>THIS IS **THE WAY**

>THIS IS NORTEL™

Product Name Ethernet Switch (ES) and Ethernet Routing Switch (ERS)
Product Number

> Layer Security Solutions for ES and ERS Switches

Enterprise Solutions Engineering
Document Date: August 16, 2005
Document Version: 1.0

Copyright © 2005 Nortel Networks

Trademarks

Abstract

This document provides various solutions for Nortel's ES and ERS switches in reference to Layer 2 security.

Table of Contents

1. Overview: Layer 2 Security

The purpose of this document is to review various security attacks that could occur in a Layer 2 network. Of concern are the following attacks that could occur in a Layer 2 environment:

- MAC table overflow
- VLAN hopping
- Spanning-Tree Protocol manipulation
- MAC Address/ARP Spoofing
- Private VLAN
- DHCP starvation

2. MAC Table Overflow:

On a Layer 2 switch, a MAC table is used to record the port where a given MAC address is connected. This allows a Layer 2 switch to switch unicast traffic between two hosts without other hosts seeing the traffic.

The following illustrates how a MAC table is supposed to work:

- A MAC address table on a switch is limited in size depending on the switch from perhaps few hundred MAC entries to perhaps over 100,000 entries
- A MAC aging timer is used to age out MAC entries. Each time a frame is transmitted with a source MAC matching an entry in the MAC table, the aging timer is reset. If a given host does not send frames within the MAC age timer, it will time out the older MAC address entry.

Given how a MAC table works, the following points are what can happen during a MAC table attack:

- An attacker sends a continuous set of frames with random source MAC addresses and random destination MAC addresses.
- As the MAC table is limited in size, eventually the switch will run out of room for new MAC entries. Depending on the switch, this could use up all the entries for the whole switch or the portion assigned to the VLAN where the attacker comes in on.
- When all MAC entries have been exhausted, the switch will flood all ports with incoming traffic. This is because it cannot find the port number for a particular MAC address in the MAC table resulting in the switch, in essence, to behave like a hub. This only occurs in the VLAN used by the attacker.
- The attacker will then see all traffic from host within the VLAN; this could include user names, passwords, and so forth which allow the attacker to place the next attack.

There are several tools available that will generate multiple random source and destination MAC and IP addresses. The most common is macof written in approximately 100 lines of PERL. This was then later ported to C language code into dsniff tools. These tools are capable of generating greater than 100,000 MAC entries per minute.

Note that there are a few caveats to the attack that you should be aware of:

- Even with a full MAC table, traffic is flooded only on the local VLAN where the attack is taking place. This means if the attacker is on VLAN 10, only traffic in VLAN 10 is flooded, however, everyone with a port in VLAN 10 will see the traffic.
- Another adjacent switch could also be affected because of the flooding.
- A DoS condition in the network could also occur.

Nortel Solution

Nortel has two solutions for this type of attack, Discard Unknown MAC Security and MAC limitation.

MAC Limiting

The ERS 8600 supports a maximum MAC count on a per port basis in release 3.7.2. For ES, BOSS 3.6 and BOSS 4.2 have this feature. For the ERS 8300, release 2.2 has this feature.

Discard Unknown MAC Security and BaySecure

Unknown MAC Security, although it does limit MAC learning, it may not be the best solution for all applications. With this feature, you have to go to each port and either manually enter a MAC address or stop/start auto-lean and lock the port.

a) ERS8600/ERS8300

The ERS 8300 and ERS 8600 switches support Discard Unknown MAC Security, which ensures that any frames bearing a MAC address that is not included in the allowed list of MACs are discarded. This feature is typically implemented in high security environments, where the network administrator wants to strictly enforce which specific devices are allowed to communicate on a given port. Also, the number of MAC addresses allowed can be configured using this parameter.

Source MAC address-based VLANs can be used in security-conscious environments, where users want to allow and disallow communication based on device MAC addresses.

The following commands are used to configure unknown MAC discard:

ERS8600 or ERS8300:

CLI:

Passport-8606:5# ***config ethernet <slot/port> unknown-mac-discard ?***
or
Passport-8310:5# ***config ethernet <slot/port> unknown-mac-discard ?***

```
Sub-Context:
Current Context:

    activation <enable|disable>
    add-allow-mac <mac> [auto]
    autolearn <enable|disable>
    autolearn-mode <one-shot|continuous>
    info
    lock-autolearn-mac <enable|disable>
    max-mac-count <max MAC count>
    remove-allow-mac <mac>
    violation-downport <enable|disable>
    violation-logging <enable|disable>
    violation-sendAuthenticationTrap <enable|disable>
```

JDM:

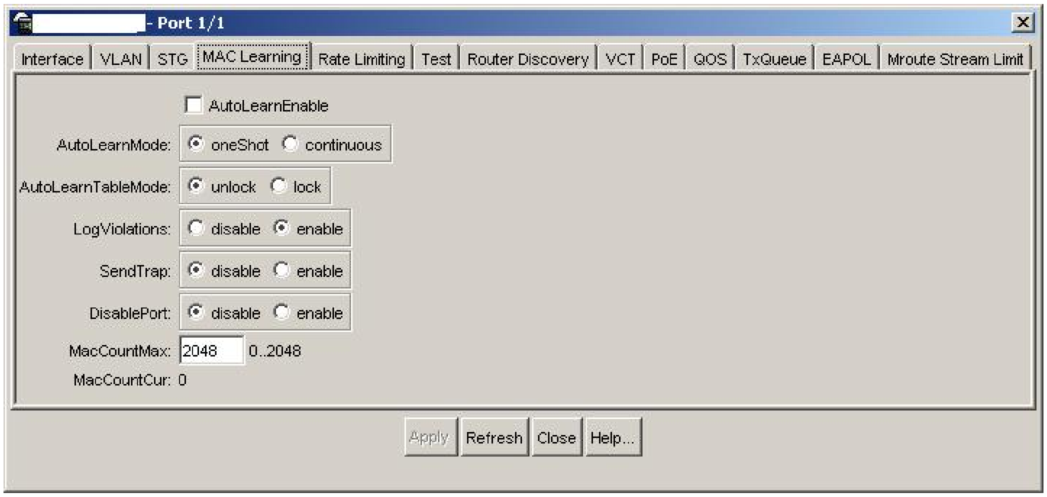

□ AutoLearnEnable

AutoLearnMode: ● oneShot ○ continuous

AutoLearnTableMode: ● unlock ○ lock

LogViolations: ○ disable ● enable

SendTrap: ● disable ○ enable

DisablePort: ● disable ○ enable

MacCountMax: 2048 0..2048

MacCountCur: 0

Field	Description
AutoLearnEnable	Sets the port to autolearn addresses for the allowed MAC table.
AutoLearnMode	Sets the autolearn mode on the port for populating the allowed MAC table.
AutoLearnTableMode	Sets the allowed MAC table to current state. When locked, no new MAC addresses will be learned.
LogViolations	Enables the system to create a system log entry when a disallowed MAC address attempts to send traffic through the selected port.
SendTrap	Indicates whether a trap should be sent to the management station when a MAC address violation is detected on the selected port. The default is disable.
DisablePort	Indicates whether the selected port should be disabled if a MAC address violation is detected. Enable means that the port should be disabled if this event occurs. The default is disable.
MacCountMax	The maximum number of MAC addresses that can be added to the selected port. The valid values are 0 to 2048.
MacCountCur	The current number of MAC addresses that have been added to the selected port.

b) ES470/460 and ERS55xx:

For the ES570/460, BOSS 3.6 supports MAC Address Security with a feature to enable the maximum number of MAC addresses per port. For the ERS5510, BOSS 4.2 supports MAC Address Security with a feature to enable the maximum number of MAC addresses per port.

