

Optimisation del ancho de banda

(Intro to Linux Traffic Control – QoS / Shaping)



Christian Benvenuti

christian.benvenuti@libero.it

La Paz, Bolivia, 11-15/5/2009



Agenda

- Survey
- Traffic Control Architecture
 - Queueing Disciplines
 - Classes
 - Filters
 - Actions
- Kernel Configuration
- *tc* command syntax
- Examples/Exercises

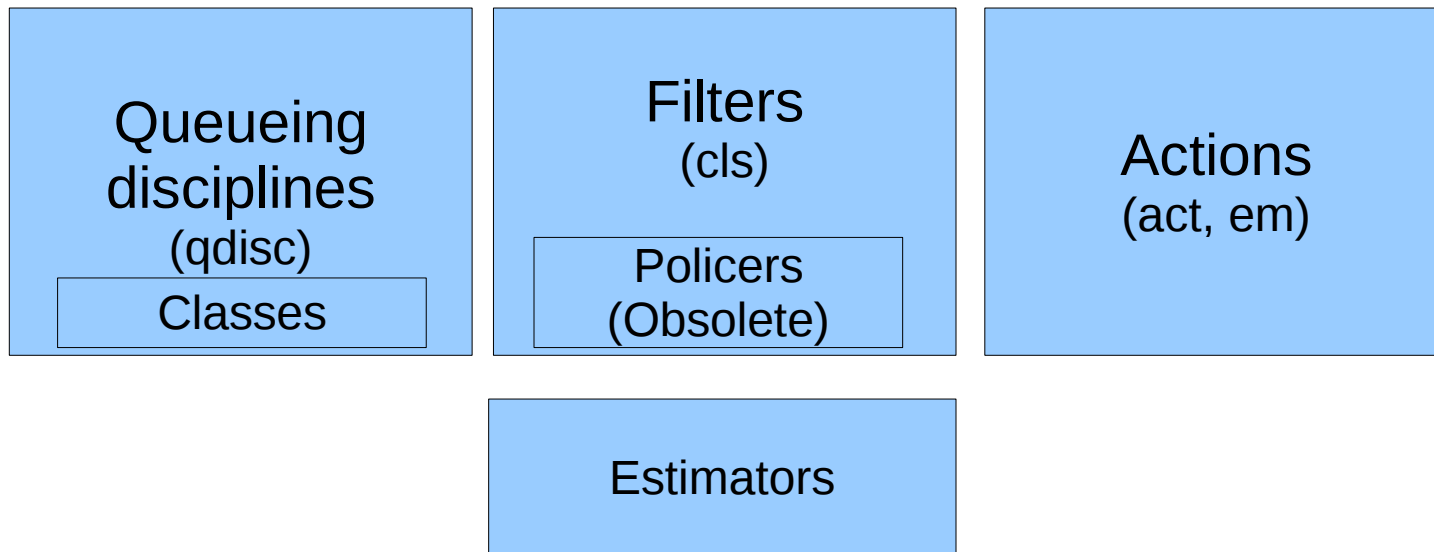
Survey

- If you already use Linux Traffic Control and you are happy with it:
 - What do you like about it ?
 - What is the most useful feature?
 - How long have you been using it?
 - Where/how did you learn how to use it?
- If you are not using it?
 - Why?
 - What are you using to configure QoS/Shaping on your network/s?


Traffic Control ...

- ... is very, very flexible ... therefore I won't be able to introduce all its functionalities today ...

