CNRS
Centre National de la Recherche Scientifique

INFN
Istituto Nazionale di Fisica Nucleare

Virgo computing status and needs for 2013.

The Virgo collaboration

VIR-0413A-12
Date : November 7, 2012

Abstract: This note describes in short the computing resources and the required computing needs for 2013. It also provides a record of what has been used by Virgo over the recent years.

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 Virgo computing overall strategy

Past reports with complementary information are available in documents VIR-016A-08, VIR-088A-08, VIR-0640A-09, VIR-0527D-10 and VIR-0595D-11.

While the two national Computing Centers (CC) of CNAF/Bologna and CCIN2P3/Lyon are used for off-line analysis, the EGO/Cascina site is the data production place in which the data are provisionally stored (less than 6 months of data) before being transferred to the final repositories of Lyon and Bologna (official Tier-1 sites). Nonetheless, data processing is also performed in Cascina to support detector commissioning, to monitor the status of the machine, to search on-line for glitch and spectral lines, and for the low-latency gravitational wave transient search (for precisely, MBTA is a CBC search pipeline).

The off-line analyses are being carried at Bologna and Lyon, which are part of the European GRID Initiative (EGI), or in LIGO Scientific Collaboration clusters (USA, UK, Germany), 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, Warsaw, APC in Paris, RMKI in Budapest, NIKHEF. So far, GRID resources are only used for continuous signals and stochastic background searches. The all-sky continuous waves searches are now ongoing on VSR2 and VSR4 data and are computationally challenging. That is the reason why computing requests for 2013 are mainly driven by these analysis. Other searches are making use of local batch schedulers (GE in Lyon and LSF in Bologna) and data stored in high mass storage system (HPSS in Lyon and GEMSS in Bologna) are accessed through local interfaces (XrootD cache in Lyon and gpfs cache in Bologna), making the whole Virgo computing resources heterogeneous.

The preparation to advanced Virgo era (commissioning will start end of 2015) is a unique occasion to revise our computing model. We have identified a series of fundamental changes and upgrades that are necessary to better use Virgo computing resources:

- Data transfer and data book-keeping between Cascina and Virgo Tier-1 sites,
- Generalization of the use of GRID for off-line data processing in Virgo Tiers,
- LSC/Virgo tiers interoperability improvement.

Data transfer and data booking upgrades has to be discussed jointly with the Virgo collaboration, the EGO computing department (in charge of data transfer) and CNAF and CCIN2P3, taking into account that both Tiers-1 are providing different interfaces.

The Grid infrastructure (in particular its "European" implementation developed by the EGI community [http://www.egi.eu/]) appears as a natural environment (for both data transfer and data retriving) to provide users with a unique and transparent interface in both centers, and one of the first priority projects for 2013 is to test the Grid solution to make the Virgo computing schema homogeneus. The LSC computing centers use a different Grid flavor (LSC Data Grid, LDG) which is based on Condor as unique batch system. Although both EGI and LDG infrastructures use common core services (for instance the Globus software), unfortunately the two Grid environments do not yet provide a common access interface to the end user.

Another bottleneck that has been identied in the past years is the difficulty to run at CNAF and CCIN2P3 pipelines that are running on LSC clusters. In fact, some of the analysis pipelines developed in the LSC explicitly require the Condor environment and therefore are not easily portable in other systems.

The first attempt to face these problems was the implementation of a submission system based on the pilot job framework, which create a virtual Condor cluster on a Virgo farm to which jobs can be submitted from a user interface. Another alternative which we are going to test is the Pegasus Workflow Management System [http://pegasus.isi.edu/], which provides a layer for the job submission in different Grid environments.

However, we would like to enforce software developments rules that would prevent pipelines to run only on LSC clusters, as it is the case right now for most search pipelines (most of these Condor-centric pipelines are developed by LSC groups, but some are developed by Virgo users in the CBC, burst and stochastic groups).

