Quality of Service (QoS) - DiffServ

EE 122: Intro to Communication Networks

Fall 2007 (WF 4-5:30 in **Cory 277**)

Vern Paxson

TAs: Lisa Fowler, Daniel Killebrew & Jorge Ortiz

http://inst.eecs.berkeley.edu/~ee122/

Materials with thanks to Jennifer Rexford, Ion Stoica, and colleagues at Princeton and UC Berkeley

Our Story So Far

- QoS = attaining some sort of reliable performance from the network
- Max-Min Fairness as concept for allocating capacity across a set of flows
- Weighted Fair Queuing as way to attain Max-Min Fairness
- Token Bucket as way to describe bounds on burstiness of a flow's packet's arriving at a queue
- Integrated Services (IntServ) as means by which flows can
  - Describe burstiness using Token Bucket descriptors
  - Set up soft-state reservations end-to-end
  - Entails admission control decision
    - Answer could be "no, you don't get it"

Problems with IntServ

- Scalability: per-flow state & classification
  - Aggregation/encapsulation techniques can help
  - Can overprovision big links, per-flow ok on small links
  - Scalability can be fixed - but no second chance
- Economic arrangements:
  - Need sophisticated settlements between ISPs
  - Contemporary settlements are primitive
    - Unidirectional, or barter
- User charging mechanisms: need QoS pricing
  - On a fine-grained basis

Differentiated Services (DiffServ)

- Give some traffic better treatment than other
  - Application requirements: interactive vs. bulk transfer
  - Economic arrangements: first-class versus coach
- What kind of better service could you give?
  - Fewer drops
  - Lower delay
  - Lower delay variation (*jitter*)
- How to know which packets get better service?
  - Bits in packet header
- Deals with traffic in aggregate
  - Provides weaker services
  - But much more scalable

Diffserv Architecture

- Ingress routers - entrance to a DiffServ domain
  - Police or shape traffic
  - Set Differentiated Service Code Point (**DSCP**) in IP header
- Core routers
  - Implement Per Hop Behavior (**PHB**) for each DSCP
  - Process packets based on DSCP

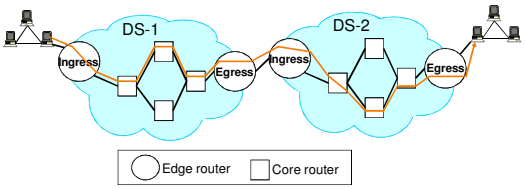

Differentiated Service (DS) Field

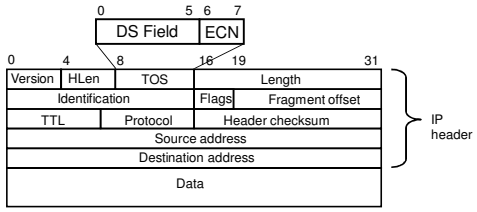

- DS field encodes Per-Hop Behavior (PHB)
  - E.g., Expedited Forwarding (all packets receive minimal delay & loss)
  - E.g., Assured Forwarding (packets marked with low/high drop probabilities)

Comparison to Best-Effort & Intserv

	Best-Effort	Diffserv	Intserv
Service	Connectivity No isolation No guarantees	Per aggregate isolation Per aggregate guarantee	Per flow isolation Per flow guarantee
Service scope	End-to-end	Domain	End-to-end
Complexity	No setup	Long term setup	Per flow steup
Scalability	Highly scalable (nodes maintain only routing state)	Scalable (edge routers maintain per aggregate state; core routers per class state)	Not scalable (each router maintains per flow state)

7

Discussion: Limited QoS Deployment

- End-to-end QoS across multiple providers/domains is **not** available today
- Issue #1: complexity of payment
  - Requires payment system among multiple parties
    - And agreement on what constitutes service
  - Diffserv tries to structure this as series of bilateral agreements …
    - … but lessens likelihood of end-to-end service
    - Architecture includes notion of "Bandwidth Broker" for end-to-end provisioning
      - Solid design has proved elusive
  - Need infrastructure for metering/billing end user

8

Limited QoS Deployment, con't

- Issue #2: prevalence of overprovisioning
  - Within a large ISP, links tend to have plenty of headroom
  - Inter-ISP links are **not** over provisioned, however
- Is overprovisioning enough?
  - If so, is this only because access links are slow?
  - What about Korea, Japan, and other countries with fast access links?
  - Disconnect: ISPs overprovision, users get bad service
- Key difference: intra-ISP vs. general end-to-end

9

Summary

- Basic mechanism for achieving better-than-best-effort performance: **scheduling**
  - Multiple queues allow priority service
  - Fair queuing provides isolation between flows
- But: still need end-to-end mechanisms
  - **Reservations** & **admission control**
  - Descriptions of bursty traffic: **token buckets**
- IntServ provides per-flow performance guarantees
  - But lacks scalability
- DiffServ provides per-**aggregate** tiers of relative perf.
  - Scalable, but not as powerful
- Neither is generally available end-to-end today
- ISPs manipulating what services receive what performance raises issues of: network neutrality

10