# IPv6 Tutorial

Marc Blanchet, CTO
Marc.Blanchet@hexago.com

# Plan

- IPv6 on FreeBSD and Windows XP
- IPv6 Transition Main Mechanisms
- Case Studies
- IPv6 Applications configuration

HEXAGO

# IPv6 on FreeBSD and Windows XP

# Enabling IPv6 on FreeBSD

- /etc/rc.conf
  - ipv6_enable= »YES »
    - generates router solicitation
- ifconfig -a

# Enabling IPv6 on WindowsXP

- ipv6 install  (or netsh interface ipv6 install)
  - Enables the stack
  - To be done once
- By default, in autoconfiguration mode
  - Generates router solicitation
- ipconfig
- netsh interface show ipv6 address

# IPv6 Transition
# Main Mechanisms

# Plan

- Strategy
- Dual Stack
- Tunnelling IPv6 over IPv4
  - Configured tunnels
  - Automatic tunnelling
    - 6to4
    - Tunnel Broker with TSP
    - ISATAP
    - Teredo
- Application-layer Gateways
  - Proxy

# Other Mechanisms (not covered)

- Tunnelling IPv6 over IPv4
  - Automatic tunnelling
    - IPv4-compatible addresses
    - 6over4
- Tunnelling IPv4 over IPv6
  - DSTM (with DHCP, RPC, TSP)
- Translation at IP or transport layer
  - SIIT
  - NAT-PT
  - Bump-in-the-stack
  - Bump-in-the-api
  - Transport-layer Translator
- Application-layer Gateways
  - Socks

# Strategies

- For end-systems:
  - Dual stack approach
- For network integration:
  - Tunnels
  - IPv6-only to IPv4-only: some kind of translation
  - Proxy

# Dual Stack Host

- Node has both IPv4 and IPv6 stacks and addresses
- IPv6-aware application asks for both IPv4 and IPv6 addresses of destination
- DNS resolver returns IPv6, IPv4 or both addresses to application
- IPv6/IPv4 applications choose the address and then can communicate
  - With IPv4 nodes using IPv4
  - Or with IPv6 nodes using IPv6

# Dual Stack Host (cont.)

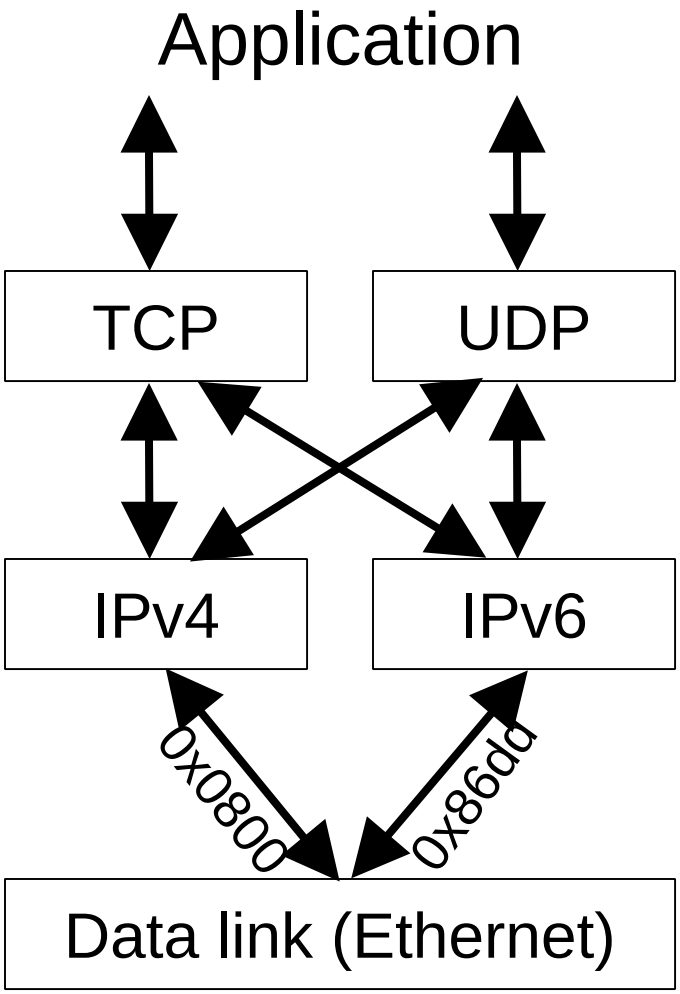

# Tunnelling IPv6 in IPv4

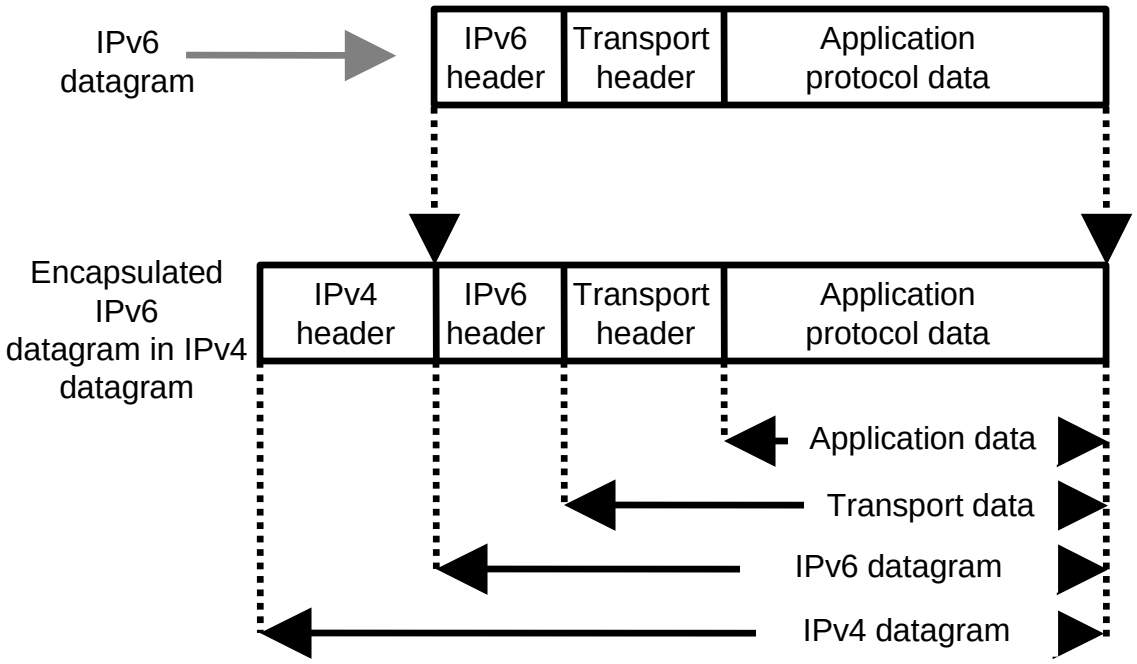

# Tunnelling IPv6 in IPv4

- IPv6 encapsulated in IPv4
  - IP protocol 41
- Many topologies possible
  - Router to router
  - Host to router
  - Host to host
- The tunnel endpoints take care of the encapsulation. This process is “transparent” for the intermediate nodes
- Tunnelling is used by most transition mechanisms
- If security gateways are present in the path, then they need to let through IP packets transporting protocol 41
- Problem if NAPT is in the path
- Can also be accomplished by GRE tunnels

