LIGO SCIENTIFIC COLLABORATION
VIRGO COLLABORATION
KAGRA COLLABORATION

Document Type	**LIGO–T2100304-v1** **VIR-0790A-21** **JGW-T2113065**

LSC-Virgo-KAGRA Operations White Paper (Summer 2021 edition)

The LSC Operations Division, Virgo Group Coordinators, KAGRA Group Coordinators

http://www.ligo.org
http://www.virgo-gw.eu
https://gwcenter.icrr.u-tokyo.ac.jp

Processed with LaTeX on 2021/10/21

Contents

1 Overview and Executive Summary

The successful operation of the LIGO, Virgo, and KAGRA detectors and key infrastructure are critical to enabling gravitational-wave astronomy. Gravitational wave detections are made possible by optimized instrument sensitivity and uptime, well-characterized noise, accurately calibrated data, and robust computational infrastructure that supports not only the detectors but also the data analysis pipelines used to identify and characterize gravitational wave signals. These activities are undertaken by the LIGO Laboratory in collaboration with the broader LIGO Scientific Collaboration (LSC), by the Virgo Collaboration in collaboration with EGO, and with the KAGRA Collaboration.

The LSC Operations Division exists to organize and manage all LSC activities related to, and necessary for, detector operation (see Section 5 of LSC bylaws: LIGO-M050172). It comprises the Detector Characterization, Calibration, Low Latency, Computing and Software, and Open Data Working Groups and the Support of Observatories and Run Planning Committees.

The Virgo Collaboration is an active participant, both in co-chairing and membership, in all the aforementioned Working Groups. The Virgo and KAGRA collaborations are represented on the Joint Run Planning Committee, which is charged with the strategic planning of engineering and observing runs. The Support of Observatories committee is unique to the LSC.

The LIGO, Virgo, and KAGRA **Commissioning** teams work toward optimizing the sensitivity and uptime of the global detector network. The LSC, Virgo, and KAGRA **Detector Characterization (DetChar)** groups interface with the detector commissioning teams and work to improve GW signal searches by identifying and mitigating noise sources that limit sensitivity to astrophysical signals. The LSC, Virgo, and KAGRA **Calibration** teams produce fast and reliable calibration of detector data. The LSC, Virgo, and KAGRA **Computing** teams support the joint LSC-Virgo-KAGRA software and computing infrastructure. The LSC, Virgo, and KAGRA **Low Latency** teams provide rapid notification of Gravitational Wave candidate events to the wider observing community. The LSC, Virgo, and KAGRA **Open Data** teams curate the observation data for public release. The LSC **Support of Observatories** committee coordinates LSC contributions to the Observatories including the LSC Fellows program.

This *LSC-Virgo-KAGRA Operations White Paper* describes the planned activities for these efforts, in the context of the LIGO Scientific Collaboration 2021 Program, M2100100.

Further details on the planned activities in collaboration between the DetChar and Calibration groups and the four LSC-Virgo-KAGRA data analysis working groups (Burst, Compact Binary Coalescence, Continuous Waves, and Stochastic Gravitational-Wave Background) can be found in the *LSC-Virgo-KAGRA Observational Science White Paper*, LIGO-T2100289. The data analysis working groups also undertake tasks related to the Operations of the LIGO-Virgo-KAGRA observatories, described therein.

This Operations White Paper also complements the *LSC Instrument Science White Paper*, LIGO-T2100298, which covers the Advanced Interferometer Configurations Working Group (AIC) including Newtonian Noise and Interferometer Simulations, the Quantum Noise Working Group (QNWG), the Lasers and Auxiliary Optics Working Group (LAWG), the Optics Working Group (OWG), the Suspensions and Seismic Isolation Working Group (SWG), and the Control Systems Working Group (CSWG).

The LIGO Laboratory operates and maintains the LIGO Hanford and Livingston observatories through a Cooperative Agreement with the US National Science Foundation. The LIGO Laboratory makes major contributions to the LIGO Scientific Collaboration, including responsibility for delivering calibrated, well-characterized gravitational strain data at a target sensitivity during designated observing runs. The broader LSC, jointly with the Virgo Collaboration, is in turn charged with producing astrophysical results, including low latency GW candidate alerts. The **LIGO Laboratory Operations** include activities which are deeply collaborative with the LSC (e.g., Calibration, Detector Characterization, Low-Latency public alerts), those with some participation from the LSC (e.g., Commissioning), and others which are largely internal to the LIGO Lab (e.g., maintenance of the vacuum equipment).

EGO operates and maintains the Virgo detector, whose various systems are developed, installed, steered and upgraded jointly by the Virgo Collaboration and EGO. The instrument is funded by CNRS (France), INFN (Italy) and NIKHEF (The Netherlands), through the the EGO Consortium.

1.1 Observatory Operations

The COVID-19 pandemic affected operations at all three observatory sites. The O3b run was suspended on March 27, 2020 in order to comply with local regulations and to ensure the health and safety of the staff.

1.1.1 LIGO Observatories

Following the decision to end the O3 run, the LIGO observatories have engaged in a phased return to normal operations. Both sites returned to limited operations, with strict safety protocols, in July of 2020. After recovering interferometer performance work began on upgrades for O4.

In the year ahead, the LIGO Laboratory science and engineering teams will continue with plans for detector improvements that have been identified as targets for sensitivity improvement. The most ambitious of these is the installation of a 300 m filter cavity to facilitate frequency-dependent squeezing.

Partly due to the impact of the pandemic, the earliest projected date for the start of the O4 observing run is now August 2022.

1.1.2 Virgo

Since the end of O3, the activity at the EGO site has focused on the installation of the detector improvements (the Advanced Virgo Plus Project Phase I) in preparation for the fourth observing run (O4).

The main upgrades have concerned the installation of the signal recycling mirror, of a higher-power laser and of the frequency-dependent squeezing, which required also the realization of a 300 m filter cavity. Other upgrades included an array of seismic sensors to allow the implementation of a Newtonian noise subtraction system, and an instrumented baffle on the Input Mode Cleaner end mirror, for stray light control.

Despite the COVID-19 pandemic, the works at the site have progressed quickly and the installations were concluded by May 2021. Meanwhile, the activities to lock the interferometer are in progress since January 2021 and are planned to end in summer 2021. A noise hunting phase will then start to reach the target sensitivity for O4.

1.1.3 KAGRA

After the end of O3, KAGRA Observatory has been focusing mainly on upgrading the equipment for O4. The impact of COVID-19 has disturbed the scheduling process for O4, so the exact details have not been finalized as of yet. However, by the time O4 starts, there are plans to upgrade the seismic isolation system, suppress stray light, install high power laser, solve the problem of frosting on mirrors and windows in cryostats, and update the DGS software, as well as upgrade the systems of Calibration, Detector Characterization, and Computing and Software groups. In parallel with those upgrades, commissioning has continued on parts of the system as they become available. Before starting O4, all the interferometer configuration will be fixed and full commissioning will be performed for the fine-tuning.

1.2 Detector Characterization

Sensitivity to gravitational-wave signals is limited by noise from the instruments and their environment. Robust detection of signals, the vetting of candidate signals, and the accuracy of parameter estimation are *crucially* dependent on the quality of the data searched and the collaboration's knowledge of the instruments and their environment. The LIGO, Virgo and KAGRA Detector Characterization groups are focused on working together with the astrophysical search groups and the detector groups to (i) deliver the data quality information necessary to avoid bad data, veto false positives, and allow candidate follow-up for gravitational-wave searches and (ii) characterize the Advanced gravitational-wave detectors to help to identify data quality issues that can be addressed in the instruments to improve future instrument and search performance. This focus leads to three core activities: 1) preparing for future observing runs, 2) supporting the upgrade and commissioning of the detectors during commissioning breaks and 3) delivering data quality information and vetting candidate gravitational-wave events during and after observing runs.

At the time of writing, O3 analyses are mostly done, with a few leftover candidates to vet still coming in, such as a few requests from the paper writing teams. Experience gained during O3 and the associated analyses is directly fed into the preparation of O4 and long-term activities.

In preparing for O4 [1], the highest priorities are improving the incorporation of data quality information into the low latency searches, further automating candidate event validation, conducting post-O3 analysis of noise sources that most affected the astrophysical analyses, and developing key tools to improve the performance of the detectors and the astrophysical searches. The highest priorities in preparing for future observing runs are improved automation of existing key tools and monitors of known data quality features, as well as maintaining and extending the software infrastructure required to provide needed data quality information to online and offline searches. Another high priority is curating data quality information for public data releases.

In support of the O3–O4 commissioning break, the highest priority is conducting on-site and off-site investigations of interferometer and environment behavior to support the upgrade and commissioning efforts. This includes the development of critical tools for commissioning and an early characterization of each sub-system, including a coherent system of monitoring web pages.

During O4, the highest priorities are to provide data quality information to the gravitational-wave searches to reduce the impact of artifacts in the data, validate the quality of the data around the time of candidate events, investigate instrumental noise which most impacts astrophysical searches, maintain key tools which allow us to monitor and investigate detector performance, and provide support to the sites.

In parallel, there are a number of key longer-term research and development tasks, including the development of improved methods to uncover the coupling paths and the sources of the noise transients which most impact the searches, and to implement automated noise characterization tools.

To accomplish these priorities, the DetChar groups require enough personpower from the Collaborations, including code developers to support and build key software infrastructure.

More details can be found in Section 3. The present document updates the 2020 Ops. Div. white paper [2].

1.3 Calibration

The mission of the LIGO, Virgo, and KAGRA calibration groups is to provide reliable and timely calibrated strain data for all detectors as well as quantified uncertainty estimates.

The raw optical power variations at the gravitational-wave readout ports of the detectors are calibrated into dimensionless strain on the detectors before use with astrophysical analyses [3, 4]. The process for doing so requires detailed modeling of the feedback control system and the interferometric, opto-mechanical response of the detectors [5]. Some model parameters vary slowly with time throughout operation of the interferometer which must be monitored and, when possible, corrected [6]. All detectors continue to use photon calibrator systems [7, 8, 9, 10] for their primary absolute displacement fiducial reference to develop each static detector model, measure parametric time dependence, and validate any strain data stream once constructed. Virgo plans to employ a secondary reference of gravitational (Newtonian) systems to enhance and compliment the primary systems [9, 11, 12].

Strain is produced in near real-time (low-latency) using methods such as those described in [3, 4]. For the purposes of this white paper, the following language will be adopted: The calibrated strain data contains a quantifiable systematic error due to imperfections in the calibration models. This systematic error is quantified to a certain of level of confidence and can therefore be quoted with, for example, a 68% confidence interval that represents a probabilistic range of the true systematic error. The combination of the systematic error and its corresponding confidence interval will be called the “total calibration uncertainty." Historically, the level of systematic error in the strain data was occasionally, directly confirmed across the detectors' sensitive frequency region, and continuously monitored at select frequencies with the photon calibrator systems. In doing so, the systematic error was confirmed to be within acceptable limits at snapshots in time, but the total calibration uncertainty at all times was not measured in low-latency. While preparing for O4, the calibration groups will investigate methods for producing much lower latency measurements of the total calibration uncertainty with the goal of prototyping a low-latency calibration uncertainty pipeline in O4. If necessary, the data will also be re-calibrated offline, with models updated and informed by careful analysis of all measurements taken throughout the observational period [13]. Such data sets have reduced systematic error and well-quantified total calibration uncertainty but at the sacrifice of weeks-to-months of latency. The low-latency calibration uncertainty pipeline in development would provide total calibration uncertainty estimates with a latency of $\sim$hours rather than weeks-to-months.

The tasks critical to the infrastructure and operations of the detector calibration groups include:

- Maintain global network of absolute displacement fiducial references, including photon calibrator (PCAL) and gravitational-field / Newtonian calibrator (GCAL/NCAL) systems.
- Upkeep of precision models of the detector response (low-latency and offline DARM models) as the detectors are commissioned in preparation for O4, including regular characterization of the detectors.
- Maintain and retain extensibility of the DARM loop model software used for workflow from calibration measurements to uncertainty estimation, including software improvement to speed up production of systematic error and uncertainty estimates. Continue efforts toward common software whenever feasible in collaboration across LIGO, Virgo, and KAGRA.
- Maintain and operate the low- and high-latency $h(t)$ data production software, including upkeep of the filters based on the DARM loop model to match the detector changes.
- Maintain and operate the production of a calibration state vector to be released along with $h(t)$ data that indicates the fidelity of the calibrated data.

- Develop and deploy real-time monitoring for the low-latency $h(t)$ data production.
- Continued coordination with observational science groups to facilitate mock astrophysical and detector noise signals injected in the DARM control loop (aka "hardware injections").
- Maintain the calibration software tools used for reviewing and diagnosing calibration issues.
- Develop, maintain, and deploy methods for removing (or "cleaning") known noise sources from the strain data.
- Continually maintain documentation of methods and results, review final calibration products and software, and publish calibration results.

The tasks that the calibration groups will be working towards with longer term R&D include:

- Improve DARM loop models and methods for characterizing the DARM loop systematic error.
- Stream-line and automate (wherever possible) the process of calibrating each interferometer.
- Improve the low-latency calibration software in accuracy, efficiency, and latency.
- Understand the impact of total calibration uncertainty on astrophysical results including detection and astrophysical parameter estimation.
- Improve the understanding of the detectors response above 1 kHz.
- Develop infrastructure for low-latency total calibration uncertainty estimation.
- Continue efforts towards simultaneously reducing systematic error and latency in $h(t)$ production.
- Improve the hardware used as the absolute calibration reference.

1.4 Timing Diagnostics and Development

The LIGO Timing Diagnostics and Development team verifies and monitores the timing performance of the detectors. This is mission-critical task for reliable interferometer operation, astrophysical data analysis and discoveries. The Advanced LIGO timing distribution system enables UTC synchronized timing between different detectors, and provides synchronous timing to sub-systems of the detector. The timing distribution system's status is monitored continuously during observation runs and is periodically tested in-depth via timing diagnostics studies. As of this writing a publication on timing diagnostic studies and on the O3 timing performance is in preparation.

Similarly to O3, the O4 critical timing tasks include:

- verifying traceable performance of the timing distribution system, validity and accuracy of the recorded time-stamp, verifying the accuracy of the distributed timing signals,
- performing in-depth timing diagnostics for exceptional GW candidates (open public alerts and confirmed catalog events) and
- measuring and documenting the timing performance.

Additional/long-term timing tasks include expanding the capabilities of data monitoring tools related to timing, expanding availability of timing diagnostics for various subsystems, and reviewing the physical/software implementation and documentation of the timing distribution and timing diagnostics components.

1.5 Low Latency

The main goals of the low-latency alerts group is the prompt identification and dissemination within the transient astronomy community of gravitational-wave detections during real-time analysis of the interferometers' data. The group provides the integration of multiple tasks, from the astrophysical searches that identify event candidates to the issuing of astronomical alerts. In the third observing run (O3) the LIGO-Virgo Collaborations has successfully provided open public alerts (OPA) to the transient observing communities. The group also provided the infrastructure for identifying coincidence of gravitational-wave triggers with gamma-ray bursts and neutrino events, and sharing them with the community in real-time. In the ensuing fourth observing run by the LIGO-Virgo-KAGRA Collaborations the group will be making incremental changes in the low-latency infrastructure that will (a) improve the efficiency of alert generation, (b) enable their dissemination beyond the traditional Gamma-ray Coordinates Network(GCN)-based platforms, (c) offer the ability to include more detectors (i.e., KAGRA), and (d) provide refinements of our data-products. The main goals of the low-latency development for O4 can be summarized as follows:

- GWCelery developments for further automation of tasks, to improve efficiency, latencies and monitoring of performance and to enable streaming platforms for the dissemination of alerts.
- Integration of KAGRA into the low-latency framework.
- Integration of MoU-based coincidence analyses that may result to public or private alerts.
- Increase automation throughout the entire alert pipeline.
- Issue and shepherd the public alerts during O4 and as they reach out the transient astronomy community via the "EM Advocates" roles.
- Interface and work with the Burst, CBC and DetChar groups in carrying out comprehensive tests of the end-to-end pipeline(s) using replay, simulated and engineering data.

1.6 Computing and Software

The KAGRA, LSC, and Virgo Computing (and Software) working groups(collectively the IGWN Computing group) provide hardware resources and cyberinfrastructure in support of collaboration business and research activities. During O3 this infrastructure was critical to the generation and distribution of all 80 transient gravitational-wave alerts in low-latency to the broader research community [14].

Much of the available personpower is consumed with maintenance and continuing operations of existing tools and services, these being critical to the daily functions of the collaborations, and to ongoing analysis of data from O3 and development in advance of O4. However, we are also looking forward to future observing runs and developing systems to meet the needs of the collaborations for the foreseeable future.

The work contained in this plan is broadly separated into the following key targets:

1. Continue to Deliver: continue successful O3 delivery of critical low and high-latency computing resources & support for Multi-Messenger Astrophysics (MMA) and discoveries, including estimating and optimizing computing needed for O4/O5 science (with a new focus on optimizing post-detection parameter estimation due to previous detection-problem optimization successes). Continue providing critical computing services in support of internal collaboration functions.

2. **Continue Shift to Distributed Computing Model**: continue strategic shift to (a) leverage large-scale external shared computing resources, (b) establish first-ever joint Lab/LSC/Virgo (and begin including KAGRA) computing architecture & vision, with shared responsibility for provisioning computing resources & computing FTEs, and (c) pursue better communication and integration on computing with the wider Physics, Astronomy, and Computer Science communities.

3. **Define Joint IGWN WBS**: develop an O4 IGWN computing WBS in order to address the consequences of growing GW computing scope and collective FTE shortfalls. Supporting the existing modes of operation while also developing and transitioning to a more efficient joint computing model for O4 has been identified as a challenge, which this document will outline a plan to address. Maintenance, operations, and support of existing infrastructure and services consume the overwhelming majority of our available computing effort. Any new software development for those existing services, and the additional work required to carry out major new projects, requires additional effort that in most cases is not available. By accepting a low level of support for existing services (e.g., tolerating single points of failure, human and computer) we can free up enough staff effort to make modest progress on some of these projects, but we cannot deliver all of them. The outcome of this will be a detailed WBS with projects and tasks that can be used to recruit, match, and manage available human resources to execute.

4. **Improve and Integrate IGWN Computing Processes & Management**: IGWN computing management now meets jointly and includes co-chairs from all three collaborations for the first time (LIGO, Virgo, KAGRA). New policies and decisions are made jointly whenever possible, and legacy policies and infrastructure from individual collaborations are gradually being reconciled and unified. This process is removing obstacles and inefficiencies in cross-collaboration science, is a major improvement on past practices, and is expected to pay additional dividends. Some of our focus areas are:

   - Ensure appropriate R&D and community outreach on new and future computing technologies and models needed in the next 2-5 years. In an understaffed situation, there is a significant risk of becoming too reactive and not forward-thinking enough. (We want to continue to be a leader in large-collaboration computing – e.g., LIGO trailblazed use of HTCondor, modern IAM – we should stay in that role.) This requires an investment of computing staff.
   - Communicate, facilitate and incentivise/reward opportunities for computing contributions by collaboration groups and members.
   - Understand and improve diversity of collaboration members working on computing. Ensure consistent, welcoming, and fair treatment of all contributors.
   - Improve policies and procedures to operate at scale.
   - Avoid "shadow IT": avoid ad-hoc, unplanned infrastructure and services being set up such that scientists rely on, which then require professionalization & continued support (unfunded) – without stifling experimentation & innovation by a collaboration of talented scientists.

Out of scope of this computing plan are:

- LIGO and KAGRA Instrument control and data acquisition systems, prior to data delivery for data analysis. For Virgo, this computing is included in the VIRGO section of this document.

- General computing and information technology (IT) for internal LIGO Laboratory, EGO, and KAGRA business functions.

- Individual data analysis pipeline software development (however this plan does include cross-pipeline infrastructure, optimization consulting, etc.)
- GW computing for third-generation (3G) gravitational-wave observatories (2030-). However, we believe that the synergies between current-generation GW computing and R&D for next-generation GW computing are many, and they should be carried out by the same people and organizations, and would like to include 3G planning into the scope of this document in the future.
- Local systems administration of computing resources that are not dedicated to IGWN collaboration computing priorities and subject to prioritization by the LIGO-Virgo-KAGRA Data Analysis Council (DAC).

1.7 Joint Run Planning

The LIGO-Virgo Joint Run Planning Committee (JRPC) is described in section 3.1 of Attachment A of the LIGO-Virgo-KAGRA MOU, M1900145-v2. Appointed by the LSC, Virgo, and KAGRA collaboration leadership, the JRPC coordinates run planning between those collaborations.

The specific coordination tasks the JRPC takes on are

- facilitating information exchange between the LIGO, Virgo, and KAGRA Collaborations to help establish and coordinate the run plans
- maintaining a schedule of EM and astroparticle observer proposal deadlines and major facility and mission availability to aid in synchronizing with the EM and astroparticle community
- maintaining the LVK-published timeline for runs and any other key communication artefacts used to communicate with the observing community
- maintaining an up-to-date online version of the Observing Scenario
- supporting the Rapid Response Teams for Virgo, KAGRA, and LIGO

1.8 Support of Observatories

The Support of Observatories (SO) committee facilitates and enhances the coordination of collaborative activities, primarily between LSC members and LIGO observatory staff. These efforts are mostly directed toward enhancing the performance and understanding of the detectors. They include detector commissioning, data mining, modeling, calibration, hardware-related investigations and implementation and myriad other opportunities to improve the performance of the interferometers.

The SO committee members work with off-site LSC members, observatory staff, and other members of the Operations Division to define projects that can be carried out remotely, in some cases with brief on-site visits by LSC members, or on-site during longer visits and/or by LSC Fellows.

LSC members have been making significant contributions to many aspects of detector performance and improvements; the SO aims to build on existing and past successful collaborations. Both observatories enjoy active LSC Fellows programs engaging 3-4 fellows at each site throughout the year in detector-related projects.

1.9 Open Data

The Open Data Working Group prepares and supports public data releases, as well as other resources available through the Gravitational Wave Open Science Center (`https://gw-openscience.org`). These include public access to gravitational-wave strain data and catalogs of events, as well as documentation, tutorials, and workshops to support these data products.

1.10 LIGO A+ Upgrade

The “A+ detector" project is a major upgrade to the existing Advanced LIGO detectors, which began in 2019 and is expected to continue through the end of 2023. The first phase of the A+ project in 2021-2022 will be carried out in parallel with preparations for the O4 observing run.

1.11 LIGO-India

LIGO-India is a project of the Government of India with primary responsibilities to build facilities and assemble, install, commission and operate an advanced LIGO detector provided by LIGO and the US National Science Foundation. Several important activities are expected to be completed in 2020-21: completion of a national testing and training facility at RRCAT, initiation of observatory construction activities, vacuum chamber prototype acceptances, beam tube prototyping and testing,

2 Observatory Operations

Op-2.1 LIGO Laboratory Operations

LIGO Laboratory operations were affected by the COVID-19 pandemic and associated shutdowns at the observatory sites and the Caltech and MIT campuses. The O3 run was suspended on March 27, 2020 and the sites adopted a configuration that maintained critical systems with a minimum of staff working on-site. In July 2020 the LIGO observatories transitioned from what was essentially a closed state, and began allowing site personnel to return in limited quantities with restrictions on activities to allow social distancing and other safe work practices. These protocols along with regular testing allowed installation and commissioning work to continue albeit at a slower pace. The sites are now planning further de-escalation and a move closer to normal operations based on the improved situation and the availability of vaccinations. Partly as a result of the pandemic O4 will now start no earlier than August 2022.

In the year ahead the observatories will continue upgrading both detectors for the O4 run. After the shutdown following O3 both observatories succeeded in recovering the detectors to full operation. A major accomplishment at Hanford saw the replacement of an Input Test Mass(ITM) optic. Point absorbers in the coating had limited H1 performance during O3. With the new point absorber free Test Mass the sensitivity improved to O3 levels before the addition of squeezing. It may be possible to align the instrument so as to avoid point absorbers on the End Test Masses (ETMs), and this skip a planned replacement of these optics. L1 will still need to replace both ETMs.

Other activities at the detectors include installation of new laser amplifiers to allow a doubling of the power in the arm cavities, enhanced squeezing by upgrading to lower loss Faraday isolaters and adding a long filter cavity to allow frequency dependent squeezing, better mode-matching, improvements to scattered light mitigation, and improving the duty cycle by damping some suspension resonances.

There are many detector-related activities at the LIGO Hanford and Livingston observatories to support Observatory Scientific Operations:

- The Commissioning team is charged with bringing the detector configuration to a state that is appropriate to meet the upcoming run goals, and to document and transfer operating knowledge to the Detector Engineering group and operators. This activity is detailed below.
- The Detector Engineering group monitors, characterizes, maintains, and repairs working detector configurations. Their goal is high-quality reliable uptime during runs. The Detector Engineering Group leaders manage Engineering Runs at the observatories, setting day-to-day priorities, scheduling work, approving interruptions, and tracking progress. The Detector Engineering Group leaders also chair the daily (or similar cadence) Engineering Run Management meetings and closely coordinate with the LSC Operations Division during the run up to an observing run.
- The Detection Coordinators work for the best science outcome and lead Observational Runs at the observatories. Together with the LSC chair of the RPC they closely plan and monitor run readiness and performance.
- The Computing, Electronics, Facilities and Vacuum teams support operations both directly and indirectly related to the detectors. In general, these groups give priority to the operational phase currently underway, be it commissioning, running, or key preparations for these. During runs these groups' activities will be carried out in close consultation with Run Management. As they also support high-priority non-run-related activities, such as the safe stewardship of Vacuum Equipment and infrastructure, cyber-security patching, and employee safety, not all of their work can be effectively overseen by Run Management.

- The LIGO Laboratory Systems group is central to the planning of all activities related to the detectors including vacuum refurbishment efforts. Typical activities that will be undertaken over the next year will focus on particulate control, stray light control, test mass point absorber R&D, improved automation of detector operation, vacuum system recovery, and the extensive modifications and installations needed for the A+ upgrade.
- The LIGO Laboratory Control and Data Systems (CDS) group maintains and updates the CDS suite of software used in real-time control and data acquisition systems deployed to the LIGO sites and R&D facilities. This includes introducing updates to the software suite based primarily on changes in software packages not developed in-house and computer technologies (software improvement) and providing general support in the area of electronics design, fabrication, test and maintenance (electronics improvements).

ACTIVITY Op-2.1-A: LIGO LAB OPERATIONS

TASK Op-2.1-A(i): LIGO LAB COMMISSIONING

TASK Op-2.1-A(ii): LIGO LAB DETECTOR ENGINEERING AND OPERATIONS

TASK Op-2.1-A(iii): LIGO LAB DETECTION COORDINATORS

TASK Op-2.1-A(iv): LIGO LAB COMPUTING AND CDS

TASK Op-2.1-A(v): LIGO LAB ELECTRONICS

TASK Op-2.1-A(vi): LIGO LAB FACILITIES

TASK Op-2.1-A(vii): LIGO LAB VACUUM

TASK Op-2.1-A(viii): LIGO LAB SYSTEMS

Op-2.2 GEO Laboratory Operations

ACTIVITY Op-2.2-A: ACTIVITIES RELATED TO GEO600 OPERATIONS

The GEO Collaboration is responsible for the operation and maintenance of the GEO600 Observatory, often taking data in “AstroWatch" mode while the LIGO detectors are being commissioned. Many technology developments to be implemented are first tested in GEO600, as was the case for the use of squeezing, now installed in Advanced LIGO.

TASK Op-2.2-A(i): CONTRIBUTION TO GEO600 OPERATIONS

Vg-2.3 Virgo Operations and Commissioning

Following the reduction of the lockdown restrictions (due to the COVID-19 pandemic), and according to the Italian legislation, it has been possible to re-start the operations at the EGO site on May 2020. The EGO Management has then prepared and implemented a protocol to guarantee the safety of the EGO staff and Virgo Collaboration members while working on the installation of the upgrades of the detector. From mid-June 2020, many Collaboration laboratory workshops become operative again and work-related travels were again allowed by the national and international laws.

The installation actitivies were carried out at the highest speed, compatible with the pandemic evolution and related restrictions, and were completed in May 2021. This first upgrade phase, coordinated through the Advanced Virgo Plus (AdV+) Project (see Section Vg-11.2), has the main goal to reduce the quantum noise of the detector by implementing:

- the signal recycling technique;
- the increase of the laser power injected into the interferometer;
- the injection of frequency dependent squeezed vacuum states.

