Virgo computing status and needs for 2013

- The full document is VIR-0413A-12
- Date : November 8, 2012
- *Pia Astone*

  (DAC and contact person for Cascina)
- *Marie-Anne Bizouard*

  (VDAS chair and contact person for CCIN2P3)
- *Alberto Colla*

  (contact person for CNAF)

for the Virgo collaboration

General policy of Virgo computing

- We actually use the farm in Cascina for on-line, in-time and low-latency searches.
- We use the CNAF and CCIN2P3 computing centers for off-line analysis.
- We are actually using also LSC computing facilities for searches developed jointly with LSC that require the Condor architecture, but this is an issue we are now planning to solve.

(1)A first attempt to face this problem 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 (2011)

(2) Another alternative which we are going to test is the Pegasus Workflow Management System, which provides a layer for the job submission in different Grid environments

(3) Another possibility we are exploring is the development of libraries based on the OpenCL framework, which might in the future make use of GPUs, thus reducing the computational time. Particularly important for All-Sky searches.

Tests for 2) and 3) are now ongoing at CNAF, with support from them. A queue and 4 machines, 96 cores, will be dedicated to tests with OpenCL

- No data have been taken during this year, as Science data taking ended in Sept. 2011
- Two ”Engineering runs” have been done during 2012 and two are scheduled for 2013. They are 1 month long.

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 new improved tools, data access, on-line data processing and low-latency g. w. searches.

Status of the Searches

- No new Science data have been taken but still

  many analysis on past data are ongoing

  and will continue during the year 2013.
- In addition to these, we are using data from ERs or **MDC** (Mock data challenges) to test new pipelines or improvements or new features of old pipelines.
- Thus, despite the fact that no new data have been taken, Data Analysis is still ongoing and in some cases the Computational Needs are now higher than in the past (”CW search”)

STORAGE in the CCs 2008-2012

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, ~ 80 TB more data were stored in GEMSS, but it was decided that these tapes cost would be reported to 2012 budget.

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.

Computing in the CCs 2007-2012

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.

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.

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 full computing power foreseen at CNAF for the year 2012 .

Computing needs for 2013 @ CCs

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.

680 kHSE06.day @ CNAF
430 kHSE06.day @ CCIN2P3

Storage needs for 2013 @ CCs

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.

30 TB on gpfs_virgo3 @ CNAF
0-80 TB on tape @ CCIN2P3