CHPC WLCG Tier2 Facility

Sean Murray
ALICE
CHPC
CSIR

July 7, 2016

Outline

What is the WLCG

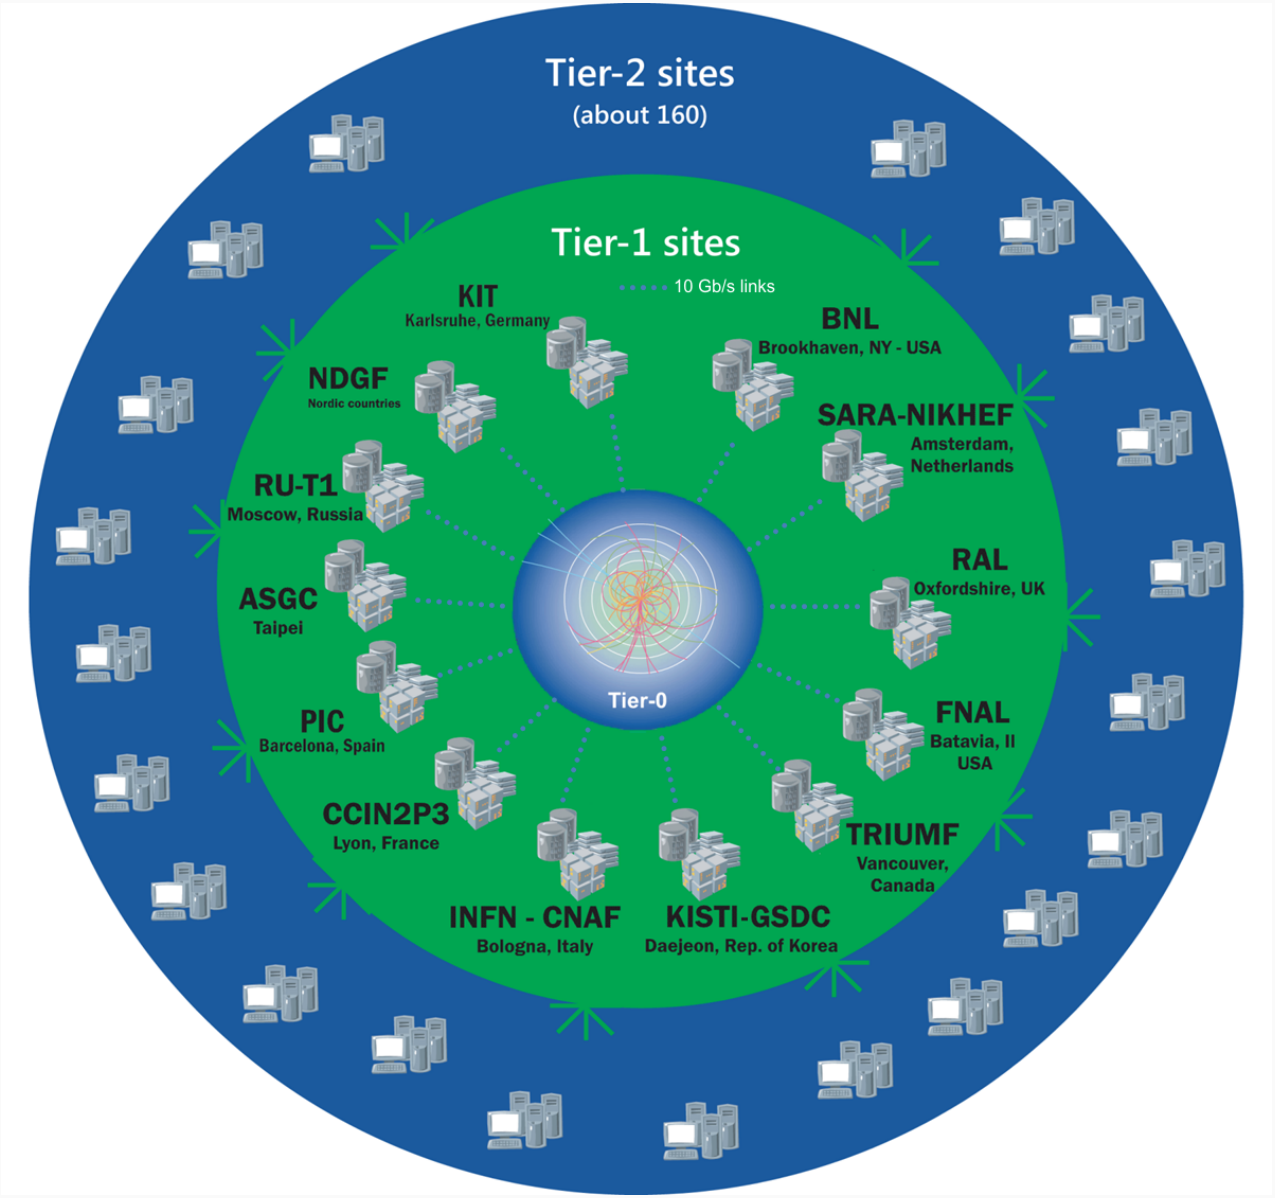

WLCG Map of Sites

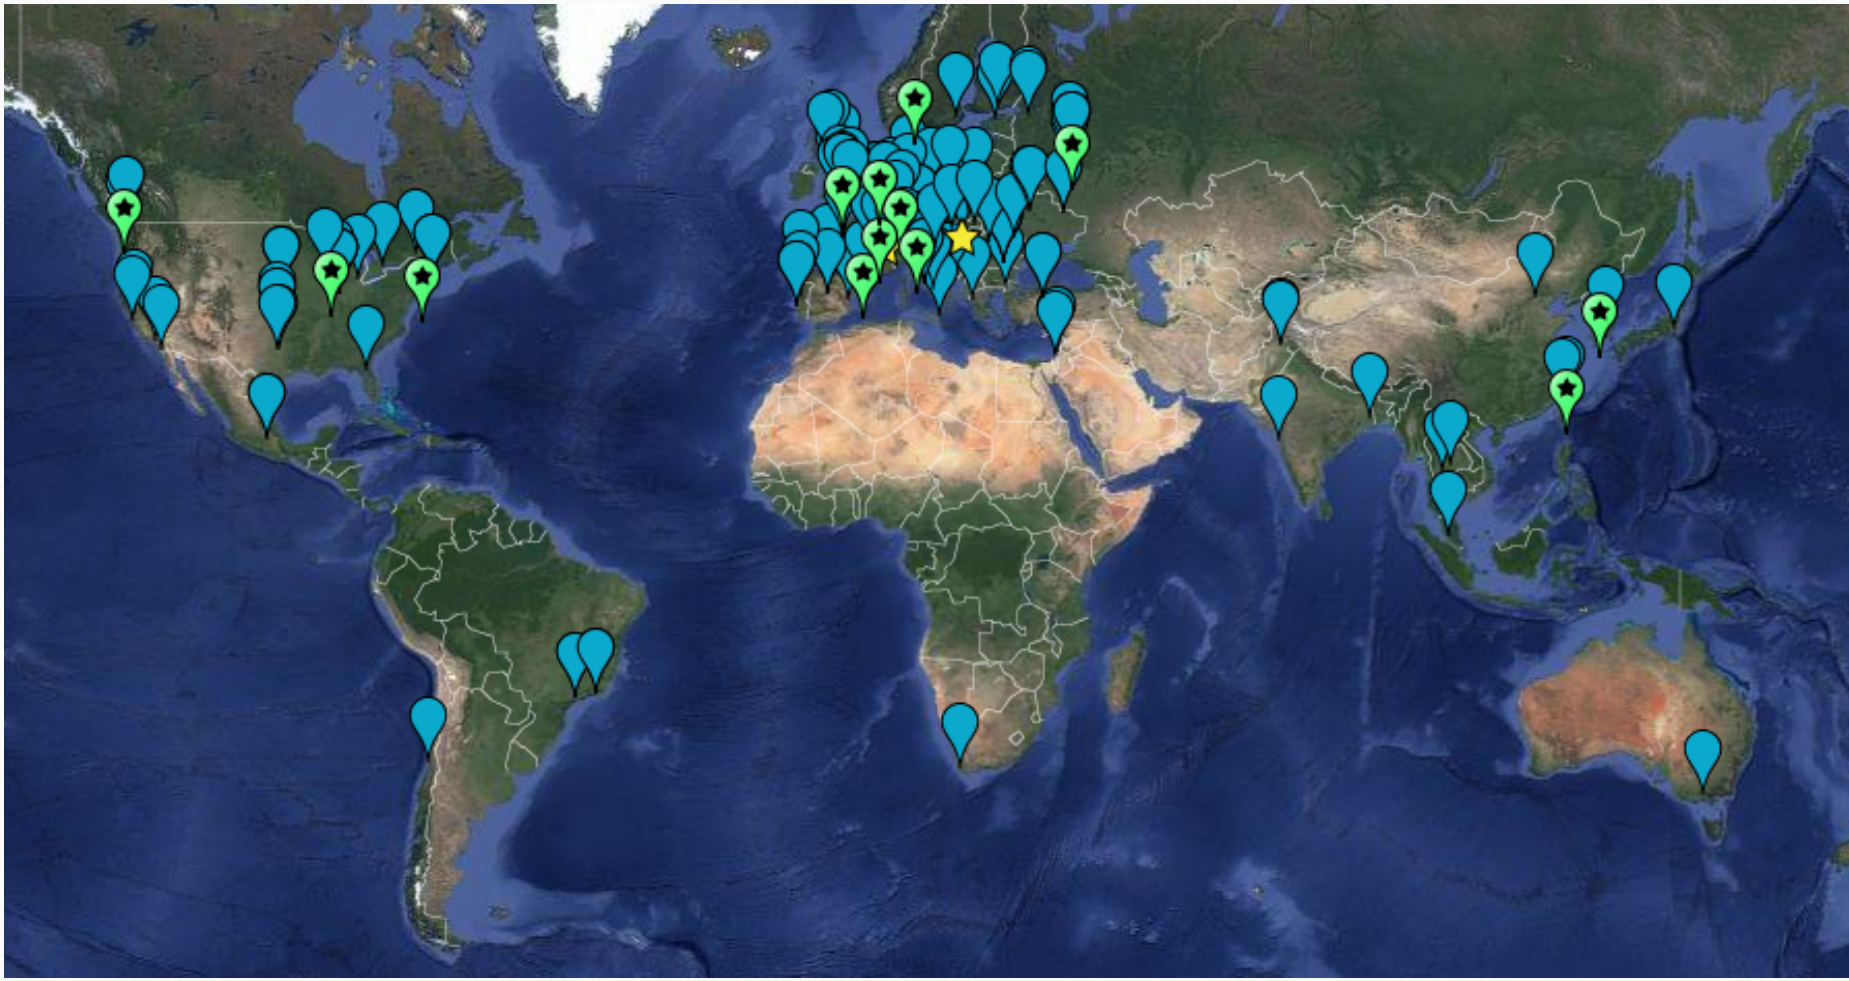

WLCG MOU Signed

CHPC
CENTRE FOR HIGH PERFORMANCE COMPUTING

WLCG
Worldwide LHC Computing Grid

28 April 2015

Commitments

According to Tender :

- ALICE 600 cores
- ATLAS 600 cores
- ALICE 400TB
- ATLAS 400TB

According to :
https://wlcg-rebus.cern.ch/apps/pledges/resources/

- 6000 HEPSPEC06 cores (560 of our cores)
- 100TB storage
- All ALICE.

