


# Choosing A Solution for Your Trial

IPv6 Implementation and IPv4 Integration



gogoNET Live 2  
2 November, 2011

# Problem Statement

- Only IPv6 available for new public IP addresses
- The vast majority of Internet content and destinations is still IPv4
- **How do we connect IPv6 users with IPv4 content?**

# Types of Transition Technologies

## ➤ Dual Stack

- IPv6 + IPv4

## ➤ Tunnels

- IPv6 ↔ IPv4 ↔ IPv6
- IPv4 ↔ IPv6 ↔ IPv4

## ➤ Translators

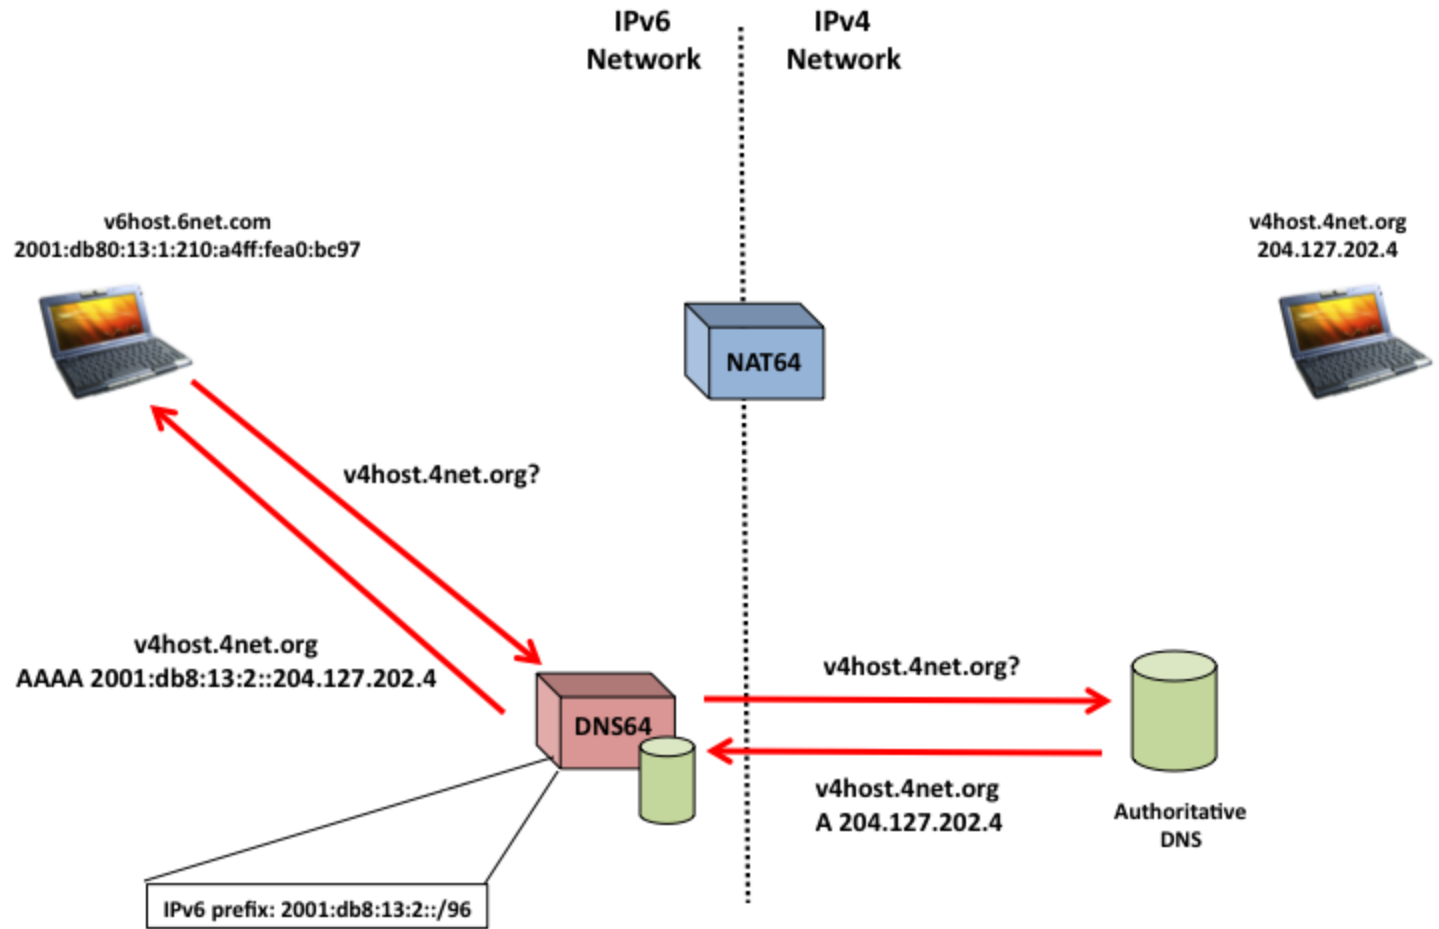
- IPv6 ↔ IPv4 (NAT64)
- Public IPv4 ↔ Private IPv4 (NAT44)



