

Securing The Network Edge

A Dell Technical White Paper



THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2011 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the *DELL* logo, and the *DELL* badge, *PowerConnect*, and *PowerVault* are trademarks of Dell Inc. *Symantec* and the *SYMANTEC* logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the US and other countries. *Microsoft*, *Windows*, *Windows Server*, and *Active Directory* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

September 2011

Contents

Introduction	2
Layer 2 Based Threats	3
Denial of Service through Traffic Storms	3
Storm Control	3
Denial of Service through MAC Learn Storms.....	4
MAC Learning Limit	4
Authenticating Users	5
Layer 3 Threats	6
Rogue DHCP Server	6
Rogue DHCP Clients.....	6
DHCP Snooping and Option 82	7
IP Spoofing.....	8
Source Address Validation	8
ARP Poisoning.....	9
Dynamic ARP Inspection	9
Network Access Policies and Controls	10
Summary	11

Figures

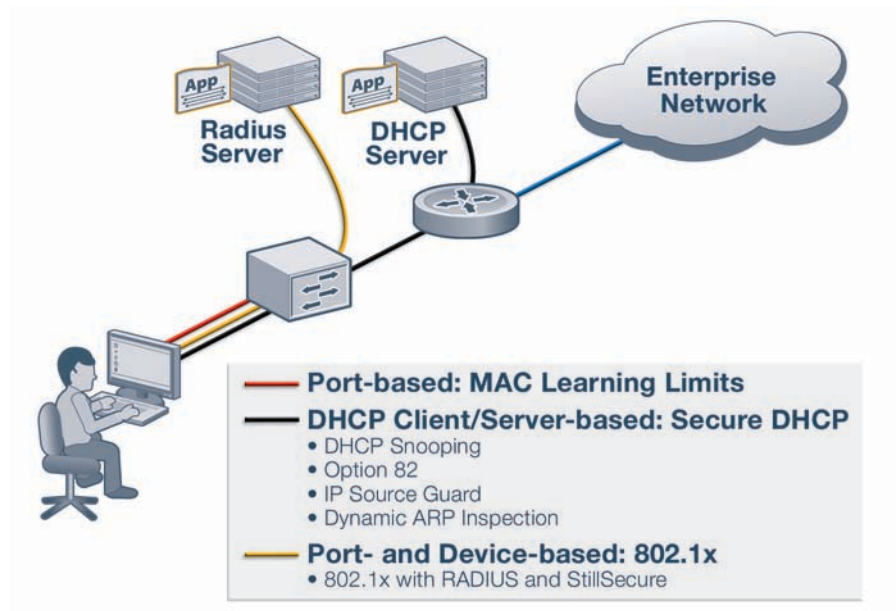
Figure 1. Security threats generally operate at one of two levels – Layer 2 or Layer 3	2
Figure 2. 802.1x uses a client-server based access control and authentication mechanism.....	5
Figure 3. DHCP servers are vulnerable to two types of attacks	6
Figure 4. DHCP snooping applies a trust model to switch ports	7
Figure 5. DHCP snooping protects pure Layer 2 devices.....	8
Figure 6. Source address validation (SAV) prevents nodes from spoofing their IP address and works in tandem with DHCP snooping	9
Figure 7. Dynamic ARP Inspection (DAI) protects against ARP poisoning.....	9
Figure 8. Dynamic ARP inspection prevents ARP spoofing	10
Figure 9. FTOS supports AAA attributes and a command to specify client Reauthentication.....	11

Introduction

Network security typically requires pulling together policies at Layer 2 and Layer 3 and at each topology layer to prevent malicious or inadvertent usage. The network boundary - where users enter the network - represents an effective security checkpoint to prevent a single endpoint device from either hijacking network resources or impacting other users' ability to access network services. Endpoint security traps the misbehaving endpoint at its closest point and minimizes the number of network links which must carry any malicious traffic.

Endpoint security also eases security policy creation since the endpoint switch should contain all required state information. In addition, it essentially distributes the state information, preventing scalability issues on interior network switches.

