

SmartSwitch Multi-layer Frame Classification

Table of Contents

Introduction	2
SmartSwitch Classification Methods	2
Layer 2 Options	2
Layer 3 Options	2
IP Options	2
IP Protocol Type	3
IPX Options	3
Layer 4 IP Options	3
Applications For Classification	4
Layer 2/3/4 VLANs	4
Traffic Containment	4
Traffic Filtering	5
Traffic Security	5
Class of Service	6
Background	6
Prioritization	6
Classification Precedence Rules	8
Conclusion	9

Application
Presentation
Session
Transport
Network
Datalink
Physical

Figure 1: OSI Model

Introduction

Network administrators can take advantage of the next-generation (firmware version 3.10+) SmartSwitch 2000 and 6000's ability to classify incoming frames based on specific Layer 2-4 information for two distinct purposes:

- To create 802.1Q VLANs based on information contained in Layers 2, 3 and 4 of received packets.
- To prioritize traffic throughout the network based on information contained in Layers 2, 3 and 4 of received packets.

Priority information is indicated throughout the network using the priority indicator embedded in an 802.1Q frame tag.

The SmartSwitch's ability to classify traffic enables four key network tasks:

Containment—Scoping or containing of frames within a specific boundary (normally referred to as a VLAN)

Filtering—Preventing protocols, applications, and/or specific users from accessing the network

Security—Securing certain resources within the network, such as specific addresses

Class of Service/Quality of Service—Associating a transmit priority to each frame based on its classification

SmartSwitch Classification Methods

The SmartSwitch supports a maximum of 1024 classification rules that can be divided between VLANs and/or priority rules based on the customer's needs. Regardless of whether a frame classification rule is to be used for assigning VLAN status or priority, the method of creation is the same. Classification is based upon a frame's Datalink, Network and Transport Layer information (Layers 2, 3 and 4, respectively, in the OSI model, shown in Figure 1). Although the SmartSwitch makes classification decisions based on Layers 2-4, its forwarding mechanism is still that of a Layer 2 store-and-forward device (a bridge or switch). This provides it with some of a router's intelligence without the complex configuration and increased cost typically associated with these devices. The specific options for each level of classification supported by the SmartSwitch are detailed below.

Layer 2 Options

At Layer 2, an administrator can classify frames into VLANs and/or priority levels based on the specific protocol type of each frame. The protocol type can be found in one of the following frame locations: MAC Address—Any 6 byte source or destination MAC Address can be specified. Ethernet Type II—The two-byte "Ether Type" field located after the source address. IEEE 802.3 (with 802.2 LLC header)—The one-byte DSAP and SSAP fields located in the frame's LLC header. The DSAP and SSAP values are one byte each; however, a combined two-byte value must be entered (these values must match, DSAP=E0 SSAP=E0). IEEE 802.3 (LLC with SNAP header)—The two-byte "Ether Type" field is located in the SNAP header. Both local and remote management provide a number of pre-defined Layer 2 frame classification options (IP, IPX, DECNET, AppleTalk etc.). However, an administrator may define any valid two-byte protocol Type value (0000-FFFF) for use as a classification rule.

Layer 3 Options

At Layer 3, an administrator can classify frames into VLANs and/or priority levels based on specific information contained within the Layer 3 header of an IP or IPX frame, as listed below:

- **IP Options**

- **IP Type of Service (TOS)**

- ToS is a one-byte field contained in the IP header of a frame. The TOS field can be used by applications to indicate priority and QoS parameters for each frame. Until recently ToS was not widely used. The IETF's Differentiated Services (Diffserv), which are gaining popularity, make use of this field. A user may define VLANs and/or priority mappings based on an exact match of the eight bits contained within the IP TOS field.

0	Hello or SAP
1	RIP
2	Echo Packet
3	Error Packet
4	Netware 386 or SAP
5	Sequenced Packet Protocol
17	Netware 286
16-31	Experimental Protocols

Table 1: IPX Packet Type Options

• **IP Protocol Type**

The protocol type of an IP frame is indicated by a two-byte field in the IP header. This field indicates the specific protocol being used. Predefined options include TCP, UDP, ICMP, and IGMP. Users can define any valid two-byte value.

Source, Destination, and Bilateral IP Address

An administrator can classify frames into VLANs and/or priorities based on the specific IP address information contained within their IP header. Each such classification rule must have a “Mask” entry, in addition to an IP address entry, which lets the switch determine whether the classification is based on a specific IP address, IP subnet, or range within an IP subnet. For example, an administrator could define a classification rule where all frames from the 134.141.28.x network are assigned to the Red VLAN by setting the IP address of 134.141.28.0 with the mask of 255.255.255.0.

