CNRS
Centre National de la Recherche Scientifique

INFN
Istituto Nazionale di Fisica Nucleare

Virgo computing status and needs for 2014.

The Virgo collaboration

VIR-0505C-13
Date : December 3, 2013

Abstract: This note describes in short the computing needs and the required computing resources by the Virgo collaboration for 2014. It also provides a record of what has been used by Virgo over the recent years. The AdV computing model, at VIR-0129E-13, gives the description and motivations for these requests, together with some previsions on the resources which will be needed in the incoming 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 Overview of Virgo computing strategy

1.1 Computing overview

Previous documents Past reports with complementary information are available in documents [1]. The overall computing strategy for Advanced Virgo (AdV), which is clearly reflected in the requests for next and following years, has been detailed in the AdV Computing Model [2].

Data distribiton model Virgo (and AdV) has a hierarchical model for data production and distribution: The different kinds of data which is produced by the interferometer and by the auxiliary enviromental sensors firstly stored at the EGO site in Cascina. There is no permanent data storage in Cascina; As for Virgo, we foresee to install a disk buffer of 6 months of data acquisition for local access also for AdV.

The external CCs receive a copy of the full data (and the h(t) and RDS data of the LIGO detectors) (see Figure 1) and provide storage resources for permanent data archiving. They must guarantee a.) fast data access b.) computing resources c.) network links to other AdV computing resources. For this goal a robust data distribution and access framework (based on file and metadata catalogs) is a crucial point.

The collaboration manages also smaller CCs used to run part of some analyses, simulations or for software developments and tests. Data - if required by the jobs running in these centers - is transferred by the job on execution from one of the big external CCs.

Analysis pipelines During science runs the Cascina facility is dedicated to data production and to different detector characterization and commissioning analysis, which have the need to run "on-line" (with a very short latency, from seconds to minutes, to give rapid information on the quality of the data) or "in-time" (with a higher latency, even hours, but which again produce information on the quality of the data within a well defined time scale). The detector characterization activity gives support to both commissioning and science analysis.

Science analyses are carried out offline at the external CCs, with the only exception of the above mentioned low-latency searches.

Due to the fact that we analyze data jointly with aLIGO for many searches, some analysis are carried on in LSC CCs.

1.2 Advanced considerations

We report here a few considerations on the AdV CM, as these have an important impact on the work we will need to do next year to be prepared for the first run of the detector, and thus in the computing requests for the next years.

Enabling existing resources, setting up new ones First of all, to face the huge computational demands of g.w. searches in advanced detectors era (ADE), there will be the need to gather the resources of many CCs into a homogeneous distributed environment (like Grids and/or Clouds) and to adapt the science pipelines to run under such distributed environment. This is the most cost-efficient way of increasing our computing resources.

In order to ensure the smooth continuation and extension of the various LIGO-Virgo joint analysis another very important need for ADE is to provide a Grid-enabled, aLIGO-compatible Condor cluster for AdV people. This will allow for more symmetric usage of the resources and will also help better planning and of the utilisation of the clusters.

Another important task, which we started to face, is the possibility in ADE to run search pipelines in many-core architectures, such as for example on GPUs. The ADE computing needs have to be expressed and defined in a form which allows for a later efficient extension/inclusion of such many-core architectures into the clusters used for data analysis.

Pipeline development Most g.w. searches require the use of a network of g.w. detectors (at least AdV and aLIGO). As a consequence, these search pipelines must be able to run either in AdV or aLIGO CCs. It is therefore important to develop pipelines adaptable to different environments or interfaces which hide the different technologies to the users.

The most important issues of the AdV CM may be summarized as follows:

- guarantee adequate storage and computing resources at Cascina, for commissioning, detector characterization and low-latency searches;
- guarantee fast communications between Virgo applications at Cascina and aLIGO CCs/other detectors for low-latency searches;
- guarantee reliable storage and computing resources for off-line analyses in the AdV CCs;
- push towards the use of geographically and administratively distributed heterogenous resources (Grid/Cloud), whenever appropriate;
- push towards a homogeneus model for data distribution, bookkeping and access.
- be able to quickly accomodate to new hardware solutions appearing in the market in the following years

Figure 1 gives a big picture of the data workflow for what concerns scientific data analysis (DA) and detector characterization (Detchar) activities for AdV.

