Marking - DiffServ

CoS, IPP, & DSCP

Differentiated Services

- The basic idea is the Per-Hop Behavior which is controlled by a QoS-tag within the IP packet header
- Stateless, connection-less QoS solution
  - Perfect scalability
  - Can be implemented end-to-end
  - Most widespread used solution

The IP Header

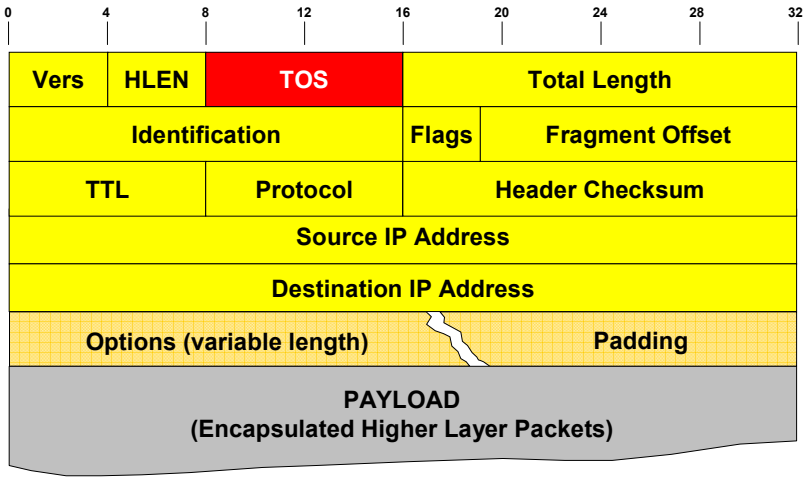

Original ToS

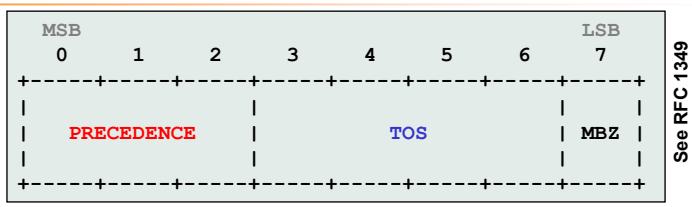

- Bits 0,1,2 – IP Precedence (IPP)
  - May carry the same information as provided in the IEEE 802.1p field
  - This allows a *mapping* between layer-2 and layer-3 and both switches and routers can process the packet as desired

Original ToS (cont.)

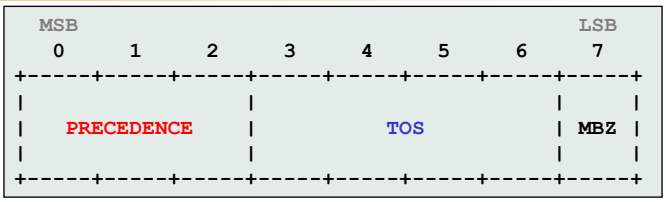

- Bit 3-6 – Type of Service (aka TOS)
  - Bit 3 specifies whether normal (0) or low (1) delay is desired
  - Bit 4 specifies normal (0) or high (1) throughput is desired
  - Bit 5 specifies whether normal (0) or high (1) reliability is desired
  - Bit 6 specifies whether normal (0) or low (1) monetary cost is desired
- Bit 7 – Must be Zero

ToS Usage (historical)

- Only 1 of these 4 bits can be turned on
- If all 4 bits are 0 it implies normal service
- Example from RFC 1340/1349
  - Telnet: 0x10 – minimize delay
  - FTP Control: 0x10
  - FTP Data: 0x80 – maximize throughput
  - SNMP and IGP: 0x04 – maximize reliability
  - NNTP: 0x02 – minimize monetary cost
- Practically not used today
  - Would require that routing protocols that differentiate routes accordingly (originally OSPF and IS-IS were designed for that!)