Traffic Control Architecture



Queueing Discipline

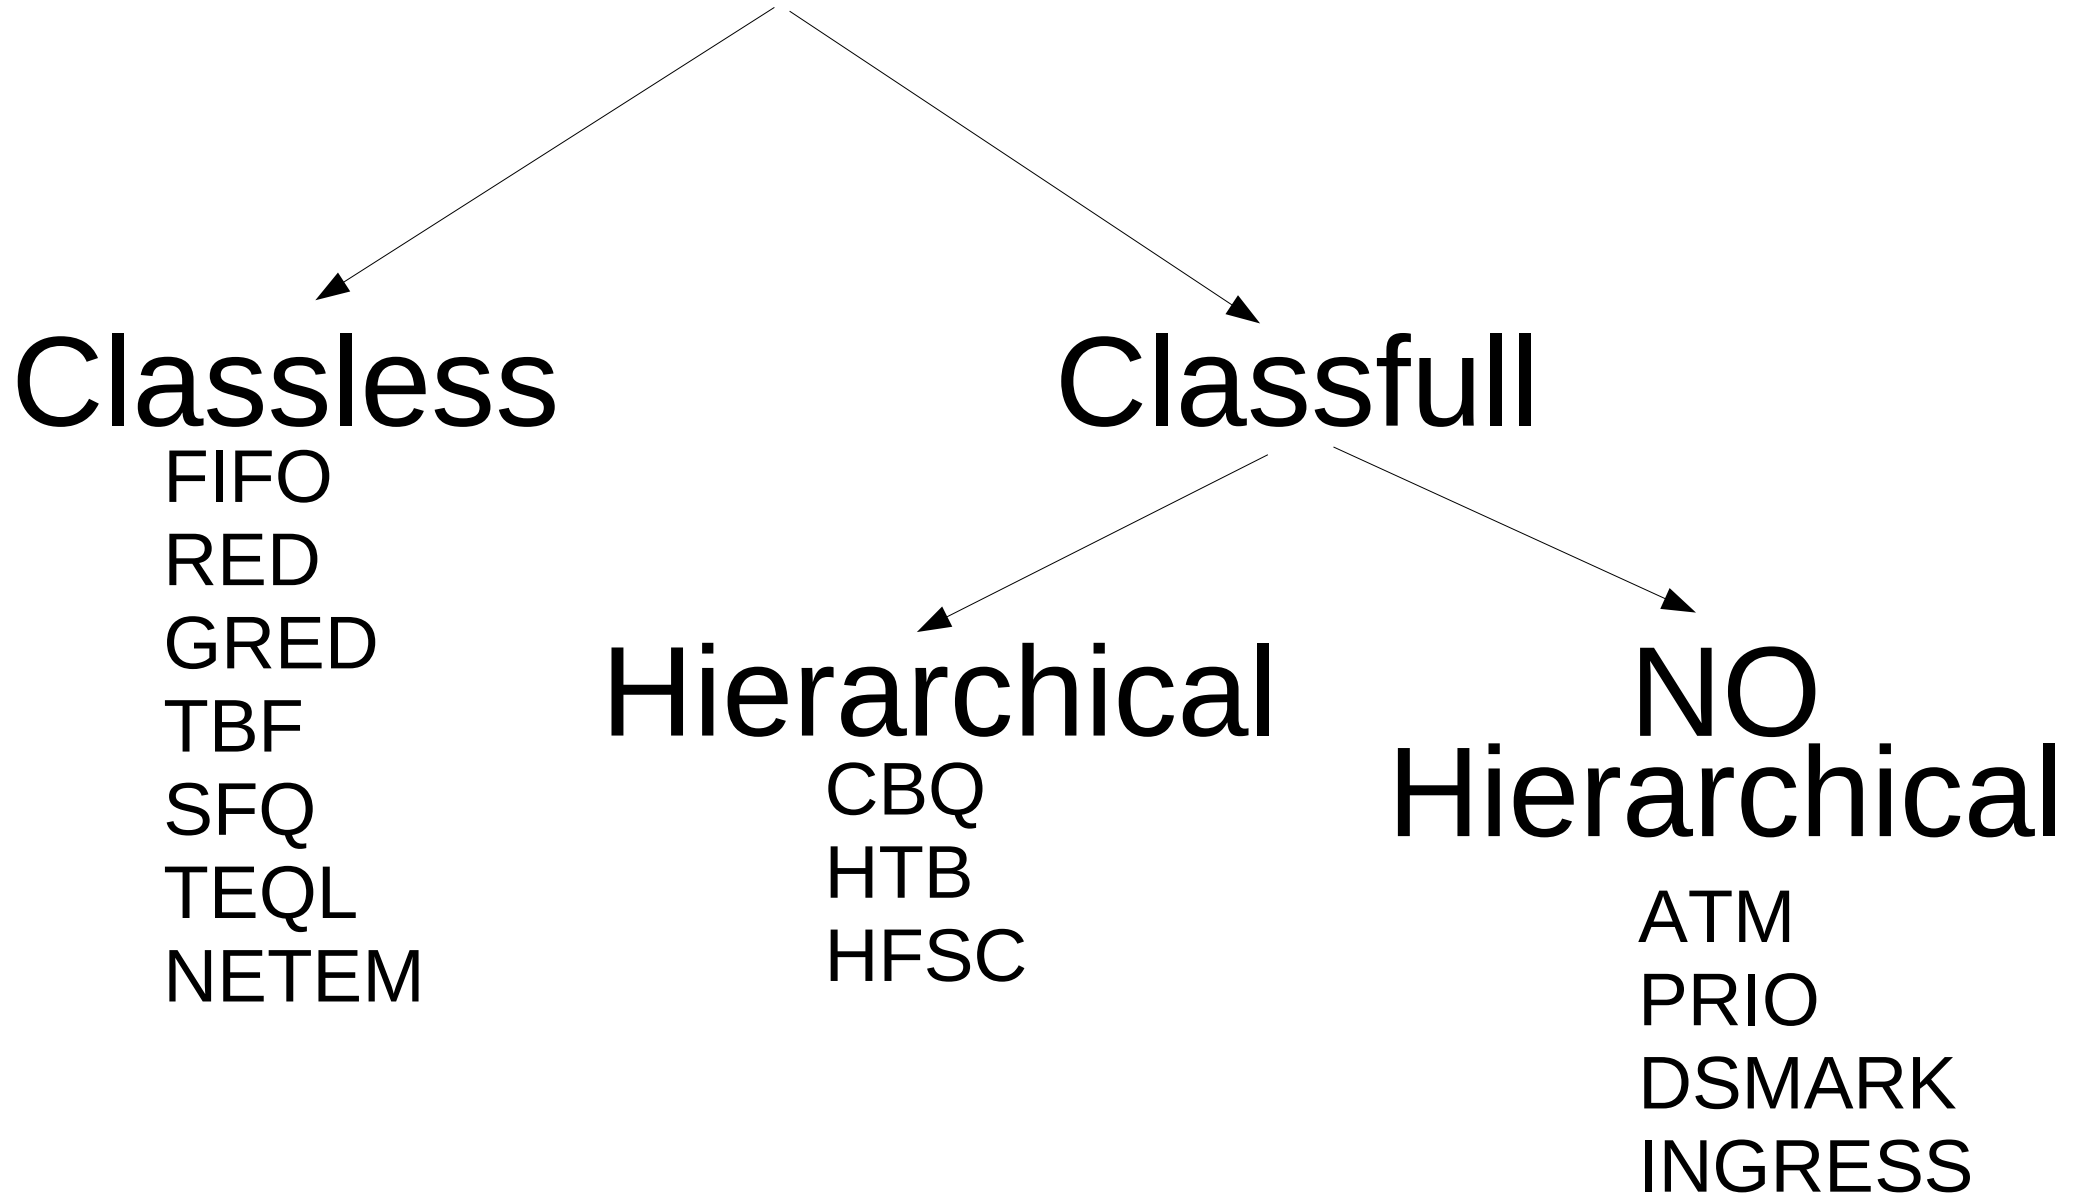


What is the simplest and most common queueing discipline?


