

The LIGO-Virgo-KAGRA Observing Run 4

Nicolas Arnaud¹, Jenne Driggers², Brian O’Reilly³ and Takahiro Sawada⁴

on behalf of the LIGO Scientific Collaboration, the Virgo Collaboration and the KAGRA Collaboration.

¹Université Claude Bernard Lyon 1 CNRS, IP2I Lyon / IN2P3, UMR 5822, F-69622 Villeurbanne, France

²LIGO Hanford Observatory, Richland, WA 99352, USA

³LIGO Livingston Observatory, Livingston, LA 70754, USA

⁴ICRR, The University of Tokyo, Hida, Gifu 506-1205, Japan

E-mail: n.arnaud@ip2i.in2p3.fr

Abstract. The Observing Run 4 (O4, May 2023 – November 2025) is the longest data-taking period to date for the LIGO-Virgo-KAGRA (LVK) network of ground-based gravitational-wave (GW) detectors. Characterized by improved performance of the instruments, this run has produced an unprecedented number of GW candidates in near-real time, publicly announced to allow astronomical multi-messenger follow-ups.

1 O4: the LIGO-Virgo-KAGRA Observing Run 4

The fourth LIGO-Virgo-KAGRA (LVK) Observing Run (O4) started more than two years ago and it is still ongoing at the time of writing – July 2025. To date, it is the longest data-taking period for the global network of ground-based gravitational-wave (GW) interferometric detectors: the two Advanced LIGO detectors [1], Advanced Virgo [2] and KAGRA [3]. It is characterized by a record rate of GW candidates, thanks to the improved performance of the different instruments on the one hand, and of the low-latency framework producing public alerts on the other.

The O4 run has been divided into three sets of roughly equal durations, whose main characteristics are summarized in Table 1. The data of each period form a single block and will all be released publicly at the same time, according to a predefined schedule [4]. They will be published on the GWOSC website <https://gwosc.org>. There is a three-month break between the end of O4a and the beginning of O4b and the data taking conditions are different between the two periods, with the Virgo detector joining the two LIGO instruments for the second one. By contrast, the O4b and O4c periods are adjacent: the change was primarily made for data release scheduling purposes. Yet, that switch took place at a time when none of the detectors was observing, thus making a clear separation between the O4b and O4c data: the former set ending *before* 1700 UTC on January 28th, 2025 and the latter starting *after* that date. Contrary to O4a and O4b, O4c includes a planned downtime of about 10 weeks, from April 1st to June 11th, during which various detector works took place – including an intervention at LIGO Livingston to fix an issue with its vacuum system prior to the 2025 hurricane season. The restart of O4c was marked by the addition of the KAGRA detector that started taking data alongside the other three observing instruments: for the first time, four advanced GW detectors are taking long stretches of data together – and five in total, including GEO600 [5] in Germany.

There is a significant imbalance in terms of sensitivity among the detectors of the LVK network. The BNS range¹ figure-of-merit shows a factor-of-three difference between the LIGO detectors and Virgo: our

¹The average distance up to which a detector can detect the merger of a binary $1.4 M_{\odot}$ neutron stars merger with a signal-to-noise-ratio of 8.

O4 period	O4a	O4b	O4c ^(*)
Dates	2023/05/24 to 2024/01/16	2024/04/10 to 2025/01/28	2025/01/28 to 2025/11/18
Main detector configuration	H1 + L1	H1 + L1 + V1	H1 + L1 + V1 + K1
Full network up duty cycle	53% (H1-L1)	31% (H1-L1-V1)	37% (H1-L1-V1)
All detectors down time fraction	17%	11%	10%
Number of public alerts	81	105	25
Data release date	2025/08/26	2026/05/26	2026/12/16

Table 1: Main characteristics of the three O4 periods.

H1 ↔ LIGO Hanford; L1 ↔ LIGO Livingston; V1 ↔ Virgo; K1 ↔ KAGRA. ^(*): O4c data end of July 14th, 2025 and the 10-week detector work period has been excluded from the duty cycle computations.

experience from GW170817 and other subsequent events has shown that a third detector, even one that is less sensitive, can significantly improve the sky localization of transient GW sources. Therefore, a joint LVK strategy was applied during O4b and confirmed for O4c: to maximize the LIGO-Virgo detector uptime in order to maximize the number of events localized by the three instruments. This has been achieved by aligning all the known weekly downtimes of the detectors: maintenance, calibration and commissioning slots. The sole exception is the Virgo maintenance, scheduled on Tuesday mornings local time, corresponding to nighttime in the USA. The consequence of this alignment is a network downtime close to 10%, when no detector is taking data. Although no well-localized electromagnetically-bright GW event has been observed so far in O4, that strategy has achieved its goal: Virgo observing data are available for about 80% of the O4b+O4c public alerts, a percentage much larger than if the detectors of the network were operating independently. As the KAGRA BNS range remains below 10 Mpc, it has been decided not to use the data of this fourth detector for low latency alerts. Yet, they flow in real time within the LVK network and are available for near real-time and offline analyses.

