
# Site operations

Costin.Grigoras@cern.ch

# Outline

- Central services
- VoBox services
- Monitoring
- Storage and networking

# Central Services

# Core central services

- Central MySQL databases + hot backups
  - Catalogue: 1B files
  - Task Queue: 250K jobs/day (avg) (up to 2x in analysis periods)
  - Scheduled transfers
- AliEn services
  - Authen: *alice-authen.cern.ch*:8080
  - PackMan: *alice-packman.cern.ch*:9991
  - Job Broker: *alice-jobbroker.cern.ch*:8050
  - Job Manager: *aliendb8.cern.ch*:8083
  - Job Info Manager: *alice-jobinfomanager.cern.ch*:8081
  - Information Service: *alice-is.cern.ch*:8099
  - API: *alice-apiserv1.cern.ch*:10000, *alice-apiserv2.cern.ch*:10000
  - LDAP: *alice-ldap.cern.ch*:8389
  - Various optimizers and internal services
- Transfer agents
  - Third party copying or store and forward of the data

# Monitoring services

- MonALISA central repository: *alimonitor.cern.ch*:80,443
  - 2 independent PostgreSQL backends
- MonALISA proxy service: *alimlproxy.cern.ch*:6001

# Build servers

- AliEn and AliROOT build systems for
  - SLC5, 32b and 64b
  - SLC6, 32b and 64b
  - Ubuntu, 64b
  - Mac OSx
- Daily analysis tags + 2 revisions weekly
  - Automatically deployed on CVMFS
  - Also available for users to install
    - *wget* directly from build servers via *alienbuild.cern.ch*:80,8880,8888,8889
    - *alitorrent.cern.ch*:80,8088,8092

# Various other services

- Automatic revision testing of QA and refiltering code
- LEGO train wagon testing machinery
- AliRoot code checkers
- Shifter and detector construction databases
- ALICE public web site
- Backup service and software repository

# VoBox services

- CE
  - Submitting generic Job Agents to the local BQ when something in the central task queue matches the site resources
- Cluster Monitor: TCP/8084
  - Message proxy between job agents and the central services
- CMReport
  - Periodic message buffer flushes to the central services
- MonALISA:
  - Collects and aggregates all site-produced monitoring data
  - Periodic tests of VoBox services health
  - ApMon listener: UDP/8884
  - Xrootd monitoring: UDP/9930
  - Bandwidth tests: TCP/1093, ICMP, UDP/33434

# Job Agent monitoring

- Instrumented with ApMon
- Full host monitoring parameters
  - CPU, load, network traffic, number of processes and sockets in each state, disk and swap IO, CPU type and spec power, OS
- Self monitoring
  - Proxy time left, CPU and memory utilization, status, current job ID, number of jobs picked up so far
- Current job monitoring
  - CPU, memory and disk utilization, number of open files, job meta information (queue ID, master job ID, owner name)

# Job monitoring

- Root is compiled with ApMon support as well, so jobs can use TMonaLisaWriter
  - Used eg. for grid-wide CPU benchmarking using Root stress benchmark
- *xrdcp* command reports transfer details to the VoBox
  - Source and destination, amount of data, time it took etc

# Storage monitoring

- Xrootd and EOS data servers publish two monitoring streams
  - ApMon daemon reporting the data server host monitoring and external Xrootd params
    - Node total traffic, load, IO
    - Version, total and used space
  - Xrootd internal reporting on file close
    - xrootd.monitor all flush 60s window 30s dest files info user MONALISA_HOST:9930
    - Client IP, read and written bytes, speed

# Site monitoring data aggregation

- Monitoring data is aggregated in real time by the VoBox ML service
- Summaries are publishes along side the individual values
  - Total traffic on the Xrootd servers
    - And split by remote site, LAN/WAN
  - Aggregated resource consumption by jobs
    - By queue, by user name
  - Count jobs in each state
  - Various aggregation functions available
    - min/max/avg/sum
    - Top jobs in terms of allocated memory

