

# Performance Improvements of an Event Index Distributed System

Á. Fernández Casaní<sup>1</sup>, Juan Orduna, S. González de la Hoz<sup>4</sup>,



## Introduction and Objectives

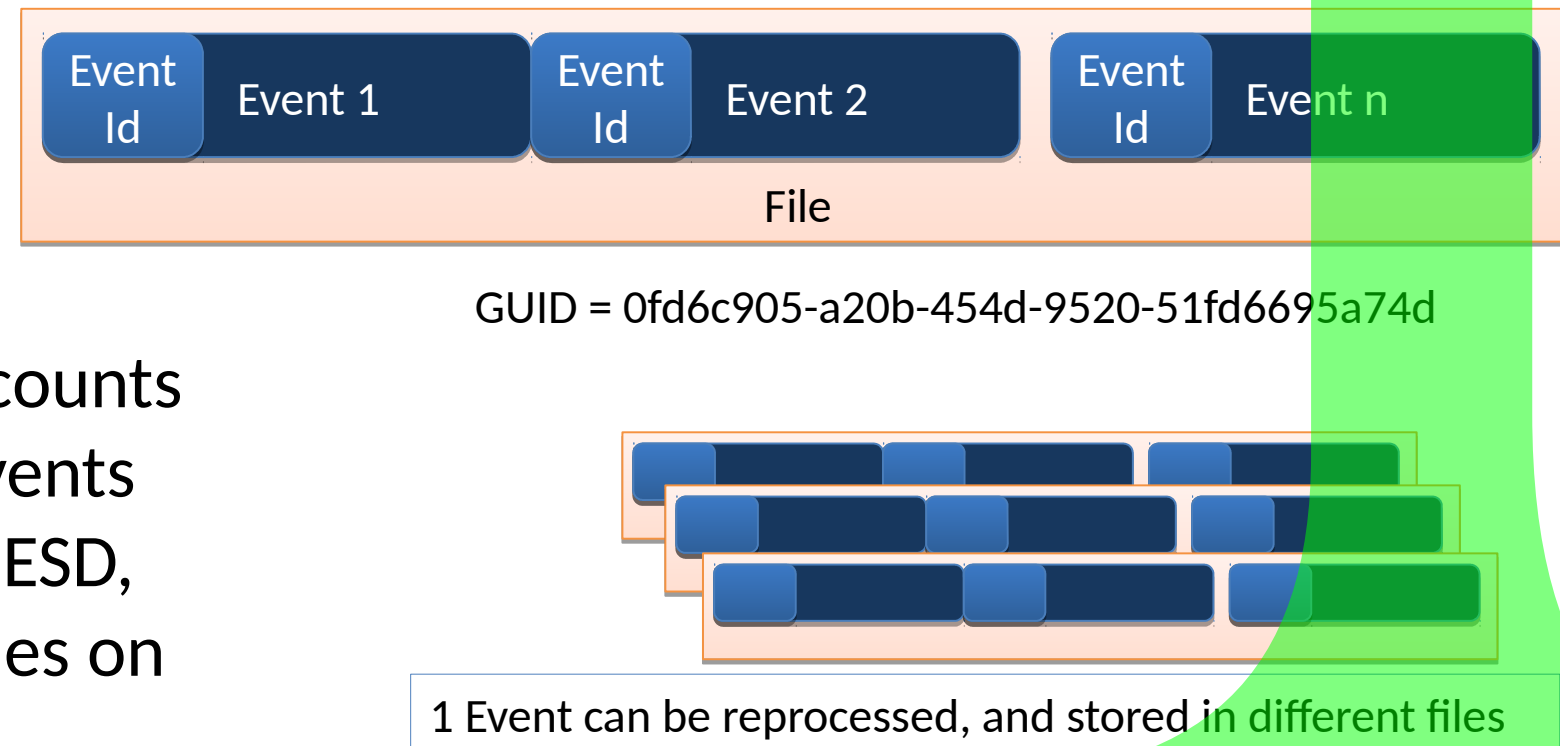
- Experiments like ATLAS produce large amounts of data that need cataloguing according to different points of view to meet multiple use cases and search criteria.
- The *Event Index project* consists in the development and deployment of a catalogue of events for experiments with large amounts of data, and the eventual distribution of the final software for the community.
- In this poster we present the current prototype of the project, including the worldwide generation of Event Indices and the usage of a Object Store (OBS) system to transmit the information to the central CERN Event Index Server. With the usage of structured storage (NoSQL) technologies for the backend database, we achieve scaling for increased data rates with commodity hardware, and fixed cost/unit.

