LIGO SCIENTIFIC COLLABORATION
VIRGO COLLABORATION

Document Type	**LIGO–T2000294-v3** **VIR-0551B-20**

LSC-Virgo-KAGRA Operations White Paper
(Summer 2020 edition)

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

WWW: http://www.ligo.org/ and http://www.virgo.infn.it

Processed with LaTeX on 2021/02/06

Contents

1 Overview and Executive Summary

The successful operation of the LIGO and Virgo 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), and by the Virgo Collaboration in collaboration with EGO.

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 and Virgo **Commissioning** teams work toward optimizing the sensitivity and uptime of the global detector network. The LIGO and Virgo **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 LIGO and Virgo **Calibration** teams produce fast a reliable calibration of detector data. The LIGO and Virgo **Computing** teams support the joint LIGO-Virgo software and computing infrastructure. The LIGO and Virgo **Low Latency** teams provide rapid notification of Gravitational Wave candidate evenst to the wider observing community. The LIGO and Virgo **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 Operations White Paper* describes the planned activities for these efforts, in the context of the LIGO Scientific Collaboration Program 2020-2021, LIGO-M2000130.

Further details on the planned activities in collaboration between the DetChar and Calibration groups and the four LIGO-Virgo data analysis working groups (Burst, Compact Binary Coalescence, Continuous Waves, and Stochastic Gravitational-Wave Background) can be found in the *LSC-Virgo White Paper on GW Data Analysis and Astrophysics*, LIGO-T2000424. The data analysis working groups also undertake tasks related to the Operations of the LIGO-Virgo experiments, described therein.

This Operations White Paper also complements the *Instrument Science White Paper,* LIGO-T2000407, 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 the EGO Consortium, consisting of CNRS and INFN, plus NIKHEF as an associated member.

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

On June 18, 2020 the LIGO Observatory Management Team announced that O3 observations would not resume and the lab would proceed with detector improvements in preparation for O4. At the time of writing the full impact of the pandemic is not yet clear, so what follows is our best guess as to how the next year will unfold.

During the period 2020-2021 the LIGO Laboratory will initially focus on recovering detector sensitivity once the pandemic-related shutdown ends. This will be followed by detector improvements in preparation for the O4 observation run. The LIGO Laboratory science and engineering teams will proceed with plans for the 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.

1.1.2 Virgo

EGO went out of strictly minimal operations on May 4, 2020 and became available for the all possible scenarios, ranging from restarting the O3 run to concentration on AdV+ installation and preparation of the O4 run.

Since the suspension of O3, the main Virgo target has been the installation of the detector improvements in preparation for the O4 observation run. The main improvements will be the installation of the signal recycling mirror, the installation of a higher-power laser and the realization of a 300 m filter cavity to implement the frequency-dependent squeezing. An array of sensors will also be installed to implement a Newtonian noise subtraction system. Once any upgrade is installed, its commissioning will start. When all the installation are completed, a period of commissioning and noise hunting will start to reach the sensitivity target for O4.

1.1.3 KAGRA Observatory

In 2019 the KAGRA observatory and the KAGRA collaboration performed installation and commissioning work to attain the minimum required sensitivity to join O3. The KAGRA interferometer achieved the sensitivity of 1 Mpc in March 2020 and the first international collaborative observation with GEO 600 was carried out in April 2020. The KAGRA observatory is working on several improvements in the 2020-2021 timeframe with the goal of enhancing the sensitivity for O4.

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 Detector Characterization (DetChar) 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 two core activities: 1) preparing for future observing runs and 2) supporting the upgrade and commissioning of the detectors during commissioning breaks.

In preparing for O4, 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 contributed to 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 to maintain and extend of 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.

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.

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 is calibrated into dimensionless strain on the detectors before use with astrophysical analyses [1, 2]. The process for doing so requires detailed modeling of the feedback control system and the interferometric, opto-mechanical response of the detectors [3]. Some model parameters vary slowly with time throughout operation of the interferometer which must be monitored and, when possible, corrected [4]. All detectors continue to use photon calibrator systems [5, 6, 7, 8] 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 and KAGRA plan to employ a secondary reference of gravitational (Newtonian) systems to enhance and compliment the primary systems [7, 9, 10].

Strain is produced in near real-time (low-latency) using methods such as those described in [1, 2]. The level of systematic error in the strain data is 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 is confirmed to be within acceptable limits, but the accuracy and precision at all times remains relatively uncertain in low-latency due to evolving detector configurations, computational flaws, and unexpected events. Later, if necessary, the data is re-calibrated offline, with models updated and informed by careful analysis of all measurements taken throughout the observational period [11]. Such data sets have reduced systematic error and well-quantified uncertainty but at the sacrifice of weeks-to-months of latency.

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 statevector 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.
- Continually maintain documentation of methods and results, review final calibration products and software, and publish calibration results.

The tasks that the calibration will be working towards with longer term R&D are described in the Observational Science and Instrument Science white papers in more detail, but are briefly summarize here:

- Improve the understanding of the detectors response above 1 kHz.
- Continue investigations of any potential systematic error in the overall scale of the calibration.
- Understand the impact of calibration systematic error and uncertainty on astrophysical results including detection and astrophysical parameter estimation.
- Develop infrastructure for low-latency systematic error estimation.
- Continue efforts towards simultaneously reducing systematic error and latency in h(t) production.
- Advance and improve the low- and high-latency calibration software production tools.

1.4 Low Latency

The role of the low-latency subgroup is prompt identification and dissemination within the transient astronomy community of gravitational wave detections during real-time analysis of the interferometer data. In the third observing run (O3) the LIGO-Virgo collaboration has successfully provided open public alerts (OPA) to the transient observing communities. This effort has led to generation of numerous GCN alerts, and updates to these alerts. We have fairly consistently sent out automated alerts (GCN notices) within seconds of the detection. The notices provided valuable information for astronomers in the form of estimates of false alarm rate of the event, real-time sky-localization, source-classification, and source-characterization. In O3, we have also provided an infrastructure for identifying coincidence of gravitational wave triggers with gamma-ray bursts and high-energy neutrino detections, and sharing them with the community in real-time. In the ensuing fourth observing run by the LIGO-Virgo-Kagra collaboration we will be making several changes in the low-latency infrastructure that will (a) improve the efficiency of alert generation, (b) give us the ability to include more detectors, and (c) provide refinements of our data-products. We delineate the main goals of low-latency development for O4.

1. Integration of KAGRA into the low-latency framework.
2. GWCelery developments for further automation of tasks, and to improve latencies and monitoring of performance.
3. GraceDB development, improvement of latencies, further integration into cloud-native system.
4. CBC injection campaign to determine a uniform standard for readiness of pipelines (detection, parameter estimation, inference) before they can be approved for running in production. This can be done by conducting an injection campaign, which all pipelines interested to participate in low-latency operation will have to sign-up for. The injection campaign will also be used to train pipelines that requires injection campaigns for operation like p-astro and EM-Bright.
5. Pipeline automation for nominal BBH events: With an order of magnitude more events than O3, it is important that the vast majority of BBH events that we expect during O4, are analyzed and released for the public in an automated fashion. This will require automation of preferred-event-selection, detector characterization, vetting, preliminary and initial notice and circular generation and dissemination.
6. Integration of PE
7. Integration of different flavors of source-classification
8. Integration of fast PE data-products to low-latency data-products
9. Alert generation and GCN
10. Preferred event definition

1.5 Computing and Software

We successfully provided sufficient hardware resources and cyberinfrastructure during O3 to generate and distribute all 80 transient gravitational wave alerts in low-latency to the broader research community (https://gracedb.ligo.org/s
This plan supports the increased computing demands of O4 while continuing to provide sufficient computing resources for the offline deep searches of O3 data, and detailed follow-up studies of O3 transient events.

In addition, we provide a growing set of critical computing services in support of internal collaboration functions.

PLANS FOR O4:

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, 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.
3. **Define Joint LIGO Lab + LSC + Virgo + KAGRA (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.

We have started developing a new bottom-up estimate of the FTE requirements and WBS for this integrated computing plan, but have not completed that yet. The scope of this new WBS will incorporate all IGWN data analysis and collaboration computing, including:

- The entire scope of the 2016 LIGO Data Grid proposal (NSF Grant No. PHY-1700765)
- IGWN computing that has emerged post-2016 to support the era of plentiful detections and low-latency public alerts of gravitational wave transient events.
- All LIGO Laboratory WBS elements supporting LSC data analysis and collaboration support computing.
- MOU commitments of LSC groups relating to computing.
- AdVirgo Computing and Data Processing Infrastructure WBS.
- Updated policies, management processes, and forward-looking R& D to continue to adapt to the current era and prepare for the next 2-5 years of GW computing.

Out of scope 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.
- General computing and information technology (IT) for internal LIGO Laboratory, EGO, and KAGRA business functions.
- Individual data analysis pipeline development (but we include multi-pipeline infrastructure, optimization consulting, etc.)
- GW computing for third-generation (3G) gravitational-wave observatories (2030-).
- 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).

1.5.1 KAGRA

The KAGRA instruments records the strain induced by astrophysical gravitational wave signals or noises as a time-stream sampled at 16,384Hz and associated state information. They are transferred in low-latency from the observatory site in the tunnel at Kamioka to the KAGRA Main Data Server (KMDS) at Kashiwa. The raw data including the strain data and auxiliary information from a large number of sensors and digital control systems are recorded at a rate of about 20 MB/s. The latency of 1-second frame data from the laser interferometer at Kamioka site to KMDS is about 3.5 seconds, including not only for the transfer speed but also the calibration filter process of strain reconstruction. KMDS is continuously receiving data of LIGO and Virgo. These data are stored and analyzed on KMDS. The KAGRA data management subsystem and the KAGRA Computing and Software Working are jointly dedicated to the development and maintenance of an infrastructure for data transfer and analysis.

Highest priority tasks

- guarantee adequate storage and computing resources at Kashiwa, for commissioning, detector characterization and low-latency searches
- guarantee reliable bulk data transfer to storage in Tier-0 KMDS and Tier-1 Computer Centers in Taiwan
- guarantee reliable storage and computing resources for off-line analyses in KMDS

High priority tasks

- the development and maintenance of KAGRA data analysis software, KAGALI (KAGRA algorithmic library)

1.6 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.

1.7 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.

1.8 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.

1.9 LIGO A+ Upgrade

The “A+ detector" project is a major upgrade to the existing Advanced LIGO detectors, beginning in 2019 and expected to continue through the end of 2023. The A+ project in 2020-2021 will be carried out in parallel with preparations for the O4 observing run.

