Grid Oversight Activity: Status, Issues, Solutions

Version 0.3
Date: 17-12-2012
Author: COD team

Table of Contents

Introduction

Now that there is still a year and a half to go towards the end of the project, we felt that it was good to look back and examine the status of grid oversight and to see how effective our follow-up activities have been so far. In this document we present the result of this investigation, the conclusions we have drawn from them and our plans for the future.

Status overview

Availability and Reliability

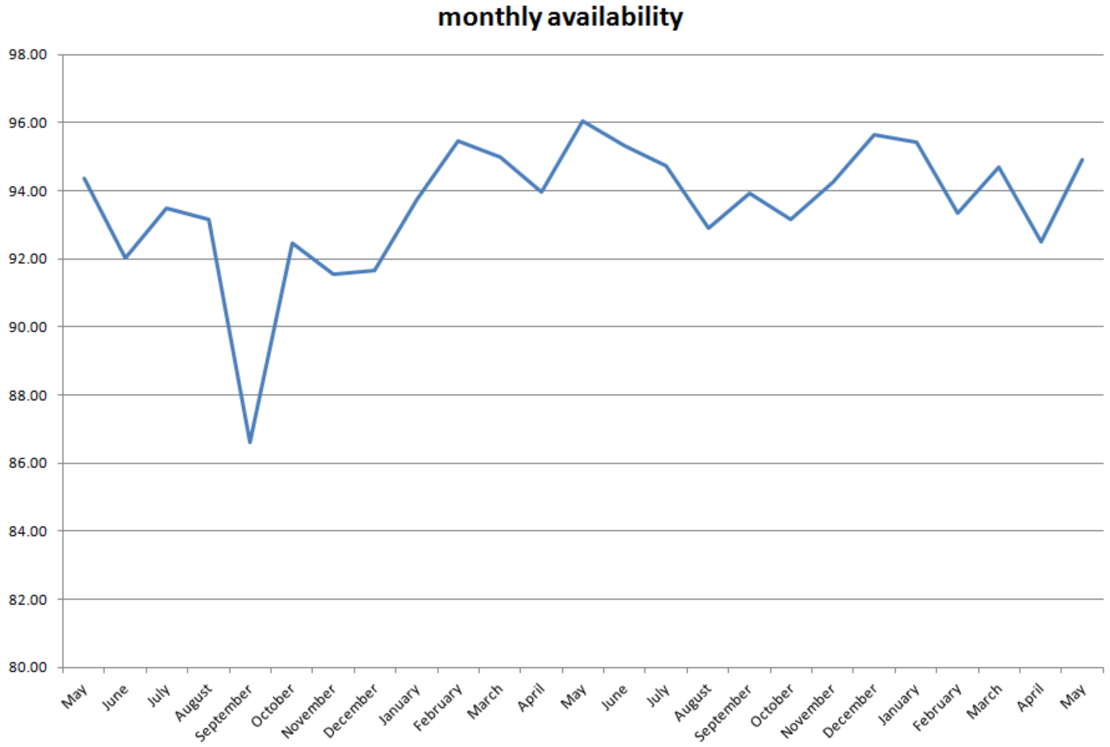

Figure 1: Average site availability

Since the beginning of the project we have been involved with the follow up on availability and reliability. Every month, GGUS tickets were issued to NGIs for every site that has had an availability that was below the 70%/75%. For sites that were below this target for three consecutive months were liable for suspension. Figure 1 shows the availability from the start of the project until May this year. It seems to be oscillating around the 94% and does not increase or decrease. The same holds for the number of tickets generated each month which is usually somewhere between 30 and 40.

The monthly generation of tickets does not seem to be an effective means to increase the average availability and reliability of sites.

Since October 2012, a nagios probe that shows an alarm when a site is below the 70% availability over the last 30 days has replaced the monthly generation of tickets. It remains to be seen how effective this will be.

ROD Performance Index

Figure 2 shows the ROD Performance Index since January 2011. A decaying trend is clearly visible which seems to stop a few months ago. Since January 2012 there has been a follow-up by means of submitting GGUS tickets to any ROD team that has had an index of higher than 10 over a given month.