Moreover, during the AdV+ Phase I, an array of seismic sensors in the central and end buildings has been deployed to test Newtonian noise cancellation techniques. Furthermore, an instrumented baffle, for stray light control, has been installed on the Input Mode Cleaner end mirror. In parallel with the other upgrades, an effort has been done to reduce the environmental noise coming for the air conditioning system, and to tackle several other technical noises.

Since January 2021, following the progress of the new hardware installation, it has been possible to re-start the commissioning activities. The plan until Summer 2021 is to focus on stabilizing the working point of the interferometer. A first milestone has been met on March 2021, when the Dual Recycled Michelson Interferometer has been locked. During next Summer, the whole interferometer is planned to be locked at 25 W of input power, which will be brought to 40 W (the foreseen value for O4).

In parallel, the quantum noise reduction system is being commissioned, with the plan to perform the first measurement of frequency dependent squeezing during the Summer, and to complete the integration with the detector in Fall 2021.

After the full interferometer has been locked, the main tasks for the detector preparation in view of O4 will concern:

- optical characterization and tuning;
- scattered light mitigation;
- noise hunting:
  - first without squeezing;
  - then with squeezing.

The Virgo DetChar, Calibration, Low-Latency and Computing & Software Groups are in the process of designing and implementing upgrades in view of O4. The groups, depending on personpower availability, will support the commissioning and noise hunting activities. Efforts are ongoing to strengthen the connections between the groups and to establish priorities, in order to optimize the use of the available resources.

Kg-2.4 KAGRA Operations

The operations of the KAGRA Observatory were affected by the COVID-19 pandemic. The on-site activities were temporarily suspended after completing the joint observation with GEO600 in April 2020. The restrictions were gradually relaxed, and activities resumed in mid-July 2020. However, some restrictions are still applied and may remain for some time in the future. We first attempted the DRFPMI trial from mid-July 2020 to mid-September 2020. After that, we have been working on upgrading the instruments for O4.

The impact of the COVID-19 pandemic has shortened the actual time available for preparation for O4 and has severely limited the human resources available to work on-site. As a result, it is challenging to keep the schedule and targets as planned initially, and the balance between the sensitivity that can be achieved and the processes that can be cut down is under discussion. Under such circumstances, we are trying to accomplish by the time O4 start include the following items:

- Solve the problem of frosting on mirrors and windows in cryostats.
- Improve the Type-A/B/Bp seismic noise isolation system.
- Update the DGS system, from 1st generation to 3rd generation.
- Mitigate the stray light by installing the additional baffles.
- Install the flow mater to investigate the Newtonian noise caused by groundwater.
- Install a High power laser (but with low power operation).
- Upgrade the systems of Calibration, Detctor Characterization, and Computing and Software groups for O4 (see each section).

In parallel with the hardware upgrade, we have been commissioning parts of the system as it becomes available. Before the start of O4, all the interferometer configuration will be fixed, and full commissioning will be performed, where the noise hunting and the fine-tuning are carried out as well as the calibration measurements and the detector characterization investigation. Then, through the Engineering Run, we will join O4 with LIGO and Virgo. In parallel, the introduction of larger mirrors and Filter Cavity in O5 is being considered and discussed.

The KAGRA Collaboration has two main lineages of organizations, which will lead the O4 observations and their preparation. The first one is the Executive Office (EO) which handles all issues regarding the experiment. Under EO, System Engineering Office (SEO) is organized, managing the hardware installation/upgrade, commissioning, and operations. SEO or Run Coordinator holds a bi-weekly chief meeting to discuss and coordinate the short- and medium-term directions and schedules among the subsystems. Commissioning is led by the Commissioning Team to optimize the interferometer configurations. The task scheduler works together with each team to coordinate the weekly and detailed daily schedules. The second one is the KAGRA Scientific Congress (KSC) which is the KAGRA collaborators' decision-making body regarding its scientific direction and strategy. The KSC board organizes the collaboration meetings and coordinates the joint observation with LIGO and Virgo, and other groups. Under KSC, Data Analysis Committee (DAC) is organized, which leads the data analysis and preparation of the pipelines. CAL, DET, and Comp&Soft have branches in SEO and DAC, respectively, and perform both on-site related work and data analysis related work in each branch. They also support commissioning.

3 Detector Characterization

The LIGO, Virgo and KAGRA detector characterization groups have the dual responsibilities of investigating and mitigating misbehavior in the instrument, and providing data quality (DQ) information to the gravitational-wave searches to reduce the impact of artifacts in the data. In addition, the detector characterization groups must help to validate the quality of the data around the time of candidate detections. Support and development of existing software infrastructure key to these tasks is vital to success.

In parallel, there are a number of research and development tasks key to DetChar's charge. Longer-term goals are development of new methods, or improvement of existing methods, for noise identification and mitigation. New methods should be supported with opportunities for performance testing with a well-understood data set and feedback from the DetChar groups.

Ahead of O4, the DetChar groups will evaluate the products and services required for the coming observing run and develop a coherent plan for development and review, in close collaboration with the data analysis groups, the commissioning groups, and the calibration groups. Complementary activities are described in the Observational Science and Instrument Science white papers.

Op-3.1 LIGO Detector Characterization Operational Plans

The following sections outline the priorities for LIGO detector characterization work in 2021-2023 in terms of preparation for **O4**, tasks necessary during the fourth LIGO-Virgo-KAGRA observing run, or tasks required for **future observing runs** beyond O4.

LIGO detector characterization priorities are also specified as **central** or **critical research and development** tasks. **Central** tasks are any task required for the delivery of DQ products to the astrophysical searches and the public, automation of key DetChar activities, event validation, and engagement in detector commissioning. **Critical research and development** tasks are those undertaken to address noise sources problematic for a particular astrophysical search or search method, techniques that don't contribute directly to generating DQ products, or exploratory research and development which is not yet certain to have a direct application to the central DetChar scope.

Op-3.1.1 Preparation and delivery of the fourth aLIGO observing run - O4

Prioritized list of central tasks

Highest priority tasks

ACTIVITY Op-3.1-A: DQ PRODUCTS FOR THE ASTROPHYSICAL SEARCHES

TASK Op-3.1-A(i): DEVELOPING, TESTING, AND DOCUMENTING OFFLINE DQ FLAGS

TASK Op-3.1-A(ii): PRODUCING AND REVIEWING OFFLINE VETO DEFINER FILES, WHICH DEFINE THE DQ FLAGS THAT ARE USED TO VETO DATA FOR EACH INDIVIDUAL ASTROPHYSICAL SEARCH

TASK Op-3.1-A(iii): SUPPORT FOR ADDITIONAL OFFLINE DQ PRODUCTS SUCH AS IDQ TIME-SERIES

TASK Op-3.1-A(iv): SPECTRAL ARTIFACTS (LINES OR COMBS) MONITORING AND STUDIES TO CHARACTERIZE, IDENTIFY, AND PRODUCE LISTS OF LINE ARTIFACTS THAT IMPACT SEARCHES FOR PERSISTENT GRAVITATIONAL WAVE SIGNALS

TASK Op-3.1-A(v): SUPPORT OF ONLINE DQ PRODUCTS

This includes development and maintenance of a git repository with configuration files and products delivered with low latency h(t) frames: state information, data quality flags, iDQ timeseries, and veto definitions

TASK Op-3.1-A(vi): REVIEW OF OFFLINE AND ONLINE DQ PRODUCTS .

TASK Op-3.1-A(vii): INVESTIGATION OF WORST OFFENDER NOISE SOURCES

These are known to include:

- Scattering
- Blip glitches
- Anomalous environmental coupling (e.g., thunderstorms, periscope motion, beam clipping)
- Spectral Artifacts
- *See LIGO DCC P2000495 for more information on the noise sources most prevalent in O3.*

TASK Op-3.1-A(viii): PRODUCE, TEST AND REVIEW GLITCH MITIGATED FRAMES FOR PARAMATER ESTIMATION ANALYSES

For reference, the O3a data mitigation review document is at [15] and the O3b data mitigation status is at [16]

TASK Op-3.1-A(ix): DEVELOP AND PRODUCE GATED H(T) FRAMES WITH LOUD SHORT-DURATION GLITCHES SUBTRACTED FOR USE BY SEARCHES FOR PERSISTENT GRAVITATIONAL WAVE SIGNALS (CW AND STOCHASTIC)

TASK Op-3.1-A(x): ASSIST THE CALIBRATION GROUP WITH INVESTIGATING AND TESTING WHICH AUXILIARY CHANNELS ARE NEEDED TO PRODUCE BROADBAND NOISE SUBTRACTED FRAMES, IF NEEDED BY THE COLLABORATION

ACTIVITY Op-3.1-B: VETTING GW EVENT CANDIDATES (INCLUDING THE DATA QUALITY REPORT AND DATA MITIGATION)

TASK Op-3.1-B(i): CONTRIBUTING TO CONDUCTING OR MENTORING GW EVENT CANDIDATE VALIDATION

Using the automatically generated Data Quality Report (DQR) [17, 18] to vet the data quality around gravitational wave event candidates, including evaluating environmental couplings

TASK Op-3.1-B(ii): FIELD A DETCHAR RAPID RESPONSE TEAM TO VET THE DATA QUALITY FOR LOW-LATENCY GRAVITATIONAL WAVE CANDIDATE EVENTS

TASK Op-3.1-B(iii): DEVELOPMENT, MAINTENANCE AND REVIEW OF DQR TASKS
The end goal is full automation (note, work related to the DQR codebase is listed in Op-3.1-E)

TASK Op-3.1-B(iv): PRODUCE TRAINING MATERIAL TO EVALUATE THE DQR FOR GW EVENT VALIDATION

TASK Op-3.1-B(v): DEVELOPMENT AND AUTOMATION OF DATA MITIGATION PROCESS

TASK Op-3.1-B(vi): SUPPORT VALIDATION OF CW GW CANDIDATES, INCLUDING DETCHAR-RELATED ITEMS IN THE PROTOCOL FOR ESTABLISHING CONFIDENCE IN A CW CANDIDATE [19]

TASK Op-3.1-B(vii): SUPPORT VALIDATION OF STOCHASTIC GW CANDIDATES, INCLUDING DETCHAR-RELATED ITEMS IN THE STOCHASTIC DETECTION CHECKLIST [20]

ACTIVITY Op-3.1-C: MONITORING KNOWN OR NEW DQ ISSUES

TASK Op-3.1-C(i): CONTRIBUTING TO CONDUCTING OR MENTORING DQ SHIFTS
Data quality shifts will be the primary means of ensuring full coverage of h(t) data quality analysis for both detectors during the next observing run, including limiting factors to interferometer performance such as weather or earthquakes. Data quality shifters must invest first in training, and a qualified mentor must be identified for new volunteers.

TASK Op-3.1-C(ii): MENTORING AND TRAINING SCIENTISTS PARTICIPATING IN THE LSC FELLOWS PROGRAM
This task supports LSC scientists working at the site to improve the detector data.

TASK Op-3.1-C(iii): MAINTENANCE AND CHARACTERIZATION OF THE PHYSICAL ENVIRONMENT MONITOR (PEM) SENSORS
This includes evaluating the environmental couplings of the interferometers.

ACTIVITY Op-3.1-D: COMMISSIONING SUPPORT

TASK Op-3.1-D(i): TRACKING ISSUES THAT AFFECT INTERFEROMETER UPTIME, SUCH AS SEISMIC MOTION

TASK Op-3.1-D(ii): INVESTIGATING NOISE SOURCES THAT LIMIT DETECTOR SENSITIVITY
For example, hour-scale correlations between h(t) and auxiliary channels or jumps in detector binary neutron star range.

ACTIVITY Op-3.1-E: MAINTAINING KEY TOOLS

TASK Op-3.1-E(i): SUPPORT OF KEY TOOLS, INCLUDING UPGRADE WORK SUCH AS PYTHON 3 COMPATIBILITY, AS WELL AS USER FEEDBACK AND DOCUMENTATION FOR KEY INFRASTRUCTURE AS LISTED BELOW .

TASK Op-3.1-E(ii): REVIEW OF KEY TOOLS, PARTICULARLY THOSE USED TO GENERATE VETOES
.

List of key infrastructure and tools pertaining to Op-3.1-E:

- Fundamentally necessary services:
  - The summary pages [21]; an invaluable set of webpages containing key plots that describe the state and behavior of the LIGO detectors and their environment.
  - The Data Monitoring Tool (DMT) [22], including the low-latency DMT DQ vector infrastructure.
  - The segment database [23, 24, 25]; which stores state and DQ flag information used by the astrophysical searches.
  - Omicron triggers [26], which identify transient noise triggers, including in low latency, delivered with very high reliability.
  - Data Quality Reports [17, 18], which automatically produce and display results from all tools necessary to validate the data quality surrounding candidate events. Many of the highest priority software tools are needed for this necessary service.
  - Safety information [27], [28], which flags auxiliary channels that witness the gravitational wave strain readout, such that a passing gravitational wave might also induce a response in an *unsafe* auxiliary channel in addition to h(t). Accurate and up-to-date safety information is necessary for interpretation of any tools that correlate auxiliary channels with h(t) to model or infer noise couplings.
- Highest priority software and services
  - iDQ [29]
  - SNAX toolkit [30]
  - GWpy [31]
  - GW-DetChar [32] (which contains Omega scans, LASSO, automated monitoring of scattering, ADC/DAC overflows and software saturations)
  - GWSumm [33] (which is used to generate the summary pages)
  - Stochmon [34]
  - STAMP-PEM [35]
  - Hveto [36]
  - UPV [37]
  - ligoDV web [38]
  - Channel Information System [39]
  - Pointy Poisson [40], [41]
  - VET [42]
  - Offline noise subtraction code (e.g., [43], [44], [45])
  - Suite of remote access tools (remote MEDM and EPICS [46], remote DataViewer [47]
  - LigoCAM [48]
  - Automated safety studies [49], [28]
  - GravitySpy [50]
  - FScans and dependent programs / scripts (spectral monitoring, daily ratio, comb tracker, etc.)[51]

– Coherence tool [52]

This list relies on software dependencies maintained by the LIGO Laboratory (e.g., the Guardian [53],), the LSC Computing and Software Committee (e.g., low-latency data distribution), the LSC Remote Access group (e.g., NDS2), and the Virgo Collaboration (e.g., Omicron, NoEMi). While those software elements are not in the scope of DetChar, continued maintenance of the software, adaptation for use on LIGO data, and operations on LIGO data are of the highest priority to enable LIGO science.

List of development standards for DetChar codebase:

- Code is open source
- Code is hosted on github.com or git.ligo.org to enable a github-flow-style development cycle
- Code includes web-accessible documentation
- Code has been through review
- Code includes unit testing
- Code includes clear and complete installation instructions
- Code configuration files are available and up to date
- Python is recommended for development to maximize compatibility with existing tools, reducing duplication-of-effort and redundancy

ACTIVITY Op-3.1-F: INTERFACING WITH COMMISSIONERS, SITE STAFF, AND SEARCH GROUPS

TASK Op-3.1-F(i): INTERFACING WITH OTHER GROUPS

This can include the following:

- Interfacing with commissioners and instrument experts to propagate instrument changes and developments to detector characterization investigations and monitoring.
- Using the Fault Reporting System (FRS) and the electronic logs (alogs) to communicate results and request tests.
- DQ liaisons identified by each pipeline should identify and report sensitivities in the pipelines to data defects.

Prioritized list of critical Research and Development tasks

Highest priority tasks

ACTIVITY Op-3.1-G: INVESTIGATION OF THE SEARCH BACKGROUNDS

TASK Op-3.1-G(i): STUDYING HOW INSTRUMENTAL ARTIFACTS AFFECT THE SENSITIVITY OF A SPECIFIC SEARCH OR SEARCH METHOD .

TASK Op-3.1-G(ii): DEVELOPING SEARCH-SPECIFIC TECHNIQUES FOR NOISE MITIGATION. THIS INCLUDES CONTRIBUTIONS TO DEVELOPING / INCORPORATING LOW-LATENCY DQ INFORMATION INTO THE ONLINE SEARCHES .

TASK Op-3.1-G(iii): INVESTIGATING THE LOUDEST BACKGROUND OUTLIERS FOR A SPECIFIC SEARCH OR SEARCH METHOD .

The standardized metric for assessing the impact of DQ information on a particular search will be search volume-time (VT), measured by the effect on the background of each search and on recoverability of signals. Population model assumptions in these studies should be clearly documented.

ACTIVITY Op-3.1-H: MACHINE LEARNING FOR O4

TASK Op-3.1-H(i): GRAVITY SPY MACHINE-LEARNING CLASSIFICATION, REGRESSION, AND DATA MINING TO IDENTIFY INSTRUMENTAL CAUSES OF GLITCH CLASSES, AND WORKING WITH GRAVITY SPY CITIZEN SCIENTISTS TO IDENTIFY NEW GLITCH CLASSES AND TO IMPROVE THE QUALITY OF THE TRAINING SET FOR THE MACHINE-LEARNING ALGORITHM .

TASK Op-3.1-H(ii): MACHINE LEARNING [54] CLASSIFICATION, REGRESSION AND DATA MINING STUDIES TARGETING KNOWN NOISE SOURCES, E.G., SCATTERING .

Op-3.1.2 Future observing runs

Prioritized list of central tasks for future observing runs

Highest priority tasks

ACTIVITY Op-3.1-I: ORGANIZATION OF PLANNING AND DEVELOPMENT FOR NEXT OBSERVING RUN

TASK Op-3.1-I(i): CONTRIBUTIONS TO DESIGN AND EVALUATION OF THE COHERENT PLAN FOR DQ PRODUCTS AND DETCHAR SERVICES FOR UPCOMING OBSERVING RUN .

TASK Op-3.1-I(ii): CONTRIBUTIONS TO ORGANIZATION AND UPKEEP OF THE LIGO DETCHAR WORKPLAN AND REQUIREMENTS DOCUMENT .

ACTIVITY Op-3.1-J: AUTOMATION OF KEY TOOLS

TASK Op-3.1-J(i): AUTOMATION OF KEY TOOLS .

ACTIVITY Op-3.1-K: CHARACTERIZATION OF INTERFEROMETER AND AUXILIARY CHANNELS DURING A+ INSTALLATION AND FOR FUTURE OBSERVING RUNS

TASK Op-3.1-K(i): DOCUMENTATION OF PLANNED OR NEWLY INSTALLED INTERFEROMETER SUBSYSTEMS AND ENVIRONMENTAL MONITORS .

TASK Op-3.1-K(ii): MAINTENANCE OF LISTS OF AUXILIARY CHANNELS USEFUL FOR DETCHAR STUDIES TO INCLUDE NEW SUBSYSTEMS AND ENVIRONMENTAL MONITORS .

TASK Op-3.1-K(iii): MAINTENANCE OF SUMMARY PAGE CONTENT TO INCLUDE NEW SUBSYSTEMS AND ENVIRONMENTAL MONITORS .

TASK Op-3.1-K(iv): SIGNAL FIDELITY STUDIES OF NEWLY INSTALLED AUXILIARY CHANNELS .

TASK Op-3.1-K(v): AUXILIARY CHANNEL SAFETY STUDIES FOR NEW SUBSYSTEMS AND ENVIRONMENTAL MONITORS .

TASK Op-3.1-K(vi): DEVELOPMENT AND IMPROVEMENT OF PEM SENSORS AND SENSOR CHARACTERIZATION FOR A+ AND FUTURE OBSERVING RUNS, E.G., MAGNETOMETERS TO MONITOR SCHUMANN RESONANCES IN THE EARTH'S ELECTROMAGNETIC FIELD .

ACTIVITY Op-3.1-L: IMPROVE MONITORS OF KNOWN DQ FEATURES

TASK Op-3.1-L(i): IMPROVING MONITORING AND REPORTING OF DIGITAL AND ANALOG OVERFLOWS, REACHING SOFTWARE LIMITS, AND OTHER KINDS OF SATURATIONS; MONITORING AND REPORTING OF REAL-TIME DATA HANDLING ERRORS (TIMING, DROPPED DATA, ETC.) .

TASK Op-3.1-L(ii): IMPROVING MONITORS FOR EXCESS MIRROR MOTION LEADING TO SCATTERED LIGHT .

TASK Op-3.1-L(iii): SCHUMANN RESONANCE STUDIES .

TASK Op-3.1-L(iv): DEVELOP TOOL TO QUERY STOCHASTIC MONITORS TO FIND WHICH AUXILIARY CHANNELS ARE COHERENT WITH THE GRAVITATIONAL WAVE STRAIN DATA AT A GIVEN FREQUENCY .

TASK Op-3.1-L(v): OPTIC SUSPENSION RESONANCE 'VIOLIN MODE' MONITORING .

ACTIVITY Op-3.1-M: CURATION OF DQ INFORMATION FOR PUBLIC DATA RELEASES THROUGH GWOSC

TASK Op-3.1-M(i): CURATION, DOCUMENTATION, AND REVIEW OF DQ VETOES FOR RELEASE BY THE GW OPEN SCIENCE CENTER (GWOSC) [55] .

TASK Op-3.1-M(ii): DEVELOPMENT AND DOCUMENTATION OF THE "DETECTOR STATUS" PUBLIC SUMMARY PAGES HOSTED BY THE GWOSC [21] .

LT-3.2 LIGO Detector Characterization Long Term Plans

High priority tasks

ACTIVITY LT-3.2-A: PREDICT DETECTOR PERFORMANCE BASED ON INSTRUMENT STATE USING MACHINE LEARNING

TASK LT-3.2-A(i): LOCK LOSS PREDICTION .

TASK LT-3.2-A(ii): PREDICTION OF NOISE CHARACTERISTICS OR OTHER DETECTOR PERFORMANCE METRICS .

Prioritized list of critical Research and Development tasks for future observing runs

Highest priority tasks

ACTIVITY LT-3.2-B: INTEGRATION OF KEY TOOLS TO BE CROSS-COMPATIBLE

TASK LT-3.2-B(i): INTEGRATION OF KEY TOOLS TO BE CROSS-COMPATIBLE

- Wherever possible, all tools in common use (i.e. excluding those in the early stages of development) should share a well-maintained, well-documented, and accessible codebase.
- All triggers and data products will be stored in appropriate common data formats [56] and will be discoverable with common tools (see key tools listed in Section Op-3.1-E). For instance, any data product should be accessible in a single function call on a site cluster.
- Improve documentation and support of key tools: All DetChar tools in common use should be fully and centrally documented, accessible on the LDAS clusters (or easy to install), and well supported by responsive experts.

High priority tasks

ACTIVITY LT-3.2-C: QUANTIFY THE IMPACT OF TRANSIENT NOISE ON PARAMETER ESTIMATION

TASK LT-3.2-C(i): TEST THE EFFECTS OF TRANSIENT NOISE ON RECOVERED SOURCE PROPERTIES

TASK LT-3.2-C(ii): DEVELOP AND TEST METHODS TO RECONSTRUCT AND REMOVE FROM H(T) ISOLATED GLITCHES AND OTHER NOISE TYPES WITHOUT AUXILIARY WITNESSES

ACTIVITY LT-3.2-D: RESEARCH AND DEVELOPMENT OF NEW METHODS FOR NOISE IDENTIFICATION/MITIGATION

TASK LT-3.2-D(i): RESEARCH AND DEVELOPMENT OF NEW METHODS FOR NOISE IDENTIFICATION/MITIGATION (18 MONTHS)

Any new methods are to be tested and validated on recent Advanced LIGO data in a performance test outlined by the DetChar group.

Additional priority tasks

ACTIVITY LT-3.2-E: DEVELOPMENT OF IMPROVEMENTS TO EXISTING TOOLS/PIPELINES FOR NOISE IDENTIFICATION/MITIGATION

TASK LT-3.2-E(i): DEVELOPMENT OF IMPROVEMENTS TO EXISTING TOOLS/PIPELINES FOR NOISE IDENTIFICATION/MITIGATION

For example, exploration of supplementary machine learning techniques for glitch classification, data mining or regression, and supplementary event trigger generators outside of software listed in Op-3.1-E. Any new methods are to be tested and validated on recent Advanced LIGO data in a performance test outlined by the DetChar group.

LT-3.2.1 LIGO DetChar Roles

There are many active roles within the LIGO detector characterization group, and often some people have more than just one role. There are two appointed DetChar chairs at present who oversee and steer the entire group. Working alongside them is a small committee who lead the low latency data quality, instrument characterization, DetChar-specific computing, and event validation efforts. This committee is structured by the DetChar co-chairs, and members are appointed by the DetChar co-chairs. There are five liaisons between search groups and DetChar. The data quality shift coordinator oversees the staffing of the data quality shifts and serves as a point of contact for data quality shift mentors. The safety studies coordinator coordinates with the Hardware Injection team to perform regular DetChar Hardware Injections, oversees the analysis of these safety studies and maintains the channel list. The review chair of the LIGO DetChar group manages the review of critical DetChar code and coordinates code configuration control with other working groups for observing runs. A small group of people also oversee, maintain and develop the key software required by the DetChar group.

3.3 Virgo Detector Characterization

This section is an update of the one published in last year's white paper [2], that the reader can refer to in order to have more information about how the Virgo DetChar work has been categorized. The same task division is reused: three main areas, each associated with a different time period (post-O3, O3-O4 upgrades, O4 and beyond) and two timescales: key tasks directly targeting O4 and R&D tasks that make take longer or whose scope is likely to evolve over time. Already existing items or projects have been updated while a few new ones have been added, based on the past year Virgo DetChar activities: they are *italicized*. Another source of input is the LVK DetChar requirement document, written earlier this year [1].

Vg-3.3.1 Operational Plans

- **Completion of O3 tasks and lessons learned**
  - Offline flags and vetoes
    Science segments and CAT1 vetoes have been produced for O3b as they were for O3a. A description of these O3 flags has been written [57] and made available to the GWOSC website for the O3a public data release. Dedicated checks of the V1 online frame integrity have been developed to catch the very rare instances of corrupted strain data. With an emphasis of maintaining the latency of these checks low for O4, in order to have the corresponding online frames be regenerated prior to the start of the offline analysis.

    The code developed by the Virgo DetChar group to compute the duty cycle of any configuration of the LV(K) network has been validated by comparing its outputs to those of a similar framework, developed independently in LIGO: their results are trully identical. This code will be reused and further automated during O4.
  - Spectral lines catalog
    The list of lines in the Virgo spectrum for the whole O3 run [58] has been compiled thanks to a joint effort between the Virgo DetChar group, the Virgo environmental noise monitoring group and the CW & Stochastic data analysis groups. The *procedure used* will be documented and reused (and automated where possible), not only to produce the Virgo O4 final list of lines, but also during data taking to provide such information more quickly, both to the instrument

team and to the data analysis groups. While doing this study, several important points have been identified (need to run additional NoEMi runs on a regular basis; production of the list of lines in various formats, including LIGO's and the one used for the Virgo known lines database; comparison of the final list of lines with what is already stored in the lines database, etc.) and will be addressed more in depth prior to the start of O4.

- Data quality for the O3b / full O3 event catalogs

  Data Quality Reports (DQRs) have been (and are still being) run for the new GW candidates reported by studies. These DQRs are still based on the Virgo O3 DQR framework, with minor improvements. Two issues have been identified in this context.

  The first one is the need to keep access to raw data for a while, in particular for the candidates that have been identified with a very high latency – ~ 1.5 years or more – for which the raw frame files may have been removed from EGO due to the finite storage capabilities. To address this problem, a *dedicated framework* has been developed to download back data files from computing centers, store them temporarily at EGO, run the DQR and finally remove these data. The full sequence has been automated to ease its usage and it should be straightforward to port it to the foreseen DQR 2.0 when available. In addition, a parallel mitigation is being pursuited by the EGO computing department that has requested a significant ($\sim 50\%$) increase of the raw data storage buffer. Finally, an additional requirement for the DQR 2.0 is to allow running at computing centers – where raw data from the runs are permanently stored. This is not critical to start O4 but will be useful later on; one of the difficulties there is that the Virgo computing centers do not have (yet!?) and HTCondor farm similar to the EGO one and for which the Virgo DQR has been developed.

  The second issue is the requirement to have a global software environment at EGO that will include Virgo-specific software in addition to versions of LIGO software that are recent enough to benefit from all their functionalities and to allow connecting to joint LVK resources like GraceDB that are regularly updated. This is being addressed by the Virgo computing team, within the new IGWN framework..