The following is an example how to configure MAC address limit using port 18 on an ES470 with a limit of three MAC addresses. There will be a TCG on this topic when BOSS 3.6 is released.

- 470-24T(config)#***interface fastEthernet 18***
- 470-24T(config-if)#***mac-security port 18 enable***
- 470-24T(config-if)#***mac-security auto-learning port 18 max-addrs 3***
- 470-24T(config-if)#***mac-security auto-learning port 18 enable***

- 470-24T(config-if)#***exit***
- 470-24T(config)#***mac-security enable***

The ES470 also supports BaySecure, via CLI or WEB, which is similar to the ERS 8600/8300 Unknown MAC Discard feature.

The following command is used to configure BaySecure on the ES470:

BS470_48(config)#***mac-security ?***

```
  disable           Disable MAC Address Security.
  enable            Enable MAC Address Security.
  filtering         Enable/disable DA filtering
  intrusion-detect  Enable/disable partitioning on intrusion detection
  intrusion-timer   Set temporary partition time for intrusion detection.
  learning          Enable/disable MAC address learning
  learning-ports    Modify ports participation in MAC address learning.
  mac-address-table Add addresses to MAC security address table
  mac-da-filter     Add/delete MAC DA filtering addresses
  security-list     Modify security list port membership.
  snmp-lock         Enable/disable SNMP lock on MAC address security
                    parameters.
  snmp-trap         Enable/disable SNMP trap generation on intrusion
                    detection.
```

BS470_48 (config-if)#***mac-security ?***

```
  disable   Disable MAC security for port(s)
  enable    Enable MAC security for port(s)
  learning  Enable MAC security address learning for port(s)
  port      Port number(s) on which to enable/disable MAC-based security
```

where:

Parameters and Variables	Description
disable|enable	Disables or enables MAC address-based security.
filtering {enable|disable}	Enables or disables destination address (DA) filtering on intrusion detected.
intrusion-detect {enable|disable|forever}	Specifies partitioning of a port when an intrusion is detected: enable—port is partitioned for a period of time disabled—port is not partitioned on detection forever—port is partitioned until manually changed
intrusion-timer <*1-65535*>	Specifies, in seconds, length of time a port is partitioned when an intrusion is detected; enter the number of you want.
learning-ports <portlist>	Specifies MAC address learning. Learned addresses are added to the table of allowed MAC addresses. Enter the ports you want to learn; it can be a single port, a range of ports, several ranges, all, or none.
learning {enable|disable}	Specifies MAC address learning: enable—enables learning by ports disable—disables learning by ports
snmp-lock {enable|disable}	Enables or disables a lock on SNMP write-access to the BaySecure MIBs.
snmp-trap {enable|disable}	Enables or disables trap generation upon intrusion detection.

3. DHCP Attacks:

There are two types of attacks that can occur with DHCP:

- An attacker could request multiple IP address from a DHCP server by spoofing its source MAC address. This can be achieved by using a tool such as gobbler: http://www.networkpenetration.com/downloads.html. If the attack is successful, all leases on the DHCP server will be exhausted.
- The second method is where the network attacker sets up a rogue DHCP server and responds to new DHCP requests from clients on the network. The attackers DHCP server could be setup to send DHCP responses using its address for the default gateway and DNS server. This would allow the attacker to sniff out the client's traffic and allowing for a 'man-in-the-middle' attack.

Nortel Solution

As is the case with MAC Table Overflow type of attack, setting a MAC address limit can be used to solve this problem.

Another solution would be to use filters to prevent spoofing of the DHCP server. This could be accomplished by simply creating an IP filter with a src-ip of the DHCP server with an action of drop and apply the filter to all the user ports only.

Listed below are configuration examples assuming the actual DHCP server address is 172.30.30.50.

ERS8600:

```
#
# TRAFFIC-FILTER CONFIGURATION
#

ip traffic-filter create global src-ip 0.0.0.0/0.0.0.0 dst-ip 0.0.0.0/0.0.0.0 id 2
ip traffic-filter filter 2 action mode drop
ip traffic-filter filter 2 match dst-port 68 dst-option equal
ip traffic-filter filter 2 match protocol udp
ip traffic-filter filter 2 name "ROGUE-DHCP"
ip traffic-filter create global src-ip 172.30.30.50/255.255.255.255 dst-ip 0.0.0.0/0.0.0.0 id 1
ip traffic-filter filter 1 action mode forward
ip traffic-filter filter 1 match dst-port 68 dst-option equal
ip traffic-filter filter 1 match protocol udp
ip traffic-filter filter 1 name "ROGUE-DHCP"
ip traffic-filter global-set 1 create name "ROGUE-DHCP_set"
ip traffic-filter global-set 1 add-filter 1
ip traffic-filter global-set 1 add-filter 2

#
# PORT CONFIGURATION - PHASE II
#

ethernet 1/1-1/48 ip traffic-filter create
ethernet 1/1-1/48 ip traffic-filter add set 1
ethernet 1/1-1/48 ip traffic-filter default-action forward
```

ES470:

```
!
! *** QOS ***
!
qos ip-filter 1 create src-ip 172.30.30.50 255.255.255.255 dst-ip 0.0.0.0 0.0.0.0  protocol udp dst-port 68
qos ip-filter 2 create src-ip 0.0.0.0 0.0.0.0 dst-ip 0.0.0.0 0.0.0.0  protocol udp  dst-port 68
qos ip-filter-set 1 create set 1 name "IP_class_1" filter 1 filter-prec 1
qos ip-filter-set 2 create set 2 name "IP_class_2" filter 2 filter-prec 1
qos policy 1 create name "Rogue_1" if-group allBPSIfcs filter-set-type ip filter-set 1 in-profile-action 65527
order 1
qos policy 2 create name "Rogue_2" if-group allBPSIfcs filter-set-type ip filter-set 2 in-profile-action 65526
order 2
```

ERS5510:

In release BOSS 4.2 for the ERS5510, you can enable DHCP Spoofing Detection via the WEB interface or by the CLI interface as shown below.

CLIexit

- 5510-48T(config)#***interface fastEthernet <ALL|port|slot/port>***
- 5510-48T(config-if)#***qos dhcp {snooping | spoofing} port <port_list> enable interface-type {access | core}***

where

Parameter	**Description**
{snooping | spoofing}	The type of QoS DHCP application to enable.
port <port_list>	The ports to enable the selected QoS DHCP application on.
interface-type {access | core}	The interface type to use.

