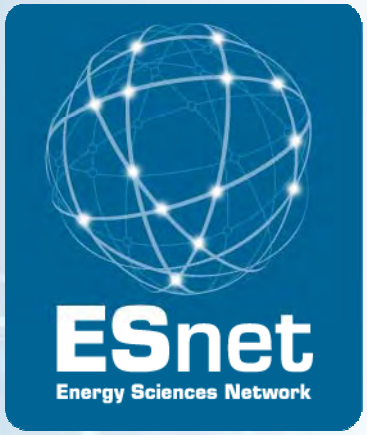

IPv6 Implementation at a Network Service Provider

2010 Inter Agency IPv6 Information Exchange
August 4, 2010

R. Kevin Oberman
Sr. Network Engineer

Who Are We?

ESnet is the network provider for the Department of Energy's Office of Science

- ESnet is a networking pioneer with nearly a quarter century of networking
  - Began as MFEnet in 1976
  - Became ESnet with broader mission in 1986
  - Started support of BGP4 and modern peering in 1994
  - Multicast support since 1995
- Provide network connectivity to DOE Science funded laboratories and research projects
- Provides full commercial connectivity with over 100 commercial peers
- ESnet is transit free

Pioneered IPv6

ESnet has pioneered IPv6 since its inception

- ESnet started working on IPv6 in 1996
  - Tony Hain and Bob Fink chaired the main IPng IETF working groups
  - ESnet worked closely with Sun, Digital, Kame, and Cisco in the development and testing of IPv6 developmental code
  - Instrumental in the development of the 6-Bone
  - Partnered with Viagenie to create 6Tap, the first IPv6 Internet Exchange
  - Received the first production IPv6 address allocation from ARIN
    - First production addressed system, hershey.es.net still sits in my office in Berkeley.

IPv6 Status

IPv6 is a fully supported production service of ESnet

- Since 2004
- Available to all sites and peerings
- ESnet web services, NTP, DNS, and mail use IPv6
- Currently we are our own best customers
  - That is changing
  - Sites are adding IPv6 connectivity
  - Even a couple of IPv6 services