Queueing Disciplines

- Class Based Queueing (CBQ)
- Hierarchical Token Bucket (HTB)
- Hierarchical Fair Service Curve (HFSC)
- Multi Band Priority Queueing (PRIO)
- Random Early Detection (RED)
- Stochastic Fairness Queueing (SFQ)
- True Link Equalizer (TEQL)
- Token Bucket Filter (TBF)
- Generic Random Early Detection (GRED)
- Differentiated Services Marker (DSMARK)
- Network Emulator (NETEM)
- Ingress Qdisc (INGRESS)


Queueing Disciplines




Queueing Disciplines & Classes



CLASSLESS

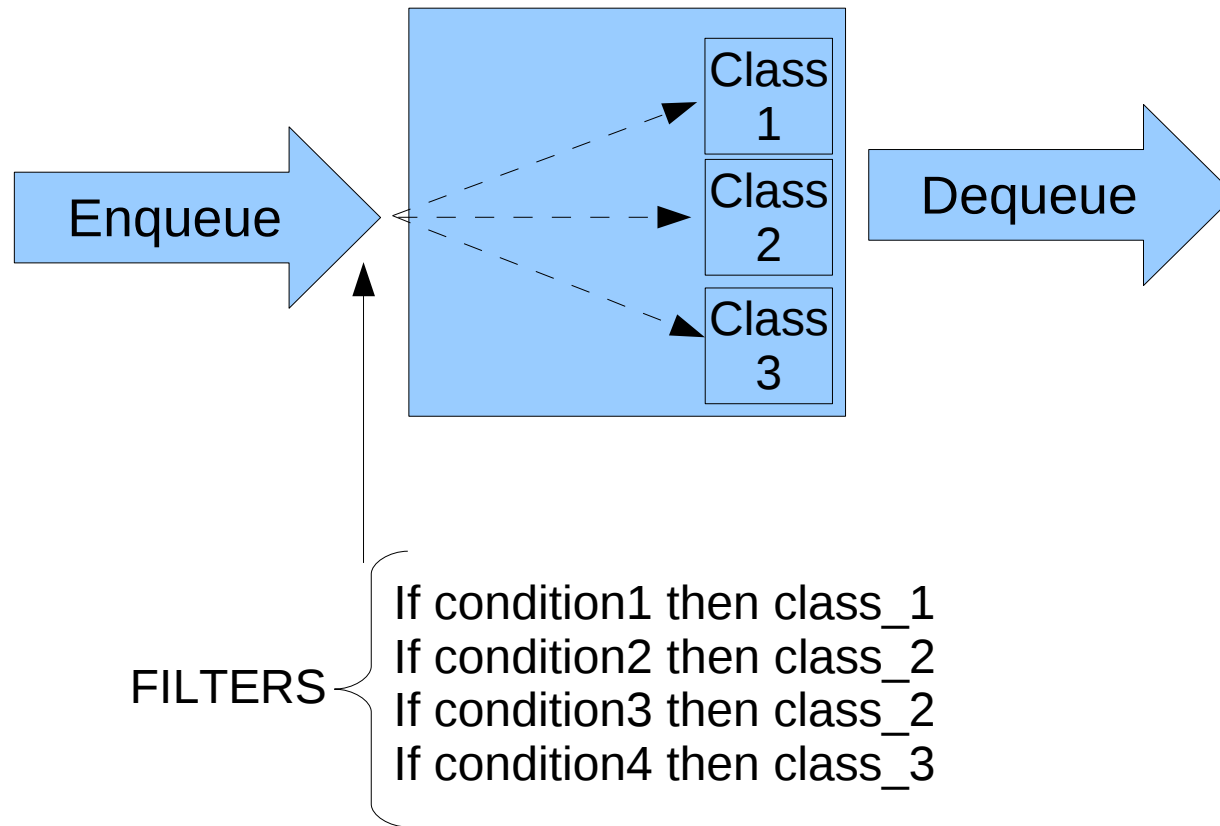


CLASSFULL



