



deploy

IPv6 Addressing case studies



Contributions

Main authors

- Bernard Tuy, RENATER - France
- János Mohácsi, NIIF/HUNGARNET - Hungary

Contributors



Updates

B. Tuy 19/05/2010

J. Mohacsi 12/03/2009



Outline of Presentation

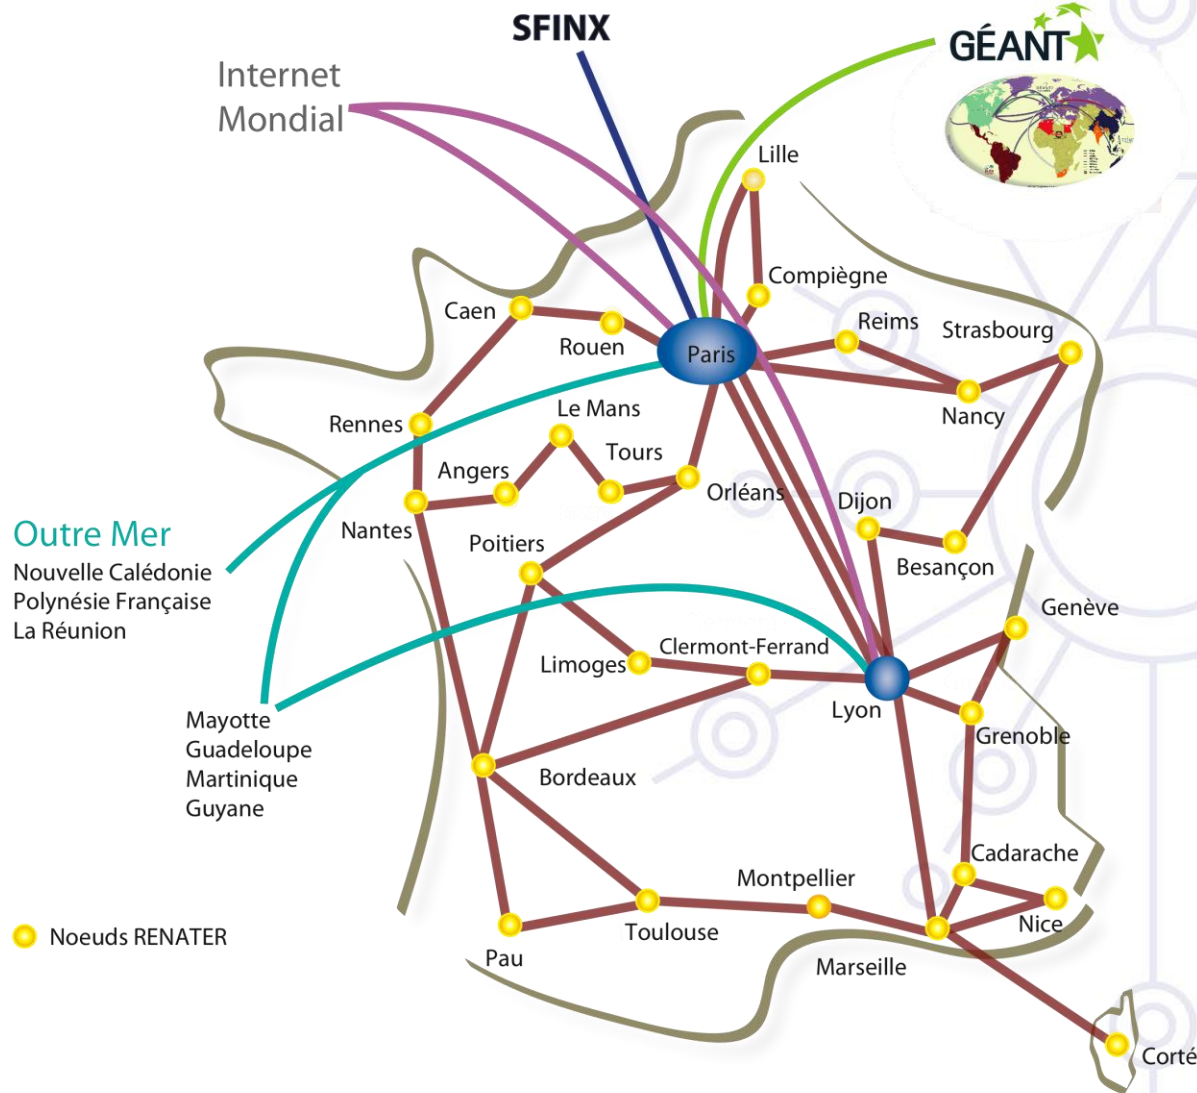
- **Overview of RENATER's network**
- **Case study of IPv6 address allocation at RENATER**
- **Overview of NIIF/Hungarnet's network**
- **Case study of IPv6 address allocation at NIIF/Hungarnet**



