Networking for LHC Data Analysis

Ein Überblick

Kars Ohrenberg

6th Annual Workshop of the Helmholtz Alliance "Physics at the Terascale"

Hamburg, 04.12.2012

HELMHOLTZ GEMEINSCHAFT

DESY

The LHC Optical Private Network

- The LHCOPN (from http://lhcopn.web.cern.ch)
  - The LHCOPN is the private IP network that connects the Tier0 and the Tier1 sites of the LCG.
  - The LHCOPN consists of any T0-T1 or T1-T1 link which is dedicated to the transport of WLCG traffic and whose utilization is restricted to the Tier0 and the Tier1s.
  - Any other T0-T1 or T1-T1 link not dedicated to WLCG traffic may be part of the LHCOPN, assuming the exception is communicated to and agreed by the LHCOPN community
- Very closed and restricted access policy

LHCONE Network Map

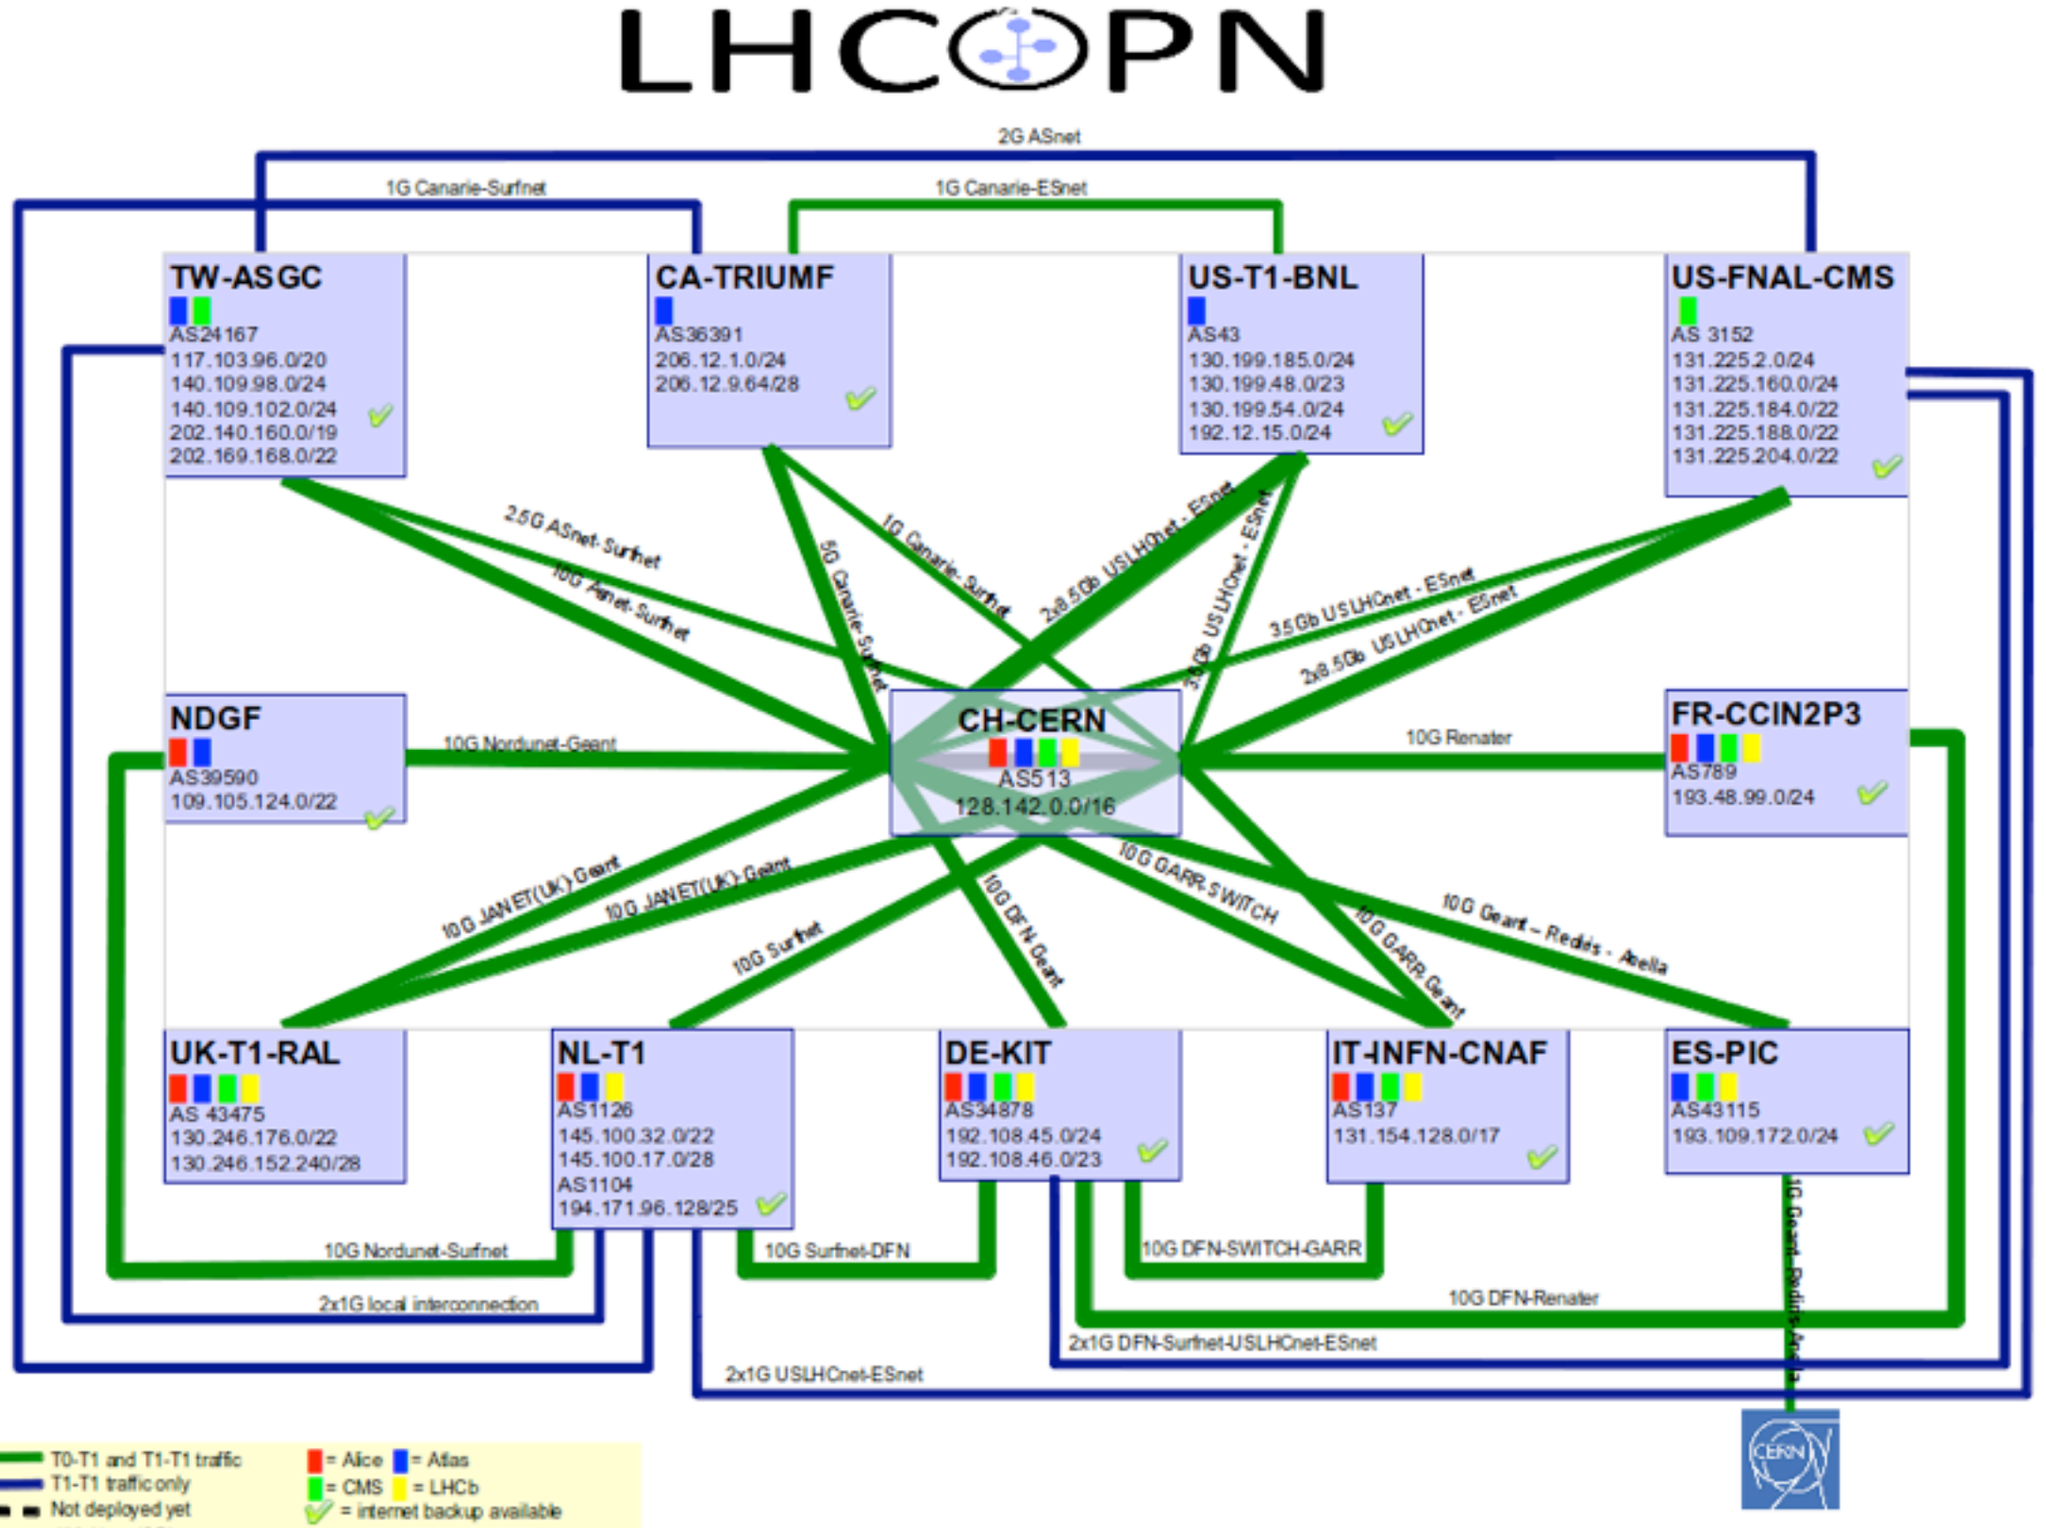

Computing Models Evolution

- The original MONARC model was strictly hierarchical
- Changes introduced gradually since 2010
- Main evolutions:
  - Meshed data flows: Any site can use any other site as source of data
  - Dynamic data caching: Analysis sites pull datasets from other sites „on demand“, including from Tier-2s in other regions
  - Remote data access
- Variations by experiment
- Increased reliance on network perfomance!

LHC Open Network Environment

- With the successful operation of the LHC accelerator and the start of the data analysis, there has come a re-evaluation of the computing and data models of the experiments
- The goal of LHCONE (LHC Open Network Environment) is to ensure better access to the most important datasets by the worldwide HEP community
- Traffic patterns have altered to the extent that substantial data transfers between major sites are regularly being observed on the national and GÉANT IP backbones, often lasting for several days.
- The main principle is to separate the LHC traffic, thus avoiding degraded performance both to the LHC community and to the other users of the academic IP networks.
- The objective of LHCONE is to provide entry points into a network that is private to the LHC T1/2/3 sites.
- LHCONE is not intended to replace LHCOPN but rather to complement it

LHCONE VRF Map (from Bill Johnston, ESNet)

