CNRS
Centre National de la Recherche Scientifique

INFN
Istituto Nazionale di Fisica Nucleare

Virgo computing status 2014 and needs for 2015.

The Virgo collaboration

VIR-0502A-14
Date : November 20, 2014

Abstract: This note describes in short the computing needs and the required computing resources by the Virgo collaboration for 2015. 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 The Advanced Virgo Computing model and previous reports

Past reports with complementary information are available in documents VIR-016A-08, VIR-088A-08, VIR-0640A-09, VIR-0527D-10, VIR-0595D-11, VIR-0413A-12 and VIR-0505C-13.

The overall computing strategy for Advanced Virgo, which is clearly reflected in the requests for next and following years, has been detailed in the AdV Computing Model [1]. The Virgo computing model is described in detail in this document, and in the earlier reports mention above, as such we do not repeat here.

1.2 Computational challenges for 2015

First of all, to face the huge computational demands of gravitational wave (GW) 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 work has started and is in a well progressed state. After a short evaluation period we found that the Pegasus workflow scheduler perfectly suits for this purpose i.e. able to provide a uniform distributed job submission framework and its scheduling and accounting system is compatible with that of LIGO's one.

Computational resource sharing between LIGO and Virgo is now given much higher priority than ever before, so we have to ensure the usability of our computing resources to LIGO colleagues. An important need for ADE is to provide a Grid-enabled, LIGO-compatible Condor cluster for AdV and LIGO scientific analysis. While Pegasus can be used for distributed job submission system, due to convenience and compatibility requirements and for more efficient workflow development, this request is still very important. In this spirit, the investigation of the use of the DIRAC (Distributed Infrastructure with Remote Agent Control) utilities is continuing in Lyon to provide a common grid enabled interface for job submission on Virgo clusters, for Condor free pipelines.

Another important task, which we started to face, is the possibility in ADE to run search pipelines in GPU clusters. Due to the very high increase of computing resources which in turn has serious financial consequences, pipeline developers must be enforced to investigate the possibility of re-implementing the most compute intensive part of their algorithms to GPUs. This will either result in a much more consolidated computing budget, or - staying with the same budget - would open the possibility to examine qualitatively new physics using the same amount of computing resources. This action is a must and have to be coordinated top-down.

The most important issues of next year computing strategy may be summarized as follows:

- guarantee adequate storage and computing resources at Cascina, for commissioning, detector characterization and low-latency searches; check and ensure that the new cluster and environment does complies with the requirements of the various groups.
- guarantee fast communications between Virgo applications at Cascina and aLIGO CCs/other detectors for low-latency searches;
- guarantee reliable storage, data transfer and computing resources for off-line analyses in the AdV CCs;
- push towards the use of geographically distributed resources (Grid/Cloud), whenever appropriate;
- push towards a unified accounting/bookkeeping system which can be used together with the LIGO solution
- push towards a homogeneus model for data distribution, bookkeping and access.
- enable the sharing of our computing resources with our Collaborators
- set up an LDG compatible Condor cluster for scientific analysis

2 Status of computing in 2014

This section shortly summarize the current status as of end of November, 2014.

2.1 Data Production

As said, the AdV Computing model (CM) VIR-129E-13 [1] 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-0505C-13 [3].

In 2014 there wasn't any scientific run, as such no new data has been generated.

2.2 Storage at CNAF and CCIN2P3

2.2.1 CNAF

The numbers below refer to the state in November 2014.
Total amount of Virgo data stored at CNAF : **XX TB**

2.2.2 CCIN2P3

In 2014, no data have been transfered or generated in Lyon. The total volume of data stored in HPSS has thus remained unchanged (862 TiB[1]). Data are accessed by users through XrootD, whose cache disk is shared with other experiments. Table 1 shows the status of the different storage systems at CCIN2P3 since 2008. HPSS, Xrootd, SRB and the EGEE storage elements are dedicated to the storage, transfer and access of the data. sps and $GROUP_DIR are used by users to store any kind of data and the $THRONG_DIR space is the Virgo software installation space.

Year	HPSS [TiB]	XrootD cache [TiB] curr./max/avail.	SRB [TiB] used/all exp	EGEE disk [GiB] used/avail.	sps [TiB]	$THRONG_DIR [GiB]	$GROUP_DIR [GiB]
2014 (Nov)	862	4/7/ 308	18 / 819	70/500	14.6	15	70

Table 1: Volumes of all storage elements at CCIN2P3 in 2014. XrootD and sps available disk size is what is provided to Virgo, but XrootD is actually a shared resource.

2.3 Computing at CNAF and CCIN2P3

The computing centers have been used for various detector characterization and scientific analsysis. Successfull attempts of local batch, remote grid and remote Pegasus submission have been performed over the year.

2.3.1 CNAF

Used computing mean power	**3915 HS06**
Used computing energy	**1272 kHS06.day**
Used computing energy extrapolated to 2014	**1429 kHS06.day**
Job efficiency (CPU time/ WCT)	**c.c. 95%**

Table 2: Computing centre usage statistics for CNAF in 2014

- **Source:** `http://tier1.cnaf.infn.it/monitor`.

