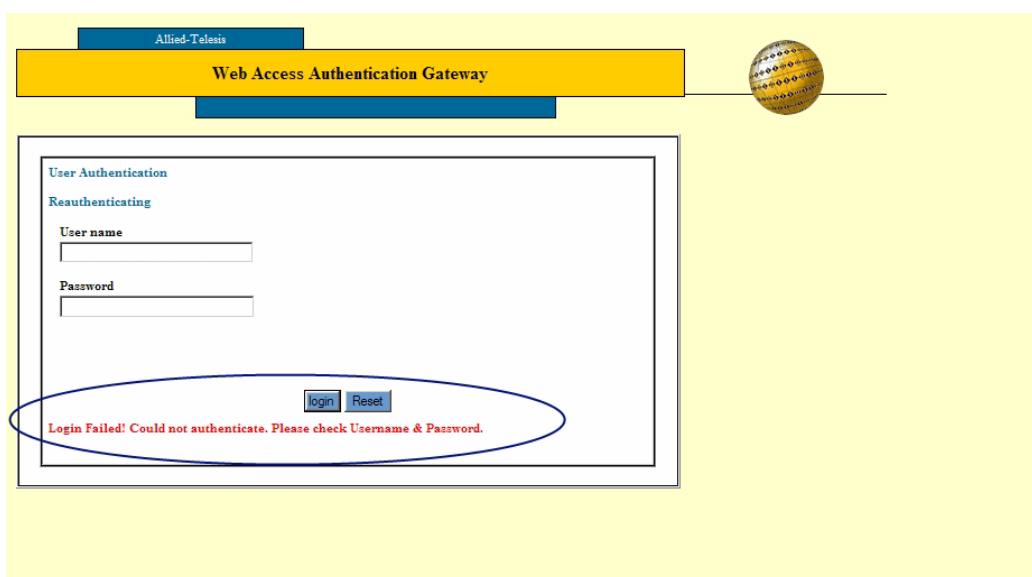


4. Once authentication is complete, the authentication result is displayed.



- If the user enters a username/password combination that is not accepted by the RADIUS server, the switch presents them with an invitation to check the username and password, and try again.

If the user enters incorrect usernames/passwords several times the authentication has failed. The number of times a user can try to login is configurable but it is set to 3 by default.



Understanding the Web-authentication Features

While the authentication process, as it has been described so far, is essentially quite simple, there are actually a number of implementation details that it glosses over.

To use Web-authentication effectively, it is necessary to understand these details – how they work and how to configure them.

We'll take a closer look at:

- Protocol support features
- Secure authentication (SSL)
- Ping-poll monitoring of supplicant presence
- Managing traffic of unauthenticated supplicants

Support for protocols underlying Web-authentication

Web-authentication does not use a dedicated protocol like 802.1x, with a standards-defined set of messages for authentication conversation. When it comes to Web-authentication, the switch is overlaying the authentication process on top of another process that was not designed for authentication.

The web browser communication process that the authentication overlays, is itself reliant on IP addressing, ARP, and DNS. The authentication needs to occur in a seamless manner for all users, irrespective of their IP and DNS setting, and before they have full access to the network.

To make this possible, the switch needs to provide facilities that enable the user's PC to access the authentication web page.

There are a few different features of Web-authentication that work together to achieve this:

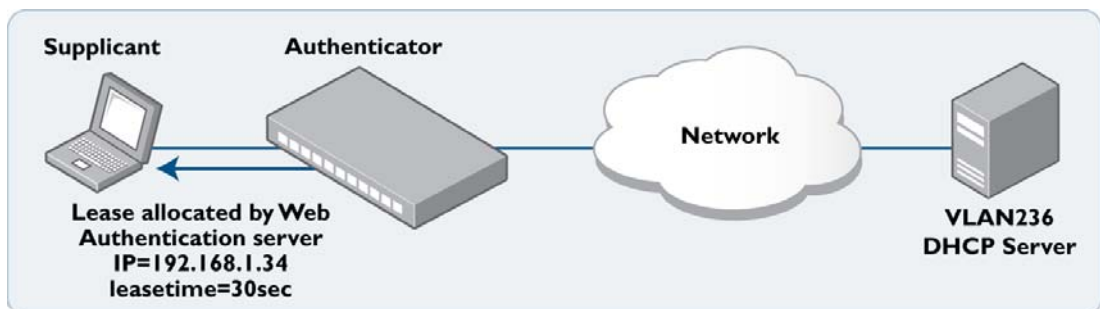

- DHCP server for Web-authentication
- Interception of clients' ARPs
- Proxy DNS response

DHCP server for Web-authentication

To initiate a web browsing session, the supplicant needs an IP address. If the supplicant has been configured to obtain its IP address by DHCP, then the authenticating switch needs to ensure that the supplicant will be served an IP address.

The simplest way to achieve this, that avoids forwarding the supplicant's DHCP requests to any other DHCP server, is to have the Web-authentication process itself act as a DHCP Server. There is a DHCP server built in to Web-authentication.

This DHCP server is dedicated to serving IP addresses to be used by Web-authentication clients.



This DHCP service is configured by the command:


```
auth-web-server dhcp ip address <ip-address/prefix-length>
```

The IP address specified in this command is the IP address of the Web-authentication service. If the Web-authentication service's IP address has not already been configured by the command **auth-web-server ip address <ip-address>**, then this command configures the service's address.



If the Web-authentication service's IP address had already been configured by the command **auth-web-server ip address <ip-address>**, then the IP address in the **auth-web-server dhcp** command must be the same as that already configured. By default, this DHCP server serves leases of 30 second's duration. The lease duration can be changed by the command **auth-web-server dhcp lease <20-60>**. The short lease is deliberate. It facilitates the transition to a new VLAN/subnet after authentication. The supplicant is unaware that the switch transitions it to another VLAN, with another DHCP server, after authentication succeeds.

Similarly, there is no mechanism by which the switch signals to the supplicant to say "I have just assigned you to VLAN 236, you now need to obtain a DHCP lease from the DHCP server on that VLAN". How can we force the supplicant to request a new DHCP lease after the completion of the authentication process? There is no mechanism by which the supplicant's web browser signals down to the DHCP client process to say "I've just completed an authentication session, you need to request a new DHCP lease"

The best way is to ensure that the lease allocated by the dedicated Web-authentication DHCP service is of a very short duration. The lease will expire within a short time from the completion of the authentication process, resulting in the supplicant requesting a new lease.



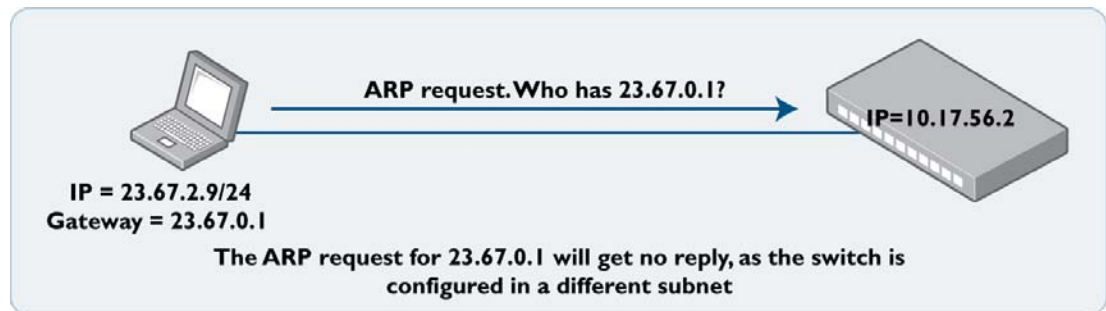
- This new request will now be serviced by the DHCP server on the supplicant's new VLAN.