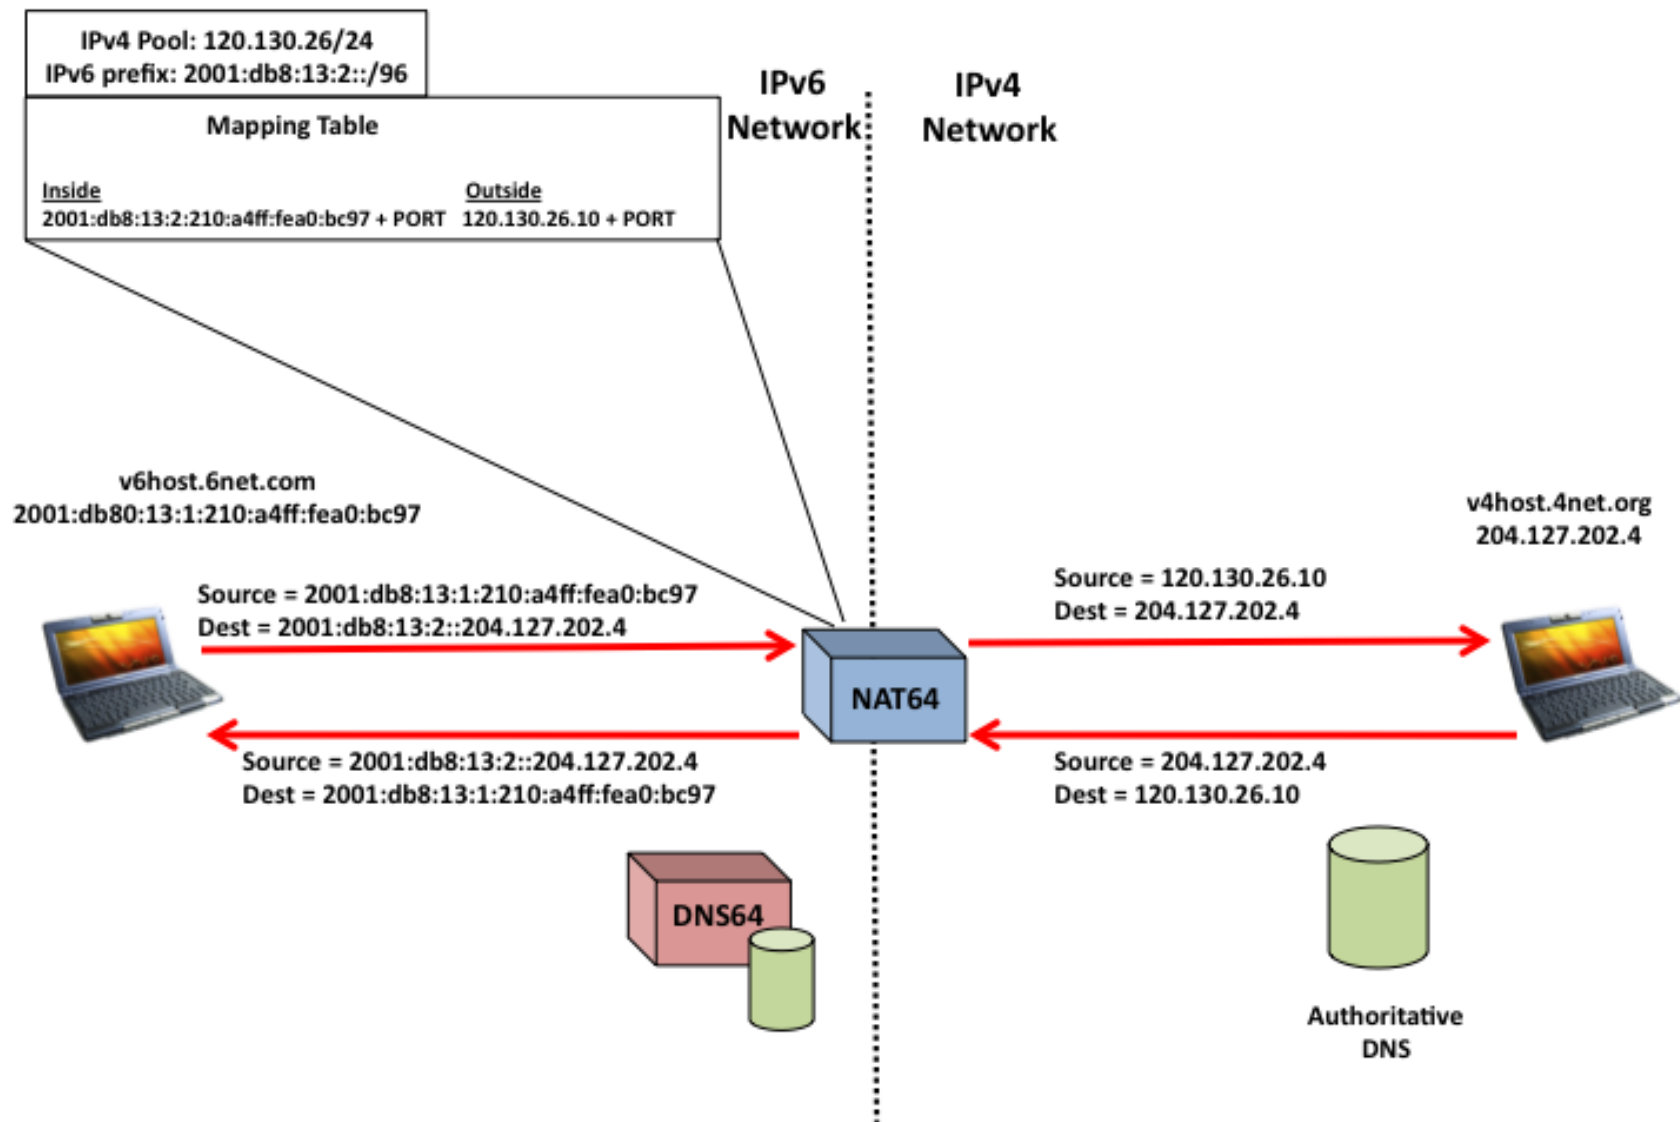
# Translators

## ➤ NAT64 + DNS64

- Replaces (deprecated) NAT-PT with DNS ALG



# NAT64 and DNS64



# Problems with NAT64


- Traffic cannot be asymmetric
  - Stateful mapping requires 2-way traffic through same translator
  - This makes the NAT64 a single point of failure
  - This makes the NAT64 an attractive attack target
- Some (**many**) applications will break
- No means to signal a session timeout
  - Can be particularly problematic in mobile networks
- Fragmented packets will not translate correctly
  - Port number only in first fragment
- No translation procedures for SCTP
  - Problematic for VoIP and mobile
- Possibility of dual AAAA Records in dual stack hosts, causing confusion
- DNS Records are not globally unique
  - Will break some applications
- Some IPsec modes do not work through translator



