Cisco.com

Is QoS Necessary?

QoS Mechanisms vs. Bandwidth Provisioning

Bruce Davie
Cisco Fellow

Agenda

- Introduction
- Overview of QoS Mechanisms
- Evaluation of mechanisms vs. bandwidth
- Deployment experience
- Conclusions

Defining QoS: Application perspective

- QoS is defined in terms of loss rate, throughput, delay
  - Throughput & loss rate tightly coupled for many apps (e.g. TCP-based)
  - Applications that need timely delivery of packets (e.g. Voice) sensitive to *delay distribution*
    - e.g. 99th percentile delay
    - “Late” packets typically useless – “early” packets can be buffered

QoS for multimedia applications

Cisco.com

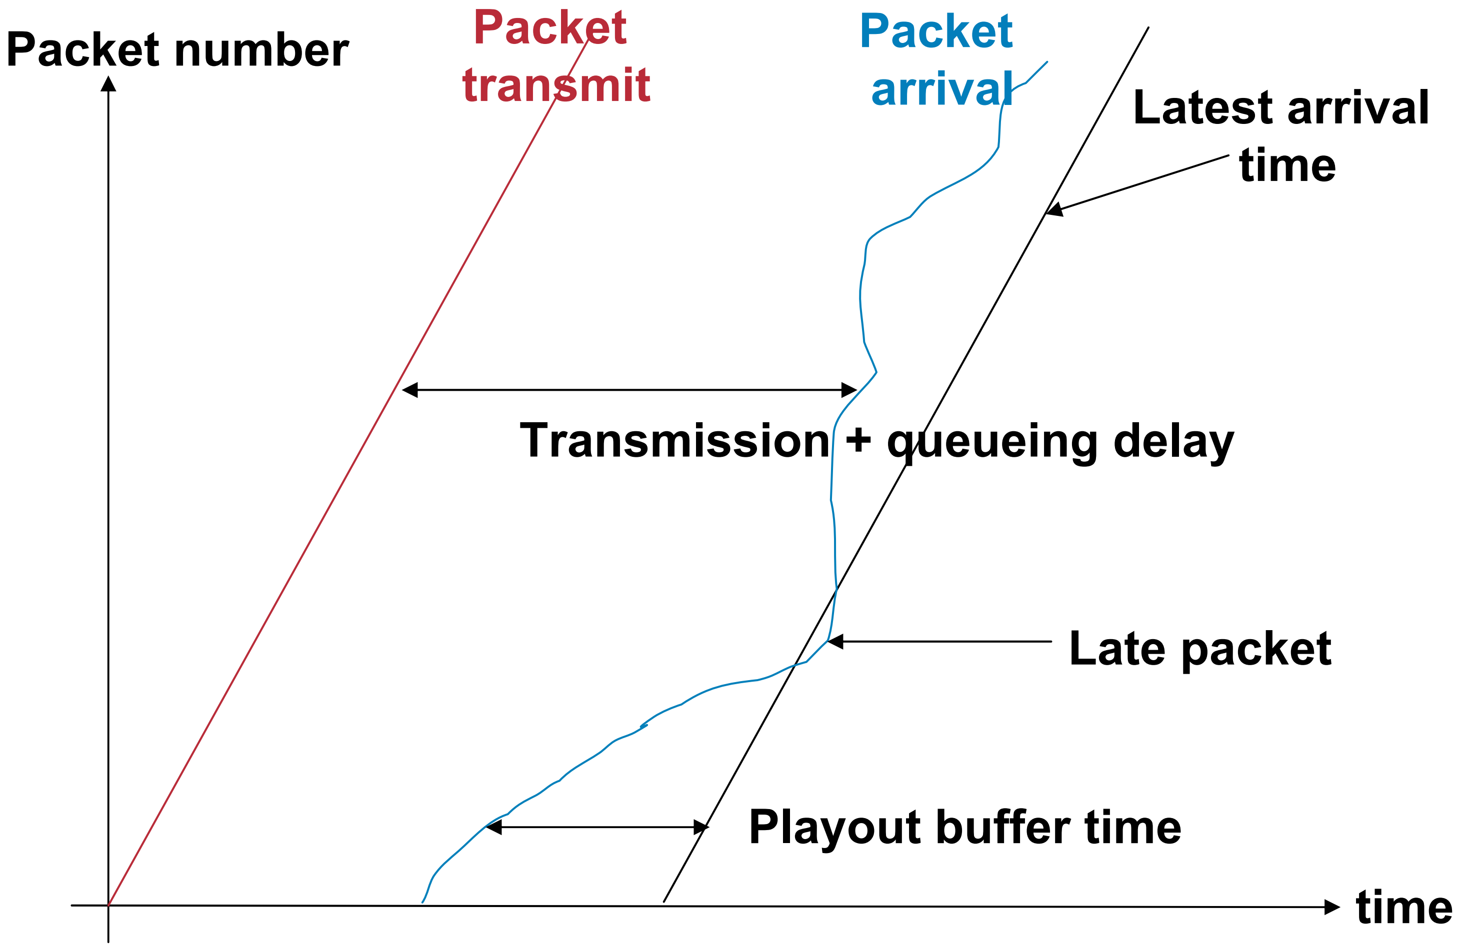

Defining QoS: Mechanism perspective

- **If bandwidth is plentiful, packets are never delayed or lost, so QoS needs of apps are met**
- **QoS often defined in terms of mechanisms deployed when bandwidth isn't plentiful**
  - **Mechanisms include classification, policing, shaping, marking, queueing, differential dropping**