deploy

RENATER IPv6 numbering

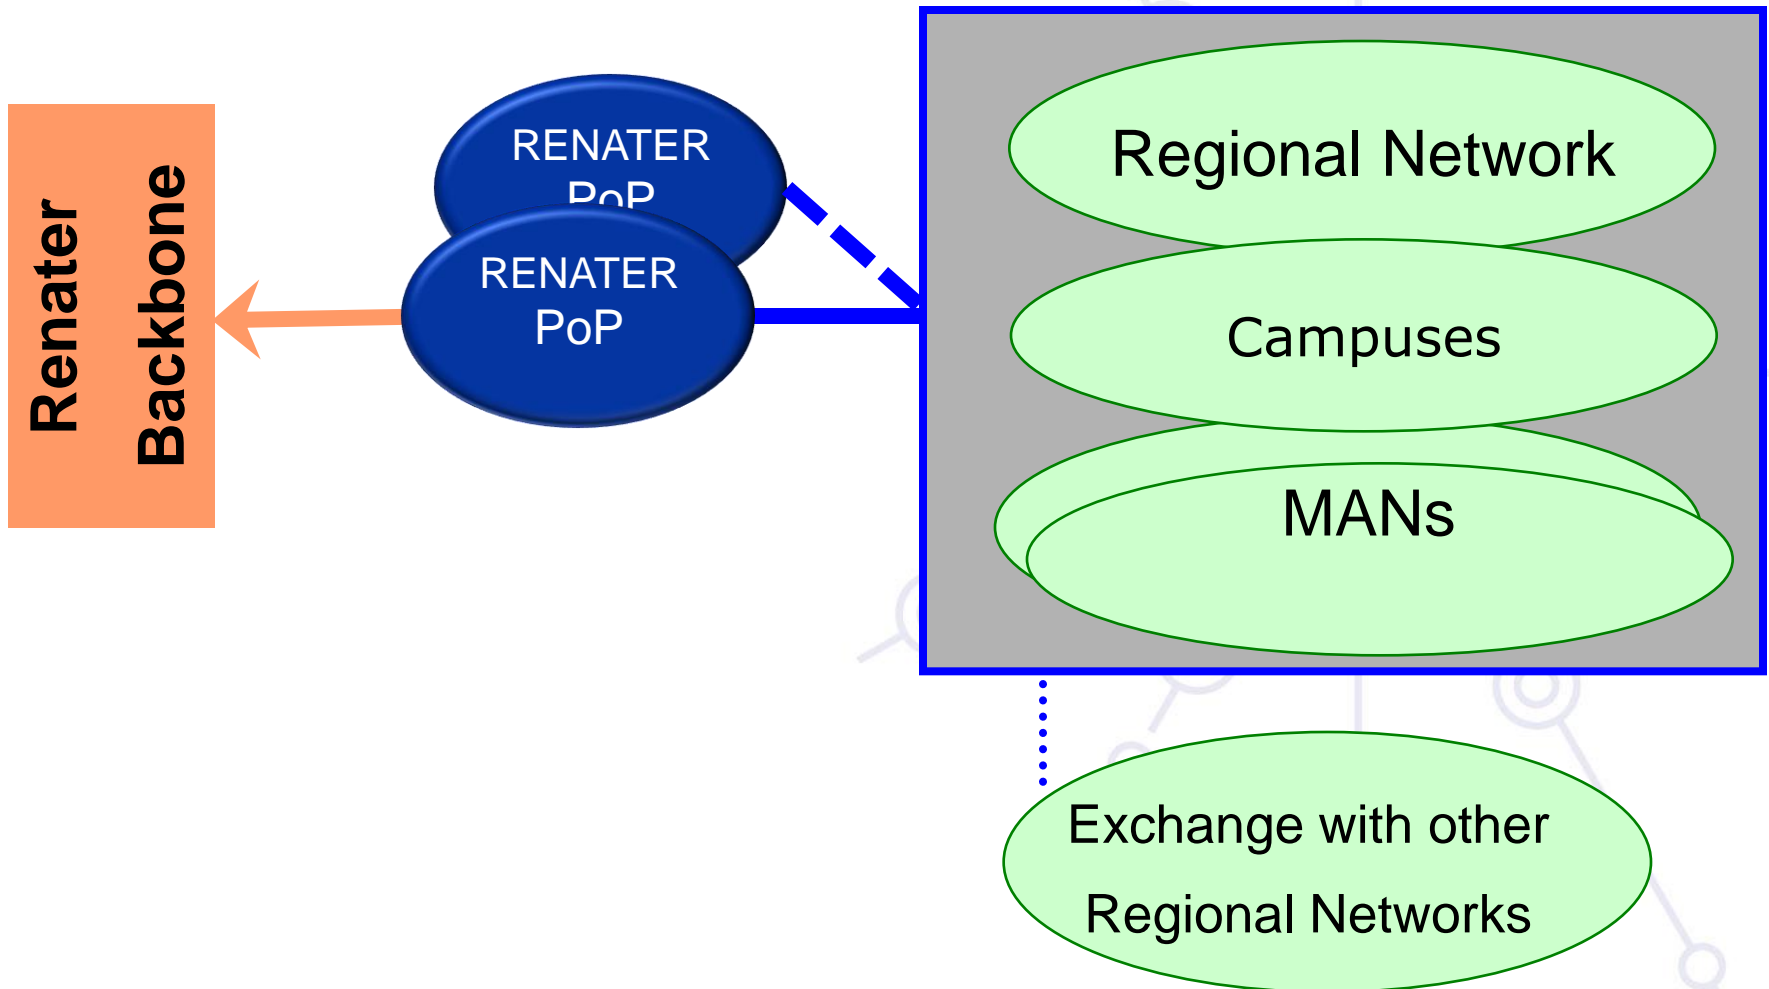
RENATER-5: national & international links