Interestingly enough, the decaying trend did not seem to be influenced by this kind of follow-up. It started before the start of the follow-up and persisted until July 2012. Analysis of the data showed that this decaying trend can be largely attributed to the so-called new NGIs, meaning not former EGEE ROCs.

Looking at the responses that we have had on our tickets, a high RPI seems not to be caused by persistent issues but mainly by transient ones. Examples of these are, people that forgot that they were on a ROD shift or people that took a day off and forgot to hand-over their shift to someone else and so on. It seems that, generally, these issues were caused by incidents rather than by structural problems. Figure 3 shows the RPI of a typical ROD team. This seems to corroborate this claim.

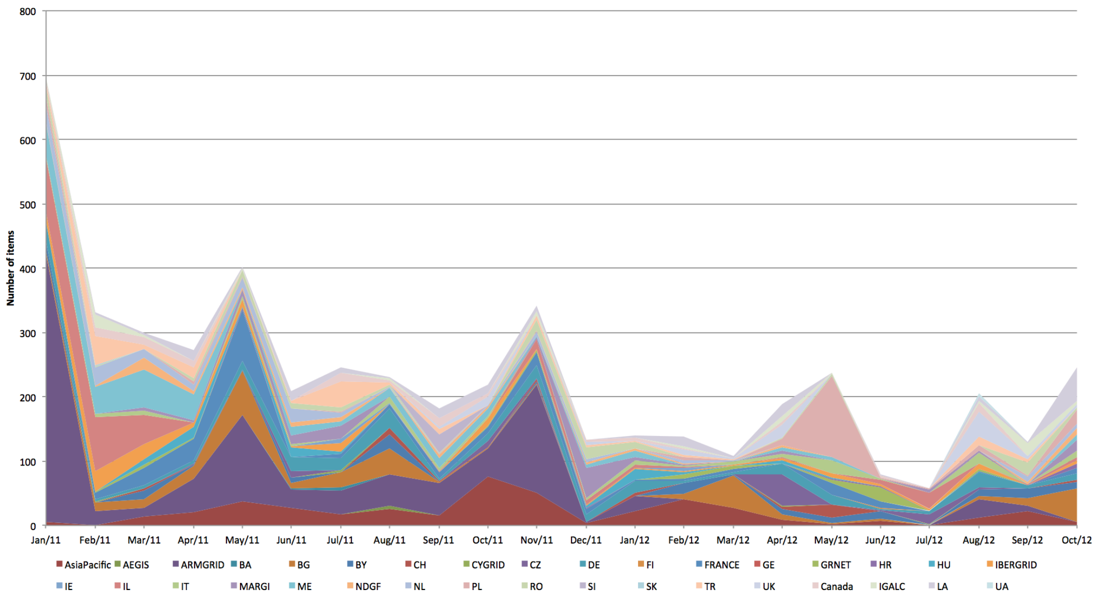

Figure 2: ROD Performance Index

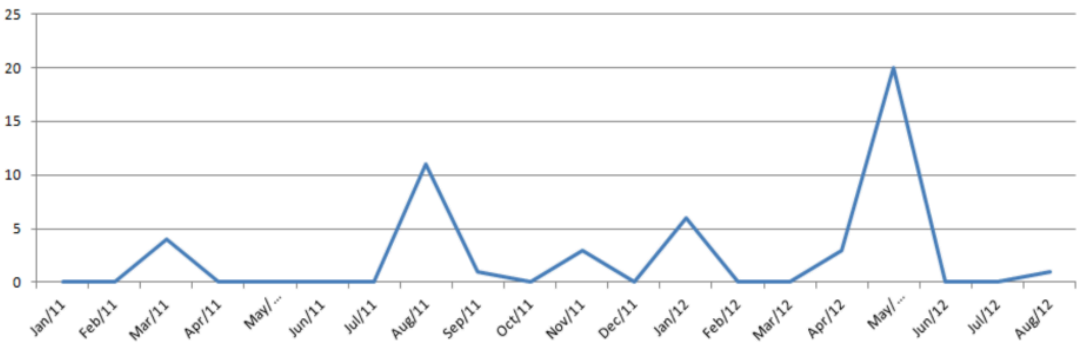

Figure 3: ROD Performance Index of a typical ROD

Toplevel BDII

Since January 2012 COD has been involved with the follow-up on toplevel BDII availability and reliability. This was done by submitting GGUS tickets to NGIs that have a toplevel BDII A/R below the 99%. Figure 4 show the number of NGIs having toplevel BDIIs with an A/R that is below target. It is clear that the follow-up had had a beneficial effect. In the third month of the follow-up, a link was put in the GGUS ticket to documentation on how to setup a reliable BDII. This may have had a positive effect as well.