Queuing Theory and QoS

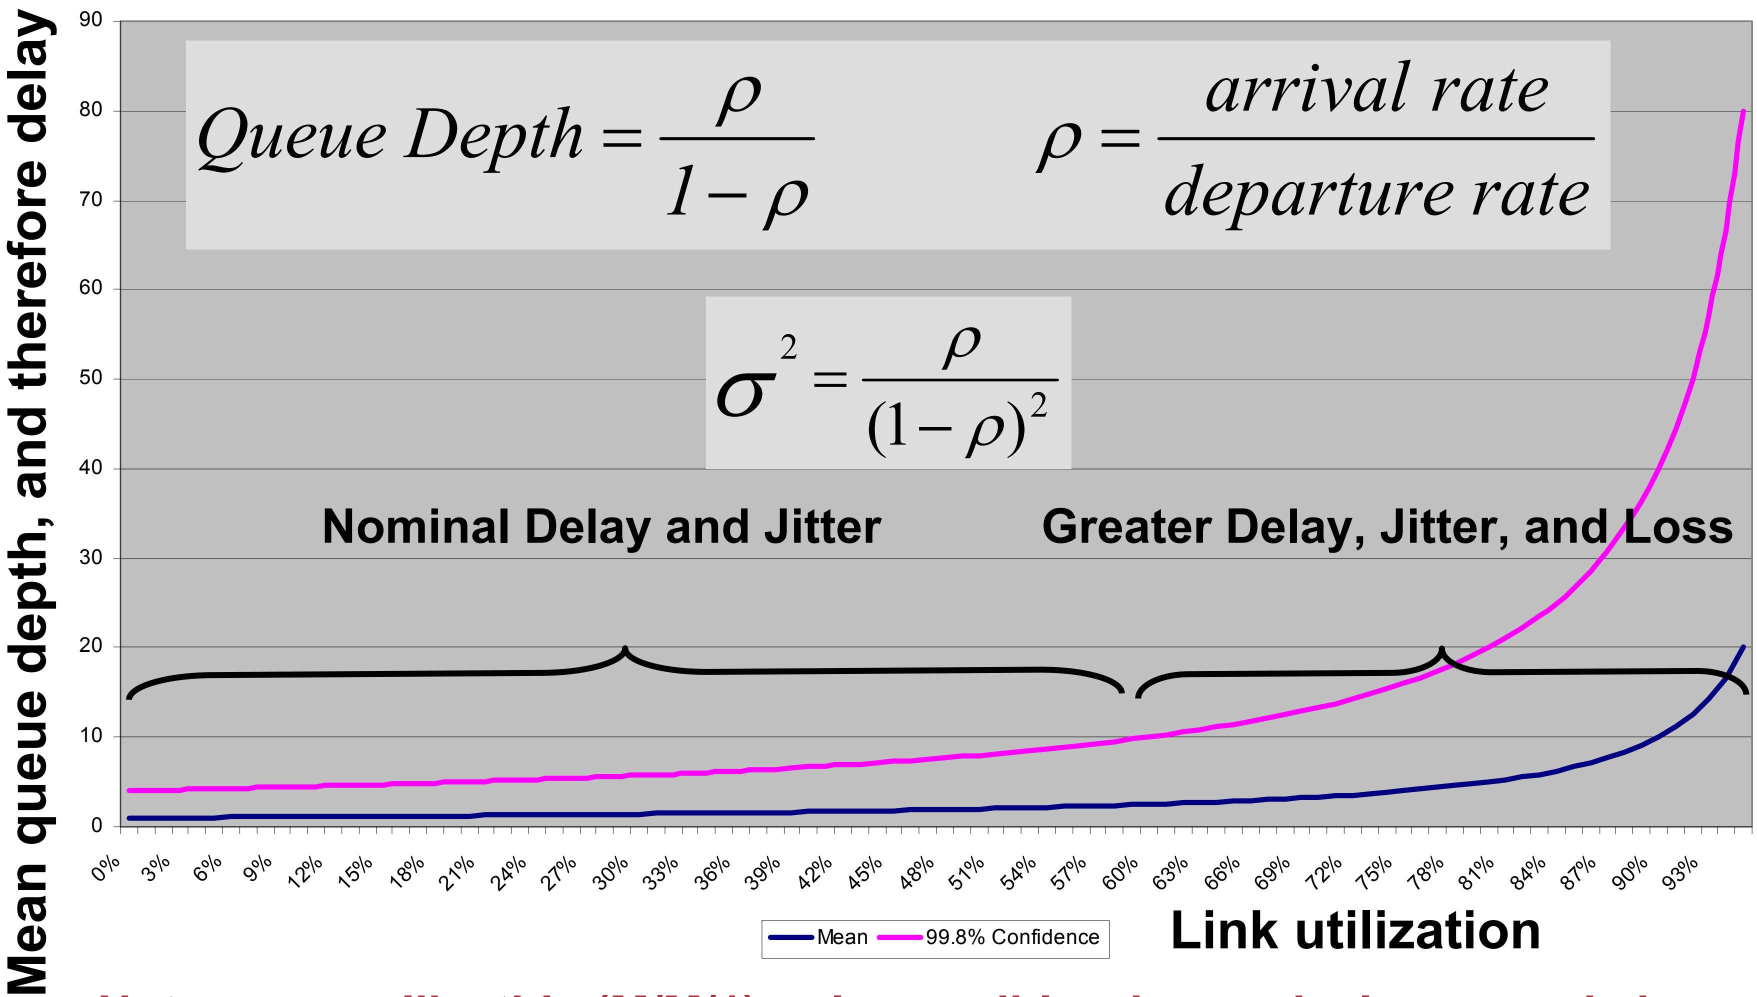

Note: curves like this (M/M/1) only possible when arrival process is known

Throwing bandwidth at the problem

Cisco.com

- As link bandwidth $\rightarrow \infty$, utilization $\rightarrow 0$, so delay, loss $\rightarrow 0$
- Problems:

  As bw $\rightarrow \infty$, cost also $\rightarrow \infty$

    When does the cost/quality tradeoff work?

