

Internet Quality of Service - Course Review

Dr Brian E Carpenter
IBM and iCAIR

Northwestern University
Winter Quarter 2000

Timeline

- Choose major assignment: before 1/18
- Mid-course quiz: 2/7 (10%)

You are almost here

- Assignment report due: 3/10 (absolutely positively no extensions) (60%)
- Exam: 3/15 9:00 a.m. (30%)

1. Basics

- Why networking is difficult
 - ▶ Communication in the presence of noise
 - ▶ More than a single hop
 - ▶ Congestion is fundamental
- Where we are today
 - ▶ End-to-end principle & triumph of pragmatism
 - ▶ Performance
- TCP versus UDP
- Measurements

The causes of loss

- (very minor) Transmission errors
- (dominant) Congestion, i.e. a router receives more packets than it can send.
 - ▶ Congestion is the main operational problem in the Internet today (but not a problem for network operators, since the routers happily discard excess traffic).
 - ▶ Congestion results from massive scale, rapid growth, and poor traffic engineering.

Congestive collapse of TCP

- Increased congestion leads to longer queues, which increases RTT
- Long queues eventually fill up, causing packets to be discarded, i.e. increased loss.
- Increased RTT and increased loss is a double whammy on throughput.
 - ▶ double RTT plus 10% loss rate reduces throughput to 16%
 - ▶ 4x RTT plus 20% loss rate reduces throughput to 5%

TCP versus UDP

- There is a fundamental difference between TCP and real time UDP traffic:
 - ▶ TCP traffic comes at an elastic rate, slows down under congestion, and can be effectively policed by routers without damaging the application.
 - ▶ Real time UDP traffic comes at a constant rate, ignores congestion, and if hit by loss or policed by routers, the application may become unusable.

Measurement

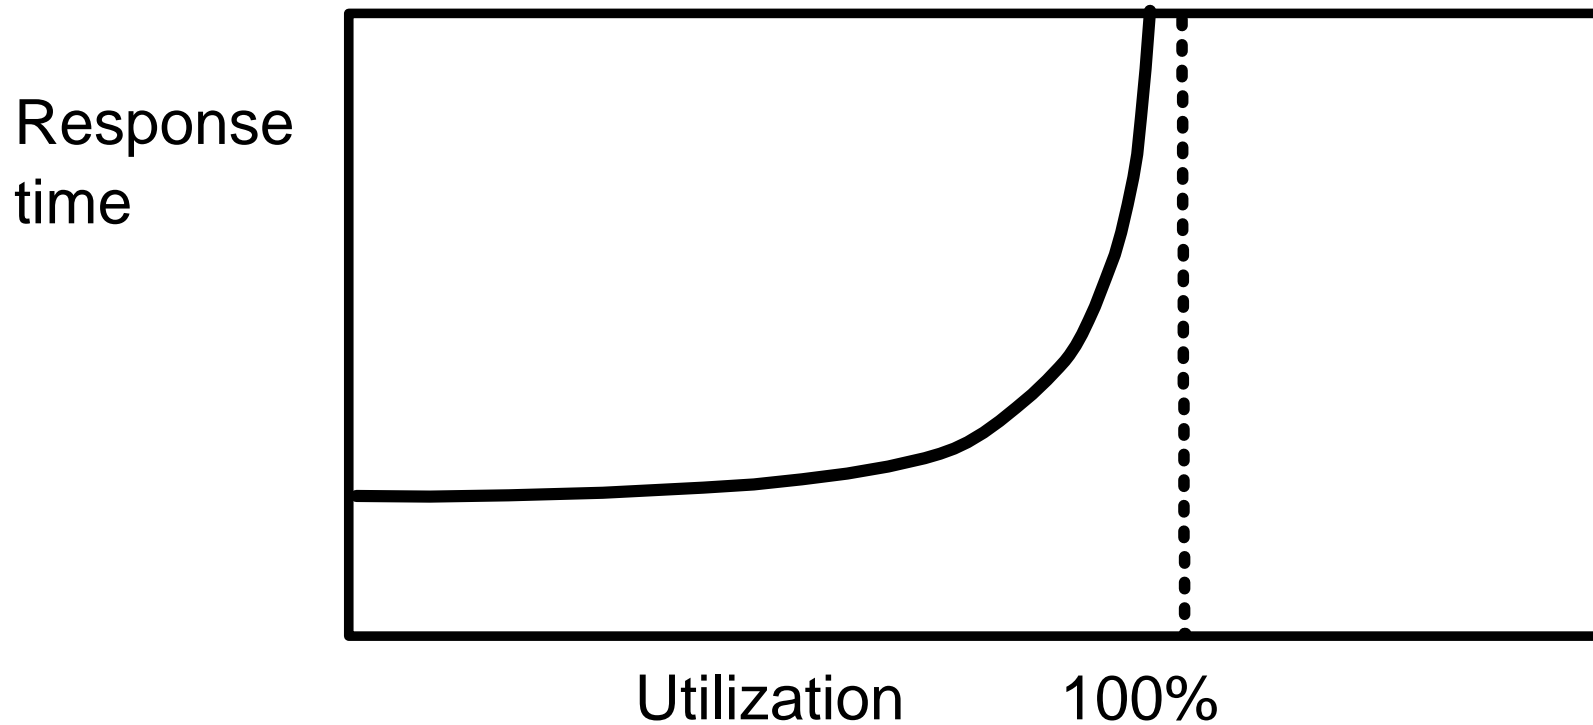
- Network traffic is too complex to understand theoretically - measurement is essential
- Passive measurement - capture and analyze real traffic.
 - ▶ measurement doesn't interfere with reality, but analysis may be very hard.
- Active measurement - generate, measure and analyze synthetic traffic
 - ▶ analysis is more straightforward, but synthetic traffic may change network behaviour.

2. Queueing theory, traffic statistics, and the real world

- Gauss, Poisson and Markov
- Queueing theory
- Traditional traffic models and Internet reality


M/M/1 behavior

- Plot response time (queueing plus processing) against utilization (% busy)





Traffic plots (sketch)

Poisson




5 seconds


IP measured



50 seconds



500 seconds



The consequence of burstiness

- If traffic is arbitrarily bursty, then any finite buffer size will sometimes be too small.
 - ▶ Corollary: no amount of money can prevent buffer overflow.
- Conclusion: either burstiness must be abolished, or packet loss must be tolerated.

3. Quality of Service versus Service Quality, and Service Level Agreements

- Why QOS is not the same thing as service quality
- Real world requirements
- The role of service level agreements (SLAs)

So, what is Service Quality?

- Users want the "*bilities*" ...
 - ▶ reliability, stability - no glitches
 - ▶ availability - 366x24
 - ▶ predictability - "the same rotten service at 3 a.m. as at 3 p.m." regardless of other users
 - ▶ applicability - supports the user's applications
 - ▶ survivability - service must survive backhoes and virtual backhoes
 - ▶ scalability - more capacity there when needed
 - ▶ responsibility - a help desk that helps

And what is Quality of Service?

- Usually quite narrowly defined to mean a specific, measurable set of characteristics of a data stream, such as
 - ▶ transit delay
 - ▶ jitter
 - ▶ loss rate
 - ▶ throughput

that may be needed by a specific application or user. We will examine this in detail later.

SLAs replace "best effort"

- Internet service provision is a competitive business. Large customers now require some assurance of service quality, which means they will require their service provider to sign up to a contract, usually called a service level agreement. It will lay down specific criteria for the "ilities".
 - ▶ Best effort is no longer good enough

4. QOS Mechanisms.

Traffic Engineering for QOS

- What are the problems to be solved?
- Mechanisms for solving them
- Making it work how you want it to work:
Traffic Engineering: network design and routing in support of QOS

