

“If We Deploy IPv6, Will It Help or Hurt Our Security?”

Internet2/ESnet Joint Techs
College Station TX, February 2nd, 2009

Joe St Sauver, Ph.D.
Manager, Internet2 Security Programs
(joe@uoregon.edu or joe@internet2.edu)

<http://www.uoregon.edu/~joe/ipv6-security/>

Disclaimer: The opinions expressed in this talk do not necessarily reflect the opinion of any other party.

Some Notes Before We Get Started

- ***Goals of This Talk:*** This talk is meant to help the community make progress getting IPv6 deployed while avoiding both:
 - paralysis associated with non-specific/speculative security worries and
 - action catalyzed solely by unfounded hopes for security improvement.We want you to get native IPv6 deployed, but security should neither be the main reason for deploying IPv6 nor a deployment roadblock.
- I'm also going to use this talk to introduce some new security topics that you might want to begin thinking about, such as what you are (or aren't) doing with IPsec, and how you should be handling IPv6 multihoming.
- ***Technical Level of This Talk:*** At least in some cases, this talk may end up being shared with less technically-oriented colleagues after Joint Techs. Therefore, I've attempted to set the technical level of this talk at a level that has something for both technical and non-technical people.
- ***Disclaimer and Acknowledgement:*** While I'm solely responsible for the content of this talk, I'd like to acknowledge the extremely helpful discussions that have occurred on the Internet2 IPv6 mailing list, as well as via private email with a number of you. Thank you very, very much!

Format of This Talk

- This talk has been prepared in my customary overly detailed format. I use that format because:
 - doing so helps to keep me on track when I have limited time
 - audience members don't need to scramble to try to take notes
 - if there are hearing impaired members of the audience, or non-native-English speakers present, a text copy of the talk may facilitate their access to this material
 - a detailed copy of the talk makes it easy for those who are not here today to go over this talk later on
 - detailed textual slides work better for search engines than terse, highly graphical slides
 - hardcopy reduces problems with potential mis-quotation.
- BUT I promise that won't read my slides to you, and I also promise that I won't go over my twenty minutes. Speaking of time...

1. IPv6: It's Time!
**(And Be Sure Your Downstream Folks
Are Ready to Do IPv6 Too!)**

The Classic Era of IPv4 Abundance Is Ending

- Geoff Huston of APNIC has done excellent work building a model forecasting the time it will take for IANA (the Internet Assigned Numbers Authority) to allocate its remaining IPv4 address blocks to the RIRs (regional internet registries, such as ARIN, RIPE, APNIC, LACNIC and AFRINIC), and for the RIRs to allocate their remaining IPv4 addresses to large ISPs and other direct customers.

- As of 28 Jan 2009 www.potaroo.net/tools/ipv4/index.html says:

Projected **IANA** Unallocated [IPv4] Address Pool Exhaustion:

24-Mar-2011 [2/2/2009 --> 3/24/2011 = 780 days or **2y 1m 22d**]

Projected **RIR** Unallocated [IPv4] Address Pool Exhaustion:

31-May-2012 [2/2/2009 --> 5/31/2012 = 1214 days or **3y 3m 29d**]

- **Obviously we don't have very much time left.**

Plus Or Minus...

- Those are **only estimates**, and twill change from day to day. I think they're good estimates, but you may end up with more or less time.
- For example, the rate of address consumption might **increase**:
 - what if there are multiple new major broadband build out efforts either here in the US or abroad, needing many more IPv4 addresses for newly connected broadband customers? (remember the Administration's Economic Stimulus packages)
 - some sites may engage in "last minute"/"panic'd" speculative requests for additional IPv4 addresses "just in case"
- Alternatively, the rate of consumption may **flatten out**:
 - sites may act responsibly, and make a real effort to consciously limit requests for additional IPv4 address space
 - the community, through the RIRs, may tighten up requirements for receiving IPv4 address space, thereby slowing the depletion
- Some **existing/already-allocated address space might be returned**, thereby increasing the pool of available addresses (but I wouldn't count on large miracles happening very often)

You've Got a Lot To Do In A Short Time


- I mention the dwindling amount of time not to scare you, but to help make the point that the classic era of IPv4 abundance is rapidly drawing to a close, and a new blended IPv4/IPv6 era is dawning.
- You and your institution should be thinking about what you're going to do when it comes to supporting IPv6, and you should be doing that brainstorming and planning **now**.
- Why now, rather than a year or two from now as the pool of IPv4 addresses gets closer to depletion? Well, if you think you have "lots of time" to get ready for IPv6, remember, you'll need time to:
 - develop a plan if you don't have one already,
 - train your staff and users,
 - order replacement gear where required,
 - get that gear installed, tested and debugged
 - get systems and applications updated for IPv6, etc.
- **At some point everyone else will also begin to do this, and there isn't a huge amount of excess capacity to handle a sudden IPv6 surge. If you wait too long, you may find you're out of luck.⁷**

Many of You *Have* Already Begun (Thank You!)

