



Proposal on advanced notice for Long Scheduled Downtimes

GDB 14th December 2016
WLCG Operations

Introduction

- Experiments would like to avoid long shutdowns without sufficient advanced notification
 - A few bad experiences in the past
 - ATLAS came up with a proposal in WLCG Ops Coord meeting in November
 - <https://indico.cern.ch/event/540423/>
- After ATLAS initiative, WLCG Operations has been collecting feedback from experiments as summarised in the backup slides
- This presentation summarises the status of the proposal giving also background information

Availability and Reliability Reports

- Official WLCG Document

- <http://cern.ch/go/P6pL>

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
 - For this reason, WLCG Operations would like to review the definition of Scheduled vs Unscheduled downtime

Downtime definition in WLCG:

Today

- Scheduled Downtime
 - Downtime of any duration in time which is declared in GOCDB/OIM at least 24hours in advance
- Unscheduled Downtime
 - All the rest
 - GOCDB:
https://wiki.egi.eu/wiki/GOCDB/Input_System_User_Documentation#Scheduled_or_unscheduled.3F
 - 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!

WLCG Ops Proposal (I)

	Scheduled Downtimes		
	D \leq 5 days	5 days < D \leq 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 & 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.*

Next Steps

- After collecting a first round of feedback from experiments, a first proposal from WLCG was presented in the Dec Ops Coord meeting
- The proposal is now open for comments until 5th January 2016
- Final proposal will be presented in January's Ops Coord meeting on 12th January
- We will then understand clashes with OIM and EGI policies
 - Also understand impact on GOCDB/OIM
 - Scheduled/Unscheduled are currently calculated automatically
- Proposal to be approved in the MB early next year

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 let them taking the appropriate actions.
 - Proposal: 1 month in advance notice for long downtimes.
 - The new policy should be respected by both Tier1s and Tier2s.
- Downtimes shorter than 1-5 days are not a problem and the policy should stay as it is, namely 1 day in 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

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