

Site usability metrics

OPS and VO-specific tests

Maria Girone
Andrea Sciabà

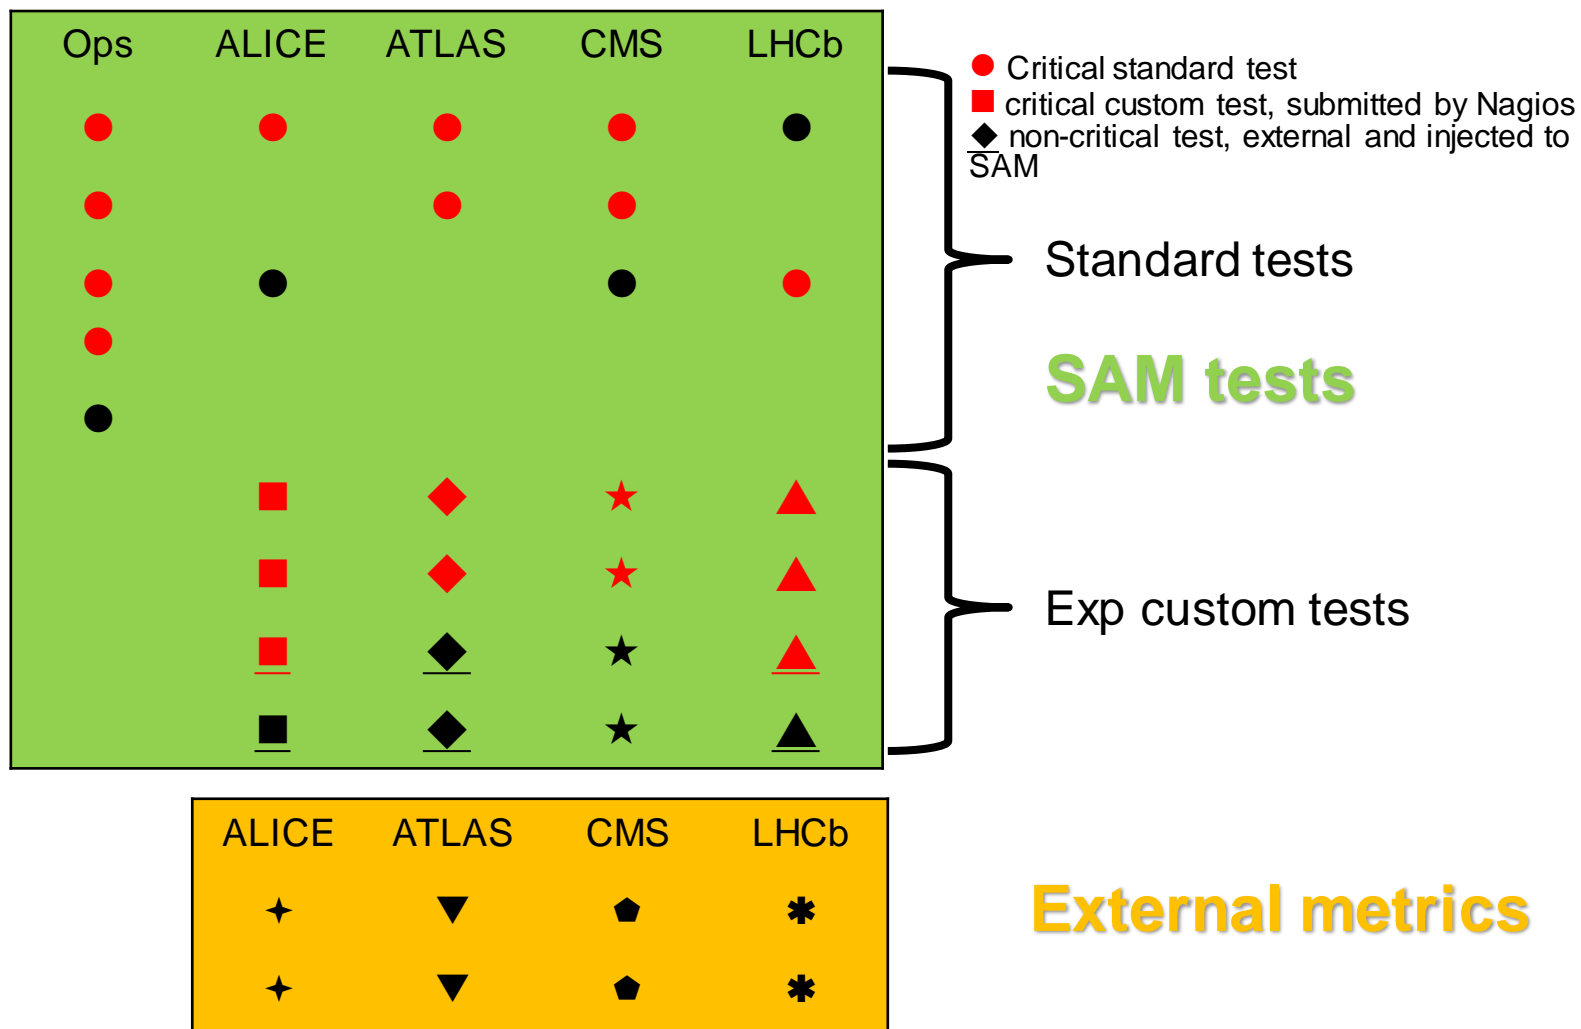
WLCG Operations TEG Workshop
12/12/2011

- Summarize the current status
- Describe limitations
- Suggest different alternative ways to measure the site usability




- **Standard tests**
 - Developed by EGI and bundled with the **official SAM probes**
- **OPS tests**
 - Tests run with an OPS proxy; they are all standard tests
- **Experiment custom tests**
 - Custom tests **developed** by an experiment and **run** with an experiment proxy
 - Submitted by a **Nagios** box, or
 - Submitted by other test systems and **published** in SAM
- **Critical tests**
 - Tests whose result is used to determine the state of a service instance (and hence the service and the site availability)
- **Other quality metrics** outside the SAM framework
 - E.g. quality of data transfers, success rate of typical jobs (Hammercloud, Job Robot, DIRAC), etc.