Figure 1. Security threats generally operate at one of two levels – Layer 2 or Layer 3



This paper describes common network threats and recommends configuration and management changes, including FTOS enterprise security features, which can be implemented on the access devices. While edge security measures are effective, threat models and operational issues required at other levels of the network also should be reviewed for completeness.

Layer 2 Based Threats

Denial of Service through Traffic Storms

Switched networks offer the advantage of isolating traffic to the required path and in turn minimizing the impact across the entire network. However, this advantage applies only to unicast traffic carrying traffic with destination addresses learned on network devices. A traffic storm sent to all paths of the network can be generated through the following types of traffic:

- Broadcast traffic
- Multicast traffic
- Unicast traffic addressed to a destination that is not present in the network i.e. a destination lookup failure (DLF)

The network devices switch these flows across all ports within the Layer 2 LAN, consuming significant bandwidth on all links.

Storm Control

The FTOS Storm Control feature limits the ability to generate such traffic flows by implementing a per-port rate limit at which such traffic is allowed. A switch which measures a received packet rate above the configurable limit starts dropping the excess flows. Depending on the Force10 Networks' platform, each port supports up to three, independent rates, one for each class of flooded traffic.

In the following sample configuration, a storm control rate of 100 packets per second is applied to unknown unicast traffic on interface GigabitEthernet 6/0.

```
Force10(conf)#interface GigabitEthernet 6/0
Force10(conf-if-gi-6/0)#storm-control unicast 100
Force10(conf)#end
Force10#show storm-control unknown-unicast gi 6/0
Unknown-unicast storm control configuration
Interface   Direction   Packets/Second
-----
Gi 6/0      Ingress     100
```

Denial of Service through MAC Learn Storms

Ethernet switches learn network nodes' locations by reading the source MAC address of incoming frames. A single node can generate frames with multiple source addresses, causing two issues in the switch and in the network:

- Multiple MAC addresses are learned on the same switch port, exhausting the hardware MAC address table memory. When this happens, the switch stops learning genuine MAC addresses, and legitimate frames become flooded.
- The switch must learn MAC addresses at a high rate, overloading the control plane and reducing the switch's ability to respond to other control and management traffic.

MAC Learning Limit

The FTOS MAC Learning Limit feature avoids the network impacts described above by placing a limit on the number of MAC addresses that each port can learn.

A switch flushes a learned MAC only after the MAC flush time expires, thus preventing the network node from contributing to an overwhelmed control plane resulting from a high MAC learning rate.

In this example, the "mac learning-limit 2" command in FTOS configures the switch to learn two MAC addresses for all VLANs on port GigabitEthernet 0/1.

```
Force10(conf-if-gi-0/1)#mac learning-limit 2
```

When the port reaches the learning limit and another frame with a new source MAC is received, depending on the configuration, the switch either logs the violation or shuts down the offending port.

```
Force10(conf-if-gi-6/0)#mac learning-limit learn-limit-violation log
```

```
Force10(conf-if-gi-6/0)#end
```

```
Force10#show mac learning-limit
```

Interface	Learning	Dynamic	Static	Unknown SA
Slot/port	Limit	MAC count	MAC count	Drops
Gi 6/0	2	2	0	0
Gi 6/47	3	3	0	0

