

Chapter 7

Overview of Layer 2 Switching

Introduction	7-2
Switch Ports	7-2
Enabling and Disabling Switch Ports	7-2
Autonegotiation of Port Speed and Duplex Mode	7-3
Port Trunking	7-3
Link Access Control Protocol (LACP)	7-3
Packet Storm Protection	7-3
Port Mirroring	7-3
Port Security	7-4
Virtual Local Area Networks (VLANs)	7-4
Creating VLANs	7-5
Summary of VLAN tagging rules	7-6
VLAN Interaction with STPs and Trunk Groups	7-6
Generic VLAN Registration Protocol (GVRP)	7-6
Quality of Service	7-7
Spanning Tree Protocol (STP)	7-8
Spanning Tree and Rapid Spanning Tree Port States	7-8
Multiple Spanning Tree Protocol (MSTP)	7-9
IGMP Snooping	7-9
Triggers	7-9

Introduction

This chapter outlines the Layer 2 and IP switching features on the switch, and how to configure some of them. For more detail, see:

- [Chapter 8, Switching](#)
- [Chapter 40, Quality of Service \(QoS\) on Switch Ports](#)

Switch Ports

Each switch port is uniquely identified by a port number. The switch supports a number of features at the physical level that allow it to be connected in a variety of physical networks. This physical layer (layer 1) versatility includes:

- Enabling and disabling of ports.
- Auto negotiation of port speed and duplex mode for all 10/100 BASE ports.
- Manual setting of port speed and duplex mode for all 10/100 BASE ports.
- Link up and link down triggers.
- Port trunking.
- Packet storm protection.
- Port mirroring.
- Support for SNMP management.

Enabling and Disabling Switch Ports

An switch port that is enabled is available for packet reception and transmission. Its administrative status in the Interfaces MIB is UP. Disabling a switch port does not affect the STP operation on the port. Enabling a switch port will allow the port to participate in spanning tree negotiation. A switch port that has been disabled by the Port Security feature cannot be enabled using the ENABLE SWITCH PORT command.

To enable or disable a switch port, use the commands:

```
ENABLE SWITCH PORT={port-list|ALL}
```

```
DISABLE SWITCH PORT={port-list|ALL}
```

Resetting ports at the hardware level discards all frames queued for reception or transmission on the port, and restarts autonegotiation of port speed and duplex mode. Ports are reset using the command:

```
RESET SWITCH PORT={port-list|ALL} [COUNTER]
```

To display information about switch ports, use the command:

```
SHOW SWITCH PORT[={port-list|ALL}]
```

Autonegotiation of Port Speed and Duplex Mode

Each of the switch ports can operate at either 10 Mbps or 100 Mbps, in either full duplex or half duplex mode. In full duplex mode a port can transmit and receive data simultaneously, while in half duplex mode the port can either transmit or receive, but not at the same time. This versatility makes it possible to connect devices with different speeds and duplex modes to different ports on the switch. Such versatility also requires that each port on the switch know which speed and mode to use.

Port Trunking

Port trunking, also known as port bundling or link aggregation, allows a number of ports to be configured to join together to make a single logical connection of higher bandwidth. This can be used where a higher performance link is required, and makes links even more reliable.

Link Access Control Protocol (LACP)

The Link Access Control Protocol (LACP) follows the IEEE Standard 802.3-2002, CSMA/CD access method and physical layer specifications. It enables trunk groups, called aggregated links, to be created automatically.

LACP operates where systems are connected over multiple communication links. It constantly monitors these links and automatically adds or removes them from trunk groups.

Packet Storm Protection

The packet storm protection feature allows you to set limits on the reception rate of broadcast, multicast and destination lookup failure packets. The software allows separate limits to be set for each port, beyond which each of the different packet types are discarded. The software also allows separate limits to be set for each of the packet types. Which of these options can be implemented depends on the model of switch hardware.

Port Mirroring

Port mirroring allows traffic being received and transmitted on a switch port to be sent to another switch port, the mirror port, usually for the purposes of capturing the data with a protocol analyser. This mirror port is the only switch port that belongs to no VLANs, and therefore does not participate in any other switching. Before the mirror port can be set, it must be removed from all VLANs except the default VLAN. The port cannot be part of a trunk group.

