VIRGO COMPUTING

Stefano Bagnasco, INFN

LOOKING FOR...

Transient sources:

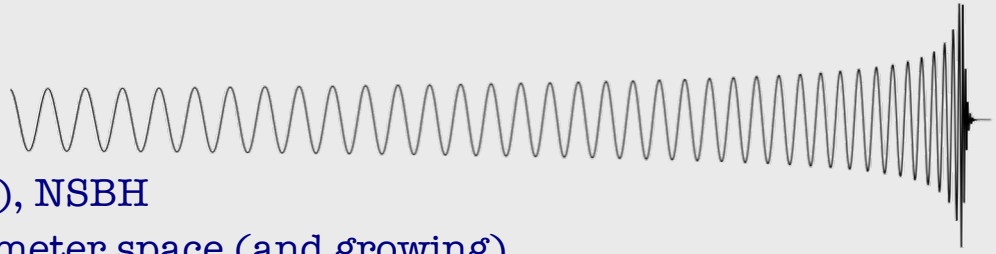

- CBC: Compact Binary Coalescence
  - Binary Black Holes (BBH), Binary Neutron Stars (BNS), NSBH
  - Strongest emitters, well modelled for much of the parameter space (and growing)
  - Matched filtering very effective

- Burst: Unmodeled transient bursts
  - E.g., Core Collapse Supernovae (CCS, and anything else)
  - Weaker and no (or poor) model, so coherence methods more effective

Continuous sources:

- CW: Continuous waves
  - E.g., Asymmetric spinning neutron stars
  - Usually well-modelled
  - All-sky and targeted searches

- SGWB: Continuous stochastic background
  - Astrophysical & cosmological

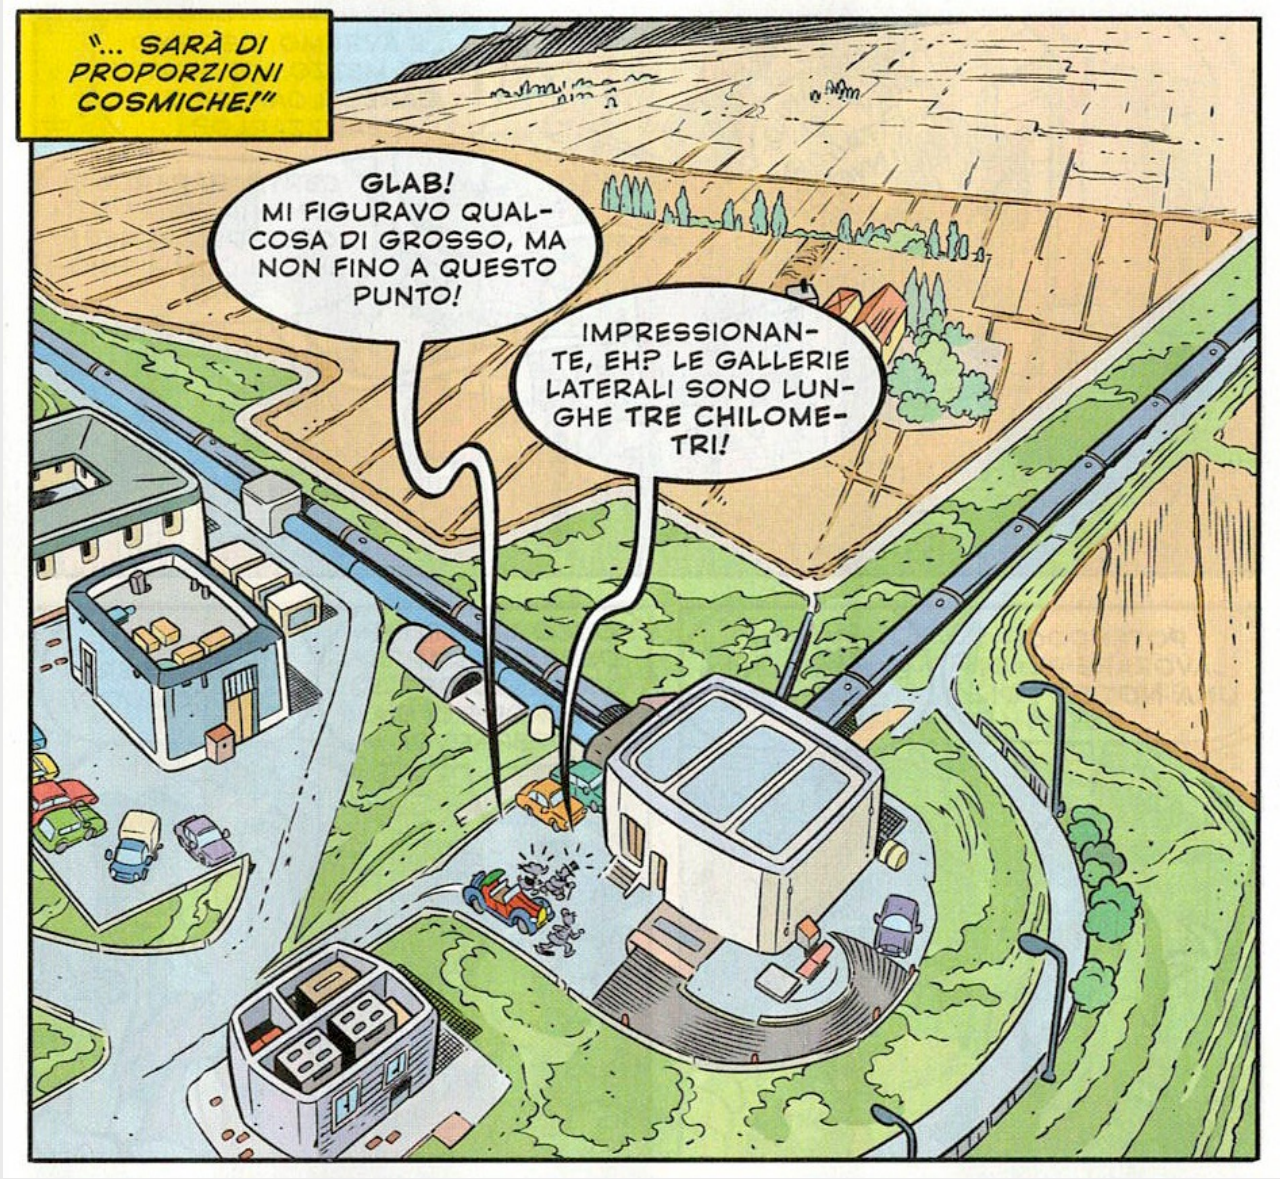

M. Bosco, G. Soldati, "Sogni d'oro zio Paperone"
Topolino **3538**:45-70 (2023)

FROM DISCOVERY TO OBSERVATION

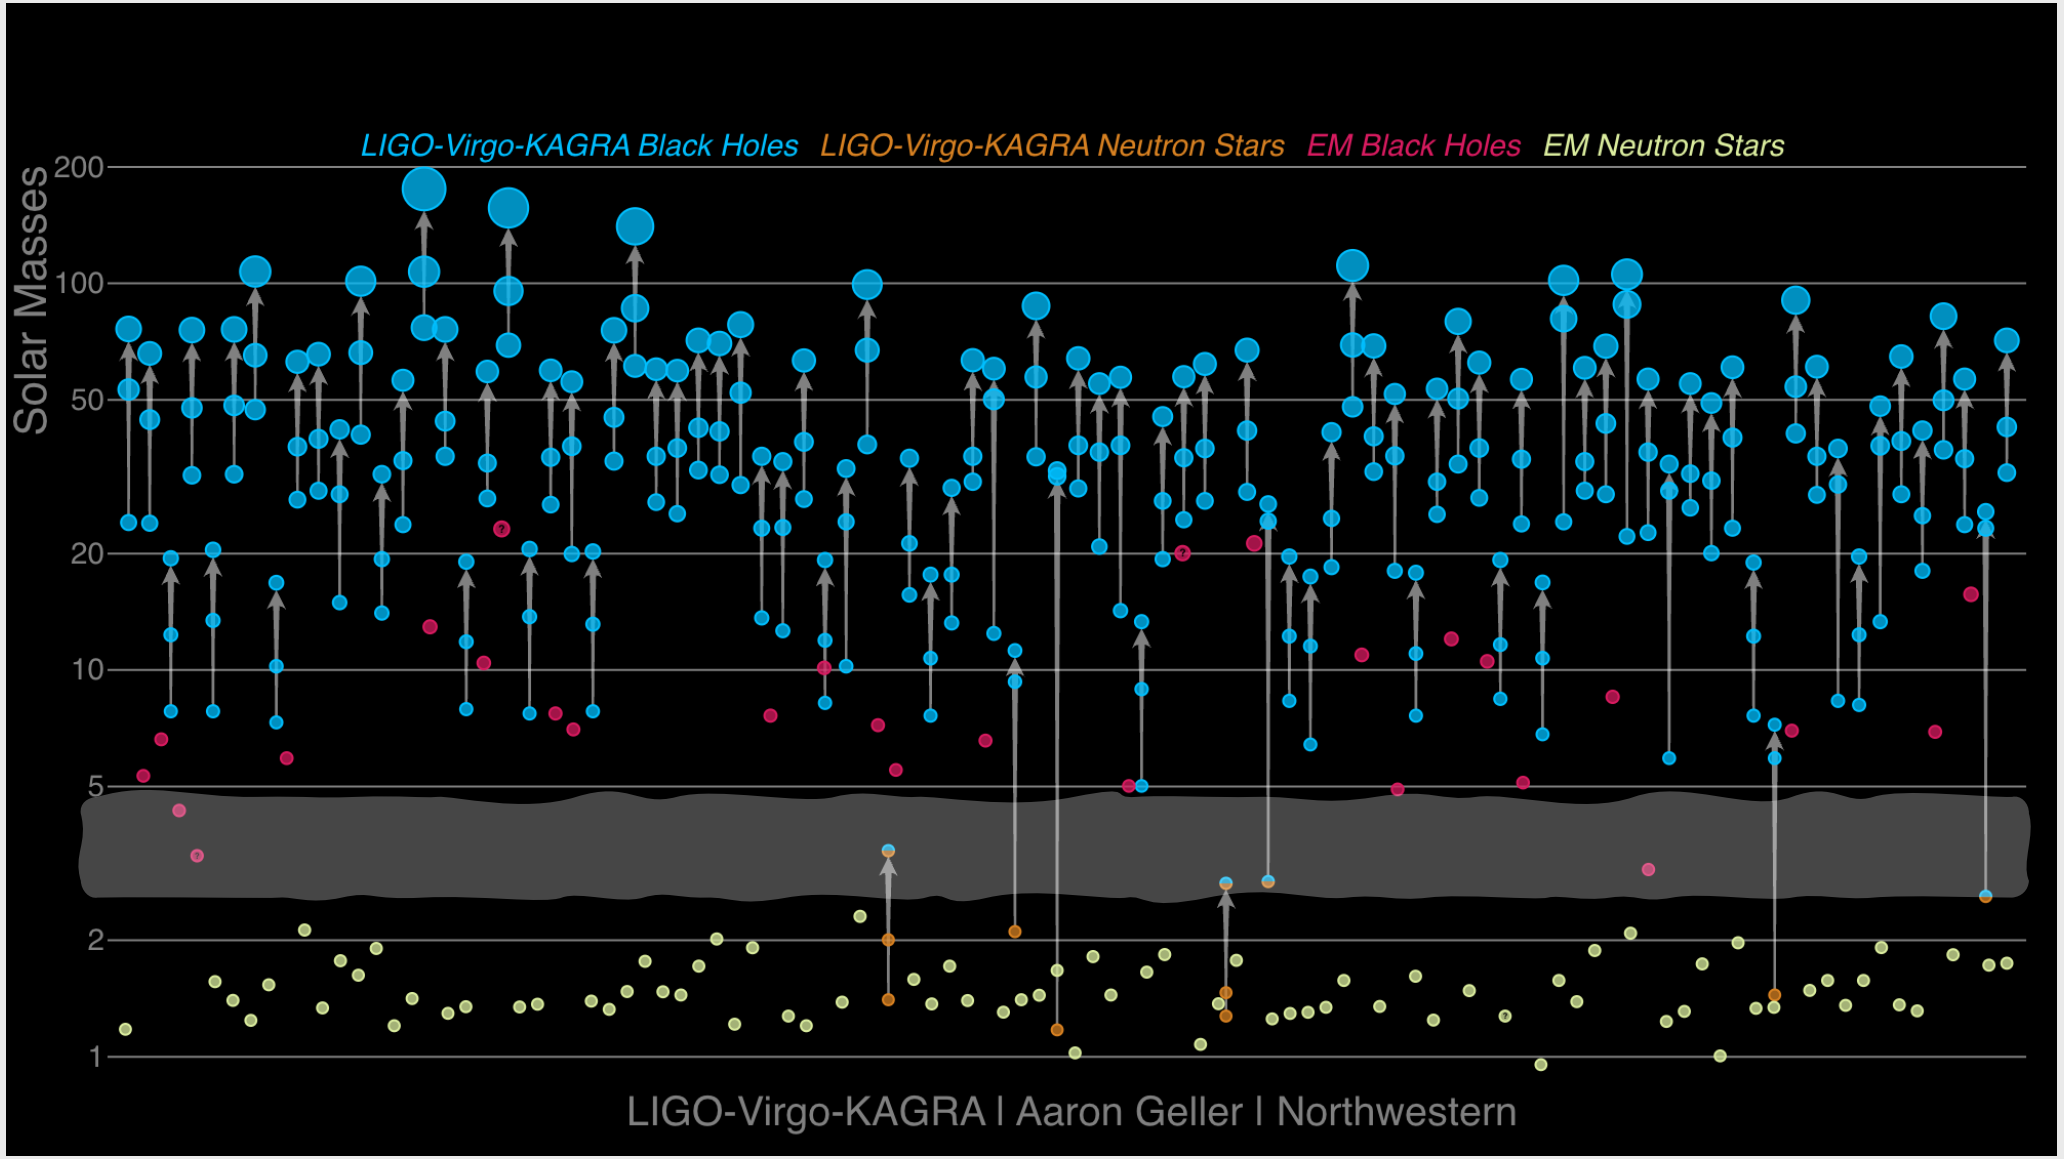

https://media.ligo.northwestern.edu/gallery/mass-plot

THREE COMPUTING DOMAINS

On-site infrastructure

Online

- Data acquisition and pre-processing
- Instrument control
- Environmental monitoring
- …

Plain old HTC (and some HPC)

Offline

- Deep searches
- Offline parameter estimation
- Detector Characterization (DetChar)
- (Template bank generation)

Here's the fun

Low-latency

- Candidate search
- Sky localization
- LL parameter estimation
- Alert generation and distribution

- A large number of common operations and coordination boards
  - Biweekly IGWN CompComm is the main coordination forum
  - Weekly/Biweekly LIGO and Virgo Dev/Ops calls (with joint participation)
  - Weekly OSG-IGWN prod calls
  - IGWN Software Change Control Board
  - ...but see also last slides on new organization!
- F2F sessions at LVK Collaboration meetings
- Semi-annual F2F IGWN High-Throughput at Georgia Tech
  - And, hopefully, in the EU sooner or later...
- Other thematic F2F meetings

