

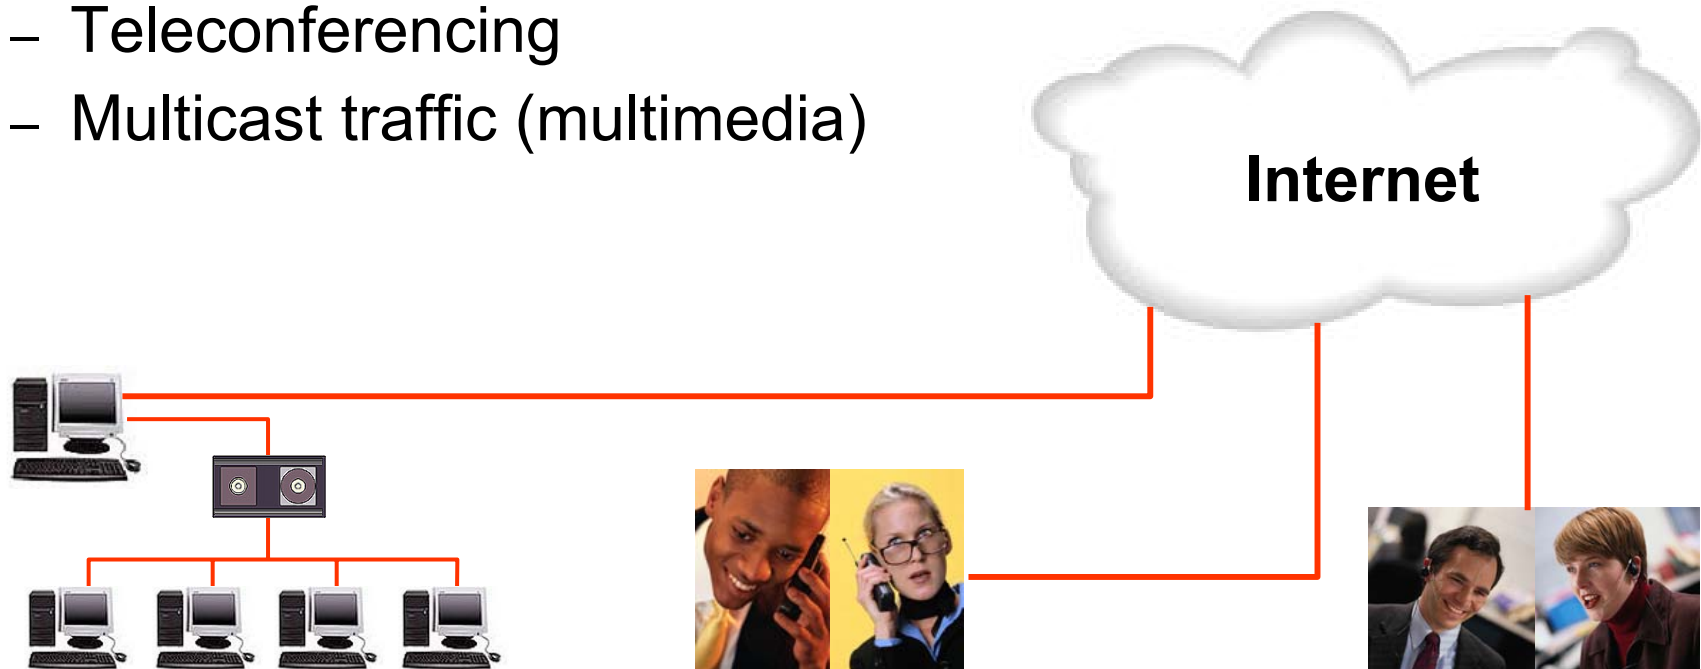
HP ProCurve Quality of Service: A Conceptual Approach

Rodney Turner
Solutions Architect
Hewlett-Packard Company



Why are we concerned about QoS?


- Time-sensitive applications are driving the move towards prioritization
 - Voice over IP
 - Teleconferencing
 - Multicast traffic (multimedia)



What is quality of service?

- QoS is the measurable degree to which an IT staff can successfully guarantee the **predictable behavior** of network services to meet customer and application requirements.
- Distinctive characteristics of the data delivery
 - Can the data be delivered?
 - How quickly is it delivered?
 - How regularly is it delivered?
 - How reliably is it delivered?

End-to-end QoS



In order to be effective,
QoS must be implemented end-to-end.

