Why there is no Network QoS and what to do about it

Abhay Parekh

IWQOS03

June 2, 2003

Why bother with YAQT?

- After all
  - People have been talking about Network QoS for 15 years
  - No Network QoS in sight!
- And..
  - ATM lost and IP won
  - Per Flow Queueing is too hard
  - Diffserv is implemented but no one seems to use it!
  - No demand for QoS applications
  - Overprovisioning is the way to implement QoS!
- And…(!)
  - ISPs don’t care about QoS
  - Router vendors don’t put new things into routers
    - Adding a feature requires vendor buy-in
    - Vendor buy-in requires customer buy-in
    - Customer buy-in requires general availability

Furthermore…

- We are all familiar with the following objections to QoS:
  - "No Business Model"
  - "QoS is harmful"
  - "No congestion problem"– current quality is fine
  - "QoS work is just incomprehensible research"
  - It's a billing problem!
  - People only like to use the internet for applications that don't need qos
    - Otherwise they prefer to use other networks
- QoS is not going to happen because…
  - It's a BUSINESS problem (from technologists)
  - It's a TECHNOLOGY problem (from business people)

OK…But

- Router vendors still talk about it as a BIG priority
  - QoS one of the four corners of Cisco's strategy
  - Delivering an assured user experience" one of the four ways Juniper wants to lead the network's transformation

Maybe because that's what customers are telling them they WANT from the network

- Users love the enhanced QoS experience
  - In an SBC survey, 96% users who have DSL connections believe that it is the most valuable technology they have their house
  - Thousands of content providers signed up with CDNs to improve the user experience
  - Microsoft has a QoS stack in XP
- QoS sensitive applications are moving to packet-based networks
  - VOIP
  - Radio
  - Sprint just started converting its voice access lines to a packetized format

OK… But

- Overprovisioning is a copout
  - Current network load only accounts for best effort applications
    - Even today, 34% of enterprise CIOs think their network is underprovisioned
  - Heterogeneity is here to stay
- The internet is not as robust, secure, or dynamically configurable as legend has it
- Can’t really scale the internet to include new applications like gaming, voice, video, interactive TV etc. without QoS

The Great Internet Deadlock

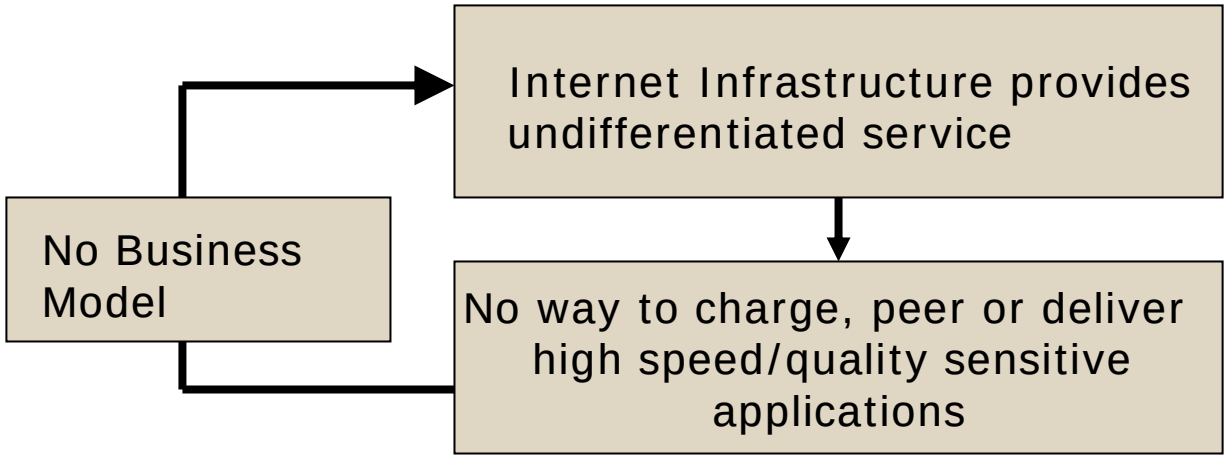

- More capacity is thrown at the undifferentiated network, and emphasis continues on “speeding up the internet”, but this just speeds up existing applications
- No future for internet media or other bandwidth intensive applications
- No future for significant high speed access penetration
- These are huge lost opportunities!!

What's the way out of this?…

- Understand the evolution of the network architecture
  - Above the IP layer
  - Below the IP layer
- Realize that the network is evolving TOWARDS network QoS and not away from it
- Provide an intellectual framework for network QoS that is consistent with the architectural changes
  - The importance of this is not to be underestimated
- Demonstrate concrete results that are directly applicable to the network infrastructure 0-5 years out

Networking Infrastructure Trends Above IP

- Layer 4-7 switching is ubiquitous
  - IP addressing and naming manipulated
- “Overlay networks” are growing in popularity
  - VPNs
  - Security
  - Multicast
  - P2P
- Middleboxes are confounding the end-to-end principles of the IP centric network
- Higher layer forwarding takes away most of the value-added functions of IP
- Let’s look back a long back time ago…

The Network in 1998: Content Provider Frustration

- Huge potential of the Internet Delivery obvious
- World Wide Wait
- Cloud confusion: no one to ensure end-to-end performance
  - Conflicting service provider incentives: uneasy equilibrium
- New functions like multicast had limited traction
- Streaming audio and video exploding but no support within the network (Broadcast.com, RealNetworks)
- Industry held view of the problem:

  Application waits for Infrastructure

  Infrastructure waits for Application

Content Networking

- Offer a service to content owners over the existing network **based on application performance**
- Change the currency of the internet service from bits to information
- Use co-location facilities at different svc providers to deploy an **overlay network** that overcomes cloud confusion
- Enable new business applications to be deployed over the internet quickly by using the overlay network
- Higher Performance allows allow media and other services

Content Internet Value Chain

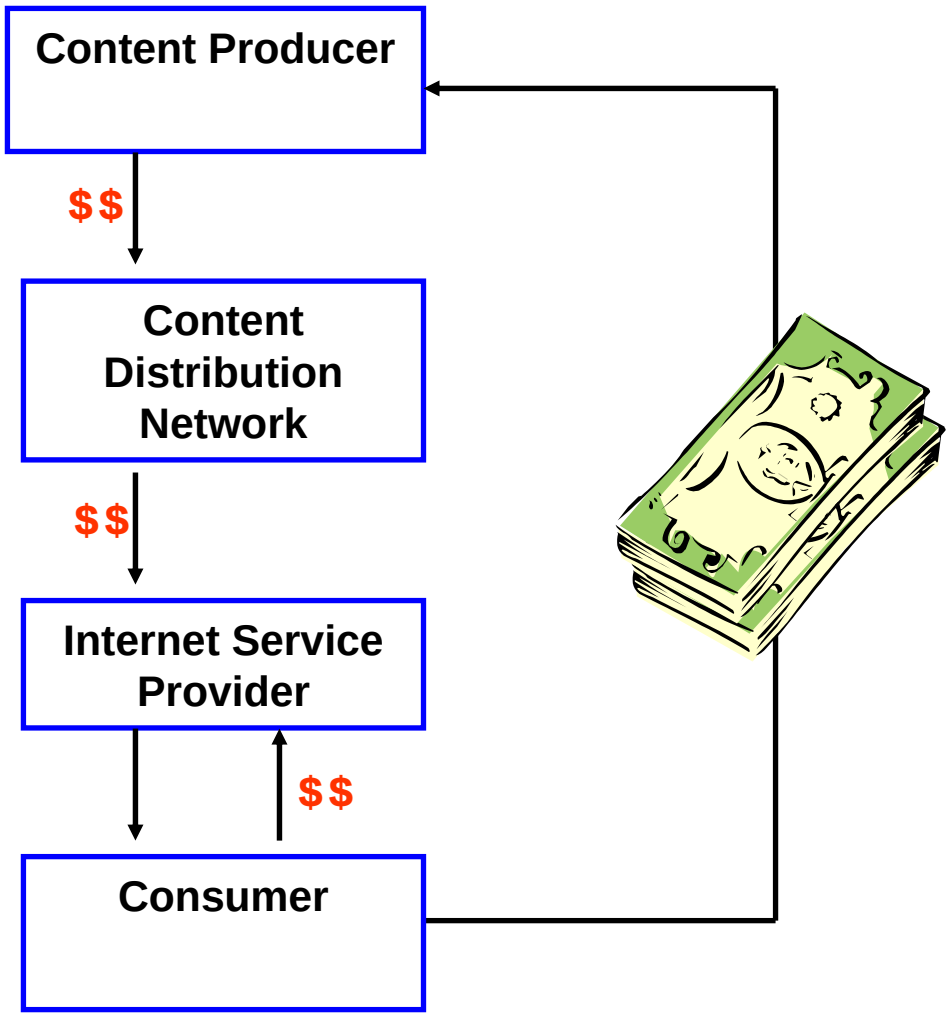

The Media Internet: 1999

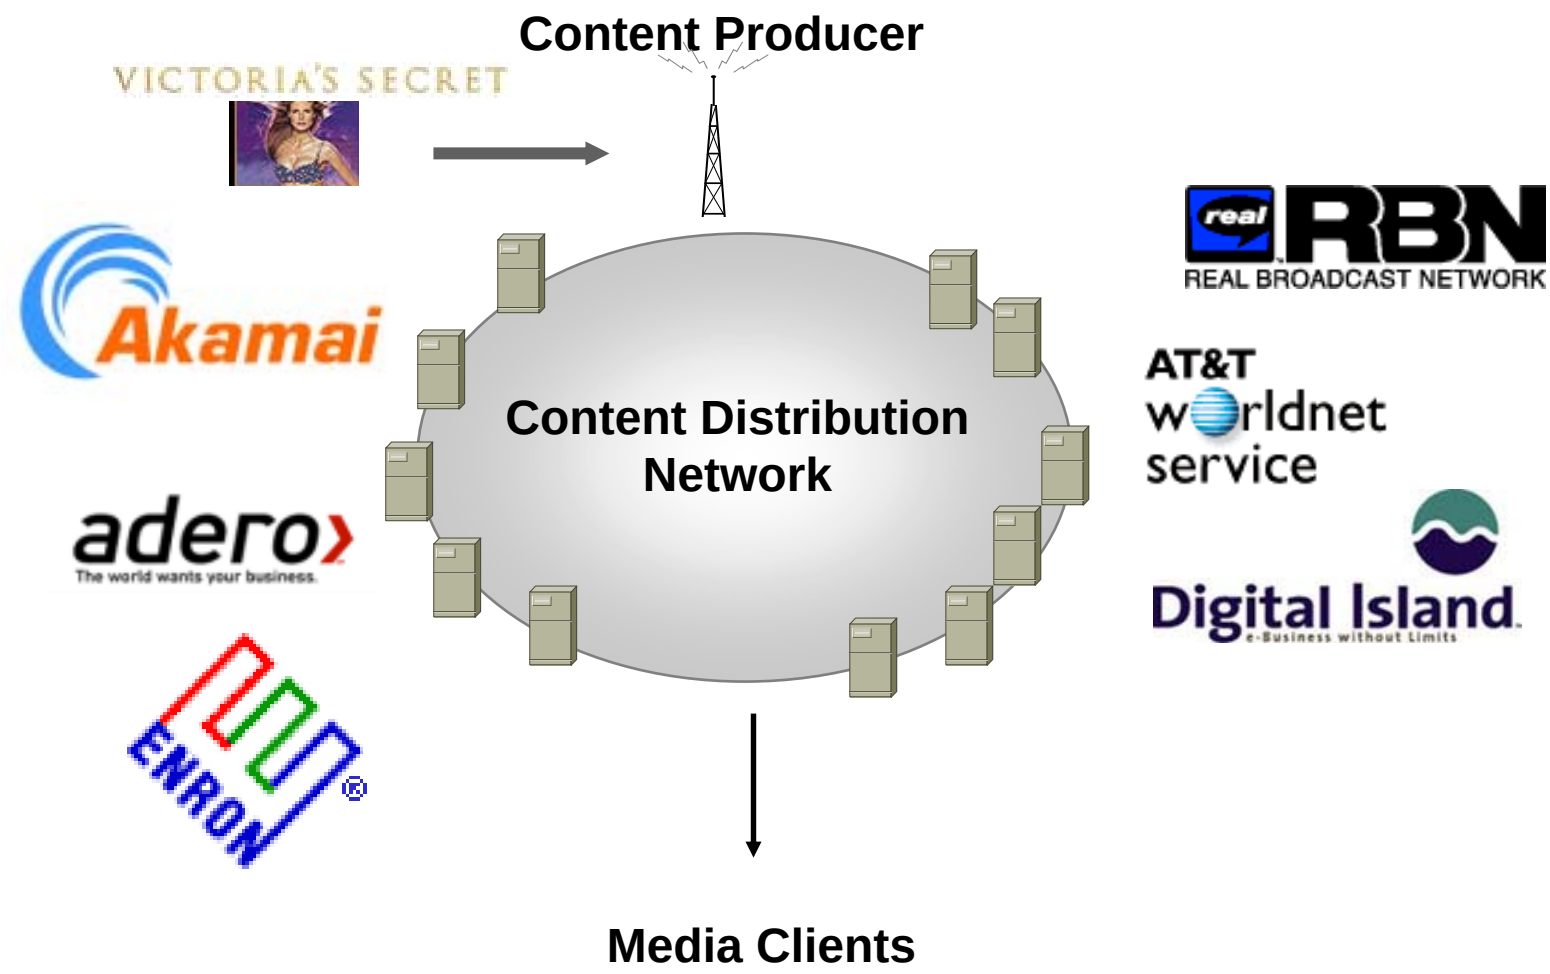

FastForward Networks Mission

To convert the Internet into the “next generation” broadcasting medium, allowing service providers and content distributors to broadcast hundreds of thousands of channels to millions of simultaneous viewers

Inktomi

fastforward networks

Why the Internet for Broadcast?

- Broadcast Today
  - Small number of players control broadcasting
  - Limited by geographical boundaries

fastforward networks

Why the Internet for Broadcast?

- Anyone can be a broadcaster
- No geographical boundaries

netnoir.com
theblacknetwork

U.S.News ONLINE

Internet

food NETWORK

with online

fastforward networks

Why the Internet for Broadcast?

- Audience Tracking and Management
  - It’s not just radios and TVs!

TiVo, TV your way.™

fastforward networks

Why not just use IP Multicast?

- Too enmeshed with layer 3 to be usable
  - Problems require router code to be debugged
- Internet-wide multicast dial-tone too much to require
- Service model gives no ISP an incentive to peer
  - Interdomain routing protocols unclear
- Need real-time audience statistics, priorities in broadcasts, billing, etc etc.
- We picked an overlay architecture

