

# IPv6 Security

Networkshop 39, Hatfield, April 2011

Tim Chown

Electronics and Computer Science (ECS)

University of Southampton

[tjc@ecs.soton.ac.uk](mailto:tjc@ecs.soton.ac.uk)

# Do you run IPv6 yet?

- You have Win7, Linux, MacOS X, Android devices, iPhones, HP printers, etc. in your network?
- These all have IPv6 support
  - And these all enable IPv6 by default
- **So there is IPv6 in your network today**
  - Your network may well not be 'IPv4 only'
  - It's a good idea to secure it!
- This talk explores some of the issues
  - *Whether you've deployed IPv6 at your site or not*

# New protocol, new issues

- You should already have security policies that are implemented in your network
  - These are probably heavily/completely IPv4-focused
- IPv6 security issues will either
  - have IPv4 equivalents (e.g. IDS payload inspection)
  - or be completely new issues (e.g. IDS packet headers)
- You should be thinking about these now
  - So you can procure equipment to handle it
  - And ask the right questions of your vendors/suppliers

# What risks does IPv6 add?

- New attack paths
  - IPv6 is a new protocol
  - Some issues can impact an 'IPv4 only' network
- New bugs/vulnerabilities
  - Relative lack of wide-scale operational experience
  - Cisco for example already have many IPv6-specific fixes
- Lack of admin staff knowledge and training
  - Do you filter ICMPv6 type 134 in your switches today?
- IPv6 incidents/issues not detected
  - Because you don't know what IPv6 traffic you have

# What might your users do?

- Unknowingly enable an IPv6 router
  - e.g. by turning on Internet connection sharing
- Consciously use an IPv6 tunnel broker
  - e.g. [www.sixxs.net](http://www.sixxs.net) or [www.tunnelbroker.net](http://www.tunnelbroker.net)
  - Traffic may be tunneled in UDP, AYIYA tunnels
  - Might be offering a /64 to a whole IPv4 subnet on your site
- Unwittingly use IPv6
  - Due to local service discovery
  - Or perhaps some automated tunneling scheme
- Configure manual tunnels through your IPv4 subnets
- Never underestimate users 😊

# IPv6 differences

	IPv4	IPv6
Address length	32 bits	128 bits
Default prefix length	Varies, typically /24	/64
Address configuration	DHCPv4	Stateless Autoconfiguration DHCPv6
Default addresses used	Private or Global	Link-local and Global
Address resolution	ARP	Neighbour Discovery (ND)
Minimum MTU	576	1280
Fragmentation	By hosts or routers	Only by hosts
Host Path MTU Discovery	Optional	Required
IPsec	Optional	'SHOULD' (draft-ietf-6man-node-req-bis-08 )
Private addressing	RFC 1918	Unique Local Addresses (ULAs) (not for use with NAT)

# New: address scopes

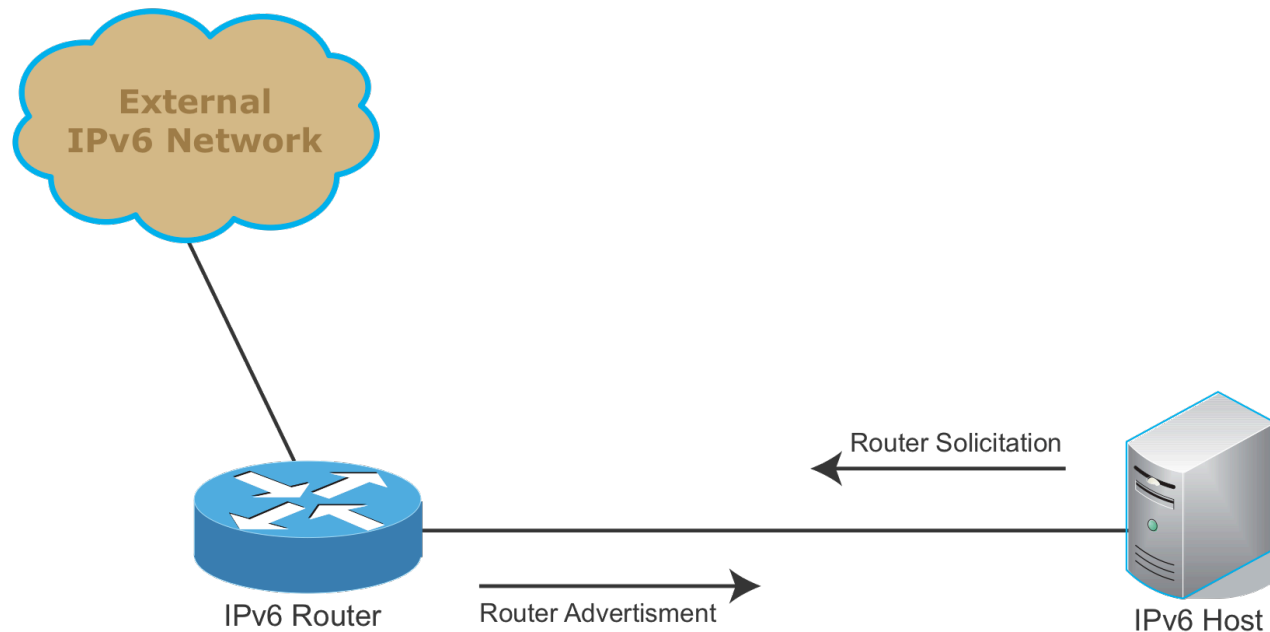
- IPv4
  - Usually just one address
  - Global, or Private (RFC 1918) + NAT
- IPv6
  - Link-local (under fe80::/10, not routed)
  - Unique Local Addresses (under fc00::/7) (RFC 4193)
  - Global (no NAT)
- **IPv6 hosts are naturally multi-addressed**
  - And when dual-stack with an IPv4 address too
  - Your management/monitoring tools must cope!

# New: address management

- Hosts can autoconfigure basic network configuration by soliciting/receiving IPv6 Router Advertisements (RAs)
  - They form their own address by combining 64-bit network prefix in the RA with MAC address + 16bits of padding, e.g.
    - MAC address is 08:00:20:9c:14:66
    - Network prefix is 2001:630:80:2::/64
    - Address is 2001:630:80:2:0a00:20ff:fe9c:1466
- Currently all campus-style IPv6 deployments are dual-stack
  - IPv4 address configuration by DHCPv4
  - IPv6 by Stateless Autoconfiguration (SLAAC) (RFC 4862)
- Dual-stack hosts rely on IPv4 DHCP to supply DNS resolver and other configuration information
  - DHCPv6 is available on many platforms, but not in wide use

# SLAAC operation

- Totally dependent on Router Advertisements
  - RA is multicast on local subnet
  - (link-local) RA source address implies default router



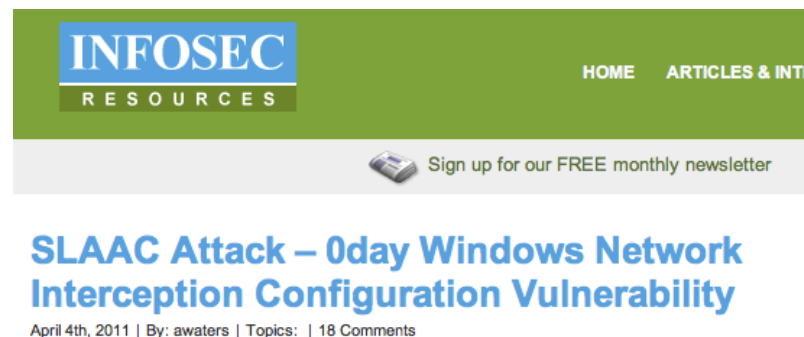
# Implications of RAs

- Host autoconfiguration can be a nice feature
- BUT hosts can **send** RAs too
  - Accidental or malicious, see RFC 6104
  - So **you should mitigate this...**
    - Use RA Guard (RFC 6105)
    - Filter ICMPv6 Type 134 on non-router switch ports
    - Demand RA Guard or equivalent capability in wireless products
    - Deploy RAmond (<http://ramond.sourceforge.net>)
- Note : all IPv6 networks must use RAs
  - There's no DHCPv6 Default Gateway option
- RAs can carry DNS resolver information (RFC 6106)
  - So could trick other hosts into using bogus DNS resolver

# Stats on campus rogue RAs

- Used RAmond on a ~50 AP dual-stack wireless network
  - RAmond issues deprecating RA against rogues
  - Rogue device may not actually be turned off for some time
- Period of 376 days (2010-02-18 to 2011-03-01)
  - Rogue RA seen on 228 of those days (60%)
  - 257,669 rogue RAs seen, **all** for 2002::/16 (connection sharing?)
  - 35 different MAC sources using 38 different link layer sources
  - Only two devices used EUI-64 addresses, one was an HTC
  - Four devices sent only one rogue RA
- Presence of a rogue RA may cause connectivity problems
  - **Even on an IPv4-only network if the hosts have IPv6 enabled**

# Recent 'Oday' IPv6 vulnerability



- Not a true Oday exploit, but it raises awareness
  - Attacks an ostensibly IPv4-only network
  - Combines Rogue RA issue with DHCPv6 and NAT-PT
  - <http://resources.infosecinstitute.com/slaac-attack/>
- NAT-PT does IPv6 to IPv4 translation, and includes a DNS proxy function, which can give evil responses
  - NAT-PT deprecated by IETF, but current variant NAT64 would probably serve the same purpose

# More multi-addressing

- So back to addresses... a dual-stack host can have
  - An IPv6 link-local, IPv6 global and an IPv4 global
- SLAAC by default has privacy concerns because the 64-bit EUI-64 host part is always the same
  - Devices become trackable across visited /64 subnets
- IETF answered this with Privacy Addresses (RFC 4941)
  - Essentially allows a host to use a random 64-bit host part
  - Can rotate privacy address periodically, e.g. every 24 hours
  - May keep previous addresses in use for up to 7 days
- Now we have even more addresses!