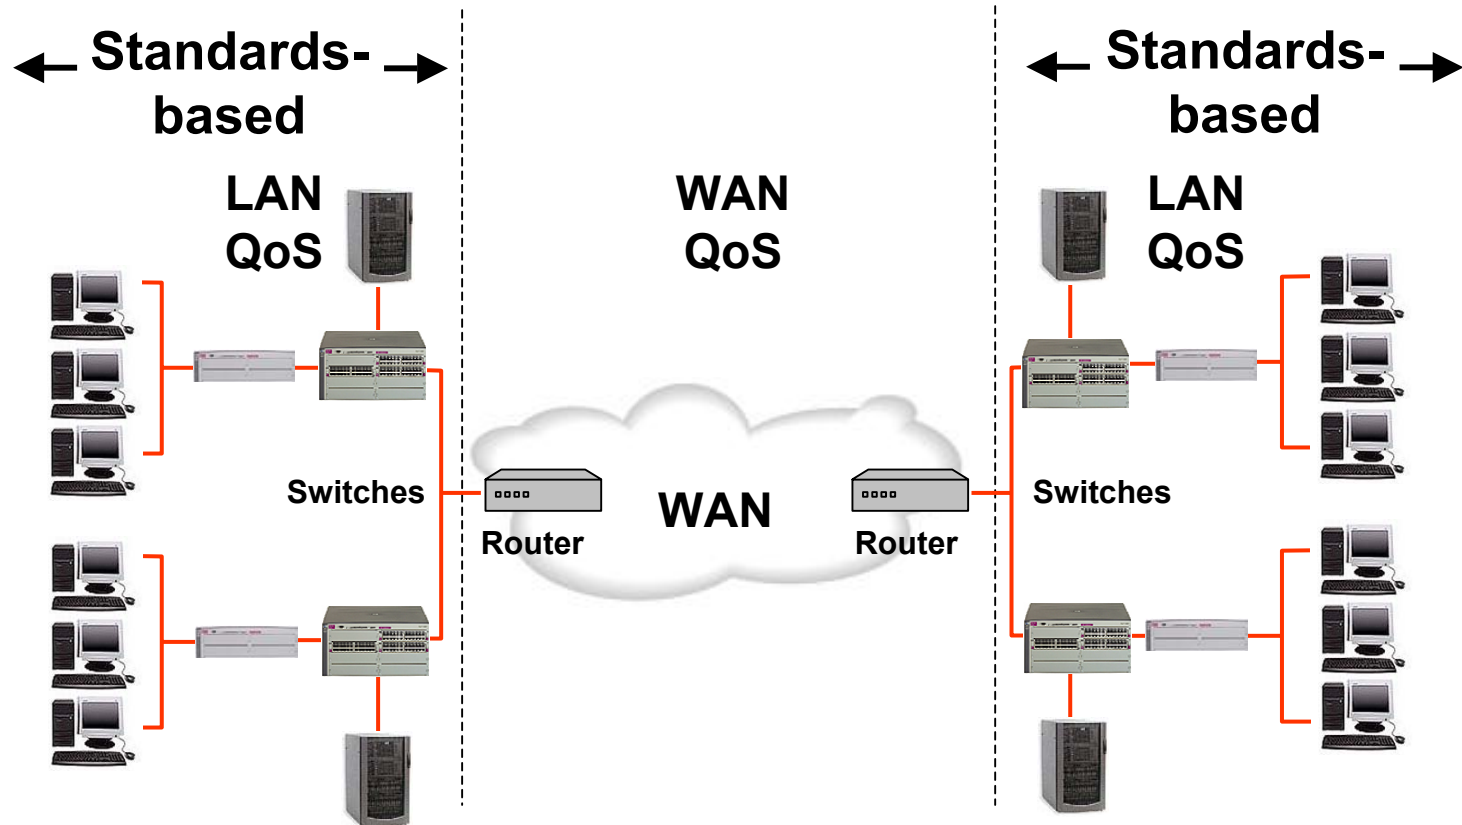
QoS on the WAN: Complicated

- Bandwidth is constrained, compared to LAN
 - If congestion occurs, it will likely be in the WAN
- A myriad of possible WAN technologies may be used
 - Point-to-point
 - ISDN
 - Frame Relay
 - ATM
 - SONET
 - And more
- Cisco focuses on the WAN

QoS on the LAN: Simple

- One network technology: Ethernet
- Not bandwidth-constrained
- Standards-based class of service
 - IEEE 802.1p
 - Differentiated Services (RFC 2474)

LAN standards allow end-to-end QoS



Standards allow LAN QoS to interface with WAN QoS.
Vendor interoperability is assured.
HP specializes in LAN QoS.

