

Deploying an IPv6 Broadband Service

Benoit Lourdelet

Cisco Systems

Cisco IOS IPv6 Technical Marketing Engineer

blourdel@cisco.com

A Today's Network Infrastructure

Cisco.com

- **MPLS technology selected as existing core infrastructure**

Current services are MPLS/VPN, L2 services over MPLS, Circuit over MPLS

- **Pure IPv4 core where a resistance to MPLS is displayed**

- **IP services portfolio**

Enterprise: Lease Lines

Home Users/SOHO: ADSL, ETTH, Dial

Data Center: Web hosting, servers,...

Deployment Activities

Cisco.com

- **IPv6 Core Network Enhancement**
 - IPv6 Provider Edge Routers (6PE) over MPLS
 - Dual Stack
- **Broadband Access Networks**
 - IPv6 over broadband data link layers
 - IPv6 address allocation guidelines
 - IPv6 AAA Radius
 - IPv6 auto-configuration – Prefix Delegation & Stateless DHCP
- **Data Center**
 - IPv6 on Layer 3 infrastructure

Dual Stack IPv4-IPv6 Infrastructure

Cisco.com

- It is generally the goal when IPv6 traffic and users will be rapidly increasing
- May not necessarily apply to the overall infrastructure. One may begin on network's portion such as Campus or Access or core networks
- Network design must be well planned
 - Memory size to handle the growth for both IPv4 & IPv6 routing tables
 - IGP options & its management: Integrated versus "Ships in the Night"
 - Full network upgrade impact
- IPv4 and IPv6 Control & Data planes should not impact each other

Pre-existing MPLS infrastructure

Cisco.com

- **If MPLS being already deployed for IPv4 services, 6PE is the preferred scenario**

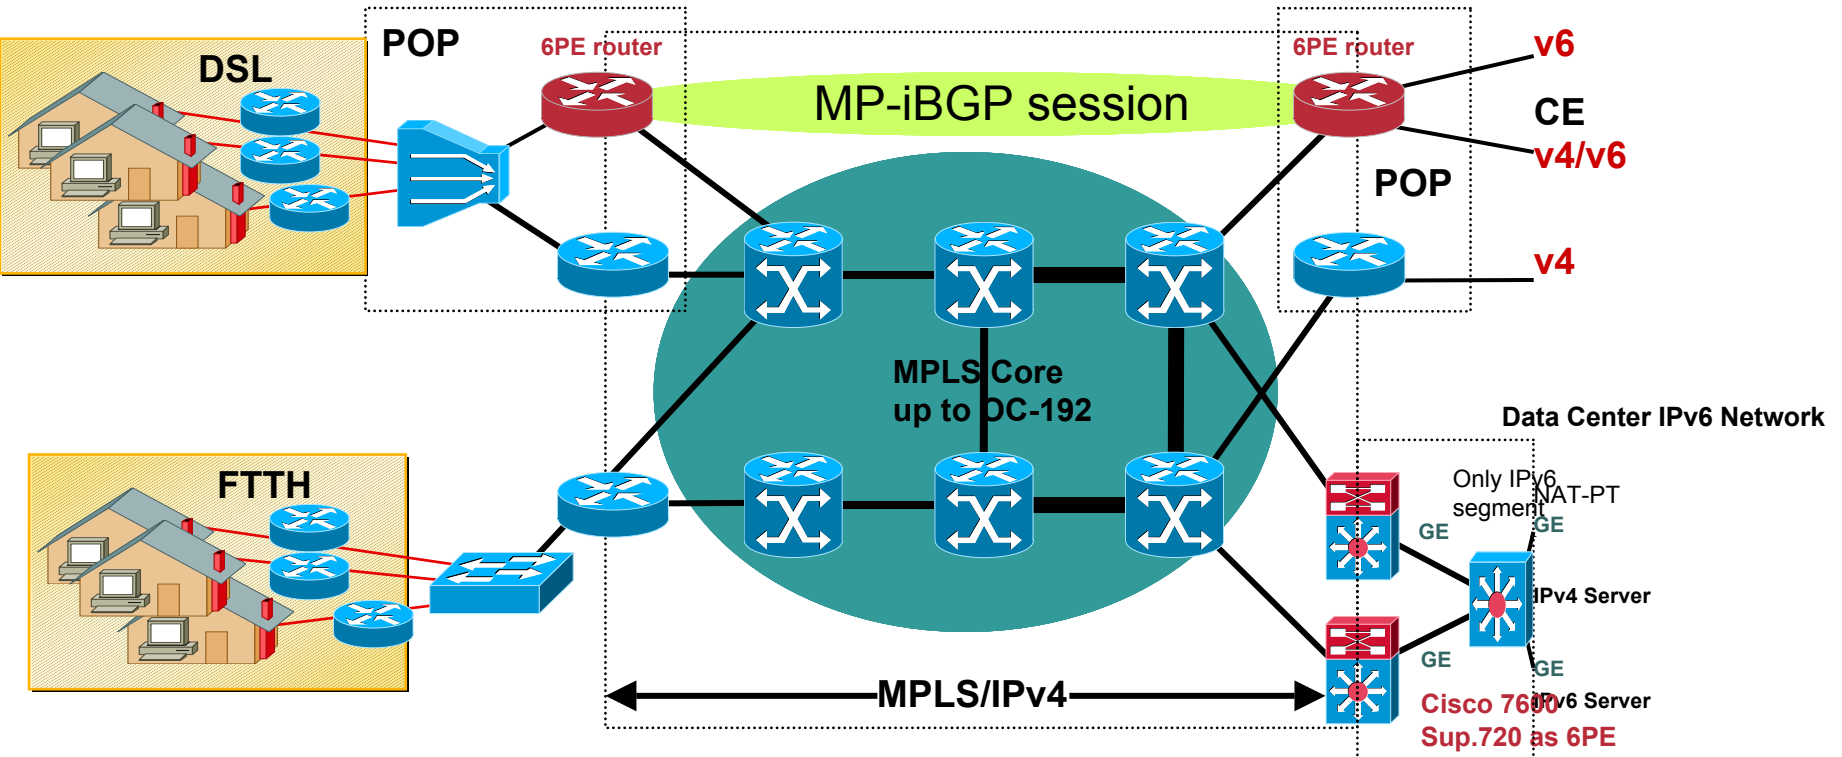
IPv6 POPs can be installed one by one (software upgrade or new PE router) – Cost of deployment is under control