# Windows XP example

## Run ipconfig

Interface 4: Ethernet: Local Area Connection

uses Neighbor Discovery

uses Router Discovery

link-layer address: 00-00-cb-68-0b-2e

preferred global 2001:630:d0:112:309e:3ba9:d0df:1afc, life 57m25s/27m25s (temporary)

deprecated global 2001:630:d0:112:cc4e:835c:7e1b:e482, life 57m25s/0s (temporary)

deprecated global 2001:630:d0:112:f4c5:398e:b5f3:bf58, life 57m25s/0s (temporary)

deprecated global 2001:630:d0:112:88bd:46d0:b997:6dc4, life 57m25s/0s (temporary)

deprecated global 2001:630:d0:112:e07c:fe6b:a58a:1608, life 57m25s/0s (temporary)

deprecated global 2001:630:d0:112:b4dc:cfc5:c6a7:3724, life 57m25s/0s (temporary)

deprecated global 2001:630:d0:112:1ca9:c9b:849e:7869, life 57m25s/0s (temporary)

preferred global 2001:630:d0:112:200:cbff:fe68:b2e, life 57m25s/27m25s (public)

preferred link-local fe80::200:cbff:fe68:b2e, life infinite

Temporary addresses are IPv6 Privacy Addresses

These change over time – default every 24 hours on Windows XP

Host also has standard global IPv6 address

Privacy addresses only used for initiating connections from host

# Privacy addresses and logging

- How many hosts do you have?
  - RFC 4941 allows an address change every 10 mins



**IPv6 host using privacy addresses**

**Generates a new privacy address every 24 hours on its subnet 2001:db8:d0:1::/64**

Web server log file:

```
Day 1: 2001:db8:d0:1:5d34:2614:f422:cb83
Day 2: 2001:db8:d0:1:69dd:23a1:7d43:90e1
Day 3: 2001:db8:d0:1:c461:83e2:f448:91b2
```

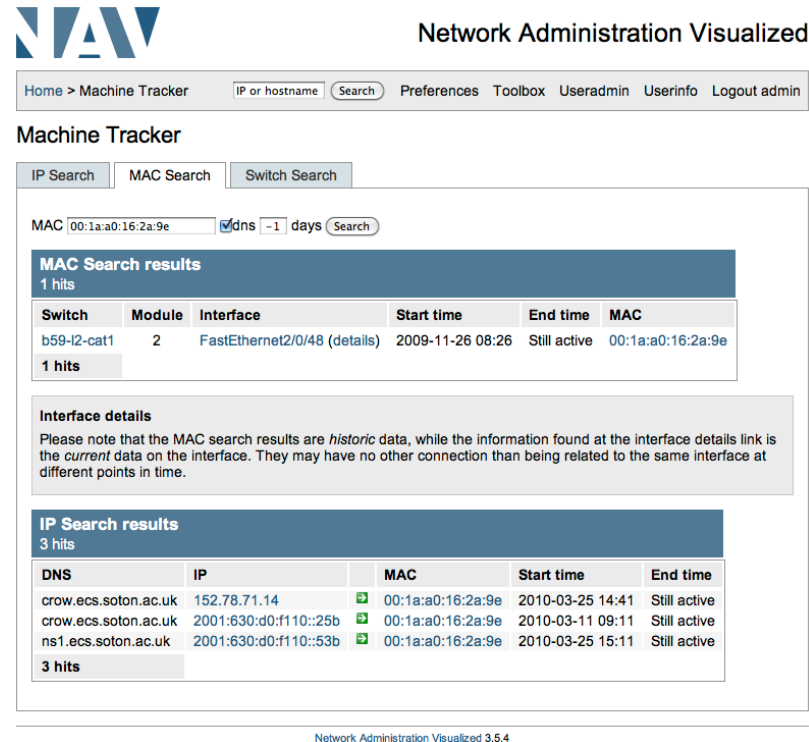
...

# New issue: accountability

- How do you tie host-generated addresses to users?
  - Use 802.1x on wired ports?
  - Scrape IP-MAC data from switch ports by SNMP?
  - Capture ND traffic? (see <http://www.digriz.org.uk/slaacer>)
- Or try to control it by only using DHCPv6?
  - Implementations are out there (Win 7, Win Server, ISC, ...)
  - But DHCPv6 uses new DHCP Unique Identifiers (DUIDs)
  - DUIDs might change with OS install – can cause problems for administrators provisioning new systems
  - A persistent DUID solution under discussion in IETF
    - <http://tools.ietf.org/html/draft-narten-dhc-duid-uuid-01>

# Living with multi-addressing?

- Accept host-generated multi-addressing and deal with it
  - You may have no choice with unmanaged hosts
- Use appropriate tools, e.g. NAV, <http://metanav.uninett.no>
  - Falls into the ‘SNMP scraping’ category



