Alert Service Requirements Document

LIGO-Virgo-KAGRA Collaboration

1 Scope/Purpose

The Scalable Cyberinfrastructure to support Multi-Messenger Astrophysics (SCiMMA) project and NASA Time-domain Astronomy Coordination Hub (TACH) are recent efforts to modernize alert distribution systems for time-domain astronomy. SCiMMA HopSkotch is a scalable, high-throughput publish-subscribe system built over Apache Kafka to distribute real-time alerts to enable Multi-Messenger Astronomy (MMA). NASA TACH is a modernization of the Gamma-ray Coordinate Network (GCN) and the High Energy Astrophysics Science Archive Research Center (HEASARC) in terms of speed, reliability, and cross-talk between missions to serve the increasing needs of time-domain astronomy.

These two projects have requested the LIGO/Virgo/KAGRA collaboration to provide requirements for the alert distribution system to be used for public/private alerts in O4 and beyond. More generally, it is useful for the collaboration to maintain a requirements document for an alert distribution system (referred to as *alert service* hereafter). Alert services used to distribute LVK alerts in O4 and beyond will need to satisfy these requirements. This is independent of requirements of the low-latency subsystem developed and maintained by the LVK.

2 System Requirements for Public Alerts

The following are the requirements of the alert service for distributing public alerts from the LVK:

- *Uptime and reliability*: The LVK requires the alert service to distribute **all production** alerts reliably during O4 operations and beyond. Scheduled downtimes may be coordinated with the LVK to ensure that this is achieved. Alternatively, the alerts maybe cached for later distribution. In any case, the uptime of the alert service will be such that 99% of the production alerts are delivered *within the latency requirement* mentioned below. In addition to production alerts, the service should also distribute mock alerts that are periodically sent by the LVK both during and outside observing

runs. The service should distribute periodic heartbeat signals to the LVK about its functional state.

- *Latency of alert delivery to consumers*: To achieve maximum success for the multi-wavelength and high-energy follow-up of GW events, the alert service should provide scalable infrastructure such that the latency of *alert delivery* is at most **1 second** irrespective of the number of consumers even in the case of peak activity mentioned below.
- *Frequency of alerts*: Based on the astrophysical rates of compact binary mergers, the discovery rate in O4 is expected to be O(1) per day. Each candidate may be associated with O(10-20) alerts based on the source properties, with a peak (though sporadic) frequency of 1 Hz. The service should be able to handle an average rate of O(100) LVK alerts every day with peak rates O(1) per second.
- *Alert contents and size*: The LVK will will adopt a new alert streaming platform to distribute alerts in O4. Depending on the GW search analysis, modeled or un-modeled, the alert content may vary in both size and schema. For example, some of the alerts may contain minimal information like the significance and the time of arrival, while others may contain the information related to source properties. The schema for the alerts will be decided by the LVK before the observing run. **In O4, the LVK plans to distribute additional data products, like the GW sky-localization (Multiorder FITS skymap files ~ 1 MB), as a part of the streaming alerts**. Alerts not containing a skymap are expected to have alert size O(1-10) kB. Skymaps will be distributed when it is first available, or when an improvement to the same is available. The expected peak frequency for sending alerts with skymaps is once per ~10 seconds. The alert service should deliver such alerts spanning different sizes and schema with the frequency and latency as mentioned above.
- *Support for legacy alerts*: In addition to streaming alerts, the LVK plans to continue sending alerts in legacy VOEvent format. Unlike the streaming alerts, these may not contain the skymap files as a part of the alert, and therefore expected to be O(1-10) kB. The format of alerts may be different between the legacy and streaming alerts. The legacy alerts may be less frequent compare to streaming alerts. The alert service should support legacy alerts in addition to streaming alerts.
- *Queryable Interface*: The alert service team will provide a web interface to query alerts that have been distributed. The interface will support multiple query types, for example, an unique identification based search, a time window based search, and so on.
- *Citations*: All alerts from the LVK will be cited as **LIGO-Virgo-KAGRA collaboration**.

3 System Requirements for Access-controlled alerts

The following are the requirements on the access-controlled alert streams between the collaboration and partner facilities.

- Access-controlled outgoing alert streams (from LVK to others) should have similar requirements as the public LVK alerts in terms of frequency, size, and latency of delivery.
- The access-controlled content will be independent of the public stream i.e., public alerts will not be a part of the access-controlled alerts.
- The alert service will provide the authentication infrastructure to send alert content to one or multiple facilities for subthreshold analysis, or other follow-up efforts.
- The LVK will have control of, and the responsibility to maintain the subscriber list that receive the access-controlled alerts. The alert service should provide an interface via which this can be done.
- Access control should be available at the user and the user group levels, for example, LIGO.ORG credentials.
- All access controlled alerts will be cited as **LIGO-Virgo-KAGRA collabortion**.
- The alert service will also provide the means to receive access-controlled alerts from partner facilities for the LVK to carry out subthreshold analysis.