

Next Generation Edge and Overlay Services



Internet2 Spring Members Mtg Tutorial

Micah Beck & Terry Moore

Center for Information Technology Research

University of Tennessee



Outline

- **Background: Edge and Overlay Services:**
- Content Delivery Networks
- Other Approaches
- Two important trends
- The Promise of the Grid
- Finding the future

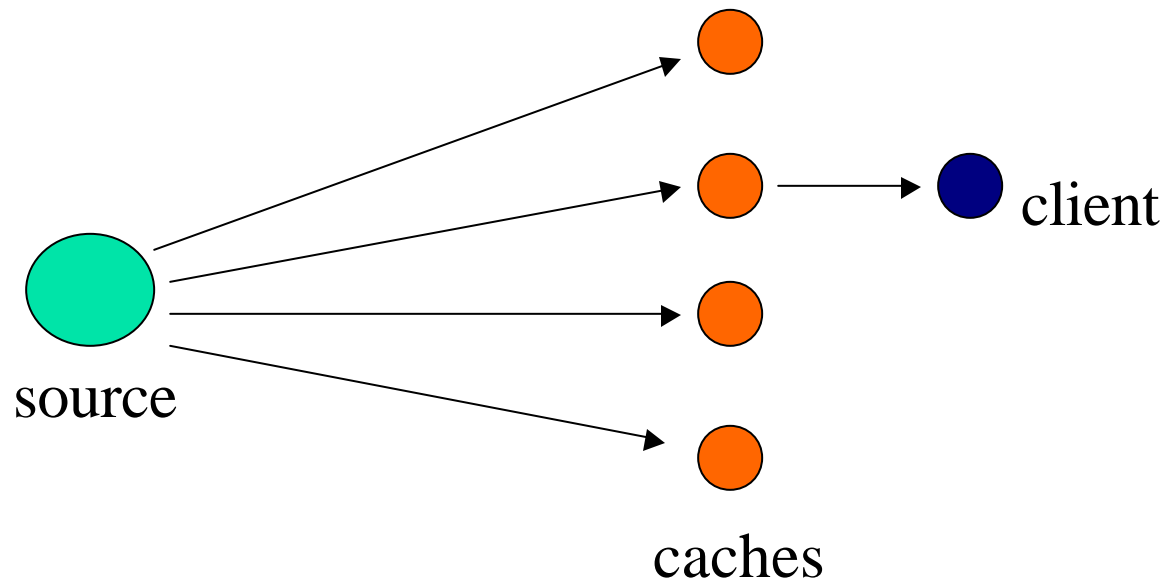


What are Edge and Overlay Services?

- IP services are defined in terms of end-to-end delivery of datagrams
- Stateful protocols only for control state
 - Multicast
 - QoS
- Edge and Overlay Services introduce application level protocol intermediaries

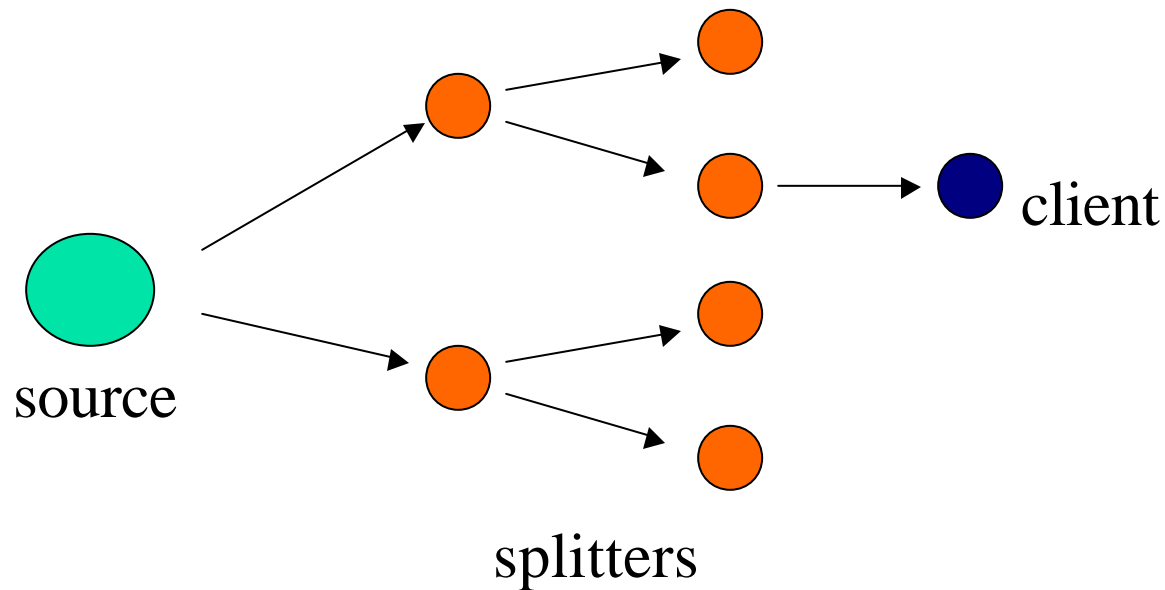
The Prototypical Edge Service: Web Caching

