

# Table of Contents

|  |          |
|--|----------|
| <b>Best practices for scheduled downtimes.....</b> | <b>1</b> |
| Tier-1 downtimes.....                              | 1        |
| Links to Tier-1 downtimes.....                     | 1        |
| Advance notifications of downtimes.....            | 1        |

# Best practices for scheduled downtimes

## Tier-1 downtimes

Experiments may experience problems when two or more of their Tier-1 sites are inaccessible at the same time. Therefore Tier-1 sites should do their best to avoid scheduling a downtime classified as "outage" in a time slot overlapping with an "outage" downtime already declared by another Tier-1 site supporting the same VO(s). The following procedure is recommended:

1. A Tier-1 should check the downtimes calendar (see below) to see if another Tier-1 already has an "outage" downtime in the desired time slot.
2. If there is a conflict, the best would be to pick another time slot.
3. In case stronger constraints do not permit another time slot, the Tier-1 should point out the existence of the conflict to the SCOD mailing list and at the next WLCG operations call, to discuss it with the representatives of the experiments involved and possibly the other Tier-1, to see if a less disruptive scenario can be arranged instead.

As an additional precaution, the SCOD will check the downtimes calendar for Tier-1 "outage" downtime conflicts at least once during his/her shift, for the current and the following two weeks; in case a conflict is found, the SCOD will follow up with the parties involved.

## Links to Tier-1 downtimes

|                       |                       |                      |                      |
|-----------------------|-----------------------|----------------------|----------------------|
| <a href="#">ALICE</a> | <a href="#">ATLAS</a> | <a href="#">CMS</a>  | <a href="#">LHCB</a> |
|                       | <a href="#">BNL</a>   | <a href="#">FNAL</a> |                      |

## Advance notifications of downtimes

The experiments have expressed how much in advance they would like to be informed of a scheduled downtime, depending on its duration. In general *the earlier, the better*. The following table summarizes what each of the experiments would appreciate as advance notifications (N) for downtimes (DT) of various lengths:

| Experiment | DT up to 5 days | DT up to 1 month | DT longer than 1 month               |
|------------|-----------------|------------------|--------------------------------------|
| ALICE      | N = 1 day       | N = 1 month      | N = 1 month                          |
| ATLAS      | N = 1 day       | N = 1 month      | N = 1 month                          |
| CMS        | N = 1 day       | N = 3 days       | Allow enough time for data migration |
| LHCb       | N = 1 day       | N = 5 days       | N = 1 month, esp. for T1 sites       |

Note that CMS would like to have sufficient time to migrate data away from a site that will have a downtime longer than 1 month, to be discussed with CMS operations.

---

This topic: LCG > BestPracticesForSchedDT

Topic revision: r2 - 2017-12-20 - MaartenLitmaath



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback