



WLCG – Worldwide LHC Computing Grid

Site Suspension: *The theory*

WLCG Workshop, Prague
March 22nd, 2009

Nick Thackray
Grid Operations





What are we trying to achieve?



- Why should we want to suspend sites?
 - For central grid operations:
 - Increase overall reliability of the production infrastructure
 - Reduce the operations support load
 - Isolating security problems
 - For the experiments:
 - Reduce the job failure rate
 - Other?



What the rulebook says



- Actually, not a lot...
 - The only place I know of is EGEE Grid Operations Manual (<https://edms.cern.ch/document/840932>)
- So what does it say?
 - When a site is unresponsive to a ticket
 - Except when that site is doing useful work for a VO (according to the VO) and is not a problem for other VOs
 - When a site is in never ending down time (>1 month)
 - When a site poses a security threat
- ! Nothing about site or service performance



Reviewing the rule book



- Should central operations continue to suspend entire sites for the reasons mentioned?
 - ? When a site is unresponsive to a ticket (*Except when ...*)
 - **NB:** *all services for all VOs are removed from BDII*
 - ✓ When a site is in never ending down time (>1 month)
 - Suggestion: **Yes** - downtime is still not generally handled well by tools, monitoring, etc.
 - Should we reduce this time? (3 weeks? 2 weeks?)
 - ✓ When a site poses a security threat
 - **Yes:** this is not a suggestion
- Should we use more restrictive criteria for the first point, basing it on some performance measure?
 - IMHO this is an experiment decision for the experiments to implement (as I hope Graeme will confirm)



And for the experiments



- Experiments already effectively suspend sites and individual services
 - They can suspend for different roles – grid operations can't
 - They can suspend independently of one another – grid operations can't
- What would really encourage sites to fix issues would be if the black-lists were published
 - This would have to be automated
 - Bonus: It would also catch when sites *are* fixed but the VO has forgotten about them 😊



Review of reasons for suspension



- For central grid operations:
 - Increase overall reliability of the grid
 - This should come anyway through the site suspension by VOs - and that is more targeted
 - Reduce the operations support load
 - The grid operations model is changing anyway
→ regionalization
 - Isolating security problems
 - Already covered
 - For the experiments:
 - Reduce the job failure rate
 - Other?
- Over to Graeme...