1.10 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 (ObsOps)

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. Planning is now in progress to make a phased return to operation. Initial steps will focus on activities that can be accomplished with "social distancing" and a minimum number of personnel.

As a first step the observatories will recover the low-noise operation of the interferometers in order to assess their condition [1]. For the remainder of the mid-2020 to mid-2021 period covered by this paper the LIGO Laboratory will undertake a series of upgrades in order to prepare for the fourth Observing run (O4). These include additional stray light control baffles to reduce noise from scattered light; high-efficiency output Faraday Isolators; a new high power laser amplifier; and Test Masses free from point absorbers. The most ambitious undertaking will be the addition of frequency-dependent squeezing, which will necessitate significant modifications to both the LIGO facilities. Prior to the pandemic the plan had been for a shutdown of 20 months with O4 beginning in late 2021 or early 2022. The exact impact of the pandemic on this schedule remains to be seen.

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

[1] The Hanford detector was affected by a 6.5 magnitude earthquake near Challis, Idaho about 400km SE of the facility. Recovering from this earthquake may require 2-3 weeks of extra effort compared to the recovery at Livingston

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).

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.

Vg-2.1 VIRGO

In the remainder of the document, VIRGO denotes the Virgo Collaboration and the European Gravitational Observatory Consortium (EGO).

Vg-2.1.1 Operations

The operations of the EGO infrastructure and of the Virgo detector, such as the activities of the Virgo Collaboration as a whole, were affected by the COVID-19 pandemic. The O3 run was suspended at EGO on March 27, 2020, simultaneously with LIGO, while many institutes were getting shut down all across Europe. On March 31, the Virgo detector was mostly turned off and put in a standby configuration, maintaining critical systems on, with a minimum of staff working on-site. Work feasible remotely kept on on a best effort given the current circumstances while all lab activities were effectively stopped for weeks.

With the relaxation of the lockdown following the improvement of the situation in Europe in general and in Italy in particular, the EGO site reopened on May 4, enforcing social distancing and safety rules according to the regulations in force. Similar evolution is occuring in European countries, at different paces depending on the local evolution of the pandemic. As of mid-June 2020, laboratory workshops are becoming available again and work-related travels start being possible as well.

Prior to the pandemic, the Virgo plan was to have a 20-month shutdown between the end of O3 and the beginning of O4 during which the Advanced Virgo Plus phase 1 upgrade would have taken place. The goal remains the same but the actual implementation of the plan will be different – see Section Vg-11.1: work is in progress, within Virgo and in coordination with LIGO and KAGRA, to assess the impact of the pandemic on the upgrade and to update the schedule.

During the AdV+ Phase 1 upgrade, each piece of equipment, new or upgraded, will be commissioned promptly by the Commissioning Team when available. A milestone will be the first long lock of the upgraded instrument, followed by noise hunting and improvements campaigns, in order to achieve the performance goals for O4 in terms of sensitivity, duty-cycle and stability. In parallel, DetChar, Calibration, Low-Latency, Computing & Software will upgrade, improve and expand their systems with O4 as target. These groups will support commissioning activities as well according to people availability. Efforts will be put on strengthening connections between groups and establishing joint priorities to optimally use the available resources.

Vg-2.1.2 Commissioning

Commissioning the Virgo interferometer means finding its best working point, understanding the limitations of the sensitivity and improving the performance of the instrument (in particular its duty cycle and stability). Both hardware modifications of some components of the instrument and of the control software can be needed to reach this target. Commissioning is carried out by the commissioning team, a transversal group composed by people coming from the different labs of the Virgo Collaboration, including EGO. During the commissioning period, visitors from LIGO and KAGRA are welcome at the Virgo site to contribute to the activities. Also, other teams, such as DetChar, Calibration, Low-Latency, Computing & Software will support Commissioning activities. Efforts will be made on strengthening connections among these groups.

Activities scheduled during a given week are managed by a "weekly-coordinator" (a Virgo instrument expert), together with the commissioning coordinator. Day-to-day activities are discussed and decided in dedicated short daily meetings, while a longer meeting is held once a week to discuss last results and to organize mid-term activities. The overall commissioning strategy is proposed by the commissioning coordinator, discussed among the Virgo coordinators and endorsed by the Virgo Steering Committee. In the preparation of an observing run, Commissioning is tasked with bringing the detector configuration to a state that matches the run goals in terms of sensitivity and duty cycle.

The beginning of the next observing run (O4, joint with LIGO and KAGRA) is scheduled in 2022 at the end of the implementation and the commissioning of the various upgrades planned for Advanced VirgoPlus phase I. During the period 2020-2021, the completion of each upgrade will be immediately followed by the commissioning of the corresponding equipment. At the end of the installation phase, an intensive period of commissioning, fine tuning and noise hunting is planned to reach the sensitivity target defined for O4.

Op-2.2 KAGRA Detector Operations

In 2019 the KAGRA observatory and the KAGRA collaboration worked on installation and commissioning to attain the minimum required sensitivity of 1 Mpc in order to join the O3 run. The KAGRA interferometer achieved this sensitivity with an interferometer configuration of Power Recycled Fabry-Perot Michelson with DC read-out in March 2020 and the first international collaborative observation with GEO 600 was carried out in April.

The KAGRA observatory is planning several activities in 2020 and 2021 to enhance the sensitivity prior to O4. The work includes improvements of the mirror suspension systems, installation of optical baffles, and exchanging of cryo-coolers. We estimate that these projects will take about half of a year after which the commissioning team will take over.

3 Detector Characterization (DetChar)

The LIGO and Virgo detector characterization (DetChar) 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 group, and the calibration group. 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 2020-2021 in terms of preparation for **O4**, or tasks necessary for the fourth LIGO-Virgo-KAGRA observing run, and **future observing runs**, or tasks required for the success of the next observing run and beyond.

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 the 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 vetoes, or exploratory research and development which is not yet certain to have a direct application to the central DetChar scope.

Op-3.1.1 O. Preparation for the fourth aLIGO observing run - O4

O.C. Prioritized list of central tasks

Highest priority tasks

O.C.1: DQ products for the astrophysical searches

- O.C.1.1. Developing, testing, and documenting offline DQ flags.
- O.C.1.2. Producing offline veto definer files, which define the DQ flags that are used to veto data for each individual astrophysical search.
- O.C.1.3. Reviewing offline veto definer files.
- O.C.1.4. Noise line studies to characterize line artifacts that impact searches for long duration gravitational waves.
- O.C.1.5. Support of online DQ products delivered with low latency h(t) frames; state information, data quality flags, iDQ timeseries, and veto definitions.
- O.C.1.6. Review of DQ products used to generate and evaluate DQ flags.
- O.C.1.7. Post-run investigation of worst offender noise sources. Particularly:
  - Scattering
  - Blip glitches
  - Mid-frequency noise (e.g. 60-200 Hz non-stationary 'blue mountains')

  - Anomalous environmental coupling (e.g. airplanes, thunderstorms, periscope motion, beam clipping)
  - Spectral Artifacts
  - *See LIGO DCC P1600110 for more information on the noise sources.*

- O.C.1.8 Produce h(t) frames that have noise contributions subtracted, when noise sources with auxiliary witnesses exist.
- O.C.1.9 Produce gated h(t) frames with loud short-duration glitches subtracted used for persistent GW searches (CW and stochastic).

O.C.2: Vetting GW event candidates

- O.C.2.1. Using the automatically produced Data Quality Report (DQR) [12, 13] to vet the data quality around gravitational wave event candidates, including evaluating environmental couplings.
- O.C.2.2. Field a DetChar rapid response team to vet the data quality for low-latency gravitational wave candidate events.

O.C.3: Monitoring known or new DQ issues

- O.C.3.1. 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.
- O.C.3.2. Mentoring and training scientists participating in the LSC fellows program, which supports LSC scientists working at the site to improve the detector data.
- O.C.3.3. Maintenance and characterization of the Physical Environment Monitor (PEM) sensors, including evaluating the environmental couplings of the interferometers.

O.C.4: Commissioning support

- O.C.4.1. Tracking issues that affect interferometer uptime, such as seismic motion.
- O.C.4.2. 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.

O.C.5: Maintaining key tools

- O.C.5.1. 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.
- O.C.5.2. Review of key tools, particularly those used to generate vetoes.
- List of key infrastructure and tools:
  - Fundamentally necessary services:

    * The summary pages [14]; 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) [15], including the low-latency DMT DQ vector infrastructure
    * The segment database [16, 17]; which stores state and DQ flag information used by the astrophysical searches
    * Omicron triggers [18], which identify transient noise triggers, including in low-latency, delivered with very high reliability
    * Data Quality Reports [12, 13], 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 [19], [20], 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 [21]
    * SNAX toolkit [22]
    * GWpy [23]
    * GW-DetChar [24] (which contains Omega scans, LASSO, automated monitoring of scattering, ADC/DAC overflows and software saturations)
    * GWSumm [25] (which is used to generate the summary pages)
    * Stochmon [26]
    * STAMP-PEM [27]
    * Hveto [28]
    * ligoDV web [29]
    * Channel Information System [30]
    * Pointy Poisson [31], [32]
    * VET [33]
    * Offline noise subtraction code (e.g. [34], [35], [36])
    * Suite of remote access tools (remote MEDM and EPICS [37], remote DataViewer [38]
    * LigoCAM [39]
    * Automated safety studies [40], [20]
    * GravitySpy [41]
    * FScans and dependent programs / scripts (spectral ratio, comb tracker)[42]

  This list relies on software dependencies maintained by the LIGO Laboratory (e.g. the Guardian [43],), 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

O.C.6: Interfacing with commissioners, site staff, and search groups

- 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.

O.RD. Prioritized list of critical Research and Development tasks

Highest priority tasks

O.RD.1: Investigation of the search backgrounds

- O.RD.1.1. Studying how instrumental artifacts affect the sensitivity of a specific search or search method.
- O.RD.1.2. Developing search-specific techniques for noise mitigation. This includes contributions to developing incorporating low-latency DQ information into the online searches.
- O.RD.1.3. 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.

O.RD.2: Machine learning for O4

- O.RD.2.1. 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.
- O.RD.2.2. Machine learning [44] classification, regression and data mining studies targeting known noise sources, e.g. scattering.