- **Gridview** availability
 - Used primarily by the WLCG **management**
 - Available for all the **WLCG VOs** and **OPS**
- **Dashboard** availability
 - Calculated by a Dashboard application (but the algorithm is the same)
 - Used by the **experiment operations** teams
 - May use **different** critical tests
- **ACE** availability
 - Not yet in production for the experiments (pending validation) but so for OPS
 - Will soon calculate **all availabilities** (possibly many per VO)




Tier-1/0 Site Availability VO:CMS (Daily Report)
(Click on the Graph below to see Availability of Individual Services at the Site)



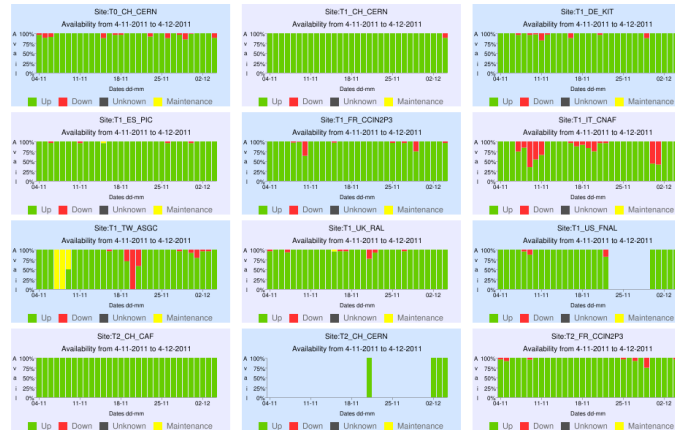
GridView

Site Availability
30 Days from 2011-11-05 to 2011-12-05



Dashboard

Site Availability Profile:CMS_CRITICAL (Daily Report)
(Click on the Graph below to see Availability of Individual Services at the Site)
Availability from 4-11-2011 to 4-12-2011