## Event Indices production

We want to index all events while they are produced and stored worldwide, in files identified by **GUID**. Each file contains typically 50-50000 events:

Index information ~=  
300bytes to 1Kbyte per Event:


- Event identifier
- Online trigger pattern and hit counts
- References (pointers) to the events at each processing stage (RAW, ESD, AOD, NTUP) in all permanent files on storage



Event Index Information is produced by Athena/python scripts running in Tier0 and GRID sites. Events are reprocessed and stored in several file formats, increasing numbers to 138 billions (10<sup>9</sup>) real events, and 93 billions simulated events in year . Current tools include POOL, RAW, TAG formats, and produce the information in a intermediate file to be sent afterwards.

265 K gridjobs/day -> 300 Million events/day -> 1Kb/event ~ = 300 GB/day. Rate of the Producers is therefore measured as 20Hz of file records containing ~3.4 kHz of event records, to be increased the following years up to 80Hz and 30kHz of event records.

## Use Cases



**Event picking:** Give me the reference (pointer) to "this" event in "that" format for a given processing cycle.

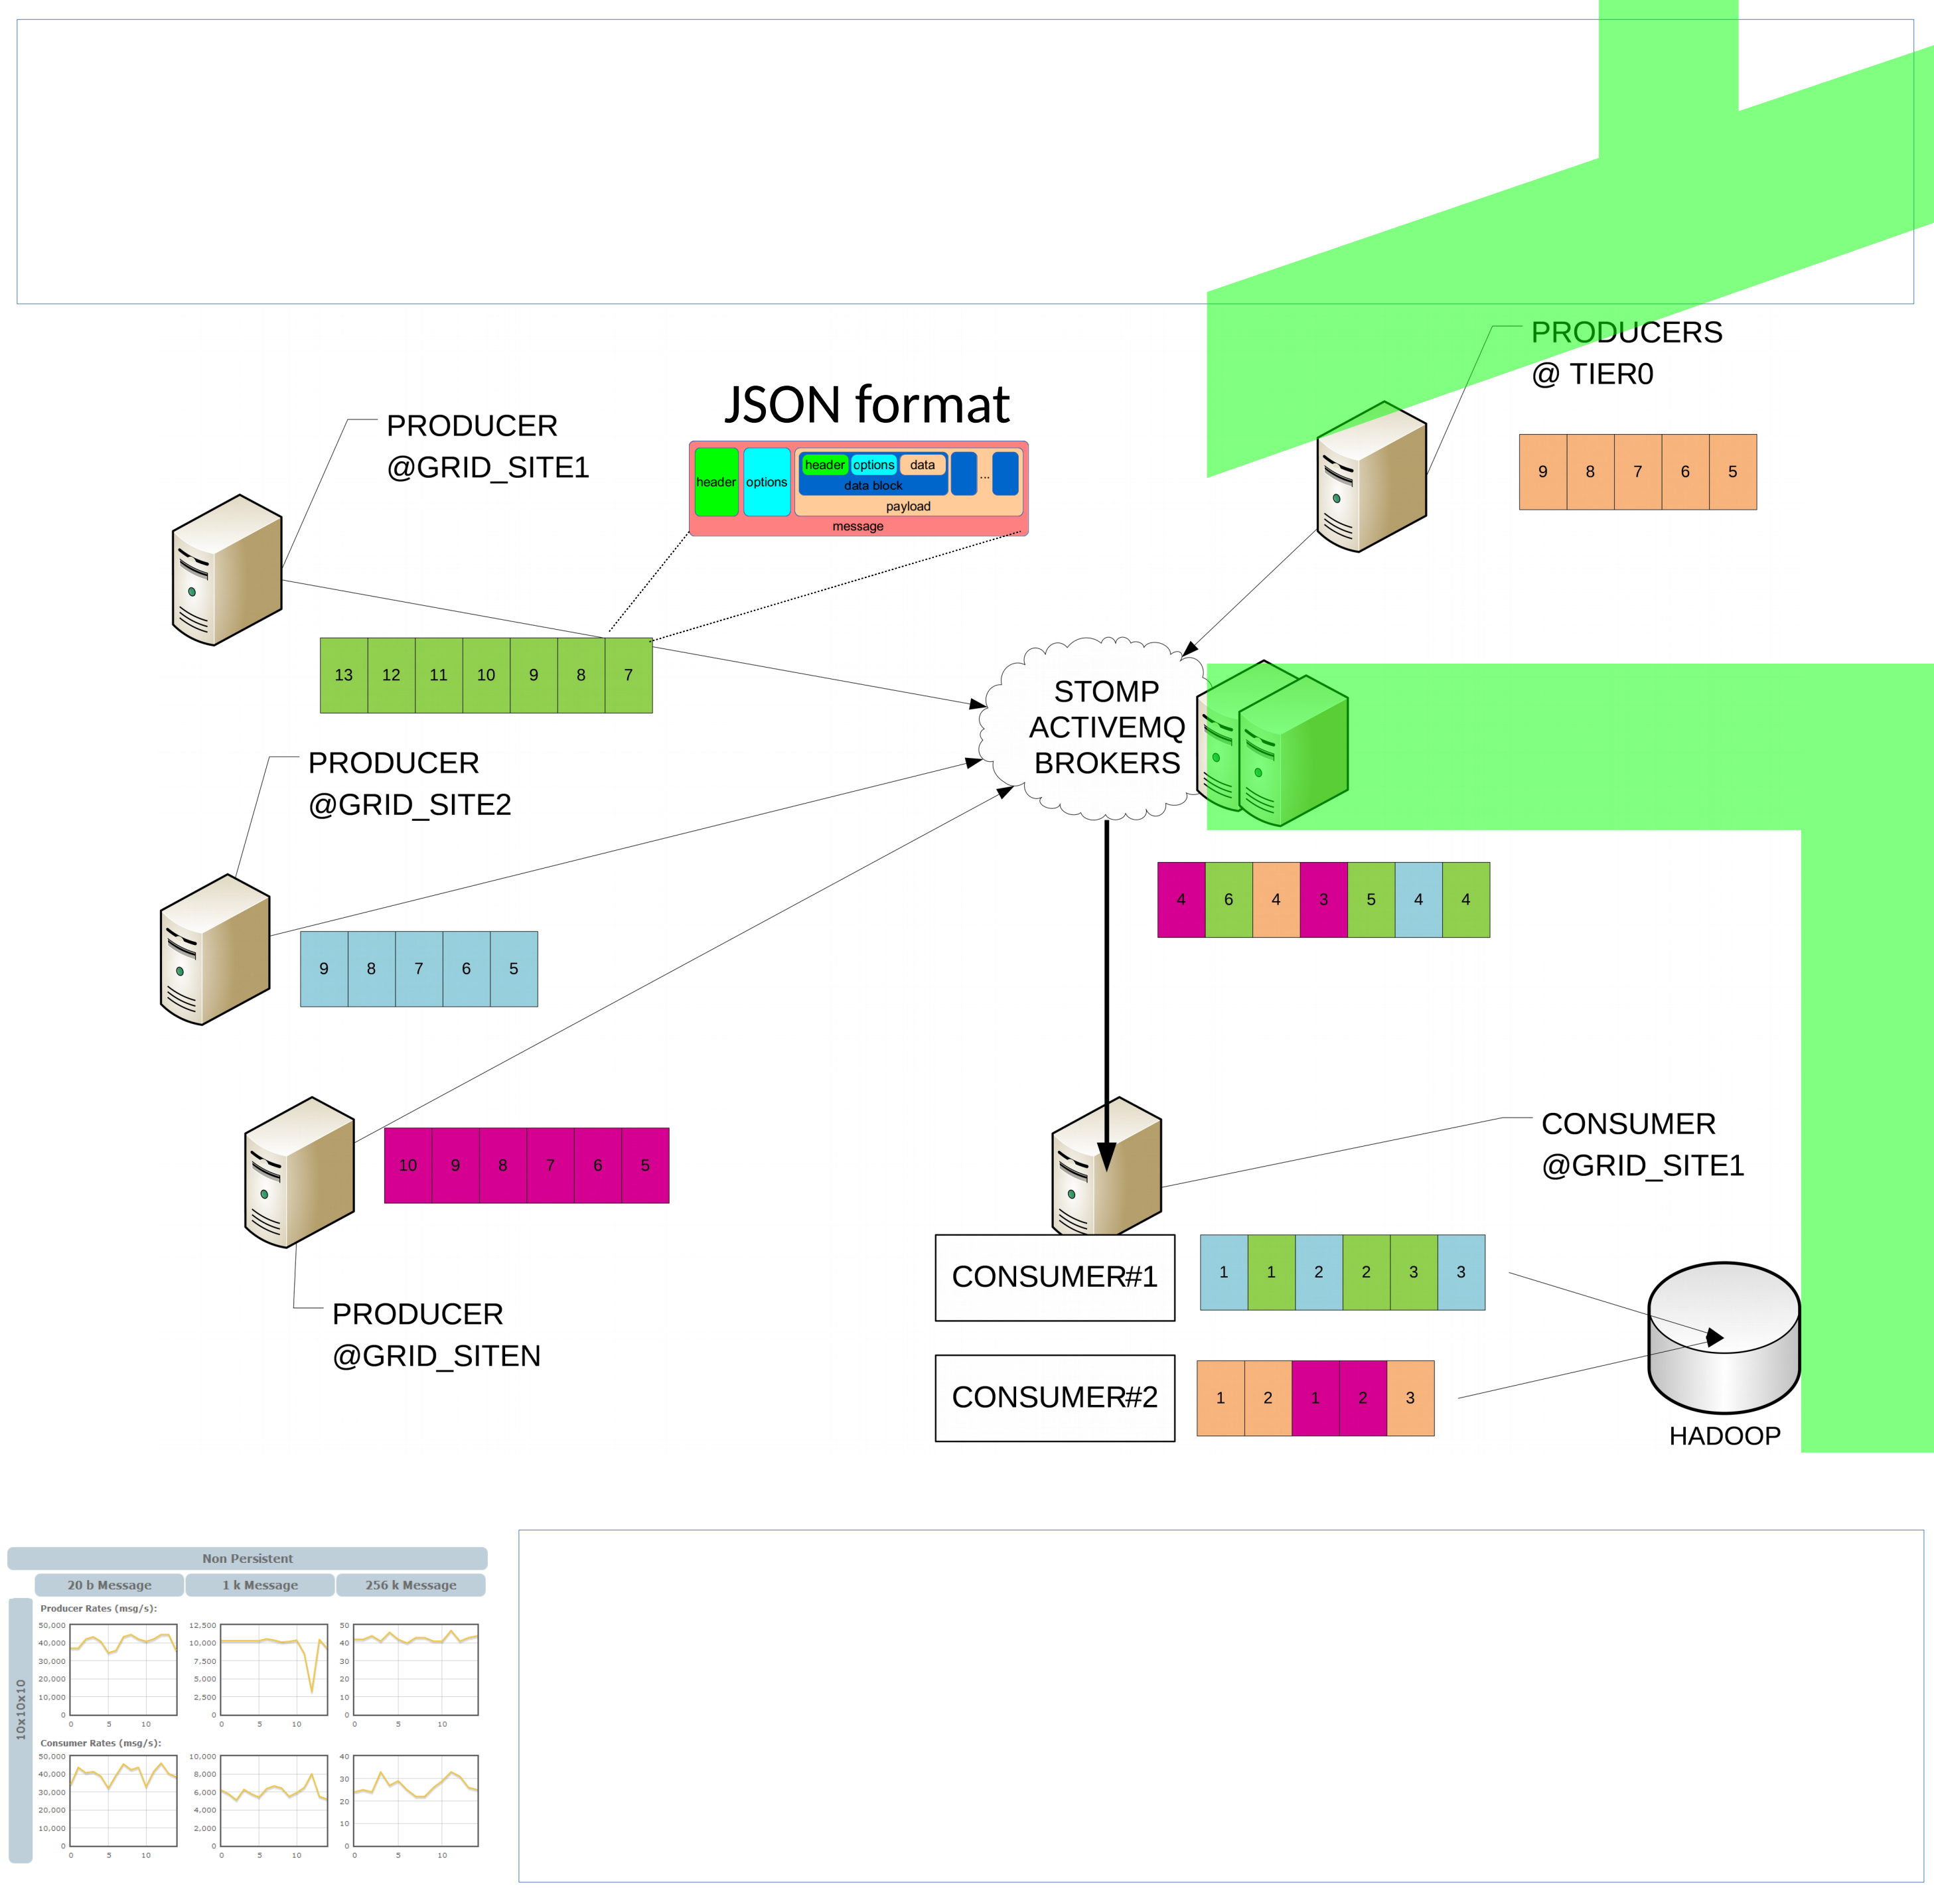
**Event skimming:** Give me the list of events passing "this" selection and their references.

**Production consistency checks:** Technical checks that processing cycles are complete.

**Panda Event service:** Give me the references for the events on this GUID file (to be distributed to HPC or cloud clusters for processing)

## Query Services

## Distributed Data Collection



## Conclusions and future steps