Enemies of a predictable network

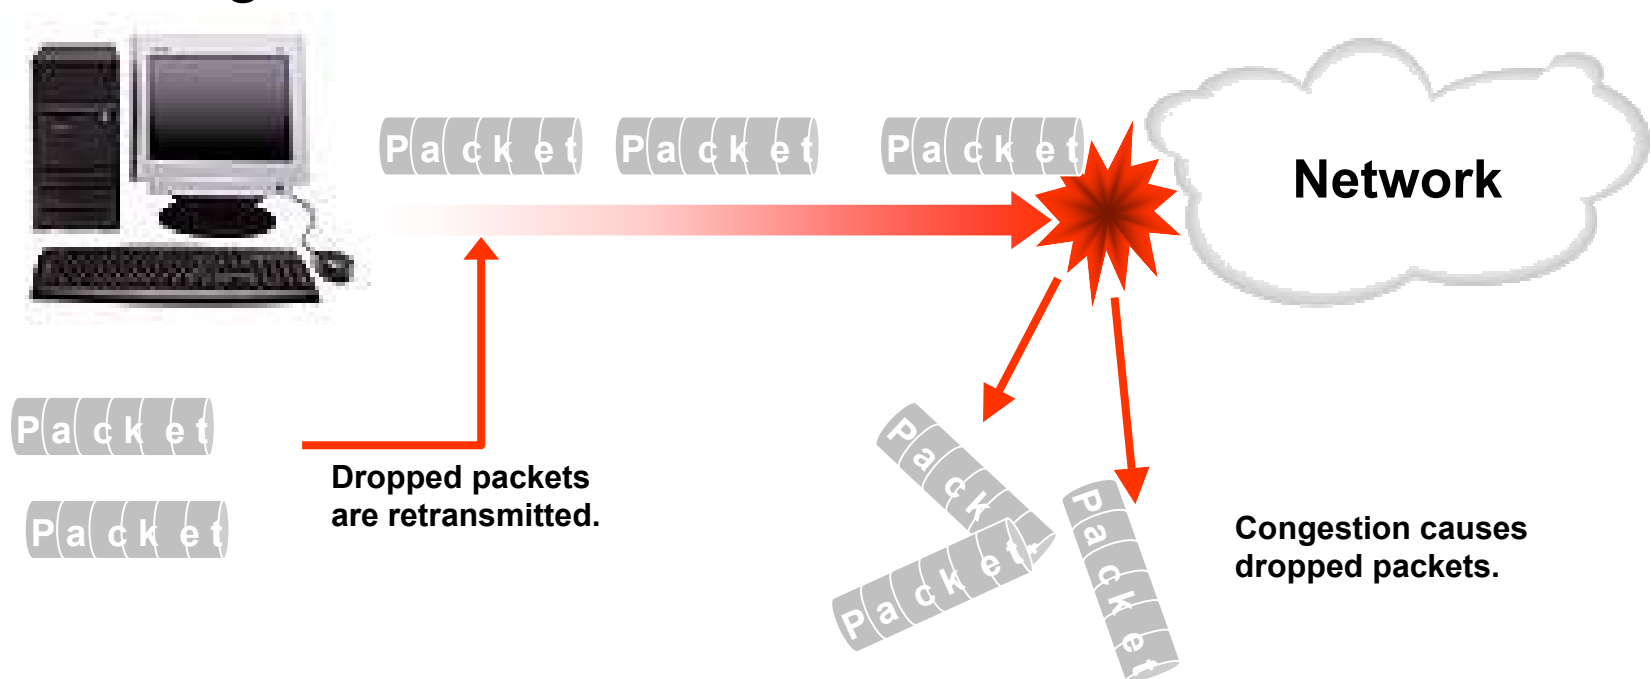
- Network failures
 - Dealt with using other technologies (redundancy)
- Network congestion
 - QoS is used to avoid or cope with network congestion



The sinister secret of network congestion

1. Too much traffic causes congestion.
2. Congestion causes dropped packets.
3. Those dropped packets are generally retransmitted.

Effect: Congestion causes increased traffic.

