



Towards a Provenance Management System for Astronomical Observatories

Mathieu Servillat, François Bonnarel, Catherine Boisson, Mireille Louys, Jose
Enrique Ruiz, Michèle Sanguillon

► To cite this version:

Mathieu Servillat, François Bonnarel, Catherine Boisson, Mireille Louys, Jose Enrique Ruiz, et al..
Towards a Provenance Management System for Astronomical Observatories. Provenance Week 2021,
Jul 2021, Charlotte, North Carolina, United States. pp.244-249, 10.1007/978-3-030-80960-7_20.
obspm-03344086

HAL Id: obspm-03344086

<https://hal-obspm.ccsd.cnrs.fr/obspm-03344086>

Submitted on 15 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Towards a provenance management system for astronomical observatories

Mathieu Servillat¹[0000-0001-5443-4128], François Bonnarel²,
Catherine Boisson¹[0000-0001-5893-1797], Mireille Louys^{2,3}[0000-0002-4334-1142],
Jose Enrique Ruiz⁴[0000-0003-3274-4445], and
Michèle Sanguillon⁵[0000-0003-0196-6301]

¹ Laboratoire Univers et Théories, Observatoire de Paris, Université PSL, CNRS, Université de Paris, 92190 Meudon, France; mathieu.servillat@obspm.fr

² Centre de Données astronomiques de Strasbourg, Observatoire Astronomique de Strasbourg, Université de Strasbourg, CNRS-UMR 7550, Strasbourg, France

³ ICube Laboratory, Université de Strasbourg, CNRS-UMR 7357, Strasbourg, France

⁴ Instituto de Astrofísica de Andalucía, Granada, Spain

⁵ Laboratoire Univers et Particules de Montpellier, Université de Montpellier, CNRS/IN2P3, France

Abstract. We present here a provenance management system adapted to astronomical projects needs. We collected use cases from various astronomy projects and defined a data model in the ecosystem developed by the IVOA (International Virtual Observatory Alliance). From those use cases, we observed that some projects already have data collections generated and archived, from which the provenance has to be extracted (provenance “on top”), and some projects are building complex pipelines that automatically capture provenance information during the data processing (capture “inside”). Different tools and prototypes have been developed and tested to capture, store, access and visualize the provenance information, which participate to the shaping of a full provenance management system able to handle detailed provenance information.

Keywords: Astronomy · Provenance · Virtual Observatory.

1 Context

Astronomical observatories and data providers are increasingly involved in the development of Open Science. The process of making data FAIR⁶ (Findable, Accessible, Interoperable and Reusable) often has to be integrated early in the development of astronomical projects. Since more than 20 years, the IVOA⁷ (International Virtual Observatory Alliance) provides various standards to foster interoperability and enable the production of FAIR data.

The Reusable principle is more subjective and requires rich metadata to demonstrate the quality, reliability and trustworthiness of the data. Detailed

⁶ <https://www.go-fair.org/fair-principles>

⁷ <https://www.ivoa.net>

provenance is thus a key information to provide along with the astronomical data. The IVOA validated in April 2020 a Provenance Data Model [9] to structure this information. It is based on the W3C PROV concepts of Entity, Activity and Agent [4] with a dedicated set of classes for activity description (e.g. method, algorithm, software) and activity configuration (e.g. parameters).

2 Requirements and current perception of provenance

Several use cases have been discussed within the IVOA and the European ES-CAPE project [8]. Astronomical projects that produce data generally develop structured pipelines, scripts and specific methodologies to prepare data products for the end-user from raw data (acquired from observations or generated by simulation).

Key information on what processes were applied and how they were performed is thus relevant to the end-user and could be captured directly during the process (capture “inside”). For older or other projects, a posteriori metadata extraction from data/metadata/logs (provenance “on top”) can also provide similar information, with the risk of missing details and links. We often realize too late that there are missing elements or links in the provenance, this is why the capture of the provenance should be as detailed as possible and as naive as possible (simply record what happens). In any case, the granularity of the provenance has to be adapted from one project to another.

2.1 Basic handling of provenance




Fig. 1. Basic handling of provenance information.

In general, the perception in the community is that provenance information is easily stored with the data, as a set of keywords recorded in the header of a data product file. This is represented in Figure 1. This perception is particularly strong in astronomy with the large adoption of the FITS (Flexible Image Transport System) file format [10], that provides a human readable header based on keywords.