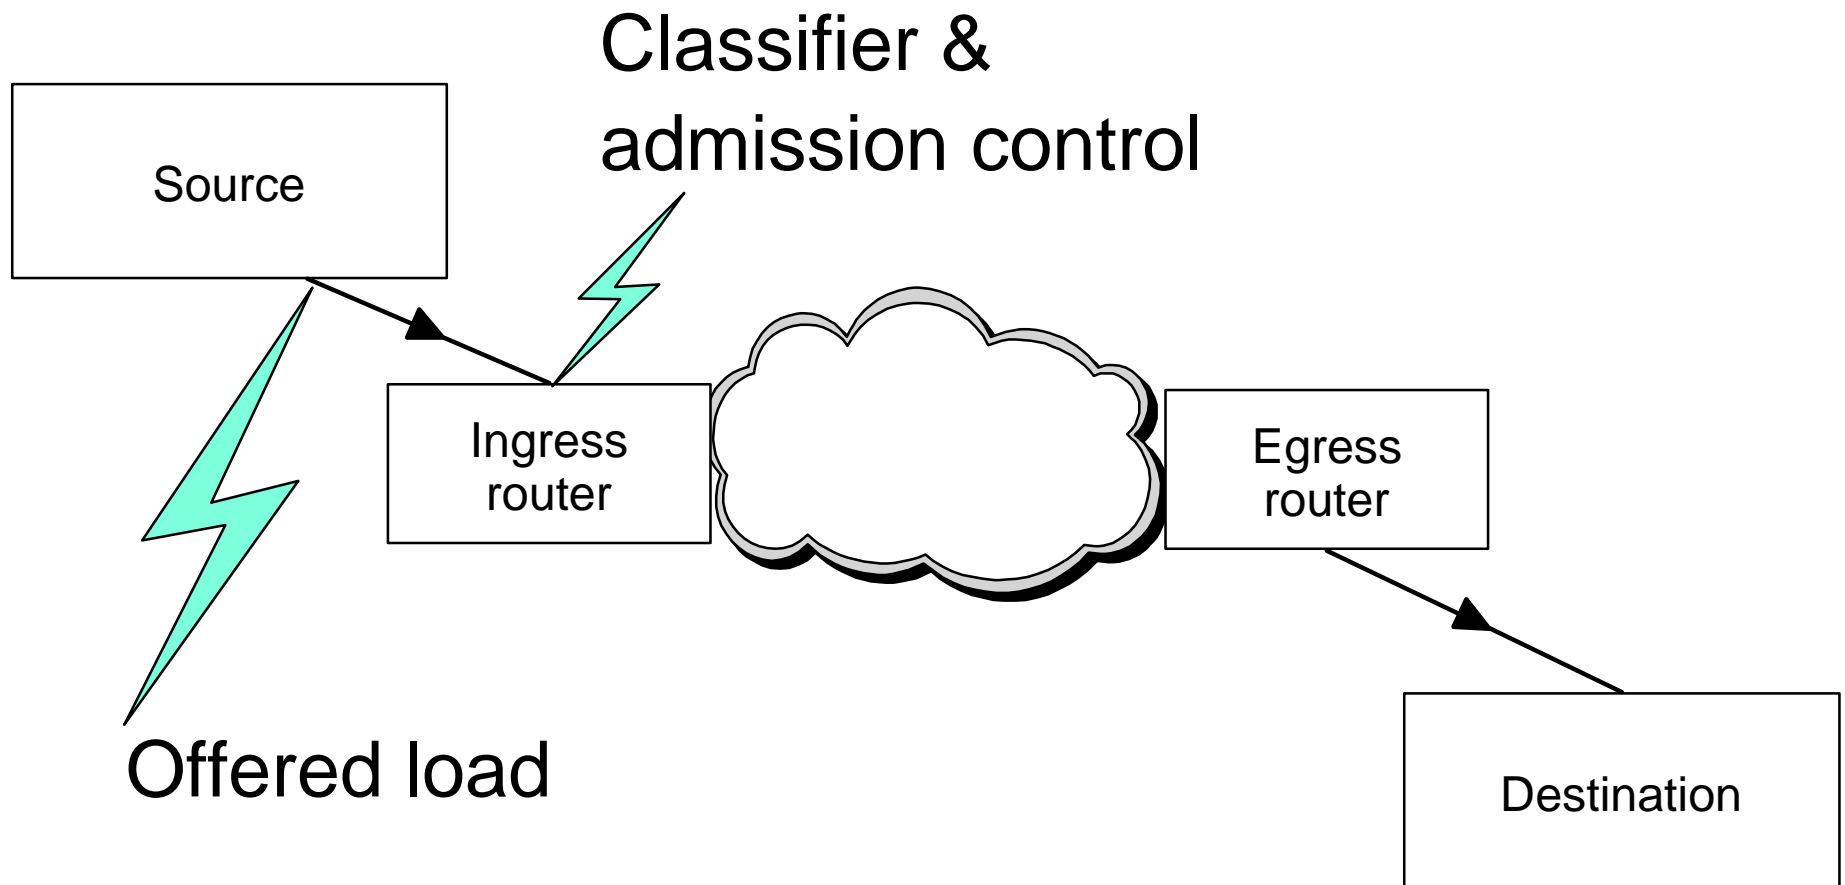
Problems to be solved

- How can QOS be described and measured?
- How do users request QOS?
- How do service providers deliver QOS?
- How do users know what QOS they get or ISPs what QOS they deliver?
- How can QOS be provided for millions of individual communications sessions?

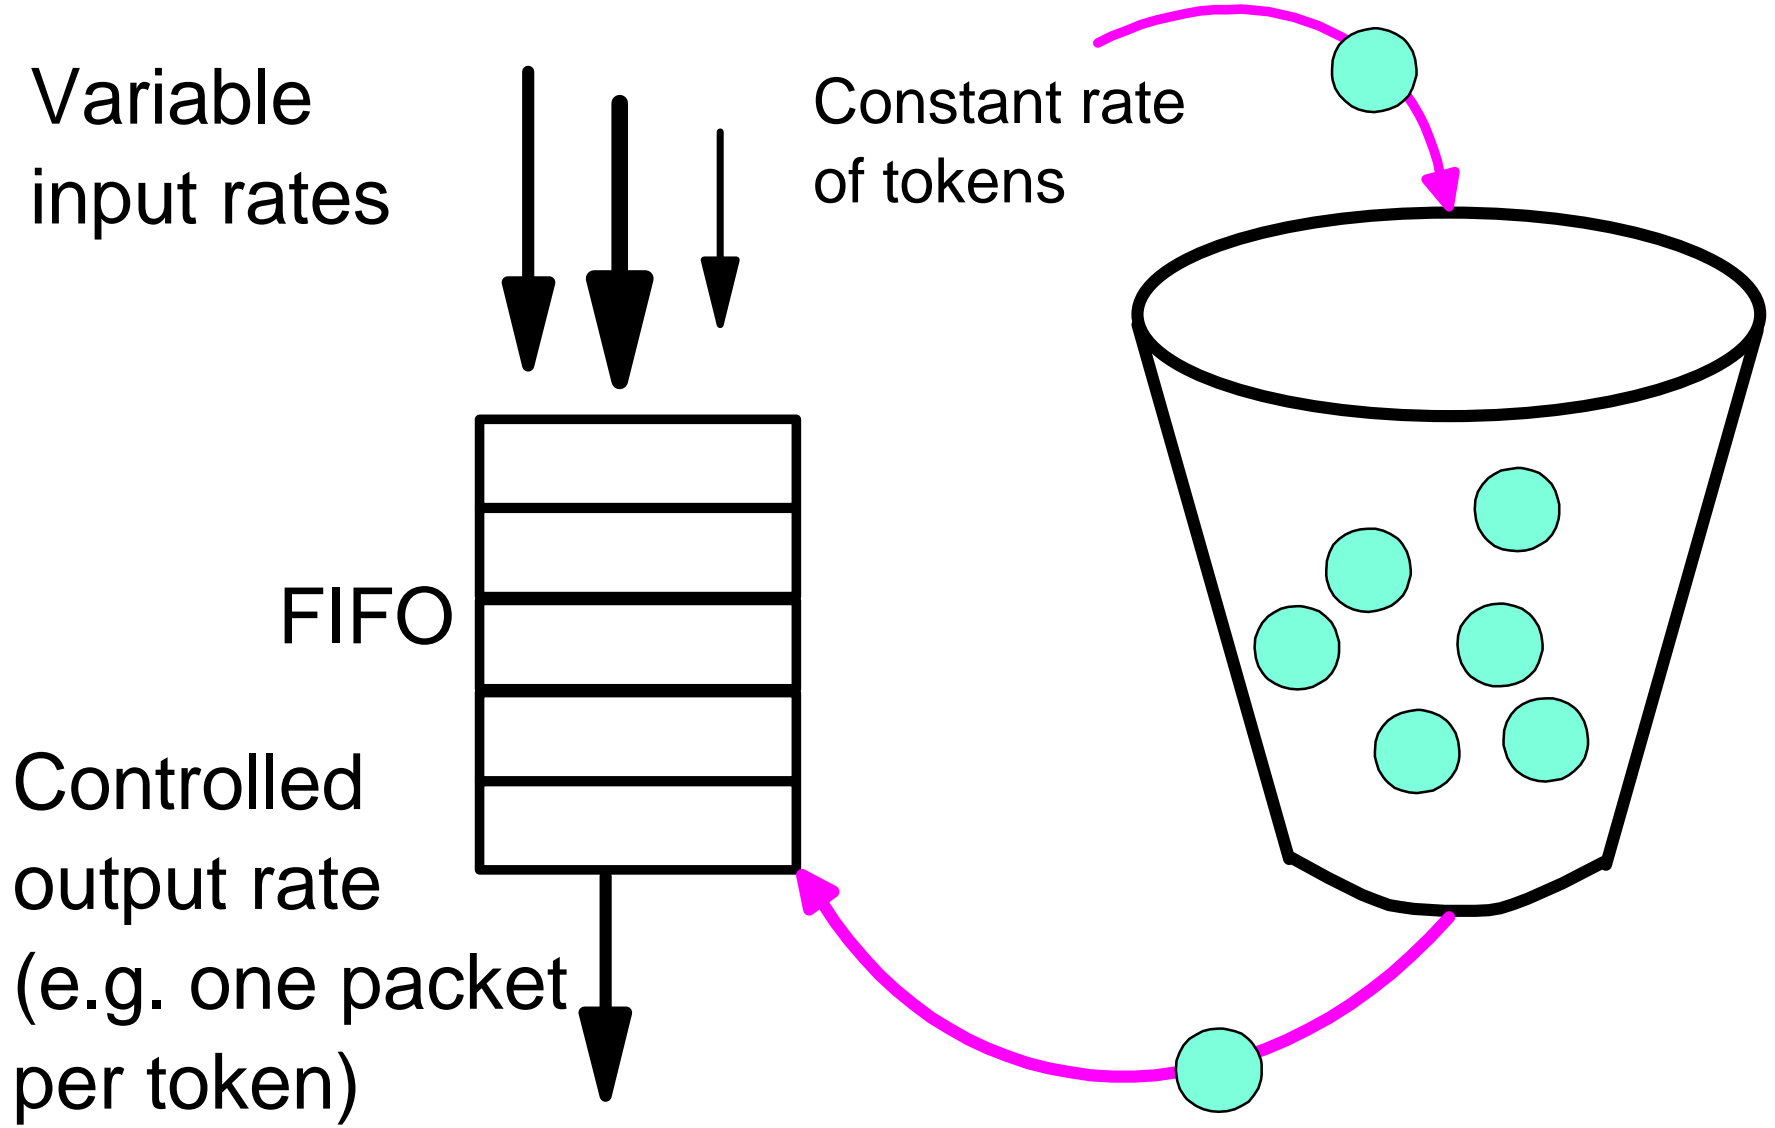
What is an "atom" of QOS?

- Useful concepts:
 - ▶ the classic 5-tuple and microflows
 - ▶ offered load
 - ▶ classifiers
 - ▶ admission control
 - ▶ traffic shaping (rate controllers, droppers, schedulers, leaky buckets, token buckets)
 - ▶ committed access rate
 - ▶ fair queueing
 - ▶ service types