- Mining the available data

  Two article projects, still ongoing at the time of writing but mature enough to expect a positive conclusion in the next few months, aim at making the best use of the O3 Virgo data. One is a DetChar/DQ article covering the whole run; the other is a *study of the impact of the external environment on the Virgo performance during the O3 run*.

- Producing information for public data releases

  As explained above, information about vetoes and spectral lines have been produced for the whole O3 run, on the occasion of the recent O3a public data release.

- Addressing personpower issues that have limited the number of tasks the Virgo DetChar group could focus on during O3

  This is the main issue for the Virgo DetChar group and it is still work in progress. Thanks to the numerous discussions that have taken place over the past year, there is now a global awareness within the Virgo collaboration that this is a serious problem and that it must be addressed prior to starting O4. But there has not been any significant, additional, personpower commitment to the Virgo DetChar group yet.

- **Support of the AdV+ phase 1 upgrade and contributions to the characterization of the instrument and to the improvement of its performance**

- Support of commissioning activities

  This activity is currently on hold, waiting for the full control of the upgraded – 'AdV+ phase I' – Virgo detector. Some light support has been provided upon request to the Virgo injection team, in order to investigate some problems of this subsystem.

  One new service that has emerged from an extensive study of the Virgo control losses ('lock losses') during O3 is the need for an automated and versatile *Lock Loss Monitor (LLM)* that will allow the commissioning team to get some information quickly each time a lock loss occurs. That new LLM will supersede the prototype that had been developed prior to O3 (and little used/maintained during the run due to a lack of available personpower). It will likely be based on the successful DQR framework: GW candidate (DQR) $\leftrightarrow$ lock loss (LLM); DQ checks (DQR) $\leftrightarrow$ algorithms testing in parallel several possible causes for the current lock loss.

- Glitches: rate, families, classification, search for their origin and help to mitigate/fix them

  The evolution of the glitch rate in Virgo during the O3 run has been studied in details. An important contribution has been identified as coming from scattered light from the suspended optical benches in the arm terminal buildings when there is bad weather and/or the sea activity seismic level is high. The control of these benches should be improved for O4, in order to produce less scattered light.

  Projects from the Reinforce E.U. project [59] – an interactive glitch web catalog and a citizen-science project to classify glitches, using not only information from the GW strain channel but also from auxiliary channels –, carried out as least partly by Virgo DetChar members, are expected to produce synergies that would benefit to both the AdV+ commissioning and the subsequent data taking phases. In particular by making new glitches available more quickly and to ease their study.

  An improved *BNS range drop monitor* – adaptative, not relying on any absolute threshold to identify drops – has been developed and tested in the past year; it will be run automatically (at least on a daily basis) during the pre-O4 noise hunting phase and later on during the run, in order to provide a list of BNS range drop GPS segments for further analysis – including correlation with glitches.

  Glitch studies will restart as soon as the full upgraded Advanced Virgo detector gets entirely controlled and that a first sensitivity curve is available..

- Investigate spectral lines and noise bumps

  See item 'Spectral lines catalog' above.

- Reinforce instrument $\leftrightarrow$ DetChar liaison

  No real progress has been made yet. This goal should get a higher priority on the instrument side when the full Virgo detector is controlled successfully. The Virgo DetChar group is eager to start that discussion as soon as possible.

- Tools

  In addition to the development described above, work is ongoing to migrate software from the Virgo SVN system to the IGWN GitLab.

• **Preparation of O4 and following runs**

- Development of a Virgo DQR 2.0

  The actual coding for the new DQR has not started yet on the Virgo side, but an EGO computing expert should be able to take part in this project, improving the coding style and making the framework more robust/performing. At the global level, a joint LIGO-Virgo-KAGRA study team has been created, with meetings every other week and a dedicated mailing list. Various requirements have arosen from these discussions, from which many functionalities were already part of the O3 Virgo DQR framework. But these had not been foreseen from the beginning: they were added as patches to fulfill needs encountered during the software development. For the DQR 2.0, they will be integrated in the project from the design phase.

- Reduce latency of online flags and vetoes production

  This will be part of a global project aiming at improving the GW strain online reconstruction, alongside the monitoring of its quality. The end-to-end replay of ∼40 days of O3b will be the natural framework to test the new codes as they become available in the coming months.

- Full automation (and latency reduction) of the production of offline flags and vetoes

  This work will consists in gathering existing macros in a coherent software framework that will access the data, run the monitoring tools, provide summary and control plots and prepare the various segment sets (SCIENCE and CAT1) to be uploaded to DQSEGDB after (human) validation. That code mainly exists already (and was used during O3); what needs to be done is to integrate it a global package and to automate its steps as much as possible.

- Deal with the new storage architecture defined at EGO

  Work is in progress to use the new DetChar storage areas that each correspond to a different usage and are thus optimized accordingly: `online`, `prod`(uction), `dev`(elopment) and `archive`. This includes moving contents from old areas to the new ones and redirecting server outputs. The next challenge will be to organize the automated migration of DetChar products during O4, from `online` (a small disk area with high reliability for low-latency generation of DetChar products) to `prod` (a much larger area, used as high-level data storage during a run), in a way transparent to the users. At east the Omicron triggers and the DQRs will have to experience this storage migration.

- Definition of a new DetChar group organization, targeting data-taking periods: factory-like and operations-oriented

  Definition of a new DetChar shift organization embedded in the global Virgo service task framework

  Construction of a core team of DetChar experts, active in the last period before O4 and during the whole run

  These three items have been grouped as they are tightly connected. As discussed previously, there has not been much progress in these areas over the past year, but this will definitely be an important topic for the months to come.

- Establish permanent data analysis groups ↔ DetChar liaisons; organize regular meetings

  On the contrary, this goal has been reached: there are now Virgo-specific liaisons from all data analysis groups who are attending regularly DetChar meetings, passing on information and requests, getting feedback, etc. We expect a similar Virgo DetChar laison from the low-latency group in the near future.

- Liaison Detector ↔ DetChar ↔ Data analysis to orientate and motivate investigations at the DetChar level

  While there have been some contacts and preliminary discussions during the past year as the O3 offline analysis were proceeding, this activity is expected to ramp up in the final months before O4 and to produce results during the run. The main actions so far have been to make the various players aware of the importance of such joint studies: in return, they should now contact DetChar more quickly when there is a need.

- Improve monitoring

  Various new checks have been developed in the past year to fufill analysis needs – in particular in connection with the two main article projects (DetChar-specific and impact of the external environment during O3) the group is working on. Most of them will be included in the set of tests that will be run automatically and regularly during O4.

- Strengthen connections with LIGO and KAGRA DetChar groups

  The connections are excellent: joint monthly meetings, joint working groups (like the O4 DQR study team) and participation of members of the Virgo DetChar group to LIGO DetChar initiatives like their semi-annual face-to-face and noise sprint meetings.

Vg-LT-3.3.2 Long-Term Plans

- **Completion of O3 tasks and lessons learned**

  - Use O3 data to test and validate new or updated tools developed for the AdV+ phase 1 upgrade or for O4.

    These have been bread-and-butter activities for the past year as there is no newer data available. Although O3 raw data are slowly disappearing from EGO storage to make way for new commissioning data, this capability will be preserved thanks to the end-to-end data replay that is using about 40 days of data from O3b that have been archived separately onsite.

- **Support of the AdV+ phase 1 upgrade and contributions to the characterization of the instrument and to the improvement of its performance**

  - Developing new tools

    In the coming months, we will be *investigating the possibility to use iDQ at Virgo during O4*. The final decision will be based on the resources needed (computing, software and, most importantly personpower) and on a comparison with the perfomance of the Virgo veto streams – that are less sophisticated but can be configured in a conservative way.

    The Virgo DetChar group is also developing a *new tool to manage, reference and provide access to the numerous lists of channels* that are needed and used by various softwares: list of safe channels, list of channels for a specific type of brute-force coherence run, list of channels for a full Omicron scan, etc. This tool will use git/gitlab functionalities to store and update these lists.

    Other than these ongoing projects and the synergies provided by the Reinforce project (see more details above), there is no plan to develop additional tools within the Virgo DetChar group, mainly because of personpower limitations. On the other hand, we are following closely development from other groups (in particular the Virgo Hrec/calibration group) and we will use their new tools in our checks when they are available and that new data come in.

- **Preparation of O4 and following runs**

  - Help improving search sensitivity

    The communication of the Virgo DetChar group has put some emphasis on the fact that we would welcome feedback from search groups, with the lowest latency possible, so that investigations could potentially lead to improvements while the data taking is still ongoing – as opposed to waiting for offline analysis to make such feedback.

- Curation and improvement of existing documentation

  This is ongoing in background mode, each expert/developer being responsible for his/her tool.
- Explore possible convergence between Virgo, LIGO and KAGRA DetChar software

  As explained above, the three DetChar groups will try to have a (partly-)common framework for the DQR in O4, in particular for what concerns its infrastructures.

  In addition, the Virgo group has completed a study of channels safety during O3(b), by analyzing data from hardware injections performed at EGO within LIGO's pointy Poisson framework [60]. *This work will be the basis for more regular Virgo channel safety studies before and during O4.* This experience confirms that despite the fact that documentation and references are available, installing such 'new' software in a different environment and making sure that it works properly takes a significant fraction of the total time spent on the project.

  Finally, a continuous effort has been ongoing at EGO since before O2 to run LIGO's Seismon framework [61, 62, 63] and benefit from its earthquake early warnings. Between the O2 and O3 runs, emphasis was put for practical reasons on making the EGO implementation more robust and providing more information to the detector crew, in particular the operator on duty. Seismon was run successfully during the whole O3, with the Virgo detector controls moved manually to a more robust configuration when the arrival of potentially strong seismic waves could be anticipated. Prior to O4, Seismon will be upgraded to its latest version and we will try to automate the decision progress and the subsequent control configuration changes. In addition, work is ongoing to explore the possibility to complement the input data received from the US Geological Survey by a stream coming from the equivalent Italian institute: the INGV. The expected benefits are an enhanced sensitivity to local or regional earthquakes and a reduction of the early warning latency, as least for a fraction of the earthquakes.

  Beyond these examples, no additional synergy has been identified, at least within the Virgo DetChar group. The amount of human resources needed to install, run and optimize 'external' (i.e. non-Virgo) packages is not to be underestimated. The software gap between the collaborations is slowly being bridged by the development of the IGWN common infrastructures and so this approach could be more productive on a longer timescale, say after O4.
- Machine learning

  No progress has been made in this area: there is currently no Virgo DetChar-related study relying on this emerging field. Again, the lack of personpower and the need to focus on core activities first is a clear showstopper for such long-term 'R&D' studies.

Kg-3.4 KAGRA Detector Characterization

KAGRA detector characterization (DetChar) group aims to enhance the reliability of the gravitational wave (GW) detections by supporting the detector commissioning and better understanding the detector.

DetChar group supports the commissioning activities to attain stable operations and sensitivity improvements of the detector. For these purposes, the GW channel as well as many auxiliary channels will be monitored and analyzed to understand the control status of the detector. Software tool developments for monitoring and analyzing various data are included in the DetChar scope. With such tools, the DetChar group will identify noise sources and their coupling mechanisms so that they can be removed from the GW channel. The DetChar group supports also efficient noise hunting activities by providing the list of identified noise sources. In addition to the detector work, DetChar group serves as an interface between the site and remote collaborators by various DetChar tools such as SummaryPages, which is a web-based monitoring tool.

For interfacing with the data analysis group, DetChar group provides characterization results of the detector. By characterizing the GW channel and various auxiliary channels, the DetChar group produces the data quality (DQ) information for the GW data. Information of auxiliary channels is also provided. The produced DQ and veto information are sent to the Rapid Response Team (RRT) when candidate events are found by various GW search pipelines.

Kg-3.4.1 Preparation for O4 observing run

Support commissioning test including remote site work

- Providing the summary pages and improving contents of the summary pages
- Developing the method for easy data access
- Providing following tools and information for commissioners
  - Summary pages [64]
  - Bruco [65]
  - Fscan [51]
  - Pastavi; WEB base data plot tool [66]
  - Noise budgetter; python base noise budget management tool [67]
- Preparing following tools for providing veto information for search groups
  - Omicron triggers [68]
  - Hveto [36]
  - Omega-scan [69]
  - CAGMon [70]
  - Independent Component Analysis; Multi-channel correlation analysis [71]
- Providing reliable and efficient PEM sensors
  - Monitoring the environment inside and outside the experimental area
  - Performing vibration and hammering test
  - Performing the acoustic and magnetic injection

Detector noise monitoring and evaluation

- Automating glitch trigger generation
- Providing spectral line information
- Providing lockloss information
- Investigating safe/unsafe channels
- Automating veto analysis and veto segment production

- Developing auxiliary channels correlation search pipelines
- Developing and testing a new method to detect (Non-) linear noise propagation by using CAGMon
- Developing and applying methods to remove nonlinearly coupled non-Gaussian noises through the independent component analysis

Kg-3.4.2 O4 observing run

Providing DQ information

- Providing online DQ state vector
- Providing Science-mode DQ segments with shorter cadence than the O3GK run
- Providing segment information by using auxiliary channels
- Providing data category information (at least CAT1) for each search group

Contribution for reliable detection

- Participating RRT
- Providing DQ information for RRT
- Providing veto segment for seach groups

Kg-3.4.3 R&D for future observing run

- Providing safe/unsafe channel information
- Improving ICA algorithm
- Update data category flag including CAT2/3/4

4 Calibration

In the time leading into O4 and throughout O4, it is imperative that each LIGO, KAGRA, and Virgo calibration group remain active in preparation and operation tasks for O4. To do so, the groups must maintain an active role in the Observational Science, Instrument Science, and Operations divisions. The following plans discuss the activities the group will perform in support of Operations. Please see interrelated sections in the parallel white papers for plans better suited therein. The sub-heading of each activity indicates which observatories the given activity applies to.

Op-4.1 Calibration Operational Plans

ACTIVITY Op-4.1-A: HARDWARE MAINTENANCE (LIGO/VIRGO/KAGRA)

TASK Op-4.1-A(i): MAINTENANCE OF PHOTON CALIBRATORS

A rigorous upkeep and monitoring schedule of all absolute displacement fiducial references must be maintained in order to retain accuracy and precision sufficient for characterization of the detectors DARM loop. We continue this maintenance schedule for the photon calibrator systems, including but not limited to: continued interaction with NIST in cross-referencing the collaborations primary power standard with theirs, and generalizing the practice of comparing transfer standards to all four detectors. The latter means shipping of transfer standards between LHO and other sites (LLO, Virgo, KAGRA) before O4, and may reveal effects of transportation on their cross-calibration.

TASK Op-4.1-A(ii): MAINTENANCE OF NEWTONIAN CALIBRATORS

In addition, recent collaboration wide R&D efforts have proven gravitational (Newtonian) calibration systems are viable complements to the photon calibrator systems. Virgo intends to work towards bringing their Newtonian calibrator system to operational quality by O4.

ACTIVITY Op-4.1-B: DETECTOR MODELING AND TOTAL CALIBRATION UNCERTAINTY SOFTWARE (LIGO/VIRGO/KAGRA)

TASK Op-4.1-B(i): DARM LOOP MODEL MEASUREMENTS AND UPDATES

Many measurements inform a carefully constructed model of the DARM control loop in order to achieve precise and accurate calibration. As the detectors are commissioned in preparation for O4, each calibration group will revisit all model parameters, remeasure if necessary, and update them to reflect the changes and upgrades to the detectors. In addition, any large-scale infrastructure changes that are required of the model, which are needed as a result of the upgrades leading into O4, must also be done in a timely fashion (e.g. the addition of signal recycling to the Virgo and KAGRA detectors, or the use of multiple DARM actuators in LIGO).

TASK Op-4.1-B(ii): DARM LOOP MODELING SOFTWARE MAINTENANCE

In LIGO, the DARM loop model and systematic error budgeting software was converted to a `Python`-based package called `pyDARM` between O2 and O3. This software package was successfully used in O3 [13]. However, work will now be done in LIGO to streamline the workflow from (1) measurement analysis, to (2) model development, to (3) installation of that model in to the low-latency and offline pipelines which produce h(t), and finally to (4) systematic error assessment and uncertainty estimation. Currently the Virgo DARM model package is in ROOT, and the intention is to keep this software in ROOT for O4. Leading up to O4, LIGO intends to coordinate with the KAGRA team to ensure the `pyDARM` software package is compatible with their loop models, such that astrophysical analysis may incorporate uncertainty in each detector in a similar way. Work is already underway to develop common methodology for determining DARM loop models and systematic error estimates, where possible, and where not, we exploit our differences as an opportunity to develop improved methods.

TASK Op-4.1-B(iii): LOW-LATENCY CALIBRATION ERROR ESTIMATION

For LIGO, the biggest development in this area will be around producing an estimation of the total calibration uncertainty on the low-latency calibrated data. This will require new methodology that could include more frequent measurements of the detector response as well as new low-latency infrastructure for constructing a calibration uncertainty estimate from these measurements.

ACTIVITY Op-4.1-C: CALIBRATION MEASUREMENTS (LIGO/VIRGO/KAGRA)

TASK Op-4.1-C(i): REGULAR MEASUREMENTS THROUGHOUT OBSERVING RUN

Regular calibration measurements will be taken throughout the O4 observing run. Transfer functions of the sensing and actuation function are obtained through swept sine measurements. The photon calibrator (at LIGO, Virgo, and KAGRA) as well as the Newtonian calibrator (at Virgo) are used to obtain these measurements.

TASK Op-4.1-C(ii): MEASUREMENTS FOR FREQUENT CALIBRATION SYSTEMATIC ERROR ESTIMATES

Additionally, LIGO will investigate new methods for performing regular calibration measurements that allow for more frequent characterization of the calibration systematic error.

ACTIVITY Op-4.1-D: STRAIN PRODUCTION SOFTWARE AND INFRASTRUCTURE (LIGO/VIRGO/KAGRA)

TASK Op-4.1-D(i): MAINTAIN SOFTWARE PACKAGE COMPUTATIONAL ENVIRONMENT COMPATIBILITY

The low- and high-latency calibration pipeline software in each of the LIGO, Virgo, (and eventually KAGRA) detectors requires constant maintenance to ensure smooth operation. Currently, the LIGO calibration pipeline software is a combination of front-end code and a `gstlal`-based code in low latency and `gstlal`-based code in high latency [3]. The Virgo pipeline is in `C` [4]. It is KAGRA's intent to mimic the LIGO pipeline. These software packages must first-and-foremost be adapted in the face of its surrounding computational environment which is, also necessarily, constantly evolving.

TASK Op-4.1-D(ii): UPDATE SOFTWARE TO BE CONSISTENT WITH DARM MODEL AS NEEDED

Further, any changes in DARM model parameters and/or infrastructure will be propagated to the respective pipelines with minimal introduction of systematic error in an extensible and robust fashion. This includes the most current and accurate ways to track time-dependence and communicate live calibration fidelity through the calibration state vector.

TASK Op-4.1-D(iii): PRODUCTION OF STRAIN DATA

Finally, the groups will do their best to produce reviewed low-latency and offline strain data (and its systematic error and uncertainty budget) in regular and predictable intervals in a timely, accurate fashion for consumption by the collaboration and beyond.

Task Op-4.1-D(iv): Cleaning of Strain Data

Certain noise sources, such as calibration lines and power mains lines, are known and can be removed from the strain data. The calibration groups will coordinate with the LIGO, Virgo, and KAGRA Detector Characterization groups to develop, maintain, and deploy the necessary noise-cleaning methods for strain data in O4.

Activity Op-4.1-E: Monitoring (LIGO/Virgo/KAGRA)

Task Op-4.1-E(i): Summary Page Maintenance and Coordination

Detector monitoring web interfaces, a.k.a. "summary pages," based on the `gwpy`, `gwsumm` (for LIGO), and `vim` (for Virgo) software packages, have been used as the primary monitoring tool for calibration outside of the control rooms. These pages need constant maintenance and upgrades to keep up with any changes and evolving checks on the calibration. The LIGO, Virgo, and KAGRA calibration groups will interface with the Computing and Software working group to keep this tool up-to-date with needed calibration monitoring checks.

Task Op-4.1-E(ii): Real-time Monitoring Tools

In addition to working with the software package development teams to maintain those detectors' general summary pages, the LIGO calibration group will also develop real-time monitoring tools for the computational status of each detector's low-latency calibration pipeline to assist in operational status tracking. The KAGRA team expects to do similar.

Activity Op-4.1-F: Hardware injections (LIGO/Virgo/KAGRA)

Task Op-4.1-F(i): Facilitation of Hardware Injections

The photon calibrator system remains the best tool for creating DARM actuation in each detector, whether this be for verification of the data produced by each calibration group, for testing of astrophysical search pipelines, or for understanding the cross-coupling of the detectors' network of auxiliary loops and sensors. LIGO, Virgo, and KAGRA plan to continue the interactive relationship with each appropriate consumer group in order to facilitate these activities.

Several artificial, simulated continuous signals are added via the photon calibrator hardware injection to the detector data stream. These signals are coherent between detectors in order to simulate a real continuous signal. They serve as known fiducial reference signals, acting as invaluable tools aiding validation of detector calibration and data cleaning procedures.

Task Op-4.1-F(ii): Early-warning Calibrated Strain Production

LIGO, Virgo, and KAGRA will modify the existing strain production methods as needed to produce an extremely low-latency calibrated strain data stream that is at least valid in the frequency range needed for early-warning transient searches. The LIGO calibration group will explore methods for modifying existing filters and filtering techniques to create an early-warning calibrated data stream. Virgo will explore using shorter FFTs and modifications to the existing strain production software to create an early-warning calibrated data stream. KAGRA will keep up with the LIGO and Virgo developments and deploy the methods most fitting for the KAGRA system.

Activity Op-4.1-G: Documentation and Review (LIGO/Virgo/KAGRA)

TASK Op-4.1-G(i): CALIBRATION REVIEW

All calibrated data, the generation and monitoring processes there-in, systematic error budgets, and uncertainty must be reviewed. The review process involves creating quantitative summary statistics of the data throughout a given observational time period, which is done using various software packages. Creating and documenting each of these statistics, reviewing the code that generated the model, data and statistics, investigating any anomalies or peculiarities that arise from this review, ensuring the results are reproducible, and fixing any issues that are identified in the process are all essential parts of the review. The review process is necessary but time consuming. All calibration groups plan to participate in any reviews held, and improve the efficiency and speed of creating the quantitative summary statistics of the data, which should assist in speeding up certain aspects of the review process that are limited by computation time.

TASK Op-4.1-G(ii): DOCUMENTATION OF CALIBRATION PROCESSES

LIGO, VIRGO and KAGRA will continue to record their efforts in the electronic logs, software repositories, technical notes, drawings, and graphical presentations document control centers, wiki pages, and peer-reviewed articles. The LIGO calibration group results are posted and documented in the LIGO Document Control Center, in the form of a related-document tree, whose trunk is [72]. Additional, organizational content may be found on the ligo.org wiki page [73].

LT-4.2 Calibration Development Plans

ACTIVITY LT-4.2-A: IMPROVEMENTS TO THE DARM LOOP AND THE CHARACTERIZATION OF SYSTEMATIC ERROR (LIGO/VIRGO/KAGRA)

It has become apparent that no detector's calibration accuracy and precision is limited by the uncertainty in its absolute reference in O3. Rather, the limitations are in understanding and quantifying unknown systematic error, and the frequency-dependent impact of that systematic error due to idiosyncrasies of the design of each detector's control system.

TASK LT-4.2-A(i): MODIFICATION OF DARM LOOP TO MINIMIZE SYSTEMATIC ERROR

The calibration groups plan to work in concert with the detector commissioning groups to modify the design of the loops to minimize contributions of systematic error and to reduce now-known, controllable, systematic error in the DARM via refinement of the interaction between other detector control loops.

TASK LT-4.2-A(ii): UPDATE TOTAL CALIBRATION UNCERTAINTY ESTIMATES

The calibration groups lan o include other now-known, but uncontrollable systematic error in the loop models a priori, and to refine uncertainty estimates of unknown systematic error.

ACTIVITY LT-4.2-B: INCREASED AUTOMATION OF CALIBRATION PROCEDURE (LIGO/VIRGO/KAGRA)

The LIGO calibration group will inventory the current procedures used to calibrate LIGO strain data. From this inventory, steps that can be automated will be identified and automation techniques will be developed. The goal of these efforts is to improve the efficiency of the calibration procedure as well as provide accurate strain data at a more constant and predicable cadence, even during commissioning time periods.

TASK LT-4.2-B(i): AUTOMATION OF CALIBRATION MEASUREMENTS

The LIGO calibration group will upgrade the current systems for automating calibration measurements, such as swept sine measurements, as O4 approaches. This has been developed in the LIGO Guardian software thusfar, and development will likely continue in Guardian. The Virgo calibration team currently uses Metatron for data taking automation. This system will need to be updated as the upgraded Virgo detector comes online in O4. KAGRA will keep up-to-date with LIGO and Virgo developments and employ automation where possible as the KAGRA calibration system is developed.

TASK LT-4.2-B(ii): AUTOMATION OF CALIBRATION MODEL UPDATES

After measurements are taken, the calibration model and filters often need to be updated based on the most recent measurements. All of the collaborations will work to improve the automation of the process from measurement to model updates. This will allow the calibration to remain more constantly accurate.

ACTIVITY LT-4.2-C: IMPROVEMENTS IN LOW-LATENCY CALIBRATION SOFTWARE (LIGO/KAGRA)

TASK LT-4.2-C(i): FRONT-END CALIBRATION PIPELINE IMPROVEMENTS

With continued demand for well-calibrated data to be delivered as soon as computationally possible, each calibration group always strives to develop ways in which low-latency data can be improved. LIGO currently uses the near-real-time processing system for an initial calibration of the strain data. However, some low-latency post-processing of this near-real-time calibrated strain data will likely still remain necessary due to the causal limitations of the real-time system. LIGO will continue to develop methods to push forward the quality of calibration that can be achieved with the near-real-time system, and make efforts to reduce data production from "low" latency to "even lower" latency.

For the LIGO near-real-time system, this includes

- implementing FIR filtering routines to improve frequency-dependent systematic error incurred with the current systems IIR filters.
- continuing development and verification of the near-real-time data product that is corrected for detector response time dependence.

KAGRA intends to keep pace with the development of the LIGO near-real-time software.

TASK LT-4.2-C(ii): LOW-LATENCY CALIBRATION PIPELINE IMPROVEMENTS

The LIGO low-latency system will be improved by

- implementing approximation free methods for calculating time-varying calibration factors that will result in better accuracy.
- improving the computational speed and resource consumption of the pipeline.
- reducing the overall latency of the pipeline.
- implementing real-time monitoring into the pipeline.

KAGRA intends to keep pace with the development of the LIGO low-latency software.

ACTIVITY LT-4.2-D: INCORPORATING CALIBRATION ERROR AND UNCERTAINTY WITH ASTROPHYSICAL RESULTS (LIGO/VIRGO/KAGRA)

Task LT-4.2-D(i): Coordination with Astrophysical Analysis Groups

The LIGO, Virgo and KAGRA groups are partnering with astrophysical analysis groups to understand the impact of systematic error and uncertainty in calibration on astrophysical results. Feedback from studies will inform how (and/or whether) to proceed in improving the existing levels of systematic error and uncertainty. See more discussion in this year's Observational Science White Paper.

Activity LT-4.2-E: Calibration Above 1 kHz (LIGO/Virgo/KAGRA)

Task LT-4.2-E(i): Modeling and Measurements at High Frequencies

Verifying the accuracy of models for the calibration above 1 kHz is challenging due to the difficulty in collecting precision measurements at high frequencies where the detector noise is high and actuator strength is low. Work is ongoing to develop and implement methods for accurately modeling the detector response at high frequencies and to more accurately determine the uncertainty in this frequency range. See further discussion in the Instrument Science White Paper.

Activity LT-4.2-F: Minimze Systematic Error (LIGO/Virgo/KAGRA)

Task LT-4.2-F(i): Correct Source of Known Systematic Errors

Maintaining precise and accurate calibrated data relies on understanding any systematic errors present in the overall scale of the calibration and resolving these errors where possible is standard part of each calibration group's effort. Depending on the results of the above mentioned understanding of the integration of calibration systematic error and uncertainty within astrophysical analysis (including those for which low systematic error above 1 kHz is important), efforts to reduce the systematic error beyond existing levels of will be appropriately matched with the needs.

