


WLCG Monitoring – some worked examples


GDB


CERN, 4th March 2008

James Casey




- What are the actual commodity solutions we can use?
- Messaging System
 - How does it perform ?
- How does the strategy affect existing systems?
- What new systems do/might we need ?





- Flexible architecture:
 - Deliver messages, either in point to point (queue)...
 - ... or multicast mode (topics)
 - Support Synchronous or Asynchronous communication.
- Reliable delivery of messages:
 - Provide reliability to the senders if required
 - Configurable persistency / Master-Slave.
- Highly Scalable:
 - Network of Brokers

- Mature open-source implementation of these ideas
 - Top-level Apache project
 - Commercial support available from IONA
- Easy to integrate
 - Multiple language + transport protocol support
- Good performance characteristics
 - See later ...
- Work done to integrate into our environment
 - RPMs, Quattor components + templates, LEMON alarms





Evaluated Parameters:

- 1) Number of Producers
- 2) Number of Consumers
- 3) Message Size
- 4) Message Number

Measurement of timestamps:

- 1) Message Sent
- 2) Message on Broker
- 3) Message Received

Results analysis:


- 1) Logs containing all information for each message
- 2) From logs, extract messages/second...
- 3) ... and messageLag

B) Connects + subscribes control;
 C) On control message, subscribes testTopic;
 D) On testTopic message, save message information to local file;

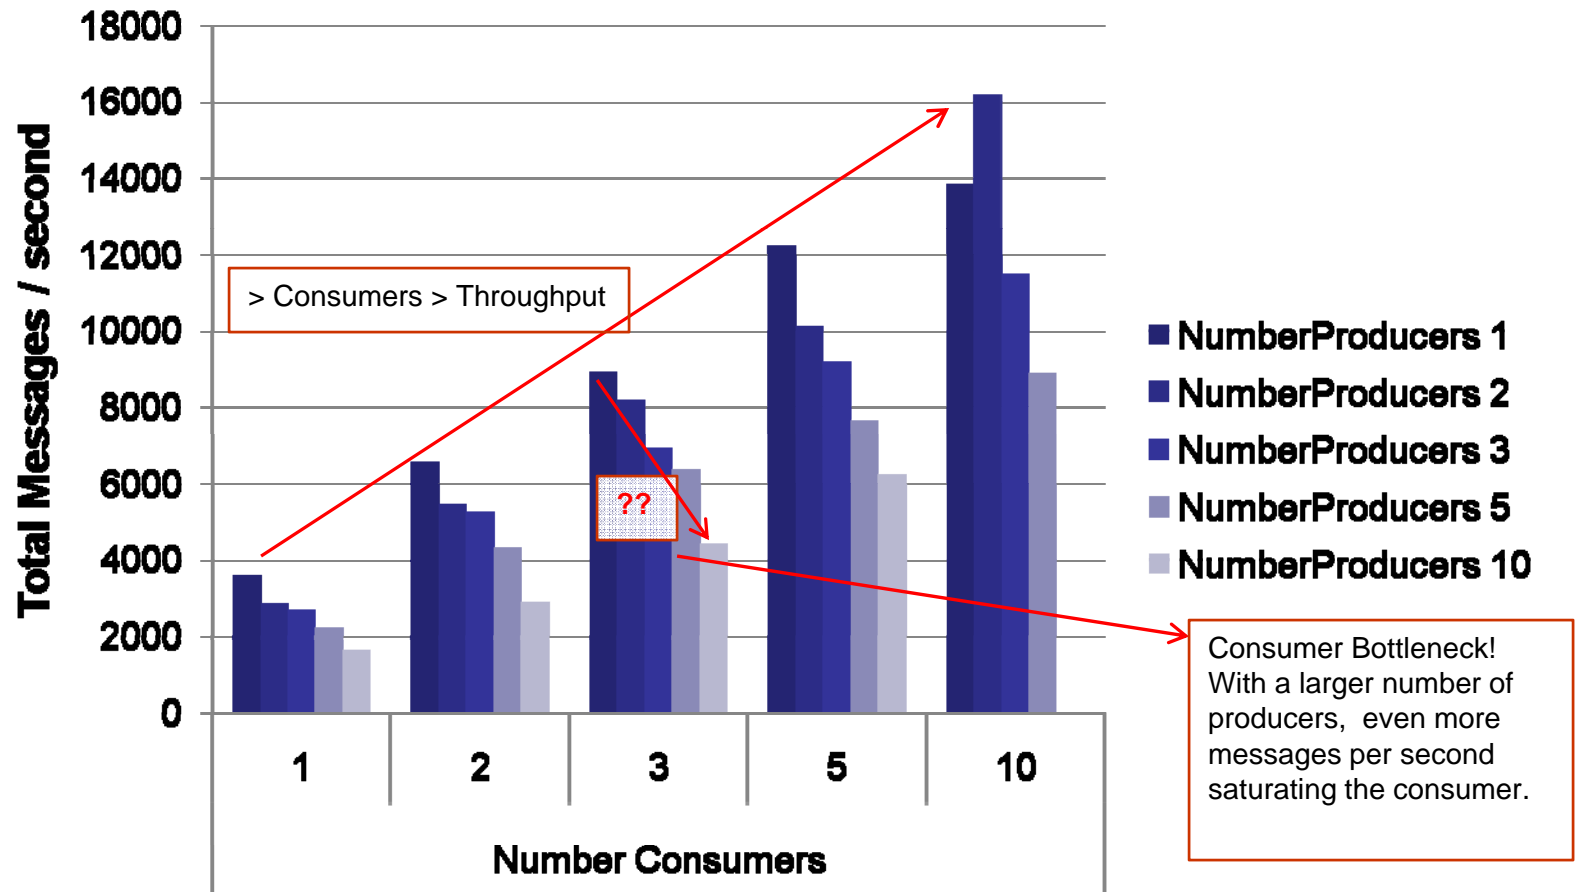
A) Manually start client;
 B) Connects + subscribes control;
 C) Sends control messages according to algorithm;

C) On control message, publishes messages to testTopic;
 D) On finish sending sends status control message;

- Results summary:
 - Running for 6 weeks with no crashes
 - 50 Million messages of various sizes (0 to 10 kB) forwarded to consumers
 - 12 Million incoming messages from producers
 - Up to 40 Producers and 80 Consumers connected at the same time
 - Stable under highly irregular test pattern:
 - Number of clients change
 - Frequent client process kills
 - Daily number of tests vary

