Address Family Transition Router - AFTR

April 22, 2010
UKNOF 16
London, UK

What is ISC?

- Internet Systems Consortium, Inc.
  - 501(c)(3) Non-profit, public benefit corporation
  - Headquartered in Redwood City, California
  - Paul Vixie, Founder & President

ISC's Mission

- Develop, maintain & support production quality open source software, such as BIND, DHCP & AFTR
- Enhance the stability of the global DNS through reliable F-root nameserver operations, support & training for our software
- Further protocol development efforts in the areas of DNS evolution, specifically DNSSEC and facilitating the transition to IPv6

We've all heard the news...

- Unallocated IPv4 is running out
  - Gone in 2012? Sooner?
  - Functionally already gone for large network operators
- Applications aren't IPv6-enabled
- Content providers are mixed at best
- ISPs/access providers have to bridge the gap

IPv4-IPv6 co-existence

- Gradual transition needed
  - Carriers need IPv6 to grow the net
  - Users need to be able to keep using IPv4 until they're ready to move on
- Extensive IETF and vendor activity on medium-term IPv4-IPv6 co-existence
  - NAT444, NAT464, 6to4, 6rd,…
  - BEHAVE, SOFTWIRE, DHCPv6,….

IPv4-IPv6 co-existence

AFTR is an implementation of dual-stack lite:

- https://datatracker.ietf.org/doc/draft-ietf-softwire-dual-stack-lite/

Dual Stack Lite Protocol

- Combines two key technologies:
  - IPv4 in IPv6 encapsulation
  - IPv4 NAT
- Intended for broadband environments
- One AFTR ("Address Family Transition Router") can handle many clients

How does DS-lite work?

- Tunnel between
  - CPE (“B4 element”) and
  - carrier-based NAT (“AFTR element”)
- B4 (“Basic Bridging BroadBand”) element needs to:
  - Find the other end of the tunnel
  - Encapsulate IPv4 in IPv6
  - No NAT!

How does DS-lite work? (2)

- AFTR ("Address Family Transition Router") needs to:
  - Establish tunnel with many B4 elements
  - Accept and decapsulate IPv6 tunnel traffic
  - NAT between customer IPv4 and its own routable IPv4 address(es)
  - Accept and encapsulate IPv4 traffic back towards B4 elements
- Manage operations, security and audit functions for multiple tunnels

AFTR 1.0 Overview

- Proof of concept implementation of Dual Stack-lite protocol
- AFTR in user mode, commodity hardware/OS
- B4 on commodity home gateway
  - Support for A+P, no-NAT pass-through