The screenshot displays the NAV web interface. At the top, the NAV logo and the text "Network Administration Visualized" are visible. Below this is a navigation bar with links for "Home > Machine Tracker", "IP or hostname Search", "Preferences", "Toolbox", "Useradmin", "Userinfo", and "Logout admin". The main content area is titled "Machine Tracker" and contains three tabs: "IP Search", "MAC Search", and "Switch Search". The "MAC Search" tab is active, showing a search for MAC address "00:1a:a0:16:2a:9e" with a "Search" button. Below the search bar, a section titled "MAC Search results" shows "1 hits" and a table with the following data:

Switch	Module	Interface	Start time	End time	MAC
b59-l2-cat1	2	FastEthernet2/0/48 (details)	2009-11-26 08:26	Still active	00:1a:a0:16:2a:9e

Below the table, there is a section for "Interface details" with a note: "Please note that the MAC search results are *historic* data, while the information found at the interface details link is the *current* data on the interface. They may have no other connection than being related to the same interface at different points in time."

At the bottom, the "IP Search results" section shows "3 hits" and a table with the following data:

DNS	IP	MAC	Start time	End time
crow.ecs.soton.ac.uk	152.78.71.14	00:1a:a0:16:2a:9e	2010-03-25 14:41	Still active
crow.ecs.soton.ac.uk	2001:630:d0:f110::25b	00:1a:a0:16:2a:9e	2010-03-11 09:11	Still active
ns1.ecs.soton.ac.uk	2001:630:d0:f110::53b	00:1a:a0:16:2a:9e	2010-03-25 15:11	Still active

The footer of the page reads "Network Administration Visualized 3.5.4".

# New IPv6 protocol issues

- IPv6 uses ICMPv6 as inherent part of protocol
  - Must be very careful filtering ICMPv6 traffic
    - See RFC 4890 for recommendations, e.g. for PMTUD
- IPv6 uses link-local multicast
  - e.g. ff02::1 is all hosts, ff02::2 is all routers
    - Be careful about multicast in older switches/hubs
- IPv6 has streamlined header format
  - But includes new Extension Header concept
    - Chain of headers may be present
    - Firewalls will need to be able to process these

# ICMPv6 (RFC 4443)

- Provides many critical IPv6 functions, including:
  - ND – Neighbour Solicitation / Neighbour Advertisement
    - Provides the equivalent of ARP for IPv4
  - SLAAC – Router Solicitation / Router Advertisement
  - PMTUD (Packet Too Big)
    - Must work for IPv6 because routers cannot fragment
  - Echo Request / Reply
  - Destination Unreachable
  - Redirects
  - Multicast Listener Discovery (MLD) (like IGMP for IPv4)
  - ...
- For a full list see
  - <http://www.iana.org/assignments/icmpv6-parameters>

# Securing ND?

- Cryptographically Generated Addresses (CGAs)
  - See RFC 3971
  - Hash { public-key, prefix, modifier } => 64-bit host part
- SEcure Neighbour Discovery (SEND) (RFC 3971)
  - Sign Neighbour Advertisements with private key
  - Can also be used to validate Ras
  - But very few implementations to date
- This is quite a 'heavyweight' approach
  - To compare: do you run Authenticated DHCPv4?

# Protecting the first-hop

- SEND is unlikely to be widely deployed in HE/FE
- Rogue RA mitigation is easier
  - Use ACLs (if your switch supports ICMPv6 filtering)
  - Use RA Guard if/when your vendor provides it
- You want DHCPv6 snooping as per DHCPv4
  - Protect against rogue DHCPv6 servers (a la '0day')
  - Again, see if your vendor has it on their roadmap
- Run monitoring daemons, see what's there
  - NDPmon (<http://ndpmon.sourceforge.net>)
  - RAmond

# Useful packages to try?

- THC toolkit
  - The most well-known suite of IPv6 attack tools
  - e.g. Duplicate Address Detection (DAD) denial of service, smurf6, ...
  - <http://freeworld.thc.org/thc-ipv6/>
- For crafting packets
  - <http://www.secdev.org/projects/scapy/>
- Scanning tools
  - Most bigger packages have some support, e.g. Nessus
- NDPmon
- Ramond
- tcpdump/Wireshark

# IPv6 firewall requirements?

- What to put into procurements?
  - See <http://www.ripe.net/ripe/docs/ripe-501>
- Be able to filter based on IPv6 devices/prefixes
  - Configure dual-stack seamlessly for consistent policy implementation
- Management over IPv6 transport
- Handling IPv6 multicast
  - Embedded-RP routing (RFC 3956) , MLDv1/v2
- Configure RAs on directly-connected links
- Support IPv6 Extension Header processing
- ND cache robustness

# Extension Headers

- Extension header ordering rules are cited in RFC 2460
- Currently about 15 headers defined
  - Some possible issues with some of these
- Hop-by-hop
  - Router alert option
  - Other options => possible attack on router CPU
- Routing header
  - RH0 - allows source routing
  - Can change point of attack, or make amplification attack
  - Should drop RH0 packets - see RFC 5095
- Check what your firewall vendor supports

