Operations in Italy and plans for INFN sites

Domenico Elia, Stefano Bagnasco,

Stefano Piano

INFN sez. Trieste

Outline of the talk

- Operations organization
- Networking status
- Resources available and evolution
- The Tier-1 at INFN-CNAF
- Status of Tier-2 sites
- Mandatory plots
- The STOA LHC project

Operations Organization

- ALICE-IT Computing Coordination:
  - Domenico Elia
- Deputy:
  - Stefano Bagnasco
- Tier-2 Operations Coordination:
  - Stefano Piano
  - Monthly phone conference for coordination and performance monitoring
  - Yearly face-to-face workshop (2012 @ Catania, 2013 @ Trieste, 2014 @ Frascati)
- Monthly Tier-1 Management Board at CNAF

Networking

- Tier-1 at 40 Gbps (LHCONE + LHCOPN) + 10 Gbps (General purpose)
- All Tier-2's connected to LHCONE with at least 10 Gbps
  - Through GARR-X
  - Most of Tier-2's easily upgradable to 40 Gbps
  - Padova-LNL already at 20 Gbps
  - Bari, Catania 40 Gbps in few months
- IPV6
  - All INFN sites will act coordinately at the same point but no exact estimate yet

Resources available for Alice

- **Tier-1 at CNAF, Bologna**
  - Shared with other LHC experiments and a large amount of others
- **4 official Tier-2 centers**
  - Bari, Catania, Padova-LNL and Torino
  - Official means directly funded by INFN according to plans and official pledges
- **Additional (minor) centers**
  - Bologna, Cagliari, Trieste
  - Local resources, different creative funding, mostly out of pledge (but Cagliari)
  - Bologna going to be removed from the list of ALICE sites (shortage of resources and manpower)
- Project providing resources in the ALICE INFN sites over the last years
  - **ReCaS (BA and CT)**, sizeable contribution to 2014 and 2015 pledges
  - **CyberSar (CA) and TriGrid (CT)**, both ended, resources becoming obsolescent

Computing resources at INFN T1/T2s

Current Status

- Tier-1

  CPU: 20900 HS06 / **DISK**: 1.9 PB / **TAPE**: 3.7 PB

- Tier-2

  CPU: 37050 HS06 / **DISK**: 3.1 PB

	Bari	Catania	LNL-Padova	Torino	Cagliari	Total
HS06	8264	10757	8264	7805	1960	37050
TB	812	683	664	814	70	3043

- **Further notes:**
  - dismissions due in 2014 fully included in the above table
  - for Bari and Catania: 2014 new resources available on site but will enter in production within ~1 month from now (infrastructure update, ReCaS funding)

Computing resources at INFN T1/T2s

Expected in 2015

- Tier-1

  CPU: 22800 HS06 / **DISK**: 3.4 PB / **TAPE**: 4.2 PB

- Tier-2

  CPU: 38600 HS06 / **DISK**: 4.4 PB

	Bari	Catania	LNL-Padova	Torino	Cagliari	Total
HS06	9104	12307	8264	7805	1120	38600
TB	1062	1183	1052	1064	20	4381

- **Dismissions and procurements for 2015:**
  - global resource budget corresponding to pledge (as in the previous years)
  - **CPU 2015:** **7710 HS06** (including 6160 HS06 for replacements)
  - **Storage 2015:** **1840 TB** (including 550 TB for replacements)

Computing resources at INFN T1/T2s

Perspectives for 2016

As from RRB October 2014

CPU (kHEPSPEC06)				
	Tier0	CAF	Tier1s	Tier2s
2015	130	45.0	120	200
2016	170	45.0	160	240
2017	200	45.0	210	270

Disk (PB)				
	Tier0[1)]	CAF	Tier1s[2)]	Tier2s
2015	11.1	0.34	15.4	22.7
2016	13.4	0.44	18.6	26.1
2017	15.7	0.54	21.8	30.7

Tape (PB)		
	Tier0	Tier1
2015	16.2	10.2
2016	21.6	15.6
2017	25.7	19.7

- INFN share ~19%
- Increase wrt previous years for Tier-1/2 resources:
  - 2016: +25% CPU, +17% DISK, +52% TAPE
  - 2017: +20% CPU, +17% DISK, +26% TAPE
  - Slightly above “flat budget”, going to be discussed to the RRB in April

Tier-1 at CNAF

- 180 kHEPSpec06 overall
- ALICE share is 22800 HS06 (about 1600 job slots)
  - LSF for queue management
  - Few WNs virtualized in a cloud-like architecture (“Worker Nodes On Demand”, WNODes) but not used by ALICE
- 18 PB-N of disk and 35 PB-N of tape overall
- ALICE share is 3.4 PB-N (T1D0+T0D1) plus tapes for 4.2 PB-N (500 TB-N used by ALICE)
  - GPFS + TSM for management
  - Xrootd as a front-end protocol

Tier-2 sites
General remarks

- All sites can allow for more resources coming in without big infrastructural investments although manpower is tight
- New Data Centers (HPC, Grid and Cloud):
  - Bari and Catania are becoming rather large
  - funding for infrastructural upgrades provided by special project (ReCas)
- All sites working to expand support to more VOs beyond LHC ones, to allow for resource optimization (e.g. Padova-LNL with OpenStack based cloud infrastructure and Torino … see Stefano's talk)