2.2 Last-step provenance

The complex modeling of provenance information makes it improper to be stored as a flat list of keywords, as provenance is better represented by a graph, based on chains of activities and entities that are used and generated. We thus define the full provenance as this graph, up to the raw data, and the last-step minimum provenance as an embedded list of keywords [8]. The last-step provenance contains information on: the entity itself, one contact agent, the last activity that generated this entity. It also contains identifiers of other used and generated entities. All this information is compatible with the IVOA Provenance data model. Such a last-step provenance can thus be stored in a file header, and should moreover enable the reconstruction of the full provenance through the recursive exploration of used entities.

3 A provenance management system

If a basic handling of provenance information may be sufficient for some projects, it is necessary to build a more advanced provenance management system that stores this information separately, as files or in a database. Such a system is composed of the following parts :

- 1/ **Capture "inside"**: provenance information is recorded during the execution of a pipeline that runs various processing steps, generates intermediate data files...
- 2/ **Ingestion**: the captured information is transported in a structured format that can be parsed and managed.
- 3/ **Storage**: the ingested information is then safely stored in a database that preserves its logic.
- 4/ **Visualization and exploration**: the full provenance can be queried and visualized.

3.1 Tools, prototypes and protocols

Several tools have been developed in relation with the IVOA Provenance data model. They are the bricks to build a full provenance management system able to handle detailed provenance information:

- **voprov**⁸: This Python package extends the W3C PROV compatible **prov** package to implement the IVOA Provenance data model. It provides a way to create a **ProvDocument** object and exchange it as an XML, JSON or graphical file.
- **logprov**⁹: This Python package captures provenance events when running Python functions or methods that are specifically decorated and defined.

⁸ <https://github.com/sanguillon/voprov>

⁹ <https://github.com/mservillat/logprov>

Those events are recorded through the logging system as structured dictionaries, and can then be transformed using `voprov`. This package was initially developed with the high level interface of the `gammapy` package [3].

- **ProvSAP**: a Simple Access Protocol that returns a W3C PROV file from a regular GET query on an HTTP endpoint. Arguments can be passed, such as: ID, DEPTH (ALL/1...), DIRECTION (FORWARD/BACKWARD), RESPONSEFORMAT (PROV-SVG/PROV-JSON...), MODEL (IVOA/W3C), AGENTS (0/1), CONFIGURATION (0/1), DESCRIPTIONS (0/1/2), ATTRIBUTES (0/1). This system is for example implemented in the OPUS job manager¹⁰ [7] and in other tools [5].
- **ProvTAP**: IVOA Table Access Protocol using ADQL for queries and a TAP Schema, itself based on the IVOA Provenance data model [1]. It's a reverse mechanism to locate data through queries on its provenance. Every feature of the model instantiated in the TAP service can then be explored. This approach enables queries to test the data quality, based on the analysis of parameters of some key activities. It is also possible to recompute datasets whose progenitors have been found erroneous.

3.2 Description of the system




Fig. 2. Provenance management system.

As shown in Figure 2, the IVOA Provenance Data Model (ProvDM) is implemented as a relational database and connected to an access service based

¹⁰ <https://voparis-uws-test.obspm.fr/provsap?ID=a9b7e2>

on the IVOA Table Access Protocol (ProvTAP) [1]. A Simple Access Protocol (ProvSAP) is also being specified within the IVOA to provide directly W3C PROV files, using the `voprov` package.

In the system, provenance information is exchanged via structured logs, W3C PROV files (XML, JSON) or graphs (SVG, PNG). The `voprov` and `logprov` packages are being developed to propose a generic solution to the implementation of the system, along with project-specific capture tools (e.g. `ctapipe`¹¹ or CTADIRAC¹² in the context of the Cherenkov Telescope Array [6]). The Visualization & Exploration subsystem is based on standards to foster interoperability and the reuse of existing tools.

Different implementations based on this schema are possible to adapt the provenance management to the needs and size of the project.

3.3 Extraction "on top"

A last block in Figure 2 (labelled 5/) indicates the use case of already existing data from which provenance can be extracted and ingested in the system. In many astronomy projects, some provenance information can be extracted from file headers, or from log files. Such an extraction would be more efficient if embedded provenance information were stored in a standard list of keywords such as the last-step provenance list (see 2.2).

4 Software and reproducibility

Depending on the project, the workflow executed to produce science ready data (the final products) can be extracted from the provenance system designed following the IVOA strategy. For each activity execution, the input and output entities and the configuration parameters are recorded, as well as a representation of the ActivityDescription class, where the software name, version, documentation, etc, are traced. To be fully reproducible, we envisage to access such coding blocks through the ActivityDescription class by pointing to a code repository. This can be set up as a dictionary of codes within a specific project, as in the CTA pipeline or other under development projects such as Euclid, LSST, etc.