Computing Infrastructure

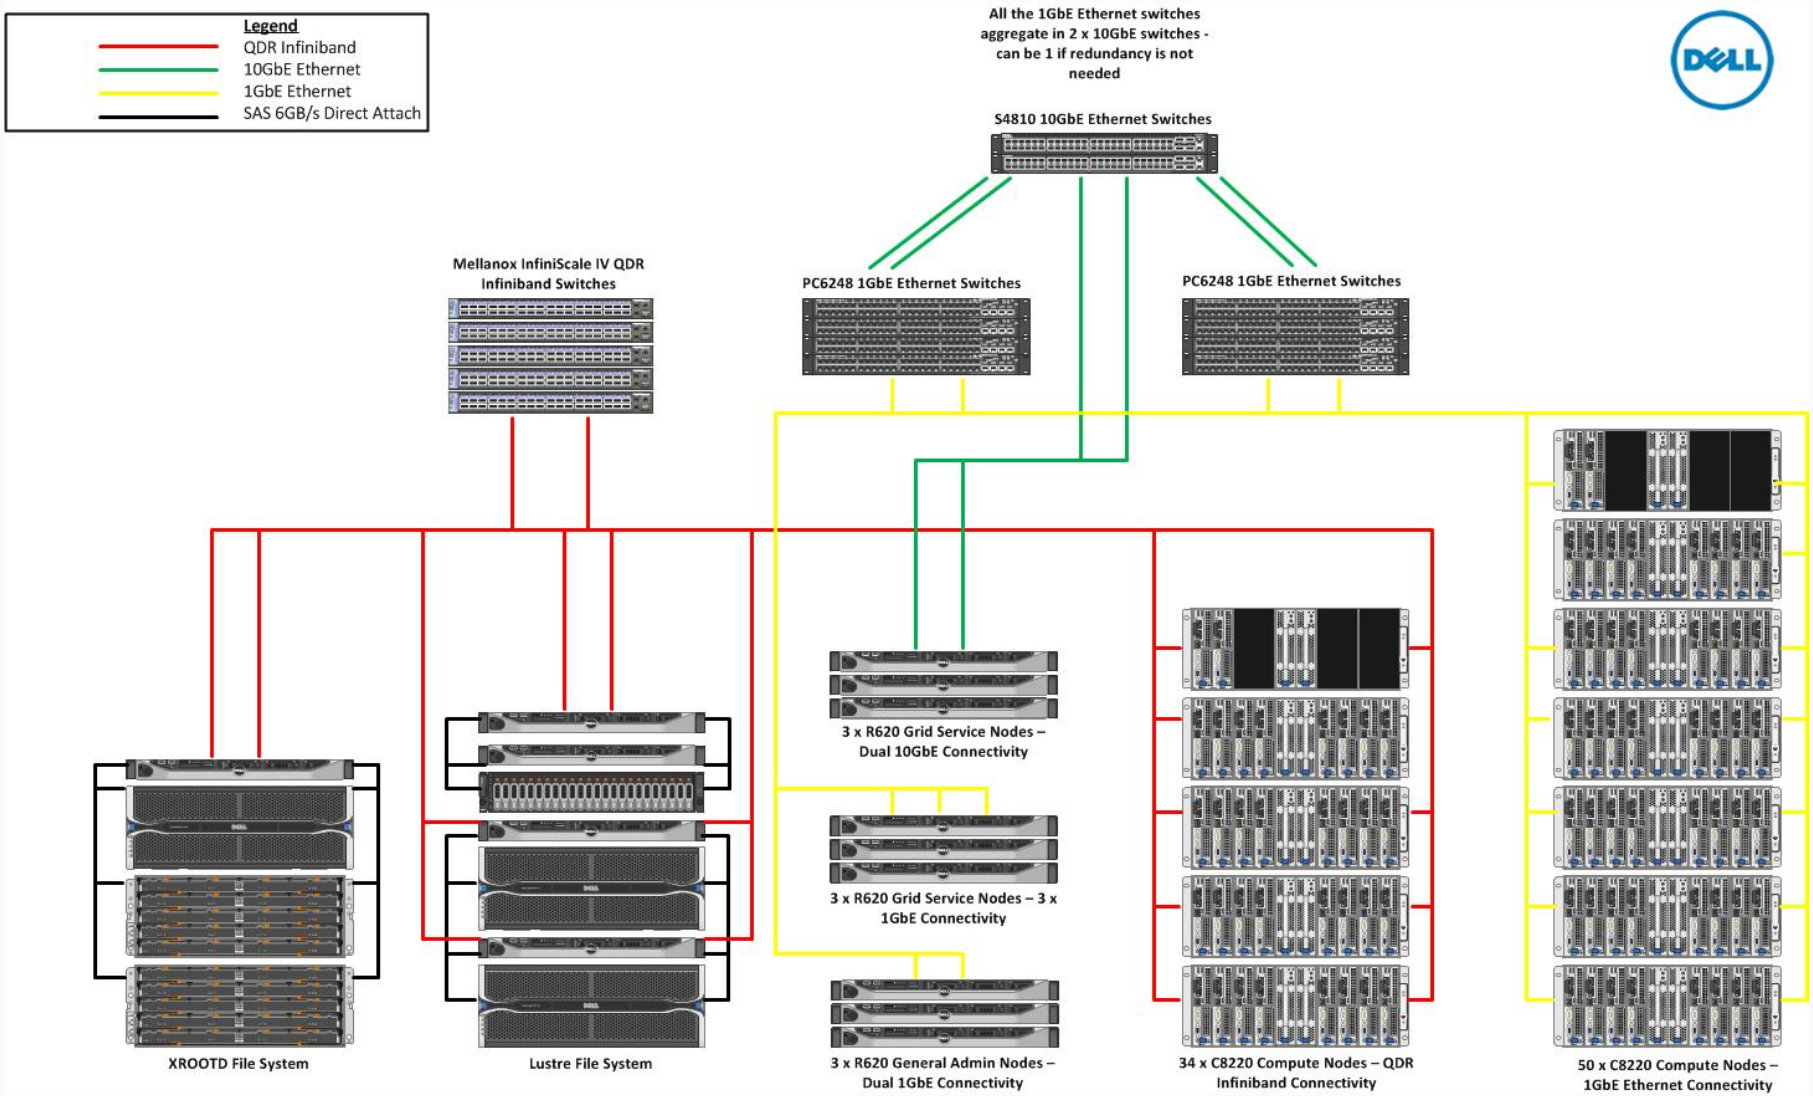

Current hardware

- 50 nodes of 48 cores 192GB RAM and 1.6TB of SSD, 1G ethernet
- 34 nodes of 48 cores 96GB RAM and 1TB, FDR infiniband, 6 “stolen”
- 100TB of Lustre on the 34 nodes with FDR infiniband.
- 9 management servers, lower spec
  - compute element (head node,ce),
  - storage element 2 redirectors, 2 storage nodes with direct attached multipath storage
  - authentication, user interface (gone), monitoring, provisioning.

Current Storage

- 383TB EOS for ALICE, down from 440TB
- 252 TB EOS for ATLAS, down from 400TB
- 107 TB lustre for 34 nodes.
- 104 TB EMC for ATLAS, not going to be used.

Reduction in data sizes is due to reorganisation for reliability.

Current Performance

- 465k ALICE jobs in last year
- Avg concurrent jobs 704.
- ALICE 358TB storage properly live 24 June, Consumed 6.4TB
- Data traffic consumed 70TB in and 60TB out in the last 3 months
- 11MB/s in and 7.5MB/s

Availability / Reliability

So who monitors us :

- WLCG
- EGI
- ALICE
- ATLAS
- me via zabbix/grafana and AAROC

There are a lot of eyes, ignoring the ones in this room.

Function	May	April	March
Availability	90	97	91
Reliability	100	90	90

Problems in monitoring infrastructure in June, pending recalculation.

Grafana ALICE

This is currently being expanded to pull in more from MonaLisa, and more experiment and storage specific metrics to help to trivialy diagnose and forewarn problems.

Problems

It has not been an easy 4 months, the big ones are :

17 Mar	Switch dies, kills everything, 4 days
14 Apr	disk dies, 15/16 April rebuild fails, 109TB of empty space lost.
21May	Site Power upgrade whole weekend.
22 Jun	General Power failure, storage out for 2 days, SAM tests in unknown for 1 week.

