Backbone transition case study

Jerome.Durand@renater.fr

Academics' story with IPv6

- G6bone
  - The first IPv6 network in France (1996)
  - One of the 3 first IPv6 nodes starting the 6bone
    - UNI-C, DK
    - WIDE, JP
    - G6, FR
  - Tunneled network (v6inv4)
  - Hierarchical addressing from the beginning
    - Two-level topology : Regional Interconnects (RIs) + IPv6 sites
  - Static routing + RIPng …

G6bone

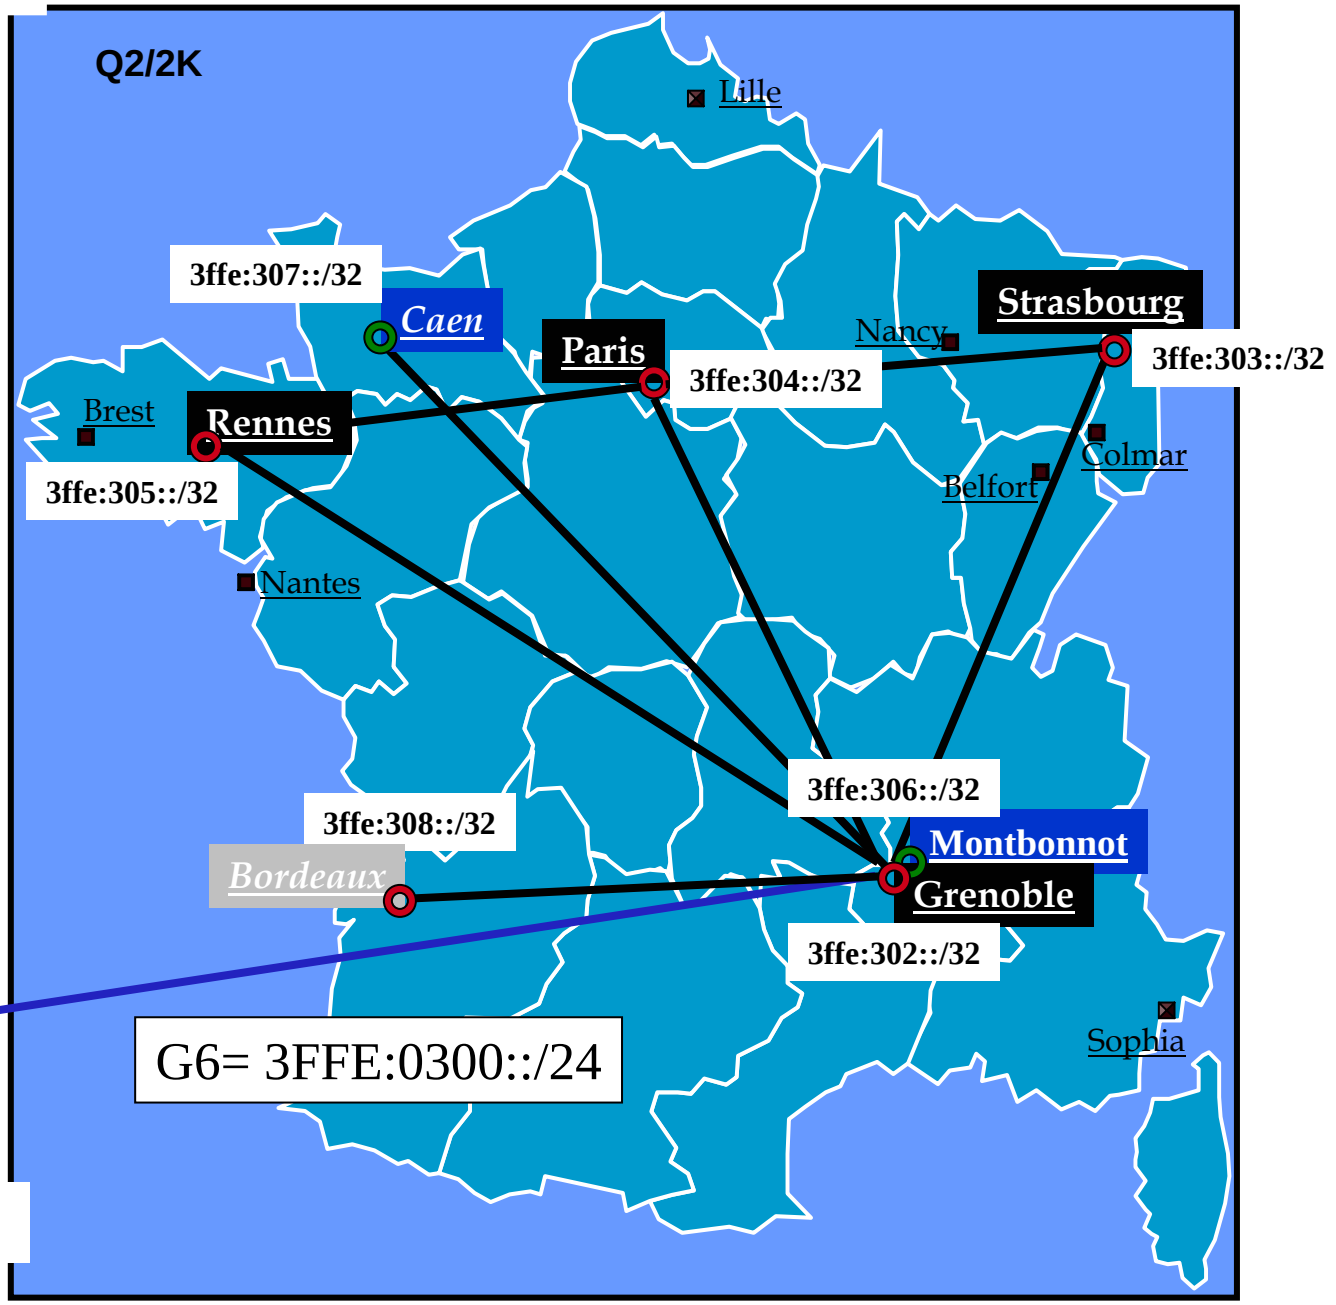
Q2/2K
Lille
3ffe:307::/32
Caen
Paris
Nancy
Strasbourg
3ffe:303::/32
3ffe:304::/32
Brest
Rennes
Colmar
Belfort
3ffe:305::/32
Nantes
3ffe:306::/32
3ffe:308::/32
Montbonnot
Bordeaux
Grenoble
3ffe:302::/32
Sophia
6Bone
G6= 3FFE:0300::/24

