IPv6 It starts TODAY!

Thomas Hedströmmer
SE
thedstrommer@infoblox.com

Global Address Space Update

- **OK, who hasn't heard that IPv4 really is running out?**
- **Reality is:**
  - IANA IPv4 pool exhausted on 3 Feb 2011
  - RIRs rapidly running out
    - APNIC reached it's last /8 on the 15th April
    - RIPE has < 5 /8s left
  - LIRs will run out last
- **So expect RIRs and LIRs to start running out in sometime from 2012 onwards and enforcing strict allocation rules**
- **Growth can only be achieved by deploying IPv6**

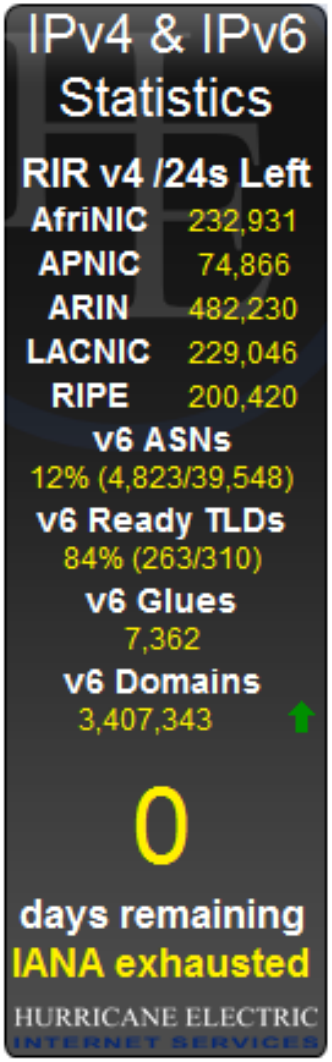

Infoblox

What's the One Thing Everyone Knows About IPv6?

128-bit addresses

Infoblox

But How Big Is That?

2^{128} addresses =

340,282,366,920,938,463,463,374,607,431,768,211,456

340 *undecillion* addresses


Say this is all of IPv6
address space...

Then all of IPv4 address space is a little bigger than this...

So How Do You Write an IPv6 Address?

- One way: 128 bits!

0010 0000 0000 0001 0000 1101 1011 1000…

- Convert to hex: 2 0 0 1 0 d b 8…

- Separate "quartets" with colons: 2001 : 0db8…

- Eight hex quartets equals 128 bits:
  2001:0db8:0000:0000:0001:0002:0003:0004
- Omit leading zeroes in a quartet:
  2001:db8:0:0:1:2:3:4
- Replace *one* series of consecutive all-zeroes quartets with "::":
  2001:db8::1:2:3:4

Infoblox

GOOD NEWS:

- Most infrastructure devices (e.g., routers) support IPv6 already
- Most modern operating systems (e.g., Windows 7, Mac OS X, Linux) support IPv6 already
- IPv6 is already routed over the backbone of the Internet

BAD NEWS:

- Much CPE doesn't support IPv6
- Very few applications support IPv6
- Very few organizations use IPv6
- There's very little organizational experience with or knowledge of IPv6

About IPv4 and IPv6

IP version	IPv4	IPv6
Deployed	1981	1999
Address Size	32-bit number	128-bit number
Address Format	Dotted Decimal Notation: 192.0.2.76	Hexadecimal Notation: 2001:0DB8:0234:AB00:0123:4567:8901:ABCD Literal (e.g. browser): [2001:DB8:1234::101]
Number of Addresses	2^{32} = 4,294,967,296	2^{128} = 340,282,366,920,938,463, 463,374,607,431,768,211,456
Examples of Prefix Notation	192.0.2.0/24 10/8	2001:0DB8:0234::/48 2600:0000::/12
Security	IPSec	IPSec Mandated, works End-to-End
Mobility	Mobile IP	Mobile IP with Direct Routing
Quality of Service	Differentiated/Integrated Service	Differentiated/Integrated Service
IP Multicast	IGMP/PIM/Multicast BGP	MLD/PIM/Multicast, BGP, Scope Identifier

Types of IPv6 Addresses and Networks

Infoblox

Prefix	Type	IPv4 Equivalent
:: (technically ::/128)	Unspecified address	0.0.0.0
::1 (technically ::1/128)	Loopback address	127.0.0.1
fc00::/7 (specifically fd00::/8)	Unique Local Addresses	RFC 1918 (e.g., 10/8)
fe80::/10	Link-local Addresses	169.254/16
2001:db8::/32	Documentation	192.0.2/24
2000::/3	Global Unicast	
ff00::/8	Multicast	224/4