Our Data's Scenic Tour

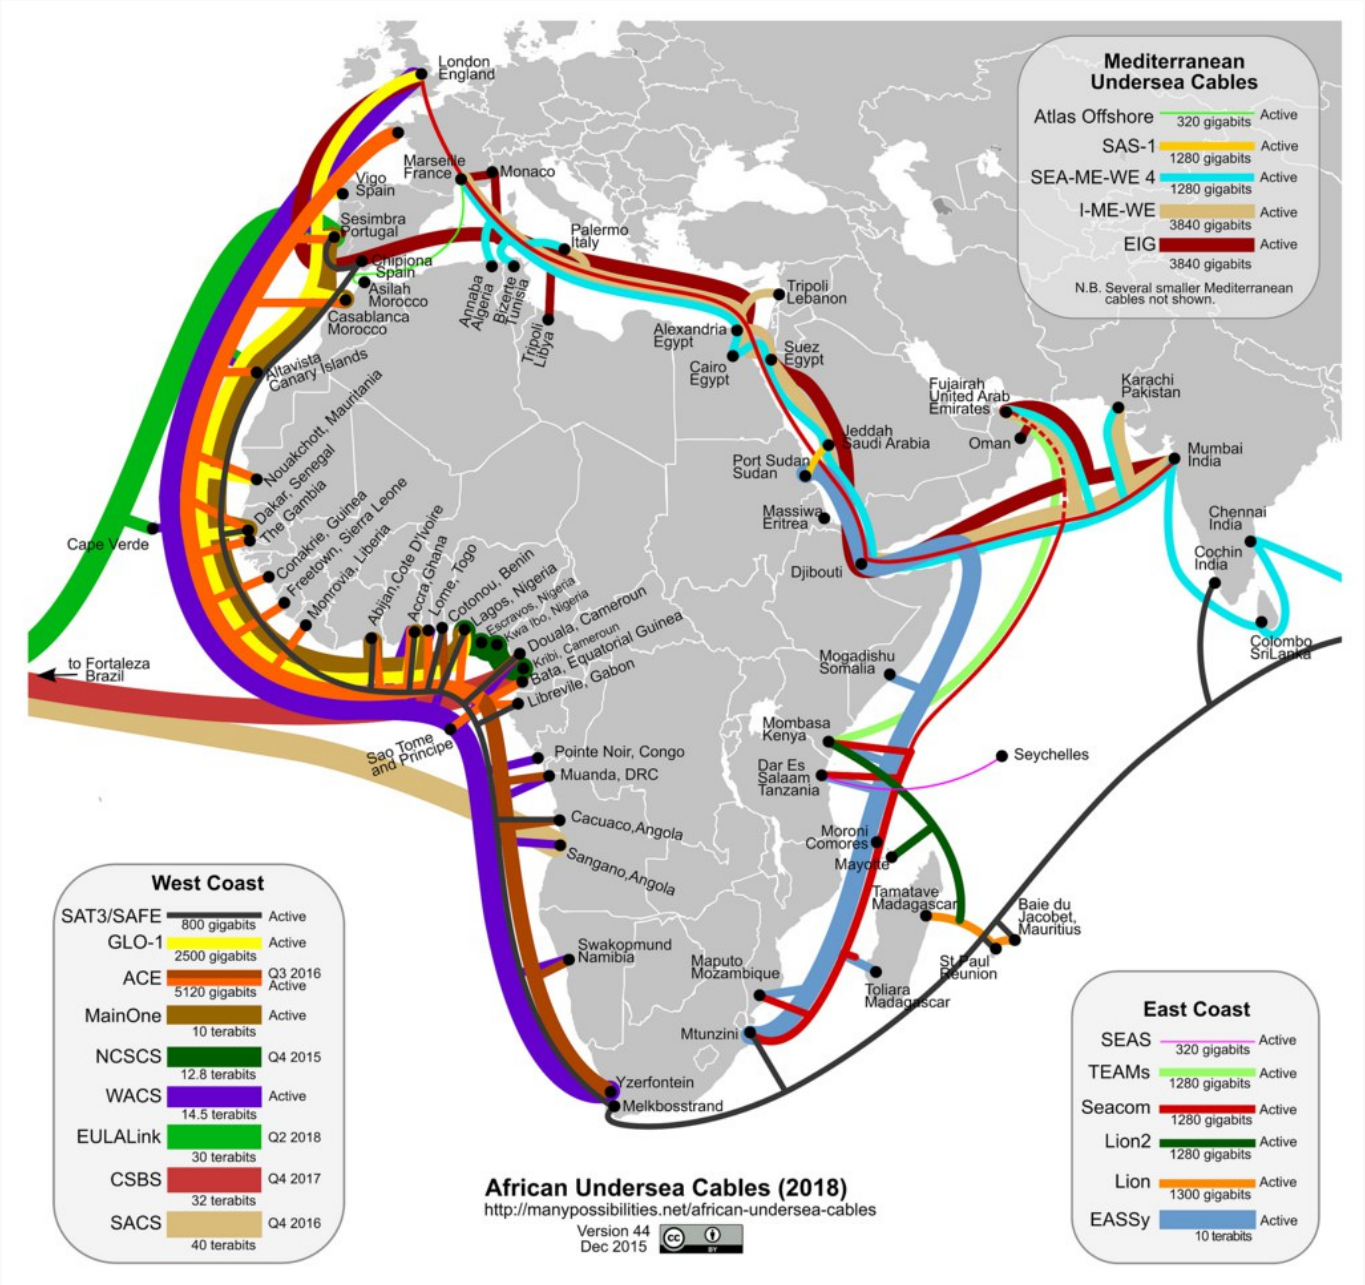

Data Traffic Mid June

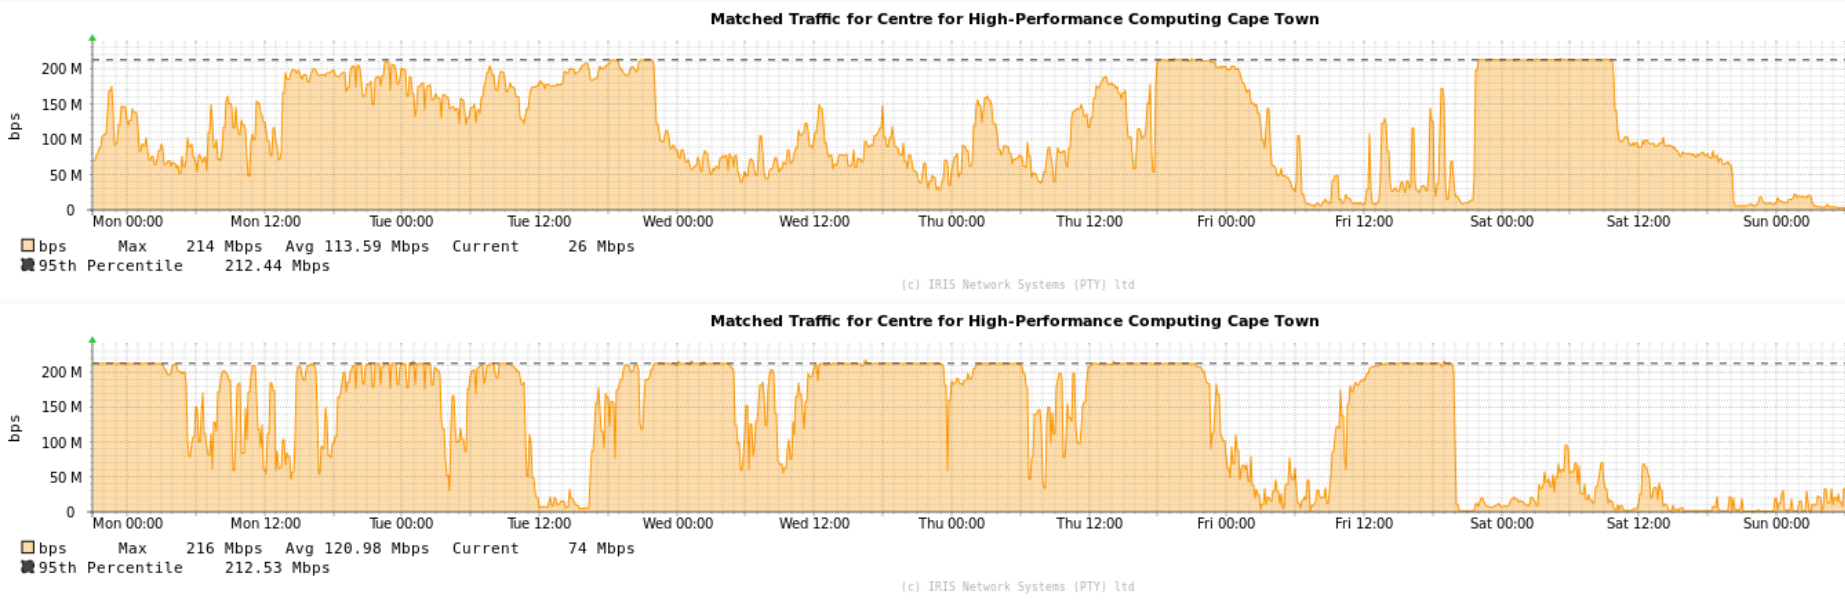

Improvements

Since my epoch . . .

- Fixed ALICE error rates, mostly.
- ALICE concurrent jobs from 700 to 1500.
- max out bandwidth now regularly.
- ATLAS running pilot jobs.
- Storage cleaned up.
- New storage quoting
- Plan in place to overhaul, and being tested.

upgrades

- attempts to delay till CC7 validation have failed.
- Puppify, to auto site deployment, r10k an issue.
- Transition to foreman from xcat.
- Fix ATLAS Storage (reinstall)
- upgrade monitoring server to zabbix3 and new grafana to technical reasons.
- Rewire whole network, re-power to monitored pdu, and monitor all.
- Add inherent redundancy into 10G interfaces on ce,se,se2.
- Reinstall while TRYING to keep A/R. problems are vobox and ce.
- Storage, we need an additional 750TB for ALICE and ATLAS.