# Network scanning

- More difficult to scan for vulnerable hosts in a /64
  - See RFC 5157
- Consider other ways attackers may gain targets
  - Publicly advertised servers (DNS, www, MX, etc)
  - Exposure of information in your DNS
  - Your hosts connecting to remote servers may have activity logged
- Hopefully should not see NAT used with IPv6
- In ECS we don't see any classic port scanning
  - Do see probes on ports on 'advertised' servers
  - Do see some badly crafted packets

# ND cache exhaustion?

- There is a concern with whether rapid scans to non-existent IPs in a /64 subnet can fill a router's ND cache before the ND operations complete
  - IETF now recommends /127 for point-to-point links
    - See draft-ietf-6man-prefixlen-p2p-01
  - No reported problems yet with /64 host subnets
    - Point to point links tend to be more 'exposed'
  - Again, check what your vendor says
- Note - no publicly reported cases of this yet
  - May wish to run your own tests?

# IP-based security

- Think about where IPv4 address controls are currently implemented
  - Probably quite a few places!
- Do you have IPv6 equivalents?
  - /etc/hosts.allow, /etc/hosts.deny
  - Web server .htaccess permissions
  - TCP wrappers
  - MX relays accepting internal connections
  - DNS zone transfers
  - ...
- Don't allow IPv6 back-doors into dual-stack systems

# Impact of IPv6 tunneling tools

- Huge range of tunneling tools out there, e.g.
  - Manually configured tunnels (GRE, Protocol 41)
  - Tunnel broker (RFC 5572)
  - 6to4 (RFC 3056)
  - ISATAP (RFC 5214, sparse tunnels intra-site)
  - Teredo (RFC 4380, IPv6 through IPv4 NATs)
  - Dual-stack Lite (IPv4 in IPv6 tunneling)
  - ...
- These have the potential to bypass your security
  - Try to detect usage of these

# Detecting IPv6 tunnels

- What properties can be used to detect IPv6 tunneling?

Protocol	Clue
6to4	IPv4: Protocol 41, IPv4 activity to 192.88.99.1 IPv6: Use of 2002::/16
Teredo	IPv4: Clients talking to server on UDP/3544 IPv6 :Use of 2001:0::/32
ISATAP	ISATAP router discovery may cause DNS lookups for isatap.<domain>
Tunnel broker	Trickier – UDP encapsulation to broker servers

- Tunneling may be required for early deployment
  - Terminate tunnel outside your firewall

# World IPv6 Day

- Organised by ISOC (Internet Society)
  - June 8<sup>th</sup> 2011
  - <http://isoc.org/wp/worldipv6day/>
  - Google, Facebook, Akamai, ...
  - Big providers adding IPv6 DNS records on primary domains
- Allows providers to measure IPv6 ‘brokenness’
- Allows IPv6 sites to test performance
  - Clients will start using IPv6 for YouTube for example
    - How much of your external traffic is that?
- Perhaps the day can be a focus for your efforts?
  - See draft-chown-ipv6-call-to-arms-01 (-02 soon)

# Summary

- IPv6 is in your network today
  - Your users may be running it, intentionally or not
- Get management buy-in to monitor and control it
- What IPv6 traffic can you find?
  - Native traffic (Ethertype 0x86dd), tunneled traffic
- Apply appropriate IPv6 controls, even if 'IPv4 only'
  - e.g. rogue RA protection, DHCPv6 snooping
- When deploying IPv6, consider security equivalence
  - Check capabilities for new IPv6 features with vendors
- Decide your plan for user accountability for IPv6
  - Hosts may be self-configured with multiple addresses

# Useful References

- “*IPv6 Security*”, Hogg S., Vyncke E., Cisco Press,  
<http://www.amazon.co.uk/IPv6-Security-Scott-Hogg/dp/1587055945>
- “*Guidelines for the Secure Deployment of IPv6*”, NIST, US Dept of Commerce,  
<http://csrc.nist.gov/publications/nistpubs/800-119/sp800-119.pdf>

# Extra bits

IPv6 performance issues:

Tunneling

IPv4 fallback delays

# Issue: IPv4 fallback performance


- What happens if IPv6 network connections fail and browsers need to fall back to IPv4?
  - Not strictly a security issue, but important
- Stats gathered by Teemu Savolainen (Nokia)
  - Presented at Prague IETF, March 2011
- Interesting to see how/if unreachable indication to host helps
  - A key point to note is Windows client performance
    - 21 second delay, whichever browser used, whether unreachable received or not

