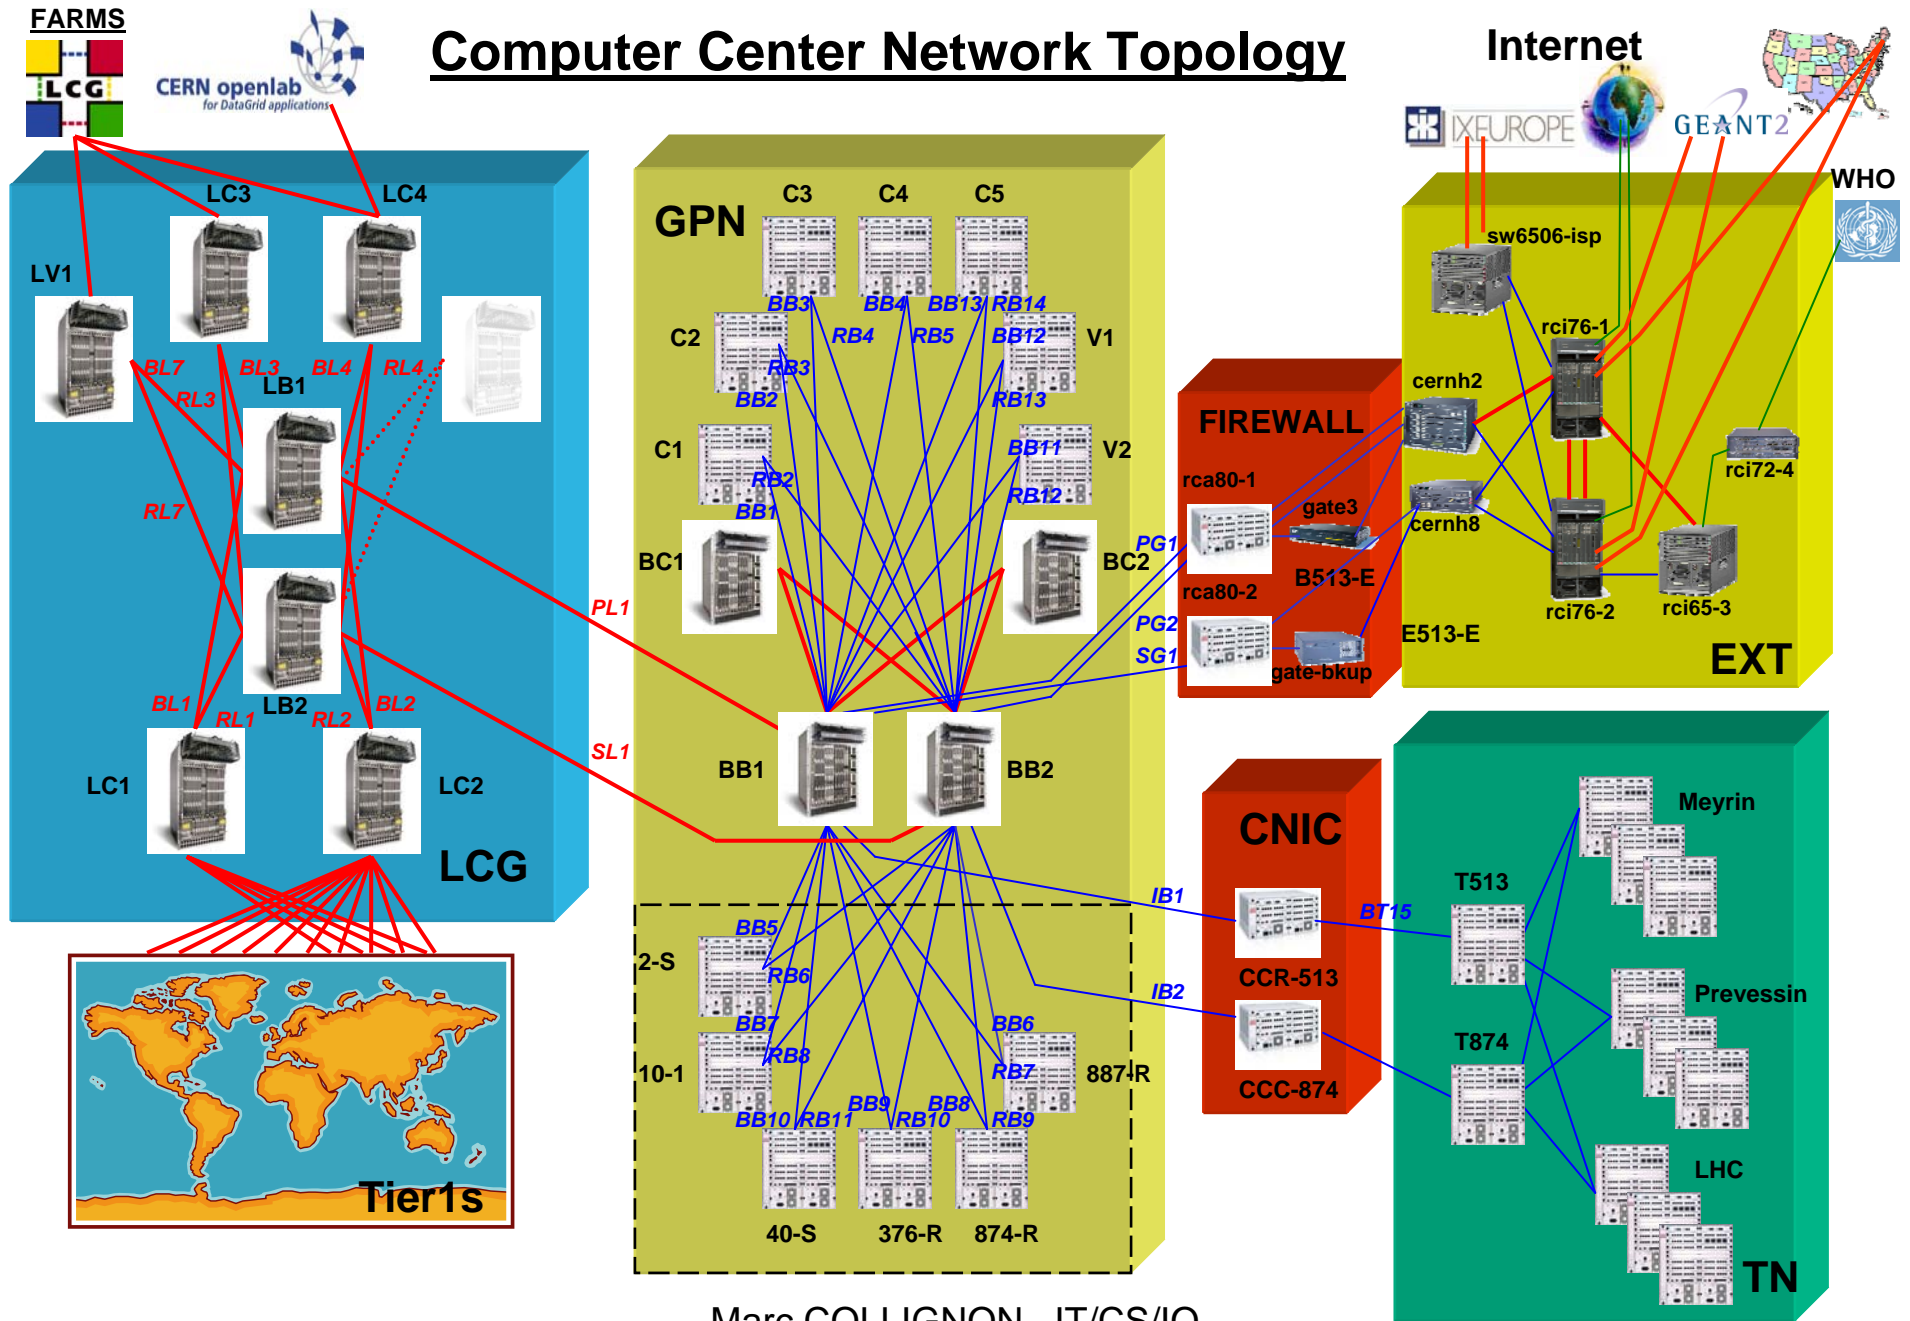


# WLCG Meeting

## LHCOPN

David Foster  
Communications Systems Group Leader  
CERN  
IT-CS

# Computer Center Network Topology



Marc COLLIGNON IT/CS/IO

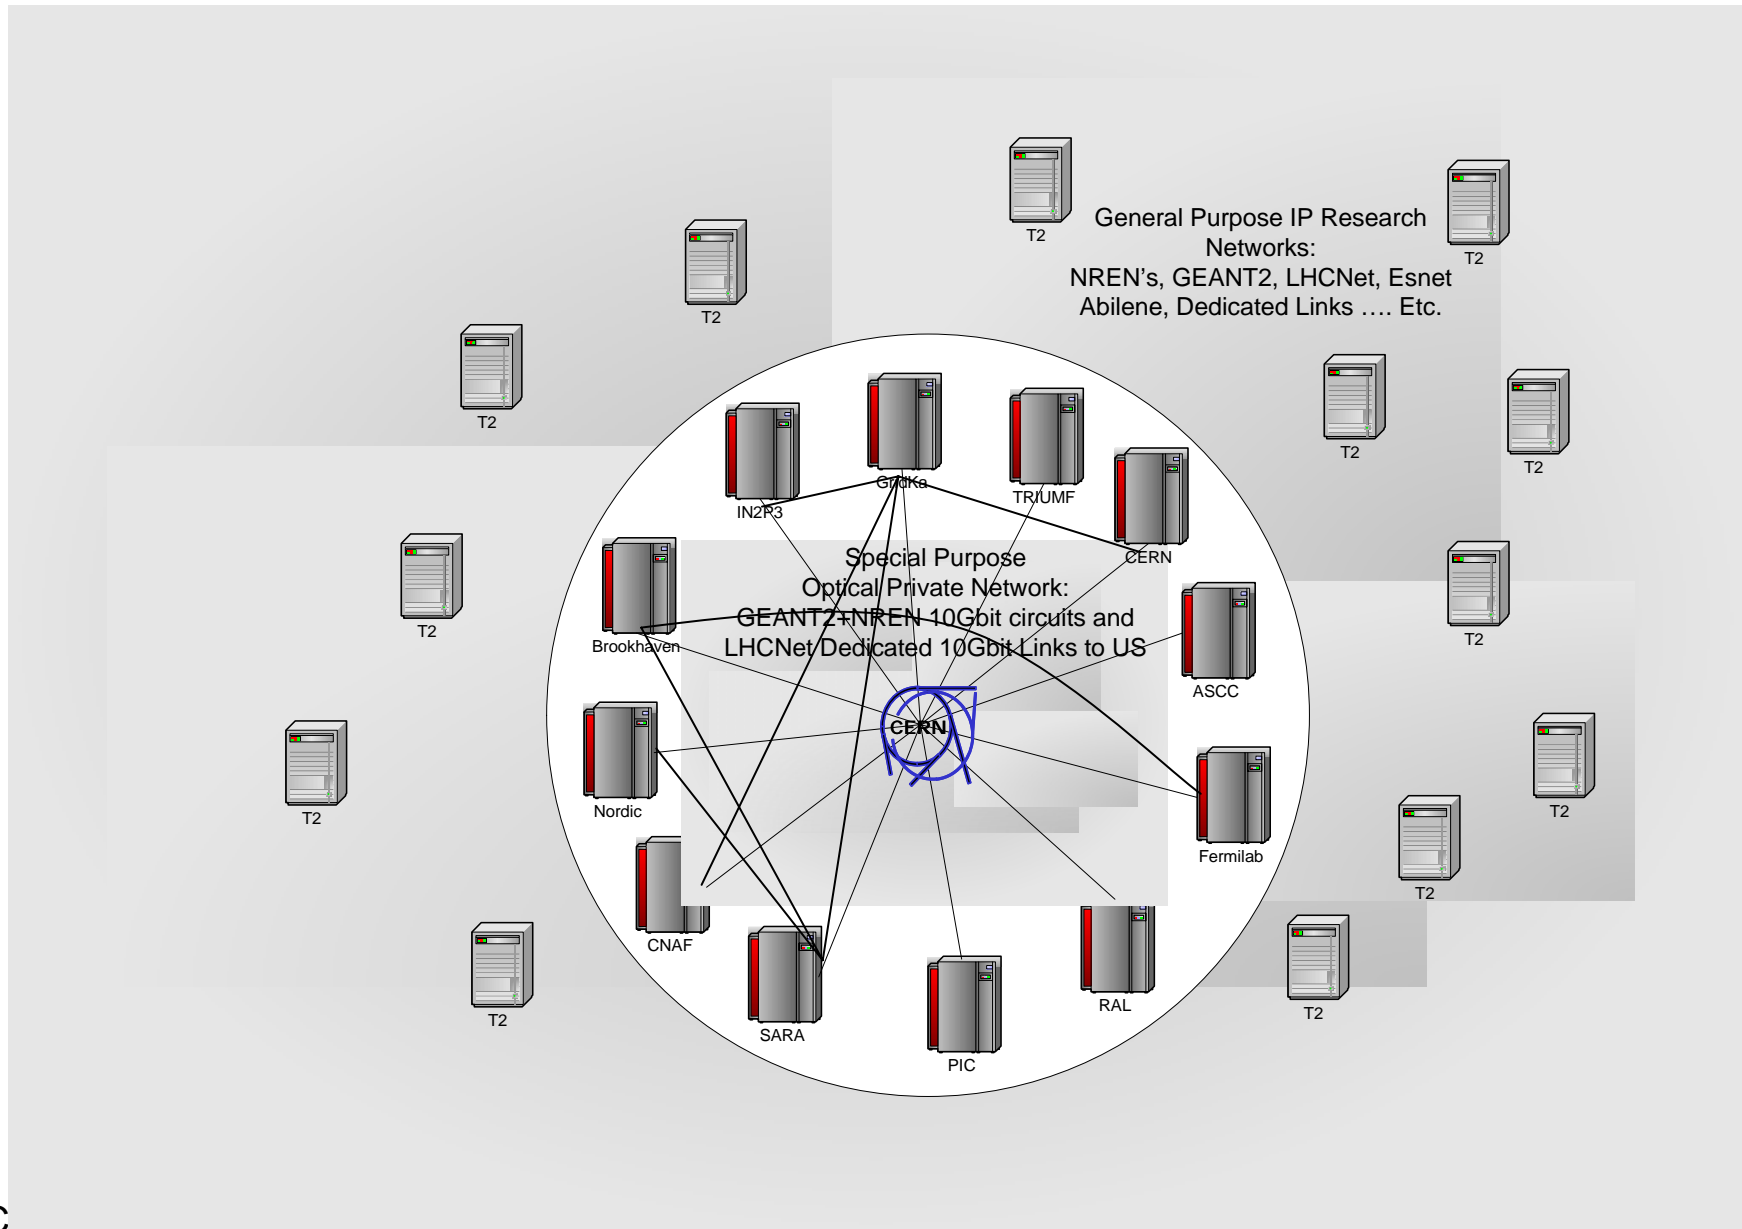
# Context

- Wide Area Networking used for all the aspects of the LHC is composed of many different, and separately managed, infrastructures
  - National Research Networks (NRENS) Worldwide (NRENS, I2, ESNet, CANARIE etc.)
  - Interconnection of NRENS in Europe (GEANT-2)
  - Transatlantic Connectivity (USLHCNet)
  - LHC Optical Private Network (LHCOPN)
- No centralised funding, no centralised management.
  - Independent domains of infrastructure and responsibility.

