ntt.net

NTT Communications

IPv6 Transition Mechanisms and Strategies

Chuck Sellers

Senior Product Engineer, CISSP

April 21, 2009

IPv6 Transition Mechanisms

Dual Stack (RFC 4213, 2893):

- IPv4 and IPv6 coexist on a host/node
- Enables a node to communicate with IPv6-only or IPv4-only nodes concurrently

Tunneling (Encapsulation) :

- Enables IPv6 islands or individual nodes to communicate over an IPv4 network
- Two modes of tunneling support are possible:

  Configured (Manual) Tunneling (RFC 4213, 2893)

  Automatic Tunneling (several types)

  - IPv4 compatible IPv6 -RFC 4213, 2893
  - IPv6 over IPv4 - RFC 2529
  - IPv6 to IPv4 – RFC 3056,
  - ISATAP (Intra-Site Automatic Tunnel Addressing Protocol)
  - Tunnel Broker – RFC 3053
  - Teredo (Microsoft proponent)

Translation:

- Enables IPv4-only networks & nodes to communicate with IPv6-only networks & nodes
- Translation at transport layer or IP layer
- Translation at application-layer gateways
- Includes IPv6 stateless address auto-configuration

DUAL STACK

Dual Stack

Description:

- Deployment and utilization of IPv4 and IPv6 concurrently
- v4 address: 123.234.20.1
- v6 address: 2001:3f0:4c02:2:320:a6ff:fe4c:350f

Benefits:

- Can be deployed on hosts, routers, on same interface as IPv4
- Deals with most address selection and DNS resolution issues
- Allows hosts to continue to reach IPv4 resources, while also adding IPv6 functionality
- Simple to deploy. Allows backwards compatibility (IPv4 support)
- Available on most platforms. Easy to use, flexible.

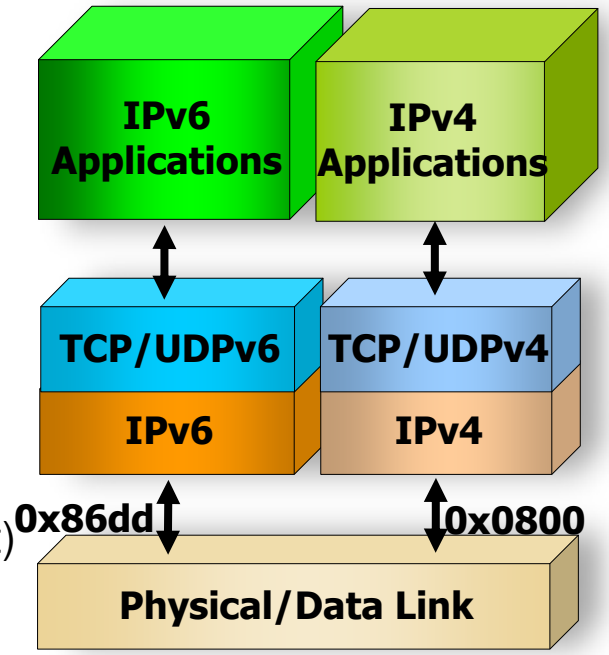

Issues:

- May require 2 routing tables & routing processes
- Additional CPU, memory
- IPv6 network security requirements are same as IPv4 networks today – don't overlook enforcing security on parallel protocol

Application Deployment:

- Easy way to deploy, flexible, bring up 2nd protocol in parallel with first

Cisco Dual Stack Example

```
ipv6 unicast-routing
ipv6 multicast-routing
ipv6 cef
!
interface Loopback0
 ip address 200.100.1.3 255.255.255.255
 ipv6 address 2001:db8:10:10::10/128
!
interface Ethernet 0
 ip address 192.168.100.1 255.255.255.0
 ipv6 address 2001:db8:1:1::1/64
!
ipv6 route ::/0 2001:db8:1:1::100
```

Tunneling (Encapsulation)

Configured (Manual) Tunnels

Description – RFC 4213, 2893:

- Tunneling allows IPv6 traffic to be carried across an IPv4 network
- Tunnel destination address is specified in the tunnel source configuration creating a P2P topology
- The tunnel acts as 1 hop for a IPv6 packet whereas an IPv4 encapsulation packet may take many hops

Benefits:

- Simple to deploy
- Allows transport of IPv6 packets over an IPv4 network
- Available on most platforms

Issues:

- Must be manually configured
- Potential (unknown) issues with delay and latency through the tunnel
- Additional CPU load for encapsulation/de-capsulation
- Single Point of Failure (tunnel endpoint)

Deployment Applications:

- Cost-effective method of obtaining IPv6 connectivity
- NTT commercial product/service, LONG Network (experimental)

Cisco Configuration Example: 6to4 Tunnel

```
hostname Router1
ipv6 unicast-routing
ipv6 multicast-routing
ipv6 cef
interface Tunnel 0
 ipv6 address 2001:db8:100:1::1/126
 tunnel source 192.168.100.1
 tunnel destination 192.168.200.2
 tunnel mode ipv6ip
ipv6 route 2001:db8:c:1::/64 tunnel0


hostname Router2
ipv6 unicast-routing
ipv6 multicast-routing
ipv6 cef
interface Tunnel 0
 ipv6 address 2001:db8:100:1::2/126
 tunnel source 192.168.200.2
 tunnel destination 192.168.100.1
 tunnel mode ipv6ip
ipv6 route 2001:db8:d:2::/64 tunnel0
```

Juniper Configuration Example: GRE Tunnel

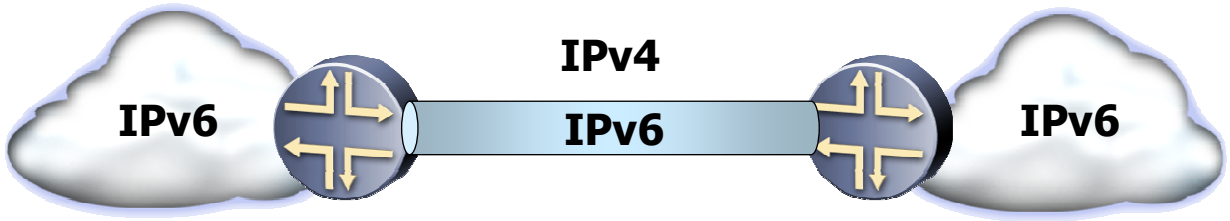


```
gr-0/0/0 {
    unit 0 {
        tunnel {
            source 172.16.1.1;
            destination 192.168.2.3;
        }
        family inet6 {
            address 2001:240:13::1/126;
        }
    }
}
```

```
gr-1/0/0 {
    unit 0 {
        tunnel {
            source 192.168.2.3;
            destination 172.16.1.1;
        }
        family inet6 {
            address 2001:240:13::2/126;
        }
    }
}
```

Automatic Tunneling: IPv4 Compatible IPv6

Description - RFC 4213, 2893:

- Automatic tunneling allows IPv6 traffic to be carried across an IPv4 infrastructure without the need for tunnel destination pre-configuration (P2MP)
- v4 address: 123.234.20.1
- v6 address: 0:0:0:0:0:123:234:20:1 or ::123:234:20:1

Benefits:

- Simple to deploy – no pre-configuration required
- Can use BOOTP, DHCP, RARP or manual configuration to obtain IPv4 address
- Allows transport of IPv6 packets over an IPv4 network using v4 tunnel endpoints
- Available on most platforms such as Cisco IOS and Microsoft XP

Issues:

- Requires a globally unique IPv4 address (No NAT tunnel endpoint allowed)
- Tunnel must not send IPv4 packets to: broadcast, multicast, loopback addresses
- 0:0:0:0:0:0/96 static route is required for automatic tunneling

Deployment Applications:

- A potentially cost-effective method of obtaining IPv6 connectivity
- Not recommended – currently being deprecated

Automatic Tunneling – IPv6 Over IPv4 (6OVER4)

Description – RFCs 4213, 2893, 2529:

- Allows isolated IPv6 hosts to become fully functional v6 hosts without direct v6 connectivity on a IPv4 physical link
- Often used in conjunction with configured tunneling: embedding the nodes IPv4 address into an IPv6 address 123.234.20.1 becomes 0:0:0:0:0:0:123:234:20:1
- Site local using IPv4 multicast as virtual link layer

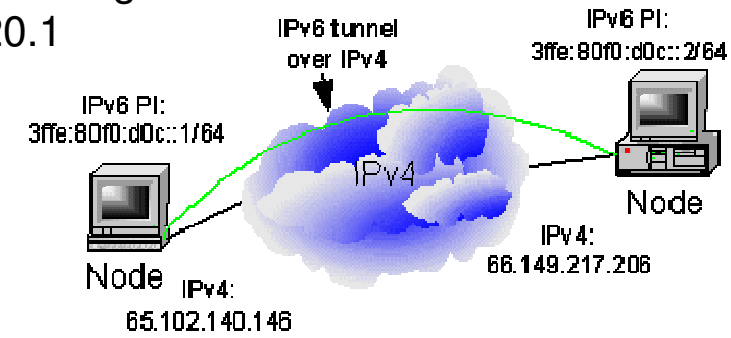

tunnel example

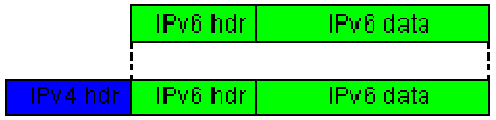

encapsulation example

Benefits:

- Simple to deploy
- Allows transport of IPv6 packets over an IPv4 network
- Available on most platforms
- Existing standard (RFC 2373 and RFC 2529)

Issues:

- Not Scalable
- 0:0:0:0:0:0/96 static route is required
- IPv6 multicast is implemented over IPv4 multicast

