MonALISA

MONitoring Agents using a Large Integrated Services Architecture

Monitoring and Workflow management in large distributed systems

March 2011

The MonALISA Framework

- **MonALISA is a Dynamic, Distributed Service System capable to collect any type of information from different systems, to analyze it in near real time and to provide support for automated control decisions and global optimization of workflows in complex grid systems.**

- **The MonALISA system is designed as an ensemble of autonomous multi-threaded, self-describing agent-based subsystems which are registered as dynamic services, and are able to collaborate and cooperate in performing a wide range of monitoring tasks. These agents can analyze and process the information, in a distributed way, and to provide optimization decisions in large scale distributed applications.**

The MonALISA Architecture

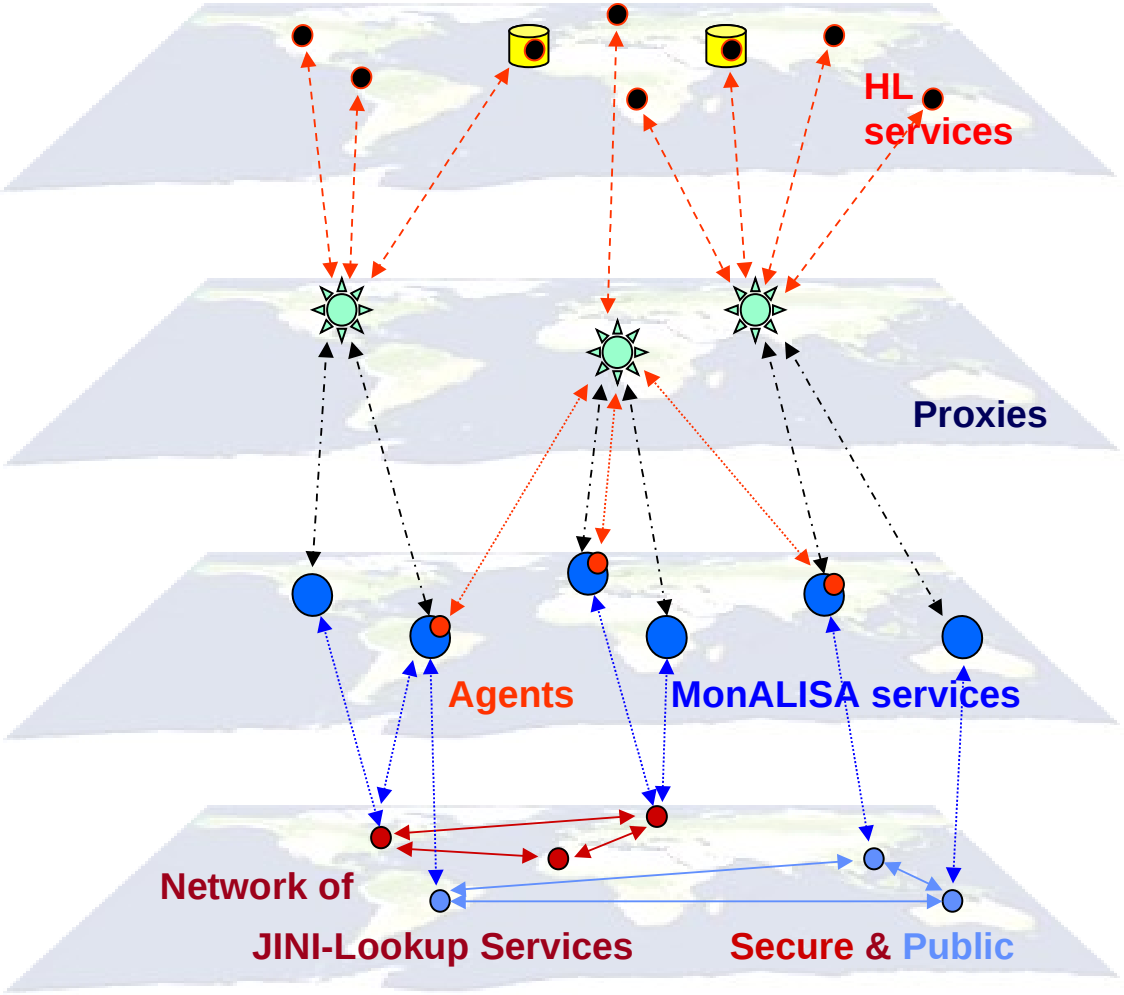

Regional or Global High Level Services,
Repositories & Clients

Secure and reliable communication
Dynamic load balancing
Scalability & Replication
AAA for Clients

Distributed System for gathering and analyzing information based on mobile agents:
Customized aggregation, Triggers, Actions

Distributed Dynamic Registration and Discovery- based on a lease mechanism and remote events