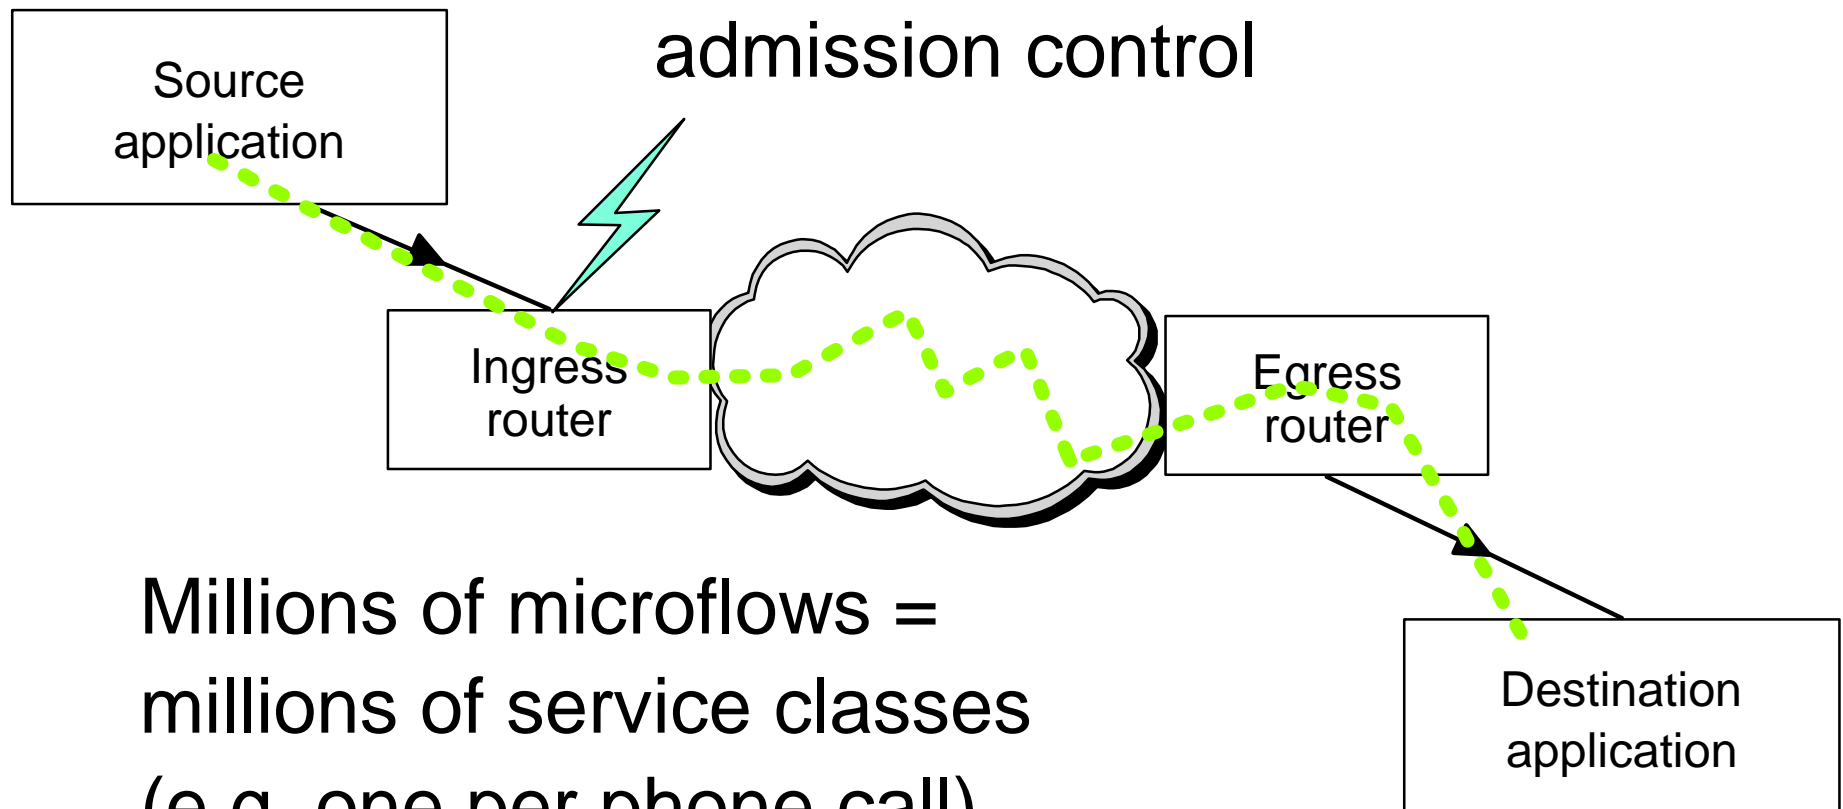
A picture




Token bucket shapers



Service class per microflow



Pre-existing service classes



Many microflows share
one service class
(e.g. all phone calls in one class)

