ATLAS activities on FR cloud

欢迎

歡迎

bun venit
bienvenue
Welcome

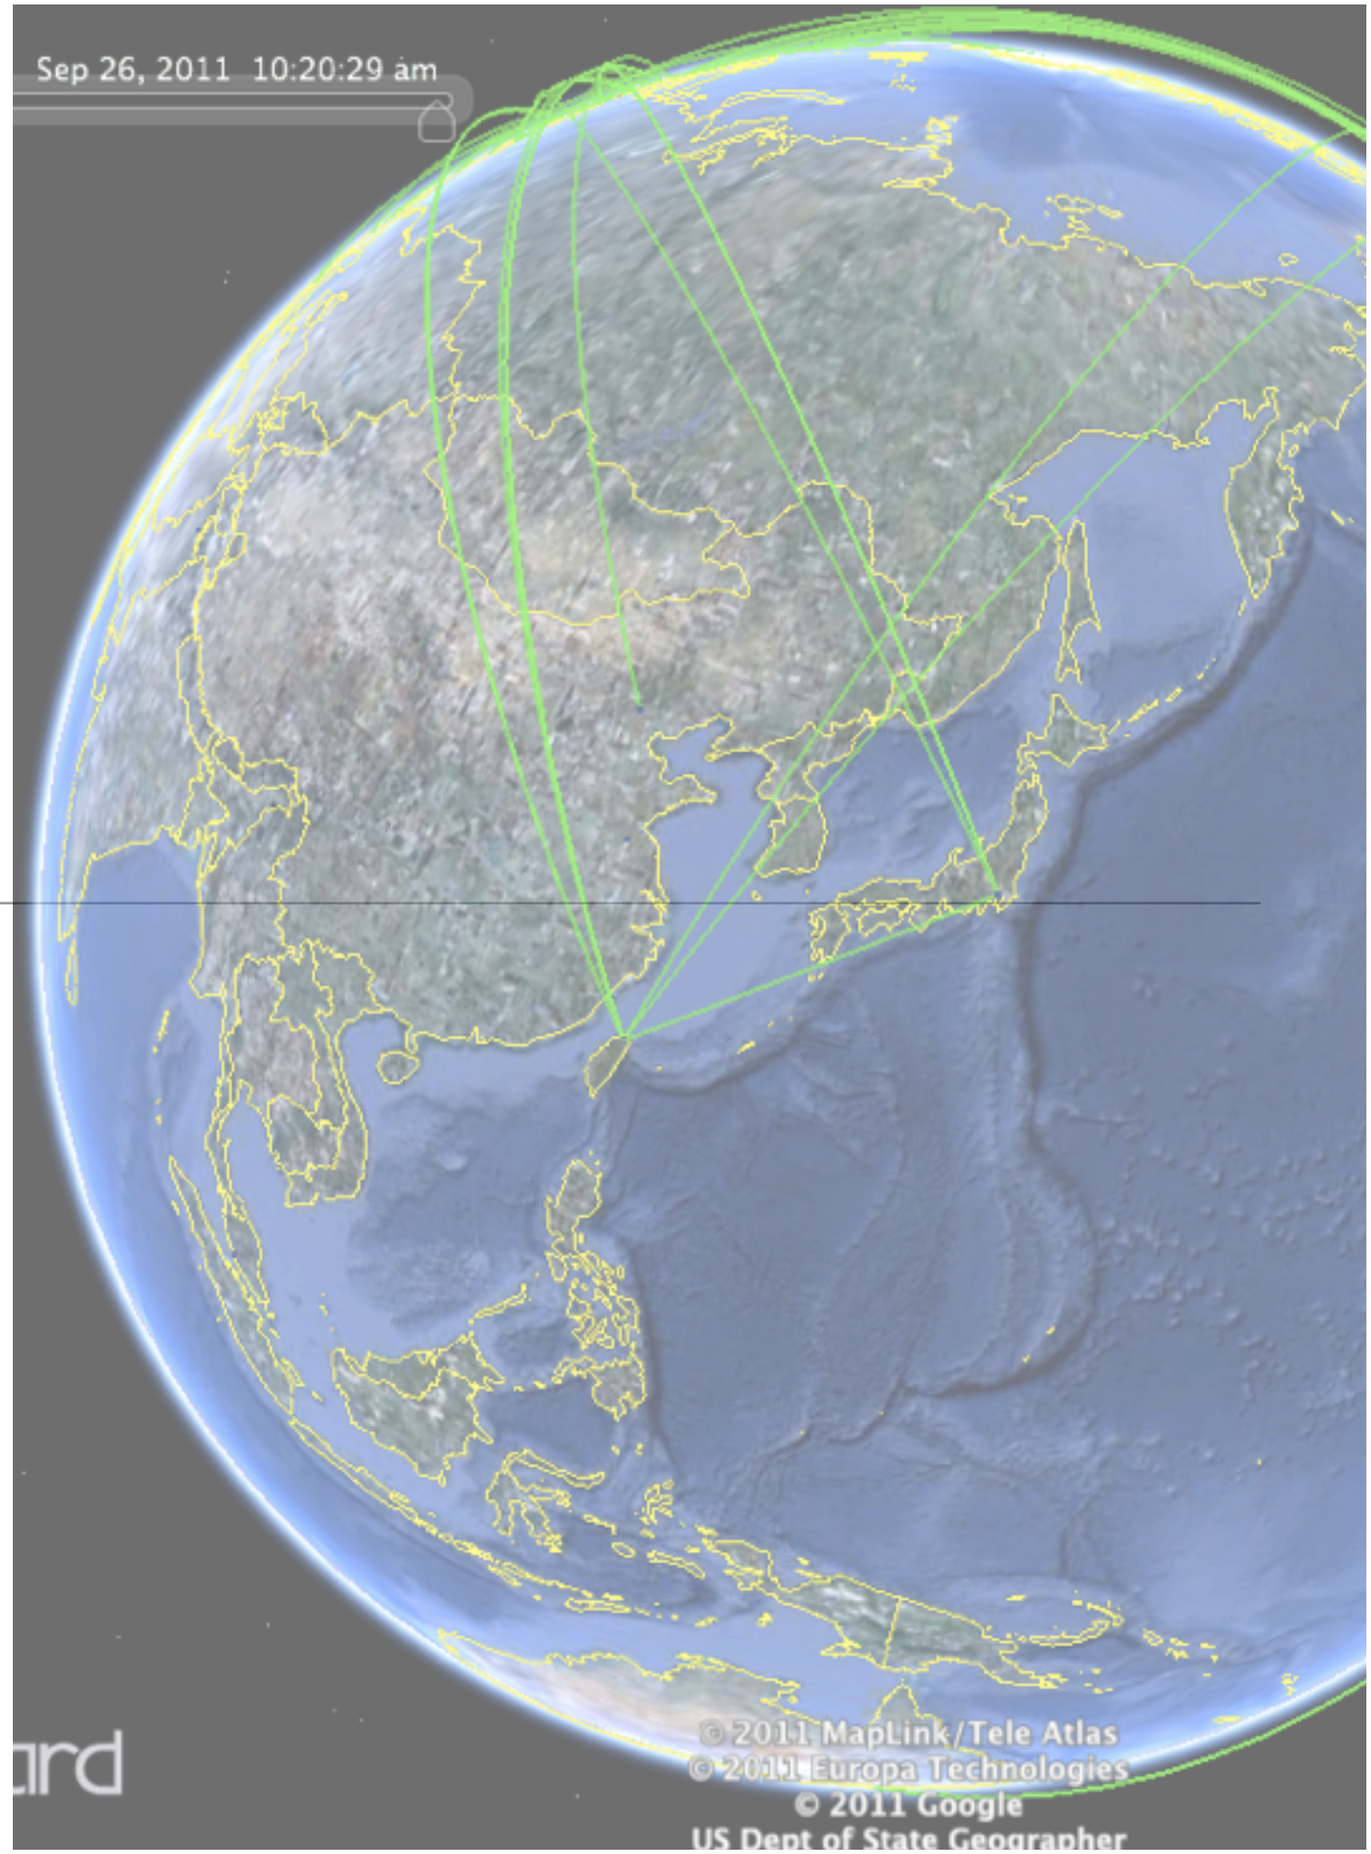

Some items about FR-cloud

- Who we are
- What we did over the last year
- Operation issues
- Performances
- Common procedures / communication (Squad/Luc)
- Network (Monday)

Items for discussions (random order) during the workshop

- ATLAS related site issues
- Network monitoring and tests
- T2Ds
- multi-core
- Xrootd federation
- cloud computing
- Monitoring
- CVMFS
- Squid/FronTier

irfu

cea

saclay

ATLAS activities [May 2011 - May 2012]

- Heavy data & MC processing

Data & MC processing jobs

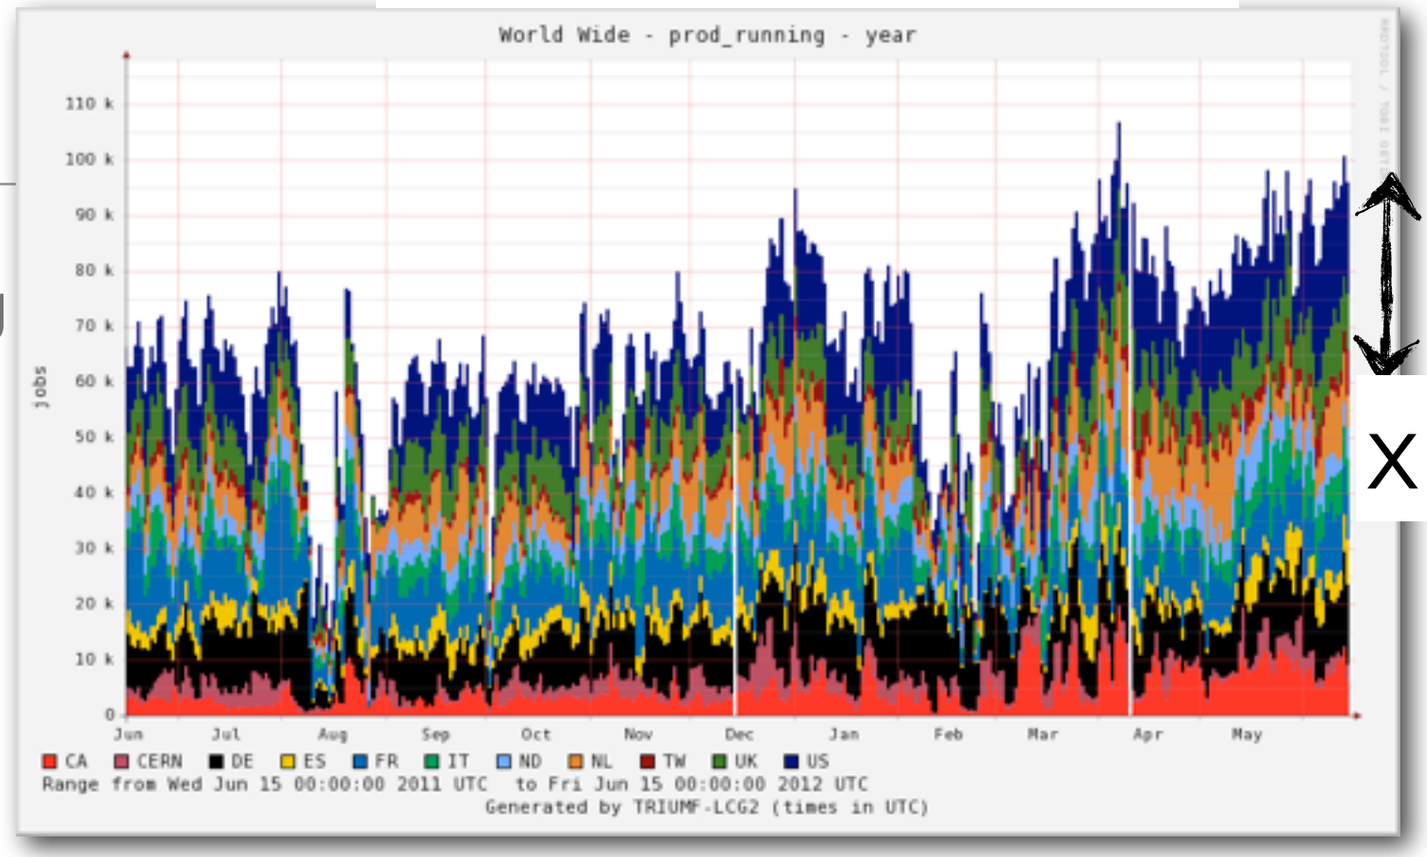

x1.5

- Increasing load of analysis

Analysis jobs

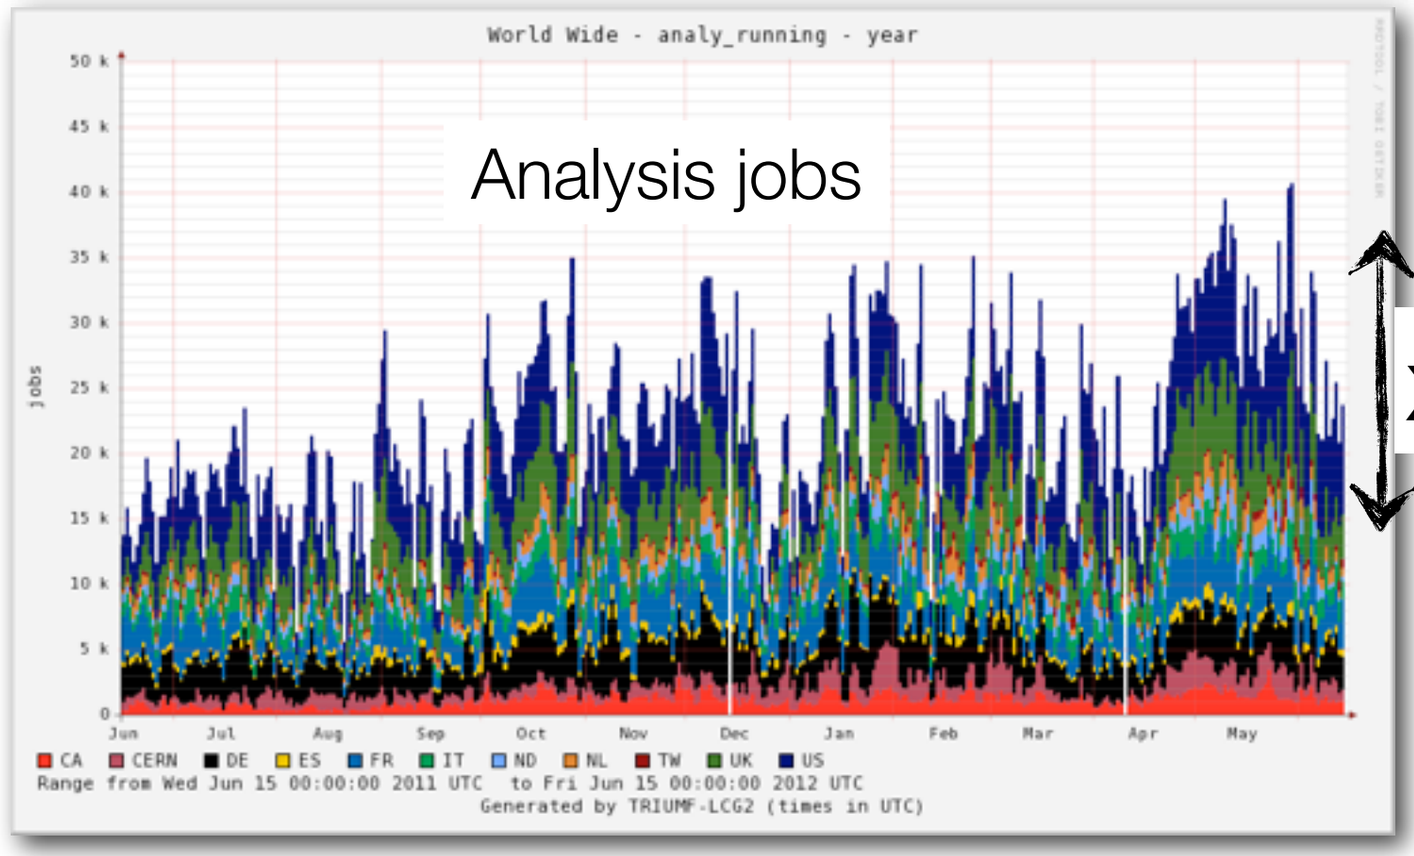

x3

irfu

saclay

The “French” cloud

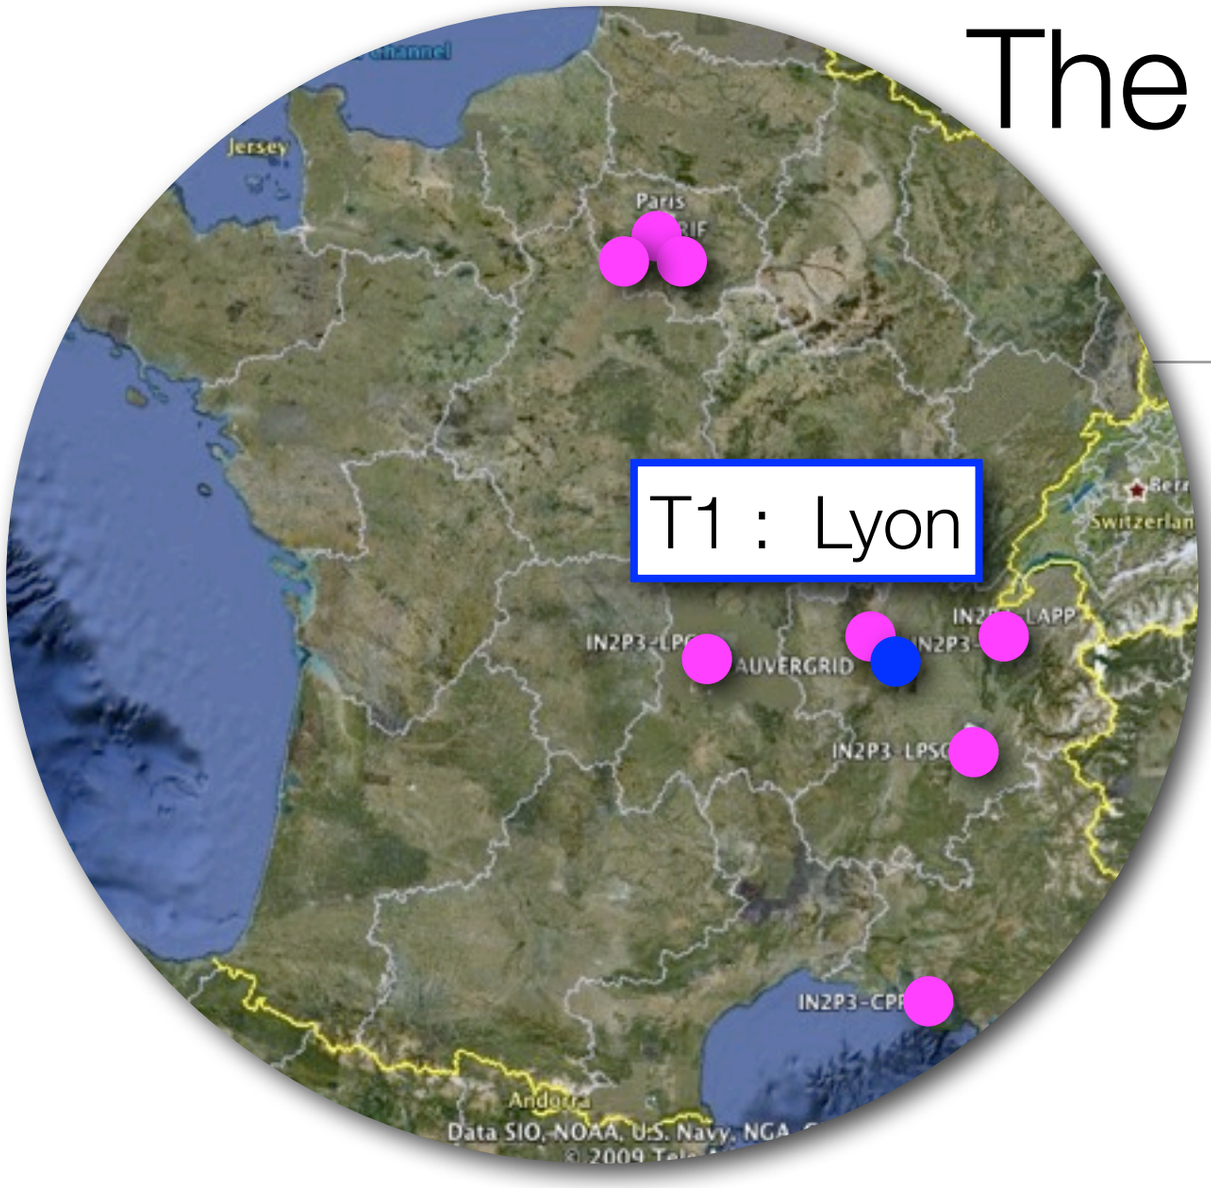

T2s : 14 sites

- Annecy
- Clermont
- Grenoble
- Grif (3 sites)
- Lyon
- Marseille
- Beijing
- Romania x4
- Tokyo

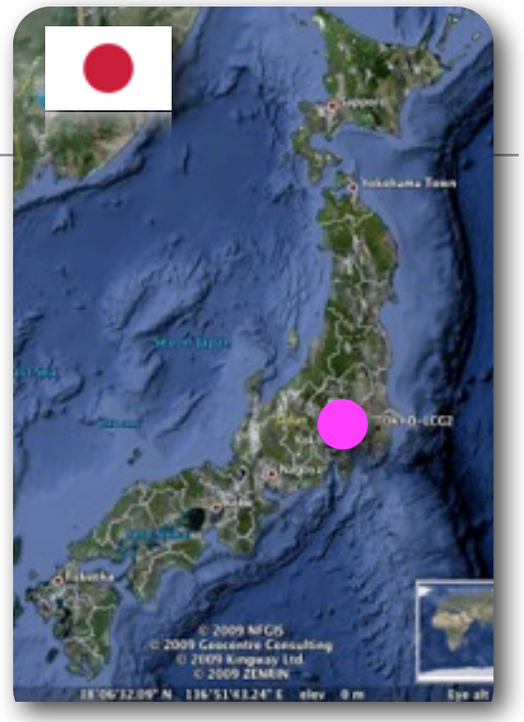

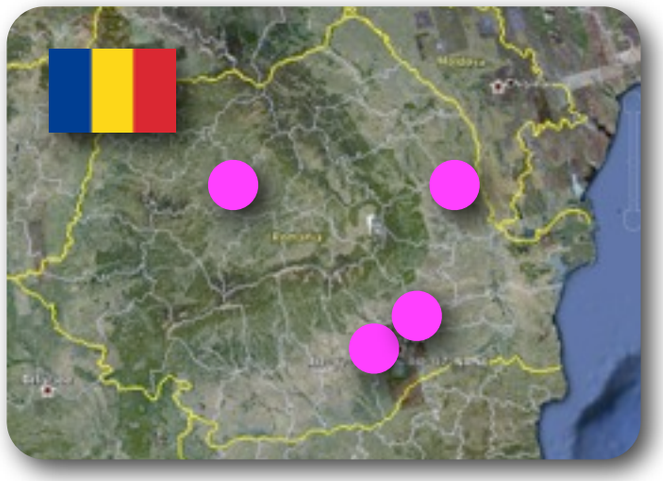

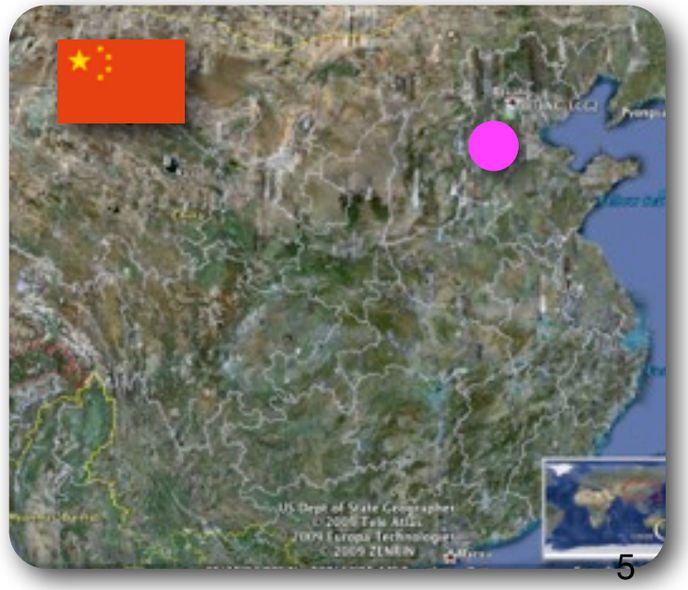

T3 :

- 2 in 2010 (LPSC, CPPM)
- 0 in 2011, T3s became T2s

irfu

cea

saclay

Most of FR T2s are connected to LHCONE

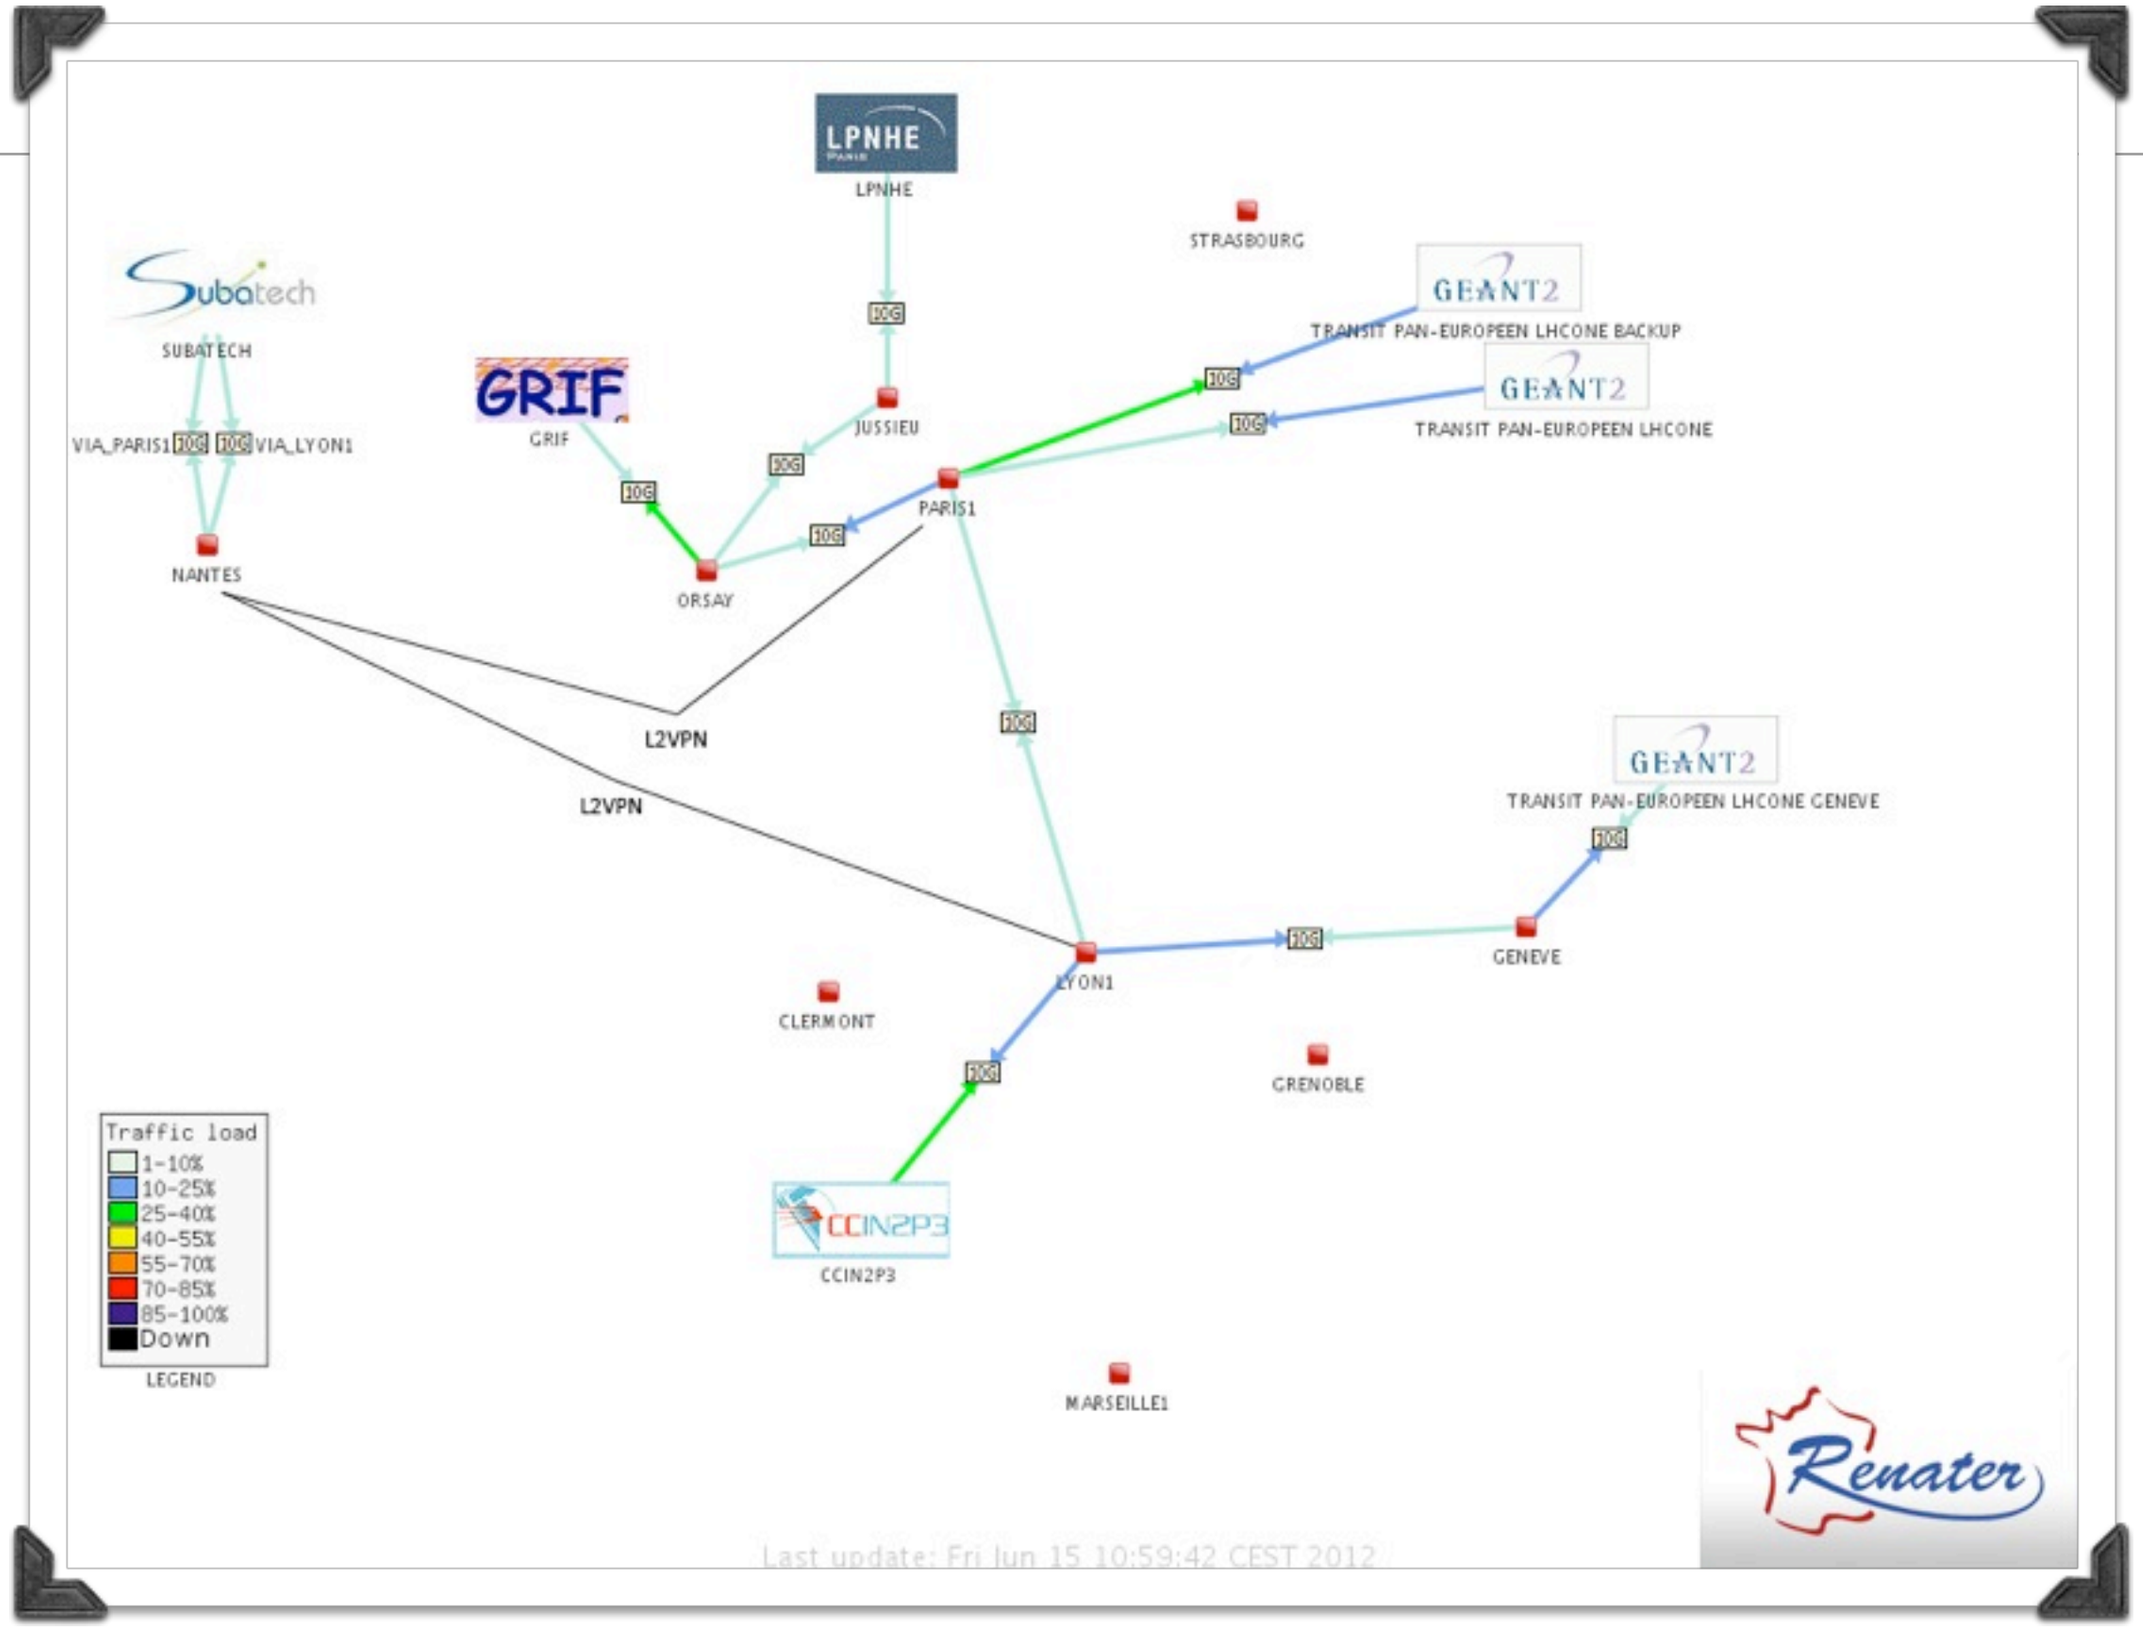

Cloud main features

- 4 countries
- Several time zone
- Large RTT spread

communication

- Beijing ~190 ms
- Tokyo ~290 ms
- Romanian sites ~55ms
- Grif ~6 ms

Network

irfu

cea

saclay

Cloud organization

- **CAF** (ATLAS France Computing)
  - One (at least) representative per French laboratory
  - ‘Experts’ from software & distributed computing
- ATLAS T1 team (Ghita, Manoulis)
- Squad-FR
  - Team : 4 people, rotating ~each week
  - S. Crepe (LPSC), Ch. Beau (LPNHE), E. Leguiriec (CPPM), L. Poggioli (LAL)
- Regular CAF meetings :
  - ≈ monthly meeting at Lyon
  - Review cloud activities
  - Part of the meeting dedicated to T1
  - Part dedicated to T2s (with T2 representatives)
- Monthly LCG-France technical meetings

irfu

cea

saclay

ATLAS resources at Lyon

- ATLAS : **45%** of LCG-France budget
- 2012 : LCG-France budget cut again...