O.RD.3 Gating requirements for the astrophysical searches

- O.RD.3.1. Contribution to a study team to evaluate the gating requirements across all astrophysical searches and recommend potential common solutions (if any)
- O.RD.3.2. Development of low-latency self-gated h(t)

Op-3.1.2 Future observing runs

F.C. Prioritized list of central tasks for future observing runs

Highest priority tasks

F.C.0: Organization of O4 planning and development

- F.C.0.1. Contributions to design and evaluation of the coherent plan for O4 DQ products and DetChar services
- F.C.0.2. Contributions to organization and upkeep of the LIGO DetChar workplan

F.C.1: Automation of key tools

- F.C.1.1. Automation of DQ veto performance testing.
- F.C.1.2. Development of DQ products necessary for fully-automated EM alerts.
- F.C.1.3. Review and further automation of results displayed within the DQR, or 'DQR tasks', with the end goal of full automation (no or limited human intervention required for interpretation)

F.C.2: Characterization of interferometer and auxiliary channels during A+ installation and for future observing runs

- F.C.2.2. Documentation of planned or newly installed interferometer subsystems and environmental monitors.
- F.C.2.3. Maintenance of lists of auxiliary channels useful for DetChar studies to include new subsystems and environmental monitors.
- F.C.2.4. Maintenance of summary page content to include new subsystems and environmental monitors.
- F.C.2.5. Signal fidelity studies of newly installed auxiliary channels.
- F.C.2.6. Auxiliary channel safety studies for new subsystems and environmental monitors.
- F.C.2.7. 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.

F.C.3: Improve monitors of known DQ features

- F.C.3.1. 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).
- F.C.3.2. Improving monitors for excess mirror motion leading to scattered light.
- F.C.3.3 Schumann resonance studies.

- F.C.3.4 Develop tool to query stochastic monitors to find which auxiliary channels are coherent with the gravitational wave strain data at a given frequency.
- F.C.3.5. Optic suspension resonance 'violin mode' monitoring.

F.C.4: Curation of DQ information for public data releases GWOSC

- F.C.4.1. Curation, documentation, and review of DQ vetoes for release by the GW Open Science Center (GWOSC) [45].
- F.C.4.2. Development and documentation of the "Detector status" public summary pages hosted by the GWOSC [14]

LT-3.2 LIGO Detector Characterization Long Term Plans

High priority tasks

F. C. 5: Predict detector performance based on instrument state using machine learning

- F.C.5.1. Lock loss prediction
- F.C.5.2. Prediction of noise characteristics or other detector performance metrics

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

Highest priority tasks

F.RD.1: Develop improved clustering for Omicron

- F.RD.1.1. Improve Omicron clustering scheme to more accurately report timing, frequency, SNR of excess power.

F.RD.2: 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 [46] and will be discoverable with common tools (see key tools listed in Section O.C.5). 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

F.RD.3: Quantify the impact of transient noise on parameter estimation

- F.RD.3.1 - test the effects of transient noise on recovered source properties

- F.RD.3.2 - develop and test methods to reconstruct and remove from h(t) isolated glitches and other noise types without auxiliary witnesses

F.RD.4: Research and development of new methods for noise identification/mitigation
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

F.RD.5: 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 O.C.5. 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 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. 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.

Vg-3.2 Virgo Detector Characterization

Three main areas have been identified by the Virgo DetChar group for the period going from the end of the O3b run to the beginning of O4. Each points to a different time period (post O3, O3-O4 upgrades, O4 and beyond) and, while they are listed in chronological order below, they should all proceed in parallel. Moreover, they are often interlinked: regular meetings, reviewing progress or difficulties and re-assessing priorities, will take place to make sure the whole project moves forward and that the Virgo DetChar group is 'ready for battle' for the next data taking periods. For each line, tasks have been divided into two categories: *key tasks* that are necessary for the O4 run and *R&D tasks* that are desirable but may take longer to complete or which scope is likely to evolve as the DetChar group receives inputs from the instrument on the one hand, from the data analysis group on the other.

Vg-3.2.1 Completion of O3 tasks and lessons learned

The O3 run has been the first long data-taking period for Virgo in the Advanced detectors era. Therefore, getting the most out of the accumulated dataset while building on the experience acquired in various areas, are among the priorities of the Virgo DetChar group for the year to come.

- **Key tasks**
  - Offline flags and vetoes
    Science segments and CAT1 vetoes are key DetChar inputs to the offline analysis: these tasks have been completed months ago for O3a and will soon be for O3b as well. During O3, work has been done to partially automate the production of these deliverables and to ease the monitoring of these tasks, such as the review of their output, by producing simultaneously a set of summary plots. The associated goal is to reduce the latency of the whole progress. More work on that topic will be performed during the O3-O4 upgrade shutdown.
  - Spectral lines catalog
    The catalog of lines from a continuous data-taking period (for instance O3a and O3b must be dealt with separately as many lines appear to be specific to one of the two sub-runs) is needed for analysis from the CW and stochastic groups. Work is currently (June 2020) in progress to complete this catalog, using inputs from the NoEMi tool and from preliminary analysis of the Virgo O3 data by the CW group; this is a joint project between the DetChar, data analysis and environmental noise monitoring groups. In addition to listing the different lines – classified by strength/persistance/presence – finding their origin is a challenge, in order to identify and remove the data analysis candidates that are of instrumental origin.
  - Data quality for the O3b / full O3 event catalogs
    The LIGO-Virgo event catalogs contain both triggers indentified in low-latency and events that have been identified offline by rerunning the latest version of the analysis software onto data that may have benefited from an improved calibration. Both types of events need to pass data-quality vetting in order to be published. This is done using primarily the same framework, the Data Quality Reports (DQRs), that has been developed for O3 and used extensively during the run to process automatically a set of data quality checks in response to interesting triggers. The Virgo DQRs run on the EGO HTCondor farm and the results, produced following the standards defined in the LIGO data-quality-report package, are copied to GraceDB when available.
  - Mining the available data
    O3 data are a mine of information about the Virgo detector, the 11-month data taking and the overall quality of the data. With the limited personpower available during the run – see below for more details – it has not been possible to carry out all needed studies on top of the mandatory tasks performed by the group. Therefore, we should try to use the time available between O3 and O4 to do more of these analysis – personpower-permitting of course.
  - Producing information for public data releases
    Data quality information and relevant monitoring quantities (duty cycles, etc.) are needed to accompany the release of parts of our dataset to the scientific community and the public.
  - Addressing personpower issues that have limited the number of tasks the Virgo DetChar group could focus on during O3
    There is a significant discrepancy betwee the public-facing side of the Virgo DetChar group – a successful O3 run, with all its main tasks completed without major failure nor delay – and the way the group has worked internally, with too few people overcommitted. This organization is not sustainable in the future and will have to change significantly: more personpower and more commitments – individually or at the level of Virgo groups – are mandatory. That issue has been identified ahead of O3 but there was no time nor ressource to fix it at that time or while the data taking was ongoing. Thererefore, now is the right time to tackle this issue and to find out solutions that will last until O4 and beyond. Work has started following a bottom-up approach, from the DetChar group to the Collaboration and its management. Although delayed by the

covid-19 pandemic, this retrospect analysis and the associated conclusions drawn from it are the priority of the group for the months to come.

- **R&D tasks**
  - Use O3 data to test and validate new or updated tools developed for the AdV+ phase 1 upgrade or for O4. The large O3 dataset, which data have been taken in numerous different conditions, is the natural playground to test and validate new code developed for the AdV+ upgrade and the following runs.

Vg-3.2.2 Support of the AdV+ phase 1 upgrade and contributions to the characterization of the instrument and to the improvement of its performance

While O3 has been a success, Virgo is expected to reach a new level with the two phases of the AdV+ upgrade. The first one will take place in between the O3 and O4 runs and will include the installation of signal recycling, of a high-power laser and possibly of frequency-dependent squeezing. The Virgo DetChar group is committed to support the detector team during the commissioning and noise-hunting phases that will follow the completion of the hardware upgrades and the successfull locking of the new dual-recycled configuration. A key part of this work will be to define priorities, jointly with the teams working on the detector upgrade and its commissioning.

- **Key tasks**
  - Support of commissioning activities
    As exemplified by the O2 and O3 preparation periods, commissioning is a key phase of any upgrade project as the future detector performance during the data taking (sensitivity, duty cycle and stability) is driven by the results of the commissioning. The Virgo DetChar group will contribute to this effort with all its available resources at the time.
  - Glitches: rate, families, classification, search for their origin and help to mitigate/fix them
    Transient noise bursts, known as glitches, impact the search for gravitational waves in many ways: sensitivity drops, unstabilities that could lead to a lock loss, false alarm triggers for the searches, etc.
  - Investigate spectral lines and noise bumps
    Spectral lines (constant in frequency or moving with time, permanent or present from time to time) and bumps (broader peaks in the sensitivity curve) are important features that must be investigated by the Virgo DetChar group, jointly with the commissioning team.
  - Reinforce instrument <-> DetChar liaison
    A first attempt was made during the final months leading to O3 and at the beginning of the run but that link stretched thin with time, in particular because people were overcommitted. In addition to daily contacts (with cross-participation to technical meetings), dedicated joint meetings should be organized every few months (following the successful model of the LIGO DetChar f2f meetings, organized twice a year) such as commong training sessions, both for newcomers looking for topical introductions and for experts willing to follow the latest developments.
  - Tools
    Depending on each tool's maturity, usage and usefulness, several tasks could be performed: upgrade, automation, maintenance or documentation.

- **R&D tasks**
  - Developing new tools
    Depending on available personpower and on the priorities set jointly with the commissioning team on the one hand and with the low-latency group (including pipelines) on the other, new tools could be developed, either by the Virgo DetChar group alone, or in collaboration with other Virgo groups. This could include adapting tools developed originally from LIGO or KAGRA.

Vg-3.2.3 Preparation of O4 and following runs

The reason for being of the tasks listed above is to ensure that the Virgo DetChar group is ready to embark onto the O4 run – that will last at least as long as O3 – and, more generally, onto the following data-taking periods. Specific developments are needed to meet this goal, both on the technical side and also to the way the group is organized and interact with other teams: instrument, observational science and the LIGO and KAGRA DetChar groups.

- **Key tasks**
  - Development of a Virgo DQR 2.0
    The Virgo DQR has proven very reliable during O3; in particular, it has provided quick inputs to all the alerts that the LIGO-Virgo rapid response team had to deal with in near real time during the data taking. Moreover, the framework has been significantly improved during the run, with the addition of several new checks that have been validated before being added to production. Yet, a significant amount of work is needed before O4 to make the code more professional and likely more performing. Moreover, this rewriting could lead to a larger automation of the trigger validation: currently, the majority of checks requires a human input. This twofold goal is motivated by the prediction that the number of alerts and thus detections will increase significantly as the sensitivity and duty cycle of the instruments improve. A skilled software developed should be part of that project.
  - Reduce latency of online flags and vetoes production
    In collaboration with the DAQ team and the calibration/h-reconstruction experts.
  - Full automation (and latency reduction) of the production of offline flags and vetoes
    This is needed to improve and fasten the feedback from pipelines and analysis groups.
  - Deal with the new storage architecture defined at EGO
    DetChar products will have to move from a small high-reliability online area where they are produced to a large offline area where they will be stored permanently. This migration should be automated and transparent for users accessing such products. Procedures and tools of generic use will have to be developed with the EGO computing department and other groups interested in benefiting from the same functionalities.
  - Definition of a new DetChar group organization, targeting data-taking periods: factory-like and operations-oriented
    O3 has taught us that the organization of the Virgo DetChar group should be different during runs. Experts should be rotating to alternate on-duty periods (during which they will be the primary contact for the tasks they are responsible for) and off-duty periods, useful to recuparate and to review ongoing actions taking a step back.