WEB

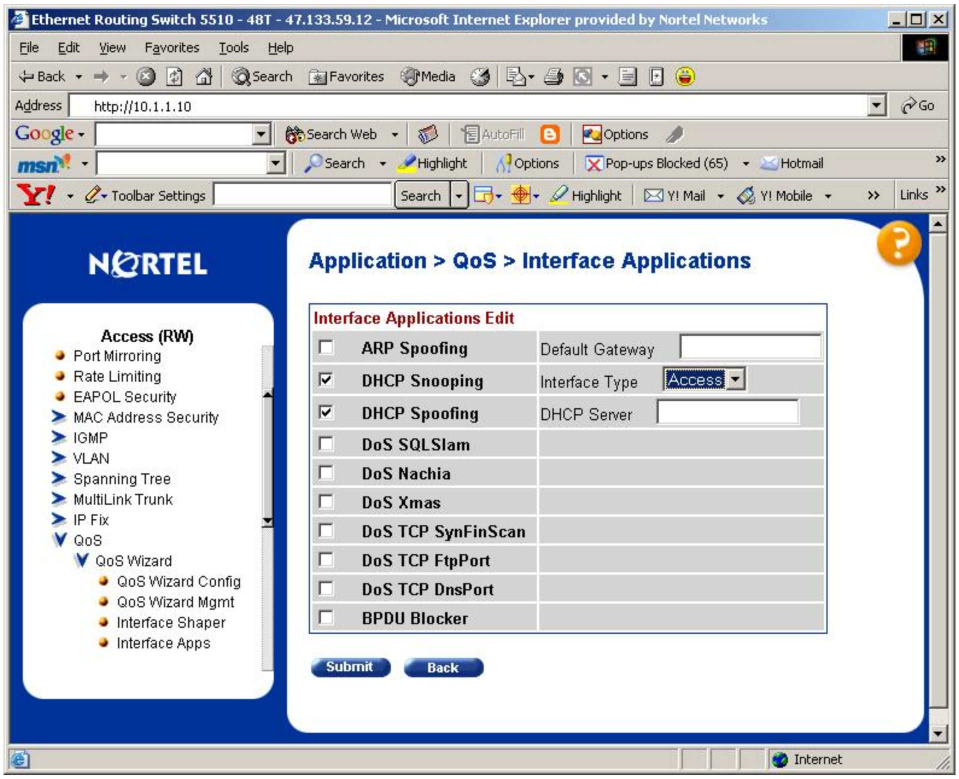
Ethernet Routing Switch 5510 - 48T - 47.133.59.12 - Microsoft Internet Explorer provided by Nortel Networks
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://10.1.1.10 Go
Google Search Web AutoFill Options
msn Search Highlight Options Pop-ups Blocked (65) Hotmail
Y! Toolbar Settings Search Highlight Y! Mail Y! Mobile Links
Internet
NORTEL
Access (RW)
Port Mirroring
Rate Limiting
EAPOL Security
MAC Address Security
IGMP
VLAN
Spanning Tree
MultiLink Trunk
IP Fix
QoS
QoS Wizard
QoS Wizard Config
QoS Wizard Mgmt
Interface Shaper
Interface Apps
Submit
Back

Application > QoS > Interface Applications

Interface Applications Edit

☐ ARP Spoofing	Default Gateway	
☑ DHCP Snooping	Interface Type	Access
☑ DHCP Spoofing	DHCP Server	
☐ DoS SQLSlam		
☐ DoS Nachia		
☐ DoS Xmas		
☐ DoS TCP SynFinScan		
☐ DoS TCP FtpPort		
☐ DoS TCP DnsPort		
☐ BPDU Blocker		

In release BOSS 4.1, you have to manually configure a filter. Shown below is an example on how to setup the filter on the ERS55xx.

```
!
! *** QOS ***
!
qos action 10 name "Standard_match" drop-action disable update-1p 0
qos ip-element 1 src-ip 172.30.30.50/32 protocol 17 dst-port-min 68 dst-port-max 68
qos ip-element 2 protocol 17 dst-port-min 68 dst-port-max 68
qos classifier 1 set-id 1 name c1 element-type ip element-id 1
qos classifier 2 set-id 2 name c2 element-type ip element-id 2
qos policy 1 name "Rogue_1" if-group allBayStackIfcs clfr-type classifier clfr-id 1 in-profile-action 10 precedence 10 track-statistics individual
qos policy 2 name "Rogue_2" if-group allBayStackIfcs clfr-type classifier clfr-id 2 in-profile-action 1 precedence 9 track-statistics individual
```

ERS 8300:

```
#
# TRAFFIC-FILTER CONFIGURATION
#

filter acl 1 ip  acl-name "acl_1"
filter acl 1 action 1 permit "acl_1"
filter acl 1 ip-hdr 1 src-ip 172.30.30.50 ipfragment non-fragments
filter acl 1 protocol 1 udp eq any
filter acl 1 port 1 dst-port bootpd-dhcp
filter acl 1 action 2 deny "acl_2"
filter acl 1 debug 2 match-count bytes-pkts
filter acl 1 ip-hdr 2 ipfragment non-fragments
filter acl 1 protocol 2 udp eq any
filter acl 1 port 2 dst-port bootpd-dhcp
filter acl 1 action default permit "default ace"
filter acg 1 1 acg-name "acg_1"

#
# PORT CONFIGURATION - PHASE II
#

interface FastEthernet 1/1  filter 1
interface FastEthernet 1/2  filter 1
```

4. Spanning Tree (STP):

Spanning Tree, available in several versions such as 802.1d, RSTP, or MSTP, all serve the same purpose and that is to avoid forwarding loops in Layer 2 networks.

As far as security is concerned, STP offers no provisioning for authentication for the bridge protocol data units (BPDUs) that are sent between switches and bridges. Knowing this, an attacker could easily send BPDUs into a network and cause any number of undesirable effects. For example, an attacker could connect to two switches in the network, send out BPDUs to cause recalculations, and become a root bridge. The attacker simply has to send BPDUs with a low bridge priority and if it indeed has the lowest bridge priority or if tied, has the lowest MAC address; it will become the root bridge. The end result is the traffic must pass between the switches in the network and the attacker's device causing a man-in-the-middle attack.

To prevent an STP attack, you could simply disable STP on all end-user ports. The only problem with this solution is an attacker could still cause an attack by introducing a loop.

Nortel Solution

Split-MLT

The easiest solution would be to not use Spanning Tree at all. With Split-MLT, you can deploy fully redundant, multi-homed, load balanced networks without using Spanning Tree at all. Compare this to Spanning Tree using a single instance, 802.1d or RSTP, where at least half of all available links are blocked to prevent loops. Or compared to Spanning Tree with multiple instances, MSTP, where you can only provide VLAN load balance. VLAN load balance is possible with MSTP by playing around with bridge and possibility link priorities on each switch for each VLAN making network configuration very complex and harder to trouble-shoot.

Spanning Tree

Disabling Spanning Tree on a per port basis for all user access port is an easy to implement solution. The only problem being is that an attacker could create a loop by making a connection between two ports on the same switch or another switch.

Another solution could be to use a different STG multicast address, i.e. a separate STG group with a new BPDU multicast address. If an attacker places a switch on the network running normal 802.1d Spanning Tree, this would allow the attacker STG instance from the host or switch connected to the Nortel switch to pass through the Nortel switched network as normal data.

Finally, MAC filters could be set up on a Nortel ES switch to filter on normal STP BPDU's and on any proprietary Cisco STP protocols.