Computing Infrastructure

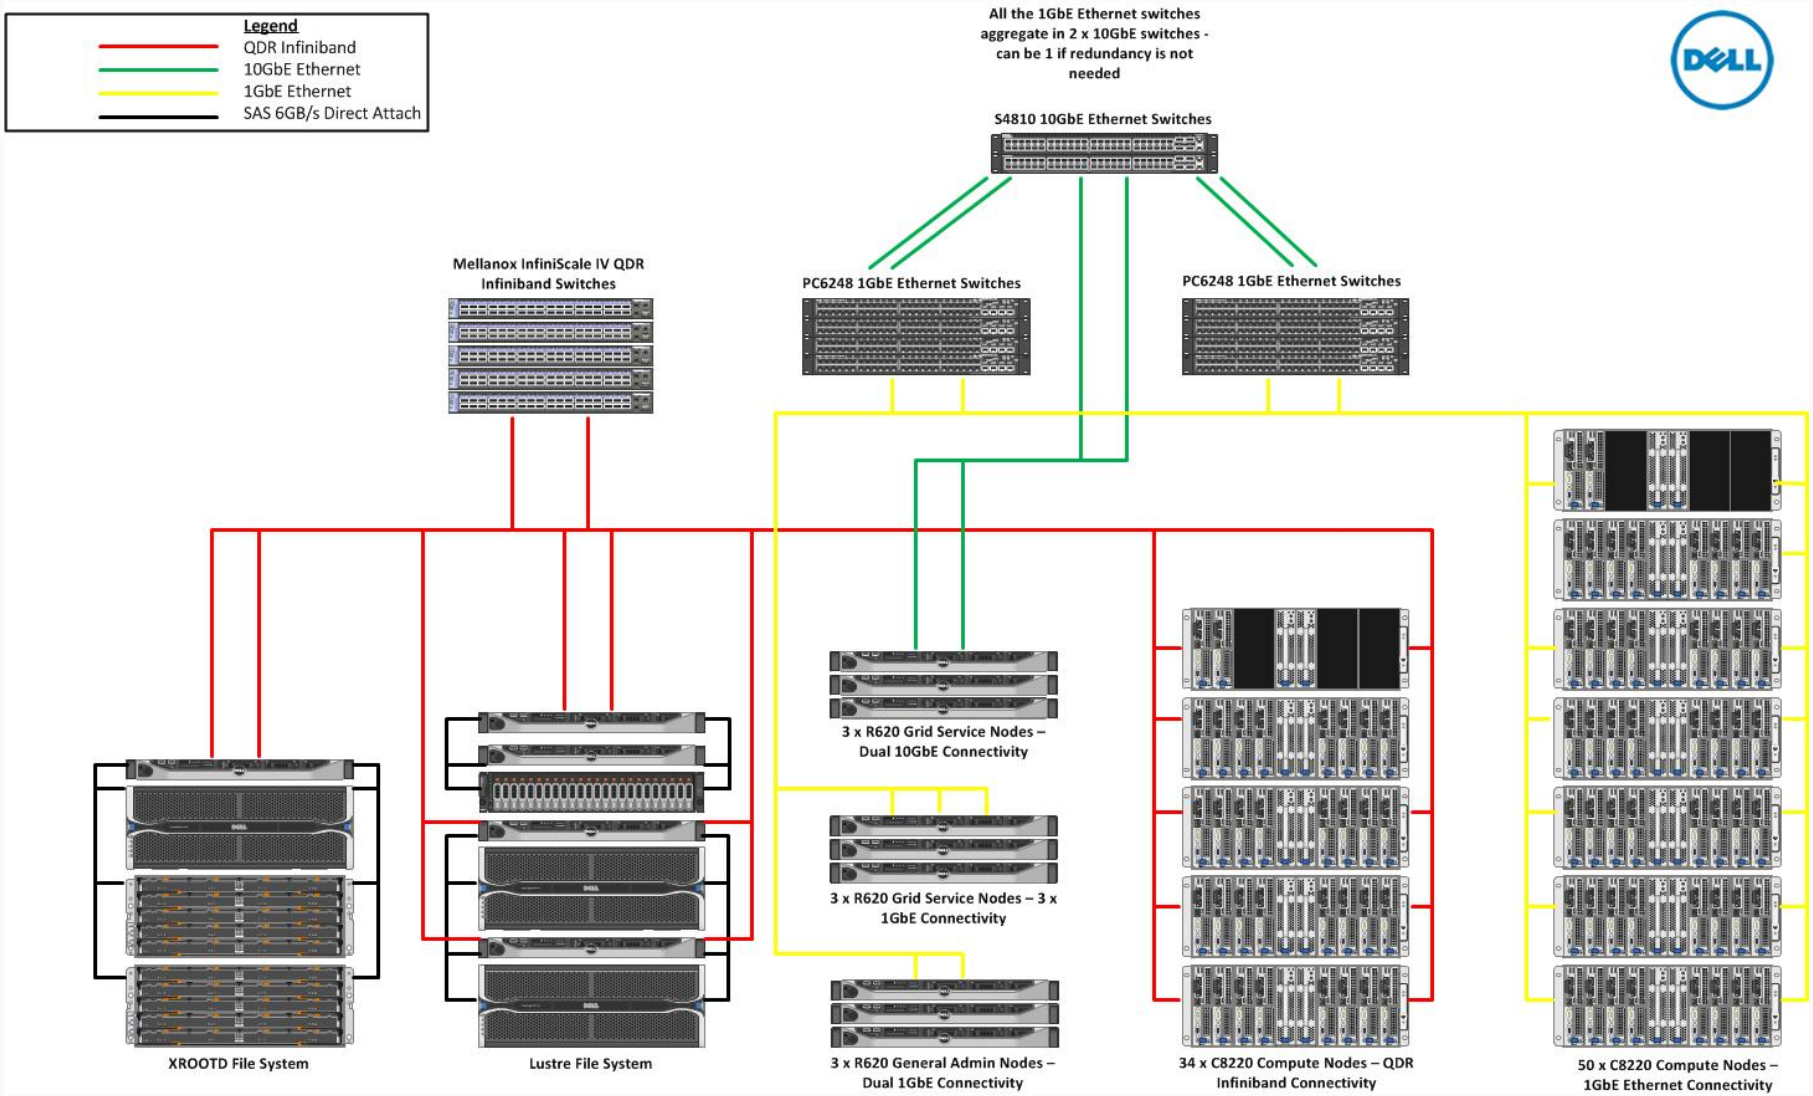

Transparency

Historically this has not been great, so . . .

- Federated logins to zabbix and grafana, i.e. You
- All code on github in line with AAROC
- All issues PUBLIC on github
- Still have GGUS for normal tickets.
- Some training on the user analysis facility (hopefully online)
- A couple of things remain private like network diagrams, obviously, and passwords.
- AAROC slac channels.


ALICE
Zoom Out
Job Agents Running
1800
Batch Jobs Running
1800
ALICE Jobs Running
1797
Jobs Running/Queued
Batch Jobs Running
Batch Jobs Queued
Running Jobs
Jobs Agents
Job Statuses
alice_failed_jobs
alice_saved_jobs
alice_error_v_jobs
alice_error_jb_jobs
alice_error_e_jobs
alice_zombie_jobs
ALICE EOS
No datapoints
ALICE Zombies Running
ALICE Running


ATLAS
Zoom Out
Total ATLAS Jobs Running
3
Jobs Running/Queued
Total ATLAS Jobs Running
Total ATLAS Jobs Queued

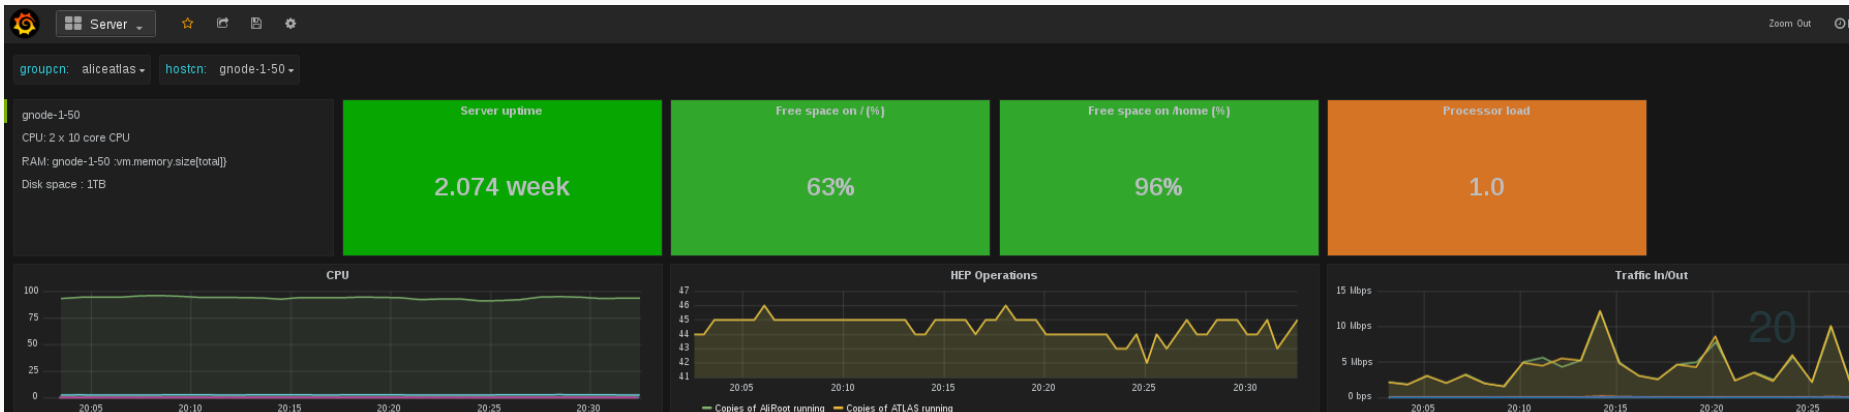
Server
Zoom Out
groupcn: aliceatlas
hostcn: gnode-1-50
gnode-1-50
CPU: 2 x 10 core CPU
RAM: gnode-1-50 :vm.memory.size[total]}
Disk space : 1TB
Server uptime
2.074 week
Free space on / (%)
63%
Free space on /home (%)
96%
Processor load
1.0
CPU
HEP Operations
Copies of AliRoot running
Copies of ATLAS running
Traffic In/Out

Outline

28 nodes

I got the go ahead 3 weeks ago to claim 28 of the 34 nodes back for SAGrid and HEP user analysis.

- Go back to SAGrid to support anybody on SAGrid VO.
- hep user analysis, based on federated identities, no user account admin.
- code based on CODE-RADE, or LHC experiments.
- Local Storage for users, eos and lustre.

AAROC and SAGrid

- A collaboration spanning Africa and Arabia.
- Everything is on github under AAROC.
- Share the resources of a massively disparate collection of computing resources transparently to the user.
- Idea of a automated site, via sound deployment tools, and constant integration.