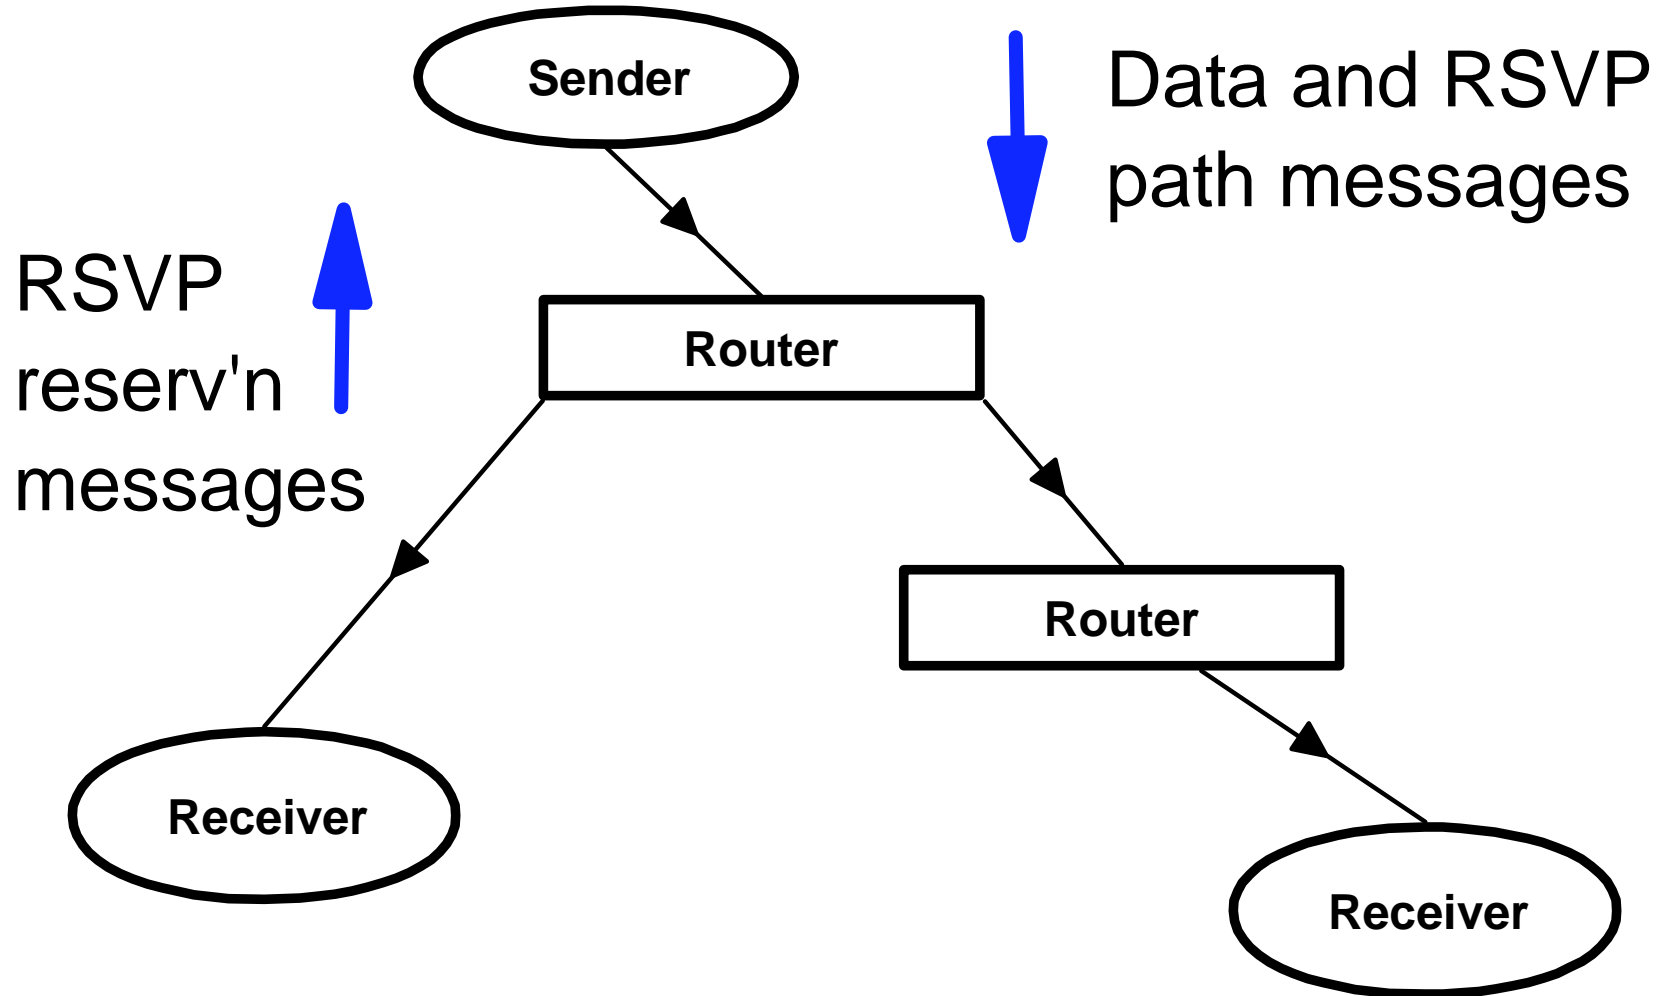
Traffic engineering

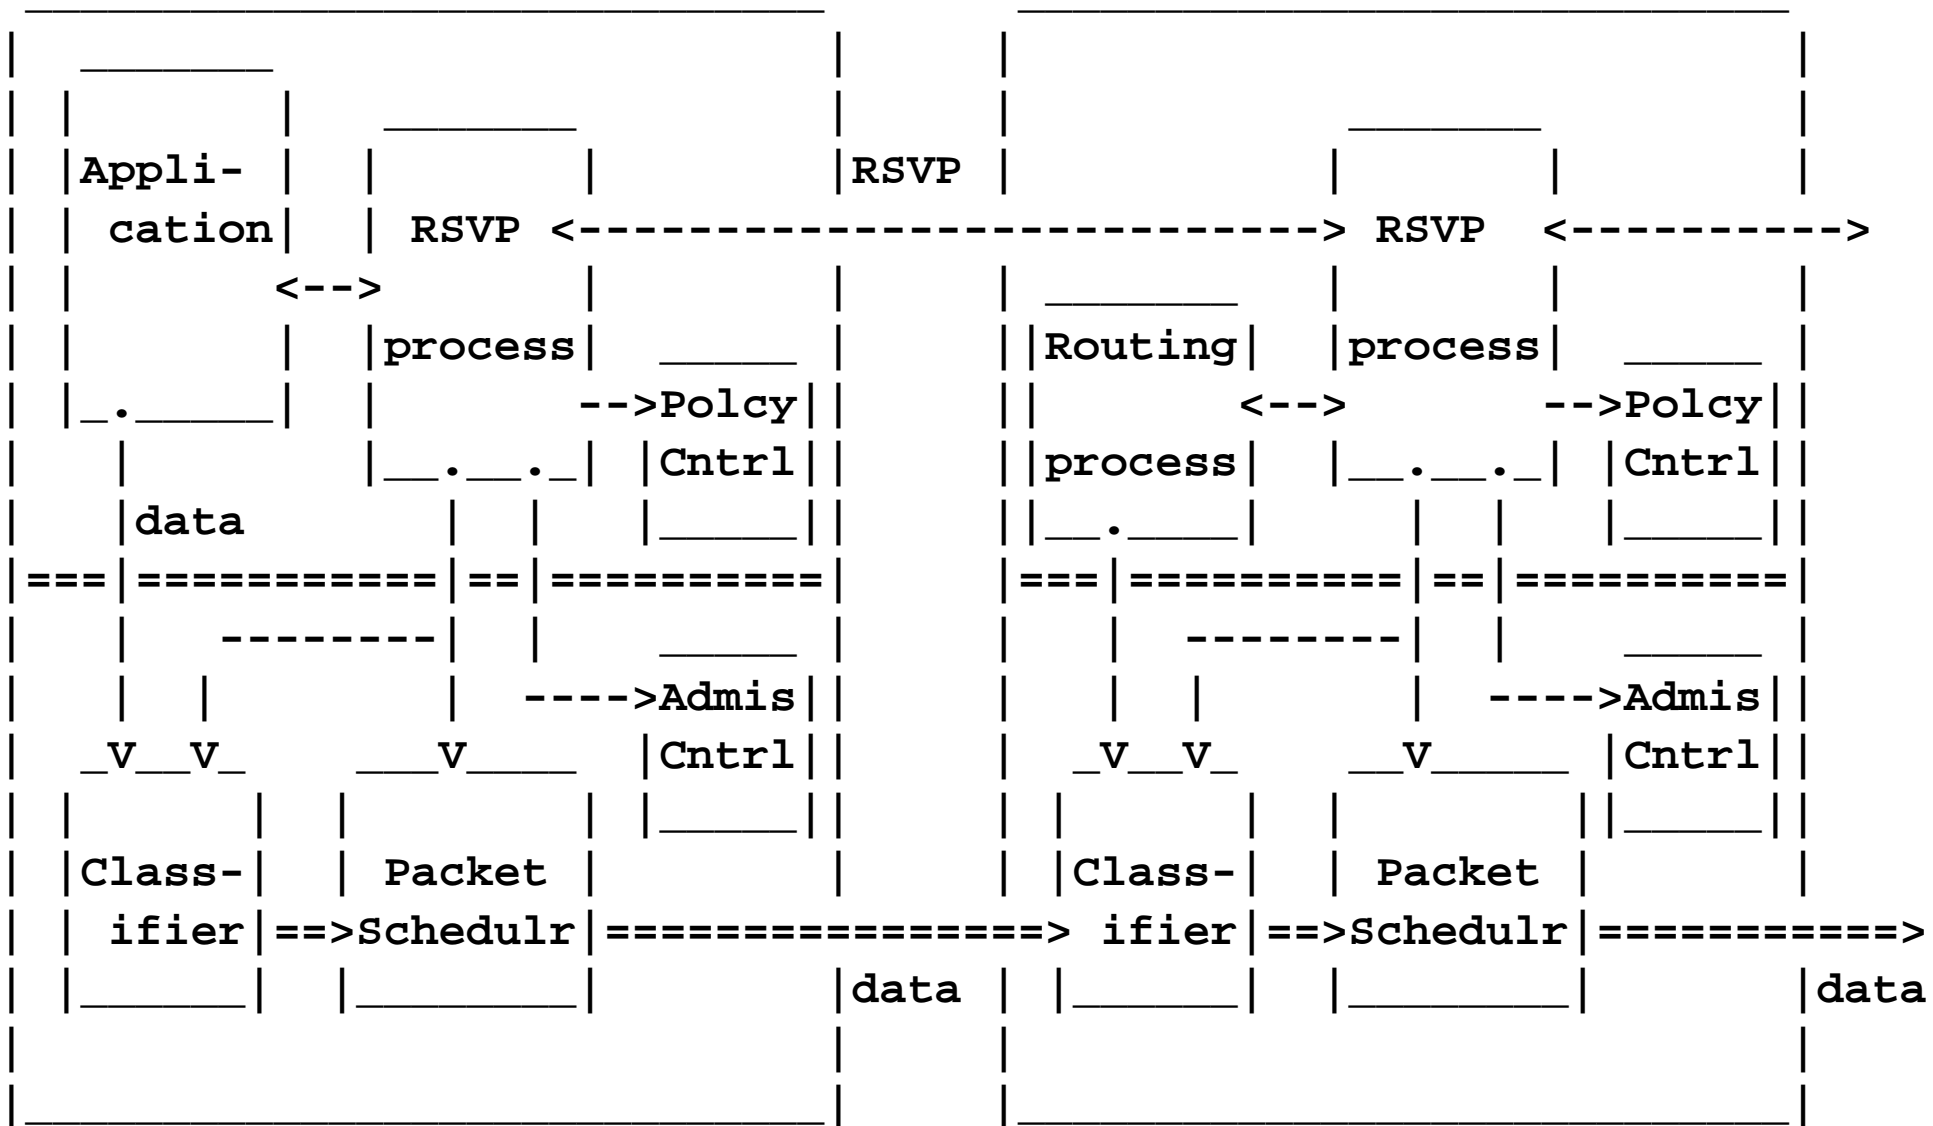
- This is the art of providing adequate resources and ensuring that traffic follows the desired paths through the network.
 - ▶ designing network topology
 - ▶ buying adequate routers and renting adequate links between them
 - ▶ configuring routing and queueing so that things work as intended

5. Integrated Services and RSVP Signalling

- When voice and video applications were first envisaged, researchers realised that service guarantees were needed (especially throughput, delay, and jitter). Clearly, classical IP does not offer such guarantees.
- An approach called "Integrated services" was developed.
- Integrated Services relies on a signalling protocol known as RSVP.

RSVP scenario





Integrated Services routers at ISP backbone scale


- Each backbone router must
 - ▶ hold state for millions of communications
 - ▶ classify each packet on {source addr, source port, destination addr, dest port, protocol} to see if they are Int Serv packets
 - ▶ apply controlled load or guaranteed service checks to each Int Serv packet.
- Products are appearing, but it is widely believed that Int Serv cannot be deployed beyond campus scale or intranets.

6. Differentiated Services

- With Integrated Services and RSVP, we saw two major problems
 - ▶ the amount of per-flow state information and the amount of per-packet processing in each router would overwhelm ISP backbone routers carrying millions of flows.
- Clearly another solution is needed for wide area ISP-scale QOS

Basic idea

- ▶ Get rid of the per-flow state and complex processing by marking each packet with a simple flag indicating how to treat it.
- ▶ Make it a one-way process so that asymmetric routes don't matter



Classifier

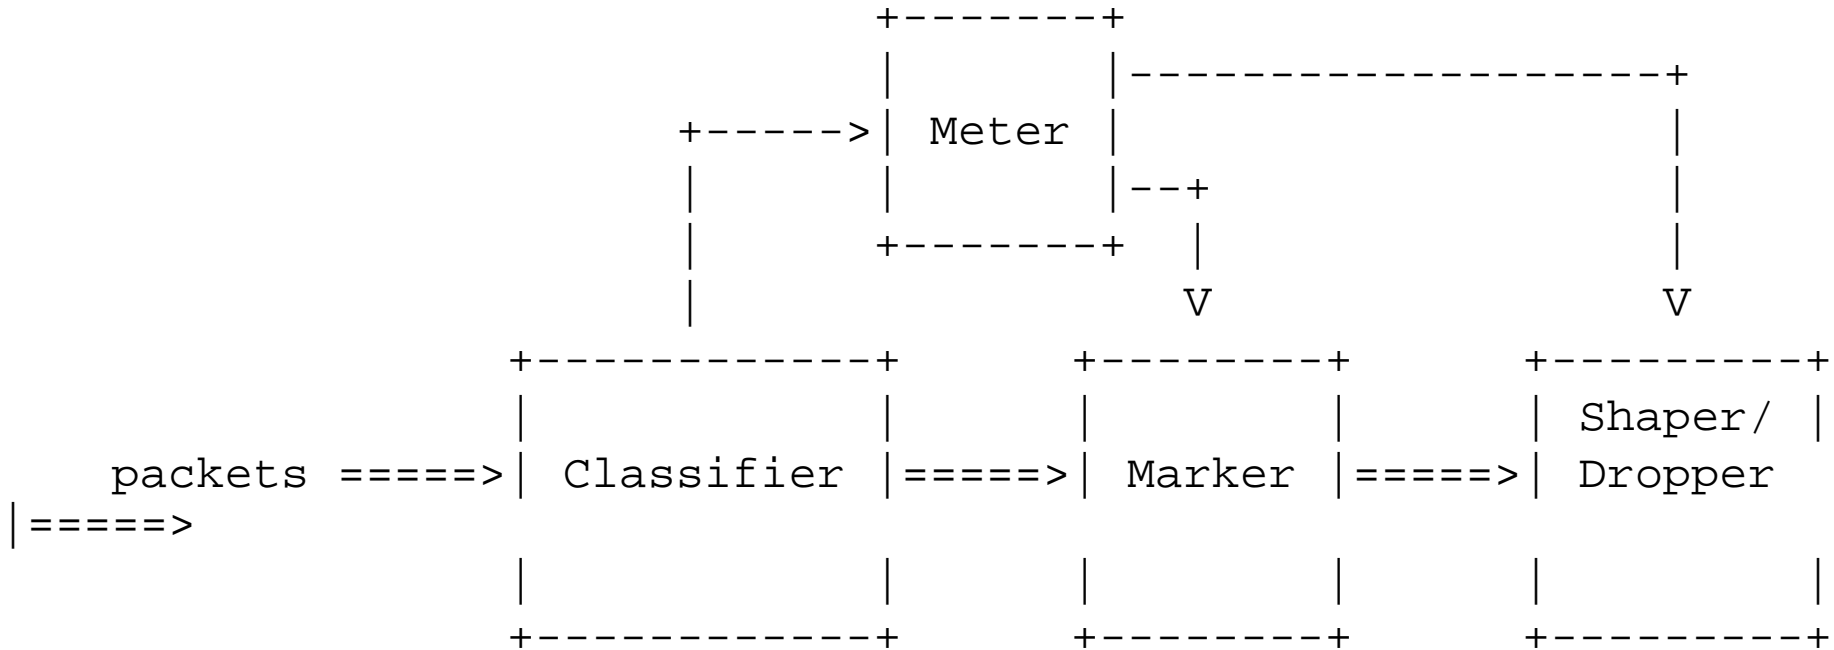



Fig. 1: Logical View of a Packet Classifier and Traffic Conditioner

The DS Field layout



PHB: per-hop behavior

CU: currently unused

(bit 0 = most significant if relevant)

Uses TOS octet in IPv4 and Traffic Class octet in IPv6.

The EF (expedited forwarding) PHB

0	1	2	3	4	5	6	7
+-----+	+-----+	+-----+	+-----+	+-----+	+-----+	+-----+	+-----+
1	0	1	1	1	0	CU	
+-----+	+-----+	+-----+	+-----+	+-----+	+-----+	+-----+	+-----+

The departure rate of EF packets from any diffserv node must equal or exceed a configurable rate. EF traffic should receive this rate independent of the intensity of any other traffic attempting to transit the node.