In the BOSS 4.2 release for the ERS55xx switch, Spanning Tree (both 802.1d and PVST+) can be blocked by checking the BPDU Blocker security button via the WEB interface or by using the CLI interface as shown below.

CLI Command:

- 5510-48T(config)#***interface fastEthernet <ALL|port|slot/port>***
- 5510-48T(config)#***qos bpdu blocker port <port_list> enable***

WEB Interface :

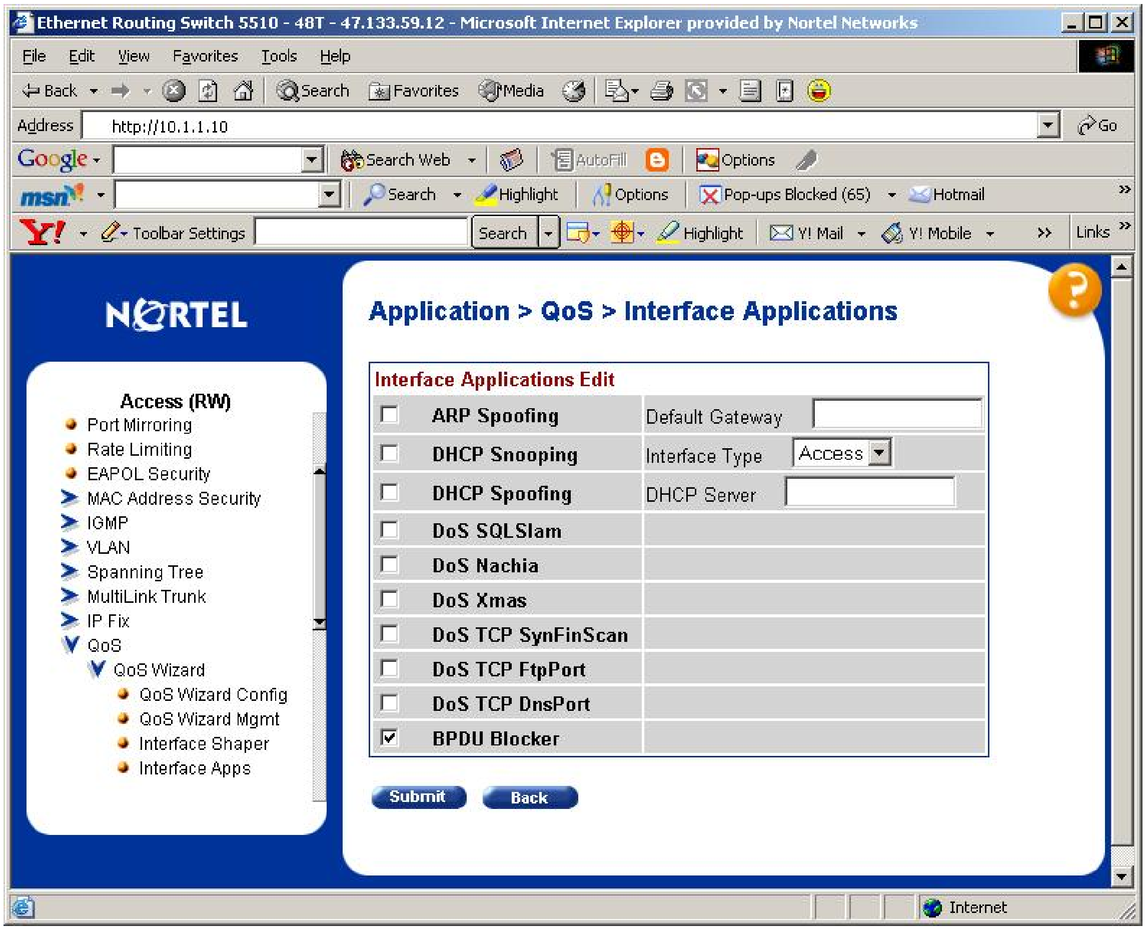
Ethernet Routing Switch 5510 - 48T - 47.133.59.12 - Microsoft Internet Explorer provided by Nortel Networks
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://10.1.1.10 Go
Google Search Web AutoFill Options
msn Search Highlight Options Pop-ups Blocked (65) Hotmail
Y! Toolbar Settings Search Highlight Y! Mail Y! Mobile Links
NORTEL
Access (RW)
Port Mirroring
Rate Limiting
EAPOL Security
MAC Address Security
IGMP
VLAN
Spanning Tree
MultiLink Trunk
IP Fix
QoS
QoS Wizard
QoS Wizard Config
QoS Wizard Mgmt
Interface Shaper
Interface Apps
Submit
Back
Internet

Application > QoS > Interface Applications

Interface Applications Edit

☐ ARP Spoofing	Default Gateway	
☐ DHCP Snooping	Interface Type	Access
☐ DHCP Spoofing	DHCP Server	
☐ DoS SQLSlam		
☐ DoS Nachia		
☐ DoS Xmas		
☐ DoS TCP SynFinScan		
☐ DoS TCP FtpPort		
☐ DoS TCP DnsPort		
☑ BPDU Blocker		

5. VLAN Hopping

VLAN hopping is the process where traffic from one VLAN can be seen by another VLAN. This could occur if the attacker sends traffic with a VLAN ID of the target VLAN. The attacker could try to behave like a switch and attempt to send and receive traffic between other VLANs.

Nortel recommends the following:

On all active ports interfacing end-user devices or nodes:

- Do not use policy based VLAN selection. MAC and IP addresses are easily spoofed.
- Disable auto-negotiation of VLAN trunking, so that the switch ignores spoofed VLAN tags.
- Set the port to be statically bound to a specific VLAN, so that the switch overwrites spoofed VLAN tags.
- Set the ports to filter tagged frames, so that the switch drops traffic with VLAN tags.
- Disable any unnecessary dynamic protocols (STP, VTP, DTP, OSPF, RIP, etc.) to prevent attacks against these protocols. VTP and DTP are particularly vulnerable to VLANs.
- Enable 802.1x to control access, so that an outsider must authenticate to connect to the layer 2 network. This also defends against DHCP attacks.

In general:

- Do not use native or default VLANs. They may have special treatment or at least, they have well known VLAN IDs.
- Disable any unused ports.

6. ARP/MAC Spoofing Attack

MAC spoofing simply involves spoofing a known MAC address of another host to make the target switch forward frames destined for the remote host to be forwarded to the attackers host. By sending frames with the other host's MAC address, the attacker is telling the Layer 2 switch to forward traffic now to the attacker's port. To correct this, the host must send out frames to tell the switch to relearn most of the host MAC address. This type of attack is confined to the switch itself within the MAC/CAM address table

The attacker can perform ARP spoofing so that it can use an IP address of an attacked host and inform the remote systems to send traffic now to the attacker's MAC address. Gratuitous ARPs (gARP) can be used maliciously by an attacker to spoof the IP address of a host on a LAN segment. It can be used to spoof the identity between two hosts or all traffic from a default gateway in a Man-in-the-middle attack.

