

CNRS

Centre National de la Recherche Scientifique

INFN

Istituto Nazionale di Fisica Nucleare



Virgo Computing Needs for 2006 and Future

VIR-LIS-ROM-7000-130

Issue: Draft

Date: October 23, 2005

VIRGO * A joint CNRS-INFN Project
:Traversa H di Via Macerata - Santo Stefano a Macerata, -56021 Cascina, Italia.
Secretariat: Telephone (+39) 050 752 521 * FAX (+39) 050 752 550 * e-mail: virgo@virgo.infn.it

	Virgo computing: Needs for 2006 and Future	DOC: VIR-LIS-ROM-7000-130 Issue Draft Date: October 23, 2005 Page: 2/7
---	---	---

CHANGE RECORD

<i>Issue/Rev</i>	<i>Date</i>	<i>Section affected</i>	<i>Reason/ remarks</i>
issue 1	23/10/2005		
Issue 2	24/10/2005		

Authors: F. Ricci S. Cortese M.A. Bizouard Approved by:	Date 23/10/2005	Signature
---	---------------------------	------------------

TABLE OF CONTENTS

1	INTRODUCTION	4
2	OVERALL COMPUTING STRATEGY	4
2.1	EGO/CASCINA.....	4
2.2	CNAF/BOLOGNA AND CCIN2P3/LYON	4
2.3	LABORATORIES.....	4
3	DATA PRODUCTION	4
3.1	REAL DATA	4
3.2	SIMULATED DATA AND USER DISK SPACE.	5
4	COMPUTING EXPENSES IN THE PAST	5
4.1	BOLOGNA.....	5
4.2	LYON.....	5
5	COMPUTING NEEDS FOR 2006	6
5.1	DATA ACQUISITION AND IN-TIME PROCESSING (CASCINA)	6
5.2	DATA TRANSFER TO COMPUTING CENTERS	6
5.3	CONTINUOUS SIGNALS	6
5.4	BURST SOURCES	6
5.5	STOCHASTIC BACKGROUND.....	6
5.6	COALESCING BINARIES	6
5.7	DATA STORAGE AND DATA ACCESS IN COMPUTING CENTERS	6
5.8	SUMMARY OF CPU REQUESTS.....	7
5.8.1	<i>Cascina</i>	7
5.8.2	<i>CNAF/Bologna</i>	7
5.8.3	<i>CCIN2P3/Lyon</i>	7
6	TRENDS FOR THE FUTURE.....	7

	Virgo computing: Needs for 2006 and Future	DOC: VIR-LIS-ROM-7000-130 Issue Draft Date: October 23, 2005 Page: 4/7
---	---	---

1 Introduction

This note describes in short the computing resources and the required computing needs for 2006. We state that during 2006 we plan to review this request after six months for tuning better the provisions reported here. At that time on the base of the Virgo commissioning status we will have a more robust estimation of the extension of the longer data taking period foreseen for the second half of 2006.

The main Virgo computing activities covered in this document are:

- Data Simulation
- Noise source characterization of the interferometer
- Search of continuous signals.
- Search of transient signal for short impulsive events ('burst')
- Search of transient signal for longer events like binary coalescing systems.
- Search of stochastic background
- Data analysis of network constituted by the various G. W. detectors in the world

2 Overall computing strategy

We recall here the purpose of each computing site involved in the process of the Virgo data analysis.

2.1 EGO/Cascina

The Cascina site is the data production place. The data acquisition system produces data that are stored and buffered locally. Detector commissioning, operation and on-line/in-time analysis for all transient signals are performed. Data conditionings like calibration, re-sampling, subtraction of instrumental artifacts (the so called h-reconstruction) are also part of the Cascina computing activities. A backup copy on tapes of all data is also performed in Cascina for crash recovery as long as the data are on the Cascina disks.

2.2 CNAF/Bologna and CCIN2P3/Lyon

The two national computing centers are the archive sites for the Virgo data. They are the primary distributors for these data to the laboratories. They provide resources for the extensive ('production') offline analysis as well as for the simulation effort. The reprocessing (in the sense of performing again in-time analysis when some bugs are found or procedures improved) will also take place in the computing centers. The most demanding activities are the search of continuous signals for which the estimated required computing power is in the Teraflops range and the search of transient signals like coalescing binaries of a network of detectors. We have demonstrated that both the typical PC farm implemented in the national computing centers and the Grid environment are suited for these kinds of searches.

