



Services reliability and availability SAM3



ALICE

A JOURNEY OF DISCOVERY

Pablo Saiz
IT/SDC

ALICE Tier-1/Tier-2 Workshop
23-25 February 2015



IT-SDC : Support for Distributed Computing

Table of contents

- SAM description
 - Definition
 - Main concepts
- SAM for ALICE
 - Current usage
 - How to improve it

Service Availability Monitoring (SAM)

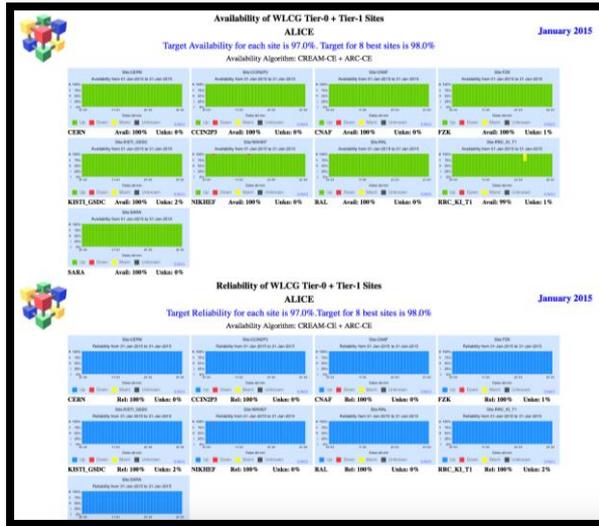
- System to verify **status of services and sites**
 - By submitting tests at regular intervals
- Multiple **sets of tests**:
 - Job submission, Storage, worker node, cvmfs
 - Generic/Experiment specific
- Two infrastructures:
 - Production (<https://wlcg-sam.cern.ch>)
 - Preproduction (<http://wlcg-sam-dev.cern.ch>)
- ATLAS/CMS use SAM results for blacklisting sites
- Monthly WLCG Availability and Reliability reports:
 - <https://espace2013.cern.ch/WLCG-document-repository/ReliabilityAvailability/Forms/AllItems.aspx>
- **SAM3** in production since **November 2014**

Availability/Reliability reports

- Created shortly after the end of the month
 - Draft, 10 natural days for correction requests, final version
- **T0-T1 Summary** (1 page):
 - Availability plots for each site - Tier-0/1 sites
- **T0-T1 6 month History** (1 page)
 - Mean reliability table for 8 best sites - Tier-0/1 sites
 - Mean reliability table for all sites - Tier-0/1 sites
 - Mean reliability table each month/each site - Tier-0/1 sites
- **VOs 6 month details** (25 pages):
 - Availability/reliability plots for each site/month - Tier-0/1 sites
 - Mean availability/reliability tables for each site/day - Tier-0/1 sites
- **All Sites (T0-T1-T2)** (9 pages):
 - Availability/reliability plots for each site Tier-0/1 sites
 - 2 Tables with federation availabilities and reliabilities: sorted by name or availability)
 - Tabl2 with T2 capacities and 3 month historical availability and reliability

Report examples (January 2015)

T1



T2

Tier-2 Availability and Reliability Report

ALICE
Federation Summary - Sorted by Name

Color coding: N/A <30% <60% <90% >=90%

Availability Algorithm: CREAM-CE + ARC-CE

Federation	Availability	Reliability
CZ-Prague-T2	86%	94%
FR-GRIF	99%	99%
FR-IN2P3-IPHC	100%	100%
FR-IN2P3-LPC	100%	100%
FR-IN2P3-LPSC	99%	99%
FR-IN2P3-SUBATECH	100%	100%
HU-HGCC-T2	100%	100%
IN-DAE-KOLKATA-TIER2	100%	100%
IT-INFN-T2	100%	100%