myWLCG

Andrea Sciabà

- Availability: technical definition
 - The fraction of time a site passes all the SAM critical tests
 - Provided for OPS and the experiments
 - Can be very different as it depends on choice of critical tests
- Usability
 - The fraction of time a site is usable by an experiment
 - It should be **easy** to interpret
 - It should be a **realistic** representation of the site usability

- OPS tests run on all EGI sites
 - \Rightarrow OPS availability
- VO tests run on all WLCG sites
 - Some **standard**, some VO **custom**
- Gridview calculates the VO availability
 - According to a set of **critical** tests
- The Dashboard calculates another VO availability
 - According to **another set** of critical tests
- **Why so many availabilities? Is it really necessary?**

Test name	Service
org.sam.CE-JobSubmit	CE
org.sam.CREAMCE-JobSubmit	CREAM-CE
hr.srce.CREAMCE-CertLifetime	CREAM-CE
hr.srce.CADist-Check	CREAM-CE, CE
hr.srce.GRAM-CertLifetime	CE
hr.srce.SRM2-CertLifetime	CE
org.sam.WN-Bi	CREAM-CE, CE
org.sam.WN-Csh	CREAM-CE, CE
org.sam.WN-Rep	CREAM-CE, CE
org.sam.WN-SoftVer	CREAM-CE, CE
org.sam.SRM-GetURLs	SRMv2
org.sam.SRM-GetTURLs	SRMv2
org.sam.SRM-LsDir	SRMv2
org.sam.SRM-Put	SRMv2
org.sam.SRM-Ls	SRMv2
org.sam.SRM-Get	SRMv2
org.sam.SRM-Del	SRMv2
org.bdii.Entries	Site-BDII
org.gstat.SanityCheck	Site-BDII

Test name	Service
org.osg.certificates.cacert-expiry	OSG-CE
org.osg.general.ping-host	OSG-CE
org.osg.globus.gram-authentication	OSG-CE
org.osg.general.osg-directories-CE-permissions	OSG-CE
org.osg.general.osg-version	OSG-CE
org.osg.srm.srmcping	OSG-SRMv2, OSG-BestmanXrootd
org.osg.srm.srmcp-readwrite	OSG-SRMv2, OSG-BestmanXrootd
org.osg.globus.gridftp-simple	OSG-GridFtp
org.arc.ARC-STATUS	ARC-CE
org.arc.csh	ARC-CE
org.arc.GRIDFTP	ARC-CE
org.arc.LFC	ARC-CE
org.arc.SRM	ARC-CE
org.arc.Jobsubmit	ARC-CE
org.arc.AUTH	ARC-CE
org.arc.CA-VERSION	ARC-CE
org.arc.SW-VERSION	ARC-CE

- Only standard tests (by definition)
- Only critical tests in these tables (other tests are submitted)

Test name	Service	Critical in GV	Critical in DB
org.sam.CREAMCE-DirectJobSubmit	CREAMCE	Y	Y
org.sam.CREAMCE-JobSubmit	CREAMCE	Y	Y
org.sam.WN-sft-vo-swdir	CREAMCE	Y	Y
VOBOX-org.alice.vobox-PM	VOBOX	Y	Y
VOBOX-org.alice.vobox-PR	VOBOX	Y	Y
VOBOX-org.alice.vobox-PSR	VOBOX	Y	Y
VOBOX-org.alice.vobox-SA	VOBOX	Y	Y
VOBOX-org.alice.vobox-UPR	VOBOX	Y	y

- Standard tests in red
- Only one set of critical tests
- All tests are critical