Task LT-4.2-F(ii): Compensating for Known Systematic Errors

Additionally, if the systematic error is quantified, the calibration groups will develop and implement methods to correct for this systematic error in the strain data, therefore eliminating it to within the level of confidence it was measured. This is unique from eliminating the source of the systematic error. New methods for calibrating the data by frequent, direct measurement of the detector response function will also be explored and developed in parallel to the current calibration methods during O4.

Activity LT-4.2-G: Improve Absolute Reference Hardware (LIGO/Virgo/KAGRA)

Task LT-4.2-G(i): Upgrades to Photon Calibrator Systems

For discussion of hardware improvements to the LIGO Pcal systems, see the LSC Instrument Science White Paper. KAGRA has the newest PCAL system, which includes a 20 W laser, dual-AOM actuation, and steering mirrors for spot-position control. Improvements to reduce laser power noise and beam quality are planned. Virgo plans to replace a few optics and photodiodes to further improve the stability of the power calibration of their PCAL systems with respect to the environment.

TASK LT-4.2-G(ii): UPGRADES TO NEWTONIAN CALIBRATOR SYSTEMS

As Virgo is leading the charge in secondary, gravitational (Newtonian) calibrators, LIGO and KAGRA hope to learn lessons and consider continuing development of an NCAL system (see further discussion in LSC Instrument Science White paper).

5 Timing Diagnostics and Development

Op-5.1 LIGO and KAGRA Timing Diagnostics and Development Plans

Traceable and closely monitored timing performance of the detectors is mission-critical for reliable interferometer operation, astrophysical data analysis and discoveries. The Advanced LIGO timing distribution system, also adopted by KAGRA, provides synchronized timing between different detectors, as well as synchronization to an absolute time measure, UTC. Additionally, the timing distribution system must provide synchronous timing to sub-systems of the detector. The timing distribution system's status is monitored continuously and is periodically tested in-depth via timing diagnostics studies.

ACTIVITY Op-5.1-A: CRITICAL TIMING TASKS

TASK Op-5.1-A(i): VERIFYING TRACEABLE PERFORMANCE OF THE TIMING DISTRIBUTION SYSTEM

TASK Op-5.1-A(ii): VERIFYING THE VALIDITY AND ACCURACY OF THE RECORDED TIME-STAMP

TASK Op-5.1-A(iii): VERIFYING THE ACCURACY OF THE DISTRIBUTED TIMING SIGNALS

TASK Op-5.1-A(iv): PERFORMING IN-DEPTH TIMING DIAGNOSTICS FOR EXCEPTIONAL GW CANDIDATES (OPEN PUBLIC ALERTS AND CONFIRMED CATALOG EVENTS)

TASK Op-5.1-A(v): MEASURING AND DOCUMENTING THE TIMING PERFORMANCE

ACTIVITY LT-5.1-B: LONG-TERM TIMING TASKS

TASK LT-5.1-B(i): EXPANDING THE CAPABILITIES OF DATA MONITORING TOOLS RELATED TO TIMING

TASK LT-5.1-B(ii): AVAILABILITY OF TIMING DIAGNOSTICS FOR VARIOUS SUBSYSTEMS

TASK LT-5.1-B(iii): REVIEWING THE PHYSICAL/SOFTWARE IMPLEMENTATION AND DOCUMENTATION OF THE TIMING DISTRIBUTION AND TIMING DIAGNOSTICS COMPONENTS

6 Low Latency

Op-6.1 Overview and goals for O4

The low-latency alerts group provides support for conducting astrophysical searches close to real-time for the purpose of connecting gravitational-wave observations with their electromagnetic and neutrino counterparts. Furthermore, it also facilitates the diagnosis of issues in the detectors via real-time analyses. The effort builds on expertise in low-latency searches, real-time data acquisition, processing and communications, as well as couplings with the electromagnetic astronomy community, understanding its needs and expectations. The work plan of the group over the August 2021 to January 2023 period is predominantly driven by the needs for public alerts over the upcoming observing run of O4, due to commence in the summer of 2022, the corresponding infrastructure to enable them and continued support for its operation over the course of the O4 run.

The overarching goals of the low-latency group over the upcoming 20 months of activities-planning can be summarized below and will be expanded in corresponding sections that follow:

- Provide a robust and redundant system that will process broadband transient event alerts as well as early warning (pre-merger) ones in real-time. We are aiming for an intrinsic (to the low-latency system) latency of up to 5 seconds. The main workflow will remain the one adopted for O3 (GWcelery: https://gwcelery.readthedocs.io/en/latest). Two alert mechanisms are envisioned to be delivered with the O4 system: a Gamma-ray Coordinates Network(GCN)-style one as used in O3 and based on pre-set session-windows, and a streaming version that will enable going beyond such pre-set session-windows and allow faster cadence of outgoing information from the low-latency pipelines. The streaming platform will be based on *kafka*. Selective/critical services like the event aggregation/grouping ("superevent") and GCN-sender will likely need to be re-worked in order to operate more efficiently.

- Incorporate KAGRA data into the low-latency operations as directed by the Operations Division. This effort will need to be planned and streamlined in the earliest stages of the group's operations.

- Provide functionality to accommodate potential MOU-based collaborations with non-gravitational-wave observatories, including interaction of offline searches and results (i.e., identified events and related data products) with event database (GraceDB) and corresponding alert generation. Like with the planning of incorporating KAGRA into the low-latency searches, such functionality will need to be planned and streamlined in the earliest stages of the group's operations and based on specifications and requirements established by the Observational Science Division.

- Increase the level of automation throughout the entire low-latency alert pipeline. This will be a necessity given the three times (by 04), and ultimately ten times (by O5), increase in the rate of detections that will need to be processed. Automation and quality assurance will be required on all tasks that come together in order to communicate public alerts. This includes but is not limited to detector characterization information, parameter estimation, source classification and characterization. Additionally, we need to aim for streamlined human vetting of events which entails faster decision-making without sacrificing rigor. This effort will be ongoing throughout the year and will be coordinated with the search groups.

- Facilitate prototyping, testing, review and deployment of the low-latency computing software and hardware infrastructure via multiple and flexible environments where these can be undertaken. As part of that, we will streamline the review and overall protocol for deployment of the low-latency-specific software infrastructure. This effort will build on the lessons-learned from O3.

- Examine and define the role of "Electromagnetic (EM) Advocates" in the O4 and beyond era and maintain good coupling with the electromagnetic astronomy community throughout the year via regular email/telecon updates, test data streaming.
- Deliver a low-latency alert pipeline and overall system that will address the needs of the collaborations with minimal changes for O5.
- In coordination with the Burst, CBC and DetChar groups, establish the testing and acceptance criteria for all low-latency search pipelines and related tasks.
- Finalize and carry out search, annotation and alert-triggering and issuing pipelines-testing in a continual form using O3 replay data, simulated data as well as engineering data as soon as they become available. All ingredients of the end-to-end pipeline should be fully tested well in advance of the start of the O4, including establishing of the operational parameters to be used and final acceptance criteria and tests of the end-to-end system.

The overall timeline of the low-latency work in the year ahead includes the following key milestones:

- Early July, 2021: O3 data replay and Mock-Data-Challenge (MDC) streaming is delivered for pipeline to process.
- July 31, 2021: Full roll-out of search and annotation pipelines processing on O3 data replay and MDCs takes place.
- September 30, 2021: First MDC run results are collected and evaluated; examine if more h(t) data is needed (beyond the O3 data replay).
- December 1, 2021: New low-latency infrastructure is commissioned for end-to-end use; review commences.
- December 1, 2021: Weekend "temperature"-checking engineering run (real data flowing, even if garbage) commences and repeated at least once per month from this point on.
- February 1, 2022: Begin testing connection with any newly introduced event brokers (e.g., GCN/TACH, SCiMMA, etc)
- March 15, 2022: MDC studies completes; low-latency alerts pipeline is finalized (one last data replay commences).
- April 30, 2022: Final replay completes; all reviews complete; assessment of what pipelines meet the acceptance criteria for roll-out in O4 takes place.
- May 31, 2022: One month-long engineering run commences; full system in deployed and operated as if in-production.
- July 1, 2022: End-to-end system freezes and effective observing commences.
- July 1, 2022 through March 31, 2023: system operations, monitoring and maintenance throughout O4, including delivery of human-driven aspects of the system.

It is understood that multiple inter-dependencies exist among all the elements that come together to form the low latency alerts pipeline. Given such inter-dependencies, there may be varying degree of completeness of the corresponding tests and milestones for the various elements of the pipeline and during its early developement stage.

ACTIVITY Op-6.1-A: OVERVIEW AND GOALS FOR O4

TASK Op-6.1-A(i): SET AND REVIEW GROUP GOALS THROUGH O4

Op-6.2 Low Latency system architecture

The low-latency system has supported many discoveries over the past three observing runs, during the course of which the architecture and tooling has gone through three distinct phases of evolution, of which the latest is GWCelery. We now have a mature understanding of the requirements from both internal and external stakeholders, and should be progressing toward a durable architecture and software stack that will serve the community through at least the next several years of observing runs.

The unifying constant across observing runs is that we wish to have a simple and reliable system for annotating and orchestrating LIGO/Virgo/KAGRA alerts, built from widely used open source components. Its responsibilities include:

- Merging related candidates from multiple online LVK transient searches
- Correlating LVK events with transient events in the electromagnetic spectrum as well as in neutrinos
- Launching automated follow-up analyses including data quality checks, rapid sky localization, automated parameter estimation, source classification, and source properties inference
- Generating and sending machine-readable alerts
- Sending updated alerts after awaiting human input
- Automatically composing templates for any human-readable prose (e.g. “Circulars”)

As part of assessing how to proceed with the O4 system architecture, we have reviewed the deficiencies in the O3 architecture, identified the requirements for the O4 system and examined the choices for frameworks and tooling to be used. A comprehensive system requirements document for low-latency alerts was made available as part of this study[1] The key aspects of the O4 system include the following:

- ability to provide two mechanisms/platforms for the distribution of the alerts; one based on the GCN protocol and overall approach as in O3 and a new one that will be based on *kafka* streams,
- target a few-seconds ($\sim 5s$) intrinsic latency of the system, including ability to send preliminary alerts with only basic information like the type and significance of the event (i.e., without the complete event payload),
- incorporate under the same infrastructure MoU-based analyses and resulting events that will require public or private alert-issuing,

[1] `https://dcc.ligo.org/LIGO-T2000740`

- maintain multiple development and testing computing environments for deployment of production code,
- retain GWcelery as the main workflow manager but rework few key services like the event aggregator and GCN alert sender.

The realization, maintenance and operation of the O4 alert system is the highest priority task for the low-latency alerts group in this White Paper cycle.

ACTIVITY Op-6.2-A: LOW LATENCY SYSTEM ARCHITECTURE

TASK Op-6.2-A(i): MEET THE PLANS AND NEEDS THROUGH O4

Op-6.3 GraceDB and LVAlert

The Gravitational-Wave Candidate Event Database (GraceDB) saw an expansion in scope and capability in O3, namely moving from a service hosted at UWM to a highly-available cloud service.

At its core, GraceDB is, architecturally, a standard Web/API application. A MySQL (MariaDB) backend is powered by a Django web framework. External requests are served by Apache acting as a reverse-proxy for a Gunicorn-based WSGI HTTP server. Files are stored on an NFS filesystem. Low-latency analyses stream data from the detectors and upload "candidate events" to GraceDB via a representational state transfer (REST) API.

State changes in GraceDB (which may take the form of new event uploads/annotations, new superevent uploads/annotations, log/file updates, etc.) are communicated to outside parties by the LIGO-Virgo Alert (LVAlert) system. LVAlert is an XMPP-based publish-subscribe (PubSub) system powered by the Java-based Openfire XMPP server software. Client-side tools maintained by LSCSoft allow users to listen and respond to LVAlert messages. LVAlert messages are machine-readable JSON-formatted so as to be read by automated follow-up processes.

LVAlert listeners act on notifications from GraceDB and are used to launch follow-up analyses (*e.g.,* Superevent creation, parameter estimation, sky localization, etc.). Results from follow-up analyses are then uploaded and stored in GraceDB.

As such, GraceDB/LVAlert has grown from serving as a queryable database of candidate events to the orchestrator and source-of-truth for external observers and follow-up processes.

Op-6.3.1 Plans and needs for O4

The goals moving into subsequent observation runs for the GraceDB/LVAlert ecosystem include:

- Maintain and support the Gravitational Wave Candidate Event Database (GraceDB) and associated client tool.
- Develop, maintain, and support the LIGO-Virgo Alert (LVAlert) service and client tool for wider IGWN use.
- Expand LVAlert to greater IGWN (and public) community to reflect the growth of the GW observation community.

- Transition LVAlert to high-availability, cloud-native deployment in line with other low-latency services.
- Explore and analyze database backend solutions for latency, throughput, and global availability.
- Explore and analyze low-latency file system improvements.
- Retool GraceDB event ingestion to read from low-latency streaming framework.
- Develop, deploy, maintain public-facing view for curated GW events in GraceDB.
- Explore streaming technologies developed natively and external to IGWN to replace XMPP for LVAlert.
- Explore and analyze authentication and identity access management technologies for GraceDB and LVAlert.

ACTIVITY Op-6.3-A: GRACEDB AND LVALERT

TASK Op-6.3-A(i): MEET THE PLANS AND NEEDS THROUGH O4

Op-6.4 Low Latency review and testing

The system architecture document for the low-latency alerts has been the main activity on the review front over the past 12 months. The main goals for the coming 12 months are:

To provide test instances of the complete low-latency service bundle (GraceDB, LValert, and GWCelery), or only some of its components according to the needs, in high availability on any cloud infrastructure at our disposal. This leads us to consider a deployment strategy based on the usage of Docker containers and their orchestration using Kubernetes. This deployment strategy fits smoothly with the GitLab built-in tools for continuous integration and continuous delivery. Therefore, we plan to apply the same deployment strategy in the longer term to the low-latency pipelines. In this way, it will be possible to automatically deploy self-contained instances of the entire bundle (services and pipelines) reflecting the code changes pushed to GitLab, thus facilitating the process of code testing, reviewing, and stress testing.

Ensure that Standard software development practices like code unit-testing, providing documentation, and examples should be the modus operandi going head. Unit-tests should capture the behavior both at a modular level and the interaction of different components. Runtime error notifications and tracking via Sentry proved helpful in O3, and should be carried over to future runs. As in O3, the full integration will be tested through the periodic issue of exhaustive MDC alerts to the final endpoint (in O3 was through GCN alerts), their listening, and validation. Such continuous integration functionality of MDC alerts will be deployed in the production environment. This will allow any external users to develop and test their listening and integration continuously.

Ensure, through extensive load test, that the system can correctly process the submission of events at least twice the expected load of events uploaded by low-latency pipeline.

To fulfill these goals, we will need 3+1 fully operational instances of the low-latency pipeline: (1) Production. (2) Playground. (3) Testing and Validation (to perform extensive acceptance and stress tests by the review team). (4) Backup production instance (to be activated on-demand in case of failure of the production instance) to achieve high availability and resilience to computer center failure.

ACTIVITY Op-6.4-A: LOW LATENCY REVIEW AND TESTING

TASK Op-6.4-A(i): LOW LATENCY DOCUMENTS REVIEW

Review the design and implementation documents of the low latency alerts system and the "User Guide of the alerts" for the non-LVK community.

TASK Op-6.4-A(ii): REVIEW OF MAIN SOFTWARE STACK OF O4 SYSTEM

The review will also include the execution of extensive stress test of the whole infrastructure including the Gracedb/LValert system that will be deployed for the O4 operations.

TASK Op-6.4-A(iii): PRODUCTION ALERT SYSTEM REVIEW

Continuous review an testing maintenance and updates of the production alert (GWCelery) system.

Op-6.5 Low Latency interfacing with LVK working groups

Over the first three observing runs of the Advanced LIGO and Virgo detectors the low-latency group provided both technical infrastructure and science analyses specific to searches, all of which are required in order to communicate transient detections as public alerts. The focus of this past year was on both the assessment of how these tasks fared over the most recent observing run (O3) and the onset of work toward O4; we are working towards establishing clear expectations and the specifications required in order to accomplish the successful delivery of such infrastructure and analyses. It should also improve w/r/t O3 communication with the LVK search group. Some of the top-level questions for collectively the LVK search groups to address in coordination with the low-latency group include:

- Provide a complete list and documentation of analysis pipelines, including all their "flavors", expected to perform low-latency searches.
- Establish a standard checklist for inclusion of a search pipeline into low-latency operations, including their interface to the ongoing mock data challenges.
- Revisit the definition of "superevent" and its attributes, and the choice of "preferred-event" based on the O3 experience.
- Investigates avenues to improve the computation of trials-factor used during O3.
- Establish which (if any) tasks currently in the workflow should migrate to search pipelines: detector status and prompt vetoes, astrophysical source classification and characterization, sky localization and parameter estimation.
- Develop infrastructure necessary for fully automated dissemination of binary black hole events.
- Better understand the redundancy and complementarity among multiple search pipelines contributing to low-latency alerts, including through the use of quantitative metrics derived from the mock data challenges.
- Streamline the parameter estimation tasks and automate the event/alert updates.

- Investigate which is the best sky localization area to communicate attached to the superevent. In order to maximise our scientific discoveries, we should always send updated localizations from automated parameter estimation, regardless of overlap with the rapid localization. If the localization changes significantly, it is important to send an update; however, if the localization does not change significantly, then it just as important to send an update to confirm that the localization is accurate. We note that one cannot make localization updates dependent upon measurement of sky area or overlaps. A "better" localization is not necessarily a smaller one; a "better" sky map could be broader if it captured uncertainties more realistically. It is well known that bad data can produce very small sky areas. Realistic modeling of detector and waveform systematics tends to broaden localizations. Nevertheless, we need to understand upon release of a new sky-map the overlap between the original and revised skymaps to understand how to present the skymap in a circular for updating ongoing observations. This could simply take the form of an automated calculation of the overlap of the 90% areas and volumes, with some way to establish for astronomers when ongoing observations should be interrupted to change to a new skymap.
- Investigate which (if any) additional external experiments should be added to the joint searches and what additional filtering of parameters in these searches would be useful
- Establish the sequence of event updates within the event database system and public alerts, from real-time detections to offline/final analyses.
- Profile latency of all tasks and implications to the target global latency of <10 s for regular, non-merger, alerts as well as pre-merger ones.

All of the above should be discussed with the search groups and documented in order to set the specifications for the O4 operations. We expect this to start from the results of the mock data challenges, where we can identify shortcomings and improvements needed. The first iteration of the mock data challenge is expected to complete by early Fall 2021, and in parallel implementation of the updated low-latency system is already proceeding. As part of this, we call for the Burst and CBC groups, in coordination with the Low-Latency group to define quantitative metrics based on the mock data challenges to address key questions of the end-to-end low-latency pipeline, including using it as part of the final acceptance test. Such data challenge is to be commonly decided and rolled out by July 2021.

ACTIVITY Op-6.5-A: LOW LATENCY INTERFACING WITH LVK WORKING GROUPS

TASK Op-6.5-A(i): MEET THE PLANS AND NEEDS THROUGH O4

Op-6.6 EM advocates

Brief review of the O3 EM follow-up organization. The day-by-day operations of the Low Latency alerts group during observing runs are mostly driven by the astrophysical (and non-) events that are identified by the online searches and the low latency pipeline. A central role in such operations during observing runs are the electromagnetic (EM) advocates. The EM advocates coordinate with colleagues on the Rapid Response Team (RRT), which included assigned individuals from the detector characterization, parameter estimation (PE), search groups, and others to support event characterization and ultimately the communications with the non-LVK community mostly in the form of GCN Circulars. During O3 operations, the EM advocates used a shift-based scheme to assign dates and responsibility of mostly volunteers that were recruited by the Low Latency group chairs and others low latency experts. Google documents and corresponding calendars were used to organize such shifts. The task was benefited by a corresponding operations guide prepared by the group, including the availability of template GCN circulars for the various types of sources and corresponding functionality (e.g., preliminary, initial, retration, and updates). The email, phone and text/SMS communications enabled by GraceDB were also essential for such mode of operation, as were higher latency tasks and communications that were carried out on chat.ligo.org . The process was assisted by experts on a broad range of topics beyond the ones already making part of the RRT; this included experts in search and annotation pipelines, computing, low latency infrastructure etc. On a longer time scale, it also was useful for EM advocates to be in communication with the PE rota team,

O4 EM follow-up organization. As we look into the Low Latency alerts group operations up until March 2023, the EM advocates may be a significant part of the preparations and running during O4. The O4 EM-Advocates organization strongly depends on the global organization made by the Operation division for the Rapid Response Policy for Event Candidates (actually discussed in LIGO-L2100046). Compared to previous observing runs, one needs to transition from "rare events" to handle multiple alerts per day. We will need to keep high purity in our alerts in order to enable multi-messenger discoveries by the follow-up community. So far, a proposal (May 2021) consists on the low-latency vetting (within less four hours) by the full RRT membership of notable candidates (non-BBH triggers with p_terrestrial smaller than 50%). Candidates of less urgency, e.g. BBHs, and updates will be discussed in a separate semi-regular meeting (daily at most). The EM advocate is expected to retain the role for editing and sending GCN circulars. As discussed by the responses of the survey distributed to the astronomers (LIGO-G2000799), more information and more updates for the notables candidates should be prioritized. This includes information for their (updated) sky maps and possible re-evaluation of their false alarm rate.

As EM advocates are a core component of the group's operations, we expect there will contributions from all the collaborations, LIGO, Virgo, and KAGRA and among groups that are active Low Latency alerts. This would also potentially enable time-zone appropriate rotas and minimize middle-of-the-night operations by advocates. To enable the broadest possible participation from the LVK in the EM advocates role, we expect training calls will be organized before and early in O4, during which online tutorials reflecting the expected activities can be performed. We also expect to have opportunities for the EM advocates to perform dry-runs in advance. While details on the role and organization of the EM advocates during O4 are still to be determined, there are several things that are considered and which will make operations easier. For example, having a central location accessible to all and linked off the event portal under GraceDB, that indicates all EM advocate shifts, including detchar, PE and other rotas would facilitate communications and expedite processing. This could also be a location to indicate that an EM rota member has gotten online for a particular event, indicating that others are not necessarily required to join, especially for inconvenient event times for a local timezone. Another proposal (see LVK-G2101363) would be to automate the processing of GCN update circulars. The role of the EM advocate will no longer be an essential one and instead the PE advocate may be able to address such GCN updates as well.

ACTIVITY Op-6.6-A: EM ADVOCATES

TASK Op-6.6-A(i): MEET THE PLANS AND NEEDS THROUGH O4

Op-6.7 Low Latency documentation and outreach

Brief review of the internal documentation of Low Latency

Participation of the low-latency development: for efficient development of the low-latency infrastructure, there should be adequate documentation and code snippet examples provided for the core components in the lines of GWCelery [2] and GraceDB [3]. Additionally, to increase the internal participation, tutorials/bootcamps like the 2019 MIT low-latency bootcamp [4] should be held more often and communicated ahead of time to bring starting graduate students and post-docs up to speed.

Participation of the EM advocate shifts: A documentation for guiding EM advocates during their shifts was developped by the low-latency group[5]. However, in O4, we suggest to have a userguide that includes all the different aspects of the RRT.

Brief review of the public outreach. There were a number of ways the low-latency operations were documented and communicated to the broader scientific public during O3. The first and most comprehensive is the Public Alerts User Guide[6]. Among other things, this guide provides the broader community information on how to receive and interpret the real-time gravitational-wave information. The guide and its update was presented regularly at the LV-EM telecons LV-EM telecons, monthly meetings between gravitational-wave and the interested community.

An IGWN documentation hub for scientific outreach

LIGO/Virgo is as a major international astronomy facility and a public scientific resource, **similar** to astrophysics space missions like the Fermi Gamma-Ray Space Telescope, the Neil Gehrels Swift Observatory, or the James Webb Space Telescope; or analogous to major ground-based observatories like Gemini Observatory, W. M. Keck Observatory, or the European Southern Observatory.

Facilities in this class are expected to provide high quality and up to date public documentation for scientists. Documentation generally includes observing capabilities, operations plans, explanation of data flow, description of data products and data formats, sample code, and tables of sensitivity and sensitivity calculators. We hold JWST JDox[7] as the exemplar that LIGO/Virgo should emulate.

The LIGO/Virgo Public Alerts User Guide that debuted in O3 partly satisfies the need for public end-user documentation. In O4, we will establish a formal editorial team tasked with curating the User Guide, soliciting and collecting contributions from other LIGO/Virgo working groups, and making timely updates.

In O3, there were many different formats and modes of communication for public documentation including wikis, Google docs, the DCC, LaTeX/PDF documents, and email messages. In O4, we will limit the number of different formats and establish a single public IGWN documentation hub. The documentation will have the following characteristics, which are shared with both the present Public Alerts User Guide and the IGWN Computing group's documentation projects:

- Documentation sources will be under public version control in git.ligo.org (or, more appropriately, git.igwn.org).
- Documentation sources will be in one of the two dominant (ca. 2010-2020) markup/Wiki languages, reStructuredText or Markdown.
- Documentation will be typeset using either of the two dominant (ca. 2010-2020) technical documentation systems, Sphinx[8] or MkDocs[9].

[2] https://gwcelery.readthedocs.io
[3] https://gracedb.ligo.org/documentation
[4] https://emfollow.docs.ligo.org/bootcamp/
[5] https://ldas-jobs.ligo.caltech.edu/~emfollow/followup-advocate-guide/introduction.html
[6] https://emfollow.docs.ligo.org/
[7] https://jwst-docs.stsci.edu
[8] https://www.sphinx-doc.org/
[9] https://www.mkdocs.org

- There may be one or several cross-linked documentation sets presented at on a public documentation hub (likely hosted at docs.igwn.org).
- An IGWN Sphinx and MkDocs theme will be provided for all IGWN projects.

During O3, there was a lack of adequate public documentation on interpretation of event significance and classification. We expect that the User Guide can be expanded to support this understanding. Tying into the previous subsection, updated event properties, perhaps with some emphasis from the LVC, may help the selection of "good" events to follow-up, which otherwise can remain opaque from the outside. We note that some of the astronomers are confused by the "Mass-Gap" category or the category change (from Mass-Gap to BBH, or BNS to BHNS). Again, a better communication with clear statement that these indicators are only preliminary results will be beneficial.

Similarly, compared to the current version of the Observing Scenarios Document, there is a need of an extension study to convert predictions with realistic observational constraints and consistent with the open public alert policy. localizations for the candidates with at least one NS present are much larger than expected based on the observing scenarios. While it is understood that the assumptions that are in the current version are appropriate for high significance events, we need a document that accurately reflects the sky areas expected based on the false alarm rates using in public alerts. We can also include the proportions of alerts that are not expected to be observable by observatories due to overlap with the Sun or their location in the Galactic Plane, and report their sky localization areas actually accessible.

In case of update to the alert content, it would likely be appropriate to have some beta testers for the new alerts (not only for gamma-ray and neutrinos but also with some optical groups), and giving the community extra time to make use of it may be appropriate. Training our collaborations to have EM advocates fully capable of assisting across all timezones (such as is done on the EM side), similarly for the PE rota, is also likely to be useful. This aspect was clearly improving by the end of O3.

ACTIVITY Op-6.7-A: LOW LATENCY DOCUMENTATION AND OUTREACH

TASK Op-6.7-A(i): MEET THE PLANS AND NEEDS THROUGH O4

Op-6.8 Low Latency public alert dissemination and external communication

The suggestions detailed above are related to internal performance of the low-latency group. They are both for our internal benefit (i.e. guidance on how we will make LVK operations better) but also to specify actions for a broader impact for the external community.

- General Communications to the Scientific Public (in particular, the follow-up community): During the O3 observational campaign, different services developed and organized by the LVK were used to communicate and transfer information in case of detection of a new GW alert. There were a variety of alert system delivery methods: the GraceDB public page, GCN service with notices and circulars, Twitter, mobile applications such as Chirp. Communication with the external community was managed by regular LV-EM telecons, the LVK user-guide, GWOSC workshops, conference talks or one-to-one exchange of external groups with the spokespersons.