Tier-2 Availability and Reliability Report

ALICE
Federation Details

Color coding: N/A <30% <60% <90% >=90%

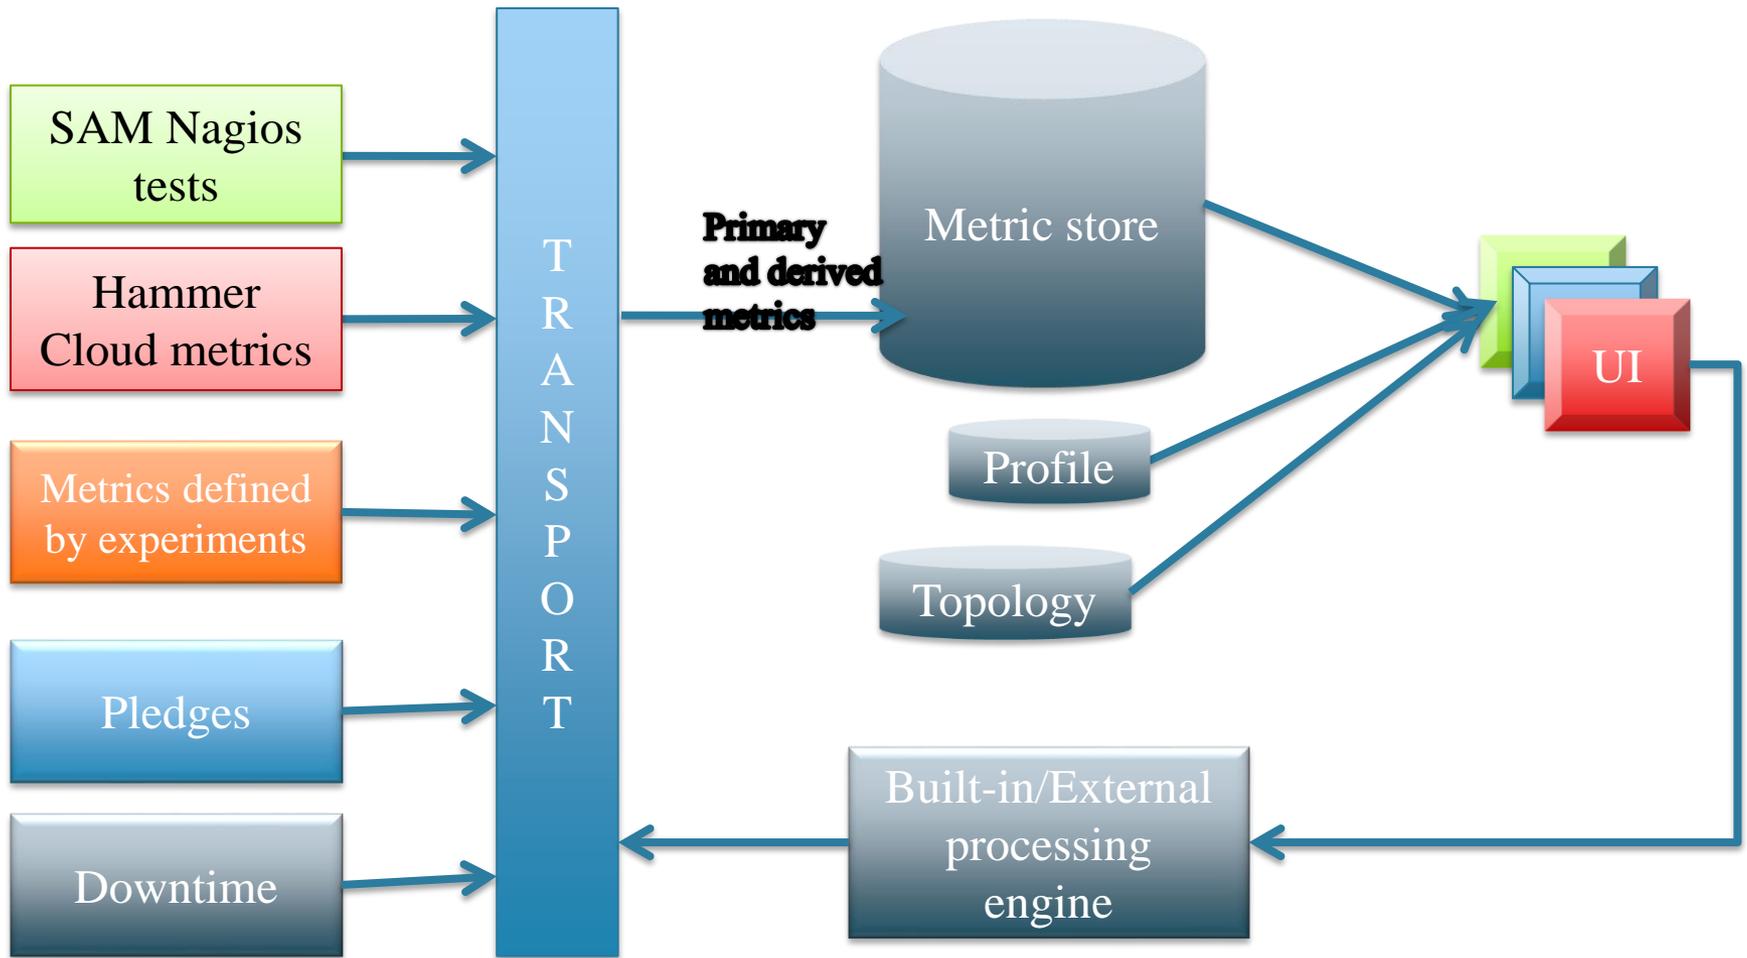
Availability Algorithm: CREAM-CE + ARC-CE

