



Answers to ECC recommendations (May 2013)

The Virgo collaboration and the EGO IT Department

VIR-0196A-14
Date : April 13, 2014

Abstract: We answer to the ECC Recommendations from the May 2013 report

VIRGO Collaboration
EGO - Via E. Amaldi - I-56021 S. Stefano a Macerata, Cascina (Pisa)
Secretariat: Telephone (39) 050 752 511 - FAX (39) 050 752 550 - e-mail perus@ego-gw.it

Contents

1 Recommendations and answers	2
1.1 Recommendation 7	2
1.1.1 Recommendation: The schedule of the work needs to be accelerated.	2
1.1.2 Answer	2
1.2 Recommendation 8	2
1.2.1 Recommendation: Importance of agile data access...	2
1.2.2 Answer	3
1.3 Recommendation 9	3
1.3.1 Recommendation: Integration of computing farms with DAQ in Cascina	3
1.3.2 Answer	3
1.4 Recommendation 10	4
1.4.1 Recommendation: Unifying databases	4
1.4.2 Answer	4
1.5 Recommendation 11	4
1.5.1 Recommendation: Multi- or multiplebi-lateral agreements with CCs	4
1.5.2 Answer	4
1.6 Recommendation 12	4
1.6.1 Recommendation: Communication and cooperation between EGO, AdV, CNAF, CC-IN2P3	4
1.6.2 Answer to Recommendation 12	5
1.7 Recommendation 13	5
1.7.1 Recommendation: Need for Remote Data Recording requirements	5
1.7.2 Answer	5
1.8 Recommendation 14	5
1.8.1 Recommendation: Evaluation of Remote Data Recording	5
1.8.2 Answer	5
1.9 Recommendation 15	5
1.9.1 Recommendation: Importance of good network connectivity	5
1.9.2 Answer	6

1 Recommendations and answers

These are answers for all what concerns data analysis and computing of the scientific and detector characterization workflows.

1.1 Recommendation 7

1.1.1 Recommendation: The schedule of the work needs to be accelerated.

The schedule of work on the AdV Computing Model needs to be accelerated. In particular, the following should be completed by the fall 2013 committee meetings:

- 1. The initial version of the complete AdV Computing Model , with all the components listed in section 4 of the ECC 2012 review.*
- 2. A short document detailing the management process for the AdV Computing Model, including the persons or bodies responsible for driving each part of the process, the chains of approval for considering each step completed, and the timing of each part of the process.*
- 3. A draft version of an implementation plan up to first AdV science operation , focusing on the methodological and technological choices and outlining the software development (both services and algorithms) which will be necessary.*
- 4. An early draft of an estimate of the budgets necessary to execute the implementation plan, including human resources and quantifying monetary as well as in-kind contributions.*

1.1.2 Answer

We have presented the version 1.0 of the Computing Model (CM) to ECC (April 2013) and STAC (October 2013), after having also received the approval by the VSC. The Management Plan (VIR-....) is ready and presented together with this document. The AdV Implementation Plan (IP) (VIR-....) has been recently sent to the ECC. After a first round of comments by the ECC and STAC (May 2014) we will submit it to the VSC for official approval. The final section of the IP contains manpower tables, with names and FTEs, for all the projects in the CM. We haven't yet discussed the impact of requests in terms of monetary needs, as this process require to involve the management of the collaboration, of EGO and, in some cases, tender notices handled by EGO IT Department.

1.2 Recommendation 8

1.2.1 Recommendation: Importance of agile data access...

The importance of agile access to a variety of datasets by the experts on commissioning and characterization of the interferometer is being clarified, and good progress in documenting the details of these datasets is being made. The method proposed for granting this access is to provide a disk buffer and associated CPU farm at Cascina. The size of the buffer is dominated by the requirement of keeping 0.5 years of Raw Data on site, with the corresponding impact on CPU. These data, however, will also be archived and accessible for analysis at CNAF and CC-IN2P3. The ECC requests a thorough evaluation of an alternative solution where an agile mechanism for access and analysis of Raw Data at CNAF and CC-IN2P3 is made available in a transparent manner to the experts working on commissioning at the interferometer site. This would reduce very significantly the size of the disk buffer space and accompanying CPU in Cascina without necessarily increasing the resources needed at CNAF and CC-IN2P3. The ECC does acknowledge that a modest amount of disk space and CPU is needed in Cascina, beyond what is strictly needed for the AdV Data Acquisition system, so that small datasets and, particularly, databases, can be quickly accessed for commissioning and characterization.