# Tunnelling IPv6 over IPv4

- Configured tunnels
- 6to4
- Tunnel broker with TSP
- ISATAP
- Teredo

# Configured Tunnels

- Tunnel endpoints are explicitly configured
- Tunnel endpoints must be dual stack nodes
  - The IPv4 address is the endpoint for the tunnel
  - Require a reachable IPv4 address (no NAPT)
- Tunnel config implies:
  - Manual configuration of:
    - Source and destination IPv4 address
    - Source and destination IPv6 address
- Between:
  - Two hosts
  - One host and one router
  - Two routers (for two networks)

# Configured Tunnels

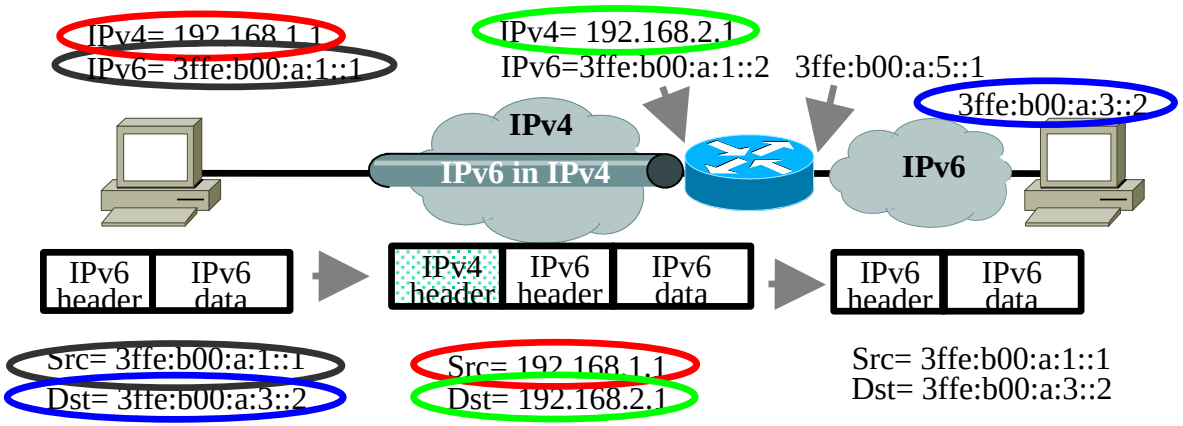

# Configured Tunnels

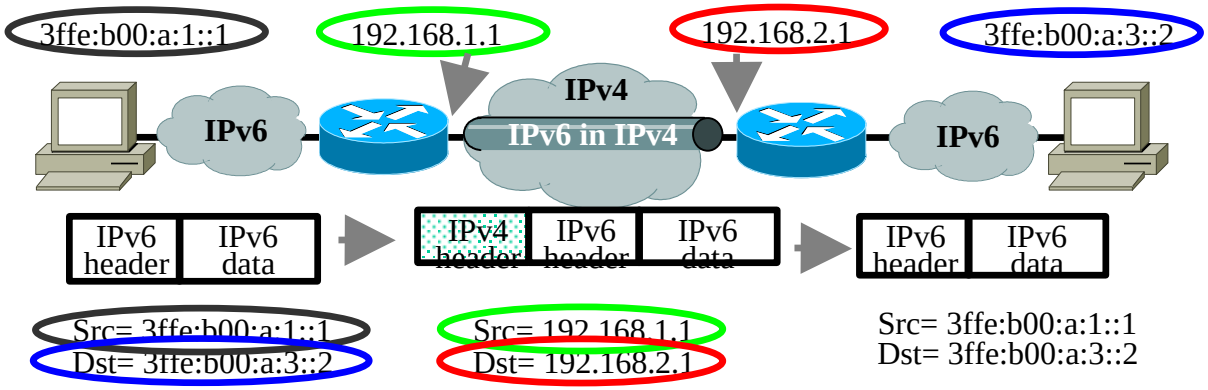

# Configured Tunnels Considerations

- Tunnels cannot go through a NAPT
- If site uses a NAPT, then one scenario might be to end the tunnel at the NAPT box

# 6to4

- Applicability: interconnection of isolated IPv6 domains over an IPv4 network
- Automatic establishment of the tunnel
  - No explicit tunnels
  - By embedding the IPv4 destination address in the IPv6 address
  - Under the 2002::/16 reserved prefix. (2002::/16 = 6to4)
- Gives a full /48 to a site based on its external IPv4 address
  - IPv4 external address embedded: 2002:<ipv4 ext address>::/48
  - Format: 2002:<ipv4add>:<subnet>::/64

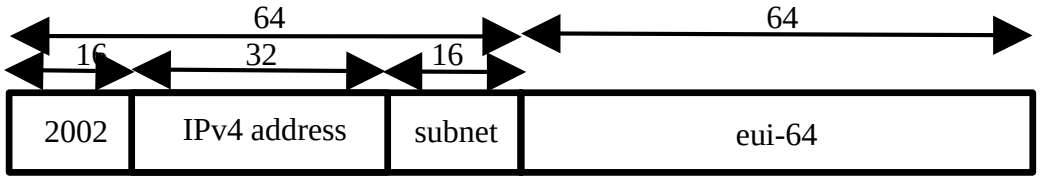

# 6to4 Network to Network

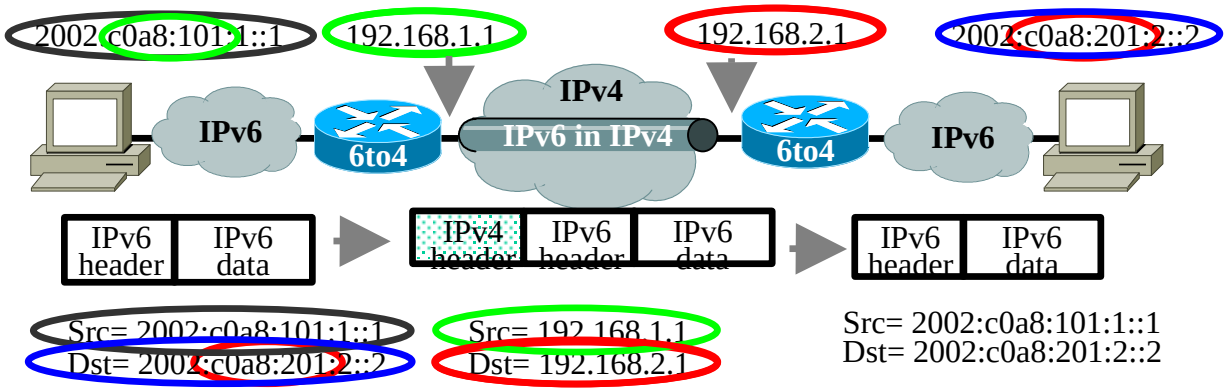

# Who Needs to Support 6to4?

- Egress router:
  - Implements 6to4
  - Must have a reachable external IPv4 address
    - Often configured using a loopback interface address
  - Is a dual-stack node
- Individual nodes:
  - Nothing needed for 6to4 support. 2002 is an "ordinary" prefix that may be received from router advertisements
  - Doesn't need to be dual-stack
- 6to4 relay

# Issues with 6to4

- Bound to the IPv4 external address:
  - If egress router changes its IPv4 address, then it means that you need to renumber the full IPv6 internal network
  - Only one entry point (no easy way to have multiple network entry points for redundancy)
- If everybody is using 6to4, it effectively puts IPv4 Internet host routes into IPv6 routes
- Needs a 6to4 relay which is an open-relay

# ISATAP

- Intra-Site Automatic Tunnel Addressing Protocol
- Automatic tunnelling from ISATAP nodes to the ISATAP routers in a private network
- Creates a virtual IPv6 link over the IPv4 network
- Special bits in the Node identifier part of the IPv6 address identify an ISATAP address.