Outline

Tier 1 Technical Requirements

Its a long list, I wont bore you or me with the intricate details of the MOU.

First and foremost :

Tier 1 Technical Requirements

Its a long list, I wont bore you or me with the intricate details of the MOU.

First and foremost :

A STABLE, RELIABLE PROVEN Tier2

The criticality of that can not be under estimated.

Tier 1 Technical Requirements

- Custodial storage of raw data.
- O(10k) cores
- single digit PB disk
- single digit PB online Tape library, custodial raw data.
- redundant links on LHCOPN (light paths to cern) 10Gbps.

This is the easy part its just a question of money.
The human and process requirements are more onerous.

Tier 1 SLA

- 99% uptime when beam on.
- 4 hour response to failures or degraded service (20%).
- A long term commitment to be a tier1.
- 12 h max delay to responding to operational problems.
- 12 h for network degradation (20%)

People Requirements

- 24x7 operation.
- storage specific experts
- experiment specific experts
- tiered support.
- operators on call.
- wlcg membership, meetings, meetings, meetings.

Outline

LHCOPN

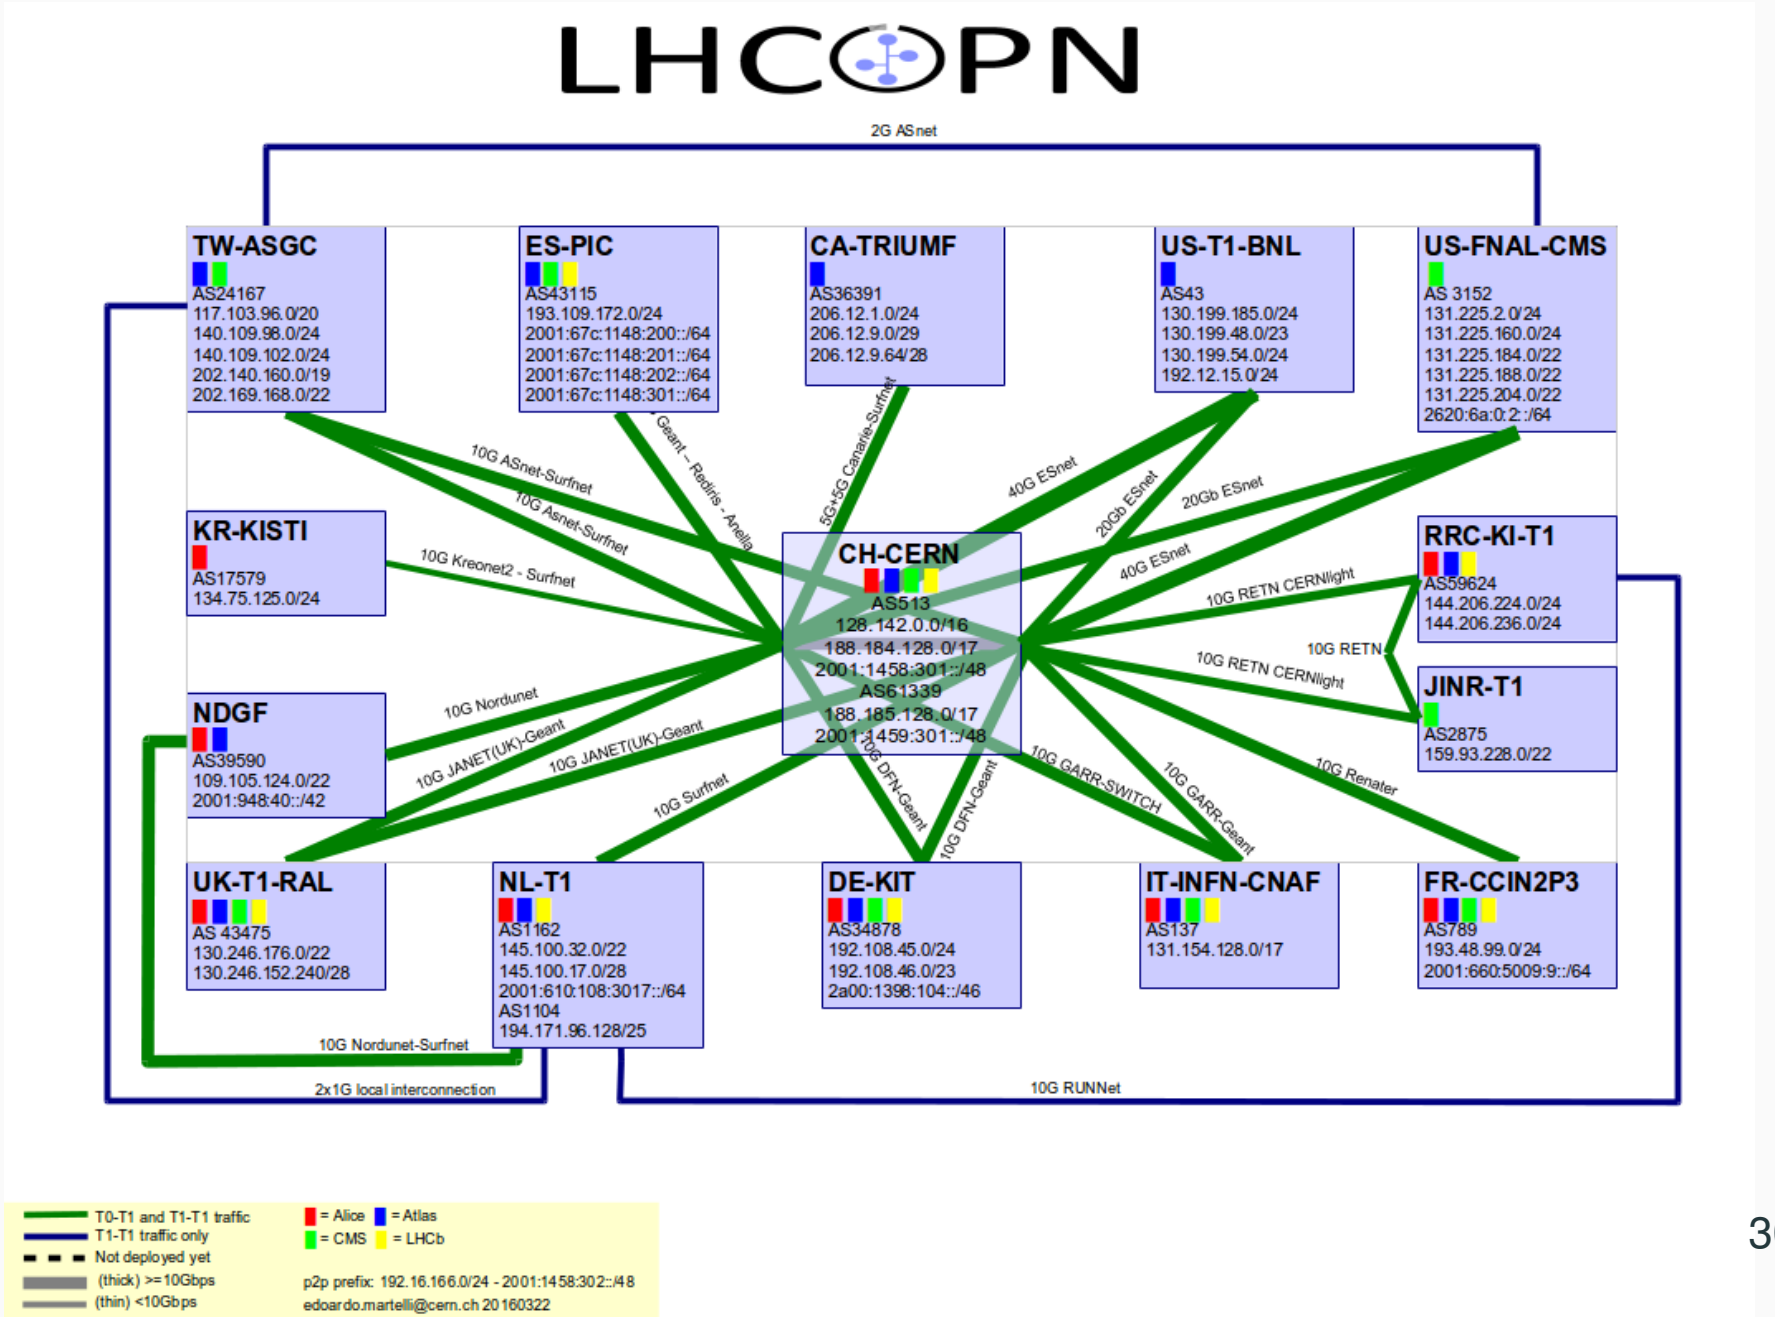
LHCOPN
2G ASnet
TW-ASGC
AS24167
117.103.96.0/20
140.109.98.0/24
140.109.102.0/24
202.140.160.0/19
202.169.168.0/22
ES-PIC
AS43115
193.109.172.0/24
2001:67c:1148:200::/64
2001:67c:1148:201::/64
2001:67c:1148:202::/64
2001:67c:1148:301::/64
CA-TRIUMF
AS36391