1.2.2 Answer

We agree that we will need to organize a reliable and fast access to the data (including raw and RDS) stored in the CCs, even when these are stored on tapes. This work is part of the re-organization process on Data Access detailed in the IP (see Sect. xx) In any case we need to maintain our request of 6 months of data at the EGO site, on a circular buffer, given that:

- the actual organization of having six months of data stored on site for commissioning activities has shown to work and our colleagues working on the apparatus have developed many tools to access and analyze the data for their specific needs and to display on site the results, that it would now require a huge work to organize the same in the CCs.
- Raw data are foreseen to be copied to CCs only during Science runs, as done for Virgo, thus it is needed to have them in the circular buffer onsite for commissioners and experimentalist people. The CCs architecture is mostly “job-oriented” than “interactive-displays”, which is what is needed for commissioning.
- Let’s note that, as clarified in the computing resource requests section of the CM, the storage request does not have any impact in the computing power needed for the Cascina farm, mainly due to the DAQ and on-line, in-time analysis needs.

1.3 Recommendation 9

1.3.1 Recommendation: Integration of computing farms with DAQ in Cascina

The ECC re-emphasizes the importance of its recommendations 1 and 2 and encourages the Disk and CPU farms in Cascina to be seamlessly integrated with the AdV Data Acquisition system, operationally as well as technically. Responsibilities for funding, deployment, operation and technical support of the Cascina Disk and CPU farm should be clearly and urgently defined between the EGO Laboratory and the AdV Collaboration. Note: Recommendation 1 and 2 are reported below.

1.3.2 Answer

The Virgo experiment has since the beginning split the data workflow from the collection at the instrument up to the access by data analysis applications at the computing centers in two subsystem: 1- Data Acquisition (DAQ): prompt collection, assembly and on-line access of the data frames from the interferometer real-time servers up to the first writing of files on disk, performed on special circular buffers with limited access from users.

2-Data Distribution or Data Management: local handling, migration to mass storage, tape archiving, transfer to the computing centers of the file produced in order to satisfy the needs from various subsystems either on-site (in-time analysis, commissioning, characterization, simulation) and off-site (off-line analysis). This is the scheme we are proposing for AdV.

We understand that the IT architecture, in particular for computing and storage, should integrate on one side with DAQ, the other ITF subsystems and the on-line data analysis pipelines for the on-line needs and on the other side with the Data Distribution framework to interface with external resources.

Indeed, this process on the ITF side has been always effective in meeting the Virgo needs. We have now given in the CM definitions and requirements of the workflows and thus the detailing of the computing resources at Cascina is progressing. Nonetheless, the technical and operational management of the storage and computing systems and their harmonization is placed on-site in the EGO IT.

The definition of responsibilities for fundings, deployment, operation and technical support is a task for the Virgo and EGO management, so the answer can’t be given in this document.

1.4 Recommendation 10

1.4.1 Recommendation: Unifying databases

Various workflows involve databases at the interferometer site, and in many cases there are previous Virgo implementations based on specific database packages. In some cases, these databases need to be exported or replicated to other sites. One key database contains data from both Virgo and LIGO. A process should be started for the evaluation of choices, risk assessment, testing and prototyping, in order to reach an agreement on a single underlying database platform which serves all workflows and fulfills the offsite export/replication requirements, as well as those of the Virgo/LIGO database. This database platform should be implemented and maintained as a logical component of the Cascina Disk and CPU farm.

1.4.2 Answer

We do agree that this is a very important question. To answer, we have put section in the IP and there are open Redmine tasks associated to the relevant persons, however there are multiple issues here which makes the problem not so trivial to solve, including:

- For lightweight analysis workflows even is using some (reduced) dataset, it is not really necessary to use anything more serious than simple text files with list of information. Any use of a more serious database client would just complicate the software.
- The development of the analysis workflows which are heavily using various datasets are dominated by our LIGO colleagues. In many cases the solution choosed are not compatible with the standard possibilities available in Virgo, as such would require extra (and continous) bi-directional translation, migration and synchronisation of the databases which is errorprone and requires to much manpower. Having this said we are not completely sure, that the database and lookup solutions applied in LIGO pipelines should really be transformed into a common database solution. (See for example ligo_data_find in ahope vs LFC or Dirac file catalgs..)
- Even in the case where it would be technically easy there is no well defined decision mechanism which would ensure that the implementation of the new common solution would happen in all the relevant places.

1.5 Recommendation 11

1.5.1 Recommendation: Multi- or multiplebi-lateral agreements with CCs

The management of the EGO Laboratory and of the AdV Collaboration should discuss and agree with the management of CNAF and CC-IN2P3 a framework agreement, either multi-lateral or multiple bi-lateral, to regulate the use of these computing centers for the data processing of AdV. The Computing Model management plan should take into account these agreements.

