ScandiNode

IPv6 in LAN environments

Stephan Lagerholm

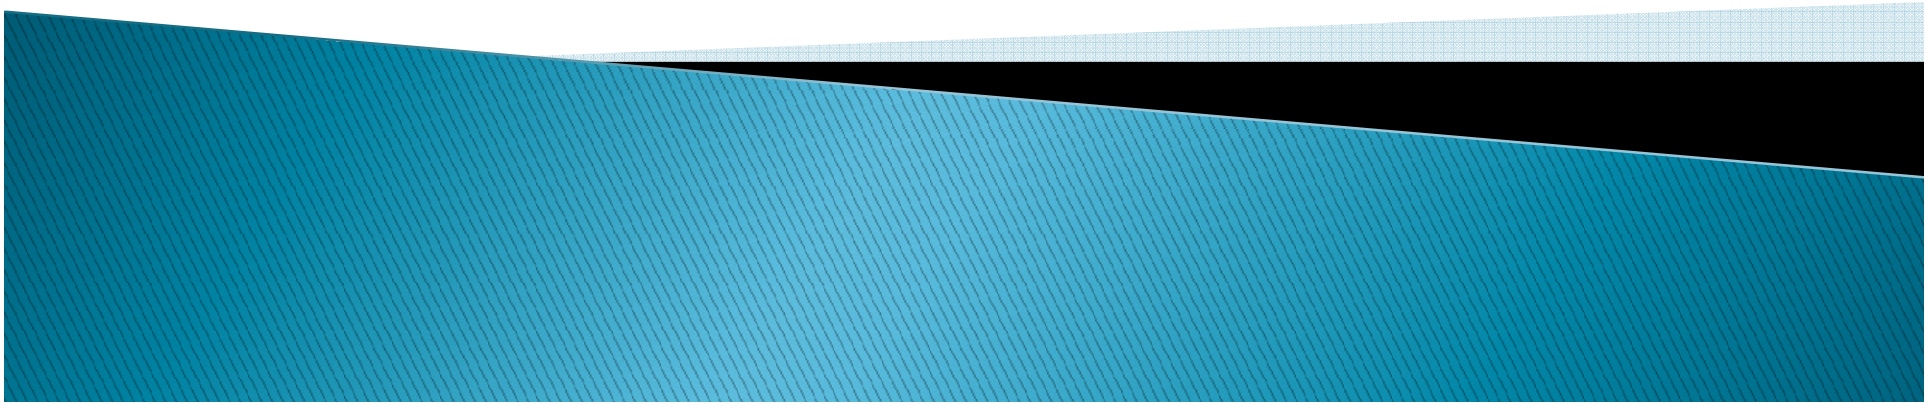

Introduction

- About me:
  - Independent consultant IPv6, Security, DNS, DHCP
  - M.Sc. from Uppsala University, Sweden, CISSP certified

  - www.scandinode.com
  - www.txv6tf.org
  - www.ipv4depletion.com
  - stephan@lagerholm.com

Top 3 reasons why migrate to IPv6

1. We are running out of IPv4 addresses.
2. There are some new cool features in IPv6.
3. Everybody else is migrating to IPv6 and we can not connect to them unless we migrate.

There is a misbelieve that IPv4 depletion is a problem for Asia and not US/EU.