- Definition of a new DetChar shift organization embedded in the global Virgo service task framework
  DetChar shifts had been correctly identified prior to O3 as a key component of the strategy of the group. Yet, practical difficulties have reduced the impact of shifters: these should be addressed and fixed for O4 and beyond. There should probably be less individual shifters but they should contribute more to DetChar activities – 2 shifters / week seems a good conpromise given the need to have a 24/7 coverage and the limited personpower available. Having shifters more involved into the group will have several advantages, among which the possibility for experienced shifters to become mentors and train newcomers and the hope that some shifters could become long-term members of the DetChar group once their shift duties are over. A Virgo DetChar shift manager will be needed for each data-taking period, starting from O4.
  One issue during O3 was that the DetChar shifts were singled out as the only collaboration-wide service task. Groups should be able to choose to contribute to DetChar shifts (or other DetChar tasks) among a wide list of service tasks, rather than being forced to contribute to them in proportiong to their size. The list of service tasks such as the mechanisms ensuring that all needs identified as key for the running of the experiment are fulfilled will have to be defined by the Virgo collaboration and its management.
- Construction of a core team of DetChar experts, active in the last period before O4 and during the whole run
  Another lesson from O3 is that the only sustainable organization of the Virgo DetChar group is to rely on a set of people that (i) are numerous enough to allow a continuous coverage while preserving good working conditions, (ii) have been around for enough time to have become experts and (iii) can keep that commitment during a long period – ideally from a few months before the beginning of the new data-taking period until its final end. A handful of people, post-docs or staff, commited to the Virgo DetChar at the 0.5 FTE level at least, will be needed to fulfill this goal.
- Establish permanent data analysis groups <-> DetChar liaisons; organize regular meetings
  During O3, there has not been enough connections between the Virgo DetChar group and the data analysis groups (pipelines nor offline analysis). This must be improved for O4, with the need to get feedback from the observational science teams about the data segments that are bad or of limited quality and about the factors that limit the sensitivity or the accuracy of the analysis. Without this information, it is difficult for the DetChar group to assess priorities and separate what is important from what is not.
- Liaison Detector <-> DetChar <-> Data analysis to orientate and motivate investigations at the DetChar level
  As explained above, the Virgo DetChar group must be the link between the instrument and the data analysis, helping to improve the transmission of information from one end to the other, such as optimizing the resource available to study the global performance of Virgo.
- Improve monitoring
  Monitoring Virgo from the data taking to the vetting of the events found by the searches, both online and offline, is another duty of the Virgo DetChar group. In addition to providing as much information as possible with a latency as low as possible, monitoring is the first level of documentation that grows as the run goes on. Information must be complete, clear and easy to find. A lot of work has been done before and during O3 to develop and improve this framework. More will be needed before O4: topics include the Virgo overall performance, the impact of the environment on the instrument and how level the low-latency infrastructure is running.
- Strengthen connections with LIGO and KAGRA DetChar groups

Th exchange of tools already exists – for instance Omicron or bruco. Moreover, the different DetChar groups should consider the possibility to adopt or review common data format. Efforts should be made to share common software environment as well.

- **R&D tasks**
  - Help improving search sensitivity
    Searches should tell the Virgo DetChar group which improvements are needed – if any given the existing differences between the detectors: sensitivity, glitch rate, etc. If some searches require DetChar inputs, the corresponding data quality products should be defined jointly and regularly reviewed.
  - Curation and improvement of existing documentation
    These include tools and training sessions.
  - Explore possible convergence between Virgo, LIGO and KAGRA DetChar software
    Using frameworks as close as possible would be ideal. Yet, the collaborations have developed software independently for years and each one has a unique knowledge of its data and instrument(s). Therefore, there will always be some differences between software packages. Moreover, having time to study and compare various codes, possibly developed fully by another collaboration, would require some additional personpower as these activities have no immediate impact on the group and should thus proceed parallel to the development of the current default software tools.
  - Machine learning
    This is a broad area developing quickly. R&D projects are welcome, provided that they aim at providing tools useful during data-taking periods and for DetChar-related analysis.

Kg-3.2 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-3.2.1 GEO-KAGRA joint observing run - O3GK

Providing Data Quality (DQ) information

- Providing and managing online DQ state vector
- Providing and managing DQ segments

Detector noise monitoring and evaluation

- Providing glitch trigger information
- Providing veto information for GW searches
- Improving the sensitivity curve by extracting noises using the independent component analysis

Kg-3.2.2 Preparation for future observing run

Providing Data Quality (DQ) information

- Improving online DQ state vector
- Improving DQ segments
- Providing data quality report (DQR)

Detector noise monitoring and evaluation

- Developing and automating glitch trigger generation
- Providing and managing spectral line information
- Providing lockloss information
- Developing auxiliary channels correlation search pipelines
- Investigating safe/unsafe channels
- Developing and automating veto analysis pipelines
- Developing noise subtraction method
- Developing noise classification by using machine learning
- Developing and applying methods to remove nonlinearly coupled non-Gaussian noises through the independent component analysis

Support commissioning test including remote site work

- Providing the summary pages and improving contents of the summary pages
- Providing and managing auxiliary channel information
- Developing the method for easy data access
- Providing following tools and information for commissioners and data analysis group
  - Summary pages [**?**]
  - Bruco [47]
  - Fscan [42]
  - Omicron triggers [48]
  - Hveto [28]
  - Omega-scan [49]
  - NoEMi [50]
  - Channel Information System [30]
- Developing new tools
  - Pastavi; WEB base data plot tool [51]
  - Noise budgetter; python base noise budget management tool [52]
  - CAGMon: (Non-) linear noise propagation monitor using multi-channel correlation analysis [53]
  - EtaGen: Event Trigger Generation Tool using Hilbert-Huang Transform [54]
  - Independent Component Analysis; Multi-channel correlation analysis [55]

Contribution for reliable detection

- Participating RRT
- Providing DQ information for RRT
- Providing veto information

4 Calibration (Cal)

In the time between O3 and O4, it is imperative that each LIGO, KAGRA, and Virgo calibration groups remain active and vigilant with an eye towards O4. To do so, the groups must maintain an active role in Observational Science, Instrument Science, and with their structures of Operations. 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. Content in all white papers are intended to be inclusive of, and generically representing, the LIGO, KAGRA, and Virgo group activity, as a unified approach is necessary to be ready for O4.

Op-4.1 Calibration Operational Plans

Hardware Maintenance: 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. In addition, recent collaboration wide R&D efforts have proven gravitational (Newtonian) calibration systems are viable complements to the photon calibrator systems. Both Virgo and KAGRA intend to work towards bringing their Newtonian calibrator systems to operational quality by O4.

Detector Modeling and Systematic Error / Uncertainty Software: 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 between O3 and 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).

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 [11]. However, work will now be done 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. Leading up to O4, LIGO intends to coordinate with the KAGRA and Virgo teams 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.

Regular Calibration Measurements 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 and KAGRA) are used to obtain these measurements.

Strain Production Software:. 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 [1]. The Virgo pipeline is in `C` [2]. 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. 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. 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. This includes a continued interactive relationship with the gravitational wave open science center.

Monitoring:. 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. In addition to working with the software package development teams to maintain those detectors' general summary pages, the calibration groups 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.

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.

Documentation and 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.

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 [56]. Additional, organizational content may be found on the ligo.org wiki page [57].

LT-4.2 Calibration Development Plans

Improvements to the DARM loop and the characterization of Unknown Systematic Error Therein. 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. 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,
- to reduce now-known, controllable, systematic error in the DARM via refinement of the interaction between other detector control loops,
- to include other now-known, but uncontrollable systematic error in the loop models a priori, and
- to refine uncertainty estimates of unknown systematic error.

Improvements in low-latency calibration software. 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. While each LIGO, Virgo, and KAGRA detector has near-real-time processing code compiled from `C`-code, some low-latency post-processing will always remain necessary due to the causal limitations of each real-time system. Each group 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

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 near-real-time and low-latency software, and Virgo intends to explore options for updating its software packages.

Incorporating calibration error and uncertainty with astrophysical results:. 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. The remaining development plans described below are based on assumed needs to-date.

Developing infrastructure for low-latency systematic error estimation: The LIGO, Virgo, and KAGRA calibration groups will explore and develop methods for determining accurate calibration systematic error estimates in low-latency. One possibility is to follow the method used by the Virgo calibration group of using many low-SNR calibration lines. These lines could then be used to construct a complete systematic error estimate across frequencies in low latency.

Calibration above 1 kHz, if necessary: 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.

Reduce systematic error, if necessary: 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.

Improve absolute reference hardware: 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 environment. As KAGRA and Virgo are leading the charge in secondary, gravitational (Newtonian) calibrators, LIGO hopes 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 (Tim)

Op-5.1 LIGO 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 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 subsystems of the detector. The timing distribution system's status is monitored continuously and is periodically tested in-depth via timing diagnostics studies.

Critical timing tasks include:

- verifying traceable performance of the timing distribution system,
- verifying the validity and accuracy of the recorded time-stamp,
- verifying the accuracy of the distributed timing signals,
- expanding the capabilities of data monitoring tools related to timing,
- availability of timing diagnostics for various subsystems,
- measuring and documenting the timing performance,
- reviewing the physical/software implementation and documentation of the timing distribution and timing diagnostics components.

6 Low Latency (LoLa)

Op-6.1 Overview and goals for O4

The low latency group provides support for conducting astrophysical searches close to real-time for the purpose of connecting gravitational-wave observations with multi-messenger ones as well as facilitating the diagnosing of detector problems 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 June 2020 to May 2021 year is predominantly driven by the needs for public alerts over the upcoming observing run of O4, due to commence in late 2021/early 2022 Aside from O4, and beyond, goals, we anticipate spending of O(1) month (primarily in June 2020) in order to enlist and diagnose shortcomings of the low latency alert infrastructure during the O3 observing run and additionally document and possibly publish in a journal publication the Early Warning experiment expected to take place in late May to early June.

