September 25, 2009

Designed for Change: End-to-End Arguments, Internet Innovation, and the Net Neutrality Debate

Presenter:
Richard Bennett, Research Fellow, ITIF

Respondents:
Dr. David Farber, Carnegie Mellon University
Christopher Yoo, University of Pennsylvania Law School
John Day, Boston University Metropolitan College
Dr. William Lehr, Massachusetts Institute of Technology

ITIF
The Information Technology & Innovation Foundation

Overview

- Net Neutrality arguments offer views of the Internet architecture.
- At best, the neutralists' architecture is overly simplistic.
  - **The claim that the Internet loves all its children the same is false: content is privileged over communication.**
  - **No changeless technology is future-proof.**
  - **Primary sources don't agree with popularizers.**
- Understanding why the Internet is the way it is helps us to differentiate good regulations from bad ones.
- Let's look critically at Internet architecture, regulatory frameworks, and challenges.

The Internet's Foundation Principle

- **"Design for Change" is the paramount notion.**
  - This enabled the Internet to move from the research setting to the general use scenario.
- RFC 1958 says it best:
  - *The principle of constant change is perhaps the only principle of the Internet that should survive indefinitely.*
- End-to-End arguments are secondary.
  - They should be about flexibility.
  - Architecture is only important as an application enabler.
- End-to-End fundamentalists have stood the Internet's dynamism on its head.
  - The design job is far from over.
  - Don't make the Internet the rotary phone of tomorrow.

How the Internet Architecture Evolved

- Let’s look at the two major precursors:
  - ARPANET
  - CYCLADES

ARPANET

- Plan devised in 1967 by Larry Roberts
  - Roberts had a number of inspirations: Licklider, Baran, Davies
  - Developed by BBN
  - Deployed in 1969 on four nodes
- Connection-oriented Network Service
  - X.25 was one successor
  - ISO CONS was another
- Meant to be a production network, not a research tool
  - Prove the packet-switching concept
  - Frustrating for network researchers
- Original home of the Internet
  - TCP replaced ARPANET NCP
  - ARPANET controls covered over TCP/IP weaknesses.

