AlliedWare Plus™ OS

How To | Configure Policy-based Routing

Introduction

In today's high performance internetworks, organizations need the freedom to implement packet forwarding and routing according to their own defined policies in a way that goes beyond traditional routing protocol concerns. Policy-based routing provides a means to route particular packets to their destination via a specific next-hop.

Using policy-based routing you can mark packets so that certain kinds of traffic receive differentiated, preferential service. The specific path that these packets will take can be based on configurable parameters such as priority, address, protocol, or VLAN membership.

The Allied Telesis SwitchBlade™ x908, x900 series and x600 series switches all support policy-based routing at wire speed, using hardware-based policy routing facilities.

These switch models provide flexible options for configuring policy-based routing, and can support multiple different policy routes, with different nexthops, simultaneously. Possible benefits include both QoS and cost savings:

- QoS by using dedicated links for certain types of traffic.
- cost savings by splitting traffic between low-bandwidth, low-cost permanent paths and high-bandwidth, high-cost switched paths.

List of terms:
Ping Poll
Ping polling is used to ensure that a device is still present, live, and contactable in the network by periodically sending a packet to an IP address and waiting for a response. Configurable actions can be performed if responses are no longer arriving.
DSCP value
The Differentiated Services Code Point within the TOS field of an IP packet header. This is a 6-bit number in the range 0-63. The premark-dscp map allows you to assign packets to egress queues based on the DSCP value in the incoming packet.
Next Hop
IP routing involves forwarding packets from one router to the next, until they reach their destination. Routers do not need to know the full path to a packet's destination, they just need to know the next router to forward the packet on to. This 'next router' is referred to as the Next-Hop of an IP route.

What information will you find in this document?

This How To Note begins with the following information:

Then it describes how to configure policy-based routing using three examples:

Which products and software version does it apply to?

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

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

This configuration applies to switches running Software Version 5.3.1 or later.

Configuring Policy-based routing

Policy-based routing adds flexibility and provides network administrators with the ability to route traffic among multiple paths based on the traffic characteristics.

- From a configuration point of view, Policy-based routing is presented as an action in a QoS policy. It can be combined with other QoS actions performed in the same policy. Other QoS actions applied to the traffic will be carried out before the traffic is sent to its next-hop.
- The selection of packets to be policy routed is carried out by the standard class-map packet matching rules.
- If the switch doesn't have the configured next-hop in its ARP table, it will send an ARP request for it. If it does not receive a reply, the switch will send an ICMP destination unreachable message to the originating host. The switch does not use the configured default route if the policy-based routing next-hop is unavailable.

The following three examples will provide you with a good understanding of how policy-based routing can be implemented. Each example has step-by-step instructions and some suggested configuration commands.

Example 1 - managing VoIP

VoIP is being used to provide voice communications within an organisation. To ensure high voice quality, dedicated data circuits are leased between remote sites, and a central site. All VoIP data is routed via these dedicated circuits from the remote sites to a central site. The VoIP is then distributed from the central site to its eventual destinations. By using this hub-and-spoke arrangement of low-bandwidth circuits for transporting VoIP, the organisation is provided with good-quality voice communications in a cost-effective manner. All other data communication between the sites is transported over the Internet. The VoIP traffic is marked with a DSCP value of 46.

Policy routing is used to ensure that the VoIP packets are sent via the dedicated circuit, whilst all other data is sent over the Internet. The VoIP traffic is also put into the highest queue on the egress port, to ensure minimum packet loss and delay.