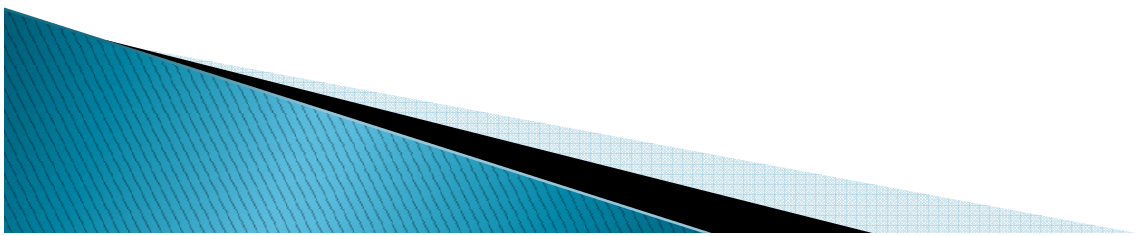

The IPv4 depletion dashboard

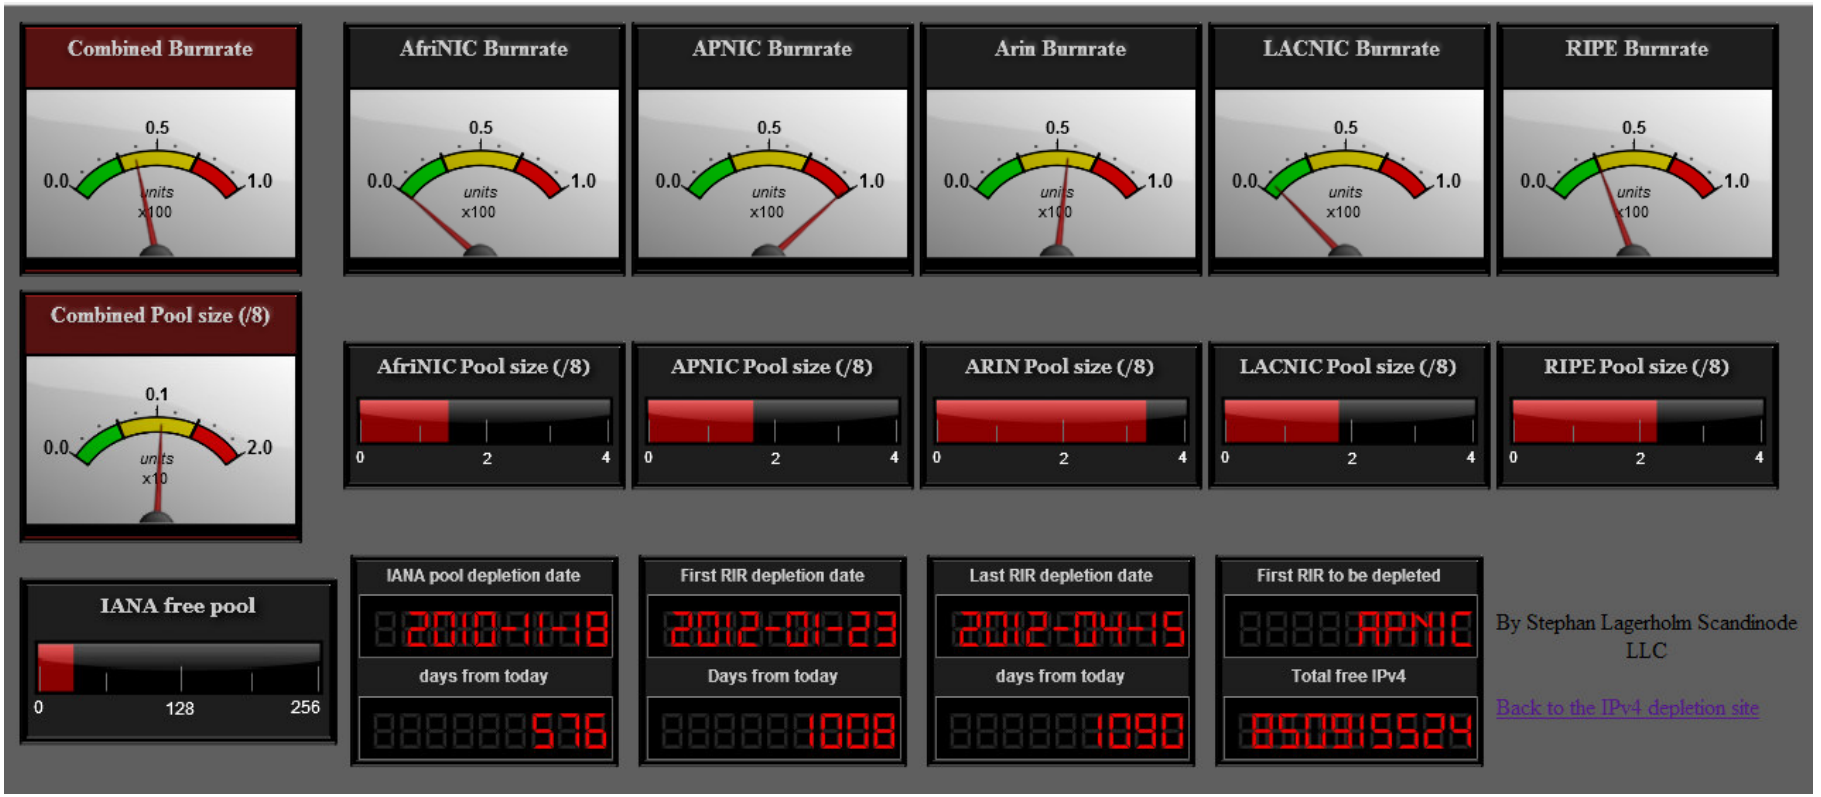

The IPv4 end game (with/without rush)