The overarching goals of the low latency group over the upcoming 12 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 normal (non-pre-merger) real time (<10 second latency) alerts as well early warning (pre-merger) ones. As part of this goal, a thorough review of the fundamental (and non-) limitations of the existing system architecture will be required, including the publication and subscription technologies used and software framework. This effort has already started as part of reviewing the low latency operations over O3.
- Incorporate KAGRA data into the low latency operations as directed by the Operations Division/Joint Run-Planning Committee. 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-GW observatories, including interaction of offline searches and results (i.e., identified events and related data products) with event database (GraceDB). 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.
- Explore alert technologies (for public and private ones) beyond GCN (e.g. VRO/LSST). This public-facing functionality of the low latency system will require a study of the current options, especially with an eye for what other survey telescopes will be using in the 5-10 year horizon.
- Revisit individual tasks that need to be accommodated under the low latency pipeline and overall software infrastructure. Will need to work closely with the LVK search and detector characterization groups in order to identify needs and specifications. The review of O3 operations and existing system under which these are implemented will be the starting point.

- Examine and define the role of "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 and CBC groups, establish the testing and acceptance criteria for all low latency search pipelines and related tasks. This will include definition of O3 data replay and simulated data over which individual search pipelines will need to run in a timely fashion in order to instruct the analysis cuts to be used in low latency as well as in order to establish the final acceptance test of the end-to-end system.

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

- June 30, 2020: O3 review and lesson learned completes.
- August 31, 2020: analysis tasks as part of the low latency infrastructure and overall automation needs and specifications overlapping scope and operations of the LVK search and detector characterization groups identified.
- August 31, 2020: expectations and specifications regarding KAGRA inclusion into low latency operations fully identified.
- August 31, 2020: All input, specifications and requirements from run planning, data analysis groups and broader collaboration collected.
- August 31, 2020: System architecture design completes. Document is handed off to review.
- August 31, 2020: computing hardware and software infrastructure to enable simultaneous prototyping, testing, review and deployment is in place.
- August 31, 2020: Specifications on data challenges is to be commonly decided and rolled out.
- September 30, 2020: low latency design document completes review by independent experts.
- December 31, 2020: specifications on alert technology choice identified.
- May 31, 2020: delivery of O4 low latency alert system for end-to-end testing.

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 alerts, built from widely used open source components. Its responsibilities include:

- Merging related candidates from multiple online LIGO/Virgo transient searches
- Correlating LIGO/Virgo events with gamma-ray bursts, neutrinos, and supernovae
- Launching automated follow-up analyses including data quality checks, rapid sky localization, automated parameter estimation, and source classification
- Generating and sending machine-readable alerts
- Sending updated alerts after awaiting human input
- Automatically composing templates for any human-readable prose (e.g. "Circulars")

Rather than conduct a fresh top-down design requirements discussion like we did at the beginning of O3[2], we instead commission a team to study the O3 design design and O4 objectives in the following manner:

1. Describe and analyze the architecture that was adopted in O3.
2. Study the changes that were required during the observing run and the reasons for them.
3. Identify any deficiencies in the O3 architecture, and propose paths forward to address them.
4. Review choices of frameworks, tooling, and DevOps technologies.

The end result of this process will be a short DCC document (no more than about 10 pages).

Due to the scope and timescale of this whitepaper, this exercise cannot be done here. However, we can identify a short list of the highest priority goals now:

1. It is important to have full participation of stakeholders from LIGO, Virgo, and Kagra collaborations in all phases: design, maintenance, and deployment.
2. Establish latency and reliability service level requirements and put in place monitoring to track them.
3. Transition to high availability, redundant, service-oriented, cloud-native deployment.
4. Study next-generation streaming technologies as possible replacements for LVAlert, VOEvent, and internal pipeline glue.
5. Adopt agile practices for project management, software testing, and deployment, making full use of collaboration tools in GitLab.

Op-6.3 GraceDB and LVAlert in O3

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

[2] `https://dcc.ligo.org/LIGO-T1800068`

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 GraceDB in O3a

Prior to the start of O3, GraceDB was migrated from its VM hosts at UWM to Amazon Web Services (AWS) in order give the service room to grow and handle more internal users and events, as well as more external users and public alerts.

The service consisted of containerized Django GraceDB apps in a Kubernetes clustered configuration. The cluster lived across multiple availability zones (AZ) in the AWS /textitus-west-2 zone in Oregon, US. Given the opportunity loss of a missed electromagnetic (EM) targetted observation, the service provides high-availability that is robust against failures.

Data is served by Amazon's Extensible File System (EFS) multi-AZ network storage. The database backend is a MariaDB version of Amazon's Relational Database Server (RDS), also designed for high-availability.

Op-6.3.2 GraceDB in O3b

GraceDB's configuration in O3b was largely unchanged from O3a. The number of GraceDB worker nodes was modified, and were clusted in a Docker Swarm configuration. Improvements were made on the architecture and application layer to improve EFS and RDS throughput, latency, and reliability.

Op-6.3.3 GraceDB and LVAlert in O4 and Beyond

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.
- 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.

Op-6.4 Low latency review and testing

The review will be performed by the appointed review team on the main milestones, namely:

1. System architecture review and design. (due end of August, 2020)
2. Low-latency design document. (due end of September, 2020)
3. Specification of the alert technology document and publication of User Guide of the alert technology. (due end of December, 2020)

The results of this milestone will be documented and reviewed in order, and after the validation of the result of the previous milestone. Once all these documents are approved by the reviewers and the **collaboration**, the finalization of the low-latency system will start. A final review phase will have the main purpose of ensuring that the delivered system conforms to the procedure approved by the **collaboration**. It will provide the validation that the system meets the expected availability goal of 99% uptime. The review step will be completed before the start of O4. To fulfill this goal we will need 3+1 fully operational instance 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.

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 useful 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 production environment. This will allow any external users to continuously develop and test their listening and integration.

Our goal for O4 is to be able to easily instantiate 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, in the longer term we plan to apply the same deployment strategy also to the low-latency pipelines. In this way, it will be possible to automatically deploy self contained instances of the full bundle (services and pipelines) reflecting the code changes pushed to GitLab, thus facilitating the process of code testing and reviewing.

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. While the process for assessing how these tasks fared over the most recent observing run (O3) is still ongoing, the onset of work toward O4 and beyond should establish clear expectations and 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.
- Revisit the definition of “superevent” and its attributes, and the choice of “preferred-event” based on the O3 experience.
- 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.
- 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 $<$10s 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 assessing the O3 low latency operations, identify shortcomings and improvements needed. The process is expected to complete by August 2020 in order to allow implementation to proceed. As part of this, we call for the Burst and CBC groups, in coordination with the Low Latency group to define a dataset involving both real and simulated data that will be used 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 August 2020.

Op-6.6 EM-advocates

Brief review of the O3 EM follow-up organization. The EM advocate operations during O3 relied on a shift-based system, using a Google calendar to organize those shifts, to house the dates of given shifts (under the responsibility of the low latency chairs). The EM advocates coordinate with colleagues on the Rapid Response Team (RRT), which included assigned individuals from the detector characterization, parameter estimation, search groups, and others to support event characterization. The assigned EM advocate, prepared by a how-to user guide for first time advocates, benefited significantly from GraceDB-based phone and e-mail alerts to notify them in case of events. The preparation of pre-defined circulars for GCN-based preliminary, initial, and updates to send at pre-defined intervals also helped significantly. The tasks were supported by assistance on the call from the chair of the RRT as well as a phone book to call search experts as required. On a longer time scale, it also was useful to be part of the PE rota team, communicating on chat.ligo.org, to be part of the PE rota.

O4 EM follow-up organization. As EM advocates are a core component of the group's operations, we expect there to be a split of shifts between the collaborations, LIGO, Virgo, and KAGRA **(weighted by collaboration size)** and among the active groups. This would also potentially enable time-zone appropriate rotas, to minimize middle-of-the-night operations by advocates. We would advocate that each group of the LVK that has low latency/"EMfollow" scope in their MOU/scope of work should provide at least one shifter for the EM advocate rota, which is already the case for detector characterization shifts in Virgo. To enable this, we expect that regular training calls will be organized during O4, during which online tutorials (including basic quizzes) reflecting the expected activities can be performed. **The advocates are expected to be able to perform dry-runs in advance**. It would help to have an online location, such as GraceDB, indicating when the next shift is, who is on shift that week, etc., with information like the PE and detchar rota member. 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.

Given the relatively low interest of follow-up of BBHs, relieving the duties of the RRT teams for those events classified entirely in the BBH category (p(BBH)$>$ 99) is warranted. We should also examine the response to events with significant Terrestrial probability; in this consideration, we should understand how the CBC focused probabilities affect unmodeled sources. On the other hand, more information and more updates for those "golden events," i.e. those containing NS components, should be prioritized for updated sky maps and false alarm rate information if possible. This process would be simplified by giving the power to any search pipeline advocate, given a clearly real event, to quickly give a sign-off and allow things to proceed quickly. This speed would be supported by the rapid (and perhaps automated) notifications of updated sky maps being available from parameter estimation; a tag on GraceDB would perhaps suffice. Similarly, organized regular calls at a fixed time for circular updates can be planned for those events that do not have significant time pressure (i.e. non-golden events). During the operations call, the EM advocate should provide a quick report to follow-up news from the events as well as a summary of what worked (or not). Finally, in case of infrastructure failure, in particular during "golden" events, a designated "maintenance" expert, available during different time zone slots, should be identified.

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 [3] and GraceDB [4]. Additionally, to increase the internal participation, tutorials/bootcamps like the 2019 MIT low-latency bootcamp [5] 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[6]. However, in O4, we suggest to have a userguide that include 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[7]. 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[8] as the exemplar that LIGO/Virgo should emulate.

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

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[9] or MkDocs[10].
- 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.

[9] `https://www.sphinx-doc.org/`
[10] `https://www.mkdocs.org`

Op-6.8 Low-latency public alert dissemination and external communication

O3, Lessons Learned 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. 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 sensitvity and publication plans were regular, we should evaluate if the external community made the correct use of our information. For example, the 190814bv event was misinterpreted by various groups as a clear NS-BH collision (see numbers of GCNs and observations related to this event and title arxiv publications). In addition, from the follow-up community, we should determine what the various trigger parameters implied as whether a candidate was astrophysical or non-astrophysical. The main one was a detailed explanation of how the different pipelines were combined to make the decision on whether a given event would be released for distribution to the public. **The understanding of the meaning of the various "parameters" measured for each event, and what they indicate (real vs false alarm), influences the decision to engage follow-up observations**. Finally, we do not have a complete picture if/how all the services set in place by our collaboration is understood and used by the external community (already discussed in section 6).