# Central monitoring repository

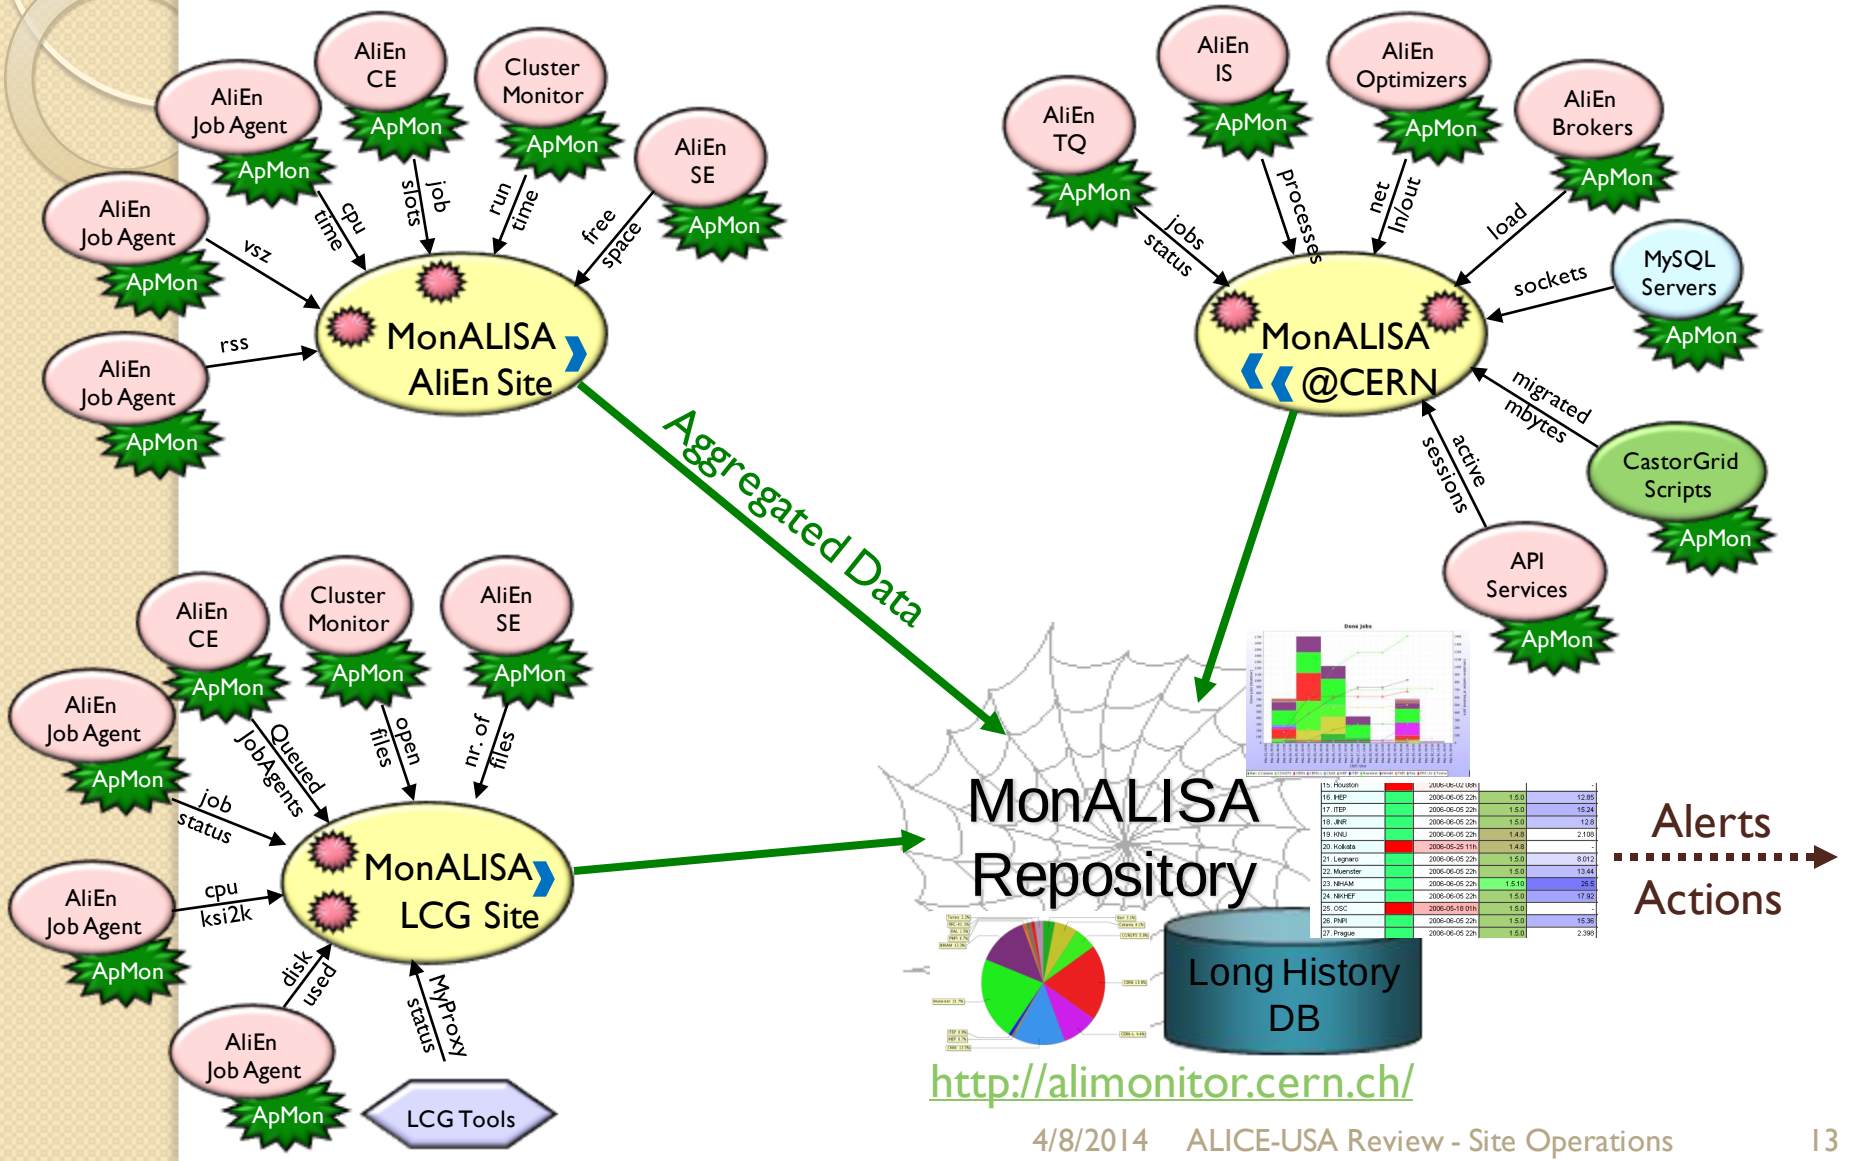

# What you can see centrally

- Current status
  - Of all services, central and site local
  - Of all jobs and ongoing productions, analysis or user activity
  - Catalogue browser
  - Various test results: storage, network
- Aggregated history data
  - Job accounting: running time, efficiency, consumed spec power
    - Per site and per user
  - Storage status
  - Network utilization
- Overview of current issues

# Network monitoring

- Periodic one TCP stream throughput test between all VoBoxes in ALICE
  - Similar to what the jobs would experience
- Pairs of VoBox machines selected by the repository
- Very important for debugging network connectivity for new sites or after major changes
- Also records *traceroute/tracepath* result along with the test for later comparison
- And VoBox kernel network parameters
- See the earlier firewall requirements

# Topology map (AS level)

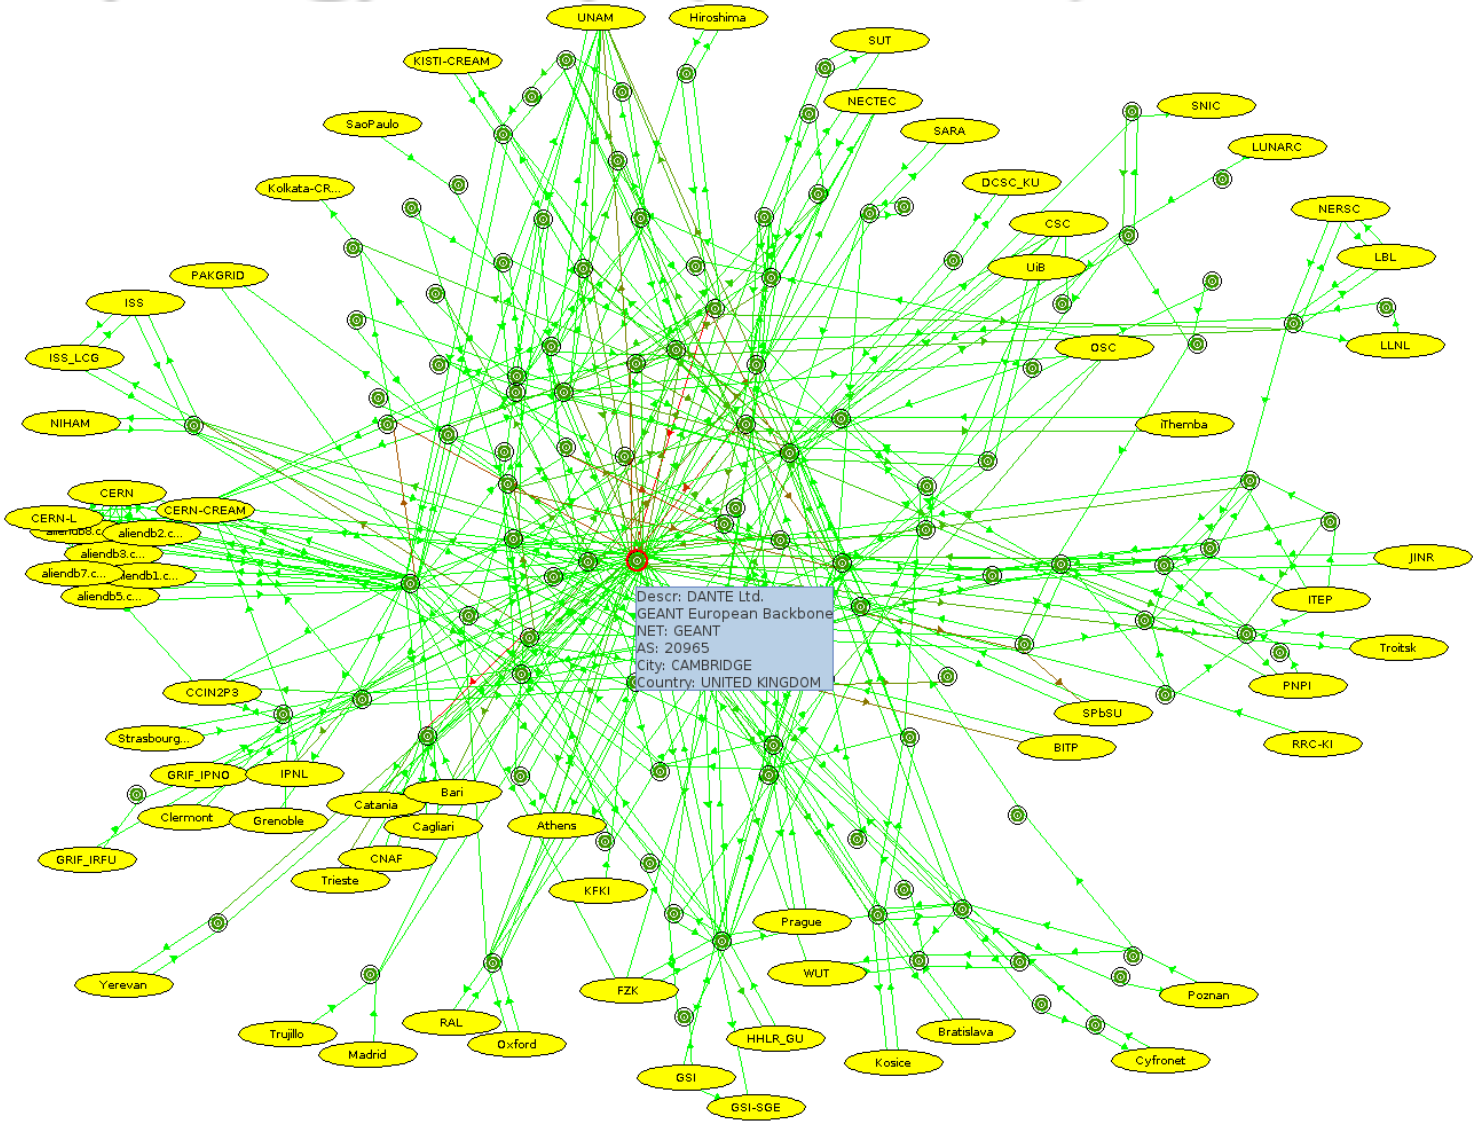

# Storage monitoring

- Every 2h a full *add*/*get*/*rm* test suite from the repository machine
  - Storage functional status
  - Remote access to it
- If the storage is full only a *get* operation is performed, but it is still marked as bad for writing
- For xrootd: individual server testing with a similar test suite
- Alarms raised on reported size different from LDAP declared size
  - Sometimes data servers are not seen by the redirector any more – restart usually cures it