IPv6 prefix `::/48` can be assigned from `::/32`

`draft-ietf-ngtrans-bgp-tunnel`

Minimum Infrastructure Upgrade for 6PE

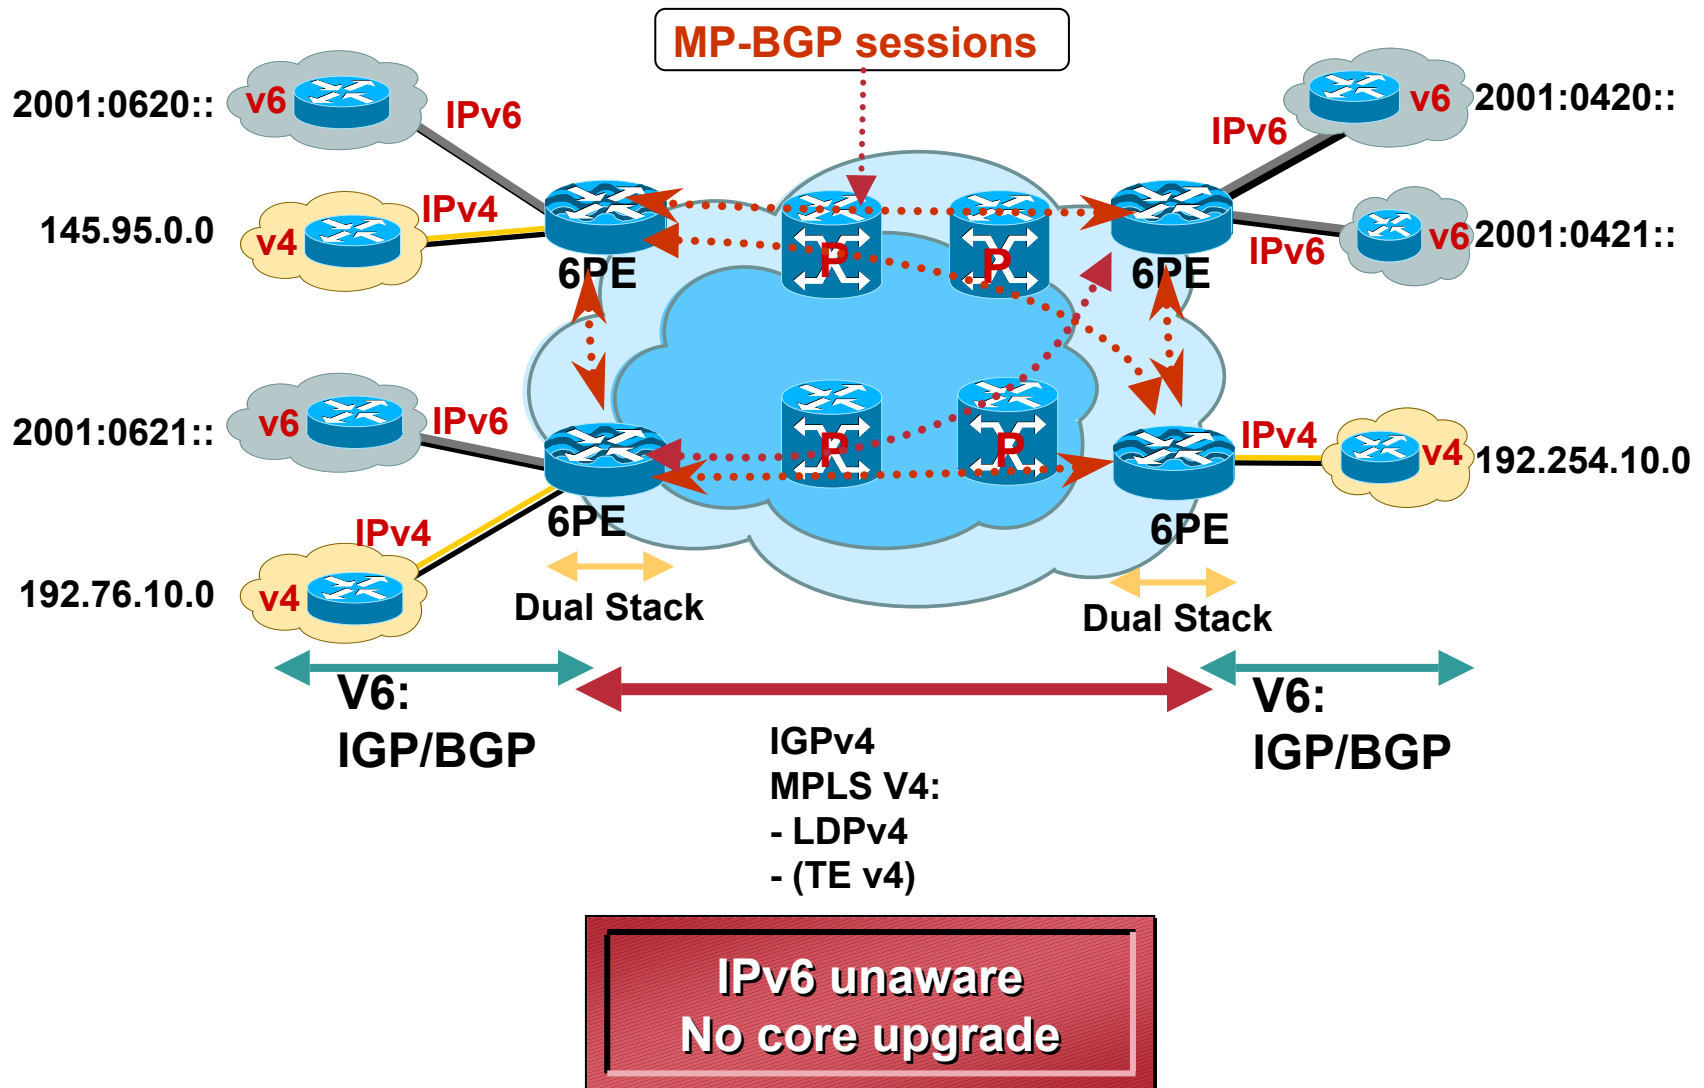
Cisco.com



- **MPLS/IPv4 Core Infrastructure is IPv6-unaware**
- **PEs are updated to support Dual Stack/6PE**
- **IPv6 reachability exchanged among 6PEs via iBGP (MP-BGP)**
- **IPv6 packets transported from 6PE to 6PE inside MPLS**

6PE Overview

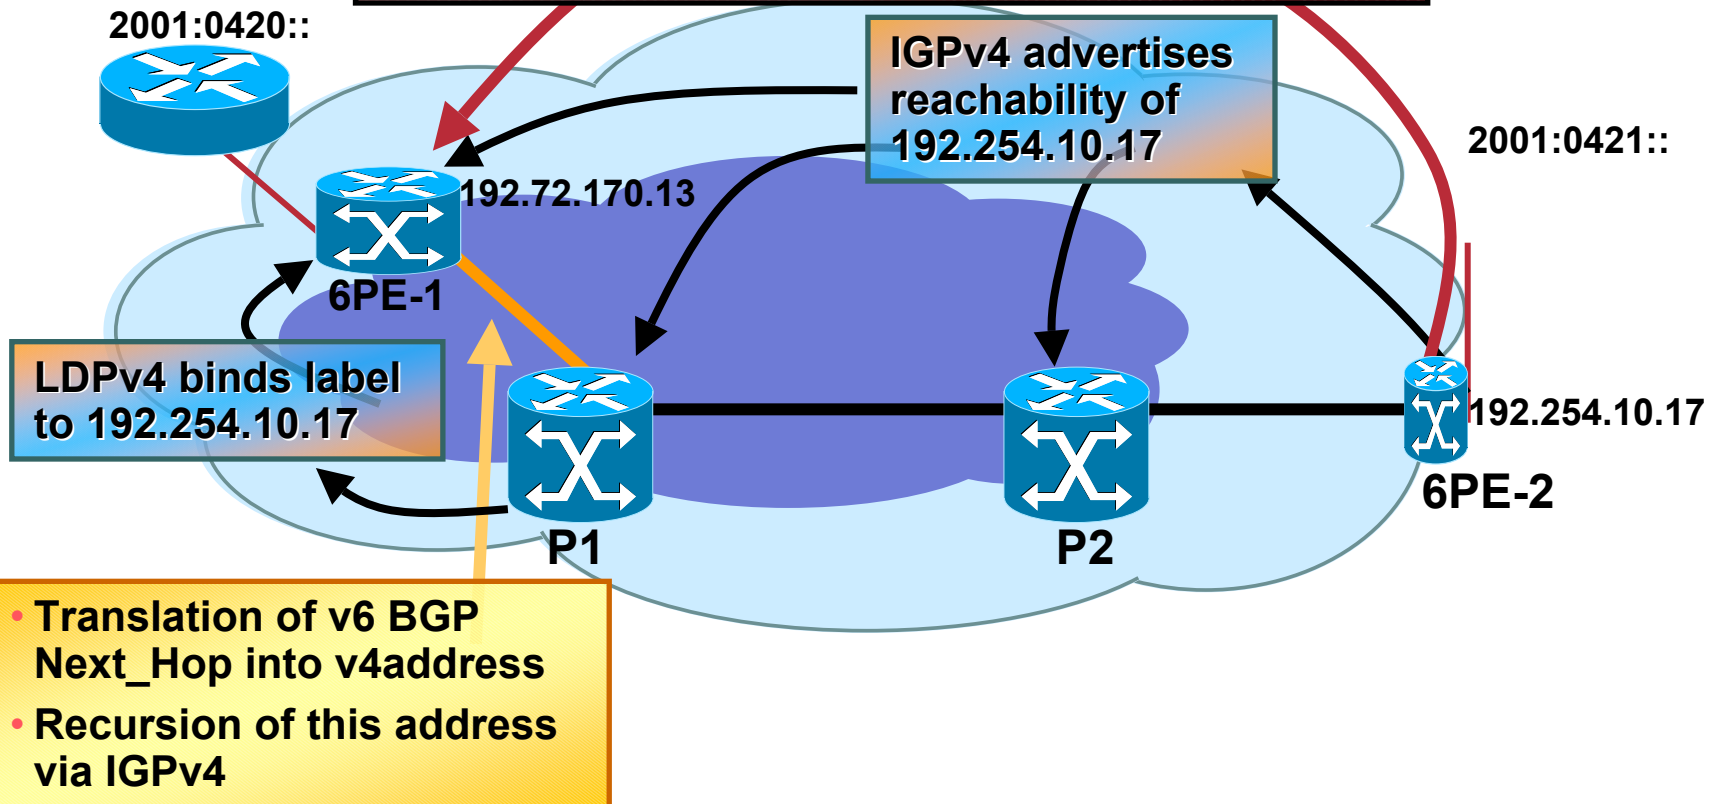
Cisco.com



6PE Routing


Cisco.com

MP-BGP advertises 2001:0421:::
and binds a (2nd level) label
IPv6 Next Hop is an IPv4 compatible IPv6 address
built from 192.254.10.17



6PE Forwarding

Cisco.com



6PE Forwarding (6PE-1)

