University of Glasgow

Internet Protocols Laboratory

Should the Internet Adopt Named Data and Services?

Colin Perkins

Outline

- What are some trends pushing the network to evolve?
- What should an internet provide?
- How should the Internet evolve?

Trends

The Changing Internet

- Politics
  - Government pressure to compromise privacy and security
  - Business and control possibilities of splintering the network

The Changing Internet

- Changes in application demand
  - Increased demand for content and services, APIs
  - Video and real-time multimedia
  - Capacity challenges – video, software updates

The Changing Internet

- Changes in technology
  - Software defined networking
  - Increasing use of virtualisation
  - Increasingly programmable infrastructure

The Changing Internet

- Increased demand for content and services, APIs
- Video and real-time multimedia
- Capacity challenges – video, software updates

- Software defined networking
- Increasing use of virtualisation
- Increasingly programmable infrastructure

- Network architecture doesn't match the way it's used – we use a host-centric network to support content-centric applications and services
- Technology changes give more programmability and flexibility than the protocols can support – **how to evolve to a network that can use such features?**

What should an internet provide?

The Internet Way of Networking

- The Internet Society tried to define the critical properties of the Internet*
  - An Accessible Infrastructure with a Common Protocol
  - An Open Architecture of Interoperable and Reusable Building Blocks
  - Decentralised Management and a Single Distributed Routing System
  - Common Global Identifiers
  - A Technology Neutral, General-Purpose Network

* https://www.internetsociety.org/resources/doc/2020/internet-impact-assessment-toolkit/critical-properties-of-the-internet/

The Internet Way of Networking

- I would add:
  - The network should be secure
  - It should be private by design, and avoid tracking users

Evaluating the Internet

- Does the Internet meet these goals?

- Yes, largely – although there are challenges
  - The Internet governance model is broadly open and accessible
  - There is an interoperable set of protocols and commonly used building blocks
    - IP, TCP, QUIC, TLS, BGP, HTTP, DNS, …
    - Although routing, addressing, and naming systems have never been entirely global
    - Security and privacy remain a work in progress
  - The system embodies a particular worldview, that is not universally accepted

Evolving the Internet

- Could another design meet the same criteria, but better support emerging applications and technologies?

Named Data, Named Services

- We **care** about providing content and services
  - Ideally with low latency to support real-time

- We **don't care** where that content or those services are located
  - Provided they can be accessed securely
  - Provided we can be confident the data is not tampered with

- The Internet is architected around named devices, but it could focus on named data and named services

Named Data, Named Services

- Does this mean we change the Internet addressing and routing model?
- Maybe – or perhaps we change the internetworking point and content routing approach to simplify the network?

Named Data, Named Services

- What we want might be an evolution of HTTP, rather than a new IP
  - Host name → authority/trusted source for data or services
    - Don’t name the host where the data is to be found, name the data or service and the source of trust that data
    - Don’t secure the connection, secure the data/service
  - Send IP packet → express interest in named data/service
    - Route request towards named data or service, rather than a named host
    - A name more clearly exposes intent than does an IP address
    - Pervasive caching pulls data and services towards consumers

Why Evolve HTTP Rather Than IP?

