

The VAO Transient Facility

Matthew J. Graham¹, S. G. Djorgovski^{1,3}, Andrew Drake¹,
Ashish Mahabal¹, Roy Williams¹ and Rob Seaman²

¹California Institute of Technology, Pasadena CA USA
email: mjpg, ajd, ashish, roy@caltech.edu, george@astro.caltech.edu

²National Optical Astronomical Observatory, Tucson AZ USA
email: seaman@noao.edu

³Distinguished Visiting Professor, King Abdulaziz Univ., Jeddah, Saudi Arabia

Abstract. The time-domain community wants robust and reliable tools to enable the production of, and subscription to, community-endorsed event notification packets (VOEvent). The Virtual Astronomical Observatory (VAO) Transient Facility (VTF) is being designed to be the premier brokering service for the community, both collecting and disseminating observations about time-critical astronomical transients but also supporting annotations and the application of intelligent machine-learning to those observations. Two types of activity associated with the facility can therefore be distinguished: core infrastructure, and user services. We review the prior art in both areas, and describe the planned capabilities of the VTF. In particular, we focus on scalability and quality-of-service issues required by the next generation of sky surveys such as LSST and SKA.

Keywords. Standards, methods: miscellaneous, astronomical data bases: miscellaneous

1. Introduction

The endorsement of the Large Synoptic Survey Telescope (LSST) in the National Research Council's Decadal Survey of Astronomy and Astrophysics (Decadal Survey 2010) identified transient astronomy as a vital area of astronomical research for the next decade and beyond. With its stated goals of enabling and facilitating science, the VAO is well placed to deliver the tools and services required by the community to exploit that new domain. Yet it is not a trivial undertaking. The transient astronomy community is a broad church—a unique, burgeoning, multi-wavelength group consisting of both professional and non-professional astronomers. Fortunately, it is already VO cognizant, with ~47000 VOEvents having been sent at the time of writing, including ~26000 from the Catalina Real-Time Transient Survey (Drake *et al.* 2009). However, it faces daunting data challenges, with LSST projecting event production rates of between ~10⁵ and ~10⁷ notifications per night and predictions that SKA will be capable of detecting 1 core-collapse supernova per second over the whole sky originating somewhere in the redshift range $0 < z < 5$.

Over the past 6 years two NSF-funded projects—VOEventNet and SkyAlert—have already contributed many aspects of the current event infrastructure. The time is now ripe to define the next-generation service—the VAO Transient Facility—which will continue that work and put it onto a sound operational footing, capable of addressing the needs of the community. Below we describe a review study which we have carried out concerning the development of such a facility within the VAO, built upon the existing SkyAlert product. We have identified the short- to mid-term requirements of the community and a programme to transition the prototype SkyAlert to a production-level service—with all that that entails.

2. Community Engagement

A VAO Transient Facility aims to provide the transient astronomy community with robust and reliable tools to produce event notifications, to subscribe to event notifications, and to leverage VAO capabilities for data discovery/analysis. Responses were solicited from a cross-section of the community—LSST, LOFAR, GAIA, SKA Pathfinders (MeerKAT, ASKAP), AAVSO, Zooniverse and LIGO—to identify specific issues within those areas that require attention and to assist in planning and determining priorities. Specific questions were asked about the mode and rates of event delivery, brokering requirements including levels of service and functionality, current standards and infrastructure, repository holdings and performance, and other areas of potential interest such as EPO opportunities. The responses from the community can be broadly grouped into two categories: simplicity and programmatic access.

It is all too easy to try to make everything as feature-rich as possible in order to meet all possible user cases. There is also a strong desire to keep things manageable. Event notifications should just report the basics of a transient event. References can always be used to link to extra data such as images, light curves, etc., but the current version (2.0) of the VOEvent standard (Seaman *et al.* 2011) is perfectly adequate. The ability to retrieve stored events based solely on their sky positions (RA, Dec) and timings meets most anticipated queries. At this stage, querying just any aspect of the full VOEvent data model is not required. There are also concerns as to whether the existing infrastructure will be able to scale to the event rates of some projects, but the existing protocols (Jabber/XMPP and TCPV) seem satisfactory.