Cisco.com

IPv6 Forwarding and Label Imposition:

- 6PE-1 receives an IPv6 packet
- Lookup is done on IPv6 prefix
- Result is:

Label binded by MP-BGP to 2001:0421::

Label1 binded by LDP/IGPv4 to the IPv4 address of BGP Next Hop (6PE-2)

2001:0420::

IPv6 packet
to 2001:0421::

192.72.170.13

6PE-1

2001:0421::

6PE-2

192.254.10.17

LDP/IGPv4
label1 to 6PE-2

MP-BGP label
To 2001:0421::

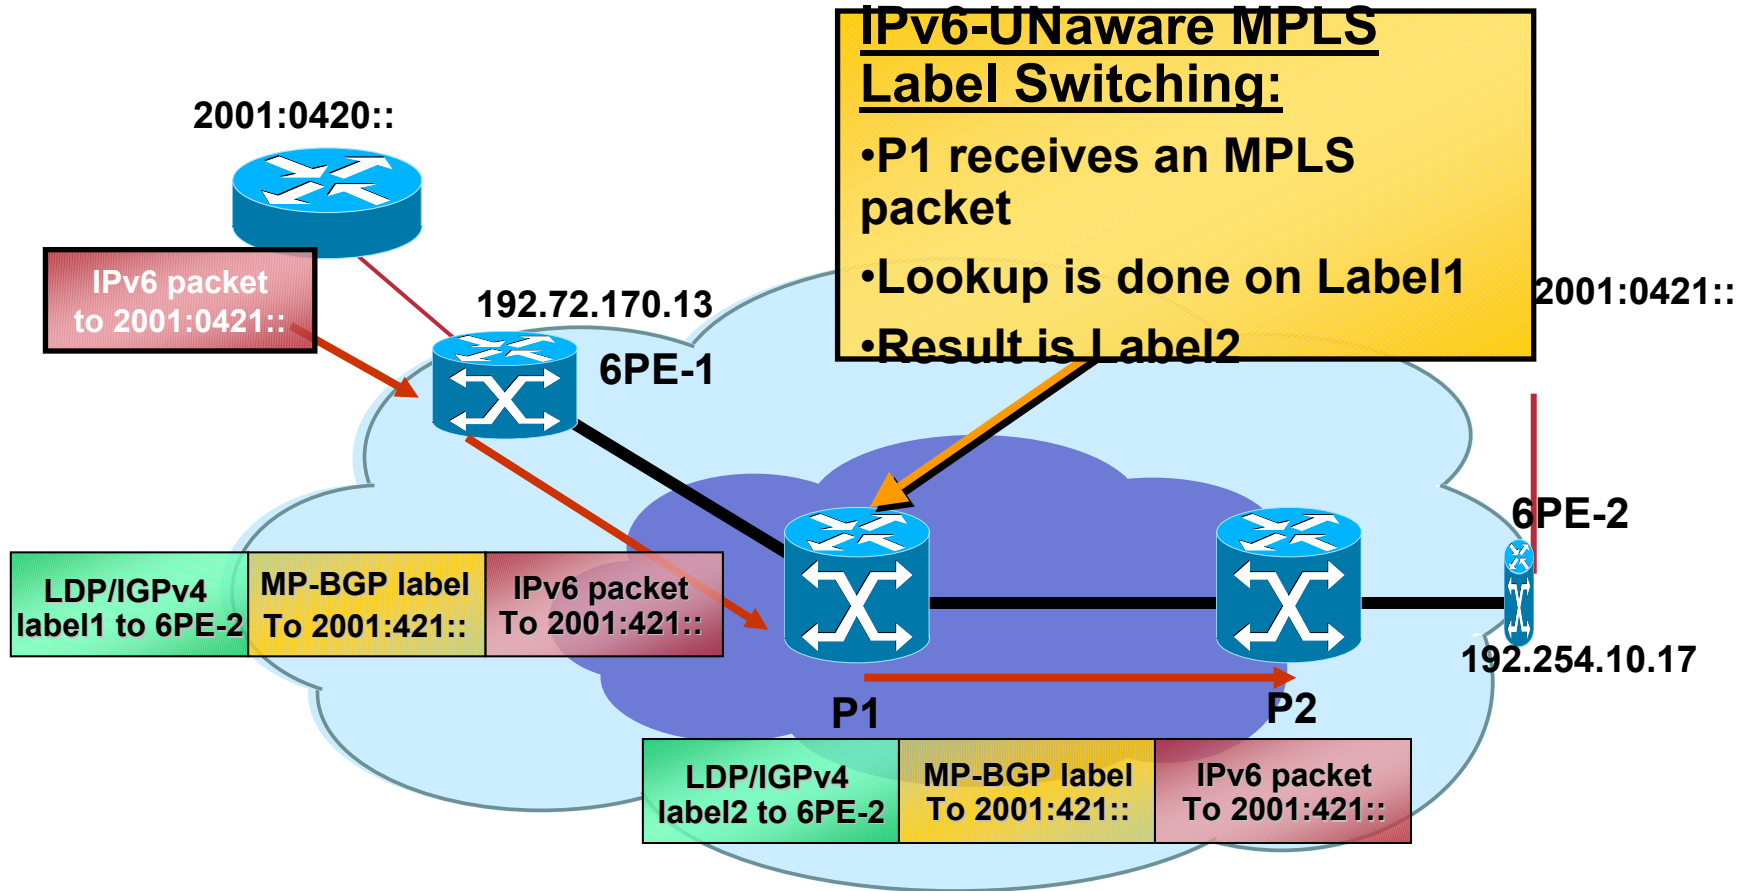
IPv6 packet
To 2001:0421::

P1

P2

6PE Forwarding (P1)

Cisco.com



6PE Forwarding (P2)

Cisco.com

IPv6-UNaware MPLS

Label Switching:

- P2 receives an MPLS packet
- Lookup is done on Label2
- Result includes Pop label (PHP)

2001:0420::

IPv6 packet
to 2001:0421::

192.72.170.13

6PE-1

LDP/IGPv4
label1 to 6PE-2

MP-BGP label
To 2001:421::

IPv6 packet
To 2001:421::

P1

P2

LDP/IGPv4
label2 to 6PE-2

MP-BGP label
To 2001:421::

IPv6 packet
To 2001:421::

M P
- V
B 6
G p
P a
c k
e t

2001:0421::

6PE-2

192.254.10.17

6PE Forwarding (6PE-2)

Cisco.com

MPLS Label Pop and IPv6 Forwarding :

- 6PE-2 receives an MPLS packet
- Lookup is done on Label
- Result is:

Pop the label & switch

IPv6

2001:0421::

packet

6PE-2

192.254.10.17

MPLS
- BGP
P
I
a
b
t

2001:0420::

IPv6 packet
to 2001:0421::

192.72.170.13

6PE-1

LDP/IGPv4
label1 to 6PE-2

MP-BGP label
To 2001:421::

IPv6 packet
To 2001:421::

P1


P2

LDP/IGPv4
label2 to 6PE-2

MP-BGP label
To 2001:421::

IPv6 packet
To 2001:421::

6PE configuration



```
ip v6 cef
!
mpls label protocol ldp
```

```
mpls ipv6 source-  
interface Loopback0
```

```
mpls ldp router-id  
loopback0  
!  
interface Loopback0  
ip address 10.10.20.2  
255.255.255.255  
ipv6 address 2003::1/64  
!
```

```
Router bgp 100  
!  
neighbor 10.10.20.1  
remote-as 100  
neighbor 10.10.20.1  
update-source Loopback0  
!  
address-family ipv6  
neighbor 10.10.20.1  
activate
```

```
neighbor 10.10.20.1  
send-label
```

```
redistribute connected  
redistribute rip ripv6CE1  
exit-address-family
```

```
ip cef  
!  
mpls label protocol ldp  
!  
interface Serial2/0  
ip address 10.10.10.2 255.255.255.252  
ip router isis  
tag-switching ip  
!
```


Show bgp ipv6 <ipv6-prefix>

Cisco.com

```
6PE> show bgp ipv6 2003:1:1:30::/64
```

```
BGP routing table entry for 2003:1:1:30::/64, version 2
```

```
Paths: (1 available, best #1, table Global-IPv6-Table)
```

```
Not advertised to any peer
```

```
Local
```

```
  ::FFFF:10.10.20.1 (metric 10) from 10.10.20.1 (192.168.254.1)
```

```
    Origin incomplete, metric 0, localpref 100, valid,  
    internal, best
```