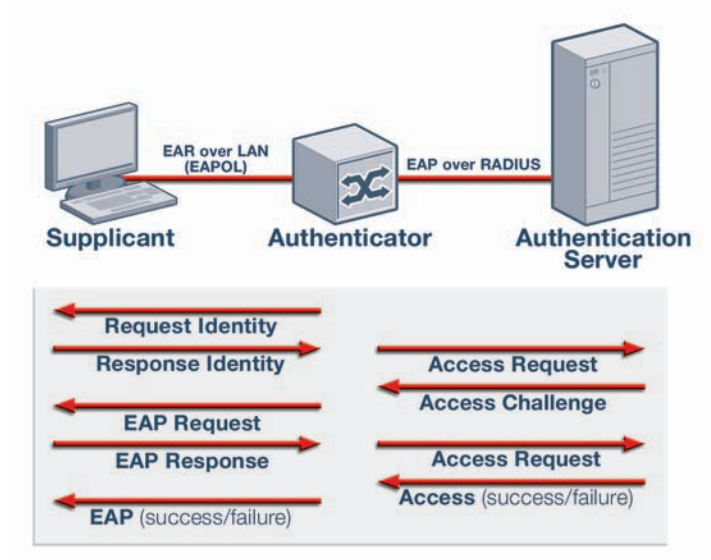
```
Force10#show mac-address-table
```

VlanId	Mac Address	Type	Interface	State
1	00:00:00:00:00:aa	Dynamic	Gi 6/0	Active
1	00:00:00:00:00:ab	Dynamic	Gi 6/0	Active
1	00:00:00:00:00:bb	Dynamic	Gi 6/47	Active
1	00:00:00:00:00:bc	Dynamic	Gi 6/47	Active
1	00:00:00:00:00:bd	Dynamic	Gi 6/47	Active

Authenticating Users

The IEEE 802.1x standard specifies a port-based authentication scheme, particularly useful for publicly accessible ports. As shown in figure 2, it uses a client-server based access control and authentication mechanism which prevents un-authorized clients from gaining access to the network.

Figure 2. 802.1x uses a client-server based access control and authentication mechanism



The role of each device in 802.1x authentication is explained below:

- Client (Supplicant) - the device attempting to receive services from the network and requiring authentication.
- AA Server (Authentication Server) - the server which stores user credentials and the specific authentication mechanism to apply.
- Switch (Authenticator) - the device through which the client gains access to services from the network. The switch should ensure that the client is authenticated before allowing it to access the resources of the network.
-

The AAA server and the switch communicate via a protocol such as RADIUS.

With 802.1x port authentication enabled, a port starts in an un-authorized state and allows only authentication (IEEE 802.1x) messages. The switch receives these messages from the client and forwards them on to the authentication server. If the authentication server requires more information, it queries the client, with the switch relaying these messages. When the message exchange completes, the authentication server determines whether the client is an authenticated user and then informs the switch via the AAA protocol. The switch uses this information to either allow the client complete access to the network or block the port, preventing network access.

Enabling 802.1x authentication on a switch running FTOS involves two steps:

1. Configuring the 802.1x parameters themselves.
2. Configuring the switch with the address of the backend AAA server and the protocol to use to reach to the server.

In the following example, interface GigabitEthernet 0/1 is configured as a Layer 2 switchport with 802.1x authentication. The switch uses RADIUS as the AAA protocol.

```
Force10(conf)#radius-server host 10.1.1.15  
Force10(conf)#interface gigabitethernet 0/1  
Force10(conf-if-gi-0/1)#switchport  
Force10(conf-if-gi-0/1)#dot1x authentication
```

Layer 3 Threats

Rogue DHCP Server

DHCP provides a common mechanism for assigning IP addresses to machines on the LAN. As it comes up, a host sends a broadcast packet to identify the DHCP servers on the network. In the event of multiple responses, it chooses one, typically the first. A node can disrupt this process and network communication by purposely giving out bogus addresses as a DHCP server.


Rogue DHCP Clients

A DHCP server serves a set of LANs. It allocates IP addresses to clients from pools it is configured to own. The server generally attempts to allocate the same address to the same client across multiple assignments through the use of the “Client Identifier” field in the DHCP Request packet.

Two types of attacks on the DHCP server are possible:

- Exhaust the IP address pool by sending multiple DHCP requests, each with a unique “Client Identifier” value. This approach forces the DHCP server to hand out multiple addresses to this node, ultimately exhausting the dynamic address pool and depriving other legitimate users of access to the network.
- Masquerade as legitimate clients by sending “Client Identifier” values of these clients. The DHCP server then allocates the address that normally would have been assigned to that client to the rogue node. In other words, the rogue node receives an IP address associated with a legitimate client on the network.

