How To | Configure the AT-8948 and AT-9900 series switches for maximum security against attack

Introduction

Increasingly we see the deployment of switched networks in the Enterprise and the use of switches in other areas of network infrastructure, such as Service Providers.

Security on these switches is equally as important as that of servers and end user computer equipment. A breach in security on a core switch can bring large networks to a complete standstill.

When installing networking equipment into an environment that exposes it to unauthorised access attempts and malicious attacks, it is important that it is configured in a way that blocks attacks.

What this document covers

This two-part document considers several aspects of secure configurations for the AT-8948 and AT-9900 series switches. Part 1 describes how to:

- securely configure management services that must be available, and how to disable ones that are not required
- use hardware filters to block undesirable traffic
- securely configure Layer 3 protocols
- configure the switch against Layer 2 attacks

This is followed by a look at other features of the switch that can enhance the security and stability of the network, for example, 802.1x port authentication and packet storm protection (see page 10).

Part 2 describes specific tests to verify the switch's invulnerability to attacks at Layer 2 (see page 12).

The information provided here applies to:

- Products: AT-8948 and AT-9900 series switches
- Software versions: 2.7.2 and above

Part 1:
Secure configuration of management services

The AT-8948 and AT-9900 series switches can be remotely managed by the following methods:

- **SNMP**
- **HTTP**
- **Telnet**
- **SSH**

Before looking at these topics individually, it is important to understand how to set up the switch with multiple access levels.

Setting up multiple levels of access privilege

There are three levels of privilege on the switch: user, manager, and security officer.

User-level privilege provides very minimal access to the switch. It allows the user to execute only one command (namely to Telnet to a set of predefined addresses).

By default, the manager-level privilege and security officer-level privilege both provide full access to all commands on the switch.

To restrict the access that a manager-level user has, it is necessary to enable system security.

Enabling system security

Before system security can be enabled, at least one securityofficer-level user needs to be defined on the switch. The commands to enable system security are:

```
Add user=<name> password=<password> priv=securityofficer
Enable system security
```

Once this has been done, security officer-level users still have full access to all commands, but manager-level users are severely restricted in the commands they can execute.

Additionally, if system security is not enabled on the switch, the encryption keys created for SSL or SSH (see page 4) are deleted upon reboot. This prevents unauthorized use of the keys.

SNMP

The SNMPv3 protocol provides the opportunity to configure SNMP in a much more secure way than has been possible with previous versions of SNMP.

In particular, with SNMPv3, it is possible to:

- Set up restricted views, i.e. limited sets of MIB variables that can be accessed by particular users (who need to enter a password to get access to their view)
- Encrypt the SNMP messages being sent across the network

A typical secure SNMPv3 configuration on the AT-8948 and AT-9900 series switch would be:

```
    enable snmp
#Enable SNMP authentication failure traps.
    enable snmp authenticate_trap
#Adds SNMP target parameters to specify a security profile for a trap
   target addresses.
    add snmp targetparams=netmonpc securitylevel=authpriv user=steve
#Add a target address where traps will be sent.
    add snmp targetaddress=nms ip=192.168.11.23 udp=162 params=netmonpc
#Create an SNMP view which will allow access to some specified  OIDs
   (Object Identifiers) onwards, restricts access to other OIDs:
    add snmp view=view1 oid=1.3.6.1 type=include
    add snmp view=view1 oid=... type=include
    ...
    ...
    add snmp view=view1 oid=1.3.6.1.6 type=exclude
    add snmp view=view1 oid=... type=exclude\
    ...
    ...
#Create an SNMP group with full read/write and notify access to view1,
   and specify authentication and privacy.
    add snmp group=group1 securitylevel=authPriv readview= view1
      writeview= view1 notifyview= view1
#Create an SNMP user and associate it with group1, and specify the
   authentication and privacy protocol and the authentication and
   privacy passwords.
    add snmp user=steve group=group1 authprotocol=MD5
      authpassword=cottonsox privprotocol=DES privpassword=woollytop
```

Note: *SNMP is not enabled on the switch by default, so if you do not wish to use SNMP, there is no need to enter any command to block SNMP access to the switch.*

HTTP

The switch can be managed via encrypted HTTP, using the SSL protocol. To enable this service, configure as follows:

```
    create enco key=0 type=rsa length=1024
    set system distinguishedname="cn=switch1,o=my_company,c=us"
#Create a self-signed certificate
    create pki certificate=cer_name keypair=0 serialnumber=12345
    add pki certificate=cer_name location=cer_name.cer trust=yes
#Set the HTTP server into secure mode, listening on TCP port 443
    set http server security=on sslkey=0 port=443
    disable HTTP server
```

Note: *The HTTP server is enabled by default on the switch, so if you do not wish to manage the switch by HTTP, then it is recommended to block HTTP access to the switch, using the command:*

Telnet

Telnet access to the switch is enabled by default. It is recommended to disable Telnet, and use SSH instead, for command line remote access. To disable Telnet, use the command:

```
    disable telnet server
```

SSH

Using SSH, it is possible to have encrypted access to the switch's command line interface.

To configure the switch for SSH access, use these commands:

```
    create enco key=0 type=rsa length=1024 description="host key"
      form=ssh
    create enco key=1 type=rsa length=768 description="server key"
      form=ssh
    enable ssh server hostkey=0 serverkey=1 expirytime=1 logintimeout=60
```

To then create a definition for a user account that is authenticated by RSA:

1. Upload the user's public RSA key onto the switch. Note that the switch requires that the file with the public key have the extension **.key**). Use one of the following:

   To load by TFTP, use:

   ```
   load server=ipadd fi=key-file-name
   ```

   To load by Zmodem, use:

   ```
   load asyn=0 meth=zmo
   ```

2. Create an ENCO key from this file on the switch.

   ```
   create enco key=7 type=rsa file=key-file-name desc="user's public
     key" form=ssh
   ```

3. Create the user account.

   ```
   add ssh user=name keyid=7
   ```

Using HW filters to protect the CPU from undesirable traffic

Malicious attacks sometimes try to overload the CPU with traffic destined for the switch. Getting the CPU to run at maximum 100% causes problems processing network control traffic, which is critical to keep a network functioning well.

