



# Local Presence Global Reach

## IPv6 transition success story

# IPv6, why oh why?

- Because we can.
- It is not going to go away.
- Give it a try, ok?
- Instant consulting!
- Please. It's 2011 already.

# We accidentally IPv6

- We were migrating to new IPv4 PI addresses anyway – adding IPv6 was no big deal. Volume discount!
- First deployed on the DNS servers
- Then routers, firewalls, and selected VPN tunnels
- Third step: the data center edge
  - Content delivery (reverse proxy/cache)
  - Load balancers (err, “application delivery”)
  - Only the untalented study. We did. Sorry.

## ...is this bad?

- Our network gear proved to support IPv6 surprisingly well: Extreme, Juniper, F5 BIG-IP
- OS support is decent, usually  
Red Hat Enterprise Linux, Ubuntu etc.
- Various applications mostly OK too, e.g.  
BIND, Postfix, Apache, Nginx
- On servers, SLAAC is usually a bad idea.  
Instead, configure role addresses manually for each service.

# All work and no play

- Must configure everything twice:
  - Routing
  - Access control in firewalls and applications
  - Measurement and monitoring systems
- Application configuration needs some attention
- Your Subnet Is Always /64
- Quite often you have to install bleeding edge versions of your server software to support IPv6.

# First blood

- IPv6 is usually enabled by default on modern OSes. The link-local address is always configured, but:  
Leak one RA and all your servers are belong to IPv6!
- Apparent Red Hat Linux kernel bug
  - IPv6 interface just stops working. [fixed]
- Juniper MX80, ae0, VLANs + IPv6 VRRP [fixed]
- Protocol conversion IPv4-IPv6 is easy with a SLB appliance or Nginx proxy, but on a firewall?

# When the edge is not enough

- IPv6 does not stop at your edge!
  - SLB appliance + HTTP header injection:  
“hide proxy” at back-end must support IPv6 too
  - Logs – analysis, statistics, forensics...
  - Any application using the source IP address to anything: your email feedback form/script may just go haywire after looking at the strange IP address.

# Diving deep

- ECMP load balancing from routers to reverse proxy servers
- No state: hashing L3 or L3+L4 headers  
and each tcp flow will not break
- IPv6 and path mtu discovery: ICMP “packet too big” will probably be sent to a different server than the tcp packets!

Why? The ICMP packet has the source address of the router somewhere along the route. Flow id is in the ICMP payload.

- MTU 1280? – and IPv4 has the same problem.
- Use DNS round-robin for LB, routing for HA. No ECMP.



# Popular IPv6 enabled web sites in Finland

- MTV3
  - `www.mtv3.fi` `2001:67c:224:81::1`
  - `static.mtv3.fi` `2001:67c:224:81::1`
  - `img.mtv3.fi` `2001:67c:224:81::1`
- Iltalehti
  - `www.iltalehti.fi` `2001:67c:224:223::1:80`
  - `static.iltalehti.fi` `2001:67c:224:223::2:80`

**Thanks and sorry!**

[cybercomgroup.com](http://cybercomgroup.com) [cybercomgroup.com](http://cybercomgroup.com)