- Collection of feedback from the external community: a survey organized by the low latency group to understand the broader community's perspective will be undertaken. The goals are 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 will focus on interactions with the alert system, including its contents and dissemination, and how they may be improved. To this end, a draft is available[11], with an associated presentation here [12]. A link to how the actual survey would look like in indico platform is here [13]. The decision and process is under the **direction** of the LIGO and Virgo spokespersons, but the content should be in some way discussed inside the observational and operation groups.

- Things to improve in the future (both between O3 and O4 and during O4):

  – Tutorials for the User Public: To address the communication issue, we suggest the following improvements: i.) Have face-to-face sessions at various scientific conferences. Given the current sitation with COVID-19, this is not directly possible, but "zoom" (or similar) sessions can/will be done, 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

[11]https://docs.google.com/document/d/1KLxwOMfaMu2Owfwapk1-MrhEGrpplbHJOEGdmOpHYVY/edit?usp=sharing

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

[13]https://indico.in2p3.fr/event/21410/surveys/61

telecons should continue and perhaps those might be increased to allow more groups to share results on interesting, recent candidates. And while it may not be always possible, continued clarification of “candidates” and “events,” perhaps with clarification 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.

– 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.

– LVK alert service. Given that i.) EM follow-up regions of the sky are large compared to the FoV for most EM-based followup instruments, ii.) tiling 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 do not waste some time chasing non-astrophysical 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).

– The NASA GCN portal can be improved as listed above: i.) Targeted Filtering of Notices as a broker: GCN has built in an 8 parameter filtering capability, but it was not used in O3 because we did not know which parameters should be assigned for these LVK types and we did know what the threshold levels should be to eliminate the majority of the non-astrophysical events. 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-dimentional filters. ii.) Filtering on the Circulars: It was true that LVK brought to the GCN a new group people who had never used it before (especially for receiving them). Some of these new people said they were receiving too many Circulars; they were only interested in the follow-up information on L-V events – and did not want to retreive all the Circulars on other sources (especially GRB sources). It is simple to achieve a 90% reduction of the GRB/SGR Circulars with a simple filter that looks for “GRB” or “SGR” in the Subject-line, and if found, it does not send that Circular to that recipient with the filter enabled. iii.) 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, etc.). LVC alerts will, of course, be part of their streams. This can be done by (a) a direct connection to the 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,...).

7 Computing and Software (CompSoft)

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 2020-21 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.

Op-7.1.1 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.

Appendix A 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 subsections of this document. All of those additional sections describe IGWN computing goals, projects and tasks *beyond* basic maintenance, operation, and support.

Op-7.1.2 Low-Latency Computing (CD.LL)

- Review 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.
  - Establish a flexible system to replay archival data to simulate new data. [CD.LL.REPLAY]
  - Eliminating remaining single points of failure in the end-to-end chain of low-latency infrastructure.
  - GraceDB enhancements.
- 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]
- 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]
- Migrate GraceDB, gwcelery, and LValert services to the new high-availability low-latency platform. [CD.LL.HAMIG]
- Add support for SciTokens to the GraceDB server & client to simplify and modernize user authentication. [CD.LL.SCIT]

Op-7.1.3 Open Data (CD.OD)

- 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]
- Add features to the new "Event Portal" , including a general query interface and additional support for multimedia files (e.g. audio).
- Update the Event Portal to support strain file locations spread across multiple directories, and to support a redesign of the strain files directory structure
- 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.
- Update the 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.

Op-7.1.4 Collaboration Data Distribution, Formats & Management (CD.DATA)

- Low-latency data
  - Drop support for O2 Data Monitoring Tool (DMT) and FrameLink solutions and harden O3 Kafka system to improve performance and reliability for O4. [CD.DATA.KAFKA]
- Bulk data
  - Migrate from LIGO Data Replicator (LDR) and Virgo legacy systems to Rucio to reduce IGWN development & support burden. [CD.DATA.RUCIO]
- Data Formats
  - 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.
  - *Subject to the outcome of a planning exercise on data formats* (see MGMT), enhance FrameCPP and FrameL in order to improve maintainability, and modularize data representation to support additional formats (e.g., HDF5).
- Data Management
  - 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.
  - 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]
  - 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]
  - 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
  - Add support for SciTokens to the gwdatafind server & client to simplify and modernize user authentication. [CD.DATA.SCIT]

Op-7.1.5 Offline computing (CD.OFFC)

- Workflow management
  - Provide a generic python-based workflow management library to generate and manage IGWN DAGs, in collaboration with the HTCondor team.
- Resource provisioning

  - 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.)

- Improved resource accounting
  - Develop improved IGWN accounting system (v2)
    - 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).
  - Enhance existing IGWN accounting system (v1)
    - Move existing instance of the IGWN accounting database and backend to a European data center in order to address GDPR issues.
    - 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.
    - Allow jobs to be tagged by event candidate or target publication.
    - Provide a web form (and semi-automated backend) to request new accounting tags.
    - 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.
- Automated workflow/resource prioritization
  - With DAC, develop an updated workflow/resource prioritization policy for O4.
  - Design and implement any infrastructure needed to support O4 workflow/resource prioritization.
- Simplify and modernize user authentication to IGWN services (x509 to SciTokens)
  - Add support for SciTokens to the IGWN Segment Database server & client.
- Usability
  - Complete and deploy new DQSegDB web interface.

Op-7.1.6 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:
  - Support IGWN developers with conda packaging for all user software.
  - Support IGWN developers with container development/deployment for pipelines.
  - Reduce Debian/RHEL packaging support to ‘critical’ and system services.
  - Support IGWN developers with migration to Python 3.6+ for all software.

- Complete Virgo transition from CMT to CMake/Meson and conda.

- Version control:
  - Complete LIGO transition from legacy version control systems (e.g., SVN) to GitLab:
    * Global Diagnostic System (GDS)
    * Document Control Center (DCC)
    * Calibration
  - Complete Virgo transition from SVN to GitLab

- Software Change Control Board (SCCB)
  - SCCB Chair Duties (chairing calls, coordinating effort, reporting, etc.)
  - Continue integration with all low-latency analysis components
  - Refine procedure to reduce time-to-approval
  - Integrate with offline analysis

- Automated Testing
  - Develop policies for where automated tests are to be required (e.g., low-latency detection pipelines and infrastructure).
  - Provide tools and services to facilitate automated testing of low-latency and offline DA pipelines, and of computing infrastructure (e.g., GitLab CI)
  - Provide support & training to assist developers in the adoption of automated testing (analogous to LIGO optimization team support of code optimization).

- Complete Transition of Python 2 Code to Python 3 to Ensure Long-Term Supportability
  - Update SegDB Server to Python 3
  - Update low-latency (kafka) data distribution to Python 3

Op-7.1.7 Detector Characterization (CD.DETCHAR)

- Enhancements to Channel Information System
  - Update user interface
  - Automate database updates from new data

- Enhancements to Data Quality Reports
  - Improve low latency response
  - Provide an easier path to add new analysis tasks

- Leverage summary page infrastructure to improve responsiveness and ease of navigation

Op-7.1.8 Computing Optimization (CD.OPT)

- O4 and O5 CPU & GPU Demand Estimation [CD.OPT.DEMAND]
- O4 and O5 CPU & GPU Supply Estimation [CD.OPT.SUPPLY]
- Trade studies for LIGO Lab Low-Latency O4 hardware procurement [CD.OPT.HW]
- Operational Optimization Development (e.g., testing & deploying completed optimizations into production) [CD.OPT.OPS]
- O4 Data Analysis (DA) Pipeline Optimization (focused on PE) [CD.OPT.O4]
- O5+ Data Analysis (DA) Pipeline Optimization [CD.OPT.O5]
- Enabling of CPU & GPU optimization efforts across IGWN (via hardware test stands, tools, consulting, etc.) [CD.OPT.ENABLE]

Op-7.1.9 Collaboration Support (CD.CSS)

- Identity and Access Management (IAM):
  - Transition from MyLigo to COmanage managed via CILogon contract and internal effort
  - 2-factor Authentication for LIGO.ORG
  - Develop mechanism/s & policy for cyberinfrastructure collaborator access
  - Reengineer igwn.org mailing-list software to reduce ongoing support burden
  - Migrate sympa.ligo.org from LLO hardware to HA platform
  - Computing Security
  - LSC Security Liaison Duties
  - Virgo Security Liaison Duties
  - KAGRA Security Liaison Duties
  - Security Committee Duties
    - Performing security reviews of critical IGWN computing infrastructure and proposals.
    - Other duties should be enumerated here.
    - Climate Change & Remote Participation Planning

Op-7.1.10 Management (CD.MGMT)

- LIGO CompSoft Chair Duties (chairing calls, coordinating effort, reporting, etc.)
- Virgo CompSoft Chair Duties (chairing calls, coordinating effort, reporting, etc.)
- KAGRA CompSoft Chair Duties (chairing calls, coordinating effort, reporting, etc.)

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 homogenous 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.

Op-7.2.1 Data Distribution, Formats & Management (SC.DATA)

- Low-Latency Data
  - 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.
  - Prototype a deployment of O4 Kafka-based low-latency data distribution system at additional data centers in Europe (beyond EGO).
- Bulk data
  - Primary aggregated *h(t)* access via CVMFS to enable distributed computing.
  - Finalize data distribution architecture (who runs XrootD Origins, CVMFS Stratum-0s, caches) to improve future performance and availability.
  - Deploy StashCache Origin for proprietary data managed by each Collaboration
  - Establish CVMFS directory layout for public and proprietary data
  - Add support for SciTokens to authenticated CVMFS for proprietary data to simplify and modernize user authentication.

- 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
  - Define and implement a derived data cataloguing and bookkeeping solution
- 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.

Op-7.2.2 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.
- Workflow management
  - Support migration of all existing non-LDG (Virgo) pipelines to use standard IGWN HTCondor interfaces and environments.
  - Support migration of additional LDG pipelines to use standard IGWN HTCondor interfaces and environments.
  - Develop a "backfill" queue and capability for IGWN resources.
- Resource provisioning
  - Integrate all LIGO, VIRGO, and KAGRA computing providers with standard IGWN services and interfaces.
- Deploy Virgo-managed HTCondor submit nodes in Europe (starting with Nikhef).
- Establish and staff a distributed IGWN Computing Operations & Support Team and process.
- Complete the process of reviewing and adapting the documentation (both for users and support team)

Op-7.2.3 Software and Environments (SC.SOFT)

- Virgo specific enablers:
  - Transition from Virgo Cascina SVN software archive to IGWN git software archive
  - Transition from CMT to CMake/Meson for software build
  - Transition to Conda for software environment definition
- Packaging:
  - Centralised conda package distributions available via CVMFS on all resources
  - Containers distributed via CVMFS on all resources

Op-7.2.4 Computing Optimization (SC.OPT)

- Improving availability & usability of shared/external computing resources via IGWN.