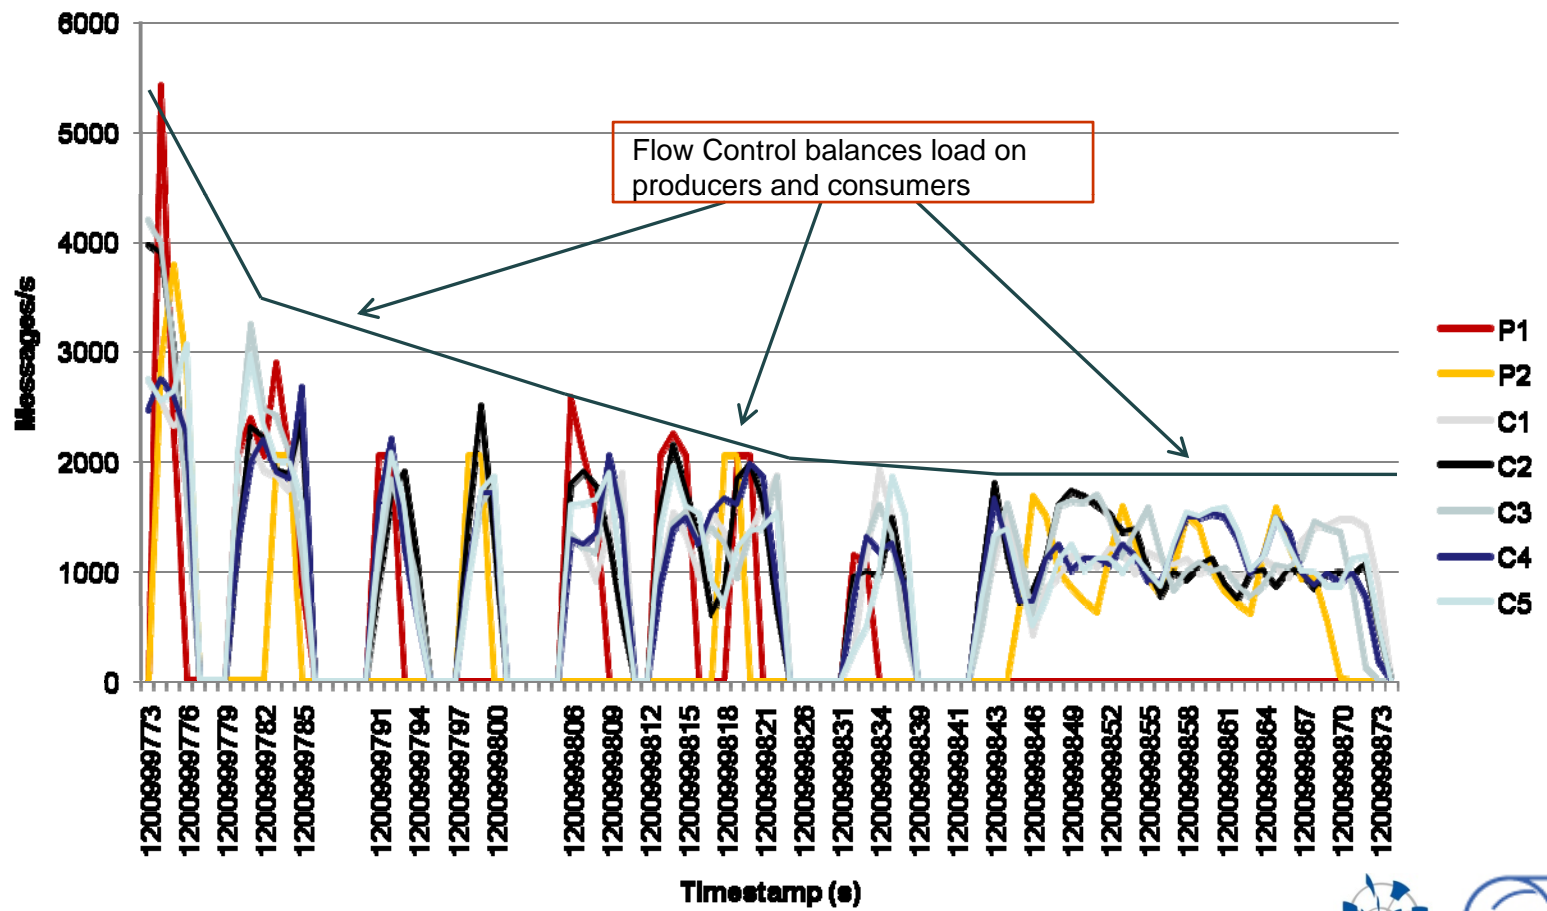



Total 100B message throughput




- Effective flow control

2Producers 5Consumers, [50k x 0B]





2 Producers, 5 Consumers
[(50k mes:0 B); (50k mes:100 B); (50k mes:1k B)]

- Currently testing failover strategies for high-availability
 - JDBC Based
 - Shared Filesystem
 - Pure Master-Slave
- Evaluation of performance characteristics and recovery operations needed
- This is an essential feature to have high-reliability service
 - First tests look promising

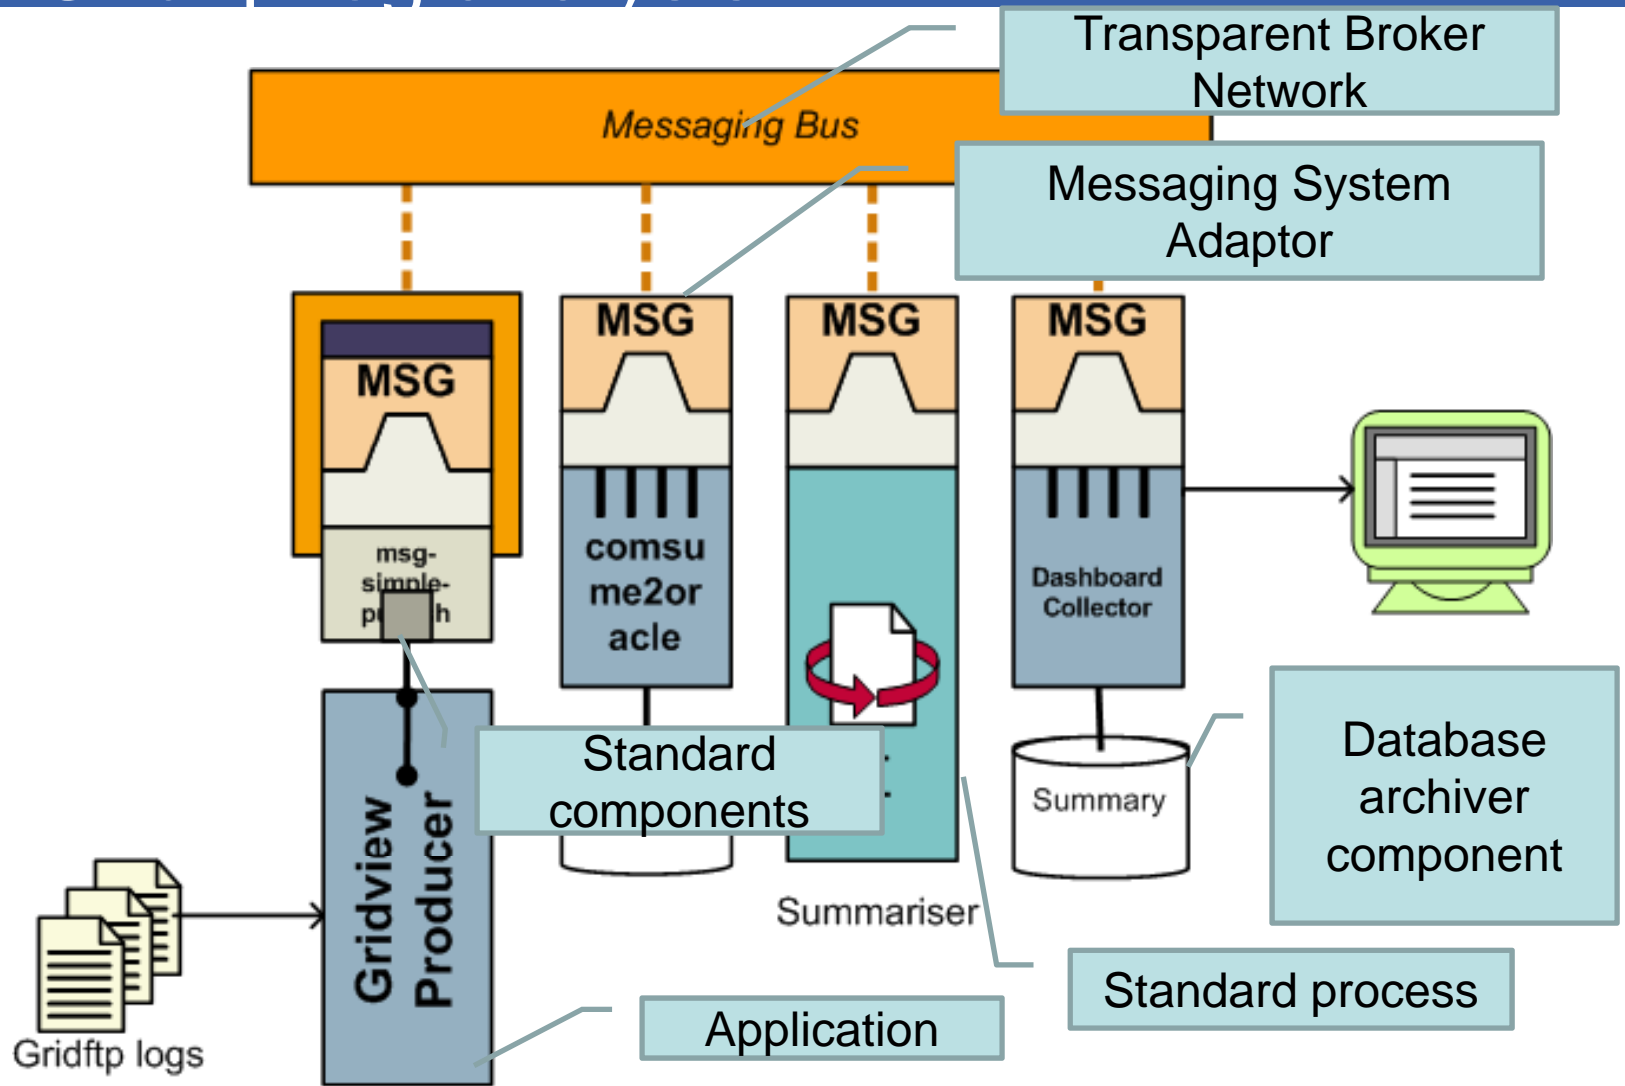


- We use it as an ‘integration bus’
 - Use when systems want to share information
 - E.g VO transfer systems publishing data rates to WLCG
- It’s another string to our bow
 - When the application model fits well, then use it
 - E.g. Async communications, broadcast messages
- Don’t force applications to use it
 - Have other solutions too
 - E.g “RESTful” web services a.la SAM Programmatic Interface