  Pent-up demand frequently drives utilization up as bw becomes available

    e.g. TCP seeking to congest the bottleneck link

    e.g. video becomes more attractive as bw increases

  "Arms race" to see who becomes the bottleneck

The growth of bandwidth

- Fashionable to talk about "effectively infinite" bandwidth solving all problems
  - WDM, Photonic networks, OC-n SONET/SDH
  - Used by research, Internet backbones, and large corporations
- Corporate practical reality
  - T-1 remains the most common access link in the US
  - Globally, 64-128 KBPS remains a "large" access bandwidth, and more is costly
- Consumer practical reality
  - "Broadband Revolution" fitful and difficult to make money on
  - Perhaps half of Internet access still via dialup
  - Even broadband connections prone to congestion

Broadband access delays

Data courtesy of Fred Baker

Cisco.com

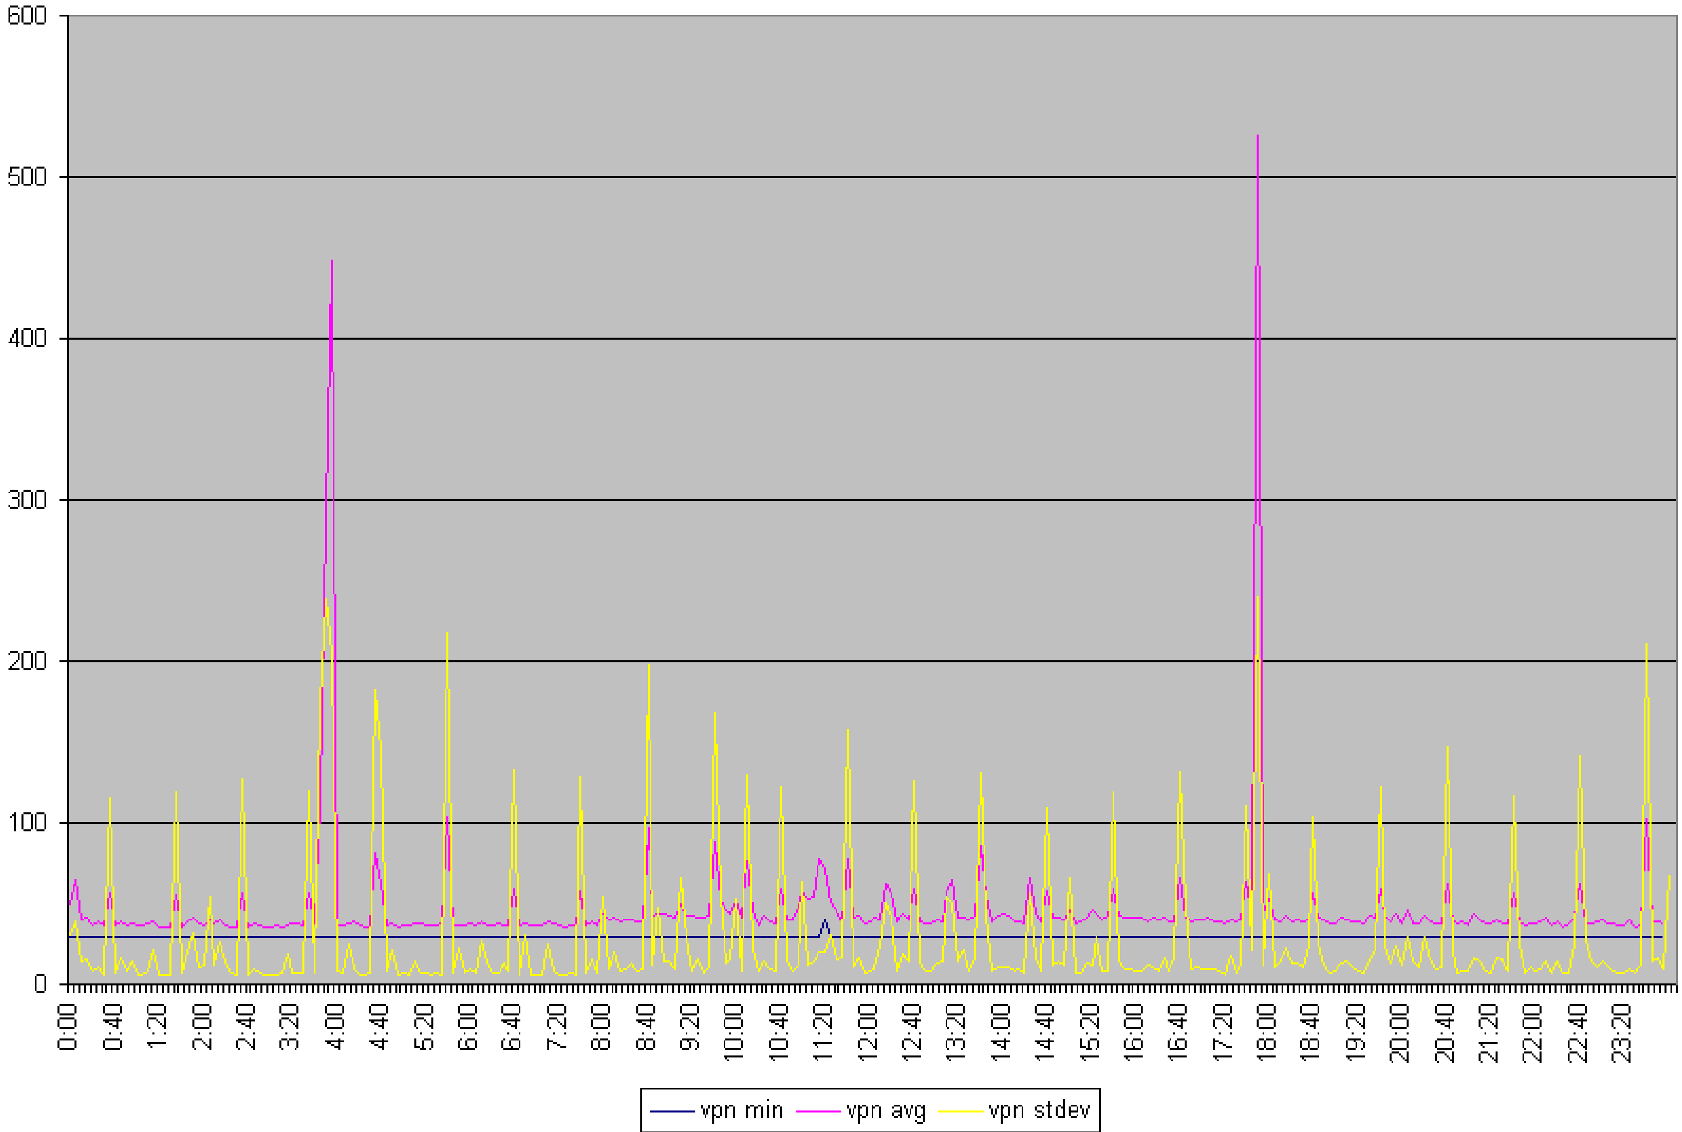

TCP behavior

- 95% of Internet traffic today runs over TCP
- The fact that the Internet hasn't collapsed can be traced to TCP congestion avoidance & "slow-start"

