

Monitoring the EGEE/WLCG Grid Services

A. Duarte, P. Nyczyk, A. Retico, D. Vicinanza
CERN – IT

CHEP 2007

Computer facilities, production grids and networking
Victoria, Canada, 4th September, 2007

www.eu-egee.org

- **SAM (Service Availability Monitoring): the platform**
- **Use of SAM: testbeds, sites, grids**
- **Focus on HEP VOs use cases**
- **Conclusions**

- 80 developers (12 research centers)
- gLite 1.0: initial version, released in April 2005
- gLite 1.5: latest LCG-independent version, released in Jan. 2006
- gLite 3.0: merging LCG 2.7 and gLite 1.5, May 2006
- since 3.0 no separate releases of LCG and gLite middleware

- **gLite services groups:**
 - Access and Security Services
 - Information and Monitoring Services
 - Data Services
 - Job Management Services
- **gLite services scopes:**
 - User
 - Site
 - Virtual Organization (VO):
 - Biomedical
 - High Energy Physics
 - etc...
 - and global (i.e. multi-VO)


- **gLite services scopes:**
 - User
 - Site
 - Virtual Organization (VO):
 - Biomedical
 - High Energy Physics
 - etc...
 - and global (i.e. multi-VO)

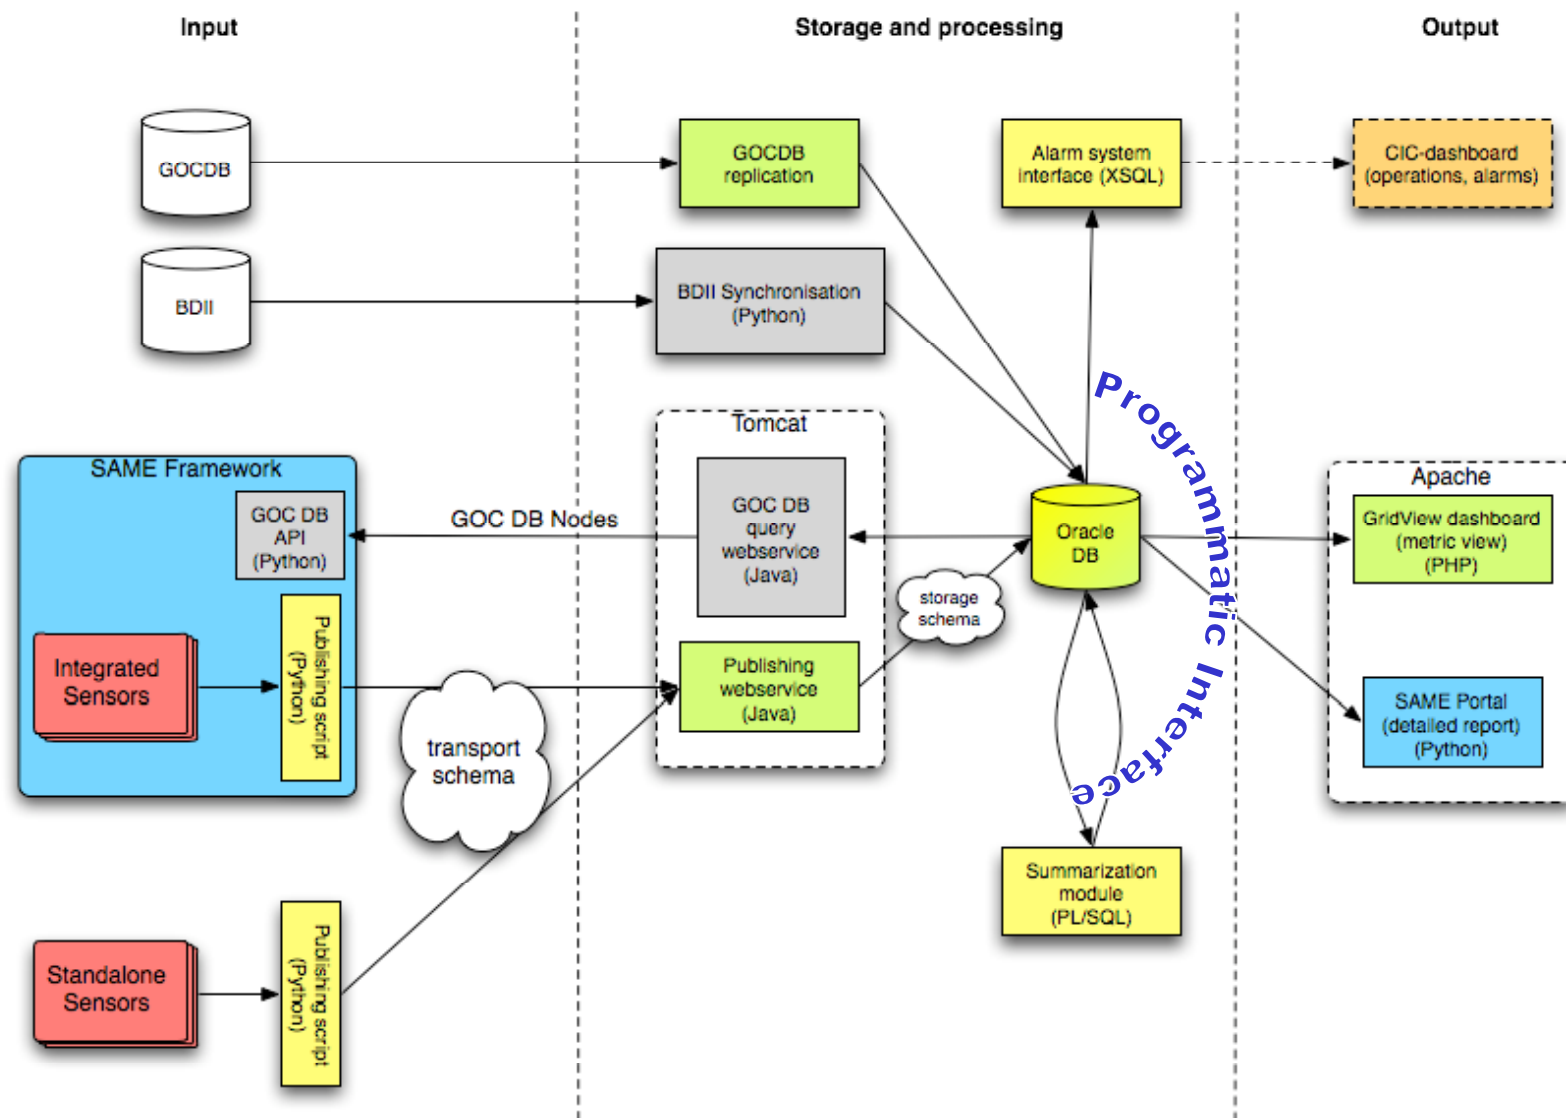
- **Access and Security Services**
 - Identifies users, allowing or denying access to services, on the basis of some agreed policies.
 - provides credentials using Public Key Infrastructure (PKI) X.509
 - Certification Authorities as trusted third parties.
- **Information Service (IS) and Monitoring:**
 - Provides information about the gLite resources and their status.
 - used to locate resources
 - and for monitoring and accounting purposes.
 - Data published to the IS conforms to a schema

- **Job Management System**
 - Computing Element (CE) service
 - computing resources localized at a site (clusters with Worker Nodes)
 - Workload Management System (WMS) - (global)
 - matching jobs to CEs according to job requirements and optimization
 - managing full life-cycle of the job across sites.
- **Data Management System**
 - **storage back-end (site)**
 - **stored files registered in a central catalogue (LFC) (global)**