- The Joint Techs crowd tends to have many technological leaders, so it isn't surprising that many of you already have IPv6 deployed in production use on your campuses or RONS/Gigapops, while others are at least be actively planning how you're going to do so.
- A copy of the Internet2 IPv6 deployment map from earlier this year can be seen on the next slide
- I do have ongoing concern, however, that IPv6 deployment at the **campus level** (and among **sponsored participants/SEGPs**), is not where it needs to be. It is critical that native IPv6 be supported by the RONS/Gigapops, but **native IPv6 *REALLY* needs to be supported end to end, all the way to servers and desktop.**
- I like to keep an eye on the IPv6 status survey that's available at http://www.mrp.net/IPv6_Survey.html -- **how does your site look on that summary table? I'm still seeing a LOT of red boxes...**
- Site of interest not listed? Try www.mrp.net/cgi-bin/ipv6-status.cgi

Internet2 Network

IPv6 Deployment
4 August 2008



IPv6 Is Neither A Magic Bullet, Nor A Poison Pill

- Some sites may be stalled wondering, “Well, if we **do** begin to deploy IPv6, will it help or hurt us when it comes to **security?**”
- As we’ll discuss today, deploying IPv6 is neither a magic bullet nor a poison pill when it comes to your site’s security. It may help in some areas, and it may make things harder in others, but it doesn’t really matter if it helps or hurts because in the final analysis, **you still need to bear down and get IPv6 deployed, and so do your downstream sites! Talk to folks at your campuses and sponsored participants/SEGP sites!**
- As you do, please don’t let them try to use “security, SECURITY!” as a reason for **not** deploying IPv6!
- At the same time, remain highly skeptical of any snakeoil claims you may hear that IPv6 will magically *improve* your network’s security (because I don’t think it will do that, either)
- Let’s look at some of the arguments you’ll hear explaining why₁₀ deploying IPv6 will somehow make you more (or less) secure.

**2. Myth: IPv6 Improves Security Because
“All IPv6 Traffic Gets Encrypted With IPSec”**

IPv6 and IPsec

- IPsec is not new with IPv6; in fact, IPsec dates to the early 1990's.
- What's different when it comes to IPv6 is that support for IPsec was made "mandatory" for IPv6 (see for example "Security Architecture for IP," RFC4301, December 2005 at section 10, and "IPv6 Node Requirements," RFC4294, April 2006 at section 8.)
- **If actually used**, IPsec would have the potential to provide:
 - authentication
 - confidentiality
 - integrity, and
 - replay protection
- All great and wonderful security objectives -- **IF** IPsec were used.
- **Unfortunately, as we'll show you, what many had expected to be the cornerstone of the Internet's security architecture has proven in fact to be widely non-used.**

How Might IPsec Be Used?

- IPsec can be used to authenticate (using AH (the Authentication Header), RFC4302), or it can encrypt and (optionally) authenticate (using ESP (the Encapsulating Security Protocol), RFC4303)
- IPsec can be deployed in three architectures:
 - gateway to gateway (e.g., securing a network segment from one router to another)
 - node to node (e.g., securing a connection end-to-end, from one host to another)
 - node to gateway (for example, using IPsec to secure a VPN connecting from a mobile device to a VPN concentrator)
- IPsec has two main encrypting modes:
 - tunnel mode (encrypting both payload and headers)
 - transport mode (encrypting just the payload)
- IPsec also supports a variety of encryption algorithms (including “null” and md5 (yech)), and a variety of key exchange mechanisms
- These alternatives obviously provides tremendous flexibility, but that flexibility also brings along a lot of complexity.


But IPsec Isn't Getting Much Use

- Raw IPsec traffic (AH+ESP, protocols 50 & 51) isn't seen much on the commercial IPv4 Internet.
- For example, Jose Nazario of Arbor Networks estimates IPsec traffic at 0.9% of octets (statistic courtesy the ATLAS project).
- CAIDA (thanks kc!) also has passive monitoring data available; see <http://www.caida.org/data/passive/monitors/equinix-chicago.xml>


You can see the protocol distribution from a couple of CAIDA's monitors for one recent day on the next couple of slides.