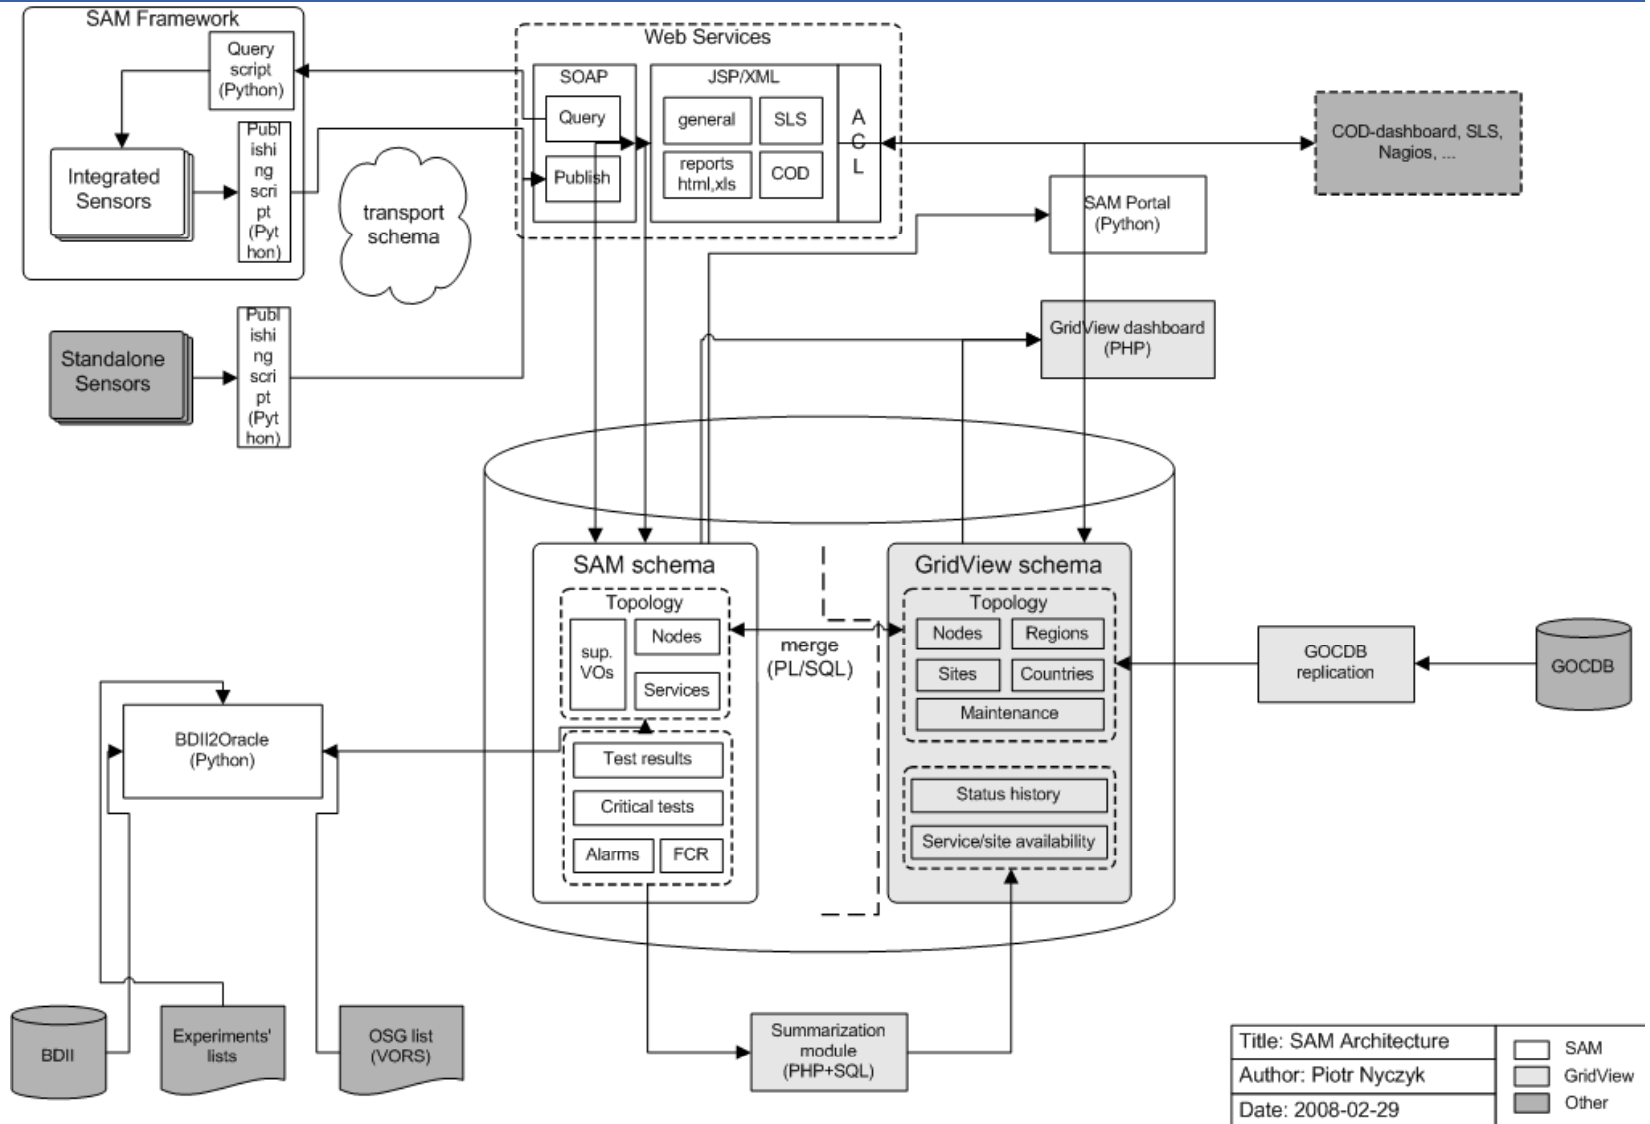


- The same patterns are repeated in many of the following examples:
 - Gather results at many points
 - Collect the raw results and store in a database
 - Perform some operation on the raw results
 - Summarisation, availability calculation, ...
 - Publish the summarised results to many clients
 - E.g. site monitoring, dashboards, ...
 - Store historical data in a database and visualize via web client
- We provide 'standard' components to make this plug'n'play for many workflows

