Astronomical Data Analysis Software and Systems XIII
ASP Conference Series, Vol. 314, 2004
F. Ochsenbein, M. Allen, and D. Egret, eds.

The ALMA Archive: A Centralized System for Information Services

A. Wicenec

European Southern Observatory, Germany

S. Farrow

University of Manchester, UK

S. Gaudet, N. Hill

Herzberg Institute of Astrophysics, Canada

H. Meuss

European Southern Observatory, Germany

A. Stirling

Jodrell Bank Observatory, UK

Abstract. ALMA will produce enormous data rates and volumes. In full operation it will generate up to 60 MB/s of scientific data and in addition auxiliary and logging data with frequencies down to 48 ms. These data have to be made persistent as early as possible after their production. Consequently the archive is placed at the very center of the ALMA data flow system and all other subsystems utilize the services provided. In addition to these services the archive subsystem has to implement the standard archive functionalities for PIs and archive researchers and it is probably the first archive to have VO compliance written in the science requirements. This paper gives an overview of the design and implementation and the current status of the ALMA archive subsystem.

1. ALMA Archive Design

The ALMA Archive design is built around two main concepts: to provide generic information services and to act as a passive archive. Generic information services are common archive functionalities like store, update, retrieve and query. These methods are implemented on the lowest level in the XMLStore and the MonitorStore using an XML(-aware) database supporting XPath. The BulkStore is implemented as a scalable file store like the ESO NGAS[1] (Wicenec, Knudstrup

