IPv6 Deployment Made Simple

Marc Blanchet

Viagénie

mailto:marc.blanchet@viagenie.ca

http://www.viagenie.ca

Presented at ClueCon, Chicago, August 2009

Viagénie Team Credentials

- Consulting and R&D in IP networking
  - Customers such as providers, enterprises, manufacturers for IPv6/VoIP/Routing deployment, application porting, ...
- 20+ years in IP networking and Unix
- 14 years on IPv6
- Team wrote IETF drafts and RFCs. Co-chair of IETF WGs (idn, vcarddav)
- Wrote the “Migrating to IPv6” book, published by Wiley
- Gave IPv6 tutorials at many conferences. Authored and delivered the Cisco IPv6 course.
- Co-founder and member of the board, IPv6Forum
- Member of steering group of North American IPv6 Task Force
- VoIP developers, ported Asterisk and Freeswitch to IPv6. (Also ported NTP, Quake, … to IPv6)
- Implemented NAT traversal server software (STUN/TURN): http://numb.viagenie.ca

Plan

- Why?
- Key Operational Concepts
- IPv6 configs in hosts, routers and PBX
- Deployment considerations
  - addressing
  - routing/switching
  - ...
- Conclusion
- Talk is not about:
  - developing applications with IPv6 (see cluecon 2008 presentation)
  - Trying to convince you about IPv6...

Why IPv6?

- Initially designed for 2 main reasons: IPv4 address space depletion and global routing table scalability.
- Opportunity to enhance IP: mobility, security, simpler address configuration, enhanced multicast, … IPv6 is a better IP compared to IPv4.
- But these improvements have not yet been sufficient to drive the general deployment of IPv6, except for some specific markets.
- Now, drive is the need for more IP address space, because of IPv4 address space depletion and large private networks:
  - Current estimate: IANA pool empty by Q2 2011. RIR pool empty by Q1 2012. Estimate will change.
- Some large scale networks are planning IPv6-only deployments because of lack of IPv4 address space already.
- Momentum: ipv6.google.com, ipv6.netflix.com, ...

IPv6 Key Operational Concepts

- IPv6 addresses:
  - Ex: 2001:1:2:3:a:b:c:d ; 128bits, 8 hex parts separated by “:”
  - Compressed forms: “::” means all zeros
  - With port number: [2001:1:2:3:a:b:c:d]:80
- Active interfaces automatically configure link-local addresses (fe80::...): link-local only used for communications on the directly connected link.
- Interfaces have multiple addresses
- Usual (and easiest) deployment: dual stack
  - Each interface has both an IPv4 address and IPv6 address.

IPv6 Key Operational Concepts (cont.)

- Address autoconfig
  - known as SLAAC: stateless Address Autoconfiguration
  - supplemental config mechanism (dhcpv6, manual also)
  - Router sends router advertisements that includes prefix, default router and other info. Host uses its mac address as host part.
- Prefix assignments:
  - Subnets/links are /64; no need for VLSM or else.
  - Enterprises receive /48 (65K subnets)
  - Homes receive either /64 (1 subnet) or /56 (256 subnets).
  - Providers receive /32 (65K enterprise customers) to start

OS and PBX IPv6 Configurations

N.B. Using 2001:db8:: as documentation prefix in examples. Do not use it on your network.

Windows

- Enabling IPv6:
  - XP: netsh interface ipv6 install (code is there, just enable it)
  - Vista: enabled by default
- By default, uses SLAAC
  - (and creates interfaces for Teredo and ISATAP: tunneling technologies)
- Network Connections->Interface->Properties->Networking->IPv6
- ipconfig; netsh interface ipv6
- Reference: http://www.microsoft.com/ipv6
  -

