

The Beginning of the End

A Quest for IPv6 only

Data Centers



Tim Crofts
Production infrastructure Engineering

A brief history

- June 6th 2012 World IPv6 Launch Day
- Email delivery was enabled in Oct 2013
- www.linkedin.com launch in June 9th 2014
- 50% of our traffic is now via IPv6
- In 2015 decided to move our Data Centers to IPv6

What was IPv6 at that point?

linkedin.com is enabled and we'd already done a bunch of work internally

- Enabled key parts of corporate network
- Enabled Edge networks
- Enabled VPNs
- Email

So the story continues from here..

What was left to enable

Our data centers.

- Production networks and servers
- Staging and Testing networks and servers
- Management networks
- Backbone network
- Intranet services.

Plan of Action

- Established working group
- Target environment: staging
- Targeted systems: infrastructure systems
 - NTP, syslog, DNS, Kerberos, SMTP
- Adding IPv6 without AAAA record
- Slowly adding AAAA to enable services.
- Expand outside of the operational area to developers
 - Only adding addresses
 - Slowly increasing the number of boxes with AAAA records
- New production data center with IPv6 from day 0



Implementation

Static IP configuration.

IP Numbering.

- Network bits
 - Mapping of /23, /24, /26 -> /64
- Host bits
 - aa.bb.cc.dd -> ::ccdd
- Gateway
 - Is always fe80::1 (%dev)

Findings

Some application just work and some don't behave as you'd expect.

- DNS just works
- Kerberos just works
- NTP well..
- Syslog just works but..
- Ssh works but...