The following are some ARP Spoofing tools:

- Ettercap – http://ettercap.sourceforge.net/
- Ac-gateway – http://packetstormsecurity/nl/UNIX/security/ae-gateway.tar.gz
- Arpspoof – http://monkey.org/~dugsong/dsniff
- Cain & Abel - http://www.oxid.it/cain.html

Figure 1 – ARP/MAC Spoofing Attack

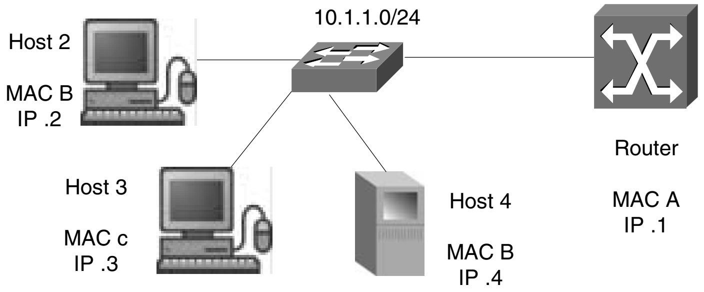

Considering the above drawing, host 4 wishes to perform an ARP spoofing man-in-the-middle (MITM) attack. When hosts 2 or 3 wish to communicate with the router, they will send an ARP request for the router's MAC address. The router (.1) will respond, but as soon as host 4 sends a gARP broadcast claiming it to be the router (.1), hosts 2 and 3 will update their ARP entry for .1 to host 4's MAC address. Also, host 4 can send a gARP to the router using its MAC address for either host 2 or host 3. Now traffic forwarded or received off the 10.1.1.0/24 for either host 2 or host 3 will go to host 4's MAC address. Host 4 could then forward the traffic to the real router, drop the traffic, sniff the traffic, or modify the contents of a packet.

Nortel Solution

MAC Limiting & Discard Unknown MAC Security/BaySecure

The ERS 8600 supports a maximum MAC count on a per port basis in release 3.7.2. For ES, BOSS 3.6 and BOSS 4.2 also have this feature. This feature can be used to limit the number of MAC address on a per port basis.

With Unknown MAC Security, you can enter a MAC or group of MAC addresses on a per port basis. BaySecure provide a similar functionality. This allows controlling what MAC addresses are allowed on a per port basis. With this feature, you have to go to each port and either manually enter a MAC address or stop/start auto-lean and lock the port. Although this feature does limit MAC learning, it may not be the best solution for all applications.

ARP Spoofing

It is possible to prevent ARP/MAC spoofing using off-set filters to block any gratuitous ARPs (gARP). Basically, you have to allow broadcast ARP, block any ARP messages using the source IP or target IP of the default gateway, and then allow ARP reply; these filters should not be applied to the router port(s), only on the user ports. The following is a configuration example using off-set filters on the ES 5510 and ERS 8600 with R-Modules.

ERS55xx:

In release BOSS 4.2 for the ERS5510, you can enable ARP Spoofing Detection via the WEB interface or by using a CLI command as shown below.

CLI

- 5510-48T(config)#***interface fastEthernet <ALL|port|slot/port>***
- 5510-48T(config)#***qos arp spoofing [port <port_list>] enable default-gateway <a.b.c.d>***

where

Parameter	**Description**
port <port_list>	The list of ports to enable QoS ARP spoofing application on.
default-gateway <A.B.C.D>	The IP address of the default gateway to use.

WEB

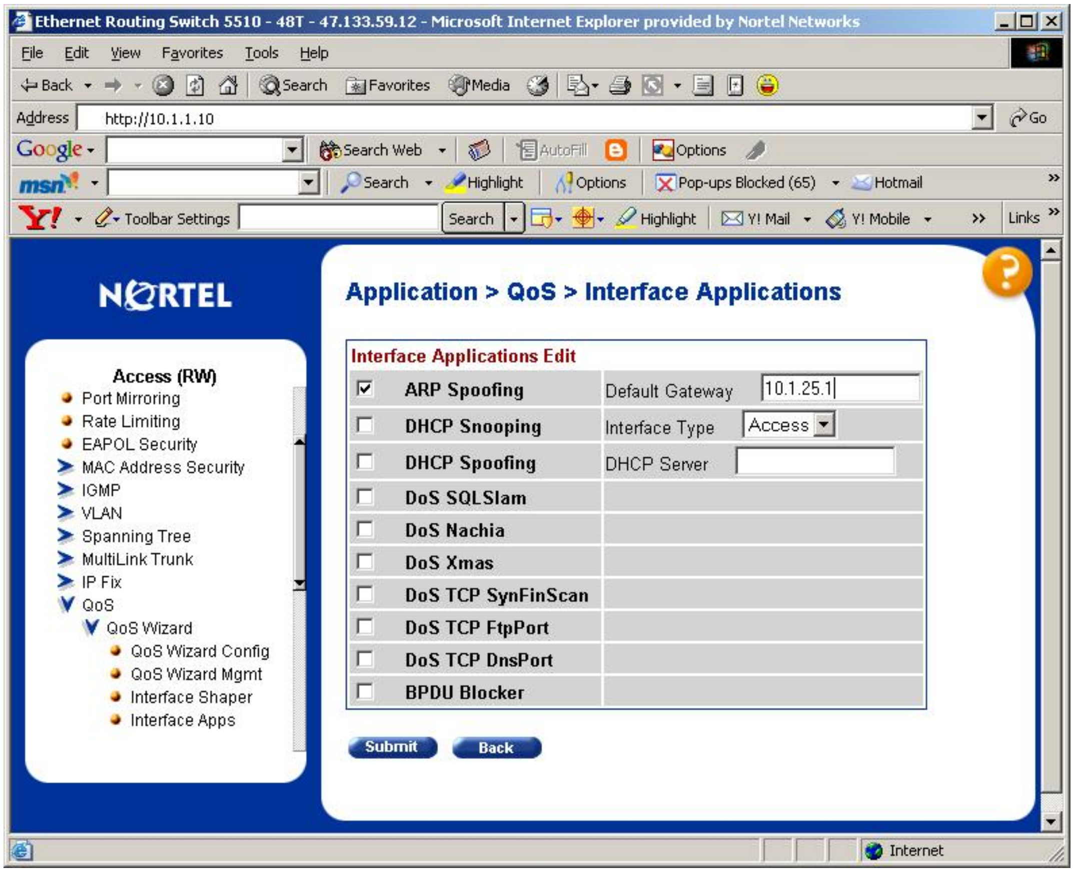

Ethernet Routing Switch 5510 - 48T - 47.133.59.12 - Microsoft Internet Explorer provided by Nortel Networks

Address http://10.1.1.10

Application > QoS > Interface Applications

Interface Applications Edit

☑	ARP Spoofing	Default Gateway	10.1.25.1
□	DHCP Snooping	Interface Type	Access
□	DHCP Spoofing	DHCP Server	
□	DoS SQLSlam		
□	DoS Nachia		
□	DoS Xmas		
□	DoS TCP SynFinScan		
□	DoS TCP FtpPort		
□	DoS TCP DnsPort		
□	BPDU Blocker		

In release BOSS 4.1, you have to manually configure a filter. Shown below is an example on how to setup the filter on the ERS55xx.