Sites connected to the G6 PoPs

Paris:
- Evry: Université INT
- Noisy-le-Gd: ESIEE
- Roquencourt: INRIA
- Saclay: CEA
- Lille: EUDIL
- Paris:
  - Aerospatiale/Matra
  - Brainstorm
  - CIE
  - CISI/ATRE
  - CNAM
  - ENST
  - Eurocontrol
  - Informatique P7
  - Institut Pasteur
  - Internatif
  - ISDnet
  - LAAS
  - LIP6 + Marocco
  - Logique P7
  - OpenTransit
  - Renater2 NOC
  - Urec/Cnrs
  - UVSQ
  - AFNIC

Grenoble :
- Echirolles: Bull
- Marseille : Ec. Sup. Mécanique
- Valbonne: Compaq
- Vanoise
- Grenoble :
  - Allied Signal
  - COSY
  - IMAG
  - MCS
  - Thomson-CSF/Detexis

6Bone :

ATT	JAnet	JOIN
INFN	Switch	Uni-C

Montbonnot:
INRIA

Strasbourg:
- Belfort: Univ Technol.
- Colmar: IUT
- Nancy: Loria
- Strasbourg:
  - Univ. L. Pasteur
  - IUT

Rennes:
- Nantes
- Brest
- Tunisia

Bordeaux:
- Univ. Reaumur

Caen:
CNET

Academics' story with IPv6

Then came Renater …

- IPv6 Pilot over Renater-2 (P6R2)
  - May 2000
  - A native IPv6 network
    - dedicated ATM VPN
  - Deploy the production addressing plan
    - July 1999 : first sTLA allocation
  - Same two-level topology as in G6bone
    - Academic sites
      - production addressing scheme
    - Industrial sites involved in research projects
      - 6bone addressing scheme
- Gain experience with a pre-production service