**CLASSFULL
HIERARCHICAL**

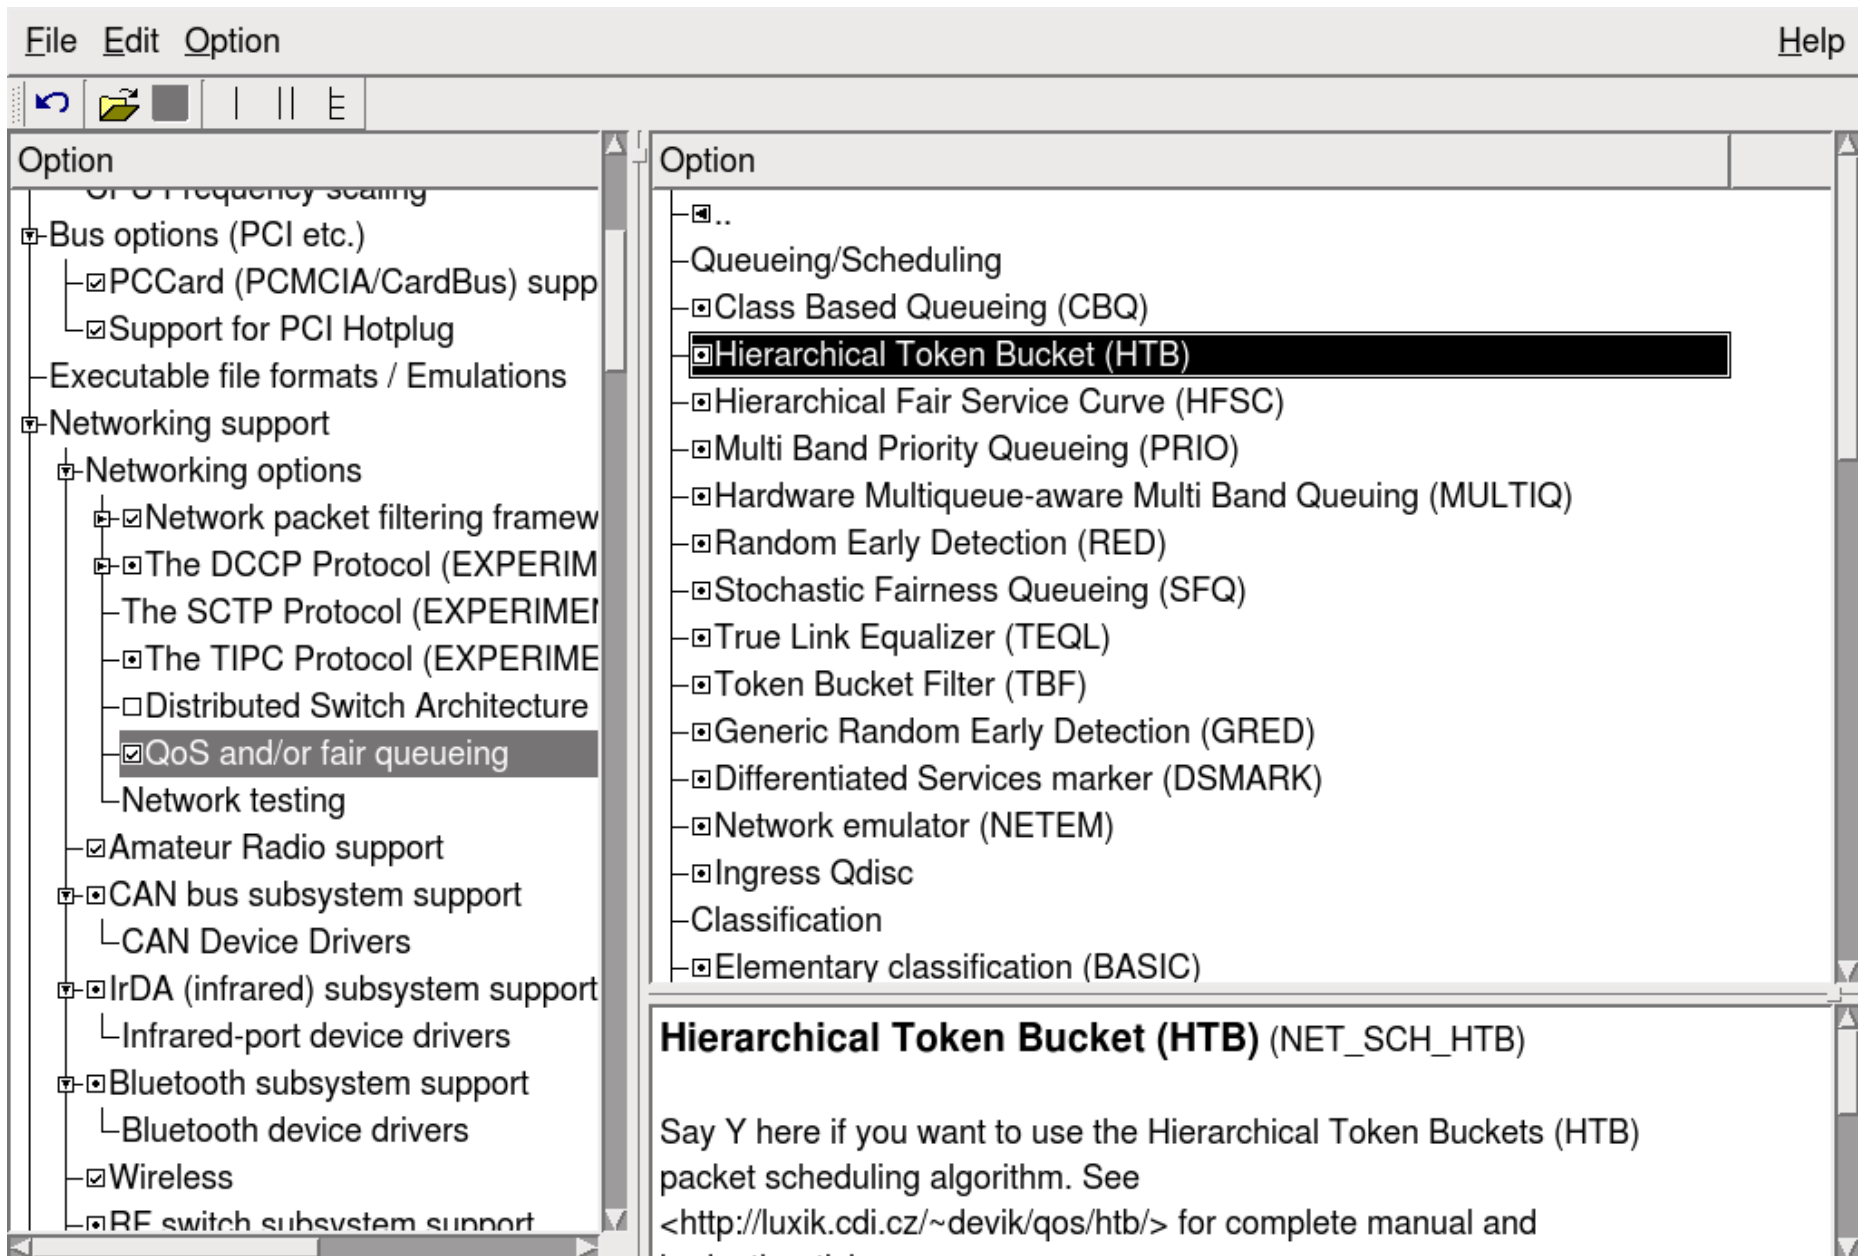
Filters



Filters

- Elementary classification (BASIC)
- Traffic-Control Index (TCINDEX)
- Routing Decision (ROUTE)
- Netfilter mark (FW)
- Universal 32bit comparison w/hashing (u32)
- IPv4 Resource Reservation Protocol (RSVP)
- IPv6 Resource Reservation Protocol (RSVP6)
- Flow Classifier (FLOW)

Kernel Configuration



tc command syntax

(*help* option)

\$ tc

Usage: tc [OPTIONS] OBJECT { COMMAND | help }

tc [-force] -batch file

where OBJECT := { qdisc | class | filter | action | monitor }

OPTIONS := { -s[tatistics] | -d[etails] | -r[aw] | -p[retty] | -b[atch] [file] }

tc command syntax

\$ tc qdisc help

Usage: tc qdisc [add | del | replace | change | show] dev STRING
[handle QHANDLE] [root | ingress | parent CLASSID]
[estimator INTERVAL TIME_CONSTANT]
[[QDISC_KIND] [help | OPTIONS]]

tc qdisc show [dev STRING] [ingress]

Where:

QDISC_KIND := { [p|b]fifo | tbf | prio | cbq | red | etc. }

OPTIONS := ... try tc qdisc add <desired QDISC_KIND> help

tc command syntax

tc qdisc add htb help

What is "help"?

