Patch Release Note

# Patch 89262-09
# For AT-8900 Series Switches

## Introduction

This patch release note lists the issues addressed and enhancements made in patch 89262-09 for Software Release 2.6.2 on existing models of AT-8900 Series switches. Patch file details are listed in Table 1.

Table 1: Patch file details for Patch 89262-09.

| | |
|---|---|
| **Base Software Release File** | 89-262.rez |
| **Patch Release Date** | 18-Feb-2005 |
| **Compressed Patch File Name** | 89262-09.paz |
| **Compressed Patch File Size** | 813281 bytes |

This release note should be read in conjunction with the following documents:

- Release Note: Software Release 2.6.2 for AT-8900 Series switches (Document Number C613-10399-00 REV A) available from *www.alliedtelesyn.co.nz/documentation/documentation.html*.
- AT-8900 Series Switch Documentation Set for Software Release 2.6.2 available on the Documentation and Tools CD-ROM packaged with your switch, or from *www.alliedtelesyn.co.nz/documentation/documentation.html*.

***WARNING: Using a patch for a different model or software release may cause unpredictable results, including disruption to the network. Information in this release note is subject to change without notice and does not represent a commitment on the part of Allied Telesyn International. While every effort has been made to ensure that the information contained within this document and the features and changes described are accurate, Allied Telesyn International can not accept any type of liability for errors in, or omissions arising from the use of this information.***

Some of the issues addressed in this Release Note include a level number. This number reflects the importance of the issue that has been resolved. The levels are:

**Level 1** This issue will cause significant interruption to network services, and there is no work-around.

**Level 2** This issue will cause interruption to network service, however there is a work-around.

**Level 3** This issue will seldom appear, and will cause minor inconvenience.

**Level 4** This issue represents a cosmetic change and does not affect network operation.

*From Patch 89262-07 onwards, issues for each patch are listed in severity order as per the levels above. Enhancement PCRs are listed after Level 4 issues.*

# Features in 89262-09

Patch 89262-09 includes all issues resolved and enhancements released in previous patches for Software Release 2.6.2, and the following enhancements:

## Level 1

**PCR: 40703** **Module: IPG** **Level: 1**

IP multicast routes were not being updated correctly when a port went down, this caused incorrect behaviour for PIM and DVMRP. This issue has been resolved.

## Level 2

**PCR: 40604** **Module: STP** **Level: 2**

There was potential for STP instances operating in rapid mode (RSTP) to malfunction. Possible impacts of this were; failure of designated ports to transition to the forwarding state, BPDUs not being transmitted and failure of the root port to age out old information. This issue has been resolved.

**PCR: 40605** **Module: IPG, IP6, SW56, SWI, VLAN** **Level: 2**

IGMP and MLD were not being updated when the master port of a trunk changed.

**PCR: 40607** **Module: IPG** **Level: 2**

IGMP routes with many downstreams were taking a considerable amount of time to update when any downstreams were deleted, added or changed. This issue has been resolved.

**PCR: 40612 Module: IPG, DNS Relay Level: 2**

There was an issue in DNS Relay that resulted in a memory leak. The leak occurred when a response to a relayed DNS request contained an authoritative nameserver or additional information and the DNS request was forwarded to one of those servers. There was also an issue whereby DNS queries handled by DNS Relay would sometime result is corrupt entries in the DNS cache. These issues have been resolved.

**PCR: 40622 Module: IP6 Level: 2**

Occasionally, if the switch was configured with an IPsec policy that allowed all IPv6 ICMP packets matching types 133,134,135 and 136, a ping to the switch might fail. This issue has been resolved.

**PCR: 40628 Module: VLAN Level: 2**

ARP packets were not being classified into subnet-based VLANs, as subnet association rules applied to IP packets only. This issue has been resolved, and ARP packets are now classified into the subnet VLANs.

**PCR: 40629 Module: FW Level: 2**

A switch reboot could occur when a large number of proxied connections were rapidly established, for example, during a SYN attack. This issue has been resolved.

**PCR: 40635 Module: PPP Level: 2**

A switch reboot could occur when **ifAdminStatus** was set to "down" on some PPPoE interfaces. This issue has been resolved.

**PCR: 40638 Module: Firewall Level: 2**

When a global interface was dynamically assigned an IP address via DHCP or PPP, NAT configurations with dynamic private interfaces (**interface=dyn-<*dyn-int-name*>**) were not updated. This resulted in the failure of sessions received on dynamic private interfaces because the global IP address was invalid. This issue has been resolved.

**PCR: 40640 Module: QOS Level: 2**

If a traffic class, which had a large number of flowgroup and classifires associated with it, was added to a port, then a reboot could occur. This issue has been resolved.

**PCR: 40646 Module: OSPF, IPG Level: 2**

The switch would sometimes add a route with it own IP address for the NEXTHOP address. This issue has been resolved.

**PCR: 40650 Module: SWMX Level: 2**

If IPv6 traffic was being passed through a switch with an IPv6 accelerator card installed, the CPU could reach 100% after some variable time, depending on the data rate. This issue has been resolved.

**PCR: 40656 Module: IPG Level: 2**

In some configurations using OSPF, the command **reset ip** could cause a switch reboot. This issue has been resolved.

**PCR: 40668 Module: FFILE Level: 2**

The switch would not respond to the setting of the **maxqueueseverity** parameter when configuring logging. This issue has been resolved.

**PCR: 40674 Module: PIM Level: 2**

When two devices had two links between them (one of which was put into a blocking state by STP) and had multicast traffic was passing between them (the upstream device had the higher STP priority), and the link that wasn't blocked by STP was removed, then replaced, multicast traffic would be sent to both ports. This issue has been resolved.

**PCR: 40685 Module: SWMX Level: 2**

The switch would not update its IP and next-hop tables correctly when receiving an ARP for a next-hop that it had already learnt on a different port. For example, when two connected VRRP devices changed states between master and backup. This issue has been resolved.

**PCR: 40690 Module: SWI, SW56 Level: 2**

Previously, when switch filters were defined, devices with matching MAC addresses could still receive Layer 3 routed packets when connected to the switch on ports other than the ports prescribed by the switch filter. This issue has been resolved.

**PCR: 40691 Module: VRRP Level: 2**

It was possible for the VRRP priority to be incorrectly decremented to 0. If this happened on both the Master and Slave, a VRRP advertisement packet storm occurred. This issue has been resolved.

**PCR: 40695 Module: VLAN, SWI Level: 2**

When an uplink port was added to a private VLAN as tagged, the tagged status was not being set correctly which was causing untagged packets to be flooded when they should have been dropped. This issue has been resolved.

**PCR: 40727 Module: IPG Level: 2**

When the firewall IDENT PROXY was disabled, the acknowledgement number in the TCP [RST, ACK] packet was not correct. This issue has been resolved.

**PCR: 40749 Module: SSH Level: 2**

SSH server would fail after a variable number of connections. This issue has been resolved.

**PCR: 50015 Module: IPG Level: 2**

When the local device receives a packet and wishes to forward it, but cannot do so immediately because the next hop has not yet been resolved, the packet is queued while the corresponding IP address is being resolved by ARP. If no ARP response was received, the switch could continue to try to ARP for the next-hop indefinitely, so the packet buffer would potentially never be freed, nor would any ICMP unreachables be sent. This issue has been resolved.

**PCR: 50030 Module: IPG Level: 2**

In very rare circumstances the switch could reboot when forwarding multicast data. This issue has been resolved.

## Level 3

**PCR: 40624 Module: QOS Level: 3**

If a traffic class was already assigned to a policy and an attempt was made to assign it to a another policy, an error would occur. This is as expected. However, the traffic class in question would then be left in such a state that its configuration could not be changed. This issue has been resolved, and a more explanatory error message has been added.

**PCR: 40645 Module: FFILE Level: 3**

The command **create conf=<filename>** would return different error messages depending on the length of the invalid string. This issues has been resolve, so that the error message returned is now "invalid file name, should be <dev>:<mod>\<fil>.<typ>.", in all cases.