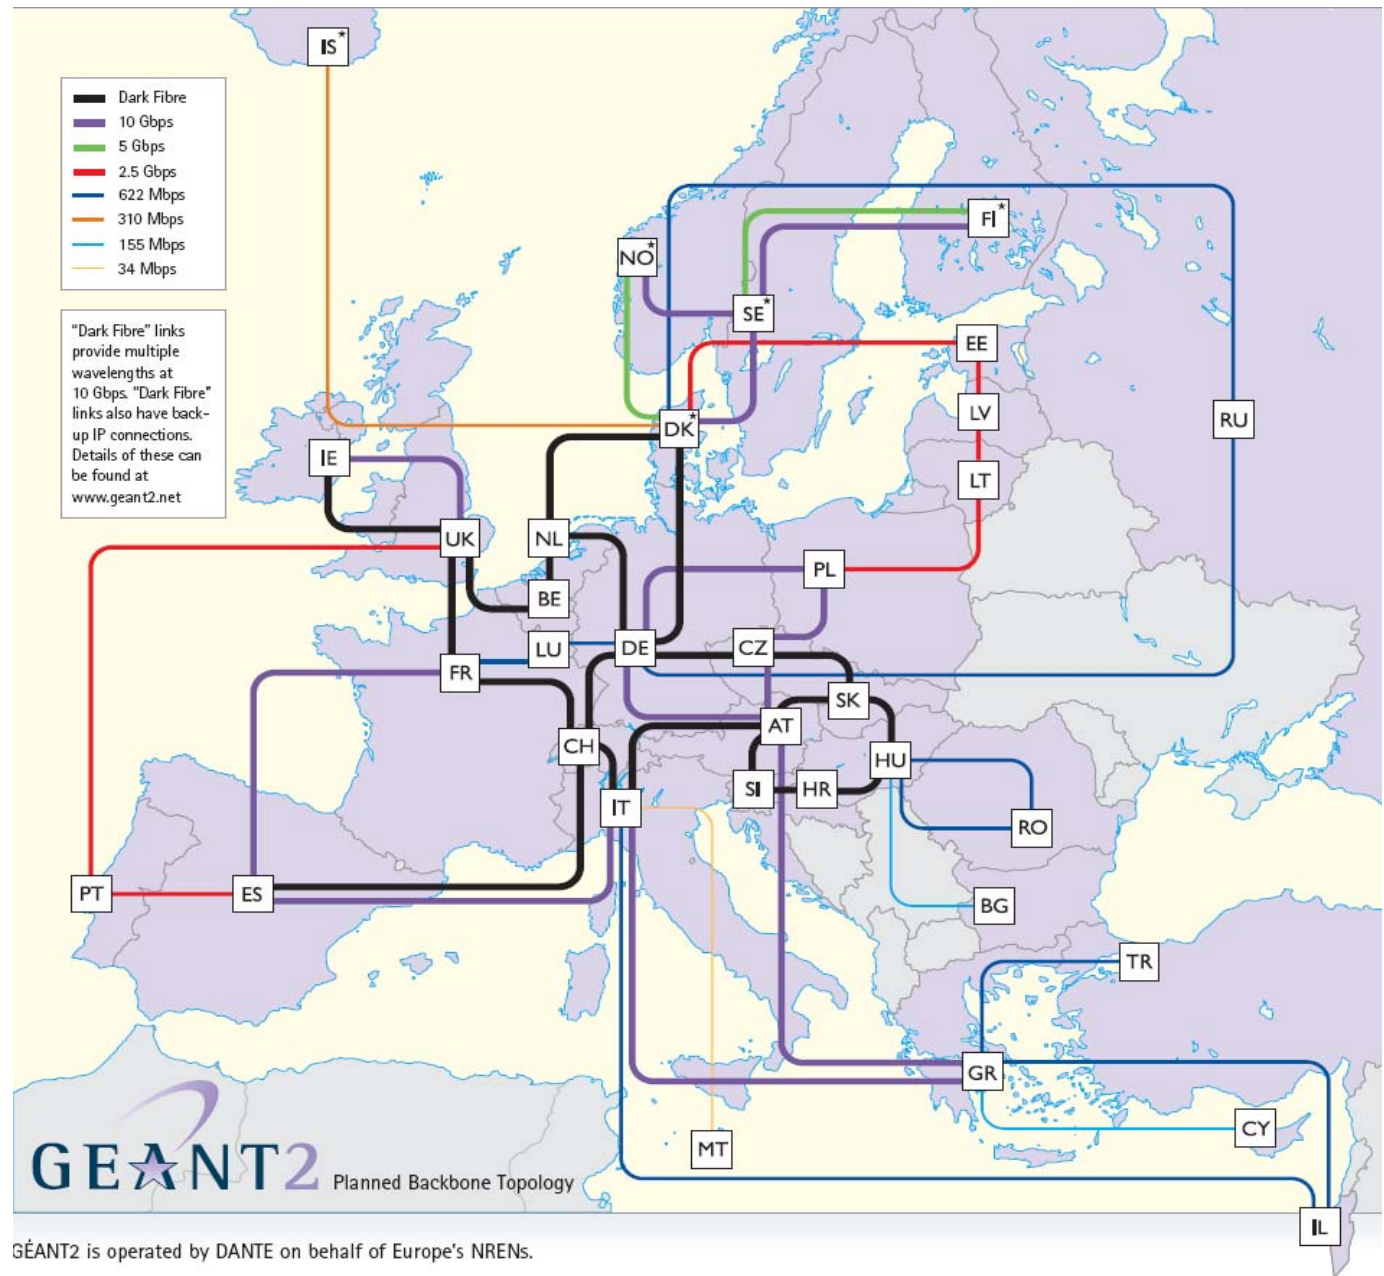
# LHCOPN Mission

- Started in 2004 to address at least one of the identifiable problems: Getting data from CERN (T0) to the T1's with a predictable performance.
- GEANT-2 Infrastructure was evolved with the LHCOPN requirements in mind.
  - Was important to have this vision in 2004 to have an infrastructure in 2007!
  - LHCOPN was identified as a “considerable achievement” at the last EU review of GEANT
- 10Gbit circuit was/is the unit level of connectivity that matches requirements and motivates development.

# LHCOPN Architecture (2004-2006)



- GEANT2 is an EU project.
- Funding comes from NREN's and the EU
- DANTE is the implementation organisation.
- Policy is decided at the NREN Policy Committee meeting (NREN PC)
- Subgroups work on different aspects. E.g. Global Connectivity Committee (GCC)