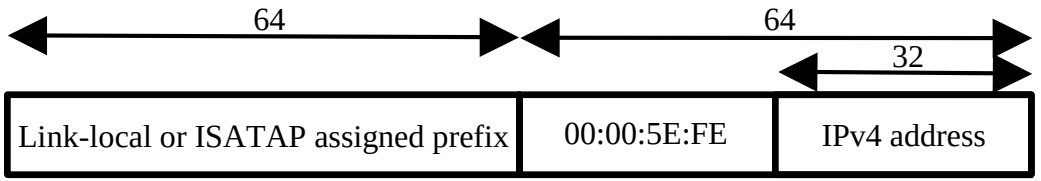

# ISATAP Example

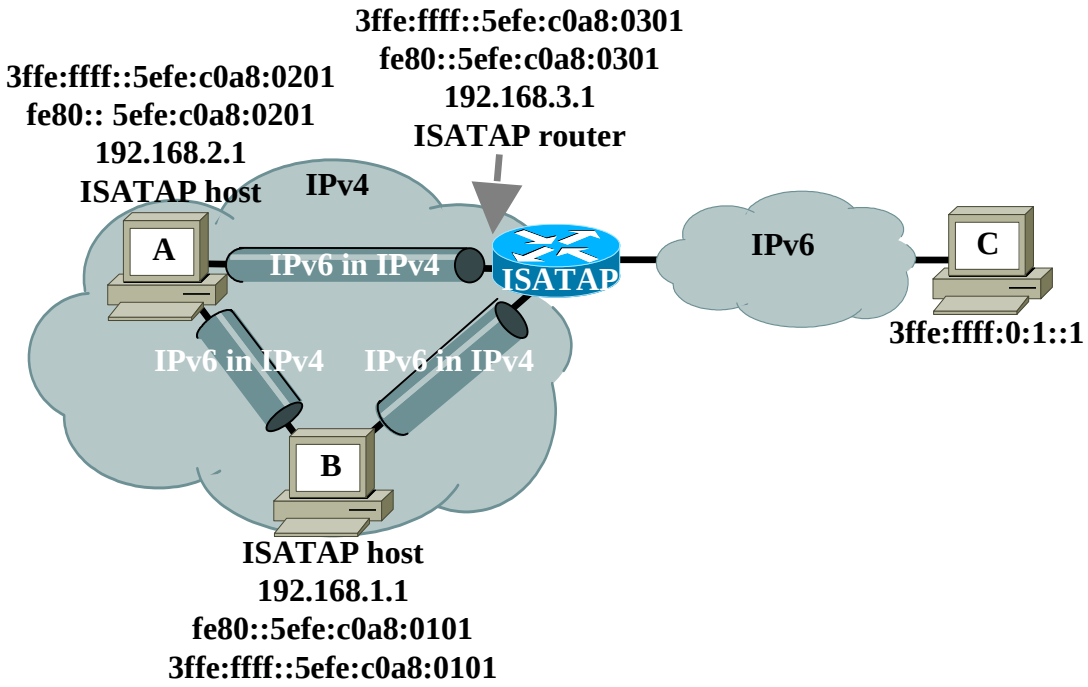

# ISATAP Applicability

- Inside a enterprise network
- No NAT in the path
- Not between providers
- Not global

# Teredo

- NAPT prohibit the use of direct tunnels.
- Uses IPv6 in UDP in IPv4
- External mapping of IPv4 address and port are discovered by the Teredo server (on the external side of NAT)

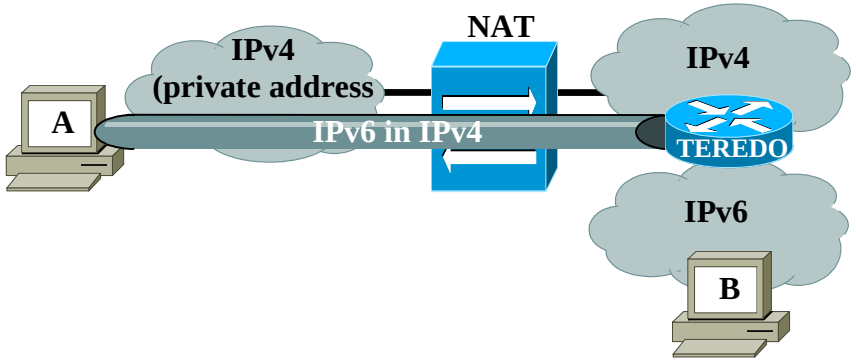


- Teredo uses a specific prefix. The address includes the IPv4 and port number of the host.

