Quality-of-Service Architectures for the Internet

QoS Service Architectures for the Internet

- Two QoS architectures have been defined for Internet.
  - **Integrated Services (IntServ)**
    - Proposed in 1994
    - Per-flow Quality of Service
    - Resource reservation/admission control
    - Can support delay guarantees
  - **Differentiated Services (DiffServ)**
    - Proposed in 1998
    - Class-based QoS
    - Resource reservation not always needed

Integrated Services

IntServ specifies two types of services:

Guaranteed Service

- Guaranteed bandwidth
- End-to-end delay bounds
- No loss due to buffer overflows

Controlled Load Service

- Provides a service that is equivalent to a best effort service in a lightly loaded netework
  - Low loss
  - Low delay
  - No absolute guarantees

Integrated Services

	in IntServ
1. At network entrance: Policing and Shaping →	FlowSpec (TSpec,RSpec)
2. Somewhere in the network: Admission Control →	Distributed
3. At switches: Classification, Scheduling →	Weighted Fair Queuing or other rate-based algorithm
4. Between hosts and routers: Signaling →	RSVP

Resource ReSerVation Protocol (RSVP)

- RSVP is a signaling protocol that enables senders, receivers, and routers of unicast or multicast sessions to communicate with each other for setting up state to support a service
  - **Receiver-driven**
    - Resource reservation is initiated by receivers
  - **Unicast and multicast** sessions
  - **Soft-state:** state information of RSVP must be periodically refreshed

- Separate mechanisms required for authorization, authentication, and charging

RSVP Functional Diagram

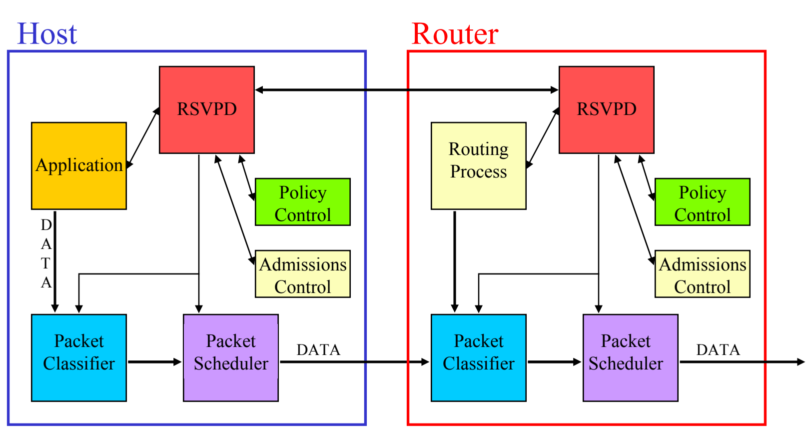

Source: *Gordon Chaffee, UC Berkeley*

Resource Reservation

- Senders advertise using PATH message
- Receivers reserve using RESV message
  - Flowspec + filterspec + policy data
  - Travels upstream in reverse direction of Path message
- Merging of reservations
- Sender/receiver notified of changes

Source: *Gordon Chaffee, UC Berkeley*

RSVP UDP Reservation (1)

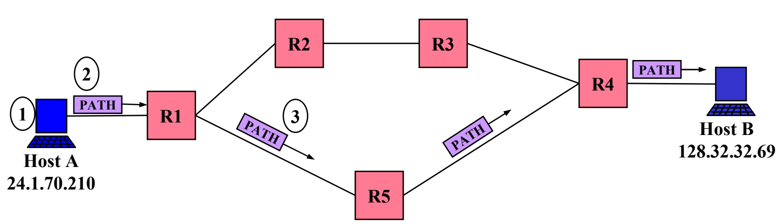

1. An application on Host A creates a session, 128.32.32.69/4078, by communicating with the RSVP daemon on Host A.

2. The Host A RSVP daemon generates a PATH message that is sent to the next hop RSVP router, R1, in the direction of the session address, 128.32.32.69.

3. The PATH message follows the next hop path through R5 and R4 until it gets to Host B. Each router on the path creates soft session state with the reservation parameters.