====> Virtual Leased Line service

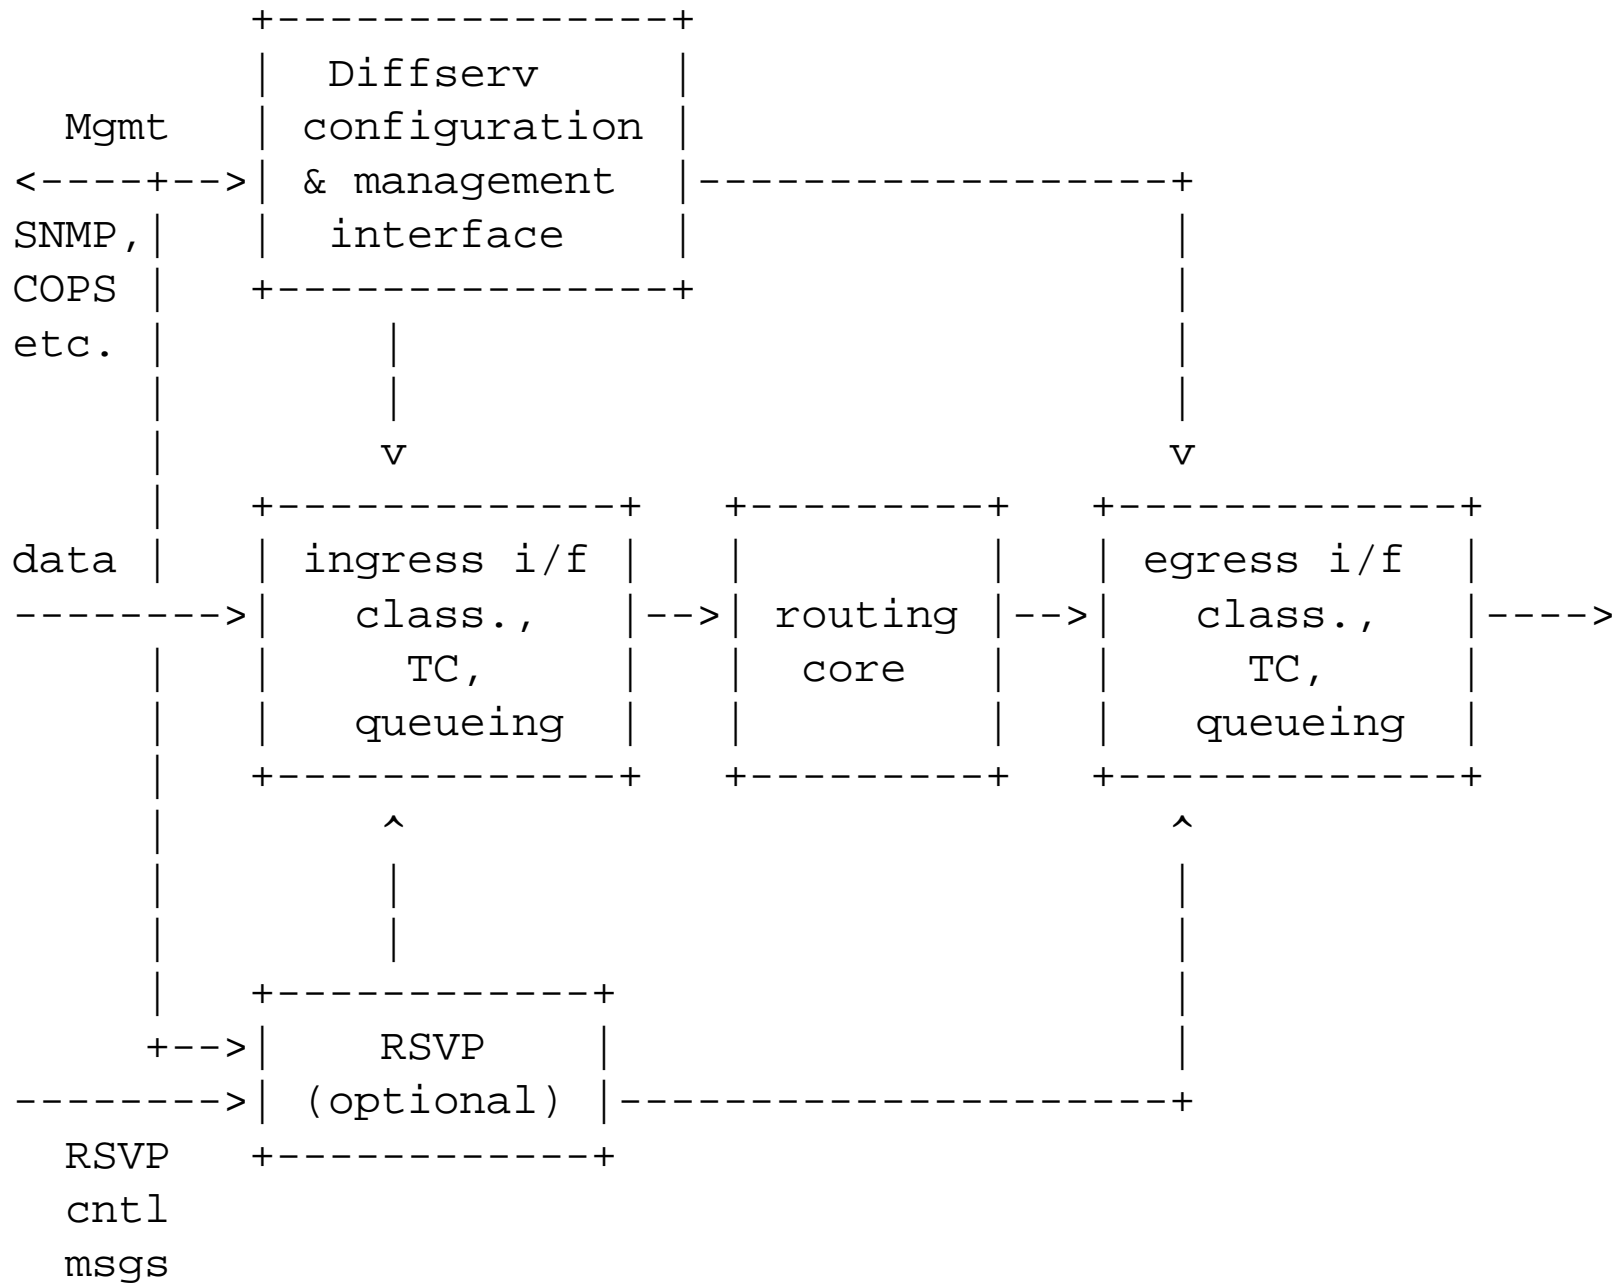


Figure 1: Diffserv Router Major Functional Blocks

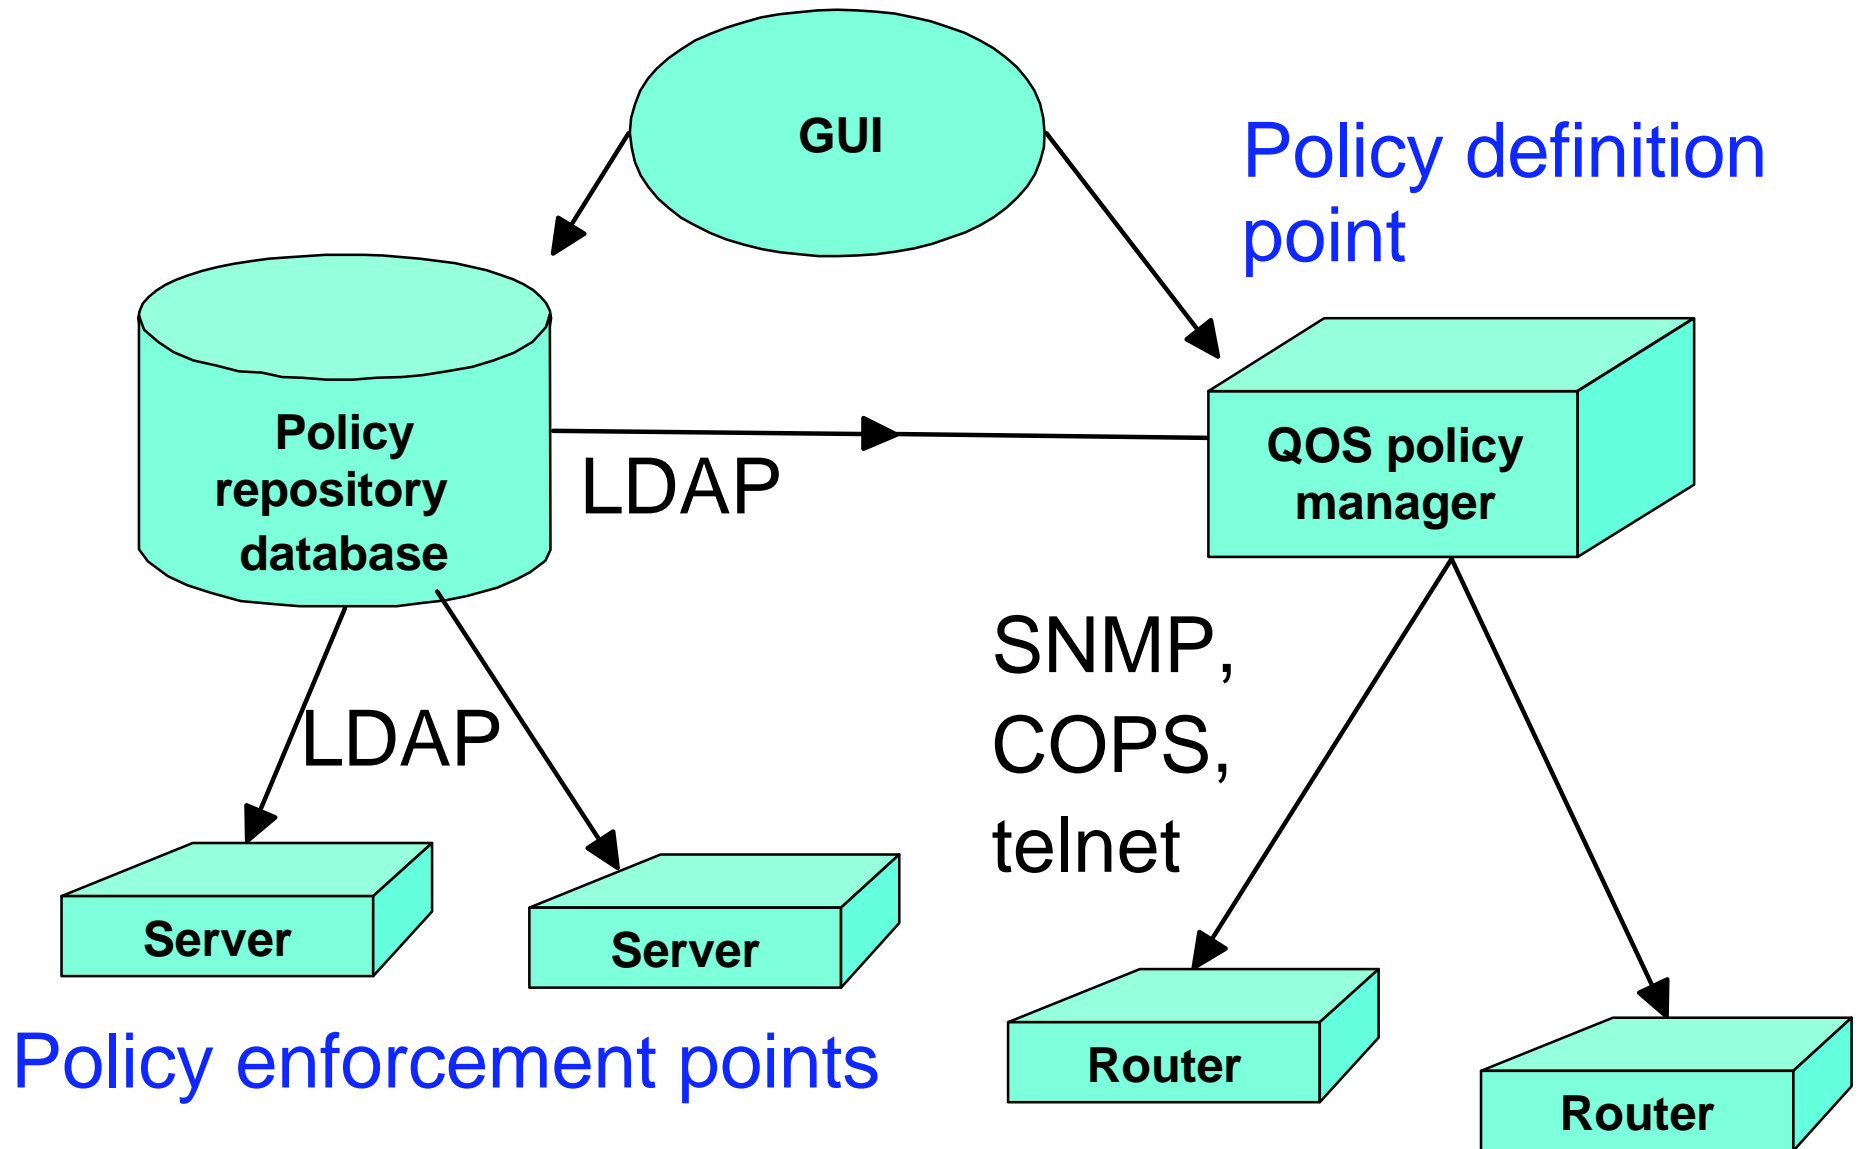
7. Defining and delivering service policy