Usage: ... qdisc add ... htb [default N] [r2q N]

default minor id of class to which unclassified packets are sent {0}

r2q DRR quanta are computed as rate in Bps/r2q {10}

debug string of 16 numbers each 0-3 {0}

... class add ... htb rate R1 [burst B1] [mpu B] [overhead O]

[prio P] [slot S] [pslot PS]

[ceil R2] [cburst B2] [mtu MTU] [quantum Q]

rate rate allocated to this class (class can still borrow)

burst max bytes burst which can be accumulated during idle period {computed}

mpu minimum packet size used in rate computations

overhead per-packet size overhead used in rate computations

linklay adapting to a linklayer e.g. atm

ceil definite upper class rate (no borrows) {rate}

cburst burst but for ceil {computed}

mtu max packet size we create rate map for {1600}

prio priority of leaf; lower are served first {0}

quantum how much bytes to serve from leaf at once {use r2q}

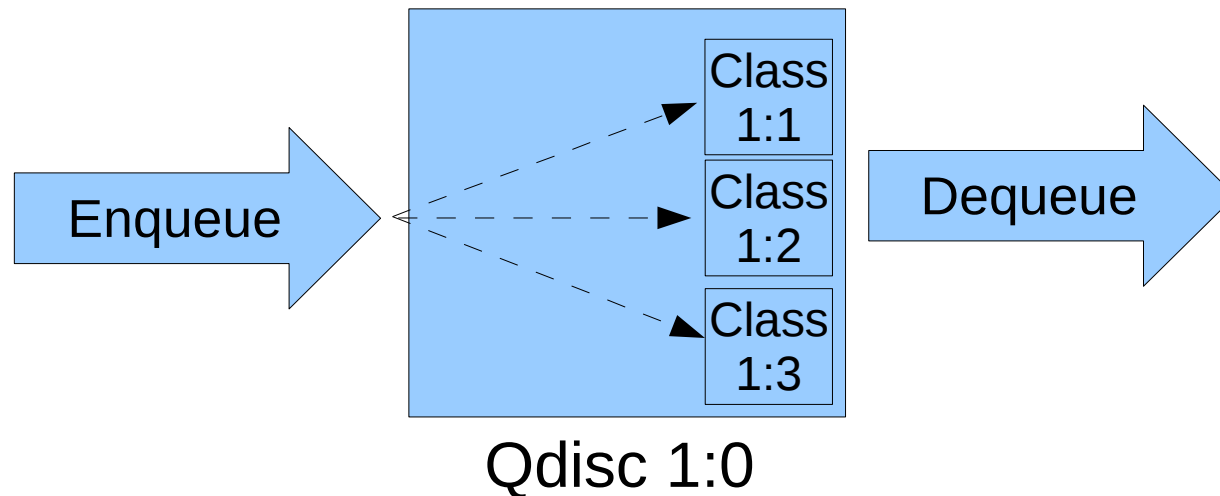
TC HTB version 3.3

Ingress/Egress QoS

- Each network device gets assigned a Queueing discipline.
 - It is also possible to configure the system so that more devices share the same Queueing Discipline/s.
- Packets go through Queueing Disciplines on the way out.
 - It is also possible to configure the system so that this happens also on the way in.

