CCIN2P3 DA report

- CCIN2P3 used by Virgo since 2001 using local batch scheduler.
  - Since 2016, OSG jobs are also running in Lyon successfully.
- CCIN2P3 is a Virgo data archival center :
  - All pre-VSR1 Virgo data and LIGO h(t) since S5.
  - All data files in HPSS, accessed through xrootd (cache).

1.9 PB in HPSS

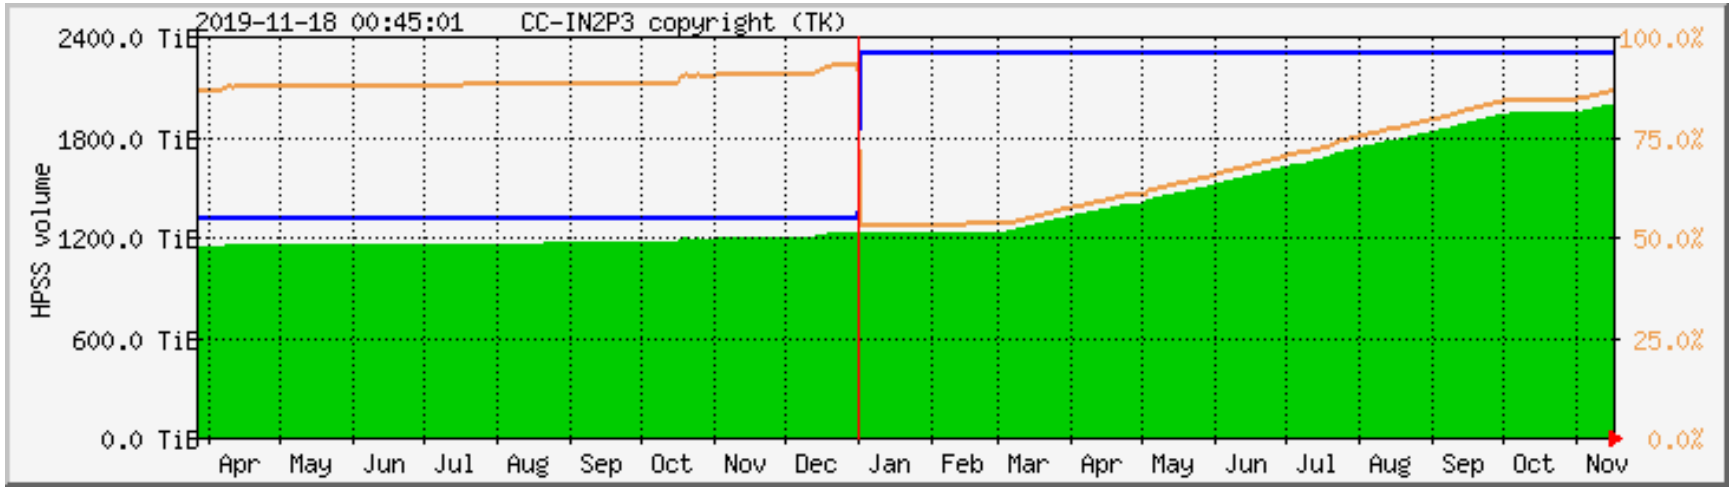

Slides prepared with the contribution of Roberto Chierici and Florent Robinet

CCIN2P3 CPU usage over years

Grid Engine jobs

OSG jobs starting in 2016

Année	Heures demandées *	Heures consommées *	CPU consommé *	Coeurs par an	Efficacité
2001			18 037		
2002	80 000		50 657		
2003	80 000		67 620		
2004	400 000		493 931		
2005	400 000		384 200		
2006	1 200 000		486 179		
2007	4 800 000		2 187 970		
2008	8 000 000		17 708 210		
2009	52 000 000		9 319 234		
2010	28 000 000		3 109 571		
2011	2 736 000		3 698 411		
2012	3 000 000	3 556 409	2 572 103	36	72 %
2013	10 320 000	2 354 232	1 290 264	24	55 %
2014	9 600 000	1 988 226	1 835 473	20	92 %
2015	24 000 000	2 800 212	2 313 862	29	83 %
2016	5 000 000	3 963 429	3 456 280	40	87 %
2017	16 000 000	10 797 240	9 371 450	110	87 %
2018	20 000 000	9 692 312	4 495 214	99	46 %
2019	20 000 000	9 859 502	7 638 138	101	77 %