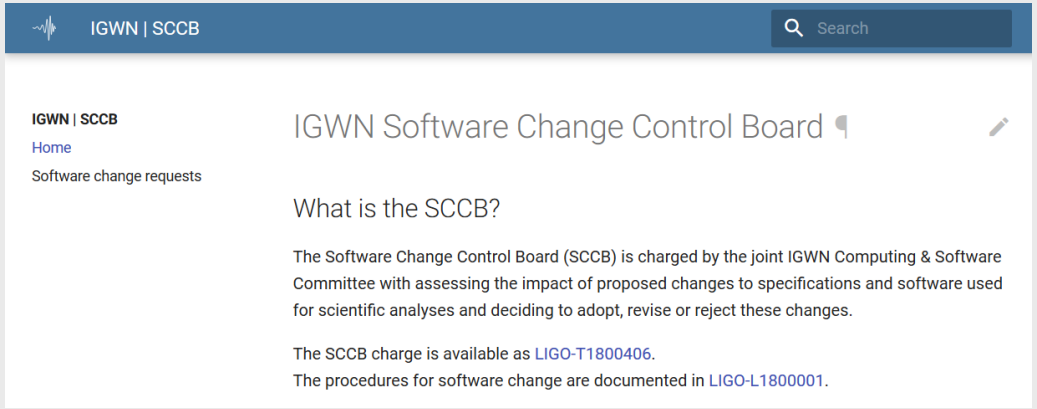

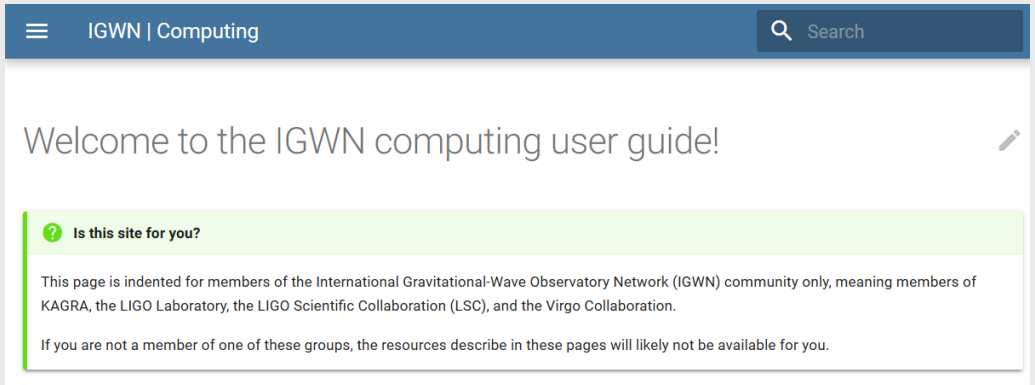

FROM 2019 CHEP PRESENTATION!

SHOPPING LIST

Bulk data transfer

Safely transfer all data to custodial storage in CCs

Software packaging and distribution

Make pipeline software available ubiquitously

Data distribution

Make h(t) data available to worker nodes anywhere

Data cataloguing and bookkeeping

Organize all data and metadata and provide querying capabilities

Workload management

Provide a uniform job submission and runtime environment

High-level workload orchestration

Keep a database of all jobs and allow the enforcement of priorities

Monitoring and accounting

Monitor distributed computing and provide reliable accounting

FROM 2019 CHEP PRESENTATION!

WHERE ARE WE?

Bulk data transfer

Rucio

Software packaging and distribution

CVMFS (+ Conda + Singularity)

Data distribution

CVMFS (but see above) + StashCache

Workload management

HTCondor and, possibly, DoDAS

Mostly decided except implementation details, under deployment or active development

Data cataloguing and bookkeeping

Probably Rucio again

Up next?

High-level workload orchestration

Possibly DIRAC, custom tools if it's an overkill

Monitoring and accounting

TBD, maybe DIRAC again?

Not an immediate priority

VIRGO

THE WORKS

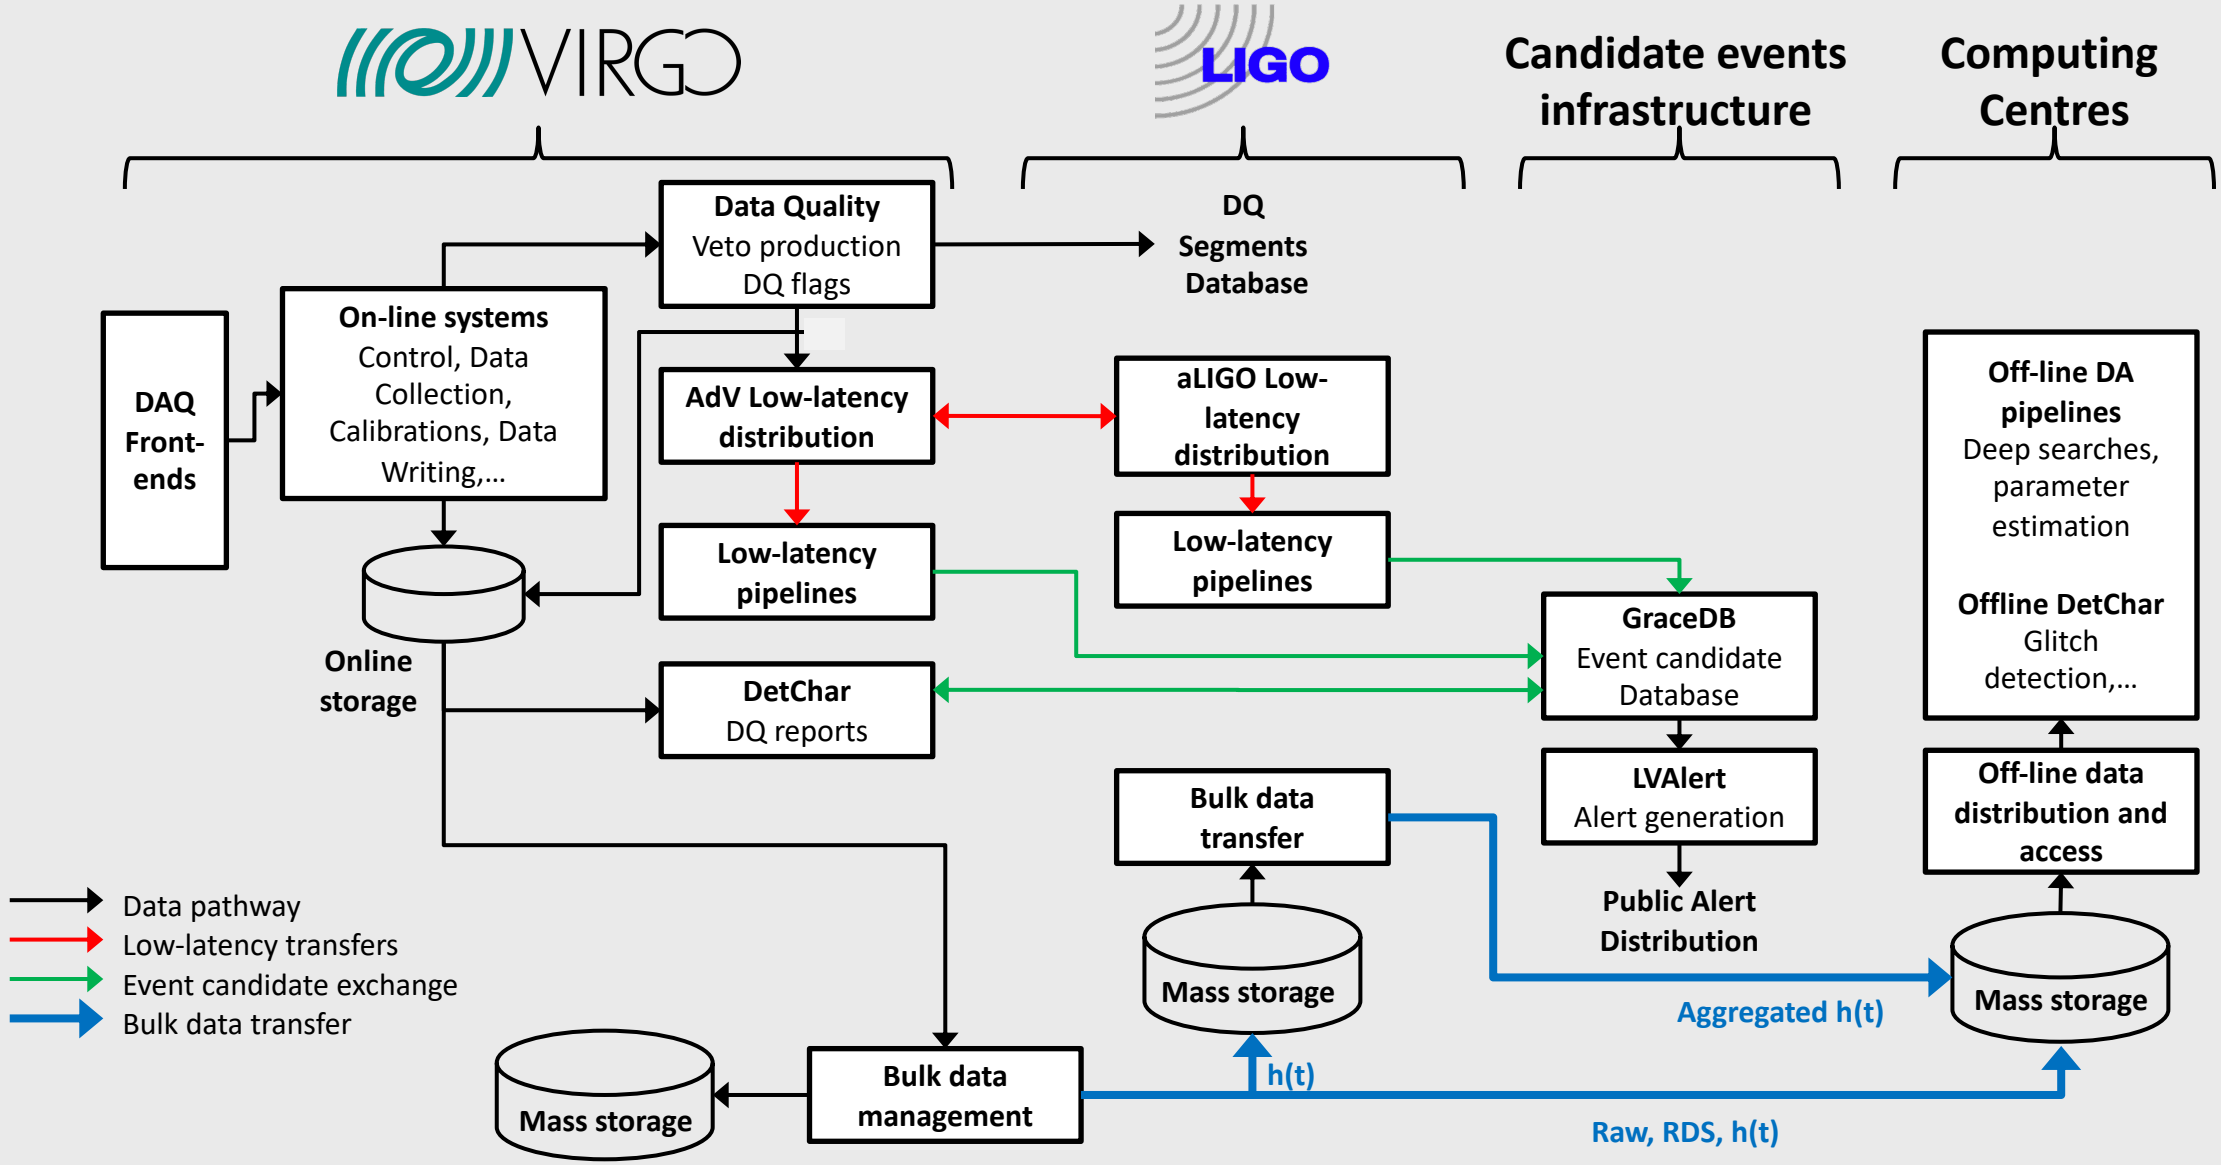

Metadata (namespace)

Actual data bytes

Open Science Data Federation, managed by the OSG

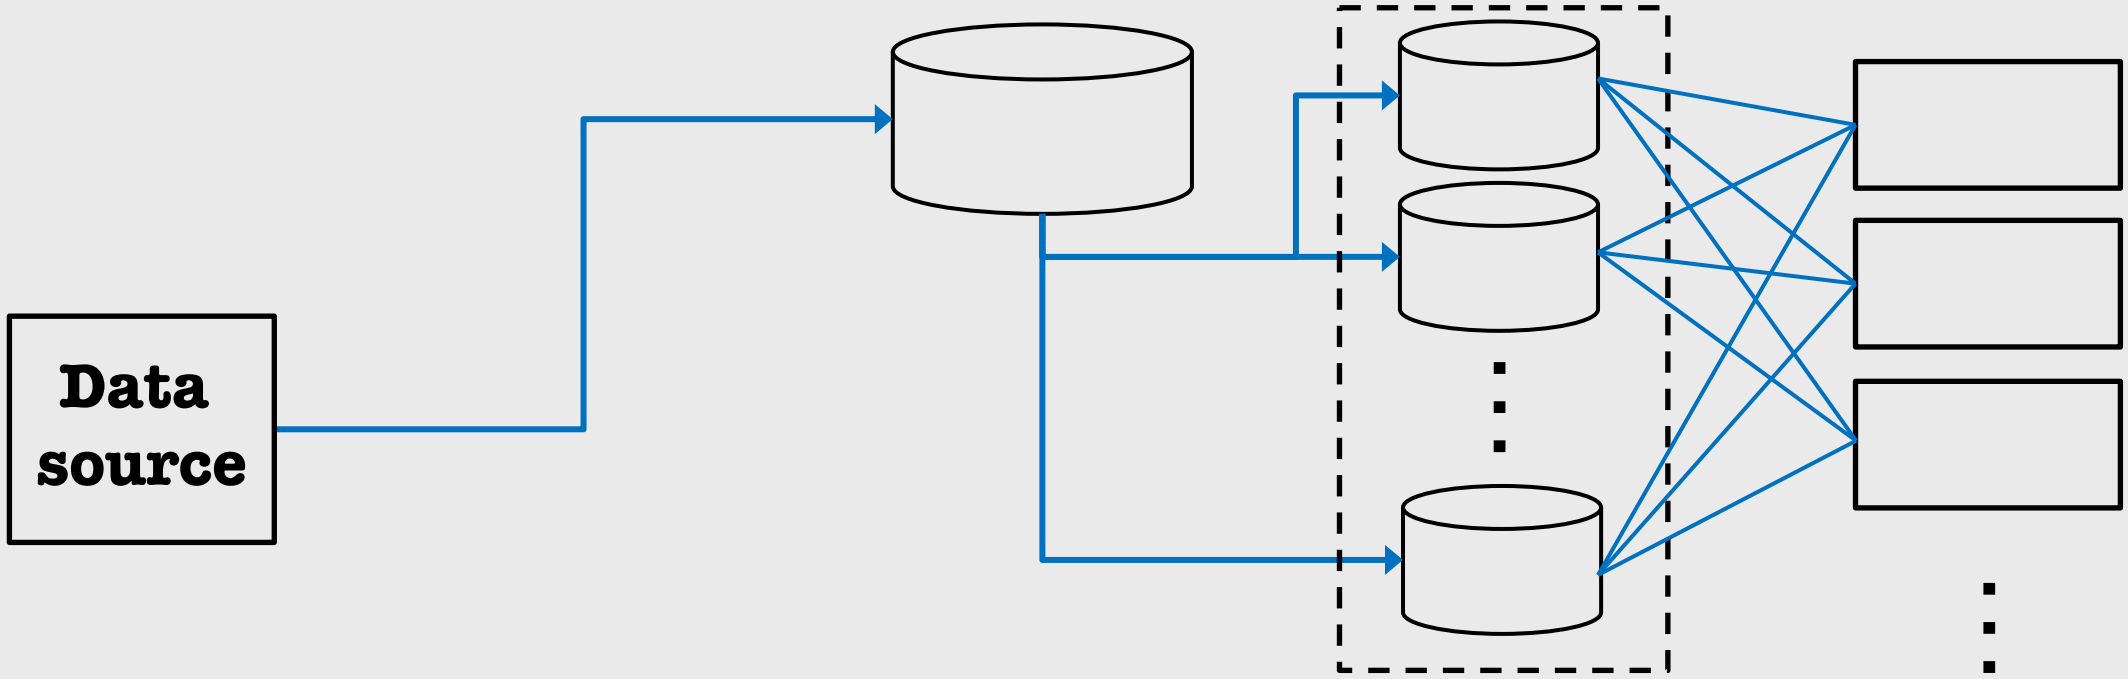

Metadata (namespace)