Test name	Service	Critical in GV	Critical in DB
org.sam.CE-JobSubmit	CE	Y	Y
org.sam.CREAMCE-JobSubmit	CREAMCE	Y	Y
org.atlas.WN-swtag	CE	Y	Y
org.atlas.WN-swspace	CE	Y	Y
org.atlas.WN-FrontierSquid	CE	N	N
org.atlas.WN-LocalFileAccess	CE	N	N
org.atlas.WN-gangarobot_panda	CE	N	N
org.atlas.WN-gangarobot_wms	CE	N	N
org.atlas.GetATLASInfo	SRM	N	N
org.atlas.SRM-VOPut	SRM	Y	Y
org.atlas.SRM-VOGet	SRM	Y	Y
org.atlas.SRM-VODel	SRM	Y	Y
org.atlas.SRMLs	SRM	N	N
org.atlas.SRMLsDir	SRM	N	N

- Standard tests in red
- Only **one set** of critical tests
- Only basic infrastructure tests are critical
- Other custom tests run but not critical

Test name	Service	Critical in GV	Critical in DB
org.sam.CE.JobSubmit	CE	Y	Y
org.sam.CREAMCE.JobSubmit	CREAMCE	Y	Y
org.sam.glexec.CE-JobSubmit	CREAMCE	N	N
org.sam.glexec.CREAMCE-JobSubmit	CREAMCE	N	N
org.cms.WN-basic	CE	N	Y
org.cms.WN-swinst	CE	N	Y
org.cms.WN-mc	CE	N	Y
org.cms.WN-analysis	CE	N	Y
org.cms.WN-frontier	CE	N	Y
org.cms.WN-squid	CE	N	Y
org.cms.glexec.WN-gLExec	CE	N	N
org.cms.SRM-GetPFNFromTFC	SRM	Y	Y
org.cms.SRM-VOPut	SRM	Y	Y
org.cms.SRM-VOGet	SRM	Y	Y
org.cms.SRM-VODel	SRM	N	N
org.cms.SRM-VOGetURLs	SRM	N	N
org.cms.SRM-VOLs	SRM	N	N
org.cms.SRM-VOLsDir	SRM	N	N

- Standard tests (or equivalent to) in **red**
- **Two sets** of critical tests: minimal for Gridview, complete for Dashboard
- Only very basic infrastructure tests are critical in Gridview
- Other custom tests run but critical only for Dashboard

