CNRS

Centre National de la Recherche Scientifique

INFN

Istituto Nazionale di Fisica Nucleare

Virgo Computing needs for 2007 and beyond

VIR-PLA-DIR-7000-131

Issue: 2

Date: November 17, 2006

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

CHANGE RECORD

Issue/Rev	*Date*	*Section affected*	*Reason/remarks*
Issue 1	10/11/2006		First version
Issue 2	17/10/2006	All	Updated version

Authors:	**Date**	**Signature**
A. Viceré	10/11/2006	
F. Ricci	14/11/2006	
C. Palomba	15/11/2006	
P. Hello	14/11/2006	
S. Cortese	14/11/2006	
F. Marion	14/11/2006	
B. Mours	16/11/2006	
M-A. Bizouard	16/11/2006	
A.Viceré	17/11/2006	

TABLE OF CONTENTS

1 Introduction

This note describes in short the computing resources and the required computing needs for 2007, with reasonable plans for 2008 and trends for beyond. We plan to review these requests after six months, according to the progress of the Virgo detector duty cycle, which is planned to be higher in the second half of the year.

The main Virgo computing activities covered in this document are:

- Monte Carlo Simulation
- Noise source characterization of the interferometer
- Search of continuous signals.
- Search of transient signal for short impulsive events ('bursts')
- Search of transient signal for coalescing binary events.
- 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. The data analysis has the purpose to support Detector commissioning, to monitor the status, to implement and perform the on-line/in-time analysis for all transient signals. The on-line h-Reconstruction, which consists of calibration, re-sampling and subtraction of instrumental artifacts, is also part of the computing activities performed in Cascina. Data are stored in a Storage Buffer and also saved on tapes for the purpose of crash recovery. These data are not archived in Cascina, since the permanent archival is performed at the Computing centers.

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') off-line analysis as well as for the Monte Carlo simulation. The data reprocessing, which can include improved versions of h-Reconstruction, starting again from the raw data, but which mainly consists of the search for Gravitational Wave signals, 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 TFlop/s range and the search of transient signals like coalescing binaries of a network of detectors, for which the estimated computing power lies in the hundreds of GFlop/s range. It is also to be mentioned that the production of Monte Carlo data (the so-called software injections) and the consequent analysis of Virgo data + simulated events, which is needed to estimate the efficiency of the GW searches, requires a significant increase of the computing costs.

The Virgo analysis software is suited to run on typical PC farms implemented in the national computing centers. For some of the pipelines, like the parallel implementation of matched filtering, the functionalities of MPI are requested. A large part of the pulsar search code is suitable to run in the Grid environment, while for other codes an adaptation would be needed.

At CNAF/Bologna we store the data on large disks to provide fast access for two years of data; we are also experimenting more advanced file systems (like GPFS) to cope with problems we have met with the availability of NFS mounted volumes; also the possibility to use the front-end disks of the CASTOR system, with the advantage of an automatic backup of the data, is being experimented. The use of CASTOR would be advantageous also for the possibility to simplify the data access over the GRID.

At CCIN2P3/Lyon Virgo data are stored in 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 labs resources, for off-line activities.

3 Data Production

3.1 Real Virgo Data

The progress of the Commissioning has allowed to restore, in Summer 2006, stable working conditions for the interferometer, which can now be routinely locked for long periods, thus marking the start of a noise hunting and detector characterization period, which is expected to progressively bring the Virgo sensitivity and stability to a level allowing to perform a Science Run.

To mark the progress of the interferometer, a series of Week end Science Runs (WSR), each one lasting 56 hrs, has been planned and is taking place on a monthly basis; data produced during these runs are treated as Science data, and are used not only for investigating the behavior of the detector, but also for exercising the analysis pipelines, on-line and off-line.