Federation & Sites	Pledge CPU	Pledge Disk	Availability	Reliability	Unknown	Oct-2014	Nov-2014	Dec-2014
CZ-Prague-T2	3500	1030	86%	94%	14%	93%	80%	95%
FR-GRIF	6120	676	100%	100%	0%	99%	100%	97%
FR-IN2P3-IPHC	3500	220	100%	100%	0%	100%	100%	100%
FR-IN2P3-LPC	3030	178	100%	100%	0%	95%	100%	100%
FR-IN2P3-LPSC	2611	296	99%	99%	0%	98%	98%	100%
FR-IN2P3-SUBATECH	3000	310	100%	100%	0%	100%	100%	100%
HU-HGCC-T2	1600	120	100%	100%	0%	100%	99%	100%
IN-DAE-KOLKATA-TIER2	15000	240	100%	100%	2%	95%	100%	91%

SAM3 benefits

- Experiments have **more power** (and more responsibility!)
 - Definition of sites and services
 - Definition of profiles
 - Injecting their own metrics
 - Possibility to overwrite (recalculation)
- **More flexibility** in the algorithm for profiles
 - Site metrics
 - 'any'/'all' services
 - Multiple FQAN
- Layered design
- **Common schema** to other Site Status Board
- Combining different UI: mywlcg and SUM
- Creation of reports from the UI
- For site administrators
 - Possibility to integrate in **local nagios**

SAM3 structure



SAM3 concepts (I)

- **Metrics:**
 - Measurement taken over a period of time
 - (Entity, Value, timestamp, validity)
 - Several types:
 - Status (e.g. JobSumit, status of CE a a site). Possible values: OK, WARN, CRITICAL, DOWNTIME, UNKNOWN
 - Numerical (e.g. # finished jobs per hour, pledged resources)
 - Mapping (e.g. name of the site for a given service, tier of a site)
- **Vofeed:**
 - XML provided by the experiments with the site names and services
 - Used to aggregate services into sites and naming conventions
- **Profile:**
 - Combination of metrics following an algorithm
 - One critical per experiment (used for the reports)

SAM3 concepts (II)

- **Downtime:**
 - Declaration of a site intervention
 - Entries collected from GOCDB and OIM
 - **ONLY SCHEDULED OUTAGES** are considered for Avl/Rel
- **Availability:**
 - Percentage of time that an instance is working over all known states
 - $Avl = (OK + WARN) / (OK + WARN + CRITICAL + DOWNTIME)$
- **Reliability:**
 - Percentage of time in working state over all known states except downtime
 - $Rel = (OK + WARN) / (OK + WARN + CRITICAL)$

Combination of metrics

OPERATION	PRIORITY
OR	     
AND	     
AND IF DATA	     
<m> OVERWRITE <n>	IF <N> !=  THEN <N> ELSE <M>
ANY <m>, <t>	OR of all instances in <m> that have the same value in <t>
ALL <m>, <t>	AND IF DATA of all instance in <m> that have the same value in <t>
FILTER <m>, <t>='v'	Take only the instances of <m> that have a value of 'v' in metric <t>