Tier-2 sites
Specific remarks

- Bari:
  - Probably migration from Lustre FS to pure Xrootd
  - 10 Gbps link GARR-X/LHC-ONE (upgrade to 40 Gbps)
  - Total Storage: 812 TB, coming soon new resources
- Catania:
  - Added 950 TB Total Storage 1200 TB (GPFS)
  - 10 Gbps link GARR-X/LHC-ONE (upgrade to 40 Gbps)
  - Coming soon:
    - New core switch HP 10508
    - 34 Supermicro with 4 CPUs x 16 Cores (2x10Gbps)
    - Openstack Juno (Distributed Virtual Routing)
    - Compute Node with 4/8 WN - 16/8 Cores – 64/32 GB RAM

Tier-2 sites
Specific remarks

- Padova-LNL:
  - Split across two physical sites (INFN site at Padova University and INFN's Legnaro National Laboratories)
  - Dedicated 10 Gbps link between the sites
  - 2 x 10 Gbps link GARR-X/LHCONE
  - Total Storage 664 TB (pure xrootd)
  - Opportunistic utilization of CMS resources (20200 HS06)
- Torino:
  - Private Cloud Infrastracture hosts full-fledged Tier2 site
  - OpenNebula software stack
  - "Elastic provisioning" of Virtual Analysis Facility
  - 9(+1) Gbps link GARR-X/LHC-ONE
  - Total Storage 550 TB (+300 TB soon) managed with GlusterFS