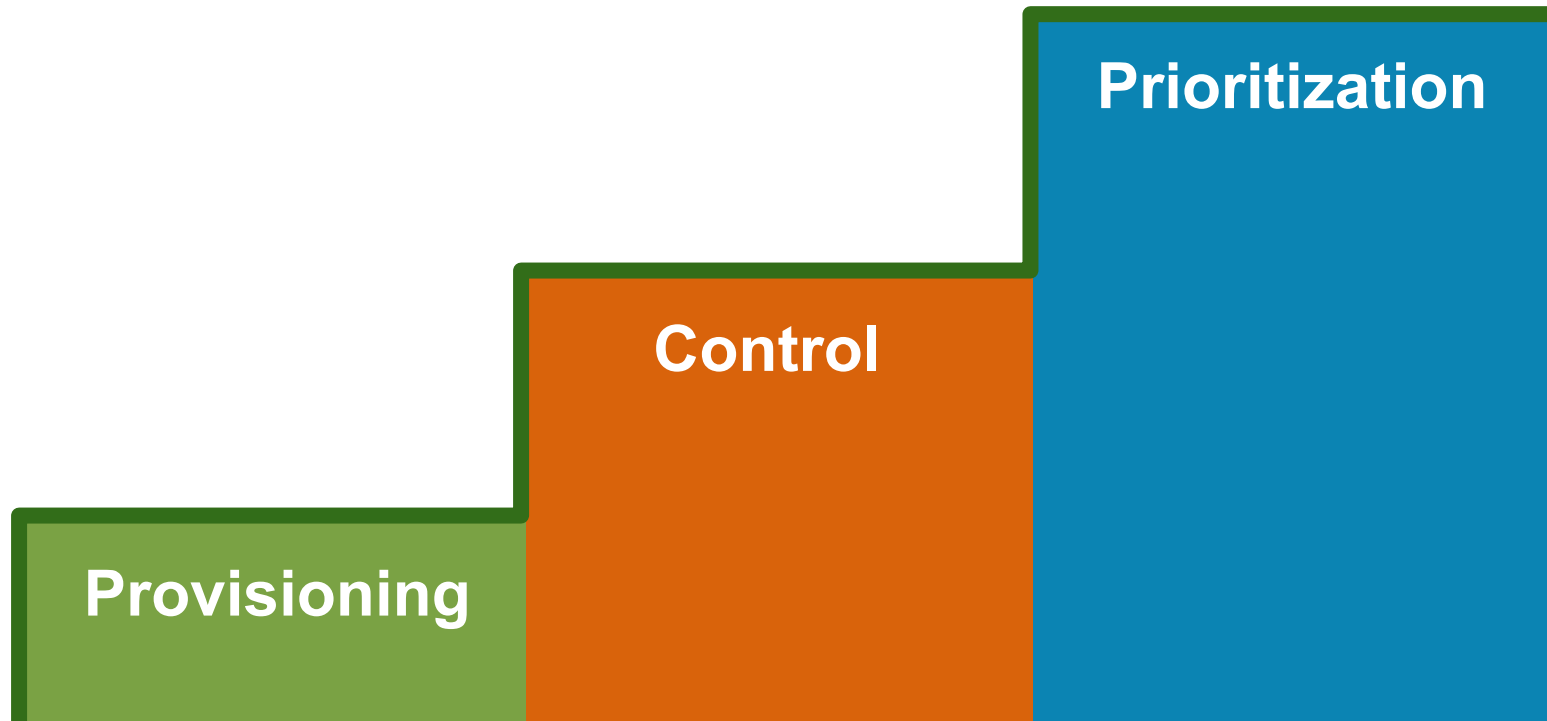


Sources of congestion

- More traffic than a link can transmit (oversubscription)
- More traffic than the switch can forward (poor performance)




HP LAN QoS philosophy




Step 1: Provisioning

- Provision the network so that congestion is eliminated in all but the most extreme cases
- HP ProCurve switches provide more “lanes” than any of our competitors




Step 2: Control

- Use current networking techniques to eliminate unwanted traffic
 - IGMP (multicast)
 - Traffic and protocol filtering
 - VLANs



Step 3: Prioritization

- Prioritize critical traffic in order to provide predictable throughput even under extreme conditions

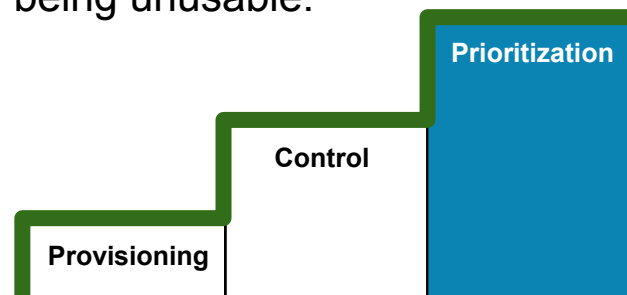


Cisco's QoS approach focuses on prioritization

“QoS refers to the ability of a network to **provide better service to selected network traffic** over various underlying technologies including Frame Relay, Asynchronous Transfer Mode (ATM), Ethernet and 802.1 networks, SONET, and IP-routed networks.”

–*Cisco Web site
January 2003*


“The key enabling technology for network convergence in the Architecture for Voice, Video, and Integrated Data (AVVID) is QoS. This is because voice, video, and mission-critical data have stringent service requirements from the network infrastructure. These requirements supersede the requirements of generic data traffic. If voice, interactive-video, and mission-critical data are not given priority service from network devices, then the quality of these important applications would quickly degrade to the point of being unusable.”



–*Cisco Enterprise Quality of Service Design
White Paper
August 2002*

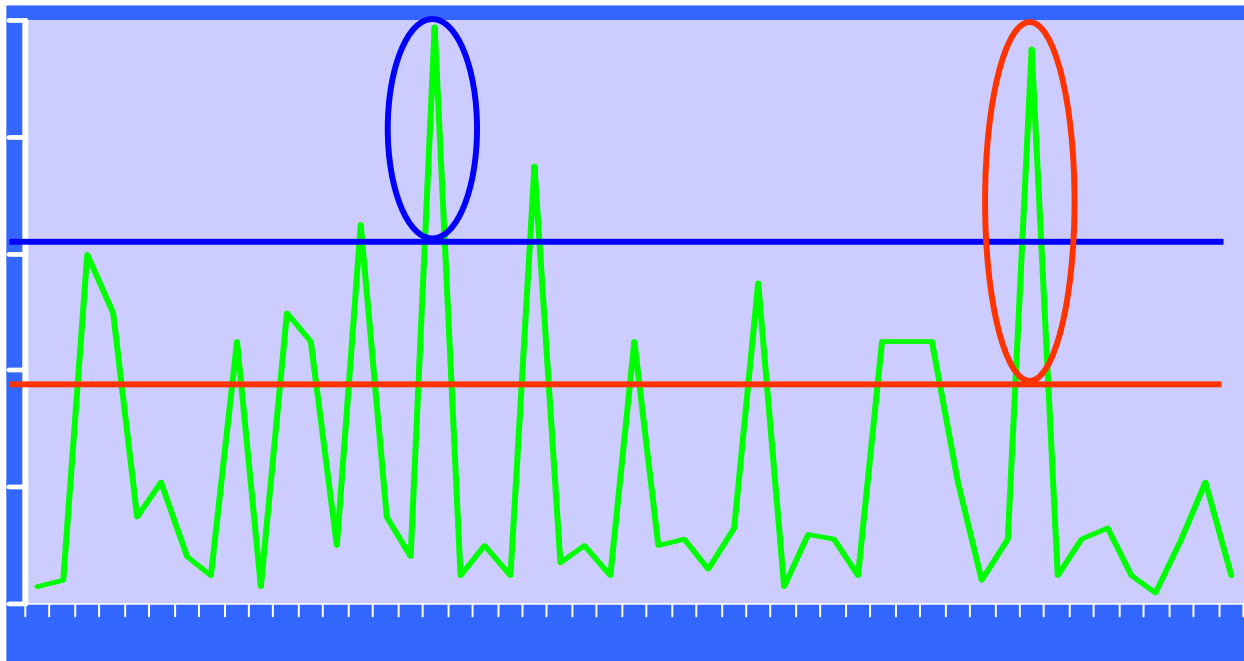
Fundamental axiom of prioritization

- If congestion does not exist, prioritization has **absolutely no** effect.




