



Networking Considerations For LHC

David Foster

Communication and Network Systems Group Leader

Information Technology Department

CERN – European Organization for Nuclear Research

Geneva, Switzerland

david.foster@cern.ch





What we have - December 2004

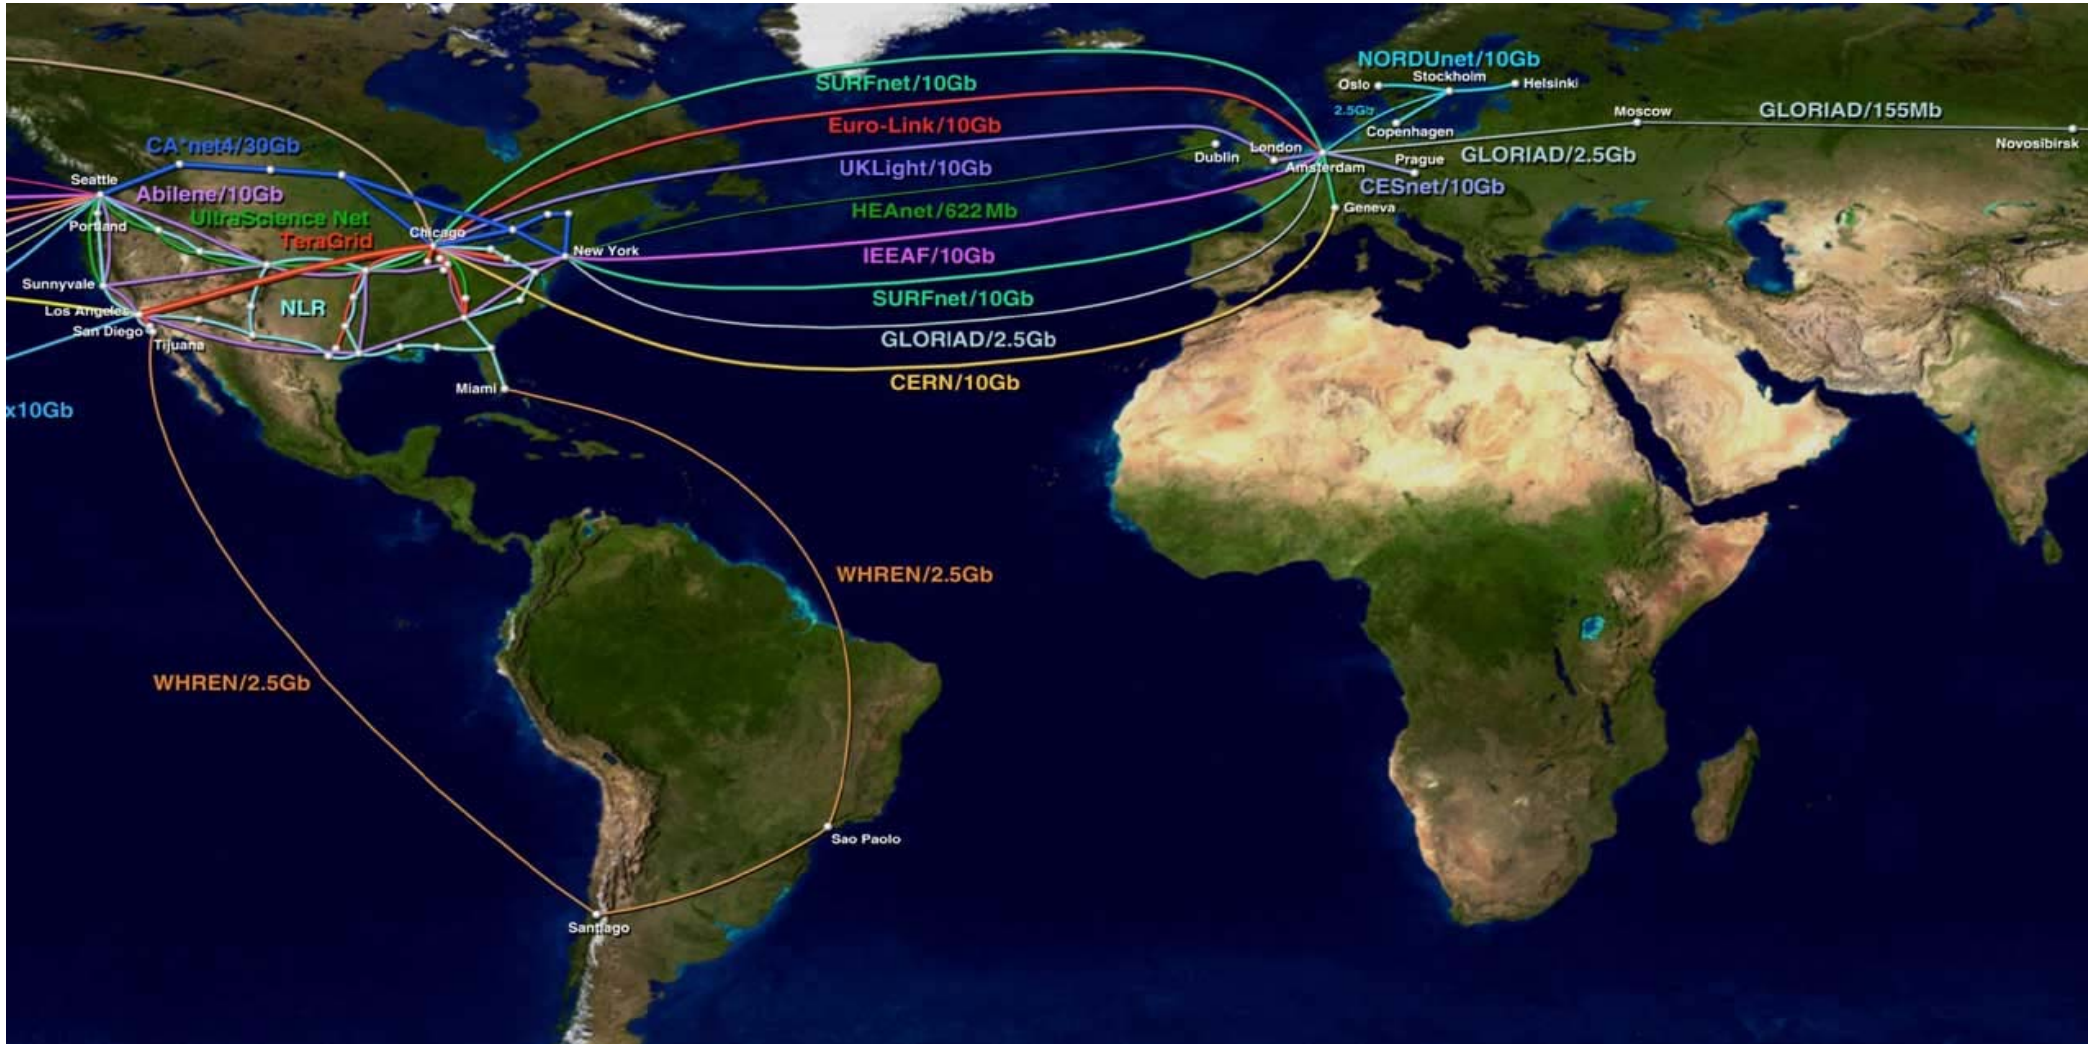
- Commodity Internet Connectivity
 - Multiple redundant connections
- Multiple 10Gb/s connections
 - To Starlight/Chicago (and from there into the US via Abilene)
 - Paid by the consortium USLiC (large contribution from the DOE).
 - Allows us to perform 10Gbit tests to Fermilab
 - To Amsterdam (and into the SurfNET NREN)
 - Provided by SurfNet
 - Divided for the most part into 8x1Gb ethernets
 - To Geant
 - Geant-2 will have a pop at CERN that will provide level-2 ethernets at 10Gb/sec
 - Today we have Level-3 routed IP network with restrictions on the achievable end-end bandwidth due to design issues with the Junipers routers in Geant-1





Global Lambda Integrated Facility

Predicted Bandwidth for Scheduled Experiments, December 2004



www.glif.is

last update 23/11/2004 12:11

Visualization courtesy of
Bob Patterson, NCSA.