Show bgp ipv6 neighbor

Cisco.com

```
6PE> show bgp ipv6 neighbors 10.10.20.1
```

```
BGP neighbor is 10.10.20.1, remote AS 100, internal link  
BGP version 4, remote router ID 192.168.254.1  
BGP state = Established, up for 00:04:07  
Last read 00:00:07, hold time is 180,  
Neighbor capabilities:
```

```
Route refresh: advertised and received(old & new)
```

```
Address family IPv6 Unicast: advertised and received
```

```
ipv6 MPLS Label capability: advertised and received
```

```
For address family: IPv6 Unicast
```

```
BGP table version 2, neighbor version 2
```

```
Index 1, Offset 0, Mask 0x2
```

```
Route refresh request: received 0, sent 0
```

```
Sending Prefix & Label
```

```
2 accepted prefixes consume 144 bytes
```

```
Prefix advertised 1, suppressed 0, withdrawn 0
```

```
Number of NLRI in the update sent: max 1, min 0
```


Show ipv6 route

Cisco.com

```
6PE> show ipv6 route
```

IPv6 Routing Table - 4 entries

Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP

I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea

```
B  2003:1:1:30::/64 [200/0]
    via ::FFFF:10.10.20.1, IPv6-mp1s
L  2003::205:32FF:FEC3:40E1/128 [0/0]
    via ::, Loopback0
C  2003::/64 [0/0]
    via ::, Loopback0
L  FE80::/64 [0/0]
    via ::, Null0
```


Deployment Activities

Cisco.com

- **IPv6 Core Network Enhancement**
 - IPv6 Provider Edge Routers (6PE) over MPLS
- **Broadband Access Networks**
 - IPv6 over broadband data link layers
 - IPv6 address allocation guidelines
 - IPv6 AAA Radius
 - IPv6 auto-configuration – Prefix Delegation & Stateless DHCP
- **Data Center**
 - IPv6 on Layer 3 infrastructure

Topology


Cisco.com



Data Link Layers

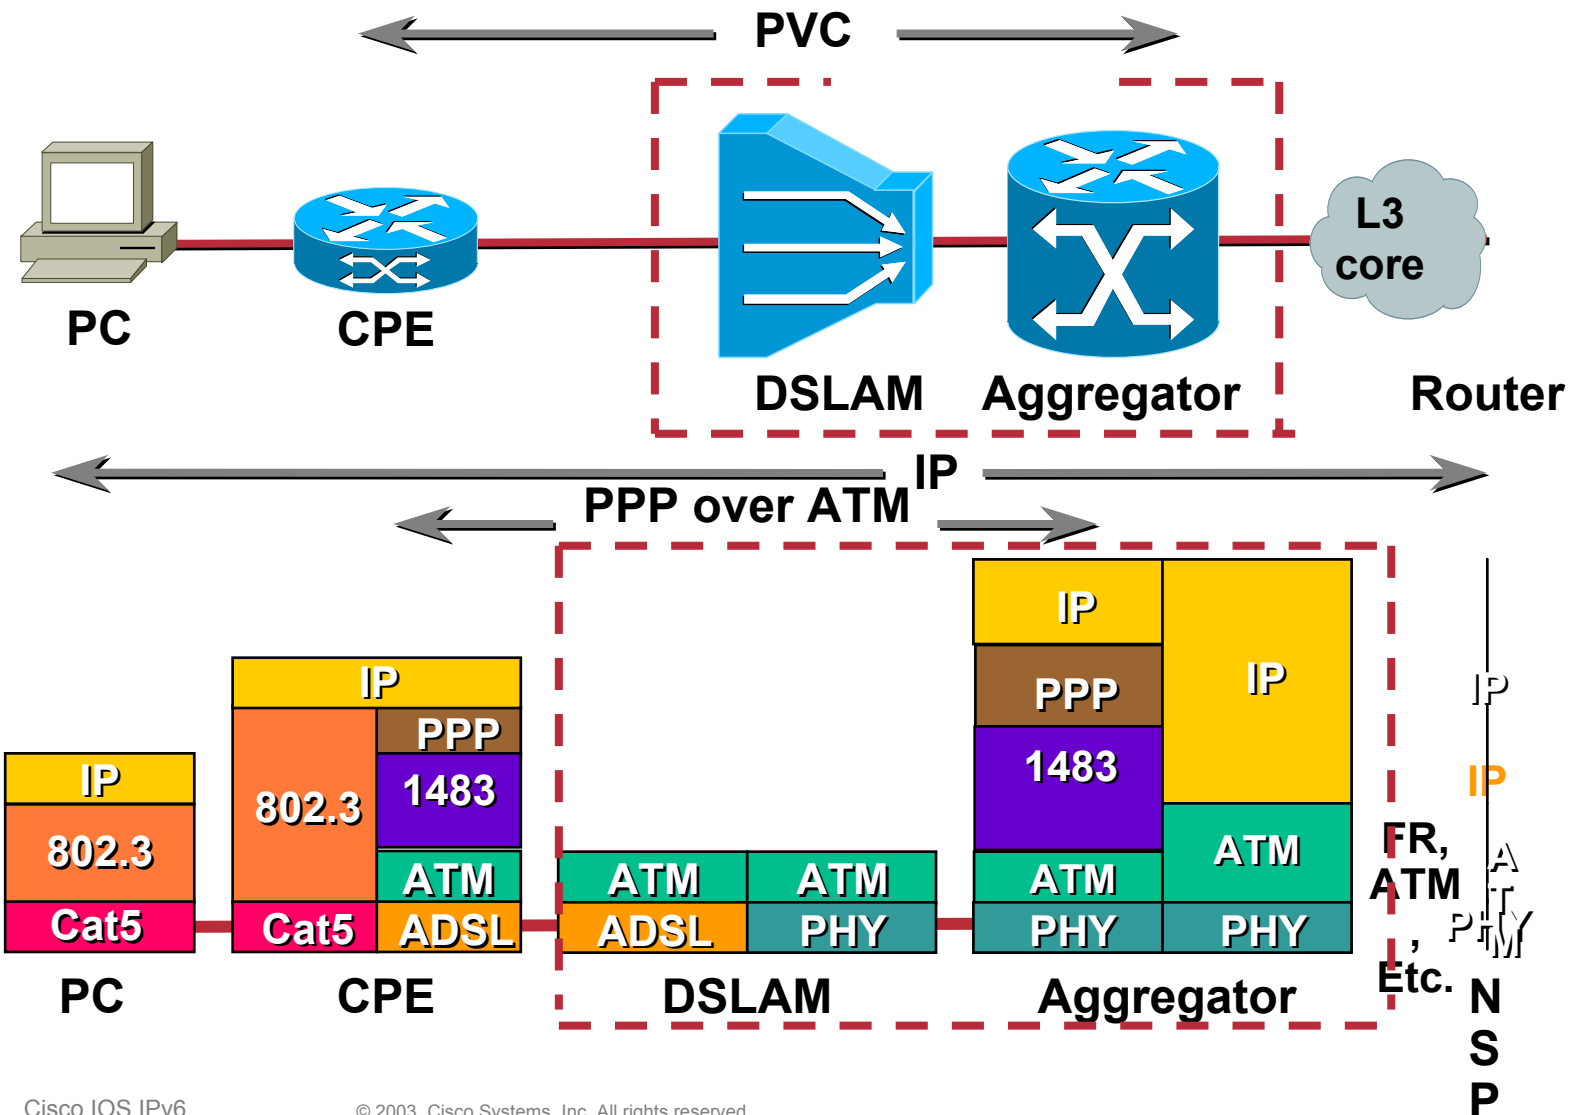
Cisco.com

- **Dial/ISDN**
PPP
- **Ethernet-To-The-Home**
Ethernet
- **802.11 (WiFi) Hot Spots**
Ethernet like
- **ADSL**
ATM RFC 1483 Routed
ATM RFC 1483 Bridged (RBE)
PPPoA
PPPoE
- **Available from Cisco IOS routers running 12.2T and 12.2B releases**




Protocol Stack - PPP over ATM

Cisco.com



PPP over ATM configuration



```

interface FastEthernet0
 ipv6 address 3FFE:ffff:123:1999::1/64
!
interface Atm 0
 pvc 1/23 encapsulation aal5mux ppp dialer
 dialer pool-member 1
!
interface dialer1
 encapsulation ppp
 dialer pool 1
 ipv6 address autoconfig
 ppp authentication chap foo
 ppp chap hostname user@domain.net
 ppp chap password 7 1111111111
 ppp ipcp address accept
!
ipv6 route ::/0 Dialer1
!
    
```

```

!
 vpdn enable
!
 vpdn-group 1
  accept dialin l2tp virtual-
  template 1 remote sp_lac
  local name lns
!
 interface Loopback0
  ipv6 address
  3ffe:ffff:bbbb::1/64
!
 interface Virtual-Template1
  ipv6 enable
  ipv6 mtu 1480
  no ipv6 nd suppress-ra
  ppp authentication chap default
!
 radius-server host 172.22.66.16
!
    
```

AAA/RADIUS

Cisco.com

- **Cisco Vendor Specific Attributes**
IPv6 Prefix, IPv6 Route, IPv6 ACL (Input & Output)
- **RADIUS and IPv6 (RFC3162)**
Framed-IPv6-Prefix
Framed-IPv6-Route
Framed-IPv6-Pool
NAS-IPv6-Address
Login-IPv6-Host
Framed-Interface-Id
- **On Cisco IOS, RADIUS transport is IPv4 as today most Radius server are used for both protocols**
IPv6 should be added later
- **IPv6 AAA available on Cisco IOS**
Cisco VSA available now from Cisco IOS 12.2T and 12.2B
RFC 3162 available from upcoming Cisco IOS 12.3T

AAA per-user attributes

Cisco.com

- **Route#**

Installs a per-user static route in the RIB

cisco-avpair="ipv6:route=3ffe:c00:1::/48"

- **Prefix#**

Adds the prefix to RA's sent out the interface, and adds a route in the RIB.

