



The Coming Internet Crisis in Routing and Addressing: An Overview

ESCC/Internet2 Joint Techs Workshop

January 24, 2008

R. Kevin Oberman
Senior Engineer

Energy Sciences Network
Lawrence Berkeley National Laboratory

oberman@es.net, www.es.net

Networking for the Future of Science



Presentations

- Kevin Oberman – Overview of area
- Michael Sinatra – Campus transition experiences
- Brent Sweeny – BGP Hinting
- David Meyer - Routeviews
- Caren Litvanyi – Commercial IPv6 on CPS
- David Meyer – LISP (not the language)
- Xing Li – IPv6 rollout in China
- Chris Heermann – Breakable Experimental Network

Kevin Oberman - Overview

- Exhaustion of IPv4 space and its impact
- Hardware limitations (RIN, FIB, and CPU)
- IPv6 transition (lack of) progress
- IPv6 transition problems
 - Possible problems made worse by successful transition
 - Call for ideas to get there

Michael Sinatra - Campus Transition

- Many problems
 - Poor dual stack implementations
 - DNS Issues, especially simple glue records
 - Educause did not handle IPv6 addresses correctly
 - Most registrars don't register IPv6 glue
 - DHCPv6 Servers???
 - No server in Windows
 - No server in Mac OSX
 - Almost ready in most Unixen
 - DHCPv6 Clients???
 - Not in Mac OSX (and not coming)
 - Vista has some issues, but mostly works
 - Most Linux still MIA
 - Ubuntu works if the dibbler-client is added
 - BSDs work
- Even trial jabber server had major problems

Brent Sweeny – BGP Path Hinting

- An update of Brent's proposal to use BGP communities to “hint” to providers about which path is preferred
- Proposal more complete
- More work needed and volunteers to try things are needed

David Meyer – Routeviews

- History and update on routeviews
- Examples of accessing the data and ways it has been used
- Overview of available facilities and archived data
- New capabilities and tools
 - BGPmon looks especially interesting
- Ongoing issues and problems

Caren Litvanyi – IPv6 on CPS

This was a discussion of the implementation of IPv6 commercial peering.

Very limited (because not many commercial nets do IPv6)

Probably not of interest to ESCC

David Meyer – LISP: Keeping the FIB Under Control

- LISP is an encapsulation (jack-up under the network layer)
- Provides locator and end-system ID
- Allows routing based on locators through DFZ
- End-system IDs used only by enterprise and local provider (hopefully)
- Read the papers. This is a complex proposal not easily summarized.
- David is looking for volunteers to try the very green code and see what they think.

Xing Li – IPv6 in China

- The size of this effort is simply on a different scale from what we are looking at
- China has far more IPv6 implementation than most of the world
- Much internal networking is over IPv6
- Done largely by building new nets as opposed to transitioning existing networks

Chris Heermann – Breakable Experimental Network

- BEN is a fiber network that may be used for testing and development and may be broken as needed
- Based out of Research Triangle and connects Duke, UNC and NCSU
- Testing and experimentation can be done at all network layers (1-7)

Where do we go from here?

- We need a viable transition to IPv6 now
 - See <http://www.civil-tongue.net/clusterf/>
 - Contribute ideas
 - Prepare to feel Randy's wrath :-)
- Look for ways to limit FIB growth (e.g. LISP)
- Look around for unused address space
 - Maybe you can sell it to get the budget for Sup750-3BXL upgrades
- **Don't panic!** *(Note the large, friendly letters)*
 - The answer is 42

Recommended Reading

- Talks from NANOG41
 - <http://www.nanog.org/mtg-0710/bush.html>
 - <http://www.nanog.org/mtg-0710/bicknell.html>
 - <http://www.nanog.org/mtg-0710/farinacci.html>
 - <http://www.nanog.org/mtg-0710/meyers.html>
 - Both slides and RealMedia recordings available
- RAM mailing list
 - <http://www.ietf1.org/mailman/listinfo/ram/>