In CNAF/Bologna we store the data on large disks to provide fast access for two years of data, while CCIN2P3/Lyon provides the very long-term archive using the HPSS mass storage system.

2.3 Laboratories

The collaboration labs support the software development, analysis prototyping, interactive activities and data visualization. They are the front-end for the off-line analysis. The laboratories can play the Tier 2/ Tier 3 role in the Data Grid architecture. This basic sharing is not expected to change with the introduction of the Grid tools. Great benefit will follow from the integration of national computing centers and labs resources, for off-line activities.

3 Data Production

3.1 Real Data

Following the commissioning plan discussed recently (October 2005) the assembly of the new injection bench should end in December 2005 and the commissioning will restart immediately. The detector will be operated tentatively in its standard configuration, but additional time is needed to tune the multiple controls and adjust the various components in order to reach the design sensitivity. In this phase, the data acquisition will run continuously by the teams of people tuning the machine. Then, once a commissioning milestone is achieved, in a stable configuration we will perform a commissioning run.

These runs produce the data that are archived and require extensive offline analysis, while the data produced daily are used by the commissioning team to learn and improve the machine. Moreover, once the safety conditions for

the interferometer are fulfilled, we plan to leave the machine unattended taking data in Science Mode during the week-ends.

The plan is to have 14 days of commissioning runs in the first half of 2006 plus some week-ends starting from March 2006 (12 days) and 90 day run during the second half + the weekend (12 days).

In addition to the run data, we plan to archive the trend data and 50Hz data. This corresponds to 2 Terabytes every 6 months. Therefore the amount of data to be archived is expected to be:

Period	Integrated run time * [day]	Total data produced [TByte]
First half of 2006	14+12	15
Second half of 2006	90 +10	52
First half of 2007	100	52
Second half of 2007	100	52

Remarks:

- The amount of data is computed assuming a 6 MBytes/s data rate (compressed data).
- The total data produced include the monitoring and network.

3.2 Simulated data and user disk space.

In the case of simulated data, only a small fraction of these channels are needed, especially for off-line analysis. Therefore the flux of simulated data will not exceed 10% of the flux of real data. Moreover specific procedures require storing the filtered data for further analysis. In particular the pulsar search require to construct from the h reconstruction channel the short FFT data base and, starting from it, it creates peak maps which are the starting point for the Hough transform step. This means that users will require disk space to host these processed data: in particular we ask to increase progressively the user disk space up to 5 Terabytes in Bologna.

4 Computing expenses in the past

In all the following, it has been considered that a single CPU of 1GHz of clock corresponds to a power of 1 kSPECINT2000. In the following the CPU request is given as energy. The unity is the kSPECINT2000.day. The conversion factor with CCIN2P3 UI is: 1kSPECTINT200.day = 480 UI.

4.1 Bologna

During 2005 the access to the resources of the computing center was available for 12 users. We opened also two specific account dedicated to the network activity LIGO-Virgo and Virgo- resonant detectors.

However the use of the computing resources in Bologna was low because of the long sequence of problems experienced by the CNAF storage system where the Virgo data are located. At present, all the runs are stored in Bologna, but only the C6 and C7 data are available for the Virgo users.

The main computation load was due to the tests performed by the Periodic Sources group. The main part of this work was carried on using the INFN Grid infrastructure without significant number of jobs sent to the CNAF site. The standard computing activity leads to an energy of 255 kSPECINT2000.day performed in Bologna.

4.2 Lyon

In Lyon all the VIRGO data are archived in the mass storage of the computing center. Until 2004 we stored 17 TBytes of data. During 2005 we archived 10 TBytes of the C6 and C7 runs.

Concerning the computing power, the following table gives the energy used in Lyon for the past four years. Note that for 2005, the request was 4580 kSPECINT2000.day.

Finally, during 2005, Lyon became successfully a member of the GRID Virgo Virtual Organization.

