CNRS

Centre National de la Recherche Scientifique

INFN

Istituto Nazionale di Fisica Nucleare

Virgo Computing
status and needs for 2012

VIR-0595B-11

Release: 3

Date: 10 November, 2011

VIRGO * A joint CNRS-INFN Project
Via E.Amaldi, I-56021 Santo Stefano a Macerata, Cascina (Pisa), Italy
Secretariat: Telephone (+39) 050 752 521 * FAX (+39) 050 752 550 * e-mail: W3@virgo.infn.it

CHANGE RECORD

Issue/Rev	Date	Section affected	Reason/remarks
1	12/10/2011	all	Initial release by MA Bizouard
2	01/11/2011	6	Inputs from P. Astone, G. Prodi and F. Marion
3	03/11/2011	4-5-6	S. Cortese, A. Bozzi, MAB
4			
5			

1 Introduction

This note describes in short the computing resources and the required computing needs for 2012. It also provides a record of what has been used by Virgo over the recent years.

2 Overall computing strategy

The purpose of each computing site involved in the process of the Virgo data analysis is recalled in the Virgo notes VIR-016A-08, VIR-088A-08, VIR-0640A-09 and VIR-0527D-10.
While the two national Computing Centers (CC) of CNAF/Bologna and CCIN2P3/Lyon are mainly used for off-line analysis, the EGO/Cascina site is the data production place in which the data are provisionally stored before being transferred to the final repositories of Lyon and Bologna. Nonetheless, data analysis is also performed in Cascina to support detector commissioning, to monitor the status of the machine and to perform the on-line/in-time analysis for all transient signals.

Most of the analyses are being carried at Bologna and Lyon or in LIGO Scientific Collaboration clusters (which can also be accessed via GRID), but also other computing resources are used. In particular, the use of GRID allows to transparently access other resources both completely or partially dedicated to Virgo. Among these we mention the Roma 416 cores farm, which is the official Virgo Tier-2 site, Pisa, Perugia, APC in Paris, NIKHEF and RMKI in Budapest. At the moment, GRID resources are mainly used for continuous signal and stochastic background searches.

In order to ease the access to data, the most recent data (and the most read) are stored on disk (cache disk of mass storage system). Not recently accessed are available from mass storage system. Data access to the whole set of the Virgo data is an important aspect of this computing model.

2011 has been the last opportunity for initial Virgo to take science data (run VSR4). Advanced Virgo installation is scheduled for the end of 2011 and will span few years before the first science run of advanced Virgo. In 2012 storage demands should not be as important as in the past and we do not anticipate large CPU demands neither.
In a future document we will focus on the demands for the data analysis carried out with advanced Virgo and advanced LIGO data at the horizon of 2015.

3 Data production

3.1 Real Virgo Data

Since 2007 Virgo has alternated periods of Science data taking (Science run) with commissioning or shutdown periods. Table 1 reminds the different periods. Raw data, trend data, 50Hz data, h(t) data and Reduced Data Sets (RDS). For VSR2 RDS containing ~ 20 Em channels plus Pr_B1_Acp has been produced offline to allow the Continuous Wave group to run Line Search alogorithm on LSC clusters (RDS have been transfered to LIGO). For VSR3 and VSR4, the number of channels has increased to 55. These RDS are now used both by LIGO and Virgo and we decided to store them in our Ccs. Usually raw data and h(t) data are stored in CC only when Virgo is in Science Mode. This happens during runs or astrowatch periods. On demand, rawdata data outside run or astrowatch periods can be transfered to CCs in order to save interesting data for commissioning (Cascina circular buffer is only 48 TB large). Trend data and 50 Hz data are transfered and stored permanently in Lyon and Bologna computing centers.