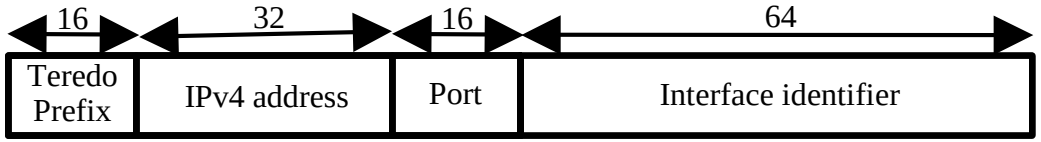

# IPv6 in UDP in IPv4

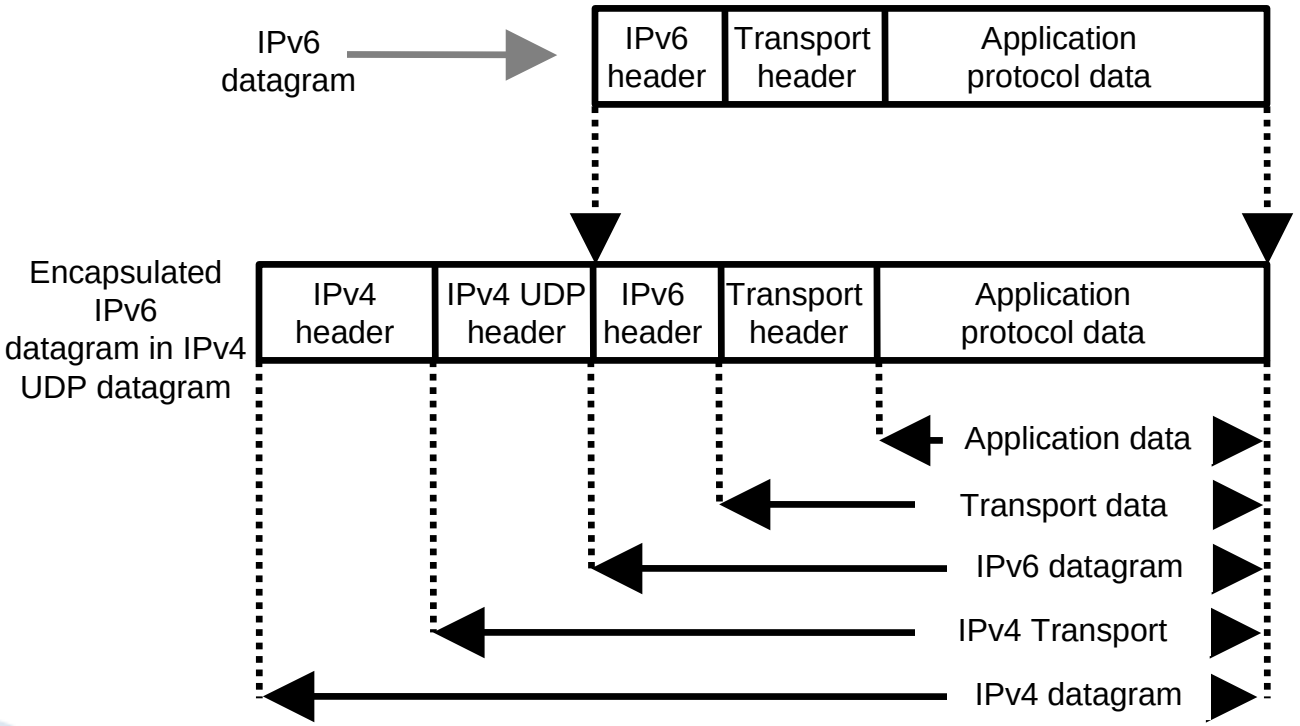

# Teredo Applicability

- Hosts behind an IPv4 NAT that wants traverse the IPv4 NAT.
- Tax for UDP encapsulation. Effective MTU is smaller.
- Automatic tunnels: Teredo servers and relays are subject to attacks.

# Tunnel Broker

- Automation of configured tunnels
- Uses a Control protocol for the tunnel setup
  - Tunnel Setup Protocol (TSP)
  - Client sendS a request for tunnel
  - Broker,
    - based on policies,
    - sends the appropriate tunnel info
    - and configures its tunnel end
  - Client then configures its tunnel end
  - Client receives:
    - a stable IPv6 address
    - a stable IPv6 prefix
- Well known service: http://www.freenet6.net

# Tunnel Broker Scenarios

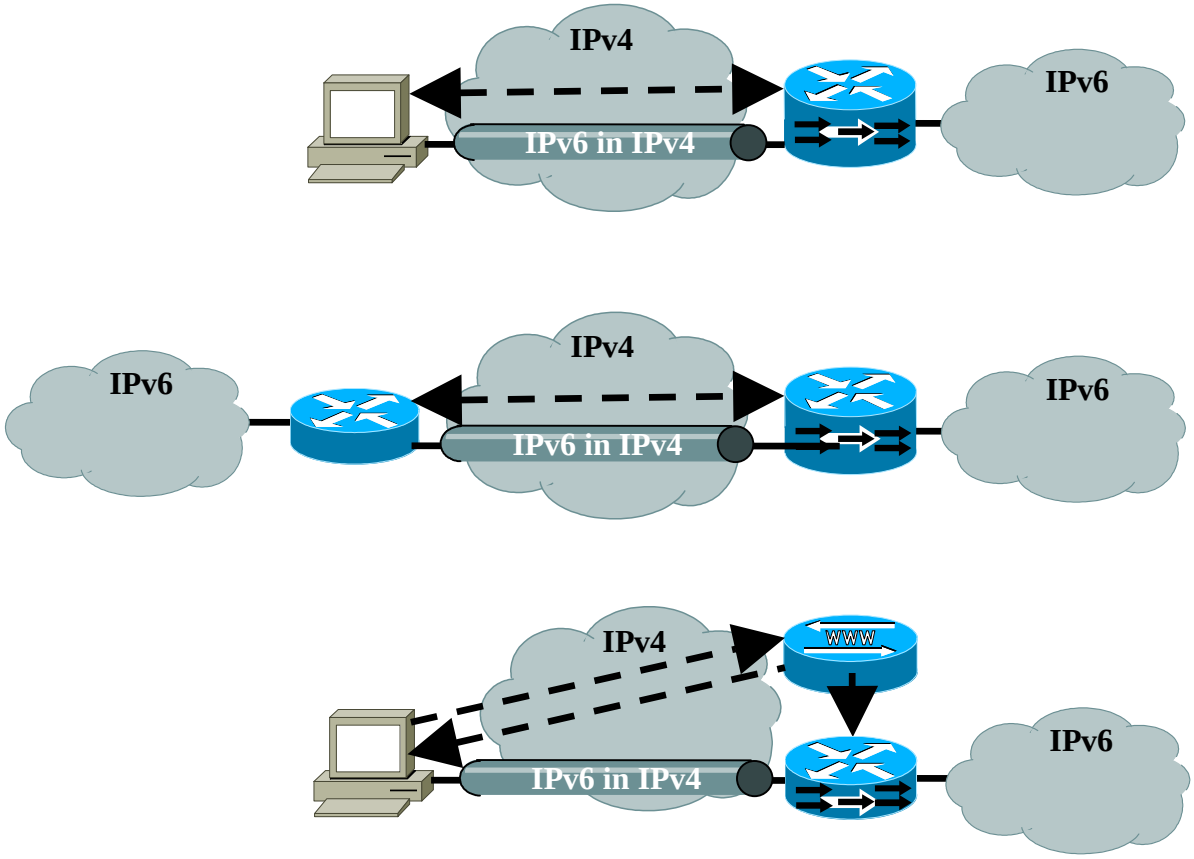

# Tunnel Broker with TSP

- Client requests:
  - A tunnel for one host
  - A tunnel for a network (routing implied)
    - With a prefix delegation
    - Without a prefix delegation (I have mine, but please announce it)
  - With routing information
    - I use RIP, BGP, OSPF…
  - With domain name information
    - The host name will be:
    - Inverse delegation
- Server can respond:
  - Here are the requested info
  - I'm full, cannot give you a new tunnel, please go to this other tunnel server (referer)
  - I can give you a host tunnel, but not a network tunnel
  - Here is the prefix, here is my bgp info (as number) …

# Tunnel Broker with TSP

- Enables:
  - Changes of data to be negotiated between the two parties: i.e. IPv4 client address change
    - IPv4 mobility but IPv6 stability of address and prefix
  - Generic Authentication: use a password, a Securid card, a public-key…
- Is implemented
  - In the boot sequence of a host
  - In host OS and router OS
- Combined with other tunnelling protocols:
  - DSTM as the Tunnel Setup Protocol for IPv4 in IPv6 tunnels
  - With UDP encapsulation for authenticated/controlled/billable way to cross NATs

# Tunnel Broker with TSP Applicability

- Host or Routers (networks) over an IPv4 network
- With or without NAT in the path
- Provides permanent address space to hosts and routers
- A change in IPv4 address do not change the IPv6 address
- Provides mobility on the IPv4 link-layer while preserving the IPv6 address
- In a corporate network
- In a provider network
- On the Internet
- Provides control, authentication, security, billing
- Provides a permanent address space

