

# LHCONE European design & implementation

Roberto Sabatino, DANTE

LHCONE Meeting, Washington, June 13-14 2011

- Overview European Approach
- Implementation status
- US setup and transatlantic connectivity for LHCONE
- Lessons learnt
- Conclusions

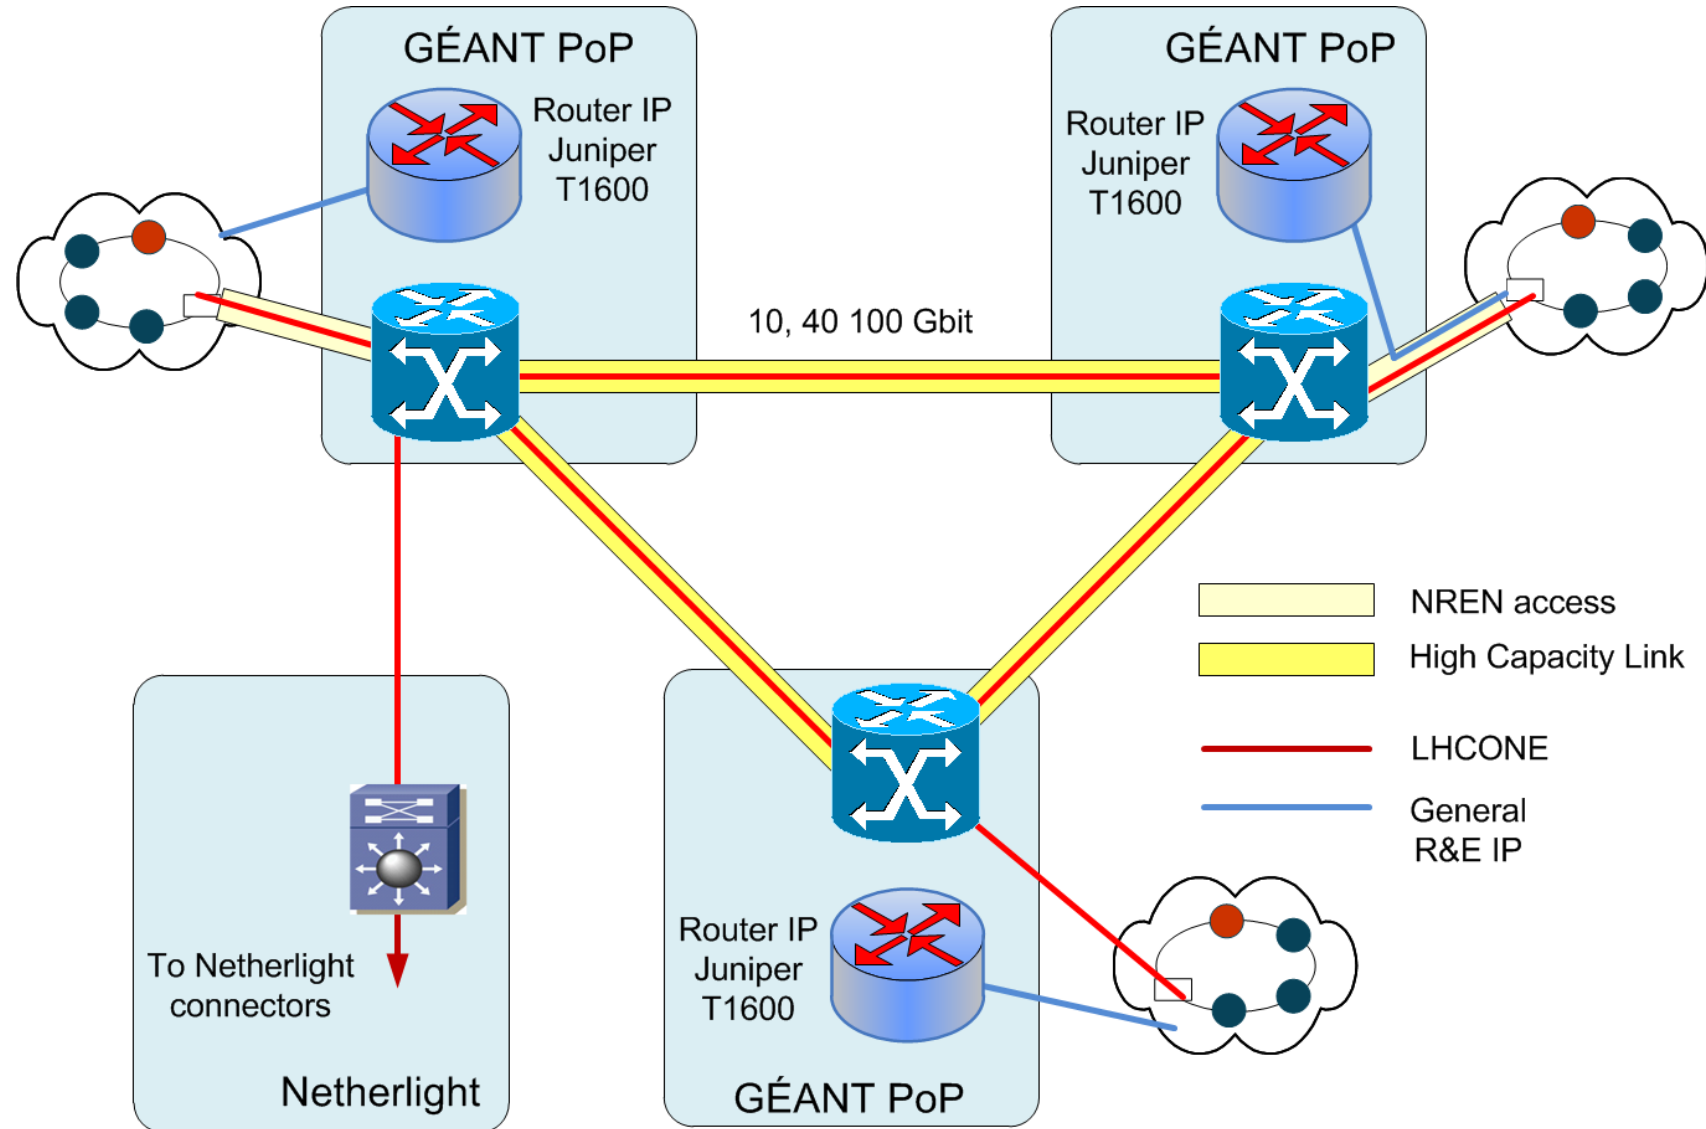
# Overview European Approach



- Designed according to European user's requirements
  - European workshop on 5<sup>th</sup> April agreed way forward
- Set of Distributed Exchange Points based on GEANT infrastructure
- Interconnection with Netherlight