Using hardware filters to block traffic destined for the switch's own IP address can protect the CPU. The following example shows how to configure classifiers to match on traffic destined for the switch's IP address. Hardware filters are then used to allow SNMP and Web traffic, for monitoring and Web (GUI) control of the switch, but block all other traffic to the CPU.

The first step in using hardware filters to protect the CPU is to create classifiers that match on traffic destined for the switch's own IP address. This example assumes that the IP address on the switch is 192.168.1.254. We must create the following three classifiers:

A classifier to match SNMP traffic, which uses UDP destination port 161.

```
create classifier=1 ipdaddress=192.168.1.254 udpdport=161
```

A classifier to match Web (GUI access) traffic, which uses TCP destination port 80.

```
create classifier=2 ipdaddress=192.168.1.254 tcpdport=80
```

A classifier to match on all other traffic destined for the switch's IP address.

```
create classifier=3 ipdaddress=192.168.1.254
```

Next we must create hardware filters to allow SNMP and Web traffic, but block other traffic.

```
add switch hwfilter=1 classifier=1,2 action=forward
add switch hwfilter=2 classifier=3 action=discard
```

Note: *Packet filters with a low filter-id have precedence over packet filters with ones. For example, a packet is matched against entries in hwfilter 1 before being matched against entries in hwfilter 2. So the SNMP and Web traffic is allowed because it matches the first filter. All other traffic destined for the switch's own IP address will match against hwfilter 2 and be discarded.*

Securely configuring Layer 3 protocols

One form of attack against Layer 3 switches is to send deliberately misleading routing protocol packets that put invalid routes in the switch's routing table. If the switch uses VRRP to choose among redundant gateway devices, the attacker can also spoof VRRP packets, and make the correct gateway device go into the backup state, which stops packets from being forwarded out of the LAN. To protect against these sorts of attacks, configure authentication on routing protocols and VRRP.

OSPF

To configure OSPF with MD5 authentication, use the commands:

```
#First configure an OSPF area with MD5 authentication.
    add ospf area={backbone|area-number} authentication=md5
#Then configure individual interfaces within the area to use MD5
   authentication.
    add ospf interface=interface authentication=md5
#Add an MD5 key that can be used for interface authentication.
    add ospf md5key=key id=1...255 interface=interface
```

RIP

To configure an IP interface to send/receive RIP using MD5 authentication, the command is:

```
    add ip rip int=<interface> auth=MD5 password=<password>
```

BGP

To configure BGP to add an MD5 digest to every BGP packet sent over the TCP connection to a particular peer, configure the peer as follows:

```
    set bgp peer=<address> authentication=MD5 password=<password>
```

VRRP

To enable authentication on a VRRP instance, configure as follows:

```
    set vrrp=<ID> authentication=plaintext password=<password>
```

How to secure the switch against Layer 2 attacks

This section describes how to configure the switch to block the following types of attack:

- **MAC flooding attack**
- **ARP attack**
- **Private VLAN attack**
- **Random frame stress attack**
- **STP attack**
- **Multicast brute force attack**

MAC flooding attack

Attack description

The MAC flooding attack is an attempt to exploit the way switches work. A network switch has a hardware-learning table, which stores source addresses of packets received. When the table is full, the switch cannot learn new packet addresses until some entries in the table expire; so packets are flooded within the ingress VLAN.

A MAC flooding attack looks like traffic from thousands of computers arriving at a single port. However, traffic is actually coming from one host spoofing the MAC address of thousands of bogus hosts. The aim is to fill the hardware table and make the switch flood subsequent traffic on the VLAN. This allows attackers to 'sniff' traffic that would not usually be available to them.

Prevention

The simplest method to guard against a MAC flooding attack is to configure port security on the switch. Port security lets you specify the number of MAC addresses that can be learned on a given port to stop a malicious user from swamping the switch with bogus host addresses.

If the configured number of MAC addresses is exceeded, the port can be configured to do the following:

- discard the packet and take no further action
- discard the packet and notify management with an SNMP trap
- disable the port

Configuration commands

```
set switch port=<number> learn=<max address to learn>
  intrusionaction=discard
```

ARP attack

Attack description

ARP is responsible for managing the relationship between MAC addresses and IP addresses for network devices. ARP information can be spoofed, or faked, to facilitate the control of all network data.

ARP poisoning occurs when an intruder changes the packet header of a network ARP and poses as a node with a valid MAC address. All subsequent communications intended for that IP address are delivered to the imposter's MAC address, intercepting traffic flowing between stations.

Prevention

Blocking communication between the imposter and attacked device at Layer 2 can prevent this type of attack. Using the private VLAN feature of the AT-9924 stops communication between devices in the same VLAN.

Configuration commands

```
create vlan=vlan2 vid=2 private
add vlan=2 port=49 uplink
add vlan=2 port=1-24
```

Private VLAN attack

Attack description

A private VLAN contains switch ports that cannot communicate with each other, but can access another network. These ports are called private ports. Each private VLAN contains one or more private ports, which cannot communicate at Layer 2, and a single uplink port for access to other networks.

An attack on a private VLAN occurs when traffic is sent from one host to another at Layer 3, via the gateway device, to get around the imposed Layer 2 private VLAN security. The attacker sends a packet with a rogue destination MAC address (the MAC of the gateway device), but with the destination IP of the victim, which resides in the same private VLAN. The gateway router forwards the packet to the victim; so the intended private VLAN security is bypassed.

Prevention

To prevent against private VLAN attacks, ACL type functionality can be implemented by using hardware filters on the AT-9924 switch to stop traffic being relayed via the gateway router.

Configuration commands

```
create vlan="vlan2" vid=2 private
add vlan="2" port=1 uplink
add vlan="2" port=2-8
#Assuming that 192.168.1.0/25 is the subnet being used on the hosts
   connected to the private VLAN
create class=1 ipsa=192.168.1.0/25 ipda=192.168.1.0/25
add switch hwf class=1 action=discard
```

Random frame stress attack

Attack description

The random frame stress attack is a brute force attack, which randomly varies several fields in a packet while keeping destination constant so traffic gets to the receiver. The aim is to cause VLAN leakage (packets seen on a different VLAN than where they should be) on the switch.

Prevention

The switch should not allow any frames to be seen on ports that are not in the receiving VLAN.

Configuration commands

No configuration required - the switch is inherently robust against this attack.