Op-7.2.5 Virgo/KAGRA support for IGWN computing infrastructure (SC.VK)

- Transition some LSC-managed services to Virgo primary responsibility
  - GitLab
  - (possible) production SegDB services
  - Negotiate KAGRA contribution to CILogon contract (currently Lab+LSC+Virgo)

Op-7.2.6 Collaboration Support (SC.CSS)

- Establishing a unified IGWN computing help desk.
- Identity and Access Management (IAM)
  - Transition services to a Federated Identity model
    - Transition from myLIGO and gw-astronomy to IGWN COmanage for Identity Management
      - Design org charts, identifiers, enrollments, etc for IGWN entities
      - Write cakePHP plugins for custom logic (Efforts, Authorship, Council, etc)
      - Migrate existing data into COmanage
    - Transition from LIGO IdP to SAML proxy for web AuthN
  - Provide support for IGWN developers adding SciTokens for data services (CVMFS, GraceDB, DQSegDB, ...)
  - Define an interoperable way of managing VO membership, AuthN and AuthZ across collaborations and Grid resources.
  - Tighter integration between LIGO, Virgo and KAGRA
  - Shibboleth development to enable integration of IGWN and external IAM infrastructures (initially to support IGWN HTC submit node at Nikhef).
- IGWN naming migration
  - Consolidating previously-independent LIGO, Virgo, and KAGRA computing infrastructure and processes.
  - ligo.org -> igwn.org for DNS where appropriate
  - Define procedure to obtain dual-domain host certificates (igwn.org and hosting site domain) to allow for reverse-dns resolution

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.

Op-7.3.1 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.

Op-7.3.2 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.
- Assign needed-effort estimates to each WBS item.
- Populate current effort assignments for each WBS item.

Op-7.3.3 Identify Risks and Opportunities to Improve (WBS.ROI)

- Prioritize IGWN computing projects, tools, and services according to science impact, and identify & communicate science costs and risks of allocating insufficient effort.
- Propose reductions or dropped support for computing infrastructure that requires the most FTE effort relative to its science impact.
- Identify opportunities to outsource non-GW-specific services, to adopt or share common cyberinfrastructure with other science collaborations.
- Identify opportunities to actively solicit additional InfraOps effort towards computing.
- 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)

Op-7.4.1 Management & Future Planning

- Develop a computing "Risk Registry" to be communicated to and accepted by the Collaborations, along with a "Liens List" of deferred projects.
- 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.)
- Update, create, synchronize, and/or merge important LSC, Virgo, KAGRA, and IGWN computing Policies and Procedure documents
  - Computing Acknowledgements policy
  - Open Source Licensing policy
  - Tier-N Document
  - Privacy policy
  - LIGO.ORG credential Acceptable Use Policy (AUP) (link)
  - Other Acceptable Use Policies (AUPs) (link)
  - Establish Service Level Agreement (SLA) for individual services
  - Data Management Plan (DMP) (link)
- Long-Term Data Format (gwf, hdf5, ligo_lw, etc.) Planning
- Establish lightweight MoU with external computing collaborators, e.g., HTCondor, OSG, and Pegasus to clarify expectations and manage risk.
- Develop a charge for LSC, Virgo, and KAGRA *Security Liaisons*, reporting to the Computing Committee of each collaboration, and fill the roles.
- Data Analysis Optimization Planning (including collaborating with HEP and MMA communities, prioritizing optimization opportunities, etc.)
- Seek external funding opportunities, e.g., NSF Cyber Infrastructure solicitations.
- Establish a database identifying computing projects and responsibilities available to LSC, Virgo, and KAGRA groups for InfraOps contributions.
- Maintain & manage important collaborations with outside computing experts (HTCondor, Open Science Grid, Internet2, IAM, XSEDE, EGI, EOSC, possibly WLCG, etc.) via formal liaisons and regular updates.
- Identify & explore the science potential of new computing models, architectures, technology, tools, optimizations, etc. for GW computing over a 2-5 year future timeframe.
- Identify opportunities for publication of IGWN computing results.

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

Op-7.5.1 Continued Maintenance, Operation, and Support of Existing Data Analysis Tools & Services (CD.MOS.DA)

- Data-Quality Segment Database (SegDB) Service (Production, Development, and Testing instances)
- Data-Quality Segment Database (SegDB) Client
- LALSuite
- GraceDB Service (Production, Development, and Testing instances)
- GraceDB Client
- LVAlert Server
- LVAlert Client
- LIGO DMT Service
- KAGRA DMT Service
- LIGO Aggregated h(t) generation Service
- LIGO USA Bulk Data Distribution
- LIGO India Bulk Data Distribution
- Virgo h(t) Aggregation & Delivery
- Virgo Bulk Data Aggregation & Delivery
- KAGRA h(t) Aggregation & Delivery
- KAGRA Bulk Data Aggregation & Delivery
- Public GWOSC Website Service
- Public GWOSC Tools
- GWOSC Open Data CVMFS Origin Service
- Workflow Management Toolkit (aka pipeline.py)
- GWDataFind CIT Service
- GWDataFind GWOSC Service
- GWDataFind LHO Service
- GWDataFind LLO Service
- GWDataFind UWM Service
- GWDataFind CVMFS Service

- NDS2 Service (CIT)
- NDS2 Service (LHO)
- NDS2 Service (LLO)
- NDS2 Client
- Software Source Distribution Archive (software.ligo.org)
- IGWN Conda Service
- Koji RHEL Build Platform
- FrameL Library
- LDAS Tools
- LIGO Summary Page Service
- LIGO LDVW (areeda 0.3 FTE - maintenance, new requests, tech support)
- LigoDV (Matlab)
- GWpy Library
- Authenticated h(t) CVMFS Origin Service
- LALSuite Extra Tools
- IGWN Submit/Dev nodes
- Low-Latency EMfollow Virtual Machines
- Low-latency Event Processing (gwcelery) Service
- Rapid CBC localization service and public sky map visualization and manipulation package (ligo.skymap)
- Public documentation and sample code for public alerts (LIGO/Virgo Public Alerts User Guide)
- LIGO Low Latency h(t) Distribution Service
- Virgo Low Latency h(t) Distribution Service
- KAGRA Low Latency h(t) Distribution Service
- Electronic Notebook Service (alog)
- LIGO Channel Information System (CIS) Service
- HTCondor Service
- CIT Cluster Computing Element (CE) Service
- LLO Cluster Computing Element (CE) Service
- LHO Cluster Computing Element (CE) Service

- GTech Cluster Computing Element (CE) Service
- InFlux Database Service
- Low-Latency Kubernetes Platform Service
- LLO Segment Generation (SegGener) Service
- LHO Segment Generation (SegGener) Service
- KAGRA Segment Generation Service
- Virgo Segment Generation Service
- Data-Quality Segment (DQSegDB) Publisher Service
- Inspiral Range Calculator Tool
- Data Quality Report (DQR) Service
- Channel Information System (CIS) Service

Op-7.5.2 Continued Maintenance, Operation, and Support of Existing Collaboration Support Tools & Services (CD.MOS.CSS)

- LIGO DCC Service
- LIGO-India DCC Service
- KAGRA DCC Service
- EPO Website Service
- MyLIGO Service
- LIGO Kerberos Service (KDCs)
- LIGO Shibboleth IdP Service
- LIGO Backup Shibboleth IdP Service
- LIGO LHO Shibboleth IdP Service
- LIGO LLO Shibboleth IdP Service
- COmanage Service
- SAML Proxy (SaToSa)
- LIGO LDAP Master Service
- LIGO robotic/automatic credentials
- LIGO data grid cluster account management system
- LIGO LDAP CIT Replica Service

- LIGO AddressBook Replica Service
- CIT gw-astronomy LDAP Replica Service
- UWM gw-astronomy LDAP Replica Service
- LIGO UWM LDAP Replica Service
- LIGO LLO LDAP Replica Service
- LIGO LHO LDAP Replica Service
- LIGO Grouper Service
- LIGO Remote Participation (teamspeak) Service
- Virgo Remote Participation (teamspeak) Service
- LIGO.ORG Gsuite Services
- LSC MOU Software
- IGWN GitLab Service
- LIGO FosWiki Service
- Virgo Wiki Service
- Jupyter Hub (CIT) Service
- Jupyter Hub (LHO) Service
- Jupyter Hub (LLO) Service
- Jupyter Hub (Cardiff) Service
- IGWN Mattermost Service
- { ligo| igwn} .org Mail Transport Agent (MTA) Service
- Usage Accounting Data Collection Client
- Usage Accounting Database Service
- Usage Accounting Reports Service
- { ligo| igwn} .org DNS Service
- ligo.org mailing lists Service (sympa.ligo.org)
- CILogon Grid Certificate Authority Service
- gw-astronomy.org COmanage Service
- gw-astronomy.org mailing list Service
- gw-astronomy.org Wiki Service

- LIGO India COmanage Service
- LIGO India IdP Support Service
- LIGO India Kerberos Service
- LIGO India LDAP Service
- LIGO Voting/Elections System (vote.ligo.org)
- LIGO P& P Database Service
- Service monitoring (icinga) Service

Kg-7.5 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, we 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-7.5 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 ↔ Kashiwa low latency links
- O.C.3.2. Monitor the CIT ↔ 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-7.5 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 (JRPC)

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 the inevitable changes due to COVID-19 are correctly communicated to the greater scientific community
- Organize LVEM/scientific community telemeetings and Town Halls as needed
- 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

Op-8.2 JRPC Activities beyond the next 12 months related to O4 preparations

- Roughly 1 year before the next Observing Run starts, currently planned to be O4 starting in early 2022, 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.

9 Support of Observatories (SOO)

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.

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/LSCFellowsProjectList.

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 January 2021 due to the Covid pandemic. While plans are somewhat uncertain, reservations are now being accepted for both observatories for Fellows visits starting January 2021.

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 2020-2021 (https://dcc.ligo.org/LIGO-M2000130), specifically 2.1 LIGO Observatory Operations, 2.2 LSC Detector Commissioning and Detector Improvement activities, 2.3 LSC Fellows Program, 2.4 Detector Calibration and Data Timing, and 2.6 Detector Characterization. Earlier sections of this LSC-Virgo Operations White Paper also list specific aspects of the work planned for 2021 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 until 2022. The remainder of 2020 and all of 2021 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 2020 and all of 2021 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 beween June 2020 and January 2022 shold provide ample opportunities for LSC members to be involved in activities related to interferometer hardware, including in-vacuo components. The SO committee chairs will help to coordinate participation of interested LSC members in these and other pre-O4 activities.