Prefixes and Subnetting

- **IPv6 is aim at using public address space – no need for NAT**
  - Guaranteed unique address space
  - Simplifies company acquisition and merger
  - 2000::/3
- **Alternatively use "Unique Local Addresses" equivalent of RFC1918**
  - But use correctly with randomized 40 bits to create uniqueness
    - http://www.sixxs.net/tools/grh/ula/
  - FD00::/8
- **Typical Prefix Allocation = /48**
  - This can be further locally sub-netted typically directly to /64s
    - providing 65535 /64 prefixes!
- **/64 leaves 64-bits for EUID-64 based host IPs**

Multicast in IPv6

- **IPv6 does not use Broadcasts**
- **Multicast used instead for internal services**
  - Set of 'well known' service multicast addresses

Address	Description
ff02::1	All nodes on the local network segment
ff02::2	All routers on the local network segment
ff02::5	OSPFv3 AllSPF routers
ff02::6	OSPFv3 AllDR routers
ff02::9	RIP routers
ff02::a	EIGRP routers
ff02::d	PIM routers
ff02::16	MLDv2 reports (defined in RFC 3810)
ff02::1:2	All DHCP servers and relay agents on the local network site (defined in RFC 3315)
ff02::1:ffxx:xxxx	Solicited-Node Address (last 24-bits of IPv6 unicast address to be resolved)
ff05::1:3	All DHCP servers on the local network site (defined in RFC 3315)
ff0x::fb	Multicast DNS
ff0x::101	Network Time Protocol
ff0x::108	Network Information Service
ff0x::114	Used for experiments

Wikipedia

Address Assignment Mechanisms

IPv6 Provisioning Mechanism	And What Sort of a Thing Is That?
Stateless Address Auto configuration (SLAAC)	Client derives IPv6 address from router-advertised prefix and self-generated suffix
Stateful DHCPv6	Like DHCP, but for IPv6
Stateless DHCPv6	Client uses SLAAC for IP address, DHCP for DNS, etc.
RFC 5006/6106 (RDNSS/DNSSL)	Client uses router advertisements for DNS, etc.