[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.

The page displays the **mean computational power, units of HS06** used over the year by the Virgo Collaboration. The numbers are shown in Table 2.

2.3.2 IN2P3

Used computing mean power	**251 HS06**
Used computing energy	**76 kHS06.day**
Used computing energy extrapolated to 2014	**91 kHS06.day**
Job efficiency (CPU time/ WCT)	**c.c. 92%**

Table 3: Computing centre usage statistics for IN2P3 in 2014

- **Source:** `http://cctools.in2p3.fr/mrtguser/info_sge_rqs.php?group=virgo`

The page displays the list of resources used at any instant, or averaged over a period of time. Among these resources one can thus see the number of slots allocated for Virgo jobs. The following numbers extracted from the monthly report are given in units of HS06.day. They measure the integrated computing power of "CPU time" and not in "Wall Clock Time" for sake of comparability with CNAF. The numbers are shown in Table 3.

It is very good to see that compared to some of the previous years (2011, 2012) job efficiency significantly increased and could be regarded as optimal. This is due to the type of jobs that has been run this year: search pipeline accessing h(t) streams, while 2011 and 2012 have been dominated by detcahr pipelines accessing large volumes of raw data streams.

2.4 Cascina

3 Overview of Virgo computing resources

3.1 Free and/or in-kind resources

The available computing resources for Virgo has been listed in Table 4. The overall total numbers mean that **7283 kHS06.day computing energy** could be available for Virgo for free in 2015. This can be translated in the following way: Virgo will be able to use c.c 1500 CPU core in average for free during the year 2015. (Here a conversion factor of 1 CPU core = 13 HS06 has been used).

3.2 Additional resources

4 Computing and storage needs for 2015

4.1 Computing requests from the analysis groups

Computing need requests from the various analysis groups. Units are understood **average number of cores over the year 2015**. Depending on the type of CPU core this can be easily translated to HS06.day using

https://w3.hepix.org/benchmarks/doku.php?id=bench:results_sl6_x86_64_gcc_445

When do the translation please be aware that in some cases there is hyperthreading turned on which results a number smaller with a factor of 2 for pysical cores. In case of doubt please contact you site administrator or VDASC group.

The following listing shows the **total demands of the analysis pipelines (not the incremental!)** for 2015. The target site in parantheses are just indications of preferences not strict requirements. In any case

Lab.	Proc.	#cores	HS06	Avail.	kHS06.day	Contact	Price
Rome	Intel Xeon E5410	416	9	80%	1093		Free
Wigner-1	AMD Opteron 6376	192	17	100%	1191		Free
Polgrav	Intel Xeon E5-2665	1024	16	100%			Free
	AMD Opteron 6172	2880	16	100%			Free
	Intel Xeon E5645	2880	15	100%			Free
	For Virgo	500	15	100%	2737		Free
Nikhef	Intel Xeon L5520, 2.27GHz	1500	15	X	X		Free
	Intel Xeon E5645, 2.40GHz	1000	15	X	X		Free
	Intel Xeon E5-4620, 2.20GHz	600	10	X	X		Free
	For Virgo	250	12	100%	1095		Free
Pisa	Opteron dual core 2.4/2.6 GHz	900	10	X	X		Free
	SUN X4600 4U 8P	160	10	X	X		Free
	Dell PowerEdge 1950	128	10	X	X		Free
	For virgo	100	10	100%	365		Free
Trieste	?	100	10	100%	365		Free
Paris	?	320	10	?%	?		Free
Bari	?	?	?	?%	?		Free
Perugia	?	?	?	?%	?		Free
Total		145	13	100%	7283		Free

Table 4: Available free and/or in-kind resources in for the Virgo Collaboration.

Lab.	Proc.	#cores	HS06	Avail.	kHS06.day	Contact	Price
CNAF	?	c.c 400	4000	100%	1460	-	Already bought

Table 5: Additional resources for Virgo for the year 2015. At CNAF we have already payed for the amount of computing resources above, so it can be used free in 2015.

we have make better balance between the usage of CNAF and IN2P3 computer centers, so whenever possible analysis code should be able executed in Lyon.

4.1.1 CW

- Polinomial: **250 cores** (Lyon)
- Stroboscopic: **30 cores** (CNAF)
- Time Domain F-stat, all-sky: **500 cores** (CNAF, Polgraw)
- Time Domain F-stat, targeted: **20 cores** (CNAF)
- Frequency Hough, all-sky: **600 cores** (CNAF, Grid)
- 5-vector, Targeted: **20 cores** (CNAF)

CW total request: **1420 core**

4.1.2 CBC

- online MBTA: **300 cores** (at EGO)
- offline MBTA MDC testing: **300 cores** (CNAF, Nikhef, Grid)

- offline Tiger: **600 cores** (CNAF)

CBC total request: **1200 core**

4.1.3 Burst

- cWB: **100 cores** (CNAF)
- cWB/wavegraph: **5 cores** (CNAF)
- Cosmic string: **10 cores** (Lyon)
- STAMP AS: **30 core** (Lyon)
- GW_EM studies: **10 cores** (Lyon)

Burst total request: **170 cores**

4.1.4 Stochastic

No request has been received, but can be included later on any time since computing needs are negligible. Stochastic total request: **0 cores**

4.1.5 Detector characterisation

No request for offline computing resources. Since Virgo is not taking data in 2015, detector characterization activities can be easily performed in Cascina cluster.