• **IPX Options**

Note: Layer 3 IPX classification rules are supported for all IPX frame types.

IPX Class of Service

IPX Class of Service is a one-byte field, located within the IPX header of a frame, which is used for transmission control (hop count) by IPX routers.

IPX Packet Type

IPX Packet Type as shown in Table 1 is one of the following one-byte options, as defined by Novell:

IPX Network Number

This user-defined, four-byte value represents the IPX network number. Choices include Source, Destination, or Bilateral Network Number.

IPX Socket Number

This user-defined, two-byte value describes an IPX socket number. These values are used by higher-layer protocols to target specific applications running among hosts. Administrators may classify frames based on Source, Destination, or Bilateral Socket Numbers.

• **Layer 4 IP Options**

At Layer 4, an administrator can classify IP frames into VLANs and/or priorities based on the specific Layer 4 TCP or UDP port numbers contained in the IP frame header. Administrators can configure a classification rule based on the following:

UDP Port

Options include UDP Source Port, UDP Destination Port, or UDP Bilateral Port. Each of these options are two-byte values.

TCP Port

Options include TCP Source Port, TCP Destination Port, or TCP Bilateral Port. Each of these options are two-byte values. Cabletron’s local and remote management for the SmartSwitch provides some common TCP and UDP port values, including HTTP, FTP and BOOTP. Network administrators can enter any valid two-byte value (as defined in RFC 1700) for the TCP and UDP port numbers.

Note: SmartSwitches do not support Layer 4 classification for the following IP frames:

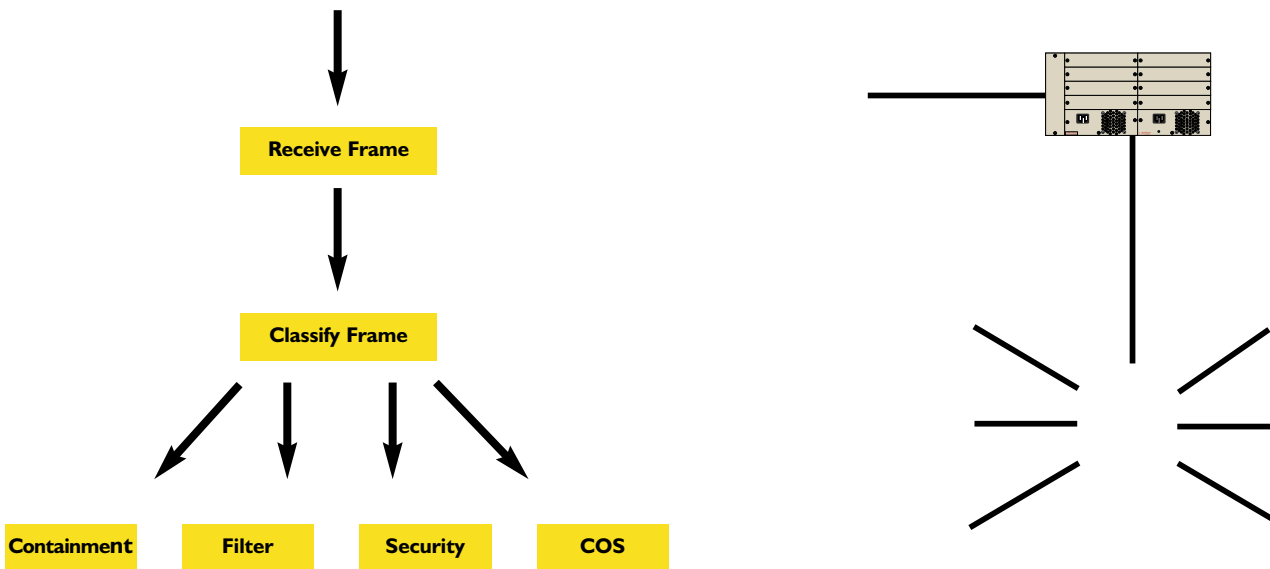
1. Frames where the options field is greater than 4 bytes.
2. Frames that have been fragmented, as the Layer 4 information is not present in these frames.
3. If a SmartSwitch has an FDDI HSIM installed, Layer 4 classification will not be supported for any frames larger than 1500 bytes. Frames larger than 1500 bytes are fragmented internally in the switch.