ARPANET Architecture

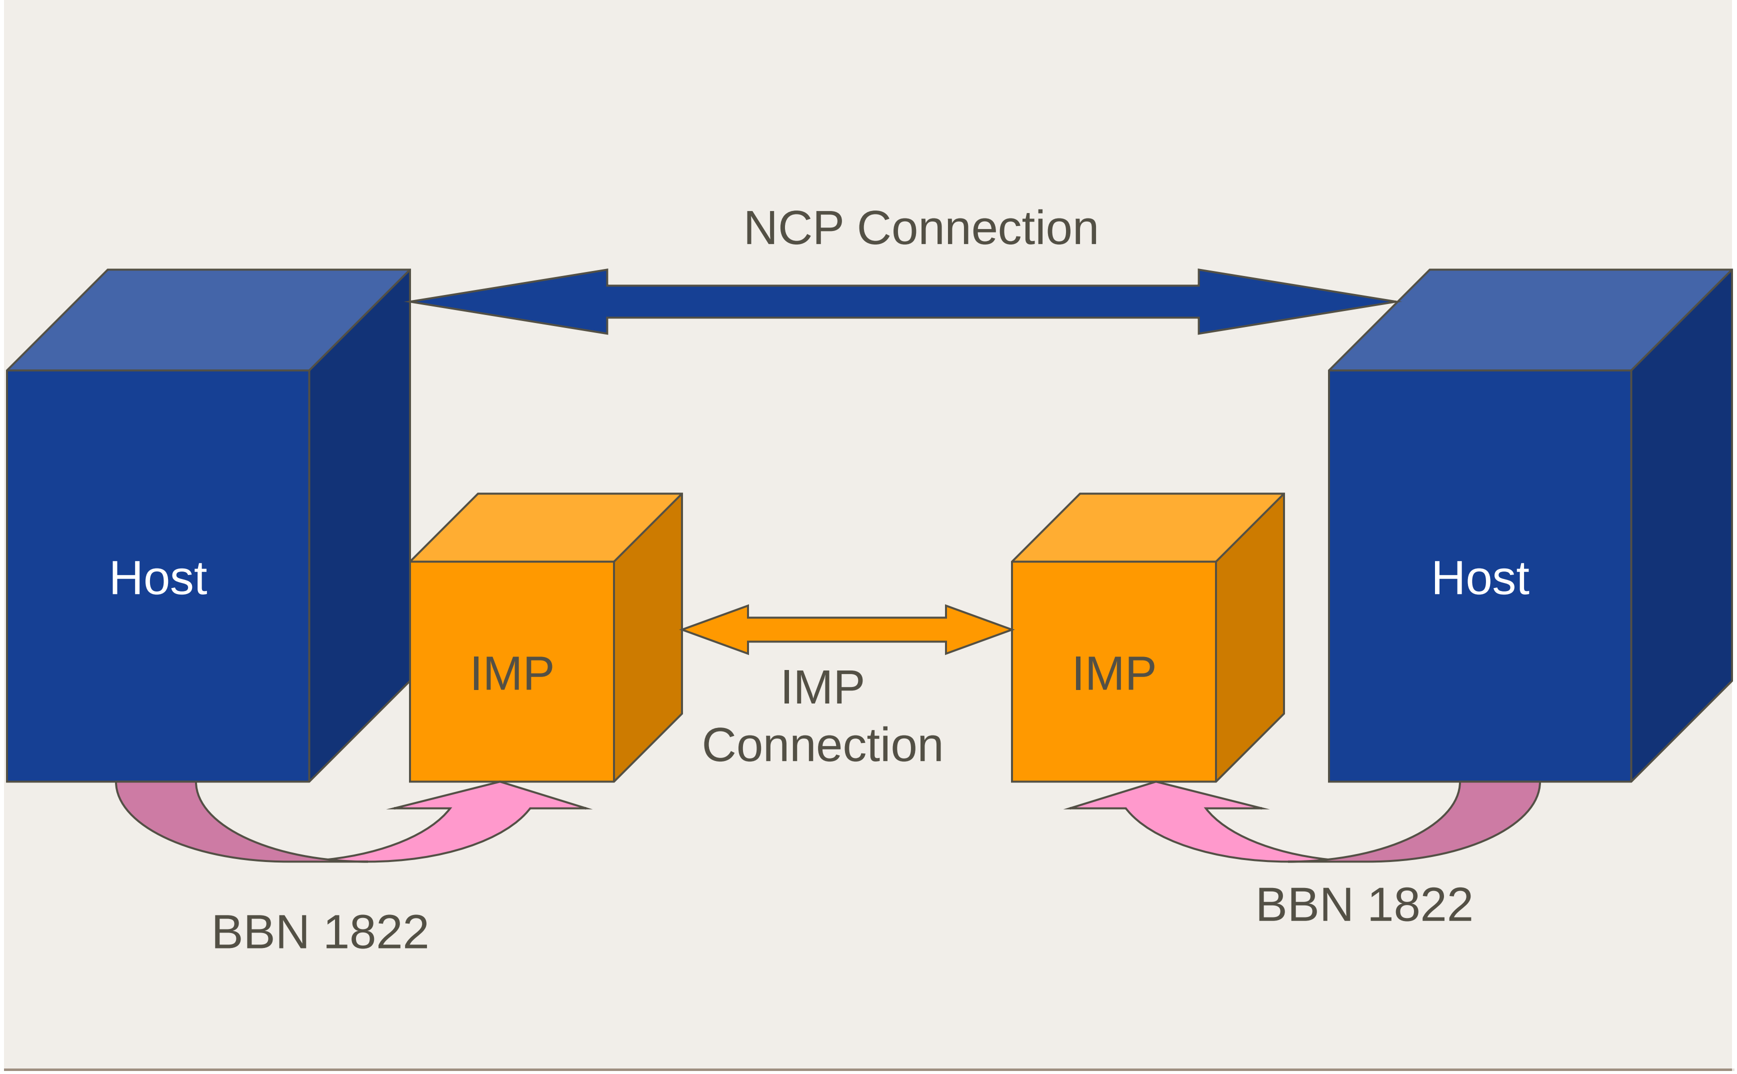

CYCLADES

- Designed by Louis Pouzin and colleagues at IRIA in France starting in 1972, with help from BBN
  - A genuine research network.
- Pioneered the structure of the modern wide-area network
  - Structured protocol design
  - Datagram Network layer
  - End-to-End Virtual Circuit Transport layer
- Heavily influenced the design of the networks of the 1970s
  - DECNet, XNS, ISO/OSI, and Internet
  - Each of these put an end-to-end transport over a datagram network layer
- Strong influence on TCP design
  - Gerard LeLann worked on TCP during 1974 sabbatical from IRIA
  - BBN’ers also contributed design assistance to the TCP team