**PCR: 40741 Module: PING Level: 3**

If after enabling a ping poll, the command **purge ping totally** was entered, a reboot could occur. This issue has been resolved.

**PCR: 50045 Module: SWMX Level: 3**

When attempting to retrieve the counters for the Gigabit switch ports (ports 49-52) via SNMP, all counters would always be all reported as zero and were not consistent with the values output by the show switch port count. This issue has been resolved.

## Level 4

**PCR: 40748 Module: STP, RSTP, SWI Level: 4**

The following dot1dStp MIB counters were not set or updated correctly,

- dot1StpTimeSinceTopologyChange
- dot1StpTopologyChanges
- dot1StpDesignatedRoot
- dot1StpRootCost
- dot1StpMaxAge
- dot1StpHelloTime
- dot1StpHoldTime
- dot1StpForwardDelay
- dot1StpPortPriority
- dot1StpPortPathCost
- dot1StpPortDesignatedRoot
- dot1StpPortDesignatedCost
- dot1StpPortDesignatedBridge
- dot1StpPortDesignatedPort

Also, the number of STP/RSTP topology changes that had occurred since a restart and the time since the last topology change occurred were not displayed as part of the **show stp** command output. These issues have been resolved.

## Enhancements

No enhancements.

# Features in 89262-08

Patch file details are listed in Table 2.

Table 2: Patch file details for Patch 89262-08.

| | |
|---|---|
| **Base Software Release File** | 89-262.rez |
| **Patch Release Date** | 15-Oct-2004 |
| **Compressed Patch File Name** | 89262-08.paz |
| **Compressed Patch File Size** | 324452 bytes |

Patch 89262-08 includes all issues resolved and enhancements released in previous patches for Software Release 2.6.2, and the following enhancements:

## Level 1

No issues.

## Level 2

**PCR: 40570 Module: IP6 Level: 2**

IPv6 multicast packets were not being forwarded when coming from a 6-over-4 tunnel. This issue has been resolved.

**PCR: 40573 Module: LOG Level: 2**

If the log module was configured to store a very large number of messages (for example, more than 3000 messages), a watchdog timeout could occur when the **show debug** command was executed. This issue has been resolved.

Please note that this problem would not occur when the **show log** command was executed. A temporary work-around would be to disable the log module before executing the **show debug** command.

**PCR: 40591 Module: SWMX Level: 2**

In a multicast setup, there was a possibility that all multicast and broadcast packets were not being forwarded as expected. This caused symptoms such as RIPv1 or RIPv2 losing routes, PIM neighbourhoods being lost, and other protocols using broadcast or multicast packets would fail to register any peers/neighbours. Also, multicast packets are not switched. This issue has been resolved.

**PCR: 40596 Module: SWMX Level: 2**

When STP is enabled, when frames with a multicast destination MAC address were sent, a loop was observed within the network. This issue has been resolved.

**PCR: 40601 Module: SWMX Level: 2**

In network configurations with multiple paths to neighbours on an interface, if the interface changed state from up to down, then up again, the

interface route may have been erroneously deleted by a route update. This issue has been resolved.

**PCR: 40619 Module: IPG Level: 2**

The **valid** and **preferred** parameters were incorrectly added to the dynamic **set ipv6 prefix** configuration. The default **onlink** and **autonomous** parameters were also being included. This issue has been resolved.

## Level 3

**PCR: 40589 Module: IPG Level: 3**

The **counter** parameter did not exist in the **show igmpsnooping** command. This issue has been resolved.

**PCR: 40603 Module: SWI, SWMX Level: 3**

An error message is now shown if the user tries to assign more than the allowable number of traffic class (or default traffic class) entries to active QOS policies.

**PCR: 40606 Module: VLAN Level: 3**

When a Core port was added to a nested VLAN, its Ingress Filtering attribute was set to be On and Acceptable Frame Type was set to be Admit Only VLAN tagged Frames. If the port was deleted from that VLAN, its Ingress Filtering attribute must be set to Off and Acceptable Frame Type must be set to Admit All Frames, however, the attributes were not changed when the port was deleted from the nested VLAN. This issue has been resolved.

**PCR: 40617 Module: TTY Level: 3**

The manager prompt did not appear when using a telnet session until the [Enter] key was pressed several times. This issue has been resolved.

**PCR: 40618 Module: SWI Level: 3**

An FTP server located on a private interface of a firewall with NAT enabled may have had its ftp-data (tcp/20) source port translated to another port. This could lead to a firewall rejecting the data packets, as they do not strictly conform to RFC 959. This issue has been resolved, ensuring that ftp-data packets are sent from port 20 on the firewall, even when NAT is enabled.

**PCR: 40618 Module: SWI Level: 3**

Fixing speed on a fibre SPF to 1000mfull in a configuration script showed up incorrectly as 10mfull actual speed after reboot. This issue has been resolved.

## Level 4

No issues.

## Enhancements

No issues.

# Features in 89262-07

Patch file details are listed in Table 3.

Table 3: Patch file details for Patch 89262-07.

| | |
|---|---|
| **Base Software Release File** | 89-262.rez |
| **Patch Release Date** | 29-Sept-2004 |
| **Compressed Patch File Name** | 89262-07.paz |
| **Compressed Patch File Size** | 316588 bytes |

Patch 89262-07 includes all issues resolved and enhancements released in previous patches for Software Release 2.6.2, and the following enhancements:

## Level 1

No issues.

## Level 2

**PCR: 40272 Module: IPG Level: 2**

The switch learned an ARP entry for an IP address that was already configured on one of its interfaces. This issue has been resolved, and the receipt of spoofed ARP packets will now generate a log message.

**PCR: 40356 Module: BGP Level: 2**

1) A switch reboot could be observed if an IP interface was deleted while BGP was learning routes.

2) A switch reboot could be observed if a BGP peer was disconnected while the other peer was learning routes.

3) Excessive log messages were generated when the switch was low on memory.

4) Not all routes were removed from the BGP route table when a peer was disabled.

These issues have all been resolved.

**PCR: 40419 Module: OSPF, IPG Level: 2**

If OSPF was configured using the command **set ospf dyninterface=stub,** to advertise dynamic interfaces such as PPPoE interfaces as stub links, the links were not being advertised as expected. This issue has been resolved.

**PCR: 40420 Module: BGP, IGP Level: 2**

In some circumstances, when a BGP peer became physically disconnected, subsequent deletion of routes could cause a switch reboot to occur.

**PCR: 40457 Module: STP Level: 2**

If an STP disabled port was moved from one VLAN to another, it caused STP to reconverge. This issue has been resolved.

**PCR: 40460 Module: SWMX Level: 2**

Incorrect behaviour would occur if the switch was configured with a hardware filter that used a classifier matching on a VLAN, plus some additional IP Layer 3 or Layer 4 information (e.g. UDPDPort); and another filter using a classifier matching on just a VLAN. The result would be that only non-IP packets would be matched by this second hardware filter. IP packets not matching the first filter would not match the second, and therefore would not be classified by the switch. This symptom could also occur for classifiers using MACSA and MACDA. This issue has been resolved.

**PCR: 40465 Module: PIM6, PIM4 Level: 2**

The switch could reboot when a user changed the Rendezvous Point Candidate (RPC) priority in the PIM6 module. This issue has been resolved.

**PCR: 40470 Module: BGP Level: 2**

When BGP redistributed routes, locally imported routes were selected rather than peer learnt routes. This issue has been resolved.

**PCR: 40473 Module: IPG Level: 2**

When IP filters are configured on IP interfaces of the switch, all IP routing must occur in software. To achieve this, the IP forwarding table in the ASIC must be emptied. There was an issue whereby Interim IP route entries (IP routes for which there is not yet an ARP entry for the nexthop address) were not removed from the hardware IP forwarding table when an IP filter was added to an IP interface. This issue has been resolved.

**PCR: 40478 Module: IPG Level: 2**

If the switch received many packets to Layer 3 route, but did not have ARP entries for the destination address, the switch memory would deplete, leading to a switch reboot. This issue has been resolved.