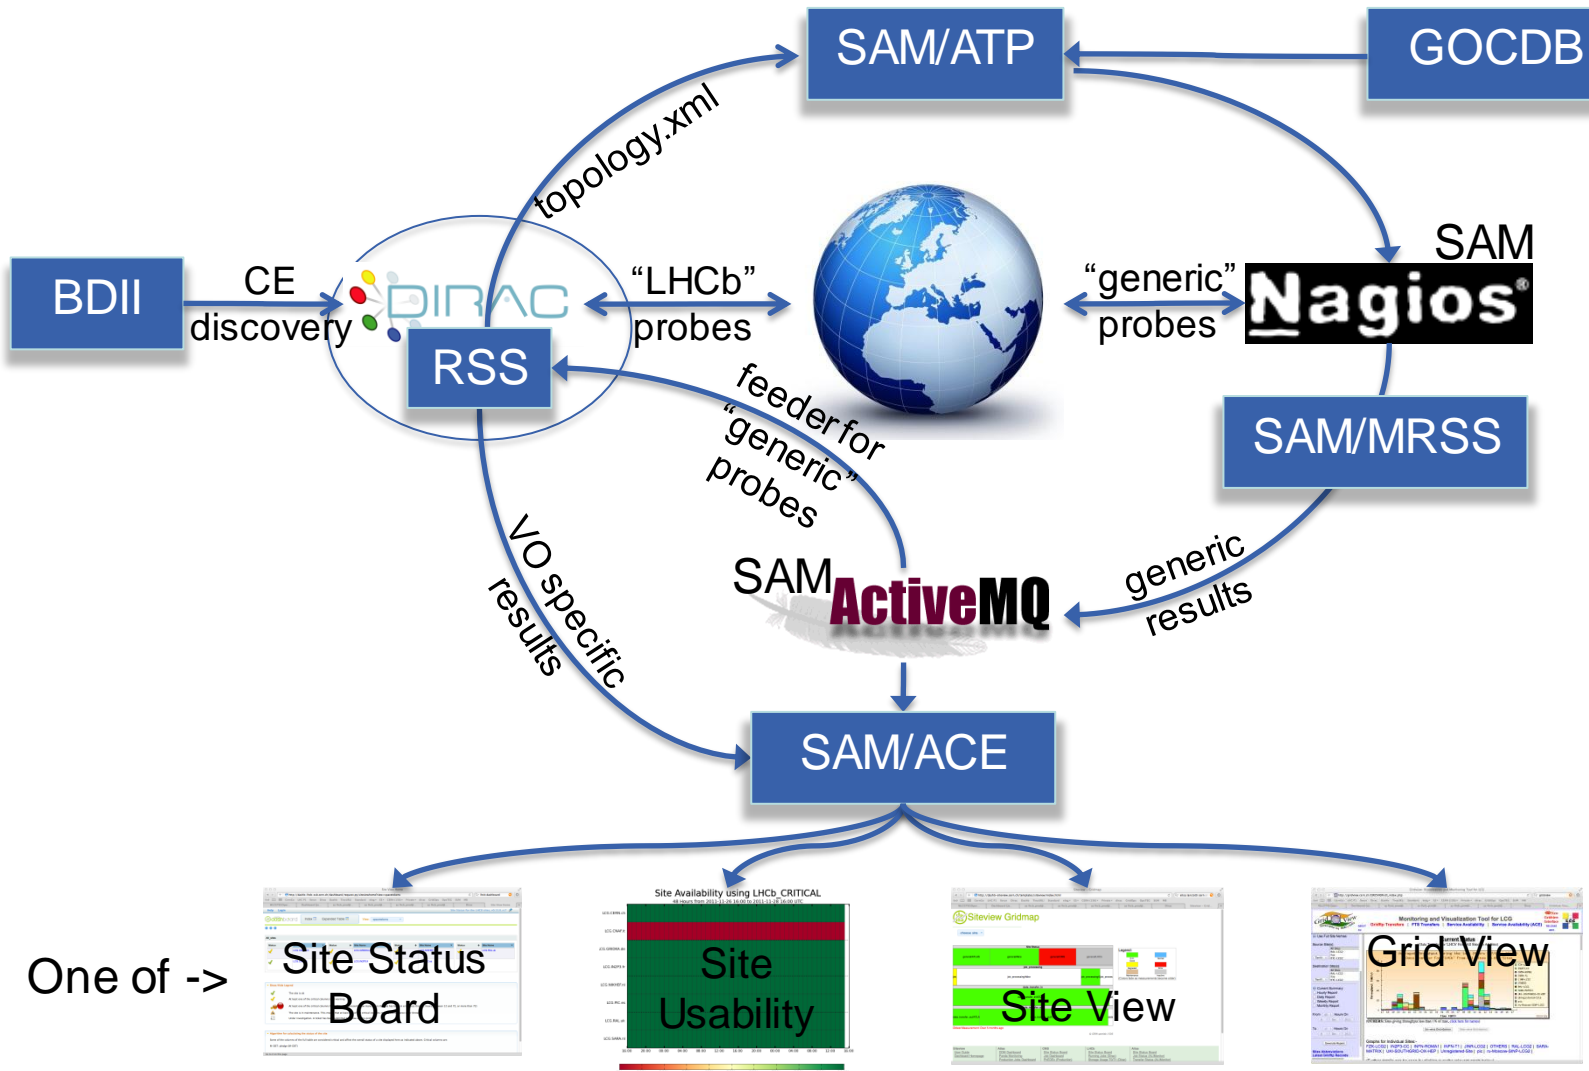
Andrea Sciabà

Test name	Service	Critical in GV	Critical in DB
org.lhcb.LFC-Ping	LFC	N	N
org.lhcb.LFC-Read	LFC	N	N
org.lhcb.LFC-Readdir	LFC	N	N
org.lhcb.LFC-Replicate	LFC	N	N
org.sam.CE-JobSubmit	CE	Y	Y
org.sam.CREAMCE-JobSubmit	CREAMCE	N	N
org.sam.glexec.CE-JobSubmit	CE	N	N
org.sam.CREAMCE.DirectJobSubmit	CREAMCE	N	N
org.lhcb.WN-CondDB	CE	N	N
org.lhcb.WN-lhcb-FileAccess	CE	N	N
org.lhcb.WN-sft-brokerinfo	CE	Y	Y
org.lhcb.WN-sft-csh	CE	Y	Y
org.lhcb.WN-sft-lcg-rm-gfal	CE	Y	Y
org.lhcb.WN-sft-voms	CE	Y	Y
org.lhcb.WN-sft-vo-swdir	CE	Y	Y
org.lhcb.WN-CE-lhcb-job-Boole	CE	N	N
org.lhcb.WN-CE-lhcb-job-Brunel	CE	N	N
org.lhcb.WN-CE-lhcb-job-DaVinci	CE	N	N
org.lhcb.WN-CE-lhcb-job-Gauss	CE	N	N
org.lhcb.WN-CE-lhcb-job-install	CE	N	N
org.lhcb.WN-CE-lhcb-job-os	CE	N	N
org.lhcb.WN-CE-lhcb-job-queues	CE	N	N
org.lhcb.SRM-Dirac-Unit	SRM	N	N
org.lhcb.SRM-GetLHCbInfo	SRM	Y	Y
org.lhcb.SRM-VODel	SRM	Y	Y
org.lhcb.SRM-VOGet	SRM	Y	Y
org.lhcb.SRM-VOLs	SRM	Y	Y
org.lhcb.SRM-VOLsDir	SRM	Y	Y
org.lhcb.SRM-VOPut	SRM	Y	Y




- OPS tests
 - Sites are very familiar with them and look at them
 - Not useful to experiments as they may test different services than those used by them and test results may depend on the VO of the proxy
- Gridview availabilities are too good
 - Too few tests are critical!
 - Largely uncorrelated to the “real” usability
 - Included in a monthly report to the WLCG
- On the other hand VO custom tests may be obscure for sites