# Long-term solution

## Layout of (generic) Distributed Exchange



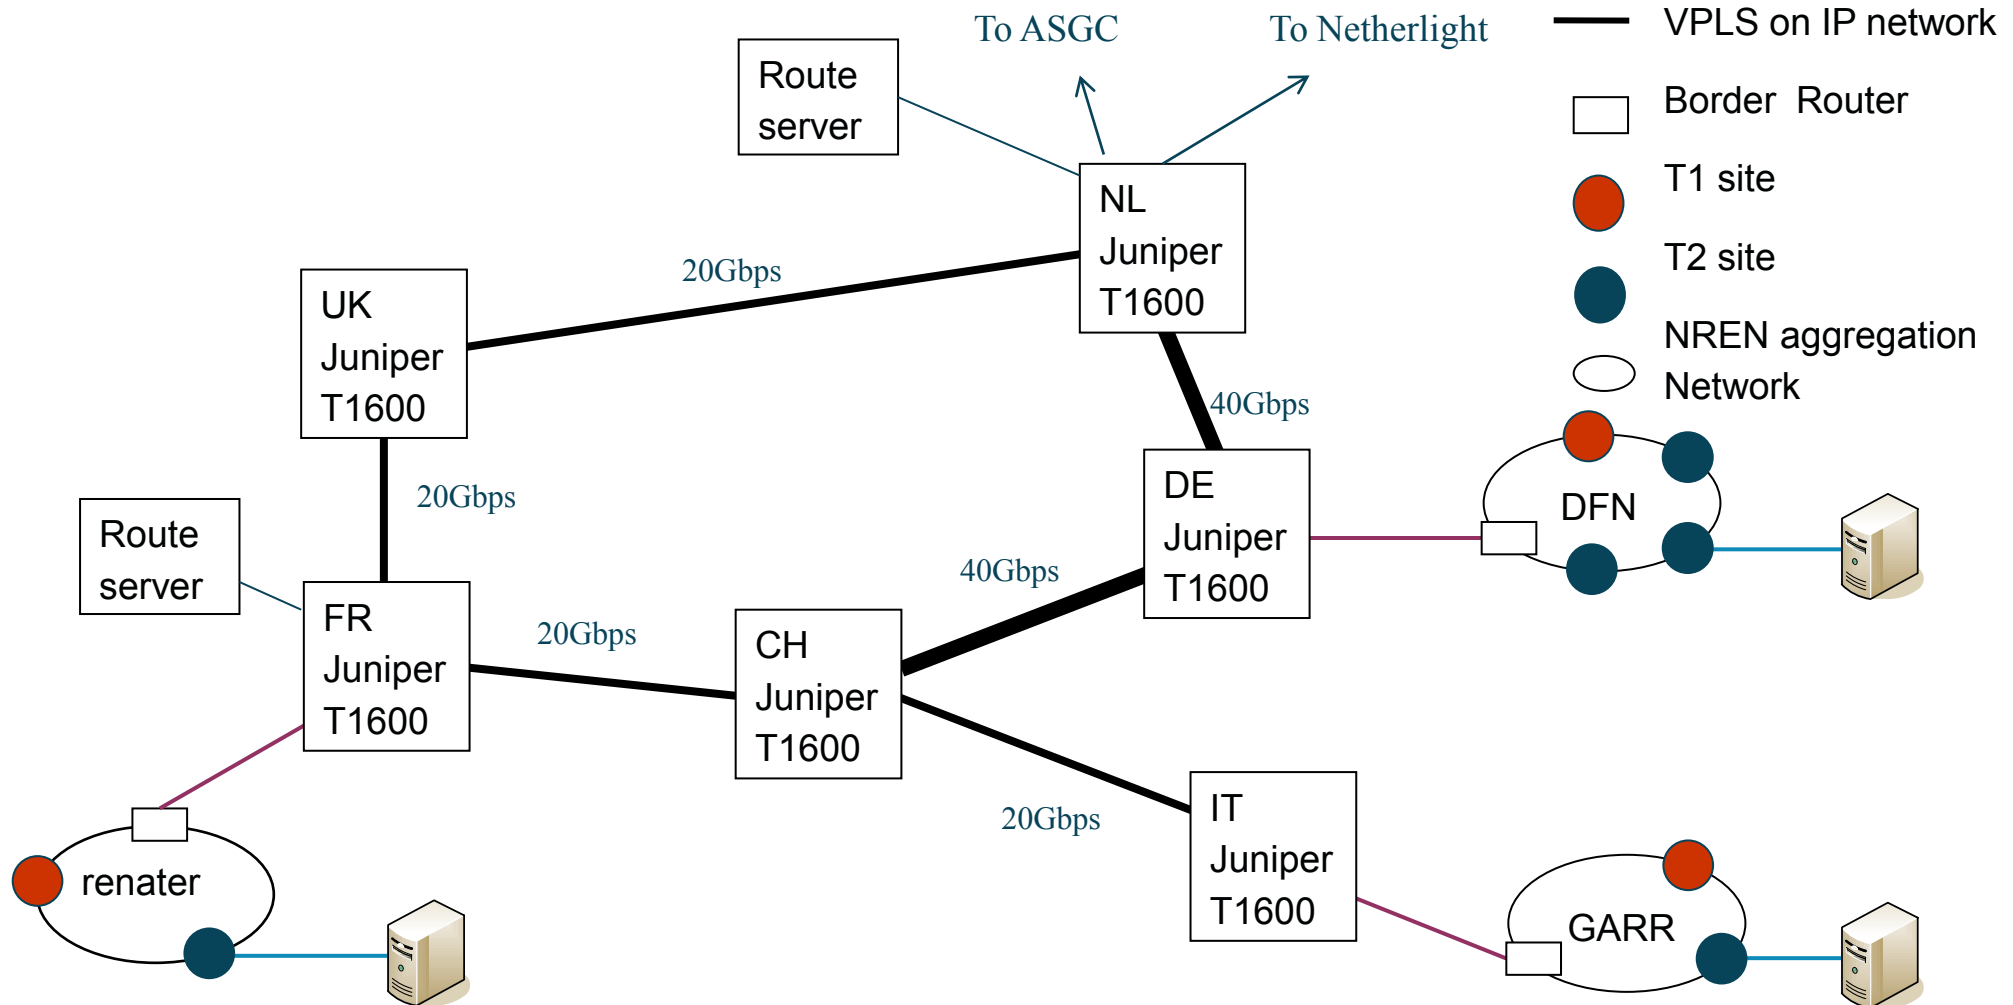
- Logically identical to long term solution but
  - smaller in scope (# of sites connected)
  - Different technology, in core and in NREN connectors
  - Shared vs dedicated capacity
- Use VPLS technology on existing Juniper T1600 routers to emulate a distributed LAN
- Make use of existing shared (IP) capacity, with logical separation from general R&E IP traffic (with separate monitoring)
  - Confidence there is enough capacity on existing backbone links
  - 40Gbps CH-DE-AMS
  - 20Gbps NL-UK-FR-CH and IT-CH
  - May upgrade links if demonstrated necessary

# Access options for NRENs to LHCONE



- During prototype
  - Via a VLAN or LSP on existing port
  - Via additional separate physical port
  - Use backup IP port
- Long Term
  - Via a VLAN on main access (e.g if main access is > 10Gbps)
  - Via a VLAN on additional access
  - ...other...TBD