Restricted to /64 prefixes.

cisco-avpair="ipv6:prefix=3ffe:c00:2::/64"

- **ACL**

cisco-avpair="ipv6:inACL=permit 3ffe:c00:2::/64"

- **Framed-Interface-Id**

Framed-Interface-Id=0:0:0:0001

Included in accounting records

IPv6 prefix-pools

Cisco.com

- **Normal prefix pools:**

ipv6 prefix-pool foo 3ffe:c00:1::/48 64

A Separate /64 is assigned each user/interface. The prefix is advertised in RA's and a route is installed in the RIB.

- **Shared prefix pools:**

ipv6 prefix-pool foo 3ffe:c00:2::/64 128 shared

/64 prefix is shared between all users of the pool. The same /64 prefix is advertised in RA's out all interfaces. The user gets an /128 based on the prefix and his Interface-Identifier. A route in the RIB is installed only for the /128.

IPv6 Address Allocation Guidelines

“...recommends the assignment of /48 in the general case, /64 when it is known that one and only one subnet is needed...”

Cisco.com

RFC3177

IAB/IESG Recommendations on IPv6 Address Allocations to Sites

Policy Implementation

Cisco.com

- **Give Home/SOHO a permanent /64 – single link**
- **Give Home/SOHO a permanent /48**
- **Short-lived /64 from a prefix-pool**

A Separate /64 is assigned each user/interface. The prefix is advertised in RA's and a route is installed in the RIB.

- **Short-lived /128 from a shared prefix-pool**

/64 prefix is shared between all users of the pool. The same /64 prefix is advertised in RA's out all interfaces. The user gets an /128 based on the prefix and his Interface-Identifier. A route in the RIB is installed only for the /128.

- **For some users set the Interface-ID explicitly**

Give home users a permanent /64 – single link

Cisco.com

- **Use:** for single PC or network with only one link
- **AAA static prefix attribute.** Interface-Id attribute to specify the complete address
- **CPE:** single PC, proxy RA, or configured router

AAA config:


Auth-Type = Local, Password = “foo”

User-Service-Type = Framed-User,

Framed-Protocol = PPP,

cisco-avpair = “ipv6:prefix=3ffe:c00::/64

Framed-Interface-Id = 0:0:0:1



Address Assignment – short-lived /128

Cisco.com

- **Use:** for single PC only. Allows one address
- **/64 prefix shared** between all users of the pool
- **AAA interface-id attribute** can be used to specify complete address
- **NAS:** IPv6 shared prefix pools
- **CPE:** Single PC


AAA config:

Auth-Type = Local, Password = “foo”

User-Service-Type = Framed-User,

Framed-Protocol = PPP,

cisco-avpair = “addr-pool=“foo-shared”




Address Assignment – short-lived /64

Cisco.com

- **Use:** for single PC or very simple network
- **NAS:** IPv6 prefix pool
- **CPE:** Proxy-RA/multi-link subnet/bridging
Renumbering issues

AAA config:

Auth-Type = Local, Password = “foo”
User-Service-Type = Framed-User,
Framed-Protocol = PPP,



Cisco.com




Address Assignment – permanent /48

Cisco.com

- Use: whole site -supports multiple links
- AAA prefix-pool attribute
- Use DHCP-PD to configure the CPE

```
interface Atm 0
pvc 1/23 encapsulation aal5mux ppp dialer
dialer pool-member 1
!
interface dialer1
ipv6 dhcp client pd Foo
!
interface FastEthernet0
ipv6 address Foo 0:0:0::1/64
!
```

Auth-Type = Local, Password = “foo2”
User-Service-Type = Framed-User,
Framed-Protocol = PPP,
cisco-avpair = “ipv6:route=3ffe:c00:1::/48

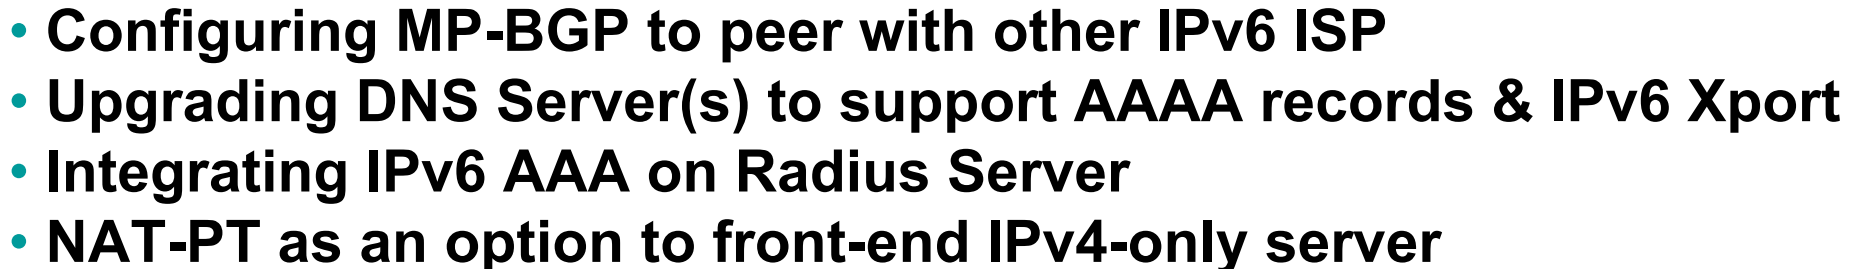


Deployment Activities

Cisco.com

- **IPv6 Core Network Enhancement**
 - IPv6 Provider Edge Routers (6PE) over MPLS
- **Broadband Access Networks**
 - IPv6 over broadband data link layers
 - IPv6 address allocation guidelines
 - IPv6 AAA Radius
 - IPv6 auto-configuration – Prefix Delegation & Stateless DHCP
- **Data Center**
 - IPv6 on Layer 3 infrastructure

Cisco.com



A simple BGP session


Cisco.com



```
Roucer1#  
interface Ethernet0  
  ipv6 enable  
  ipv6 address 3FFE:B00:C18:2:1::F/64  
!  
router bgp 65001  
  bgp router-id 10.10.10.1  
  no bgp default ipv4-unicast  
  neighbor 3FFE:B00:C18:2:1::1 remote-as 65002  
  address-family ipv6  
    neighbor 3FFE:B00:C18:2:1::1 activate  
    neighbor 3FFE:B00:C18:2:1::1 prefix-list bgp65002in in  
    neighbor 3FFE:B00:C18:2:1::1 prefix-list bgp65002out out  
  exit-address-family  
!  
ipv6 prefix-list bgp65002in seq 5 permit 3FFE::/16 le 24  
ipv6 prefix-list bgp65002out seq 5 permit 3FFE::/16 le 24
```

