This slide is intentionally blank

(The AV guys always keep getting upset with me if I don't put this on my blank slide before my talk. They think their projector is broken.)

IBM Internet Security Systems

The Brave New World of IPv6.

Michael H. Warfield

mhw@linux.vnet.ibm.com

01/20/2011

Outline

Introduction to IPv6

State of IPv6 Deployment

IPv6 Addressing

Transition Mechanisms and Tunnels

System and Application Support

Summary and Conclusion

Introduction

- **IPv6 - the “next generation” Internet protocol**
- **Under development for many many years**
- **In production for many years**
- **Largely ignored in areas rich in IPv4 addresses**
- **IPv6 addresses limitations in IP version 4 (IPv4)**
  - IPv4 addresses are limited to 32 bits
  - Routing tables are taxing routers
  - Networks and subnetworks are ad-hoc
  - Allocations are disorganized
  - Initially no security features on the IP layer with IPv4

IPv6 Overview

- **Expands addresses to 128 bits**
- **Formalized address boundaries**
- **IPSec (backported to IPv4)**
- **Quality of Service (QoS) typing**
- **Stateless as well as stateful autoconfiguration**
- **Provides for dynamic network address renumbering**
- **Rich set of transition tunnels and translators**
- **Robust resistance to brute force scanning**
- **Has no broadcast addresses**

Paradigm Shift

- **Contrary to popular belief - IPv6 is NOT merely IPv4 with fat addresses**
- **IPv4 allocations were a paradigm of scarcity**
  - Use of dense allocations to optimize utilization
- **IPv6 allocations are a paradigm of abundance**
  - Use of sparse allocations to optimize versatility
- **Best practices in IPv4 may not be best practices in IPv6**
- **Best practices for IPv6 may not have been best practices for IPv4**
- **Even if IPv6 were IPv4 with fat addresses (which it's not) it couldn't be because of the paradigm change**

IPv6 Deployment

- **Provider deployment in North America & Australia was relatively slow but has been getting a lot better**
  - A few tunnel brokers
  - Some ISPs provide native support
- **Very Common in Europe**
  - Many native ISP's plus some tunnel brokers
  - RIPE has had more allocations than the rest of the world
- **Widespread adoption in APAC**
  - Many IPv6-only networks
  - At least one IPv6-only ISPs China

Transition Mechanisms

- **Promote IPv6 adoption and interoperability**
- **Mapped addresses aid IPv4 – IPv6 communications**
- **IPv6 can be tunneled over many other protocols**
  - 6in4 / SIT (Six in Tunnel)
  - 6to4 Automatic 6in4 / SIT tunnels
  - 6over4 and multicast
  - IPv6 over UDP (Teredo, TSP, AYIYA, OpenVPN)
- **Proxy Servers, Services, Protocol Bouncers**
- **DSTM and 4in6 (IPv4 tunneled over IPv6)**
- **Translators (NAT-PT / TRT)**

Providers

- **Tunnel brokers provide IPv6 access across IPv4 networks**
  - FreeNet6 <freenet6.net> - North America
  - Hurricane Electric <tunnelbroker.net> - NA, EU, APAC
  - OCCAID <occaid.org> - NA, EU (now partnered with SixXs in US)
  - SixXS <sixxs.net> - EU, NA (with OCCAID)
  - AARNet <broker.aarnet.net.au> - Australia
- **Some providers supply native IPv6 in North America**
- **Comcast is using IPv6 to manage cable devices**
  - End customer delivery is in beta
  - Beta had to be closed due to demand!