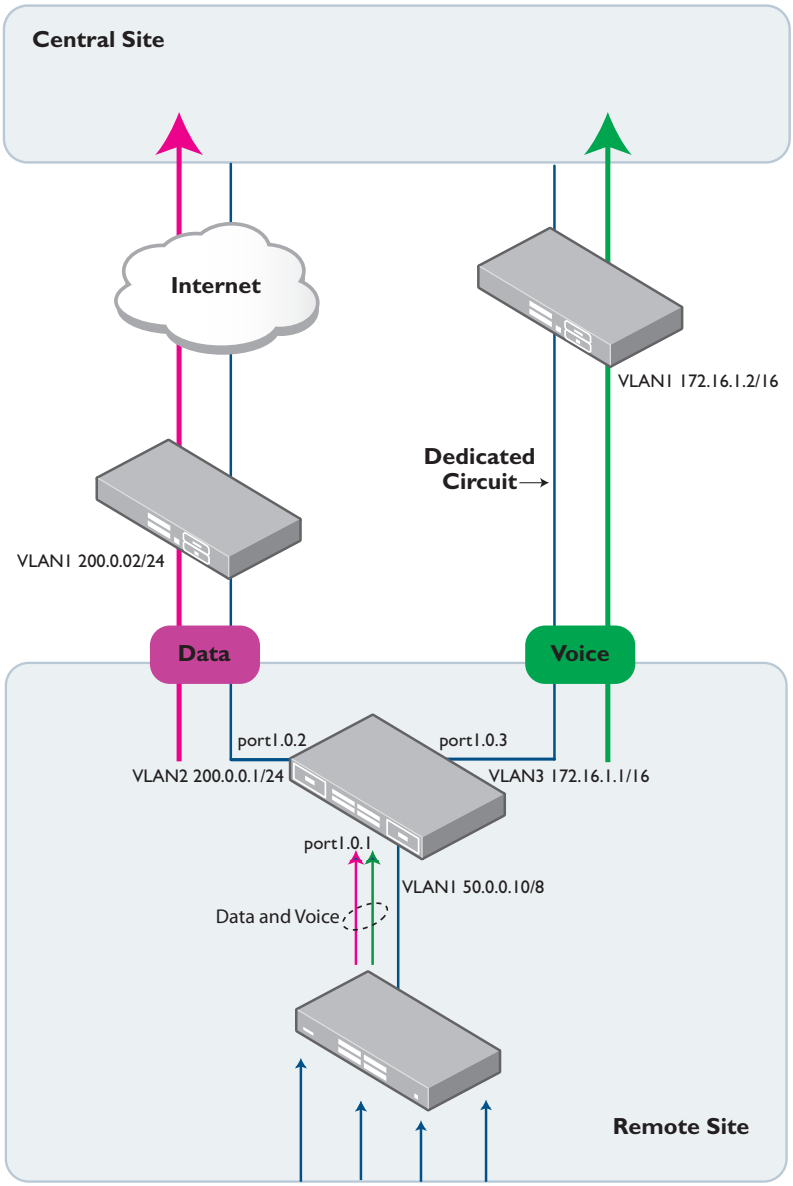

Configuring the x900 at the remote site

1. Create the VLANs and assign IP addresses to them

```
Awplus(config)#vlan database
Awplus(config-vlan)#vlan 2-3 state enable

#VLAN1 connects to the local LAN
Awplus(config)#interface vlan1
 Awplus(config-if)# ip address 50.0.0.10/8

#VLAN2 connects to the Internet
Awplus(config)#interface vlan2
 Awplus(config-if)#ip address 200.0.0.1/30

#VLAN3 is directed towards the dedicated circuit
Awplus(config)#interface vlan3
 Awplus(config-if)# ip address 172.16.1.1/30
```

2. Assign the VLANs to the ports (port 1.0.1 is already in VLAN1 by default)

```
Awplus(config)#interface port1.0.1
 Awplus(config-if)#switchport
 Awplus(config-if)#switchport mode access

Awplus(config)#interface port1.0.2
 Awplus(config-if)#switchport
 Awplus(config-if)#switchport mode access
 Awplus(config-if)#switchport access vlan2

Awplus(config)#interface port1.0.3
 Awplus(config-if)#switchport
 Awplus(config-if)#switchport mode access
 Awplus(config-if)#switchport access vlan3
```

3. Configure the default route. This is the route to the Internet

```
Awplus(config)#ip route 0.0.0.0/0 200.0.0.2
```

4. Enable QoS globally

```
Awplus(config)#mls qos enable
```

5. Create the class-map which will match on traffic with a DSCP value of 46

```
Awplus(config)#class-map test
Awplus(config-cmap)# match dscp 46
```

6. Create the policy-map which will specify the actions to be taken on the classified traffic. The first action instructs the switch to send this packet to queue 7 on the egress port. The second action sets the next-hop for this traffic to 172.16.1.2, so that it will be directed towards the dedicated circuit.

```
Awplus(config)#policy-map pbr
 Awplus(config-pmap)#class test
 Awplus(config-pmap-c)#set queue 7
Awplus(config-pmap-c)#set ip next-hop 172.16.1.2
```

7. Finally, the policy-map must be attached to the ingress port

```
Awplus(config)#interface port1.0.1
Awplus(config-if)#service-policy input pbr
```

Example 2 - reducing data transmission costs

In this scenario (using the same diagram as in example 1 above) we have a faster Internet connection on port1.0.2, which costs the company more to use. The ISP providing this faster connection charges on the basis of the amount of data sent over the connection.

The company has decided that traffic from their Web server will be sent to the Internet via this connection on Monday to Friday during business hours only (9.00am to 5.30 pm). This provides good web service during business hours, whilst keeping some limit on the total amount of data sent over the faster (more expensive) connection. Outside these times, traffic from the web server is sent via the default route. Any other traffic is always sent via the default route.

Additionally a ping poll can be configured that will regularly check that the Policy Route next-hop is reachable. If the ping poll fails to get a response from the next-hop, a trigger will be run which removes the policy from the ingress port. This will ensure that if the Policy Route next-hop fails, the configured default route will be used, ensuring no loss of connectivity. Once the next-hop is reachable again, another trigger adds the policy back onto the ingress port.