Renater's IPv6 Pilot topology

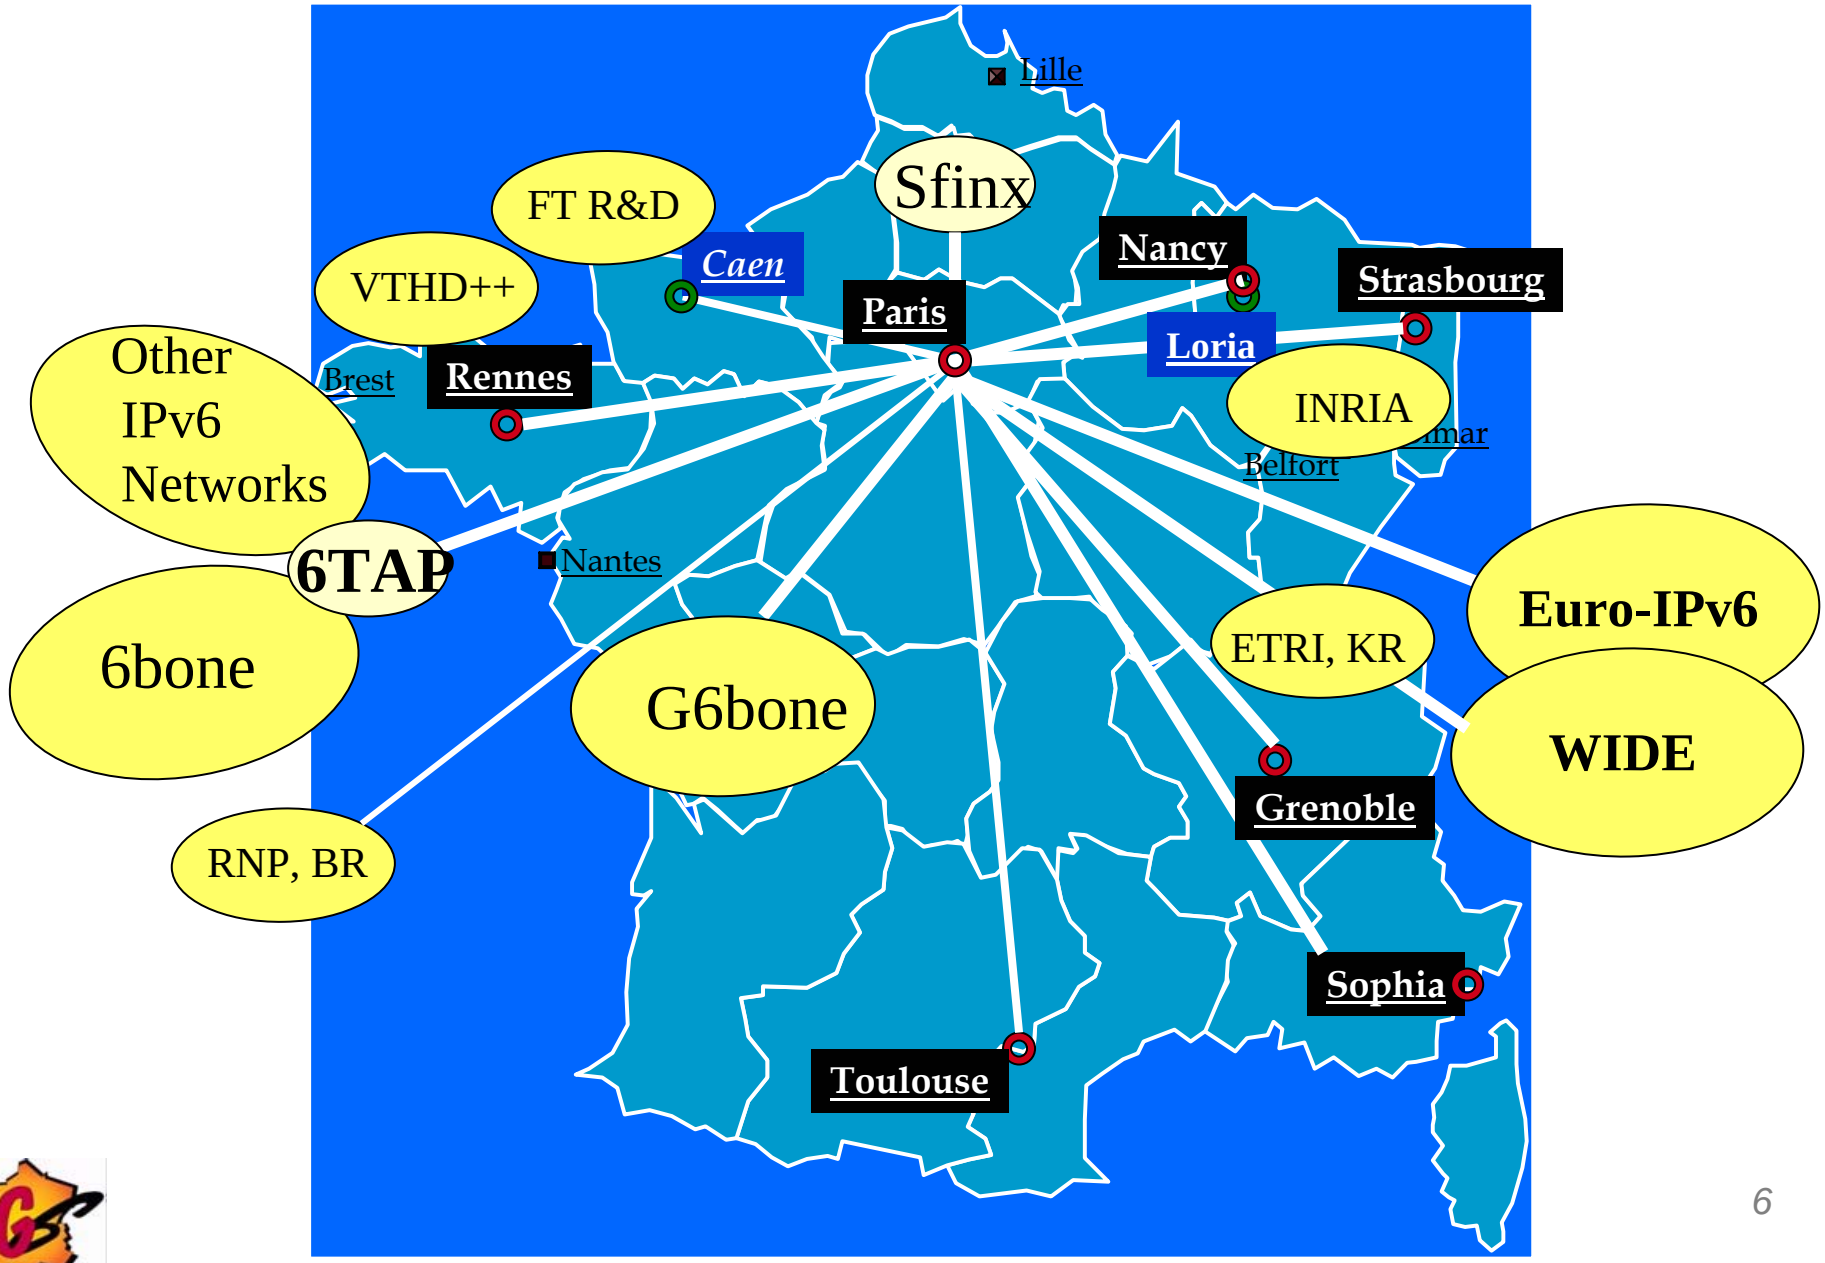

The Pilot experience

- Experience Using the protocol
  - Equipment
    - Cisco partnership
  - Addresses
    - Deploying a consistent scheme (/35) for the core and the sites
  - Routing
    - ISIS and BGP4+
- IPv6 resources allocation
  - Procedures and management
- IPv6 DNS
  - Deployment of the DNS service
  - Reverse zones delegation to RIs and end-users sites
- Management
  - IPv6 NOC within Renater-2 NOC
  - Management and monitoring tools
    - Set of looking glasses at the RIs

Academics' story with IPv6

- Summary
  - Understand the technology
  - Deploy the network
  - Manage the whole thing
    - Technical resources
    - Human resources
    - Financial resources

Towards a native IPv6 network

- G6bone was an overlay tunneled network
  - v6 traffic encapsulated in v4 packets
- « independent » from Renater's underlying infrastructure
- P6R2, IPv6 pilot was a VPN of ATM PVCs
- Goals
  - Have a production IPv6 network
    - In the core
    - Allow Regional and Metro Nets to deploy IPv6