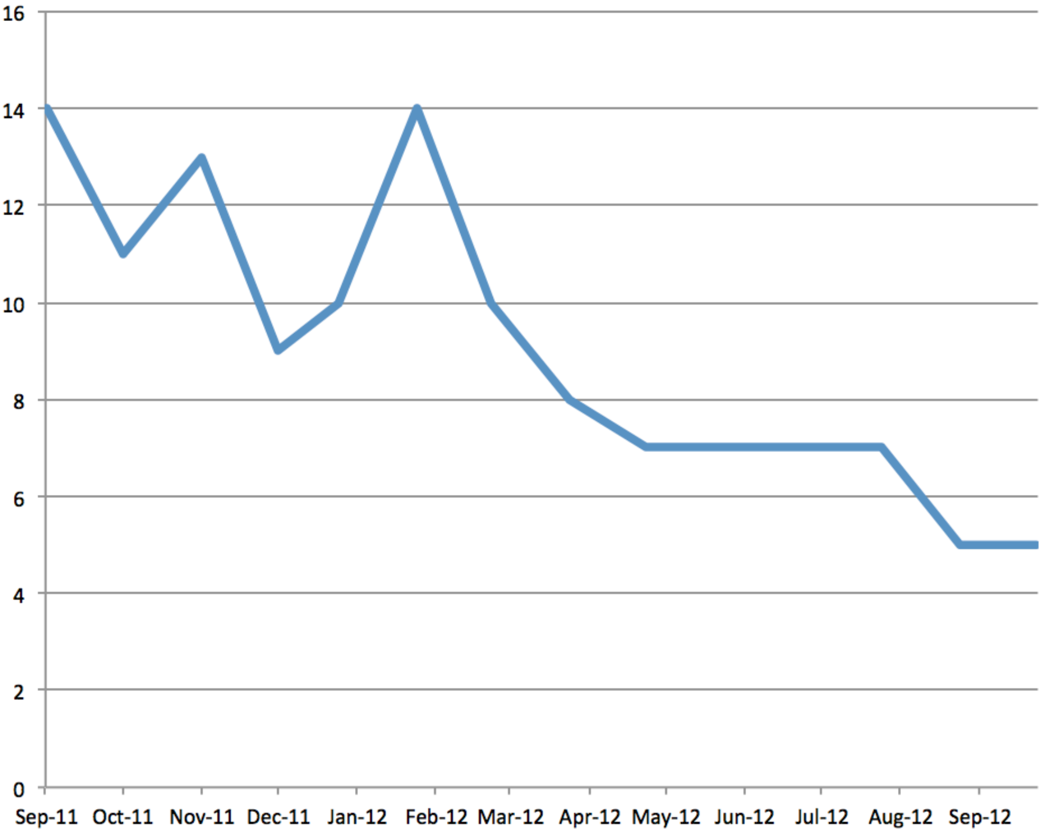

Figure 4: The number of NGIs having toplevel BDIIs below target

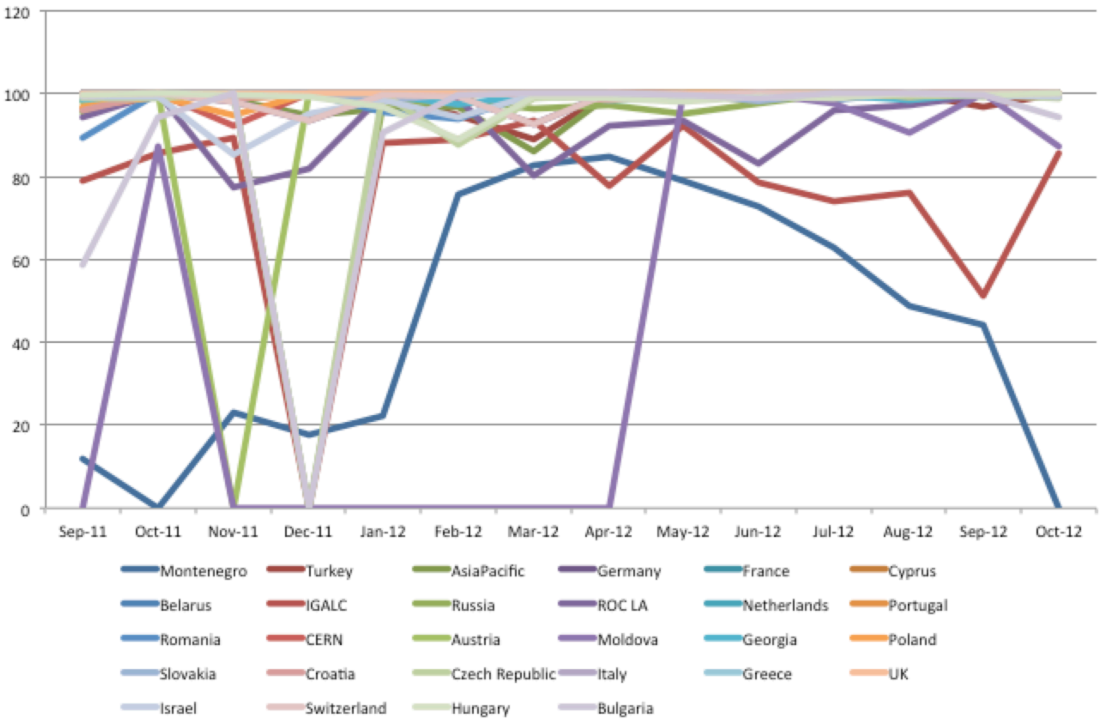

Figure 5: Toplevel BDII availability

Figure 5 shows the availability of the toplevel BDIIs of various NGIs. This figure shows that at the time of writing there are three NGIs that have problems with their toplevel BDII. One of them is in the process of decommissioning, one is using the catchall toplevel BDII now and the third one has just recently been certified and needs to be closely monitored to see how this develops. The other NGIs seem to be doing fine the last six months and if one of them is below target, it is only just below target and it is only incidental.