ESnet4 Backbone Topology

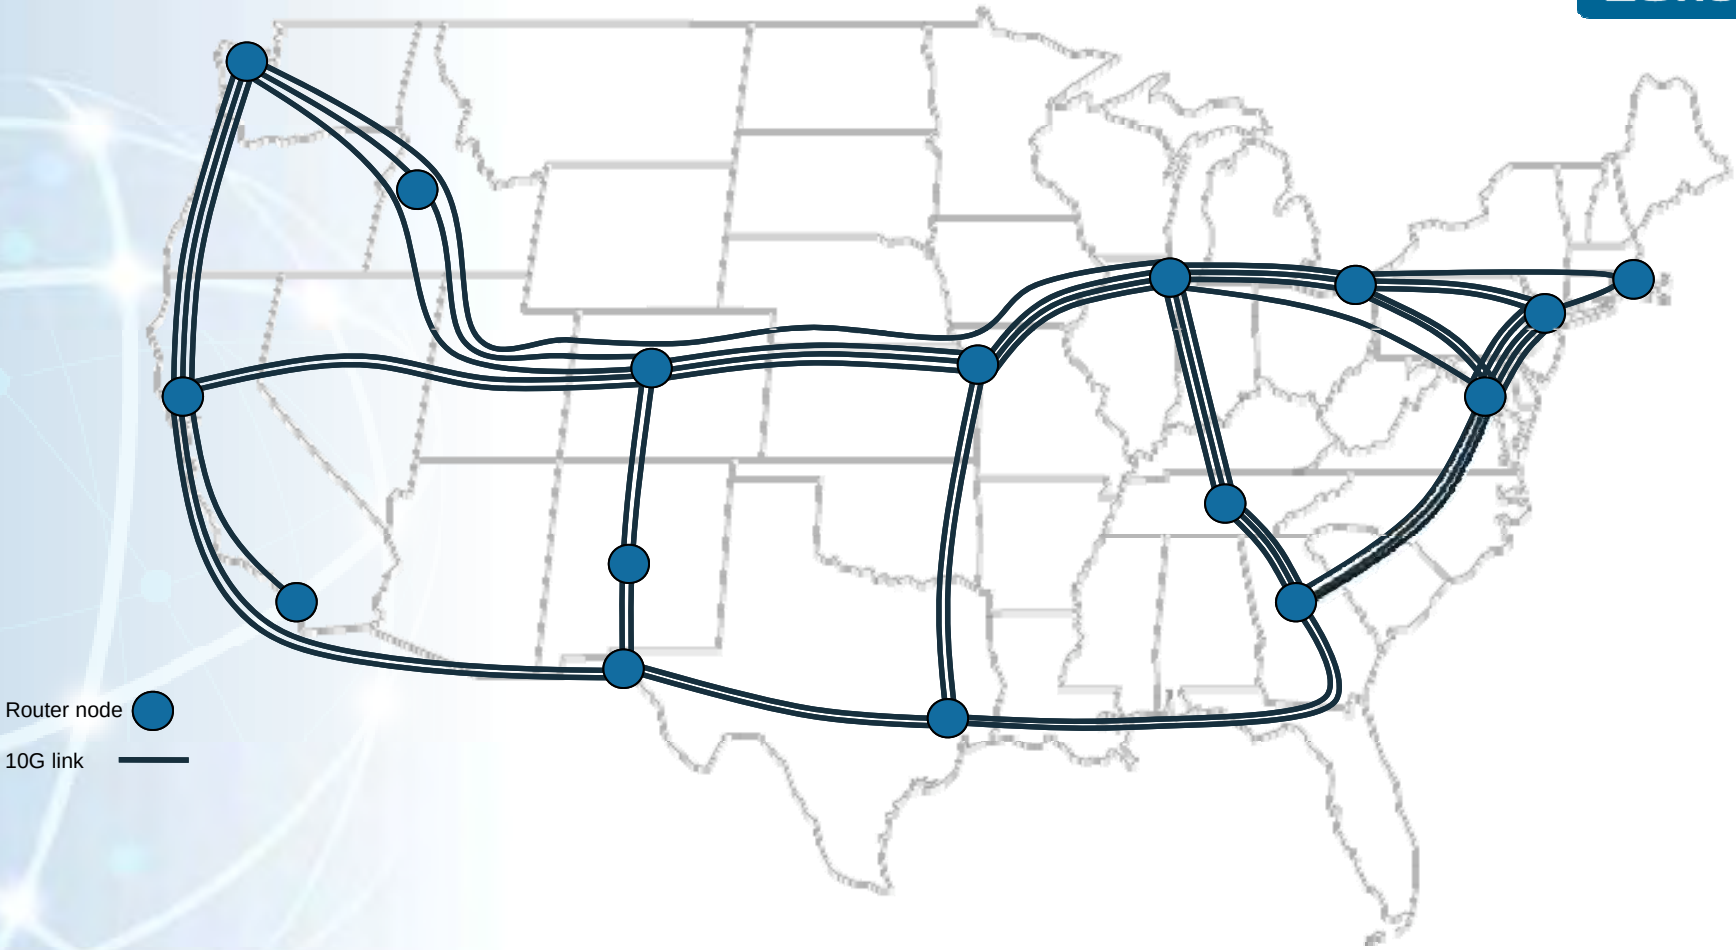

IPv6 Implementation

- IGP is ISIS
  - Common implementation for many protocols
  - Security advantages
- iBGP advertizes IPv6 over a common mesh with IPv4
  - Be careful of next-hop self
  - Not all route vendors support this
- eBGP is all native
  - ESnet does not use tunnels

Internal use

ESnet uses IPv6 whenever possible (and it usually is)

- Our mail and web services are IPv6
- DNS is IPv6
- NTP is IPv6
- Network management uses IPv6 (SNMP)
  - Fully implemented on CA Spectrum Network Management system
- Console access to most systems is ssh over IPv6

Site support

- Provide technical assistance for sites implementing IPv6
- Provide address space (Provider aggregable) in /48 chunks

IPv6 Peering

- No significant differences between IPv6 and IPv4 peering
- Most major providers now have some level of IPv6 capability
- Some run full dual stacks on peering routers
- Some still depend on tunnels to reach a limited number of dual-stack routers
- Some provide IPv6 only at a limited number of locations

The situation has improved significantly in the past 12 month for commercial providers

Issues with IPv6 support

- Many management tools do not yet support IPv6
  - This is changing, but rather slowly
  - Will change must more quickly when customers start demanding FULL IPv6 feature parity
    - (you all do that already, don't you?)
- Not much of a registry for IPv6 routing information
- Many peering monitoring tools have limited or buggy IPv6 support
- You need an addressing plan (or two or three)

IPv6 Addressing Plans

- Addressing plans are crucial to successful deployment
  - The are seldom easy
  - Will need occasional adjustment
  - May even require full replacement
    - This can almost always be avoided
- Design the addressing plan for you logical topology
- Always allocate more bits than you need!
  - Addresses are plentiful and cheap
  - Don’t be penny wise and pound foolish
- Assignments should be on nibble boundaries

The problem that is SLAAC

- SLAAC is StateLess Address Auto Configuration
  - SLAAC seemed like a good idea
    - Simplifies readdressing
    - Does not need a DHCPv6 server, only a router
- SLAAC is a bad idea
  - Removes control
  - Adds vulnerabilities
  - Lack ability to provide added information like:
    - DNS servers
    - NTP servers
    - Fallback gateways
- Eats the last 64 bits of the address

You Can't Say "NO!" to SLAAC

- Inherent in IPv6 design
- Systems often become RAs by accident
- Turning it off essentially turns off IPv6
  - Demand RA-Guard to block rogue RAs

IPv6 Security

It was often claimed that IPv6 has better security than IPv4

There is little or no basis for this!

- IPv6 implementations have far less testing to find vulnerabilities
- IPv6 is often not treated correctly by firewalls and filters
- IPv6 has the dread Next Header system which allows “hiding” malicious headers beyond the reach of most routers
- Hacker have been using IPv6 for some time and know it well
  - Often not used for hacking, but as a means of hiding activities
- Ping-pong DOS attacks are often easy
  - But they are easy to prevent/fix

Summary

- IPv6 generally works well on modern routing equipment
- Extra fees to run IPv6 are vanishing
- IPv6 is easy to set up in a backbone
- Mostly can be handled exactly like IPv4
- Your Address plan is important
- SLAAC is evil (Did I mention RA-Guard?)
- Many management and security tools are weak or simply don't support IPv6
- IPv6 presents security concerns (though most are similar to IPv4)
- The hard part of IPv6 is the services
  - Network folks have the easy part