LCG MB Meeting

LHCOPN
(Summary of the WLCG presentation)

David Foster
Communications Systems Group Leader
CERN
IT-CS

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)

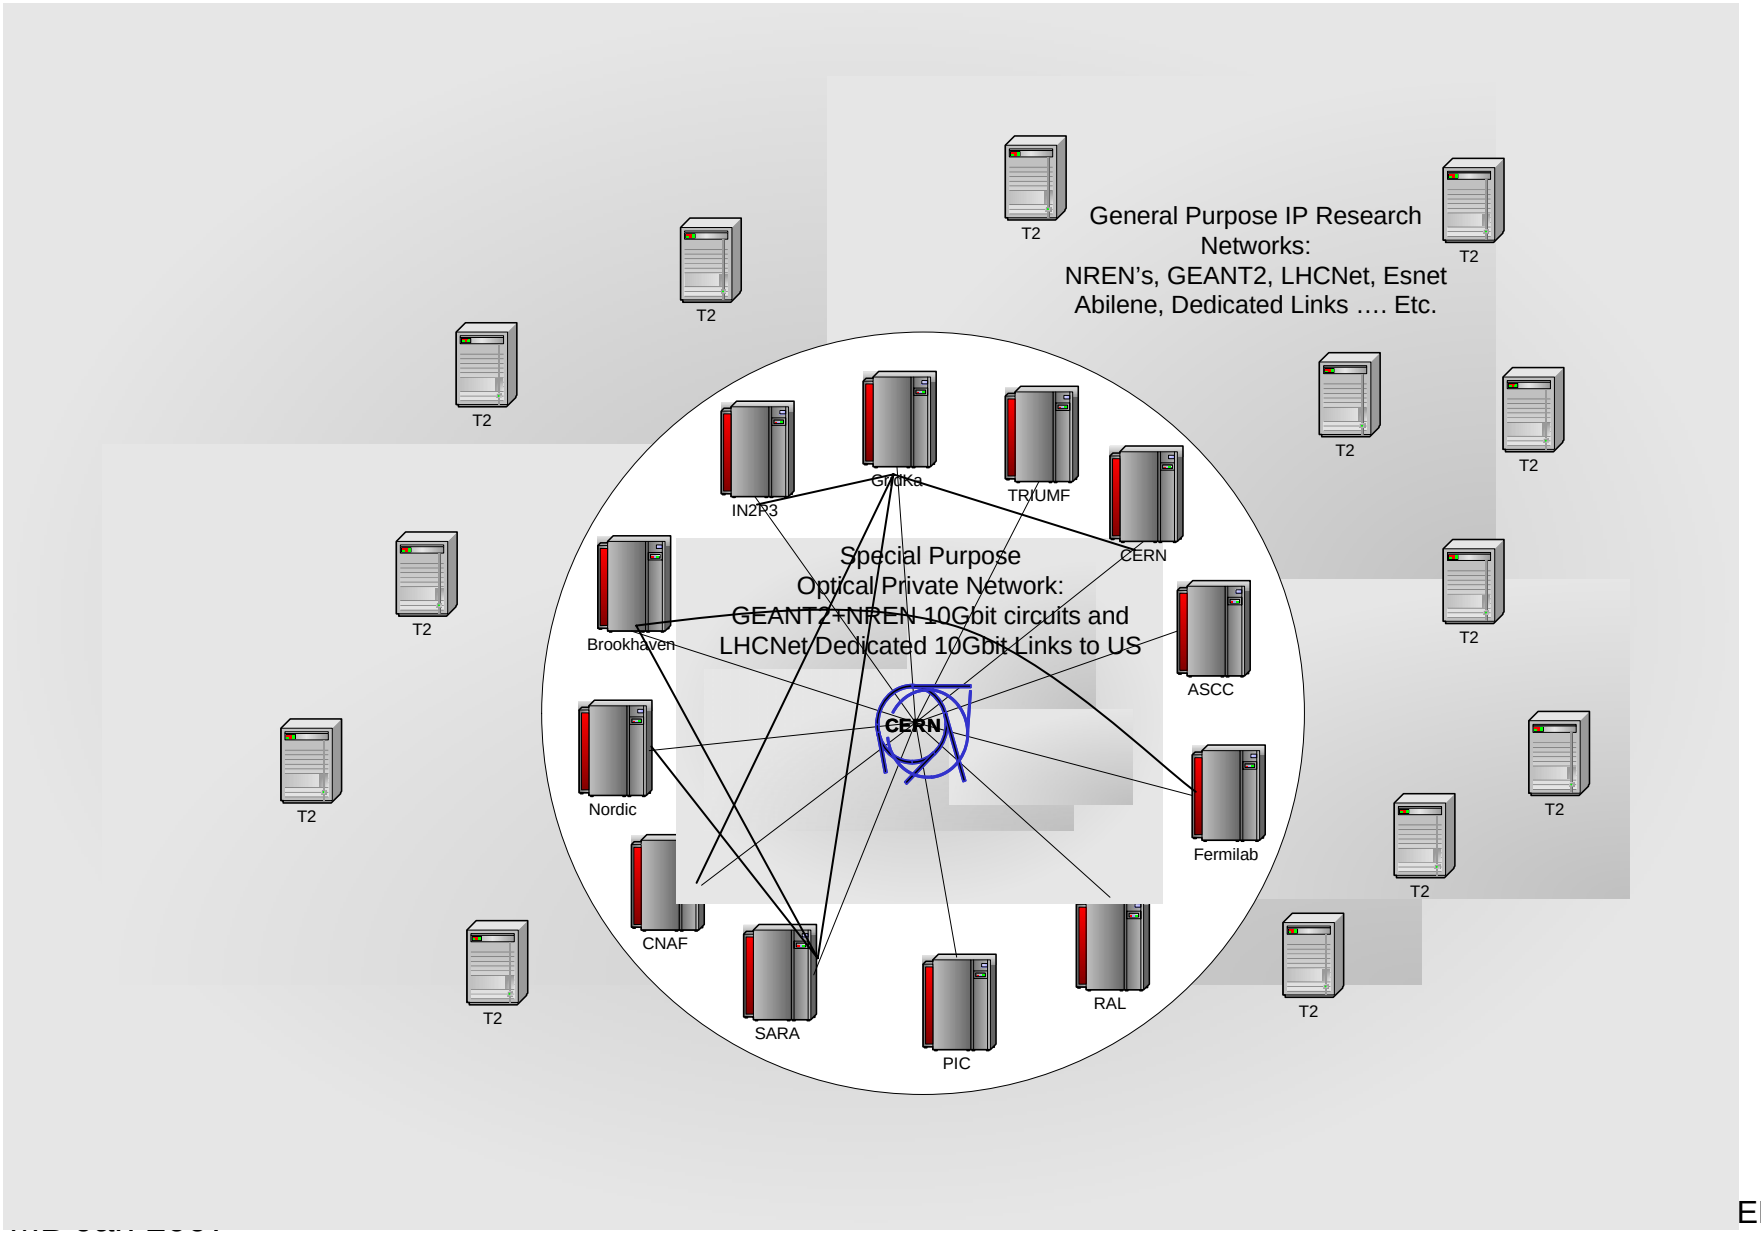

US LHC Network Working Group Mission Statements

- **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

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

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	***0.44***	133.0	***1.06***
BNL		92.4	***0.74***	225.7	***1.81***
CERN		49.4	***0.40***	71.4	***0.57***
CNAF		68.1	***0.54***	155.7	***1.25***
FNAL		30.0	***0.24***	248.0	***1.98***
FZK		85.4	***0.68***	191.2	***1.53***
IN2P3		179.0	***1.43***	215.8	***1.73***
NDGF		3.9	***0.03***	14.8	***0.12***
NIKHEF		41.0	***0.33***	69.9	***0.56***
PIC		35.6	***0.28***	91.6	***0.73***
RAL		94.9	***0.76***	113.1	***0.90***
TRIUMF		14.3	***0.11***	35.6	***0.28***
ALICE US T1		32.5	***0.26***	13.2	***0.11***