- Because changing the lower layers is not feasible

Internetworking Challenges

- What does internetworking look like in future?
  home → edge → hypergiant

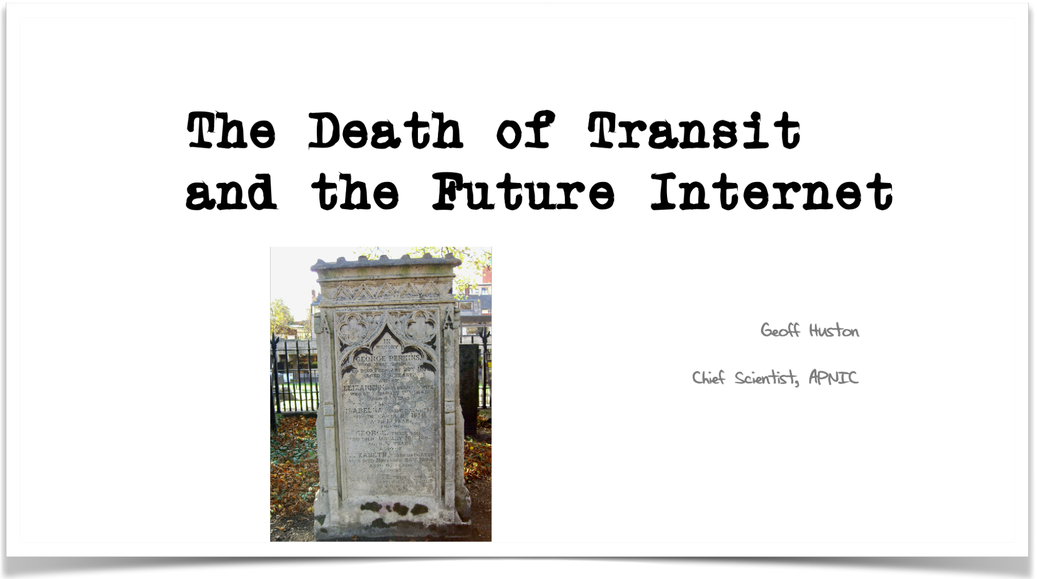

Internetworking Challenges

- What does internetworking look like in future?
  home → edge → hypergiant

- The home runs IP (too ossified to change)
  - IP is a useful link layer: 192.168.0.0/16 with headers rewritten at every domain boundary

Internetworking Challenges

- What does internetworking look like in future?
  home → edge → hypergiant

- The home runs IP (too ossified to change)
  - IP is a useful link layer: 192.168.0.0/16 with headers rewritten at every domain boundary
  - But IPv6 will solve this!

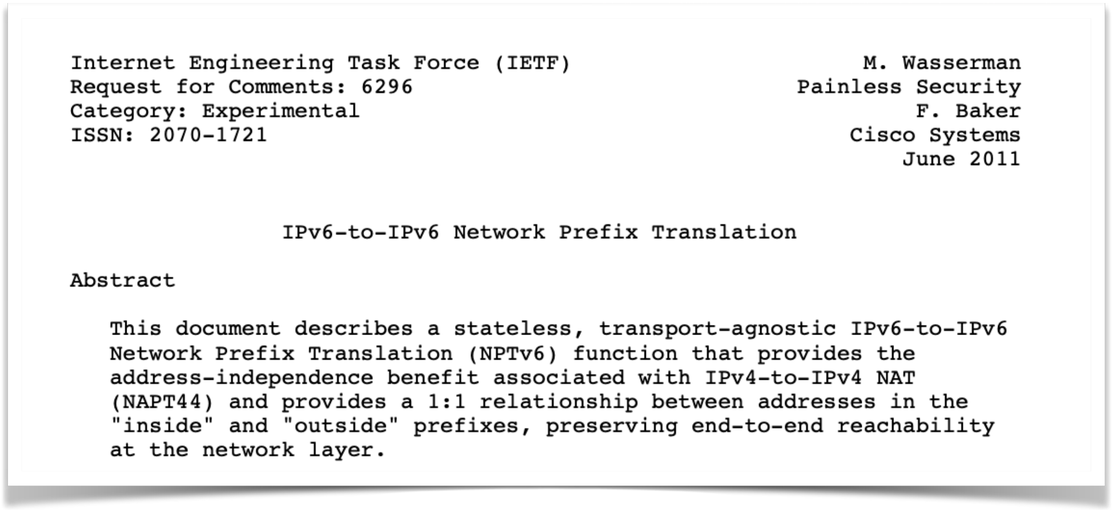

Internet Engineering Task Force (IETF) — M. Wasserman
Request for Comments: 6296 — Painless Security
Category: Experimental — F. Baker
ISSN: 2070-1721 — Cisco Systems
June 2011

IPv6-to-IPv6 Network Prefix Translation

Abstract

This document describes a stateless, transport-agnostic IPv6-to-IPv6 Network Prefix Translation (NPTv6) function that provides the address-independence benefit associated with IPv4-to-IPv4 NAT (NAPT44) and provides a 1:1 relationship between addresses in the "inside" and "outside" prefixes, preserving end-to-end reachability at the network layer.

Internetworking Challenges

- What does internetworking look like in future?
  home → edge → hypergiant

- The home runs IP (too ossified to change)
- The edge runs IP to talk to the home, and to the long tail

Internetworking Challenges

- What does internetworking look like in future?
  home → edge → hypergiant

- The home runs IP (too ossified to change)
- The edge runs IP to talk to the home, and
  to the long tail
- The hypergiants do whatever they choose internally, but
  expose it via the latest HTTP variant over IP – because
  HTTP is the only protocol that can change

Internetworking Futures?