We are now preparing the computing plan for Advanced Virgo [1] according to which new tools and Cascina computing resources upgrade should be ready by end of 2015 when advanced Virgo starts taking data. In parallel

to this, advanced detectors data preparation passes through an active participation to the engineering runs (1 month long organized each 6 months). In 2013 ER3 is scheduled January 23rd - February 20th and ER4 will take place in June. The goal for these joint LIGO-Virgo runs is to test important software infrastructure, establish procedures for software release/maintenance during advanced detector era, test detector characterization early using real subsystem data, measure progress of analysis groups on software for key science goals. During these runs we would test data transfer with GRID tools, on-line data processing and low-latency gravitational wave searches.

2 2012 status

2.1 Data Production

Data production, and the evolution during the years, has been described in the document VIR-0595D-11 [2]. In particular during the last run, the raw-data rate (compressed) was 11.1 MB/s. Science data acquisition stopped in September 2011. LIGO data (only h(t) data stream) are also transfered to CNAF and CCIN2P3. No new LIGO data have been transfered since October 2010. Recently (October 2012), Virgo users have made the demand of using LIGO S6 RDS data in Virgo computing centers. These datasets ($\sim$ 80 TB) should have been transfered earlier. The transfer will start soon and should be completed before end of 2012. Table 1 gives information about data taken by Virgo since 2007 and store at CNAF and CCIN2P3. Table 2 shows the volume of 50 Hz and trend data transferred and stored in CCs. These streams are stored in HPSS since 2000 at Lyon. They are stored at CNAF since a more recent date. We also receive at Cascina and then transfer to CNAF and CCIN2P3 LIGO h(t) data. Table 3 gives the LIGO data volumes archived in Virgo CCs so far. In the next months, we foresee to clean up Virgo data archive after a complete survey of the different volumes at CNAF and CCIN2P3.

Run name	Dates	Numb. days	Raw-data rate [MB/s]	Raw-data [TB]	h(t) [TB]	RDS [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	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
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	1.8
VA4	Sep 5 2011 - Oct 24 2011	50	11.1	56	0.2	0
Total				670	14	4.4

Table 1: Information on past (from 2007) VIRGO runs. Cx are commissioning runs; VSRx are Science runs, VAx are Astrowatch runs.

2.2 Storage at CNAF and CCIN2P3 (October 2012 status)

- CNAF

  gpfs_virgo4. Size= 368 TB. Currently used=325 TB (*)

  gpfs_virgo3. Size= 48 TB. Currently used=33 TB.

  GEMSS = 826 TB

  In 2011 data from Castor migrated to GEMSS (Grid enabled mass storage system), which uses gpfs_virgo4 as cache disk. (*) We have *temporarily* lent 62 TB to other experiments running at CNAF, which we are

Year	Trend data [TB]	50 Hz data [TB]
2007	0.45	4.3
2008	0.43	3.8
2009	0.57	5.2
2010	0.61	5.2
2011	0.67	5.1
Total	2.7	24

Table 2: 50 Hz and trend Virgo data stored at CCIN2P3 and CNAF. At CCIN2P3 these data sets have been archived since 2001.

Run name	Dates	Numb. days	Volume [TB]
S5 (H1,H2,L1)	May 18 2007 - Oct 1 2007	136	4.5
S6 (H1,L1)	Jul 7 2009 - Oct 20 2010	471	10.2
Total			14.7

Table 3: Storage of LIGO h(t) data transferred at our CCs. This does not include RDS volumes that have not yet been transfered.

now going to have back. These 62 TB have been included in the disk size and usage reported in Table 4 that shows the different storage system capacity at CNAF since 2009.