**PCR: 40479 Module: OSPF Level: 2**

For OSPF-originated routes, it was possible for a route to be deleted from the IP routing table, but still be referenced by OSPF. This could cause a switch reboot when later generating a summary LSA that contained the old route. This occurred using the **reset ip** command. This issue has been resolved.

**PCR: 40481 Module: IPG Level: 2**

If the configuration script on the switch contained several commands for creating static arp entries, the switch could reboot on startup. This issue has been resolved.

**PCR: 40487 Module: CORE Level: 2**

A memory leak could occur when the accessing of the environmental monitoring chip failed. This issue has been resolved.

**PCR: 40488 Module: IPG Level: 2**

When a BGP peer was physically disconnected, the best routes for all prefixes learned were written to the silicon. The memory that was used to contain these routes was sometimes not freed, resulting in a memory leak. This issue has been resolved.

**PCR: 40496 Module: DHCP Level: 2**

When DHCP is enabled, it reclaims IP addresses at switch startup to determine if the addresses are in use or not. If, during this process, DHCP was disabled then re-enabled, the switch would not attempt to reclaim the remaining IP address ranges. This would lead to the rejection of DHCP requests for IP addresses that were still being reclaimed. This issue has been resolved.

**PCR: 40500 Module: BGP Level: 2**

When doing AS-Path regular expression matching in the **show bgp route** command, the router could reboot if there were withdrawn routes in the BGP table. This issue has been resolved.

**PCR: 40510 Module: VRRP Level: 2**

A configuration generated with the **create config** or **show config dynamic** commands could under some circumstances, include a **disable vrrp = <vrrpid>** command even if VRRP was enabled. This issue has been resolved.

**PCR: 40516 Module: DHCP Level: 2**

While initialising a range, the switch acting as a DHCP server may release a dynamic entry incorrectly. This issue has been resolved.

**PCR: 40519 Module: SWI Level: 2**

When the STP mode was changed, all the static arps on the ports belonging to the STP were deleted. This issue has been resolved.

**PCR: 40520 Module: DVMRP Level: 2**

Multicast data could not flow from PIM to DVMRP on a PIM/DVMRP border switch. This issue has been resolved.

**PCR: 40522 Module: CLASSIFIER Level: 2**

If a **create classifier** command contained the **ethformat** parameter, the resulting entry in a script created using the **create config**, or **show config dynamic** commands could sometimes be incorrect. This issue has been resolved.

**PCR: 40530 Module: IPG Level: 2**

When both Load Balancer and Firewall were configured, the very first TCP session was established after rebooting. Subsequent TCP session startup packets may have been routed out to an incorrect interface causing sessions to not be established. This issue has been resolved.

**PCR: 40531 Module: VLAN Level: 2**

A customer port in one nested VLAN could be set to be a core port in another nested VLAN. This issue has been resolved.

**PCR: 40535 Module: SWMX Level: 2**

When a particular model of SFP is fitted, and a switch reboot occurs, the switch could get locked in a reboot cycle, and never manage to finish booting. This issue has been resolved.

**PCR: 40537 Module: BGP Level: 2**

When the status of an interface changed, the BGP revaluation of IP routes for redistribution (via the **add bgp import** or **add bgp network** commands) was incorrect. This gave inconsistent BGP route tables depending on the order of events. This issue has been resolved.

**PCR: 40538 Module: IP6, SWI Level: 2**

Multicast data failed to be forwarded by PIM-SM if an MLD report was received on the switch before the corresponding multicast stream had arrived. This issue has been resolved.

**PCR: 40540 Module: SWI Level: 2**

Problems could occur if hardware filters or QOS policies were created using a classifier matching on: a Layer 2 attribute other than MAC Destination (e.g. MACSA or VLAN ID), and one or more other parameters that match on IPv4 frames (e.g. IP address, UDP/TCP parameters), and another classifier matching on Layer 2 attributes only.

The possible problems were; a generic Layer 2 match only succeeding if the frame was not of IP type, or false-positive matches when a frames had the same MAC Destination as a classifier designed to match on MAC Source address. This issue has been resolved.

**PCR: 40541 Module: CLASSIFIER Level: 2**

If a classifier was created that used the **ethformat** and **protocol** parameters, the resulting entry in a script created using the **create config** command or **show config dynamic** commands was not always correct. This issue has been resolved.

**PCR: 40543 Module: SWMX Level: 2**

When a Novell IPX packet was received by a port in a a protocol-based VLAN, it was being flooded out all ports on the default VLAN. This issue has been resolved.

**PCR: 40544 Module: VLAN Level: 2**

1.) A port was remaining in the default VLAN after having been added to a nested VLAN.

2.) A port was being returned to the default VLAN when deleted from one nested VLAN even if it was still a member of another nested VLAN.

3.) A port associated with a protocol or subnet rule was not being returned to the default VLAN when deleted from a nested VLAN.

These issues have all been resolved.

**PCR: 40549 Module: SWI Level: 2**

The receipt of two IP packets whose destination IP addresses were subnet addresses caused the switch to reboot. This issue has been resolved.

**PCR: 40550 Module: SWMX Level: 2**

Wrong bits were being set in the ASIC rule table for classifiers that matched on IPX source socket. As the result, the classifiers would match on incorrect values of IPX source socket. This issue has been resolved.

**PCR: 40554 Module: QoS Level: 2**

Some MIB values for the AT-QOS MIB, have been changed so that the output of the switch is now compatible with the latest version of the MIB.

**PCR: 40561 Module: SWMX Level: 2**

After the **disable switch learn** command had been executed successfully, automatic MAC learning was still operating. This issue has been resolved.

**PCR: 40562 Module: SWNP Level: 2**

If the command **enable switch accelerator function=icmpredirect** had been executed, there was no resulting entry in a script subsequently created by the **create config** or **show config dynamic** commands. This issue has been resolved.

**PCR: 40565 Module: SWMX Level: 2**

If two protocol VLAN association rules were added to two different VLANs, the second protocol association rule would not work. This issue has been resolved.

**PCR: 40571 Module: SWMX Level: 2**

When PIM or DVMRP was enabled, if IP multicast packets were received on the non-RPF (Reverse-Path-Forwarding) interface, i.e. the wrong ingress interface, the CPU could become highly utilised, and the packets were not correctly Layer 2 switched. This issue has been resolved.

**PCR: 40574 Module: SWMX Level: 2**

Adding, or deleting, port=42 to, or from, a VLAN could cause the entire Layer 2 Multicast Table to be cleared. This would subsequently cause high CPU utilisation under heavy multicast traffic. This issue has been resolved.

**PCR: 40586 Module: SWMX Level: 2**

If VRRP had been enabled on the switch, then the routing of any packets that entered the switch via a an interface on which VRRP was operating would be performed in software, rather than using the L3 switching process in the ASIC. This issue has been resolved.

**PCR: 40592 Module: BOOTP Level: 2**

If a timed-out ARP entry was renewed by BOOTP, the new entry be created with no port association. This issue has been resolved.

## Level 3

**PCR: 40471 Module: SWI Level: 3**

When an accelerator card is installed, and ports had been configured for mirroring, the **enable switch mirror** caused an unnecessary warning message to be displayed. This issue has been resolved.

**PCR: 40474 Module: IPG Level: 3**

When an accelerator card is installed, the **set switch mirror** command caused the switch to display an incorrect message saying that the maximum port number is 54. This issue has been resolved.

**PCR: 40493 Module: DHCP Level: 3**

In certain scenarios when acting as a DHCP server, the switch would send a DHCP ACK to an invalid MAC address. This issue has been resolved.

**PCR: 40498 Module: OSPF Level: 3**

When a virtual link end point is no longer reachable, the virtual interface is not brought down, and the virtual neighbour is not removed. This issue has been resolved.

**PCR: 40515 Module: QoS Level: 3**

Setting switch enhanced mode to “none” in order to disable QoS counters did not disable the QoS counters properly. A debug error message was shown when attempting to view the traffic class counters. This debug error message has now been removed and an appropriate error message is now displayed.