# Storage discovery

- Closest working replicas are used for both reading and writing
  - Sorting the SEs by the network distance to the client making the request
    - Combining network topology data with the geographical location
  - Leaving as last resort only the SEs that fail the respective functional test
  - Weighted with their recent reliability and remaining free space
- Writing is finally slightly randomized for more 'democratic' data distribution

# Distance metric function

- *distance*(IP, IP)
  - Same C-class network
  - Common domain name
  - Same AS
  - Same country (+ function of RTT between the respective AS-es if known)
  - If distance between the AS-es is known, use it
  - Same continent
  - Far, far away

0 ↓ 1

- *distance*(IP, Set<IP>): Client's public IP to all known IPs for the storage

# Weight factors

- Free space contributes with
  - f (ln(free space / 5TB))
- Recent history contributes with
  - 75% * last day success ratio +
  - 25% * last week success ratio
- *add* test result used for write discovery, *get* test result used for reading
- Resulting value added to the distance

# Impact on analysis jobs

- Local SE problems makes the jobs read remotely

| Site | Job eff. | HepSpec06 | All files | Local files | Remote files |
|---|---|---|---|---|---|
| **CERN**<br>9771 jobs (51.29%) | 69.76% | 10.25 | 58109 files<br>2.667 MB/s | 58100 (99.98%)<br>2.668 MB/s | 9 (0.015%)<br>0.889 MB/s |
| **FZK**<br>2966 jobs (15.57%) | 13.16% | 9.837 | 17560 files<br>0.471 MB/s | 1523 (8.673%)<br>2.963 MB/s | 16037 (91.33%)<br>0.44 MB/s |

- In this particular case the SE tests are all fine
  - Under investigation why the jobs cannot access local data
- Remote access can severely impact the jobs efficiency

| | | | | | |
|---|---|---|---|---|---|
| **LEGNARO**<br>268 jobs (1.407%) | 42.09% | 10.08 | 1496 files<br>1.285 MB/s | 636 (42.51%)<br>2.628 MB/s | 860 (57.49%)<br>0.95 MB/s |

# Remote access efficiency

| Storage / WNs | CERN | LEGNARO | TORINO | CNAF | FZK |
|---|---|---|---|---|---|
| CERN | 2.668 MB/s | | 0.27 MB/s | | |
| FZK | 0.486 MB/s | 0.161 MB/s | 0.213 MB/s | | 2.963 MB/s |
| LEGNARO | 1.611 MB/s | 2.628 MB/s | 0.673 MB/s | | 0.749 MB/s |
| TORINO | 1.848 MB/s | 1.609 MB/s | 0.684 MB/s | | 0.891 MB/s |
| CNAF | 2.193 MB/s | | 0.623 MB/s | 2.126 MB/s | |

- Problems can come from both network and the storage
- IO performance seen by jobs doesn't always match the VoBox-to-VoBox throughput measurements
  - Congested firewall / network segment, different OS settings, saturated storage IO
- Reflected in the overall efficiency

# Focus on US

- http://alimonitor.cern.ch?1163

IO stats for **2** master jobs executed between **02.04.2014 11:40** and **03.04.2014 01:35**

| Site activity | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| **Site** | **Job eff.** | **All files** | **Local files** | **Remote files** | **CERN EOS** | **LBL SE** | **NIHAM FILE** | **CCIN2P3 SE** | **FZK SE** | **PRAGUE SE** |
| **CERN**<br>137 jobs (61.99%) | **44.44%** | 7739 files<br>2.735 MB/s | 7702 (99.52%)<br>2.736 MB/s | 37 (0.478%)<br>2.453 MB/s | **7702 (99.52%)**<br>**2.736 MB/s** | | | 21 (0.271%)<br>2.616 MB/s | 8 (0.103%)<br>4.006 MB/s | 1 (0.013%)<br>1.513 MB/s |
| **LBL**<br>31 jobs (14.03%) | **25.44%** | 1979 files<br>4.531 MB/s | 1972 (99.65%)<br>5.438 MB/s | 7 (0.354%)<br>102.1 KB/s | | **1972 (99.65%)**<br>**5.438 MB/s** | | | | 5 (0.253%)<br>0.13 MB/s |
| **NIHAM**<br>27 jobs (12.22%) | **72.77%** | 1864 files<br>6.151 MB/s | 1864 (100%)<br>6.151 MB/s | | | | **1864 (100%)**<br>**6.151 MB/s** | | | |
| **LLNL**<br>26 jobs (11.76%) | **24.38%** | 1848 files<br>4.759 MB/s | | 1848 (100%)<br>4.759 MB/s | | 1840 (99.57%)<br>4.959 MB/s | | | 4 (0.216%)<br>0.313 MB/s | 1 (0.054%)<br>0.602 MB/s |

