

# Transatlantic Networking Workshop

David Foster IT

June 2010

# Logistics

- Tea/Coffee at regular intervals!
- Lunch in restaurant 2 - 1.5hrs foreseen.
- Cocktail tonight in Restaurant 2 (whenever we finish, but hopefully before 18:00)
- Please upload your talks Pwd:TN100610
- Any Issues let myself or Kristina know.

# Format

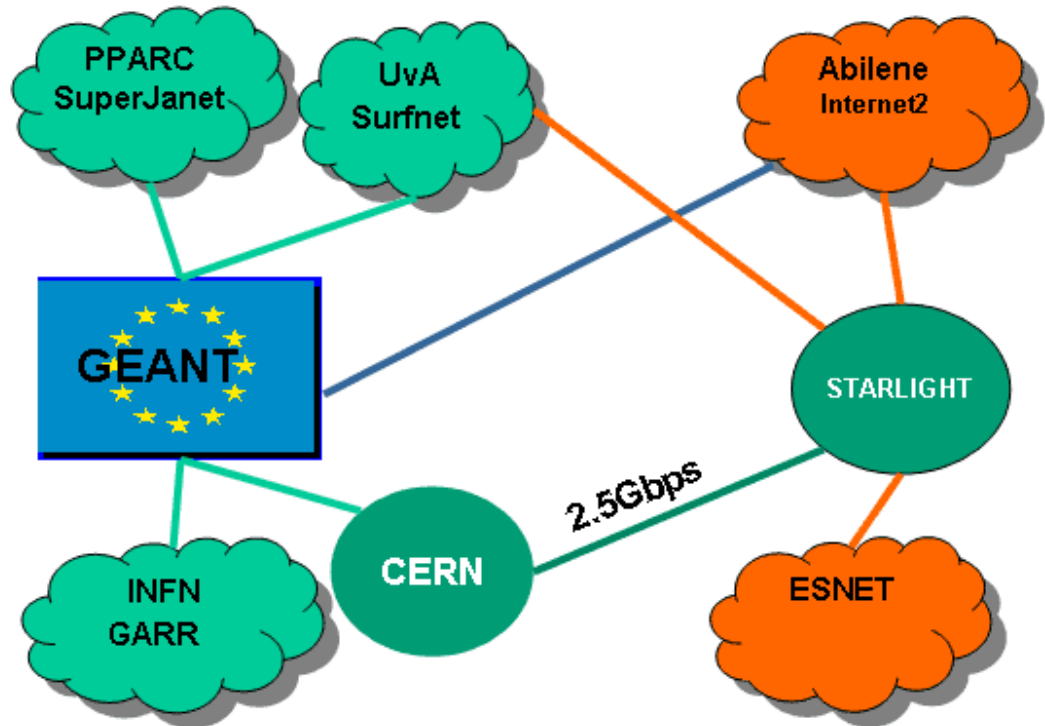
- A slight departure from the “norm” (this is supposed to be a workshop!)
  - Some introductory presentations (30 minutes each)
  - Panel – each member has 5 minutes to say what is most “on their mind” about the topics
  - Discussion – this involves everyone else!

# Programme

- 4 Themes (1/2 Day each)
  - **Architectural Approaches: Best architectures, L1 resiliency, MPLS, Carrier Ethernet, Operations, PoP's etc**
    - What will be the dominant technologies and paradigms in the next 5 years?
  - **Network as a Resource - GLIF, BoD, Oscars+DRAC, AutoBAHN, experiment interfaces and opportunities etc**
    - What are the existing and future advanced network services, and what can they offer to HEP experiments?
  - **Experiment Use Cases - Bandwidth requirements, Models, T0/1/2/3**
    - What are the data management models, bandwidth and service requirements and what would experiments expect from advanced network architectures?
  - **Networking Environment and Future Directions- Commercial activities, Opportunities, New Projects, Technology Changes, 100G etc**
    - What will be the impact of network growth, internet technologies and cost on transatlantic connectivity?

# Recent History

- In some ways, a follow on of a long tradition
  - DataTag: set up of a transatlantic test-bed 2002-2004.