CCIN2P3 usage in 2019

- 2019 CPU demands (HS06 hours unit): 20 Mhours HS06

Group name	Unix group	1st quarter			2nd quarter			3rd quarter			4th quarter			Overall 2019		
		Objectives	Consumed	Ratio	Objectives	Consumed	Ratio	Objectives	Consumed	Ratio	Objectives	Consumed	Ratio	Objectives	Consumed	Ratio
Virgo	virgo	5 000 000	309 665	6,2 %	5 000 000	1 181 748	23,6 %	5 000 000	0		5 000 000	0		20 000 000	1 491 413	7,5 %

- What is running in Lyon ?
  - OSG jobs launched from LSC clusters (load has increased since september)
  - O3 data quality flags reprocessing (raw data access)
  - MBTA offline for a contribution to the O3a/O3b catalog papers.

Temps de résidence normalisé en HS06.h	
0	
4 237 719	MBTA
1 829 590	OSG
3 792 196	Detchar
9 859 505	Total

Consumption : 10 Mhours HS06
Requested : 20 Mhours HS06

MBTA offline for Compact Binary Coalescence catalog paper

- New comer in 2019
- MBTA is a multi-threaded algorithm optimised for low latency processing using multi-core machines @ Cascina.
- MBTA offline pipeline still under development to run @ CCIN2P3 (use of GE batch scheduler).
- Science/program :
  - Search for CBC signals in BNS, NSBH and BBH mass range in H1, L1 and V1 data.
  - Contribution to the O3 CBC catalogs.

Slides from R. Chierici

MBTA : analysis principle

- **A signal in data is searched on the basis of known waveforms (« templates ») that mainly depend on coalescence parameters like masses and spins.**
  - The search is using matched filtering (each template applied to data)
  - This is done, typically, in (two) different frequency bands : technique us to reduce the needed CPU by one order of magnitude

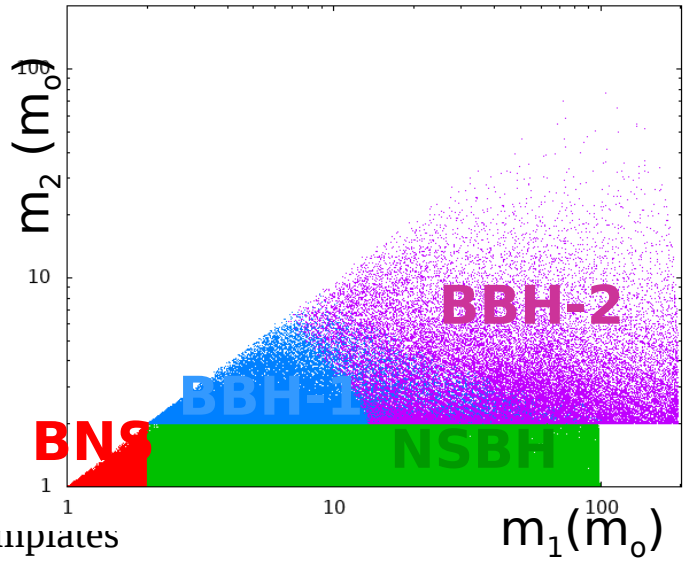

- **Three regions are identified in the (m_1, m_2) plane**
  - **BNS**, analyzed in 2 frequency bands, 26170 templates
  - **NSBH**, analyzed in 2 frequency bands, 522934 templates
  - BBH
    - **BBH-1**, analyzed 2 frequency bands: 133198 templates
    - **BBH-2** (m_{tot} > 110 $M_\odot$ or Δt < 0.2 s), analyzed in 1 band: 36277 templates

- **A multitude of jobs are needed to chop out the phase space (grand total of about 750K templates)**
  - For one pass over data typically 500 jobs cover the full parameter space. multi-threading (4 threads) is implemented
  - Typical CPU per job: 50-100 CPU hours (need a factor 5 for CC HS06 CPU units)