Figure 3. DHCP servers are vulnerable to two types of attacks



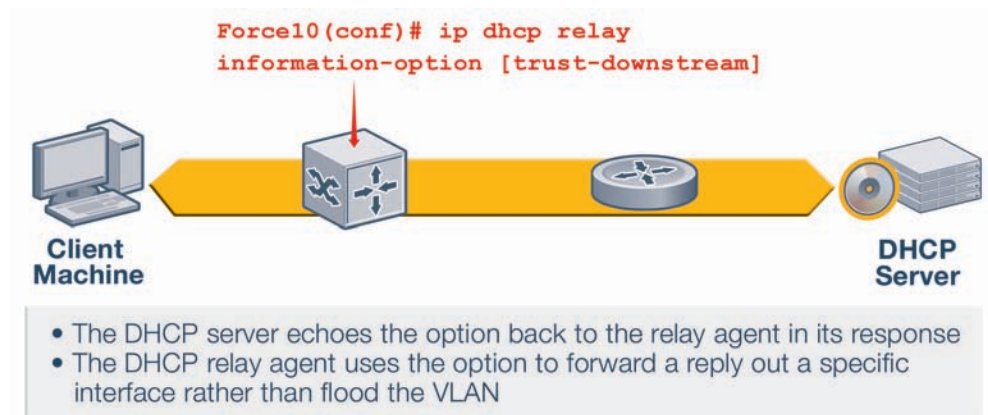
DHCP Snooping and Option 82

The FTOS DHCP snooping and DHCP relay option 82 features prevent DHCP server attacks. These features are enabled on an FTOS switch acting as a DHCP relay agent, which traps the DHCP packets and forwards them as unicast packets to a DHCP server on a different IP network.

As described earlier, the DHCP server uses the “Client Identifier” field as a mechanism to identify the client and assign the appropriate IP address. With the option 82 feature, the relay agent appends the client’s DHCP packet with the option 82 TLV, which describes the client’s network point of access. The server can use this information to allocate the IP address, rather than the “Client Identifier”.

DHCP snooping applies a trust model to switch ports. By default, all ports are untrusted, and the switch denies such ports to send DHCP server packets. Thus, network nodes connected to these ports are prevented from masquerading as a DHCP server. The actual DHCP server port and the inter-switch links are configured as trusted ports to allow DHCP server packets to ingress the network as required.

Figure 4. DHCP snooping applies a trust model to switch ports



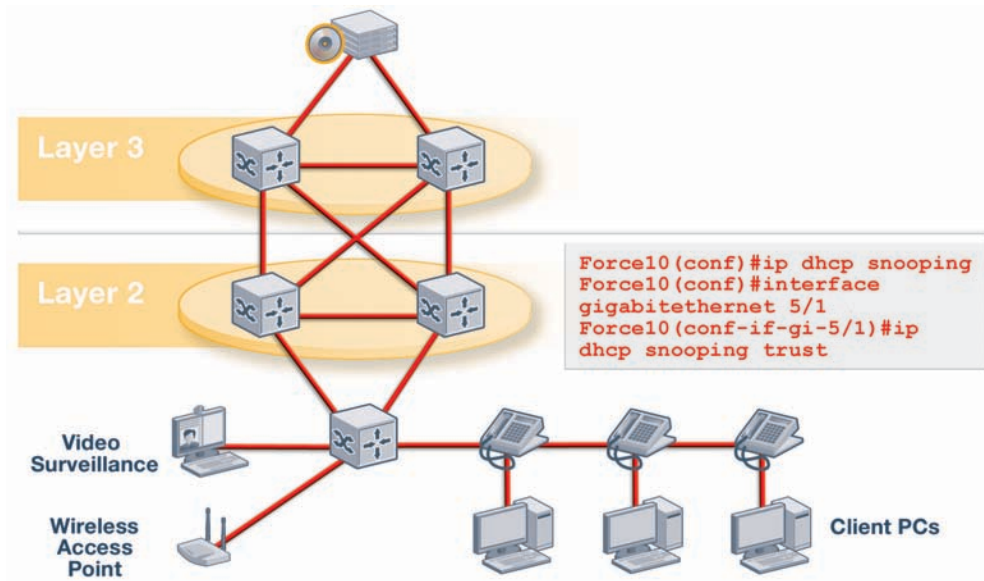
With DHCP snooping enabled, the switch builds a snooping table, which it populates with MAC Address, Port Number, and IP Address. This table is built as follows:

- When the node comes up, it sends a DHCPDISCOVER message. One or more DHCP servers respond with a DHCPOFFER message.
- The client chooses an offer and sends the DHCPREQUEST message. The server responds with a DHCPACK message.
- During this exchange, the switch snoops the DHCP messages. When the DHCPACK message is received, the switch creates an entry in the DHCP snooping table, tying a MAC address and IP address pair to a port.
- The switch removes this entry when the node sends a DHCPRELEASE message or when the entry times out after the lease time granted by the DHCP server.

Other security features, including dynamic ARP inspection and source address validation, use the snooping table information.

In the following diagram, multiple client devices connect to the pure Layer 2 switches, which are enabled with DHCP snooping and DHCP relay with option 82. These Layer 2 switches connect to Layer 3 switches, which connect to the DHCP server.

Figure 5. DHCP snooping protects pure Layer 2 devices



IP Spoofing

IP spoofing describes the technique in which a client sends packets with a non-legitimate source IP address. This technique allows the client to perform various types of attack both on the destination node and the source node that it is masquerading. Such IP spoofing is difficult to detect at places further from the source of the packet.

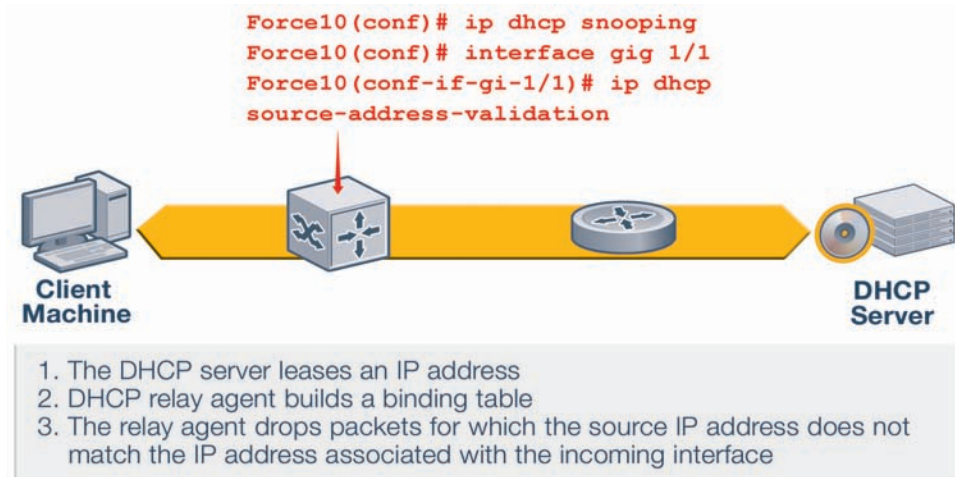
Source Address Validation

Source address validation (SAV) prevents nodes from spoofing their IP address and works in tandem with DHCP snooping. With SAV, the switch allows only incoming DHCP packets on a port. When the client receives an IP address assigned by the DHCP server, the switch uses DHCP snooping to learn this address and adds an ACL entry to the IP to MAC + port number database. The ACL allows only packets with the source IP address on the port and denies all other source IPs.

FTOS supports two types of SAV:

- **IP only** - the ACL entry filters based on the source IP address and port.
- **IP+ MAC** - the ACL entry filters based on the source IP, source MAC and port.