STP attack

Attack description

Spanning Tree Protocol (STP) is used in switched networks to stop loops from forming between switches and therefore prevents broadcast storms. An STP attack occurs when the attacker sends out STP BPDUs announcing that there is a new root bridge and therefore an STP reconvergence takes place. While this is happening VLAN leakage might occur.

Prevention

Frames should be contained within the VLAN that they were sent into and not 'leaked' to any other VLAN. This should still be the case even when STP is re-converging.

Configuration commands

No special configuration required. The switch is inherently robust against this attack.

Multicast brute force attack

Attack description

The multicast brute force attack involves a storm of Layer 2 multicast packets being sent into the switch to see if these are constrained within the ingress VLAN.

Prevention

The Layer 2 multicast packets should be constrained within the ingress VLAN. No packets should be 'leaked' to other VLANs.

Configuration commands

No special configuration required. The switch is inherently robust against this attack.

Other features that enhance the security and stability of the network

802.1x port authentication

802.1X port authentication is a good way to secure your network from unauthorized access. This functionality allows a network controller to restrict external devices from gaining access to the network behind an 802.1x-controlled port.

External devices that wish to access services via a port under 802.1x control must firstly authenticate themselves, and gain authorisation from an external authentication server, before any packets are allowed to pass through the 802.1x controlled port to the network.

There are several options for setting up 802.1x on your switch, which include single and multiple supplicants (the devices wishing access) per port. You can also use the switch as a supplicant and authenticator, using 802.1x to control access between two LANs.

The basic 802.1x setup is configured as follows:

```
#Enable 802.1x

enable portauth

#Enable 802.1x on the port you want to configure as an
  authenticator.

enable portauth port=port-name type=authenticator [other-
  options...]

#Add a RADIUS server to process authentication requests from the
  supplicant sent via the port #configured as an authenticator.

add radius server=ipadd secret=secret port=port-number
```

Packet storm protection

Broadcast storms

Excessive broadcast traffic can be crippling to a switch as it goes to the CPU for processing and is software switched, causing performance degradation.

Ethernet frames (Layer 2 packets) have no TTL field. So when a broadcast storm occurs for some reason in a network, these packets can 'loop' around the network indefinitely, until a user intervenes. This severely limits the bandwidth available to the applications running on the network.

Multicast storms

Some viruses use large numbers of multicast frames to bring networks to a standstill. By implementing multicast limiting on the switch, this can be avoided.

To help control and prevent broadcast and multicast storms, the AT-9924 allows limits to be set on a per switch and per port basis.

Configuring storm protection

You can set port limits for broadcast and multicast packets. Broadcast limits are set for an individual port, and then multicast limiting is either enabled or disabled for the switch. If it is enabled, the broadcast limit set per port also applies to multicast packets.

So multicast rate limiting is either on for all ports that have broadcast limiting enabled, or off for all ports. You cannot set up multicast limiting only.

To set broadcast storm control and the reception rate limit for a port, use the command:

```
set switch port={port-list|all} [bclimit={none|limit]
```

This limits the maximum data rate of reception of Layer 2 broadcast packets by the specified port. Packets beyond this limit are discarded.

To enable and disable the multicast storm control, which works with broadcast limiting, use the command:

```
enable switch mclimiting
```

Destination lookup failure protection

Destination lookup failure packets have a Layer 2 destination address that the switch has not learned. The switch does not know where to forward the packets, so the packets are broadcast to all ports on the switch, and to the CPU. You can limit the rate at which destination lookup failure packets are received.

To set destination lookup failure rate limiting on the switch, use the command:

```
set switch dlflimit={none|limit}
```

This specifies the maximum data rate at which destination lookup failure packets are received by the whole switch. Packets beyond this limit are discarded.

Other switch security settings

Depending on what protocols are used, and the network set up of your switch, there are other security settings worth considering. This section contains examples.

Private VLANs: If you want Layer 2 separation of hosts, private VLAN functionality can be implemented. It is a good idea to include a filter to guard against a private VLAN attack as discussed in the testing section earlier.

STP: If you are running STP, consider disabling it on ports not connected to other network devices also running STP. For RSTP, we recommend configuring edge ports for workstations and non-RSTP clients.

Filtering: Filters for Layers 2 and 3 can be employed in hardware to control network access with no performance impact because they operate at wire speed.

Viruses commonly infect a host and then propagate themselves to as many attached devices as possible. If a network administrator notices large amounts of traffic suddenly going through the switch using a specific TCP port number, this may indicate virus activity. Implementing a hardware filter on the switches in the network so that they drop all traffic to that TCP port number could stop the propagation of the virus.

Part 2:
Tests to verify the switch's invulnerability to Layer 2-based attacks

This section describes tests that have verified the switch is robust against some common VLAN attacks (Layer 2). The switch was tested against the following types of attack:

Test 1: MAC flooding attack

Network setup

The AT-9924 switch was setup in the following network:

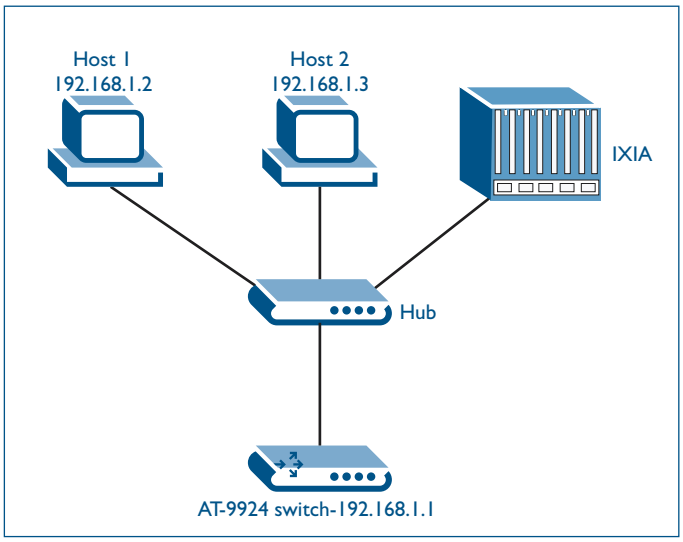

Switch configuration

```
set system name="9924_Switch"
set switch port=1 learn=1 intrusionaction=discard
enable ip
add ip int=vlan1 ip=192.168.1.1
```

Attack results

The Ixia was set to send a large number of packets with random MAC addresses, with the learn limit on the port set to 10. The command to set the learn limit is:

```
set switch port=1 learn=10 intrusionaction=discard
```

Once 10 MAC addresses were learned on the port, the port was locked (the state of the port can be seen with the command 'sh switch port=x') and the switch could learn no further MAC addresses. This stops the MAC address table being filled and subsequent traffic flooded on the VLAN. A second test was done with the learn limit set to 1, where Host 1 would try and ping the AT-9924 and then Host 2 would try and ping the AT-9924.