Deployment Applications:

- Used rarely

Juniper Configuration Example: MPLS Tunnel

PE Router:

```
mpls {
    ipv6-tunneling;
    label-switched-path v6-tunnel1
{
        to 192.168.2.3;
        no-cspf;
    }
}
bgp {
    group IPv6-neighbors {
        type internal;
        family inet6 {
            labeled-unicast {
                explicit-null;
            }
        }
        neighbor 192.168.2.3;
    }
}
```

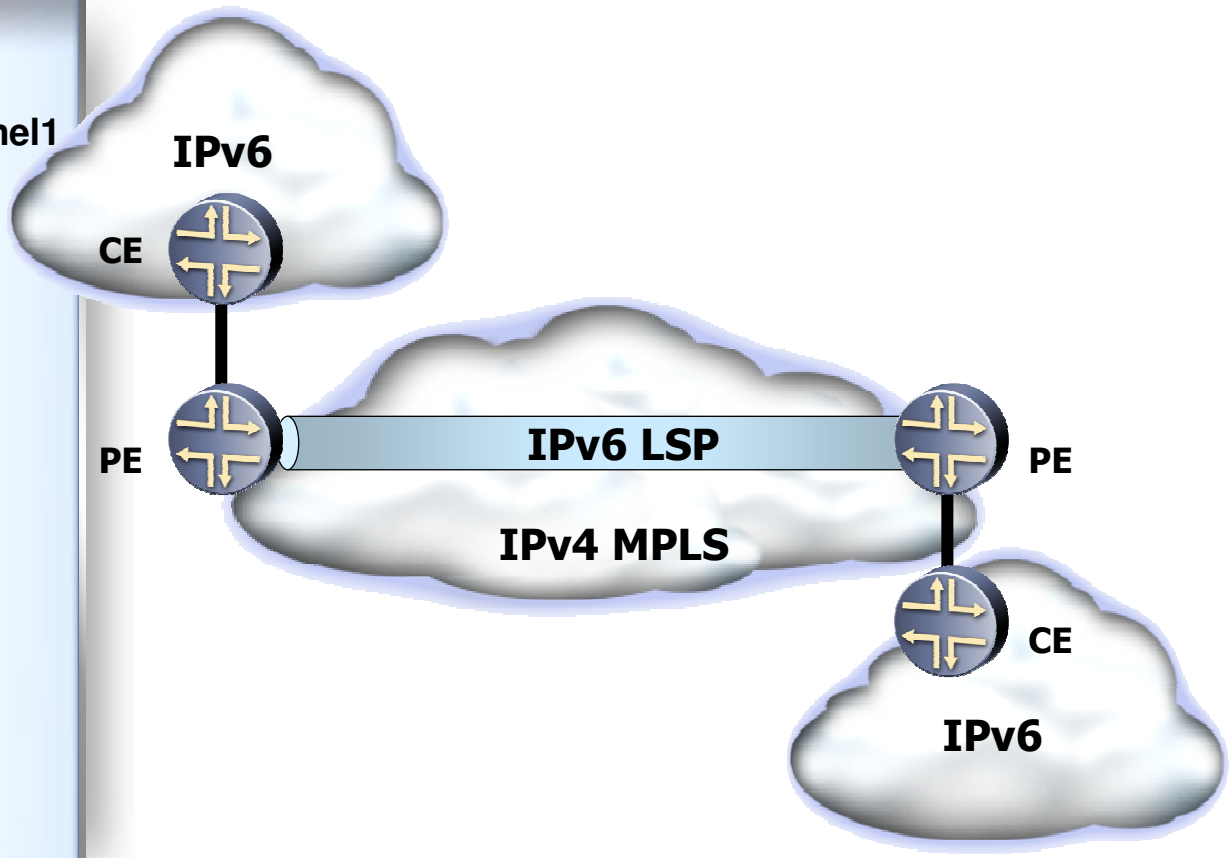

Automatic Tunneling – IPv6 to IPv4 (6to4)

Description – RFC 3056, 3068:

- Stateless automatic tunneling
- Encapsulation of IPv6 addresses automatically into IPv4 address (P2MP)
- IANA defined 2002:: /16 (improvement over use of IPv4 addresses)
- 123.234.20.1 becomes 2002:7bea:1401:: /48 (6to4 32-bit IPv4 GW address)

Benefits:

- Extremely easy for IPv6 "islands" located in IPv4 network to communicate
- Provides for 65,536 /64 networks with 2^{64} nodes/network
- Creates a globally unique /48 IPv6 prefix for use within the AS
- Good for IPv6 domains/sites with no IPv6 support

Issues:

- Requires one globally unique IPv4 address (No NAT) and 6to4 relay router
- Scales well for sites, NOT for individual hosts
- Number of problems remain for communication between an isolated IPv6 network and the IPv6 Internet
- Tunnel must not send IPv4 packets to: Broadcast, multicast, loopback addresses

Deployment Applications:

- 6Bone, switch.ch – gateway to other IPv6 clouds, APANA Melbourne, Australia

6to4

- Designed for site-to-site and site to existing IPv6 network connectivity
- Site border router must have at least one globally-unique IPv4 address
- Uses IPv4 embedded address

Example:

Reserved 6to4 TLA-ID:	**2002::/16**
IPv4 address:	**138.14.85.210 = 8a0e:55d2**
Resulting 6to4 prefix:	**2002:8a0e:55d2::/48**

- Router advertises 6to4 prefix to hosts via RAs
- Embedded IPv4 address allows discovery of tunnel endpoints

Automatic 6to4 Tunneling

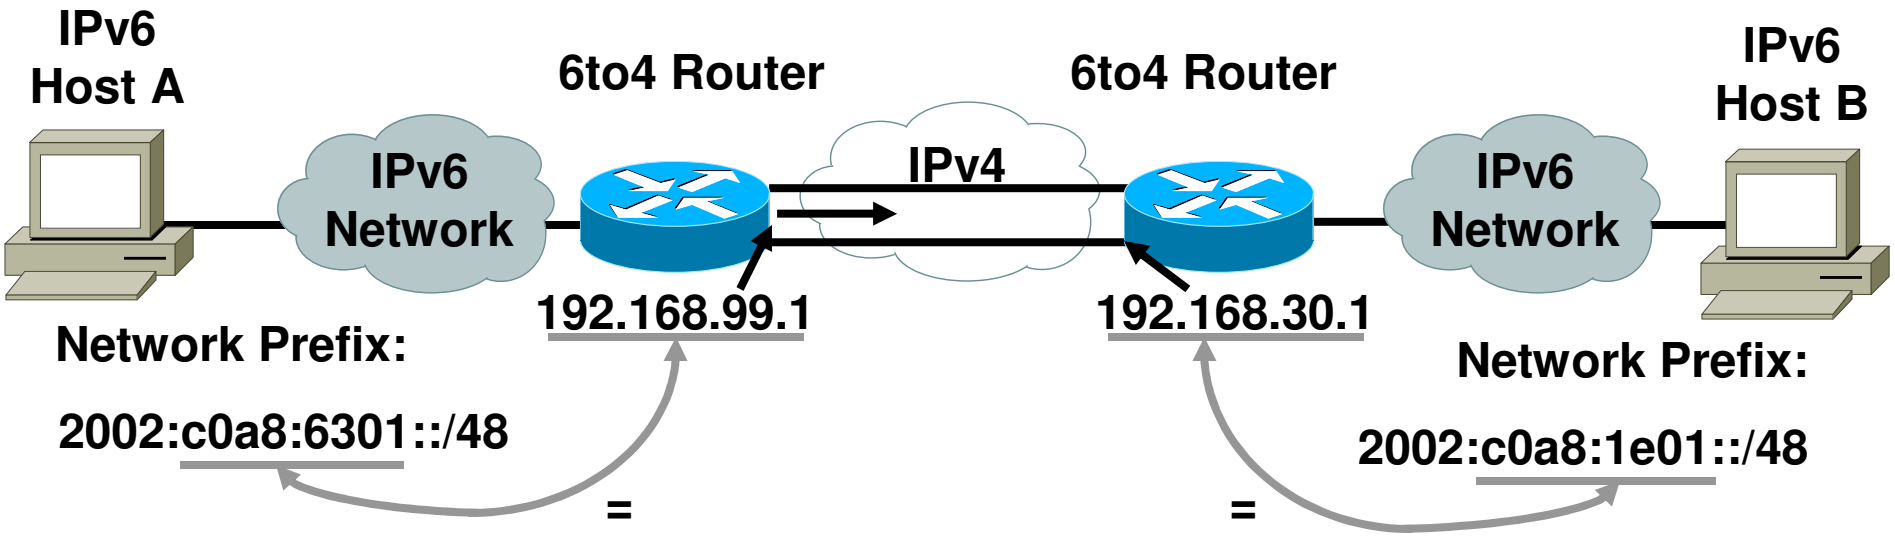

- 6to4:
  - Is an automatic tunnel method
  - Gives a prefix to the attached IPv6 network

Cisco Configuration Example: 6 to 4

```
hostname BorderRouter
interface Ethernet0
 ip address 200.168.100.1 255.255.255.0
interface Tunnel0
 no ip address
 ipv6 address 2002:c8a8:6401:1::1/128
 tunnel source Ethernet0
 tunnel mode ipv6ip 6to4
ipv6 route 2002::/16 Tunnel0
ipv6 route ::/0 2002:c8a8:c802:2::2
```

```
hostname 6to4RelayRouter
interface Ethernet0
 ip address 200.168.200.2 255.255.255.0
interface Tunnel0
 no ip address
 ipv6 address 2002:c8a8:c802:2::2/128
 tunnel source Ethernet0
 tunnel mode ipv6ip 6to4
ipv6 route 2002::/16 Tunnel0
```