Tunnelbroker.net

- **Hurricane Electric – tunnelbroker.net, he.net**
- **One of the earliest US adopters and tunnel brokers**
- **Based in California**
- **Now sports an international IPv6 backbone**
- **POPs all over US, Europe, plus Singapore and Australia**
- **Provides static 6in4 tunnels (manual reconfiguration)**
- **Free /48 and /64 networks**
- **BGP feeds and peering available**
- **Excellent free “Certification” quiz open to anyone!**
- **Extensive IPv6 support forums**

Freenet6

- **Hexago – freenet6.net, hexago.com, gogo6.com**
- **Based in Canada**
- **Provider of major tunnel broker servers for dynamic tunnels**
  - TSP – Tunnel Setup Protocol
  - Good for road warrior tunnels from anywhere
- **Static and dynamic 6in4 tunnels**
- **IPv6 over UDP**
  - Free OpenSource tunnel client for many platforms
  - Best tunnel broker in NA if you are behind a NAT!
- **Free /56 networks**

OCCAID

- **Consortium of IPv6 developers and networks - www.occaid.net**
- **Over 10 countries in EU plus North America**
- **Points of presence throughout the US**
  - Several POPs in Atlanta
- **Static 6in4 tunnels (manual endpoint reconfiguration)**
- **Free /64 and /48 networks**
- **BGP peering available (and may be required for some tunnels)**
- **Now partnered with SixXs to provide end user service**

Some IPv4 Guesstimates

- **IPv4 host addresses – 4 billion total 0.0.0.0 – 255.255.255.255**
- **IPv4 networks (pure guesswork)**
  - If all of IPv4 space were /24 nets - 16 million
  - If all allocated space were /24 nets - ~ 9 million
  - Old estimate of broadband and DSL accounts - 30 million
  - Best wild guess – 8 million to 30 million
    - Do you count grandma's wireless router as a network?
    - Not all IPv4 addresses have NAT networks
- **IPv4 core routes (from BGP) - > 340,000 (varies with view)**
- **Routable IPv4 addresses (1/20/2011) – 2.4 billion**
- **APNIC CIDR Report**

Some IPv6 Statistics (recent)

- **IPv6 advertised networks (non-transition)**
  - Ignore 2001::/32 (Teredo)
  - Ignore 2002::/16 (6to4)
  - Ignore 3FFE::/16 (6 Bone – now defunct)
- **IPv6 core routes – <4000 routes in BGP**
- **IPv6 routable /48 networks – 180 trillion**
- **10 trillion networks in /32 provider blocks (ISP blocks)**

IPv4 Address Exhaustion

- **Only 7 /8 provider blocks left at IANA**
- **Next request will trigger allocation of everything left**
- **Current IANA Exhaustion Estimate: February 2011**
- **Current RIR Exhaustion Estimate: October 2011**
- **UK expects to be out of customer IPv4 addresses in early 2012**
- **Efforts to extend IPv4 has created NAT and now CGN**
  - Carrier Grade NAT aka NAT444
  - NAT at the ISPs
  - Customers will not even have global IPv4 addresses
  - VPNs and some games and protocols will break!

Addresses

- **IPv4**
  - 32 bits - 4 billion addresses
  - 4 8-bit decimal octets, 0-255
    - www.wittsend.com: 130.205.32.64
  - Variable size subnets
- **IPv6**
  - 128 bits - 3.4 * 10^38 addresses
  - 8 16-bit hex fields, 0-FFFF
    - www.ip6.wittsend.com: 2001:4830:3000:2:260:8ff:40ce:7322
  - Fixed subnets (/64), and networks (/48)

TLA / NLA / SLA / EUI

- **TLA: Top Level Aggregator**
  - First 16 bits
- **NLA: Next Level Aggregators (Sub-TLA / NLA)**
  - Second and third 16 bit fields
  - Variable field spliting between RIRs (registries) and ISPs
- **SLA: Site Level Aggregator – IPv6 subnet ID**
  - Fourth 16 bit field
  - Some providers are spliting this field for suballocations
- **EUI: End Unit Identifier – Host identifier**
  - Lower 64 bits
- **tttt:nnnn:nnnn:ssss:eeee:eeee:eeee:eeee**