The WSR runs data are therefore archived at the Computing Centers, and require extensive off-line analysis, while the data produced outside the WSRs are used by the commissioning team to characterize and improve the machine behavior and are mostly analyzed at the site, although some data stretches can be marked as important and are archived.

The WSR runs will continue until the start of a continuous Science Run is possible.

For the first half of 2007, a total run-time of three months can be foreseen, including both the WSRs and the start of a Science Run; including possible limitations in the duty cycle, we could assume that this run-time will result in a data production equivalent to a two-months continuous run.

During the second half of 2007 we expect to be in a Science Run, and we can plan tentatively to run for 5 months, which could total to 4 months of data. Of course these figures will need to be reviewed in 6 months from now.

The planning for 2008 is not yet fully defined, however will have to keep into account the Virgo+ upgrades, which could start during the first half of the year, and continue for part of 2008. A possibility is that during about three months of 2008 we keep taking data continuously, then we go in a Detector upgrade and then in Commissioning mode for about six months, with monthly WSRs during the Commissiong; depending on the start of the upgrade, it could be possible (but not likely) to acquire continuous data at the end of 2008.

Overall, we could expect to acquire an integrated amount of three months of data during 2008, mostly during the first half of the year.

In addition to the run data, acquired at a compressed rate of 6 MByte/s[1], we plan to archive trend data and 50Hz data, which correspond currently to 2 Terabyte/s every 6 months, acquired at all times, not just during the runs.

Further, the search groups are currently in the process of defining Reduced Data Sets (RDS) which could be transferred with a priority to the Computing Centers, and serve for a first analysis, thus allowing to lower the requirements on the data transfer speed for the Raw Data.

The precise size of the RDS is not defined yet, but an upper limit is 10% of the total raw data, that is about 600kByte/s. this figure has been included in the estimate, keeping it separate, in the Total Virgo data column

Period	Integrated run time * [day]	Total Virgo data produced [TByte]
First half of 2007	60	31 + 1 +3
Second half of 2007	120	62 + 1 +6
First half of 2008	70	36 + 1 +4
Second half of 2008	20	10 + 1+1

In the total Virgo data column the three addenda represent respectively the raw data, the trend data and the upper limit on the Reduced Data Set.

3.2 LSC data

The start of the Science Run will also mark the start of a data exchange with the LSC. We have therefore to take into account those data in the requirements for data storage, and also in the provisions for the computing.

We do not have yet a list of the channels which will be exchanged; we presume that they will include h-Reconstructed channels for each of the 4 LSC interferometer (GEO, H4K, H2K and L4K), plus auxiliary information needed for data quality, plus vetoes. A first rough estimate could assume that each detector requires 2 channels sampled

[1] The current rate is close to 6,8 MByte/s, but methods to improve the compression factor have been performed, showing that it is possible to reduce the rate below 6 MByte/s.

at 16384 sample/s. Assuming float (4 bytes) numbers, and a compression factor of 2, this corresponds to a flux of about 64 KByte/s (5.5 GByte/day) of compressed data for each detector.

Then we can assume that all the 4 LSC detector run up to August 2007, when the Livingston 4K shut downs for Enhanced LIGO upgrade; then the two detectors at Hanford shut down in January 2008, while GEO keeps running for the whole 2008.

We can further assume that the data exchange starts with the start of our Science Run, and not during the WSRs of 2007. From all these assumptions, the following table results:

Period	LSC run time exchanged [day*#detectors]	Total LSC data exchanged [TByte]
First half of 2007	90*4	2,0
Second half of 2007	60*4+120*3	3,3
First half of 2008	180*1	1,0
Second half of 2008	180*1	1,0

These numbers should be taken with care, but show that the resulting requirements on disk space are not very heavy compared with the raw Virgo data.

3.3 Simulated data and user disk space.

In the case of simulated data, only a small fraction of these channels are needed. For the event based searches, it will be necessary to produce several instances of data (in the order of several tenths), which can however be compressed very efficiently because they consist mostly of zeroes, interspersed with relatively short events.