    Q. What happens if TCP becomes less dominant?

- TCP "avoids" congestion by causing it, then reacting

    Attempts to find BW of the "bottleneck link" and send at that rate

TCP behavior

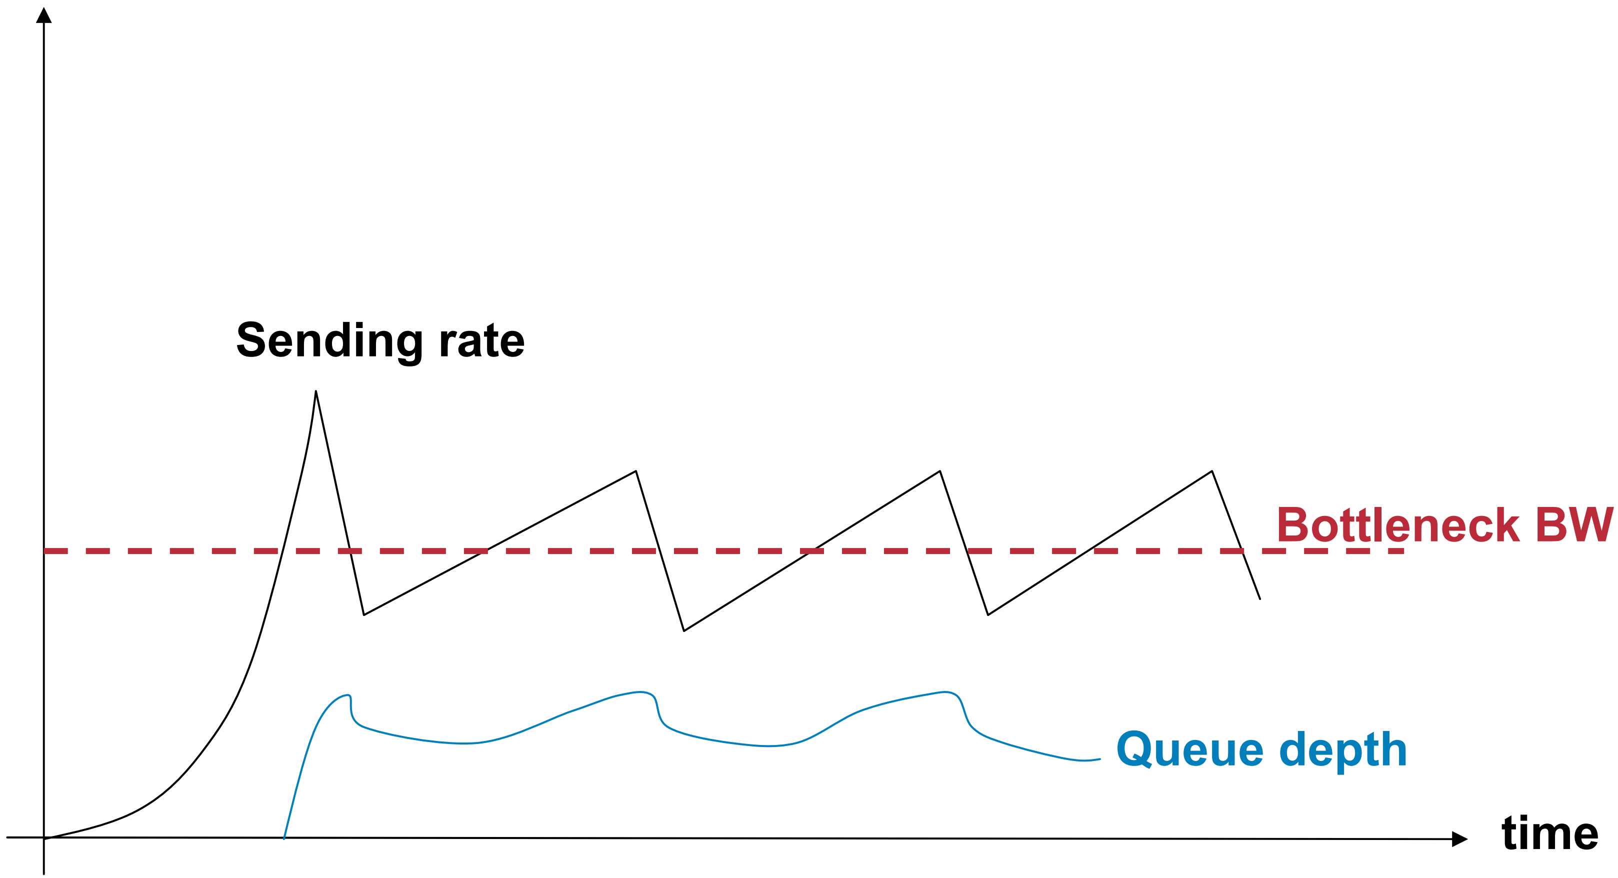

Timescales

- **A 10ms delay per hop may be too much for voice**
- **It takes < 10ms to create a 10ms queue**