Device	DNS query sending style	IPv6 broken, time until fallback to IPv4			Comments
		Black hole	No route	Address unreachable	
Symbian^3 on Nokia N8 (11.012)	<b>A</b> first and used if possible. <b>AAAA</b> if no IPv4.	N/A	N/A	N/A	Symbian^3 prefers IPv4 hence tested fallback scenarios are N/A. The DNS query order is a configuration parameter.
Windows 7 Starter Edition on HP IE 8.0.7600 & Google Chrome 8.0.552.224 & Safari 5.0.2	<b>A</b> and after reply <b>AAAA</b> . Uses IPv6 if both available.	~21s	~21s (after 3 SYN & ICMPv6 errors)	~21s (after 3 SYN & ICMPv6 errors)	Same initial delay with those browsers. <b>The 21 seconds is TCP timeout after 3rd SYN failed.</b>
iOS4 4.2.1 on Apple iPhone4 Safari	<b>A</b> first and <b>AAAA</b> immediately after. Uses IPv6 if both available.	<b>No fallback</b>	~4s (After 5 SYN & ICMPv6)	~4s (After 5 SYN & ICMPv6)	Lucky observation: waits ~350 ms for AAAA to arrive after A is received before going for IPv4
Apple OS/X 10.6.6 on iMac Safari 5.0.3 Firefox 3.6.13	<b>A</b> first and <b>AAAA</b> immediately after. Uses IPv6 if both available.	~75s	~4s (After 5 SYN & ICMPv6)	~4s (After 5 SYN & ICMPv6) <b>Firefox: no fallback at all!</b>	Special note that Firefox did not fallback on address unreachable error.
Android 2.3.1 on Samsung Nexus S Native browser	<b>AAAA</b> and after reply <b>A</b> . Uses IPv6 if both available.	~21s	~0s (acts on first ICMPv6)	~0s (acts on first ICMPv6)	<b>The 21 seconds is TCP timeout after 3rd SYN failed.</b>
Maemo5 IPv6 enabled version on Nokia N900 Firefox & native	<b>AAAA</b> and after reply <b>A</b> . Uses IPv6 if both available.	~189s	~0s (acts on first ICMPv6)	~0s (acts on first ICMPv6)	189s is after 6th SYN failed. Kernel: 2.6.28-based
Ubuntu 10.04 / 10.10 on "PC" Firefox 3.6.13	<b>AAAA</b> and after reply <b>A</b> . Uses IPv6 if both available.	~21s	~0s (acts on first ICMPv6)	~0s (acts on first ICMPv6)	Note: immediate fallback to IPv4 happens also during complex page load (i.e. minimizes damage when IPv6 is always preferred) Kernel (10.04): 2.6.32-27, (10.10): 2.6.35-24