**PCR: 40525 Module: SWI, SWX Level: 3**

When MIB counters relating to packet flows reached their maximum possible value (0xFFFFFFFF), they should have returned to 0, and counted up from 0 again. However, they were remaining stuck at 0xFFFFFFFF.

This issue has been resolved, so that the counters will correctly roll over from 0xFFFFFFFF to 0 and start counting up again.

## Level 4

No issues.

## Enhancements

**PCR: 40511 Module: RSTP**

The RSTP module has been enhanced to detect simple loop scenarios downstream of an RSTP enabled edge port. If a loop is detected, the port is

placed into a Backup/Discarding/LoopbackDisabled state. In this state, all packets are discarded. The port transitions to a Designated/Discarding state after 3 x helloPeriod. If the loop still exists, the Backup/Discarding/LoopbackDisabled state is repeated.

**PCR: 40521 Module: TACACS+**

The new command **show tacplus** has been added. This command shows the module status, number of servers, and number of logged in users.

# Features in 89262-06

Patch file details are listed in Table 4.

Table 4: Patch file details for Patch 89262-06.

| | |
|---|---|
| **Base Software Release File** | 89-262.rez |
| **Patch Release Date** | 12-Aug-2004 |
| **Compressed Patch File Name** | 89262-06.paz |
| **Compressed Patch File Size** | 274040 bytes |

Patch 89262-06 includes all issues resolved and enhancements released in previous patches for Software Release 2.6.2, and the following enhancements:

**PCR: 31225 Module: IPG Level: 3**

While the switch was set with a CIDR interface address, when it received an ECHO request with a network broadcast destination address for a class C network, the switch sent the ECHO reply packet. Also, the switch forwarded the ECHO request packet using a broadcast MAC address. These issues have been resolved.

**PCR: 40008 Module: NTP Level: 3**

When the device operated in NTP Client mode, the SHOW TIME command sometimes displayed the incorrect time. This issue has been resolved.

**PCR: 40075 Module: OSPF Level: 2**

Total exception errors occurred when the OSPF DEFAULTROUTE was set to from ON to OFF. This problem has been resolved.

**PCR: 40123 Module: OSPF Level: 2**

OSPF did not refresh a network LSA when it received a LSA with errors from another vendor's device. This has now been fixed.

**PCR: 40261 Module: PIM4 Level: 4**

PIM counters were not totalling up correctly for erroneous packets if the type of PIM packet was not known. This issue has been resolved.

**PCR: 40266 Module: IPSEC Level: 2**

Out of sequence IPSEC packets could cause a switch reboot. This issue has been resolved.

**PCR: 40284 Module: PIM Level: 2**

When PIM-SM was configured and a very large number of IGMP v2 joins were received, a switch reboot could occur. This issue has been resolved.

**PCR: 40321 Module: ENCO, IPSEC Level: 2**

When **expirykbytes** was set to a low value in the IPsec policy, it was possible that a memory leak could occur if heavy IPsec traffic was being

transmitted while the IPsec SA renegotiation took place. This issue has been resolved.

**PCR: 40340 Module: IPG Level: 3**

The IP Options fields were being processed multiple times if a Firewall or NAT were enabled. This resulted in two Timestamps or Record_Route fields being added at each hop, instead of one. This issue has been resolved.

**PCR: 40350 Module: OSPF Level: 2**

All OSPF packets sent had an IP Precedence of 0 rather than 110. This issue has been resolved.

**PCR: 40372 Module: IPG Level: 2**

A slow memory leak was observed in some circumstances when adding and deleting routes in the routing table. This issue has been resolved.

**PCR: 40378 Module: PPP Level: 2**

If the remote PPPoE client was not responding to LCP Configure Requests, the PPPoE access concentrator would continually send configure requests, as the PPP template could not be configured to change this default setting. This issue has been resolved.

**PCR: 40399 Module: IPv6 Level: 3**

The **add ipv6 nd** command did not work when the port parameter was specified. This issue has been resolved.

**PCR: 40402 Module: IPSEC Level: 2**

When two devices (A and B) had an IPsec tunnel connecting them and the default route of device A was to device B, device B had a fatal error. If A lost a link, any packets for that link were delivered to B unencrypted. If these packets were routed through device B to device A, then B recognised the packets as needing to be decrypted and attempted it. This caused a fatal error. This issue has been resolved.

**PCR: 40403 Module: BGP Level: 2**

Procedures for handling **bgp update** messages which contained an invalid **next_hop** attribute specified in Section 6, RFC1711 were incorrect. This issue has been resolved.

**PCR: 40405 Module: ENCO Level: 2**

If the ENCO process used to encrypt an ISAKMP packet failed, a switch reboot could occur. This issue has been resolved.

**PCR: 40408 Module: SWMX,SWI Level: 2**

When a nested VLAN core port received a packet, it could trigger some debug to the console port, and prevent further communication via that console port. This issue has been resolved.

**PCR: 40411 Module: IP6 Level: 2**

In certain cases where static routes in a multi-path environment were used and routes were changed on the switch, the IPv6 flow table of the switch wasn't refreshed correctly, i.e. an entry in the flow table had no outgoing interface and ND entry. This issue has been resolved.

**PCR: 40413 Module: QOS Level: 3**

The command **reset qos accel** was generating an unexpected error message. This issue has been resolved.

**PCR: 40415 Module: VRRP Level: 2**

When a master VRRP router was configured from a bootup script, the transition to the MASTER state occurred before the Layer 2 interface had been initialised, preventing the gratuitous ARP from being sent. This issue has been resolved.

**PCR: 40416 Module: QOS Level: 3**

The output of the **show qos trafficclass=x** command was not correctly indicating the state of the parameter "Ignore BandwidthClass". This issue has been resolved.

**PCR: 40417 Module: OSPF Level: 3**

When LS Acks (Link State Advert acks) were received, they were compared against the transmitted LSA (Link State Advert). If it was the same, the LSA was removed from the re-transmission list. The algorithm used in this check has been changed to be compliant with the algorithm specified in section 13.1 of RFC2328, to determine if the LS Ack received is the instance as the LSA.

**PCR: 40418 Module: SWMX Level: 3**

When the IPv6 Accelerator card was enabled in loopback mode, and you used an SNMP management station to display the forwarding database, a random value was displayed for one CPU MAC entry. This issue has been resolved.

**PCR: 40422 Module: FIREWALL Level: 3**

A problem existed when setting non-default Firewall attack trigger levels for SMTP attacks. The **show firewall policy attack** output and dynamically generated configuration scripts were incorrect. This issue has been resolved.

**PCR: 40425 Module: VLAN Level: 2**

When a private VLAN had a tagged uplink, and at least one untagged private port, a configuration generated using the **create conf** command would contain incorrect information. This issue has been resolved.

**PCR: 40431 Module: SWMX Level: 2**

When the IPV6 accelerator card was present but disabled, the switch was not transmitting CPU-initiated packets after receiving non-reserved multicast packets. This issue has been resolved.

**PCR: 40433 Module: VLAN Level: 2**

When the nested VLAN feature was disabled using the **disable feature** command, the switch did not remove, from software and hardware, all of the nested VLAN's associated ports, protocols and subnets as well as the nested vlans. This issue has been resolved.

**PCR: 40440 Module: CLASSIFIER Level: 3**

For those classifiers that specified the IP protocol as a match criterion, the IP protocol number was being stored and displayed in a configuration file as a hexadecimal value rather than a decimal value. This issue has been resolved.

**PCR: 40441 Module: IPG, VRRP Level: 4**

If VRRP was enabled and a **reset ip** command was issued followed by a **disable vrrp** command, then the device would still reply to pings, even though the device was no longer the VRRP master. Duplicate echo replies were seen on the device sending the pings. This issue has been resolved.

**PCR: 40446 Module: DHCP Level: 2**

In certain situations, if a DHCP client used a DHCP relay agent to request IP addresses from the switch acting as the DHCP server on a different subnet, it was not be able to renew the IP address allocated to it. This issue has been resolved.

**PCR: 40453 Module: IPG Level: 2**

Particular IP packets (unicast destination IP, but multicast destination MAC) could result in a memory leak, which in some cases could cause the device to stop responding to the command line. This issue has been resolved.