The Broadcast Internet

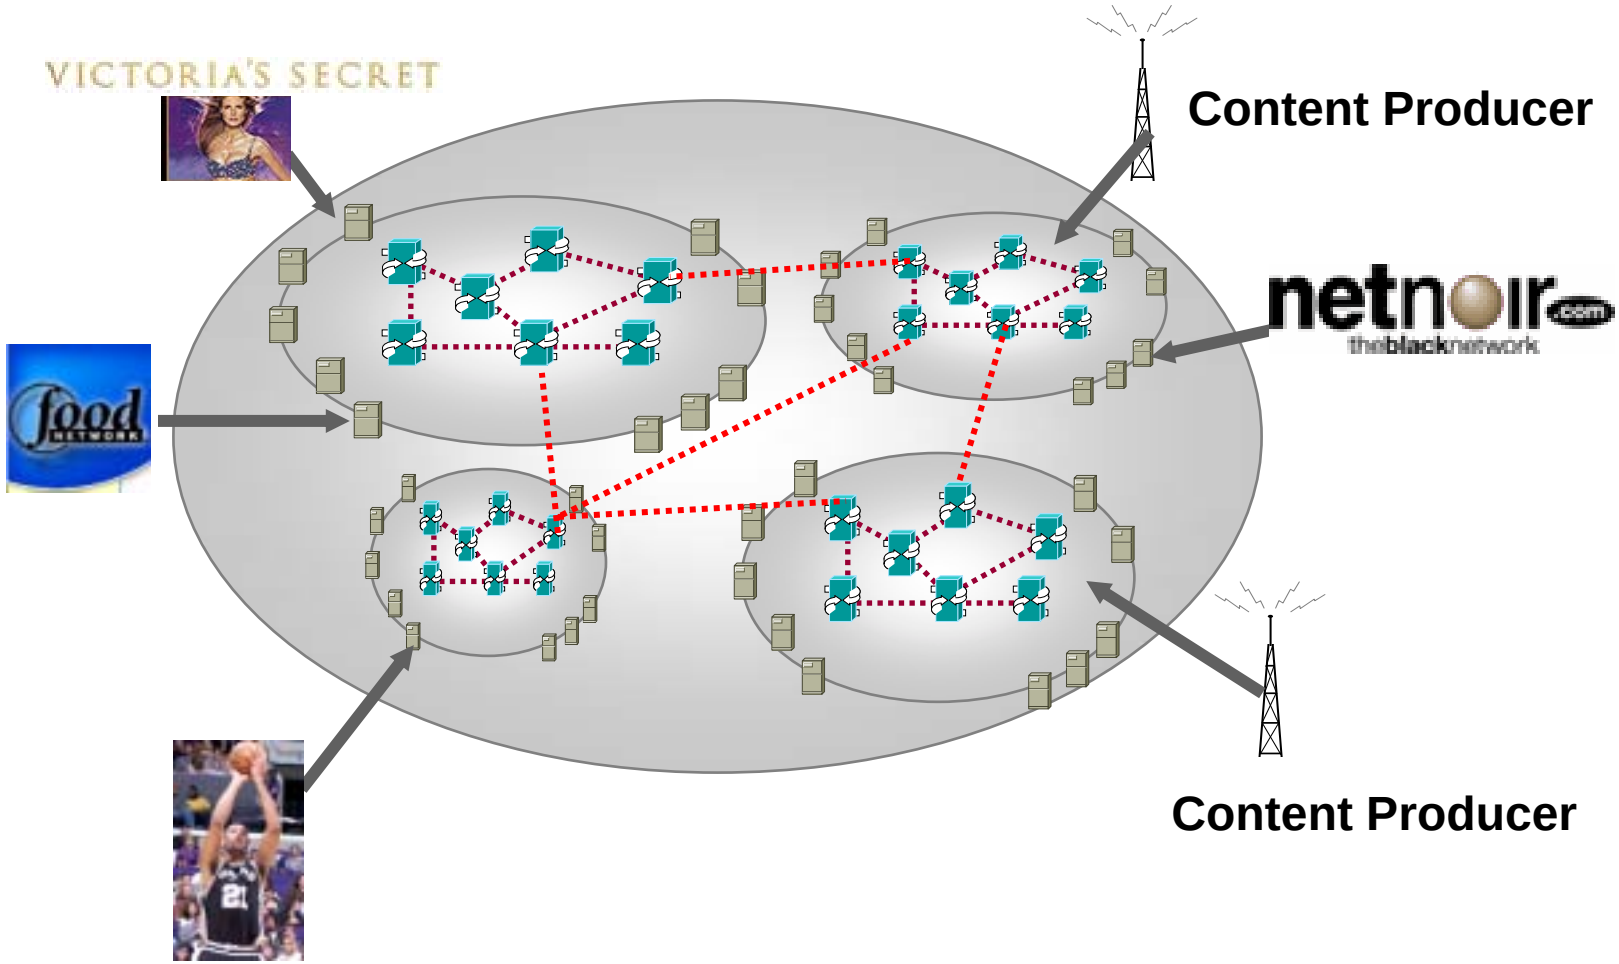

fastforward networks

Broadcast Overlay Architecture

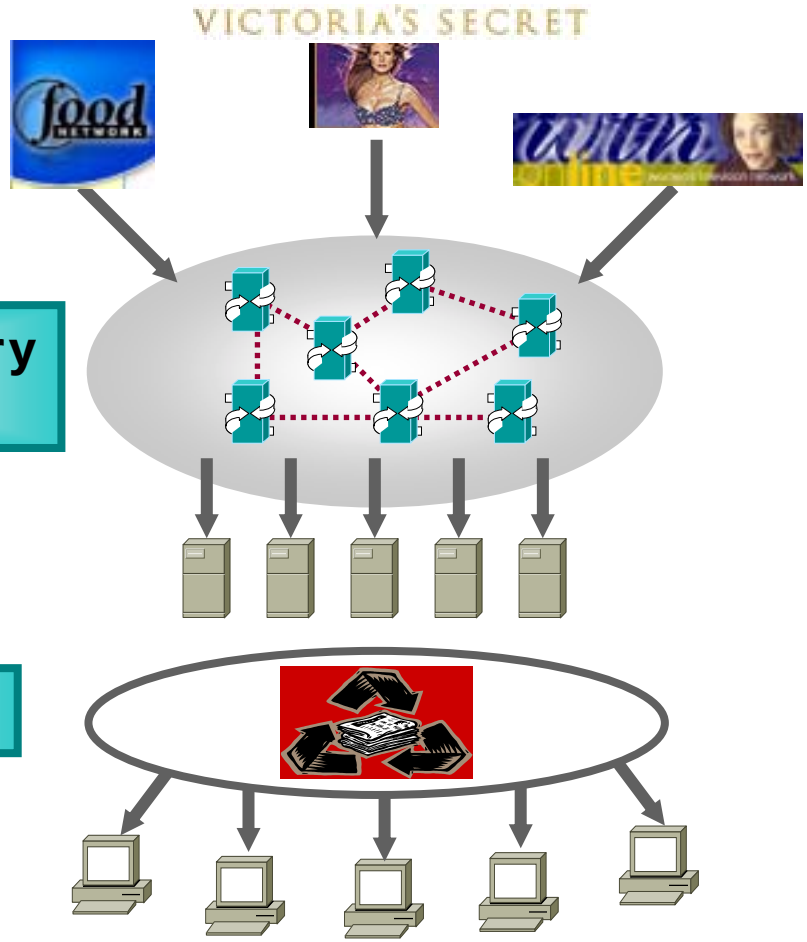

Management Platform

content management
injection & real-time control

network management
monitoring & provisioning

server management

redirection management
load balancing
system availability