The first host was able to successfully ping the AT-9924, and its MAC address was added to the MAC table. At this point the port was locked from learning any more MAC addresses as the limit was set to one. This can be seen in the output of 'show switch port=1" below:

```
Manager Mustang> sh swi port=1

 Switch Port Information
--------------------------------------------------------------------------------
 Port ........................... 1
   Description ................... -
   Status ........................ ENABLED
   Link State .................... Up
   UpTime ........................ 00:32:04
   Port Media Type ............... ISO8802-3 CSMACD
   Configured speed/duplex ....... Autonegotiate
   Actual speed/duplex ........... 100 Mbps, half duplex
   MDI Configuration (Polarity) .. Automatic (MDI)
   Loopback ...................... Off
   Configured master/slave mode .. Not applicable
   Actual master/slave mode ...... Not applicable
   Acceptable Frames Type ........ Admit All Frames
   Disabled Egress Queues ........ -
   BCast & MCast rate limit ...... -
   BCSC rate Limiting ............ disabled
   Egress rate limit ............. -
   Learn limit ................... 1
   Intrusion action .............. Discard
   Current learned, lock state ... 1, locked
   Relearn ....................... OFF
   Mirroring ..................... Disabled
   Is this port mirror port ...... No
   Enabled flow control(s) ....... -
   Advanced Flow Control length .. -
   Ingress Filtering ............. Off
   Trunk Group ................... -
   STP ........................... default
   Cable Length .................. -
```

The second host was unable to successfully ping the AT-9924 as the port was now locked from accepting any packets from other source MAC addresses.

MAC flooding attacks can be successfully negated using port security as demonstrated to stop a malicious user from swamping the switch with bogus host addresses.

Test 2: ARP attack

In this test Host 2 is a malicious user who wishes to intercept traffic destined for the gateway router. By changing his IP address to the IP of the gateway, and sending an ARP broadcast, the ARP table of Host 1 becomes corrupted and traffic is sent to Host 2 instead of the 'real' gateway.

Network setup

The AT-9924 switch was setup in the following network:

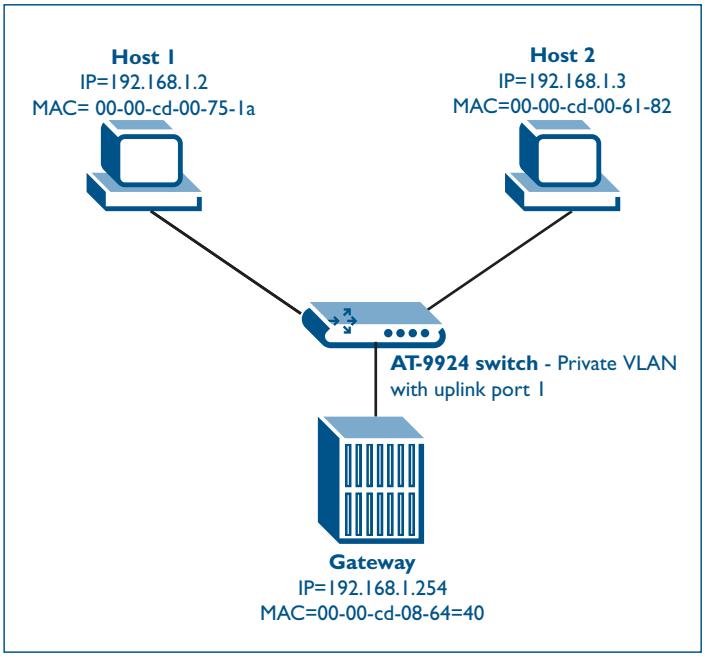

Switch configuration

```
set system name="9924_Switch"
create vlan=vlan2 vid=2 private
add vlan=2 port=1 uplink
add vlan=2 port=2-24
```

Attack results

The following test was carried out with no configuration on the AT-9924 switch.

Both Host 1 and Host 2 send traffic to the gateway router. The ARP table on Host 1 looked as follows:

```
Manager Host_1> sh ip arp

  Interface        IP Address       Physical Address      ARP Type    Status
  ------------------------------------------------------------------------
  eth0             192.168.1.2      00-00-cd-00-61-82     Dynamic     Active
  eth0             192.168.1.254    00-00-cd-08-64-40     Dynamic     Active
  eth0             192.168.1.255    ff-ff-ff-ff-ff-ff     Other       Active
  eth0             255.255.255.255  ff-ff-ff-ff-ff-ff     Other       Active
```

The correct MAC address for Host 2 (192.168.1.2) of xx-xx-61-82 is seen and the correct MAC address for the gateway (192.168.1.254) is xx-xx-64-40.

Then the IP address on Host 2 was changed to the IP of the gateway (192.168.1.254). Host 2 then sent an ARP broadcast to Host 1. The ARP table on Host 1 now looked like this:

```
Manager Host_1> sh ip arp

  Interface        IP Address       Physical Address      ARP Type    Status
--------------------------------------------------------------------------------
  eth0             192.168.1.2      00-00-cd-00-61-82     Dynamic     Active
  eth0             192.168.1.254    00-00-cd-00-61-82     Dynamic     Active
  eth0             192.168.1.255    ff-ff-ff-ff-ff-ff     Other       Active
  eth0             255.255.255.255  ff-ff-ff-ff-ff-ff     Other       Active
```

At this point the gateway IP of 192.168.1.254 was associated with the MAC address of Host 2 instead of the gateway router. Traffic from Host 1 to the 'gateway' went to the malicious user on Host 2.

ARP attack protection

Implementing a private VLAN on the AT-9924 switch successfully stopped this attack. A private VLAN stops communication between hosts within the VLAN. They can only communicate with the configured uplink port. In this case the uplink port is port 1 connected to the gateway.

By using the above configuration, the malicious user on Host 2 was unable to send the ARP broadcast to Host 1 and corrupt his ARP table. All traffic from Host 1 to the gateway continued to reach the 'real' gateway.