Tier-1 and Tier-2 sites
Running jobs profile

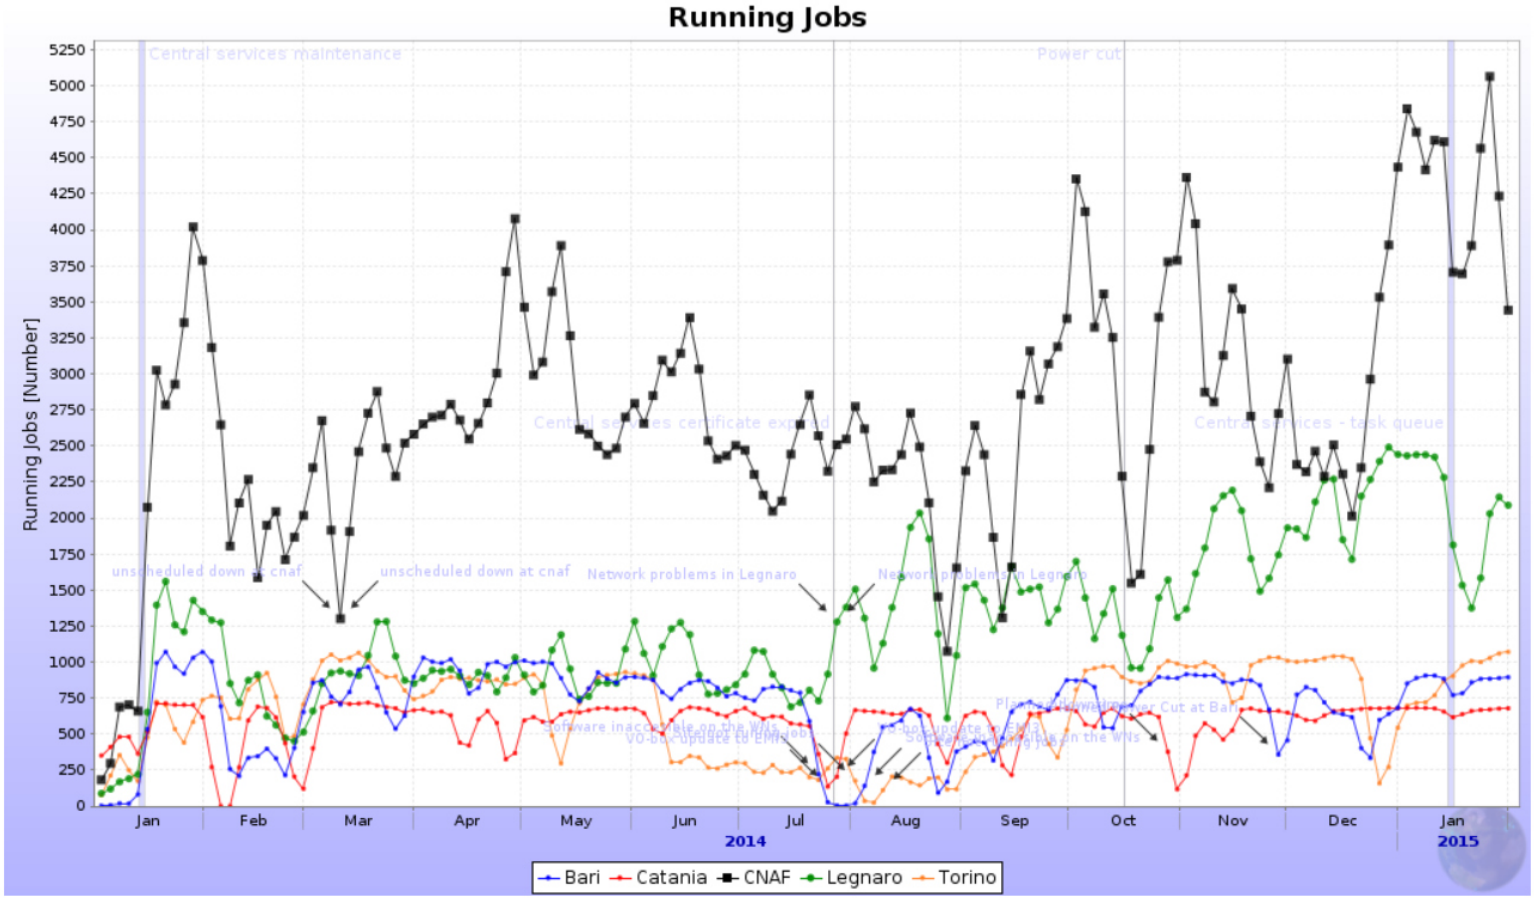

Tier-2 sites
Running jobs profile

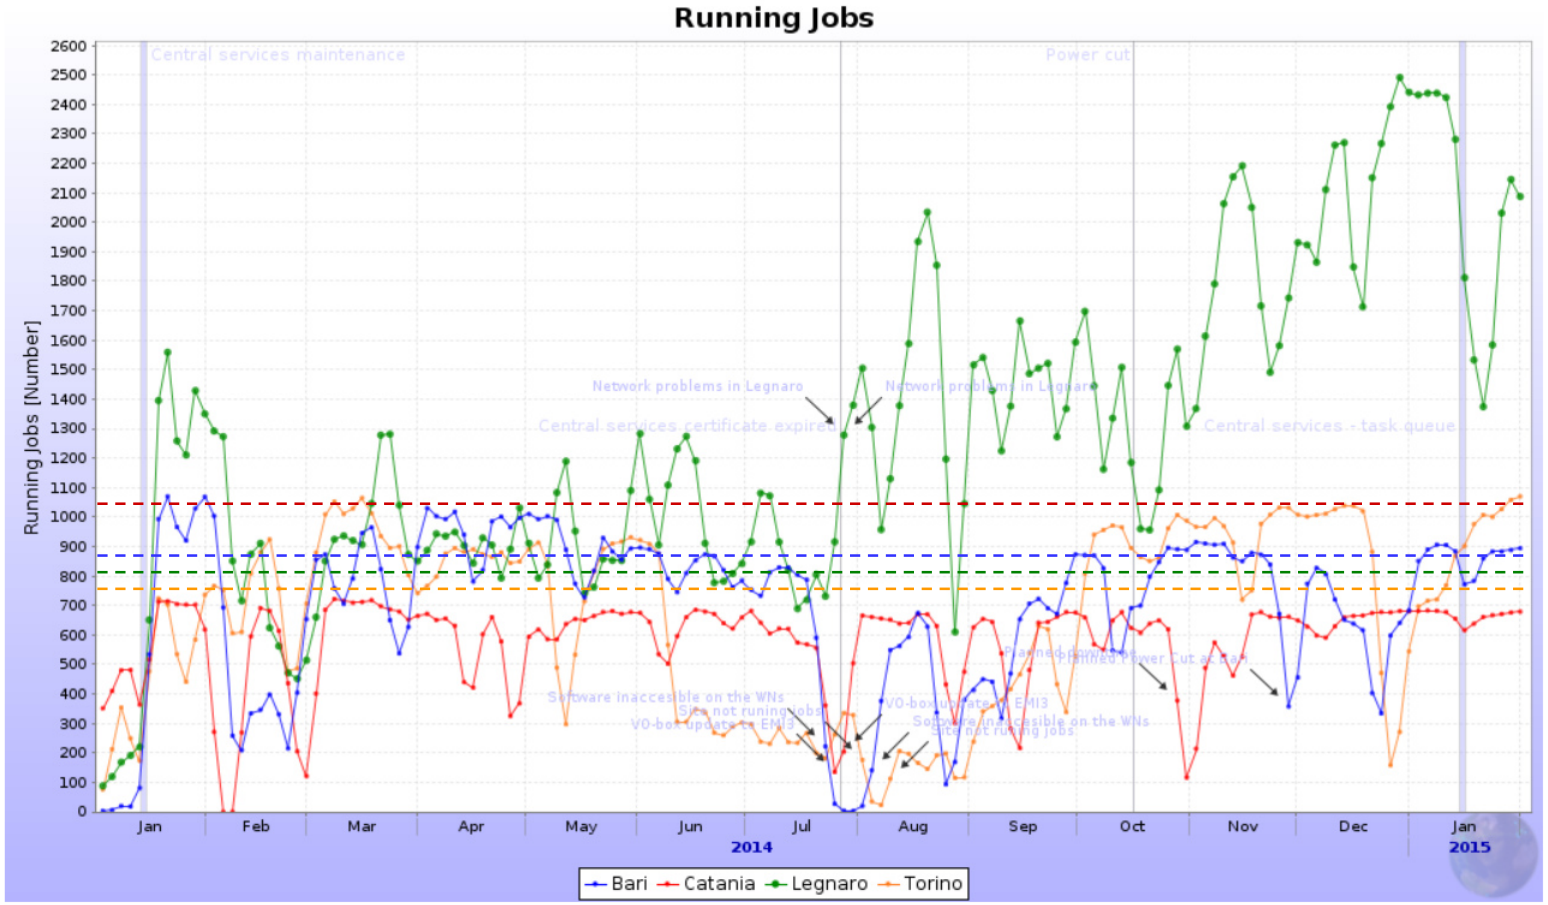

Pledge:

Catania
Bari
PD-LNL
Torino

Tier-1 and Tier-2 sites
WallTime PieChart

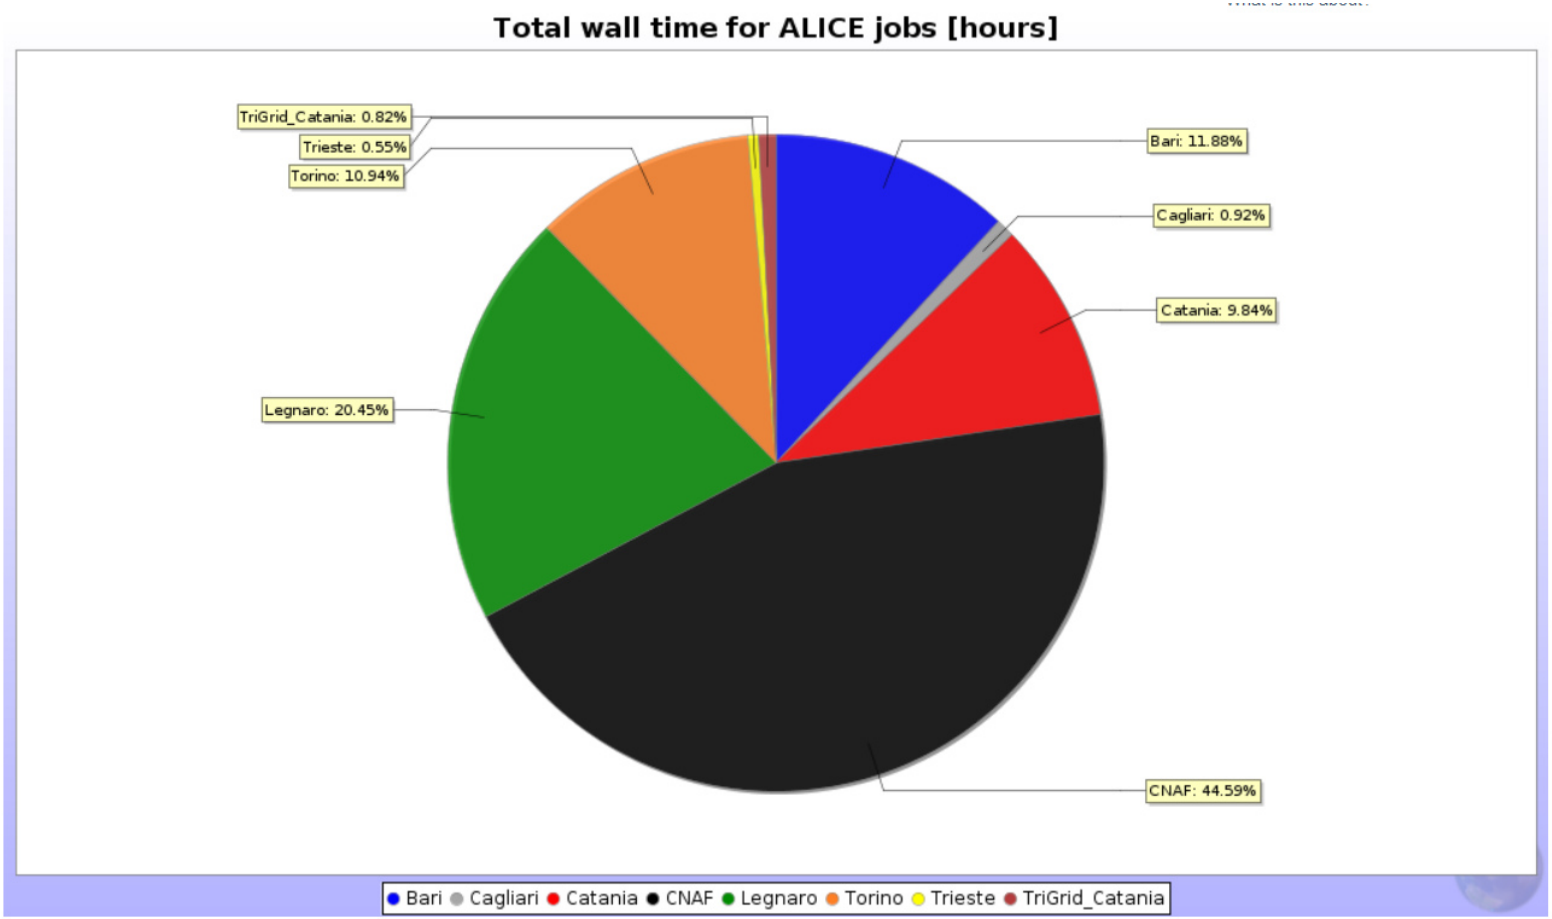