Year	Used Energy (kSPECINT2000.day)	Used Energy (IN2P3 Units)	Number of BQS job submitters**
2002	520	250,000	4
2003	700	338,000	4
2004	5200	2,500,000	4
2005*	3100	1,500,000	7

* Estimation for end of 2005 December

** GRID counts for 1 job submitter

	Virgo computing: Needs for 2006 and Future	DOC: VIR-LIS-ROM-7000-130 Issue Draft Date: October 23, 2005 Page: 6/7
---	---	---

5 Computing needs for 2006

We stress that in 2006 the analysis will deal with real data that have significant physics content. This implies that we expect to have a significant increase in the computing activity.

5.1 Data acquisition and in-time processing (Cascina)

In 2005, 80 TByte were available to buffer the data taken during the commissioning activity. We don't foresee to increase this amount. Concerning the CPUs, the on line farm of 64 Opteron processors will be expanded in order to fulfill the request of the physics group. We refer to the technical document on the 300Gflops farm in preparation for the action details.

5.2 Data transfer to computing centers

In order to avoid saturating the bandwidth between Cascina and GARR-B, data are transferred from Cascina to Bologna and then forwarded from Bologna to Lyon.

As far as the network is concerned, the Cascina site is transferring data to Bologna at maximum speed up to 60 Mbits/s. The performance was degraded by at least one third when transfer from Bologna to Lyon was performed (use of the same server with only one Ethernet interface). In order to maintain the data transfer above the data production rate, we need at minimum to be able to transfer data at 60Mbit/s. That means that in Bologna, a dedicated machine with 2 Ethernet interfaces, or two machines, are requested for 2006. The maintenance of this/these machine by Bologna staff is mandatory to avoid obsolescence. In addition, at the end of 2006 it could be necessary to increase by one third the maximal transfer rate during the long run scheduled at the end of the year.

5.3 Continuous signals

For 2006, the Periodic Sources group plan to perform the analysis of the real data collected during C6, C7 and the following runs. This will implies a significant change in the computing consumption compared to that of the previous years. On the base of our theoretical evaluations which are in agreement with the results obtained during the tests carried on in 2005, for the analysis of C6 And C7 the needs are 8,000 kSPECINT2000.day.

The energy will be used in the first semester of 2006. The group makes use of the GRID framework: the VIRGO Virtual Organization includes both the access to the INFN Grid resources and the Lyon computing resources.

5.4 Burst Sources

The Burst group will continue the simulation work to develop new algorithm, make coincidence and coherent analysis in a network of detectors in the context of the collaborative data analysis effort with LIGO and the resonant detector groups. The data analysis of the real data will increase compared to 2005. That's why we foresee an increase of the CPU consumption: 4000 kSPECTINT2000.day is requested for 2006.

5.5 Stochastic Background

During the first half of 2005, the Stochastic Background group performed data analysis essentially using the facilities available in Cascina. The integration in the computing centers is foreseen in the first semester of 2006. The needed power is still at a negligible level and can be handled with the total CPU request.

5.6 Coalescing Binaries

During the first half of 2006, the CB group will continue exploiting the resources in Cascina to perform the in-time analysis, and will run on the PC farm according to the deployment of the 300 GFlops machine

In 2006 the CB group plans to start performing reprocessing in the computing centers mid-scale Monte Carlo simulations and network analysis. To do that we need to access to farm nodes where MPI is available.

However, seen the fact that for 2006 this activity is mainly concentrated in Cascina, for the off line it seems realistic to request just 1000 kSPECINT2000 .day .

5.7 Data storage and data access in computing centers

The data in Bologna are stored on disk. It is mandatory to provide a reliable access to the Virgo data disks. In Lyon, the data are stored in HPSS. We have noticed at end of 2005, that the access to real data stored in HPSS was limiting the speed of data analysis. We are accessing less than 3% of the data contained in a file, but we copy on scratch disk the full file saturating the HPSS servers. Lyon is proposing to Virgo to use a XrootD server which should avoid such a bottleneck when accessing real data in HPSS. In the meantime in October 2005, Lyon has provided Virgo with an additional HPSS server managing 2TB of cache disk. In total, we have now 2 servers managing a 4TB cache disk.