- HTTP response capture and replay
- TCP proxy intermediaries



The Prototypical Overlay Service: Application Level Multicast

- Streaming video splitters
- Statically configured UDP relays





Why are Edge and Overlay Services Important Today?

- Commodity backbone congestion
- Server overload
- Latency
- Flexibility/openness
- Deployability
- Resource scalability



Issues with Edge Services

- Request redirection
 - IP hijacking
 - unreliable DNS resolution hacks
- Non-standard protocols
 - proprietary streaming (Real, QT, MS)
 - proprietary edge/overlay services
- IETF Mantra:
Scalability, mangability, interoperability!



Edge Services and Advanced Networks

- Edge networks are very capable
- Backbones support massive aggregate load
- High performance applications can overwhelm servers
- TCP problems at high bandwidth x delay
- Latency: Complaints to Prof. Einstein
- Flexibility, deployability, resources



Outline

- Background: Edge and Overlay Services:
Edge and Overlay Services
- **Content Delivery Networks**
- Other Approaches
- Two important trends
- The Promise of the Grid
- Finding the future



Commodity CDNs

Akamai, Digital Island, Mirror Image

- Content Delivery \approx Content Distribution
- Cache resources leased to content providers
 - disk space partitioned
 - bandwidth metered
- Colocation cheap or free
 - Local network saves bandwidth
 - Reselling outbound bandwidth
- Complex DNS redirection algorithms



Infrastructure Vendors

Inktomi, Cisco, Novell

- Service providers have taken the lead
 - Service providers serve as content brokers
 - ISPs are not experienced at running CDNs
 - No interoperable infrastructure
- Infrastructure vendors serve enterprises
 - Self contained content provider and ISP
 - Single infrastructure provider
- Digital Island Uses Inktomi Caches



If Knowledge is Power, What do you know?

- Control of communication has always been mission critical for Universities
 - University Presses controlled their own printing for centuries
 - Campuses buy and operate telephony switches
 - Campuses control IP services through Internet²
- Why cede control of wide area publication to CDN service providers?



IETF activities

- Web Replication and Caching (WREC)
 - Taxonomy, requirements
- Content Delivery Internetworking (CDI)
 - Settlements, SLAs, property rights,
- Web Intermediaries (WEBI)
 - Content Invalidation Protocol
- Open Pluggable Edge Services (OPES)
 - Rules-based invocation of proxy services



Outline

- Background: Edge and Overlay Services
- Content Delivery Networks
- **Other Approaches**
- Two important trends
- The Promise of the Grid
- Finding the future



Peer-to-Peer: Cooperative Overlay Services

- Edge resources on client workstations
- Cooperative sharing of resources
- No (or minimal) settlements
- Very specific services
- High end-user value
- Whoever organizes a service takes responsibility for it (e.g., Napster)



