



Proposal on advance notice for Long Scheduled Downtimes

v2

Ops Coordination meeting, Jan 26th 2017

Introduction

- Experiments would like to avoid long downtimes without sufficient advanced notification
 - A few bad experiences in the past
 - ATLAS came up with a proposal in the Ops Coord meeting in November
- Ops Coordination collected feedback from experiments as summarised in the backup slides
- A first comprehensive proposal was presented at the Dec GDB
- Feedback from sites has been incorporated in an updated version presented here

Availability and Reliability Reports

- Official WLCG Document

The WLCG [Memorandum of Understanding](#) defines availability targets for CERN (Annex 3.1), Tier-1s (Annex 3.2) and Tier-2s (Annex 3.3) as follows:

$$\text{availability} = \text{time_running} / \text{scheduled_up_time}$$

The site availability metrics are calculated by the [Site Availability Monitor](#), which runs a range of different tests at regular intervals throughout the day. A site is considered to be *available* if a defined set of critical tests complete successfully. The daily averages are calculated for each site and [reported monthly](#) to the WLCG management. These metrics distinguish between *availability* and *reliability*, with the following definitions:

$$\text{availability} = \text{time_site_is_available} / \text{total_time}$$

$$\text{reliability} = \text{time_site_is_available} / \{\text{total_time} - \text{time_site_is_scheduled_down}\}$$

The Management Board considers that reporting both of these measures is necessary, and has selected a terminology that follows the normal meanings of the two terms. This has the unfortunate consequence that the term *availability* used in the MoU refers to the *reliability* metric.

- WLCG availability == SAM Reliability
 - Not affected if the site is in Scheduled Downtime

Downtime definition in WLCG today

- Scheduled Downtime
 - Downtime of any duration in time which is declared in GOCDB/OIM at least 24 hours in advance
- Unscheduled Downtime
 - All the rest
- GOCDB algorithm:
 - All downtimes declared less than 24h in advance will be automatically classified as UNSCHEDULED
 - All other downtimes will be classified as SCHEDULED
- OIM
 - No documentation found!

Dec GDB WLCG Ops Proposal (I)

	Scheduled Downtimes		
	D \leq 5 days	5 days < D \leq 1 month	1 month < D
WLCG OPS	Declaration in GOCDB/OIM at least 1 day in advance	Declaration in GOCDB/OIM at least 1 month in advance	Declaration \geq 1 month in advance & Data migrated

Existing policy

New policy

D = Downtime length

WLCG Ops Proposal (II)

- Availability and Reliability calculations of sites that do not respect the proposed policy will be affected
- If declarations in GOCDDB/OIM are not done within the proposed timeline, downtimes will be considered as unscheduled
- *Exactly as it happens today when a site declares a downtime e.g. 23 hours in advance.*

Feedback from sites & considerations

- What if a long downtime needs to be adjusted at the last moment?
 - E.g. due to late delivery, illness of an expert, ...
 - We should allow affected downtimes to be extended or rescheduled (e.g. a week later) without penalty
- Migrating all data out of a site may take many weeks
 - To plan the downtime, the site would have to “guestimate” and work with the affected experiment(s) to try and make it happen
 - What if the time was underestimated?
 - We need to give sites workable recipes
- Sites usually do their best to serve their experiments
 - We should not over-engineer these procedures

WLCG Ops Proposal v2

	Scheduled Downtimes		
	D ≤ 5 days	5 days < D ≤ 1 month	1 month < D
WLCG OPS	Declaration in GOCDDB/OIM at least 1 day in advance	Declaration in GOCDDB/OIM at least 1 month in advance	Declaration ≥ 1 month in advance & Engage with affected experiments to allow (sufficient) data to be migrated
		Allow the downtime to be adjusted as needed: extended or postponed by a few days or weeks	

Existing policy

New policy

New in v2

D = Downtime length

Next Steps

- Incorporate further feedback, if any
- Understand differences with OIM and EGI policies
 - Impact, if any, on GOCDB/OIM
 - Scheduled/Unscheduled are currently calculated automatically
 - Impact on SAM
 - Implementation of the new policy
- Present the proposal in the MB

Backup slides

Experiments Input

ATLAS and ALICE

- Scheduled “Long” downtimes:
 - “Long” means 5 working days or more.
 - Downtimes should be agreed with the experiments to allow them to take the appropriate actions.
 - Proposal: 1-month advance notice for long downtimes.
 - The new policy should be respected by both Tier1s and Tier2s.
- Downtimes shorter than 5 days are not a problem and the policy should stay as it is, namely 1 day of advance notice.

CMS

- For downtimes lasting:
 - **Less than one week:**
 - The site should just make sure CMS Operations team is aware of it
 - **One week or more:**
 - At least three working days notice
 - **Exceeding a month:**
 - Provide enough time to migrate all data (which could take some time)
 - The site should do this in coordination with CMS

LHCb

- OK to keep 24h warning for downtimes less than 5 working days
- At least 5 days warning for downtimes of 5 working days to one month
- No specific recommendation for downtimes over one month
- OK for counting as unscheduled any downtimes not respecting the advanced warning
- No two T1 sites should be down for the same VO (at least, for LHCb) at the same time – A site should not announce a downtime if another site has already scheduled one
 - *Note: it has been pointed out that such a policy cannot be enforced, because sites are independent entities; but we can try to approximate it*

Summary of Experiments Input

	Scheduled Downtimes		
	D ≤ 5 days	5 days < D ≤ 1 month	1 month < D
ALICE	Notification 1 day in advance	Notification 1 month in advance	
ATLAS		Notification 1 month in advance	
CMS		Notification 3 days in advance	Data migrated
LHCb		Notification 5 days in advance	Notification 1 month in advance, especially T1s

Existing policy

New policy

D = Downtime length