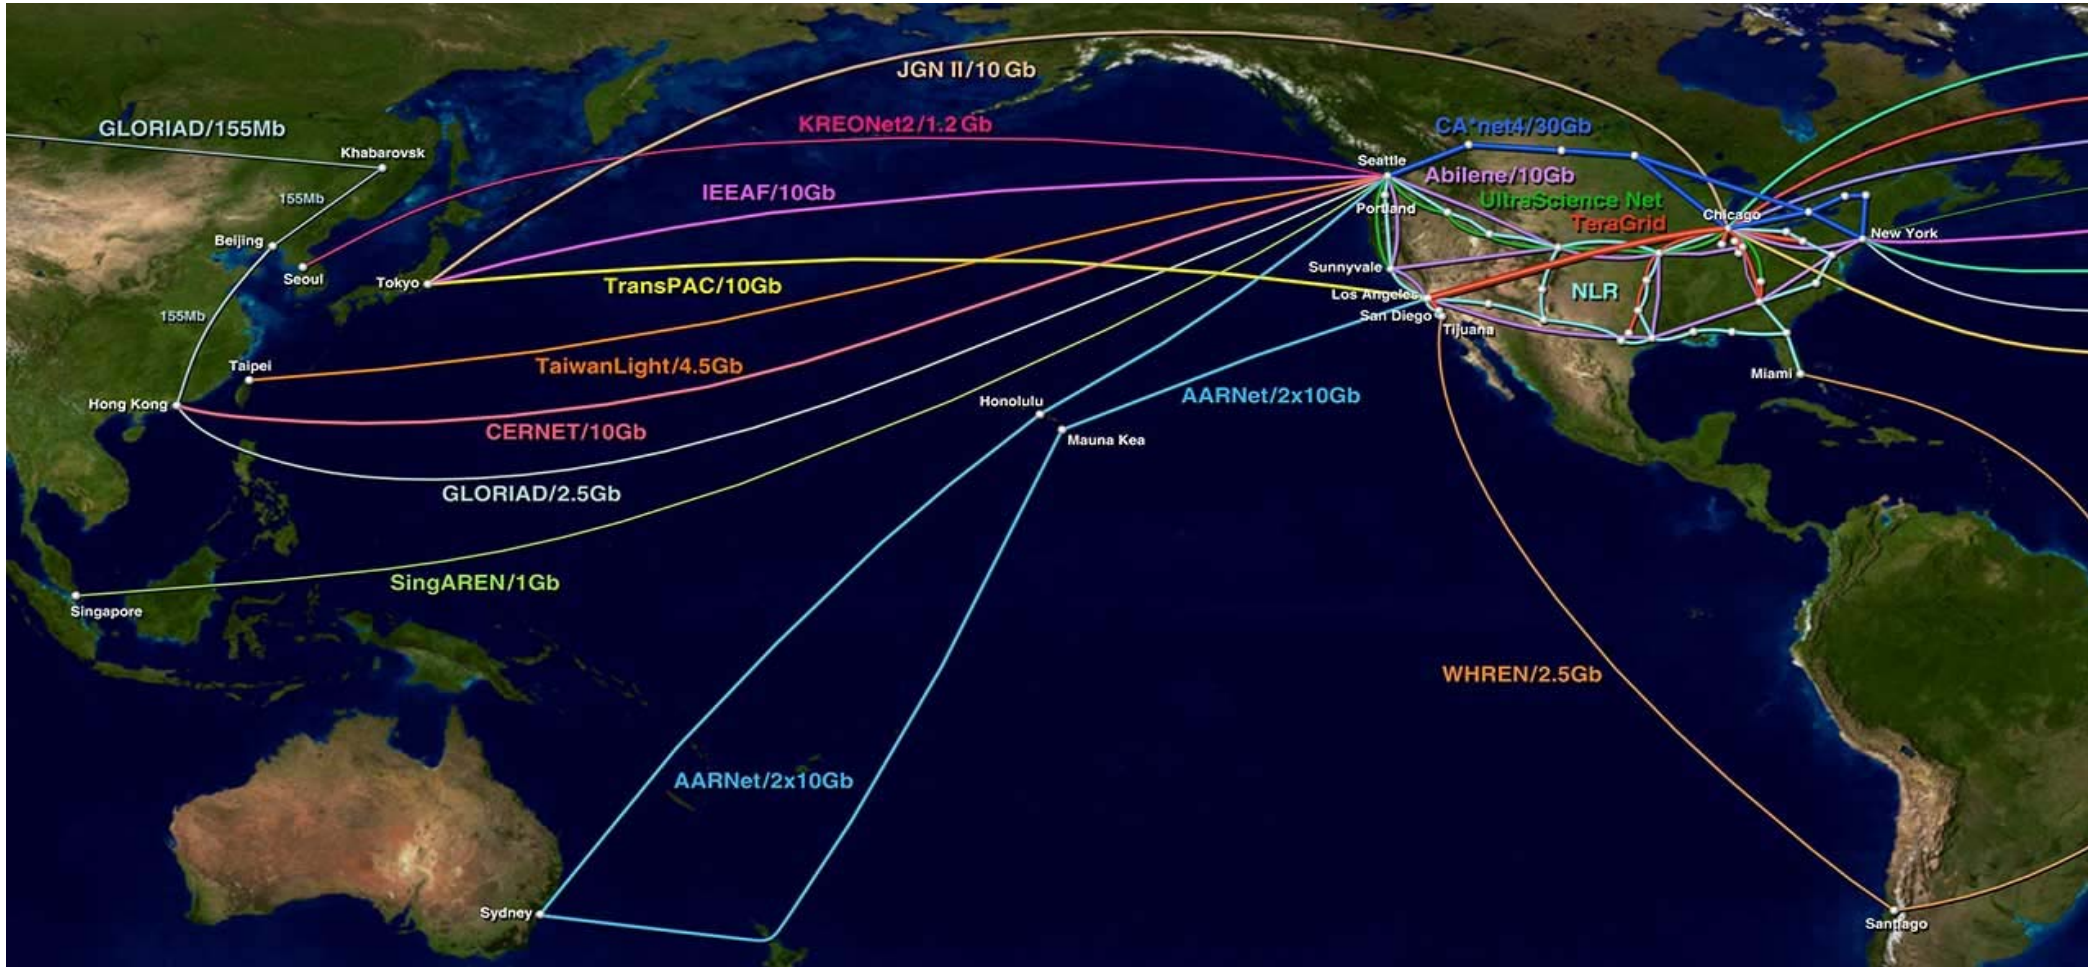
David Foster - cern-it-3