A browser interface is great for trying out things and getting a feel for what something does, but for operational usage having an API to code against is much more preferable. In particular, APIs were requested so that it should be possible to define subscription filters using an expression syntax to limit the events received to those of interest. Something akin to the current Python-based syntax used by SkyAlert would be broadly acceptable. Users would also like to be able to provide specific pieces of code (“plugins”) to execute when matching events are received by the broker. They might add extra information to an event report before it is passed on to the subscriber.

3. SkyAlert

SkyAlert (<http://www.skyalert.org>) is an excellent prototype of the kind of requested broker service. Instances have been deployed by various projects around the world, and it has a strong user base (~300 subscribers). The VAO has its own set of requirements to which products and services must comply, covering development aspects such as version control and testing as well as operational and user support issues. A review of SkyAlert has been carried out to ensure that it can be integrated into the VAO with minimal effort, and also to identify potential problems that we can address: points of failure and bottlenecks. Two main activities have been identified as being required: the development of tests, and documentation. Additional tasks are then required for a production-level transient facility, such as addressing scalability issues—migrating the code base away from a single server instance—and potential integration with VAO security mechanisms.

4. Event Infrastructure

The current event infrastructure is an *ad hoc* mixture of unwritten agreements, ideology, and hacks that work. There are no inherent conceptual issues with this framework

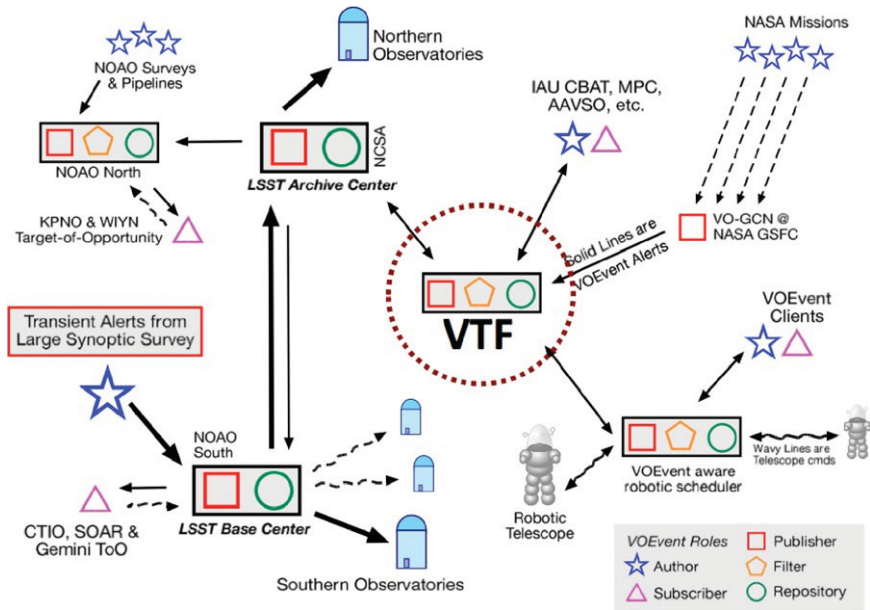


Figure 1. This shows a possible architecture for the dissemination of LSST events.

(see Fig. 1 as to how LSST events could be disseminated by it), but with its increasing uptake and usage by new projects, its various components need to be standardized to ensure interoperability. Specifically, the format for describing infrastructure metadata, a protocol for accessing and querying stored event notifications, and the transport mechanisms used to distribute event notifications need to be specified.

5. Recommendations and Schedule

The VAO Transient Facility can meet the specific needs of the transient community, in terms of event notification and brokering capabilities, by offering services based on the SkyAlert system and by ensuring that the event infrastructure is completely formalized. The development schedule for such a venture is driven by consideration of a number of potential external science collaborations, specifically the LSST 2012 Data Challenge (June 2012), in which event notifications will be circulated at a 25% expected data rate, and the GAIA Event Challenge (March 2013). Obviously such activities would serve as a good test and demonstrator of the capabilities of the facility. To meet them, we propose to have a cluster-based fully-tested installation by mid-late 2012, compliant with a fully-specified event infrastructure.

Acknowledgements

This work has been funded through NSF grants AST-0834235, AST-0909182 and OCI-0915473.

References

- Drake, A. J., *et al.* 2009, *ApJ*, 696, 870
 NAS Committee, *New Worlds, New Horizons*, Decadal Survey of Astronomy and Astrophysics (Washington, DC: NAS), 2010
 Seaman, R., *et al.* 2011, *IVOA Recommendation*, arXiv:1110.0523