Rapid IPv6 Deployment

Skeeve Stevens, CEO of eintellego Pty Ltd

What is stopping you from implementing IPv6?

What you need to get past, before you can rapidly deploy

- ✓ It is too expensive to implement
- ✓ It is just too hard to implement right now
- ✓ I don't know where to start
- ✓ My customers aren't asking for it
- ✓ So much doesn't support it
- ✓ What is IPv6?
- ✓ Who can I turn to for help?

eintellego

Where to start with IPv6

IPv6 isn't hard – but it is big... There are a lot of OS systems, network devices, operating systems, and other devices to look into. You do need to plan your migration properly

- Start where you can't break things (hopefully!)
  - In the lab (bit boring, but safe)
  - External Co-location Linux Box (US is easy and cheap. <US$100 for dual stack server with 1TB+ traffic. HE apparently do the same, but no pricing on website or easy sign-up that I could find)
- Prepare a IPv6 readiness report - allocate a small amount of time – few hours a week to analyse where IPv6 will most impact your network.

eintellego

Where to start with IPv6… (continued)

- Build a lab – You don't need much to test BGP, OSPFv3, interfaces and so on – Dynamips (Cisco) or Olive (Juniper) are great simulators to replicate your core network
- Start at the border – bring BGP to your edge and then pause and reflect (don't forget security – mentioned later)
- If you are large company with multiple departments (Network Build, NetOps, ServerOps, Security, etc), you need to work together to make this happen. Help each other get the job done – GET OVER YOUR POLITICS. Present a unified front and co-ordinated proposition to your management.

Scenario – The Hosting Provider

Provider of Colocation, Cloud Services, Dedicated and Shared Hosting

For existing members, IPv6 is easy.

For new members, plan for $4000ex in setup

Cisco switches need rebooting to enable IPv6.

Outages need to be planned and executed

VLAN's are used to bypass legacy equipment which doesnt support IPv6 – i.e. Cisco 3550's

Initially, bringing up IPv6 BGP on Pipe Peering was safest – with no DNS using it yet

Start issuing IPv6 addresses to end customers

eintellego

Rapid IPv6 Deployment

How rapid is rapid?

If you want to do it quickly…. You need to be prepared

- APNIC members automatically get IPv6 allocation (Kickstart)
- Transit Provider with Dual-stack Transit – if not, tunnel (HE)
- Cisco/Juniper/Vyatta/Quagga edge with BGPv4
- Cisco/Juniper/HP/Brocade Switching Infrastructure
- Engineer familiarity with vendor hardware above & BGP/OSPF
- Transit/Peers/MPLA IX providers have allowed announcements (1-4 days if not)

Rapid IPv6 Deployment continued…..

- Build test boxes
  - Most ISPs use Linux of some kind
  - A windows 2008 box is also useful
  - Details provided later
- Ethernet based end-user IPv6 assignments
  - Colocation, VMs, MetroE, Wireless, Hosting
- Access Technology – xDSL (L2TP) design Discussion -
  - The hard part if you haven't been working with IPv6 long.

eintellego

Rapid IPv6 Deployment continued…..

Conduct an IPv6 readiness assessment:

- Network Infrastructure – Routers and Switches
- Servers & PCs (i.e. operating systems)
- Network Devices – Appliances, KVM, OoB and so on
- Network management tools (HP, Cisco, etc)
- Security – everywhere you have it now – needs to be replicated
- Applications – dealing with IPv6 addresses
- OSS systems – Billing, Accounting, Radius, etc
- In-house skills

eintellego

Simplified Addressing

2001:db8::F:127:0:0:1

We have developed a strategy which helps network and server admins be able to understand & deploy IPv6 rapidly while not requiring a huge time investment in training.

This strategy uses the network's existing IPv4 addressing topology so that the IPv6 addresses can be instantly recognised and built upon over time.

This strategy is all about the rapid deployment of IPv6, getting it into the network and being used day to day.

Our experience with our customers shows us that this is the quickest and most rapid method of building and educating 'resource and time poor' organisations with the fast approaching IPv4 exhaustion.

Simplified Addressing is a short to medium term strategy – and is not for long term use.

eintellego

Simplified Addressing continued…

Simplified address format:

2406:9800::f:110:76:128:158

(www.eintellego.net aka 110.76.128.158)

This allows you to represent ANY IPv4 address – Public or RFC1918

- This means you can even use it internally: 2001:db8::f:10:255:0:16, and with overlaps you could just use ::f0:….. and keep reusing the same ranges.
- In IPv4 we refer to each number as an 'octet' (in 8 bit terms). IPv6 has no official name we can find – so we refer to it as a Chazwazza ;-) Thanks to Nathan Ward and Kurt Bales.