Port Security

The port security feature allows control over the stations connected to each switch port, by MAC address. If enabled on a port, the switch will learn MAC addresses up to a user-defined limit from 1 to 256, then lock out all other MAC addresses. One of the following options can be specified for the action taken when an unknown MAC address is detected on a locked port:

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

Virtual Local Area Networks (VLANs)

A Virtual LAN (VLAN) is a logical, software-defined subnetwork. It allows similar devices on the network to be grouped together into one broadcast domain, irrespective of their physical position in the network. Multiple VLANs can be used to group workstations, servers, and other network equipment connected to the switch, according to similar data and security requirements.

Decoupling logical broadcast domains from the physical wiring topology offers several advantages, including the ability to:

- Move devices and people with minimal, or no, reconfiguration
- Change a device's broadcast domain and access to resources without physically moving the device, by software reconfiguration or by moving its cable from one switch port to another
- Isolate parts of the network from other parts, by placing them in different VLANs
- Share servers and other network resources without losing data isolation or security
- Direct broadcast traffic to only those devices which need to receive it, to reduce traffic across the network
- Connect 802.1Q-compatible switches together through one port on each switch

Devices that are members of the same VLAN only exchange data with each other through the switch's switching capabilities. To exchange data between devices in separate VLANs, the switch's routing capabilities are used. The switch passes VLAN status information, indicating whether a VLAN is up or down, to the Internet Protocol (IP) module. IP uses this information to determine route availability.

The switch has a maximum of 255 VLANs, ranging from a VLAN identifier (VID) of 1 to 4094. When the switch is first powered up, a "default" VLAN is created and all ports are added to it. In this initial unconfigured state, the switch will broadcast all the packets it receives to the default VLAN. This VLAN has a VID of 1 and an interface name of `vlan1`. It cannot be deleted, and ports can only be removed from it if they also belong to at least one other

VLAN. The default VLAN cannot be added to any STP, but always belongs to the default STP. If all the devices on the physical LAN are to belong to the same logical LAN, that is, the same broadcast domain, then the default settings will be acceptable, and no additional VLAN configuration is required.

Creating VLANs

To briefly summarise the process of creating a VLAN:

1. Create the VLAN.
2. Add tagged ports to the VLAN, if required.
3. Add untagged ports to the VLAN, if required.

To create a VLAN, use the command:

```
CREATE VLAN=vlan-name VID=2..4094
```

Every port must belong to a VLAN, unless it is the mirror port. By default, all ports belong to the default VLAN as untagged ports.

To add tagged ports to a VLAN, use the command:

```
ADD VLAN={vlan-name|1..4094} PORT={port-list|ALL}  
FRAME=TAGGED
```

A port can be tagged for any number of VLANs.

To add untagged ports to a VLAN, use the command:

```
ADD VLAN={vlan-name|1..4094} PORT={port-list|ALL}  
[FRAME=UNTAGGED]
```

A port can be untagged for zero or one VLAN. A port can only be added to the default VLAN as an untagged port if it is not untagged for another VLAN. A port cannot transmit both tagged and untagged frames for the same VLAN (that is, it cannot be added to a VLAN as both a tagged and an untagged port).

To remove ports from a VLAN, use the command:

```
DELETE VLAN={vlan-name|1..4094} PORT={port-list|ALL}
```

Removing an untagged port from a VLAN will return it to the default VLAN, unless it is a tagged port for another static VLAN. An untagged port can only be deleted from the default VLAN if the port is a tagged port for another static VLAN.

Ports tagged for some VLANs and left in the default VLAN as untagged ports will transmit broadcast traffic for the default VLAN. If this is not required, the unnecessary traffic in the switch can be reduced by deleting those ports from the default VLAN.

To change the tagging status of a port in a VLAN, use the command:

```
SET VLAN={vlan-name|1..4094} PORT={port-list|ALL}  
FRAME=TAGGED
```

To destroy a VLAN, use the command:

```
DESTROY VLAN={vlan-name|2..4094|ALL}
```

VLANs can only be destroyed if no ports belong to them.

To display the VLANs configured on the switch, use the command:

```
SHOW VLAN [= { vlan-name | 1 . . 4094 | ALL } ]
```

To view packet reception and transmission counters for a VLAN, use the command (see [Chapter 11, Interfaces](#)):

```
SHOW INTERFACE=VLANn COUNTER
```

Summary of VLAN tagging rules

When designing a VLAN and adding ports to VLANs, the following rules apply.