Configuration steps

1. Create the VLANs and assign IP addresses to them

```
Awplus(config)#vlan database
Awplus(config-vlan)#vlan 2-3 state enable

Awplus(config)#interface vlan1
Awplus(config-if)# ip address 50.0.0.10/8

Awplus(config)#interface vlan2
Awplus(config-if)# ip address 200.0.0.1/30

Awplus(config)#interface vlan3
Awplus(config-if)# ip address 172.16.1.1/30
```

2. Assign the VLANs to the ports (port 1.0.1 is already in VLAN1 by default)

```
Awplus(config)#interface port1.0.1
Awplus(config-if)#switchport
Awplus(config-if)#switchport mode access

Awplus(config)#interface port1.0.2
 Awplus(config-if)#switchport
 Awplus(config-if)#switchport mode access
 Awplus(config-if)#switchport access vlan 2

Awplus(config)#interface port1.0.3
 Awplus(config-if)#switchport
 Awplus(config-if)#switchport mode access
 Awplus(config-if)#switchport access vlan 3
```

3. Configure the default route

```
Awplus(config)#ip route 0.0.0.0/0 172.16.1.2
```

4. Enable QoS globally

```
Awplus(config)#mls qos enable
```

5. Create the access-list which classifies on traffic from the Web server

```
Awplus(config)#access-list 3001 permit tcp <web-server-IP
   address> eq 80 any
```

6. Apply this access-list to a class-map

```
Awplus(config)#class-map web-server
 Awplus(config-cmap)#match access-group 3001
```

7. Apply this class-map to a Policy-map and configure the policy to send the traffic matching class web-server to a specific next-hop

```
Awplus(config)#policy-map pbr
 Awplus(config-pmap)#class web-server
 Awplus(config-pmap-c)#set ip next-hop 200.0.0.2
```

8. Create the script which will add the policy to the ingress port when trigger 1 is run

```
Edit policy-on.scp
enable
conf  t
int port1.0.1
service-policy input pbr
```

9. Create the script which will remove the policy from the ingress port when trigger 2 is run

```
Edit policy-off.scp
enable
conf t
int port1.0.1
no service-policy input pbr
```

10. Configure the trigger which will add the policy to the ingress port at the required day/time

```
Awplus(config)#trigger 1
 Awplus(config-trigger)#type time 09:00
 Awplus(config-trigger)#day monday tuesday wednesday thursday
   friday
 Awplus(config-trigger)#script 1 policy-on.scp
```

11. Configure the trigger which will remove the policy from the ingress port at the required day/time

```
Awplus(config)#trigger 2
 Awplus(config-trigger)#type time 17:30
 Awplus(config-trigger)#day monday tuesday wednesday thursday
   friday
 Awplus(config-trigger)#script 1 policy-off.scp
```

12. Configure the ping poll which will regularly check that the next-hop is reachable

```
awplus(config)#ping-poll 1
awplus(config-ping-poll)#description “check policy route next
   hop”
awplus(config-ping-poll)#ip 200.0.0.2
 awplus(config-ping-poll)#source-ip 200.0.0.1
awplus(config-ping-poll)#active
```

13. Configure the trigger which will be activated when the ping poll fails

```
awplus(config)#trigger 3
 awplus(config-trigger)#type ping-poll 1 down
awplus(config-trigger)#script 1 policy-off.scp
awplus(config-trigger)#active
awplus(config-trigger)#time after 09:00:00 before 17:30:00
awplus(config-trigger)#day monday tuesday wednesday thursday
   friday
```

14. Configure the trigger which will be activated when the next-hop is reachable

```
awplus(config)#trigger 4
awplus(config-trigger)#type ping-poll 1 up
awplus(config-trigger)#script 1 policy-on.scp
awplus(config-trigger)#active
awplus(config-trigger)#time after 09:00:00 before 17:30:00
awplus(config-trigger)#day monday tuesday wednesday thursday
   friday
```