**PCR: 40454 Module: SWNP Level: 3**

The **enable/disable switch accelerator** commands could be executed even though no IPv6 accelerator card was installed. This issue has been resolved.

**PCR: 40458 Module: IPG Level: 2**

The switch was accepting network RIP packets from foreign subnets. This issue has been resolved.

**PCR: 40463 Module: IPG Level: 2**

Under the IPv6 multipath environment, e.g. when both a static route and a RIPng route were available, if the static route was disconnected, the switch used the CPU to transmit outgoing packets on the RIPng Route, i.e using software routing. This issue has been resolved.

**PCR: 40509 Module: SWMX Level: 3**

NetBios responses to NetBeui packets were not being classified by a protocol based VLAN. This was because both are represented by 0xF0, but in the packet NetBios is 0xF0F1 and NetBeui is 0xF0F0. Now, when adding a VLAN classification rule for NetBeui (0xF0), two classification rules are added to the hardware. One for NetBeui (0xF0F0), and one for NetBios(0xF0F1).

# Features in 89262-05

Patch file details are listed in Table 5.

Table 5: Patch file details for Patch 89262-05.

| | |
|---|---|
| **Base Software Release File** | 89-262.rez |
| **Patch Release Date** | 6-Jul-2004 |
| **Compressed Patch File Name** | 89262-05.paz |
| **Compressed Patch File Size** | 222160 bytes |

Patch 89262-05 includes all issues resolved and enhancements released in previous patches for Software Release 2.6.2, and the following enhancements:

**PCR: 40304 Module: IPv6 Level: 2**

Occasionally a fatal error might be observed when the switch was removing IPv6 multicast downstream interfaces (for example, when downstream clients left a multicast group). This issue has been resolved.

**PCR: 40371 Module: VLAN Level: 2**

If the switch was configured with nested VLANs, and a command was entered to disable the nested VLAN feature, then a fatal error was observed as the switch removed the configured nested VLANs. This issue has been resolved.

**PCR: 40397 Module: IPv6 Level: 2**

A fatal error was observed after entering the command **restart reboot** or **restart switch**, when there were 1000 IPv6 interfaces configured on the device with all links up. This issue has been resolved.

**PCR: 40409 Module: SWNP Level: 2**

A Watchdog fatal error was observed when many (e.g. >1000) IPv6 interfaces configured on the device join a multicast group all at once. This issue has been resolved.

# Features in 89262-04

Patch file details are listed in Table 6.

Table 6: Patch file details for Patch 89262-04.

| | |
|---|---|
| **Base Software Release File** | 89-262.rez |
| **Patch Release Date** | 25-Jun-2004 |
| **Compressed Patch File Name** | 89262-04.paz |
| **Compressed Patch File Size** | 221932 bytes |

Patch 89262-04 includes all issues resolved and enhancements released in previous patches for Software Release 2.6.2, and the following enhancements:

**PCR: 40279 Module: IPG Level: 2**

Occasionally the device suffered a fatal error if it received a large number of directed broadcast packets. This issue has been resolved.

**PCR: 40313 Module: SWNP Level: 3**

Adding the maximum number of allowed filters to an IPv6 Accelerator hardware filter set (filling the hardware filters to capacity) would delete all existing filters and fail. This issues has now been resolved. The device now correctly accepts a full set of hardware filters.

**PCR: 40344 Module: IP6 Level: 2**

When multiple RIPng routes existed, the correct route was not chosen by the device. This issue has been resolved.

**PCR: 40349 Module: IPv6 Level: 3**

Attempts to ping a site-local address from a global unicast address would fail. This issue has been resolved.

**PCR: 40351 Module: VLAN Level: 3**

GVRP added tagged ports to dynamic VLANs as static entries. As a result, these member ports were not timed out or deregistered properly when the ports were disabled, the link went down, or GARP was disabled. This issue has been resolved. In addition, GVRP no longer operates on private or nested VLANs.

**PCR: 40352 Module: SWNP Level: 3**

IPv6 Accelerator MIB counters could not be reset. This issue has been resolved.

**PCR: 40363 Module: SWNP Level: 3**

The terminal session would freeze when a large number of IPv6 Accelerator hardware filters were added. This issue has been resolved.

**PCR: 40365 Module: EZDRV, IPv6, SWNP Level: 2**

When a large number of IPv6 interfaces were configured:

- a boot script could take several minutes to process
- processing multicast updates could take several minutes
- high packet loss sometimes occurred when updating hardware for a change in IPv6 multicast membership

These issues have been resolved.

**PCR: 40367 Module: SWMX Level: 3**

Enhancements to tuning and buffer configuration settings have improved reliability at extreme temperatures and performance.

**PCR: 40376 Module: QOS Level: 3**

The QoS MIB has been restructured to separate generic switching and AT-8948 specific MIB variables. Traffic class counters were always returned as 0. This PCR modifies PCR 40213 (see "Features in 89262-02" on page 26). You should obtain the latest revision of the QoS MIB from your authorised Allied Telesyn distributor, reseller or customer service representative.

**PCR: 40380 Module: CLASSIFR Level: 3**

The **create config** command generated duplicate entries for the **protocol** parameter in IPv6 classifiers. This issue has been resolved.

**PCR: 40391 Module: SWMX Level: 4**

The *ifJackType* MIB object (RFC 2239) always returned the value BNC, regardless of the actual GBIC/SFP installed. The correct value for the installed GBIC/SFP (e.g. Fiber, LC, BNC) is now returned.

**PCR: 40396 Module: SWMX Level: 3**

The command:

```
set qos port defaultqueue=value forcedefqueue=yes
```

failed to enforce the use of the default queue. This issue has been resolved.

**PCR: 40116 Module: FIREWALL Level: 2**

When the firewall was used on a NAT interface in conjunction with IP policy filters, Telnet to this interface was not possible. This issue has been resolved.

**PCR: 40355 Module: VRRP Level: 3**

When VRRP was enabled and an IP interface on which VRRP was operating went down, VRRP was not being disabled, preventing VRRP from transitioning to the Initial state. This issue has been resolved.

**PCR: 40364 Module: IPG Level: 3**

When IGMP Snooping was disabled the DVMRP forwarding database was not updated correctly. This issue has been resolved.

**PCR: 40382 Module: SWMX Level: 3**

The IPv6 Accelerator Card port (port 53) can only be associated with a QOS policy containing classifiers that specify **ethii-tagged** as the Ethernet encapsulation and **IPv6** as the protocol:

```
create classifier=rule-id ethformat=ethii-tagged
   protocol=ipv6 [other-parameters...]
```

The **add qos port** command now checks that all classifiers associated with the QoS policy being assigned to the port specify **ethformat=ethii-tagged** and **protocol=ipv6**.

# Features in 89262-03

Patch file details are listed in Table 7:

Table 7: Patch file details for Patch 89262-03.

| | |
|---|---|
| **Base Software Release File** | 89-262.rez |
| **Patch Release Date** | 02-Jun-2004 |
| **Compressed Patch File Name** | 89262-03.paz |
| **Compressed Patch File Size** | 181136 bytes |

Patch 89262-03 includes all issues resolved and enhancements released in previous patches for Software Release 2.6.2, and the following enhancements:

**PCR: 03420 Module: IPG, SWI Level: 3**

It is now possible to prevent specified ports from acting as IGMP all-group ports, and specify which ports are allowed to behave as all-group entry ports. This is enabled with the ENABLE IP IGMP ALLGROUP command, and disabled with the DISABLE IP IGMP ALLGROUP command.

For details, see "*IGMP Snooping All-Group Entry*" on page 28.

**PCR: 03890 Module: IGMP, SWI Level: 2**

The switch was adding a router port for multicast packets to destinations with an address in the range 224.0.0.x. Switch port entries are now only created for special router multicast addresses.

**PCR: 31133 Module: IPG**

This PCR introduces an enhancement that extends an issue that was resolved in PCR 03890, in which switch port entries are only created for special router multicast addresses. It is now possible to specify reserved multicast addresses that will be treated as multicast packets from routers. For details, see "*IGMP Snooping*" on page 30.