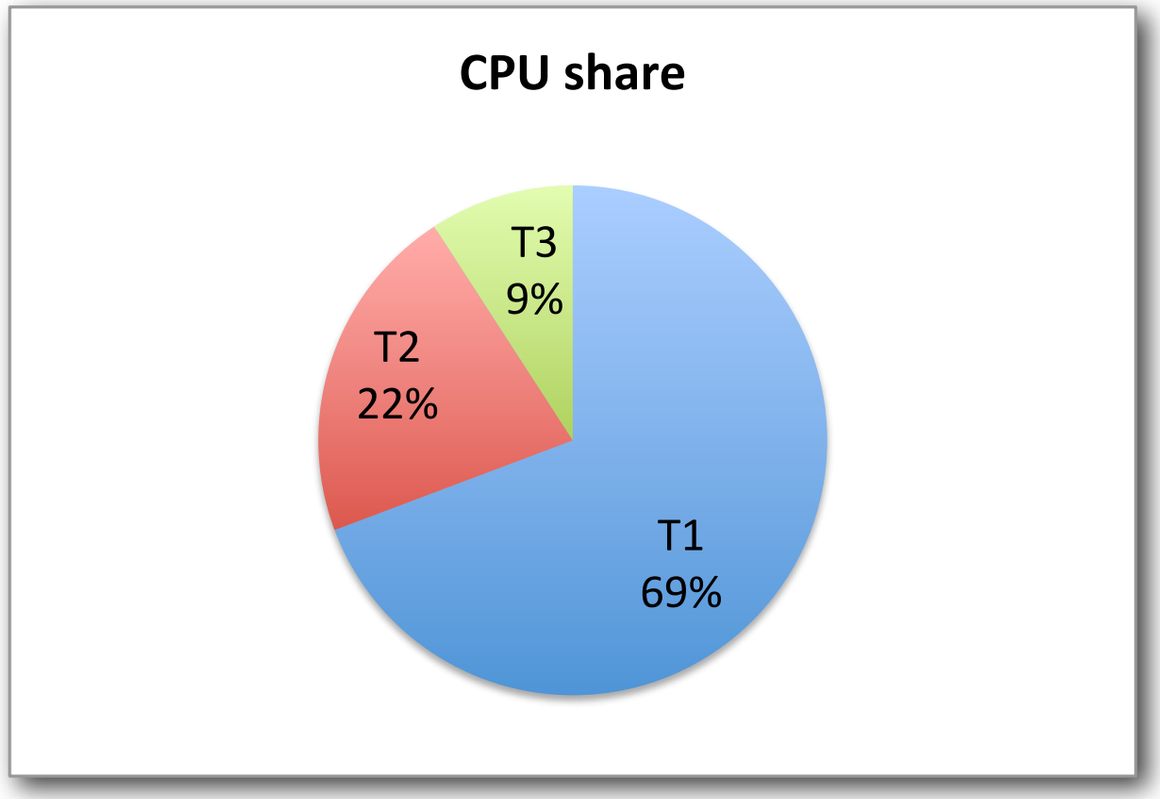

- **T1** (reprocessing, MC processing):
- **T2** (MC production, analysis)
- **T3** (grid and non grid resources for French users) : 50/50

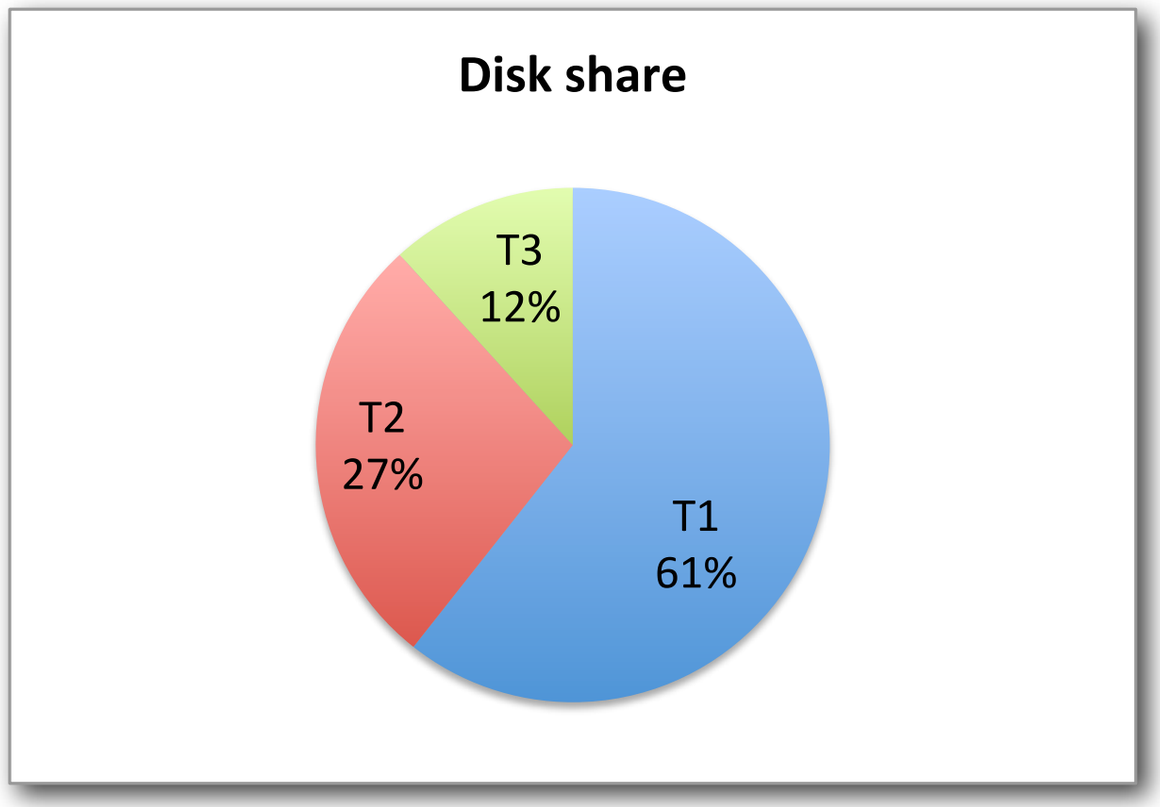

irfu

cea

saclay

T1 2012 disk pledges

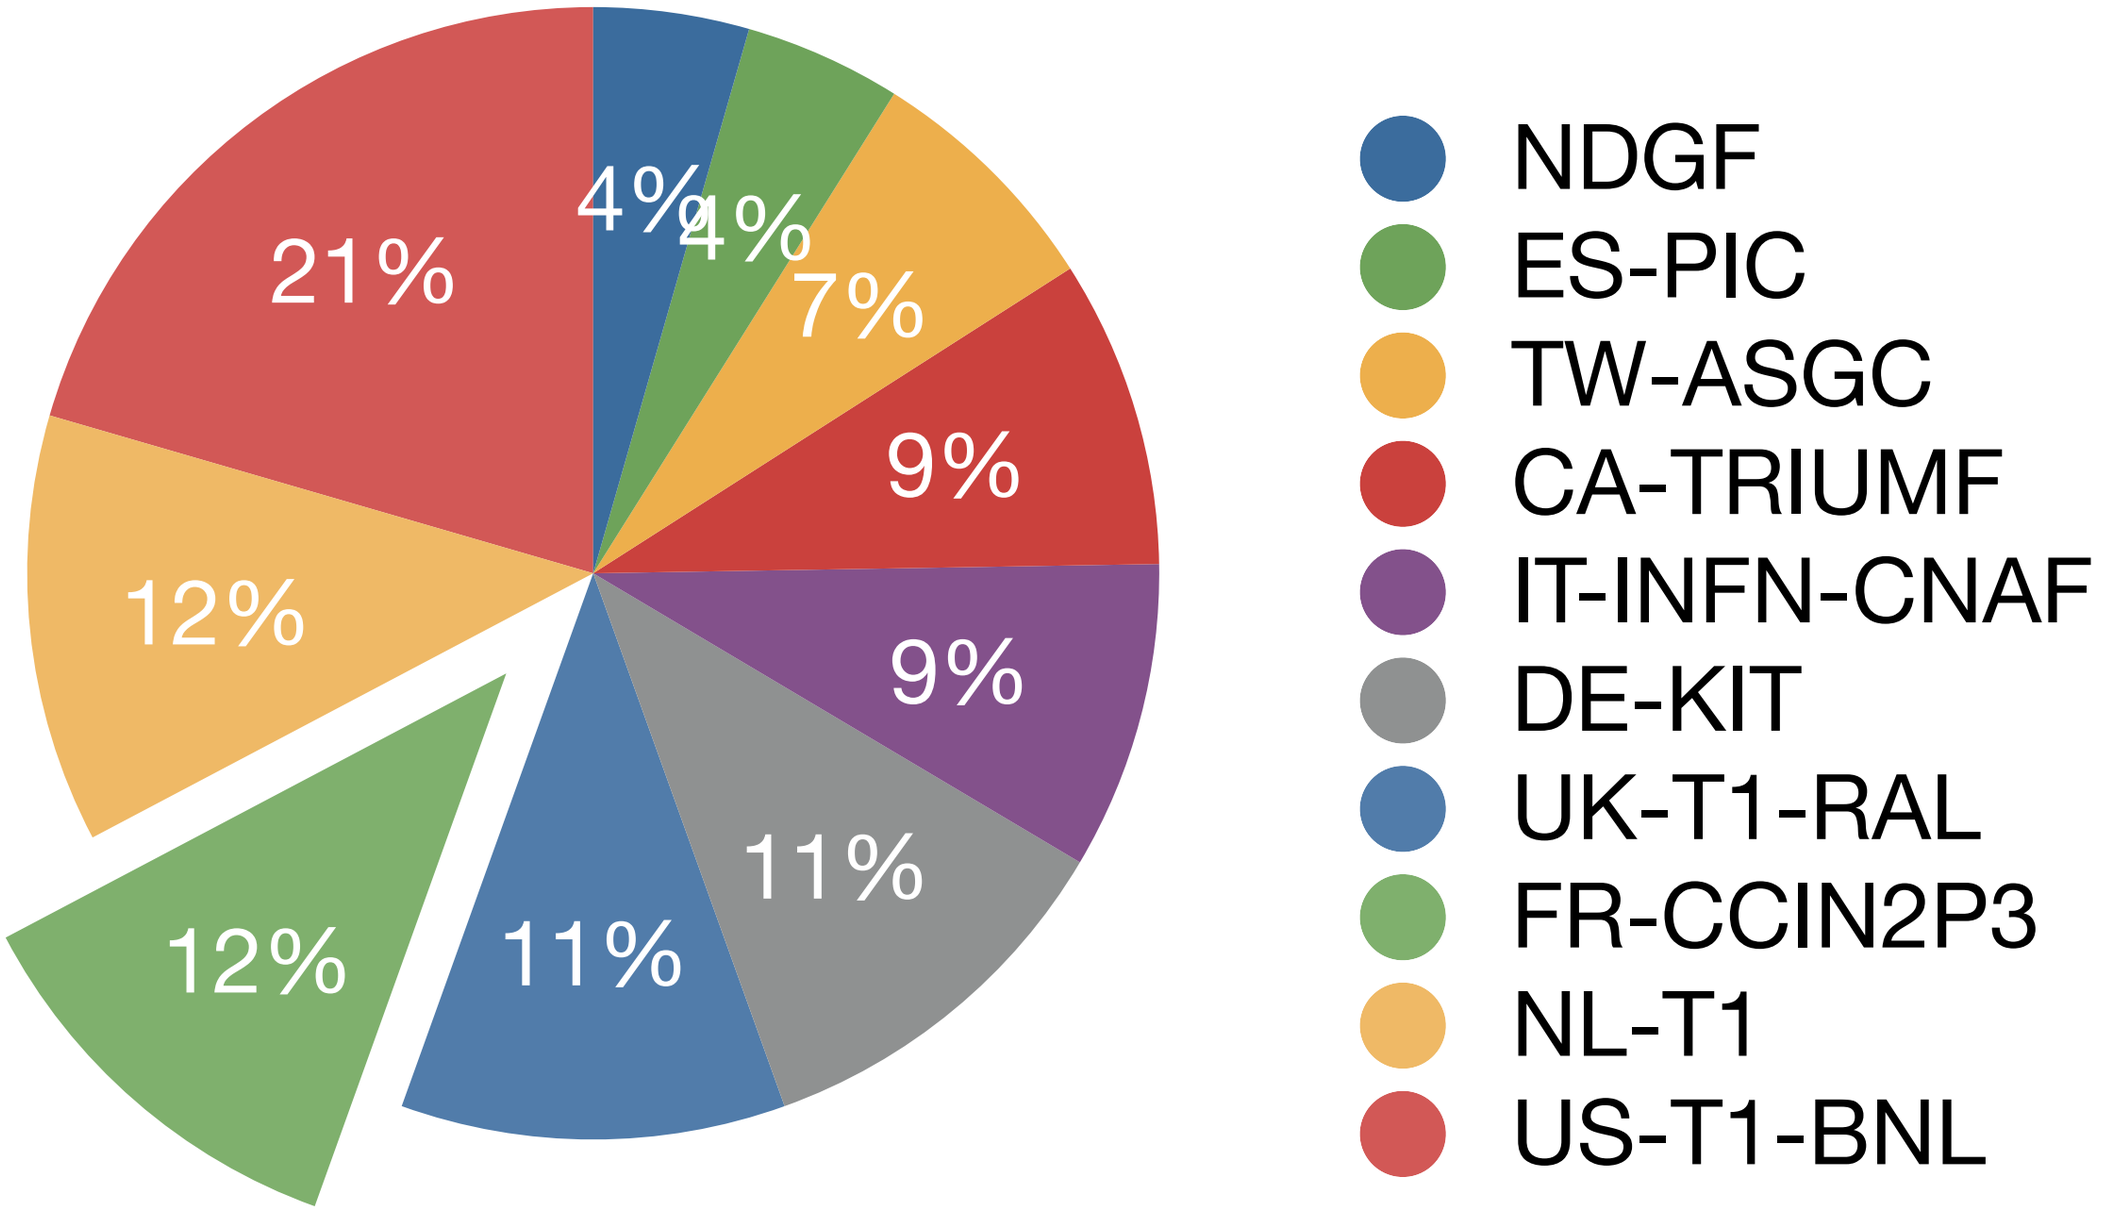

irfu

cea

saclay

2012 T2 disk pledges

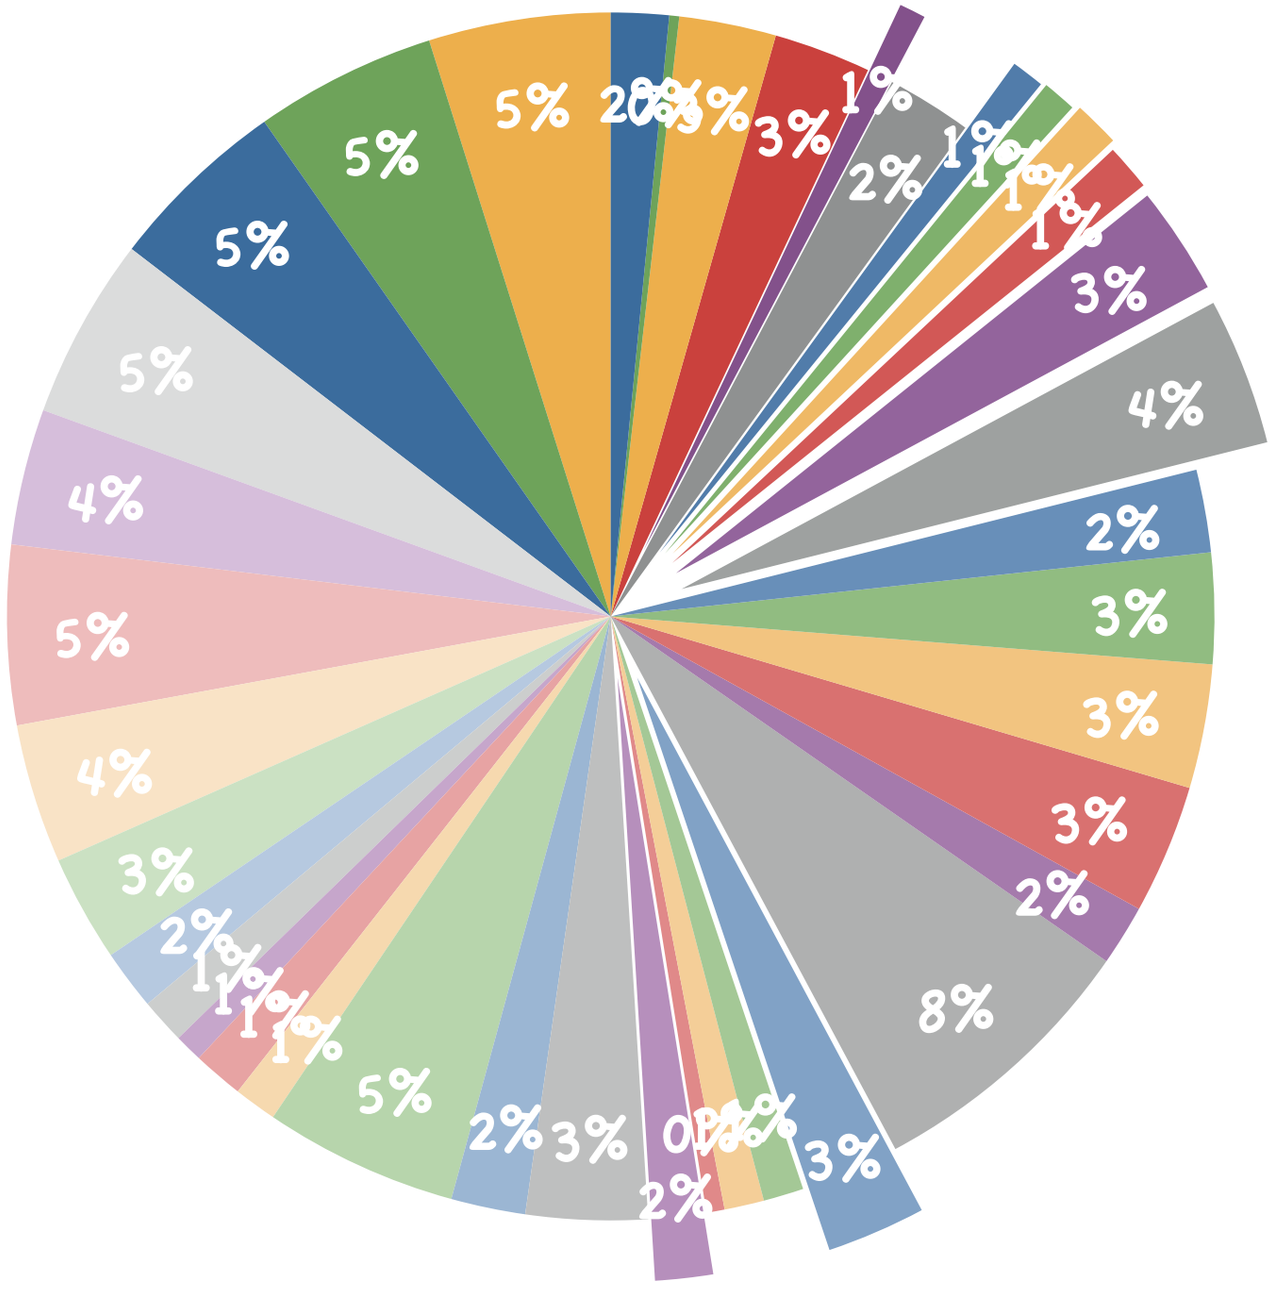

- University of Melbourne Australia
- Austrian Tier-2 Federation Austria
- Canada-East Federation Canada
- Canada-West Federation Canada
- IHEP, Beijing China
- FZU AS, Prague Czech Republic
- CPPM, Marseille France
- LPSC Grenoble France
- LPC, Clermont-Ferrand France
- LAPP, Annecy France
- CC-IN2P3 AF France
- GRIF, Paris France
- ATLAS Federation, HH/Goe Germany
- ATLAS Federation, Munich Germany
- ATLAS Federation DESY Germany
- ATLAS Federation FR/W Germany
- IL-HEP Tier-2 Federation Israel
- INFN T2 Federation Italy
- ICEPP, Tokyo Japan
- UNINETT SIGMA Tier-2 Norway
- Polish Tier-2 Federation Poland
- LIP Tier-2 Federation Portugal
- Romanian Tier-2 Federation Romania
- Russian Data-Intensive GRID Russian Fe
- SiGNET Slovenia
- ATLAS Federation Spain
- SNIC Tier-2 Sweden
- CHIPP Switzerland
- Taiwan Analysis Facility Federation Taipe
- Turkish Tier-2 Federation Turkey
- SouthGrid UK
- ScotGrid UK
- London Tier 2 UK
- NorthGrid UK
- Northeast ATLAS T2 USA
- Great Lakes ATLAS T2 USA
- Midwest ATLAS T2 USA
- SLAC ATLAS T2 USA
- Southwest ATLAS T2 USA

16% of ATLAS

irfu

cea

saclay

T2 2012 disk pledge by cloud

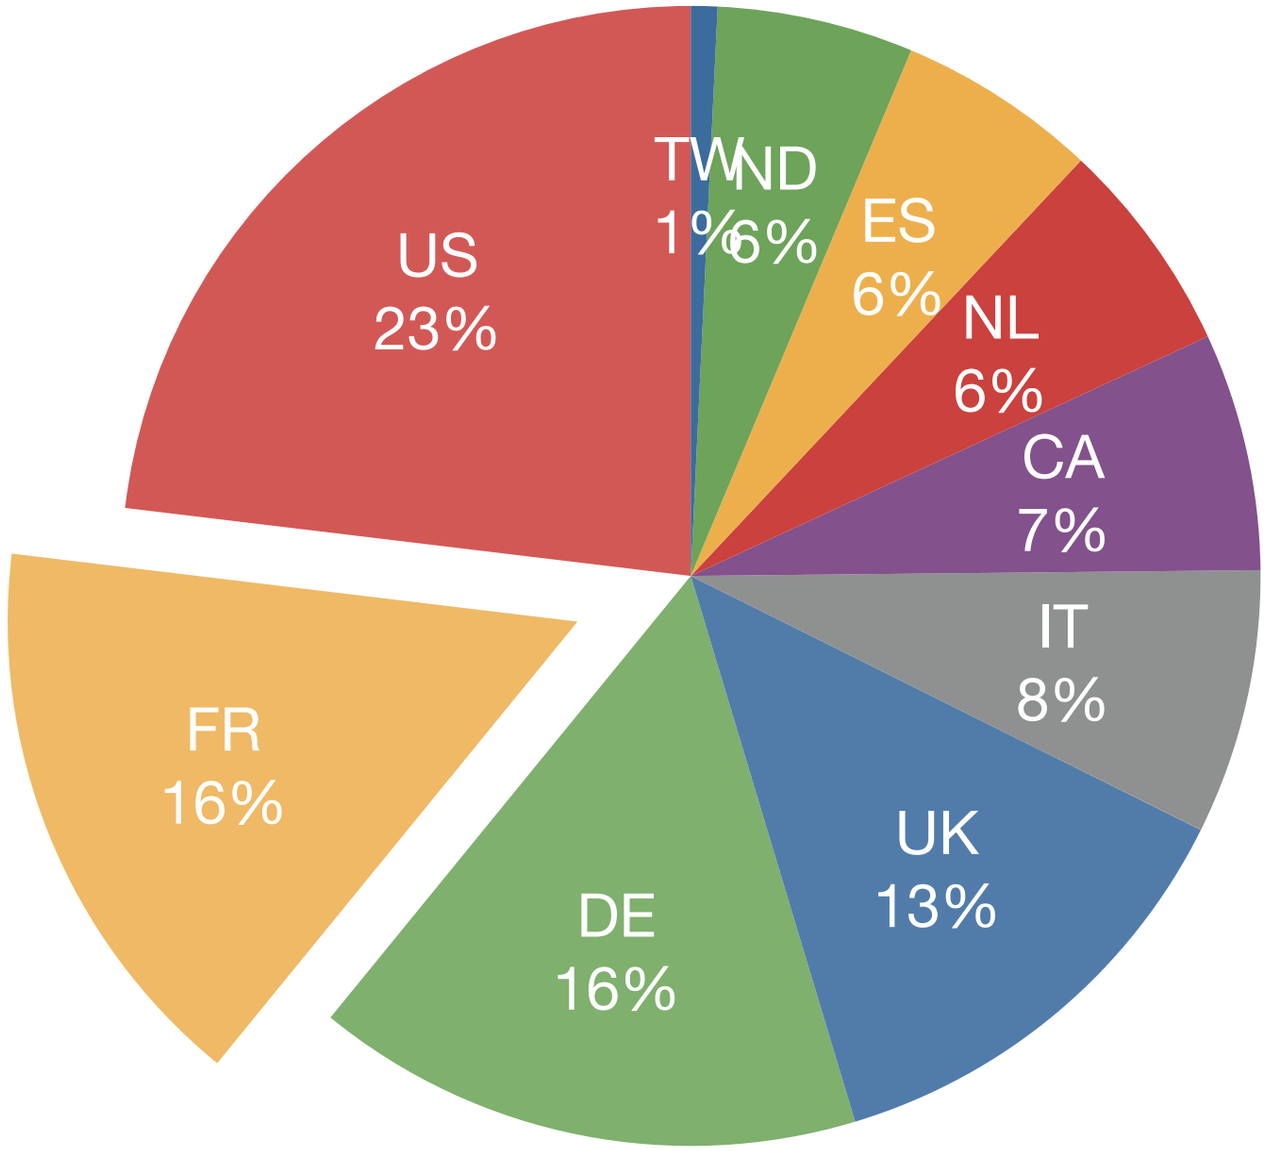

FR-cloud share :
12 - 16 % of ATLAS activities

irfu
cea
saclay

Performances over last year
[May, 2011 - May,2012]

Analysis jobs [May, 2011 - May, 2012]

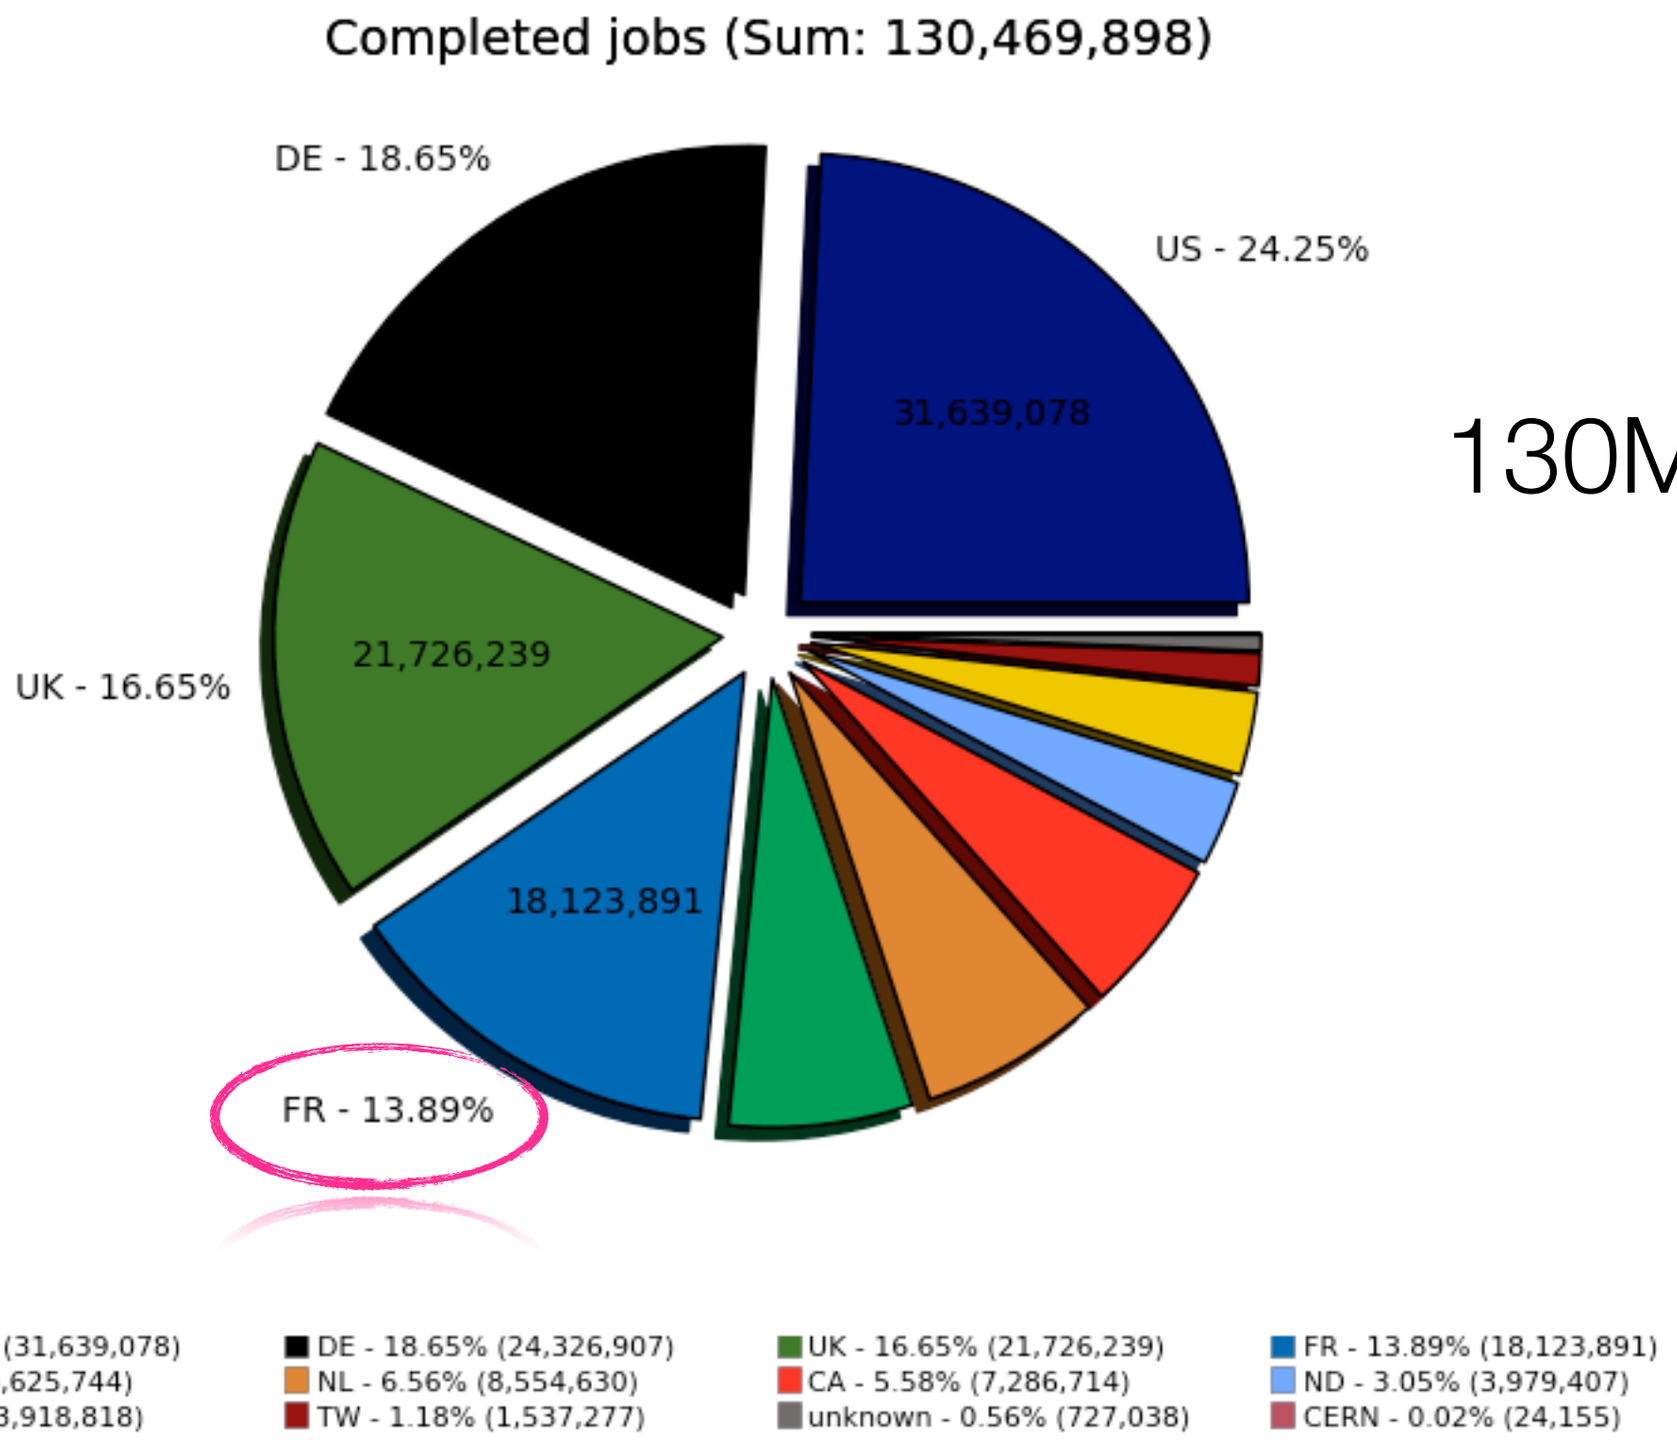

130M jobs

Production jobs [May, 2011 - May, 2012]