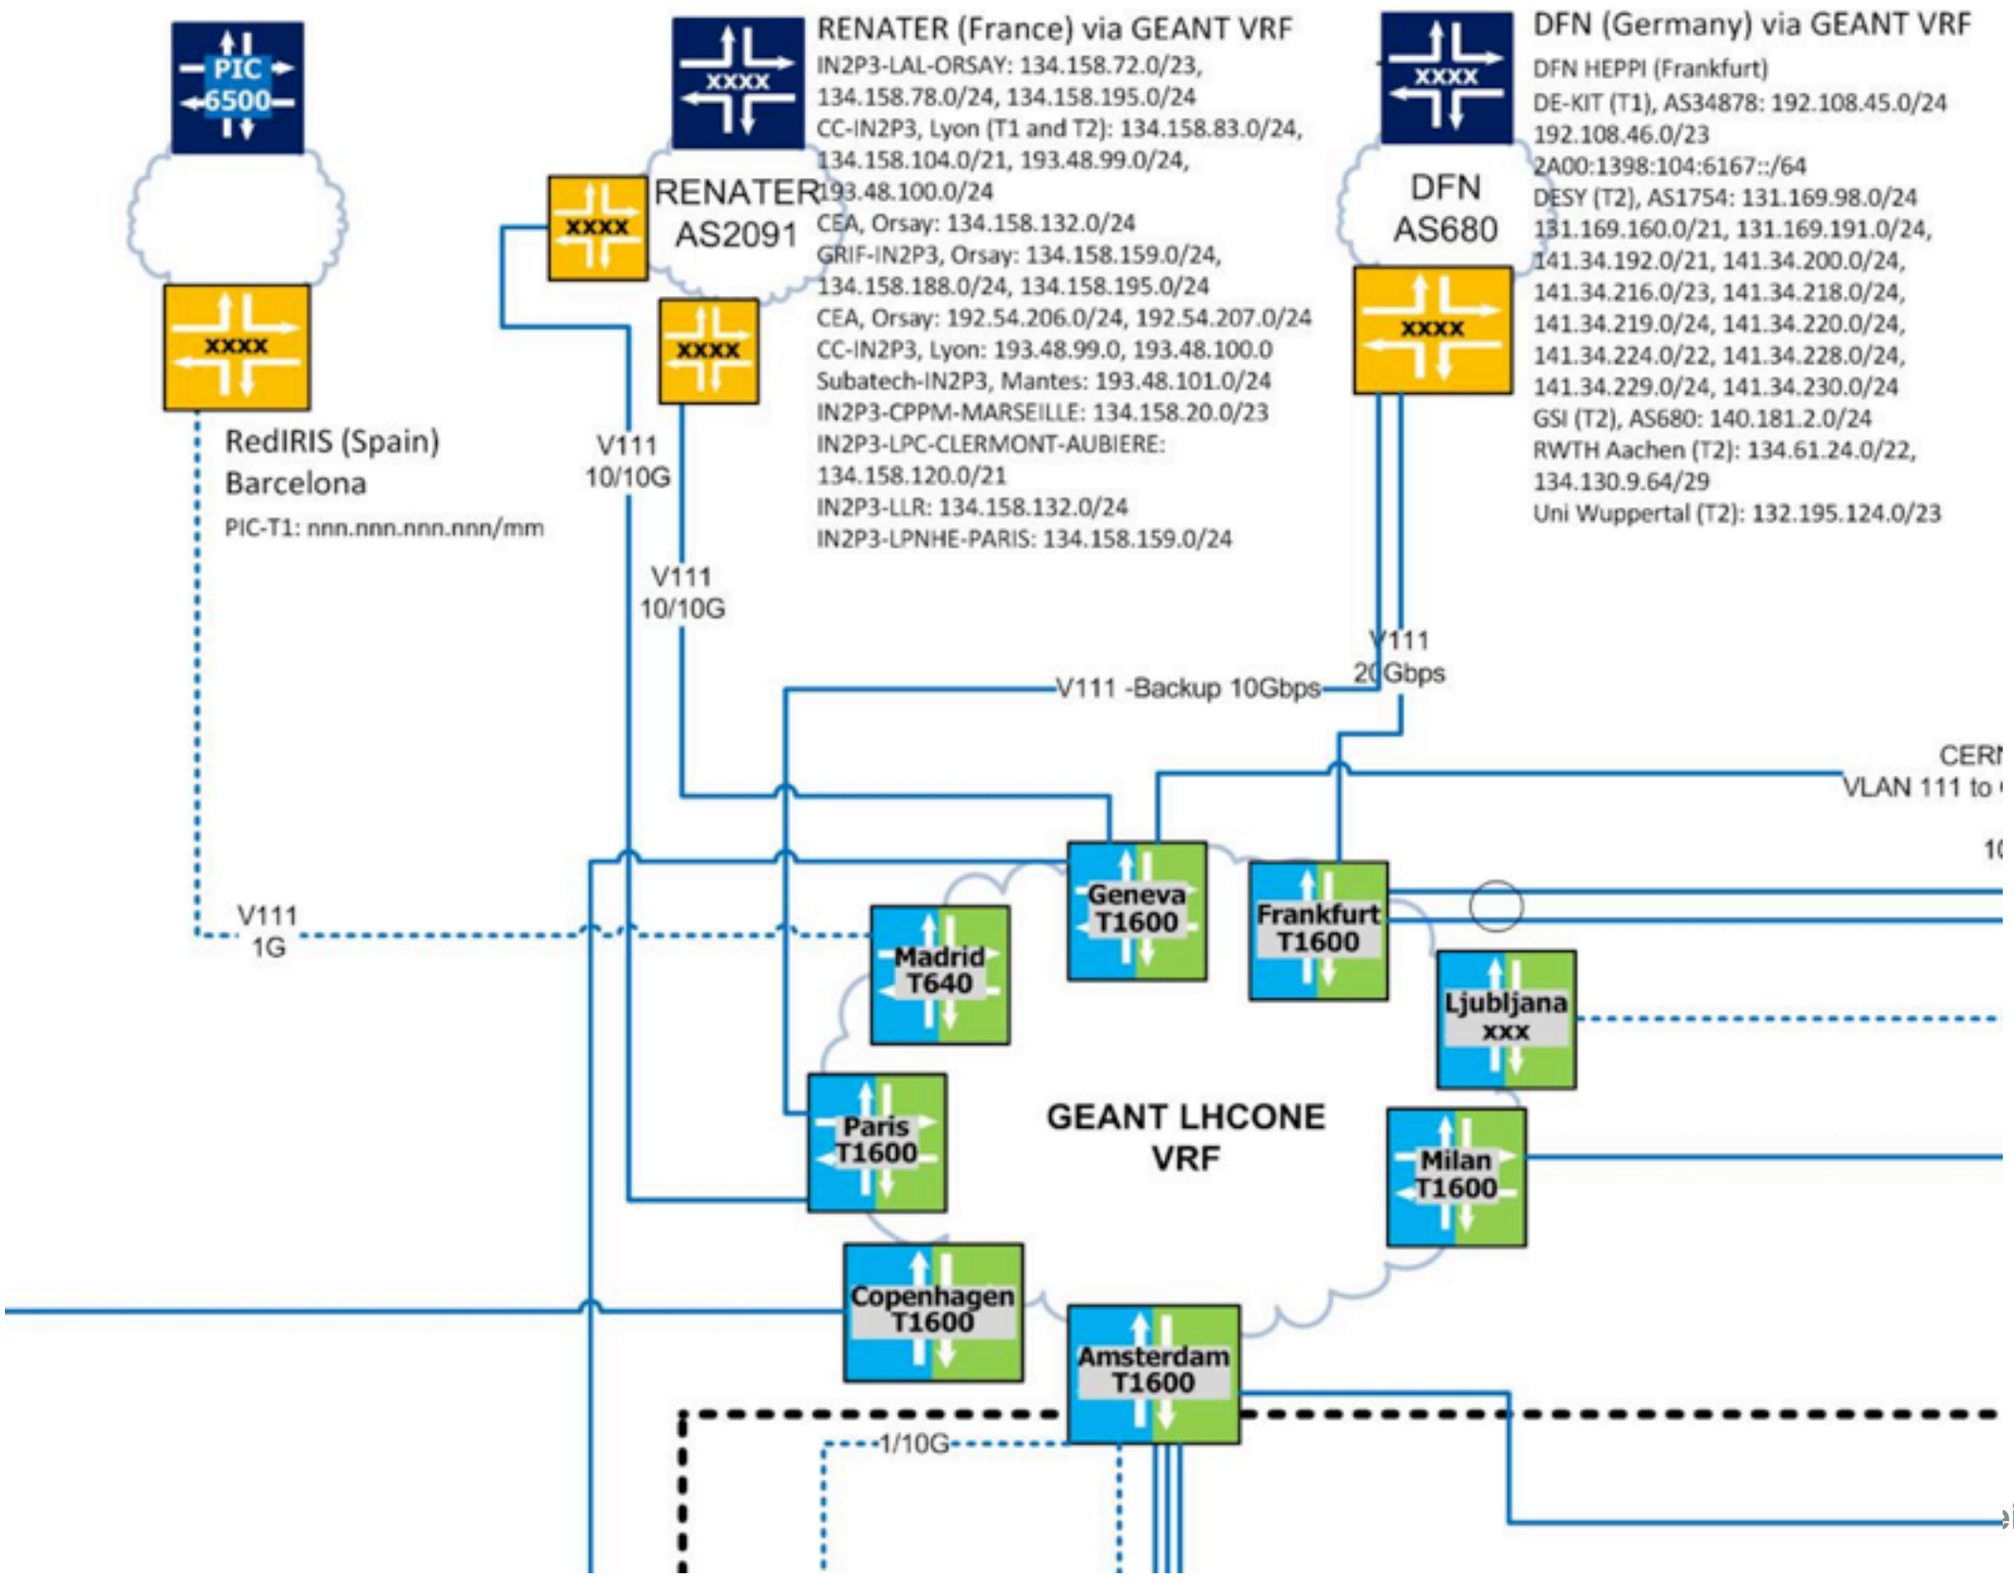

LHCONE VRF Map (from Bill Johnston, ESNet)