Global Lambda Integrated Facility

Predicted Bandwidth for Scheduled Experiments, December 2004



www.glif.is

last update 23/11/2004 12:11

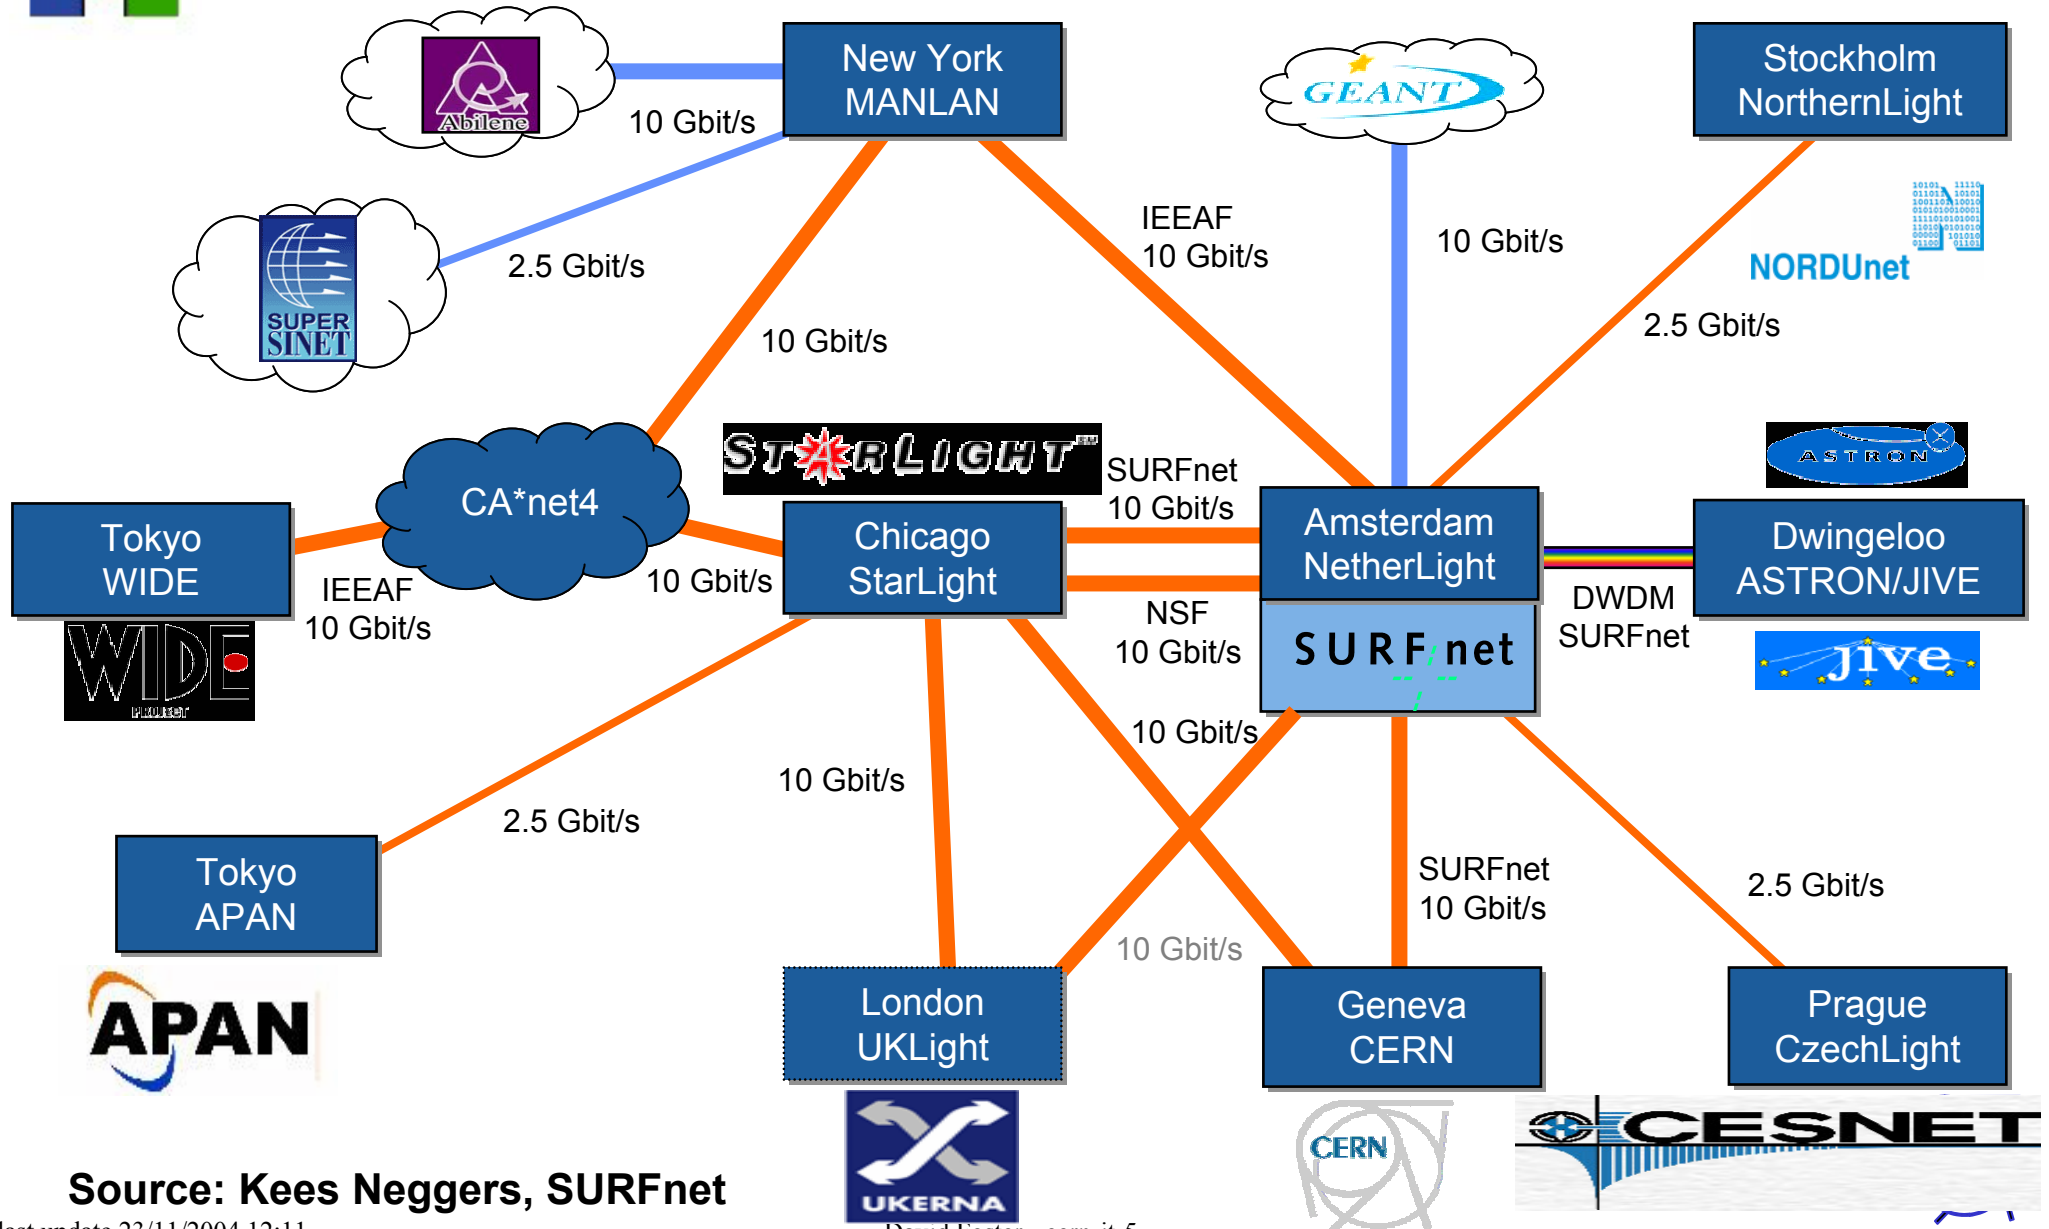
Visualization courtesy of
Bob Patterson, NCSA.