  ⇒Traffic bursts at the millisecond level can matter

Average utilization measured over seconds or longer doesn’t capture this

Timescales

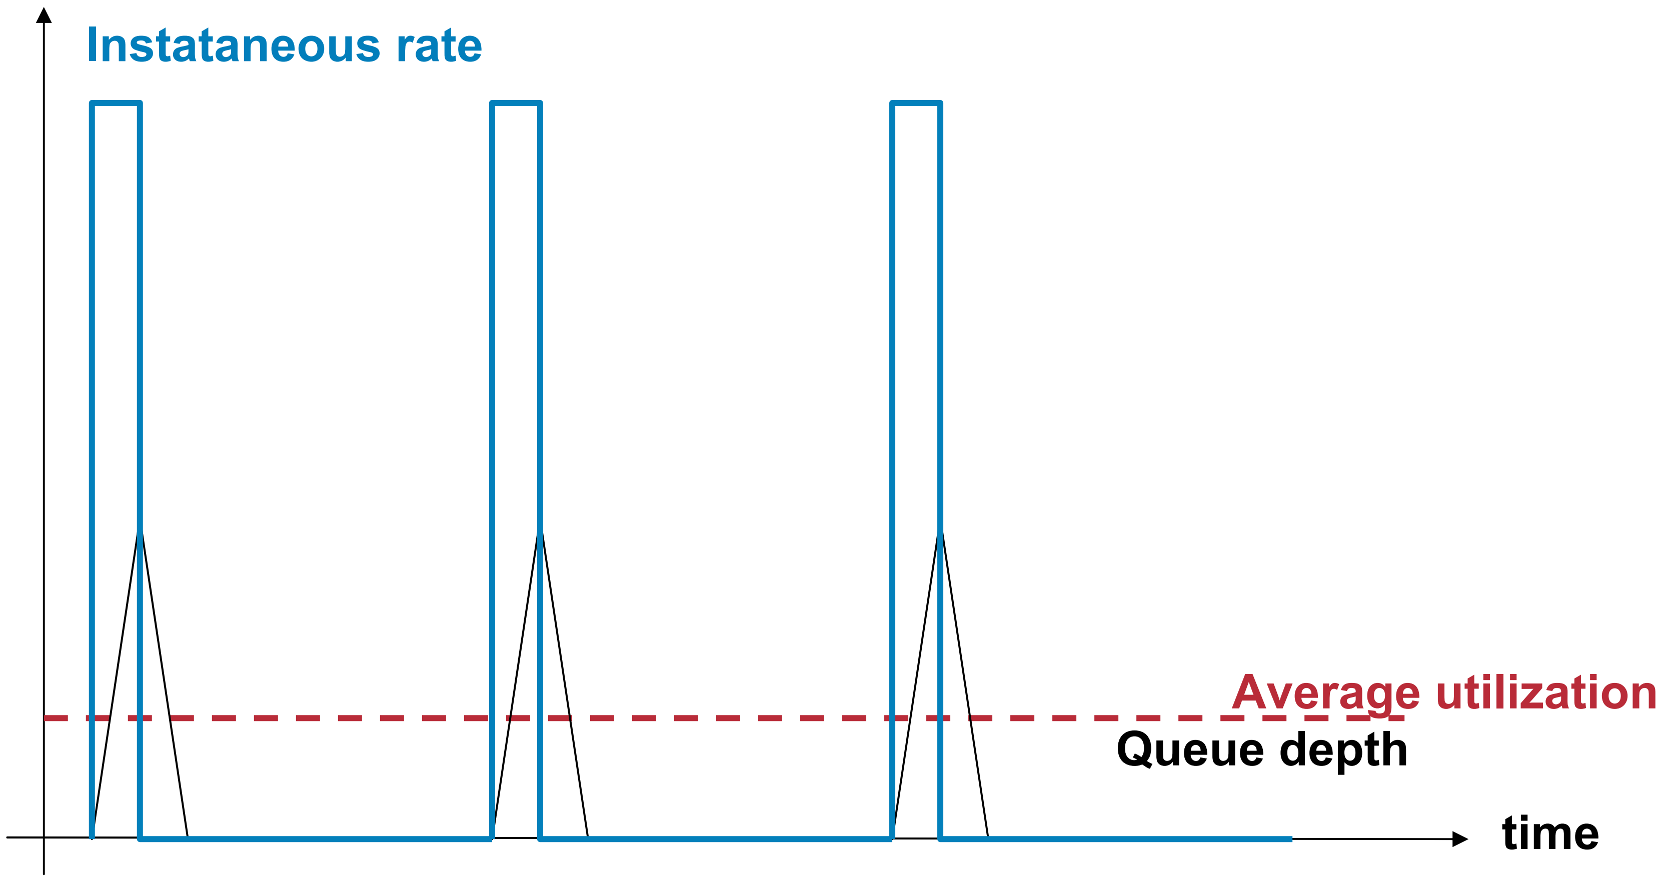

QoS mechanisms

- **If BW allocation is the only tool, then need to provide enough BW to meet the most stringent QoS needs of any app**
  - provide low delay to all traffic, not just voice/video
  - low loss to all apps (maybe overkill for voice)
- **Most QoS mechanisms aim to allocate the delay & loss to apps that are less sensitive to it**
  - Can’t create BW, just control who gets access to it and when
  - Sorting traffic into distinct queues creates different utilizations for each class