	Virgo computing: Needs for 2006 and Future	DOC: VIR-LIS-ROM-7000-130 Issue Draft Date: October 23, 2005 Page: 7/7
---	---	---

5.8 Summary of CPU requests

Taking into account all the needs of the physics groups, a total of 13,000 kSPECINT2000.day is requested by Virgo in 2006. This number is dominated by the Periodic Sources request (8,000) while Bursts and Coalescing Binaries groups (300 GFlops machine excluded) are asking roughly 5000 kSPECINT2000.day. In order to insure a good balance between the two centers and saving manpower, it has been decided to share the Periodic Sources activity in the two centers.

5.8.1 Cascina

The requests are the ones given in the In-Time processing section:

- Installation of the second part of the 300 GFlops machine
- Upgrade of the network connection in the second half of the year

5.8.2 CNAF/Bologna

- Concerning the CPU, the request is 7,000 kSPECINT2000.day.
- We ask to increase the disk space the processed data produced by the users bringing it up to 5 TB at the end of 2006.
- About storage, the 2006 request is leading to a total amount of 67 TByte in 2006 (15 TBytes in the first half and 52 in the second half). We assume also that the policy of Bologna includes to back-up the data stored on disk.
- We ask to provide a permanent back-up of the transferred data.
- One machine with 2 Ethernet interface or 2 machines dedicated to the data transfer from Cascina up to Lyon.

5.8.3 CCIN2P3/Lyon

- The requested CPU in Lyon is the same as in Bologna: 6,000 kSPECINT2000.day, which roughly corresponds to $6 \cdot 10^6$ CCIN2P3 units. Some support will be necessary to insure a good interaction with Bologna, Rome and Naples for GRID.
- About the storage the requested space is similar to the Bologna one: 67 TByte (15 +52) in HPSS to be added to the data already stored there. In order to provide a good access to data stored in HPSS, it may be needed to provide Virgo with a new HPSS server during 2006. Besides, test of using a XrootD server will be done in the following weeks in order to check if XrootD can improve our data access. In addition to this in order to cope with the amount of data stored in HPSS, we need to have an additional tape driver. We need also 0.5 TBytes of NFS mounted disk space for storing the intermediate results of pulsar search analysis.
- We ask to provide a permanent back-up of the transferred data.

6 Trends for the future

The needed storage will obviously evolve according to the advancement of the ITF commissioning. The foreseen production is about 100 TByte per year corresponding to 200 days of scientific running of the instrument.

For Periodic Sources, the aim is to reach to 1 TFlops machine needed to perform the full sky search with a reasonable number of the spwin-down parameters. At that time, it represents about 300-400 CPUs during the whole year. It corresponds 110,000-150,000 kSPECINT2000.day. It should be underlined that the CPU power is a limiting factor for this analysis and if available more power can be used. Moreover, in order to overcome the problem of the data access limits, we need to have disk space dedicated to this analysis.

For Coalescing Binaries and Bursts group, the main effort will be to run the reprocessing of data. By reprocessing, it is meant to redo some or all the operations made in the in-time chain (calibration, h reconstruction ,data quality, triggering for Coalescing Binaries and Bursts).

It seems necessary to run faster than real-time and a factor two is a minimal value. It means that a 600GFlops machine (roughly 200 CPU) is needed for the reprocessing. If we keep the number of 200 days of scientific running, it leads to 40,000 kSPECINT2000.day for one year. It represents about one third of the Periodic Sources request.

It is obvious that the request for Stochastic Background sources will increase in the future. Unfortunately, no deep investigations have been performed up to now and the tests made in 2005 will give some indications for the future.

The last point concerns the network analysis (coherent analysis or coincidence), for which we know that the needed power for a coherent analysis goes at least as the number of involved detectors. If it is a blind search (you need to add at least the position of the source in the sky as free parameter), the needed power can rapidly grow. This field is under investigation in the context of the present collaborative activities LIGO-Virgo and Resonant Bars-Virgo. Thus, we don't report here any figure of merit about the computing needs of this task. However, it should be underlined that this kind of analysis (at least coincidence) is mandatory to insure the discovery of gravitational waves.