[1] Next Generation Archive System

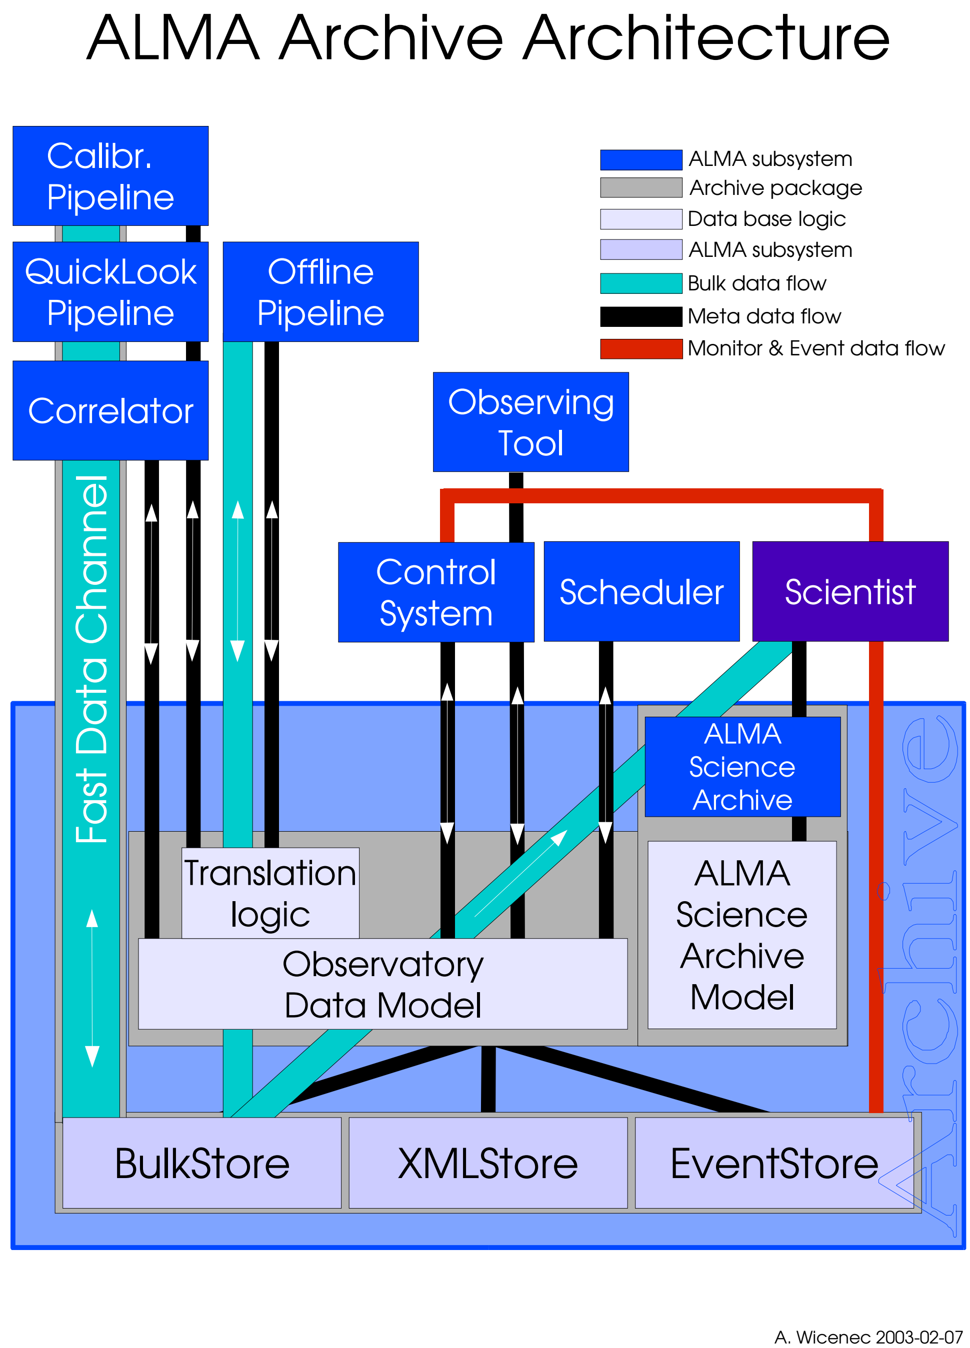

Figure 1. Schematic view of the ALMA Archive subsystem. The 'Fast Data Channel' Bar on the left connects the three subsystems Correlator (main data provider), QuickLook pipeline and Telescope Calibration pipeline with the archive. As depicted here the ALMA archive foresees access to the data through different data models, depending on whether the user is part of the observatory data flow system or an external archive user or the VO system.

& Johnston 2002) where the 'store' method is implemented as a direct streaming interface using VOTable[2] based multipart/related messages. Retrieval of files from the BulkStore by other subsystems will only be done through the XML-Store. A passive archive does not carry out any 'business' logic on the data items it holds and in particular it does not know about the semantics of the data. The semantics (cross references, relations) of the business logic reside in the data models (Observatory Data Model and Science Archive Model). The data model layer of the archive is implemented as second level meta data, i.e. a data model is kept in the database as a document containing references to meta data which in turn are describing data objects in the BulkStore or the MonitorStore. The active logic (e.g. program tracking, scheduling, archive request handling) are solely responsibility of the other subsystems, where the ALMA Science Archive is seen as a separate subsystem here.

2. The Archive as a Central Repository

Centralization in the sense of the ALMA Archive does not mean that there is exactly one place where the archive is located, but rather that all ALMA subsystems are using the Archive as an area for persistent storage. A schematic view of this is shown in the figure below.

For performance reasons there are interfaces directly between subsystems and to the archive. The meta data describing observing projects is kept in a hierarchical structure of XML documents. Some levels of this structure are not referencing any 'real' data, but are necessary to describe the project correctly. Every document will be stored in the archive as an entity with a unique entity ID. Some of the leaf nodes of this tree contain the actual correlator data, which is also stored as an entity with a unique ID.