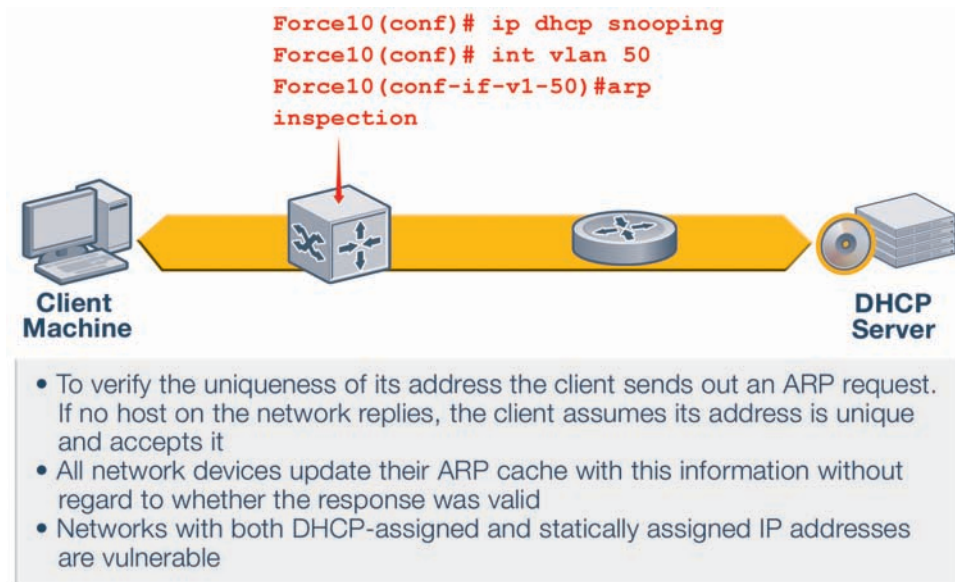
Figure 6. Source address validation (SAV) prevents nodes from spoofing their IP address and works in tandem with DHCP snooping



ARP Poisoning

ARP is the mechanism by which network nodes learn the mapping between the IP address and the Ethernet address of the corresponding node on the LAN. It is possible that a node can capture packets meant for another node by sending spurious ARP packets.

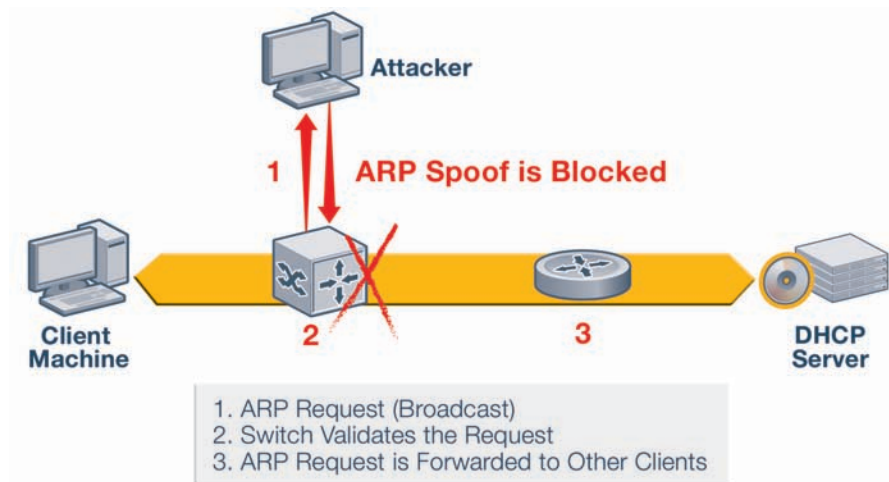
Figure 7. Dynamic ARP Inspection (DAI) protects against ARP poisoning



Dynamic ARP Inspection

Dynamic ARP Inspection (DAI) protects against ARP poisoning and works in tandem with DHCP snooping. On untrusted ports, the switch captures all ARP packets (both request and reply) and then validates the Source Protocol and Source Hardware address values against the snooping table database for that port. If the MAC address and IP address and the corresponding port do not match the snooping database entry, the ARP packets are dropped. DAI thus prevents the node from specifying a non-legitimate IP-MAC address binding which differs from what was given by the DHCP server.

Figure 8. Dynamic ARP inspection prevents ARP spoofing



Network Access Policies and Controls

The above mechanisms provide protection against a device which is trying to perform malicious operations on the network. However, there are circumstances where the authentication used by the above mechanism passes, but the endpoint remains a threat to the network. For example:

1. The endpoint does not have the latest security patches installed.
2. The endpoint is not running any anti-virus software or the latest virus signature files are not installed.
3. The endpoint is currently infected and may spread that infection.