Applications For Classification

As stated earlier, frame classification can be used for two distinct purposes: to create 802.1Q VLANs based on Layer 2-4 information; and to enable Class of Service based on the priority indicator embedded in an 802.1Q-tagged frame. Conceptually, you should envision frames as undergoing two separate classification processes within the switch—once for VLAN assignment and once for priority assignment—as shown in Figure 2.

Layer 2/3/4 VLANs

Earlier SmartSwitch firmware classified frames into 802.1Q VLANs based on the VLAN assignment of the receiving port, regardless of protocol type. With the latest firmware, network administrators can now intelligently classify frames into a specific VLAN based on their Layer 2, 3, and 4 information. This gives administrators the ability to perform Containment. Each of these is discussed in more detail below.




Traffic Containment

Network Administrators can logically group users of a given protocol, subnet, or application together and control the flow of their traffic on the network.

Figure 3 shows a configuration where the network administrator wants to separate end-user traffic into VLANs based on the assigned IP subnet of each department. He or she can easily accomplish this by setting up two Layer 3 classification rules based on the IP subnet range of the respective departments:

- **Rule 1**—Engineering, which uses the 134.141.28.x subnet, will be assigned to the Red VLAN.
- **Rule 2**—Sales, which uses the 134.141.29.x subnet, will be assigned to the Blue VLAN.

Based on these two Layer 3 classification rules, traffic from the Engineering VLAN will be isolated from that of the Sales VLAN, and vice versa. Also, remember that since these rules are based on Layer 3 information, an Engineering user could enter the network from a connection in the Sales department, and that user would still be contained within the proper VLAN.



Layer 3 Rule-IP Protocol type 520(OSPF) = Null VLAN
 Layer 4 Rule-UDP port 89(RIP) = Null VLAN
 Result: All RIP and OSPF frames are filtered from end-users

Figure 4: Traffic Filtering Configuration

- **Rule 1 (Layer 3)**—Any frame received on SmartSwitch 1 with an IP Protocol Type of 89 (OSPF), will be classified into the Null VLAN. This step would be duplicated on SmartSwitch 2.
- **Rule 2 (Layer 4)**—Any frame received on SmartSwitch 1 with a Bilateral UDP port number of 520 (RIP), will be classified into the Null VLAN. This step would be duplicated on SmartSwitch 2. Based on this configuration, as long as the Null VLAN is not configured to exit the switch, all RIP and OSPF frames will be filtered from the end users.




Figure 5: Traffic Security Configuration

Traffic Filtering

Network Administrators can filter specific unwanted traffic, such as broadcast routing protocols, traffic from specific IP addresses, or even application traffic such as HTTP or SMTP. To perform a filtering operation, the SmartSwitch must be configured so that a particular VLAN is not on the forwarding (egress) list of any port on the switch. This VLAN can be referred to as a “Null VLAN.” Any traffic to be filtered should be classified to the Null VLAN; since it is not configured on any of the switch ports’ egress lists, the traffic will be discarded. Figure 4 shows a common configuration in which a routed backbone is using both RIP and OSPF as its routing protocols. The network administrator does not want the multicast OSPF and broadcast RIP frames being propagated to the end stations. The network is designed so that only end users are attached to the SmartSwitches. To implement filtering in this scenario, a Layer 3 rule and a Layer 4 rule will be needed.

Traffic Security

The SmartSwitch can provide many levels of security using Layer 2/3/4 classification. Configuring security is similar to configuring Traffic Filtering, described previously. Figure 5 illustrates a network configuration that includes a router and a SmartSwitch 6000. In this configuration, end-users connect to ports 1-24 of the SmartSwitch 6000. Some of these users have been “hacking” into the router and altering its configuration. A simple configuration rule can be put in place to prevent these types of occurrences. Since the end users should never need to communicate directly to the router using the router’s IP address, a Layer 3 IP classification rule will be used.

- **Rule 2**—Any frames received by the switch with a destination IP address of the router (192.168.1.2) will be classified to the Null VLAN. As described in the Traffic Filtering example above, the Null VLAN should be configured so that it does not exit the switch.

Priority Indicator	Transmit Queue
7	3
6	3
5	2
4	2
3	1
2	0
1	0
0	1

Priority 7=Highest

TX Queue 3 is serviced first

Table 2: Default Priority to Transmit Queue Mapping for Smartswitch 2000/600

The end result is that any frames from a user trying to hack into the router will be discarded before they ever reach the router.