Avoiding congestion with provisioning

Network provisioning can minimize traffic that must be controlled using QoS



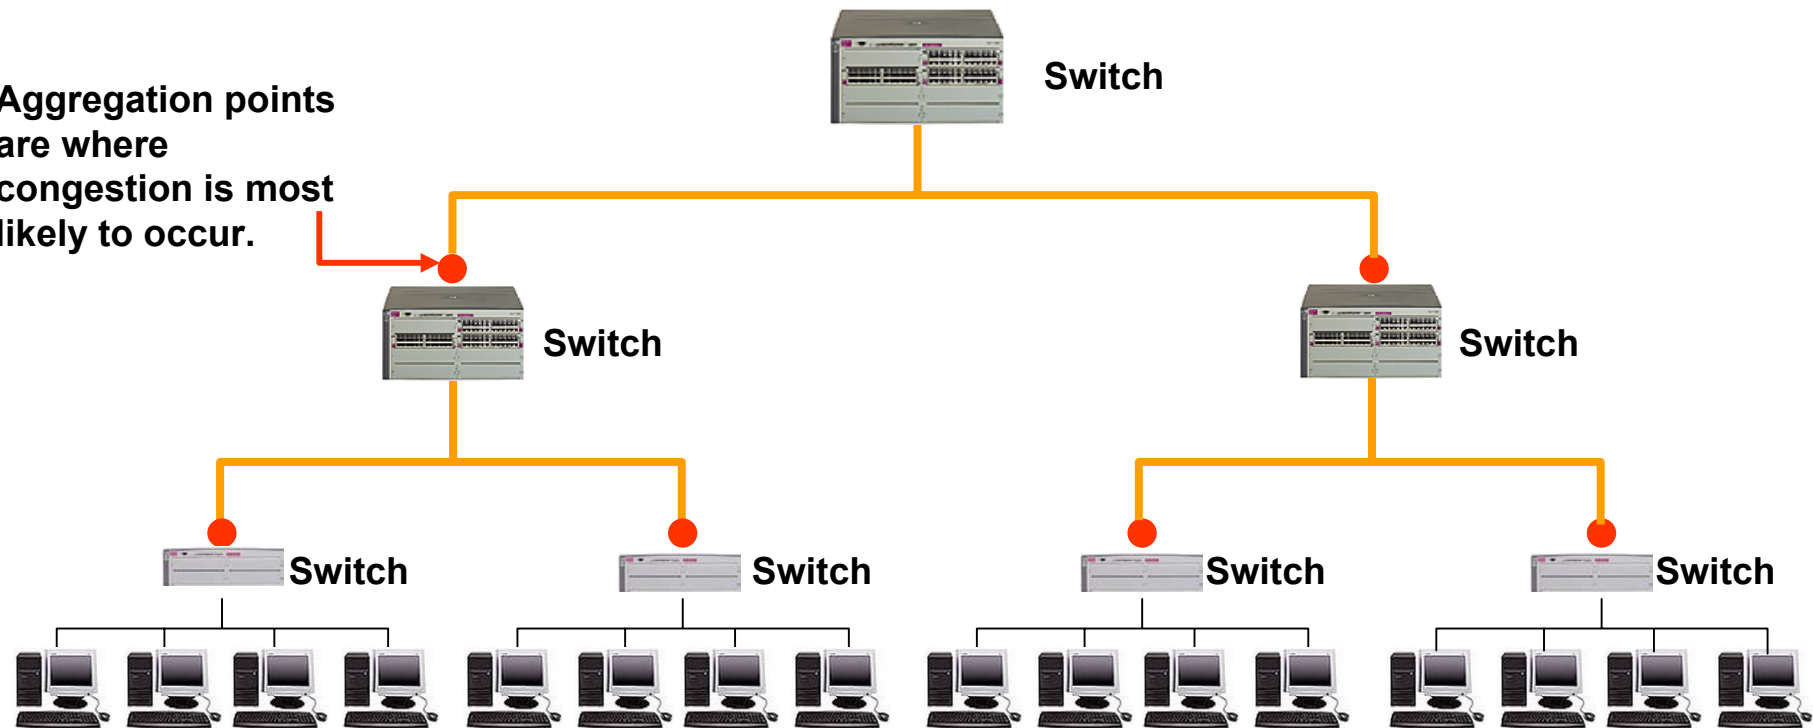
- Network provisioning with HP ProCurve switches
- Network provisioning with a competitor product



Provisioning and oversubscription


- Wire speed is not required to every end node on the network
- Oversubscription is used in some parts of the network

Aggregation points are where congestion is most likely to occur.



Getting traffic information for management

How do you get this kind of information about your network traffic?



Extended RMON in HP switches provides detailed knowledge of your network traffic.