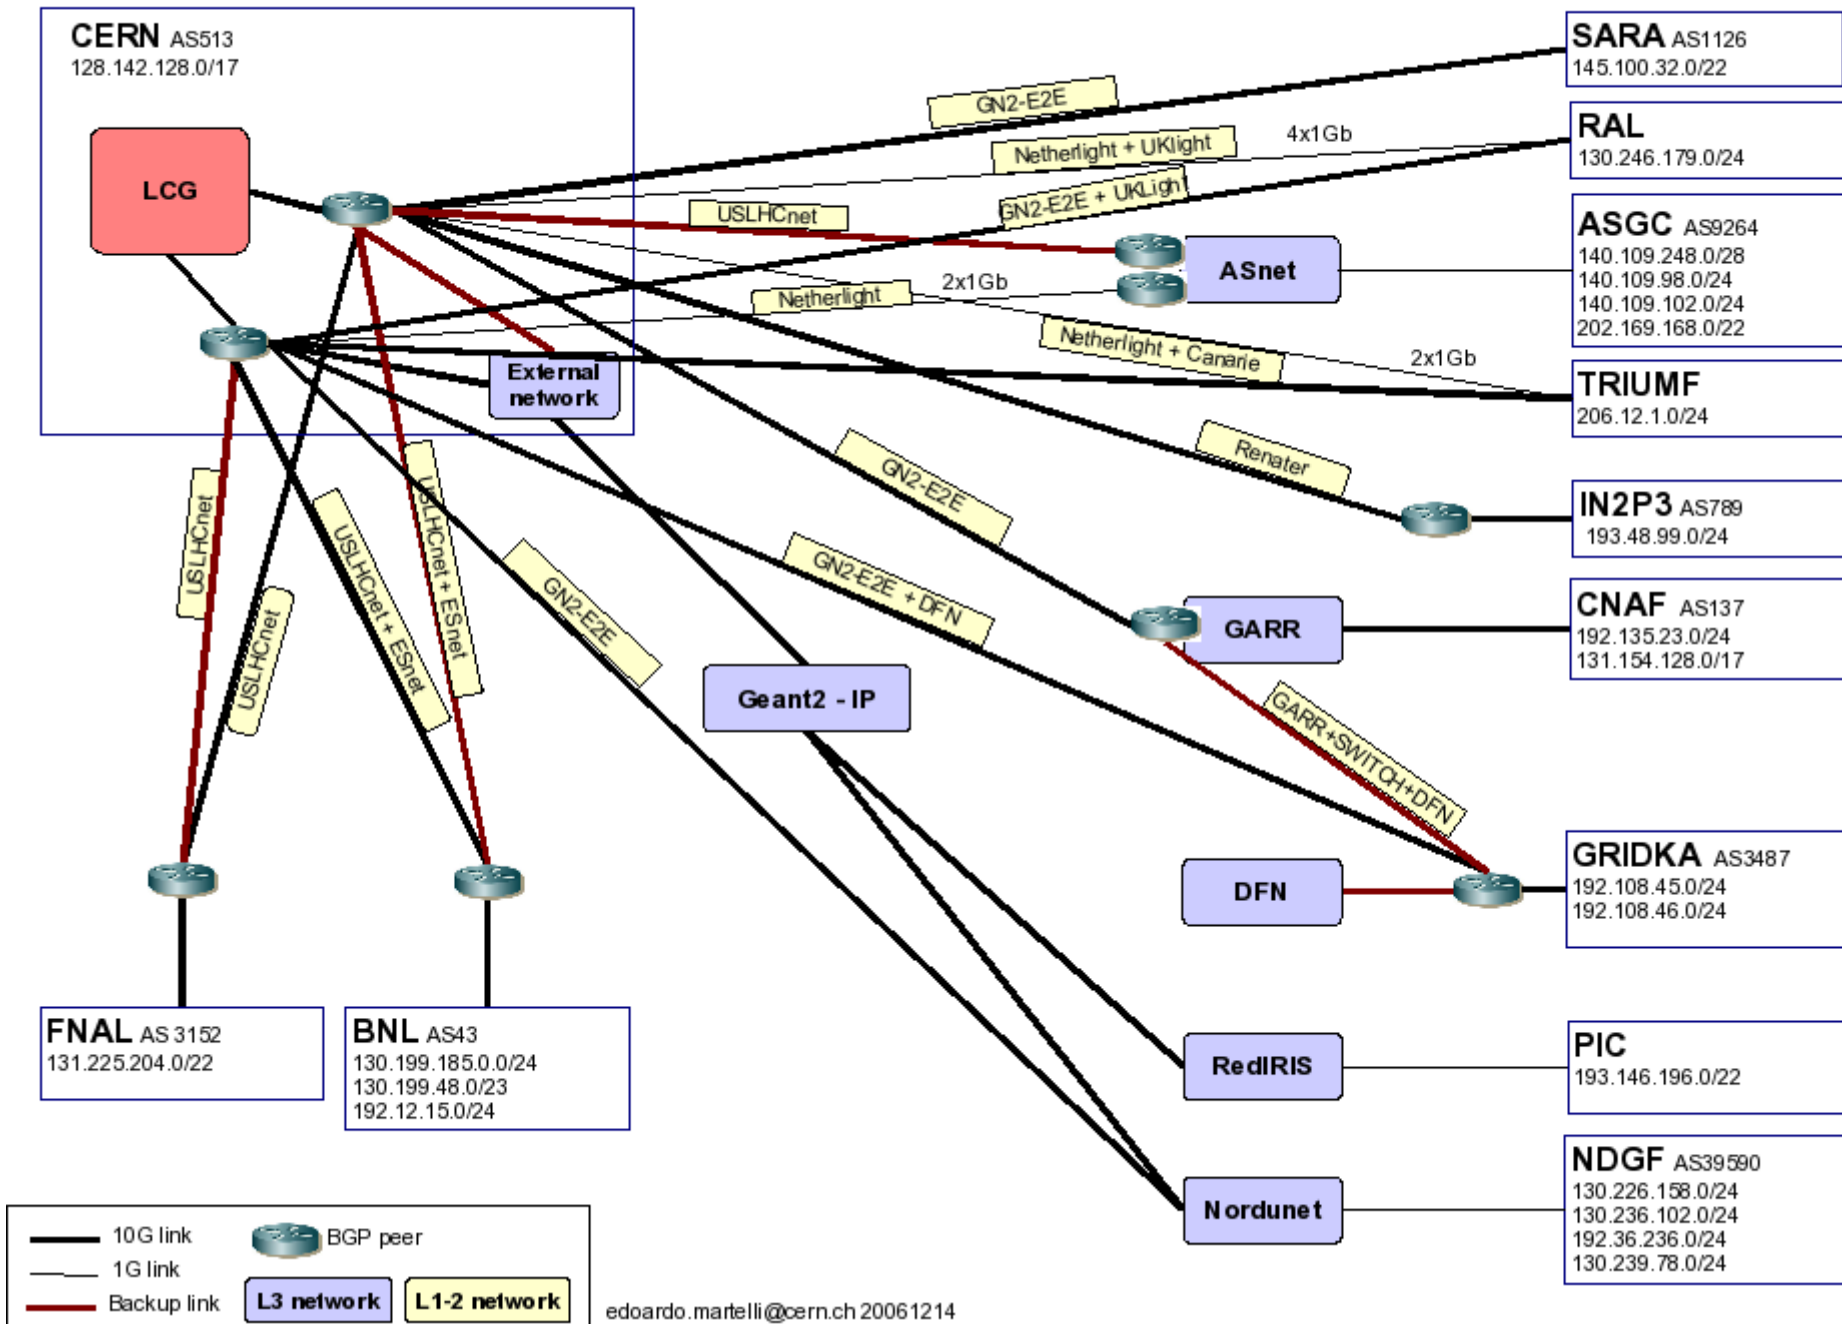


GEANT2 is operated by DANTE on behalf of Europe's NRENs.

AT Austria	CZ Czech Republic	ES Spain	HR Croatia	IS Iceland*	LV Latvia	PL Poland	SE Sweden*
BE Belgium	DE Germany	FI Finland*	HU Hungary	IT Italy	MT Malta	PT Portugal	SI Slovenia
BG Bulgaria	DK Denmark*	FR France	IE Ireland	LT Lithuania	NL Netherlands	RO Romania	SK Slovakia
CH Switzerland	EE Estonia	GR Greece	IL Israel	LU Luxembourg	NO Norway*	RU Russia	TR Turkey
CY Cyprus							UK United Kingdom

\*Connections between these countries are part of NORDUnet (the Nordic regional network)