viewer management
subscriptions, PPV, monitoring, Neilson ratings, targeted advertising

fastforward networks

Broadcast Management

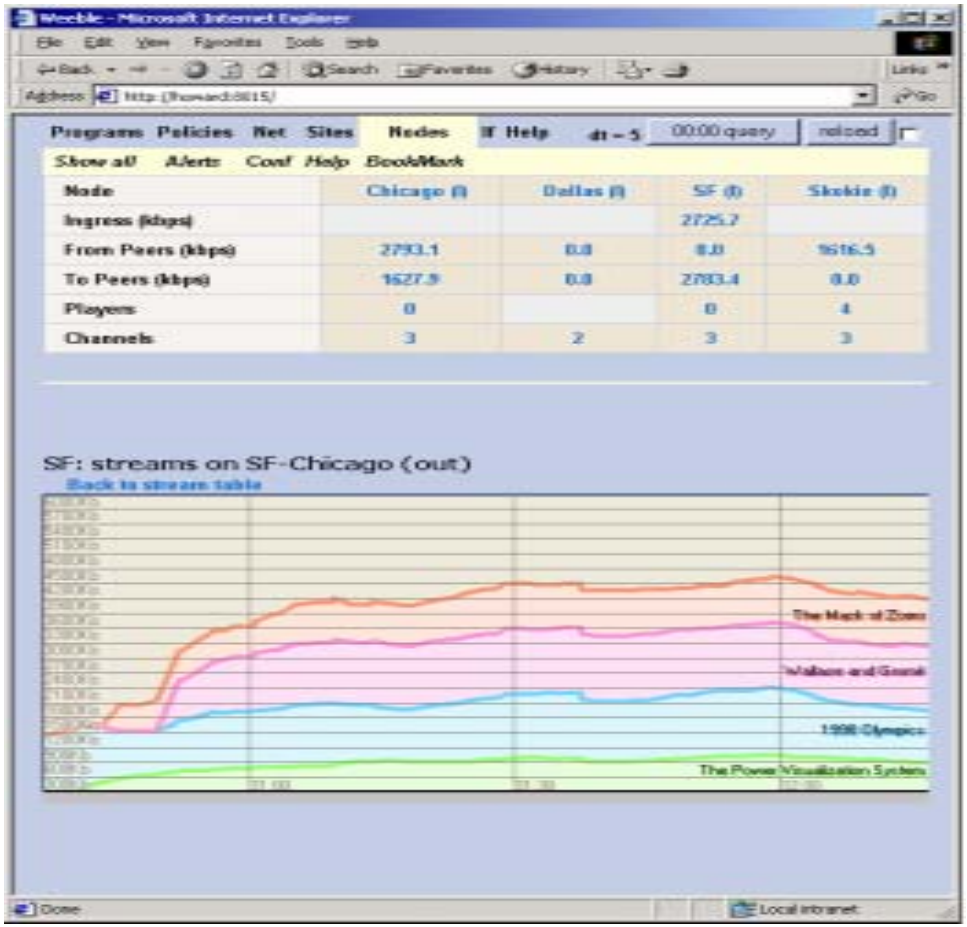

- Application-level information for management and tracking
- Works across multiple networks
- Content Producer event programming with ad-hoc query audience statistics

