


# Update on Event Broker Policy Development

Eric Bellm, Alert Production Science Lead  
LSST SAC Meeting



February 21, 2018

In August 2017 we presented the SAC early plans for community broker selection.



Included:

- An overview of technical and human capacity needed by brokers
- Expectations around LSST bandwidth constraints
- LSST plans for developing alert stream technologies and providing resources for community testing
- A high-level overview of the LSST mini-broker
- A proposed timeline for community broker selection

We are incorporating the suggestions from the SAC minutes of that meeting.



Clarify data rights issues around broker access to the LSST DAC

Clarify capabilities of LSST to enable archival studies of transients

Provide more detail on the planned proposal and evaluation process for community brokers

Explore LSST ability to service small-scale, specialized brokers

There's a range of progress on these goals to report today.



Early drafting of a community-facing document, “Plans and Policies for LSST Alert Distribution”

Activities of the LSST Data Rights Working Group (see earlier presentation)

Discussion of LSST mini-broker requirements and capabilities in conjunction with the LSST Science Platform

Further refinement and testing of the technical tools baselined for LSST alert distribution

# We are developing the “Broker Policy Document” as a broad resource for understanding LSST’s alert plans.



Goal: In narrative language, outline LSST plans and policies for alert distribution for a broad audience.

Audience: LSST science users, community broker developers, funding agencies, LSST Project personnel

Venue: LDM document (LDM-612), to be posted on the arxiv after CCB acceptance.

Potentially a living document

Timeline: draft to SAC by June 2018; CCB acceptance by AHM 2018?

Links to existing requirements documents (SRD, DMSR, DPDD, etc.) as well as new Alert Production Requirements document, new Alert Production Test Plan, new Alert Distribution Design document

Should answer all the questions posed by the SAC!

The Broker Policy document has four sections.



Components and Capabilities of the LSST Alert Distribution System

Data Rights to Alert Stream Components

Guidelines for Community Brokers

Guidelines for Users of the LSST Filtering Service and Science Platform

# The Broker Policy document has four sections.



## Components and Capabilities of the LSST Alert Distribution System

- Summarizes LSST data flow, the Data Access Center (DAC), data processing, and data products relevant to alert-driven science
- Provides links to governing design and requirements documents for interested readers
- *Answers questions like:*
  - *What's in an alert packet?*
  - *What happens in crowded fields?*
  - *What are the bandwidth limits on alert distribution?*
  - *Can DAC users query the Prompt and/or Data Release Production Difference Image Analysis (alert) databases?*  
*(yes)*

# The Broker Policy document has four sections.



## Data Rights to Alert Stream Components

- Summarizes who can access which parts of the alert stream
- Broadly: alert packet contents are world-public; access to all other components requires data rights.
- *Answers questions like:*
  - *How does a scientist without data rights access the alert stream? (through a community broker)*
  - *Can community brokers access DAC resources that require data rights? (Yes, but only if they can limit redistribution of LSST-proprietary data to users with data rights.)*



# The Broker Policy document has four sections.



## Guidelines for Community Brokers

- Outlines a process by which community brokers will be chosen
- Criteria and timeline outlined in August 2017 SAC presentation
- *Answers questions like:*
  - *Are community brokers required to rebroadcast the alert stream to all? (not required, but proposals that do so will be stronger)*
  - *Is specific or minimum functionality (e.g., filtering, archiving) required to be a community broker? (Other than the technical capability to process the stream, no. We suggest an “open” call to allow proposers to put forward brokers best suited to their aims, recognizing that the number of brokers will be limited and more capable and open offerings will be favored.)*
  - *Do community brokers have to receive the full alert stream? (Proposers can request pre-filtering using the technology employed by the LSST mini-broker if desired.)*

# The Broker Policy document has four sections.



## Guidelines for Users of the LSST Filtering Service

- Describes capabilities and limitations of the LSST mini-broker
- Users may upload simple filters that operate on single LSST packets and receive the output, subject to capacity constraints
- Many details are being resolved—see next slides
- *Answers questions like:*
  - *How many alerts can a mini-broker filter pass?*  
(Currently 20 per user per visit)
  - *Can exceptions be made to this limit?*  
(A proposal process could allow for additional resources to be allocated.)
  - *Can users make the outputs of their filters accessible to others, potentially in a private group?*  
(Yes—group permissions are an expected feature of the Science Platform)

# The mini-broker is not fully defined by existing documents.



Formal requirements: (LSE-61):

- Accept user-defined filters that act on single alert packets
- Provide pre-defined filters
- Provide up to 20 full alerts per user per visit for up to 100 simultaneously-connected users
- No latency requirement

Not prepared to adjust these performance requirements today.

Our envisioned implementation may provide improvements here but place new limits elsewhere (e.g., the total number of filters)

We envision a mini-broker more tightly coupled to the Science Platform.



*PRELIMINARY!*

Potentially a choice of output channels for user filters:

- A VOEvent stream and/or similar real-time feed
- Storage in user databases in the Science Platform
- Potentially, enable connections to other channels: email, GCN, ATEL, Transient Name Server, Twitter... ?

Share visualization tools, user groups and permissions, database storage, compute, etc. with other aspects of the Science Platform.

# Performance testing of the baselined alert distribution technologies continues.



Testing a system using open source technologies:

Apache Avro (binary serialization format for alerts)

Apache Kafka (realtime streaming platform; used for alert distribution)

Apache Spark (cluster computing framework, used here for filtering)

LSST-scale experiments detailed in Data Management Technical Note #28 ([dmtn-028.lsst.io](https://dmtn-028.lsst.io))

*acceptable performance at LSST scale, although packets with stamps in them are large enough to impact performance*

A fork of this technology stack has recently been put into production by the Zwicky Transient Facility collaboration

We expect to make substantial progress this year.



March: improve design of the LSST mini-broker with LSST Data Management

April: Alert System presentation to Science Collaboration chairs—  
solicit feedback on these plans

May: NSF milestone test for alert distribution performance

May: alert standardization discussions at the IVOA Interop

June? Draft of Broker Policy Document to SAC for review

August: Public release of Broker Policy Document