Fully Distributed System with no Single Point of Failure

MonALISA Service & Data Handling

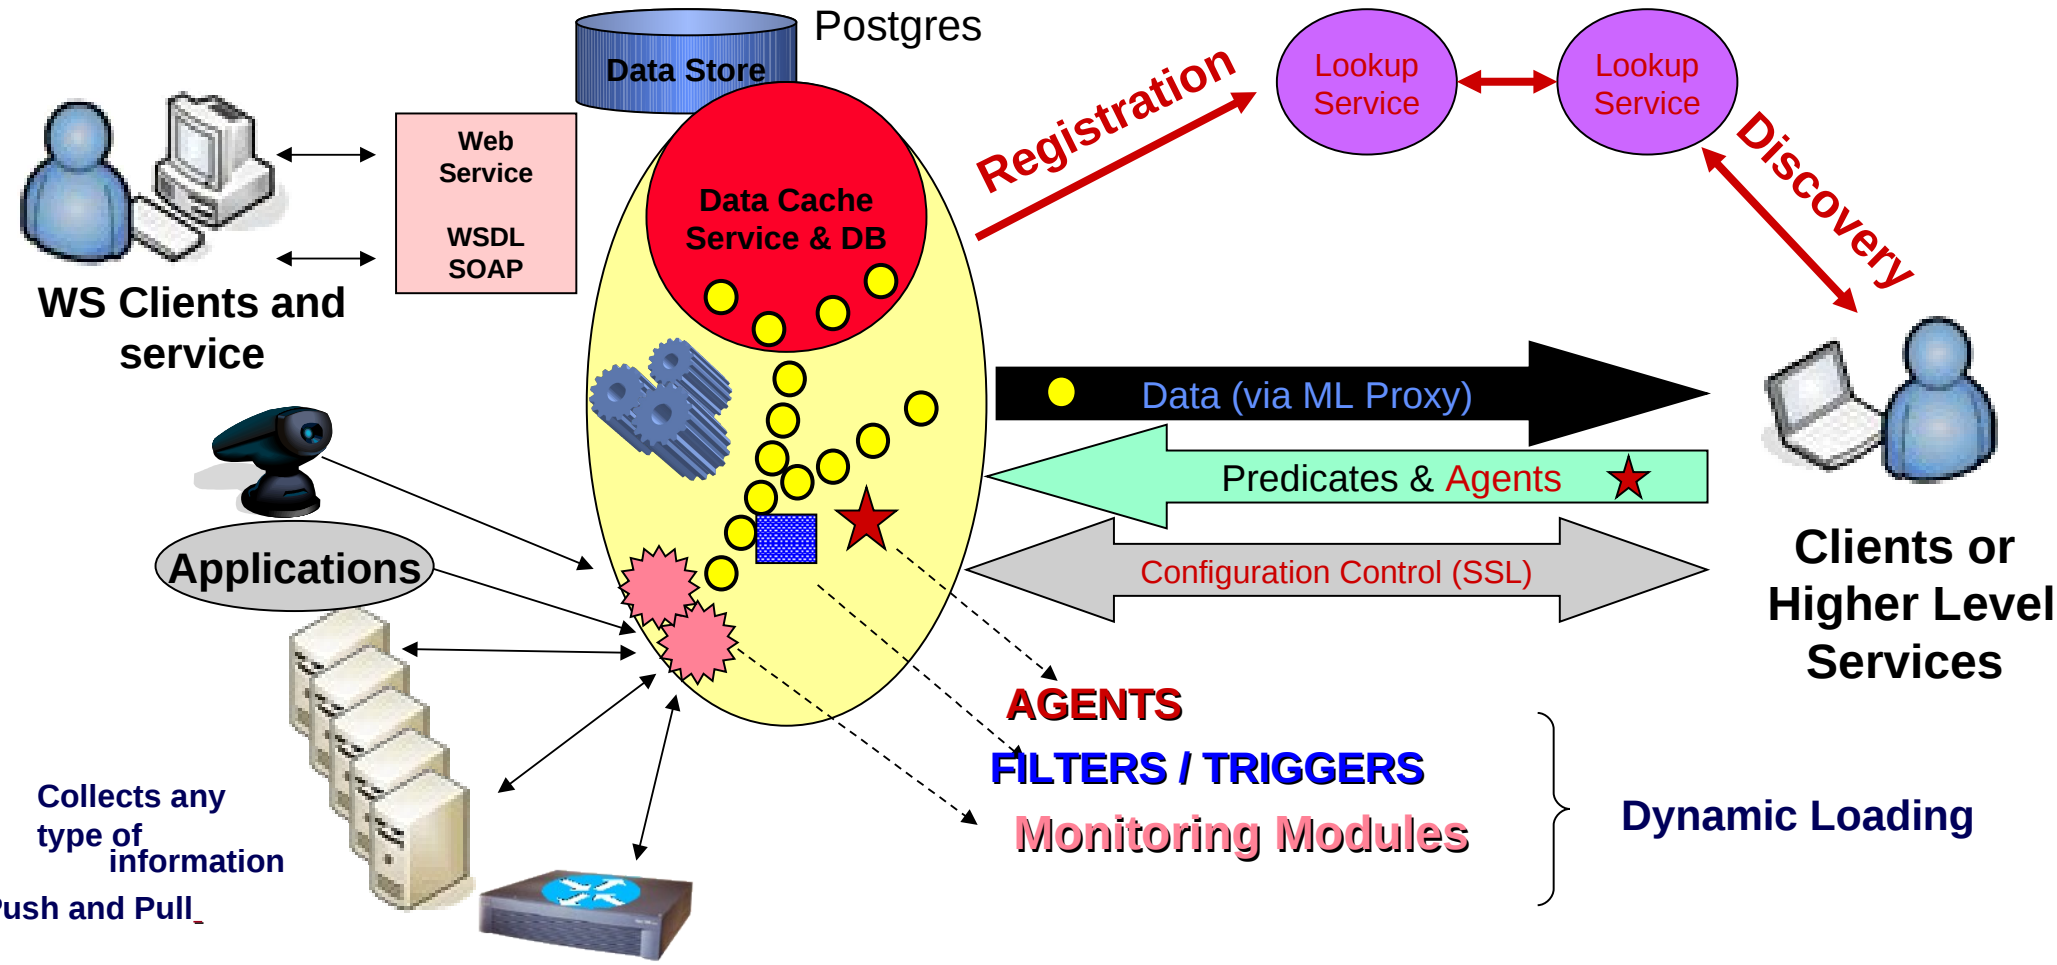

Registration / Discovery
Admin Access and AAA for Clients

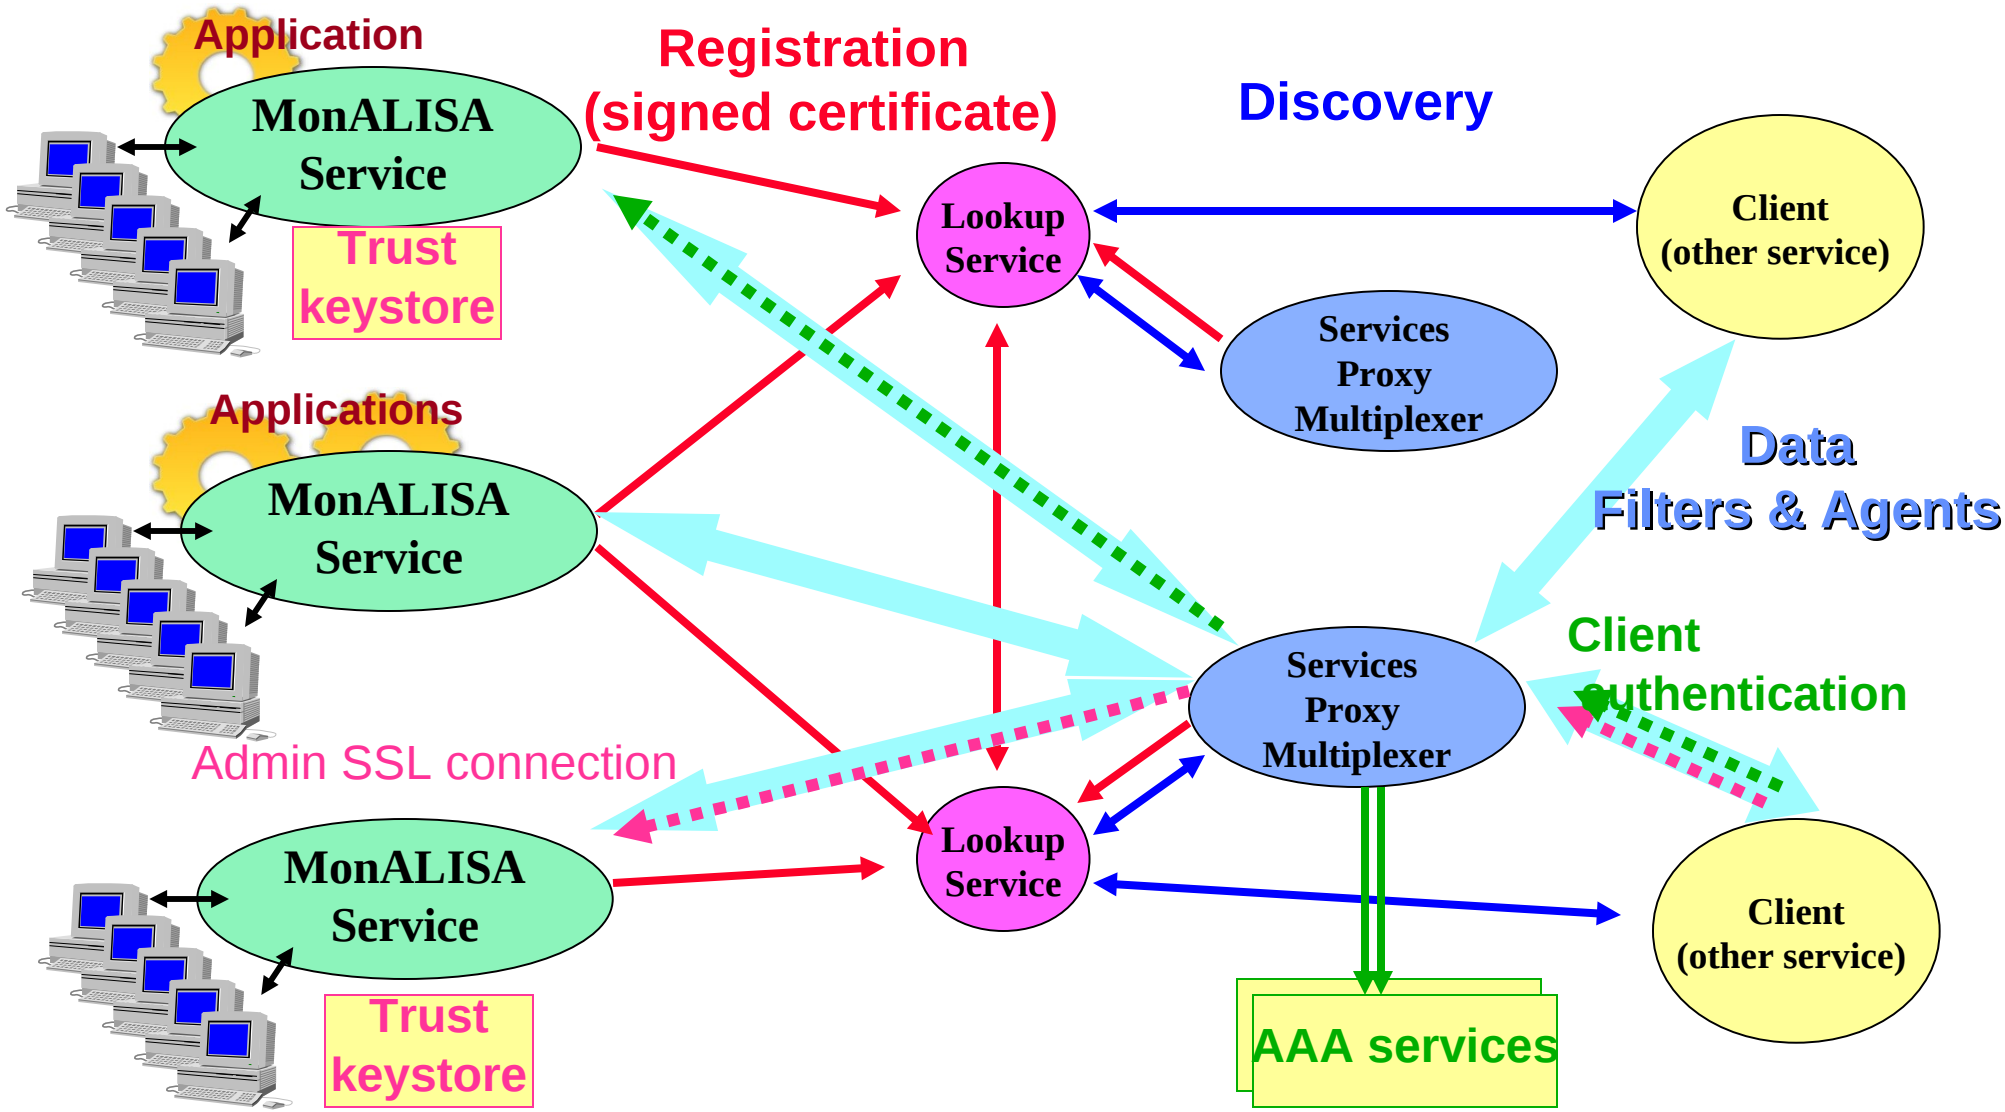

Monitoring Grid sites, Running Jobs, Network Traffic, and Connectivity

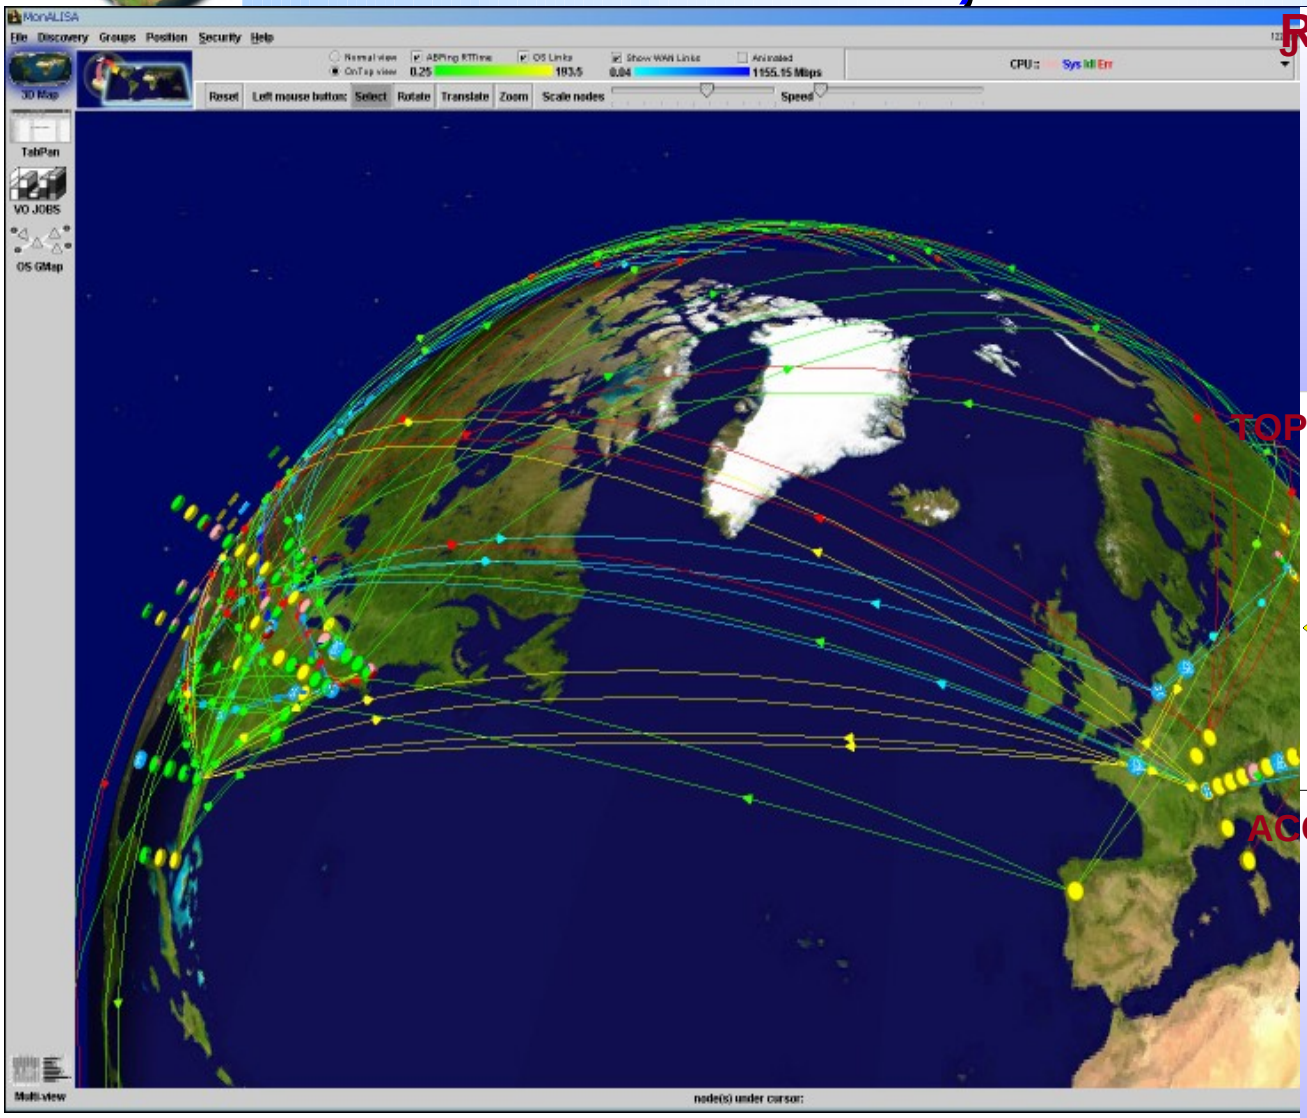

Running Jobs

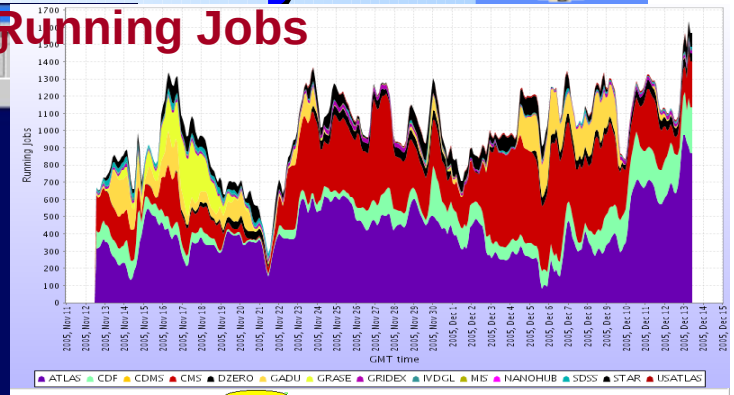

TOPOLOGY

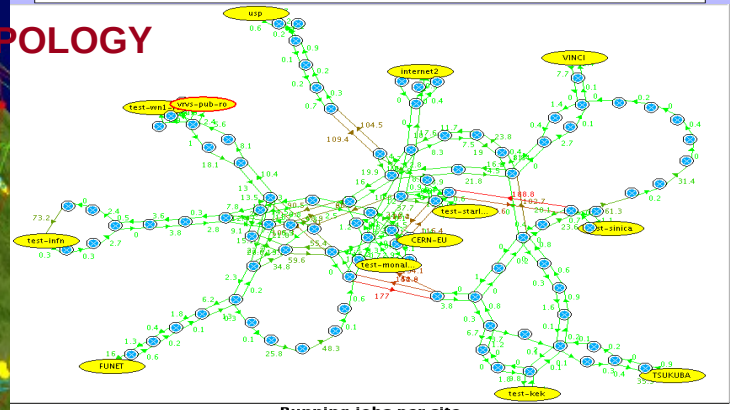

ACCOUNTING

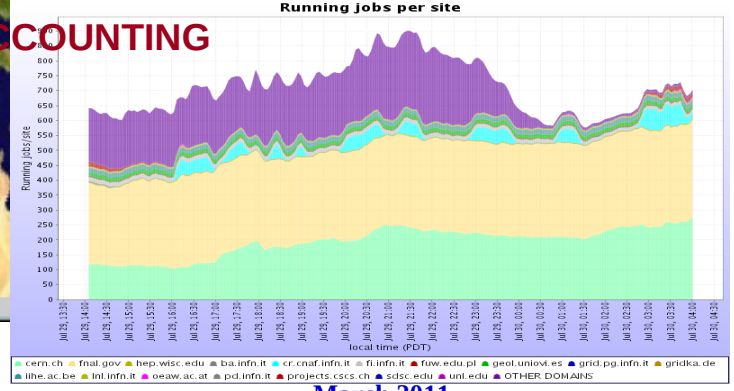

Monitoring CMS Jobs Worldwide

CMS is using MonALISA and ApMon to monitor all the production and analysis jobs. This information is than used in the CMS dashboard frontend

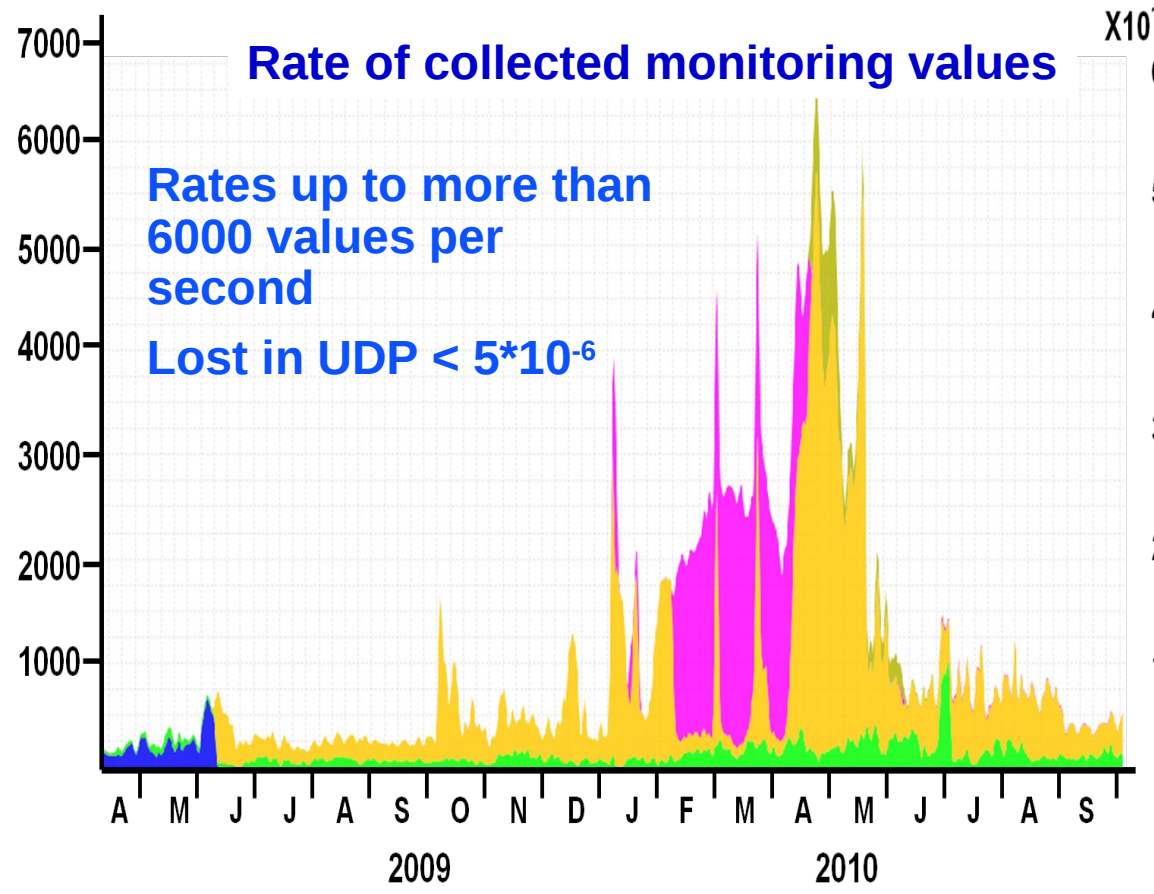

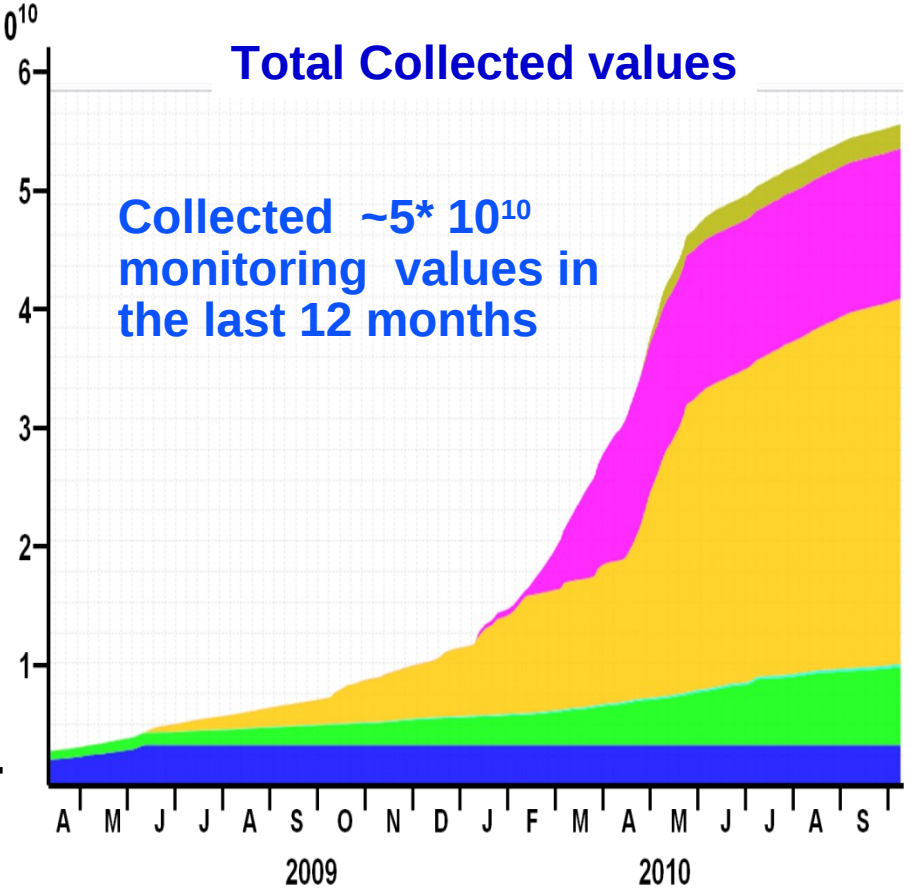

More than 3 years continuous operation without any problems

Monitoring CMS Jobs Worldwide

File Edit View History Bookmarks Tools Help

http://dashb-cms-job-task.cern.ch/dashboard/request.py/taskmonitoring#action=tasksTable&usergridname=YousiMa&timera

Most Visited | Release Notes | EVO | docs | Caltech T2 | Fedora | ultralight/~wart | Index of /u/aglo/pro... | Google Calendar | ahm.sbgrid.org

Task Monitoring

TASK MONITORING

Select a User: YousiMa Select a Time Range: Last Week Refresh: 5 Minutes Help User Support

Job Processing is not completed unless job GRID status is DONE. This page does not track further steps inside CRAB Server. Please ignore the GRID status for local submissions.

TaskMonitorId	Num of Jobs	Pending	Running	Appl Successful	Failed	Unknown	Completed Successfully	Consumed Time	Plots
yma_crab_0_110303_100839_dw4b06	1	0	0	1	0	0	DONE	Time Info	Plot Selection
yma_crab_0_110303_113157_wje871	3	0	0	3	0	0	2 out of 3	Time Info	Plot Selection
yma_crab_0_110303_113851_q681bp	3	0	0	2	1	0	2 out of 3	Time Info	Plot Selection
yma_crab_0_110303_114424_23dv8z	2	0	0	2	0	0	DONE	Time Info	Plot Selection
yma_crab_0_110303_114221_14rq2c	3	0	0	3	0	0	DONE	Time Info	Plot Selection
yma_crab_0_110303_114858_23n9gz	113	0	0	11	102	0	11 out of 113	Time Info	Plot Selection
yma_crab_0_110303_114615_d6y34i	3	0	0	3	0	0	DONE	Time Info	Plot Selection
yma_crab_0_110303_115037_bzq371	19	0	0	8	11	0	1 out of 19	Time Info	Plot Selection
yma_crab_0_110303_115310_j983py	6	0	0	5	1	0	4 out of 6	Time Info	Plot Selection
Sum Total	153	0	0	38	115	0	-	-	-

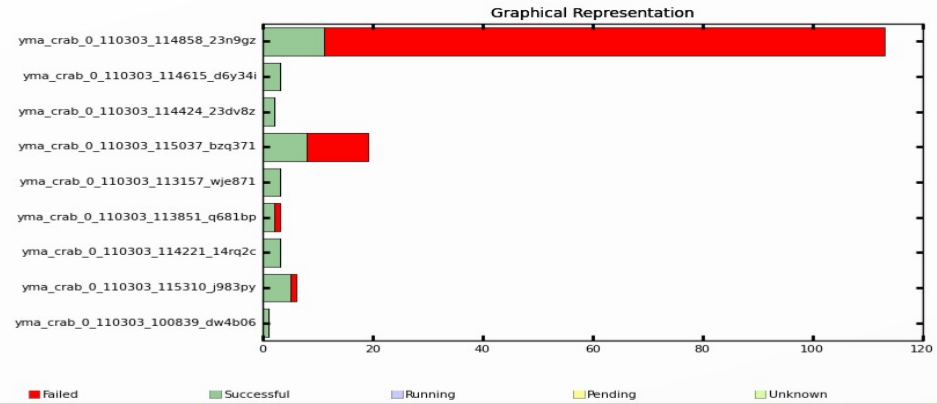

Transferring data from dashb-cms-job-task.cern.ch... Connect Error While Signing In FoxyProxy: Patterns