Interception of clients' ARPs

If the supplicant has been configured with a static IP address, then it is more than likely that the supplicant's IP configuration bears no relation to the Web-auth server address. A computer's IP communications will always be preceded by sending out ARP requests for host addresses in its local subnet, or for its gateway address.

If the IP address and gateway address have been statically configured on the computer, and the subnet used in this static configuration is different to that on the authenticating switch, then the ARP requests will receive no reply, and the PC will not begin IP communication.

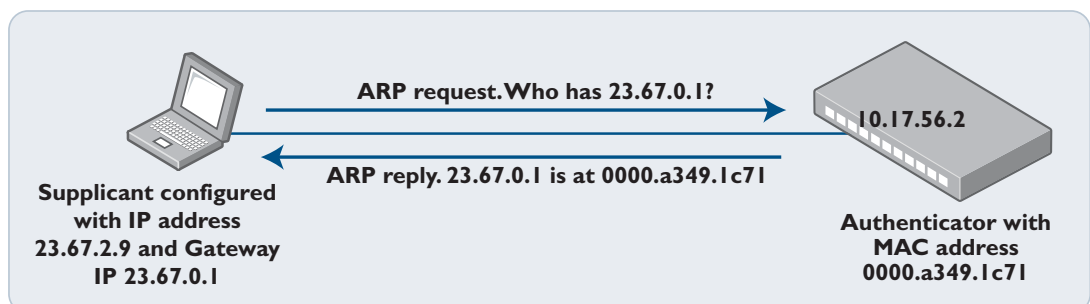


To deal with any arbitrary IP configuration on the supplicants, Web-authentication needs a method for replying to arbitrary ARP requests. This is the **ARP interception** feature.

ARP interception, can operate in three modes:

```
auth-Web-server mode {intercept|none|promiscuous}
```

1. **Intercept** – will respond to ARP requests for any IP address that is in the same subnet as the switch's own IP address. Will provide its own MAC address in the ARP reply, irrespective of what IP address (within its own subnet) was being requested.
2. **None** – will only respond to ARP requests for its own IP address.
3. **Promiscuous** – will respond to **any** ARP request. Will provide its own MAC address in the ARP reply, irrespective of what IP address was being requested. When this mode is configured, the Web-authentication server can interoperate with **any** static IP configuration on a supplicant.




In promiscuous mode, the switch will send its own MAC address in response to an ARP request for ANY address, no matter whether the requested address bears any relation to the switch's own IP address on the interface where the ARP is received.

Proxy DNS response

Typically, an HTTP session from a web browser is preceded by a DNS request for the IP address of the web site the user wishes to browse to. If the DNS request receives no reply, the web browser will never progress on to connecting an HTTP session.

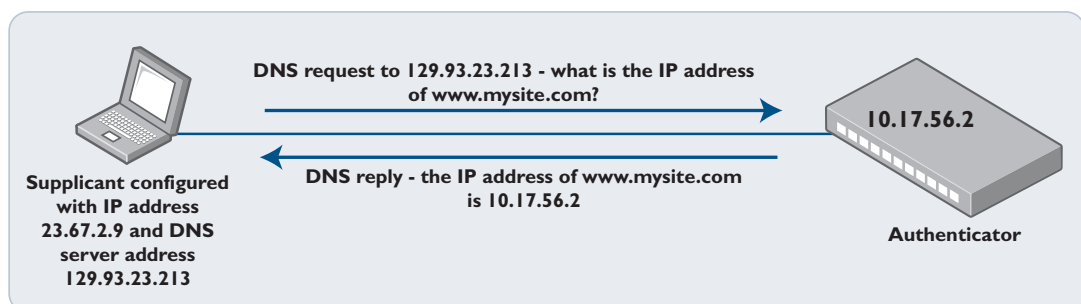
The Web-authentication server needs a mechanism to reply to DNS requests, so that the Web-authentication session can begin.