# LHCOPN – current status

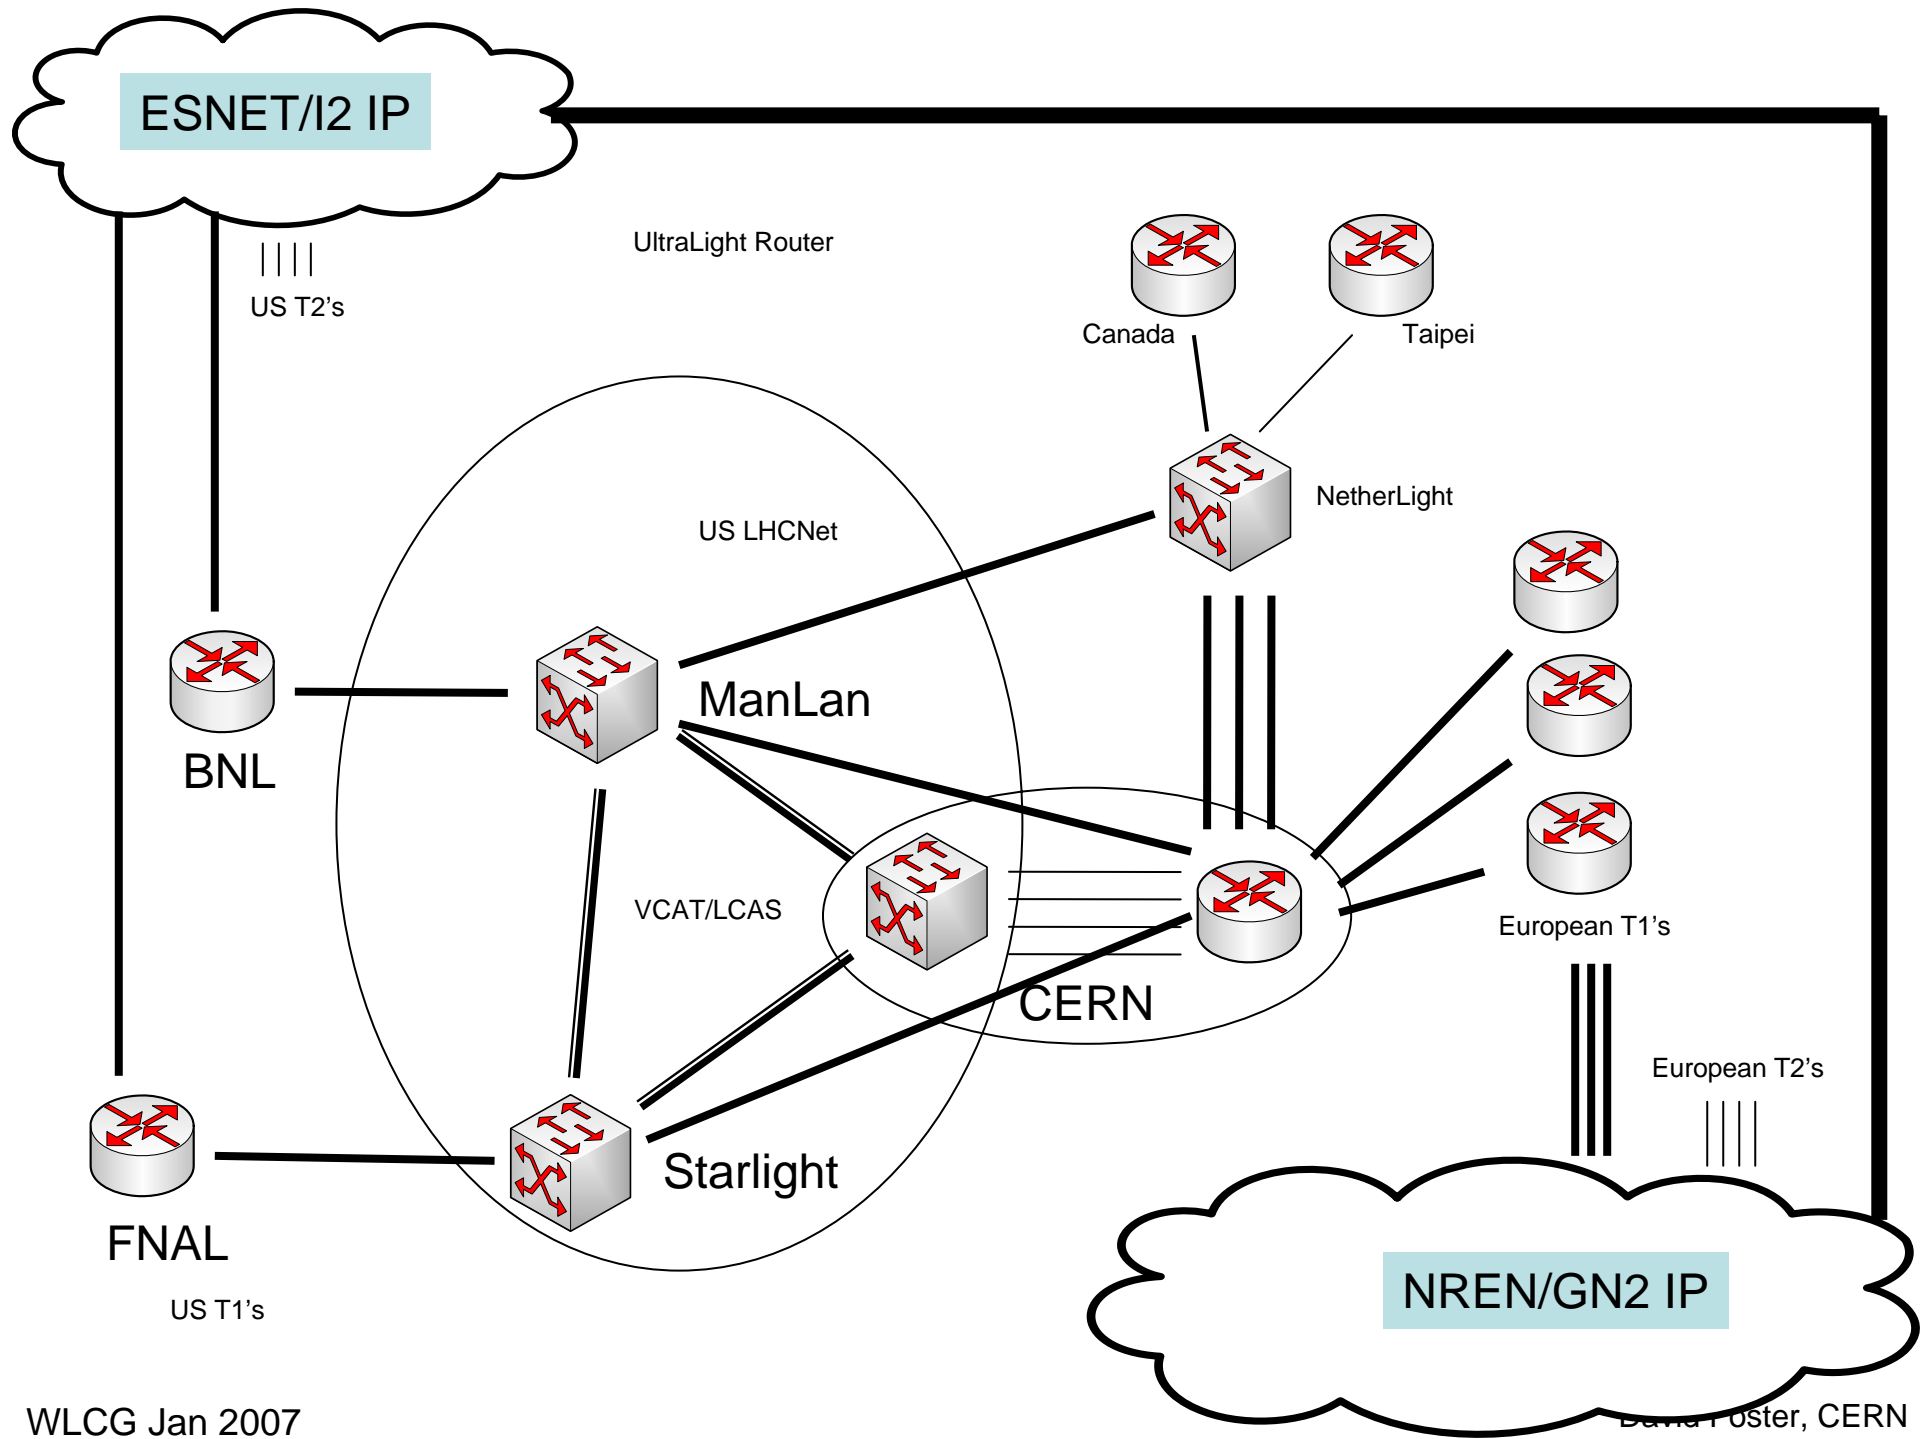


# US LHC Network Working Group Mission Statements

- To support the LHC Physics program by continuing to provide US and Transatlantic networks with the capacity and capabilities required for the experiments to take full advantage of the LHC's unique potential for physics discoveries
  - To provide this capability in a manner compatible with, and generally beneficial to the needs of other major programs in high energy physics, as well as other fields of science supported by the funding agencies
- To develop a worldwide partnership among the major mission-oriented and research and education networks, in the US, Europe, Asia-Pacific, Latin America and across the Atlantic and Pacific, as well as the HEP laboratories and other Tier1 and Tier2 sites, to ensure compatible network operations fulfilling the needs of all sectors of the LHC Collaborations
- To cooperatively develop an operations and management paradigm, network provisioning and management methods, and associated software systems, to make the full capabilities of the networks provided available to the LHC community, and to other sectors of HEP and other scientific communities as appropriate