The A, M, O? Router Advertisement Flags

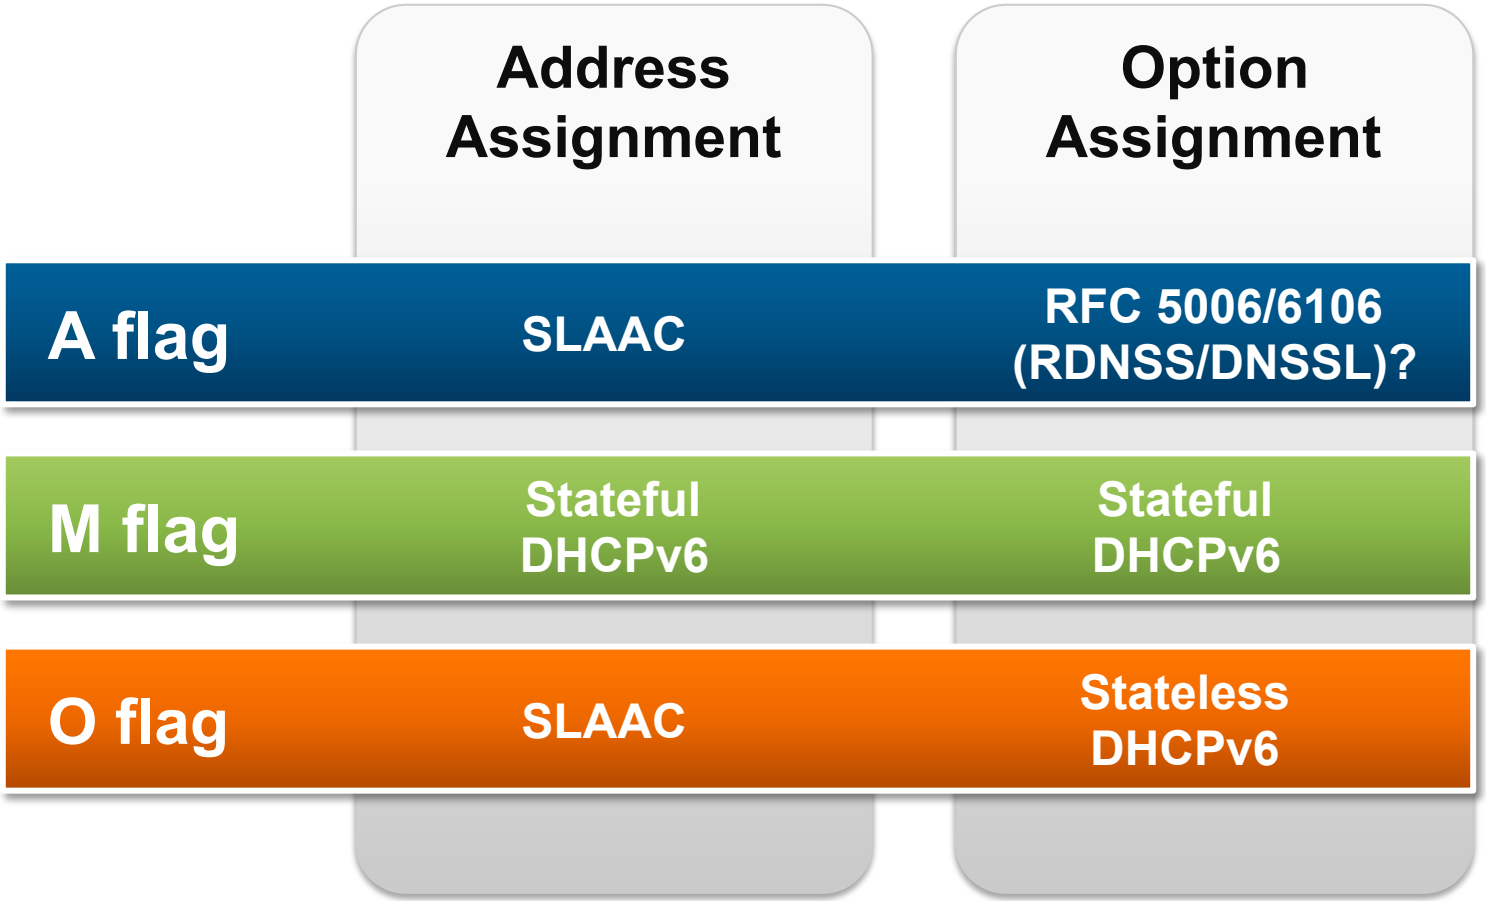

Not as simple as DORA!

IPv4 DHCP	IPv6 DHCP
DHCPOFFER	ADVERTISE (2)
DHCPREQUEST	REQUEST (3), RENEW (5), REBIND (6)
DHCPACK/DHCPNAK	REPLY (7)
DHCPRELEASE	RELEASE (8)
DHCPINFORM	INFORMATION-REQUEST (11)
DHCPDECLINE	DECLINE (9)
--	CONFIRM (4)
DHCPFORCERENEW	RECONFIGURE (10)
--	RELAY-FORW (12), RELAY-REPLY (13)

DHCPv6 (Typical Flow)

