# Quality of Service in the Internet:

# Fact, Fiction, or Compromise?

*Paul Ferguson*,
Cisco Systems, Inc.

*Geoff Huston*,
Telstra

# What is the Expectation?

- Today’s Internet is plagued by sporadic poor performance
  - This is getting worse, not better

# Customers want….

- Customers want access to an Internet service which provides ***consistent & predictable high quality service levels***

# QoS

- Network mechanisms intended to meet this demand are categorized within the broad domain of ***Quality of Service***

***But can the Internet deliver?***

# QoS is not…

- QoS is ***not*** a tool to compensate for inadequacies elsewhere in the network
  - Massive over-subscription
  - Horrible congestion situations
  - Poor network design

# QoS is not…

QoS is not magic

- QoS will not alter the speed of light
  - On an unloaded network, QoS mechanisms will not make the network any faster
    - Indeed, it could make it slightly worse!
- QoS does not create nonexistent bandwidth
  - Elevating the amount of resources available to one class of traffic decreases the amount available for other classes of traffic
- QoS cannot offer cures for a poorly performing network

# QoS is…

QoS is unfair damage control

- QoS mechanisms attempt to preferentially allocate resources to predetermined classes of traffic, when the resource itself is under contention
- Resource management only comes into play when the resource is under contention by multiple customers or traffic flows
  - Resource management is irrelevant when the resource is idle, or not an object of contention