DSCP

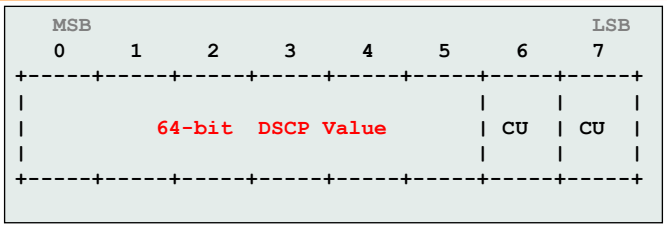

Practically only 5 bits are used:

Only zero if bits 3,4 should be interpreted as drop precedence

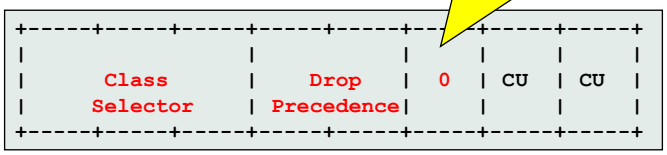

Since RFC 2474 (1998) the ToS field had been redefined and renamed to DiffServ Code Point (DSCP) or DS Field. As specified in RFC 2474 only six bits are used.

DSCP Details

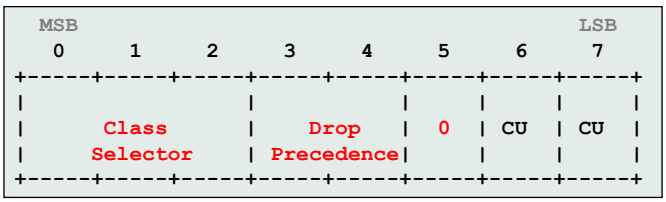

- Bits 0,1,2 – Class Selector
  - Again carries the 0-7 IP precedence value
  - However there is a new definition of these values…
- Bits 3,4 – Drop Precedence
  - Carry the drop precedence which allows a further differentiation of traffic within the same class level
  - Note that this requires bit 5 to be zero
- Bits 6,7 – Currently Unused (or ECN…)

Class Selector Values

Level	Meaning
7	unchanged (link layer and routing protocol keep alive)
6	unchanged (used for IP routing protocols)
5	Express Forwarding (EF)
4	AF Class 4
3	AF Class 3
2	AF Class 2
1	AF Class 1
0	Best effort

- The legacy IP Precedence values (0-7) can be directly mapped into the three Class Selector bits (0,1,2) *with the three LSBs (3,4,5) set to zero*
- This results in the seven CSx values
  - CS0 = DSCP 00 = 000000
  - ...
  - CS7 = DSCP 56 = 111000

14 Recommended Code Points

- Expedited Forwarding (EF)
  - DSCP 46 = 101 110 binary
  - For low delay, low loss, and low jitter
  - Defined in RFC 3246
- Assured Forwarding (AF)
  - 12 codepoints: 4 classes and 3 drop precedences each
  - Defined in RFC 2597
- Best Effort (BE)
  - 000000 binary

Assured Forwarding (AF)

- Guarantees a certain bandwidth to a traffic class
  - If the traffic exceeds the committed bandwidth the drop probability is raised according to the specified drop precedence
- There are 12 different AF behavior code points
  - Consisting of 4 classes (AF1y to AF4y)
  - And 3 drop probabilities (AFx1 to AFx3) for each class (low/med/hi)

Drop:	Class 1			Class 2			Class 3			Class 4		
Low	AF11	10	001010	AF21	18	010010	AF31	26	011010	AF41	34	100010
Medium	AF12	12	001100	AF22	20	010100	AF32	28	011100	AF42	36	100100
High	AF13	14	001110	AF23	22	010110	AF33	30	011110	AF43	38	100110

decimal | binary

Notations

- Note that there are two common notations:
  - Either "DSCP zz" with zz a decimal number
  - Or referring to a *well-defined* code point, e. g. BE, CS0..7, AFxy, or EF
- Example:
  - DSCP 46 means EF
  - DSCP 22 means AF23