# European prototype overview



# Services supported in prototype



- P2mp and static p2p services
- Dynamic services not yet available. Expected 2012
  - GEANT-NREN pilot started now
- perfSONAR release available May 2011
- Note: p2p circuits (with BW guarantees) are already available via GEANTPlus.
  - No need to establish new infrastructure
  - May need to install additional capacity, depending on demand

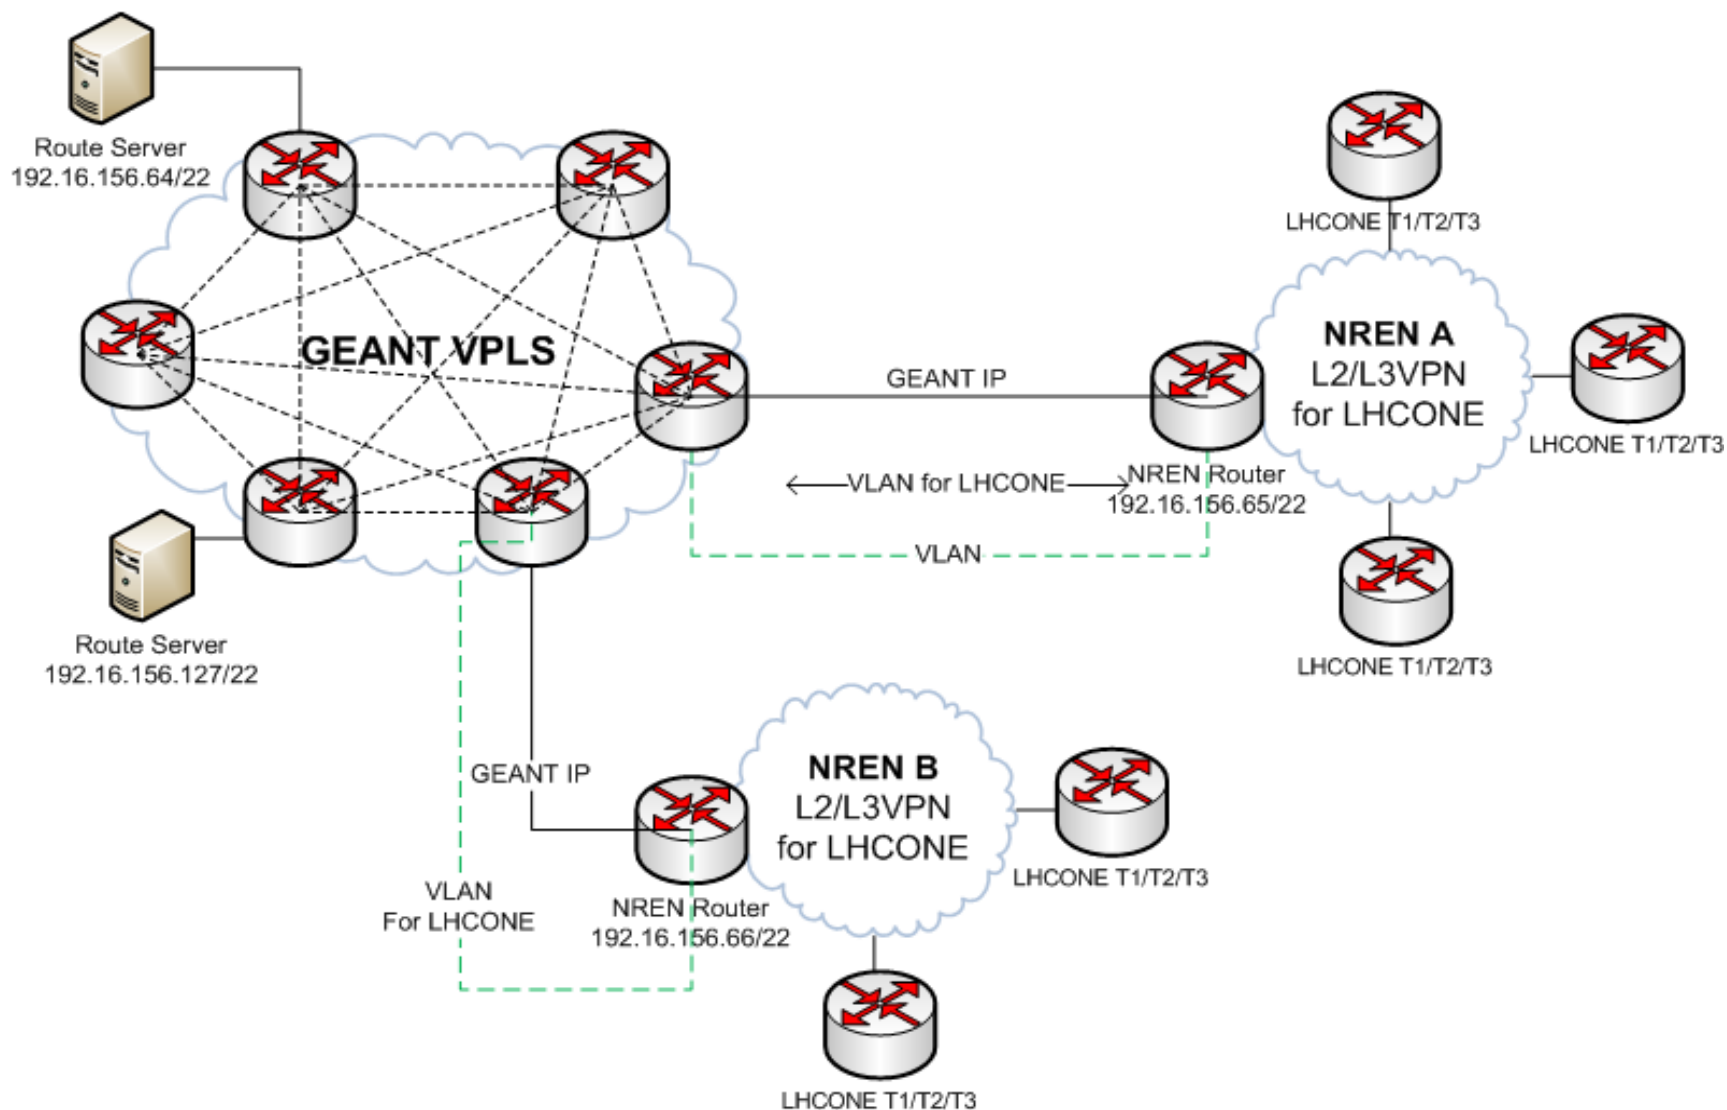


# Prototype costs within Europe




- Depending on how NREN aggregator networks (T2 clouds) connect, there may be an additional 10GE port in some locations
- We do not expect to have to add additional backbone capacity, for the prototype
  - There may be exceptions, depending on traffic growth
- Route Server