# Objectives


- Now we have first LHC data, an open discussion with the community of users and providers to look at:
  - Strengths
    - What's working well and what directions we should continue in?
  - Weaknesses
    - What needs optimising, what can be done better?
  - Opportunities
    - What the LHC experiments are looking at doing in the future that will affect the networking?
    - What technology trends will help us going forward?
  - Threats
    - Where industry is going and what environmental factors (e.g. Cost) could affect what we are doing?
- Summarise the strategic ideas that are either agreed or that need further work, consideration or information.

# LHCOPN Design Model

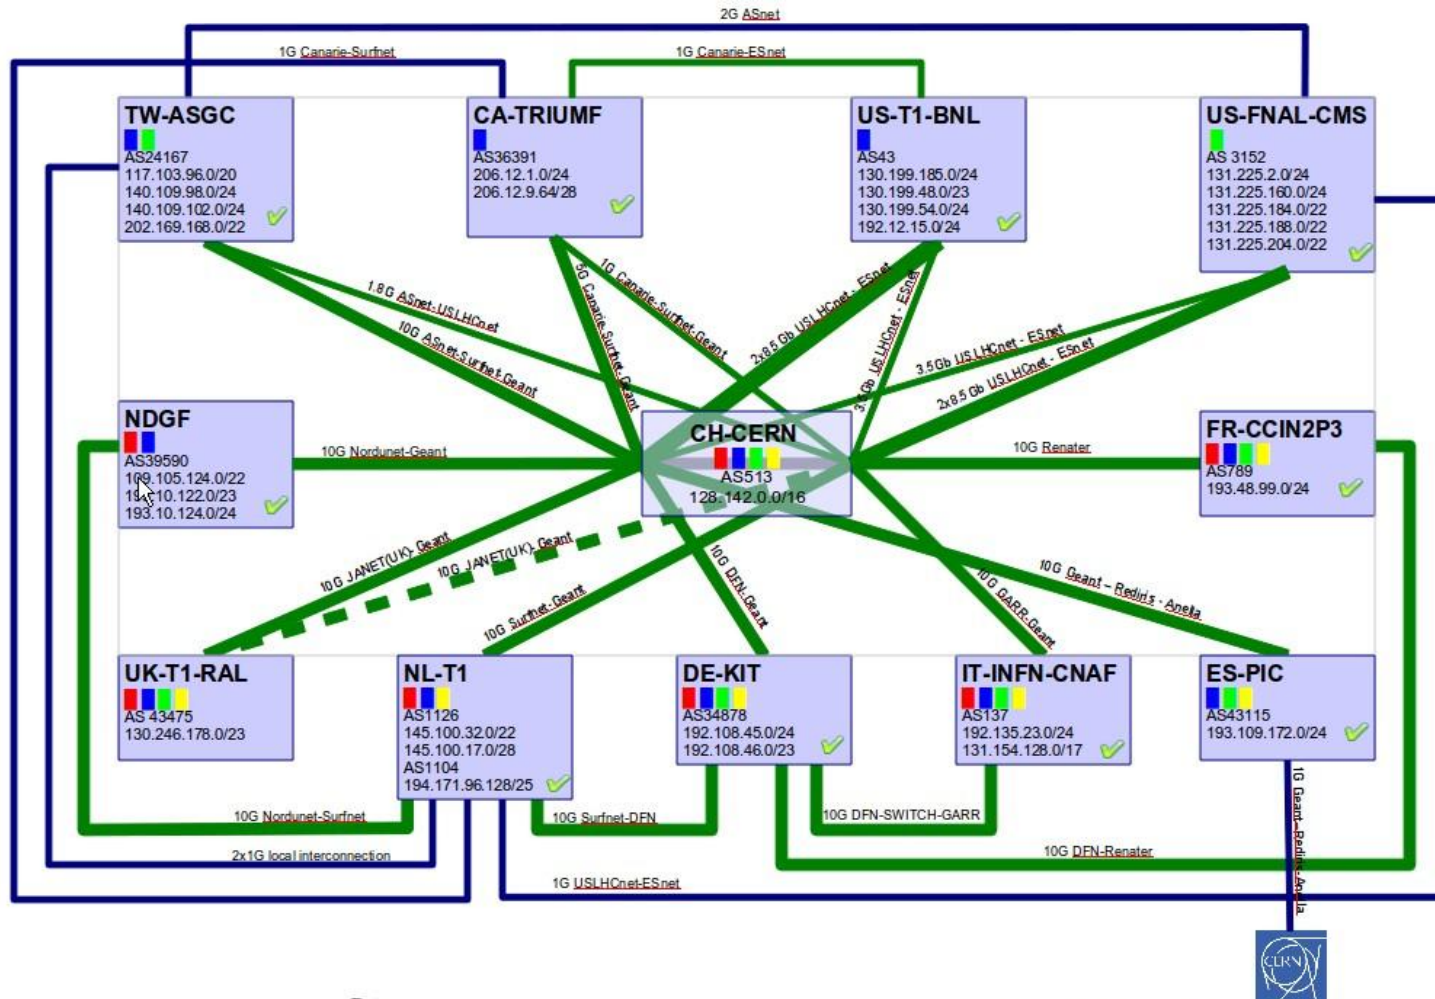
## Tier0 center, CERN

very large center; tape storage; 1st-level processing and meta-data storage; quality service (24\*7)

- Tier1 center, **11** world-wide  
large capacity; tape storage; quality service
- Tier2 center, **129** world-wide  
medium size, some large;  
no 24\*7 service;  
no custodial storage
- Tier3 center  
very small to medium size;  
no availability guarantee;  
focus on end-user analysis activity



# LHCOPN L2 Network Map



# LHCOPN


— T0-T1 and T1-T1 traffic  