Actual data bytes

Open Science Data Federation, managed by the OSG

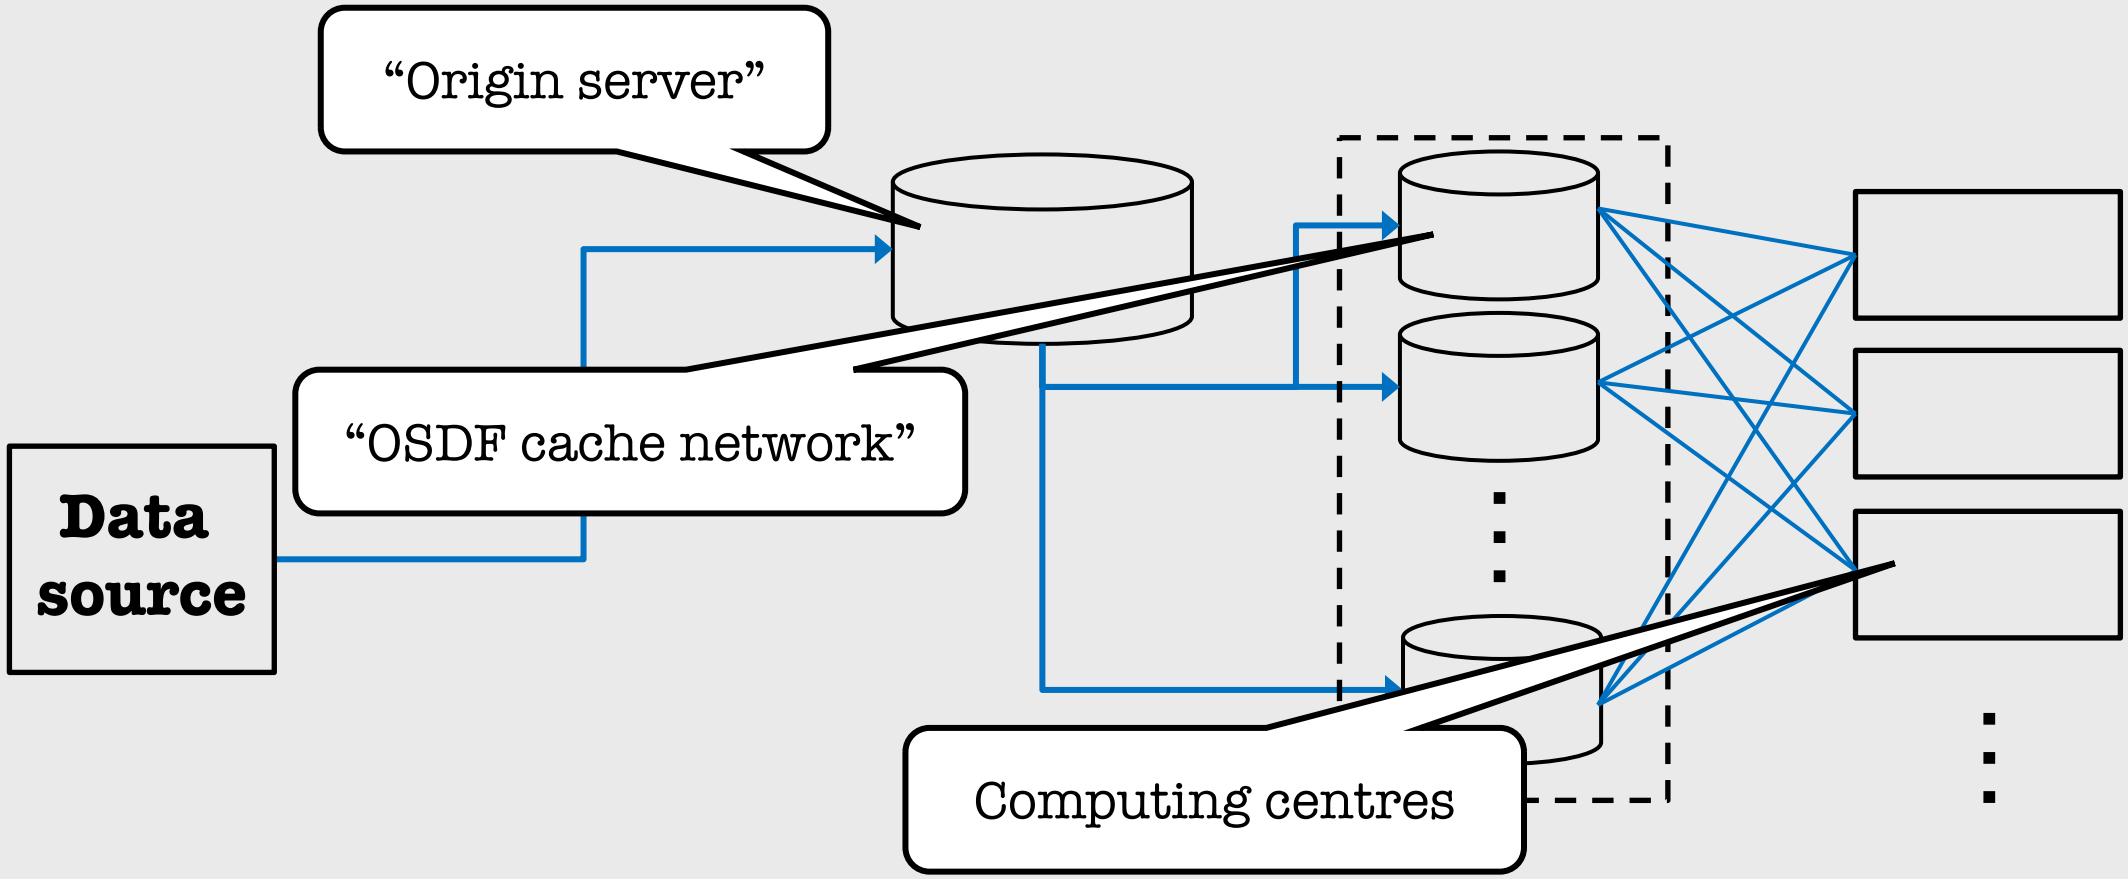

TECHNOLOGY RECAP: OSDF

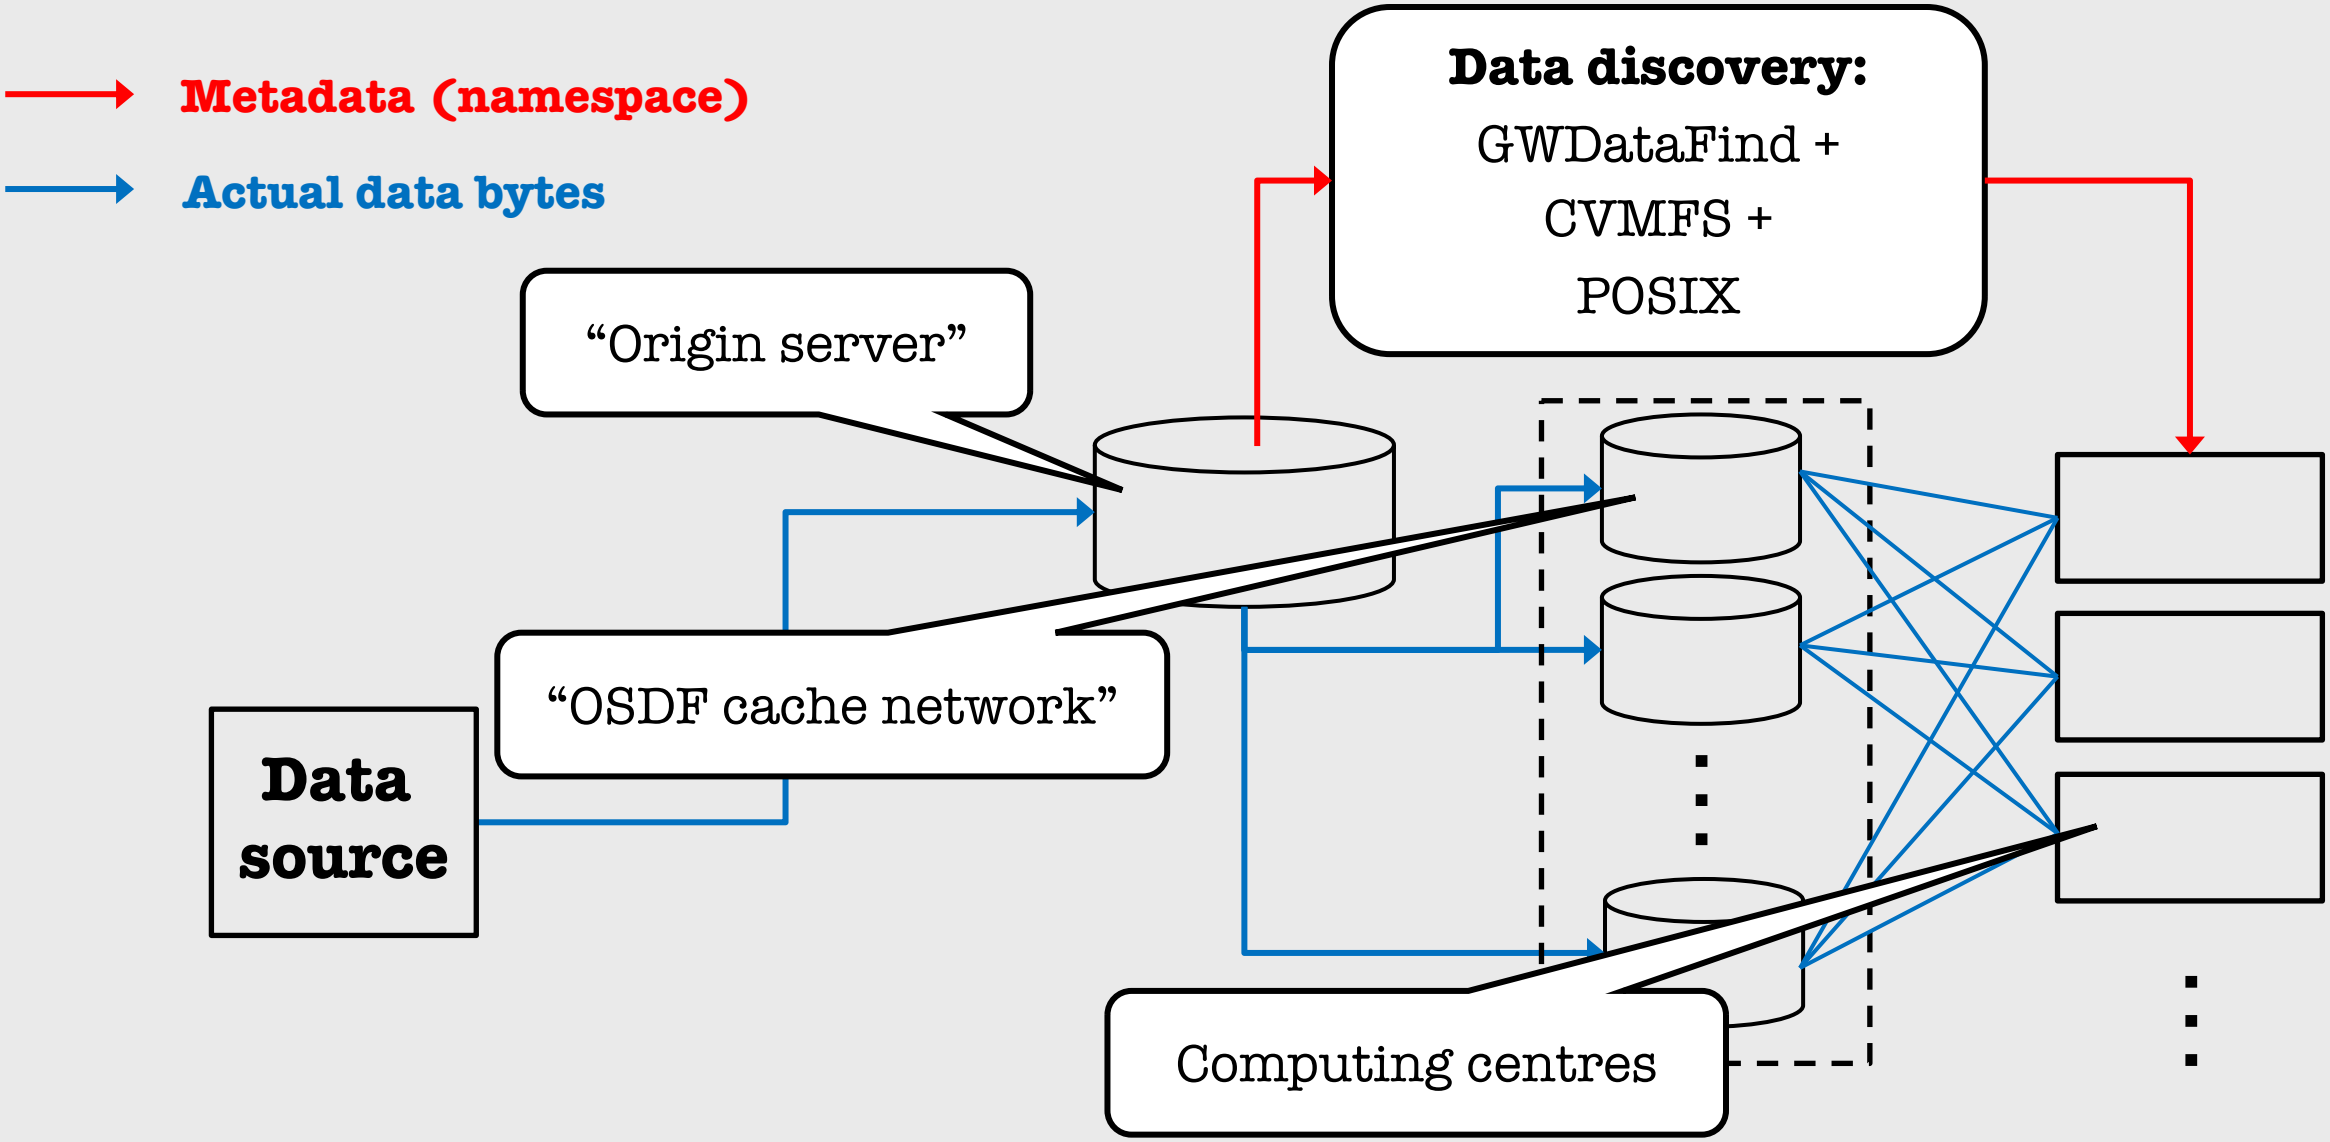

TECHNOLOGY RECAP: OSDF

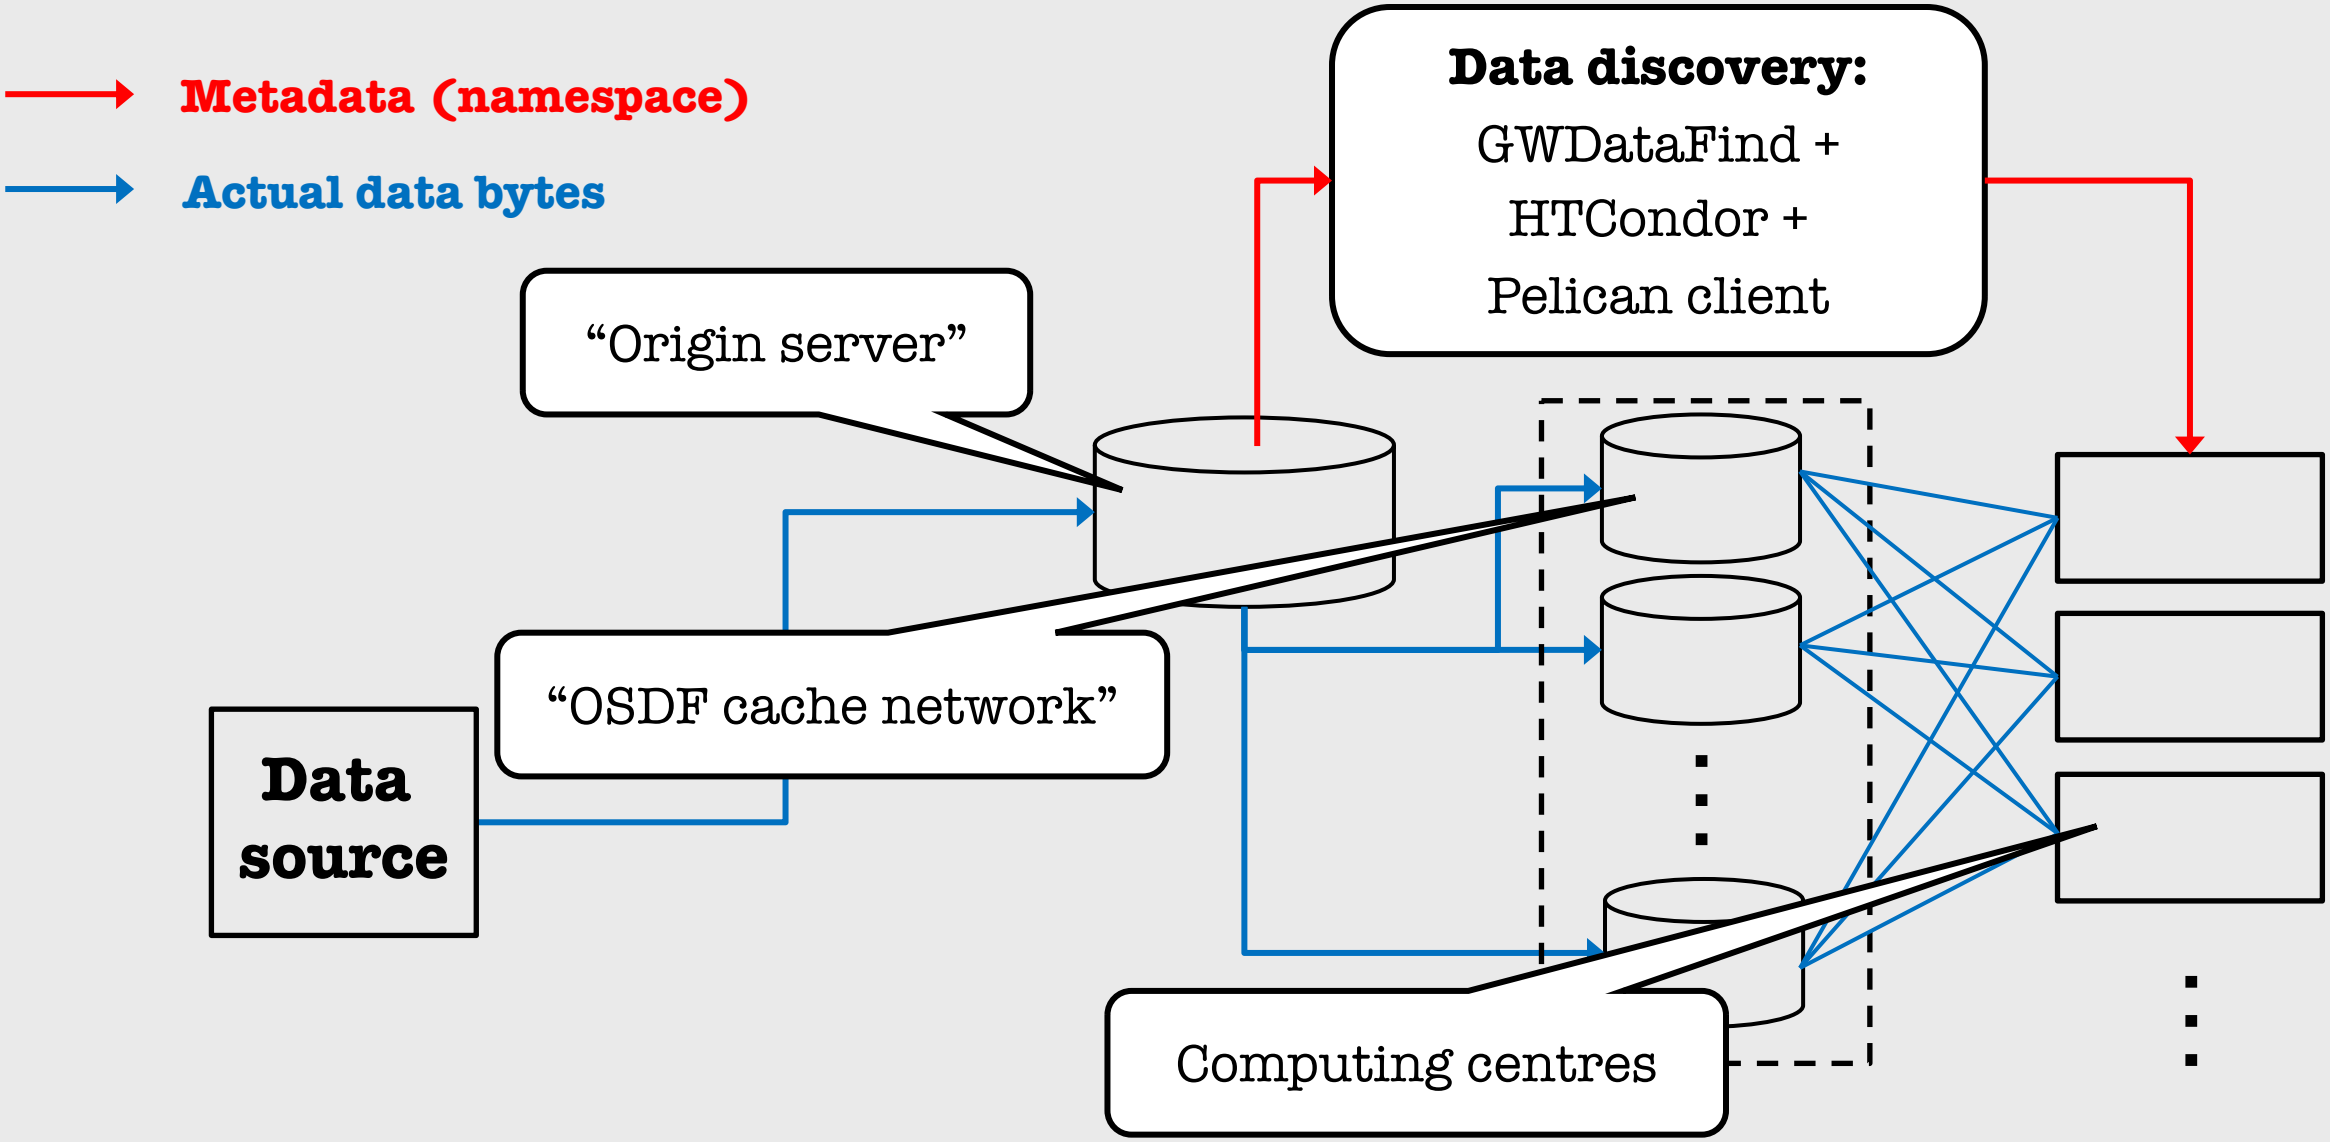

OFFLINE PATHS

Proprietary frame and "derived" data:

- <prefix>/igwn/ligo/…
- <prefix>/igwn/virgo/…
- <prefix>/igwn/kagra/…
- <prefix>/igwn/shared/…

Where <prefix> **is:**

/cvmfs/<ifo>.storage.igwn.org/ for CVMFS access

osdf:// for OSDF access and with transfer_input_files

Public open data:

- /cvmfs/gwosc.osgstorage.org/gwdata

Software:

- /cvmfs/software.igwn.org/ (i.e., CONDA environments)

Reference and details:

- https://computing.docs.ligo.org/guide/htcondor/data/

OFFLINE DATA DISTRIBUTION

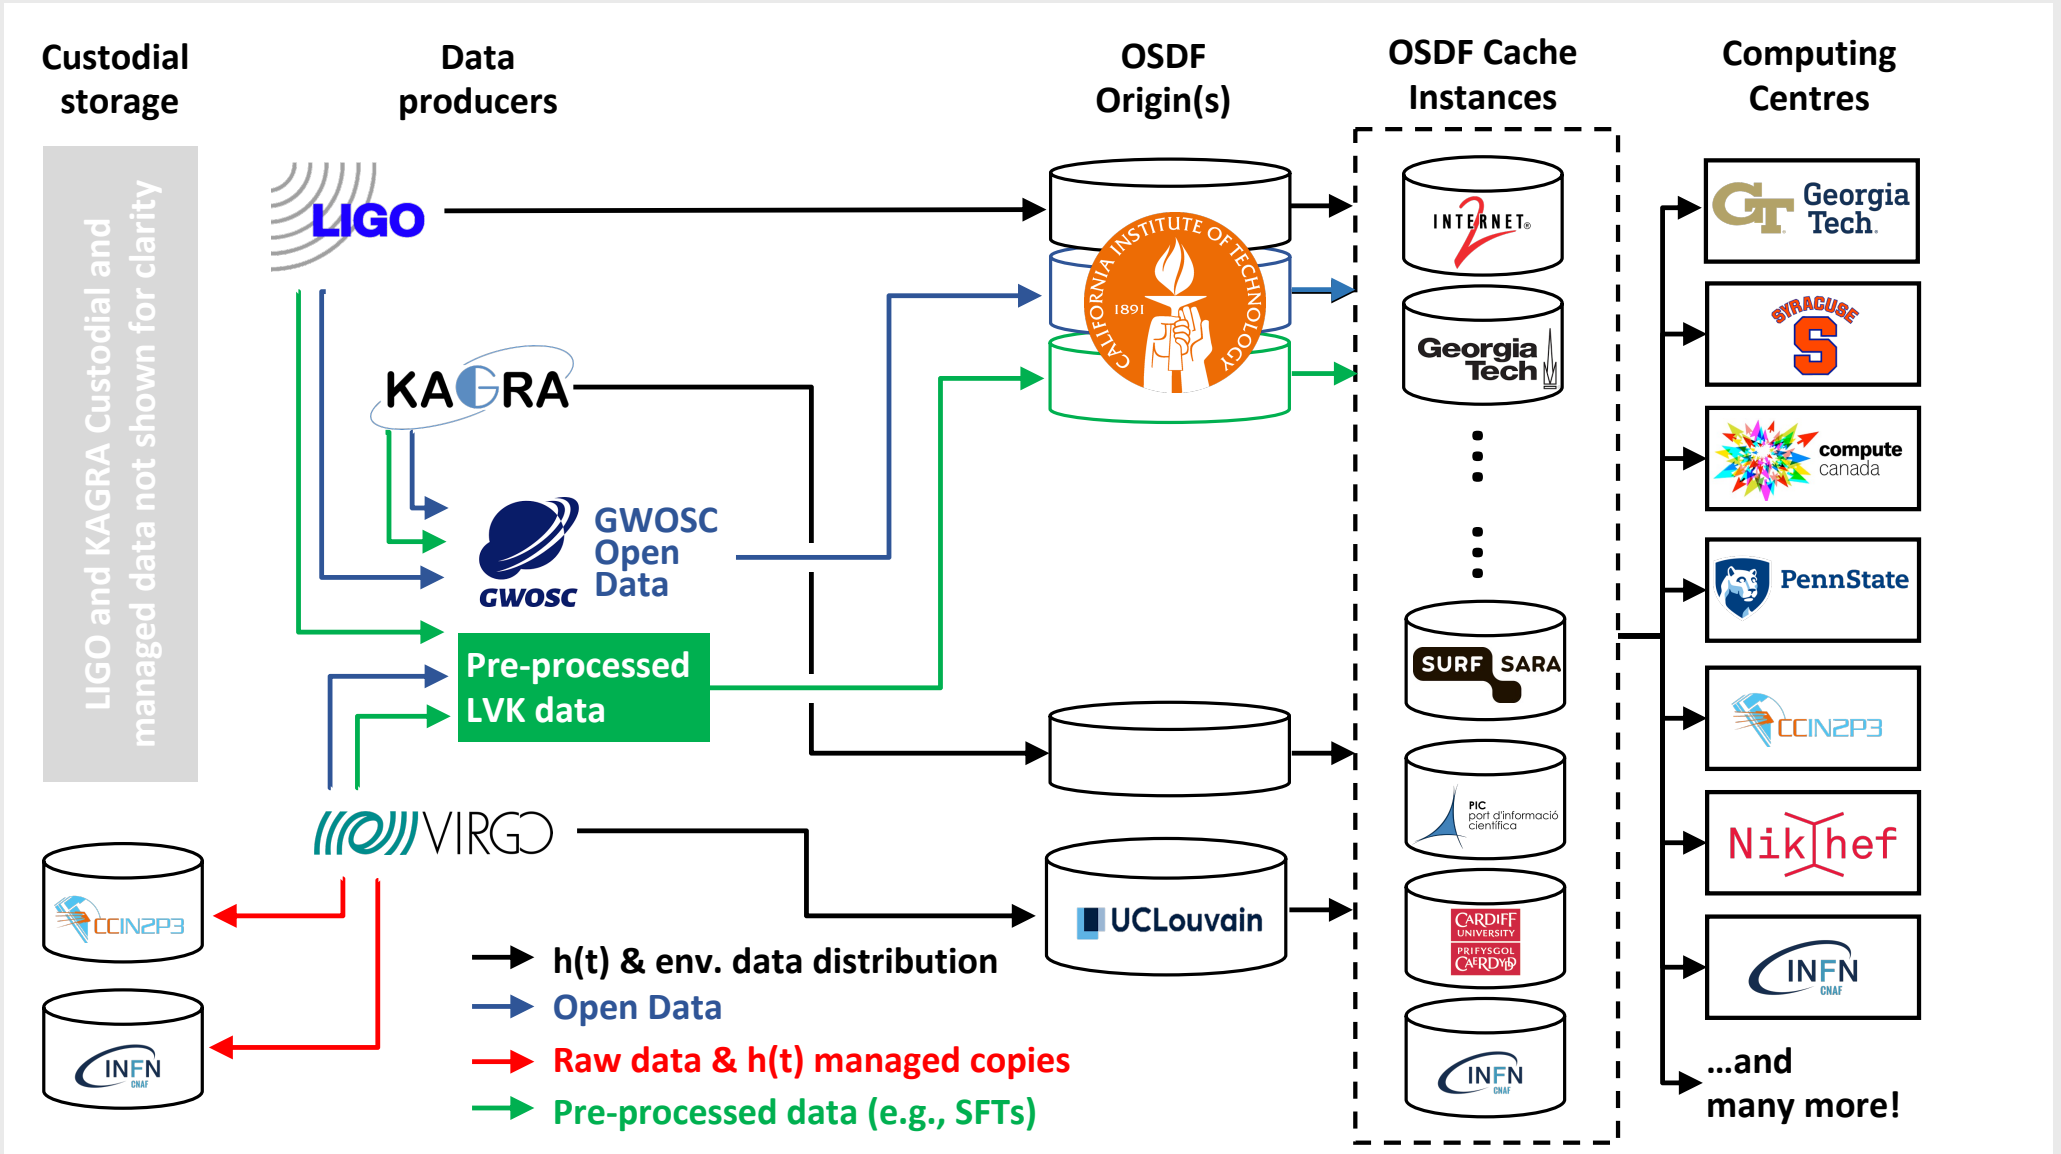

OSDF VS CVMFS (FOR DATA)

- In the original design, CVMFS used to distribute data namespace, OSDF (AKA StashCache) to distribute actual data bytes
- The CVMFS layer works but is fragile
  - E.g., when several jobs access same files at the same time (notably, CW preparatory workloads)
- Direct usage of the OSDF client is always recommended
  - Also allows smooth use of HTCondor file transfer capabilities using `transfer_input_files`
  - At the cost of being more explicit in describing job's needs in the submit file
- GWDataFind client able to provide OSDF paths
  - Use the https://datafind.igwn.org server and include `-u osdf` option

EUROPEAN ALTERNATIVE

- The ESCAPE Data Lake
  - Based on federation of RUCIO Storage Elements
  - INDIGO IAM as AuthN/AuthZ
- RucioFS to replace CVMFS
  - Same UX, different under-the-hood technologies
- ESCAPE/OSCARS proposal submitted for ET
  - See tomorrow's discussion!

OSDF TOPOLOGY

WORKLOAD MANAGEMENT

- HTCondor as a uniform job submission layer
  - To local resources
  - To the IGWN pool through GlideinWMS
  - To other available pools (e.g., OSG)
- No central task queue or job database
  - Priorities enforced “by hand” after a posteriori checking of the accounting
  - Difficult to do better without central job DB because of heterogeneity iof resources…

GlideinWMS in a nutshell

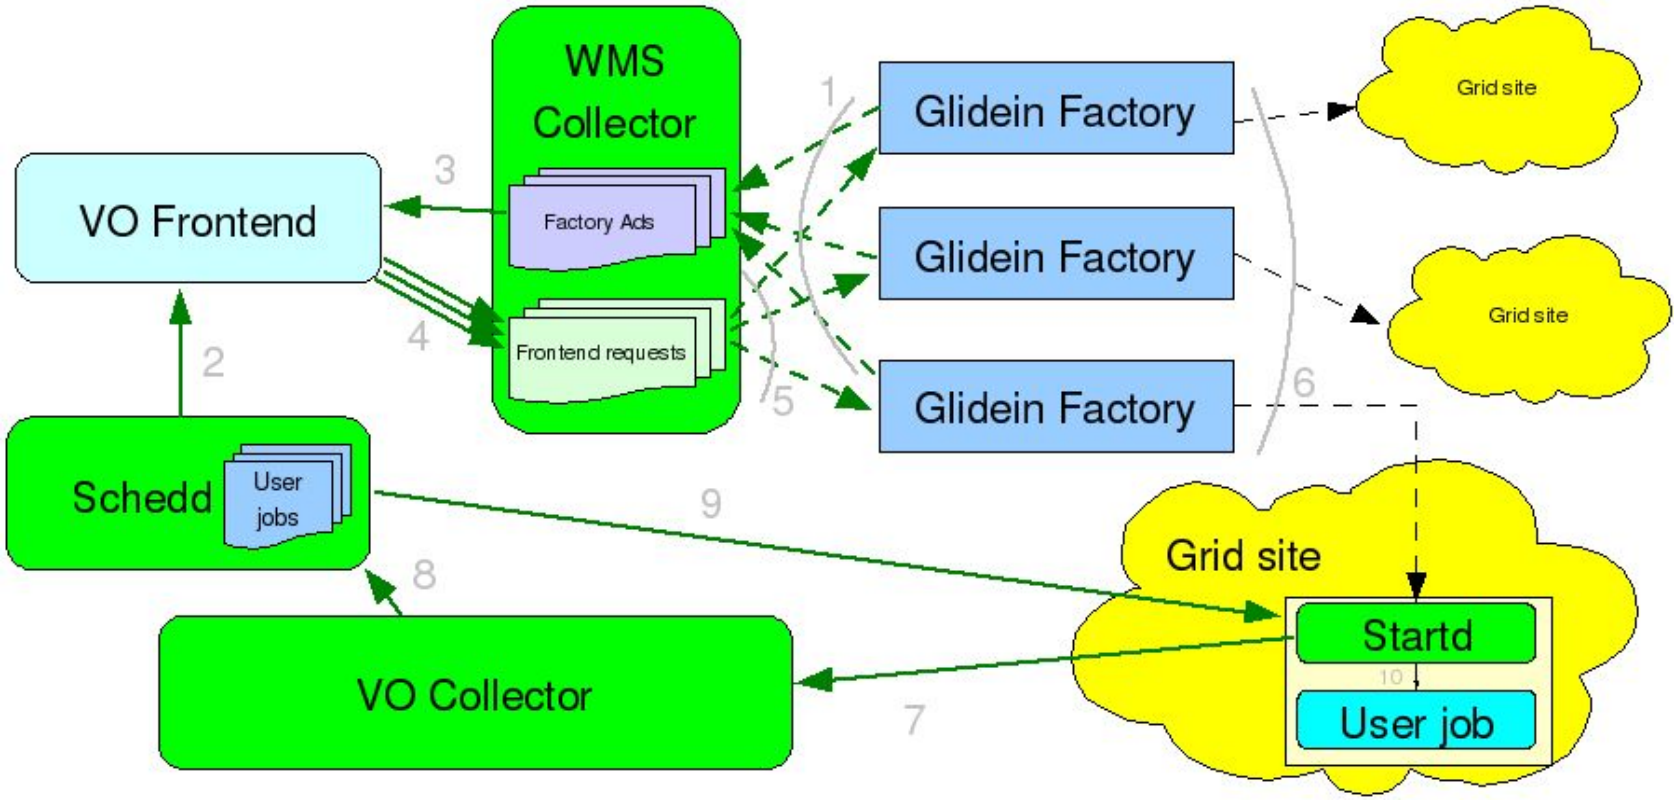

O4B READINESS

- End-user experience will remain largely the same, with differences during O4b
  - Except for a few changes, see Josh's presentation @ LVK [LIGO-G2400618-v1]
- Virgo low-latency, OSDF and bulk data infrastructures tested and exercised
  - Some last minor issues in bulk data transfer being ironed out
- Ongoing issues with CVMFS/OSDF performance at several sites. Some improvements to caching infrastructure in Europe:
  - New caches in Amsterdam and Lyon (upcoming)
  - Cache hierarchy issues being investigated and fixed
  - CVMFS layer for data remains fragile and OSDF should be preferred for this purpose in all cases.
- New grid access point(s) being deployed in Europe
  - General use, but with some bespoke features to cater to "local" groups

OFFLINE PATHS

Proprietary frame and "derived" data:

- <prefix>/igwn/ligo/…
- <prefix>/igwn/virgo/…
- <prefix>/igwn/kagra/…
- <prefix>/igwn/shared/…

Where <prefix> **is:**

/cvmfs/<ifo>.storage.igwn.org/ for CVMFS access

osdf:// for OSDF access and with transfer_input_files

Public open data:

- /cvmfs/gwosc.osgstorage.org/gwdata

Software:

- /cvmfs/software.igwn.org/ (i.e., CONDA environments)

Reference and details:

- https://computing.docs.ligo.org/guide/htcondor/data/

IAM ACTIVITIES

- Work on Virgo IAM integration very slowly but steadily progressing
  - INDIGO IAM as “glue” between VMD and CILogon
  - VIM is the first instance
- Two-factor authentication to be deployed for ligo.org credentials
  - Commercial solution: Cisco DUO
  - Timeline TBD, but will not impact O4b
  - Starting from high-risk but low-impact (for “everyday physicist” logins), apps, e.g., grouper
  - IAM integration and ligo.org -> igwn.org transition will ensure also Virgo and KAGRA will eventually be covered
- KAGRA switch to common IAM (CILogon) platform
  - Confirmed for as soon as possible, but some technical work needs to be done first