DS-lite Topology

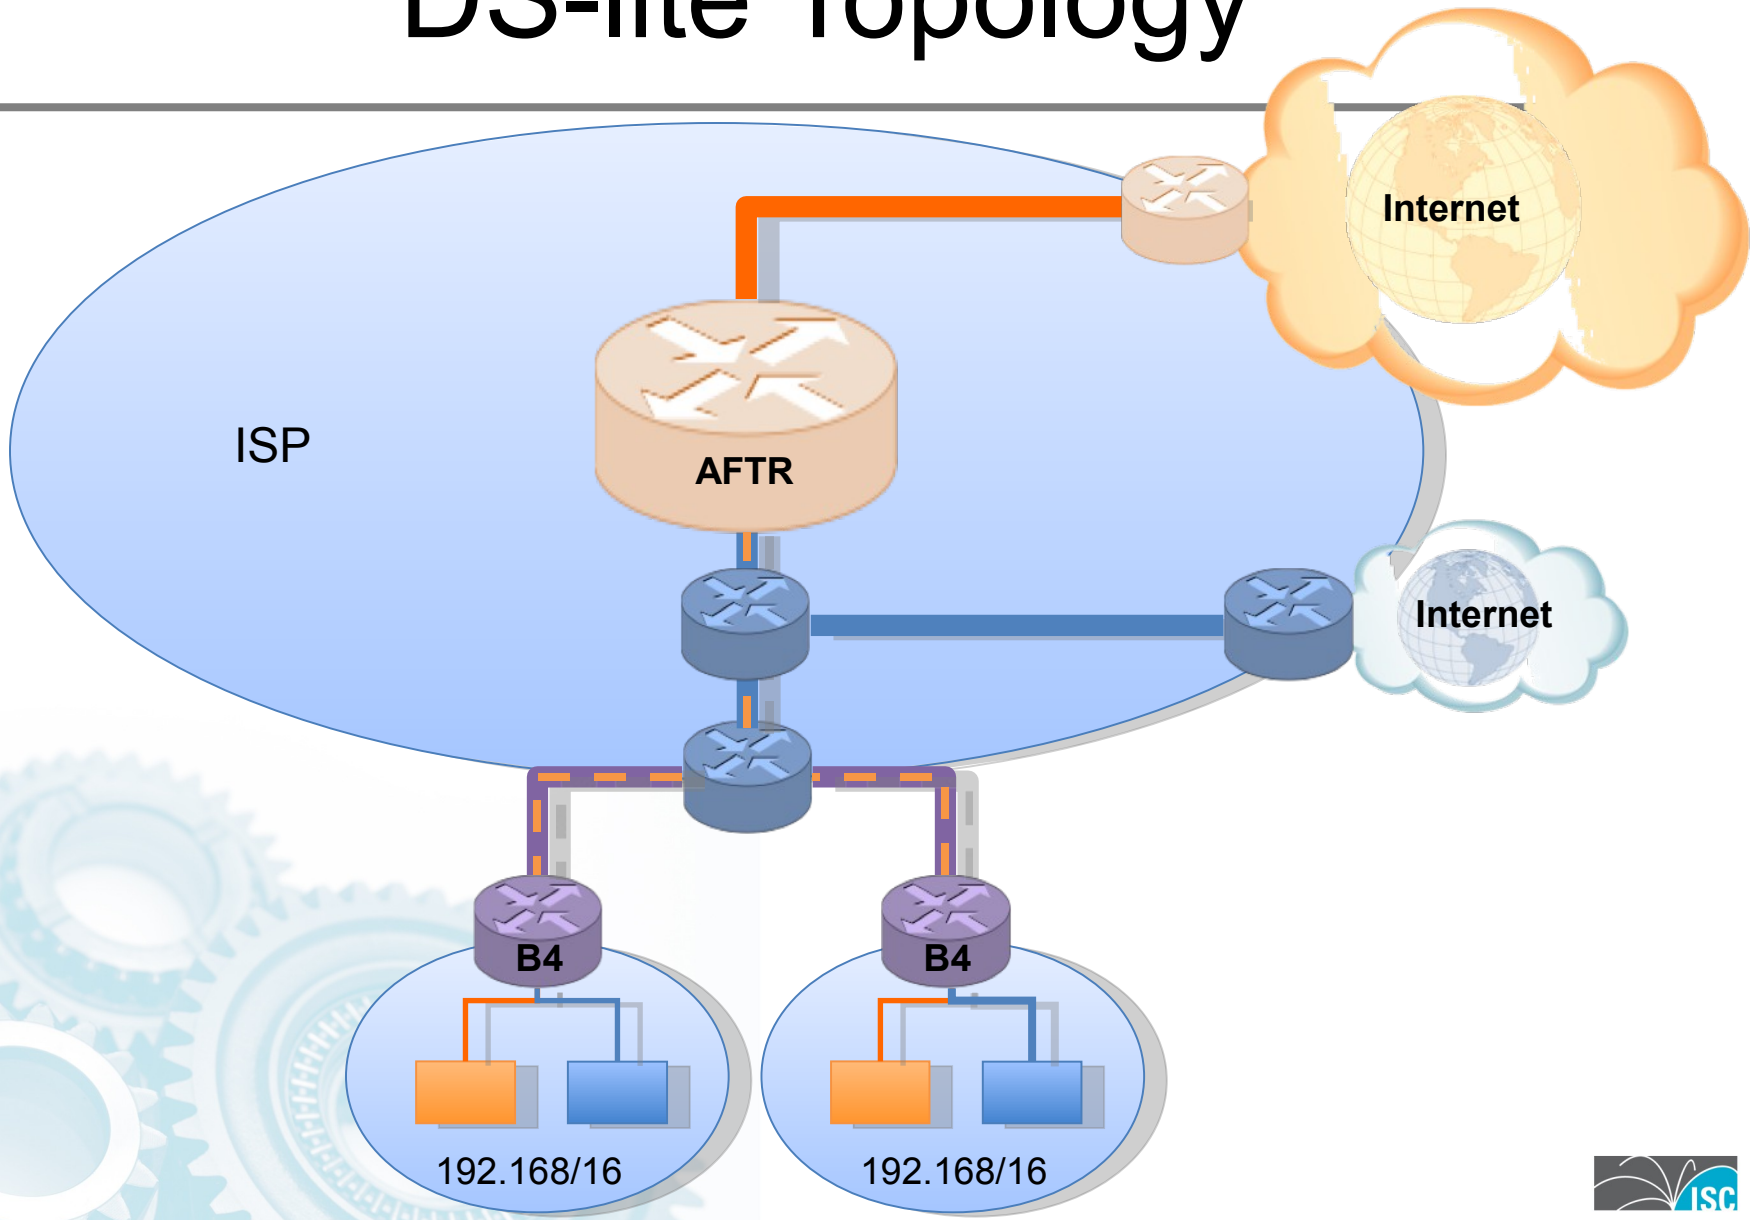

AFTR 1.0 Limitations

- No UPnP
- No customer-initiated port forwarding
- Configuration

AFTR 1.1 Roadmap

Features anticipated for AFTR 1.1 (to be released late in Q2 2010):

- UPnP
- Web portal, XML back-end
- Admin log tools
- Improved documentation

AFTR 1.2 Planning

- Comcast leading Field Trial late 2010
- 1.2 Release to follow
- Focus on optimization and performance improvement based on issues found in field trial

AFTR in the Community

AFTR Trac Site

Trac site with:

- Subversion code repository for community viewing
- Bug submission, with visible, searchable bug repository
- AFTR community Wiki

http://aftr.isc.org

AFTR in the Community

AFTR Mailing Lists

Open community email lists:

- aftr-announce@lists.isc.org
- aftr-users@lists.isc.org

AFTR Development Model

Managed Open Source

- ISC dedicated engineers
- ISC Product Manager
- Community participation
  - Product planning
  - Testing & Feedback

AFTR Funding Model

- Comcast initial project sponsor
  - Funding
  - Project direction
- AFTR now maintained by ISC
  - Developers
  - Resources
- Community Participation
  - Code
  - Resources
  - Funding
- *AFTR Forum*

ISC AFTR Forum

- Primary funding vehicle
  - Donations
  - Membership
  - Directed development
- Membership Levels with different fees and benefits