User-level task monitoring

Monitoring CMS Jobs Worldwide

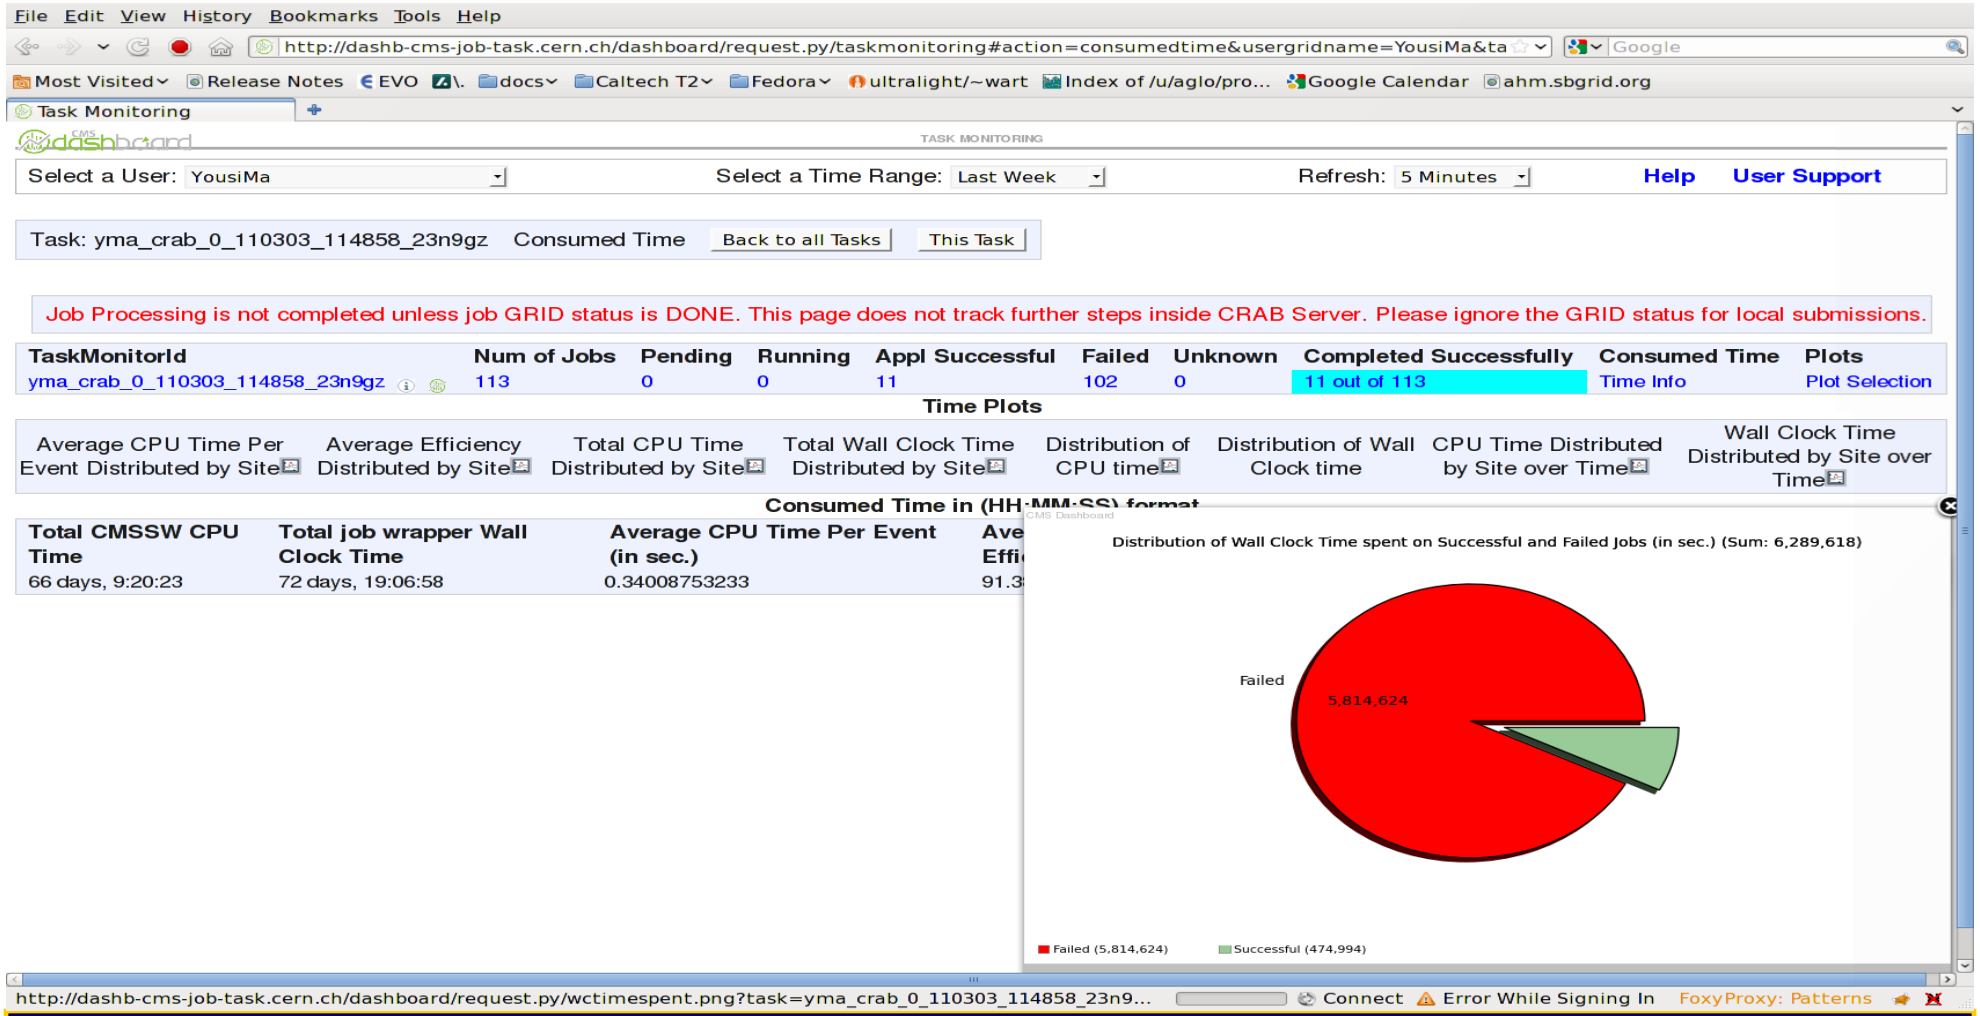

User-level task monitoring

Monitoring architecture in ALICE

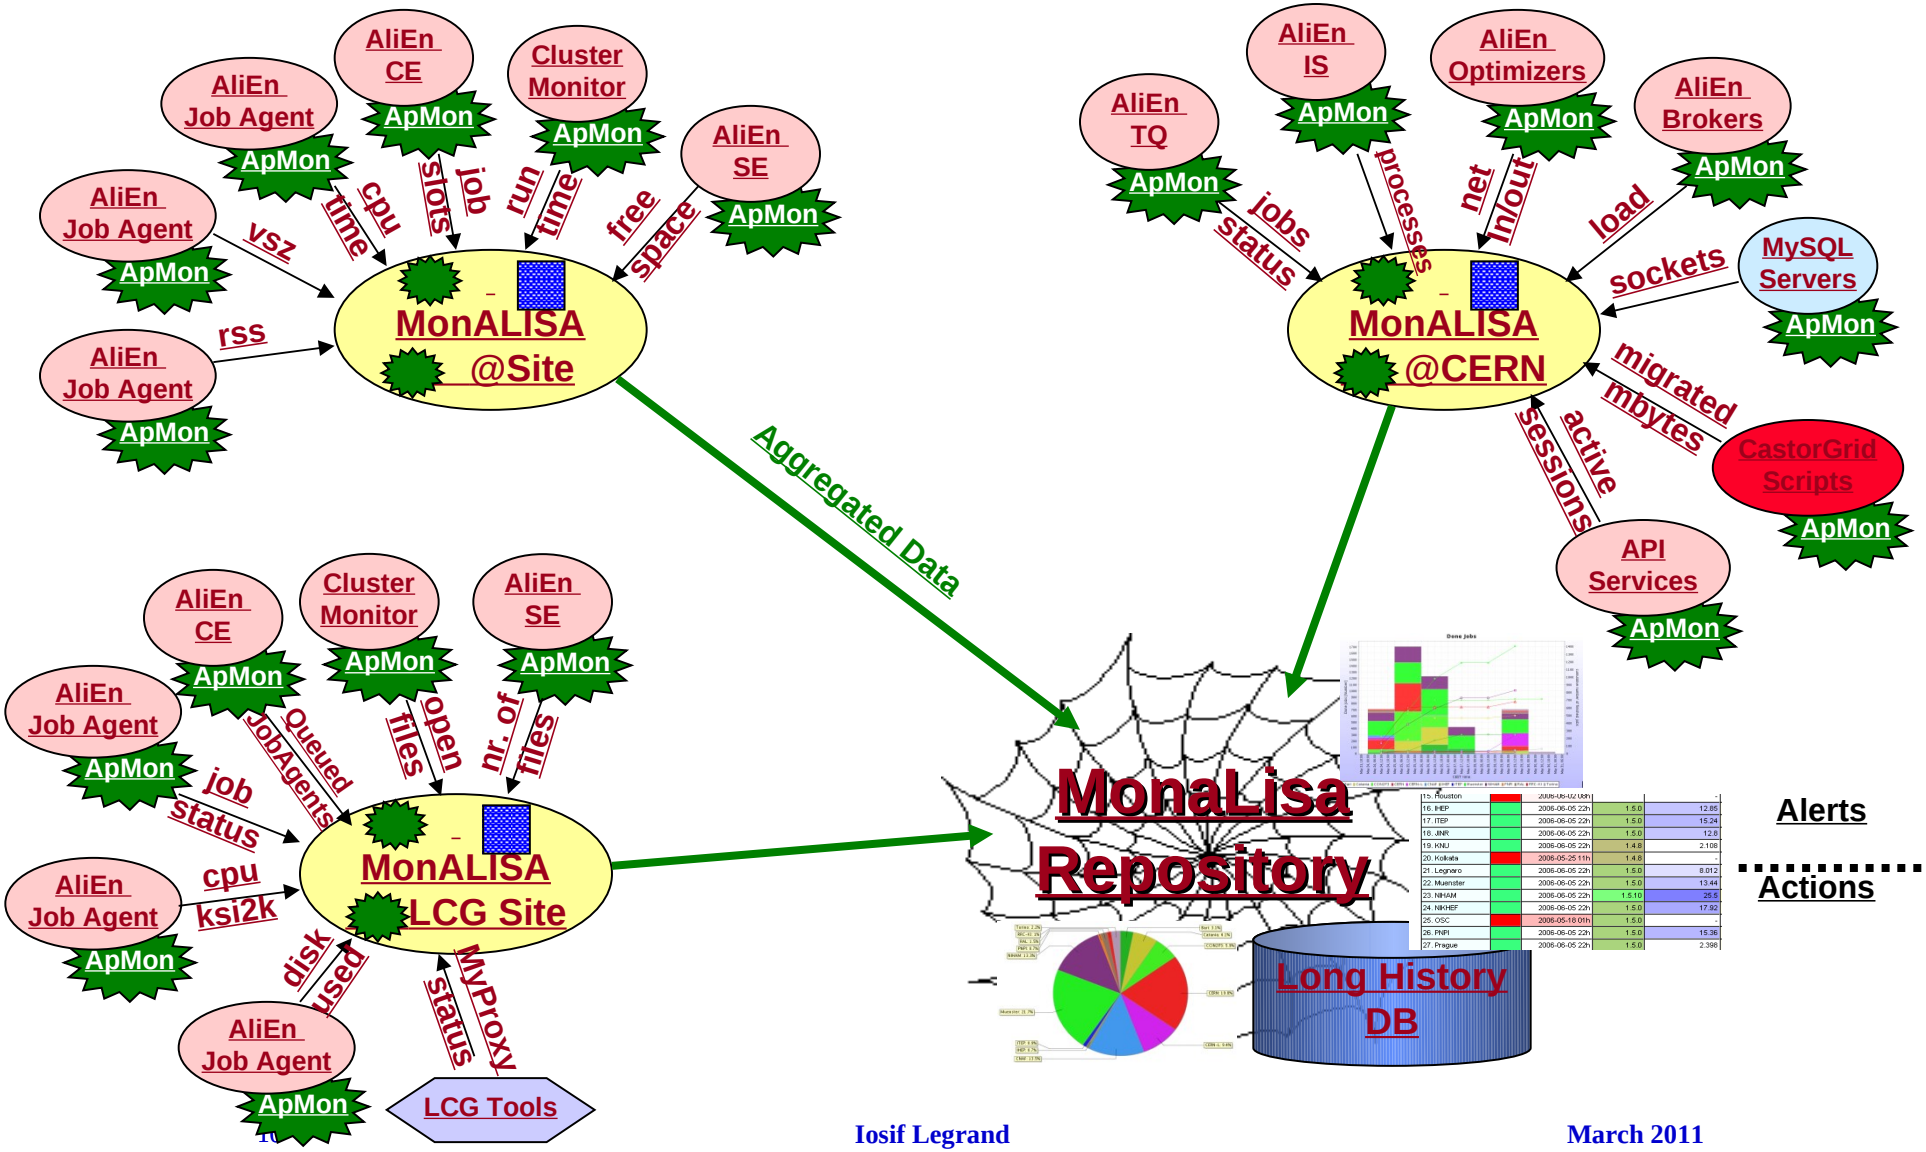

ALICE : Global Views, Status & Jobs

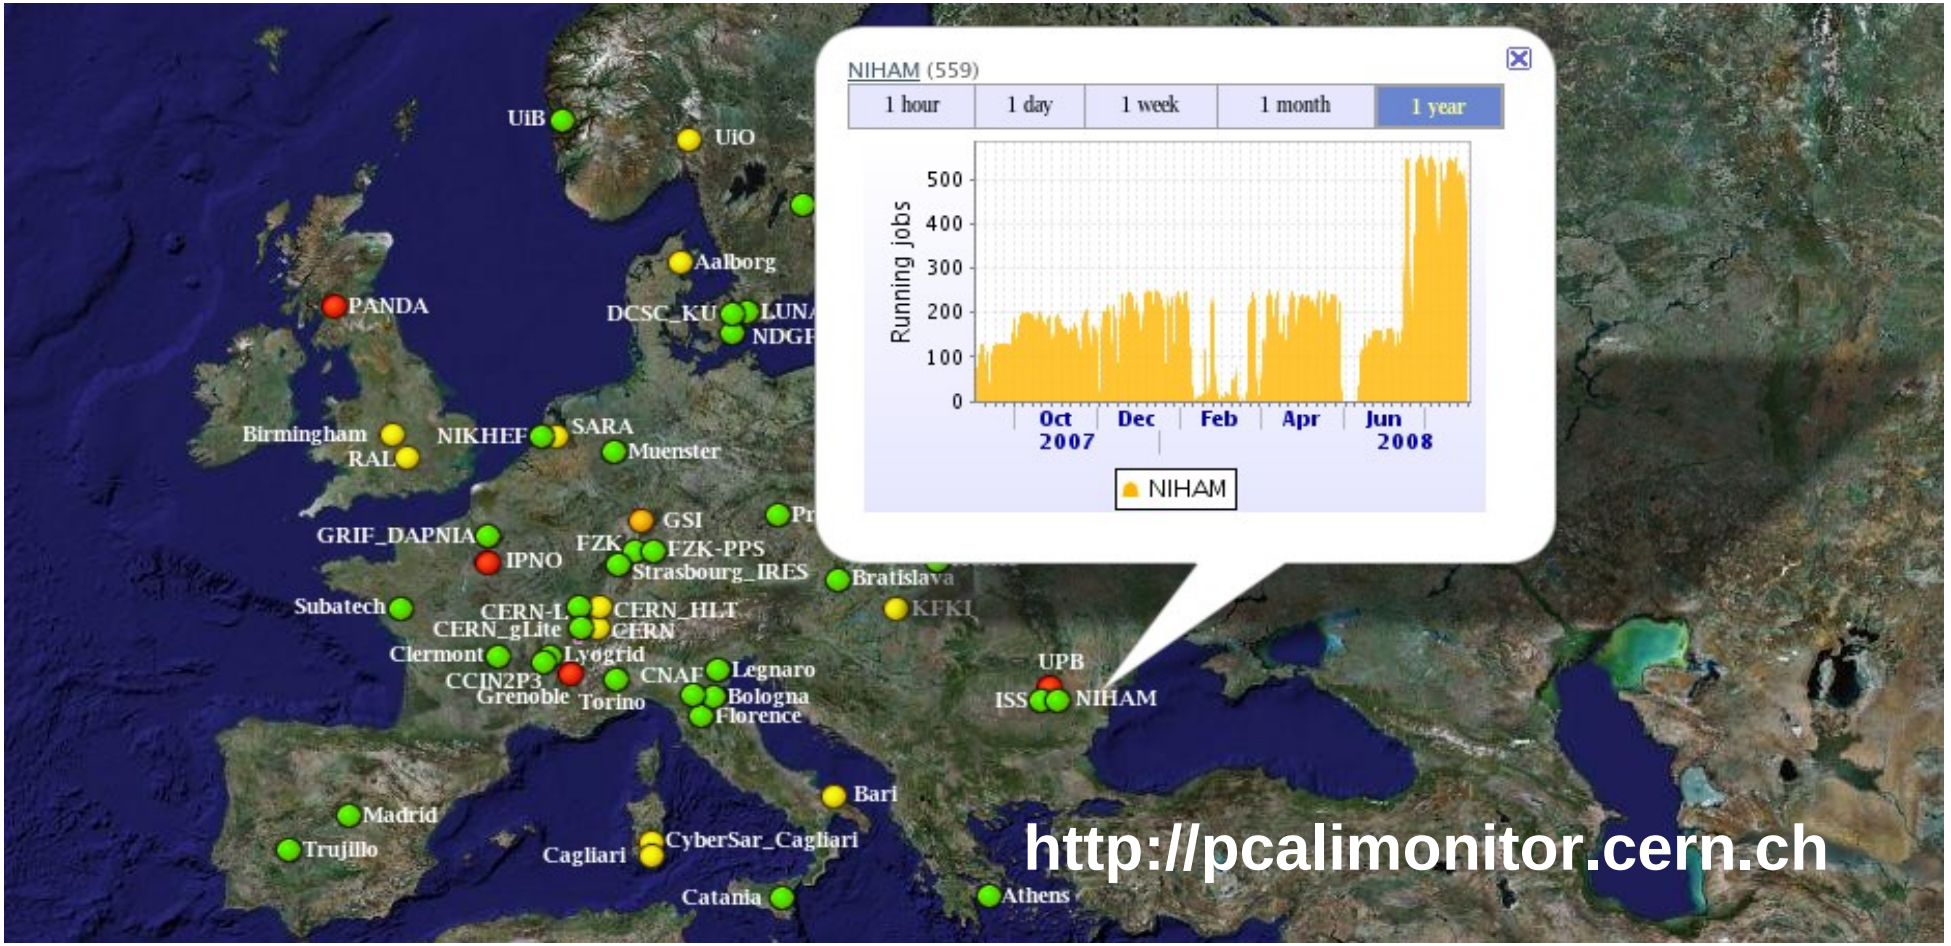

Monitoring in ALICE: jobs, resources, services

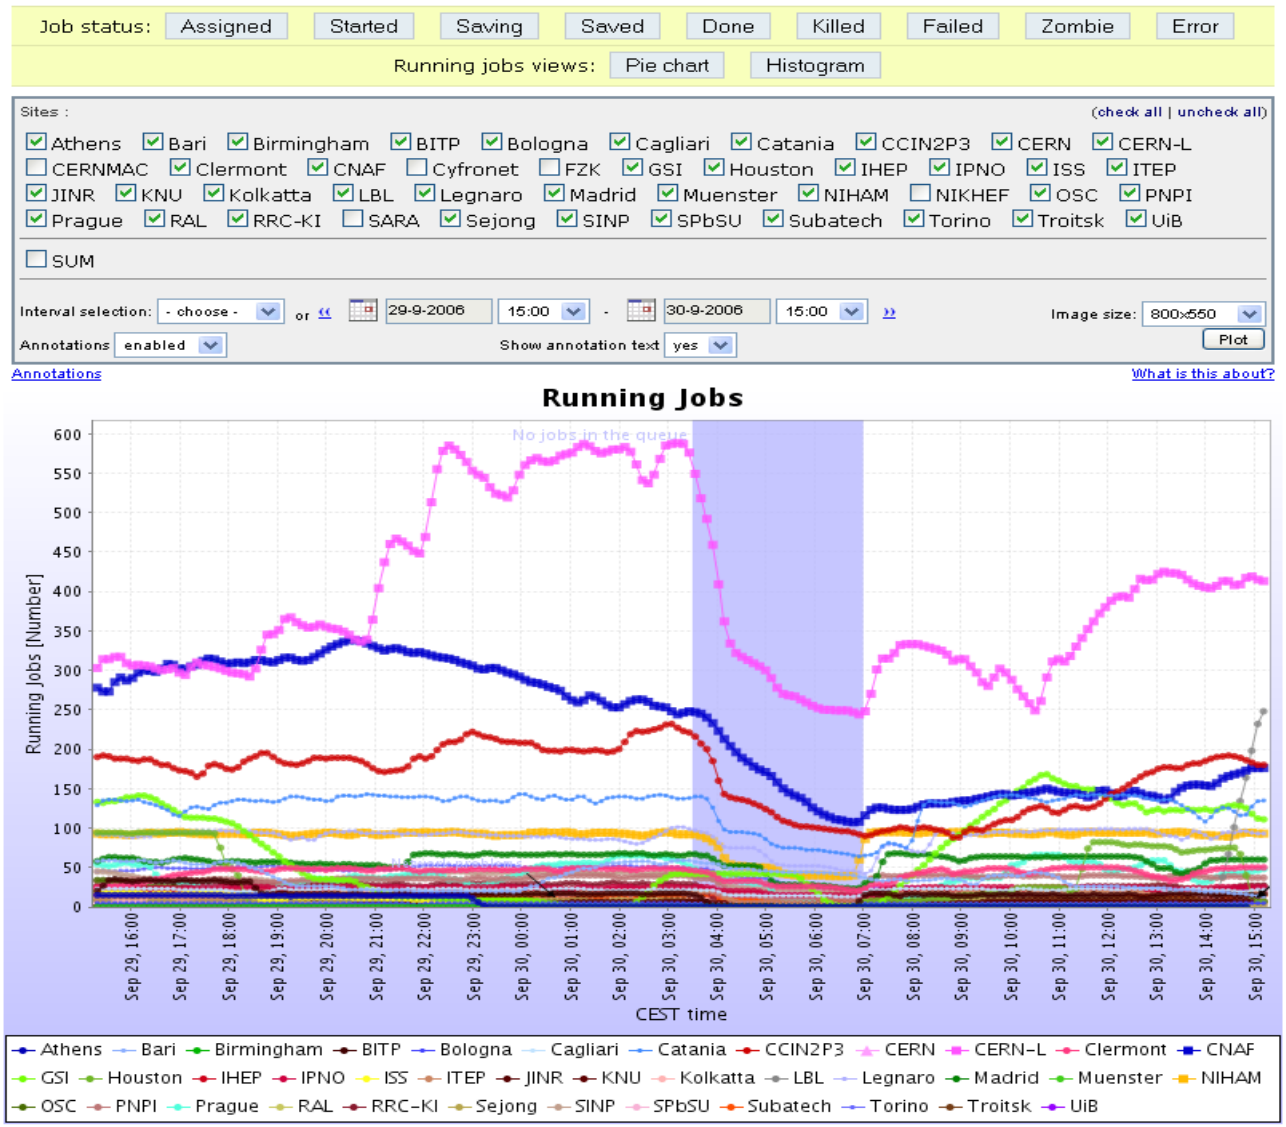

Running Jobs				
Farm	Last value	Min	Avg	Max
Athens	0	0	4.746	15
Bari	13	13	37.05	60
Birmingham	0	0	0	0
BITP	16	0	16.02	35
Bologna	5	0	4.291	5
Cagliari	22	13	21.13	23
Catania	135	63	125.2	144
CCIN2P3	180	84	164.1	234
CERN	0	0	0.041	1
CERN-L	413	238	390.8	591
Clermont	50	20	42.56	50
CNAF	176	106	224.7	340
GSI	111	1	66.9	172
Houston	7	0	35.44	94
IHEP	0	0	0	0
IPNO	25	13	23.29	25
ISS	0	0	4.508	20
ITEP	18	3	16.25	20
JINR	7	0	4.365	8
KNU	0	0	0	0
Kolkatta	20	0	15.47	20
LBL	248	0	133	262
Legnaro	93	41	86.98	102
Madrid	60	21	57.56	68
Muenster	0	0	0	0
NIHAM	93	39	87.2	95
OSC	0	0	6.076	34
PNPI	37	9	33.99	40
Prague	48	0	37.88	67
RAL	0	0	0	0
RRC-KI	27	1	21.67	31
Sejong	2	1	6.375	10
SINP	27	12	30.19	45
SPbSU	3	1	2.995	4
Subatech	15	7	13.8	15
Torino	22	17	37.53	53
Troitsk	6	5	7.72	10
UiB	19	4	17.04	24
Total	**1898**		**1776**	