Test 3: Private VLAN attack

In this test a AT-9924 switch was configured with a private VLAN and connected via the 'uplink' port to a gateway router. There were two hosts in the private VLAN. The user on Host 2 could not send data to Host 1 as the AT-9924 switch was stopping communication between hosts in the private VLAN, only allowing traffic to pass to the gateway router connected to the uplink port.

The malicious user on Host 2 instead sent data to the IP address of Host 1, but with the MAC address of the gateway router. The gateway router then forwarded this data to Host 1, bypassing private VLAN Layer 2 security.

To stop this Layer 3 'workaround' to bypass private VLAN security, we configured a hardware filter on the switch to block the traffic relayed by the gateway router. The filter blocked traffic with source and destination addresses in the private VLAN's subnet.

Note: *We need to be sure the subnet mask of the filter excludes the gateway router, so desired traffic can still be passed between the Hosts and the gateway.*

Network setup

The AT-9924 switch was setup in the following network:

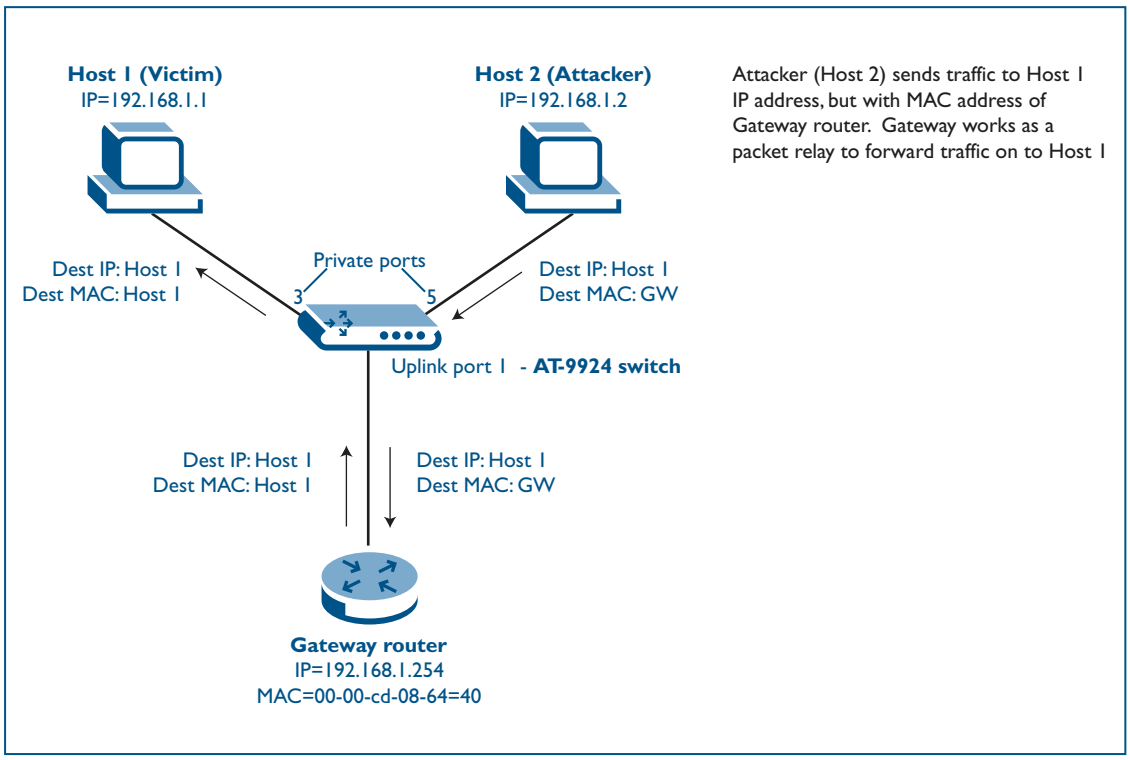

Switch configuration

```
Set system name="9924"

create vlan="vlan2" vid=2 private
add vlan="2" port=1 uplink
add vlan="2" port=2-8

create class=1 ipsa=192.168.1.0/25 ipda=192.168.1.0/25
add switch hwf class=1 action=discard
```

Attack results

The network was set up as above without the hardware filter in the configuration. Testing showed that under normal circumstances the two hosts could not communicate, as private VLAN functionality prevented this as expected.

When the 'malicious' Host 2 sent traffic with a destination IP address of Host 1, and a destination MAC address of the gateway, the traffic was relayed by the gateway and received by Host 1. The private VLAN security was circumvented.

Host 2 sent 10 packets as above. The output (next page) shows that the AT-9924 switch has received the 10 packets from Host 2 on port 5 and sent them out the uplink to the gateway on port 1. The gateway has received and then relayed these 10 packets back to the AT-9924 switch. Once relayed back from the gateway they have then been transmitted out port 3 to Host 1. We can then see that Host 1 has received the 10 packets.

```
Manager Gateway> sh eth cou

ETH instance 0:       5281 seconds    Last change at:       5160 seconds
Interface MIB Counters

  Receive:                              Transmit:
ifInOctets                     792     ifOutOctets                    1018
ifInUcastPkts                   10     ifOutUcastPkts                   10
ifInNUcastPkts                   0     ifOutNUcastPkts                   1
```

```
Manager Mustang> sh swi port=1,3,5 cou

 Switch Port Counters
--------------------------------------------------------------------------------
Port 1. Ethernet MAC counters:
 Combined receive/transmit packets by size (octets) counters:
  64                               22 512 - 1023                          0
  65 - 127                          0 1024 - MaxPktSz                     0
  128 - 255                         0
  256 - 511                         0

 General Counters:
 Receive                           Transmit
  Octets                          704 Octets                            704
  Pkts                             10 Pkts                               10

Port 3. Ethernet MAC counters:
 Combined receive/transmit packets by size (octets) counters:
  64                               11 512 - 1023                          0
  65 - 127                          0 1024 - MaxPktSz                     0
  128 - 255                         0
  256 - 511                         0

 General Counters:
 Receive                           Transmit
  Octets                           0 Octets                            704
  Pkts                             0 Pkts                               10

Port 5. Ethernet MAC counters:
 Combined receive/transmit packets by size (octets) counters:
  64                               0 512 - 1023                           0
  65 - 127                         0 1024 - MaxPktSz                      0
  128 - 255                        0
  256 - 511                        0
 General Counters:
 Receive                           Transmit
  Octets                           0 Octets                              0
  Pkts                            10 Pkts                                0
  CRCErrors                        0
  MulticastPkts                    0 MulticastPkts                       0
```

```
Manager Host_1> sh eth cou

ETH instance 0:       5296 seconds    Last change at:       5161 seconds

Interface MIB Counters

  Receive:                                Transmit:
ifInOctets                     864     ifOutOctets                       0
ifInUcastPkts                   10     ifOutUcastPkts                    0
ifInNUcastPkts                   1     ifOutNUcastPkts                   0
```