6to4: Imbedding Tunnel Endpoints

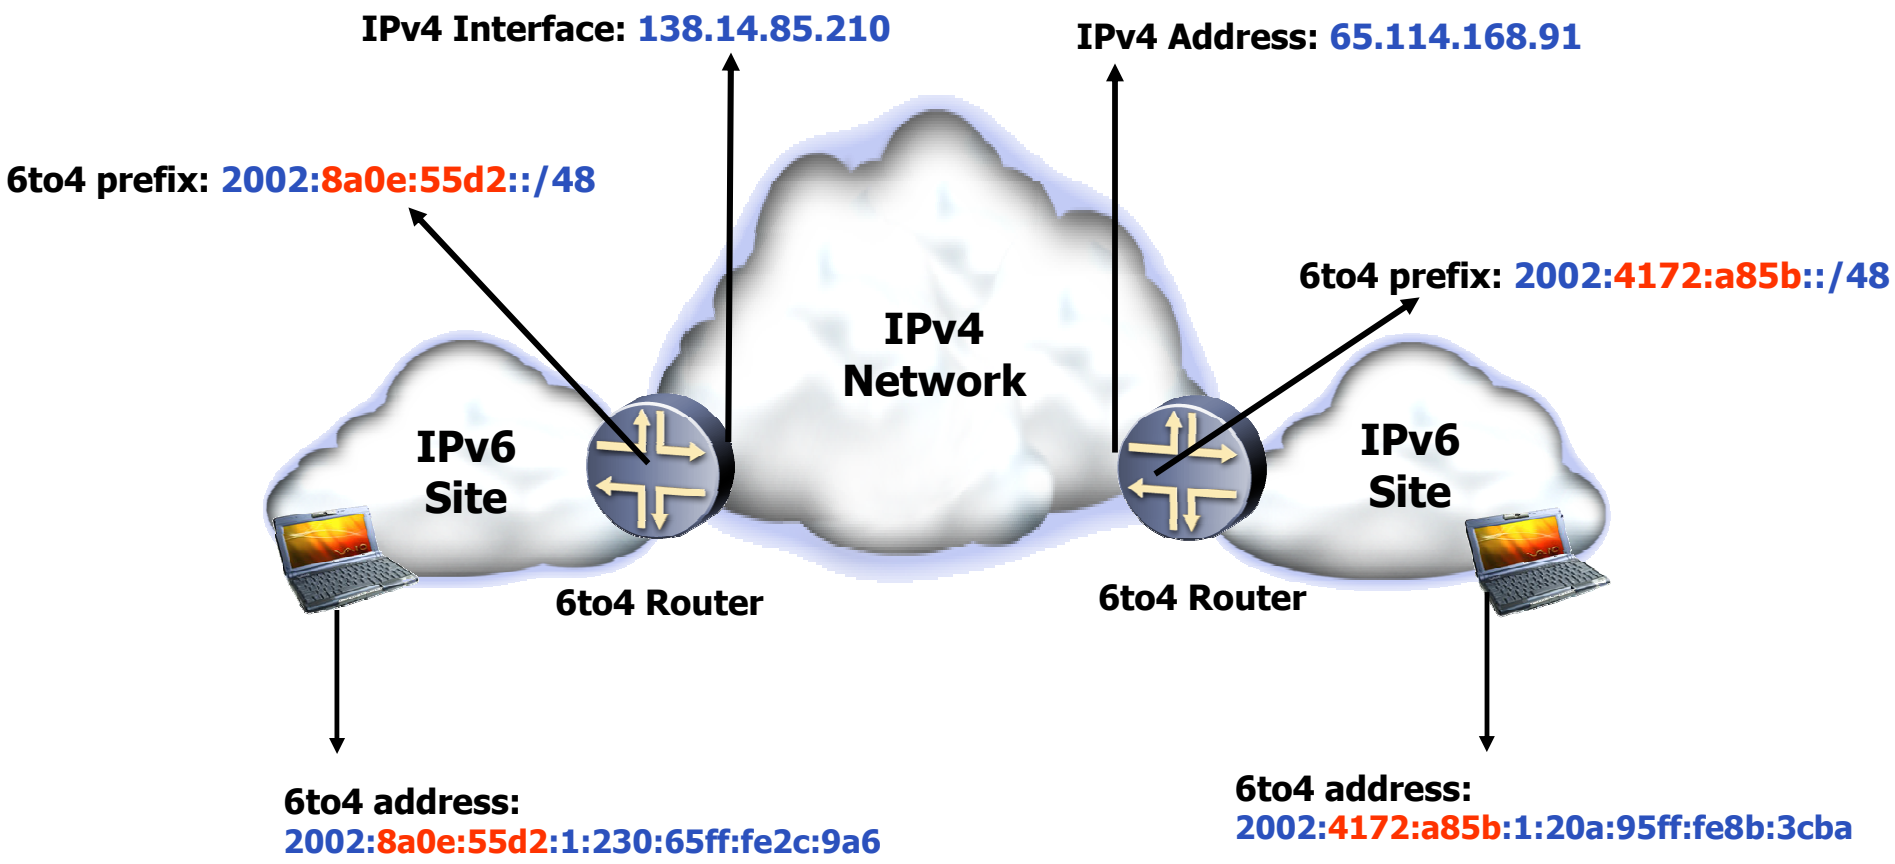

6to4: Tunnel Setup

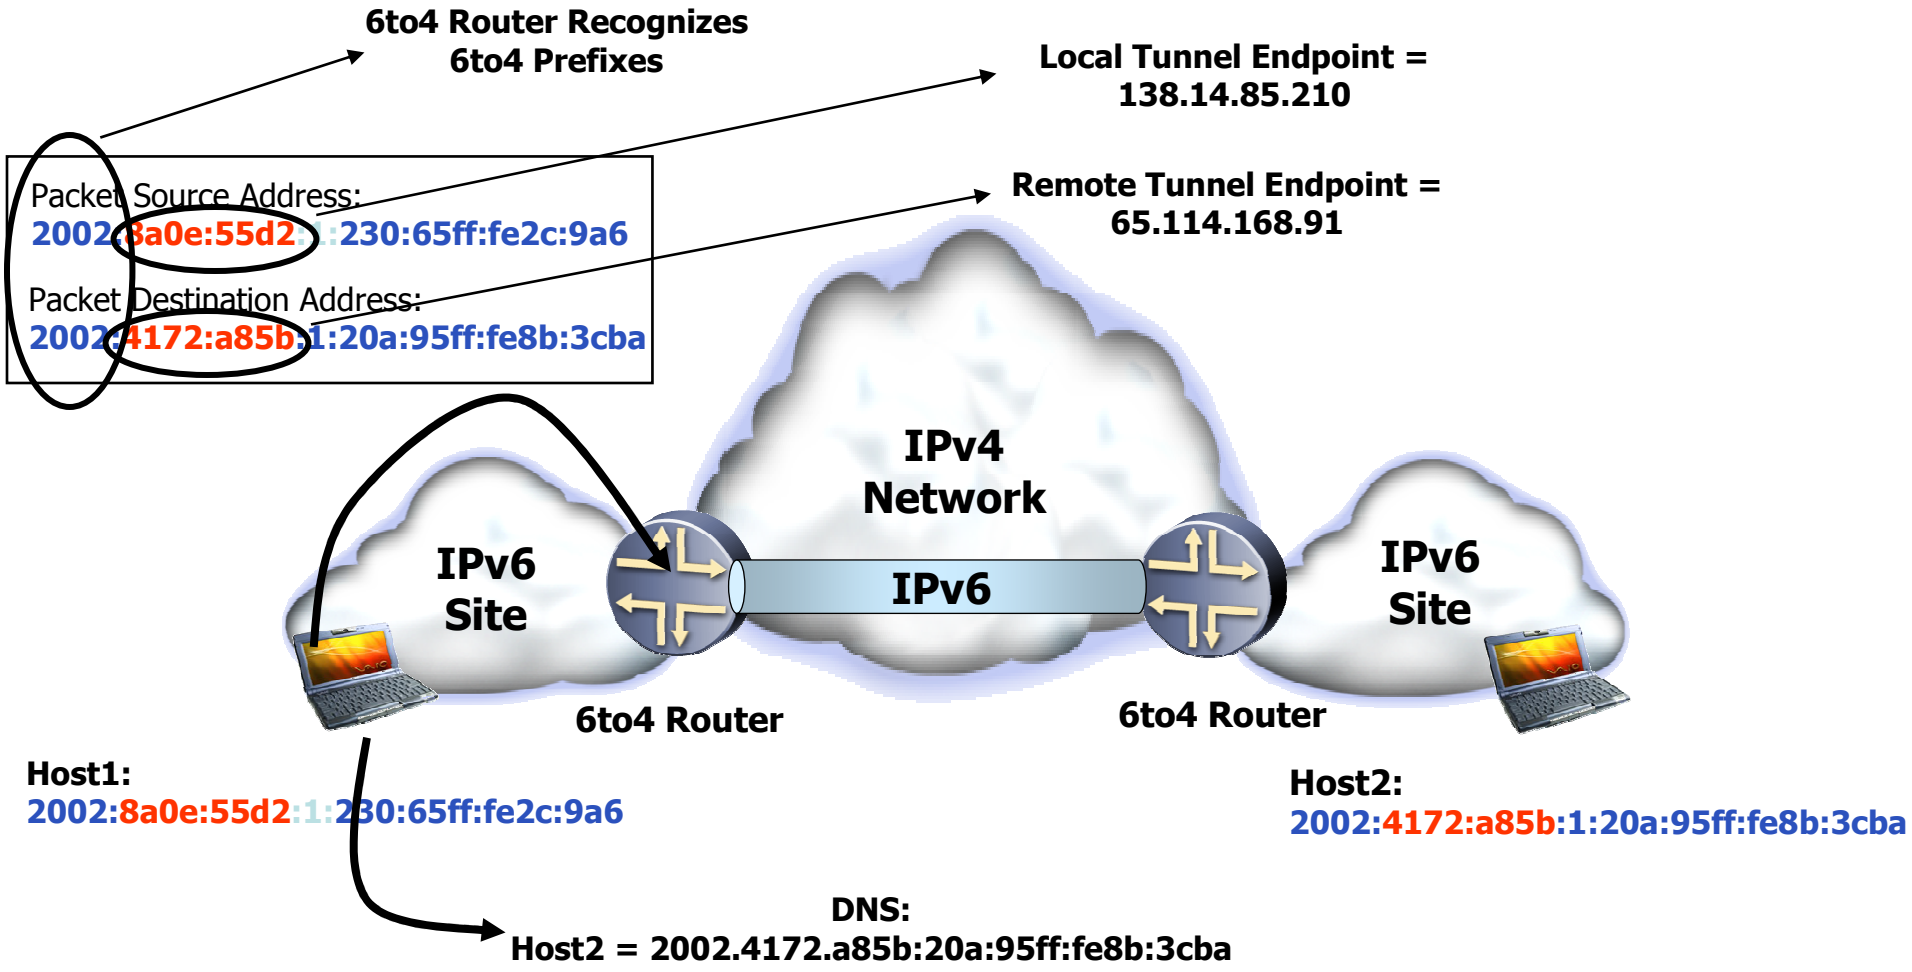

Configuration Example: Windows XP 6to4 Interface

```
C:\Documents and Settings\c.sellers>ipv6 if 3
Interface 3: 6to4 Tunneling Pseudo-Interface
  does not use Neighbor Discovery
  does not use Router Discovery
    preferred global 2002:4172:a85b::4172:a85b, life infinite
  link MTU 1280 (true link MTU 65515)
  current hop limit 128
  reachable time 23000ms (base 30000ms)
  retransmission interval 1000ms
  DAD transmits 0
```

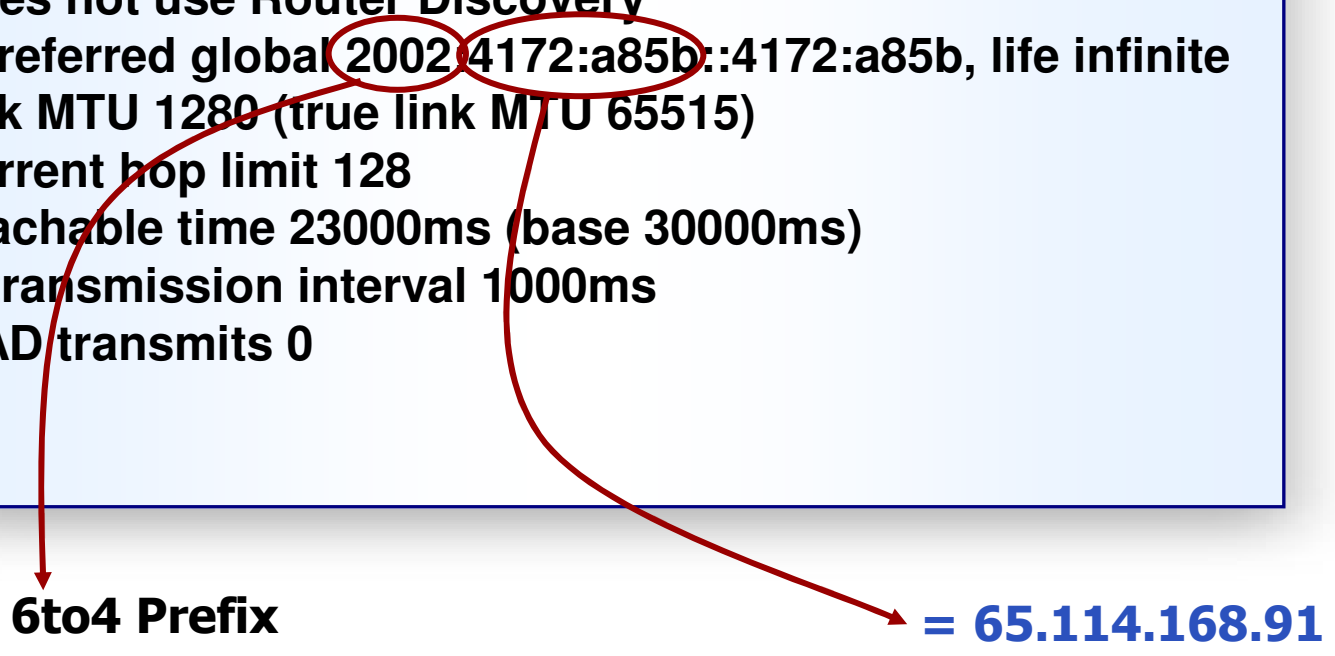

Automatic Tunneling – ISATAP

Description (Intra-Site Automatic Tunnel Addressing Protocol) RFC 5214, 4214:

- ISATAP embeds IPv4 gateway address in the EUI-64 interface identifier
- 123.234.20.1 becomes 2001:abcd:1234:5678:0000:5efe:7bea:1401
- Globally unique IPv6 address created via auto-configuration and DNS lookup
- Combines dual stack and automatic tunneling techniques
- Views IPv4 network as link layer for IPv6, views other nodes on network as potential IPv6 hosts/routers

Benefits:

- Easy incremental deployment of IPv6 to disparate nodes within AS (intra-site)
- Supported on many platforms

Issues:

- Can require more setup than other methods
- RFC draft status
- Some security issues
- Designed for Intra-site use, not Inter-site connectivity

Deployment Applications:

- Unknown

ISATAP

- Forms 64-bit Interface ID from IPv4 address + special reserved identifier
  - Format: ::0:5efe:W.X.Y.Z
  - 0:5efe = 32-bit IANA-reserved identifier
  - W.X.Y.Z = IPv4 address mapped to last 32 bits

Example:

IPv4 address:	**65.114.168.91**
Global IPv6 prefix:	**2001:468:1100:1::/64**
Link-local address:	**fe80::5efe:65.114.168.91**
Global IPv6 address:	**2001:468:1100:1::5efe:65.114.168.91**