IPsec ESP traffic is conveniently colored red;
IPsec AH traffic is too low to be seen on the graph.

Protocol Distribution From One of CAIDA's Passive Monitors




And The CAIDA Distribution Seen From Another Monitored Link




IPv4 IPsec Traffic on Internet2

- Raw IPv4 IPsec traffic is quite rare on Internet2 as well, usually running well under <1% of octets (see table 7, <http://netflow.internet2.edu/weekly/20090112/>).
- Raw IPv4 IPsec traffic has been (gradually) growing, however. See Internet2 IPv4 IPsec ESP traffic levels as a percent of all octets over time by way of example:




IPv4 IPsec Traffic May Be Associated With A Limited Number of Users/Systems (IPsec VPNs?)

- Protocols that are used by a **small** number of users or systems tend to exhibit “spikey” or rapidly varying patterns while protocols which are in ubiquitous use tend to “average out” or “smooth out.”
- The appearance of these new (thanks Brent!) firewall-based graphs (see <http://vixen.gnoc.iu.edu/jfirewall-viz/index-bits.html>) are consistent with IPsec traffic from a few users (IPsec VPN traffic?)



V4-AH (LOSA)



V4-ESP (LOSA)

IPv6 Traffic Visibility on The Backbone

- Ideally, for production IPv6 traffic, one would want **full IPv6 SNMP support** and **full IPv6 Netflow (V9) support**. Regretably, native IPv6 SNMP support and IPv6 V9 Netflow support remains elusive. That's increasingly unfortunate for IPv6 as a production protocol that is, or should be, on par with IPv4.
- One way to improve IPv6 visibility on the Internet2 backbone, at RONS/Gigapops, and on our campuses would be to deploy at least a limited number of dedicated, IPv6-aware, passive measurement appliances. For instance, one IPv6 working group participant expressed pleasure on the mailing list about IPv6 support available from InMon Corporation's Traffic Sentinel product (e.g., see <http://www.inmon.com/products/trafficsentinel.php>)
- In the mean time, we at least have the ability to get some sense of traffic volumes (packets and octets) from logging firewall ACLs, see <http://dc-snmp.wcc.grnoc.iu.edu/i2net/>

Backbone IPv6 IPsec Traffic on Internet2

- Brent Sweeny was also kind enough to send me raw IPv6 interface firewall byte statistics for some interfaces, so I didn't have to try to interpolate values from the graphical displays. Looking at the statistics for one of the interfaces he sent me:

Protocol	IPv6 Interface "X" (Input)	IPv6 Interface "X" (Output)
TCP	309451387384 (~70.1%)	129179860235 (~43.2%)
UDP	130213671989 (~29.5%)	168081323562 (~56.1%)
Other	1786083367 (~0.4%)	2119234653 (~0.7%)
	-----	-----
Total	441451142740	299380418450

- If we assume that all IPv6 IPsec protocol traffic lands in the "other" bucket, at least in the case of this example interface, at this point in time, we can bound the amount of IPv6 IPsec traffic as being no more than 0.7% This is consistent with the commercial Internet IPv4 IPsec traffic levels which we mentioned earlier.
- On the other hand, w/o DPI we can't rule out the possibility that some of that large amount of UDP traffic MAY be **UDP encapsulated IPsec ESP traffic** [e.g., see "UDP Encapsulation of IPsec ESP Packets," RFC3948, from Jan 2005] Then again, this could just as easily be something else, like IPv6 Bittorrent traffic (for more on IPv6-ified Bittorrent see <http://www.sixxs.net/tools/tracker/>).²⁰

Why *Aren't* We Seeing More IPSec Traffic?

- Sites may not be deploying IPsec because IPsec (like many crypto-based security solutions) has developed a reputation as:
 - not completely baked/still too-much under development
 - too complex
 - hard to deploy at significant scale
 - less than perfectly interoperable
 - firewall issues
 - potentially causing a performance hit (crypto overhead issues)
 - congestion insensitive (UDP encapsulated IPsec traffic)
 - something which should be handled as an end-to-end matter by interested system admins (from a network engineer perspective)
 - something to be handled at the transport layer router-to-router (from an overworked system administrator's perspective)
 - duplicative of protection provided at the application layer (e.g., encryption is already being done using ssh or ssl)
 - complicating maintaining/debugging the network, etc., etc., etc.
- Regardless of whether those perceptions are correct (some may be, some may **not** be), IPsec adoption hasn't happened much to date.