 OK
  WARNING
  CRITICAL
  DOWNTIME
  KNOWN
  DATA

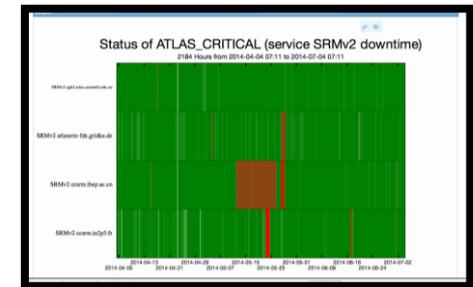
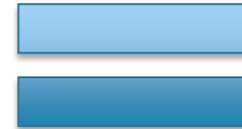
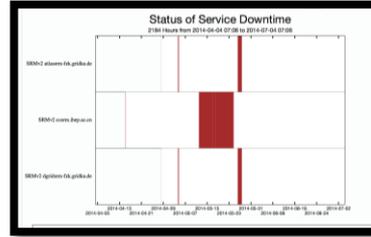
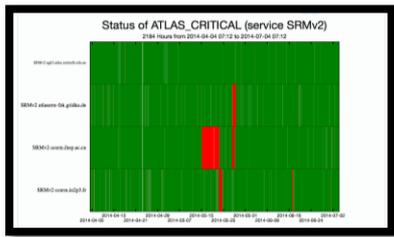
* Will be converted to  if there is at least one more metric

Profiles

- Combines status metrics of services and sites, given a new metric
 - And calculating several intermediate metrics
- These new metrics could be the input for other profiles
- Profiles can be corrected by the experiments
 - <https://twiki.cern.ch/twiki/bin/view/ArdaGrid/ProfileCorrections>