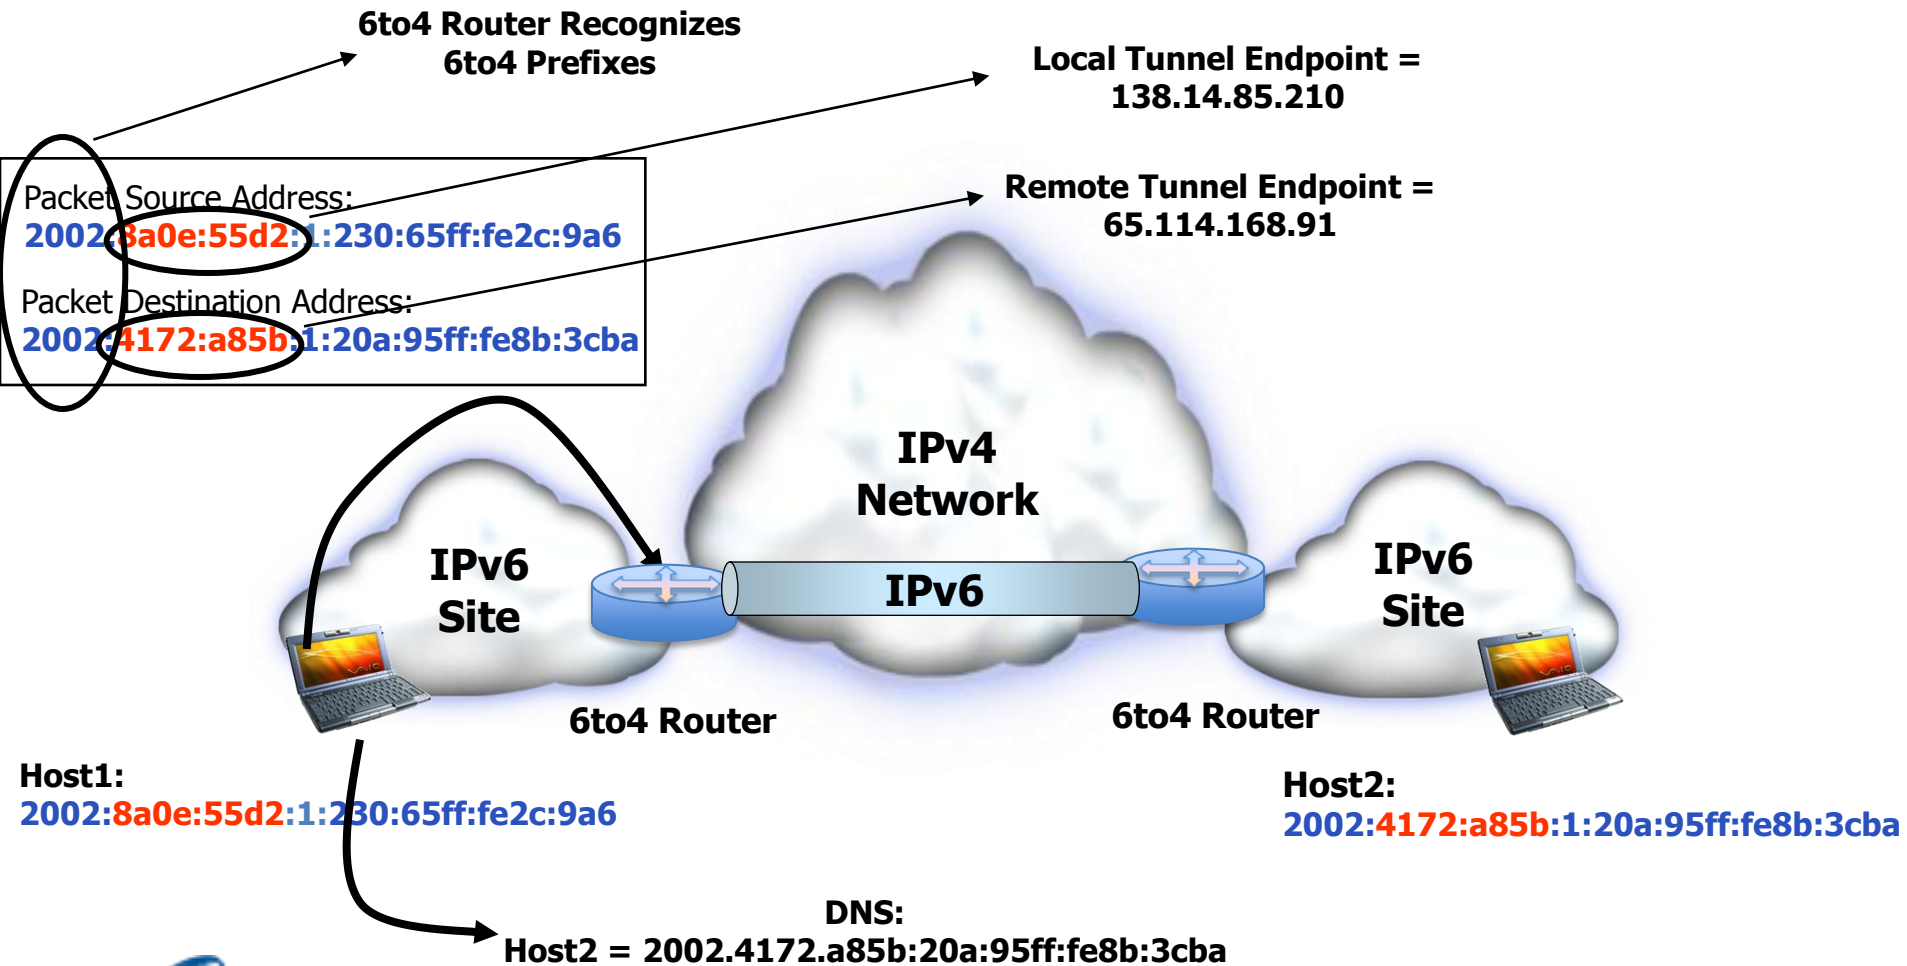
# Tunnels

- Used when some segments of the network cannot support IPv6
- Manual Tunnels
  - Tunnel endpoints manually configured
  - Useful for permanent site-to-site connectivity
  - IP-in-IP, GRE
  - MPLS is an excellent tunneling technology
    - 6PE: Native IPv6 over LSPs
    - 6VPE: IPv6 VPNs
- Automatic Tunnels
  - Useful for transient connectivity
    - Site-to-site
    - Device-to-device
  - Tunnel endpoints must be derived automatically
    - IPv4 endpoints embedded in IPv6 address
    - IPv4 endpoints assigned by an authoritative server

# 6to4: Embedded Tunnel Endpoints



# 6to4: Tunnel Setup



# 6rd

- **The problem with 6to4:**
  - Reserved 6to4 prefix: 2002::/16
  - Outside sources cannot locate a 6to4 destination without a 6to4 relay
- **6rd (IPv6 Rapid Deployment) is a superset of 6to4**
  - Allows 6rd prefix to be taken from local address block
  - Outside sources route to address block normally
  - 6rd border relay recognizes locally reserved block and tunnels accordingly
  - Eliminates translation difficulties of 6to4 relay

# 6rd Components

IPv4 Interface (Tunnel Destination)  
192.88.99.100 (c058:6364)