- There are site quality metrics that are not measured by SAM
 - E.g. CMS Job Robot, HammerCloud, data transfer quality etc.

- Experiments run a subsample of the “standard” tests with their credentials
 - Note that the only standard tests which make sense to be run with an experiment proxy are
 - `org.sam.*CE-JobSubmit`
 - `org.sam.SRM-*`
 - (LFC tests)
- Limitations
 - Success of the standard tests is necessary but **not** sufficient for the site to be usable
 - Easy? **Yes**, sites are very familiar with them
 - Realistic? **No**

- To be a realistic estimate all relevant metrics must be expressed as SAM tests
 - In the real world many relevant metrics come from other sources
- The resulting set of SAM tests will strongly depend on experiment-specific functionality
- There will be many tests not trivial to interpret for a non-expert
- ⇒ Realistic? **Yes**
- ⇒ Easy to understand? **Yes**, if everything is properly documented



One of ->



- LHCb Dirac Resource Status System (RSS)
 - Aggregates many different monitoring sources
 - Decides with LHCb policies how to operate
 - LHCb Nagios Ops (generic) tests are testing the basic site functionality / availability
 - Results are fed back to the RSS for further decisions
 - LHCb RSS decisions/results are fed into ACE
 - Further usage for availability calculation
 - Two result types for “LHCb” and “Sites”
 - Visualization tool to be decided – should be one
- ⇒ A **realistic** view of LHCb grid operations