But That's All Moot Relative to The Key Point...

- **It would be foolhardy to expect IPsec to provide any material improvement to your site's security since the vast majority of your aggregate traffic (including virtually all your IPv4 traffic) will NOT be IPsec secured.**
- On the other hand, the “good news” is that a lack of IPsec usage in the IPv6 world is substantively no worse than a lack of IPsec usage in the IPv4 world.
- Let's look at another potential security issue.

3. Another Myth: “If We Don’t Deploy Native IPv6, We’ll Be Able to Control Whether Our Users Are Able to Get At IPv6-Served Content”

Even If Your Site “Officially” Foregoes IPv6, Your Users May Decide to “Informally” Try It...

- Some sites which rely heavily on firewalls and perimeter security may decide to forego or postpone deployment of native IPv6. Having made the decision to do so, folks may emit a big relieved sigh, believing that by “sitting this dance out,” they will have foreclosed any possibility of user access to IPv6-only resources.
- Unless that policy is **very** carefully enforced on a technical basis, you may be in for a surprise or two because users may be able to easily work their way around your non-implementation or filters.
- This is particularly important if you’re relying primarily on perimeter filtering to control either the **infiltration** of malware (or other unwanted content, e.g., “adult entertainment” concerns in K12 school environments, c.f. <http://www.ipv6experiment.com/>), or the **exfiltration** of site-sensitive information (as at some federal sites).
- BTW, a very cool IPv6 web hack is sixxs.org’s IPv6 web gateway: try www.cnn.com.sixxs.org (for example), from an IPv6-ified box

My Point? Your Users Will Be Fulfilled

- It is natural and entirely appropriate that your users will want to try new things, such as things they may hear about from their friends and colleagues. One of those things may be IPv6.
- If a technology they're interested in (such as IPv6) isn't one that you're currently supporting, they may search for and find "ad hoc" approaches which they can try without "having to bother you."
- Sometimes there's a hope that obscurity or technical difficulty will keep users from trying some work-arounds, but I wouldn't count on "security through obscurity" in the case of IPv6.
- For example, if a user is on a Mac at a "non-IPv6 site" and that site also doesn't have a perimeter or interior firewall, one option would be for him to enable "6to4." How hard would that be?
- As another example, assume a user is behind a firewall and is using a PC running Windows XP at a "non-IPv6 site." How hard would it be for her to enable Teredo as a way to get IPv6 access?

Enabling 6to4 on a Mac

- ***N.B.:*** 6to4 (RFC3056) usually **won't work** behind a firewall
- -- Apple Menu ==> System Preferences ==> Network ==> Show: Network Port Configuration
- -- If no 6 to 4 port already exists, click “New”
- -- Select 6 to 4 for the port from the pull down list of ports
- -- Enter “6 to 4” for the port's name
- -- Click OK
- -- Make sure “6 to 4” is checked as “On”
- -- Click “Apply Now”

[the above details may vary on some versions of OS X]

- If you're using Firefox 2.x on a Mac, you may also need to tell Firefox to allow IPv6 DNS resolution to occur
- -- In Firefox go to the URL `about:config`
- -- Filter on the string IPv6
- -- Set `network.dns.disableIPv6` to be false
- -- Try going to `http://ipv6.google.com/` (the logo should dance)
- To disable 6to4, use System Preferences to set 6to4 to be “Off”

Enabling Teredo on a Windows XP SP2 PC

Teredo (RFC4380) **will work** even behind a firewall/NAT box (unless the firewall blocks outgoing IPv4 traffic on 3544/UDP).

To set up IPv6 and Teredo on a Windows XP SP2 system, do:

Start ==> Accessories ==> Command Prompt

```
netsh interface ipv6 install
```

```
netsh interface ipv6 set teredo client
```

In Firefox 3.x, try going to <http://ipv6.google.com>

You should see the Google logo dance

If you're running something other than Windows and you're behind a firewall and you want a Teredo-like solution, check out

Miredo (<http://www.remlab.net/miredo/>)

Neither of Those Tasks Were Very Tough

- In my opinion, **pretty much any “reasonably motivated” semi technical user will be able to successfully enable 6to4 or Teredo on their desktop or laptop**, even if they don’t fully understand the technology or the implications of having done so.
- And even if you block 6to4 and Teredo, users can still use RFC3053 IPv6 tunnel brokers (there’s a nice list of them at http://en.wikipedia.org/wiki/List_of_IPv6_tunnel_brokers for example), and so on and so forth.
- On the other hand (and for interesting reasons), there is still no IPv6 version of Tor (the onion routing protocol) yet, see <http://wiki.noreply.org/noreply/TheOnionRouter/TorFAQ> at section 8.13
- Anyhow, rather than playing “IPv6 cat and mouse” with your users, why not just buckle down and run native IPv6 instead? Trying to fight transition mode IPv6 traffic will ultimately become really trickier and trickier over time, particularly if users encrypt.