- **Signal search aside, the analysis must be characterised (and tuned) over « Monte-Carlo »**
  - Need injection campaigns -and relative runs- over the data for all regions of mass parameters (BNS, BHNS, BBH), sometimes for other parameter scans (e.g. spins) and for different generators

- **Need, therefore, to account for several passes over the same data period (as defined in LVC, typically ~10 days-2 weeks)**
  - One pass for signal searches and background/fake alarm characterization, one per injection set (details in next slide)

MBTA

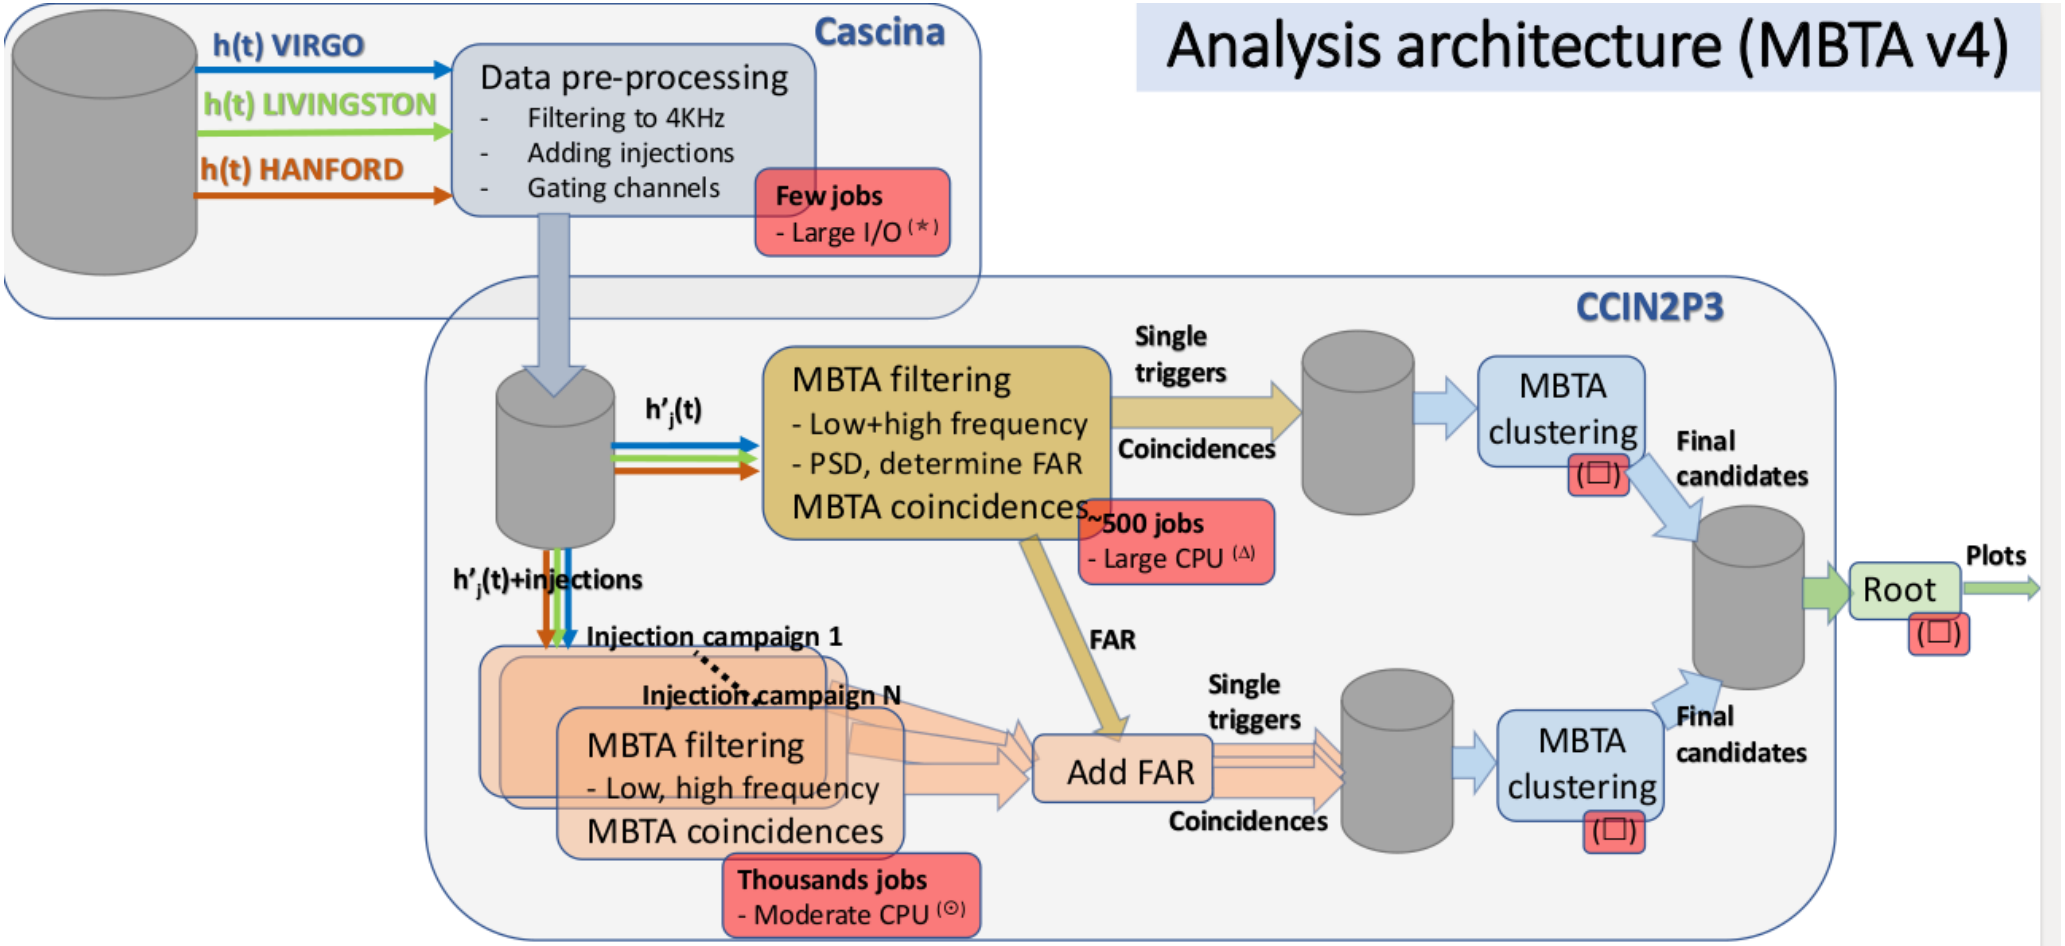
Analysis architecture (MBTA v4)
Cascina
h(t) VIRGO
h(t) LIVINGSTON
h(t) HANFORD
Data pre-processing
- Filtering to 4KHz
- Adding injections
- Gating channels
Few jobs
- Large I/O (*)
CCIN2P3
h'j(t)
MBTA filtering
- Low+high frequency
- PSD, determine FAR
MBTA coincidences
Single triggers
Coincidences
~500 jobs
- Large CPU (Δ)
MBTA clustering
(□)
Final candidates
Root
(□)
Plots
h'j(t)+injections
Injection campaign 1
Injection campaign N
MBTA filtering
- Low, high frequency
MBTA coincidences
FAR
Add FAR
Single triggers
Coincidences
MBTA clustering
(□)
Final candidates
Thousands jobs
- Moderate CPU (⊙)

Slides from R. Chierici

MBTA : batch scheduler needs

- **The first step in the analysis consists in pre-processing the data**
  - Few jobs, but large I/O needed.

→ Can this be handled in a computing center with access to data?