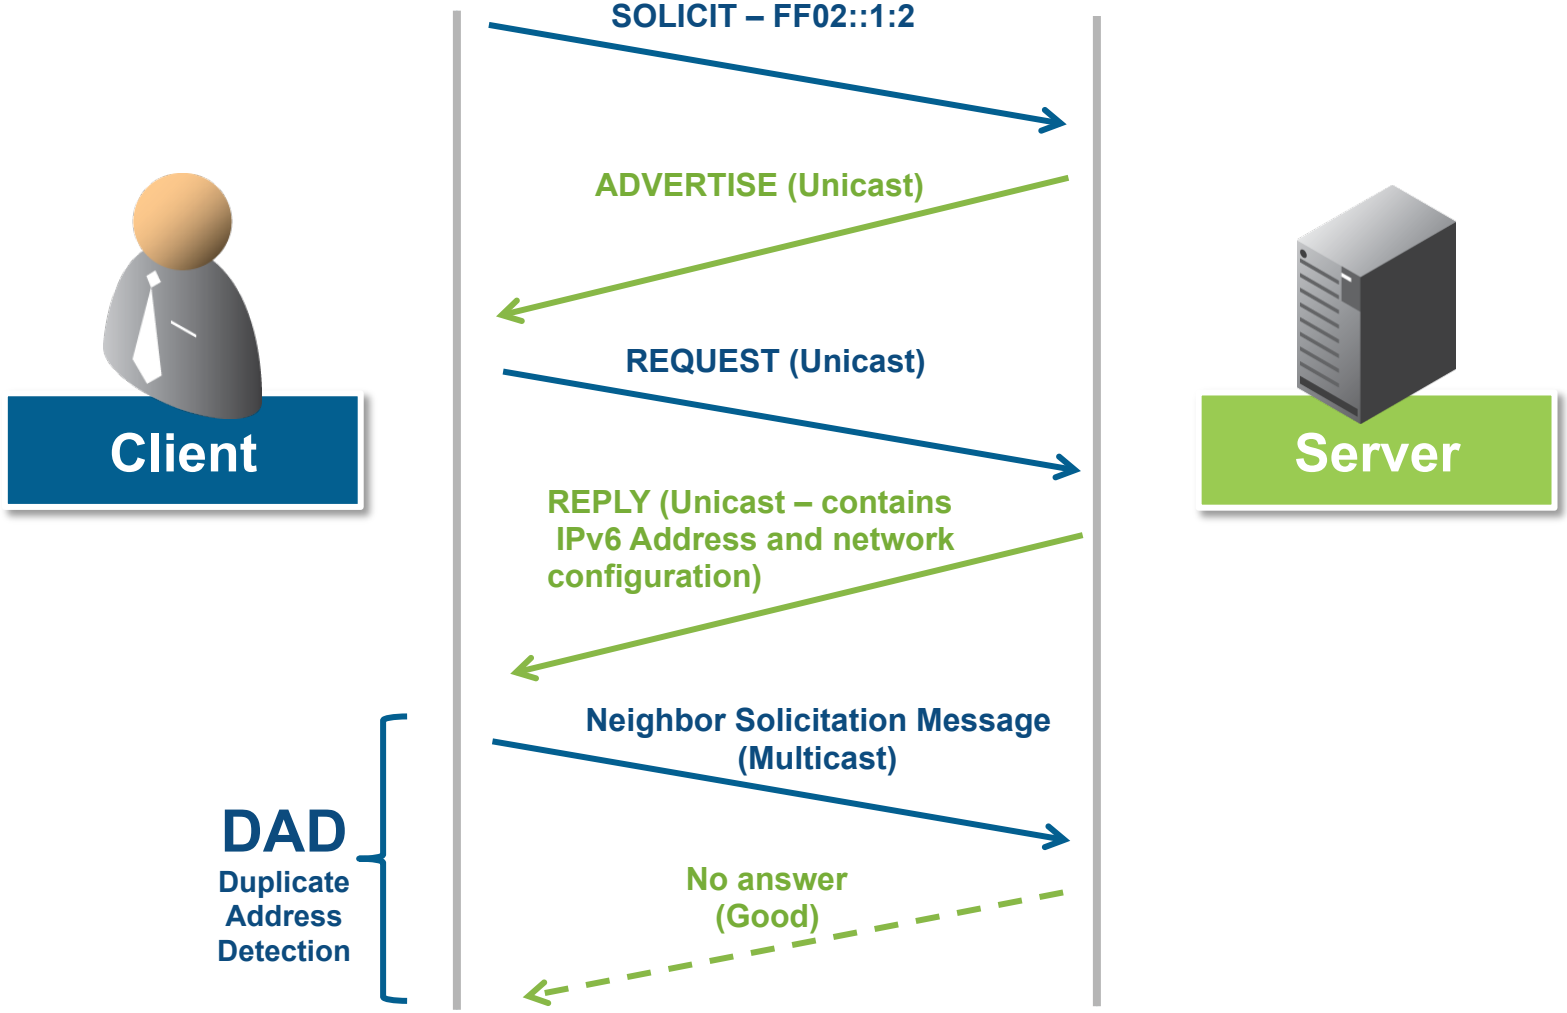

Spot the IPv6 Address!

```
% ifconfig lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> mtu 16384
        options=3<RXCSUM,TXCSUM>    inet6 fe80::1%lo0 prefixlen 64 scopeid 0x1    inet
127.0.0.1 netmask 0xff000000        inet6 ::1 prefixlen 128 en0:
flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500
        options=27<RXCSUM,TXCSUM,VLAN_MTU,TSO4>       ether 00:23:df:a9:98:14
        inet6 fe80::223:dfff:fea9:9814%en0 prefixlen 64 scopeid 0x4
        inet6 2001:db8:1f05:1a2f:223:dfff:fea9:9814 prefixlen 64 autoconf
        inet6 2001:db8:1f05:1a2f:e1d1:33a6:981f:de48 prefixlen 64 autoconf temporary
        inet6 2001:db8:1f05:1a2f::64 prefixlen 64     inet 192.168.0.244 netmask
0xffffff00 broadcast 192.168.0.255   media: autoselect (100baseTX <full-duplex,flow-
control>) status: active
en1: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500    ether 00:23:6c:
97:bc:bf inet6 fe80::223:6cff:fe97:bcbf%en1 prefixlen 64 scopeid 0x5
        inet6 2001:db8:1f05:1a2f:223:6cff:fe97:bcbf prefixlen 64 autoconf
        inet6 2001:db8:1f05:1a2f:dc3a:80d7:5ea:f0c5 prefixlen 64 autoconf temporary
        inet6 2001:db8:1f05:1a2f::6b prefixlen 64     inet 192.168.0.236 netmask
0xffffff00 broadcast 192.168.0.255   media: autoselect  status: active
```

DNS and IPv6: Forward Mapping

- Forward mapping (mapping domain names to IPv6 addresses) uses the AAAA record

```
charlie      IN    AAAA  2001:db8:cafe:1::2
jessie       IN    AAAA  2001:db8:cafe:1::3
```

- Can use the abbreviated form of the IPv6 address
- AAAA records can be added to any zone data file