Year CNAF	gpfs4 [TB] used / available Virgo	gpfs3 [TB] used / available Virgo	Castor or GEMSS [TB]	Castor disk [TB] used / available all exp.
2009	190 / 256	9 / 16	145 (Castor)	(+)
2010 (Oct. 1)	261 / (256+186)=442	16 / 16	163 (Castor)	17 / 36
2011	345 / 384	26 / 32	750 (**)	0
2012 (Oct. 29)	325 / 368	33 / 48	826	0

Table 4: Storage at CNAF since 2009. (+) means that we don't know the exact number. In 2011 data from Castor have migrated to GEMSS, which uses gpfs_virgo4 as cache disk. (**) by the end of 2011, $\sim$ 80 TB more data were stored in GEMSS, but it was decided that these tapes cost would be reported to 2012 budget.

- CCIN2P3

  The total volume of data stored in HPSS amounts to 791 TiB[1]. Data are accessed by users mainly through XrootD, whose cache disk contains 89 TB of Virgo data. Semi-permanent users' disk space (sps) is regularly full and requires users to clean up old files. Table 5 shows the status of the different storage system at CCIN2P3 since 2008.

Table 6 shows the demands of increase of the storage at the CCs since 2009.

2.3 Computing at CNAF and CCIN2P3 (October 2012 status)

CNAF accounting system is providing[2] information in wct_hep_day and cpt_hep_day[3]. Current computing consumption at CCIN2P3 are reported[4] in HSE06.hours (CPU time not wall clock time). From now one, we

[1] 1TB=10^{12} Bytes following IEEE 1541-2002 standard. The unit TiB (2^{40} Bytes) is still used by CCIN2P3 by many monitoring tools of storage space.

[2] http://tier1.cnaf.infn.it/monitor

[3] 1 hep_day = 1 HSE06.day

[4] http://cctools.in2p3.fr/mrtguser/info_sge_rqs.php?group=virgo

Year (CCIN2P3)	HPSS [TiB]	XrootD cache [TiB] used / available Virgo	sps [TiB] used / available Virgo	SRB [TiB] used / available all exp.
2008	162	108 / 140	2 / 3.4	26 / ♯
2009	317	109 / 184	1.1 / 5.4	32 / 106
2010	497	162 / 184+124	3.6 / 5.4	32 / 203
2011	790	145 / 308	4.0 / 5.4	44 / 203
2012 (oct 2012)	791	89 / 308	3.4 / 7.8♣	27 / 270

Table 5: Storage at CCIN2P3 since 2008.♯ means that we don't know the exact number. ♣ 2.4 TB have been added on October 2012 while no Virgo request has been made for 2012.

year	CNAF gpfs disk / CASTOR-GEMSS / user disk [TB]	CCIN2P3 XrootD cache / HPSS / user disk [TB]
2009	120 / 90 / 0	44 / 190 / 2
2010	186 / 20 / 0	124 / 140 / 0
2011	0 / 160 / 25	0 / 200 / 0
2012	-16 / 100 / +16 from gpfs4	0 / 0-10 / 0

Table 6: Requests of increase of storage at CCIN2P3 and CNAF since 2009. What is requested is sometimes different from what Virgo really gets as needs sometimes change.

are reporting CPU and wall clock time for CCIN2P3. Current consumption for 2012 is given below, while Table 7 shows the evolution since 2007 of the CPU consumption at the CCs.

- **CNAF (date: January, 1st - October, 29 2012)**

  ```
  Wall clock time (wct_hep_day): 89642.8 (Total)   295.85 (Average)
  CPU time (cpt_hep_day): 47122.208 (Total)   159.736 (Average)
  ```

  At CNAF in 2012, CPU resources are used for the all-sky CW search since this summer. The searches have started but end of 2012 (Nov., Dec.) and 2013 will correspond to a processing peak. Some detector characterization work is also performed at CNAF on VSR2 & VSR3 raw-data set.

  The request at CNAF for the year 2012 was 400 kHSE06.day. This number, given that 1 core = 7.5 HS06, corresponds to 150 cores for 1 year or to 900 cores for two months, which is what we need now. We have agreed with CNAF of the use of this computing power in November and December (in particular CNAF agreed to assign us 1000 CPUs for two months). **We would thus use by the end of the year the computing power foreseen last year at CNAF for the year 2012**.