6to4 & Teredo May Rely on “Remote Resources”

- In addition to things like 6to4 and Teredo traffic posing surprises for things like border filtering and traffic monitoring, tunneled traffic may also rely on comparatively **remote resources**.
- Depending on how far away some of those resources may be, the additional **latency** associated with reaching those gateways may impact the performance of untuned network connections.
- Remote resources may also be a sign that there’s only a **limited pool** of available gateways (if the pool was large and well distributed, presumably you’d be using a nearby gateway rather than a remote one). When the pool of available resources is constrained, it may eventually get “**loved to death**” (overloaded).
- One could also imagine a site run by a cyber criminal, kindly offering free gateway services in an effort to attract your customer’s traffic for surreptitious **MITM**-ish monitoring.
- Services such as 6to4 and Teredo which do not require any sort of registration or authentication may also end up being **abused** by bad guys just as **open SMTP relays** once were.

Magic Addresses

- 6to4 uses 192.88.99.1 as a magic address, anycast via the magic prefix 192.88.99.0/24 (see RFC3068 at 2.3 and 2.4)
- Do you know where your 192.88.99.1 traffic is going? (simple test: traceroute to 192.88.99.1 from a machine at your home site) [Maybe you want to *routinely* monitor the path to 192.88.99.1?]
- When I looked at some examples from public traceroute servers, (examples which I'll omit here), I've seen:
 - large academic sites whose customers may end up using anycast 6to4 relays located clear across the country,
 - government mission networks whose customers may rely on 6to4 anycast relays hosted on the campus of academic sites
 - commercial providers whose customers may rely on anycast 6to4 relays hosted by some of their competitors.
- Or consider Teredo -- Teredo relies on Teredo servers and Teredo relays. Do you know which ones *your* folks may be using?
<http://technet.microsoft.com/en-us/library/cc722030.aspx>
mentions the Teredo *server* `teredo.ipv6.microsoft.com`

But What About Teredo *Relays*, Where the Bandwidth Intensive “Heavy Lifting” Happens?

IPv6 for Microsoft Windows: Frequently Asked Questions

http://www.microsoft.com/technet/network/ipv6/ipv6faq.aspx

Q. Why can't I reach locations on the IPv6 Internet using the Teredo client in Windows?

A. To reach the IPv6 Internet from behind a network address translator (NAT) using the Teredo client included with Windows, there must be an operational Teredo relay attached to the IPv4 and IPv6 Internets. At this time, Microsoft is not providing any operational Teredo relays for reachability to locations on the IPv6 Internet for Windows-based Teredo clients.

Sites Which Are Advertising 2001:0::/32

- RFC4380 at 2.6 specifies 2001:0::/32 for the Teredo relay service. Martin Levy recently presented “IPv6 Traffic Levels on Hurricane Electric’s Backbone,” (see www.nanog.org/meetings/nanog45/presentations/Tuesday/Levy_traffic_level_hurricane_N45.pdf):
“[Teredo] traffic is all eastward across the Atlantic
Flows toward teredo.bit.nl AS12859 via AMS-IX
2001::/32 announce by other networks including
AS12637 Seeweb, AS1257 Tele2, etc.” [emphasis added]
- If you telnet to one of the IPv6 aware routeviews.org nodes (such as route-views.linx.routeviews.org), you can see sites advertising 2001:0::/32 by using the command “show ipv6 bgp 2001:0::/32”
- When I check the Routeviews IPv6 nodes, I’m seeing 2001:0::/32 from AS1257 (Tele2), AS1741 (FUNet), AS7019 (NTT America), AS12637 (Seeweb), AS12859 (Bit.NL) and AS21155 (ProServe).
- If you are globally advertising 2001:0::/32, but for some reason your ASN isn’t listed here, I’d love to hear from you.

‘So Are You Telling Me That I Should Try To “Break” or “Disable” 6to4 and/or Teredo?’