# IMPLEMENTATION: LHCONE Setup in GEANT

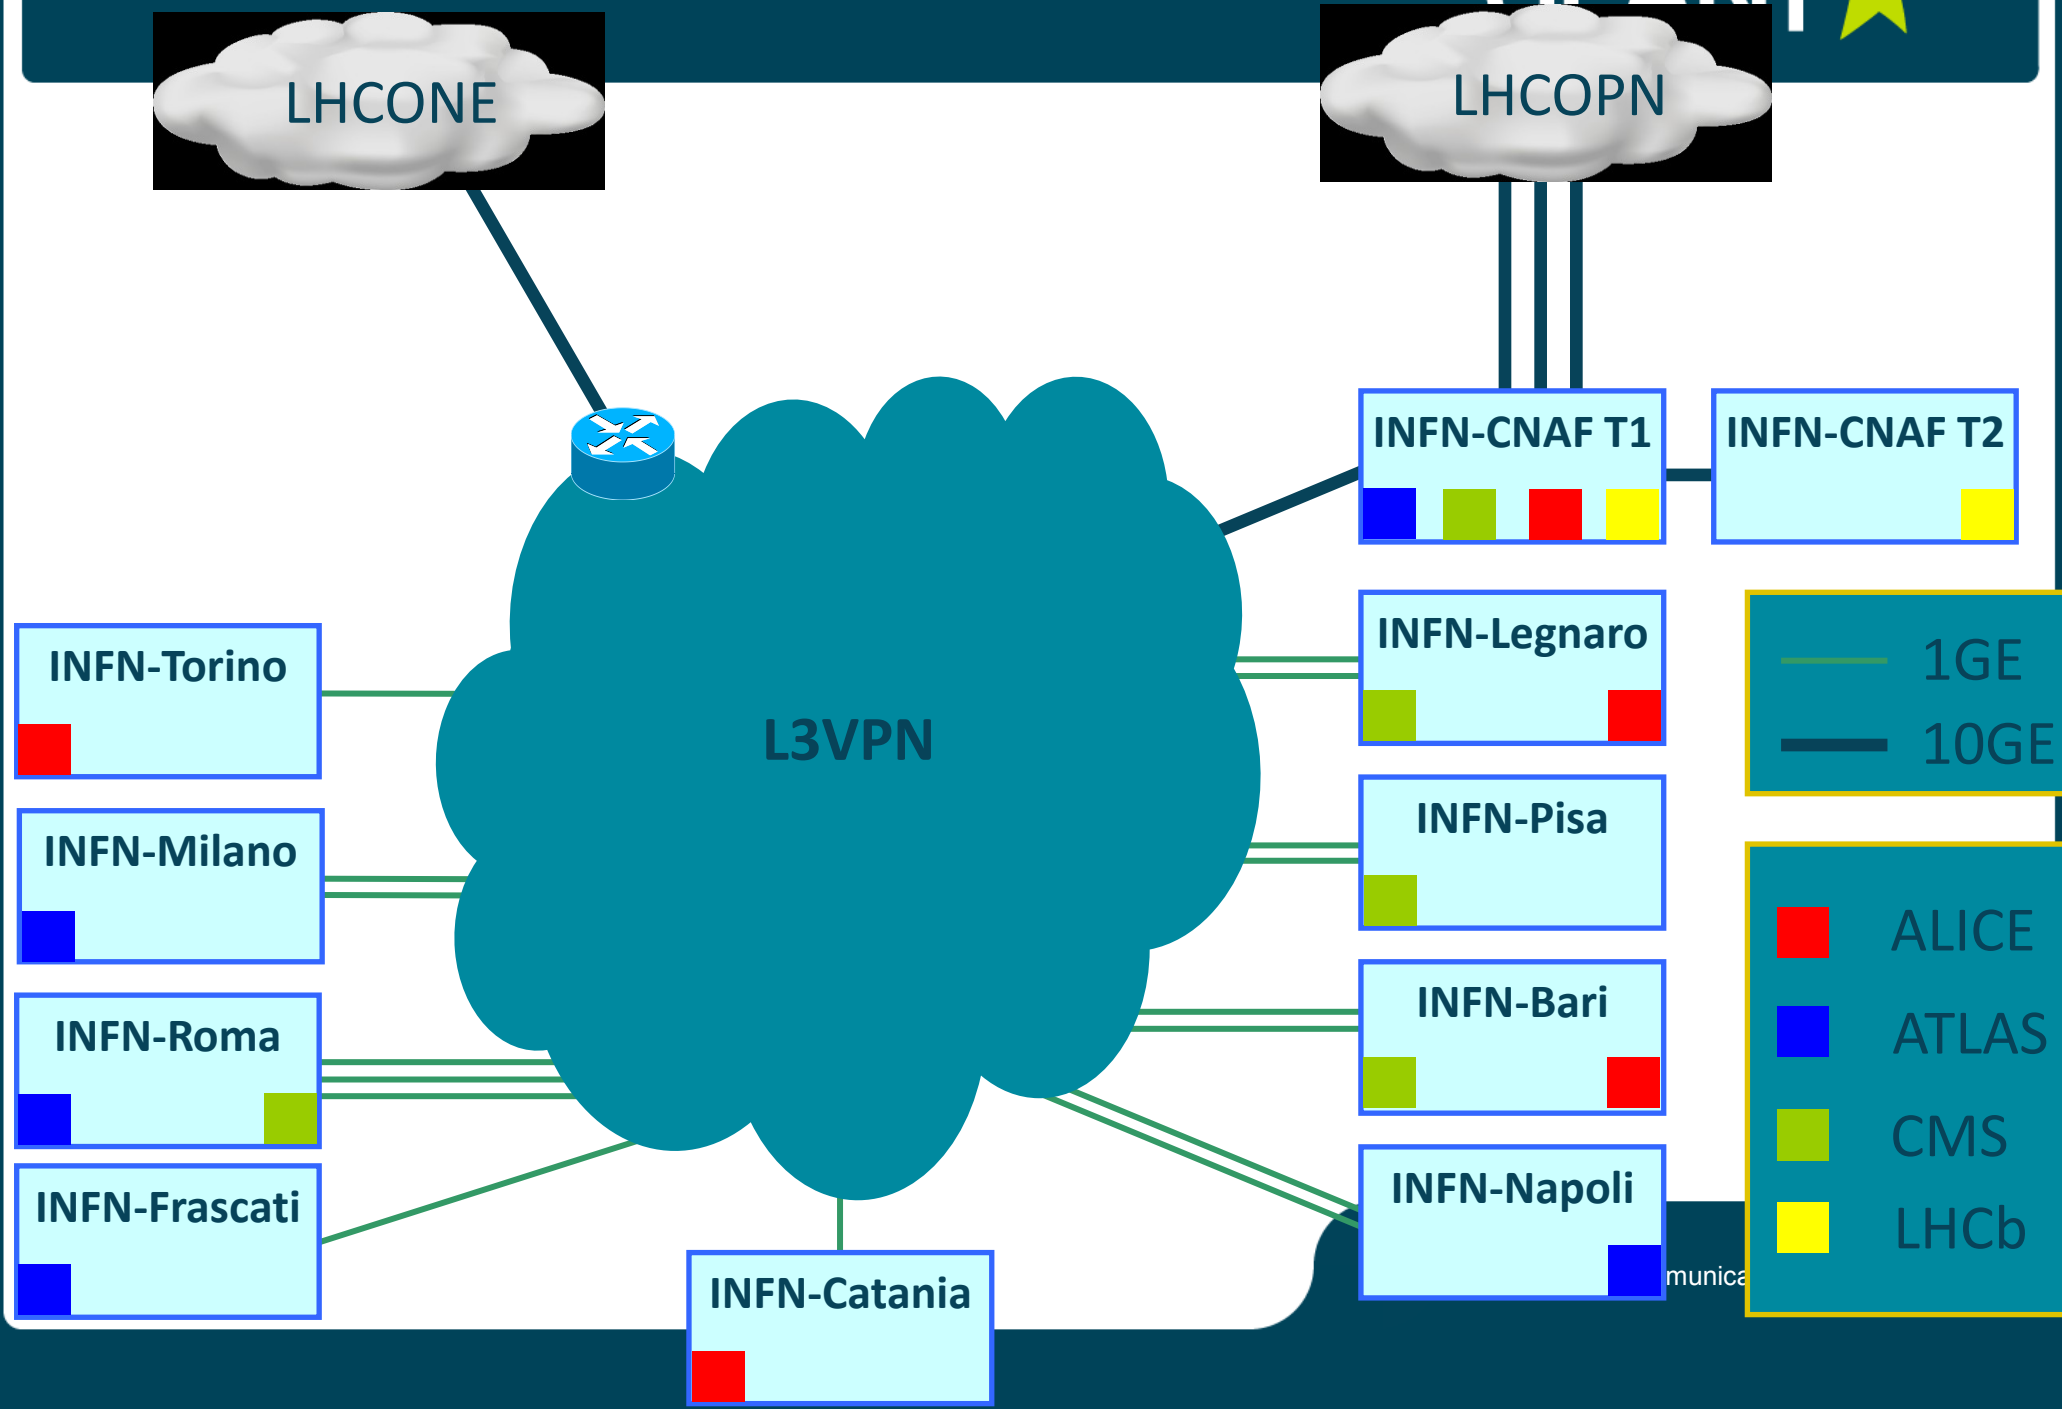


# LHCONE Setup in GEANT

- GEANT IP Network will be invisible to T2 sites (via LHCONE)
- All LHC connectors will be within the same IP subnet
- All sites should peer with all route servers
  - Looking for a workaround
- Technical service guidelines agreed with participants