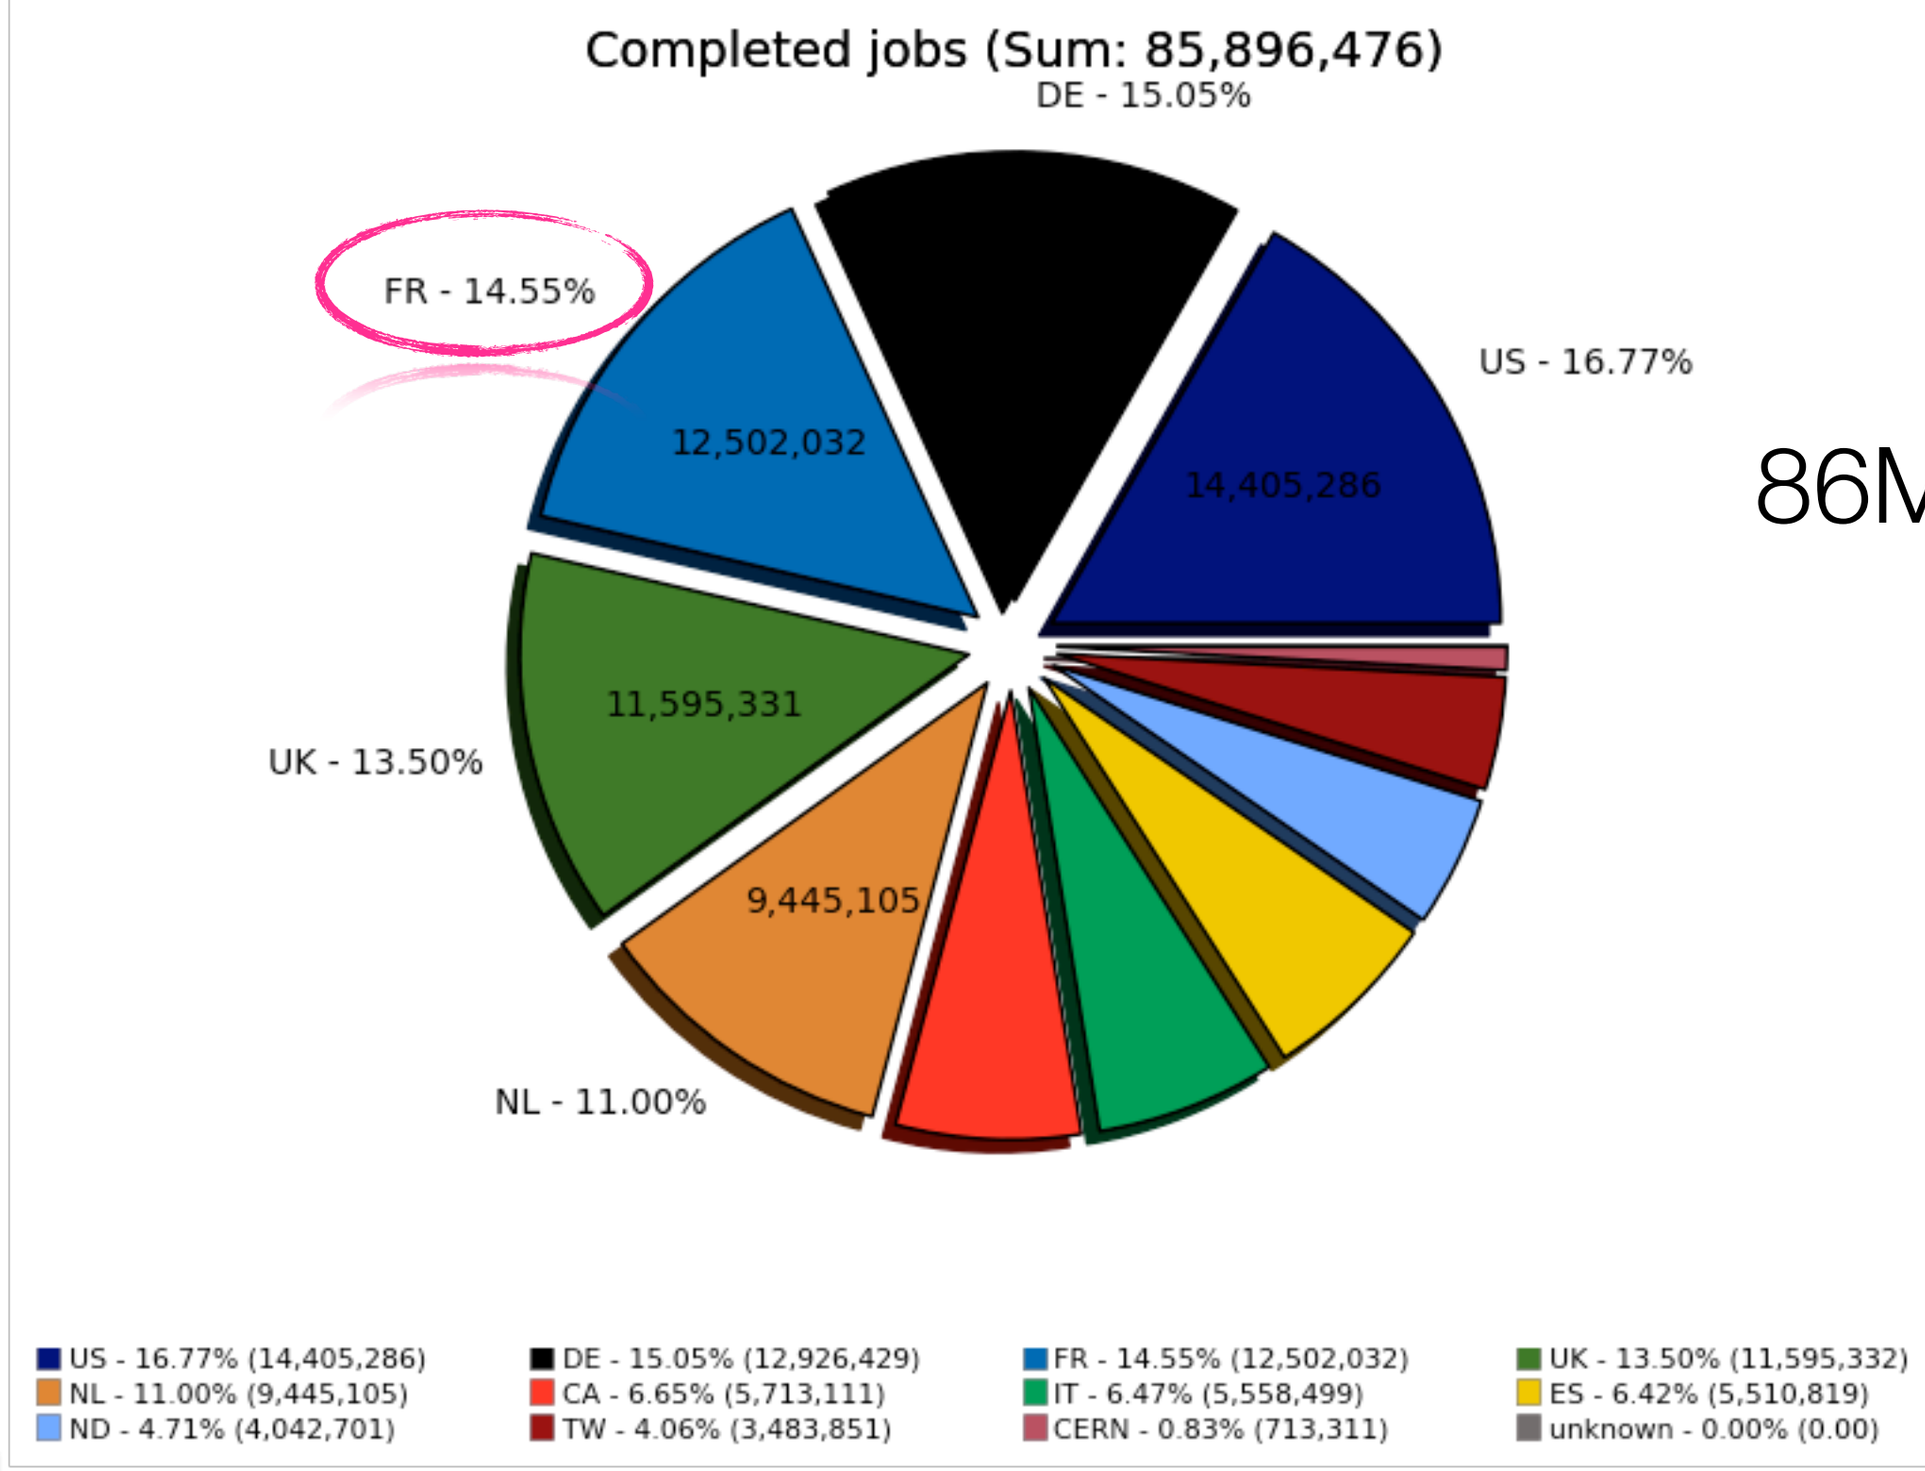

86M jobs

Production jobs [May, 2011 - May, 2012]

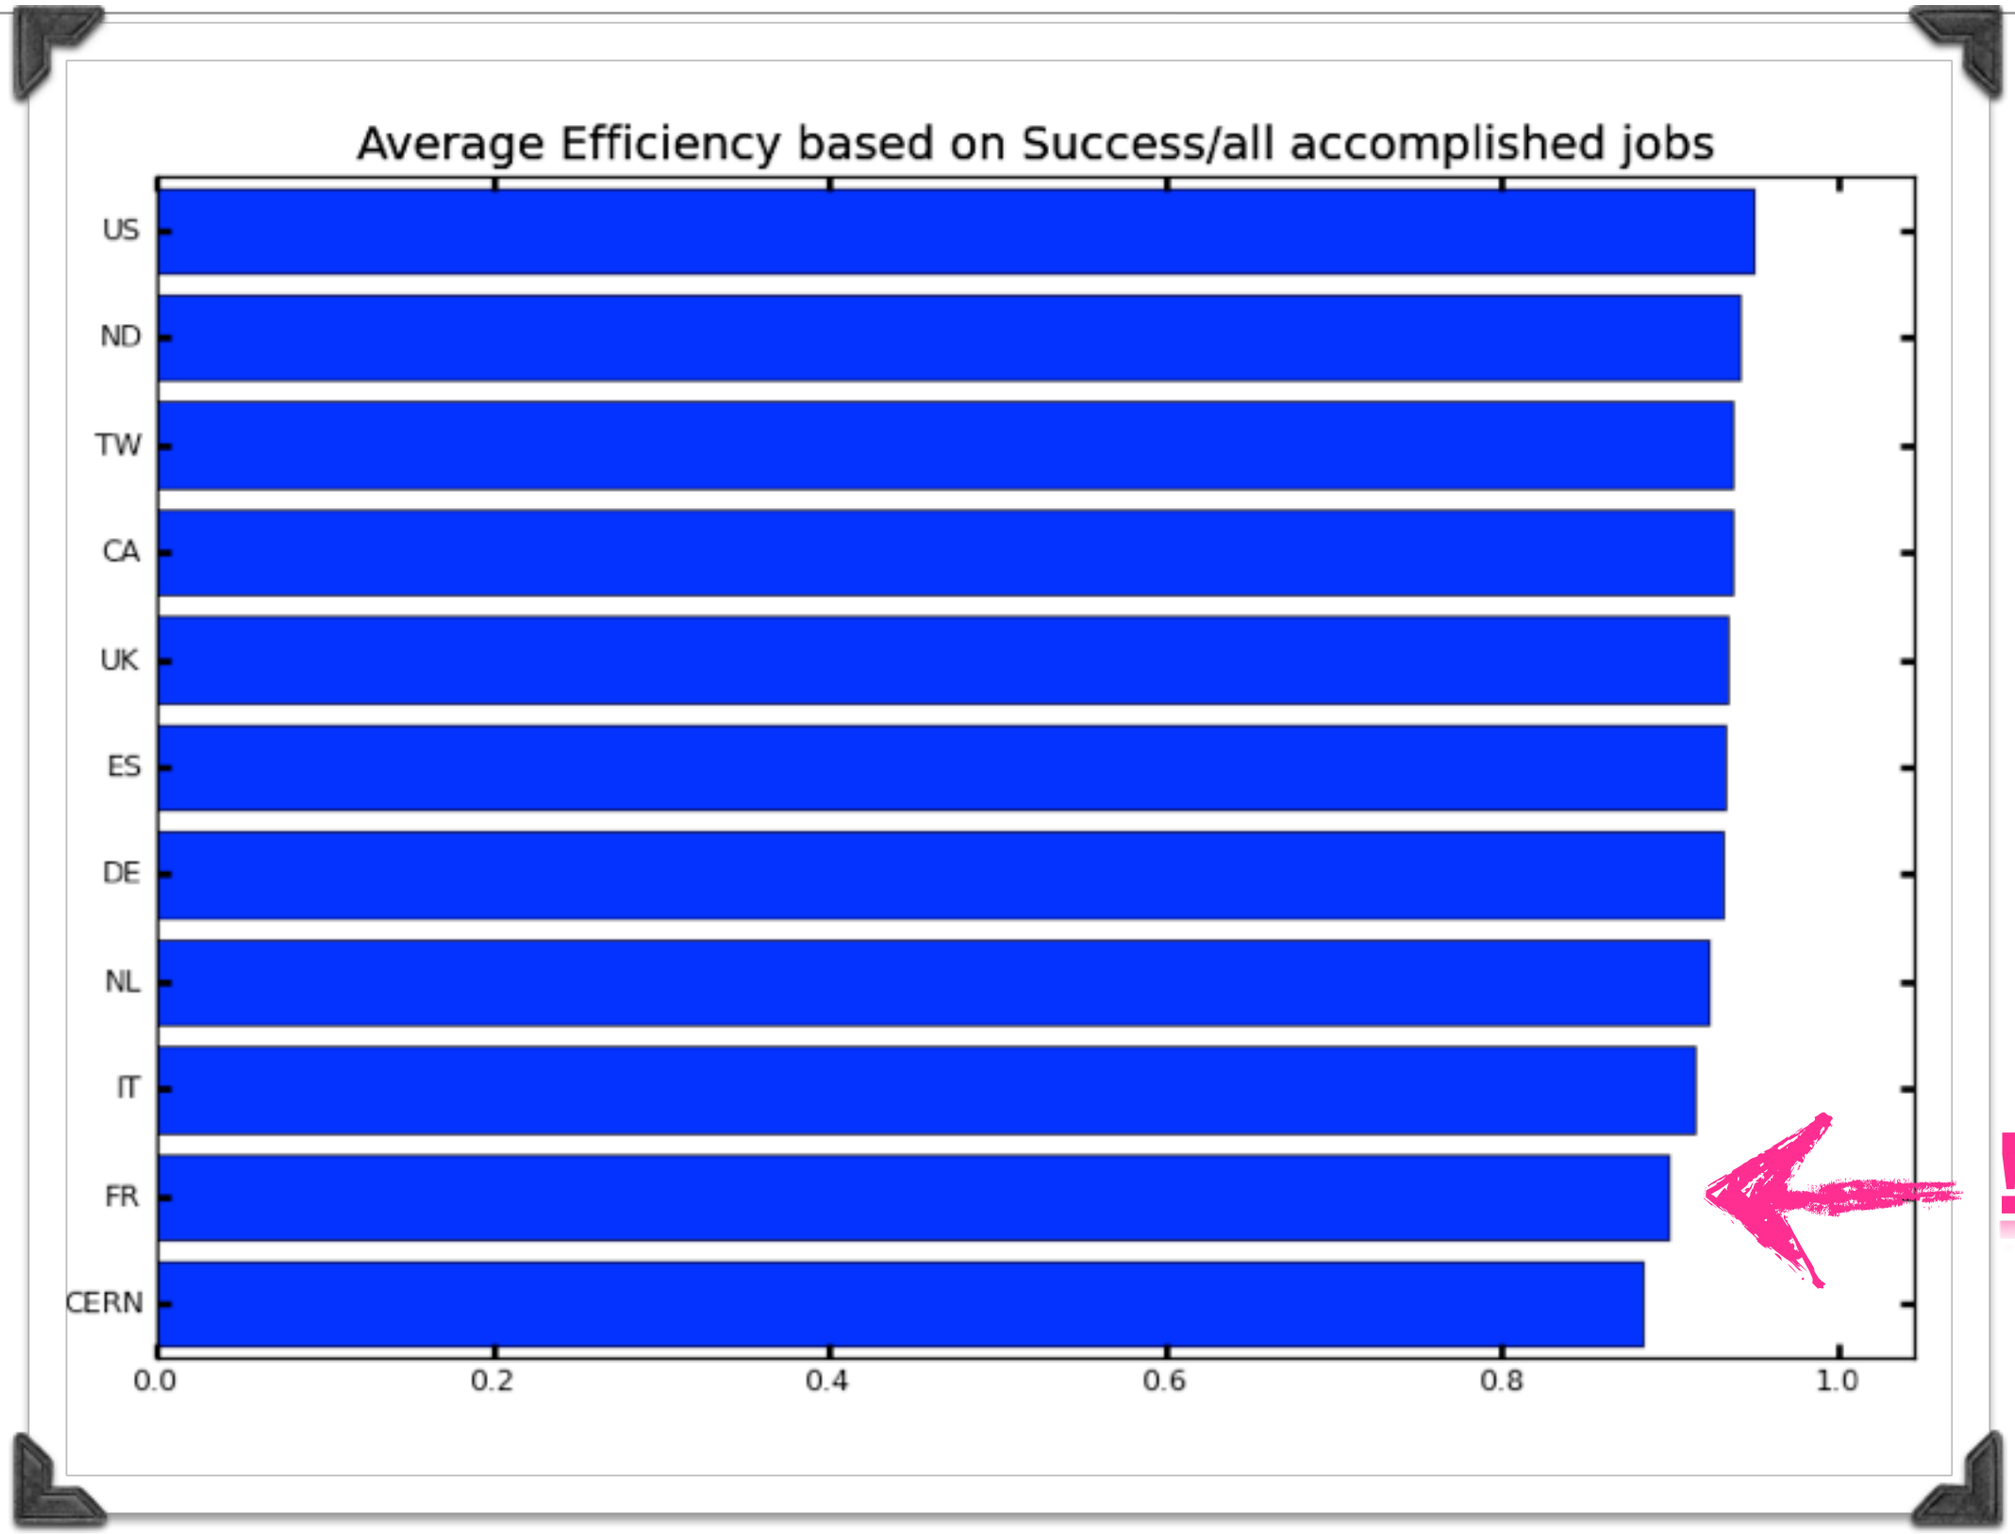

Production jobs [May, 2011 - May, 2012]

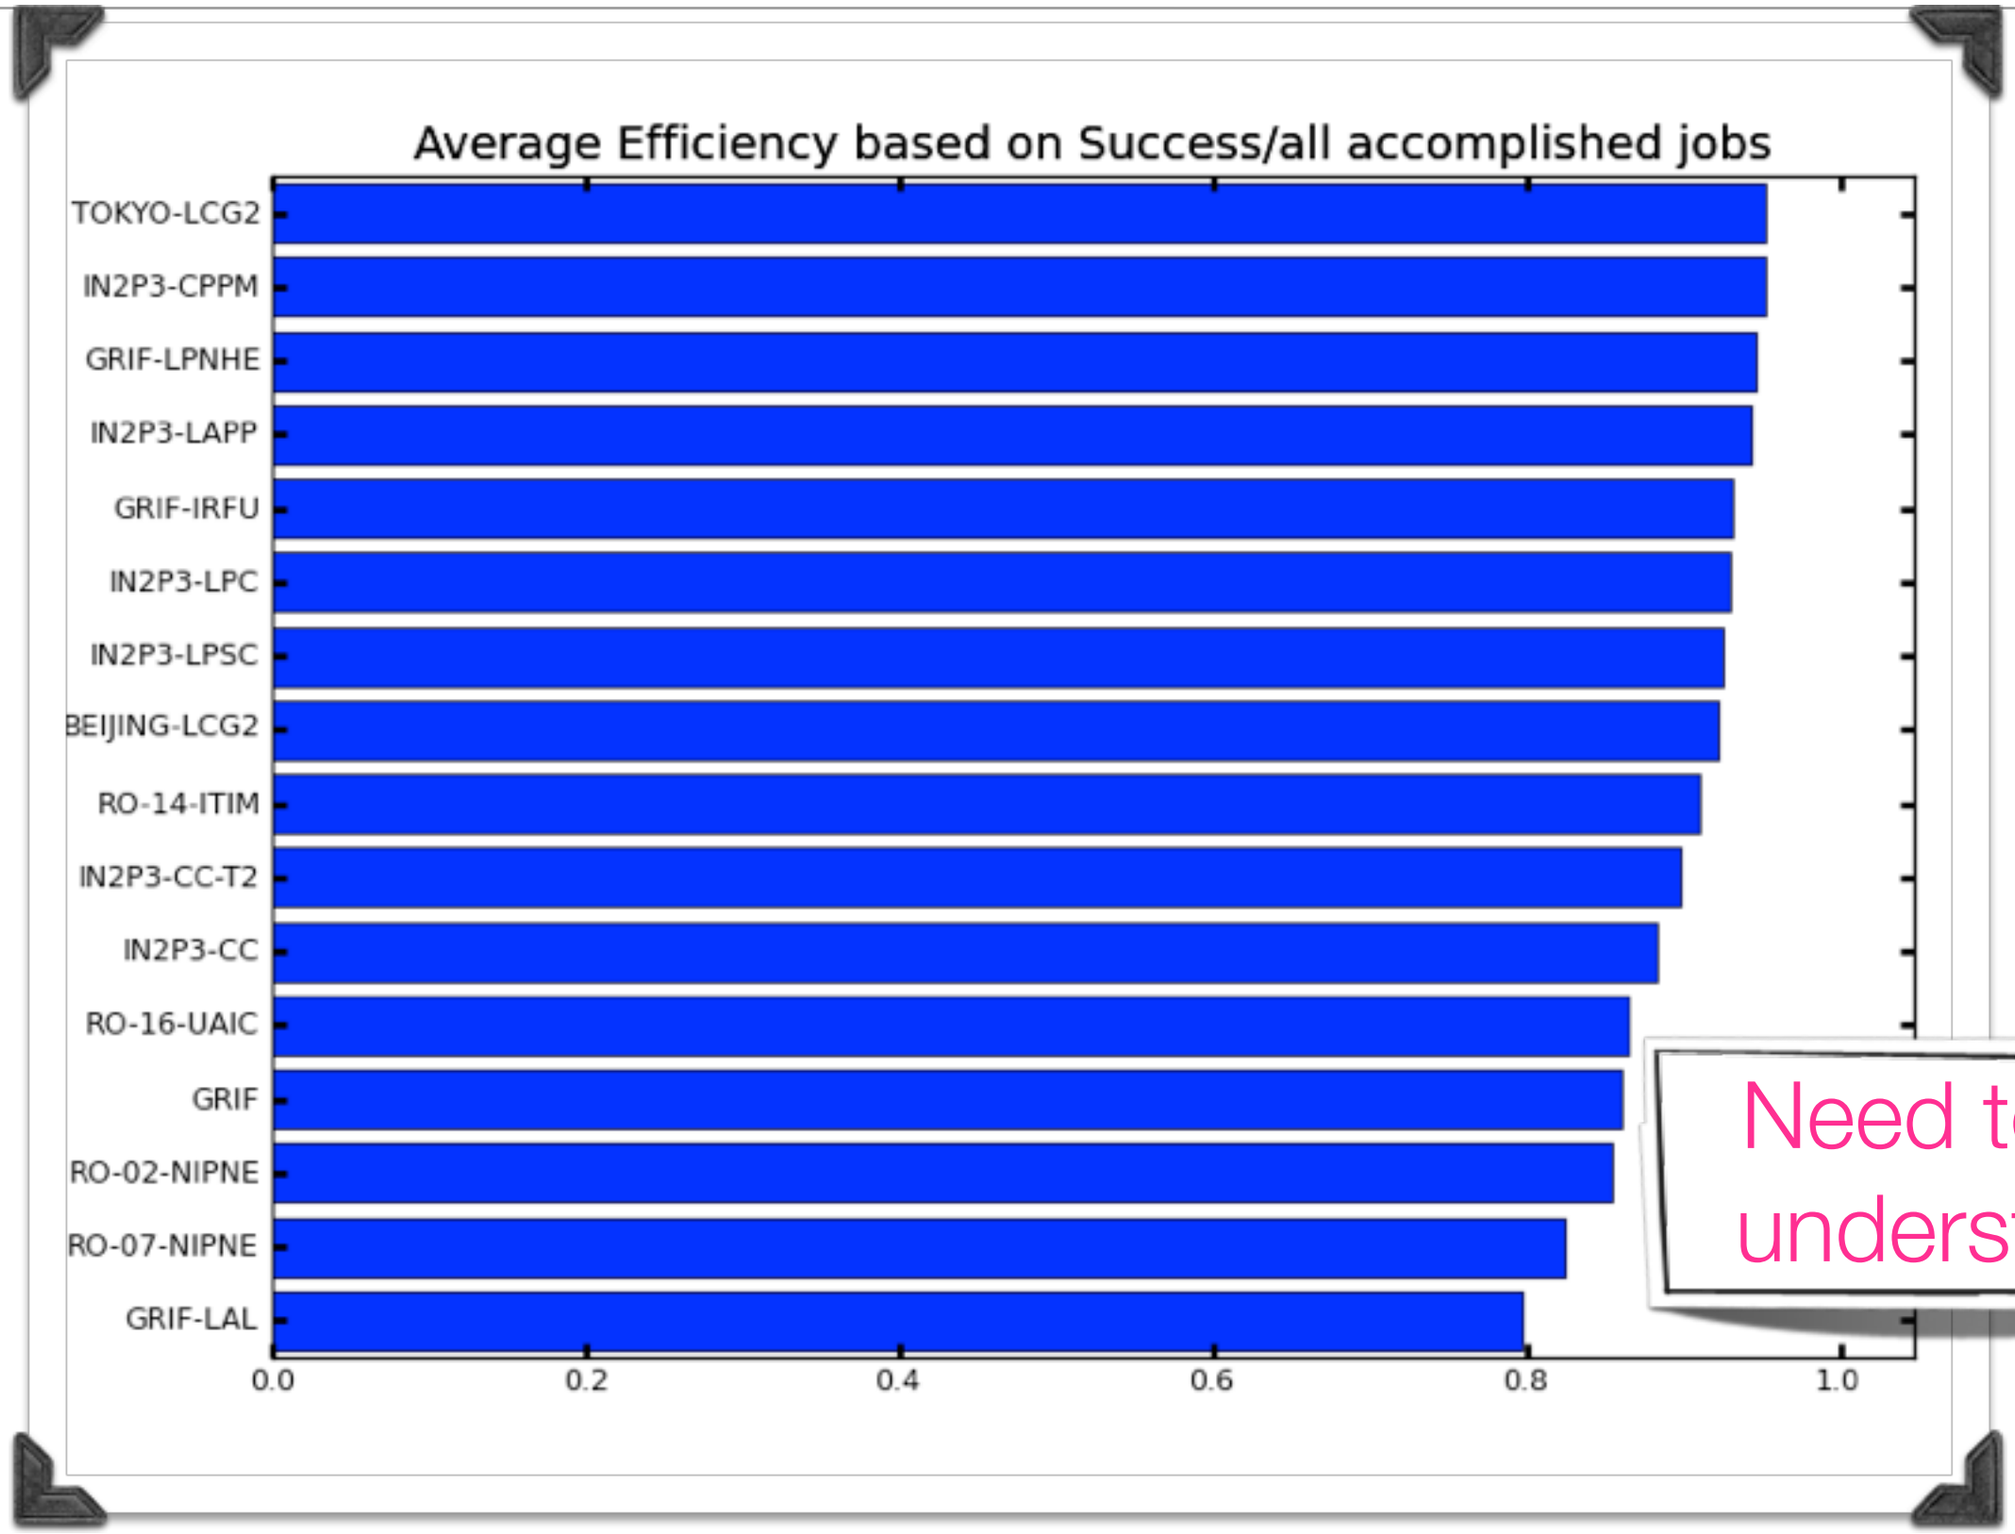

Production efficiency vs time by site

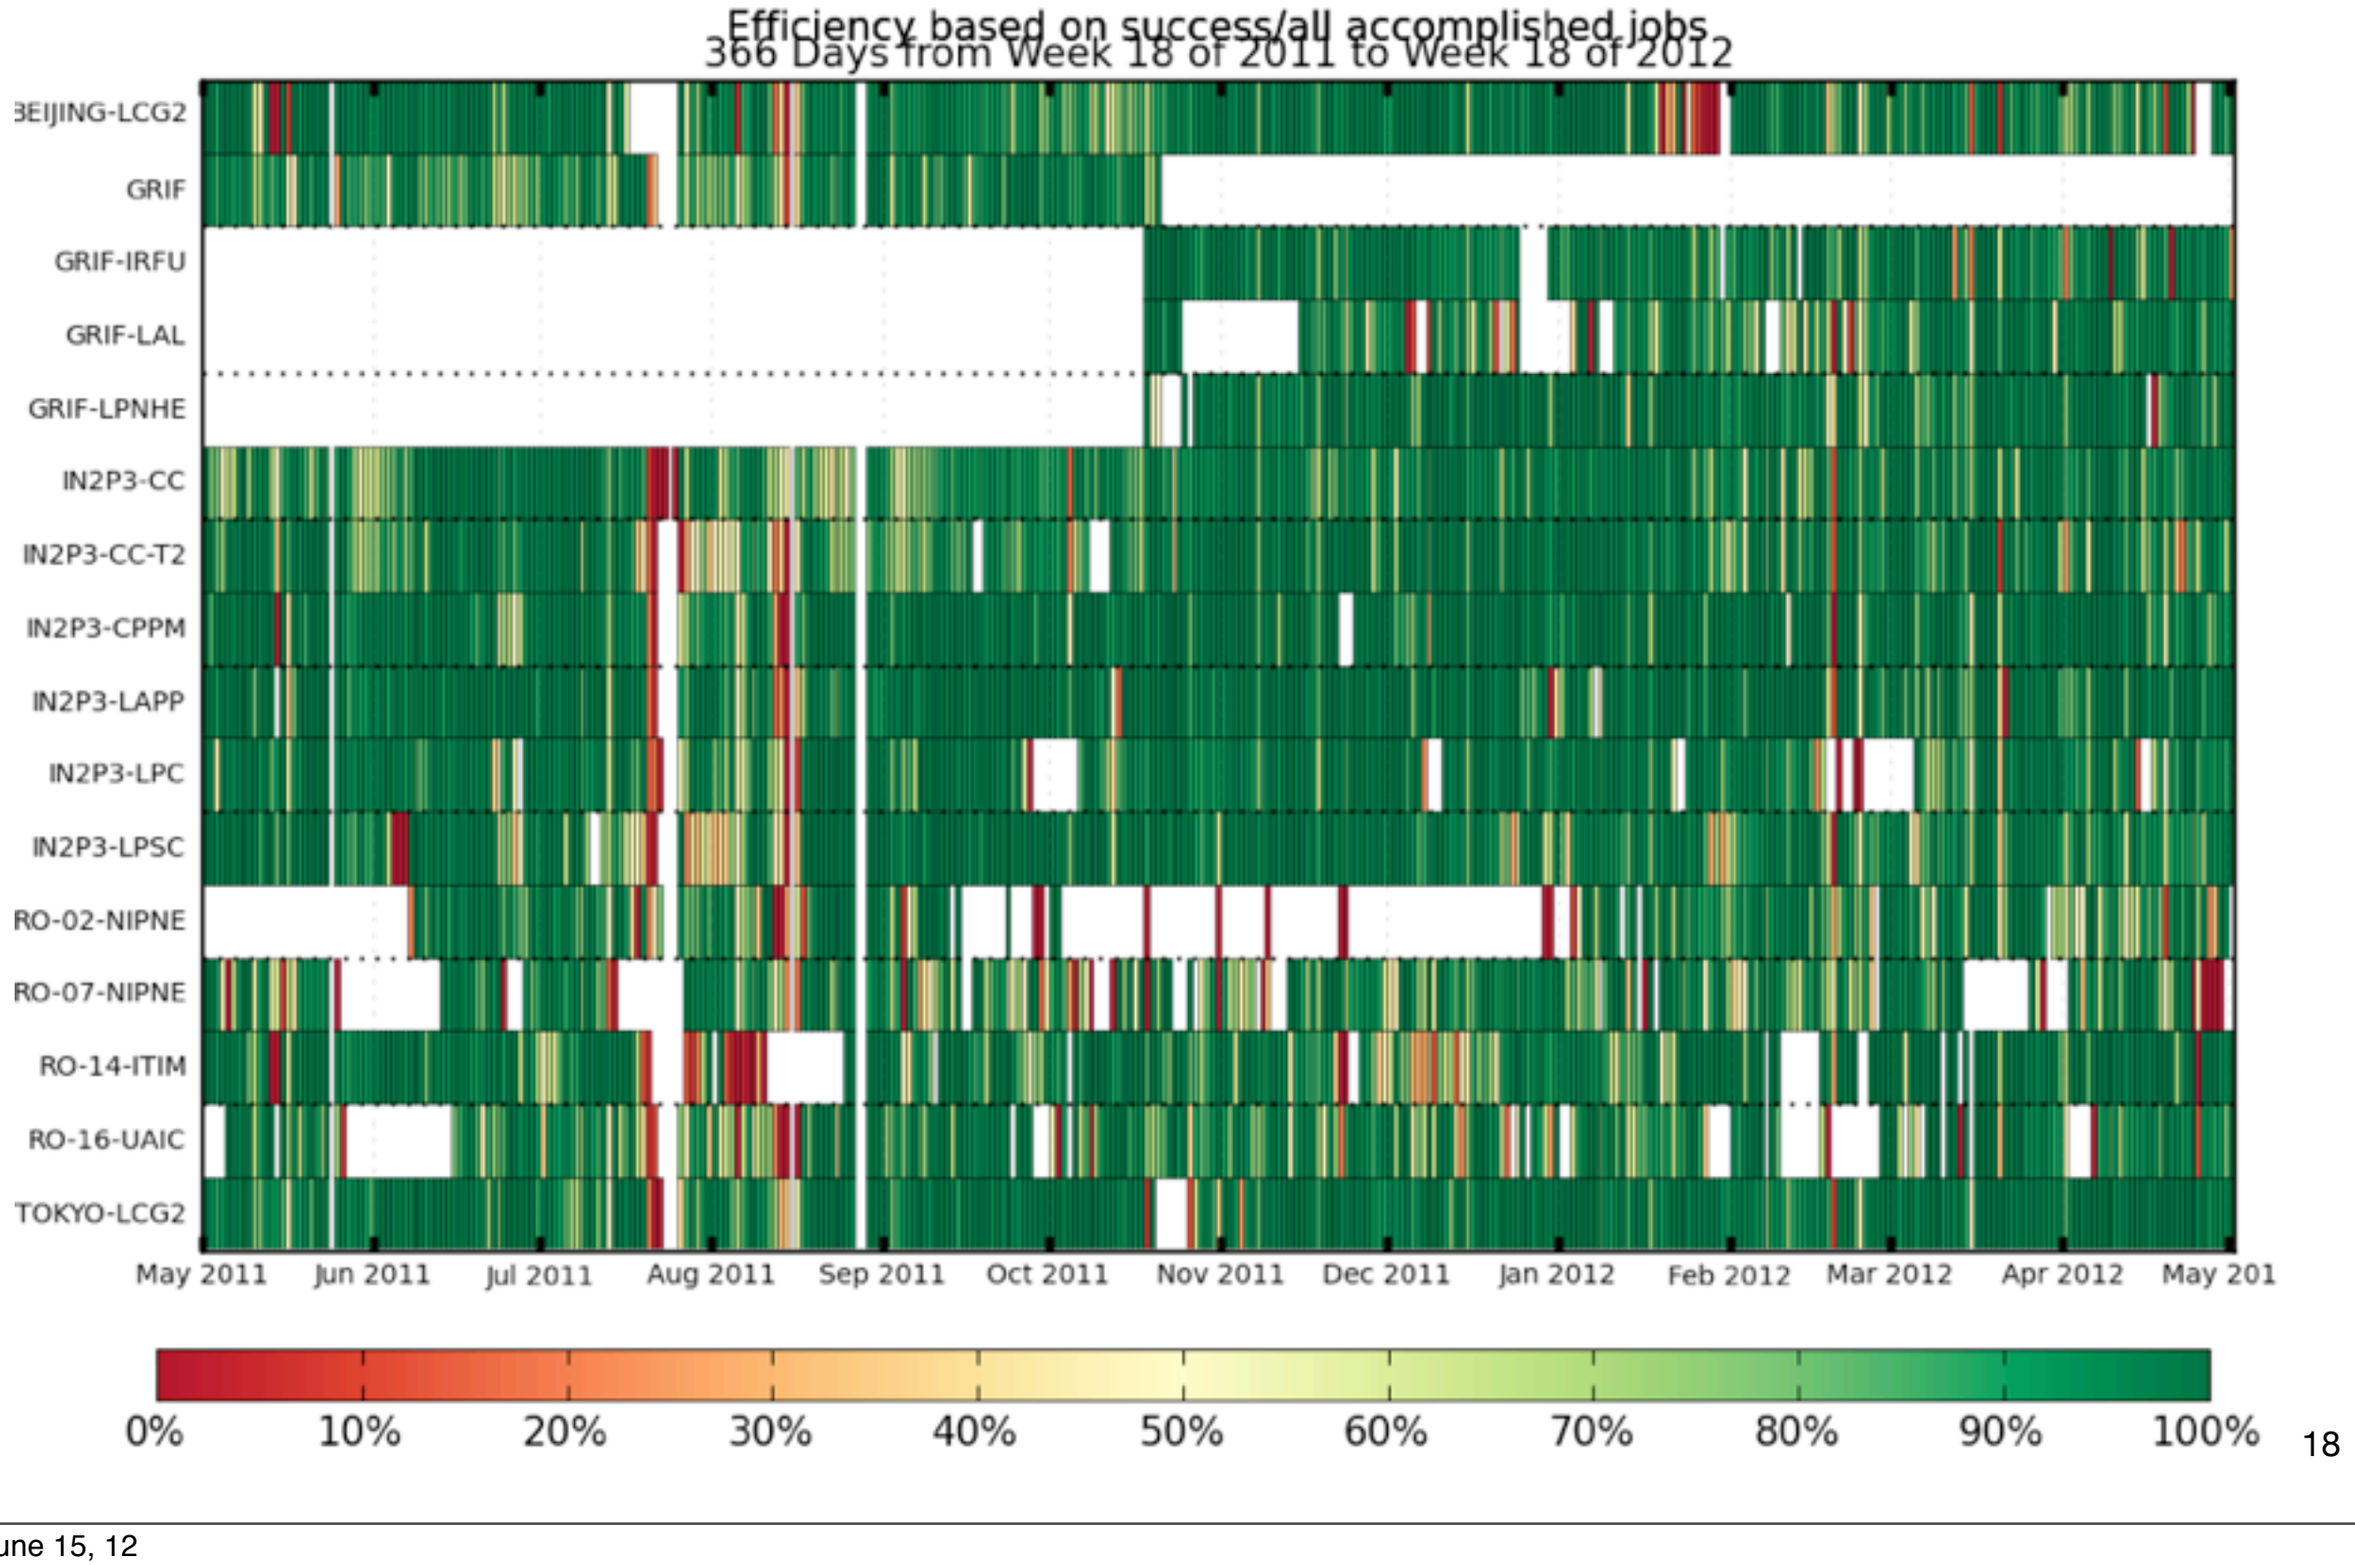

irfu

cea

saclay

errors by site (T1 excluded)