2 Low latency alerts and rapid response team vetting

A clear commitment of the LVK network is to deliver public alerts with the lowest possible latency when transient GW candidates are detected by the real-time processing of the data from the LIGO and Virgo detectors. The corresponding framework, that ranges from the production of the calibrated GW strain streams to the publication of a new alert via GCN (The General Coordinates Network: <https://gcn.nasa.gov>), has been developed during the early observing runs and has expanded and matured since then – details can be found in the “LIGO/Virgo/KAGRA Public Alerts User Guide” (<https://emfollow.docs.ligo.org/userguide>). The first five minutes following the detection of a new candidate see the automated sending of two *preliminary* GCN notices, a first one only a few tens of seconds after the trigger and a second, more refined, later, aggregating all the information available from multiple low-latency pipelines. Then, a human-operated Rapid Response Team (RRT) takes over and vets the event, using automatically-generated data quality reports [6, 7] to make a decision: to confirm the public alert or to retract it. That decision is communicated to the astronomy community by sending an *initial* machine-readable notice and its accompanying circular. *Updated* notices and circulars can be sent at later times to provide additional information, such as an updated estimation of the main astrophysical parameters of the source of the GW signal. Figure 1 displays a cumulative count of the O4 public alerts and shows how the rate of such events has increased over time.

The RRT plays a key role in the smooth and timely processing of the low latency alerts. This is a three-tiered system: two level-0 (lv0) shifters are online 24/7 to vet the triggers found by the online system; they are aided by level-1 (lv1) experts who bring their expertise of a particular component of the framework when needed; finally, if a complex issue arises or a major decision is required, a level-2 (lv2) forum is convened on short notice to discuss the event. Lv0 shifts are organized via a rota over three geographical regions, each managed by an LVK collaboration: Europe/Africa for Virgo, Asia/Pacific for KAGRA and the Americas for LIGO. This allows lv0 shifters to work during daylight hours in their respective time zones. Site advocates (the LVK run coordinators or their delegates) attend and steer the RRT meetings, with support from the RRT shift managers.

At the time of the Amaldi16 conference, the RRT has vetted more than 200 public alerts – plus about 10% of low latency triggers that had to be retracted due to instrumental problems or because of a confirmed terrestrial origin: the fraction of retractions is significantly lower than during the O3 run. That

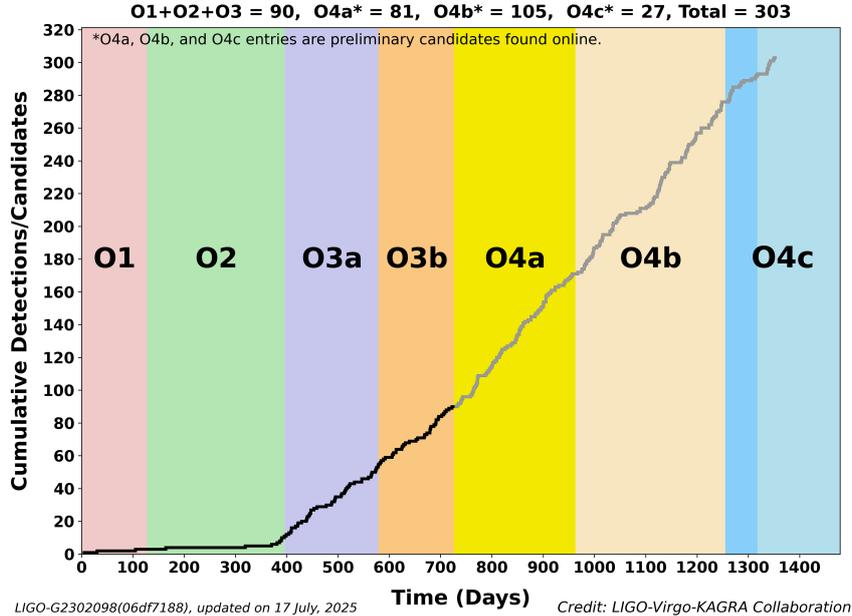


Figure 1: Cumulative plot of the LVK GW events versus the number of days of data taking, with the different observing runs identified by color bands. This snapshot was taken during the Amaldi16 conference: as this plot mixes published detections from O1, O2 and O3 with public alerts from O4, it evolves with time, not only by accumulating new low-latency triggers, but also when results from a past run are published: the preliminary public alerts are replaced by the final confirmed detections.