Customizing Delay/Load Trade-Off

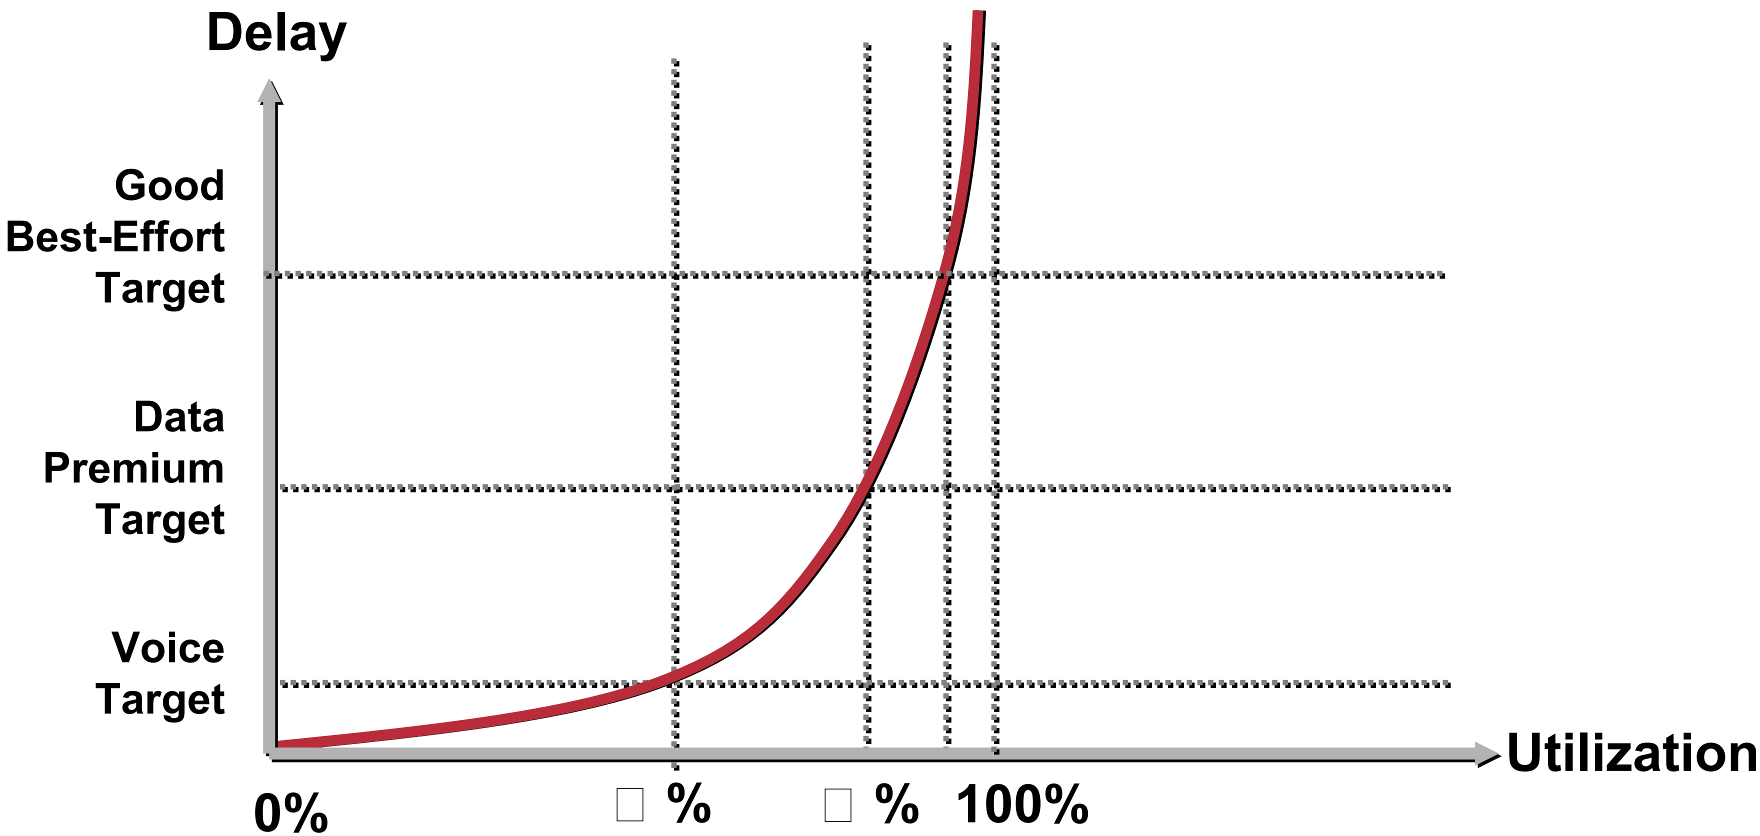

If I Can Keep Voice Traffic < □ %, I Will Keep Delay Under *M1* ms
If I Can Keep Premium Traffic < □ %, I Will Keep Delay Under *M2* ms

QoS mechanisms – other reasons

- **Service differentiation**

  Providers would like to extract money from customers willing to pay, e.g., corporate users with mission critical data, voice

  Requires mechanisms to make some traffic see "better" service

- **Protection among classes**

  Interaction between bursty & non-bursty sources

  Interaction between TCP-friendly & others

Cisco.com

QoS Mechanisms

QoS Mechanisms

- Classification: sorting packets into categories
- Marking
- Queueing/scheduling: determining which order packets get sent
- Policing/metering: counting packets & taking some action (e.g. remarking, dropping)
- Shaping: limiting the bandwidth to a class
- Dropping
- Admission control

Classification

- Sort packets into categories

  e.g. voice, video, premium, best effort

  Can use almost any criteria

    packet contents – IP header, application data

    point of entry (representing a customer)

  Per-flow or aggregate classification

- Necessary precursor to most other QoS functions

Queueing/scheduling

- **Determine which order packets get sent**
- **Determines how much bandwidth a class receives**
- **Basic approaches:**

  Priority – send higher priority traffic always

    Creates the illusion that other traffic is absent, i.e. emulates a less loaded network

  (Weighted) Fair Queue – like round robin, allocates a share of the link to a class

    Enables utilization to be controlled on per-class basis

  Work-conserving algorithms avoid bandwidth wastage

Policing/Metering

- Count the rate of packet arrivals in a class
  - Typically uses a "token bucket" with a rate and burst size
    - Tokens accumulate at some rate, up to max burst, and are "spent" as each packet arrives
    - If packet arrives at empty bucket, either drop it or downgrade its marking
  - Provides a means to limit the number of packets in a class
  - Limits can be on a per-customer basis (or finer)

Dropping

- Traditionally, packet drops occurred only when queues were exhausted (tail-drop)
- Active queue management drops more intelligently
  - Send the message to a TCP app with just one packet drop, not several
  - Drop preferentially from less sensitive/less important applications
  - Drop preferentially from traffic that was previously policed as "out of contract"

Random Early Detection

- **Average queue length is monitored to detect onset of congestion**

  Averaging time constant is O(RTT)

- **As small percentage of packets is dropped, TCP backs off, congestion is averted**