- **CCIN2P3 (date: January, 1st - September, 30 2012)**

  ```
  Wall clock time: 2,466,000 HS06.hours
  CPU time: 1,909,000 HS06.hours
  ```

  At CCIN2P3 in 2012, CPU has been mainly used for burst searches and detector characterization studies on VSR2 & VSR3 raw-data set. 103 kHS06.day have been used until October 1st while 140 kHS06.day was requested for the entire year.

year	CNAF (WCT) [kHSE06.day]	CCIN2P3 (WCT/CPU) [kHSE06.day]
2007	60	+ / 91
2008	240	+ / 740
2009	453	+ / 388
2010	162	+ / 130
2011	674	+ / 142
2012 (Oct. 29^{th})	90	103 / 80

Table 7: Evolution since 2007 of the CPU used at the CCs. + means that wall clock time numbers are not known for all years before 2012.

2.4 Cascina status

In 2012, Cascina computing resources have been used for the first two engineering runs (ER1 and ER2). Data (h(t) stream only) have been generated by re-coloring VSR3 data, simulating glitches, duty cycle of real data. These data have been kept on disk at Cascina (less 1 TB). The CBC low latency search MBTA has run, as well as few pipelines dedicated to detector characterization for testing their online implementation for Advanced Virgo era. As reminded in Section 1, at Cascina priority is given to low latency searches, commissioning activities and detector characterization processing related to commissioning and low latency searches. All other activities, demanding large computing resources are encouraged to migrate to CNAF and CCIN2P3.

3 2013 computing and storage needs

3.1 Requests for 2013 at CNAF

3.1.1 Burst

The all-sky cWB pipeline is testing a new method to suppress noise in h(t) using auxiliary channels available in raw-data. To test the performance of the method and more generally the enhancement of cWB, a full all-sky search reprocessing (covering all S5/S6/VSRx datasets) is foreseen at CNAF in 2013: 30 kHSE06.day will be needed.

3.1.2 CBC

The CBC group is now going to begin tests at CNAF, to test how their pipeline "ihope" runs under GRID (so far they have run almost all their jobs under LSC clusters). They will begin also some work with a new library "GWtools". Tests have already started, with help from CNAF personnel who installed the required libraries on 4 worker nodes (corresponding to 90 CPUs), dedicated to the test. Then, some computing power is needed to run pipelines dedicated to compact binary parameter estimation. The overall estimation, a part the dedicated nodes and queue which is part of a preliminary test, is 100 kHSE06.day

3.1.3 CW

The CW group is the most demanding group. One reason is due to the fact that the "all-sky" blind searches are computationally a major challenge. The second reason is that these searches are now processing the whole VSR2 and VSR4 data sets, generating results for publication. Thus during 2013 the Rome and Polgraw groups will both need resources to complete these searches. The scientific goals are described in VIR0026B-12.

To estimate the required CPU needs for 2013 we are using the search which is running on limited bandwidths (20 – 128 Hz in Rome) and with an under-resolved grid of the sky, to reduce the computing time needed. Extrapolating the needed computing time from this search, we have estimated the following requests for CNAF:

```
Polgraw group at CNAF: 100 kHSE06.day
Rome group: 450 kHSE06.day.
```

The total needed for the Rome all-sky CW search is 10^3 kHSE06.day, but we consider we could equally distribute the jobs at CNAF and in the Rome local cluster (450 at CNAF and 550 in Rome) The new search we will run next year will exploit a different parameter space compared to the search we are going to complete this year.

Additional user space (gpfs_virgo3) is needed: The rome group is asking 20 TB and the Polgraw groups requires 10 TB to complete the all-sky searches.

3.1.4 SGWB

No demand at CNAF

3.2 Requests for 2013 at CCIN2P3

3.2.1 Burst

For 2013, we anticipate to have at least 3 different pipelines[5] processing past Virgo and LIGO h(t) data for either finishing analyses or new developments for the advanced detectors era. Based on past experience and consumption, the different projects/teams require a total of 400 kHS06.day. As some of the activities are related to detector characterization (Omicron), 50 kHSE06.day are actually accounted for the detchar group.