means more than 4,000 lv0 shifts, performed by more than 600 lv0 shifters, with the help of more than 50 lv1 experts, 13 site advocates and 4 RRT shift managers. Figure 2 summarizes the performance and the progress of the RRT over time, by showing a measurement of its “latency” for each O4 low-latency trigger. Not only does the spread of the distribution decrease with time, but its median value goes down as well. In average, a few tens of minutes at most are now needed to confirm or vet an alert and distribute the decision about its origin – likely astrophysical or not.

In addition to decreasing its latency, the LVK public alert framework has been enriched. Since the beginning of O4, a new class of alerts, the “early warning” (EW) triggers, has completed the existing “full bandwidth” (FB) ones. Compact binary coalescences involving at least one neutron star spend tens of seconds if not more in the sensitive frequency band of the detectors. Therefore, if the signal is strong enough, it could be detected confidently *early*, that is before the merging time. While validated in simulation [8] and operating for more than a year now, the first public alert combining EW and FB triggers is yet to come. Another novelty, put in production for the second part of O4c, is the addition of some source chirp mass information for each relevant public alert. A binned histogram of the source chirp mass probability density function is provided, with the bins ranging from sub-solar systems to very heavy ones, by way of GW170817-like binaries. More information about the performance of the LVK low-latency system can be found in the Amaldi16 proceedings [9].

3 Conclusions

The LVK O4 run is significantly better than O3. The integrated observing Volume \times Time spanned by the network is significantly larger than for the previous run, which leads to a proportional increase in the number of public alerts. This, plus the duration of O4, makes this dataset unprecedented. This overall major success is the result of the dedication and effort of all the scientists worldwide involved in the global network of ground-based interferometric GW detectors. Almost 10 years after the discovery of GW with the event GW150914, the LVK continues making discoveries, showing the power of its network and of the related collaboration work. The LVK is committed to continue this forward, after O4 and for the years to come, in spite of the uncertainties in funding that prevail at the moment. Figure 3 that