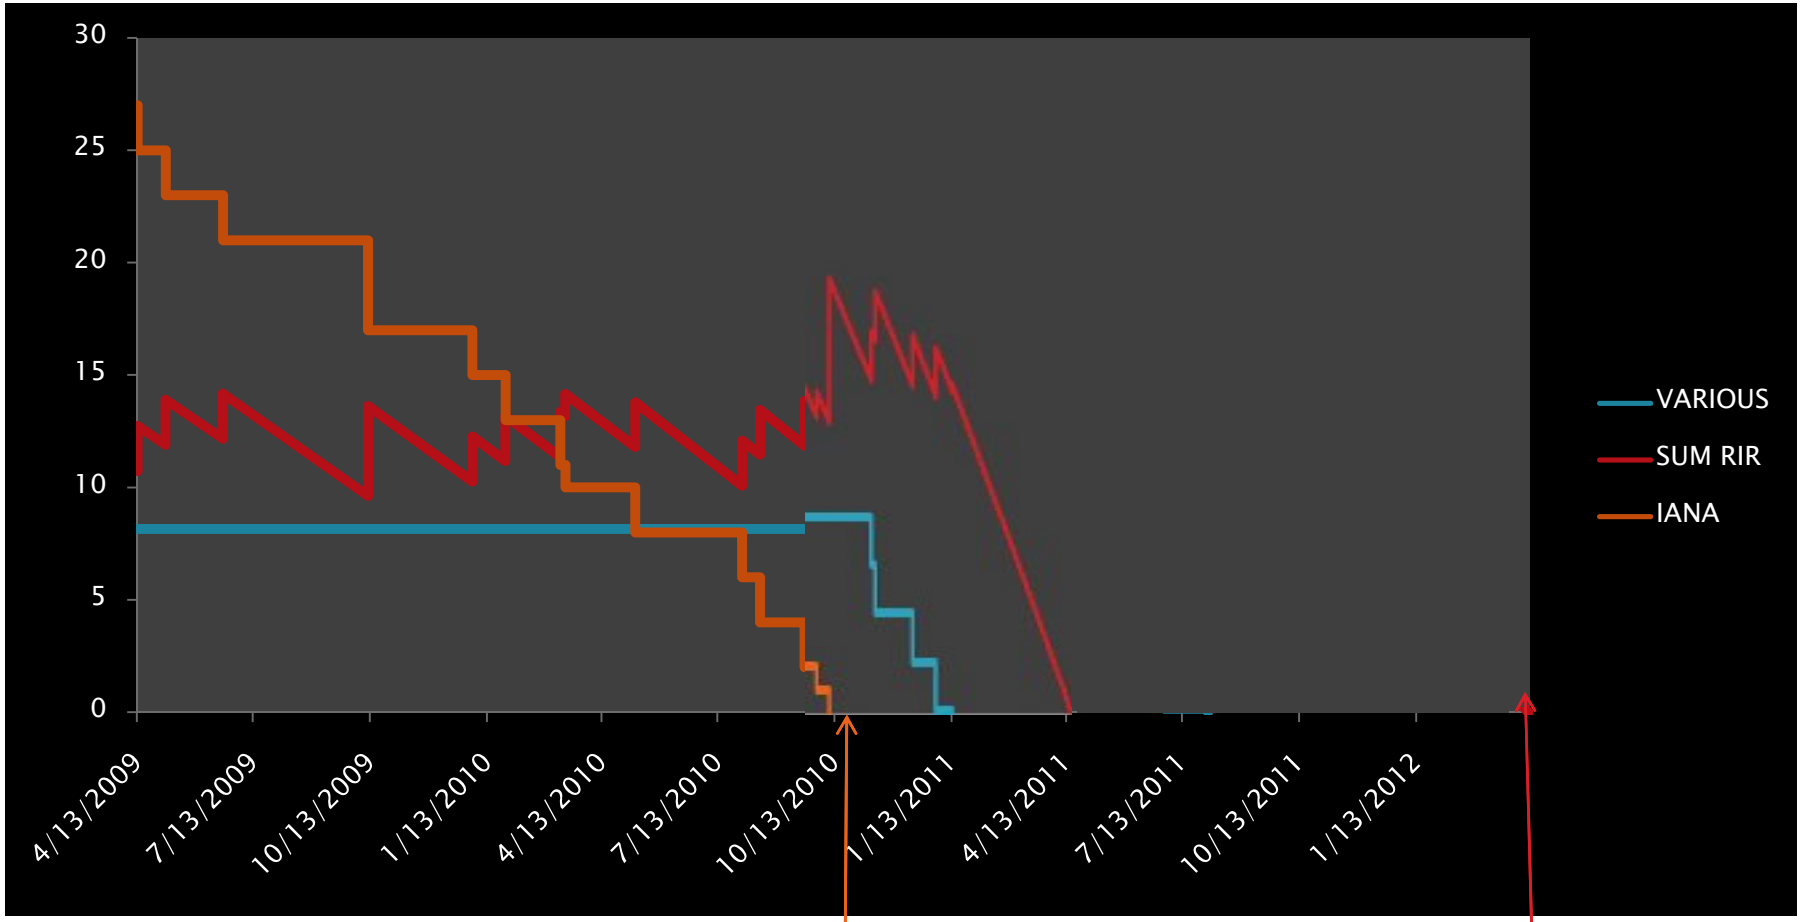

Facing the facts!

- This is a major upgrade.
- All equipment > L2 must be upgraded.
- First upgrade for most Network engineers.

- We are running out of time
- IPv6 adoption will explode at some point.
  - We are already behind.
  - The growth of IPv6 Internet doesn't require any digging.
  - The value of IPv4 will decrease quickly.

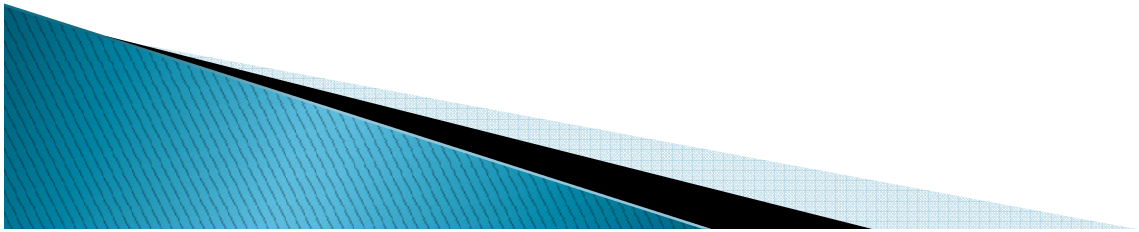

How long does it take?

- Google
  - Started slowly in March 2005.
  - Most work done in 2008.
  - Google over IPv6 launched in January 2009.
  - Alerts, Picasa, Maps and a bunch of others in March 2009.
  - Still just a handful of services.
- Bechtel
  - Started in October 2004
  - Completed migration by the end of 2008.
- Scandinode
  - One Linux Laptop, One Vista, One XP, Nintendo Wii
  - 10 hosted web domains.

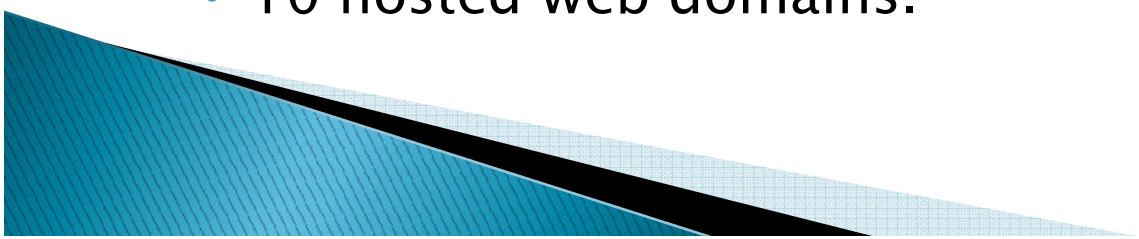

Planning for IPv6 in your LAN

- Address planning and management
- Naming (DNS)
- IP-address assignment (DHCP)
- Security