- A domain name can point to both A and AAAA records (for hosts with dual stacks)

DNS and IPv6: Reverse Mapping

- **Reverse mapping (mapping IPv6 addresses to domain names) uses the PTR record**
  - Just like IPv4 reverse mapping
- **Unfortunately, the *entire* (i.e., unabbreviated) IPv6 address is embedded in the domain name**

```
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.e.f.a.c.8.b.d.0.1.0.0.2.ip6.arpa.
      IN       PTR      charlie.foo.example.
3.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.e.f.a.c.8.b.d.0.1.0.0.2.ip6.arpa.
      IN       PTR      jessie.foo.example.
```

- **These records obviously go in different reverse-mapping zones**

This is transition time….

- **The preferred transition technology is IPv6 to IPv4 (not IPv4 to IPv6) therefore enterprise needs to take IPv6 strategy seriously**
  - Dual Stack
  - External Services
  - IPv6 Only Networks

- **NAT46 is bad**
  - Standards tried and deprecated
  - Performance Issues

- **NAT64 (with DNS64)**
  - Standardized
  - Works
  - Available

DNS64/NAT64 Use Cases

Carrier/ISP/Mobile

- **Lack of IPv4 address space limits rollout of new services**
  - WiFi / WiMAX
  - Smart Phones
  - Consumer Broadband
- **Need to connect customers who are IPv6 only to "legacy" IPv4 sites that have yet to adopt IPv6**

Enterprise IT

- **Not an immediate need for most enterprise IT**
- **Can be used to provide IPv6 at the edge (web properties)**
- **Will be more important when IPv6 rollout is near completion**
  - Customer will eventually wish to disable IPv4 on most infrastructure
  - Some legacy systems will not "speak" IPv6
  - Customer can use DNS64/ NAT64 to create islands of legacy IPv4 services