# Italian Tier2 prototype

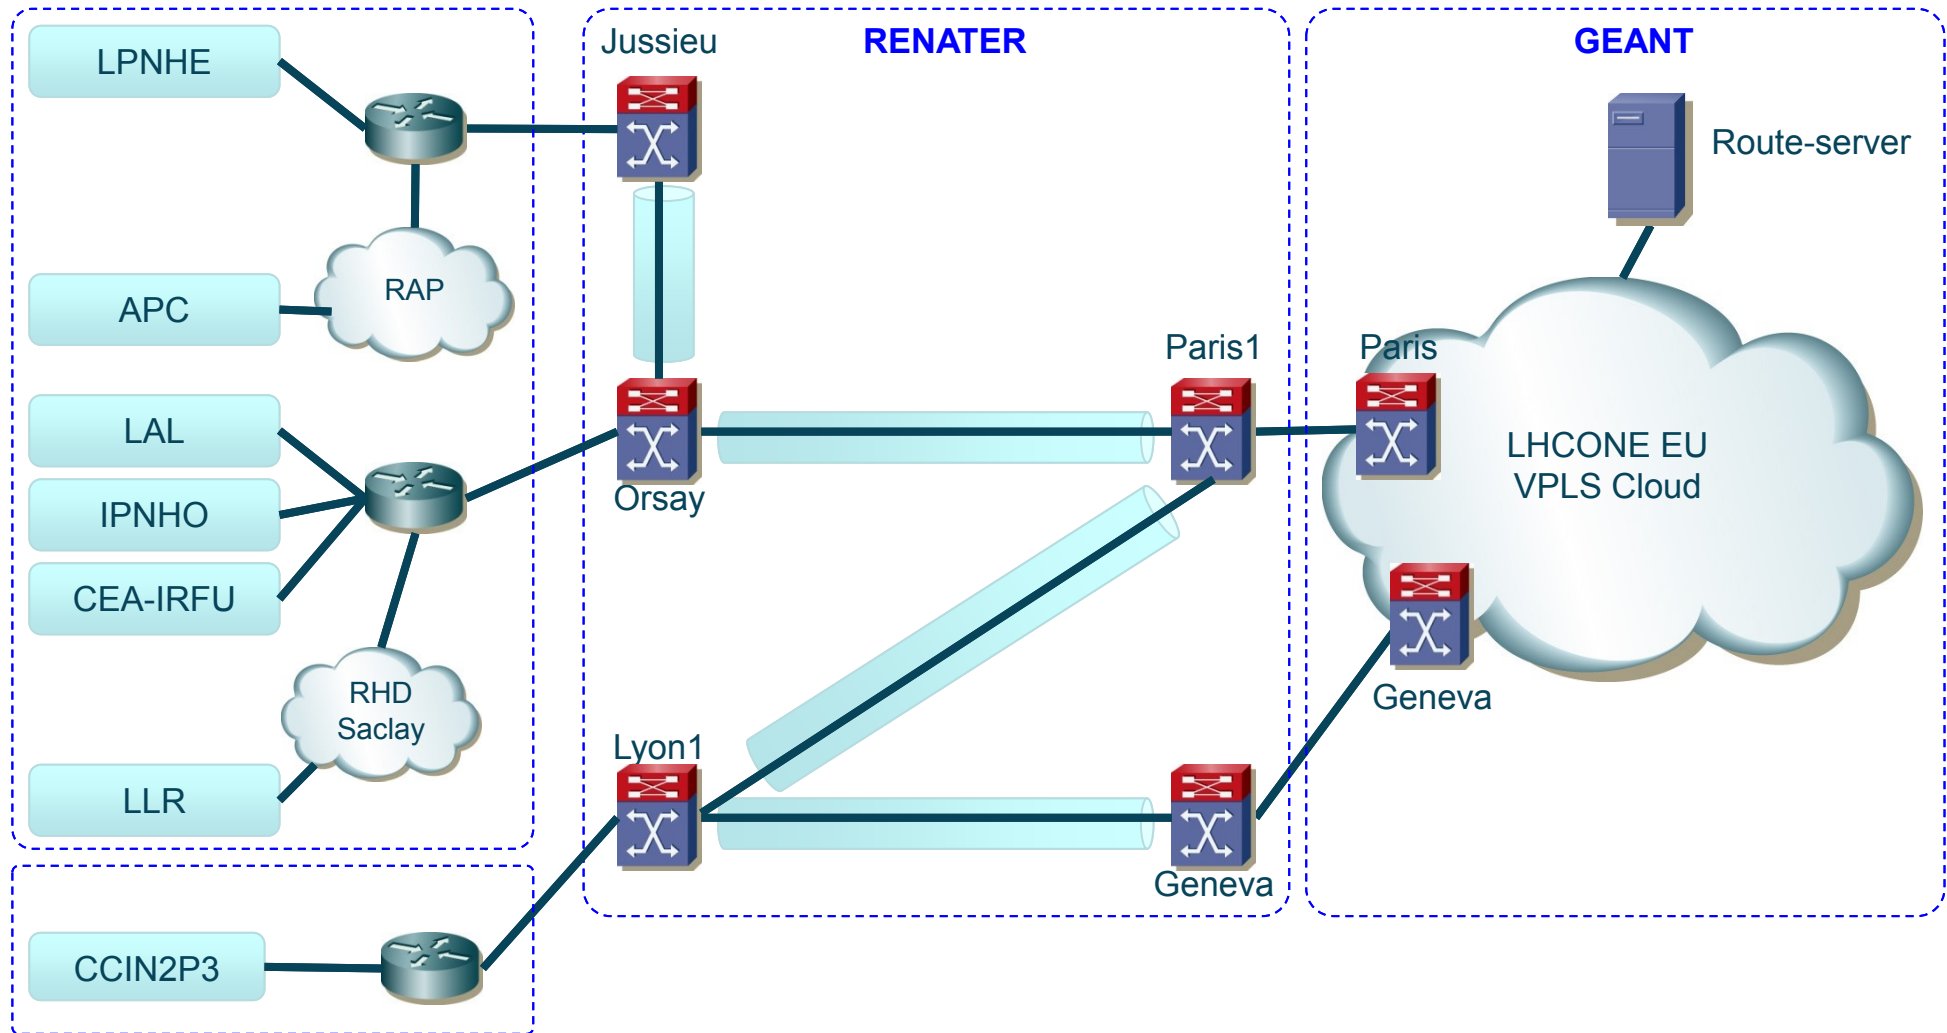


# LHCONE FR Prototype target architecture



GRIF - Distributed T2)

T1




We are also currently investigating if a VRF LITE implementation would be more appropriate

connect • communicate • collaborate



## X-WiN and HEPPI (schematic)



- green: inner core with 10GE
- all traffic in L3VPN
- Traffic flow HEP1<->HEP2 via HR
  - except if connected to same T2R
- additional Loopback interfaces on all routers that are part of HEPPI
  - serve as BGP next-hops in L3VPN
- separate OSPF process to distribute next-hops
- TE-Tunnels: OSPF not possible -> static routing for Loopbacks
- GN access as VLAN on existing 2x10GE channel

# Timescale

- VPLS Implementation
  - Routers configured 10<sup>th</sup> June-11
  - Tests will be performed till 17<sup>th</sup> June-11
- Route Server Implementation
  - Route servers deployed 10<sup>th</sup> June-11
  - Tests will be performed till 17<sup>th</sup> June-11
- NREN Connections
  - NRENS will be able to connect LHC T2/3s from week starting 20<sup>th</sup> June-11
- Trans-Atlantic Connectivity
  - GEANT ready to connect on 20<sup>th</sup> June-11