Monitoring in ALICE: Xrootd servers

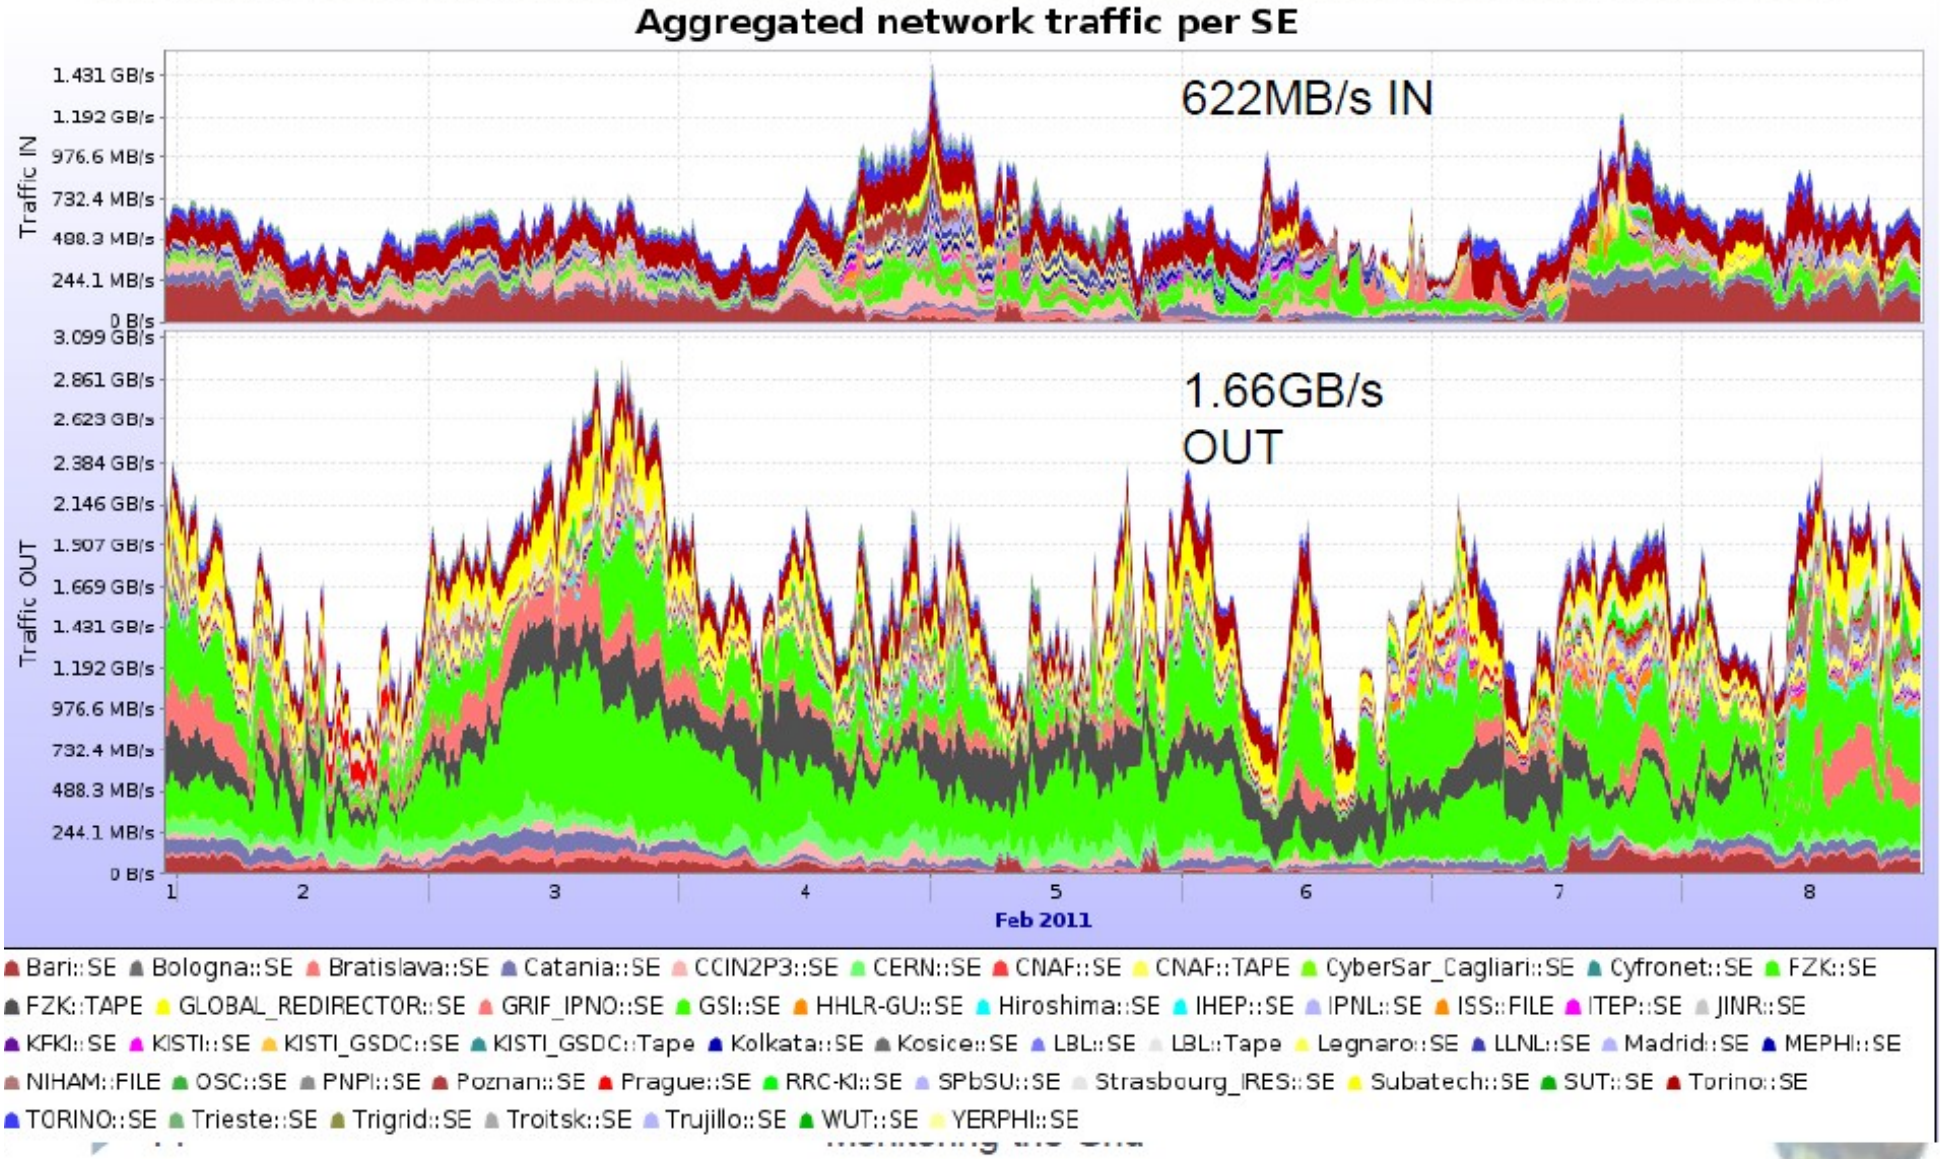

Active Available Bandwidth measurements between all the ALICE grid sites

Aalborg »

Links: FDT, Kernel parameters tuning

‹Aalborg›

Chart view »

IN from

No.	ID	Site	Speed (Mbps)	Hops	RTT (ms)	Streams
1.	126976	NDGF	685.81	11	6.87	1
2.	131876	DCSC_KU	430.88	6	6.61	1
					27.88	1
					34.49	1
					38.15	1
						1
					26.09	1
					59.96	1
					29.27	1
					23.93	1
						1
					21.91	1
					29.97	1
					40.07	1
					42.44	1
					43.52	1
					40.54	1
					24.97	1
						1

OUT to

No.	ID	Site	Speed (Mbps)	Hops	RTT (ms)	Streams
1.	127538	UiB	679.24	16	33.91	1
2.	128970	IPNO	662.03	17	36.19	1
3.	129355	NDGF	627.51	11	6.78	1
4.	127195	DCSC_KU	564.75	7	6.38	1
5.	126998	LUNARC	314.01	14	31.54	1
6.	130490	ISS	162.100	19	49.94	1
7.	129827	CSC				
8.	130994	CNAF				
9.	128512	CNAF-CR				
10.	130365	OSC				
11.	126963	SARA				
12.	130267	NIHAM				
13.	127450	Kolkata-C				
14.	129399	RAL				
15.	128153	CERN-L				
16.	131295	Prague				
17.	131055	Kolkata				
18.	127177	PNPI				
19.	130170	GSI				
20.	129558	Grenoble				
21.	129903	Catania				
22.	127138	SINP				
23.	131236	Trujillo				
24.	92520	UPB				
25.	131713	Madrid				
26.	126729	TriGrid				
27.	129296	Legnaro				
28.	131748	ITEP				
29.	129381	KPI				

Nodes configuration for test 128970

‹Aalborg›	Source
IP	130.225.192.122
OS	Ubuntu 8.04.1
Kernel	2.6.24-17-server
TCP algo	reno
Write buffers	8388608 (4096 1875000 8388608)
Suggestions	

‹IPNO›	Target
IP	134.158.78.52
OS	Scientific Linux SL release 4.6 (Beryllium)
Kernel	2.6.9-67.0.4.ELlargesmp
TCP algo	
Receive buffers	8388608 (4096 87380 8388608)
Suggestions	

Tests from **Aalborg** to IPNO

No.	ID	Speed (Mbps)	Hops	RTT (ms)	Streams
1.	128970	662.03	17	36.19	1
2.	123260	523.89	19	36.23	1
3.	117348	324.43	19	36.17	1
4.	112041	445.69	16	36.19	1
5.	107523	384.84	17	36.04	1

Tracepath for test 128970

Tracepath from **Aalborg** to **IPNO**

Hop	IP	RTT (ms)	Domain
0	**130.225.192.122**	0	aau.dk
1	**130.225.192.126**	0.57	aau.dk
2	**130.225.192.126**	0.47	aau.dk
3	**192.38.59.54**	0.59	
4	**192.38.59.213**	6.33	
5	**130.225.242.34**	6.28	fsknet.dk
6	**130.225.244.145**	6.93	fsknet.dk
7	**130.225.244.218**	6.72	fsknet.dk
8	**193.10.68.121**	6.68	nordu.net
9	**62.40.124.45**	6.68	geant2.net
10	**62.40.112.78**	19.66	geant2.net
11	**62.40.112.138**	27.71	geant2.net
12	**62.40.112.105**	35.11	geant2.net
13	**62.40.124.70**	35.73	geant2.net
14	**193.51.179.90**	35.74	
15	**193.51.188.161**	35.98	
16	no_reply		
17	**193.51.188.161**	36.19	

Target was not reached

Active Available Bandwidth measurements between all the ALICE grid sites (2)

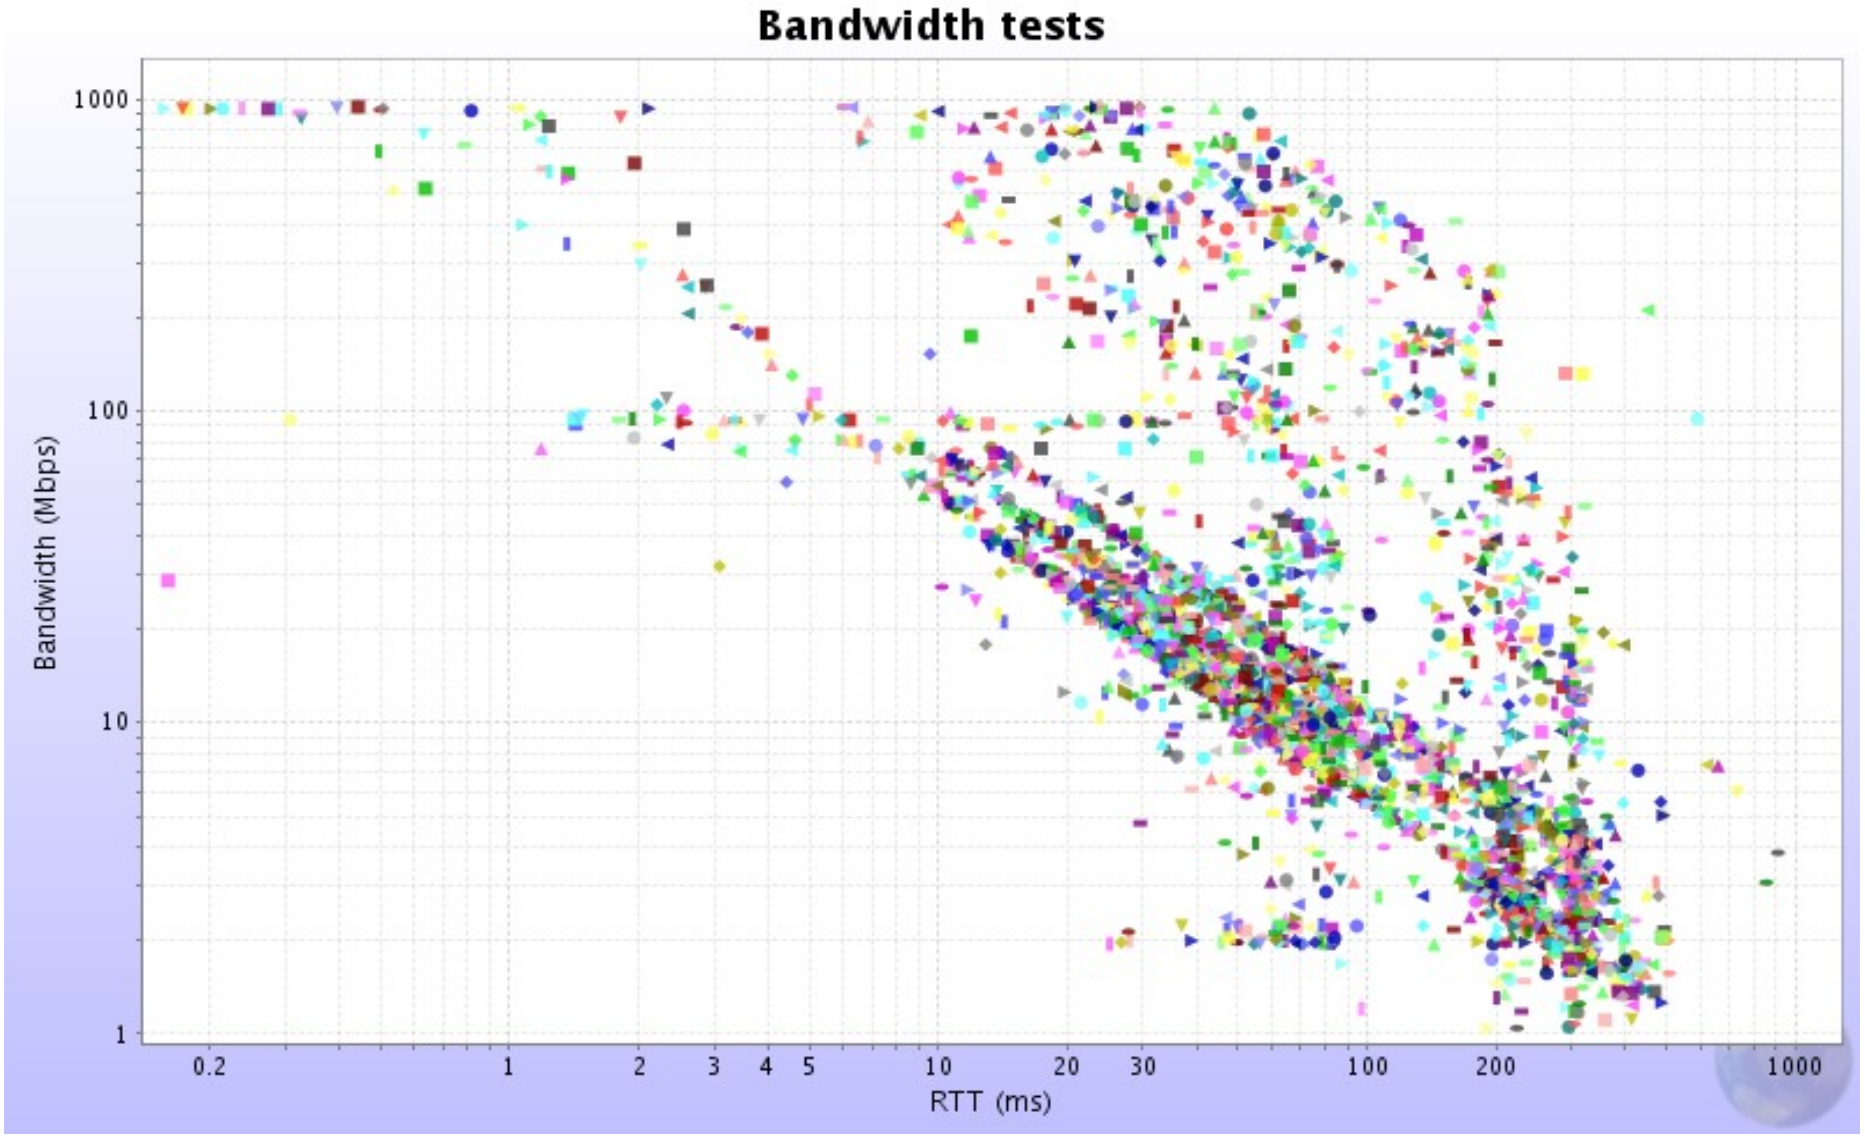

Local and Global Decision Framework

- **Two levels of decisions:**
  - local (autonomous),
  - global (correlations).
- **Actions triggered by:**
  - values above/below given thresholds,
  - absence/presence of values,
  - correlations between any values.
- **Action types:**
  - alerts (emails/instant msg/atom feeds),
  - automatic charts annotations in the repository,
  - running custom code, like securely ordering MLs service to change connectivity – optimize traffic, submit jobs, (re)start global service.