IPv6 Global Addresses

- **IPv4 Compatible: ::0000:n.n.n.n**
  - IPv6 node to IPv6 node over IPv4 tunnel
  - Now Deprecated
- **IPv4 Mapped: ::FFFF:n.n.n.n**
  - IPv4 node to IPv6 node over IPv4
- **Global Unicast: 2000::/3 (2000: -> 3fff:)**
  - Each /16 has as many IPv6 networks as there are IPv4 addresses
  - V6 Internet: 2001::/16, 2003::/16 – 2fff::/16
  - Teredo: 2001:0::/32
  - 6to4: 2002::/16 (6in4 protocol 41)

IPv6 Local Addresses

- **Unique Local Addresses (ULA): FC00::/7**
  - FC00::/8 – Centrally administered assignments (CULA)
  - FD00::/8 – Locally administered assignments (LULA)
  - 40 bit “random” global id + 16 bit SLA + EUI
  - Should not be propagated between networks / sites
  - Replaces deprecated Site Local (FEC0::/10) addresses
- **Link Local: FE80::/10**
  - Must not be propagated across subnets
  - Not unique within site
  - Multi-homed devices also must specify interface
  - Use for link local discovery and advertisements

IPv6 Multicast Addresses

- **Multicast** **FF00::/8**
  - Interface local: FF01::/16
  - Link local: FF02::/16
  - Site local: FF05::/16
  - Global: FF0E::/16
- **Services:**

All nodes:	FF0[12]::1	Interface(If)/Link
All routers:	FF0[125]::2	If/Link/Site
NTP:	FF0[125E]::101	If/Link/Site/Global
DHCP:	FF0[25]::2:2	Link/Site

- **Never allowed as a source address**

“Standard” allocations

- **IANA / IETF recommendations**
- **Standards? We have lots of standards...**
- **/65 – /128 – P2P or internal peering (special cases)**
- **/64 – Sites with a single subnet**
- **/56 – Non-standard allocations of 256 subnets**
- **/48 – Sites with multiple subnets**
- **/32 – ISP Provider block - “minimum” routing granularity**
- **/23 – /16+ – RIR allocations**

EUI-64

- **EUI is the lower 64 bits of an IPv6 address**
- **EUI-64 Based on the interface MAC address**
- **EUI-64 Remains constant over renumbering**
- **Remains constant across subnets**
- **Potential privacy issues**
- **Potential network mapping issues**
- **::mmMM:MMff:feMM:MMMM (M=Mac address)**
  - Invert one bit
  - Split address in half and insert “fffe”

Privacy Enhanced Addresses

- **Random non-conflicting EUI addresses**
- **EUI changes from boot-up to boot-up**
- **EUI may change over time**
- **Multiple EUIs may be assigned and overlap**
- **Network mapping prevented (sort of)**
- **Node tracking prevented (even site local node tracking)**
- **Troubleshooting and tracing is very difficult**
- **P2P users will love privacy enhanced addresses**

Other EUI Addressing Schemes

- **They're your addresses. Do with them what you will.**
- **Standards like EUI-64 are options**
- **Some are more scannable than others**
- **Addresses can be picked at random**
  - Neighbor discovery detects any extremely rare collisions
- **Addresses can be changed periodically**
- **Can use mixed / different methods on different subnets**
- **Addresses can be assigned by cryptographic formula**
  - Client authentication by EUI check?
  - Filtering on source address by hash code?

Stateless Autoconfiguration

- **Allows for auto configuration of IPv6 addresses**
- **Allows for dynamic renumbering of prefixes**
- **Subnets may have multiple perimeter routers**
  - Different prefixes
  - Different lifetimes
  - Different preferences
  - Route lifetime can be used to set prioritized default routes
- **Interfaces may have multiple global addresses**
- **Rogue routers may inject IPv6 routes on IPv4 nets**
- **Rogue routers may interfere with IPv6 routers**