# The QoS Margin is small

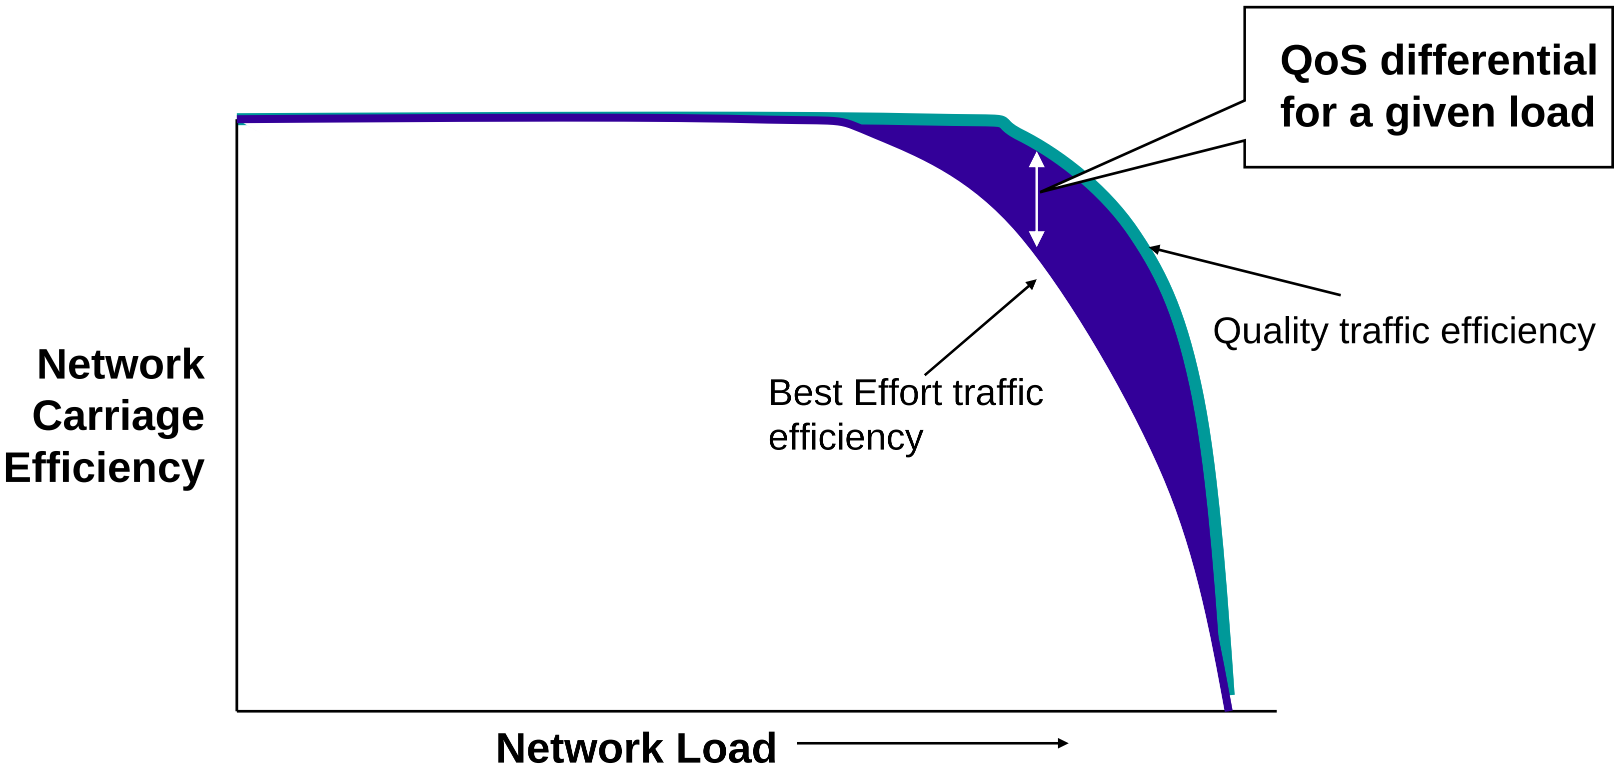

# QoS is…

QoS is relative, not absolute

- QoS actively discriminates between preferred and non-preferred classes of traffic at those times when the network is under load (congested)
- Qos is the relative difference in service quality between the two generic traffic classes
  - If every client used QoS, then the net result is a zero sum gain

# QoS is…

QoS is intentionally elitist and unfair

- The QoS relative difference will be greatest when the preferred traffic class is a small volume compared to the non-preferred class
- QoS preferential services will probably be offered at a considerable price premium, to ensure that quality differentiation is highly visible for a small traffic component

# Expectation setting

- QoS does not work for all types of traffic
  - TCP flows use a ‘network clock’ to adapt the transfer rate to the current network condition
    - This ‘dynamic equilibrium’ takes time to establish
    - Short Flows do not adapt to full speed in time
  - UDP flows use external signal clocking
    - UDP cannot transfer faster than the external data clock

# What is *Quality*?

- *Quality* cannot be measured on an entire network.
  - Flow bandwidth is dependant on the chosen transit path.
  - Congestion conditions are a localized event.
  - Quality metrics degrade for those flows which transit the congested location.
- *Quality* can be measured on an end-to-end traffic flow, at a particular time.

# *Quality* metrics

- Quality metrics are amplified by network load.
  - **Delay** increases due to increased queue holding times.
  - **Jitter** increases due to chaotic load patterns.
  - **Bandwidth** decreases due to increased competition for access.
  - **Reliability** decreases due to queue overflow, causing packet loss.
- Quality differentiation is only highly visible under high network path load.

# Approaches

# Network State

- Per flow traffic management to undertake one of more of the following service commitments:
  - Place a preset bound on jitter.
  - Limits delay to a maximal queuing threshold.
  - Limit packet loss to a preset threshold.
  - Delivers a service guarantee to a preset bandwidth rate.
  - Deliver a service commitment to a controlled load profile.
- Challenging to implement in a large network.
- Relatively easy to measure success in meeting the objective.