ISATAP

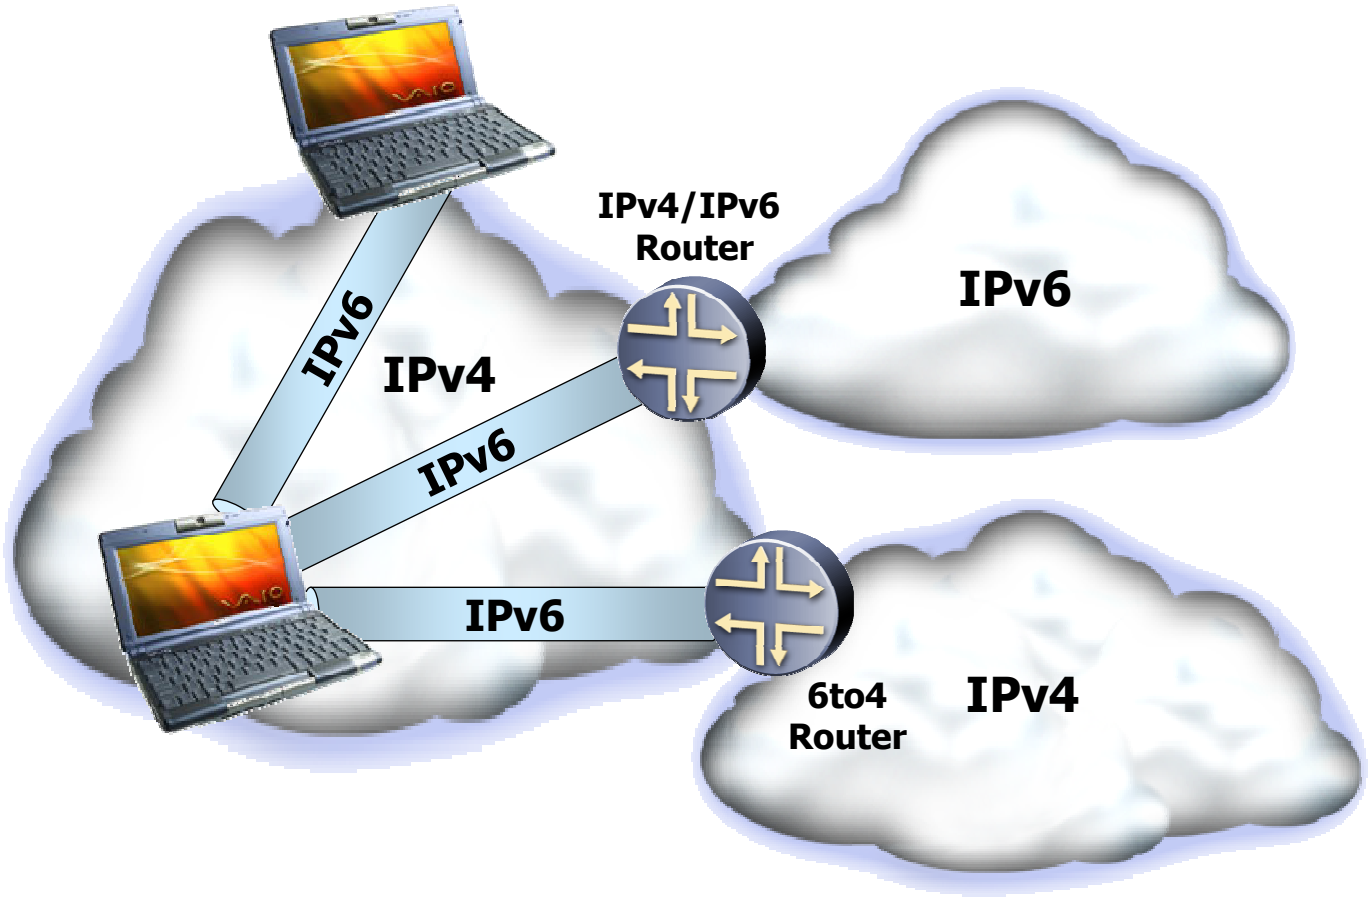

Configuration Example: Windows XP ISATAP Interface

```
C:\Documents and Settings\c.sellers>ipv6 if 2
Interface 2: Automatic Tunneling Pseudo-Interface
  does not use Neighbor Discovery
  does not use Router Discovery
  router link-layer address: 0.0.0.0
  EUI-64 embedded IPv4 address: 0.0.0.0
    preferred link-local fe80::5efe:169.254.113.126, life infinite
    preferred link-local fe80::5efe:65.114.168.91, life infinite
    preferred global ::65.114.168.91, life infinite
  link MTU 1280 (true link MTU 65515)
  current hop limit 128
  reachable time 24000ms (base 30000ms)
  retransmission interval 1000ms
  DAD transmits 0
```

- **Link-Local IPv6 Address**
- **ISATAP Identifier**
- **IPv4 Address**

Tunneling – Tunnel Brokering

Description – RFC 3053:

- An alternative approach involving dedicated servers which automatically configure tunnels on behalf of users
- IPv4-to-IPv6 transition tool developed through CSELT's participation in the IETF ngtrans work group.

Benefits:

- Allows management of IPv6 tunnel requests from users
- Allows ISPs to easily perform access control on the users enforcing their own policies on network resources utilization

Issues:

- Requires special attention to implementing security
- Mechanism may not work if host using private IPv4 addresses behind a NAT box

Deployment Applications:

- Particularly suitable for connections between small users such as traditional users of dial-up Internet connectivity and an IPv6 Service Provider

Tunneling – Tunnel Brokering

- Three basic components:
  - Client: Dual-stacked host or router, tunnel end-point
  - Tunnel Broker: Dedicated server for automatically managing tunnel requests from users, sends requests to Tunnel Server
  - Tunnel Server: Dual-stacked Internet-connected router, other tunnel end point
- A few tunnel brokers:
  - Freenet6 [Canada] (www.freenet6.net)
  - CERNET/Nokia [China] (www.tb.6test.edu.cn)
  - Internet Initiative Japan (www.iij.ad.jp)
  - Hurricane Electric [USA] (www.tunnelbroker.com) * Sponsor
  - BTexacT [UK] (www.tb.ipv6.btexact.com)
  - Many others…

Tunnel Broker

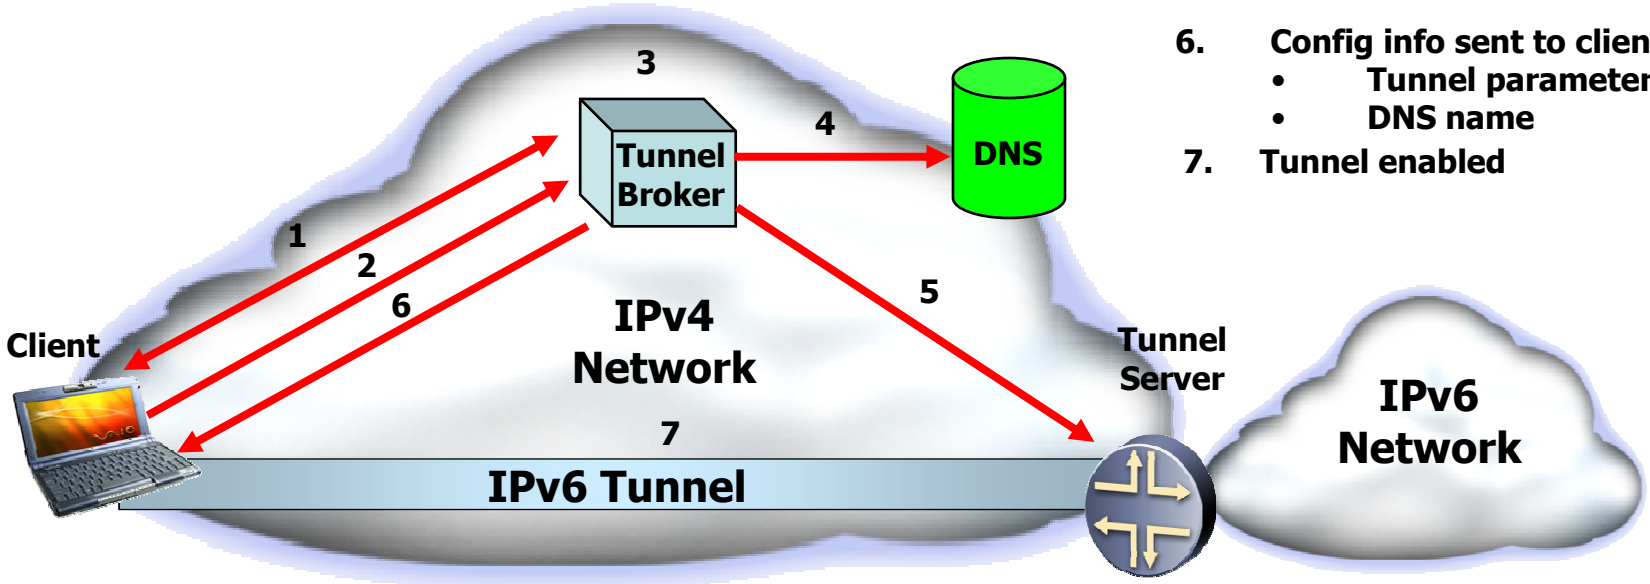

1. AAA Authorization
2. Configuration request
3. TB chooses:
   - TS
   - IPv6 addresses
   - Tunnel lifetime
4. TB registers tunnel IPv6 addresses
5. Config info sent to TS
6. Config info sent to client:
   - Tunnel parameters
   - DNS name
7. Tunnel enabled