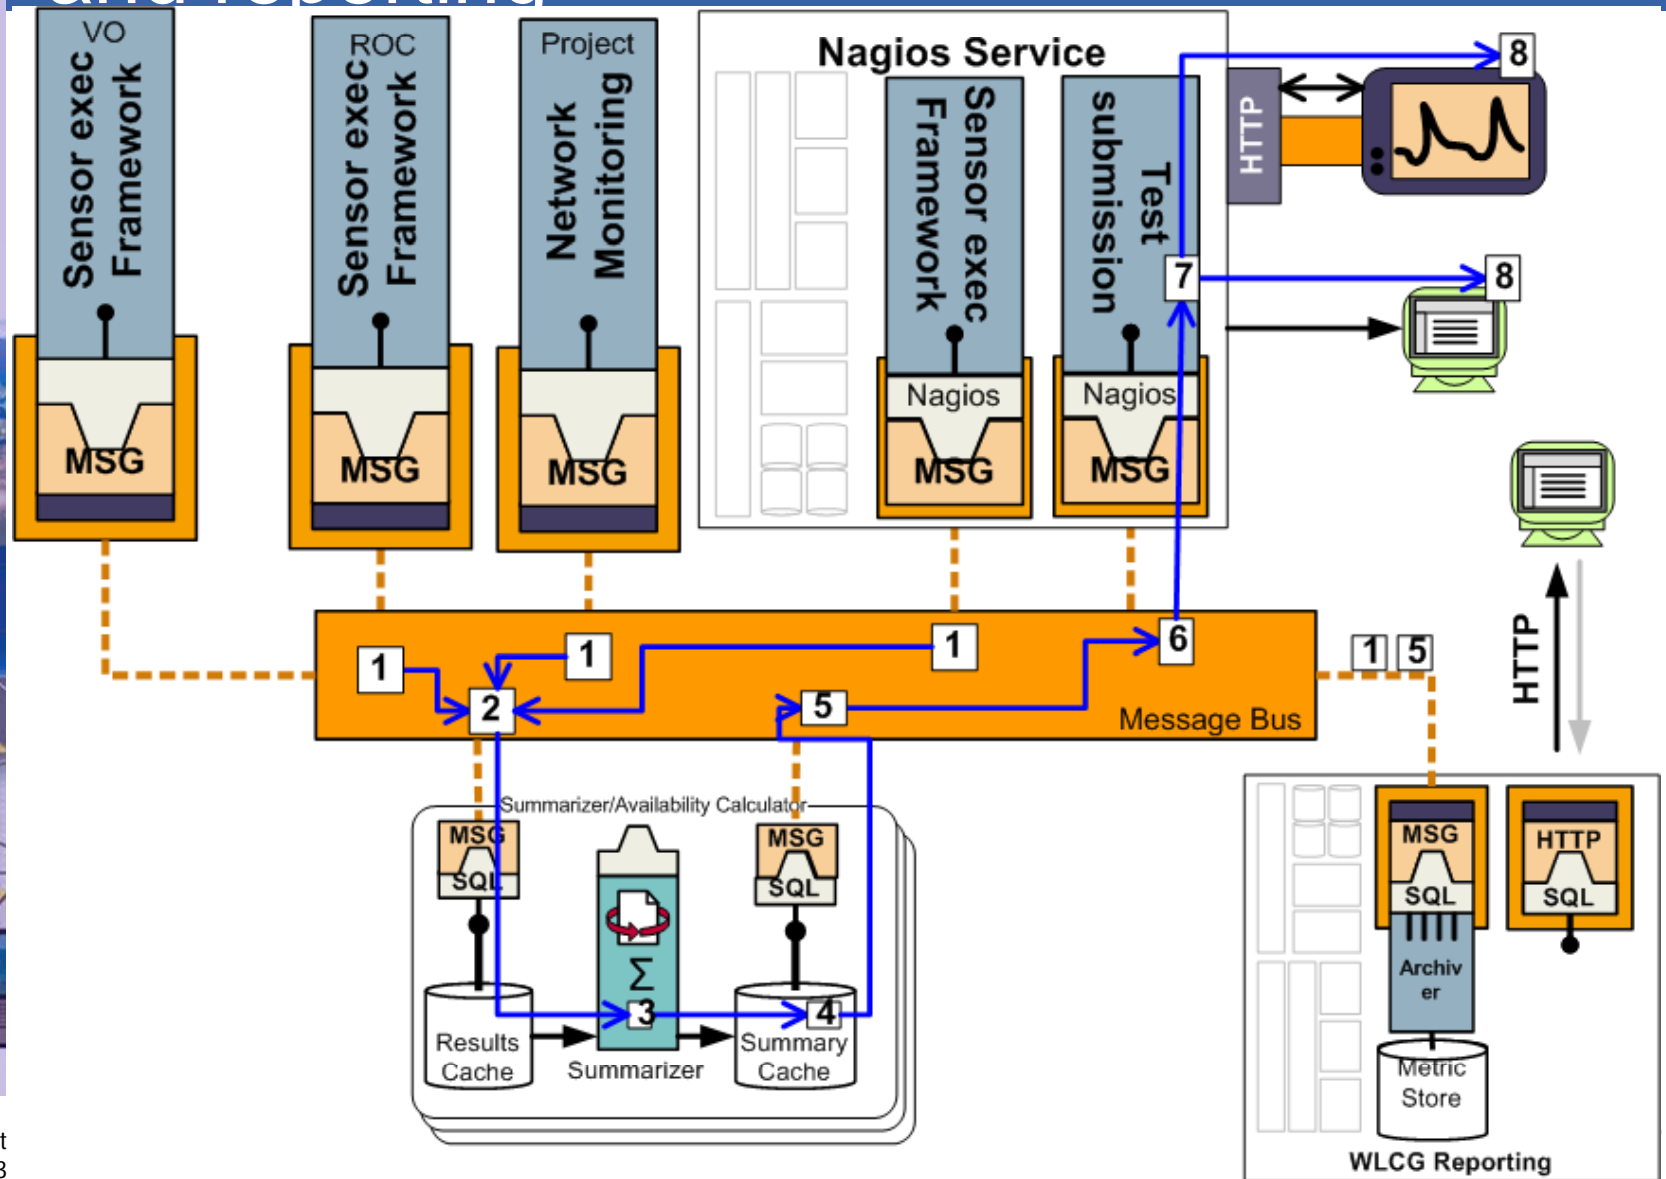


- Existing Systems
 - Service availability monitoring
 - gStat
 - Site Monitoring
 - Usage reporting
- New systems
 - Gridview report publication
 - MoU reporting
 - ServiceMap – VO metric integration
- All fit into the standard pattern mentioned before
 - Now for details...

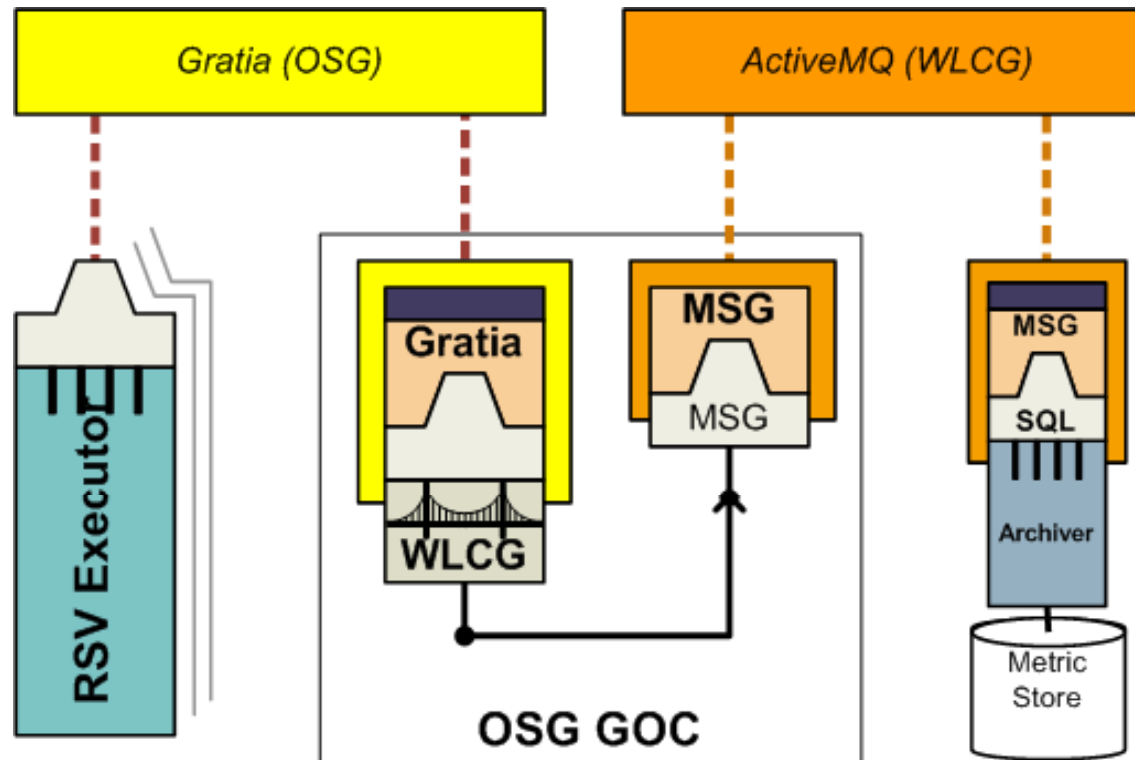




Messaging based archiving and reporting



- RSV – Resource and Service Validation
 - Uses Gratia as native transport
 - And OSG GOC provide a bridge to SAM



- Port ‘useful’ SAM tests into WLCG probe format
 - And supplement where needed
- Move with 1 or 2 nagios-expert ROCs to have them run the SAM tests for their region
 - Setting up a standard execution environment at ROC will be key
- Slowly migrate ROCs over the course of EGEE-III
- Provide VO specific nagios at CERN for experiment specific tests (where needed)

- WLCG working group focus on site monitoring
 - Problem reporting close to the source
- SAM becoming regional doesn't change this
 - Hybrid monitoring strategy
- Site tests often only relevant for site admin
 - Not for availability/reliability calculation


Source	# checks / service	Type
Central	1-2	Network monitoring, Service 'Ping'
Regional	5-10	User-oriented actions (e.g existing SAM tests)
Site local	10-30	Detailed functional tests

- EGEE CEE work on Nagios-based testing for grid services
 - Basis for future SAM execution framework
 - Uses WLCG probe format
 - Compatible with tests from OSG RSV
 - Publishes into ActiveMQ for subset of tests
 - Can be used to augment availability
- New version (ncg v2) is ready
 - Meet needs of sites to integrate into existing nagios installs
 - No need for gLite UI on Nagios server
 - Partitioned config file generated – include what you need