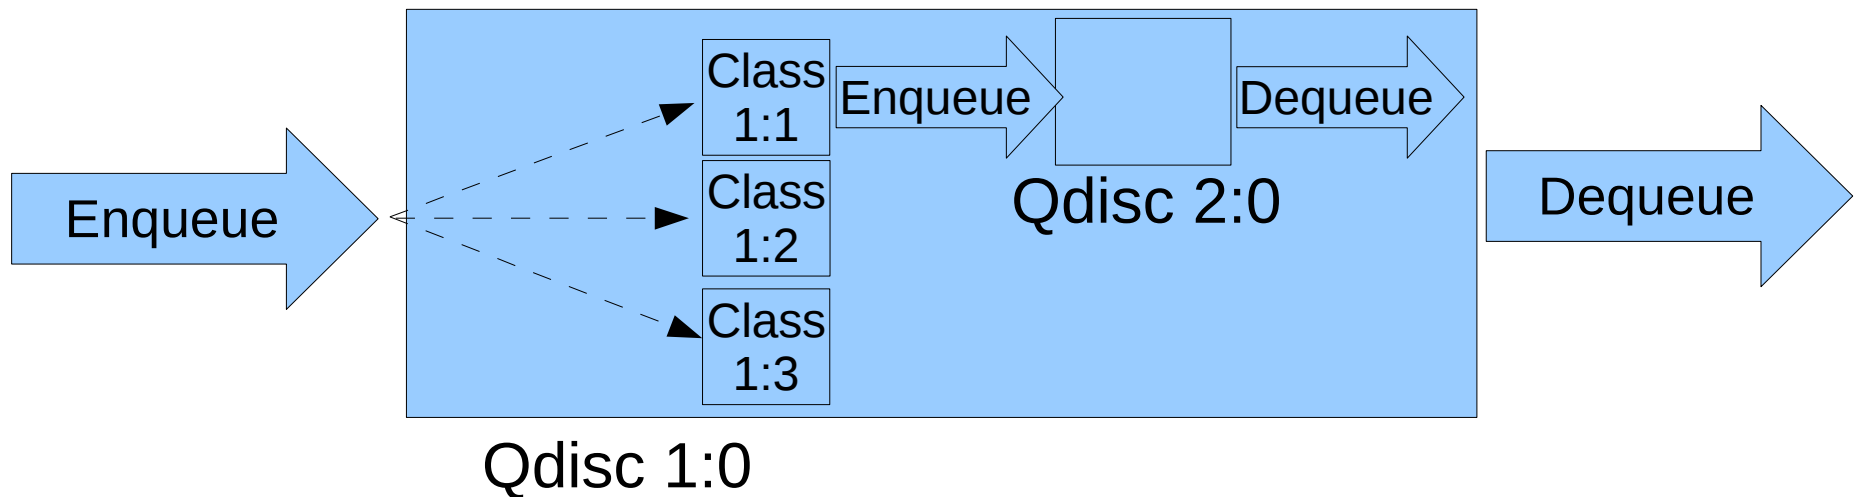
Qdisc/Class IDs

- ID of a Queueing Discipline:
 - X:0 (example: 1:0)
 - ID of its classes (if classfull):
 - X:Y (examples: 1:1 1:2 1:3 etc)



Nested Qdisc

- I can attach a Queueing Discipline to a class



Exercise 1

- Each group gets assigned one queueing discipline. Using the Internet (or its knowledge/experience), the group will investigate on the assigned qdisc and later share with the other groups the following information:
 - How the qdisc works (ie, the algorithm)
 - The meaning of its configuration parameters

Group 1: **RED**

Group 2: **TBF**

Group 3: **SFQ**

Group 4: **NETEM**

Exercise 2

- Use the ***tc show*** command to check the default QoS configuration of the local interfaces.

Exercise 3

- Download Wondershaper and try to understand what configuration it applies.
 - <http://lartc.org/wondershaper/>

Exercise 4

- Configure HTB on the group router so that:
 -
 -
 -
 -
 -

Documentation

- <http://lartc.org>
 - Good documentation (but not very updated)

Copyright



- This presentation is released under the Creative Common License:
 - Attribution, Noncommercial, Share Alike 2.5
 - (<http://creativecommons.org/licenses/by-nc-sa/2.5/>)
- Attribution
 - You must attribute the work in the manner specified by the author or licensor.
- Noncommercial.
 - You may not use this work for commercial purposes.
- Share Alike.
 - If you alter, transform, or build upon this work, you may distribute the resulting work only under a license identical to this one.