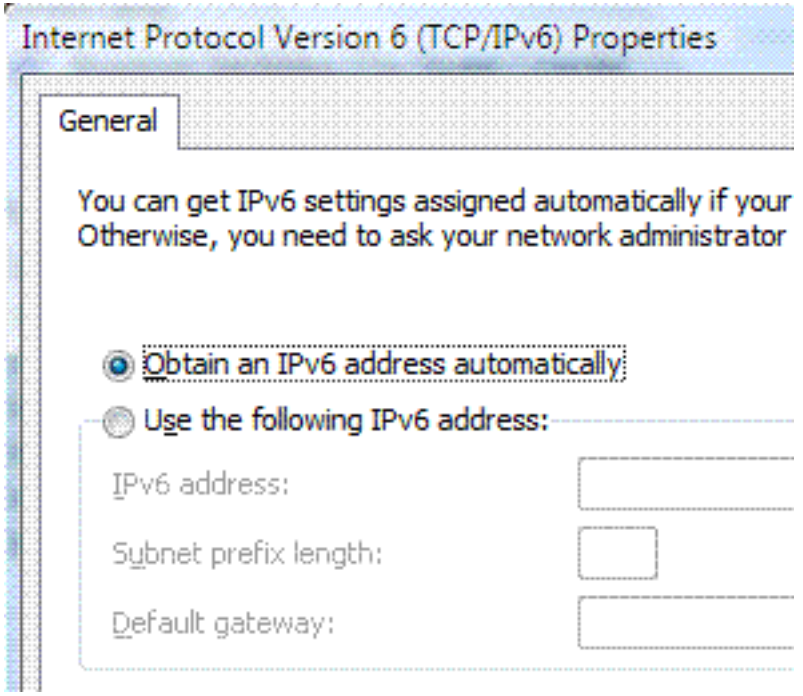

MacOSX

- Network Preferences->Advanced
- Supports: manual or SLAAC

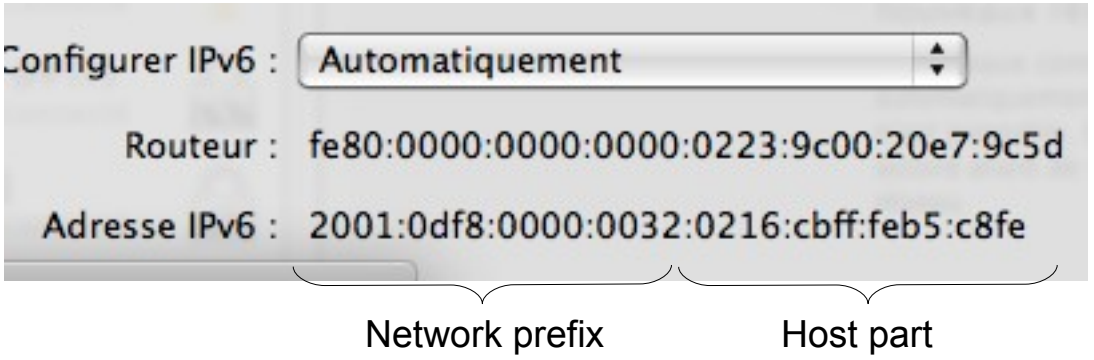


```
# ifconfig en1
  inet6 fe80::216:cbff:feb5:c8fe%en1 prefixlen 64 scopeid 0x6
  inet6 2001:df8::32:216:cbff:feb5:c8fe prefixlen 64 autoconf
  ether 00:16:cb:b5:c8:fe
```

Linux

- Fedora as example
- /etc/sysconfig/network
  - NETWORKING_IPV6=yes
- /etc/sysconfig/network-scripts/ifcfg-eth0
  - IPV6INIT=yes
    - activate IPv6 on interface
    - listen to RA (i.e. SLAAC)
  - Manual address:
    - IPV6ADDR=2001:db8:1:1::1/64
    - IPV6_DEFAULTGW=2001:db8:1:1::2/64

FreeBSD (and *BSD)

- /etc/rc.local
  - ipv6_enable=yes
    - # enables ipv6 on all interfaces
    - #listens to RA (ie. SLAAC)
  - Manual address config:
    - ipv6_ifconfig_fxp0=”2001:db8:1:1::1/64”
    - ipv6_defaultrouter=”fe80::2”

Cisco IOS

- ipv6 unicast-routing
  - enables ipv6 forwarding.
- ipv6 cef/dcef
  - Enables Cisco express forwarding for IPv6
- interface gigabitethernet0
  - ipv6 address 2001:db8:1:1::1/64
  - ipv6 nd prefix-advertisement 2001:db8:1:1::/64
    - (not needed. By default, sends RA when interface has ipv6 address)
- ipv6 route ::/0 gigabitethernet0 (::/0 = default route)
- ipv6 access-list ...

Juniper JunOS


```
interfaces fe-0/0/1 { unit 0 {
  family inet6 { address 2001:db8:1:1::1/64; } } }
  # enables IPv6 on the FE interface

protocols { router-advertisement {
  interface fe-0/0/1 { prefix 3ffe:b00:0:1::/64; } } }
  # configures RA on FE interface
```

Asterisk IPv6

- IPv6 feature not yet in trunk.
- Get code from http://www.asteriskv6.org
- Only SIPv6 available (no IAXv6)
- sip.conf
  - Bindaddr:
    - If none specified, asteriskv6 binds to both v4 and v6 wildcards
    - Or specify v4 or v6 address with/without ports
    - bindaddr=2001:db8:1::1
    - bindaddr=[2001:db8:1::1]:5061
    - Can specify many bindaddr lines as needed

Freeswitch IPv6

- Since 1.0.1
  - (announced at Cluecon 2008. see presentation)
- Only SIP is IPv6 capable. (not yet IAX,...). IPv6 ACL mostly coded, not yet committed.
- Conf:
  - $${local_ip_v6} refers to the binded IPv6 address
  - <param name=””> may be set to IPv6 addresses:
    - rtp-ip
    - sip-ip
  - Ex: <param name=”sip-ip” value=”2001:db8:1:1::1”>

Phones

- Many available as alpha-beta-or-soon-to-be-productized versions: Counterpath, Snom, Dlink, Polycom, ...
- Ask your preferred vendor...