Name	Dates	Number of days	Rawdata rate (compressed) MB/s	Raw data volume (TB[1])	h(t) volume (TB)	RDS volume (TB)
VSR1	May 18 2007 – Oct 1 2007	136	6.2	76	4.94	0
VA1	Aug 5 2008 – Aug 18 2008	14	6.8	7.6	0.08	0
C8	Dec 15 2008 – Dec 18 2008	4	6.7	1.7	0	0
VA2[2]	Apr 10 2009 – Apr 13 2009	4	10.4	1.9	0	0
VSR2	Jul 7 2009 – Jan 8 2010	192	10.4	164	6.1	1.6
VSR3	Aug 11 2010 – Oct 20 2010	72	11.3	67	1.5	1

[1]1TB = 10^{12} Bytes following IEEE 1541-2002 standard. The unit TiB (2^{40} Bytes) still used by CCIN2P3 and EGO.
[2]Available only at CNAF (need to be transfered)

VA3	Oct 20 2010 – Jun 3 2011	224	11.3	200	0.4	0
VSR4	Jun 3 2011 – Sep 5 2011	95	11.1	96	0.5	1(*)
VA4	Sep 5 2011 – Oct 24 2011	50	11.1	40(*)	0.1(*)	0

Table 1: Virgo Science (VSR) and commissioning (C) and astrowatch (VA) runs since 2007. (*) means that we give a prediction for the end of the period. VSR3 RDS are not yet generated.
Table 2 shows the volume of 50 Hz and trend data transfered and stored in CCs. These streams are stored in HPSS since 2000 at Lyon. They are stored on disk at CNAF since a more recent date.

	Trend data (TB)	50Hz data (TB)
2007	0.45	4.3
2008	0.43	3.8
2009	0.57	5.2
2010	0.61	5.2
2011	0.62 (*)	4.6 (*)

Table 2: 50 Hz and trend data stored at CCIN2P3 and CNAF. (*) means that we give a prediction for the end of the period.

3.2 LSC data

We receive and store in the CCs, h(t) from LIGO detectors. The volume is very similar to Virgo h(t) data stream.
Numbers are given in Table 3. LIGO detectors have been shut down on 2010 oct 20th for few years of installation of Advanced LIGO. For 2011 no data from LIGO has been transfered to the CCs. The simulated data of the Software Engineering Runs (SERs) foreseen between 2012 and 2015 will have to be transferred and stored in the Ccs for offline studies. The volume will be at maximum few TB.

Name	dates	Number of days	Volume (TB)
S5 (H1,H2,L1)	May 18 2007 – Oct 1 2007	136	4.5
S6 (a+b+c+d) (H1,L1)	Jul 7 2009 – Oct 20 2010	471	10 (*)

Table 3: LSC run data volume. (*) means numbers are based on prediction. Note that we plan to delete first C00 LIGO data files that have been transfered at the beginning of the run. That represents 1-2 TB that do not appear in this table.

3.3 Data storage for 2011

After October 24th, Virgo will shut down for advanced Virgo installation. That means that no more data will be acquired in 2011 after this date. In this scenario, only a few TB of data will be transferred from now until the end of the year. The volume of rawdata and h(t) data produced in 2011 is 280 TB. The volume of trend and 50Hz data should be of 6 TB at the end of 2011.

Total: ~ 286 TB.

In VIR-0527D-10, ~200 TB of data were supposed to be produced in 2011 assuming 6 months of Science run, but no astrowatch program.

4 Data transfer

During data taking periods, the Virgo and LSC data need to be transferred to the CCs in a timely and reliable way, so as to enable the users to begin the off-line analysis work as soon as possible. The data transfer process are also in charge of producing the .ffl files in each CC as soon as files are transfered to allow Virgo users to access data.

During Science run, Virgo data (raw, trend, 50Hz and h(t) data) are transferred to the CCs in a quasi-continuous way, meaning with a delay of 1-2 days maximum. The data are transferred simultaneously from Cascina to the CCs according a star architecture owing to capacity of the EGO geographical link. Tools like bbftpPro and SRB are used for the data replication software developed for the data transfers in 2010. For advanced Virgo, the use of other GRID tools like *lcg-*

utils is considered, but the Virgo collaboration and EGO have not arrived to a final decision. On demand, other set of data, such as calibration or "interesting" data set for commissioning are transferred to CCs where they can be analysed.
To transfer LIGO data to the Cascina site and Virgo h(t) data to a LIGO site (AEI Hannover) we use Lightweight Data Replicator (LDR) of the LIGO collaboration. LIGO data need then to be transferred to the CCs with the same priority of Virgo data.