Tunneling for RMv6TF

IPv6 Summit Demo Lab
2009-04-21

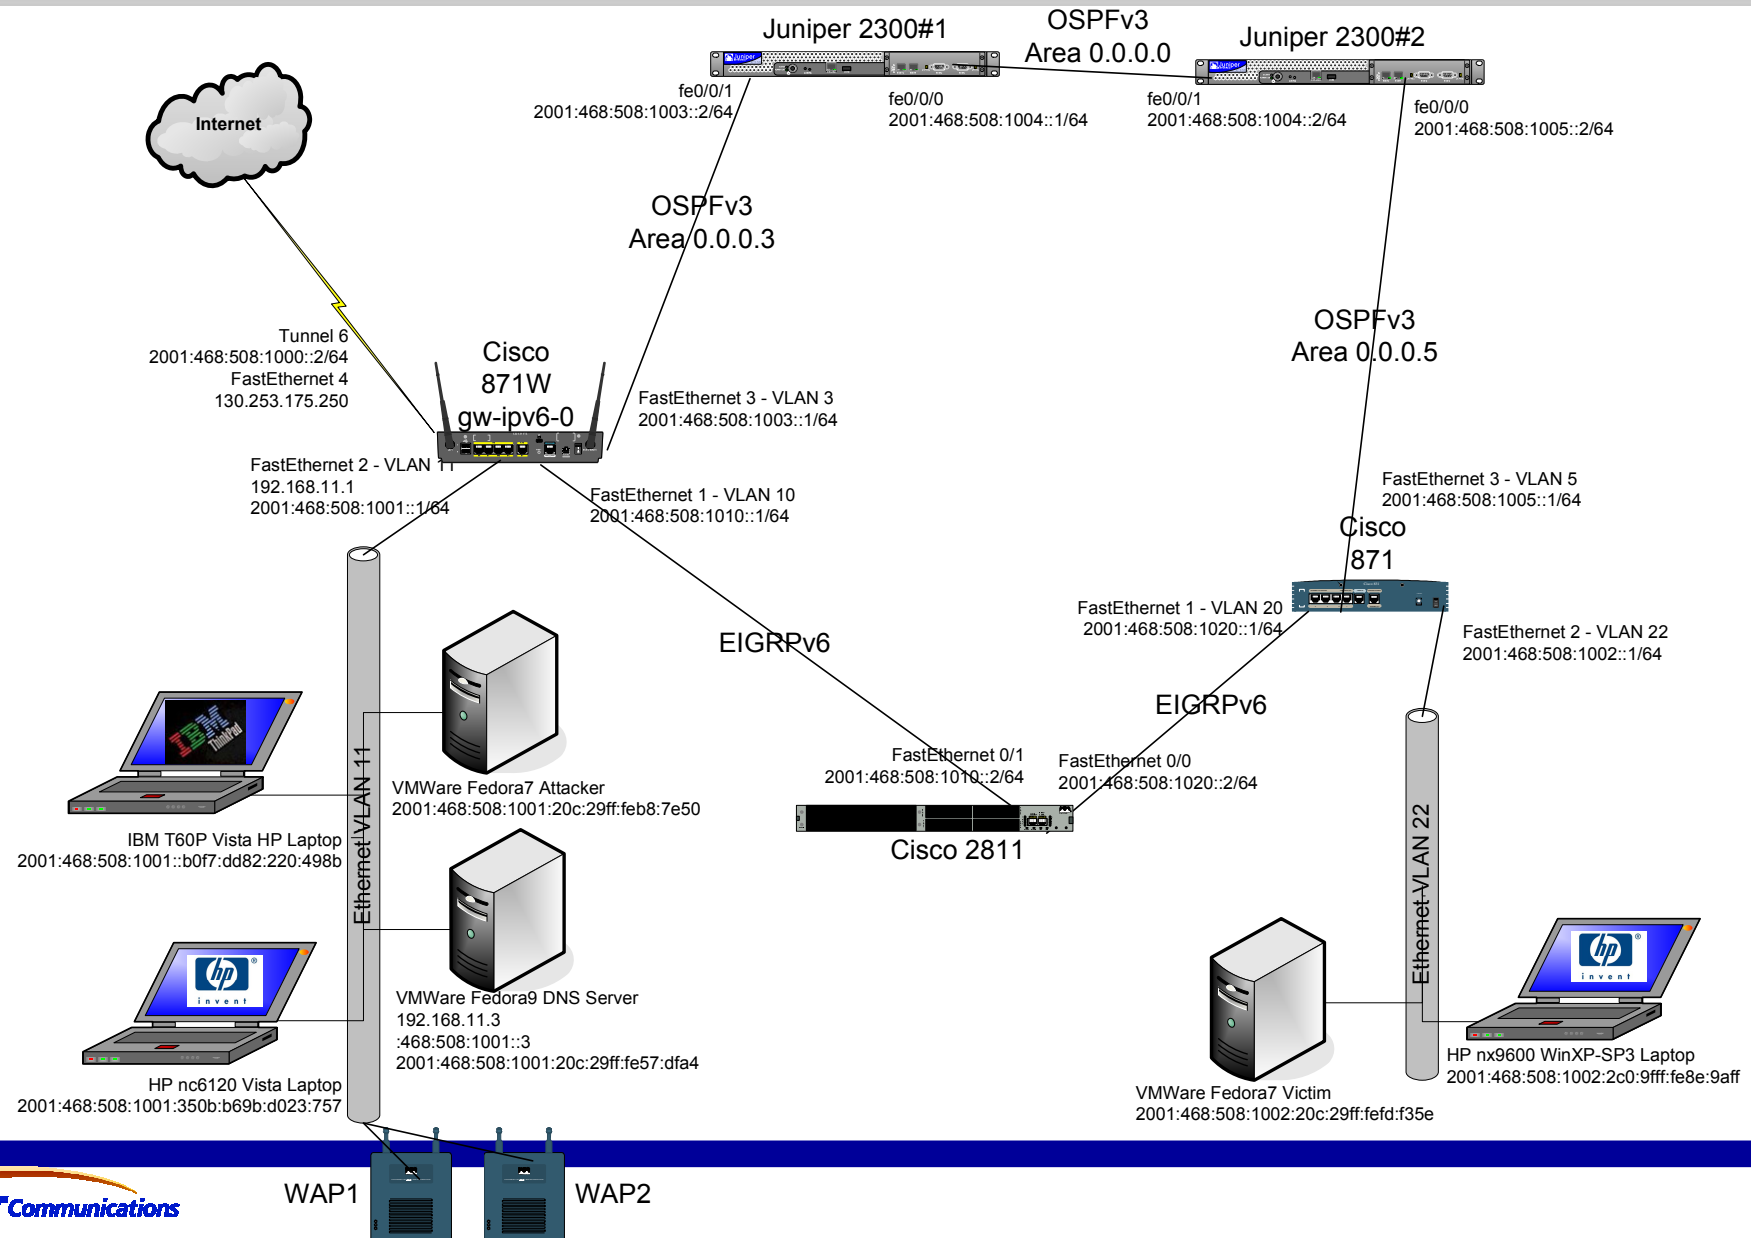

Network Address and Protocol Translation Techniques

Automatic Tunneling – Teredo (aka Shipworm)

Description – Encapsulation with automatic tunnels / dual stack:

- Known as IPv4 NAT traversal for IPv6
- Allows users in an IPv4 NAT environment with private addressing IPv6 connectivity
- Encapsulates IPv6 packets in an IPv4 UDP packet and tunnels them to a Teredo server on the IPv4 Internet

Benefits:

- Easy to implement on a “one-off” basis
- Works with private address space (NAT) configured nodes

Issues:

- Requires a significant amount of configuration
- Complicated to administer
- Some security risks

Deployment Applications:

- Microsoft is developing Teredo (RFC 4380)

Teredo

Teredo navalis

- For tunneling IPv6 through one or several NATs
  - Other tunneling solutions require global IPv4 address, and so do not work from behind NAT
  - Can be stateless or stateful (using TSP)
- Tunnels over UDP (port 3544) rather than IP protocol #41
- Basic components:
  - Teredo Client: Dual-stacked node
  - Teredo Server: Node with globally routable IPv4 Internet access, provides IPv6 connectivity to client
  - Teredo Relay: Dual-stacked router providing connectivity to client
  - Teredo Bubble: IPv6 packet with no payload (NH #59) for creating mapping in NAT
  - Teredo Service Prefix: Prefix originated by TS for creating client IPv6 address

More info at: http://en.wikipedia.org/wiki/Teredo_tunneling

Network Address and Protocol Translation

Description – RFC 2765, 2766:

➢ Provides routing between IPv4 and IPV6 network

- Network Address Translation (NAT)
  Translates IP address, IP, TCP, UDP, and ICMP header checksums

- Network Address Port Translation (NAPT)
  NAT + translation of TCP, UDP port numbers, ICMP message types

- Network Address Translation and Protocol Translation (NAT-PT)
  Translates IPv6 packet into equivalent IPv4 packet

- Network Address Port Translation and Protocol Translation (NAPT-PT)
  Allows IPv6 hosts to communicate with IPv4 hosts using single IPv4 address

Network Address and Protocol Translation (Cont'd)

Benefits:

- Allows IPv6 hosts to communicate directly with IPv4 hosts

Issues:

- Temporary solution – only use as last resort if no other translation options are viable
- Does not support IPv6 advanced features such as end-to-end security
- Limited design topology
- NAT device Single Point of Failure
- Flexible routing mechanisms cannot be used

Deployment Applications:

- Avoid if possible

Acquisition of IPv6 Service

How can I start with IPv6? Will it take a lot of capital to migrate to IPv6?

- Enablement of existing software features
- Acquiring upgrades through Maintenance & Support

Can I purchase IPv6 commercially?

- NTTA only commercially, globally available Tier 1IPv6 service, since 2002
- Non-commercial (R&D) service can be obtained through Moonv6, FreeNet6 and other test-beds
- Several other providers now able to provide IPv6, e.g. Verizon Wireless

IPv6 Network Operations

- Increased operational costs due to running dual stack
  - Dual stack is not the Point of Arrival
  - Dual-stacking will increase CPU and memory utilization by 15 to 25%
  - Performance issues with equipment that is optimized for IPv4 but not IPv6
  - Overhead caused by maintaining IPv4 and IPv6 routing tables, firewalls, DNS servers, etc.
- Operational teams need IPv6 troubleshooting skills
  - Tunnels are more difficult to troubleshoot than physical links
- Configuration management systems will help monitor the transition
- Regular operational checks needed to insure operations

Social Issues Regarding Transition

Denial - Perspectives

- *Most Network Managers will not ask for IPv6 until they run into a problem getting IPv4 space.* It is simple human nature to ignore a problem until it becomes a crisis.

- *Consumers will not ask for IPv6,* until the price they pay for a single IPv4 address exceeds the cost of a new home gateway, and the press tells them what to call it.

- *Application Developers will not make the necessary changes to deliver version-agnostic code until they see that the market is changing.* Network managers need to drive awareness as the application community will not directly feel the impact of IPv4 free-pool exhaustion, and consumers will not know what to ask for.

Courtesy: Tony Hain, Cisco Systems, ahain@cisco.com

Anger – IPv6 Only Events

- Distraught "IPv4 experts" are having difficulties:
  - IPv6-only exposes IPv4 dependencies in applications and middleware.
    - "Thunderbird and Firefox disable IPv6 dns by default"
- Failures when translating between versions exposes the invalid assumptions that some ISPs have been making.
  - "Linux NAT-PT (naptd) has stability issues and wedges"
- Provisioning model assumptions are exposed by new ways of handling addressing.
    - "It's a real pain in the ass to get DHCPv6 working."
    - "Why doesn't the RA include the DNS service?"
  - **Typing ":" instead of "." in a literal address exposes how resistant people are to change. (banging on keyboard and yelling "@#$***&!!!"**
    - **"Why do we have to type colon instead of dot like in a real address?"**

Courtesy: Tony Hain, Cisco Systems, ahain@cisco.com

Anger – Unbalanced Impacts

- The IPv4 address shortage will disproportionately harm the access providers relative to the content providers due to their imbalanced needs for additional addresses.
- If content providers require growth beyond the availability of IPv4, they can deploy IPv6, and then wait for the access providers to connect the content customers.

- Trading smaller and smaller blocks will cause the global IPv4 routing table to explode, and numerous small blocks make it difficult for large service providers to acquire enough space to sustain the business needs.
- Shortage driven IPv4 address block hijackings will become routine, which in turn will result in the routing table being politicized and access to content sites will be problematic.