Overwritten values in SAM3

- Combining two metrics:
 - Used for the Downtimes and corrections

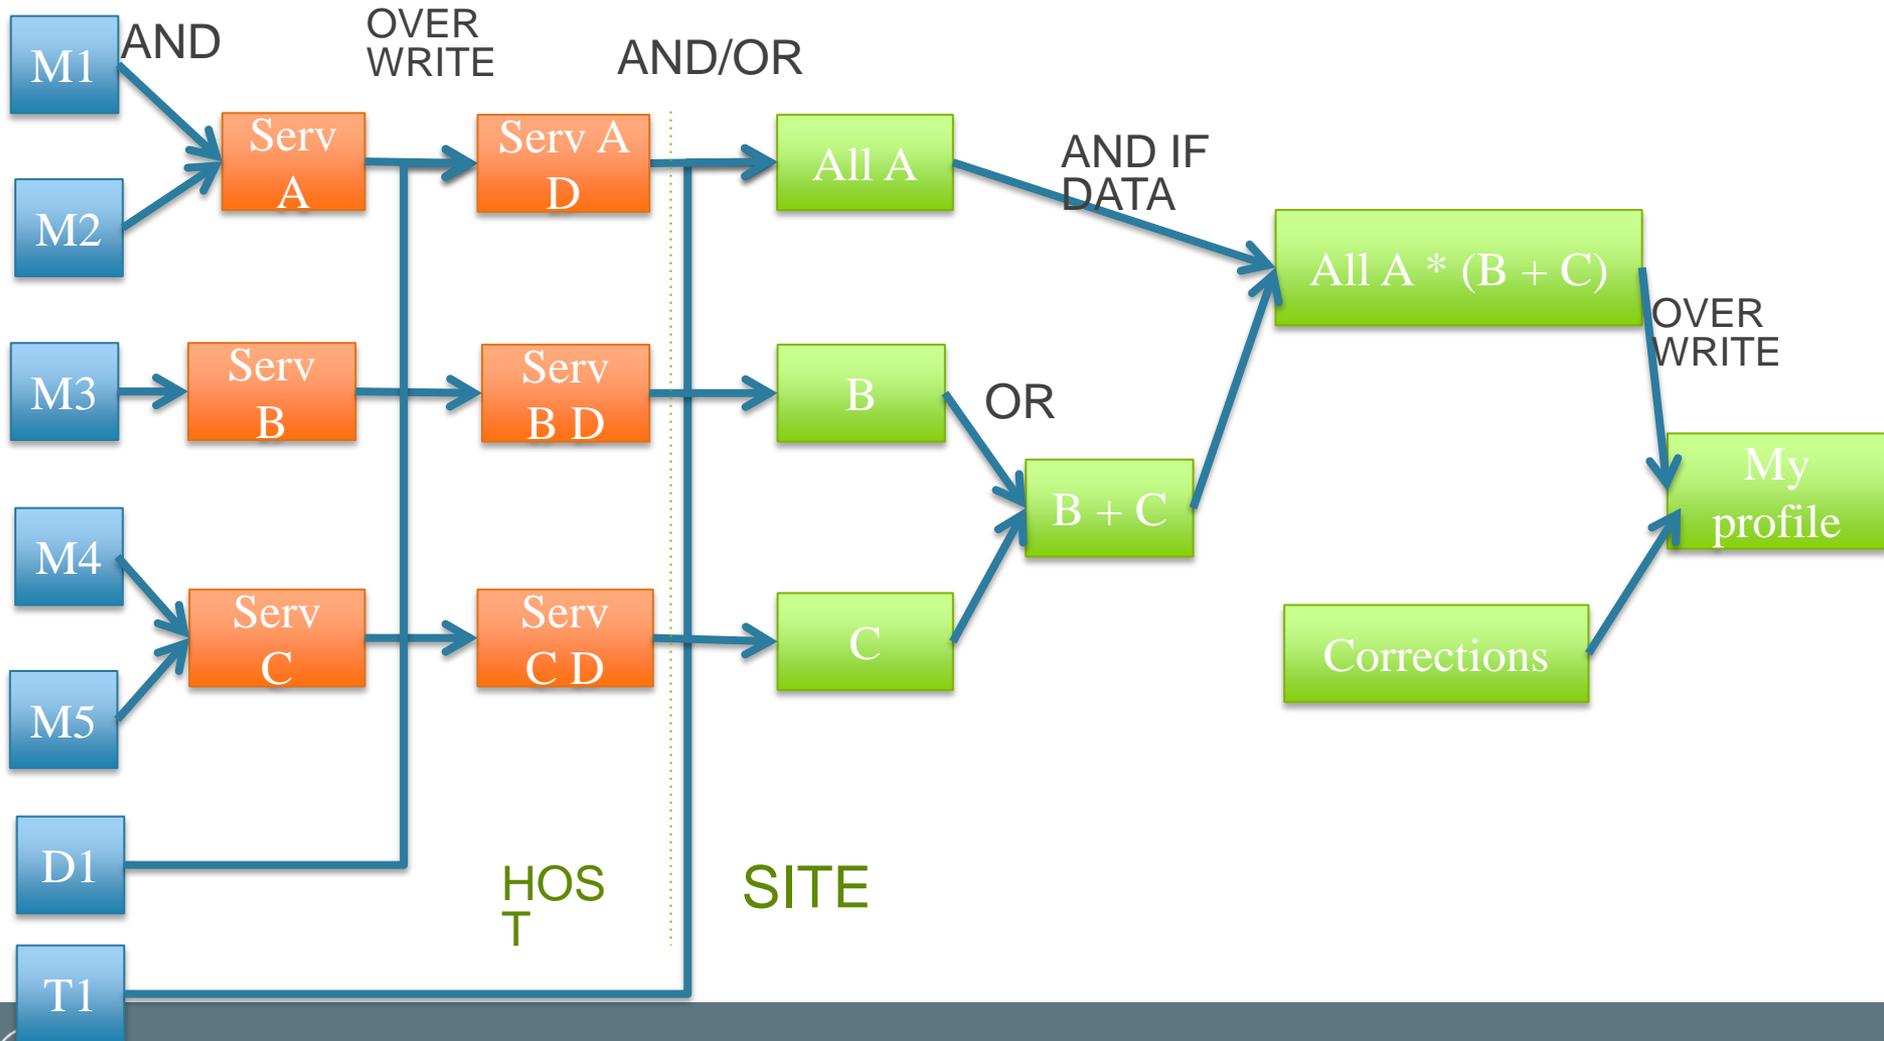


- Advantages:
 - Simple
 - Record of before/after & overwritten
 - Standard way of combining metrics
 - Can be visualize at any moment in time

Example of profile

- Name: My profile
- Algorithm: $\text{all } A * (B + C)$
- Metrics: $A \rightarrow M1, M2, B \rightarrow M3, C \rightarrow M4, M5$
- Topology: T1
- Service Downtime: D1

All A * (Any B + Any C)



Still under development...