- All SAM tests for all VOs supported at your site


Host ↑↓	Service ↑↓	Status ↑↓	Last Check ↑↓	Duration ↑↓	Attempt ↑↓	Status Information
ce107.cern.ch	CE-ATLAS-GngRbt-Atlas	OK	03-04-2008 23:58:20	1d 21h 42m 37s	1/1	SAM status: ok
	CE-host-cert-valid-OPS	OK	03-05-2008 09:15:12	14d 0h 39m 11s	1/1	SAM status: ok
	CE-lhcb-install-LHCb	OK	03-05-2008 07:20:56	9d 7h 19m 15s	1/1	SAM status: ok
	CE-lhcb-job-Boole-LHCb	OK	03-05-2008 07:21:02	9d 7h 19m 13s	1/1	SAM status: ok
	CE-lhcb-job-Brunel-LHCb	OK	03-05-2008 07:21:06	9d 7h 19m 12s	1/1	SAM status: ok
	CE-lhcb-job-Gauss-LHCb	OK	03-05-2008 07:21:01	9d 7h 19m 14s	1/1	SAM status: ok
	CE-lhcb-os-LHCb	OK	03-05-2008 07:20:50	9d 7h 19m 18s	1/1	SAM status: ok
	CE-lhcb-queues-LHCb	OK	03-05-2008 07:20:54	0d 4h 20m 13s	1/1	SAM status: ok
	CE-sft-brokerinfo-OPS	OK	03-05-2008 09:43:07	14d 1h 27m 33s	1/1	SAM status: ok
	CE-sft-caver-Alice	OK	02-18-2008 12:04:00	15d 23h 37m 7s	1/1	SAM status: ok
	CE-sft-caver-Atlas	OK	03-05-2008 00:13:37	14d 1h 4m 30s	1/1	SAM status: ok
	CE-sft-caver-OPS	OK	03-05-2008 09:43:13	14d 1h 27m 18s	1/1	SAM status: ok
	CE-sft-csh-OPS	OK	03-05-2008 09:43:07	14d 1h 27m 31s	1/1	SAM status: ok
	CE-sft-job-Alice	OK	02-18-2008 12:53:47	15d 22h 47m 20s	1/1	SAM status: ok
	CE-sft-job-Atlas	CRITICAL	03-05-2008 07:38:21	0d 10h 6m 0s	1/1	SAM status: error
	CE-sft-job-CMS	OK	03-05-2008 09:18:26	0d 2h 22m 41s	1/1	SAM status: ok
	CE-sft-job-OPS	OK	03-05-2008 10:18:42	0d 1h 22m 25s	1/1	SAM status: ok
	CE-sft-lcg-rm-OPS	OK	03-05-2008 09:43:39	5d 19h 23m 48s	1/1	SAM status: ok
	CE-sft-sofver-OPS	OK	03-05-2008 09:43:07	14d 1h 27m 33s	1/1	SAM status: ok
	CE-sft-vo-swdir-Alice	OK	02-18-2008 12:03:55	15d 23h 37m 12s	1/1	SAM status: ok
	CE-sft-vo-swdir-Atlas	OK	03-05-2008 00:13:31	14d 1h 4m 36s	1/1	SAM status: ok
	CE-sft-vo-tag-Atlas	CRITICAL	03-05-2008 00:13:32	14d 1h 4m 35s	1/1	SAM status: error


- Documentation to be finished
- Then deployment is key
 - Do we need YAIM for Nagios to help small sites?
- Add features as requested from site admins
 - E.g. Nagios acknowledgements become entries in a related GGUS ticket




- Other main part of monitoring – Usage statistics
 - Gridftp transfers, FTS transfers, job records, ...
- Used to calculate throughput and reliability
- Currently handled in GridView, Dashboards
 - Use messaging system to unite these efforts

- Delegate parsing/routing of specific information back to experts
 - L&B, FTS, ...






- Dashboard team working on getting non-RB job information into L&B direct from Condor submission
- Gridview converting gridftp publishers to messaging system
- Message formats defined for records
 - First records have been sent
- Will add new message types, such as SRM transfers



- Currently a post-processing of results and graphs in Excel
 - Trying to implement it directly on the GridView DB
- Using a mature open-source reporting toolkit – JasperReports
 - UI Report builder – iReports
 - Web-based report server - OpenReports



GS

OpenReports

CERN IT Department

Reliability of WLCG Tier-1 Sites + CERN July 2007 - December 2007

Data from SAM Monitoring. Plots show Reliability for last 3 Months

Reliability is calculated as $\frac{\text{time_site_is_available}}{\text{total_time} - \text{time_site_is_scheduled_down}}$

Target reliability for each site is 91% and Target for 8 best sites is 93% from June, 2007

Target reliability for each site is 93% and Target for 8 best sites is 95% from December, 2007

Site	Average Reliability
CERN-PROD	100%
CA-TRIUMF	94%
DE-KIT	85%
ES-PIC	96%
FR-CCIN2P3	88%
IT-INFN-CNAF	95%
NDGF	99%
NIL-T1	71%
TW-ASGC	65%

Document structure:

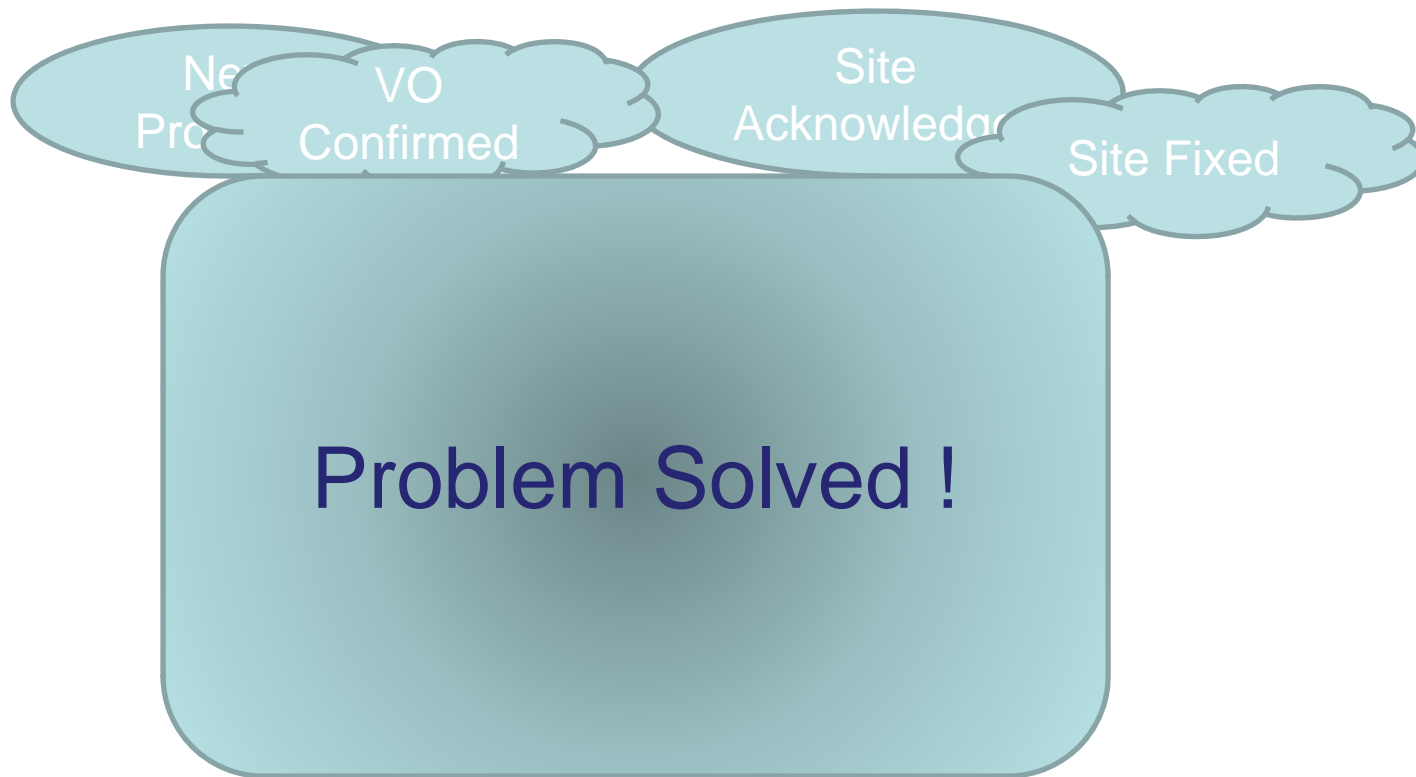
- Document
- Parameters
- Fields
- Variables
- background
- title
- pageHeader
 - static
 - static
 - static
- subreport
- columnHeader
- detail
 - textField [76,98]
 - image-1 [5,0]
- columnFooter
- pageFooter

CERN IT Department
CH-1211 Genève
Switzerland
www.cern.ch/


- We agreed to try and measure MoU metrics during CCRC'08
 - To evaluate if we can actually do it !