CYCLADES Architecture

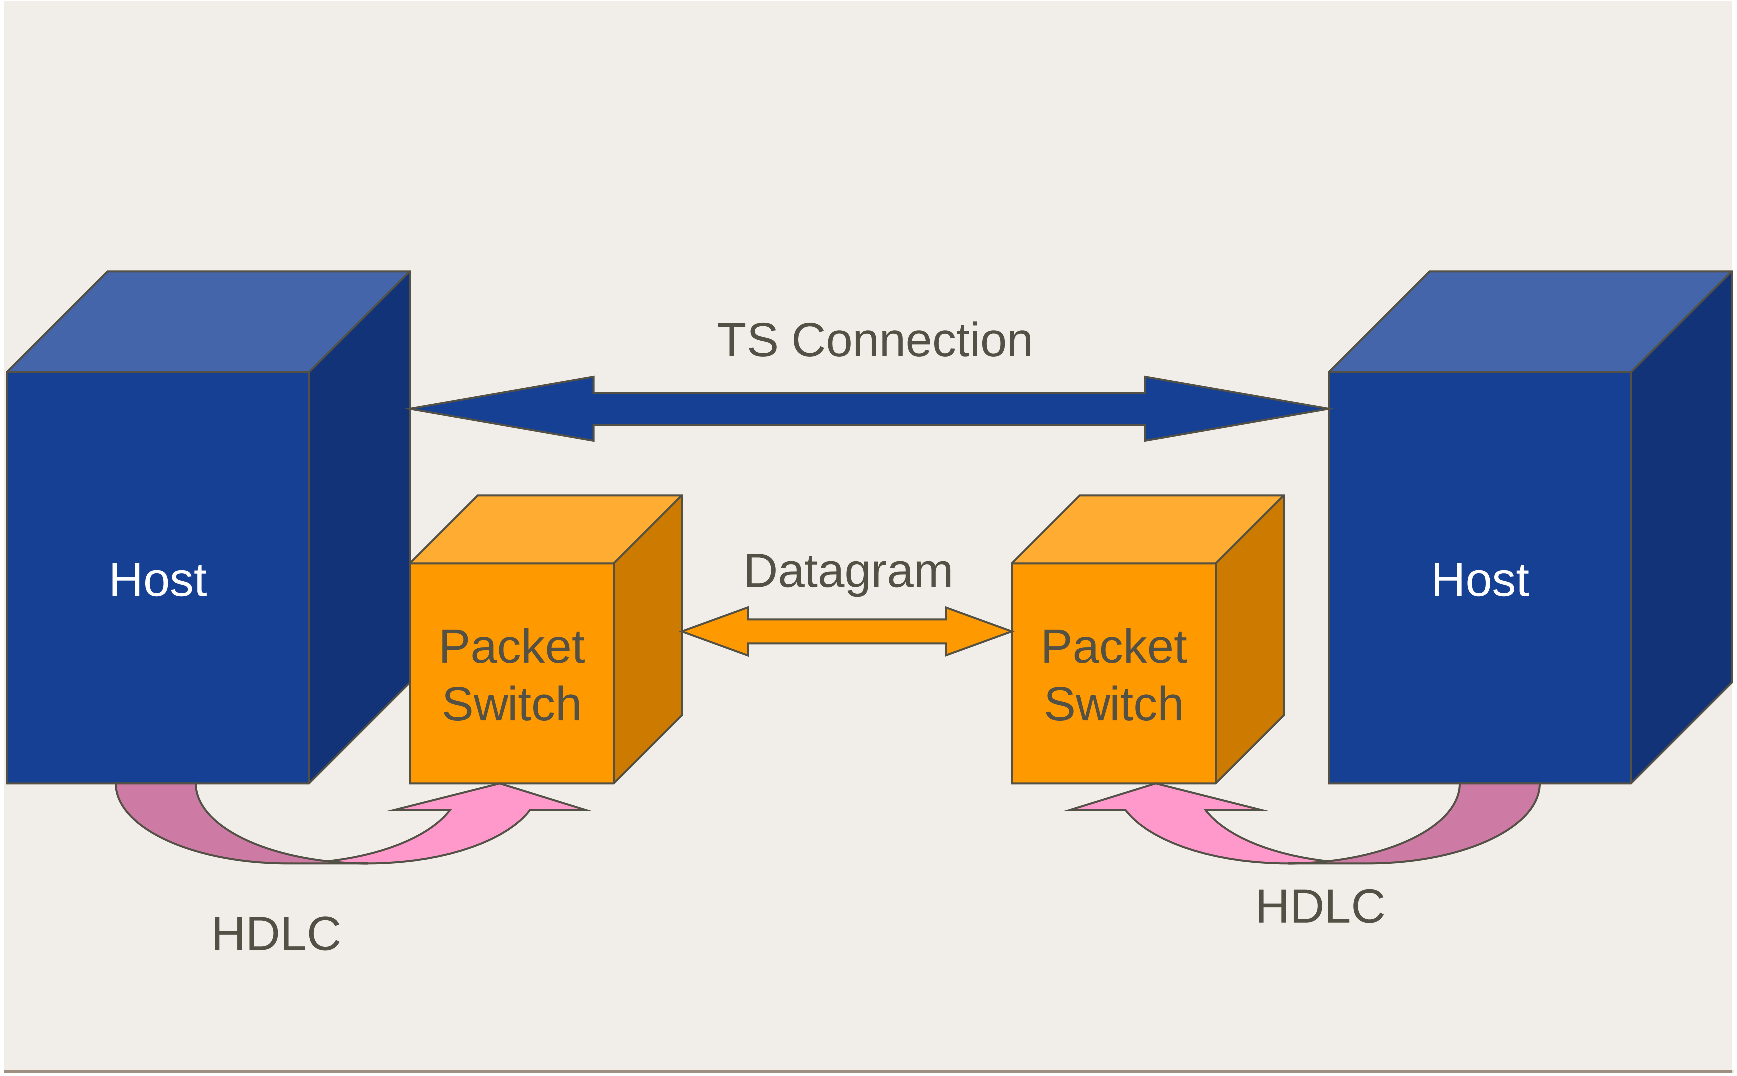

CYCLADES Motivation

- Testbed for research on **both** protocols and applications
- Particularly interested in “Multiple Gates”
  - Divide traffic from source to destination along multiple paths
  - Needed to move control of stream from network to transport
  - Peer-to-peer does this today
- Test protocols in the end point, and optimize in the core
  - Moves network control to the end point
  - Architecture driven by practical goals, not social philosophy
  - Not the most efficient design
- Is this a practical way to build production networks?
  - *CYCLADES structure provides a good testbed for trying out various protocols, as was its intent; but it requires a more cooperative and coordinated set of hosts than is likely to exist in a public environment* – Larry Roberts

Internet

- High-level architecture same as CYLCADES
  - Datagram network layer
  - End-to-End transport layer
  - Pipelined ("Sliding Window") Packet Stream
- Cerf acknowledged debt to Pouzin and LeLann
  - Particularly with respect to "sliding window"
  - Sliding Window implementation implies common structure.
- 1974 spec combined TCP and IP.
  - Later, these were split into three pieces
- Parallel work at DEC, Xerox, and ISO
- Modular structure was very appealing to computer people.
- "*…requires a more cooperative and coordinated set of hosts than is likely to exist in a public environment.*"