1.5.2 Answer

Yes, we have put a section in the Computing Management Model document

1.6 Recommendation 12

1.6.1 Recommendation: Communication and cooperation between EGO, AdV, CNAF, CC-IN2P3

Organizational and technical communication and cooperation between EGO, AdV, CNAF and CC-IN2P3 needs to be urgently streamlined. The ECC recommends the immediate setup of an AdV Computing Coordination Committee (AdV-CCC) composed of one contact each from EGO, CNAF and CC-IN2P3 and chaired by the AdV Data Analysis Coordinator. The existence of this coordination committee should be incorporated into the

framework agreements of recommendation 11, but given the urgency it should be setup immediately even if the framework agreements are not finalized.

1.6.2 Answer to Recommendation 12

The Committee (CTCC) has been created. It has a membership described in the recommendation in addition the chair of the VDASC group is also a member. A regular communication channel has been established in a form of a restricted membership email-list: ctcc@ego-gw.it

1.7 Recommendation 13

1.7.1 Recommendation: Need for Remote Data Recording requirements

The needs for Remote Data Recording, at both CNAF and CC-IN2P3, of 0.9 PB per year of AdV science operation have been clearly established. Early budget planning should incorporate the need for deployment of 0.3 PB of tape archive by the start of AdV commissioning and 0.9 PB per year of AdV science run at each center. Efficient use of this tape archive will require a disk buffer, data access methods and access to the CPU farms at the centers. Requirements for these elements remain to be estimated.

1.7.2 Answer

We do agree this is very important. The IP defines the ongoing work to re-organize data access at the external CCs.

1.8 Recommendation 14

1.8.1 Recommendation: Evaluation of Remote Data Recording

A process for evaluation of choices, testing and prototyping, and reaching agreement on a solution for the Remote Data Recording should be started, ideally driven by the AdV-CCC. In this context, the ECC re-emphasizes the importance of its recommendation 4. The search for a solution should give maximum priority to choices which maximize the economies of scale of both CNAF and CC-IN2P3, presumably those which will be adopted for the second run of LHC. In order to maximize reliability, ease operations at the interferometer, and minimize technical manpower, emphasis should be given to finding a single solution with a common implementation for sending AdV data from the EGO site to both CNAF and CC-IN2P3.

1.8.2 Answer

Yes, we agree this is very important and infact we have been doing an extensive investigation about possible solutions. Section 4.3.2 of the IP address the data transfer question. EGO has found that their custom developed transfer framework is the best solution, since it maximises the available support from the CCs and already proved to work reliably.

1.9 Recommendation 15

1.9.1 Recommendation: Importance of good network connectivity

An adequate network connection of the EGO Laboratory to the European Academic Research Networks, presumably through the Italian NREN, should be ensured in order to implement Remote Data Recording and, possibly, Remote Data Access for commissioning and characterization (see Recommendation 8). The EGO Laboratory, in coordination with the AdV Collaboration, should start dialogs with the relevant network providers and funding agencies in order to ensure that the technical planning and corresponding funding for the network is in place by the start of AdV commissioning and maintained throughout the AdV science run.

1.9.2 Answer

The EGO IT Department has conducted a preliminary survey concluding that there are no main technical obstacles in increasing the bandwidth of the Cascina-GARR link using the full 1Gbit/s physical bandwidth or even to create virtual dedicated links to the remote data centers at Bologna and Lyon in order to have guaranteed bandwidth for data transfer or data access (in this case increasing the yearly Internet link cost, with an increase of 130K). Therefore we consider that the main condition to proceed to the upgrade is the advancement of the model for the data access layer of the Advanced Virgo applications, above all the interactive ones.

We report below the two Recommendations (April 2012) 1,2

Recommendation 1 *A complete, unified survey should be made immediately of the IT installations in Cascina, classifying items into sub-systems with a clearly defined purpose. The identified purposes should be prioritized, and any sub-systems whose purpose does not meet a minimum priority should be targeted for decommissioning. The remaining sub-systems should be organized into a unified, coherent, Cascina IT installation management plan specifying capacity planning and renewal strategies for each sub-system and an associated budget planning, approval and funding process.*

Recommendation 2 *The computing and data processing cluster at Cascina should not be for general data analysis use. It should be specifically targeted to the commissioning, calibration, monitoring, safety and operation of the experimental facilities. The renewal of the Cascina cluster should be included in the overall process described in Recommendation 1*