Assuming:

- The default gateway is 10.1.25.1
- The user ports are port 26 to 30; we will create an interface group named vlan10 for these ports

NOTE: The items high-lighted in red in the configuration file below is the default gateway address.

```
!
! *** QOS ***
!
qos agent reset-default
qos if-group name vlan10 class unrestricted
qos if-assign port 26-30 name vlan10
qos action 10 name "Null_2" drop-action disable
interface FastEthernet ALL
exit
qos system-element 1  pattern-data
FF:FF:FF:FF:FF:FF:00:00:00:00:00:00:00:00:00:00:08:06:00:01:08:00:06:04:00:01:00:00:00:00:00:00:00:00:00:00:00:00:
00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00
:00:00:00 pattern-mask
FF:FF:FF:FF:FF:FF:00:00:00:00:00:00:00:00:00:00:FF:FF:FF:FF:FF:FF:FF:FF:FF:FF:00:00:00:00:00:00:00:00:00:00:00:0
0:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:
00:00:00:00
qos system-element 2  pattern-data
00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:08:06:00:01:08:00:06:04:00:01:00:00:00:00:00:00:00:00:00:00:00:00:00
:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:0
0:00:00 pattern-mask
00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:FF:FF:FF:FF:FF:FF:FF:FF:FF:FF:00:00:00:00:00:00:00:00:00:00:00:00:
00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00
:00:00:00
qos system-element 3  pattern-data
00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:08:06:00:01:08:00:06:04:00:00:00:00:00:00:00:00:0A:01:19:01:00:00:0
0:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:
00:00:00 pattern-mask
00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:FF:FF:FF:FF:FF:FF:FF:FF:00:00:00:00:00:00:00:00:FF:FF:FF:FF:00:00
:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:0
0:00:00:00
qos system-element 4  pattern-data
00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:08:06:00:01:08:00:06:04:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00
:00:00:00:0A:01:19:01:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:0
0:00:00 pattern-mask
00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:FF:FF:FF:FF:FF:FF:FF:FF:00:00:00:00:00:00:00:00:00:00:00:00:00:00:
00:00:00:00:FF:FF:FF:FF:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:0
0:00:00:00
qos system-element 5  pattern-data
00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:08:06:00:01:08:00:06:04:00:02:00:00:00:00:00:00:00:00:00:00:00:00:00
:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:0
0:00:00 pattern-mask
00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:FF:FF:FF:FF:FF:FF:FF:FF:FF:FF:00:00:00:00:00:00:00:00:00:00:00:00:
00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00
:00:00:00
qos classifier 1 set-id 1 name c1 element-type system element-id 1
qos classifier 2 set-id 2 name c2 element-type system element-id 2
qos classifier 3 set-id 3 name c3 element-type system element-id 3
qos classifier 4 set-id 4 name c4 element-type system element-id 4
qos classifier 5 set-id 5 name c5 element-type system element-id 5
qos policy 1 name P1 if-group vlan10 clfr-type classifier clfr-id 1 in-profile-action 10 non-match-action 9 precedence 10
qos policy 2 name P2 if-group vlan10 clfr-type classifier clfr-id 2 in-profile-action 1 non-match-action 9 precedence 9 track-
statistics individual
qos policy 3 name P3 if-group vlan10 clfr-type classifier clfr-id 3 in-profile-action 1 non-match-action 9 precedence 8 track-
statistics individual
qos policy 4 name P4 if-group vlan10 clfr-type classifier clfr-id 4 in-profile-action 1 non-match-action 9 precedence 7 track-
statistics individual
qos policy 5 name P5 if-group vlan10 clfr-type classifier clfr-id 5 in-profile-action 10 non-match-action 9 precedence 6
```

ES 8600 with R-Modules:

Assuming:

- The default gateway is 10.1.25.1
- The user ports are ports 4/25-4/27

NOTE: The items high-lighted in red in the configuration file below is the default gateway address

```
#
# R-MODULE FILTER CONFIGURATION
#

filter act 1 create
filter act 1 ethernet dstMac
filter act 1 arp operation
filter act 1 pattern p1 add ether-begin 224 32
filter act 1 pattern p2 add ether-begin 304 32
filter act 1 apply
filter acl 1 create inPort act 1
filter acl 1 port add 4/25-4/27
filter acl 1 ace 1 create
filter acl 1 ace 1 action permit
filter acl 1 ace 1 debug count enable
filter acl 1 ace 1 ethernet dst-mac eq ff:ff:ff:ff:ff:ff
filter acl 1 ace 1 arp operation eq arprequest
filter acl 1 ace 1 enable
filter acl 1 ace 2 create
filter acl 1 ace 2 action deny
filter acl 1 ace 2 debug count enable
filter acl 1 ace 2 arp operation eq arprequest
filter acl 1 ace 2 enable
filter acl 1 ace 3 create
filter acl 1 ace 3 action deny
filter acl 1 ace 3 debug count enable
filter acl 1 ace 3 advanced custom-filter1 p1 eq a011901
filter acl 1 ace 3 enable
filter acl 1 ace 4 create
filter acl 1 ace 4 action deny
filter acl 1 ace 4 debug count enable
filter acl 1 ace 4 advanced custom-filter2 p2 eq a011901
filter acl 1 ace 4 enable
filter acl 1 ace 5 create
filter acl 1 ace 5 action permit
filter acl 1 ace 5 debug count enable
filter acl 1 ace 5 arp operation eq arpresponse
filter acl 1 ace 5 enable

#
# R-MODULE FILTER CONFIGURATION
#


filter act 1 create
filter act 1 arp operation
filter act 1 pattern p1 add ether-begin 224 32
filter act 1 apply
filter acl 1 create inPort act 1
filter acl 1 port add 4/25
filter acl 1 ace 1 create
filter acl 1 ace 1 action deny
filter acl 1 ace 1 debug count enable
filter acl 1 ace 1 advanced custom-filter1 p1 eq a011901
filter acl 1 ace 1 enable
```

ERS8300

In regards to the ERS 8300, it will support off-set filters in the ?? release.

7. Private VLAN Attacks

Private VLANs are used to limit communication between ports within a VLAN. This feature disallows all communications between user ports on the same VLAN. It only allows bi-direction traffic between a user port and the trunk port.

A network attack capable of bypassing the security of private VLANs involves the use of a proxy attack. A proxy attack simply uses ARP, such as a static ARP entry where the victims host is reachable by the router's MAC address. In figure 2, the attacker sends a packet with the source IP and MAC of his or her machine, a destination IP address of the target system and with a destination MAC address of the router. When the frame arrives at the router, the router will notice the packet is destined for the victim and will rebuild the frame with the correct MAC address and send it to the victim as shown in figure 2 (b). This network attack works only for unidirectional traffic as the Private VLAN configuration will block any traffic from the victim to the attacker. If both hosts are compromised, via static ARP entries, then bidirectional traffic could occur.