- Include numerical metrics in algorithm
 - E.g. : efficiency > 50%
- ‘Not’ operator
 - E.g.: ‘not’ a T1
- Different weight for components:
 - 80% CE + 20% SRM

How does it look for ALICE

- <https://wlcg-sam-alice.cern.ch>



SAM VISUALIZATION | ALICE

[VO view](#) | [Latest Results](#) | [Historical View](#) | [Generate Report](#) | [Help](#)

Site-Groups	Profiles	Algorithm:
All Sites	ALICE_CRITICAL	CREAM-CE + ARC-CE
Sites	Service Flavours	Metrics
All Sites Athens Bari Birmingham	All Service Flavours ARC-CE CREAM-CE	All Metrics emi.ce.CREAMCE-JobSubmit (alice) emi.cream.CREAMCE-DirectJobSubmit (alice)

Show site metrics

Algorithm for calculating the Site and Service Availability

Legends for Metric Result Status

Status:	OK	WARNING	CRITICAL	UNKNOWN	MAINTENANCE
Legend:	OK	W	C	U	MT

Note: darkest colors: test is 0 - 12 hours old, ... lightest colors: test is more that 12 hours old

Legend	Metric Name
1	emi.ce.CREAMCE-JobSubmit (alice)
2	emi.cream.CREAMCE-DirectJobSubmit (alice)

[Link to data](#)

Sitename	Flavour	Host status in profile	Hosts	
Athens	CREAM-CE	CRITICAL	cream.afroditi.hellasgrid.gr	C
Bari	CREAM-CE	CRITICAL	cream-ce-2.ba.infn.it	C
		OK	cream-ce-3.ba.infn.it	OK
		OK	cream-ce-4.ba.infn.it	OK

Site Availability using ALICE_CRITICAL

12 hours from 2015/02/23 19:00:00 to 2015/02/24 06:00:00



23:02 23:02 23:02 23:02 23:02 24:02 24:02 24:02 24:02 24:02 24:02
 19:00 20:00 21:00 22:00 23:00 00:00 01:00 02:00 03:00 04:00 05:00 06:00



Current ALICE profiles

Profile Name	Algorithm	Flavour Metrics
ALICE_CRITICAL	CREAM-CE + ARC-CE	CREAM-CE: emi.cream.CREAMCE-DirectJobSubmit (alice) ; ARC-CE: emi.ce.CREAMCE-JobSubmit (alice) ;
ALICE_CRITICAL	CREAM-CE	CREAM-CE: emi.cream.CREAMCE-DirectJobSubmit (alice) ;
ALICE_MONALISA	(ARC-CE + CREAM-CE) * AliEn-CE * AliEn- VoBox-Test * AliEn-SE	AliEn-VoBox-Test: alien proxy (None) ; AliEn-CE: ServiceStatus (None) ; CREAM-CE: emi.cream.CREAMCE-DirectJobSubmit (alice) ; AliEn-SE: ADD (None),GET (None) ; ARC-CE: emi.ce.CREAMCE-JobSubmit (alice) ;
ALICE_PILOT	CREAM-CE + ARC-CE	CREAM-CE: emi.cream.glexec.CREAMCE-DirectJobSubmit (/alice/Role_pilot),emi.cream.glexec.WN-gLExec (/alice/Role_pilot) ; ARC-CE: emi.ce.glexec.CREAMCE-JobSubmit (/alice/Role_pilot),emi.cream.glexec.WN-gLExec (/alice/Role_pilot) ;
ALICE_PILOT	ARC-CE+CREAM-CE	CREAM-CE: emi.cream.CREAMCE-DirectJobSubmit (alice) ; ARC-CE: emi.ce.CREAMCE-JobSubmit (alice) ;