# Application-level gateways

- For applications using ALGs
  - ALG is dual-stack, connected to both IP networks
  - Clients can use one or the other IP protocol
  - Peers can use one or the other IP protocol
- Limitations of ALGs
  - Not good if the application protocol embed IP addressing
  - One application at a time.
  - Many applications are not designed for ALGs
- Examples
  - Email
  - Web

# Summary

- Many transition tools exists
  - To tunnel between IPv6 islands
  - To translate between IPv4 and IPv6
- Others are available, others will be defined
- None is for all possible scenarios
- Not all will succeed on the market
- Choose the right one for your scenario

# References


- RFC2766, Network Address Translation - Protocol Translation (NAT-PT), G. Tsirtsis, P. Srisuresh, ETF, 2000-02-01, http://www.normos.org/ietf/rfc/rfc2766.txt
- RFC2767, Dual Stack Hosts using the "Bump-In-the-Stack" Technique (BIS), K. Tsuchiya, H. Higuchi, Y. Atarashi, IETF, 2000-02-01, http://www.normos.org/ietf/rfc/rfc2767.tx
- RFC2893, Transition Mechanisms for IPv6 Hosts and Routers, R. Gilligan, E. Nordmark, 2000-08-01, http://www.normos.org/ietf/rfc/rfc2893.txt
- RFC3053, IPv6 Tunnel Broker, A. Durand, P. Fasano, I. Guardini, D. Lento, IETF, 2001-01-01, http://www.normos.org/ietf/rfc/rfc3053.txt
- RFC3056, Connection of IPv6 Domains via IPv4 Clouds, B. Carpenter, K. Moore, IETF, 2001-02-01, http://www.normos.org/ietf/rfc/rfc3056.txt

# References


- RFC3068, An Anycast Prefix for 6to4 Relay Routers, C. Huitema, IETF, 2001-06-01, http://www.normos.org/ietf/rfc/rfc3068.txt
- draft-vg-ngtrans-tsp-xx.txt, Tunnel Setup Protocol (TSP), M. Blanchet, R. Desmeules, A. Cormier,
- draft-vg-ngtrans-tsp-v6v4profile-xx.txt, IPv6 over IPv4 profile for Tunnel Setup Protocol (TSP), M. Blanchet, R. Desmeules, A. Cormier
- draft-blanchet-ngtrans-tsp-dstm-xx.txt, DSTM IPv4 over IPv6 tunnel profile for Tunnel Setup Protocol (TSP), M. Blanchet

# Case Studies

# Plan

- Isolated host without IPv4 NAT
- Isolated host with IPv4 NAT
- Mobile host
- Small network without IPv4 NAT
- Small network with IPv4 NAT
- Mobile network
- Large corporate network without IPv4 NAT
- Large corporate network with IPv4 NAT
- Provider network
- Addressing Plans
- Routing

# Considerations

- Always use native IPv6 if you can. (together with IPv4: dual stack approach)
- If you can’t, then transition mechanisms are used
- Tunnelling IPv6 in IPv4 works if no NAT in the path. A user or a network manager do not necessarily know in advance if there is one NAT or not in any possible path. UDP encapsulation is important considering the NAT installed base.

# Isolated Host without IPv4 NAT

- Dual-stack host on an IPv4 network without IPv4 NAT

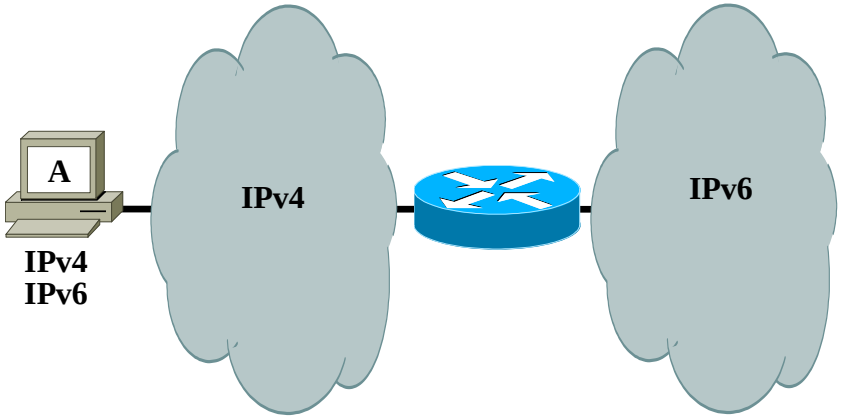

# Isolated Host without IPv4 NAT

- IPv6 in IPv4 Tunnelling is used:
  - Configured tunnels:
    - need manual changes everytime IPv4 address change.
  - 6to4:
    - IPv6 address will change everytime IPv4 address change
    - Need to find a 6to4 relay. May be far.
  - ISATAP:
    - Yes if on a corporate network. No if mobile. Will renumber when full IPv6 deployment is done. Not for servers.
  - TSP Tunnel Broker:
    - If available, yes. Could be mobile later or behind NAT and keep its IPv6 address (space).

# Isolated Host with IPv4 NAT

- Dual-stack host on an IPv4 network with IPv4 NAT

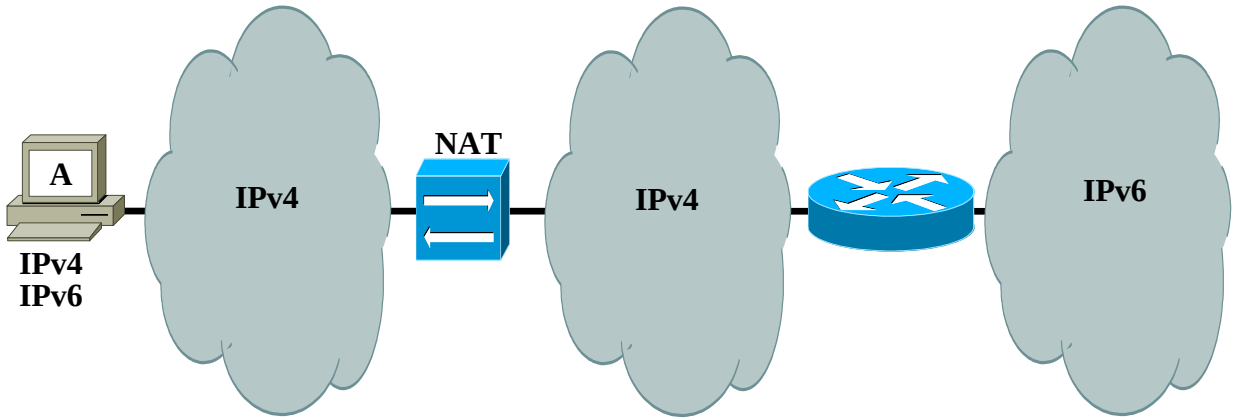

# Isolated Host with IPv4 NAT

- IPv6 in IPv4 Tunnelling is used:
  - Configured tunnels/6to4/ISATAP do not work.
  - Teredo:
    - If available, yes. If mobile and not behind a NAT, encapsulation tax and will change IPv6 address (renumbering). No prefix for any subnet behind the host.
    - For the provider of the service, no authentication/control of service, no possible billing.
  - TSP Tunnel Broker:
    - If available, yes. Could be mobile later or not behind NAT and keep its IPv6 address (space).