3.2.2 CBC

The group will carry out here only part of the tests they plan to do at CNAF. For this reason the request is lower, $\sim$ 30 kHSE06.day.

3.2.3 CW

No requests from the CW group, if the request at CNAF are satisfied.

3.2.4 SGWB

No demand at CCIN2P3.

3.2.5 Detector characterization

Most of the data processing concerning past data sets has been done. However, new algorithms are under development for advanced Virgo. It is unlikely these developments will require large computing resources [6]. On the contrary, they might need, like in the past, an access to large volume of raw-data ($\sim$ 200 TB) stored in HPSS and additional user disk space (/sps/ semi-permanent disk) to store triggers. In addition to this, LIGO S6 RDS dataset have been required at Lyon for studies with Omicron. These dataset ($\sim$ 80 TB total) should have been transfered after S6, but were never required by Virgo users so far.

3.3 Requests for 2013 at Cascina

At Cascina, 2013 will be a transition year as no major data processing is foreseen. The main activity will concern the preparation for advanced Virgo and especially the participation to two 1-month long engineering runs (ERs). These ERs are used to test pipelines that will be used when advanced Virgo takes data. That mainly concerns detector characterization and low latency GW searches. The current computing resources (olnodes) and system is able to cope with the 2013 data processing needs. We would need only an **additional storage of 30 TB**: 20 TB for tests with NoEMi and 10 TB for tests with D-NMAPI.

[5] Omicron all-sky, STAMP all-sky, image follow-up, IPN GRB, etc ...

[6] The most demanding activity in terms of CPU is Omicron which has been mentioned in the Burst section

In parallel, the online farm and servers need to be upgraded for the next science run (the most recent machines are from 2008 and advanced Virgo is expecting to take science data not before end of 2015). A complete requirements description of the new advanced Virgo online farm and servers is required for 2013. However to cope with the sensitivity increase of advanced detectors, low latency MBTA pipeline requires new architecture machines. In accordance with the DAQ group, we would like to test multi-core machine architecture with the scheduled ERs. That's the reason why we are asking **a new machine with 32 cores with at least 4GB of memory per core** to be purchased as soon as possible.

4 Summary of requests: computing and storage in the CCs

Here a summary of needs for the year 2013. Table 8 reports the computing needs in kHSE06.day. Table 9 reports the storage needs.

Group	CNAF [kHSE06.day]	CCIN2P3 [kHSE06.day]
Burst	30	350
CBC	100	30
CW	550	0
SGWB	0	0
detchar	0	50
TOTAL	680	430

Table 8: Computing needs for 2013 in kHSE06.day.

Group	CNAF [TB] gpfs_virgo4 / tapes / gpfs_virgo3	CCIN2P3 [TB] XrootD cache / tapes / sps
Data production	0/0/0	0/0-80/0
Burst	0/0/0	0/0/0
CBC	0/0/0	0/0/0
CW	0/0/30	0/0/0
SGWB	0/0/0	0/0/0
detchar	0/0/0	0/0/0
TOTAL	0/0/30	0/0-80/0

Table 9: Storage needs (in TB) at the CCs in the year 2013. For Lyon the volume of RDS LIGO data to be transfered is 80 TB, but we also plan to save space by deleting obsolete datasets.

5 Summary of requests in Cascina

As explained before, we ask for 30 TB user disk space for application results and the purchase of a multi-core server (32 cores with at least 4GB of memory per core) for the CBC low-latency tests.

References

[1] The Virgo Collaboration, The advanced Virgo computing plan, Tech. Rep. VIR-0026B-12, Virgo (2012). URL `https://tds.ego-gw.it/ql/?c=8828`

[2] The Virgo Collaboration, Virgo computing status and needs for 2012, Tech. Rep. VIR-0595D-11, Virgo (2011).
URL https://tds.ego-gw.it/ql/?c=8652