And under development

Profile Name ▲	Algorithm ◇	Flavour Metrics
ALICE_CRITICAL	CREAM-CE	CREAM-CE: emi.cream.CREAMCE-DirectJobSubmit (alice) ;
ALICE_TEST	ARC-CE+CREAM-CE	CREAM-CE: emi.cream.CREAMCE-DirectJobSubmit (alice) ; ARC-CE: emi.ce.CREAMCE-JobSubmit (alice) ;
ALICE_TEST2	(ARC-CE + CREAM-CE) * ((AliEn-CE * Vobox-test) + ARC-CE + CREAM-CE) * AliEn-SE	Vobox-test: emi.ce.glexec.CREAMCE-JobSubmit (/alice/Role_pilot) ; AliEn-CE: emi.cream.CREAMCE-DirectJobSubmit (alice) ; CREAM-CE: emi.cream.CREAMCE-DirectJobSubmit (alice) ; AliEn-SE: emi.ce.CREAMCE-JobSubmit (alice) ; ARC-CE: emi.ce.CREAMCE-JobSubmit (alice) ;
ALICE_TEST4	@SitesWithCE *AliEn-SE	AliEn-SE: emi.ce.CREAMCE-JobSubmit (alice) ; SitesWithCE: ALICE/TEST ;

An example from ATLAS

Profile Name ▲	Algorithm ▼	Flavour Metrics
ATLAS_AnalysisAvailability	@SiteAvailability	SiteAvailability: ASAP ;
ATLAS_CRITICAL	(OSG-CE + CREAM-CE + ARC-CE) * (&SRMv2 + &OSG-SRMv2)	CREAM-CE: org.sam.CONDOR-JobSubmit (/atlas/Role_lcgadmin),org.atlas.WN-swspa SRMv2: org.atlas.SRM-VOGet (/atlas/Role_production),org.atlas.SRM-VOPut (/atlas/Ro OSG-SRMv2: org.atlas.SRM-VOPut (/atlas/Role_production),org.atlas.SRM-VOGet (/atl OSG-CE: org.atlas.WN-swspace (/atlas/Role_lcgadmin),org.atlas.WN-swspace (/atlas/R ARC-CE: org.atlas.WN-swspace (/atlas/Role_pilot),org.sam.CONDOR-JobSubmit (/atlas
ATLAS_DDM_CRITICAL	(OSG-CE+CREAM-CE+ARC-CE)*(&SRMv2+&OSG-SRMv2)	SRMv2: org.atlas.DDM-Get (/atlas/Role_production),org.atlas.DDM-Put (/atlas/Role_pro CREAM-CE: org.atlas.WN-swspace (/atlas/Role_lcgadmin),org.sam.CONDOR-JobSubn OSG-CE: org.sam.CONDOR-JobSubmit (/atlas/Role_lcgadmin),org.atlas.WN-swspace (r OSG-SRMv2: org.atlas.DDM-Put (/atlas/Role_production),org.atlas.DDM-Get (/atlas/Rol ARC-CE: org.sam.CONDOR-JobSubmit (/atlas/Role_lcgadmin),org.atlas.WN-swspace (
ATLAS_GENERAL	(OSG-CE+CREAM-CE+ARC-CE)*(&SRMv2+&OSG-SRMv2)	CREAM-CE: org.atlas.WN-FrontierSquid (/atlas/Role_lcgadmin),org.atlas.WN-swspace (/atlas/Role_lcgadmin),org.atlas.HC-Analysis (atlas),org.atlas.WN-cvmfs (/atlas/Role_lcg SRMv2: org.atlas.SRM-VOGetATLASDATADISK (/atlas/Role_production),org.atlas.SRM VODEIATLASGROUPDISK (/atlas/Role_production),org.atlas.SRM-VOGet (/atlas/Role_ /Role_production),org.atlas.SRM-VOPutATLASSCRATCHDISK (/atlas/Role_production) /Role_production),org.atlas.SRM-GetATLASInfo (/atlas/Role_production),org.atlas.SRM- VOGetATLASLOCALGROUPDISK (/atlas/Role_production) ; OSG-SRMv2: org.atlas.SRM-VODEIATLASSCRATCHDISK (/atlas/Role_production),org /Role_production),org.atlas.SRM-VOGetATLASSCRATCHDISK (/atlas/Role_production) /Role_production),org.atlas.SRM-GetATLASInfo (/atlas/Role_production),org.atlas.SRM- VOPutATLASLOCALGROUPDISK (/atlas/Role_production),org.atlas.SRM-VOPut (/atlas VOGetATLASLOCALGROUPDISK (/atlas/Role_production),org.atlas.SRM-VOGet (/atlas OSG-CE: org.atlas.WN-swspace (/atlas/Role_lcgadmin),org.atlas.CE-Pilot-submission (; (atlas),org.sam.CONDOR-JobSubmit (/atlas/Role_pilot),org.sam.CONDOR-JobSubmit (r ARC-CE: org.atlas.CE-Pilot-submission (atlas),emi.cream.glexec.WN-gLExec (/atlas/Ro /Role_lcgadmin),org.sam.CONDOR-JobSubmit (/atlas/Role_lcgadmin),org.sam.CONDC