— T1-T1 traffic only  
- - - Not deployed yet  
 (thick) >= 10Gbps  
 (thin) < 10Gbps  
■ = Alice ■ = Atlas  
■ = CMS ■ = LHCb  
✓ = internet backup available  
 p2p prefix: 192.16.166.0/24  
 edoardo.martelli@cern.ch 20100525



# LHCOPN Monthly Traffic Bandwidth

Traffic statistics on the CERN LHCOPN routers. Period: 25-04-2010 to 25-05-2010  
 (source: <http://network-statistics.web.cern.ch/network-statistics/ext/?p=sc&q=LHCOPN%20Total%20Traffic&m=LHCOPN-Total>)

LHCOPN TOTAL Traffic Flow (Out-bound)




To CNAF	Avg	987.52 M	Max	4.07 G	Peak	6.60 G
To DE-KIT	Avg	1.62 G	Max	8.31 G	Peak	9.62 G
To IN2P3	Avg	1.67 G	Max	6.06 G	Peak	7.57 G
To NDGF	Avg	359.15 M	Max	1.74 G	Peak	3.26 G
To ES-PIC	Avg	1.56 G	Max	6.58 G	Peak	7.55 G
To RAL	Avg	1.81 G	Max	5.52 G	Peak	7.08 G
To NLT1	Avg	786.67 M	Max	6.97 G	Peak	7.75 G
To TRIUMF	Avg	646.92 M	Max	2.70 G	Peak	3.30 G
To BNL	Avg	1.51 G	Max	5.08 G	Peak	5.75 G
To FNAL	Avg	351.68 M	Max	1.89 G	Peak	3.49 G
To ASGC	Avg	1.20 G	Max	4.76 G	Peak	6.71 G

Total to Tiers1 - average 12.71 G  
 Total to Tiers1 - maximum 39.77 G

SPECTRUM Report Gateway  
 Last Updated: Tue May 25 17:48:18 2010

LHCOPN TOTAL Traffic Flow (In-bound)