The same test was run with the following hardware filter added in:

```
create class=1 ipsa=192.168.1.0/25 ipda=192.168.1.0/25
add switch hwf class=1 action=discard
```

Note: *The subnet mask on the source and destination IP's is /25 instead of /24 so it does not cover the whole class C subnet. This is so communication between the hosts and the gateway at the .254 address will not be blocked, but traffic between the Hosts (which fall in the /25 subnet) will be blocked from Layer 3 relaying.*

With the AT-9924 switches' hardware filter in place, Host 2 was unable to relay traffic through the gateway to Host 1 because the switch discarded it. However, the hosts could communicate with the gateway for access out of the network.

Test 4: Random frame stress attack

The network in this example was set up with two IXIA ports sending random data to the MAC address of the 'Receiver' on the same VLAN (VLAN 1). At the same time the two hosts in VLAN 2 were sending a continuous ping between them. We can observe that the ping continues to be successful and that none of the attack data is 'leaked' across to VLAN 2.

The two IXIA ports sent data with the same destination MAC address so it would always reach the 'Receiver' host. One was a UDP stream and the other a TCP data stream. The following fields were random in packets from both streams:

- Frame size
- Inter-packet gap
- Source MAC address
- Source IP address
- Destination IP address

Network setup

The AT-9924 switch was setup in the following network:

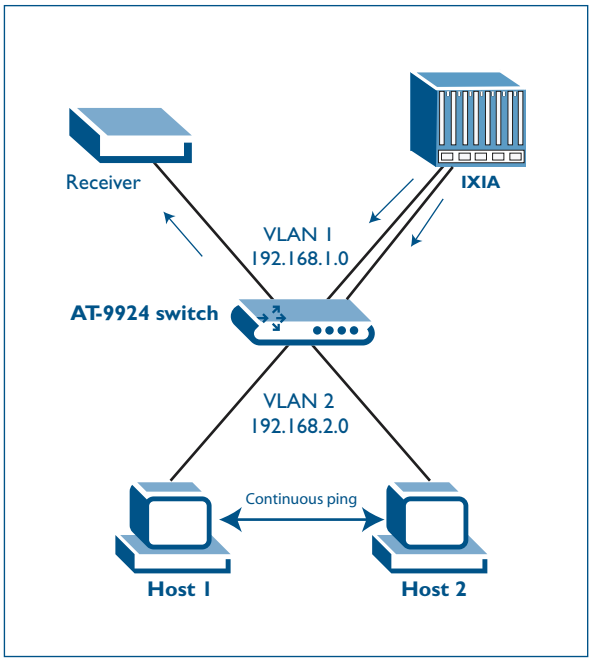

Switch configuration

```
set system name=Mustang
create vlan=vlan2 vid=2
add vlan=2 port=2,4
enable ip
add ip int=vlan1 ip=192.168.1.254
add ip int=vlan2 ip=192.168.2.254
```

Attack results

The continuous ping was started between the hosts in VLAN 2. The two data streams were observed coming into the switch and were transmitted on the 'Receivers' port. The ping continued to be successful. Once the ping and data streams were stopped, the port counters were examined.

These showed that the ping traffic was all that had been seen on the Hosts and their connected switch ports in VLAN 2. The random frame data had been received and transmitted out the port to the Receiver with no sign of VLAN leakage.

The following output shows the counters for the VLAN 2 hosts and their switch ports:

```
Manager Host_1> sh eth cou

ETH instance 0:       8070 seconds    Last change at:       6759 seconds

Interface MIB Counters

  Receive:                              Transmit:
ifInOctets                   53430     ifOutOctets                  53430
ifInUcastPkts                  685     ifOutUcastPkts                 685



Manager Host_2> sh eth cou

ETH instance 0:       8106 seconds    Last change at:       6786 seconds

Interface MIB Counters

  Receive:                              Transmit:
ifInOctets                   53430     ifOutOctets                  53430
ifInUcastPkts                  685     ifOutUcastPkts                 685



Manager Mustang> sh swi port=2,4 cou

Port 2. Ethernet MAC counters:
 Combined receive/transmit packets by size (octets) counters:
  64                                0 512 - 1023                        0
  65 - 127                       1370 1024 - MaxPktSz                   0
  128 - 255                         0
  256 - 511                         0

 General Counters:
 Receive                            Transmit
  Octets                        47950 Octets                       47950
  Pkts                            685 Pkts                           685


Port 4. Ethernet MAC counters:
 Combined receive/transmit packets by size (octets) counters:
  64                                0 512 - 1023                        0
  65 - 127                       1370 1024 - MaxPktSz                   0
  128 - 255                         0
  256 - 511                         0

 General Counters:
 Receive                            Transmit
  Octets                        47950 Octets                       47950
  Pkts                            685 Pkts                           685
```

Note that no other 'leaked' traffic from the Random Frame attack data is seen on the VLAN 2 switch ports, only the pings between the two VLAN 2 Hosts.

Test 5: STP attack

Three AT-9924 switches were setup running STP. There were two VLANs (VLAN 2 and VLAN 3) with a separate STP instance on each. The STP priorities were set as follows:

Switch ASTP priority = 10000

Switch BSTP priority = 20000

Switch CSTP priority = 30000

So we know that switch A was the root bridge and switch B the backup root bridge.

A stream of 2Mbps of Broadcast traffic was sent into the root bridge (switch A) on VLAN 2. This was forwarded to the other two switches on VLAN 2. Switch C with the lowest (STP) priority in VLAN 2 had its STP priority changed to 100 so that it would become the 'new' root bridge and force an STP reconvergence.

The VLAN 3 ports were monitored to see if any broadcast traffic was seen before, during and after the STP reconvergence.