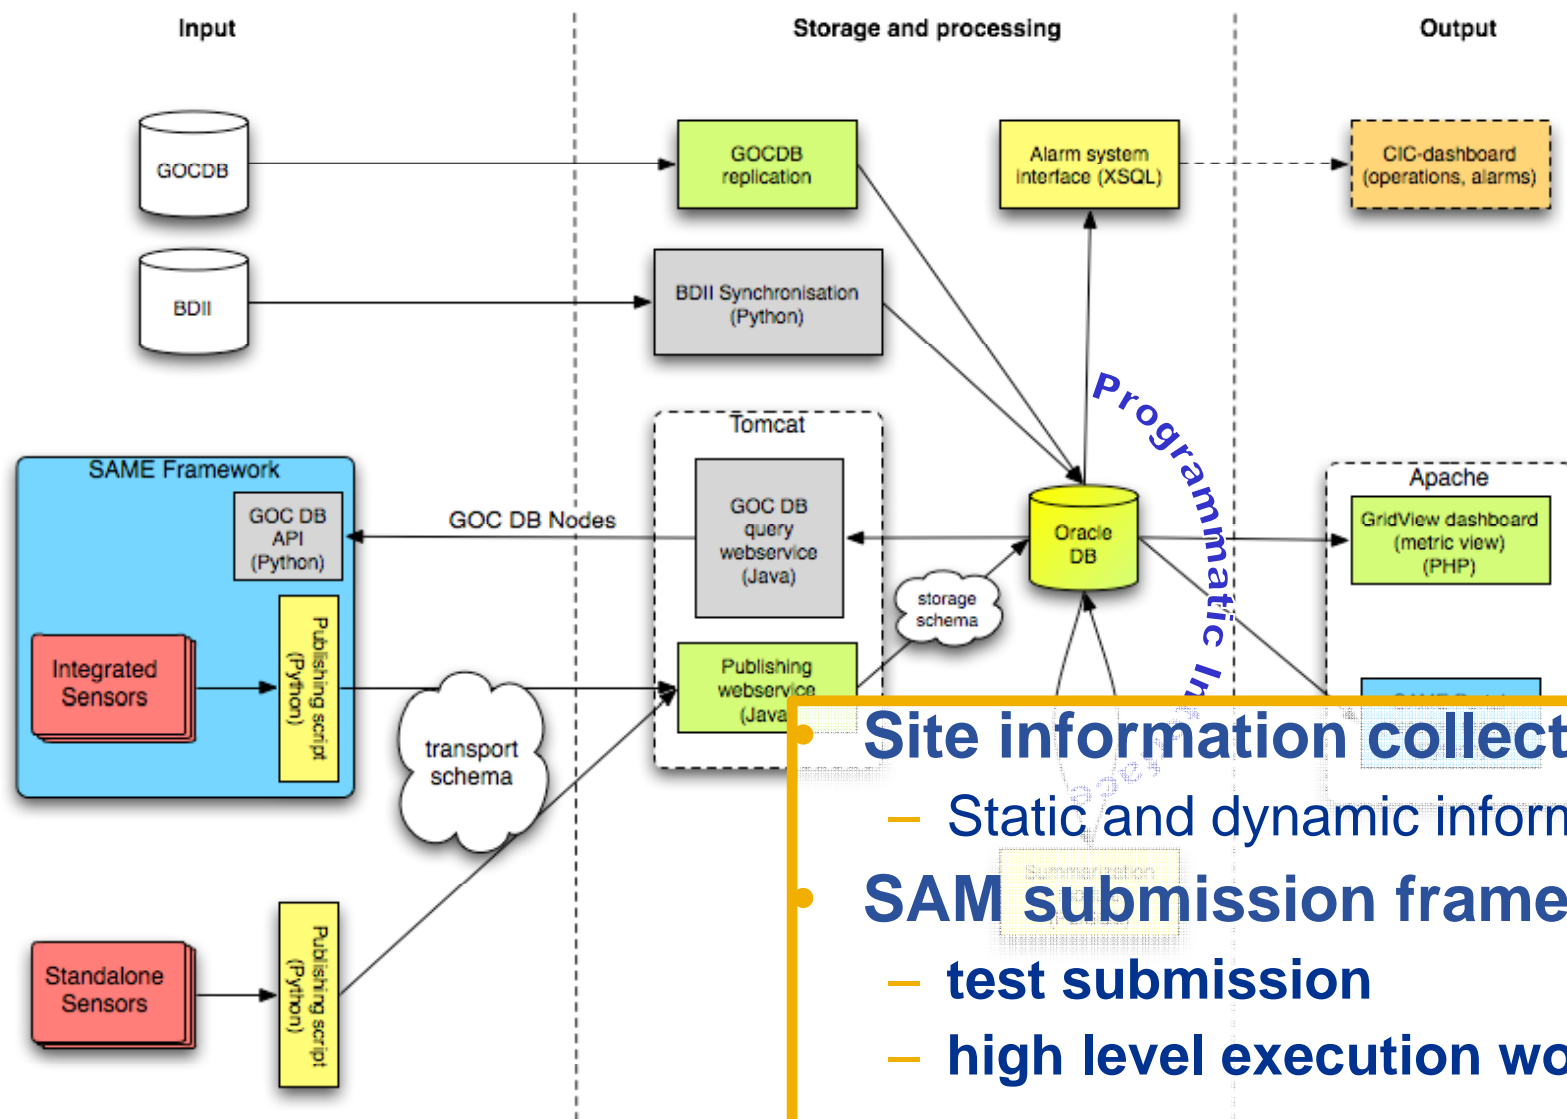
- **Monitoring EGEE/WLCG grid infrastructure**
- **Service level monitoring**
 - Service availability (and functionality) checked by launching tests on the monitored sites
- **In production since June 2006**
- **Managing a growing infrastructure**
 - 20 sites --> 60 sites --> 200 sites (in four years)
- **Main source of information for Grid Operations**
- **Basis for Availability**

- SAM Framework structure
 - **Submission framework**
 - **Oracle DB**
 - **Web Services**
 - **Visualization part (SAM displays)**





SAM framework: input



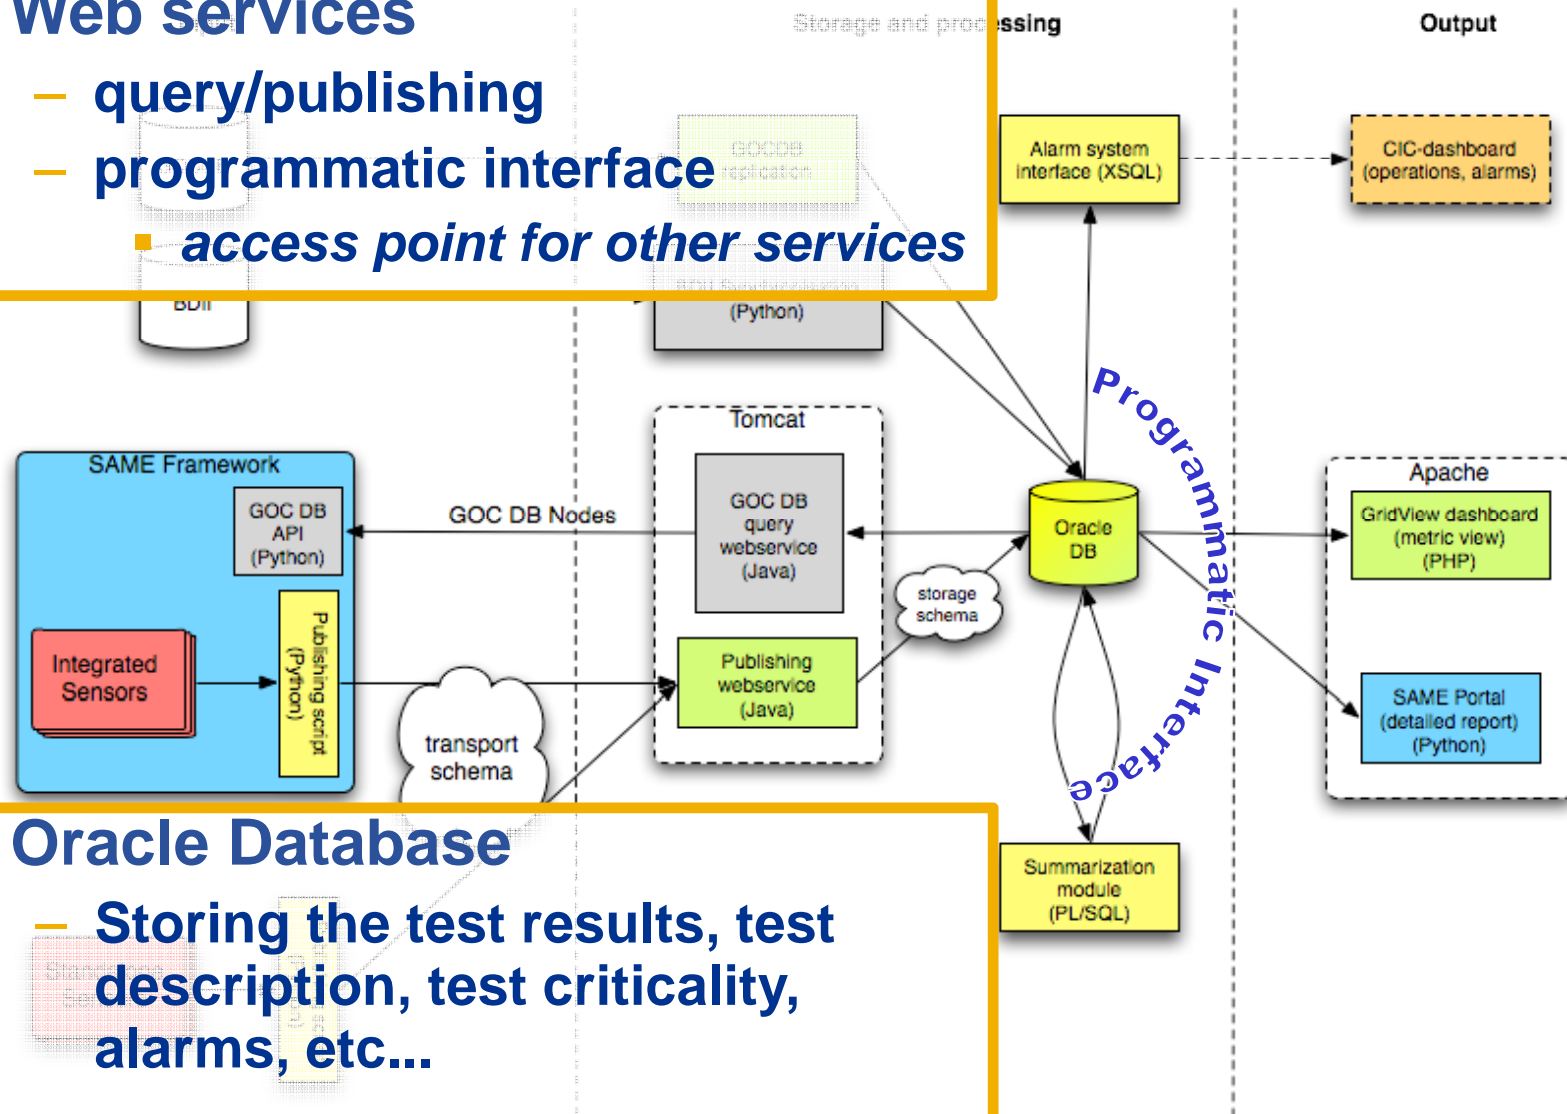
Site information collection tools

- Static and dynamic information

SAM submission framework

- test submission
- high level execution workflow

- **Web services**
 - query/publishing
 - programmatic interface
 - *access point for other services*



- **Oracle Database**
 - Storing the test results, test description, test criticality, alarms, etc...