Figure 2 – Private VLAN Proxy Attack

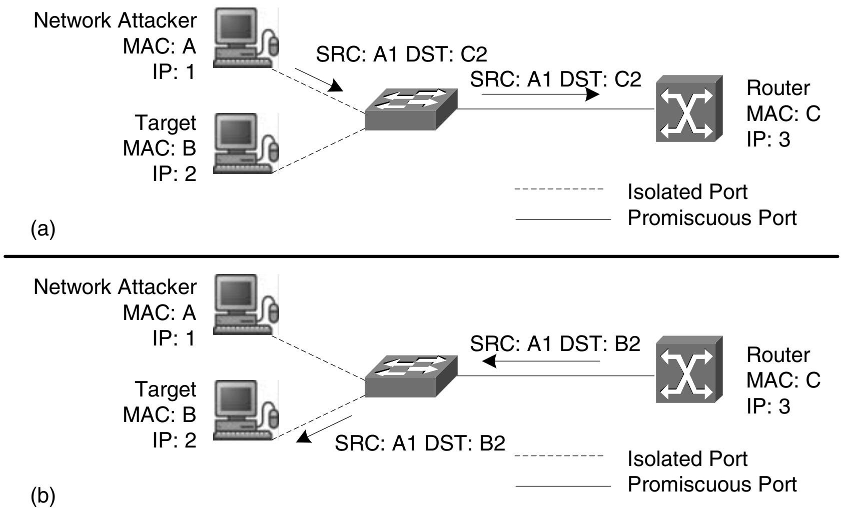

This type of attack is not a private VLAN vulnerability as the private VLAN rules are enforced. The problem is simply the way the router works. Stopping this type of attack is very easy. All that is required is an inbound ACL on the router to stop all traffic on the local subnet.

Nortel Solution

The Nortel Private VLAN Edge feature, available on all ES and ERS55xx switches, provides isolation between hosts. This feature provides separation between access ports by assigning one or more ports to a different VLAN with one common VLAN uplink port. For example, let's assume we wish host isolation between access ports 1, 5, and 10 with port 25 being the uplink port. The configuration would look something like the following as shown in the chart below. It is irrelevant what VLANs are used as port members are untagged.

VLAN	VLAN Member	PVID Assignment	
		Port	PVID
10	1 and 25	1	10
11	5 and 25	5	11
12	10 and 25	10	12
50	1, 5, 10, and 25	25	50

In regards to the router configuration, the following is a filter example for the ERS 8600 assuming the local subnet is 172.1.2.0/24 on Ethernet interface 3/19.

```
#
# TRAFFIC-FILTER CONFIGURATION
#
ip traffic-filter create destination dst-ip 172.1.2.0/255.255.255.0 src-ip 0.0.0.0/0.0.0.0 id 4
ip traffic-filter filter 4 action mode drop
ip traffic-filter filter 4 action statistic enable
ip traffic-filter filter 4 name "Private_VLAN"
ip traffic-filter global-set 1 create
ip traffic-filter set 300 create name "Private"
ip traffic-filter set 300 add-filter 4

#
# PORT CONFIGURATION - PHASE II
#
ethernet 3/19 ip create 172.1.2.1/255.255.255.0 2210 mac_offset 1
ethernet 3/19 ip traffic-filter create
ethernet 3/19 ip traffic-filter add set 300
ethernet 3/19 ip traffic-filter default-action forward
```

8. IEEE 802.1x

In regards to EAP, a couple of flaws with the 802.1x standard in a Wireless LAN are pointed out by Arunesh Mishra and William Arbaugh of the University of Maryland's Computer Science Department - http://www.cs.umd.edu/~waa/1x.pdf. In this paper, the two security attacks are:

- session hijacking
- establishment of a man-in-the-middle (MITM) attack.

In a MITM attack, an attacker takes advantage of the fact that 802.1x is an asymmetric protocol. This allows the network to authenticate the user, but does not allow the user to authenticate the network. Here, an attacker can act as an access point to the supplicant and as a client to the real network access point. The attacker sends the supplicant a forged EAP-Success message that appears to have come from the authenticator. If successful, data will pass between the victims station to the legitimate access point via the attacker.

In the second type of attack, an attacker attempts to hijack an existing session. Here the attacker poses as an access point to the victim's supplicant and poses as a supplicant to the network access point. First, it sends a fake *dissociate* message forged with the authenticator's MAC address telling the victim's supplicant to drop its connection. Then the attacker hijacks that connection using the victims MAC address to fool the access point into exchanging data with it. The authenticator state is still in an authenticated and associated state.

Solution:

Although 802.1X was conceived of as asymmetric, there are a number of protocols that were developed to address these and other concerns about 802.1x. EAP-TLS, EAP-TTLS, and EAP-PEAP are examples of such protocols. These protocols provide for strong mutual authentication between mobile station and access point. These protocols use master keys that are bound to the authentication and are used during the session. This prevents a MITM attack as the attacker was not authenticated to either the supplicant or authenticator as it has no master key and can't get one. Hence, it could not decipher data even if it was able to pass data between it and the victim's supplicant and network access point.

Nortel Solution

The Nortel WLAN 2300 Series supports traditional 802.1x pass-through or normal operation in addition to 802.1x acceleration.

When configured for 802.1x acceleration, the WLAN 2300 Series offloads RADIUS authentication from the RADIUS server. This feature offloads the AAA servers and accelerates up to 90% of EAP processing resulting in faster authentication and roaming. Without this feature, the exiting AAA servers are burdened with all EAP processing resulting in slower authentication and increased roaming latency. EAP-TLS, PEAP, and EAP-MD5 are supported.

Overall, the following is a summary of the overall 802.1x offerings available from Nortel.

2300 products

802.1x passthrough (normal 802.1x "authenticator" behavior):
EAP-TLS
EAP-TTLS
PEAP-MS-CHAP-V2

802.1x Acceleration/EAP Offload:
PEAP-MS-CHAP-V2
EAP-TLS
EAP-MD5

2270/2230 products

802.1x
PEAP
EAP-TLS
EAP-TTLS
LEAP (Steelbelt and Oddysey client have been tested successfully)

2201 (NIC)

The following is what the AEGIS client that ships with the NIC supports:
LEAP
EAP-TLS
EAP-TTLS
PEAP-MS-CHAP-V2

2202 (NIC)

The following is what the AEGIS client that ships with the NIC supports:
EAP-TLS
EAP-TTLS
PEAP-MS-CHAP-V2
PEAP (has not been tested, but should work)

2210/11/12 (handsets)

Cisco FSR (which uses LEAP for authentication)

9. Appendix

9.1 ARP

ARP Request/Reply Frame

6 bytes	6 bytes	2 bytes	28 bytes for IP	4 bytes
Ethernet destination address	Ethernet source address	Frame type	ARP Request / Reply	CRC

- Ethernet Destination Address
  - ff:ff:ff:ff:ff:ff (broadcast) for ARP

- Ethernet Source Address
  - MAC address of ARP requester

- Frame Type
  - ARP request/reply: 0x0806
  - RARP request/reply: 0x8035
  - IP datagram: 0x0800

ARP Request/Reply Format

0 7	8 15	16 31	
Hardware Type		Protocol Type	28 bytes
Hardware len	Protocol len	ARP operation	
Sender MAC address (bytes 0-3)			
Sender MAC address (bytes 4-5)		Sender IP address (bytes 0-1)	
Sender IP address (bytes 2-3)		Destination MAC address (bytes 0-1)	
Destination MAC address (bytes 2-5)			
Destination IP address (bytes 0-3)			