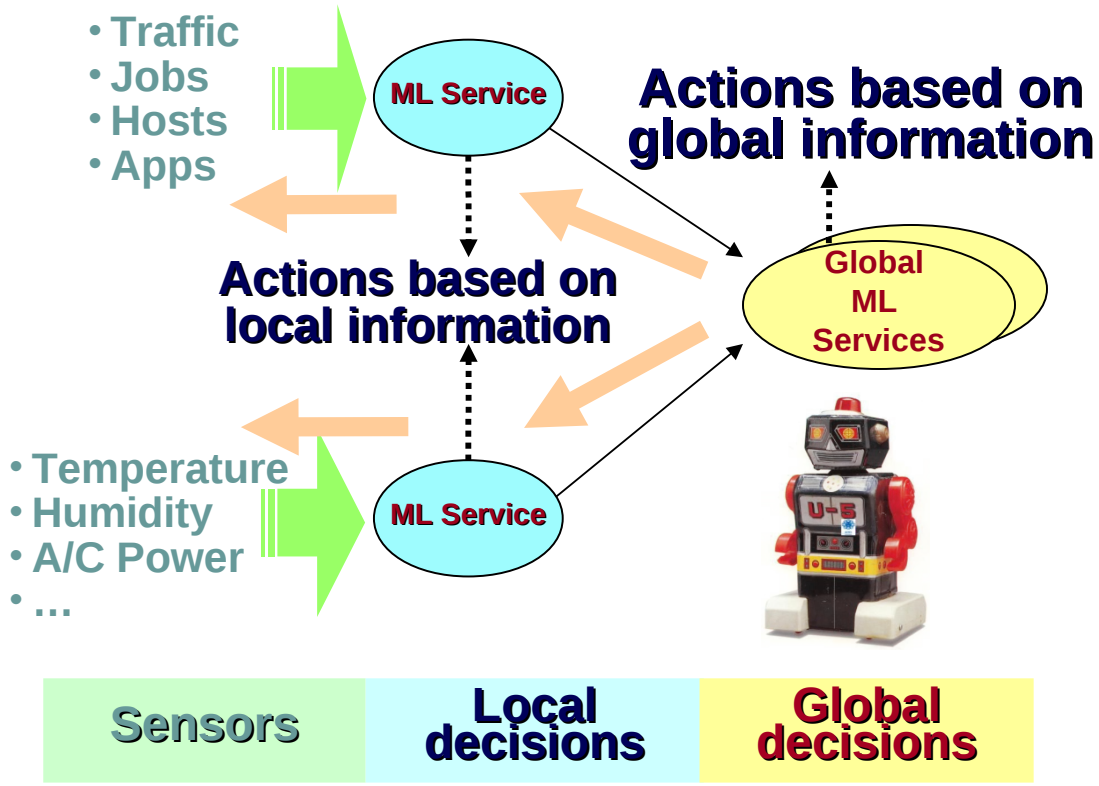

ALICE: Automatic job submission Restarting Services

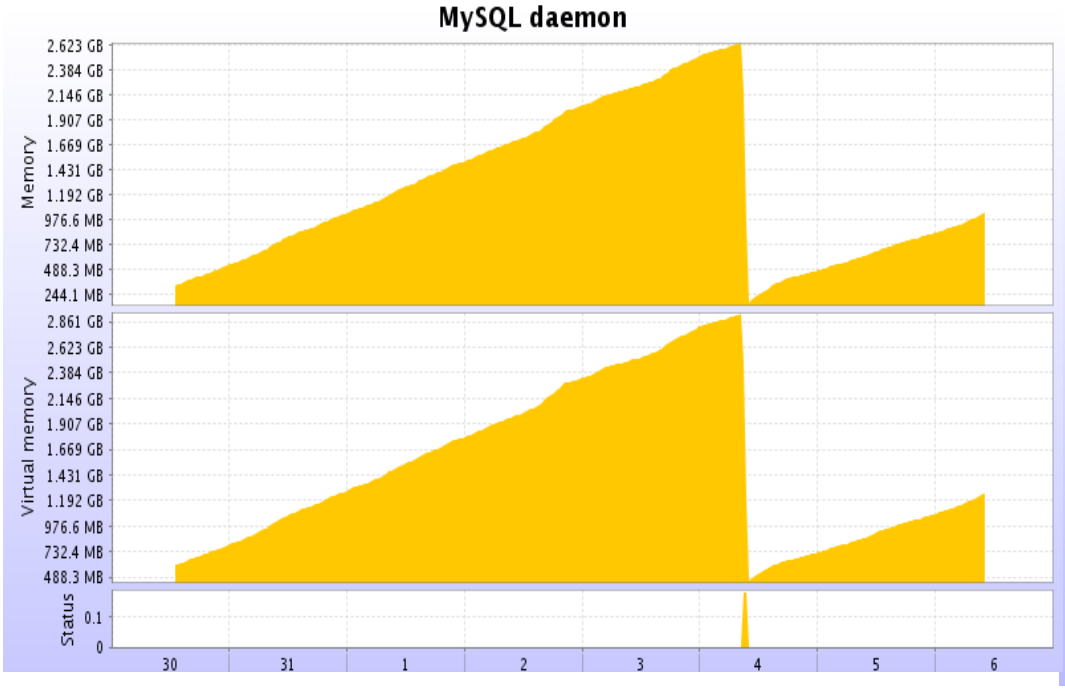

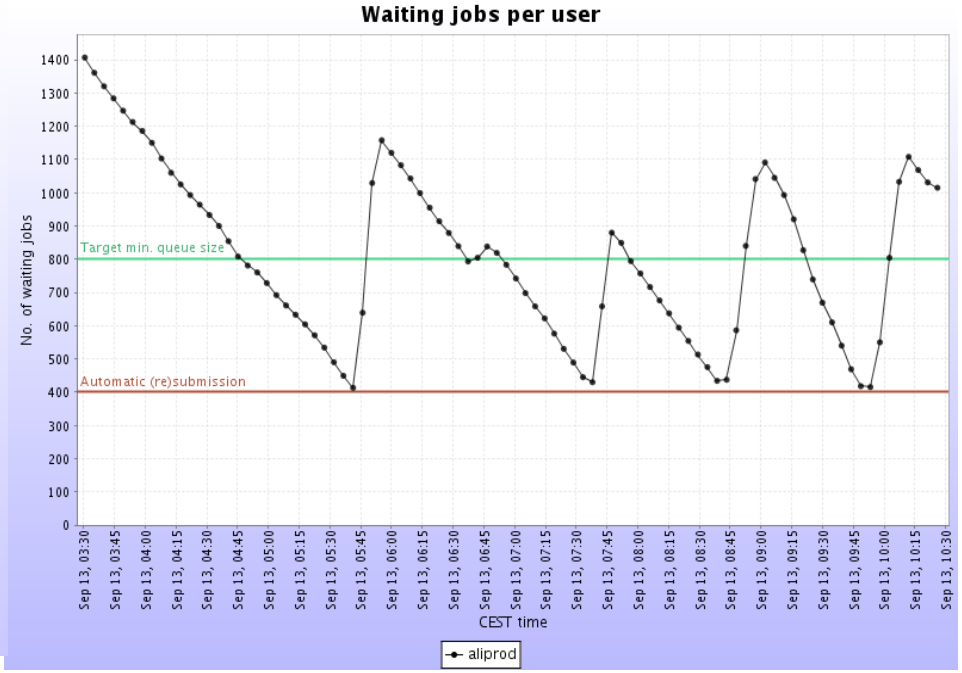

MySQL daemon is automatically restarted when it runs out of memory
Trigger: threshold on VSZ memory usage

ALICE Production jobs queue is kept full by the automatic submission
Trigger: threshold on the number of *aliprod* waiting jobs

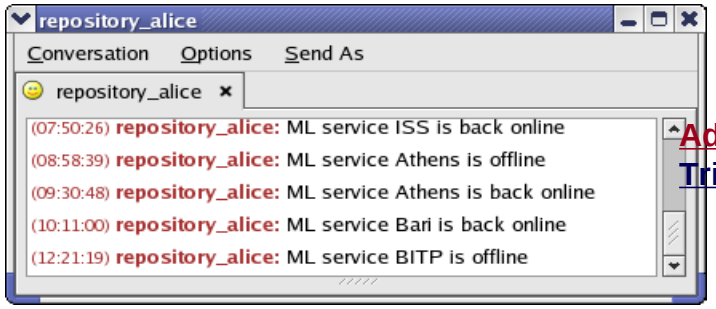

Administrators are kept up-to-date on the services' status
Trigger: presence/absence of monitored information

Automatic actions in ALICE

- ALICE is using the monitoring information to automatically:
  - resubmit error jobs until a target completion percentage is reached,
  - submit new jobs when necessary (watching the task queue size for each service account)
    - production jobs,
    - RAW data reconstruction jobs, for each pass,
  - restart site services, whenever tests of VoBox services fail but the central services are OK,
  - send email notifications / add chart annotations when a problem was not solved by a restart
  - dynamically modify the DNS aliases of central services for an efficient load-balancing.
- Most of the actions are defined by few lines configuration files.

Monitoring USLHCNet Topology

Topology & Status & Peering

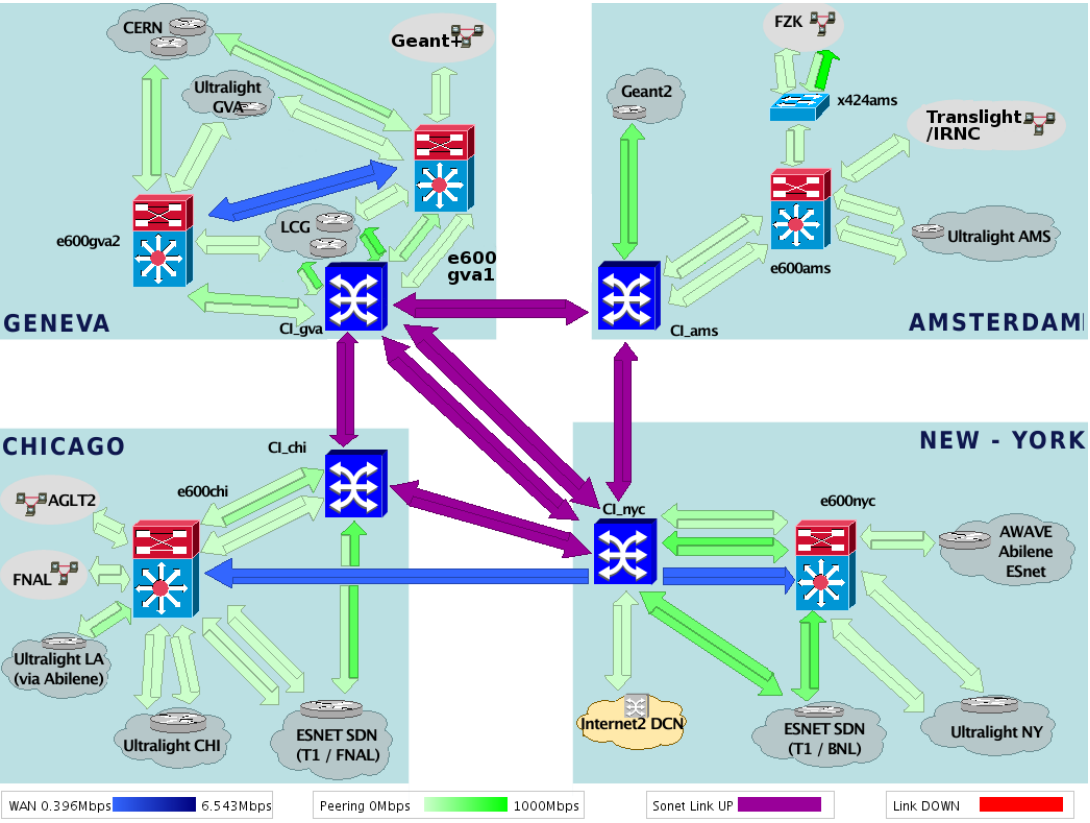

Real Time Topology for L2 Circuits

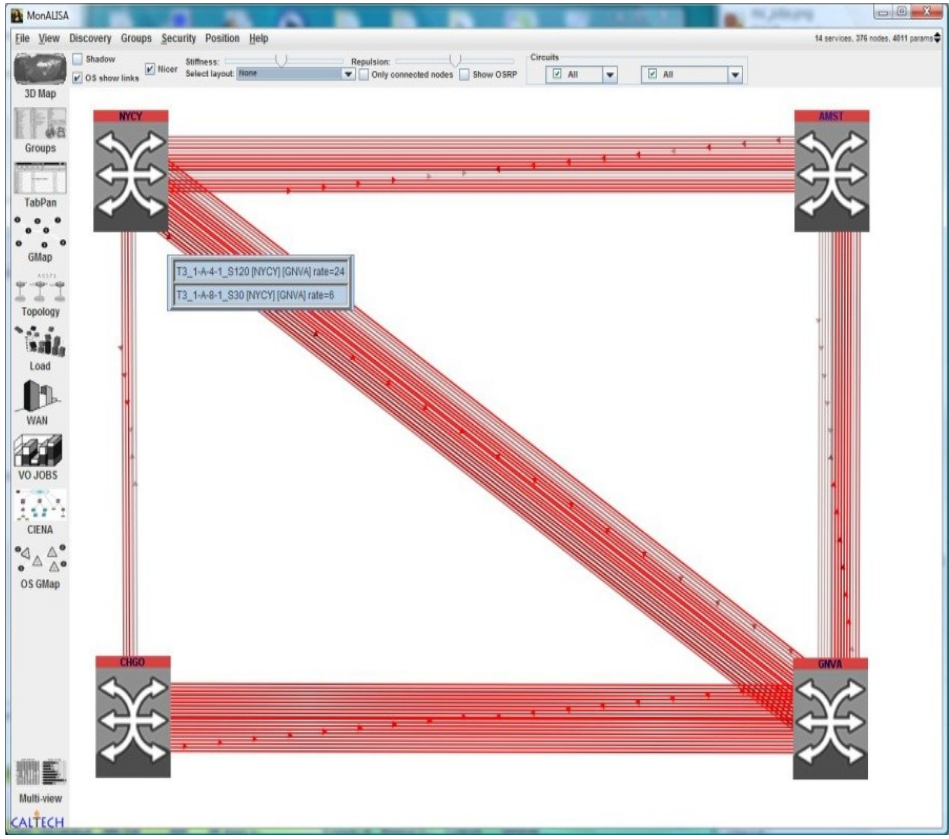

Monitoring Links Availability
Very Reliable Information

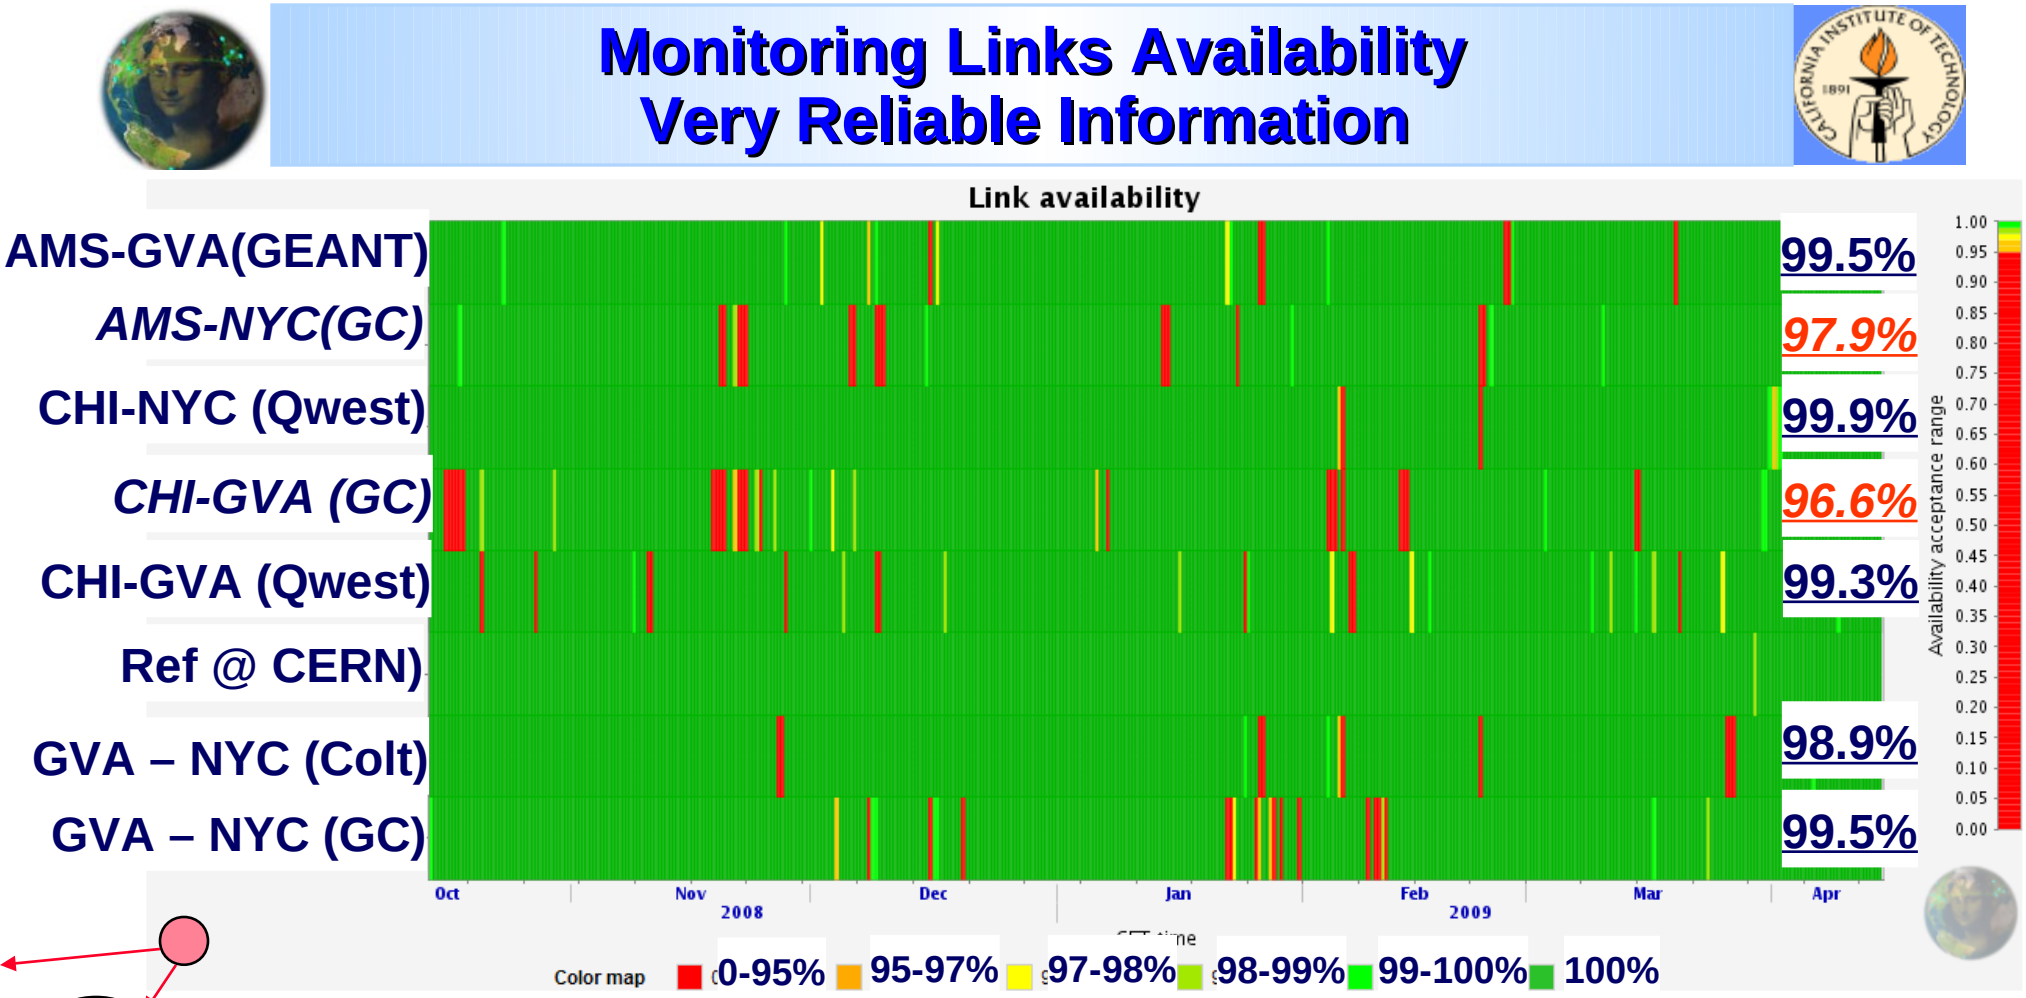

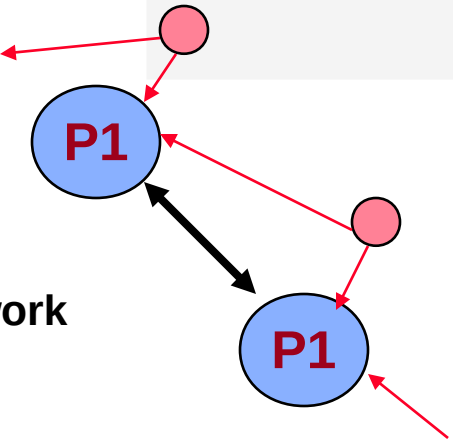

ork

Statistics

Link name	Data		Monitoring		Link
	Starts	Ends	Availability(%)	Gaps	Availability(%)
AMS-GVA (Geant)	14 Oct 2008 12:22	14 Apr 2009 12:21	99.100%	4m 30s	99.53%
AMS-NY (GlobalCrossing)	14 Oct 2008 12:22	14 Apr 2009 12:21	100%	-	97.87%
CHI-NY (Qwest)	14 Oct 2008 12:22	14 Apr 2009 12:21	99.93%	2:59	99.90%
CHI-NY (GlobalCrossing)	14 Oct 2008 12:22	14 Apr 2009 12:21	99.62%	16:40	96.59%
CHI-GVA (Qwest)	14 Oct 2008 12:22	14 Apr 2009 12:21	99.100%	4m 31s	99.29%
GVA1-GVA2 (USLHCNet)	14 Oct 2008 12:22	14 Apr 2009 12:21	100%	-	99.100%
GVA-NY (Colt)	14 Oct 2008 12:22	14 Apr 2009 12:21	100%	-	98.91%
GVA-NY (Geant & GlobalCrossing)	14 Oct 2008 12:22	14 Apr 2009 12:21	99.99%	13m 28s	99.47%