**PCR: 40112 Module: PIM6 Level: 2**

PIM Dense Mode *Graft* and *GraftAck* messages were not being sent. This issue has been resolved.

**PCR: 40249 Module: SWI, SWMX**

Fibre SFPs can now be set to a fixed speed. Copper SFPs still can not have their speed parameters modified.

**PCR: 40270 Module: PIM6**

The maximum number of PIM6 interfaces has been increased from 255 to 2000. The number of MLD interfaces has also been increased to 2000 to match the maximum number of PIM interfaces.

**PCR: 40274 Module: VLAN, RMON Level: 3**

In SNMPv2c the *histroyControlBucketsGranted* value was outside the acceptable range. This issue has been resolved.

**PCR: 40277 Module: IPG Level: 2**

BGP routes were not always readvertised to peers after a next-hop route was lost and then reactivated. This issue has been resolved.

**PCR: 40289 Module: IPG, IPV6 Level: 2**

The device's own MAC address was accepted as the next hop MAC address when adding a static IP ARP entry with the **add ip arp** command, or an IPv6 neighbour with the **add ipv6 nd** command. This issue has been resolved.

**PCR: 40290 Module: IPG Level: 3**

An incorrect nexthop for a RIP route was advertised in RIP messages. This issue has been resolved.

**PCR: 40299 Module: IPG Level: 2**

When the most recently added VLAN was deleted, the switch did not respond correctly to subsequent DHCP *Discover* messages. This issue has been resolved.

**PCR: 40301 Module: SWMX Level: 3**

The value for the *Device rule space limit* in the output of the **show switch** command was incorrect. This issue has been resolved.

**PCR: 40302 Module: SWI, SWNP Level: 3**

Incorrect options were returned with "?" help for the **show switch accelerator** command. This issue has been resolved.

**PCR: 40307 Module: SSL Level: 2**

A memory leak sometimes occurred when many HTTPS sessions (using SSL) were load balanced. This issue has been resolved.

**PCR: 40310 Module: IPV6 Level: 3**

The incorrect parameter **ipaddress** was returned with "?" help for the **delete ipv6 host=name** command. This issue has been resolved.

**PCR: 40312 Module: IPV6 Level: 2**

IPv6 packets travelling through a 6-to-4 tunnel were sent to an incorrect IPv4 end address, although the first packet was sent correctly. This issue has been resolved so that all packets are sent to the correct address.

**PCR: 40318 Module: IPV6 Level: 2**

The next hop of an IPv6 RIPng *Response* message was not assigned with a link local address when an invalid next hop was specified in the prefix field of the route table entry. This issue has been resolved.

**PCR: 40322 Module: TM Level: 3**

An error message was not returned if an unavailable interface was specified for the **enable test interface** command. This issue has been resolved.

**PCR: 40334 Module: VLAN Level: 2**

A dynamic VLAN created by GVRP should contain a tagged port, but it did not. This issue has been resolved.

**PCR: 40338 Module: VLAN Level: 2**

When an untagged port was added to a non-default VLAN, and then set as tagged, it could not be added back to the default VLAN as an untagged port. This issue has been resolved.

**PCR: 40341 Module: SWMX, SWI**

Hardware filtering has been enhanced so that the traffic class can be remapped using classifiers and hardware filtering. This can increase the chance of packets that match the classifier reaching the CPU because they are transmitted to the CPU on a different queue and DMA channel. To configure this type of filtering, use the command **add switch hwfilter action=setl2qos**.

**PCR: 40342 Module: SWI Level: 2**

It was not possible to set a QoS policy to a port if the corresponding classifier with IPv6 parameters was changed. This issue has been resolved.

**PCR: 40343 Module: IPV6, VLAN Level: 2**

If there were more than one port connected on the same IPv6 interface, traffic stopped if a cable was unplugged from the egress port of the switch and then plugged into a different port. This issue has been resolved.

**PCR: 40348 Module: VLAN Level: 2**

STP restarted when GVRP sent a *Join* message for a newly created dynamic VLAN. This issue has been resolved.

# Features in 89262-02

Patch file details are listed in Table 8:

Table 8: Patch file details for Patch 89262-02.

| | |
|---|---|
| **Base Software Release File** | 89-262.rez |
| **Patch Release Date** | 20-May-2004 |
| **Compressed Patch File Name** | 89262-02.paz |
| **Compressed Patch File Size** | 154860 bytes |

Patch 89262-02 includes all issues resolved and enhancements released in previous patches for Software Release 2.6.2, and the following enhancements:

**PCR: 40196 Module: SWNP, EZDRV**

A new command has been added to enable the switch to issue ICMP redirect messages for unicast IPv6 traffic, as recommended by RFC 2461 "Neighbor Discovery for IP Version 6 (IPv6)". See "*ICMP Redirect Messages for IPv6*" on page 33.

**PCR: 40213 Module: QOS**

A MIB for QoS has been added. A more recent version may be available from your authorised distributor or reseller.

**PCR: 40224 Module: EZDRV**

Network processor performance has been improved.

**PCR: 40245 Module: SWNP**

It is no longer necessary to add a filter number to the HWFILTER parameter in the ADD SWITCH ACCELERATOR HWFILTER command.

**PCR: 40315 Module: SWMX**

Some hardware settings for routing (narrow) RAM on the AT-ACC01 accelerator card have been modified.

# Features in 89262-01

Patch file details are listed in Table 9:

Table 9: Patch file details for Patch 89262-01.

| | |
|---|---|
| **Base Software Release File** | 89-262.rez |
| **Patch Release Date** | 20-May-2004 |
| **Compressed Patch File Name** | 89262-01.paz |
| **Compressed Patch File Size** | 29200 bytes |

Patch 89262-01 includes the following enhancements for Software Release 2.6.2:

**PCR: 40155 Module: OSPF**

The switch has been enhanced to enable up to 300 routes to be imported from BGP to OSPF. See "*Importing BGP routes into OSPF*" on page 34 for details.

**PCR: 40189 Module: IPV6**

IPv6 is now available when a software feature licence for IPV6 is enabled, or the AT-ACC01 network processor accelerator card is present.

# IGMP Snooping All-Group Entry

Because IGMP is an IP-based protocol, multicast group membership for VLAN aware devices is on a per-VLAN basis. If at least one port in the VLAN is a member of a multicast group, multicast packets will be flooded onto all ports in the VLAN by default.

*IGMP snooping* enables the switch to forward multicast traffic intelligently on the switch. The switch listens to IGMP membership reports, queries and leaves messages to identify the switch ports that are members of multicast groups. Multicast traffic will only be forwarded to ports identified as members of the specific multicast group.

This enhancement allows network managers to prevent specified ports from acting as IGMP all-group ports, and specify which ports are allowed to behave as all-group entry ports, by using the ENABLE IP IGMP ALLGROUP command.

For example, consider a video streaming service which has 15 channels. When the switch receives IGMP membership reports destined for the address 239.0.0.2 from an unauthorised user, all 15 channels of multicast data floods to that port, which may affect the service of the network. In order to avoid this, the network manager decides whether or not to allow a particular port to behave as an IGMP all-group port, e.g. port 8. Then, whenever the above IGMP membership report is sent, the switch will not automatically add port 8 as one of the egress ports for any IGMP membership report group, so video streaming will not get forwarded to disabled all-group ports selected by the network manager.

## Command Reference

This enhancement modifies one command:

- SHOW IP IGMP

and has two new commands:

- ENABLE IP IGMP ALLGROUP
- DISABLE IP IGMP ALLGROUP

# show ip igmp

**Syntax** `SHOW IP IGMP [COUNTER] [INTERFACE=`*`interface`*`]`

**Description** This command displays information about IGMP, and multicast group membership for each IP interface.

This enhancement includes the line "**Disabled All-groups ports**" on the output of this command, as shown in Figure 1 on page 29. Ports that are disabled have a "#" symbol next to the port number.