# Mobile Host

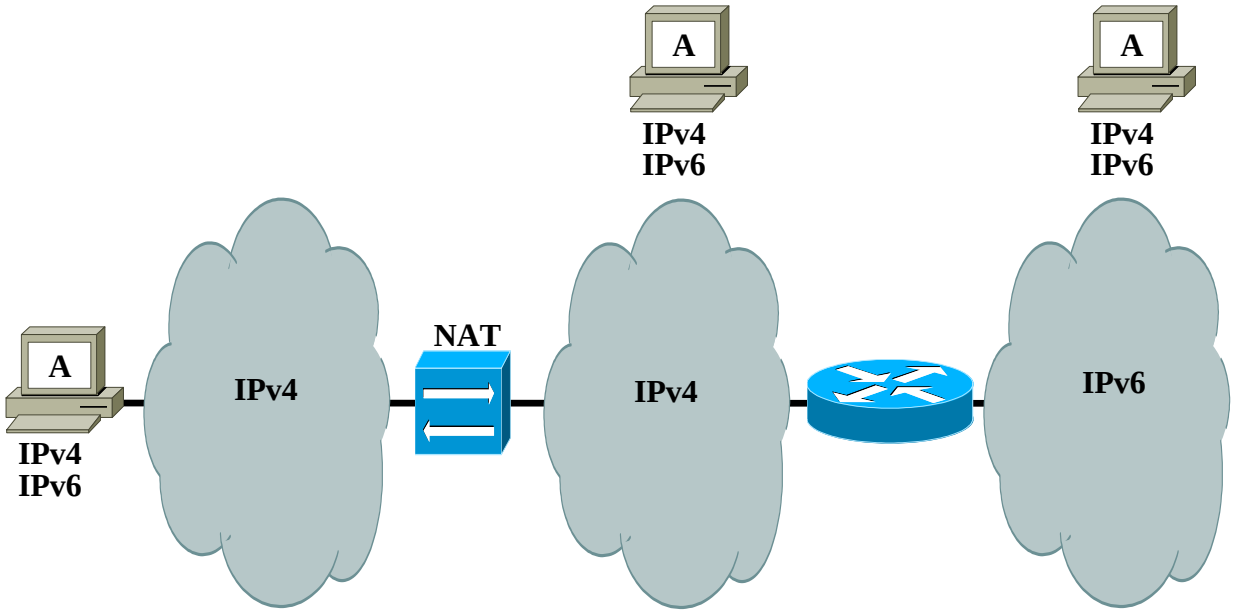

# Mobile Host

- Use IPv6 if home network is IPv6 native.
- Use MobileIPv6 if visiting network is IPv6 native
  - If implemented and configured in stack
  - Not efficient if source routing is prohibited in the visiting network
  - Needs MobileIPv6 configuration on the home network
- No solution with native and MobileIPv6 if the visiting network is not IPv6
- Since high probability of IPv4 NAT in some visiting network, then need a NAT-friendly transition mechanism:
  - Teredo: same issues as before.
  - TSP Tunnel Broker: If available, yes. Keeps its IPv6 address (space).

# Small Network without IPv4 NAT

- Network could be IPv6 only or IPv4/IPv6
- Gateway is dual-stack

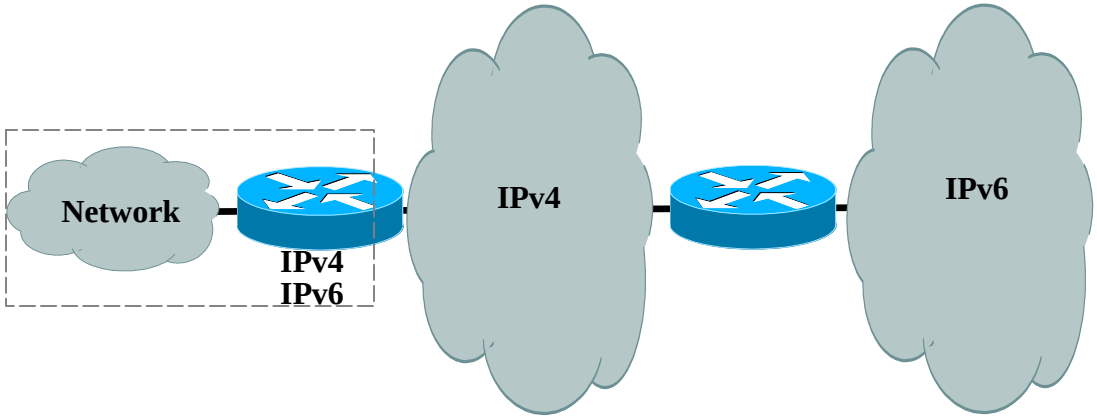

# Small Network without IPv4 NAT

- IPv6 in IPv4 Tunnelling is used if external network is IPv4:
  - Configured tunnels
    - need manual changes everytime IPv4 address change.
    - no IPv6 prefix for the network behind
  - 6to4:
    - IPv6 address will change everytime IPv4 address change
    - need to find a 6to4 relay. May be far.
  - ISATAP: only for inside. Does not help for outside.
  - TSP Tunnel Broker
    - if available, yes; could be mobile later or behind NAT and keep its IPv6 address (space)
- Hosts inside the network could be IPv6 only if gateway sends router advertisements
- If external network is IPv6, then native IPv6 is used

# Small Network with IPv4 NAT

- Network could be IPv6 only or IPv4/IPv6.
- Gateway is dual-stack.
- If gateway is the IPv4 NAT.

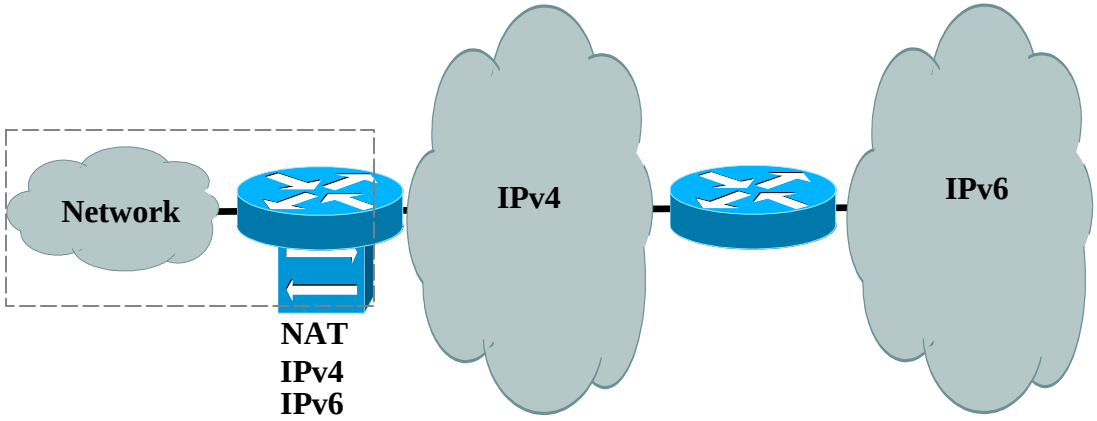

# Small Network with IPv4 NAT

- Gateway is dual-stack.
- If gateway is the IPv4 NAT, then it could be an IPv6 router.
  - IPv6 in IPv4 Tunnelling is used if external network is IPv4:
    - Configured tunnels:
      - need manual changes everytime IPv4 address change.
      - no IPv6 prefix for the network behind
    - 6to4:
      - IPv6 address space will change everytime IPv4 address change
      - need to find a 6to4 relay. May be far.
    - ISATAP: only for inside. Does not help for outside.
    - TSP Tunnel Broker:
      - if available, yes; could be mobile later or behind NAT and keep its IPv6 address (space)
  - Hosts inside the network could be IPv6 only if gateway sends router advertisements.