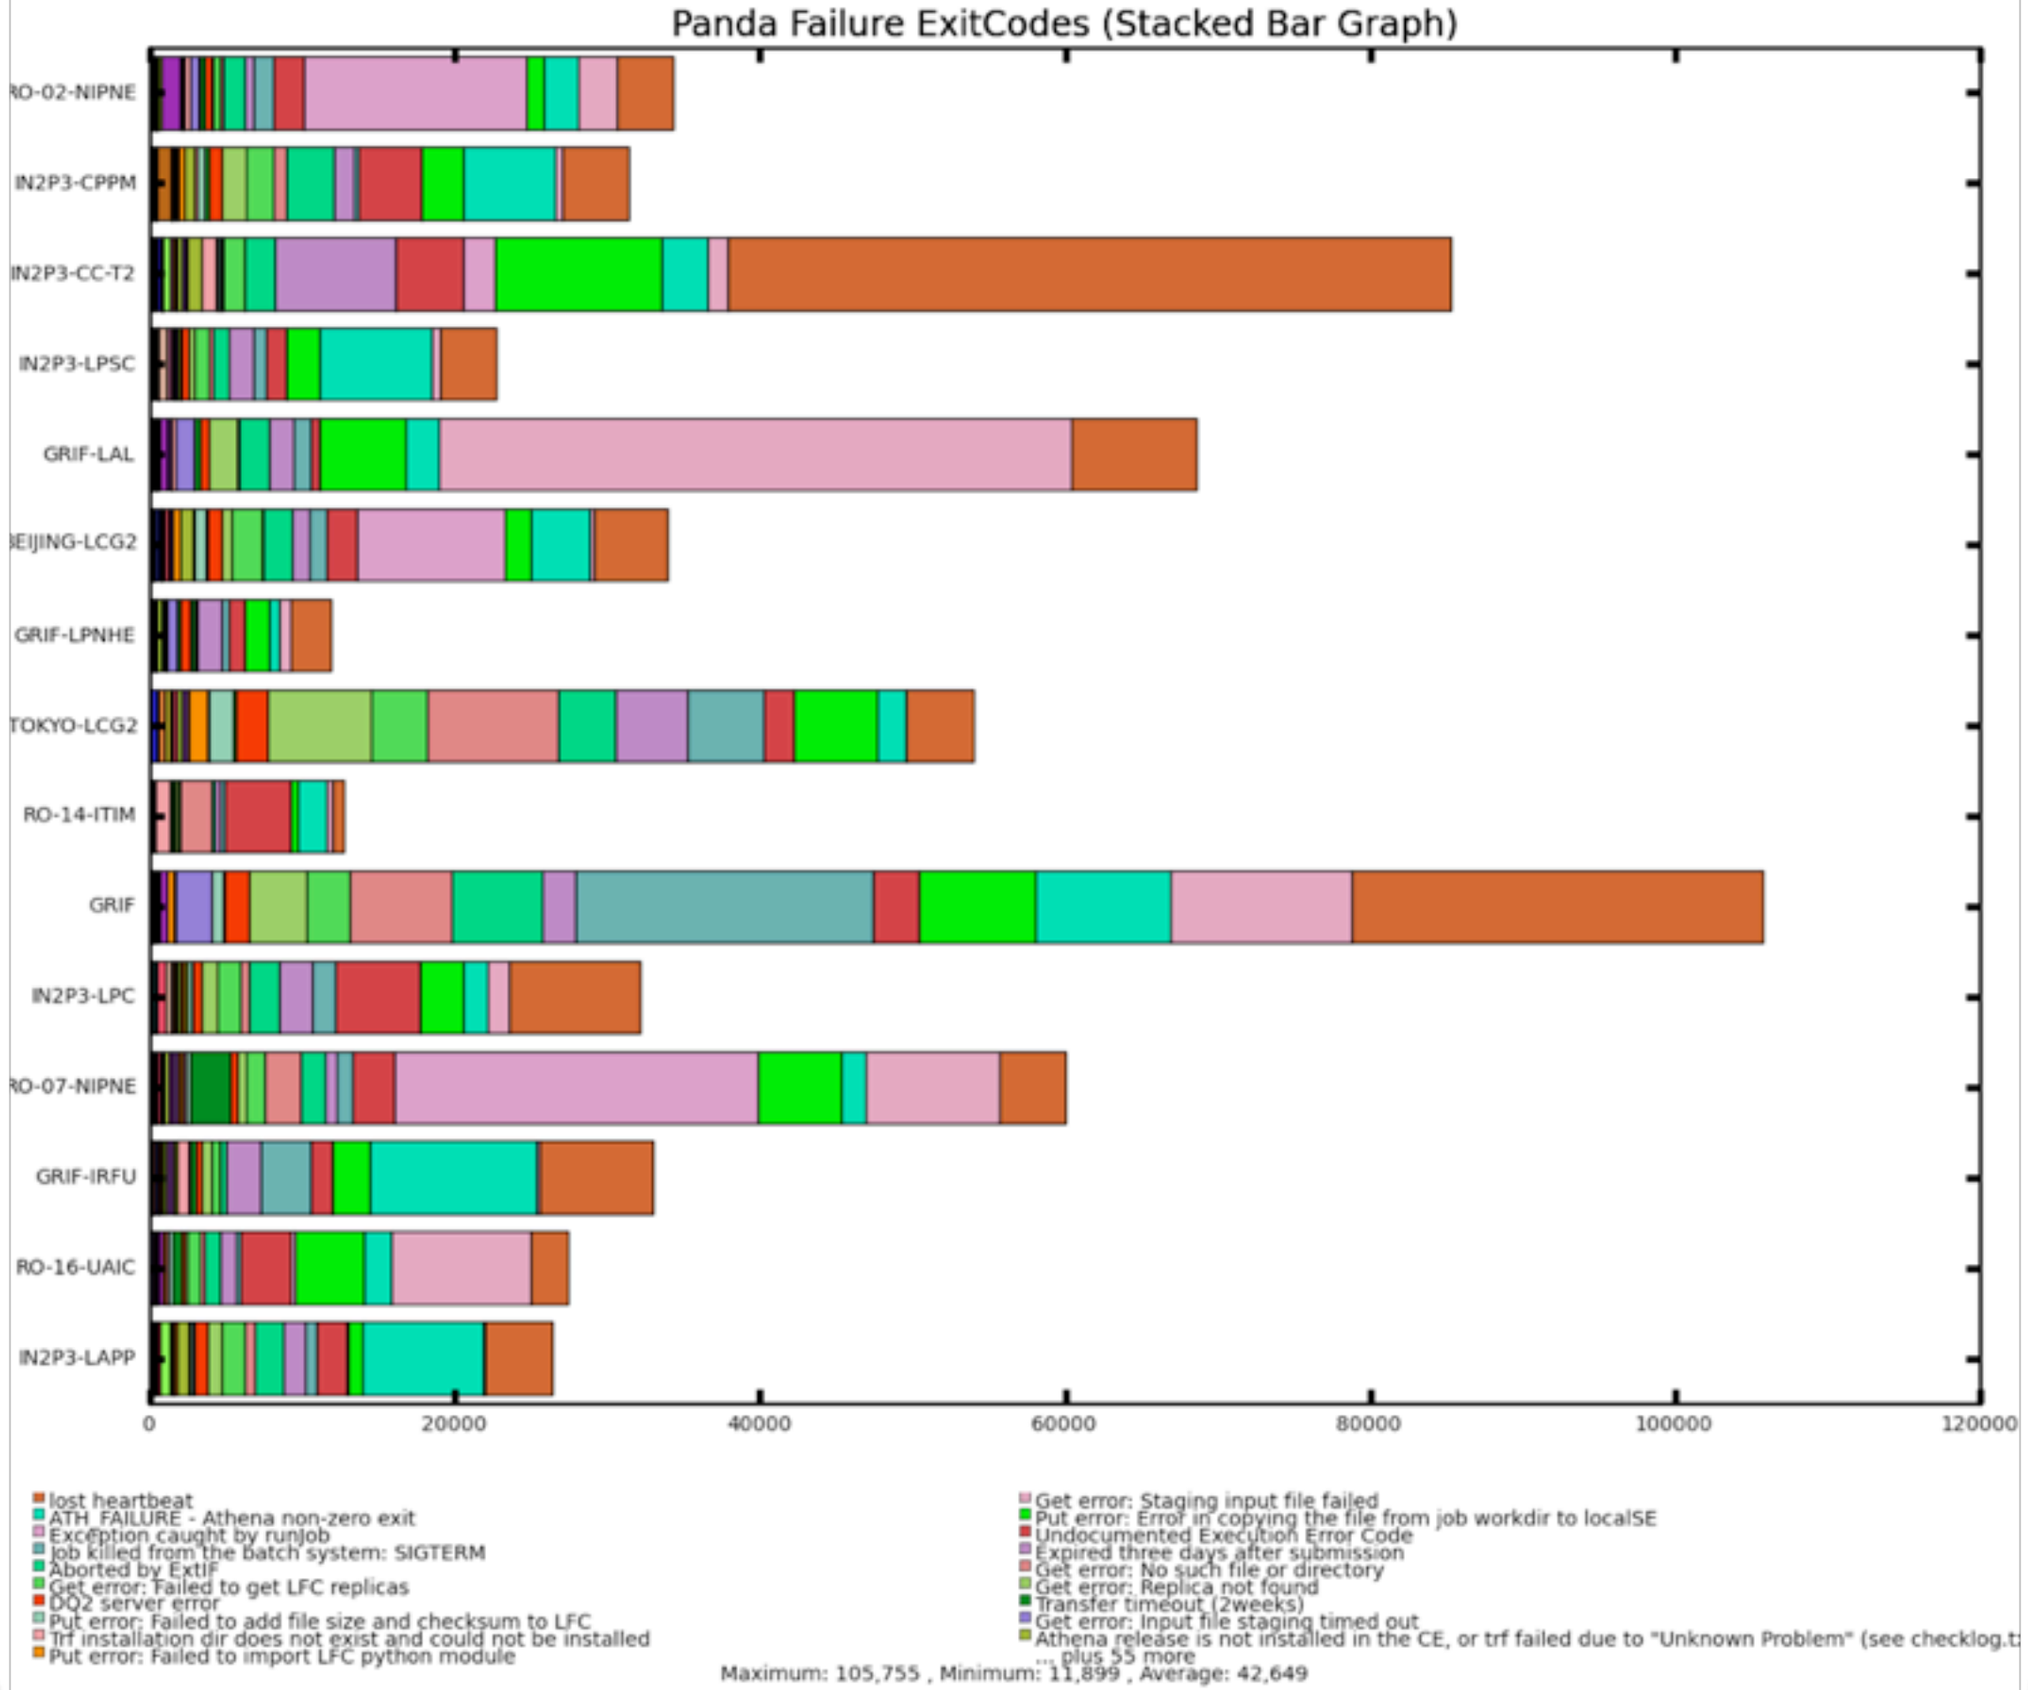

Transfer to FR cloud [May 2011 - May 2012] : **8.1 PB**

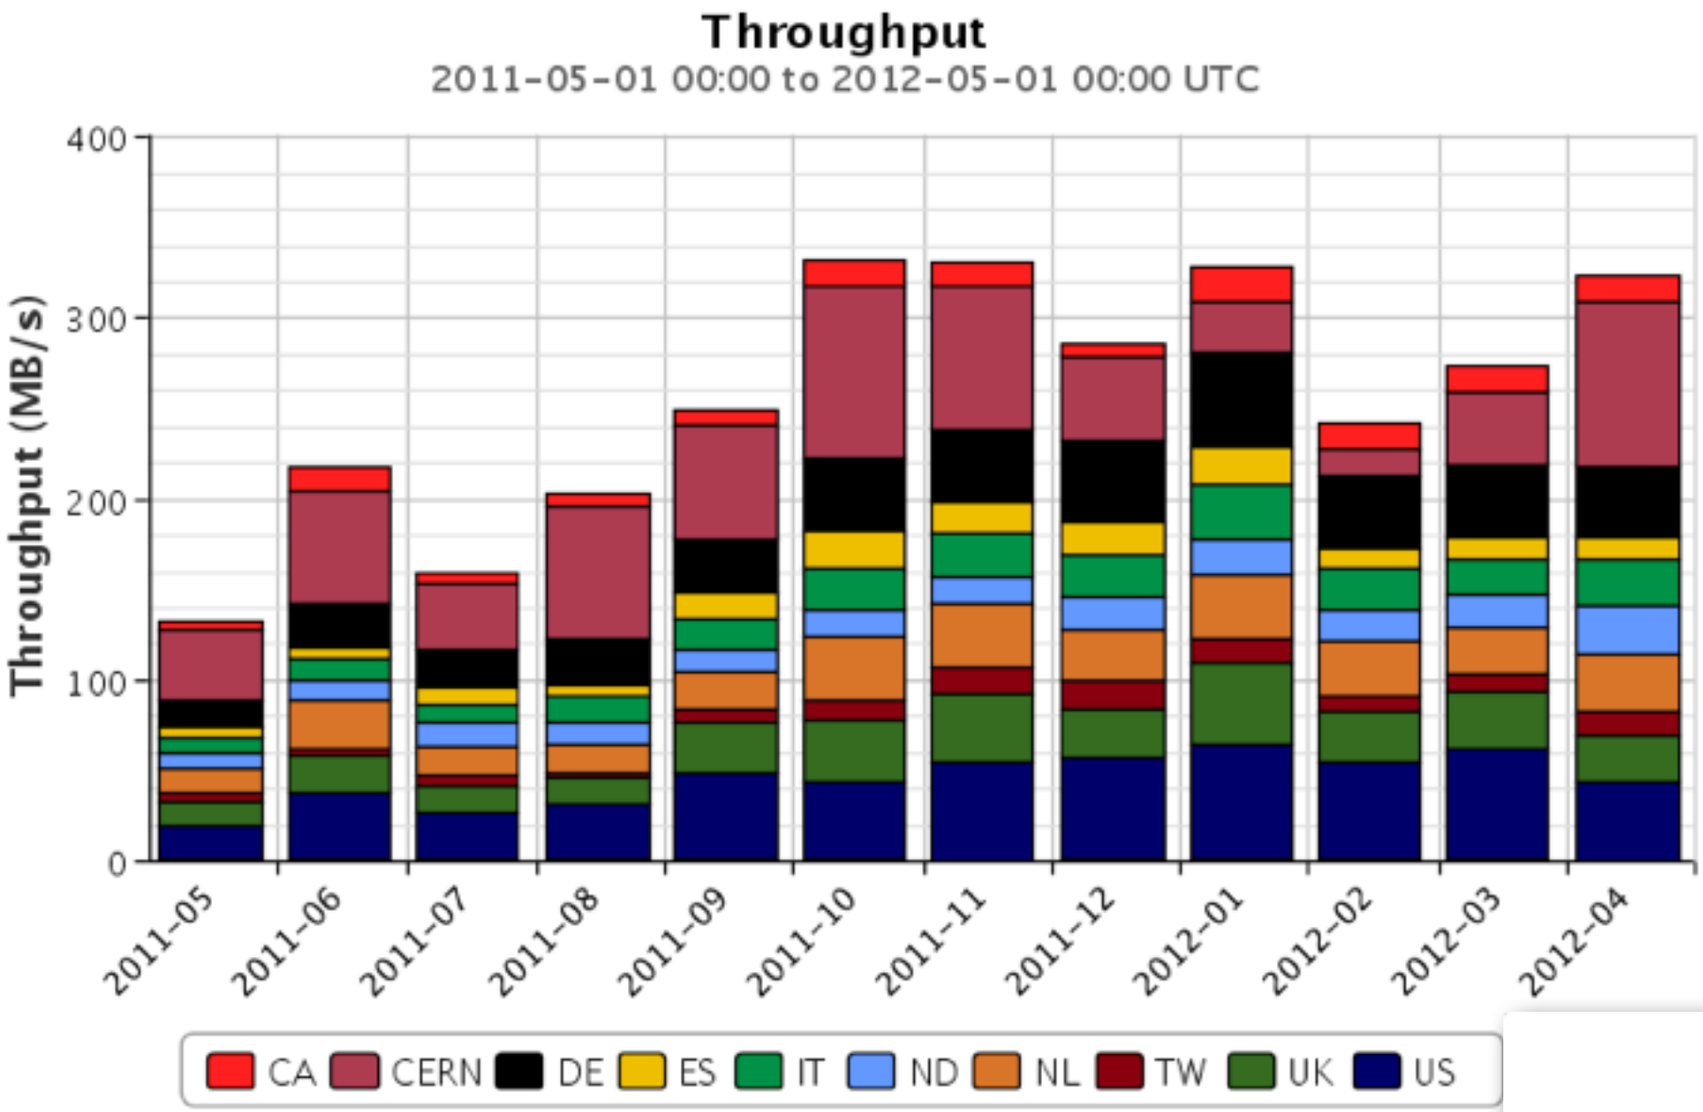

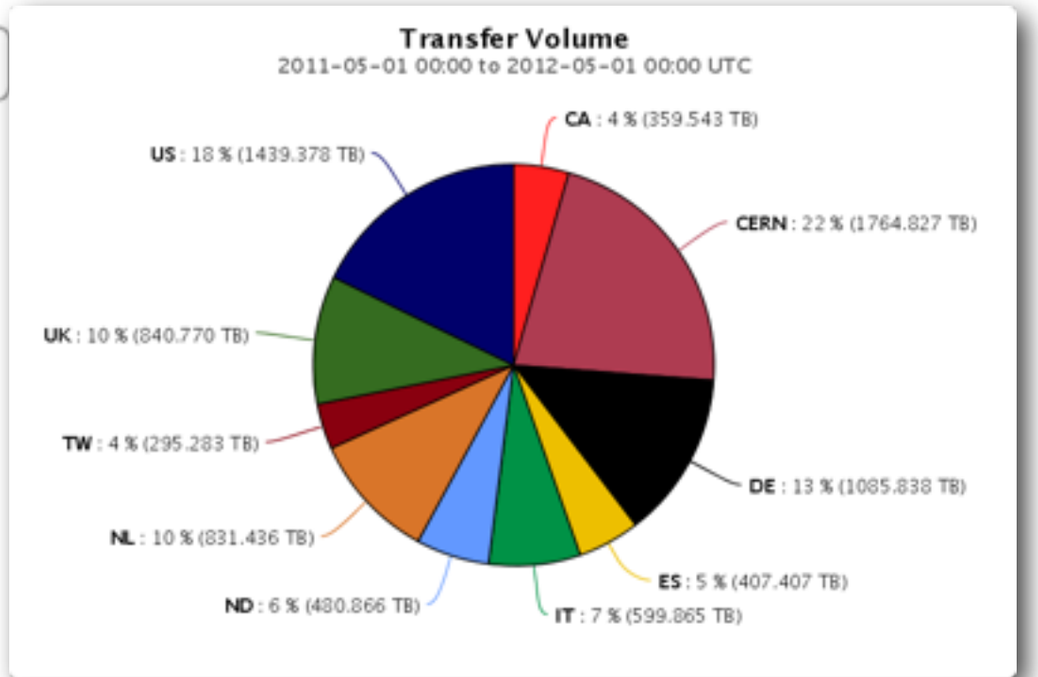

irfu

cea

saclay

import destination : **70% to T1**

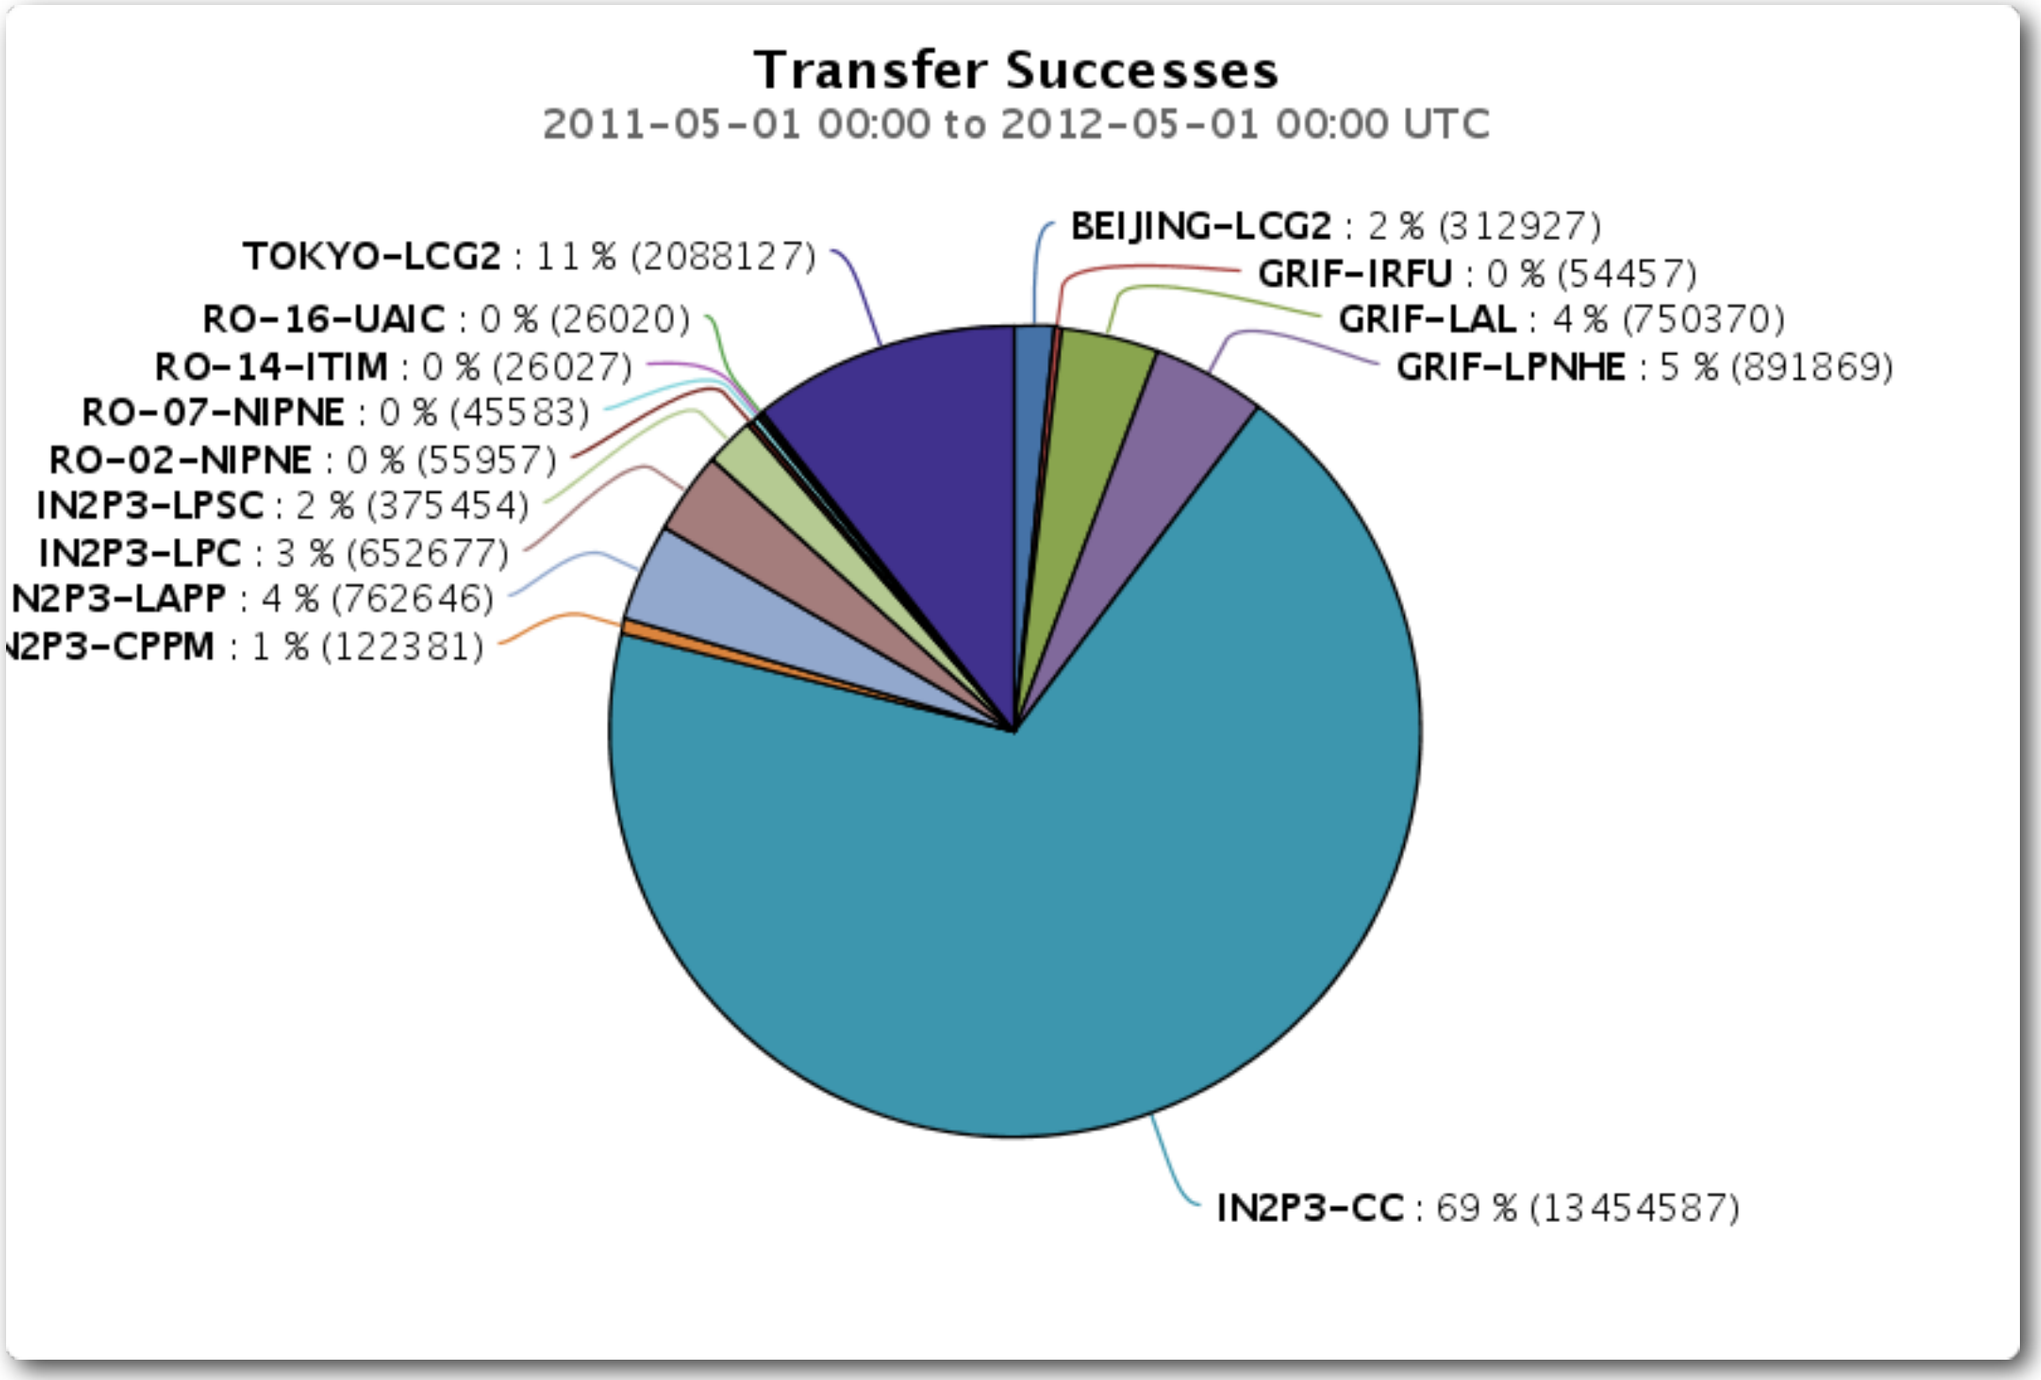

Export from FR-cloud [May, 2011 - May, 2012] : **5.6 PB**

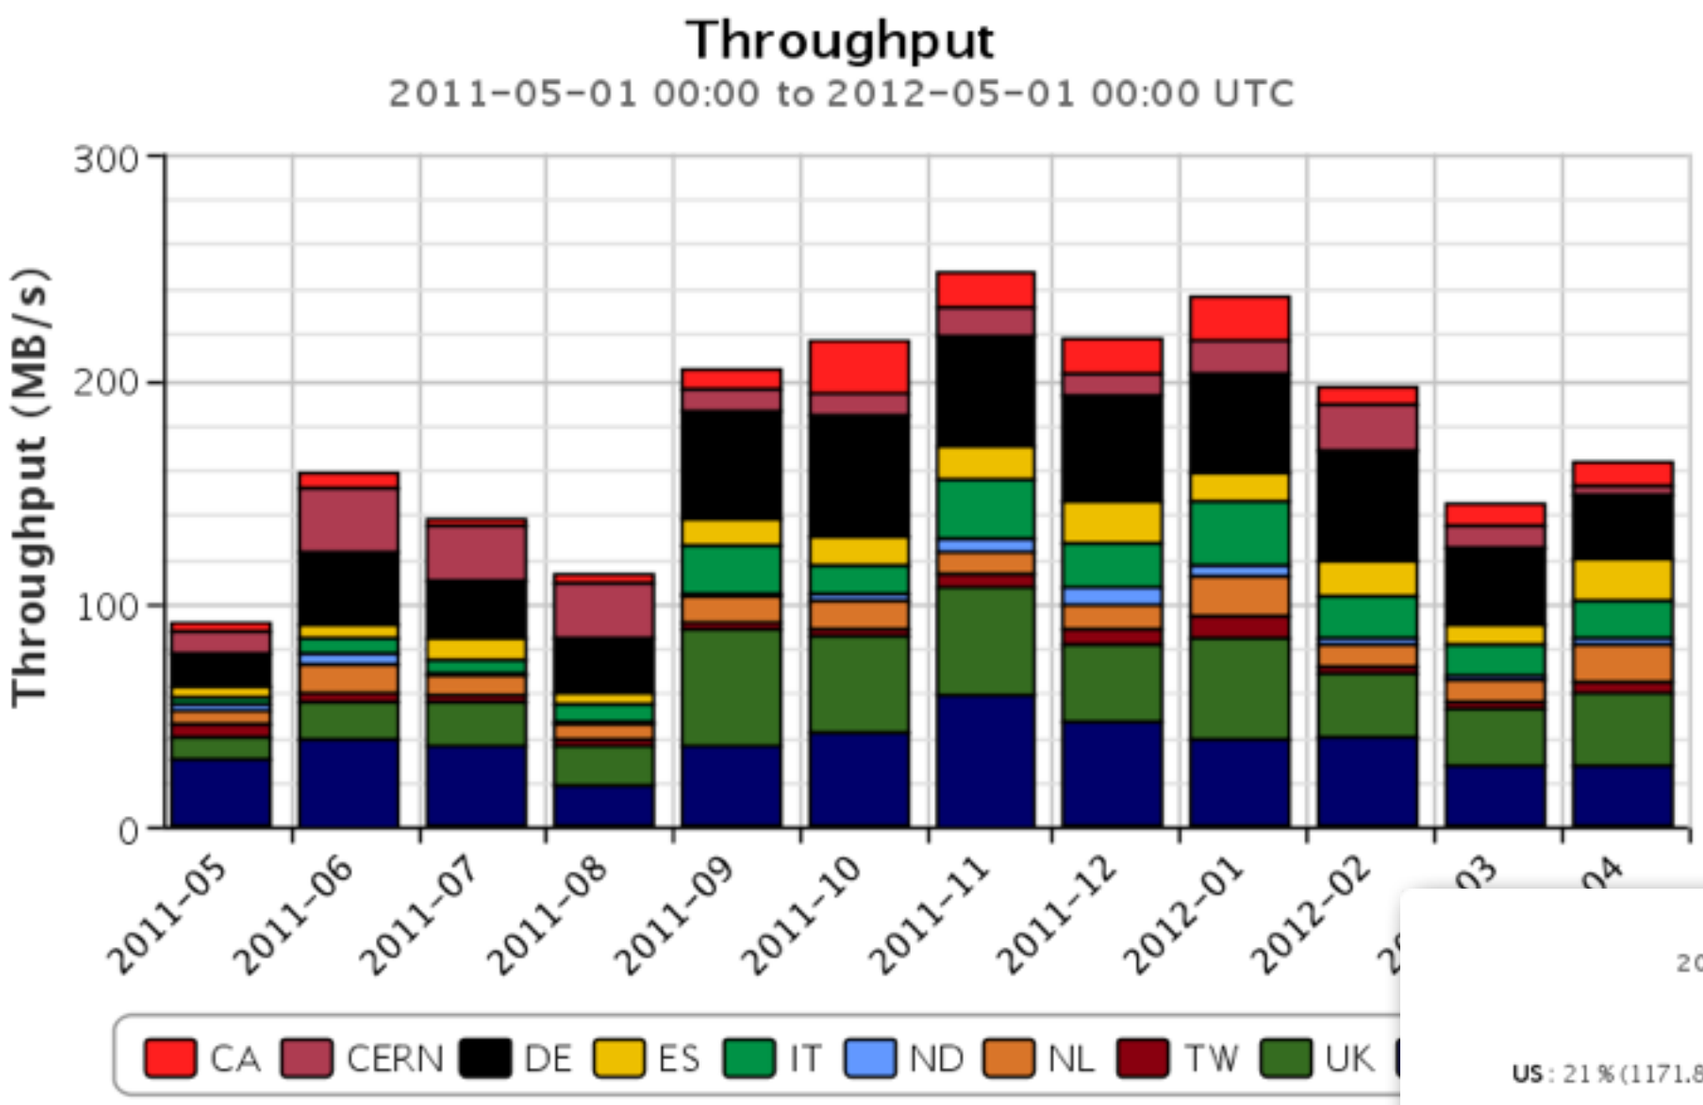

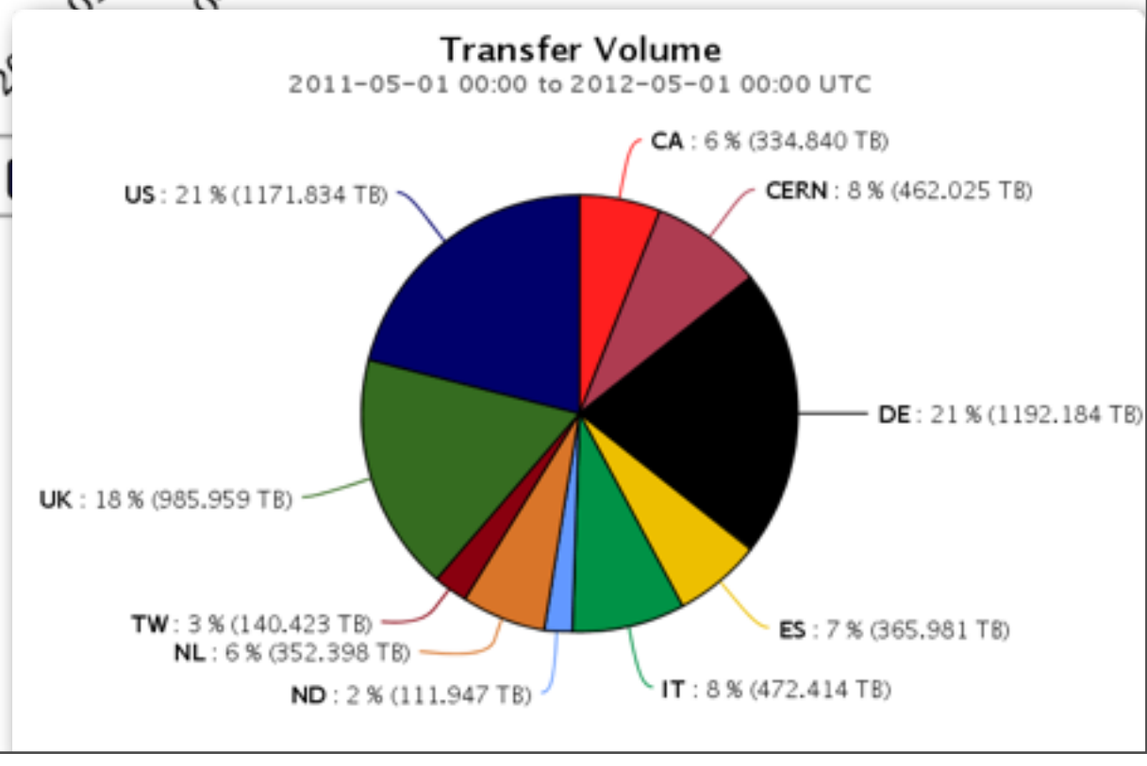

irfu

cea

saclay

export : source **87% from T1**

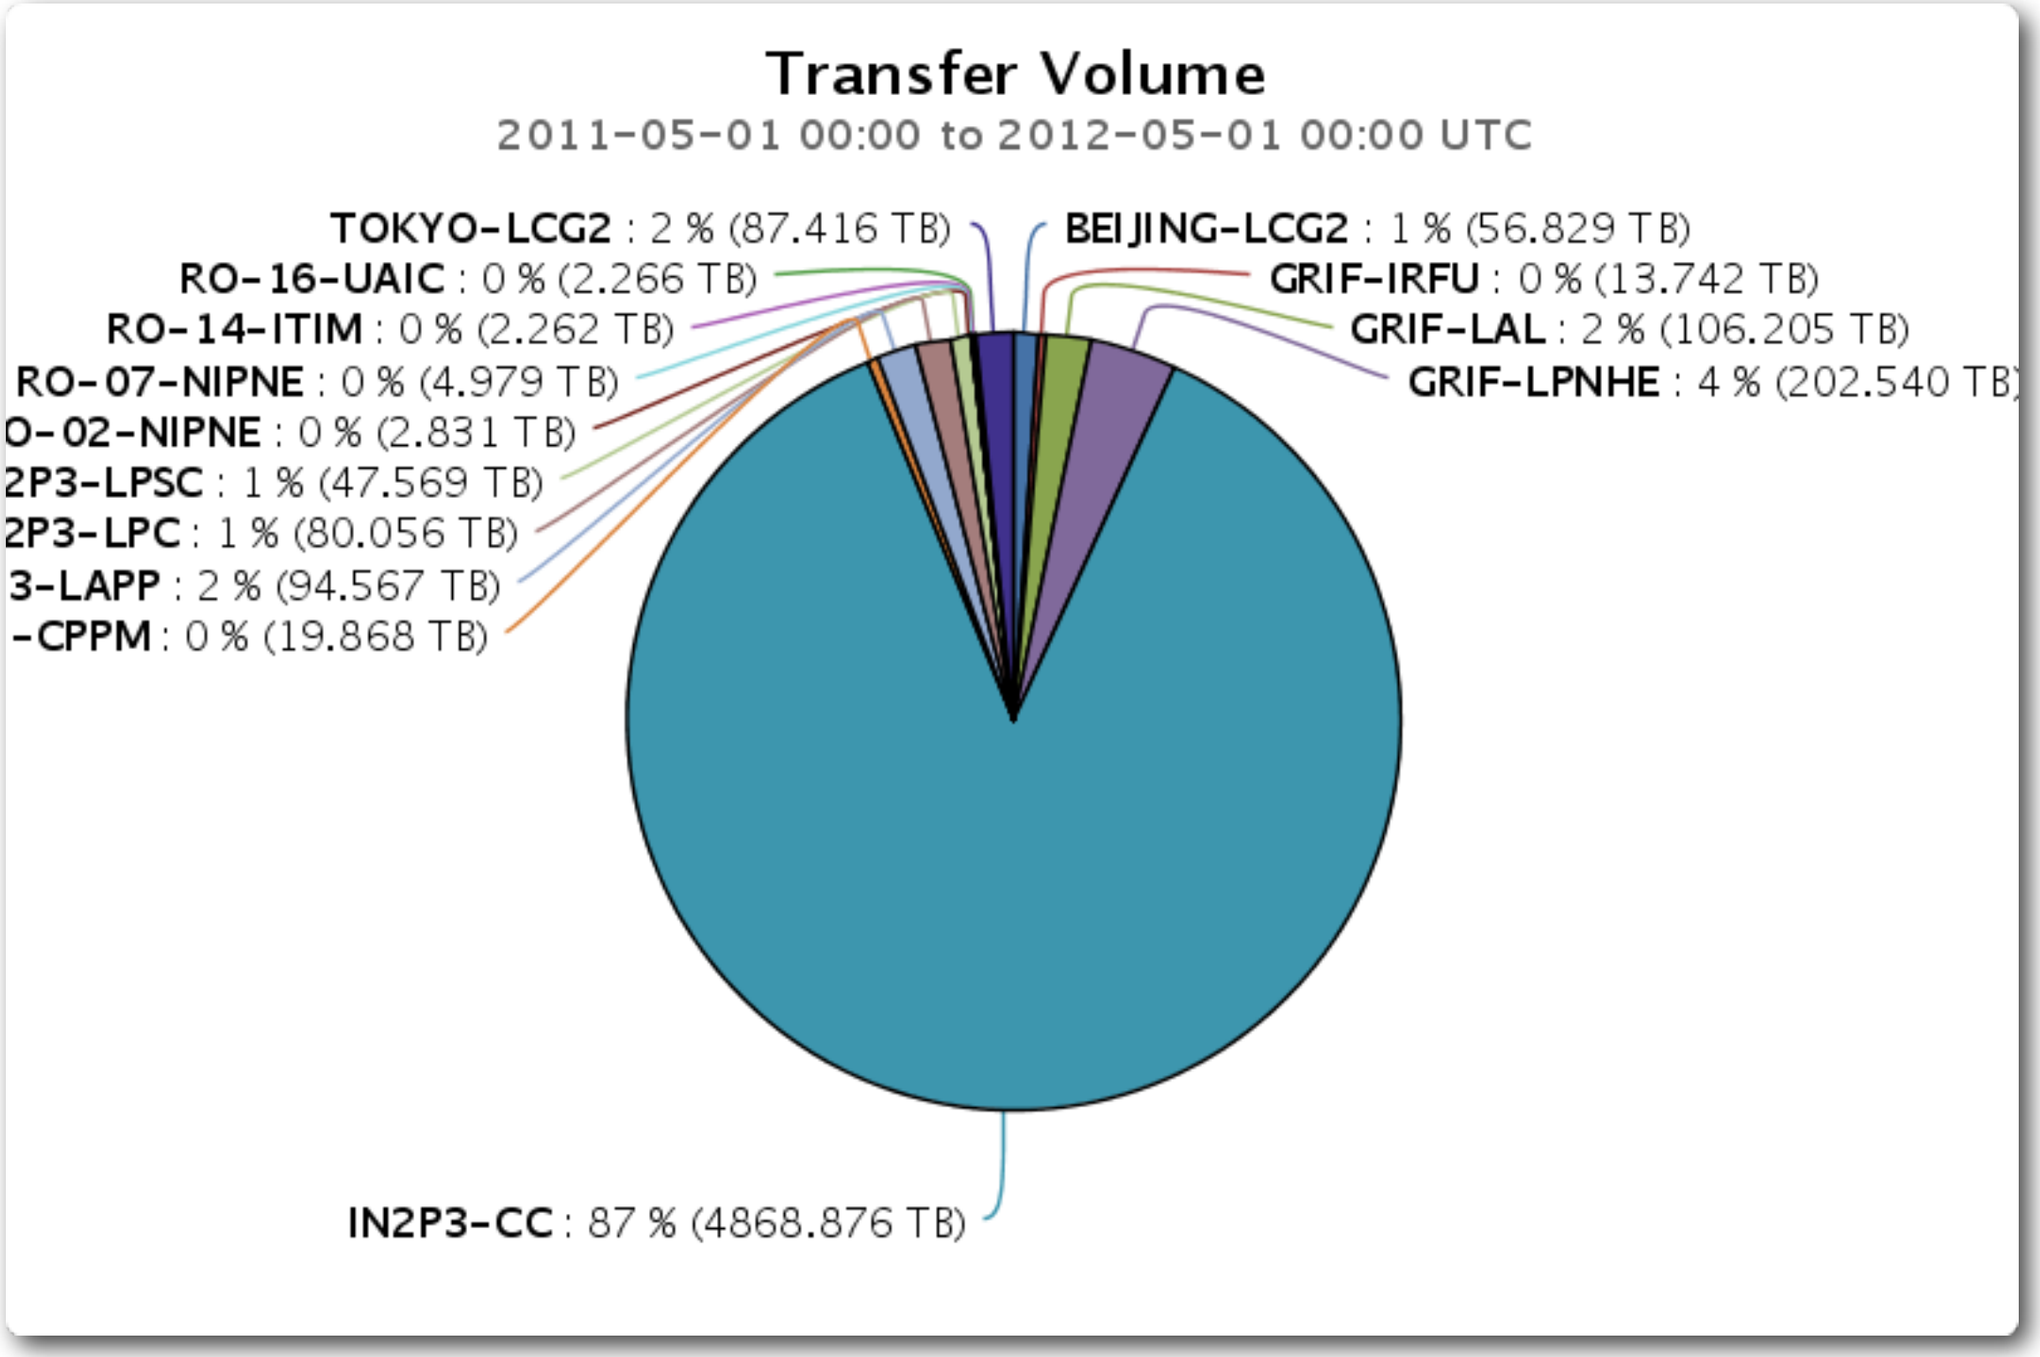

Intra cloud transfers [May,2011 - May,2012]

Not so uniform

Efficiency
Throughput
Successes
Errors
Registration:
Datasets
Files
Errors

0 % 100 %

100 %
0 %

SOURCES

DESTINATIONS

DESTINATIONS \ SOURCES	TOTAL-	BEIJING-LCG2+	GRIF-IRFU+	GRIF-LAL+	GRIF-LPNHE+	IN2P3-CC+	IN2P3-CPPM+	IN2P3-LAPP+	IN2P3-LPC+	IN2P3-LPSC+	RO-02-NIPNE+	RO-07-NIPNE+	RO-14-ITIM+	RO-16-UAIC+	TOKYO-LCG2+
TOTAL-	93 % 170 MB/s	90 % 3 MB/s	98 % 8 MB/s	89 % 4 MB/s	94 % 8 MB/s	92 % 117 MB/s	96 % 4 MB/s	97 % 6 MB/s	95 % 5 MB/s	93 % 3 MB/s	87 % 771 kB/s	86 % 2 MB/s	96 % 365 kB/s	94 % 466 kB/s	96 % 9 MB/s
BEIJING-LCG2+	96 % 8 MB/s	100 % 0 kB/s	94 % 72 kB/s	94 % 160 kB/s	86 % 378 kB/s	96 % 6 MB/s	99 % 50 kB/s	100 % 300 kB/s	94 % 212 kB/s	98 % 137 kB/s	91 % 17 kB/s	70 % 14 kB/s	100 % 1 kB/s	98 % 1 kB/s	97 % 476 kB/s
GRIF-IRFU+	93 % 32 MB/s	64 % 178 kB/s	100 % 191 kB/s	89 % 539 kB/s	91 % 1 MB/s	96 % 27 MB/s	90 % 327 kB/s	86 % 645 kB/s	72 % 464 kB/s	80 % 375 kB/s	82 % 125 kB/s	77 % 112 kB/s	54 % 0 kB/s	72 % 1 kB/s	68 % 238 kB/s
GRIF-LAL+	89 % 8 MB/s	83 % 44 kB/s	71 % 233 kB/s	100 % 6 kB/s	92 % 235 kB/s	89 % 6 MB/s	80 % 54 kB/s	86 % 177 kB/s	94 % 113 kB/s	76 % 75 kB/s	83 % 7 kB/s	75 % 25 kB/s	67 % 1 kB/s	72 % 1 kB/s	96 % 377 kB/s
GRIF-LPNHE+	94 % 11 MB/s	91 % 147 kB/s	95 % 776 kB/s	96 % 533 kB/s	85 % 163 kB/s	95 % 8 MB/s	82 % 87 kB/s	95 % 512 kB/s	80 % 286 kB/s	97 % 133 kB/s	86 % 23 kB/s	79 % 32 kB/s	99 % 1 kB/s	95 % 1 kB/s	97 % 291 kB/s
IN2P3-CC+	95 % 43 MB/s	94 % 2 MB/s	99 % 5 MB/s	97 % 2 MB/s	95 % 3 MB/s	88 % 13 MB/s	96 % 3 MB/s	98 % 3 MB/s	98 % 2 MB/s	94 % 2 MB/s	88 % 490 kB/s	86 % 1 MB/s	97 % 353 kB/s	94 % 450 kB/s	98 % 6 MB/s
IN2P3-CPPM+	84 % 16 MB/s	28 % 169 kB/s	91 % 174 kB/s	20 % 203 kB/s	94 % 396 kB/s	89 % 14 MB/s	100 % 62 kB/s	96 % 349 kB/s	51 % 232 kB/s	92 % 132 kB/s	94 % 27 kB/s	79 % 35 kB/s	99 % 1 kB/s	100 % 1 kB/s	88 % 218 kB/s
IN2P3-LAPP+	97 % 12 MB/s	99 % 103 kB/s	96 % 277 kB/s	94 % 167 kB/s	86 % 360 kB/s	97 % 10 MB/s	99 % 100 kB/s	95 % 103 kB/s	97 % 223 kB/s	100 % 95 kB/s	80 % 10 kB/s	72 % 36 kB/s	100 % 1 kB/s	99 % 1 kB/s	99 % 212 kB/s
IN2P3-LPC+	92 % 13 MB/s	93 % 168 kB/s	92 % 589 kB/s	83 % 265 kB/s	91 % 959 kB/s	93 % 10 MB/s	64 % 159 kB/s	95 % 490 kB/s	100 % 23 kB/s	94 % 145 kB/s	68 % 18 kB/s	80 % 90 kB/s	69 % 1 kB/s	70 % 1 kB/s	91 % 534 kB/s
IN2P3-LPSC+	92 % 8 MB/s	96 % 74 kB/s	99 % 39 kB/s	91 % 154 kB/s	94 % 252 kB/s	92 % 6 MB/s	96 % 76 kB/s	93 % 421 kB/s	96 % 228 kB/s	85 % 8 kB/s	92 % 11 kB/s	91 % 14 kB/s	88 % 2 kB/s	91 % 1 kB/s	93 % 103 kB/s
RO-02-NIPNE+	65 % 3 MB/s	73 % 102 kB/s	45 % 46 kB/s	64 % 28 kB/s	70 % 118 kB/s	66 % 2 MB/s	34 % 17 kB/s	69 % 82 kB/s	44 % 32 kB/s	54 % 29 kB/s		70 % 24 kB/s	69 % 1 kB/s	80 % 2 kB/s	65 % 44 kB/s
RO-07-NIPNE+	81 % 6 MB/s	72 % 79 kB/s	52 % 163 kB/s	55 % 84 kB/s	64 % 207 kB/s	84 % 5 MB/s	70 % 96 kB/s	71 % 175 kB/s	69 % 183 kB/s	80 % 90 kB/s	72 % 29 kB/s		81 % 1 kB/s	68 % 5 kB/s	72 % 228 kB/s
RO-14-ITIM+	91 % 817 kB/s	89 % 3 kB/s	100 % 3 kB/s	98 % 5 kB/s	99 % 18 kB/s	92 % 769 kB/s	99 % 3 kB/s	99 % 2 kB/s	69 % 3 kB/s	100 % 3 kB/s	86 % 2 kB/s	51 % 5 kB/s		83 % 1 kB/s	91 % 2 kB/s
RO-16-UAIC+	88 % 1 MB/s	91 % 3 kB/s	90 % 3 kB/s	88 % 5 kB/s	89 % 11 kB/s	88 % 1 MB/s	89 % 2 kB/s	94 % 4 kB/s	56 % 3 kB/s	89 % 3 kB/s	69 % 3 kB/s	53 % 3 kB/s	80 % 1 kB/s		94 % 6 kB/s