Input

Storage and processing

Output

SAM portal

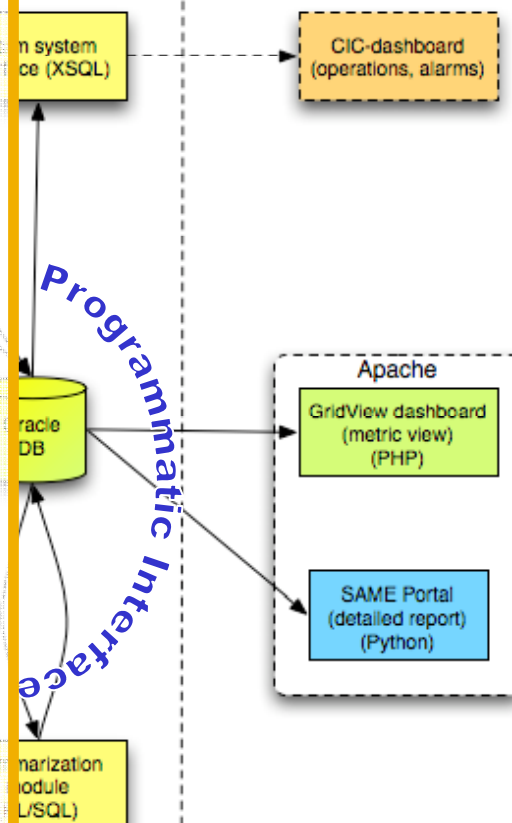
- to be phased but still widely used
- SFT-like presentation → aging

GridView: superset of the functionalities of Sam portal

- availability graphs
- historical test results
- detailed test results

Solutions for an alternative display

- Under discussion
- Many advanced users started developing displays





GridView visualization of SAM results

Enabling Grids for E-science

GridView: Visualization and Monitoring Tool for LCG - Windows Internet Explorer provided by CERN

http://gridview.cern.ch/GRIDVIEW/same_index.php

Monitoring and Visualization Tool for LCG
Data Transfer | Job Status | Service Availability
(Version: gridview-3.0.2, Installation Date: Jan 08, 2007)

Detailed SAM Test Results

Timestamp	Status ID	Status Info	Test Env	Result Summary	Detail Result
at 00:09:05 on 23/01/2007	10	ok	Test Env		Detail Result
at 00:13:43 on 23/01/2007	10	ok	Test Env		Detail Result
at 00:42:44 on 23/01/2007	10	ok	Test Env		Detail Result
at 00:58:09 on 23/01/2007	50	error	Test Env	timeout	Detail Result
at 04:55:02 on 23/01/2007	50	error	Test Env	timeout	Detail Result
at 05:13:34 on 23/01/2007	10	ok	Test Env		Detail Result

Generating JDL LHCb file:

```
Executable = "/bin/sh";
Arguments = "-c 'tar xzf testjob.tgz ; export SAME_WORK='pwd'/work ; bin/same-exec -c same.conf -
StdOutput = "testjob.out";
StdError = "testjob.out";
InputSandbox = {"testjob.tgz","same.conf"};
OutputSandbox = {"testjob.out","testjob-results.tgz"};
Requirements = other.GlueCEInfoHostName == "ce107.cern.ch";
```

Submitting a LHCb job


- **Service Availability is computed**
 - Per Service Instance
 - Per Service Type (eg. CE) for a site
 - Per Site
 - Over various periodicities like Hourly, Daily, Weekly and Monthly

Status of node N = $\bigwedge_{t \in \text{CriticalTests}} \text{TestResult}(N, t)$

\wedge = boolean AND
 \vee = boolean OR

Status of central service C = $\bigvee_{N \in \text{instances}(C)} \text{Status}(N)$

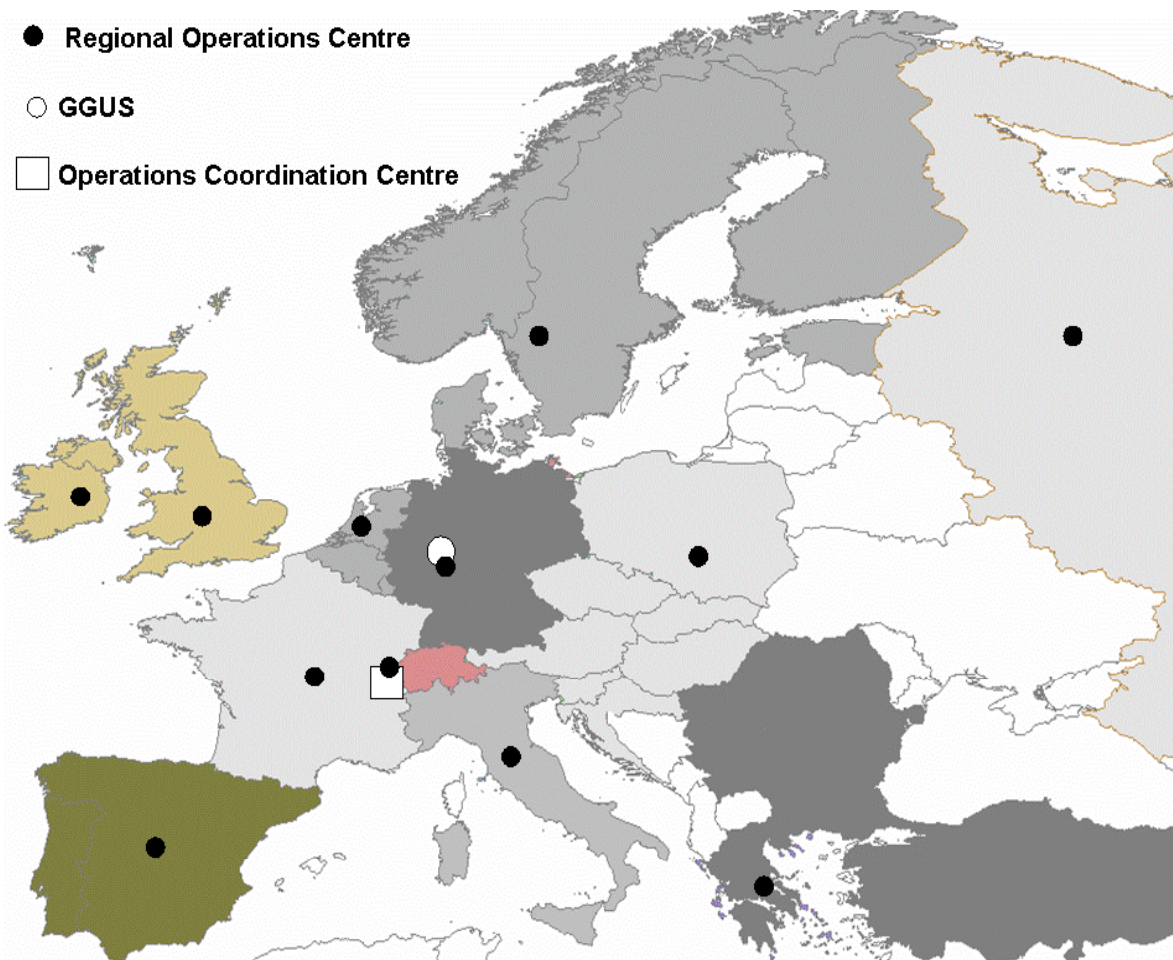
Status of site S = AND



Everything is calculated for each VO
that defined critical tests in FCR
Results make sense only if VO
submits tests!!!

- EGEE/WLCG infrastructure;
 - **~200 sites**
 - **11 federations or regions**
- ROC:
 - **responsibility for the services within its region**
 - **conformity to a set of agreed operation procedures.**
- Grid Operators (COD):
 - **monitoring the availability and performance of the grid services.**

- COD is Operator on Duty
- **global WLCG/EGEE GRID monitoring**
- **SAM tests raise alarms about site failures which are reported to COD**
- **Then COD:**
 - detect issues affecting the grid services
 - provides a first analysis
 - reports existing problems to the relevant ROCs
 - validates the solution provided
- **1 (2) ROCs responsible for the whole GRID operations at a time**
 - 11 ROCs involved
 - weekly rotation



Regional Operations Centres (ROC)

- One in each region (incl. Asia-Pacific)
- Front-line support for user and operations issues
 - point of contact for sites in the region
- Provide local knowledge and adaptations
- Manage daily Grid operations – oversight, troubleshooting
- Run infrastructure services

for Asia-Pacific region

- Asia-Pacific
 - roc@lists.grid.sinica.edu.tw
 - Jason Shih, Min-Hong Tsai, Shu-Ting Liao
- CERN (catch-all ROC)
 - egee-roc-cern@cern.ch
 - Nicholas Thackray

- **SAM platform in use in EGEE-SA3 (Integration/Testing/Release) for middleware certification purposes**
- **Standard tests used as basic functionality and regression tests**
- **Additional tests (e.g. LB, BLAH) integrated locally**
- **Lightweight display interfaced directly SAM DB**
 - (<https://lxb0714.cern.ch/easysam/perl/easysam.cgi>)

- **Grid Operations**
 - Grid Operator-on-Duty (COD)
 - Alarms shown by COD Dashboard are generated by SAM
- **Site Certification**
 - Technical suitability, convenient level of quality
 - SAM test results are crucial in the certification procedures of most EGEE/WLCG ROCs.
 - On demand submission (web interface, Poznan)
 - Official hourly submission (CERN)
- **Availability**
 - ROC reports
- **Site monitoring**
 - site admins, ROC, etc...

- A number of grid infrastructures are currently monitored by SAM. Major examples:
 - **EGEE/WLCG**
 - **SEE-Grid**
 - **EELA**
 - **Health-e-Child**
 - **EuMedGrid**
 - **EuChinaGrid**
 - **BalticGrid**
- SAM platforms were deployed for those projects in slightly different configurations, according to the number of sites monitored, hardware and software resources.