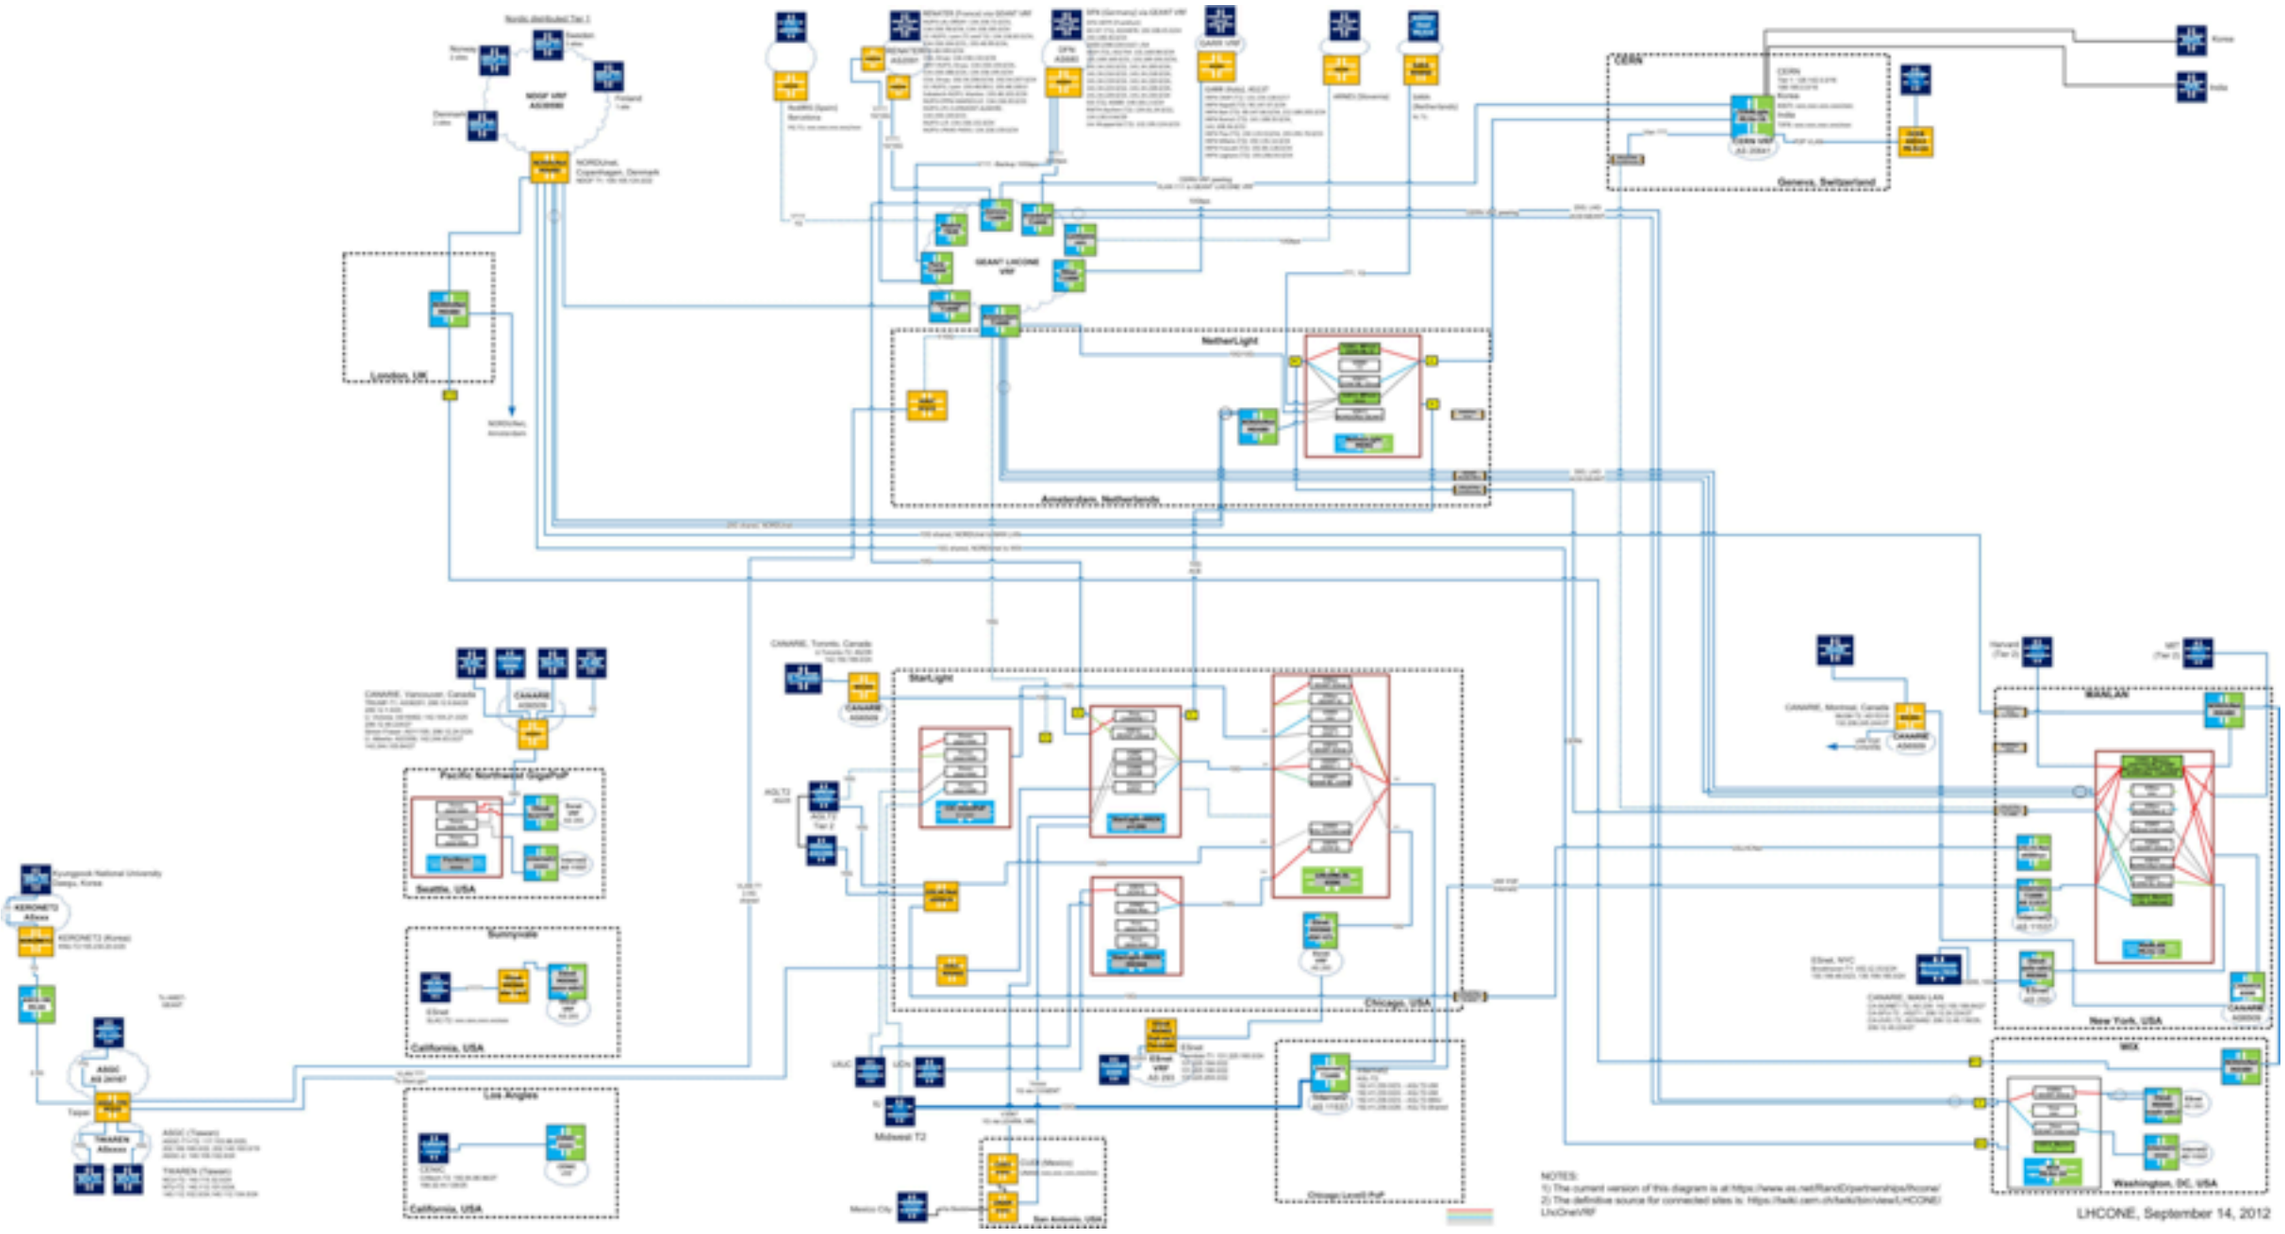

LHCONE L3 Map

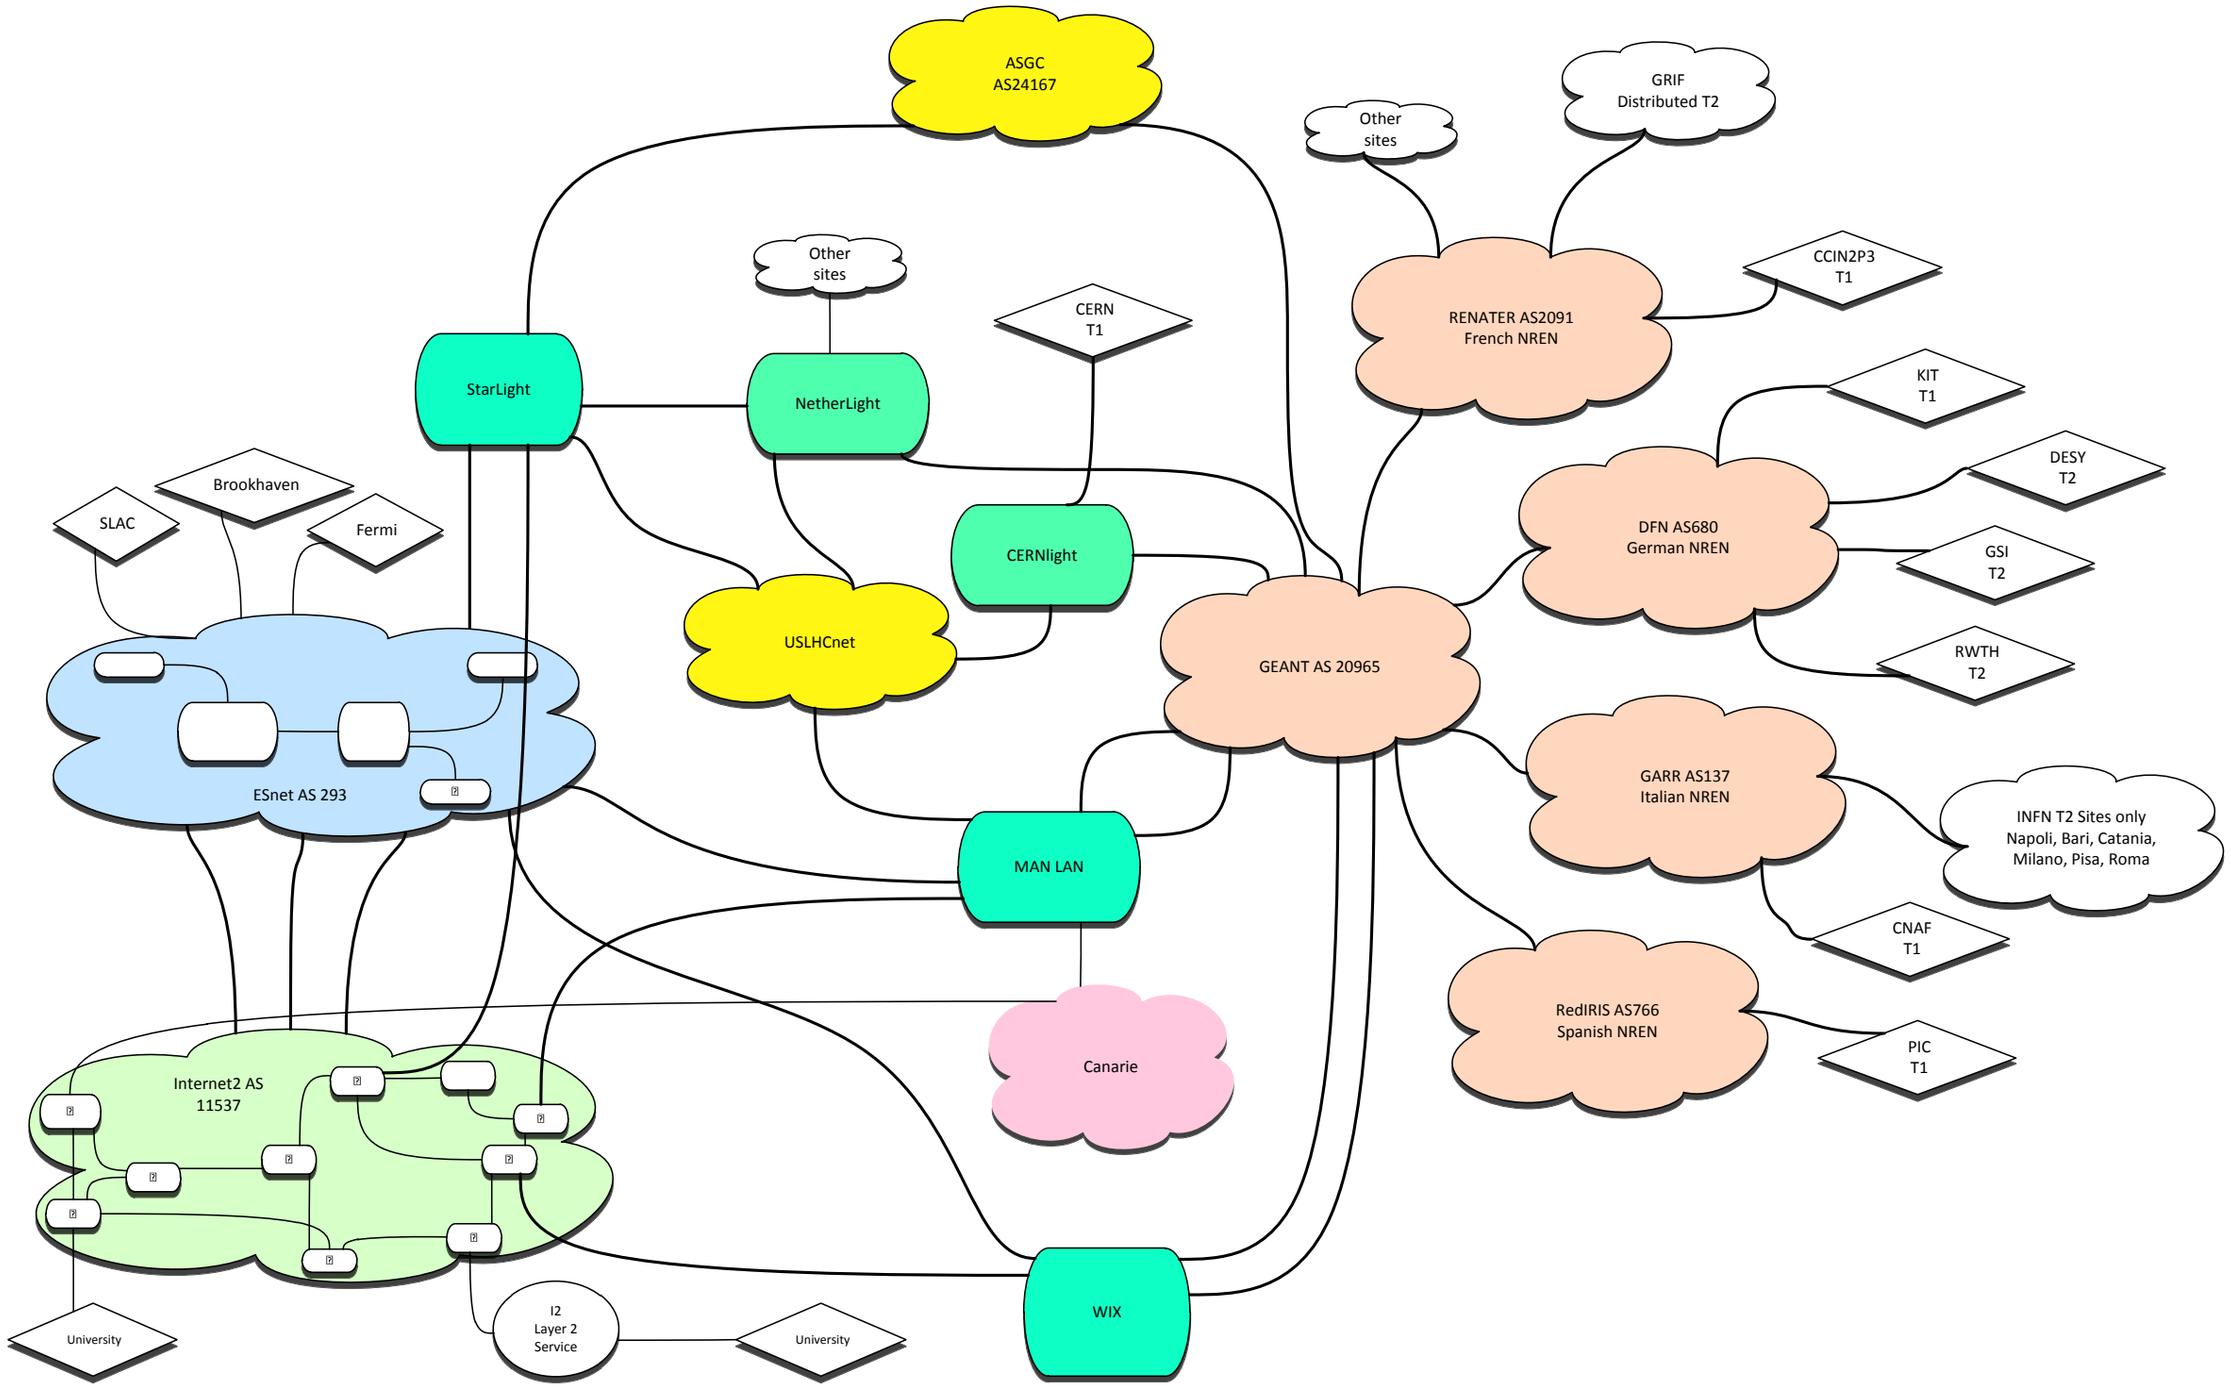

LHCONE Routing Policies

- Asymmetric routing should be avoided as this will cause problems for traffic passing (public) firewalls
- Only the networks which are announced to LHCONE are allowed to reach the LHCONE
- Only these networks will be reachable via the LHCONE
- Other traffic has to use the public networks

LHCONE - the current status

- Currently ~100 network prefixes
  - German sites currently participating at LHCONE
    - KIT
    - DESY
    - GSI
    - RWTH Aachen
    - Uni Wuppertal
    - DFN connected via 2 x 10 Gbps to GEANT
  - International:
    - CERN
    - BNL
    - SARA
    - TRIUMF
    - ....
- Detailed monitoring via perfSONAR

LHCONE Monitoring

d (old)
t)
d (old,test)