Additional goals

As production addresses became available

And sTLA expanded from /35 to /32

- Renumber the IPv6 pilot using a new addressing scheme
  - much simpler to be aligned on nibble boundaries !
- Keep a two-level hierarchy
  - A core backbone of Regional Interconnects (RI)
  - User sites connect to one or more RIs

Additional goals (2)

- Transition period
  - Offer IPv6 connectivity via the new/native infrastructure
  - Keep the old infrastructure in place
  - Move step by step : no D day
- Gather non academic organizations in the G6bone addressing plan (3FFE:0300::/24)
  - Allow them to gain experience with IPv6 until commercial ISPs are ready
  - Have full IPv6 connectivity to the evolving Internet v6
- Connect the pilot to the Sfinx (Renater's IX)
  - Peer with ISPs and non academic organisms
- Provide IPv6 connectivity to
  - National projects (RNRT/RNTL)
  - European projects (IST, Esprit)
  - ...

Toward a Production IPv6 service

And now Renater-3 …

- Why a production-like IPv6 service ?
- ATM removed …
  - Move all network services on a unique topology
  - Do we want to forget about IPv6, IPv4 multicast … ?
- Need of IPv6 transport
  - Research projects using IPv6
  - Sites with native IPv6 network

  →install a native IPv6 core

  →run both versions of IP the same way
- Manage the IPv6 service with the same operational quality as for IPv4