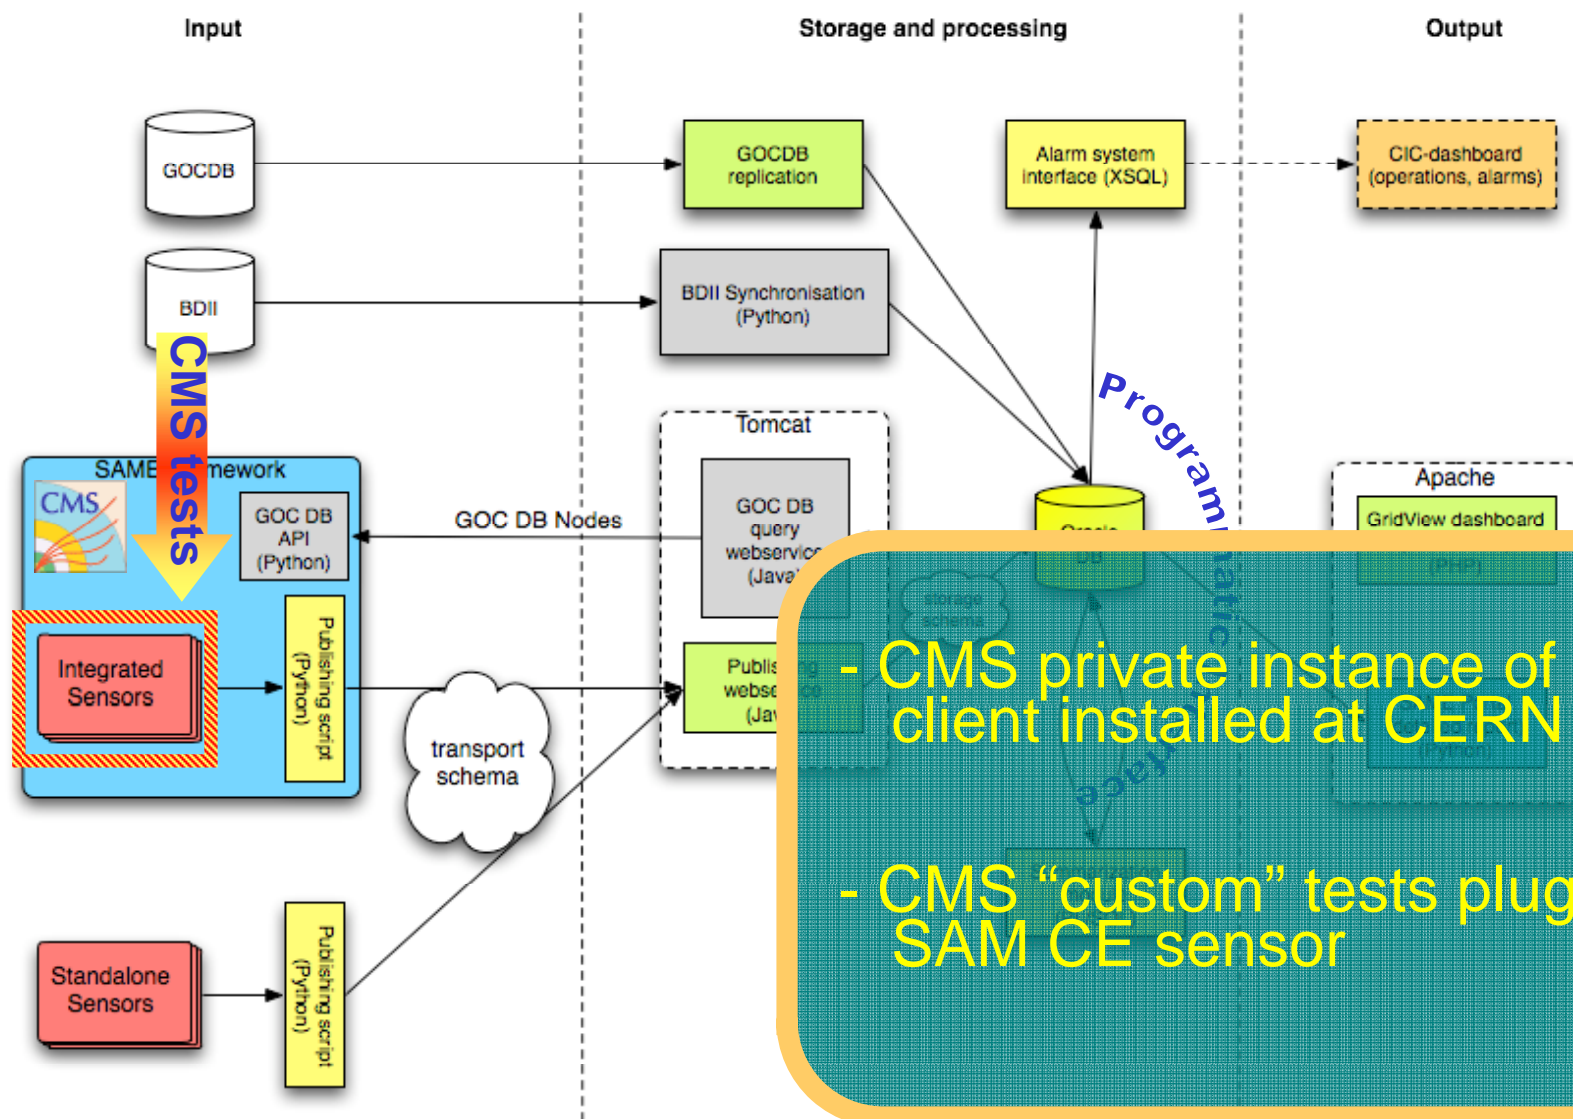
- All the four LHC experiments are running (or planning to run) custom tests using the production instance of SAM
- Goal: sanity checks against selected **grid** and **application** services.
 - **CMS, Alice, LHCb**
 - running custom tests in production using
 - two different submission approaches
 - **Atlas**
 - running standard tests in production using Atlas proxy.
 - preparing to submit custom tests
- The production SAM platform is supporting the four VOs
 - **Only minor changes were needed to support Alice**

Two different approaches

1. Advanced usage of the Submission Framework
 - **CMS, (Atlas)**
2. Hybrid submission methods
 - **ALICE, LHCb**

Both approaches successful and interesting ...


A Clean Integration: CMS



- CMS private instance of SAM client installed at CERN

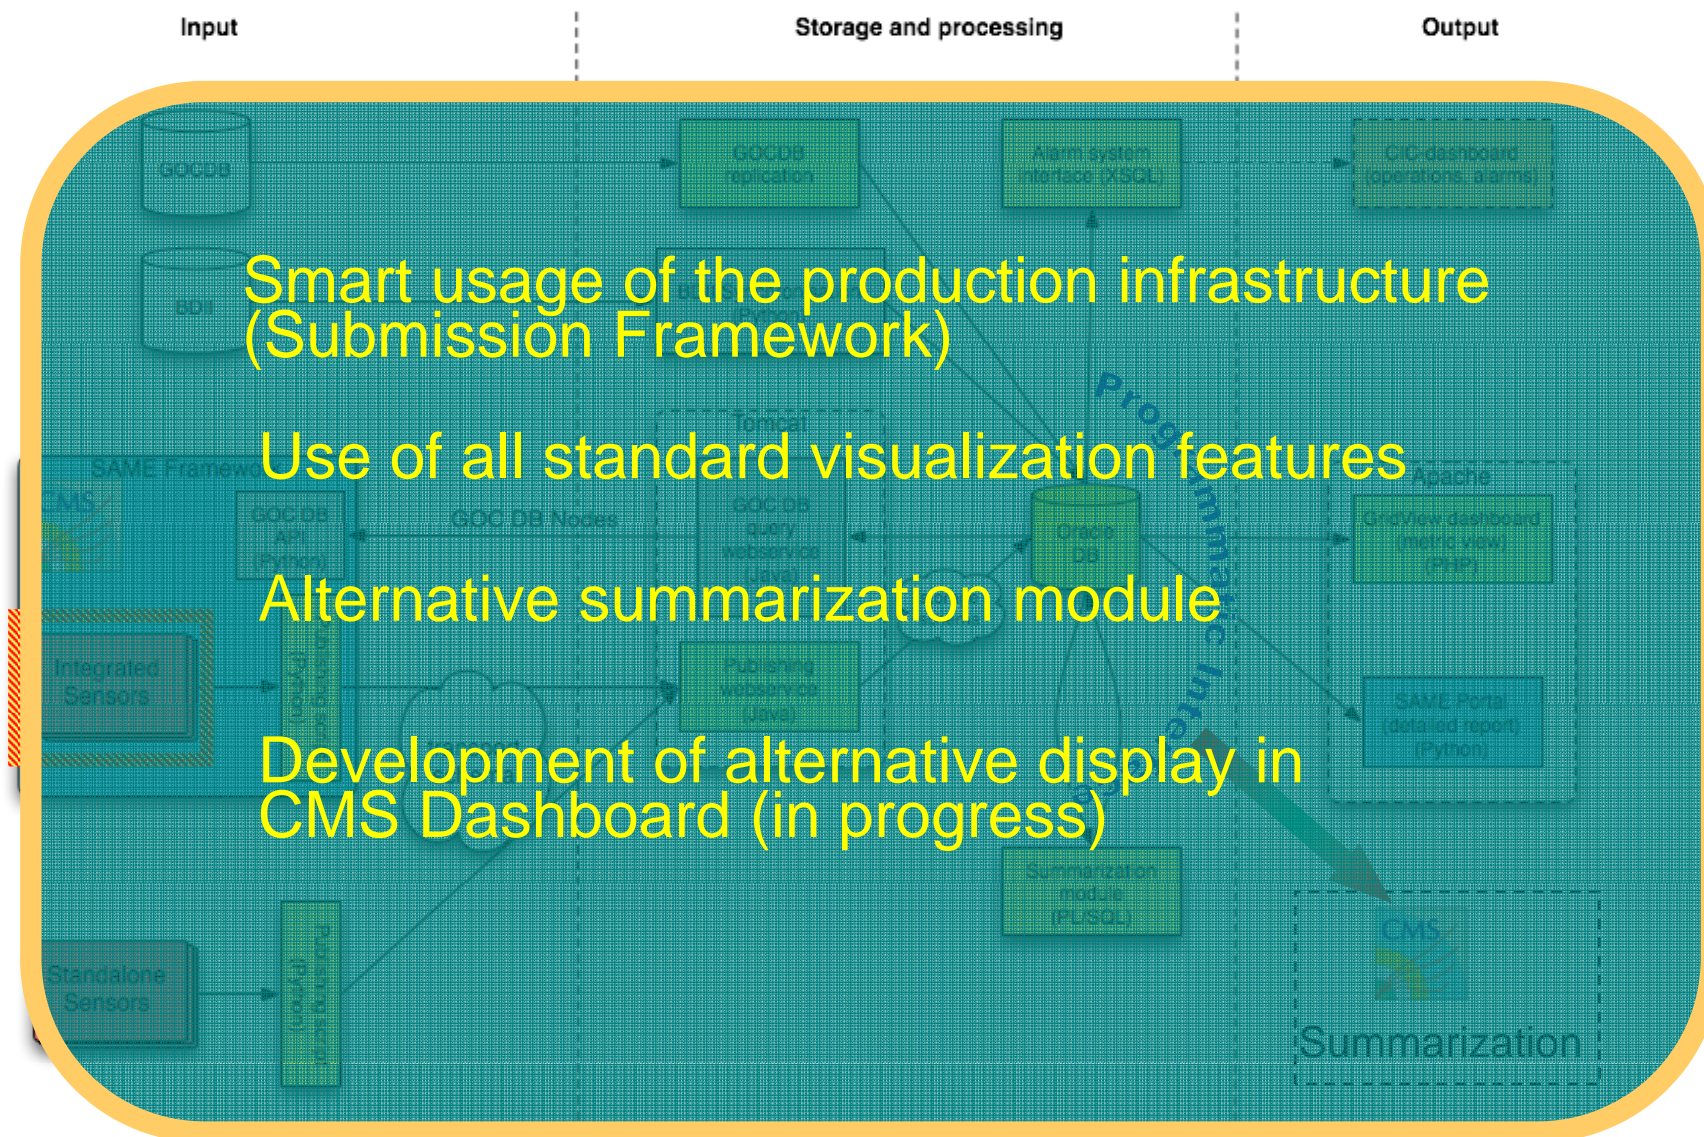
- CMS "custom" tests plugged in SAM CE sensor