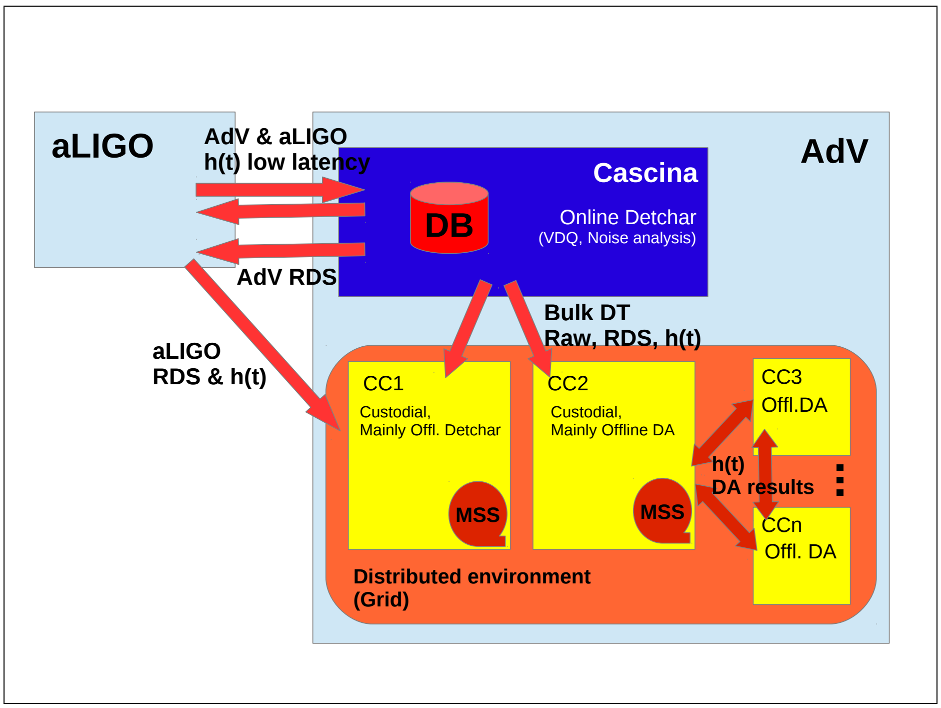

Figure 1: Data workflow for DA and Detchar activities in AdV.

2 2013 status

2.1 Data Production

As said, the AdV Computing model (CM) VIR-129E-13 [2] describes the data model and the data analysis workflows for AdV. Detailed descriptions of all the data have been put in Part II of the CM and summarized in Part V, which contains the "Computing facilities resource requirements" in Cascina and in external CCs. We recall here that data production, and the evolution during the years, have been described last year in the document VIR-0413A-12 [?]. In particular during the last Virgo 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, and RDS when needed) have also been transferred to CNAF and to CCIN2P3.

Table 1 gives information about data taken by Virgo since 2007 and stored 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 LIGO h(t) data at Cascina and transfer them to CNAF and CCIN2P3. Table 3 gives the LIGO data volumes archived in Virgo CCs so far.

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.

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.

2.2 Storage at CNAF and CCIN2P3 (November 2013 status)

Storage at CNAF The occupancy of the file systems used by the collaboration. See Figure 4 for the numbers.

In 2011 data from Castor migrated to GEMSS (Grid enabled mass storage system), which uses gpfs_virgo4 as cache disk.

Run name	Dates	Numb. days	Volume [TB]
S5 h(t) (H1,H2,L1)	May 18 2007 - Oct 1 2007	136	4.5
S6 h(t) (H1,L1)	Jul 7 2009 - Oct 20 2010	471	10.2
S6 RDS (H1,L1)	Jul 7 2009 - Oct 20 2010	471	83
Total			97.7

Table 3: Storage of LIGO h(t) data transferred at our CCs. RDS volumes (83 TB) have been transferred during this last year (2013) at CCIN2P3 only for detector characterization developments.

Name	Size (TB)	Used (TB)
Common area (gpfs_virgo4)	379	254
User disk (gpfs_virgo3)	48	47
GEMSS		826 TB

Table 4: Occupancy of various file systems at CNAF

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
2013 (Nov. 18)	254/ 379	47 / 48*	826	0

Table 5: 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. (*): 19 TB of the 30 TB we have asked for 2013 are going to be assigned and thus the available quota on gpfs3 will be 67 TB. (**) 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.