IPv4 Interface (Tunnel Source)  
199.15.4.2 (c70f:0402)

2001:db8:25:ac::ef12:5678



IPv6  
Internet

6RD  
Border  
Relay

IPv4  
Service Provider

6RD  
CPE

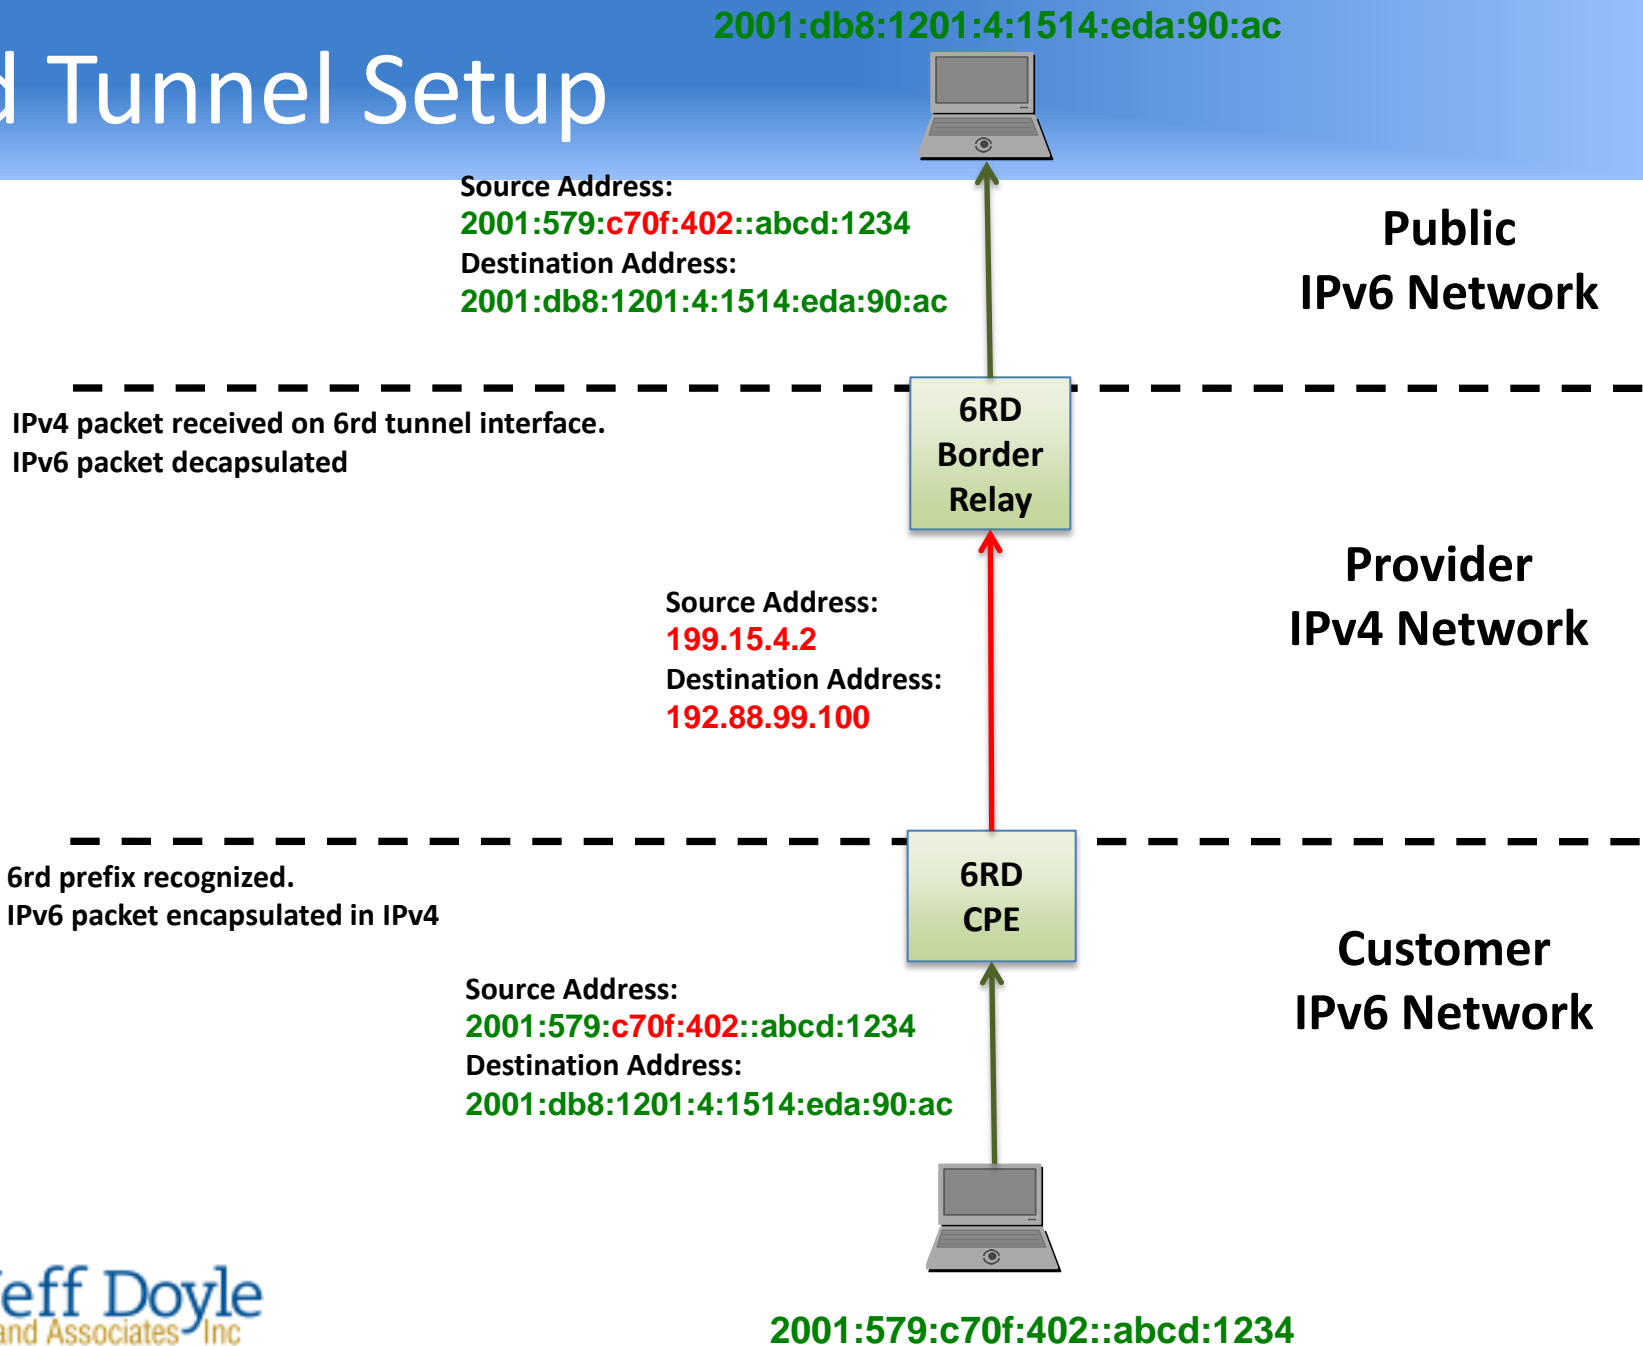
IPv6  
Customer



2001:579:1:2::abcd:1234



# 6rd Tunnel Setup



# Other Automatic Tunnels

## ➤ ISATAP


- Used for device-to-device connectivity

## ➤ Tunnel Brokers

- Many implementations
- Server-assigned tunnel endpoints
- Device or router tunneling to public IPv6 Internet

## ➤ Teredo

- Device to public IPv6 Inte



ISATAP Tunnels

# Dual Stack


- Simplest solution
  - Every interface “speaks” both IPv4 and IPv6
- Transition is driven by DNS
  - If destination address in A Record, speak IPv4