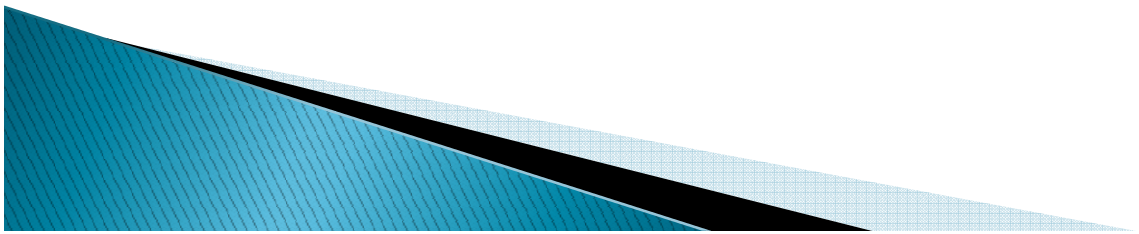

Address assignment

Noah:Maliah:Isaiah:Nariah:McCai:Josiah:Jonah:Jeremiah

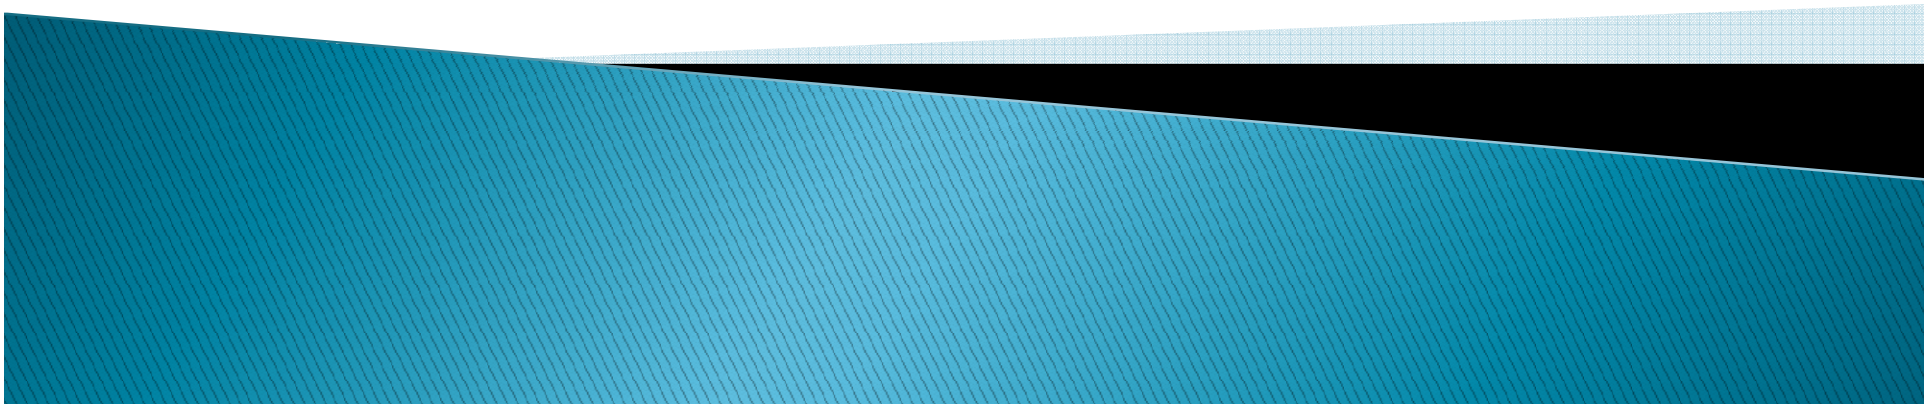

IP Address Management

- Basically the same as in IPv4 but…
- Each subnet is a million times bigger than the US deficit.
- Think “Post-it” note
- No need to reuse.
- No need to be frugal.

IPv6 subnet design

- Example
  - 2001:DB8:CAFE::/48 allocated from your ISP
  - 65k subnets to play with

Nariah group

	Possibilities	Digit position
Campus	0-F	xxxx:xxxx:xxxx:Xxxx
Building	0-F	xxxx:xxxx:xxxx:xXxx
Floor	0-F	xxxx:xxxx:xxxx:xxXx
Subnet	0-F	xxxx:xxxx:xxxx:xxxX

- 2001:DB8:CAFE:**1A40**::/64 Campus 1, Building 10 (A), Floor 4, Subnet 0
- 2001:DB8:CAFE:**42FF**::/64 Campus 4, Buildning 2, Floor 15 (F), Subnet 15 (F)

Address assignment

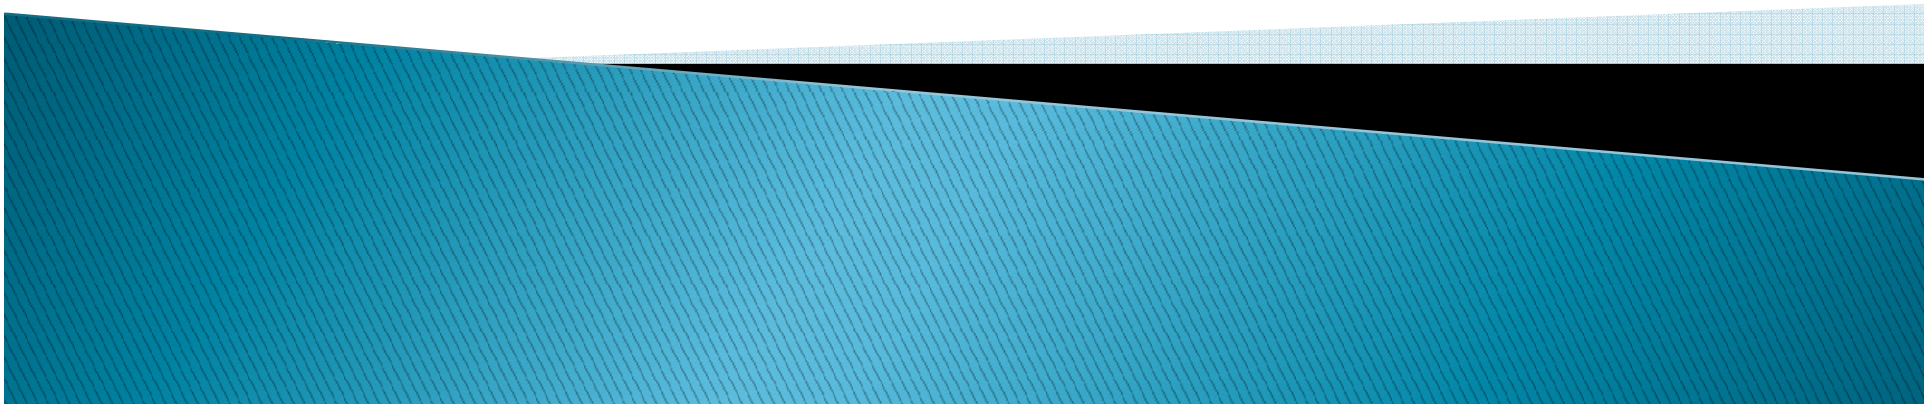

Address assignment in IPv6

- Static/Stateless/Dynamic
- SLAAC
  - Default gateway
  - no DNS server (experimental in RFC 5006)
- DHCPv6
  - New protocol, new ports (UDP 546, 547)
  - DNS server
  - no default gateway
- So, you will need them both.