Source: *Gordon Chaffee, UC Berkeley*

RSVP UDP Reservation (2)

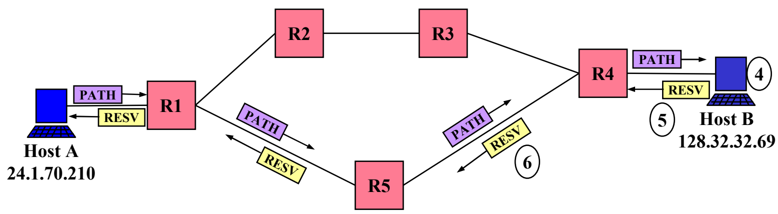

4. An application on Host B communicates with the local RSVP daemon and asks for a reservation in session 128.32.32.69/4078. The daemon checks for and finds existing session state.

5. The Host B RSVP daemon generates a RESV message that is sent to the next hop RSVP router, R4, in the direction of the source address, 24.1.70.210.

6. The RESV message continues to follow the next hop path through R5 and R1 until it gets to Host A. Each router on the path makes a resource reservation.

Source: *Gordon Chaffee, UC Berkeley*

RSVP Flowspecs

Sender TSpec, Controlled Load Flowspec

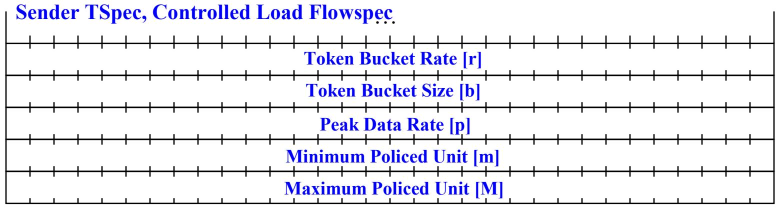

Guaranteed Flowspec

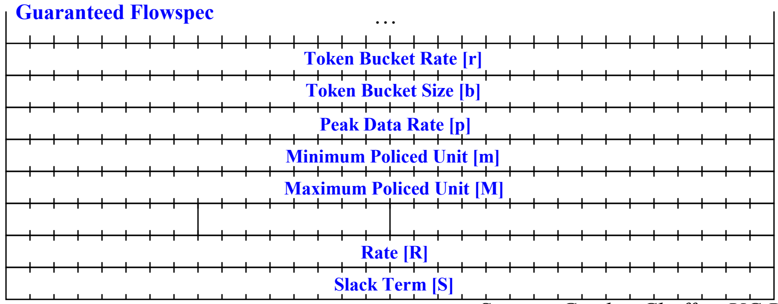

Source: *Gordon Chaffee, UC Berkeley*

Reservation Merging

Reservations merge as they travel up tree.

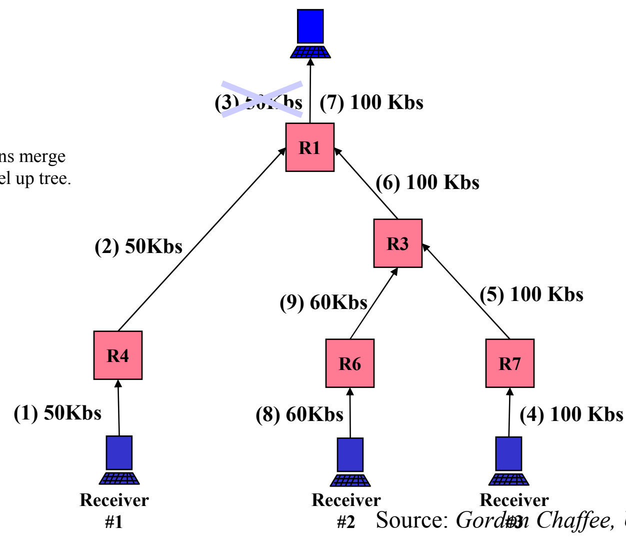

Source: *Gordon Chaffee, UC Berkeley*

Fair Queueing

- **Processor Sharing: Each flow receives a "fair share" of bandwidth**
  - r_i – flow arrival rate
  - f –fair rate of a link
  - C – link capacity
- If link congested, compute fair share f such that $\sum_i \min(r_i, f) = C$