  While channels of communication and updates of our sensitivity and publication plans were regular, we should evaluate if the external community made the correct use of our information. For example, GW190814 event was misinterpreted by various groups as a clear NS-BH collision (see numbers of

GCNs, observations, and publications related to this event). In addition, from the follow-up community, we should determine how the community is determining whether a candidate is astrophysical or non-astrophysical. We require detailed explanations of how pipeline output is aggregated to decide whether a given event is released to the public. **In particular, how various released "parameters" indicate whether an event is real or a false alarm is not well-understood, despite being critical in influencing the decision to engage follow-up observations**.

- Collection of feedback from the external community: a survey organized by the low-latency group to understand the broader community's perspective has been undertaken. The goals were to i.) inspire new ideas for fresh capabilities useful for the broader community and ii.) prioritize plans among the LVK low-latency group looking towards O4 improvements. Therefore, it focused on interactions with the alert system, including its contents and dissemination, and how they may be improved. To this end, the survey document with an associated presentation live here [10]. Decisions about communication and scientific choices should be in some way discussed inside the observational and operation groups with this community effort in mind. Further efforts to engage in ongoing discussions with the community, including through potential future surveys, will be required in the run up to and into O4.
- Things to improve in the future (both between O3 and O4 and during O4):
  - Tutorials for the Public Users: To address the communication issue, we suggest the following improvements: i.) Face-to-face sessions at various scientific conferences and/or workshops. Given the current situation with COVID-19, this is not currently possible, but "zoom" (or similar) sessions can/will be done and by the end of the year this situation may have improved, ii.) LV-EM telecon and communications : we should evaluate how to increase their efficiency and benefit, as they are currently mostly one way (LVK to external community), and iii.) better documentation (e.g. in the "User's Guide"): see section 6. We expect that the LV-EM telecons should continue and perhaps those might be increased to allow more groups to share results on interesting, recent candidates. While it may not be always possible, continued clarification of "candidates" and "events" is important, perhaps with explanation from the LVC when events have been elevated within our community to warranting single object papers or similar, including with expected timelines by the GW collaborations. With the introduction of early warning alerts in O4, we need to establish how multiple early warning (and subsequent full bandwidth alerts) from the same event will be communicated to the public promptly.
  - Alert content and latency: The alert content (public information delivered and science properties of the alert such as FARs, p(astro), EM-Bright numbers, etc.), reliability of the sky maps and latency expectations will not be discussed here as they are under the responsibility of the observational science division. However, we immediately require documentation and methods for collating Early Warning events, in particular for how they will interact with full bandwidth candidates, for the external community.
  - LVK alert service. Given that i.) EM follow-up regions of the sky are large compared to the FoV for many EM-based followup instruments (although this is not the case for certain radio and γ-ray instruments), ii.) tiling for those instruments takes a long time to cover the region determined by the LVK network and iii.) the occurrence rate will be increased by a factor of 10 due to the increased sensitivity in O4, it is important to provide a real vs. not real filtering service to the EM follow-up user so they can make informed decisions with regards to the astrophysical nature of events. As discussed in the general section, we need to explore alert technologies (public and private) beyond GCN (e.g. those developed for VRO/LSST).

[10] https://dcc.ligo.org/LIGO-G2000799

- i.) Targeted Filtering of Notices as a broker: GCN has built-in an 8 parameter filtering capability (easy to add more if needed). These parameters are used to filter out reception by the LVK operators and reception by the public for "cleaner" events to the public. Operations in O3 have given insight as to which parameters to use and what threshold levels to use in O4. This helps to know which triggered events are astrophysical and which are noise. Further, these 8 dimensions of phase space operate independently. It might be such that 2 or 3 parameters can/should be combined to make a more effective multi-dimensional phase filtering capability. This will be studied to see (a) what the optimum settings are for the 1-dimensional filter, and (b) if there is a 2-or-3 dimensional trigger that works better than the 1-dimensional filters. ii.) Filtering on the Circulars: O3 brought to GCN a new group of people with no prior experience of GCN nor of LVC events. Some thought that there were too many Circulars for each event. GCN and TACH are working to create a system to distribute Circulars with criteria (eg. parameters on type of event (LVK events vs GRB events (short and long) vs SGR events). iii.) Contribution to Faster-Broader Brokers: Given that by the time of the start of O4, there will be one to a few large broker systems collecting and distributing astrophysical alerts (i.e. LSST-level-2 brokers, the Transient Name server, AMON, SCiMMA, GCN-Kafka, etc.). LVK alerts will, of course, contribute their streams. This can be done by (a) a direct connection to the given broker, (b) the GCN publisher to the large broker to be combined into their large streams, or possibly (c) using both with GCN as a backup in case of failure of a direct LVK connection to the broker. At a larger scale, the LVK alert service can be connected with broader multi-messenger meta-pipelines/ databases/ communication platforms / multi-observatory schedulers (e.g. AMON, AMPEL, ICARE, STARS,...).

ACTIVITY Op-6.8-A: LOW LATENCY PUBLIC ALERT DISSEMINATION AND EXTERNAL COMMUNICATION

TASK Op-6.8-A(i): MEET THE PLANS AND NEEDS THROUGH O4

7 Computing and Software

Op-7.1 Continue to Deliver (CD)

The success of IGWN data analysis computing rests on a broad and ever-growing base of existing infrastructure and processes which we must maintain, operate, support, and continue to develop to stay current and support new use-cases. In 2021-23 this includes continued delivery of tools, services, and computing for O3 offline computing, and highly-available operations of critical O4 low-latency computing infrastructure for MMA and new discoveries. This also includes estimation and optimization of codes and resources required for O4/O5 science, with an evolving focus on optimizing post-detection parameter estimation (PE) as our detection optimization over the past five years has shown success.

The LIGO, Virgo, and KAGRA collaborations also critically rely on a substantial and growing set of Collaboration Support Services (CSS) to support internal collaboration functions, including integrated identity and access management (e.g., single sign-on, groups, etc.), computing security, and productivity tools and services including mailing-lists, wikis, and remote participation.

In addition, IGWN is committed to sustaining the key deliverable of Open Data for the broader research community through the Gravitational Wave Open Science Center (GWOSC), with an increasingly effective suite of tools, services, and support for the public analysis of IGWN data.

The LIGO Laboratory and KAGRA largely manage internal laboratory computing work necessary for IFO control, upgrades, commissioning, and data acquisition separately from the data analysis and collaboration support computing represented in this IGWN computing plan, and that work, while critical, is not included here. EGO and Virgo do not have as sharp a separation as LIGO and KAGRA between computing for IFO control, commissioning, and data acquisition vs. computing for post-delivery data analysis, and see this work as part of a single holistic computing plan included here.

ACTIVITY Op-7.1-A: MAINTAIN, OPERATE, AND SUPPORT CORE IGWN SERVICES (CD.MOS)

This one enterprise absorbs the *vast majority of the total available computing FTE effort* across LIGO, Virgo, and KAGRA and still is under-resourced.

Op-7.5 enumerates the large set of existing Data Analysis & Collaboration Support Services IGWN currently must maintain, operate, and support.

Maintenance includes updates to OS packages needed to keep a service current with upstream supported OS platforms, and ensuring functioning backups, logging, monitoring, etc.

Operation includes watching and responding to monitoring data, coordinating with system maintenance (restarts etc.), debugging problems, etc.

Support includes interactions with both downstream users and operators/developers of inter-dependent services.

All new software **development** or integration work to adapt to changes in data analysis requirements, add or enhance functionality, etc. should be represented as a distinct goal, project, or task in the other sections of this document. All of those additional sections describe IGWN computing goals, projects and tasks *beyond* basic maintenance, operation, and support.

ACTIVITY Op-7.1-B: LOW-LATENCY COMPUTING (CD.LL)

TASK Op-7.1-B(i): REVIEWS OF ARCHITECTURE/INTERFACE/DESIGN PROPOSALS AND SOFTWARE DEVELOPMENT/OPERATIONS PLANS DEVELOPED AND MANAGED BY THE LOW-LATENCY WORKING GROUP, E.G.

Capacity planning (CPUs, GPUs, data xfer, etc.), including how many and which low-latency pipelines will run in O4.

New/improved infrastructure to support early warning alerts.

Quality of Server (QoS) metrics for low-latency infrastructure and plans for improvements needed to meet them.

Eliminating remaining single points of failure in the end-to-end chain of low-latency infrastructure.

GraceDB enhancements.

TASK Op-7.1-B(ii): ESTABLISH A FLEXIBLE SYSTEM TO REPLAY ARCHIVAL DATA TO SIMULATE NEW DATA.

[CD.LL.REPLAY]

TASK Op-7.1-B(iii): INTEGRATE SERVICE UPTIME MEASUREMENTS ACROSS LOW-LATENCY SYSTEM INTO MONITOR.LIGO.ORG

[CD.LL.MON]

TASK Op-7.1-B(iv): DEVELOP & IMPLEMENT SYSTEM FOR PRIORITIZING LOW-LATENCY JOBS ON HARDWARE RESOURCES THAT ARE ALSO AVAILABLE TO OFFLINE ANALYSES IN ORDER TO INCREASE AVAILABLE LOW-LATENCY COMPUTING CAPACITY WITHOUT DEGRADING OVERALL RESOURCE UTILIZATION.
[CD.LL.PRIO]

TASK Op-7.1-B(v): DEPLOY A SINGLE, CONSISTENT DISTRIBUTED PLATFORM ENABLING HIGH-AVAILABILITY DEPLOYMENTS OF SERVICES NEEDED FOR LOW-LATENCY ALERTS (E.G., KUBERNETES FOR GRACEDB, GWCELERY, ETC.) TO REPLACE MULTIPLE, HETEROGENEOUS PLATFORMS AND SERVERS MANAGED BY DIFFERENT GROUPS.
[CD.LL.HAPLAT]

TASK Op-7.1-B(vi): MIGRATE GRACEDB, GWCELERY, AND LVALERT SERVICES TO THE NEW HIGH-AVAILABILITY LOW-LATENCY PLATFORM.
[CD.LL.HAMIG]

TASK Op-7.1-B(vii): ADD SUPPORT FOR SCITOKENS TO THE GRACEDB SERVER & CLIENT TO SIMPLIFY AND MODERNIZE USER AUTHENTICATION.
[CD.LL.SCIT]

ACTIVITY Op-7.1-C: OPEN SCIENCE (CD.OD)

TASK Op-7.1-C(i): DEVELOP A PROPOSAL TO UNIFY PROPRIETARY AND PUBLIC DATA FORMATS (E.G., DECOMMISSION THE PROPRIETARY IGWF FORMAT FOR O5 AND SWITCH TO AN OPEN STANDARD FORMAT LIKE HDF5), AND UNDERSTAND DEVELOPMENT NECESSARY TO IMPLEMENT IT (E.G., IN FRAMECPP AND FRAMEL).
[CD.OD.DF]

TASK Op-7.1-C(ii): ADD FEATURES TO THE NEW "EVENT PORTAL", INCLUDING A GENERAL QUERY INTERFACE AND ADDITIONAL SUPPORT FOR MULTIMEDIA FILES (E.G. AUDIO).
[CD.OD.EPQ]

TASK Op-7.1-C(iii): UPDATE THE EVENT PORTAL TO SUPPORT STRAIN FILE LOCATIONS SPREAD ACROSS MULTIPLE DIRECTORIES, AND TO SUPPORT A REDESIGN OF THE STRAIN FILES DIRECTORY STRUCTURE.
[CD.OD.EPS]

TASK Op-7.1-C(iv): EXPLORE POSSIBILITIES FOR AUXILIARY CHANNEL RELEASES, POSSIBLY INCLUDING A PROTOTYPE AUXILIARY RELEASE SUPPORTED BY RANA ADHIKARI. THIS MAY INCLUDE A NEW PUBLIC NDS2 SERVER TO MAKE THESE DATA AVAILABLE.
[CD.OD.AUX]

TASK Op-7.1-C(v): UPDATE THE (GWOSC?) DATABASE SCHEMA FOR EVENTS AND CATALOGS TO HANDLE THE NEW NAMING SCHEME - BOTH FOR EVENTS AND CATALOGS - IN DEVELOPMENT BY THE DATA ANALYSIS WORKING GROUPS. THIS MAY INCLUDE SUPPORT FOR EVENTS HAVING MULTIPLE NAMES AND/OR LISTED IN MULTIPLE CATALOGS.
[CD.OD.SCHEMA]

ACTIVITY Op-7.1-D: COLLABORATION DATA DISTRIBUTION, FORMATS & MANAGEMENT (CD.DATA)

Low-latency data:

TASK Op-7.1-D(i): DROP SUPPORT FOR O2 DATA MONITORING TOOL (DMT), FRAMELINK AND VIRGO SPECIFIC (CM) SOLUTIONS AND HARDEN O3 KAFKA SYSTEM TO IMPROVE PERFORMANCE AND RELIABILITY FOR O4.
[CD.DATA.KAFKA]

TASK Op-7.1-D(ii): REPLACE PROPRIETARY KAFKA GLUE WITH CONFLUENT KAFKA
[CD.DATA.CONFKAF]

Bulk data:

TASK Op-7.1-D(iii): MIGRATE FROM LIGO DATA REPLICATOR (LDR) AND VIRGO LEGACY SYSTEMS TO RUCIO TO REDUCE IGWN DEVELOPMENT & SUPPORT BURDEN.
[CD.DATA.RUCIO]

Data Formats:

TASK Op-7.1-D(iv): DEFINE SHORT-TERM ROADMAP FOR THE LIGO_LW DATA FORMAT AND ITS SUPPORT. ADDRESS INCOMPATIBILITIES BETWEEN THE TWO LIGO_LW FORMAT VARIATIONS IN CURRENT USE.
[CD.DATA.LIGOLW]

TASK Op-7.1-D(v): ENHANCE FRAMECPP AND FRAMEL IN ORDER TO IMPROVE MAINTAINABILITY, AND MODULARIZE DATA REPRESENTATION TO SUPPORT ADDITIONAL FORMATS (E.G., HDF5).
Subject to the outcome of a planning exercise on data formats (see MGMT), [CD.DATA.FRAME]

Data Management:

TASK Op-7.1-D(vi): DEVELOP A DB / REGISTRY OF DATA IN THE PERMANENT LIGO, VIRGO, AND KAGRA DATA ARCHIVES, WITH METADATA (MAY USE RUCIO) IN ORDER TO IMPROVE LONG-TERM CURATION AND TRUST IN THE ARCHIVES.
[CD.DATA.REG]

TASK Op-7.1-D(vii): CONSULT WITH THE CBC AND BURST SCIENCE GROUPS, AS WELL AS THE LOW-LATENCY AND OPEN DATA OPERATIONS GROUPS, TO UNDERSTAND THE SCIENTIFIC AND TECHNICAL REQUIREMENTS FOR STORING GW TRIGGERS AT A SCALE INAPPROPRIATELY LARGE FOR GRACEDB, IDENTIFY POSSIBLE SOLUTIONS, AND EVALUATE THEIR COSTS & BENEFITS.
[CD.DATA.GWTRIG]

TASK Op-7.1-D(viii): CONSULT WITH THE CBC AND BURST GROUPS TO UNDERSTAND THE REQUIREMENTS FOR STORING EXTERNAL EM TRIGGERS (GRBS, FRBS, ETC.), IDENTIFY POSSIBLE SOLUTIONS, AND EVALUATE THEIR COSTS & BENEFITS.
[CD.DATA.EXTRIG]

TASK Op-7.1-D(ix): DEVELOP A VERSIONED *DAC dataset* SCHEME THAT MAY INCLUDE DIFFERENT VERSIONS AND CALIBRATIONS OF DATA FROM DIFFERENT IFOS AT DIFFERENT TIMES (E.G., LIGO VS VIRGO).
[CD.DATA.DACDS]

Data Access:

TASK Op-7.1-D(x): ADD SUPPORT FOR SCITOKENS TO THE GWDATAFIND SERVER & CLIENT TO SIMPLIFY AND MODERNIZE USER AUTHENTICATION.
[CD.DATA.SCIT]

ACTIVITY Op-7.1-E: OFFLINE COMPUTING (CD.OFFC)

Workflow management:

TASK Op-7.1-E(i): PROVIDE A GENERIC PYTHON-BASED WORKFLOW MANAGEMENT LIBRARY TO GENERATE AND MANAGE IGWN DAGS, IN COLLABORATION WITH THE HTCONDOR TEAM.
[CD.OFFC.WF]

Resource provisioning:

TASK Op-7.1-E(ii): IMPROVE THE STABILITY AND EFFICIENCY OF IGWN HTCONDOR CLUSTERS BY DEVELOPING AND DEPLOYING MECHANISMS FOR PARTITIONING USER COMMUNITIES INTO GROUPS GRANTED PROGRESSIVELY INCREASING CAPABILITIES, IN COLLABORATION WITH THE HTCONDOR TEAM (E.G., LIMITING THE MAXIMUM SCALE OF WORKFLOWS THAT CAN ACCESS HOME DIRECTORIES DIRECTLY FROM WORKER NODES, LIMITING THE NUMBER OF JOBS THAT CAN RUN CONCURRENTLY, ETC.)
[CD.OFFC.LIMIT]

Improved resource accounting:

TASK Op-7.1-E(iii): PERFORM A REVIEW OF THE O3 USAGE DATA COLLECTION SYSTEM, STORAGE BACKEND, AND REPORTING INTERFACE AND DEVELOP A PLAN TO DELIVER IMPROVEMENTS BEFORE O4 (LINK).
[CD.OFFC.AREVIEW]

TASK Op-7.1-E(iv): MOVE EXISTING INSTANCE OF THE IGWN ACCOUNTING DATABASE AND BACKEND TO A EUROPEAN DATA CENTER IN ORDER TO ADDRESS GDPR ISSUES.
[CD.OFFC.AMOVE]

TASK Op-7.1-E(v): DEVELOP A POLICY AND MORE SUSTAINABLE/AUTOMATED MECHANISM FOR COLLECTING LOCAL ACCOUNTING DATA FOR IGWN DATA ANALYSIS COMPUTING WORKFLOWS SUBMITTED FROM NON-IGWN SUBMIT NODES TO VIRGO CLUSTERS IN EUROPE, WITH AS MUCH AUTOMATION AND AS LOW A LATENCY AS POSSIBLE.
[CD.OFFC.AV]

TASK Op-7.1-E(vi): ALLOW JOBS TO BE TAGGED BY EVENT CANDIDATE OR TARGET PUBLICATION.
[CD.OFFC.ATAG]

TASK Op-7.1-E(vii): PROVIDE A WEB FORM (AND SEMI-AUTOMATED BACKEND) TO REQUEST NEW ACCOUNTING TAGS.
[CD.OFFC.AFORM]

TASK Op-7.1-E(viii): REMOVE LIGO.* AND VIRGO.* PREFIXES FROM ALL ACCOUNTING TAGS FOR JOINT IGWN DATA ANALYSES (OR REPLACE THEM WITH IGWN.* PREFIXES) TO REFLECT THEIR JOINT NATURE.
[CD.OFFC.APF]

Automated workflow/resource prioritization:

TASK Op-7.1-E(ix): WITH DAC, DEVELOP AN UPDATED WORKFLOW/RESOURCE PRIORITIZATION POLICY FOR O4.
[CD.OFFC.DACPRIO]

TASK Op-7.1-E(x): DESIGN AND IMPLEMENT ANY INFRASTRUCTURE NEEDED TO SUPPORT O4 WORKFLOW/RESOURCE PRIORITIZATION.
[CD.OFFC.PRIO]

Simplify and modernize user authentication to IGWN services (x509 to SciTokens):

TASK Op-7.1-E(xi): ADD SUPPORT FOR SCITOKENS TO THE IGWN SEGMENT DATABASE SERVER & CLIENT
[CD.OFFC.SEGST].

Usability:

TASK Op-7.1-E(xii): COMPLETE AND DEPLOY NEW DQSEGDB WEB INTERFACE.
[CD.OFFC.SEGWEB]

ACTIVITY Op-7.1-F: SOFTWARE AND ENVIRONMENTS (CD.SOFT)

We plan to continue the gradual transition from a monolithic LIGO Reference Operating System (RefOS) model to a more flexible container- and/or Conda-based model.

Packaging:

TASK Op-7.1-F(i): SUPPORT IGWN DEVELOPERS WITH CONDA PACKAGING FOR ALL USER SOFTWARE.
[CD.SOFT.SUPCONDA]

TASK Op-7.1-F(ii): REDUCE DEBIAN/RHEL PACKAGING SUPPORT TO 'CRITICAL' AND SYSTEM SERVICES. SUPPORT IGWN DEVELOPERS WITH RPM PACKAGING ONLY AS NECESSARY.
[CD.SOFT.SUPRPM]

TASK Op-7.1-F(iii): SUPPORT IGWN DEVELOPERS WITH MIGRATION TO PYTHON 3.6+ FOR ALL SOFTWARE.
[CD.SOFT.SUPPY]

TASK Op-7.1-F(iv): COMPLETE VIRGO TRANSITION FROM CMT TO CMAKE/MESON AND CONDA.
[CD.SOFT.CMT]

Version control:

TASK Op-7.1-F(v): DELIVER A VERSION CONTROL SOLUTION FOR LARGE DATA PRODUCTS (UNSUITABLE FOR GIT).
[CD.SOFT.BIGVC]

TASK Op-7.1-F(vi): MIGRATE LIGO LAB SVN REPOSITORIES & TRAC SERVICES (OUTSIDE OF CDS) TO GITLAB (E.G., TRAC.LIGO.CALTECH.EDU)
[CD.SOFT.LSVN]

TASK Op-7.1-F(vii): COMPLETE VIRGO TRANSITION FROM SVN TO GITLAB
[CD.SOFT.VSVN]

Software Engineering standards and Change Management (SCCB):

TASK Op-7.1-F(viii): SCCB CHAIR DUTIES (CHAIRING CALLS, COORDINATING EFFORT, REPORTING, ETC.)

- Refine procedure to reduce time-to-approval.
- Ensure SCCB management of all offline data analysis pipelines used in collaboration publications.

[CD.SOFT.SCCBC]

TASK Op-7.1-F(ix): IMPROVE SCCB MANAGEMENT OF LOW-LATENCY ANALYSIS COMPONENTS
[CD.SOFT.SCCBLL]

Automated Testing:

TASK Op-7.1-F(x): DEVELOP POLICIES FOR WHERE AUTOMATED TESTS ARE REQUIRED COUPLED TO THE REVIEW PROCESS.
e.g., for low-latency detection pipelines and infrastructure. [CD.SOFT.TESTPOL]

TASK Op-7.1-F(xi): IDENTIFY NEW & IMPROVED TOOLS, SERVICES, AND POLICIES TO FACILITATE AUTOMATED TESTING AND DEVELOP A 1-3 YEAR PLAN TO DELIVER THEM.
[CD.SOFT.TESTPLAN]

TASK Op-7.1-F(xii): DELIVER A DOCUMENTED, PRODUCTION CAPABILITY TO PERFORM AUTOMATED SYSTEM TESTING (FOR CORRECTNESS, PERFORMANCE, ETC.) OF ENTIRE DATA ANALYSIS WORKFLOWS.
[CD.SOFT.SYSTEST]

TASK Op-7.1-F(xiii): DELIVER A DOCUMENTED, PRODUCTION CAPABILITY TO OFFLOAD RESOURCE-INTENSIVE GITLAB CI JOBS TO HTCONDOR (AND IDEALLY THE IGWN GRID RATHER THAN A LOCAL CONDOR POOL).
[CD.SOFT.IGWNCI]

TASK Op-7.1-F(xiv): PROVIDE SUPPORT & TRAINING TO ASSIST DEVELOPERS IN THE ADOPTION OF AUTOMATED UNIT & SYSTEM TESTING (ANALOGOUS TO LIGO OPTIMIZATION TEAM SUPPORT OF CODE OPTIMIZATION).
[CD.SOFT.TESTSUP]

Complete Transition of Python 2 Code to Python 3 to Ensure Long-Term Supportability:

TASK Op-7.1-F(xv): UPDATE SEGDB SERVER TO PYTHON 3
[CD.SOFT.PYSEG]

TASK Op-7.1-F(xvi): UPDATE LOW-LATENCY (KAFKA) DATA DISTRIBUTION TO PYTHON 3
[CD.SOFT.PYKAF]

TASK Op-7.1-F(xvii): GUARDIAN/METATRON AUTOMATION IN VIRGO
[CD.SOFT.PYVDAQ]

ACTIVITY Op-7.1-G: DETECTOR CHARACTERIZATION (CD.DETCHAR)

Enhancements to Channel Information System

TASK Op-7.1-G(i): UPDATE USER INTERFACE
[CD.MOS-DA]

TASK Op-7.1-G(ii): AUTOMATE DATABASE UPDATES FROM NEW DATA
[CD.MOS-DA]

Enhancements to Data Quality Reports (DQR)

TASK Op-7.1-G(iii): IMPROVE LOW LATENCY RESPONSE
[CD.MOS-DA.DQRLL]

TASK Op-7.1-G(iv): PROVIDE AN EASIER PATH TO ADD NEW ANALYSIS TASKS
[CD.MOS-DA.DQRNA]

TASK Op-7.1-G(v): LEVERAGE SUMMARY PAGE INFRASTRUCTURE TO IMPROVE RESPONSIVENESS AND EASE OF NAVIGATION.
[CD.MOS-DA.DQRSP]

ACTIVITY Op-7.1-H: COMPUTING OPTIMIZATION (CD.OPT)

TASK Op-7.1-H(i): O4 AND O5 CPU & GPU DEMAND ESTIMATION
[CD.OPT.DEMAND]

TASK Op-7.1-H(ii): O4 AND O5 CPU & GPU SUPPLY ESTIMATION
[CD.OPT.SUPPLY]

TASK Op-7.1-H(iii): TRADE STUDIES FOR LIGO LAB LOW-LATENCY O4 HARDWARE PROCUREMENT
[CD.OPT.HW]

TASK Op-7.1-H(iv): VIRGO LOW-LATENCY COMPUTING REDEPLOYMENT
[CD.OPT.VLLR]

TASK Op-7.1-H(v): OPERATIONAL OPTIMIZATION DEVELOPMENT (E.G., TESTING & DEPLOYING COMPLETED OPTIMIZATIONS INTO PRODUCTION)
[CD.OPT.OPS]

TASK Op-7.1-H(vi): O4 DATA ANALYSIS (DA) PIPELINE OPTIMIZATION (FOCUSED ON PE)
[CD.OPT.O4]

TASK Op-7.1-H(vii): O5+ DATA ANALYSIS (DA) PIPELINE OPTIMIZATION
[CD.OPT.O5]

TASK Op-7.1-H(viii): ENABLING OF CPU & GPU OPTIMIZATION EFFORTS ACROSS IGWN (VIA HARDWARE TEST STANDS, TOOLS, CONSULTING, ETC.)
[CD.OPT.ENABLE]

ACTIVITY Op-7.1-I: COLLABORATION SUPPORT (CD.CSS)

Identity and Access Management (IAM):

TASK Op-7.1-I(i): ARCHITECT IAM SYSTEM AND MANAGE EFFORTS AND PRIORITIES
[CD.CSS.IAMMANAGE]

TASK Op-7.1-I(ii): TRANSITION FROM MYLIGO AND GW-ASTRONOMY TO IGWN COMANAGE FOR IDENTITY MANAGEMENT

- Design org charts, identifiers, enrollments, etc for IGWN entities [CD.CSS.COmArch]
- Write cakePHP plugins for custom logic (Efforts, Authorship, Council, etc) [CD.CSS.COmPI]
- Migrate existing data into COmanage

TASK Op-7.1-I(iii): TRANSITION FROM LIGO IDP TO SAML PROXY FOR WEB AUTHN
[CD.CSS.SATOSA]

TASK Op-7.1-I(iv): DEVELOP A PLAN TO MIGRATE THE LIGO IDP FROM AN IGWN-MANAGED SERVICE TO A LIGO LAB-MANAGED SERVICE
[SC.CSS.LABIDP]

TASK Op-7.1-I(v): PROVIDE EXPERT SUPPORT FOR IGWN DEVELOPERS ADDING SCITOKENS FOR DATA SERVICES
CVMFS, GraceDB, DQSegDB, gwdatafind, ... [CD.CSS.STSUP]

TASK Op-7.1-I(vi): 2-FACTOR AUTHENTICATION FOR WEB ACCESS TO LIGO.ORG
[CD.CSS.MFASAML]

TASK Op-7.1-I(vii): 2-FACTOR AUTHENTICATION FOR SHELL ACCESS TO LDG
[CD.CSS.MFASSH]

TASK Op-7.1-I(viii): DEVELOP MECHANISM/S & POLICY FOR CYBERINFRASTRUCTURE COLLABORATOR ACCESS
[CD.CSS.CICOLAB]

TASK Op-7.1-I(ix): REENGINEER IGWN.ORG MAILING-LIST SOFTWARE TO REDUCE ONGOING SUPPORT BURDEN
[CD.CSS.LISTAUTO]

TASK Op-7.1-I(x): MIGRATE SYMPA.LIGO.ORG FROM LLO HARDWARE TO HA PLATFORM
[CD.CSS.LISTHA]

TASK Op-7.1-I(xi): DEVELOP A LIGO-INDIA IAM STACK FOR INTERNAL WORK THAT MUST BE HOUSED AND MANAGED ENTIRELY ON INDIAN SOIL.
[CD.CSS.LIIAM]

Computing Security:

TASK Op-7.1-I(xii): IGWN SECURITY MANAGEMENT DUTIES (POLICIES, REPORTING, INCIDENT RESPONSE)
[CD.CSS.SECL]

TASK Op-7.1-I(xiii): LSC SECURITY MANAGEMENT DUTIES (POLICIES, REPORTING, INCIDENT RESPONSE)
[CD.CSS.LSECL]

TASK Op-7.1-I(xiv): VIRGO SECURITY MANAGEMENT DUTIES (POLICIES, REPORTING, INCIDENT RESPONSE)
[CD.CSS.VSECL]

TASK Op-7.1-I(xv): KAGRA SECURITY MANAGEMENT DUTIES (POLICIES, REPORTING, INCIDENT RESPONSE)
[CD.CSS.KSECL]

TASK Op-7.1-I(xvi): LSC SECURITY COMMITTEE DUTIES

- Performing security reviews of critical IGWN computing infrastructure and proposals.
- Defining and updating security-related policies of IGWN and its individual member collaborations.

