CNRS

Centre National de la Recherche Scientifique

INFN

Istituto Nazionale di Fisica Nucleare

Virgo Computing
status and needs for 2011

VIR-0527D-10

Release: 5

Date: 19 October, 2010

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	28/09/2010	all	Initial release by MA Bizouard
2	4/10/2010	all	Inputs from Pia Astone, Alberto Colla, Fulvio Ricci, Gianluca Guidi, Giancarlo Cella, Tania Regimbau, Giovanni Prodi, Frderique Marion, Elena Cuoco, Didier Verkindt, Loic Rolland.
3	13/10/2010	all	MA Bizouard
4	17/10/2010	all	Antonella Bozzi, MA Bizouard
5	19/10/2010	All	J. Colla, A. Bozzi, MA Bizouard

1 INTRODUCTION....3

2 OVERALL COMPUTING STRATEGY....3

3 DATA PRODUCTION....3

3.1 REAL VIRGO DATA....3
3.2 LSC DATA....4
3.3 DATA STORAGE FOR 2010....4

4 DATA TRANSFER....4

5 STORAGE AT CNAF AND CCIN2P3 : PRESENT STATUS....5

6 COMPUTING AND STORAGE IN 2010 AND 2011....6

6.1 COMPUTING....6
6.2 STORAGE....6
6.2.1 CNAF/Bologna and CCIN2P3/Lyon present situation - summary....6
6.2.2 EGO/Cascina site....8

7 ANNEX A: LOCATION OF VIRGO DATA AT LYON AND BOLOGNA....8

1 Introduction

This note describes in short the computing resources and the required computing needs for 2011. 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.
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 (gpfs disk or cache disk of mass storage system). Older data (or not recently accessed) are available only from mass storage system. The challenge of this document for what concerns storage is to be able to estimate the volume of data that Virgo will produce and the volume of data that will be requested by data analysis in the coming year.

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). 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 6-month 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	74 (*)	1.5 (*)	1 (*)
VA3	Oct 20 2010 – xxx	Xx	11.3	Xx	Xx	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.

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

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	0.44 (Oct 1st)	3.8 (oct 1st)
2010 forecast	0.7 (*)	5.5 (*)

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. A shutdown of the LIGO detectors will occur on 2010 oct 20th for few years of installation of Advanced LIGO. For 2011 no data from LIGO will be transfered to the CCs. Onthe other hand, mock data challenge data set produced by LIGO might be available regularly from 2011 up to Advanced LIGO first science run. The volume is expected to be small.

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 2010

Contrary to the initial plans, VSR3 run will stop on October 20th 2010. After this date, Virgo will focus on commissioning activities until at least Christmas. However, Virgo will take Science data during nights and week-ends when the detector is stable enough (astrowatch mode) with minimal support. The Science duty cycle would be thus largely reduced. However rawdata and h(t) data will be saved hence transfered to CCs as for VA3 (Virgo Astrowatch 3). Trend and 50 Hz data will be transfered to CCs regardless the status of the interferometer

It is also foreseen to reprocess VSR3 h(t) data before the end of 2010. This will double the volume of VSR3 h(t) data. VSR3 RDS data stream (1TB) will also be produced and transfered to CC as soon as VSR3 is over.

In this scenario the volume of rawdata and h(t) data produced from january 1st up to october 20th will be 80 (Virgo)+10 (LSC) TB (end of VSR3/S6 runs). The volume of trend and 50Hz data should be of 6 TB at the end of 2010. For the period after october 20th – december 31th 2010, assuming a duty cycle of 60%, the volume of data transfered to CCs will be ~40 TB (raw) + ~1 TB (h(t)).

Total: ~ 100 + 40 TB.

In VIR-0640A-09, 237 (rawdata+h(t)) + 31 (astrowatch data) TB of data were supposed to be produced in 2010 assuming 8 months of Science run.

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. The collaboration is analysing different alternatives for

the tools to be used for the transfer. Tools like bbftpPro and SRB are to be used like the engines for the data replication software developed for the data transfers in 2010. For 2011, 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 Appendix A. 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. This task is under the responsability of Benoit Mours and Alain Masserot.

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 Bologna, the most recent data (corresponding to approximately 2 years of data) are stored on gpfs disks. Older data are then moved to CASTOR. Users access data stored on disk.
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
2009	317	109 / 184	32 / 106	1.1 / 5.4
2010 (oct 1st)	399	162 / 184+124	32 / 203	3.6 / 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 (+) / 254	9 / 16	145 / (+)	(+)
2010 (oct 1st)	261 / 254+188 = 442	16 / 16	163 / 165	17 / 36

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.

6 Computing and storage in 2010 and 2011

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 2010, outlined in the VIR-0640A-09, 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. 320,000 and 305,000 kSI2k.day for CNAF and CCIN2P3 respectively were demanded. On October 1st 2010, much less has been consumed: ~ 54000 kSI2k.day and 28000 kSI2k.day. The consumption over the last years is reported in Table 4. 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 4 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
2010 (forecast)	280000 (*)	150000 (*)

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 2010 has been the Continuous Wave search and at Lyon the all-sky search for neutron stars ring-down on S5/VSR1 and MBTA Compact Binary Coalescing search in VSR2/VSR3/S6 data 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 2010 that concerned VSR2 and VSR3). No major overload is expecting before the end of the year at both centers.