- Encountering 6to4 or Teredo is like encountering extra-terrestrial intelligence. Squelch any immediate reptilian instinct to smash/kill/eat anything which is new/different/potentially threatening. :-)
- At the same time, let’s avoid philosophically overanalyzing this. We should not let “the perfect” get in the way of the “adequate.” While I **really** want to see native IPv6 deployed end-to-end, 6to4 or Teredo (at least as long as it works and isn’t being abused), is better for many users than no IPv6 service at all.
- **Thus, notwithstanding some of the issues mentioned on previous slides, please refrain from breaking 6to4 or Teredo.**
- **You should consider fielding a carefully monitored version of those services, accessible only by your local users, thereby soaking up the local demand for those services (and if you do see folks using ’em, nudge them toward native IPv6 instead)**

**4. Myth: “Wide Area Native IPv6 Connectivity
Is Just Like Wide Area IPv4 Connectivity”**

Remember, “Security” Includes “Availability”...

- Is IPv6 architected as robustly as production IPv4 services?
- In IPv4, the “standard of care” for provisioning high availability wide area Internet connectivity is multihoming. Sites which are multihomed buy upstream transit connectivity from multiple providers (and/or peer with other networks), announcing their own “provider independent” (PI) address space via BGP. This approach works well. By multihoming, if one upstream provider has an outage, experiences “peering wars,” imposes outrageous terms and conditions or has other issues, the customer’s other provider(s) can “pick up the slack.”
- When IPv6 was being designed, however, great attention was paid to the problem of growth in the size of the global Internet routing table. Therefore, in architecting IPv6, significant emphasis was placed on hierarchically assigning IPv6 addresses so that providers could announce just a single (yes, just one!) aggregated prefix rather than hundreds or thousands of customer routes.

Multiple IPv6 Addresses Per Host

- In that idealized hierarchical IPv6 address assignment model, address space which is obtained from one IPv6 upstream provider can't also be announced by other upstream providers.
- So what was the IPv6 solution to this issue? Simple: if a site wanted to multihome using multiple IPv6 providers, assign multiple IPv6 addresses per host, one for each upstream provider.
- The “tricky bit” <cough> is getting outbound traffic written with the “right” IPv6 address chosen from a slate of several possibilities, and handling things like rapidly responding to link failures (and other topology changes). See, for example, <http://www.shim6.org/> Fortunately the IPv6 routing table is still small, so we still have some slack, and work on scalable IPv6 multihoming can continue.
- In the mean time, many sites have transferred the classic IPv4 PI multihoming approach over to IPv6, obtaining and announcing their own PI IPv6 space across multiple IPv6 transit providers and/or IPv6 peering points.

Internet2 Participants With PI IPv6 Allocations

To see these, at <http://routerproxy.grnoc.iu.edu/internet2/> select a node and then do:
show route table inet6.0 community 11537:950 terse (you're looking for non 2001:468 prefixes)
[BGP community info is at www.abilene.iu.edu/i2network/maps--documentation/cookbooks.html]

2001:4d0:9c00::/40	NASA
2001:4e0::/32	WiscNet
2001:5e8::/32	Pittsburgh Supercomputing Center
2001:1458::/32	CERN
2001:1860::/34	Pacific Northwest Gigapop
2001:1860:c000::/34	Pacific Northwest Gigapop
2001:18e8::/32	Indiana U
2001:1948::/32	Utah Education Network
2001:4898::/32	Microsoft
2001:48a8::/32	Merit
2001:48d0::/32	San Diego Supercomputer Center
2001:4930::/32	State of North Dakota ITD
2001:49d0::/32	Kansas Research and Education Network
2001:49d8:40::/42	Commonwealth of PA - OA / Integrated Network Management Services
2002::/16	6to4
2607:f010::/32	UCLA
2607:f140::/32	Berkeley
2607:f290::/32	UC Riverside
2607:f320::/32	U Nebraska-Lincoln
2607:f378::/32	UC Santa Barbara

I2 Participants With PI IPv6 Allocations (cont.)