- Use an **estimator** of the site **usability** which is a function of arbitrary metrics (possibly including availability)
 - $U(site) = F(M_1, \dots, M_n)$
 - $U(site) \in \{good, bad\}$
 - Calculate U daily
 - Use metrics related to high-level functionality (e.g. transfers, job success, ...)
 - Show U together with the original metrics (M, \dots, M_n)
 - Provide adequate documentation

- Requires to read the documentation to fully understand all the metrics
- Metrics may be different for different experiments
- Easy? **Yes**, if high level information is enough and documentation is good
- Realistic? **Yes**





T2_AT_Vienna																						
Site Readiness Status: R R R R W R R R R R R W W R R																						
Daily Metric: O O O O O O O O E E E O O O O O O O E E O O																						
Maintenance:	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up		
Maintenance (Topology):	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up		
Job Robot:	99%	100%	100%	100%	100%	98%	97%	100%	66%	16%	100%	99%	98%	99%	100%	100%	100%	73%	100%	91%	100%	
SAM Availability:	100%	100%	100%	100%	100%	100%	100%	72%	12%	100%	84%	100%	100%	100%	100%	100%	100%	100%	100%	98%	100%	
Good T2 links from T1s:	0/8	8/8	8/8	8/8	8/8	8/8	8/8	0/8	0/8	0/8	0/8	0/8	0/8	0/8	8/8	7/8	8/8	2/8	8/8	8/8	8/8	
Good T2 links to T1s:	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	6/6	
Active T2 links from T1s:	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	
Active T2 links to T1s:	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	
	18	19	20	21	22	23	24	25	26	27	28	29	30	01	02	03	04	05	06	07	08	09
	Nov							Dec														

Report made on 2011-12-09 03:30:01 (UTC)

* = Due to operational errors, the metric has been corrected manually (!= \$SB).

--- = Errors on weekends are ignored on Site Readiness computation for T2s [-info]

"Site Readiness Status" as defined in Site Commissioning Twiki:

- R = READY
- W = WARNING
- NR = NOT-READY
- SD = SCHEDULED-DOWNTIME

"Daily Metric" as boolean AND of all individual metrics:

- O = OK (All individual metrics above Site Commissioning Thresholds; "n/a" ignored)
- E = ERROR (Some individual metrics below Site Commissioning Thresholds)
- SD = SCHEDULED-DOWNTIME

- INDIVIDUAL METRICS -

"Scheduled Downtimes": site maintenances

- Up = site is not declaring Scheduled-downtime
- SD = full-site in SD OR all CM \$ SE(s) in SD OR all CM \$ CE(s) in SD
- ~ = Some SE or CE services (not all) Downtime
- UD = full-site in UD OR all CM \$ SE(s) in UD OR all CM \$ CE(s) in UD

"SAM Availability":

- O = SAM availability is ≥ 90%
- E = SAM availability is < 90%

"Active T2 links from T1s":

- O = Site has ≥ 4 DDT-commissioned links from T1 sites
- E = Otherwise

"Job Robot":

- O = Job success rate is ≥ 80%
- E = Job success rate is < 80%
- = Jobs submitted but not finished
- n/a = Job success rate is n/a

"Active T2 links to T1s":

- O = Site has ≥ 2 DDT-commissioned links to T1 sites
- E = Otherwise

"Good Links":

- O = at least half of links have "good" transfers (i.e. with transfer quality > 50%)
- E = Otherwise

Custom display


Index Expanded Table

Show 200 entries Copy Print Save view: commission Search...