Courtesy: Tony Hain, Cisco Systems, ahain@cisco.com

Negotiation -Alternatives

"The Class-E space (240/4) …"

Existing deployed end systems will not accept configuration into, or even talk to that space if a new endpoint tried to use it. Only useful for an entirely new walled-garden deployment.

"Carrier based nat(CGN) …"

natat the edge „works" because protocols like UPNP open holes to allow applications through. UPNP is a link-scope protocol with no security functions by design, so natin the core will fail for applications that rely on UPNP. the private IPv4 address space is not big enough for some CGN deployment models which include multiple addresses per device and coordination with external content partners.

Observed first hand at IETF 74 v6ops BOF in March 2009

Courtesy: Tony Hain, Cisco Systems, ahain@cisco.com

Depression -CGN breaks AJAX Apps

- Google Maps opens ~ 70 parallel connections

- IPv4/natmultiplexes multiple users through the port range, so 64k divided by 300 parallel connections results in ~200 customers per ISP based nataddress (assuming each customer is only allowed to run one simultaneous instance of iTunes or similar apps). Restricting the number of connections impacts utility of the app. Consensus wisdom before deployment is to plan on at most 8 customers per public IPv4 address

- Services generally don't allow connections from the same host to span multiple public side addresses, so when a port pool is exhausted, the subsequent connections on another address will cause the application to fail.

- Reuse of port pairs can't be guaranteed with a high rate of churn in the port pool, so the likelihood of matching src/dstport pairs to popular sites will expose the probability of TCP sequence number overlap between unrelated connections, and/or a port sitting in TCP Time-Wait at the server.

Google Maps limited to 10 parallel connections

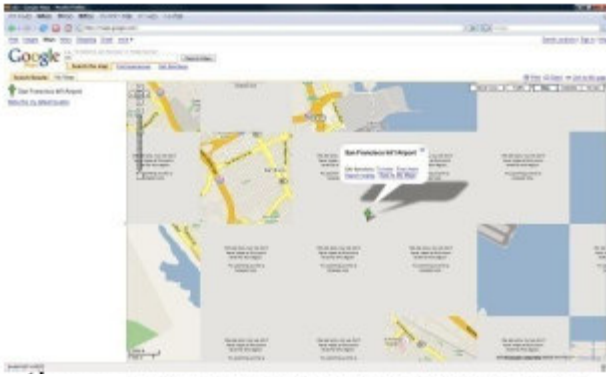

Source: Shin Miyakawa , Ph.D. NTT Communications Corporation

Courtesy: Tony Hain, Cisco Systems, ahain@cisco.com

Acceptance – IPv6 Enabled Web Sites

(growing list at sixy.ch)

ipv6.google.com

http://[2001:470:d:2ed::1]

http://[2001:4830:20e0:1::5]

http://[2001:4f8:fff6::21]

http://[2001:b48:12:1::2]

http://[2001:da8:200:200::4:28]

http://[2405:5000:1:2::99]

http://[2001:240:400::8591:e013]

http://[2001:b48:10::3]

http://[2001:252:0:1::2008:6]

http://[2001:49f0:1000::3]

http://[2001:218:2001:3005::8a]

http://[2001:2040:2000::6]

http://[2001:4830:2480:11::137]

http://[2607:f0d0:1000:11:1::2]

http://[2001:470:0:64::2]

http://[2a02:250::6]

http://[2001:1890:1112:1::20]

http://[2620:0:2d0:1::193]

http://[2a01:e0c:1:1599::1]

http://[2a01:298:3:1::abcd]

http://[2001:610:240:11::c100:1319]

http://[2001:dc0:2001:0:4608:20::]

http://[2001:9b0:1:104:230:48ff:fe56:31ae]

http://[2001:470:1:3a::13]

http://[2001:48a8:6880:95::21]

http://[2001:500:4:13::81]

http://[2a01:48:1:0:2e0:81ff:fe05:4658]

http://[2001:440:fff9:100:202:b3ff:fea4:a44e]

http://[2a01:a8:0:5::26]

http://[2001:630:200:4240:203:baff:fe87:14ed]

http://[2001:838:1:1:210:dcff:fe20:7c7c]

Courtesy: Tony Hain, Cisco Systems, ahain@cisco.com

Take-Away

- The good news is that there are IPv6 enabled content sites. The bad news is that they are mostly niche & fit on one slide.

- Content and application developers need to be aware that the carriers will be connecting the eyeballs via IPv6, and/or breaking IPv4 connectivity to all but basic web & email through CGN deployments.

- Business continuity requires the ability to operate and grow after the IPv4 free-pool is exhausted. Self-defensive technologists are resisting change, while CIOs are indecisive due to conflicting viewpoints. Recognizing that "fear of losing market value" is behind this impasse, will allow everyone to confront reality and take steps to move forward.

Transition Issues: Security

- Many transition technologies open security risks such as DoS attacks
- Examples:
  - Abuse of IPv4 compatible addresses
  - Abuse of 6to4 addresses
  - Abuse of IPv4 mapped addresses
  - Attacks by combining different address formats
  - Attacks that deplete NAT-PT address pools
- IPv6 security vulnerabilities not yet well-understood
  - By security personnel
  - By attackers

Transition Planning

- Assumption: Existing IPv4 network
- Easy Does It
  - Deploy IPv6 incrementally, carefully
- Have a master plan
- Think IPv4/IPv6 interoperability, not migration
- Evaluate hardware/software support
- Evaluate application porting
- ** KEEP TRANSITION SIMPLE **
- Limit scope and interaction of mechanisms
- Make sure normal humans can fully understand the interactions and implications of all mechanisms
- Monitor IETF v6ops WG

Transition Strategies: Dual-Stacked IPv4/IPv6 Backbone

- Migration Direction:
  - All at once
    - If all hardware and software is IPv6 capable
    - Migration is then controlled by DNS and address assignment
  - Edge-to-core
    - The edge is the killer app!
    - When services are important
    - When addresses are scarce
    - User (customer) driven
  - Core-to-edge
    - Reduced complexity (no tunneling)
- What hardware/software must be upgraded?
  - If “a lot,” cost/ complexity of migration is significantly increased
  - Watch for dependencies!
  - Perform careful regression testing