Internet Architecture

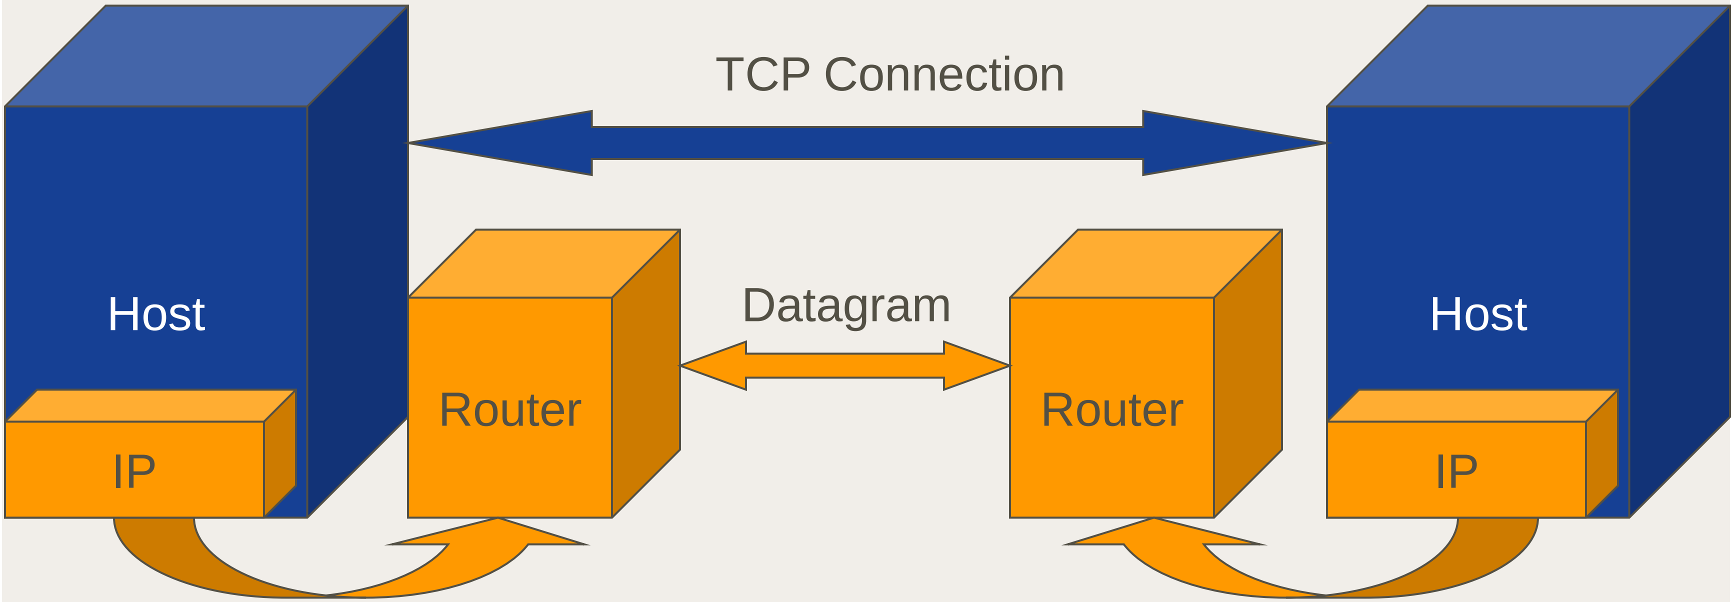

End-to-End Arguments

- End to End Arguments in System Design
  - Devised in 1981, nine years after CYCLADES.
  - An attempt to reverse engineer the Internet's architecture
  - Circular reasoning:
    The function in question can completely and correctly be implemented only with the knowledge and help of the application standing at the end points of the communication system. Therefore, providing that questioned function as a feature of the communication system itself is not possible.
- Paper states the obvious.
- Nonetheless, the Internet has an important End-to-End component.
- A series of RFCs has tried to capture its usefulness.

Design for Change

- RFC 1958 by Brian Carpenter in 1996:
  - Emphasizes reliability over “correctness”.
  - Adds “Design for Change” back to the discussion:

    Over the 25 years since the ARPANET started, various measures of the size of the Internet have increased by factors between 1,000 (backbone speed) and 1,000,000 (number of hosts). In this environment, some architectural principles inevitably change. Principles that seemed inviolable a few years ago are deprecated today. Principles that seem sacred today will be deprecated tomorrow. ***The principle of constant change is perhaps the only principle of the Internet that should survive indefinitely***

A Critical Review

- Tim Moors, "A Critical Review of End-to-end Arguments in System Design," 2002.
- Shows that E2E Args doesn't correctly describe the file transfer process: "*According to the end-to-end arguments, applications, not transport layers, should check integrity*."
- Shows that TCP does error checking for FTP because it's a trusted part of the system: ", *the decision to implement reliable transfer in the transport layer is not justified on the basis of end-to-end arguments, but rather on the basis of trust*.
- E2E had never been as important as we thought it was.

Layering Considered Harmful

- RFC 3439 by Bush and Meyer in 2003 scales E2E down:
  - *End-to-End Argument does not imply that the core of the Internet will not contain and maintain state. In fact, a huge amount coarse grained state is maintained in the Internet's core (e.g., routing state). However, the important point here is that this (coarse grained) state is almost orthogonal to the state maintained by the end-points (e.g., hosts). It is this minimization of interaction that contributes to simplicity.*
- Introduces the notion of horizontal layering:
  - *The first order conclusion then, is that horizontal (as opposed to vertical) separation may be more cost-effective and reliable in the long term.*
- Horizontal Layering alters the meaning of the end point.