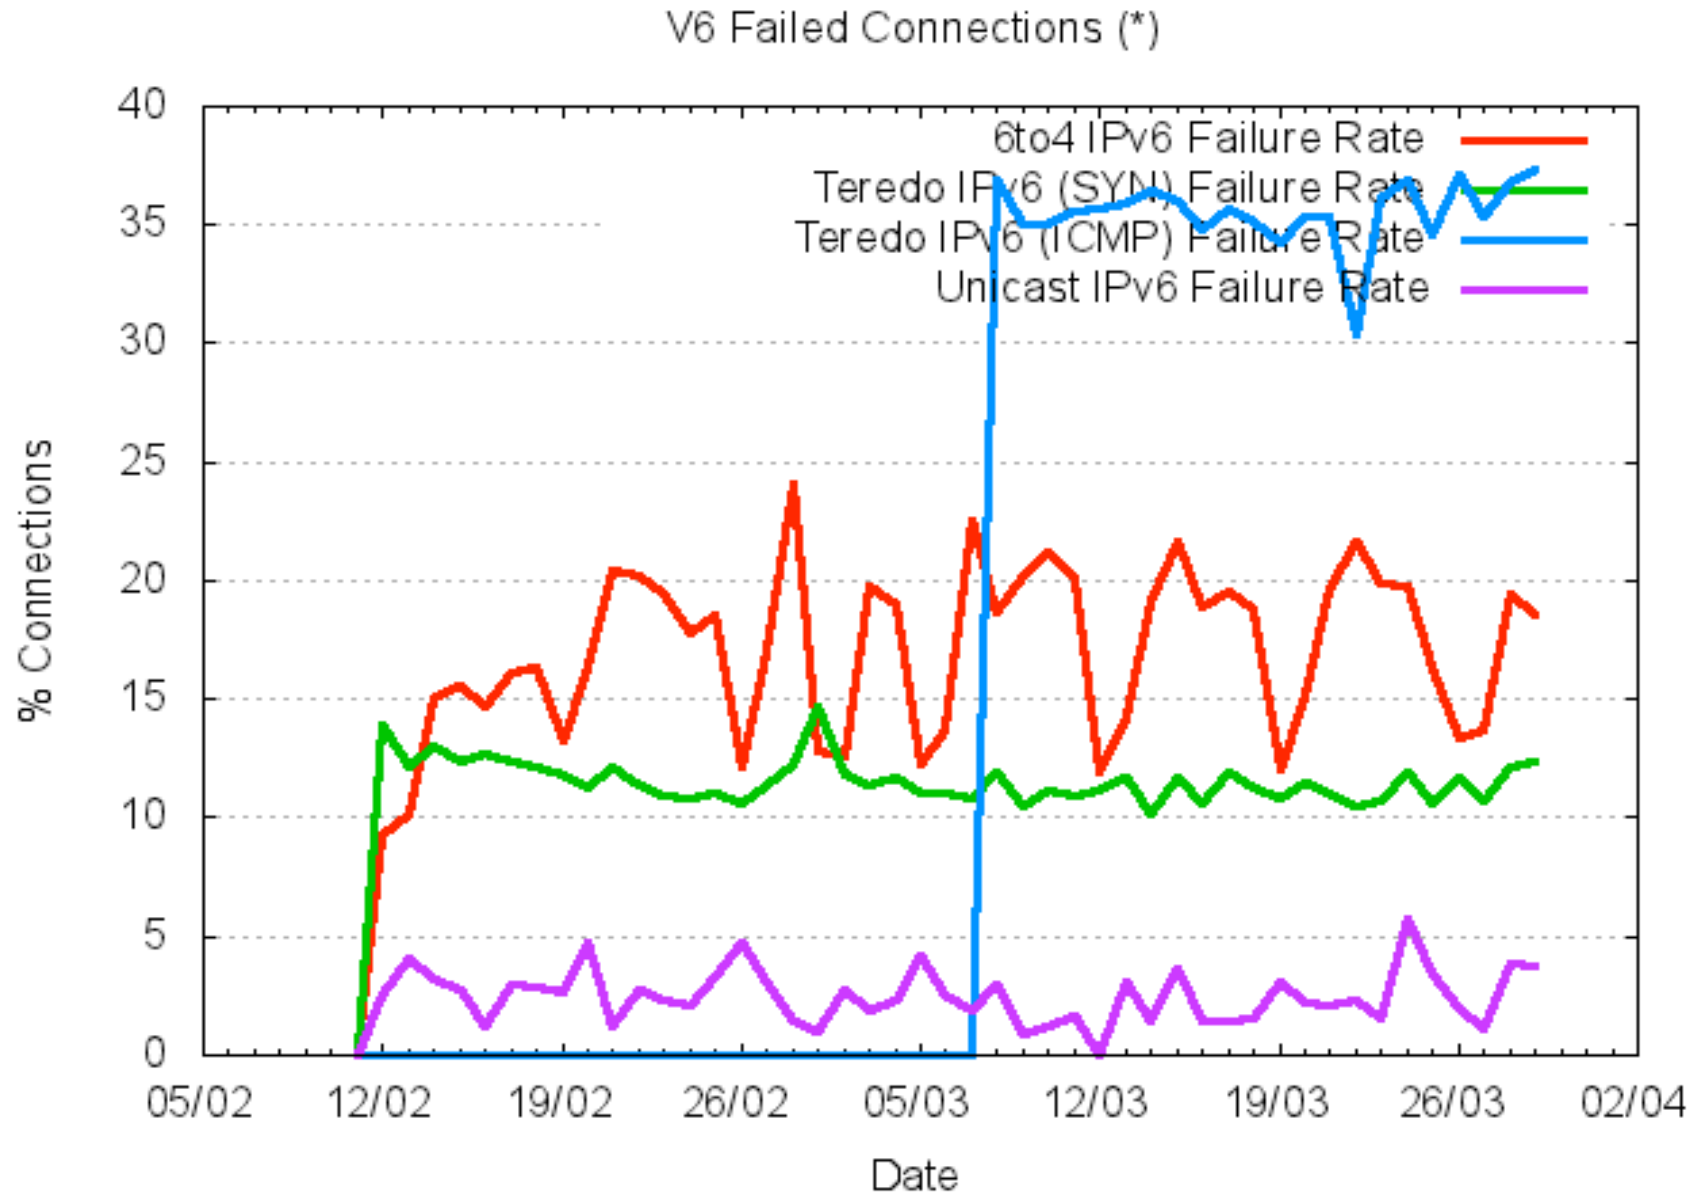
# Issue: Tunnel performance

- How do 6to4 and Teredo perform?
- Do we want users using these methods?
- Stats gathered by Geoff Huston and presented at Prague IETF, March 2011
  - Shows increased latencies for 6to4/Teredo
  - Shows really bad connection failure rates when IPv6 literals tested (forcing IPv6 choice)

