Enterprise IPv6 Deployment
Security and other topics

6. Slo IPv6 Summit
8 Nov, 2011
Ljubljana, Slovenia

Ron Broersma
DREN Chief Engineer
SPAWAR Network Security Manager
Federal IPv6 Task Force
ron@spawar.navy.mil

IPv6 and Security

- 2006 – DREN sponsored security study
  - IPv6 is basically no more or less secure than IPv4
- NSA studies and recommendations
- Milestone objectives 1, 2, 3
  - MO3 signed in Sept
  - U.S. DoD operational networks fully approved for operating IPv6
- Other sources:
  - http://thc.org/thc-ipv6/
  - http://www.si6networks.com/presentations/hacklu2011/fgont-hacklu2011-ipv6-security.pdf
- Basic approach is to secure IPv6 network infrastructure in equivalent or better way than IPv4 network.
  - until new architectures and policies are developed, and implementations mature
  - don't want IPv6 to be the weakest link

Maturity of implementations

- Significant security concern is maturity of implementations
  - We have 30 years of maturity with IPv4 implementations
  - Much of the IPv6 code is VERY new
  - We haven't had enough time and operational experience to find all the bugs
  - How many will be discovered and exploited by adversaries?

Operational Complexity

- Added complexity increases security risk
- dual-stack can be more complex than IPv4 alone
- example: firewalls
  - are all your policies equivalent?
  - how do you keep them in sync?
  - twice as much work?

This may incentivize us to shut down IPv4 sooner than later

Rogue Router Advertisements See RFC 6104

- Router Advertisements (RAs) inform hosts of the default router/gateway
- Windows systems with Internet Connection Sharing (ICS) enabled, and IPv6 enabled, will announce itself as the default router using RAs ("Rogue RAs").
  - VERY common problem
- Hosts then start sending all their default traffic to the Windows system
- Workaround: set router preference to "high" (RFC 4191)
  - Doesn't work on JunOS
- Long term: "RA Guard" (RFC 6105) or SeND (RFC 3971)

Privacy Addresses (RFC 4941)

- Incompatible with many Enterprise environments
  - Need address stability for many reasons
    - Logging, Forensics, DNS stability, ACLs, etc.
- Enabled by default in Windows
  - Breaks plug-n-play because we have to visit every Windows machine to disable this feature.
- Just added in Mac OS X "Lion".
- Ubuntu thinking about making it default.
- Need a way for the network to inform systems about proper default on managed enterprise networks
  - new flag in RA prefix information option?

[Privacy addresses] are horrible and I hope nobody really uses them, but they're better than NAT.
… Owen DeLong, Hurricane Electric

Living with Privacy addresses

- What if the privacy address thing is a losing battle, and we have to live with it?
- We've debated the topic in various forums.
- New initiative:
  - created subnet where we allow privacy (temporary, random) addresses, and moved a bunch of machines there (Windows, Mac).
  - disabled the alarms (warning about privacy addresses).
  - modified our NDT scanner and auto-DNS-update tool to keep things updated in DNS (PTR records).
    - some argue that this should not be necessary, but some anti-spam tools will reject email from originating hosts that aren't in DNS.
  - going to generate historical database of MAC address to IPv6 address mapping, for use in IDS and forensics tools.

Other security issues

- Linux < 2.6.20 iptables dropped IPv6 frags, breaking some DNSSEC functions
  - RHEL5 uses 2.6.18
- many VPN products don't support IPv6
  - only IPv4 goes through the tunnel, not IPv6
- Symantec Endpoint Protection (SEP) breaks IPv6
  - now being fixed
- DISA STIG says to disable IPv6 in Windows
  - but Microsoft does not test this configuration
- Brocade: extended IPv6 ACLs not supported
- JunOS ACL – no "fragments" keyword for IPv6
- JunOS IPv6 IPSEC implementation flaws
  - ICMP from tunnel endpoint used wrong address

Addressing and security

- Addressing plan can be structured to align with security topology and policy
  - can greatly simplify ACLs and firewall policies

Updates from previous talk

Network Management

- We've been trying to do ALL network management using IPv6, so we can remove IPv4 from the management networks.
- Most products cannot be fully managed over IPv6

Management over IPv6 in some products

- Previously (June)…

	SSH HTTPS	DNS	Syslog	SNMP	NTP	RADIUS	Unified MIB RFC4293	Flow export	TFTP FTP	CDP LLDP
Cisco										
Brocade				1				2	3	4
Juniper								5		

- Now…

	SSH HTTPS	DNS	Syslog	SNMP	NTP	RADIUS	Unified MIB RFC4293	Flow export	TFTP FTP	CDP LLDP
Cisco[6]										
Brocade				1				2	3	4
Juniper										
ALU	5							7		

World IPv6 day

- For DREN and SPAWAR, nothing new to turn on for the day
  - every day is IPv6 day for us
- What does it look like from an enterprise perspective, where ALL clients (users) are dual-stack?