10 Open Data (OpDa)

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 data releases are described in the 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. So far, we have hosted three 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[14], and additional notes may be found in DCC documents G2000859, G1900798, and G1800778.

Goals for the Open Data Working Group in FY2021 include:

- Maintain the GWOSC web server and quality of service to the users

[14] gw-osc.org/s/workshop1 - gw-osc.org/s/workshop2 - gw-osc.org/s/workshop3

- 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
- Plan and host an Open Data Workshop
- Respond to user questions that come in through GWOSC ticket system (gwosc@igwn.org)
- Develop new features of the GWOSC Event Portal, including improved queries and documentation
- Develop capability for auxiliary channel releases, including documentation and channel selection
- Evolve data archive directory structure to a format more consistent with internal data archive, both for events and bulk data
- Promote and contribute to the standardization of confidence intervals for the astrophysical parameters and of their production from PE samples

In this document, 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 be counted 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 2020 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.
- Plan and host an Open Data Workshop, or a “mini-workshop” re-using existing materials.
- Answer questions in the GWOSC ticket system
- 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

11 Upgrades

Op-11.1 LIGO A+ Upgrade

The “A+ detector" project is a major upgrade to the existing Advanced LIGO detectors, beginning in 2019 and expected to continue through the end of 2023. The A+ project in 2019 will be carried out in parallel with the O3 observing run.

Design and procurement for A+ is underway, led by the LIGO Lab and international partners and with support by members of the LSC. The implementation of A+ has design, fabrication, assembly, and test of components in parallel with the observing runs, and with intensive installation during commissioning breaks. In the measure possible, A+ elements will be integrated with the existing aLIGO instruments to make incremental improvements in the performance. In particular, the A+ project plans to install the filter cavity between O3 and O4. This involves civil construction, vacuum system enlargements, installation of new seismic and suspension components, and commissioning.

Activities related to A+ operations are: testing frequency dependent squeezing at 1064 nm; designing measurement and implementation methods for Newtonian noise reduction; testing low noise control of the homodyne readout; reliability testing for higher stress silica fibers; and studying production of fused silica suspension fibers to ensure that frequencies of violin modes are sufficiently matched. Substantial efforts are underway to develop new optical coatings for A+ with improved mechanical loss. These coatings are expected to be amorphous oxide coatings deposited with ion-beam-sputtering techniques. Parallel efforts are underway to understand the fundamental loss mechanisms for these coatings, and to improve the loss with different compositions, nano-layered coatings, and modified deposition and annealing processes.

Results from these tests are expected to be implemented in A+, using frequency dependent squeezing with a 300 m filter cavity, balanced homodyne readout, implementation of lower loss coatings when developed, and installation of new test masses from upgraded pulling and welding systems for fused silica fibers. Details on A+ can be found in the LSC Instrument white paper [58].

Vg-11.1 Virgo AdV+ Upgrade

Advanced Virgo Plus (AdV+) is an upgrade of Advanced Virgo to be realized in two phases named Phase I and Phase II. The installation of AdV+ Phase I will take place between the observation runs O3 and O4. The installation of AdV+ Phase II will take place between the observation runs O4 and O5. Broadly speaking the main goal of AdV+ Phase I is to reduce the interferometer sensing noise (the limitation coming from quantum noise) by implementing the following improvements:

- Implementation of the signal recycling technique to enlarge the detector bandwidth by reshaping the quantum noise.
- Increase of the laser power injected into the interferometer by a factor of two to decrease the effect of shot noise above a few hundred Hertz.
- Injection of frequency dependent squeezed vacuum states to reduce further the shot noise without increasing the radiation pressure noise at low frequencies.

In addition, AdV+ Phase I will be the occasion to deploy an array of seismic sensors in the central and end buildings to test Newtonian noise cancellation techniques. In parallel with the deployment of the array of sensors, an effort will be done to reduce the environmental noise coming for the air conditioning system as well as several technical noises.

The combination of the improvements outlined above will allow improving the sensitivity of the detector and reaching a range of the order of 100 Mpc for coalescing binary neutron stars.

AdV+ Phase I will use the same mirrors currently installed and so the beam geometry in the interferometer arms will remain the same. Therefore, the mirror thermal noise will remain at the current level.

AdV+ Phase II instead will be focused on the reduction of the mirror thermal noise. To do so the beam on the end mirrors will be enlarged to about 20 cm diameter while keeping the same size in the central area of the interferometer. In order to do so larger mirrors will have to be used for the end mirrors. These mirrors will be 55 cm in diameter and 100 kg in weight. The change in the beam divergence will require to change also the two cavity input mirrors, the power recycling mirror and the signal recycling mirror. These mirrors will use substrates of the same size as the ones currently in use (35 cm in diameter). To further reduce the mirror thermal noise the best available coatings available at the time of O4 will be used. To this purpose, the development of coating with lower thermal noise is part of AdV+ Phase I.

The combination of the improvements outlined above should allow pushing the binary neutron star range of Virgo around 200 Mpc with the exact level depending on the reduction of coating thermal noise.

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.

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, etc. 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.

References

[1] 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.

[2] 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.

[3] 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.

[4] 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.

[5] 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.

[6] Dripta Bhattachargee et al. Fiducial displacements with improved accuracy for the global network of gravitational wave detectors. Technical Report P2000113.

[7] 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.

[8] Virgo Collaboration. Virgo photon calibrators for O3. *In preparation*.

[9] 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.

[10] Technical report.

[11] Ling Sun et al. Characterization of systematic error in Advanced LIGO calibration. 5 2020.

[12] LIGO-Virgo Data Quality Report for Event Candidates Design Document. `https://dcc.ligo.org/LIGO-T1800123`.

[13] Data Quality Report Documentation. `https://docs.ligo.org/detchar/data-quality-report/`.

[14] The LIGO-Virgo Summary Pages. `https://www.gw-openscience.org/detector_status/`.

[15] Data Monitor Tool Project . `https://labcit.ligo.caltech.edu/~jzweizig/DMT-Project.html`.

[16] The LIGO-Virgo segment database project wiki.
https://wiki.ligo.org/Computing/DQSEGDB.

[17] The LIGO-Virgo segment database.
https://segments-web.ligo.org/.

[18] PyOmicron.
https://doi.org/10.5281/zenodo.2616570.

[19] B.P. Abbott et al. Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914. *Class. Quant. Grav.*, 33(13), 2016.

[20] DetChar Hardware Injections.
https://dcc.ligo.org/LIGO-T1900555.

[21] 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.

[22] SNAX toolkit: Stream-based Noise Acquisition and eXtraction toolkit.
https://lscsoft.docs.ligo.org/gstlal/feature_extraction.html.

[23] GWpy: a python package for gravitational-wave astrophysics.
https://gwpy.github.io
https://doi.org/10.5281/zenodo.597016.

[24] GW Detector Characterization python utilities.
https://github.com/gwdetchar/gwdetchar
https://doi.org/10.5281/zenodo.2575786.

[25] Gravitational-wave interferometer summary information.
https://github.com/gwpy/gwsumm
https://doi.org/10.5281/zenodo.2647609.

[26] Stochmon: A LIGO Data Analysis Tool.
https://dcc.ligo.org/LIGO-T1400205.

[27] STAMP-PEM.
https://ldas-jobs.ligo-wa.caltech.edu/~stochastic/stamppem_docs/.

[28] 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.

[29] 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.

[30] The LIGO Channel Information System.
https://github.com/lscsoft/cis.server.

[31] Pointy Poisson.
https://docs.ligo.org/detchar/data-quality-report/tasks/pointy-dqr.html.

[32] September 2019 O3a Saefty Studies.
https://dcc.ligo.org/LIGO-T1900697.

[33] The Gravitational-wave Veto Evaluation and Testing suite.
https://github.com/gwpy/vet.

[34] 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.

[35] 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.

[36] 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.

[37] 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.

[38] DataViewer User Guide.
https://labcit.ligo.caltech.edu/~aivanov/Dataviewer/helpG2.html.

[39] LIGO Physical Environment Monitoring.
http://pem.ligo.org/.

[40] DetChar Hardware Injections.
https://git.ligo.org/detchar/detchar-hwinj.

[41] Michael Zevin et al. Gravity Spy: Integrating Advanced LIGO Detector Characterization, Machine Learning, and Citizen Science. *Class. Quant. Grav.*, 34(6):064003, 2017.

[42] Michael Coughlin. Noise line identification in LIGO S6 and Virgo VSR2. *J. Phys. Conf. Ser.*, 243:012010, 2010.

[43] J. Rollins et al. Distributed state machine supervision for long-baseline gravitational-wave detectors. *Review of Scientific Instruments*, 87(094502), 2016.

[44] Elena Cuoco et al. Enhancing Gravitational-Wave Science with Machine Learning. *arXiv:2005.03745*, 2020.

[45] The Gravitational Wave Open Science Center.
https://www.gw-openscience.org.

[46] Proposal for common data formats for LIGO-Virgo data products.
https://dcc.ligo.org/LIGO-T1800014.

[47] Bruco.
https://github.com/mikelovskij/bruco.

[48] Omicron.
F. Robinet, *VIR-0545C-14*
https://tds.virgo-gw.eu/?content=3&r=14693.

[49] Omega-scan.
https://github.com/gwdetchar/gwdetchar.

[50] NoEMi.
T. Accadia, et. al, **1** 012024, 2012.

[51] Pastavi.
H. Yuzurihara, *JGW-G1910938*
https://gwdoc.icrr.u-tokyo.ac.jp/cgi-bin/private/DocDB/ShowDocument?docid=10938.

[52] NoiseBudgetter.
A. Shoda, *JGW-T1910917*
https://gwdoc.icrr.u-tokyo.ac.jp/cgi-bin/private/DocDB/ShowDocument?docid=10917.

[53] CAGMon.
http://gwwiki.icrr.u-tokyo.ac.jp/JGWwiki/CAGMon.

[54] E. J. Son. EtaGen. *arXiv preprint arXiv:1810.07555*.

[55] S. Morisaki. ICA. *Proc. Jpn. Acad. Ser.*, B92:336, 2016.

[56] LIGO calibration DCC Document Tree.
https://dcc.ligo.org/LIGO-E1500281.

[57] LIGO calibration web page.
https://wiki.ligo.org/Calibration/WebHome.

[58] The LIGO Scientific Collaboration. LSC Instrument Science White Paper - 2020. *LIGO Document: LIGO-T2000407*, 2020.