4.1.6 Overall request from the analysis groups:

Summing up all the above numbers the total request for computing resources from the analysis groups results:
2800 core in average over the year in our CCs + what is offered for LSC observations runs! This more than in the previous years, beacuse so far the Virgo analysis was running on LSC resources for free.

4.1.7 Comparing resources and requests

The overall requests is 2800 core in average for 2015, while Virgo will be able to use 1500 + 400 = 1900 core for free. The difference, i.e. the incremental need for 2015 is c.c 1000 core. Because of the fact that smaller fragmented grid resources cannot be used as efficiently as bigger clusters there is a need for some safety margin, so **we estimate the total incremental need for 2015 as 1100 CPU core**.

In the view of the above numbers we foresee 3 action item

- enforce anaysis groups to reconsider their computing needs and to consider re-implementation of the computationally intensize part of their algorithms on more efficient architectures, such as GPUs
- we do have to make more serious financial investment in computing resources and support, since just as lasers and mirrors, computers are part of the detection process and cannot be ignored or neglected.
- improve the usability of our CCs for and the compatibility of our job submission systems with LIGO so that the use of computer resources become better balanced and less physical resourse is able to serve the same computing request. This requires investment in manpower.

4.2 Requests for 2015 at CNAF

No new data will be taken next year by Virgo. The possible short comissioning runs will be stored in the Cascina circular buffer and very probably deleted after some period of time.

However LIGO will have 3 month of science run and the h(t) and RDS data set has to be stored. Preliminary estimate is that we need **40 TB additional storage space** to store LIGO data.

CPu requests for CNAF (part of them should be moved to Lyon).

- Continue the Freq-hough CW searches: increasing the range of the parameter space as much as possible, furthermore Mock data challanges and and various development runs **600 core**
- Time domain F-statistics search (Polgraw group) **50 core**
- Offline Tiger analysis. (Nikhef Group) **600 core**
- Offline MBTA analysis (Annecy, Urbino Group), **300 core**
- Burst group is starting to regularly execute cWB pipeline in CNAF (Padova Group), **100 core**
- Stroboscopic search (Pisa Group) **20 core**

Besides this, a number of cores O(50) will be used for detector characterization work, including h-reconstructed pre-conditioning work (see details in the CM).

Thus for 2015, the total number of cores requested (by the analysis groups) at CNAF 1700 core, however a significant fraction of this request should be moved to our gGrid resources. We beleive that at least 40% (c.c 700) of this request should be able to run on other Grid resources or moved to the Lyon computing center, so only 1000 core in total is required for CNAF. Since Virgo already has c.c 400, **increment of 600 is requested for the year 2015 for CNAF**.

4.3 Requests for 2015 at CCIN2P3

No new data will be taken next year by Virgo. The possible short comissionning runs will be stored in the Cascina circular buffer and very probably deleted after some period of time.

However LIGO will have 3 month of science run and the h(t) and RDS data set has to be stored. Preliminary estimate is that we need **40 TB additional HPSS storage space** to store LIGO data.

- Polynomial search (Nikhef group) **250 core**
- Comic strings **10 core**
- STAMP AS **30 core**
- GW_EM studies **10 core**

Thus for 2015, the total number of cores requested for Lyon is c.c 300 cores.

4.4 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, ER6, will be in December 2014.

That mainly concerns detector characterization and low latency GW searches. The current computing resources (olnodes) and system is able to cope with the 2015 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 is ongoing and non-negligible testing activity is required to ensure that the new cluster and environment complies with the requirements of the groups.

5 Summary of requests

Here we give a summary of needs for the year 2015.

- 40 TB of additional storage space in CNAF.
- 40 TB of additional HPSS storage space in Lyon + 10 GB in $THRONG_DIR.
- 1000 kHS06.day (300 cores) at Lyon.
- 3650 kHS06.day (10 kHS06 mean power, 1000 cores) in total, as such 600 cores (6 kHS06 mean power) in addition at CNAF.
- 4 commodity type GPU at Cascina online farm from beginning of 2015.
- 30 commodity type GPU at CNAF from the Q3 2015.

Bibliography

References

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

[2] The Virgo Collaboration,

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

[3] The Virgo Collaboration,

Virgo computing status and needs for 2014, ”VIR-0505C-13”, url https://tds.ego-gw.it/itf/tds/file.php?callFile=VIR-0505C-13.pdf