From CNAF	Avg	253.76 M	Max	3.08 G	Peak	4.05 G
From DE-KIT	Avg	1.48 G	Max	7.82 G	Peak	9.31 G
From IN2P3	Avg	1.25 G	Max	4.66 G	Peak	8.07 G
From NDGF	Avg	320.66 M	Max	4.25 G	Peak	7.15 G
From ES-PIC	Avg	843.88 M	Max	6.73 G	Peak	6.73 G
From RAL	Avg	1.71 G	Max	8.46 G	Peak	9.54 G
From NLT1	Avg	852.69 M	Max	3.72 G	Peak	5.48 G
From TRIUMF	Avg	298.17 M	Max	2.33 G	Peak	3.01 G
From BNL	Avg	638.23 M	Max	2.87 G	Peak	3.62 G
From FNAL	Avg	174.93 M	Max	3.60 G	Peak	6.07 G
From ASGC	Avg	190.62 M	Max	1.89 G	Peak	3.65 G

Total from Tiers1 - average 8.10 G  
 Total from Tiers1 - maximum 28.45 G

SPECTRUM Report Gateway  
 Last Updated: Tue May 25 17:48:18 2010

# T0-T1 links availability

## Servicegroup 'BGP'

2009-07-01 00:00:00 to 2010-05-18 08:49:50

Duration: 321d 8h 49m 50s

### CERN LHCOFN routers

Host	% Time Up	% Time Down	% Time Unreachable	% Time Undetermined
<a href="#">1513-c-rftec-1</a>	100.000% (100.000%)	0.000% (0.000%)	0.000% (0.000%)	0.000%
<a href="#">1513-c-rftec-2</a>	99.995% (99.995%)	0.000% (0.000%)	0.005% (0.005%)	0.000%

### Servicegroup 'BGP' Service State Breakdowns:

Host	Service	% Time OK	% Time Warning	% Time Unknown	% Time Critical	% Time Undetermined
<a href="#">1513-c-rftec-1</a>	<a href="#">lhcopn-bgp-asgc</a>	99.688% (99.688%)	0.000% (0.000%)	0.000% (0.000%)	0.312% (0.312%)	0.000%
	<a href="#">lhcopn-bgp-bnl</a>	99.719% (99.719%)	0.000% (0.000%)	0.000% (0.000%)	0.281% (0.281%)	0.000%
	<a href="#">lhcopn-bgp-bnl-sec</a>	64.467% (99.908%)	0.000% (0.000%)	0.000% (0.000%)	0.059% (0.092%)	35.473%
	<a href="#">lhcopn-bgp-fnal-bu</a>	99.135% (99.135%)	0.000% (0.000%)	0.000% (0.000%)	0.865% (0.865%)	0.000%
	<a href="#">lhcopn-bgp-gridka</a>	99.755% (99.755%)	0.000% (0.000%)	0.000% (0.000%)	0.245% (0.245%)	0.000%
	<a href="#">lhcopn-bgp-ndgf</a>	99.356% (99.356%)	0.000% (0.000%)	0.000% (0.000%)	0.644% (0.644%)	0.000%
	<a href="#">lhcopn-bgp-pic</a>	98.182% (98.182%)	0.000% (0.000%)	0.000% (0.000%)	1.818% (1.818%)	0.000%
	<a href="#">lhcopn-bgp-ral</a>	99.975% (99.975%)	0.000% (0.000%)	0.000% (0.000%)	0.025% (0.025%)	0.000%
	<a href="#">lhcopn-bgp-triumf</a>	98.316% (98.316%)	0.000% (0.000%)	0.000% (0.000%)	1.684% (1.684%)	0.000%
<a href="#">1513-c-rftec-2</a>	<a href="#">lhcopn-bgp-asgc-starlight</a>	99.839% (99.839%)	0.000% (0.000%)	0.003% (0.003%)	0.158% (0.158%)	0.000%
	<a href="#">lhcopn-bgp-bu-bnl</a>	99.452% (99.452%)	0.000% (0.000%)	0.000% (0.000%)	0.548% (0.548%)	0.000%
	<a href="#">lhcopn-bgp-cnaf</a>	99.970% (99.970%)	0.000% (0.000%)	0.003% (0.003%)	0.027% (0.027%)	0.000%
	<a href="#">lhcopn-bgp-fnal</a>	99.577% (99.577%)	0.000% (0.000%)	0.000% (0.000%)	0.423% (0.423%)	0.000%
	<a href="#">lhcopn-bgp-fnal-sec</a>	50.752% (99.216%)	0.000% (0.000%)	0.000% (0.000%)	0.401% (0.784%)	48.847%
	<a href="#">lhcopn-bgp-in2p3</a>	99.996% (99.996%)	0.000% (0.000%)	0.000% (0.000%)	0.004% (0.004%)	0.000%
	<a href="#">lhcopn-bgp-sara</a>	99.485% (99.485%)	0.000% (0.000%)	0.003% (0.003%)	0.511% (0.511%)	0.000%
	<a href="#">lhcopn-bgp-triumf-bu</a>	93.958% (93.958%)	0.000% (0.000%)	0.000% (0.000%)	6.042% (6.042%)	0.000%