- Results of CMS tests are used to calculate availability according to CMS specific metrics
- Data accessed through the Programmatic Interface
- Daily site availability calculated according CMS metrics
- CMS Availability = $\frac{\text{running time}}{\text{total time}}$
- CMS Reliability = $\frac{\text{running time}}{\text{total time} - \text{scheduled downtime}}$



- **CE tests submitted by CMS since early 2007**
 - CMS software area
 - Site local configuration
 - CMS version test
 - local stage out (WN → SE)
 - Discovery of local Squid server
 - Read Calibration data via Squid server
- **SRMv1 and v2 tests also in production**
 - Verify translation LFN → SURL
 - Test data access UI → remote SE
 - push, pull, delete file
 - get file metadata

Credit: Andrea Sciaba' - CMS




- **Off-the-shelf SAM client used to submit standard tests with an Atlas proxy**
- **Developing custom tests and new sensors**
 - SE Sensor: access to DQ2 directories with new lcg-utils
 - Custom SRM: to run low-level test on SRMs
 - Sanity check of software installation
- **Planning an “orthodox” use of Submission Framework**
- **Developing visualization on ARDA dashboard**

An Hybrid Integration: Alice

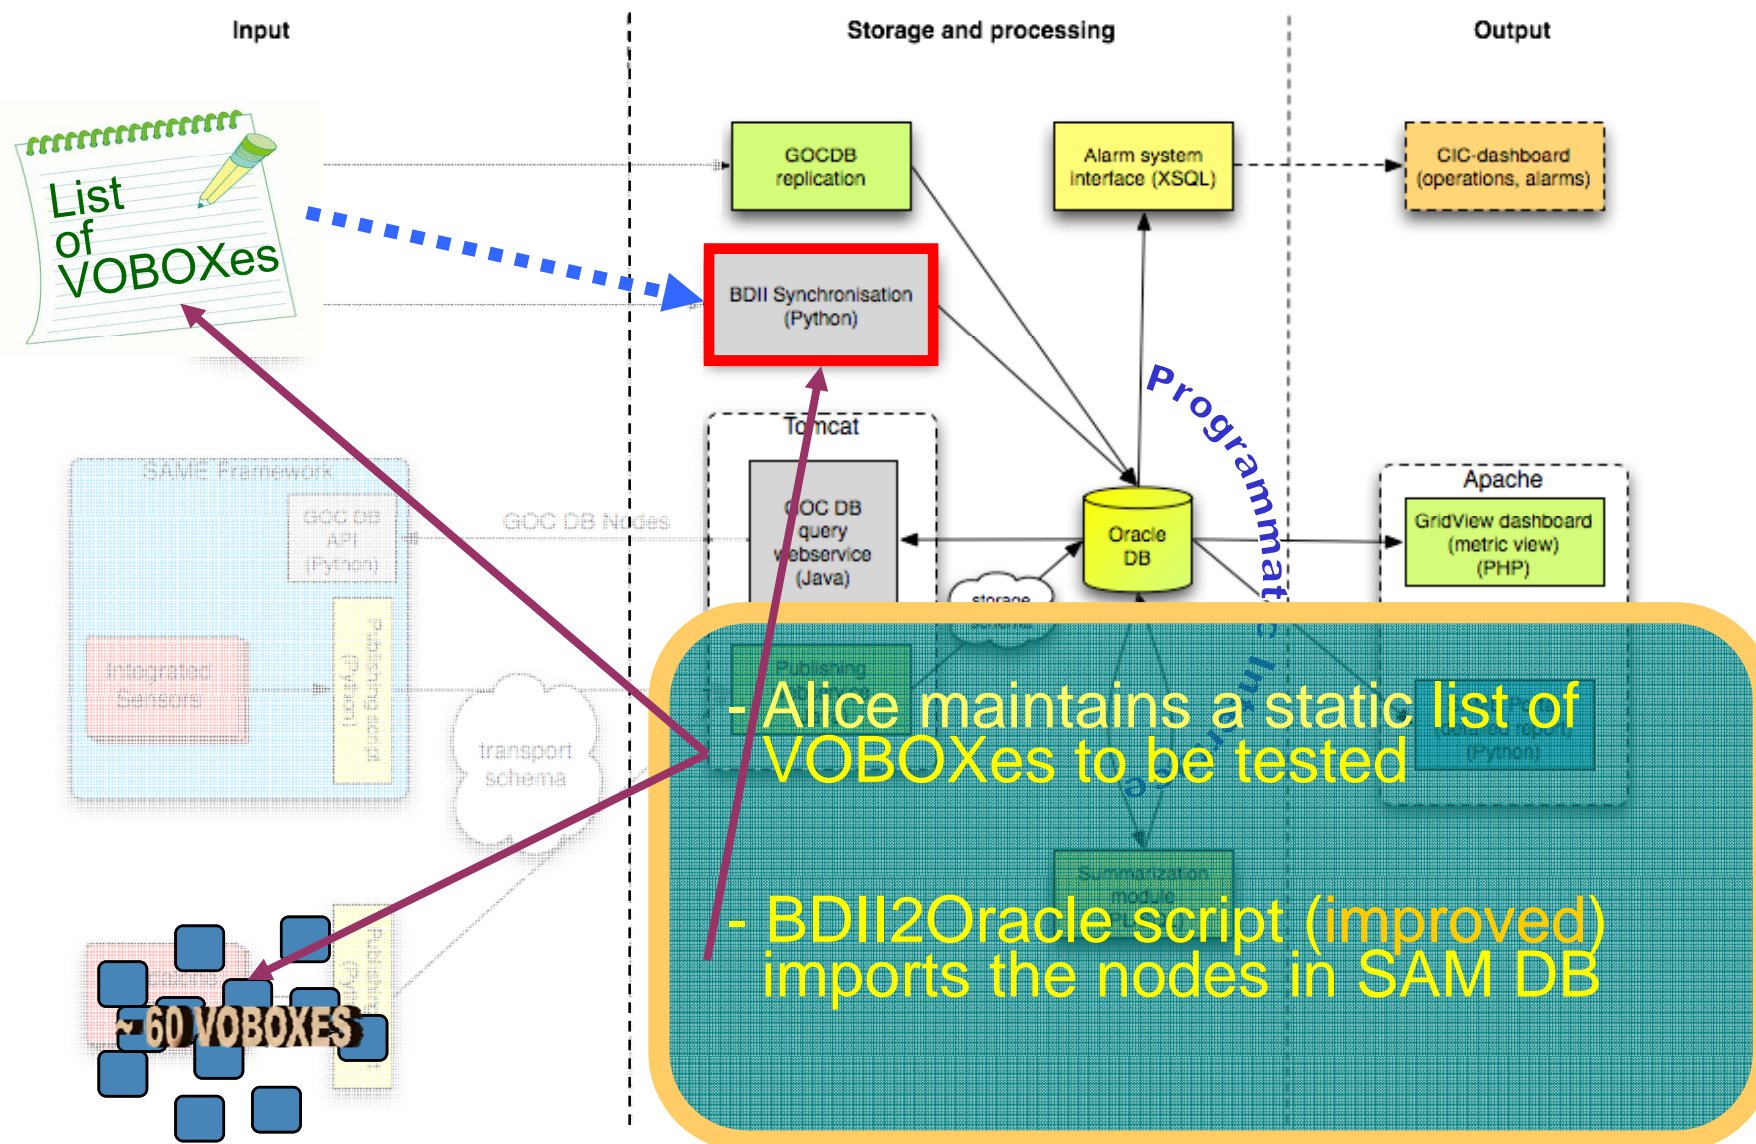
Alice: “Custom” Use Cases

- Alice has developed test scripts to be run on VOBOXES
- Results to be visible in *MonAllsa* ...
- Test to be repeated at all Alice sites (~60)...