- Network Management - state of the art.
- Policy management and configuration management
- One model for QOS policy management
- Current work in the IETF
- Bandwidth brokers

Policy management

- Increasingly, especially with quality of service support and new security tools coming, companies want to centrally manage policy issues such as
 - ▶ which users may initiate video-conferences?
 - ▶ which users may request better quality, and for which applications?
 - ▶ what happens when demand exceeds capacity?
 - ▶ security policy also needs to be managed

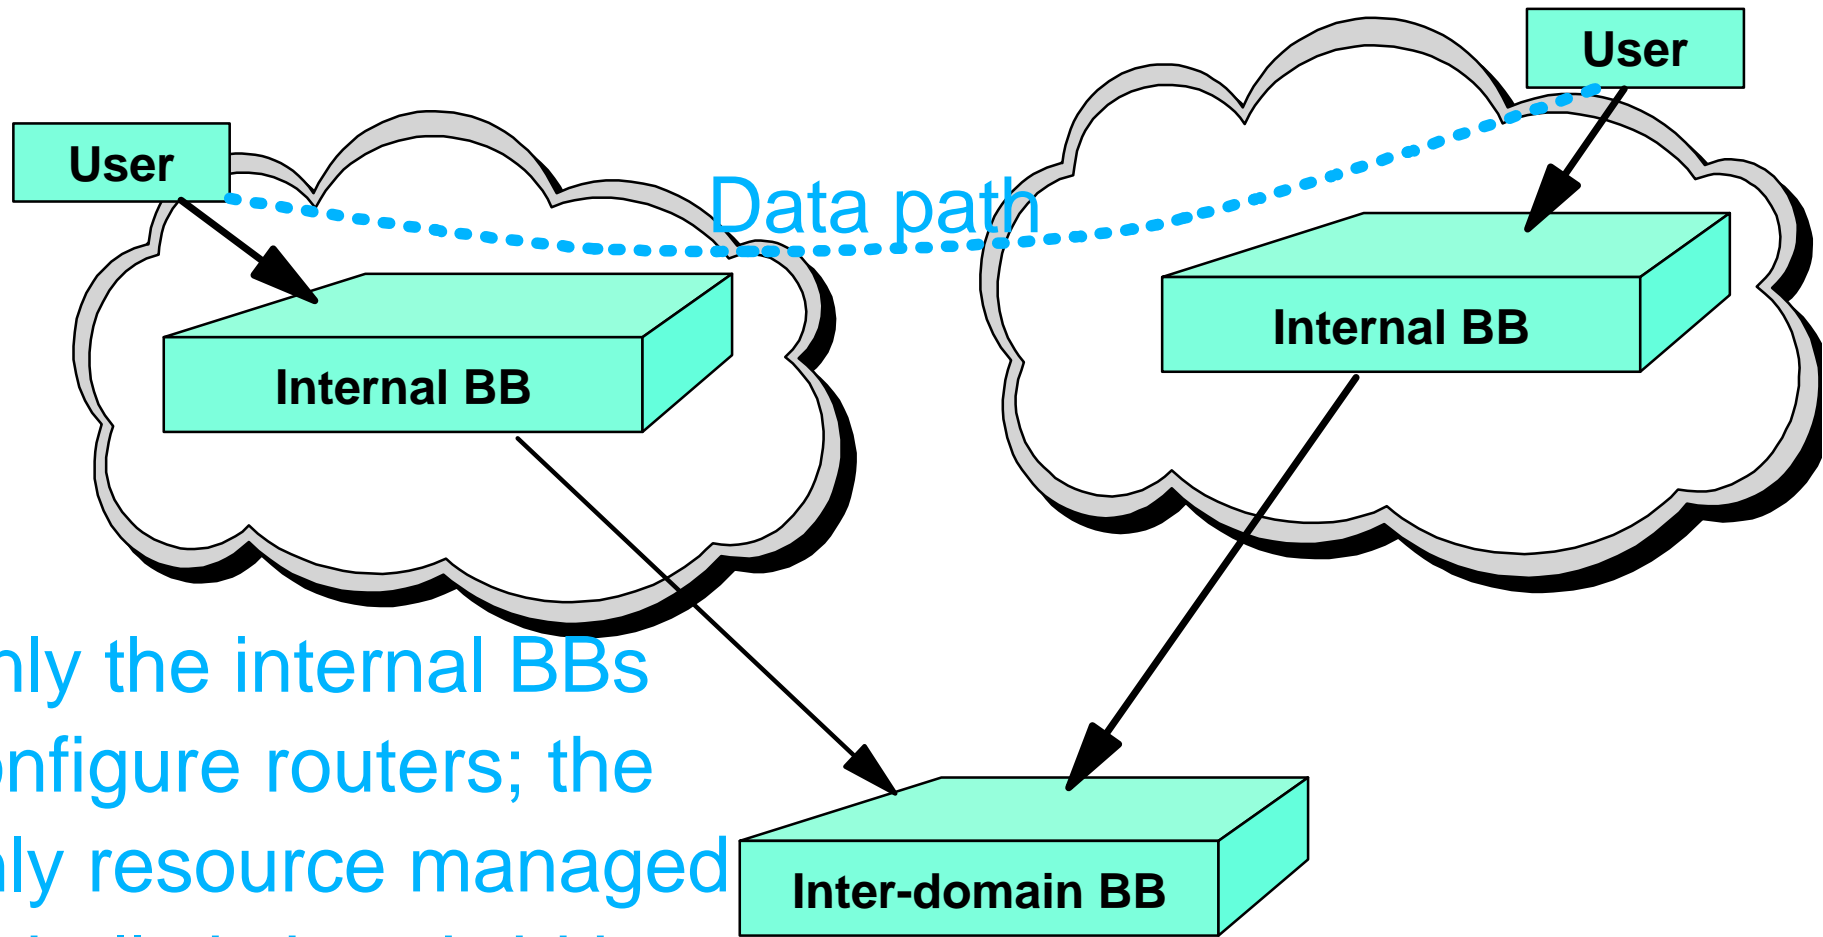
A policy management system



What goes in a DiffServ router MIB?

- Example only - still evolving in the IETF. Think of it as a linked-list data structure with multiple elements:
 - ▶ Behavior Aggregate Classification Table
 - ▶ Multi-Field Classification Table
 - ▶ Classifier Table
 - ▶ Meter Table
 - ▶ Action Table
 - ▶ Queue Table

Bandwidth broker picture



Only the internal BBs
configure routers; the
only resource managed
globally is bandwidth