## Explanation

Link availability calculated on the uptime of the BGP routing protocol relationship with each neighbor. It says for how long the T0's routers were able to route traffic to/from every T1 over the LHCOFN link. It doesn't say anything about Data-Centre to Data-Centre connectivity.

FNAL-Sec and BNL-sec are the additional 10G links recently deployed, thus the high "Time Underdetermined".

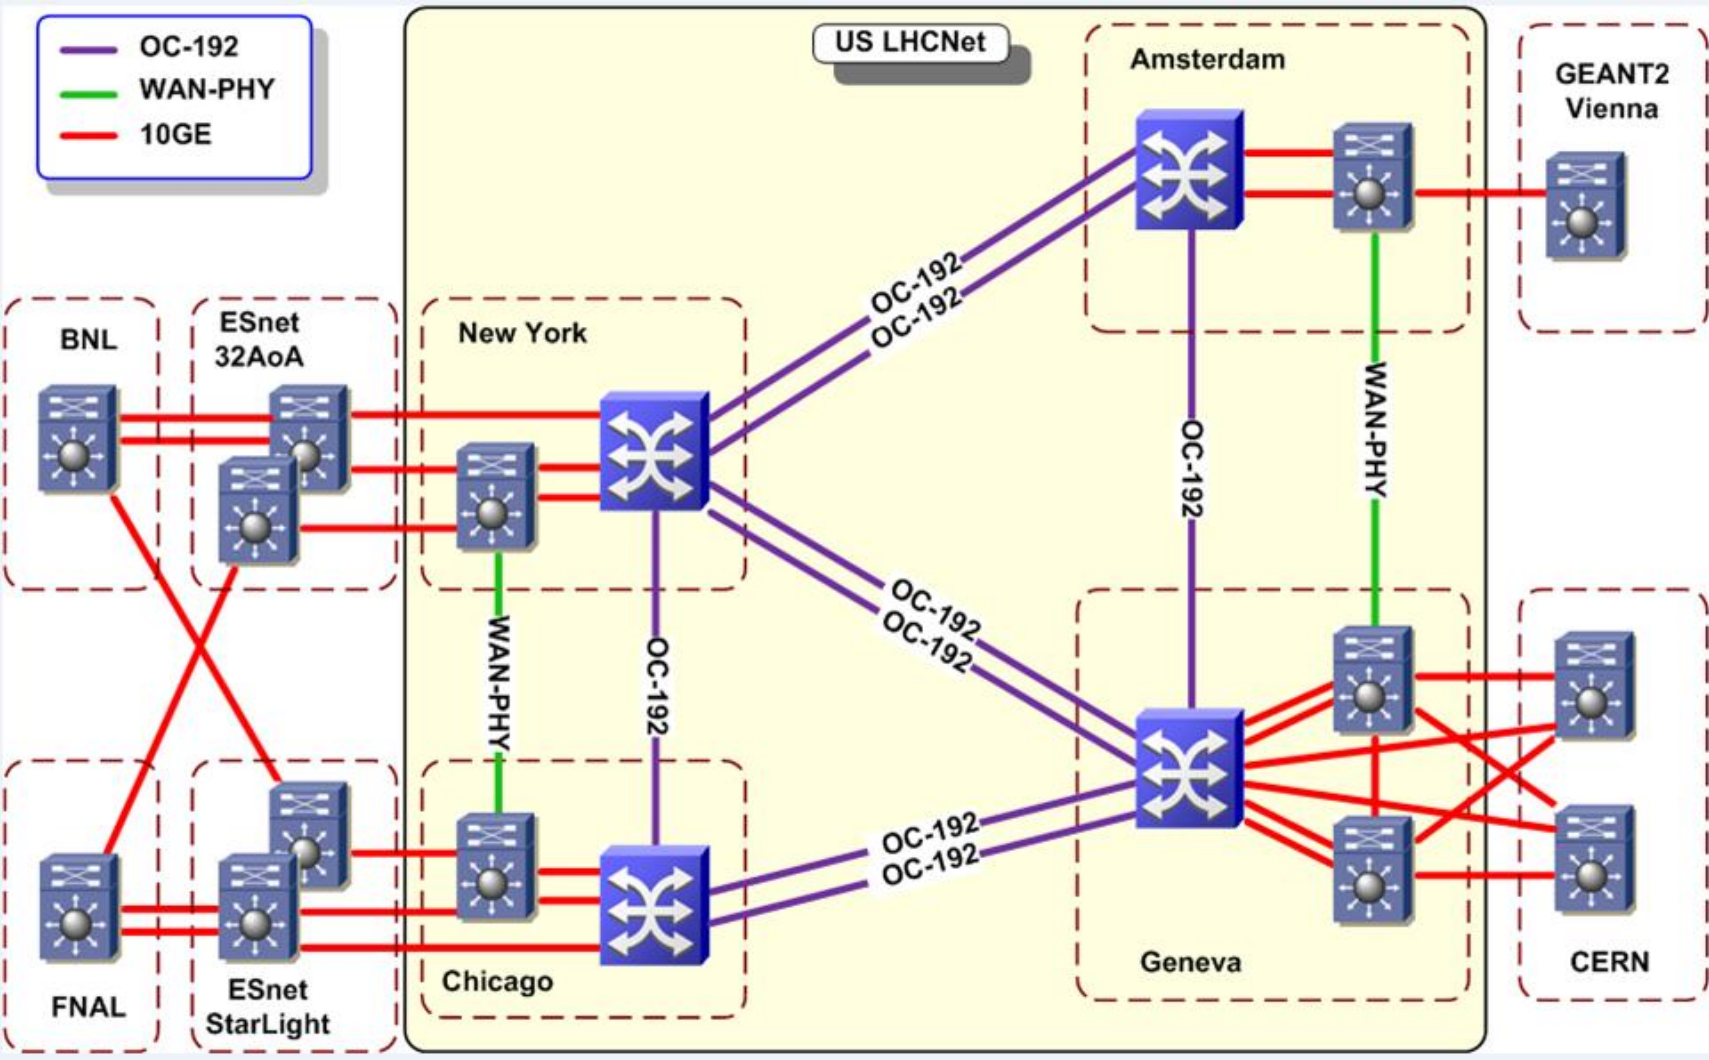
# LHCOPN, AS A Whole

- *Has good parentage ....*
  - Tier-0/1 network managers, GEANT, USLHCNet, Commercial Services.
  - Allows a distributed funding model.
  - Federated approach giving close collaboration between users and service providers.
- *... is 5 years old this year ...*
  - Stable.
  - Achieving its design goals.
- *... knows its limits ...*
  - Bounded community (T1 centers)
  - T0-T1 and T1-T1
- *... but is still an infant ...*
  - Capacity planning.
  - Monitoring
  - Operational Robustness.
  - SLA's becoming increasingly important
- *... and can grow ...*
  - New types of services, dynamic circuits etc.
  - Help coordinate other activities, e.g. For T2's?

# USLHCNet

- Is a network infrastructure in its own right dealing with the special issues of transatlantic networking, broader than just the LHCOPN:
  - Provides services for the LHCOPN.
  - Provides services for Tier-2 connectivity.
  - Engages in forward planning and feasibility testing for future requirements.

# USLHCNet



# Thanks!

- To everyone for attending.
- To Artur Barczyk for helping in all aspects and lobbying everyone.
- To Kristina Gunne for all the logistics.
- To Ludwig Pregernig, John Shade and Artur Barczyk for helping with the running of the meeting.