Site Name	SiteComm JR	JR (Dashboard)	HC	Site Nagios availability	Commissioned Links (expand this column)	Good links combined (expand this column)	SiteReadiness Status	Is SiteInSiteDB?	TopologyMaintenances
T0_CH_CERN	100	100.0	100.0	96	n/a	n/a	n/a	true	1 SRMv2 . 6 CREAM-CE . 4 CE
T1_CH_CERN	100	100.0	100.0	96	2/5 combined	5/7 combined	R	true	1 SRMv2 . 6 CREAM-CE . 4 CE
T1_DE_KIT	100	100.0	100.0	100	3/5 combined	5/7 combined	R	true	1 SRMv2 . 5 CREAM-CE
T1_ES_PIC	100	100.0	97.4	100	3/5 combined	5/7 combined	R	true	1 SRMv2 . 4 CREAM-CE
T1_FR_CCIN2P3	100	100.0	99.9	88	3/5 combined	5/7 combined	SD	true	1 SRMv2 . 2 CREAM-CE
T1_IT_CNAF	77	76.3	86.1	100	3/5 combined	5/7 combined	W	true	OUTAGE SCHEDULED 2/6 CREAM-CE down
T1_TW_ASGC	100	100.0	98.5	96	3/5 combined	5/7 combined	R	true	1 SRMv2 . 4 CREAM-CE . 4 CE
T1_UK_RAL	100	100.0	100.0	100	3/5 combined	5/7 combined	R	true	1 SRMv2 . 3 CREAM-CE
T1_UK_RAL_Disk	n/a	n/a	n/a	n/a	3/5 combined	n/a	n/a	true	n/a
T1_US_FNAL	100	100.0	100.0	100	3/5 combined	5/7 combined	R	true	1 SRMv2 . 3 CE
T1_US_FNAL_Disk	n/a								
T2_AT_Vienna	100								2 . 1 CREAM-CE
T2_BE_IHE	100								1 CREAM-CE
T2_BE_UCL	98	98.5	99.7	100	2/5 combined	2/7 combined	R	true	1 SRMv2 . 1 CE
T2_BR_SPRACE	100	100.0	95.6	100	2/5 combined	2/7 combined	R	true	1 SRMv2 . 1 CE
T2_BR_UERJ	99	100.0	99.7	100	2/5 combined	2/7 combined	R	true	1 SRMv2 . 2 CE
T2_CH_CAF	n/a	n/a	n/a	n/a	n/a	n/a	n/a	true	6 CREAM-CE . 4 CE
T2_CH_CERN	n/a	n/a	n/a	n/a	2/5 combined	n/a	n/a	true	1 SRMv2 . 6 CREAM-CE . 4 CE
T2_CH_CSCS	94	95.4	91.7	100	2/5 combined	2/7 combined	R	true	1 SRMv2 . 2 CREAM-CE
T2_CN_Beijing	100	100.0	99.0	100	2/5 combined	2/7 combined	R	true	OUTAGE UNSCHEDULED 1/1 SRMv2 down 1/1 CREAM-CE

Showing 1 to 123 of 123 entries DB query took 0.0288 s First Previous 1 Next Last

Trivial to create similar tables for all experiments



Input metrics

Usability

- There is **no subset** of the standard tests which is representative of the site **usability** for an experiment. Therefore:
 - The standard tests must be **expanded**, or
 - The VO **custom** tests must be used
- But new standard tests should be general and in most cases VO custom tests **cannot be generalized**
 - The largest common set would be too small
- Therefore, **VO custom tests are necessary**

- Is availability the best metric to measure usability or something else can be defined?
 - Must be both **realistic** and **understandable**

Advantages	Disadvantages
Availability	
Compatible with GridView	SAM developed for functional tests (OK/WARNING/ERROR): not all metrics fit in (e.g. job success rate)
All metrics published in Nagios	Tests of complex workflows must be artificially converted into tests of “service instances”
Usability	
Easier to combine heterogeneous metrics	Requires a visualisation tool other than GridView (e.g. the Site Status Board)
Easier to drill down to input metrics	Combination algorithms may differ from one experiment to another