Tier-1 and Tier-2 sites
WallTime PieChart

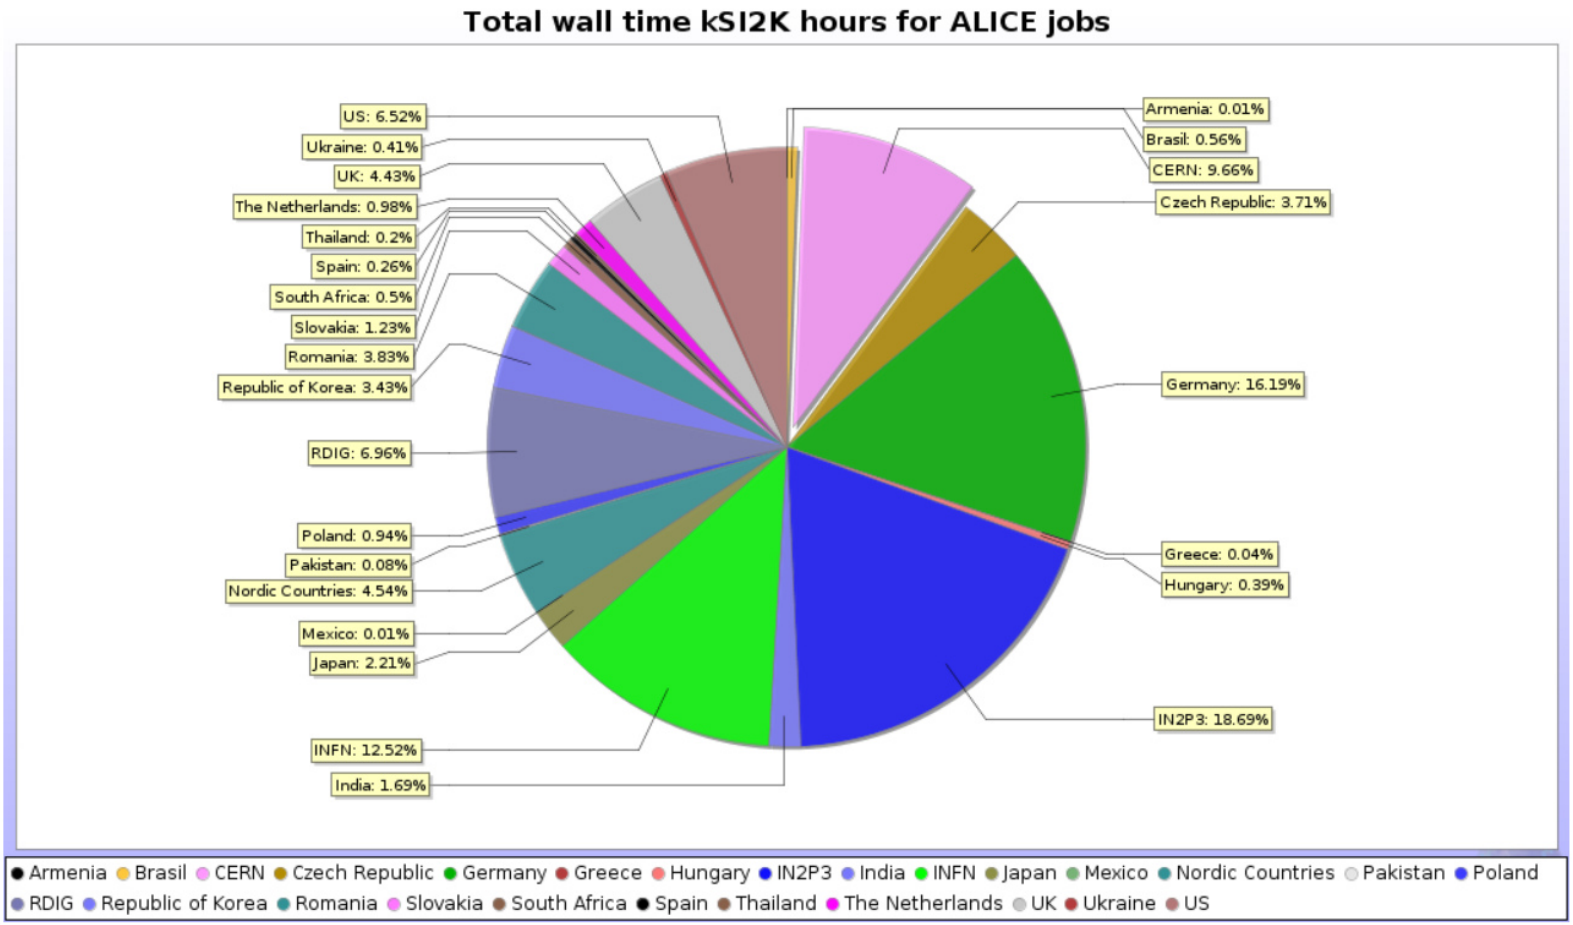

Tier-1 and Tier-2 sites