A web browser must request a DNS Server for the IP address corresponding to a URL. But the switch will not forward the request if the supplicant is not yet authenticated

The three modes listed also control the operation of the proxy DNS replies.

1. **Intercept** – responds to DNS requests whose source IP address is within the same subnet as the IP address on the switch. The IP address provided as the resolution of the DNS lookup is the switch’s own IP address, so that the subsequent HTTP traffic will be directed to the switch.
2. **None** – does not respond to DNS requests.
3. **Promiscuous** – responds to DNS requests from any source IP address. The IP address provided as the resolution of the DNS lookup is the switch’s own IP address, so that the subsequent HTTP traffic will be directed to the switch.



In promiscuous mode, the switch will reply to ANY DNS request from an authenticated supplicant, regardless of whether the destination IP address of the DNS server bears any relation to the switch’s own IP address. The DNS reply from the switch will always specify its own IP address as the URL that was being requested.

Secure Authentication

The Web-authentication service can be configured to use a secure HTTPS connection. This ensures that the username and password are sent from the supplicant to the switch in encrypted form, and cannot be snooped by anyone eavesdropping on the session.

Secure Web-authentication requires two steps:

1. Create an SSL certificate for the switch using the command:

```
crypto pki enroll local
```

2. Configure the Web-authentication service to use HTTPS instead of HTTP:

```
auth-web-server ssl
```

Once the Web-authentication service has been put into secure mode, the service will always use HTTPS for Web-authentication sessions, irrespective of whether or not the user had initially directed their browser to an HTTPS session.

- Even if the initial session that the user directed their browser to was <http://mysite.com>, the Web-authentication service will automatically redirect them to <https://<Web-authentication server address>>.

By default, the Web-authentication service uses TCP port 443 for HTTPS sessions, but it can be configured to use a different port, using the command:

```
auth-web-server sslport <1-65535>
```

Copying a certificate onto the switch

As well as using the self-created certificate, it is also possible to create a certificate elsewhere, and copy that certificate onto the switch to be the SSL certificate for the Web-authentication service.

The command to copy the certificate onto the switch is:

```
copy tftp://<tftp server address>/<certificate file name>  
Web-auth-https-file
```

Note: that the file that is copied onto the switch must:

- be in PEM format
- contain both the certificate and the corresponding Private key

Such a file could be created, for example, by using `openssl`, which is available for multiple different operating systems.

The **openssl** commands to create a key pair and a certificate are:

- Create the private key

```
openssl genrsa -out privkey.pem 2048
```

- Create a self-signed certificate for this key

```
openssl req -new -x509 -key privkey.pem -out cacert.pem
```

- This will result in you being prompted for a number of parameters, like organisation name, email address, etc. Enter whatever values you want for these parameters.

Privkey.pem and cacert.pem are text files. Use a text editor to combine the content of these files together into a single file. The order within the file does not matter – the key could be first, or the certificate could be first.

Once the file has been copied onto the switch to be the Web-authentication HTTPS file, the output of the command **show auth-Web-server** will show:

```
awplus#show auth-Web-server
Web-authentication server
Server status: enabled
<SNIP>
Certification: user <----->
<SNIP>
```

If you are not using a certificate that was copied onto the switch, but using one generated by the switch itself, then this is reported as "Certification: Default". If you wish to remove the certificate that you have copied onto the switch, and go back to using the switch's self-generated certificate, use the command:

```
erase web-auth-https-file
```