NAT-PT Concept

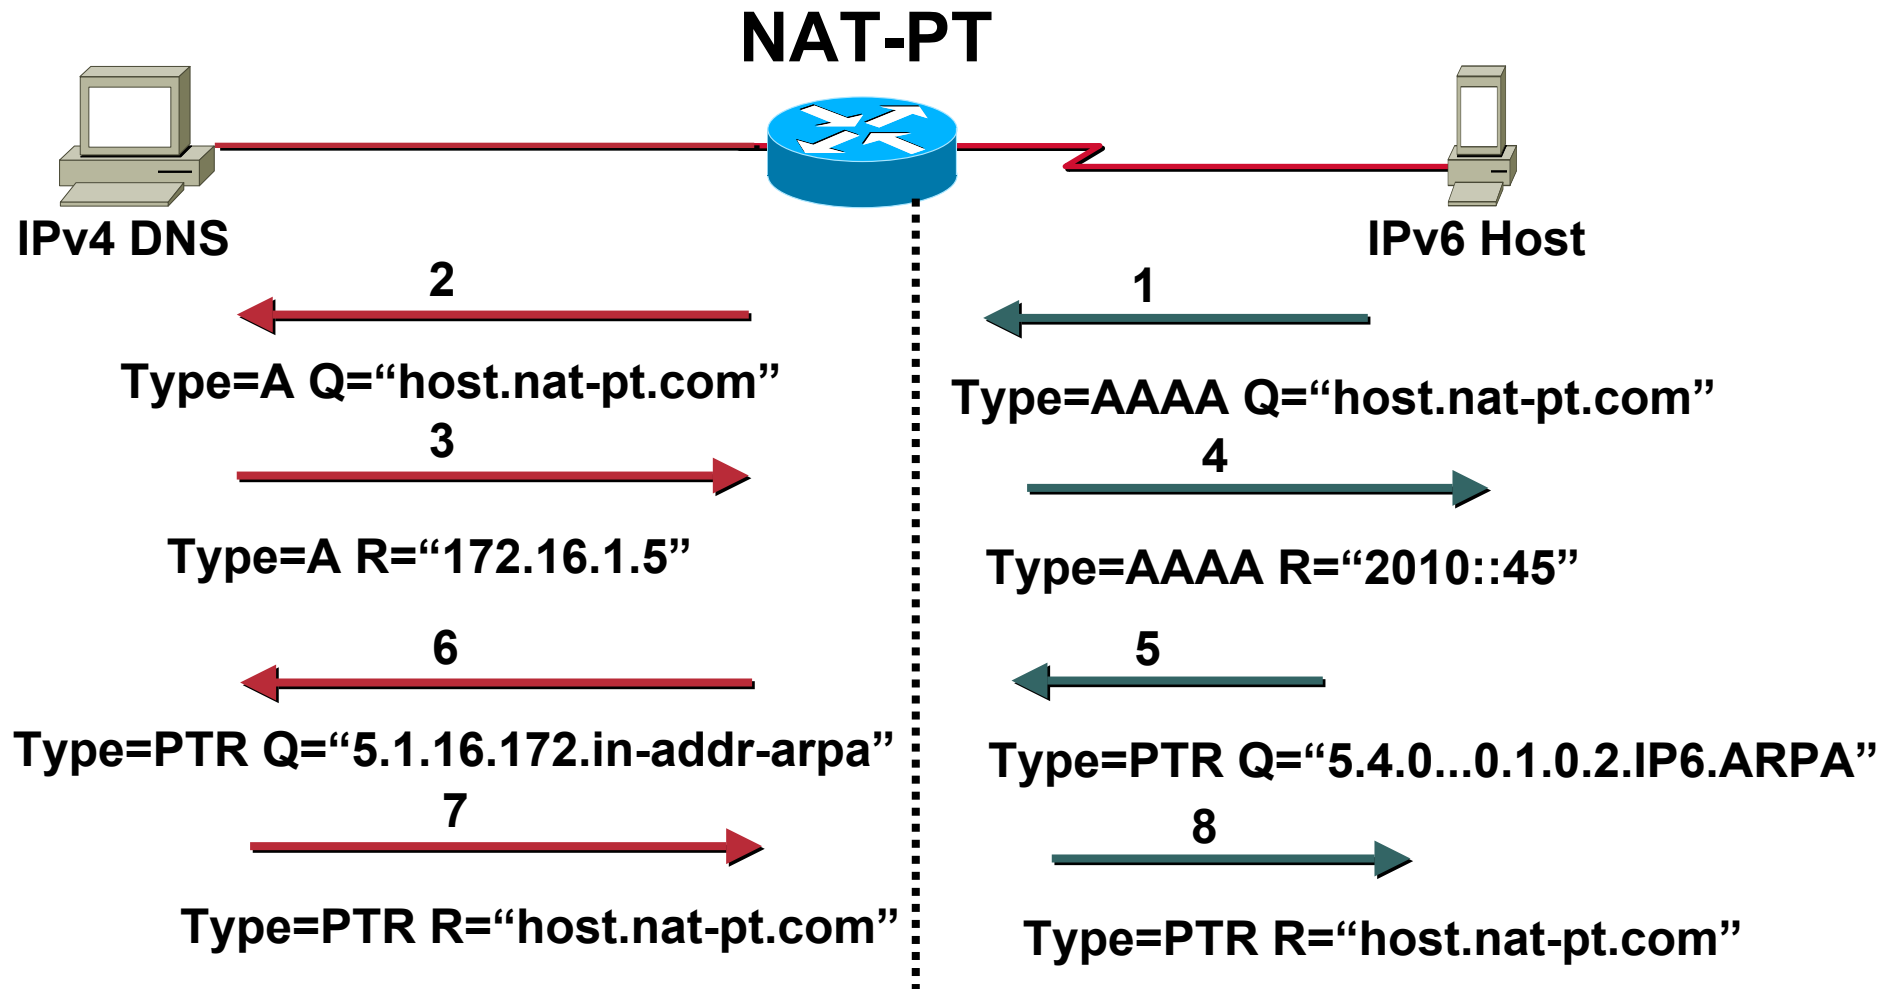
Cisco.com



- **PREFIX** is a 96-bit field that allows routing back to the NAT-PT device

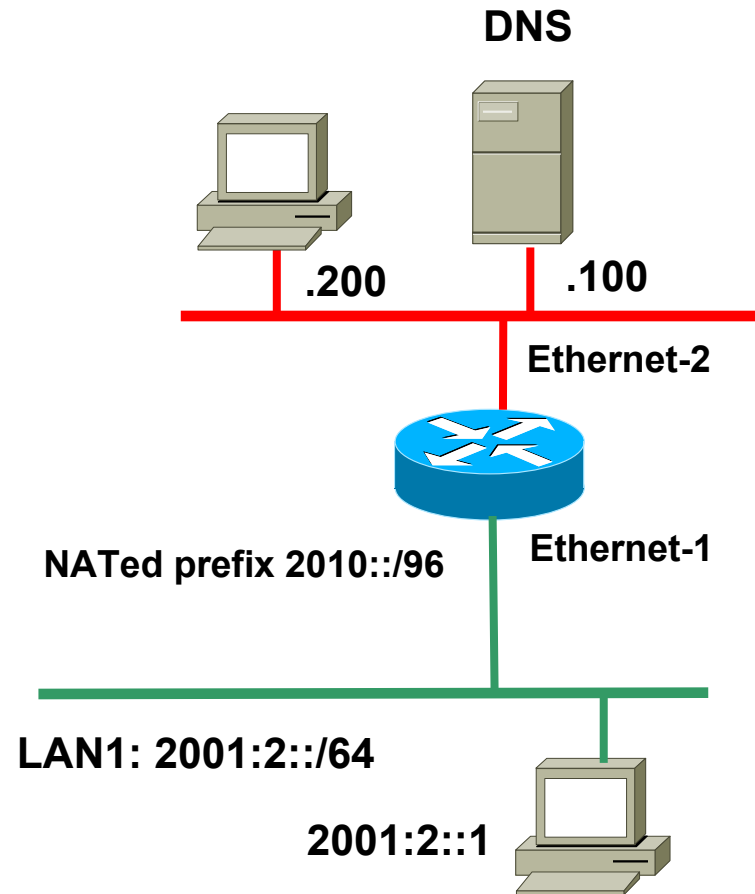
DNS Application Layer Gateway

Cisco.com



Cisco IOS NAT-PT w/ DNS ALG Configuration

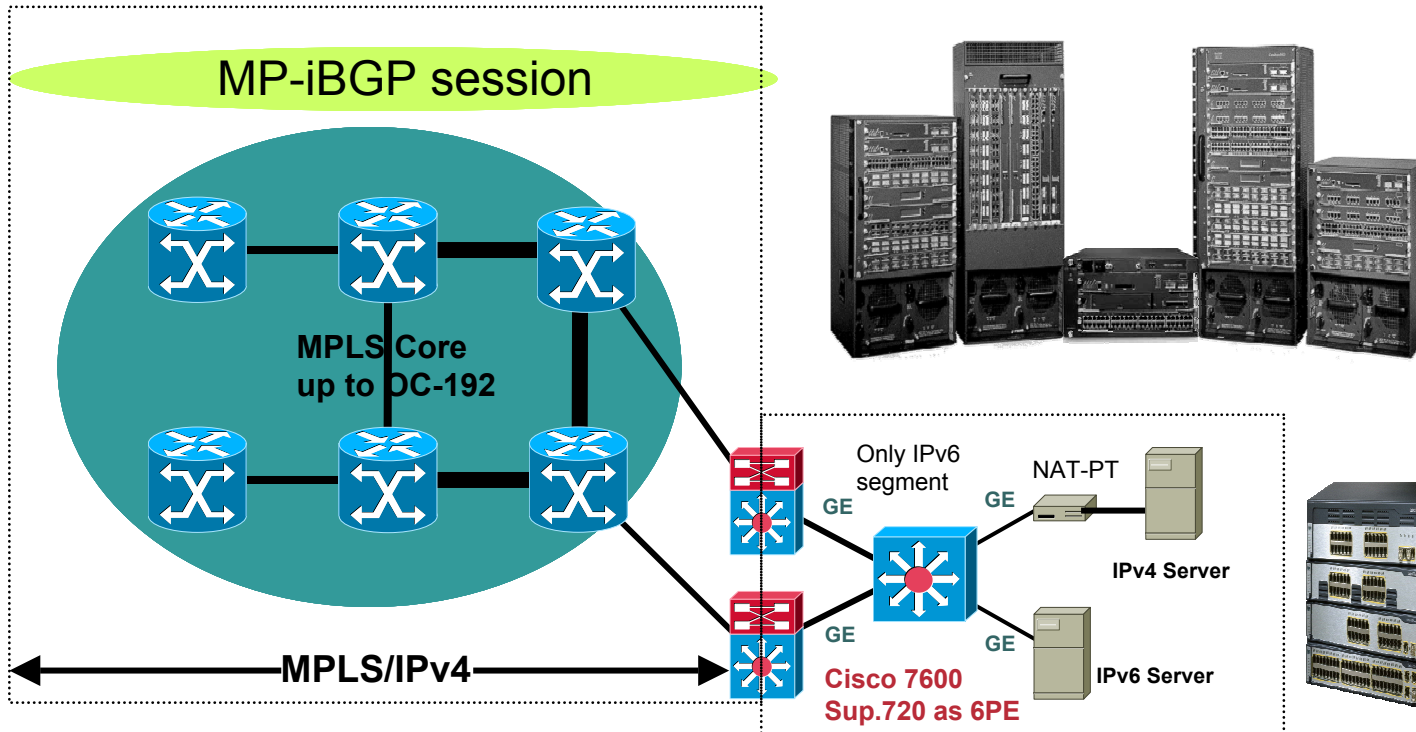
Cisco.com



```
interface ethernet-1
  ipv6 address 2001:2::10/64
  ipv6 nat
!
interface ethernet-2
  ip address 192.168.1.1
  255.255.255.0
  ipv6 nat
  ipv6 nat prefix 2010::/96
!
ipv6 nat v4v6 source 192.168.1.100
2010::1
!
ipv6 nat v6v4 source route-map map1
pool v4pool1
ipv6 nat v6v4 pool v4pool1
192.168.2.1 192.168.2.10 prefix-
length 24
!
route map map1 permit 10
  match interface Ethernet-2
```

IPv6 Deployment in Data Center

Cisco.com






- **Best-in-class IPv6 performances on Catalyst series**

- Supervisor Engine 720 & distributed PFC3 modules, 10GE HW FW

- IPv6 hardware assistance for IPv6 native and IPv6 over IPv4 tunnels (configured. 6to4. ISATAP)

Campus Deployment Scenario

Cisco.com