1. Each port, except for the mirror port, must belong to at least one static VLAN. By default, a port is an untagged member of the default VLAN.
2. A port can be untagged for zero or one VLAN. A port that is untagged for a VLAN transmits frames destined for that VLAN without a VLAN tag in the Ethernet frame.
3. A port can be tagged for zero or more VLANs. A port that is tagged for a VLAN transmits frames destined for that VLAN with a VLAN tag, including the numerical VLAN Identifier of the VLAN.
4. A port cannot be untagged and tagged for the same VLAN.
5. The mirror port, if there is one, is not a member of any VLAN.

VLAN Interaction with STPs and Trunk Groups

VLANs may have ports in more than one STP, when the ports belong to multiple VLANs. VLANs can belong to multiple STPs.

All the ports in a trunk group must have the same VLAN configuration: they must belong to the same VLANs and have the same tagging status, and can only be operated on as a group.

Generic VLAN Registration Protocol (GVRP)

The GARP application GVRP allows switches in a network to dynamically share VLAN membership information, to reduce the need for statically configuring all VLAN membership changes on all switches in a network. For detailed information see [Chapter 10, Generic Attribute Registration Protocol \(GARP\)](#).

Quality of Service

Quality of Service (QoS) enables you to prioritise traffic and/or limit the bandwidth available to it. The concept of QoS is a departure from the original networking protocols, which treated all traffic on the Internet or within a LAN the same. Without QoS, every different traffic type is equally likely to be dropped if a link becomes oversubscribed. This approach is now inadequate in many networks, because traffic levels have increased and networks transport time-critical applications such as streams of video data. QoS also enables service providers to easily supply different customers with different amounts of bandwidth.

Configuring Quality of Service involves two separate stages:

1. Classifying traffic into flows, according to a wide range of criteria.

Classification is performed by the switch's packet classifier. It is described in [Chapter 39, Generic Packet Classifier](#).

2. Acting on these traffic flows.

Approaches, methods and commands for this are described [Chapter 40, Quality of Service \(QoS\) on Switch Ports](#).

Rapier i Series switches include full QoS functionality, including

- policies, to provide a QoS configuration for a port or ports
- traffic classes, for bandwidth limiting and user prioritisation
- maximum bandwidth limiting on a traffic class
- flow groups within traffic classes, for user prioritisation
- control of the egress scheduling algorithm
- priority relabelling of frames, at Layer 2, by replacing the VLAN tag User Priority field
- class of service relabelling of frames, at Layer 3, by replacing the DSCP (DiffServ Code Point) or the TOS precedence value in the IP header's Type of Service (TOS) field

Table 7-1: The different QoS-type controls available on the switch

Command set	Use for	Do not use for
QoS	Bandwidth limiting of classified traffic flows. Priority queuing of classified traffic flows. Replacing TOS or DSCP byte of IP header. Replacing User Priority in VLAN tag header. Providing a coordinated QoS solution for a port or ports, using the QoS policy model. Configuring a DiffServ domain.	Limiting total bandwidth on a port, unless other QoS controls are also required.

Table 7-1: The different QoS-type controls available on the switch (cont.)

Command set	Use for	Do not use for
Hardware packet filters	Priority queuing of classified traffic flows. Replacing TOS or DSCP byte of IP header. Replacing User Priority in VLAN tag header. Forwarding a flow that is marked to be dropped (for example, because bandwidth allocation is exceeded). Specifying actions for packets which match the ingress and egress ports of a classifier (if set), but do not match the classifier's other parameters. Configuring a simple DiffServ domain.	Bandwidth limiting. Configuring most DiffServ domains.
Layer 3 switch filters	Priority queuing of up to 16 distinctly different types of traffic flow. Replacing TOS or DSCP byte of IP header. Replacing User Priority in VLAN tag header.	QoS for non-IP traffic (e.g. IPX). QoS based on VLANs. Bandwidth limiting.
SET SWITCH PORT	Limiting total ingress and/or egress bandwidth on ports.	Limiting bandwidth of particular types of traffic. Limiting egress bandwidth if priority queuing is also required.

Spanning Tree Protocol (STP)

The Spanning Tree Protocol (STP) makes it possible to automatically disable redundant paths in a network to avoid loops, and enable them when a fault in the network means they are needed to keep traffic flowing. A sequence of LANs and switches may be connected together in an arbitrary physical topology resulting in more than one path between any two switches. If a loop exists, frames transmitted onto the extended LAN would circulate around the loop indefinitely, decreasing the performance of the extended LAN. On the other hand, multiple paths through the extended LAN provide the opportunity for redundancy and backup in the event of a bridge experiencing a fatal error condition.

