



WLCG Tier1 [Performance] Metrics

~ ~ ~

Points for Discussion

Jamie.Shiers@cern.ch

~ ~ ~

WLCG GDB, 8th July 2009

The Perennial Question

- During this presentation and discussion we will attempt to sharpen and answer the question:
- ***How can a Tier1 know that it is doing OK?***
- We will look at:
 - What we can (or do) measure (automatically);
 - What else is important – but harder to measure (at least today);
 - How to understand what “OK” really means...

Resources

- In principle, we know what resources are pledged, can determine what are actually installed(?) and can measure what is currently being used;
- If installed capacity is significantly(?) lower than pledged, this is an anomaly and site in question **“is not doing ok”**
- But actual utilization may vary – and can even exceed – “available” capacity for a given VO (particularly CPU – less or unlikely for storage(?))
- **This should also be signaled as an anomaly to be understood (it is: poor utilization over prolonged periods impacts future funding, even if there are good reasons for it...)**

Services

- Here we have extensive tests (OPS, VO) coupled with production use
 - A “test” can pass, which does not mean that experiment production is not (severely) impacted...
 - Some things are simply not realistic or too expensive to test...
- But again, significant anomalies should be identified and understood
- Automatic testing is one measure: GGUS tickets another (# tickets, including alarm, time taken for their resolution)
 - This can no doubt be improved iteratively; additional tests / monitoring added (e.g. tape metrics)
- A site which is “green”, has few or no tickets open for > days | weeks, and no “complaints” at operations meeting is doing ok, surely?
- **Can things be improved for reporting and long-term traceability? (expecting the answer YES)**

The Metrics...

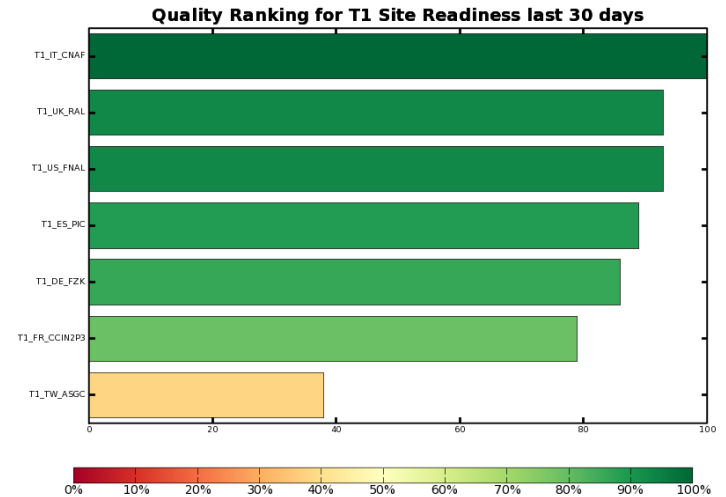
- For STEP'09 – as well as at other times – explicit metrics have been set against sites and for well defined activities
- Can such metrics allow us to “roll-up” the previous issues into a single view?
- If not, what is missing from what we currently do?
- Is it realistic to expect experiments to set such targets:
 - During the initial period of data taking? (Will it be known at all what the “targets” actually are?)
 - In the longer “steady state” situation? Processing & reprocessing? MC production? Analysis?? (largely not T1s...)
- Probable answer: only if it is useful for them to monitor their own production (which it should be..)

So...

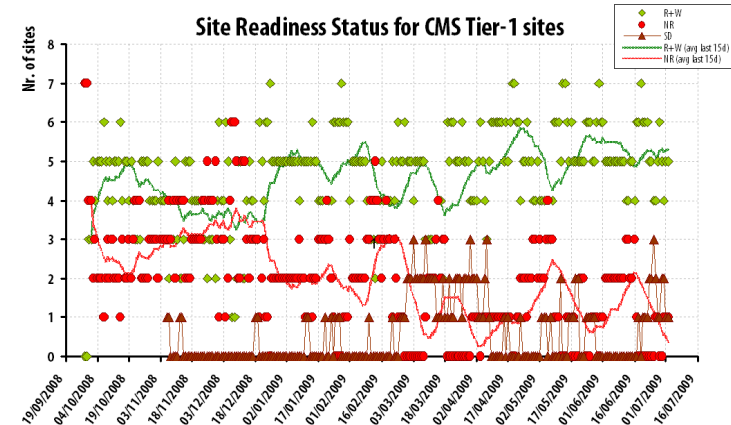
- **Definition: a site is “doing ok” if (and only if?):**
 1. It has provided (usable) resources that match those pledged & requested;
 2. The services are running smoothly, pass the tests and meet reliability and availability targets;
 3. “WLCG operations” metrics on handling scheduled and unscheduled service interruptions and degradations are met (more automation / reporting required here);
 4. It is meeting or exceeding metrics for “functional blocks” (this latter will require some work: is it “fair” to expect experiments to define these? Must be as simple as “autopilot settings”)
- **And in practice? Some examples from CMS...**

Example from CMS: Site Readiness

- Regularly used to evaluate CMS sites, results are shown at the weekly Facilities meeting
- Several metrics taken into account
 - SAM availability must be > 90% for Tier-1
 - Job Robot success rate must be > 90% for Tier-1
 - Submitting about 600 jobs / (day*site)
 - Number of commissioned links must be high enough
- Other metrics, not included but measured
 - "backfill" jobs, prod/analysis success rates from Dashboard, link quality, etc.



		T1_DE_FZK																					
Site Readiness Status:		R	W	R	W	NR	NR	R	R	R	R	W	R	R	R								
Daily Metric:		E	E	O	O	O	O	E	O	E	E	O	O	O	O	E	O	O	O				
Maintenance:		Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up				
Job Robot:		84%	75%	97%	100%	100%	100%	95%	95%	95%	95%	95%	95%	95%	95%	95%	95%	95%	95%				
SAM Availability:		84%	82%	98%	100%	100%	100%	99%	100%	44%	66%	100%	100%	100%	100%	100%	100%	100%	100%				
T1::downlinkT0:		1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1				
T1::downlinksT1s:		1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1				
T1::uplinksT1s:		1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1				
T1::uplinksT2s:		1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1				
		15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	01	02	03	04	05	06
		Jun															Jul						




So...

- **Definition: a site is “doing ok” if (and only if?):**
 1. It has provided (usable) resources that match those pledged & requested;
 2. The services are running smoothly, pass the tests and meet reliability and availability targets;
 3. “WLCG operations” metrics on handling scheduled and unscheduled service interruptions and degradations are met (more automation / reporting required here);
 4. It is meeting or exceeding metrics for “functional blocks” (this latter will require some work: is it “fair” to expect experiments to define these? Must be as simple as “autopilot settings”)
- **If it is as simple as this, all we have to do is:**



- Make it so...

The Dashboard Again...



WLCG Tier1 (Performance) Metrics

#	Metric
1	Site is providing (usable) resources that match those pledged & requested;
2	The services are running smoothly, pass the tests and meet reliability and availability targets;
3	"WLCG operations" metrics on handling scheduled and unscheduled service interruptions and degradations are met;
4	Site is meeting or exceeding metrics for "functional blocks".