Alice Test Case: VOBOXES



Alice test case deposited to the MonALISA VOBOXES



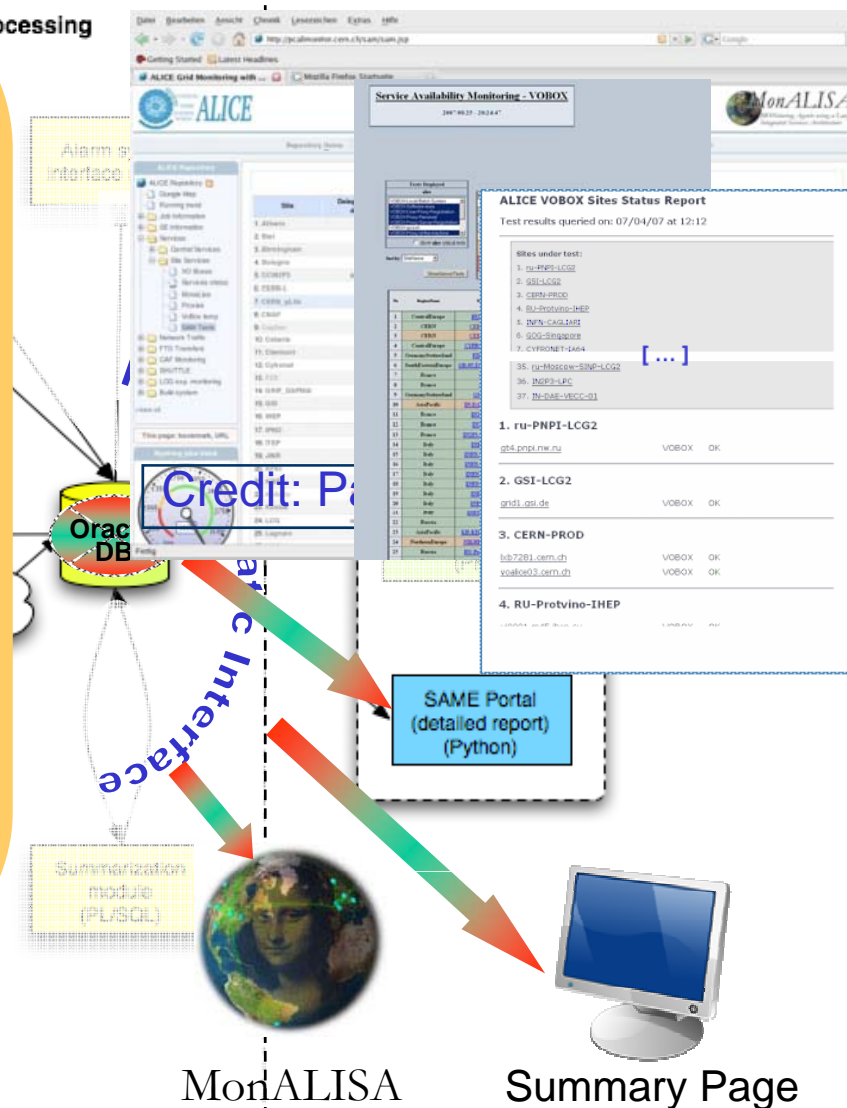


- VOBOS test results available to MonALISA via the SAM Programmatic Interface

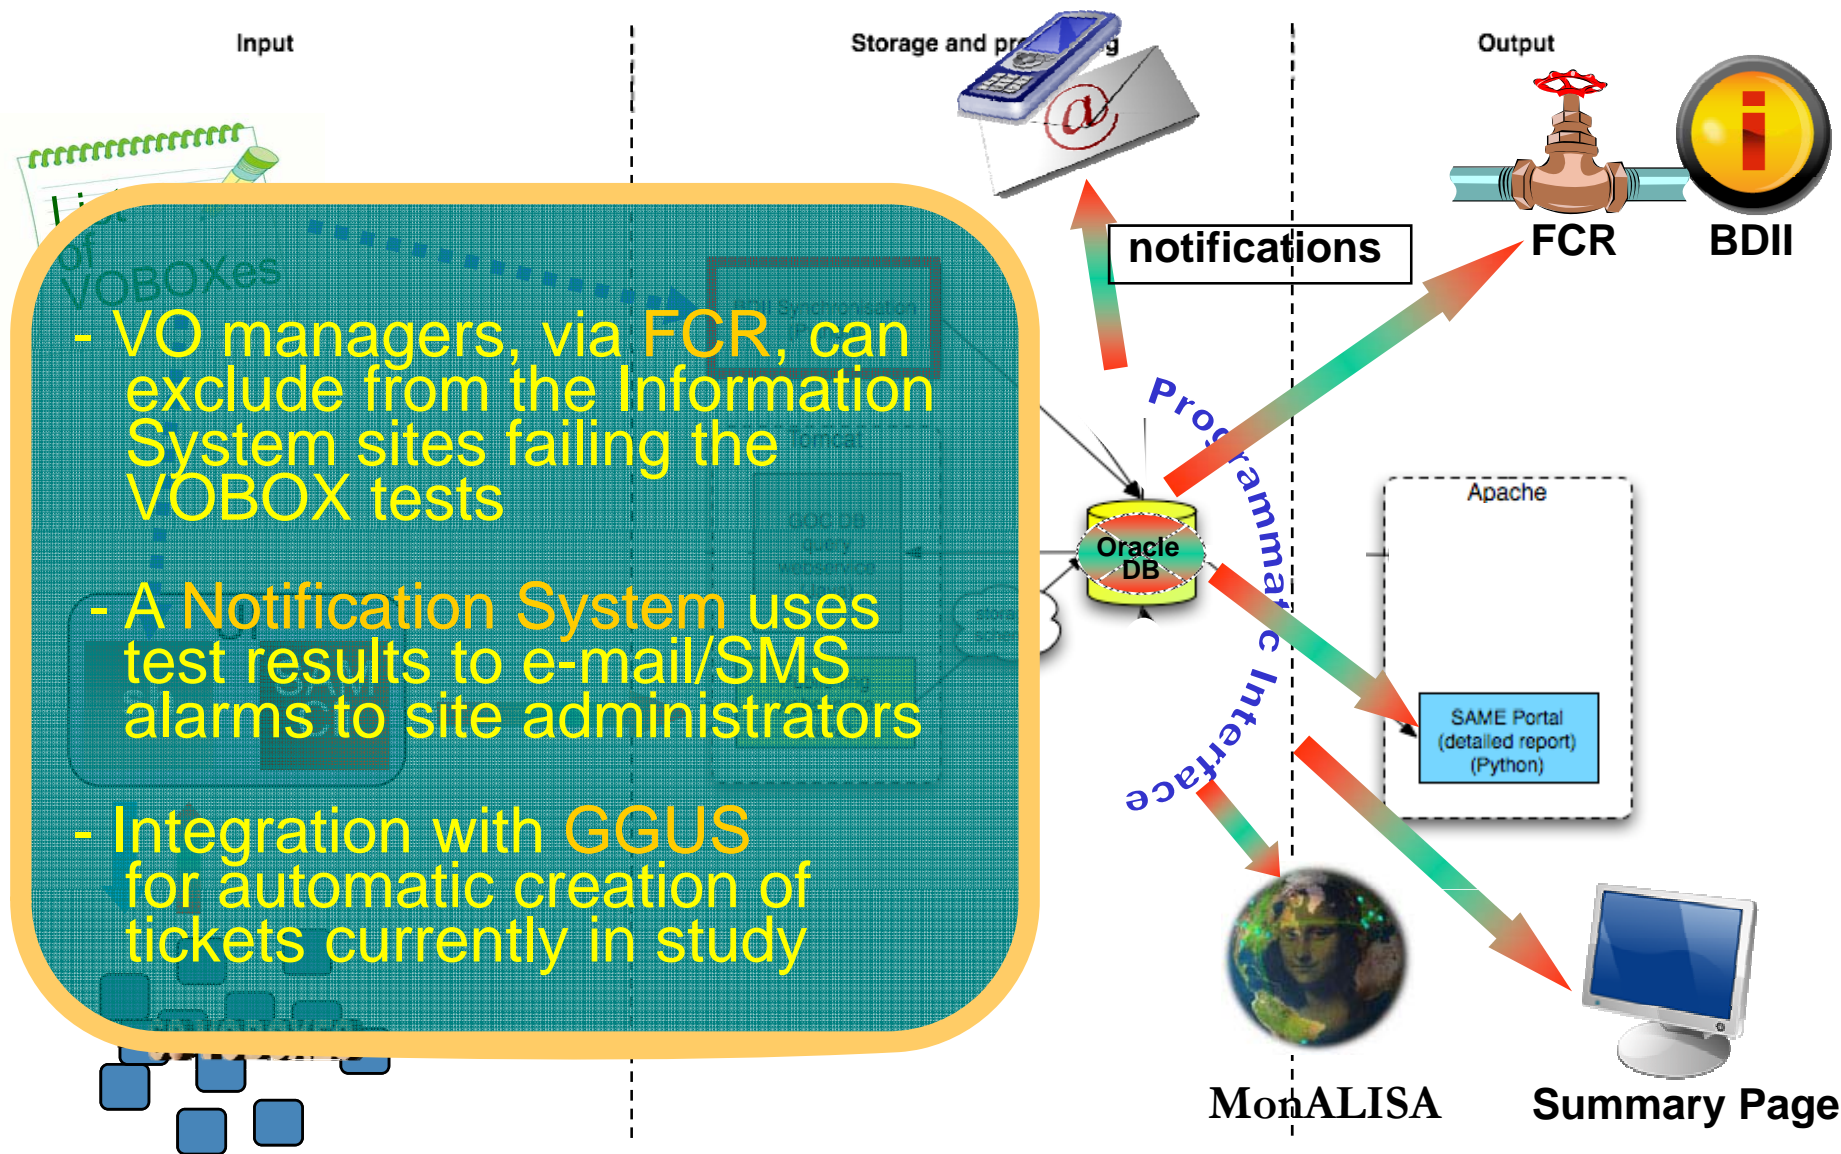
- Results are also visible in the standard SAM Portal display