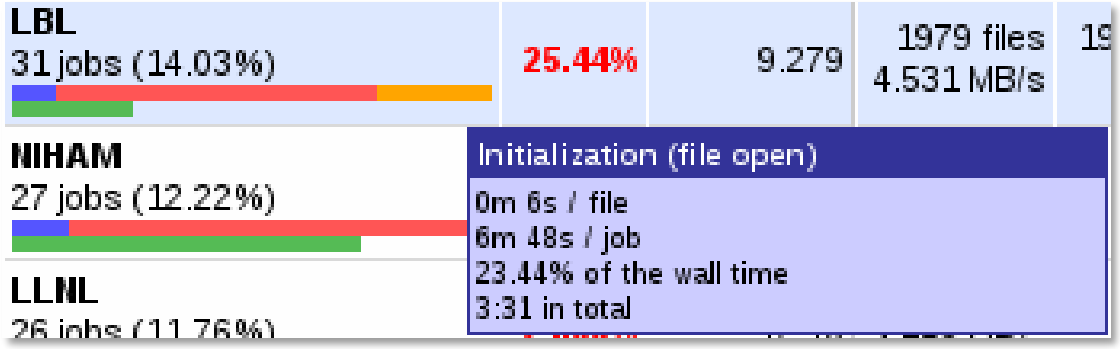

# LBL::SE traffic during that time

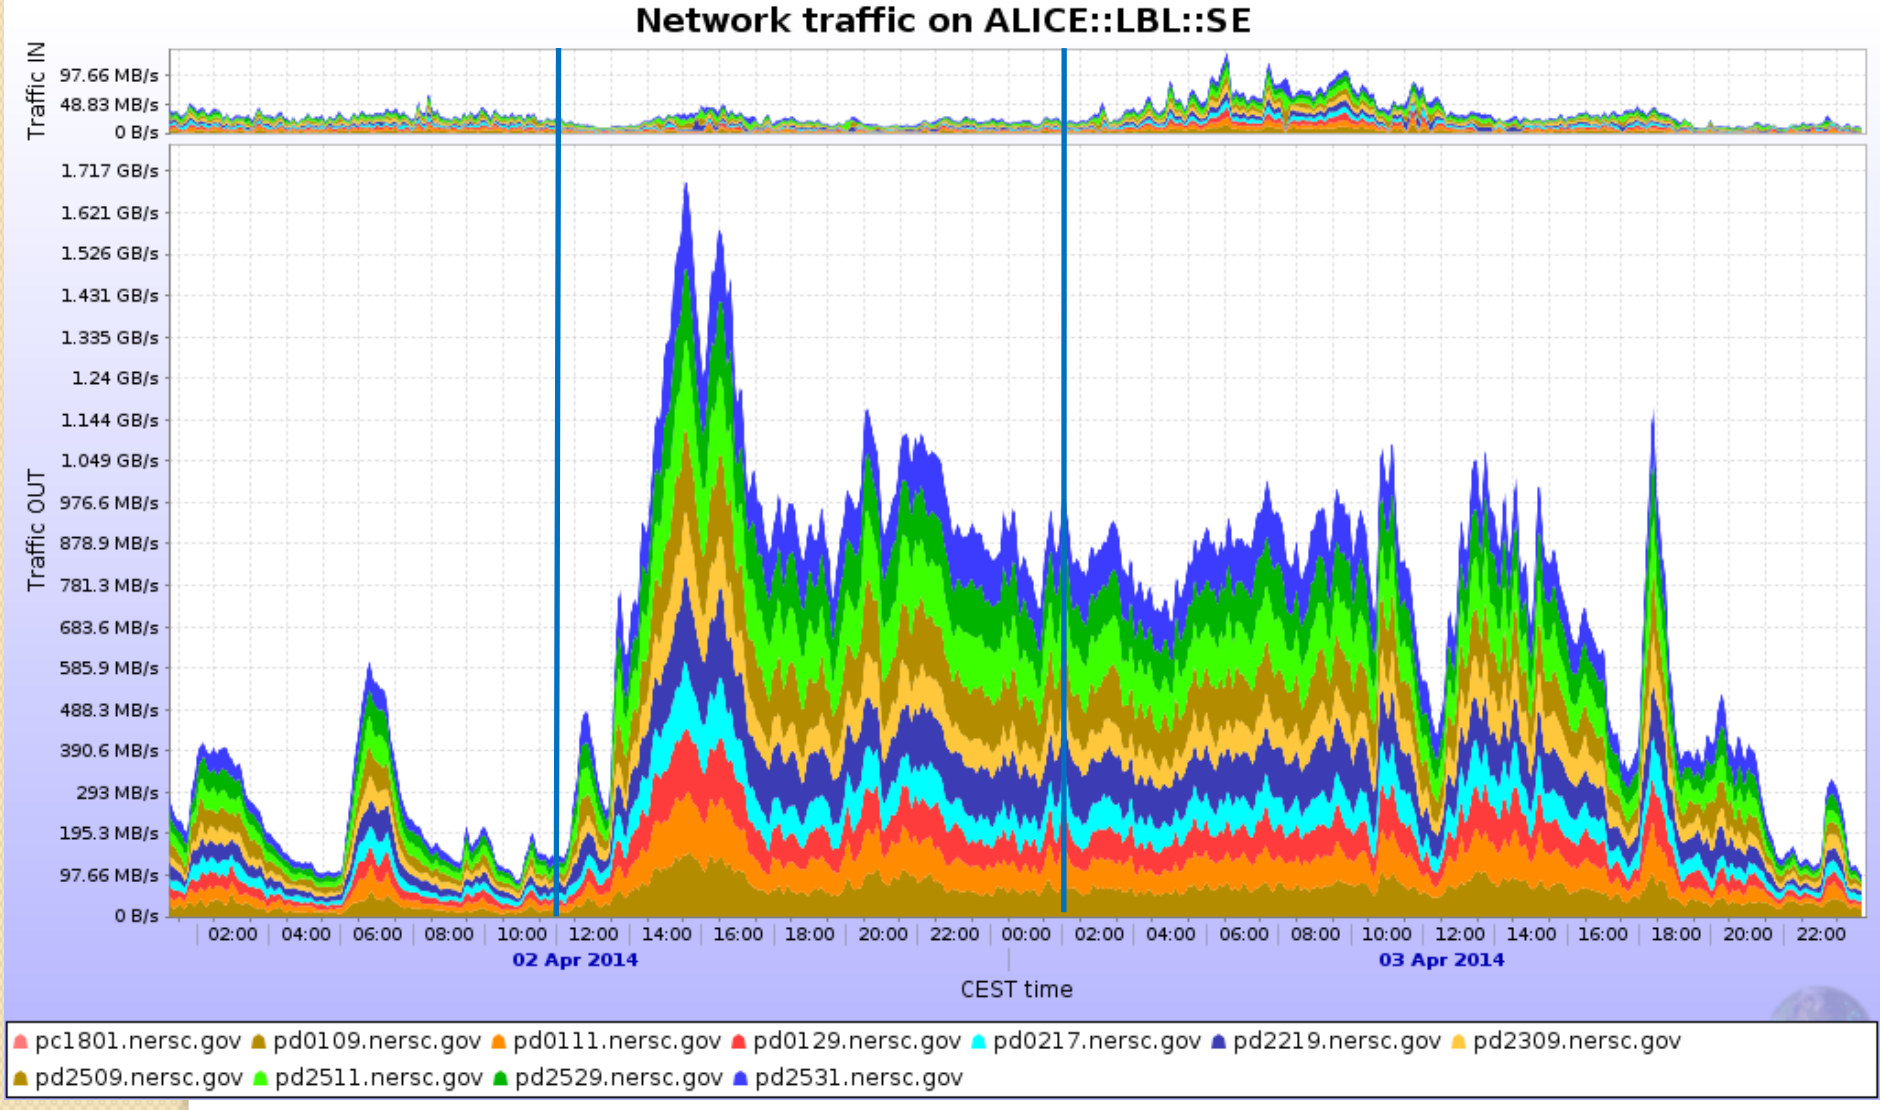

# LBL::SE server load

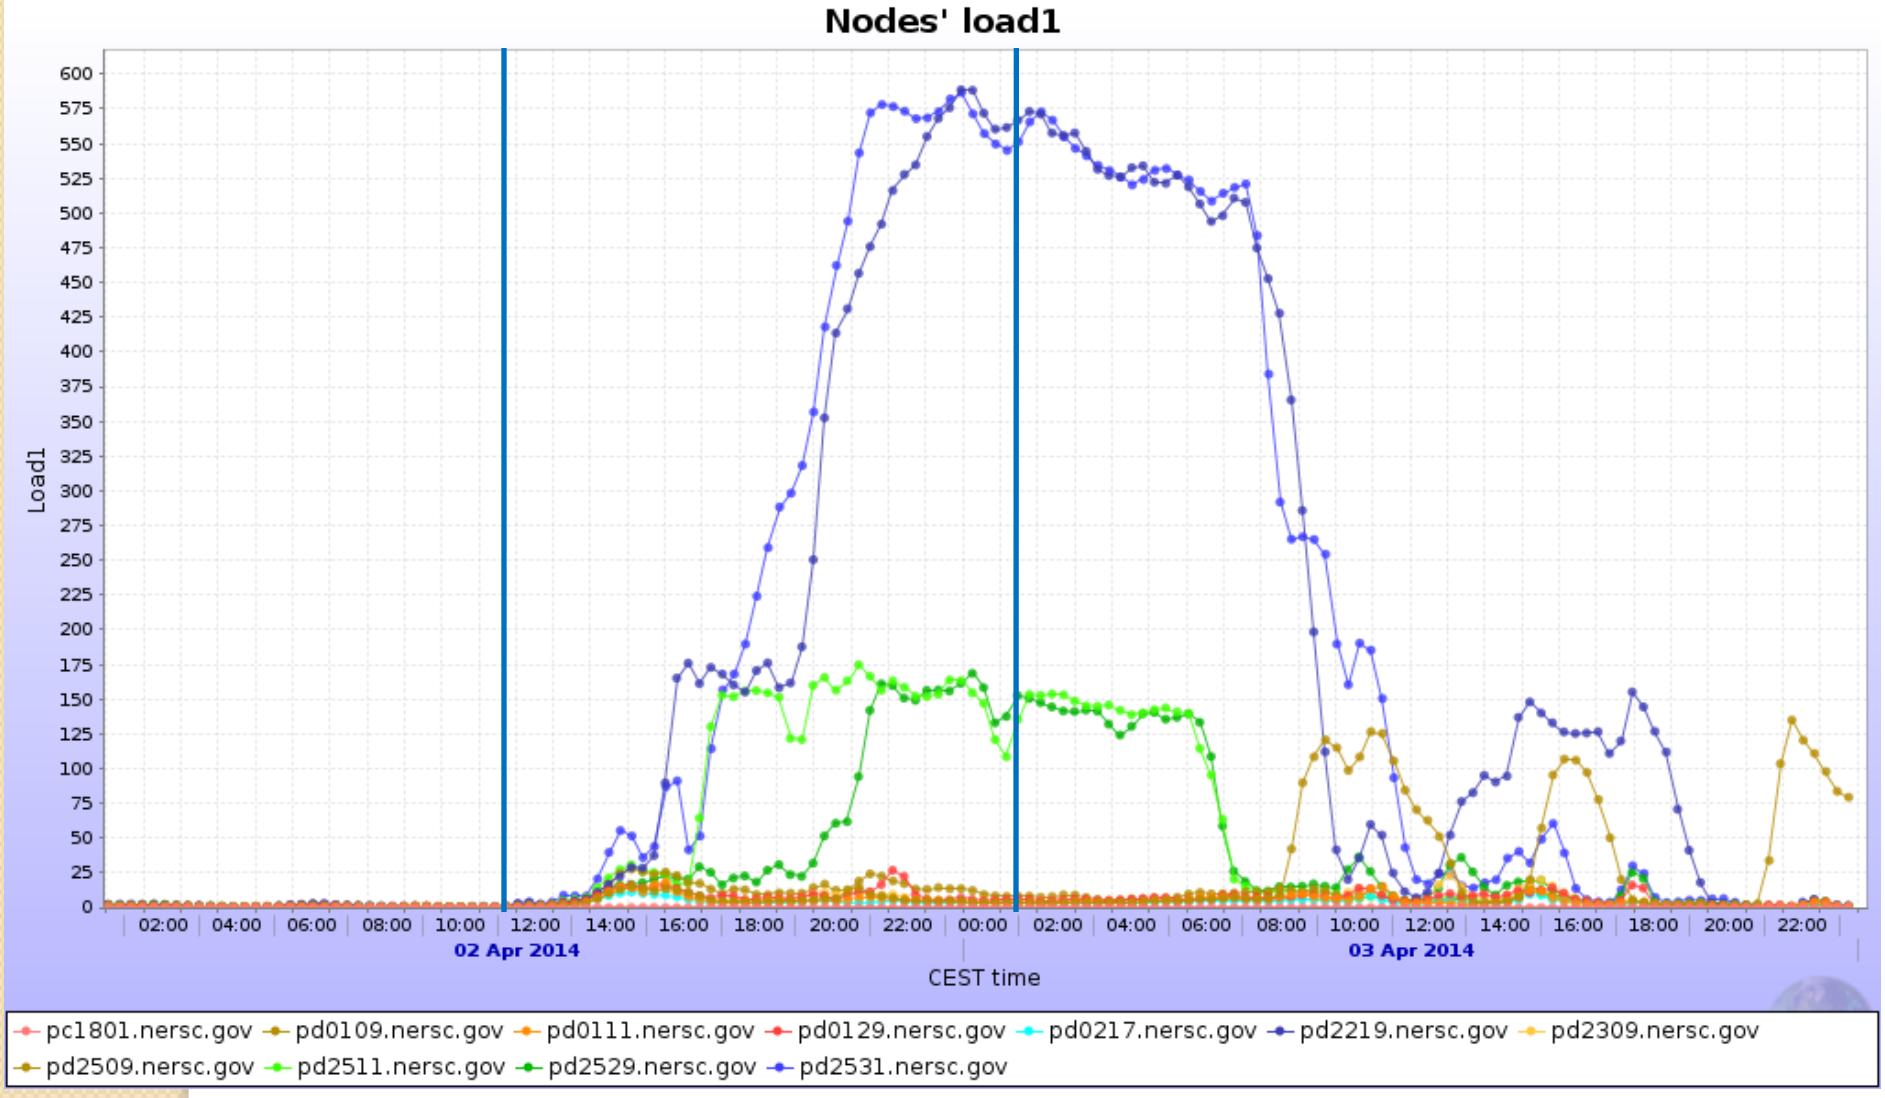