  - If destination address in AAAA Record, speak IPv6
  - If both A and AAAA Records are returned, prefer IPv6
- **Problem:** How do we dual stack if we do not have enough IPv4 addresses?
- **Solution:** Large-Scale NAT (LSN)

# Large-Scale NAT: Concepts

- aka Carrier Grade NAT (CGN)
- Normal NAT44, but located in provider network
  - What defines “carrier grade” is questionable
- Moves globally unique IPv4 address pool from customer edge to more centralized location
  - Allows a single IPv4 address to serve many more users
  - Address + Port translation (NAPT)
  - 65,535 TCP ports + 65,535 UDP ports
- 3 LSN architectures:
  - NAT444
  - NAT464
  - DS-Lite

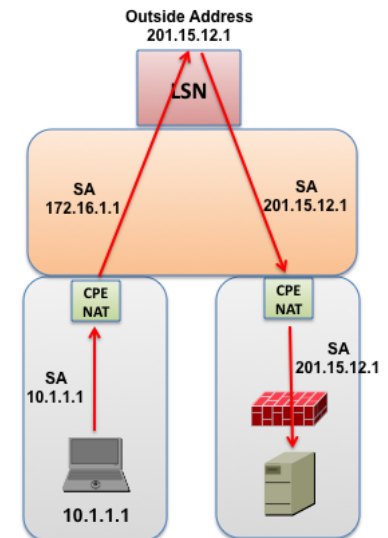
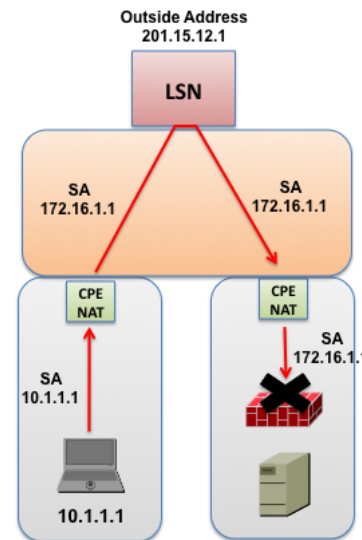
# NAT444