Percentage of Internet traffic over IPv6

- 1% (2009, before Google whitelisting)
- 2.5% (Google whitelisted)
- 10% (late Jan 2010, Youtube added)
- World IPv6 day… (peak at 68%)

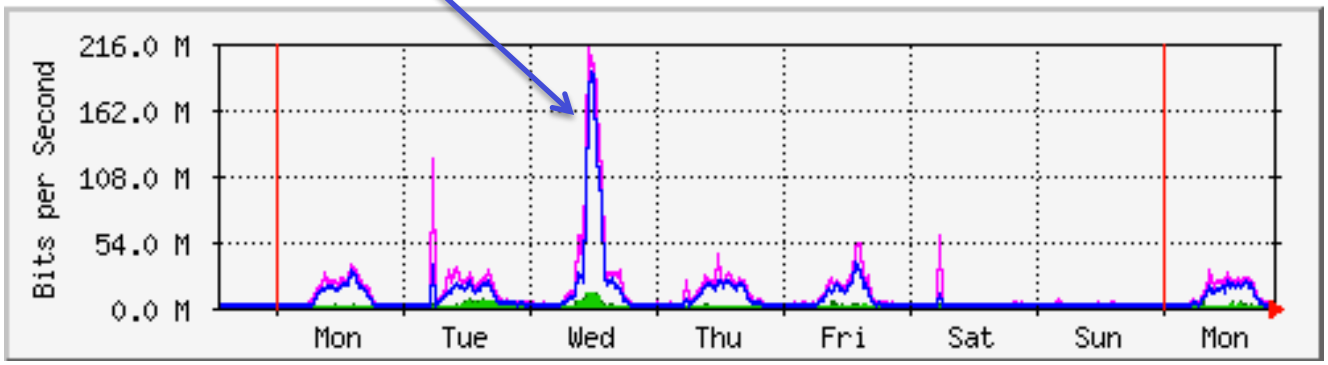

After IPv6 day

- Percentages across a day (5 min averages):

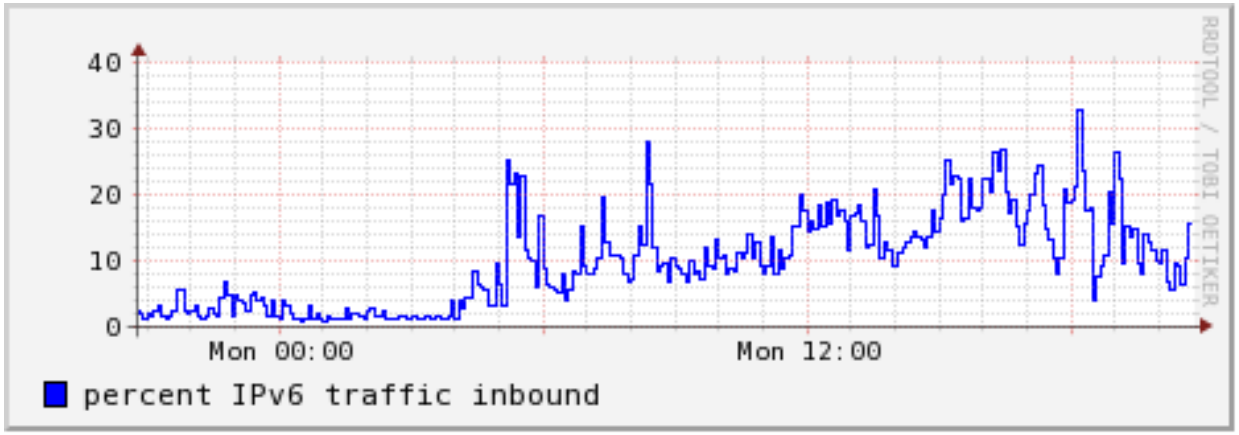

- Why higher during the work day?