USLHCnet: Accounting for Integrated Traffic

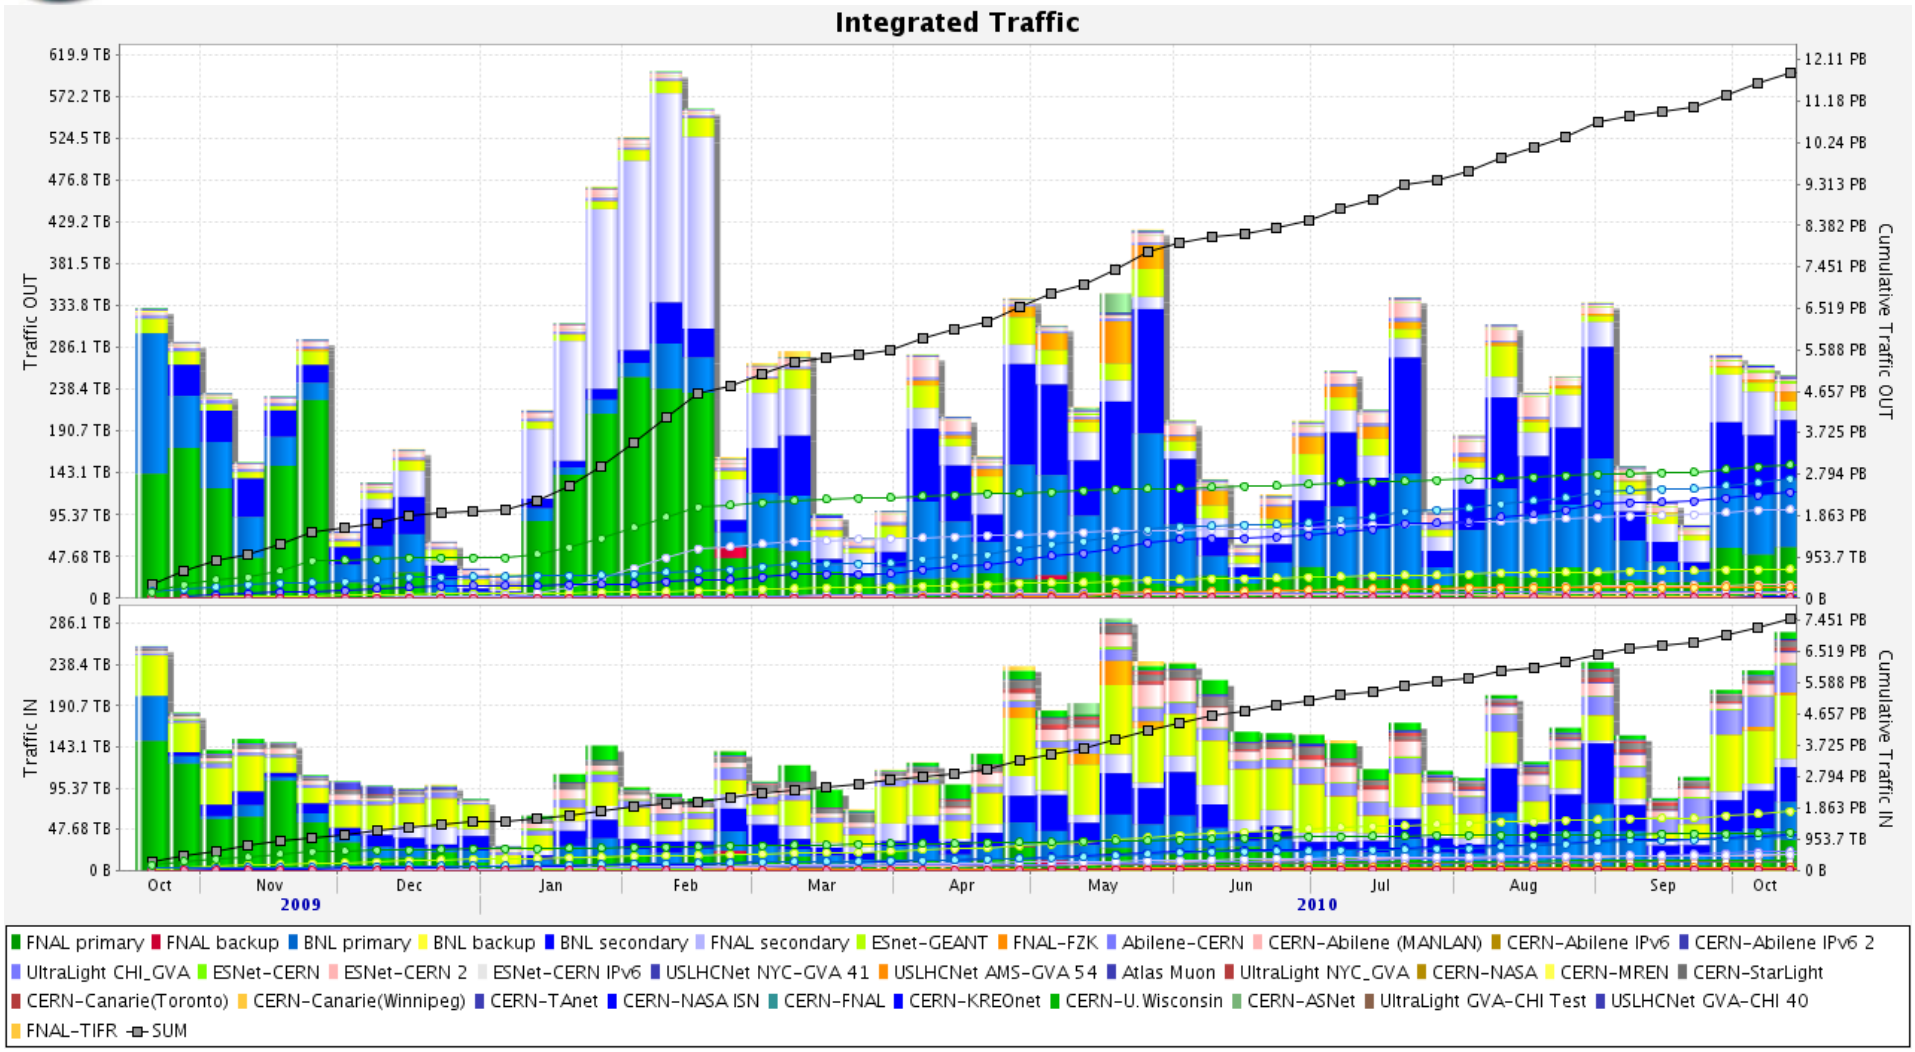

ALARMS and Automatic notifications for USLHCnet

CIENA Alarms for USLHCNet

Date (GMT)	Site	Node IP	Alarm	Remarks
last week				Filter
18.10.2010 12:09	AMS_USLHCNET_CDS	192.65.197.40	"1-A-3-1-1,GIGE:CR,LOS,SA,2010-10-18,10:09:00,,:\"Loss of signal\","	
18.10.2010 12:06	AMS_USLHCNET_CDS	192.65.197.40	"1-A-3-1-1,GIGE:CR,LOS,SA,2010-10-18,10:05:33,,:\"Loss of signal\","	
17.10.2010 03:21	CHI_USLHCNET_CDS	192.65.196.107	"TimingInput_LINE_1,REF:MN,SYNCCLK,NSA,2010-10-17,01:20:56,,:\"Frequency offset ...	
17.10.2010 03:20	CHI_USLHCNET_CDS	192.65.196.107	"TimingInput_LINE_1,REF:MN,SYNCCLK,NSA,2010-10-17,01:20:56,,:\"Frequency offset ...	
17.10.2010 03:19	AMS_USLHCNET_CDS	192.65.197.40	"1-A-2-2,OC192:MN,RFI-L,NSA,2010-10-17,01:19:16,,:\"Line RFI\","	
17.10.2010 03:10	AMS_USLHCNET_CDS	192.65.197.40	"1-A-2-2,OC192:MN,AIS-L,NSA,2010-10-17,01:09:12,,:\"Line AIS\","	
17.10.2010 03:09	AMS_USLHCNET_CDS	192.65.197.40	"1-A-2-2,OC192:MN,RF	
17.10.2010 03:09	GVA_USLHCNET_CDS	192.65.196.172	"1-A-8-1,OC192:MN,AI	
17.10.2010 03:06	GVA_USLHCNET_CDS	192.65.196.172	"gva-chi-S1-2,SNC:CR,	
17.10.2010 03:06	GVA_USLHCNET_CDS	192.65.196.172	"gva-chi-S1-3,SNC:CR,	
17.10.2010 03:06	GVA_USLHCNET_CDS	192.65.196.172	"gva-chi-S1-6,SNC:CR,	
17.10.2010 03:06	GVA_USLHCNET_CDS	192.65.196.172	"gva-chi-S1-7,SNC:CR,	
17.10.2010 03:06	GVA_USLHCNET_CDS	192.65.196.172	"gva-nyc-3513-9,SNC:C	
17.10.2010 03:06	GVA_USLHCNET_CDS	192.65.196.172	"gva-nyc-3524-6,SNC:C	
17.10.2010 03:06	GVA_USLHCNET_CDS	192.65.196.172	"gva-nyc-S1-1,SNC:CR,	
17.10.2010 03:06	GVA_USLHCNET_CDS	192.65.196.172	"gva-nyc-S1-4,SNC:CR,	
17.10.2010 03:06	GVA_USLHCNET_CDS	192.65.196.172	"gva-chi-3500-4,SNC:C	
17.10.2010 03:06	GVA_USLHCNET_CDS	192.65.196.172	"gva-chi-3500-6,SNC:C	
17.10.2010 03:06	GVA_USLHCNET_CDS	192.65.196.172	"gva-chi-3506-5,SNC:C	
17.10.2010 03:06	GVA_USLHCNET_CDS	192.65.196.172	"gva-chi-3506-6,SNC:C	

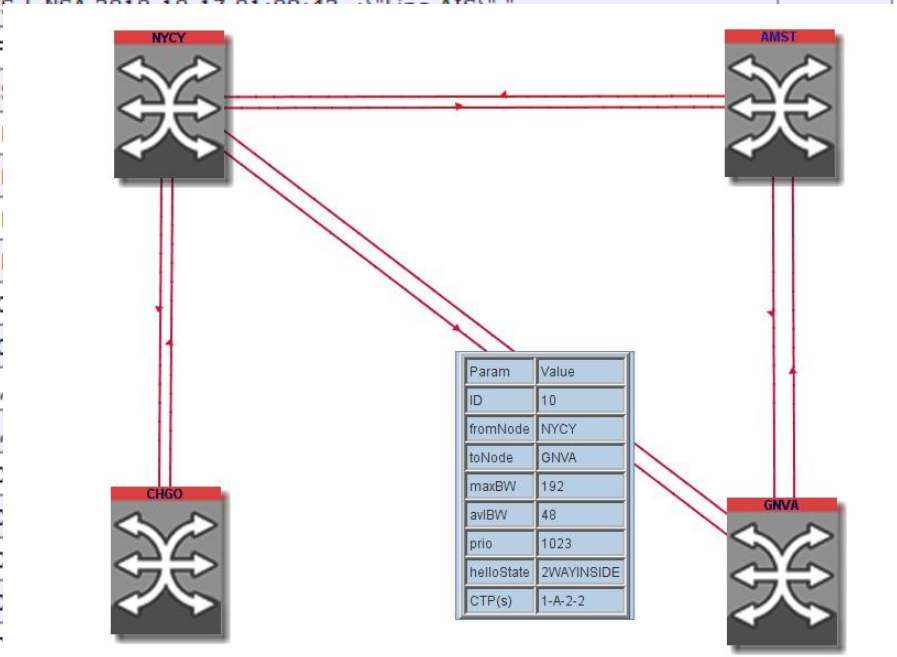

Monitoring Network Topology (L3), Latency, Routers

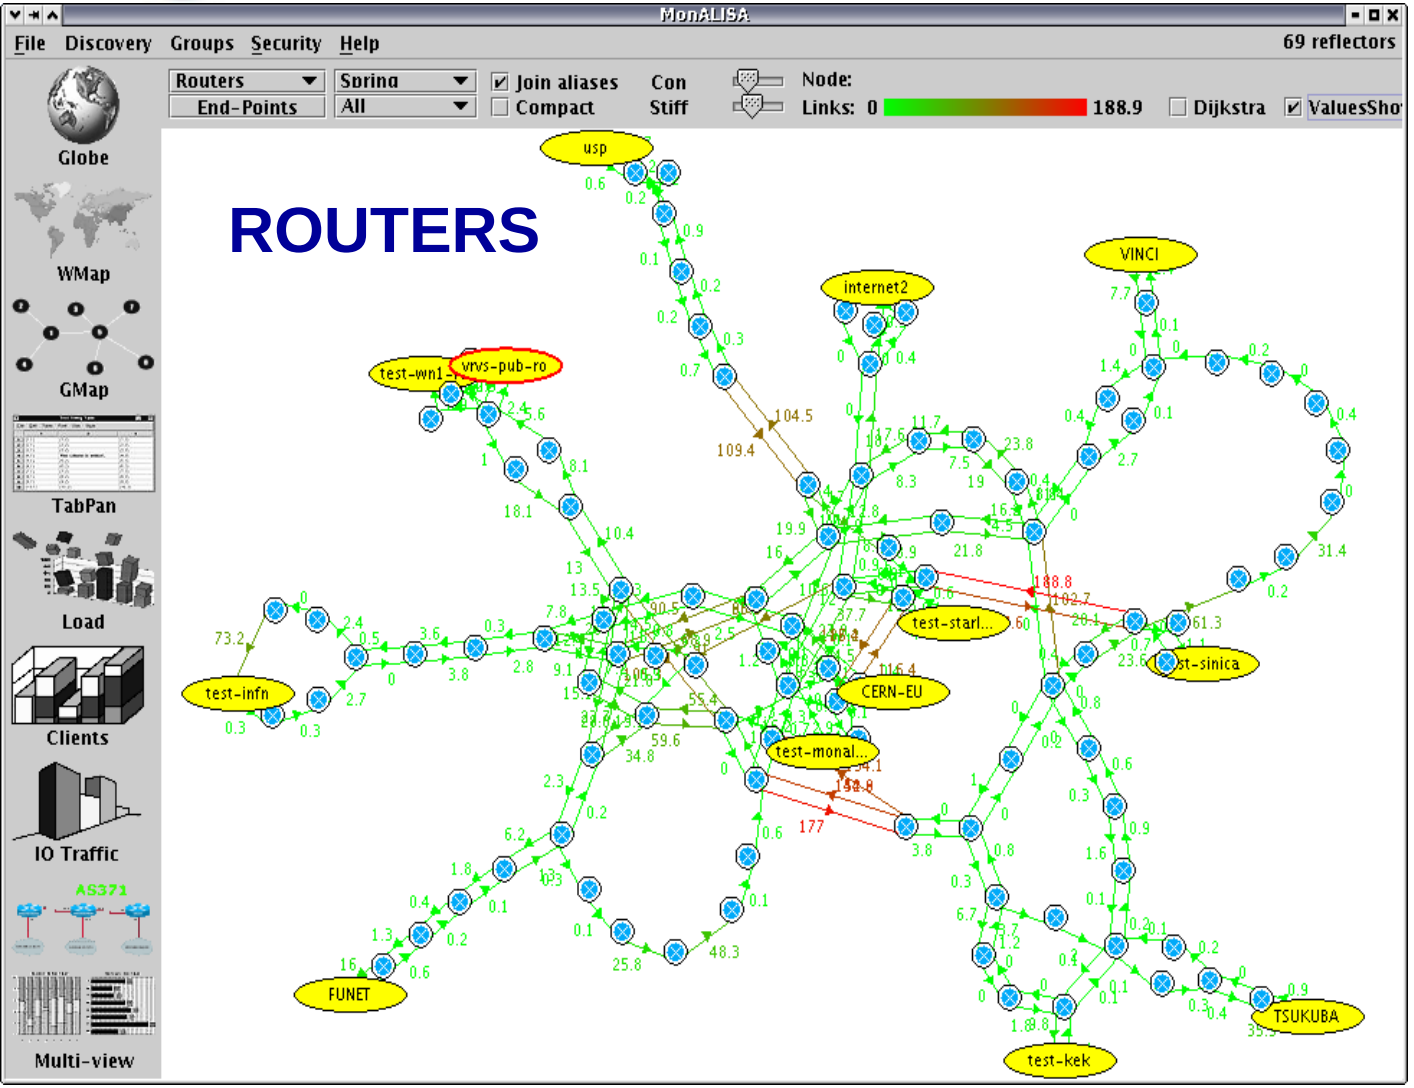

NETWORKS

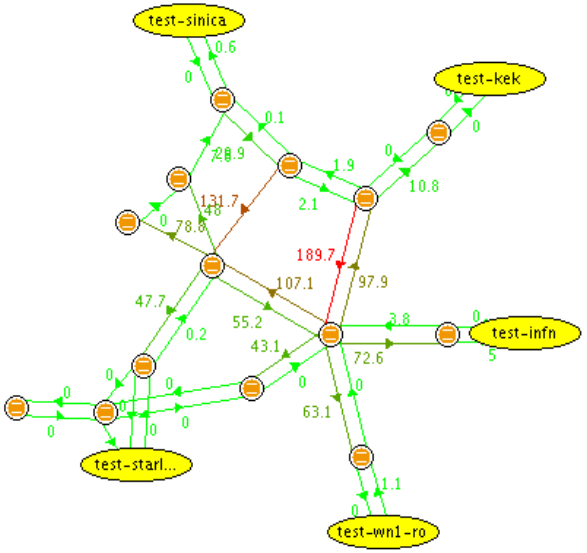

AS

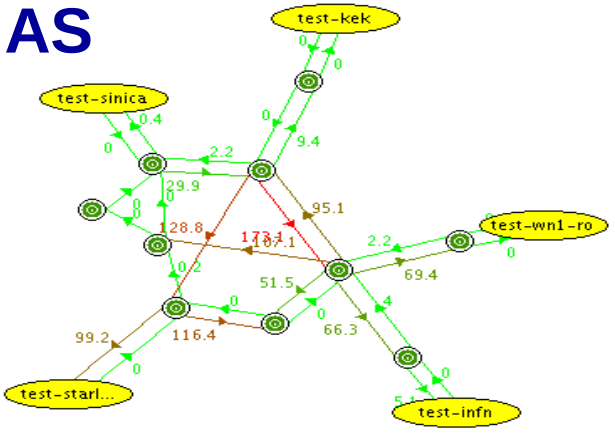

Real Time Topology Discovery & Display

EVO : Real-Time monitoring for Reflectors and the quality of all possible connections