fastforward networks

Broadcast Manager

Node Information

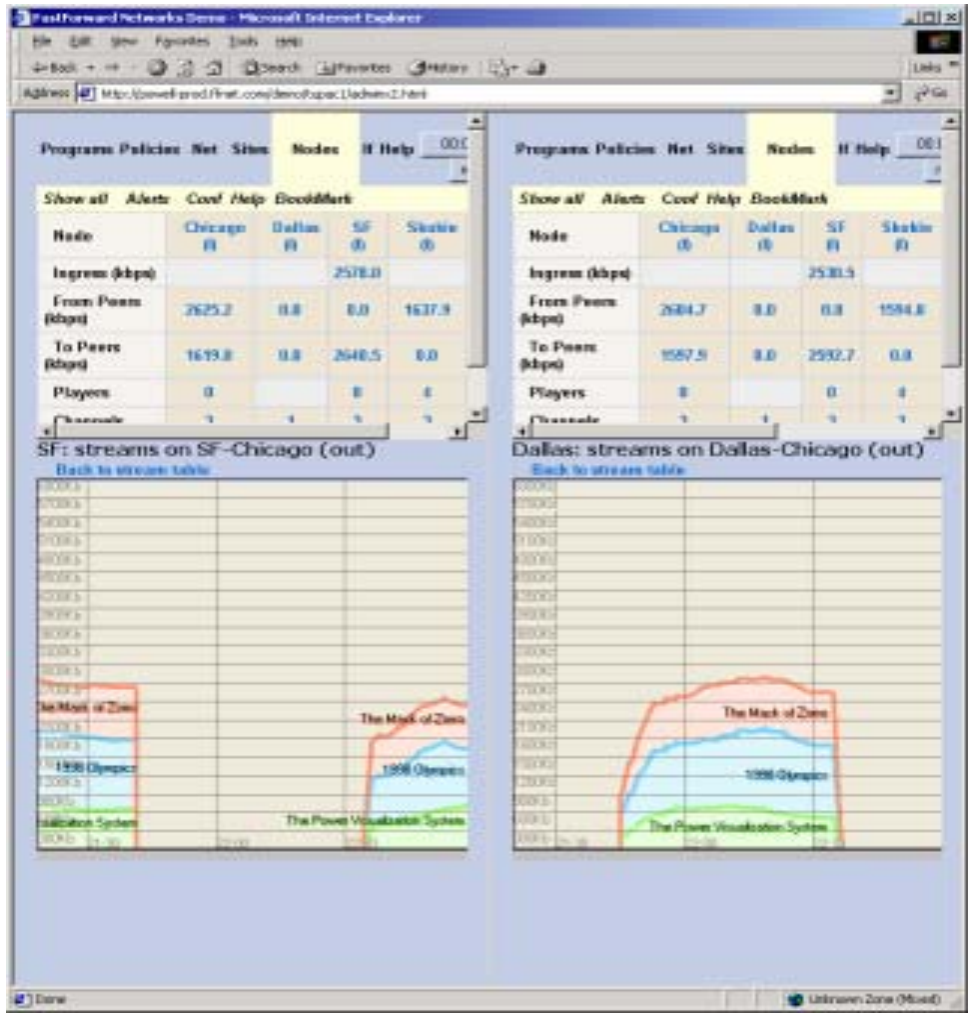

fastforward networks

Policy Management

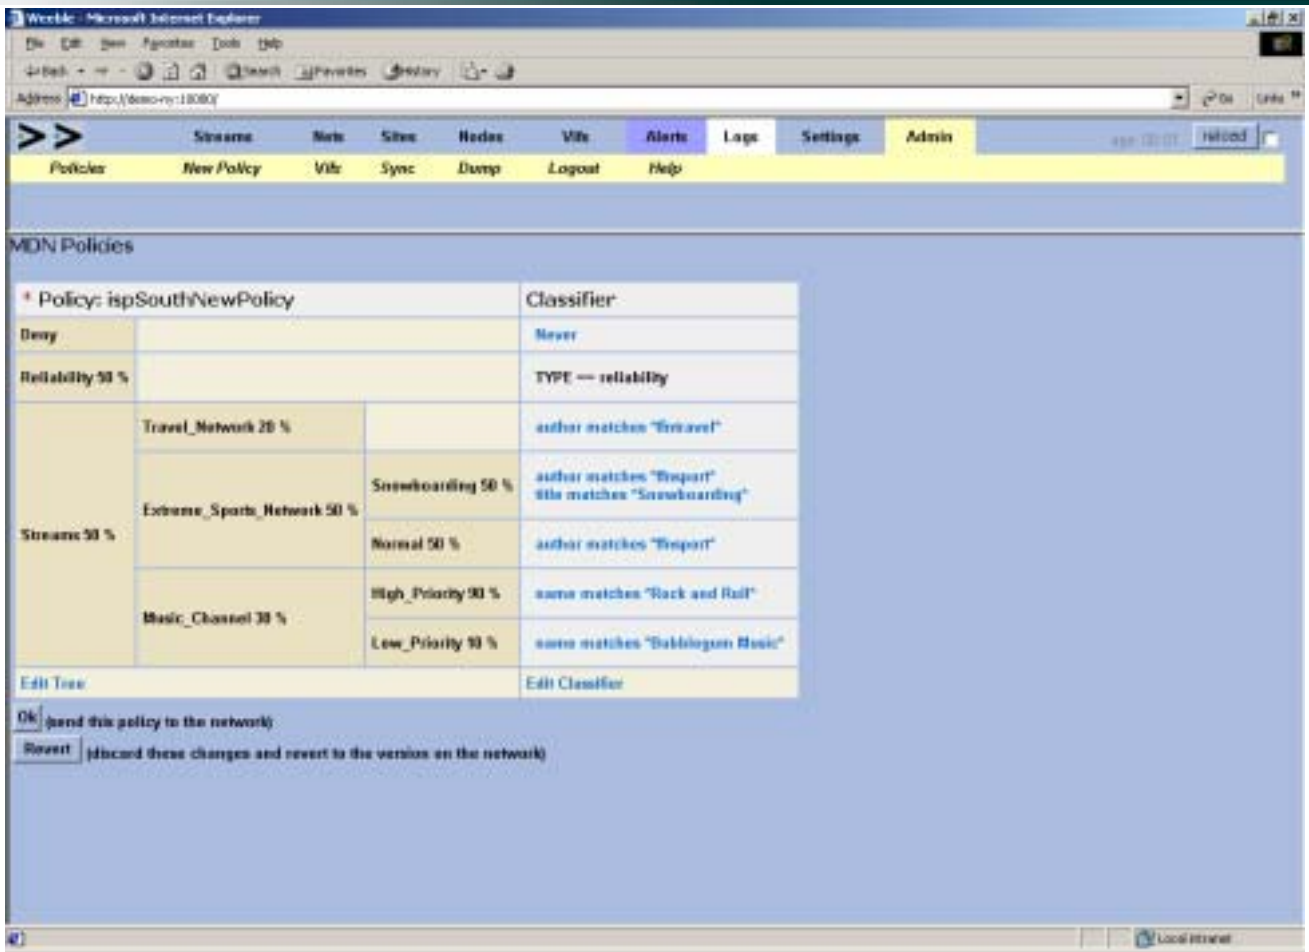

fastforward networks

What happened: 1998-2000?