6in4 / SIT Tunnels

- **6in4 (aka SIT) Transition Tunnels**
  - On *BSD these are referred to as GIF tunnels
- **Simple Internet Transition / Six In Tunnel**
- **IP Protocol 41 (ipv6) IPv6 encapsulated in IPv4**
- **Basis for several IPv6 tunnel schemes**
  - Static SIT tunnels use preconfigured endpoints
  - 6to4 automatic tunnels employ formated v6 addresses
- **Can pass “many” IPv4 NAT devices (proto 41 forwarding)**
- **Many tunnel brokers provide IPv6 through SIT tunnels**
- **Some tunnel brokers adapt to dynamic addresses**

6to4

- **Automatic 6in4 / SIT tunnels**
- **2002::/16 Prefix**
- **Uses TLA/NLA/SLA/EUI scheme**
- **An IPv6 network assigned to each IPv4 address**
- **No tunnel broker required**
- **2002:{IPv4_ADDR}::/48 Network**
- **Gateway IPv4 address is the NLA**
- **Autorouted on IPv4 by the NLA address**
- **192.88.99.1 Anycast Gateway to other TLAs**

Teredo / Shipworm

- **IPv6 over UDP (default - port 3544/udp)**
- **Intended to provide IPv6 tunnels over IPv4 NAT devices**
- **Both endpoints may be NATed and/or firewalled!**
  - Can bypass most firewalls (uses outbound UDP sockets)
  - Uses a robust NAT traversal similar to STUN (RFC 3489)
  - Provides peer-to-peer IPv6 connectivity for clients over NAT devices
- **Clients requires a Teredo server and relay on public IPv4**
- **Miredo project provides Teredo on Linux and FreeBSD**
- **IANA assigned address prefix 2001:0::/32**
- **IETF Standard RFC 4380**

OpenVPN

- **VPN over UDP (assigned port 1194/udp – old 5000/udp)**
- **Popular VPN that works over NAT and through firewalls**
- **Was used by the "Join Project" tunnel brokers in Germany**
  - Join Project was disbanded with the availability of native IPv6
- **SSL/TLS authentication**
- **ESPinUDP (IPSec- NAT-T encapsulation)**
- **Direct VPN of IPv6 over IPv4**
  - Currently in peer-to-peer mode only
  - Multi-client server mode expected soon
- **Clients for Windows, Linux, *BSD, Mac OS/X, Solaris**

IPv6 Only Networks

- **IPv6-only networks are possible and are even deployed**
- **DNS munging handled by totd**
  - "Trick or Treat Daemon"
  - Translates DNS A records into AAAA records
  - Early 2004 showed over 600 instances of totd as DNS servers
    - Each server represents at least one IPv6 only network
    - Represented over 2000 domains
- **Some older MS protocols still require some private IPv4**
- **IPv6 to IPv4 handled by proxies and translators**
- **Can mix IPv4 private with IPv6 global**
- **Registrars must provide IPv6 glue in nameserver records**

Microsoft Windows Support

- **Windows Vista & Windows 7 – Got it – Cannot disable it**
- **Windows XP - Native support – Got it – Just turn it on**
  - –No need to reboot after installing (most of the time)
- **Windows 2003/2008 Server - Native support**
- **Windows 2000 (SP1 and above) - Patch from MS**
- **Windows NT - 3rd party patches**
- **Windows 95 & 98 - 3rd party support**

Unix / Linux Support

- **Linux (most modern distributions)**
  - All kernels since 2.1.8
  - Firewall support for IPv6 in 2.4
  - Fedora Core 2 Enabled IPv6 BY DEFAULT (by accident)
  - Fedora 8 and above – very difficult to disable (by intent)
  - Major recent Linux distros certified for OMB IPv6 compliance
- **Unix**
  - FreeBSD / OpenBSD / NetBSD
  - Solaris / Solaris x86 version 8 and higher
  - AIX 4.3 and up
  - HP/UX 11i and up

Other Systems and Devices

- **Apple - Mac OS X**
  - Enabled by default
- **Airport Extreme Wi-Fi Basestations**
- **Linux based Wi-Fi Basestations & DD-WRT firmware**
- **Novell - Netware 6**
- **Routers**
- **Cell Phones**
  - Android on WiFi!