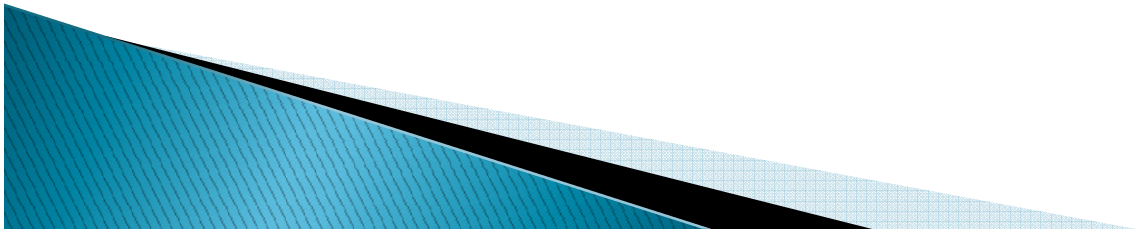

Split brain DHCP and lease times

DHCP A

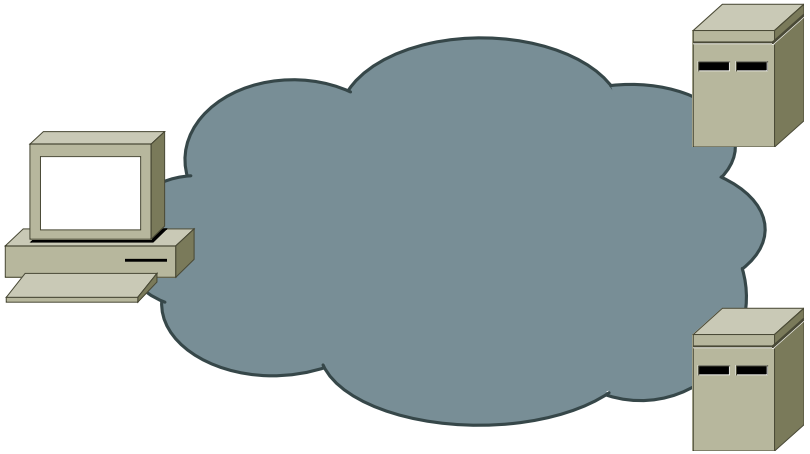

2001:db8:1:2:0000:0000:0000:0000
To
2001:db8:1:2:7FFF:FFFF:FFFF:FFFF

2001:db8:1:2:8000:0000:0000:0000
To
2001:db8:1:2:FFFF:FFFF:FFFF:FFFF

DHCP B

- Conservation of IP addresses no longer a design goal.
- Lease times can be much longer.

Rapid commit in DHCPv6

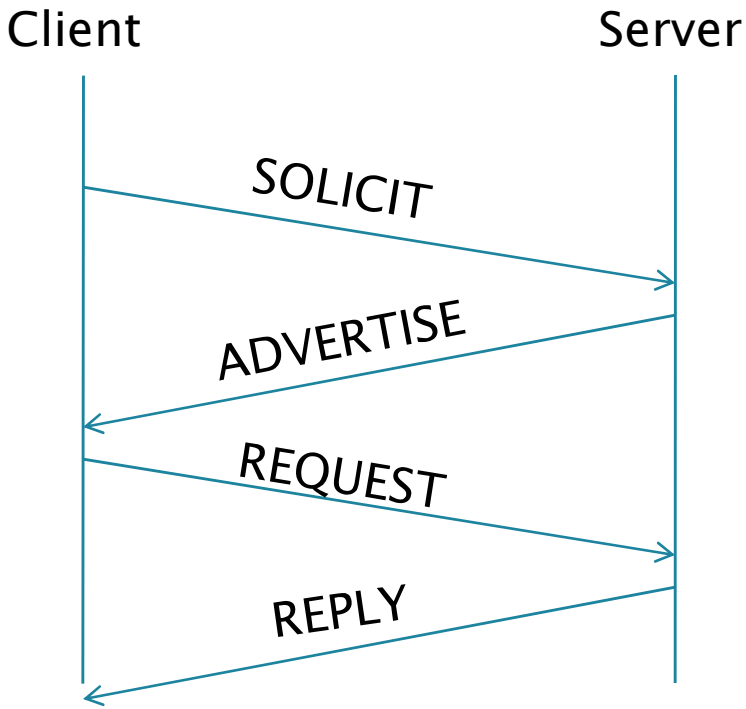

Prefix delegation

- In DHCP for IP version 4, you can only assign addresses to clients, not networks.
- In DHCPv6 you can assign networks instead of addresses with the prefix delegation feature.
- For consumers with Broadband connections this features can be used by the ISP to delegate a /48 to each customer.
- How to update the routing table is still under discussion.