Dealing with congestion

- Long-term or short-term problem?
- If a highway is consistently congested, how do we solve this issue?
 - More lanes or alternate routes (increased capacity)
 - Diamond lanes (prioritization)



VoIP traffic volumes


- Number of industry-standard G.711 voice codec (uncompressed) conversations
 - 10 Mbps: 106
 - 100 Mbps: 1,068
 - 1 Gbps: 10,684
- Network utilization for 200 VoIP conversations
 - 100 Mbps: 18.7%
 - 1 Gbps: 1.9%

Prioritization terms

Classification

- The ability of a network device to distinguish between traffic flows based on some characteristic of the flow and then forward traffic using the appropriate priority


- Application
- Device
- Protocol
- Port
- Other



Prioritization terms (continued)

Marking

- The ability of a network device to mark traffic with its priority after the traffic has been classified



Standard prioritization marks

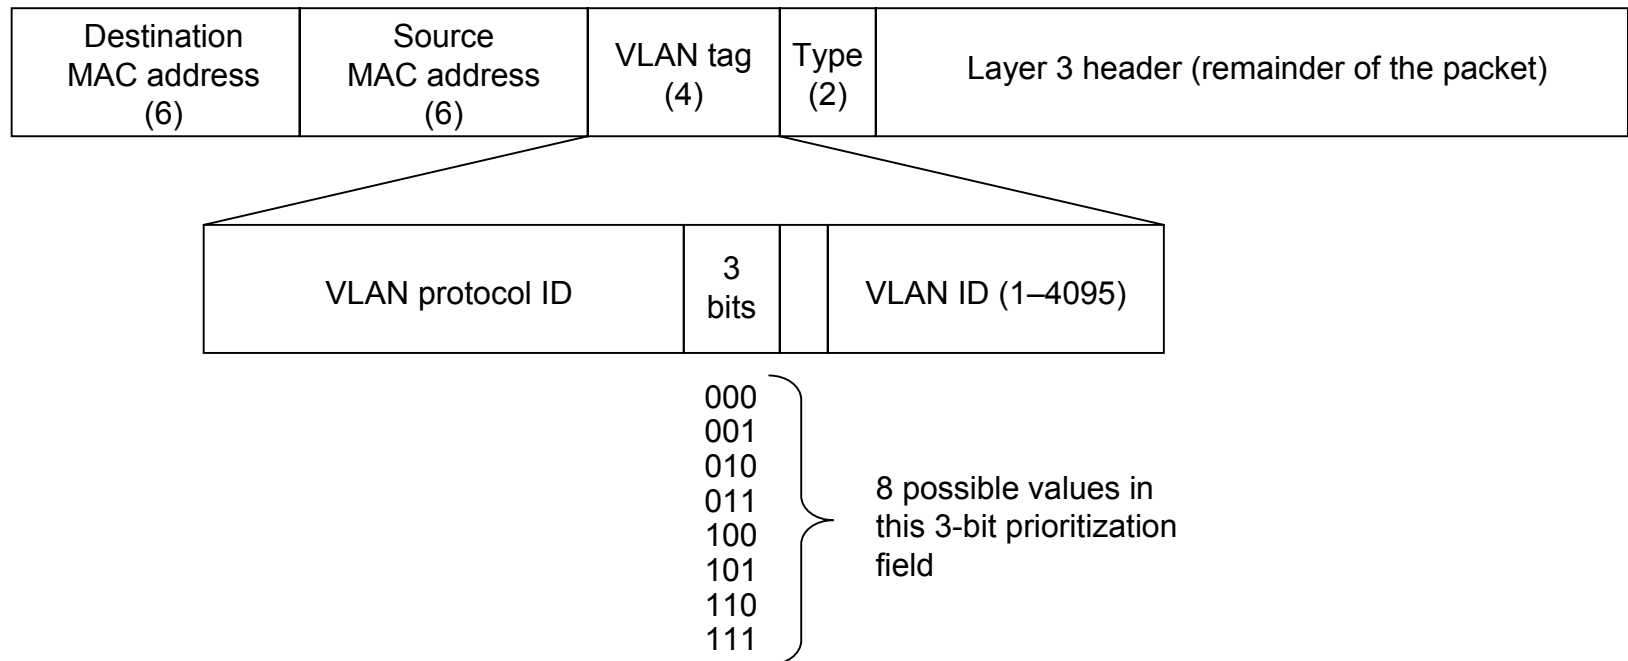
- Layer 2: IEEE 802.1p
- Layer 3: IP Type of Service (IP ToS)
 - IP Precedence
 - Differentiated Services
- Same priority markers that Cisco uses

Traffic classes

- Traffic classes establish relative priority:
 - 7 (high)
 - 6
 - 5
 - 4
 - 3
 - 0 (normal) *
 - 2
 - 1 (low)

Traffic prioritization using 802.1p


- A 3-bit field in the 802.1Q tag is reserved for prioritization. There are 8 possible values (0–7).
- Q-compliant end stations can set priorities for their traffic.
- Q-compliant switches honor priority settings in forwarded frames.