Dual Stack Where You Can, Tunnel Where You Must

Dual Stack Networks Run Both Protocols

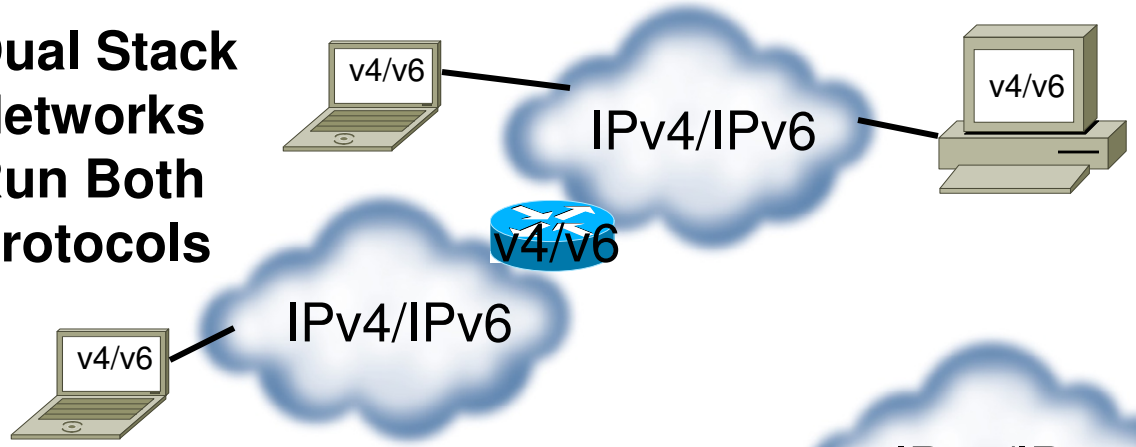

Tunneling IPv6 Across IPv4 Requires Careful Management

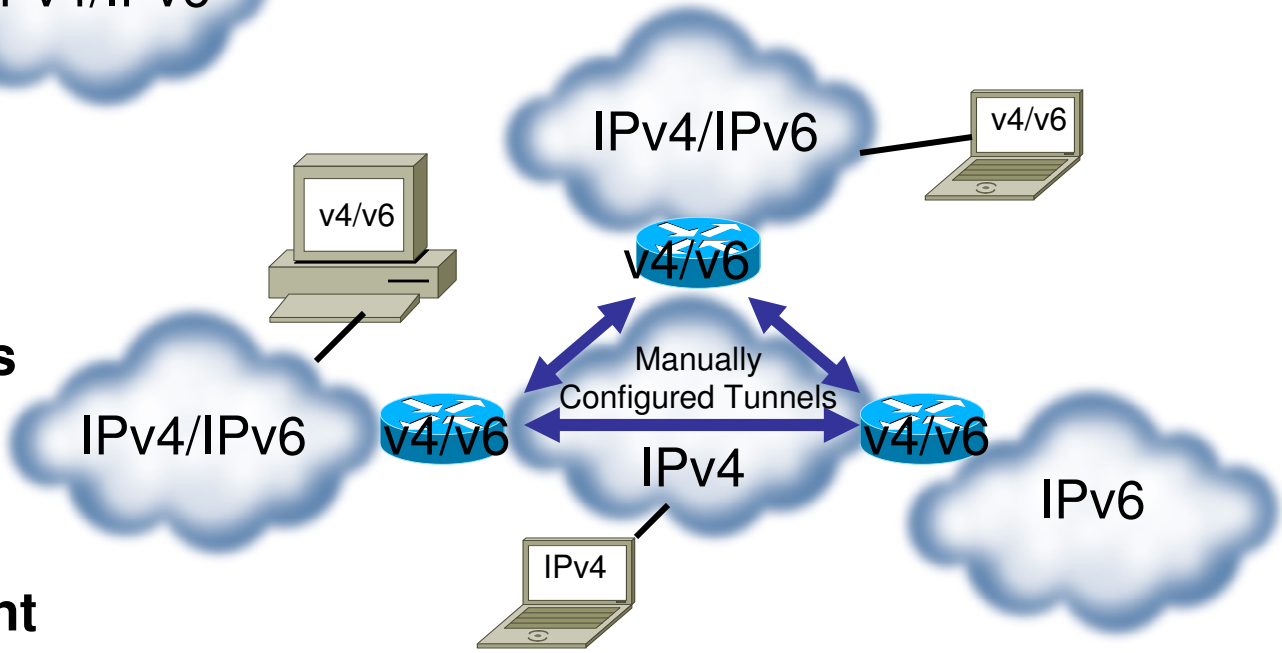

NTT's IPv6 Microsite: http://v6atwork.com

Countdown to Depletion: T -578 Days

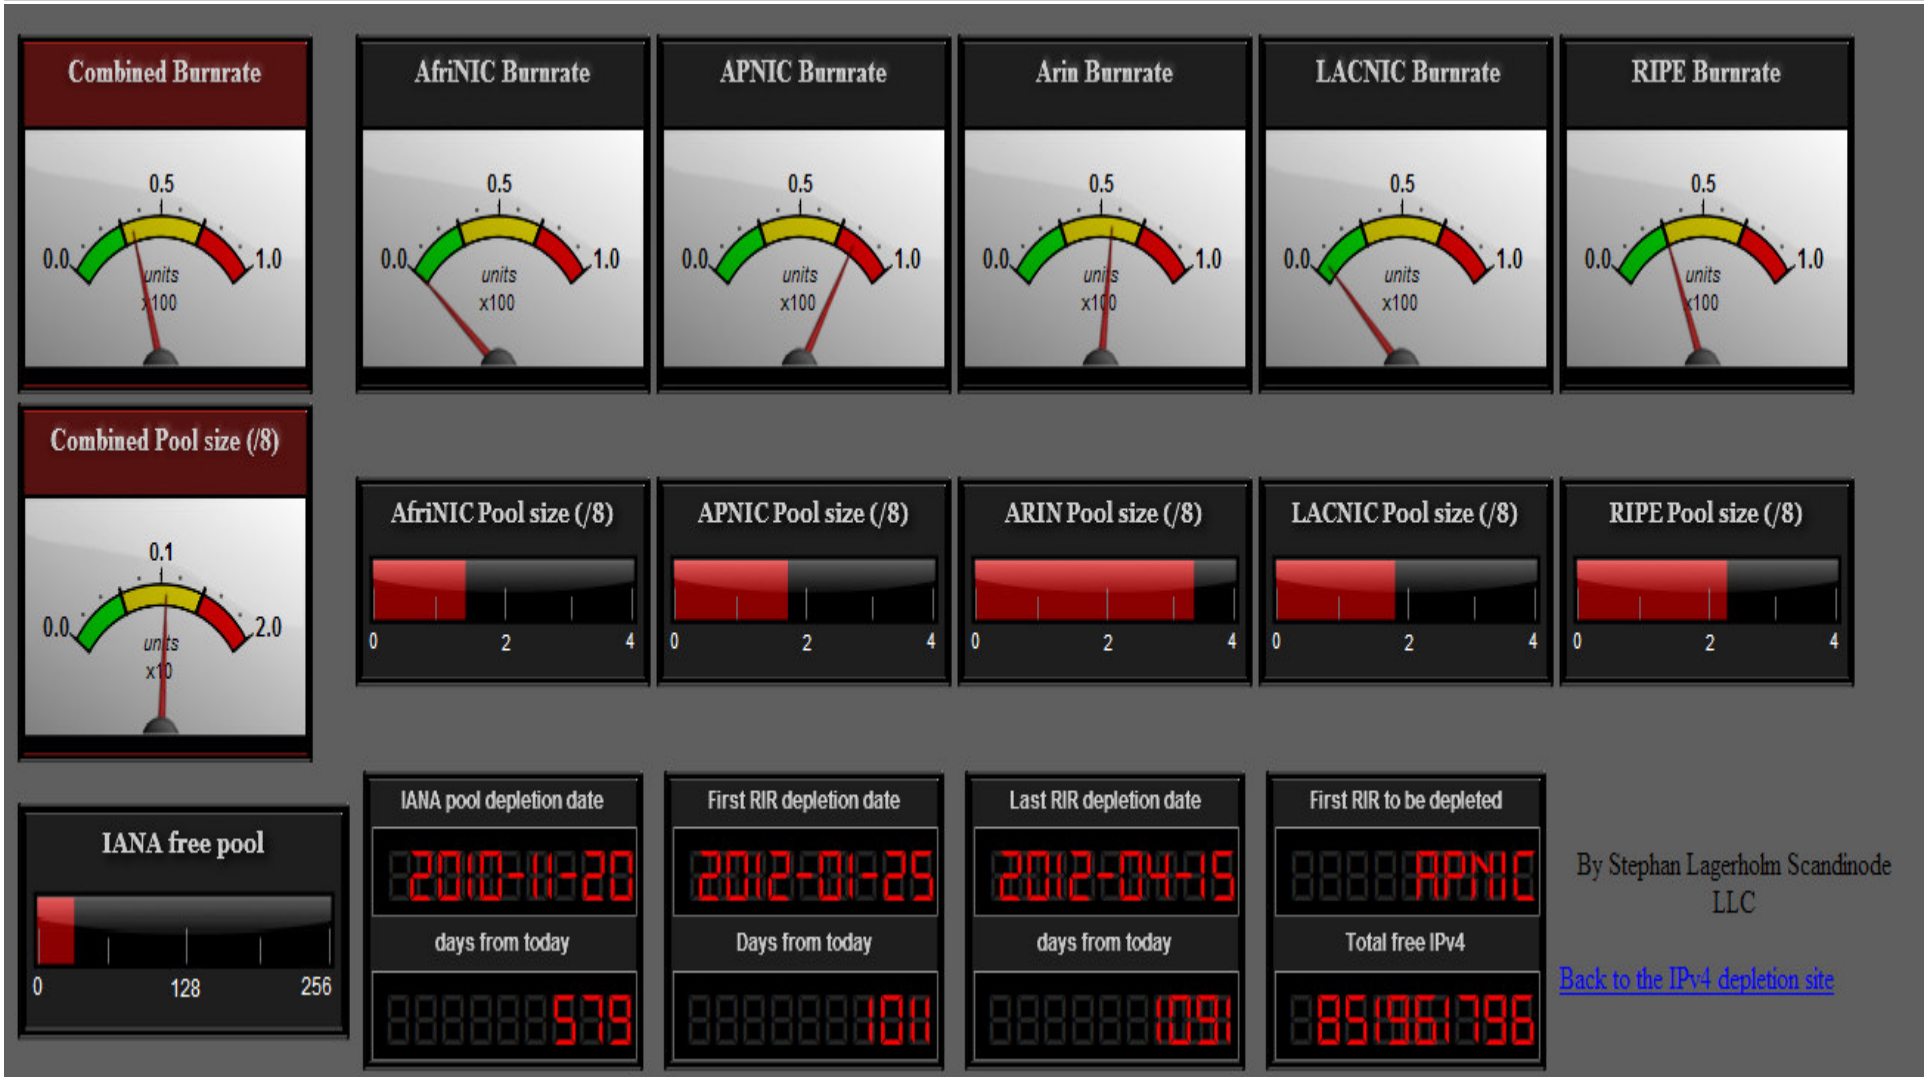

Questions?

Additional questions and comments can be addressed to:

csellers@us.ntt.net

303-919-6649 Mobile

Backup

Draft RFCs

- "IPv6 Tunnel Broker with the Tunnel Setup Protocol (TSP)", Marc Blanchet, Florent Parent, 6-May-08, <draft-blanchet-v6ops-tunnelbroker-tsp-04.txt>
- "Link Adaptation for IPv6-in-(foo)*-in-IPv4 Tunnels", Fred Templin, 14-May-07, <draft-templin-linkadapt-06.txt>

RFCs with IPv6 Tunnel Capability

- RFC 5214, Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)
- RFC 4891, Using IPsec to Secure IPv6-in-IPv4 Tunnels
- RFC 4380, Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)
- RFC 4213, Basic Transition Mechanisms for IPv6 Hosts and Routers
- RFC 3068, An Anycast Prefix for 6to4 Relay Routers
- RFC 3053, IPv6 Tunnel Broker
- RFC 2529, Transmission of IPv6 over IPv4 Domains without Explicit Tunnels
- RFC 2473, Generic Packet Tunneling in IPv6 Specification

- Microsoft ISATAP link
- http://www.microsoft.com/downloads/details.aspx?FamilyID=0f3a8868-e337-43d1-b271-b8c8702344cd&displaylang=en

IPv6 link-local over IPv4

Description – RFC2529

- IPv6 link-local addresses over IPv4 using IPv4 multicast
-IPv4 multicast becomes “virtual Ethernet link”
- IPv6 sees the entire network as a single LAN
- Interface Identifier is 32-bit IPv4 address padded with 32 leading ‘0’s
-123.234.20.1 becomes FE80:0:0:0:0:0:7bea:1401

Benefits

-Allows transport of IPv6 packets over IPv4 network
- Uses IPv6 auto-configuration

Issues

- Currently deprecated
- Not scalable
-IP multicast service not yet generally available

Deployment Applications

-Not recommended