- Full Layer 3 infrastructure upgrade to implement Dual Stack
 - Software or Hardware/Software dependent of expected performances
- Native IPv6 Router(s) on a Stick
 - VLANs are terminated on one or more IPv6 interfaces/routers
- Native IPv6 Routers on dedicated LANs interconnected using configured tunnels.
- ISATAP tunnels

ISATAP - Intra-Site Automatic Tunnel Addressing Protocol

Cisco.com

- **Tunnelling of IPv6 in IPv4**
- **In a single administrative domain**
- **Creates a virtual IPv6 link over the full IPv4 network**
- **Automatic tunnelling is done by a specially formatted ISATAP address which includes**
 - **An ISATAP special identifier**
 - **The IPv4 address of the node**
- **ISATAP nodes are dual-stack**

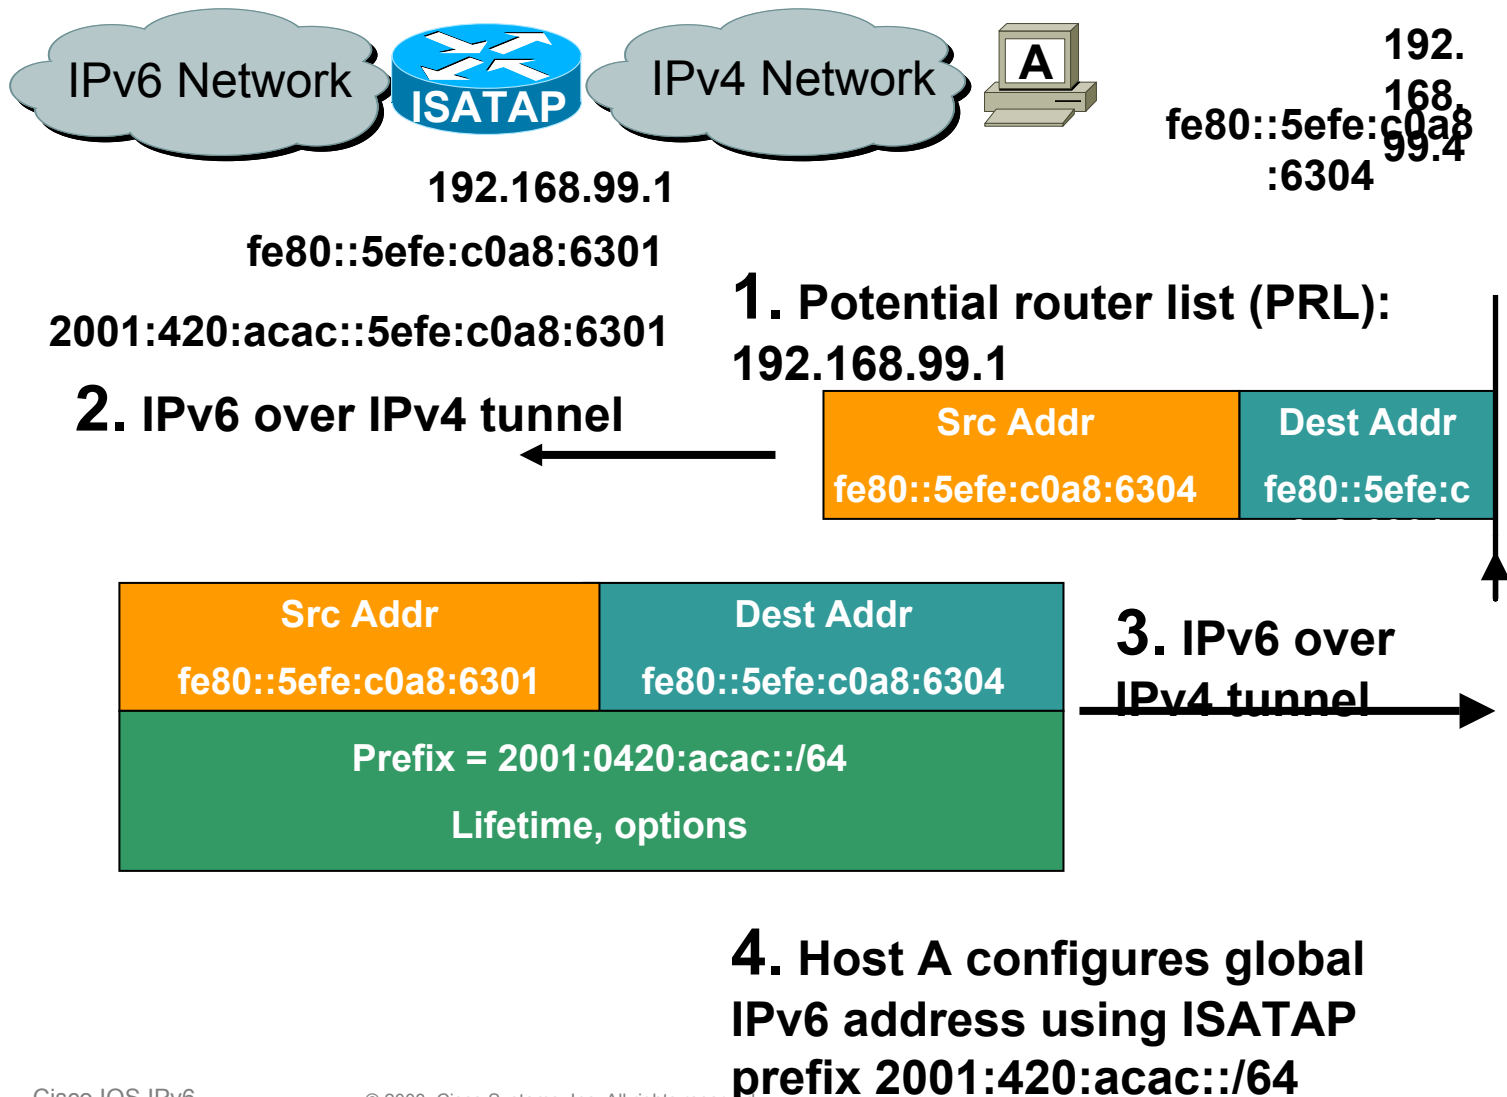
ISATAP address format

- An ISATAP address of a node is defined as:
- A /64 prefix dedicated to the ISATAP overlay link
- Interface identifier:
 - Leftmost 32 bits = 0000:5EFE:
 - Identify an ISATAP address
 - Rightmost 32 bits = <ipv4 address>
 - The IPv4 address of the node