For the continuous searches, on the contrary, the simulated data match the flux of the h-Reconstructed channel.

In total, we anticipate that the flux of simulated data will not exceed 10% of the flux of real raw data.

Moreover specific procedures require storing the filtered data for further analysis. In particular the pulsar search requires 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 the computing centers.

The coalescing binaries search will need also disk space at the Cascina site, to cope with a transient during which it is not possible to run the off-line analysis in Bologna, until limitations on the MPI configuration are resolved.

In Cascina, the request is for 2 TByte of disk space for simulated data and user disk space.

At the CNAF/Bologna, the request is for 5 TBytes for the users.

At the CCIN2P3/Lyon, taking into account the cost and limitations of the AFS volumes, the request is for 300 GBytes. This request will be reviewed in six months and may be increased according to the needs of the users.

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 0,33 kSPECINT2000. In the following the CPU request is given as *energy*, and the unit is the kSPECINT2000.day. The conversion factor with CCIN2P3 UI is: 1kSPECINT2000.day = 480 UI.

4.1 Bologna

Computational resources have been used mainly through the submission of grid jobs and to a lesser extent using directly the local batch system.

The main use has been the analysis of the Virgo Commissioning Runs C6 and C7 for the search of periodic gravitational signals and the ‘calibration’ of the software through the injection of simulated signals in real data.

In the following we report some figures concerning the usage of resources (to Nov. 3).

Jobs completed	CPU time used (h)	Average CPU time per job (s)	KSPECINT2000*day
42635	41198	3284	1088

The extrapolation to the end of the 2006 year yields about 1300 KSPECINT2000.day energy consumption.

The usage has been lower than the request mainly because of a misunderstanding about the conversion factor between 1GHz CPUs and kSPECINT2000. The configuration problems in Bologna made also impossible to run MPI software there; this limitation affects both the pulsar group and the coalescing binaries group, and needs to be resolved in order to allow efficient off-line processing, both for real and Monte Carlo data.

Concerning data storage, about 27TByte of data are stored in Bologna, most on disks (both SAN and NAS systems) and part on CASTOR. A migration of the data older than about 2 years (i.e., run E*, C0-2) on CASTOR has been foreseen to free space needed for the WSRs.

4.2 Lyon

In Lyon all the VIRGO data are archived in the mass storage of the computing center; as of today about 36 TByte are stored in HPSS.

Since 2006, the data access has been improved with the use of XrootD.

Concerning the computing power, the following table gives the energy used in Lyon for the past five years Since 2005, Lyon is a member of the GRID Virgo Virtual Organization, and about 15% of the computing power used in Virgo for pulsar searches is obtained in Lyon by way of GRID.

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	3700	1,800,000	7
2006*	5500	2,600,000	11

* Estimate to the end of the year
** GRID counts for 1 job submitter

5 Computing needs for 2007

We stress that in 2007 the analysis will deal with real data of scientific interest, also considering the collaboration with the LSC. This implies that we expect to have a significant increase in the computing activity.

5.1 Data acquisition and in-time processing (Cascina)

In 2006, the storage buffer has been brought to a capacity of 123 TByte, with the addition of new disks and the replacement of obsolescent ones. The additional request for this year is 10 TByte, and for additional 10 TByte next year.

Concerning the CPUs, the on line farm has been upgraded, and the current configuration consists of two blocks of nodes:

the 32 old nodes consist of 2 AMD CPUs, running at 2.2 GHz; they share 8GB of RAM (4 GB/core).

The 64 new nodes carry 2 Dual Core AMD CPUs (4 cores in total), running at 2.2 GHz; they share 4GB of RAM (1 GB/core).

The farm is being re-configured right now for the needs of the Physics Groups, and will be operated mainly for Coalescing Binaries and Burst searches, as well as for Detector characterization.