It is now official!!(not)!! www.urbandictionary.com/define.php?term=chazwazza

eintellego

IPv6 is not hard…. (on Cisco IOS)

Initial Example using a Loopback interface – Cisco

```
rtr#conf te
rtr(config)# ipv6 unicast-routing    (won't work until the SDM profile is set correctly)
rtr(config)# interface loopback426
rtr(config-if)# ipv6 enable     (turns ipv6 on and generates the link local)
rtr(config-if)# ipv6 address 2001:db8::F:10:0:0:1/128
rtr(config-if)# end
rtr#ping 2001:db8::F:10:0:0:1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:db8:0:F:10::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/0 ms
```

- **On some Cisco switches you have to set SDM and reboot before you can use IPv6**

IPv6 is not hard…. (on Juniper JUNOS)

Initial Example using a Loopback interface – Juniper

root> configure

[edit]

root# set interfaces lo0 unit 0 family inet6 address 2001:db8:0:f:10:0:0:1/128

[edit]

root# commit and –quit

```
root> ping 2001:db8:0:F:10:0:0:1
PING6(56=40+8+8 bytes) 2001:db8:0:f:10:0:0:1 --> 2001:db8:0:f:10:0:0:1
16 bytes from 2001:db8:0:f:10:0:0:1, icmp_seq=0 hlim=64 time=0.252 ms
16 bytes from 2001:db8:0:f:10:0:0:1, icmp_seq=1 hlim=64 time=0.132 ms
16 bytes from 2001:db8:0:f:10:0:0:1, icmp_seq=2 hlim=64 time=0.185 ms
^C
--- 2001:db8:0:F:10:0:0:1 ping6 statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 0.132/0.191/0.252/0.029 ms
```

eintellego

Simplified Addressing – Network Equipment

Example: Interconnect on Cisco router (with dynamic routing)

```
interface VlanXXX
 description Layer3_to_Router
 ip address 10.76.132.65 255.255.255.252
 ipv6 address 2001:db8:0:F:10:76:132:65/128 (this is mainly for TS purposes)
 ipv6 enable
 ipv6 ospf 2 area 0
```

We rely on the link-local addressing for OSPF to function and establish neighbor relationships.

Simplified Addressing – Network Equipment

Example: Interconnect on Cisco router (without dynamic routing)

```
interface VlanXXX
 description Layer3_to_Router
 ip address 10.76.132.65 255.255.255.252
 ipv6 address 2001:db8:0:F:10:76:132:65/128
 ipv6 address 2001:db8:0:1C::4001:2/112
 ipv6 enable
!
ipv6 route ::0/0 2001:db8:0:1C::4001:1
```

With no dynamic routing you need a default route out of the device

Simplified Addressing – End User Connections

Example: End User connecting to a Cisco switch

```
interface VlanXXX
 description VMHEAD02.samplenetwork.net
 ip address 10.76.128.233 255.255.255.252
 ipv6 address 2001:db8:0:F:10:76:128:233/128 (Isn't needed, but useful for TS)
 ipv6 address 2001:db8:0:4019::1/64
!
ipv6 route 2001:db8:0:F:10:76:128:234/128 2001:db8:0:4019::2
```

We use /64 for all end customer assignments (even though we don't think a /64 is needed for the data centre)

Static route needed on the interconnection device to make v4-in-v6 work.

eintellego

Simplified Addressing – End User Connections

Example: End User connecting to a Juniper switch

set interfaces vlan unit 100 family inet address 10.76.128.**233**/30
set interfaces vlan unit 100 family inet6 address 2001:db8:0:4019::1/64

set routing-options rib inet6.0 static route 2001:db8:0:F:10:76:128:**234/**128 next-hop 2001:db8:0:4019::2

Static route still needed on the interconnection device to make v4-in-v6 work.