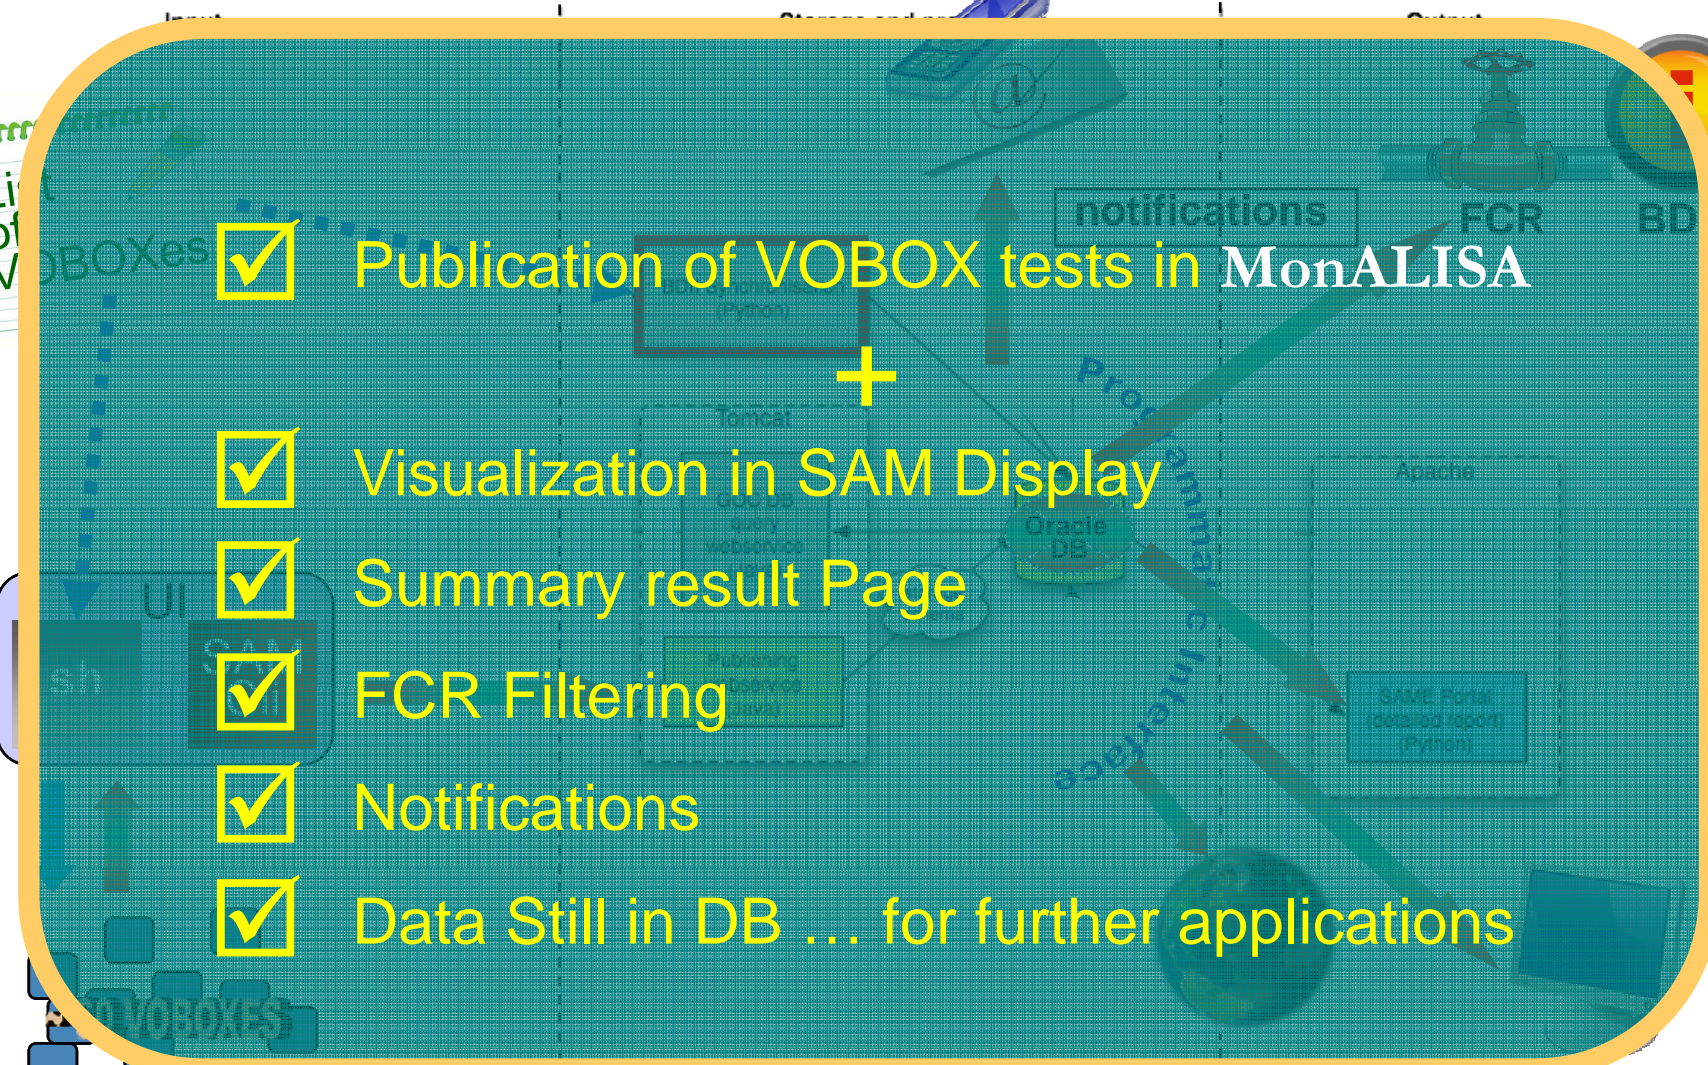
- A Web Page with a summary of VOBOS test results at all Alice sites is generated. Data accessed through the Programmatic Interface

~ 60 VOBOS



Alice Integration: operations



- 
- ✓ Publication of VOBOX tests in MonALISA
 - +
 - ✓ Visualization in SAM Display
 - ✓ Summary result Page
 - ✓ FCR Filtering
 - ✓ Notifications
 - ✓ Data Still in DB ... for further applications

MONALISA

Summary Page

DIRAC Monitoring (Test)

http://lhcb.pic.es/DIRAC/Monitoring/Test/

Accounting Overview

Transfer Accounting

Dashboard

Details

Production ID: All

Site: ANY

Job Status: All

App Status: All

Owner: sam_manager

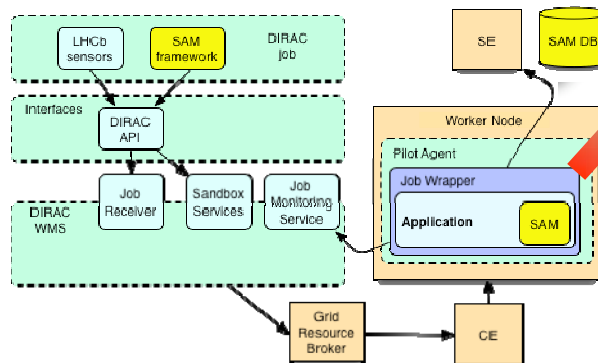
Max results: 100

Job IDs:

Jobs updated after: 10/08/2007

Submit Renewal

Job ID	State	Unknown	Submission	ANY	mailame.cnh.nam.es	2007-08-10 06:01:02	sam_manager
985480	waiting	Unknown	PilotAgent Submission	ANY	mailame.cnh.nam.es	2007-08-10 06:01:02	sam_manager
985488	waiting	Unknown	PilotAgent Response	ANY	prep-cc-02.pd.infn.it	2007-08-10 13:55:21	sam_manager
985490	waiting	Unknown	PilotAgent Response	ANY	thn20.nikhef.nl	2007-08-10 13:59:06	sam_manager
985426	running	Unknown	Starting the application	LCG.NCP.pk	penep04.aep.edu.pk	2007-08-10 13:51:34	sam_manager
985427	running	Unknown	Starting the application	LCG.PAKGRID.pk	CE.pakgrid.org.pk	2007-08-10 13:39:19	sam_manager
985428	outputready	Unknown	Job finished successfully	LCG.LPC.fr	cdlgec02.in2p3.fr	2007-08-10 09:33:26	sam_manager
985429	outputready	Unknown	Job finished successfully	LCG.LPC.fr	cdlgec01.in2p3.fr	2007-08-10 10:23:51	sam_manager
985430	running	Unknown	Starting the application	LCG.Dixton.uk	tzcd07.physics.ox.ac.uk	2007-08-10 13:44:08	sam_manager
985431	running	Unknown	Starting the application	LCG.Napoli-Atlas.it	atlasc01.na.infn.it	2007-08-10 13:54:09	sam_manager
985432	waiting	Unknown	PilotAgent Response	ANY	hognd5.bo.infn.it	2007-08-10 13:59:06	sam_manager
985433	outputready	Unknown	Job finished successfully	LCG.HHP.su	ce0001.m45.thep.su	2007-08-10 11:45:39	sam_manager
985434	outputready	Unknown	Job finished successfully	LCG.Imperial.uk	ce00.hep.ph.ic.ac.uk	2007-08-10 10:49:20	sam_manager
985435	outputready	Unknown	Job finished successfully	LCG.FORTH.gr	ce01.ariadni.hellasgrid.gr	2007-08-10 11:02:55	sam_manager
985436	outputready	Unknown	Job finished successfully	LCG.FTE-RTH.lv	ce01.grd.est.m.lv	2007-08-10 10:12:11	sam_manager
985437	waiting	Unknown	PilotAgent Response	ANY	ce02.esc.gmu.ac.uk	2007-08-10 13:13:13	sam_manager
985438	outputready	Unknown	Job finished successfully	LCG.ACAD.hg	ce02.grd.acad.hg	2007-08-10 12:06:10	sam_manager
985439	outputready	Unknown	Job finished successfully	LCG.CNAF-sic4.it	ce054g.er.cnaf.infn.it	2007-08-10 11:10:05	sam_manager
985440	outputready	Unknown	Job finished successfully	LCG.CERN-sic4.ch	ce109.cern.ch	2007-08-10 11:06:10	sam_manager
985441	outputready	Unknown	Job finished successfully	LCG.CERN-sic4.ch	ce112.cern.ch	2007-08-10 11:41:11	sam_manager

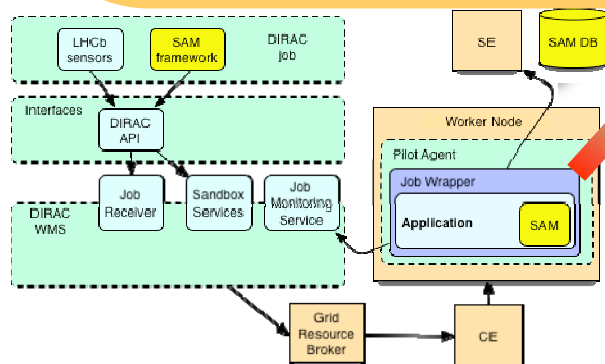


- Submission method conceptually very similar to Alice's
- A DIRAC jobs is used to run tests locally on the WNs
- Results are published with a SAM client shipped on the WN by the DIRAC job
- The progress of the SAM test jobs can be followed with the DIRAC monitoring system

- FCR used to filter sites in downtime
- Test results are used to calculate site availability according to LHCb policies
- Finally a Production Mask is produced with sites eligible for production