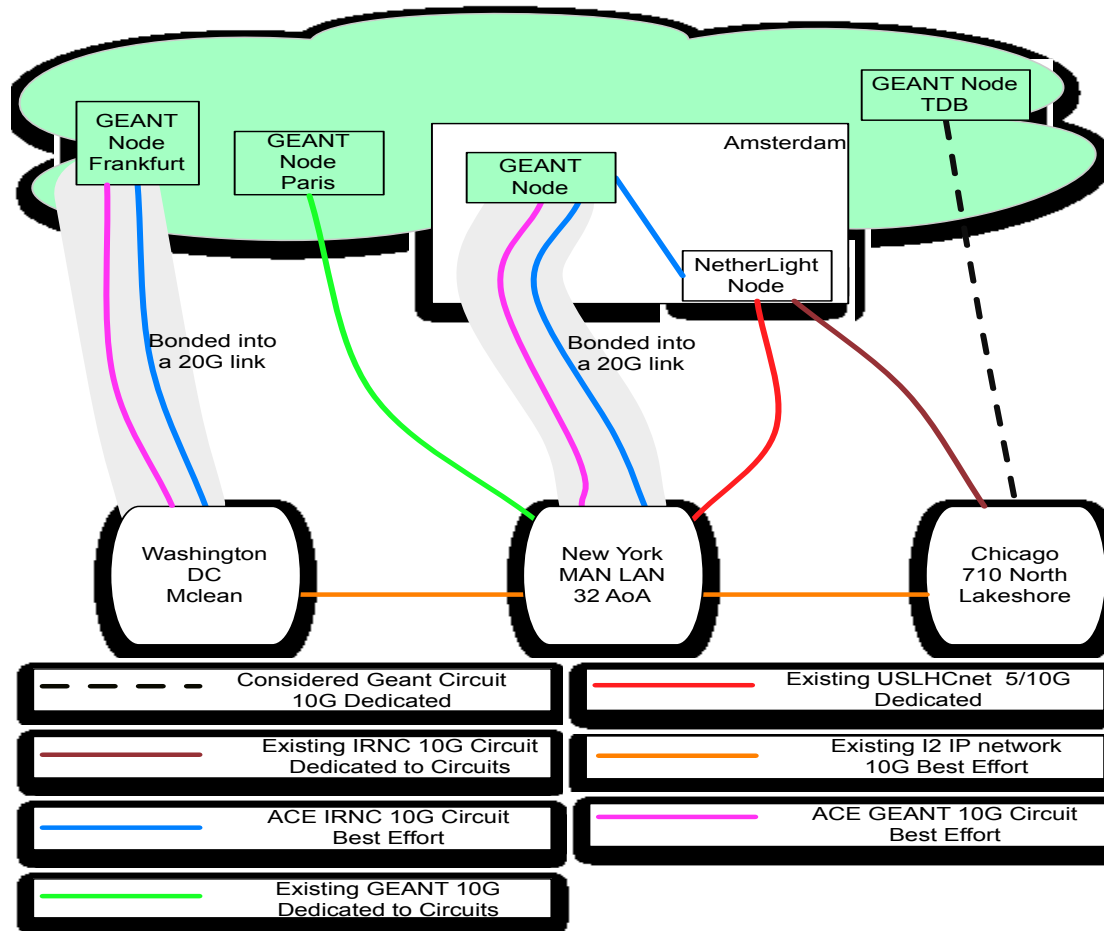
# Conclusions Intra European set-up



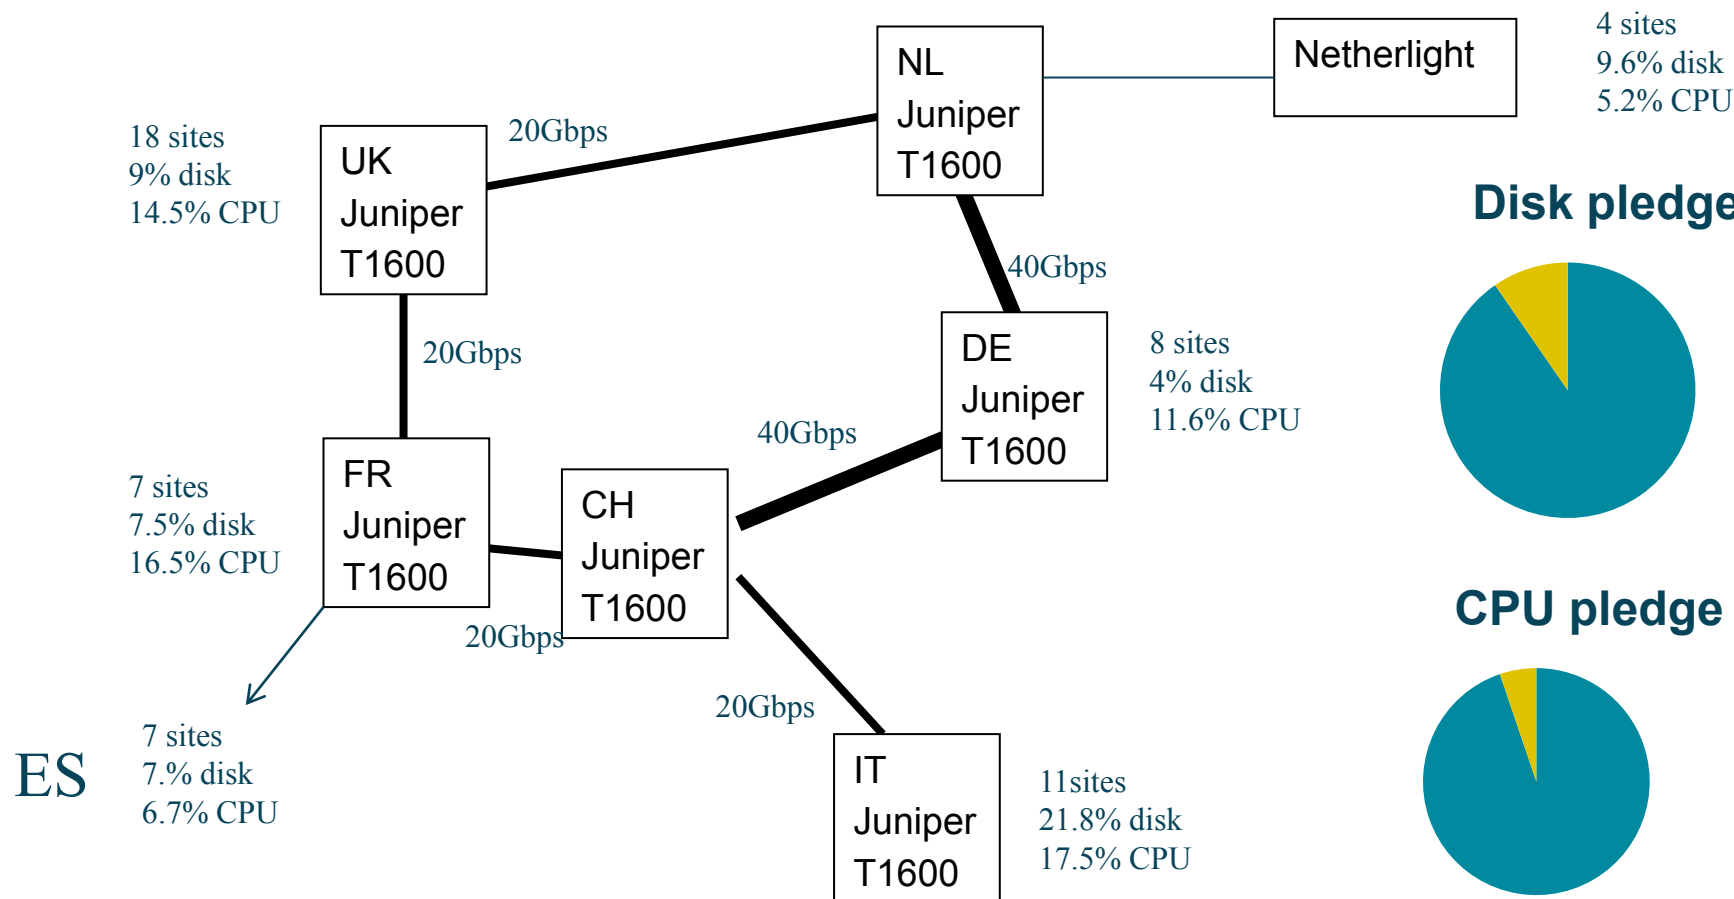
- Conforms to LHCONE architecture
- LHCONE “core” in Europe provided by a set of distributed exchange points in GEANT PoPs, interconnected to Netherlight
  - New equipment in 2012, with allocated capacity from 10G up to 100G
  - VPLS on existing Juniper T1600 for prototype. No need for additional capacity (bar possible exceptions)
- NREN aggregator networks
- Costs:
  - Minimal during prototype
- Prototype implementation well under way