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.
- An “All Hands (Caltech, DOE, CERN, Fermilab, CMS, ESnet, I2, GEANT)” USLHCNet Meeting in October concluded:
  - Sufficient T1-T2 connectivity will be provided by the general purpose IP infrastructures.
    - Some 20Gb/sec of little used IP peering (ESnet/I2 to GEANT) is available.
    - An initial extra 5Gb/sec will be provided by the USLHCNet link from NY to AMS providing the DOE agrees.

Misunderstandings

- “We have 1G to CERN”

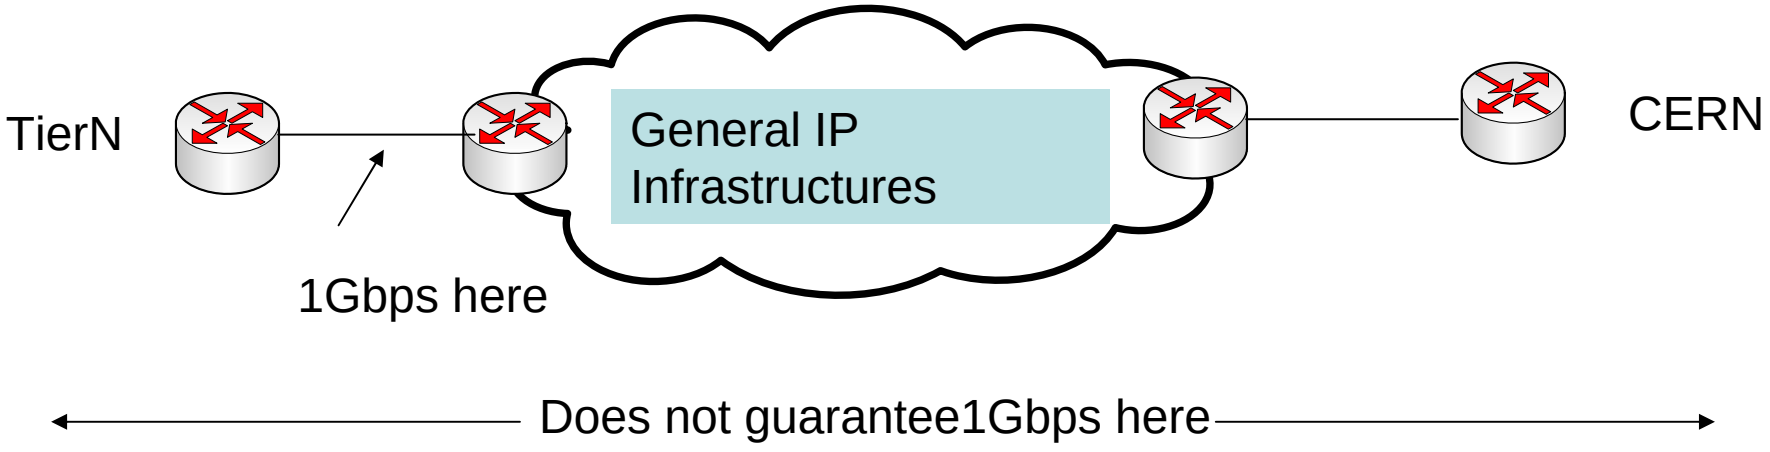

This is a specific example, but applies to any connection between centers.

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

Experiments must test their actual connectivity and decide if it is “good enough”

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 to be provisioned as needs arise.

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