- Hardware type: 1 for Ethernet
- Protocol type: 0x0800 for IP
- Hardware length: 6 bytes for Ethernet
- Protocol length: 4 bytes for IP
- ARP operation: 1 = request, 2 = reply, 3/4 = RARP req/reply

ARP request vs. Gratuitous ARP Format

Assuming the host MAC address is 00:00:00:3d:01:0a and the host IP is 172.1.20.11, the ARP request and gratuitous ARP format will look like the following

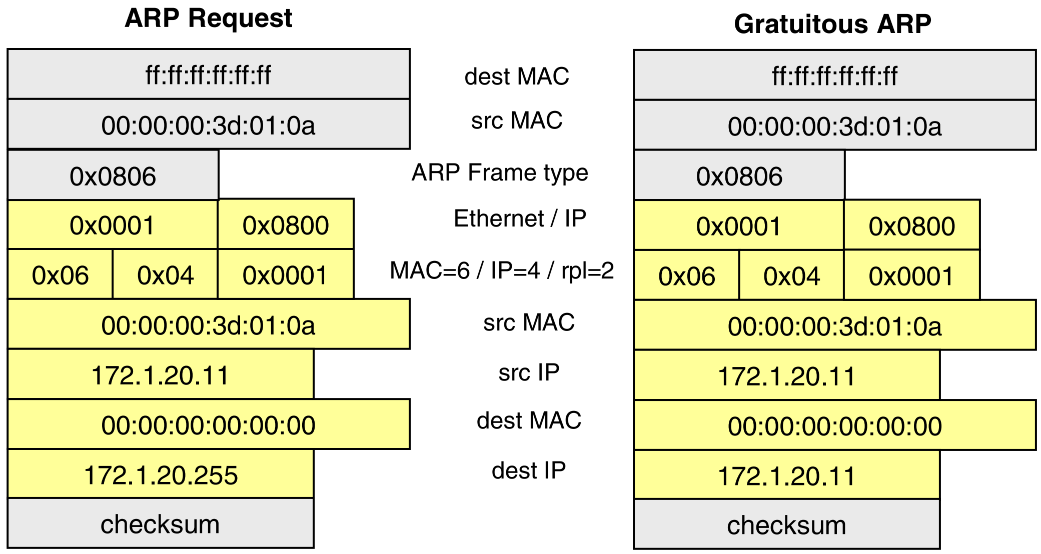

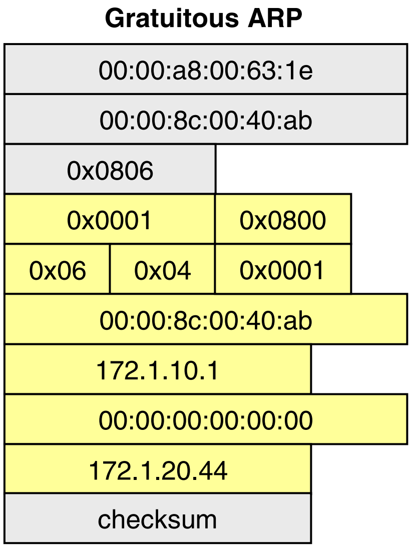

When a host sends an ARP request to resolve its own IP address, it is called gratuitous ARP. In the ARP request packet, the source IP address and destination IP address are filled with the same source IP address itself. The destination MAC address is the Ethernet broadcast address (FF:FF:FF:FF:FF:FF).

9.2 DHCP Message Format

```
0                   1                   2                   3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     op (1)    |   htype (1)   |   hlen (1)    |   hops (1)    |
+---------------+---------------+---------------+---------------+
|                            xid (4)                            |
+-------------------------------+-------------------------------+
|           secs (2)            |           flags (2)           |
+-------------------------------+-------------------------------+
|                          ciaddr  (4)                          |
+---------------------------------------------------------------+
|                          yiaddr  (4)                          |
+---------------------------------------------------------------+
|                          siaddr  (4)                          |
+---------------------------------------------------------------+
|                          giaddr  (4)                          |
+---------------------------------------------------------------+
|                                                               |
|                          chaddr  (16)                         |
|                                                               |
|                                                               |
+---------------------------------------------------------------+
|                                                               |
|                          sname   (64)                         |
+---------------------------------------------------------------+
|                                                               |
|                          file    (128)                        |
+---------------------------------------------------------------+
|                                                               |
|                          options (variable)                   |
+---------------------------------------------------------------+
```

Field	Octets	Description
Op	1	***Operation Code***. Specifies the message type. A value of 1 indicates a request while a value of 2 is a reply message. A DHCP client sending a request to a server uses an OP code of 1 while the server reply's with a code of 2.
HType	1	***Hardware Type***. This field specifies the type of hardware used for the local interface. Some examples are a value of 1 for Ethernet (10Mb), 6 for IEEE 802 Networks, 15 for Frame Relay, 20 for Serial Line. Please see RFC 2131 for more details.
HLen	1	***Hardware Address Length***. Specifies how long the hardware addresses are in this message. For Ethernet, the value is 6.
Hops	1	***Hops***. Set to zero by a client prior to transmitting a request and used by relay agents to control the forwarding of BOOTP and/or DHCP messages.
XID	4	***Transaction Identifier***. A 32-bit identification field generated by the client to associate or match messages and responses between a client and a server.
Secs	2	***Seconds***. For DHCP, the seconds elapsed since the client began an attempt to acquire or renew a lease. This may be used by the server to prioritize replies when multiple clients requests are outstanding
Flags	2	***Flags***. This corresponds to the formerly empty field for BOOTP

Field	Octets	Description
		message format defined in RFC 951. It is redefined in RFX 1542 as the Flags field. Only one flag subfield is used, subfield B, size 1/8 or 1 bit that is set when a client does not know its own IP address at the time it sends a request. It indicates to the DHCP server or relay agent that receives the request that it should send its reply back by broadcast
CIAddr	4	***Client IP Address***. This field is filled in only if the client has a valid IP address while it is in the BOUND, RENEWING, or REBINDING states. Otherwise, this field is set to zero. The client can only use this field when its address is valid and useable, not during the process of acquiring a particular IP address in a lease.
YIAddr	4	Your (client) IP address. The IP address that the server assigns to the client
SIAddr	4	***Server IP Address***. The IP address of the server that the client of the next server to use for the next step in the bootstrap process. This address may or may not be the server sending this reply.
GIAddr	4	***Gateway IP Address***. This field is the relay agent IP address used to route BOOTP messages with BOOTP relay agents are used to facilitate the communications of BOOTP requests and replies between a client and a server of different subnets or networks
CHAddr	16	***Client Hardware Address***. The client hardware (layer 2) address used for identification and communication.
SName	64	***Server Name***. Optional server host name of the server sending a DHCPOFFER or DHCPACK. This is a simple text name or fully qualified DNS domain name.
File	128	***Boot Filename***. Optional boot file name used by a client to request a particular type of boot file in a DHCPDISCOVERY message. Used by the server in a DHCPOFFER to fully specify a boot file directory path and filename.
Option	Variable	***Options***. This field can be used by both client and server that includes several parameters required for basic DHCP operations.

Contact Us:

For product support and sales information, visit the Nortel Networks website at:

http://www.nortel.com

In North America, dial toll-free 1-800-4Nortel, outside North America dial 987-288-3700.