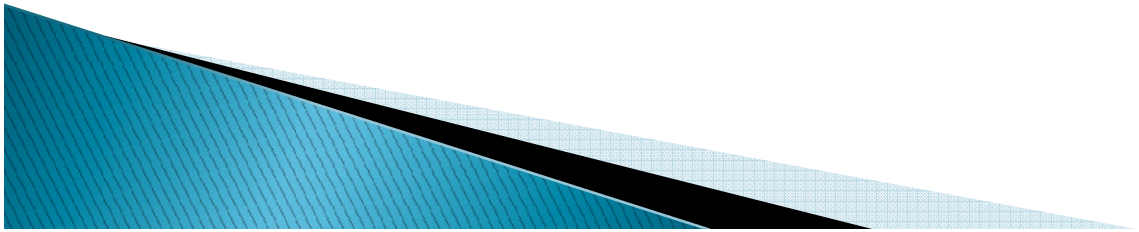

DNS

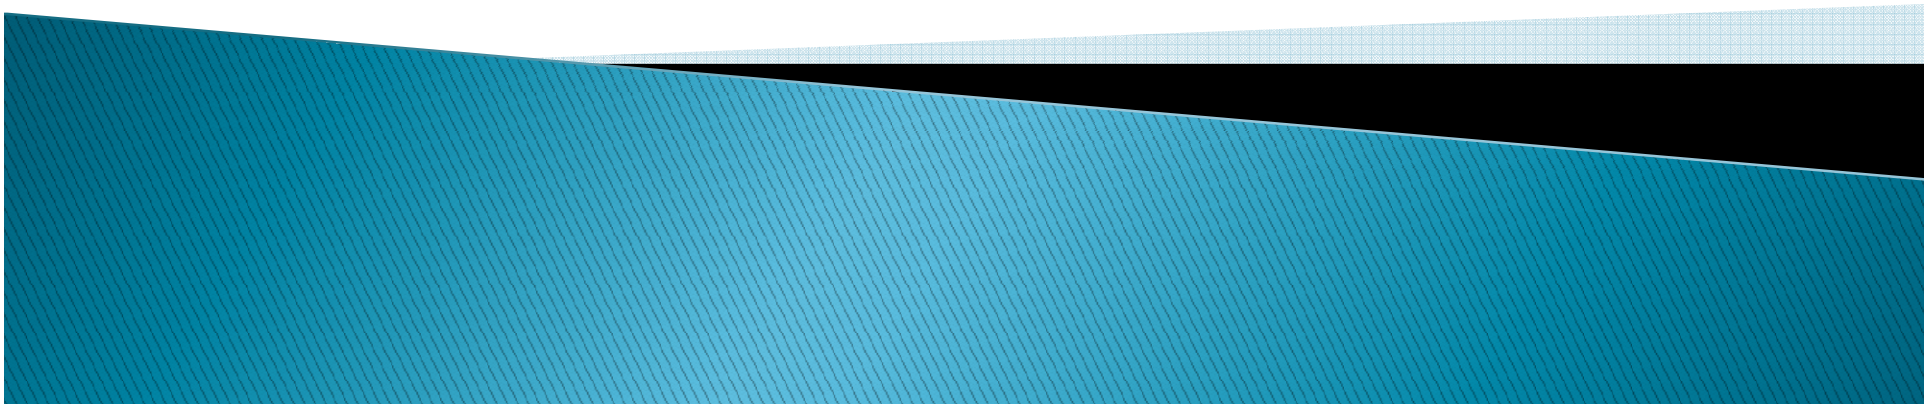

DNS in IPv6

- DNS maps a
  - hostname to an IP-address
  - An IP-address to a hostname
  - Helps hosts find other servers (Mail, AD, etc)

- DNS mandatory in IPv6, even for internal hosts, router, switches, etc.
- IPv6 is 128 bits = Impossible to manually work with.

- In the simplest form, just add AAAA records.

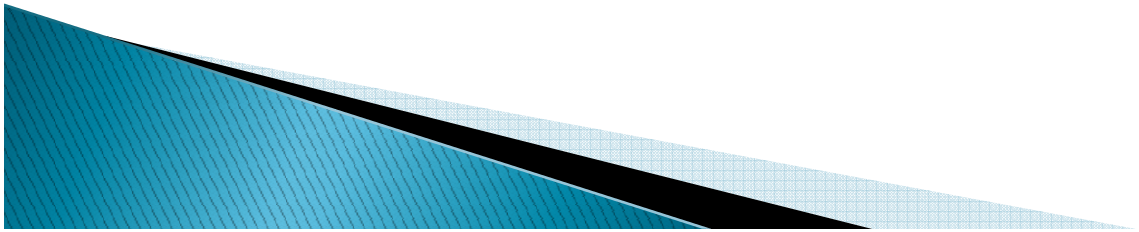

DNS and dual stack

- DNS is key for a dual stack implementation.

- The lookup of A versus AAAA records is independent of whether the DNS packets are carried over IPv4 or IPv6
- There is no assumption that the DNS servers know:
  - The IPv4/IPv6 capabilities of the requesting node.
  - The IPv4/IPv6 capabilities of the intermediate network
- Most modern OS tends to try IPv6 if there is an AAAA record.