Like this every data item, be it meta-data a data model or correlator bulk data is treated the same way and the core archive can be implemented to provide very generic functionality only. The core archive is depicted by the lowest level in Figure 1. The interfaces to the three stores are very similar, the XMLStore and the MonitorStore are based on the same code. 'Normal' subsystems do not interact directly with this layer, the exception is the Correlator subsystem which needs to stream data into the archive at a very high rate. The other subsystems interact with the data model layer above the core archive as this layer provides more specific interfaces which are usually even implemented using type safe XML binding classes in Java automatically generated from XMLSchema files using Castor[3].

Our very basic prototype archive browser is using an Apache Tomcat[4] application server and Java servelets and we are looking into IBM WebSphere and database integrated webservices as well. This kind of technology is also dis-

---

[2]VOTable: XML Format for Astronomical Tables, `http://cdsweb.u-strasbg.fr/doc/VOTable/`

[3]`http://castor.exolab.org/`

[4]`http://jakarta.apache.org/tomcat/`

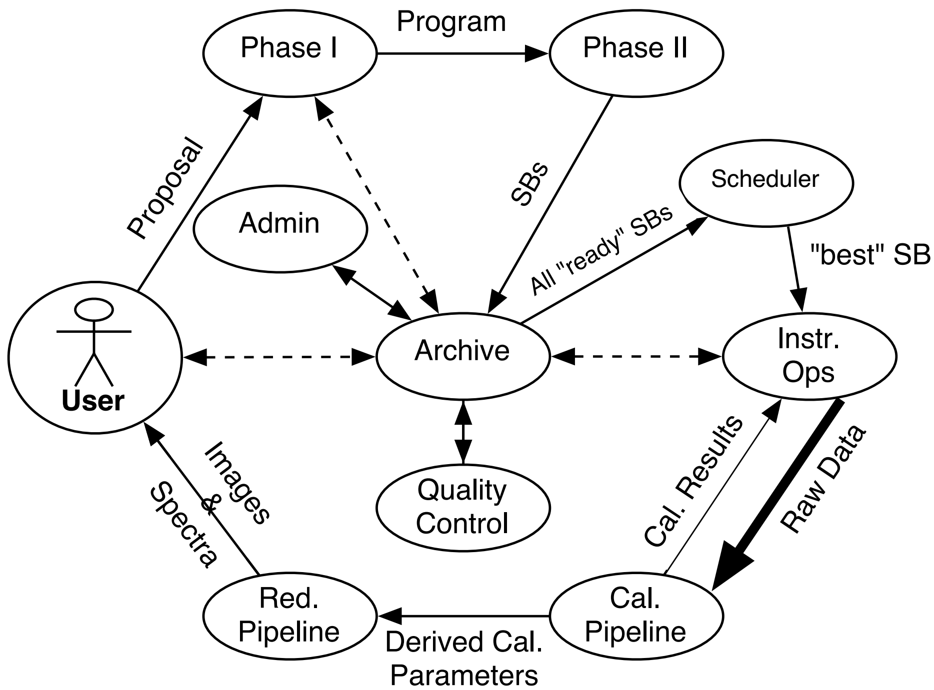

Figure 2. Schematic view of the ALMA data flow. The proposal and program preparation follows a standard two phase process. SB stands for 'Scheduling Block' which is the smallest entity handled and scheduled by the ALMA system.

cussed in the various VO working groups as it is capable of providing external interfaces for distributed systems.

The ALMA Archive science requirements list VO compliance as a very generic term, while international VO efforts converge on standards and approaches. We are actively involved in the VO discussions and development, in particular in the area of the VOTable standard where we are trying to expand the VOTable definition to be useful for interferometric data. In addition we are involved in the definition and implementation of data models for interferometric (UV-plane) data in general and radio/sub-millimeter data in particular.

References

Wicenec, A., Knudstrup, J., & Johnston, S. 2002, in ASP Conf. Ser., Vol. 281, ADASS XI, ed. D. A. Bohlender, D. Durand, & T. H. Handley (San Francisco: ASP), 'ESO's Next Generation Archive System', 95