The spanning tree algorithm ensures that the extended LAN contains no loops and that all LANs are connected by:

- Detecting the presence of loops and automatically computing a logical loop-free portion of the topology, called a *spanning tree*. The topology is dynamically pruned to a spanning tree by declaring the ports on a switch redundant, and placing the ports into a 'Blocking' state.
- Automatically recovering from a switch failure that would partition the extended LAN by reconfiguring the spanning tree to use redundant paths, if available.

Spanning Tree and Rapid Spanning Tree Port States

If STP is running in STANDARD mode, then each port can be in one of five Spanning Tree states, and one of two switch states. If STP is running in RAPID mode, then each port can be in one of four states. The state of a switch port is

taken into account by STP. To be involved in STP negotiations, STP must be enabled on the switch, the port must be enabled on the switch, and enabled for the STP it belongs to.

Multiple Spanning Tree Protocol (MSTP)

Multiple Spanning Tree Protocol (MSTP) was developed to address limitations in the existing protocols, Spanning Tree Protocol (STP) and Rapid Spanning Tree Protocol (RSTP). These limitations apply mainly to networks that use multiple VLANs with topologies employing alternative physical links. MSTP is defined in IEEE Standard 802.1Q 2003.

For detailed information, see [“Multiple Spanning Tree Protocol \(MSTP\)” on page 9-13 of Chapter 9, Spanning Trees.](#)

IGMP Snooping

IGMP (*Internet Group Management Protocol*) is used by IP hosts to report their multicast group memberships to routers and switches. IP hosts join a multicast group to receive broadcast messages directed to the multicast group address. IGMP is an IP-based protocol and uses IP addresses to identify both the multicast groups and the host members. For a VLAN-aware devices, this means multicast group membership is on a per-VLAN basis. If at least one port in the VLAN is a member of a multicast group, by default multicast packets will be flooded onto all ports in the VLAN.

IGMP snooping enables the switch to forward multicast traffic intelligently on the switch. The switch listens to IGMP membership reports, queries and leave 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.

IGMP snooping is performed at Layer 2 on VLAN interfaces automatically. By default, the switch will only forward traffic out those ports with multicast listeners, therefore it will not act as a simple hub and flood all multicast traffic out all ports. IGMP snooping is independent of the IGMP and Layer 3 configuration, so an IP interface does not have to be attached to the VLAN, and IGMP does not have to be enabled or configured.

IGMP snooping is enabled by default. To disable it, use the command:

```
DISABLE IGMP Snooping
```

For more information, see [Chapter 27, IP Multicasting.](#)

Triggers

The Trigger Facility can be used to automatically run specified command scripts when particular triggers are activated. When a trigger is activated by an event, global parameters and parameters specific to the event are passed to the script that is run. For more information, see [Chapter 60, Trigger Facility.](#)

The switch can generate triggers to activate scripts when a fibre uplink port loses or gains coherent light. To create or modify a switch trigger, use the commands:

```
CREATE TRIGGER=trigger-id MODULE=SWITCH
  EVENT={LIGHTOFF|LIGHTON} PORT=port [AFTER=hh:mm]
  [BEFORE=hh:mm] [DATE=date|DAYS=day-list] [NAME=name]
  [REPEAT={YES|NO|ONCE|FOREVER|count}] [SCRIPT=filename...]
  [STATE={ENABLED|DISABLED}] [TEST={YES|NO|ON|OFF}]

SET TRIGGER=trigger-id PORTS={port-list|ALL} [AFTER=hh:mm]
  [BEFORE=hh:mm] [DATE=date|DAYS=day-list] [NAME=name]
  [REPEAT={YES|NO|ONCE|FOREVER|count}]
  [TEST={YES|NO|ON|OFF}]
```

The following sections list the events that may be specified for the EVENT parameter, the parameters that may be specified as *module-specific-parameters*, and the arguments passed to the script activated by the trigger.

Event LINKDOWN
Description The port link specified by the PORT parameter has just gone down.
Parameters The following command parameter(s) must be specified in the CREATE/SET TRIGGER commands:

Parameter	Description
PORT= <i>port</i>	The port on which the event will activate the trigger.

Script Parameters The trigger passes the following parameter(s) to the script:

Argument	Description
%1	The port number of the port which has just gone down.

Event LINKUP
Description The port link specified by the PORT parameter has just come up.
Parameters The following command parameter(s) must be specified in the CREATE/SET TRIGGER commands:

Parameter	Description
PORT= <i>port</i>	The port on which the event will activate the trigger.

Script Parameters The trigger passes the following parameter(s) to the script:

Argument	Description
%1	The port number of the port which has just come up.