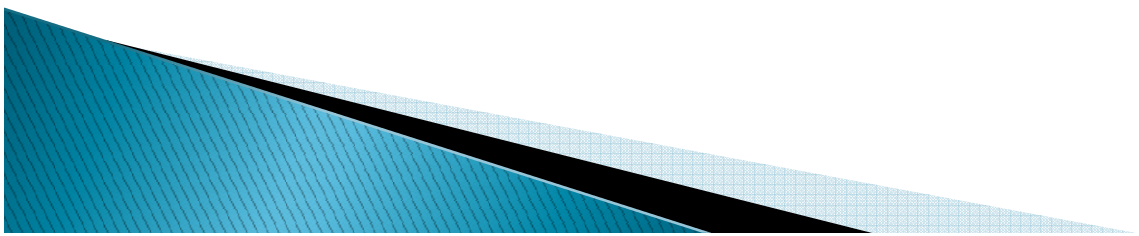

Some Security Pitfalls in IPv6

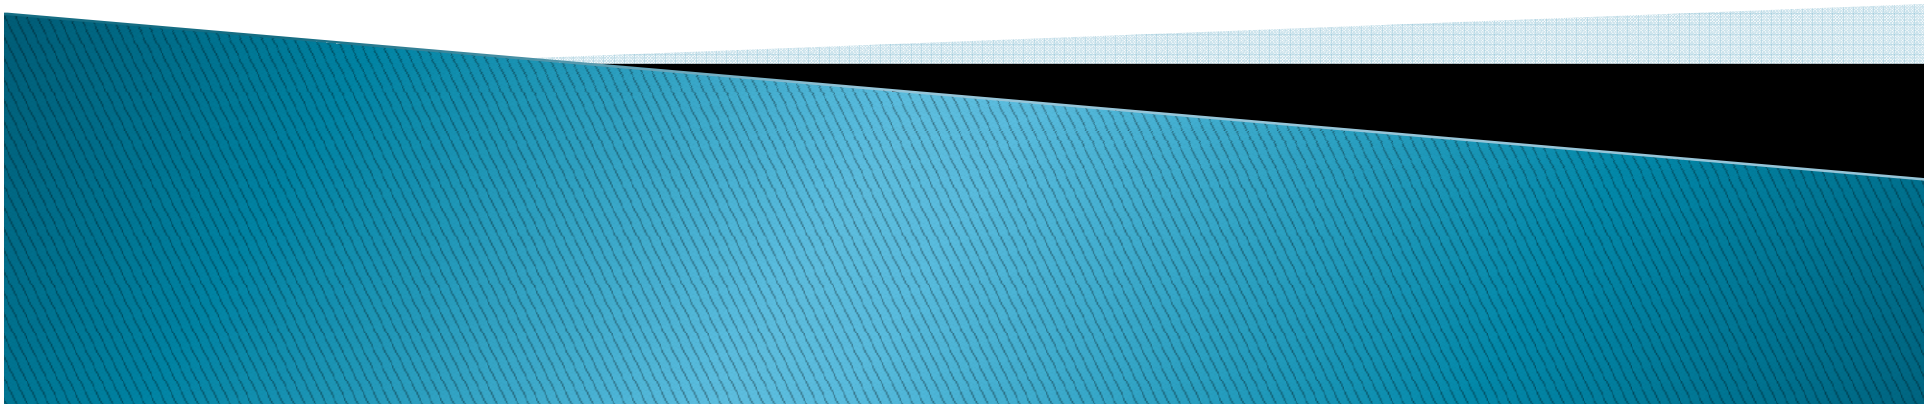

DNS and Firewalls in IPv6

- A DNS packet with IPv6 can be bigger than 512 bytes.

- DNS implemented by many firewalls as:
  - UDP 53 is used for queries
  - TCP 53 for zone transfers
- RFC says
  - UDP is used for queries less than 512 bytes
  - TCP for all other packets

- Next Idea
  - Implement EDNS0 that an extension that allowed up to 4096 bytes DNS packets.
- However, many firewalls just throws away DNS packets larger than 512 bytes.

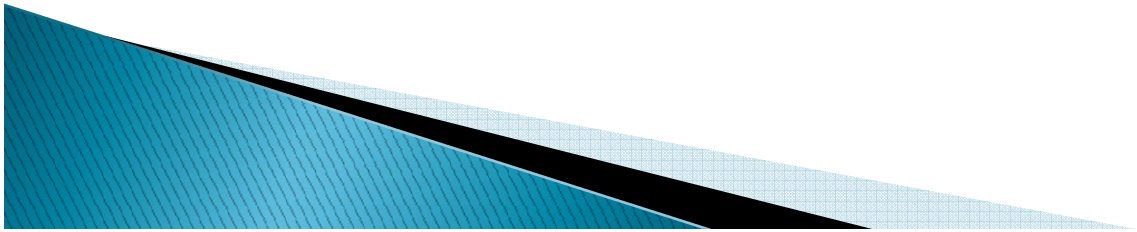

Root zone operators are being nice

```
stephan@pi:~$ dig @127.0.0.1 . ns

; <<>> DiG 9.3.4-P1.1 <<>> @127.0.0.1 . ns
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 25893
;; flags: qr rd ra; QUERY: 1, ANSWER: 13, AUTHORITY: 0, ADDITIONAL: 15

;; ADDITIONAL SECTION:
l.root-servers.net.     120942  IN      A       199.7.83.42
m.root-servers.net.     120942  IN      A       202.12.27.33
a.root-servers.net.     120942  IN      A       198.41.0.4
b.root-servers.net.     120942  IN      A       192.228.79.201
c.root-servers.net.     120942  IN      A       192.33.4.12
d.root-servers.net.     120942  IN      A       128.8.10.90
e.root-servers.net.     120942  IN      A       192.203.230.10
f.root-servers.net.     120942  IN      A       192.5.5.241
g.root-servers.net.     120942  IN      A       192.112.36.4
h.root-servers.net.     120942  IN      A       128.63.2.53
i.root-servers.net.     120942  IN      A       192.36.148.17
j.root-servers.net.     120942  IN      A       192.58.128.30
k.root-servers.net.     120942  IN      A       193.0.14.129
l.root-servers.net.     120942  IN      AAAA    2001:500:3::42
m.root-servers.net.     120942  IN      AAAA    2001:dc3::35

;; Query time: 2 msec
;; MSG SIZE  rcvd: 492
```