[CD.CSS.SECCOM]

TASK Op-7.1-I(xvii): IGWN SECURITY OPERATIONS (LOGGING, MONITORING, INTRUSION DETECTION, TECHNICAL INCIDENT RESPONSE, ETC.)
[CD.CSS.SECOPS]

TASK Op-7.1-I(xviii): LSC SECURITY OPERATIONS (LOGGING, MONITORING, INTRUSION DETECTION, TECHNICAL INCIDENT RESPONSE, ETC.)
[CD.CSS.LSECOPS]

TASK Op-7.1-I(xix): VIRGO SECURITY OPERATIONS (LOGGING, MONITORING, INTRUSION DETECTION, TECHNICAL INCIDENT RESPONSE, ETC.)
[CD.CSS.VSECOPS]

TASK Op-7.1-I(xx): KAGRA SECURITY OPERATIONS (LOGGING, MONITORING, INTRUSION DETECTION, TECHNICAL INCIDENT RESPONSE, ETC.)
[CD.CSS.KSECOPS]

TASK Op-7.1-I(xxi): MIGRATE UNSUPPORTABLE/INSECURE LEGACY LAAC THESIS DB TO A NEW, SUSTAINABLE PLATFORM.
[CD.CSS.LAACDB]

TASK Op-7.1-I(xxii): CLIMATE CHANGE & REMOTE PARTICIPATION PLANNING
[CD.CSS.REMPLAN]

TASK Op-7.1-I(xxiii): COMPUTING REVIEW COMMITTEE CHAIRPERSON DUTIES
[CD.CSS.REVLEAD]

TASK Op-7.1-I(xxiv): COMPUTING REVIEW COMMITTEE DUTIES

- Performing reviews of critical IGWN computing infrastructure and proposals.
- Defining and updating review processes and policies.

[CD.CSS.REVCOM]

TASK Op-7.1-I(xxv): DEVELOP NEW LSC MOU SOFTWARE SYSTEM
[CD.CSS.MOUDEV]

TASK Op-7.1-I(xxvi): COMPUTING TRAINING, DOCUMENTATION, SUPPORT

TASK Op-7.1-I(xxvii): IGWN HELP DESK SUPPORT
[CD.CSS.HDESK]

TASK Op-7.1-I(xxviii): IAM GENERAL HELP DESK SUPPORT
[CD.CSS.IAMHD]

TASK Op-7.1-I(xxix): IGWN COMPUTING USER GUIDE - TOP-LEVEL
including maintaining the top-level web pages of the computing guide, identifying and coordinating the writing & maintenance of all other computing documentation, amd organizing tutorial workshops, trainings, etc. [CD.SUP.DOCUSER]

TASK Op-7.1-I(xxx): IGWN GRID COMPUTING (OSG) USER TUTORIAL
[CD.SUP.DOCIGWN]

TASK Op-7.1-I(xxxi): WORKING EXAMPLE IGWN DA WORKFLOW LIBRARY
illustrating IGWN/OSG/Condor best practices and modern features (e.g., workflow construction via python library, dataflow jobs, checkpointing, input/intermediate/output data discovery/transfer/archiving, gwdata URIs, etc.) [CD.SUP.WFLIB]

TASK Op-7.1-I(xxxii): LSC/VIRGO/KAGRA/IGWN ONBOARDING GUIDE/S - COMPUTING SECTION/S
[CD.SUP.ONBDOC]

TASK Op-7.1-I(xxxiii): IGWN COMPUTING SECURITY POLICIES AND BEST PRACTICES TUTORIAL
[CD.SUP.SECTUT]

TASK Op-7.1-I(xxxiv): IGWN RESOURCE PROVIDER DEPLOYMENT & ADMIN GUIDE
[CD.SUP.CEDOC]

TASK Op-7.1-I(xxxv): SHIBBOLETH SP ADMIN GUIDE
[CD.SUP.SPDOC]

ACTIVITY Op-7.1-J: MANAGEMENT (CD.MGMT)

TASK Op-7.1-J(i): LIGO COMPSOFT CHAIR DUTIES
chairing calls, coordinating effort, reporting, etc.
[CD.MGMT.LIGO]

TASK Op-7.1-J(ii): VIRGO COMPSOFT CHAIR DUTIES
chairing calls, coordinating effort, reporting, etc.
[CD.MGMT.VIRGO]

TASK Op-7.1-J(iii): KAGRA COMPSOFT CHAIR DUTIES
chairing calls, coordinating effort, reporting, etc.
[CD.MGMT.KAGRA]

Op-7.2 Continue Shift to Shared Computing Model (SC)

LIGO and VIRGO have agreed on a strategic shift from predominantly dedicated computing environments (including the homogeneous LIGO Data Grid and more heterogeneous native environments of Virgo computing centers such as CNAF and IN2P3) to a joint Lab/LSC/VIRGO (and soon KAGRA) computing environment based on the Open Science Grid platform, with shared responsibility for provisioning computing resources & computing FTEs.

Support for shared computing resources is critical to exploit resources outside of the LIGO Laboratory and LIGO Data Grid (LDG), including international GW partners (e.g., EGO/Virgo, KAGRA, LIGO-India, KISTI, OzGrav), the Open Science Grid, HPC centers (e.g., XSEDE, Blue Waters, DOE systems), GPU systems, LSC groups' institutional resources (Cardiff, Syracuse, Georgia Tech), and public and commercial clouds (e.g., via OSG or HTCondor Annex).

A unified computing environment will also make more efficient use of the collective computing effort of each Collaboration, and enable more efficient development and operation of data analysis software by Collaboration scientists.

This shift will rely on increasing communication and integration on computing with the wider Physics, Astronomy, and Computer Science communities, e.g., the Open Science Grid (OSG), the HTCondor Project and Center for High-Throughput Computing (CHTC), High-Energy Physics Software Foundation (HSF), Scalable CyberInfrastructure to support Multi-messenger Astrophysics (SciMMA), etc.

A unified environment for dedicated and shared computing resources also provides breathing room for unexpected increases in need, human mistakes, surprises by nature, and error bars on initial demand estimates.

ACTIVITY Op-7.2-A: DATA DISTRIBUTION, FORMATS & MANAGEMENT (SC.DATA)

TASK Op-7.2-A(i): FINALIZE DATA DISTRIBUTION ARCHITECTURE (WHO RUNS XROOTD ORIGINS, CVMFS STRATUM-0S, CACHES) INCLUDING DIRECTORY LAYOUT TO IMPROVE FUTURE PERFORMANCE AND AVAILABILITY
[CD.SC.DATA.STCA].

TASK Op-7.2-A(ii): DEPLOY STASHCACHE ORIGIN FOR PROPRIETARY DATA MANAGED BY EACH COLLABORATION
[CD.SC.DATA.STCA]

TASK Op-7.2-A(iii): ADD SUPPORT FOR SCITOKENS TO AUTHENTICATED CVMFS FOR PROPRIETARY DATA TO SIMPLIFY AND MODERNIZE USER AUTHENTICATION.
[SC.DATA.CVMFSSciT]

TASK Op-7.2-A(iv): IN CONJUNCTION WITH DETCHAR AND SCIENTIFIC WORKING GROUPS (E.G., CBC), ESTABLISH THE SCIENTIFIC AND TECHNICAL REQUIREMENTS FOR IMPROVED ACCESS TO RAW DATA FOR OFFLINE ML INVESTIGATIONS, AND ADD THE RESULTING TASKS & DELIVERABLES TO THIS COMPUTING PLAN.
[SC.DATA.RAW]

TASK Op-7.2-A(v): PROVIDE STANDARD WAY TO PUBLISH, DISCOVER, AND ACCESS DERIVED/INTERMEDIATE DATA PRODUCTS BETWEEN DISTRIBUTED WORKFLOWS (FOR SOME PIPELINES THIS IS STILL A SERIOUS HINDRANCE FOR IGWN ADOPTION)

- Technical implementation
- Policy and administrative procedure
- Standardisation of intermediate data format
- Define and implement a derived data cataloguing and bookkeeping solution

[SC.DATA.INTDATAMP, SC.DATA.INTDATAMI]

TASK Op-7.2-A(vi): DEVELOP A GWDATA:// FILE TRANSFER PLUGIN

Improve the portability and scalability of data analysis workflows by delivering a gwdata:// URI enabling HTCondor IGWN workflows to specify GW input data without filenames (in the same manner as gw_data_find) so that the computing system can make smarter data transfer decisions at runtime.

ACTIVITY Op-7.2-B: OFFLINE COMPUTING (SC.OFF)

Port workflows to IGWN: There are likely to be no future LIGO Lab purchases for dedicated offline (non-low-latency) computing clusters; future supply will be increasingly provided via shared resources, so workflows will need to be adapted to use distributed ("grid") IGWN computing resources.

TASK Op-7.2-B(i): SUPPORT MIGRATION OF ALL EXISTING NON-LDG (VIRGO) PIPELINES TO USE STANDARD IGWN HTCONDOR INTERFACES AND ENVIRONMENTS.

[SC.OFF.PORTVIRGO]

TASK Op-7.2-B(ii): SUPPORT MIGRATION OF ADDITIONAL LDG PIPELINES TO USE STANDARD IGWN HTCONDOR INTERFACES AND ENVIRONMENTS.

[SC.OFF.PORTLDG]

TASK Op-7.2-B(iii): DEVELOP A "BACKFILL" QUEUE AND CAPABILITY FOR IGWN RESOURCES.

Resource provisioning:

TASK Op-7.2-B(iv): INTEGRATE LIGO CLUSTERS WITH STANDARD IGWN SERVICES AND INTERFACES

[SC.OFF.CLUSTLIGO]

TASK Op-7.2-B(v): INTEGRATE VIRGO CLUSTERS WITH STANDARD IGWN SERVICES AND INTERFACES

[SC.OFF.CLUSTVIRGO]

TASK Op-7.2-B(vi): INTEGRATE KAGRA CLUSTERS WITH STANDARD IGWN SERVICES AND INTERFACES

[SC.OFF.CLUSTKAGRA]

TASK Op-7.2-B(vii): DEPLOY VIRGO-MANAGED HTCONDOR SUBMIT NODES IN EUROPE (STARTING WITH NIKHEF).

[SC.OFF.SUBMITVIRGO]

TASK Op-7.2-B(viii): DEPLOY IGWN SUBMIT NODES IN EUROPE

TASK Op-7.2-B(ix): ESTABLISH AND STAFF A DISTRIBUTED IGWN COMPUTING OPERATIONS & SUPPORT TEAM AND PROCESS.

TASK Op-7.2-B(x): COMPLETE THE PROCESS OF REVIEWING AND ADAPTING THE DOCUMENTATION (BOTH FOR USERS AND SUPPORT TEAM)

ACTIVITY Op-7.2-C: SOFTWARE AND ENVIRONMENTS (SC.SOFT)

Virgo specific enablers:

TASK Op-7.2-C(i): TRANSITION FROM CMT TO CMAKE/MESON FOR SOFTWARE BUILD
[CD.SOFT.CMT]

TASK Op-7.2-C(ii): TRANSITION TO CONDA FOR SOFTWARE ENVIRONMENT DEFINITION
[CD.SOFT.CMT]

Packaging:

TASK Op-7.2-C(iii): DELIVER A SOLUTION ENABLING DATA ANALYSTS TO QUICKLY DEPLOY DEVELOPMENT CONTAINERS OR CODE TO CVMFS TO ENABLE AN EFFICIENT WORKFLOW TEST/DEBUG CYCLE ON IGWN.
[SC.SOFT.CONTCI]

ACTIVITY Op-7.2-D: COMPUTING OPTIMIZATION (SC.OPT)

TASK Op-7.2-D(i): IMPROVE AVAILABILITY & USABILITY OF SHARED/EXTERNAL COMPUTING RESOURCES VIA IGWN.

ACTIVITY Op-7.2-E: VIRGO/KAGRA SUPPORT FOR IGWN COMPUTING INFRASTRUCTURE (SC.VK)

Transition some LSC-managed services to Virgo primary responsibility:

TASK Op-7.2-E(i): (POSSIBLE) MIGRATE GITLAB (GIT.LIGO.ORG) TO VIRGO PRIMARY RESPONSIBILITY
[SC.VK.GITLAB]

TASK Op-7.2-E(ii): (POSSIBLE) MIGRATE PRODUCTION SEGDB SERVICES TO VIRGO PRIMARY RESPONSIBILITY
[SC.VK.SEGDB]

TASK Op-7.2-E(iii): MIGRATE MANAGEMENT OF SOME LL ALERT SYSTEM (GWCELERY + LVALERT + GRACEDB) TIERS IN SEPARATE DEPLOYMENTS TO VIRGO PRIMARY RESPONSIBILITY
[SC.VK.LLMGMT]

TASK Op-7.2-E(iv): NEGOTIATE KAGRA CONTRIBUTION TO CILOGON CONTRACT (CURRENTLY LAB+LSC+VIRGO) [SC.VK.CILOGONCONTRACT]

ACTIVITY Op-7.2-F: COLLABORATION SUPPORT (SC.CSS)

TASK Op-7.2-F(i): DHTC (AKA OSG, "GRID") USER SUPPORT
[SC.CSS.USUP]

Identity and Access Management (IAM):

TASK Op-7.2-F(ii): INTEGRATE EVOLVING EUROPEAN IAM INFRASTRUCTURE WITH IGWN SHIBBOLETH DEVELOPMENT TO ENABLE INTEGRATION OF IGWN AND EXTERNAL IAM INFRASTRUCTURES (INITIALLY TO SUPPORT IGWN HTC SUBMIT NODE AT NIKHEF).
[CD.CSS.SHIBEXT]

IGWN naming migration:

TASK Op-7.2-F(iii): INVENTORY EXISTING `*.LIGO.ORG` SERVICES TO IDENTIFY THOSE TO MIGRATE TO `*.IGWN.ORG`.
[SC.CSS.IGWNDOM]

TASK Op-7.2-F(iv): CONSOLIDATING PREVIOUSLY-INDEPENDENT LIGO, VIRGO, AND KAGRA COMPUTING INFRASTRUCTURE AND PROCESSES.

TASK Op-7.2-F(v): `LIGO.ORG` → `IGWN.ORG` FOR DNS WHERE APPROPRIATE
[SC.CSS.DNS]

TASK Op-7.2-F(vi): DEFINE PROCEDURE TO OBTAIN DUAL-DOMAIN HOST CERTIFICATES (IGWN.ORG AND HOSTING SITE DOMAIN) TO ALLOW FOR REVERSE-DNS RESOLUTION
[SC.CSS.CERTIFICATE]

Op-7.3 Define Joint IGWN WBS (WBS)

We are developing a detailed, joint LIGO+Virgo computing WBS. The European Gravitational-Wave Observatory (EGO) defined a Virgo-specific Computing WBS in 2019 which will feed into this joint IGWN WBS.

Lessons Learned during O3:

- ***The LSC, Virgo and KAGRA do not collectively have enough full-time professional computing staff to continue normal operations at an adequate level of service.***
- We need to outsource non-GW-specific services where possible to focus IGWN computing effort on GW-specific computing.
- We need to reduce or drop support for services that aren't high priority.
- We know that we do **not** have enough effort to deliver all the work in this Computing Plan for O4.

ACTIVITY Op-7.3-A: ESTABLISH INITIAL WBS (WBS.INIT)

Inventory IGWN computing projects, tools, and services, integrating LIGO Lab, LDG, LSC MOU, Virgo WBS, and KAGRA work-plans as appropriate.

TASK Op-7.3-A(i): INTEGRATE ADVIRGO COMPUTING AND DATA PROCESSING INFRASTRUCTURE WBS (DRAFT 2020-21 VIRGO WBS).
[WBS.INIT.VWBS]

TASK Op-7.3-A(ii): INTEGRATE THE KAGRA COMPUTING WHITE PAPER INTO IGWN DOCUMENT.
[WBS.INIT.KWP]

ACTIVITY Op-7.3-B: IDENTIFY RISKS AND OPPORTUNITIES TO IMPROVE (WBS.ROI)

TASK Op-7.3-B(i): PRIORITIZE IGWN COMPUTING PROJECTS, TOOLS, AND SERVICES ACCORDING TO SCIENCE IMPACT, AND IDENTIFY & COMMUNICATE SCIENCE COSTS AND RISKS OF ALLOCATING INSUFFICIENT EFFORT.

TASK Op-7.3-B(ii): PROPOSE REDUCTIONS OR DROPPED SUPPORT FOR COMPUTING INFRASTRUCTURE THAT REQUIRES THE MOST FTE EFFORT RELATIVE TO ITS SCIENCE IMPACT.

TASK Op-7.3-B(iii): IDENTIFY OPPORTUNITIES TO OUTSOURCE NON-GW-SPECIFIC SERVICES, TO ADOPT OR SHARE COMMON CYBERINFRASTRUCTURE WITH OTHER SCIENCE COLLABORATIONS.

TASK Op-7.3-B(iv): IDENTIFY OPPORTUNITIES TO ACTIVELY SOLICIT ADDITIONAL INFRAOPS EFFORT TOWARDS COMPUTING.

TASK Op-7.3-B(v): IDENTIFY OPPORTUNITIES TO PROVIDE COLLABORATION-LEVEL FUNDING FOR COMPUTING INFRASTRUCTURE OR SERVICES.

The outcome of this process may result in substantial changes to the O4 plans in this document.

Op-7.4 Improve and Integrate IGWN Computing Processes & Management (MGMT)

ACTIVITY Op-7.4-A: MANAGEMENT & FUTURE PLANNING [MGMT.MGMT]

TASK Op-7.4-A(i): DEVELOP A COMPUTING "RISK REGISTRY" TO BE COMMUNICATED TO AND ACCEPTED BY THE COLLABORATIONS, ALONG WITH A "LIENS LIST" OF DEFERRED PROJECTS.
[MGMT.MGMT.RISKREG]

TASK Op-7.4-A(ii): DEFINE A LIGHTWEIGHT PROCESS FOR PROPOSING, DESIGNING, AND DEPLOYING NEW IGWN INFRASTRUCTURE
(including arch, design, and security reviews, plans for ongoing maintenance, operations, and support, risk assessment, etc.)
[MGMT.MGMT.COMPPROP]

TASK Op-7.4-A(iii): DEFINE AN UPDATED IGWN-WIDE DATA ANALYSIS SOFTWARE PRACTICES POLICY BASED ON LIGO-T040245.
[MGMT.MGMT.DASPP]

TASK Op-7.4-A(iv): LONG-TERM DATA FORMAT (GWF, HDF5, LIGO_LW, ETC.) PLANNING
[MGMT.MGMT.DATAFMT]

TASK Op-7.4-A(v): ESTABLISH LIGHTWEIGHT MOU WITH EXTERNAL COMPUTING COLLABORATORS, E.G., HTCONDOR, OSG, AND PEGASUS TO CLARIFY EXPECTATIONS AND MANAGE RISK.
[MGMT.MGMT.EXTMOU]

TASK Op-7.4-A(vi): DEVELOP A CHARGE FOR LSC, VIRGO, AND KAGRA *Security Liaisons*, REPORTING TO THE COMPUTING COMMITTEE OF EACH COLLABORATION, AND FILL THE ROLES.
[MGMT.MGMT.SECLIAISON]

TASK Op-7.4-A(vii): LIAISON WITH LOW-LATENCY OPERATIONS WORKING GROUP

- Gather requirements for a) where low-latency data will be delivered and b) latency goals for each location to inform the requirements on the Kafka data-distribution system.

[MGMT.MGMT.LOWLATLIAISON]

TASK Op-7.4-A(viii): DATA ANALYSIS OPTIMIZATION PLANNING (INCLUDING COLLABORATING WITH HEP AND MMA COMMUNITIES, PRIORITIZING OPTIMIZATION OPPORTUNITIES, ETC.)
[MGMT.MGMT.OPTIMIZATION]

TASK Op-7.4-A(ix): SEEK EXTERNAL FUNDING OPPORTUNITIES, E.G., NSF CYBER INFRASTRUCTURE SOLICITATIONS.
[MGMT.MGMT.FUNDING]

TASK Op-7.4-A(x): ESTABLISH A DATABASE IDENTIFYING COMPUTING PROJECTS AND RESPONSIBILITIES AVAILABLE TO LSC, VIRGO, AND KAGRA GROUPS FOR INFRAOPS CONTRIBUTIONS.
[MGMT.MGMT.PROJDB]

TASK Op-7.4-A(xi): OPERATE AND MAINTAIN A DATABASE IDENTIFYING COMPUTING PROJECTS AND RESPONSIBILITIES AVAILABLE TO LSC, VIRGO, AND KAGRA GROUPS FOR INFRAOPS CONTRIBUTIONS.
[MGMT.MGMT.PROJDBOPS]

Maintain & manage important collaborations and contracts with outside computing experts via formal liaisons and regular updates:

TASK Op-7.4-A(xii): MAINTAIN & MANAGE IGWN COLLABORATION WITH PATH (HTCONDOR AND OPEN SCIENCE GRID) VIA A LIAISON
[MGMT.MGMT.EXTLIASONPATH]

TASK Op-7.4-A(xiii): MAINTAIN & MANAGE IGWN COLLABORATION WITH CILOGON VIA A LIAISON
[MGMT.MGMT.EXTLIAISONCILOGON]

TASK Op-7.4-A(xiv): MAINTAIN & MANAGE IGWN COLLABORATION WITH INTERNET2 VIA A LIAISON

TASK Op-7.4-A(xv): MAINTAIN & MANAGE IGWN COLLABORATION WITH IAM VIA A LIAISON

TASK Op-7.4-A(xvi): MAINTAIN & MANAGE IGWN COLLABORATION WITH XSEDE VIA A LIAISON

TASK Op-7.4-A(xvii): MAINTAIN & MANAGE IGWN COLLABORATION WITH EGI VIA A LIAISON

TASK Op-7.4-A(xviii): MAINTAIN & MANAGE IGWN COLLABORATION WITH EOSC VIA A LIAISON

TASK Op-7.4-A(xix): MAINTAIN & MANAGE IGWN COLLABORATION WITH WLCG VIA A LIAISON

TASK Op-7.4-A(xx): MAINTAIN & MANAGE IGWN COLLABORATIONS WITH THE SCITOKENS TEAM VIA A LIAISON.
[MGMT.MGMT.EXTLIAISONSciTokens]

TASK Op-7.4-A(xxi): IDENTIFY & EXPLORE THE SCIENCE POTENTIAL OF NEW COMPUTING MODELS, ARCHITECTURES, TECHNOLOGY, TOOLS, OPTIMIZATIONS, ETC. FOR GW COMPUTING OVER A 2-5 YEAR FUTURE TIMEFRAME.

TASK Op-7.4-A(xxii): IDENTIFY OPPORTUNITIES FOR PUBLICATION OF IGWN COMPUTING RESULTS.

TASK Op-7.4-A(xxiii): DEVELOP A STRATEGY FOR CLOUD SERVICE PLATFORMS.

TASK Op-7.4-A(xxiv): IDENTIFY CONTRIBUTIONS TO IGWN COMPUTING OPERATIONS AND/OR DEVELOPMENT THAT KAGRA CAN PROVIDE, AND WORK WITH KAGRA TO PLAN TO DELIVER THEM
[MGMT.KCOMP].

Update, create, synchronize, and/or merge important LSC, Virgo, KAGRA, and IGWN computing Policies and Procedure documents:

TASK Op-7.4-A(xxv): COMPUTING ACKNOWLEDGEMENTS POLICY
[MGMT.POLICY.ACK]

TASK Op-7.4-A(xxvi): OPEN SOURCE LICENSING POLICY
[MGMT.POLICY.OPEN]

TASK Op-7.4-A(xxvii): TIER-N DOCUMENT
[MGMT.POLICY.TIERN]

TASK Op-7.4-A(xxviii): LIGO.ORG CREDENTIAL ACCEPTABLE USE POLICY
[MGMT.POLICY.CREDAUP]

TASK Op-7.4-A(xxix): PERSONAL INFORMATION POLICY
(link)
[MGMT.POLICY.PIP]

TASK Op-7.4-A(xxx): DATA GRID ACCEPTABLE USE POLICIES (AUPS)
(link)
[MGMT.POLICY.LDGAUP]

TASK Op-7.4-A(xxxi): IGWN COMPUTING SERVICE REQUIREMENTS POLICY
including SLA; code; release processes (e.g., SCCB); developer, admin and user documentation; maintenance & operations and support expectations; etc.
[MGMT.POLICY.SERVICE]

TASK Op-7.4-A(xxxii): IGWN COMPUTING TOOL REQUIREMENTS POLICY
including code; release processes (e.g., SCCB); developer & user documentation; maintenance and support expectations; etc.
[MGMT.POLICY.TOOL]

TASK Op-7.4-A(xxxiii): LIGO DATA MANAGEMENT PLAN (DMP) (LINK)
[MGMT.POLICY.LDMP]

TASK Op-7.4-A(xxxiv): VIRGO DATA MANAGEMENT PLAN (DMP) (LINK)
[MGMT.POLICY.VDMP]

Op-7.5 Appendix A: Existing Data Analysis and Collaboration Support Tools & Services (CD.MOS)

ACTIVITY Op-7.5-A: CONTINUED MAINTENANCE, OPERATION, AND SUPPORT OF EXISTING GW DATA ANALYSIS TOOLS & SERVICES (CD.MOS-DA)

TASK Op-7.5-A(i): DQSEGDB SERVICE (PRODUCTION, DEVELOPMENT, AND TESTING INSTANCES)
[CD.MOS-DA.SDBS]

TASK Op-7.5-A(ii): DQSEGDB CLIENT
[CD.MOS-DA.SDBC]

TASK Op-7.5-A(iii): LALSUITE
[CS.MOS-DA.LALSMA, CS.MOS-DA.LALSS]

TASK Op-7.5-A(iv): GRACEDB SERVICE (PRODUCTION, DEVELOPMENT, AND TESTING INSTANCES)
[CD.MOS-DA.GdbSMO,CD.MOS-DA.GdbSu]

TASK Op-7.5-A(v): GRACEDB CLIENT
[CD.MOS-DA.GdbCMO,CD.MOS-DA.GdbCSu]

TASK Op-7.5-A(vi): LVALERT SERVER
[CD.MOS-DA.LVASMO,CD.MOS-DA.LVASSu]

TASK Op-7.5-A(vii): LVALERT CLIENT
[CD.MOS-DA.LVACMO,CD.MOS-DA.LVACSu]

TASK Op-7.5-A(viii): LIGO DMT SERVICE

TASK Op-7.5-A(ix): KAGRA DMT SERVICE

TASK Op-7.5-A(x): LIGO AGGREGATED H(T) GENERATION SERVICE

TASK Op-7.5-A(xi): LIGO BULK DATA DISTRIBUTION (LEGACY / LDR)
[CD.MOS-DA.LBDDL]

TASK Op-7.5-A(xii): BULK DATA DISTRIBUTION (RUCIO)
[CD.MOS-DA.BDD]

TASK Op-7.5-A(xiii): LIGO-INDIA BULK DATA DISTRIBUTION

TASK Op-7.5-A(xiv): VIRGO H(T) AGGREGATION & DELIVERY (LEGACY / LDR)

TASK Op-7.5-A(xv): KAGRA H(T) AGGREGATION & DELIVERY

TASK Op-7.5-A(xvi): KAGRA BULK DATA AGGREGATION & DELIVERY

TASK Op-7.5-A(xvii): PUBLIC GWOSC WEBSITE SERVICE

TASK Op-7.5-A(xviii): PUBLIC GWOSC TOOLS

TASK Op-7.5-A(xix): GWOSC CVMFS ORIGIN SERVICE FOR OPEN DATA

TASK Op-7.5-A(xx): WORKFLOW MANAGEMENT TOOLKIT (AKA PIPELINE.PY)

TASK Op-7.5-A(xxi): GWDATAFIND CLIENT TOOL
[CD.MOS-DA.GWDFC]

TASK Op-7.5-A(xxii): CIT GWDATAFIND SERVICE

TASK Op-7.5-A(xxiii): GWOSC GWDATAFIND SERVICE

TASK Op-7.5-A(xxiv): LHO GWDATAFIND SERVICE

TASK Op-7.5-A(xxv): LLO GWDATAFIND SERVICE

TASK Op-7.5-A(xxvi): UWM GWDATAFIND SERVICE
[CD.MOS-DA.DFSUWM]

TASK Op-7.5-A(xxvii): CVMFS GWDATAFIND SERVICE
[CD.MOS-DA.DFSCVM]

TASK Op-7.5-A(xxviii): PSU GWDATAFIND SERVICE?