The Google Survey

- **Goggle conducted study to test for IPv6 clients**
- **“Enrolled” small fraction of visitors to www.google.com**
- **.238% of clients would use IPv6 when offered**
- **Half of IPv6 clients are Mac**
- **Country rankings:**
  - Russia 0.76%
  - France 0.65%
  - Ukaine 0.64%
  - Norway 0.49%
  - United States 0.45%

DNS

- **Domain Naming Service / Bind**
- **IPv4 has “A” records**
- **IPv6 has “AAAA” records**
- **Hosts may have mix of A and AAAA records**
- **IPv4 uses reversed octets for reverse lookups**
- **IPv6 uses reversed hex nibbles**
- **6to4 (2002::) reverse lookups are not available**

```
# host alcove.wittsend.com
alcove.wittsend.com has address 130.205.12.10
# host 130.205.12.10
10.12.205.130.in-addr.arpa domain name pointer alcove.wittsend.com.
# host -t AAAA www.ip6.wittsend.com
www.ip6.wittsend.com has AAAA address 2001:4830:3000:2:204:8ff:fe00:1151
# host 2001:4830:3000:2:204:8ff:fe00:1151
1.5.1.1.0.0.e.f.f.f.8.0.4.0.2.0.2.0.0.0.0.0.0.3.0.3.8.4.1.0.0.2.ip6.arpa domain name pointer www.ip6.wittsend.com.
```

IPv6 Now!

- **IPv6 is in active production and utilization right now!**
- **Many root name servers have IPv6 addresses**
  - Now published in the root zone.
- **Registrars are now supporting IPv6 nameserver glue records**
- **Regular IPv6 DNS server-to-server traffic even for IPv4 queries**
- **Regular E-Mail delivery over IPv6**
- **Regular traffic to IPv6 web servers over IPv6**
- **Lots of Linux repositories on IPv6**
- **IPv6 enabled bittorent clients are in service and IPv6 bittorent servers and seeders are deployed**

IPv6 Everywhere!

- **IPv6 is available anywhere IPv4 is**
  - Native
  - Tunneled directly over IPv4
  - Tunneled over UDP (over NAT) or other VPNs
- **Vacations, resorts, hotels**
- **Conferences all over the world**
- **5 Cruise ships at sea**
- **1 Aircraft in flight**
- **Over cell phones and MiFi tethers**

IPv6 Today!

- **Get on IPv6 tonight!**
- **Autotunnels**
  - Turn on 6to4 on almost anything
  - Teredo on Windows or Miredo on Linux and BSD
  - Airport Extreme, Linux routers, DD-WRT firmware
- **Brokers**
  - Freenet6 (NAT & UDP)
  - Hurricane Electric / Tunnelbroker.net
    - Get your certification!
  - OCCAID / SixXs
    - Local to Atlanta (low latency)

Providing IPv6

- **To provide IPv6 to a network, you must support it**
- **Tunnels should be terminated security perimeters (firewalls)**
- **6to4/6in4 should be prohibited within a corporate network**
- **Native IPv6 should be provided within the corporate network**
- **Router advertisements should be monitored for anomalies**
- **Prefixes should be monitored for expected changes**
- **Unusual router advertisements should be investigated**
- **IDS systems should detect rogue routers and prefixes**
- **Avoid trivial EUI addresses where and when possible**

Avoiding IPv6

- **To avoid having IPv6 on a network, you must support it**
- **Tunneling protocols and transports should be blocked**
  - At all security perimeters
  - At routers and subnet boundaries
  - Across all VPNs
- **IDS / IPS systems should monitor for IPv6 protocols**
  - Neighbor discovery
  - Router advertisements
  - NIDS systems should detect IPv6 – native and tunneled
  - Host systems should be monitored for IPv6
- **Most new systems will have IPv6 enabled by default!**