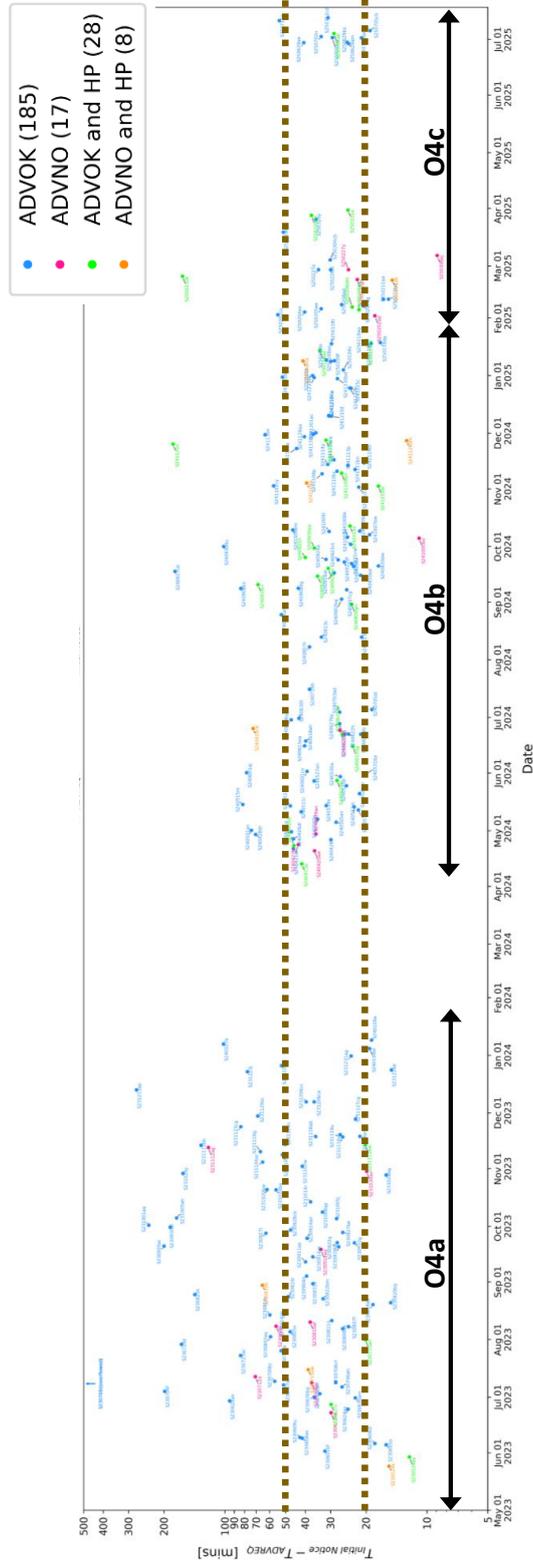


Figure 2: Latency of the RRT: that is, for each GW candidate identified online, the time elapsed between the generation of the trigger and the human decision to validate the public alert (ADVOK), or retract it (ADVNO). The horizontal dashed lines are drawn at 20 and 50 minutes respectively. A further categorization of the events is done by highlighting the few that are considered as high-profile (HP) by the LVK, meaning that a larger pool of experts is immediately involved in the event vetting.

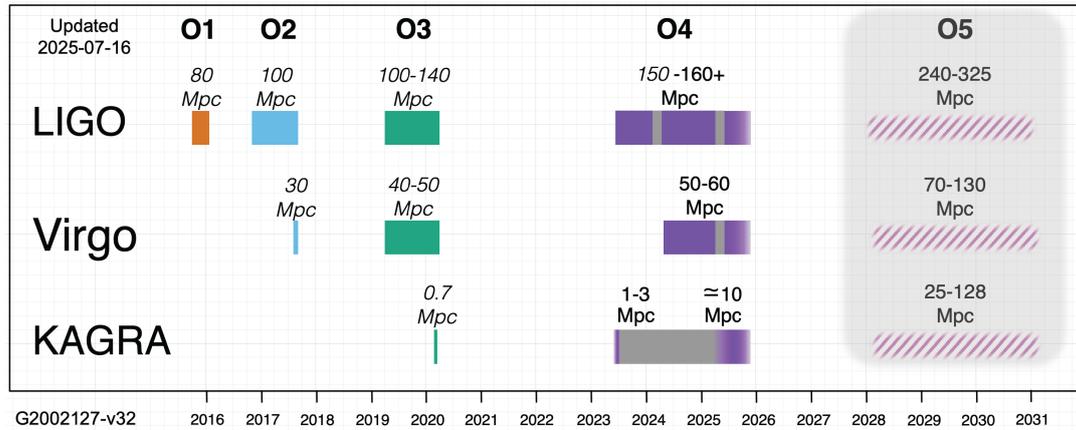


Figure 3: Timeline of the LVK Observing plans, as of July 16th, 2025 – the monthly update that was released during the Amaldi16 conference. The results from the O1-O3 runs have been published and the corresponding data publicly released, while the O4 run is still ongoing –with an end date of Tue. November 18th, 2025. Plans and timelines for O5 are being reassessed.

shows a snapshot of the LVK Observing Plans at the time of the Amaldi16 conference will be updated as plans for future operations and upgrades become available.

Acknowledgments

This material is in part based upon work supported by NSF’s LIGO Laboratory which is a major facility fully funded by the National Science Foundation. Detailed acknowledgments may be found in <https://dcc.ligo.org/P2100218>.

References

- [1] The LIGO Scientific Collaboration 2015 *Classical and Quantum Gravity* **32** 074001 URL <https://dx.doi.org/10.1088/0264-9381/32/7/074001>
- [2] The Virgo Collaboration 2014 *Classical and Quantum Gravity* **32** 024001 URL <https://dx.doi.org/10.1088/0264-9381/32/2/024001>
- [3] The KAGRA Collaboration 2020 *Progress of Theoretical and Experimental Physics* **2021** 05A101 ISSN 2050-3911 URL <https://doi.org/10.1093/ptep/ptaa125>
- [4] LIGO Data Management Plan, <https://dcc.ligo.org/M1000066>
- [5] The GEO Collaboration 2016 *Classical and Quantum Gravity* **33** 075009 URL <https://dx.doi.org/10.1088/0264-9381/33/7/075009>
- [6] The Virgo Collaboration 2023 *Classical and Quantum Gravity* **40** 185005 URL <https://dx.doi.org/10.1088/1361-6382/acdf36>
- [7] Soni, S et al 2025 *Classical and Quantum Gravity* **42** 085016 URL <https://dx.doi.org/10.1088/1361-6382/adc4b6>
- [8] Magee, R et al 2021 *The Astrophysical Journal Letters* **910** L21 URL <https://dx.doi.org/10.3847/2041-8213/abed54>
- [9] Toivonen A 2025 *J. Phys.: Conf. Ser* Amaldi16, submitted