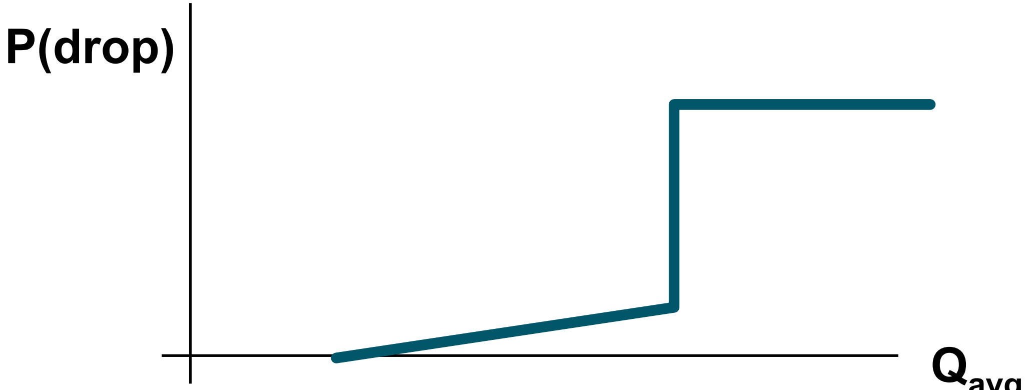

Weighted RED

Cisco.com

- Just like RED, but with different drop profiles for different classes
- Can carry extra traffic opportunistically & shed load as needed

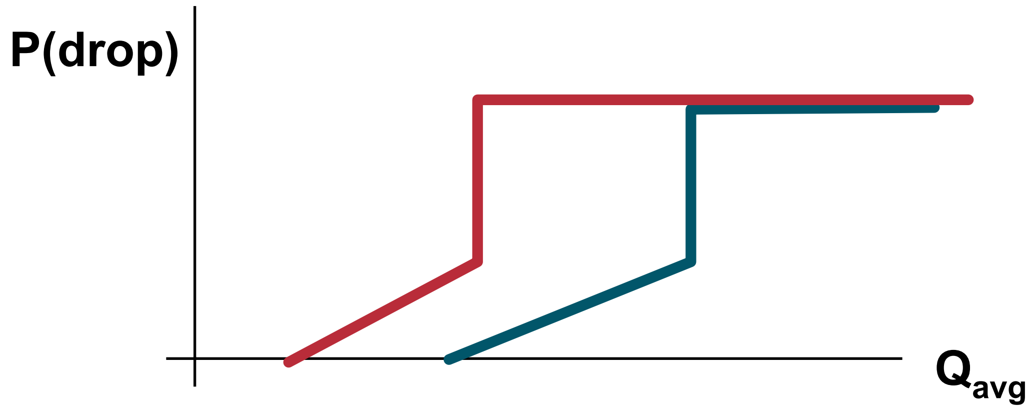

Explicit Congestion Notification

- A simple enhancement to RED
- Packets are marked rather than dropped
- TCP congestion avoidance responds as if drop occurred – other transports may also react appropriately
- Congestion avoidance without loss
- Host participation required
  - Packets marked “ECN capable”
  - Receiver conveys marking back to sender at transport layer

Admission Control

- **All mechanisms so far are in “forwarding plane” – admission control is the “control plane”**
- **Determining ahead of time if the resources are available**

  Avoiding the “everyone is high priority” syndrome

  Most valuable if the application would prefer no service to impaired service, e.g. “fast-busy” for voice

  Better to have 99 good quality calls and 1 fast busy than 100 poor calls

  Applications with steep “utility functions” the best candidates

Focusing on IP QoS

- IP the dominant networking protocol today

    Global reach

    Heterogeneity of link layers and applications

- IP QoS is not an oxymoron (as some had thought)
- Mechanisms are in fact generic

Differentiated Services Overview

- "Diff-serv" is a set of Internet standards
- Clearly the preferred QoS technology for ISPs & enterprises today
- Near-minimal complexity

  e.g. can deploy DS with just 1 header bit and 2 "per-hop behaviors" (PHBs)

- Edge behavior (classification, marking, policing etc.) + core behavior (PHBs) provides services

  variety of services from a single PHB

Diff-serv Contributions

- Standardized definition of the “Diff-serv Code Point” (DSCP) in IP header
  - (after years of confusion about TOS)
- Small set of standard PHBs
  - Expedited Forwarding (like a priority queue)
  - Assured forwarding (WFQ & WRED)
- An overall architecture for DS
  - Mostly formalizing ideas already in use