CPU Efficiency

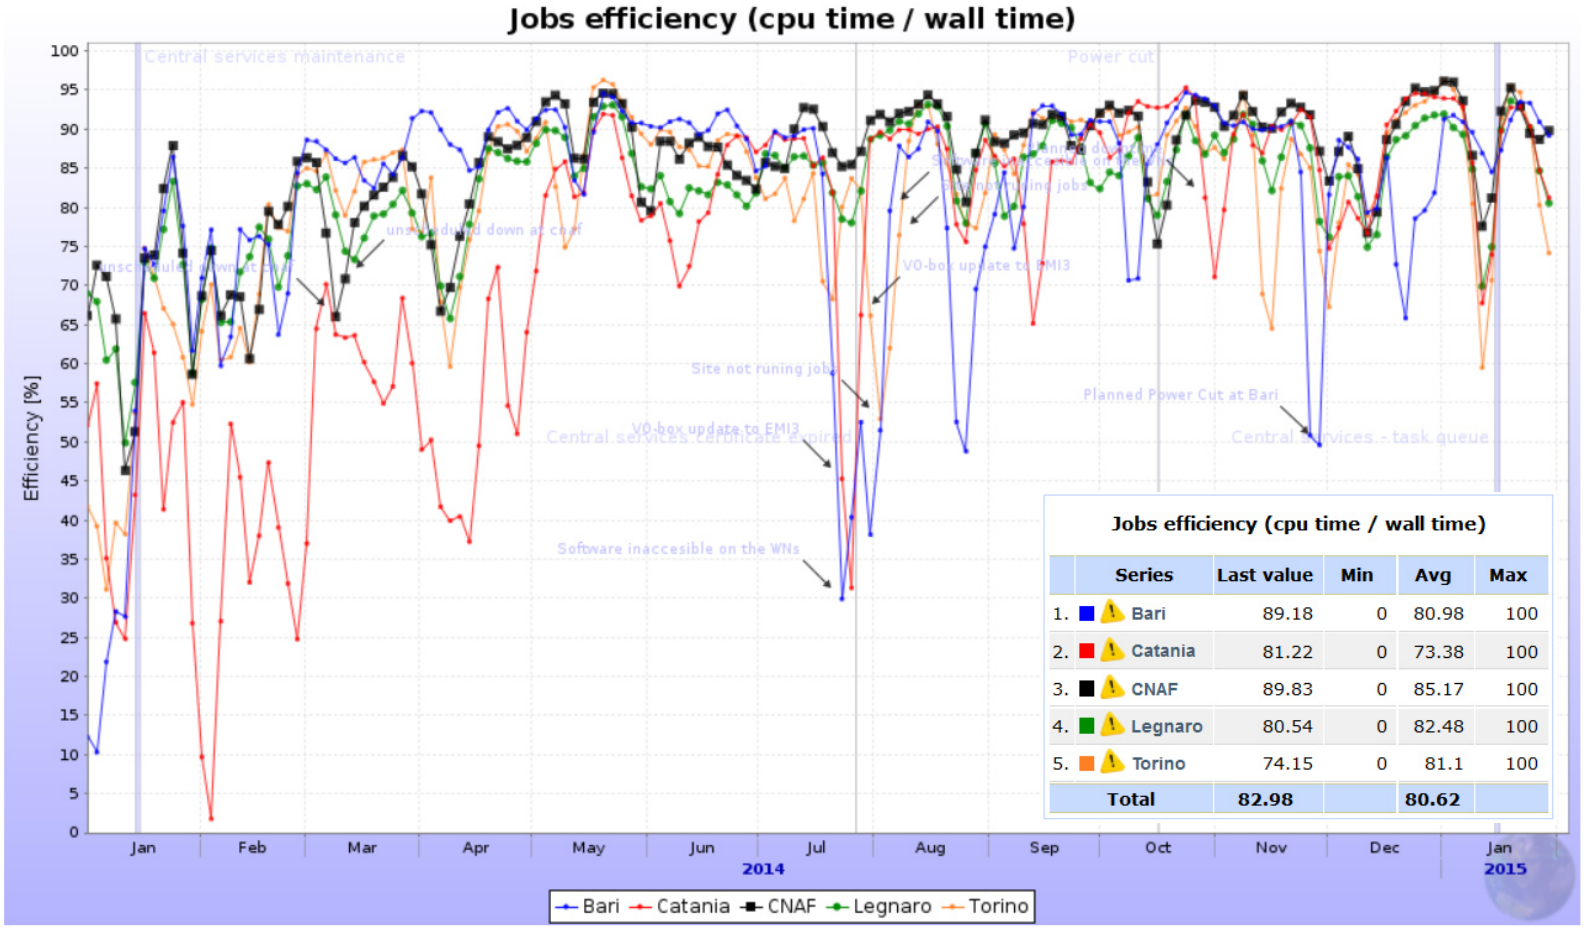

Jobs efficiency (cpu time / wall time)