Software can also be shared within the community and curated in code registries, such as the Software Heritage [2], or the astronomy dedicated software published in ASCL¹³ (Astrophysics Source Code Library), or for multi-messenger astronomy, the future ESCAPE OSSR project¹⁴.

Many astronomical projects deal with large amounts of data and require increasing computation power. This has pushed forward the development of science platforms that implement the code-to-the-data strategy. In this new computing and distributing architecture, rich metadata profiles to describe the provenance

¹¹ <https://cta-observatory.github.io/ctapipe>

¹² <https://gitlab.cta-observatory.org/cta-computing/dpps/CTADIRAC>

¹³ <http://ascl.net>

¹⁴ <https://wiki.escape2020.de/index.php/WP3.-.OSSR>

of datasets and the code applied to process them, is a key for reproducibility and interoperability.

Acknowledgements

We acknowledge support from the ESCAPE project funded by the EU Horizon 2020 research and innovation program (Grant Agreement n°824064). Additional funding was provided by the INSU (Action Spécifique Observatoire Virtuel, ASOV), the Action Fédératrice CTA at the Observatoire de Paris and the Paris Astronomical Data Centre (PADC).

References

1. Bonnarel, F., Louys, M., Mantelet, G., Nullmeier, M., Servillat, M., Riebe, K., Sanguillon, M.: ProvTAP: A TAP Service for Providing IVOA Provenance Metadata. In: Teuben, P.J., Pound, M.W., Thomas, B.A., Warner, E.M. (eds.) ADASS XXVII. ASP Conf. Ser., vol. 523, p. 313 (Oct 2019)
2. Di Cosmo, R., Zacchioli, S.: Software heritage: Why and how to preserve software source code. In: iPRES 2017: 14th International Conference on Digital Preservation. Kyoto, Japan (2017), <https://hal.archives-ouvertes.fr/hal-01590958>
3. Lefaucheur, J., Deil, C., Donath, A., Jouvin, L., Khélifi, B., King, J.: Gammapy - an Open-source Python Package for γ -Ray Astronomy. In: Ballester, P., Ibsen, J., Solar, M., Shortridge, K. (eds.) ADASS XXVII. ASP Conf. Ser., vol. 522, p. 525 (Apr 2020)
4. Moreau, L., Missier, P., Belhajjame, K., B'Far, R., Cheney, J., Coppens, S., Cresswell, S., Gil, Y., Groth, P., Klyne, G., Lebo, T., McCusker, J., Miles, S., Myers, J., Sahoo, S., Tilmes, C.: PROV-DM: The prov data model. W3C Recommendation (Apr 2013), <http://www.w3.org/TR/prov-dm>
5. Sanguillon, M., Bonnarel, F., Louys, M., Nullmeier, M., Riebe, K., Servillat, M.: Provenance Tools for Astronomy. In: Ballester, P., Ibsen, J., Solar, M., Shortridge, K. (eds.) ADASS XXVII. ASP Conf. Ser., vol. 522, p. 545 (Apr 2020), <https://arxiv.org/abs/1812.00878>
6. Sanguillon, M., Arrabito, L., Boisson, C., Bregeon, J., Kosack, K., Servillat, M.: Storing Provenance information in a data processing workflow: one CTA use case. In: Ruiz, J.E., Pierfederici, F. (eds.) ADASS XXX. ASP Conf. Ser., vol. TBD, p. TBD (2021)
7. Servillat, M., Aicardi, S., Cecconi, B., Mancini, M.: OPUS: an interoperable job control system based on VO standards. In: Ruiz, J.E., Pierfederici, F. (eds.) ADASS XXX. ASP Conf. Ser., vol. TBD, p. TBD (2021), <https://arxiv.org/abs/2101.08683>
8. Servillat, M., Bonnarel, F., Louys, M., Sanguillon, M.: Practical Provenance in Astronomy. In: Ruiz, J.E., Pierfederici, F. (eds.) ADASS XXX. ASP Conf. Ser., vol. TBD, p. TBD (2021), <https://arxiv.org/abs/2101.08691>
9. Servillat, M., Riebe, K., Boisson, C., Bonnarel, F., Galkin, A., Louys, M., Nullmeier, M., Renault-Tinacci, N., Sanguillon, M., Streicher, O.: IVOA Provenance Data Model Version 1.0. IVOA Recommendation (Apr 2020), <https://www.ivoa.net/documents/ProvenanceDM>
10. Wells, D.C., Greisen, E.W., Harten, R.H.: FITS - a Flexible Image Transport System. *A&AS* **44**, 363 (Jun 1981)