Unknown

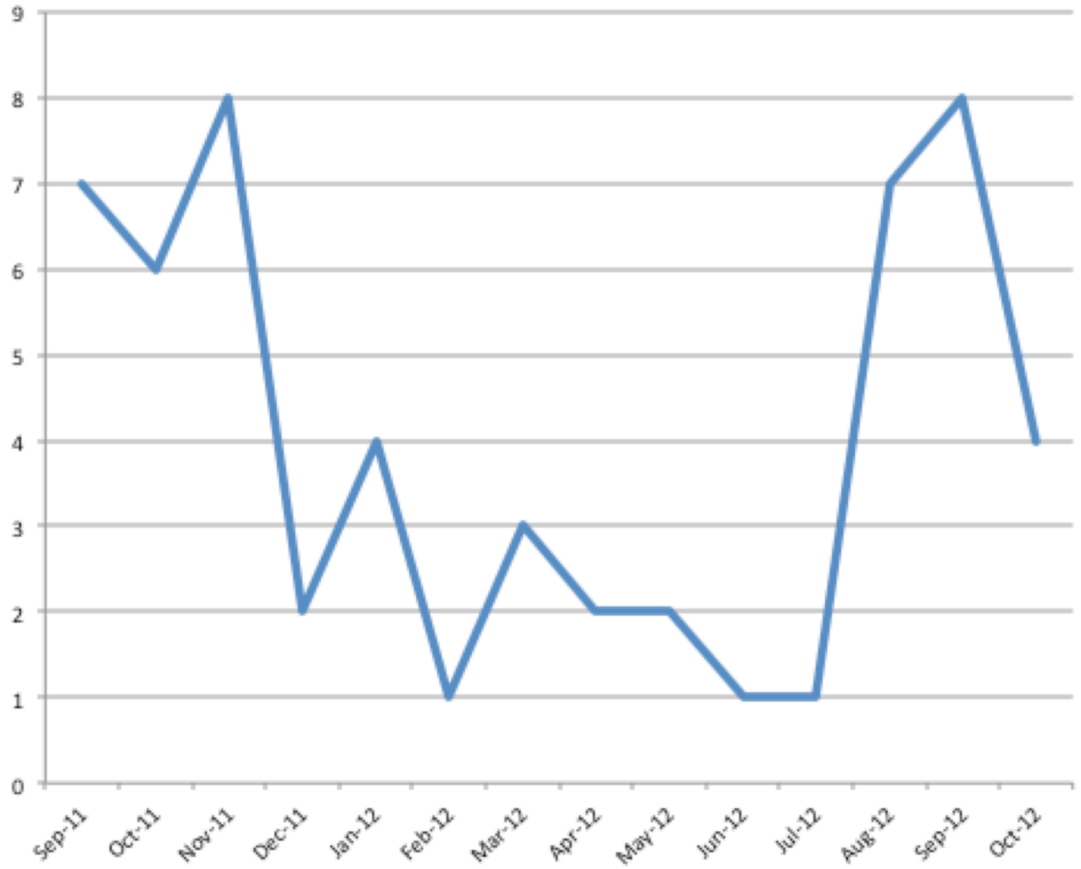

Figure 6: Number of NGIs that have an unknown percentage of higher than 10% for all their sites

The follow-up of the unknown issue has also started in January 2012. This follow-up was done by submitting GGUS tickets to NGIs having sites with an unknown percentage higher than 10. In the GGUS tickets NGIs were informed of this fact and asked to look into the matter.

Figure 6 shows the number of NGIs that have an unknown percentage about 10 for all their sites. Figure 7 shows the average unknown percentage of various NGIs. It seems that the follow-up by means of GGUS tickets did not have any beneficial effect at all. But it seems that there are about three major contributors to the unknown percentage over the past year or so.