The precise location of the data in each CC is given in VIR-0566A-10. This location scheme has to be respected to avoid further data migration

The transfer of LIGO h(t) data to Cascina site and of Virgo h(t) data to LSC sites for network low latency analysis is performed also with a different tool (Cm) with a latency of some minutes.

5 Storage at CNAF and CCIN2P3 : present status

In Lyon, all data are stored in HPSS. When data are transfered from outside, they are written directly into HPSS. Some are produced locally (for instance reprocessed h(t)). In that case they are temporary stored on disk and then moved to HPSS from which users access them. The access to data files stored on HPSS is provided through 2 means: most of the jobs are using XrootD, but an access through SRB is also granted. Cache disks are necessary to provide a reliable and fast access to data
In 2011, Virgo data stored at CNAF have been migrated to GEMSS (GRID Enabled Mass Storage System). Similarly to Lyon a cache disk (gpfs disk) where the most recently read data are stored provides an easy and fast access to users. With GEMSS and HPSS, Bologna and Lyon Ccs are proposing similar storage system. Table 4 and 5 summarize the volume of data stored on the different storage systems in Lyon and Bologna.

Year	HPSS (TiB)	XrootD cache (TiB) used / available for Virgo	SRB cache (TiB) used /available for all experiments	sps (TiB) used /available for Virgo
2008	162	108 / 140	26 / (+)	2 / 3.4
2009	317	109 / 184	32 / 106	1.1 / 5.4
2010	497	162 / 184+124	32 / 203	3.6 / 5.4
2011 (oct 10th)	707	96 / 308	35 / 203	4.0 / 5.4

Table 4: CCIN2P3 storage resource utilization by Virgo. The volume "available" represents the volume that is guarantied to Virgo (resources shared by many experiments) or that is shared by a pool of experiments. For 2010 the "available" volume guarantied to Virgo takes into account what has been requested and is explicitly written.

Year	Gpfs 4 (TB) used / available for Virgo	Gpfs 3 (TB) used / available for Virgo	CASTOR (TB) used / available for Virgo	CASTOR disk (TB) used / available for all experiments
2009	190 (+) / 256	9 / 16	145 / (+)	(+)
2010 (oct 1st)	261 / 256+186 = 442	16 / 16	163 / 165	17 / 36

Year	Gpfs 4 (TB) used / available for Virgo	Gpfs 3 (TB) used / available for Virgo	GEMSS (volume on tapes)	CASTOR
2011	345 / 384	26 / 32	?	0

Table 5: CNAF storage resource utilization by Virgo. The volume "available" represents the volume that is assigned to Virgo or that is shared by a pool of experiments. For 2010 the "available" volume assigned to Virgo takes into account what has been requested. (+) means that the exact number is not known. In 2011, data in Castor have been migrated to GEMSS which is using gpfs_virgo4 as cache disk.

6 Computing and storage in 2011 and 2012

We describe in the following the computing and storage use in the previous years and give an estimation for 2011.

6.1 Computing

The computing needs for 2011, outlined in the VIR-0527D-10, were based on the actual use of computing resources during previous years and on the basis of the possible planned searches on S5/VSR1 data and S6/VSR2 data. 434,000 and 114,000 HSE06.day for CNAF and CCIN2P3 respectively were demanded. On October 1st 2011, a bit more has been consumed: ~ 695000 HSE06.day and 141000 HSE06.day. The consumption over the last years is reported in Table 6. Note that since now on we use the energy unit HES06.day instead of kSI2k.day. The conversion factor is: 1kSI2k.day = 4 HS06.day. Numbers in Table 6 have been updated accordingly.