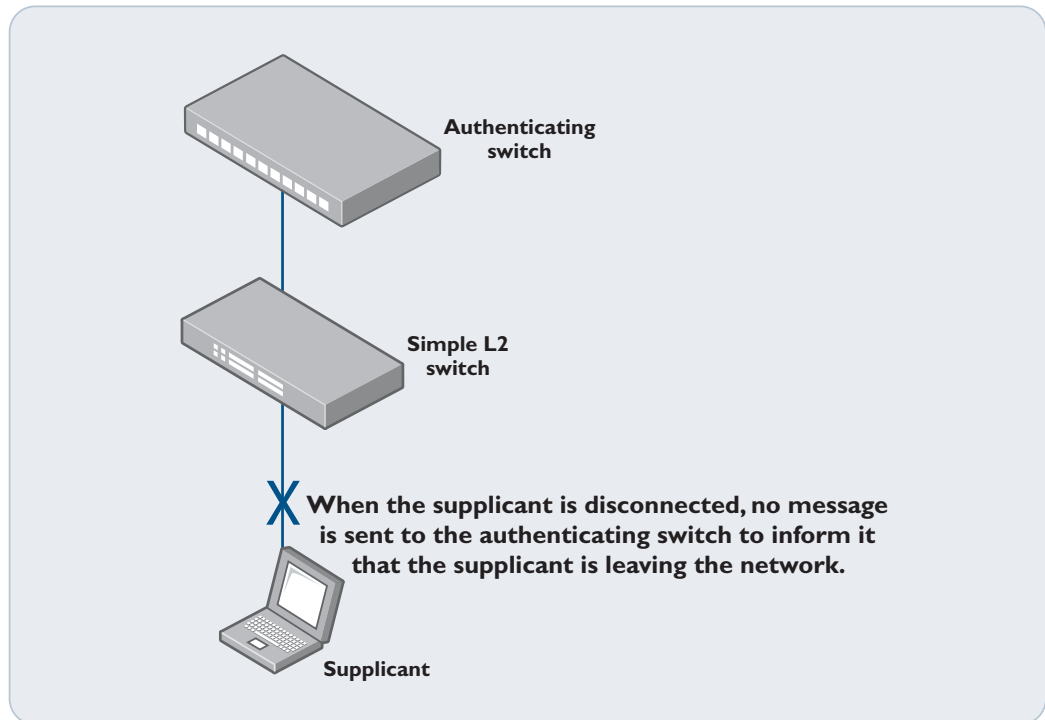
Ping-poll Monitoring of Supplicant Presence

A supplicant's authenticated session on the network must eventually come to an end. How does the authenticator decide that a supplicant's session has ended, and so remove it from the list of authenticated supplicants?

Sometimes it is obvious when the supplicant's session has ended, if the:

- supplicant unplugs from a port
- user clicks the **logout** button they were provided with on the "Authentication Success" as described in "[Starting a Web-authentication Session](#)" on page 7

Consider the case that a supplicant is not directly connected to the authenticating switch, but is connected to another switch that lies between itself and the authenticating switch, and the user simply disconnects their workstation.



If the network administrator wishes to ensure that the authenticating switch detects the supplicant's disconnection quickly, rather than waiting for the next expiration of the re-authentication period, then they can use **ping polling** to monitor the supplicants. This feature is enabled by the command:

```
auth-web-server ping-poll enable
```

Once ping polling has been enabled, the Web-authentication service will automatically ping-poll every Web-auth supplicant once they have been authenticated.

By default the ping-poll has a:

- polling **interval** of 30 seconds
- **timeout** of 1 second (i.e. the switch waits 1 second for the ping response before deciding the ping has failed)
- **failcount** of 5 (i.e. if a given supplicant fails to respond to five pings in a row, its authenticated session is terminated)

These default values can be altered by using the commands:

```
auth-web-server ping-poll interval <1-65535>
```

```
auth-web-server ping-poll timeout <1-30>
```

```
auth-web-server ping-poll failcount <1-100>
```

The currently configured values of these parameters can be seen by using the command:

```
show auth-web-server
awplus#show auth-web-ser
Web-authentication server
  <SNIP>
HTTP Redirect: enabled
  Session keep: enabled
  PingPolling: enabled
PingInterval: 30
  Timeout: 1
  FailCount: 5
  ReauthTimerRefresh: disabled
```