DNS64 – Infoblox Feature

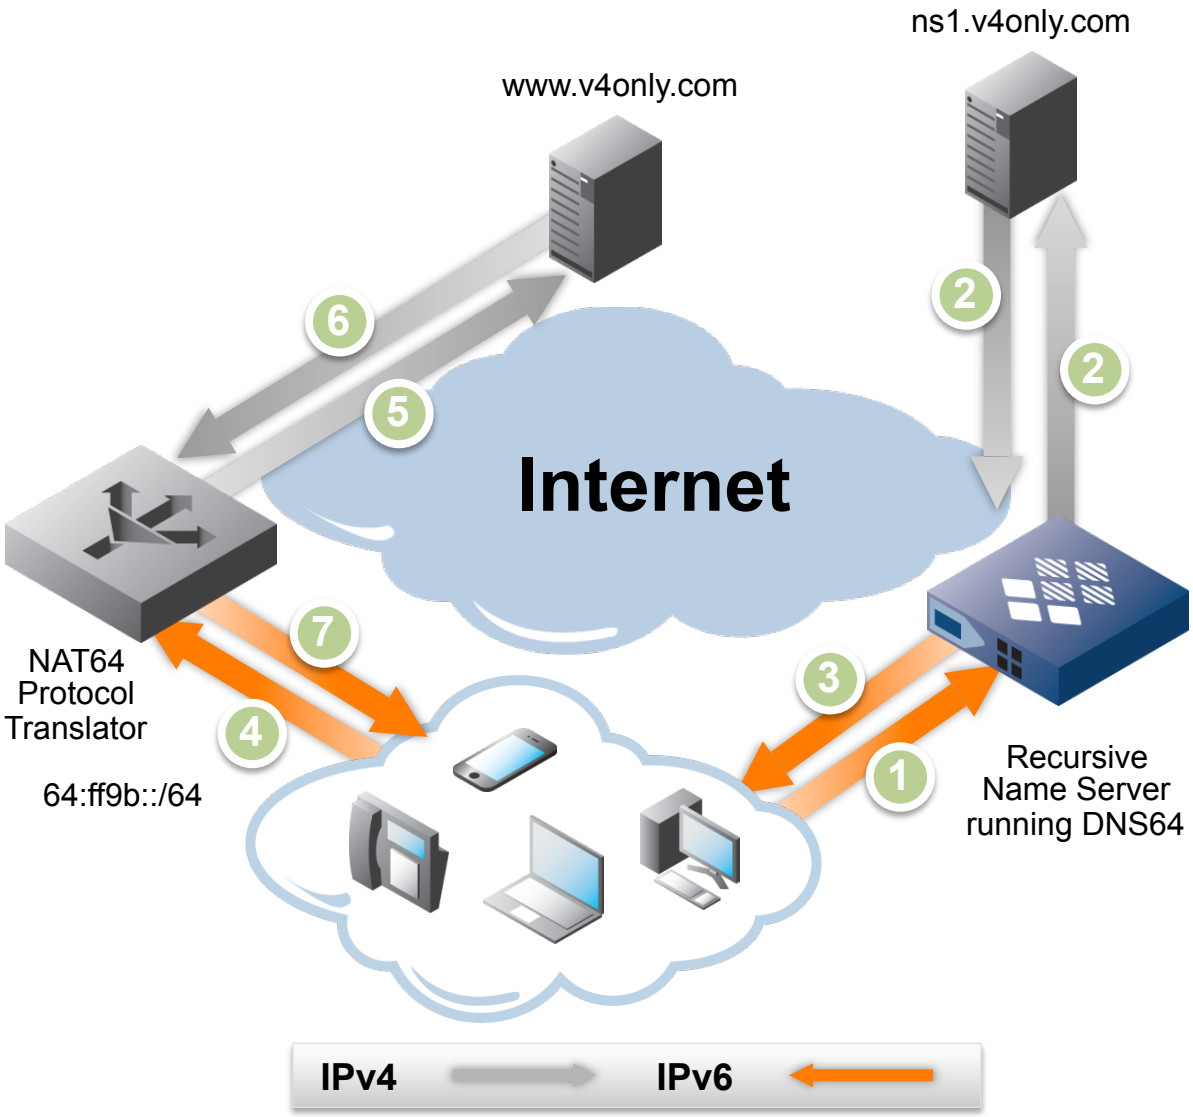

1. Client queries www.v4only.com AAAA to local recursive name server
2. Recursive name server queries name server for v4only.com and gets no AAAA response
3. Recursive name server synthesizes a IPv6 address to return to client – using /96 prefix
4. Client sends packet to synthesized IPv6 address which routes to the NAT64 protocol translator
5. NAT64 device translates packet to destination IPv4 address
6. IPv4 only web server returns the response over IPv4 to NAT64 device
7. NAT64 device converts the packet to IPv6 to return to the originating client

DNS64 / NAT64 - Integration

The Infoblox DNS64 solution requires a NAT64 device

- Tested and certified by Infoblox QA with LTM product
- Working on joint press
- LTM product – version 11.0

http://www.f5.com/solutions/resources/solution-profiles/managing-ipv6.html

- Tested and certified by Infoblox
- AX Series – version 2.6.4

http://www.a10networks.com/products/axseries-NAT64_DNS64.php

- May work with JunOS 10.4
- Not currently certified – working on integration testing