TASK Op-7.5-A(xxix): IUCAA GWDATAFIND SERVICE

TASK Op-7.5-A(xxx): IGWN GRID (OSG) GWDATAFIND SERVICE

TASK Op-7.5-A(xxxi): NDS2 SERVICE (CIT)

TASK Op-7.5-A(xxxii): NDS2 SERVICE (LHO)

TASK Op-7.5-A(xxxiii): NDS2 SERVICE (LLO)

TASK Op-7.5-A(xxxiv): NDS2 CLIENT TOOL

TASK Op-7.5-A(xxxv): SOFTWARE SOURCE DISTRIBUTION ARCHIVE (SOFTWARE.LIGO.ORG)
[CD.MOS-DA.SOURCEARCH]

TASK Op-7.5-A(xxxvi): IGWN CONDA SERVICE
[CD.MOS-DA.IGWNCONDA]

TASK Op-7.5-A(xxxvii): RHEL PACKAGE REPOSITORY SERVICE
[CD.MOS-DA.RHELREPO]

TASK Op-7.5-A(xxxviii): KOJI RHEL BUILD PLATFORM
[CD.MOS-DA.KOJI]

TASK Op-7.5-A(xxxix): FRAMEL LIBRARY

TASK Op-7.5-A(xl): LDAS TOOLS

TASK Op-7.5-A(xli): LIGO SUMMARY PAGE SERVICE
[CD.MOS-DA.SumPageMO,CD.MOS-DA.SumPageSu]

TASK Op-7.5-A(xlii): LIGO LDVW (MAINTENANCE, NEW REQUESTS, TECH SUPPORT)
[CD.MOS-DA.LDVWMO,CD.MOS-DA.LDVWSu]

TASK Op-7.5-A(xliii): LIGODV (MATLAB)
[CD.MOS-DA.LDVMO,CD.MOS-DA.LDVSu]

TASK Op-7.5-A(xliv): VIRGO DATA DISPLAY
[CD.MOS-DA.DyMO,CD.MOS-DA.DySu]

TASK Op-7.5-A(xlv): GWPY LIBRARY
[CD.MOS-DA.GWpyMO, CD.MOS-DA.GWSu]

TASK Op-7.5-A(xlvi): LALSUITE EXTRA TOOLS

TASK Op-7.5-A(xlvii): CIT IGWN GRID (OSG) SUBMIT NODE & LOCAL STORAGE SERVICES

TASK Op-7.5-A(xlviii): CIT IGWN GRID (OSG) SUBMIT NODE & LOCAL STORAGE SERVICES

TASK Op-7.5-A(xlix): LHO IGWN GRID (OSG) SUBMIT NODE & LOCAL STORAGE SERVICES

TASK Op-7.5-A(l): NIKHEF IGWN GRID (OSG) SUBMIT NODE & LOCAL STORAGE SERVICES

TASK Op-7.5-A(li): IUCAA IGWN GRID (OSG) SUBMIT NODE & LOCAL STORAGE SERVICES

TASK Op-7.5-A(lii): UWM IGWN GRID (OSG) SUBMIT NODE & LOCAL STORAGE SERVICES

TASK Op-7.5-A(liii): GTECH IGWN GRID (OSG) SUBMIT NODE & LOCAL STORAGE SERVICES

TASK Op-7.5-A(liv): CARDIFF IGWN GRID (OSG) SUBMIT NODE & LOCAL STORAGE SERVICES

TASK Op-7.5-A(lv): PSU IGWN GRID (OSG) SUBMIT NODE & LOCAL STORAGE SERVICES

TASK Op-7.5-A(lvi): LOW-LATENCY EMFOLLOW VIRTUAL MACHINES

TASK Op-7.5-A(lvii): LOW-LATENCY EVENT PROCESSING (GWCELERY) SERVICE

TASK Op-7.5-A(lviii): RAPID CBC LOCALIZATION SERVICE AND PUBLIC SKY MAP VISUALIZATION AND MANIPULATION PACKAGE (LIGO.SKYMAP)

TASK Op-7.5-A(lix): PUBLIC DOCUMENTATION AND SAMPLE CODE FOR PUBLIC ALERTS (LIGO/VIRGO PUBLIC ALERTS USER GUIDE)

TASK Op-7.5-A(lx): LIGO LOW LATENCY H(T) DISTRIBUTION SERVICE
[CD,MOS-DA.LLDMO, CD.MOS-DA.LLSu]

TASK Op-7.5-A(lxi): VIRGO LOW LATENCY H(T) DISTRIBUTION SERVICE

TASK Op-7.5-A(lxii): KAGRA LOW LATENCY H(T) DISTRIBUTION SERVICE

TASK Op-7.5-A(lxiii): ELECTRONIC NOTEBOOK SERVICE (ALOG)

TASK Op-7.5-A(lxiv): LIGO CHANNEL INFORMATION SYSTEM (CIS) SERVICE
[CD.MOS-DA.CISMO, CD.MOS-DA,CISSu]

TASK Op-7.5-A(lxv): CIT CLUSTER HTCONDOR SERVICE

TASK Op-7.5-A(lxvi): LLO CLUSTER HTCONDOR SERVICE

TASK Op-7.5-A(lxvii): LHO CLUSTER HTCONDOR SERVICE

TASK Op-7.5-A(lxviii): IUCAA CLUSTER HTCONDOR SERVICE?

TASK Op-7.5-A(lxix): UWM NEMO CLUSTER HTCONDOR SERVICE

TASK Op-7.5-A(lxx): CARDIFF CLUSTER HTCONDOR SERVICE?

TASK Op-7.5-A(lxxi): PSU CLUSTER HTCONDOR SERVICE

TASK Op-7.5-A(lxxii): LDG HTCONDOR SERVICE
[CD.MOS-DA.IGWNHTCMO,CD.MOS-DA.IGWNHTCSu]

TASK Op-7.5-A(lxxiii): CIT CLUSTER COMPUTING ELEMENT (CE) SERVICE

TASK Op-7.5-A(lxxiv): LLO CLUSTER COMPUTING ELEMENT (CE) SERVICE

TASK Op-7.5-A(lxxv): LHO CLUSTER COMPUTING ELEMENT (CE) SERVICE

TASK Op-7.5-A(lxxvi): IUCAA CLUSTER COMPUTING ELEMENT (CE) SERVICE

TASK Op-7.5-A(lxxvii): UWM NEMO CLUSTER COMPUTING ELEMENT (CE) SERVICE

TASK Op-7.5-A(lxxviii): GTECH CLUSTER COMPUTING ELEMENT (CE) SERVICE

TASK Op-7.5-A(lxxix): CARDIFF CLUSTER COMPUTING ELEMENT (CE) SERVICE

TASK Op-7.5-A(lxxx): PSU CLUSTER COMPUTING ELEMENT (CE) SERVICE

TASK Op-7.5-A(lxxxi): NIKHEF CLUSTER COMPUTING ELEMENT (CE) SERVICE

TASK Op-7.5-A(lxxxii): CIT CLUSTER JOB RESULTS WEB SERVER (LDAS-JOBS)

TASK Op-7.5-A(lxxxiii): LLO CLUSTER JOB RESULTS WEB SERVER (LDAS-JOBS)

TASK Op-7.5-A(lxxxiv): LHO CLUSTER JOB RESULTS WEB SERVER (LDAS-JOBS)

TASK Op-7.5-A(lxxxv): IUCAA CLUSTER JOB RESULTS WEB SERVER (LDAS-JOBS)

TASK Op-7.5-A(lxxxvi): UWM NEMO CLUSTER JOB RESULTS WEB SERVER (LDAS-JOBS)

TASK Op-7.5-A(lxxxvii): GTECH CLUSTER JOB RESULTS WEB SERVER (LDAS-JOBS)

TASK Op-7.5-A(lxxxviii): CARDIFF CLUSTER JOB RESULTS WEB SERVER (LDAS-JOBS)

TASK Op-7.5-A(lxxxix): PSU CLUSTER JOB RESULTS WEB SERVER (LDAS-JOBS)

TASK Op-7.5-A(xc): INFLUX DATABASE SERVICE

TASK Op-7.5-A(xci): LOW-LATENCY KUBERNETES PLATFORM SERVICE

TASK Op-7.5-A(xcii): LLO SEGMENT GENERATION (SEGGENER) SERVICE
[CD.MOS-DA.SEGGEN]

TASK Op-7.5-A(xciii): LHO SEGMENT GENERATION (SEGGENER) SERVICE
[CD.MOS-DA.SEGGEN]

TASK Op-7.5-A(xciv): KAGRA SEGMENT GENERATION SERVICE

TASK Op-7.5-A(xcv): VIRGO SEGMENT GENERATION SERVICE

TASK Op-7.5-A(xcvi): DATA-QUALITY SEGMENT (DQSEGDB) PUBLISHER SERVICE
[CD.MOS-DA.SEGPUB]

TASK Op-7.5-A(xcvii): INSPIRAL RANGE CALCULATOR TOOL

TASK Op-7.5-A(xcviii): DATA QUALITY REPORT (DQR) SERVICE
[CD.MOS-DA.DQRMO,CD.MOS-DA.DQRSu]

TASK Op-7.5-A(xcix): CHANNEL INFORMATION SYSTEM (CIS) SERVICE

TASK Op-7.5-A(c): CIT CLUSTER HARDWARE/OS

TASK Op-7.5-A(ci): CIT CLUSTER LDG USER ENVIRONMENT, SUBMIT NODE, AND SHARED FILESYSTEM SERVICES

TASK Op-7.5-A(cii): LLO CLUSTER HARDWARE/OS

TASK Op-7.5-A(ciii): LLO CLUSTER LDG USER ENVIRONMENT, SUBMIT NODE, AND SHARED FILESYSTEM SERVICES

TASK Op-7.5-A(civ): LHO CLUSTER HARDWARE/OS

TASK Op-7.5-A(cv): LHO CLUSTER LDG USER ENVIRONMENT, SUBMIT NODE, AND SHARED FILESYSTEM SERVICES

TASK Op-7.5-A(cvi): IUCAA CLUSTER HARDWARE/OS (25% PROPORTION DEDICATED TO IGWN COMPUTING AND SUBJECT TO DAC PRIORITIZATION)

TASK Op-7.5-A(cvii): ICTS CLUSTER HARDWARE/OS (50% PROPORTION DEDICATED TO IGWN COMPUTING AND SUBJECT TO DAC PRIORITIZATION)

TASK Op-7.5-A(cviii): IUCAA CLUSTER LDG USER ENVIRONMENT, SUBMIT NODE, AND SHARED FILESYSTEM SERVICES (25%)

TASK Op-7.5-A(cix): UWM NEMO CLUSTER HARDWARE/OS (90% PROPORTION DEDICATED TO IGWN COMPUTING AND SUBJECT TO DAC PRIORITIZATION)

TASK Op-7.5-A(cx): UWM NEMO CLUSTER LDG USER SERVICES (SUBMIT NODE, SOFTWARE ENVIRONMENT, AND SHARED FILESYSTEM) (90% PROPORTION DEDICATED TO IGWN COMPUTING AND SUBJECT TO DAC PRIORITIZATION)
[CD.MOS-DA.UWMLDG]

TASK Op-7.5-A(cxi): GTECH CLUSTER HARDWARE/OS (90% PROPORTION DEDICATED TO IGWN COMPUTING AND SUBJECT TO DAC PRIORITIZATION)

TASK Op-7.5-A(cxii): CARDIFF CLUSTER HARDWARE/OS (67% PROPORTION DEDICATED TO IGWN COMPUTING AND SUBJECT TO DAC PRIORITIZATION)

TASK Op-7.5-A(cxiii): CARDIFF CLUSTER LDG SERVICES (SUBMIT NODE, USER ENVIRONMENT, AND SHARED FILESYSTEM) (67%)

TASK Op-7.5-A(cxiv): PSU CLUSTER HARDWARE/OS (25% PROPORTION DEDICATED TO IGWN COMPUTING AND SUBJECT TO DAC PRIORITIZATION)

TASK Op-7.5-A(cxv): PSU CLUSTER LDG USER ENVIRONMENT, SUBMIT NODE, AND SHARED FILESYSTEM SERVICES (??%)
[CD.MOS-DA.PSULDG]

TASK Op-7.5-A(cxvi): OZSTAR CLUSTER HARDWARE/OS (30% PROPORTION DEDICATED TO IGWN COMPUTING AND SUBJECT TO DAC PRIORITIZATION)

TASK Op-7.5-A(cxvii): EGO CLUSTER HARDWARE/OS

TASK Op-7.5-A(cxviii): CNAF CLUSTER HARDWARE/OS (ALLOCATED CYCLES - NO M&O - SUPPORT ONLY)

TASK Op-7.5-A(cxix): IN2P3 CLUSTER HARDWARE/OS (ALLOCATED CYCLES - NO M&O - SUPPORT ONLY)

TASK Op-7.5-A(cxx): NIKHEF CLUSTER HARDWARE/OS (ALLOCATED CYCLES - NO M&O - SUPPORT ONLY)

TASK Op-7.5-A(cxxi): OTHER VIRGO CLUSTERS HARDWARE/OS (ALLOCATED CYCLES - NO M&O - SUPPORT ONLY)

TASK Op-7.5-A(cxxii): KISTI CLUSTER HARDWARE/OS (PROPORTION DEDICATED TO IGWN COMPUTING)

TASK Op-7.5-A(cxxiii): OTHER KAGRA CLUSTER HARDWARE/OS (PROPORTION DEDICATED TO IGWN COMPUTING)?

TASK Op-7.5-A(cxxiv): IGWN GRID (OSG) GLIDEINWMS SERVICE

TASK Op-7.5-A(cxxv): IGWN GRID (OSG) HTCONDOR POOL SERVICE

TASK Op-7.5-A(cxxvi): CVMFS ORIGIN SERVICE FOR SOFTWARE DISTRIBUTION

TASK Op-7.5-A(cxxvii): CVMFS ORIGIN SERVICE FOR PROPRIETARY DATA
[CD.MOS-DA.cvmfsOriServMO]

TASK Op-7.5-A(cxxviii): STASHCACHE (XROOTD) ORIGIN SERVICE
[CD.MOS-DA.StCaOrigMO,CD.MOS-DA.StCaOrigSu]

TASK Op-7.5-A(cxxix): OSG GLOBAL STASHCACHE (XCACHE) SERVICE

TASK Op-7.5-A(cxxx): LSCSOFT METAIO TOOL

TASK Op-7.5-A(cxxxi): LSCSOFT GRID LSC USER ENVIRONMENT (GLUE) TOOL

TASK Op-7.5-A(cxxxii): LSCSOFT PARAMETER ESTIMATION SUMMARY PAGE BUILDER (PE-SUMMARY) TOOL

TASK Op-7.5-A(cxxxiii): LSCSOFT LIGO-SEGMENTS TOOL

TASK Op-7.5-A(cxxxiv): LIGO DATA GRID (LDG) CLIENT TOOL

TASK Op-7.5-A(cxxxv): MATLAB APPLICATIONS (MATAPPS) TOOL

TASK Op-7.5-A(cxxxvi): LIGO TOOLS (LIGOTOOLS) TOOL

TASK Op-7.5-A(cxxxvii): LOW-LATENCY DATA REPLAY SYSTEM
[CD.MOS-DA.ReplayMO,CD.MOS-DA.ReplaySu]

TASK Op-7.5-A(cxxxviii): IGWN CONDA
[CD.MOS-DA.CondaMO,CD.MOS-DA.CondaSu]

TASK Op-7.5-A(cxxxix): CONTAINER BUILD AND DEPLOY INFRASTRUCTURE
[CD.MOS-DA.ContMO,CD.MOS-DA.ContSu]

ACTIVITY Op-7.5-B: CONTINUED MAINTENANCE, OPERATION, AND SUPPORT OF EXISTING GW COLLABORATION SUPPORT TOOLS & SERVICES (CD.MOS-CSS)

TASK Op-7.5-B(i): LIGO DCC SERVICE

TASK Op-7.5-B(ii): LIGO-INDIA DCC SERVICE

TASK Op-7.5-B(iii): KAGRA DCC SERVICE

TASK Op-7.5-B(iv): EPO WEBSITE SERVICE

TASK Op-7.5-B(v): MYLIGO SERVICE
[CD.MOS-CSS.myLIGOMO, CD.MOS-CSS.myLIGOSu]

TASK Op-7.5-B(vi): LIGO KERBEROS (KDC) SERVICE

TASK Op-7.5-B(vii): LIGO SHIBBOLETH IDP SERVICE

TASK Op-7.5-B(viii): LIGO BACKUP SHIBBOLETH IDP SERVICE

TASK Op-7.5-B(ix): LHO SHIBBOLETH IDP SERVICE

TASK Op-7.5-B(x): LLO SHIBBOLETH IDP SERVICE

TASK Op-7.5-B(xi): SHIBBOLETH SP TOOL ("SECOND-LEVEL" SUPPORT OF SHIBBOLETH SP INFRASTRUCTURE FOR SP ADMINS, RATHER THAN OF ANY SPECIFIC SP INSTANCE FOR ITS USERS)

TASK Op-7.5-B(xii): COMANAGE SERVICE
[CD.MOS-CSS.COmSu]

TASK Op-7.5-B(xiii): SAML PROXY (SATOSA) SERVICE

TASK Op-7.5-B(xiv): LIGO LDAP MASTER SERVICE

TASK Op-7.5-B(xv): LIGO ROBOTIC/AUTOMATIC CREDENTIALS SERVICE

TASK Op-7.5-B(xvi): LIGO DATA GRID CLUSTER ACCOUNT MANAGEMENT SERVICE

TASK Op-7.5-B(xvii): CIT LIGO LDAP REPLICA SERVICE

TASK Op-7.5-B(xviii): LIGO ADDRESSBOOK REPLICA SERVICE

TASK Op-7.5-B(xix): LIGO ROSTER SERVICE

TASK Op-7.5-B(xx): CIT GW-ASTRONOMY LDAP REPLICA SERVICE

TASK Op-7.5-B(xxi): UWM GW-ASTRONOMY LDAP REPLICA SERVICE
[CD.MOS-CSS.LDAPGWAUWMMO,CD.MOS-CSS.LDAPGWAUWMSu]

TASK Op-7.5-B(xxii): UWM LIGO LDAP REPLICA SERVICE
[CD.MOS-CSS.LDAPLIGOUWMMO,CD.MOS-CSS.LDAPLIGOUWMSu]

TASK Op-7.5-B(xxiii): UWM IGWN LDAP REPLICA SERVICE
[CD.MOS-CSS.LDAPIGWNUWMMO,CD.MOS-CSS.LDAPIGWNUWMSu]

TASK Op-7.5-B(xxiv): LLO LIGO LDAP REPLICA SERVICE

TASK Op-7.5-B(xxv): LHO LIGO LDAP REPLICA SERVICE

TASK Op-7.5-B(xxvi): OTHER CLUSTERS LDAP REPLICA SERVICES?

TASK Op-7.5-B(xxvii): LIGO GROUPER SERVICE

TASK Op-7.5-B(xxviii): LIGO REMOTE PARTICIPATION (TEAMSPEAK) SERVICE

TASK Op-7.5-B(xxix): VIRGO REMOTE PARTICIPATION (TEAMSPEAK) SERVICE

TASK Op-7.5-B(xxx): LIGO.ORG GSUITE SERVICES

TASK Op-7.5-B(xxxi): IGWN GSUITE SERVICES

TASK Op-7.5-B(xxxii): LSC MOU SOFTWARE [CD.MOS-CSS.MOU]

TASK Op-7.5-B(xxxiii): IGWN GITLAB SERVICE (GIT.LIGO.ORG)
[CD.MOS-CSS.GitLabMO, CD.MOS-CSS.GitLabSu]

TASK Op-7.5-B(xxxiv): LSC FOSWIKI SERVICE (WIKI.LIGO.ORG)
[CD.MOS-CSS.wikiMO,CD.MOS-CSS,wikiSu]

TASK Op-7.5-B(xxxv): VIRGO FOSWIKI SERVICE (WIKI.VIRGO-GW.EU)

TASK Op-7.5-B(xxxvi): CIT JUPYTER HUB SERVICE

TASK Op-7.5-B(xxxvii): LHO JUPYTER HUB SERVICE

TASK Op-7.5-B(xxxviii): LLO JUPYTER HUB SERVICE

TASK Op-7.5-B(xxxix): CARDIFF JUPYTER HUB SERVICE

TASK Op-7.5-B(xl): IGWN MATTERMOST (CHAT.LIGO.ORG) SERVICE

TASK Op-7.5-B(xli): IGWN {LIGO|IGWN}.ORG MAIL TRANSPORT AGENT (MTA) SERVICE

TASK Op-7.5-B(xlii): IGWN USAGE ACCOUNTING DATA COLLECTION CLIENT TOOL

TASK Op-7.5-B(xliii): IGWN USAGE ACCOUNTING DATABASE SERVICE
[CD.MOS-DA.UseAccMO,CD.MOS-DA.UseAccSu]

TASK Op-7.5-B(xliv): IGWN USAGE ACCOUNTING REPORTS (ACCOUNTING.IGWN.ORG) SERVICE

TASK Op-7.5-B(xlv): IGWN DNS ({LIGO|IGWN}.ORG) SERVICE

TASK Op-7.5-B(xlvi): VIRGO DNS (VIRGO-GW.EU) SERVICE

TASK Op-7.5-B(xlvii): LSC MAILING LISTS (SYMPA.LIGO.ORG) SERVICE
[CD.MOS-CSS.listMO,CD.MOS-CSS.listSu]

TASK Op-7.5-B(xlviii): VIRGO MAILING LISTS SERVICE

TASK Op-7.5-B(xlix): CILOGON GRID CERTIFICATE AUTHORITY SERVICE

TASK Op-7.5-B(l): GW-ASTRONOMY.ORG COMANAGE SERVICE
[CD.MOS-CSS.GWACOm,CD.MOS-CSS.GWACOmSu]

TASK Op-7.5-B(li): GW-ASTRONOMY.ORG MAILING LIST SERVICE
[CD.MOS-CSS.GWAMLMO,CD.MOS-CSS.GWAMLSu]

TASK Op-7.5-B(lii): GW-ASTRONOMY.ORG WIKI SERVICE
[CD.MOS-CSS.GWAWikMO,CD.MOS-CSS.GWAWikSu]

TASK Op-7.5-B(liii): LIGO-INDIA COMANAGE SERVICE

TASK Op-7.5-B(liv): LIGO-INDIA IDP SUPPORT SERVICE

TASK Op-7.5-B(lv): LIGO-INDIA KERBEROS SERVICE

TASK Op-7.5-B(lvi): LIGO-INDIA LDAP SERVICE

TASK Op-7.5-B(lvii): LSC VOTING/ELECTIONS (VOTE.LIGO.ORG) SERVICE

TASK Op-7.5-B(lviii): LSC P&P DATABASE SERVICE

TASK Op-7.5-B(lix): SERVICE MONITORING (ICINGA) SERVICE
[CD.MOS-CSS.ServMonMO,CD.MOS-CSS.ServMonSu]

TASK Op-7.5-B(lx): CI ECP CLIENT (CIECPLIB) TOOL

TASK Op-7.5-B(lxi): CLOUD PLATFORM SERVICE (LOW-LATENCY / UWM AWS)
[CD.MOS-CSS.LLCLD]

TASK Op-7.5-B(lxii): CLOUD PLATFORM SERVICE (GENERAL / CIT AWS)
[CD.MOS-CSS.CLD]

TASK Op-7.5-B(lxiii): CLOUD PLATFORM SERVICE (DATA / PRP K8S)
[CD.MOS-CSS.DCLD]

TASK Op-7.5-B(lxiv): CLOUD PLATFORM SERVICE (FUTURE / NSF AWS)
[CD.MOS-CSS.NCLD]?

TASK Op-7.5-B(lxv): CBC PROJECT PLANNING SERVICE
[CD.MOS-CSS.CBCProjMO, D.MOS-CSS.CBCProjSu]

Kg-7.6 KAGRA Computing

The computing and networking facilities of KAGRA consist of two parts. The first one, which is referred to as the iKAGRA system, was installed at Kamioka and Kashiwa sites in 2014. It is still dedicated to data transfer from Kamioka to Kashiwa now. The second one, which is referred to as the KAGRA Main Data Server (KMDS), was installed at Kashiwa in 2017. It is dedicated to data storage as the KAGRA Tier-0 center, data transfers to Tier-1 and -2 centers, low-latency and bulk data transfers between LIGO/Virgo, detector characterization and commissioning analysis. The low-latency searches and offline analyses are shared with the other systems at the Kashiwa site, the Osaka City University, RESCEU of the University of Tokyo, etc. We are planning to upgrade of KMDS and integrate the iKAGRA system into KMDS by the beginning of O4. Since the financial resources for the computing facilities of KAGRA is limited, there is a need to gather resources of many computing centers. For a robust data distribution and access framework, stable network links between these computing resources must be provided.