DSCP Values Overview

Code Point Name	DSCP		Whole IP TOS byte		
	hex	dec	binary	hex	dec
EF	0x2e	46	10111000	0xb8	184
AF41	0x22	34	10001000	0x88	136
AF42	0x24	36	10010000	0x90	144
AF43	0x26	38	10011000	0x98	152
AF31	0x1a	26	01101000	0x68	104
AF32	0x1c	28	01110000	0x70	112
AF33	0x1e	30	01111000	0x78	120
AF21	0x12	18	01001000	0x48	72
AF22	0x14	20	01010000	0x50	80
AF23	0x16	22	01011000	0x18	24
AF11	0x0a	10	00101000	0x28	40
AF12	0x0c	12	00110000	0x30	48
AF13	0x0e	14	00111000	0x38	56
CS7	0x38	56	11100000	0xe0	224
CS6	0x30	48	11000000	0xc0	192
CS5	0x28	40	10100000	0xa0	160
CS4	0x20	32	10000000	0x80	128
CS3	0x18	24	01100000	0x60	96
CS2	0x10	16	01000000	0x40	64
CS1	0x08	8	00100000	0x20	32
CS0 = BE	0x00	0	00000000	0x00	0

Ethernet QoS Marking

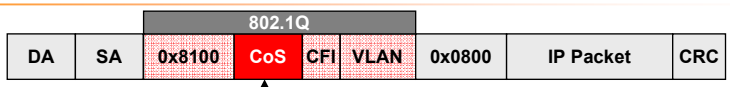

	Priority	Traffic type
highest	7	Network management
	6	Voice
	5	Video
	4	Controlled load
	3	Excellent effort
default	0	Best effort
	2	Spare (undefined)
lowest	1	Background

- 802.1p aka "Class of Service" (CoS)
  - aka "User Priority" (UP)
- Note that CoS 0 is *better* than CoS 1 and 2

Marking and Classification

- We need both L2 and L3 marking
  - L2 markings are examined by switches
  - L3 markings are preserved through WAN

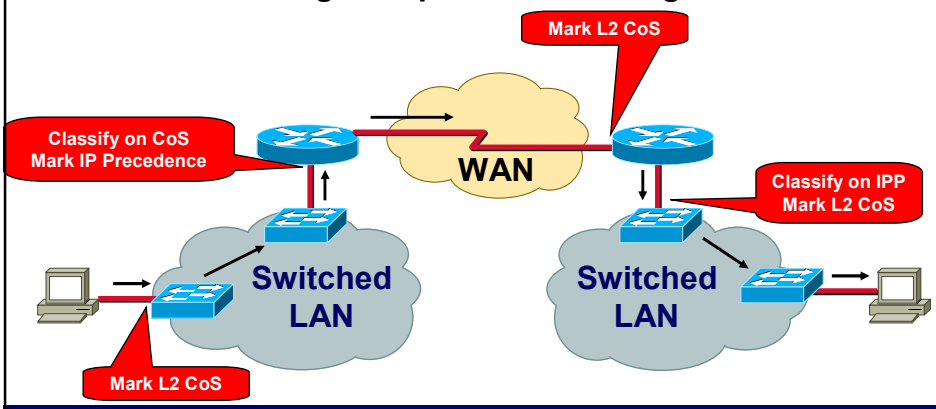

Modular QoS Commandline (MQC)

Three intuitive steps

1) Identify interesting traffic

```
class-map match-all MATCH_WWW
  match access-group ACL_WWW_TRAFFIC
class-map match-any MATCH_DATA
  match access-group ACL_FTP_TRAFFIC
  match protocol icmp
```

2) Do something with that traffic

```
policy-map MARK_DATA_APPS
  class MATCH_WWW
   set dscp default
  class MATCH_FTP
   set dscp af11
```

3) Apply this rule on an interface

```
interface fa0/0
 service-policy input MARK_DATA_APPS
```