- aka 2-layer or dual NAT
- 3 address layers, all IPv4
- **Advantage:** Simple (well-understood) NAT44 used on both sides
- **Advantage:** CPE NAT does not need to be changed
- **Advantage:** Re-use of same private address blocks behind each LSN provides scaling




# Problems with NAT444

- Many applications break through dual-layer NAT
- Potential address conflicts between customer's private addresses and provider's private link addresses
- Potential filter problems for traffic between customers behind same LSN
  - Firewall and ACL policies often block private source addresses
- Potential Solution: “Hairpinning” through LSN
  - But does not solve address conflict problem
- Potential Solution: “ISP Shared Addresses”:
  - Reserved address block from remaining IPv4 space
  - But, no designated shared space yet
- Potential Solution: NAT464



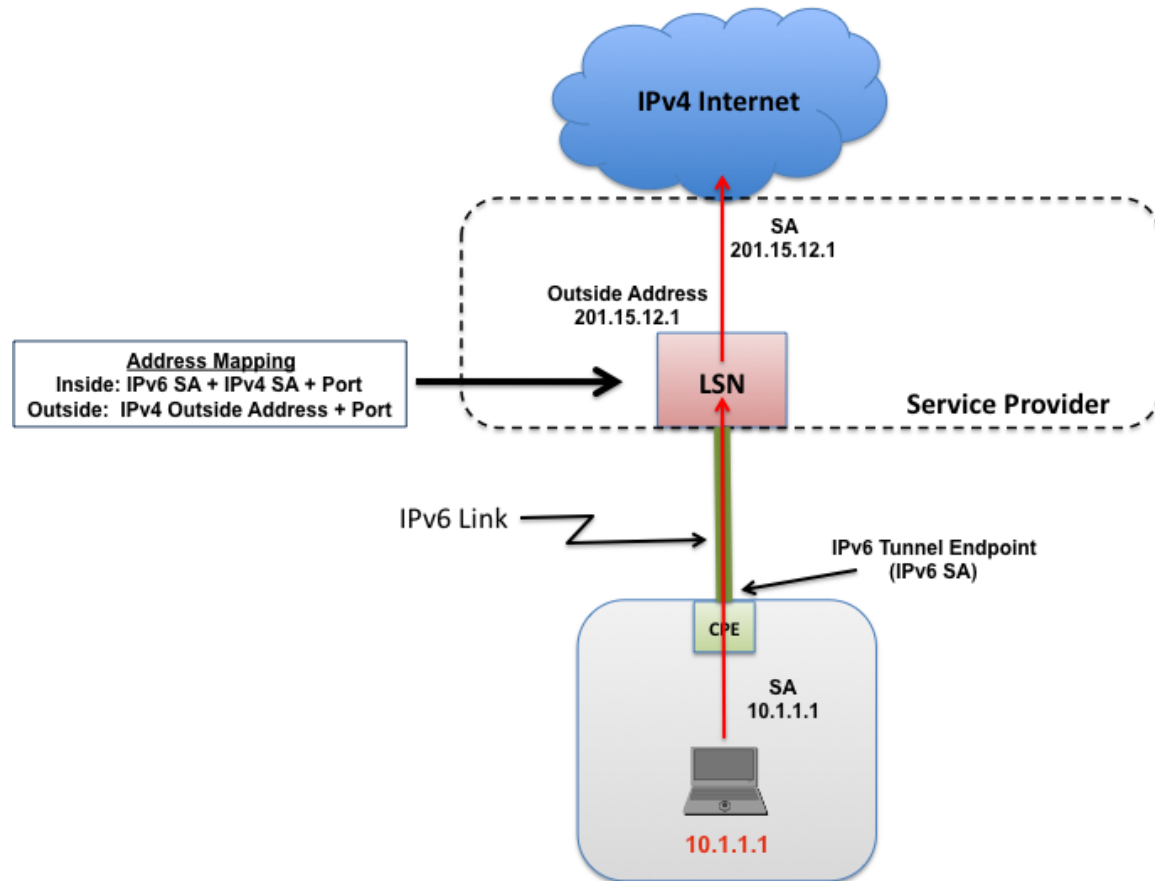
# NAT464



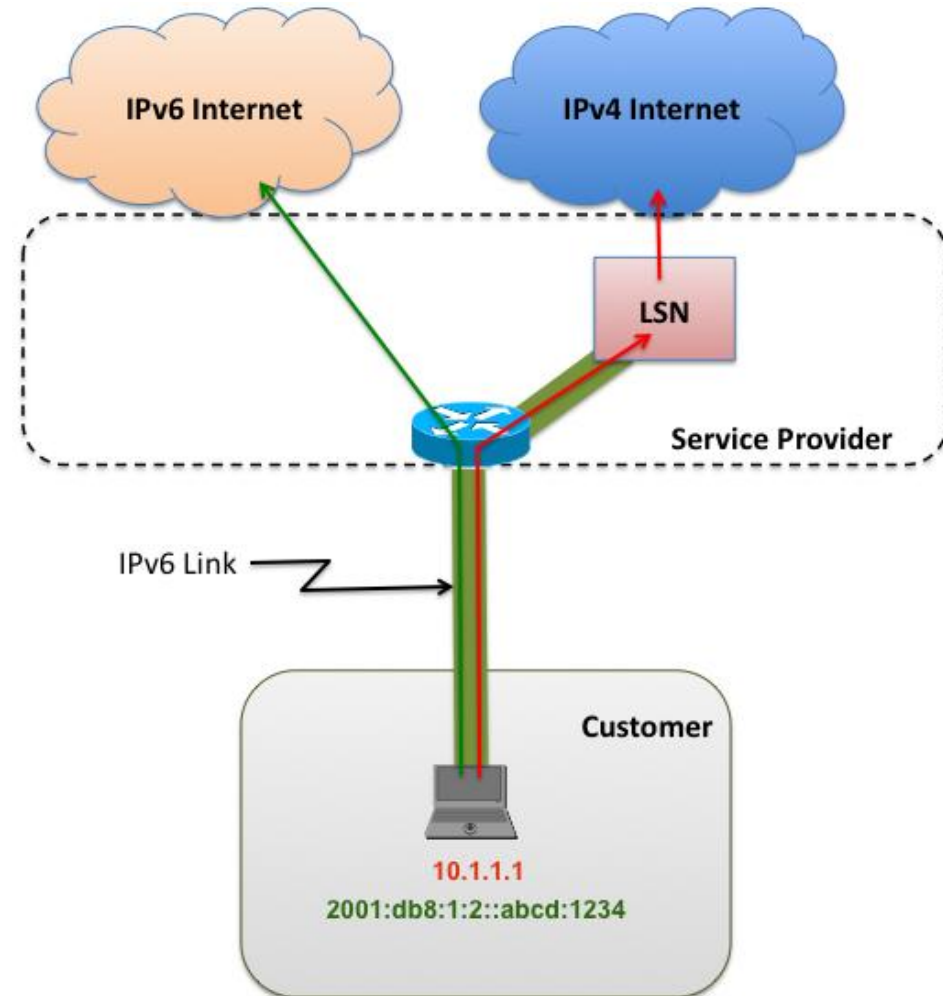

- **Advantage:** Solves address conflict and intra-LSN filtering problems
- **Advantage:** Uses IPv6-only links between provider and customer
- **Disadvantage:** All the problems of NAT64, doubled

# Dual-Stack Lite

- IPv4-in-IPv6 tunneling to LSN
- **Advantage:** Solves address conflict and intra-LSN filtering problems
- **Advantage:** Uses IPv6-only links between provider and customer
- **Advantage:** Only one NAT layer
- **Disadvantage:** Requires specialized, DS-Lite capable CPE NAT



# DS-Lite Scenarios



# LSN Issues: Can LSN Save IPv4?

- **Scaling**
