

## IPv6 @ LinkedIn



**Franck Martin**


Senior Staff Engineer – Chair of LinkedIn IPv6 in the DC WG




**2012**

The world is going to run out of IPv4 addresses in 3-5 years time.

### Geoff Huston IPv4 exhaustion by RIR




Data flows between LI and its members has been **IPv6-enabled since 2014**. Apache Traffic Server listens on IPv6 and sends the internal request over IPv4, with the IPv6 address in a special field



Email: 2013 Web: 2014

To allow testing, in 2014, we enabled IPv6 in our offices.


### IPv6 to LinkedIn website (worldwide)



**IPv6 is faster than IPv4\*.**  
Most mobile providers in the USA are at more than 75% deployment.

\*Mileage may vary and devil is in the details


### LinkedIn IPv6 Heat Map



2015

LinkedIn is going to run out of RFC1918 (Internal) IPv4 addresses in 3-5 years time.

### RFC1918 /16 exhaustion @ LinkedIn (measured late 2016)



2015

IPv6-DC-WG established  
Let's build Oregon DC dual stack with no AAAA  
Let's build next data center after it IPv6-only

### IPv6 in the Data Center: Oregon Dual Stack



IPv6 in the DC?

- **Scale** – From dense (x10) to virtual compute (x100)
- **Opportunities** – New technical solutions not constrained by addressing space
- **End to End connectivity** – No NAT44 between DCs or offices.



For traffic to go on IPv6

Client: IPv6 global address → Server: IPv6 global address + DNS AAAA

### 3 Pillars

#### Network

- Static vs Dynamic
- ACLs and security rules
- VIPs, Anycasts
- Edge Networks
- No more deep NATs

#### Hardware

- UEFI Network boot over IPv6
- BMC/IPMI over IPv6
- Server Build Automation

#### Software

- Listening over IPv6
- Discovering Services
- Connect strategy (fail, failover, HE)
- Upgrade, Upgrade, Upgrade
- Support legacy IPv4 software

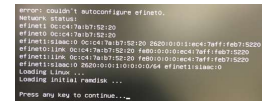


#### Network

- **Static** – IPv6 address decided at build time
- **Gateway** – FE80::1
- **Tools** – Convert IPv4 configuration to IPv6 automatically
- **Mapping** – IPv6 can be deduced from IPv4 for dual stack machines but without technical debt – Each IPv4 network has a paired IPv6 network



#### Hardware



- **UEFI** – Supports IPv6 boot (SLAAC vs DHCPv6, TFTP vs HTTP)
- **Grub** – IPv6 support is weak
- **BMC/IPMI** – Redfish standard but IPMI tools are lagging
- **Firmware** – Always dangerous to flash at scale



#### Software

- **Listener** – Listens on IPv4 and IPv6 – every language is special
- **Connect strategy** – hardfail, fallback, Happy Eyeballs
- **Java** – control in java settings
- **Deploy** – and redeploy till right
- **IPv6 ready** – “should work” or limited support – Test, don’t believe.

## Key Takeaways

## Where are we?



- Staging environment entirely dual stack with A and AAAA
- Retrofitting production environment
- Working on building IPv6 only servers

**We don't want to deploy IPv6, we want to remove IPv4!**

## IPv6 Takeaways



Zaid Ali, Sritam Avelia, Andrey Bilib, Donaldo Carvalho, Tim Crofts, Bo Feng, David Fontaine, Prakash Gopinadham, David Hoa, Sanaldas KB, Henry Ku, Prasanth Kumar, Vikas Kumar, Tommy Lee, Leigh Maddock, Navneet Nagori, Marijana Novakovic, Ved Prakash Pathak, Stephanie Schuller, Chintan Shah, Harish Shetty, Andrew Stracner, Veerabahu Subramanian, Shawn Zandi, Andreas Zaugg, David Paul Zimmerman...

- Know all the unknowns very early
- Engage vendors straight away (takes one year to deliver)
- Don't relax or forget you have IPv6
- Engage Top Level Management
- Need Software Engineers more than Network Engineers

Thank you