Harvey Newman October 2006



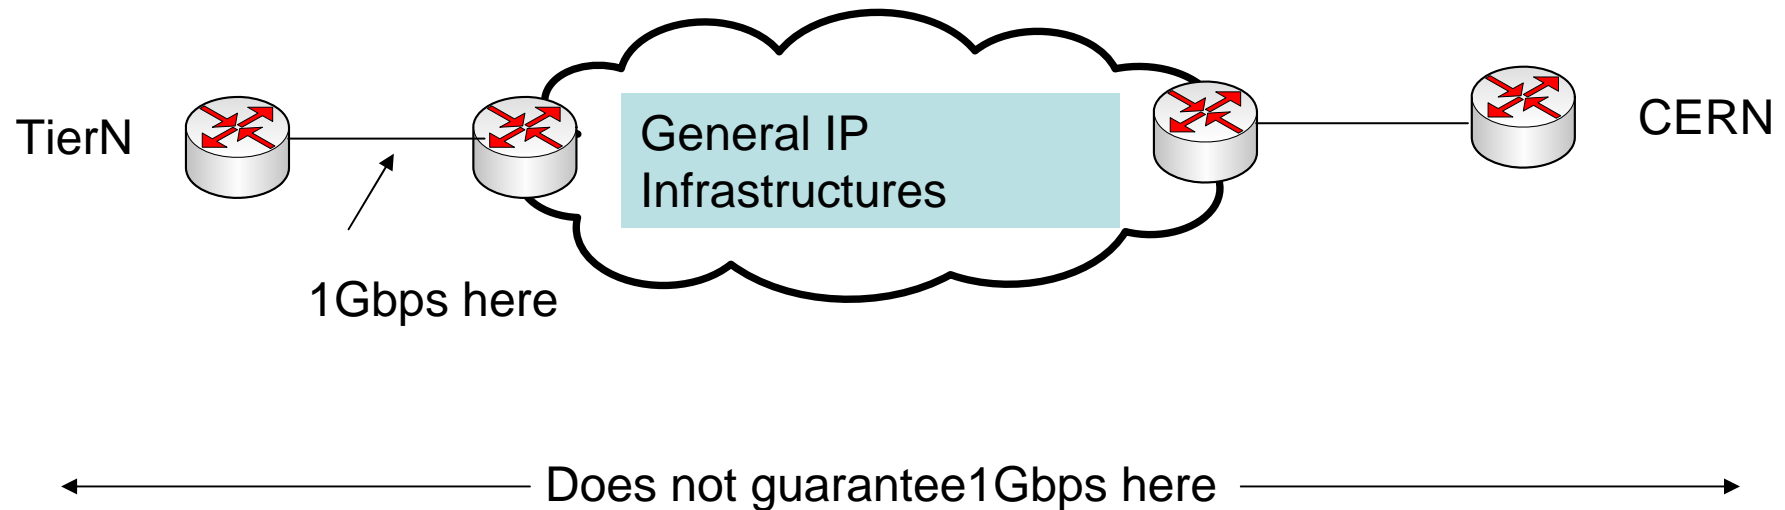


# Summary 2007-2008

- According to what is known, and what has been tested, the starting situation is:
  - LHCOPN will support the T0-T1 connectivity requirements.
  - LHCOPN will be able to support a (large) fraction of T1-T1 requirements.
  - T1-T2 connectivity will be provided by the general purpose IP infrastructures.
    - Some 25Gb/sec of little used IP peering (ESnet/I2 to GEANT) is available.

# Misunderstandings

- “We have 1G to CERN”



This is a specific example, but applies to any connection between centers.

The above statement is only true when a dedicated circuit has been provisioned

# Mega Words of Caution

- The LCG “Megatable” activity aims to provide a (useful) “bottom up” view of the network requirements.
- Some figures appear to be peak (T0/T1/T1) and some average (T1/T2).
- All figures probably have been generated from simple models of data movement for the “standard” data movement cases.
- Conclusion: Whilst it gives “order of magnitude” requirements for the most basic network needs, network provisioning (which takes a long time) needs to work to a model based more on future predicted behaviour, capability and availability.