MONITORING DASHBOARD

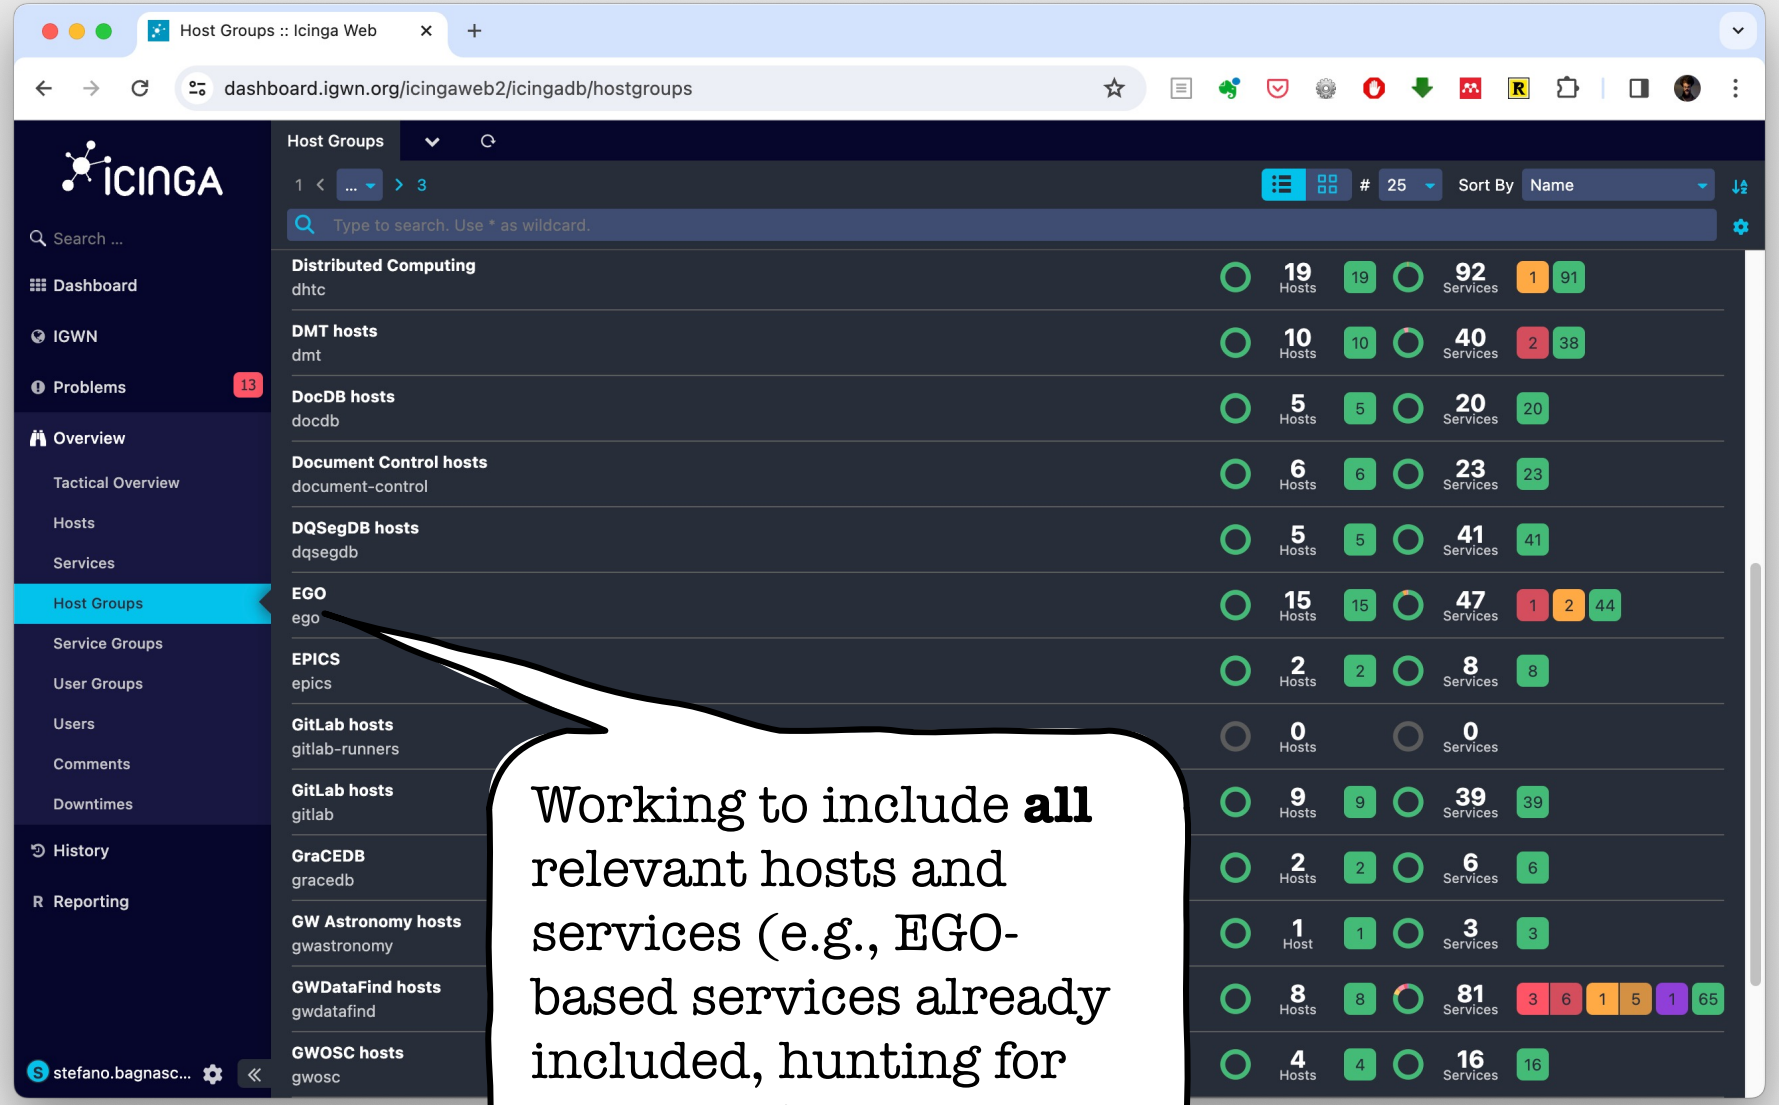

Working to include **all** relevant hosts and services (e.g., EGO-based services already included, hunting for runaways)

ACCOUNTING SYSTEM

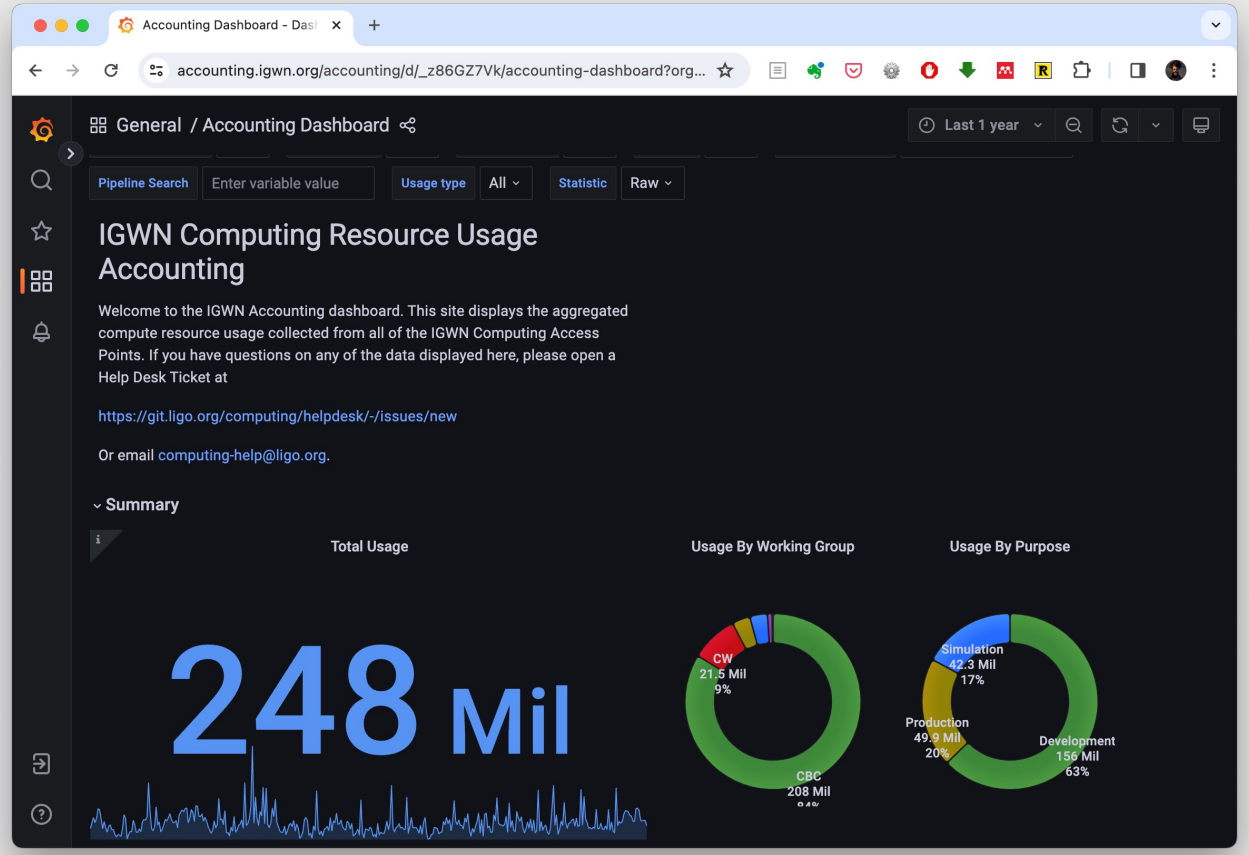

- New dashboard!
  - accounging.igwn.org
  - Collection and database backend not changed
- Aim to account all computing usage
  - Data collection script to be deployed for local usage in non-grid sites (not yet everywhere)
- Still some inconsistencies to be checked
  - E.g., wrt local or EGI accounting

HOW IT WORKS

- Prerequisite: all jobs need to be associated with a tag
  - **`ligo.prod.o3.cbc.imbh.gstlaloffline`**
- "IGWN" jobs: accounted for at the source
  - The HTCondor schedd provides information wherever the job actually ran
- Local jobs: accounting script to be installed
  - Publishes accounting records on web page
  - Database backend collects them daily
  - HTCondor client available
  - Slurm client?

- Open issue: can we enforce mandatory tags in EU sites?