Storage at CCIN2P3 The total volume of data stored in HPSS amounts to 862 TiB[1]. Data are accessed by users mainly through XrootD, whose case disk is shared with other experiments. The maximal Virgo use of the XrootD cache in 2013 has been of 152 TiB. Semi-permanent users' disk space (sps) is regularly full and requires users to clean up old files. Table 6 shows the status of the different storage system at CCIN2P3 since 2008.

Table 7 shows the demands of increase of the storage at the CCs since 2009.

2.3 Computing at CNAF and CCIN2P3 (November 2013 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 on, we are reporting CPU and wall clock time for CCIN2P3.

Current consumption for 2013 is given below, while Table 8 shows the evolution since 2007 of the CPU consumption at the CCs. We mainly use "Wall-clock time" as this is the quantity we use to account to the CCs.

[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 HS06.day
[4] http://cctools.in2p3.fr/mrtguser/info_sge_rqs.php?group=virgo

Year (CCIN2P3)	HPSS [TiB]	XrootD cache Virgo [TiB] current / max / available	sps Virgo [TiB] current / max / available	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
2013 (Nov 2013)	862	12 / 152 /308	5.9 / 7.1 / 14.6	5 / 653

Table 6: Storage at CCIN2P3 since 2008.♯ means that we don't know the exact number. XrootD and sps available disk size is what is provided to Virgo, but XrootD is actually a shared resource. ♣ 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
2013	+11/ 0 / +19*	0 / 0 / 10

Table 7: Requests of increase of storage at CCIN2P3 and CNAF since 2009. (*): considering the 19 TB which have not been assigned yet

- **CNAF (date: January, 1st - November, 20 2013)**

  ```
  Wall clock time (kHS06.day):  850 (Total)
  ```

  At CNAF in 2013, CPU resources have mainly been used for the all-sky CW search and for the CBC Parameter Estimation and General Relativity work, in preparation of advanced detector era (ADE).

  Some detector characterization work is also performed at CNAF on VSR2 & VSR3 raw-data set. Very recently some resources are being used at CNAF to develop new improved features of the CBC low-latency pipeline MBTA.

  The request at CNAF for the year 2012 was 400 kHSE06.day. This number, using a conversion factor such that 1 core = 7.5 HS06, corresponds to using 150 cores for 1 year or to using 900 cores for two months. We have used the requested energy in the period November and December 2012, as CNAF has agreed to assign us O(1000) CPUs for two months.

  During the year 2013, an important work has been done to do the porting of CBC pipelines from an LSC related submission method to an architecture complaint also with our CCs and in particular with GRID. This has solved the limit to run CBC analyses only on LSC clusters, opening new possibilities in particular in view of ADE, to both the Virgo and LSC collaborations. Tests on real data have begun in September 2013 and to allow them CNAF has granted to Virgo a number of cores O(1000), which is the minimum needed to prepare the CBC analysis in view of ADE and to run some new CW analysis (on VSR2/VSR4 data) enlarging the parameter space covered so far.

- **CCIN2P3 (date: January, 1st - December, 31 2013)**

  ```
  Wall clock time: 115 kHS06.day
  CPU time: 63 kHS06.day
  ```

These numbers are derived from end of October comsumption extrapolated to 12 months. In 2013, CPU has been mainly used for burst searches (1 pipeline) and detector characterization studies with access to VSR2, VSR3 and S6 RDS data. Note that 430 kHS06.day was asked for 2013 and we will be using only one third of our demand in 2013.

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	669	103 / 80
2013 (Nov. 20^{th})	850	115 / 63♣

Table 8: Evolution since 2007 of the CPU used at the CCs. + means that wall clock time numbers are not known for all years before 2012. ♣ 2013 CCIN2P3 numbers are derived from end of october consumption extrapolated to 12 months.

2.4 Cascina status

In 2013, Cascina computing resources have been used for the engineering runs (ER3 and ER4). 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. As clarified in the CM, all other activities demanding large computing resources run only in external CCs.

3 2014 computing and storage needs

3.1 Requests for 2014 at CNAF

No new data will be taken next year. So no additional disk space is needed at CNAF.

The plan for next year is to continue the CW searches, increasing the range of the parameter space as much as possible: we will run on VSR2 and VSR4 two different analyses each requiring 300/600 cores. Besides this, we will run tests and simulations on Mock Data Challenge with the CBC pipeline. To continue the work the group will need a minimum of 1000 cores.

Thus for 2014, the plan is to share the (1000) cores we have asked @CNAF, cores at the Rome Tier-2 (400), cores available at Wigner (order of 200) and some resources in Poland among these two groups. In 2015 with LIGO starting data taking (3 months) analysis on new data will start, in particular in the CBC and Burst groups, and we will need more resources.