# Megatable OPN Rates

	T0-T1 (MB/sec)	T1-T1 In (MB/sec)	Total In Gb/sec	T1-T1 Out (MB/sec)	Total Out Gb/sec
ASGC	91.3	158.3	<b>2.00</b>	128.8	<b>1.03</b>
BNL	287.2	274.1	<b>4.49</b>	218.5	<b>1.75</b>
CERN	1343.0	208.7	<b>1.67</b>	104.3	<b>11.58</b>
CNAF	136.2	208.0	<b>2.75</b>	209.4	<b>1.68</b>
FNAL	105.0	63.4	<b>1.35</b>	214.9	<b>1.72</b>
FZK	132.6	220.1	<b>2.82</b>	193.6	<b>1.55</b>
IN2P3	157.2	229.5	<b>3.09</b>	263.7	<b>2.11</b>
NDGF	54.4	51.9	<b>0.85</b>	54.2	<b>0.43</b>
NIKHEF	121.7	134.3	<b>2.05</b>	172.2	<b>1.38</b>
PIC	63.7	167.5	<b>1.85</b>	88.2	<b>0.71</b>
RAL	137.2	218.3	<b>2.84</b>	219.4	<b>1.76</b>
TRIUMF	48.3	50.1	<b>0.79</b>	47.4	<b>0.38</b>
ALICE US T1	8.2	10.8	<b>0.15</b>	3.7	<b>0.03</b>

# Megatable GP IP Rates

	T2-T1 In (MB/sec)	Total In Gb/sec	T1-T2 Out (MB/sec)	Total Out Gb/sec
ASGC	54.8	<b>0.44</b>	133.0	<b>1.06</b>
BNL	92.4	<b>0.74</b>	225.7	<b>1.81</b>
CERN	49.4	<b>0.40</b>	71.4	<b>0.57</b>
CNAF	68.1	<b>0.54</b>	155.7	<b>1.25</b>
FNAL	30.0	<b>0.24</b>	248.0	<b>1.98</b>
FZK	85.4	<b>0.68</b>	191.2	<b>1.53</b>
IN2P3	179.0	<b>1.43</b>	215.8	<b>1.73</b>
NDGF	3.9	<b>0.03</b>	14.8	<b>0.12</b>
NIKHEF	41.0	<b>0.33</b>	69.9	<b>0.56</b>
PIC	35.6	<b>0.28</b>	91.6	<b>0.73</b>
RAL	94.9	<b>0.76</b>	113.1	<b>0.90</b>
TRIUMF	14.3	<b>0.11</b>	35.6	<b>0.28</b>
ALICE US T1	32.5	<b>0.26</b>	13.2	<b>0.11</b>

# Continuing Evolution

- Provisioning additional bandwidth intra-europe or intra-us should remain cost effective.
- Provisioning additional bandwidth transatlantic will remain relatively costly.
- GEANT Cost Sharing and AUP policies require caution
  - LHCOPN services will remain very constrained for the moment, but this is compatible with the stated mission and technical implementation issues.
- Additional circuits between centers T2/T1/T1 will need provisioning as needs arise.

# OPN Organisation

“Even simple things are hard”

- LHCOPN Meetings 4 times a year
  - Organised as a sub-activity of the GDB
- Current Working groups
  - Operations (Dante)
  - Monitoring (USA)
  - Routing (CERN)
  - Security (UK)
- Working group evolution
  - Routing will become a long term technical body
  - Monitoring will evolve to become more technical network management including tool selection and deployment
  - Operations will continue to be the problem determination and resolution
  - Security will continue to be an advisory and policy body.



# OPN Status Summary

Link	Status	Nominal E2e Capacity	Provider Changes	Expected
BNL	OPN Production	10G	Colt->Colt	Done
FNAL	OPN Production	10G	GC->Qwest	Done
TRIUMF	OPN Production	2G (limited CERN-AMS)	GN2 Lambda CERN-AMS	
ASGC	OPN Production	2G (2.5G to AMS)	GN2 Lambda CERN-AMS	
NDGF	GN2 IP/OPN Production	10G (Copenhagen)	GN2 Lambda	Routing Issues. GN2 IP used for now
SARA	GN2 IP/OPN Production	10G	GN2 Lambda	Routing Issues. GN2 IP used for now
RAL	OPN Production	10G		Waiting for public AS number
FZK	OPN Production	10G		
CNAF	OPN Production	10G		
IN2P3	OPN Production	10G		
PIC	GN2 IP		GN2 lambda	Expected Q1 2007

# CBF Status Summary

Link	Status	Nominal E2e Capacity	Provider Changes	Expected
SARA - NDGF		10G		Q1/07
SARA - FZK	In Place. Unused	10G		
FZK - CNAF	In Place. Unused	10G		
FZK - CERN		10G	DFN/Switch	Mid 07
BNL - FNAL	In Place (from GC)	10G	To ESnet	Q2 07
FZK - IN2P3		10G		End 06

# Other Links Summary

Link	Status	Nominal E2e Capacity	Provider Changes	Expected
ManLan - Netherlight	Ordered	10G	GC	1/1/07
Netherlight - CERN	Surfnet to make request to GN2 Exec	10G	GN2	1/1/07

# Operational Status

- Several links “in production” but coherent operational management across organisational domains must be organised.
  - Agreement has been reached to deploy one initial monitoring tool “Perfsonar” across all domains.
  - Workshops have been held (Dante)
    - Munich 19 July 06: DANTE, DFN, REDIRIS, GARR, SURFnet, NORDUnet, RENATER, CERN (and LRZ-Munich)
    - Toronto 18-19 September 06: DANTE, I2, ESnet, TRIUMF, Canarie, FNAL, USLHCNET
- End-to-End coordination unit (E2ECU) being implemented by Dante as part of the overall NOC – Full Operation January 2007.
- ENOC providing information integration with Grid Operations (EGEE-SA2) working closely with E2ECU
- All operational information documented on the LHCOPN Twiki:  
<http://lhcopn.cern.ch>

# Issues and Activities

- Backup remains a major issue
  - Availability of circuits to make a logical backup feasible.
  - Modeling of single point failures
  - Understanding how single point physical failures affect the logical model
    - Lots of fibers actually occupy the same physical trunking e.g. both NREN and GEANT fibers.
    - A single trunk failure could lead to multiple simultaneous logical topology failures
  - Dante with the NRENs are taking the lead in this.
- Operational procedures still being refined
  - Monitoring
  - E2ECU/ENOC collaboration.
- Capacity planning
  - Real usage and experience is important
- Requirements for US-ALICE T1