TOKYO-LCG2+	91 % 9 MB/s	39 % 161 kB/s	56 % 61 kB/s	41 % 98 kB/s	95 % 685 kB/s	95 % 8 MB/s	89 % 77 kB/s	99 % 215 kB/s	22 % 69 kB/s	99 % 96 kB/s	99 % 10 kB/s	83 % 47 kB/s	100 % 1 kB/s	98 % 1 kB/s	99 % 59 kB/s

irfu
cea
saclay

FR-cloud **intra** volume of data transfer : **5.4 PB**

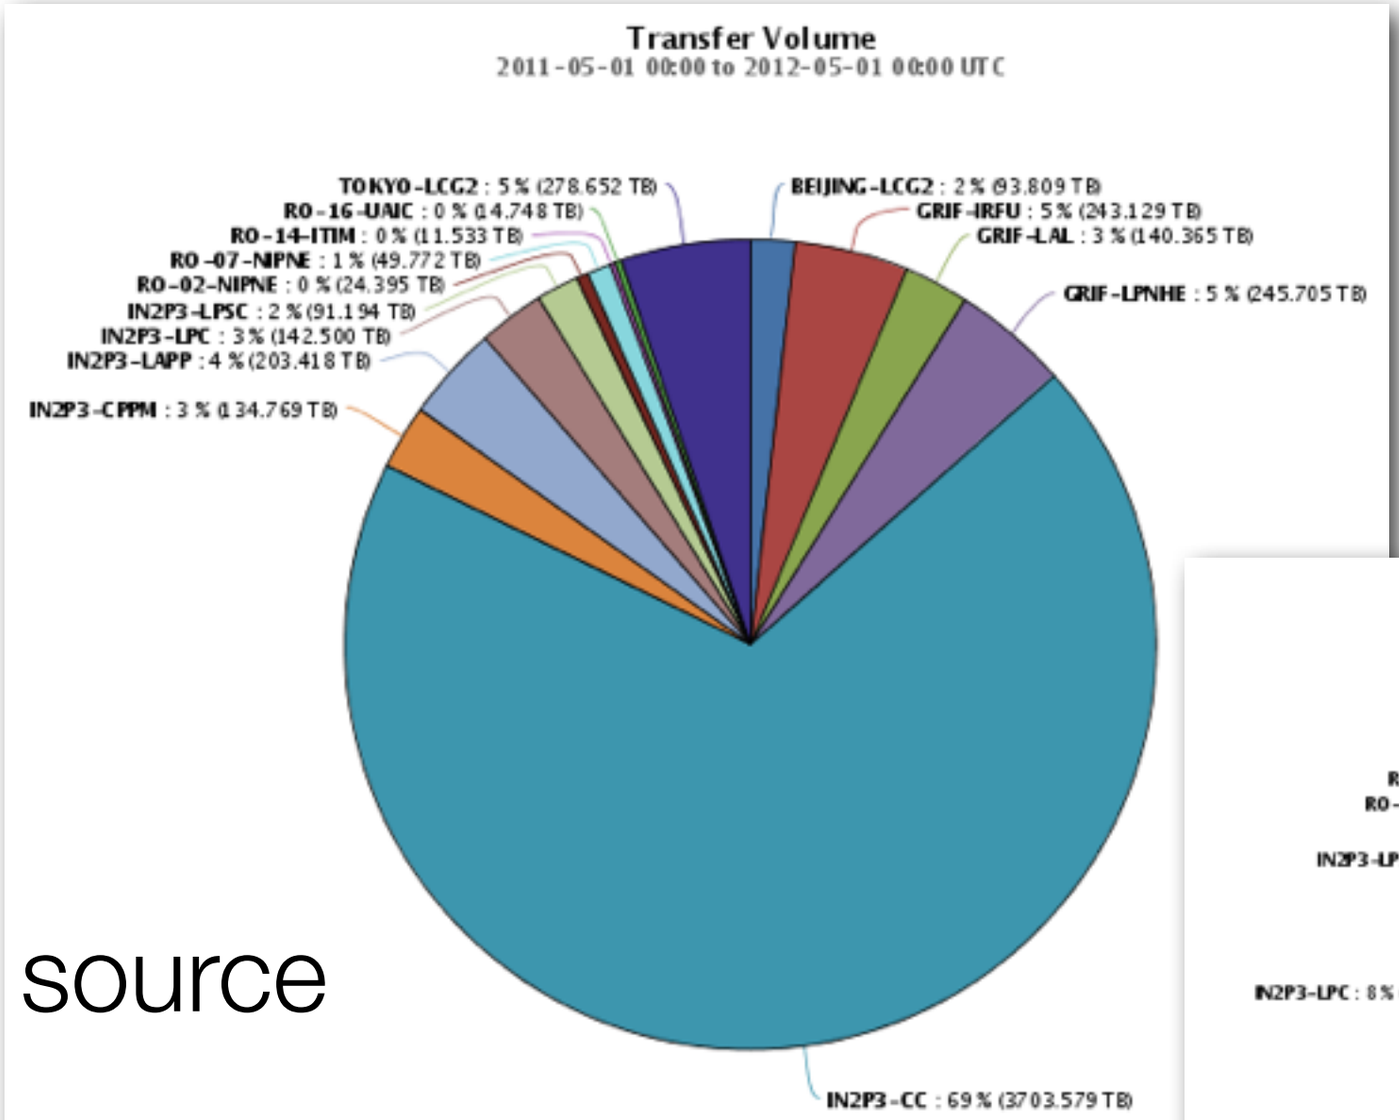

source

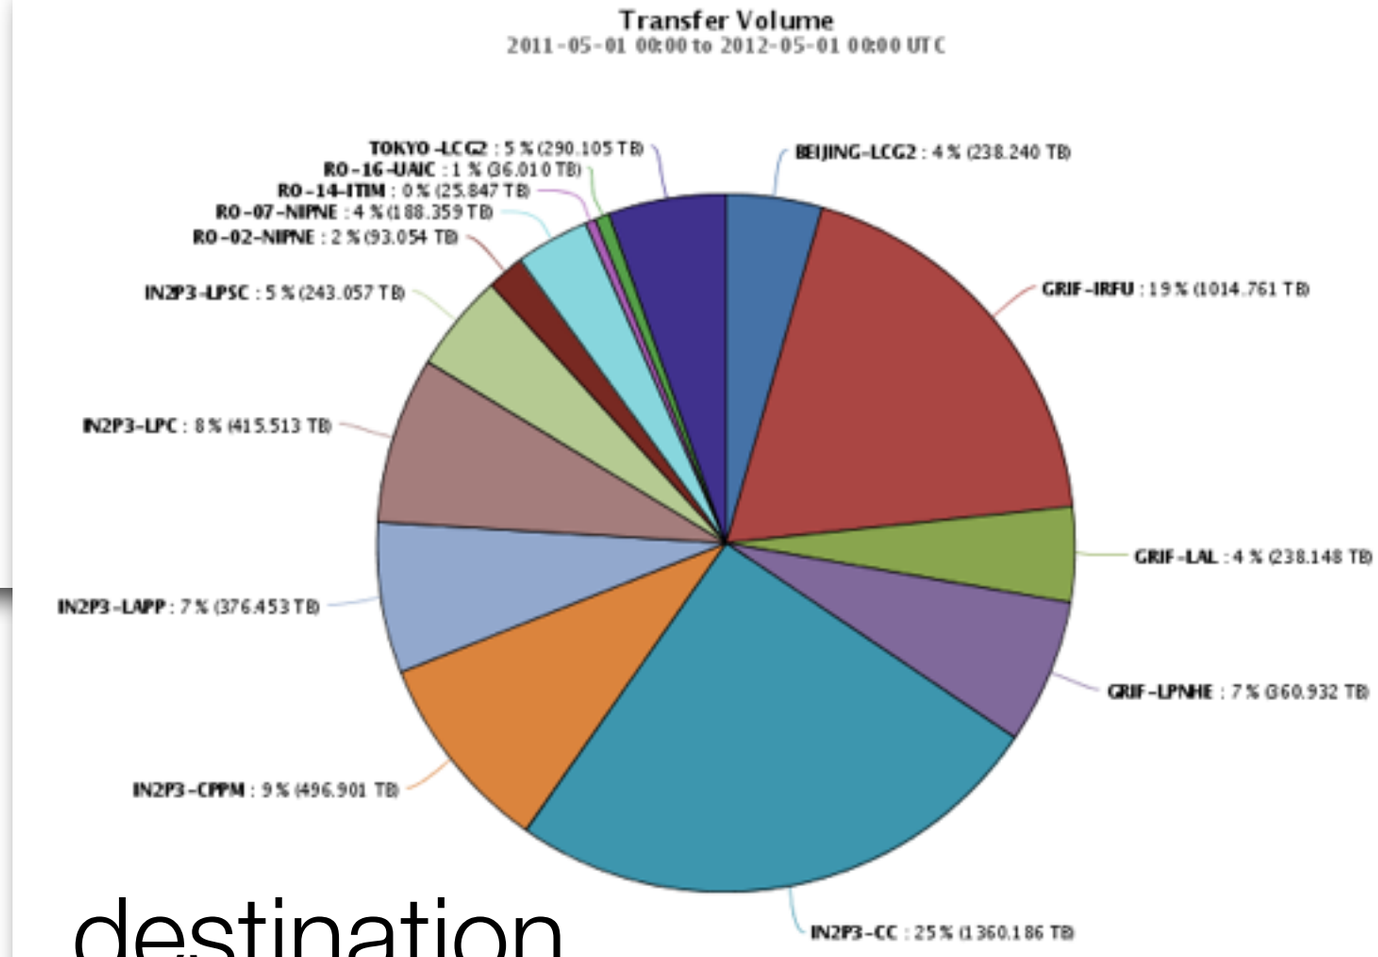

destination

irfu

cea

saclay

user subscription destination on FR-cloud : **1.5 PB**

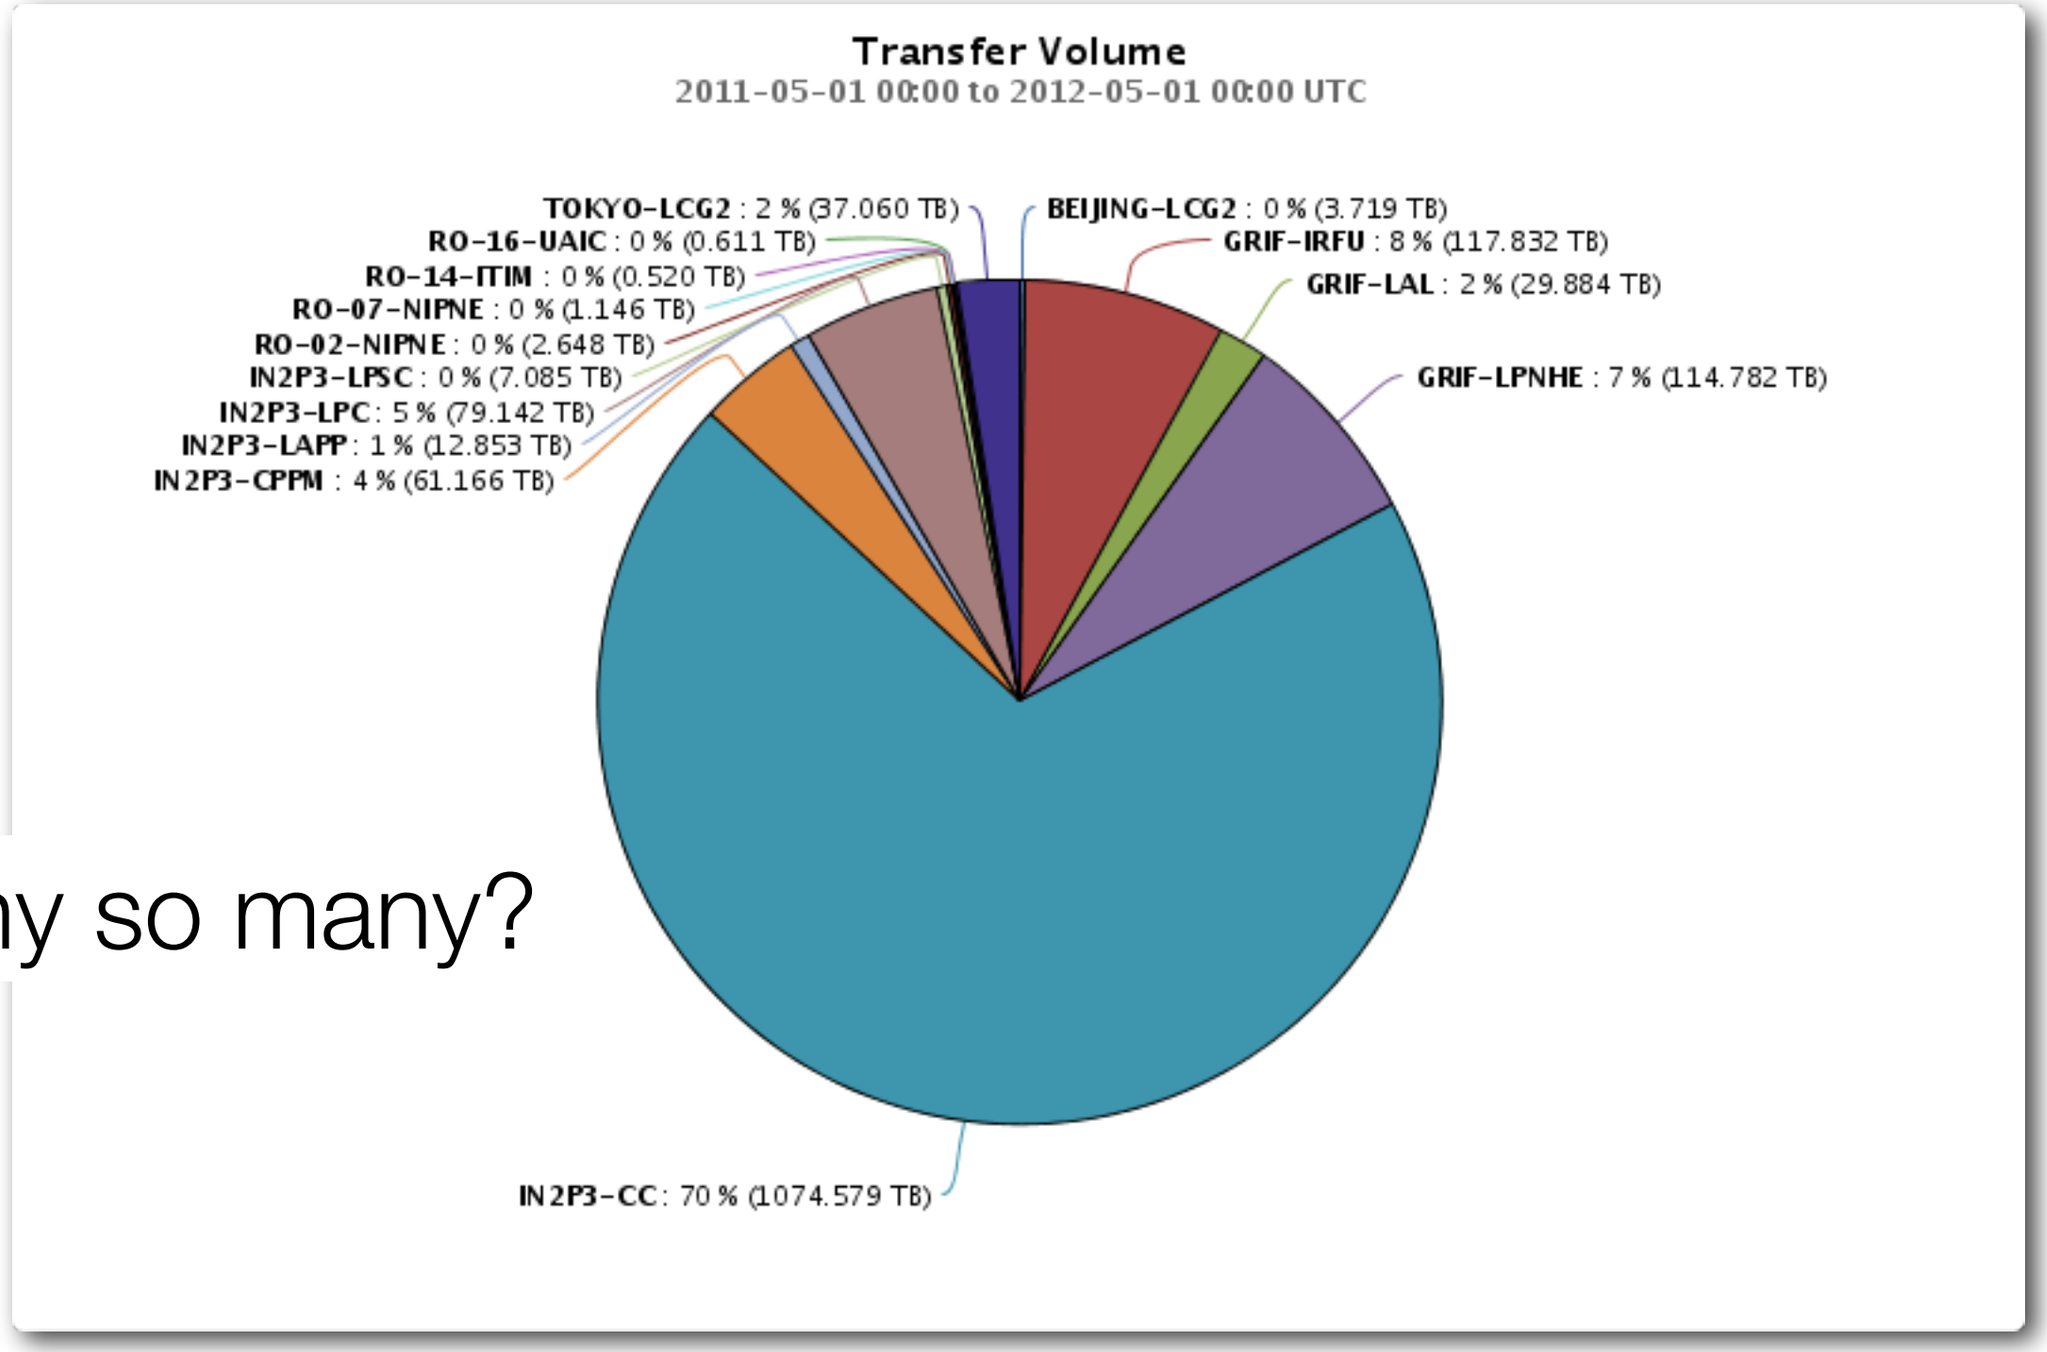

why so many?

correlated to ‘local’ activity at sites ?

irfu

cea

saclay

The ATLAS Data Model has changed

USE THE NETWORK

- Moved away from the historical model
- 4 recurring themes:
  - **Flat(ter) hierarchy**: Any site can replicate data from any other site
  - **Multi Cloud Production**
    - Need to replicate output files to remote Tier-1
  - **Dynamic data caching**: Analysis sites receive datasets from any other site "on demand" based on usage pattern
    - Possibly in combination with pre-placement of data sets by centrally managed replication of datasets
  - **Remote data access**: local jobs accessing data stored at remote sites

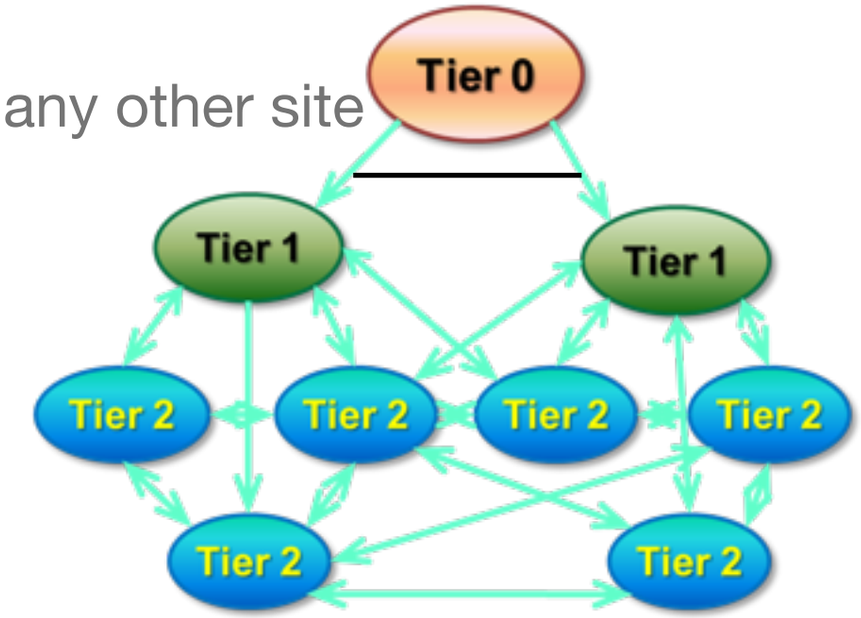

- **ATLAS is now heavily relying on multi-domain networks and needs decent e2e network monitoring**

Thank you Dan v.d. Steer

Data Processing Model Revised

CERN Analysis Facility

10 Tier-1 centers + CERN
Aggregation of produced data
Source for further distribution

Tier-1

Tier-2

Tier-2

Tier-2-D

job input

Tier-1

Tier-2

Tier-2-D

Tier-1

Tier-2

Tier-2

Tier-2-Direct : 'Super' Tier-2s

ATLAS sites and connectivity

- ATLAS computing model is evolving
  - Experience of one year of data taking
  - Tools and monitoring are getting more mature

Aim :
all FR-cloud T2s
should be T2D

- New category of sites : Direct T2s (**T2Ds**)
  - primary hosts for datasets (analysis) and for group analysis
  - get and send data from different sites
- Connectivity is important for T2Ds (not only…)

validated T2Ds

- CA : CA-SCINET-T2, CA-VICTORIA-WESTGRID-T2, SFU-LCG2_DATADISK, CA-MCGILL
- DE : DESY-HH, DESY-ZN, MPPMU, LRZ-LMU, CSCS-LCG2, GOEGRID, UNI-FREIBURG,
- ES : IFIC-LCG2, IFAE, UAM-LCG2, LIP-LISBON, NCG-INGRID-PT
- FR : GRIF-LPNHE, GRIF-LAL, IN2P3-LAPP, IN2P3-LPC, IN2P3-LPSC, BEIJING-LCG2
- IT : INFN-NAPOLI-ATLAS, INFN-MILANO-ATLASC, INFN-ROMA1
- UK : UKI-LT2-QMUL, UKI-NORTHGRID-MAN-HEP, UKI-SCOTGRID-GLASGOW, UKI-NO
- US : AGLT2, MWT2_UC, NET2, SLACXRD, SWT2 CPB

Network performance monitoring

- Monitoring tools developed in 2011
  - ATLAS ‘sonar’ : ‘calibrated’ file transfers by the ATLAS Data Distribution system, from storage to storage
    - perfSONAR : network performance (throughput, latency, traceroute, ...)
  - Non stable network performances : also side effects on storage system at CCIN2P3

more on Monday

ES

ATLAS T2Ds

CERN IT Department

- T2Ds: Tier2s directly connected to Tier1s of different clouds
  - "directly" in the DDM topology

ATLAS transfers from T2D to any T1

- T2Ds should
  - demonstrate "good" connectivity from/to every T1/T2D
  - provide a certain level of commitment

Good connectivity

- T2Ds could be de-commissioned if they degrade

Performances monitored

- There is no "maximum number" of T2Ds
  - We should not create too many channels (FTS performance)

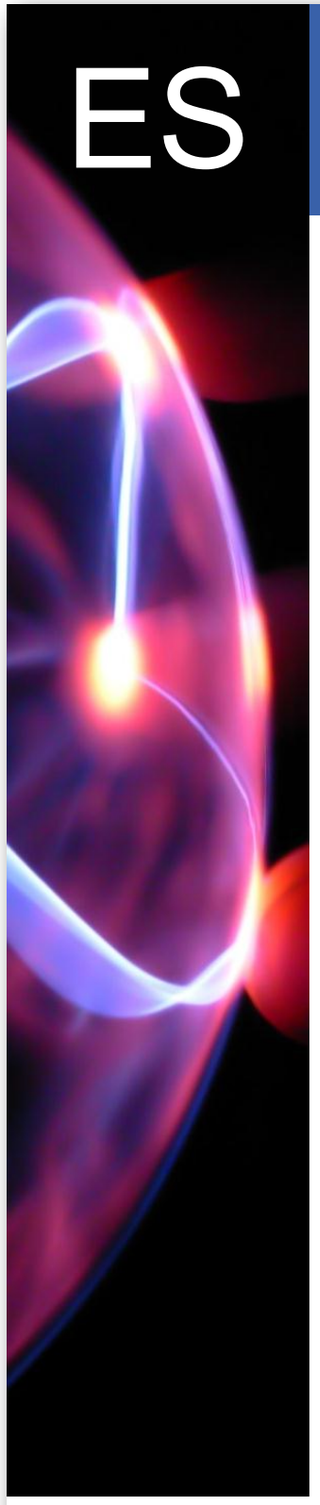

CERN IT Department
CH-1211 Geneva 23
Switzerland
www.cern.ch/it

T2D: revising the criteria

Current criteria

- All transfers from the candidate T2D to 10/12 T1s for big files (‘L’) must be above 5 MB/s during the last week and during 3 out of the 4 last weeks.
- All transfers from 10/12 T1s to the candidate T2D for big files must be above 5 MB/s during the last week and during 3 out of the 4 last weeks