5.2 Data transfer to computing centers

In order to avoid saturating the network bandwidth between Cascina and GARR-B, the Virgo data was up to Summer 2006 transferred from Cascina to Bologna and then forwarded from Bologna to Lyon.

During 2006, because of volume availability problems at Bologna, Virgo has decided to do the transfer first to Lyon, and from there forward the data to Bologna.

The data transfer to Lyon has been rendered fully automatic, using BBFTP + PERL scripts. The same development is under way for the transfer Lyon -> Bologna.

Further, Virgo is developing the use of GRID for these transfers, which should make the system more robust and flexible, but will not change our requirements for bandwidth and availability.

The Virgo collaboration has discussed the need of a prompt transfer to the Computing Centers of data to be processed off-line. However, since requesting a short maximum delay in the transfer would result in expensive requirements on the bandwidth, it has been decided that a prompt transfer (with a maximum delay of 1day) is requested only for the Reduced Data Set (RDS) which will consist of at most 10% of the raw data. For the raw data the requirement can be relaxed instead, to 1 week of maximum delay.

5.3 Continuous signals

For 2006, the Periodic Sources group plans to perform the analysis of the real data collected during the WSR runs and the science run; this implies a significant change in the computing, compared with the previous years.

On the basis of the computing "energy" spent for the analysis of C6 and C7 data, which was about 2000 KSPECINT200.day for analyzing about 3 weeks of data, we anticipate a need for 2007 of 10,000 kSPECINT2000.day in the first half of the year, to analyze about 2 months of integrated data, including WSR and the start of the Science Run. In the second half of the year we expect to increase this request to about 40,000 kSPECINT2000.day, for a total of **50,000 kSPECINT2000.day** over the 2007 year.

The group makes use of the GRID framework: the VIRGO Virtual Organization allows access to the INFN Grid resources and the Lyon computing resources, and the plan is to share the load equally between the two centers.

5.4 Burst Sources

The Burst group will continue the analysis work on WSR data, continue the collaboration with the LSC on testing coincidence and coherent analysis methods, and then start the analysis of Science Run data.

The data analysis of the real data will increase significantly compared to 2006, and we foresee an increase of the CPU consumption: **20,000 kSPECINT2000.day** is requested for 2007, to be spent in Lyon.

5.5 Stochastic Background

The Stochastic Background analysis needs have been up to now satisfied using computing facilities available in Cascina. Generally speaking the "classical" search for uniformly distributed stochastic background has small costs and its requirements also for the off-line are negligible compared with the other groups.

However, in the next year work will start on the targeted searches (the so called "radiometer" or "sky map") which are more expensive. The resulting needs will be re-evaluated in 6 months from now.

5.6 Coalescing Binaries

During 2006, the Coalescing Binaries group has exploited the resources in Cascina to perform the in-time analysis, and has implemented the algorithms on-line.

In 2006 the CB group has made an experiment to use the Computing Center in Bologna, as planned, to perform reprocessing and Monte Carlo simulation. However the impossibility to run MPI in Bologna, because of a wrong configuration of the computing farm, made it impossible to exploit the facilities there for production run.

We expect in 2007 to ramp up significantly the computing requests, for the off-line, and therefore to request **30,000 kSPECINT2000.day**, distributed between Lyon (1/3 of the request) and Bologna.

5.7 Data storage and data access in computing centers

The data in Bologna are currently stored mostly on disk, while older runs are being moved to the CASTOR tape system. The access to the volumes has to be made more reliable, and a progressive migration towards GPFS is planned, while it is being experimented the use of CASTOR, which would make the data access more uniform with the Lyon model.

In Lyon, the data are stored in HPSS, and XrootD has been successfully experimented and is now being used for the analysis. This system permits to access the data in a faster way, since it allows to retrieve only the fraction of channels requested, avoiding to saturate the HPSS servers.