* Juniper cannot do a /128 on a host interface as we do on a Cisco (for reference only)

"set interfaces vlan unit 100 family inet6 address 2001:db8:0:F:10:76:128:**233/**128"

Results in: "Cannot have just a host address on a multipoint interface"

eintellego

Simplified Addressing – End User Connections

Example: The Linux End User configuration (RHEL 5.*x*)

NETWORK=10.76.128.156
GATEWAY=10.76.128.157
IPADDR=10.76.128.158
BROADCAST=10.76.128.159
NETMASK=255.255.255.252

IPV6INIT=yes
IPV6_AUTOCONF=no
IPV6ADDR=2001:db8:0:4019::2/64
IPV6_DEFAULTGW=2001:db8:0:4019::1

IPV6ADDR_SECONDARIES="2001:db8:0:F:10:76:128:234/128”

BGPv4 with IPv6

```
 interface GigabitEthernet0/1.651
  description IPv6 Transit Provider
  encapsulation dot1Q 651
  ipv6 enable
  ipv6 address 2001:DB8::FFFF:FFFF:F/112
!
 router bgp 65501
 neighbor Transit-v6 peer-group
 neighbor Transit-v6 remote-as 65500
 neighbor Transit-v6 update-source GigabitEthernet0/1.651
 neighbor 2001:DB8::FFFF:FFFF:1 peer-group Transit-v6
!
```

BGPv4 with IPv6

```
address-family ipv6
  neighbor Transit-v6 send-community
  neighbor Transit-v6 soft-reconfiguration inbound
  neighbor Transit-v6 route-map TRANSIT-V6-IN in
  neighbor Transit-v6 route-map TRANSIT-V6-OUT out
  neighbor Transit-v6 maximum-prefix 10000 95
  neighbor 2001:DB8::FFFF:FFFF:1 activate
  network 2001:DB8::/32
exit-address-family
!
ipv6 route 2001:DB8::/32 Null0 250
ipv6 route ::/0 2001:DB8::FFFF:FFFF:1
!
ipv6 prefix-list V6-BLOCKS seq 10 permit 2001:DB8::/32
```

BGPv4 with IPv6

```
ip as-path access-list 2 permit .*
ip as-path access-list 10 permit ^65500$
!
route-map TRANSIT-V6-IN permit 10
 match as-path 10
 set weight 500
!
route-map TRANSIT-V6-IN permit 20
 match as-path 2
 set weight 100
!
route-map TRANSIT-V6-OUT permit 10
 match ipv6 address prefix-list V6-BLOCKS
!
alias exec rectransitv6 show ip bgp ipv6 unicast neighbors 2001:DB8::FFFF:FFFF:1 received-routes
alias exec sendtransitv6 show ip bgp ipv6 unicast neighbors 2001:DB8::FFFF:FFFF:1 advertised-route
alias exec resettransitv6 clear ip bgp ipv6 unicast peer-group Transit-v6
alias exec bgpsum6 show ip bgp ipv6 unicast summary
```

BGPv4 with IPv6

```
RTR-A#bgpsum6
BGP router identifier 10.76.128.1, local AS number 65501
BGP table version is 979052, main routing table version 979052
2573 network entries using 401388 bytes of memory
5147 path entries using 391172 bytes of memory
55020/2026 BGP path/bestpath attribute entries using 9243360 bytes of memory
49574 BGP AS-PATH entries using 2077220 bytes of memory
25 BGP community entries using 652 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Bitfield cache entries: current 5 (at peak 6) using 156 bytes of memory
BGP using 12113948 total bytes of memory
2573 received paths for inbound soft reconfiguration
BGP activity 7438494/7124077 prefixes, 17613820/16969579 paths, scan interval 60 secs

Neighbor                 V  AS      MsgRcvd MsgSent  TblVer   InQ      OutQ     Up/Down  State/PfxRcd
2001:DB8::FFFF:FFFF:1    4  65500   1276188 442582   979052   0        0        7w0d          2572
```

Cisco Switch Problems

Oh oh.

- 3750 SDM Gotcha - The 3750 (along with the 3560) is Cisco's entry level IPv6 L3 switch.
- While generally a very versatile switch, changing the SDM template on a 3750 greatly reduces some of its TCAM availability in other areas.
- SDM "Routing" profile allows for 3000 directly connected IPv4 hosts and 8000 indirect subnets
- SDM "Dual-IPv4-and-IPv6 Routing" limits this to:
  - - IPv4: 1500 directly connected and 1250 indirect subnets
  - - IPv6: 1500 directly connected and 1250 indirect subnets
  - - Check your current TCAM usage:- *switch# show platform tcam utilization*
- WARNING: If you are not summarizing in your routing protocols and you have a large number of subnets floating around, this **can (and has) caused massive problems** for one of our clients
- WARNING 2: If you are also using 3750's in stacking mode with SDM Dual-IPv4-and-IPv6 Routing configured, there are scenarios where TCAM entries will fail to be installed (typically 0.0.0.0 just to hurt you!) and cause the switch to process switch all traffic.- cf: Cisco TAC ID CSCse38079

Juniper Switch Problems

More oh oh.