T2D: revising the criteria

New criteria - under evaluation

- All transfers from the candidate T2D to **9**/12 T1s for big files ('L') must be above 5 MB/s during the last week and during 3 out of the **5** last weeks.
- All transfers from **9**/12 T1s to the candidate T2D for big files must be above 5 MB/s during the last week and during 3 out of the **5** last weeks

http://gnegri.web.cern.ch/gnegri/T2D/t2dStats.html

ES

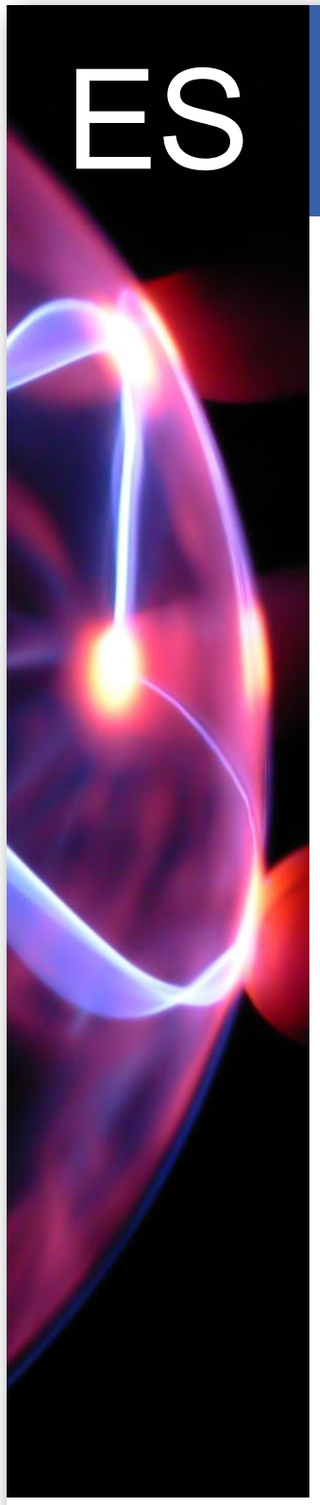

T2D evaluation

CERN IT Department

- New T2D candidates
  - Sites/clouds should contact Central Operations if they want to be a candidate T2D
- Site will start receiving more Sonar Tests (large files from/to all T1s and T2Ds)
  - If the performance is acceptable, results are reported to the ADC Weekly Meeting and the site is declared a T2D
  - Need for FTS configuration of channels at T1s and this is done approx every 3 months
- So far, T2Ds candidates are monitored by squads
  - Cloud and sites themselves should take the necessary actions (monitoring, improving performance, reporting results to Central Operations)

CERN IT Department
CH-1211 Geneva 23
Switzerland
www.cern.ch/it

irfu
cea
saclay

perfSONAR(-PS) and ATLAS

Needed on every FR-cloud site

- Being deployed on T2D sites
- To measure latency and throughput (2 machines) between sites (matrix)
- Details in Shawn Mckee's (BNL) talk at last GDB
- Matrix monitoring available

http://psps.perfsonar.net/

Summary for LHCONE

- Our specific goal in setting up perfSONAR-PS for LHCONE is to acquire before and after network measurements for the selected early adopter sites. This is **not** the long-term network monitoring setup for LHCONE…that is TBD
- Details of which sites and how sites should setup the perfSONAR-PS installations is documented on the Twiki at: https://twiki.cern.ch/twiki/bin/view/LHCONE/SiteList
- In the next few slides I will highlight some of the relevant details

ATLAS transfers to/from Beijing over one year

Beijing → CCIN2P3
CCIN2P3 → Beijing

asymmetry (why?) in transfer rate performance reversed over last year

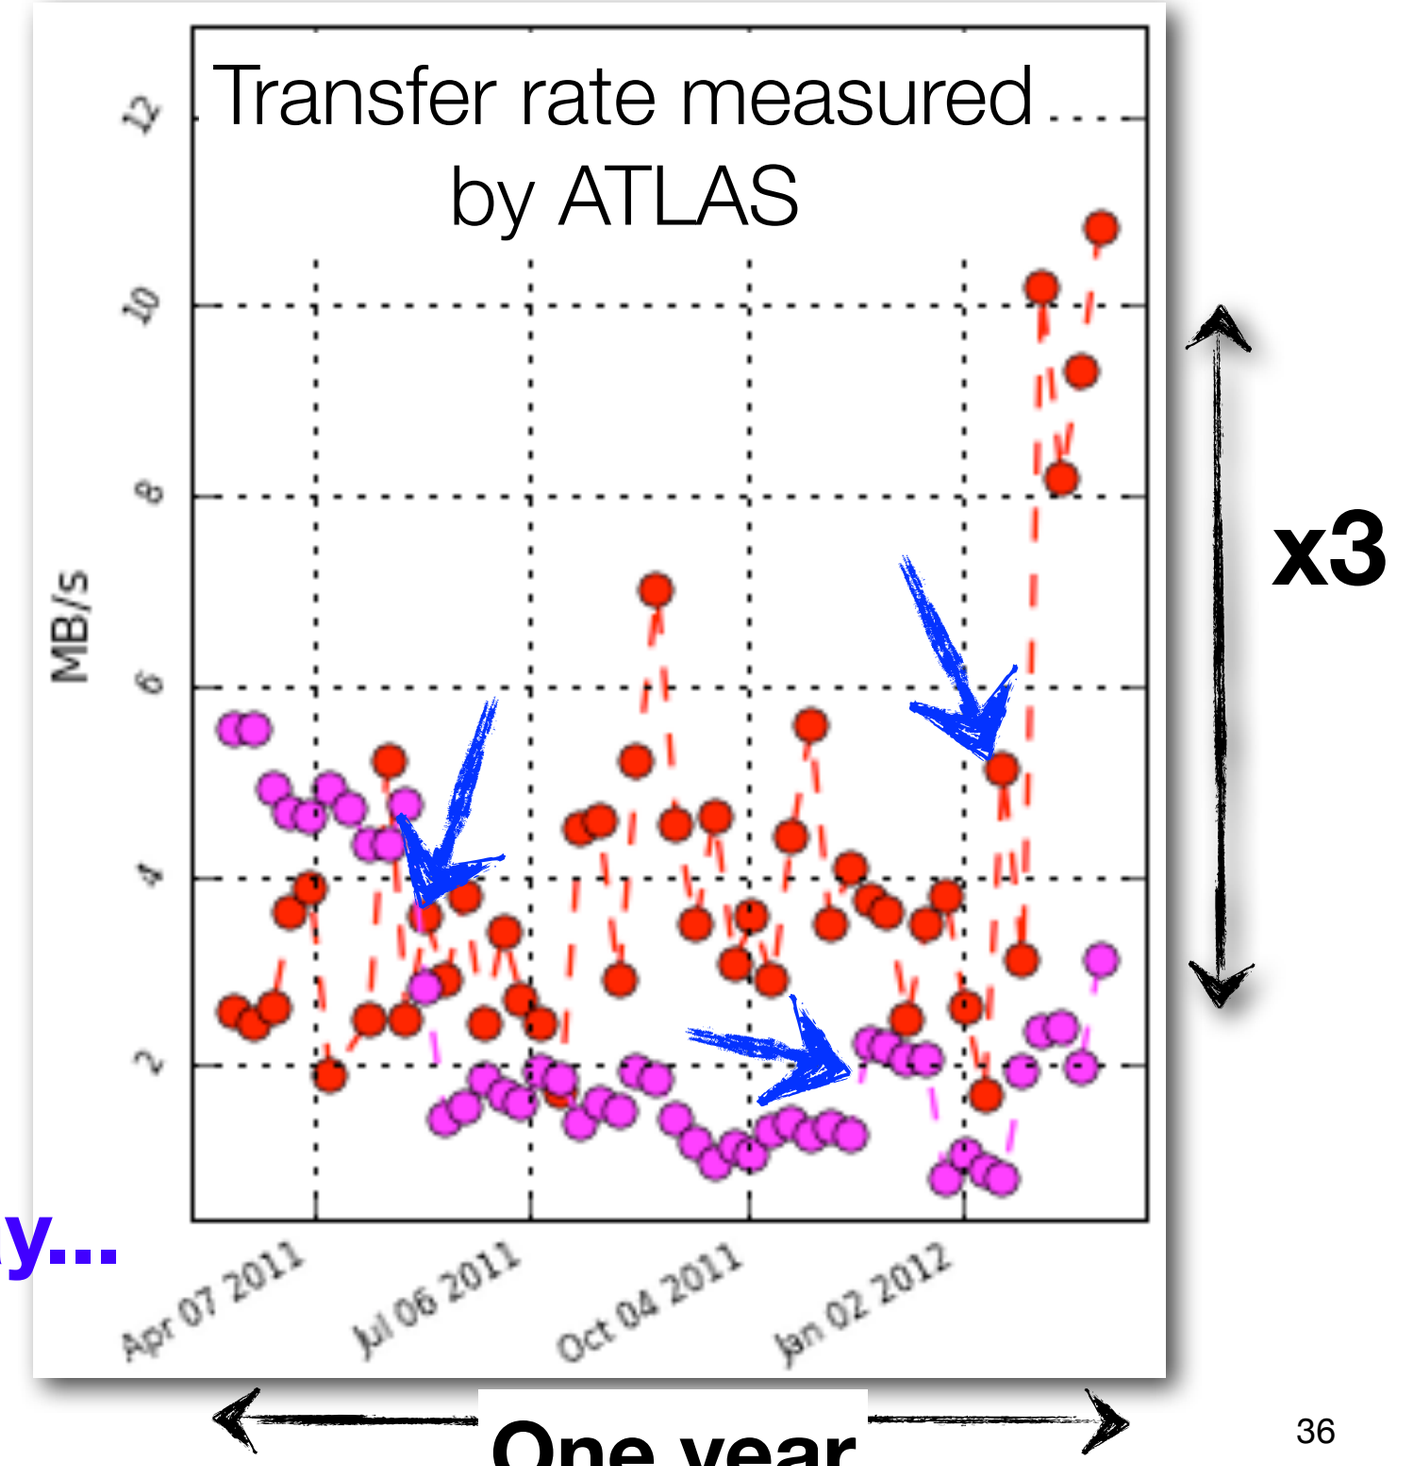

Each 'event' explained sometime after some delay...

irfu
cea
saclay

Network throughput measured with perfsonar

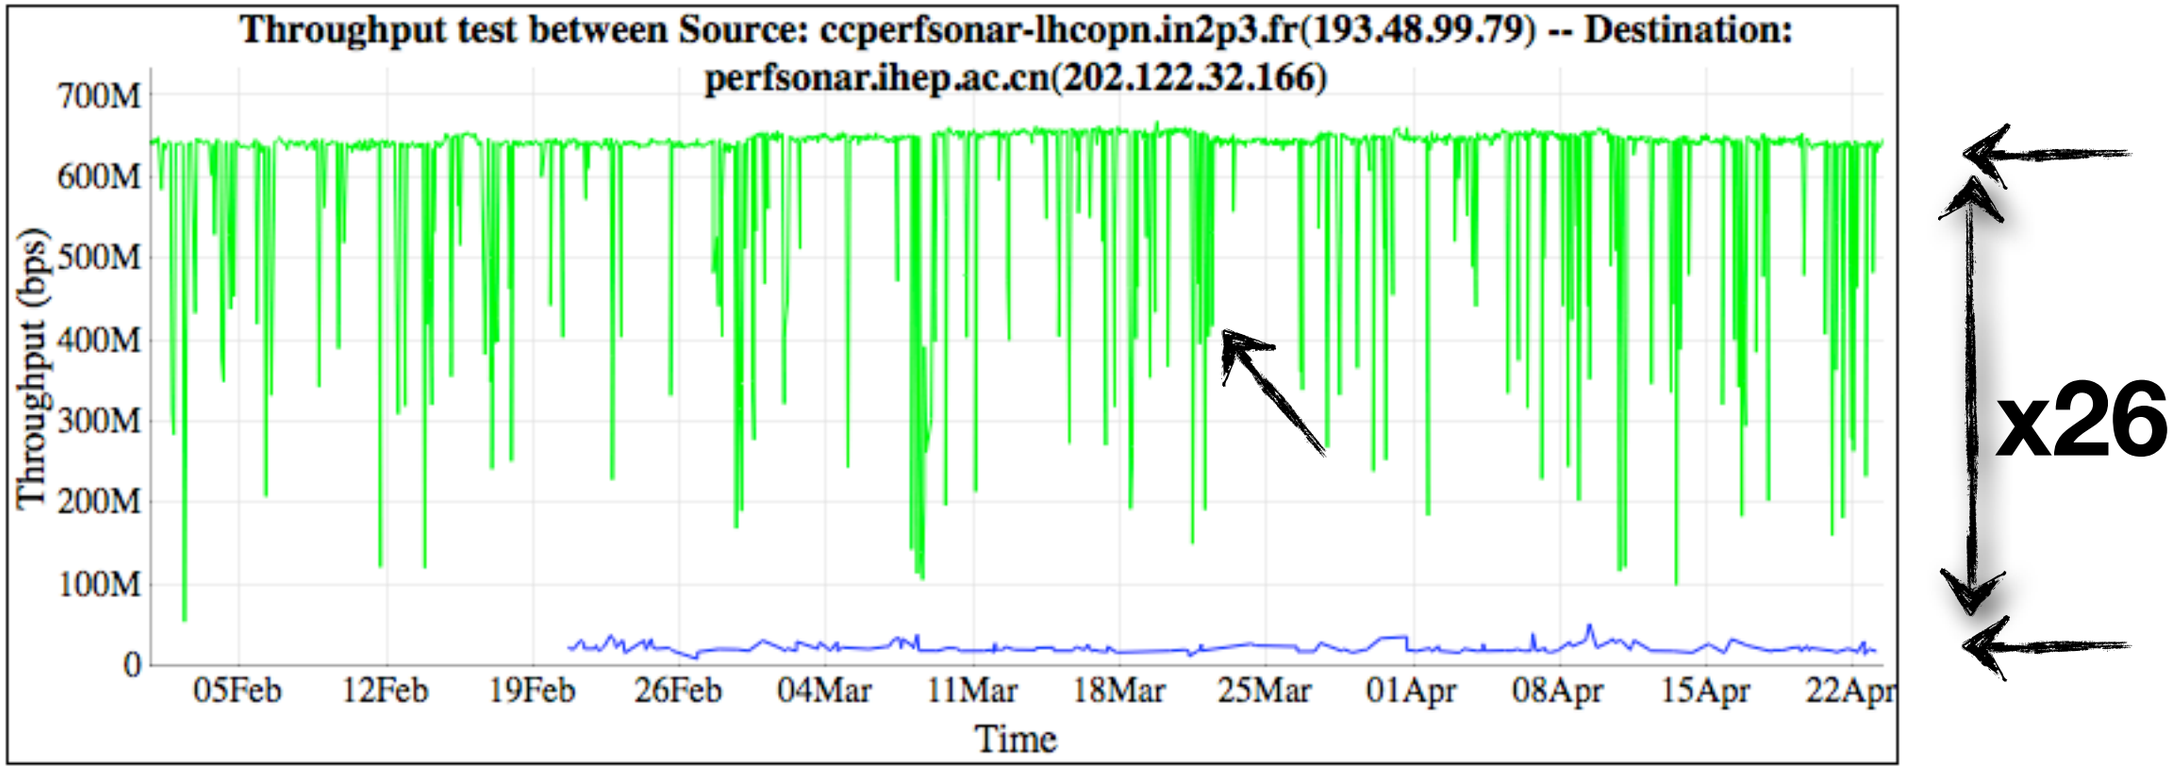

<- 1 month

1 month ->

Timezone: GMT+0200 (CEST)

Beijing → CCIN2P3 **CCIN2P3 → Beijing**

Direction	Max throughput(bps)	Mean throughput(bps)	Min throughput(bps)
Src-Dst	52.76M	23.19M	9.49M
Dst-Src	669.33M	616.84M	55.11M