Protecting Innovation

- *RFC 3724 by Kempf and Austen in 2004:*
  - *Does the end-to-end principle have a future in the Internet architecture or not? If it does have a future, how should it be applied? Clearly, an unproductive approach to answering this question is to insist upon the end-to-end principle as a fundamentalist principle that allows no compromise.*
  - *One desirable consequence of the end-to-end principle is protection of innovation. Requiring modification in the network in order to deploy new services is still typically more difficult than modifying end nodes.*
- "Protecting innovation" has always been the main concern. How does end-to-end do that?

Clark's Tussle

- Tussle in Cyberspace, 2005, David Clark et. al. supports new E2E:
  - *If the core of the network must be modified to deploy a new application, this puts a very high hurdle in front of any unproven idea, and almost by definition, a new idea is unproven.*
- What about the applications the core doesn't support very well?
  - *"Protocol design, by creating opportunities for competition, can impose a direction on evolution."*
- What about reliability?
  - *"The more simple the core of the network, the more reliable it is likely to be."*
- Is this right?
  - Systems actually achieve reliability by redundancy: replicated routes allow end points to recover from router failures.

Tussle's Implications

- Innovation is enabled by new protocols & services:
  - Markey Bill is at odds with this discovery.
- Communications applications can't thrive in the current system:
  - The Internet that we have is best for content.
- We've lowered our expectations of the Internet so that its failure to support quick response and high bandwidth at the same time is not a problem.
  - We need to expect more of the Internet.
  - We can't solve all problems with more bandwidth.
- Managing congestion is the central issue.

The Congestion Problem

- Managing congestion is the central design problem of any multi-user communication network.
  - Packet Switched networks like the Internet and the Public Switched Telephone Network (PSTN) take very different approaches.
- PSTN allocates bandwidth in fixed increments for long periods of time to avoid congestion:
  - Unused bandwidth goes to waste.
- Packet Nets manage millisecond-by-millisecond:
  - “Congestion” is simply high utilization.
  - High utilization is the goal.
- On the PSTN, congestion is a pathology, but on a packet net it’s a goal.

Approaches to Congestion Control

- CYCLADES and early Internet sent Quench packet to user whose packet was dropped.
  - Inconsistent implementation doomed this approach for Internet.
  - Increases network traffic.
  - Dropped packet became Congestion Signal in Jacobson patch.
- Better approach signals congestion by setting a flag in the packet.
  - Explicit Congestion Notification (ECN) and Re-ECN do this.
  - ECN is an Early Warning Signal that congestion is building.
  - ECN leads to more efficiency and higher utilization of links.
- What about applications that don't want to back off?
  - This is a problem in any scheme.
  - Applications aren't created equal.
  - Pay for priority is not a novel concept.

Congestion Scenario

A congestion resolution algorithm has to be implemented for traffic going from Router 1 to Router 2.