Period	CNAF (HSE06.day)	CCIN2P3 (HSE06.day)
2007	60000	91000
2008	240000 (estimation?)	740000
2009	452600	388000
2010 (oct 1st)	220000	122000
2011 (oct 1st)	695000	141000

Table 6: CPU consumed at CNAF and CCIN2P3 over the recent years. 2010 values are for the first 9 months of 2010. (*) means numbers are based on prediction.

At CNAF, the main activity in 2011 has been the Continuous Wave search and at Lyon the all-sky search for neutron stars ring-down on S5/VSR1, MBTA Compact Binary Coalescing search in VSR2/VSR3/S6 data and the post-processing of the optical images taken by telescopes during VSR3 have been the main consumers. These searches access mainly to h(t) data streams. On the other, detchar/DQ jobs have been run in Lyon. These jobs mainly access to raw data files (in 2011 that concerned VSR3 and VSR4). No major overload is expecting before the end of the year at both centers.

Based on 2011 activity and inputs from physics groups, we have estimated the needs for 2012 as reported in Table 7. Note that Virgo Stochastic Background searches are performed on GRID Pisa cluster and/or LSC cluster. A large fraction of burst and cbc searches run by Virgo members are using LSC clusters.

	CNAF/Bologna [HSE06.day]	IN2P3/Lyon [HSE06.day]
Continuous signals	400000	0
Burst sources	0	90000
Stochastic Background	0	0
Coalescing Binaries	30000	30000
Detector Characterization	4000	4000
Total	434000	124000

Table 7: Computing needs for 2011 in HSE06.day units.

6.2 Storage

6.2.1 CNAF/Bologna and CCIN2P3/Lyon present situation - summary

The storage situation at CNAF/BOLOGNA is (October 1st):

- gpfs4 disk: 384 TB total, 0 TB free -- storage;
- gpfs3 disk: 32 TB total, 6 TB free -- user space.
- GEMSS: 754 TB total – tape

The storage situation at CCIN2P3/LYON is (October 1st):

- XrootD cache: 96 TiB total;
- HPSS: 707 TiB,
- SRB cache: 32 TiB
- sps disk: 5 TiB total, 1.4 TiB free

The storage situation by the end of 2011 is satisfactory. See Table 8 for the requested increase in storage for 2011 and 2012. An increase of 10-20 TB of data stored on tapes is foreseen before the end of 2011 to finish VA4 data storage and

few additional data production (h(t) and RDS processing). We do not anticipate any additional storage request in 2012 as both Virgo and LIGO will not produce any data. Less than 1 TB of simulated data for the advanced detectors engineering runs might be transferred. We make the following demands:

CCIN2P3/Lyon:

- no XrootD cache disk increase. The present size of what is available for all experiments is enough to access VSR4 and older dataset.
- Marginal few TB in HPSS (tapes)
- no sps disk space increase

CNAF/Bologna:

- No gpfs4 (GEMSS cache) disk increase
- Marginal few TB on tapes (GEMSS)
- gpfs_3 (user) disk: 40 TB demanded by the CW group (20 TB for Roma and 20 TB for Polgrav).

Period	CNAF/Bologna gpfs disk / CASTOR-GEMSS / user disk [TB]	IN2P3/Lyon XrootD cache / HPSS / user disk [TiB]
2009	120 / 90 / 0	44 / 190 / 2
2010	186 / 20 / 0	124 / 140[3] / 0
2011	0 / 160 / 25 (requested)	0 / 200 / 0 (requested)
2012	0 / 0 -10 / 40	0 / 0-10 / 0

Table 8: increase of storage space in the Virgo CCs since 2009.

6.2.2 EGO/Cascina site

In Cascina, the CW group asks for:

- no disk increase: the study of LIGO S6 RDS data and/or new ER data will request less than 1 TB and existing space should accommodate needs for 2012. Let's remind that for 1 year of data (1 detector), NOEMI needs 1 TB of space disk in the database.
- 8 olnodes for NOEMI. Existing but not used olnodes will be reassigned for NOEMI

[3]This is what is expected to be used in HPSS by the end of 2010.