# LBL::SE socket count

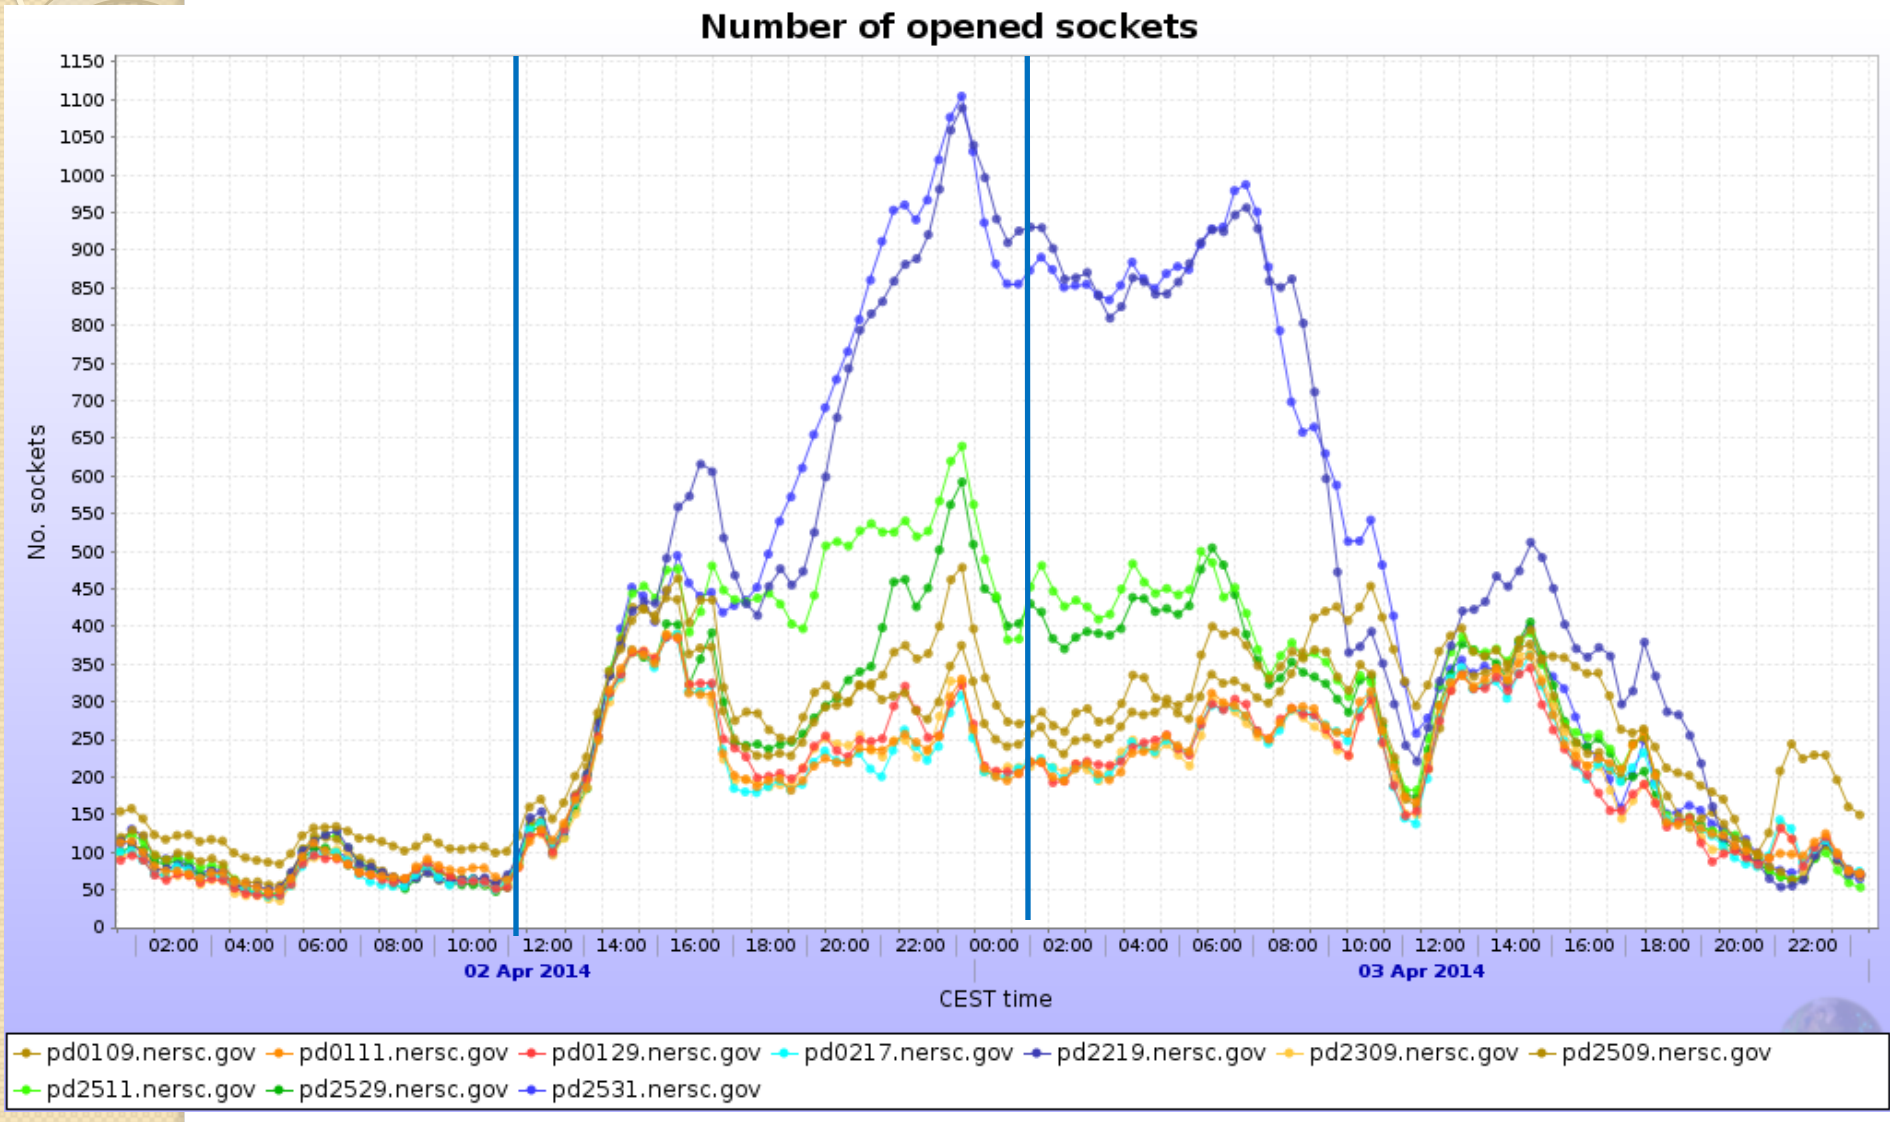

# LBL::SE top client sites

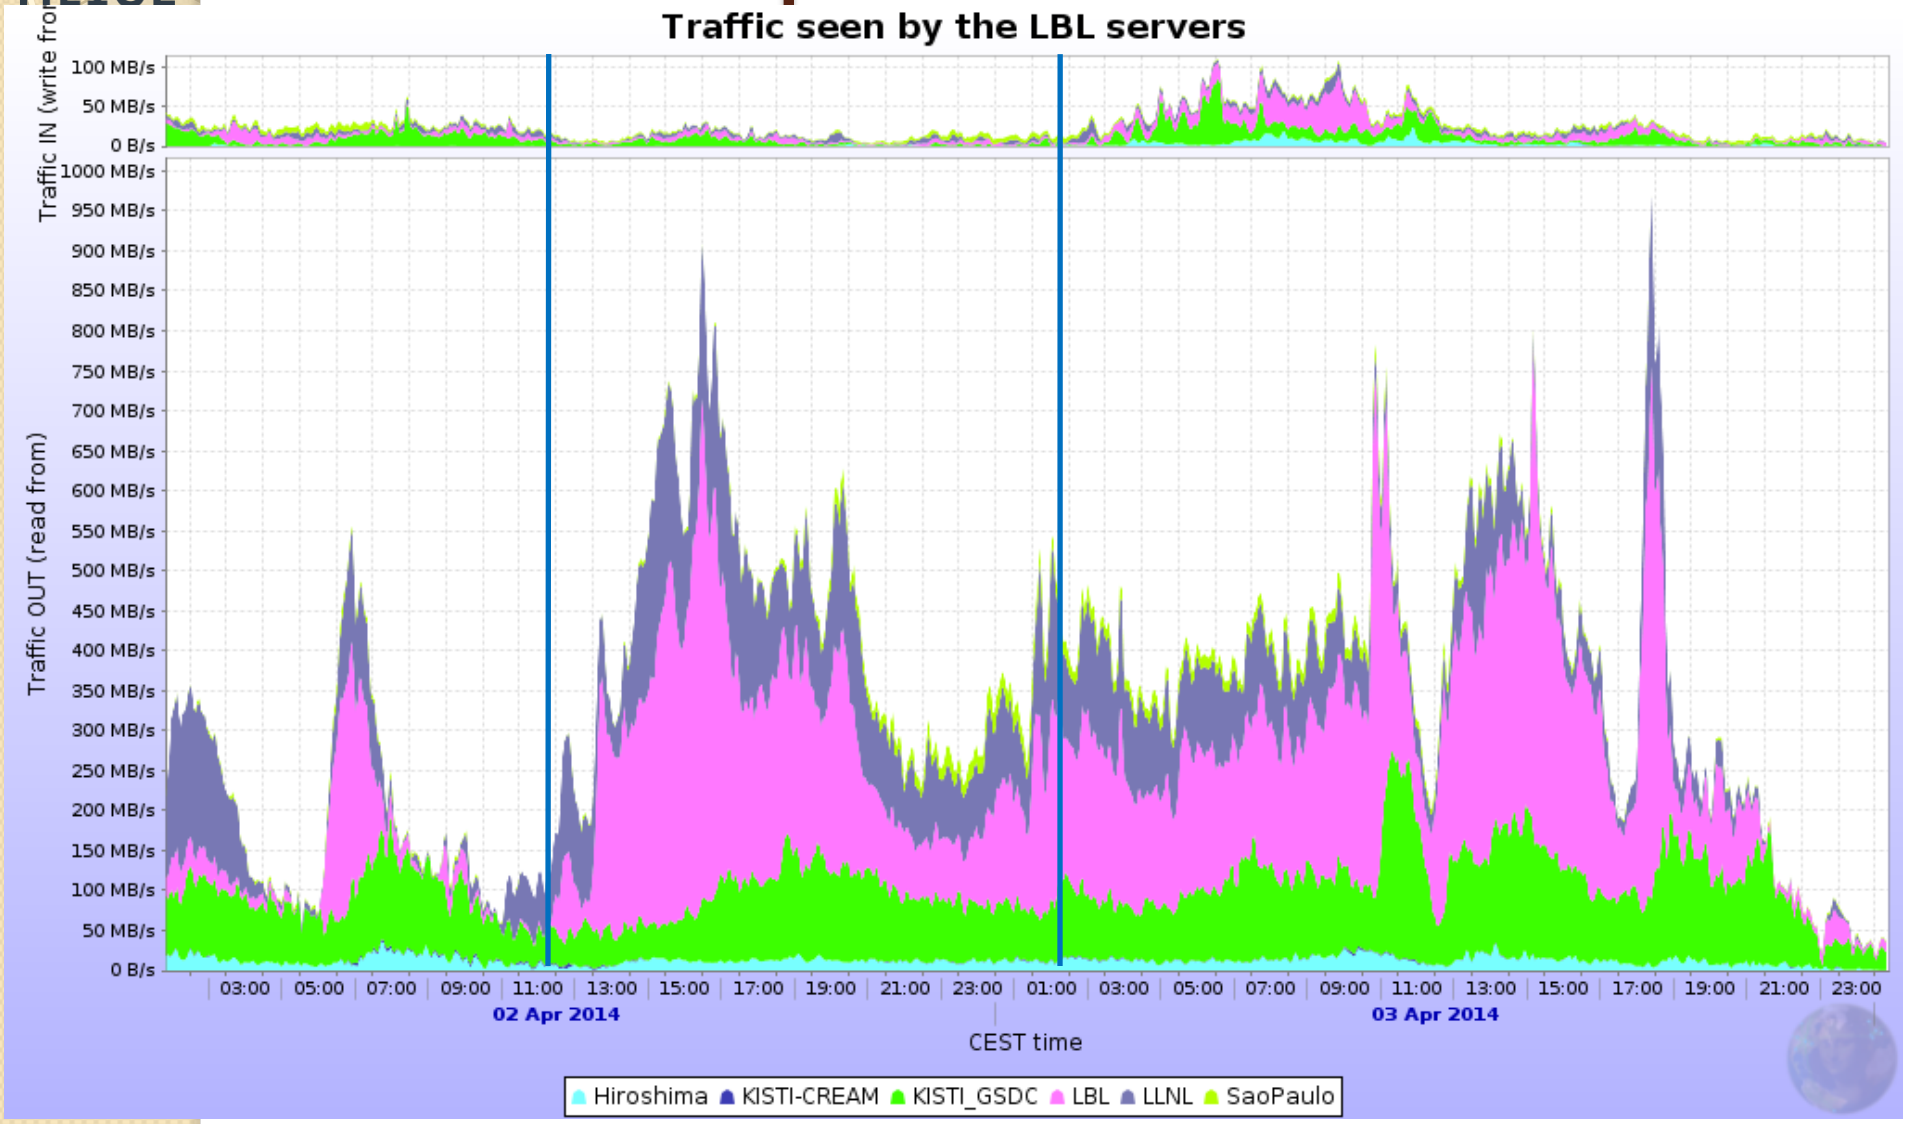