Based on 2010 activity and inputs from physics groups, we have estimated the needs for 2011 as reported in Table 5. 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 in 2010 and 2011 are using LSC clusters.

	CNAF/Bologna [HSE06.day]	IN2P3/Lyon [HSE06.day]
Continuous signals	400000	0
Burst sources	0	80000
Stochastic Background	0	0
Coalescing Binaries	30000	30000
Detector Characterization	4000	4000
Total	434000	114000

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):

- Disk: 442 TB total, 182 TB free -- storage;
- Disk: 16 TB total, 0 TB free -- user space.
- CASTOR: 163 TB total, 2 TB free – tape
- CASTOR: 36 TB total, 19 TB free – buffer disk

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

- XrootD cache: 308 TiB total, 146 TiB free;
- HPSS: 400 TiB,
- SRB cache: 32 TiB
- sps disk: 5 TiB total, 1.4 TiB free

The storage situation by the end of 2010 should be rather comfortable (less data transferred so far). See Table 8 for the requested increase in storage for 2010 and 2011. 2011 requests are based on an estimate of free space by the end of 2010. Following the scenario described in section 3.3 for the " after oct 20th " period we will have:

- ~ 25 TB of data produced by the end of VSR3 (including RDS and reprocessed h(t))
- ~ 40 TB of data produced in astrowatch mode.

Let's say that ~ 70 TB will be transfered in the CCs between oct 1st and the end of the year. The expected free space should then approximately be:

Estimated free space situation at CNAF/BOLOGNA by the end of 2010:

- Disk: 110 TB free - storage;
- Disk: 0 TB free - user space.
- CASTOR: 2 TB free – tape (if nothing changes wrt oct 1st)
- CASTOR: 19 TB free – buffer disk (if nothing changes wrt oct 1st)

Estimated free space situation at CCIN2P3/Lyon by the end of 2010:

- XrootD cache: ~100 TiB free;

In 2011, it's likely we will have a Science run (VSR4) of a few months before the Advanced Virgo shutdown. The exact number of months is still unknown. We will give predictions in the scenario of a 6 months[3] run with VSR2 equivalent duty cycle and rawdata acquisition rate (11.3 MB/s):

- 190 TB of rawdata
- 5 TB of h(t)
- 6TB of trend and 50 Hz data (assuming 2009 volumes)
- 1 TB of mock data challenge from LIGO.

Total: ~ 200 TB which would represent the largest Science run undertaken by Virgo.

Given the expected free space we expect to have at the end of 2010, and given the requests from the physic groups we make the following demands:

CCIN2P3/Lyon:

- no XrootD cache disk increase, because VSR2 raw data won't be fully accessed anymore, but only a fraction.
- 200 TB in HPSS to store VSR4 data
- no sps disk space increase

CNAF/Bologna:

- No new disk increase to store VSR4 data if we start to move VSR2 rawdata (160TB) to CASTOR/GEMSS during 2011 (can't be done in the first months of 2011 as VSR2 rawdata will be used by CW and CBC in the next 6 months). But the extra 110 TB free disk should be enough to manage the transition where we still have VSR2 rawdata on disk and VSR4 rawdata coming.
- 160 TB in CASTOR to move VSR2 rawdata by the end of 2011
- gpfs_3 (user) disk space increase of 25 TB

Period	CNAF/Bologna gpfs disk / CASTOR / user disk [TB]	IN2P3/Lyon XrootD cache / HPSS / user disk [TiB]
2009	120 / 90 / 0	44 / 190 / 2
2010	184 / 20 / 0	124 / 160[4] / 0
2011	0 / 160 / 25 (requested)	0 / 200 / 0 (requested)

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:

- 1 TB of disk per year to store NOEMI results.
- NOEMI has 6 olnodes with 4GB each. An increase of 4 GB RAM (total 8 GB) is required.

[3]Which seems to be an upper limit if Adv Virgo shutdown starts in July 2011.

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

7 Annex A: location of Virgo data at Lyon and Bologna

In order to have a rational storage, ease management and avoid to move data after they have been transferred, we propose to follow the following convention for data files storage at each CC. The first part (prefix) of the path depends on location and media.

Location/media	Prefix name
CNAF	/storage/gpfs_virgo4/virgo/data/
Lyon/HPSS	cchpssvirgo:/hpss/in2p3.fr/group/virgo/
Lyon/SRB	/ccin2p3/virgo/

Table 9: prefix names of the storage path of data at CNAF and CCIN2P3. Note that files in Lyon that appear in SRB and HPSS have exactly the same path. Only the prefix is different.

Then, depending on the data files content, we have the following path convention (example given for Lyon. See Table 10 for the exact full path both for Lyon and Bologna).