	Series	Last value	Min	Avg	Max
1.	Bari	89.18	0	80.98	100
2.	Catania	81.22	0	73.38	100
3.	CNAF	89.83	0	85.17	100
4.	Legnaro	80.54	0	82.48	100
5.	Torino	74.15	0	81.1	100
Total		82.98		80.62	

Tier-1 and Tier-2 sites
SE Availability

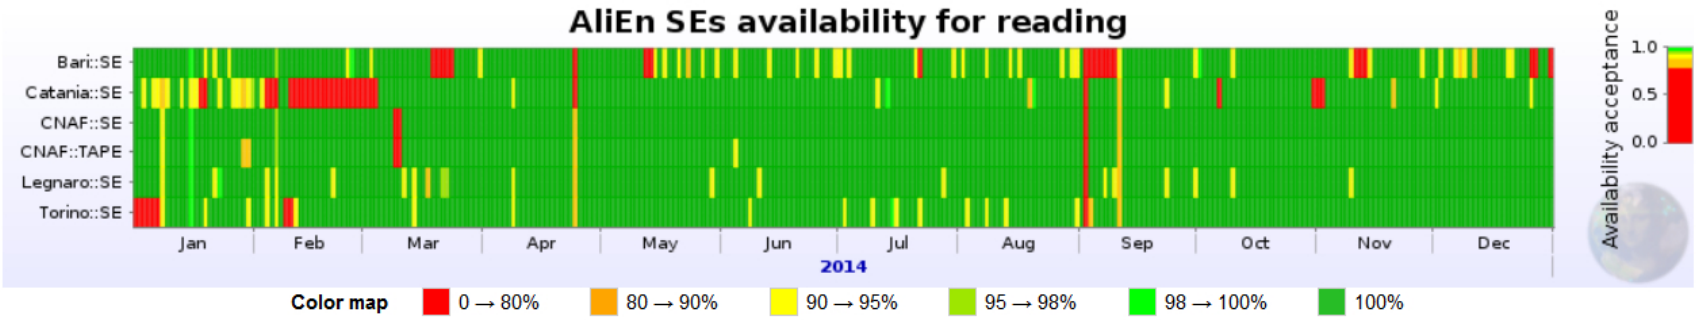

Statistics

Link name	Data		Individual results of reading tests			Overall
	Starts	Ends	Successful	Failed	Success ratio	Availability
Bari::SE	01 Jan 2014 00:00	01 Jan 2015 00:00	4082	283	93.52%	93.66%
Catania::SE	01 Jan 2014 00:03	01 Jan 2015 00:03	4067	298	93.17%	93.24%
CNAF::SE	01 Jan 2014 00:07	01 Jan 2015 00:07	4334	30	99.31%	99.40%
CNAF::TAPE	01 Jan 2014 00:08	01 Jan 2015 00:07	4327	37	99.15%	99.24%
Legnaro::SE	01 Jan 2014 00:20	01 Jan 2015 00:18	4330	33	99.24%	99.39%
Torino::SE	01 Jan 2014 00:37	31 Dec 2014 22:32	4243	119	97.27%	97.42%