206.12.1.0/24
206.12.9.0/29
206.12.9.64/28
US-T1-BNL
AS43
130.199.185.0/24
130.199.48.0/23
130.199.54.0/24
192.12.15.0/24
US-FNAL-CMS
AS 3152
131.225.2.0/24
131.225.160.0/24
131.225.184.0/22
131.225.188.0/22
131.225.204.0/22
2620:6a:0:2::/64
KR-KISTI
AS17579
134.75.125.0/24
CH-CERN
AS513
128.142.0.0/16
188.184.128.0/17
2001:1458:301::/48
AS61339
188.185.128.0/17
2001:1459:301::/48
RRC-KI-T1
AS59624
144.206.224.0/24
144.206.236.0/24
JINR-T1
AS2875
159.93.228.0/22
NDGF
AS39590
109.105.124.0/22
2001:948:40::/42
UK-T1-RAL
AS 43475
130.246.176.0/22
130.246.152.240/28
NL-T1
AS1162
145.100.32.0/22
145.100.17.0/28
2001:610:108:3017::/64
AS1104
194.171.96.128/25
DE-KIT
AS34878
192.108.45.0/24
192.108.46.0/23
2a00:1398:104::/46
IT-INFN-CNAF
AS137
131.154.128.0/17
FR-CCIN2P3
AS789
193.48.99.0/24
2001:660:5009:9::/64
10G ASnet-Surfnet
10G Asnet-Surfnet
Geant - Rediris - Anella
5G+5G Canarie-Surfnet
40G ESnet
20Gb ESnet
20Gb ESnet
40G ESnet
10G Kreonet2 - Surfnet
10G RETN CERNlight
10G RETN
10G RETN CERNlight
10G Nordunet
10G JANET(UK)-Geant
10G JANET(UK)-Geant
10G Surfnet
10G DFN-Geant
10G DFN-Geant
10G GARR-SWITCH
10G GARR-Geant
10G Renater
10G Nordunet-Surfnet
2x1G local interconnection
10G RUNNet
T0-T1 and T1-T1 traffic
T1-T1 traffic only
Not deployed yet
(thick) >=10Gbps
(thin) <10Gbps
= Alice
= Atlas
= CMS
= LHCb
p2p prefix: 192.16.166.0/24 - 2001:1458:302::/48
edoardo.martelli@cern.ch 20160322