Network setup

The AT-9924 switch was setup in the following network:

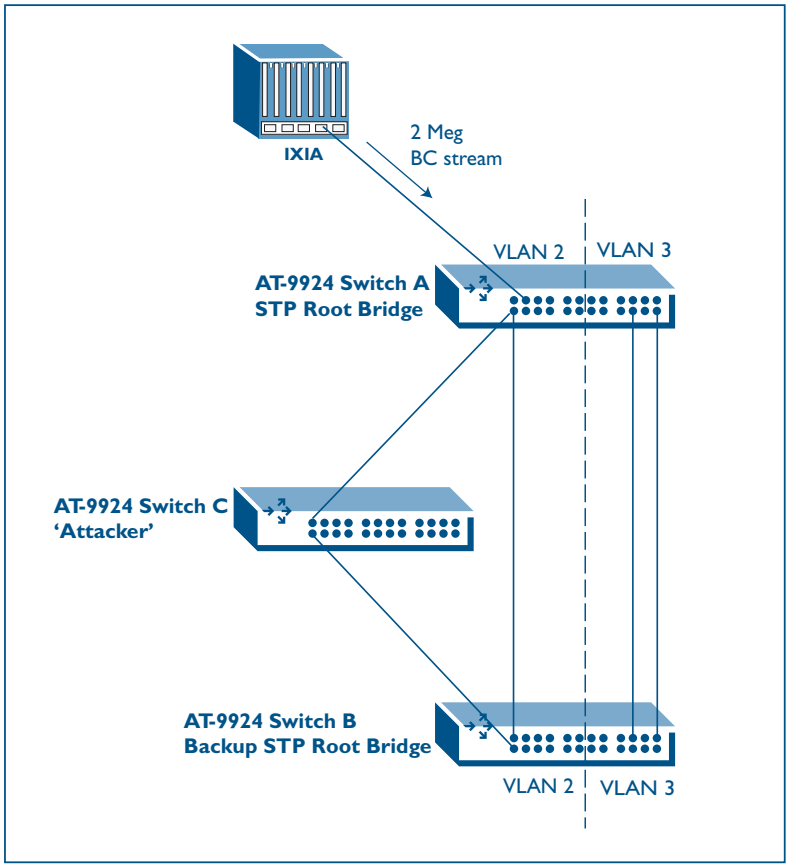

Test 6: Multicast brute force attack

Attack description

The multicast brute force attack involves a storm of Layer 2 multicast packets being sent into the switch to see if they are constrained within the ingress VLAN.

Prevention

The Layer 2 multicast packets should be constrained within the ingress VLAN. No packets should be 'leaked' to other VLANs.

Attack test

Two 100Mbps streams of multicast frames were sent from the IXIA into VLAN 1 on the AT-9924 switch. IGMP snooping was disabled on the switch so that the multicast packets would be flooded within the VLAN. These multicast packets should be flooded out of all VLAN 1 ports and be seen by the 'Multicast Receiver'.

While the multicast traffic was being flooded in VLAN 1, there was a continuous ping from VLAN 2 to the 'Host' switch and from VLAN 3 to the Host switch. These pings should continue successfully. There should be no 'leaked' multicast packets seen on the VLAN 2 and 3 ports.

Network setup

The AT-9924 switch was setup in the following network:

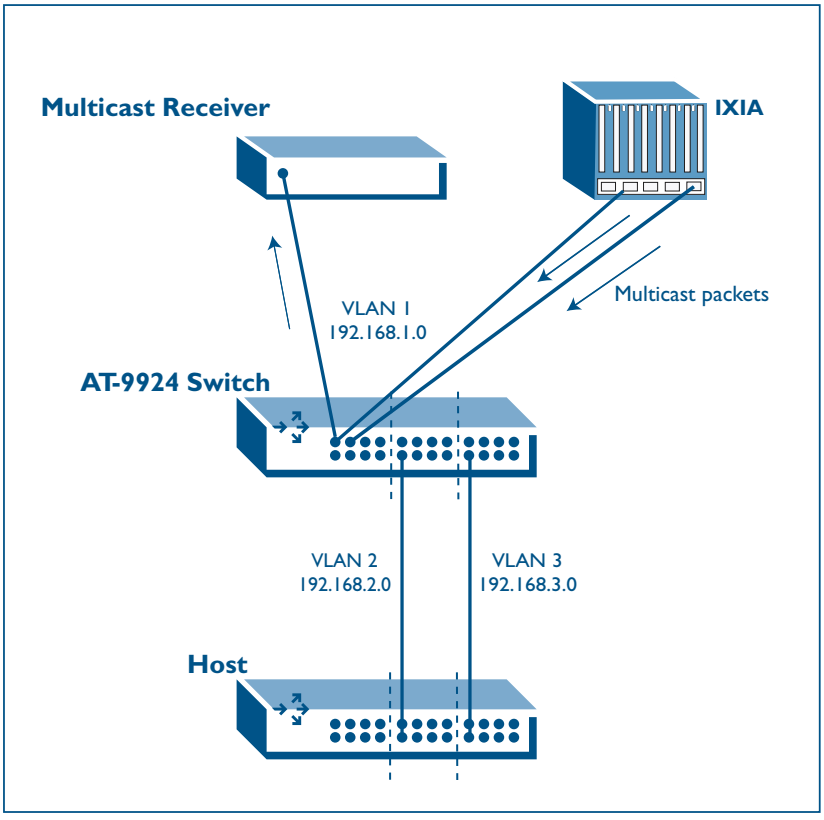

Switch configuration

```
set system name="9924"
disable igmpsnooping
create vlan="vlan2" vid=2
create vlan="vlan3" vid=3
add vlan="2" port=9-16
add vlan="3" port=17-24
enable ip
add ip int=vlan1 ip=192.168.1.254
add ip int=vlan2 ip=192.168.2.254
add ip int=vlan3 ip=192.168.3.254
```

Attack results

The test was run and the pings between the VLAN 2 / 3 ports and the 'Host' switch continued successfully. The multicast packets were flooded as expected out VLAN 1 ports and seen by the 'Multicast Receiver'.