ation

LHCONE Throughput Matrix

---	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
0:BNL (lhcmon.bnl.gov)	--- ---	1.39 1.67	1.79 1.97	0.06 0.23	0.00 0.00	0.38 0.58	0.06 0.07	0.00 0.13	0.14 0.15	0.17 0.10	0.02 0.03	1.32 1.64	0.03 0.06	0.00 0.00	0.42 0.50	0.11 0.13	0.29 0.39	0.04 0.05
1:AGLT2 (psmsu02.aglt2.org)	1.33 1.32	--- ---	2.96 3.20	0.10 0.10	0.28 0.47	0.26 0.20	0.05 0.05	0.00 0.01	0.14 0.13	0.12 0.10	0.01 0.01	1.91 2.61	0.23 0.18	0.00 0.00	0.00 0.00	0.00 0.03	0.35 0.34	0.02 0.03
2:AGLT2 (psum02.aglt2.org)	1.77 1.05	3.61 3.22	--- ---	0.07 0.07	0.35 0.56	0.39 0.45	0.05 0.05	0.00 0.00	0.08 0.14	0.13 0.13	0.02 0.02	2.55 1.87	0.24 0.24	0.00 0.00	0.00 0.00	0.00 0.00	0.40 0.47	0.03 0.03
3:ASGC (lhc-bandwidth.twgrid.org)	0.33 0.15	0.00 0.32	0.28 0.13	--- ---	0.38 0.12	0.22 0.24	0.00 0.03	0.00 0.14	0.10 0.07	0.00 0.03	0.00 0.01	0.21 0.08	0.00 0.00	0.00 0.00	0.00 0.00	0.15 0.15	0.70 0.54	0.35 0.35
4:CERN (perfsonar-ps.cern.ch)	0.00 0.00	0.47 0.59	0.40 0.43	0.01 0.00	--- ---	1.88 1.80	0.63 0.00	0.00 0.79	0.78 0.80	0.38 0.46	0.11 0.00	0.00 0.12	0.00 0.49	0.00 0.00	0.87 0.92	0.36 0.42	0.27 0.14	0.04 0.00
5:DESY-HH (perfsonar-ps-02.desy.de)	0.64 0.59	0.30 0.33	0.22 0.41	0.02 0.06	0.76 1.35	--- ---	0.34 0.00	0.00 0.01	0.32 0.45	0.00 0.71	0.07 0.00	0.00 0.16	0.15 0.15	0.00 0.00	0.00 0.00	0.00 0.01	0.21 0.24	0.01 0.00
6:GRIF-LAL (psonar2.lal.in2p3.fr)	0.54 0.51	0.45 0.41	0.51 0.24	0.08 0.19	0.00 0.82	0.00 0.85	--- ---	0.00 0.94	0.87 0.93	0.00 0.78	0.38 0.00	0.00 0.47	0.00 0.89	0.00 0.00	0.00 0.00	0.01 0.00	0.25 0.24	0.17 0.15