2607:f388::/32	U Wisconsin Madison
2607:f390::/32	Louisiana Board of Regents/Louisiana Optical Network Initiative
2607:f3b0::/32	NJEDge.Net, Inc.
2607:f470::/32	U Pennsylvania
2607:f600::/32	NYU
2610:8::/32	Penn State
2610:20:8000::/35	US Dept of Commerce
2610:28::/32	NCREN
2610:48::/32	U Maine System
2610:58::/32	Boston U
2610:a8::/32	OARnet
2610:d0::/32	Cisco
2610:e0::/32	U Missouri - dba the Missouri Research and Education Network
2610:130::/32	Iowa Communications Network
2610:148::/32	Georgia Tech
2610:1e0::/32	Kentucky Educational Computing Network
2620:0:270::/48	U Texas Health Science Center at Houston
2620:0:bc0::/48	George Mason U
2620:0:c30::/48	U South Florida
2620:0:c80::/48	NCSA
2620:0:df0::/48	Bryant U, RI
2620:0:e50::/48	U Iowa

Internet2 Members with Non-PI IPv6 Address Assignments From Internet2's IPv6 Allocation

- Internet2's IPv6 address block is 2001:0468::/32

You can see documented assignments from within that block via whois. If you have a Linux box or Mac, pop up a terminal window and then enter...

```
% whois -h whois.arin.net \>\ 2001:468::
```

- There's also an HTML page that lists all assignments from within that block: http://ipv6.internet2.edu/Abilene_Allocations.shtml

If Your Site Wanted To Get PI IPv6 Space

- In the ARIN region, the Number Resource Policy Manual describes the minimum requirements which a LIR (e.g., a service provider such as a RON/Gigapop) must meet in order to receive an initial minimum allocation of an IPv6 /32
(see <http://www.arin.net/policy/nrpm.html#six>):

6.5.1.1. Initial allocation criteria

To qualify for an initial allocation of IPv6 address space, an organization must:

1. be an LIR;
2. not be an end site;
3. plan to provide IPv6 connectivity to organizations to which it will assign IPv6 address space, by advertising that connectivity through its single aggregated address allocation; and
4. be an existing, known ISP in the ARIN region or have a plan for making at least 200 end-site assignments to other organizations within 5 years.

Some Networks May Not Need IPv6 Multihoming

- Some networks may only see limited IPv6 traffic volumes to date, or may be treating IPv6 as an experimental service and therefore may decide to forego IPv6 multihoming at least for now.
- In those cases, sites will normally use Internet2-supplied IPv6 address space and rely exclusively on Internet2 for all their IPv6 bandwidth, including their IPv6 commodity transit bandwidth.
- This is somewhat different from the model that is employed by Internet2 sites for their commodity IPv4 transit connectivity.
- I don't want people to infer that I'm saying this alternative IPv6 commodity transit model is an "issue" or a "problem" (because at least for now I **don't** think that it is an issue or a problem), but because this model is different, we should be thinking now about (a) how growth in IPv6 transit requirements will be handled, and (b) whether there are any operational implications to this sort of non-multihomed IPv6 architecture (for example, should it make scheduling of IPv6 maintenance windows more of a "big deal?"⁴¹)

Gratuitous Provision of IPv6 Transit

- Another example of how IPv6 connectivity can be at times less robust than IPv4 can be seen in problems associated with things like the “gratuitous provision of global transit.”
- While offering to route anyone’s IPv6 transit traffic at no charge and without prearrangement may seem like an incredibly generous thing to do, it can cause problems when production IPv6 traffic suddenly follows a “shorter” (BGP) path that flows indirectly via geographically remote parts of the world (or attempts to flow via circuits not provisioned to carry a material fraction of the whole world’s IPv6 transit bandwidth). Fortunately, better BGP filtering has largely reduced or eliminated this issue today.
- A set of IPv6 BGP filters meant to provide a nice start at reducing the number of “problematic” global IPv6 routes is available at <http://www.space.net/~gert/RIPE/ipv6-filters.html>
As always, the more strictly you filter, the more carefully/closely you’ll need to work at keeping your filters updated.

Mitigating DDoS Attacks Against IPv6 Sites

- Another example of a security-related routing issue that may arise in conjunction with IPv6 sites is mitigating distributed denial of service (DDoS) attacks. In the IPv4 world, a common option to avoid having DDoS traffic saturate downstream links is the use of blackhole routes.
- For example, Internet2's IPv4 BGP policy allows connectors to advertise BGP discard routes tagged with the BGP Community 11537:911 and a mask length from /24 to /32, in which case all packets arriving for that route will be discarded by all Internet2 Network routers, before those packets can saturate downstream customer links.
- The Internet2 community should consider whether or not a comparable policy, obviously adjusted for IPv6 address lengths and prefix usage patterns (e.g., perhaps accepting masks from /64 to /128) should be implemented for IPv6 on Internet2.