# RSVP

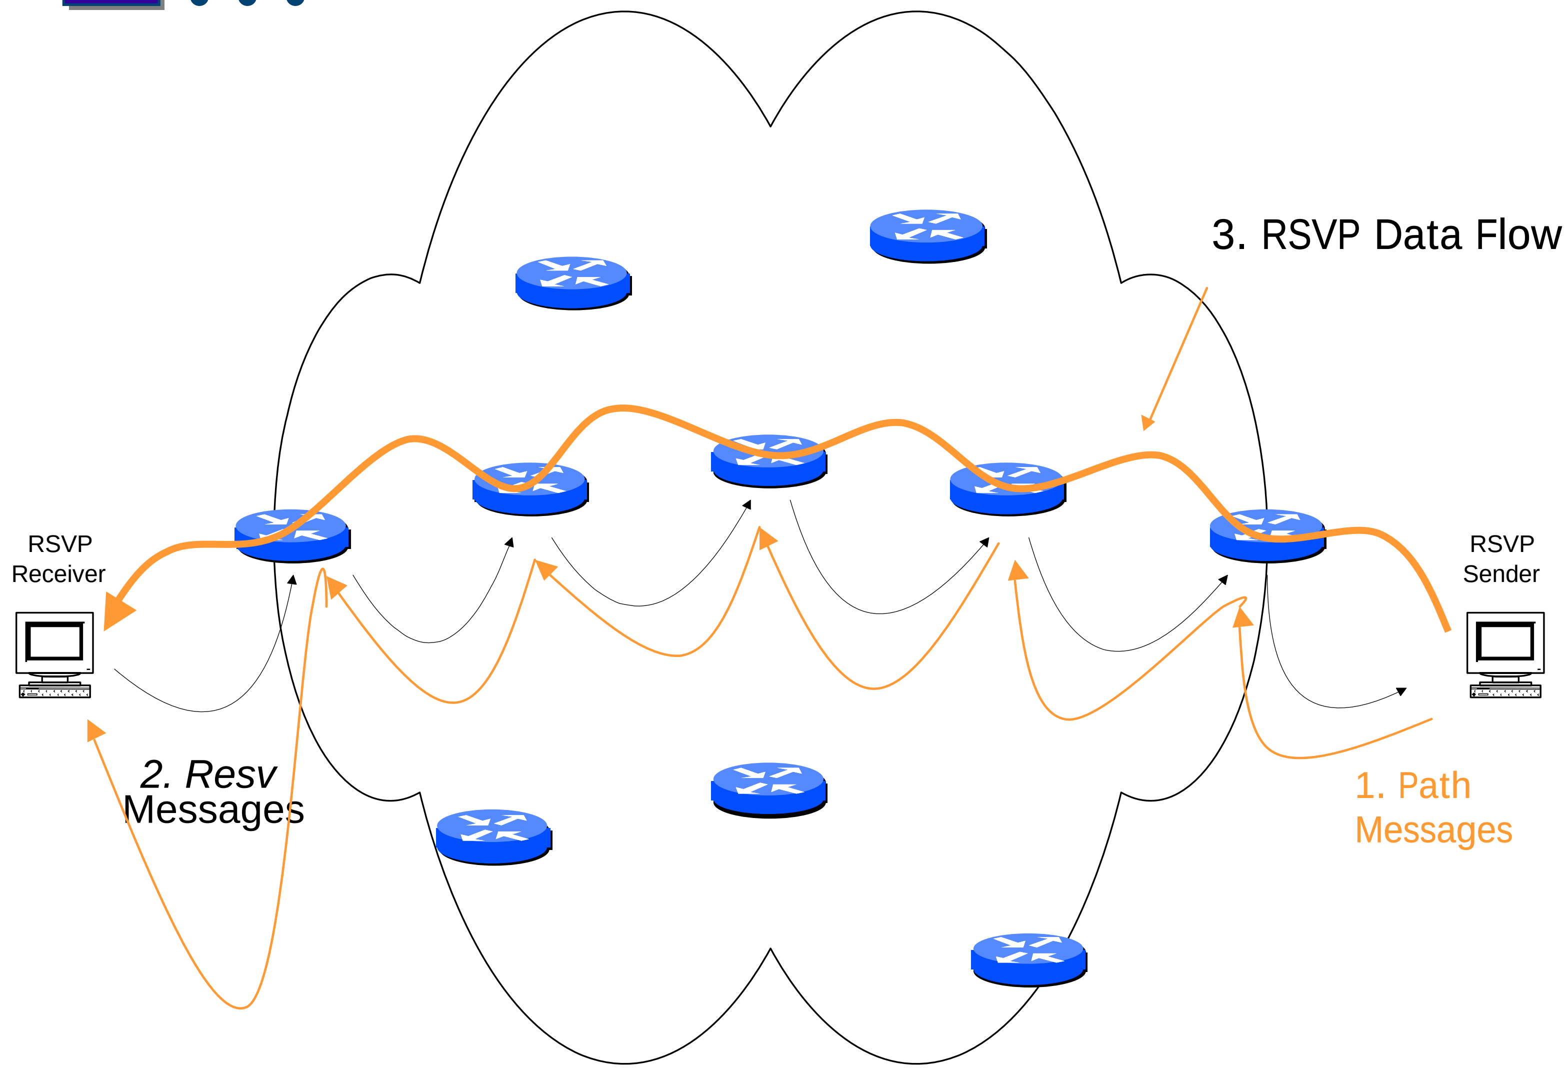
3. RSVP Data Flow
RSVP
Receiver
RSVP
Sender
2. Resv
Messages
1. Path
Messages

# Network State and the Internet

- Integrated Services requires the imposition of flow-based dynamic state onto network routers in order to meet the stringent requirements of a service guarantee for a flow.
- Such mechanisms do not readily scale to the size of the Internet.