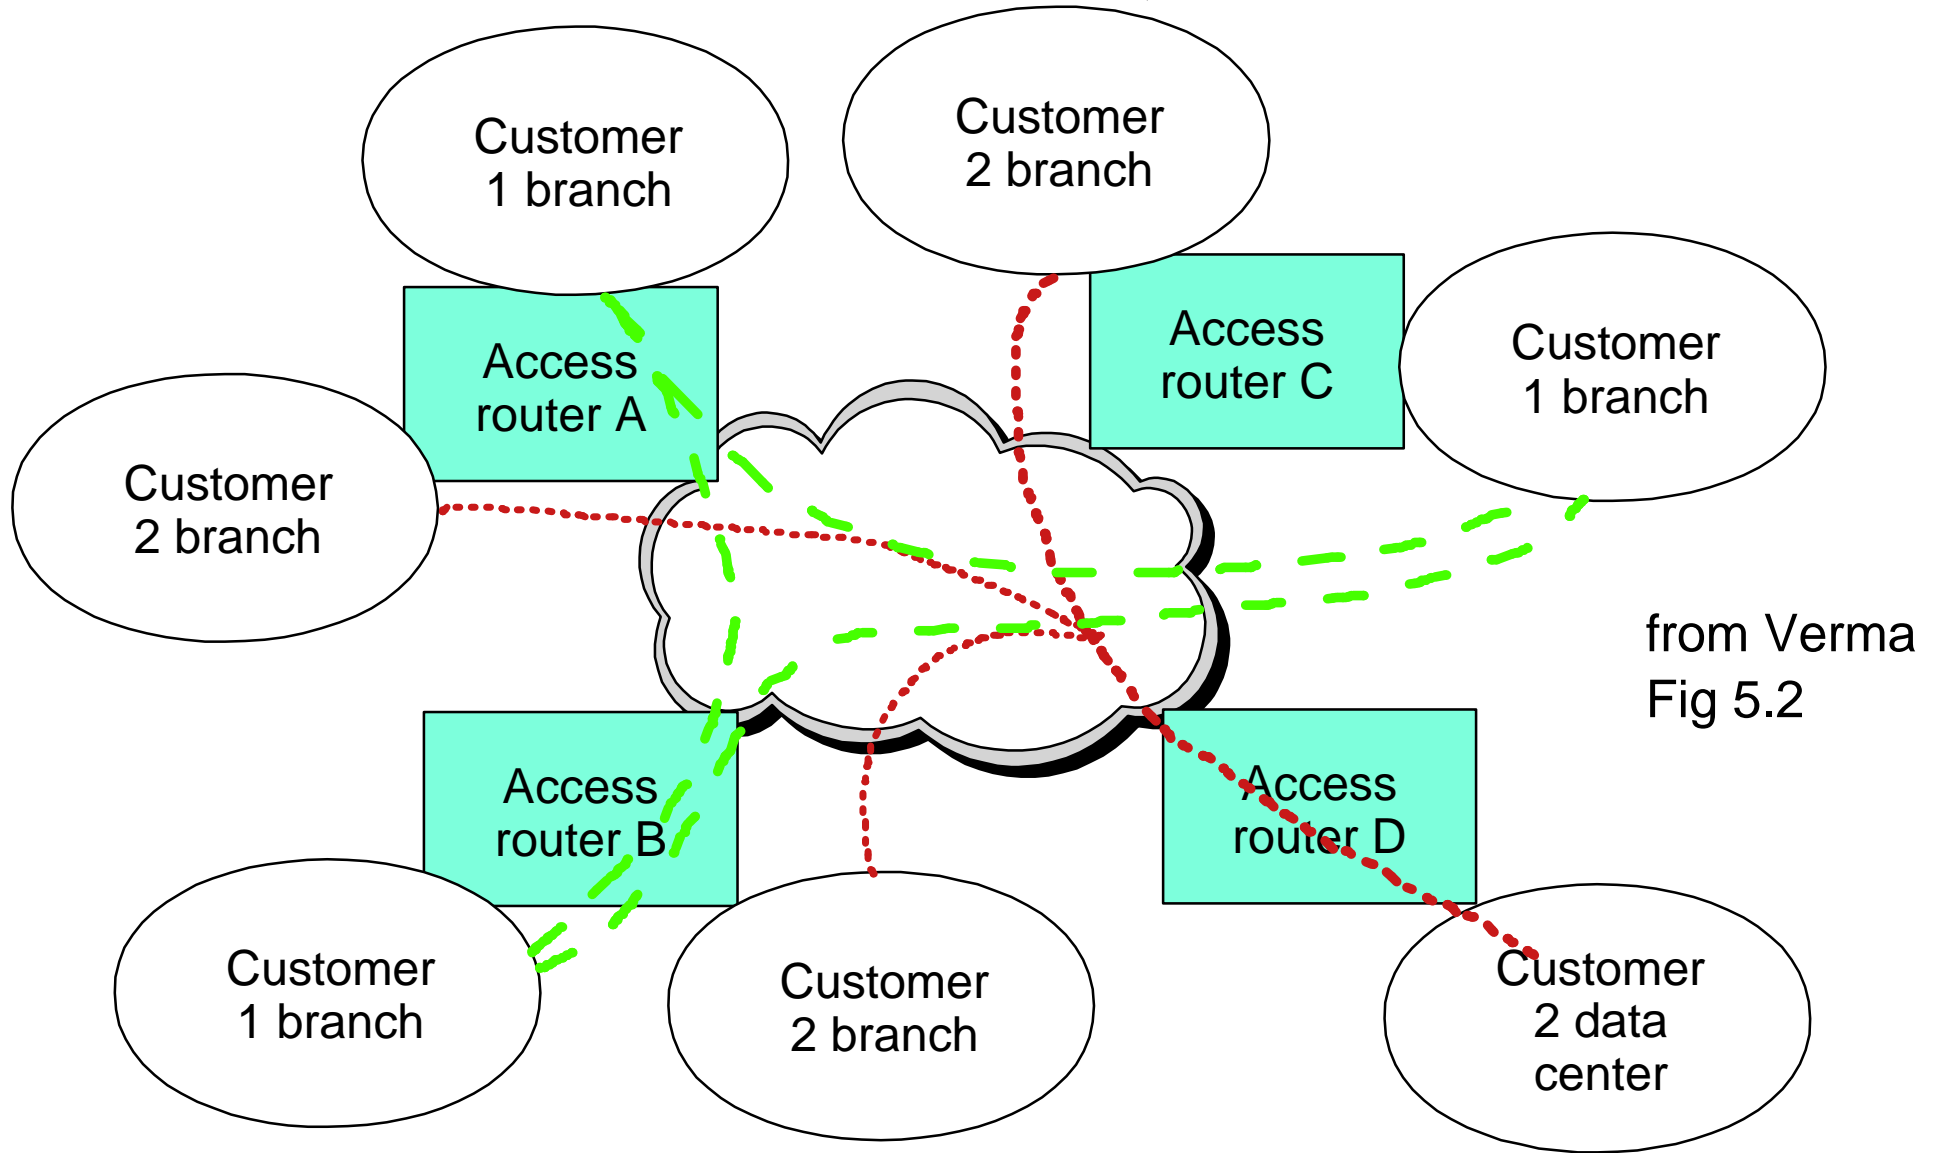
8. Service Level Agreements

- Review of need for SLAs & content of SLAs
- Three types of SLAs
- A generic SLA architecture
- Network Connectivity SLAs with Best Effort and more
- Network monitoring for SLAs

Typical contents of an SLA


- Expected performance level (metrics)
- Reporting process
- Time frame for problem resolution
- Monitoring and reporting
- Price of service, penalties for poor service
- Legal stuff

Several customers, one network




Satisfying the SLA

- Assuming each customer access has a fixed access rate (e.g. a T1 or T3 line), then we know the absolute limit on the traffic load. We can meet the required average delays by using the response curve to calculate the required backbone link capacities.



Components of a "commit" type of transaction (simplified)

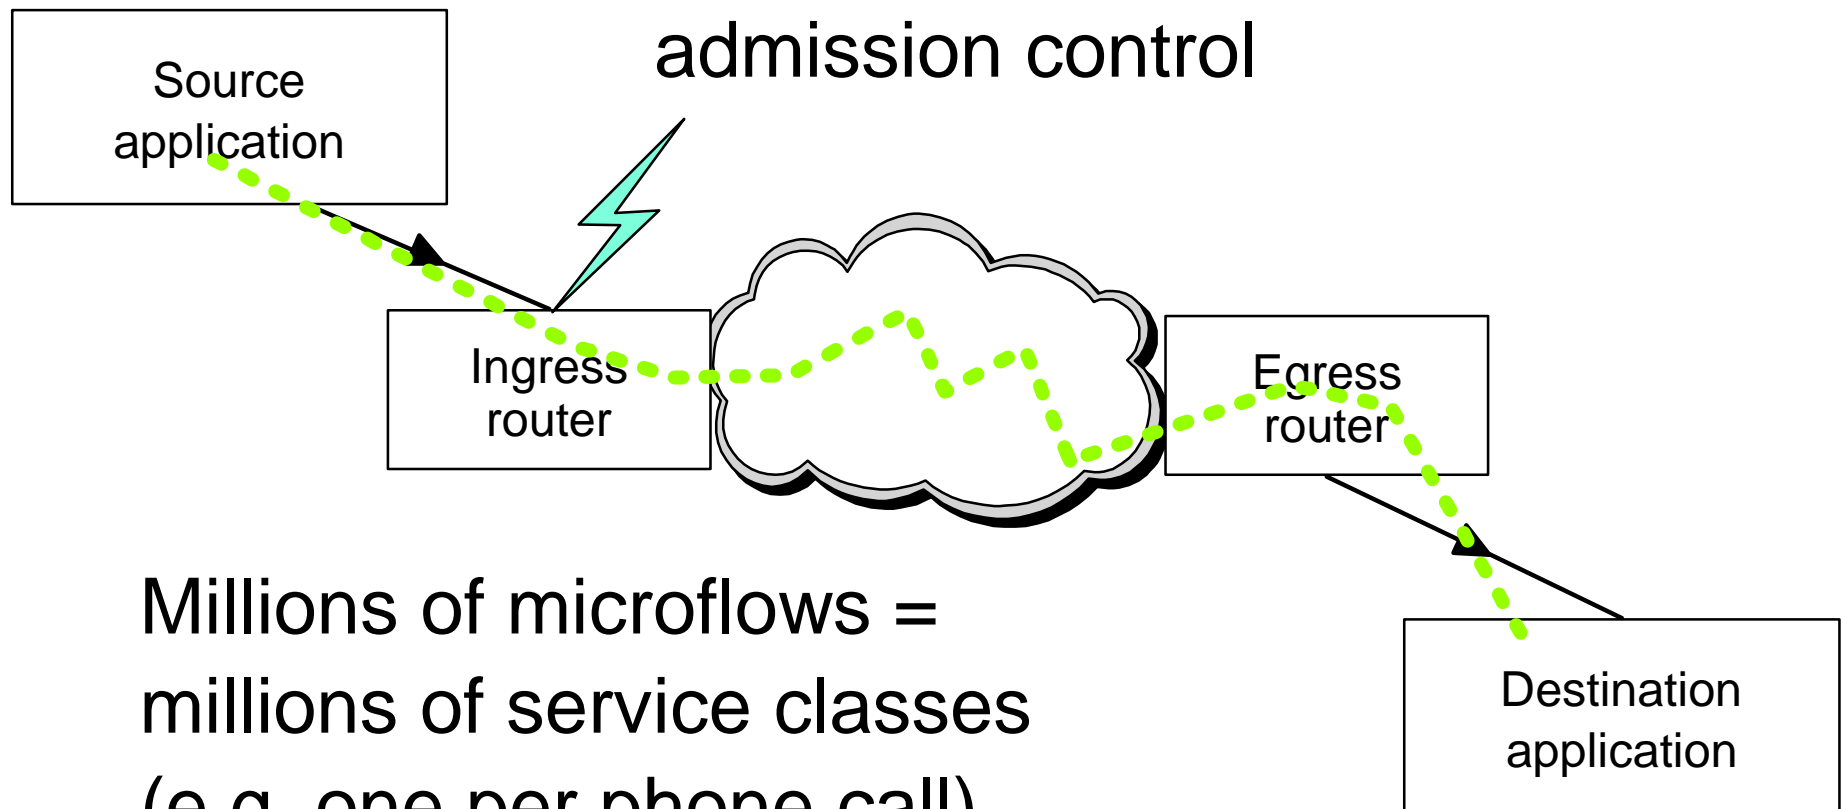


Elements of a generic SLA architecture

- Edge devices - the boxes that define the boundaries within which an SLA applies, for example access routers and servers.
Outside those boundaries: no guarantees.
- Policy server - the system that delivers the Service Level Specification (configuration data) to edge devices & other boxes.
- Performance monitor
 - ▶ Policy server and performance monitor get SLS data from policy repository