Class of Service

• **Background**

In 1998, Traffic Class Expediting functionality was added to the IEEE 802.1D standard. During its development stages, this technology was referred to by the IEEE as 802.1p. There are two parts to IEEE 802.1p: Dynamic Multicast Filtering, and Traffic Classification (which is the focus of this discussion). Traffic Classification defines a method of prioritizing packets based on a Layer 2 tag which is inserted into the frame by an end station, a switch, or a router. The standard defines eight different priority levels. Each priority level is mapped into a specific transmit queue by the switch or router (the standard states that devices must support at least two transmit queues per port). The insertion of the priority value (0-7) allows all tag-aware devices in the network to make intelligent forwarding decisions based on their own level of support for prioritization. The next-generation SmartSwitches support four transmit queues (0-3) per port. These queues can be serviced based on a strict method (meaning that all frames in Queue 3 will be transmitted before the frames in Queue 0), or based on a fair weighted method (which allows the network administrator to give a certain percentage or weight to each queue, preventing a lower priority queue from being starved). Traffic Classification allows an individual switch port, user or application to have a higher priority than another switch port, user or application. This scheme provides Class of Service functionality for 802.1D devices.

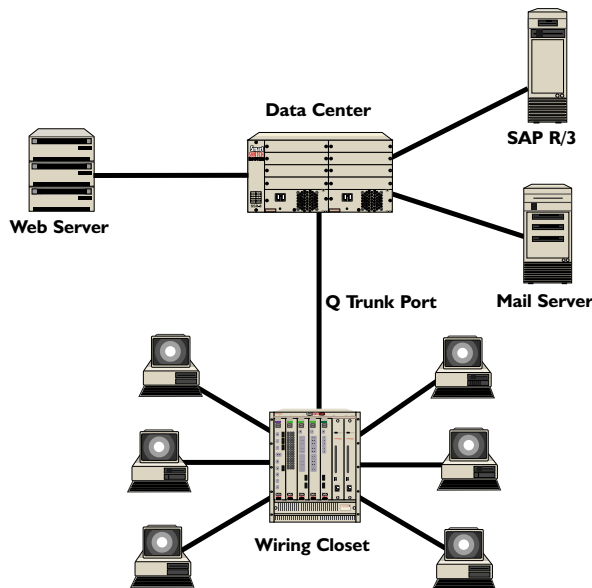


Figure 6: Prioritization for Class of Service

Prioritization

The Priority feature on the SmartSwitch takes this concept a step further to allow for better defined Class of Service configurations. Network administrators can now, on a port-by-port basis, classify a frame based on its Layer 2-4 information with higher or lower priority status than other frames received on that port. A good use of this feature would be a configuration where a network administrator wants to assign priority to three network applications—SAP R/3, web traffic, and e-mail, in that order—as shown in Figure 6.

Note: The configuration for this example is the same as described in the Containment section of this document.

There are two main steps required to accomplish this: configuring the classification rules, and configuring the Priority-to-Transmit Queue mapping for the switch.

Classification Rules

- **Rule 1 (SAP R/3)**—All frames to or from the IP address of the SAP R/3 server will be tagged with a priority indicator of 7 (highest).

Note: An IP address classification was selected for Rule 1 because it has been observed that SAP R/3 dynamically negotiates the TCP/UDP port used, and the port number selections vary from session to session. If this were not the case, a Layer 4 UDP classification could be used.

- **Rule 2 (web)**—All frames with a UDP port number of 80 (HTTP) will be tagged with a priority indicator of 5.
- **Rule 3 (e-mail)**—All frames with a UDP port number of 25 (SMTP) will be tagged with a priority indicator of 3.

Priority Queuing Configuration

Based on the default SmartSwitch Priority-to-Transmit queue mapping, the values selected above will work out so that each frame classification type will be mapped to the desired transmit queue. This means that no user configuration of the priority-transmit queue mapping would be required. It should be noted, however, that this mapping can be configured by the administrator.

Table 3 shows the default SmartSwitch 2000/6000 Priority-to-Queue mappings.

Result


With the classification rules for the network shown in Figure 6, the SmartSwitches would provide advanced Class of Service functionality for individual network applications, but still forward at Layer 3. Table 4 summarizes the resultant QoS configuration.

Application	Classification Type	Desired Priority	802.1 Tag Value	SS6000 Queue Mapping
SAP R/3	Bilateral IP	High	7	3
Web	UDP Port number	Medium	5	2
e-mail	UDP Port number	Low	3	1

Table 3: Sample QoS Configuration Summary

VLANs versus Priority

As stated earlier, classification rules for VLANs and priority are treated separately. The configuration shown in Figure 7 is the combination of Figure 3, which shows VLANs being defined by subnet, and Figure 6 which shows applications being prioritized according to Layer 3/4 information. For the sake of this discussion, let's say that the three network applications shown in Figure 6 are used by both the Engineering and Sales departments. Even though the traffic from Engineering is in a different VLAN than the Sales traffic, the priority rules apply to both VLANs. The key point to remember is that VLAN classification is treated separately from priority classification. The Engineering group does not have separate priority rules assigned to its VLAN.