ACCOUNTING.IGWN.ORG

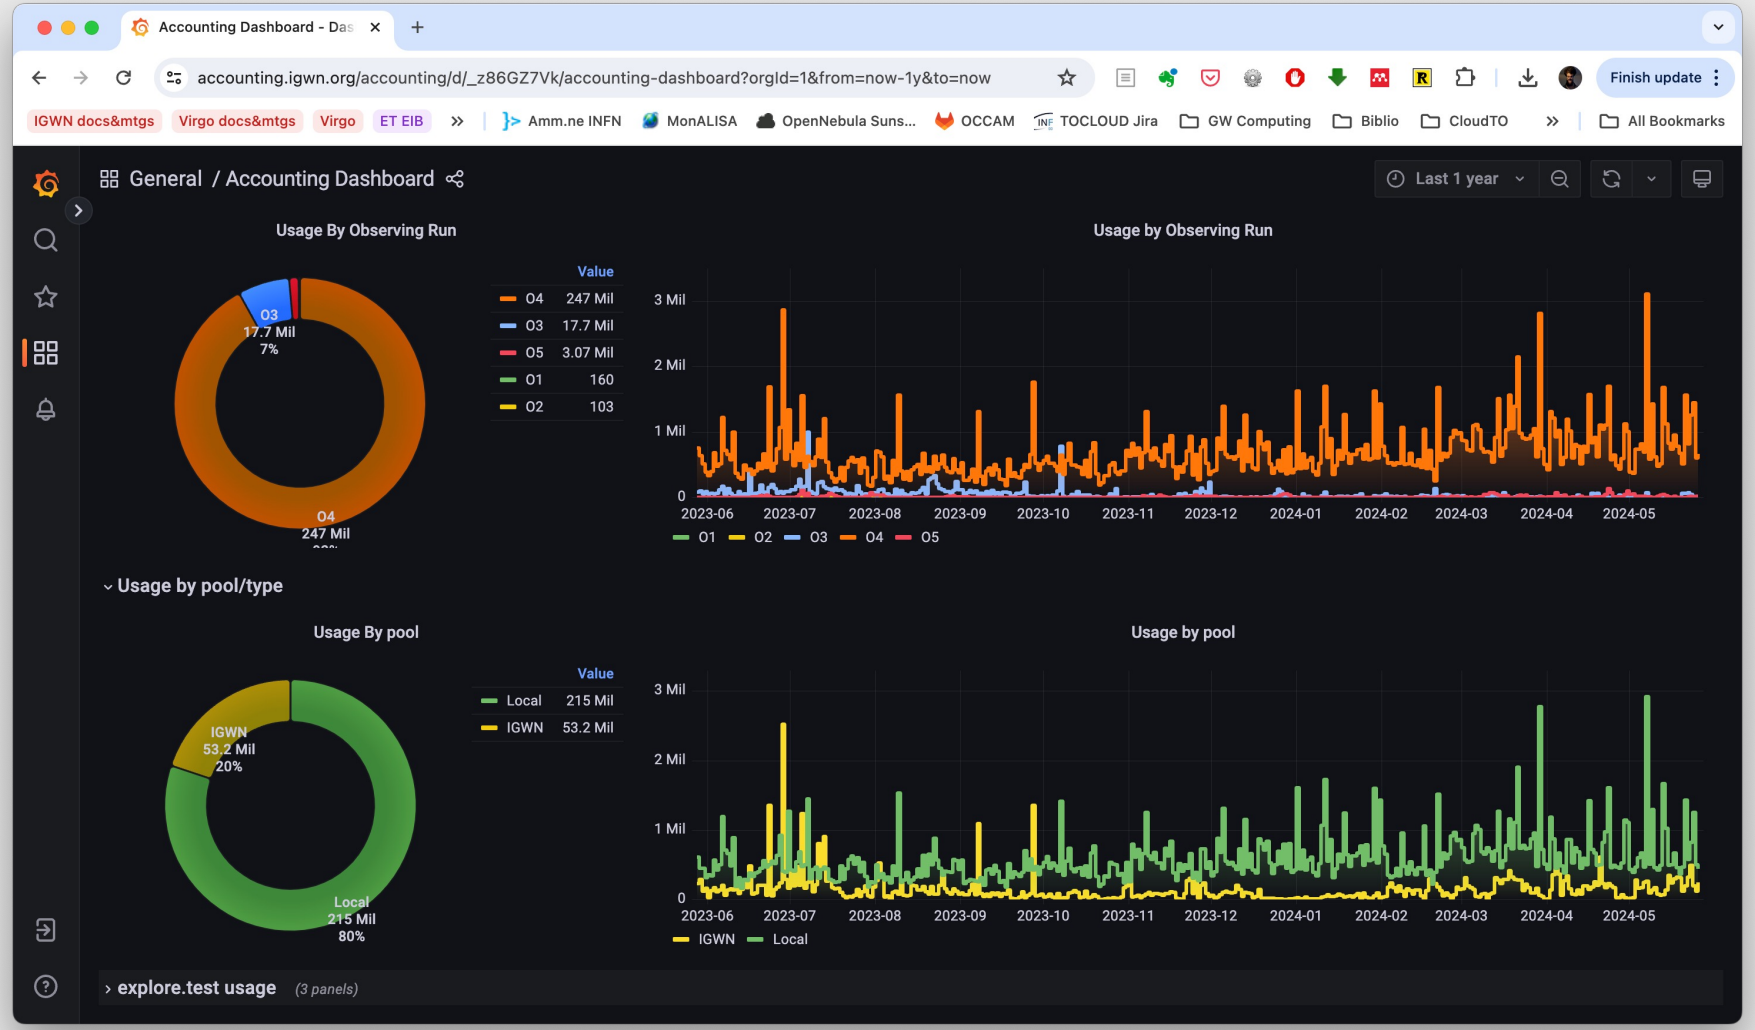

TWO YEARS OF PROCESSING

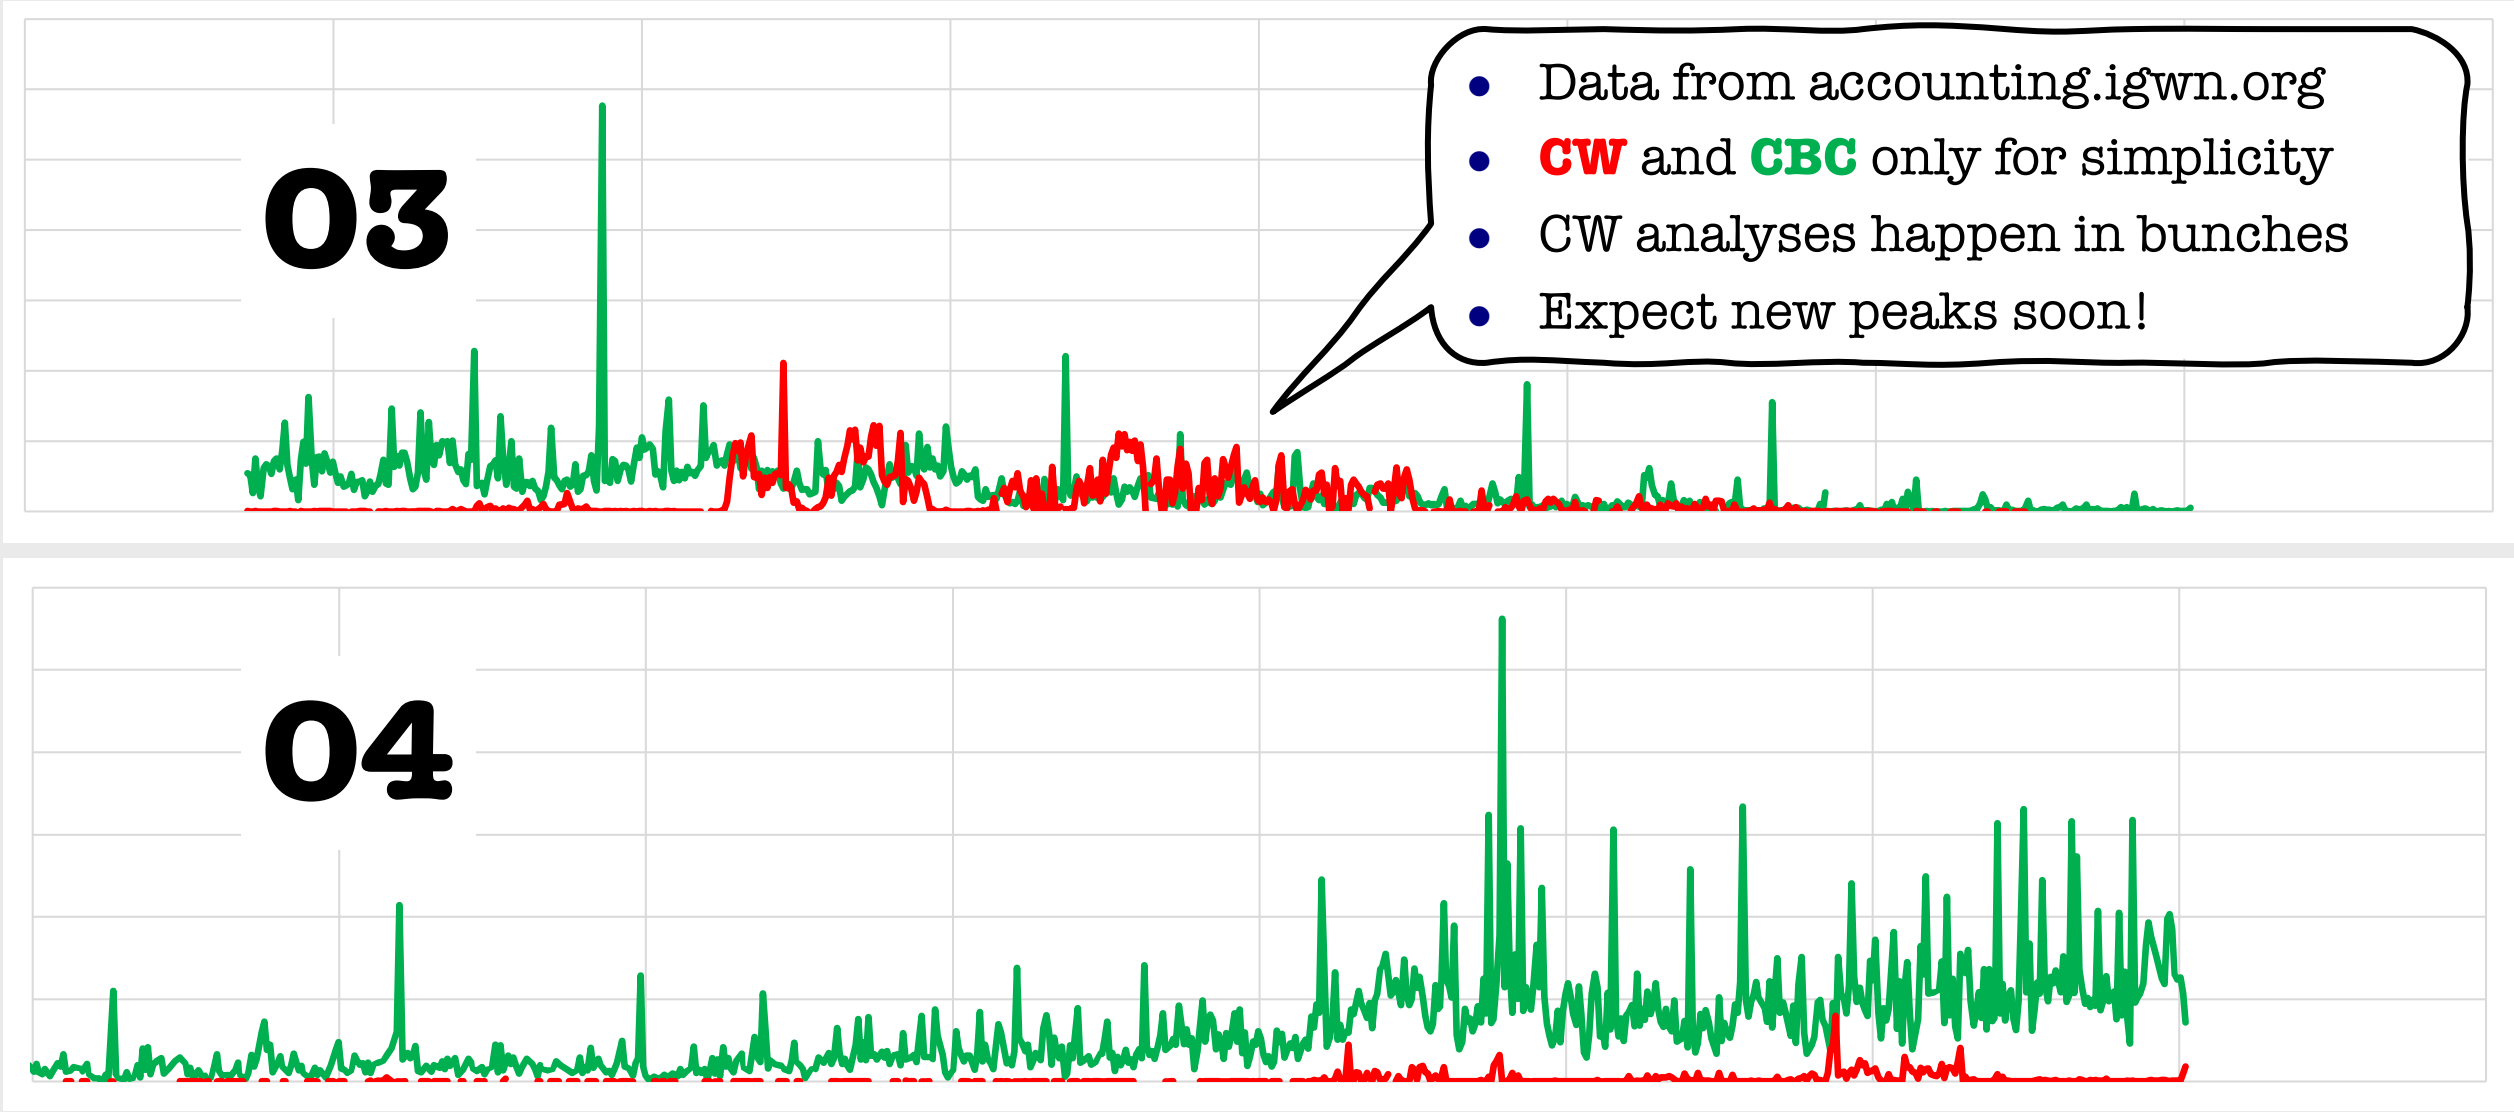

IGWN SUPPORT

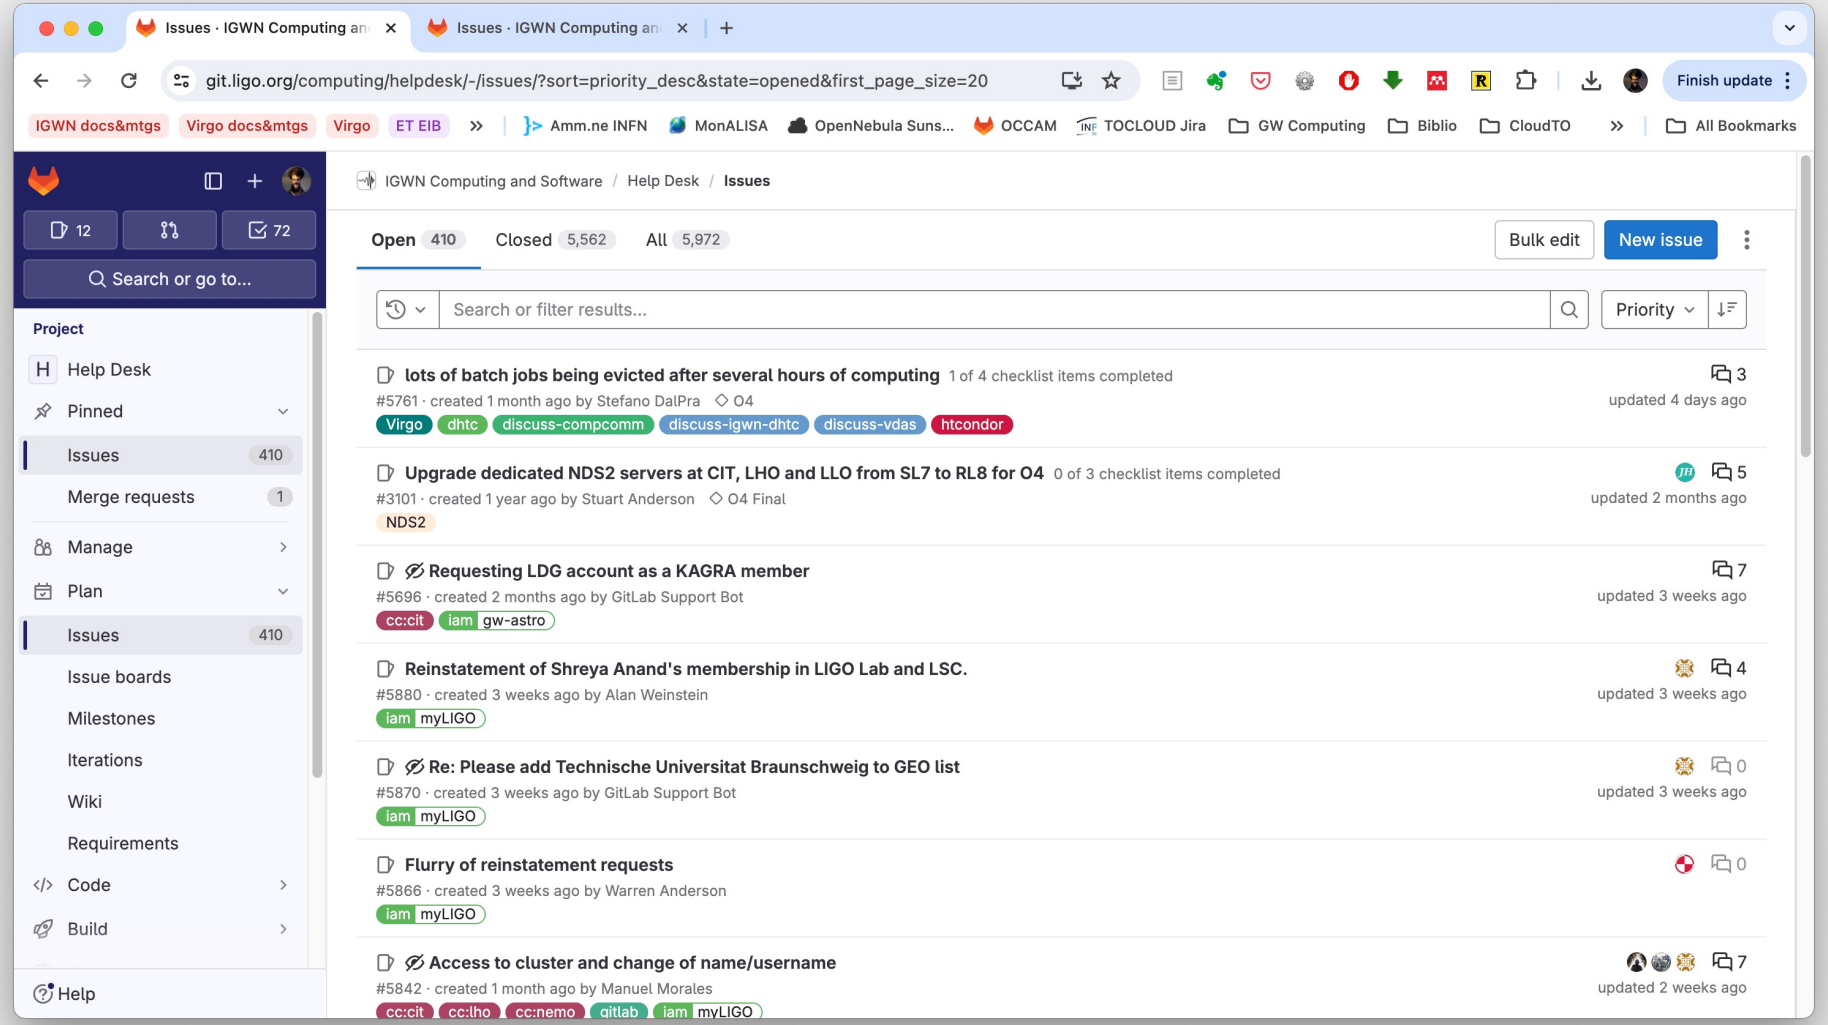

ROADMAP

- Continue integration into IGWN infrastructure
- Continue to ramp up Virgo's contribution
- Move some processing away from EGO
  - LL searches; engineering cluster @ CIEMAT
- Try to onboard some collaboration services
  - We were unsuccessful with GitLab…
  - LL tier @ CNAF
  - …
- Medium term: exploit ET synergies to gain some leverage
  - And try to guide new developments
  - See tomorrow's discussion

COMPUTING REORGANIZATION

- Our computing activities are described in the Virgo Computing Work Breakdown Structure
  - VIR-0019E-19
  - Old and not completely reflecting actual current organization of work
- However, activities in the LVK collaboration are organized through the Operations Whitepaper
  - VIR-0790C-21
- Need to keep the two consistent
  - FTE counting
  - Coordination

VIRGO COMPUTING WBS

- WP3 (On-line System Processes) and WP6 (Low-latency Alerts infrastructure) managed differently and omitted from the schema below