# Packet State

- Active differentiation of network traffic to provide a *better than best effort* performance for a defined traffic flow, as measured by one of more of:
  - Packet jitter
  - Packet loss
  - Packet delay
  - Available peak flow rate
- Implementable within a large network.
- Relatively difficult to measure success in providing service differentiation.

# Packet State and the Internet

- Differentiated Services can be implemented through the deployment of differentiation router mechanisms triggered by per-packet flags, preserving a stateless network architecture within the network core.
- Such mechanisms offer some confidence to scale to hundreds of millions of flows per second within the core of a large Internet

# Mechanisms

# Virtual Circuits

- Segmented bandwidth resource for QoS states:
  - Virtual circuits & statistical muxing (e.g. ATM, Frame Relay) with ingress traffic shaping
  - RSVP admission control & reservation state
- Segmentation mechanisms by themselves are unrealistic in a large scale heterogeneous Internet which uses end-to-end flow control.

# QoS Paths

- Alternate path selection
  - Alternative physical paths
    - E.g., cable and satellite paths
  - QoS Routing v. administrative path selection
- Must be managed with care
- Can lead to performance instability
- Prone to inefficient use of transmission
- May not support end-to-end path selection

# QoS Paths

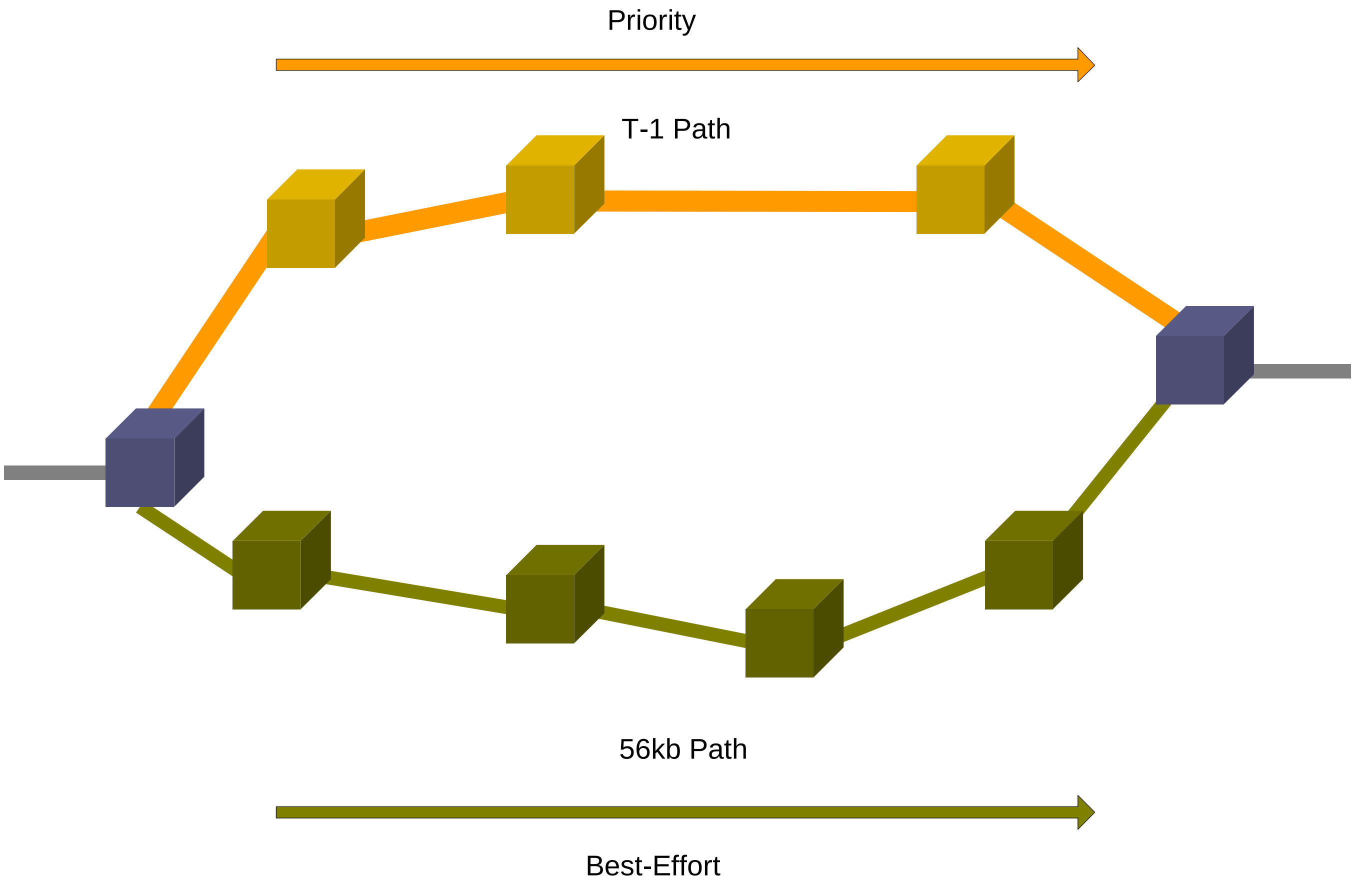

# QoS Service Mechanism

- Admission traffic profile filter
  - In-Profile traffic has elevated QoS, out-of-profile uses non-QoS

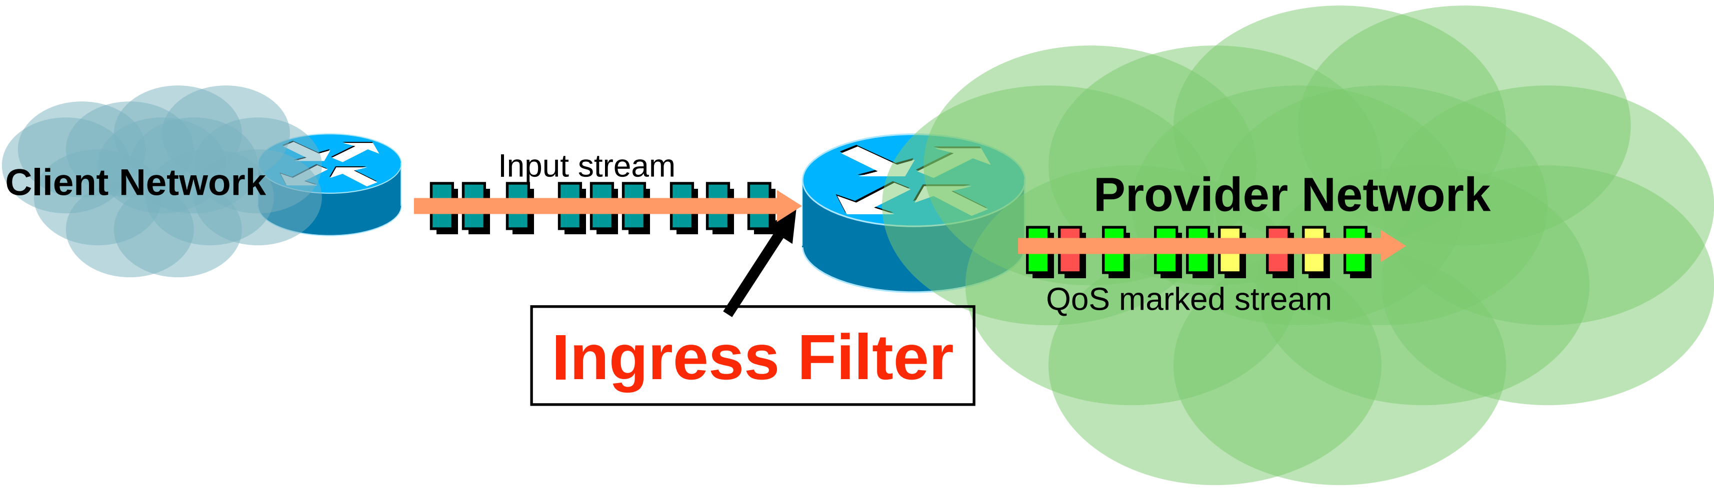

# QoS per packet indicators

- Explicit per packet signaling of:
  - Precedence indication (delay)
  - Discard indication (reliability)

  As an indication of preference for varying levels of best effort
- Routers configured to react to per packet indicators through differentiated packet scheduling and packet discard behaviours
- This is deployable - today

# QoS WFQ Precedence

- Schedule traffic in the sequence such that a equivalent weighted bit-wise scheduling would deliver the same order of trailing bits of each packet