how/Hide Link

Beijing to T1s

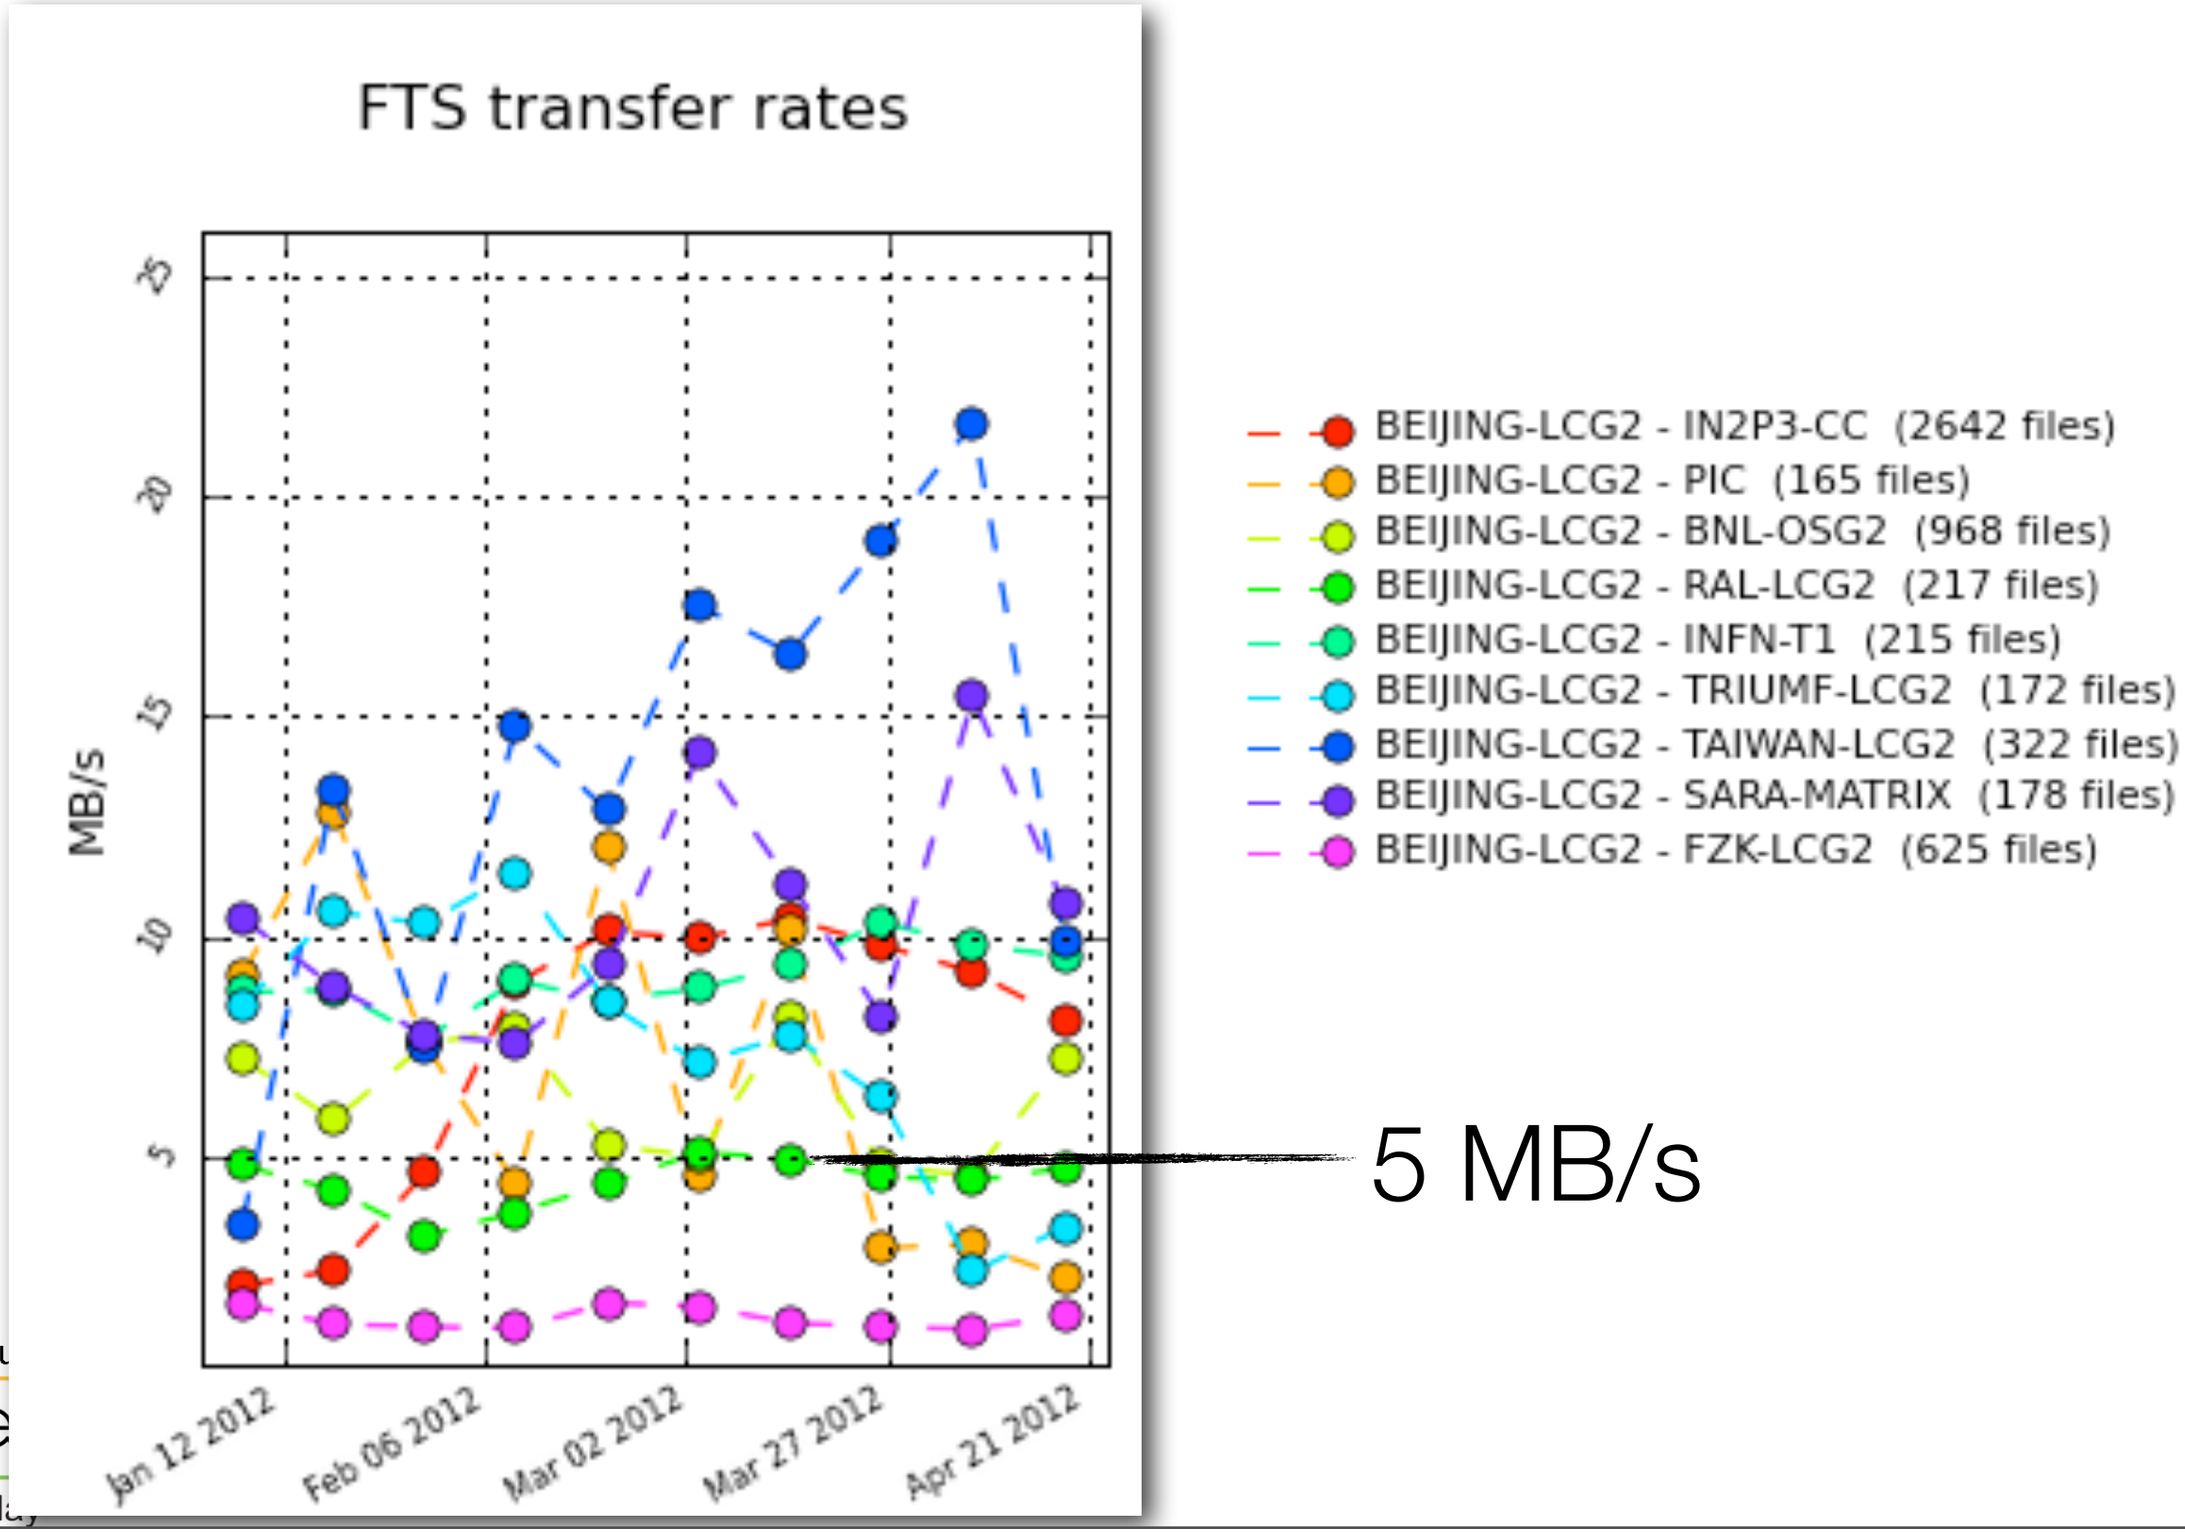

5 MB/s

T1s to Beijing

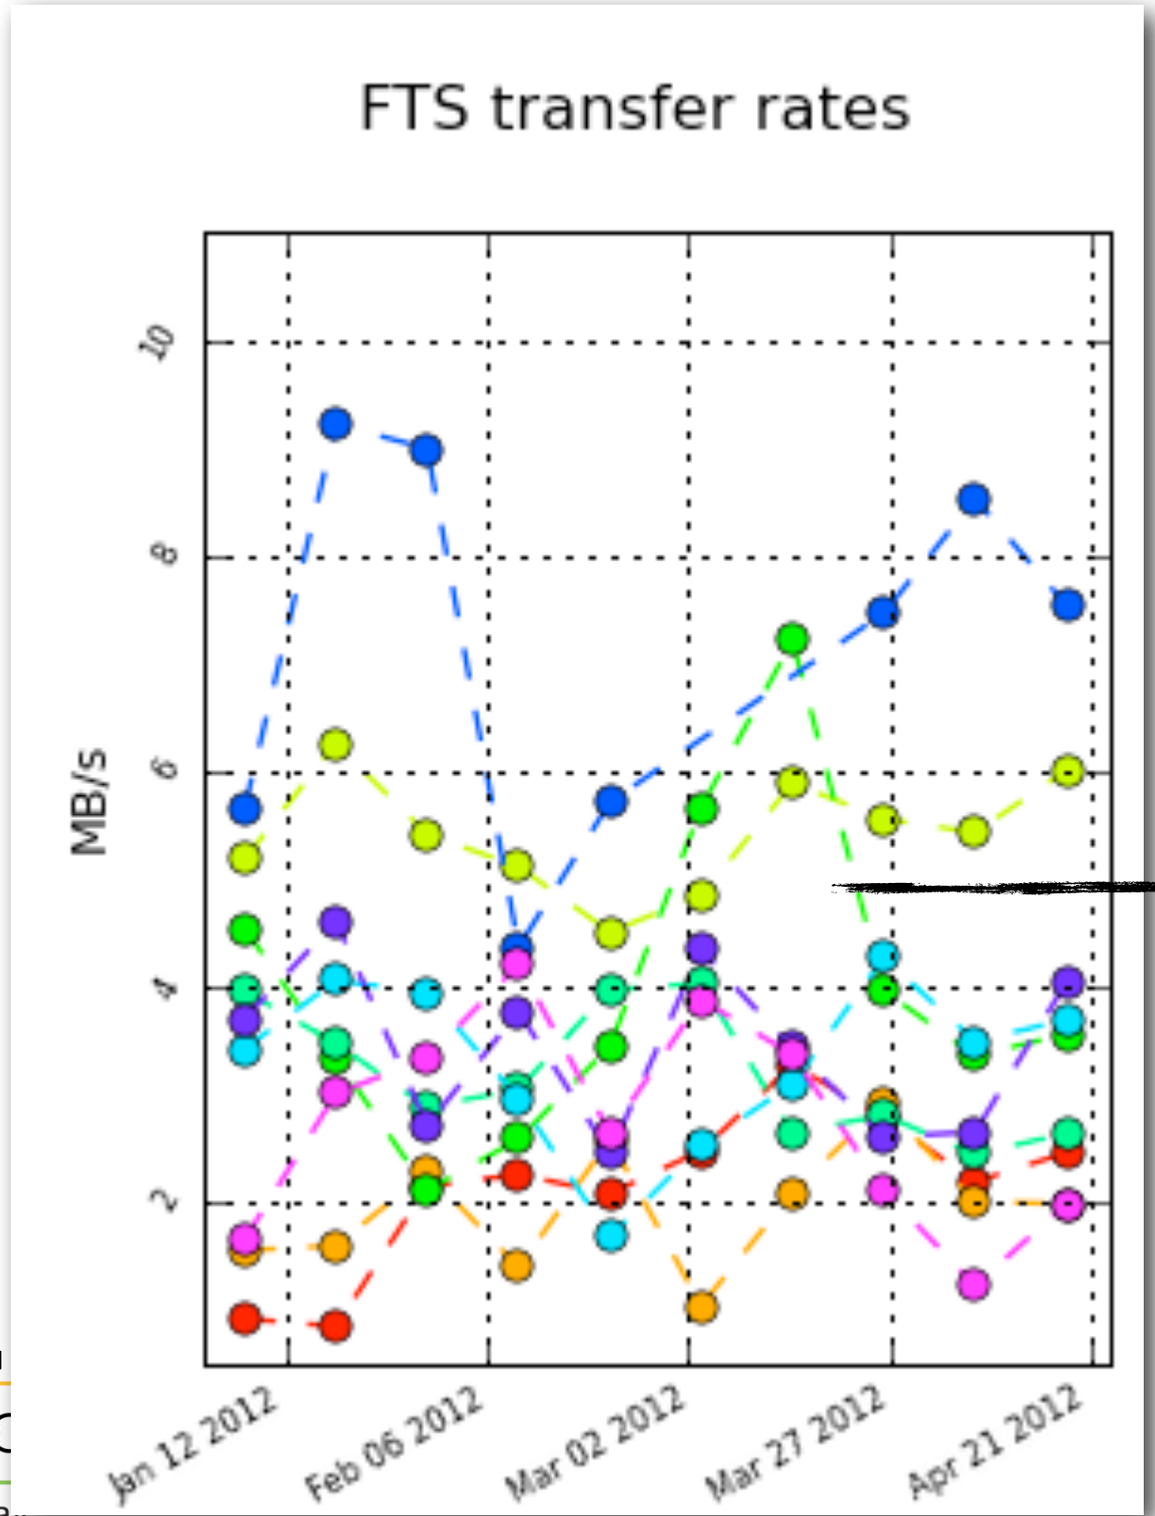

5 MB/s

each site is different

ATLAS transfers to/from Tokyo over one year

CCIN2P3 → Tokyo

Tokyo → CCIN2P3

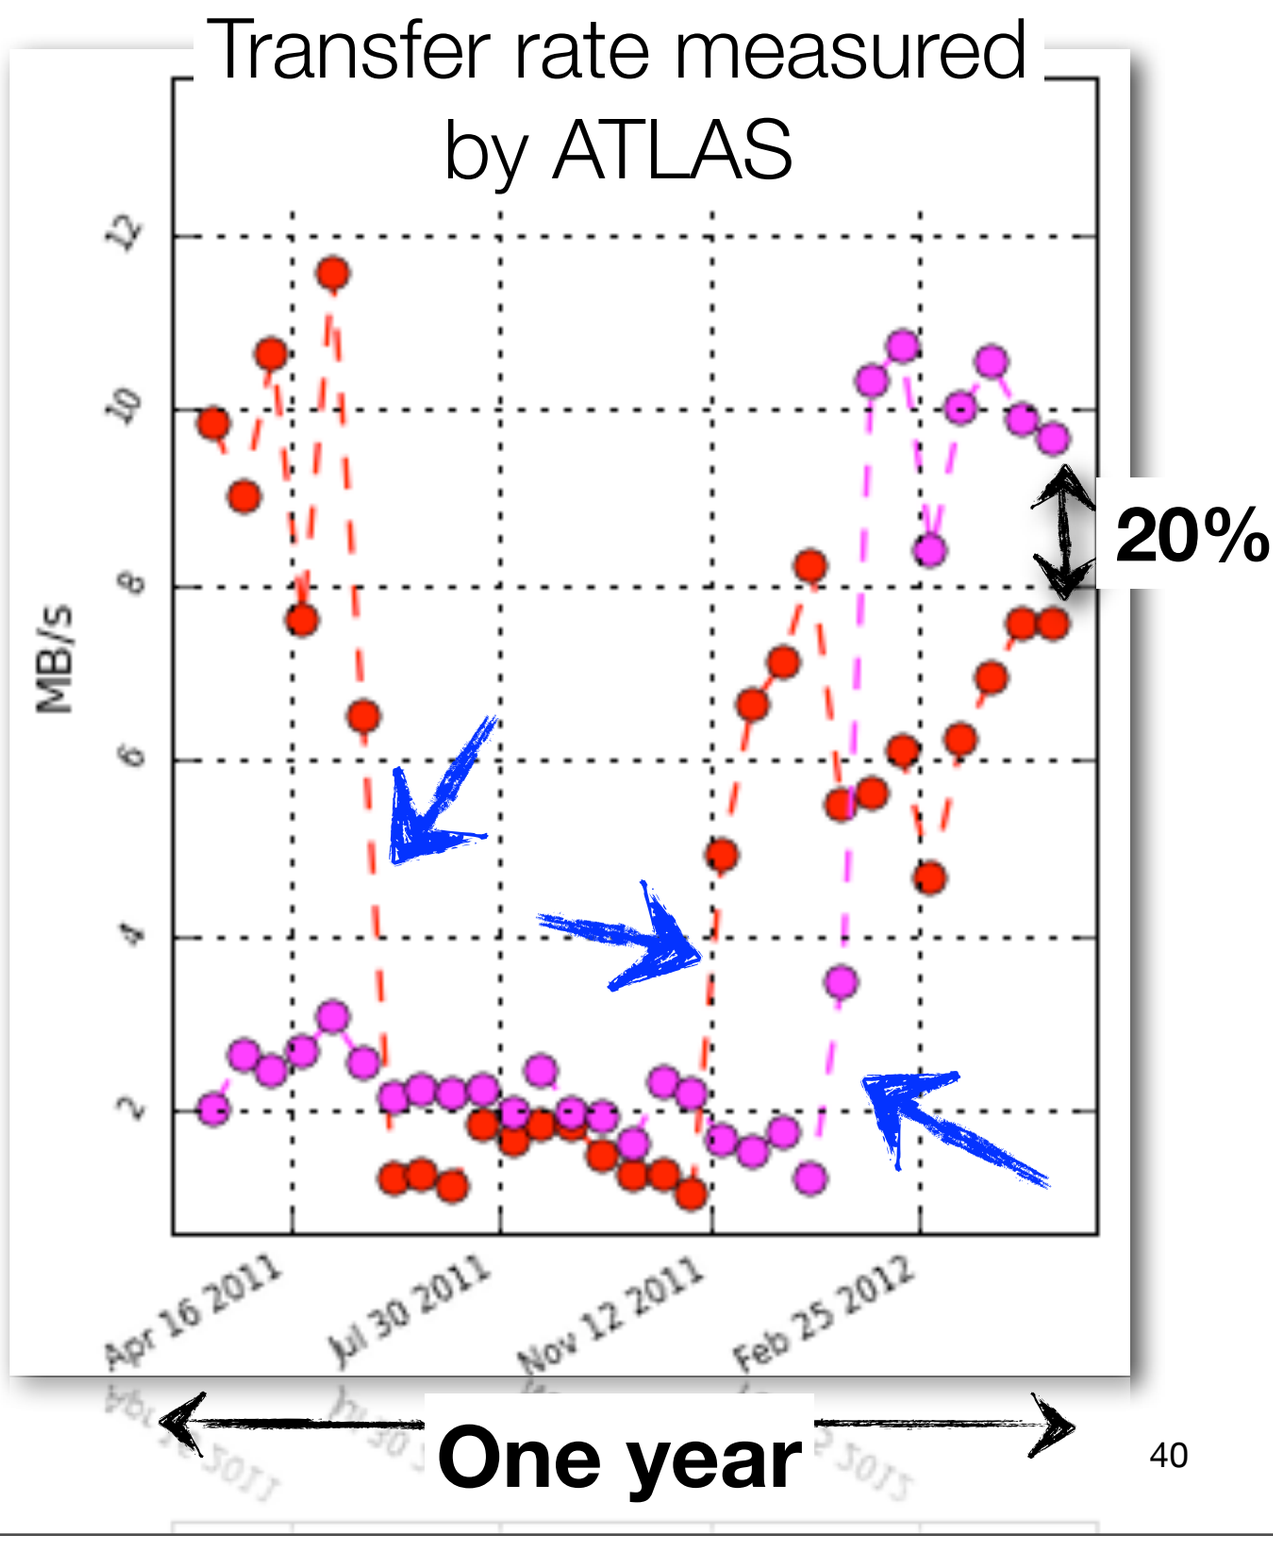

Network throughput measured with perfSONAR

CCIN2P3 → Tokyo **Tokyo → CCIN2P3**

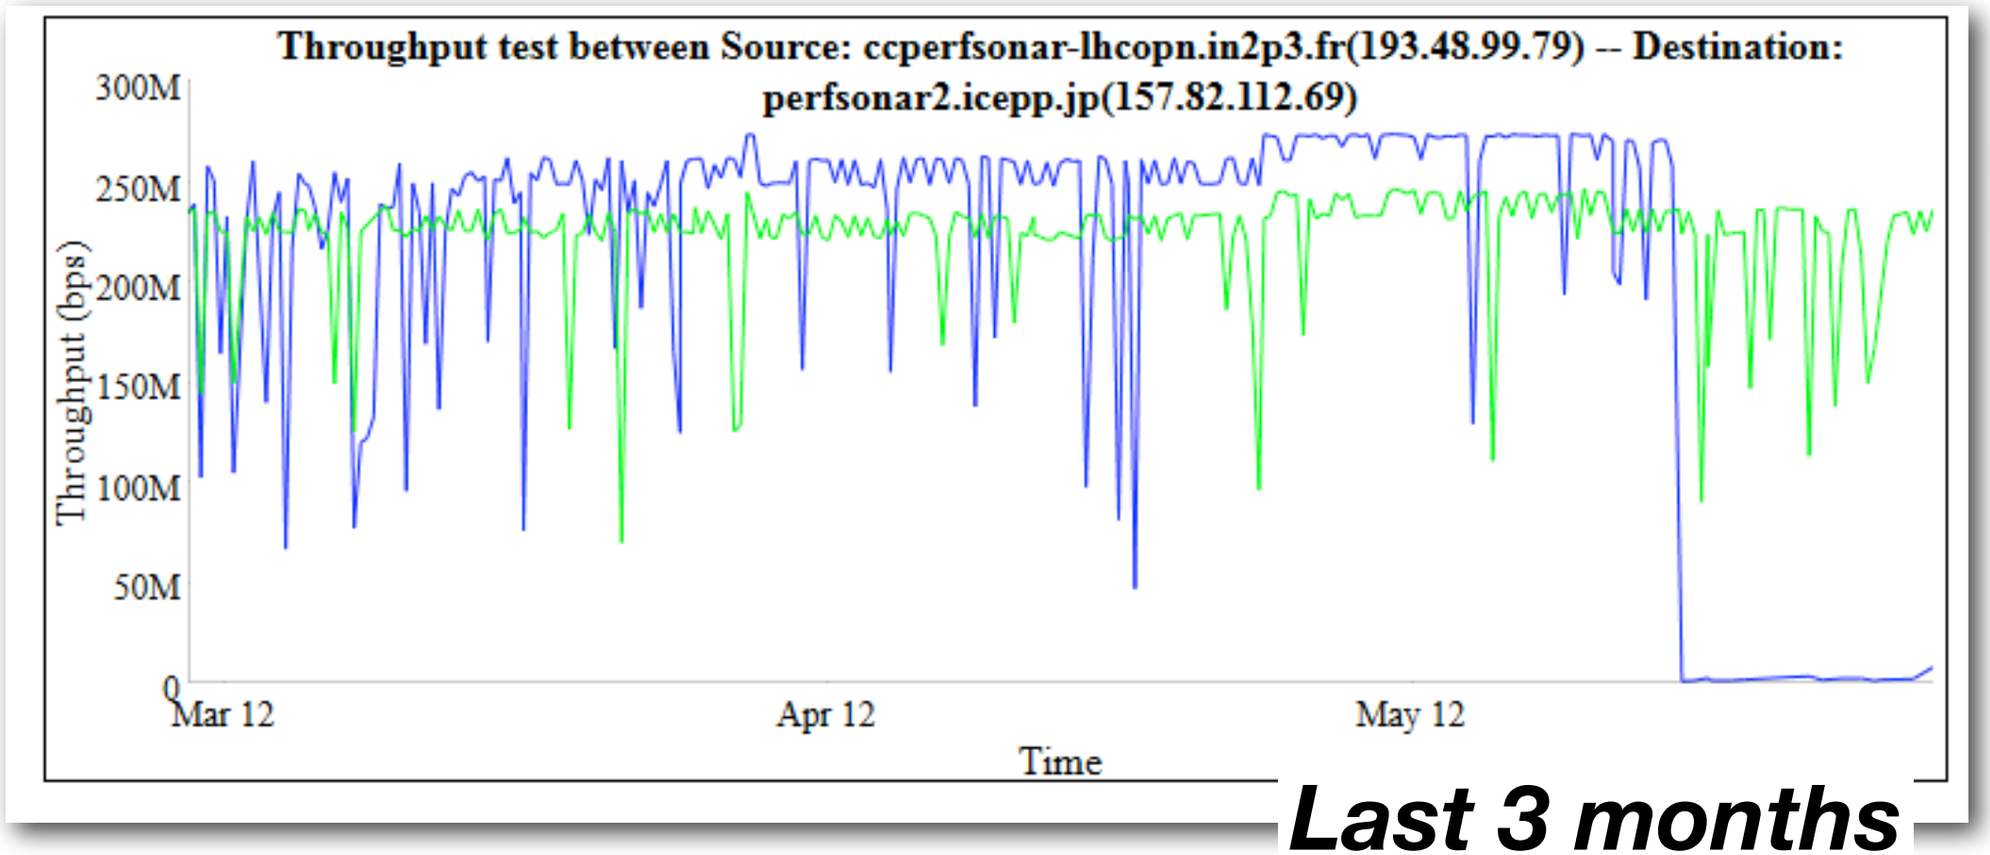

Last 3 months

No so stable

better by ~5% for CCIN2P3 → Tokyo

Different pattern for another French site

GRIF-LPNHE → Tokyo

Tokyo → GRIF-LPNHE

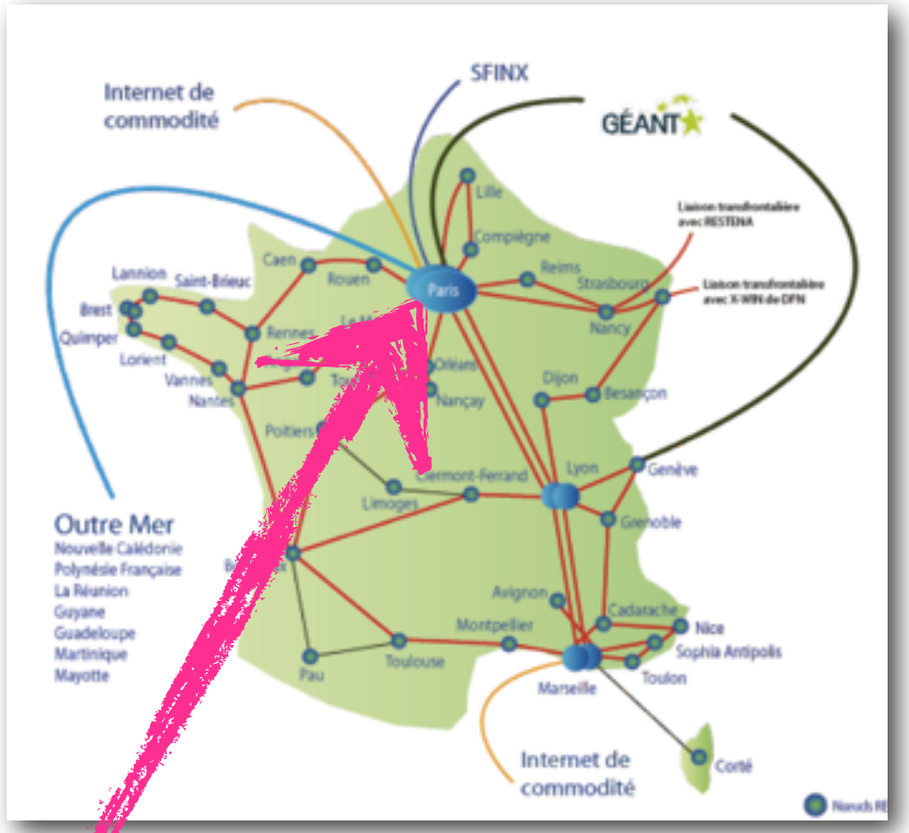

GRIF-LPNHE = Paris

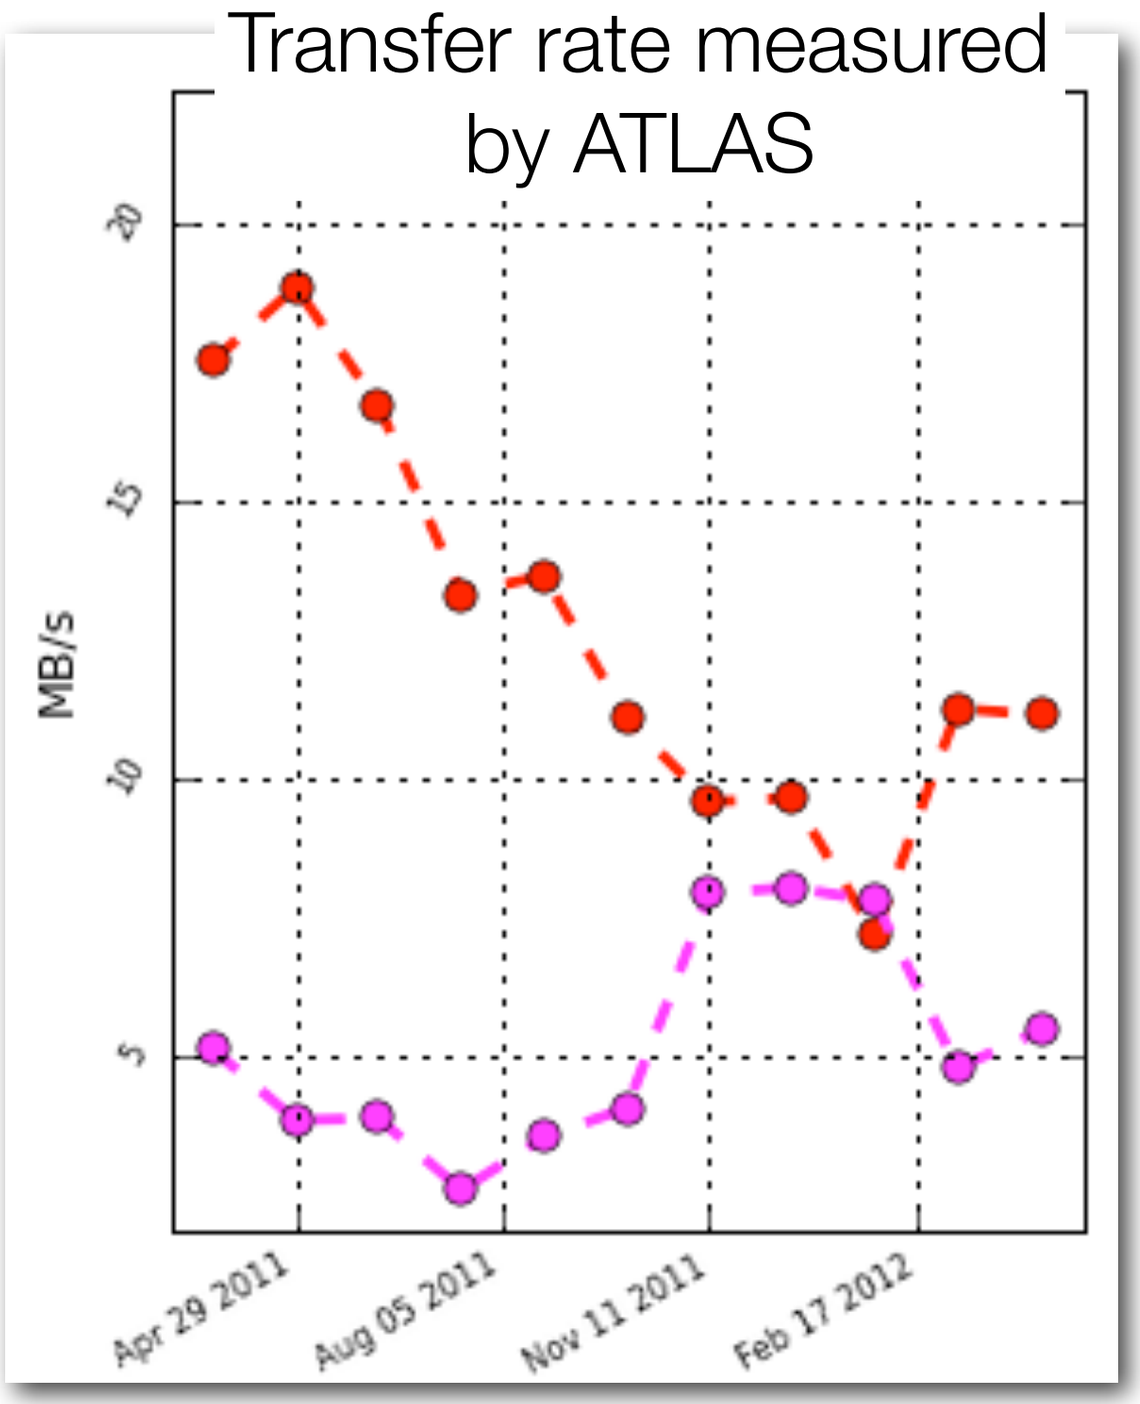

Efficiency of data transfer to Tokyo from FR-cloud sites

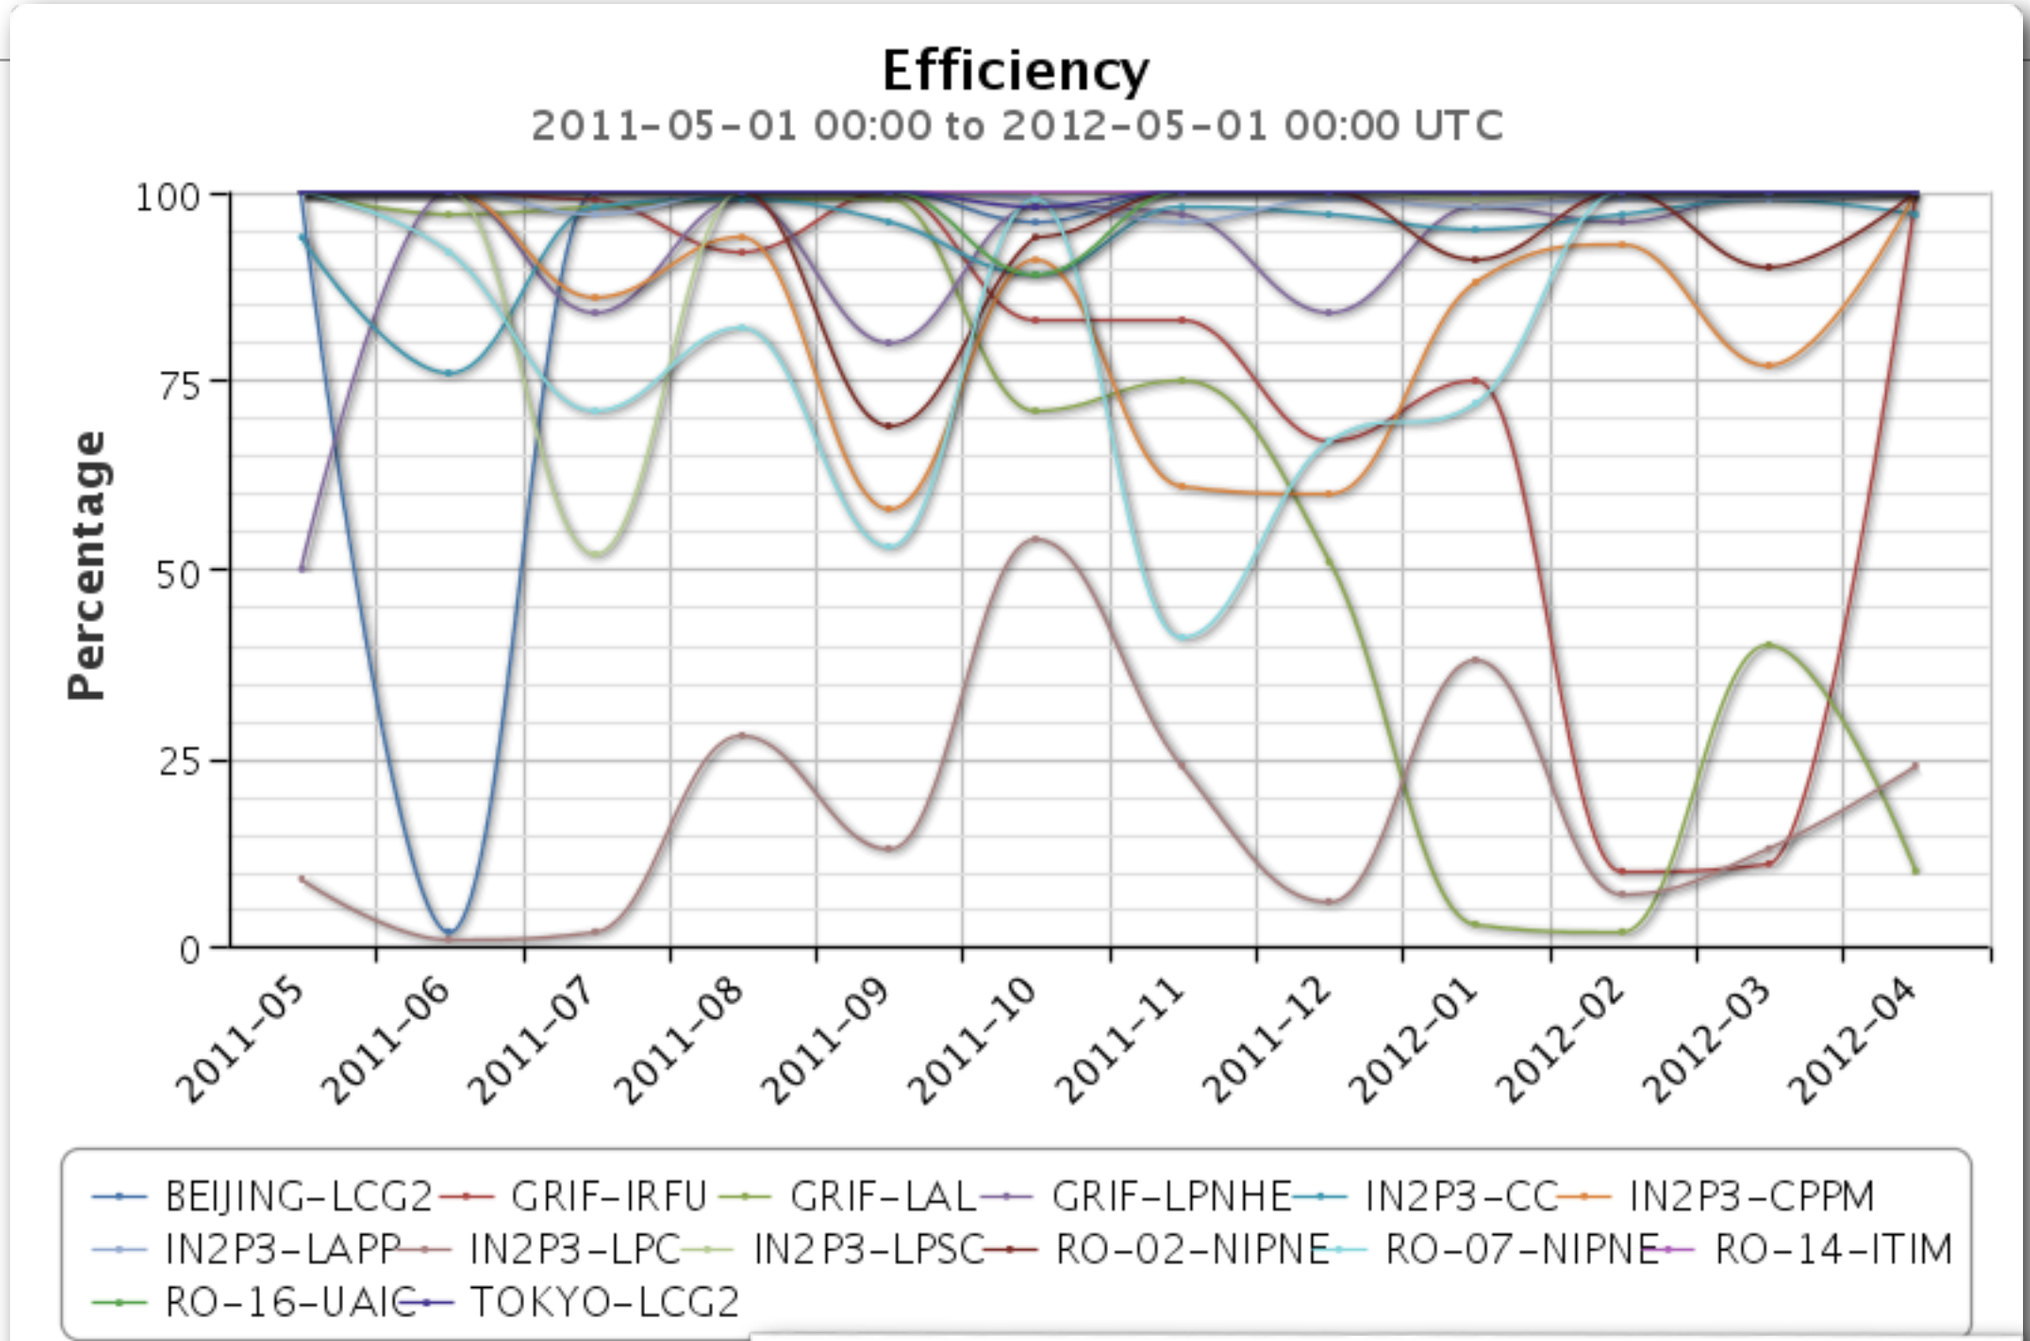

Each site is different, room for improvement everywhere...

FR-cloud view in ATLAS perfSonar dashboard

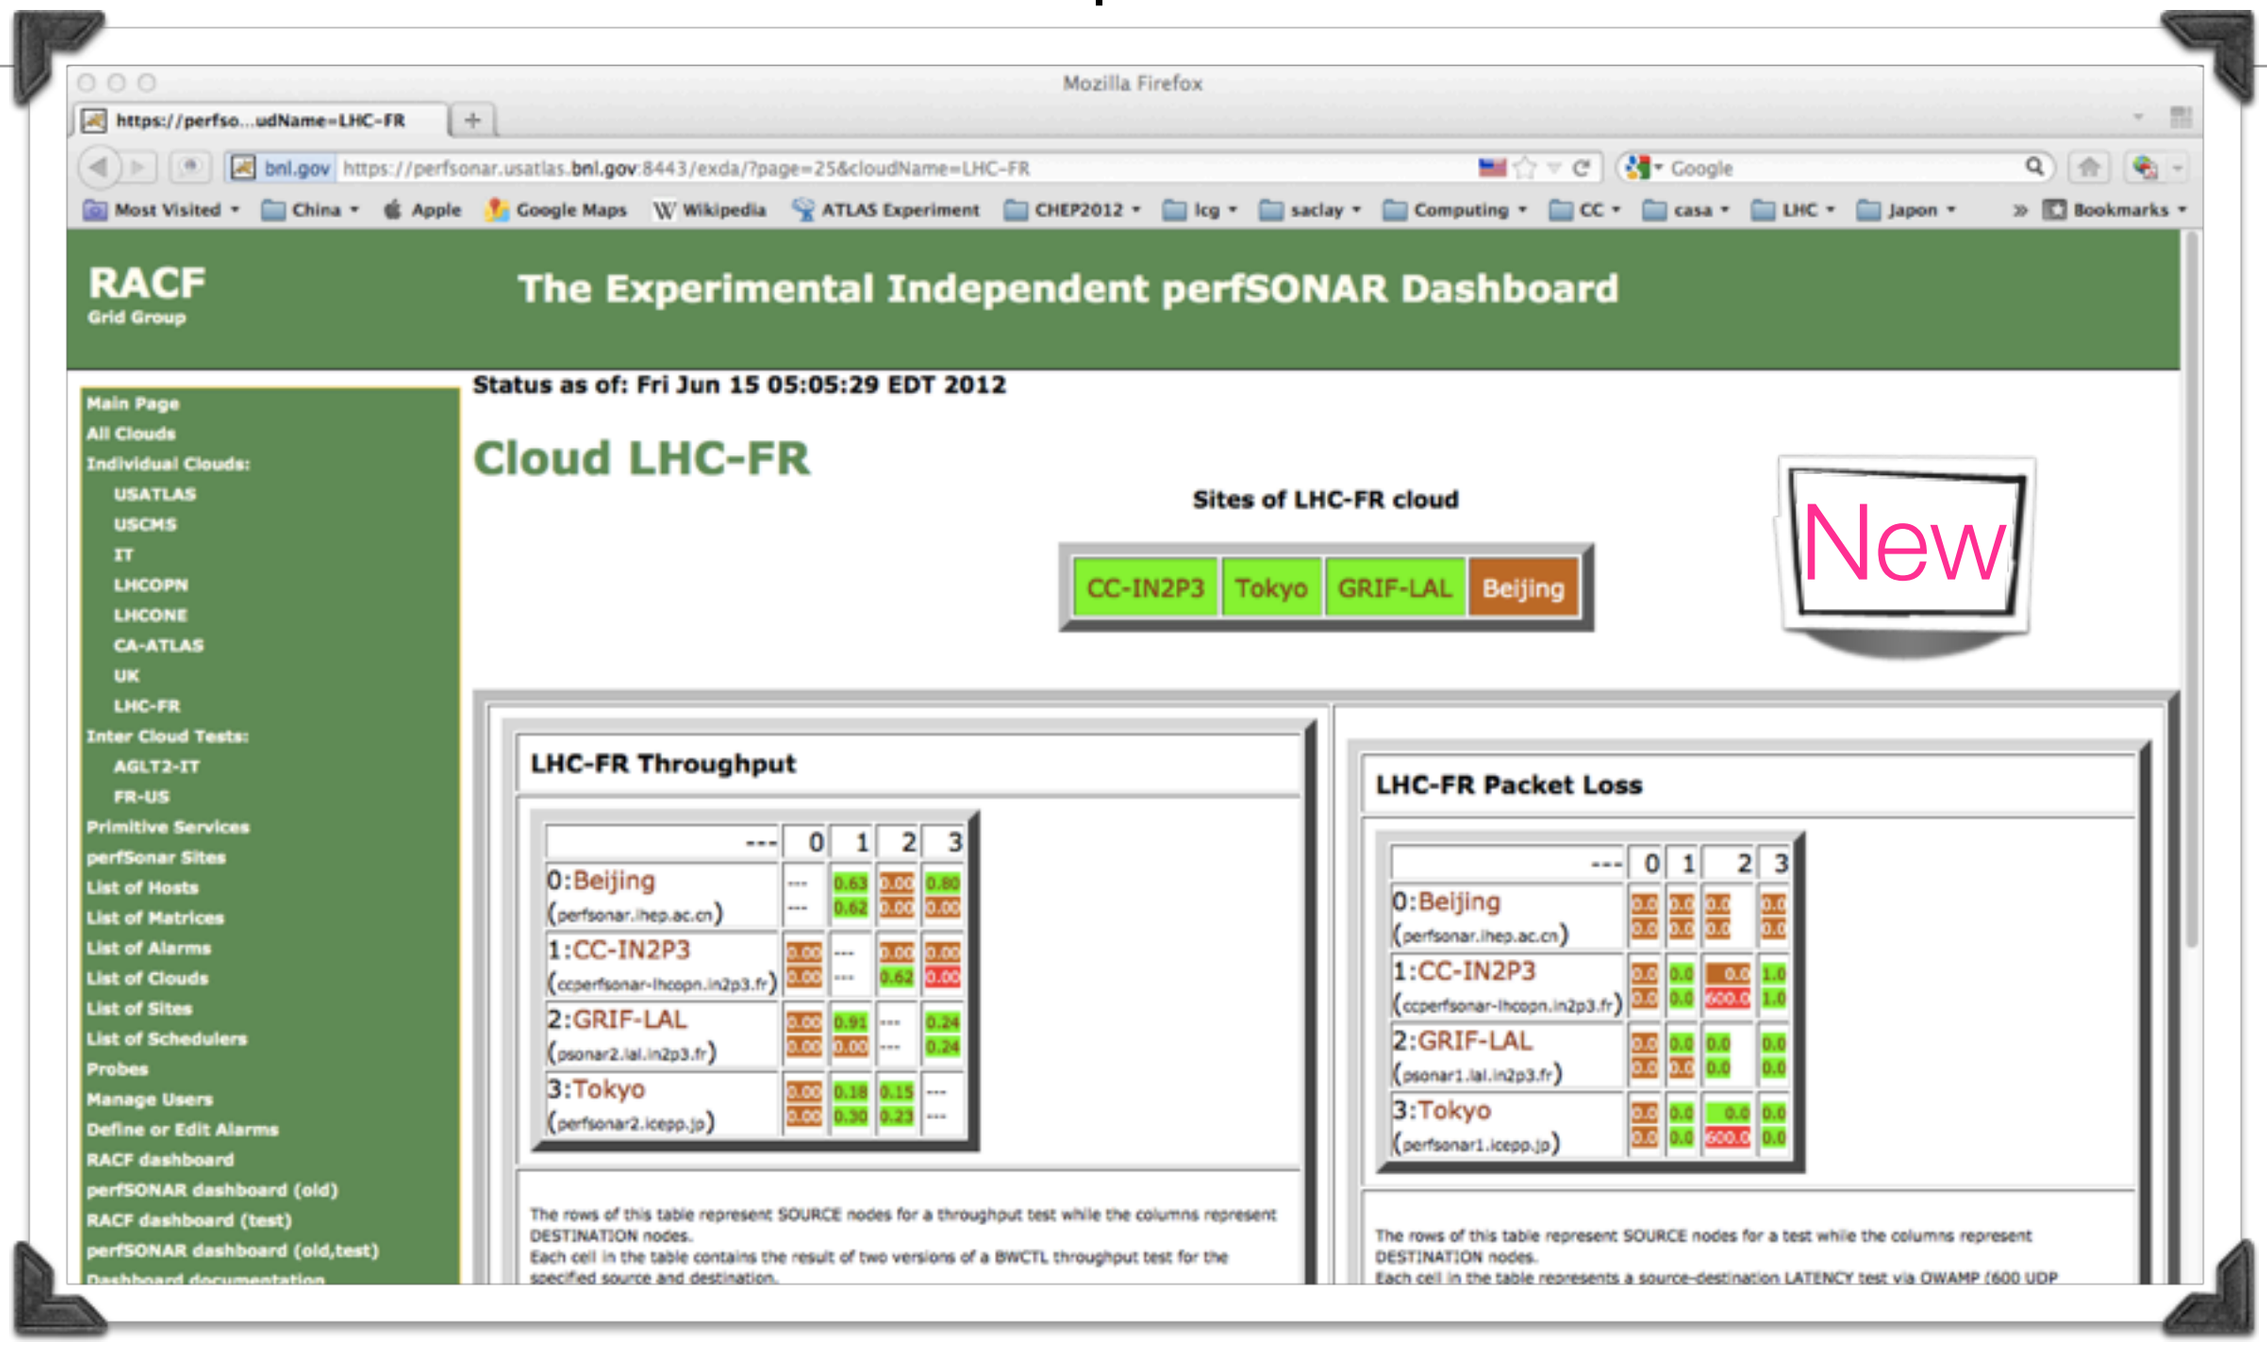

https://perfsonar.usatlas.bnl.gov:8443/exda/?page=25&cloudName=LHC-FR

ATLAS monitoring tools

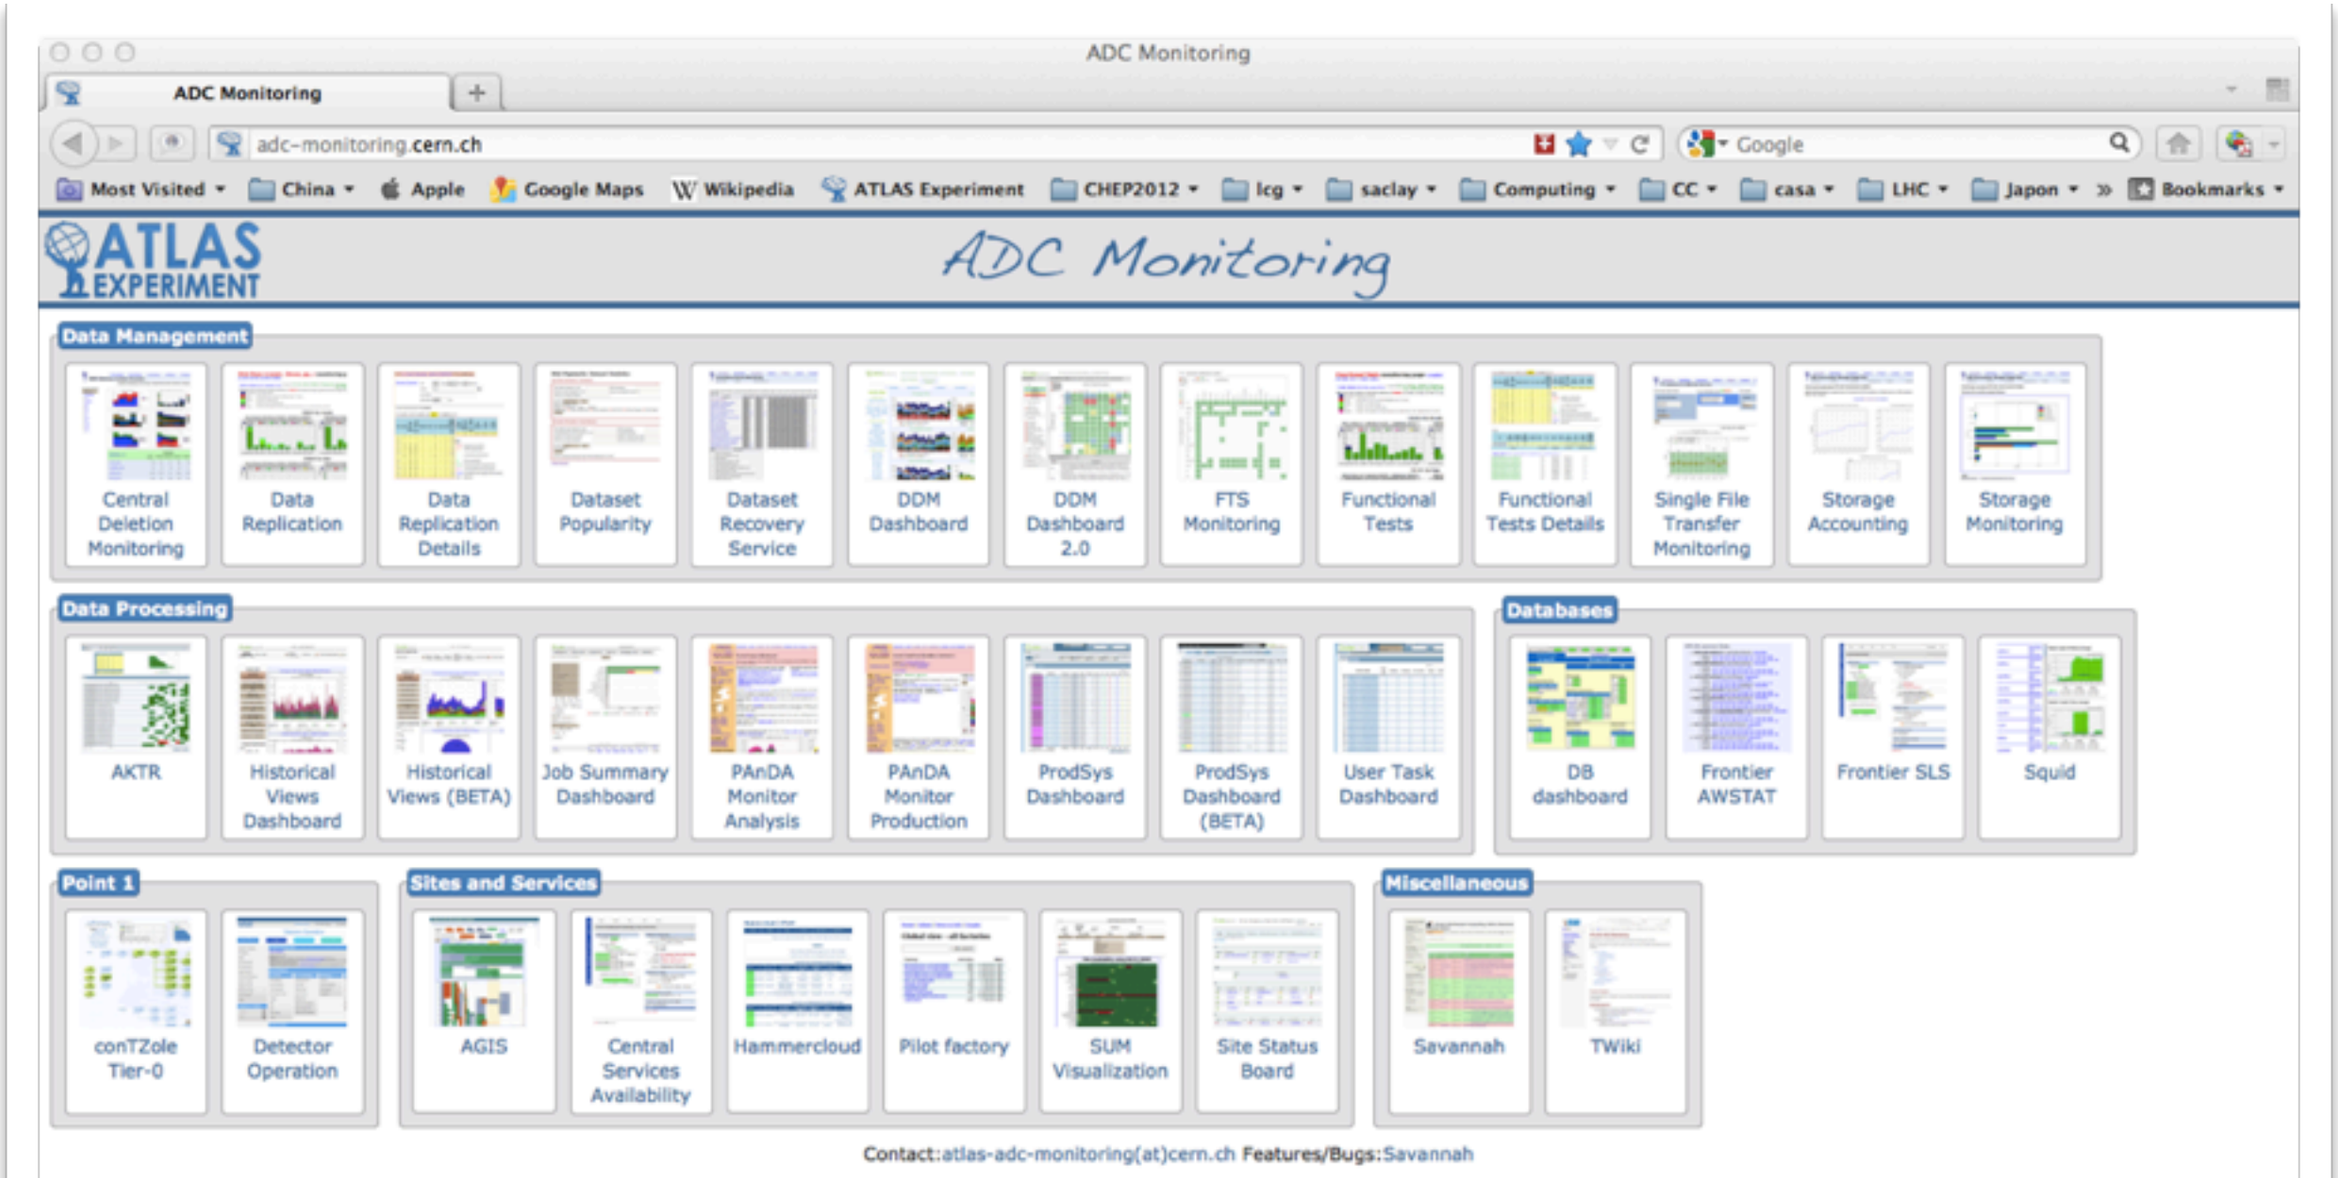

Lot of new tools over last year

Too many entry points?