Ramping up

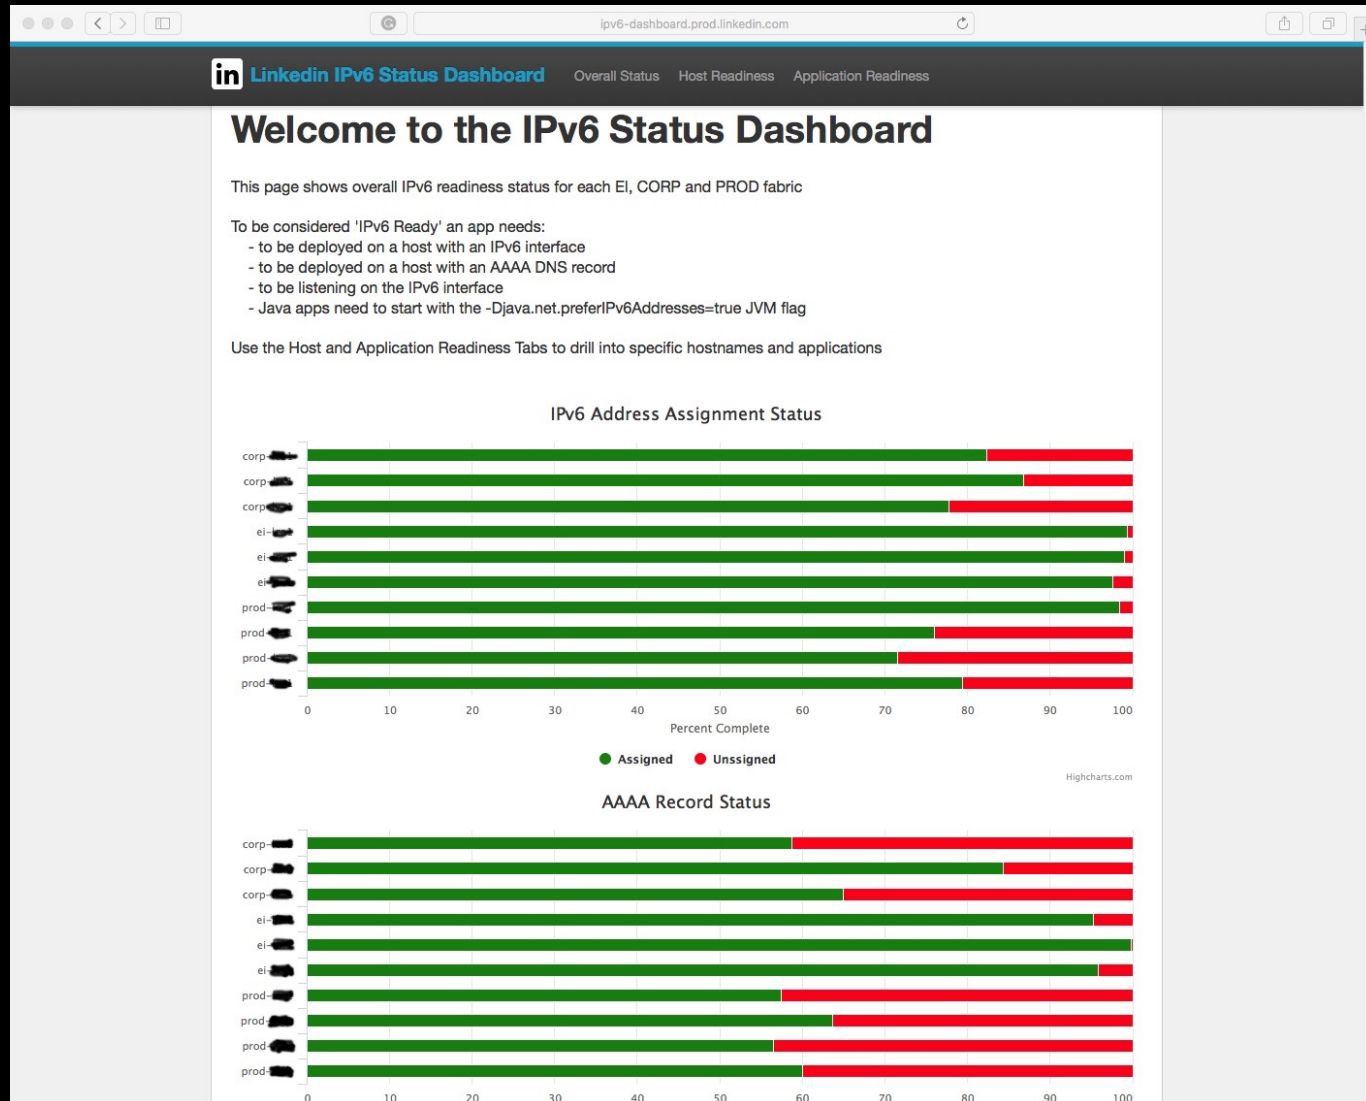
Various languages using IPv6

- Python
- Java
- Ruby

Still catching up

- Hadoop
- Couchbase
- JIRA/Confluence

Measure it



Doubling up

Dual stack consumes more. Memory, resource & time.

- Resource
- Maintaining ACLs
- Debugging

So treat as a stepping stone.

Removing of V4

- Can we remove v4
- The challenges of removing IPv4
- Get your developers to think V6 first

```
tcrofts — -bash — 117x21

from ps output
tcrofts 6608 6602 0 20:34 ? 00:00:00 nc -w 28807 app1234 22

from lsof
sshd 6602 tcrofts 3u IPv6 1455547370 0t0 TCP bastion01:ssh->[2600:19:5000:20c:6c00:2e:baad:82]:49887 (ESTABLISHED)
nc 6608 tcrofts 3u IPv4 1455548022 0t0 TCP bastion01:51794->app1234:ssh (ESTABLISHED)

snipit from sshd_config
Match Address 172.16.0.0/12,!172.31.0.0/16
    RequiredAuthentications2 publickey
    Banner none
    ForceCommand /etc/issue.sh

file /etc/issue.sh
#!/bin/bash
echo "connecting using legacy IPv4"
echo "You need to use IPv6 to connect to this host"
echo "Raise a support ticket if your client doesn't support IPv6"
exit 1
```




The IPv6 Frontier: Provisioning

- PXE over IPv6
- Legacy BIOS will not support IPv6
- UEFI last spec has IPv6
- ONIE
- Firmware support is sketchy

Why provisioning is important now

At LinkedIn it takes us about year to build a new Data Center

- Location
- The design
- The hardware
- If it doesn't support IPv6 we are stuck with it for several years

Whats required

- 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

```
error: couldn't autoconfigure efinet0.  
Network status:  
efinet1 0c:c4:7a:b7:52:20  
efinet0 0c:c4:7a:b7:52:20  
efinet1:slaac:0 0c:c4:7a:b7:52:20 2620:0:0:1:ec4:7aff:feb7:5220  
efinet0:link 0c:c4:7a:b7:52:20 fe80:0:0:0:ec4:7aff:feb7:5220  
efinet1:link 0c:c4:7a:b7:52:20 fe80:0:0:0:ec4:7aff:feb7:5220  
efinet1:slaac:0 2620:0:0:1:0:0:0:0/64 efinet1:slaac:0  
Loading Linux ...  
Loading initial ramdisk ...  
  
Press any key to continue..._
```

Conclusion so far

- Be prepared to ask vendors
- Server provisioning works.
- Get developers on board with IPv6.
- Measure your transition progress.
- You can control the transition.
- Core services work.
- Have a plan.
- It would ideal if you could do a clean switch.

Q & A



<https://engineering.linkedin.com/blog/2016/07/ipv6-at-linkedin-part-i--chippin-away-at-ipv4>