https://www.isc.org/software/guild/aftr

More Information

AFTR the Fire: IPv4-IPv6 Co-Existence Technology

Google Techtalk

January 21, 2010

Paul Selkirk, ISC

https://www.isc.org/community/presentations/video#21jan-techtalk

Questions?

The B4 Element: Behavior

- The B4 element is the end-user side
- Needs to run in a low-cost, limited-resource, commodity environment
  - i.e. home gateway
- Simple functionality:
  - Set up the IPv4-in-IPv6 tunnel
  - Encapsulate/decapsulate IPv4 traffic

The B4 Element: Implementation

- OpenWRT module as initial platform
  - Linux distro for home gateways
- Provisioned with IPv6 address only
- Requests DHCPv6 option for AFTR tunnel endpoint
  - IPv6 address or DNS name
  - draft-ietf-softwire-ds-lite-tunnel-option-01.txt
  - Encapsulation/decapsulation is the easy part (except for fragmentation, which we'll get to later)

The AFTR element: basic functionality

- Decapsulate IPv4 packet from B4
  - NAT IPv4 packet to public IPv4 address
- Encapsulate IPv4 from Internet to B4
  - Reverse NAT translation
- Port reservation/management
- Configuration and Logging
- Resource management
  - Randomized port buckets per B4
  - Tunnel characteristics

The AFTR element: Implementation

- Commodity hardware, operating systems
  - Linux netbook
- Simple architecture
- Runs in user space