FR-Cloud synthesis

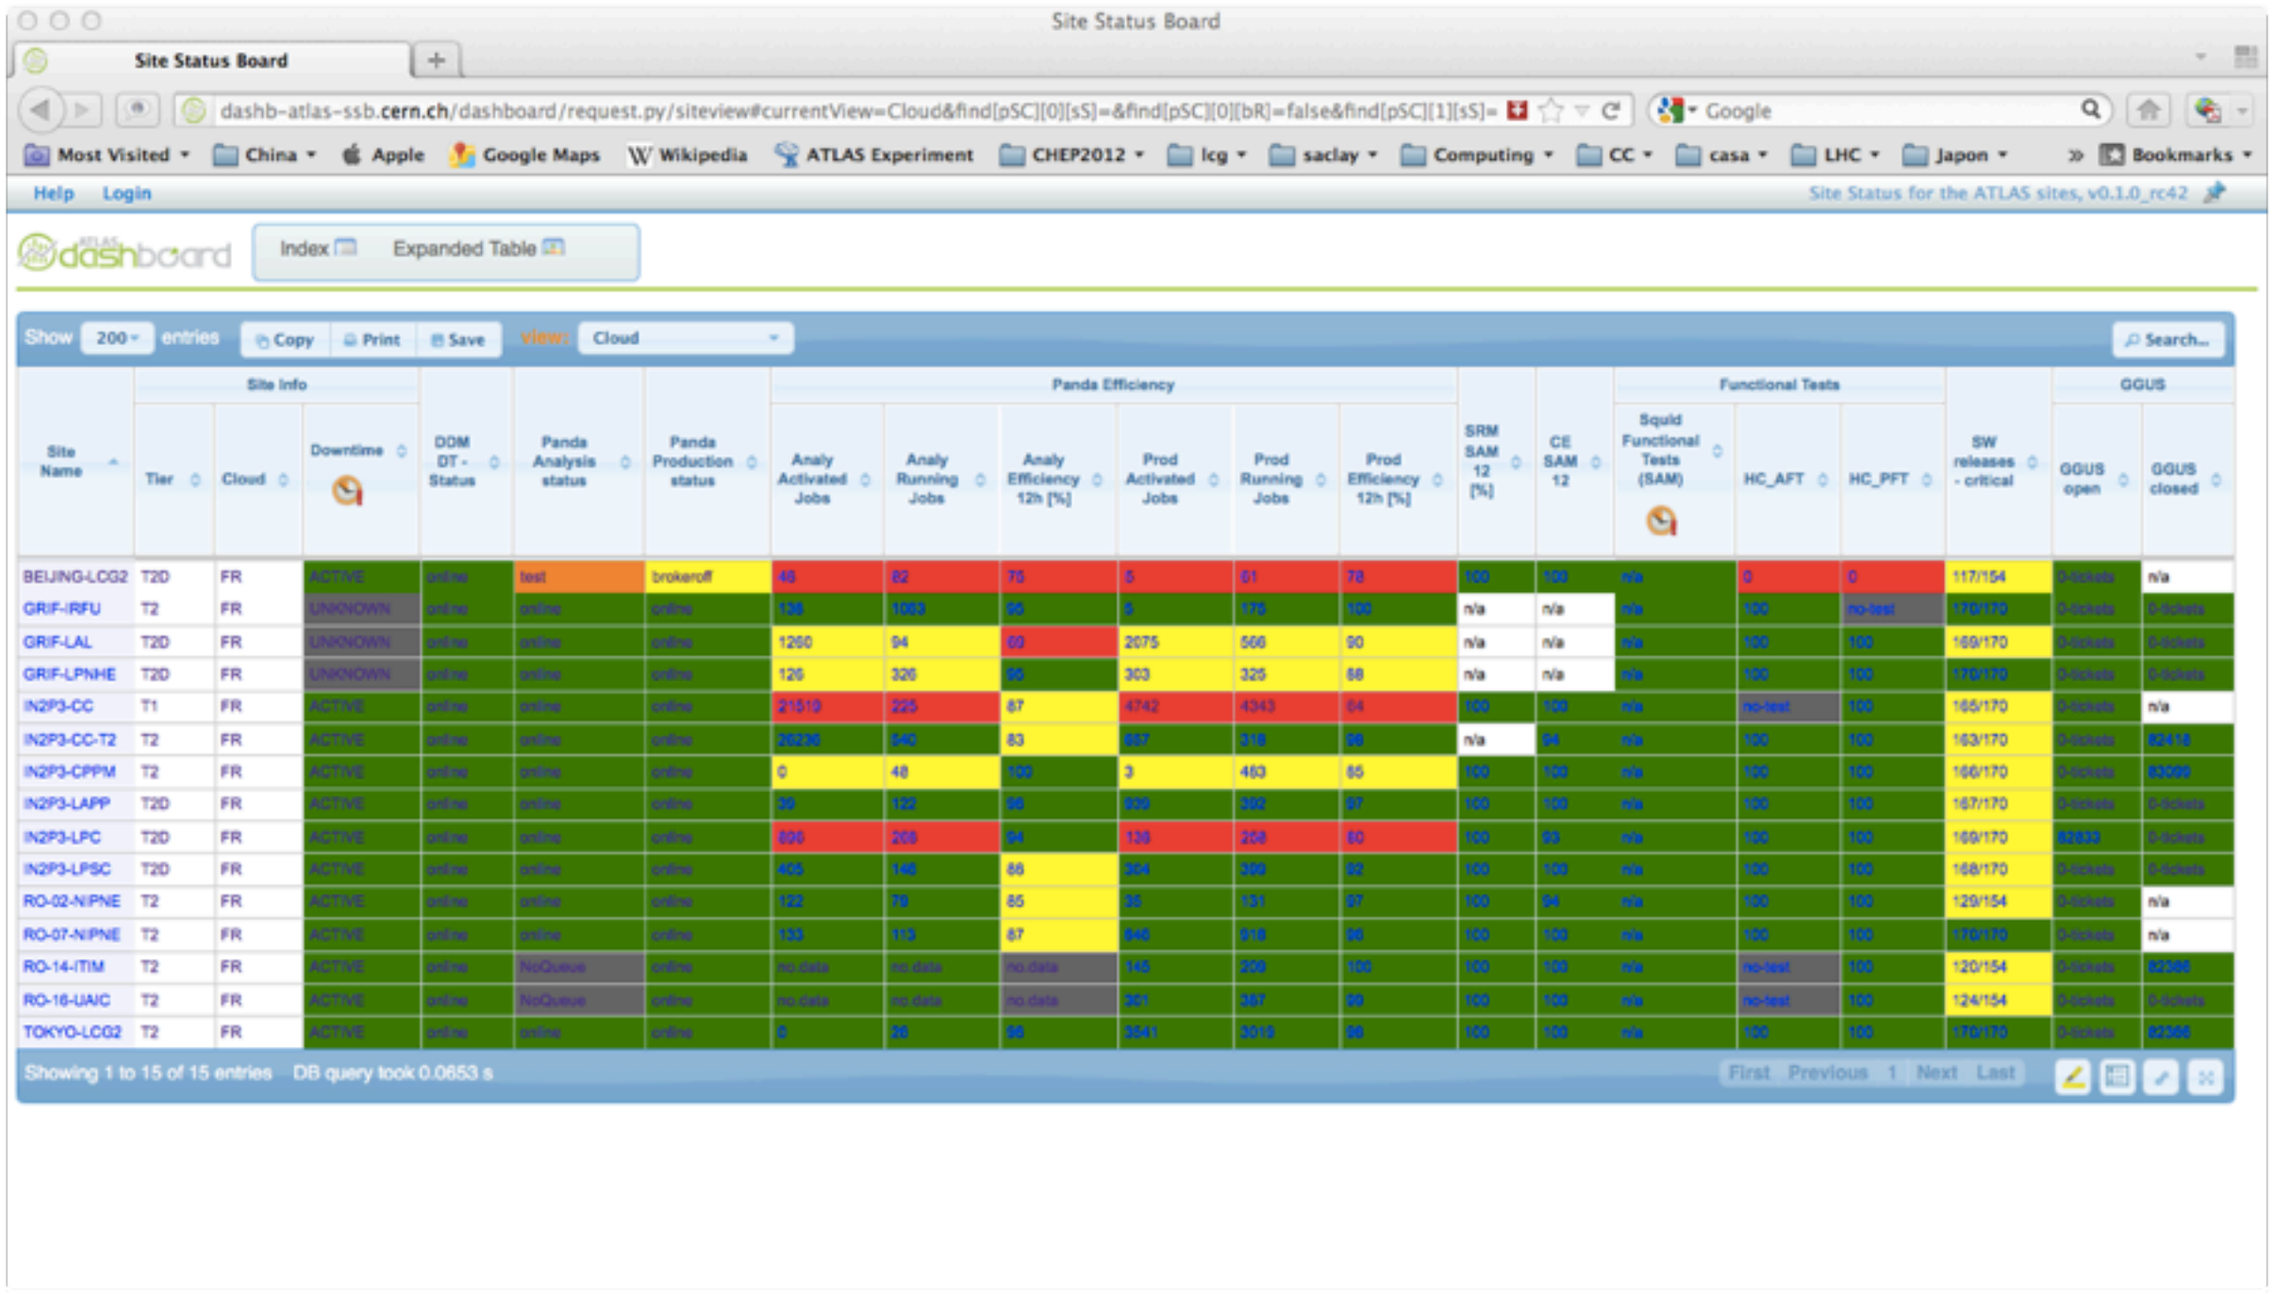

CVMFS

Deployed on every site of FR-cloud?

- CVMFS (CERN Virtual Machine File System)
- Used by ATLAS (https://twiki.cern.ch/twiki/bin/view/Atlas/CernVMFS) for
  - software distribution
  - conditions data distribution
- Advantages
  - no need to worry about increasing size of NFS software area and bottleneck of shared area
  - increase of CPU efficiency
- Ongoing deployment on ATLAS sites (http://dashb-atlas-ssb.cern.ch/dashboard/request.py/siteview#aaSorting[0][]=1&aaSorting[0][]=asc&currentView=cvmfs&highlight=false)

Xrootd federations

Use Case #2 Sharing storage amongst nearby Tier 2 sites

- In US - The 5 sites in ANALY_AGLT2 & ANALY_MWT2 are all within 7 ms RTT
- Midwest Tier 2 (MWT2) internal federation amongst 3 sub sites
- AGLT2 has storage federated between two sites
- ~ 4.4 PetaBytes of storage amongst sites
- Direct Read mode
- Most efficient for jobs that read part of data file (like Analysis jobs)

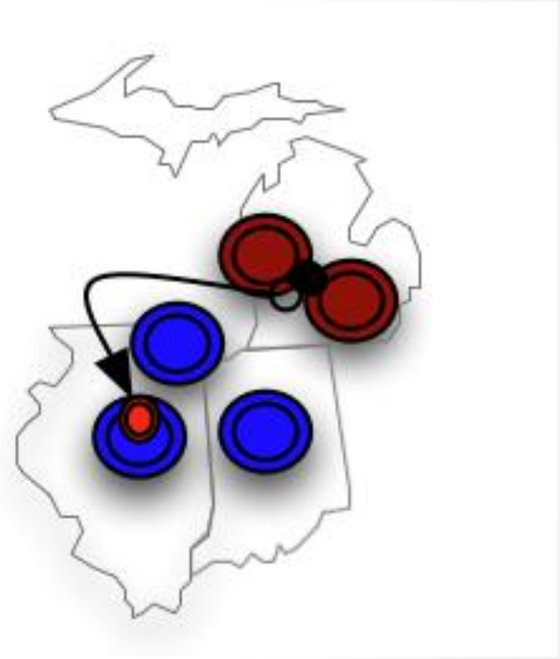

Applicable for some FR-cloud sites?

Multi-core processing

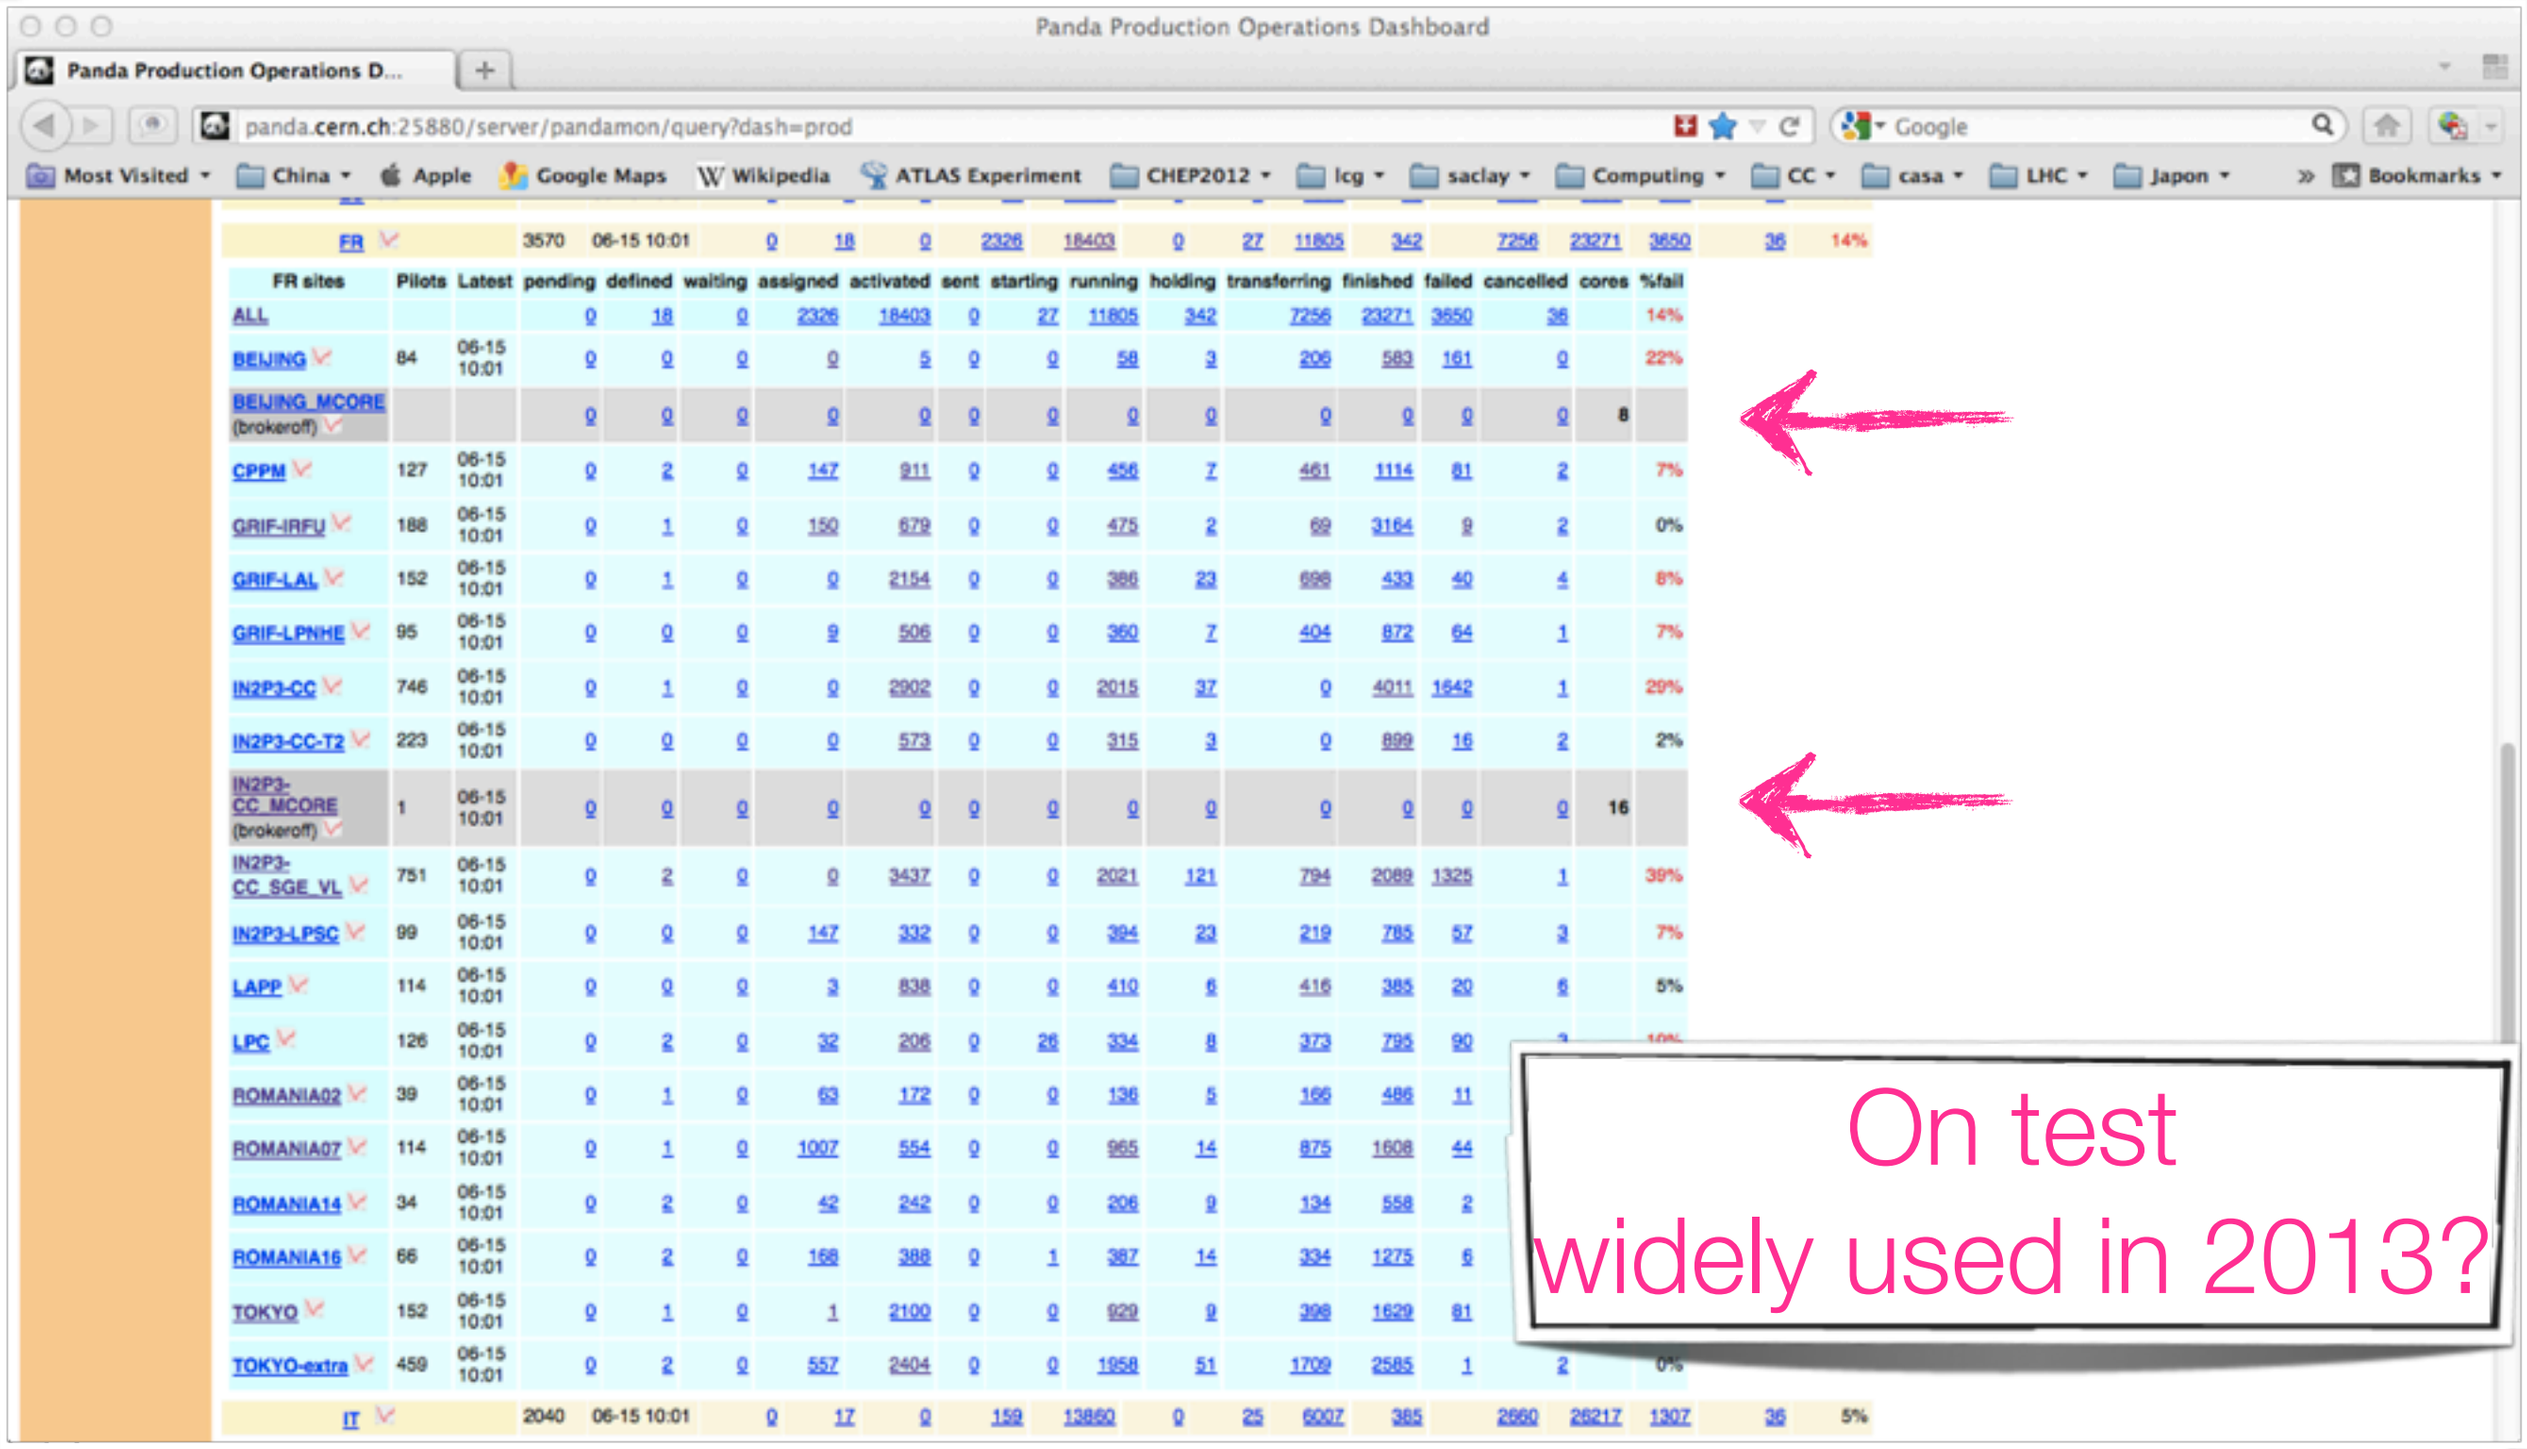

FR		3570	06-15 10:01	0	18	0	2326	18403	0	27	11805	342		7256	23271	3650	36	14%
FR sites	Pilots	Latest	pending	defined	waiting	assigned	activated	sent	starting	running	holding	transferring	finished	failed	cancelled	cores	%fail	
ALL			0	18	0	2326	18403	0	27	11805	342	7256	23271	3650	36		14%	
BEIJING	84	06-15 10:01	0	0	0	0	5	0	0	58	3	206	583	161	0		22%	
BEIJING_MCORE (brokeroff)			0	0	0	0	0	0	0	0	0	0	0	0	0	8		
CPPM	127	06-15 10:01	0	2	0	147	911	0	0	456	7	461	1114	81	2		7%	
GRIF-IRFU	188	06-15 10:01	0	1	0	150	679	0	0	475	2	69	3164	9	2		0%	
GRIF-LAL	152	06-15 10:01	0	1	0	0	2154	0	0	386	23	698	433	40	4		8%	
GRIF-LPNHE	95	06-15 10:01	0	0	0	9	506	0	0	360	7	404	872	64	1		7%	
IN2P3-CC	746	06-15 10:01	0	1	0	0	2902	0	0	2015	37	0	4011	1642	1		29%	
IN2P3-CC-T2	223	06-15 10:01	0	0	0	0	573	0	0	315	3	0	899	16	2		2%	
IN2P3-CC_MCORE (brokeroff)	1	06-15 10:01	0	0	0	0	0	0	0	0	0	0	0	0	0	16		
IN2P3-CC_SGE_VL	751	06-15 10:01	0	2	0	0	3437	0	0	2021	121	794	2089	1325	1		39%	
IN2P3-LPSC	99	06-15 10:01	0	0	0	147	332	0	0	394	23	219	785	57	3		7%	
LAPP	114	06-15 10:01	0	0	0	3	838	0	0	410	6	416	385	20	6		5%	
LPC	126	06-15 10:01	0	2	0	32	206	0	26	334	8	373	795	90				
ROMANIA02	39	06-15 10:01	0	1	0	63	172	0	0	136	5	166	486	11				
ROMANIA07	114	06-15 10:01	0	1	0	1007	554	0	0	965	14	875	1608	44				
ROMANIA14	34	06-15 10:01	0	2	0	42	242	0	0	206	9	134	558	2				
ROMANIA16	66	06-15 10:01	0	2	0	168	388	0	1	387	14	334	1275	6				
TOKYO	152	06-15 10:01	0	1	0	1	2100	0	0	929	9	398	1629	81				
TOKYO-extra	459	06-15 10:01	0	2	0	557	2404	0	0	1958	51	1709	2585	1	2		0%	
IT		2040	06-15 10:01	0	17	0	159	13860	0	25	6007	385		2660	26217	1307	36	5%

And now...

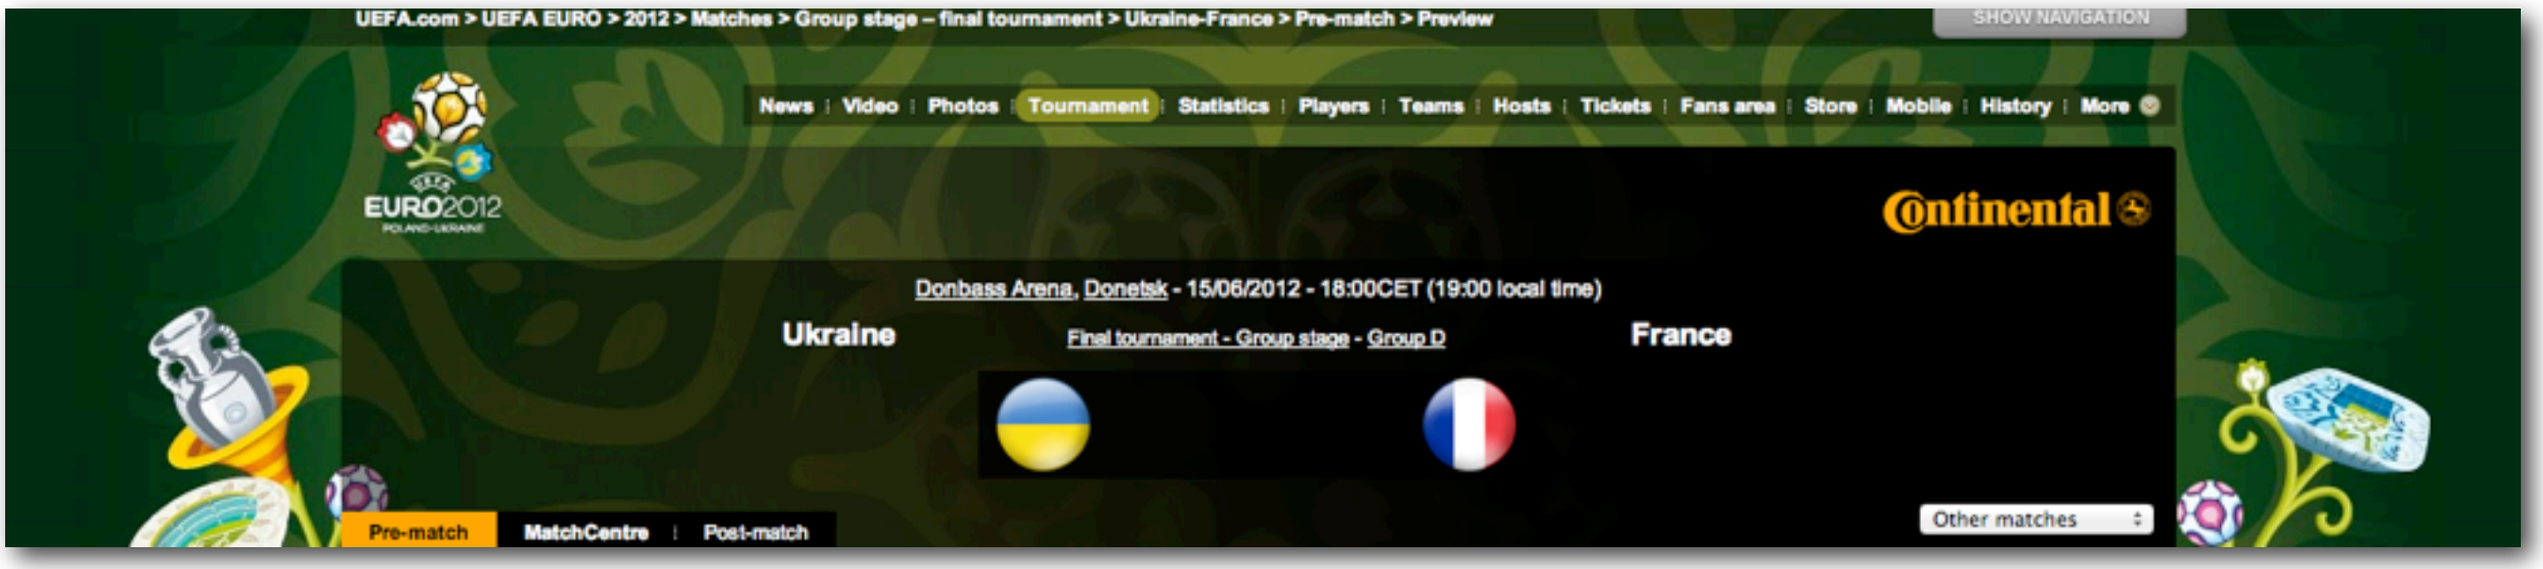
UEFA.com > UEFA EURO > 2012 > Matches > Group stage – final tournament > Ukraine-France > Pre-match > Preview
SHOW NAVIGATION
News | Video | Photos | Tournament | Statistics | Players | Teams | Hosts | Tickets | Fans area | Store | Mobile | History | More
EURO2012
POLAND-UKRAINE
Continental
Donbass Arena, Donetsk - 15/06/2012 - 18:00CET (19:00 local time)
Ukraine
Final tournament - Group stage - Group D
France
Pre-match
MatchCentre
Post-match
Other matches

ATLAS Computing Model: T0 Data Flow

200Hz
RAW: ~1.6MB/evt

Calibration

CERN Analysis Facility

Event Summary Data (ESD): ~1 MB/evt
Analysis Object Data (AOD): ~100 kB/evt
derived data (dESD, dAOD, NTUP,...)
distributed over the Grid

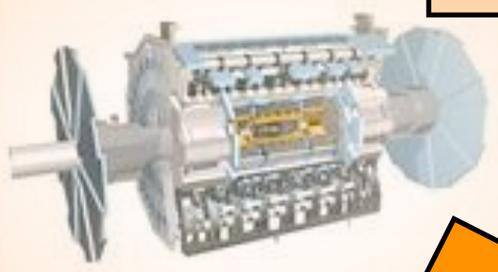

Tier-0

Data Recording to tape
First Pass Processing

Tier-1

Tier-1

Tier-2

Tier-1

Tier-2

10 Tier-1 centers
RAW data copy on tape
Analysis data on disk
Reprocessing

38 Tier-2 centers
Analysis data on disk
User Analysis

Static environment, fear of the **networks**!

- Tiered site/cloud architecture
- Pre-planned data distribution
- Jobs-to-data brokerage

ATLAS Computing Model: MC Data Flow

CERN Analysis Facility

10 Tier-1 centers + CERN
Aggregation of produced data
Source for further distribution

Tier-1

Tier-2

Tier-2

Tier-2

job input

Tier-1

Tier-2

Tier-2

Tier-2

Tier-1

Tier-2

Tier-2

Tier-2

38 Tier-2 centers
MC Simulation
Group/User Analysis