Deployment Considerations

Addressing

- As enterprise, get IPv6 prefix (/48) from upstream, unless multihomed where you can get PI from RIR
- As provider, get IPv6 prefix (/32) from RIR
- Internal network:
  - Each subnet has /64
  - Only consideration for subnet assignments is IGP route prefix aggregation
- Efficient address plan method: RFC3531
-
- RIR = (ARIN, RIPE, APNIC, LACNIC, AFRINIC)

IPv6 Private Address Space

- Known as ULA (Unique Local Address)
- fd<40 bits random number>::/48
- (Shall be) non-overlapping private space.
- In addition to global address space:
  - ie. A link may have a global prefix and ULA prefix.
- May be used to number your network before you have an IPv6 upstream provider
- RFC4193

Routing/Switching

- OSPFv3:
  - different process/database/... than OSPFv2.
- ISIS:
  - same process; integrated topology. Be careful if v4 and v6 topologies are different: extension exists.
- BGP
  - Same process, new address family
- MPLS:
  - 6PE: PE are IPv6-enabled. P are unaware of IPv6.
  - 6VPE: for IPv6 VPN across MPLS core.

IPv4 and IPv6

- Dual-stack:
  - all links, hosts, routers run IPv4 and IPv6.
  - usual recommended approach for deployment
- IPv6-only:
  - Some starting to plan IPv6-only networks because depletion of IPv4 address space coming soon.
    - Ex: Large cable operators in North America
  - Easier for network operation: only one IP network to manage.
  - However, for application/service perspective, it is more difficult: need proxy or “smart” translators between the two address families. Scaling considerations.
  - IPv4 legacy nodes and applications.

Provisioning

- Infrastructure devices (servers, routers, vpn, switches, ...) shall have fixed (ie. manual) IPv6 addresses.
- Endpoints may use SLAAC. DHCPv6 is used when more control is needed.
- DNS IPv6 server addresses are sent either:
  - Through RA [RFC5006] (not implemented by all yet)
  - Through DHCPv6
  - Not so much a concern if you still run IPv4 DNS servers and dual-stack network
- While IPv6 stack supports multiple addresses per interface, try to provision the least number; makes operations simpler (logs, ACL, ...)

Upstream Provider

- If your upstream is not providing IPv6, you may use:
  - Tunnel brokers:
    - NAT traversal
    - Provides real IPv6 address and prefix
    - see: http://freenet6.net
  - 6to4: (v4 address embedded in v6 with automatic tunnelling). Be careful about:
    - RTT
    - dependency on your IPv4 external address
    - No NAT traversal
  - Teredo: for single hosts. NAT traversal.

Sample LAN

- Link: 2001:db8:1:1::/64
- Router: 2001:db8:1:1::1
- PBX: 2001:db8:1:1::2
- Phone1: 2001:db8:1:1::3 (or SLAAC or dhcpv6)
- Phone2: 2001:db8:1:1::4

- NB. 2001:db8::/32 is a documentation prefix. Don't use it on your network. Use ULA or assigned prefix.

VoIP IPv4 - IPv6 Interoperability

- IPv4 and IPv6 UAs can communicate via a relay.
- Usually relay is a B2BUA (e.g. FreeSWITCH)
- Consider using a cross-protocol TURN server.
- SIP Protocol/Deployment considerations:
  - Three protocols: SIP, SDP, RTP. Different states. Sometimes on different boxes.
  - Negociating different address families when not all pieces are at the same implementation levels may have “interesting” side effects.
  - Might want to consider single address family VoIP deployments for operational simplicity.
  - Not many people can claim to have good experience in mixed VoIP deployments. Be careful...

Conclusion

- Discussed:
  - IPv6 key operational concepts
  - Configurations for hosts, routers, PBXs
  - Deployment considerations
- Try IPv6 now!
  - http://www.freeswitch.org
  - http://www.asteriskv6.org
  - http://freenet6.net

Questions?

Contact info:

Marc.Blanchet@viagenie.ca

This presentation is available at http://www.viagenie.ca/publications/

References

- http://www.asteriskv6.org, http://www.freeswitch.org, http://freenet6.net
- IPv4 address depletion estimate: http://ipv4.potaroo.net
- [RFC3531] Blanchet, M., "A Flexible Method for Managing the Assignment of Bits of an IPv6 Address Block", RFC 3531, April 2003.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.
- Migrating to IPv6, Marc Blanchet, Wiley, 2006, ISBN 0-471-49892-0, http://www.ipv6book.ca
- IPv6 Network Programming, Junichiro itojun Hagino, Elsevier, 2004, ISBN 1555583180.