Kg-7.7 KAGRA Computing operations for O4

O. The fourth KAGRA observing run - O4

O.C. Prioritized list of central tasks

Highest priority tasks

O.C.1: KAGRA Platform and Services

- O.C.1.1. Finalize configuring Rusio on KMDS for bulk data transfer between KAGRA, LIGO and Virgo
- O.C.1.2. Upgrade of storage, networking and computing resources of KMDS
- O.C.1.3. Monitor the existing bulk data transfer from Kamioka site to Tier-0 center at Kashiwa in order to maintain the quasi-real time performances achieved
- O.C.1.4. Monitor the K1 aggregated hoft files transfer to CIT
- O.C.1.5. Improve user documentation and tutorials for data analysis computing

O.C.2: Software Management

- O.C.2.1. Maintain new software package version deployments and servers configurations under control
- O.C.2.2. Transition to Conda for software environment definition and software packaging
- O.C.2.3. Maintain KAGALI (KAGRA Algorithmic Library) and increase the support to its and code review

O.C.3: Low Latency data distribution

- O.C.3.1. Transition to Kafka for the CIT $\leftrightarrow$ Kashiwa low latency links
- O.C.3.2. Monitor the CIT $\leftrightarrow$ Kashiwa low latency links

O.C.4: Offline Computing Services

- O.C.4.1. Define and implement the data cataloging

O.C.5: Data Analysis Tools And Pipelines

- O.C.5.1. Progress with the capability to run LV pipelines into KMDS and other computing facilities at Kashiwa, RESCEU and Osaka
- O.C.5.2. Development of KAGRA pipelines using KAGALI

Kg-7.8 KAGRA Computing preparation for future operations

F.C. Prioritized list of central tasks for future observing runs

Highest priority tasks

F.C.1: KAGRA Platform and Services

- F.C.1.1. Integration of the computing resources in Japan, Taiwan, Korea, etc. to provide a homogeneous distributed environment
- F.C.1.2. Consider providing grid computing environment
- F.C.1.3. Upgrade or replacement of data transfer from Kamioka to Kashiwa

F.C.2: Software Management

- F.C.2.1. Full deployment of KAGALI for software distribution from KAGRA

F.C.3: Data Analysis Tools

- F.C.3.1. Development of Machine Learning tools for low-latency and offline analysis

8 Joint Run Planning Committee

The LIGO-Virgo Joint Run Planning Committee (JRPC) is appointed by the LSC, Virgo, and KAGRA collaboration leadership. The JRPC facilitates the coordination of run planning between those collaborations. It is also responsible to ensure good communication with the greater scientific community on run planning, low-latency alerts, and more generally observatory operations.

Op-8.1 JRPC Activities related to planning for O4

For the period covered by this White paper, the JRPC will

- Update and maintain public-facing documentation on run plans, ensuring they are correctly communicated to the greater scientific community. An online approach via the IGWN Editorial Teams (see Op-6.7) is planned. Re-assessment of any network capabilities as a function of projected sensitivity will be needed, drawing in modeling expertise.
- Organize LVEM/scientific community telemeetings and Town Halls as needed
- Facilitate the organization and work of the RRT
- Work with the Virgo, KAGRA, and LIGO organization to further define the best role for the JRPC as the collaborations evolve
- Work with KAGRA on the scope of run planning to facilitate its full integration into the structures of the LVK

ACTIVITY Op-8.1-A: JRPC AND RRT PREPARATION FOR O4

TASK Op-8.1-A(i): JRPC AND RRT SUPPORTING O4 PREP

ACTIVITY Op-8.1-B: JRPC PUBLIC OBSERVING SCENARIO

TASK Op-8.1-B(i): JRPC OBSERVING SCENARIO BACKGROUND RESEARCH

TASK Op-8.1-B(ii): JRPC OBSERVING SCENARIO WEB PRESENCE

Op-8.2 JRPC Activities beyond the next 12 months related to O4 preparations

- Roughly 1 year before the O4 Observing Run starts, the JRPC will re-assess its responsibilities and activities to ensure its mission is fulfilled to support a smooth start to O4.
- If Observing plans are changed, JRPC will arrange to support them.

ACTIVITY Op-8.2-A: JRPC AND RRT O4 OPERATIONS IN O4

TASK Op-8.2-A(i): JRPC AND RRT SUPPORTING O4 OPS

9 Support of Observatories

Op-9.1 LIGO Committee

The Support of Observatories Committee (SO) is, by design, a resource to be used by both LIGO Laboratory personnel and LSC members outside the LIGO laboratory to facilitate collaborative work aimed at improving detector performance.

The SO chairs, Anamaria Effler at LLO (aeffler@caltech.edu) and Rick Savage at LHO (rsavage@caltech.edu) serve as liaisons between the LIGO observatories and the broader LSC community. LSC members (including LIGO Lab. staff) and others who either have a need for support from the LSC for a particular investigation, or who would like to get involved in efforts to support the observatories, can contact the SO chairs for assistance in identifying and organizing efforts to satisfy the needs.

ACTIVITY Op-9.1-A: ACTIVITIES RELATED TO SUPPORT OF OBSERVATORIES

One of the existing avenues for organizing and supporting on-site contributions by LSC members is the LSC Fellows program. It provides logistical support (airfare, housing, rental cars, etc.) for long-term (three months or more) visits to the observatory sites. Descriptions of past projects that have engaged LSC fellows can be found on the LSC Fellows wiki page: https://wiki.ligo.org/LSC/Fellows/LSCFellowsProjectl Another existing avenue for shorter stays is the LIGO Visitor Program. The SO chairs could assist with organizing support from these programs, if appropriate. Note that The LSC Fellows program was suspended until April 2021 due to the Covid pandemic. While plans are somewhat uncertain, reservations are now being accepted for both observatories for Fellows visits starting July 2021.

TASK Op-9.1-A(i): CONTRIBUTIONS TO THE LSC FELLOWS PROGRAM

The range of avenues by which LSC members can contribute to the observatories spans much of the scope of the LSC. Relevant areas of focus are detailed in Section 2 of LIGO Scientific Collaboration Program (2021 Edition) (https://dcc.ligo.org/LIGO-M2100100), specifically 2.1 LIGO Observatory Operations, 2.2 LSC Detector Commissioning and Detector Improvement activities, 2.3 A+ Upgrade Project, 2.6 LSC Fellows Program, 2.7 Detector Calibration and Data Timing, and 2.8 Detector Characterization. Earlier sections of this LSC-Virgo Operations White Paper also list specific aspects of the work planned for 2021 and 2022 that require support from the LSC.

LT-9.1 Long-term plans

With the delays associated the Covid pandemic, the O4 observing run is not expected to begin before August 2022. The remainder of 2021 and the beginning of 2022 are expected to be dedicated to improvements in detector hardware - test masses, lasers, squeezers, output modecleaners, etc. Installation and shake-down periods will be interspersed with commissioning intervals during which the impact of the hardware changes will be assessed. Contributions by LSC members during the latter half of 2021 and the first half of 2022 are thus expected to involve a broad range of activities, except those specifically relevant to observing runs. This latter set of activities, that includes vetting of candidate events, serving as on-site liaison for the Detector Commissioning shifters, etc., won't resume until the start of the O4 observing run.

The period between August 2021 and August 2022 should provide ample opportunities for LSC members to be involved in activities related to interferometer hardware, including in-vacuo components. During the O4 observing run there will be ample opportunities for LSC members to be involved at the observatories with detector calibration and detector characterization activities as well as activities related to the vetting of event candidate. The SO committee chairs will help to coordinate participation of interested LSC members in these and other activities.

TASK LT-9.1-(i): GENERAL CONTRIBUTIONS TO SITE PREPARATIONS FOR O4

10 Open Data

Op-10.1 Deliver and Support Public Data Products

The Open Data Working Group is responsible for public release of instrumental data and associated documentation. The entry point for discovering and accessing these data products is the Gravitational Wave Open Science Center (GWOSC) web site, available at `gw-openscience.org`. The scope and timing of LIGO data releases are described in the LIGO Laboratory Data Management Plan (DCC: M1000066).

In addition to developing and maintaining a range of resources on the GWOSC web site, the Open Data Working Group helps educate the scientific community about GW data analysis though workshops and email support. We have hosted a number of Open Data Workshops, where mentors from across the LVK share their expertise with students and scientists getting started working with LIGO/Virgo data. Past workshop materials are available on the GWOSC web site[11], and additional notes may be found in DCC documents G2000859, G1900798, and G1800778.

GWOSC web services are maintained by a "development team" that includes a small dedicated staff. In addition, any member of the LVK is invited to join the GWOSC development team and contribute. Goals for the GWOSC development team include:

[11] gw-osc.org/workshops

- Release O3b data set
- Develop, review, and release O3c GEO/KAGRA data
- Add segments and query features to GWOSC client
- Evolve data archive directory structure to a format more consistent with internal data archive, both for events and bulk data
- Plan Open Data Workshop #5
- Release an interactive app to display PE samples
- Streamline PE ingestion, maybe directly from samples, and promote standardization of confidence intervals
- Develop a policy for releasing auxiliary channels (or not)
- Inestigate use of serializers to streamline JSON / HTML interface
- Maintain the GWOSC web server and quality of service to the users
- Prepare, review, and release strain data from events, catalogs, and bulk data, as described in the data management plan
- Coordinate with data analysis working groups and paper writing teams to support public release of some analysis results, including parameter estimation samples
- Maintain and update the GWOSC Event Portal
- Collaborate with calibration and detchar teams to ensure that for O4, the public version of the final, calibrated strain data will be delivered to working groups within a few months of data collection. This will enable identical files to be released to the public as were used for astrophysical searches.

In addition, we stress that there are a number of important areas related to open data where any member of the LVK may make significant contributions. These activities provide a valuable service both to the LVK and the larger scientific community, and should will count towards InfraOps service. LVK members are encouraged to contact the Open Data Working Group chairs if they wish to volunteer. Some areas where additional volunteers are still needed for FY21 include:

- Develop tutorials and mentor at an Open Data Workshop. The 2021 Open Data Workshop was supported by a team of 20 mentors, representing a diverse cross-section of the LVK. We hope that future workshops will receive similarly broad support.
- Serve as a mentor in GWOSC Office hours
- Plan and host an Open Data Workshop
- Answer questions in the GWOSC ticket system (gwosc@igwn.org)
- Prepare, review, and document parameter estimation sample data releases
- Develop and maintain the `gwosc` python client

- Prepare and document potential auxiliary channel releases
- Prepare and document data quality and segment information for O3b and O4
- Serve on the GWOSC review team

ACTIVITY Op-10.1-A: GWOSC ACTIVITIES

TASK Op-10.1-A(i): GWOSC DEVELOPMENT TEAM

TASK Op-10.1-A(ii): PLAN AND HOST AN OPEN DATA WORKSHP

TASK Op-10.1-A(iii): MENTORS FOR OPEN DATA WORKSHOPS AND GWOSC OFFICE HOURS

TASK Op-10.1-A(iv): CALIBRATION TEAM ADDITIONAL EFFORT TO PREPARE PUBLIC STRAIN DATA FOR WORKING GROUPS

TASK Op-10.1-A(v): REVIEW

11 Upgrades

Op-11.1 LIGO A+ Upgrade

The “A+ detector" project is a major upgrade to the existing Advanced LIGO detectors, which began in 2019 and expected to continue through the end of 2023. Design, procurement and installation for A+ is well underway, led by the LIGO Lab with UK and Australian hardware contributions, as well as support by other members of the LSC.

ACTIVITY Op-11.1-A: ACTIVITIES CONNECTED TO THE A+ UPGRADE

A number of A+ Project deliverable elements have been fast-tracked for installation prior to the O4 run. These fast-tracked (accelerated development) elements of the A+ Project include a high-transmission Faraday isolator in the output optics section, adaptive wavefront control for improved mode-matching to the output mode cleaner and frequency-dependent squeezing. With concurrent detector improvements (including mitigation of point absorbers and higher laser power), the BNS range for O4 could be as high as 192 Mpc; This would effectively triple the O3 event detection rate.

In order to accomplish this A+ scope for O4 a number of vacuum system and facility modifications must be accomplished. Facility construction of a Filter Cavity End Station (FCES) and Filter Cavity Enclosure (FCE) have recently begun at both observatories. The filter cavity vacuum system elements have begun arriving at the observatories as well. In particular two HAM chambers (from decommissioning the H2 detector) have been shipped from LHO to LLO to serve as chambers for the Frequency-Dependent Squeezing (FDS) system. Efforts are underway to ensure particulate and hydrocarbon cleanliness of these chambers.

The FDS components (seismic isolation, suspensions, electronics, photodetectors, scattered light baffling, etc.) have all been delivered to both observatories and are ready for installation. The high-transmission (low loss) Faraday isolators have been delivered and are in the process of being installed.

The adaptive wavefront control elements (both thermally and piezo-actuated) are in production and testing phases.

The full A+ expected BNS design range of 325 MPc is not expected until the O5 run with the addition of lower thermal noise coatings on the test mass optics, a larger beamsplitter and suspension and the addition of a Balanced Homodyne Detection (BHD) readout system.

TASK Op-11.1-A(i): CONTRIBUTIONS TO THE A+ UPGRADE

Vg-11.2 Virgo AdV+ Upgrade

Advanced Virgo Plus (AdV+) is the upgrade of Advanced Virgo to be realized in two steps: named Phase I and Phase II. The installation of AdV+ Phase I started in May 2020 and was concluded in May 2021. The installation of AdV+ Phase II will take place after the end of O4. While the main goal of AdV+ Phase I is to reduce the interferometer quantum noise, AdV+ Phase II will be focused on the reduction of the mirror thermal noise. The main changes will concern:

- increase of the beam size on the end test masses, from 6 cm to 10 cm radius;
- increase of the size (from 35 cm diameter to 55 cm diameter) and mass (from 40 kg to 100 kg) of the end mirrors;
- improvement of mirror coatings: ower mechanical losses, less point defects, better uniformity;
- new suspensions/seismic isolators for large mirrors;
- further increase of the interferometer input power (from 40 W to 80 W).

To achieve the planned increase of the beam size on the ETMs, while keeping the same size in the central area of the interferometer, the geometry of the arm cavities and of the recycling cavities has to be modified as well, requiring the replacement also of the two cavity input mirrors, the power and signal recycling mirrors. The reduction of the mirror thermal noise also relies on the development of better coatings, which is part of AdV+ Phase I. Other detector subsystems are involved in AdV+ Phase II: vacuum; thermal compensation system; injection; detection; pre-stabilized laser and stray light control (instrumented baffles).

The list of deliverables for AdV+ Phase II is available and the Technical Design Report is in the process of being written and reviewed.

The combination of the improvements outlined above should allow increasing BNS range of Virgo to around 200 Mpc.

12 LIGO-India

Op-12.1 LIGO-India

LIGO-India is a project of the Government of India with primary responsibilities to build facilities and assemble, install, commission and operate an advanced LIGO detector provided by LIGO and the US National Science Foundation. "In-principle" approval by the Cabinet of the Government of India for LIGO-India was granted on February 17, 2016.

ACTIVITY Op-12.1-A: ACTIVITIES RELATED TO LIGO-INDIA

Some important activities that were expected to be completed in 2020-21 were negatively impacted by the pandemic. Completion of a national Testing and Training Facility (TTF) at RRCAT, and the building of a construction office at the Aundha site were impacted by lockdowns, but should be completed soon. The vacuum chamber prototype fabrication at VTPL in Bangalore suffered some delay but has successfully completed all factory acceptances before shipping to RRCAT. Procurement has begun on prototyping and testing of a long cryopump and gate valve system. Procurement has also begun for a 10-m prototype interferometer at the RRCAT TTF. Planning has begun to ship some PSL and CDS equipment to India to support training in the new TTF. Travel bans prevented India participation in the detector upgrades at LHO and LLO, but optical simulation work has continued. LIGO has begun work to extend remote operations capabilities to India, as well as creating video documentation of some detector upgrade work.

Important activities to be completed in 2021-2022 include the final vacuum testing of the prototype vacuum chambers at RRCAT, testing of the protoype long cryo pump, development of the 10-m prototype at RRCAT and extension of remote operations capability to India. Once the DAE-CIT-MIT MoU is signed, shipments of LIGO hardware can begin to the new TTF. Following "final" approval for the LIGO-India Project by the Cabinet of India, full construction activities can proceed for the LIGO Aundha Observatory. The LSC is also engaged in developing and training the LIGO-India scientific workforce and planning the integration of LIGO-India data into the full detector network.

TASK Op-12.1-A(i): CONTRIBUTION TO THE LIGO-INDIA PROJECT

References

[1] LIGO-Virgo-KAGRA Detector Characterization Requirements Document. *LIGO Document: LIGO-T2100043; Virgo Document: VIR-0402A-21*, 2021.

[2] The LSC Operations Division, Virgo Group Coordinators, KAGRA Group Coordinators. LSC-Virgo-KAGRA Operations White Paper (Summer 2020 edition). *LIGO Document: LIGO-T2000294; Virgo Document: VIR-0551B-20*, 2020.

[3] A. D. Viets, M. Wade, A. L. Urban, S. Kandhasamy, J. Betzwieser, D. A. Brown, J. Burguet-Castell, C. Cahillane, E. Goetz, K. Izumi, S. Karki, J. S. Kissel, G. Mendell, R. L. Savage, X. Siemens, D. Tuyenbayev, and A. J. Weinstein. Reconstructing the calibrated strain signal in the Advanced LIGO detectors. *Class. Quantum Grav.*, 35(9):095015, apr 2018.

[4] D. Estevez, B. Mours, L. Rolland, and D.Verkindt. Online h(t) reconstruction for Virgo O3 data: start of O3. Technical Report VIR-0652B-19, Virgo, 2019.

[5] B. P. Abbott et al. Calibration of the Advanced LIGO detectors for the discovery of the binary black-hole merger GW150914. *Phys. Rev. D*, 95(6):062003, 2017.

[6] D. Tuyenbayev, S. Karki, J. Betzwieser, C. Cahillane, E. Goetz, K. Izumi, S. Kandhasamy, J. S. Kissel, G. Mendell, M. Wade, A. J. Weinstein, and R. L. Savage. Improving LIGO calibration accuracy by tracking and compensating for slow temporal variations. *Class. Quantum Grav.*, 34(1):015002, dec 2016.

[7] S. Karki, D. Tuyenbayev, S. Kandhasamy, B. P. Abbott, T. D. Abbott, E. H. Anders, J. Berliner, J. Betzwieser, C. Cahillane, L. Canete, C. Conley, H. P. Daveloza, N. De Lillo, J. R. Gleason, E. Goetz, K. Izumi, J. S. Kissel, G. Mendell, V. Quetschke, M. Rodruck, S. Sachdev, T. Sadecki, P. B. Schwinberg, A. Sottile, M. Wade, A. J. Weinstein, M. West, and R. L. Savage. The Advanced LIGO photon calibrators. *Review of Scientific Instruments*, 87(11):114503, 2016.

[8] Dripta Bhattachargee et al. Fiducial displacements with improved accuracy for the global network of gravitational wave detectors. Technical Report P2000113.

[9] Yuki Inoue, Sadakazu Haino, Nobuyuki Kanda, Yujiro Ogawa, Toshikazu Suzuki, Takayuki Tomaru, Takahiro Yamanmoto, and Takaaki Yokozawa. Improving the absolute accuracy of the gravitational wave detectors by combining the photon pressure and gravity field calibrators. *Physical Review D*, 98(2):022005, 2018.

[10] Virgo Collaboration. Virgo photon calibrators for O3. *In preparation*.

[11] D Estevez, B Lieunard, F Marion, B Mours, L Rolland, and D Verkindt. First tests of a newtonian calibrator on an interferometric gravitational wave detector. *Classical and Quantum Gravity*, 35(23):235009, 2018.

[12] Technical report.

[13] Ling Sun et al. Characterization of systematic error in Advanced LIGO calibration. 5 2020.

[14] The LIGO Scientific Collaboration and the Virgo Collaboration. Gracedb | public alerts.

[15] Review of data mitigation for the O3a event catalog. `https://dcc.ligo.org/T2000251`.

[16] O3b data mitigation review status update.
https://dcc.ligo.org/LIGO-G2100185.

[17] LIGO-Virgo Data Quality Report for Event Candidates Design Document.
https://dcc.ligo.org/LIGO-T1800123.

[18] Data Quality Report Documentation.
https://docs.ligo.org/detchar/data-quality-report/.

[19] Establishing Detection Confidence for a Continuous Wave Signal in the Advanced Detector Era.
https://dcc.ligo.org/LIGO-T1500151.

[20] Detection checklists for stochastic searches.
https://git.ligo.org/stochastic/detection_checklist/-/blob/master/checklists/checklists.pdf.

[21] The LIGO-Virgo Summary Pages.
https://www.gw-openscience.org/detector_status/.

[22] Data Monitor Tool Project .
https://labcit.ligo.caltech.edu/~jzweizig/DMT-Project.html.

[23] The LIGO-Virgo segment database project wiki.
https://wiki.ligo.org/Computing/DQSEGDB.

[24] The LIGO-Virgo segment database.
https://segments-web.ligo.org/.

[25] R. P. Fisher, G. Hemming, M.-A. Bizouard, D. A. Brown, P. F. Couvares, F. Robinet, and D. Verkindt. DQSEGDB: A time-interval database for storing gravitational wave observatory metadata. *SoftwareX*, 14:100677, 2021.

[26] PyOmicron.
https://doi.org/10.5281/zenodo.2616570.

[27] B.P. Abbott et al. Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914. *Class. Quant. Grav.*, 33(13), 2016.

[28] DetChar Hardware Injections.
https://dcc.ligo.org/LIGO-T1900555.

[29] Reed Essick, Patrick Godwin, Chad Hanna, Lindy Blackburn, and Erik Katsavounidis. iDQ: Statistical Inference of Non-Gaussian Noise with Auxiliary Degrees of Freedom in Gravitational-Wave Detectors. *arXiv preprint arXiv:2005.12761*, 2020.

[30] SNAX toolkit: Stream-based Noise Acquisition and eXtraction toolkit.
https://lscsoft.docs.ligo.org/gstlal/feature_extraction.html.

[31] GWpy: a python package for gravitational-wave astrophysics.
https://gwpy.github.io
https://doi.org/10.5281/zenodo.597016.

[32] GW Detector Characterization python utilities.
https://github.com/gwdetchar/gwdetchar
https://doi.org/10.5281/zenodo.2575786.

[33] Gravitational-wave interferometer summary information.
https://github.com/gwpy/gwsumm
https://doi.org/10.5281/zenodo.2647609.

[34] Stochmon: A LIGO Data Analysis Tool.
https://dcc.ligo.org/LIGO-T1400205.

[35] STAMP-PEM.
https://ldas-jobs.ligo-wa.caltech.edu/~stochastic/stamppem_docs/.

[36] J. R Smith, T. Abbott, E. Hirose, N. Leroy, D. MacLeod, J. McIver, P. Saulson, and P. Shawhan. A hierarchical method for vetoing noise transients in gravitational-wave detectors. *Classical and Quantum Gravity*, 28(23):235005, 2011.
https://doi.org/10.5281/zenodo.2584615.

[37] Tomoki Isogai. Used percentage veto for LIGO and virgo binary inspiral searches. *J. Phys. Conf. Ser.*, 243:012005, 2010.

[38] Joseph S. Areeda, Joshua R Smith, Andrew P Lundgren, Edward Maros, Duncan M. Macleod, and John Zweizig. LigoDV-web: Providing easy, secure and universal access to a large distributed scientific data store for the LIGO Scientific Collaboration. *Astron. Comput.*, 18:27–34, 2017.

[39] The LIGO Channel Information System.
https://github.com/lscsoft/cis.server.

[40] Pointy Poisson.
https://docs.ligo.org/detchar/data-quality-report/tasks/pointy-dqr.html.

[41] September 2019 O3a Saefty Studies.
https://dcc.ligo.org/LIGO-T1900697.

[42] The Gravitational-wave Veto Evaluation and Testing suite.
https://github.com/gwpy/vet.

[43] D. Davis, T. J. Massinger, A. P. Lundgren, J. C. Driggers, A. L. Urban, and L. K. Nuttall. Improving the Sensitivity of Advanced LIGO Using Noise Subtraction. *Class. Quant. Grav.*, 36(5):055011, 2019.

[44] Gabriele Vajente, Yiwen Huang, Maximiliano Isi, Jenne C. Driggers, Jeffrey S. Kissel, Marek J. Szczepanczyk, and Salvatore Vitale. Machine-learning nonstationary noise out of gravitational-wave detectors. *Phys. Rev. D*, 101(4):042003, 2020.

[45] Rich Ormiston, Tri Nguyen, Michael Coughlin, Rana X. Adhikari, and Erik Katsavounidis. Noise Reduction in Gravitational-wave Data via Deep Learning. *arXiv:2005.06534*, 5 2020.

[46] Providing remote access to the LIGO interferometers with EPICS based tools .
https://indico.fnal.gov/event/9718/session/3/contribution/54/material/slides/0.pdf.

[47] DataViewer User Guide.
https://labcit.ligo.caltech.edu/~aivanov/Dataviewer/helpG2.html.

[48] LIGO Physical Environment Monitoring.
http://pem.ligo.org/.

[49] DetChar Hardware Injections.
https://git.ligo.org/detchar/detchar-hwinj.

[50] Michael Zevin et al. Gravity Spy: Integrating Advanced LIGO Detector Characterization, Machine Learning, and Citizen Science. *Class. Quant. Grav.*, 34(6):064003, 2017.

[51] Michael Coughlin. Noise line identification in LIGO S6 and Virgo VSR2. *J. Phys. Conf. Ser.*, 243:012010, 2010.

[52] Michael Coughlin et al. Noise Line Identification in LIGO S6 and Virgo VSR2. *Journal of Physics: Conference Series*, 243:012010, 2010.

[53] J. Rollins et al. Distributed state machine supervision for long-baseline gravitational-wave detectors. *Review of Scientific Instruments*, 87(094502), 2016.

[54] Elena Cuoco et al. Enhancing Gravitational-Wave Science with Machine Learning. *arXiv:2005.03745*, 2020.

[55] The Gravitational Wave Open Science Center.
https://www.gw-openscience.org.

[56] Proposal for common data formats for LIGO-Virgo data products.
https://dcc.ligo.org/LIGO-T1800014.

[57] Category 1 (CAT1) data quality vetoes applied to the analysis of the O3 run Virgo data. *Virgo Document VIR-0560A-21*, 2021.

[58] Irene Fiori Cristiano Palomba Kevin Turbang Lorenzo Pierini Maria Concetta Tringali Nicolas Arnaud Agata Trovato Ornella Piccinni, Kamiel Janssens. Virgo O3 List of Lines (v2). *Virgo Document VIR-0291A-21*, 2021.

[59] The REINFORCE (Research Infrastructures FOR Citizens in Europe) project.

[60] Reed Essick, Geoffrey Mo, and Erik Katsavounidis. A coincidence null test for poisson-distributed events. *Phys. Rev. D*, 103:042003, Feb 2021.

[61] Michael Coughlin, Paul Earle, Jan Harms, Sebastien Biscans, Christopher Buchanan, Eric Coughlin, Fred Donovan, Jeremy Fee, Hunter Gabbard, Michelle Guy, Nikhil Mukund, and Matthew Perry. Limiting the effects of earthquakes on gravitational-wave interferometers. *Classical and Quantum Gravity*, 34(4):044004, feb 2017.

[62] S Biscans, J Warner, R Mittleman, C Buchanan, M Coughlin, M Evans, H Gabbard, J Harms, B Lantz, N Mukund, A Pele, C Pezerat, P Picart, H Radkins, and T Shaffer. Control strategy to limit duty cycle impact of earthquakes on the LIGO gravitational-wave detectors. *Classical and Quantum Gravity*, 35(5):055004, jan 2018.

[63] Nikhil Mukund, Michael Coughlin, Jan Harms, Sebastien Biscans, Jim Warner, Arnaud Pele, Keith Thorne, David Barker, Nicolas Arnaud, Fred Donovan, Irene Fiori, Hunter Gabbard, Brian Lantz, Richard Mittleman, Hugh Radkins, and Bas Swinkels. Ground motion prediction at gravitational wave observatories using archival seismic data. *Classical and Quantum Gravity*, 36(8):085005, apr 2019.

[64] KAGRA SummarPages.
https://gwdet.icrr.u-tokyo.ac.jp/~controls/summary/today/.

[65] Bruco.
https://github.com/mikelovskij/bruco.

[66] Pastavi.
H. Yuzurihara, *JGW-G1910938*
https://gwdoc.icrr.u-tokyo.ac.jp/cgi-bin/private/DocDB/ShowDocument?docid=10938.

[67] NoiseBudgetter.
A. Shoda, *JGW-T1910917*
https://gwdoc.icrr.u-tokyo.ac.jp/cgi-bin/private/DocDB/ShowDocument?docid=10917.

[68] Omicron.
F. Robinet, *VIR-0545C-14*
https://tds.virgo-gw.eu/?content=3&r=14693.

[69] Omega-scan.
https://github.com/gwdetchar/gwdetchar.

[70] CAGMon.
http://gwwiki.icrr.u-tokyo.ac.jp/JGWwiki/CAGMon.

[71] S. Morisaki. ICA. *Proc. Jpn. Acad. Ser.*, B92:336, 2016.

[72] LIGO calibration DCC Document Tree.
https://dcc.ligo.org/LIGO-E1500281.

[73] LIGO calibration web page.
https://wiki.ligo.org/Calibration/WebHome.

List of Tasks