Checking the IP addresses of the supplicants

To verify the IP addresses of the supplicants that the switch is ping-polling, use the command **show auth-web supplicant brief**.

```
awplus#show auth-web supplicant brief
Interface port1.0.19
authenticationMethod: Web
  <SNIP>
```

| Interface | VID | Mode | MAC Address | Status | IP Address | Username |
|------------|-----|------|----------------|---------------|---------------|----------|
| port1.0.19 | 150 | W | 001c.7e95.d6bb | Authenticated | 192.168.150.9 | andrewr |

Ping-poll and promiscuous mode

Note that if you are using promiscuous mode to enable workstations with arbitrary IP addresses to be authenticated, there is no guarantee that the ping poll process will successfully send pings to those supplicants.

If the supplicant's IP address is not in the same subnet as the IP address the switch has on the VLAN that the supplicant is allocated into, then the ping poll process will not be able to send pings to that supplicant.

We recommend that you **do not** use the ping poll feature if you are using promiscuous mode, it risks the possibility of certain supplicants' authentication being terminated on a regular basis.

```
2010 Jun 4 18:46:09 user.notice awplus 802.1X[1044]: port1.0.19: Supplicant
andrewr logoff, Mac 001c.7e95.d6bb
2010 Jun 4 18:46:09 user.notice awplus 802.1X[1044]: port1.0.19: Supplicant
andrewr unauthorized, Mac 001c.7e95.d6bb
```

Managing Traffic of Unauthenticated Supplicants

The forwarding, blocking, and VLAN classification of traffic that arrives at the switch from unauthenticated supplicants is not entirely straightforward, and is subject to configuration. For the most part, the switch does not make a distinction between supplicants who have not yet attempted authentication, and those that have tried to authenticate, but failed.

The case in which there is a distinction drawn between those two classes of unauthenticated supplicant is when the auth-fail VLAN has been configured. Even when the auth-fail VLAN is not configured, the treatment of unauthenticated supplicants' traffic will differ, depending on whether or not the guest VLAN is configured.

We will take four different configuration combinations in turn, and look at the treatment of unauthenticated supplicants' traffic in each case of the cases where the port where the traffic arrives is configured with:

1. No Guest VLAN or Auth-fail VLAN
2. Guest VLAN, but no Auth-fail VLAN
3. Auth-fail VLAN, but no Guest VLAN
4. Auth-fail VLAN, and Guest VLAN

No Guest VLAN or Auth-fail VLAN

| Traffic Type | How Traffic is Processed |
|---|---|
| HTTP packets to Web-auth server address | Sent to CPU <ul style="list-style-type: none"> • Processed by Web-authentication |
| DHCP | Sent to CPU <ul style="list-style-type: none"> • Processed by Web-auth DHCP server (if configured) or • Processed by switch's standard DHCP server or • Relayed to another DHCP server |
| DNS | Sent to CPU <ul style="list-style-type: none"> • Intercepted by Web-authentication (if intercept configured) or • Forwarded to another DNS server or • Dropped |
| ARP | Sent to CPU. Supplicant's ARP is learnt. <ul style="list-style-type: none"> • Normal reply to ARP requests for switch's own IP address • Intercepted by Web-authentication (if intercept configured) or • Dropped |

| Traffic Type | How Traffic is Processed |
|---------------|--|
| Other packets | <p>If auth-Web forwarding configured</p> <pre>auth-web forward {arp dhcp dns ...}</pre> <p>packets matching criteria will be forwarded with native VLAN (not routed to other VLANs).</p> <p>All other packets dropped.</p> |