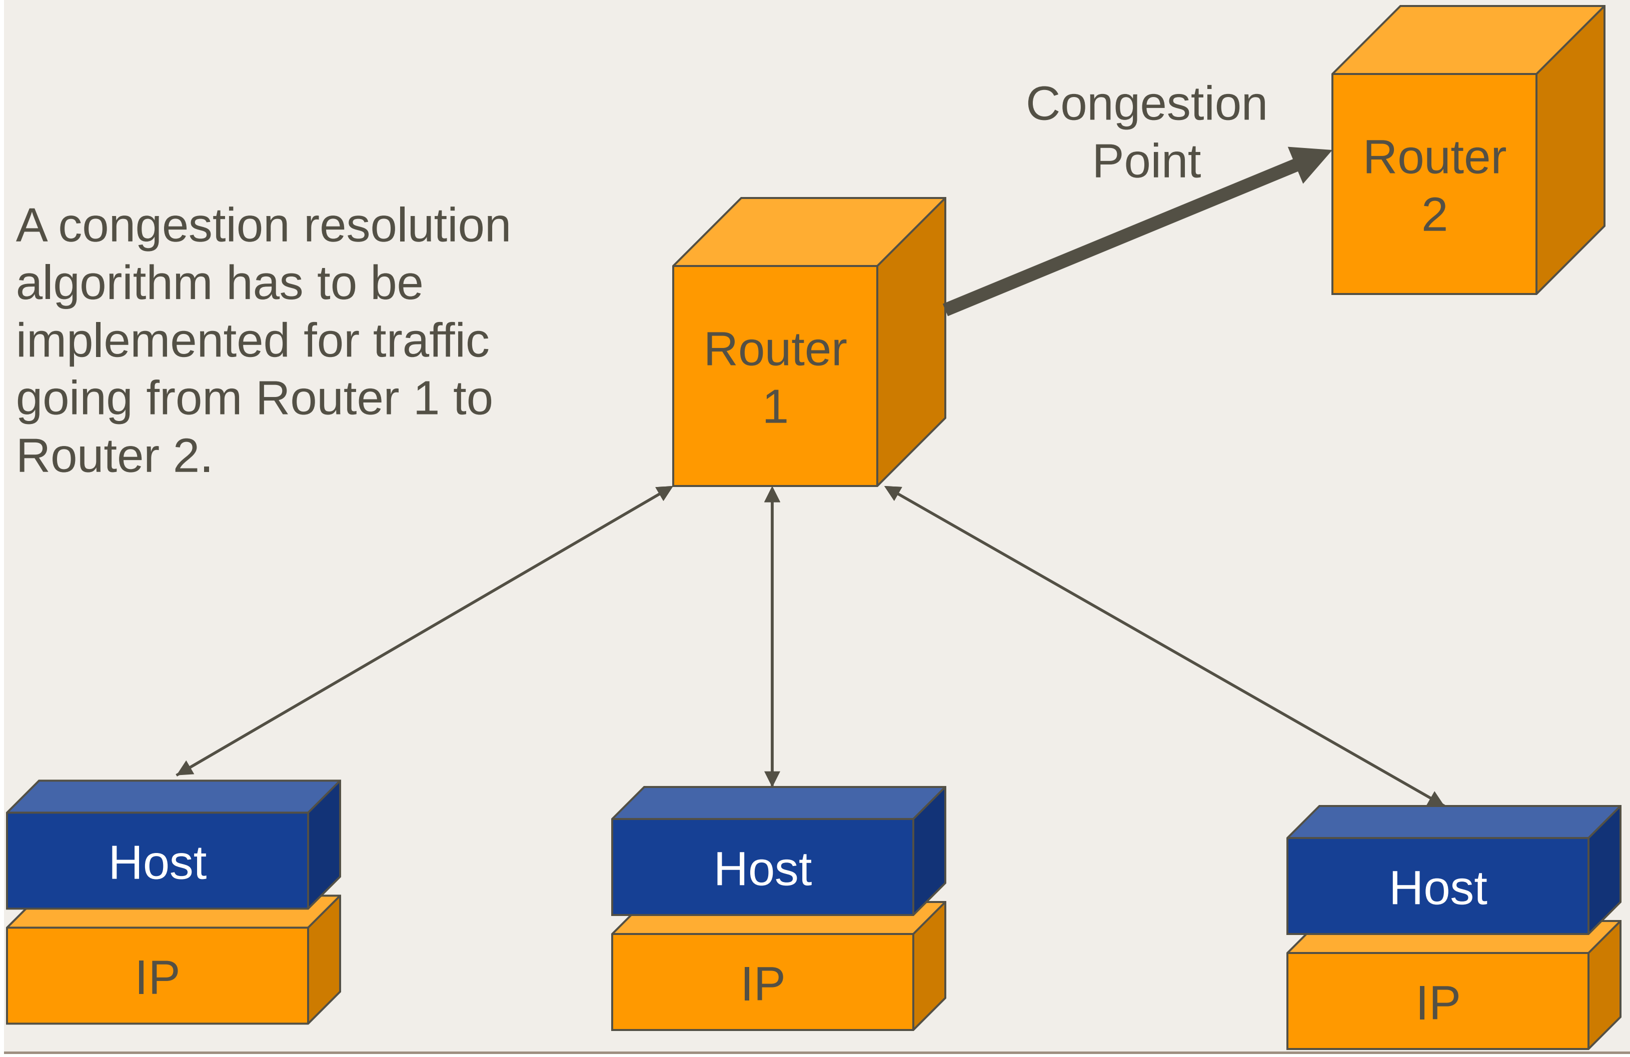

Congestion Algorithms

- Queuing Disciplines
  - First-in, First-out
  - Token Bucket
  - Numerous variations
- Adding Priorities to Queuing Disciplines
  - Class-based approaches
  - Reservation-based approaches
- Economics turn out to be vital
  - Absent a price, all packets are highest-priority.
  - Packets in excess of quota can be demoted to a lower class.