7:GRIF/LPNHE (lpnhe-psb.in2p3.fr)	0.30 0.00	0.28 0.16	0.28 0.00	0.08 0.00	0.45 0.00	0.52 0.51	0.94 0.00	--- ---	0.61 0.52	0.45 0.00	0.37 0.00	0.46 0.14	0.46 0.00	0.00 0.00	0.00 0.00	0.01 0.00	0.22 0.25	0.01 0.00
8:INFN Napoli (perfsonar.na.infn.it)	0.47 0.49	0.39 0.39	0.38 0.36	0.06 0.13	0.73 0.82	0.87 0.74	0.78 0.67	0.89 0.73	--- ---	0.92 0.92	0.48 0.06	0.00 0.40	0.88 0.87	0.00 0.00	0.00 0.00	0.00 0.00	0.25 0.26	0.00 0.29
9:KIT (perfsonar-de-kit.gridka.de)	0.35 0.47	0.37 0.44	0.10 0.36	0.01 0.16	0.93 0.75	0.84 0.00	0.87 0.00	0.00 0.42	0.92 0.93	--- ---	0.58 0.00	0.00 0.43	0.00 0.81	0.00 0.00	0.93 0.89	0.32 0.30	0.23 0.25	0.22 0.00
10:LRZ-LMU (lcg-lrz-perfs2.grid.lrz.de)	0.44 0.35	0.24 0.00	0.38 0.24	0.06 0.21	0.00 0.77	0.00 0.46	0.00 0.38	0.00 0.44	0.69 0.61	0.00 0.79	--- ---	0.00 0.33	0.00 0.65	0.00 0.00	0.00 0.00	0.00 0.00	0.14 0.17	0.07 0.09