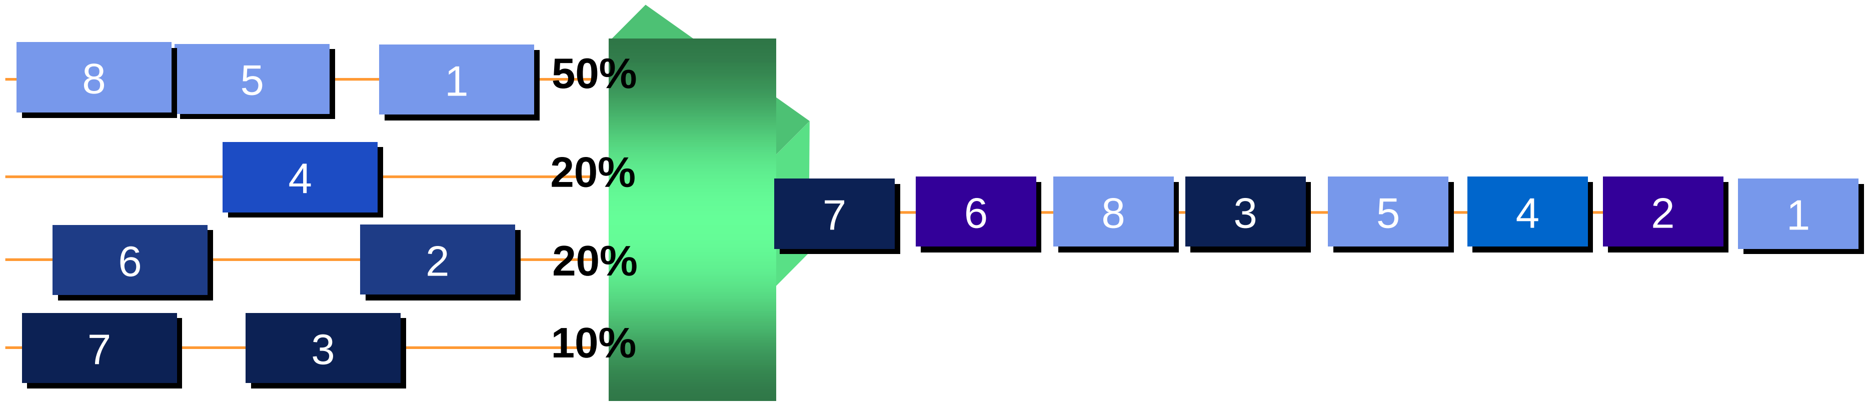

# Considerations

# Pervasive homogeneity - Not in the Internet!

- Reliance on link-layer mechanisms to provide QoS assumes pervasive end-to-end, desktop-to-desktop, homogenous link-layer connectivity
- This is simply not a realistic assumption for the Internet

# State and Scale

- To undertake firm commitments in the form of per-flow carriage guarantees requires network-level state to be maintained in the routers
- State adds to the network cost
- State is a scaling issue

- Wide-scale RSVP deployment will not scale in the Internet
  - (See: RFC2208, RSVP Applicability Statement).

# Network Layer Tools

- Traffic shaping and admission control
- Ingress IP packet marking for both delay indication and discard preference
- Weighted Preferential Scheduling algorithms
- Preferential packet discard algorithms

  (e.g. Weighted RED, RIO)
- End result: Varying levels of service under load

- Of Course: No congestion, no problem

# QoS Implementation Considerations

- Complexity: If your support staff can't figure it out, it is arguably self-defeating
- Delicate balance between good network design and engineering and QoS damage control

# Yet to be Resolved

- Long held adaptive flows are susceptible to network layer shaping
- Short held flows (WWW transactions)
  - Are not very susceptible to network layer shaping
- UDP flow management
  - Unicast flow control model
  - Multicast flow control model
- Inter-Provider semantics for differentiated services multi-provider QoS support

# Unanswered Questions

- How does the provider measure QoS?
- How does the customer measure QoS?
- How do you tariff, account, and bill for QoS?
- How will QoS work in a heterogeneous Internet?
  - QoS across transit administrative domains which may not participate or use different QoS mechanisms?

# Summary

- There is no magic QoS bullet

  Sorry
- There are no absolute guarantees in the Internet

  Sorry
- There is possibly a “middle ground” somewhere between traditional single level best effort and guaranteed customized services

# References

- *Differential Services in the Internet*
  http://diffserv.lcs.mit.edu/
- *Quality of Service: Delivering QoS in the Internet and the Corporate Network*
  http://www.wiley.com/compbooks/ferguson/

# Questions?

# Thank you.

**Paul Ferguson**
**Consulting Engineer, Internet Architecture**
**ferguson@cisco.com**

**Geoff Huston**
**Technology Manager**
**Telstra Internet**
**gih@telstra.net**