- Virgo rawdata, h(t) and RDS are stored following the generic path:

raw: <prefix>/Run/<run_name>/raw/<period>/
proc: <prefix>/Run/<run_name>/proc/[<processing_nb>]
RDS: <prefix>/Run/<run_name>/RDS/

- Virgo trend and 50 Hz data are stored following the generic path:

50Hz <prefix>/DATA/50Hz/<year>/
trend: <prefix>/DATA/trend/<year>/

Run	Data type	CCIN2P3 location	CNAF location
VSR1	raw	Run/VSR1/raw/863-865 Run/VSR1/raw/866-868 Run/VSR1/raw/869-871 Run/VSR1/raw/872-875	VSR1/rawdata/
	proc	Run/VSR1/proc/ Run/VSR1/proc/V2 Run/VSR1/proc/V3	VSR1/procdata VSR1/VSR1-HrecV1 VSR1/hrec-v2
	RDS	Run/VSR1/RDS (not there)	VSR1/VSR1/RDS (not there)
VSR2	raw	Run/VSR1/raw/931-933 Run/VSR1/raw/934-936 Run/VSR1/raw/937-939 Run/VSR1/raw/940-943 Run/VSR1/raw/943-946 Run/VSR1/raw/946-948	VSR2/raw/931-933 VSR2/raw/934-936 VSR2/raw/937-939 VSR2/raw/940-943 VSR2/raw/943-946 VSR2/raw/946-948
	proc	Run/VSR2/proc/ Run/VSR2/proc/V2 Run/VSR2/proc/V3	VSR2/proc/ VSR2/hrec_v2 VSR2/hrec_v3
	RDS	Run/VSR2/RDS/ (not yet produced)	VSR2/RDS/ (not yet produced)
VSR3	raw	Run/VSR3/raw/970 Run/VSR3/raw/966	VSR3/raw/970 VSR3/raw/966

		Run/VSR3/raw/965 Run/VSR3/raw/969 Run/VSR3/raw/968 Run/VSR3/raw/967 Run/VSR3/raw/971	VSR3/raw/965 VSR3/raw/969 VSR3/raw/968 VSR3/raw/967 VSR3/raw/971
	proc	Run/VSR3/proc Run/VSR3/proc/V2 (not yet produced)	VSR3/proc VSR3/hrec_v2 (not yet produced)
	RDS	Run/VSR3/RDS (not yet produced)	VSR3/RDS (not yet produced)
VSR4	raw	Run/VSR4/raw/978 (not yet produced) ...	VSR4/raw/978 (not yet produced) ...
	proc	Run/VSR4/proc (not yet produced)	VSR4/proc (not yet produced)
	RDS	Run/VSR4/RDS (not yet produced)	VSR4/RDS (not yet produced)
VA1	raw	Run/VA1/raw	VA1/raw (not there)
	proc	Run/VA1/proc	VA1/proc (not there)
VA2	raw	Run/VA2/raw (not there)	VA2/raw
	proc	Run/VA2/proc (not there)	VA2/proc (not there)
VA3	raw	Run/VA3/raw/971 (not yet produced) Run/VA3/raw/972 Run/VA3/raw/973 Run/VA3/raw/974 Run/VA3/raw/975 Run/VA3/raw/976 Run/VA3/raw/977 Run/VA3/raw/978	VA3/raw/971 (not yet produced) ...
	proc	Run/VA3/proc (not yet produced)	VA3/proc (not yet produced)
S5	H1	Run/S5/H-H1_RDS_C03_L2	LIGO/S5/H-H1_RDS_C03_L2 (not there)
	H2	Run/S5/H-H2_RDS_C03_L2	LIGO/S5/H-H2_RDS_C03_L2 (not there)
	L1	Run/S5/L-L1_RDS_C03_L2	LIGO/S5/L-L1_RDS_C03_L2 (not there)
S6	H1	Run/S6/H1_LDR_C02_L2/<xxxx>/	LIGO/S6/H1_LDR_C02_L2/<xxxx>/ (not there)
	L1	Run/S6/L1_LDR_C02_L2/<xxxx>/	LIGO/S6/H1_LDR_C02_L2/<xxxx>/ (not there)
trend		DATA/trend/2002 DATA/trend/2003 DATA/trend/2004 DATA/trend/2005 DATA/trend/2006 DATA/trend/2007 DATA/trend/2008 DATA/trend/2009 DATA/trend/2010 DATA/trend/2011	Apr-2006 May-2006 Jun-2006 ....
50Hz		DATA/50Hz/2002	50Hz_archive/2006

		DATA/50Hz/2003 DATA/50Hz/2004 DATA/50Hz/2005 DATA/50Hz/2006 DATA/50Hz/2007 DATA/50Hz/2008 DATA/50Hz/2009 DATA/50Hz/2010 DATA/50Hz/2011	50Hz_archive/2007 50Hz_archive/2008 50Hz_archive/2009 50Hz_archive/2010 50Hz_archive/2011

Table 10: list of locations of Virgo data files at CNAF and Bologna.