# LBL WNs data access

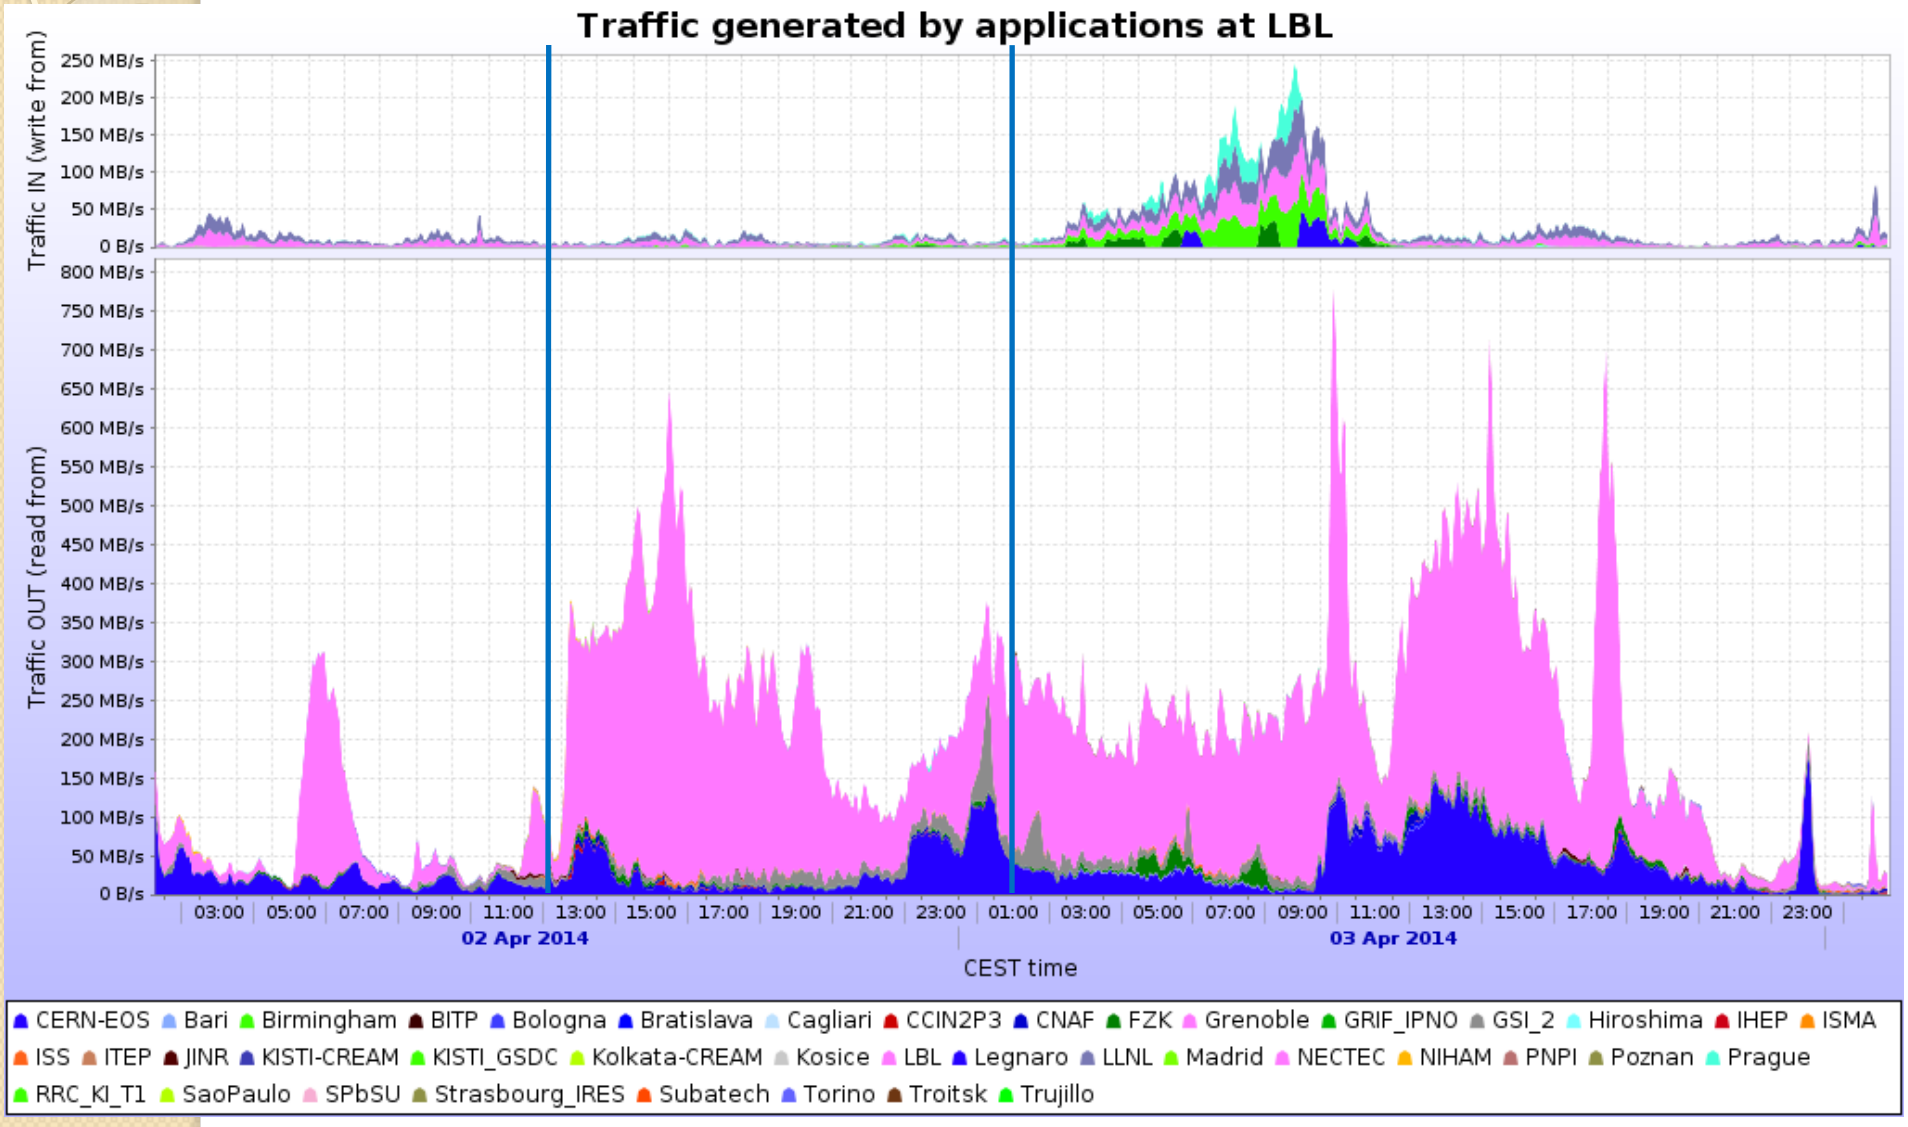

# LLNL WNs data access

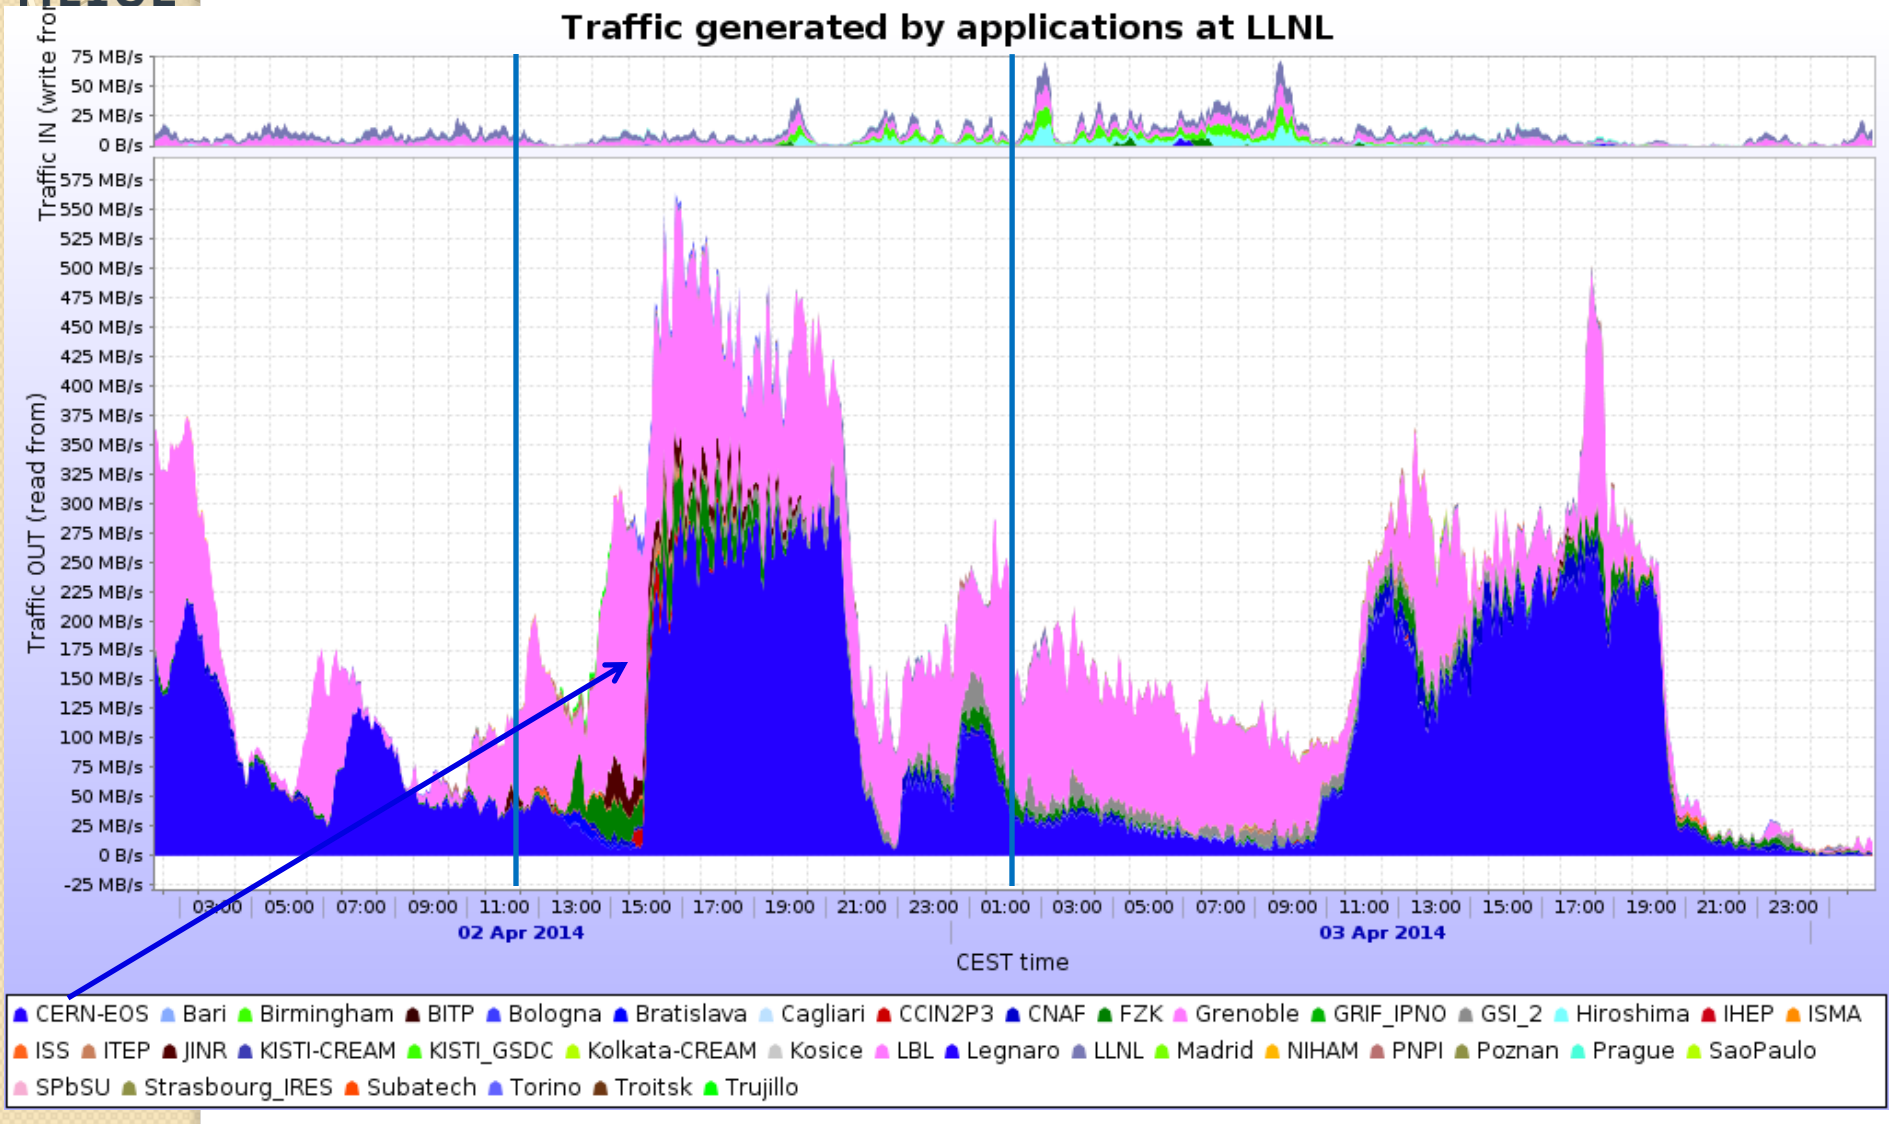

# Remote data access is significant

- Remember to tune **all** machines in your clusters for large average RTT (WNs, data servers, and use same values on the VoBox for reference)
- Kernel parameters as seen here: http://monalisa.cern.ch/FDT/documentation_syssettings.html
- Or even better the ESNet recommended values: http://fasterdata.es.net/host-tuning/linux/

# Network I TCP stream throughput