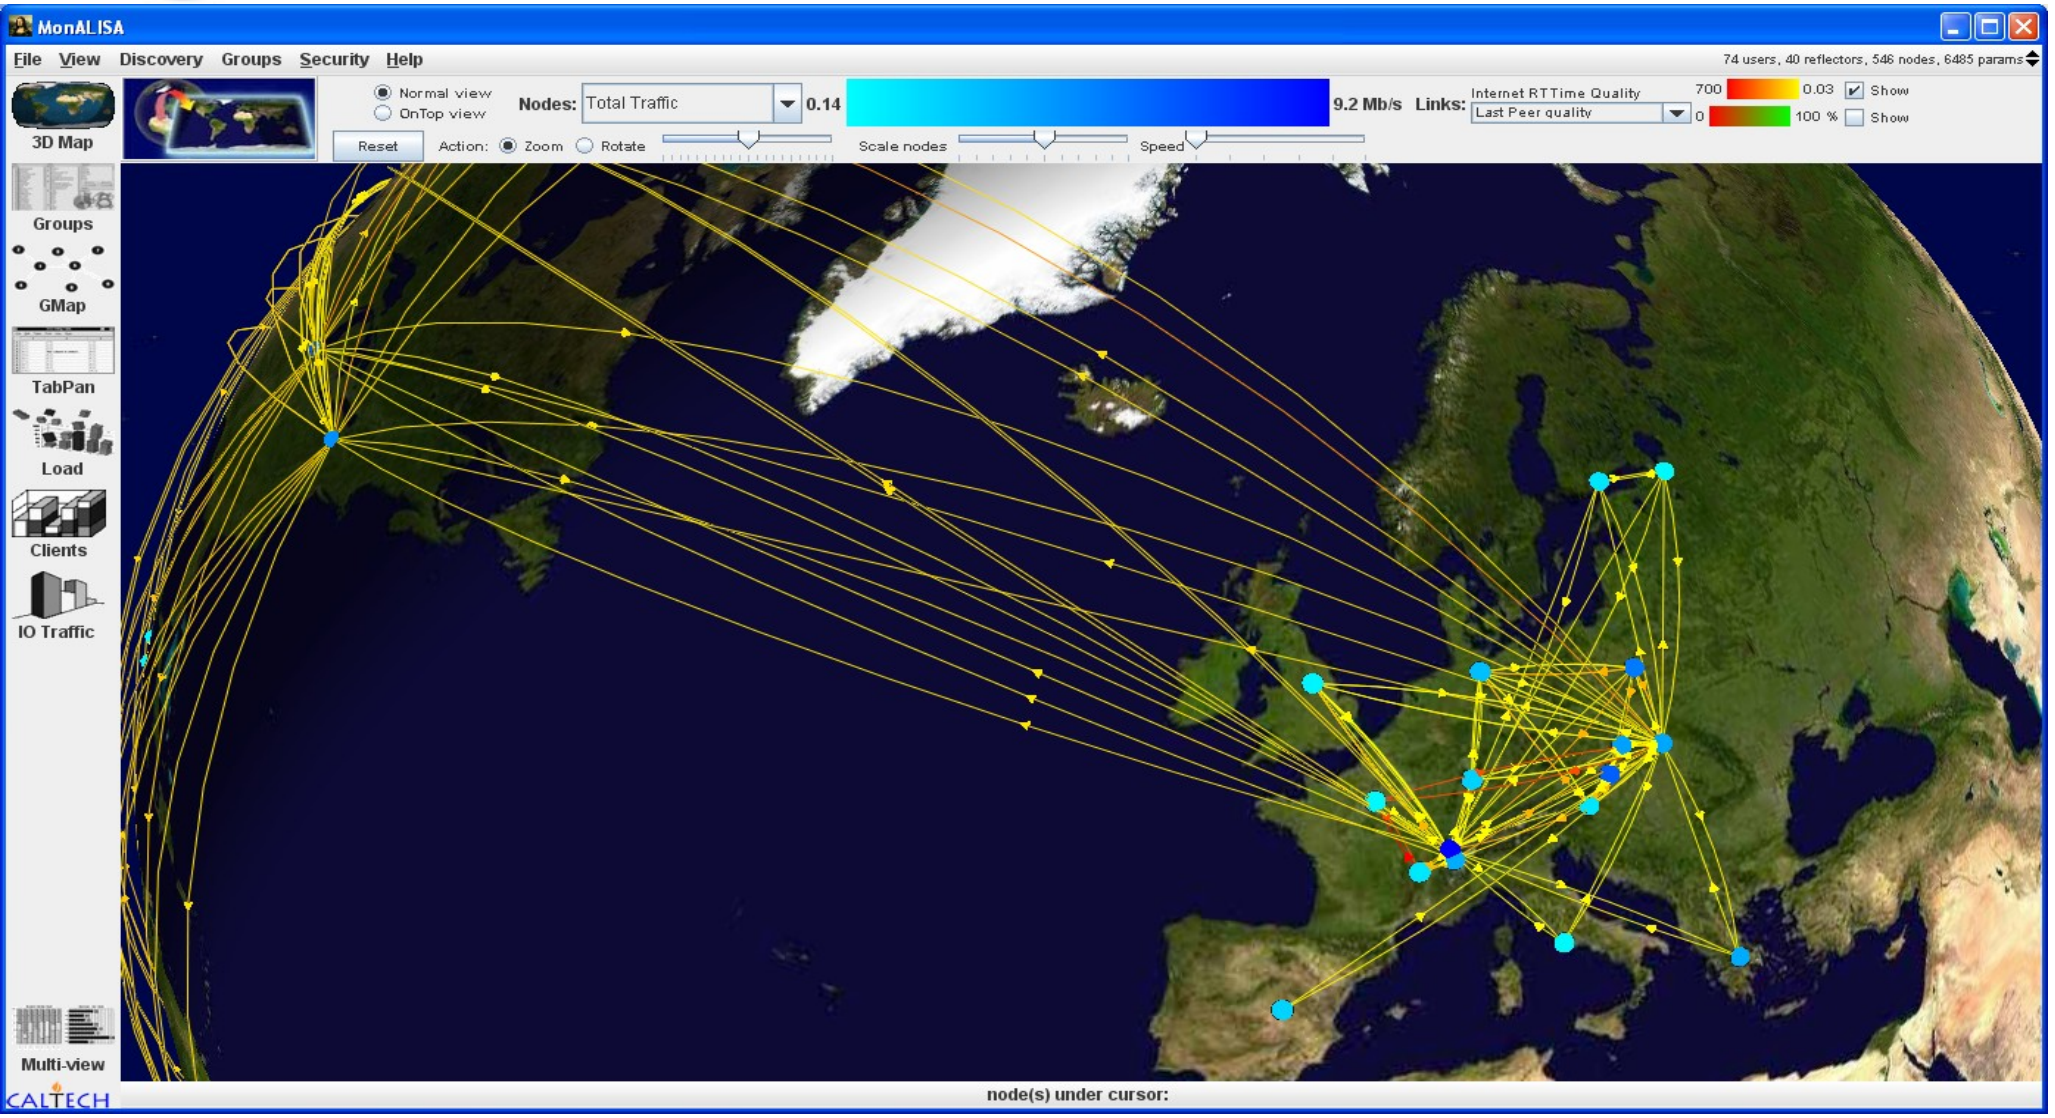

EVO: Creating a Dynamic, Global, Minimum Spanning Tree to optimize the connectivity

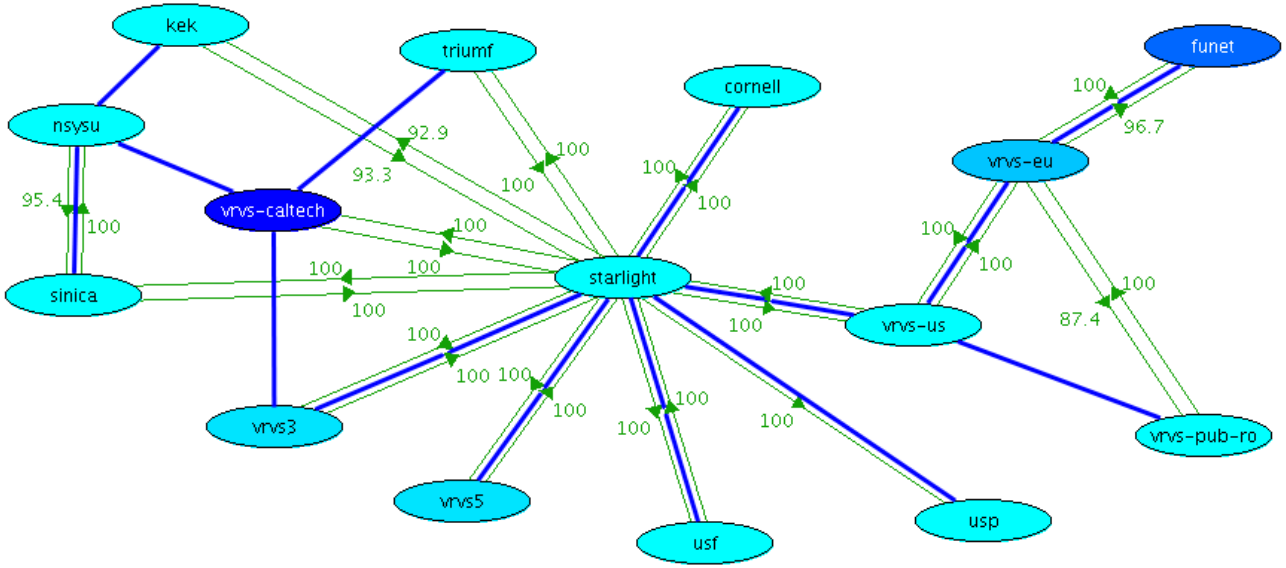

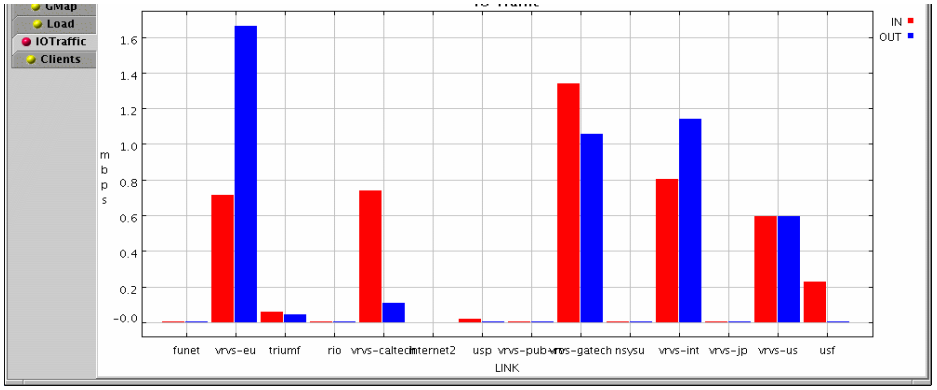

A weighted connected graph $G = (V,E)$ with n vertices and m edges. The quality of connectivity between any two reflectors is measured every second. Building in near real time a minimum-spanning tree *with addition constrains*

$$w(T) = \sum_{(v,u)\in T} w((v,u))$$

Resilient Overlay Network that optimize real-time communication

Dynamic MST to optimize the Connectivity for Reflectors

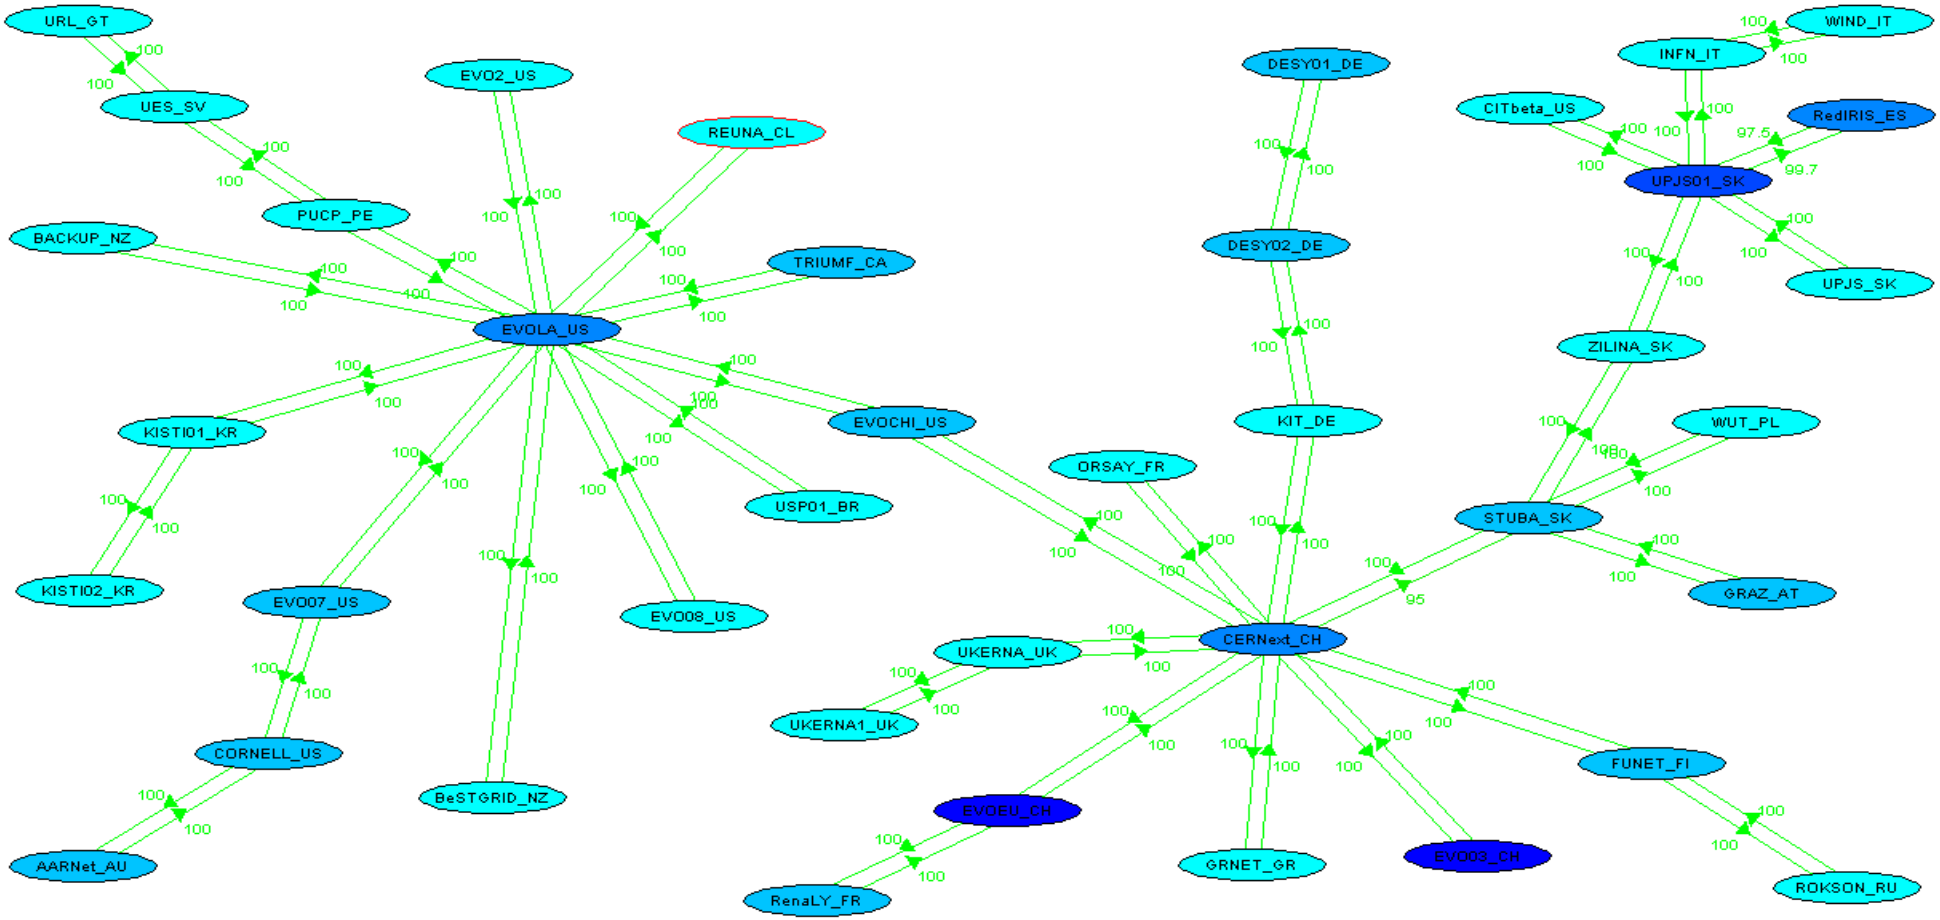

Frequent measurements of RTT, jitter, traffic and lost packages
The MST is recreated in ~ 1 S case on communication problems.

EVO: Optimize how clients connect to the system for best performance and load balancing