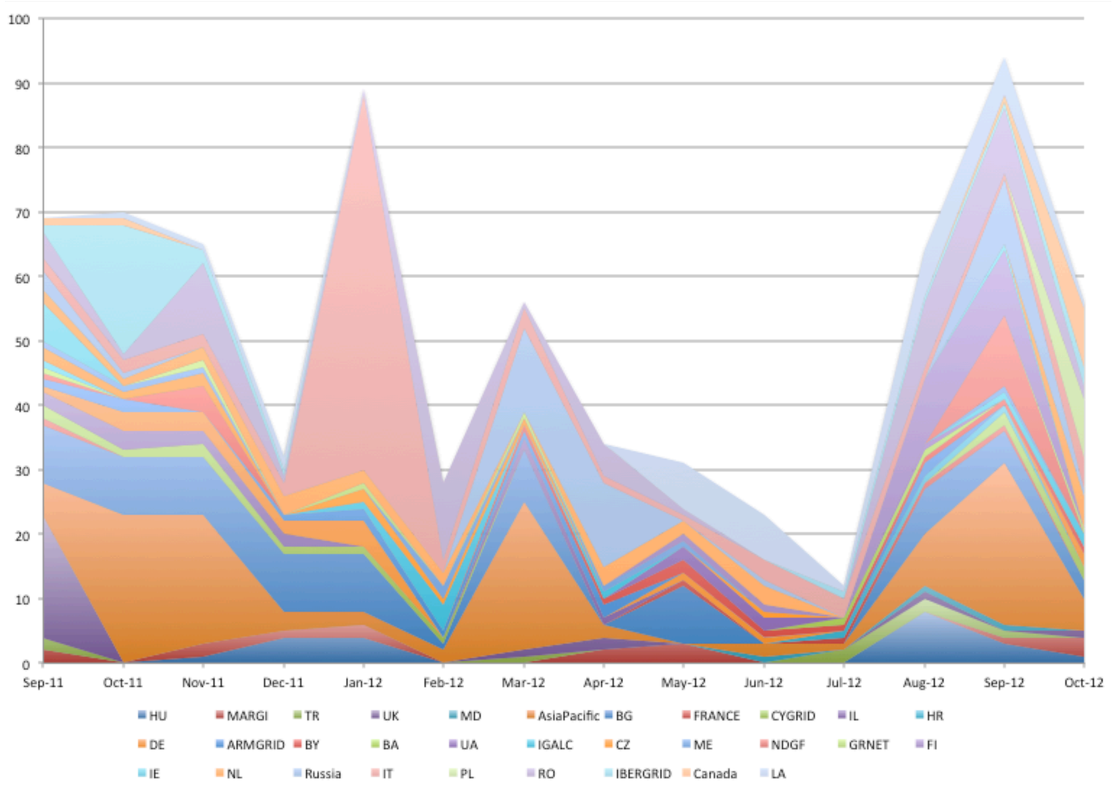

Figure 7: Unknown percentage for the various NGIs

Observations and Conclusions

First we would like to note is that getting the data into a format so that it is easy to work with and extract information from is a painful process. Sometimes there are excel sheets available but, for example, when you want to investigate the evolution in time for the "unknown" percentage you can only get the data from a pdf file. It would be good if acquiring and processing of data of this kind would become easier. Another thing is that in the excel files naming is not consistent. For example, in one excel file NGIs are denoted by their country names, like Germany, while in another file you have NGI_DE. It would be good to have consistent naming. In the excel file for the toplevel BDII availability and reliability there are even NGIs that have not undergone the certification process. This is also not correct we believe.

The results in the previous section show that follow-ups by means GGUS ticket have stopped to be effective. For the follow-up of availability and reliability, ROD performance index and the unknown issue there was no beneficial effect visible. The follow-up of the toplevel BDII availability and reliability has had a positive effect in the beginning but the number of NGIs having toplevel BDIIs that are below the targets seems to level off now. So, continuing the follow-ups by means of tickets is likely not going to result in any improvement. The reason behind the follow-up with tickets is that it is annoying for sites and NGIs to get tickets that would hopefully then result in improvement. This effect, if it was there at all, has worn off. It seems that answering and closing tickets become a day-to-day routine and automated tickets result in automated answers that are impossible to verify.

Solutions

Since we doubt that the continuation of follow-up by means of automated generation of GGUS tickets will be effective, we have the following proposal.

Apart from the role of guardians of OLAs, we believe that COD should put more emphasis on providing or arrange for support. Moreover, it is our experience that when you approach people directly, people are more inclined to do what you ask from them than when the get automatically generated tickets. If people know that someone cares about the work that they do and how they do it, they will be more inclined to put effort into it.

With this in mind we propose to stop with the monthly follow-up of unknown percentage, ROD performance index and toplevel BDII availability and reliability in a different way. We would propose to notify NGIs about OLA violations by means of emailing pdf's to the noc-managers list. We believe that it is good to address OLA violations at management level since this is where such a thing belongs.

On top of this we propose to collect the data and investigate it and try to identify persistent issues with NGIs and address that. It is not very sensible to ask an NGI for an improvement plan when their toplevel BDII is at 98.9% A/R when it has always been at 100%. Sometime sites are in a rough spot and in a big infrastructure, bad things do happen. Rather than reacting on incidents, we believe it makes more sense to focus on structural problems and approach those NGIs and give them support if possible or organise support for them. The idea is to notify them about OLA violations and at the same time put more emphasis on helping them to meet OLA requirements.

We feel that our efforts can be more effectively used if we focus ourselves on structural issues rather than incidents. We believe that this could strengthen our bonds with our RODs. Of course, this strategy will only work if people are willing. If they are not willing, nothing will help whatever we do.