- Be aware that on the EX series switch (lab example was an EX3200-24) you need to be running at least 9.3 for IPv6 routing.
- Our lab tests showed that prior versions will allow you to configure IP addresses in all the appropriate locations, and even ping them from the device itself - but no device outside the switch was able to talk to it, even those devices connected to the switch could ping each other.
- Save yourself the head ache and run at least 9.3 according to Juniper's current recommendations
- https://www.juniper.net/customers/csc/software/junos_software_versions.jsp
- Junper's Current (as of March 2010) recommendation is:
  - JUNOS 10.0S1.1 for EX Series
  - JUNOS 9.6R3 for J and SRX series (even though 10 has been released)

Security

Oh oh (again)

- Enabling IPv6 leaves you wide open
- Every aspect of security needs to be replicated to IPv6
- SSH, Telnet, Access Lists, SNMP, CoPP – All are immediately open and accessible when you turn on IPv6.
- It isn't hard to do the security – you just HAVE to do it – or else
- Nothing has changed with the basic rules of security – just all new commands for some platforms – and in strange places
- The only consideration is that IPv6 requires ICMP for PMTU (Path MTU Discovery) – disabling it WILL break things (in ways that you can't easily troubleshoot)

The Example Linux Test bed

RedHat/CenOS 5.x

- Networking
- SSH
- Ip6tables (IPTables v6)
- Postfix/Sendmail
- Bind
- FTP
- NTP
- Apache (WWW)
- SQL
- SNMP
- Virtualisation

The Example Windows Test bed

Windows 2008

- Networking
- SSH (Remote Management)
- Exchange 2010
- Active Directory
- FTP
- NTP
- IIS
- MSSQL
- SNMP
- Microsoft Server Products -

Not Just Routers, Switches, Servers and Apps

- CPEs / Home Gateways
- ILO / DRAC
- Blade Management
- Storage Systems (SAN/NAS/etc) Management interfaces
- Printers / MFC / Photocopiers
- VOIP handsets / ATA
- Hardware Firewalls
  - Cisco ASA – from 8.2
  - Netscreen – ScreenOS 5
  - Juniper SRX – JUNOS platform supports IPv6
- Time-clock/Biometric Scanners
- IP Cameras & DVRs
- Cell Phones / PDAs
- Access Points
- Media Players
- Game consoles
- Video Conferencing
- Security Systems
- Building Automation
- UPS (with network support)

Who can help you with IPv6?

- Commercially - few companies in the region – most expertise is either in-house, especially with ISP's and Vendors at the moment
- Help each other – community? Do you help others?
- Training courses – APNIC (whole region), IPv6Now (AU/NZ), Fast Lane, Men & Mice (Not sure about .au/.nz), Dimension Data & New Horizons offer Cisco Cert module for IPv6 training (IP6FD) – DD was AU$4600 for a 5 day course - ouch
- Consulting and/or Implementation – eintellego (AU, NZ, FJ, AP), IPv6Now (AU), Braintrust (NZ), Prophecy (NZ), Avonsys (FJ, PAC), Cisco Professional Services in some countries, Juniper Professional Services if you use them – talk to Damien Holloway in HK.
- But it is growing all the time. If you have IPv6 skills… sell them!

eintellego

Advice

- If you have no access to IPv6 transit you may need to tunnel - talk to HE (Hurricane Electric) – Martin Levy and Owen DeLong are here at the conference
- If your transit provider doesn't do IPv6 today – start turning up the pressure and start getting alternative quotes if you can
- Same applies to your vendors (hardware, software, etc) – start demanding!
- 6 months after all our implementations, most of the technical ISP staff were fully conversant with IPv6 and had started to deploy other services – they were using it every day!
- Convincing management isn't that hard – Explain how much it will cost them later as opposed to now
- Some parts won't happen overnight – it takes time to migrate some services such as IPv6 DNS Servers to clueful registrars which support glue records for IPv6.

Predictions – Why you should do IPv6 NOW!

- The resource gold rush will happen – we're seeing it now. People are already doing IPv4 resource requests for no reason other than they can.
- Once this ramps up – the hype and outrage will accelerate it even faster.
- APNIC (and other RIR's) will have a large surge in membership of people not wanting to be left behind – and also from those hoping to capitalise on the resource shortage. There is very little the RIR's can do to prevent this happening – expect enforcement of tighter policies which we don't yet have
- This may bring forward exhaustion by 6-9 months – but we will see it coming
- A secondary market will appear for IPv4. For sale, for lease? Is $20,000 for a /24 a lot? Will it be more or less?
- IPv4 will have a cash value for 10 years with the height of value being around 2011-13

eintellego

Thanks for listening – Questions?

www.eintellego.net

eintellego