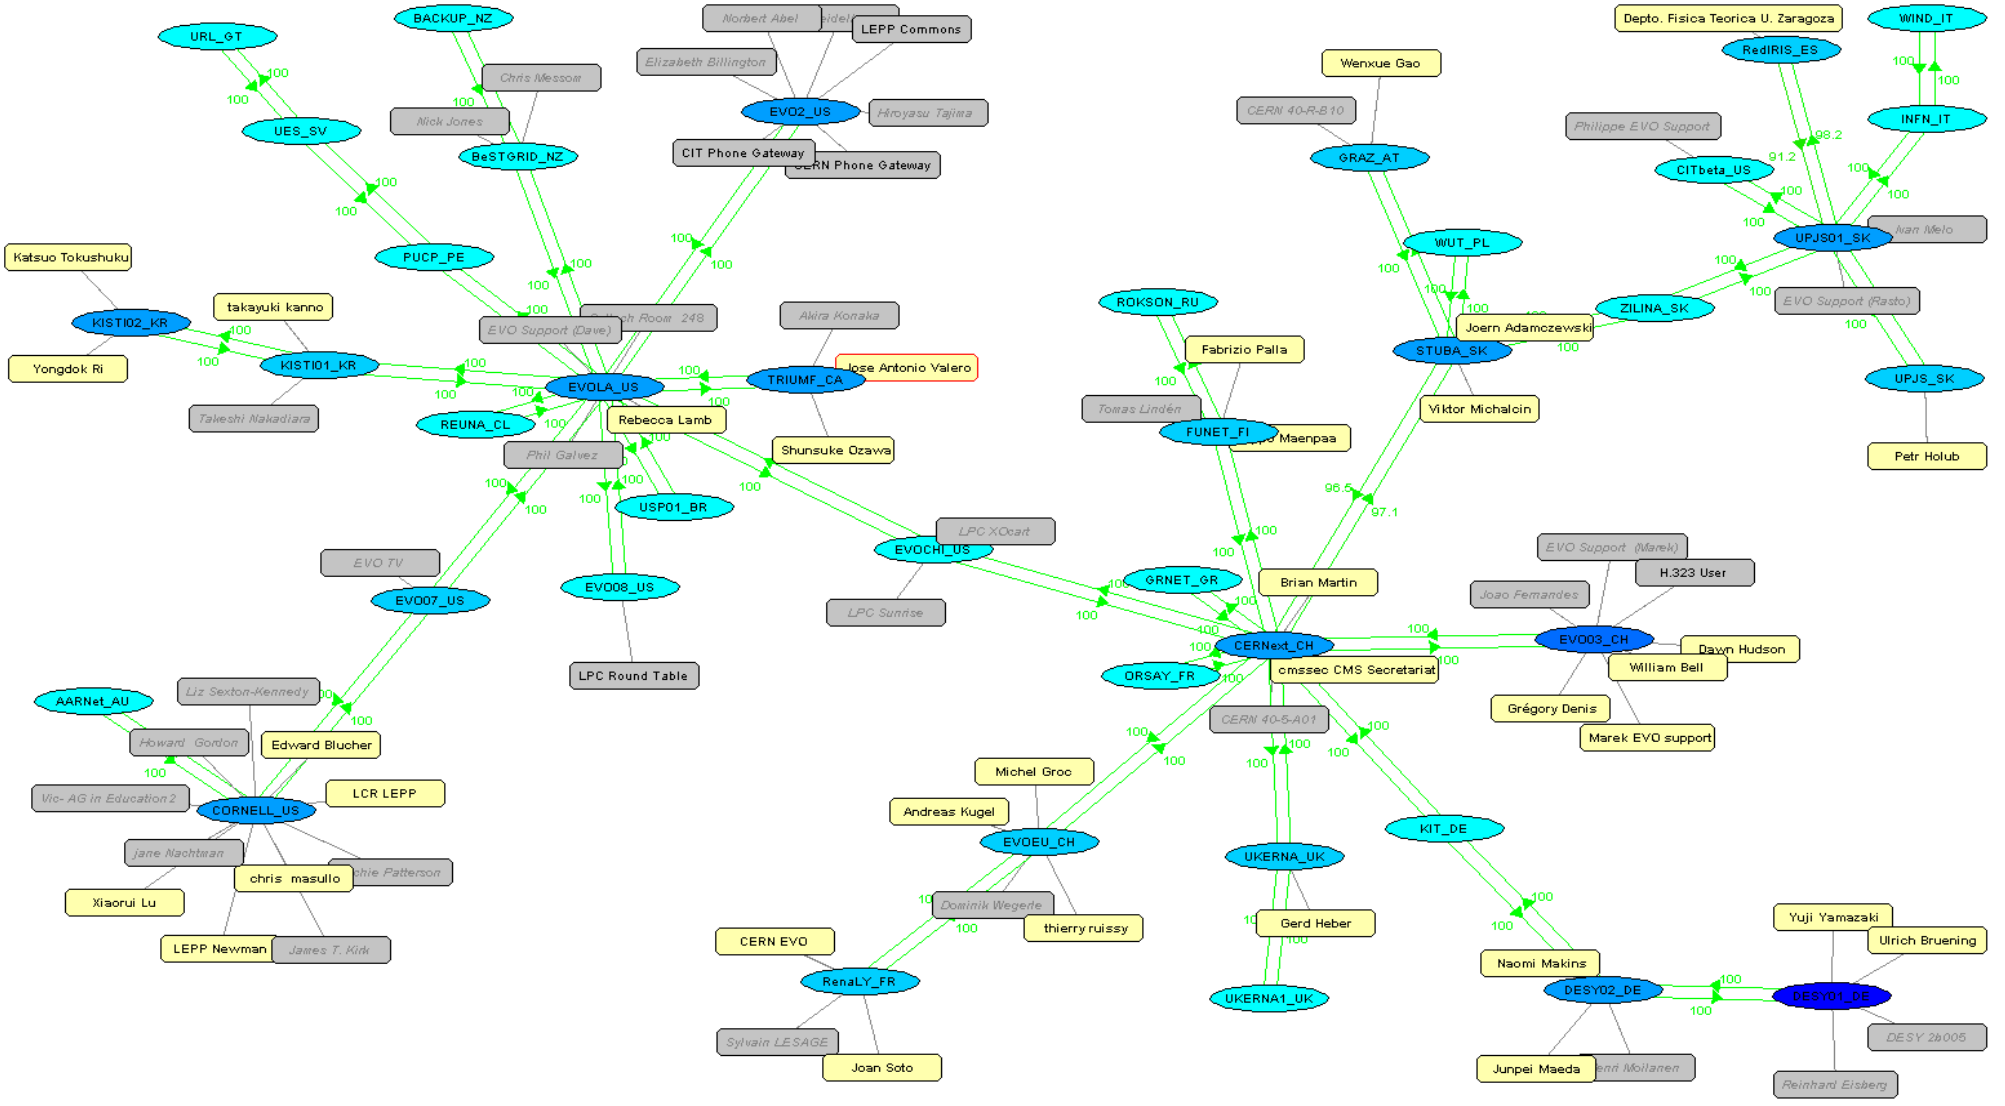

Monitoring the Topology and Optical Power on Fibers for Optical Circuits

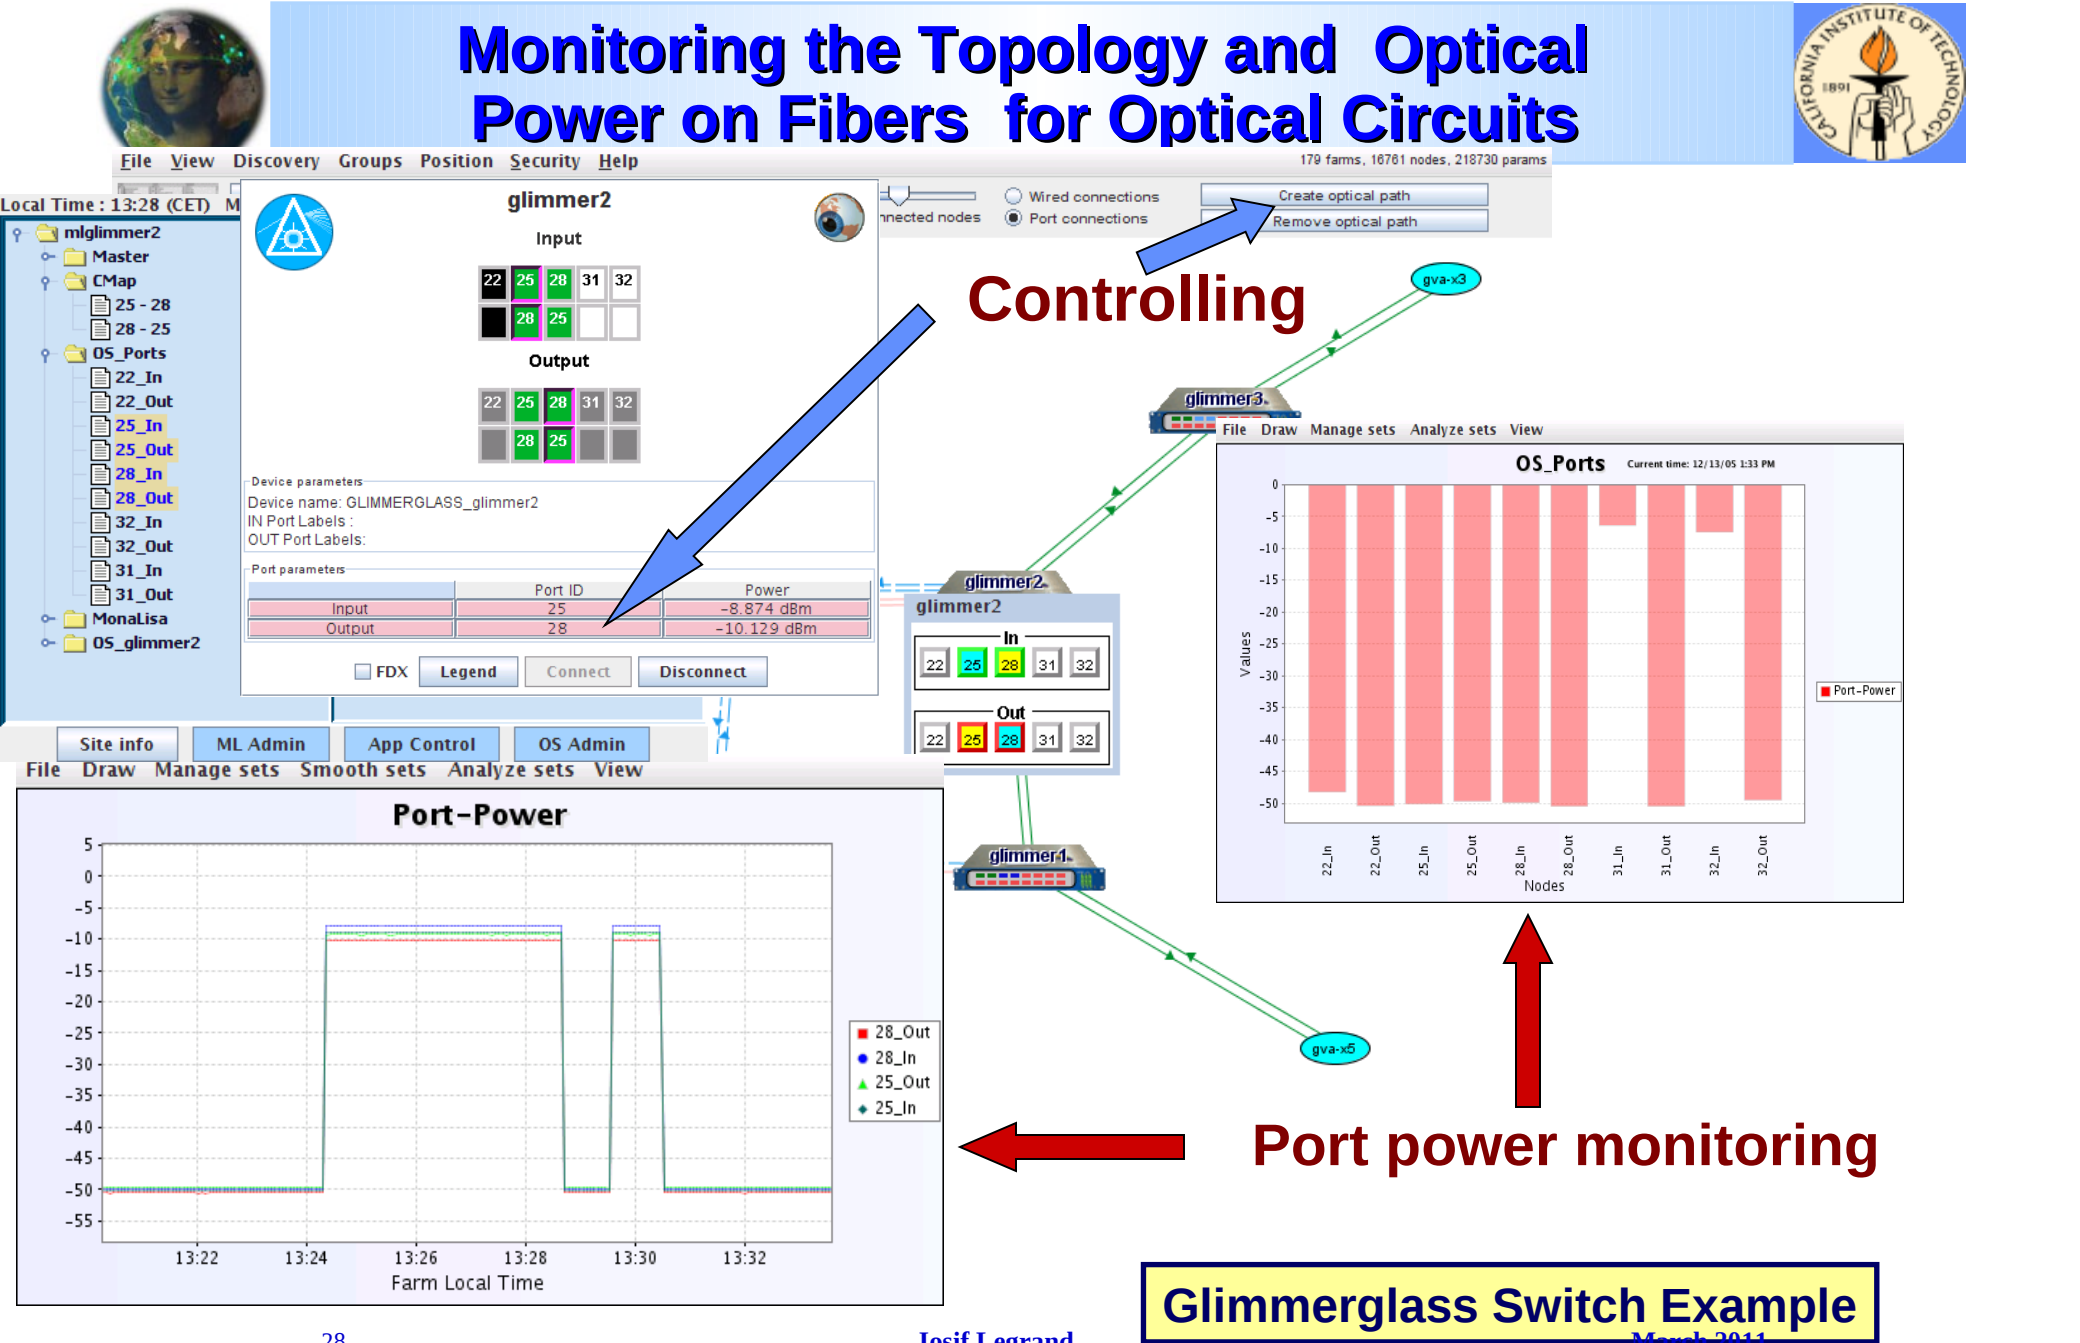

Glimmerglass Switch Example

"On-Demand", End to End Optical Path Allocation

CREATES AN END TO END PATH < 1s

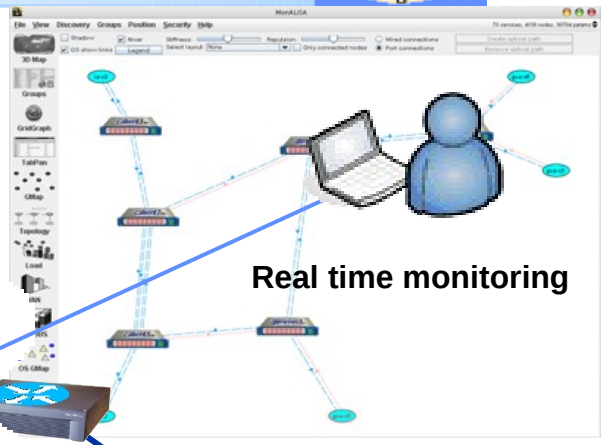

Internet

MonALISA Distributed Service System

Regular IP path

APPLICATION

A

Control

Monitor

MonALISA Service

OS Agent

DATA

B

LISA Agent

TL1

Optical Switch

Active light path

Detects errors and automatically recreate the path in less than the TCP timeout

LISA AGENT

LISA sets up
- Network Interfaces
- TCP stack
- Kernel parameters
- Routes

LISA ↑ APPLICATION
"use eth1.2, …"

Controlling Optical Planes
Automatic Path Recovery

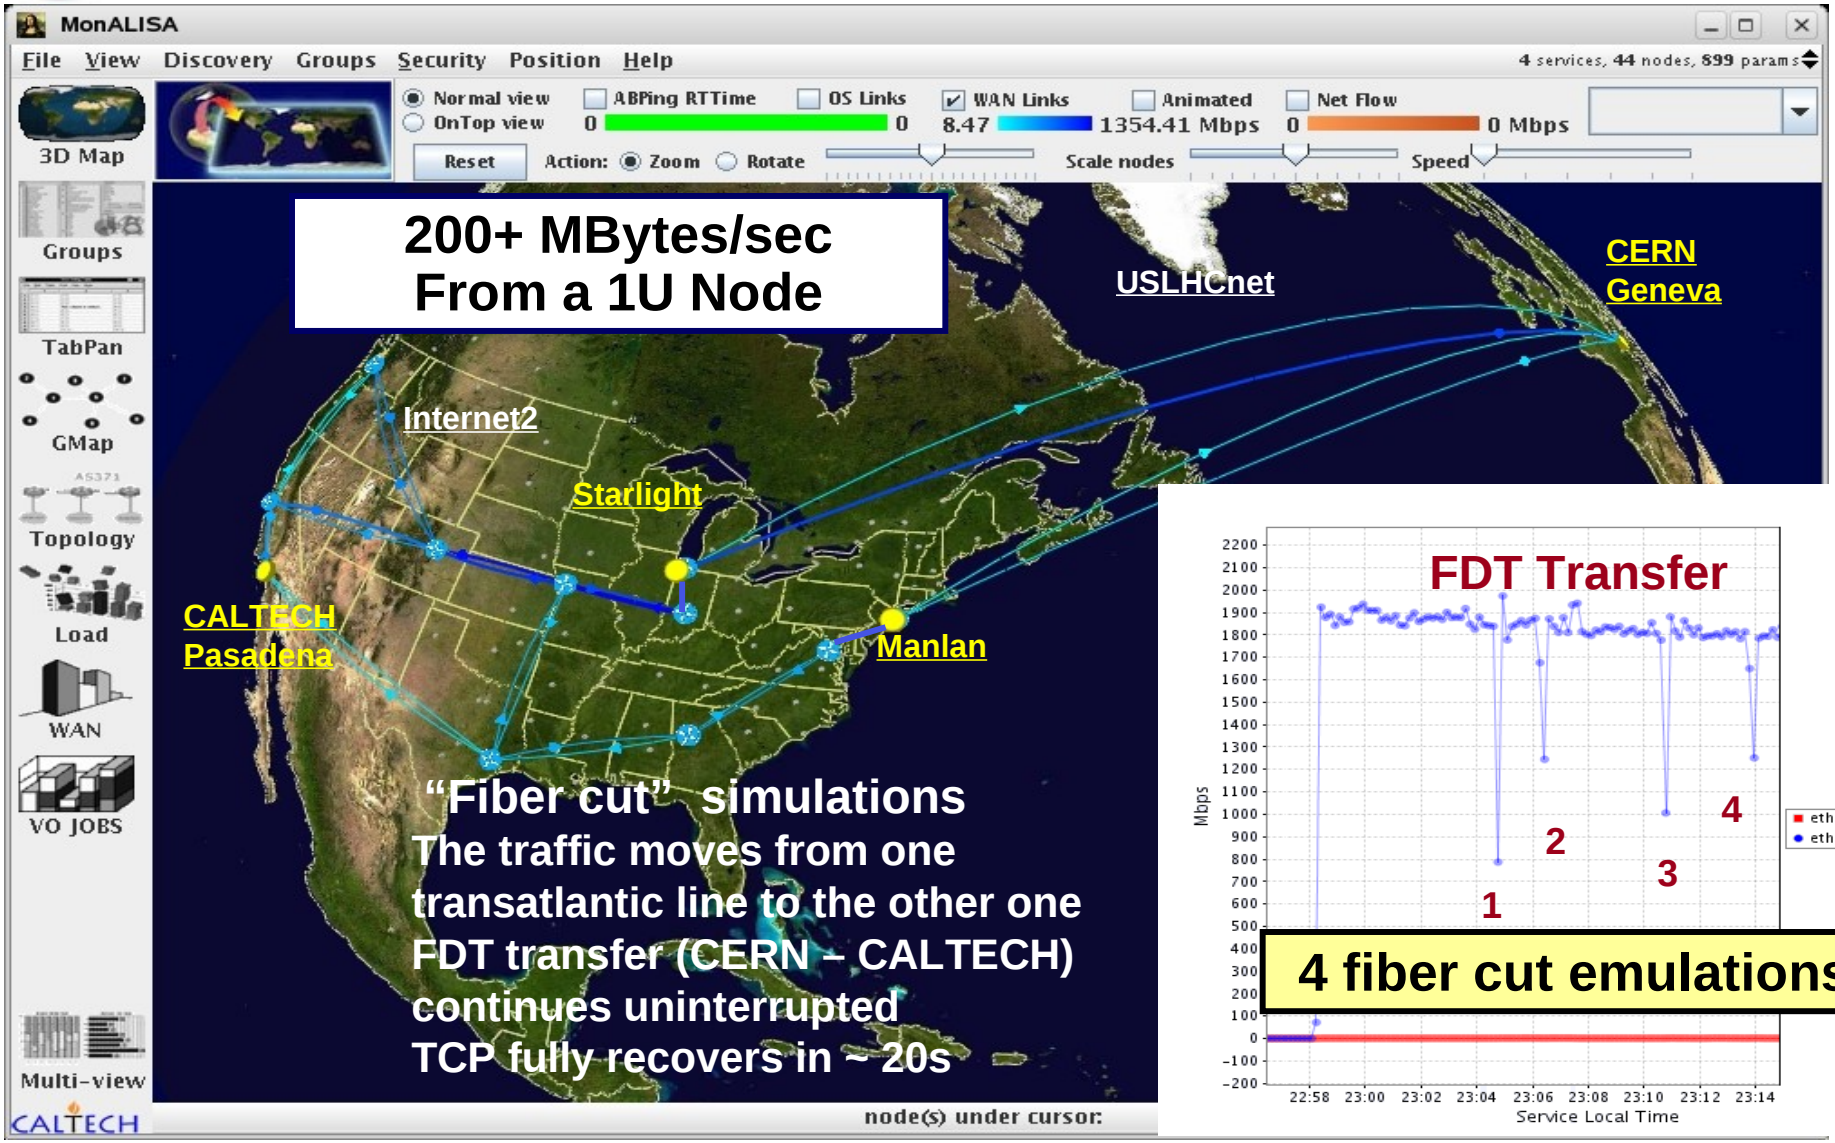

MonALISA collects any type of monitoring information in distributed systems

The MonALISA package includes:

- Local host monitoring (CPU, memory, network traffic , processes and sockets in each state, LM sensors, APC UPSs), log files tailing
- SNMP generic & specific modules
- Condor, PBS, LSF and SGE (accounting & host monitoring), Ganglia
- Ping, tracepath, traceroute, pathload and other network-related measurements
- TL1, Network devices, Ciena, Optical switches
- Calling external applications/scripts that return as output the values
- XDR-formatted UDP messages (such as ApMon).

New modules can be easly added by implementing a simple Java interface.

Filters can be used to generate new aggregate data.

The Service can also react to the monitoring data it receives (actions alarms).

MonALISA can run code as distributed agents for global optimization

- Used by Evo to maintain the tree of connections between reflectors
- On demand end to end optical paths
- Controls distributed data transfers

MonALISA Summary

MonALISA Today
Running 24 X 7 at ~360 Sites

- Collecting ~ 2 million "persistent" parameters in real-time
- 80 million "volatile" parameters per day
- Update rate of ~25,000 parameter updates/sec
- Monitoring
  - 40,000 computers
  - > 100 WAN Links
  - > 8,000 complete end-to-end network path measurements
  - Tens of Thousands of Grid jobs running concurrently
- Controls jobs summation, different central services for the Grid, EVO topology, FDT …
- The MonALISA repository system serves

Major Communities

- ALICE
- CMS
- ATLAS
- PANDA
- EVO
- LGC RUSSIA
- OSG
- MXG
- RoEduNet
- USLHCNET
- ULTRALIGHT
- Enlightened

http://monalisa.caltech.edu