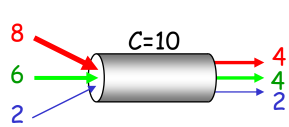

f = 4:
min(8, 4) = 4
min(6, 4) = 4
min(2, 4) = 2

Generalized Processor Sharing

- Associate a weight w_i with each flow i
- If link congested, $\sum_i \min(r_i, f \times w_i) = C$

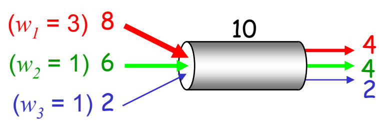

f = 2:
min(8, 2*3) = 6
min(6, 2*1) = 2
min(2, 2*1) = 2

- Flow i has a rate guarantee g_i of $g_i \geq \frac{w_i}{\sum_k w_k} C$

Generalized Processor Sharing (GPS)

- Red session has packets backlogged between time 0 and 10
- Other sessions have packets continuously backlogged

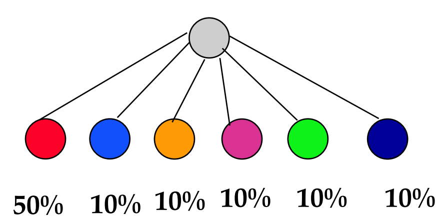

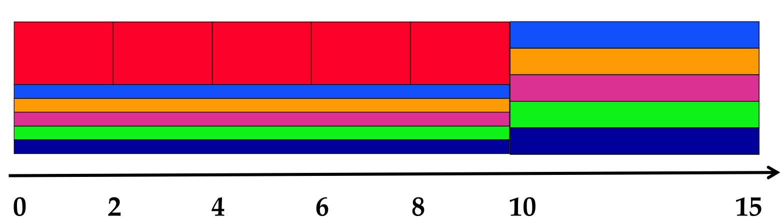

Weighted Fair Queueing

- GPS assumes that traffic behaves like a fluid flow system, that can transmit one bit at a time

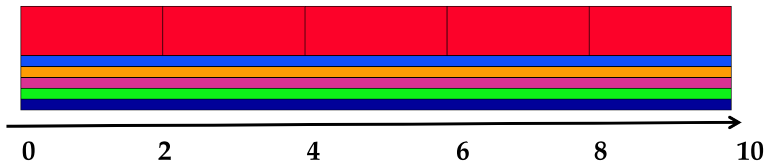

- In practice, GPS can only be approximated
- **Weighted Fair Queueing:** Select the first packet that finishes in GPS

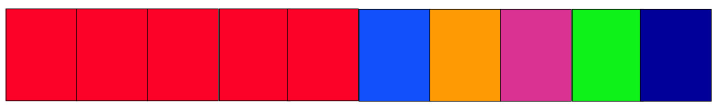

- Can be proven: WFQ can approximate GPS closely

Fluid Flow System: Example 1

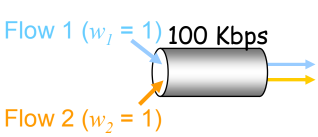

	Packet Size (bits)	Packet inter-arrival time (ms)	Rate (Kbps)
Flow 1	1000	10	100
Flow 2	500	10	50