<https://prod-grid-logger.cern.ch/eelog/CCRC'08+Logbook/>

Service	Maximum delay in responding to operational problems			Average availability 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	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	24 hours	48 hours	48 hours	98%	98%
All other services – prime service hours	2 hour	2 hour	4 hours	98%	98%
All other services – other times	24 hours	48 hours	48 hours	97%	97%

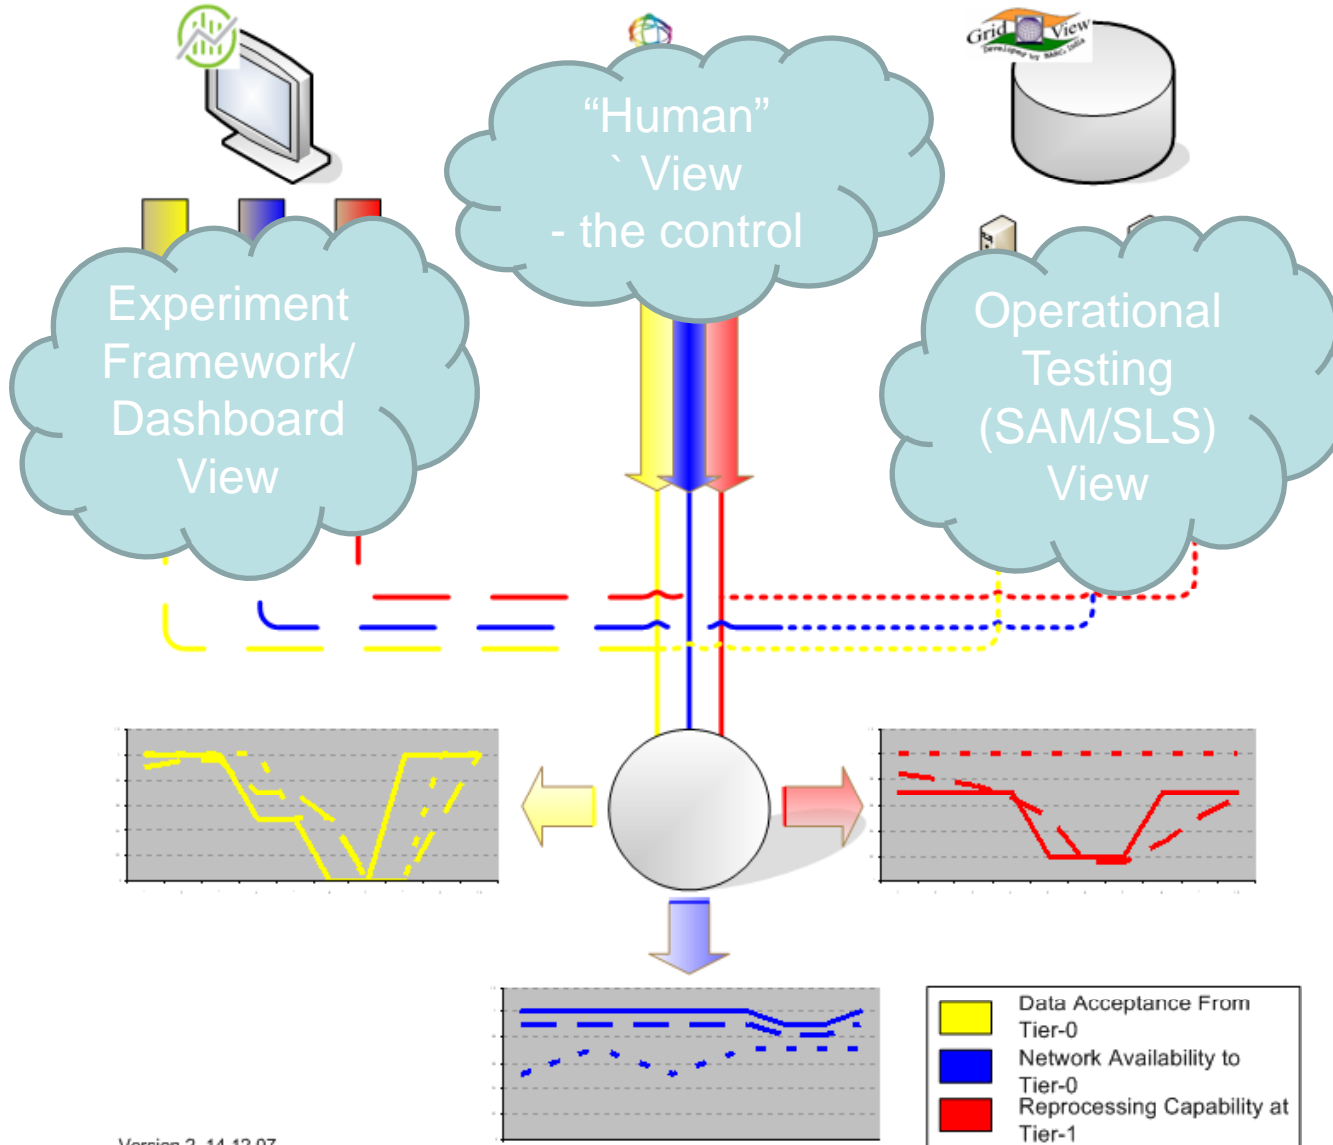


Problem Report: Issue ID #42 : 2008-02-01 10:30 :
MoU Area: CERN-PROD/ Distribution of data to Tier-1 Centres
Time to First Response : 1:00
Time to Problem resolved : 1:29
Time to VO confirmation : 2:23



ier-1 centres,
 ted

Comparing Metrics from Dashboard and SAM/Gridview against the User Experience



Tier-1		Grid Service															
		ArcCE	BDII	CE	FTS	LFC	MYPX	OSGCE	RB	RGMA	SE	SRM	SRMV2	VOBOX	gCE	gRB	sBDII
MoU Category	Acceptance of data from Tier-0 *											•	•				
	Networking Services to Tier-0 *																
	Data-intensive analysis service, including networking to Tier-0	•		•	•			•							•		•
	All Other Services		•			•	•		•	•				•		•	

- Map grid services status (from SAM) to MoU categories
 - These are “custom” service availability calculations
- Use the CMS SAM portal framework as basis for implementing this ?
 - And send results direct to Tier-1 Nagios

Legend: NA OK MAINTENANCE ERROR WARNING INFO NOTE CRITICAL


Note: brightest colors: test is 0 - 6 hours old, ... lightest colors: test is more than 24 hours old

Link to the table

Sitename	Service Type	Service Name	getmeta	del	get	getpfn	put
FZK-LCG2	SRM	gridka-dCache.fzk.de	ok	ok	ok	ok	ok
		pps-srm-fzk.gridka.de	warn	warn	warn	warn	warn
IN2P3-CC	SRM	ccsrm.in2p3.fr	ok	ok	ok	ok	ok
INFN-T1	SRM	castorsrm.cr.cnaf.infn.it	ok	ok	ok	ok	ok
		sc.cr.cnaf.infn.it	warn	warn	warn	warn	warn


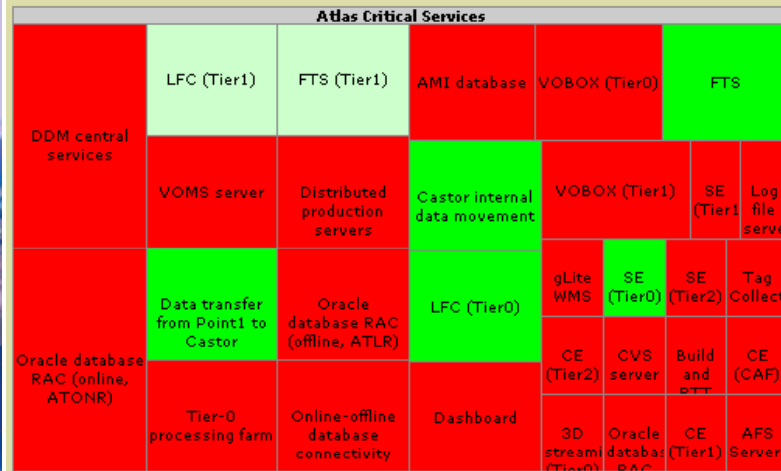
Site	Data Acceptance	Networking	Analysis Services	All Other Services
FZK-LCG2	ok	ok	ok	ok
IN2P3-CC	ok	ok	ok	ok
INFN-T1	ok	ok	ok	error
RAL-LCG2	warn	warn	ok	ok
Taiwan-LCG2	error	ok	ok	ok
pic	error	ok	ok	ok
uscms-fnal-wc1	ok	ok	ok	ok

- What's a ServiceMap?
 - It's a gridmap with many different maps, showing different aspects of the WLCG infrastructure
- What's the CCRC'08 ServiceMap?
 - Service 'readiness'
 - Service availability
 - For VO critical services
 - Experiment Metrics
- A single place to see both the VO and the infrastructure view of the grid




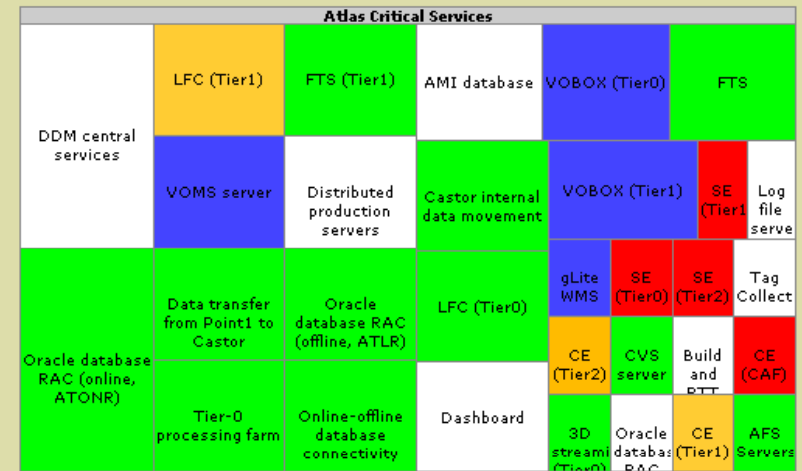
WLCG CCRC'08 Critical Services "GridMap"

Ticklist Status (updated manually)



Alice Atlas CMS LHCb


Test Status (live data)



...Demo...