After IPv6 day

- Past week (hourly averages):

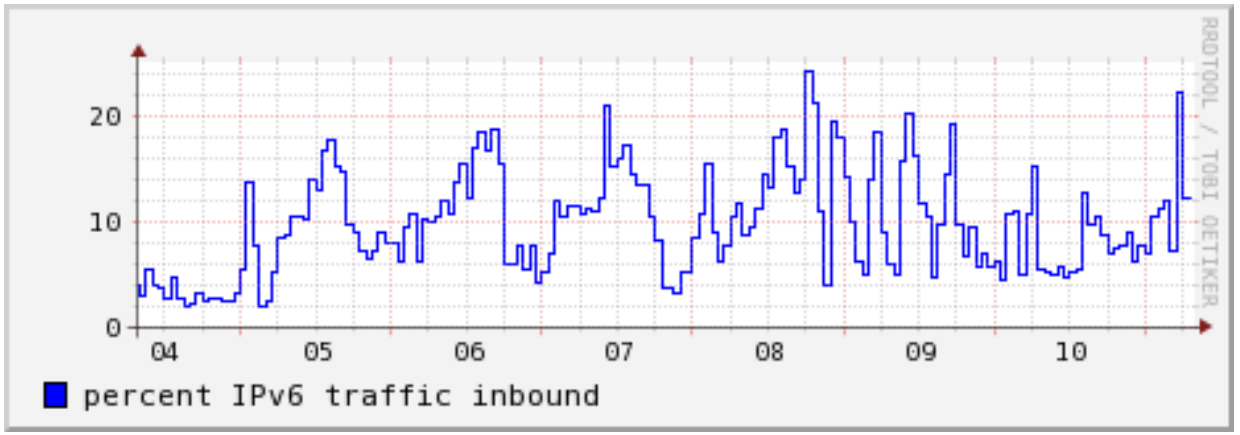

- Month (daily averages):

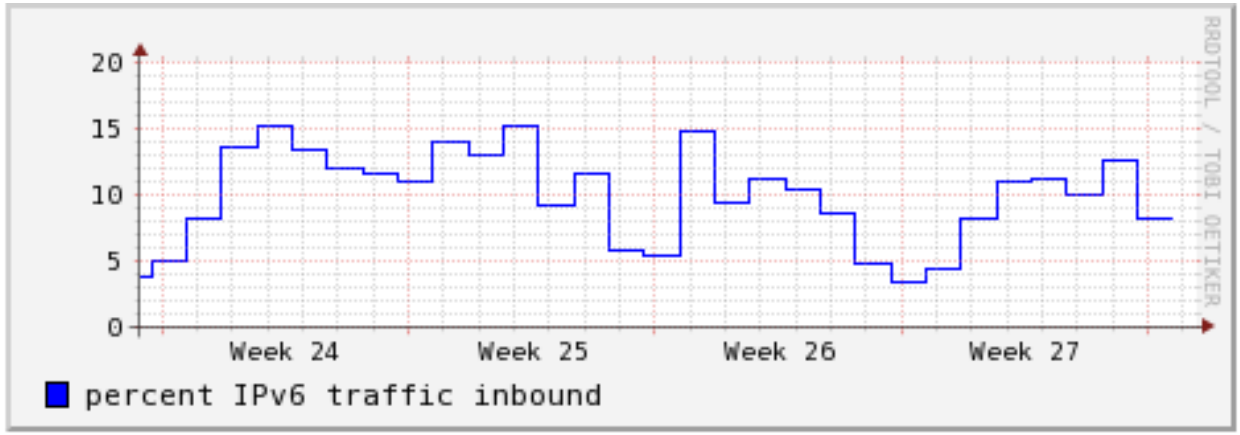

A note on Addressing

Addressing plans

- Without sufficient operational experience with IPv6 deployment, you WILL get it wrong at first.
  - usually takes 3 attempts to get it right
- Planners are hindered by IPv4-thinking
  - being conservative with address space
  - thinking “hosts” instead of “subnets”
- Typical mistakes:
  - suggesting other than /64 for standard subnet size
    - Didn’t read RFC 4291 nor 5375
  - thinking a /48 is wasteful for some small sites
  - thinking a /64 is wasteful for point-to-point links
  - request-up instead of pre-allocate-down

Addressing plans

- After operational experience, you realize:
  - you never have to "grow" subnets, so you don't need to accommodate that situation
  - if you don't use /64's for subnets, you can't do SLAAC, DHCPv6, Multicast with Embedded-RP, etc.
  - huge opportunity to align addressing with security topology, to simplify ACLs
  - can better align subneting and aggregation with existing topology
  - bad idea to embed IPv4 addresses in IPv6
  - nibble (4 bit) boundaries align better with PTR records
  - every interface has multiple IPv6 addresses
  - internal aggregation is not as important as you initially thought
  - you can do a lot of pre-allocation

Are there any near-term benefits to IPv6?

Benefits of IPv6 today (examples)

- Addressing
  - can better map subnets to reality
  - can align with security topology, simplifying ACLs
  - sparse addressing (harder to scan/map)
  - never have to worry about "growing" a subnet to hold new machines
  - auto-configuration, plug-n-play
  - universal subnet size, no surprises, no operator confusion, no bitmath
  - shorter addresses in some cases
  - at home: multiple subnets rather than single IP that you have to NAT
- Multicast is simpler
  - embedded RP
  - no MSDP

End