  - How many users can a single public IPv4 address support?
  - What are the performance (mapping) characteristics?
  - ~10,000 users per CMTS or DSLAM: Can one LSN support this?
- **Security**
  - Moving address pool from customer edge to provider network creates a depletion target
- **Identifying users by IP address becomes impractical**
  - Potential security problems
  - Interesting social abuse issues
- **White- and Black-Listing is problematic**
  - Need to study traceback solutions
- **Architecture**
  - LSN in forwarding path becomes single point of failure
  - LSN away from forwarding path may require source routing or other undesirable complexities
- **UPnP does not work behind LSN**
  - Solutions currently being studied
- **Lawful Intercept will introduce intense logging activity**

# LSN Conclusions

- LSN must not be viewed as a permanent solution
  - Issues and complexities make it undesirable
  - No accurate projections yet when IPv4 can be de-commissioned
- Almost no production experience with NAT444 or DS-Lite
  - Therefore no reliable scaling or performance statistics
  - Expect the unexpected
- Lab testing is essential
- DS-Lite appears to be the least-undesirable solution
  - Most broadband providers appear to be moving in that direction
  - But there may be applications for NAT444

# Questions?

[jdoyle@doyleassociates.net](mailto:jdoyle@doyleassociates.net)

[www.doyleassociates.net](http://www.doyleassociates.net)

+1-303-428-4680