R&D Network Trends

- Increased multiplicity of 10Gbps links in the Major R&E networks: GEANT, Internet2, Esnet, European NREN, ...
- 100G Backbones in place and transition now underway
  - GEANT, DFN, ...
  - CERN - Budapest 2 X 100G for LHC Remote Tier- 0 Center
- OpenFlow (Software-defined switching and routing) taken up by much of the network industry and R&E networks

Bandwidth Evolution @ DFN

- DFN is upgrading the optical platform of the X-WiN
  - Contract awarded to ECI Telecom (http://www.ecitele.com)
  - Migration work is currently underway
- High Bandwidth Capabilities
  - 88 wave length per fiber
  - Up to 100 Gbps per wave length
    - thus 8.8 Tbps per fiber!
  - 1 Tbps Switching Fabric (aggregation of 10 Gbps lines on single 100 Gbps line)
- Significant cheaper components for 1 Gbps and 10 Gbps components -> reduced cost for VPN connections, new DFN conditions to be announced tomorrow

Kapazitätsentwicklung des WiN

Übertragungskapazität der Kernnetzverbindungen des Wissenschaftsnetzes:

X-WiN 2012: 8.800 Gbit/s

X-WiN 2006: 400 Gbit/s

G-WiN: 10 Gbit/s

B-WiN: 622 Mbit/s

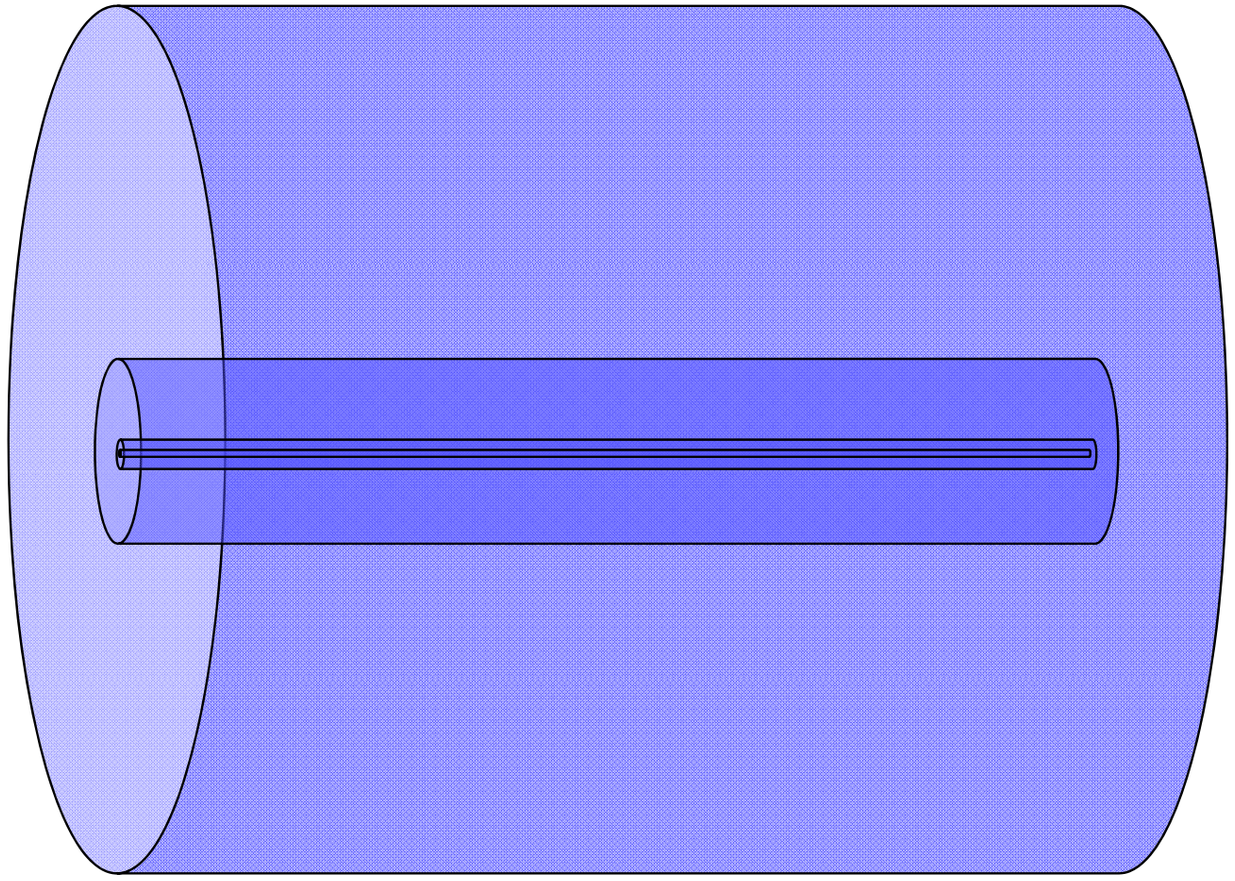

Summary

- The LHC computing and data models continue to evolve towards more dynamic, less structured, on-demand data movement thus requiring different network structures
- LHCOPN and LHCONE may merge in the future
- With the evolution of the new optical platform VPNs will get more affordable