Mask for mode : production
Date : 30 Aug 2007
Time : 15:44:52

DIRAC.Bologna-T2.it	DIRAC.Bristol-HEP.uk	DIRAC.IF-UFRO.br	DIRAC.LHCbONLINE.ch	DIRAC.Lyon.fr
DIRAC.ScotGrid.uk	DIRAC.UCD.ie	DIRAC.Zurich-MH.ch	DIRAC.Zurich-spz.ch	DIRAC.Zurich.ch
GLITE.BRISGrid.hk	GLITE.GR-04.gr	GLITE.PDC.se	GLITE.Padova.it	GLITE.ULAKBIM.tr
LCG.ACAD.bg	LCG.ADVES.fr	LCG.Alberta.ca	LCG.BHAM-HEP-slc4.uk	LCG.BHAM-HEP.uk
LCG.BIFI.es	LCG.Barcelona.es	LCG.Bari.it	LCG.Bologna.it	LCG.Bristol.uk
LCG.Brunel-slc4.uk	LCG.Brunel.uk	LCG.CERN-slc4.ch	LCG.CERN.ch	LCG.CESGA.es
LCG.COG.fr	LCG.CIEMAT.es	LCG.CNAF-GRITIT.it	LCG.CNAF-slc4.it	LCG.CNAF.it
LCG.CNB.es	LCG.CNIC.cn	LCG.CPFR-slc4.fr	LCG.CPFR.fr	LCG.CSCS-slc4.ch
LCG.CY01.cy	LCG.Cagliari.it	LCG.Cambridge.uk	LCG.Catania.it	LCG.Darmstadt-slc4.de
LCG.Darmstadt.de	LCG.Durham.uk	LCG.EELA-CIEMAT.es	LCG.EELA-UFRO.br	LCG.EELA.br
LCG.EFDA.uk	LCG.ELTE.hu	LCG.ETP-RTH.lv	LCG.Edinburgh.uk	LCG.FESB.hr
LCG.FORTH.gr	LCG.Ferrara.it	LCG.Firenze.it	LCG.GOC.org	LCG.GR-01.gr
LCG.GR-03.gr	LCG.GR-04.gr	LCG.GR-05.gr	LCG.GRIDKA-slc4.de	LCG.GRNET.gr
LCG.Glasgow.uk	LCG.HG-02.gr	LCG.HG-04.gr	LCG.HG-04.gr	LCG.HP.pr
LCG.HPC2N.se	LCG.HellasGrid.gr	LCG.IC1.ro	LCG.IFH.de	LCG.IFIC.es
LCG.IHEP.eu	LCG.INEP3-slc4.fr	LCG.IN2P3.fr	LCG.INR.ru	LCG.INTA.es
LCG.IPP.bg	LCG.IPSL-IPGP-slc4.fr	LCG.IPSL-IPGP.fr	LCG.IPB.br	LCG.ITP.ru
LCG.IYHE.de	LCG.Imperial-old.uk	LCG.Imperial.uk	LCG.Iowa.us	LCG.JINR-slc4.ru
LCG.KFI.es	LCG.KFI.hu	LCG.KIAE.ru	LCG.KIAM.ru	LCG.KNU.kr
LCG.Krakow-fails.pl	LCG.Krakow.pl	LCG.LAL-slc4.fr	LCG.LAL.fr	LCG.LAPP.fr
LCG.LECCE.it	LCG.LEUVEN.be	LCG.LIP.pt	LCG.LISA.nl	LCG.LPC.fr
LCG.LPN-fails.fr	LCG.LPN-slc4.fr	LCG.LPN.fr	LCG.Lancashire.uk	LCG.LeSC.uk
LCG.Leparc.it	LCG.Liverpool.uk	LCG.MIF.it	LCG.MPI-BGG.de	LCG.Manchester.uk
LCG.Milano.it	LCG.Montreal.ca	LCG.MCP.pk	LCG.NCU.tw	LCG.NIIF.hu
LCG.NIKHEF.nl	LCG.NIPHE.ro	LCG.NOVSU.ru	LCG.NSC.se	LCG.Napoli-Atlas.it
LCG.Napoli.it	LCG.OU.il	LCG.Oxford.uk	LCG.PAKGRID.pk	LCG.PAMUKALE.tr
LCG.PDC.se	LCG.PIC-slc4.es	LCG.PIC.es	LCG.PNPI.ru	LCG.Padova.it
LCG.Pisa.it	LCG.QMUL.uk	LCG.RAL-HEP.uk	LCG.RAL-slc4.uk	LCG.RAL.uk
LCG.RBUL.uk	LCG.SARA.nl	LCG.SINP.ru	LCG.SPACI.it	LCG.SRCE-slc4.hr
LCG.Sheffield.uk	LCG.Sofia.bg	LCG.TAO-slc4.il	LCG.TAU.il	LCG.TCD.ie
LCG.TIFR.in	LCG.Torino.it	LCG.Toronto.ca	LCG.Trieste-slc4.it	LCG.Trieste.it
LCG.UAM-FT.es	LCG.UCL-CCC.uk	LCG.UCL-HEP.uk	LCG.UBER.br	LCG.UIB-BCCS.no
LCG.ULAKBIM.tr	LCG.UNI-KA.de	LCG.UNC.us	LCG.WARSAN.pl	LCG.WCSS-slc4.pl
LCG.WCSS.pl	LCG.WEIZMANN.il	PFS.BHAM.uk	PFS.CERN.ch	PFS.CESGA.es
PFS.CERN.ch	PFS.CNAF.it	PFS.CVP-RH.pl	PFS.DLGE.it	PFS.EEA.int
PFS.FRANHOFFER.de	PFS.GRIDKA.de	PFS.IC.uk	PFS.IIETA.pt	PFS.IJS.pl
PFS.IN2P3.fr	PFS.ISTI.it	PFS.LIT.br	PFS.PB.it	PFS.PIC.es



Production Mask

Summary Page

- SAM client run in a single Dirac task including critical tests for **LHCb applications** ...
 - Length of LHCB queue on the CE
 - Verify OS and architecture
 - Whole MC chain of LHCb applications
- ... and **Grid Services** ...
 - SRM
- Installation of LHCb software and publishing tags

Credit: Roberto Santinelli - LHCb

- **It provides:**
 - status and utilization information at site and resource level
 - basic statistics
 - real-time alerts
 - geographic map
- **Main server based on Nagios (open source, host and network service monitor)**
- **Centralized architecture**
 - a main server periodically queries a set of nodes to extract information about the status of grid and network services, and the utilization of resources.
- **Collected information is stored in a DBMS and used to build aggregate statistics and trigger alerts**

- Information System monitoring web interface
- Analysing data published by the sites
 - **sanity of the data**
 - **reliability of the data**
 - **aggregated and detailed graphs**
 - **history plots**
- Provides information to SAM
- Gathers information the site publishes about the services running there

File
Edit
View
Go
Bookmarks
Tools
Help

http://goc.grid.sinica.edu.tw/gstat/INFN-BARI/

GStat: 11:48:17 06/14/06 GMT

[home](#)
[alert](#)
[table](#)
[service](#)
[regional](#)
[service](#)
[metrics](#)
[links](#)
[?](#)
[prod](#)
[pps](#)
[test](#)
[baltic](#)
[eela](#)
[euchina](#)
[eumed](#)
[seegrid](#)

INFN-BARI Status: OK
ok
GOC
graphs

GOCDB Configuration information:
status: Certified, type: Production
giis url: ldap://gridba2.ba.infn.it:2170/mds-vo-name=inf-n-bari,o=grid

To test site GIIS:: ldapsearch -x -H ldap://gridba2.ba.infn.it:2170 -b mds-vo-name=inf-n-bari,o=grid

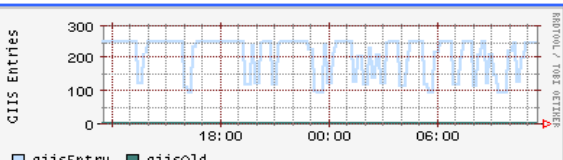
BDII Node Check: .
alert_history ?

CERN SE Check .
alert_history ?

No BDII Node to check in GOCDB
test: ldapsearch -xLLL -l 15 -h bdiidhostname -p 2170 -b 'GlueSEUniqueID=lxn1183.cern.ch,mds-vo-name=CERN-CIC,mds-vo-name=local,o=grid' '(&(|GlueSEUniqueID=lxn1183.cern.ch)(objectclass=GlueSA))' GlueSEUniqueID

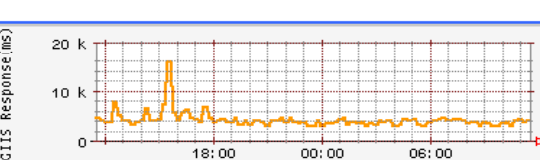
GIIS Perf Check: ok
alert_history ?

Query Response Time (ms): 4097.1 - OK
GIIS Entries Found: 248 - OK
GIIS Old Entries Found: 0 - OK

GIIS Entries


giisEntry
giisOld

giisEnt max: 252.00 avg: 208.49 cur: 248.00
giisOld max: 0.00 avg: 0.00 cur: 0.00

GIIS Response (ms)


giisTime

giisTim max: 15.72 k avg: 4.08 k cur: 4.04 k

GIIS Sanity Check: ok
alert_history ?

Passed

To test site GIIS:: ldapsearch -x -H ldap://gridba2.ba.infn.it:2170 -b mds-vo-name=inf-n-bari,o=grid

Service Check: ok
alert_history ?

hostname	monitor	nodetypes	missing services	history
gridba2	Y	CE,None	none missing	alert_history
gridba6	Y	SE,None	none missing	alert_history
gridba6	Y	MON,SE	none missing	alert_history

- Service Availability Monitoring or SAM, is currently used to
 - **Monitor some of the largest production grids available nowadays**
 - **Improve the **reliability** of the monitored grid services**
- Discussed SAM use for
 - **Middleware Certification**
 - **Grid Operations**
 - **Site Certification**
 - **VO Application Monitoring**
- HEP VO use cases detailed
 - **Clean Integration**
 - **Hybrid Integration**

Thanks for the attention! ☺

Related Contributions at CHEP'07:

232 – Joel CLOSIER - Ensuring GRID resource availability with the SAM framework in LHCb