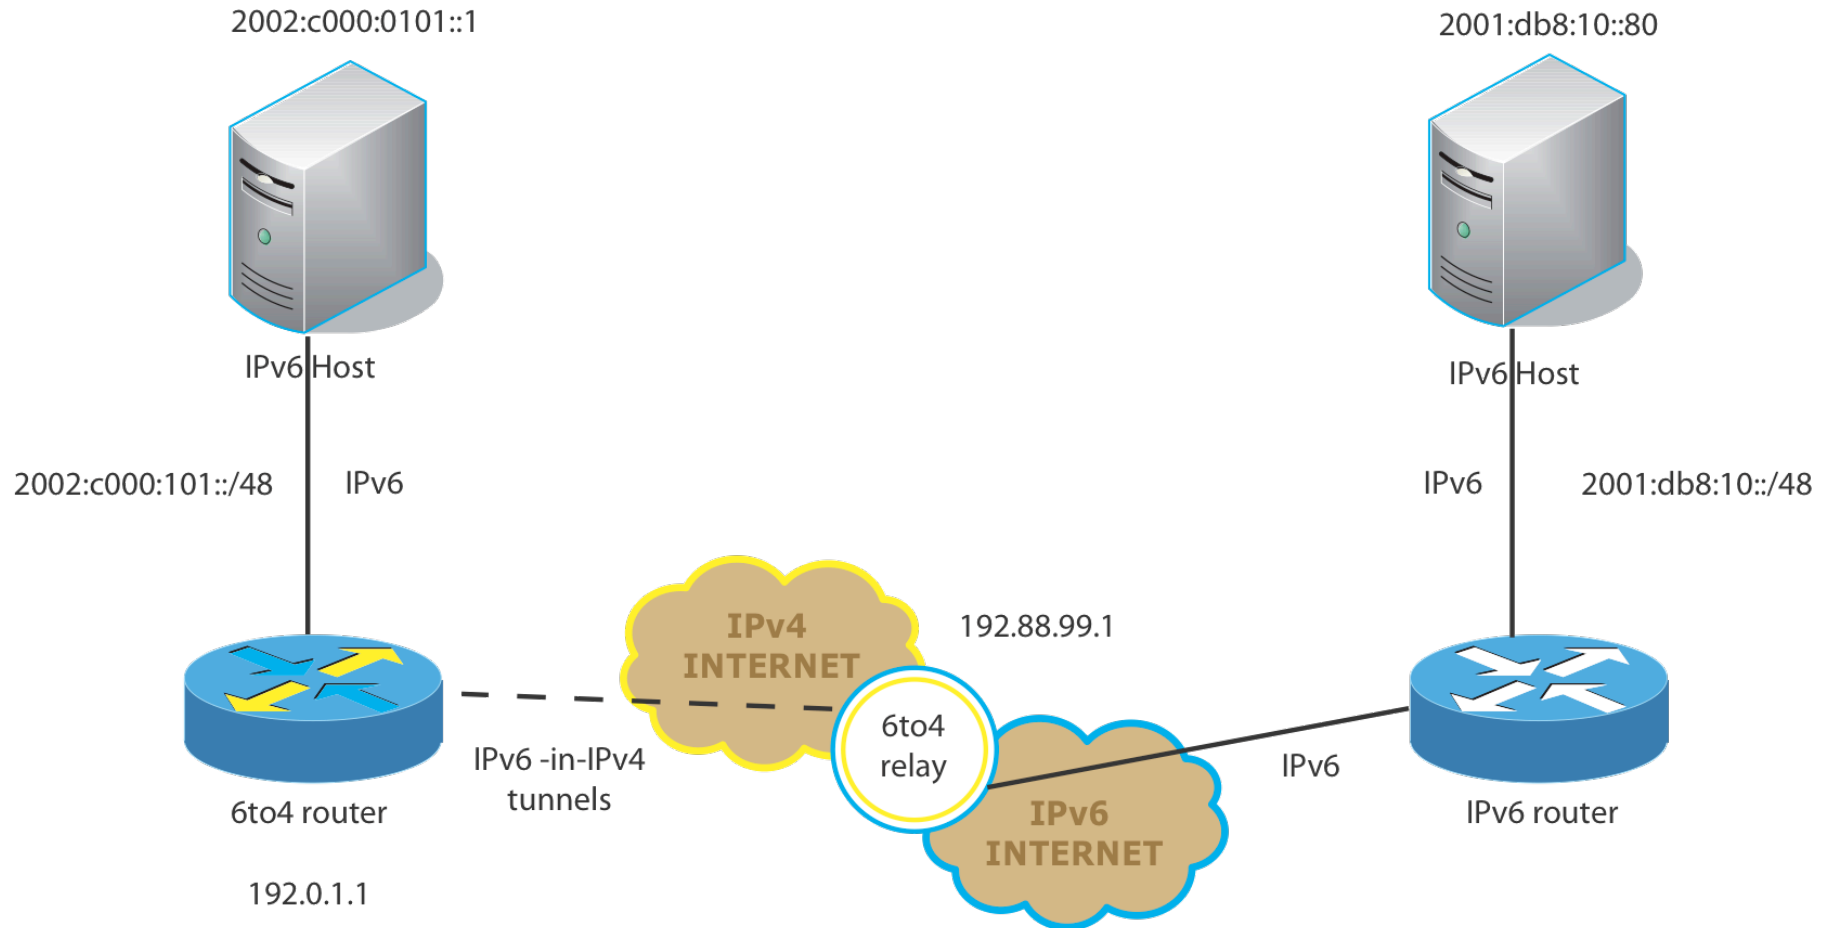
# 6to4/Teredo Performance



# IPv6 Connection Failure using V6 Literal



# 6to4 highly dependent on relay



# 6to4 at the IETF

- In the Prague IETF in March 2011 the IETF had consensus to progress two drafts:
- Moving 6to4 to Historic
  - draft-ietf-v6ops-6to4-to-historic-00
  - Should be no new usage of 6to4, deprecate prefix
- Advisory on use of 6to4
  - draft-ietf-v6ops-6to4-advisory-00
  - Calls for ISPs to deploy relays
  - How to make the best of 6to4 if absolutely needed