Putting it all together

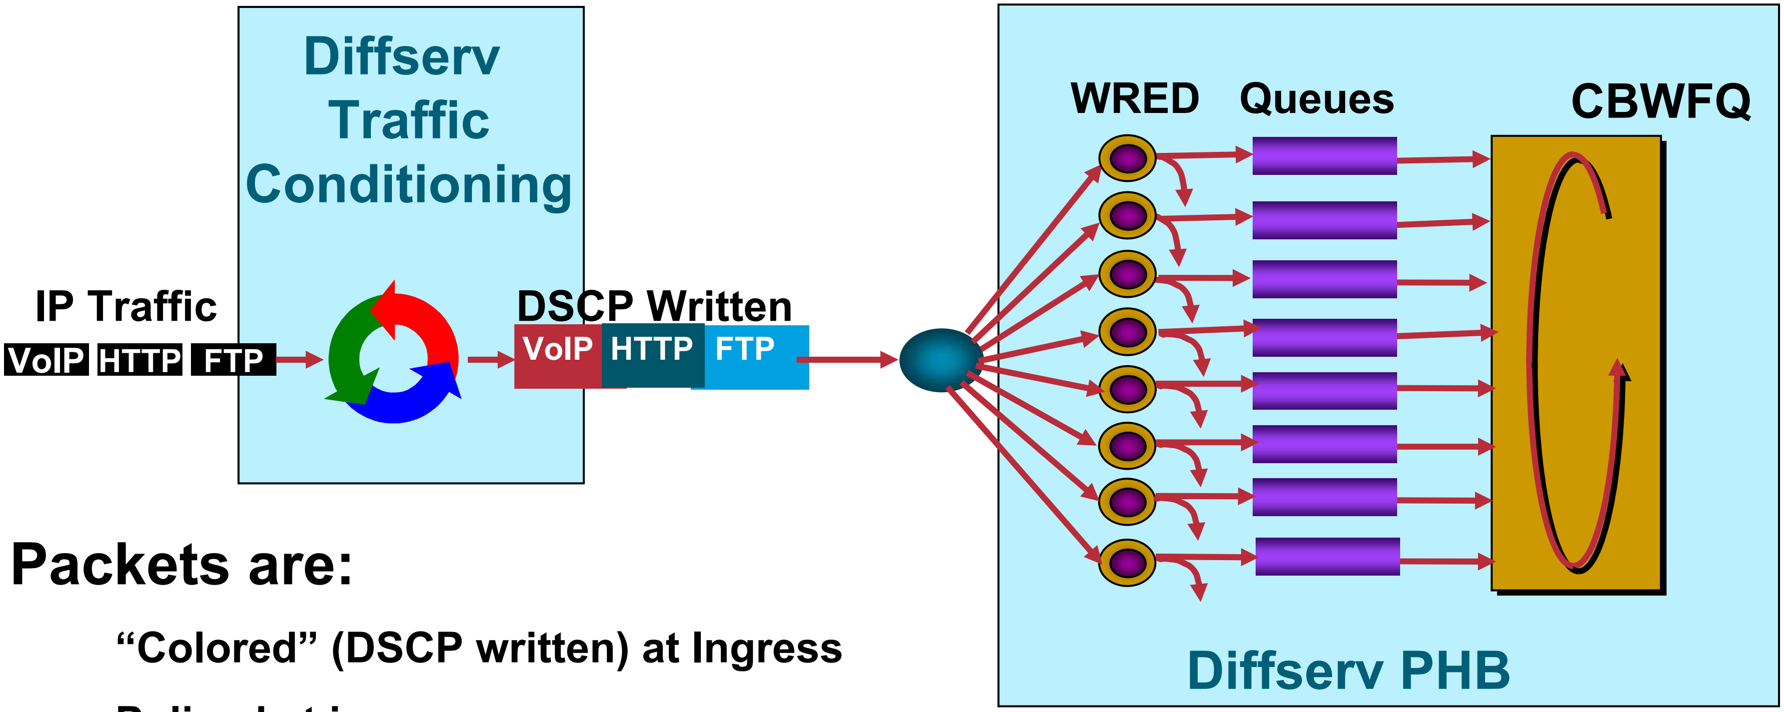

Packets are:

- "Colored" (DSCP written) at Ingress
- Policed at ingress
- Classified before queuing
- Potentially discarded by WRED
- Placed in queues based on DSCP
- Scheduled by CBWFQ (and/or LLQ)

Example Service

- Classifier + token bucket policer at network edge

  recognize & meter traffic in need of isolation; set DSCP = x

- Dedicate a queue (and some bandwidth) to DSCP = x

- Effect is to run this traffic on its own logical network (with controlled utilization)

Example Service (2)

- Customers buy an SLA involving an allowed amount of “premium” data service

    Specified as a token bucket

- Policer at provider ingress marks traffic below the rate as “premium”, excess is marked as “best effort”

- These markings are used by WRED to ensure that excess traffic is sent if possible, but dropped under congestion

- Note that WRED is implemented on a single queue, so no intra-flow misordering

    Good news for TCP

IPv6 and QoS

Cisco.com

- It was widely perceived that IPv6 would enhance the QoS of IP
- In fact, QoS mechanisms for IPv6 are identical to those of IPv4

  Both use 6 bits in header for Diffserv, 2 bits for ECN

  Both supported by RSVP

Evaluating Mechanisms vs. Bandwidth

Evaluation – It's all about cost

- **Bandwidth comes at some price**

    Very much dependent on environment (LAN/WAN, last mile/backbone, own/lease, regulatory factors)

    Or not…for many consumers, higher speed access isn't there at any price

- **QoS mechanisms incur operational costs**

    Train personnel to configure correctly

    Negotiate and bill for (more complex) SLAs with customers

    Monitor network for incorrect or outdated configurations, risk of SLA violations

    Possible performance impact on equipment

Other Cost/Benefit Factors

Cisco.com

- **If QoS mechanisms make more applications work on common infrastructure (e.g. VOIP) then there are cost savings**

  The driving reason behind enterprise VOIP is avoiding toll costs, leveraging existing capacity

Benefits of traffic isolation

- Separating "premium" customers from others makes stronger (and more lucrative) SLAs possible
    - e.g. may continue to deliver SLA even during failure/overload scenarios
- Separating traffic types can be good for all
    - TCP performs worse when sharing queues with "non-responsive" flows
    - Voice performs worse when sharing queues with more bursty traffic types

Asian perspective

- **Application needs**
  - Probably not too different from anywhere else
  - VOIP more attractive wherever TDM voice is costly relative to IP bandwidth
- **Link costs**
  - Glut of capacity less apparent than in the US
  - Transoceanic links inherently costly (and time-consuming to deploy/upgrade)
  - Regulation frequently drives costs up
  - Providers who don’t own their links much more sensitive to cost

Deployment Experience

QoS Deployment

- Backbone ISPs
- VPN Service Providers
- Enterprise

Backbone ISPs

- Most today have built in sufficient capacity to make QoS mechanisms superfluous
- Hard to start selling premium services when
  - Your basic service is already good enough for voice
  - So is your competitors'
- Some are carrying VOIP and using basic priority queueing mechanisms
- Interest in QoS mechanisms higher when failures are considered
- Interest may rise as capacity glut subsides

VPN service providers

Cisco.com

- Over 100 SPs today offer “Layer 3” VPN services (MPLS/BGP VPNs)
- Mostly selling as a frame-relay replacement
- High customer demand for “CIR”-like services
- Many SPs using PQ for VOIP and WRED for premium data services

Enterprise

- Bandwidth in the campus is cheap enough to overprovision
- Inter-site links are the bottlenecks
- Plenty of interest in PQ mechanisms for VOIP
    - Admission control often needed also
- Mission critical apps with response-time requirements also common
- Video getting more common

Conclusions

- In theory, bandwidth could solve all QoS problems
- In practice, plenty of need for QoS mechanisms
  - BW too expensive/not available
  - Protection from “greedy” applications
  - Protection against failures
  - Premium services to enhance revenue
- There are many useful IP QoS technologies deployed today