RENATER-3
IPv6 native support

40 NR
2,5 Gbits/sec
Global IP service

Open Transit

Same level of:
Performance
Availability
Management
Support

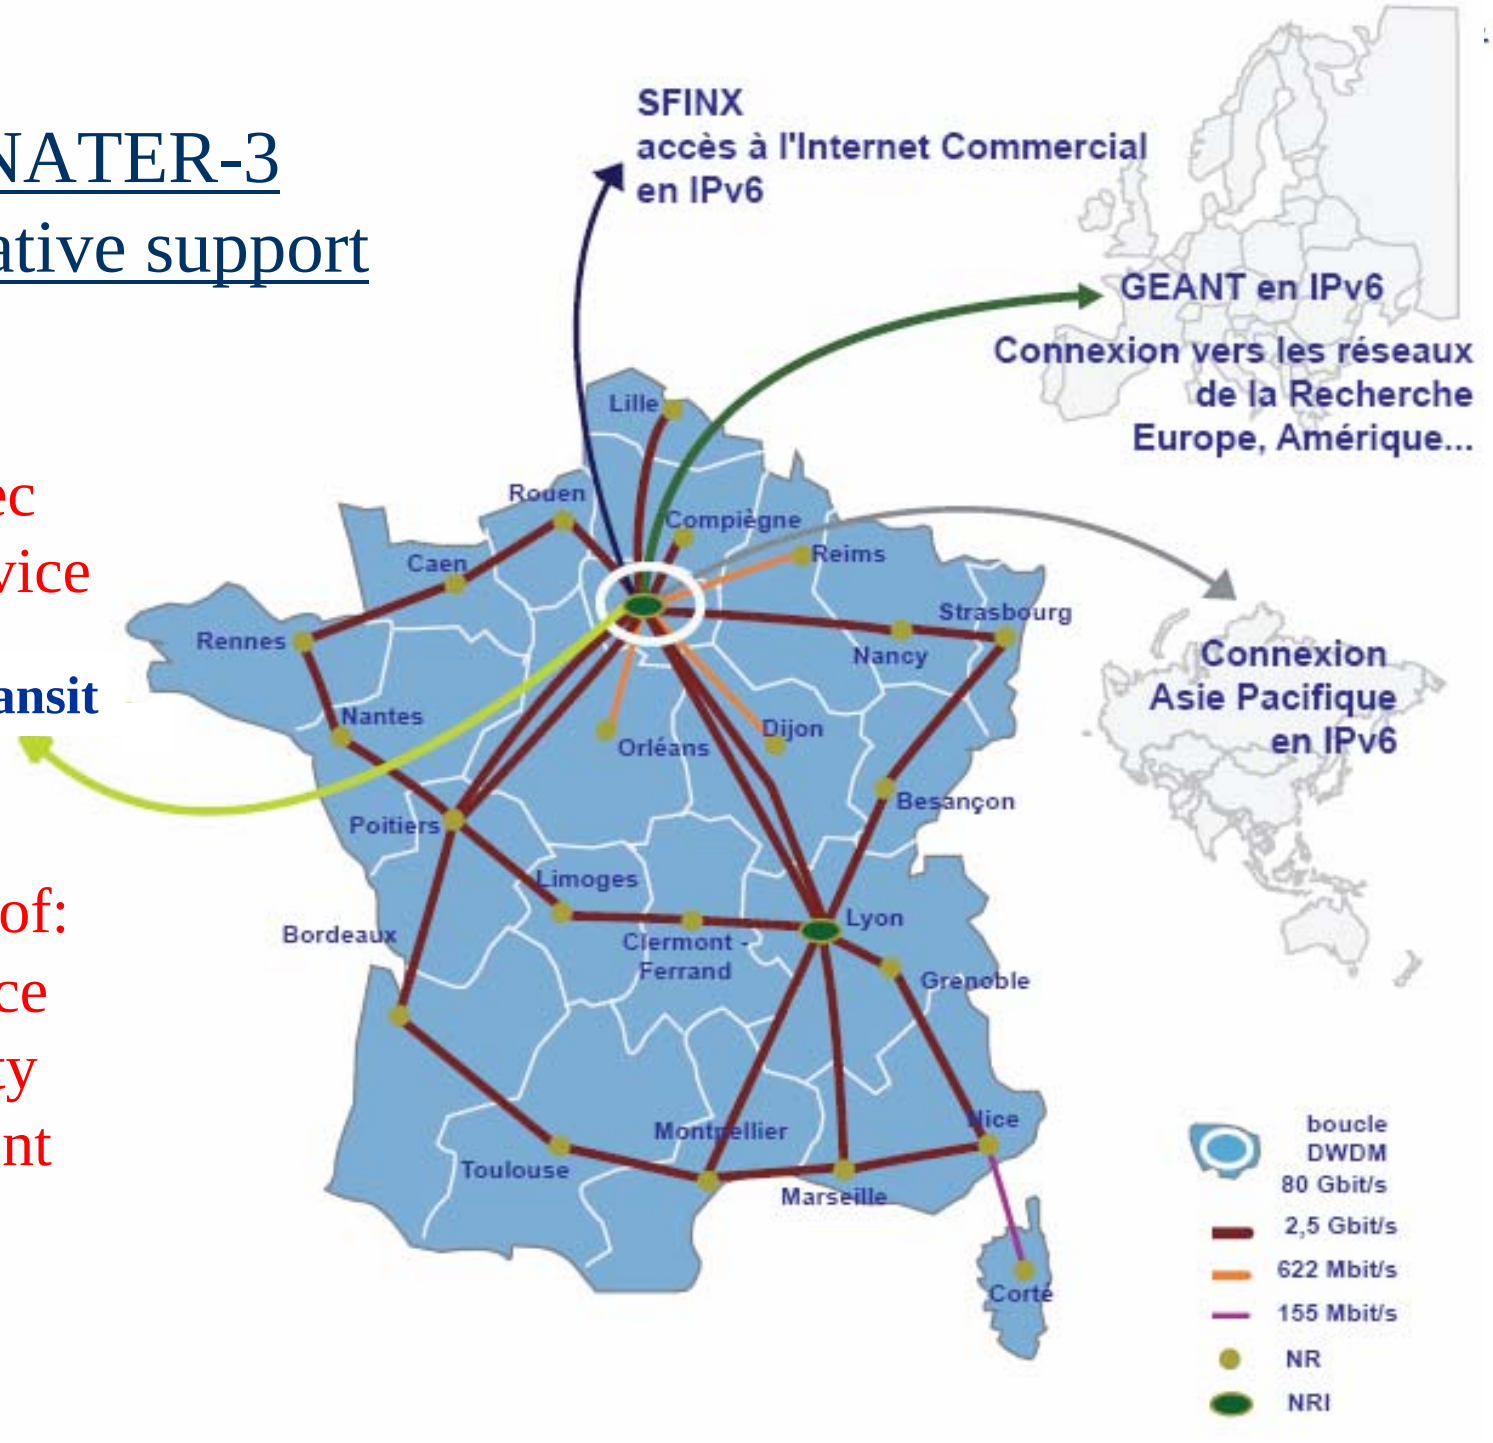

Renater 3: Native support

- 2.5 Gbits/s backbone
- 30 Regional Interconnects (RI)
- Native IPv6 support on all RIs
  - Dual stack backbone → IPv4 and IPv6
- Global IP Service
  - IPv4 unicast and multicast
  - IPv6 unicast
  - IPv6 and IPv4 carried without any distinction
- Experimental IPv6 multicast network
- Goal : achieve an equal level of
  - Performance
  - Availability
  - Management
  - Support

Renater IPv6 addressing scheme

IPv6 service in Renater-3

- Based on experience gained with the IPv6 Pilot deployment
- Principles for 6R3
  - /35 expands to /32 (2001:0660::/32)
  - Two-level hierarchy : core + access
  - Core are /40 allocated (easier to manage)
    - Each PoP identified with a Reg-ID
  - Sites are /48 (as recommended)
    - Identified with a NLA-ID

Addressing scheme

- What do we need to number?
  - Regional Interconnects: /40
    - Reg-IDs allocation
  - Sites (labs, campuses …): /48
    - NLA-IDs allocation
    - 16 bits are reserved for the site topology
  - Interconnection networks
    - RI – sites
    - Renater – other IPv6 networks
      - Operational
      - Projects

Addressing scheme (2)