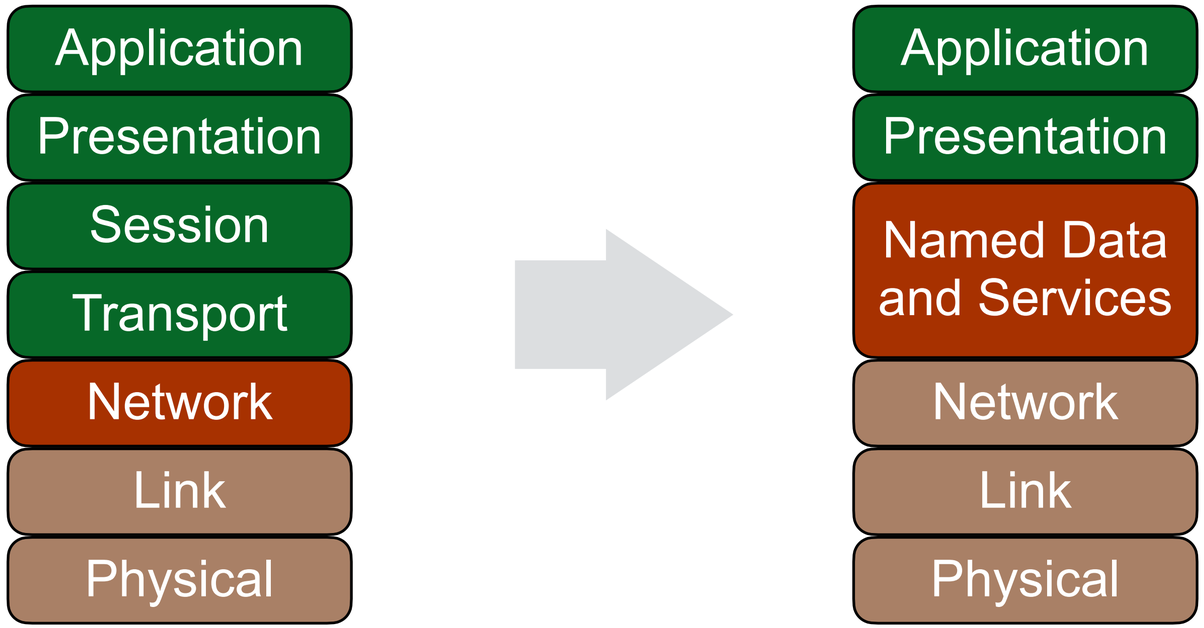

Is this still the Internet way?

- Critical features
  - An Accessible Infrastructure with a Common Protocol – yes!
  - An Open Architecture of Interoperable and Reusable Building Blocks – yes!
  - Decentralised Management and a Single Distributed Routing System – yes, although the details need to be defined
  - Common Global Identifiers – yes, although the details need to be defined
  - A Technology Neutral, General-Purpose Network – yes!
  - The network should be secure – yes!
  - It should be private by design, and avoid tracking users – to be defined

University of Glasgow

Challenges

Challenges

- Applications request named data and invoke named services via HTTP today
  - The IP layer isn't important to applications or service providers
  - The de-emphasis has already happened
- Research challenges:
  - How to transition to, and scale, name-based routing rather than name-to-address mapping at the content distribution layer?
  - How to use the increasingly programmable infrastructure to support this?
  - What structure is required in names?
  - What is the trust model for named data and services?
  - How to avoid user tracking based on named data requests?

What Don’t we Need?

- Limited domains
  - The concept of internetworking is correct – we need common protocols
  - We just need to refine the Internetworking point, accept that it is moving up the stack
- Congestion control, quality of service, enhanced packet switching
  - If you have capacity, these are irrelevant and the infrastructure only serves to help track users; if you don’t have capacity, your users will get fed-up and leave anyway
  - Application-agnostic mechanisms like L4S are important; service scheduling matters
- Blockchain
  - Just no.

Conclusion

- The Internet already uses named data and services – **it just doesn't admit it**
- The capacity and flexibility of the increasingly programmable infrastructure **could** permit a transition to name-based routing supporting pervasive caching of data and services – saving complexity by aligning the network with its uses

- Interesting challenge to secure the expected simplifications while maintaining privacy