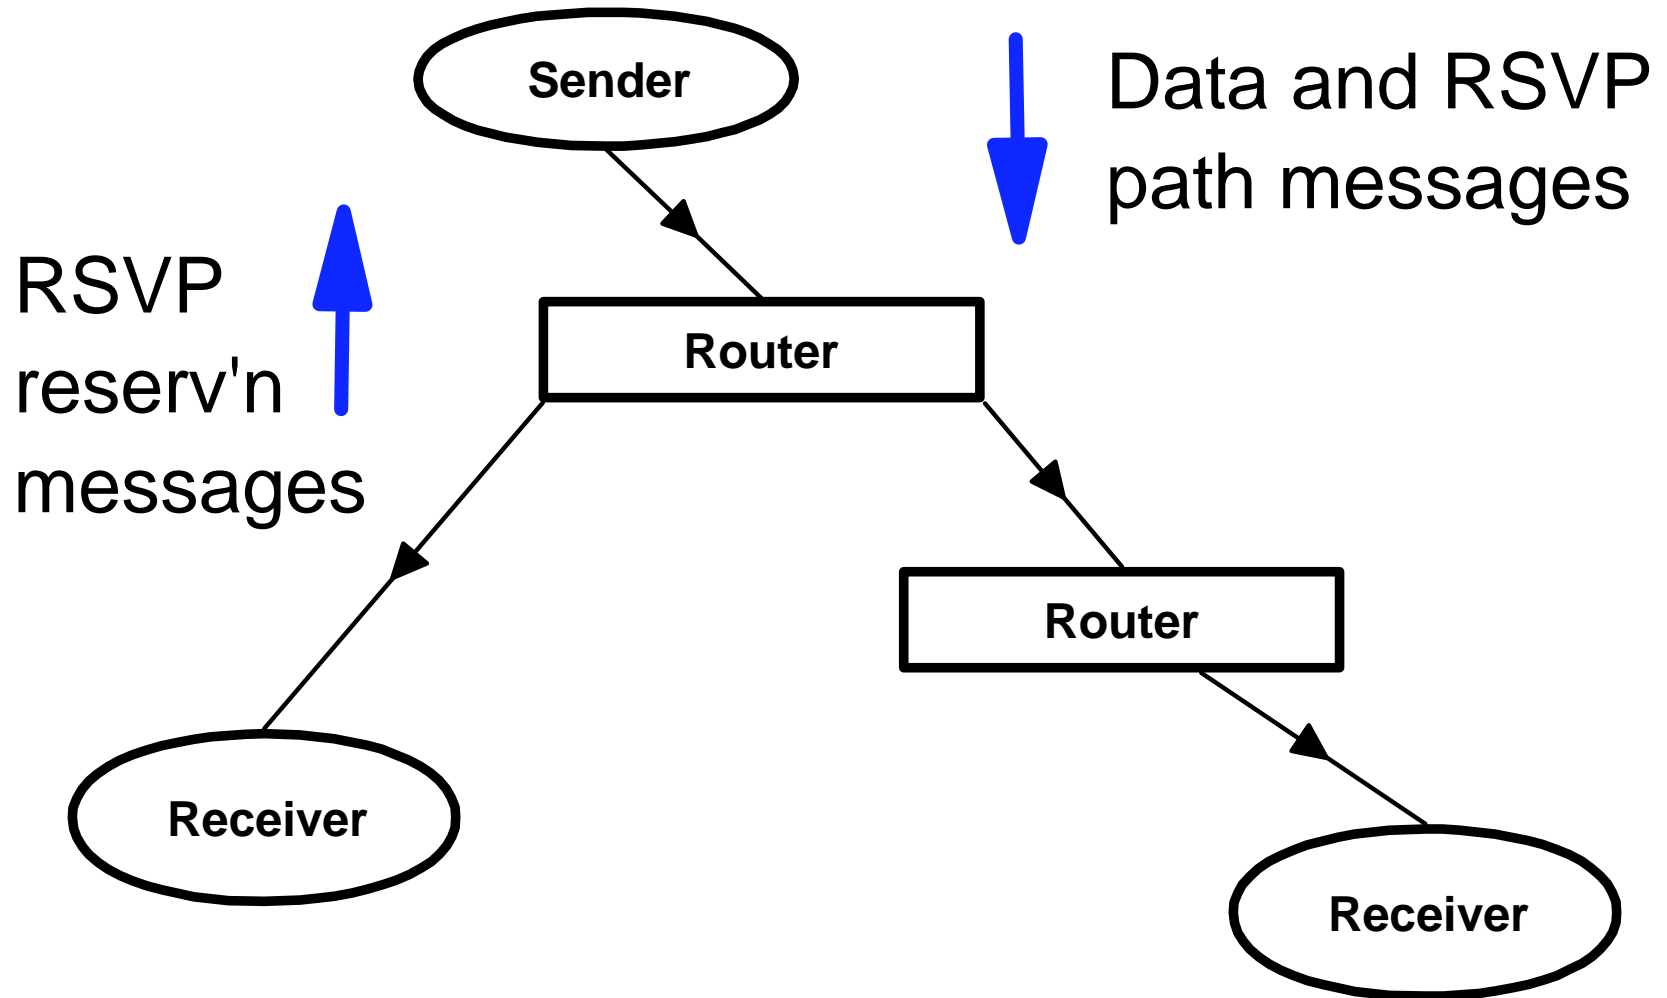
Final summary

Service class per microflow

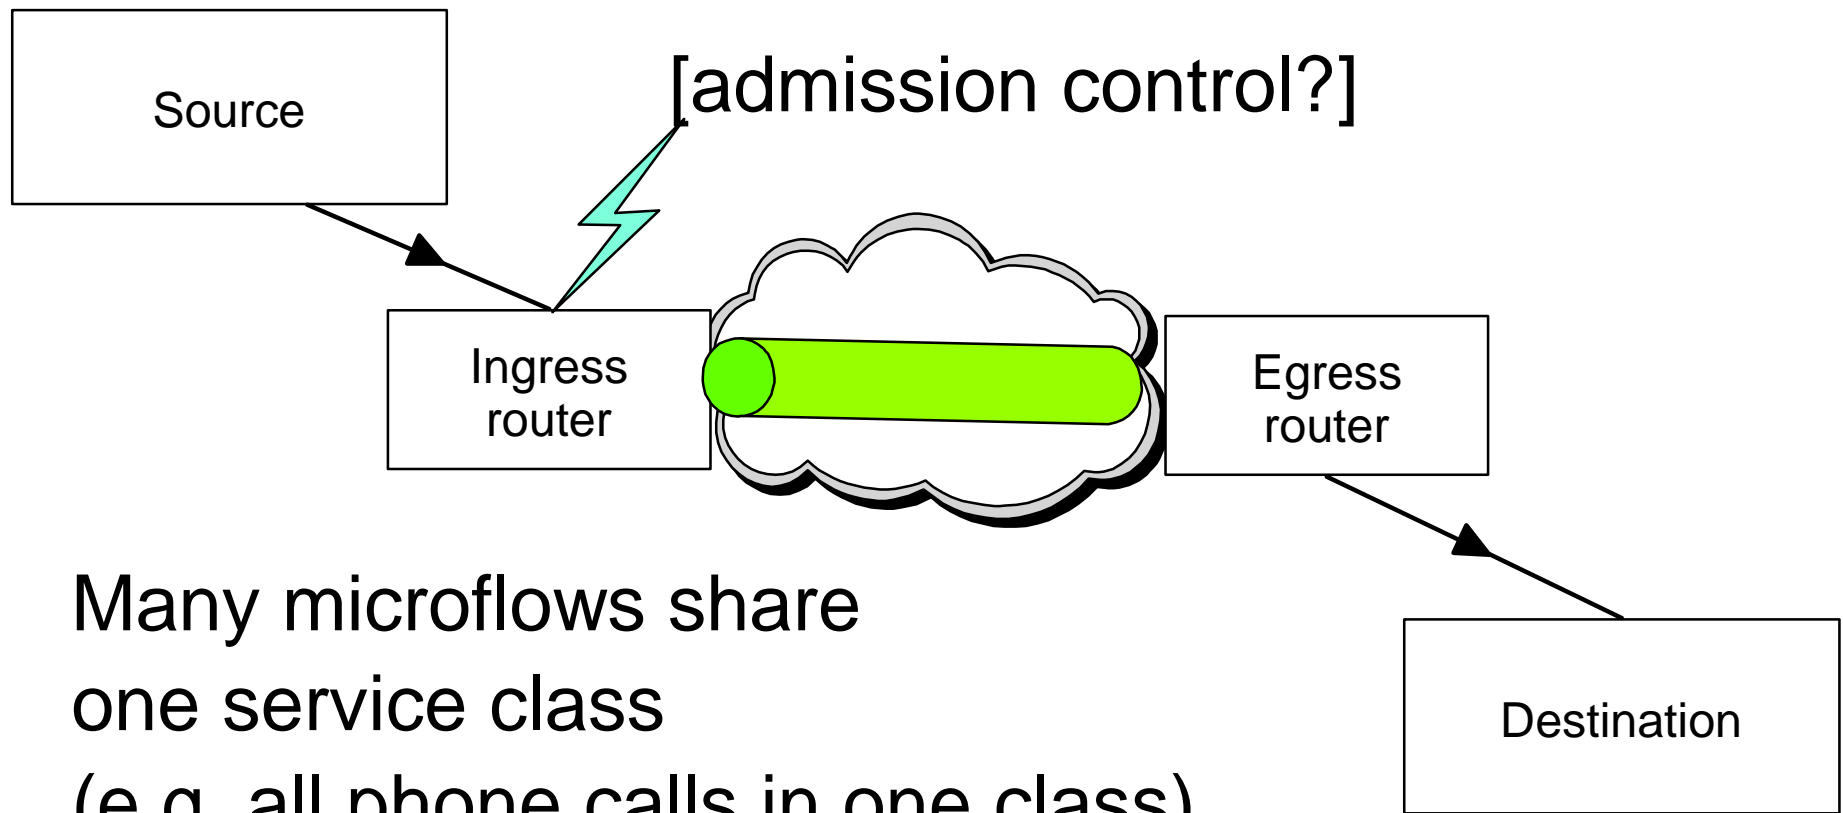


Millions of microflows =
millions of service classes
(e.g. one per phone call)

RSVP scenario




Pre-existing service classes



Basic idea

- ▶ Get rid of the per-flow state and complex processing by marking each packet with a simple flag indicating how to treat it.
- ▶ Make it a one-way process so that asymmetric routes don't matter



A policy management system