Thus, a network security policy may require endpoint verification of conformance to these policies. Any non-conforming endpoints are placed in a quarantine network (VLAN) and the system administrator informed of the same for corrective actions.

StillSecure from Safe Access provides endpoint verification services, including:

1. Pre-entry authentication
2. Pre-entry testing and validation
3. Post-entry monitoring and validation

When an endpoint requests network access, it goes through the following sequence:

1. The endpoint sends a DHCPREQUEST for an IP address.
2. StillSecure traps the request and sends a temporary IP address in the quarantine VLAN.
3. StillSecure authenticates the endpoint using the IP address in the Quarantine VLAN.
4. If authentication passes, StillSecure then tests the endpoint for conformance to network policies.
5. If all tests pass, the endpoint is asked to renew the IP address again.
6. The endpoint sends a fresh DHCPREQUEST for an IP address.
7. StillSecure receives the request and passes it the DHCP server, which assigns an IP address, as expected.
8. This address is passed back by StillSecure to the endpoint and also saved for later periodic validation and monitoring of the endpoint (if so configured).

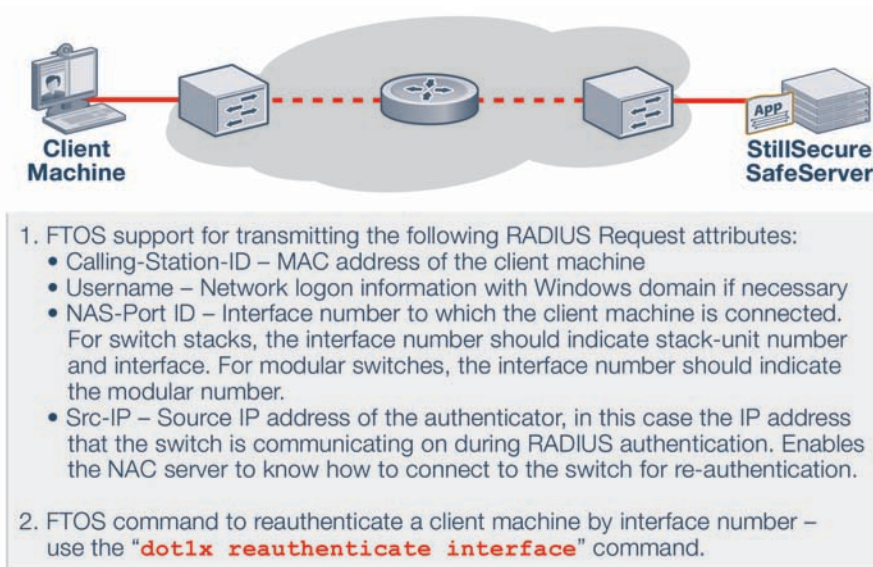
Safe Access provides connectivity options with both DHCP and IEEE 802.1x. It also supports periodic endpoint testing to ensure that it detects a device which at first is authenticated as safe but then gets infected at a later point of time.

FTOS integrates with SafeAccess products for a complete endpoint verification to network policies solution. StillSecure works by trapping the AAA request, processing it, and, if necessary, passing it to the actual RADIUS server. FTOS adds the following AAA attributes to the request:

1. Calling-Station-ID - the attribute which carries the MAC address of the client.
2. User-name and User-Password - the client's network login information, such as the Windows login details
3. NAS-IP - the IP address of the switch to which the client is connected
4. NAS-Port - the switch port identifier to which the client is connected

These attributes identify the endpoint in detail, and the network administrator then leverages StillSecure to define a specific policy for that endpoint.

Figure 9. FTOS supports AAA attributes and a command to specify client Reauthentication



In addition, the administrator can force the clients to re-authenticate periodically by controlling the session timeout parameter. If the client does not authenticate successfully within this interval, the switch disables access to the client. The FTOS command "dot1x reauthenticate interface" specifies that the connected endpoint must re-authenticate.

Summary

The right network security mechanism depends on the required security level and on the types of available network services. This paper identified common Layer 2 and Layer 3 security threats at the network access point and the associated features in FTOS to safeguard against such threats. Several FTOS features may be necessary depending on the required security level.