Ignoring IPv6

- **If you don't provide or prevent IPv6, you will have IPv6**
  - You won't control it
  - You won't recognize it
  - You won't be managing it
  - It will still be globally addressible
  - It will still be fully routable (independent of IPv4 routing)
  - Others will be providing IPv6 routes and routers, not you
- **Others providing IPv6 will not have your best interest at heart**
  - Users bypassing restrictions
  - Intruders securing backdoors
- **You might even be broken, now!**

Welcome to the Brave New World of IPv6

- **Ready or not, here it comes^H^H^H^H^H is!**
- **IPv6 is supported on most common platforms**
- **IPv6 can be used over most existing networks**
- **IPv6 is easy to set up**
- **IPv6 is easier and cheaper to provide than prevent**
- **IPv6 is ready for you**
- **Are you prepared for IPv6?**

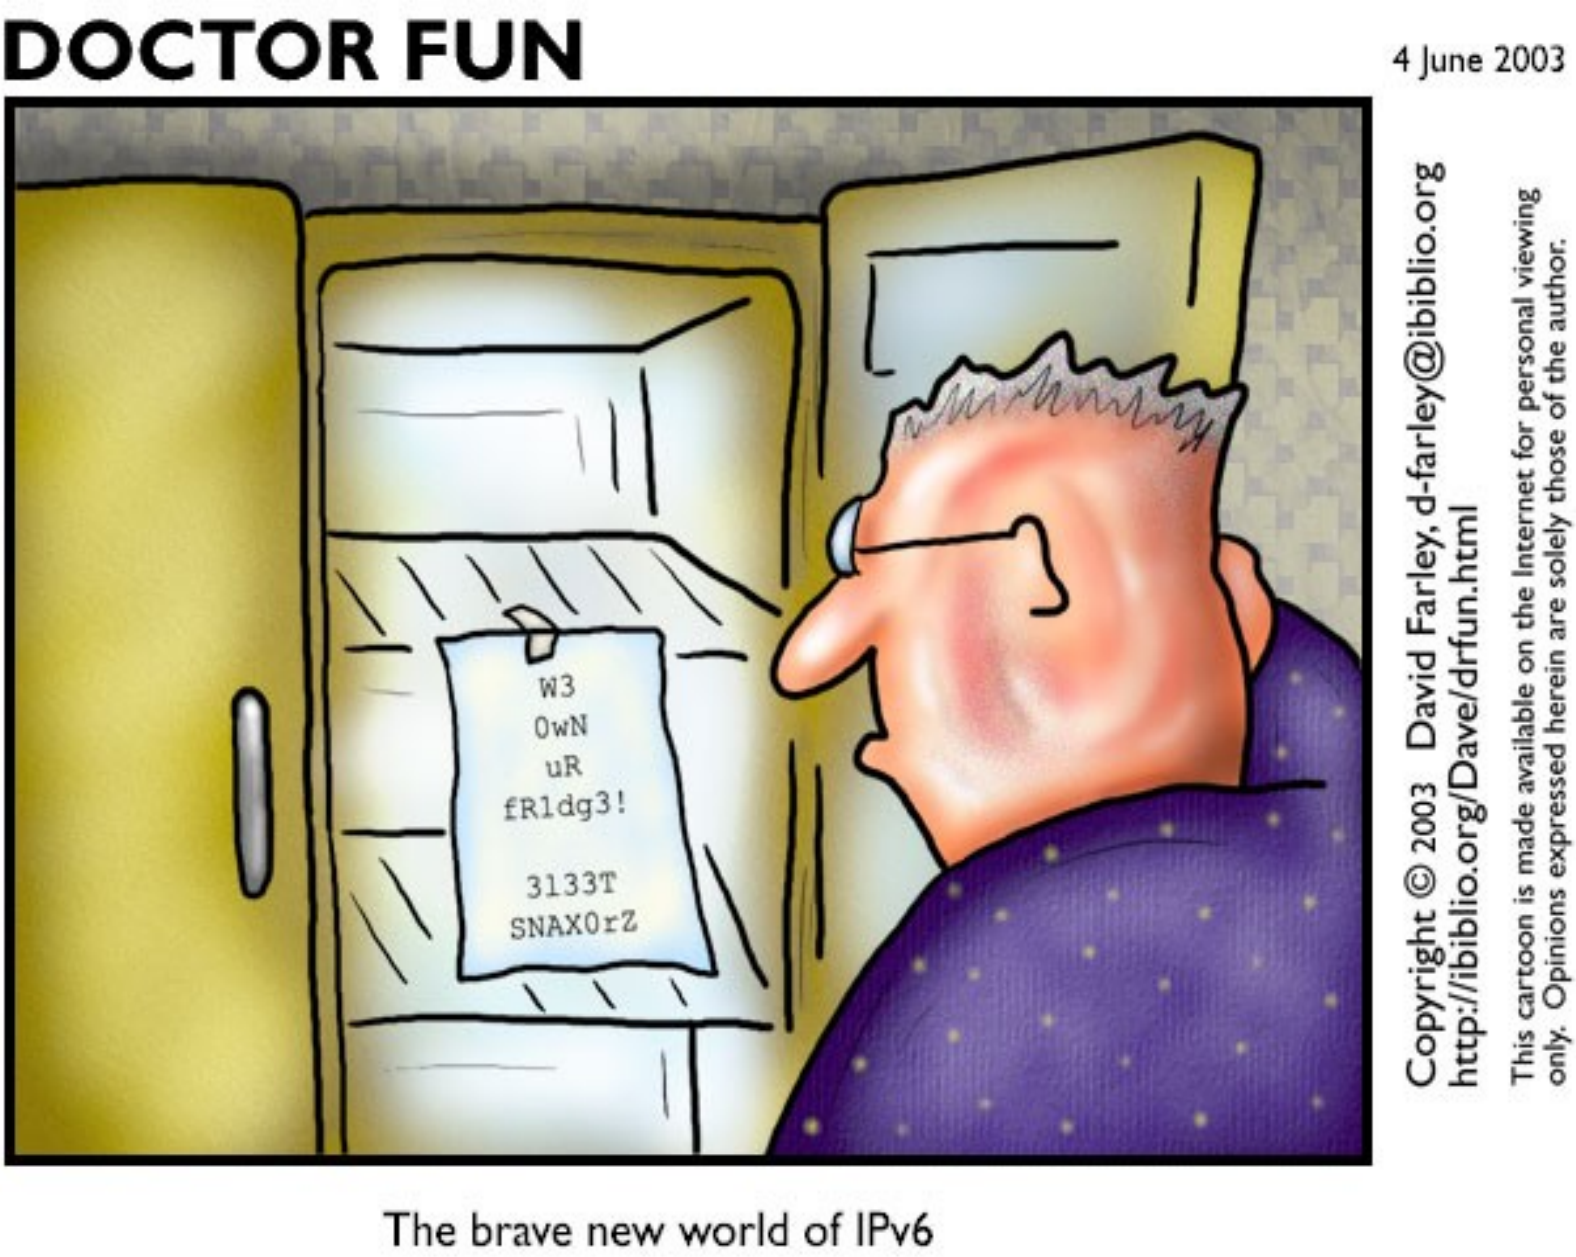

And he didn't even know it was IPv6 enabled...

Thank you very much!

IPv6 Resources

- **<http://www.ipv6style.jp>** **(English)**
- **<http://www.ipv6.org>**
- **<http://www.6bone.net>**
- **<http://www.nav6tf.org>**
- **<http://www.linux-ipv6.org>**
- **<http://www.sixxs.net>**
- **<http://www.tunnelbroker.net>**
- **<http://www.gogo6.com>**
- **<http://www.occaid.org>**

IBM Internet Security Systems

The Brave New World of IPv6.

Michael H. Warfield

mhw@linux.vnet.ibm.com

mhw@WittsEnd.com

01/20/2011