Economics of Congestion

- As a link approaches Congestion, applications need to make their preferences known:
  - How valuable is low latency to you?
  - Communication-oriented apps answer differently than Content-oriented ones.
  - The user's decision has economic consequences.
- Real-time bidding for latency solves the problem.
  - Implemented by a "Bandwidth Broker."
  - A fair auction doesn't discriminate.
  - Assign costs to users creating the congestion.
- Why wasn't this done in the past?
  - The difference between high-load and light-load apps wasn't as great as it is today.
  - Flat-rate pricing is discrimination by design.

Explaining the Internet to the Policy Community

- Berkman Center scholars interpret the technology for us:
  - Barlow: *"You have no moral right to rule us nor do you possess any methods of enforcement we have true reason to fear"*
  - Isenberg: *"This would be a true Stupid Network – one treatment for all kinds of traffic."*
  - Lessig: *"This minimalism in design is intentional. It reflects both a political decision about disabling control and a technological decision about the optimal network design. The designers were not interested in advancing social control; they were concerned with network efficiency."*
- Each of these statements is factually incorrect and a-historical
  - They form the basis of the net neutrality arguments.

More Sophisticated Net Neutrality Arguments

- Tim Wu: Internet discriminates by design:

  “Proponents of open access have generally overlooked the fact that, to the extent an open access rule inhibits vertical relationships, it can help maintain the Internet’s greatest deviation from network neutrality. That deviation is favoritism of data applications, as a class, over latency-sensitive applications involving voice or video.”

- Functional Layering theorists:
  - Whitt, Solum & Chung, Benkler, Werbach, Frieden
  - What about “Layering Considered Harmful?”
  - OSI Reference Model was a conjecture
  - A question, not an answer

Why the Internet Succeeded

- In the right place at the right time
  - Personal Computer explosion created a need
  - Deployment advantage
- Massive funding by US Government and research community
  - Open protocol implementations
  - NSF Backbone
  - Overseas connections in Europe and Asia
- Good enough to get the job done
  - Corporations wanted e-mail instead of fax
  - Personal Computer users wanted the Web

Internet Innovation Cycle

1. Internet is in equilibrium.
2. New application alters the mix:
   - Web
   - P2P File Transfer
3. Internet operators adjust to application:
   - Fatter Pipes
   - Traffic Engineering
4. Application adjusts to the network:
   - Change the application to match the network
5. Equilibrium is regained:
   - Internet favors content over communication
   - People favor communication over content

Challenges to the Internet

- Attacks
  - DDoS is a business model.
  - Hundreds of these attacks happen every day.
- Traffic Growth
  - Private peering is growing faster than public peering; Odlyzko doesn’t measure this
- Proliferating Prefixes
  - BGP routes will explode soon.
  - Result of weak Internet Routing Architecture
  - Mobility and multi-homing aggravate
  - IPv6 PI addressing aggravates
  - Moore’s Law doesn’t operate on low-volume parts.

IAB Routing and Addressing Workshop

- Meyer, Zhang, and Fall, “Report from the IAB Workshop on Routing and Addressing,” RFC 4984, 2006:
  - *“The clear, highest-priority takeaway from the workshop is the need to devise a scalable routing and addressing system, one that is scalable in the face of multihoming, and that facilitates a wide spectrum of traffic engineering (TE) requirements.”*
- This system has not been developed
  - Probably can’t be developed around IPv4 or IPv6.
- At last, people are talking about the need for a New Architecture
  - NSF funding NEWARCH and FIND
  - FCC Big Ideas Workshop
- Many testbeds, no new architecture
  - We’re not there yet

Regulating an Architecture in Crisis

- Net Neutrality regulators seek to preserve "the Internet as we know it."
  - But they don't know it very well…
  - …and preserving it is not an option.
- Technology Silos and Functional Layers both have it wrong
- Preserve the freedom to develop a New Architecture.
- Social value is important:
  - We protect that by regulating commercial interactions with users, not by micro-managing routing policies.
- The Internet shows that social utility comes from engineering freedom:
  - The experiment is just beginning

Conclusion

- The real meaning of end-to-end is dynamism:
  - Innovation in fundamental network services supports innovation in applications.
- The Internet was designed for testing new protocols and services.
- Nobody can predict the future, but anyone can react constructively.
- Making systems work better should not be a crime:
  - Innovators embrace change.

Thank you!

rbennett@itif.org
www.itif.org