- Some devices support NAT64
- No certification at this time

DHCPv6 – Customer Use Cases

Stateless Deployments

- Infoblox DHCP server provides IPv6 DHCP options
- Other client settings are received from router or auto-configured

Stateful Deployments

- Infoblox DHCP server provides IPv6 address and IPv6 DHCP options
- Much like IPv4 works today

- **Which one will customers deploy? Stateless or Stateful?**
- **Delegation of IPv6 prefixes**
  - In addition to leasing a single address an IPv6 DHCP client can now lease a whole network
  - IPv6 prefixes are delegated to other devices
  - Prefix address space may differ from client address location

DHCPv6 – Stateful vs. Stateless

Infoblox

Stateless Deployments

- Client receives all required information from the router to configure default gateway and address
- Generally no DHCP options – However, additional options provided by DHCP if available

Stateful Deployments

- Client receives address from the DHCP server
- Client receives options from the DHCP server

- **"M" and "O" bits in RA (router advertisement)**
  - "M" == "Managed address configuration" flag – tells the client that DHCPv6 is available for address configuration
  - "O" == "Other configuration" flag – tells the client that DHCPv6 is available for other (non-address) information
    - For example: DNS server options

DHCPv6 – Stateful with Infoblox

Infoblox
IPv6 Enabled
DHCP Server

IPv6
Router

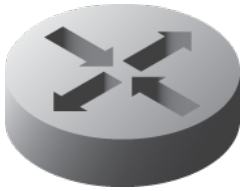

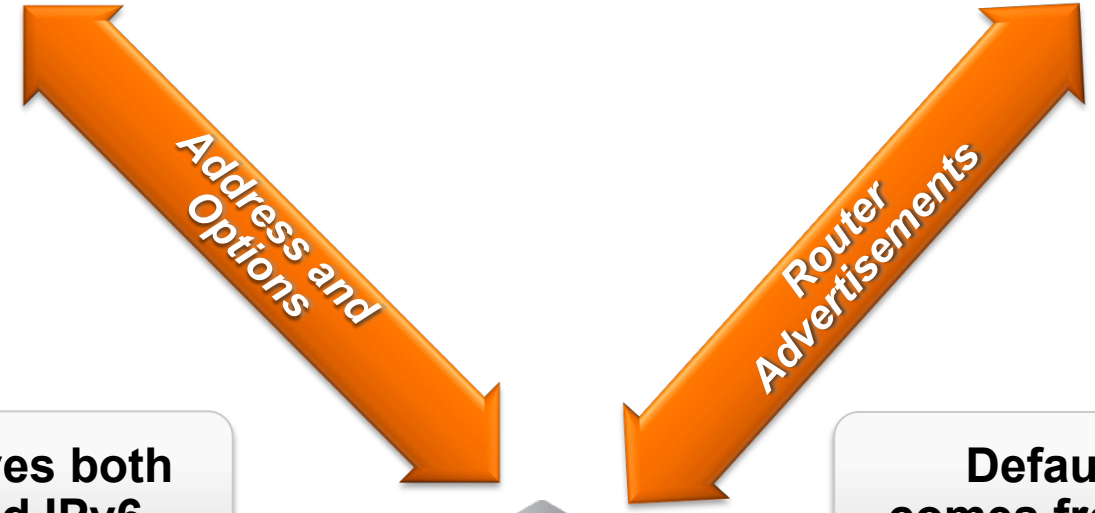

Client receives both address and IPv6 DHCP Options from the DHCP server

Default router comes from the IPv6 router via router advertisements

IPv6 Client

DHCPv6 – Stateless with Infoblox

Infoblox
IPv6 Enabled
DHCP Server

IPv6
Router

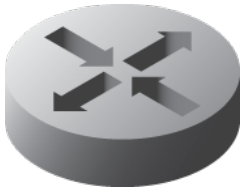

Default router comes from the IPv6 router via router advertisements

Options

Router and Address

Client receives only IPv6 DHCP Options from the DHCP server

Client auto-configures IPv6 address based on network prefix

IPv6 Client

Resources

- http://www.getipv6.info/
- http://www.tunnelbroker.net/
- http://ipv6.he.net/certification/
- http://www.6deploy.eu/

Questions?

Infoblox

Thank you!