802.1p

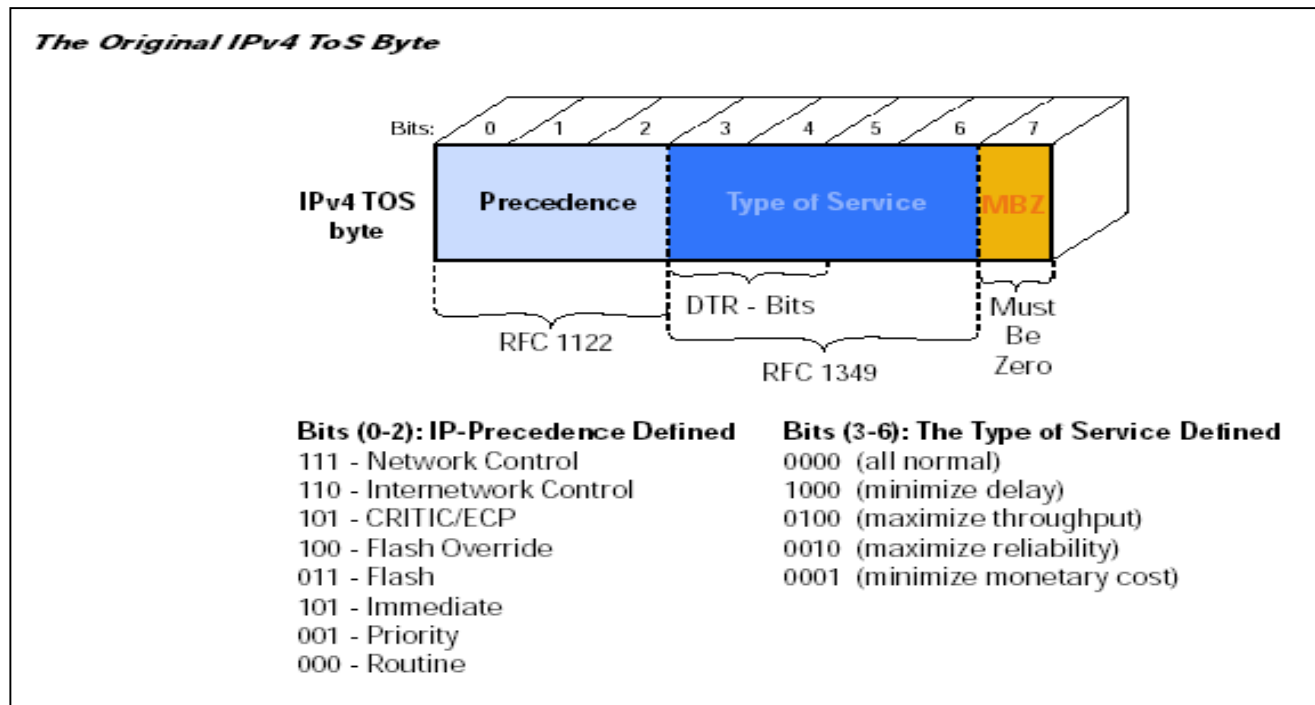
Priority for different types of traffic

- IEEE 802.1p standard
- Priority field included in 802.1Q frame tag
- Priority traffic causes delay in other traffic