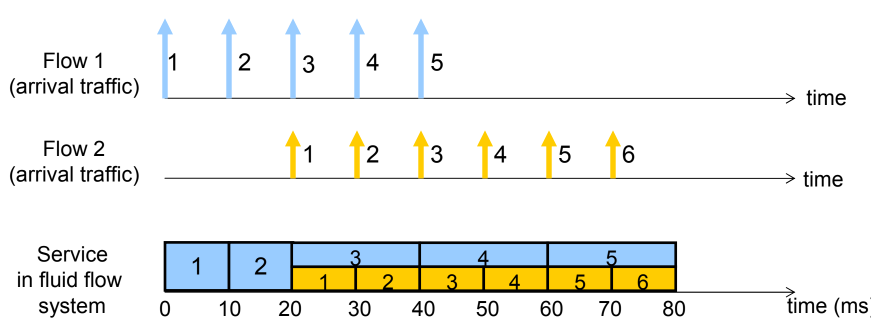

Slide from Ion Stoica

Packet System: Example 1

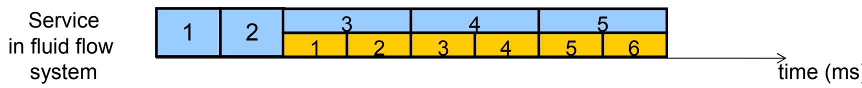

- Select the first packet that finishes in the fluid flow system

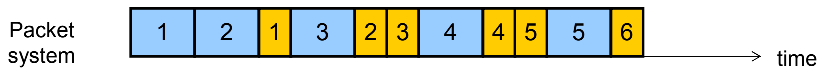

- **WFQ is not easy to implement.**
- Need to keep track of the finishing times of all packets in the GPS system
- Solution: Maintain a sorted queue that sorts packets according to finishing times in the GPS system
- Admission control in guaranteed service is based on maximum delay in a WFQ scheduler with guaranteed rate R and burst size b.

Slide from Ion Stoica

Summary of IntServ

– **Advantages:**
  - Strong guarantees (bounded delays)

– **Disadvantages:**
  - Requires that all routers implement IntServ
  - Scalability concerns since routers must maintain state information
  - Charging and authentication of reservations must be solved
  - Interdomain issues are difficult to resolve

DiffServ

- **Motivation**:
  – The Integrated Services (IntServ) model is not scalable since it requires per-flow state in each node

Goal:

- Push complexity to the network edge and keep network core simple
- Avoid per-flow state within the network as much as possible

Differentiated Service Mechanisms

- Definitions:
  – Mechanisms that allow providers to allocate different levels of service to different users of the Internet

  – broad view: Any mechanism that treats different users differently, including signaling (RSVP), per-session scheduling, etc.
  – Internet context: Simple and lightweight mechanisms that do not depend entirely on per-flow reservation

Components of Differentiated Services

(1) Service profile **between user and network** defines commitment of the network to the user

(2) Aggregate traffic from each user is policed at the network entrance according to profile

(3) Node behavior: network nodes implement a variety of forwarding, scheduling, buffer management techniques

(4) Bits in packet header trigger action at nodes

Common to Most Proposed Services

- Traffic marking (in-profile, out-profile) and enforcement is done only at network boundaries
- Inside the network: Only differentiate a few service classes, based on marking of the packets

Operational Model

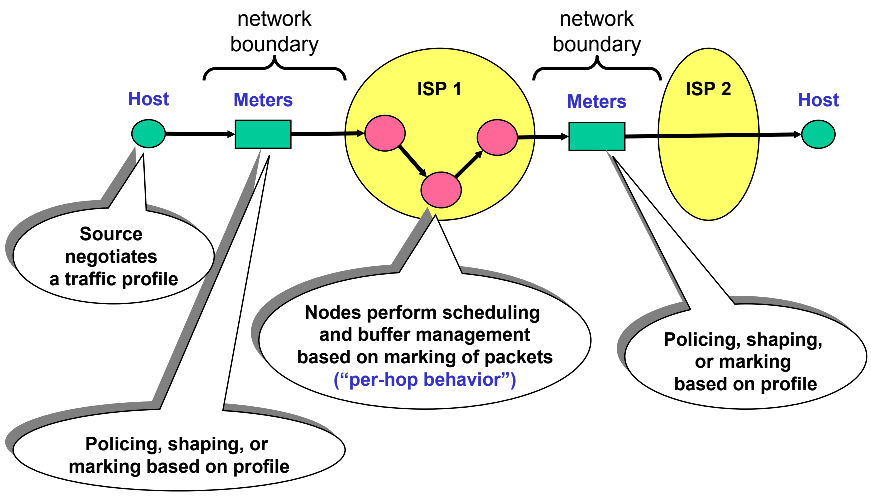

Aspects of a Differentiated Service

(1) Semantics of the service:
Which service is given to in-profile traffic of a user?

(2) Spatial Granularity:
Is the profile applied to a single destination, a subset of destinations, or all destinations?

(3) Assurance Level:
What is the level of certainty that an in-profile packet will be delivered?