# Transatlantic connectivity - general




# Distribution of European Resources for experiments (from WSLCG)



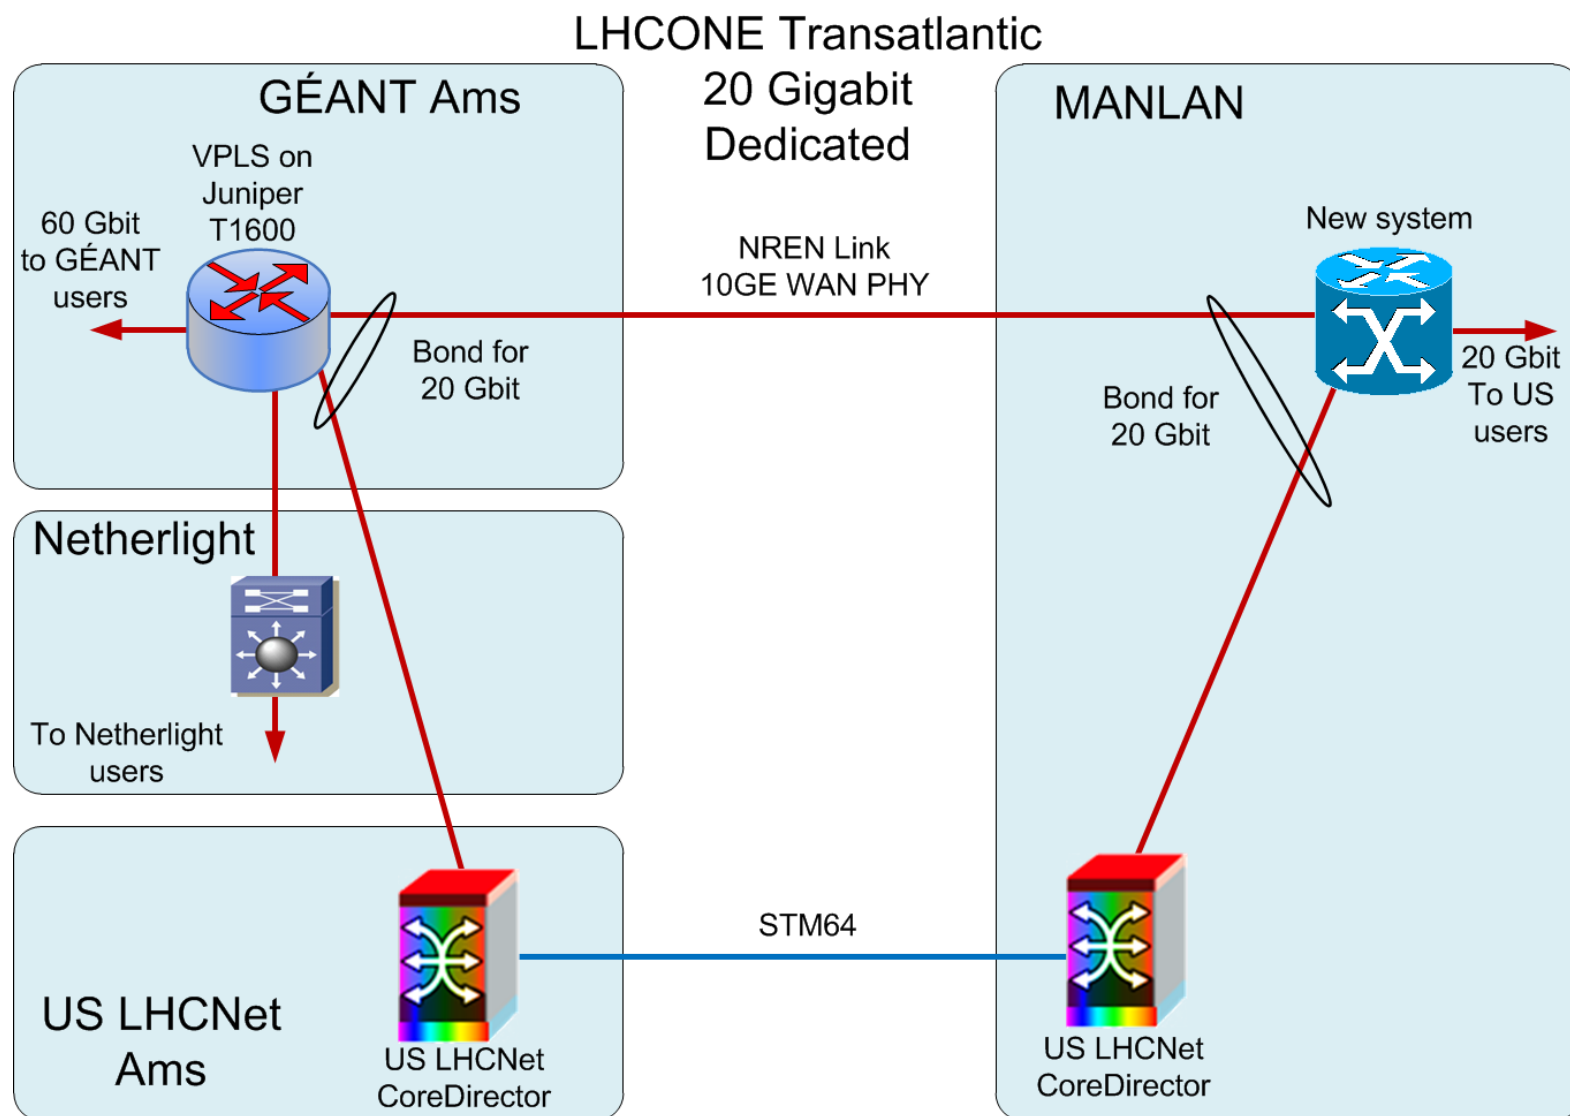
## Disk pledge 2011



## CPU pledge 2011



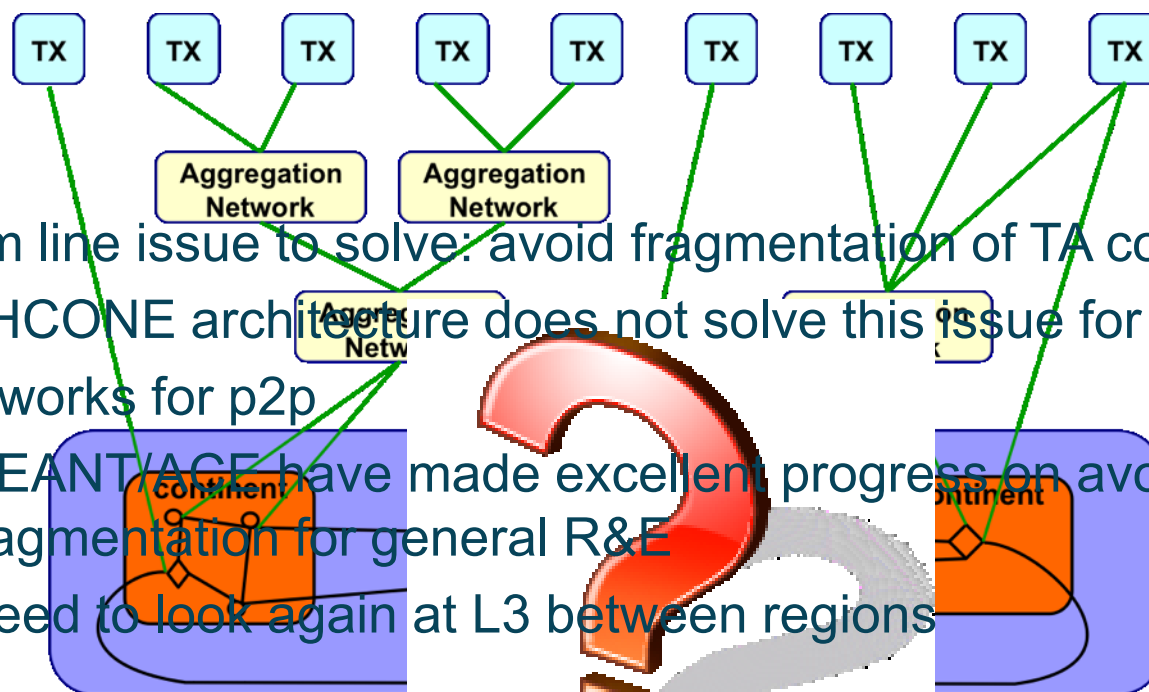
# Transatlantic – p2mp service



# Prototype lesson N.1: Architectural issues



- L2 architecture does not support multi-continent/region p2mp service well:
  - Loop avoidance
  - Resilience
  - NA – Europe – Asia or Asia – NA – Europe Or Europe – Asia – NA ?



- Bottom line issue to solve: avoid fragmentation of TA connectivity
  - LHCONe architecture does not solve this issue for p2mp
  - It works for p2p
  - GEANT/AGE have made excellent progress on avoiding fragmentation for general R&E
  - Need to look again at L3 between regions

◇ distrib  
○ single

connect • communicate • collaborate

- Progress within Europe/GEANT on prototype implementation
- Initial proposal for T.A setup that meets the connectivity needs of the users (***for the prototype, within its limitations***)
  - Partnership between EU NREN funded connectivity and USLHCNET
  - Based on GEANT PoPs in Europe, US to choose PoPs in US
  - GEANT ready from June 20<sup>th</sup>
- Study lessons learnt from prototype, feed into next architecture iteration
  - Recognition of limitation of LHCONE architecture to support p2mp services
    - *Start prototype with one path between EU-N.A*
  - Review architecture
- Objectives & success criteria of prototype need to be clear