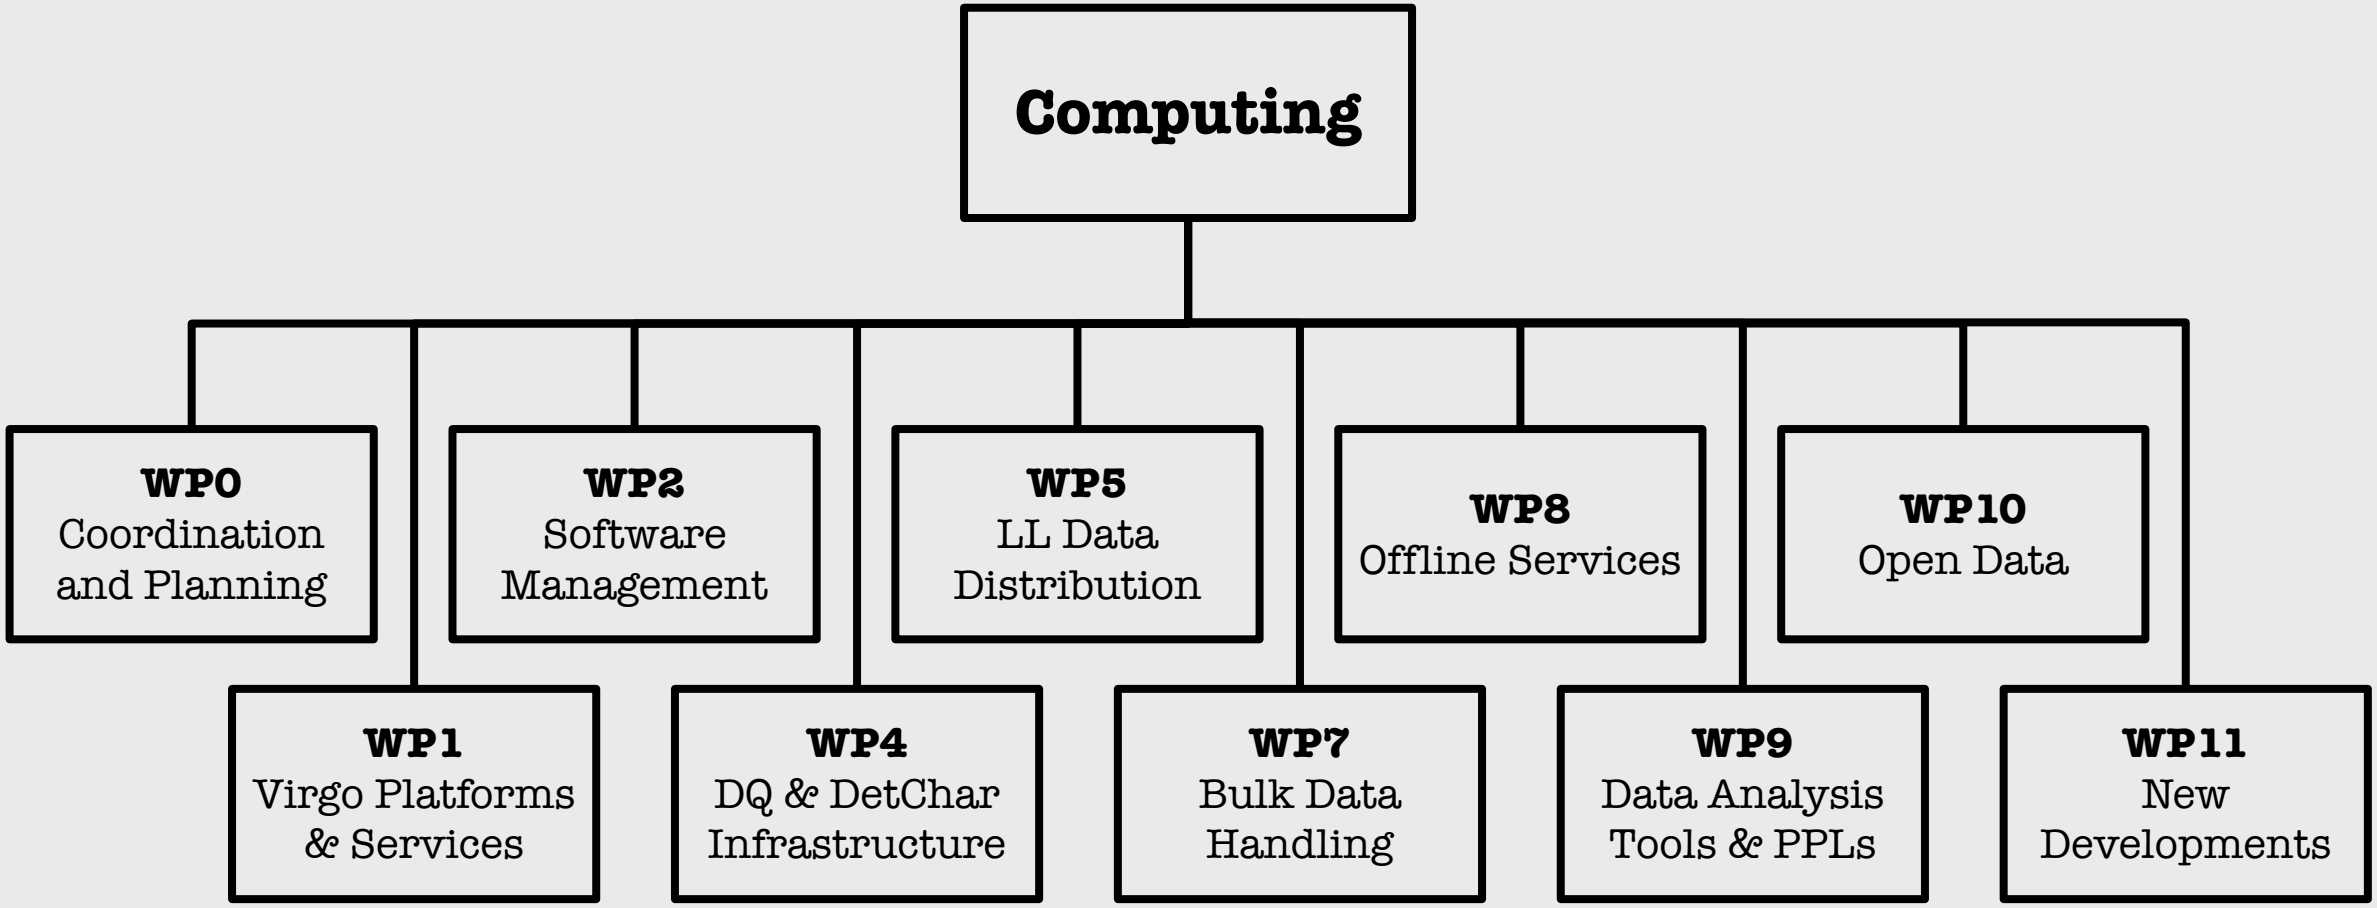

LVK COMPSOFT IN THE OPS WP

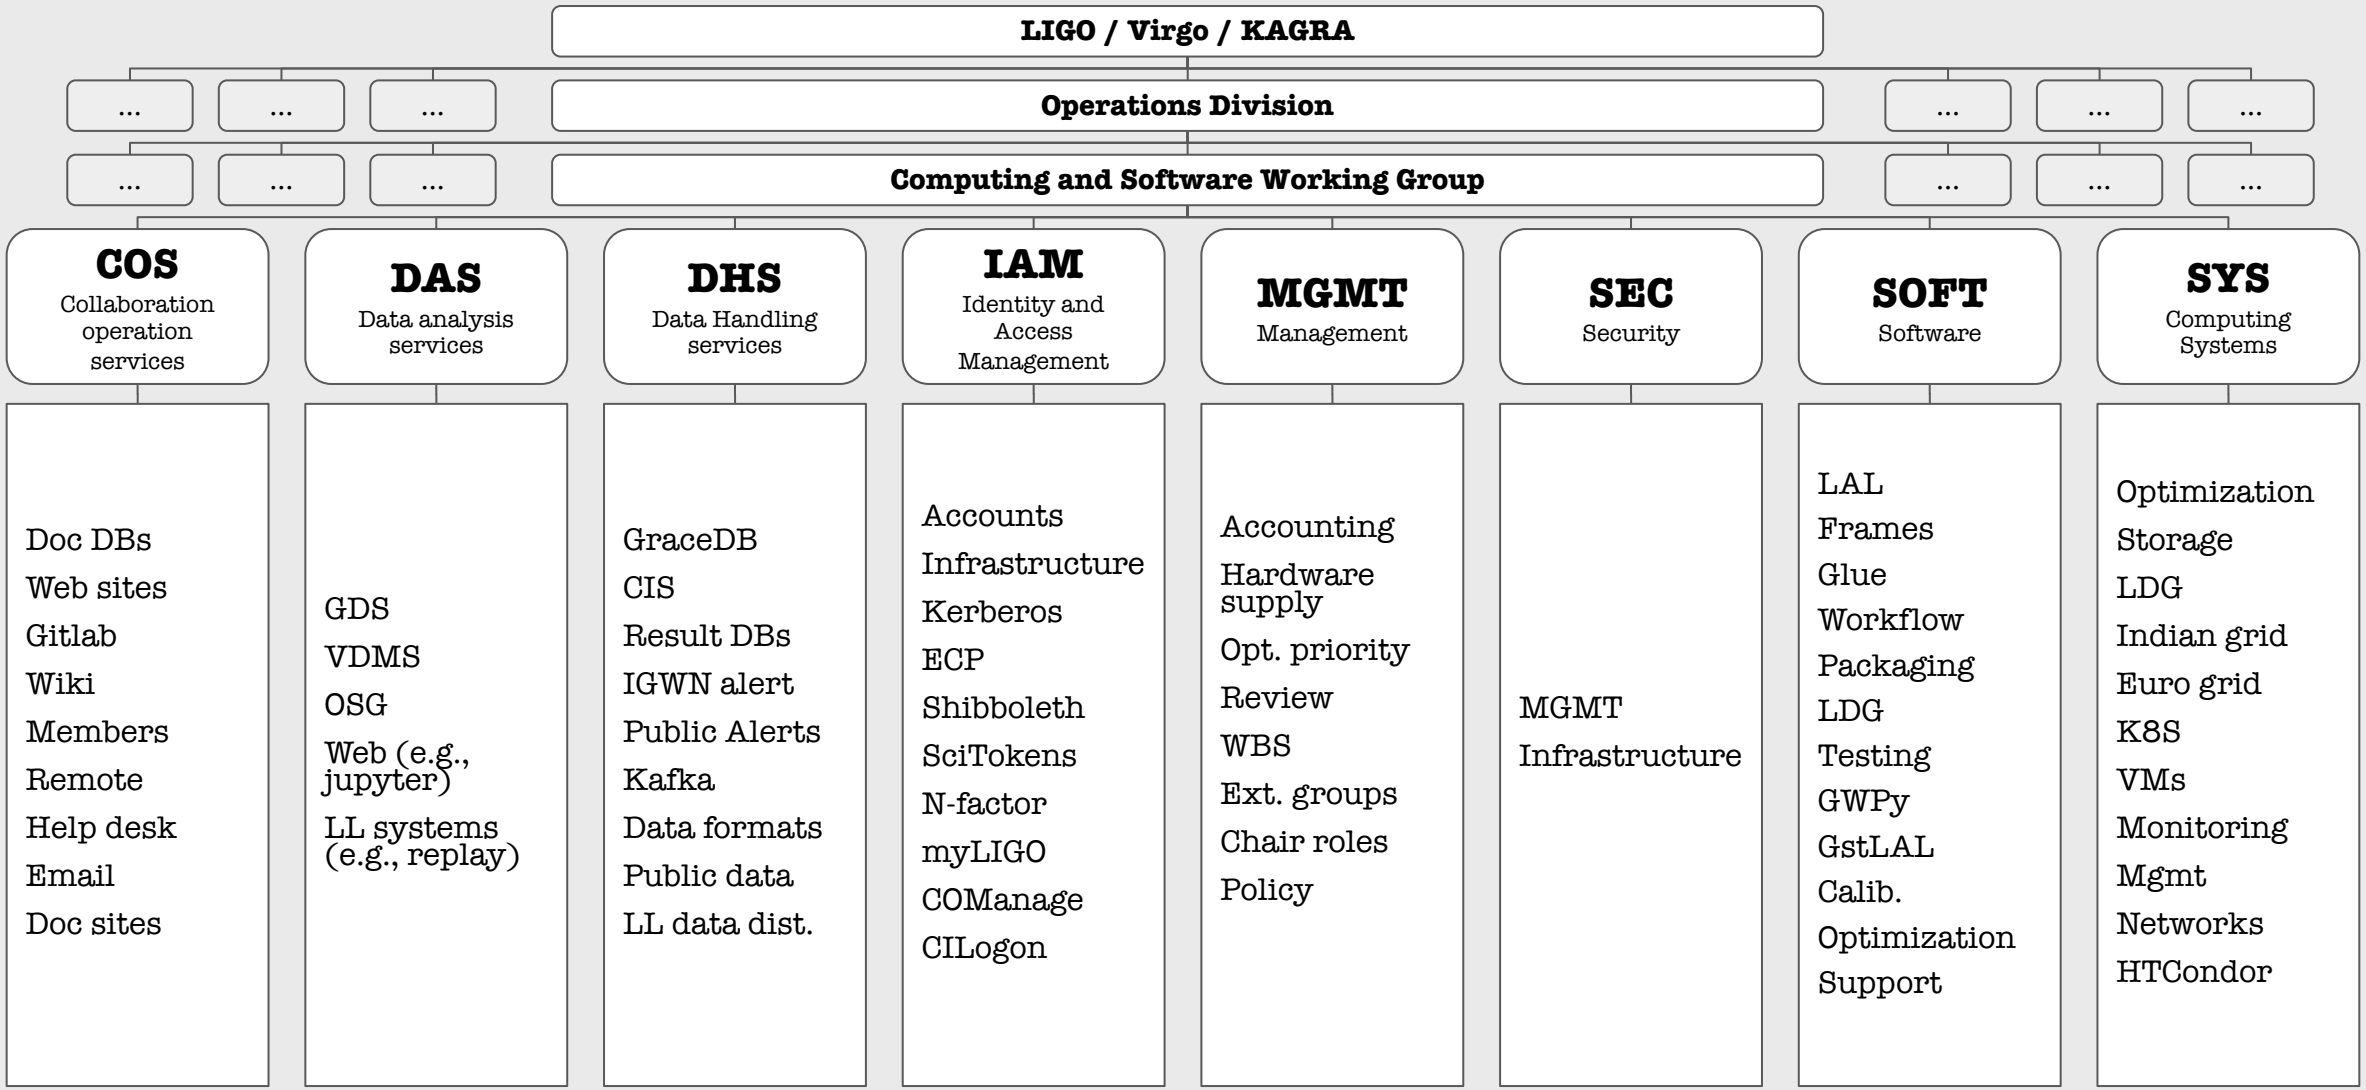

COMPSOFT SUBGROUPS

9 FTE	15 FTE	15 FTE	14 FTE	20 FTE
COS Collaboration operation services	**DASH** Data Analysis & Handling services	**SIAM** Security, Identity and Access Management	**SOFT** Software	**SYS** Computing Systems
Doc DBs Web sites Gitlab Wiki Members Remote Help desk Email Doc sites	GDS VDMS OSG Web (e.g., jupyter) LL systems GraceDB CIS Result DBs IGWN alert Public Alerts Kafka Data formats Public data LL data dist.	Accounts Infrastructure Kerberos ECP Shibboleth SciTokens N-factor myLIGO COManage CILogon	LAL Frames Glue Workflow Packaging LDG Testing GWPy GstLAL Calib. Optimization Support	Optimization Storage LDG Indian grid Euro grid K8S VMs Monitoring Mgmt Networks HTCondor

Of which Virgo needs to provide o(20%)

2-3 cochairs each

PLANS

- Assign the relevant people to the subgroups
- Try to appoint some Virgo cochairs
- For the time being, keep alive the Virgo-wide communication channels
  - Online-offline Thursday calls
  - VDAS mailing list

- ...and we need to appoint a Virgo Cybersecurity Officer!

PERSONPOWER ISSUE

- All the same, personpower issue needs to be addressed
  - Otherwise something may have to go “best-effort” or unsupported
  - Detailed task inventory allows for fine-grained prioritization
- This is worsened by the difficulty in recruiting (and keeping) people with the right skills
  - Scientific Computing infrastructure is not IT and is not pipeline software development; our (few) computing positions are not competitive with industry
  - This is a problem shared by all the broader physics community
  - The choice of “cloud” technologies (e.g. Kubernetes) and mainstream tools (e.g. CVMFS) should shift some of the burden to CC staff
- Under-supported services and missed opportunities for computing improvements will come at a science cost to the Collaborations
  - IGWN CompComm will seek to communicate those risks to management for acknowledgement, mitigation, and/or acceptance via the IGWN Computing Risk Registry