ALICE issues

- **Critical profile: ARC-CE OR CREAM-CE**
 - Sites without those services do not appear (at the moment, 19 sites)
 - **Solution:** include AliEn-CE or VOBOx in algorithm
- **Storage ignored from availability**
 - **Solution:** include AliEn-SE in algorithm
- **NDGF does not appear as a single T1**
 - The individual sites appear as 'Nordugrid Sites', and are ignored in the reports
 - **Solution:** modify Vofeed to define NDGF-T1
- **Timeout in JobSubmit**
 - If the site is busy, SAM jobs might wait too long on the queue
 - **Solution:** Use different credentials for SAM test
- **Unused profiles:**
 - **Solution:** delete unused profiles: ALICE_GENERAL, ALICE_TEST

Conclusion

- SAM calculates the site and service availability and reliability
- Creates WLCG monthly Availability reports
- SAM3 offers new functionality and flexibility
 - Definition of algorithm
 - Metrics and topology defined by experiments
- We need to tune the new functionality!
- <https://wlcg-sam.cern.ch>



New Nagios Plugin

- Bash script downloads JSON files for each of the experiment's SAM (SUM) dashboard
- After parsing it, it get status of each individual check in a given profile
- It sends status to local Nagios system

Slide from
Jordi Casals
(PIC)
<http://cern.ch/go/N8r7>

Link to data

Dashboard JSON

```
{
  "data": {
    "siteName": [
      "T1_ES_PIC"
    ],
    "results": [
      {
        "sitename": "T1_ES_PIC",
        "flavours": [
          {
            "hosts": [
              {
                "metric": [
                  {
                    "status": "OK",
                    "timestamp": "2014-04-09T09:43:15Z",
                    "age": 0,
                    "Abbr": 2,
                    "metric_name": "emi.cream.CREAMCE-JobSubmit"
                  }
                ]
              }
            ]
          }
        ]
      }
    ]
  }
}
```

Nagios Plugins LCG SAM

```
...
...
# Get JSON content using input
args
json=$(wget -qO - "http://dashb-
$vo-
sum.cern.ch/dashboard/request.py/1
atestresultssmry-sumjson?
profile=$profile&flavour=$flavour&
site=$site")
...
...
```

PIC Nagios

```
WARNING (for 0d 23h 55m 53s)
CREAM-CE: WARNING

HOSTS
=====
CREAM-CE
ce07.pic.es: WARNING
ce08.pic.es: WARNING
ce09.pic.es: WARNING
ce10.pic.es: WARNING
ce11.pic.es: WARNING

PROBLEMS
=====
CREAM-CE - ce07.pic.es - org.cms.WN-xrootd-access: W
CREAM-CE - ce08.pic.es - org.cms.WN-xrootd-access: W
CREAM-CE - ce09.pic.es - org.cms.WN-xrootd-access: W
CREAM-CE - ce10.pic.es - org.cms.WN-xrootd-access: W
CREAM-CE - ce11.pic.es - org.cms.WN-xrootd-access: W
```