ISATAP dedicated prefix	0000:5EF E	IPv4 address
----------------------------	---------------	--------------

ISATAP prefix advertisement

Cisco.com



ISATAP Router



Cisco IOS 12.2(15)T, Net
12.2(14)S

Supported in Windows XP Pro SP1

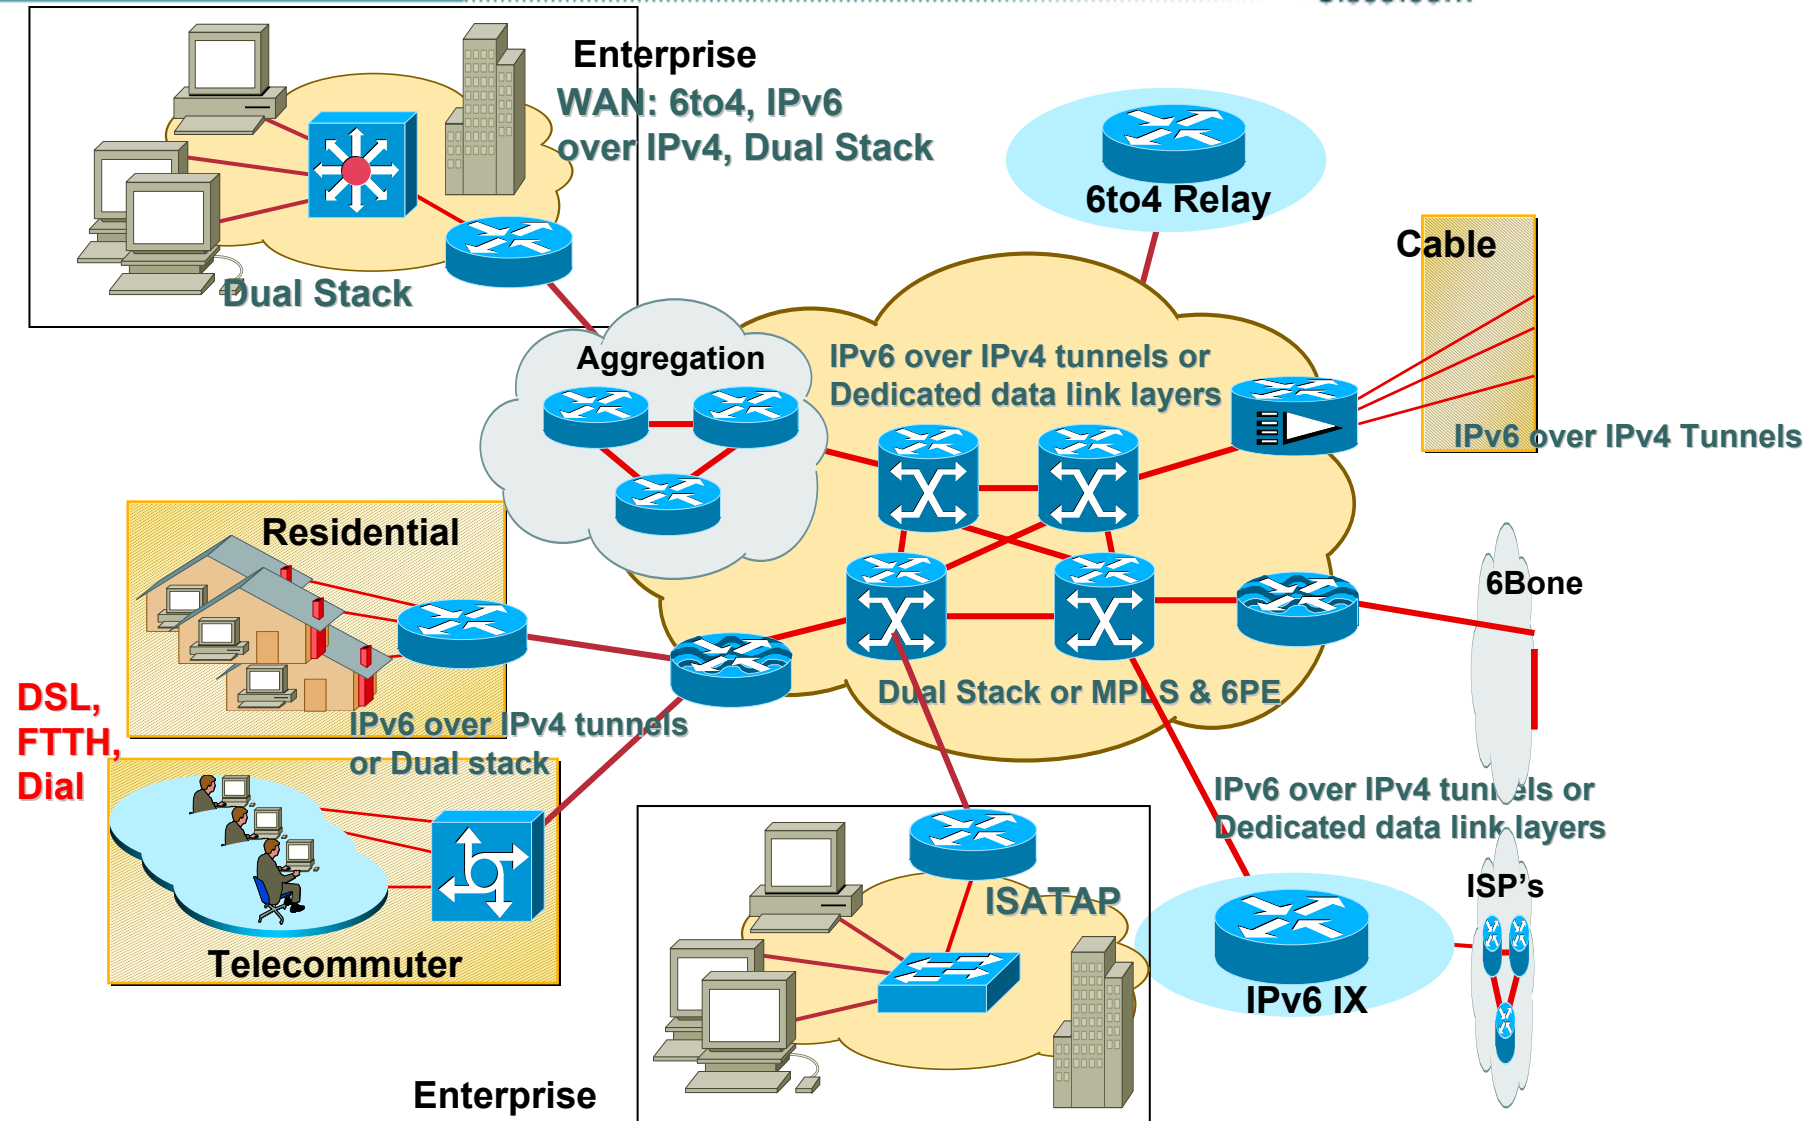
The tunnel source command must point to an interface with an IPv4 address configured

Configure the ISATAP IPv6 address, and prefixes to be advertised just as you would with a native IPv6 interface

```
ISATAP-gw#
!
interface Ethernet0
 ip address 192.168.99.1 255.255.255.0
!
interface Tunnel0
 ipv6 address 2001:420:ACAC::/64 eui-64
 no ipv6 nd suppress-ra
 tunnel source Ethernet0
 tunnel mode ipv6ip isatap
```

Moving IPv6 to Production, running Cisco IOS

Cisco.com



Conclusion

Cisco.com

Moving IPv6 to Production?


- **Core IPv6 specifications are well-tested & stable**
Some of the advanced features of IPv6 still need specification, implementation, and deployment work
- **Application, middleware and Scalable Deployment scenario are IPv6 Focus and Challenge.**
- **Plan for IPv6 integration and IPv4-IPv6 co-existence**
Training, applications inventory, and IPv6 deployment planning
- **Cisco is committed to deliver advanced IPv6 capabilities to the Internet industry**
IPv6 Solutions, ABC of IPv6, e-Learning/Training, ISD,...
See <http://www.cisco.com/ipv6>

Questions?

Cisco.com



CISCO SYSTEMS



EMPOWERING THE
INTERNET GENERATION