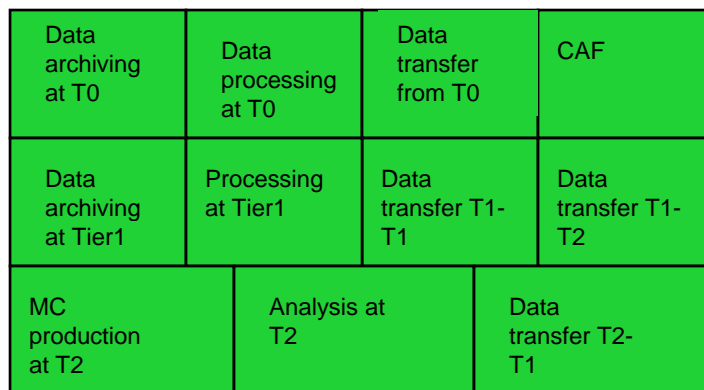
<http://gridmap.cern.ch/ccrc08/servicemap.html>

- Use messaging as interface between VO and the ServiceMap
- VO publishes statistics from dashboard /framework from their perspective
 - Job reliability
 - Data transfer rates
 - Links back to their application for more information
- ServiceMap becomes launchpad for getting to all operational monitoring data
 - For the infratructures and the VOs

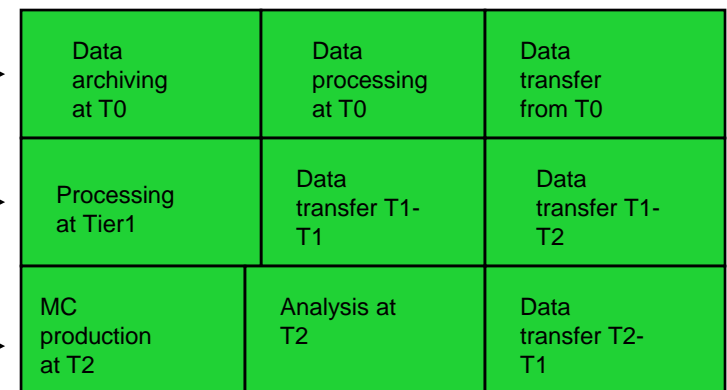


- Functional blocks for LHC experiments are similar to a large extent
 - Allows for a site to compare the service they provide for different experiments
 - e.g - functional blocks for ATLAS and CMS for CCRC08