Application Service Providers: Data Centers at the Edge

- Applications installed on servers in distributed data centers
 - semistatic installation
 - location transparency within cluster
- Flexibility & licensing advantages
 - usage fees enable business model
- Clients access via remote user interface
 - some OS specific mechanisms (Citrix)



Rule-Based Active Edge Services

- OPES distributes proxylets
 - Intermediary Rule Markup Language
 - Filtering, ad insertion, transcoding, translation
- Internet2 Distributed Storage Infrastructure (I2-DSI) distributes static content and servelets Portable Channel Representation
 - Database access, high interactivity, localization
 - WWW10 paper with Lokomo Systems



Outline

- Background: Edge and Overlay Services
- Content Delivery Networks
- Other Approaches
- **Two important trends**
- The Promise of the Grid
- Finding the future



Storage Will Rock Your World!

New IBM Storage Products

- **Ultrastar 36Z15**
- 35 GB, 4MB cache
- 52.8 MBs sustained xfer
- 2 ms average latency
3.4 ms seek time
- 10.7 Gb/in² inch density
- 26H/102W/146D (mm), .75 kg
- 6 disks, 12 heads
- **Microdrive DSCM 11000**
- 1 GB, 128K buffer
- 2.6 MBs sustained xfer
- 8 ms avg latency
12 ms avg seek time
- 15.2 Gb/in² areal density
- 5H, 43W, 36D (mm), 16 g
- 1 disks



Reliance on Network Services

- Critical services: Real Time Market Info.
- Workflow: Publishing, Graphics
- Organizational Coordination: Calendar, Directory, Scheduling, etc
- Distance Learning: Lectures, Exams
- Subscription Content: Law, Medicine, Porn
- Napster: I burned my CDs!



Outline

- Background: Edge and Overlay Services
- Content Delivery Networks
- Other Approaches
- Two important trends
- **The Promise of the Grid**
- Finding the future



Computing Grids

- Remote access to computational resources
- Operating system style: Legion
- Distributed Machine Room: Globus
- Other approaches:
 - Condor (UW)
 - NetSolve (UT)
- Standardization: Grid Forum



Data Grids

- Data Intensive Computing Environments
 - Digital Libraries and Large Data Sets
 - SDSC Storage Resource Storage Broker
- CERN Large Hadron Collider (Gryphon)
 - 1 PB/year data distribution
 - Continental/regional/campus centers
- Internet Backplane Protocol (UT)



Access Grids

- Advanced Telemedicine
- Remote Instrumentation
- You know the story...

Video's the one thing we **can** do!



Logistical Computing and Internetworking (UT)

- Logistical Internetworking: Flexible coscheduling of storage and bandwidth
- Logistical analogy: Movement of industrial materiel using trucking and warehouses
- LoCI: Add computational resources
- Middleware components
 - Internet Backplane Protocol
 - Network Weather Service
 - NetSolve
- Common resources for communities, campuses