RENATER-5: links speed

- **Backbone links : 10 Gbps (mostly)**
 - But Corsica : 2,5 Gbps
 - Paris-Lyon: 2 x 10 Gbps
- **International links:**
 - IP Transit North: 10 Gbps
 - IP Transit South: 10 Gbps
 - GEANT (EU NRENs Backbone): 2 x 10 Gbps
 - CBFs with
 - DFN (Kehl)
- **SFINX (Internet eXchange): 2 x 10 Gbps**

RENATER architecture



RENATER's Production IPv6 service

Why a production-like IPv6 service ?

- **Needs for an IPv6 transport**

- Research projects using IPv6
- Sites with native IPv6 network
- → install a native IPv6 core
- → run both versions of IP on the same equipments

⇒ **Monitor the IPv6 service in the same operational way than IPv4**

Renater : IPv6 Native support

10 Gbps backbone links

50 Regional Nodes (NR)

Native IPv6 on all regional nodes

- Dual stack backbone → IPv4 and IPv6

Global IP Service

- IPv4 unicast and multicast
- IPv6 unicast
- IPv6 and IPv4 carried without any distinction

Goal : achieve for both versions of IP an equal level of

- Performance
- Availability
- Management
- Support

Addressing

Hierarchical addressing RENATER

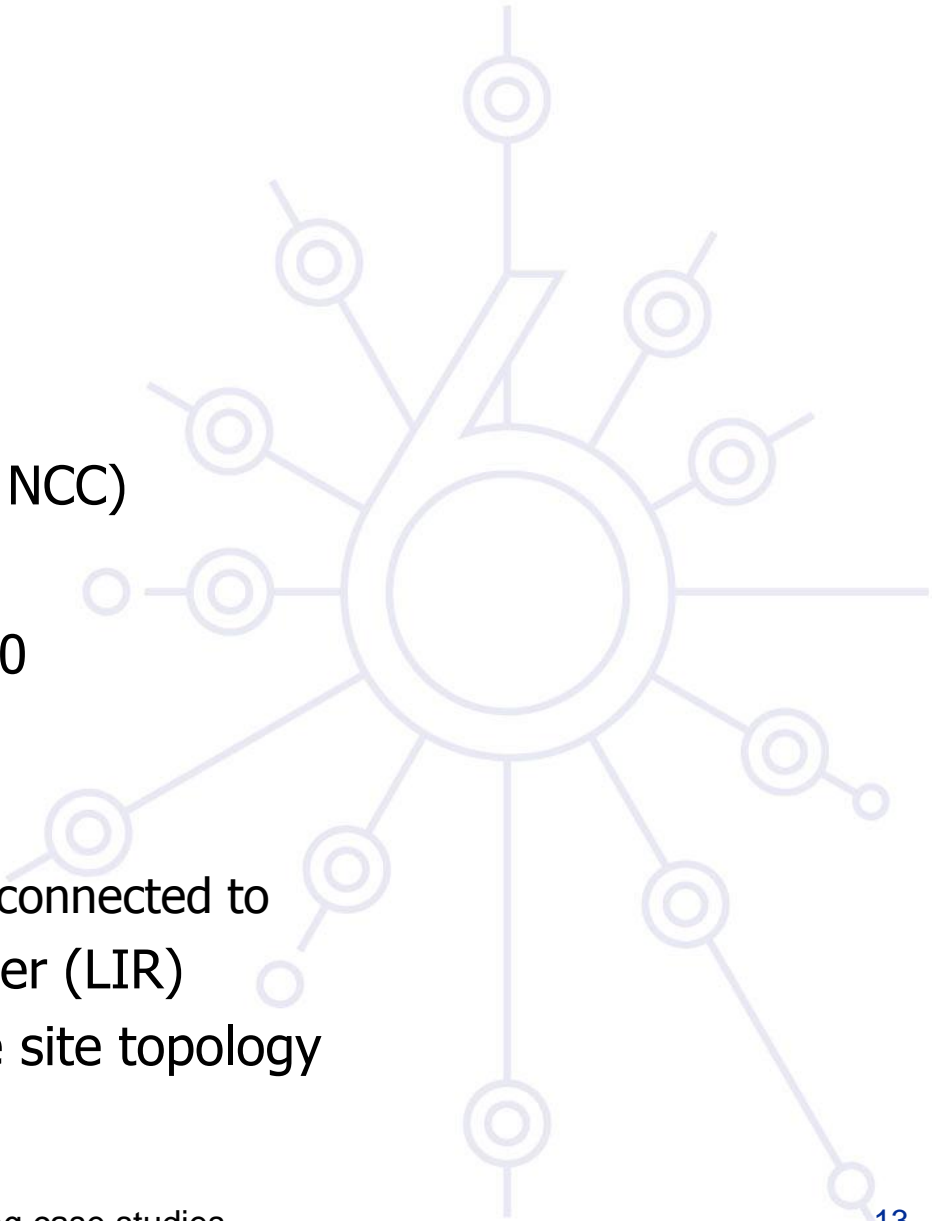
- Prefix = 2001:0660::/32
- Allocated by the RIR (RIPE NCC)

Regional Nodes


- POP-ID = 2001:0660:xy::/40

Site

- Site-ID : a /48
 - from RN's prefix (/40) it's connected to
- Site-IDs allocated by Renater (LIR)
- 16 bits are reserved for the site topology



Addressing



2001:0660:



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

2001:0660:300x:/48

Example

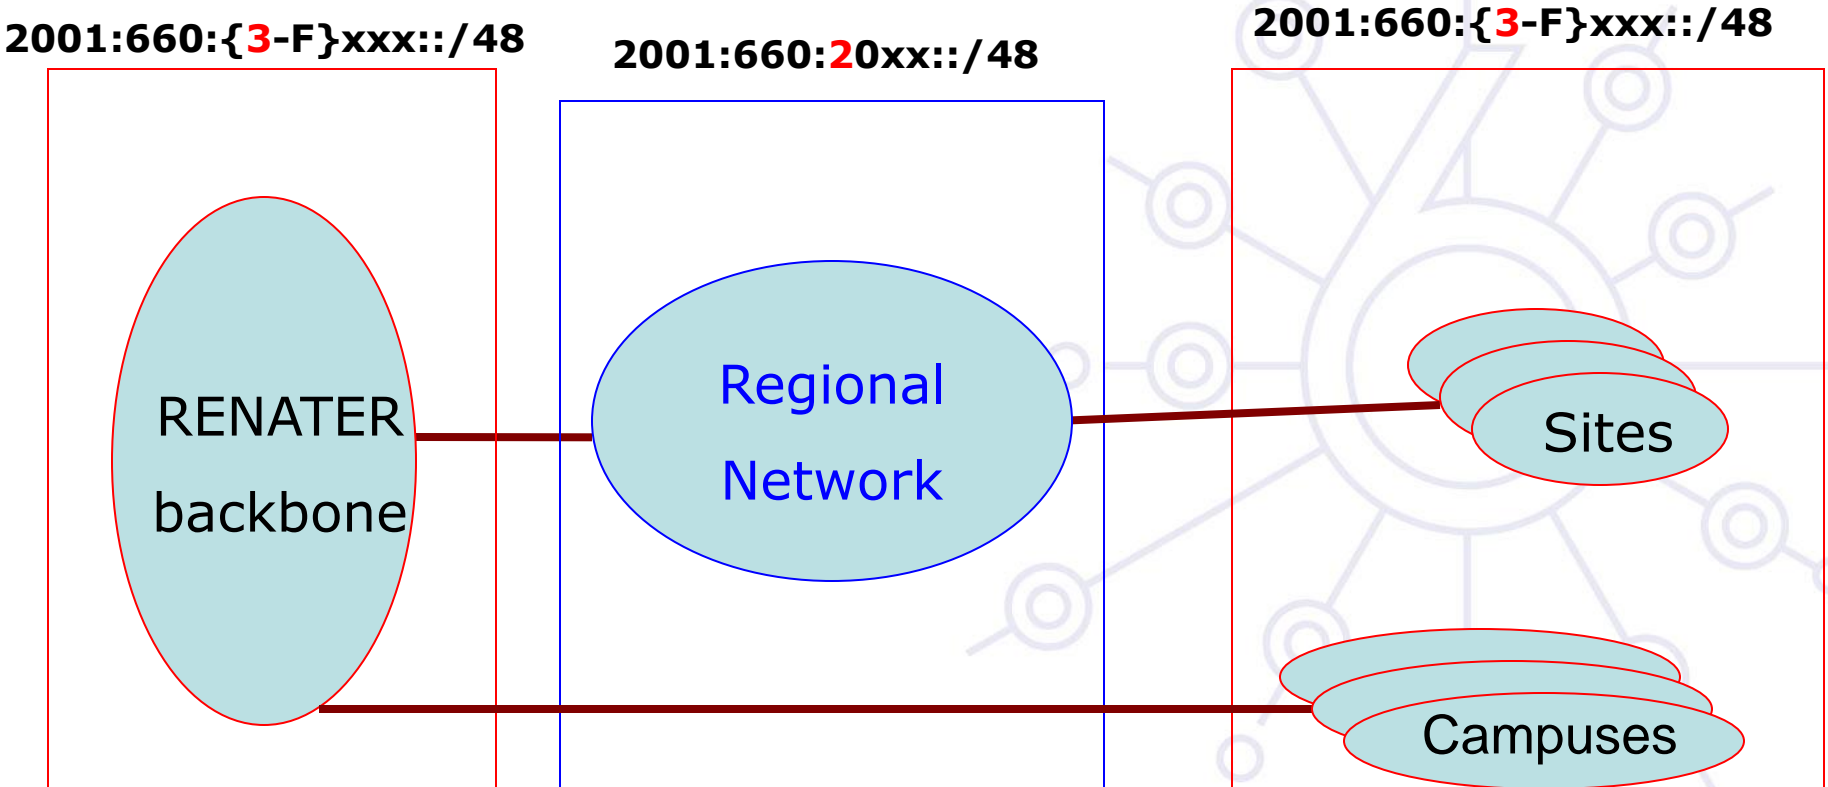
RENATER prefix	2001:0660::/32
POP-ID Strasbourg	2001:0660:4700::/40
Sites connected to Strasbourg's RI	2001:0660:4701::/48 2001:0660:4702::/48 ...