- **At every pass over the data in order to get a sound answer, all jobs MUST be recovered**
  - No holes in phase space, hence resubmitting dead jobs is mandatory
- **There is some hierarchization between the signal search without and in presence of injections**
  - Need to complete a full pass over data for the characterization of the false alarm rate to be used in the injection studies

→ Resources should be stable over time to avoid delays in queueing

- **The aim is to be able to fully process N days of data taking in ~N/2 CPU days**
  - The rest of the time will be needed for the understanding/elaboration of the results and their presentation

→ Resources should be enough to grant constant process of data (assessment in next slide)

- **Disk space**
  - Having h(t) data locally is helpful if the pre-processing can be handled in situ
  - Having pre-processed data requires ~500 GB/week (including injection channels). Results from the full analysis approximately the same amount → 50 TB for 2020 should be ok (shared resources preferred to mass storage)
  - Need a disk area with LVC credential access for regularly publishing the results of the analysis

-->Disk space should likely not be an issue. If needed, and if no latency introduced in the analysis, we can adapt to any I/O policy

Slides from R. Chierici

MBTA (v3) CPU estimation for analyzing O3

- **Results for one pass over the data at the CC, according to the MBTA output**
  - Just one week of data taking

	Jobs	Total/per job Wallclock (hours)	Total/per job CPU (hours)	Band 0 total/per job	Band 1 total/per job	Search total/per job
BNS	30	678 / 22.6	3251 / 108	413 / 14	2558 / 85	166 / 6
BBH	126	1975 / 15.6	5416 / 43	435 / 3.5	4244 / 34	425 / 3.4
BHNS	150	2348 / 15.7	6824 / 45	1539 / 10	4510 / 30	393 / 2.6

- **Grand total**
  - 15500 CPU hours, but need different correction factors
    - Need of a factor ~2 to get this done in half the time
    - Need another factor >2 for the injection passes over the data
  - = 62 khours HS05/week. It makes a total « energy » of 15.5 Mhours (with 50 weeks/year and a factor 5 for CC HS06 CPU units)
  - No further safety factor accounted for
    - several data processing (e.g. new calibrations), more Monte-Carlo tests, dead-time at the CC, unexpected problems/errors….

Omicron reprocessing at CC-IN2P3 – Context –

Slides from F. Robinet

→ Omicron is an event trigger generator based on the Q-transform
= short Fourier transforms using Gaussian windows with variable widths

→ Omicron is running with low latency (~10s) in Cascina over *h(t)* and ~400 auxiliary channels

→ Omicron triggers are used:
– for noise investigation
– to generate online data quality flags
– to characterize gravitational-wave candidates

→ The online trigger production is not perfect:
– missing input data,
– failure in the online architecture: data flow, processing, disk access, etc.
– the available disk space is limited
– the number of channels is limited
– the omicron resolution is limited (frequency, SNR) to minimize the latency

→ Online Omicron trigger files are short-lived. They are deleted after ~1 year

Omicron reprocessing at CC-IN2P3 – Offline analysis –

Slides from F. Robinet

→ Virgo raw data is reprocessed @CC-IN2P3

→ All Virgo channels with a sampling frequency larger 500 Hz are processed
  → a total of ~3000 channels are processed

→ The configuration of Omicron covers a larger parameter space
  – higher resolution
  – lower SNR threshold
  – larger frequency range

→ Omicron triggers are used
  – to generate CAT1 data quality flags (when the *h(t)* trigger rate is very high)
  – to run UPV and generate event-based vetoes
  – for offline noise investigation

→ Omicron trigger files are archived on tapes

→ Current status: O3a data (April 1st – October 1st) is reprocessed

Omicron reprocessing at CC-IN2P3 – Workflow –

Slides from F. Robinet

Input data

→ Virgo raw data are transfered to CC-IN2P3 via iRods and saved on tapes (hpss)
→ Continuous and automatic pre-staging of the raw data in xrootd cache disks (Rachid Lemrani @CC-IN2P3)
→ The data is analyzed by chunks of 1–1.5 months: 5 chunks to cover O3a

~600 TB rawdata

Chunk analysis

→ Download the analysis segments (nominal lock) from the LIGO-Virgo segment database (DQSEGDB)
→ Run a script to prepare the jobs for this chunk:
  – extract channel list from a frame file for that chunk
  – slice up the analysis segments in segments of ~1 week
  – generate omicron job bash scripts: (number of channels) x (number of weeks) = 18000 jobs
→ Submit the bash scripts via the CC batch Grid Engine system
→ Check the log files for completed jobs
→ Merge resulting trigger files (root format) to bigger files (at least 500 MB) to meet hpss requirements
→ Trigger files are transfered to tapes (hpss)