Application Priority for Engineering and Sales
 SAP R/3 = High Web = Medium Email = Low

Figure 7: VLANs by Subnet and Applications Prioritized According to Layer 3/4 Information

Classification Precedence Rules

When there are multiple classification rules assigned to a switch, the SmartSwitch must determine which rule takes precedence. The order of precedence is predefined in the SmartSwitch, and is currently not configurable. The SmartSwitch will make a classification based on one of the classification options. Table 4 details the order of precedence.

Classification Type (IP)	Precedence Level
802.1Q frame tag received	1
Source MAC Address	2
Destination MAC Address	3
Source IP Address exact match	4
Source IP best match (Subnet)	5
Destination IP exact match	6
Destination IP best match (Subnet)	7
UDP/TCP source port	8
UDP/TCP destination port	9
IP TOS	10
IP type	11
Protocol type (Ether type or DSAP/SSAP)	12
Receive port	13
Classification Type (IPX)	Precedence Level
802.1Q frame tag received	1
Source MAC Address	2
Destination MAC Address	3
Source IPX network number	4
Destination IPX network number	5
IPX Source socket	6
IPX Destination socket	7
IPX Class of service	8
IPX Type	9
Protocol type (Ether type or DSAP/SSAP)	10
Receive port	11
Classification Type (non IP/IPX)	Precedence Level
802.1Q frame tag received	1
Source MAC Address	2
Destination MAC Address	3
Protocol type (Ether type or DSAP/SSAP)	4
Receive port	5

- 1A = Highest Precedence 4 = Lowest Precedence
- Exact Match indicates a match of an explicitly defined address.
 - Best Match indicates a match of an entire subnet, or range within a subnet.

Table 4: Precedence Table

The precedence table concept is illustrated in the following scenarios:

Scenario 1

A network administrator has defined two classifications:

- All frames with a UDP port number of 55 are assigned to the Red VLAN.
- All frames sourced from the 134.141.28.x subnet are assigned to the Blue VLAN.

If a frame is received with a source address of 134.141.28.99 and contains a UDP port number of 55, the frame will be assigned to the Blue VLAN because, as shown in Table 4, a Layer 3 IP Address rule takes precedence over a Layer 4 rule.

Scenario 2

A network administrator defines two classifications:

- All frames with an IP TOS value of AA are assigned a priority of 7.
- All frames with a TCP port number of 80 are assigned a priority of 3.

If a frame is received with a TOS value of AA, and a TCP port number of 80, the frame will be assigned a priority of 3, because as shown in Table 4, TCP port number classifications take precedence over IP TOS classifications. It is very important that network administrators have a comprehensive understanding of the precedence concept before configuring the SmartSwitch, as these rules can significantly impact the operation of the network.

Conclusion

The standards-based multilayer frame classification abilities of the SmartSwitch products provide network administrators with a powerful set of utilities that allow a more intelligent confirmation and management of today's switched networks.

Corporate

35 Industrial Way
 Rochester, NH 03867
 U.S.A.
 (603) 332-9400

Europe/Middle East/Africa

Network House
 Newbury Business Park
 London Road, Newbury
 Berkshire, England RG13 2PZ
 44-1635-580000

Asia Pacific

85 Science Park Drive
 #03-01/04
 The Cavendish
 Singapore 118259
 65-775-5355

Unit 8, Allambie Grove Estate
 25 Frenchs Forest NSW 2086
 Sydney, Australia
 61-29950-5900

Latin America

Periferico Sur No. 3642
 Piso 6
 Colonia Jardines del Pedregal
 Mexico City DF 01900
 Mexico
 525-490-3400
 Av Jurubatuba, 73-3º andar
 Brooklin-São Paulo
 04583-100-Brazil
 55-11-5508-4600

Copyright © 2000 Cabletron Systems, Inc. All Rights Reserved. SmartSwitch Router is a trademark or registered trademark of Cabletron Systems, Inc. All other products or services mentioned are identified by the trademarks or service marks of their respective companies or organizations. NOTE: Cabletron Systems, Inc. reserves the right to change specifications without notice. Please contact your representative to confirm current specifications.

Lit.# 9011639-1 5/00