DiffServ Services

- Two services defined:

- Assured Forwarding (AF)
  - customers sign service agreements with ISPs
  - Edge routers mark packets as being “in” or “out” of profile
  - core routers run RIO: RED with in/out
  - Distinguishes different classes:

- Expedited Forwarding (EF)
  - Hard guarantee on the delay and delay variations

Assured Forwarding - 1

- User defines traffic profile (token bucket)
- Profile meter at network entrance tag packets as *in-profile* or *out-profile*
- *Service guarantee: In-profile packets are unlikely to be dropped*
- *Out-profile* packets have higher drop preference at routers

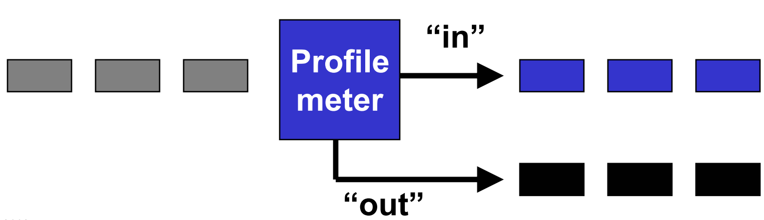

Assured Forwarding - 2

Mechanisms

- Mechanisms Needed :
  - Dropping Mechanisms at routers
  - Mechanism for tagging packets ("Meters")
  - Method to classify packets

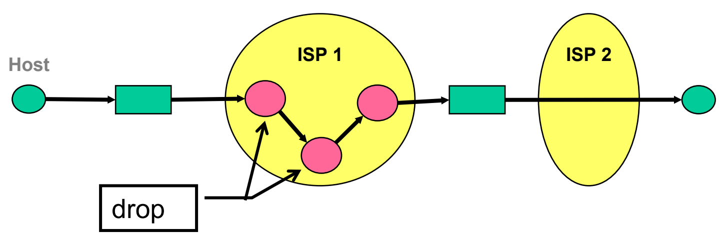

Assured Forwarding - 3

RIO

- Routers have different dropping mechanism:
  RIO = RED with `in' and `out'
- Routers do not perform separate queueing
- RED (Random Early Detection):

  When the avgerage queue size exceeds a threshold drop each packet with a certain probability (P_{drop})

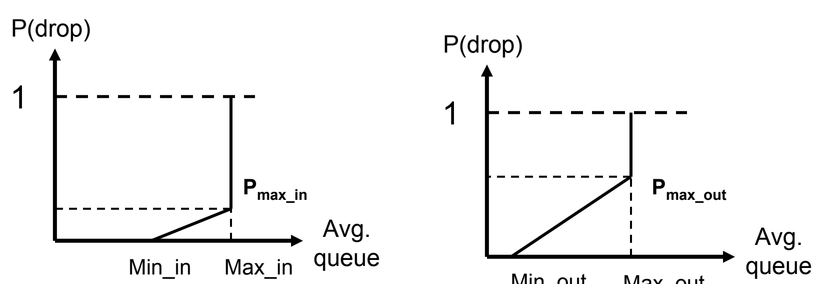

Expedited Forwarding - 1

- EF traffic must be served at a configured rate of R or faster, independent of the load
- Service is equivalent to a "virtual leased line"
- Routers have two priority levels (premium and best effort)
- Admission Control via Bandwidth Brokers

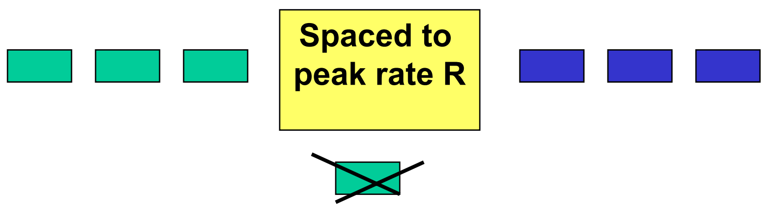

Expedited Forwarding - - 2
Admission Control

- "Bandwidth Brokers" perform admission control at ingress router
- Only the ingress router differentiates flows