The output below shows that multicast packets were seen only on the VLAN 1 ports (1- 3). The multicast packets were received on ports 2 and 3 from the IXIA and flooded out of ports 1- 3 in VLAN 1. Port 1 was connected to the 'Multicast Receiver'.

```
Manager 9924> sh switch port=1-3 count
 Switch Port Counters
------------------------------------------------------------------------------
Port 1. Ethernet MAC counters:
 Combined receive/transmit packets by size (octets) counters:
  64                          118416127 512 - 1023                          0
  65 - 127                            0 1024 - MaxPktSz                     0
  128 - 255                           0
  256 - 511                           0

 General Counters:
 Receive                                Transmit
  Octets                              0 Octets                   7578632128
  Pkts                                0 Pkts                      118416127
  CRCErrors                           0
  MulticastPkts                       0 MulticastPkts             118416127
  BroadcastPkts                       0 BroadcastPkts                     0
  FlowCtrlFrms                        0 FlowCtrlFrms                      0
  OversizePkts                        0
  Fragments                           0
  Jabbers                             0
  UpsupportOpcode                     0
  UndersizePkts                       0
                                        Collisions                        0
                                        LateCollisions                    0
                                        ExcessivCollsns                   0
Miscellaneous Counters:
  MAC TxErr                           0
  MAC RxErr                           0
  Drop Events                         0
```

```
Manager 9924> sh switch port=1-3 count (continued)

 Port 2. Ethernet MAC counters:
  Combined receive/transmit packets by size (octets) counters:
  64                        118416127 512 - 1023                        0
  65 - 127                          0 1024 - MaxPktSz                   0
  128 - 255                         0
  256 - 511                         0

 General Counters:
 Receive                            Transmit
  Octets                   3821644736 Octets                  3756987392
  Pkts                       59713199 Pkts                      58702928
  CRCErrors                         0
  MulticastPkts              59713199 MulticastPkts             58702928
  BroadcastPkts                     0 BroadcastPkts                    0
  FlowCtrlFrms                      0 FlowCtrlFrms                     0
  OversizePkts                      0
  Fragments                         0
  Jabbers                           0
  UpsupportOpcode                   0
  UndersizePkts                     0
                                      Collisions                       0
                                      LateCollisions                   0
                                      ExcessivCollsns                  0
Miscellaneous Counters:
  MAC TxErr                         0
  MAC RxErr                         0
  Drop Events                       0

Port 3. Ethernet MAC counters:
 Combined receive/transmit packets by size (octets) counters:
  64                        118416127 512 - 1023                        0
  65 - 127                          0 1024 - MaxPktSz                   0
  128 - 255                         0
  256 - 511                         0
General Counters:
 Receive                            Transmit
  Octets                   3756987392 Octets                  3821644736
  Pkts                       58702928 Pkts                      59713199
  CRCErrors                         0
  MulticastPkts              58702928 MulticastPkts             59713199
  BroadcastPkts                     0 BroadcastPkts                    0
  FlowCtrlFrms                      0 FlowCtrlFrms                     0
  OversizePkts                      0
  Fragments                         0
  Jabbers                           0
  UpsupportOpcode                   0
  UndersizePkts                     0
                                      Collisions                       0
                                      LateCollisions                   0
                                      ExcessivCollsns                  0
Miscellaneous Counters:
  MAC TxErr                         0
  MAC RxErr                         0
  Drop Events                       0
```

Port 10 in VLAN 2 and port 20 in VLAN 3 continued to successfully ping the 'Host' switch while the multicast traffic was being flooded in VLAN 1. Counters for ports 10 and 20 show (below) that the only traffic seen on the VLAN 2 and 3 ports is ping traffic; that is, no multicast traffic has 'leaked' into these VLANs.

```
Manager 9924> sh switch port=10,20 count
 Switch Port Counters
-------------------------------------------------------------------------------
Port 10. Ethernet MAC counters:
 Combined receive/transmit packets by size (octets) counters:
  64                              0 512 - 1023                          0
  65 - 127                      850 1024 - MaxPktSz                     0
  128 - 255                       0
  256 - 511                       0
General Counters:
 Receive                          Transmit
  Octets                      29750 Octets                        29750
  Pkts                          425 Pkts                            425
  CRCErrors                       0
  MulticastPkts                   0 MulticastPkts                     0
  BroadcastPkts                   0 BroadcastPkts                     0
  FlowCtrlFrms                    0 FlowCtrlFrms                      0
  OversizePkts                    0
  Fragments                       0
  Jabbers                         0
  UpsupportOpcode                 0
  UndersizePkts                   0
                                    Collisions                        0
                                    LateCollisions                    0
                                    ExcessivCollsns                   0
 Miscellaneous Counters:
  MAC TxErr                       0
  MAC RxErr                       0
  Drop Events                     0

Port 20. Ethernet MAC counters:
 Combined receive/transmit packets by size (octets) counters:
  64                              0 512 - 1023                          0
  65 - 127                      916 1024 - MaxPktSz                     0
  128 - 255                       0
  256 - 511                       0
General Counters:
 Receive                          Transmit
  Octets                      32060 Octets                        32060
  Pkts                          458 Pkts                            458
  CRCErrors                       0
  MulticastPkts                   0 MulticastPkts                     0
  BroadcastPkts                   0 BroadcastPkts                     0
  FlowCtrlFrms                    0 FlowCtrlFrms                      0
  OversizePkts                    0
  Fragments                       0
  Jabbers                         0
  UpsupportOpcode                 0
  UndersizePkts                   0
                                    Collisions                        0
                                    LateCollisions                    0
                                    ExcessivCollsns                   0
Miscellaneous Counters:
  MAC TxErr                       0
  MAC RxErr                       0
  Drop Events                     0
```

Conclusion - Layer 2 VLAN attacks

The AT-9924 has proven to be very robust with regards to VLAN security. No VLAN attacks were successful against the switch when appropriate measures were in place. These methods used appropriate configurations and some relied on inherent security features built into the switch.

USA Headquarters | 19800 North Creek Parkway | Suite 200 | Bothell | WA 98011 | USA | T: +1 800 424 4284 | F: +1 425 481 3895
European Headquarters | Via Motta 24 | 6830 Chiasso | Switzerland | T: +41 91 69769.00 | F: +41 91 69769.11
Asia-Pacific Headquarters | 11 Tai Seng Link | Singapore | 534182 | T: +65 6383 3832 | F: +65 6383 3830
www.alliedtelesyn.com

C613-16050-00 REV B

Connecting The IP World

Allied Telesyn®