Figure 1: Example output from the **show ip igmp** command.

```
IGMP Protocol
--------------------------------------------------------------------------------
Status ........................... Enabled
Default Query Interval ........... 125 secs
Default Timeout Interval ......... 270 secs
Disabled All-groups ports ........ 1,5,7

Interface Name .......... vlan2 (DR)
IGMP Proxy .............. Off
Group List ..............

  Group. 238.0.1.2          Last Adv. 172.50.2.1         Refresh time 34 secs
  Ports 3,11,23

  Group. 224.1.1.2          Last Adv. 172.50.2.1         Refresh time 130 secs
  Ports 2,11,23

  All Groups                Last Adv. 172.50.1.1         Refresh time 45 secs
  Ports 1#,11,23


Interface Name .......... vlan4             (DR)
IGMP Proxy .............. Off
Group List ..............
  No group memberships.

--------------------------------------------------------------------------------
```

Table 10: New parameter in the output of the **show ip igmp** command.

| Parameter | Meaning |
|---|---|
| Disabled All-groups ports | A list of ports that are prevented from behaving as IGMP all-group ports. |

**Examples** To show information about IGMP, use the command:

```
SHOW IP IGMP
```

**See Also** ENABLE IP IGMP ALLGROUP
DISABLE IP IGMP ALLGROUP

# enable ip igmp allgroup

**Syntax** `ENABLE IP IGMP ALLGROUP=[{`*`port-list`*`|ALL}]`

where:

- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 ad end at m, where m is the highest numbered Ethernet switch port, including uplink ports.

**Description** This command enables the specified port(s) to behave as a multicast all-group ports.

The ALLGROUP parameter specifies the list of ports able to behave as all-group entry ports. If ALL is specified, all ports are able to behave as all-group entry ports. The default is ALL.

**Examples** To enable ports 1, 5 and 7 to behave as all-group entry ports, use the command:

```
ENABLE IP IGMP ALLGROUP=1,5,7
```

**See Also** DISABLE IP IGMP ALLGROUP
SHOW IP IGMP

# disable ip igmp allgroup

**Syntax** `DISABLE IP IGMP ALLGROUP=[{`*`port-list`*`|ALL}]`

where:

- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.

**Description** This command disables the specified port(s) from acting as a multicast all-group entry ports. Ports that are disabled have a "#" symbol next to the port number in the output of the SHOW IP IGMP command.

**Examples** To prevent ports 1, 5 and 7 from behaving as all-group entry ports, use the command:

```
DISABLE IP IGMP ALLGROUP=1,5,7
```

**See Also** ENABLE IP IGMP ALLGROUP
SHOW IP IGMP

# IGMP Snooping

You can now specify the mode of operation when IGMP Snooping is enabled with the command:

```
SET IGMPSNOOPING
   ROUTERMODE=[ALL|DEFAULT|IP|MULTICASTROUTER|NONE]
```

If ALL is specified, all reserved multicast addresses (i.e. 224.0.0.1 to 224.0.0.255) are treated as router multicast addresses.

If DEFAULT is specified, the following addresses are treated as router multicast addresses:

- IGMP Query, 224.0.0.1
- All routers on this subnet, 224.0.0.2
- DVMRP Routers, 224.0.0.4
- OSPFIGP all routers, 224.0.0.5
- OSPFIGP designated routers, 224.0.0.6

- RIP2 routers, 224.0.0.9
- All PIM routers, 224.0.0.13
- All CBT routers, 224.0.0.15

If IP is specified, addresses treated as router multicast addresses are specified by the user using the ADD IGMPSNOOPING ROUTERADDRESS and the DELETE IGMPSNOOPING ROUTERADDRESS commands. When in this mode, the switch retains previous addresses that have already been specified.

If MULTICASTROUTER is specified, the following addresses are treated as router multicast addresses:

- DVMRP Routers, 224.0.0.4
- All PIM routers, 224.0.0.13

If NONE is specified, the switch does not create router ports at all.

To add and delete reserved IP multicast addresses to and from the list of router multicast addresses that are specified by the SET IGMPSNOOPING ROUTERMODE command when the IP parameter is selected, use the commands:

```
ADD IGMPSNOOPING ROUTERADDRESS
DELETE IGMPSNOOPING ROUTERADDRESS
```

The IP addresses specified must be from 224.0.0.1 to 224.0.0.255.

To display information about the current list of configured IP multicast router addresses configured on the switch, use the command:

```
SHOW IGMPSNOOPING ROUTERADDRESS
```

# add igmpsnooping routeraddress

**Syntax** `ADD IGMPSNOOPING ROUTERADDRESS=`*`ipaddr`*`[,...]`

**Description** where:

- *ipaddr* is a reserved IP multicast address in dotted decimal notation.

This command adds reserved IP multicast addresses to the list of router multicast addresses. The IP address specified must be within the range 224.0.0.1 to 224.0.0.255. This command is only valid if the IGMP snooping router mode is set to IP with the SET IGMPSNOOPING ROUTERMODE command.

# set igmpsnooping routermode

**Syntax**

```
SET IGMPSNOOPING ROUTERMODE=
   {ALL|DEFAULT|IP|MULTICASTROUTER|NONE}
```

**Description**

This command sets the mode of operation for IGMP Snooping.

If ALL is specified, all reserved multicast addresses (i.e. 2240.0.1 to 224.0.0.255) are treated as router multicast addresses.

If DEFAULT is specified, the following addresses are treated as router multicast addresses:

- IGMP Query: 224.0.0.1
- All routers on this subnet: 224.0.0.2
- DVMRP Routers: 224.0.0.4
- OSPFIGP all routers: 224.0.0.5
- OSPFIGP designated routers: 224.0.0.6
- RIP2 routers: 224.0.0.9
- All PIM routers: 224.0.0.13
- All CBT routers: 224.0.0.15

If IP is specified, addresses that are treated as router multicast addresses are specified with the ADD/DELETE IGMPSNOOPING ROUTERADDRESS command. In this mode, the switch will retain previous addresses that have already been specified.

If MULTICAST is specified, the following addresses are treated as router multicast addresses:

- DVMRP Routers: 224.0.0.4
- All PIM routers: 224.0.0.13

If NONE is specified, no router ports are created.

# delete igmpsnooping routeraddress

**Syntax**

```
DELETE IGMPSNOOPING ROUTERADDRESS=ipaddr[,...]
```

where

- *ipaddr* is a reserved IP multicast address in dotted decimal notation.

**Description**

This command deletes reserved IP multicast addresses from the list of router multicast addresses. The IP address specified must be within the range 224.0.0.1 to 224.0.0.255. This command is only valid if the IGMP snooping router mode is set to IP with the SET IGMPSNOOPING ROUTERMODE command.

# show igmpsnooping routeraddress

**Syntax** `SHOW IGMPSNOOPING ROUTERADDRESS`

**Description** This command displays information about the list of configured IP multicast router addresses currently configured on the switch (Figure 2 on page 33).

Figure 2: Example output for the **show ipmpsnooping routeraddress** command.

```
IGMP Snooping Router Address
------------------------------------------------------------------------
IGMP Snooping Router Mode ..... IP

Router Address List
-------------------------------
224.0.0.4
224.0.0.6
224.0.0.80
224.0.0.43
224.0.0.23
224.0.0.15
224.0.0.60
------------------------------------------------------------------------
```

# ICMP Redirect Messages for IPv6

A new command has been added to enable the switch to issue ICMP redirect messages for unicast IPv6 traffic, as recommended by RFC 2461 "Neighbor Discovery for IP Version 6 (IPv6)". To enable the new function, use the command:

```
ENABLE SWITCH ACCELERATOR FUNCTION=ICMPREDIRECT
```

The switch accelerator does not issue ICMP redirect messages. When ICMPREDIRECT is enabled, any unicast IPv6 packet that has the same ingress and egress VLAN is routed in software rather than hardware, and ICMP redirect messages may be generated.

The purpose of an ICMP redirect is to replace layer 3 routing with layer 2 switching. On a layer 3 switch, switching and routing are both performed in hardware, so ICMP redirects offer no performance improvement. The ICMPREDIRECT function may impact performance in particular (unusual) network configurations, and is not recommended for use in a general network environment. The ICMPREDIRECT function is disabled by default.