To summarize, as also detailed in the CM, the computing request for next year is to have 1000 cores, which amounts to a power of 10 kHS06, using the conversion factor 1 core = 10 HS06.

These are needed for new analysis and some reanalysis in the CW group (both Rome and Polgraw groups have plans for new analysis) and for the work on preparation of ADE in the CBC group. In particular, the CW group will be analyzing data from the Virgo run, exploiting smaller regions of the Sky and/or small frequency bandwidths for low values of τ_{min}.

The CBC group needs to run tests and analysis on MDC, to optimize the algorithms in view of ADE.

Besides this, a number of cores O(100) will be used for detector characterization work, including h-reconstructed pre-conditioning work (see details in the CM). We clarify that the power, and thus cores, to which we are referring here is the total we will need next year, that is it includes our actual pledge at CNAF (we will need only to add the needed power to our actual pledge)

3.2 Requests for 2014 at CCIN2P3

No new data will be taken, and thus no additional HPSS space (tapes) is needed at CCIN2P3. On the other hand, an increase (+5 TB) of the users' disk space is needed to store results of burst searches that will be performed in 2014. The disk space that hosts Virgo and LSC software ($THRONG_DIR) needs an increase of 10 GB to cope with the installation of new software.

The computing request for next year amounts to 400 kHSE06.day for burst searches (2 pipelines) and detector characterization activities. The developement of the main online triggers generator is done at CCIN2P3. It does not require lots of CPU but requires access to rawdata or RDS data sets.

3.3 Requests for 2014 at Cascina

At Cascina, 2014 will again be a transition year as no major data processing is foreseen. The main activity will concern the preparation for advanced Virgo and 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. Next one, ER5, will be in January 2015.

That mainly concerns detector characterization and low latency GW searches. The current computing resources (olnodes) and system is able to cope with the 2014 data processing needs.

Besides this, we need to note that the online farm needs to be upgraded for the next science run, as the most recent machines are from 2008 and advanced Virgo is expecting to take science data not before end of 2015. The process of upgrading the farm will need to begin soon.

4 Summary of requests: computing and storage in the CCs

Here we give a summary of needs for the year 2014. Table 9 reports the computing needs in kHS06.day. Table 10 reports the storage needs.

Group	CNAF [kHSE06.day]	CCIN2P3 [kHSE06.day]
Burst	0	300
CBC/CW	3600	0
SGWB	0	0
detchar	O(360) *	100
TOTAL	3600	400

Table 9: Computing needs for 2014 in kHSE06.day. the request at CNAF amounts to a number of cores O(1000) assigned to Virgo for the whole year. (*) = not added to the total as they will be taken from the cores principally assigned to CBC/CW. At CNAF, the power, and thus cores, to which we are referring here is the total we will need next year, that is it includes our actual pledge at CNAF (we will need to add to the actual pledge the power needed to arrive to the request)

5 Summary of requests in Cascina

No new requests have been done to run data analysis next year. In any case it is important to prepare the new storage and computing farm which will be operative for AdV.

Group	CNAF [TB] gpfs_virgo4 / tapes / gpfs_virgo3	CCIN2P3 [TB] XrootD cache / tapes / sps / $THRONG_DIR
Data production	0/0/0	0/0/0/0
Burst	0/0/0	0/0/5/0.01
CBC	0/0/0	0/0/0/0
CW	0/0/0	0/0/0/0
SGWB	0/0/0	0/0/0/0
detchar	0/0/0	0/0/0/0
TOTAL	0/0/0	0/0/5/0.01

Table 10: Storage needs (in TB) at the CNAF and IN2P3 CCs in the year 2014.

Bibliography

References

[1] See the following documents: VIR-016A-08, VIR-088A-08, VIR-0640A-09, VIR-0527D-10, VIR-0595D-11 and VIR-0413A-12.

[2] The Virgo Collaboration, The AdV Computing Model, 2013, "VIR-129E-13", url = https://tds.ego-gw.it/ql/?c=10325

[3] The Virgo Collaboration,

Virgo computing status and needs for 2012, "VIR-0595D-11", url ="https://tds.ego-gw.it/ql/?c=8652,

[4] The Virgo Collaboration,

Virgo computing status and needs for 2013, "VIR-0413A-12", url =https://tds.ego-gw.it/itf/tds/file.php?callFile=VIR-0413A-12.pdf