Example 3 - routing TCP traffic

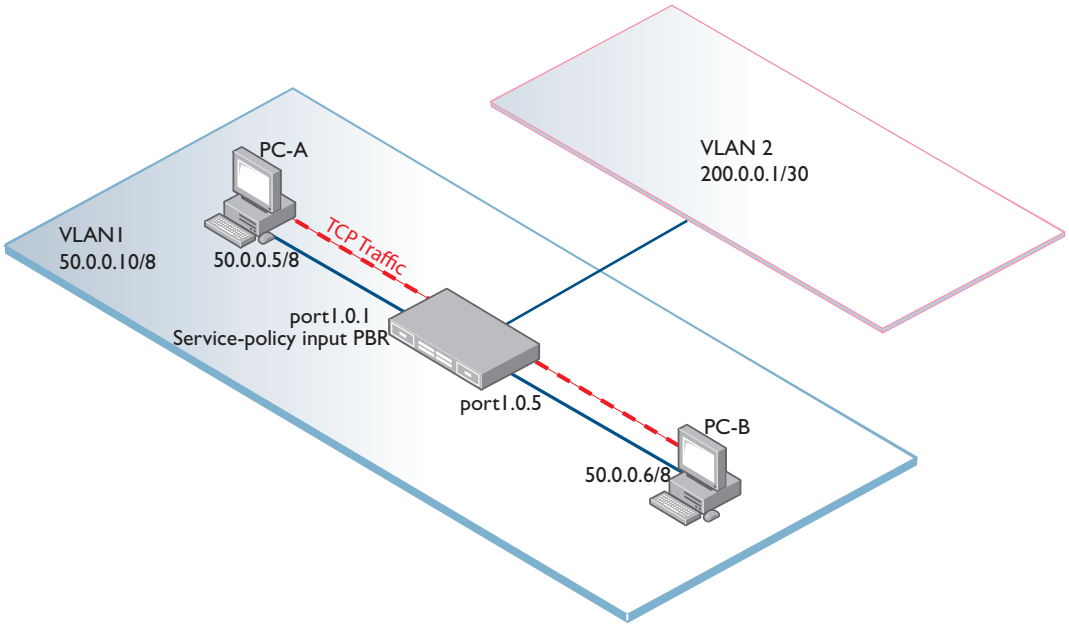

When using Policy-based Routing, traffic that's destined for the switch itself can be unexpectedly routed to the PBR next-hop. This includes TCP traffic such as telnet. In the example below, all TCP traffic is to be Policy Routed. However, this will also affect telnet to the switch itself. To avoid this, the following configuration can be used:

1. Create two IP interfaces - VLAN1 and VLAN2

```
Awplus(config)#interface vlan1
Awplus(config-if)# ip address 50.0.0.10/8

Awplus(config)#interface vlan2
Awplus(config-if)# ip address 200.0.0.1/30
```

We do not want traffic destined to the addresses of either of these interfaces to be Policy Routed.

2. Create ACLs for any TCP traffic destined for these networks, as well as the PBR ACL

```
Awplus(config)#mls qos enable
Awplus(config)#access-list 3001 permit tcp any 50.0.0.0/8
Awplus(config)#access-list 3002 permit tcp any 200.0.0.1/32
Awplus(config)#access-list 3003 permit tcp any any
```

Note: In the diagram above, any traffic matching the Policy-Based Routing access-list 3003, which classifies on TCP traffic ingressing port1.0.1, will be sent to the next-hop address of 200.0.0.2 on VLAN2. Even L2 traffic in the same VLAN through the switch will be sent to the PBR next-hop if it matches the PBR access-list.

This would mean that any TCP traffic from PC-A destined for another device connected to the switch in the same VLAN (e.g. PC-B) would not be forwarded to its destination but sent to the next-hop of 200.0.0.2. So, we configure ACL 3001 to match on any TCP traffic destined for the entire subnet in use on VLAN1, not just the IP address of the switch itself.

Allied Telesis recommends that the PBR policy always includes a class to allow traffic. This class must precede the Policy Routing class.

3. Apply these access-lists to class-maps

```
Awplus(config)#class-map local1
 Awplus(config-cmap)#match access-group 3001

Awplus(config)#class-map local2
 Awplus(config-cmap)#match access-group 3002

Awplus(config)#class-map tcp
 Awplus(config-cmap)#match access-group 3003
```

4. Apply the class-maps to a Policy-map and configure the policy to send the matching TCP traffic to a specific next-hop

The effect of the policy map is that:

- Any traffic matching classes local1 or local2 (i.e. ACLs 3001 or 3002) will simply pass through to the normal forwarding process.
- Any traffic matching class tcp will be policy routed.

```
Awplus(config)#policy-map pbr
  Awplus(config-pmap)#class  local1
  Awplus(config-pmap)#class  local2
 Awplus(config-pmap)#class tcp
  Awplus(config-pmap-c)#set ip next-hop 200.0.0.2
```

5. Finally, the policy-map must be attached to the ingress port

```
Awplus(config)#interface port1.0.1
```

```
Awplus(config-if)#switchport
Awplus(config-if)#switchport mode access
 Awplus(config-if)#service-policy input pbr
```

Note: Further reading - more detail of the configuration of policy-based routing can be seen in the QoS chapter of the x900 and x600 series software reference manuals.

USA Headquarters | 19800 North Creek Parkway | Suite 100 | 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.alliedtelesis.com

C613-16147-00 REV A

Connecting The IP World

Allied Telesis™