- Thousands of Content Providers signed up with CDNs
- Even mainstream content producers got excited about internet broadcasting
- Our stuff worked
  - We did huge broadcasts e.g. Big Brother
  - We improved lossy links and rerouted a lot of video
  - We had Digital Island peering with AOL
  - We made a multi-million dollar sale
  - We merged with Inktomi for about $1.3B

What happened: 2000-2002?

- Most content distribution networks went bankrupt
- The major caching companies were unable to build businesses
- Streaming is still prevalent but has not taken off the way we expected it to.

- The content distribution vision had an impact, but did not materially affect QoS.
- Layer 4-7 switching matured into a new product category

Why did this happen?

- Overlay Networks didn’t work as well as advertised in improving performance
  - Always at the mercy of underlying facilities
  - Too Expensive to build a ubiquitous network
  - Local caches are cheap and available
  - Overlay Networks that did not depend on network performance did fine
- Bottom Line: Can’t solve QoS problems just with Overlays!!
- But Overlay’s are here to stay!

Network Infrastructure Trends Below IP

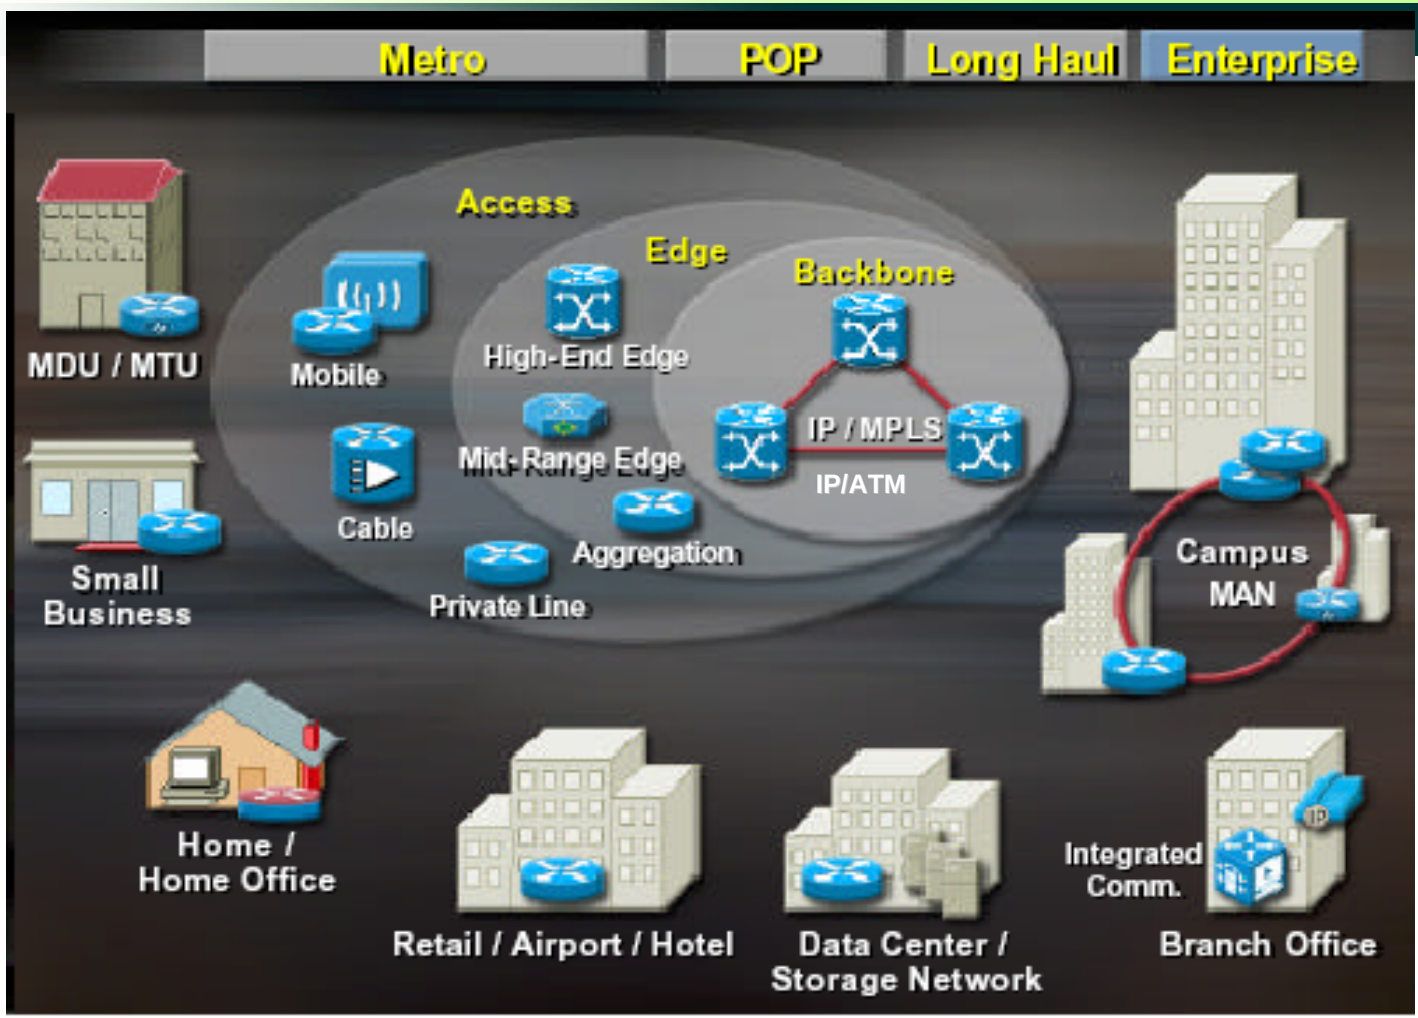

(Figure adapted from a slide on the Cisco website)

Access

- DSL is growing faster than people think
  - 10% penetration in 2001
  - 100x growth from 1997-2001
- Most of the US cable plant is unicast capable
- Over 40% of cable boxes in the US are digital
  - No need to “broadcast” same set of channels to all

- More users are accessing the internet for the first time using broadband than using narrow band
- Some countries are making broadband deployment a national priority (Korea)
- In an SBC survey, 96% users who have DSL connections believe that it is the most valuable technology they have their house

The great optical buildout

- Large amounts of fiber have been laid in the Metro and Long Haul
  - Much of it has not been “lit up”
- Metro bottleneck is largely an issue of the costs of lighting up fiber coming down
  - Moore’s law will solve the Metro bottleneck
- Will Ethernet replace SONET?
  - Maybe, but not for a while
- Enough bandwidth in the network infrastructure to support high bandwidth, qos-sensitive applications

New Label Switching Layer

- IP is not THE forwarding layer
  - Virtual circuit “underlays” using ATM and increasingly, MPLS are the way to go
- Virtual circuits at the “trunk” level provide SLAs and a basis for traffic engineering
  - Traffic engineering circumvents poor IP routing protocol performance
  - GMPLS integrates optical devices into traffic engineering
- Current thinking: All provider backbones will converge to either IP/MPLS or IP/ATM
- DSL uses ATM transport even though the host interfaces use Ethernet
- Why not just use IP forwarding?
  - Flows are essential to networking
  - Not sufficient to manage flows end-to-end
- Isn’t this what QoS researchers have been saying for 15 years?

QoS in the Label Switching Layer

- Currently focused on providing better trunks and VPNs
  - Signaling is not as dynamic as required
  - Call admission not an issue