**5. Some Additional IPv6 Security Topics,
One Slide Per Topic, If We Have Time....**

IPv6 Security Hardware Device Support

- DREN has been pushing very hard to get IPv6 deployed on their network (given a June 2008 deadline for doing so), and Ron Broersma has been doing a great job keeping the community up-to-date (see a list of his briefings at <http://www.v6.dren.net/>)
- Unfortunately, as mentioned in his January 2008 briefing, many security appliances (such as firewalls) still aren't IPv6 clean.
- This may or may not be a problem for your site, depending on:
 - your architecture (e.g., no firewalls? firewall issues obviously aren't going to be very relevant for you)
 - your vendor (some vendors have many IPv6 issues, others none)
 - the nature of your traffic mix (if you have "exotic" traffic, you will likely trigger more corner case bugs than if you're vanilla)
 - you and your site's willingness to be an unpaid beta tester :-)
- Bottom line, there really isn't much of an option except to proceed, working with vendors on the bugs which will crop up

IPv6 DNS Blocklist Support

- Many higher ed sites rely on DNS-based blocklists for things like spam control. Support for IPv6 address listings by blocklist providers still doesn't exist, and even attempting to query DNS block lists for IPv6 addresses may result in undesirable consequences (see, for example, Randy Bush's experiences with an IPv6 enabled server, a copy of Exim and one anti-spam DNSBL at <http://ran.psg.com/~randy/ipv6-westin.html>)
- If you are aware of a block list that **is** now listing IPv6 IP addresses or netblocks, I'd really love to hear about it.
- Until IPv6-aware DNSBLs are available, you may want to handle IPv6 SMTP abuse problems on a case by case basis (since there are relatively few IPv6 providers right now, the IPv6 world is a lot more like the "good old days" when people actually took care of their abuse issues than our current miscreant overrun IPv4 Internet). Eventually, IPv6 SMTP whitelists may be a solution.

Pre-Attack Network Reconnaissance

- It is common for miscreants to remotely scan IPv4 network addresses in an effort to identify active addresses, operating systems in use, open ports, etc., intelligence which may help them plan an attack against you. An increasingly common (if unfortunate) response to that threat has been to insert a firewall between the Internet and local users, thereby deflecting some scans and probes, albeit at the cost of a loss of transparency.
- Because IPv6-connected sites typically have a far larger number of addresses than IPv4-only sites, and end-to-end connectivity was another key objective of IPv6's architecture, some have suggested that it might be harder for attackers to do exhaustive scans of IPv6 sites simply because of the vastly larger number of addresses involved. That's true, as far as it goes, but that's not the whole story. If you haven't seen RFC 5157 ("IPv6 Implications for Network Scanning," March 2008), I'd urge you to look it over.

IPv6 and Privacy

- Privacy is another important issue that came up while IPv6 was being designed, particularly in conjunction with the use of an IPv6 node's ethernet MAC address as part of the generated stateless autoconfigured IPv6 address assigned to that system. Because that MAC address would typically be invariant, it might serve as a basis for persistent tracking and traffic analysis. Solutions for this problem have been proposed: see RFC4941 (“Privacy Extensions for Stateless Address Autoconfiguration in IPv6,” Sep 2007).
- At the same time, each site has a legitimate need to be able to associate network traffic on a particular IP with a particular customer, so that they can deal with abuse-related issues, etc.
- Over time, the availability of DHCPv6 has largely ameliorated this issue, putting it on par with the privacy of dynamically assigned IPv4 addresses (which are more or less untrackable by external sources, but which can be mapped to customers by administrators)

IPv6 Long Term, CGN, A+P, Etc.

- One of my colleagues, Dave Meyer, did an excellent talk for the Jan '09 NANOG: “It’s The End of the World As We Know It (aka “The New Internet Architecture”), see www.nanog.org/meetings/nanog45/presentations/Monday/Meyer_iteotwawki_N45.pdf
Dave’s a very sharp guy and it is well worth your time to read and think about his very provocative perspective on how the IPv6 rollout has gone so far, and where it may be going in the future. One of his conclusions, from slide 14, is that “Carrier Grade NAT (et al) will be deployed” [I wouldn’t assume that DMM likes or dislikes it from that statement, just that he believes it will happen]
- For info on an alternative “A+P” (address plus port) approach see Maennel et. al.’s “A Better Approach Than Carrier-Grade-NAT,” <http://mice.cs.columbia.edu/getTechreport.php?techreportID=560> which mentions, among other things, “CGNs pose a security threat and/or an administrative nightmare...” Read the paper to see why!

Thanks for the Chance to Talk!

- Are there any questions?