All Italian sites
SE Availability

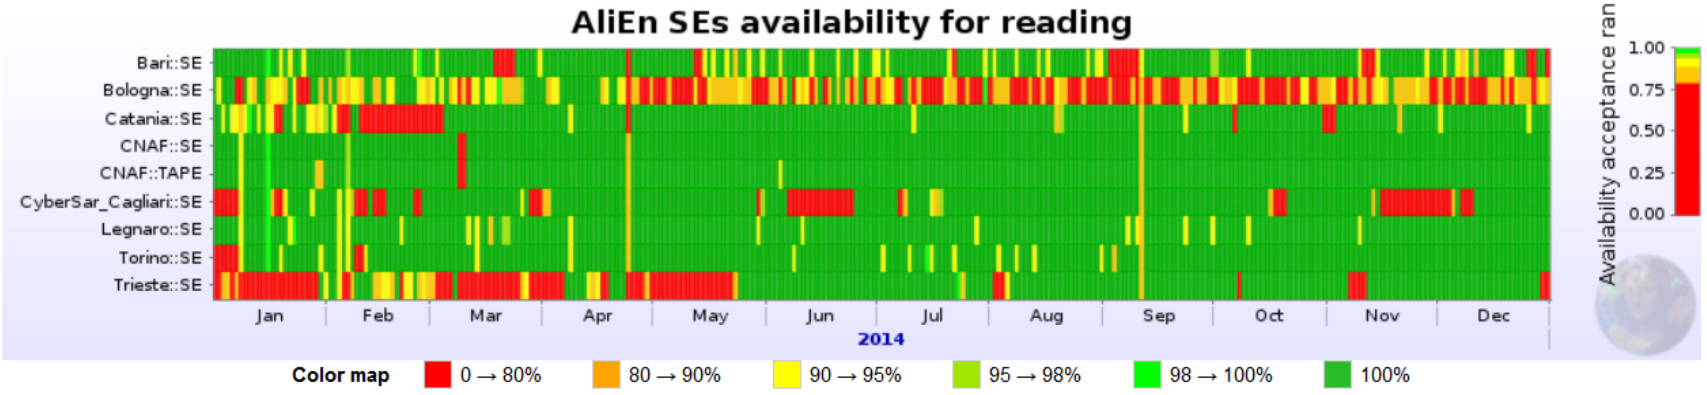

Statistics

Link name	Data		Individual results of reading tests			Overall
	Starts	Ends	Successful	Failed	Success ratio	Availability
Bari::SE	01 Jan 2014 00:00	01 Jan 2015 00:00	4081	281	93.56%	93.66%
Bologna::SE	01 Jan 2014 00:02	01 Jan 2015 00:02	3536	825	81.08%	81.81%
Catania::SE	01 Jan 2014 00:03	01 Jan 2015 00:03	4066	294	93.26%	93.31%
CNAF::SE	01 Jan 2014 00:07	01 Jan 2015 00:07	4333	27	99.38%	99.47%
CNAF::TAPE	01 Jan 2014 00:08	01 Jan 2015 00:07	4326	34	99.22%	99.31%
CyberSar_Cagliari::SE	01 Jan 2014 00:13	01 Jan 2015 00:08	3664	697	84.02%	84.17%
Legnaro::SE	01 Jan 2014 00:20	01 Jan 2015 00:18	4331	30	99.31%	99.46%
Torino::SE	01 Jan 2014 00:37	31 Dec 2014 22:32	4243	116	97.34%	97.49%
Trieste::SE	01 Jan 2014 00:33	31 Dec 2014 22:37	3301	1058	75.73%	75.64%

All Italian sites
SE Availability

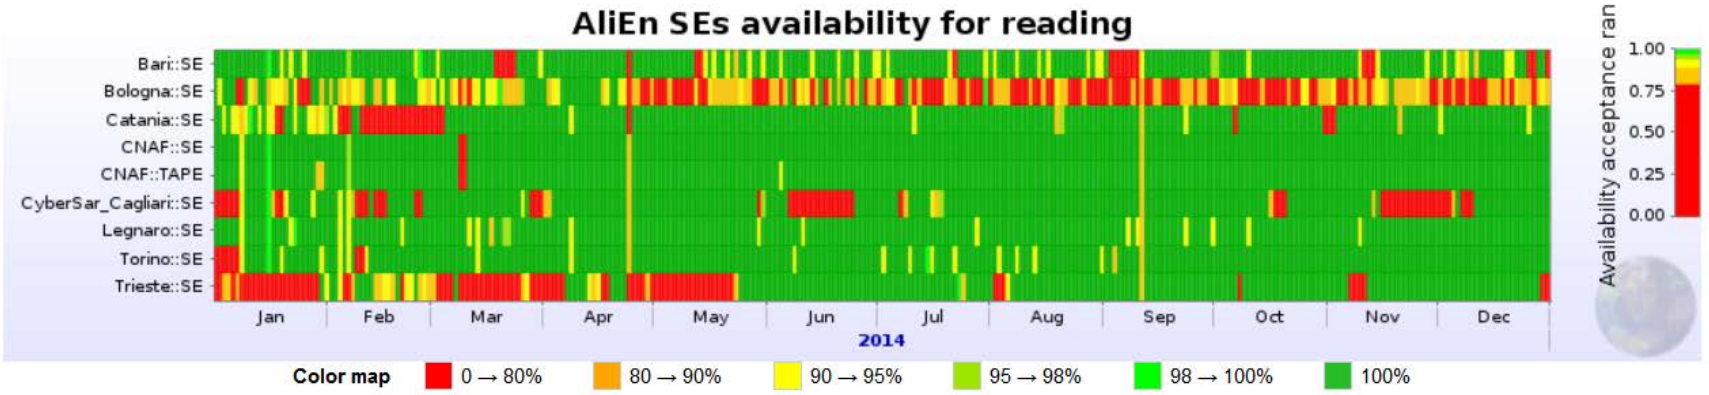

Bologna (13 TB) : no longer maintained
proposal to remove it from the list of ALICE sites

Cagliari (71 TB) : suffered from the power instabilities of the campus

Trieste (26 TB) : test failed because SE was full (100%) and resolved a long-standing metadata issue

Almost all Italian sites The STOA LHC Project

- Three-year National Research Project (“PRIN”) approved for 2013-2015
  - STOA-LHC: Optimization of data access, network and interactive data analysis for LHC experiments
- Most ALICE-relevant activities focused on resource federation (xrootd and clouds) and interactive analysis (PROOF)
  - Activities on interactive analysis on cloud infrastructures (Torino, Bari, PD-LNL, Trieste), optimization of data access (Bari), development of a Science Gateway for ALICE (Catania)
- Post-doc positions in Torino, Bari, Legnaro and Trieste
- See Sara’s talk for the current achievements

Conclusions

Thanks

for

your attention !

Resources available for Alice

- **Tier-1 at CNAF, Bologna**
  - Shared with other LHC experiments and a large amount of others
- **4 official Tier-2 centers**
  - Official means directly funded by INFN according to plans and official pledges
  - **Bari, Catania, Padova-LNL and Torino**
- **Bologna, Cagliari, Trieste**
  - Local resources, different creative funding
- **CyberSar (CA) and TriGrid (CT)**
  - Both projects ended, resources becoming obsolescent