  - Interdomain issues not addressed
- But
  - MPLS uses RSVP to create tunnels
  - Possible to improve the granularity of control
  - Easier to address interdomain issues with virtual circuit routing

Bottom line on network infrastructure

- Lots of available bandwidth
- Forwarding occurs at virtually all levels of the protocol stack
- Two layers are emerging
  - “Infrastructure Services” Layer above IP and
  - “Label Switching” Layer below IP
- Mechanisms deployed to support QoS – but below the IP layer in the Switching layer
- Overlay networks can add new functionality without depending on router upgrades
- Many of the ingredients required to enable QoS are already available
- Two cavaets
  - Interdomain problems have become worse (everyone wants to be an AS)!
  - Strict layering hides QoS mechanisms from applications

Interdomain problems

- Accountability for traffic is not as a big a problem as it is made out to be
  - Look at the circuit switched world
- Service providers are not indifferent to QoS -- there aren't too many ways for them to make money if they are.
  - However in the current environment there isn't much incentive to support QoS
- Virtual circuits make peering a lot easier
  - Especially when bits have differential value
  - Peering should be done at the Label Switching Layer
- The harmful explosion of /24 AS's is a consequence of the inefficiencies of an IP centric network model

10 Ways to Make Network QoS Happen

1. Don’t try to come up with a “complete” definition of QoS
   - Sometimes generality dilutes relevance
   - How about focusing on voice and video as concrete applications?
2. Rethink the Diffsrv approach to QoS
   - This is too vague to justify the cost of deployment/use
3. Peering cannot just be in the form of undifferentiated bits
   - How about voice “trunk” peering?
   - Adding differential value to bits will alleviate peering pathologies
4. Routing Protocol pathologies can’t be ignored for ever
   - Serious problems with dynamic routing
   - Fortunately things get better in virtual circuit networks
5. View the IP layer as a “connectivity layer” NOT the QoS forwarding layer

6. Don’t waste time putting QoS mechanisms in the IP layer
   - IP will be the pre-eminent connectivity layer but not much more
7. Infrastructure Services Layer should communicate directly with the Label Switching layer directly (by reading each other’s headers).
   - Keep the IP layer simple, i.e. out of QoS
8. No need to oversell QoS
   - The network is evolving in a way that is conducive to its success
9. Refocus pricing work away from the undifferentiated, best effort model
10. Continue to “Re-think the fundamentals” -- the network architecture is not what it once was.

Conclusion

- Network QoS is not a lost cause.
- The network is evolving to support it but NOT at the IP layer
- Now is the time to provide the intellectual leadership to create models and to move thinking towards a new level of network performance which can unleash a new set of applications
- The true potential of the internet will only emerge when the ultimate modes of communication are an integral part of it: voice and video