Omicron reprocessing at CC-IN2P3 – Conclusions –

Slides from F. Robinet

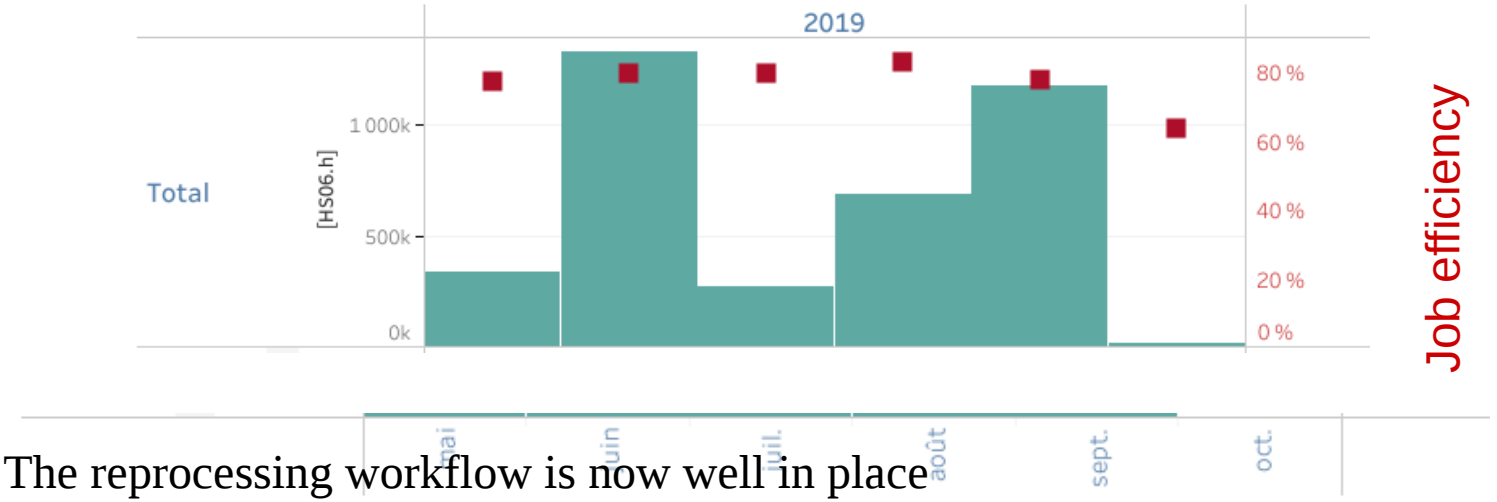

→ The reprocessing workflow is now well in place

→ The workflow has been automatized as much as possible

→ Remaining manual actions
- Check data integrity and completeness (first processing with h(t) data only)
- Submit the jobs by small amount at a time
- Run the merge script interactively

→ Recent troubleshooting:
- some cache data (xrootd) may disappear before the end of the reprocessing
- CC-IN2P3 updates: new xrootd environment
- CC-IN2P3 maintenance break: some jobs are lost

Conclusion : needs for 2020

- CPU :
  - MBTA : 30 M HS06 hours
  - Omicron : 4 M HS06 hours
  - OSG jobs : 4 M HS06 hours ?
    - → Grand total : 38 M HS06 hours
- Sps disk (user)
  - MBTA : increase of 40 TB
- Tapes : volume for the end of O3 : ~600 TB

Conclusion : trends

- Medium term :
  - More users/ new pipelines ?
    - Few pipelines mean less flexibility to accomodate peak CPU request for a given pipeline (yearly average energy request).
    - Next pipeline candidate ? *stampas* to search for long duration GW transients
  - More OSG jobs
    - Need more information from the LSC to better estimate the growth
- Longer term :
  - Keep CCIN2P3 as one of the data archival center for rawdata reprocessing
  - Harmonize the access to resources over CCs using the IGWN proposed architecture