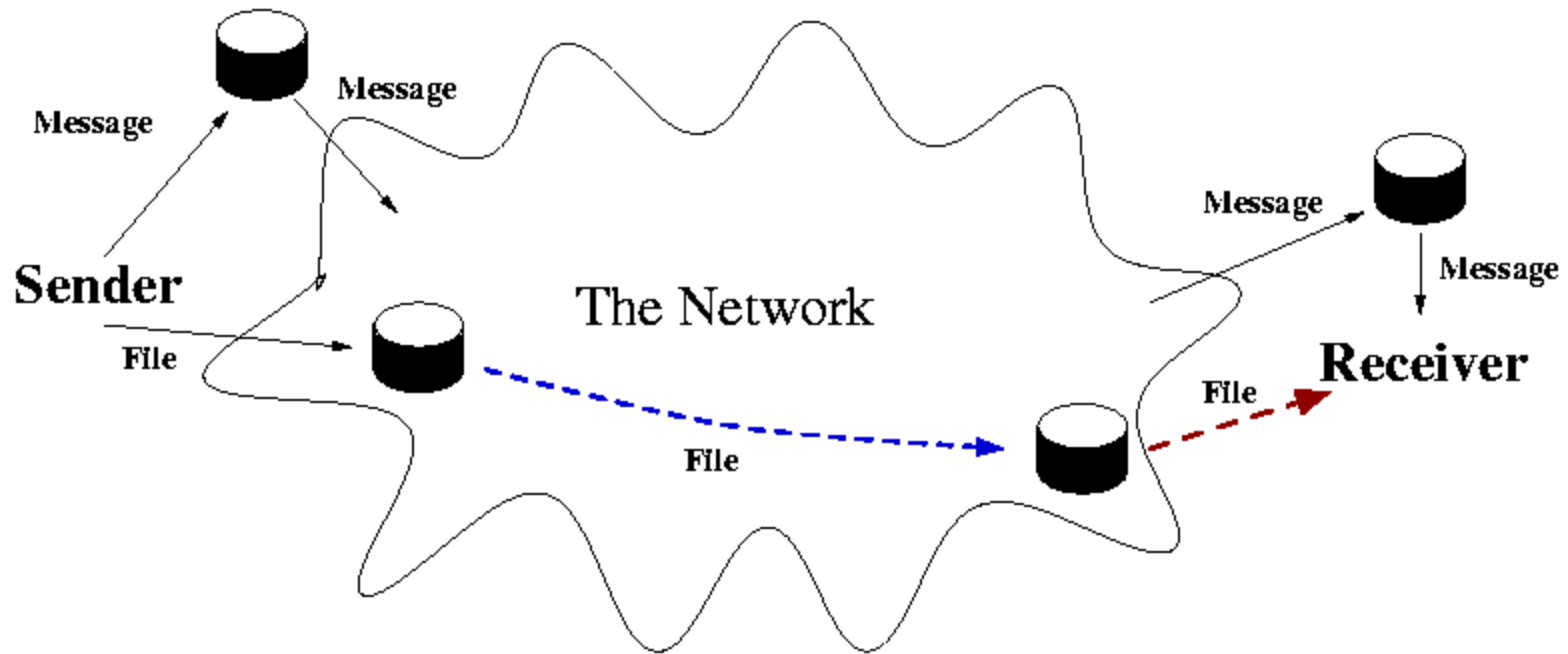


Internet Backplane Protocol

<http://icl.cs.utk.edu/ibp>

- Storage analogy to IP
- Lightweight allocation
- Weak semantics for a storage service
 - Unreliable
 - Time limited
 - Volatile
- Model close to physical resource
- Analogy: processor backplane

IBP Mail





Outline

- Background: Edge and Overlay Services
- Content Delivery Networks
- Other Approaches
- Two important trends
- The Promise of the Grid
- **Finding the future**



Searching Where the Light Is

- Scene: Drunk on hands and knees under street light
- Cop: What's up?
- Drunk: Lost m' keys!
- Cop: I'll help you look. Did you drop 'em right here?
- Drunk: No, I dropped 'em over there (points to darkness in the middle of the block)
- Cop: So why are you looking over here?
- Drunk: The light's much better!



The Loneliness of the Long Distance Fiber

- Stateful, reliable, ubiquitous services are shaping the future of network apps (It's just not what Internet² does!)
- These attributes are not defined as “advanced”
- Exciting commercial apps are ruled out
- We've locked ourselves in our campuses and we won't come out!



What we do: High bandwidth wide area TCP to the desktop

- Without easy-to-use reliable, high bandwidth connections the application community will remain restricted
- High bandwidth x delay product makes TCP very hard to tune automatically
- Many valuable edge and overlay services do not require end-to-end connectivity
- Internet2's Procrustean bed: **sleep tight!**



A Vision of Inclusive Advanced Networking

- Criterion: New uses of the network enabled by the Internet² project
- Criterion: Advance the mission of Internet² member universities
- View IP connectivity as a means, not an end
- Serve user communities, not router jockeys
- We belong to the future, not it to us