IOU table

three or more other institutes providing amongst them 52-week coverage).

The following parameters define the minimum levels of service. They will be reviewed by the operational boards of the WLCG Collaboration.

Service	*Maximum delay in responding to operational problems*			*Average availability[6] measured on an annual basis*	
	Service interruption	Degradation of the capacity of the service by more than 50%	Degradation of the capacity of the service by more than 20%	During accelerator operation	At all other times
Raw data recording	4 hours	6 hours	6 hours	99%	n/a
Event reconstruction or distribution of data to Tier-1 Centres during accelerator operation	6 hours	6 hours	12 hours	99%	n/a
Networking service to Tier-1 Centres during accelerator operation	6 hours	6 hours	12 hours	99%	n/a
All other Tier-0 services	12 hours	24 hours	48 hours	98%	98%
All other services[7] – prime service hours[8]	1 hour	1 hour	4 hours	98%	98%
All other services[7] – outwith prime service hours[8]	12 hours	24 hours	48 hours	97%	97%

Annex 3.2. Tier-1 Services

Each Tier1 Centre[9] forms an integral part of the central data handling service of the LHC Experiments. It is thus essential that each such centre undertakes to provide its services on a long-term basis (initially at least 5 years) and to make its best efforts to upgrade its installations steadily in order to keep pace with the expected growth of

mou table

reviewed by the operational boards of the WLCG Collaboration.

Service	*Maximum delay in responding to operational problems*			*Average availability[6] measured on an annual basis*	
	Service interruption	Degradation of the capacity of the service by more than 50%	Degradation of the capacity of the service by more than 20%	During accelerator operation	At all other times
Acceptance of data from the Tier-0 Centre during accelerator operation	12 hours	12 hours	24 hours	99%	n/a
Networking service to the Tier-0 Centre during accelerator operation	12 hours	24 hours	48 hours	98%	n/a
Data-intensive analysis services, including networking to Tier-0, Tier-1 Centres outwith accelerator operation	24 hours	48 hours	48 hours	n/a	98%
All other services[7] – prime service hours[10]	2 hour	2 hour	4 hours	98%	98%
All other services[7] – outwith prime service hours[10]	24 hours	48 hours	48 hours	97%	97%

The response times in the above table refer only to the maximum delay before action