Regional Networks Addressing

Two possibilities

- Uses its own prefix (Commercial ISP)
- Uses RENATER's address space
 - 2001:0660:2---::/48
- In both cases
 - **Sites** are addressed in Renater's prefix
 - 2001:0660:{3-F}---::/48
 - Interco Network (site – Regional / MAN)
 - First /64 from the Site-ID

Addressing scheme






deploy

NIFF/HUNGARNET IPv6 numbering

NIIF/HUNGARNET backbone

April 2006



IPv6 deployment at NIIF/Hungarnet

Initial IPv6 deployment:

- MPLS based backbone: 6PE with additional dual stack routers + sometimes tunnels at connected institutions

Second phase (2004):

- Router upgrade for HW based IPv6 forwarding
- Used features
 - Routing: IPv4 (unicast, multicast), IPv6 (unicast only), OSPFv2, OSPFv3, BGP, MPLS VPNs
 - Netflow, minimal QoS
 - IPv6 multicast with additional dual stack routers with tunnels

Third phase (2008):

- Software upgrade for IPv6 multicast support
- Netflow v9 support

IPv6 address space – based on flexible address allocation RFC3531

Location	IPv6 POP addressing:
CNTRL (Central)	2001:0738:0::/36
Gödöllő (Szent István University)	2001:0738:58::/44
BME (Budapest University of Technology and Economics)	2001:0738:2000::/44
KFKI (Research Institute on Physics)	2001:0738:5000::/44
SZEGED (University of Szeged)	2001:0738:7000::/44
MISKOLC (University of Miskolc)	2001:0738:6000::/44
PECS (University of Pécs)	2001:0738:7800::/44

Site addressing

Each site (including site infrastructure) gets /48:

- each NIIF managed site the 16 bit SLA is allocated based on the following convention: <SLA> = Address segmentation within the POP
- Where for <SLA>:
 - Range: 0000 till 00FF: Loopback addresses
 - Range: 0100 till 01FF: Intra-pop point-to-points (if it necessary to number it)
 - Range: 0200 till 02FF: connections to HUNGARNET member of institution
 - Range: 0300 till 03FF: external IPv6 connectivity (e.g. local IPv6 peering)
 - Range: 0400 till 04FF: POP Local Ethernets

IPv6 loopback addresses

loopback address will also be used for operational and management actions on the equipment, and for routing protocols like iBGP, which will use these addresses for terminating the peering-sessions.

Loopback addresses have typically a prefix mask of /128. This will avoid unnecessary unused addresses although address conservation is not really an issue in IPv6.

Link IPv6 addresses?

Not necessary!

- OSPFv3 is working with link-local
- IS-IS not necessary

IGP table can be quite small!

- Reduces the convergence time

Customer network is propagated into BGP (even if static routes are used)

- not with redistribute
- with network statement

Drawback:

- Traceroute can pick up arbitrary IPv6 address as a reply source -
- Avoid - configure on each point-to-point links:
 - `ipv6 unnumbered loopback0`

Link IPv6 addresses -other options

/127: not a good idea

- the all-zeros address is supposed to be the any router anycast address although this is not widely implemented today - see more RFC 3627

/126: works

- although the top 128 addresses are reserved for anycast stuff

/120: no clashes with top 128 anycast addresses

/112: alignment is on a nice colon boundary

/64: based on RFC 3513

- Allows to use EUI-64 addressing
- advisable for point-multipoint and broadcast link scenarios

Customers' Nets Addressing

Two possibilities

- Uses its own prefix (Commercial ISP)
- Uses NIIF/Hungarnet's address space
 - 2001:0738:<customer id>::/48 (/44 pre-allocated, /48 assigned)

Conclusion

Preparing an IPv6 addressing plan is a bit complex Plan it in advance ...

- Not forgetting your PoPs equipment (loopbacks, admin LANs, interconnects ...)

Draw benefit from aggregation

- Smaller routing tables to manage (even in the core)
- Less prefixes to advertise to BGP peers

Lot of people have an experience yet ...

- Not necessary to reinvent the wheel ;)



deploy

Questions ...