Guest VLAN but no Auth-fail VLAN

| Traffic Type | How Traffic is Processed |
|--|--|
| HTTP packets to Web-auth server address | <p>Sent to CPU</p> <ul style="list-style-type: none"> Processed by Web-authentication |
| DHCP | <p>Sent to CPU</p> <ul style="list-style-type: none"> Processed by Web-auth DHCP server (if configured) <p>or</p> <ul style="list-style-type: none"> Processed by switch's standard DHCP server <p>or</p> <ul style="list-style-type: none"> Relayed to another DHCP server if guest VLAN 'routing' enabled |
| ARP | <p>Sent to CPU. Supplicant's ARP is learnt.</p> <ul style="list-style-type: none"> Normal reply to ARP requests for guest VLAN IP address Intercepted by Web-authentication (if intercept configured) <p>or</p> <ul style="list-style-type: none"> Dropped |
| Other packets, not destined to switch's own IP address | <ul style="list-style-type: none"> L2 switched within guest VLAN L3 switched to other VLANs if guest VLAN 'routing' enabled <p>Note: If guest VLAN 'routing' option is configured, take care to use ACLs to constrain what packets can be routed where for example, you would probably configure ACLs to allow the traffic to be L3 forwarded to a DHCP server, a DNS server and possibly a NAC remediation server, or network domain controller.</p> |
| Other packets destined for switch's own IP address | <ul style="list-style-type: none"> Dropped |

Auth-fail VLAN, but no Guest VLAN

In this case traffic from supplicants who are deemed to have failed authentication is treated differently to traffic from supplicants who are deemed to not yet have fully tried authentication.

The definition of a supplicant having “failed authentication” is when the supplicant’s number of failed authentication attempts has reached the value configured by the command:

```
auth-web max-auth-fail <0-10>
```

By default, the value is 3.

Traffic from as-yet unauthenticated supplicants

This traffic is associated with the Native VLAN on the port on which the traffic arrives. It is treated in exactly the same way as described for the case of no Guest VLAN or auth-fail VLAN on [page 18](#).

Traffic from supplicants which has failed authentication

- This traffic is associated with the auth-fail VLAN configured on the ingress port
- The traffic is L2 switched within the auth-fail VLAN

Auth-fail VLAN, and Guest VLAN

Traffic from as-yet unauthenticated supplicants

This traffic is associated with the guest VLAN on the port on which the traffic arrives. It is treated in exactly the same way as described for the case of guest VLAN and no auth-fail VLAN.

Traffic from supplicants which has failed authentication

This traffic is associated with the auth-fail VLAN configured on the ingress port. It is treated in exactly the same way as described for the case of auth-fail VLAN and no guest VLAN.

Monitoring the Operation of Web-authentication

There is no specific debugging available for Web-authentication. The conversation between Web-authentication and a RADIUS server can be output by the command:

```
debug RADIUS all
```

An audit trail of Web-authentication events is kept in the system log. Successful and unsuccessful login attempts; and logoffs all generate entries in the system log.

```
2010 Jun  4 18:50:54 daemon.notice awplus radiusd[1712]: Login OK:
[andrewr] (from client 127.0.0.1 port 5019 cli 00-1c-7e-95-d6-bb)

2010 Jun  4 18:50:56 user.notice awplus 802.1X[1044]: port1.0.19: Web-
authentication successful for andrewr, IP 10.32.4.78, Mac 001c.7e95.d6bb

2010 Jun  4 18:52:31 daemon.notice awplus radiusd[1712]: Login incorrect:
[tester] (from client 127.0.0.1 port 5019 cli 00-1c-7e-95-d6-bb)

2010 Jun  4 18:52:33 user.notice awplus 802.1X[1044]: port1.0.19: Web-
authentication failed for tester, IP 192.168.101.6, Mac 001c.7e95.d6bb

2010 Jun 15 18:35:00 user.notice awplus 802.1X[1046]: port1.0.19:
Supplicant and rewr unauthorized, Mac 001c.7e95.d6bb
```

A list of all currently authenticated auth-Web supplicants can be seen from the commands:

```
show auth-web supplicant
show auth-web supplicant brief
```