sTLA = **2001:0660::/32**

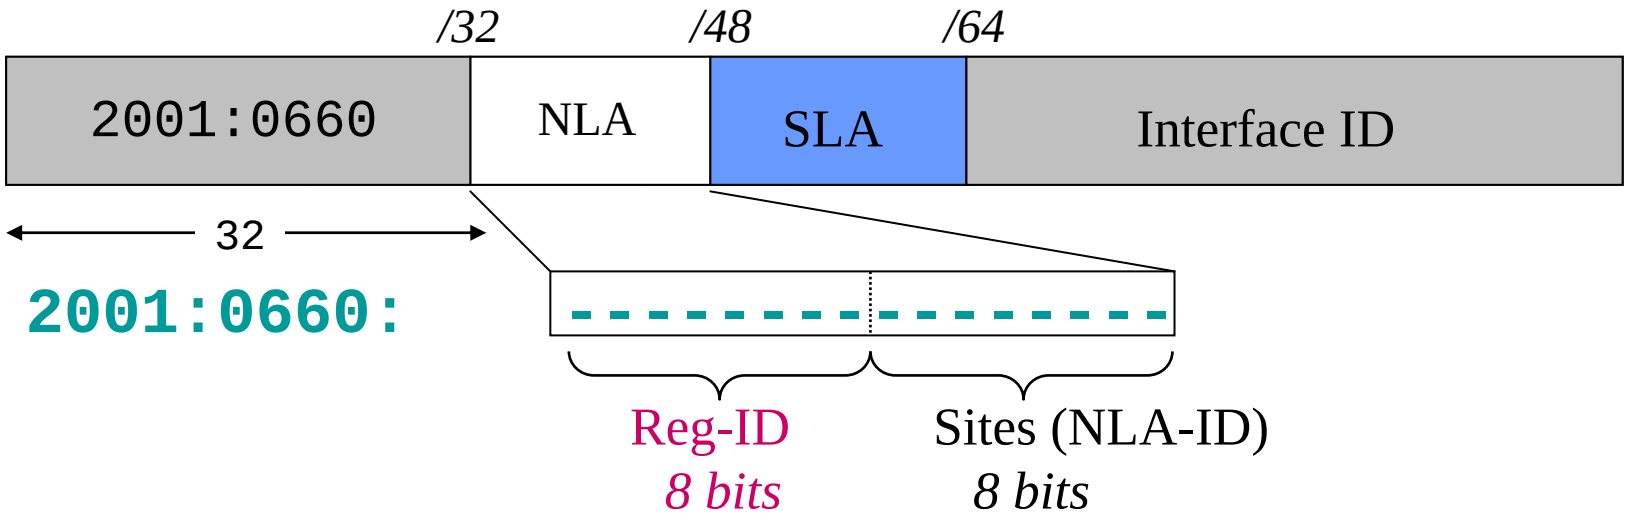


```
2001:0660:3000:/40   Paris NRI
2001:0660:3300:/40   Paris Jussieu
2001:0660:4400:/40   Lille
2001:0660:5400:/40   Marseille (…)
```

```
2001:0660::/48
```

Addressing scheme (3)

- Hierarchical addressing
- Renater: 2001:0660::/32 from RIR
- Regional RIs: /40 (reg-ID)
- Sites: /48 from /40 of RIs
  - NLA-IDs allocation
  - /48s aggregation to a single /40 for all sites connected to the same PoP
    - Now we decided to keep /48's everywhere for monitoring purpose.
  - 16 bits are reserved for the site topology ("subnets")

Example

- Renater's sTLA: 2001:0660::/32
- RI Rennes : 2001:0660:7300::/40
- RI's local network : 2001:0660:7300::/48
- Sites connected to the RI
  - 2001:0660:7301::/48
  - 2001:0660:7302::/48
  - (…)

Multihomed domains

- In IPv4, create lots of entries in default free routing tables
- In IPv6, interface will have several IPv6 addresses
  - Problem of source address selection is still under study

Naming

- Direct DNS
  - Same domain name for IPv6 and IPv4
  - Ex : site.fr for IPv4 and IPv6
  - Just add an IPv6 entry for IPv6 addresses
- Reverse DNS
  - 0.6.6.0.1.0.0.2.ip6.int from the beginning
    0.6.6.0.1.0.0.2.ip6.arpa under deployment
  - Reverse zone's delegation of /48 allocated to the sites

Routing & routing policy

- IGP: ISIS + iBGP
- EGP: e-BGP4+
- Route Reflectors
  - At each NRI
- In the backbone
  - /48 of sites aggregated in /40
- International advertisements
  - Announce Renater /32 sTLA
  - Accept /32 (or shorter) or /35 from ISPs
- Prefixes not allowed are filtered out
- Client sites connections
  - Their own choice: static, BGP4+
  - Not allowed to advertise more specific prefixes than /48s

Transition

- Renater's Backbone is native IPv6
- Some sites too
- BUT most of regional networks are not IPv6 capable yet …