5.8 Summary of the requests for 2007

Taking into account all the needs of the physics groups, a **total of 100,000 kSPECINT2000.day** is requested by Virgo in 2007. This number is shared between the Periodic Sources request (50,000 kSPECINT2000.day) and the request by the Bursts and Coalescing Binaries groups, on-line farm excluded.

In order to insure a good balance between the two centers and saving manpower, the following sharing of the resources has been decided: the Periodic Sources and the Coalescing Binaries activities will be shared between the two centers, while the Burst activity will take place mostly in Lyon.

5.8.1 Cascina requests

The request by the physics groups is to dedicate 2 TByte to simulated data and user space

5.8.2 CNAF/Bologna

- The request for the CPU usage is 50,000 kSPECINT2000.day.

- The request for Virgo data storage amounts to a total of 110 TByte in 2007 (37 TBytes in the first half and 73 in the second half), in addition to data already stored.
- The request for the users disk space, is for a total of 5 TBytes.
- It is requested to modify the farm configuration in order to allow the use of MPI.
- It is requested to allow access to data on CASTOR using the XrootD interface.
- Support is requested also for setting up the GRID based data transfer.

5.8.3 CCIN2P3/Lyon

- The requested CPU in Lyon is the same as in Bologna: 50,000 kSPECINT2000.day, corresponding to $2.4*10^7$ CCIN2P3 units.
- The request for Virgo data storage is also the same as in Bologna: 110 TByte (37 +73) in HPSS to be added to the data already stored there.
- The request for the users disk space, is for a total of 300 GBytes. This request will be reviewed in six months from now , and possibly increased according to the user needs.
- Support is requested for setting up the GRID based data transfer.

6 Trends for the future

The 2008 tentative requests that we have placed reflect the current plans for detector upgrade and re-commissioning, and are of course to be taken with some care.

After the Virgo+ upgrade, the foreseen scenario is to take data until it is possible to start the Advanced Virgo upgrades, presumably at the end of 2010.

This means that for 2009 and 2010 we expect an high duty cycle, which could lead to about 200 integrated days of science data in 2009, and roughly the same or more in 2010.

The resulting need for storage can be estimated to about 120 TByte per year, at each computing center.

Still concerning the data storage, we remind the reader that in Lyon the plan is to keep all the Virgo data, while in Bologna we keep the most recent two years worth of data. This means that eventually, the space in Bologna will start to be recycled for new data, and Virgo will have only one copy of the raw data in Lyon.

This constitutes a risk, therefore in one year from now we will re-evaluate the possible need to start vaulting an extra copy, to keep the raw data duplicated for all the science runs.

For Periodic Sources, the aim is to exploit a computing power of the order of 1 TFlop, full time, to perform the full sky search with a reasonable number of the spin-down parameters. This could correspond to the use of about 1000 1GHz CPUs during the whole year, that is about 120,000 kSPECINT2000.day. We underline that the CPU power is a limiting factor for this analysis; if more power is available, it can be used to enlarge the parameter space searched.

For Coalescing Binaries and Bursts group, the main effort will be to perform an extensive off-line analysis, which includes the reprocessing of data, and also the assessment of the search efficiency. 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). By assessing the search efficiency, it is meant to perform extensive software injections to measure the capability of the analysis to recover the signals, as a function of the signal strength.

The scaling of the event analysis cost is roughly linear with the acquisition time, and considering also the increased costs related with the collaboration with the LSC, one can foresee a yearly need of the order of 80,000 kSpecINT2000.day.

The request for Stochastic Background sources is currently very small with respect to the other searches, but could increase in the future with the implementation of the “map of the sky” search. It is impossible however to quantify the need as of today.

All the figures provided for event searches are based on the hypothesis that the collaboration with the LSC is mainly based on a coincidence analysis, plus a coherent follow-up. It is still under investigation whether it is advantageous to perform fully coherent analysis for event based searches, which would lead to significant increases in the computational cost that we are not able to quantify now.