# Small Network with IPv4 NAT

- Network could be IPv6 only or IPv4/IPv6.
- Gateway is dual-stack.
- If gateway is not the IPv4 NAT

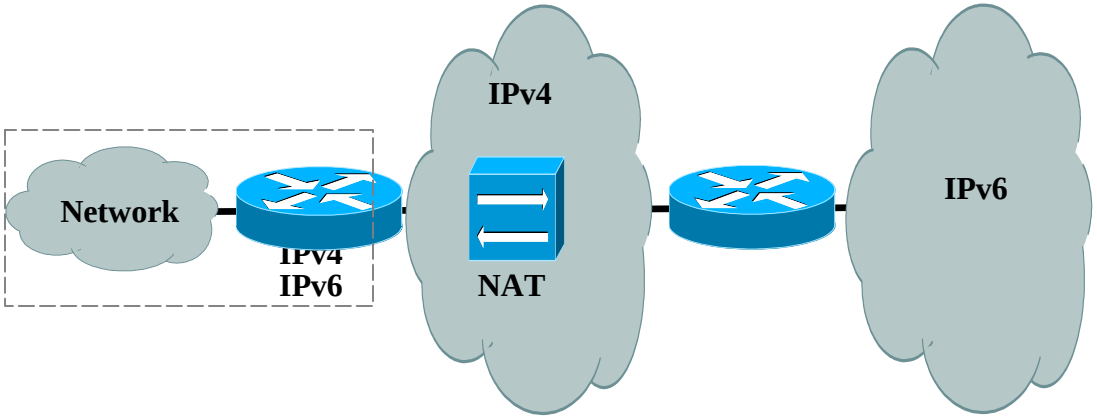

# Small Network with IPv4 NAT

- Gateway is dual-stack.
- If gateway is not the IPv4 NAT, then it needs NAT-friendly mechanism.
  - IPv6 in IPv4 Tunnelling is used if external network is IPv4:
    - Configured tunnels/6to4/ISATAP does not work
    - Teredo: do not provide prefix for the network behind. A host-based solution, not for the network.
    - TSP Tunnel Broker:
      - if available, yes; could be mobile later and keep its IPv6 address (space).
  - Hosts inside the network could be IPv6 only if gateway sends router advertisements.

# Small Network - Inside

- If possible, deploy native IPv6
- If not, then need tunnelling
  - Configured tunnels: manual, do not scale
  - 6to4: rely on IPv4 external address for IPv6 prefix. Renumbering when IPv4 address change.
  - ISATAP: good
  - TSP Tunnel Broker: would handle all cases (for example, if a NAT is used inside the network).

# Mobile Network

- Network could be IPv6 only or IPv4/IPv6.
- Gateway is dual-stack.

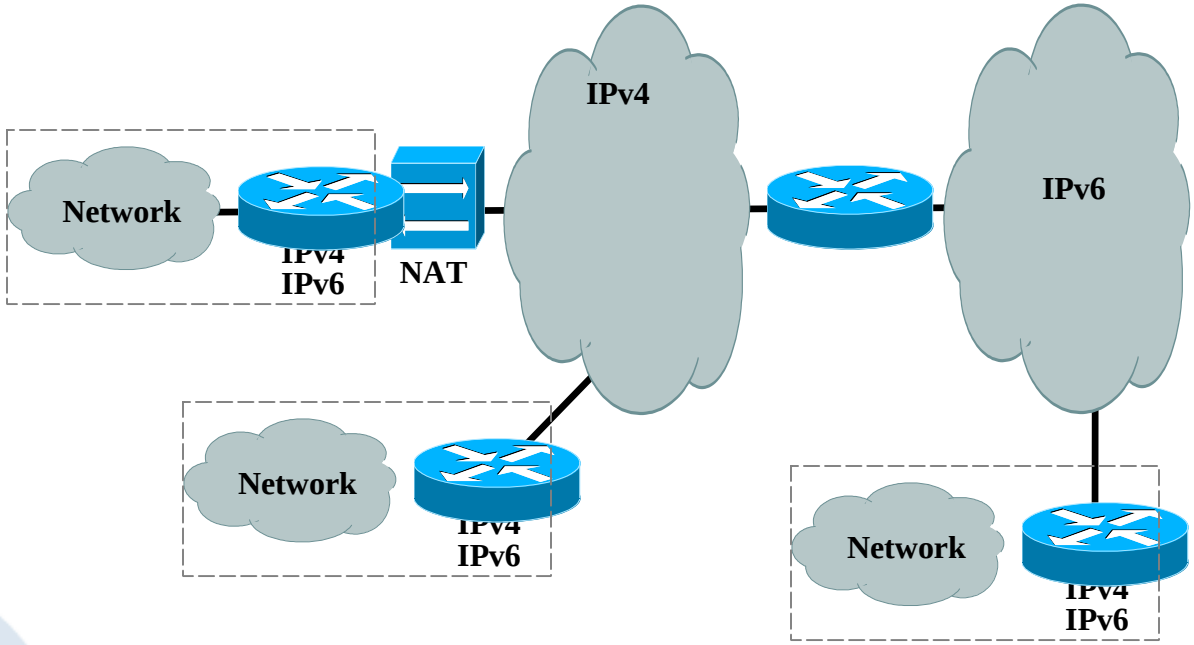

# Mobile Network

- Use IPv6 if home attached network is IPv6 native
- Use MobileIPv6 if attached network is IPv6 native
  - If implemented and configured in stack
  - Not efficient if source routing is prohibited in the visiting network
  - Needs MobileIPv6 configuration on the home network
  - No solution if the visiting network is IPv4
  - Do not give IPv6 prefix for the network behind.
  - No known implementations of MobileIPv6 networks.
- Since high probability of NAT in some visiting network, then need a NAT-friendly transition mechanism:
  - Teredo: same issues as before. No prefix given.
  - TSP Tunnel Broker: If available, yes. Keeps its IPv6 address (space).

# Large Corporate Network without IPv4 NAT

- Similar to small network without IPv4 NAT
- However:
  - 6to4 have many scaling limitations for large networks (eg: only one exit/entry point, large network renumbering).
  - Multiple kinds of islands if incremental deployment:
    - Isolated hosts
    - Isolated networks
  - ISATAP is only for hosts.
  - TSP Tunnel Broker works in these multiple configurations.