To disable the ICMPREDIRECT function, use the command:

```
DISABLE SWITCH ACCELERATOR FUNCTION=ICMPREDIRECT
```

To see whether the function is enabled or disabled, use the command:

```
SHOW SWITCH ACCELERATOR
```

# Importing BGP routes into OSPF

## Introduction

With this enhancement you can import routes from BGP into OSPF. OSPF will then redistribute these routes. This enhancement adds three parameters to the **set ospf** command, and modifies the output of the **show ospf** command. The new parameters are **bgpimport**, **bgpfilter** and **bgplimit**.

BGP can learn thousands of routes, so it's important to consider the network impact of importing these routes. Routing devices in the OSPF domain may become overloaded if they store too many routes. You can prevent this by limiting the number of routes that will be imported.

***Do not enable the importing of BGP routes into OSPF unless you are sure about the consequences for the OSPF domain.***

## Enabling BGP route import

To enable importing BGP routes into OSPF, use the command:

```
set ospf bgpimport=on
```

## Limiting the number of routes

There are two ways to limit the number of BGP routes imported into OSPF. One way is to specify a maximum number of routes with the command:

```
set ospf bgplimit=1...300
```

When the limit is reached, the importing of routes will stop until existing routes are removed. Because they are BGP routes, actions of BGP control when the routes disappear.

The other way to limit the imported routes is to configure a routing filter. This filter is used in conjunction with the **bgpfilter** parameter in the **set ospf** command to control the passing of routing information in and out of the device. To configure a filter, use the **add ip filter** command:

```
add ip filter=filter-number {action=include|exclude}
   source=ipadd [smask=ipadd] [entry=entry-number]
```

Use this filter to limit imported BGP routes with the command:

```
set ospf bgpfilter=300...399
```

where the filter number is the previously configured filter.

Take care when configuring the IP filter. If the number of imported routes reaches the **bgplimit** parameter, you may not have imported all the routes specified with the **bgpfilter** parameter.

## Advertising desired routes

The order in which routes are added is arbitrary. This means that to have desired BGP routes advertised by OSPF, you must take care setting the **entry** number for the route filter with the **add ip route** command. Assign a low entry number to a filter used to import preferred BGP routes. Alternatively, set the **bgplimit** parameter above the total number of routes that BGP will ever add to the routing table.

**Configuration example**

This example supposes that you want to import the route 192.168.72.0 into the OSPF routing domain, but no other routes. This route is received on the gateway router as a BGP route. The following steps show the sequence of commands to use in this scenario.

1. Set up the IP filter:

```
add ip filter=300 source=192.168.72.0 smask=255.255.255.255
   action=include
```

2. Set up OSPF BGP import parameters:

```
set ospf bgpimport=on bgpfilter=300 bgplimit=1
```

3. Check that BGP has added the route to the IP route table:

```
show ip route=192.168.72.0
```

   The route should be visible in the output of the command.

4. Check that OSPF has imported the route:

```
show ospf lsa=192.168.72.0
```

   The output should show that there is an AS external LSA with this ID.

## Command Reference

This section contains details about the commands used to configure the BGP route import feature.

*Only the syntax for the BGP route import feature is shown here. For the full syntax of these commands, see the Software Reference on the Documentation and Tools CD-ROM bundled with your switch, or at www.alliedtelesyn.co.nz/documentation/manuals.*

# set ospf

**Syntax**

```
SET OSPF [BGPFilter={NOne|300...399}]
   [BGPImport={ON|OFF|True|False|YES|NO}]
   [BGPLimit=1...300] [other-parameters]
```

**Description**

This command sets general OSPF routing configuration parameters. Use this command to configure the importing of BGP routes into OSPF. See Table 11 on page 35 for details about each parameter.

Table 11: Parameters for the BGP route import feature in the **set ospf** command.

| Parameter | Option/Range | Description |
|---|---|---|
| BGPFilter | NOne | No filters are defined so all routes from BGP will be imported into OSPF. The default is **none**. |
| | 300...399 | The route filter that will be used when importing BGP routes into OSPF. Route filters are created with the **add ip filter** command. If a route filter is defined, the entries for the filter will include or exclude routes for importation. If routes have not been included by a previous entry, they will be excluded from the import. |
| BGPImport | ON\|True\|YES | Importing BGP routes into OSPF is enabled. |
| | OFF\|False\|NO | Importing BGP routes into OSPF is disabled. The default is **off**. |

| Parameter | Option/Range | Description |
|---|---|---|
| BGPLimit | 1...300 | The maximum number of BGP routes that can be imported into OSPF at a time. Once this limit is reached, importing stops until existing routes are removed. The default is **300**. |

*Caps denote command shortcuts

# show ospf

**Syntax** `SHow OSPF`

**Description** This command displays information about the general configuration of OSPF routing (Figure 3 on page 36, Table 12 on page 36). New entries for the BGP route import feature are in bold.

Figure 3: Example output from the **show ospf** command

```
Router ID ..................... 123.234.143.231
OSPF module status ............ Enabled
Area border router status ..... Yes
AS border router status ....... Disabled
PTP stub network generation ... Enabled
External LSA count ............ 10234
External LSA sum of checksums ... 1002345623
New LSAs originated ........... 10345
New LSAs received ............. 34500
RIP ........................... Off
BGP importing:
  Enabled ..................... Yes
  Import filter ............... 301
  Routes imported/limit ....... 214 / 300
Export static routes .......... Yes
Dynamic interface support ..... None
Number of active areas ........ 10
Logging ....................... Disabled
Debugging ..................... Disabled
AS external default route:
  Status ...................... Disabled
  Type ........................ 1
  Metric ...................... 1
```

Table 12: Parameters for the BGP route import feature in the output of the **show ospf** command.

| Parameter | Meaning |
|---|---|
| BGP importing | Information about the importing of BGP routes into OSPF. |
| Enabled | Whether or not the importing of BGP routes into OSPF is enabled; one of "Yes" or "No". |
| Import filter | The IP filter number used to filter routes before they are imported into OSPF, or "None" if no filters are used. |
| Routes imported/limit | The number of BGP routes imported into OSPF, and the maximum number of routes that can be imported at a time. |

# add ip filter

**Syntax** `ADD IP FILter=`*`filter-number`* `{ACTion=INCLude|EXCLude} SOurce=`*`ipadd`* `[SMask=`*`ipadd`*`] [ENTry=`*`entry-number`*`] [`*`other-parameters`*`]`

**Description** This command adds a pattern to a routing filter. For details about the command parameters, see Table 13 on page 37.

Table 13: Parameters for the BGP route import feature in the **add ip filter** command.

| Parameter | Option/Range | Description |
|---|---|---|
| filter-number | 300...399 | Filters in the range 300 to 399 are treated as routing filters, and use the **action** parameter to specify the action to take with a route that matches the pattern. |
| ACTion | | The action to take when the filter pattern is matched. |
| | INCLude | Route information matching the filter will be included. |
| | EXCLude | Route information matching the filter will be excluded. |
| SOurce | | The source IP address, in dotted decimal notation, for the filter pattern. |
| SMask | | The mask, in dotted decimal notation, to apply to source addresses for this pattern. The mask is used to determine the portion of the source IP address in the IP packet that is significant for comparison with this pattern. The values of **source** and **smask** must be compatible. For each bit in **smask** which is set to zero (0) the equivalent bit in **source** must also be zero (0). If **source** is not 0.0.0.0, then **smask** can not be 0.0.0.0. The default is 255.255.255.255, unless **source** is 0.0.0.0. |
| ENTry | entry-number | The **entry** parameter specifies the entry number in the filter which this new pattern occupy. Existing patterns with the same or higher entry numbers are pushed down the filter. The default is to add the new pattern to the end of the filter. |

*Caps denote command shortcuts

# Availability

Patches can be downloaded from the Software Updates area of the Allied Telesyn web site at *www.alliedtelesyn.co.nz/support/updates/patches.html*. A licence or password is not required to use a patch.