=> Install an equipment in each RN to connect IPv6

- Between regional router and sites:
  - VLANS
  - Tunnels
  - ATM PVC

Scenario 1:
Peering router is IPv6 capable

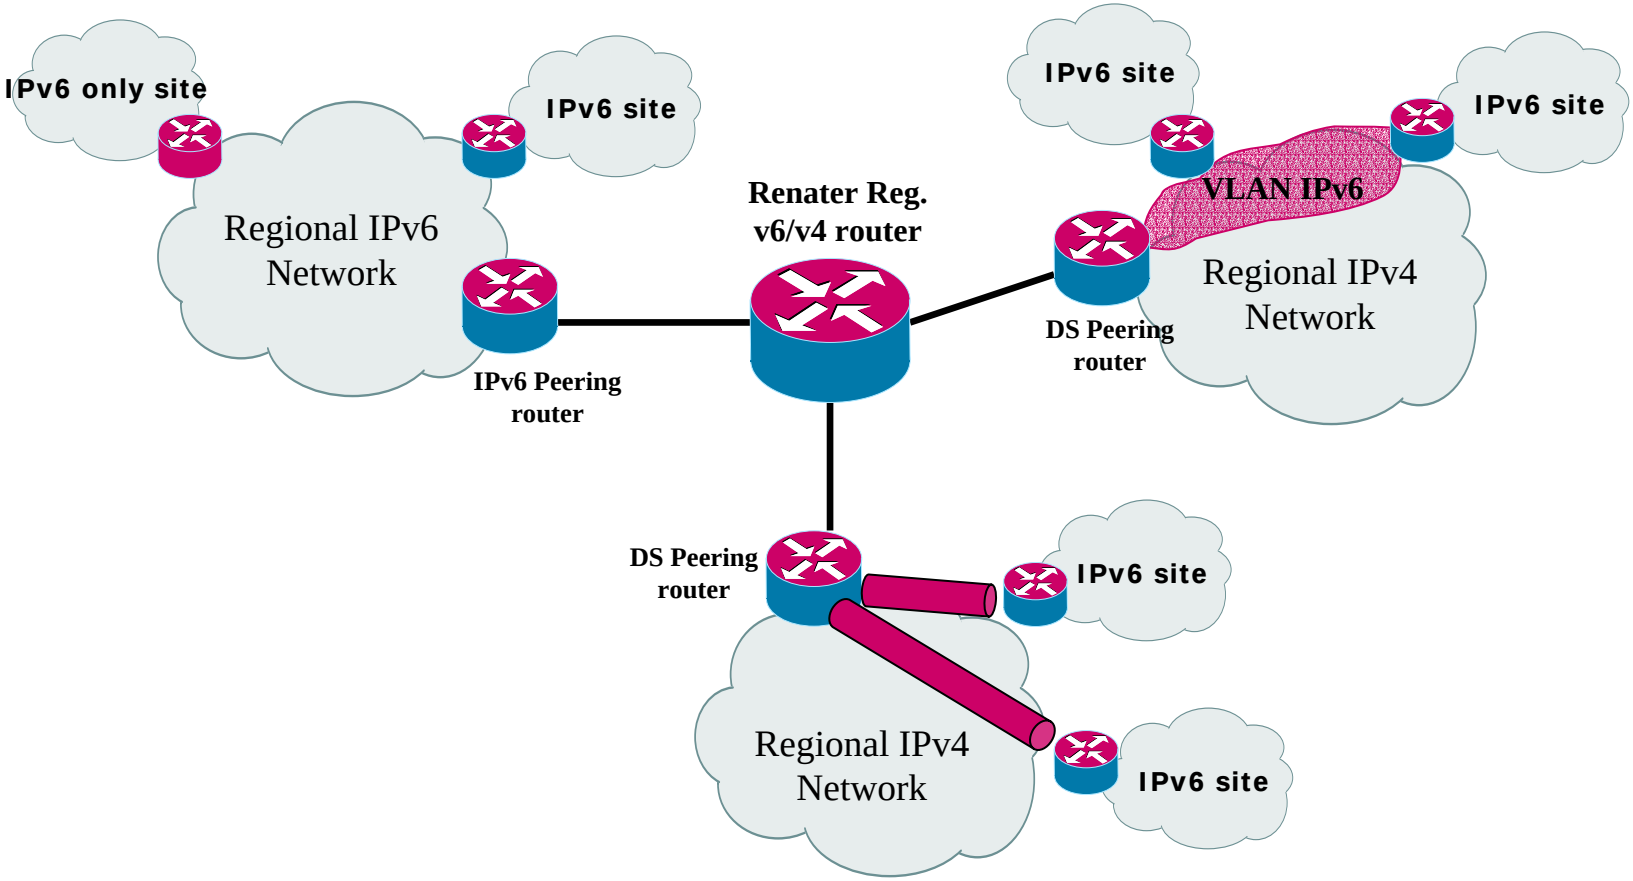

Scenario 2:
Peering router is IPv4 only

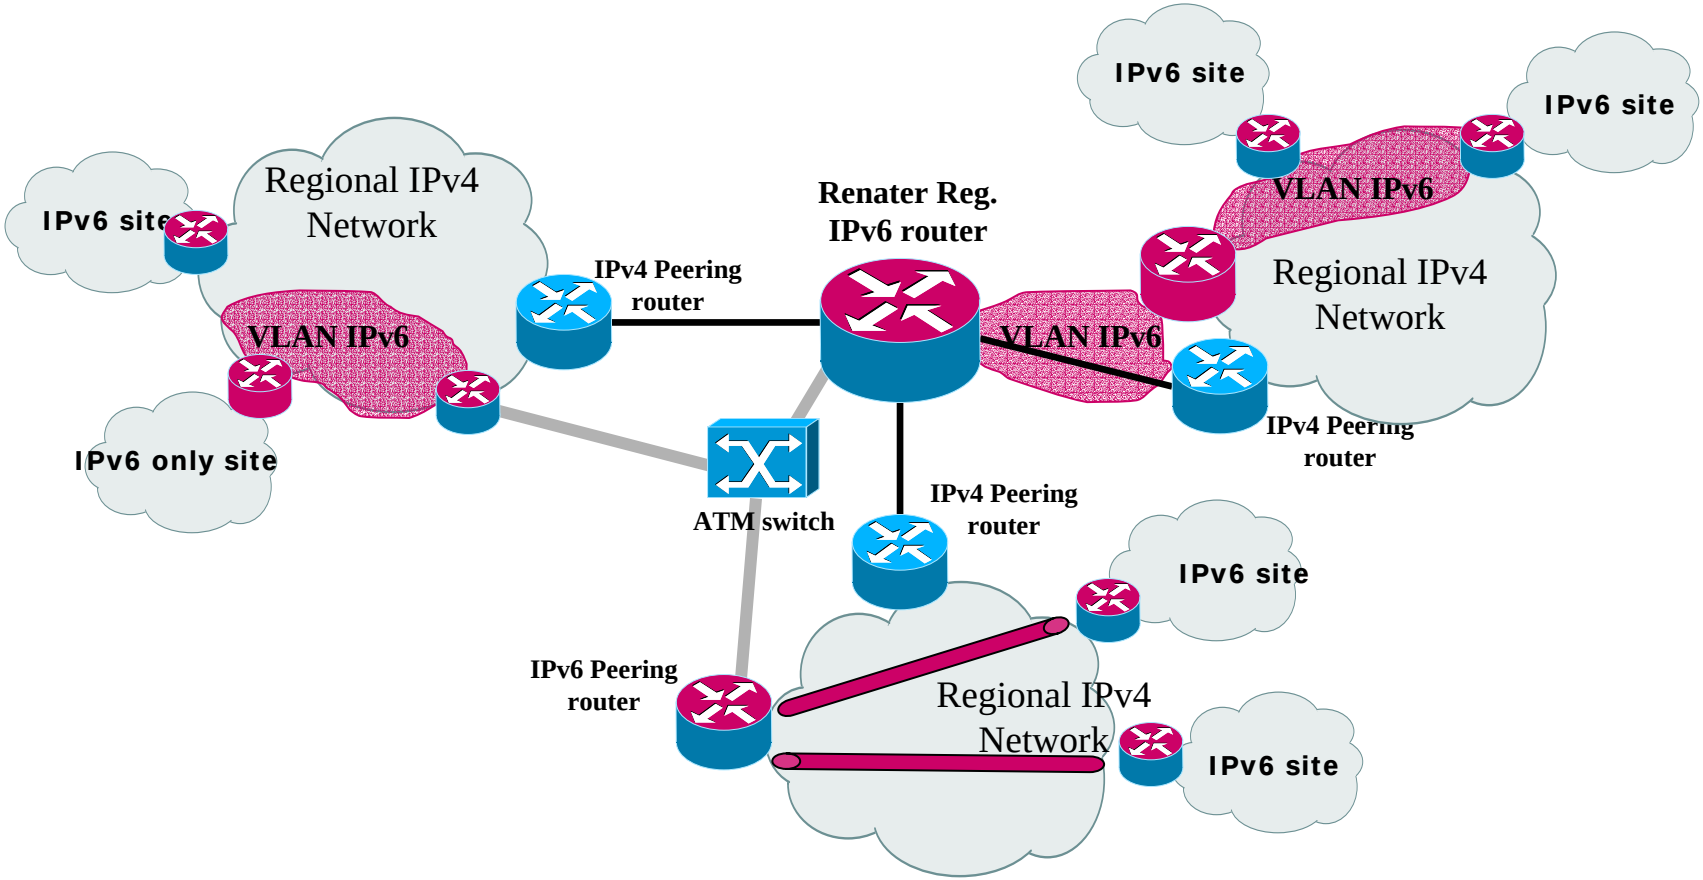

Equipment

- Core routers are Cisco C124xx
  - POS + GEth interfaces …
- Edge routers are
  - Mainly Cisco's (C7xxx, C36xx, C65xx, …)
  - But also Juniper's M5, M10 …
  - 6WIND 6200
  - …

Latest update: Tunnel broker configuration

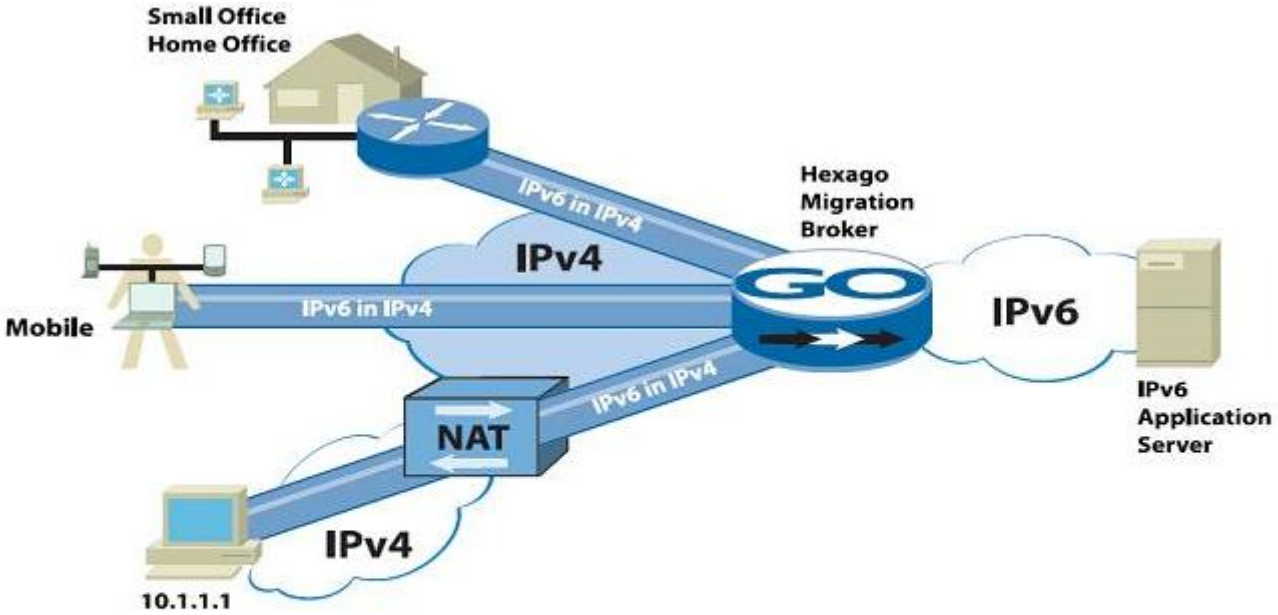

Hexago Migration Broker™ provides:

- IPv6 connectivity over IPv4 networks

Why ? We get native !

- Lack of IPv6 enabled regional networks
  - 16 networks IPv6 enabled today
- Connect non-RENATER sites:
  - Projects
  - Industry R&D
- Events, show rooms
- Connect directly end-stations so that they can get anytime/anywhere IPv6 connectivity

How it works

- TSP (Tunnel Setup Protocol) automatically creates the tunnels
  - Prefixes automatically allocated
    - Sites
    - SOHO…
- Need to allocate a prefix for the tunnel broker, that will then allocate prefixes to sites/clients:
  - Ex:
    - Tunnel broker = /40
    - Allocated prefixes = /48 ou /64

IPv6

Because I’m worth it