CMS



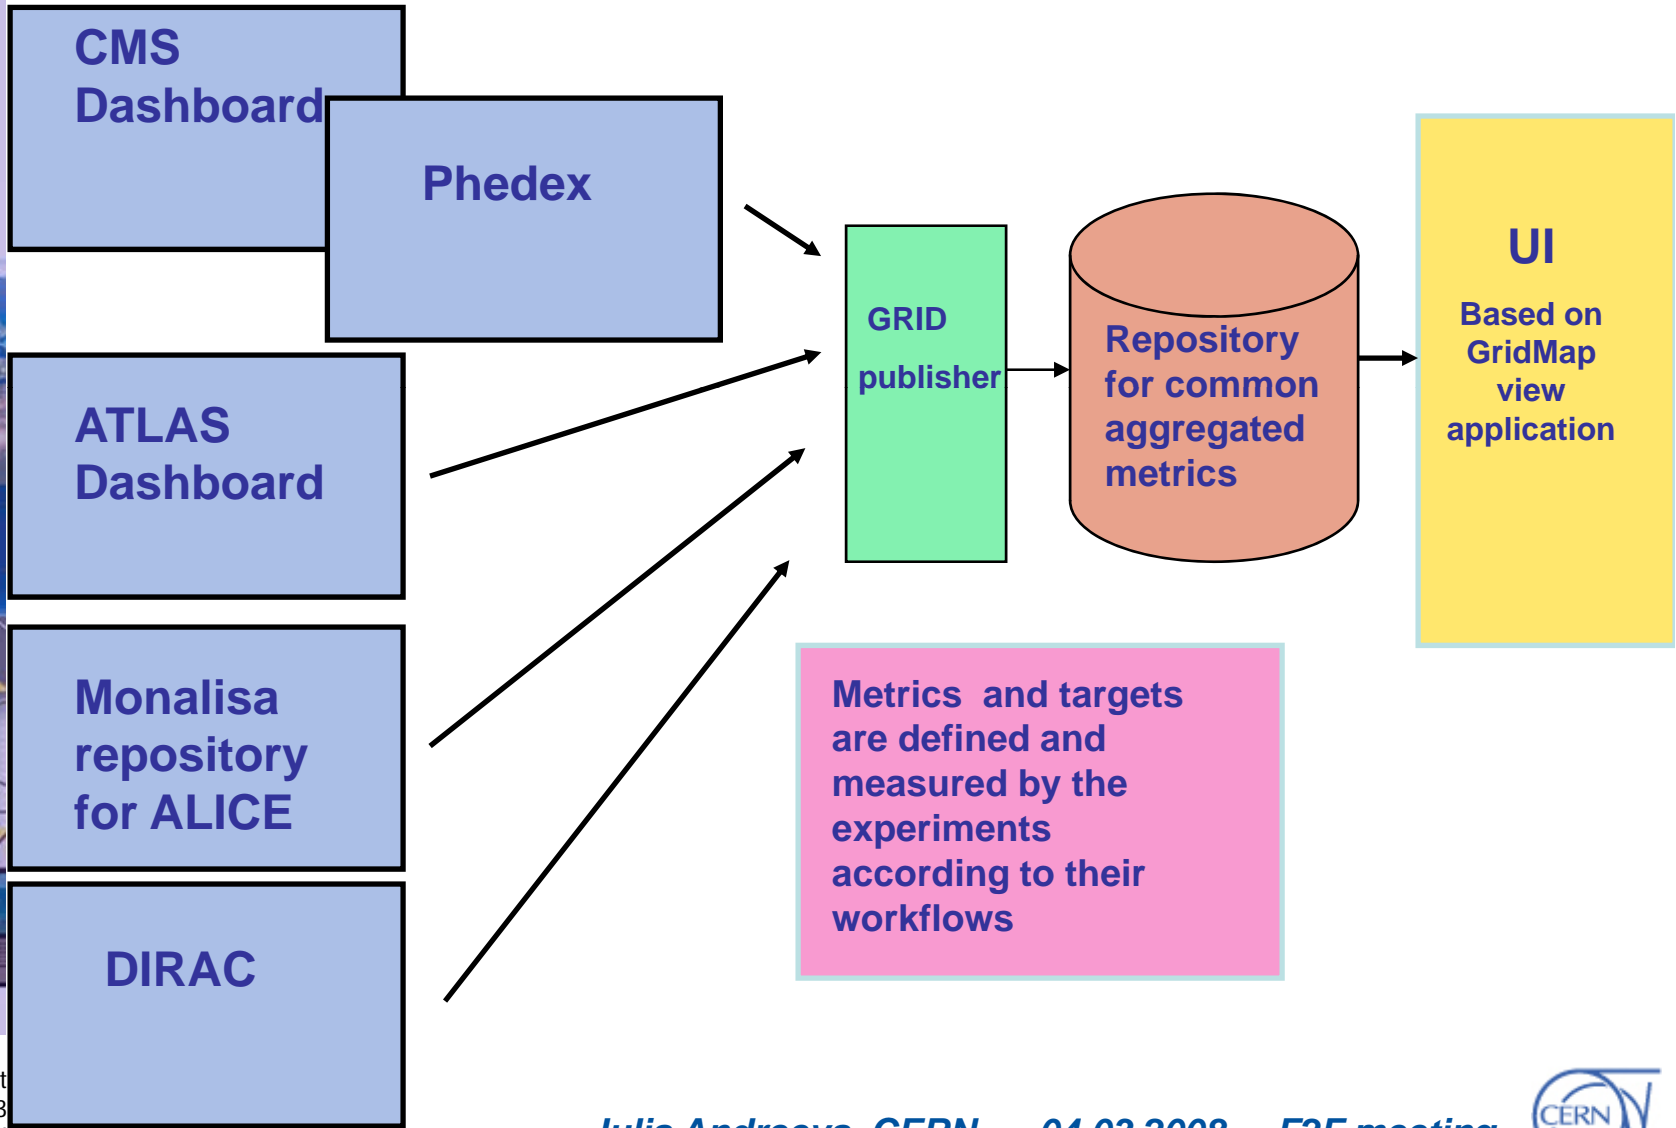
ATLAS

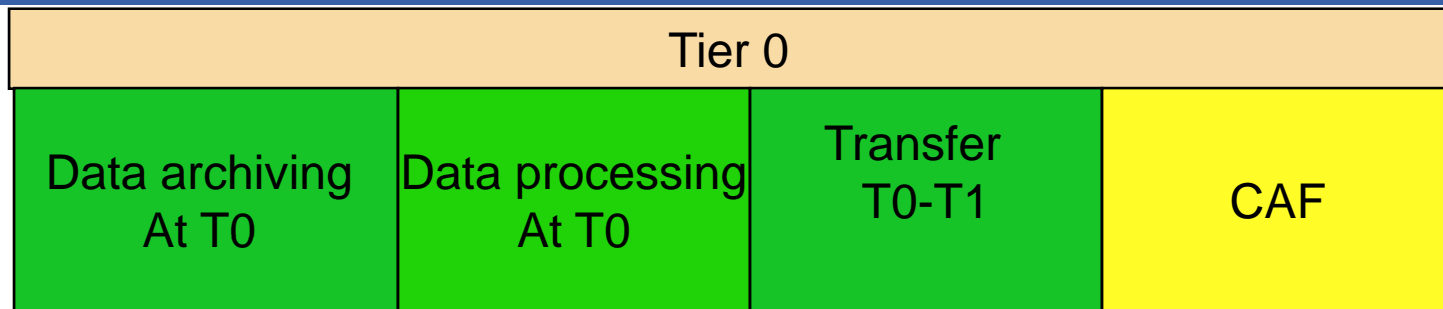


← T0 →


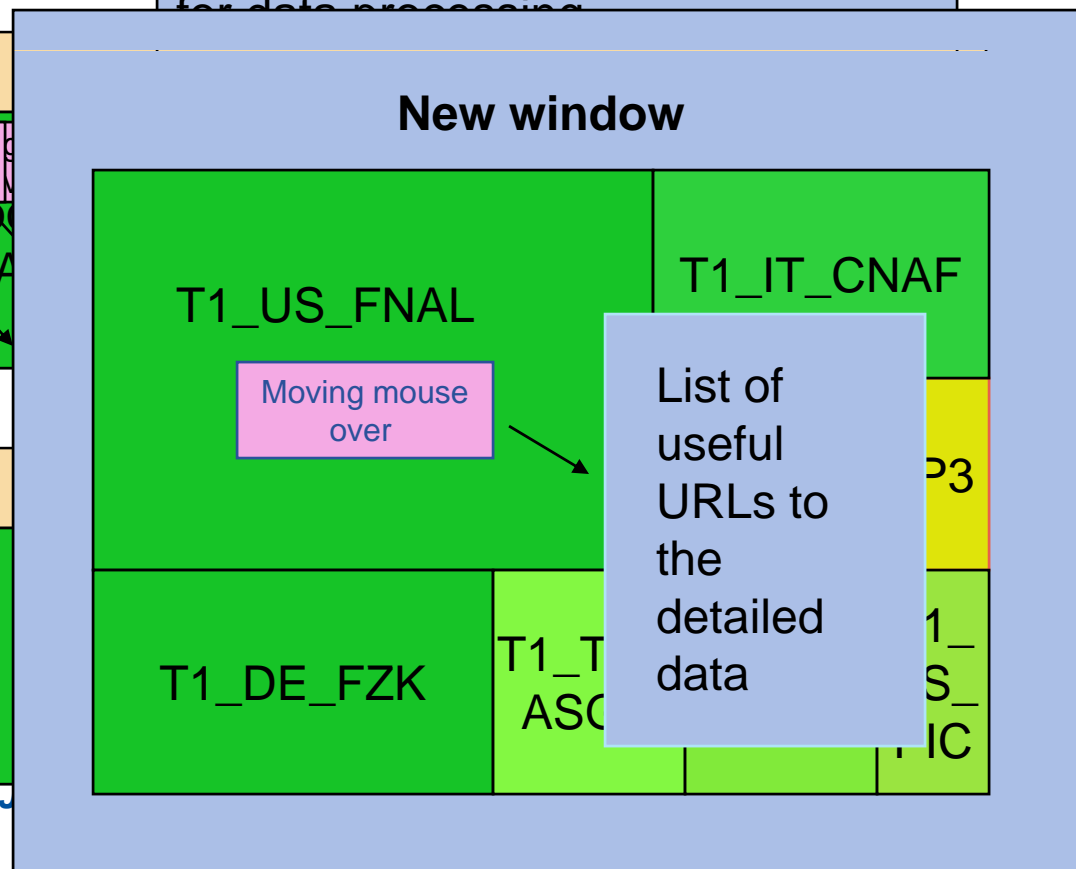

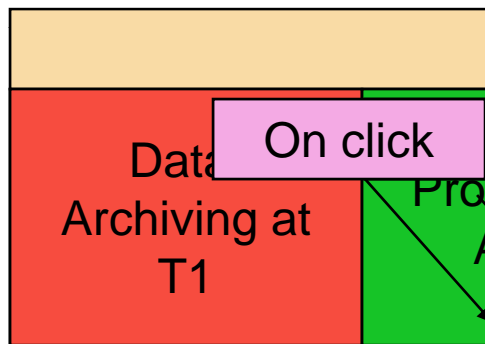
← T1 →

← T2 →






CCRC08 target for number of slots for data processing




- Have a standard model for many of the monitoring data/workflows
 - And components that can help us
- Now it's an integration exercise
 - Augment existing components to talk together
- A lot of work to do
- But it will benefit us in the long term
 - Better sharing of knowledge/resources
 - Reduce operational effort
 - Easier to acquire personnel who know the commodity technologies



- ActiveMQ
<http://activemq.apache.org/>
- Nagios
<http://www.nagios.org/>
- JasperReports
http://www.jasperforge.org/jaspersoft/opensource/business_intelligence/jasperreports/
- OSG RSV
<http://rsv.grid.iu.edu/documentation/>
- CCRC'08 ServiceMap
<http://gridmap.cern.ch/ccrc08/servicemap.html>

Extra Slides



- gStat focuses on monitoring at Site BDII level

Good example of using the commodity software approach to speed up development and deployment !

- Use Nagios as execution framework
- WLCG probe-format compatible probes (in progress)
- Custom presentation layer

Top-Level BDII summary

bdii	goc	pub	hosts	clusters	ldap	services	response
PAKBDII.pakgrid.org.pk	OK	WARNING	1	1	2.013766	1441	24.294298
agh4.atlas.unimelb.edu.au	OK	WARNING	1	2	1.284725	1690	12.27494
atlas-bdii.cern.ch	OK	WARNING	7	1	0.051946	1738	0.249604
bdii-egee.bifi.unizar.es	OK	WARNING	1	1	0.19924	1737	3.341326
		WARNING	3	7	0.130733	1689	0.793513
		WARNING	1	1	0.294597	1689	1.392606
		WARNING	3	19	0.208244	1747	3.356223
bdii.egee-see.org	OK	WARNING	3	2	0.256456	1689	1.300983
		WARNING	1	1	0.157759	3442	9.405109
		WARNING	1	10	0.119491	1738	0.50696
		WARNING	1	1	0.194341	1700	0.736088
bdii.ipp.acad.bg	OK	WARNING	1	3	0.126946	1808	2.607374
bdii.isabella.grnet.g					0.275714	1701	1.725631
bdii.itep.ru					0.278783	1805	1.680495
bdii.marie.hellasgri					0.207686	1689	1.287298
bdii.mif.vu.lt	OK	WARNING	1	7	0.219237	1799	1.24752
bdii.phy.bg.ac.yu					0.22	1755	1.092903
bdii.pic.es					0.379	1689	1.135153
bdii.scotgrid.ac.uk	OK	WARNING	1	3	0.167522	1738	0.780702
bdii.ulakbim.gov.tr	OK	WARNING	1	7	0.309674	1796	1.562301

How many hosts behind this load-balanced alias ?

How many CE clusters reference this top-level BDII?

How many services in this BDII?

Response time ?