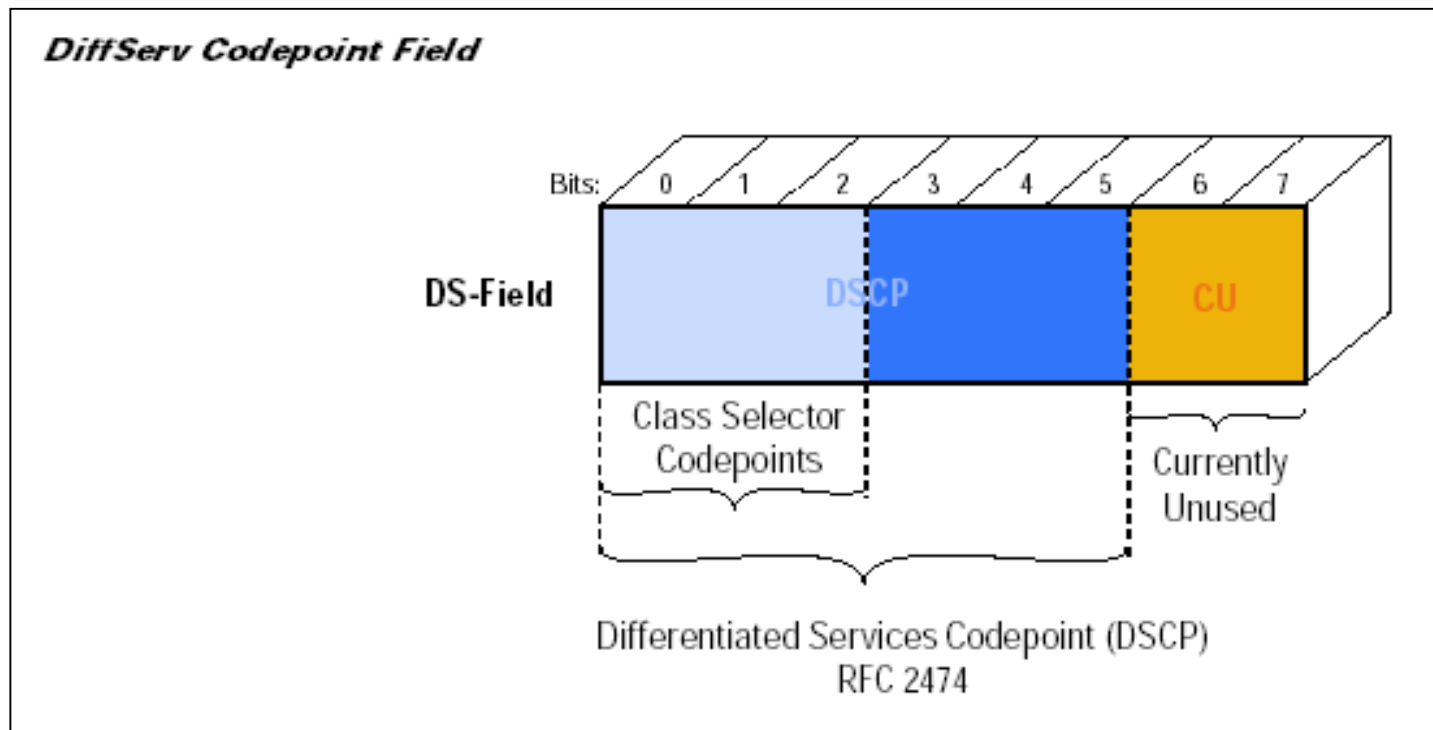
Type of Service

- The Type of Service (ToS) field (8 bits) was defined in the IP header to provide information to an infrastructure device on how a packet should be forwarded.
 - Original definition (IP Precedence):



Type of Service (continued)

- The Type of Service (ToS) field (8 bits) was defined in the IP header to provide information to an infrastructure device on how a packet should be forwarded.
 - New definition (Differentiated Services):



HP ProCurve Class of Service features



- HP ProCurve switches* have the capability to differentiate traffic based on the following characteristics:
 - IP address
 - Source address or destination address
 - IP ToS: IP Precedence or Differentiated Services
 - (old interpretation of IP ToS field)
 - (new interpretation of IP ToS field)
 - Protocol
 - IP, ARP, IPX, DEC LAT, AppleTalk, SNA, NetBEUI
 - IEEE 802.1Q VLAN ID

* HP ProCurve Switches 8000M, 4000M, 2424M, 2400M, and 1600M with firmware revision C.07.XX

Who should prioritize traffic?

The ideal world


- Ideally, client applications should classify and mark their traffic with the appropriate priority (the application knows best).

The real world

- If client applications cannot mark their own traffic, then the network infrastructure must take this responsibility.
 - Case 1: Infrastructure devices can only classify traffic (per-hop behavior: undesirable).
 - Case 2: Infrastructure device can classify and mark traffic (edge classification: desirable).
- HP ProCurve switches classify and mark traffic.


Case 1: Per-hop behavior

- Each network device on the path between the end nodes
 - Classifies all traffic
 - Prioritizes appropriate traffic through this device
 - Sends traffic out unmarked
- Next device starts all over again
- Process continues up to and including busy core devices



Case 2: Edge classification

- Network edge switch
 - Classifies and marks all traffic with traffic's priority
 - Prioritizes traffic through this switch
- Next switch in path
 - Traffic is carrying its own priority marker, so classification overhead is eliminated.
 - All later devices prioritize traffic based on the marker.



HP ProCurve 5300 QoS classification capabilities



- The HP ProCurve 5300xl can classify traffic based on the following characteristics:
 - Source and/or destination IP address (device)
 - Source or destination TCP/UDP port (application)
 - Protocol (IP, ARP, IPX, DEC LAT, AppleTalk, SNA, NetBEUI)
 - VLAN ID
 - Multicast traffic
 - Physical port
 - Existing 802.1p can be used or re-marked (overridden)
 - Existing IP ToS can be used or re-marked (overridden)

Other HP ProCurve QoS classification capabilities

- The HP ProCurve 2500, 2600, 4100, and 6108 can classify traffic based on the following characteristics:
 - Multicast traffic
 - Physical port
 - Existing 802.1p marks

HP ProCurve QoS marking capabilities



- The HP ProCurve 5300xl can mark traffic using
 - The IP ToS field (IP Precedence or DiffServ)
 - The IEEE 802.1p field
- The HP ProCurve 4100gl, 6108, 2650, and 2500 can mark traffic using
 - The IEEE 802.1p field

Order of precedence

- When traffic to be classified matches with multiple criteria, the switch enforces the following precedence:
 - UDP/TCP application type (port number)
 - Device priority (destination or source IP address)
 - IP Type of Service field
 - Protocol (Ethernet Type field)
 - VLAN ID
 - Incoming source-port on the switch
 - Incoming 802.1p value

Two divergent philosophies for QoS

Philosophy of other vendors

- Stress prioritization—the more a vendor talks about the importance of prioritization, the more likely it is that the vendor cannot supply a non-congesting network.
- Provide lots of controls to ensure that critical traffic gets through.
- “It’s going to be congested. Deal with it.”

HP philosophy

- Engineer a network design that will avoid congestion in all but extreme cases.
- Then employ switches with easy-to-use, standards-based, interoperable prioritization to ensure critical traffic gets through in extreme cases.

HP ProCurve QoS advantages

- **Provisioning:** The price/performance of HP ProCurve switches make it easy to properly provision your network.
- **Control:** HP ProCurve switches provide industry-standard traffic controls to keep unwanted traffic off your network.
- **Prioritization:** HP ProCurve 5300xl's outstanding class-of-service features will future-proof your network.



HP WORLD 2003

Solutions and Technology Conference & Expo

Interex, Encompass, and HP bring you a powerful new HP World.