Size 492 without EDNS0

```
stephan@pi:~$ dig +bufsize=4096 @127.0.0.1 . ns

; <<>> DiG 9.3.4-P1.1 <<>> +bufsize=4096 @127.0.0.1 . ns
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 21599
;; flags: qr rd ra; QUERY: 1, ANSWER: 13, AUTHORITY: 0, ADDITIONAL: 21

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; ADDITIONAL SECTION:
i.root-servers.net.     121031  IN      A       192.36.148.17
j.root-servers.net.     121031  IN      AAAA    2001:503:c27::2:30
j.root-servers.net.     121031  IN      A       192.58.128.30
k.root-servers.net.     121031  IN      AAAA    2001:7fd::1
k.root-servers.net.     121031  IN      A       193.0.14.129
l.root-servers.net.     121031  IN      AAAA    2001:500:3::42
l.root-servers.net.     121031  IN      A       199.7.83.42
m.root-servers.net.     121031  IN      AAAA    2001:dc3::35
m.root-servers.net.     121031  IN      A       202.12.27.33
a.root-servers.net.     121031  IN      AAAA    2001:503:ba3e::2:30
a.root-servers.net.     121031  IN      A       198.41.0.4
d.root-servers.net.     121031  IN      A       128.8.10.90
e.root-servers.net.     121031  IN      A       192.203.230.10
f.root-servers.net.     121031  IN      AAAA    2001:500:2f::f
g.root-servers.net.     121031  IN      A       192.112.36.4
h.root-servers.net.     121031  IN      AAAA    2001:500:1::803f:235
h.root-servers.net.     121031  IN      A       128.63.2.53

;; Query time: 2 msec
;; MSG SIZE  rcvd: 643
```

Size 643 with EDNS0

Path MTU discovery

- IPv6 uses a different method to figure out the Maximum Transfer Unit on a path.
- ICMPv6 “Packet too big” sent from the intermediate router back to the sender.

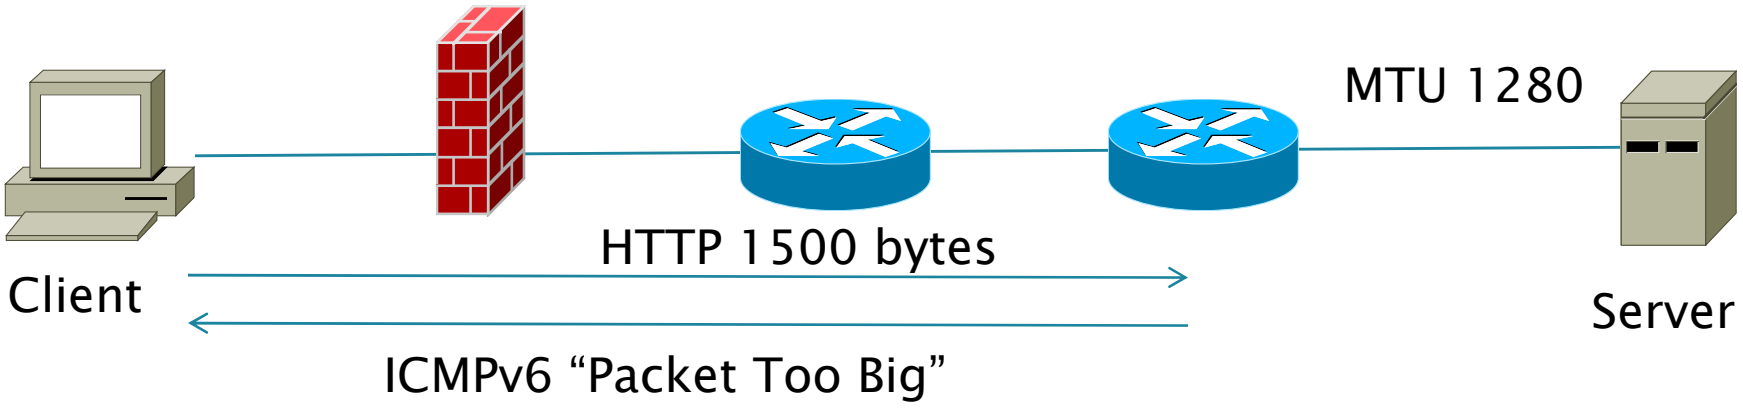

Path MTU discovery

- Need to open a src:ANY rule in the firewall.
  - Typically requires: IPS, OS hardening, controls of patching etc.
- There is a payload field in the ICMP packet.
- Very hard to secure this for all internal nodes.
- pMTU might not happen, perhaps just use 1280 bytes.

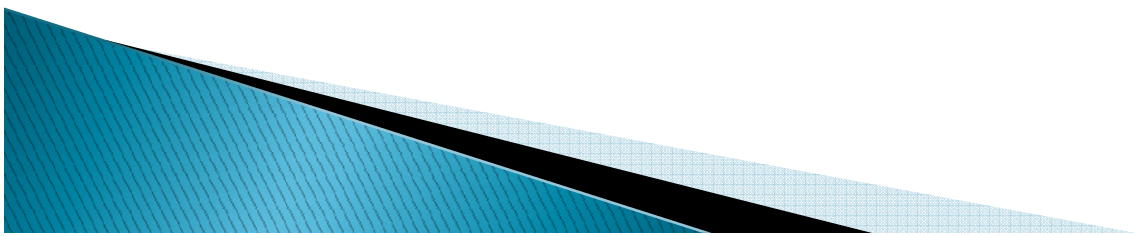

IPSEC

- IPSEC Mandatory in IPv6 (Optional in IPv4).
- This does not mean that you can start encrypting sessions left and right.
- Can only do this when encryption is solved.

- Direct Access in Windows 7.

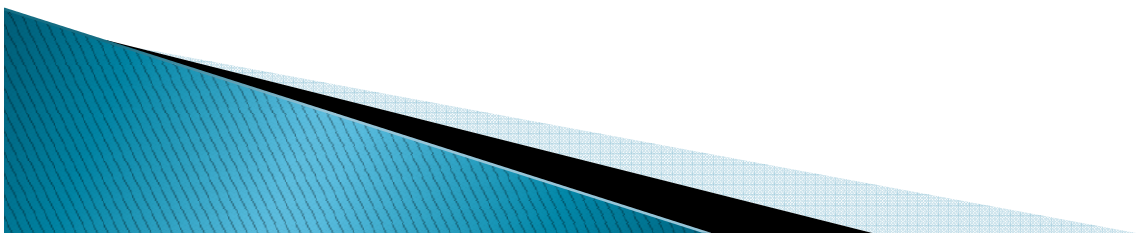

The IPv6 ready program

Questions ???

- www.scandinode.com
- www.txv6tf.org
- www.ipv4depletion.com
- stephan@lagerholm.com