# Large Corporate Network with IPv4 NAT

- Similar to Small network with NAT
- However
  - 6to4 have many scaling limitations for large networks (eg: only one exit/entry point, large network renumbering).
  - Multiple kinds of islands if incremental deployment:
    - Isolated hosts
    - Isolated networks
  - ISATAP is only for hosts and would not help to reach external (since NAT)
  - TSP Tunnel Broker works in these multiple configurations.

# Provider Network

- Deploy native IPv6 if possible
- If not, incremental deployment with tunnelling
- To reach customers
  - Configured tunnels works but very painful for large scale
  - 6to4/ISATAP/Teredo could not be used
  - TSP Tunnel Broker enables the incremental deployment
- Could offer IPv6 services
  - TSP Tunnel Broker for customer hosts and networks
  - Teredo server/relay for customer hosts
  - 6to4 relay

# Nodes Configuration

- If native IPv6
  - Use autoconfiguration (with router advertisements) unless need to push more information to the node (through DHCPv6)
    - Set the valid and preferred lifetimes to non-infinite. (order of days/weeks)
    - Use some DNS discovery mechanisms, such as site-local well-known addresses.

# Address Plan

- If you are a top-level (tier-1) for assignments, start bits from left to right.
  - 0000, 8000, 4000, C000, ...
- If you are an intermediate level for assignments:
  - If you want maximum flexibility with maximum aggregation in the future, use centermost bits
    - identify the best target prefix length (ex: /40)
    - start using bits growing on both sides of the boundary
- If you are a end-level for assignments, start bits from right to left.
  - Subnets always have /64; should be able to allocate all bits. routing aggregation is the only concern, if it is in your case.
- See RFC3531

# Address Plan for a Tier-1 Provider Network

- Maximum aggregation of customers
- Different sizes of customers (Tier-X providers, Enterprise, Users)
- Start numbering with bits from left to right
- Example: 3ffe:0b00::/32
  - 3ffe:0b00:0000::
  - 3ffe:0b00:8000::
  - 3ffe:0b00:4000::
  - 3ffe:0b00:C000::
  - ...

# Address Plan for an Enterprise Provider Network

- Topology defines number of leaf networks
  - Identify the maximum size at distribution level
- Use centermost bits allocation
  - If size of distribution level grows, still have room without breaking aggregation
  - If more distribution level networks grows, still have room without breaking aggregation
- Example: 3ffe:0b00:0100::/40
  - Max size at distribution level = /44  (i.e. 44-40 = 4 => 16 leaf networks)
  - 3ffe:0b00:0100::/44
  - 3ffe:0b00:0108::/44
  - 3ffe:0b00:0110::/44
  - 3ffe:0b00:0118::/44

# Address Plan for an Enterprise Network

- Topology defines number of leaf subnets
  - Identify the maximum size at distribution level
- Use centermost bits allocation
  - If size of distribution level grows, still have room without breaking aggregation
  - If more distribution level networks grows, still have room without breaking aggregation
- Example: 3ffe:0b00:0001::/48
  - Max size at distribution level = /60  (i.e. 64-60 = 4 => 16 leaf subnets)
  - 3ffe:0b00:0001:0000::/64
  - 3ffe:0b00:0001:0008::/64
  - 3ffe:0b00:0001:0010::/64
  - 3ffe:0b00:0001:0018::/64
  - ...

# Routing in an Enterprise Network

- If RIP is fine for your network, then RIPng should be fine too.
- OSPFv3 manages a different Link-state database than OSPFv2. Ships in the night. Easy to deploy incrementally.
- ISIS uses the same link-state database. Topologies must be synchronized. Unless vendor support.

# Routing in a Provider Network

- BGP
  - Peering
    - separate
    - combined
      - transport
    - assess use of link-local
  - Filtering
    - non valid addresses (fec0::, fe80::, ...)
    - 6bone prefixes? (3ffe::/16)
    - relays? (6to4 2002::/16 address)

# DNS

- Records
  - Only put AAAA records when the server is IPv6 reachable for all cases
  - Must be an IPv6 stable address
- Transport
  - IPv6 transport needed
    - if your users have IPv6 resolvers
    - IPv6 only nodes
- Inverse mapping
  - Defines both ip6.int and ip6.arpa

# Servers and Routers

- Configuration of servers
  - Use static addresses for interfaces
  - Do not use eui-64, in order to avoid binding of the address and the interface card, in case of hardware interface change
  - Make sure the IPv6 part is reachable from all the users in all cases who are using this server
- Configuration of routers
  - Use static address for interfaces
  - Do not use eui-64, in order to avoid binding of the address and the interface card, in case of hardware interface change
  - Memory usage for multiple topology databases
  - RA with non-infinite valid and preferred lifetimes

# Summary

- Choose the right transition tool
- Plan the deployment
  - Nodes configuration
  - Address plan
  - DNS
  - Routing
  - Interface configuration

HEXAGO

# DNS, Apache and tunnel broker demos

# Building DNS files

- Using BIND as example
- www.example.org = 192.0.2.1 and 3ffe:b00:1::1

- Name to IP address
  - name to IPv6 address = AAAA record
  - name to IPv4 address = A record
  - in zone file example.org:
    - www.example.org. IN A 192.0.2.1
    - www.example.org. IN AAAA 3ffe:b00:1::1
- Works in Bind since 4.9.X

# Building DNS files (cont.)

- IP address to name using PTR record
  - ipv4 address to name with in-addr.arpa
  - ipv6 address to name with ip6.arpa

  - 1.2.0.192.in-addr.arpa IN PTR www.example.org.
  - 1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.0.0.0.0.0.b.0.e.f.f.3.ip6.arpa IN PTR www.example.org.

# Building DNS files (cont.)

- Additional considerations
  - named.conf
    - listen-on {192.0.2.1;};
    - listen-on-v6 {3ffe:b00:1::1; };
    - masters {3ffe:b00:1::1;};
    - allow-transfer {3ffe:b00:1::1;};
  - IPv6 transport available in
    - Bind 9.X
    - Bind 8.4
    - Bind 8.X with ipv6 patch

# Building DNS files (cont.)

- Client
  - IPv6 transport of DNS requests
    - /etc/resolv.conf
      - nameserver 3ffe:b00:1:1::2
- Testing
  - dig name a
  - dig name aaaa
  - dig name any

# Enabling IPv6 on Apache Web Server

- Apache 2.0 or Apache 1.3 with ipv6 patch
- httpd.conf
  - ipv4: Listen 192.0.2.1:80
  - ipv6: Listen [3ffe:b00:1:1::1]:80
  - <VirtualHost 192.0.2.1 3ffe:b00:1:1::1>

# Using a Tunnel Broker

- IPv6 over an IPv4 cloud
- Need a dual stack and the TSP client
- Using Freenet6.net (http://www.freenet6.net)
  - with TSP client on freebsd (available for windows, *bsd, linux, solaris, qnx...)
  - tspc.conf
    - client_v4=auto  # use the IPv4 address on the interface
    - userid=anonymous|myusername
    - passwd=mypasswd  #if not using anonymous
- # tspc -v
- Whenever your ipv4 address change, do tspc again.
  - It keeps your IPv6 address and prefix
  - It re-establishes the IPv6 tunnel
  - Looks like dhcp for IPv6 (over IPv4 link-layer)