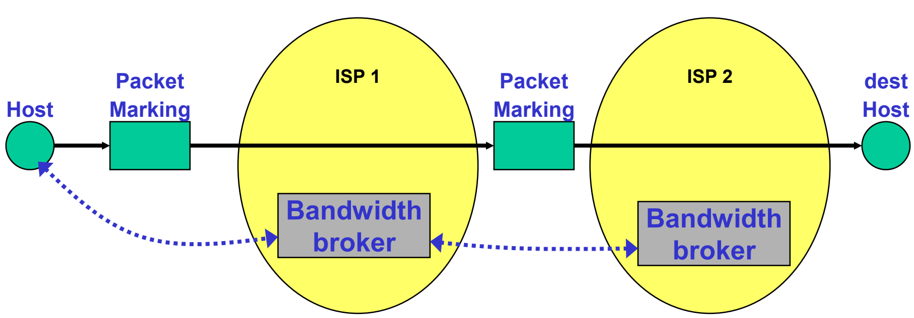

Summary of DiffServ

– **Advantages:**
  - No per-flow processing in network core
  - Per-flow processing only at the network edge
  - Simpler to implement than IntServ
  - No signaling protocol

– **Disadvantages:**
  - AF has weaker service guarantees
  - EF service raises same issues with charging and authentication as IntServ services

Leaf Router: Input

(Leaf router = the router closest to the source)

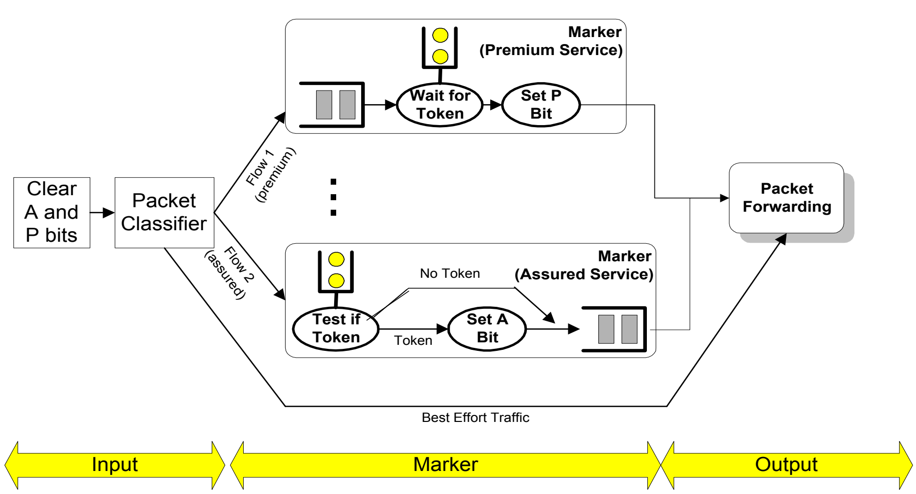

Border Router: Input

(Border router = ingress router of a network)

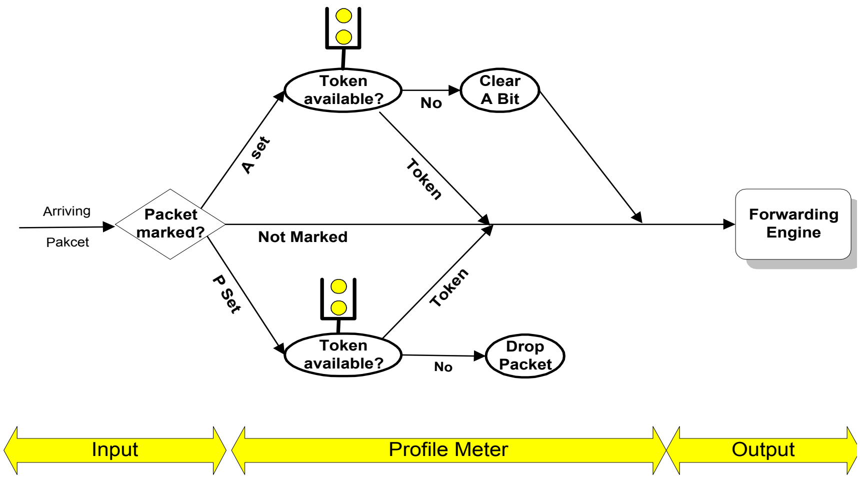

Router: Output