David Foster - cern-it-4





Global Participation: GLIF at StarLight



Source: Kees Neggers, SURFnet

last update 23/11/2004 12:11

David Foster - cern-it-5



What is the connectivity architecture for LHC 2004-2006?

- **Geant**
 - An EU Project, Managed by Dante spun off from Terena (an EU project) !
 - **Geant-2 lifetime (4 years from 2004).**
 - What Geant-2 will be able to deliver is not clear, yet.
 - Pop will be installed H1 2005, equipment, connectivity being selected now.
 - Requirements being passed to Dante through CERN-NREN-Dante meetings and EGEE SA2 via the Technical Network Liaison Committee (TNLC)
- **SurfNet**
 - Funded by the Dutch Government managed the GigaPort and NetherLight projects.
 - **Well funded to 2008**
 - Amsterdam is a major exchange point.
 - **Optical switch project NetherLight**
 - Will look at improving connectivity to Amsterdam to multiple wavelengths (dark fiber?)
- **US**
 - NSF and DOE both funding connectivity from [NY & Chicago] to [Amsterdam & CERN]
- **Any Direct Connections Requested by the T1's**
- **Will work with all parties to evolve connectivity.**





T1's, T2's and all that

- Not all T1 locations are completely clear.
 - "Distributed" T1's have no real meaning from a network SLA view.
- We have some idea of the T0-T1 bandwidth requirements.
 - Seems a good starting point to get 10Gb/sec end-end, CERN to the T1's. A bottom up analysis has been done as part of phase II planning.
 - But, is it clear to the T1's that they need to provision the connectivity?
 - European T1's need to negotiate with NRENs who in turn need a clear pricing model from Dante for GEANT-2.
 - Direct connections need to be provisioned from commercial providers
 - The end-end circuit has no single management responsibility (unlike commercial circuits). E.g. A T1 might connect to ESNNet, Peer with Abeline and GEANT to a European NREN. Who does problem determination and followup?
- We have no real idea of T1-T1 or T1-T2 bandwidth.
 - A 200TB T2 disk cache, refreshed in 2 days would need a dedicated 10G link
 - Less than required connectivity will give slower than expected analysis but we don't want to be too far away.
 - Computing models at the end of 2004 may shed more light on this.





Risks and Uncertainties

- To provide a 10G link to each T1 is not difficult technically if we could pay for provisioned end-end circuits. Instead we will rely on transiting interconnected networks that have no overall end-end management.
 - ESN Net initiative "ITECHS" and performance monitoring initiatives between GEANT and Abilene are steps in the right direction
 - CERN has started an end-end monitoring activity in CS Group.
- Some discussions recently indicate that centers may not yet understand their responsibility in providing the connectivity to CERN.
 - E.g. Atlas ideas around remote level-3 trigger farms.
 - GDB will have a working group to coordinate networking status with T1's and raise the visibility.
 - Should have a T0-T1 network workshop early 2005 in preparation for the TDR.
- US has many initiatives for European Connectivity but clearly identified production links and funding delivering data to Fermilab and Brookhaven has yet to appear.





Roadmap

- 2004
 - Service challenges start
- 2005
 - Q1 (January) Experiment requirements and computing models available
 - Translated into networking expectations agreed at T0-T1 workshop
 - Needs the participation of the T1's, NRENS and GEANT to create implementation plan.
 - Q2 Networking plans clarified as part of the TDR.
 - Understand the plans of the T1's, NRENS and GEANT which should be clarified as part of the MOU.
 - Q2 Will see the build up of GEANT-2 and end-end tests will be made
 - Q4 Service challenges should reach "production levels".
 - Q4 End-end monitoring technologies and limitations should be known.
- 2006
 - Q4 Full end-end connectivity in place for all T1's.
 - Q4 Data transfers at production levels.

